Artificial General Intelligene via Finite Covering with Learning

Yong K. Hwang, Samuel B. Hwang and David B. Hwang
Think Life, Inc.
Palo Alto, CA, USA

Abstract. This position paper claims that the combination of solutions to a finite collection of problem instances and an expansion capability of those solutions to similar problems is enough to achieve the artificial general intelligence comparable to the human intelligence. Learning takes place during expansion of existing solutions using various methods such as trial and error, generalization, case-based reasoning, etc. This paper also looks into the amount of innate problem solving capability an artificial agent must have and the difficulty of the tasks the agent is expected to solve. To illustrate our claim examples in robotics are used where tasks are physical movements of the agent and objects in its environment.

Keywords. human intelligence, machine learning, knowledge expansion, robotics

Introduction

Achieving human-level intelligence has been an illusive goal due to the vast amount of knowledge humans build up and use to solve many different tasks. Since it is impossible to preprogram solutions to all the problems an agent will face in the real world, many attempts are made to endow the agent with learning capabilities so that it can increase its knowledge base from a relatively small knowledge base. This approach has some success in a well controlled environment, but cannot yet handle complex problems arising in the real-world situations like humans can.

Although humans can solve many problems, they are functional only in familiar environments. If a human is put in a world with irregular gravity he will have a hard time navigating the world. It is our claim that the kinds of worlds humans can function efficiently are not that many compared to the set of all possible variations of the world. From this claim follows that a successful AGI can be built with a finite number of experiences. A positive (negative) experience is a pair of a problem and a successful (unsuccessful) action for it. One more requirement is capability of generalizing or expanding an experience to similar situations. The resulting knowledge base is then capable of handling a subset of all possible worlds, where the subset is a union of blobs around the points representing positive experiences minus the blobs representing negative experiences (see Fig. 1). This subset is a finite covering of the world in which the AGI system can function well to suit its purposes and avoid failures.

This paper concentrates on the task planning in the physical world, which occupies a middle of the intelligence hierarchy [1]. We now present the process of building up a knowledge base by expansion capability to generate a finite covering of the world. Examples in robotics are used to illustrate this process.

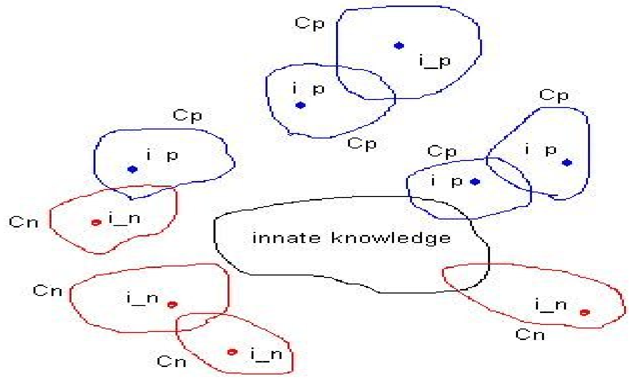

Fig. 1. The knowledge set for solving problems is the union of neighborhoods of positive experiences while the knowledge about failure is that of negative experiences.

1. Knowledge Representation

Our agent starts out with an innate knowledge, and expands it via generalization and learning. Our main claim is that an adequate knowledge base for agent's well being can be built as a *finite, i.e., manageable*, union of problem-solution pairs and that of problem-failure pairs. A problem, *q*, is defined as the pair of agent's current and desired states of the agent and its environments, *q(e_current, e_desired)*. A positive experience instance, *i_p* is a pair of a problem and an agent's action resulting in a success, written as *i_p(q, a_p)*. Similarly, a negative experience instance is written as i*_n(q, a_n)*. When the agent has a capability to recognize a problem, *q~*, similar to *q* and a capability to modify the action *a_p* to solve *q~*, the set of problems near *q* solvable by the agent form a set covering *q*, written as *Cp(i_p)*. Similarly, the set of problem instances around *i_n* resulting in an undesirable outcome is called *Cn(i_n)*. Let *Kp* and *Kn* be the knowledge sets of the positive experience and negative experience, respectively. *Kp* (*Kn*) is a union of *Cp*'s (*Cn*'s). Then the set of problem instances solvable by the agent with 100% confidence is *Kp – Kn*. If we are willing to accept some failure, the agent might try to solve problems belonging to both *Kp* and *Kn*.

