Auton Agent Multi-Agent Syst (2009) 19:272–296
DOI 10.1007/s10458-009-9082-0

A grounding framework

Mary-Anne Williams · John McCarthy · Peter Gärdenfors · Christopher Stanton · Alankar Karol

Abstract In order for an agent to achieve its objectives, make sound decisions, communicate and collaborate with others effectively it must have *high quality representations*. Representations can encapsulate objects, situations, experiences, decisions and behavior just to name a few. Our interest is in designing high quality representations, therefore it makes sense to ask of any representation; *what* does it represent; *why* is it represented; *how* is it represented; and importantly *how well* is it represented. This paper identifies the need to develop a better understanding of the grounding process as key to answering these important questions. The lack of a comprehensive understanding of grounding is a major obstacle in the quest to develop genuinely intelligent systems that can make their own representations as they seek to achieve their objectives. We develop an innovative framework which provides a powerful tool for describing, dissecting and inspecting grounding capabilities with the necessary flexibility to conduct meaningful and insightful analysis and evaluation. The framework is based on a set of clearly articulated principles and has three main applications. First, it can be used at both theoretical and practical levels to analyze grounding capabilities of a single system and to evaluate its performance. Second, it can be used to conduct comparative analysis and evaluation of grounding capabilities across a set of systems. Third, it

---

M.-A. Williams (✉) · C. Stanton · A. Karol
Innovation and Enterprise Research Laboratory, University of Technology, Sydney, Australia
e-mail: mary-anne@it.uts.edu.au

C. Stanton
e-mail: cstanton@it.uts.edu.au

A. Karol
e-mail: alankar@it.uts.edu.au

J. McCarthy
Department of Computer Science, Stanford University, Stanford, CA, USA
e-mail: jmc@cs.stanford.edu

P. Gärdenfors
Lund University Cognitive Science, Lund, Sweden
e-mail: peter.gardenfors@lu.se

offers a practical guide to assist the design and construction of high performance systems with effective grounding capabilities.

Keywords Knowledge representation · Cognitive robotics · Grounding · Perception · Artificial intelligence

1 Introduction

This paper addresses the fundamental issue of grounding system representations which is an important, fundamental and challenging problem in Artificial Intelligence (AI). Roughly speaking, grounding involves building and maintaining representations, and the higher the quality of the representations, the higher the system's performance because the quality of representations influences the effectiveness of perception, conceptualization, decision making, behavior, and the learning of new knowledge and new skills.

The lack of a comprehensive understanding of grounding is a major impediment to the development of genuinely intelligent systems that can generate their own representations and that know what they are doing [17,64]. The main reason grounding still presents a major barrier to the progress of AI is that the key concept of *representation* is defined and used in a wide variety of different, and sometimes confusing, ways across the literature. According to [64] the term *representation* is regularly conflated with model and interpretation, and sometimes simultaneously; it is used to refer to a system's internal world model, to external artifacts like system architecture diagrams, and it is sometimes used to refer to the process of interpreting models as well as the outcome of that process. Williams goes on to delineate representation (the idea) from information (the expression of the idea), and defines the process of grounding as the making and management of representations.

In this paper we endeavor to address the need to understand grounding better by developing a new framework that allows us to evade the need to formally define *representation*, while at the same time enabling insightful analysis and evaluation of grounding capabilities across a wide class of systems from Labrador puppies to Sony AIBO robots (see Fig. 1).

The framework is a novel analysis tool designed to help describe, dissect and inspect grounding capabilities. It has several objectives which include: helping to understand the process of grounding and its implications for intelligence and cognition, guiding and informing theoretical and practical analyzes of grounding capabilities of a single system, evaluating system performance, providing a means to conduct comparative analysis and evaluation of grounding capabilities across several (different) systems, and to offer a practical guide to the design and construction of high performance systems with more effective grounding capabilities.

Our main interest is in the grounding capabilities of intelligent agents however, systems from airline reservations to autonomous mobile robots rely on *grounded* representations. An airline reservation system must manage information about flights and passengers in a way that corresponds to real flights and real passengers over time. Likewise, an autonomous mobile robot that navigates a physical space will be more effective in achieving its objectives if its representations of physical barriers correspond to real physical barriers in its environment. A sound grounding capability provides basic infrastructure for cognition and intelligence. Consequently, *how*, and *how well* an agent is grounded is of significant interest and crucial importance in AI.

Fig. 1 A Sony ERS-210 AIBO robot, a personal entertainment robot. The robot's sensors include a color camera, microphone, joint angles, and accelerometer. The head and legs possess a combined 15 degrees of freedom

1.1 Contents

The rest of the paper is organized as follows:

- In Sect. 2 we discuss grounding as it is described in the literature.
- In Sect. 3 we describe our broad notion of representation.
- In Sect. 4 we describe grounding capabilities of systems and provide a set of principles that guide the Grounding Framework which is developed in Sect. 5.
- Section 5 introduces the Grounding Framework.
- In Sect. 6 we illustrate the power, usefulness and impact of the framework by demonstrating its use in (i) analyzing a specific system, (ii) comparing the grounding capabilities of several systems, and (iii) developing a quality ranking for the system development life-cycle.
- In Sect. 7 we highlight the major benefits and applications of the framework.
- In the final section we discuss the potential of the framework and our plans for future work using the framework.

2 Grounding

2.1 What is it?

In everyday parlance, to say someone or something, is (or is not) grounded can have different meanings. It could mean that someone or something is physically confined (literally or figuratively), such as a plane that cannot take off, or a child that must stay at home. Alternatively, to say an agent, such as a person, is not grounded could mean that their understanding or beliefs about the world (or a particular topic) are out-dated, wrong, irrelevant, or even delusional. In contradistinction, a grounded person is the opposite—for example, "the mechanic has a solid grounding in truck engines" implies the mechanic has experience, knowledge and/or a thorough understanding of the mechanics of truck engines. It is this type of grounding—

loosely described as an "understanding" of the world—that we are concerned about in this paper.

2.2 Grounding and classical AI

The classical knowledge representation paradigm and research agenda, led by [37,40], embraces Tarski's [59] notion of truth to relate system representations to real world "things", such as objects, fluents, events and situations. It focuses on building powerful, and in some cases elaborate and exotic, logic based representations. This stream of research has been one of the most valuable and productive in the field of AI, and as a result we have a good understanding of the power of a wide variety of logic systems, like databases and situational calculus, which are grounded via truth-based representational semantics.

2.3 "Symbol" grounding

In 1990, Harnad coined the term "symbol grounding" to describe a problem related to the Chinese room thought experiment [47] and "symbolic" (logic-based) systems. Harnad likens the symbol grounding problem to trying to learn Chinese from a Chinese dictionary alone, where every word is defined in terms of other Chinese words. Thus, without any knowledge or experience of Chinese, as each Chinese word is defined in other *meaningless* words, the task is seemingly impossible due to an infinite regress. At some point, Harnad argues, the meaning of at least some words must be grounded in experience. Neural networks are suggested as a means of learning the relationships between (arbitrary) symbols and the subsymbolic, "invariant" sensory features to which they relate. Thus, Harnad attempts to solve the grounding problem by finding the patterns of sensory features that correlate to symbols, or vice versa—the apparent argument being that because the machine's symbols are "connected" to (or influenced by) sensory experience the symbols are now (somehow) intrinsically meaningful to the machine.

2.4 "Physical" grounding

Whereas Harnad [30] advocated solving the grounding problem by the use of hybrid logical-connectionist systems, the behavior-based robotics camp, ostensibly led by Brooks [6–9], rejected the notion and the need for logic based approaches at all. Nouvelle AI embraced the doctrine of embodiment and situatedness, claiming that so-called symbolic representations were not needed at all because the world/environment can be sensed directly when required. The physical grounding hypothesis states that "to build a system that is intelligent it is necessary to have its representations grounded in the physical world". This nouvelle AI emphasised that physically grounded systems require real robots, built bottom up, so that "high level abstractions have to be made concrete" [6]. Brooks argues that because of the physical grounding hypothesis, "traditional symbolic representations" are no longer necessary, and the symbol grounding problem is avoided. See [64] for a critique of this stance with regards to systems that anticipate future environmental states.

2.5 Grounding as correspondence

A large body of research treats grounding as a problem of reference. A theory of reference [14]—or alternatively "correlational semantics" [44]—involves relating internal representations with external entities (e.g., somehow connecting a symbol "John" with the real-world *John*)—in other words, understanding how the atomic units of a language come to have meaning. Symbol grounding has been described as establishing the "direct correspondence between internal symbolic data and external real world entities" [2], the problem of how "symbols should acquire their meaning from reality" [61], or the association of a symbol "with a pattern of sensory data that is perceived when the entity that the symbol denotes is seen, or tasted etc" [36]. Likewise, "anchoring" [18], a variation of the grounding problem, embraces the problem of reference—anchoring involves "maintaining the correspondence between symbols and sensor data that refer to the same physical objects" [19].

