Joint project team proposal
Long description

Evolving knowledge

mOeX

Carrier:	Jérôme Euzenat
Current team:	Exmo
UMR:	Laboratoire d'Informatique de Grenoble (LIG)
Partner institutions:	INRIA & Univ. Grenoble Alpes

Human being are apparently able to communicate knowledge. However, it is impossible for us to know if we share the same representation of knowledge.

mOeX addresses the **evolution of knowledge representations** in individuals and populations. The ambition of the mOeX project is to answer, in particular, the following questions:
- how do agent populations **adapt** their knowledge representation to their environment and to other populations?
- how must this knowledge **evolve** when the environment changes and new populations are encountered?
- how can agents preserve knowledge **diversity** and is this diversity beneficial?

We will study them chiefly in a well-controlled computer science context.

For that purpose, we **combine knowledge representation and cultural evolution** methods. The former provides formal models of knowledge; the latter provides a well-defined framework for studying situated evolution.

We will consider knowledge as a culture and study the properties of adaptation operators applied by populations of agents by jointly:
- **experimentally** testing the properties of adaptation operators in various situations using experimental cultural evolution, and
- **theoretically** determining such properties by modelling how operators shape knowledge representation.

We aim at acquiring a precise understanding of knowledge evolution through the consideration of a wide range of situations, representations and adaptation operators.

This document is readable as such; its PDF version contains links to external sources.

Contents

A. Scientific position: cultural knowledge evolution

We first motivate the need and relevance of the proposed project (§A.1), before stating explicitly its objectives (§A.2). We then describe the state of the art on subjects related to cultural knowledge evolution (§A.3) and position the proposed project in this landscape (§A.4).

A.1 Motivation

Our societies produce knowledge and data at an ever increasing pace. These knowledge and data are generated in an independent manner by autonomous individuals or companies. They are heterogeneous and their joint exploitation requires connecting them.

However, data and knowledge have to evolve, facing changes in what they represent, changes in the context in which they are used and connections to new data and knowledge sources. These sources are currently mostly maintained by hand. As they grow and get more interconnected, this becomes less sustainable. But if knowledge does not evolve, it will freeze leading to sure obsolescence.

Beyond the production of knowledge on the semantic web and linked data, this problem applies to any domain in which knowledge is produced in a way usable by computers. For instance, smart cities or the internet of things produce a wealth of changing data. The knowledge about this data has to evolve continuously to remain up-to-date as new data sources are encountered and conditions are changing. Knowledge must evolve organically with the life of its users.

This problem lies in the lack of autonomous evolution of heterogeneous knowledge. No one waits for knowledge to be perfect before using it and agents and societies cannot be interrupted for upgrading their knowledge. Hence, knowledge has to be situated, i.e., considered with respect to its use (called situation), and evolve continuously, i.e., without interruption.

A.2 Research objectives and questions

mOeX addresses the seamless evolution of knowledge representations in individuals and populations. The question at the core of our proposal is to understand how to make knowledge representation continuously evolve in presence of environment changes and new knowledge sources. Currently, no satisfactory solution to this problem exists.

To tackle this problem, we start from two specific hypotheses: (*i*) Knowledge is shaped by both experience and communication, (*ii*) which act as selective pressure. More precisely, knowledge is subject to internal constraints, imposed by logical coherence, and external constraints, imposed by the environment and communication with others.

Based on such hypotheses, we will study populations of agents sharing knowledge through interaction. The interactions may be carried out through precisely specified modalities (which may involve direct knowledge exchange, talking, acting together or in presence). After interacting, when they discover that constraints have changed, agents will not relearn knowledge from scratch. Instead, adaptation operators, taking into account the current knowledge and other constraints, will adapt it to the new constraints. We will study how knowledge evolve when these populations (*i*) experience changes in the environment in which they operate, or (*ii*) encounter other populations with which they have to communicate. Our goal is to establish the properties satisfied by the resulting knowledge at the scale of a population of agents. Hence, mOeX aims at establishing the global properties achieved by local operators.

The highly difficult problem is not to have procedures allowing such agents to converge towards a common state of knowledge, but to characterise this state by the properties satisfied by the resulting knowledge. Such properties may, for instance, be:

- Agents converge to a common knowledge representation (a convergence property).
- Agents converge towards different but compatible (logically consistent) knowledge (a logical

epistemic property), or towards closer knowledge (a metric epistemic property).

- That under the threat of a changing environment, agents which have operators that preserve diverse knowledge recover faster from the changes than those which have operators that converge towards a single representation (a differential property under environment change).

Moreover, this has to be guaranteed in the long term involving both environment change and population encounter. Hence, it is critical that convergence towards a particular state does not turn into a handicap when the environment changes. This means that a delicate balance has to be found between adapting optimally to the current situation and interlocutors and evolvability in the long run.

What is radically new here is that these problems are approached from the standpoint of the resulting knowledge representations.

mOeX work will contribute to answer the following questions:

- How can populations with different knowledge (representation) communicate?
- How is their representation influenced by their environment and communication with others?
- How may knowledge diversity be preserved and is it useful?

In their whole generality, such questions apply to human beings, possibly animals, as well as software agents. We propose to study them chiefly in a well-controlled computer science context.

Our ambition is to spark a new approach to knowledge evolution that we call *cultural knowledge evolution*. It designs, studies, and experiments with mechanisms for making knowledge representations serendipitously evolve through their use. This should enable developing and sharing complex knowledge in a more robust way.

Now is the right time to start such a research programme: on the one hand, developments on the semantic web provide us with proven knowledge representation formalisms and tools which have been designed for sharing knowledge; on the other hand, work on experimental cultural evolution provides a solid methodology for carrying out this type of research. This approach has not been applied yet to knowledge representation directly. Both fields are mature enough to be associated.

A.3 Cultural knowledge evolution = Cultural evolution + Knowledge representation

We consider how existing disciplines deal with the problem of knowledge evolution. This state of the art is divided into three broad parts: Knowledge representation (§A.3.1), cultural evolution (§A.3.2) and multi-agent systems (§A.3.3).

A.3.1 Knowledge representation

Knowledge representation has been studied in artificial intelligence for decades leading to semantically well-defined formalisms. Shared representations have been promoted in the context of the semantic web as ontologies. We will consider ontologies and data expressed using the vocabulary of these ontologies, alignments between these ontologies and links between their data. This provides a well integrated framework, designed for sharing knowledge at a wide scale and for which both formal semantics and tools exist [5, 2]. In particular, they are available for reasoning with these ontologies [2], for matching them [22] and for revising ontologies and alignments [20, 29]. Nonetheless, our results should apply to wider classes of representation formalisms.

Sharing knowledge leads to problems of heterogeneity: different individuals or organisations will develop and share different representations. For dealing with these problems, we have developed ontology matching [22]. Ontology matching finds relations between ontology entities and expresses them as sets of correspondences, called *alignments*.

When different knowledge representations are confronted, they may be compatible or incompatible; in logical terms, merging them may be consistent or inconsistent. In the latter case, belief and knowledge revision [1] has been developed for adopting a consistent theory minimising the changes brought for recovering consistency. This has recently been adapted to alignment repair [29] and network of ontology revision [20]. However, revision only characterises the set of possible solutions: it does not take into account the situation in which representations are used for selecting the revision to apply. Hence, applying revision blindly may lead to consistent knowledge irrelevant to the context in which agents live. Through the introduction of adaptation operators, we will allow for selecting the revision according to the situation [19].

Taking advantage of the context has been studied for ontology matching: Interaction-situated semantic alignment [3] considers ontology matching as framed by interaction protocols that agents use to communicate. Agents induce alignments between the different ontologies that they use depending on the success expectation of each correspondence with respect to the protocol. Failing dialogues lead them to revise their expectations and associated correspondences. This approach has recently been studied under the angle of cultural evolution providing encouraging results [9].

Other approaches, such as Bayes networks, neural networks, or Markov decision processes, address this problem by introducing non symbolic processing. They hardly suffer from inconsistency due to smoother conditions. However, their numerical basis hinders the extraction of an explicit shareable knowledge representation that may be communicated across agents.

