Game-based Language Tutoring: A Case Study for Colour

Katrien Beuls
Vrije Universiteit Brussel
Brussels, Belgium
katrien@arti.vub.ac.be

Joris Bleys
Vrije Universiteit Brussel
Brussels, Belgium
jorisb@arti.vub.ac.be

Abstract

Since computational agents tend to be used more and more outside research labs, a serious effort should be put into their ability to learn new skills to adapt to changing environments. These skills also include language processing and language learning. How do we teach natural language to agents so that their linguistic knowledge is robust enough to deal with new situations? We present a case study in which a software agent needs to acquire the English colour prototypes and their corresponding terms, purely based on interactions with a human tutor. In order to provide such a system, we have extended the existing Babel framework [Steels and Loetzsch, 2009], which has been designed to model grounded communication and cognitive learning processes in a multi-agent population, so that a human can take up the role of a tutor and provide input for the agent learner through a GUI. Hereby, an agent's language skills are constantly being shaped due to the variation that is present in the situations he encounters and the corresponding linguistic input he receives from the tutor.

1 Introduction

It is becoming increasingly important to equip natural language processing agents (robotic/software) with deep language processing skills that go beyond statistical predictions based on previous input in order to communicate with people in a humanlike way. Scenes can be conceptualised in different ways depending on the perspective you take or the language you speak. Take for instance a tool such as the screwdriver. In English, it is conceptualised as an instrument to drive screws into an object, in French it only turns the screws (*tournevis*) whereas in German it pulls screws out of an object (*Schraubenzieher*).

We have developed a first prototype of a so-called *tutoring game*, in which a human tutor can teach a particular language system, that is a functional subpart of a full language such as the way to speak about space or time, to a software agent. A tutoring game is a particular kind of language game [Steels, 1995] in which a human tutor replaces one of the two agents that interact in the game. The tutor plays a crucial role in the game, since he can constrain the situation, scaffold complexity, and provide pragmatic feedback for the learner. The importance of the tutor already became visible in the world learning experiment carried out by Steels and Kaplan [2000] where a robotic agent (Sony's AIBO) had to acquire words for a predefined number of objects.

This paper shows a demonstration of the *Color Tutoring Game* (CTG), that allows a human tutor to teach colour terms to a software agent by selecting colour chips and naming them accordingly. Since we are dealing with colour language, the final number of words is open-ended and depends on the tutor's colour classification. The fact that we are dealing with a continuous input space makes it an interesting case for tutoring since it is the tutor's task to help the learner to classify and name instances drawn from this input space. The system we are presenting is fully operational and has been tested on one German speaking tutor. We are planning an experiment with more participants to get a better evaluation of the system.

2 Language Games

Previous research has shown how a population of agents can reach a consensus on how they name or describe particular situations they encounter in terms of spatial language [Spranger *et al.*, 2010], verbal argument structure [van Trijp, 2011], quantification [Pauw and Hilferty, 2011], colour terms [Bleys *et al.*, 2009], etc. All these experiments are embedded in the Babel framework [Steels and Loetzsch, 2009], an open-ended system through which agents learn and shape their knowledge through repeated peer-to-peer interactions, or language games [Steels, 2011].

A language game follows a particular script in which two players try to reach a certain communicative goal they share. Usually a language game involves a sequence of interactions, and successful communication takes place when the communicative goal has been achieved. This can normally be observed or inferred by speaker and hearer.

The set of cognitive processes underlying a language game are depicted in Figure 1. This Figure depicts the so-called *semiotic cycle* and represents the flow of information within a single interaction. The left side of the Figure represents the speaker, the right side shows the internal script of the hearer. The round boxes in the middle of the Figure are shared among speaker and hearer: the physical world they are grounded in and observe and the utterance that is produced by the speaker

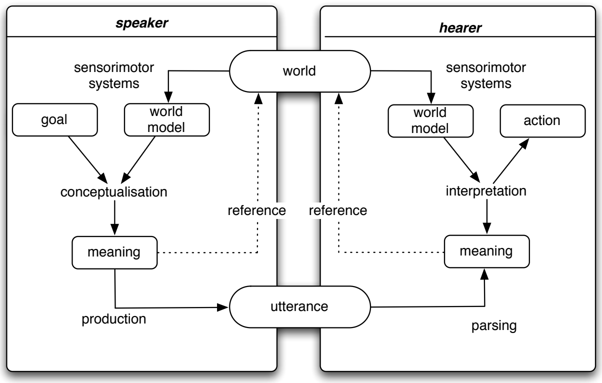

Figure 1: A full semiotic cycle: the basic architecture of every language game.

