Symbol Grounding in Connectionist and Adaptive Agent Models

Angelo Cangelosi

Adaptive Behaviour and Cognition Research Group,
School of Computing, Communications and Electronics,
University of Plymouth, UK
acangelosi@plymouth.ac.uk

Abstract. This paper presents the Cognitive Symbol Grounding framework for modeling language in neural networks and adaptive agent simulations. This approach is characterized by the hypothesis that symbols are directly grounded into the agents' own categorical representations, whilst at the same time having syntactic relationships with other symbols. The mechanism of grounding transfer is also introduced. This is the process by which the grounding of basic words, acquired via direct sensorimotor experience, is transferred to higher-order words via linguistic descriptions. Various simulations are briefly reviewed to demonstrate the use of the Cognitive Symbol Grounding approach.

1 The Cognitive Symbol Grounding Framework

The grounding of symbols in computational cognitive systems requires that the simulated cognitive agent is able to access the meaning of its symbols (words) directly, without the intervention of an external viewer such as the experimenter. This has been a significant shortcoming of cognitive models only based on symbolic architecture. Such a limit is commonly referred to as the Symbol Grounding Problem [6].

Recent cognitive models based on connectionist agents and robots use the Cognitive Symbol Grounding framework to intrinsically link symbols to the agents' own cognitive system [1]. This approach is characterized by the hypothesis that symbols are directly grounded into the agents' own categorical representations, whilst at the same time having logical/syntactic relationships with other symbols. First, each symbol is directly grounded into internal categorical representations. These representations include perceptual categories (e.g. the concept of blue color, square shape, and male face), sensorimotor categories (the action concept of grasping, pushing, carrying), social representations (individuals, groups and relationships) and other categorizations of the agent's own internal states (emotions, motivations). Secondly, these categories are connected to the external world through our perceptual, motor and cognitive interaction with the environment.

Two main modeling approaches to the symbol grounding are presented here: (1) the connectionist approach, based on artificial neural network simulations of

S.B. Cooper, B. Löwe, and L. Torenvliet (Eds.): CiE 2005, LNCS 3526, pp. 69–74, 2005.

categorization and linguistic tasks, and (2) the embodied agent modeling method based on multi-agent simulations and robotic studies. Both approaches share an integrative view of the cognitive systems, in which vision, action and language are intrinsically linked with one another. This permits the design of cognitive systems where language is directly grounded in the agent's own sensorimotor and cognitive abilities.

The use of an integrated language/cognition/action system is particularly important for the mechanism of grounding transfer. This is the process by which the grounding of basic words, acquired via direct sensorimotor experience, is transferred to higher-order words via linguistic descriptions (i.e. indirect grounding). For example, I can learn by direct experience that the word `horse` is grounded on my sensorimotor experience of seeing (and/or riding) a horses and that the word `horn` is grounded on the vision of the horn of an animal. Subsequently, via the linguistic description "`The unicorn is an animal similar to a horse with a horn`" I can learn the new word, and concept, `unicorn` and indirectly ground it in my experience of horses and horns.

The following sections will briefly review some modeling work on the Cognitive Symbol Grounding hypothesis developed by the author and collaborators. These examples will demonstrate the use of connectionist and adaptive agent models for the direct grounding of symbols and the transfer of grounding from low level words to higher-level symbols.

2 Connectionist Simulations

The connectionist approach to symbol grounding is based on simulations of artificial neural networks for category learning and naming tasks. In particular, this has been possible through the use of dual-route neural networks architectures [8] that permit the link (a) between perceptual and sensorimotor representations and (b) between these sensorimotor representations and symbolic knowledge. Typically, a neural network will have visual and linguistic input units indirectly link, via hidden units, to motor and linguistic output units. The process of language understanding can be simulated with the link from linguistic input to motor outputs, while language production links visual input to linguistic output units.

A seminal paper on categorization and symbol grounding with neural networks is that proposed by Harnad et al. [7]. This specifically focused on grounding symbols in categorical perception. The authors used a multi-layer perceptron to categorizing lines according to their length. Training consisted of two sequential backpropagation learning tasks. The first was an autoassociation task to train networks to discriminate between different stimuli. In the second task, the networks were trained to categorize stimuli by sorting lines into three categories: short, middle, long. The comparison of the pre- and post-categorization hidden activations showed the well known categorical perception effects, i.e. within-category compression and between-category expansion of category members. The hidden categorical representations constituted the grounding of categorization names (labels).

Subsequent symbol grounding models have focused on the mechanisms of grounding transfer. For example, Riga et al. [9] proposed a connectionist models in which basic symbols, such as names of shapes and colors, are first intrinsically connected to the categories being acquired through direct interaction with the environment. These basic symbols are successively used to construct descriptions of new categories of stimuli consisting of individual objects made by specific combinations of a shape and color). New categories (and their symbols) are in this way defined without the need of a direct experience of their referents. This process of grounding transfer enables the system to express meanings that go beyond immediate experience. New symbols, acquired exclusively from symbolic descriptions, are ultimately grounded in the interaction of the system with its environment.