A.3.2 Cultural evolution

The notion of cultural evolution applies an idealised version of the theory of evolution to culture. Culture, in this context refers to a "patrimony of knowledge accumulating over generations"[1] [8]. It is somewhat related to the notion of meme [14] which follows more closely the genetic evolution analogy. It has been introduced, in ethology [14, 25], population dynamics [8] and anthropology [34]. In such fields, culture may be bird song melodies, food regime or psychological dispositions. Work in cultural evolution is usually based on the observation of long-term behaviours: it relies on the long-term observation of populations or the study of archeological artefacts. In its quantitative form, it is modelled as dynamic systems and compared to observations [8, 34]. Computers have been recently used for exploring small scale phenomena, e.g., the influence of population size on artefact complexity [15].

Cultural evolution experiments are performed through multi-agent simulation: a society of agents adapts its culture through a precisely defined protocol [4, 7]. 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 [37, 35]. Agents play *language games* and adjust their vocabulary and grammar as soon as they are not able to communicate properly, i.e., they misuse a term or they do not behave in the expected way. It showed its capacity to model various settings in a systematic framework and to provide convincing explanations of linguistic 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 [37] or position [35], their properties are measured through language. So far, the closest works have only considered the terminological aspects of ontologies, i.e., associations between terms and concepts [36, 33]. This is the goal of the well-known naming game where agents learn to associate terms to objects or concepts [36]. Experiments have focussed on the way agents agree on *terms* for

[1] http://www.balzan.org/fr/laureats/luigi-luca-cavalli-sforza/a-panoramic-synthesis-of-my-research-anglais-sforza

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?). Only recently, we [19] and others [9] started to deal with elaborate symbolic knowledge representation using a cultural evolution approach.

A.3.3 Multi-agent systems

BDI (Beliefs, Desires, Intentions) is the dominant paradigm in multi-agent systems. Agents attribute beliefs, desires and intentions to themselves and other agents. On a theoretical side agent knowledge is expressed in a modal logic allowing them to reason about such beliefs, desires and intentions (and in particular to compute plans which allow agents to fulfil their desires) [39]. Focusing more on agent's knowledge, epistemic logics [23] are very appealing to reason about what others know and its dynamic version accounts for events [16]. So, these logics may be useful to abstract what occur in the situation we consider; less useful for agent implementations as they often adopt a global view of phenomena.

As mentioned previously, experimental cultural evolution uses directly multi-agent simulations. Such social simulations are also related to the artificial life field, which may include evolutionary simulations. For instance, Aevol simulates the evolution of bacteria colonies over hundred thousands of generations [6].

Cultural evolution suggests connections with various bioinspired approaches, such as evolutionary computation [17], including memetics [14], or biological models, such as evolutionary game theory [28]. An important difference with cultural evolution is that, because it leads to faster adaptation, cultural evolution focuses on horizontal rather than vertical, i.e., genetic, transmission [8]: agents manipulate directly their culture, through adaptation operators, instead of depending on random mutations. So we will not attempt to combine them in the mOeX project.

Moreover, our goal is to study knowledge evolution and what knowledge properties are satisfied by specific operators, thus we can find only limited inspiration in these works.

A.4 Scientific positioning

To investigate the foundations of situated knowledge evolution we need an approach that:

- is general enough so that results can encompass various cases;
- is flexible enough so that many specific settings may be established;
- provides an explicit representation of knowledge in order to communicate it to others;
- allows for continuous local adaptation depending on the situation;
- allows for both theoretical and experimental work.

No single approach offers all these features.

Thus, mOeX will develop the unique combination of knowledge representation and experimental cultural evolution methods. Knowledge representation provides formal models of knowledge; experimental cultural evolution provides a well-defined framework for studying situated evolution. We do not intend to replace symbolic representation, but to complement it.

The reasons why these approaches are well adapted are the following:

- Agents usually cannot wait for slow processes to terminate: they will apply adaptation operators allowing them to communicate;
- Agents need an explicit representation of knowledge in order to communicate it to others;
- Agents do not need that all knowledge is correct before acting;
- Agents do not have a global view of other agents' knowledge: they need a distributed solution with partial knowledge.

The mOeX project thus builds on a variety of fields:

	Language	Knowledge	Space	Robots	Agents	Origin	Diversity	Φ Grounded	Motivation	Selection (ext)	Sem. web. rep.	Fluid grammar/IRL
Axelrod [4]	○	○	●	○	●	○	●	○	○	○	○	○
Steels [37]	●	◐	◐	●	●	●	○	●	○	○	○	●
Oudeyer [31]	○	○	◐	●	●	●	◐	●	●	○	○	○
mOeX	○	●	○	○	●	◐	●	○	○	●	●	○

Table 1: Relations with other experimental cultural evolution efforts. The domain is very broad and each team usually focusses on specific aspects of the domain.

Knowledge representation and reasoning for manipulating the considered knowledge;

Multi-agent systems for simulating the behaviour of agents;

Machine learning for learning knowledge from the environment;

Ontology matching for representing relations between knowledge;

Revision and machine learning for implementing repair operations;

Logic and category theory for establishing properties of adaptation operators.

What is new is that we bring these approaches together to deal with *knowledge evolution in situations* on which we put a strong emphasis.

Although we do not exclude to contribute to these fields if the opportunity occurs, this is not our primary goal.

mOeX departs from the computational cultural evolution work with respect to the search for the origin (of language, speech) [32]. We also do not investigate the important problem of knowledge grounding and embodiment. Finally, we assume perfect language communication between agents. The reasons for this is that we want to simplify the setting and concentrate on the specific problem of knowledge evolution.

We study how entities (called agents) can adapt their knowledge in a decentralised way and evolve their knowledge from very few data. We do not aim at learning the optimal problem solving method sticking to gigantic amounts of data. Understanding which conditions warrant or prevent specific knowledge properties is especially important when this knowledge has to be communicated.

With respect to the state of the art, the most comparable effort is that of Luc Steels and his colleagues [37], though there are some diverging features, the salient one being that we concentrate on knowledge instead of language.

Beside knowledge representation, at the moment, the field which is chiefly concerned by such work is that of multi-agent systems.

Table 1 shows the features taken into account by mOeX and other related projects. Of course, we have selected projects and features, but it seems clear that our focus on knowledge representation is mOeX's distinguishing feature.

More generally, there is a wide, sparsely explored, space of problem related to cultural knowledge evolution and each team occupies one singular point in that space.

B. Methodological approach: experimental and theoretical knowledge evolution

Our methodology relies on the workflow pattern of Figure 1. It involves the following three tasks interacting together in a constant feedback:

- **Designing games** and knowledge adaptation operators;
- **Performing experiments** for testing potential properties of such operators;
- **Analytically proving** such properties or their conditions.

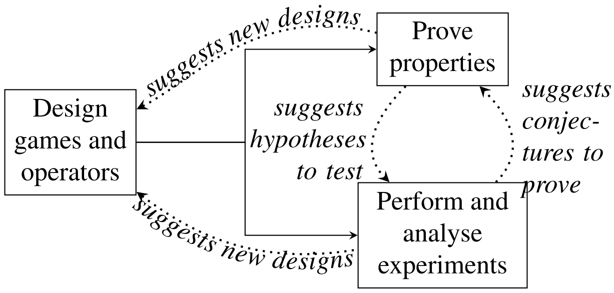

Figure 1: Articulation between our three main types of activities. Both theoretical and experimental studies apply to the same game settings and operators. They provide feedback to each other.

Thus, mOeX adopts a dual theoretical and experimental strategy: determining which adaptation operators provide expected knowledge properties and designing actual mechanisms such that results can be assessed experimentally. Both approaches will cross-fertilise: theory by suggesting adequate operators and properties to test and experiments by providing results to explain.

Each of these tasks follows a different methodology. Hence, we first present our object of study to fix the vocabulary and define what must be provided by the design task (§B.1). Then, we present in more detail the experimental methodology deployed to deal with cultural knowledge evolution (§B.2), and the theoretical tools considered to assess directly its properties (§B.3). Finally, we also introduce the software platform to be developed to support this approach and, in particular, ensure the repeatability and reusability of experiments (§B.4).

