

Interoperable AI: Evolutionary Race Towards Sustainable Knowledge Sharing

Stefan Sarkadi, Andrea G. B. Tettamanzi, Fabien Gandon

▶ To cite this version:

Stefan Sarkadi, Andrea G. B. Tettamanzi, Fabien Gandon. Interoperable AI: Evolutionary Race Towards Sustainable Knowledge Sharing. IEEE Internet Computing, In press. hal-03807461v1

HAL Id: hal-03807461 https://inria.hal.science/hal-03807461v1

Submitted on 10 Oct 2022 (v1), last revised 11 Oct 2022 (v2)

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.

INTEROPERABLE AI: EVOLUTIONARY RACE TOWARDS SUSTAINABLE KNOWLEDGE SHARING *

Ștefan Sarkadi, Inria, Université Côte d'Azur, CNRS, I3S Sophia-Antipolis, France stefan.sarkadi@inria.fr Andrea G.B. Tettamanzi Université Côte d'Azur, CNRS, I3S, Inria Sophia-Antipolis, France andrea.tettamanzi@unice.fr

Fabian Gandon Inria, Université Côte d'Azur, CNRS, I3S Sophia-Antipolis, France fabien.gandon@inria.fr

ABSTRACT

The advancement and deployment of artificial intelligent agents brought numerous benefits in knowledge and data gathering and processing. However, one of the key challenges in deploying such agents in an open environment like the Web is their interoperability as they currently mostly run in silos. In this paper we report on a simulation and evaluation based on evolutionary agent-based modelling to empirically test how sustainable different strategies are for knowledge sharing in open multi-agent systems (MAS). Our results show the importance of translation-based approaches and the need for incentives to support these.

1 Introduction

We are still waiting for a large scale deployment of intelligent agents on the World *Wild* Web [1]. One of the key challenges in deploying such agents in an open environment like the Web is their interoperability.

To ensure interoperability, approaches in distributed AI (MAS) range from enforcing a global schema to allowing as many schemas as existing components, as well as every possible hybrid solution in between these extremes. Smith [2] hypothesized that a universal ontology is unlikely to be realizable. Inversely, as soon as more than one schema exists, the question of mapping and translating between these schemas arises, with an associated cost.

Our objective in this paper is to understand, from the perspective of economics of deployment, the advantages of choosing one approach over another according to the costs and benefits of each to reach interoperable MAS. Each approach is treated as a knowledge-sharing strategy, and more precisely a strategy for adopting one vs a strategy to adopt several ontologies. The simulation remains at the level of strategies and does not consider specific ontologies. We define knowledge-sharing games powered by an evolutionary mechanism for social learning that selects agents of a population to interact over long periods of time.

2 Related Work & Motivation

Early work on interoperability in MAS has mostly focused on the standardisation of Agent Communication Languages, like FIPA and KQML.

More recently, game-theoretical and mechanism design agent-based approaches have been applied in the area of ontologies [10, 11]. These multi-agent approaches are very different from ours both conceptually and implementation-wise, as they do not aim to study the economics of knowledge-sharing under evolutionary constraints.

^{*} Citation: Authors. Title. Pages.... DOI:000000/11111.

All these approaches share a similar conceptual foundation for ensuring interoperability, namely that two interacting agents share or produce a common representation of the world [3].

Some argue that if a common ontology had been established for artificial agents to interact on the Web, then we would have seen those seamless interactions. But, perhaps Smith [2] was right suggesting that such a universal ontology is unlikely to be realized. Alternatively, it might be the case that mediation services could instead be more sustainable to overcome the challenge of interoperability.

To test Smith's intuition, we do not engineer a specific MAS. Instead, we derive our approach from the evolutionary agent-based modelling of cooperation, communication, and knowledge-sharing in complex social systems as the ones in [6, 7, 9]. None of these previous studies have looked at the question of shared language vs. translation strategies in the context of ensuring sustainable communication. This is exactly the research question we set for this paper.

3 Research Questions & Method

We know that the diversity of resources in modern infospheres like the Web calls for artificial ecosystems with a diversity of interacting agents ranging from reactive to deliberative paradigms, maintaining the information ecology [5] and supporting their hybrid communities in sustainable ways. We also know that all approaches in AI come with associated costs of resources and maintenance, from classical knowledge representation methods such as ontologies [4] to latest machine learning ones.

In particular, alternative agent communication strategies can differ in terms of software maintenance and in terms of costs incurred to store and share knowledge.

At the two extremes, the burden could be put either on the agents to pay the price of having different ontologies when they need to interact, or on a shared ontology provider to ensure a *lingua franca*. Each extreme has its costs and benefits and there is a continuum of hybrid options in-between. For instance, letting agents have their own ontologies means they will pay the cost of whatever transformation they have to perform when they communicate, whereas maintaining a shared ontology comes with the price of achieving consensus, commitment and versioning every time it needs to be changed. But how sustainable would these different strategies be from an economy of knowledge perspective?

In terms of economy of knowledge sharing, the Web can be seen as an open MAS, where knowledge is mostly public but can also be private or restricted to a group or even a couple of users or agents. We can also say that Web agents that contribute to Wikipedia are building a public good. The over-arching goals of Web agents would be to share as much knowledge as necessary to ensure they achieve their goals without disclosing their entire knowledge base (there might be some knowledge that needs to remain private). From an evolutionary perspective, this search for interoperability makes Web agents compete with each other, and assuming that the knowledge shared is relevant, the agents with the fittest strategies are able to share the most knowledge.

