HAL
open science

The Role of Intrinsic Motivation in Artificial Language Emergence: a Case Study on Colour

Miquel Cornudella Gaya, Thierry Poibeau, Remi van Trijp

To cite this version:

Miquel Cornudella Gaya, Thierry Poibeau, Remi van Trijp. The Role of Intrinsic Motivation in Artificial Language Emergence: a Case Study on Colour. 26th International Conference on Computational Linguistics (COLING 2016), Dec 2016, Osaka, Japan. pp.1646 - 1656. hal-01422478

HAL Id: hal-01422478

https://hal.science/hal-01422478

Submitted on 26 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The Role of Intrinsic Motivation in Artificial Language Emergence: a Case Study on Colour

Miquel Cornudella[1,2] **Thierry Poibeau**[2] **Remi van Trijp**[1]

[1]**Sony Computer Science Laboratory Paris**
6 rue Amyot, 75005
Paris, France
{cornudella,remi}@csl.sony.fr

[2]**Laboratoire LATTICE-CNRS**
ENS & U. Paris 3 & CNRS
PSL and USPC
1 rue Maurice Arnoux, 92120
Montrouge, France
thierry.poibeau@ens.fr

Abstract

Human languages have multiple strategies that allow us to discriminate objects in a vast variety of contexts. Colours have been extensively studied from this point of view. In particular, previous research in artificial language evolution has shown how artificial languages may emerge based on specific strategies to distinguish colours. Still, it has not been shown how several strategies of diverse complexity can be autonomously managed by artificial agents . We propose an intrinsic motivation system that allows agents in a population to create a shared artificial language and progressively increase its expressive power. Our results show that with such a system agents successfully regulate their language development, which indicates a relation between population size and consistency in the emergent communicative systems.

1 Introduction

Over the past two decades, language evolution studies have attracted the attention of researchers working on domains such as biology, anthropology, artificial life or linguistics. This multitude of perspectives provides a rich variety of techniques on how to address this issue, including including agent-based modelling, which consists in studying the emergence and evolution of *artificial languages*, i.e. human-like communicative systems, in a population of artificial agents through recurrent peer-to-peer interactions (Smith et al., 2003; Steels, 2012). Results using this approach have shed light on the emergence of spatial relations (Spranger, 2013), case systems (van Trijp, 2013), colour categories (Bleys, 2010) or syntax (Garcia-Casademont and Steels, 2016).

In most of these experiments the control of the complexity relies on the experimenter, who carefully selects the stages of the experiment, constraining the language development. Insights from different models dealing with complexity come from research in AI and robotics, where a number of studies where a number of studies tried to specify how agents regulate the complexity of their actions in an autonomous way. Several models have been proposed, including error reduction (Andry et al., 2001), prediction (Marshall et al., 2004), interest (Merrick and Maher, 2009) or curiosity (Oudeyer et al., 2007).

These systems have been deeply inspired by psychological studies on the role of *motivation* (see Graham (1996) for an overview). According to Ryan and Deci (2000), motivation can be defined as "to be *moved* to do something". Psychologists further distinguish two types of motivation, depending on the reasons to perform an action: *extrinsic*, when the interest relies on the outcome of the action, or *intrinsic*, when the action itself results inherently enjoyable. Both types of motivation can incite to take part in the same activity. For example, a tennis player can work on improving her slice shots because she wants to win her next tournament or because she enjoys improving her technique.

This paper presents a simulation experiment to study how a population provided with the *Autotelic principle* (2004), a computational motivation system inspired on the *Flow theory* (1990), is able to self-organize and extend the expressive power of a shared communication system for the continuous domain of colour. Agents play a language game[1] and they use this motivation system to decide on the complexity

This work is licensed under a Creative Commons Attribution 4.0 International Licence. Licence details: http://creativecommons.org/licenses/by/4.0/

[1]Interested readers in the methodology used in this experiment to study the emergence of artificial language are referred to Steels (2012).

Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers, pages 1646–1656, Osaka, Japan, December 11-17 2016.

of both the context of the interaction and the utterances they formulate and comprehend.

In the next section of the paper, we describe Flow Theory and an operational computational version, the Autotelic Principle, that allow agents to autonomously regulate their development. Section 3 briefly reviews artificial language evolution research on the domain of colour, the case study in which the motivational system is tested. Section 4 presents the experiment design in detail: how agents interact, the different communicative tasks, how the context of interactions is chosen and the different operational mechanisms agents use to emerge and align a shared communicative system. Finally, sections 5 and 6 presents the experimental results and conclusions.

2 Flow theory and architecture

Csíkszentmihályi wanted to understand what moves people to be absorbed in complex activities that do not provide an external reward, such as rock climbing, painting or sculpting. He found that the reason was that participants found these activities inherently enjoyable. He called these activities *autotelic*, as the motivational driving force (*telos*) comes from the individual itself (*auto*).