(production) and subsequently received by the hearer (parsing). There are four major stages in the semiotic cycle where the speaker and hearer have to run through [Steels, 2011]:

1. Grounding: the connection between internal factual memory and the current situation model.
2. Conceptualisation/Interpretation: the link between the situation model and a higher order semantic categorisation over the current situation.
3. Production/Parsing (FCG): the mapping between semantic structure (representation of categories and their relations) and surface form (syntax).
4. Speech Articulation/Speech Recognition: the rendering or de-rendering of a sentence.

3 Tutoring Games

In a tutoring game, the human and the agent take up both roles in the semiotic cycle. It is however, the human who can mediate the flow of the game by selecting the topic and giving feedback to the agent at the end of the game to signal success or failure. A tutoring game typically involves a world W consisting of objects. Each object is characterised by a point in an n-dimensional feature space. For example, the colour of an object is a point in the three-dimensional colour feature space with the dimensions red-green, yellow-blue, and lightness. In more complex games, the world is filled with scenes that contain events, event participants, properties and relations. The human is in the initial stages always the speaker, while the agent is the hearer. A context C is established which contains a subset of the objects in the world W. Then the following interaction takes place (adapted from the agent-agent game script [Steels, 2011]):

1. The human selects one object out of the context, further called the topic T.
2. The human finds the distinctive category for the object and names this category.
3. The agent looks up which object is associated with this category in his memory and examines the context to find out whether there is an object which has this distinctive characteristic.
4. The agent then signals to the human which object was intended according to him, for example by pointing.
5. The human checks whether the agent selected the same object as the one he had originally chosen. (a) If they are the same, the game is a success, and the human signals this outcome to the agent. (b) If they are different, the game is a failure. The human signals this outcome and then points to the topic he had originally chosen.

4 The Colour Tutoring Game

This section demonstrates the workings of the Colour Tutoring Game (CTG) with a human tutor. The system can also be used with reversed roles, where the human is learning the colour terms that have been previously taught to the software agent by another human user. This scenario falls outside the scope of this workshop and will be explored in other papers.

The script of a single interaction between a human tutor and a software agent in the CTG is included below. Figures 2 and 3 show how one can interact with the agent through a graphical user interface.

1. The human tutor selects one chip out of a context with three chips (see Figure 2). When there is no chip in the context the human can name, he can decide to skip to context and generate a new one.
2. When the agent does not have a colour prototype whose values are close to the chip, the human names the chip (Figure 3a). Else the agent tries to come up with a name (Figure 3b).
3. When the agent was able to come up with a name, the human is asked to give feedback whether the agent was right or wrong (Figure 3b). When the agent's answer was wrong, the human corrects him (Figure 3c).
4. The agent stores a newly heard name in his memory together with the values of the colour chip in his three dimensional colour space. If he had already heard the name before, he aligns the pre-existing category to the current chip by means of the alignment operator (see Section 4.1).

Two example interactions of the Color Tutoring Game proceed as follows:

```
1. Human: What colour is this chip?
2. Agent: I don't know...
3. Human: It is blue.
4. Agent: Blue?
5. Human: Yes.

1. Human: What colour is this chip?
2. Agent: Pink.
3. Human: No, it is blue.
4. Agent: The chip is blue.
```

It is important to note that the game prototype we are presenting in this paper does not include the operationalisation of stages 1 and 4 in the semiotic cycle. The learner agent

