Intelligent Systems: Perspectives and Research Challenges

Rajendra Akerkar

Technomathematics Research Foundation

TMRF Report 10- 2012

Intelligent Systems

Perspectives and Research Challenges

TMRF-report-10-2012

TMRF White Paper

By Rajendra Akerkar

October 2012

Contents

Section 1

Summary

The human brain is an incredible system. It is a very sophisticated processing unit – the brain – that can perform some amazing tasks. The brain lets us to accomplish several complex tasks, such as recognizing individuals we know, performing hard computations, and making critical decisions based on experiences. In spite of their speed and memory capacity, today's computers struggle to emulate human brain – still computers lack the ability to perceive, reason, and learn as well as we do. The branch of computer science known as Artificial Intelligence tries to narrow the gap.

AI has been focused on narrowing the gap between human brains and computers by endeavouring to develop machines with the ability to act intelligently. Sometimes the brain does not do the action immediately but uses its imagination. It selects a response rule and determines what situation results from the action. Then it selects again an action for this new situation and determines the probable result. Thus it can choose not only one response rule but a complete plan of action. This allows us to create *intelligent systems,* which operate autonomously, interact naturally with their environment and the humans therein, and be adaptive to changing situations and contexts, including the user's preferences and needs. Examples of such intelligent systems in operation today include mobile devices that can translate and interpret foreign languages, a social emotional *bot* as an edutainment tool, systems that supports the selection, configuration and operation of strategies and tools in the bioinformatics, and machines that can automatically analyse medical images such as CAT scans to discover tumours or bone fractures. However, popular attempts in creating intelligent systems are still largely restricted to systems designed for environments having limited scope and performing simple tasks. In the future, research efforts must be devoted to intense *cognitive* challenges which are measurable and scalable in open ended scenarios under changing conditions.

This report will be published in CSI Communications | December 2012

Section 2

Intelligence

The We define *intelligence* [1] as the competence of a system to achieve a target or sustain desired behaviour under conditions of uncertainty. Our definition is fundamentally based on the phenomenon of intelligence in biological systems where, one can say that intelligence helps them to deal with unpredictable changes in the environment.

Intelligent behaviour is demonstrated by artefacts and biological systems capable of achieving definite goals or sustaining anticipated behaviour under conditions of uncertainty even in feebly structured environments, for instance, a situation where an mobile robot must distinguish between a person and an equipment at a workplace in which it operates.

Intelligent behaviour can be characterized by clearly identifiable features as shown in the Table 1.

Table 1: Intelligent behaviour characterisation

Feature	**Capability**
Adaptability	To achieve specified goals or sustaining desired behaviour in an environment characterised by unpredictable external changes
Self-Maintenance	To maintain its own state of operational readiness
Communication	To exchange information with other systems
Autonomy	To act independently from other systems, including human operators
Learning	Being trained to carry out certain tasks
Self-Improvement	To improve its own future performance based on past performance combined with learning from other agents or human operators
Anticipation	To predict changes in its environment which may affect its operation
Goal-Seeking	To formulate and modifying tactical sub-goals with a view to achieving planned goals
Creativity	To generate new useful concepts, theories, testing methods and methodologies
Replication	To create replicas of itself

Let us try to understand some of the ways AI scientists and engineers have organized their programs to achieve intelligent behaviour. Some of them were inspired mainly by engineering and computational considerations and some by cognitive science in its attempt to model psychological data. Some were even influenced by ideas about how various brain regions function. Parallel operation is assumed in many of these architectures, even though it is often of the simulated variety.

Section 3

Three Basic Functions

To exhibit autonomous intelligent behaviour a machine must be capable of performing three fundamental functions named in as: Perception, Cognition and Execution (see Figure 1).