Based on these observations he developed the *Flow theory* (1990). Autotelic activities can be explained based on the relation of two dimensions (Figure 1a): *challenge*, a certain task to be done, and *skill*, the abilities that a person has to tackle that task. This relation accounts for the range of different mental states that people experience when they are involved in an autotelic activity: *boredom*, when the skills are too high for the current challenge, *anxiety*, when the challenge is too difficult given the current skills, and *flow*, when there is a balance between both. He identified the latest as the optimal state of experience. This state provides the best scenario for further enlarging their skills. As a consequence, the state of flow is in in continuous motion, as skills evolve over time. Participants seek to stay in flow state, therefore becoming self-motivated.

Inspired by the work of Csíkszentmihályi, Steels (2004) proposed the *Autotelic principle*, an operational version of the Flow theory to provide artificial agents with a system to self-regulate their development. As in the Flow theory, the balance between challenge and skills acts as the motivational driving force in agents. Agents use this relation to identify their internal state (*boredom*, *anxiety* or *flow*) and consequently react to it increasing or decreasing their challenges.

2.1 Challenge management

The principle establishes two different phases, *operational* and *shake-up*. The first one corresponds to the state of flow: agents explore a particular challenge and try to develop the abilities required to cope with it. The latest is reached when agents are either in a state of anxiety or boredom. It acts as a trigger to adjust the challenge to be addressed in order to look for a more balanced challenge-skills relation. In case of boredom a more demanding challenge should be attempted. In contrast, a more accessible challenge should be tackled in an anxiety state.

Challenges are characterised as a set of parameters: given a multi-dimensional parameter space P, a challenge p_i is defined as a vector $< p_{i,1}, p_{i,2}, ..., p_{i,n} >$, where $p_{i,j}$ corresponds to the value of the parameter j in the challenge i. Agents are able to generate more complex or manageable challenges by changing the specific configuration of a challenge. The space of possible challenges depends on the number of parameters of the set and the different values each parameter can have.

2.2 Skill evaluation

Agents indirectly measure their skills based on their performance. Each challenge is monitored with a value in the range [0,1] where 1 represents optimal performance for that task. This value is used to compute the *confidence* an agent has in its skill level to accomplish a particular challenge. A confidence value of 1 is interpreted as boredom as it indicates that the agent has acquired the skills to handle the current task. On the other hand, a steady value of 0 is interpreted as anxiety as it shows that the agent is not succeeding in expanding its abilities to cope with the actual challenge.

After each interaction, the confidence on the current challenge is updated taking into account the individual competences of the agent and the outcome of the interaction. When the interaction is a

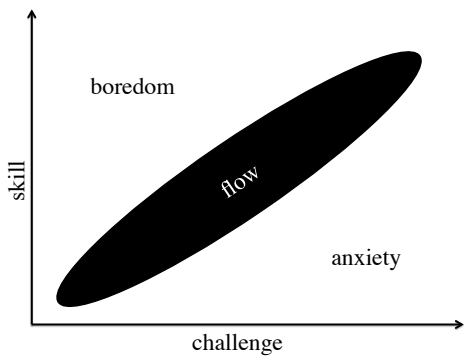

(a) Flow theory

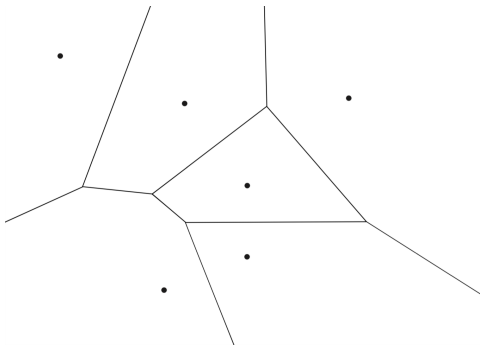

(b) Visual representation of the partition of the colour space into basic colour categories

Figure 1: According to Csíkszentmihályi (Figure 1a), individuals enter a state of *flow* when they correctly balance their skills against the selected challenges when performing particular activities. A mismatch of this balance may lead to anxiety (i.e. the challenge is too big) or boredom (i.e. the challenge is too easy). Figure 1b illustrates the division of the colour space into basic categories. Image extracted from Bleys (2010).

success, the confidence value is updated as follows: $conf_i = conf_{i-1} + \delta increase$, where $conf_i$ and $conf_{i-1}$ are the current and previous confidence values and $\delta increase$ is set to 0.005. On the contrary, when the outcome of an interaction is a failure the confidence value is updated in this way: $conf_i = conf_{i-1} - \delta decrease + ind_{comp}$, where $\delta decrease$ is set to 0.02 and ind_{comp} is the value in the range [0,0.015] that corresponds to the evaluation of the individual competences of the agent.

