Knowledge diversity under socio-environmental pressure

Jérôme Euzenat

Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LIG, F-38000 Grenoble France
http://moex.inria.fr

Abstract. Experimental cultural evolution has been convincingly applied to the evolution of natural language and we aim at applying it to knowledge. Indeed, knowledge can be thought of as a shared artefact among a population influenced through communication with others. It can be seen as resulting from contradictory forces: internal consistency, i.e., pressure exerted by logical constraints, against environmental and social pressure, i.e., the pressure exerted by the world and the society agents live in. However, adapting to environmental and social pressure may lead agents to adopt the same knowledge. From an ecological perspective, this is not particularly appealing: species can resist changes in their environment because of the diversity of the solutions that they can offer. This problem may be approached by involving diversity as an internal constraint resisting external pressure towards uniformity.

Cultural evolution experiments are performed through multi-agent simulation: a society of agents adapts its culture through a precisely defined protocol [2]. Agents perform repeatedly and randomly a specific task, called game, and their evolution is monitored. This protocol aims to discover experimentally the state that agents may reach and the properties of that state.

Experimental cultural evolution has been successfully and convincingly applied to the evolution of natural language [12; 10]. Agents play language games and adjust their vocabulary and grammar as long as they are not able to communicate properly, i.e., they misuse a term or they do not behave in the expected way. This experimental set-up has been used to test various models in a systematic framework and to provide convincing explanations of linguistic evolution phenomena. Such experiments have shown how agents can agree on a colour coding system or a grammatical case system.

This approach has not been applied yet to knowledge representation directly. Although language experiments involve modifying cognitive representations, e.g., of colour [12] or position [10], their properties are measured through language. So far, the works closest to this have only considered the terminological aspects of ontologies, i.e., associations between terms and concepts [11; 9]. This is the goal of the well-known naming game where agents learn to associate terms to objects or concepts [11]. Experiments have focussed on the way agents agree on terms for naming concepts (chair is the same as seat) and not on the way concepts are organised (through subsumption or disjointness relations for instance, e.g., what is the relation between a chair and a seat with four legs?).

We apply experimental cultural evolution to knowledge. Indeed, knowledge can be thought of as a shared artefact among a population influenced through communication with others. More precisely, it can be seen as resulting from contradictory forces: internal consistency, i.e., pressure exerted by logical constraints, against environmental and social pressure, i.e., the pressure exerted by the world and the society agents live in (see Figure 1). The application of this pressure associated to adaptation operators shapes the knowledge that agents use to reason. This line of work leads to investigate a reformulation of the Sapir-Whorf hypothesis: does and how much the environment and society determine knowledge?

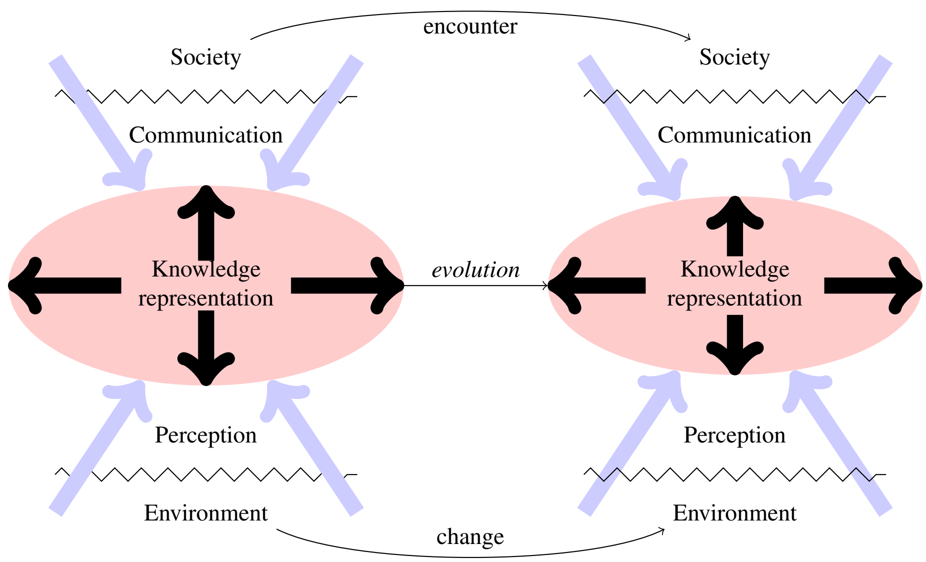

Fig. 1. Pressure exerted on knowledge by society, environment and internal constraints.

This work is only in its infancy and a wide open research playground is facing us. Recently, several teams have started to deal with elaborate knowledge representations using a cultural evolution approach [8; 1; 3; 4; 6; 5]. One specificity of our own work is to represent knowledge in logical formalisms and to actually use the knowledge during interaction.