The simulation consisted of three sequential training stages and a grounding transfer test phase. During the first training stage, an unsupervised network learns to discriminate between different stimulus categories by constructing a feature self-organizing map of different shape and color categories. The network acquires analogue sensorial representations of their environment that enables it to categorize the stimuli along the dimensions of shapes and colors. In the second training phase, symbolic stimuli (the names of the colors and shapes) are presented to the network, together with the images. These symbols are directly grounded in the sensorial representations acquired during the first phase. In the third training phase, the training input is exclusively symbolic. Linguistic descriptions of new higher-order categories are presented. These contain the previously acquired symbols in combination with a new symbol that denotes a new category (e.g. `Red + Square = DAX`). Finally, in the test phase, the network is presented with images of the previously unseen objects, such as `DAX`, to check if these are recognized and named. The successful naming of these previously-unseen images would demonstrate that the grounding transfer has occurred. Simulation results consistently showed that networks are able to recognize and name the images of new objects, therefore demonstrating that the grounding has been transferred from basic order categories to higher order concepts. Thus, the proposed connectionist simulation provides the basis of a working model for the implementation of an autonomous cognitive system able to use combinations of previously-grounded symbols to expand its knowledge of the world.

Other neural network models of language have focused on the grounding of special types of symbol, that is function words. These includes linguistic terms such as spatial prepositions (e.g. in, on, over, under) and quantifiers (e.g. few, some, many). Recently, Coventry et al. [5] have developed a neural network model of the spatial prepositions over, under, above, below. The model addresses the integration of functional and object knowledge factors (”what”) with geometric factors such as the relative position of objects (”where”). The model processes movies of a located object (teapot) pouring a liquid (water) into a reference object (cup). The task of the network is to name the objects and to select the most appropriate spatial preposition describing the scene. The model consists of three modules: (1) a neurally-inspired vision module to process the visual scenes,

(2) a simple recurrent neural network to learn compressed representations of the dynamics of interacting objects, and (3) a dual-route network for producing the names of objects and the spatial prepositions. The dual-route network plays the core function of the grounding process by integrating visual and linguistic knowledge. The activation values of the linguistic output nodes correspond to rating values given by subjects in language comprehension experiments. The multi-layer perceptron is trained via error backpropagation, by converting the rating data into stimulus presentation frequencies. Simulation results consistently show that the networks produce rating values similar to that of experimental subjects. It can also accurately predicts new experimental data on the ratings of scenes where only the initial frames are shown and the subjects must "mentally replay" the scene and predict its end frame (i.e. where the liquid ends). Such a model is currently being extended to deal with further linguistic terms, such as the vague quantifiers few, a few, some, many, a lot of. The underlying hypothesis is that this grounded connectionist approach will permit the identification of the main mechanisms responsible for quantification judgment and their linguistic expression.

3 Adaptive Agent Simulations

The adaptive agent approach includes multi-agent simulations of the evolution of language and cognitive robotics experiment on communication and language learning. The multi-agent approach uses populations of simulated agents that interact with each other to develop a shared set of symbols (lexicon) to describe their interaction with the world. The robotic studies uses embodied evolutionary and/or epigenetic robotic agents that interact in a simulated (or real) physical environment and build a linguistic representation of this interaction.

In a grounded multi-agent model of language evolution [2], neural networks were used to simulate learning and the genetic algorithm to simulate evolution. The model considers two ways of acquiring categories and language which are in direct competition with one another: In "sensorimotor toil," new categories are acquired through feedback-corrected, trial and error experience in sorting input stimuli. In "symbolic theft," new categories are acquired by hearsay from propositions (i.e. language) based on boolean combinations of symbols. In competition, symbolic theft always beats sensorimotor toil. This is hypothesized to be the basis of the adaptive advantage of language, after basic categories are learned by toil, to avoid an infinite regress (the symbol grounding problem). Changes in the internal representations of categories must take place during the course of learning by toil. These changes were analyzed in terms of the compression of within-category similarities and the expansion of between-category differences. Such compression/expansion effects, called Categorical Perception, have previously been reported with categories acquired by sensorimotor toil. This simulation also shows that they can also arise from symbolic theft alone.

Studies with adaptive robotic agents include simulations of robots that learn to imitate actions and to communicate linguistically about such motor abilities. For example, in an epigenetic robotic model [4], agents learn to perform actions on objects using imitation and grounding transfer mechanisms. The model is based on a simulation of two robots embedded in a virtual environment that accurately models physical constraints present in real-world situations (using the physics software engine Open Dynamics Engine). Each robot has twelve Degrees of Freedom (DoF) and consists of two 3-segment arms attached to a torso and a base with 4 wheels. The teacher robot has preprogrammed behavior to manipulate objects. The imitator agent learns to perform the same actions by observing the teacher executing them and using an on-line backpropagation algorithm. This effectively enables the agent to mimic a movement in real-time and provides the agents with a mechanism to approximate movements without need for prior learning. The robot's neural network memorizes action patterns related to objects and enables the autonomous execution of the movement associated with an object in absence of the teacher agent. The neural controller receives in input sensorial data on the object's visual properties and the proprioceptive information on the imitator's joint angles. In output it produces the motor force applied to each joint. Overall, the simulation results show that it takes just few online training epochs to obtain a satisfactory performance. Typically, the agents produce a movement that is very similar to the original after having mimicked it once, and optimize it during successive training cycles.

