Roger F. Gay: A Developer's Perspective

How and why did you get into the field of robotics, how has it changed in the last ten years, and what are the goals of your company?

I did not get involved in robotics until 2003 or 2004 while looking for applications of some ideas I had about improving AI in the 1980s. The problems I addressed back then were still with us in 2004 and I noticed that the technology available now makes application of my old ideas much easier; thus, commercially interesting. Helping to make robot systems work better and to make robots smarter seemed a logical place to start. My direct participation in the robotics industry started with my association with Peter Nordin, whose work in AI and learning systems lies at the heart of our commercial activities. Much of my work since then has been devoted to business development, although I have been involved in conceptualization and some high-level design. Such a fate awaits many engineers after a certain age.

It is clear that a great deal has happened in the past 10 years. iRobot's famous autonomous vacuum sweeper, Roomba only came on the market in 2002 and the company went public in 2005 due to its overwhelming success. I'm sure this will be part of the historians' account of robot history – a major turning point for the industry. Analysts have been saying that the robotics industry will grow to be larger than the automotive industry. I'm one of the greater optimists who thinks that we don't need to wait too long to see that happen.

Many of the mobile robots in use today are still largely controlled remotely by human operators. In activities such as mine clearing and some surveillance work, they are tools that allow workers to keep their distance while doing dangerous jobs. Over the past 10 years, governments around the world have been pouring a great deal of investment into robotics research, initially driven and still with heavy involvement from the military. This was well-timed and has resulted in steady progress in the related science and technology. Particularly when it comes to progress in the technology for making robots smarter and capable of performing a greater range of tasks, even insiders who shouldn't be surprised cannot help but be a little amazed. It seems to me that the expanding public interest in robot ethics is a direct result of this rapid progress. There are various estimates about how fast progress will occur in the future – how soon we'll have certain kinds of intelligent robots in our living rooms etc. – but whether or not such progress will occur seems now only debatable at the outermost fringes.

The Institute of Robotics in Scandinavia AB (iRobis) served as a technology transfer unit that brought Peter's work and that of others out of university laboratories and into first commercial form for complete robot software systems development. Peter and I are now committed to putting the software in the hands of end-product developers. This will likely involve a new company start-up. We face an educational challenge in that our software is used and performs much differently than traditional ("old-fashioned") software systems. Interest in learning systems and genetic programming in particular, including their application in robotics has grown exponentially, which is helpful. During the last couple of years, some of the largest companies in the world have started R&D programs in the field. We also keep noticing a wealth of other possible applications for a powerful "cognitive" system. How much we can do is a matter of time and money.

What are your goals for your cognitive software system "Brainstorm"?

One of our goals is to decide on a new name for the commercial product. I'll take advantage of any part of the readership that has maintained interest this far and ask that they may send suggestions if they wish.

Our initial vision has been to provide an advanced learning and adaptive software system for robots. We will provide that to companies that want to create robots and take them to market. Our primary goal at this point is to get "Brainstorm" into the hands of end-use developers. We can make arrangements for training or even joint development efforts. In some special cases, we may even be able to develop initial prototypes for potential customers.

In the near term, I've mentioned that we have an educational goal to achieve. It's still a little odd for many engineers and business decision-makers go accept the idea of letting their robots learn behavior rather than having it rigidly programmed in line-by-line. It can also be difficult to imagine letting a machine adapt – change its behavior – while in operation (optional). What if its behavior changes in a bad way? I do not see fear of the technology in the people I speak with. But the approach is new to many of them and these are perfectly reasonable issues. The simple answer – and there is one – is that it's not yet time to fire all the engineers. Although development can be much faster and robots smarter, it still takes capable people to design and develop and test before sending a product to market. Developers will still have much more than sufficient control over what is created, not just to assure product quality, but to use their own creative energies to produce useful and interesting machines.

Much of the history of machine learning actually lies outside of robotics. Genetic programming (GP) in particular has been applied to many "thought" problems. For example, GP systems read Internet material and provide specialized summaries of interest to their users and have even created patentable inventions in electronics. This has created one of our nicer challenges, although it still is an educational challenge. When we first tell people about our robotics software, they often want to know what specific tasks it has been developed to perform. I often respond by asking – what do you want it to do? In the world of traditional robotics, where advanced behavior can take years to develop, this can seem like an odd question. We are crossing a threshold from a situation in which technical capabilities drive product development decisions to one in which we are ready to ask what people want.