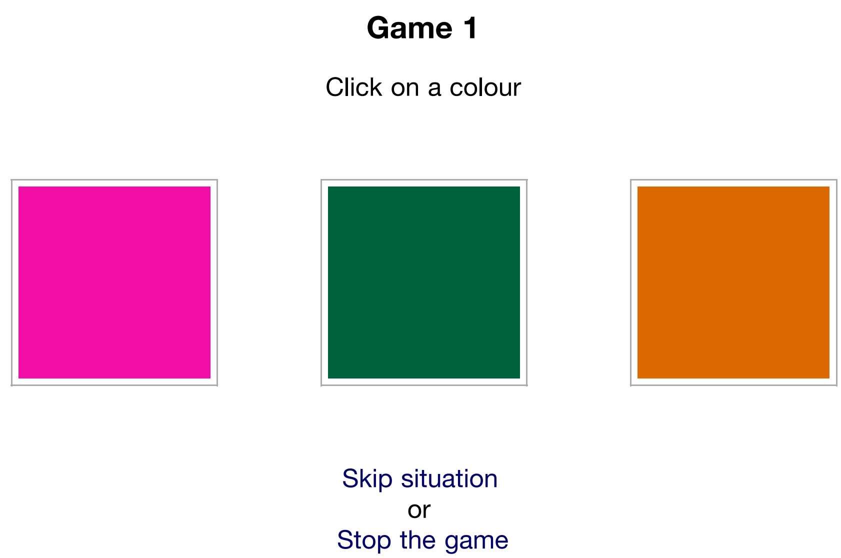

Figure 2: Screenshot of Step 1 in the interaction script: the human tutor is asked to select one colour chip.

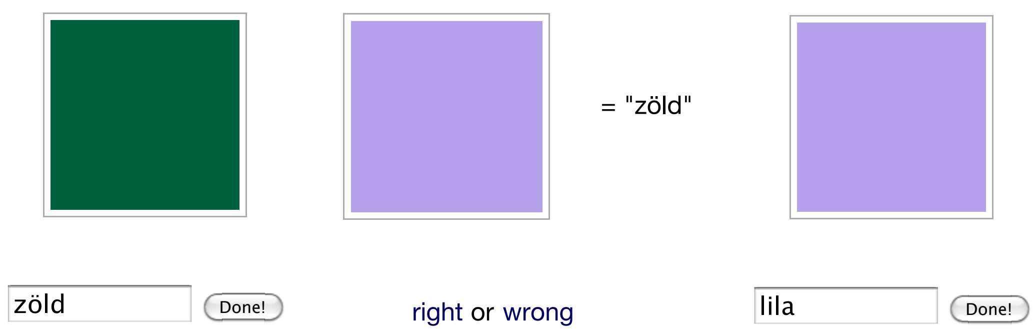

(a) The human names (b) The agent names (c) The human corrects

Figure 3: Screenshots of Steps 2 and 3: The human interacts with the agent either by naming the colour chip he selected (a) or by assessing the name the agent used to name the chip (b) and possibly correcting the chosen name (c).

is thus only concerned with conceptualisation and production as a speaker and parsing and interpretation as a hearer. The grounding of a colour language game has been shown in previous research [Bleys *et al.*, 2009] and a speech system could be added in a later stage.

4.1 Behind the Scenes

The colour chips that are used to generate the contexts correspond to the Munsell chips, which are also used by anthropologists and psychologists in their colour surveys [Kay *et al.*, 2010; Sturges and Whitfield, 1995]. The colours of these chips are converted to RGB values so they can be reproduced on graphical interfaces using standard formulae [Fairchild, 1998]. Great care has been taken to incorporate device specific information of the display to make this reproduction as accurately as possible. More details on the exact techniques that have been used can be found elsewhere [Bleys, 2010].

The colour categories of the agent are represented by a single point. This point represents is the prototype of the colour category and corresponds to the best sample of this category, as commonly assumed in psychological studies [Rosch, 1973]. The conceptual space that is used to represent these prototypes is the three dimensional CIE $L^*a^*b^*$ colour space: L^* represents lightness, a^* the red-green and b^* the blue-yellow opponent cannel. CIE $L^*a^*b^*$ has been especially designed to reflect human colour perception in which colour difference can be calculated by taking the Euclidean distance between two points.

Whenever an agent needs to name a colour chip, it maps the colour of that chip in the CIE *L*a*b** colour space and assigns it to the category of which the prototype is the most similar to that color (one nearest-neighbour classification). Next, it looks up the unique name that is assigned to this category and utters it. New categories are added to the repertoire of the agent whenever it encounters a name it had never encountered before, as also observed in psychological experiments [Xu, 2002].

At the end of a successful interaction, the agent updates its repertoire of colour categories to better reflect the colour system of the human tutor. This process is called alignment and is implemented by shifting the existing prototype of the category that was evoked through a colour term in the direction of the colour values of the colour chip that was selected by the human. The rate by which this shift happens is controlled by *colour category alignment rate* (r_a) which linearly specifies the new location of the prototype (c_{n+1}) on the line segment between the old location of the prototype (c_n) and the topic (t). The exact formula is shown in Equation 1 (adapted from [Bleys, 2011]). If the alignment rate is 0, the prototype does not shift at all, whereas at a rate of1, the new location would be the topic of the last interaction. The alignment rate is for the current experiment fixed to 0.05 .

