Grounded Models

Josefina Sierra Santibáñez
Escuela Técnica Superior de Informática
Universidad Autónoma de Madrid
28049 Madrid, Spain
Email: Josefina.Sierra@ii.uam.es

Abstract

We introduce *grounded models* and compare them to axiomatic models of mathematics. Grounded models are constructed by an autonomous agent connected to its environment through sensors and actuators using some conceptualization mechanisms described in (Steels 1999). They support a form of ***intuitive reasoning***, which is based on conceptualization and it is argued to be the basis of axiomatization. This is illustrated with a simple example of spatial reasoning.

Introduction

The concept of *grounded model* introduced in this paper is based on some ideas and mechanisms described in *The talking heads experiment* (Steels 1999). The experiment involves a set of robotic "Talking Heads" playing *language games* with each other about scenes perceived through their cameras on a white board in front of them. In particular, we focus on the conceptualization part of a language game, called the discrimination game, which generates the meaning of the verbal hint transmitted by the speaker to the hearer in a language game.

The *discrimination game* (Steels 1996) is played by a single agent, and consists of the following steps. First, the agent perceives an image on a white board through his camera, segments the image into coherent units, and computes various sensory characteristics about each image segment, such as its color, horizontal or vertical position. Then, the agent chooses an image segment as topic, and tries to find a combination of categories that distinguishes the topic from the other objects in the image. The game succeeds if the agent finds a combination of categories that is true for the topic, but it isn't true for the other objects in the image. If the game fails, the agent adapts its internal structures to become more successful in future games.

The rest of this section describes, in some detail, the different processes and cognitive structures involved in the discrimination game.

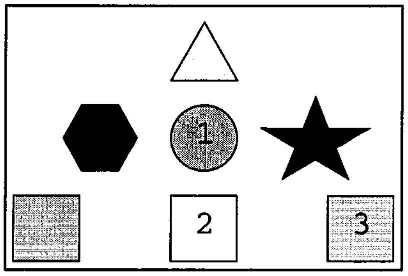

Figure 1: Each object in the scene is characterized by values on three primitive sensory channels: HPOS, VPOS and GRAY.

Perception

The agent captures an image on a white board in front of him. Geometric figures of various sizes, shapes and colors can be pasted on the white board. Then, the agent segments the image into coherent units. Next low level visual processes gather information about each segment, such as its color, horizontal or vertical position. Each process outputs its information on a *sensory channel*. We assume that there are only three primitive sensory channels.

- HPOS(obj) contains the x-midposition of object obj.
- VPOS(obj) contains the y-midposition of object obj.
- GREY(obj) contains the average gray-scale of obj.

The values on the sensory channels HPOS, VPOS and GREY are scaled so that its range is the interval (0.0 1.0). Consider the three objects numbered in figure 1, object 1 has the values HPOS(Obj1)=0.5, VPOS(Obj1)=0.5, GRAY(Obj1)=0.5, object 2 the values HPOS(Obj2)=0.5, VPOS(Obj2)=0.2, GRAY(Obj2)=0.1, and object 3 the values HPOS(Obj3)=0.8, VPOS(Obj3)=0.2, GRAY(Obj3)=0.3.

In addition to three primitive sensory channels, there are some sensory channels constructed from them.

- HPOS-DIFF(obj1,obj2) contains the difference of the x-midpositions of objects obj1 and obj2, i.e. HPOS-DIFF(obj1,obj2)=HPOS(obj1) − HPOS(obj2).

- VPOS-DIFF(obj1,obj2) contains the difference of the y-midpositions of segmented objects obj1 and obj2.
- GREY-DIFF(obj1,obj2) contains the difference of the average gray-scale of objects obj1 and obj2.
- EQUAL(obj1,obj2) is defined as a predicate, i.e. a function which takes the value 1 (i.e. true) if the values of the primitive sensory channels are equal for obj1 and obj2, and 0 (i.e. false) otherwise.

For example, the sensory channel VPOS-DIFF has the value 0.3 when it is applied to the pair of objects formed by object 1 and object 2, i.e. VPOS-DIFF(Obj1,Obj2)=0.3. The range of the sensory channels HPOS-DIFF, VPOS-DIFF and GREY-DIFF is (-1.0 1.0). The range of the sensory channel EQUAL is the discrete set of Boolean values $\{0, 1\}$.