How can we imagine "genetic programming"? What is it used for in the development of robots? What is the difference to other approaches of AI programming?

The idea was taken from the concept of evolution. In the genetic programming approach (GP), a "population" of programs is created. All the programs run, and feedback is used to measure performance. The best performers are allowed to "survive" and are modified by processes that were inspired by genetics. One of them is a recombination of elements of two "parent" programs into a single child program. Just enough random changes are made to keep the possibilities open. This approach has been quite successful in guiding improvement in each successive generation, which is one of the reasons it is practical for use in the real world. Randomly creating and testing all possible programs for example, until one that does what you want it to do is created, would be impractical.

It is a very powerful technical approach. It is used to create "Turing complete" programs, which is to say that there are no logical limitations to the programs that can be created. It is capable of creating "arbitrarily complex" programs – in a good way. That is, there are no limitations on the complexity of the program that is needed.

Peter Nordin has been a pioneer in genetic programming for decades and much of his work is related to robotics. Starting in the 1990s, he had the opportunity to consolidate this effort in The Humanoid Project at Chalmers University in Sweden. One of the developments was the basic architecture for GP robotics software systems used in Brainstorm. Brainstorm is not simply a GP processing engine. It is the mind of a robot, capable of dealing with many things. It consists of several layers to deal rapidly and directly with low-level processing through higher level "thinking" and problem-solving processes. Built-in simulation allows the GP system to build and test its programs without physically carrying out tasks. It can first imagine how it will do something before doing it. (This also means that robots do not need to physically perform through generations of populations of programs to produce a working one.)

Within The Humanoid Project, GP was applied in hundreds of robot projects on a variety of hardware platforms. Humanoid robots learned to walk, four-wheeled robots learned to drive themselves, the world's first flapping wing robot took flight. A four-legged robot learned to walk, broke a leg, and automatically learned to walk efficiently with three legs. Robots have learned hand-eye coordination, grasping movements, to mimic human behavior, and to navigate, among other things. Higher level cognitive processes take care of such tasks as planning, evaluating safety, and problem solving.

GP is the only approach capable of producing arbitrarily complex, Turing complete programs. Brainstorm is the first robotics software system to carry the label "complete cognitive system." When we install Brainstorm on a robot (notice I don't mention a specific physical type of robot), it learns about itself, learns to move, and wanders through its environment learning about it as it goes. By knowing about itself and its environment, it is able to first determine how to deal with its environment and then carry out the programming that has been created in its own imagination.

Where are the major differences between your work and other approaches like the artificial brain projects, which also use evolutionary algorithms?

I do not know a lot about the artificial brain projects. From the name, and what I have read, it seems clear that we are much less interested in modeling the brain. Our interest is in a working robotic mind instead; a cognitive system for machines. There can be some incidental overlap in structure because nature is often quite logical in its designs, but modeling the brain is not our goal. I've heard some good things about some of the brain projects. There are some very smart people involved. But, I think even by their estimates, it will be a very long time before a functional artificial humanoid brain exists.

I should also mention that from the time of Darwin to the present day, some very smart people have theorized that our minds might use an evolutionary process when we think. People naturally rationalize. Thoughts that don't make sense do not survive (even if they're right). Thoughts emerge that make sense to the person doing the thinking (even when they're wrong). For simpler things this process seems effortless – or at least to some extent it is "subconscious." In higher-level problem solving, you might be aware of simpler ideas growing in complexity as they are examined in your imagination. You consciously discard ideas that seem like they won't work and add ideas that might.

Could genetic programming be a step towards a recursive self-improving artificial intelligence (Seed AI)?

Yes, I think so. I do not sense at present, any general consensus on what a step toward strong AI is supposed to look like, but since it hasn't happened yet, I think the floor is still open to the widest range of opinion. In GP, we still tell the system what we want it to do – at least how to measure results. This is the basis for determining whether a program is better or worse than others – which survive and which perish. Above, I stated that engineers still control the character and quality of the results. The specification of what results are considered good is a very important part of that control. This is a valuable characteristic of the software for companies that want to assure that their robots aren't going to go out of control or start a rebellion and try to take over the world, etc. The current necessity of it is also one of those things that seem to put a wall between where the technology is and strong AI.