One way to identify more 'sustainable' knowledge-sharing is to allow agents to optimise between the strategies that they choose to share knowledge and the cost of doing so in various contexts. If these agents were able to choose between different knowledge sharing strategies, and able to learn from their peers the costs and benefits of these strategies, then how would such a society of agents behave in the long-run? What kind of strategy would they choose, from an evolutionary perspective? If they had the option to choose a common universal ontology as a knowledge sharing strategy, would these 'sustainability-aware' agents choose it over other strategies?

As a result, We have formulated a set of research questions to test the different alternatives:

- 1. Does the choice of a common universal ontology emerge as an evolutionary stable strategy in populations of self-interested knowledge-sharing agents that learn from each other through imitation?
- 2. Do voluntary interactions between agents of a population have an impact on the choice of strategy?
- 3. Does the number of agents actively engaged in knowledge-sharing between themselves affect the choice of strategy?
- 4. Do mediating agents/services (that translate knowledge between agents) play a role in the sustainability of such systems?

We answer these questions empirically by performing extensive simulations with evolutionary agent-based models.

4 Agent-Based Modelling of Knowledge Sharing

We use the word strategy, to refer to a strategy for sharing knowledge the context of an evolutionary game. Agents are able to change their strategies by switching to other existing strategies through social learning. We consider that strategies for knowledge sharing correspond to strategies for choosing one or more languages, or one or more ontologies for knowledge-sharing in MAS. We will use language and ontology interchangeably in this context of the knowledge sharing games and we will not be considering a specific language or ontology, only global strategies.

4.1 Evolutionary Knowledge-Sharing Games

When sharing knowledge publicly, e.g. on the Wild Web, it is reasonable to assume that self-interested intelligent agents would do so rationally and strategically. In particular, they take into account the existing incentives for sharing knowledge as well as what knowledge to keep private to preserve some of their advantages. The Knowledge-sharing Game (KSG) is a race between a set of knowledge sharing strategies that pushes agents to pick the best one. Notably, KSG optimises cooperation. What the agents in the KSG are actually doing is to compete with each other to find the best cooperative strategy, and even though KSG is a competitive game, it is actually a race between agents towards selecting the best mechanism for cooperative knowledge-sharing.

We define a KSG as a tuple $(S, N, M, c, r, \theta, s, \mu, \alpha, \beta)$, whose parameters are defined in Table 1.

Р	Value	Definition
$S \neq \emptyset$		a non-empty set of strategies
N	100	the number of agents in a population to play a KSG
N_{S_i}	$\leq N$	the number of agents using strategy $S_i \in S$
M	2 or 5	the number of agents selected to play the KSG at a given iteration from a
		population of N
k	10	the entire amount of knowledge of any agent
\overline{q}	2	the amount of knowledge every agent decides to keep private
c = k - q	8	the amount of knowledge that every agent contributes to a KSG
r	1.5	a multiplication factor that acts as an incentive for sharing a particular amount
		of knowledge c
$\theta \in [0, 1]$	0.5	the compatibility rate between any two or more agents that interact
$s \ge 0$	10^{3}	the social learning (imitation) strength
$\mu \ge 0$	0.001	the exploration rate i.e. the inclination of agents to randomly adopt another
		strategy
ϵ	2	the cost of using one own's ontology/language
α	2	the cost of translating knowledge to another ontology/language
β	2	the cost for an agent of adopting a common ontology/language
Т	10^{4}	number of iterations of a KSG
		Table 1: KSG parameters

Table 1: KSG parameters.

For each strategy, a payoff is computed that takes into account the costs and benefits of choosing the respective strategy. We explore the changes in frequencies over time of agents who are able to learn strategies for knowledge sharing from other agents. For each KSG, we perform explicit computations of what payoffs the agents will receive given a sub-population that is selected to play the game at each iteration. The relative differences between the payoffs obtained by the agents with different strategies in the sub-population determine the probability that an agent will adopt a different strategy.

Social learning s represents the tendency of an agent that is selected to adopt a strategy that, compared to its current strategy, maximises the agent's payoff. The higher the value of s, the stronger the tendency of adopting the better strategy. If s = 0, this gives a random outcome, while if $s \xrightarrow{\infty}$ this will always make the agent select the stronger strategy. The exploration rate μ represents random mistakes in actions as well as purposeful exploration regardless of relative payoffs. While μ represents the probability of an agent to randomly adopt any other strategy, the probability of an agent *i* with a payoff P_i to switch to another agent's strategy *j* with the payoff P_j is given by $P_{ij} = \frac{1}{1 + exp[-s(P_i - P_i)]}$. This stochastic approach allows us to dynamically represent how the frequencies of the different types of agents evolve over time.