This motivation system has been used before in experiments of language emergence in discrete domains (Steels and Wellens, 2007; Cornudella et al., 2015). The work presented here differs from previous experiments in that it tests the autotelic principle in a continuous domain. In this experiment agents need to self-organize a communicative system but also agree on the meanings associated with their lexicons: colours are no longer discrete values but rather points in a three dimensional feature space and membership prototypes are values in the range [0,1].

3 A case study on colour

Research on the domain of colour has been of great interest to a lot of researchers, due to the differences observed in how colours are described in human languages (Berlin and Kay, 1969). It is commonly accepted that a *colour space*, the space of colours that can be perceived, is organised in different *colour categories*, subsets of this space (Figure 1b). *Colour prototypes* are the best representation of a particular colour category in a colour space (Rosch, 1973). Formally, a colour space is composed of a set of colour prototypes $\{c_1, c_2, ..., c_n\}$. Given the colour prototype c_k, its associated cell R_k, which determines the associated colour category, contains every point whose distance to c_k is shorter or equal to the distance to any other prototype c_i.

Although the research in this domain is extensive, most studies have focused on the use of single terms to describe colours. Experiments by Simpson and Tarrant (1991) and Lin et al. (2001) showed that only 15% of colour samples were described using a single colour term when human subjects were asked to describe colour samples without any restriction. These results provide evidence to the fact that usually people prefer to express more information about the colours they are describing instead of only employing single terms.

Colour has been of particular interest to researchers working on artificial language evolution. The majority of models have focused on the emergence of single colour terms (Steels et al., 2005; Belpaeme and Bleys, 2007; Baronchelli et al., 2010; Baronchelli et al., 2015), but some attempts to model more complex descriptions also exist. In this respect, the most advanced contribution is the work of Bleys (2009; 2009; 2012). In his doctoral thesis (2010) different *language strategies*, a particular method to express one area of meaning, are explained and studied. These language strategies are then tested on

artificial language evolution experiments, showing how these models can emerge and be learned by a population of artificial agents.

4 Experiment

In this section the design of the experiment is explained: the particular language game agents play, the different communicative challenges of the experiment, how the contexts of interactions are determined and the operational mechanisms agents use to develop and align their language. The experiment is implemented in Babel2[2], a multi-agent experiment framework (Loetzsch et al., 2008).

4.1 Language Game

The experiment consists in recurrent communicative interactions in a population of artificial agents equipped with the autotelic principle situated in a particular context. In every interaction a randomly selected pair of agents is picked from the population. One of them assumes the role of *speaker* and the other the role of *hearer*. The goal of the interacting agents is to communicate about one colour sample from the context.

The specific language game that agents play is called multi-word guessing game. The speaker selects the context of the interaction, based on the challenge it is currently addressing, and randomly picks a colour sample as topic[3]. When the speaker is able to discriminatingly conceptualise the topic into a meaning predicate it uses its language component to formulate an utterance which is transmitted as text to the hearer. The hearer tries to comprehend the utterance and constructs hypotheses about the topic. If the hearer has only one hypothesis, it points to the interpreted topic.

If the hypothesis corresponds to the topic, the speaker gives positive feedback and the interaction ends. On the other hand, if the hypothesis does not correspond to it, the speaker gives negative feedback to the hearer and points to the intended topic. When the hearer has no or multiple hypotheses it signs to the speaker that it could not identify the topic. The speaker then gives feedback by pointing to the intended topic. The interaction is a success only when the hearer has one hypothesis about the topic that corresponds with the topic selected by the speaker. Otherwise, the result of the interaction is a failure.

4.2 Challenges

Agents can use different language strategies to communicate about colour samples. These strategies are identified as the three communicative challenges of the experiment and were previously analysed in Bleys (2010): *basic colour*, *graded membership* and *graded category combination*[4]. A parameter *level* is associated to each challenge, according to its complexity. Agents use this parameter to move between challenges, depending on their internal state. Agents are able to perform three operations on the colour space: add a colour prototype, compute the distance between a colour sample and its closer colour prototype (Figure 2a) and transform the colour space towards a colour prototype (Figure 2b).

In the basic colour strategy a single term is used to describe a colour sample. Agents use this term to refer to the closest colour prototype. An example for English would be to use a term as “green” to describe a colour sample. In the experiment, agents are initialised with an empty vocabulary and colour space. This means that they need to converge both on a classification of the colour space into colour prototypes and on the terms associated to each colour prototype.

The graded membership strategy characterises a colour sample by expressing both the closest colour prototype and the distance between the colour sample to it. This strategy is observed, for instance, in English, where it is possible to describe a colour sample by combining a basic colour term with adverbs such as “very” or the postfix “-ish”, as in “very blue” or “greenish”. When addressing this challenge, agents also need to agree on both the membership prototypes and the terms associated with them.