2.6 "Autonomous" grounding

For some, grounding is a technical problem which involves connecting symbols with experience (and vice-versa) *autonomously*. For example, Prem [44] argues that symbol grounding is "automated model construction". Likewise, Steels [54] argues if "someone claims that a robot can deal with grounded symbols we expect that this robot autonomously establishes the semiotic map that it is going to use to relate symbols with the world". For Taddeo and Floridi [58] symbol grounding is the problem of how an artificial agent can "autonomously elaborate its own semantics" through interacting with its environment. Williams [64] characterises autonomous grounding as the process of making sense that results in representations.

2.7 Grounding terminology

With "meaning" and "understanding" being such broad, far-reaching concepts, there is not surprisingly a large body of "grounding" related literature. In fact, a study of the literature reveals the term "grounding" has a range of different intuitive and technical meanings. For example, "representation grounding" [12], "theory grounding" [45], "physical symbol grounding" [61]; the grounding of "beliefs" [46], "double grounding" [41,42], and "anchoring" [18]. Grounding has also been use to describe a large body of work concerned with the grounding of language and how words get their meaning [11,27,31], and how meaningful languages evolve [50–53,55,56]. The term "grounding" has also been used to describe how human participants in conversation reach shared meaning and understanding [15,16], and similarly how participants in human-computer interaction reach shared meaning and understanding [5,10]. For reviews of the literature, see Ziemke [48]; Taddeo and Floridi [57]; Sun [58]; and Stanton [65].

2.8 An important, open problem

The need for intelligent systems to attribute meaning and conduct commonsense reasoning continues to be highlighted in the literature [47,20]. The grounding problem has been described as one of the major outstanding problems facing artificial intelligence [23]. At a practical level, grounding effects all systems which rely upon representations such as beliefs (however implicit or explicit) about the state, nature or behavior of the world for the purposes of decision making or action. As noted earlier systems from airline reservations to autonomous mobile robots rely on *grounded* representations.

For agents like robots to respond appropriately to novel and unforseen situations (situations for which they have not been explicitly programmed) they need sophisticated grounding capabilities. Novel situations need to be interpreted meaningfully by the agent—and not just by the designer a priori—but in real time. However, even correctly classifying a situation, event or "thing" as being new or unknown is a difficult task in itself [34]. The need to cope in novel and/or unexpected circumstances is an imperative design requirement, as for artificial agents to act intelligently they must be able to behave appropriately in complex and dynamic environments in which they should "expect the unexpected".

In summary, grounding still remains an open problem and our current understanding is preliminary and immature; grounding continues to parade around in too many deceptive guises. It is lurking behind every debate on the nature of intelligence and cognition. We argue that there is significant value and potential payoff in pursuing a better understanding of grounding. In fact, we believe that a major development in our understanding of grounding could lead to a major breakthrough in the field of Artificial Intelligence. Scientific breakthroughs are difficult to craft but the development of a practical Grounding Framework designed to enhance our understanding is an important step forward. Due to its concern with meaning and understanding, grounding is a difficult and complex concept, perhaps akin in difficulty to the problem of cognition itself. However, due to its central and increasing importance in AI we cannot simply put it into the "too hard basket". Instead, we must delve into its richness and develop a better understanding, in order to move the field of Artificial Intelligence forward.

3 Representations

3.1 What are representations?

Representations for our purposes are completely open, provided of course that they represent information. They can range from low level sensorimotor representations (an image from a camera) all the way up to high level logic and linguistic representations. Other common examples of representation in robotics (and software in general) include variables, classes, databases, and so forth—i.e., any data structure, internal state or memory. Representations in our framework include low level sensorimotor information such as YUV or RGB values of pixels in a digital image (see Fig. 2) through to information about entities that can no longer be experienced like dinosaurs and melted ice cubes and imaginary entities like Tolkein's hobbits [60].

3.2 The need for representation

Robotic agents require an appropriate set of instructions, describing how to act in the world—but not necessarily an explicit description of the world itself. As such, the value of representations has been greatly debated in the literature. For example, Brooks' (1990) response to the symbol grounding problem (and other problems related to traditional AI) was to argue that traditional, explicit symbolic representation "just gets in the way". According to Brooks, "the world is its own best model…the trick is to sense it appropriately and often enough". Instead of building top-down models of the world, Brooks believes that intelligent behaviour can *emerge* in a bottom-up direction from a collection of cooperating "behaviours", with each behaviour tightly coupled to sensors and effectors. While the behaviour based approach showed early promise, it has failed to scale. Most importantly, future world states are

```
,137,129)(98,135,125)(100,137,135)(96,140,133)(87,136,141)(85,136,140)(82,144,146)(86,140,133)(90,142,133)(93,14
0,142)(44,134,147)(47,136,143)(45,126,143)(46,126,150)(46,129,149)(47,131,135)(44,129,140)(45,131,139)(45,130,14
(49,134,147)(48,128,157)(50,132,153)(50,130,149)(48,131,145)(48,134,144)(47,126,149)(45,138,152)(48,130,160)(45,1
6,143,127)(96,138,134)(100,133,129)(97,136,131)(89,131,135)(87,134,143)(85,137,131)(82,144,133)(88,140,126)(89,1
32,138)(44,128,139)(46,130,141)(43,131,139)(45,131,140)(45,128,146)(45,131,147)(47,132,135)(45,129,142)(45,132,1
(51,135,144)(48,127,153)(52,131,152)(52,126,138)(50,130,139)(48,130,158)(48,128,155)(47,132,139)(50,126,145)(48,1
,137,123)(97,141,130)(99,144,128)(95,141,133)(90,133,137)(86,134,137)(83,139,138)(84,135,134)(84,139,129)(81,142
,146)(47,134,133)(44,131,144)(45,135,145)(47,137,140)(46,131,144)(44,127,141)(46,133,133)(48,135,134)(48,134,133
(47,131,151)(46,126,158)(48,128,156)(48,132,154)(46,126,152)(48,126,155)(47,126,150)(47,124,149)(50,126,155)(49,1
,142,132)(99,139,126)(101,135,131)(97,134,131)(84,137,141)(80,137,148)(78,140,147)(80,139,140)(82,135,131)(78,13
3,135)(47,134,132)(46,132,135)(46,131,137)(50,138,141)(48,128,136)(46,121,139)(46,129,134)(46,134,146)(46,129,15
(46,130,149)(48,122,150)(52,126,148)(48,128,146)(52,128,141)(48,131,151)(49,124,148)(53,121,140)(48,115,140)(48,1
9,139,128)(99,137,120)(104,134,128)(98,136,122)(84,137,137)(84,138,144)(75,137,140)(76,141,137)(78,144,129)(80,1
32,136)(46,135,131)(48,134,133)(47,133,137)(45,131,138)(48,133,140)(49,130,135)(50,133,140)(46,134,136)(47,138,1
(49,131,141)(48,130,153)(51,128,153)(51,126,142)(49,130,143)(47,131,151)(45,127,151)(45,133,153)(46,131,153)(50,1
,142,124)(97,135,132)(100,137,129)(95,139,132)(85,134,134)(81,131,140)(73,135,144)(72,143,142)(79,141,138)(73,13
1,136)(46,133,135)(47,131,141)(46,134,136)(46,135,137)(47,134,140)(47,135,138)(49,133,136)(49,132,132)(49,135,13
(48,131,139)(50,126,138)(49,131,144)(48,126,148)(49,134,147)(49,126,147)(45,128,144)(49,133,137)(49,126,145)(46,1
,141,127)(97,137,125)(99,140,129)(94,137,130)(84,136,140)(77,135,142)(75,139,138)(78,135,135)(75,141,126)(77,142
,138)(45,134,139)(48,132,136)(47,132,138)(49,133,136)(46,132,135)(46,127,149)(46,133,151)(49,135,141)(49,132,136
(50,131,144)(52,128,149)(50,125,148)(47,127,150)(44,134,160)(46,127,153)(45,125,160)(44,125,156)(48,125,157)(48,1
6,140,130)(96,140,127)(100,134,132)(94,136,132)(83,131,140)(73,132,146)(71,145,147)(74,140,137)(73,143,139)(74,1
31,136)(44,135,145)(49,131,147)(49,131,142)(52,132,129)(48,133,137)(50,134,148)(52,134,139)(48,133,148)(51,134,1
(48,134,142)(49,128,147)(48,126,150)(49,133,144)(52,139,143)(52,127,138)(45,128,146)(50,132,153)(48,120,143)(48,1
,138,128)(100,135,124)(101,136,126)(96,133,126)(85,129,135)(75,125,143)(72,142,144)(74,143,128)(74,141,126)(72,1
33,136)(50,130,137)(52,128,133)(47,134,143)(48,134,139)(48,130,143)(49,130,145)(50,136,139)(50,136,139)(50,132,1
(49,129,146)(46,128,150)(47,126,152)(48,129,149)(49,130,150)(50,121,149)(47,124,143)(48,135,140)(45,130,155)(45,1
,140,131)(97,139,128)(99,136,126)(94,135,131)(81,131,140)(77,126,139)(74,138,139)(72,140,132)(71,139,122)(70,139
,130)(50,133,141)(50,128,129)(50,131,136)(49,136,133)(49,128,140)(52,131,138)(51,133,141)(51,139,140)(48,133,133
```

Fig. 2 Raw image data represented as a string of triples, as processed by a robot's vision processing system (*top*), and the corresponding data represented as an RGB image for humans to understand (*bottom*)

generally are not a feature of the current world state, consequently systems that need to plan effectively and anticipate future world states require representations.