$$c_{n+1} = (1 - r_a)c_n + r_a t \quad (1)$$

4.2 Preliminary Results

Figure 4 plots the success scores of the agent in naming the colour samples. At the end of the game, the human is asked whether the agent was right or wrong in naming the colour chip that was selected. When the agent was right, the game succeeds, which translates into score of 1. In case of failure, the agent receives a score of 0. The success scores rise rapidly within the first 20 games and the agent always reaches over 80% success after half of the games have been played. Due to the large variations in the data set, the agent's success score could never reach 100%. Some chips are too much in the middle between two colour categories

At the end of each series of games, the resulting colour data set can be saved in the system and used for further processing. Figure 5 shows the positions of the colour prototypes in the two-dimensional *(L*)a*b** space after an agent has played 100 games. This agent has memorised nine colour prototypes, which would correspond to the English basic colour terms (from left to right): green, blue, yellow, brown, black, pink, orange, purple and red. Two basic colour terms are uncovered: white and grey. As mentioned above, this is due to the high degree of saturation of the colour samples.

4.3 Discussion

The results presented in this paper are only a demonstration and were generated by a single German speaking subject teaching English colour terms in three series of 100 interactions. Besides the obvious extension to include more human subjects, it would be interesting to make a more thorough comparison to the techniques commonly used in psychological experiments. The prototypes of the colour categories of the subject could be determined by indicating them directly in set of all Munsell chips which in turn could be compared

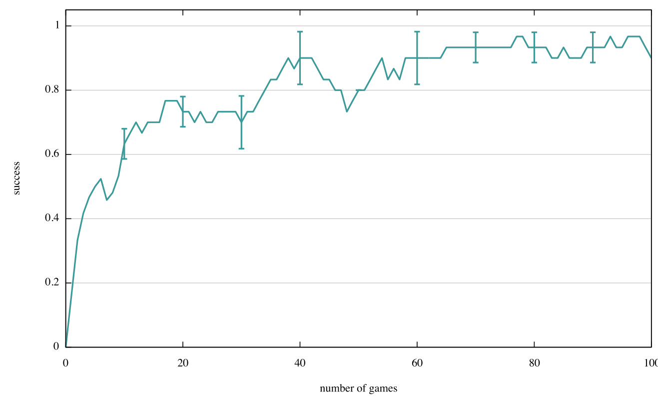

Figure 4: The average success scores of the learning agent over three game series of 100 games each. The x-axis plots the number of games played, the y-axis plots the success score (0 or 1). The success score is one when the agent could correctly name a colour sample.

to the prototypes of the tutored agent to see how closely they correspond.

The preliminary results of the proposed tutoring system are however very promising. A previous study has shown that when an agent is pre-programmed with an English colour system [Sturges and Whitfield, 1995], it could name 83% of the chips that are consistently named by English speakers correctly [Belpaeme and Bleys, 2009]. Although the measure used in this paper is not exactly identical, it measures a similar quantity: the number of chips which are named correctly according to the human tutor. As the human tutor can choose to ignore difficult contexts in which there is no chip that can be clearly named, it also only measures the chips that are likely to be consistently named by the human tutor. The tutored agent in this study achieves a success of around 90% which is higher than the baseline established in the previous study. The tutoring system that is presented in this paper achieved a very promising result.

The tutoring system presented in this paper only covered the domain of colour, but could easily be extended to other domains based on results of other language-game research, such as quantifiers [Pauw and Hilferty, 2011], case systems [van Trijp, 2011] or spatial language [Spranger *et al.*, 2010]. For each of this domain, different contexts and corresponding learning operators have been presented which could be converted to tutoring systems similar to the one presented in this paper for colour.

5 Conclusions and Future Work

This paper has shown the potential usability of the Babel framework as a tool for interactively tutoring software agents. By replacing one agent in a language game by a human (speaker), we were able to focus on the parsing and interpretation processes (speech has been excluded) and the formation of prototypes in the memory of the agent. The extension of the language game paradigm for human - agent interaction

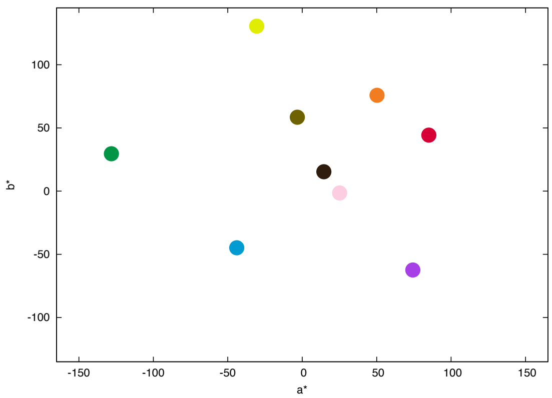

Figure 5: The colour prototypes of one learner agent after 100 language games projected on the hue plane of the CIE *L*a*b** colour space.

is especially promising due to its inherent focus on communication. Moreover, the dual processing architecture opens up possibilities to diagnose problems that occur in processing and generate possible repair strategies whose viability can be tested in future interactions.