Lastly, the graded category combination strategy describes a colour sample by referring to two colour and one membership prototypes. Agents first identify the closest colour prototype to the colour sample

[2]Babel2 is available as open-source software at `www.emergent-languages.org`.

[3]Interested readers in the impact of active selection of the topic are pointed to Schueller and Oudeyer (2015).

[4]The language strategies used in this paper have been replicated from Bleys (2010), who granted us access to the original implementation.

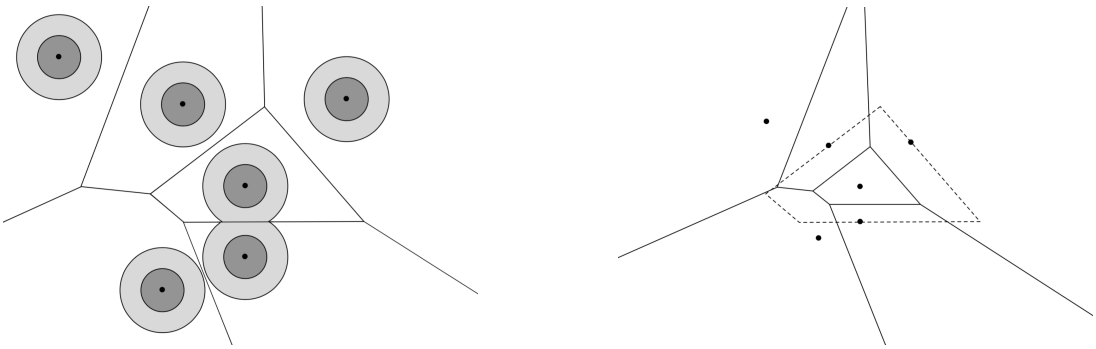

(a) Graded membership operation

(b) Category combination operation

Figure 2: Visual representation of available operations. In a graded membership operation (Figure 2a) the distance between the colour sample and the closer colour prototype is computed. In a category combination operation (Figure 2b) the colour space is transformed towards the main colour prototype of the colour sample in order to perform a second classification. Images extracted from Bleys (2010).

and transform the colour space towards that prototype. Agents classify again the colour sample on the transformed colour space, obtaining a second colour prototype[5]. Finally, they express how close the colour sample is to the identified colour prototype in the transformed colour space using a graded membership term. "Very dark green" or "blueish purple" are examples of this strategy in English.

The complexity of a communicative task is determined by its number of *cognitive operations*: algorithms that encode a particular cognitive function used in conceptualisation and interpretation. This number differs among the different challenges and is used to determine its level. Moreover, more complex colour descriptions can reuse skills developed on earlier stages.

4.3 Context

The world consist of 268 different colour samples in the CIE 1967 $L^*A^*B^*$ colour space. Colour samples are represented in three dimensions: the L^* dimension represents lightness, the A^* dimension roughly redness-greenness and the B^* approximately yellowness-blueness. The difference between two colour samples is determined by their Euclidean distance. The world contains the focal colours and the consensus samples[6] for English and colour samples created when combining two focal colours in different percentages: 25%, 45%, 55% and 75%, respectively.

(a) Example of a context for the basic colour challenge

(b) Example of a context for the graded membership challenge

(c) Example of a context for the graded category combination challenge

Figure 3: Example of the different contexts speakers can create, depending on their current challenge.

In each interaction the speaker selects the context, which is a subset of the colour samples present in the world. It chooses both the size of the context and the different colour samples that are part of it. The choice depends on the current challenge of the speaker. In the basic colour challenge the context is created by randomly picking three focal colours of English. In the graded membership challenge the speaker chooses five random samples from the consensus samples for English. Finally, in the graded category combination challenge the speaker picks six colour samples that correspond to the combination of two focal colours for English. Figure 3 provides an example of each context.

[5]In the experiment agents can classify the colour sample to the same colour prototype twice, before and after transforming the colour space. For instance, they can create utterances as "blueish blue" if the colour sample is very close to a certain colour prototype.

[6]Colour samples that were consistently named in English by all participants. See Sturges & Whitfield (1995).

4.4 Operational mechanisms

Agents create and learn *constructions*, which can be seen as form-meaning pairs. Constructions are stored in the *construction inventory* of the agent, which defines its vocabulary and grammar. Agents make use of their construction inventory to *formulate*, verbalise a conceptualised meaning, and *comprehend*, extract the meaning representation of an input utterance. Agents start the experiment with an empty vocabulary and enlarge it using different *diagnostics*, used to identify problems during an interaction, and *repairs*, processes to solve diagnosed problems.

Invention: the speaker cannot find a discriminating colour or membership prototype in formulation, caused by the lack of a relevant colour or membership prototype.

