

Representation = Grounded Information

Mary-Anne Williams

Innovation and Enterprise Research Laboratory
University of Technology, Sydney
NSW 2007, Australia
Mary-Anne@it.uts.edu.au

Abstract. The grounding problem remains one of the most fundamental issues in the field of Artificial Intelligence. We argue that representations are grounded information and that an intelligent system should be able to make and manage its own representations. A perusal of the literature reveals much confusion and little progress in understanding the grounding problem. In this paper we delineate between information and representation where a representation is grounded information; as a result we provide much needed clarity and a new base from which to conduct an innovative analysis of grounding that delivers a novel and insightful understanding of intelligence that can be used to guide and inform the design and construction of robust autonomous intelligent systems with intention and that know what they are doing and why they are doing it.

Keywords: Robotics, Perception, Knowledge Representation, Intelligent Systems.

1 Introduction

Knowledge representation is a key pillar of knowledge management, artificial intelligence, cognitive science and software engineering. The field of Artificial Intelligence (AI) explores ways to build representations of information from the world's richness and to manage these representations over time for a range of purposes from decision making to actuation. Like humans, a truly intelligent system should be able to build its own representations. This paper argues that the field of AI should devote more effort and resources to the development of a practical understanding of how to design systems that can build their own representations, possess intention and know what they are doing and why.

Despite its fundamental importance there are few formal frameworks for exploring and evaluating knowledge representations. Even worse there is no common agreement on what a representation actually is or should be, not even an informal description of a representation. Terms like information, representation, and models are often used interchangeably which only fuels the confusion. In this paper we distinguish information from representation, and classify models as either informational or representational.

In the early 1990s Brooks challenged the prevailing assumption that representations are needed for intelligence. His work generated much lively debate and argument. His so-called behavior-based architecture is founded on the idea that a system

need not build and maintain models of the world because the world is its own best model and an intelligent system can simply sense the world directly, instead of investing in building and maintaining a so called world model [2, 3]. Despite giving the status quo a shakeup and putting robot behavior and experience alongside rationality on the research agenda, Brooks' approach has not lead to major advances or scientific breakthroughs. Furthermore, whilst his robot designs were not based on logical representations they were not devoid of symbolic representation e.g. they used entities in computer memory such as numbers.

In this paper we put forward the proposal that representations are grounded information, and that human intelligence stems from our astounding ability to build representations, i.e. to ground information. In fact, we ground most information so effortlessly that we take the process and the capability to do so for granted, and as a result we often conflate information, models and representations, and even worse our *perception of reality* with *reality* itself. When we try to build artificially intelligent systems like robots and design their grounding capabilities we come up against the complexity and challenge that the grounding problem presents. This paper takes an important step towards clarifying the underlying concepts and terminology that we can use to advance our understanding of the grounding problem.

AI researchers have designed and built impressive robotic systems that demonstrate complex skills and that can perform sophisticated tasks; however these systems are limited in their capacity for autonomous representation construction, because the robot's mechanisms to ground information are carefully crafted by system designers and the robots operate with significant human assistance and intervention over their whole lifetime. Few robots can build representations on the fly by themselves. As a result most systems perform poorly in novel or unexpected situations because they cannot create effective representations of their new experiences.

The idea that representations are grounded information provides a crucial distinction between the basic concepts of *representation* and *information*. We believe that the conflation of these concepts has lead to unnecessary confusion and a lack of clarity in the literature and the grounding debate, and as a result our understanding of grounding has been hampered and led to a lack of progress towards the development of genuinely intelligent systems that know what they are doing and that can build their own representations autonomously via their own experiences, rather than relying on human designers to interpret the world and craft their representations for them.