Perceptually Grounded Categories

The data on sensory channels are values from a continuous domain (except for sensory channel EQUAL). To be the basis of natural language communication, these values must be transformed into a discrete domain. One means of categorization is to divide up each domain of values on a particular sensory channel into regions and assign a category to each region. For example, the HPOS channel can be cut in two halves leading to a distinction between [LEFT] ($0.0 < \text{HPOS} < 0.5$) and [RIGHT] ($0.5 < \text{HPOS} < 1.0$). Object 3 in figure 1 has the value HPOS=0.8 and would therefore be characterized as [RIGHT]. Similarly, the VPOS-DIFF channel can be cut in two halves as well leading to a distinction between [ABOVE] ($0.0 < \text{VPOS-DIFF} < 1.0$) and [BELOW] ($-1.0 < \text{VPOS-DIFF} < 0.0$).

It is always possible to refine a distinction by dividing its region. Thus an agent could divide the bottom region of the HPOS channel (categorized as [LEFT]) in two subregions [TOTALLY-LEFT] ($0.0 < \text{HPOS} < 0.25$), and [MID-LEFT] ($0.25 < \text{HPOS} < 0.5$). The categorization networks resulting from these consecutive binary divisions form *discrimination trees.*

We label categories using the sensory channel from which they operate, followed by the upper and lower bound of the region they carve out. Thus [TOTALLY-LEFT] is labeled as [HPOS 0.0 0.25], because it is true for a region between 0.0 and 0.25 on the HPOS channel. We assume that perceptually grounded categories correspond to n-ary predicates of first order logic. For example, the category [HPOS 0.0 0.25] corresponds to the unary predicate [HPOS 0.0 0.25](obj), and the categories [VPOS-DIFF 0.0 1.0] and [EQUAL] to the binary predicates [VPOS-DIFF 0.0 1.0](obj1,obj2) and [EQUAL](obj1,obj2).

There are other ways to move from the continuous domain of sensory channels to the discrete domain of categories. We could introduce focal values and associate a category with each of them. In this case, the categorization process consists in identifying the focal point that is closest to an object's value.

Concepts

Perceptually grounded categories can be combined to construct *concepts.* We use a set of concepts which can be defined by induction as follows.

1. If p is an n-ary category and $x_1, \ldots, x_n$ are variables, then $p(x_1, \ldots, x_n)$ is a concept.
2. If c is a concept, then the negation of c (written $\neg c$) is a concept.
3. If c_1 and c_2 are two concepts, then the disjunction of c_1 and c_2 (written $c_1 \vee c_2$) is a concept.

The symbols $\wedge$, $\rightarrow$ and $\leftrightarrow$ are introduced as abbreviations: (1) $c_1 \wedge c_2$ is an abbreviation of $\neg(\neg c_1 \vee \neg c_2)$; (2) $c_1 \rightarrow c_2$ is an abbreviation of $\neg c_1 \vee c_2$; (3) $c_1 \leftrightarrow c_2$ is an abbreviation of $(c_1 \rightarrow c_2) \wedge (c_2 \rightarrow c_1)$. Notice that the set of concepts is the set of free-quantifier formulas that can be constructed from the predicates associated with perceptually grounded categories.

For example, the concept $\neg$[VPOS-DIFF 0.0 1.0]$(obj1, obj2) \wedge$ $\neg$[VPOS-DIFF 0.0 1.0]$(obj2, obj1)$ is true for a pair of segmented objects obj1 and obj2 if neither obj1 is above obj2, nor obj2 is above obj1. That is, they have the same value on the VPOS-channel.

Some concepts have the property of being true for every possible tuple of segmented objects. We will call them *theorems.* For example, the concept [VPOS-DIFF 0.0 1.0]$(obj1, obj2) \rightarrow$ $\neg$[VPOS-DIFF 0.0 1.0]$(obj2, obj1)$ is a theorem.

Other concepts have the property of being false for every possible tuple of segmented objects. We will call them *inconsistencies.* For example, the concept [VPOS-DIFF 0.0 1.0]$(obj1, obj1)$ is an inconsistency.