- *Diagnostic*: the speaker cannot come up with a discriminative conceptualization of the topic.
- *Repair for lack of relevant colour prototype*: the speaker creates a colour prototype C and sets the colour sample of the topic as its colour prototype. Additionally, the speaker invents a new term t for the colour prototype and creates a new construction relating C with t.
- *Repair for lack of relevant membership prototype*: the speaker creates a new membership prototype M and sets its value to the distance between the colour sample and the prototype of its closest colour category. Additionally, the speaker invents a new term t for the membership prototype and creates a new lexical construction relating M with t.

Adoption: the hearer cannot identify the topic due to an unknown word t, which can refer to either a colour or a membership prototype.

- *Diagnostic*: the hearer encounters an unknown word in the input utterance.
- *Repair for unknown word that refers to a colour prototype*: the hearer uses the feedback from the speaker to create a colour prototype C with the colour sample of the topic as its colour prototype. Additionally, the hearer creates a new construction relating C with t.
- *Repair for unknown word that refers to a membership prototype*: the hearer uses the feedback from the speaker to create a new membership prototype M and sets its value to the distance between the colour sample and the prototype of its closest colour category. Additionally, the hearer creates a new lexical construction relating M with t.

Moreover, agents can also create and learn *grammatical constructions*. These constructions allow agents to express meaning predicates not captured by lexical constructions and restrict the ambiguity of multi-word sentences by imposing form constraints.

4.5 Alignment

In the previous subsection we have introduced *adoption*, which allows hearers to learn new word-meaning associations for both colour and membership prototypes. When adopting an unknown word agents have to decide between adding the observed colour or membership prototype as a new prototype to their inventory or associate the unknown word to an existing one. The decision is based on how close the observed prototype is from the closest prototype in the inventory. A new prototype will be added when the Euclidean distance between both prototypes is bigger than 0.05. When the hearer associates the word to an already existing prototype it introduces competition in its construction inventory, as at least two constructions refer to the same prototype. For instance, the terms “blue” and “azul” would be competitors if they are associated to the same colour prototype.

Agents are provided with a mechanism called *alignment* to manage the competition between constructions in their construction inventory. Each construction has a score with a value between 0.0 and 1.0, and is initialized at 0.5. Scores are used by agents to decide which constructions apply to express one meaning, selecting the one with the highest score. After each interaction the scores of the constructions of the interacting agents are updated using an alignment method called *lateral inhibition* (De Vylder and

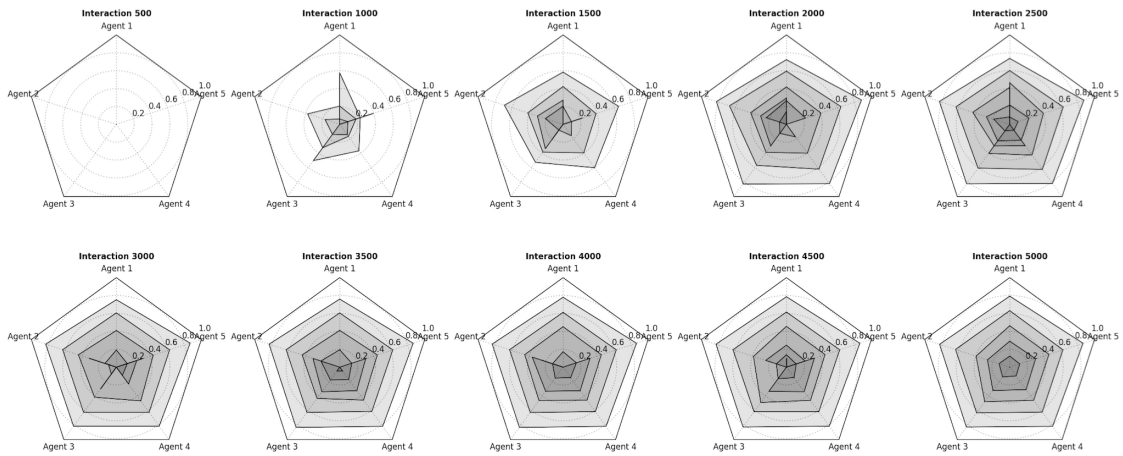

Figure 4: Example of the alignment of membership prototypes for a population of 5 agents. Initially each agent has different prototype values. After each successful interaction the involved membership values of the interacting agents are adjusted. At the end of the simulation the population converges to similar prototypes.

Tuyls, 2006). When a construction has reached a score of 0.0 is removed from the construction inventory of the agent.

Alignment takes into account the outcome of the interaction (i.e. communicative success or failure) to update the scores of the constructions. In a successful interaction both speaker and hearer increase the constructions used by a score $\delta_{increase}$ and punish their competing constructions by a score of $\delta_{decrease}$. If the result of the interaction is a failure, the speaker punishes the constructions used by a score of $\delta_{decrease}$. In the experiment both $\delta_{increase}$ and $\delta_{decrease}$ are set to 0.1.