It is fair to say that the reason logic-based methods have come to dominate the field of Knowledge Representation is because they provide mechanisms that help "make sense" of information by providing guidelines that minimize the number of interpretations, the number of representations that can be legally constructed, or provide a truth-based semantics which can be used to interpret information using a straightforward procedure. In addition significant strides have been made in areas like databases and the Semantic Web because the representation-information dichotomy underlies conceptual schema that embed "meaning" into information so that it can be made sense of and converted to representations during subsequent processing. In these highly constrained relatively simple systems there is often a 1:1 relationship between information and representation, but this relationship is lost when we move outside simple domains into more open, complex and dynamic environments of the kind that mobile robots inhabit. There is little work that explicitly recognizes the difference between information and representation in the field of robotics.

The diagram shows the Schrödinger equation:
$$\frac{\partial^2 \psi}{\partial x^2} + \frac{8\pi^2 m}{h^2} (E - V) \psi = 0$$
 with arrows pointing to various parts: 'Second derivative with respect to X' points to the first term; 'Shrodinger Wave Function' points to ψ ; 'Position' points to x ; 'Energy' points to E ; and 'Potential Energy' points to V .

Fig. 1. Schrödinger's Equation

To emphasize the need to recognize and address the issues that the information-representation dichotomy raises consider Schrödinger's Equation (in Fig. 1.), the source of a rich representational model that predicts the future behavior of a dynamic system; written on a piece of paper it is merely a source of information. In most of the literature no distinction is made between the model as information in the external world and the model as representation in the mind of a cognitive agent like a person. There are many levels, perspectives and lens through which we can ground the information encapsulated in Figure 1; the pixel level, the character level, the general mathematical lens, or the physics lens. The information conveyed in Figure 1 would almost certainly be grounded differently and therefore converted into different representations, by a child, an AI researcher, a mathematician, a chemist, a physicist, an historian, and a cat.

2 The Grounding Problem in Retrospect

In this section we provide a brief history of the progress towards an understanding of grounding, and we argue that the next best constructive step is the development of a framework based on well defined concepts which can be used to design and analyse grounding capabilities, and to explicitly compare grounding across systems and agents. The key purpose being to develop cognitive agents, like robots, that can build and manage their own representations all by themselves, i.e. autonomously.

Grounding has been occurring at least as long as there have been living systems on the Earth. One could take the view that evolution has been busy at work pursuing the challenge of inventing increasingly better designs and technologies for grounding, which so far has culminated in the human body and mind.

In contrast to Brooks we consider cockroaches to have representations that help them determine how to respond to information captured by their senses, e.g. light, chemicals, and physical surfaces. There is no evidence that cockroaches develop and maintain world models, in the same way that people or a robot does, however, they do have representations that allow them to interpret information impinging on their senses so as to support their behaviour and survival. Even simple nematodes such as *Caenorhabditis Elegans* have a rudimentary memory which they use to represent chemical gradients and to control chemotaxis [8]. Clearly the *better* a systems' grounding capability relative to its environment, the more successful it will be and the more likely it will multiply and pass its DNA to the next generation. Evolution rewards and propagates well grounded systems over poorly grounded systems where a

well grounded system has or creates high quality representations that help it make sense of its inner and outer worlds [8]. If a system has a memory, it probably has representations that need to be grounded. Representations can take a wide range of forms, we are not suggesting they are necessarily symbolic, and our approach makes no particular assumptions about the form or content of representations. Our approach can be used to analyse grounding capabilities of cockroaches just as well as the Mars Rovers, Spirit and Opportunity, or to compare a chimpanzee's grounding capabilities with that of a human's [22].

The Ancient Greeks were fascinated by grounding; they pondered the relationship between ideas and reality incessantly and obsessively. We have made some progress in understanding grounding over the ensuing millennia but a thorough and comprehensive understanding still alludes us today. Socrates, Plato, and Aristotle were interested in the relationship between physical reality and the reality of ideas; they gave ideas the highest value and saw physical reality as the imperfect reflection of ideal reality. From this base scientists and philosophers across the ages have explored grounding. Undoubtedly, Galileo had a sophisticated and unique understanding of the interplay between concrete abstract ideas and reality due to his broad knowledge and experience in mathematics, theoretical physics, experimental physics, and astronomy. So it is a great pity that he lived during the Roman Inquisition which meant that he had to obfuscate many of his theories and thoughts on such matters. A younger contemporary Rene Descartes, however, was able to propose a dualistic' understanding of reality, and later Leibnitz considered reality as mental phenomena and insightfully wrote "... matter and motion are not substances or things as much as they are the phenomena of perceivers, the reality of which is situated in the harmony of the perceivers with themselves (at different times) and with other perceivers."