The rest of the concepts, called *regular concepts* or simply *concepts,* can be used to discriminate those tuples of objects that satisfy them from those tuples of objects which do not make them true.

Grounded Models

Categorizers

A *categorizer* is a cognitive procedure capable of determining whether a category applies or not. For example, the behavior of the categorizer for category [VPOS 0.0 0.5](obj) can be described by a function that takes the value 1 if $0.0 < \text{VPOS(obj)} < 0.5$, and 0 otherwise. The relation between categories and categorizers is a mapping from syntax to semantics. A category, such as [VPOS 0.0 0.5](obj), is a syntactic expression we use to refer to the fact that its categorizer (i.e. its meaning) holds for object obj. As categories, categorizers can be organized in discrimination trees.

The categorizers of concepts are compositions of the categorizers of categories. For example, the behavior of the categorizer for concept [VPOS-DIFF 0.0 1.0]$(obj1, obj2)\} \wedge$ $\neg$[VPOS-DIFF 0.0 1.0]$(obj2, obj1)$ can be described by a function that takes the value 1 if 0.0 <

[VPOS-DIFF 0.0 1.0]$(obj1, obj2) < 1.0$, and [VPOS-DIFF 0.0 1.0]$(obj2, obj1) \leq 0.0$.

We distinguish again between concepts and categorizers. Concepts are syntactic expressions we use to refer to non-primitive categorizers. The systematic relation between concepts and categorizers (i.e. their meanings) is established by the following interpretation function G.

1. If p is an n-ary category and $x_1, \ldots, x_n$ are variables, then $G(p(x_1, \ldots, x_n))$ is the cognitive procedure that applies the categorizer associated with category p to the segmented objects associated with the variables $x_1, \ldots, x_n$.
2. If c is a concept of the form $\neg c_1$, then $G(c)$ is the cognitive procedure that applies the logical operator of negation to the meaning of c_1, i.e. $G(c) = \neg G(c_1)$.
3. If c is a concept of the form $c_1 \vee c_2$, then $G(c)$ is the cognitive procedure that applies the logical operator of disjunction to the meanings of c_1 and c_2, i.e. $G(c) = G(c_1) \vee G(c_2)$.

Grounded Models

A *grounded model* is the set of categorizers constructed by an agent at a given time. These categorizers are organized as follows.

1. A discrimination tree for every sensory channel containing the categorizers which use the values of that sensory channel.
2. A set of axioms containing the categorizers of a set of "independent" theorems, i.e. concepts which are true for every possible tuple of segmented objects.
3. A set of concepts containing the categorizers of a set of concepts which are neither theorems nor inconsistencies.

The axioms and concepts are organized according to their abstraction level, success rate, usage or arity.

A grounded model reflects the conceptualization of the world constructed by an agent at a given time in his development history. Grounded models are not static, but evolve and adapt as the agent is confronted with new experiences. They are nonmonotonic in the sense discussed in (McCarthy 1980), (McCarthy 1986) and (Lifschitz 1993), but go an step further by allowing the extension of the set of basic concepts (i.e. the set of non-logical symbols of the language).

Consider the grounded model G of an agent that has constructed top-level categorizers for each sensory channel, but does not have any concept or axiom yet. The categorizers of this grounded model are cognitive procedures whose behavior can be described by linear constraints. We associate a linear constraint describing the behavior of each categorizer with a predicate symbol that corresponds to the category it is capable of recognizing.

[VPOS 0.0 0.5]$(x) \equiv 0.0 < \text{VPOS}(x) \wedge \text{VPOS}(x) < 0.5$

[VPOS 0.5 1.0]$(x) \equiv 0.5 < \text{VPOS}(x) \wedge \text{VPOS}(x) < 1.0$

$\vdots$

[GREY-DIFF -1.0 0.0]$(x, y) \equiv$
$-1.0 < \text{GREY}(x)\text{-GREY}(y) \wedge \text{GREY}(x)\text{-GREY}(y) < 0.0$

[GREY-DIFF 0.0 1.0]$(x, y) \equiv$
$0.0 < \text{GREY}(x)\text{-GREY}(y) \wedge \text{GREY}(x)\text{-GREY}(y) < 1.0$

Categorizers are probably implemented by neural networks in natural agents. We only use linear constraints to model their behavior. We are not assuming therefore that agents do learn such constraints, but that they build them in their perceptual systems.