3.3 Representation grounding

In defining the symbol grounding problem, Harnad [30] describes the problem in relation to "symbolic AI". However, what is a symbol? Steels [54] comments that the use of the term "symbol" in artificial intelligence has probably created "the greatest terminological confusion in the history of science". Steels argues this confusion arises from differing uses (or meanings) of the term symbol by researchers with different backgrounds, i.e., philosophers, linguists and computer scientists use "symbol" in different ways. For example, a computer programming language is itself symbolic, yet when a neural network is implemented in a computer language using such symbols, the neural network is not considered by cognitive scientists or philosophers to be "symbolic"—rather, it is considered to be "subsymbolic". Thus, "symbols" in the context of symbol grounding literature usually refer to logical sentences that are used for reasoning about the world (e.g., loosely speaking "Sydney is in Australia", "Australia is in the southern hemisphere", so therefore "Sydney is in the southern hemisphere").

Several authors [12,25,35,43,64,65] have commented that the grounding problem is not limited to "symbols", but to "representation" more generally. For example, MacLennan [35] describes the grounding problem as "how do representations come to represent", while Pfeifer and Verschure [43] describe a "general grounding problem" which applies to knowledge "structures", rather than just "symbols". Williams [64] takes this more general view that representations are grounded information, where a mental model is an example of representation whilst the environment gives rise to electromagnetic information captured by the human eye and subsequently transduced by the cones and rods in the retina and transmitted to the visual cortex for further grounding via the optic nerve. Johnson and Williams [32] provide the first formal framework for the symbol grounding problem.

A grounded representation does not require that every entity in the representation be *linked* to a corresponding physical manifestation, but we might expect that a meaningful relationship typically exists between the entities in a representation and the entities being represented in some sense. For example, the image represented on the retina is an inverted visual reflection of the outside world. As noted in Sect. 2.5 maintaining a correspondence between representations of physical objects and the objects themselves is important but so too are representations of object functionalities and relationships between objects, as well as experiences, intentions descriptions of ways to interact with specific objects such as affordances [26,64].

3.4 Sensations, perceptions, and simulations

For the purpose of understanding grounding it is insightful to classify representations using the hierarchy of Gärdenfors [24,25], illustrated in Fig. 3, which describes relationships between three key representational entities: sensations, perceptions, and simulations. Representations in the hierarchy can take two forms: cued and detached. *Cued representations* are based on the perception of things that are present, and detached representations focus on entities that are not currently perceived, and possibly never perceived. This hierarchical classification is useful for analyzing existing systems and designing new ones.

According to Goldstein [28] *sensations* are immediate sensorimotor impressions from a single sensor, and *perceptions* are conscious sensory experience. *Simulations* are detached representations; our imagination is a good example of simulated representations, so too

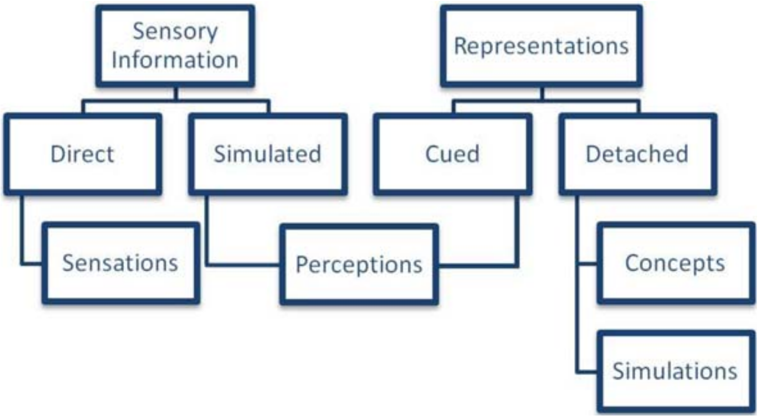

Fig. 3 Cued and detached representation hierarchy

those representations produced by mirror neurons [29], but there are many others. Hybrid representations of a combination of cued and detached representations are also possible, e.g. one could look at a chair and imagine Napolean sitting on it.

Sensations provide systems with an awareness of the external world and their internal world. They exist in the present, are localised in the body/system, and are modality specific, e.g. visual, auditory, but not both. Perceptions encapsulate more information than raw sensorimotor information. They can represent accumulated sensorimotor information and sensorimotor information reinforced with concepts, knowledge and simulations [4,24,25]. Whereas sensations are signals from sensors monitoring the environment of self (i.e., proprioception), perceptions require additional information derived from previous experiences and/or outcomes of learning. In contrast to sensation, perception is cross-modal, and perceptions can generate permanence, e.g., object permanence.

Detached representations of objects exist as well as detached representations of relationships, properties, affordances, actions, events, states, and processes.

Representations can be derived from information that has been gathered from a wide range of sources, e.g., internal and external sensors, internal and external effectors, external instruments, external systems, etc. In addition they can result from fusing sensorimotor information with high level representations such as perceptions, concepts and linguistic representations. Consider a doctor who not only grounds his own sensorimotor information, but information gained from colleagues, books, lab tests, instruments such as thermometers, and equipment used to visualize heart beat, and to measure blood pressure and oxygen content of the blood [64].

We illustrate several kinds of representations in Fig. 4 based on a Robot Soccer System [1] which are constructed from raw robot camera data captured by a Sony ERS210 AIBO robot; this robot is illustrated in Fig. 1. Figure 4a is a 2D visualisation of the ERS210's *raw* camera data, and Fig. 4b is a processed version of Fig. 4a where specific colours (YUV values) of pixels are used to determine if they *belong* to specific objects of interest—a ball, a beacon and a goal are clearly identified.

The information illustrated in Fig. 4b can be used to find the distance, heading and elevation from the robot's camera, of the various objects of interest which in turn can be used to calculate the pose of the robot in a global reference frame. Information represented in Fig. 4b can be combined with a relational representation of robot location, namely $robot(id, x, y, \phi_1)$ to create a relational representation of the location of objects given by $object(o, r, \phi_2, \theta)$ where id is a robot's identifier, x and y are coordinates of the robot, ϕ_1 the heading of the robot in a predefined world coordinate system, o can be one of [*ball, beacon, goal, team-mate, opposition-robot, obstacle*], r is the distance from the camera of the robot to the object, ϕ_2 is the heading to the object and θ is the elevation to the object relative to the robot's camera system.

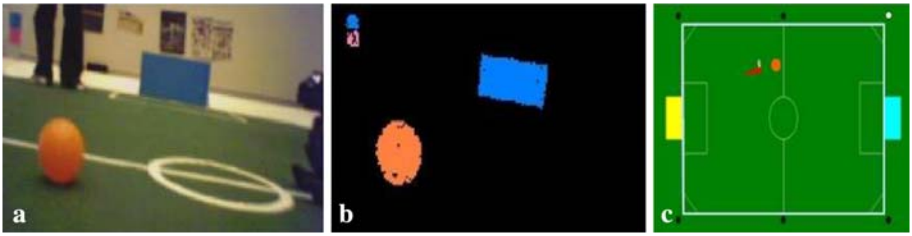

Fig. 4 Representations from a Robot Soccer System: **a** digitial image derived from a robots camera, **b** perceptual representation of the ball, beacon and goal, **c** the world model

The set of spatial relationships a robot has with objects described by *object* relations can be depicted and visualized as a 2D soccer field for ease of interpretation by humans. Visualizations of relevant soccer objects such as goals, robots, and the soccer ball can be placed in specific locations on the field image in a way that represents their position on the real field—see Fig. 4. Figure 4a illustrates a sensation, Fig. 4b illustrates a cued perception derived from the integration of sensorimotor information and a detached representation of the robot's body, and Fig. 4c is a detached representation of the robot's world model, constructed using a history of object location observations.

Detached representations are extremely powerful, but creating and maintaining them requires significant computational resources. Resources that reptiles probably do not possess, e.g., reptiles do not exhibit capabilities for object permanence in the same way mammals do. Detached representations can be manipulated independently of the external world. They can be conceived without being perceived. Some examples of detached representations are absent objects, past and potential future world states.

