Grand Challenge: Building Artificial Minds

Włodzisław Duch, Nicolaus Copernicus University, Toruń, Poland

Long-term Vision:

"Artificial Minds" (AMs), or personoids, are software and robotic agents that humans can talk to and relate to in a similar way as they relate to other humans.

WWW is full of bots: talking heads that answer questions about commercial companies (VW, DaimlerChrysler), hospitals (Harvard), and politicians (German Kanzler).
Bots lack personality, anticipation and genuine understanding of their interlocutors.
Needed: intelligence + mind, sharing many assumptions and behavioral patterns.
Beyond bots: AMs require models of human cognition, many aspects of human mind.
AMs capable of sympathetic understanding of their owner will adapt themselves to particular personality and habits of users, providing access to structured information on any subject and advice on everyday decisions.
AMs as smart interfaces for telephones, PDAs, personal computers, toys and robots.

Grand Challenge: **creating Artificial Minds.**

Building artificial minds is **an ultimate challenge** for many scientific disciplines coming from cognitive science, computer sciences and engineering.

AMs will radically change the way we live: provide information, reason about it, explain it, use well-structured databases and Semantic Web, know about the global information ecosystems to advise on trivial and serious decisions, facilitate exploration, interaction and critical evaluation of knowledge sources, changing education.

R&D Themes (broad, overlap with some FET proposals)

- Natural Language Processing (NLP) and Knowledge Engineering

Many new ideas in NLP, question answering, ontologies, data-oriented parsing, probabilistic grammars, evolving semantics, information extraction, knowledge representation and management, graphical models (semantic, belief, Bayesian networks).

- Modeling of brain functions

Recognition, semantic, episodic and working memory models, attention, executive function, recognition and expression of emotions, reasoning, chunking, perception.

- Cognitive robotics

Soft robotics: bots, talking heads, avatars, graphics and emotional expressions.
Real robots: toys and humanoids, physical embodiment, cooperative behavior.

Grand Challenge: Building Artificial Minds

Włodzisław Duch, Nicolaus Copernicus University, Toruń, Poland

Why AMs?

Building artificial minds is **an ultimate challenge** for many scientific disciplines coming from cognitive science, computer sciences and engineering.
It requires integration of many technologies proposed during this workshop.

Cognitive sciences should provide theories of cognition leading to mind models, including simulation of brain functions, various types of memory, many learning techniques, linguistic competence, executive goal directed functions of frontal cortex, attention, awareness, affective evaluation of information saliency, non-verbal cues in speech, reasoning, better understanding of creativity, intuition and many other issues necessary to make interactions with AMs similar to interactions with real minds.

Computer sciences should develop multi-agent technology, understand their cooperative behavior, specialized and common sense ontologies and knowledge bases, graphical models for huge semantic networks and analysis of semantic spaces, natural language processing algorithms, affective computing techniques, scalable algorithms for learning and understanding structures at perceptual and conceptual level, multicriterion optimization techniques, self-organizing processes and algorithms creating complex behavior from simple build-in drives.

Engineering sciences inspired by better understanding of minds should implement some of these ideas in cognitive and developmental robotics (Japan is clearly leading in these areas), sensor networks, signal analysis, control systems, building teams of cooperating robots and designing very complex systems.

AMs should serve **as** an **enabling technology** for many types of applications.
As sophisticated user interface: collecting interesting information for the user or for the project, for example building a theory of biological system, presenting relevant information from human perspective, knowing the types of arguments that a particular user will find important.
As personal avatar: software agent that will serve as “alter egos” of the user. Many routine tasks may be delegated to personal avatars working 24 hours a day.
AM systems capable of reflection on the state of their own working memory, sustaining internal dynamical states that reflect imagination constrained by extensive knowledge of possible relations between objects, may notice new interesting knowledge structures, and may claim to be aware of their own existence.

Feasibility: Very few attempts to create crude, integrated mind models were made on the symbolic level, but not yet at the sub-symbolic level. Development of cognitive science and electronic hardware leaves no doubt that in 10+ years time artificial minds are going to be one of the central topics in science and technology. Great progress has been made in low-level cognition: speech synthesis and understanding, natural language understanding in restricted domains, sub-symbolic grounding of semantics in observations and actions of the artificial mind (simulated baby minds and personalities from situated exchange of information), evolving elementary linguistic abilities, extensive specialized and common sense knowledge (ontologies).

Grand Challenge: Building Artificial Minds

Vision: software and robotic agents that humans can talk to and relate to in a similar way as they relate to other humans are called here "Artificial Minds" (AMs). Bots (conversational agents) are already proliferating on Web pages, frequently in the form of talking heads that answer questions about commercial companies, hospitals and politicians. They are not based on models of human mind, and thus have no working, semantic and episodic memory, lack personality, anticipation and genuine understanding of their interlocutors. Responding to cues found in questions bots provide irrelevant and irritating answers. Intelligence in itself is not sufficient, people relate to each other because they have similar minds, sharing many assumptions and behavioral patterns. AMs should be based on theories of human cognition, modeling various aspects of human mind. Artificial mind technologies will be used in telephones, PDAs, personal computers, in toys and robots. AMs capable of sympathetic understanding of their owner will adapt themselves to particular personality and habits of users, providing access to structured information on any subject and advice on everyday decisions.