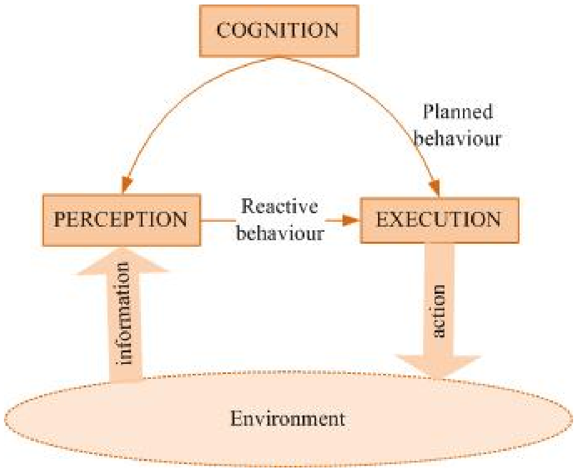

Figure 1: Functions that exhibit intelligent behaviour

The perception function provides information about the actual state of the system and its environment. The perception collects data about the world in which the system operates and processes collected data with a view to gathering reliable information to take decisions on future system behaviour.
The *cognition function* consists in planning an initiating the system's actions while taking into account information provided by perception.

The *execution function* has the role of initiating, controlling, handling and terminating the system's actions, based on instructions received from cognition and perception systems.

Obviously there are many promising ways of organising these functions to achieve autonomous intelligent behaviour. Conceptually the most simple is a centralised architecture with perception, cognition and execution functions implemented as separate but interconnected subsystems. However, from the engineering point of view a centralised architecture is not feasible. For example, the complexity of a centralised perception subsystem for an intelligent workplace would be hard to imagine. For an autonomous vehicle working in a workplace such architecture is not truly practical. Centralised architectures are on the way out even in decision support systems, which are less complex since they do not have to process sensory data.

The usual approach to reducing complexity is to adopt a multi-level hierarchical architecture with perception, cognition and execution functions distributed at various levels of the hierarchy. Many systems of this kind are under development. However, hierarchies have a major disadvantage and that is their rigidity. Evidence is mounting that hierarchies are not suitable for worlds characterised by frequent changes.

A number of very successful prototypes of intelligent systems have been constructed using the so called layered architecture.

Section 4

Intelligent Systems

Intelligent systems are concerned with the theories and techniques for building computer systems which exhibit some form of intelligent behaviour. The area has had an active and exciting history and is now a relatively mature area of computer science. Many of the research discoveries have reached the point of industrial application and products, and many companies have made and saved millions of Rupees by exploiting the research results in this area. However, many challenging research problems remain. From the perspective of computation, the intelligence of a system can be characterized by its flexibility, adaptability, memory, learning, temporal dynamics, reasoning, and the ability to manage uncertain and imprecise information. In general, intelligent systems have to deal with sources of uncertainty [2], such as the occurrence of unexpected events, and uncertain – incomplete, inconsistent or defective – information available to the system for the purpose of deciding what action to be taken next.

At large, Artificial Intelligence (AI) comprises of two key directions.

1. Humanistic AI (HAI) that studies machines that think and act like humans. HAI is the art of creating systems that perform functions that require intelligence when performed by people. It is the study of how to make computers do things at which, at the moment, people are better.
2. Rationalistic AI (RAI) that examines machines that can be built on the understanding of intelligent human behaviour. RAI is a field of study that seeks to explain and emulate intelligent behaviour in terms of computational processes. It is the branch of computer science that is concerned with the automation of intelligent behaviour.

Intelligent systems as seen nowadays have more to do with rationalistic than with humanistic AI. Intelligent systems exhibit intelligent behaviour as seen in nature as a whole. In addition, intelligent systems are motivated by the need to solve complex problems with improving efficiencies.

We distinguish between two classes of intelligent systems in the following Table 2:

Table 2: Classes of intelligent systems

Class	Definition	Example
Intelligent decision support systems	AI programs that advise and support decision making via human-computer interfaces	Intelligent decision support in healthcare (e.g. ADEMA – system for Asthma healthcare)
Intelligent machine systems	Machines and interconnected machines with embedded AI which are capable of operating autonomously	Intelligent robots, cognitive systems (e.g. Chiara - open source educational robot)

4.1Three-Layer Architectures

The constituents of intelligent system, Shakey, were organized into a three-*layer* (high-, intermediate-, and low-level) architecture. The Shakey's programs and data can be assembled into levels. Shakey's perceptual and basic (motor) programs were embedded in the low-level actions, whereas the intermediate-level actions combined the low-level ones in various ways to perform certain common tasks. The high level was in charge of planning and overall execution of plans.