At this point I suppose I should repeat that I'm one of the greater optimists. The interaction between design and evolution fascinates me. There are some interesting ways for GP to infer logic from examples and other input techniques being brought into the mix, like showing a robot what you want and teaching it words (and concepts). It makes sense to me that expansion of the information sources robots can use and the ways in which they learn combined with GP's ability to create new will lead to something more than the sum of the parts. Yes – I think GP is a step.

What impact could genetic programming have on complex artificial *intelligence in the field where robots act as moral agents and are confronted with ethical decision-making?*

There is a very optimistic discussion on using GP to approach robot ethics in Wallach and Allen's book, Moral Machines: Teaching Robots Right from Wrong. They describe a "top down" plus "bottom up" strategy in their thoughts about how autonomous moral agents might develop. Using GP, this involves the interaction between design and evolution that fascinates me. We have the possibility of experimenting with different philosophies of morality and to combine them, letting the robot evolve its own way of responding to moral issues. The "top-down" part is in the fitness functions, the programs that are designed to measure performance. This is our way of specifying good and bad.

We have already demonstrated the use of our system in evolving safe behavior. This implies that we're already in the field of ethics. Some "textbook" ethical questions involve the operation of a machine that puts human life in danger. Instead of asking what the human driver or observers should do, we let the software evolve its own reaction. We haven't tried any textbook ethical dilemmas yet, but we have looked into a robot's imagination and watched it consider the sacrifice of a jeep and itself to dispose of a bomb.

As I write this, a team is on its way to Barcelona to set-up and run an experiment with Brainstorm aimed at setting a course toward designing autonomous moral agents. There is enough interest among us to have added the experiment to the tail end of a larger project. I would be very interested in seeing the work continue. Two graduate students are involved in the effort, which may lead to some interesting thesis results if they decide to keep this focus.

The discussion around the Barcelona experiment has already become quite interesting. How, for example, might the robot's knowledge of self in combination with its ability to imagine be used to improve its moral judgments? Can we substitute a human model for the robot's self to create a form of artificial empathy? From where we are now, I can easily imagine a meaningful exploration in a larger focused project.

UPDATE: At the end of the work done at iRobis, we were able to squeeze in an initial experiment in robot ethics even though our project was not specifically funded to do that. We used a REEM-B humanoid robot at Pal Robotics in Barcelona and provided software that allowed the robot to learn how to please the human it was interacting with. The robot learned (rather than being programmed with the knowledge) to recognize a can of Coke (Coca-Cola), that Coke will quench human thirst, and that it makes a person happy to receive one when they are thirsty. It created its own set of rules for pleasing someone based on that learned knowledge.

The experiment was included in a Swedish documentary that provided a broader look at the robot ethics and RoboEthics discussion.[1] Noel Sharkey provided a pessimistic perspective against the demonstration and Peter Nordin's positive vision. Unfortunately, time ran out before the robot faced the planned ethical choice between providing a Coke to quench the human's thirst and quenching its own thirst for electrical power.

The experiment reinforced my optimism about the short-term potential for advances in learning ethical behaviour. It was rather clear and simple, as initial experiments should be; especially when there is little time to do it all. And simple isn't bad. Engineers face a lot of complexity and a simple yet powerful idea is gold. In a very short time, a general engine for learning relationships between its own behaviors and how humans are affected by it and applying the learned knowledge was created and demonstrated.

What's been demonstrated is a shift from the need for human programmers to develop and program the logic and knowledge required to create autonomous moral agents. It has been shown that there is a way for robots to learn about the effects of their behavior using simple determinants for what is a good outcome and a bad one. It seems to me that the case for optimism is extremely clear and concrete. Robots can learn ethical behavior much the way humans do; with the advantage of learning ethics in controlled circumstances and being tested sufficiently to assure the quality of outcomes.

As I explained above, the GP learning approach is Turing-complete and capable of producing arbitrarily complex programs. Logically, there is no practical limit to what can be accomplished.

For the use in robots you have put forward an "ethical regulator mechanism". How could such a system work?

In Brainstorm Responds to Robot Ethics Challenge[2], I describe something of the idea from the 1980s mentioned above, and its potential for application as an ethics regulator. I used a rather generic title for the idea – HLL (High Level Logic). It was initially envisioned as a concept for creation of more powerful expert systems and was a few years ago suggested to a large number of AI scientists and roboticists as having potential for development of a standard component for many AI systems, including autonomous robots.