However, interaction-based adaptation can easily lead to uniformity (in spite of some polarisation results [2]): adapting to environmental and social pressure tends to lead agents to adopt the same knowledge. From an ecological perspective, this is not particularly appealing: species can resist changes in their environment because of the diversity of solutions that they can propound. Usually, uniformity is detrimental to evolution and diversity provides robustness.

This raises further exciting questions such as:

- How is it possible to preserve knowledge diversity within a population of agents? Can this be achieved while preserving communication within the population?
- Under which conditions agents decide to preserve different knowledge or to merge it within common knowledge?
- Does diverse knowledge indeed provide benefits in case of environmental changes? How can agents decide between maintaining their own representation and adopting that of others?

The solution seems to be to embed diversity as internal constraints resisting external pressure towards uniformity. We plan to experiment with different techniques preserving diversity. This may involve two ingredients:

- a way to maintain diverse knowledge representations, and
- mechanisms for agents to create and delete such representations.

They will have to be compared with the other operators with respect to added robustness to environment change and population encounters.

It happens that Artificial intelligence has considered many ways to express different perspectives on knowledge. I myself have had the opportunity to work on leaving room for alternative hypotheses, considering knowledge under different granularities, building consensual knowledge while preserving diverging knowledge, or maintaining alignments across different ontologies. Ontology alignments [7] express the relations between two heterogeneous representations. Contrary to what one may think, their purpose is not to merge diverse ontologies into a single one, but instead, to preserve the diversity of ontologies though permitting interoperability.

Diversity may be achieved through the use, by the same agent, of several ontologies related by alignments in order to preserve its favourite ontology and its links to other ones. This would make individuals able to interact differently with several subgroups of their populations and with individuals from different populations. Hence we will study mechanisms by which agents are able to distance themselves from the common culture, through still mastering this culture. In our case, the success factor, communication, is not directly connected to knowledge: this leaves the door open to reach perfect communication with diverse knowledge.

However, this may lead all agents to maintain and duplicate many correspondences across ontologies. It would be more economical to have more restricted networks. For that purpose, it is necessary to understand on what basis agents may decide to adopt the knowledge of others or to maintain their own alternative.

It will be necessary to measure the long term costs and benefits of such behaviours, testing if and how agents with such representations achieve successful communication. But in the long run, another important hypothesis to test is whether maintaining knowledge diversity provides a selective advantage, in particular, in case the environment of the agents changes. This may be tested through running experiments in which lethal environment changes impair agent viability.

References

1. Michael Anslow and Michael Rovatsos. Aligning experientially grounded ontologies using language games. In *Proc. 4th international workshop on graph structure for knowledge representation, Buenos Aires (AR)*, pages 15–31, 2015.

2. Robert Axelrod. The dissemination of culture: a model with local convergence and global polarization. *Journal of conflict resolution*, 41(2):203–226, 1997.
3. Paula Chocron and Marco Schorlemmer. Attuning ontology alignments to semantically heterogeneous multi-agent interactions. In *Proc. 22nd European conference on artificial intelligence (ECAI), The Hague (NL)*, pages 871–879, 2016.
4. Paula Chocron and Marco Schorlemmer. Vocabulary alignment in openly specified interactions. In *Proc. 16th International conference on autonomous agents and multi-agent systems (AAMAS), Saõ Paolo (BR)*, pages 1064–1072, 2017.
5. Jérôme Euzenat. Crafting ontology alignments from scratch through agent communication. In *Proc. 20th International conference on principles and practice of multi-agent systems (PRIMA), Nice (FR)*, 2017.
6. Jérôme Euzenat. Interaction-based ontology alignment repair with expansion and relaxation. In *Proc. 26th International Joint Conference on Artificial Intelligence (IJCAI), Melbourne (VIC AU)*, pages 185–191, 2017.
7. Jérôme Euzenat and Pavel Shvaiko. *Ontology matching*. Springer, Heidelberg (DE), 2 edition, 2013.
8. Jérôme Euzenat. First experiments in cultural alignment repair (extended version). In *Proc. ESWC 2014 satellite events revised selected papers*, number 8798 in Lecture notes in computer science, pages 115–130, 2014.
9. David Reitter and Andrew Howes. How groups develop a specialized domain vocabulary: a cognitive multi-agent model. *Cognitive systems research*, 12:175–185, 2011.
10. Michael Spranger. *The evolution of grounded spatial language*. Language science press, Berlin (DE), 2016.
11. Luc Steels. The origins of ontologies and communication conventions in multi-agent systems. *Autonomous agents and multi-agent systems*, 1(2):169–194, 1998.
12. Luc Steels, editor. *Experiments in cultural language evolution*. John Benjamins, Amsterdam (NL), 2012.