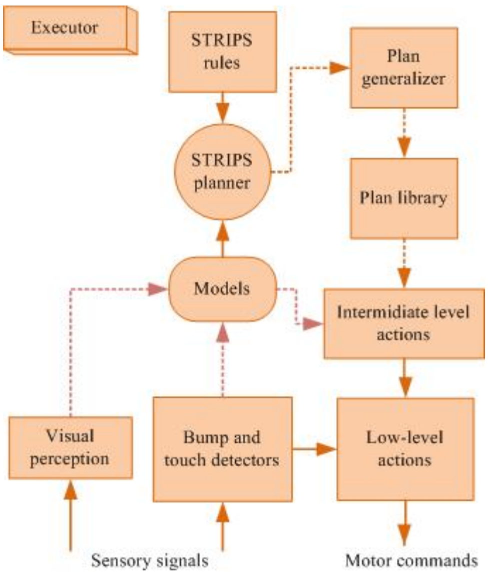

Figure 3: Shakey system

Low level :

- The low-level actions (LLAs) use a short and fast path from sensory signals to effectors.
- Important reflexes are handled by this pathway.
- Stop, Servo control of motors and so on

Intermediate level :

- Combine the LLAs into more complex behaviours
- Intermediate-level action (ILA)

High level :
Plan is expressed as a sequence of ILAs along with their preconditions and effects.

In Figure 3 interaction among programs in these levels is illustrated by connecting lines.
Three-layered architectures, such as the one used by Shakey, are used in several robot systems [3]. One example of a three-layered architecture is used in the German driverless “seeing passenger car" described by Ernst Dickmanns et al [4].

4.2Multi-layered Architectures
As an alternative to the three-layered schemes, Rodney Brooks [7] proposed architectures that controlled system actions in a way that reacted directly to changes in the environment without the need for planning. Initially called *subsumption architectures*, a way of decomposing one complex behaviour into many ‘simple’ layers of increasingly more abstract behaviours. Every layer can overrule (subsume) the decision of the overlaying layer. These architectures were later called *behaviour-based* because they were composed of programmed robot behaviours.
The different behaviours are arranged in levels, each responsive to its own set of environmental stimuli and each able to control the system depending on the sensed situation. Resulting overall behaviour of a system that is not explicitly represented in a computer system is known as “emergent behaviour."
Common principles present in behaviour-based architectures are shown in the Figure 4:

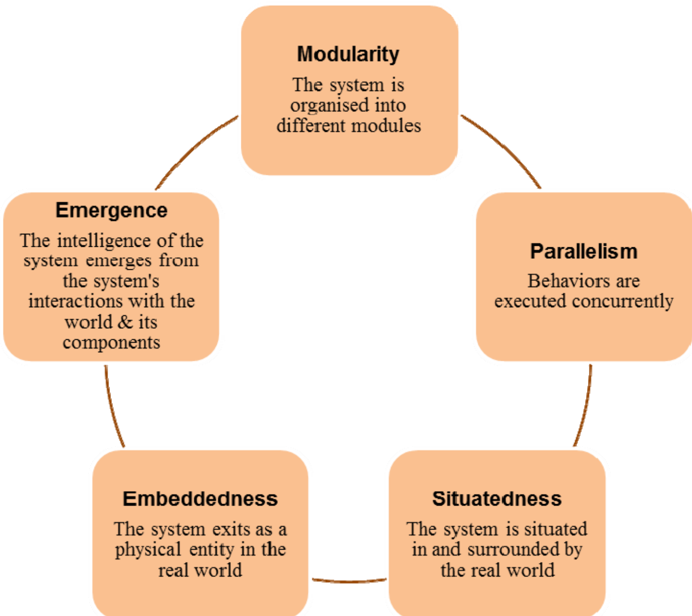

Figure 4: Fundamental principles in behaviour based architecture

Another popular way for achieving strategic behaviour in behaviour based architecture is to couple a behaviour based component to a deliberative system. The deliberative system activates behaviours in a similar way to higher layers in multi-layered system. There may be domains where hybrid systems are more suitable than pure behaviour based models and vice versa. A number of hybrid systems have been developed to overcome the perceived weaknesses of other architectures. Most hybrid architectures divide the system into layers, generally one layer for high-level planning and another for handling the details of interacting with the world. The upper layer is usually a reactive planning system and the lower layer a reactive or behaviour based system. Hybrid systems may, in general, be more complex than multi-layered system because they require two different types of architecture whereas multi-layered system uses multiple versions of the one architecture. The integration of two architectures implies that there will be issues to resolve in order to get effective behaviour.

