



Basic language learning in artificial animals

Downloaded from: <https://research.chalmers.se>, 2019-05-24 01:26 UTC

Citation for the original published paper (version of record):

Johannesson, L., Nilsson, M., Strannegård, C. (2018)

Basic language learning in artificial animals

Proceedings of the 2018 Annual International Conference on Biologically Inspired Cognitive Architectures

N.B. When citing this work, cite the original published paper.

Postproceedings of the 9th Annual International Conference on Biologically Inspired Cognitive Architectures, BICA 2018 (Ninth Annual Meeting of the BICA Society)

Basic Language Learning in Artificial Animals

Louise Johannesson^a, Martin Nilsson^b, Claes Strannegård^{a,*}

^a*Department of Computer Science and Engineering, Chalmers University of Technology, 41296 Gothenburg, Sweden*

^b*Department of Physics, Chalmers University of Technology, 41296 Gothenburg, Sweden*

Abstract

We explore a general architecture for artificial animals, or animats, that develops over time. The architecture combines reinforcement learning, dynamic concept formation, and homeostatic decision-making aimed at need satisfaction. We show that this architecture, which contains no *ad hoc* features for language processing, is capable of basic language learning of three kinds: (i) learning to reproduce phonemes that are perceived in the environment via motor babbling; (ii) learning to reproduce sequences of phonemes corresponding to spoken words perceived in the environment; and (iii) learning to ground the semantics of spoken words in sensory experience by associating spoken words (e.g. the word “cold”) to sensory experience (e.g. the activity of a sensor for cold temperature) and vice versa.

© 2019 The Authors. Published by Elsevier B.V.

This is an open access article under the CC BY-NC-ND license (<https://creativecommons.org/licenses/by-nc-nd/4.0/>)

Peer-review under responsibility of the scientific committee of the 9th Annual International Conference on Biologically Inspired Cognitive Architectures.

Keywords: generic animat; language learning; babbling; sequence learning; grounded semantics; poverty of the stimulus

1. Main Text

Darwin argued that “The difference in mind between man and the higher animals, great as it is, certainly is one of degree and not of kind.” [1]. Several differences have been found between the human brain and the brains of other primates, e.g. in terms of relative cortex volume, neuron packing density, and axonal conduction velocity, but they all seem to be differences in degree rather than kind [2, 3]. According to Darwin’s perspective, human cognition is a special case of animal cognition and psychology is a special case of ethology. Moreover, not even the ability to use language or mathematics counts as a clear-cut dividing line between *homo sapiens* and the other animals.

On the other hand, Noam Chomsky among others, argued that humans are born with a predisposition for language -a universal grammar- that is unique in the animal kingdom [4]. His evidence for this claim includes the “poverty of

* Corresponding author. Tel.: +46-31-772-6036.

E-mail address: claes.strannegard@chalmers.se

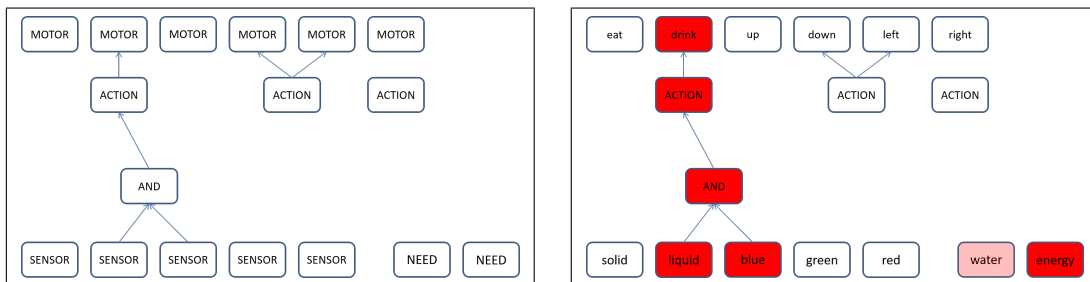


Fig. 1: The images to the left and right show the same network. To the left the node types are shown. To the right the names of the nodes are shown. The intensity of the red color reflects the level of activity of the corresponding node.

the stimulus”, the idea that the linguistic input received by children is insufficient to explain their detailed knowledge of their first language.

Machine learning has been used for learning natural language syntax and semantics from examples. Grammar induction models are used for learning grammars from examples [5, 6], e.g. by the use of syntactic equivalence classes [7]. Vector space models [8] are used for learning semantic relations via statistical analysis of large corpora. These models have been used to construct powerful search engines, recommendation systems, and summary generators. A key component of this method is the use of linear algebra and metrics, often the cosine function, that enable word contexts to be numerically compared. For instance, the words “cold” and “ice” may more similar to each other than “cold” and “grass”.

Reinforcement learning algorithms can be used by agents whose objective is to accumulate reward over time [9]. In standard reinforcement learning, the reward signal is one-dimensional; in multi-objective reinforcement learning, it is multidimensional [10]. Local Q-learning, where Q-values collected from multiple agents are merged into a global Q-value are considered in [11].