From a more formal perspective Tarski invented a theory of truth based on syntax and semantics [17], and for his system *truth is reality*. Peirce wanted more and developed his theory of semiotics with the extra feature of a referent. According to Davidson, a Tarskian *theory of truth* for a language is a *theory of meaning* for that language. The Deflationists believe that to assert a statement is true is the same as asserting the statement itself [18].

Albert Einstein, like Leibnitz, was also a phenomenalist and playfully stated "Reality is merely an illusion, albeit a very persistent one." Feynman viewed the world as teeming with information like energised waves and particles whizzing around us of which only a tiny percentage is captured by our sensors. He often noted that there was a close relationship between the laws of physics and reality, and that the laws of physics were only ever mere conjectures, since they could never be proved correct, only proven false when contradicted by experimental evidence. One of the more inventive approaches to describing reality is the Clifford Pickover's Hawking Reality Scale [14] where reality as people experience it is 0 on the exponential Reality Scale, and successive 10 units on the scale account for a reality one hundred times more distorted than the last. Hawking himself, continues to maintain that something as incomprehensible as quantum mechanics, "should concern us all, because it is a completely different picture of the physical universe and of reality itself".

Searle reinvigorated the grounding debate in Artificial Intelligence with the now well known and provocative Chinese Room thought experiment [15] where a non-Chinese speaking man trapped in a room, answers questions in Chinese by following

stepwise instructions with access to a Chinese questions and answers database. Apart from the fact that, such a database does not exist, Searle convincingly demonstrated that the man gave the impression he understood Chinese, but in actual fact he was unable to attribute meaning to the questions encoded in Chinese directly. Instead he was merely following instructions and had absolutely no understanding of the content of the questions or the answers. It is noteworthy that although he was unable to attribute the same meaning to the Chinese characters as say a Mandarin speaker might, he was able to ground the characters in order to match the correct response.

Brooks [2, 3] attacked the high level symbolic approach, the dominant approach at the time, in an attempt to make a case for his layered activity based bottom-up approach to intelligence. According to Brooks “the world is its own best model” [3]. He is a strong proponent of the need to ground systems in the physical world, and he continues to maintain that it is only through physical grounding that a system can give meaning to its processing. Harnad introduced the term “symbolic grounding” and appears to have taken the debate off in the direction of symbolic processing and something he calls “intrinsic meaning” [10]. Steels [19] has argued that labelling the grounding problem, the symbol grounding problem, has led to widespread confusion because there is no unique or widely accepted interpretation of “symbol” and a broad range of interpretations span researchers from AI to Linguistics. In particular, the nonsymbolic systems described by Harnad, namely neural nets, are implemented as computer programs and one could argue that they are as symbolic as a logic-based knowledge base or an Arabic dictionary.

Rapid recent progress in the field of robotics has raised the grounding problem once more [5, 9], since intelligent robots need to ground information and build representations autonomously. The main challenge is how to discover how to make representations from information so that we can build artificially intelligent systems that can do it. For humans, a representation is the result of grounding information where grounding involves *making sense* of information. Not only do humans make sense of raw information from sensors, but we are aware we do it, and are even conscious of some of our representation making. Unfortunately, we are not typically aware of *how* we do it and as a result trying to replicate it in artificial systems is a challenge.