4.3The Beliefs, Desires, and Intentions (BDI) Architecture

Michael Georgeff and Anand Rao [8] have proposed perhaps the first agent architecture based on the philosophical concepts of beliefs, desires, and intentions. This architecture is proved to be the most durable agent architecture developed to date.

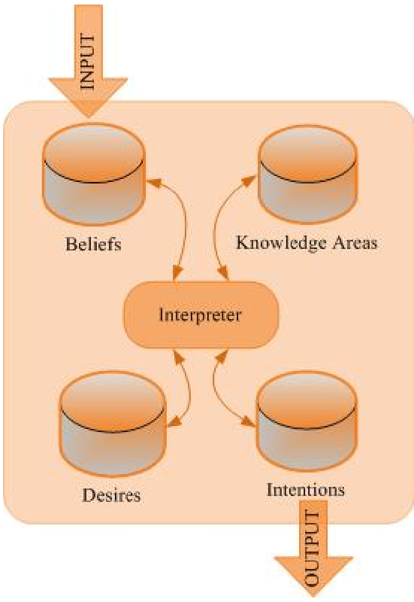

Figure 5: The Procedural Reasoning System

An agent's beliefs represent its knowledge about its environment (including itself and other agents), usually expressed in some kind of logical language, such as the first-order predicate calculus. An agent's desires represent the agent's goals – situations that it wants to achieve. An agent's intentions represent those desires that the agent has actually chosen to begin to achieve. That is, it has begun executing a plan to achieve them. BDI architectures, as distinct from behaviour-based, reactive ones for example, explicitly represent beliefs, desires, and intentions as actual data structures.

Here, knowledge of the environment is held as *beliefs* and the overall goals are *desires*. Together, these shape the *intentions*, i.e., selected options that the system commits itself toward achieving.

Readers may refer to [5, 6] to know more on how the architecture works.

4.4Architectures for Agents Alliances

An intelligent agent possesses some intelligence grounded on its *knowledge base*, *reasoning mechanisms* and *learning capabilities*. Depending on the assignment of a particular agent, there are differences in types of information contained in its knowledge base. However, generally this information can be divided into two parts – *the agent's knowledge about its principal contained in owner's profile* and *the agent's knowledge about its environment*. Intelligent agents exist in environments containing other intelligent agents, both humans and machines. Many of these agents collaborate or compete in the performance of their tasks. Agent-to-agent communication strategies and multi-agent architectures have become important intelligent systems topics.

A very attractive alternative is to assemble a system from a number of autonomous intelligent agents connected in a network and capable of collectively generating desirable system behaviour. Intelligent agents may be designed to operate in collectives, organisations similar to colonies of ants, in which every constituent element obeys precisely defined rules of collaboration, or in societies, organisations similar to human societies, in which artificial intelligent agents negotiate, collaborate or compete among themselves.

Section 5

Learning in Intelligent Systems

The learning process in intelligent systems involves acquiring information about its environment, and deploying the information to establish knowledge about the environment, and, consequently, generalizing the knowledge base so that it can handle uncertainty in the environment. A number of machine intelligence techniques have been developed to introduce learning in machines [9], e.g. imitation learning and reinforcement learning. In case of robot learning, a multi-learning method is being used. The divide and conquer rule is also applied to the learning tasks. Each algorithm is given a specific task to handle. The learning algorithms are chosen carefully after considering the characteristics of the specific task. Alternative latent solution to learning is intelligent agents. Agents collect data and learn about the surrounding environment, and adapt to it. The learning process in agents also requires a self-organizing mechanism to control a group of autonomous agents. The task of divulging learning into intelligent systems is not a simple; however the learning capability is what makes a system intelligent.