We assume that the incompatibility between two or more agents that interact can vary, hence the compatibility label between M agents that interact $\theta \in [0, 1]$ and the compatibility difference given by $1 - \theta$. For some strategies, the

Strategy	Payout & Cost
Own Language <i>O</i> : The agents share knowledge in their own language	Payout: $c \times r \times \theta - c$
only.	Cost: ϵ
Common Language CL: The agents communicate in a common lan-	Payout: $c \times r - c$
guage shared by all the agents in a population. The common language	Cost: β^N
has an exponential cost of maintenance.	
Common Language v2 CL_2 : idem but the cost of maintaining the	Payout: $c \times r - c$
common language is proportional to the number of agents in the system	Cost: $\beta \times N_{CL_2}$
that actually choose to use it.	
Translation TR : The agents choose to pay a cost for translating the	Payout: $c \times r - c$
knowledge that they wish to share with other agents.	Cost: $\alpha \times c$
Translation Per Knowledge TR_2 : <i>idem</i> but agents only pay for the	Payout: $c \times r - c$
translation of knowledge that they do not understand i.e. for the exact	Cost: $(1 - \theta) \times \alpha \times c$
pieces of knowledge for which they do not share exactly the same	
ontologies.	
Hybrid <i>H</i> : Whenever possible, agents communicate in a common	Payout: $c \times r - c$
language the knowledge that they wish to share. For other pieces of	Cost: $\alpha \times c \times (1 - \theta) + \theta \times \beta^{N_{CL} + N_H}$
knowledge they pay the cost for translation.	
Hybrid 2 H_2 : The agents that communicate in a partially known	Payout: $c \times r - c$
common language that is maintained for agents that actually choose	Cost: $\alpha \times c \times (1-\theta) + \theta \times \beta \times (N_{CL_2} + N_{H_2})$
to use it. For knowledge that they cannot express in the common	
language, they pay the cost of translating that knowledge, similarly to	
TR_2 .	
*Loner L: The agents that choose not to interact in the KSG. We call	The payout and costs of the Loner always
the agents that choose no interaction a Loner agent. The role of the	amount to a fixed payoff $L = \sigma$
Loner is to give a chance to other strategies to invade the population.	

Table 2: Different strategies for knowledge sharing with their payout and cost. All parameters were defined in Table 1.

compatibility between agents determines how much knowledge can be exchanged without additional cost. For an agent using a translation service the cost α of communicating knowledge is multiplied by c, which measures the amount of knowledge it intends to share while β is the cost for the adoption and maintenance of a common language (ontology).

4.2 Strategies for Knowledge Sharing

In Table 2 we define 7 strategies that agents can use to play a KSG, and one strategy (Loner) for agents that choose not to play the KSG in voluntary games. These strategies are selected through mutation or social learning at every iteration of a KSG, as described in the previous section.

Additionally, KSGs can be played under two mutually-exclusive conditions regarding interactions. The first condition is that interactions between agents of a population are assumed to be compulsory. KSGs played under this condition are called *Non-Voluntary*. Each agent that plays the game receives a reward (a payout) which is multiplied by (M - 1) (the number of other sampled agents playing the same game), whereas the costs are proportional to the fraction of other sampled agents playing the entire population $\frac{M-1}{N-1}$.

The second condition relaxes this assumption and agents are allowed not to participate in KSGs. Hence the latter KSGs are called *Voluntary*. In this paper we simulate KSGs under both of these conditions. To do this, we have defined two different payoffs for each strategy, found in Table 4-Appendix.

While the payoffs for non-voluntary KSGs follow straight from the strategies defined in Table 2, voluntary KSGs require an additional strategy. A voluntary KSG means that there is a probability that the agents are not willing to exchange knowledge with others in the population. To model this phenomenon, we introduce an additional strategy to the KSG, namely, the Loner strategy L.

To compute the payoffs of the other strategies in voluntary KSGs, we need to take into account the probability that all other M - 1 sampled individuals are Loners, namely:

$$P_{\sigma} = \frac{N_L!(N-M)!}{(N_L - M + 1)!(N-1)!} = \frac{\binom{N_L}{M-1}}{\binom{N-1}{M-1}}$$
(1)

Table 3: Experimental Setups. These interactions happen under different conditions identified by the prefix "P" if they are public and/or suffix "-V" if they are voluntary (agents have the possibility to choose not to interact at all).

Type: KSG1	Type: KSG2	Type: KSG*
KSG1, KSG1-V, PKSG1, PKSG1- V	KSG2, KSG2-V, PKSG2, PKSG2- V	KSG*, KSG*-V, PKSG*, PKSG*- V
Four KSG1 games that we designed to represent interactions where agents can use their own predefined ontology (O), can try to adopt an exponentially costly common ontology, can choose to pay for a service that translates everything they are trying to communicate, or can adopt a strategy that mixes the adoption of a common ontology with paying for a translation service. Under the assumptions of KSG1, KSG1-V, PKSG1, and PKSG1-V and different setups we can test how the strategies CL and TR perform against each other and also against other strategies such as H (the hybrid between CL and TR) and OL (and L where relevant).	Four KSG2 games that we de- signed to represent interactions where agents can use their own pre- defined ontology (OL), can try to adopt a common ontology that has a cost proportional to the number of agents that adopt it (contrast- ingly to KSG1 games where this cost is exponential), can choose to pay for a service that tranlsates only the incompatible information that agents are communicating (in con- trast to KSG1 games where the ser- vice translates everything that is communicated), or can adopt a strat- egy that mixes the two previous strategies. We test the hypotheses on CL_2 for assumptions of KSG2, KSG2-V, PKSG2, and PKSG2-V and strategies CL_2 , TR_2 and H_2	Four KSG* games that represent interactions where both sets of as- sumptions about costs of adopting translation and common ontologies from the previous two game types are present. The aim of this setup is to directly compare the two types of strategies for knowledge sharing. Under the assumptions of KSG*, KSG*-V, PKSG*, and PKSG*-V, which we can test under different setups how the strategies CL , TR , CL_2 , TR_2 , perform against each other and against OL (and L where relevant).
Strategies: Common Language, Translation, Own Language, Hybrid, and No In- teraction* (Loner)	Strategies: Common Language 2, Translation 2, Own Language, Hybrid 2, and No Interaction* (Loner)	Strategies: Common Language, Common Lan- guage 2, Translation, Translation 2, Own Language, and No Interaction* (Loner)
Tested Hypotheses: $TR > CL, CL > TR$	Tested Hypotheses: $TR_2 > CL_2, CL_2 > TR_2$	Tested Hypotheses: $CL_2 > CL, CL > CL_2, TR_2 >$ $TR, TR > TR_2, TR_2 > CL_2,$ $CL_2 > TR_2$
Verified Hypotheses: TR > CL	Verified Hypotheses: $TR_2 > CL_2$	Verified Hypotheses: $CL_2 > CL, TR_2 > TR, TR_2 > CL_2$