4 Grounding capabilities

Intuitively, a grounded agent has a good grasp of their environment, themselves, their own capabilities, other agents and their capabilities. We constantly make complex assessments about the grounding of agents that we engage with, e.g., we all know people who are down to earth, naive, unrealistically optimistic, hopelessly impractical, or irritatingly rational. These assessments about the grounding capabilities of others are important because they influence what we think and do. It is noteworthy that most of these assessments are largely based on a judgement relative to our own grounding capabilities and representations.

Representations represent information, and a grounding capability creates representations [64]. The process of grounding plays an important role, providing critical infrastructure for cognition and intelligence. It produces representations, and our framework attempts to discern the degree of groundedness representations possess. Grounding underpins grounded systems; groundedness (a noun) refers to the property possessed by grounded things—in our case representations and systems.

The framework described in the next section focuses on grounding capabilities, namely, the ability to make representations. It helps to make judgments about the quality of those representations, and thus the grounding capability itself. A grounding capability can capture information and manage information exchanges to and from the external world, it can also create, interpret, manage and maintain internal and external world representations.

Grounding capabilities support system goals and objectives, and therefore measuring the quality of a grounding capability should be conducted with respect to the system goals and objectives. The purpose of a grounding capability is to construct and maintain representations that correspond meaningfully to the things being represented so that the system can achieve its aims and objectives effectively. Clearly the *quality* of a system's grounding capability will have an crucial impact on the success of the system, and on what it can achieve.

4.1 Grounding in traditional system development

In most computer systems, the determinations of what to represent and the establishment of a correspondence between representations and the entities they represent is typically established and maintained by human designers. As a result, the systems produced are (at least partially) grounded externally via the human mind. In other words, the human mind (i.e., designers, developers, users, etc) plays a major role in creating, subsequently interpreting, and maintaining the correspondence between entities in representations and the actual entities themselves throughout a systems lifetime. The computer based system's role is to conduct data processing, of which the final result is typically interpreted by humans.

Computer based systems are reviewed by (a team of) people during the various phases of the system development life-cycle. Through those reviews, the groundedness of the system is evaluated, with the evaluation often resulting in modifications of the system so as to improve it. System evaluations amount to a determination of the system's ability to satisfy the system requirements—i.e., an assessment of the system's grounding capability.

Depending on the level of sophistication of the system under review, once humans have established the correspondence between the behaviors of entities in the "external world" and the internal system's representations (such as a database conceptual schema), systems such a DBMS can manage the correspondence relationship over time, but only in restricted ways. For example, a DBMS can add, remove, modify and validate relational tuples via application programs without human intervention, but should the database conceptual schema require modification, a human will likely determine how best to accommodate the changes and ensure that the database remains grounded, i.e. there exists an appropriate correspondence with its *external world*.

Changes to grounding requirements due to changes in the environment and/or systems requirements, for example, result in changes to the system grounding capabilities and those changes are typically made by human designers, not the system itself for simple applications like database management systems.

4.2 Principles for grounding and groundedness

In this section we present the principles that will be used to guide the development of the Grounding Framework in Sect. 5.

- *Grounding is a process* that creates representations.
- *Grounding is conducted by systems with grounding capabilities* which create, manage and maintain representations. It can involve a single system or extend beyond the boundaries of a single system, e.g., it may rely on components beyond the system such as external sensors, resources, tools, instruments, other systems or agents.
- *Groundedness is a property of representations and grounded systems.*
- *Grounding can be contextual*, and when it is, it should be measured relative to system goals.
- *Systems can ground their representations in a variety of different ways*: top-down via the process of design; bottom-up via sensors, effectors, and interfaces, and through information obtained via external objects (e.g., physical tools), external sensors (e.g., radar), and external systems (e.g., medical monitoring system); or a complex combination.
- *Groundedness is graded.* The concept of groundedness is rich. A representation or system is not simply grounded or ungrounded. There are *degrees* of groundedness.
- *Groundedness is multidimensional.* As groundedness should be measured relative to system goals, the salient dimensions will vary depending on what is determined to be important for the task at hand.
- *Measures of groundedness can be qualitative or quantitative, continuous or discrete.* A grounding framework should not impose restrictions on the form and measure of assessment, only its capacity to support the systems goals.

- *The nature of a representation is open*. A grounding framework should not place undue restrictions on what could or should be grounded. Anything can be grounded: physical things, abstract things, nonexistent things, and things that have never been experienced. Things can include: objects, relationships, states, actions, processes, events, etc.
- *Grounding occurs in biological and artificial systems*. A grounding framework should cater for a wide range of systems from artificial to biological.

These principles clearly articulate our key underlying assumptions upon which the Grounding Framework in the next section builds on.

5 The grounding framework

The Grounding Framework is motivated by the need to understand grounding better and to understand how to build sophisticated systems that can ground information to form representations by themselves. AI systems and agents need to be able to create their own representations so that they can cope with novel, dangerous and unexpected situations.

We build on a wide variety of pioneering and important work on grounding, some of which was described in Sect. 2, to develop a new understanding of grounding and the first framework for evaluating *how well* system representations are grounded. The framework is designed to describe, evaluate and in some cases formally measure the quality of representations and grounding capabilities which can be system specific, domain specific, and context specific. Our framework strongly supports the idea that when it comes to assessing grounding capabilities there are few absolute measures. Typically *groundedness* of a system is measured relative to the groundedness of other systems, e.g., it is common to evaluate the grounding of systems using human mental representations as *reality* or human grounding capabilities augmented with additional sensors and instruments such as a radio telescope. The framework we develop can be used to understand grounding capabilities in existing systems and to support the design and implementation of intelligent agents whose representations need to be grounded in order for them to achieve their design goals.

The Grounding Framework comprises five essential elements which can be as detailed as required for the purpose of the analysis at hand:

1. System Objectives
2. Architecture of Grounding Capability
3. Purpose and Scope of the Analysis
4. Nature of the Grounding Capability
5. Groundedness Qualities.

The Grounding Framework provides a structure and some guidelines to assist the analyst to understand and describe the grounding capabilities and groundedness of a system and its representations. The analyst drives an analysis of their systems(s) and gathers and documents information about the system along the lines suggested by the Grounding Framework. In subsequent sections will illustrate how to use the Grounding Framework and we highlight the benefits it brings.

All five components of the Grounding Framework are related; the objectives and the scope often determine how the qualities of groundedness are chosen, interpreted, and assessed. For example, a system designed to locate objects within nanometers clearly requires representations with that level of accuracy.

The rest of the paper provides a description of the framework's components and several examples of how to use it.

5.1 System objectives

The first part of the framework involves developing a description of the system(s) objectives, goals, tasks and activities. The level of detail will be determined by the nature of the system grounding analysis being undertaken.

5.2 Architecture of grounding capability

The second component of the framework is a description of the underlying system architecture that supports or implements the grounding capability.

An architecture defines the structure and organization of the main system components and their channels of communication. There are a wide range of potential architectures, e.g., layered, embodied, cognitive, etc. Furthermore, a grounding capability can be described in terms of a number of different architectures. The architecture should be described so as to maximally expose the grounding capability. The description of the underlying system architecture should focus on the components and processes involved in systems' grounding activities, this may involve describing the role human designers play in the crafting and creation of representations.

If systems are being compared then it is desirable to describe the architectures using similar concepts and components. Often representations are translated from one form to another. Details about the relationships among the representations can also be given including details of elements of representations that are preserved and those that are changed during such transformations.

5.3 Purpose and scope of the analysis

The third component of the framework is a detailed description of the purpose and scope of the analysis. The purpose states why a grounding analysis is being conducted, and the scope describes what parts of the system will be analyzed. For example, the scope of the analysis might be restricted to a specific component of the system or grounding capability, a specific set of interfaces or system activities, or specific grounding activities such as the creation of associations between representations and external entities.

5.4 Nature of the grounding capability

The fourth component describes the nature of the grounding capability under analysis. A useful approach to describing the nature of the grounding capability is with respect to the underlying architecture. Important characteristics of a grounding capability are described in the example given in Sect. 6.1.

5.5 Groundedness qualities

The fifth component of the framework has a "plug and play" aspect where the groundedness qualities are identified by the analyst and then described and evaluated. This component of the Grounding Framework includes a description of the pertinent groundedness qualities as determined by the analyst, and an assessment of them relative to each architectural component of the grounding capability, where appropriate. The instruments for evaluating or measuring the qualities should also be identified. It is important to be able to compare and contrast grounding capabilities in different systems, consequently lower level features of groundedness need to be determined in order to evaluate grounding capabilities in systems effectively.

As noted earlier, a key principle is that grounding is multi-dimensional and graded. The components of a grounding capability, and the dimensions/qualities of groundedness need to be identified in order to better understand and ultimately evaluate groundedness. In order to enable the evaluation of groundedness we identify a set of important features/dimensions that can be used as key performance indicators for assessing the quality of a grounding capability.