Since this first tutoring demo only concerned lexicon formation, the next step to take is to extend the framework to deal with more complex scenes that induce the use of grammatical structures. Evolutionary language game experiments on grammar [van Trijp, 2011; Pauw and Hilferty, 2011] have already shown the strength of the Babel framework in this field. The extension towards human - agent interaction is thus the next logical step to take.

Acknowledgments

This research was conducted at the Vrije Universiteit Brussel, financed by a strategic basic research grant (IWT-489) from the agency for Innovation by Science and Technology (IWT). Additional funding came from the European research project ALEAR (FP7, ICT-214856). We especially want to thank all members of the team who helped achieving this: Remi van Trijp, Vanessa Micelli, Michael Spranger, Pieter Wellens, Joachim de Beule, Simon Pauw and Luc Steels for his continuous input of extremely inspiring and interesting ideas.

References

[Belpaeme and Bleys, 2009] Tony Belpaeme and Joris Bleys. The Impact of Statistical Distributions of Colours on Colour Category Acquisition. *Journal of Cognitive Science*, 10(1):1–20, 2009.

[Bleys *et al.*, 2009] J. Bleys, M. Loetzsch, M. Spranger, and L. Steels. The grounded color naming game. In *Proceedings of the 18th IEEE International Symposium on Robot and Human Interactive Communication (Ro-man 2009)*, 2009.

[Bleys, 2010] Joris Bleys. *Language Strategies for the Domain of Colour*. PhD thesis, Vrije Universiteit Brussel, Brussels, Belgium, 2010.

[Bleys, 2011] Joris Bleys. The evolution of language strategies for color. In Luc Steels, editor, *Experiments in Language Evolution*. John Benjamins, Amsterdam, 2011.

[Fairchild, 1998] Mark D. Fairchild. *Color Appearance Models*. Addison Wesley, 1998.

[Kay *et al.*, 2010] Paul Kay, Brent Berlin, Luisa Maffi, William R. Merrifield, and Richard S. Cook. *World Color Survey*. CSLI Lecture Notes. University of Chicago Press, 2010.

[Pauw and Hilferty, 2011] Simon Pauw and Joseph Hilferty. The emergence of quantification. In Luc Steels, editor, *Experiments in Language Evolution*. Springer Verlag, 2011.

[Rosch, 1973] Eleanor H. Rosch. Natural categories. *Cognitive Psychology*, 4:328–350, 1973.

[Spranger *et al.*, 2010] Michael Spranger, Simon Pauw, and Martin Loetzsch. Open-ended semantics co-evolving with spatial language. In *Proceedings of EVOLANG 8*, 2010.

[Steels and Kaplan, 2000] L. Steels and F. Kaplan. Aibo's first words: The social learning of language and meaning. *Evolution of Communication*, 4(1):3–32, 2000.

[Steels and Loetzsch, 2009] L. Steels and M. Loetzsch. Babel: A tool for running experiments on the evolution of language. In S. Nolfi and M. Mirolli, editors, *Evolution of Communication and Language in Embodied Agents*, pages 307–313. Springer, Berlin, 2009.

[Steels, 1995] L. Steels. A self-organizing spatial vocabulary. *Artificial Life Journal*, 2(3):319–332, 1995.

[Steels, 2011] Luc Steels. Exploring cultural language evolution with language games. In *Experiments in Language Evolution*. Springer, 2011.

[Sturges and Whitfield, 1995] Julia Sturges and T. W. Allan Whitfield. Locating basic colours in the munsell space. *Color Research & Application*, 20(6):364–376, 1995.

[van Trijp, 2011] Remi van Trijp. The emergence of a case grammar. In Luc Steels, editor, *Experiments in Language Evolution*. John Benjamins, Amsterdam, 2011.

[Xu, 2002] Fei Xu. The role of language in acquiring object concepts in infancy. *Cognition*, 85:223–250, 2002.