Minimal Change in Evolving Multi-Context Systems

Ricardo Gonçalves, Matthias Knorr, and João Leite

NOVA LINCS, Departamento de Informática, Faculdade Ciências e Tecnologia, Universidade Nova de Lisboa

Abstract. In open environments, agents need to reason with knowledge from various sources, possibly represented in different languages. The framework of *Multi-Context Systems* (MCSs) offers an expressive, yet flexible, solution since it allows for the integration of knowledge from different heterogeneous sources in an effective and modular way. However, MCSs are essentially static as they were not designed for dynamic scenarios. The recently introduced *evolving Multi-Context Systems* (eMCSs) extend MCSs by also allowing the system to both react to, and reason in the presence of dynamic observations, and evolve by incorporating new knowledge, thus making it even more adequate in Multi-Agent Systems characterised by their dynamic and open nature.

In dynamic scenarios which admit several possible alternative evolutions, the notion of *minimal change* has always played a crucial role in determining the most plausible choice. However, different KR formalisms – as combined within eMCSs – may require different notions of minimal change, making their study and their interplay a relevant highly non-trivial problem. In this paper, we study the notion of minimal change in eMCSs, by presenting and discussing alternative minimal change criteria.

1 Introduction

Open and dynamic environments create new challenges for knowledge representation languages for agent systems. Instead of having to deal with a single static knowledge base, each agent has to deal with multiple sources of distributed knowledge possibly written in different languages. These sources of knowledge include the large number of available ontologies and rule sets, as well as the norms and policies published by *institutions*, the information communicated by other agents, to name only a few.

The need to incorporate in agent-oriented programming languages the ability to represent and reason with heterogeneous distributed knowledge sources, and the flow of information between them, has been pointed out in [12, 25, 30, 1], although a general adequate practical solution is still not available.

Recent literature in knowledge representation and reasoning contains several proposals to combine heterogeneous knowledge bases, one of which – Multi-Context Systems (MCSs) [6, 17, 32] – has attracted particular attention because it provides an elegant solution by considering each source of knowledge as a module and then providing means to model the interaction between these modules. More specifically, an MCS consists of a set of contexts, each of which is a knowledge base in some KR formalism, such that each context can access information from the other contexts using so-called bridge rules. Such non-monotonic bridge rules add their heads to the context's knowledge base provided the queries (to other contexts) in their bodies are successful. Managed Multi-Context Systems (mMCSs) [7] extend MCSs by allowing operations, other than simple addition, to be expressed in the heads of bridge rules, thus allowing mMCSs to properly deal with the problem of consistency management within contexts. MCSs have gained some attention by agent developers [3, 13, 33].

Whereas mMCSs are quite general and flexible to address the problem of integration of different KR formalisms, they are essentially static in the sense that the contexts do not evolve to incorporate the changes in dynamic scenarios. In such scenarios, new knowledge and information is dynamically produced, often from several different sources – e.g., a stream of raw data produced by some sensors, new ontological axioms written by some user, newly found exceptions to some general rule, observations, etc.

Evolving Multi-Context Systems (eMCSs) [19] inherit from mMCSs the ability to integrate and manage knowledge represented in heterogeneous KR formalisms, and at the same time are able to react to dynamic observations, and evolve by incorporating knowledge. The semantics of eMCSs is based on the stable model semantics, and allows alternative models for a given evolution, in the same way as answer sets represent alternative solutions to a given ASP program.

One of the main principles of belief revision is minimal change, which in case of eMCSs means that information should be maintained by inertia unless it is required to change. In dynamic scenarios where systems can have alternative evolutions, it is thus desirable to have some minimal change criteria to be able to compare possible alternatives. This problem is particularly interesting and non-trivial in dynamic frameworks based on MCSs, because of the heterogeneity of KR frameworks that may exist in an MCS – each of which may require different notions of minimal change –, and also because the evolution of such systems is based not only on the semantics, but also on the evolution of the knowledge base of each context.

In this paper, we study minimal change in eMCSs, by presenting different minimal change criteria to be applied to the possible evolving equilibria of an eMCS, and by discussing the relation between them.

The remainder of this paper is as follows. We introduce the framework of eMCSs in Sect. 2. Then, we present and study some minimal change criteria in eMCSs in Sect. 3, and conclude with a discussion of related work and possible future directions in Sect. 4.

2 Evolving Multi-Context Systems

In this section, we revisit evolving Multi-Context Systems as introduced in [19], which generalize mMCSs [7] to dynamic scenarios in which contexts are enabled to react to external observations and evolve.