B.1 Object of study

Our object of study involves situations in which agents play games. We describe situations and agents, together with events that may modify a situation and properties expected from situations.

Situations They cover the environment in which agents are, the society of involved agents and (eventually) public knowledge available to all agents. The *environment* is seen in our case as data that can be described by agent knowledge representations.

Societies are made of sets of agents called *populations* that may or may not interact together. *Agents* are characterised by their knowledge representation and how and which operators they use to adapt it. They represent the knowledge they have about their environment. These *representations* will be *ontologies*, i.e., a logical axiomatisation of concepts and relations, expressed in description logics. Data is expressed with respect to one or several such ontologies. Because agents may have different ontologies, these ontologies may be related by *correspondences* expressing the relation (equivalence, disjointness, etc.) between concepts of two ontologies. A set of such correspondences is called an *alignment* and we call *networks of ontologies* a set of ontologies related by alignments. Finally, entities found in data may be identified in different ways depending on agents, so (sameAs) *links* between entity references express that they correspond to the same entities. This enables us to use tools such as ontology learning, ontology matching, alignment repair, revision of network ontology to implement agent behaviours. The formal definition of these structures also allows us to measure and study formally the phenomena involved in evolution.

Agents use *adaptation operators* when a game is completed to adapt their knowledge to what they learnt from the game. Such operators may, for instance, come from symbolic knowledge representation (knowledge revision), probabilistic reasoning (maximising likelihood), game theory (maximising expected utility) or may even rely on random choices. However, they will have to rely only on what agents can access (locality). We will first focus on operators requiring from agents limited computational capacity and limited global knowledge.

Finally, **games** describe how agents interact through a precisely defined protocol defining turns in which agents perform actions such as asking a question or identifying an item in the environment. Games are parameterised by the agents playing them and their roles. An example is provided in §B.2.

Events Problems arise due to the occurrence of events such as:

Environment change The environment will determine what can be known and represented; it will also determine what agents can communicate about. Consequently, changes in the environment should constrain agents to adapt their representations. This environment may exert strong selective pressure by killing non adapted agents.

Population encounter Population encounter will put agents in presence of other agents whose knowledge representation is different. This will shape what they communicate (which concepts they can use and what their relations are). In order to improve communication with new encounters, agents may adopt part of their knowledge, and this knowledge may in turn pervade the culture of their population. It may also happen that the two populations merge, especially if populations are defined by their cultures.

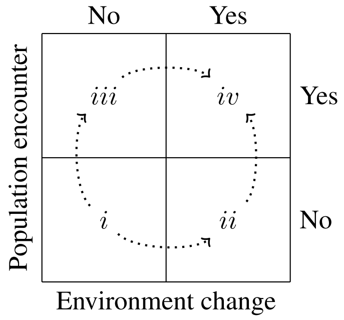

Figure 2: The four quadrants of knowledge evolution organised in two dimensions corresponding to population encounter and environment change.

These changes may be thought of as selective pressure, exerted by the situation in which agents are, that constrains them to adapt their knowledge. We call *adaptation* the action that an agent performs in a particular situation to recover from incorrect knowledge, and *evolution* the longer term effect of adaptation and selection on the knowledge of agents or populations of agents under the pressure of environment change and population encounter. The main goal of mOeX is thus to design adaptation strategies and to characterise their properties. Each type of pressure mandates different experimental modalities. However, it is also important to study them together in order to find how they may be balanced.

Our work will be structured by these two sources of changes which correspond to the two dimensions of Figure 2. It is organised in four quadrants: (*i*) the acquisition of knowledge from the environment and sharing with other agents; (*ii*) the evolution of knowledge when the environment changes; (*iii*) the evolution of knowledge when another population is encountered; (*iv*) the combination of both types of evolution.

Epistemic properties We are specifically interested in the properties satisfied by the representations with respect to their adaptation operators and the situations in which they are. Such properties are called epistemic.

These may be simple convergence properties: There is a representation towards which all agents converge, all agents converge towards the same representation (which can be different depending on the initial situation) or all agents converge towards different representations.

Independently of convergence, there are intrinsic logical properties satisfied by agents representations: They may be consistent, i.e., they have a model. They may also be non empty. Emptiness may be characterised by the state of knowing nothing. However, consistency and non emptiness may be too strong properties already.

Hence, we consider differential epistemic properties, such as monotonic increase or decrease

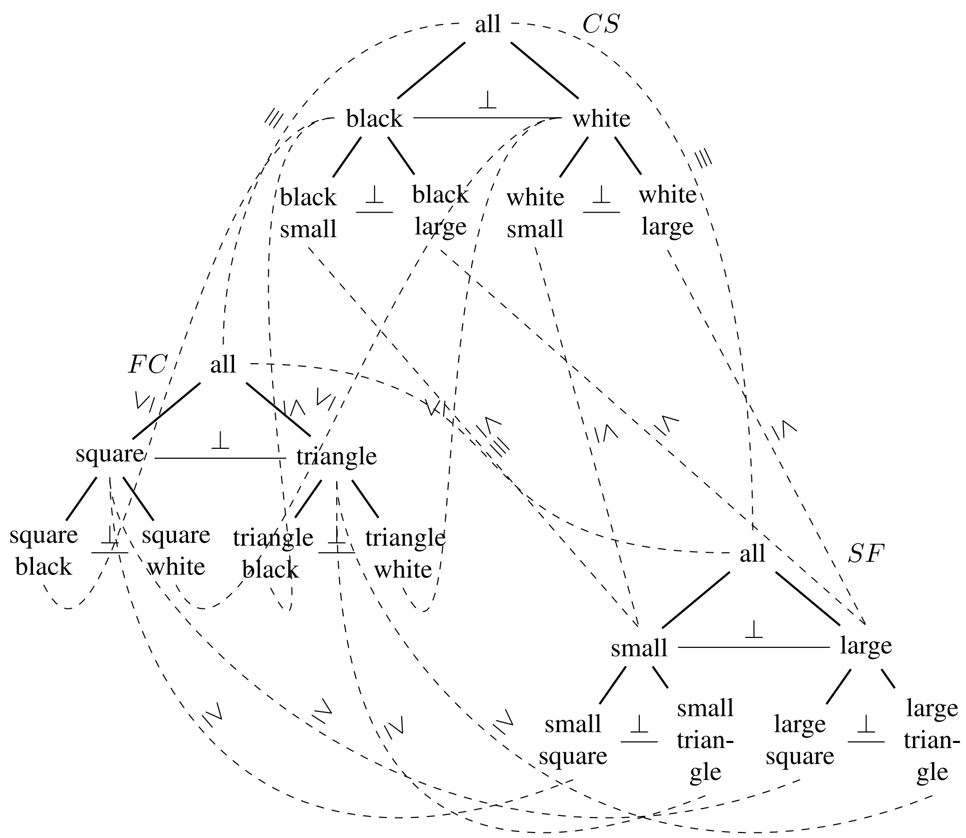

Figure 3: Example of a generated network of ontologies with the exact reference alignments.

of knowledge or decreasing distance between the representations used by the agents [10, 12]. These properties depend on the particular operator, e.g., merging or connecting two classes when they qualify the same set of objects, as well as the environment in which they are used, e.g., the topology of agent networks influences the creation of different cultures [4].

B.2 Experimental culture evolution methodology

We adopt the experimental strategy developed in [37] transposing it from natural language to knowledge representation. We illustrate the approach by summarising an experiment that we performed [19]. This is applied to a communication game independently from the environment, but an acquisition game can be defined along the same lines.

Example Consider an environment populated by objects characterised by three Boolean features: colour = {white|black}, shape = {triangle|square} and size = {small|large}. This characterises $2^3 = 8$ types of individuals: ■, ▲, □, △, ■, ▲, □, △.

Three agents have their own ontology of what is in the environment. These ontologies, shown in Figure 3, identify the objects partially based on two of these features. Here, they are a circular permutation of features: FC (shape, colour), CS (colour, size) and SF (size, shape).

In addition to their ontologies, agents have access to a set of shared alignments. These alignments comprise equivalence correspondences between their top (all) classes and other correspondences. Initially, these are randomly generated equivalence correspondences. For instance, they may contain the (incorrect) correspondence: SF:small $\equiv$ CS:black.