Grounded Models vs Axiomatic Models

A grounded model looks like an axiomatization or a theory of first order logic, but differs from it in some aspects. Like a logical theory, it has a set of basic concepts (which correspond to the categorizers of discrimination trees), defined concepts and axioms[1].

One of the most important differences between grounded models and logical theories is the form in which concepts are defined. In a logical theory, concepts are defined by simple symbols whose meaning is determined or constrained by the relationships the axioms in the theory postulate about them. In a grounded model, however, concepts are defined by cognitive procedures which given a tuple of objects return one of the Boolean values $\{0, 1\}$. For basic concepts, these cognitive procedures are the categorizers associated with perceptually grounded categories in discrimination trees. For regular concepts, the cognitive procedures are compositions of the procedures associated with basic concepts. Concepts have, therefore, precise and explicit meanings in grounded models, which determine the set of concepts that can constitute the theorems of a grounded model. This can be contrasted with the meanings of concepts in logical theories, which are implicitly defined by the relationships the axioms in the theory postulate about them.

It should be observed as well, that the meanings of basic concepts in grounded models are not arbitrary procedures from the set of possible tuples of segmented objects on the Boolean values $\{0, 1\}$, as it happens in model theory semantics. They are cognitive procedures which should make intuitively plausible distinctions (in general, relations) using data extracted by realistic sensory channels operating on real world environments. This fact has important consequences on the shape and inferential power of grounded models. First, the meanings of basic concepts are highly constrained by physical aspects of the environment and the sensory-motor apparatus of the agent. Because, in order to qualify

[1] In logical theories, basic concepts are specified by the non-logical symbols of the language, and defined concepts by mathematical abbreviations or definitions.

as a possible meaning, a cognitive procedure has to be implementable as a relatively easy computation on the range of values produced by realistic sensory channels. Second, the set of possible axioms is constrained as well by the mathematical relationships holding among the meanings of basic concepts. For example, the following relationship holds among the meanings of the concepts [VPOS 0.0 0.5]$(obj1)$, $\neg$[VPOS 0.0 0.5]$(obj2)$ and [VPOS-DIFF 0.0 1.0]$(obj1, obj2)$.

$$[\text{VPOS } 0.0\ 0.5](obj1) \land \neg[\text{VPOS } 0.0\ 0.5](obj2) \rightarrow \neg[\text{VPOS-DIFF } 0.0\ 1.0](obj1, obj2)$$

Therefore, the following concept can never be an axiom of a grounded model containing the basic concepts [VPOS-DIFF 0.0 1.0]$(obj1, obj2)$ and [VPOS 0.0 0.5]$(obj1)$.

$$[\text{VPOS } 0.0\ 0.5](obj1) \land \neg[\text{VPOS } 0.0\ 0.5](obj2) \land [\text{VPOS-DIFF } 0.0\ 1.0](obj1, obj2)$$

The mathematical relationships holding among the meanings of basic concepts are one of the most important features of a grounded model. As we will see later on, they allow agents to reason about their environment without having an explicit axiomatization of it. This form of *intuitive reasoning* is commonly used by people to reason about everyday problems, and it is also the basis of formal reasoning in the sense of mathematics. In next section, we explain how a grounded model is constructed by an agent as it interacts with the environment. The axiomatization task, which is part of this model building process consists in discovering mathematical relations that hold among the concepts of a grounded model, and storing them in the repertoire of axioms.

We can observe a dual character between grounded models and logical theories. Grounded models are constructed around the notion of concept, i.e. the meanings of concepts determine the set of theorems of a grounded model. Logical theories are instead constructed around the notion of axiom, i.e. the set of axioms determine the set of theorems of a theory, and to some extent the meanings of concepts. By looking at our example of grounded model G, one could say that grounded models are logical theories whose axiom set consists only of definitions of concepts. But this is not true, because the meanings of concepts in grounded models are cognitive procedures rather than logical formulas. Sometimes, the behavior of these procedures can be described by mathematical functions, as in the case of grounded model G, but other times it cannot be easily described this way, as in the case of categorizers for approximate concepts. In the first case, intuitive reasoning can be simulated by mathematical methods, as we will see later on, but in the second one different techniques must be applied. Natural agents manage to reach sound conclusions in both cases, and we try to propose a framework in which both cases can be described.