HLL includes "experts", manager(s), and at least one executive related in a hierarchy similar to many human organizations. Executives set goals and assign them to managers. Managers formulate plans and have the authority to approve or disapprove actions. Both executives and managers have specific responsibilities in formulating and controlling acceptable behavior. Experts with specialized knowledge can play a supportive role involving details, such as whether an action would violate the Geneva Convention.

HLL also provides a structured approach to robot-robot and robot-human interaction. For example, a human commander could modify a robot's executive orders and then allow the robot to carry out the orders autonomously. Given the same structure in a group of robots, it was easy to imagine executives assigning tasks to other robots – chain of command. Each robot's own executive would be aware of the robot's capabilities, which could include sufficiency in ethics related to a particular command. In this way, an individual robot's executive could potentially refuse to carry out an order when it is incapable of properly performing the task; instantly informing commanders that another decision is needed. A structured approach to sharing knowledge specifically as needed, automatically, is also in the vision.

There were plans to build HLL and integrate it with Brainstorm during a project that is now at its end. About the time the project started however, Microsoft offered its robotics development kit and the technical team decided to start by using it to deal with some of the lower level and service oriented mechanisms that HLL would have provided. Peter Nordin's initial design already included high level processing in ways that nicely integrated with or directly use GP processing. HLL got shifted off the table. I built an initial prototype in 2007 that includes the basic structure. But so far it's only been run independently with a very simple robot simulation.

UPDATE: I have started a 6 month project that includes making HLL available in an Open-Source project. A cleaned up version of the simple prototype built at iRobis should be online by the end of August (2010) along with a description of desired improvements. The first offering will include a very (very) simple robot simulation. I hope it will one day be used in development of ethical processing. At least small demonstrations should become very simple as the basic system matures. (Some simple demonstrations wouldn't be terribly difficult from the start.) Of course, it would also be quite nice to have HLL applying some learned ethical behavior as well.

Do you think human society is ready for autonomous systems in their daily life?

I'm sure that I want a washing machine to wash my cloths rather than doing it by hand. Same goes for the dishes. Better still if I don't need to be involved in either activity. Let someone else do it, or some thing. Humans tend to put enormous effort into making life easier and I doubt acceptance will pose an insurmountable problem. I think history can tell us much about what problems we should expect. When increased automation happens quickly for example, it can cause unemployment. But adjustments have always been made, smooth or not. When adjustments lead to higher paying jobs and lower prices, workers will run out and buy the new gadgets. Ask Henry Ford. Ask the stockholders in iRobot, which went public after only a few years due in part to acceptance of their autonomous vacuum sweepers and floor cleaners. In the broader view, I believe the age of robots will be welcomed by society in a fashion not unlike that of the acceptance of automobiles and computers. Aside from the particular benefits robots will provide, the potential for industry is enormous. Society always seems to like it when economic times are good – and the quality of life benefits that brings. What we need are plenty of good robots at reasonable prices.

Generally humans are less forgiving if machines make mistakes than if humans do. Will the human society be able to cope with robots which choose their actions according to their goal autonomously?

I'm not sure that I agree with the question's premise. Maybe it's partly because I'm an engineer. When I look at a cute, fuzzy little baby seal robot snuggling someone, I'm still very much aware of the machine parts and electronics that lie beneath the artificial skin. I could easily destroy a machine with the only consideration being cost verses benefit. Forgiveness isn't much of an issue. Not so with a fellow human. Be that as it may, I believe human society will cope in one way or another. Even in the longer term – if we're talking about – maybe even robots that are smarter than we are. There will always be those among us who will work to solve problems rather than giving up. I can however imagine recalls to fix problems, such as with automobiles, and the possibility of "grounding" robots until fixes are made – as well as investigations like the FAA conducts after airplane disasters. I also think manufactures will be aware of potential economic liabilities, which – aside from our own humanity – will help guide decisions about the products that are offered. Safety isn't a new issue in design, manufacture, and sale of machines.

The question of responsibility – who will be held responsible for actions of a (semi)autonomous robot? Is there a need for additional legislation?

I'm a bit pessimistic about the possibility of additional legislation having a positive effect (although I should mention that I don't know much about Austrian law). I think the best guidance is to look at what is already established and by working with the understanding that robots are machines. In the near future, whether a manufacturer or an operator should be held responsible depends on the details. What happened? In concrete circumstances, it should usually be much easier to determine who was at fault after something goes wrong. The difficulties will not be unlike those of centuries of liability cases in human history. Established precedents should still hold validity. The common law approach offers the benefit of dealing with new circumstances as they arise, based on a concrete view of what actually happened; whether a manufacturer delivered a faulty product, whether maintenance was performed improperly, whether an informed operator chose to take a risk, or whether something happened purely by accident – unpredictable and beyond human control.