Agents play a very simple game: a pair of agents, a and b, and an object of the environment o, are drawn at random. Agent a asks agent b the class c (source) to which the object o belongs, then it uses an alignment to establish to which class c' (target) this corresponds in its own ontology. Depending on the respective relation between c and c', a may take the decision to change the alignment.

For instance, if agent CS draws the small-black-triangle (▲) and asks agent SF for its class, this one will answer: small-triangle. The correspondence SF:small $\equiv$ CS:black and the class of ▲ in CS is black-small which is a subclass of CS:black, the result is then a SUCCESS. The fact that the correspondence is not valid is not known to the agents, the only thing that counts is that the result is compatible with their own knowledge.

If, on the contrary, the selected object is small-white-triangle ($\triangle$), SF would have made the same answer. This time, the result would be a FAILURE because $\triangle$ belongs to class CS:white-small which is disjoint from CS:black (see Figure 3).

How to deal with this failure is a matter of strategy. Different adaptation operators may be applied:

delete SF:small $\equiv$ CS:black can be suppressed from the alignment;

replace SF:small $\equiv$ CS:black can be replaced by SF:small $\leq$ CS:black;

add in addition, the weaker correspondence SF:small $\geq$ CS:all can be added to the alignment (but this correspondence is subsumed by SF:all $\equiv$ CS:all).

Over time, it is expected that the shared alignments will improve and that communication will be increasingly successful. Successful communication can be observed directly. Alignment quality may be assessed through other indicators: Figure 3 shows (in dashed lines) the correct (or reference) alignments. Reference alignments are not known to the agents but can be automatically generated and compared to the resulting network of ontologies for measuring its quality.

This example illustrates the strong entanglement of cultural evolution (game playing, success rate, experimental setting) and knowledge representation (ontologies, alignments, subsumption and consistency tests).

This experiment resorts to the first quadrant of Figure 2. In this example, the knowledge that evolves is not the ontologies, but the alignments that agents have in common. We are currently designing other games with totally different features: agents do not use alignments and enhance their ontologies when they encounter a new object in the environment or that other agents have an unexpected class.

Experimental set up Such types of experiments can be systematically described by different aspects in the style of [37]. We illustrate this by detailing the experiment of the above example.
Environment: The environment contains objects which are described by a set of n characteristics (we consider them ordered). Each characteristic can take two possible values that, in this experiment, are considered exclusive.
Population: The experiment uses n agents with as many ontologies. Each agent is assigned one different ontology. In this first setting, each agent will have an ontology based on $n-1$ of these characteristics. The ontology is a simple decision tree of size 2^{n-1} in which each level corresponds to a characteristic and subclasses are disjoint. [32] shows how such decision trees can be obtained in a numeric way.
Shared network of ontologies: A complete network of $\frac{n\times(n-1)}{2}$ invertible alignments between the ontologies is shared among agents (public). The network is symmetric (the alignment between o and o' is the converse of the alignment between o' and o) and a class is in at most one equivalence correspondence per alignment.
Initialisation: In the initial state, each alignment contains equivalence correspondences between the most general classes of both ontologies, plus 2^{n-1} randomly generated equivalence ($\equiv$) correspondences.
Game: A pair of distinct agents $\langle a, b\rangle$ is randomly drawn as well as a set of characteristic values describing an individual (equiprobable). The first agent (a) asks the second one (b) the (most specific) class of its ontology to which the instance belongs ($source$). It uses the

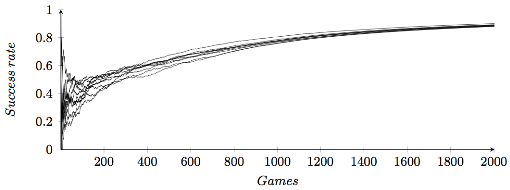

Figure 4: Ten random runs and their overall success rate, i.e., the proportion of games which were successful so far: evidence of convergence of random games (from [19]).

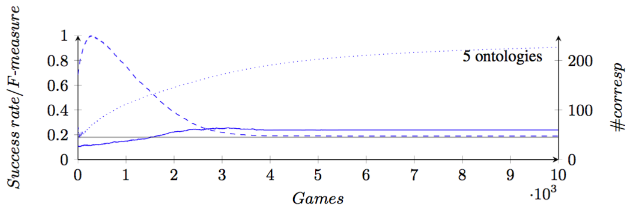

Figure 5: Average success rate (dotted), semantic F-measure (blue) and number of correspondences (dashed) for the add operator compared to the Alcomo and LogMap semantic F-measure baseline (black) (from [19]).

alignment between their respective ontologies for finding to which class this corresponds in its own ontology ($target$). This class is compared to the one the instance belongs to in the agent a ontology ($local$).

Success: Full success is obtained if the two classes ($target$ and $local$) are compatible ($target \geq local$).

Failure: Failure happens if the two classes are disjoint ($local \bot target$). In such a case, the agent a will proceed to repair.

Repair: Several adaptation operators may be used in case of failure:

delete the correspondence is simply discarded from the alignment;

replace if the correspondence is an $\equiv$ correspondence it is replaced by the $\leq$ correspondence from the target class to the source class;

add in addition to the former a new $\leq$ correspondence from the source to a superclass of the target is added. This correspondence was entailed by the initial correspondence, but would not entail the failure.

These three operators share two (local) properties: (*i*) the resulting alignment would have avoided the uncovered mistake; (*ii*) the resulting alignment is logically entailed by the initial one.

Success measure: The classical success measure is the rate of successful communication, i.e., communication without failure. It can always be computed.

Secondary success measure: Several measures may be used for evaluating the quality of the reached state: consistency, redundancy, discriminability. We use two different measures: the average degree of incoherence [30] and the semantic F-measure [18]. Indeed, this setting allows for computing automatically the reference alignment in the network (see Figure 3), so we can compute the F-measure balancing precision and recall.

External validation: The obtained result can be compared with that of other methods. We compared the results obtained with those of two directly available alignment repair algorithms: Alcomo [29] and LogMap repair [27].

Results We explored how primitive cultural adaptation operators can be applied to alignment repair [19]. Figure 4 and 5 presents some of the results. We measured:

- Converging success rate (towards 100% success);
- Coherent alignments (100% coherence);
- F-measures on a par with logical repair systems;
- Number of games necessary to repair increasing very fast.

The convergence towards coherent alignments is an epistemic property and the F-measure increase a differential epistemic property.

B.3 Theoretical approach

In parallel to this experimental work, properties satisfied by adaptation operators will be studied analytically. Although experiments provide insights on what happens in precisely defined situations, theoretical results apply generally to all situations satisfying minimal hypotheses.

Modelling the dynamics of knowledge For that purpose, we have to model agent representations and the effects of adaptation operators on this knowledge. The ability to consider together a large diversity of games will be critical. Because we mainly aim at assessing epistemic properties, the theoretical study will be mostly grounded on logics.

Logics allows to determine semantically the relationships between various knowledge states that agents may reach and isolate invariants in such states or across states. For instance, in the experiment above one can establish that consistency can only increase with the delete operator which removes a correspondence as soon as it causes inconsistent results. Moreover, since individual ontologies are consistent, this will ultimately lead to the full consistency of the network of ontology as soon as the game will uncover all inconsistency causes. This is different with the add operator because it is not restricted to remove correspondences but introduces new ones. However, the added correspondences were already entailed by the removed one. So, adopting a semantic point of view allows to restore the property. Hence, using the semantics of the logic is necessary to prove that the property holds. Differentially, one can remark that delete removes more knowledge: this explains that it converges faster than add.

Extending situated semantic alignment to knowledge games We will, in particular, study such epistemic properties through interaction-situated semantic alignment [3]. Its principle is that protocols used by agents shape possible successful conversations. Alignment permitting such conversations can be extracted from protocol composition (§A.3.1). This may be used in our case by replacing protocols using ontologies by games (or sequences of games). However, games offer more variation in their functioning. In the example given above, alignments are given at the beginning and "repaired" by the process.