The importance of the distinction between representation and information is cleverly illustrated by the Belgian Surrealist René Magritte in his painting reproduced in Figure 2 below. First the object of interest is the image/picture on the page which is just a bunch of pixels/ink dots organized in a certain way. Consider a computer-based photocopy machine which is able to capture the information in the image and reproduce it with a high degree of fidelity, and able to make an almost exact copy. Should we say it can make representations? It has a form of memory/storage where the captured information takes on meaning for the machine, so it seems reasonable to suggest that a modern photocopier can represent attributes of pixels in an image. At the other end of the spectrum people can trivially “make sense” of the pixels and “see” an object, some might recognize as a pipe, some can represent and read the French text “*Ceci n'est pas une pipe*”. Some people may ignore the text, others try to interpret it; some may conclude it is a mistake, a contradiction, a joke, or an allusion, and others may see truth in the statement, after all a picture of a pipe, is not a pipe. Presented with this picture reflected photons impinge on the rods and cones in our further processing in the brain to create a conscious representation.



Fig. 2. Ceci n'est pas une pipe. (This is not a pipe) by René Magritte 1926.

The information continuously captured by our senses is used to produce our perceptions, concepts and knowledge to produce rich and complex representations, and we are only conscious of some of it, mostly the end product of all the complex processing. We are constantly bombarded with information; it impinges on our senses in steady streams and fluctuating bursts. Our grounding capability meticulously finds patterns in the information streaming in through our senses, and we carefully (and magically) massage the information we capture from multiple sensors into (meaningful) representations of physical and abstract things that we are able to subsequently comprehend and interpret.

There is a significant difference between information and representation; one can think of them as the input and output to a grounding process which for many applications could be conceived as a service and, therefore, it makes sense to consider the quality of a grounding service. Williams *et al* develop a Grounding Framework for evaluating how well a system grounds its information in [22].

3 Information and Representation

Information is non-representational; it creates the so-called raw feels that our senses detect internally and externally, and also particles/waves/energy we cannot detect with our limited human senses, e.g. radio waves. We are immersed in a sea of information and we are able *to sense* some directly like electromagnetic radiation in the so-called visible part of the spectrum, *to conceive and imagine* information beyond our senses such as radio waves and *to create* information in our own minds that we sometimes transpose for others to sense and make sense of such as the drawing of a dragon. We place no restrictions on information; it can be anything: a photon, a particle, a wave, a physical object, an abstract object, a class, an event, an action, relationship, or a representation.

Making sense of information creates representations. Making sense can involve recognizing simple patterns like a snake from information impinging on its heat sensors or a university chemistry student trying to understand the wave function, ψ , in Schrödinger's Equation. Representations for our purposes are open and the only requirement is that they ground some information. Representations can be generated from sensorimotor information, logic-based information, linguistic expressions, artefacts like a web page, a cave painting, a wristwatch, a hammer. Representations of physical objects are important but so too are representations of abstract entities such

as object functionalities and relationships between objects, and descriptions of ways to interact with specific objects such as affordances.