where

 N_L is the number of Loners in the population.

The other 7 strategy payoffs are adjusted in order to take into account the likelihood of sampling a Loner from the agent population pool when two or more agents interact in knowledge sharing games. This is not the case in the non-voluntary model, where there are no potential non-participants to the KSG.

5 Experimental Design

The name of each experiment encodes its configuration. First, we separate two sets of experiments corresponding to two types of KSGs designed to test the impact of private vs. public knowledge exchange as these two settings exist on the Web. The first set (prefixed by "KSG") explores the direct competition between 2 agents that share knowledge

between themselves in a private way. The second set (prefixed by "PKSG") explores the idea of knowledge exchanged publicly. We have set the number of agents that interact in these games M = 5 following the study presented in [7], which uses similar evolutionary agent-based modelling techniques to ours, based on [6, 8]. If we consider that knowledge is shared publicly (e.g., on the Web) in domains where agents interact in groups, a first reasonable setup is to adopt the same value for number of agents engaged in these interactions.

The KSG contribution c was given the same arbitrary integer value for all strategies in order not to affect the strategy payoffs, and r was given the value of 1.5 to represent a 50% return on the contribution. The number of iterations $T = 10^4$ was set to show the effect of a selecting a strategy over great periods of time, whereas the population size N = 100 was selected to represent a relatively large and open MAS, but for which we also perform a sensitivity analysis to show its effect of even larger MAS in the Appendix. We adopt the fixed intermediate value for social learning, $s = 10^3$, which was picked with the purpose to represent a relatively strong capability of agents to adopt other strategies without assuming that agents are always successful at it. Contrastingly the value of 0.001 was picked for the exploration rate μ to represent a low tendency to explore. For both s and μ we perform a sensitivity analysis in the Appendix. θ was given the value 0.5 to represent a 50% chance (coin flip) of sharing knowledge successfully with another agent when both agents use their own language. ϵ , α , β were given the same fixed integer value of 2 to isolate their effect on the payoffs, but for which we also provide a sensitivity analysis in the Appendix. Loner payoff σ was adopted from [8] with the value 1.

Both pairwise and public interactions are experimented in non-voluntary and voluntary games, the latter being identified by the "-V" suffix. Moreover, there are two groups of assumptions about Common Language and Translation that we wish to explore. These assumptions correspond to the underlying language strategy assumptions that correspond to KSG1 (exponentially costly common ontology) and KSG2 (where a common ontology has a cost proportional to the number of agents that adopt it), respectively. KSG* represents an additional experiment designed to compare the two groups of underlying assumptions by putting them in the same competition.

Table 3 summarizes the different combinations of strategies and experimental conditions. In total, we have 12 experiments summarized and, for instance, the name "PKSG1-V" identifies a Knowledge Sharing Game where interactions are voluntary, public and based on the KSG1 assumptions. The games, strategies and hypotheses we will now introduce and test are summarised in Table 3 and the parameters are in Table 1.

5.0.1 Hypotheses

In order to test the sustainability of each strategy we formulate sets of hypotheses for each type of KSG, in pairs of the form X > Y that means that the long-run avg. frequency of strategy X is greater than that of strategy Y. What interpret this as X is more sustainable than Y (see Table 3).

6 Results

For each type of KSG we present the statistical results of long-run average frequencies for 100 independent runs of 10^4 iterations. The first run starts from a population of agents that use the Own Language (O) strategy. Each subsequent run starts from the population distribution generated at the end of the previous run. All the results and the sensitivity analysis for PKSG*-V are included in the Appendix (Supplementary Material). They are presented as means, standard deviations, t-tests, and One-way ANOVA results in Tables. Significant results where p < 0.01 are represented by *. The results are also visually presented in figures, where the average frequencies of strategies are presented with variance bars (that represent one standard deviation from the mean). Additionally, we plot the dynamics of a single individual-based simulation run for each KSG, where each coloured line represents the frequency of a given strategy at a given time t.

We report on each test we performed and its results in the rest of this section.