Constructing Grounded Models

Grounded models are constructed as a side effect of agents' activity. In particular, the grounded models studied in the paper are constructed by the agents as they play discrimination games. A *discrimination game* (Steels 1996) is played by a single agent. The agent perceives a scene and chooses a topic from the possible tuples of segmented objects in the scene. He then uses his current grounded model to come up with a category or concept that is valid for the topic, but not for any other tuple of objects in the context. The game succeeds if the agent can find such a category or concept. If the game succeeds, the use and success counters of the categorizers involved go up[2]. If the game fails, the use counters of the categorizers involved go up, and a repair process in which a new category or concept is generated takes place.

Initially, the agent constructs top-level categorizers for each sensory channel that has contained distinctive data in the recent past. If a channel has the same data for every segment it is not going to be possible to find a distinctive category from it. Afterwards, the agent extends his discrimination trees or constructs new concepts from existing categories or concepts.

A categorizer for a new category is constructed by taking a categorizer node in a discrimination tree and dividing its range into two new subranges. For example, if we take the categorizer [HPOS 0.0 0.5], which is true when the object is in the left most half of a scene, two new categorizers are created by dividing [0.0 0.5] into two halves, one for the range [0.0 0.25] ([HPOS 0.0 0.25] or totally left), and one for the range [0.25 0.5] ([HPOS 0.25 0.5] or mid left). A new categorizer is added to the tree for each of these halves.

A categorizer for a new concept is constructed by composition of the categorizers of existing categories or concepts. We use three composition operations: negation, disjunction and substitution. These operations allow constructing every free-quantifier formula of the first order language defined by the categories of the grounded model. In general, new concepts are constructed from categories or concepts that have been useful in previous games, i.e. which have a high success rate. Two logical compositions are preferred to relate existing categories and concepts: conjunction and implication. Conjunctions tend to be more discriminating than their components, increasing the chances of success in future games. Implications, on the other hand, are commonly used to describe causality, and world facts are often expressed as causal laws.

The final step of the discrimination game, called *assimilation*, is as follows. Whenever the agent constructs a new concept, he checks whether it is a theorem or an

[2]The use and success counters of a categorizer are used for different purposes, such as computing the categorizer rate (which is the result of dividing success by use), or choosing which categories or concepts should be used to create new concepts.

inconsistency. If the new concept is a theorem and it is independent[3] with respect to the axioms in his current grounded model, the concept is added to the axiom set of his grounded model. If it is a theorem but it is not independent of his grounded model axioms or if it is an inconsistency, the concept is ignored and it is not included in the grounded model. Finally, if the concept is a regular concept, and it is not equivalent to some concept in his current grounded model, it is added to the set of concepts of the grounded model.

Intuitive Reasoning

The process by which an agent tries to determine whether a concept is a theorem, a regular concept or an inconsistency of a grounded model is called *intuitive reasoning.* A theorem of a grounded model is a concept whose meaning is true for every possible tuple of segmented objects. A regular concept is a concept whose meaning is true for some tuples of objects, but false for others. And an inconsistency is a concept whose meaning is false for every possible tuple of segmented objects.

We hypothesize that intuitive reasoning happens by a process of simulation in natural agents. First, they construct the meaning of the concept they want to check combining the categorizers of their grounded models. Then, they try to find combinations of values of their sensory channels that can satisfy the meaning. This process can be seen as a form of constraint satisfaction by search, in which the agents generate possible combinations of values for sensory channels, and test them using the concept's meaning. Of course, the search cannot be exhaustive, because sensory channels take values on a continuous domain. The search is performed at the level of subregions in which the categorizers involved take different values. Each subregion is represented by a single value of a sensory channel, and the combinations of subregions for sensory channels by tuples of values. As soon as a value is shown incompatible with the concept's meaning or sufficient for satisfying it, all the combinations containing that value can be eliminated from the search space. This allows reducing the complexity of the search process. Sometimes, however, the agents cannot explore the entire space of possibilities, because it is too complex, and they make errors when they try to approximate the result of the search.