In what follows, we describe some groundedness qualities which are appropriate for assessing an intelligent agent. For any particular groundedness analysis, we envisage that a set of appropriate qualities will be identified based on both the objectives of the system(s) under analysis, and the scope and nature of grounding analysis. Some of the qualities, below, are so fundamental to the grounding endeavor that they could be used as candidate qualities for almost every grounding analysis.

Expressiveness is the breadth of concepts that are representable and hence need to be grounded, namely objects, relationships, processes, actions, events, states, situations, contexts, affordances etc. Expressiveness is a measure of the richness of representations. It is widely used in Knowledge Representation and there are methods that can be used to measure expressiveness. For example, it is well known that Predicate Calculus is more expressive than Propositional Calculus.

Relevance determines the degree of relevance of the entities that are represented by a system. Relevance is related to, but different from, expressiveness. It focuses on issues related to those aspects of the world that are important for a system to achieve its goals. For example, a robot soccer player may not perceive the audience, or field lines painted on the field because they are not relevant to its tasks or it can achieve its goals without specifically representing them. Changes in task, goals and environment are considered elsewhere and so in the assessment of relevance we only consider current goals; not potential or future goals. Representations are selective in terms of what they can represent—a representation cannot capture every feature or aspect of the world. Choices have to be made with regard to the entities that are important, relevant, and necessary for the system to complete its tasks and achieve its goals. Determining relevance autonomously is a major challenge for computer based systems.

Faithfulness is a property of the relationship between entities in representations and the entities they represent, e.g., the relationship between a robot's world model and the world itself. Faithfulness is a matter of degree and the pertinent question is how *closely* does a system's representations correspond to the entities being represented. Determining the degree of faithfulness, or fidelity, is sometimes achieved by measuring the ability of the system to model the world states and world state transitions in terms of prediction and explanation.

Correctness concerns the ability of a system to represent information in accordance with its specification. For example, a robot soccer player's ability to determine the location of the ball on the field would be an example of a task which has a well defined specification and whose correctness could be measured. The correctness of the task could be context dependent. For example, a robot's ability to perceive a ball's location may be better when it the ball is stationary than when it the ball is in motion.

Accuracy/Precision is related to faithfulness and correctness and involves the accuracy of the information being represented. For example, the robot soccer players perception of its position and the ball's position on the field might be required to be measured to different degrees of accuracy, e.g., to the nearest millimeter or meter. Accuracy is often measurable and as a result relatively easy to evaluate.

Robustness is the ability of a system's representations to behave appropriately to unexpected or abnormal conditions. For example, the ability of a robot soccer player to handle changes in the environment such as lighting variations, changes in background noise, changes in playing surface texture. More dramatic environmental changes would include a change of ball, e.g., different size, different color, different degree of hardness, and/or different density.

Adaptability is the ability of representations to adapt to task and goal changes. For example task changes might involve a robot's ability to change soccer positions, e.g., from Defender to Striker. More dramatic changes involve changes to the rules of robot soccer or a change in the number of robots on a team. Adaptability can be measured by determining the nature/difficulty of the changes that the system can tolerate [39]. To what extent the system can change itself, and when does it require human assistance if we introduce new objects, new relationships between objects, new actions, new events, new affordances, etc.

Timeliness is the ability of representations to be created and respond (appropriately) in a timely fashion. For example, a robot soccer player's ability to dive for a ball as the ball approaches, rather than after it has passed it by, depends on the timeliness of the construction of the representation that can be used to predict the arrival of the ball to the robot.

Efficiency is the ability of a representation to place as (few) demands as possible on hardware resources such as processor time, communication bandwidth, internal and external storage, sensors, effectors, and actuators.

Self-awareness Since systems with self-awareness requirements are typically embedded or embodied the degree of self-awareness is of interest. For example, the question of whether a robot is aware of the state of its body parts such as its forearm is cocked at a 45° angle, will be important when assessing a grounding capability. Self-awareness involves a representation that is graded from physical awareness up to intention awareness. It also raises issues concerning the role of trust in grounding, e.g., being aware of one's own sensor limitation can impact grounding capabilities.

Awareness of others Awareness of others is graded: the spectrum of awareness of others spans from the mere existence of others to the intention of others. The degree of awareness about the grounding capability of others and the intentions of others is important for communication and collaboration because such an understanding facilitates the sharing of information in meaningful ways. The issue of trust is also important, e.g., awareness of other's limitations and biases can impact grounding capabilities.

Functionality involves identifying the system abilities that require grounding. For example, some basic functionality of a robot soccer player includes the ability to recognize the ball, move to the ball, grab the ball, and kick the ball. Different robot players may have different abilities, for example a goalie may be able to dive for the ball whilst a forward may not.

Transparency is the ability of a system to represent its internal information and knowledge in a way that is accessible to itself or other systems. For example, is the representation of information explicitly represented or implicit, clearly derivable or buried in a black box processor. Transparency is a crucially important quality for some systems. A strong transparency quality allows a system to be compared with other systems across a wide range of dimensions with confidence.

Testability involves the ease of testing system grounding capabilities and associated activities such as behavior and decision making. Building more effective systems in the future will be advanced by learning from grounding capabilities in existing systems, and clearly more will be learned from transparent, easily understood and testable systems.
Uncertainty management It can be important to identify, qualify and quantify uncertainty in the grounding capability. This will involve determining the strategies used by the system to address the uncertainty. The focus is on how the system reduces the uncertainty of information gathering and internal information management rather than what techniques are used to manage uncertainty in representations.

Important interrelationships exist among the qualities described above such as faithfulness, correctness and accuracy. Transparency and testability are also clearly related. Other qualities may be derived from those listed above such as reliability which could be the likelihood of an agent to malfunction or the likelihood that a system will behave similarly in similar situations, flexibility which is related to robustness, and adaptability, and performance which captures the responsiveness of a grounding capability which could be measured by the time required to respond to stimulus or the number of events processed in some interval of time. Performance is related to a number of qualities including efficiency, expressivity, timeliness, faithfulness and relevance. The quality dimensions chosen for the analysis will depend on its purpose.

5.6 Evaluating groundedness qualities

In order to assess the quality of grounding it is helpful if the entities to be grounded are clearly identified. Such entities might include objects, relationships, affordances, actions, states, events, plans, and other processes. Objects can be physical (a ball), or abstract (a penalty), internal (a forearm angle) or external (a teammate's relative position). They can be permanent, temporary or ephemeral. Relationships typically exist between objects such as the ball is in the penalty box region on the field; the ball is located in a specific robot's half; the goalie has possession of the ball; the ball is out of bounds; the ball is in the goal area; the ball is dead, i.e., out of play.

Our approach to evaluating groundedness is to assess and/or measure constituent quality dimensions relative to system goals and architecture. A wide range of instruments can be used in concert to assess and measure specific qualities depending on the nature of the systems in question, and the cost costs and benefits of measurement. These include:

- Direct observation and analysis of working system behavior. Occam's razor is a useful heuristic in this circumstance.
- Design of test cases and scenarios that examine, expose and push the limits of system grounding capabilities.
- Analysis of artifacts produced by and for the system should they exist, e.g., design documents, software code, system output.
- Development of formal measures, e.g., the *closeness* of a soccer field configuration to the actual field configuration can be measured using case-based techniques developed in Karol et al. [33].

Some evaluation methods for certain systems are external such as direct observation, others involve internal analysis. Some qualities can be evaluated via external methods, others need to be measured internally, whilst others measured using a combination of a both modes.

5.7 Benefits of using logic driven systems

Logic-driven systems, from database applications to more sophisticated knowledge systems, form an important, privileged, and well studied class of systems. Major benefits flow from the possession of clear semantics in particular building, managing, testing, and measuring grounding capabilities can become straight forward when the representations have a fully specified unambiguous semantics. For example, the faithfulness quality often collapses to an evaluation of truth/falsity (systems can often be shown to be sound, complete, and sometimes decidable), and as a result properties of the methods and algorithms used to determine truth/falsity are at focus. A clear semantics can also enhance the qualities of expressiveness, relevance, correctness, accuracy, timeliness, understandability, transparency and testability. Accounts and measures of robustness, adaptability, self-awareness, awareness of others can sometimes be given.

Many types of logical representations and systems have been developed to enhance standard logics' ability to represent more complex, imprecise, incomplete, uncertain, and dynamic information such as nonmonotonic reasoning [38], possibility logic [21], belief revision [3,62], and languages for action and change [40].

For many applications the benefits of using logic based systems far outweigh the costs, and as a result database and related technologies will continue to flourish. Once a database/knowledge based application has been implemented the DBMS or KBMS maintains the integrity of the information's groundedness, often without human intervention through transactions (addition/deletion) or logical operations (revision/contraction/update) which keep the database or knowledge base in-step with events in the systems' external world, e.g., if a customer changes his address then the DBMS can update the appropriate tuples in the database autonomously via a pre-defined transaction.