6.1 Pairwise Interactions two at-a-time *M*=2

KSG1 and KSG2 (*Data:* Tab. 5a and Tab. 5b; Fig.2c and Fig.2d – *Confirmed Hypotheses:* TR > CL, $TR_2 > CL_2$.) **KSG1-V and KSG2-V** (*Data:* Tab. 6a and Tab. 6b; Fig.3c and Fig.3d – *Confirmed Hypotheses:* TR > CL, $TR_2 > CL_2$.)

KSG* and KSG*-V (*Data:* Tab. 7a and Tab. 7b; Fig.4c and Fig.4d – *Confirmed Hypotheses:* $CL_2 > CL$, $TR_2 > TR$, $TR_2 > CL_2$.)

What have we shown? Generally, mediation strategies in open systems in which either voluntary and non-voluntary interactions take place, are more sustainable than the implementation of common ontologies for knowledge sharing in **pairwise** interactions between rational social learning agents.

6.2 Public Interactions five at-a-time *M*=5

PKSG1 and PKSG2 (*Data:* Tab. 8a and Tab. 8b; Fig.5c and Fig.5d – *Confirmed Hypotheses:* TR > CL, $TR_2 > CL_2$.)

PKSG1-V and PKSG2-V (*Data:* Tab. 9a and Tab. 9b; Fig.6c and Fig.6d – - *Confirmed Hypotheses:* TR > CL, $TR_2 > CL_2$.)

PKSG* and PKSG*-V (*Data:* Tab. 10a and Tab. 10b; Fig.7c and Fig.7d – *Confirmed Hypotheses:* $CL_2 > CL$, $TR_2 > TR$, $TR_2 > CL_2$.)

What have we shown? Similarly to the experimental conditions from pairwise interactions, mediation strategies in open systems in which either voluntary and non-voluntary interactions take place, are more sustainable than the implementation of common ontologies for knowledge sharing in **public** interactions between rational social learning agents.

7 Discussion

In this paper we have studied the sustainability of different knowledge-sharing strategies from an evolutionary social learning perspective by running extensive individual agent-based simulations. To do so, we have defined a set of hypotheses about the sustainability of different strategies for knowledge sharing and we have designed 12 knowledge-sharing games to test these hypotheses. The previous section summarized the results of the games vs our hypotheses.

From these results we showed that under the assumption of cooperation and the possibility of decentralisation, translation strategies in open MAS in which either voluntary or non-voluntary interactions take place, are more sustainable than shared ontology strategies for both public and private knowledge sharing between rational agents. We also showed that Translation 2 (where agents pay only for the items of knowledge that need mediating) is more sustainable than Translation (where agents pay for always using mediation services) in all configurations. In other words, we observed that the translation strategies outperform common language strategies in terms of sustainable knowledge sharing. We also observe some niche cases, however unlikely these might be, where common-language strategies with non-exponential costs are more sustainable than translation strategies. Such niche cases can arise for both voluntary and non-voluntary interaction conditions (illustrated in KSG* Fig. 4c and KSG*-V Fig. 4d), as well as in cases of group interactions (illustrated in PKSG* Fig. 7c and PKSG*-V 7d). This could suggest the possibility of having different scales of social structures, e.g. a very small local sub group with a shared ontology (e.g. a factory room) vs a large and distributed group (e.g. a B2C or B2B network). Another observation we can make is that a distributed system of strategies emerges $(CL_2 + TR_2)$ instead of a single hybrid strategy (H_2) . Intelligent agents seem to prefer separate systems of communication to run in parallel within the same population.

We must note that our results are only valid for cooperative systems, where the goal of the agents is to optimise knowledge-sharing with other agents - to select the best strategy for cooperation. In specific cases, centralisation CL_2 can be the optimal strategy, for instance in an IoT factory where a single agent dictates the language that must be used for inter-operation. For such a group of agents that use the same strategy but in a larger MAS, they can be collapsed into a single agent. Another example from the real world is when a certain entity, e.g. Google or Apple dictate what kind of software code must be implemented by independent software developers in order for their apps to run on iOS or 'play nicely' with Google services. But in its current state, Web 2.0 is not a good example of a decentralised open MAS because it is a collection of silo-ed and centralised systems.

When we designed the KSGs, we had in mind the possibility of agents to self-organise as truly decentralised distributed systems with independent autonomous agents that interact and self-organise in order to share knowledge. As future work, our KSG model opens the path towards studying the interoperbility of autonomous and self-organising AI agents on existing graphs-based structures of the Web.

In a nutshell, assuming cooperation in knowledge-sharing and the possibility of a fully decentralized Web of resources and services, our results provide reasons for adopting translation strategies. However we still have a number of challenges to address even just for interoperability. We are in a chicken-and-egg problem where we don't have translation services because there is no demand, no existing agent in the wild, and inversely we have no agents in the wild because, among other problems, they have no support to interoperate.



(a) Long-run average frequencies for PKSG1-V.



(b) Long-run average frequencies for PKSG2-V.



(c) Long-run average frequencies for PKSG*-V.

Figure 1: Long-run average frequencies for a sample of 100 individual-based runs in public and voluntary interactions.

By extension, finding incentives for translation services to be provided and maintained is problematic. The question includes the study of business models for translators or opportunities for saving costs in IT systems, e.g., by supporting the outsourcing of some IT services. Classic business models on the Web include advertisement, but this kind of funding model doesn't apply to interactions between software agents. Clearly, to us, the question of incentive and business models for translation-based interoperability appears to be on the critical path of achieving MAS on the Web.