Agents simultaneously learn the words corresponding to actions, whist they are are taught to perform the basic actions by mimicking them. Robot learn the basic actions of opening and closing their left and right arms (upper arms and elbows), lifting them (shoulders), and moving forward and backward (wheels), together with the corresponding words. They also learn higher-level composite behaviors by receiving linguistic descriptions containing these previously acquired words (grounding transfer). After basic grounding, the robot receives 1st level linguistic descriptions of combined actions, consisting in a new word and two known words referring to basic actions. For example, the action of grabbing the object in front of them was described as: "`CloseLeft + CloseRight = Grab`". Grounding is successfully transferred from the basic words `CloseLeft` and `CloseRight` to the higher order symbol `Grab`. In a test phase, when the agent is given the command `Grab` it successfully executes the combined action of pushing its arms towards the object and grabbing it. Robots can also receive further higher-level descriptions, in which a defining word is itself learned exclusively from a linguistic description. For example, the grabbing and moving forward actions were combined into the higher-order action of carrying: "`MoveForward + Grab = Carry`". Grounding transfer was successfully transferred to the new word, enabling the agent to correctly perform the action of carrying on hearing the word `Carry`. The system learned several of these combined actions simultaneously, and also four-word definitions and grounding transfers of up to three levels have been realized. In addition to demonstrating the grounding transfer mechanism in robotic agents, this model also highlights the role of language as a cognitive enhancer tool, i.e. a means through which new behaviors can be acquired quickly and effortlessly, building on experience accumulated by previous generations of agents.

4 Conclusions

Overall, these studies demonstrate the feasibility of the Cognitive Symbol Grounding approach in which neural networks are used to deal with the symbol grounding problems and the grounding transfer. The "embodiment" of such connectionist architectures in either simulated agents or robots permits a deeper understanding of the relationship between linguistic/symbolic abilities and other sensorimotor and cognitive capabilities. For example, adaptive agent models of verb and noun learning have shown that linguistic and sensorimotor representations share common neural structures. In these simulations, the same hidden units are involved or the processing of nouns and sensorimotor representations, whilst separate hidden units specialize for verb and motor processing [3]. This approach also has important practical and technological implications. For example, in robotics and artificial intelligence, language grounding models can provide novel algorithms and methodologies for the development of effective interaction between humans and autonomous computer and robotic systems. As demonstrated in the epigenetic robotic model of the symbol grounding transfer, the imitation and language instruction modalities can be integrated to form a situated learning process in which higher-order linguistic representations can be autonomously grounded into the agents' own sensorimotor and cognitive abilities.

References

1. Cangelosi, A.: Grounding symbols in perceptual and sensorimotor categories: Connectionist and embodied approaches. In: Cohen, H., Lefebvre, C. (Eds), Handbook of Categorization in Cognitive Science, Elsevier (2005)
2. Cangelosi, A., Harnad, S.: The Adaptive Advantage of Symbolic Theft over Sensorimotor Toil: Grounding Language in Perceptual Categories. Evolution of Communication **4** (2000) 117–142
3. Cangelosi, A., Parisi, D.: The processing of verbs and nouns in neural networks: Insights from synthetic brain imaging. Brain and Language, **89** (2004) 401–408
4. Cangelosi, A., Riga, T.: Adaptive agent and robotic approaches to the sensorimotor ground-ing of language Artificial Intelligence Journal (in submission)
5. Coventry, K.R.,Cangelosi, A., Rajapakse, R., Bacon, A., Newstead, S., Joyce, D., Richards, L.V.: Spatial prepositions and vague quantifiers: Implementing the functional geometric framework. Spatial Cognition Conference, Germany, 11–13 October (2004)
6. Harnad, S.: The Symbol Grounding Problem. Physica D, **42** (1990) 335–346
7. Harnad, S., Hanson, S. J. and Lubin, J.: Categorical perception and the evolution of super-vised learning in neural nets. In D.W. Powers and L. Reeker (Eds.), Proceedings of Proceedings of the AAAI Spring Symposium on Machine Learning of Natural Language and Ontology (1991)
8. Plunkett, K., Sinha, C., Mller, M.F., Strandsby, O.: Symbol grounding or the emergence of symbols? Vocabulary growth in children and a connectionist net. Connection Science, **4** (1992) 293–312
9. Riga, T., Cangelosi, A., Greco, A.: Symbol grounding transfer with hybrid self-organizing/supervised neural networks. Proceedings of IJCNN04 International Joint Conference on Neural Networks. Budapest, July (2004)