My view is seasoned by engineering experience. It is first principle in product development that we create things that people want. In most of my personal experience, this has always meant creating useful things on purpose. The path is still one of deciding what useful things we want robots to do, designing, building and testing products before they go to market. I understand that your question comes from consideration of future robots with behavior that is more truly autonomous. That gives rise to our interest in robot ethics. Optimistic as always, I believe the technology of ethics for robots can grow alongside increased autonomy. We should be looking at this as part of the equation for maintaining a balance.

A lot has been written on the use of robots in elderly care and in the entertainment industry mainly concerning on how this will influence interpersonal relations. What do you think is the future of robots in and their impact on the human society?

I've spent time as a hospital patient and wouldn't rate it highly as a stimulating social experience. Some of the stories I've heard about abuse of the elderly in care facilities make my teeth curl. I look forward to the day when robots can take over many routine duties and are vigilant and capable enough to recognize when something is wrong. This is in some way an expansion on the idea of hooking people up to monitors, but may be less physically intrusive. This doesn't mean that people in care should be entirely isolated except for contact with machines. Human specialists could focus more directly on social and psychological needs. I wouldn't underestimate the positive value of psychological stimulation from machines, however. Benefits have been shown in controlled circumstances. We also need to use our imaginations to create benefits. Technology might for example, more reliably inform someone when a friend is going to a common area in an elderly care facility and wishes company. It could potentially keep family members better informed about the state of their relatives in care. Again – an important question is – what do you want it to do?

On a more general basis, what do you think about robots as moral agents? What is to be expected in the next decade?

One can imagine robots standing motionless, doing nothing, in the presence of a human in need. So long as they are not causing the problem, there would be little distinction in this regard between the robot and an automobile or washing machine. It could be better of course, if robots were capable of helping people in need, even choosing to perform "heroic" acts to save a human from tragedy.

As I said above, I think designing ethics into robots is part of the equation for balancing increasing autonomy. Put in human terms, greater autonomy should be balanced with greater personal responsibility. As robots become more intelligent, more capable of "thinking" for themselves, we need mechanisms to control the quality of their decisions that are equal to the task. I take this as part of a design philosophy, separate from the issue whom courts hold liable when things go wrong. A manufacturer can be held liable for not including sufficient ethical safeguards in a robot's design.

Some aspects of moral behavior are obviously quite necessary. For example, if we want to build robots that are physically capable of killing and put them into domestic service, we don't want them to go around killing people. In fact, we will very definitely want them to avoid behavior that could result in harm. We can't build them as simple utilitarian creatures, single-mindedly concerned about specialized tasks. If we did, we might end up with what is now only a sci-fi nightmare – robots disposing of living creatures because they get in the way.

Accurately predicting what will actually happen in the future, in a particular time period especially, is a lot harder than discussing what is possible. To a pretty large extent, what actually happens during the next 10 years will depend on who gets money to do what and how much. I will predict an increase in interest and funding for research of work on moral agents. This prediction is based only in part on what I have said so far. I believe that research into developing autonomous moral agents can yield a great deal of general value in the field of AI. After all, we use the same mind to process moral questions as we do others.

In the field of military robots ethical questions have been raised. Some questions are tackling issues at hand other issues seem decades away. How do you see your role as a developer in this context?

In modern design, we tend to create enabling technology – technology that can be used for the creation of numerous end-products for a variety of purposes. Brainstorm, and GP technology generally, is an advanced example. If you build a fitness function specifying what you want to happen, a program can automatically be built to do it. We intend to put this technology in the hands of end-product developers, where final decisions about product design will be out of our control. I would be more than happy to include the best tools for ethics in the development package. Our present effort focuses on the use of existing Brainstorm technology and demonstrating fitness functions for ethical decision making. Expanding even just this effort would be of value. I have noticed over my years in the software industry, the rapid adoption of solutions presented as examples in software tutorials. Aside from the research value in doing the work, the more examples we can produce, the better.

I also appreciate being able to address your questions, whatever they may be. The interest in interdisciplinary discussions regarding robot ethics is quite beneficial in my opinion. I am pleased to participate.