For the purpose of understanding the process of grounding it is insightful to classify representations into two important classes: *cued* and *detached* [6]. Cued representations are associated with the perception of things that are present, and detached representations are associated with entities that are not currently perceived. Sensations are immediate sensorimotor impressions, perceptions are grounded sensorimotor impressions, and simulations are detached representations. Sensations provide systems with information from the external (e.g. environment) and internal (e.g. body) world that a system can use to build representations and awareness. They exist in the present, are localised in the body/system, and are modality specific, e.g. visual, auditory, not both. Perceptions can ground more information than raw sensorimotor information, e.g. they can merge information from multiple sources to create multi-modal representations. They can represent accumulated sensorimotor information and sensorimotor information reinforced with simulations [6, 8]. Sensations are raw feels and involve signals from external sensors or from inside the system itself, but perceptions can use additional information derived from previous experiences and/or outcomes of learning. In contrast to sensation, perception can be cross-modal and can generate persistent representations which give rise to important cognitive capabilities such as object permanence. A detached representation is a simulation or abstract concept; something that can be utilized regardless of whether what it represents is present or not. A detached representation can even stand for something that does not exist at all. It turns out that humans readily use detached representations however other animals tend to use cued representations, and as a result they require a less responsive memory and they take fewer things into consideration when generating future possibilities during planning [8, 13] because they tend to plan with cued representations. A cognitive agent using detached representations will be more powerful in terms of explanation, prediction, planning, and collaboration provided they have high quality grounding capabilities and sufficient information processing power to build and manage them. Representations can be generated from information that has been gathered from a wide range of sources e.g. internal and external sensors, internal and external effectors, external instruments, and external systems. In addition they can result from fusing sensorimotor information with high level representations such as perceptions, concepts and linguistic expressions. Consider a doctor who not only grounds his own sensorimotor information, but information acquired from colleagues, books, lab tests, instruments that measure temperature, heart beat, blood pressure and oxygen content of the blood. Detached representations are extremely powerful. They can be manipulated independently of the external world, i.e. can be conceived and do not need to be directly perceived. Some examples of important detached representations are absent objects, past and potential future world states.

4 Grounding Capabilities

Grounding impacts how and what representations are created. For example, for a system to plan it needs representations of potential future needs and future world states [6, 8]. Chimpanzees can construct elaborate but limited plans where elements in the plan are usually constructed from cued representations of objects. Our current understanding suggest that chimpanzees do not tend to make representations of future

needs based on the observation that they do not tend to carry tools such as stones that can be used to crack nuts, from place to place [6].

We define the process of *grounding* to be the construction, management and interpretation of representations. Systems from airline reservations to autonomous mobile robots rely on well grounded representations [22]. A grounding capability provides basic infrastructure for cognition and intelligence. Consequently, what, how, and how well, information is grounded is of significant interest and crucial importance in AI. Grounding can be goal specific, domain specific, context specific, and system specific. Having defined grounding and grounded information we can focus on exploring the nature of the quality of the representation. A representation is created by a system's grounding capability, i.e. the capability to make sense of information.

Clearly, the ability to make representations is closely related to what we call intelligence. In order to build a truly intelligent system we need to ensure that once deployed the system can generate high quality representations on the fly without the assistance or direct intervention of humans, even in novel situations.

Cognitive agents manage representations about their internal states and the external world they inhabit; they use these representations to perform and achieve their design goals. As the complexity of the world they seek to exploit increases they have a growing need for more sophisticated methods to generate and manage representations effectively over time in the face of uncertainty, inconsistency, incompleteness, and change in complex and dynamic environments.

Significant work in the area of logic-based AI has been carried out in the quest to address these issues, however how these high level constructs are grounded so that systems can manage them autonomously over time in complex and dynamic environments remains an important and open challenge. Much of AI explores agent beliefs and many of the problems belief management faces stem from the conflation of the information and its corresponding representation.

In Intellectual Property Law ideas and their expression are treated entirely differently. This is known as the *idea-expression* dichotomy [20]. For example, one cannot copyright an *idea*, only the *expression of an idea*. In AI however, the expression of belief is often treated as the belief, but *isa(John,father)* is not a belief, in the same way the object depicted in Figure 1 is not a pipe. Beliefs are *ideas* while *isa(John,father)* and the pipe as depicted in Figure 2 are *expressions of ideas*. The distinction between a representation and information, or idea and expression, becomes crucially important when we try to design and build any intelligent system, not just AI systems on robots.

Having defined the process of grounding and an information-representation dichotomy, we can focus on the quality of the process of grounding. In order to be intelligent a system needs to be suitably grounded. Poorly grounded representations give rise to deluded systems that exhibit poor quality decision making, an inability to communicate and collaboration, and potentially dangerous behaviour. Representations are grounded information and are said to possess the property of *groundedness*. Groundedness is rich and complex, and it has a definitive impact on the possible attributes and behaviours that a system can display, and capabilities it can possess. Clearly, systems are not just grounded or ungrounded, but they can exhibit degrees of groundedness in a range of different dimensions. It makes sense to say of two grounded systems that one is more grounded than another.