In the context of the experiment carried out before, the objects may be made of descriptions of the situation. This can be the network of ontologies in which each agent knows its own ontology and the signature of the other ones and the relations with others provided by the shared alignments. We have provided elsewhere networks of ontologies with subsumption as homomorphisms [20]: a morphism between two networks preserves information. In our example, as games are played, constraints brought by the representation of the other agents will shape a global view of the valid alignments between ontologies, subsuming the initial situation. Hence, an operator can be designed such that each sequence of games, indeed converges towards a limit representation (which is coherent: each situation can be represented consistently). In the case where operators suppress a correspondence or replace it by a weaker correspondence, they will lead to a network of ontologies for which there exist morphisms to the two initial representations.

Mixing games and dynamics From here, there are two particular challenges: mixing games, i.e., having agents able to play different games during their life, and dealing with changes in the environment and encountered populations.

We plan to develop a large library of adaptation operators and game modalities. For a significant number of them, it should be possible to establish such properties as logical convergence, consistency, knowledge preservation, or the conditions for them to hold. However, as soon as agents adopt a different behaviour, e.g., depending on the state of the game add correspondences or merge ontologies, convergence and representation properties will be more difficult to prove.

There are two complementary ways to deal with this: having an operator able to combine game definitions or defining larger classes of games and their properties. Interaction-situated semantic alignment adopted the latter by abstracting from the actual interaction protocols used by agents. We will investigate both approaches.

As mOeX uses two complementary approaches, operator properties may be demonstrated by both experimental measures and theoretical proofs. Ideal results will be obtained when these two routes coincide. However, results obtained experimentally when the theory is non conclusive will be even more stimulating.

A third route to tackle knowledge evolution would be the comparison of our results with models proposed by social sciences and humanities. Although we do not plan, within the mOeX project, to run controlled experiments with people, we would like to confront our own experiments with actual models. This could be done in collaboration with other teams.

B.4 Experimental software platform

An important part of the work carried out in mOeX is experimental. We plan to develop a modular open source platform for setting up, running and publishing experiments. It will be developed with three properties in mind: *(reusability)*, so that others can develop their experiments with it and improve it, even beyond the project; *(repeatability)*, so that experiments could be routinely repeated and independently reproduced; and *(accountability)*, with the idea of publishing our results (positive or negative) continuously and supporting publications of these results seamlessly. In experimental sciences, laboratory logbooks are compulsory. Because our experiments are made *in silico*, maintaining a logbook is assumed to be easier because it can be filled automatically, for large parts.

We have the experience of maintaining the very successful Alignment API [11], used by numerous teams around the world. We have also run the long lasting Ontology Alignment Evaluation Initiative [21]. Finally, we have performed evaluation automation within the SEALS project which has provided an infrastructure to express evaluation procedures and to automate their processing. The projected platform will be more complex in its ability to support a large variety of experiments, this is the reason why we would like to devote specific resources to it.

We have started developing a prototypical environment for cultural knowledge evolution which supports our work. Its goal is to ease performing, controlling, sharing, and repeating experiments. In a first instalment, this software is based on off-the-shelve technologies: available semantic web technologies (RDF, OWL and alignments) are used for representing knowledge; git is used for tracking software versions; a wiki is used as an experiments logbook. This software combo allows to spare development time while filling minimally its purposes.

The current software is made of a Java API describing populations, agents, environments, knowledge models, game, experiments and loggers, in which we have implemented the necessary classes for the experiments in [19]. Using this minimal platform allows us to learn about what is necessary in such platform before eventually deciding, provided that we have sufficient resources, to develop more ambitious software.

We consider building on existing open-source software such as Sumatra [13] (an open source framework for logging, repeating and documenting experiments), Jupyter Notebook (another framework for documenting repeatable experiments), NetLogo [38] or Gama [24] for multi-agent simulation or the researchobject model [26] (a model for experimental research production based on semantic web technologies).

However, our effort will concentrate on the specific developments required for supporting cultural evolution experiments: describing standardised experiments covering the variations that we want to implement, performing large-scale controllable experiments and collecting directly data so that it can feed logbook entries automatically. These entries should be complete enough to repeat experiments automatically and run them again with different initial conditions.

C. Initial research workplan

The proposed research is integrated and not divided into particular axes. We describe it in two ways: first, particular topics that we plan to investigate more precisely (§C.1); second, the temporal organisation of our first four-years plan (§C.2). We will start our work with simplifying hypotheses, and complexify them at every four-years period.

C.1 Prioritary topics

We will concentrate our work on specific challenges at the core of Figure 6: *adaptation*, *evolution* on environment changes and population encounter, *diversity* and *holistic balance* of adaptation operators. There are other topics deserving investigation, e.g., what makes a population, knowledge and language coevolution, or agent reproduction and selection, but we choose to focus on those that we found more fundamental.

Adaptation: How do agents adapt their knowledge representation to their environment and to other agents? Agents have two sources from which to learn knowledge: their environment and the other agents. Knowledge is learnt incrementally as agents discover it. This requires the definition of adaptation operators to reorganise this knowledge when it becomes inadequate, i.e., when agents make a wrong prediction about the world or communication fails. The properties of these operators are assessed by the properties of the knowledge that agents build in the long run. For instance, they may all have an accurate representation of the environment, they may all have a representation that allows them to communicate, they may all have the same representation, or they may all have a consistent representation.

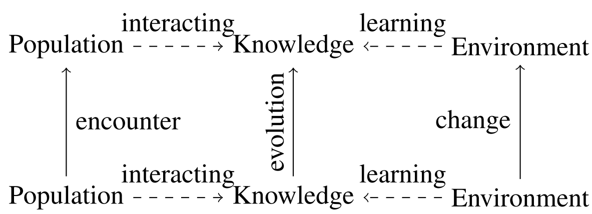

Figure 6: Simplified representation of the considered situations. Evolution is the constant result of applying (interacting and learning) adaptation operators.

The goal of the adaptation topic is to determine the influence of (*a*) the kind of knowledge representation, (*b*) the type of acquisition mechanism/adaptation operator, and (*c*) the articulation between different such mechanisms on (*i*) communication, and (*ii*) knowledge representations. This will require answering questions as: How to articulate symbolic and interaction-based techniques so that populations converge to successful communication?

Evolution with respect to environment change is concerned with how agents can react when continuous or dramatic changes occur in their environment. Such changes may be the introduction or suppression of different types of objects or a different spatial organisation. Obviously, the more agent knowledge will be fitted to the environment the less adequate it will become when the environment changes. Our goal is to assess the same long-term properties of the same type of operators as before. This is related to both action modalities (interacting with the environment and communicating with others) as the environmental change may lead to an agent changing its representation which in turn causes miscommunication. This will help understanding if operators preserving heterogeneous representations across agents are more robust to environment changes.

Evolution with respect to population encounter takes into account the introduction of new populations sharing or not sharing the same environment, but potentially having different knowledge. Again, the properties of the same operators as before, used locally to adapt knowledge, will be assessed with respect to population encounter. Across two populations, eventually able to communicate successfully, new miscommunication may occur. We aim at determining which operators lead populations to merge and which ones preserve different populations. The result of adapting knowledge to such situations may be that the knowledge of one population is imposed to the other because this eases communication or that all knowledge is preserved, but some of it is shared. This can be rephrased as epistemic properties of the operators.

Diversity: How can agents preserve knowledge diversity and is this diversity beneficial? The problem with adapting to environmental and social pressure is that it tends to have agents adopting the same knowledge. However, usually, uniformity is detrimental to evolution and diversity provides robustness. 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. This raises further exciting questions such as: How can agents decide between maintaining their own representation and adopting that of others? In our ontology matching work, we developed the concept of alignments relating two representations and allowing to use one or another [22]. We plan to use this notion in a particular way in which each agent may develop connections between knowledge representations in order to achieve communication while preserving its knowledge. This also requires studying when agents decide to adopt the knowledge of others or to maintain their own. Adaptation operators that preserve diversity will be compared with the other operators with respect to added robustness to environment change and population encounters.