Stewart Wilson introduced the term *animat* for artificial animals [12] and outlined the *animat path to AI*, which seeks to create artificial intelligence by modeling animal intelligence [13, 14]. Of particular interest to this line of research are the homeostatic agents strive to keep the status of their needs in certain intervals [15]. Increases and decreases in the status of a need can be identified with reward and punishment, respectively. Animals may have access to information about the status of their own needs via signals emanating from interoceptors. Examples of animal needs and their associated interoceptors include energy (glucose receptors), water (osmoceptors), and oxygen (CO_2 receptors).

In this paper we explore whether a generic animat model without any dedicated mechanisms for language processing is capable of three basic forms of language learning. Section 2 gives a condensed overview of the animat model. Section 3 presents the results of the language learning experiments. Section 4, finally, draws some conclusions.

2. The animat model

In this section we give an overview of the animat model. More complete descriptions can be found in [16] and [17]. The code is available at [18].

A *network* is a graph, where each node has a name and type among NEED, SENSOR, AND, SEQ (sequence), MOTOR, AAND (simultaneous action), and ASEQ (action sequence). Time is modeled in discrete time steps called *ticks*. At each tick the animat receives input signals from the environment in the form of real values in $[0, 1]$ to the NEED nodes and boolean values to the SENSOR nodes. The latter values are propagated through the network. Action nodes can be activated either via reflexes or via decision-making. Networks are equivalent to sets of formulas of temporal logic [19]. In particular we use the binary modal operator SEQ that enables the construction of sequences. The formula $p \text{ SEQ } q$ is true at time t if p is true at $t - 1$ and q is true at t . Figure 1 gives an example of a dynamic graph.

An *animat* consists of:

- A network
- An activity pattern on the network
- A set of experience values. These are real values of three kinds: (i) local Q-values $Q_i(c, a)$, which encode the value from the perspective of need i of taking action a when a given node c is active; (ii) The transition matrix entries $Trans(c, a, c')$, which encode the probability that taking action a when node c is active leads to node c' being active at the next tick; (iii) The Conditional matrix entries $Cond(b, b')$, which encode the probability of two nodes b and b' being active at the same time.
- A *shape*, which is a subset of \mathbb{R}^3 .

The animat perceives the environment via its NEED nodes and the SENSOR node activity that propagates to other nodes (or concepts) of the network. This perception together with the animat's experience values determine which action is selected. The policy of the animat is ϵ -greedy in the sense that it explores (makes a random action) with probability ϵ and exploits (makes a best action) with probability $1 - \epsilon$.

The animat learns by updating its experience values using a version of reinforcement learning (multi-objective local Q-learning). It also learns by forming new nodes and by removing existing nodes (structural learning). New nodes are formed when a large enough prediction error is made in terms of expected reward. New nodes can also be formed at random moments by joining two nodes that were just active simultaneously or consecutively by means of AND or SEQ. Nodes that are superfluous from the perspective of reward prediction are eliminated in order to keep the network as small as possible (forgetting).

3. Language learning

We conducted three experiments relating to language with the animat model.

3.1. Learning to produce phonemes

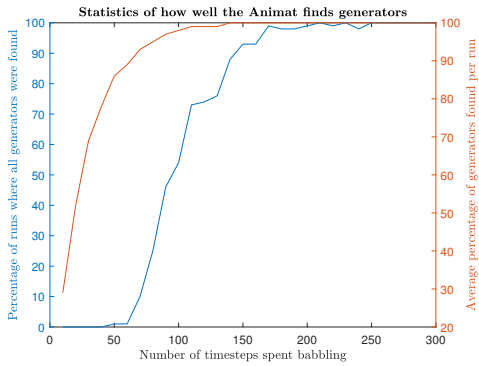
By letting the animat activate a random motor node in each time step, and then updating its transition matrix, it was capable of learning how to produce phonemes/letters that it had both sensors and motors for. This was evaluated by giving the animat sensors and motors for the letters of the english alphabet. The animat was then allowed to explore and update its transition matrix for a fixed number of time steps. Any action node a such that the value of the entry $Trans(True, a, b)$ in the transition matrix was equal to 1 was defined as a generator for the node b . The results are shown in figure 2a.

3.2. Learning to produce phonetic sequences


The learning of sequences was done through probabilistic learning. The animat would with some probability create a new sequence node, combining two nodes with probability proportional to how often in the last 100 time steps they had been active in sequence. It would then also create an action node producing that sequence by combining the generators for the two perception nodes combined. This was evaluated by exposing the animat to a set of 100 unique words, repeated randomly with equal probability on average 20 times each. The results of how many of the words that the animat had created nodes for after a certain amount of time is shown in figure 2b. As can be seen the animat quickly manages to learn the majority of the words. It also creates several nodes for parts of words, resulting in the animat creating more nodes than necessary. The animat was capable of producing all the sequences that it had learnt to recognise.

3.3. Learning to associate phonetic sequences to sensory experience


The animat was made to associate between perception nodes by using the rows of a time-extended conditional matrix (storing the probability of b and b' being active at almost the same time) as vector representations of the nodes. The cosine distance between the vectors was then used to determine the most similar node, in the same way as in vector space models. This allowed the animat to associate not only between words, but also from perception nodes representing phonetic sequences to perception nodes representing sensory experiences.