Agents also align their prototypes. After a successful interaction that involved a membership prototype *M*, both speaker and hearer update the value m associated to that prototype as follows: $m_i = m_{i-1} - \delta rate(m_{i-1} - act_i)$, where m_i and m_{i-1} are the current and previous values of M and act_i the activation of the topic. In the experiment $\delta rate$ is set to 0.05. Figure 4 illustrates the alignment of membership prototypes in a population of five agents.

5 Experimental results

All experimental results have been tested on ten runs, to ensure the consistency of the results. In each trial agents start with an empty construction inventory and colour and membership inventories. The following measures are reported:

- *Communicative success* measures the average performance of the population in the communicative task. When the communication is successful a value of 1.0 is recorded, 0.0 otherwise.
- *Alignment success* measures the average cohesion of the construction inventory on the population. A value of 1.0 is recorded when there was communicative success and both agents would use the same constructions to refer to the topic of that interaction, 0.0 otherwise.
- *Lexical stability* measures the average scores of lexical constructions of the population. A value of 1.0 means that all lexical constructions on each agent have the maximum score.
- *Confidence in challenge* measures the average confidence that the population has for a certain challenge level. It has a value between 0.0 and 1.0.

The resulting dynamics of the experiment with a population of ten agents are shown in Figure 5. Agents start addressing the first challenge, on which the population has to create and coordinate their

construction inventory for basic colour terms. Agents gain confidence for this challenge fast, as both the communicative success and the confidence value for the first challenge rapidly increase. An abrupt drop occurs around interaction 2000, when agents start to reach maximum confidence. As it corresponds to the internal state of boredom, agents enter in the shake-up phase and move to the second challenge. In the course of the second challenge agents are exposed to a bigger diversity of contexts, which makes the alignment of membership prototypes and its associated lexical constructions more difficult. This is also reflected in the evolution of lexical stability, as the average score of lexical constructions drops despite the fact that agents are converging to an optimal lexicon for basic colour terms.

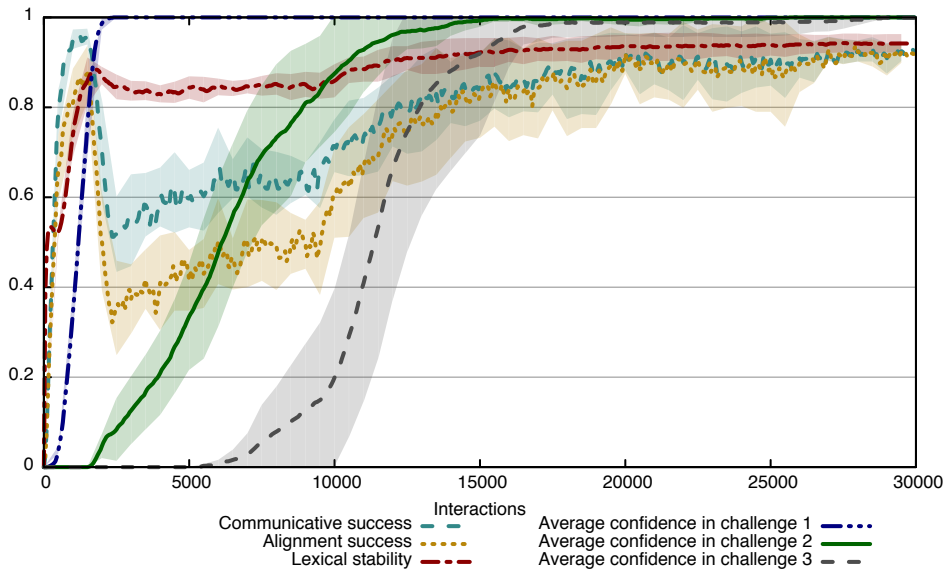

Figure 5: Resulting dynamics of the experiment for a population of 10 agents averaged over 10 runs of 30000 interactions. By the end of the simulation all agents in the population reach a steady communicative success value above 90% and maximum confidence for the three challenges. Error bars represent the maximum and minimum across the different experimental runs.

An overlap between the second and third challenge starts approximately at interaction 5000 where a fraction of the population has already reached maximum confidence for the second challenge. At this point agents identify their internal state as boredom and are motivated to attempt the third challenge. Communication success progressively improves as population succeeds in aligning their construction inventory and membership prototypes. As a consequence of this alignment, alignment success also increases and reaches the same value as communicative success. By interaction 30000 all agents in the population have reached maximum confidence for the three challenges and a steady communicative success value above 90%.

Bleys (2010) showed that agents using these strategies cannot come up with a discriminative conceptualisation in certain situations, which explains why the population does not reach a 100% communicative success even when all agents have reached the highest confidence score for all challenges. However, lexical stability settles to a value around 95%, which means that not all lexical constructions in the population have a score of 1.0. This is caused by different membership categories no longer used but still in the lexicon of some agents. Therefore, population has not fully converged to a minimal lexicon, although they manage to communicate successfully in most contexts.