8 ACKNOWLEDGMENTS

We thank the 3IA Côte d'Azur ANR-19-P3IA-0002 and the HyperAgents project ANR-19-CE23-0030 for their support.

References

- [1] J. Hendler, "Where are all the intelligent agents?". IEEE Intelligent Systems, vol. 22, no. 3, pp. 2–3, 2007. (journal)
- [2] B. Smith, "Chapter 11: Ontology" *Blackwell Guide to the Philosophy of Computing and Information, L. Floridi* (ed.), pp. 155–166, 2003. (book chapter)
- [3] Gruber, T.R., 1993. A translation approach to portable ontology specifications. *Knowledge acquisition*, 5(2), pp.199-220.
- [4] Gómez-Pérez, A., Fernández-López, M. and Corcho, O., 2006. Ontological Engineering: with examples from the areas of Knowledge Management, e-Commerce and the Semantic Web. Springer Science & Business Media.
- [5] Gandon, F.L., 2003, October. Combining reactive and deliberative agents for complete ecosystems in infospheres. In *IEEE/WIC International Conference on Intelligent Agent Technology*, 2003. IAT 2003. (pp. 297-303). IEEE.
- [6] Nowak, M.A., Plotkin, J.B. and Krakauer, D.C., 1999. The evolutionary language game. *Journal of theoretical biology*, 200(2), pp.147-162.
- [7] Sarkadi, Ş., Rutherford, A., McBurney, P., Parsons, S. and Rahwan, I., 2021. The evolution of deception. *Royal Society open science*, 8(9), p.201032.
- [8] Sigmund, K., De Silva, H., Traulsen, A. and Hauert, C., 2010. Social learning promotes institutions for governing the commons. *Nature*, 466(7308), pp.861-863.
- [9] Čače, I. and Bryson, J.J., 2007. Agent based modelling of communication costs: Why information can be free. In Emergence of communication and language (pp. 305-321). Springer, London.
- [10] Bourahla, Y., Atencia, M. and Euzenat, J., 2021, May. Knowledge Improvement and Diversity under Interaction-Driven Adaptation of Learned Ontologies. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems* (pp. 242-250).
- [11] Zhi, N., Payne, T.R., Krysta, P. and Li, M., 2019, October. Truthful mechanisms for multi agent self-interested correspondence selection. In *International Semantic Web Conference* (pp. 733-750). Springer, Cham.

A Supplementary Material - Results & Data Analysis

Payoff in Non-Voluntary Model	Payoff in Voluntary Model
$O = (M-1) \left(c \times \theta \times r - c \right) - \epsilon \frac{M-1}{N-1} $ (2)	$O^* = P_{\sigma} \times \sigma + (1 - P_{\sigma})(N - N_L - 1)$ (c \times \theta \times r - c) - \epsilon \frac{M - 1}{N - 1} (3)
$CL = (M-1)(c \times r - c) - \beta^{N} \frac{M-1}{N-1} $ (4)	$CL^* = P_{\sigma} \times \sigma + (1 - P_{\sigma})(N - N_L - 1)$ $(c \times r - c) - \beta^N \frac{M - 1}{N - 1} $ (5)
$CL_{2} = (M-1)(c \times r - c) - \beta \times N_{CL_{2}} \frac{M-1}{N-1} $ (6)	$CL_2^* = P_\sigma \times \sigma + (1 - P_\sigma)(N - N_L - 1)$ $(c \times r - c) - \beta \times N_{CL_2} \frac{M - 1}{N - 1} $ (7)
$TR = (M-1)(c \times r - c) - \alpha \times c\frac{M-1}{N-1} $ (8)	$TR^* = P_{\sigma} \times \sigma + (1 - P_{\sigma})(N - N_L - 1)$ $(c \times r - c) - \alpha \times c \frac{M - 1}{N - 1} $ (9)
$TR_{2} = (M-1)(c \times r - c) - (1-\theta) \times \alpha \times c \frac{M-1}{N-1} $ (10)	$TR_2^* = P_{\sigma} \times \sigma + (1 - P_{\sigma})(N - N_L - 1)$ (c × r - c) - (1 - θ) × α × c $\frac{M - 1}{N - 1}$ (11)
$H = (M-1) (c \times r - c)$ $- \left(\alpha \times c \times (1-\theta) + \theta \times \beta^{(N_{CL}+N_H)}\right) \frac{M-1}{N-1} $ (12)	$H^* = P_{\sigma} \times \sigma + (1 - P_{\sigma})(N - N_L - 1) (c \times r - c)$ $- \left(\alpha \times c \times (1 - \theta) + \theta \times \beta^{(N_{CL} + N_H)}\right) \frac{M - 1}{N - 1} $ (13)
$H_2 = (M-1) (c \times r - c) - (\alpha \times c \times (1-\theta) + \theta \times \beta \times (N_{CL_2} + N_{H_2}))) \frac{M-1}{N-1} $ (14)	$H_2^* = P_\sigma \times \sigma + (1 - P_\sigma)(N - N_L - 1) (c \times r - c)$ $- (\alpha \times c \times (1 - \theta) + \theta \times \beta \times (N_{CL_2} + N_{H_2})) \frac{M - 1}{N - 1} $ (15)

Table 4: Different payoffs in non-voluntary and voluntary models combining costs and payouts from Table 2 and parameters defined in Table 1. On the left-hand side of the table we can find the payoffs for Non-Voluntary KSGs, while on the right-hand side we find the payoffs for the *Voluntary* KSGs.