It has been argued that, if properly programmed, robots could behave more ethically than human soldiers on the battlefield. How can we imagine something like a machine readable copy of the Geneva Convention?

Not all weapon systems should be measured against human performance. I read an argument recently regarding an autonomous delivery system, capable of completing a bombing mission on its own. Although the record is not perfect, smart technologies have been better at finding and hitting military targets with less incidental damage.

There is a larger set of considerations to this case, but my point here is that system performance should be compared to appropriate alternatives.

But let's consider, hypothetically at least, a goal of building an ultimate autonomous humanoid soldier. We want this breed to obey the Conventions rather than being a Terminator type robot that decides to wipe out the human race. They might even play an important role in United Nations peace keeping missions. If they're going to think for themselves and decide actions related to the Conventions, then they will need to have knowledge of the rules of the Conventions, one way or another. Somewhere in the chain of research and development, this knowledge needs to get into a form that machines can use.

The suggestion made in my article assumes that putting Geneva Conventions in machine readable form is not a complex undertaking. Based on the parts I am familiar with, it does not appear that it would be. Neither would formulating rules that computers could process.

The greater technical challenge is in getting the robot to recognize circumstances well enough to reliably apply the rules. For example: Can a robot distinguish between a green military uniform and green civilian clothing? What if combatants on the other side dress to blend with the civilian population? Can it distinguish between a combatant ready to fight and one who is surrendering? The challenges have led to suggestions that robots capable of autonomous behavior should be banned from combat roles – where they could replace human combatants.

There are many possibilities for autonomous and semi-autonomous machines that are more limited in their capabilities than an ultimate humanoid soldier. What approach can be taken to create efficient and objective criteria to assure that autonomy is sufficiently ethical? ("Efficient" in form for real-world use.)

What we need is a systematic approach that integrates development of robot ethics capabilities directly with the flow of R&D work on autonomy. My idea for immediate use of machine-readable Conventions is rather basic, involving just the sort of thing that R&D engineers often do. Build, test, and focus engineering effort on solvable even if challenging problems. Continue research in areas where solutions are still farther out. Is the robot's performance better, worse, or equal to human performance? Keep track of the roles and circumstances that can be supported technically, in view of the Conventions. Maintain a balance between measurable technical capabilities and the roles and circumstances of autonomous machine behavior in deployment.

I have imagined an Open / Open Source project that would first focus effort on creating a basic or generic (sort of) machine-readable encoding of the Conventions. What I mean is that it should not be geared especially toward any particular processing system. Developers should be able to use it efficiently while making their own choices regarding how the data is processed and how it is used. One team might choose a rule processing system while another may use it in conjunction with a learning system. It could also be used to formulate tests in simulation and in quality assurance systems that would also help formulate functional and technical requirements.

Having an Open project seems a good idea. It would allow for rapid dissemination to all interested parties and could make use of feedback from a larger body of users. Of critical importance is general agreement on the validity of results. Perhaps this point becomes clearer in the next paragraph.

Extending such a project could lead to detailed answers to robot ethics issues. The project could play a central role toward developing international technical standards; standard tests and benchmarks, working out how to measure performance of systems in particular roles at particular levels of autonomy. A certain amount of technology could also be developed in support of applying the standards, running the tests. It's a safe bet that the first machines to pass tests won't pass all the tests. The complete ultimate humanoid soldier that properly handles all aspects of the Conventions is a ways off yet.

The idea responds to those who suggest banning all autonomous machine capabilities as weapons, by suggesting technical support for an acceptable deployment strategy. Bring autonomy into a fight to the extent that it is well-tested and quality assured. Will autonomous robots be better than humans? Yes, robots designed to obey the rules will be better if testing standards require it. Those that aren't, won't be deployed (by nations that accept the standards). Taking a systematic approach that integrates ethics directly into the R&D process would push development of better ethical performance. Ethics would become a systematic integral part of technical development, with results measured according to international standards.

Let me take that one step further. Let's say we do all this and the UN finds the standards and testing an acceptable way to determine whether autonomous machine technology can be used to support more dangerous potential peace-keeping missions – or even peace-creating missions. Can machine autonomy help to create and maintain peace? It's a thought that brings my own focus back to a question I've asked more than once. What do you want robots to do?

---

[1] English translation available here: http://isr.nu/robots/SVT_Barcelona_EN.doc.
[2] available on the Internet: http://mensnewsdaily.com/2008/12/10/brainstorm-responds-to-robot-ethics-challenge.