Let us give an example. In robotics field, two of the fundamental tasks are robot navigation and object manipulation. In navigation, the problem is to find a short, collision-free path from the current configuration of the robot to the desired configuration. The *e_current* (*e_desired*) is specified by the current (desired) robot configuration along with configurations of all objects. In manipulation, *e_current* (*e_desired*) is specified by the current (desired) configuration of the object to be manipulated along with those of nearby objects. Suppose the robot has a path-planning algorithm, *a_pathplan* to find a collision-free path. An arrangement of objects in its environment and robot's start/goal configuration form a problem instance *q(e_start, e_goal)*. If *a_planpath* successfully finds a solution path, then the pair *(q, a_pathplan)* forms a positive experience, *i_p*. Otherwise, it forms a negative experience, *i_n*. If in the same environment an object irrelevant to the solution path is moved away from the path, then *a_palnpath* will succeed again. Then the new experience forms a part of the positive covering around *i_p*, namely *Cp(i_p)*.

2. Expansion of Knowledge

As an agent solves problems, it expands its knowledge base in several ways. The trial-and-error method is useful when the penalty for failure is small. Another method is to learn from a teacher, whether a human or other agents. Other machine learning algorithms can be used to expand knowledge [2]. These processes increase the sets of positive and negative experiences whose sizes are of measure zero, i.e., they are just points in the space of all problems. If we have some local expansion capability, the agent can grow the point experiences into sets of finite measures, vastly increasing the size of the problems it can solve. This expansion capability enables us to claim that not infinite, but only a finite number of experiences are enough for agent's successful life. The increase of knowledge occurs in two categories: by solving new problems, and by expanding an existing piece of knowledge for a particular problem to similar problems. Let's call these global and local learning. These are elaborated below.

2.1. Global Learning

Simply put, a global learning takes place when the agent learns a *brand new* stuff. Suppose the agent learns to use a knife to cut objects, which it does not know how before. This is a kind of global learning. Another example is learning to walk in a stair case when the agent could walk only on a flat floor. When a global learning occurs, inserted in the knowledge base is a new pair, *(p, a),* of a problem, *p,* represented as the current and desired environmental and an algorithm, *a,* to solve the problem. A global learning can be achieved by various means. A trial-and-error method is simple but may have a low success rate. A more prudent way is to learn from a teacher, which can be humans or observation of successful task completion by other agents. Another method is divide-and-conquer scheme which decomposes a task into subtasks that are solvable by existing algorithms. Learning the category of grasp such as power grip or precision grip [3] can also be classified as global learning. In terms of our notation the global learning increases the number of coverings *Cp*'s or *Cn*'s in the agent's knowledge base.

2.2. Local Learning

A local learning occurs when a new problem is solved by an existing algorithm with substitution of input parameters. For paper cutting task, for example, if the agent learns to use the algorithm of using scissors for cutting a string then a local learning has taken place (substitution of object). Local learning can be achieved by various methods. A trial-and-error method can be used to apply an existing algorithm to tasks with different objects. Or the case-based reasoning can be used to recognize similar situations and the agent can use the action that succeeded. Local learning increases the size of covering, *Cp* or *Cn*. If the algorithm solves (fails to solve) the problem, the size of *Cp* (*Cn*) is increased. The key difference between global and local learning is whether a new algorithm is introduced in the knowledge base.

2.3. Quantum Learning

The most challenging type of learning is the automatic generation of new algorithms based on existing knowledge. It is at a higher level of learning than global learning, and for this reason we call it quantum learning. Suppose our agent know about depth-first and breath-first search algorithms. How do you make the agent generate A* search on its own? We have to somehow embed in the knowledge the principles of “best first” exploration, “possible option” retention and “hopeless option” deletion. One way of quantum-learning is maintaining many small knowledge pools as in Society of Mind [4], where a particular combination of pools forms a new algorithm. This is not in the scope of our paper, and we are just pointing out the difficulty of quantum learning.

3. Complexity of Building Artificial General Intelligence

We explore the amount of knowledge required to build human-level AGI. We don’t claim our list to be anywhere near a complete set, but it will give a good starting point for necessary components required by AGI systems. We first give a set of physical principles that reduces the amount of required innate knowledge and the complexities of the algorithms for problem solving. Presented next is the minimal requirement on the components of the innate knowledge from which to build a comprehensive AGI.

3.1. Physical Principles

We have come up with several principles based on the physical laws that can be used to simplify the complexity of AGI. They are proximity-in-space (PIS), proximity-in-time (PIT) and benign-environment (BE) principle. The PIS was used in [5] for generating natural-language-like robot commands. These principles are explained in detail below.