An evolving multi-context system (eMCS) consists of a collection of components, each of which contains knowledge represented in some *logic*, defined as a triple $L = \langle \mathbf{KB}, \mathbf{BS}, \mathbf{ACC} \rangle$ where $\mathbf{KB}$ is the set of well-formed knowledge bases of L, $\mathbf{BS}$ the set of possible belief sets, and $\mathbf{ACC} : \mathbf{KB} \to 2^{\mathbf{BS}}$ a function describing the semantics of L by assigning to each knowledge base a set of acceptable belief sets. We assume that each element of $\mathbf{KB}$ and $\mathbf{BS}$ is a set, and define $F = \{s : s \in kb \wedge kb \in \mathbf{KB}\}$.

In addition to the knowledge base in each component, *bridge rules* are used to interconnect the components, specifying what operations to perform on its knowledge base given certain beliefs held in the components of the eMCS. For that purpose, each component of an eMCS is associated with a *management base*, which is a set of operations that can be applied to the possible knowledge bases of that component. Given a management base OP and a logic L, let $OF = \{op(s) : op \in OP \wedge s \in F\}$ be the *set of operational formulas* over OP and L. Each component of an eMCS gives semantics to operations in its management base using a *management function* over a logic L and a management base OP, $mng : 2^{OF} \times \mathbf{KB} \to (2^{\mathbf{KB}} \setminus \{\emptyset\})$, i.e., $mng(op, kb)$ is the (non-empty) set of knowledge bases that result from applying the operations in op to the knowledge base kb. We assume that $mng(\emptyset, kb) = \{kb\}$.

In an eMCS some contexts are assumed to be *observation contexts* whose knowledge bases will be constantly changing over time according to the observations made, similar, e.g., to streams of data from sensors.[1] The changing observations will then affect the other contexts by means of the bridge rules. As we will see, such effect can either be instantaneous and temporary, i.e., limited to the current time instant, similar to (static) mMCSs, where the body of a bridge rule is evaluated in a state that already includes the effects of the operation in its head, or persistent, but only affecting the next time instant. To achieve the latter, the operational language is extended with a unary meta-operation $next$ that can only be applied on top of operations. Given a management base OP and a logic L, we define eOF, the evolving operational language, as $eOF = OF \cup \{next(op(s)) : op(s) \in OF\}$.

The idea of observation contexts is that each such context has a language describing the set of possible observations of that context, along with its current observation. The elements of the language of the observation contexts can then be used in the body of bridge rules to allow contexts to access the observations. Formally, an *observation context* is a tuple $O = \langle \Pi_O, \pi \rangle$ where Π_O is the *observation language* of O and $\pi \subseteq \Pi_O$ is its *current observation*.

We can now define *evolving Multi-Context Systems* (eMCS).

Definition 1. *An eMCS is a sequence $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$, such that, for each $i \in \{1, \ldots, \ell\}$, $O_i = \langle \Pi_{O_i}, \pi_i \rangle$ is an* observation context*, and, for each $i \in \{1, \ldots, n\}$, C_i is an* evolving context *defined as $C_i = \langle L_i, kb_i, br_i, OP_i, mng_i \rangle$ where*
- $L_i = \langle \mathbf{KB}_i, \mathbf{BS}_i, \mathbf{ACC}_i \rangle$ *is a logic*
- $kb_i \in \mathbf{KB}_i$
- br_i *is a set of bridge rules of the form*

$$H(\sigma) \leftarrow a_1, \ldots, a_k, \mathbf{not}\ a_{k+1}, \ldots, \mathbf{not}\ a_n \tag{1}$$

  such that $H(\sigma) \in eOF_i$, and each a_i, $i \in \{1, \ldots, n\}$, is either of the form $(r : b)$ with $r \in \{1, \ldots, n\}$ and b a belief formula of L_r, or of the form $(r@b)$ with $r \in \{1, \ldots, \ell\}$ and $b \in \Pi_{O_r}$
- OP_i *is a management base*
- mng_i *is a management function over L_i and OP_i.*

[1] For simplicity of presentation, we consider discrete steps in time here.

Given an eMCS $M_e = \langle C_1, \dots, C_n, O_1, \dots, O_\ell \rangle$ we denote by $\mathbf{KB}_{M_e}$ the set of *knowledge base configurations for* M_e, i.e., $\mathbf{KB}_{M_e} = \{\langle k_1, \dots, k_n \rangle : k_i \in \mathbf{KB}_i$ for each $1 \leq i \leq n\}$. A *belief state* for $M_e = \langle C_1, \dots, C_n, O_1, \dots, O_\ell \rangle$ is a sequence $S = \langle S_1, \dots, S_n \rangle$ such that, for each $1 \leq i \leq n$, we have $S_i \in \mathbf{BS}_i$. We denote by $\mathbf{BS}_{M_e}$ the set of belief states for M_e.