Holistic balance These four individual topics are themselves difficult pieces of research. However, their holistic consideration is yet another challenge by itself. It requires the combination of operators able to adapt knowledge in case of environment change and population encounter which preserves diversity. The difficulty lies in the contradictory forces of adaptation, which would tend to stabilise knowledge, and, evolution under environmental and social pressure, which mandate to break this stability when reacting to external changes. It is critical that convergence towards a particular state does not become a handicap when the environment changes. This means that a delicate balance has to be struck between adapting to the current situation and interlocutors and evolvability in the long run. Diversity comes into this as a way to mitigate the adaptation effect as it should avoid over-fitting to the current situation.

C.2 Four-years plan

The goal of our first activity will be to demonstrate the relevance and fertility of combining the cultural evolution approach and knowledge representation methods.

Because the mOeX programme is very ambitious and the design of evolution mechanisms is extremely challenging, we will first concentrate on *adaptation* and *evolution*, leaving *diversity* for later work.

In the first period of two years, we will focus on the design of atomic adaptation operators working independently and investigate their properties individually. We will progressively increase the complexity of such operators. Variations that we plan to investigate include:

- Altering ontologies instead of alignments;
- Learning ontologies and alignments from the environment;
- Introducing objects that kill the agent in case of misidentification;
- Considering the whole network of alignments, and not just one alignment;
- Allowing agents to adopt each others ontologies.

The proposed operators may work well in controlled experiments, e.g., [19], but be difficult to generalise to wider contexts. Hence, in the second period, and as soon as possible, they will be considered globally in order to address the critical challenge of identifying how these modalities interfere and which mechanisms can balance these individual operators.

The first period will quickly ensure that results on atomic operators can be achieved. Moreover, this program covers the two types of changes introduced in §C.1 (environment change and population encounter). This approach is summarised in the following (indicative) chart:

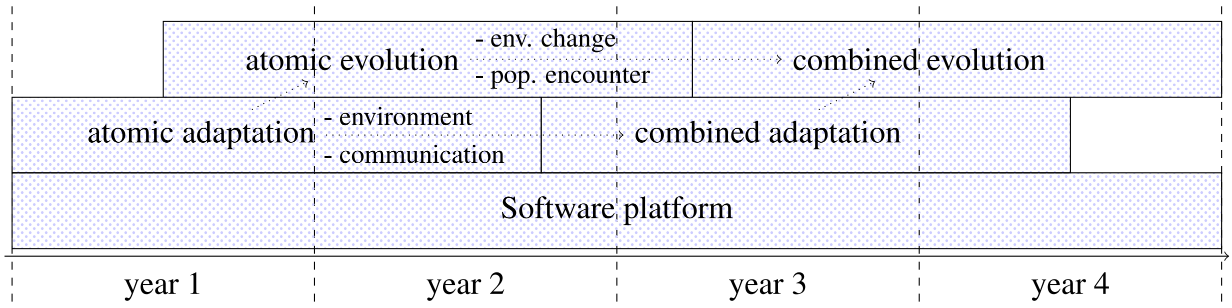

D. Expected outcomes: sparking a new discipline

We discuss the expected outcome of mOeX, through its expected direct research contributions (§D.1) wider impact (§D.2), potential applications (§D.3), and finally risks (§D.4).

D.1 Research contributions

The goal of the mOeX research programme is to acquire a broad knowledge of cultural knowledge evolution. During its life span, mOeX is expected to contribute:

- Adaptation operators for dealing with environment changes;
- Adaptation operators for supporting heterogeneity with encounters;
- Strategies for combining these operators that contribute to robust behaviours;
- Experiments showing (or disproving) the effects of these operators;
- Proofs of specific properties of knowledge representations reached by agents concerning adaptation, evolution and diversity;
- A software platform enabling repeatable experiments with various cultural knowledge evolution settings.

D.2 Impact

If successful, mOeX will open a whole new research trend in knowledge representation. It should initiate a wider movement to investigate situated knowledge evolution. In the longer term, this should lead to fluid interoperability of cognitive systems through principled continuous knowledge adaptation.

Indeed, for long, we have been lacking a way to make knowledge representations evolve in an organic manner. Belief revision and update provide useful guides but little means to choose the appropriate change with respect to the situations in which knowledge is used. Using cultural evolution to unveil properties of adaptation operators will pave the way to new research on:

- knowledge and language coevolution, in particular relaxing the perfect communication hypothesis;
- the influence of knowledge representation formations on epistemic properties and algorithms;
- stacking genetic agent/operator selection on top of knowledge selection;
- tackling directly the problem of the origin of knowledge representation, i.e., starting with empty representations, currently left aside.

More practically, this is prone to improve the sustainability of knowledge and data publication efforts. Seeing them as evolving ecosystems of representations (knowledge formalised as ontologies, data, alignments), provided by various actors (companies, individuals, automated agents), our work would contribute to make these representations evolve more seamlessly than they currently do. Avoiding hiccups in the communication of data and knowledge would further improve socio-economical performance.

D.3 Applications

At the beginning of the 2000s, when we started the work on the semantic web, we had the results of 20 years of research on knowledge representation and the experience of distributing them on the web behind us. Many of us had a clear idea of what was ready and what were the missing pieces. We had in mind very precise applications, some of which materialised; some unexpected ones appeared.

We are now in a state in which it is time to go back to the foundations in order to obtain results that, in a longer term, could have the same applicative fertility. This proposal is deliberately fundamental work. As computer scientists, we aim at understanding how to build systems with particular properties.

However, one can consider broad application areas in which this work may be exploited.

The ***semantic web*** is made of a gigantic amount of interrelated knowledge and data provided by independent actors. Because data and knowledge sources are independent, they can evolve independently; because they are related, this may hamper their use. Controlling all evolution aspects is impossible and avoiding it would freeze their development and use. Applications maintaining and using the knowledge expressed on the semantic web may detect failures when they occur and react by adapting it. Allowing more serendipitous evolution of this knowledge will contribute to increase the sustainability and relevance of the semantic web. The same semantic web technologies are used by applications with intensive need to share data such as smart cities or semantic internet of things.

Smart cities are for a large part made of data interchange. In this context, semantic web technologies are well-positioned to ensure interoperability (we contributed to this in the Ready4SmartCities European project). However, it is expected that the knowledge expressing the context of data will have to change. For instance, in smart cities, modelling knowledge about traffic profiles during winter events may be useful for organising emergency plans. However, if it is not updated as road equipment and citizen behaviour change, it becomes obsolete very quickly. Knowledge evolution must happen organically with the life of the city. As for the semantic web, services able to acquire knowledge, to detect faults and to adapt their knowledge continuously would allow for a smoother operation of such smart cities.

In ***domestic robotics,*** robots will have to adapt their knowledge representation to the situations shared with human beings. Let us take as an example a new kind of domestic robot supposed to both entertain and help an elderly person. Such a robot would have to acquire knowledge about the physical environment of the person. It would also have to interact with this person and her relatives. The robot will have to adapt its knowledge to the points of concern of the cared person, which may range from taking prescribed medicines as well as greeting relatives for their birthdays to holding a conversation. Moreover, through time, both the material environment and the intellectual capabilities of the person will evolve. The robot will have to respond to this seamlessly, by acknowledging mistakes and adapting its knowledge and behaviour. The question of adopting the same knowledge as the assisted person or to preserve heterogeneous knowledge is an important one when the mental capabilities of this person are declining.

We are currently not involved in the development of such applications.

D.4 Risks

We identify and discuss various risks tied with the mOeX proposal:

Breeding fields Coupling knowledge representation and cultural evolution has never been tackled. The fact that these two approaches come from completely different origins makes the task very risky: these techniques may not fit together, making it impossible to obtain results.

Unfruitfulness The project may fail by establishing only ridiculous trivial properties or by being unable to show anything, because there is nothing much to be shown. Although we obtained promising results [19], this cannot be fully excluded.

Mixing operators The most serious threat to this program is that the various operators that will be developed in a strictly controlled context may not work *together* (holistically) properly. The problem of finding a balance between different knowledge representations in presence of a changing world is not understood yet. In particular, it could be easy to end up with, on the one side, each agent having its own representation, i.e., no culture, or, on the other side, all agents having the same representation; this is not a desirable state with respect to diversity.

Disciplinary splits The mOeX program is not particularly part of an established domain. People working on such topics publish in venues dedicated to robotics, artificial life, linguistics and artificial intelligence. This dispersion make that there is no obvious media in which cultural evolution work is obviously welcome and, beside publication records and personal careers, this may jeopardize impact.