Table 5: Means, standard deviations, t-test, and One-way ANOVA results for KSG1 in Tab. 5a vs KSG2 in Tab. 5b. Significant results where p < 0.01 are represented by *.

	CL	TR	Н	0
mn, sd	.0; .0	.96; .004	.03; .0	.0; .005
CL	N/A	-2044.75*	-658.75*	-0.97
TR	N/A	N/A	1949.60*	1394.90*
H	N/A	N/A	N/A	74.88

(a) One-way ANOVA result (statistic=1884732.98, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=862700.36, pvalue=0.0). T-tests and standard deviations in-table.

	CL_2	TR_2	H_2	0
mn, sd	.079; .002	.919; .005	.0; .0	.0; .007
CL_2	N/A	-146.56*	298.07*	95.57*
TR_2	N/A	N/A	1800.62*	990.51*
H_2	N/A	N/A	N/A	-1.01



Figure 2: Population dynamics for non-voluntary interactions in KSG1 fig.2a vs KSG2 fig. 2b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for KSG1 in fig. 2c and for KSG2 in fig. 2d.

Table 6: Means, standard deviations, t-test, and One-way ANOVA results for KSG1-V in Tab. 6a vs KSG2-V in Tab. 6b. Significant results where p < 0.01 are represented by *.

	CL	TR	Н	0	L
mn, sd	.0; .0	.95; .09	.04; .0	.0; .09	.0; .0
CL	N/A	-98.86*	-112.39*	-1.04	-0.09
TR	N/A	N/A	94.76*	70.42*	98.86*
Н	N/A	N/A	N/A	3.14*	112.39*
0	N/A	N/A	N/A	N/A	1.04

(a) One-way ANOVA result (statistic=4937.86, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=22407.57, pvalue=0.0). T-tests and standard deviations in-table.

	CL_2	TR_2	H_2	0	L
mn, sd	.038; .003	.95; .04	.0; .03	.0; .01	.0;.0
CL_2	N/A	-191.44*	8.66*	31.17*	118.86*
TR_2	N/A	N/A	153.43*	194.12*	199.88*
H_2	N/A	N/A	N/A	0.68	0.97
0	N/A	N/A	N/A	N/A	1.001



Figure 3: Population dynamics for voluntary interactions in KSG1-V fig.3a vs KSG2-V fig. 3b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for KSG1-V in fig. 3c and for KSG2-V in fig. 3d.

Table 7: Means, standard deviations, t-test, and One-way ANOVA results for KSG* in Tab. 7a vs KSG*-V in Tab. 7a. Significant results where p < 0.01 are represented by *.

	CL	CL_2	TR	TR_2	0
mn, sd	.0; .0	.038; .003	.0; .0	.96; .005	.0;.0
CL	N/A	-101.16*	-0.35	-1733.47*	-1.01
CL_2	N/A	N/A	101.15*	-1374.54*	37.28*
TR	N/A	N/A	N/A	-1733.44*	-1.01
TR_2	N/A	N/A	N/A	N/A	887.98*

(a) One-way ANOVA result (statistic=690368.13, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=14239.41, pvalue=0.0). T-tests and standard deviations in-table.

	CL	CL_2	TR	TR_2	0	L
mn, sd	.0;.0	.038;.002	.0; .006	.95;.05	.0; .05	.0; .0
CL	N/A	-139.62*	-1.01	-165.84*	-1.0	-0.51
CL_2	N/A	N/A	55.99*	-158.98*	6.08*	139.59*
TR	N/A	N/A	N/A	-164.79*	-0.87	1.00
TR_2	N/A	N/A	N/A	N/A	120.11*	165.84*
0	N/A	N/A	N/A	N/A	N/A	0.99



Figure 4: Population dynamics for interactions in KSG* fig. 4a vs KSG*-V fig.4b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for KSG* in fig. 4c and for KSG*-V in fig. 4d.

Table 8: Means, standard deviations, t-test, and One-way ANOVA results for PKSG1 in Tab. 8a vs PKSG2 in Tab. 8b. Significant results where p < 0.01 are represented by *.

	CL	TR	Н	0
mn, sd	.0; .0	.96; .015	.038; .0	.0; .015
CL	N/A	-624.11*	-483.98*	-1.0
TR	N/A	N/A	598.38*	431.0*
Н	N/A	N/A	N/A	22.88*

(a) One-way ANOVA result (statistic=181270.76, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=118870.68, pvalue=0.0). T-tests and standard deviations in-table.

	CL_2	TR_2	H_2	0
mn, sd	.08; .015	.918; .02	.0; .0	.0; .005
CL_2	N/A	-330.52*	54.12*	50.37*
TR_2	N/A	N/A	448.5*	432.58*
H_2	N/A	N/A	N/A	-0.98



Figure 5: Population dynamics for non-voluntary interactions in PKSG1 fig.5a vs PKSG2 fig. 5b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for PKSG1 in fig. 5c and for PKSG2 in fig. 5d.