5 Evaluating Grounding

The grounding process needs to be analysed and assessed on three important aspects: the construction, the management and the interpretation of representations. For representation construction we note that most contemporary artificially intelligent systems do not construct, or even manage their own representations, instead the mechanisms to do so are crafted by human designers; little of the construction and management mechanism is autonomous even in agents that use unsupervised machine learning. Once human designers deploy the mechanisms for systems to build, manage and interpret representations, systems can perform highly sophisticated tasks and their performance reflects the abilities of the designers.

A crucial ingredient for intelligence is the ability to make sense of information, in other words to ground information into representations. Even Brooksonian robots, like Genus Khan, attempt to make sense of the world and produce representations in this sense. Many aspects of intelligence rely on representations for example:

- Knowing what you know, e.g. *I know that my son was born in 1998*, what you don't know, e.g. *where my son is right now*, what you are doing e.g. *writing*, how you are doing it e.g. *on a computer*, why you are doing it e.g. *its fun*, what your goals are, e.g. *to convey innovative ideas*, how to achieve your goals, why they are your goals, etc.
- Knowing what others know, what others don't know, what others are doing (and how and why), what interests others and what will get their attention (and why).
- Anticipating future developments; intelligent agents need to anticipate events and eventualities by exploiting their ability to represent and rehearse the future; to conduct risk assessment and risk management.

Intelligence involves various abilities including intention, motivation, adaptability, reasoning, foresight, insight, reactivity, proactivity, creativity, problem solving, judgement, decision making, planning, strategizing, learning, empathy, communication, collaboration. The quality of these powers and abilities is directly dependent on the quality of the systems' grounding capabilities, the representations they make and their ability to interpret them. High quality grounding is crucial to intelligent behaviour; moreover the ability to reason about other agents grounding capabilities is crucial for meaningful communication, and effective collaboration.

Contemporary robots do not experience the world like us because their grounding capabilities are significantly different to our own. For example, as illustrated in Figure 3 a typical soccer playing AIBO lives in a colour coded world and does not have stereo vision, and as a result the its perception of the objects like the ball is as a 2D blob in an image, and not as a 3D sphere. It is important to note that even though the soccer playing robots live in a colour coded world colours are completely dependent on ambient lighting conditions so that lowering the ambient lighting makes the ball appear a darker colour and eventually not orange at all. People tend to conceptualise colour as an intrinsic property of objects but it is not intrinsic at all.

Our perceptions of reality are uniquely human, and when we design and analyse the grounding capabilities of other systems, like soccer playing robots, we often implicitly compare them to ourselves and fail to clearly identify the assumptions we make in that comparison. Reality is experienced and represented (perceived and

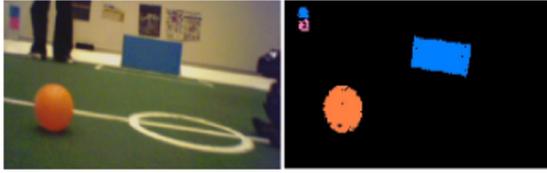


Fig. 3. (a) Raw image derived from a robots camera (b) Perceptual representation of the ball, beacon and

conceived) differently by different systems; it is dependent on sensory, motor and cognitive capabilities. Systems with similar grounding capabilities probably experience the world similarly, e.g. you and me; chimpanzees and gorillas. Humans have a tendency to anthropomorphise other agents; this essentially amounts to projecting our grounding capabilities onto other entities like Honda's humanoid Asimo, as a means to design and interpret their behaviour.

We conceptualize grounding as a capability since it can be more than a transformative step by step process; it can be a complex capability that is experienced and distributed. A grounding capability could be an agent or a group of agents like a team or organization. It is sometimes useful to formally describe grounding capabilities and we use a simple relational description because we want to illustrate the main ideas and to build agents that can reason about the grounding capabilities of other agents.