The following table summarises our analysis of these risks:

Risk	Impact	Probability	Mitigation plan
Breeding fields	High	Low	We already mitigated this risk by experimenting [19]
Unfruitfulness	High	Low	This is assumed, we can only know by trying.
Mixing operators	High	Medium	If such a risk materialises, then we will have to learn from it and understand what prevents operators to work together.
Disciplinary splits	Medium	High	This risk may be an opportunity to contribute to the organisation of a new field

Indeed, this project could fail. However, we are convinced that we could learn a lot from failure.

E. Institutional landscape: the mOeX network

Here, we situate mOeX with respects to our previous work (§E.1), the team we are starting with (§E.2), other teams we are interacting with locally, nationally and worldwide (§E.3). We finally discuss about the continuation of our current work (§E.4) and research contracts (§E.5).

E.1 Exmo work

Exmo has contributed to the take-off of the semantic web. In particular, we have initiated and led the work on ontology matching [22]. Ontology matching deals with the heterogeneity faced by the semantic web by finding relations between ontologies. Some important results Exmo has obtained include (*i*) providing a semantics for ontology alignments and compatible algebras for manipulating them, (*ii*) launching and running matcher evaluation, creating new evaluation measures, (*iii*) designing original matchers based on fixed point computation and context, (*iv*) developing and maintaining the Alignment API, the main API for manipulating alignments. We have recently extended this work to interlinking linked data.

In case it is necessary, all information about Exmo, including publications, is available from its web site (http://exmo.inria.fr).

The programme that we want to develop builds on our results, but largely departs from these activities.

E.2 The mOeX team

The mOeX proposal is set up by Exmo's core permanent researchers:

Jérôme Euzenat, PhD 1990, HDR 1999, Univ. Grenoble; DR INRIA has created and led the Exmo team since 2003. His research field is definitely knowledge representation with work of theoretical, experimental and software nature. His research has always focussed on confronting diverse representations. He brought contributions to: (*i*) context-propagation truth maintenance systems, considering a corpus of knowledge under different hypotheses, (*ii*) symbolic temporal granularity, establishing the constraints under which the granularity of a representation changes, (*iii*) collaborative knowledge base construction, defining a collaboration protocol inspired by peer review that allows confronting different knowledge corpora, (*iv*) multimedia document adaptation, adapting document execution to target environments depending on semantic constraints, and (*v*) of course, ontology matching. http://exmo.inria.fr/~euzenat/

Jérôme David, PhD Univ. Nantes 2007, MCF Université Grenoble Alpes Jérôme David's main research interest is the semantic web and more specifically ontology matching and data interlinking. In his PhD, he proposed an instanced-based ontology matcher. Since he joined Exmo, he has contributed to ontology matching, by studying and proposing ontology distances, and to data interlinking by proposing key/link key-based approaches to this problem. He will be invaluable on all experimental aspects of the project. http://exmo.inria.fr/~jdavid/

Manuel Atencia, PhD UA Barcelona 2010, MCF Université Grenoble Alpes Manuel Atencia's main research interests are knowledge representation and reasoning, and the semantic web. In his PhD he addressed the problem of semantic heterogeneity in agent communication by defining an interaction-situated semantic alignment and a protocol that agents can use to compute this alignment while interacting. This approach has been formalised using category theory. In the last years, he has been involved in different topics, like the study of models and algorithms for data interlinking in the context of linked data, and the definition of a semantics for weighted ontology alignments. His expertise will be a decisive advantage for the theoretical part of the project. http://exmo.inria.fr/~atencia/

E.3 Relation with other teams

Concerning the French research strategy 2015, the project can be considered in the "Information and communication society" major challenge and in the "men and culture" action programme.

E.3.1 In Grenoble and at LIG

mOeX feels at home in the "human digital ecosystems" motto of the LIG. Currently Exmo belongs to the "Large Scale Processing of Data and Knowledge" axis. This is quite natural and mOeX could belong to this axis as well, though the "Interactive and Cognitive Systems" axis may also be a relevant choice. We plan to globally open discussion at LIG.

The work proposed in mOeX comes at a complement with that of various LIG teams:

Tyrex, Slide, Steamer are our usual colleagues on semantic web along different approaches. We will continue to work with them, though our focus shift will lead us to also consider other teams.

Magma covers many aspects of multi-agent systems, though no work seems specifically dedicated to agent knowledge and communication.

Adèle develops research in software engineering with increasing emphasis on autonomic computing in which software resilience can be improved through the exploitation of feedback. ***CTRL-A*** (also a INRIA team candidate) works on control aspects of autonomic computing. The topic is far away from knowledge representation but shares the resilience motivation and the adaptation to feedback approach.

Polaris (also a INRIA team candidate) is connected through its work towards experimental reproducibility.

Hence, mOeX would rather be complementary to other teams.

We have sought a wider collaboration with teams from social sciences and humanities. However, it seems that no team in Grenoble is working on related topics.

E.3.2 At INRIA

Similarly to LIG, Exmo is deeply anchored by in the "Data and Knowledge Representation and Processing" INRIA theme.

Orpailleur (Nancy Grand Est), Graphik and Wimmics (Sophia Antipolis Méditérannée) are our usual partners on semantic web activities. As before their focus are different from that of mOeX.

Beagle (Grenoble Rhône-Alpes) has one of its main activities in fine grain simulation of evolution phenomena. Although their primary focus is biology, we have a lot to learn from them on the design of experimental workbenches.

Diverse (Rennes Bretagne-Atlantique) is applying evolutionary techniques to improve software robustness through diversity.

Alpage (Paris Rocquencourt) mentioned, during their last evaluation, their interest in the work on Cultural language evolution. Their primary interest is language and not knowledge, however, a long-term project to develop a gaming platform for language corpus acquisition is something that we shall look into.

Flowers (Bordeaux Sud-Ouest) is the obvious connection since Pierre-Yves Oudeyer has contributed to the work at Sony CSL. The coverage of Flowers is very broad with general topics of interest for mOeX. However, Flowers' focal points such as learning, motivation, or autonomy and applications in robotics differ from mOeX focus on symbolic knowledge and communication.

The mOeX program meets those of various teams but its very goal is not treated by any of these.

E.3.3 Worldwide

We list teams working on related topics and with which we have contacts. Some of them are related to our activities in the semantic web, but they are likely to be pursued:

Instituto d'Investigació en Intelligència Artificial, Barcelona Marco Schorlemmer works on situated ontology matching and on the categorical interpretation of evolutionary processes. Carles Sierra leads the design of auto-evolving communities mediated by agents.

Universität Otto von Guericke, Magdeburg Till Mossakovski's group works on ontology blending which could be used both as a theoretical and practical tool.

Fondazione Bruno Kessler, Trento Luciano Serafini and Chiara Ghidini contribute to the semantics of networks of ontologies. The team of Fausto Giunchiglia in Trento is also studying knowledge diversity, so we will exchange with them, albeit they consider this from the classical standpoint of ontology matching and data integration.

Universität Mannheim We have worked closely with Heiner Stuckenschmidt's group on ontology matching. They have more specifically developed alignment repair.

Sony CSL, Paris is one of the institutions that developed the work on cultural language evolution [37]. Although we are working on knowledge representation instead of language we have met some members of the team and would like to reinforce these contacts.

UNIRIO, Rio We are collaborating with Kate Revoredo and Fernanda Baião on using inductive logic programming in knowledge evolution.

We are also in contact with the Essence Marie Curie initial training network dedicated to the "evolution of shared semantics in computational environments" and involving some of these teams. Jérôme Euzenat gave a tutorial at their 2015 summer school.

E.4 Long tail activities

The goal of the mOeX project-team is to take over the activities of the Exmo project team. They have been carried out around three main types of activities:

– Ontology matching: most of our contribution to the field is behind us, Exmo has fully carried out its research programme;

– Data interlinking: we still have some bright ideas to publish and plan to do this as soon as possible;

– Dynamic knowledge: was mostly an anticipation of the mOeX activities and they naturally can be carried out there.

We plan to maintain some pieces of software which where developed in the two first activities as they are useful in some of the experiments that we will carry out.