We have studied the relation between lexical stability and communicative success by testing the same configuration on different populations. Figure 6a presents the resulting communicative success, alignment and lexical stability for a population of two, five, ten and twenty agents (3000, 10000, 50000 and 150000 interactions, respectively). Results show a slight reduction of communicative success as population size increases. More importantly, a little discrepancy between communicative success and alignment is observed in bigger populations. This gap occurs because in some interactions agents prefer distinct

discriminative conceptualisations for certain topics. In other words, different prototypes are triggered as more accurate conceptualisations of the topic in a particular context and therefore agents select different terms to describe the same colour sample.

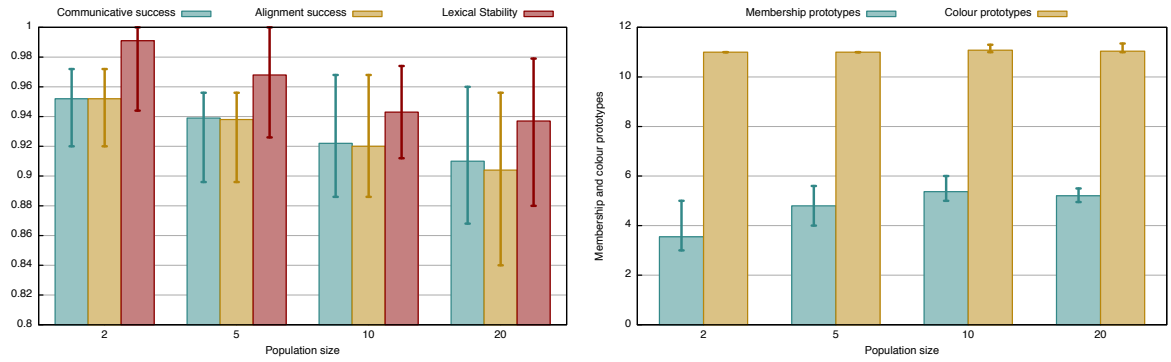

(a) Resulting communicative success, alignment and lexical stability for different population sizes averaged over 10 runs.

(b) Resulting membership and colour categories for different population sizes averaged over 10 runs.

Figure 6: Effect of population size. Figure 6a presents the resulting communicative success, alignment and lexical stability scores in a population of two, five, ten and twenty agents. The scale on the Y-axis is set to the range [0.8,1]. Figure 6b displays the average membership and colour prototypes for the same populations.

This effect can be explained by the fact that bigger populations converge to systems with more membership prototypes (Figure 6b). An increased number of membership prototypes requires more time to align: this helps prototypes which are not spread over the population to stay longer in individual inventories as they are less used. The decrease in communicative success is therefore explained by (a) a lower alignment of agents' construction inventories and membership prototypes and (b) longer presence of non spread membership prototypes among the population that are used in conceptualisation. These results suggest that smaller populations could be able to arise more consistent communicative systems for the domain of colour.

6 Conclusions

In this paper we have studied how a population of agents provided with a motivation system to regulate their complexity is able to develop an artificial language of increasing expressive power to refer to colours. Agents using this system develop to their construction inventory in three progressive stages: they first (a) converge on a language for colour prototypes and then extend its expressive power by (b) developing categories to express degrees of similarity between a colour sample and a colour prototype and (c) combining colour prototypes.

The results obtained show that a population of agents equipped with an architecture of flow successfully manages to progressively develop its communicative skills when trying to remain in a state of flow. Moreover, simulations with different population sizes show that bigger populations converge to systems with more membership prototypes on average.

Acknowledgements

This research was conducted at and funded by the Sony Computer Science Laboratory Paris. Miquel Cornudella is partially supported by a CIFRE grant (agreement no. 2013/0730). We would like to thank those who provided help and feedback, particularly Paul Van Eecke, and three anonymous referees for very constructive remarks after a careful reading of the manuscript. Special thanks go to Joris Bleys, who granted us access to the original implementation of the language strategies used in his experiments.

References

Pierre Andry, Philippe Gaussier, Sorin Moga, Jean-Paul Banquet, and Jacqueline Nadel. 2001. Learning and communication via imitation: An autonomous robot perspective. *Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on*, 31(5):431–442.

Andrea Baronchelli, Tao Gong, Andrea Puglisi, and Vittorio Loreto. 2010. Modeling the emergence of universality in color naming patterns. *Proceedings of the National Academy of Sciences*, 107(6):2403–2407.

Andrea Baronchelli, Vittorio Loreto, and Andrea Puglisi. 2015. Individual biases, cultural evolution, and the statistical nature of language universals: The case of colour naming systems. *PLoS ONE*, 10(5):1–19, 05.

Tony Belpaeme and Joris Bleys. 2007. Language, perceptual categories and their interaction: Insights from computational modelling. In *Emergence of Communication and Language*, pages 339–353. Springer.