An *instant observation* for M_e is a sequence $\mathcal{O} = \langle o_1, \dots, o_\ell \rangle$ such that, for each $1 \leq i \leq \ell$, we have that $o_i \subseteq \Pi_{O_i}$.

Given a belief state $S = \langle S_1, \dots, S_n \rangle$ for M_e and an instant observation $\mathcal{O} = \langle o_1, \dots, o_\ell \rangle$ for M_e, we define the satisfaction of bridge literals of the form $(r : b)$ as $S, \mathcal{O} \models (r : b)$ if $b \in S_r$ and $S, \mathcal{O} \models \mathbf{not}\,(r : b)$ if $b \notin S_r$. The satisfaction of bridge literal of the form $(r@b)$ depends on the current observations, i.e., we have that $S, \mathcal{O} \models (r@b)$ if $b \in o_r$ and $S \models \mathbf{not}\,(r@b)$ if $b \notin o_r$. For a set B of bridge literals, we have that $S, \mathcal{O} \models B$ if $S, \mathcal{O} \models L$ for every $L \in B$.

We say that a bridge rule σ of a context C_i is *applicable given a belief state* S *and an instant observation* $\mathcal{O}$ if its body is satisfied by S and $\mathcal{O}$, i.e., $S, \mathcal{O} \models B(\sigma)$. We denote by $app_i(S, \mathcal{O})$ the set of heads of bridges rules of the context C_i which are applicable given the belief state S and the instant observation $\mathcal{O}$. Recall that the heads of bridge rules in an eMCS may be of two types: those that contain $next$ and those that do not. The former are to be applied to the current knowledge base and not persist, whereas the latter are to be applied in the next time instant and persist. Therefore, we distinguish these two subsets.

Definition 2. *Let* $M_e = \langle C_1, \dots, C_n, O_1, \dots, O_\ell \rangle$ *be an eMCS,* S *a belief state for* M_e*, and* $\mathcal{O}$ *an instant observation for* M_e*. Then, for each* $1 \leq i \leq n$*, consider the following sets:*
- $app_i^{next}(S, \mathcal{O}) = \{op(s) : next(op(s)) \in app_i(S, \mathcal{O})\}$
- $app_i^{now}(S, \mathcal{O}) = \{op(s) : op(s) \in app_i(S, \mathcal{O})\}$

If we want an effect to be instantaneous and persistent, this can be achieved using two bridge rules with identical body, one with and one without $next$.

Similar to equilibria in mMCS, the (static) equilibrium is defined to incorporate instantaneous effects based on $app_i^{now}(S, \mathcal{O})$ alone.

Definition 3. *Let* $M_e = \langle C_1, \dots, C_n, O_1, \dots, O_\ell \rangle$ *be an eMCS, and* $\mathcal{O}$ *an instant observation for* M_e*. A belief state* $S = \langle S_1, \dots, S_n \rangle$ *for* M_e *is an* equilibrium of M_e given $\mathcal{O}$ *iff for each* $1 \leq i \leq n$*,* $S_i \in \mathbf{ACC}_i(kb)$ *for some* $kb \in mng_i(app_i^{now}(S, \mathcal{O}), kb_i)$*.*

To be able to assign meaning to an eMCS evolving over time, we introduce evolving belief states, which are sequences of belief states, each referring to a subsequent time instant.

Definition 4. *Let* $M_e = \langle C_1, \dots, C_n, O_1, \dots, O_\ell \rangle$ *be an eMCS. An* evolving belief state *of size* s *for* M_e *is a sequence* $S_e = \langle S^1, \dots, S^s \rangle$ *where each* S^j*,* $1 \leq j \leq s$*, is a belief state for* M_e*.*

To enable eMCSs to react to incoming observations and evolve, a sequence of observations (defined below) has to be processed. The idea is that the knowledge bases of the observation contexts O_i change according to that sequence.