E.5 Contracts

For information, our current contracts are finishing in 2016, we currently have only one submission:

Elker (submitted to ANR, 2017-2021) is a cooperative research project proposal on Link keys (§E.4) with colleagues from INRIA Nancy and U. of Vincennes.

References

[1] Carlos ALCHOURRÓN, Peter GÄRDENFORS, and David MAKINSON, On the logic of theory change: partial meet contraction and revision functions, *Journal of symbolic logic* 50(2):510–530, 1985.

[2] Grigoris ANTONIOU, Paul GROTH, Frank VAN HARMELEN, and Rinke HOEKSTRA, *A semantic web primer*, 3rd ed., The MIT press, Cambridge (MA US), 2012.

[3] Manuel ATENCIA and Marco SCHORLEMMER, An interaction-based approach to semantic alignment, *Journal of Web Semantics* 13(1):131–147, 2012.

[4] Robert AXELROD, The dissemination of culture: a model with local convergence and global polarization, *Journal of conflict resolution* 41:203–226, 1997.

[5] Franz BAADER, Diego CALVANESE, Deborah MCGUINNESS, Daniele NARDI, and Peter PATEL-SCHNEIDER, eds., *The description logic handbook: theory, implementations and applications*, Cambridge University Press, 2003.

[6] Bérénice BATUT, David PARSONS, Stephan FISCHER, Guillaume BESLON, and Carole KNIBBE, *In silico* experimental evolution: a tool to test evolutionary scenarios, *BMC Bioinformatics* 14(Suppl. 15):511, 2013.

[7] Joanna BRYSON, The role of stability in cultural evolution: innovation and conformity in implicit knowledge discovery, in: Virginia DIGNUM and Frank DIGNUM, eds., *Perspectives on culture and agent-based simulations*, Springer, pp. 169–187, 2014.

[8] Luigi Luca CAVALLI-SFORZA and Marcus FELDMAN, *Cultural transmission and evolution: a quantitative approach*, Princeton University Press, Princeton (NJ US), 1981.

[9] Paula CHOCRON and Marco SCHORLEMMER, Attuning ontology alignments to semantically heterogeneous multi-agent interactions, in: *Proc. 22nd ECAI Conference, The Hague (NL)*, pp. 871–879, 2016.

[10] Jérôme DAVID and Jérôme EUZENAT, Comparison between ontology distances (preliminary results), in: *Proc. 7th International Semantic Web Conference (ISWC)*, vol. 5318, Lecture notes in computer science, pp. 245–260, 2008.

[11] Jérôme DAVID, Jérôme EUZENAT, François SCHARFFE, and Cássia TROJAHN, The Alignment API 4.0, *Semantic Web Journal* 2(1):3–10, 2011.
[12] Jérôme DAVID, Jérôme EUZENAT, and Ondřej ŠVÁB-ZAMAZAL, Ontology similarity in the alignment space, in: *Proc. 9th International Semantic Web Conference (ISWC)*, vol. 6496, Lecture notes in computer science, pp. 129–144, 2010.
[13] Andrew DAVISON, Automated capture of experiment context for easier reproducibility in computational research, *Computing in science and engineering* 14:48–56, 2012.
[14] Richard DAWKINS, *The selfish gene*, Oxford University Press, Oxford (UK), 1976.
[15] Maxime DEREX, Marie-Pauline BEUGIN, Bernard GODELLE, and Michel RAYMOND, Experimental evidence for the influence of group size on cultural complexity, *Nature* 503:389–391, 2013.
[16] Hans van DITMARSCH, Wiebe van der HOEK, and Barteld KOOI, *Dynamic Epistemic Logic*, Springer, Heidelberg (DE), 2007.
[17] Agoston EIBEN and Jim SMITH, From evolutionary computation to the evolution of things, *Nature* 521:476–482, 2015.
[18] Jérôme EUZENAT, Semantic precision and recall for ontology alignment evaluation, in: *Proc. 20th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 348–353, 2007.
[19] — First experiments in cultural alignment repair (extended version), in: *Proc. ESWC 2014 satellite events revised selected papers*, Lecture notes in computer science 8798, pp. 115–130, 2014.
[20] — Revision in networks of ontologies, *Artificial intelligence* 228:195–216, 2015.
[21] Jérôme EUZENAT, Christian MEILICKE, Pavel SHVAIKO, Heiner STUCKENSCHMIDT, and Cássia TROJAHN, Ontology Alignment Evaluation Initiative: six years of experience, *Journal on Data Semantics* XV:158–192, 2011.
[22] Jérôme EUZENAT and Pavel SHVAIKO, *Ontology matching*, 2nd ed., Springer, Heidelberg (DE), 2013.
[23] Ronald FAGIN, Joseph HALPERN, Yoram MOSES, and Moshe VARDI, *Reasoning about knowledge*, The MIT press, Cambridge (MA US), 1995.
[24] Arnaud GRIGNARD, Patrick TAILLANDIER, Benoit GAUDOU, Duc An VO, Quang Nghi HUYNH, and Alexis DROGOUL, GAMA 1.6: Advancing the Art of Complex Agent-Based Modeling and Simulation, in: *Proc. International Conference on Principles and Practice of Multi-Agent System (PRIMA)*, vol. 8291, pp. 117–131, 2013.
[25] Mark HAUSER, *The evolution of communication*, The MIT Press, Cambridge (MA US), 1997.
[26] Kristina HETTNE, Harish DHARURI, Jun ZHAO, Katherine WOLSTENCROFT, Khalid BELHAJJAME, Stian SOILAND-REYES, Eleni MINA, Mark THOMPSON, Don CRUICKSHANK, Lourdes VERDES-MONTENEGRO, Julian GARRIDO, David DE ROURE, Oscar CORCHO, Graham KLYNE, Reinout van SCHOUWEN, Peter 'T HOEN, Sean BECHHOFER, Carole GOBLE, and Marco ROOS, Structuring research methods and data with the research object model: genomics workflows as a case study, *Journal of biomedical semantics* 5(1):41, 2014.
[27] Ernesto JIMÉNEZ-RUIZ, Christian MEILICKE, Bernardo CUENCA GRAU, and Ian HORROCKS, Evaluating Mapping Repair Systems with Large Biomedical Ontologies, in: *Proc. 26th Description logics workshop, Ulm (DE)*, pp. 246–257, 2013.
[28] John MAYNARD SMITH, *Evolution and the theory of games*, Cambridge University Press, Cambridge (UK), 1982.
[29] Christian MEILICKE, Alignment incoherence in ontology matching, PhD thesis, Universität Mannheim, 2011.

[30] Christian MEILICKE and Heiner STUCKENSCHMIDT, Incoherence as a basis for measuring the quality of ontology mappings, in: *Proceedings of the 3rd ISWC international workshop on Ontology Matching*, pp. 1–12, 2008.

[31] Pierre-Yves OUDEYER, *Aux sources de la parole: auto-organisation et évolution*, Odile Jacob, Paris (FR), 2013.

[32] Pierre-Yves OUDEYER and Frédéric KAPLAN, Discovering communication, *Connection science* 18(2):189–206, 2006.

[33] David REITTER and Andrew HOWES, How groups develop a specialized domain vocabulary: a cognitive multi-agent model, *Cognitive systems research* 12:175–185, 2011.

[34] Peter RICHERSON and Robert BOYD, *Not by genes alone: how culture transformed human evolution*, University of Chicago press, Chicago (IL US), 2005.

[35] Michael SPRANGER, *The evolution of grounded spatial language*, Language science press, Berlin (DE), 2016.

[36] Luc STEELS, The origins of ontologies and communication conventions in multi-agent systems, *Autonomous agents and multi-agent systems* 1:169–194, 1998.

[37] — ed., *Experiments in cultural language evolution*, John Benjamins, Amsterdam (NL), 2012.

[38] Uri WILENSKY and William RAND, *An introduction to agent-based modeling: Modeling natural, social and engineered complex systems with NetLogo*, The MIT press, Cambridge (MA US), 2015.

[39] Michael WOOLDRIDGE, *Reasoning about rational agents*, The MIT press, Cambridge (MA US), 2000.