6 The power and utility of the grounding framework

In this section we highlight the power and utility of the framework by demonstrating how it can be used to analyse several real complex multi-agent systems. In particular we will use the Grounding Framework to illustrate how to (i) measure the groundedness of the UTS Unleashed! 2003 Robot Soccer System [1], (ii) compare the groundedness of the UTS Unleashed! 2003 Robot Soccer System with the UTS Unleashed! 2004 Robot Soccer System [13], and (iii) develop a grounding quality ranking for use in systems design.

6.1 Measuring system groundedness

In this section we illustrate the use of the Grounding Framework by outlining an analysis of a sophisticated robot soccer system's grounding capability. Robots on this team can perceive the ball, search for it when it is not in view, chase it, kick it, etc. The robots build and maintain a representation of the state of the soccer field from their sensors and internal body data, and then use that representation to make decisions about the best action to perform. The system is based on the classical *sense-think-act* processing cycle [22].

1. *System objective* To play soccer in the RoboCup 4-Legged League[1] at an internationally competitive level under the rules for 2003 using Sony AIBO ERS210 Robots, illustrated in Fig. 1.

---

[1] See http://www.robocup.org for details.

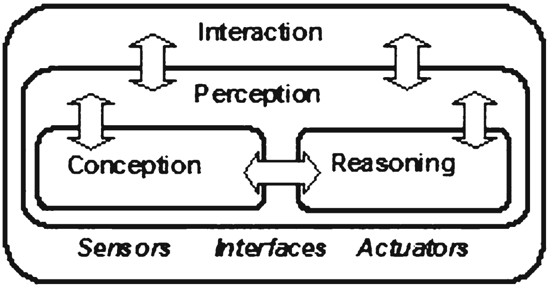

Fig. 5 2003 UTS Unleashed! Robot Soccer architecture

2. *Architecture of grounding capability* The system is the UTS Unleashed 2003 robot team [1], which is composed of four Sony AIBO ERS210 robots which are 4-legged mobile autonomous robots. Each robot has a camera, and only uses visual and auditory cues to communicate. The architecture of the system is illustrated in Fig. 5 where the grounding capability is viewed as involving four major subsystems: interaction, perception, conception, and problem solving. *Interaction* involves the exchange of information across interfaces, sensors, and actuators. *Perception* involves the creation, acquisition, management and maintenance of sensorimotor and other cued representations. *Conception* involves the creation, acquisition, management and maintenance of concepts. *Problem Solving* involves the creation, acquisition, management and maintenance of high-level representations such as declarative, procedural, and tacit knowledge used for problem solving, reasoning, and decision making activities.
   All interaction between the outside world and the internal representations takes place via the *interaction subsystem*. The *conception subsystem* and the *problem solving subsystem* are embedded in the *perception subsystem*. The conception, problem solving, and perception subsystems can communicate with each other directly. Overall robot behavior is driven by the problem solving subsystem which communicates to the actuators in the interaction subsystem via the perception subsystem.
3. *Scope of the analysis* The analysis will focus on the representation of visual and actuator information onboard the robot. All the architectural subsystems will be involved in the analysis. Only activities related to grounding within the robot are to be considered, i.e., the human designer's role in grounding is outside the scope of the analysis. Representations are considered to be data captured by the robot software system only.
4. *Nature of the grounding capability* The world model (field configuration) of the localization subsystem is grounded through visual information acquired via the robots' camera and a high level model of the robot body. The world model representation can be visualized via a 2D picture of the field with objects identified place in their perceived location, and subsequently evaluated.
5. *Groundedness qualities* Due to the lack of space we briefly describe a few of the more pertinent groundedness qualities introduced in Sect. 5.5.

   Expressiveness The robots interact with the environment through sensors and actuators. The sensor under analysis is the camera which uses YUV values for each pixel. Parameters for motion are sent and received from motors in the robot's body. Perception for the purpose of this analysis involves vision and control of actuators to achieve bodily movements such as walking and kicking. Conception creates and maintains the following concepts: *physical objects* [ball, beacons, goal, team mates, opposition robots], *abstract objects* [player positions, attack, defend, strategy], *physical relationships* [behind, inside penalty area], *actions* [search, kick, walk], *events* [game start, game restart, game end, kick-off]. The problem solving subsystem constructs a representation of the location of objects such as the ball, team mates, opposition robots, and based on it the robot determines its next action.

Relevance Only relevant soccer related entities are represented. They include ball, goals, beacons, team-mates, and opposition robots.

Faithfulness Each robot builds and maintains a representation of the field configuration. The extent to which the field configuration representation is faithful to the real configuration can be measured using the similarity measure developed in Karol et al. [33] which measures the *distance* from one field configuration to another, and it provides a means to explicitly measure the distance/similarity between the *real* configuration of the field and its representation built by the robot as it moves its body and analyzes its raw camera data.

Timeliness The robots are fairly responsive to changes in field configurations and in particular to changes of ball locations. Robot response times can be easily measured and quantified using a wide range of methods at many levels of granularity.

Transparency is low due because almost all representation management is buried in C++ code.

Robustness The grounding capability is robust to field surfaces, but not robust to minor changes in lighting. Specific measurements can be made regarding the lighting levels and the roughness of playing surfaces to determine the range of tolerance.

Adaptability The grounding capability is not adaptable. It cannot make any changes to itself, nor is the system learning "on-the-fly" during a soccer match.

Self-awareness The grounding capability is aware of some of its internal settings such as neck angles, appendages, touch button states, motor parameters. It cannot recognize its own body parts if it perceives them, i.e., it cannot recognize its own feet.

Awareness of Others is achieved through visual and auditory cues only.

6.2 Comparing systems grounding capability

In this section we briefly compare the UTS Unleashed! 2003 System described above with the UTS Unleashed! 2004 System [13]. The 2004 System upgraded the 2003 System including the actual hardware from Sony AIBO ERS210 to ER7 (see Fig. 6) and it possesses the same overall objectives and underlying grounding infrastructure with a few important extensions. The scope of the analysis is the same as for Sect. 5.1 and the nature of the grounding capability of the 2004 System has been extended to include sharing vision information about field configurations among the robots via a wireless network, so that each robot has access to information such as the ball's location from other team members.

Groundedness qualities

Expressiveness The 2004 system can represent all the entities representable in the 2003 version as well as the concept of a shared ball. Shared ball is a fusion of all the robots' estimates of where the ball is.

Interaction Robots share information from their world models via a wireless network, in other words full field configurations are represented in the interaction subsystem. In addition improvements were made to the walking engine so that machine learning techniques such as reinforcement learning with self detection and correction could be applied to improve walking speeds.

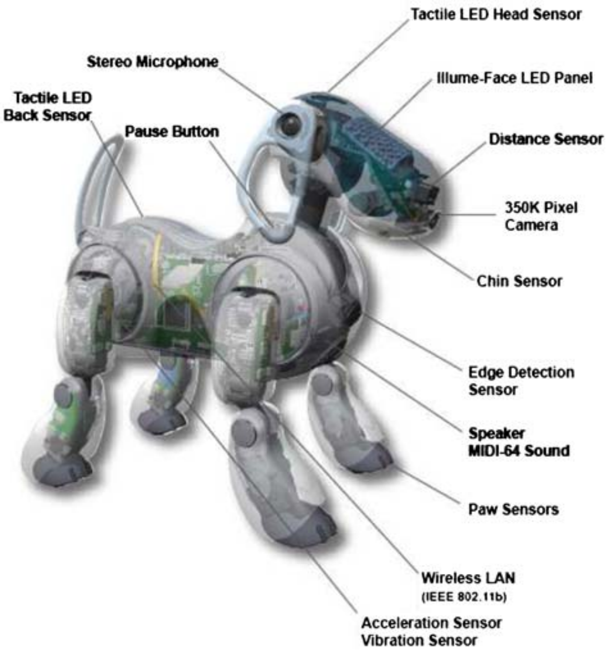

Fig. 6 A Sony ERS-7 AIBO robot (picture Sony Corporation)

Perception Major improvements in the 2004 Systems include (i) the relationship between YUV values of pixels and symbolic colors can be one to many, rather than one-to-one as in the 2003 system which allows for overlapping colors and more flexibility in identifying objects [49], (ii) the velocity of the ball is perceived which supports new high level skills such as passing, catching, and diving, and (iii) field line recognition by perception subsystem.

Conception new object recognition for field lines, new skills conceived (dodge, dive, catch, pass) and implemented as actions. New strategies that exploit the new skills and perception grounding capabilities were also developed.

Problem solving Robots in the 2004 System can *share information* derived from their world model representation such as the location of the ball and the location other robots [33]. Robots on the team who cannot perceive objects directly can be alerted to their location from team mates. For example, if the goal keeper can not see the ball with its own camera, it can face the location of the ball as perceived by its team mates. In addition, using the shared information they can localize using the ball's location and their internal body sensors.