Table 9: Means, standard deviations, t-test, and One-way ANOVA results for PKSG1-V in Tab. 9a vs PKSG2-V in Tab. 9b. Significant results where p < 0.01 are represented by *.

	CL	TR	Н	0	L
mn, sd	.0; .0	.94; .12	.04; .009	.017; .12	.0; .0
CL	N/A	-76.38*	-41.78*	-1.4	-1.48
TR	N/A	N/A	73.02*	53.46*	76.38*
Н	N/A	N/A	N/A	1.8	41.75*
0	N/A	N/A	N/A	N/A	1.39

(a) One-way ANOVA result (statistic=2873.48, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=28446.99, pvalue=0.0). T-tests and standard deviations in-table.

	CL_2	TR_2	H_2	0	L
mn, sd	.039; .005	.96; .04	.0; .03	.0; .01	.0;.0
CL_2	N/A	-205.61*	10.5*	51.3*	70.72*
TR_2	N/A	N/A	171.4*	214.53*	213.07*
H_2	N/A	N/A	N/A	0.23	0.99
0	N/A	N/A	N/A	N/A	0.98



Figure 6: Population dynamics for voluntary interactions in PKSG1-V fig.6a vs PKSG2-V fig. 6b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for PKSG1-V in fig. 6c and for PKSG2-V in fig. 6d.

Table 10: Means, standard deviations, t-test, and One-way ANOVA results for PKSG* in Tab. 10a vs PKSG*-V in Tab. 10a. Significant results where p < 0.01 are represented by *.

	CL	CL_2	TR	TR_2	0
mn, sd	.0; .0	.038; .0	.0; .007	.96; .007	.0;.007
CL	N/A	-719.52*	0.25	-1297.13*	-0.97
CL_2	N/A	N/A	718.95*	-1242*	48.54*
TR	N/A	N/A	N/A	-1297*	-0.97
TR_2	N/A	N/A	N/A	N/A	893.06*

(a) One-way ANOVA result (statistic=781654.67, pvalue=0.0). T-tests and standard deviations in-table.

(b) One-way ANOVA result (statistic=343995.38, pvalue=0.0). T-tests and standard deviations in-table.

	CL	CL_2	TR	TR_2	0	L
mn, sd	.0;.0	.038;.0	.0; .004	.96;.01	.0; .0	.0; .0
CL	N/A	-522.85*	-0.97	-755.8*	-1	0.35
CL_2	N/A	N/A	91.9*	-724.16*	40.78*	523.76*
TR	N/A	N/A	N/A	-718.96*	-0.51	0.98
TR_2	N/A	N/A	N/A	N/A	610.86*	755.8*
0	N/A	N/A	N/A	N/A	N/A	1



Figure 7: Population dynamics for interactions in PKSG* fig. 7a vs PKSG*-V fig.7b. Long-run average frequencies for a sample of 100 individual-based runs are also reported for PKSG* in fig. 7c and for PKSG*-V in fig. 7d



(b) Exploration Rate for range s = [0.001, 1] - PKSG*-V

Figure 8: 8a Effect of social learning s and 8b exploration rate μ in PKSG*-V when an individual run starts from all agents using own language O.





(b) Long-run avg. frequencies for PKSG*-V when population distribution resets to own language O. Sample size = 100.

Figure 9: 9a dynamics of an individual-based simulation run Dynamics and 9b avg. of 100 long-run avg. frequencies for PKSG*-V when an individual run starts from all agents using own language *O*.

B Supplementary Material - Sensitivity Analysis

Next, we perform a sensitivity analysis for PKSG*-V, as it is the closest the most representative KSG given a fully open heterogrenous and self-organising multi-agent system that we have modeled where Common Language and Translation compete as knowledge sharing strategies. In summary, we show here the following significant effects:

- The particular case of when population size N drives the formation of silos if population distribution resets to O after every independent run. We also show here the effects of s and μ for a typical simulation run.
- For a controlled sensitivity analysis we also set up a simulation that starts and resets to keep an equal distribution of strategies in the population of N agents after every run. Hence every strategy will start with a distribution of N/6 agents. We show the avg. of the long-run avg. frequencies, and the effects of social learning s, compatibility, and strategy costs.

B.1



Figure 10: Effect of population size $N \in [100, 1000]$ in PKSG*-V when population distribution resets to own language O after every run. Sample size 100 individual runs. Lines represents means and shaded areas represend ± 1 SD from the mean.





(a) Long-run avg. frequencies for PKSG*-V when population distribution starts from equal distribution and resets to equal distribution at every run.





Figure 11: 11a Long-run avg. frequencies and effects of 11b own language cost, 11c common language cost, and 11d translation cost, $\theta, \beta, \alpha \in [10^{-3}, 10^3]$ in PKSG*-V when every run starts from an equal distribution of strategies. Sample size 100 individual runs. Lines represents means and shaded areas represend ±1 SD from the mean.



Figure 12: Effect of social learning $s \in [10^{-4}, 10^3]$ on 12a TR_2 and 12b CL_2 in PKSG*-V when every run starts from an equal distribution of strategies. Sample size 100 individual runs. Lines represents means and shaded areas represend ± 1 SD from the mean.



Figure 13: Effect of agent compatibility $\theta \in [0, 1]$ in PKSG*-V when every run starts from an equal distribution of strategies. Sample size 100 individual runs. Lines represents means and shaded areas represend ± 1 SD from the mean.