If, after the search process, the agents cannot find any combination of values which does not satisfy the meaning, the concept is seen as a theorem. If they find some combinations that satisfy it, and others which do not, the concept is seen as a regular concept; and, otherwise, as an inconsistency.

[3]We only discuss in some detail the mechanism used by the agents to determine whether a concept is a theorem. Space limitation prevents us from explaining the mechanisms for determining independence or equivalence of concepts. These mechanisms are by necessity approximated, as intuitive reasoning is.

Spatial Reasoning

To clarify the ideas discussed in previous sections, we compare a first order theory T_S, which can be used for reasoning about spatial relations among objects on the plane, with the grounded model G described in section 2. The language of T_S consists of two binary predicates $A(x, y)$ and $R(x, y)$. Its axiom set is as follows.

$$A(x,y) \rightarrow \neg A(y,x) \quad (1)$$

$$A(x,y) \wedge A(y,z) \rightarrow A(x,z) \quad (2)$$

$$A(x,y) \wedge A(x,z) \wedge \neg R(y,z) \wedge \neg R(z,y) \wedge y \neq z \rightarrow A(y,z) \vee A(z,y) \quad (3)$$

$$R(x,y) \rightarrow \neg R(y,x) \quad (4)$$

$$R(x,y) \wedge R(y,z) \rightarrow R(x,z) \quad (5)$$

$$R(x,y) \wedge R(x,z) \wedge \neg A(y,z) \wedge \neg A(z,y) \wedge y \neq z \rightarrow R(y,z) \vee R(z,y) \quad (6)$$

In the classical approach to AI if one wants to build an agent capable of reasoning about spatial relations, such as *Above* or *Right-of*, one must construct a first order theory like this one and apply automatic theorem proving methods. Theories like this one are constructed by agent designers, not by the agents themselves, and they must be extended or updated by designers as well when the agents are faced with new challenges. These theories are normally used to determine whether a fact follows or not from the axioms of the theory, i.e. whether it is a theorem, or to extract answers in the form of sets of tuples of objects that satisfy a particular property expressed as a logic formula.

We are going to see how a simple grounded model, such as the one shown in section 2, can be used for the same tasks as T_S. The difference is that grounded models are constructed and updated by agents rather than agent designers, and that inference is done by *intuitive reasoning.*

First, we see that every theorem of T_S is also a theorem of the grounded model G. In order to do that, we need to prove that every axiom of T_S is a theorem of G, and that the theorems of G are closed under the inference rule of resolution.

When the behavior of the categorizers for basic concepts can be described by linear constraints, intuitive reasoning can be seen as a process of linear constraint satisfaction. In particular, in the grounded model G, the behavior of the categorizer of every concept can be described by a disjunction of linear constraints which can be computed by replacing every category by a linear constraint describing the behavior of its categorizer in the concept expression, and computing the disjunctive normal form of the result. Proving that the meaning of a concept is true for every possible tuple of segmented objects requires proving that the constraint system associated with the concept is true for every value of every sensory channel. And this is equivalent to proving that the negation of the system is unsatisfiable. Therefore, intuitive reasoning can be performed in that model by showing that the disjunctive normal form of the negation of the concept's meaning is unsatisfiable.

For example, it can be shown that axiom 2 is a theorem of the grounded model by checking that the following disjunction of constraints is unsatisfiable for every value of x, y and z in the interval (0.0 1.0). This system has been obtained replacing every instance of VPOS(x), VPOS(y) and VPOS(z) by x, y and z in the disjunctive normal form of the negation of the axiom's meaning.

$$\{0.5 < x\text{-}y,\ x\text{-}y < 1.0,\ 0.5 < y\text{-}z,\ y\text{-}z < 1.0,\ x\text{-}z \leq 0.5\} \vee$$
$$\{0.5 < x\text{-}y,\ x\text{-}y < 1.0,\ 0.5 < y\text{-}z,\ y\text{-}z < 1.0,\ 1.0 \leq x\text{-}z\}$$

It is easy to check that this constraint system is unsatisfiable. A disjunction of constraints is unsatisfiable if each disjunct is unsatisfiable. And, each disjunct is a linear arithmetic constraint that can be solved by a linear constraint solver, such as the one of Sicstus Prolog.