Let us explain by analogy how PIS makes problems in the physical world simpler in practice. A lone ninja fighting against many enemy fighters needs to deal with only 4 or so enemies at a time because the rest of enemies cannot get to the ninja. By the same token, although it is theoretically true that all objects need to be considered when planning collision-free motion in navigation or manipulation tasks, only several nearby relevant objects need to be taken account. Nailing two pieces of wood together requires consideration of only the wooden pieces, a nail and a hammer.

The PIT assumes that events relevant to the current problem occur in proximal time unless explicitly specified by some means. One of the basic assumptions in human reasoning is that things do not change too fast. If they do, for example, such as in a car accident, humans also make mistakes. When planning an action, it is prudent to consider events occurred in the near past or going to happen in the near future with more emphasis. To solve a crime, for example, detectives investigate people who visited the crime scene near the time of crime. The PIT, therefore, can greatly simplify the complexities of planning algorithms necessary for human-level intelligence.

The BE asserts that the environment is mostly benign and is not there to make the agent fail. Humans set up their habitat with favorable conditions of weather, food resources, etc. Furthermore, they build their houses so that everyday tasks can be performed with ease. We also assume objects are safe to interact with unless explicitly specified, such as “fire is hot.” Also, most knives are made with a handle to facilitate safe grasp. The default assumption of safety greatly reduces the amount of knowledge to be stored. This is true for many other object properties such as strength, weight, etc. Another example is the existence of hallways in buildings to expedite navigation. Yet another example is the fact that most tasks humans do can be performed with one or two hands. The development of AGI, thus, should be easier for manmade environments.

3.2. Components and Size of Knowledge Base

Humans are visionaries, i.e., their primary sensor is the eyes. To build an agent that interacts with humans, the agent must have an ensemble of good vision algorithms. Tactile sensing is of a great help in object manipulation. To recognize similar situations and similar objects, a similarity recognizer is needed. To detect changes in objects or environments, a difference operator is needed. Search algorithms are also required to find optimal solutions to various tasks. They include findpath, findplace (to put an object), findgrasp (to grab an object) to name a few. To predict the effect of an action during task planning, a qualitative simulator [6] is also necessary. Also needed are a database to store object properties and many other information such as pre- and post conditions of actions, goodness evaluation functions of situations/actions, and so on.

Now the important question is "how big is the knowledge base that is sufficient for AGI?" For home-service robots, 1000 actions seem sufficient [7]. The number of objects involved in an action is small, like 4, according to BE. Given a task, the action planning process involves a search in the space of agent actions. If we provide a template of an action sequence for each task, a search is triggered to find an alternate action only when one of the actions in the task template results in a failure. One of the main problems with the traditional reasoning system is the combinatorial explosion. A counter argument to this is the small-world phenomenon or six degrees of separation. It states that there is a relatively short chain of acquaintances to link two people in the world. By analogy, we claim that there is a short chain of actions to go from the current to a desired situation. Or relevant objects to solving a problem are linked to the agent or current situation in a few levels of depth in the search tree for a situation.

4. Conclusions

We claim in this paper that there is a hope for building a human-level AGI. This is conjectured from the fact that 1) even humans are not functioning effectively in an unfamiliar environment, and 2) by employing knowledge expansion methods a finite number of experiences can be beefed up to solve enough problems for an artificial agent. It remains to be seen that our claim is true by actually developing an AGI system. It is our belief that a system with human-level AGI is in the near future.

References

[1] M. Minsky, P. Singh, and A. Sloman, The St. Thomas common sense symposium: designing architectures for human-level intelligence, AI Magazine, Summer 2004, 25(2):113-124.

[2] T. Mitchell, Machine Learning, McGraw-Hill Education (ISE Editions), 1st edition, October 1997.

[3] M.R. Cutkosky, "On grasp choice, grasp models and the design of hands for manufacturing tasks," IEEE Trans. On Robotics and Automation, vol. 5, no. 3, pp. 269-279, June 1989.

[4] M. Minsky, Society of Mind, Simon & Schuster, March 1988.

[5] P.C. Chen and Y.K. Hwang, "Generating natural language-like robot motion commands through machine learning," IEEE Int. Conf. on System, Man and Cybernetics, vol. 1, pp. 114-121, Oct. 1994.

[6] B. Kuiper, 2001. Qualitative simulation. In R. A. Meyers (Ed.), Encyclopedia of Physical Science and Technology, Third Edition, NY: Academic Press.

[7] Y.K. Hwang, M. Lee and D. Lee, "Robots' role in the ubiquitous computing household environment," Proceedings of the International Symposium on Robotics, Paris, France, March 2004.