(a) Plot showing how often the Animat is capable of finding all generators in the different tests (blue). And showing the average number of generators found in each test (orange).



(b) Plot showing how the total number of perception nodes, and the number of perception nodes representing actual words, changes over time.



(a) Plot showing the results of associations between words. The blue plot shows the probability distribution of the Animats accuracy when the same score is given for all associations of the vector space model, along with its mean. The red plot shows the distribution and mean for when the associations were worth different amounts depending on how highly they were associated by the vector space model.



(b) Plot showing the results of associations between sensations and words. The solid lines show the probability distributions of the Animats accuracy, with the blue line representing the result of the associations from sensations to keywords, and the red line representing result of the associations from keywords to sensations. The corresponding results for the vector space model is shown with dashed lines.

The animat's ability to associate between word was tested on a simple input text and compared to a simple vector space model. When given the same input as the simple vector space model, the animat managed to return the same top ten associations as the vector space model with an average accuracy of over 80 percent as seen in figure 2a (blue line). When weighting the value of the associations made higher for the top associations of the vector space model (and decreasing the value linearly), the animat achieved an average accuracy of over 90 percent (as seen in the red plot in figure 2a). When the animat made use of a simple chunking method, it obtained an accuracy of over 99 percent.

The Animats ability to associate between words and other input was also evaluated, and compared to the vector space model. The associations of both models were compared with a dictionary between words and other sensations, that was used when ceating the training data used as input. The associations were tested both from words to other senses and from senses to words. The result is shown in figure 2b, showing the animat's accuracy as solid lines and the vector space model's accuracy as dashed lines.

4. Conclusion

We conducted three experiments with a general animat model without any dedicated mechanisms for language learning. The results obtained indicate that the animat model is capable of three kinds of basic language learning.

Acknowledgement

This research was supported by the Torsten Söderberg Foundation Ö110/17.

References

- [1] C. Darwin, *The descent of man and selection in relation to sex*, vol. 1. Murray, 1888.
- [2] G. Roth and U. Dicke, *Evolution of the brain and intelligence in primates*, in *Progress in brain research*, vol. 195, pp. 413430, Elsevier, 2012.
- [3] S. Herculano-Houzel, *The human brain in numbers: a linearly scaled-up primate brain*, *Frontiers in human neuroscience*, vol. 3, p. 31, 2009.
- [4] N. Chomsky, *Rules and representations*, *Behavioral and brain sciences*, vol. 3, no. 1, pp. 115, 1980.
- [5] E. Brill, *Automatic grammar induction and parsing free text: A transformation-based approach*, in *Proceedings of the 31st annual meeting on Association for Computational Linguistics*, pp. 259265, Association for Computational Linguistics, 1993.
- [6] A. DULIZIA, F. Ferri, and P. Grifoni, *A survey of grammatical inference methods for natural language learning*, *Artificial Intelligence Review*, vol. 36, pp. 127, Jun 2011.
- [7] A. Clark, C. Fox, and S. Lappin, *The handbook of computational linguistics and natural language processing*. John Wiley & Sons, 2013.
- [8] S. Clark, *Vector Space Models of Lexical Meaning*, pp. 493522. John Wiley & Sons, Ltd, 2015.
- [9] R. S. Sutton and A. G. Barto, *Reinforcement learning: An introduction*. MIT press, 1998.
- [10] D. M. Roijers, P. Vamplew, S. Whiteson, R. Dazeley, et al., *A survey of multi-objective sequential decision-making.*, *J. Artif. Intell. Res.(JAIR)*, vol. 48, pp. 67113, 2013.
- [11] S. J. Russell and A. Zimdars, *Q-decomposition for reinforcement learning agents*, in *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pp. 656663, 2003.
- [12] S. W. Wilson, *Knowledge growth in an artificial animal*, in *Adaptive and Learning Systems*, pp. 255264, Springer, 1986.
- [13] S. W. Wilson, *The animat path to AI*, in *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior (J. A. Meyer and S. W. Wilson, eds.)*, 1991.
- [14] E. Tuci, A. Giagkos, M. Wilson, and J. Hallam, eds., *From Animals to Animats. 1st International Conference on the Simulation of Adaptive Behavior*, Springer, 2016.
- [15] G. Konidaris and A. Barto, *An adaptive robot motivational system*, in *SAB*, pp. 346356, Springer, 2006.
- [16] C. Strannegård, N. Svängård, D. Lindström, J. Bach, and B. Steunebrink, *The animat path to artificial general intelligence*, in *Workshop on Architectures for Generality and Autonomy, IJCAI-17*, 2017.
- [17] L. Johannesson and M. Nilsson, *Simple language learning in artificial general intelligence*, Masters thesis, Chalmers University of Technology, 2018.
- [18] L. Johannesson and M. Nilsson, *Animat code*. <https://github.com/Pessimus/MasterThesisAnimatAGI>, 2018
- [19] D. M. Gabbay, I. Hodkinson, and M. Reynolds, *Temporal logic (vol. 1): mathematical foundations and computational aspects*. Oxford University Press, Inc., 1994.