Brent Berlin and Paul Kay. 1969. *Basic color terms: Their universality and evolution*. University of California Press, Berkeley.

Joris Bleys and Luc Steels. 2009. Linguistic selection of language strategies. In *European Conference on Artificial Life*, pages 150–157. Springer.

Joris Bleys, Martin Loetzsch, Michael Spranger, and Luc Steels. 2009. The Grounded Color Naming Game. In *Proceedings of the 18th IEEE International Symposium on Robot and Human Interactive Communication (Ro-man 2009)*.

Joris Bleys. 2010. *Language Strategies for the Domain of Colour*. Ph.D. thesis, Vrije Universiteit Brussel.

Joris Bleys. 2012. Language strategies for color. In Luc Steels, editor, *Experiments in Cultural Language Evolution*, pages 61 – 85. John Benjamins.

Miquel Cornudella, Paul Van Eecke, and Remi van Trijp. 2015. How intrinsic motivation can speed up language emergence. In *Proceedings of the European Conference on Artificial Life 2015*, pages 571–578.

Mihaly Csíkszentmihályi. 1990. *Flow: The psychology of optimal experience*. Harper and Row, New York.

Bart De Vylder and Karl Tuyls. 2006. How to reach linguistic consensus: A proof of convergence for the naming game. *Journal of Theoretical Biology*, 242(4):818 – 831.

Emilia Garcia-Casademont and Luc Steels. 2016. Insight grammar learning. *Journal of Cognitive Science*, 17(1):27–62.

Sandra Graham. 1996. Theories and principles of motivation. *Handbook of educational psychology*, 4:63–84.

Helen Lin, M Ronnier Luo, Lindsay W MacDonald, and Arthur WS Tarrant. 2001. A cross-cultural colour-naming study. part i: Using an unconstrained method. *Color Research & Application*, 26(1):40–60.

Martin Loetzsch, Pieter Wellens, Joachim De Beule, Joris Bleys, and Remi van Trijp. 2008. The Babel2 manual. Technical Report AI-Memo 01-08, AI-Lab VUB, Brussels, Belgium.

James B Marshall, Douglas Blank, and Lisa Meeden. 2004. An emergent framework for self-motivation in developmental robotics. In *Proceedings of the 3rd international conference on development and learning (ICDL 2004), Salk Institute, San Diego*, volume 10.

Kathryn Merrick and Mary Lou Maher. 2009. Motivated learning from interesting events: adaptive, multitask learning agents for complex environments. *Adaptive Behavior*, 17(1):7–27.

Pierre-Yves Oudeyer, Frédéric Kaplan, and Verena Vanessa Hafner. 2007. Intrinsic motivation systems for autonomous mental development. *Evolutionary Computation, IEEE Transactions on*, 11(2):265–286.

Eleanor H Rosch. 1973. Natural categories. *Cognitive psychology*, 4(3):328–350.

Richard M Ryan and Edward L Deci. 2000. Intrinsic and extrinsic motivations: Classic definitions and new directions. *Contemporary educational psychology*, 25(1):54–67.

William Schueller and Pierre-Yves Oudeyer. 2015. Active learning strategies and active control of complexity growth in naming games. In *the 5th International Conference on Development and Learning and on Epigenetic Robotics*.

Jean Simpson and Arthur WS Tarrant. 1991. Sex-and age-related differences in colour vocabulary. *Language and speech*, 34(1):57–62.

Kenny Smith, Simon Kirby, and Henry Brighton. 2003. Iterated learning: A framework for the emergence of language. *Artificial Life*, 9(4):371–386.

Michael Spranger. 2013. Evolving grounded spatial language strategies. *KI-Künstliche Intelligenz*, 27(2):97–106.

Luc Steels and Pieter Wellens. 2007. Scaffolding language emergence using the autotelic principle. In *Artificial Life, 2007. ALIFE'07. IEEE Symposium on*, pages 325–332. IEEE.

Luc Steels, Tony Belpaeme, et al. 2005. Coordinating perceptually grounded categories through language: A case study for colour. *Behavioral and brain sciences*, 28(4):469–488.

Luc Steels. 2004. The autotelic principle. In I. Fumiya, R. Pfeifer, L. Steels, and K. Kunyoshi, editors, *Embodied Artificial Intelligence*, volume 3139 of *Lecture Notes in AI*, pages 231–242. Springer Verlag, Berlin.

Luc Steels. 2012. Self-organization and selection in cultural language evolution. In Luc Steels, editor, *Experiments in Cultural Language Evolution*, pages 1 – 37. John Benjamins, Amsterdam.

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

Remi van Trijp. 2013. Linguistic assessment criteria for explaining language change: A case study on syncretism in german definite articles. *Language Dynamics and Change*, 3(1):105–132.