Grounding capabilities are those capabilities required to produce representations. We say grounding capability G grounds information I to generate representation R , and denote it as $grounds(G, R, I)$ and we define a relation *represents* such that $R = represents(I)$. Note, some special cases may require the relation *represents* to be a function and clearly I may be a label for the actual information, e.g. when I is a form of energy. Consider the pain *Agent007* feels when someone punches him, we can say $grounds(Agent007, pain, punch_impact)$, but the string "pain" hardly represents the pain he feels, and the word punch leaves much unrepresented about that action, especially if *Agent007* would like to anticipate and avoid it next time. These difficulties and deficiencies need to be addressed in system design, serious problems arise if we just say 'pain' = pain and 'punch impact' = punch_impact as the Deflationists do. In robotic system design in particular it is often crucially important to identify and describe grounding via experience and to analyse a particular grounding or representation by making it explicit. Other examples include $grounds(Agent007, "Doctor", image(doctor.jpg))$ which says *Agent007* uses the word "Doctor" to represent image information contained in a file called *doctor.jpg*. We can also use *grounds* to make statements about a common or social grounding which is shared among a group of agents such as the Australian Medical Board (AMB) $grounds(AMB, "Doctor", image(doctor.jpg))$.

Conceptualizing grounding as a capability raises the following obvious and important research questions: (i) what are the important features of a grounding capability, (ii) how can different grounding capabilities be compared and measured, (iii) how can we understand grounding capabilities better, (iv) how can we reason about the different grounding capabilities, and (v) how can we construct a grounding capability for artificially developed systems so that they can build their own representations.

Ownership of representations is important and the ability to reason about ownership and privacy for example is a powerful ability.

On this view if grounding creates representations from information then the search for better representations for humans and computers is a valuable line of research, but a more valuable challenge to pursue is a better understanding of how information can be grounded to create representations by an artificial system through its own experience all by itself. The relationship between making representations and intention also raises many important research questions.

We need to develop tools and techniques that help us explore the process of grounding, only then will we be able to develop systems that can ground themselves effectively, i.e. generate, manage and interpret their own high quality representations.

6 Discussion

A poorly grounded system will struggle to exhibit intelligence, no matter how sophisticated its decision making ability. Despite its significance, a trawl of the literature soon highlights that grounding is still a poorly understood concept and capability. Even worse confusion abounds. Similarly for intelligence, for example, there is no widely accepted definition of intelligence – in this paper we have argued that intelligence involves the ability to make representations. Given the importance of grounding and its major impact in system design it is surprising that it is rarely addressed head. Progress towards building intelligent systems with rich flexible representations has been retarded because the concept of grounding is complex and poorly understood.

The quality of representations singularly determines the choice, quality and power of problem solving approaches and the scope for creativity, innovation, and intelligence. The field of AI has taken some impressive steps towards building systems that can ground themselves but in order to take those achievements to the next level we need to have a better understanding of grounding so as to build robust autonomously grounded systems that can adapt and respond appropriately in novel and unexpected situations.

We provided an innovative conceptualization of representation that delivers significant value given there was no such conceptualization previously; it identifies an important distinction between information and representation much like the distinction between an idea and its expression in Copyright Law. Furthermore, it improves our understanding of intelligence and cognition, and has implications for systems design. In addition, it highlights the need for future research to operationalise the operator *grounds*, since it is this operator that makes representations.

We argue that without a clear and deep understanding of grounding, and practical methods for systems to use to autonomously make representations the achievement of Artificial Intelligence will remain a dream and beyond our grasp.

Acknowledgments. This paper was written while the author was funded on a Pauli Fellowship at the Technical University of Vienna which was made possible by Georg Gottlob and Thomas Eiter. The author is extremely grateful for the opportunity to be part of this prestigious fellowship program.