For the purposes of illustration we make brief comments about some of the other qualities. The *relevant* entities represented in the 2003 system are represented in the 2004 system. *Faithfulness* is evaluated using the visualization of the world model representation built by each robot. The similarity measure developed in Karol et al. [33] which measures the *distance* from one field configuration to another, allows us to explicitly measure the distance between the *real* configuration of the field and the configuration represented by the robot. Based on our experimental testing the 2004 System was more faithful than 2003. The 2004 System also turned out to be significantly more *accurate, responsive, transparent*, and *robust* to changes in lighting conditions (due to the one-to-many relationship between pixels and symbolic colors) than the 2003 System. The 2004 System was *aware* of its internal power levels and the 2003 system was not, and furthermore it had a heightened *awareness of others* because high level representations regarding the game's current state (e.g., the location and heading of each robot, the location of the ball, etc.) were communicated directly via the wireless network, and as a result it was more *adaptable* because if a robot was unable to "see" the ball then his teammates could broadcast the ball's location via the wireless network. In addition, in the 2004 System the robot's movements were more adaptable due to the incorporation of machine learning techniques in the walking engine.

6.3 Measuring groundedness in system design

Prioritized quality rankings can be generated from the Grounding Framework by attaching levels of priority to the groundedness qualities. The resultant priority rankings can then be used to evaluate grounding capabilities during system designs. Tailored rankings can be developed for each system and used to develop system requirements.

A priority ranking of groundedness qualities is developed during requirements engineering where more important grounding qualities are higher in the priority ranking. For example, adaptability may be less important than faithfulness for a specific system to achieve its design goals. Belief revision techniques can then be used to maintain the ranking over time automatically [63], and the ranking becomes the standard against which the design is evaluated.

Design decisions should respect the priority ranking. Clearly design and implementation decisions will impact on various qualities in different ways and the key idea is to ensure that high priority qualities are maintained in preference to lower ranked qualities whenever faced with a choice. Given the interrelationships that can exist between the groundedness qualities, sometimes trade-offs will be necessary. Due to the potential complexity automated tools could also be developed to assist. Increasing efficiency is well-known to negatively impact most other qualities regardless of how we choose to rank them. Identifying a priority ordering of qualities is standard practice in software quality assessments. Some groundedness qualities could be identified as so crucial that they must be part of the design and should not be sacrificed for the sake of improving other qualities including efficiency.

Dimensions of groundedness can be graded according to their importance. A ranking that reflects the importance of the qualities determined in requirements allows system developers to understand and evaluate grounding capabilities. A typical Grounding Quality Ranking captures certain values that should be preserved during development and an example is illustrated below:

Rank 1 Essential—Failure to meet the stated degree of qualities will result in complete failure of the system's faithfulness and transparency.
Rank 2 Important—Failure to meet the stated degree of qualities will result in a system with certain kinds of problems of robustness.
Rank 3 Desirable—Failure to meet the stated degree of qualities will result in less flexibility than the desired level of adaptability.

Different rankings for different systems will reflect the design goals and values. Different design goals and values will lead to different priorities. For example we would expect that a robot soccer system designed for winning would have a different ranking of grounding qualities that a system designed for innovative play!

The ranking of the qualities does not necessarily lead to a natural order of consideration in the design process for example the quality of faithfulness may be less important than accuracy but should be considered before it in the design process. In other words, the ranking of groundedness qualities is only related to the evaluation of the groundedness system, and not the order they need to be developed. The ranking of groundedness qualities can give clues as to what should be taken into consideration at design time. An explicit shared ranking of groundedness dimensions simplifies and improves decision making and design decisions within a development team.

7 Discussion

Grounding of representations is an important capability for intelligent systems. Despite its importance there has not been a practical framework that could be used to describe, evaluate or compare grounding capabilities. The Grounding Framework presented herein supports the identification and articulation of important similarities and differences in grounding capabilities across systems, and can be used to demonstrate how and why one system is grounded *better* than another. For the purpose of designing more effective intelligent systems it is important to be able to articulate why one system has a better grounding capability than another, or to say things like if system A's grounding capability had certain properties then it would have an equivalent or better grounding capability than system B.

The Grounding Framework has lead to a deeper and richer understanding of grounding capabilities. It provides guidance on how to evaluate grounding capabilities, to compare grounding capabilities across several systems, and to build more effective grounding capabilities. Moreover, by developing a better understanding of grounding the framework has allowed us to isolate new research problems, challenges, and directions. For example the framework raises the following research questions: (i) Tarski [59] developed a powerful *Theory of Truth* but what should a *Theory of Reference* look like? (ii) Is there a relationship between the hierarchy of representations in Sect. 3, the qualities of groundedness in Sect. 5, and consciousness? (iii) How can we build systems capable of reasoning about their own grounding capability and that of other systems?

8 Conclusions and future work

The Grounding Framework has provided a powerful tool that has helped us understand several generations of robot soccer systems and guided our system design. We have also used the Grounding Framework to analyze other team systems and it has assisted us to identify areas where our robot systems representations can be improved.

From 2009 onwards the NAO robot (illustrated in Fig. 7) will replace the Sony AIBO. The Grounding Framework will be the main vehicle that we will use to guide the upgrade from our current AIBO-based sytem to the NAO. We will use the Grounding Framework in a similar way to that used with the upgrade from the ERS210 to the ERS7 AIBO robots. In addition we will use the Grounding Framework to develop an innovative grounding capability for a Bear Robot (also see Fig. 7). The Bear Robot Project seeks to develop a rich inner life for Bear Robot including self-awareness, awareness of agents, motivation, intention and reasoning capabilities. We are currently focusing on self-awareness which is a prerequisite for the robot to "know what it is doing". In contrast, the AIBOs in the soccer team did not possess a grounded representation of self. For example, when saw their own body parts with

Fig. 7 The Tribotix "bear" (*left*) and the Aldebaran Nao (*right*)

their camera, the robot would not realise or recognize the body part as being part of itself. The Bear, however, will be able to "connect" entities it views via its camera with its body description. Currently, the Bear is able to learn by doing, however a grounded self-awareness will allow it to learn by watching. We will use our experiences with the NAO and Bear robots to develop, improve and enrich the Grounding Framework.

Acknowledgements This work is partially supported by grants from the Australian Research Council, and the University of Technology, Sydney and Stanford University who helped fund Professor John McCarthy's month long visits to Sydney in February and November 2004.

References

1. Agnew, N., Brownlow, P., Dissanayake, G., Hartanto, Y., Heinitz, S., Karol, A., et al. (2003). Robot soccer world cup 2003: The power of uts unleashed!
2. Albus, J., & Barbera, A. (2005). RCS: A cognitive architecture for intelligent multi-agent systems. *Annual Reviews in Control, 29*(1), 89–99.
3. Alchourron, C. E., Gardenfors, P., & Makinson, D. (1985). On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic, 50*, 510–530.
4. Barsalou, L. (1999). Perceptual symbol systems. *Behavioral and Brain Sciences, 22*, 577–609.
5. Brennan, S. (1998). The grounding problem in conversations with and through computers. In S. R. Fussell & R. J. Kreuz (Eds.), *Social and cognitive psychological approaches to interpersonal communication* (pp. 201–225). Hillsdale, NJ: Lawrence Erlbaum.
6. Brooks, R. A. (1990). Elephants don't play chess. *Robotics and Autonomous Systems, 6*, 3–15.
7. Brooks, R. A. (1991a). Intelligence without reason. In *Proceedings of 12th international joint conference on artificial intelligence*, Sydney, Australia, pp. 569–595.
8. Brooks, R. A. (1991b). Intelligence without representation. *Artificial Intelligence, 47*, 139–160.
9. Brooks, R. A. (1991c). New approaches to robotics. *Science, 253*, 1227–1232.
10. Cahn, J., & Brennan, S. (1999). A psychological model of grounding and repair in dialog. In *Proceedings of AAAI fall symposium on psychological models of communication in collaborative systems, American Association for Artificial Intelligence*, North Falmouth, MA, pp. 25–33.
11. Cangelosi, A., & Harnard, S. (2001). The adaptive advantage of symbolic theft over sensorimotor toil: Grounding language in perceptual categories. *Evolution of Communication, 4*(1), 117–142.
12. Chalmers, D. (1992). Subsymbolic computation and the Chinese room. In J. Dinsmore (Ed.), *The symbolic and connectionist paradigms: Closing the gap*. Lawrence Erlbaum.
13. Chang, M., Dissanayake, G., Gurram, D., Hadzic, F., Healy, G., Karol, A., et al. (2004). Robot world cup soccer 2004: The magic of uts unleashed!
14. Christiansen, M., & Chater, N. (1993). Symbol grounding—the emperor's new theory of meaning? In *Proceedings of the 15th annual conference of the cognitive science society*, pp. 155–160.