Section 6

Research Trends

Truly intelligent systems should be able to operate autonomously, interact naturally with their environment and the humans therein, and be adaptive to changing situations and contexts, incorporating the user's preferences and needs. Nowadays, an encouraging range of isolated elements in the area of intelligent systems is practicable, including learning, vision, speech, planning, control and decision making. However, the focus of these developments is predominantly on performance in well defined, limited domains. Popular attempts in creating artificial, intelligent systems are still largely restricted to systems designed for environments having limited scope and performing simple tasks. The ability to varying contexts and tasks without expensive redesign of specific, ad hoc solutions is still not realised.

New approaches are needed to intelligent system development, focusing more on understanding processes that lead to autonomous growth and development than on system development. We present some prominent research directions in the following subsections.

6.1Interactive Intelligent Agents

Cooperation between agents must be based on the principles of alignment, entrainment, imitation, sharing, anticipation and proactive interaction. There is a need for innovative theories of interaction to enable human-robot, human-human, and robot-robot interaction. The goal is to develop autonomous, interactive intelligent agents that operate within human environments and play a beneficial role in the daily lives of people. A key aspect in this field is multimodal interface technology, which allows humans and their environments to be perceived by recruiting signals from multiple audio-visual sensors.

6.2Bio-inspired Systems

The enhancement and emergence of cognition relies on artificial embodiments having rich perceptual and motor capabilities. Biologically inspired intelligent systems with such capabilities therefore represent the most appropriate experimental platform for studying cognition. Humanoid personal robots (see Figure 6) are examples of artificial cognitive systems and a key growth industry of the 21st century.

Figure 6: A humanoid robot

The great challenge is the advancement of robotic technology to the point where interactions between humans and robots run smoothly and robots are able to fulfil roles in the human living space. Another important aspect is to capture principles of collective operation, such as altruistic cooperation, dynamic division of labour and emerging communication that are applicable to a wide set of cognitive platforms and tasks.

It is therefore necessary to bring together biologists, control theorists and cognitive scientists to develop principles and algorithms that hold in the reality of specific cognitive systems (say, robots) and animals; and also are general enough to be easily applicable to new platforms.

6.3Intelligent Simulation Systems

For many activities, on-the-job training is very effective, providing the trainee with the chance to make real, quick decisions and see the consequences. Simulation systems that could portray realistic simulated worlds, such as the capability to produce realistic simulations of people, would enable development of training systems for such situations. Several commercial, military, educational, entertainment, and scientific applications need the capability of creating realistic simulated worlds. The systems we foresee here differ in both scale and function from those that exist today. The estimated scale of advance simulations is illustrated by the problem of providing accurate simulations of a crisis like 26/11 Mumbai terrorist attack that would be used in training crisis managers. Such simulations might require thousands of actors to play the role of victims, fire fighters, commandos, and emergency rescue squads. It might be economical to use actual people for only a few of these roles; the rest could be simulated. Intelligent simulation technology can assist people in such stressful, time-pressured situations to look further ahead in determining the consequences of proposed actions.

6.4Biological Processes & emergence

An important aspect for future intelligent systems is the notion of biological processes, called *morphogenetic* processes, which cause an organism to develop its shape, for information processing. This takes into account cooperation, stabilization, consolidation, focusing, categorization and mode selection. Autonomous, interactive and incremental learning and co-developmental approaches will be a key element in the development of processes for emergence. We, human beings, acquire sensorimotor knowledge since from our childhood. This kind of learning is the focus of developmental approaches, which have gained a lot of attention in intelligent systems community in recent years.

6.5Intelligent Decision Support

The enormous array of computational techniques and data available due to today's use of high-throughput technologies can be quite overwhelming for researchers investigating scientific problems. For any problem, there are many possible models and algorithms giving different results. New (innovative) intelligent systems need to be created that supports the selection, configuration and operation of strategies and tools in the scientific or engineering domain. For example, a system that guides researchers in building a data analysis workflow and acts as an interface by rendering transparent details regarding the implementation of the tools proposed or the configuration of on-line services. This type of system should thus be an edge between traditional decision support systems (DSS) and workflow management systems (WFMS). The knowledge, in such systems, comprises expertise on the application domain, which is composed of heuristics and strategies derived from domain literature and experiments and/or provided by one or more human experts.