AMs will radically change the way we live. AMs should not only provide information but also reason about it and explain it. Personalized AMs, knowing about preferences and goals of their users, will be able to provide them with additional information, rational advices and protection from mistakes. Many well-structured databases already can be accessed in the Internet, and the Semantic Web should be available in not too distant future. This will lead to the global information ecosystems that AMs should be able to access and use. AMs will be used in making trivial decisions, such as "to which bus stop should I go to minimize traveling time", or "I like the style of this song; find me similar stuff". It will also be used in more serious decisions, like selection of health plans, evaluation of candidates for a job, or evaluation of consequences of political choices. Education will focus on exploration, interaction and critical evaluation of knowledge sources, changing the role of teachers from knowledge providers to mentors and experts on finding and using relevant knowledge.

AMs should serve as an enabling technology for many types of applications. Some AMs may specialize in creation of very complex theories, such as theories of biological systems. At present system-level models of cells and microorganisms are created by providing a framework for knowledge base, and updating the information manually after thorough analysis of experimental literature. Thousands of proteins, genes and chemical reactions are linked together in a model that is too complex to comprehend its details by human mind. AMs may oversee the process of refining such models not only in biology, but also in system-level models of functioning of local communities and global economies. AMs should be able to explain at least some consequences of actions, such as the influence of blocking particular gene on the metabolic function of the organism, or the influence of changing some taxes on economy and environment. AM supervisors should oversee such complex projects, and grant access to personal AMs that will present relevant information from human perspective, knowing the types of arguments that a particular user will find important.

Artificial minds will also be used in personal avatars, software agents that will know preferences of their owners, acquire some of their habits and behavioral patterns, and may serve as their "alter egos". Transactions between people and service or utility providers, shops, government offices and other people in many cases will be delegated to their personal avatars working 24 hours a day. Societies are formed by strongly interacting individuals; all systems that interact strongly show emergent properties. So far very little experience has been gathered in construction of networks of complex agents, but it is quite feasible that interacting teams of AMs specializing in different subject areas will be able to solve problems that are beyond the capabilities of each individual AM. Moreover, AM systems capable of reflection on the state of their own working memory, sustaining internal dynamical states that reflect imagination constrained by extensive knowledge of possible relations between objects, may be able to imagine new things and may claim to be aware of their own existence.

Building artificial minds is **an ultimate challenge** for many scientific disciplines. **Cognitive sciences** should provide theories of cognition leading to mind models, including simulation of brain functions, various types of memory, learning, linguistic competence, executive goal directed functions of frontal cortex, attention, awareness, affective evaluation of information saliency, non-verbal cues in speech, better understanding of creativity, intuition and many other issues necessary to make interactions with AMs similar to interactions with real minds. **Engineering sciences** inspired by better understanding of minds should implement some of these ideas in cognitive and developmental robotics (Japan is clearly leading in these areas), sensors, signal analysis, control systems, building teams of cooperating robots and designing very complex systems. **Computer sciences** should develop multi-agent technology, understand their cooperative behavior, specialized and common sense ontologies and knowledge bases, graphical models for huge semantic networks, natural language processing algorithms, affective computing techniques, scalable algorithms for learning and understanding structures at perceptual and conceptual level, multicriterion optimization techniques, self-organizing processes and algorithms creating complex behavior from simple build-in drives.

Feasibility: great progress has been made in many of the fields mentioned above. Speech synthesis and understanding (conversion of speech to text), facilitating voice-based communication, is well developed. Natural language analysis is still a great challenge, because it must be based on broad knowledge about the world that is missing in current dialogue systems. It should also use sub-symbolic grounding of semantics in observations and actions of the artificial mind. Progress in this direction has been relatively slow but many promising ideas are being tested. Simulated baby minds and personalities are being developed as a result of situated exchange of information, evolving elementary linguistic abilities. Experiments with robotic toys, such as AIBO dogs, are also aimed at evolving language by situated learning.

Personalized software agents that has been recently introduced use speech-to-text techniques to talk with bots; although they are far from making an impression of being artificial minds they are already useful, being able to remember simple statements like “John P’s address is Harold Lane 23”, “John P’s email address is jp@yahoo.org” etc. Such agent may be asked to email John P, tell a joke or to find some information stored in the computer or in the Internet. Experimental search engines that access structured sources of information (such as the MIT Start system) may be queried using natural language statements. Interesting and useful talking heads performing a variety of tasks will appear in the next few years, but without models of artificial minds they will remain at the level of “smart idiots”. Computerized tutorials modeling knowledge of students have been used with success in narrow domains. Techniques that prove correctness of theorems are widely used and may be extended to answer questions of the type “is this a correct assumption”, rather than more difficult question of the type “how to solve this problem”. There are no large-scale memory models (working, semantic or episodic), but some existing approaches may be scaled up. Very few attempts to create crude, integrated mind models were made on the symbolic level, but not yet at the sub-symbolic level. Development of cognitive science and electronic hardware leaves no doubt that in 10+ years time artificial minds are going to be one of the central topics in science and technology.

Written for the May 2004 EU meeting in Brussels on challenging problems for the 7th Framework Future Emerging Technologies Program.