15. Clark, H., & Brennan, S. (1991). Grounding in communication. In L. B. Resnick, J. M. Levine, & S. D. Teasley (Eds.), *Perspectives on socially shared cognition* (pp. 127–149). Washington, DC: APA.
16. Clark, H., & Schaefer, E. (1989). Contributing to discourse. *Cognitive Science, 13*, 259–294.
17. Cohen, P., Oates, T., Beal, C., & Adams, N. (2002). Contentful mental states for robot baby. In *Eighteenth national conference on artificial intelligence, American Association for Artificial Intelligence*, Menlo Park, CA, USA, pp. 126–131.
18. Coradeschi, S., & Saffiotti, A. (2000). Anchoring symbols to sensor data: Preliminary report. In *AAAI/IAAI*, pp. 129–135.
19. Coradeschi, S., & Saffiotti, A. (2003). An introduction to the anchoring problem. *Robotics and Autonomous Systems, 43*(2–3), 85–96. Special issue on perceptual anchoring. Online at http://www.aass.oru.se/Agora/RAS02/.
20. Dreyfus, H. (1993). *What computers still can't do: A critique of artificial reason* (3rd ed.). London, England: The MIT Press.
21. Dubois, D., & Prade, H. (1994). Possibilistic logic. In D. M. Gabbay, C. J. Hogger, J. A. Robinson & J. H. Siekmann (Eds.), *Handbook of logic in artificial intelligence and logic programming* (pp. 439–513). Oxford University Press.
22. Dudek, G., & Jenkin, M. (2000). *Computational principles of mobile robotic* (1st ed.). London, England: Cambridge University Press.
23. Floridi, L. (2004). Open problems in the philosophy of information. *Metaphilosophy, 35*, 554–582.
24. Gärdenfors, P. (2003). *How Homo became Sapiens: On the evolution of thinking*. London: Oxford University Press.
25. Gärdenfors, P., & Williams, M. A. (2007). Multi-agent communication, planning, and collaboration. In A. C. Schalley & D. Khlentzos (Eds.), *Mental states: Language and cognitive structure* (pp. 197–253). Amsterdam: John Benjamins Pub Co.
26. Gibson, J. J. (1979). *The ecological approach to visual perception*. Boston: Houghton-Mifflin.
27. Glenberg, A., & Kaschak, M. (2002). Grounding language in action. *Psychonomic Bulletin and Review, 9*, 558–565.
28. Goldstein, E. B. (Ed.). (2006). *Sensation and perception*. Wadsworth Publishing.
29. Grazino, M., Hu, T., & Cross, C. (1997). Coding the locations of objects in the dark. *Science, 277*, 239–241.
30. Harnad, S. (1990). The symbol grounding problem. *Physica, 42*, 335–346.
31. Harnad, S. (2003). The symbol grounding problem. In *Encyclopedia of cognitive science*. London: Nature Publishing Group/Macmillan.
32. Johnston, B., & Williams, M. A. (2009). A formal framework for the symbol grounding problem. In *The proceedings of the second conference on artificial general intelligence*. http://www.agi-09.org/papers/paper_20.pdf.
33. Karol, A., Nebel, B., Stanton, C., & Williams, M. A. (2003). Case-based game-play in the robocup four-legged league. In D. Polani, B. Browning, A. Bonarini, & K. Yoshida (Eds.), *RoboCup 2003: Robot soccer world cup VII* (pp. 739–747). Springer.
34. Kennedy, C. (1998). A conceptual foundation for autonomous learning in unforeseen situations. In *Proceedings of the IEEE international symposium on intelligent control (ISIC/CIRA/ISIS'98)*, Gaithersburg, Maryland.
35. MacLennan, B. (1993). Grounding analog computers. *Think, 2*(1), 48–52.
36. Mayo, M. (2003). Symbol grounding and its implication for artificial intelligence. In *Twenty-sixth Australian computer science conference*, pp. 55–60.
37. McCarthy, J. (1956). The inversion of functions defined by turing machines. In C. E. Shannon & J. McCarthy (Eds.), *Automata studies. Annals of mathematical studies* (Vol. 34, pp. 177–181). Princeton University Press.
38. McCarthy, J. (1980). Circumscription—a form of non-monotonic reasoning. *Artificial Intelligence, 13*, 27–39.
39. McCarthy, J. (1998). Elaboration tolerance. In *Working papers of the fourth international symposium on logical formalizations of commonsense reasoning*.
40. McCarthy, J., & Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In B. Meltzer, & D. Michie (Eds.), *Machine intelligence* (Vol. 4, pp. 463–502). Edinburgh: Edinburgh University Press.
41. Pastra, K. (2004a). Viewing vision-language integration as a double-grounding case. In *Proceedings of the American Association of Artificial Intelligence (AAAI) fall symposium on "Achieving human-level intelligence through integrated systems and research"*, Washington D.C., USA.
42. Pastra, K. (2004b). *Vision-language integration: A double-grounding case*. PhD thesis, Department of Computer Science, University of Sheffield, UK.

43. Pfeifer, R., & Verschure, P. (1995). The challenge of autonomous agents: Pitfalls and how to avoid them. In L. Steels & R. Brooks (Eds.), *The artificial life route to artificial intelligence* (pp. 237–263). Hillsdale, New Jersey: Lawrence Erlbaum.
44. Prem, E. (1995). Dynamic symbol grounding, state construction and the problem of teleology. In J. Mira & F. Sandoval (Eds.), *From natural to artificial neural computation. International workshop on artificial neural networks. Proceedings* (pp. 619–626). Berlin, Germany: Springer-Verlag. http://citeseer.ist.psu.edu/prem95dynamic.html.
45. Prince, C. (2001). Theory grounding in embodied artificially intelligent systems. In *Proceedings of the first international workshop on epigenetic robotics: Modeling cognitive development in robotic systems*, Lund, Sweden.
46. Roy, D. (2005). Semiotic schemas: A framework for grounding language in action and perception. *Artificial Intelligence, 167*(1–2), 170–205.
47. Searle, J. R. (1980). Minds, brains, and programs. *Behavioral and Brain Sciences, 3*, 417–457.
48. Stanton, C. (2007). *Grounding oriented design*. PhD thesis, UTS.
49. Stanton, C., & Williams, M. A. (2005). A novel and practical approach towards color constancy for mobile robots using overlapping color space signatures. In *RoboCup 2005*, pp. 444–451.
50. Steels, L. (1995). When are robots intelligent autonomous agents? *Journal of Robotics and Autonomous Systems, 15*, 3–9.
51. Steels, L. (1996). Perceptually grounded meaning creation. In M. Tokoro (Ed.), *Proceedings of the international conference on multi-agent systems*. Cambridge, MA: MIT Press.
52. Steels, L. (2001). Language games for autonomous robots. *IEEE Intelligent Systems, 16*(5), 16–22.
53. Steels, L. (2003). Evolving grounded communication for robots. *Trends in Cognitive Science, 7*(7), 308–312.
54. Steels, L. (2007). The symbol grounding problem is solved, so what's next? In M. D. Vega, G. Glennberg, & G. Graesser (Eds.), *Symbols, embodiment and meaning*. New Haven: Academic Press. http://www.isrl.uiuc.edu/~amag/langev/paper/steels07symbolGrounding.html.
55. Steels, L., & Kaplan, F. (2002). Bootstrapping grounded word semantics. In *Linguistic evolution through language acquisition: Formal and computational models* (pp. 53–74). Cambridge University Press.
56. Steels, L., & Vogt, P. (1997). Grounding adaptive language games in robotic agents. In I. Harvey & P. Husbands (Eds.), *Advances in artificial life. Proceedings of the fourth European conference on artificial Life*. Cambridge, MA: MIT Press.
57. Sun, R. (2000). Symbol grounding: A new look at an old idea. *Philosophical Psychology, 13*(2), 149–172.
58. Taddeo, M., & Floridi, L. (2005). Solving the symbol grounding problem: A critical review of fifteen years of research. *Journal of Experimental and Theoretical Artificial Intelligence, 17,* 419–445.
59. Tarski, A. (1944). The semantic conception of truth and the foundations of semantics. *Philosophy and Phenomenological Research, 4*, 13–47.
60. Tolkien, J. R. R. (1937). *The hobbit*. Ballantine Books.
61. Vogt, P. (2002). The physical symbol grounding problem. *Cognitive Systems Research Journal, 3*(3), 429–457.
62. Williams, M. A. (1994). On the logic of theory base change. In *JELIA*, pp. 86–105.
63. Williams, M. A. (1998). Applications of belief revision. *Transactions and Change in Logic Databases, 1472*(1), 285–314. http://research.it.uts.edu.au/magic/Mary-Anne/publications/BeliefRevisionApplicationsM-AWilliams.pdf.
64. Williams, M. A. (2008). Representation = Grounded Information. In T. B. Ho & Z. H. Zhou (Eds.), *Trends in artificial intelligence: 10th Pacific rim international conference on artificial intelligence. Lecture Notes in Computer Science*. Germany: Springer.
65. Ziemke, T. (1999). Rethinking grounding. *Understanding representation in the cognitive sciences* (pp. 177–190). New York: Plenum Press.