The rest of the axioms of T_S can be shown to be theorems of the grounded model G by intuitive reasoning as well. It can also be proved that intuitive reasoning in grounded models in which the behavior of the meaning of categories can be described by linear constraints is closed under resolution. That is, if two concepts are theorems of a grounded model, its resolvent is a theorem of the grounded model as well. Therefore, every theorem of T_S can be shown to be a theorem of the grounded model G by intuitive reasoning. And, intuitive reasoning is much simpler than theorem proving for this domain. In fact, the grounded model G can be used to prove many more theorems than T_S, because it contains categorizers for a broader set of basic concepts including the standard notions of *up(x), down(x), dark(x)* and *darker(x,y).*

Grounded models can be used to extract answers as well. For example, we can obtain all the triples of objects in the scene of fig. 1 such that $A(x, y) \wedge A(y, z)$ by applying the categorizer of this concept to every triple of segmented objects in the scene, and picking up those triples that satisfy it.

Notice that, as soon as the agents have constructed categorizers for the basic concepts *above* and *right-of,* they are capable of proving every theorem of T_S by intuitive reasoning. This means that an agent capable of linguistic competence (Steels 1998) at the level of interpreting and generating first order logic formulas does not need an axiomatization to reach sound conclusions about its environment by intuitive reasoning in this domain. It only needs to understand the basic concepts involved, and the grammar rules by which concept expressions are translated into concept meanings.

Finally, grounded models can be used for a task that logical theories do not support, namely, *concept generation.* As a side effect of their activity playing discrimination games, the agents construct new concepts which they store and can use afterwards for different purposes, such as classification, discrimination or communication. Some of these new concepts are theorems. These are stored in the axiom set of a grounded model, and allow agents to learn general facts about their environment. Others are regular concepts, and are stored in the concept set of a grounded model. Axiomatizations and conceptualizations get constructed then as a side effect of agents' activity, and do not have to be built into them.

Acknowledgments

The author would like to thank Luc Steels for many interesting conversations on the topics of the origins of language and grounded representations.

Conclusions

We have introduced grounded models and compared them to axiomatic models of mathematics. Grounded models are based on conceptualization, and support a form of intuitive reasoning which is argued to be the basis of axiomatization. This has been illustrated with an example of simple spatial reasoning.

References

Lifschitz, V. 1993. Circumscription. In Handbook of Logic in Artificial Intelligence and Logic Programming, D. Gabbay and C.J. Hogger, Ed., Oxford University Press.

McCarthy, J. 1959. Programs with common sense. In *Mechanization of Thought Processes, Proc. of the Symp. of the National Physics Laboratory,* 77–84.

McCarthy, J. 1980. Circumscription -a form of non-monotonic reasoning. *Artificial Intelligence* 13:27–39.

McCarthy, J. 1986. Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence* 28:89–116.

McCarthy, J. 1990 Formalizing common sense. Papers by John McCarthy. Edited by Vladimir Lifschitz. Ablex Publishing Corporation, New Jersey.

Shoenfield, J.R. Mathematical Logic. Addison-Wesley Publishing Company.

Steels, L. 1996 Perceptually grounded meaning creation. Proceedings of the International Conference on Multi-Agent Systems. Tokoro, M. (ed.). AAAI Press, Menlo Park Ca. pp. 338-344.

Steels, L. 1998 The origins of syntax in visually grounded agents. Artificial Intelligence 103, 1-24.

Steels, L. 1999 The talking heads experiment. Volume 1. Words and meanings. Special pre-edition for LABORATORIUM, Antwerpen.

Steels, L. 2000 The Emergence of Grammar in Communicating Autonomous Robotic Agents. In Horn, W. (ed.) Proceedings of ECAI 2000. IOS Publishing, Amsterdam.

Steels, L. and P. Vogt 1997 Grounding Adaptive Language Games in Robotic Agents. In Proceedings of ECAL 97. The MIT Press, Cambridge Ma.