6.6Intelligence and embedded reasoning

The Intelligence in embedded systems requires systems of sensors, actuators, and processors to be adaptive, distributed, and robust. Often, a tight coupling with both the physical world and temporal requirements leads to challenges in real-time execution and in process and communication concurrency. Such systems must be able to understand their environments and act intelligently and often autonomously, with (noisy) sensor inputs and imperfect models of system behaviour. There is a need to consider the associated programming paradigms and approaches to validation, and fault tolerance. These systems must be able to interact efficiently with each other and with their human operators. Embedded reasoning methods will transform various application areas including robotics, transportation systems, defence, and industrial automation. Applications range from battlefield robots to aerospace vehicles, from prognostics to factory floors, and from printers to medical devices.

Section 7

Research Challenges

Present-day challenges in the field of intelligent systems can be summarized as follows:

- Systems must be able to foresee body dynamics of the world and, thus, develop ability to reason about it.
- Novel approaches supporting the learning of new skills, adaptation of existing skills and the ability to switch between different learning modalities. The pertinent architectures and practices should allow for capability learning as well as autonomous and cooperative skill and strategy transfer to varying contexts and tasks.
- New frameworks and models for the representation and organization of enormous knowledge for complex sensory-motor control, choice and combination of actions for handling with everyday situations.
- Approaches boosting learning, recognition and classification of objects and events.
- Methodologies supporting the development of systems that investigate their own sensorimotor primitives, body morphology, the environment and their effective interaction with it.
- Distributed complex platforms with standard or open software, which allow researchers from different fields to assess their concepts and provide a framework for the benchmarking of various algorithms.
- Efficient cognitive frameworks that allow the integration of perception, action, reasoning, learning and communication components.

Section 8

Conclusions

India has the potential to play a leading role in the analysis and design of future intelligent systems. However, the existing expertise in information theory, biological sciences and social sciences can be more intensely bundled to provide better theoretical foundations towards understanding the processes and underlying mechanisms on which intelligence builds. Let us hope that through extensive international collaboration, a fruitful synthesis will emerge, giving technical oriented scientists new inspiration from biology and providing AI scientists with new ways to prove and evaluate their biological models.

Undoubtedly, substantial progress can only be made through a demanding dialogue among researchers from the fields of natural and artificial intelligence.

Section 9

Bibliography

1) Rajendra Akerkar, “Introduction to Artificial Intelligence”, PHI Learning, New Delhi, 2005.
2) Rajendra Akerkar and Priti Sajja, “Knowledge Based systems”, Jones & Bartlett, 2010.
3) Erann Gat, “On Three-Layer Architectures," in David Kortenkamp, R. Peter Bonnasso, and Robin Murphy (eds.), pp. 195{210, Cambridge, MA: MIT Press, 1998.
4) Ernst D. Dickmanns et al., “The Seeing Passenger Car `VaMoRs-P,'" Proceedings of the IEEE 1994 Symposium on Intelligent Vehicles, pp. 68-73, 24-26 October 1994.
5) Michael P. Georgeff and F. F. Ingrand, “Decision-Making in an Embedded Reasoning System," Proceedings of the 11th International Joint Conference on Artificial Intelligence, pp. 972-978, 1989.
6) David Kinny, Michael P. Georgeff, and Anand S. Rao. A Methodology and Modelling Technique for Systems of BDI Agents. In Proceedings of MAAMAW 1996, pages 56--71, 1996.
7) R. A Brooks (1991). "Intelligence without Representation", Artificial Intelligence 47 (1991) 139-159.
8) Anand S. Rao and Michael P. Georgeff, BDI Agents: From Theory to Practice. In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pages 312-319, June 1995.
9) Stuart J. Russell, Peter Norvig: Artificial Intelligence - A Modern Approach (3. internat. ed.). Pearson Education, 2010.