References

1. Barsalou, L.: Perceptual Symbol Systems. *Behavioural & Brain Sciences* 22, 577–660 (1999)
2. Brooks, R.A.: The Engineering of Physical Grounding. In: *Proceedings of 15th Annual Meeting of the Cognitive Science Society*, pp. 153–154. Lawrence Erlbaum, Hillsdale (1993)
3. Brooks, R.A.: The relationship between matter and life. *Nature* (6818), 409–411 (2001)
4. Chang, M., Dissanayake, G., Gurrum, D., Hadzic, F., Healy, G., Karol, A., Stanton, C., Trieu, M., Williams, M.-A., Zeman, A.: *Robot World Cup Soccer 2004: The Magic of UTS Unleashed!* (2004), <http://www.unleashed.it.uts.edu.au/TeamReports>
5. Chen, X., Lui, W., Williams, M.-A. (eds.): Special Issue on Practical Cognitive Agents and Robots. *Journal of Autonomous Agents and Multi-Agent Systems* (in press, 2009)
6. Gärdenfors, P.: *How Homo became Sapiens*. MIT Press, Cambridge (2004)
7. Gärdenfors, P., Williams, M.-A.: Reasoning about Categories in Conceptual Spaces. In: *Proceedings of the IJCAI*. Morgan Kaufmann, San Francisco (2001)
8. Gärdenfors, P., Williams, M.-A.: Communication, Planning and Collaboration based on Representations and Simulations. In: Khlenhos, D., Schalley, A. (eds.) *Language and Cognitive Structure*, Benjamins, p. 56 (2007)
9. Gärdenfors, P., Williams, M.-A.: Building Rich and Grounded Robot World Models from Sensors and Knowledge Resources: A Conceptual Spaces Approach. In: *Proc. of International Symposium on Autonomous Minirobots for Research & Edutainment* (2003)
10. Harnad, S.: The symbol grounding problem. *Physica D* 42, 335–346 (1990)
11. Karol, A., Nebel, B., Stanton, C., Williams, M.-A.: Case-Based Game Play in the RoboCup Four-Legged League: Part I The Theoretical Model. In: *The Proceedings of the International RoboCup Symposium*. Springer, Heidelberg (2004)
12. McCarthy, J., Hayes, P.J.: Some Philosophical Problems from the Standpoint of Artificial Intelligence. In: Michie, D. (ed.) *Machine Intelligence*, vol. 4. American Elsevier, Amsterdam (1969)
13. Newell, A.: Physical Symbol Systems. *Cognitive Science* V4 (1980)
14. Pickover, C.: *Liquid Earth*, <http://www.pickover.com>
15. Searle, J.: Minds, brains and programs. *Behavioural & Brain Sciences* 3, 417–457 (1980)
16. Sharkey, N.E., Jackson, S.A.: Grounding Computational Engines. *Artificial Intelligence Review* 10, 65–82 (1996)
17. Tarski, A.: *The Concept of Truth in Formalized Languages, Logic, Semantics, Metamathematics*. Oxford University Press, Oxford (1956)
18. Stanford Online Encyclopaedia of Philosophy, plato.stanford.edu
19. Steels, L.: Perceptually grounded meaning creation. In: Tokoro, M. (ed.) *Proceedings of the International Conference on Multi-Agent Systems*. AAAI Press, Menlo Park (1996b)
20. Miller, A.R., Davis, M.H.: *Intellectual Property: Patents, Trademarks, and Copyright*, 3rd edn. West/Wadsworth, New York (2000)
21. Trieu, M., Williams, M.-A.: Grounded Representation Driven Robot Design. In: *Proceedings of the Robot Soccer World Cup* (2007)
22. Williams, M.-A., Gärdenfors, P., McCarthy, J., Karol, A., Stanton, C.: A Framework for Grounding Representations. In: *The Proceedings of IJCAI Workshop on Agents in Real-Time and Dynamic Environments* (2005)