Definition 5. *Let* $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$ *be an eMCS. A* sequence of observations *for* M_e *is a sequence* $Obs = \langle \mathcal{O}^1, \ldots, \mathcal{O}^m \rangle$*, such that, for each* $1 \leq j \leq m$*,* $\mathcal{O}^j = \langle o_1^j, \ldots, o_\ell^j \rangle$ *is an instant observation for* M_e*, i.e.,* $o_i^j \subseteq \Pi_{O_i}$ *for each* $1 \leq i \leq \ell$*.*

To be able to update the knowledge bases and the sets of bridge rules of the evolving contexts, we need the following notation. Given an evolving context C_i, and a knowledge base $k \in \mathbf{KB}_i$, we denote by $C_i[k]$ the evolving context in which kb_i is replaced by k, i.e., $C_i[k] = \langle L_i, k, br_i, OP_i, mng_i \rangle$. For an observation context O_i, given a set $\pi \subseteq \Pi_{O_i}$ of observations for O_i, we denote by $O_i[\pi]$ the observation context in which its current observation is replaced by π, i.e., $O_i[\pi] = \langle \Pi_{O_i}, \pi \rangle$. Given $K = \langle k_1, \ldots, k_n \rangle \in \mathbf{KB}_{M_e}$ a knowledge base configuration for M_e, we denote by $M_e[K]$ the eMCS $\langle C_1[k_1], \ldots, C_n[k_n], O_1, \ldots, O_\ell \rangle$.

We now define when certain evolving belief states are evolving equilibria of an eMCS M_e given a sequence of observations $Obs = \langle \mathcal{O}^1, \ldots, \mathcal{O}^m \rangle$ for M_e. The intuitive idea is that, given an evolving belief state $S_e = \langle S^1, \ldots, S^s \rangle$ for M_e, in order to check if S_e is an evolving equilibrium, we need to consider a sequence of eMCSs, $M^1, \ldots, M^s$ (each with ℓ observation contexts), representing a possible evolution of M_e according to the observations in Obs, such that each S^j is a (static) equilibrium of M^j. The current observation of each observation context O_i in M^j is exactly its corresponding element o_i^j in $\mathcal{O}^j$. For each evolving context C_i, its knowledge base in M^j is obtained from the one in M^{j-1} by applying the operations in $app_i^{next}(S^{j-1}, \mathcal{O}^{j-1})$.

Definition 6. *Let* $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$ *be an eMCS,* $S_e = \langle S^1, \ldots, S^s \rangle$ *an evolving belief state of size* s *for* M_e*, and* $Obs = \langle \mathcal{O}^1, \ldots, \mathcal{O}^m \rangle$ *an observation sequence for* M_e *such that* $m \geq s$*. Then,* S_e *is an* evolving equilibrium *of size* s *of* M_e *given* Obs *iff, for each* $1 \leq j \leq s$*, the belief state* S^j *is an equilibrium of* $M^j = \langle C_1[k_1^j], \ldots, C_n[k_n^j], O_1[o_1^j], \ldots, O_\ell[o_\ell^j] \rangle$ *where, for each* $1 \leq i \leq n$*,* k_i^j *is defined inductively as follows:*
- $k_i^1 = kb_i$
- $k_i^{j+1} \in mng_i(app^{next}(S^j, \mathcal{O}_i^j), k_i^j)$.

3 Minimal change

In this section, we discuss some alternatives for the notion of minimal change in eMCSs. What makes this problem interesting is that there are different parameters that we may want to minimize in a transition from one time instant to the next one. In the following discussion we focus on two we deem most relevant: the operations that can be applied to the knowledge bases, and the distance between consecutive belief states.

We start by studying minimal change at the level of the operations. In the following discussion we consider fixed an eMCS $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$.

Recall from the definition of evolving equilibrium that, in the transition between consecutive time instants, the knowledge base of each context C_i of M_e changes according to the operations in $app_i^{next}(S, \mathcal{O})$, and these depend on the belief state S and the instant observation $\mathcal{O}$. The first idea to compare elements of this set of operations is to, for a fixed instant observation $\mathcal{O}$, distinguish those equilibria of M_e which generate a minimal set of operations to be applied to the current knowledge bases to obtain the knowledge bases of the next time instant. Formally, given a knowledge base configuration $K \in \mathbf{KB}_{M_e}$ and an instant observation $\mathcal{O}$ for M_e, we can define the set:

$$MinEq(K, \mathcal{O}) = \{S : S \text{ is an equilibrium of } M_e[K] \text{ given } \mathcal{O} \text{ and there is no equilibrium } S' \text{ of } M_e[K] \text{ given } \mathcal{O} \text{ such that, for all } 1 \leq i \leq n, app_i^{next}(S', \mathcal{O}) \subset app_i^{next}(S, \mathcal{O})\}$$

This first idea of comparing equilibria based on inclusion of the sets of operations can, however, be too strict in most cases. Moreover, different operations usually have different costs,[2] and it may well be that, instead of minimizing based on set inclusion, we want to minimize the total cost of the operations to be applied. For that, we need to assume that each context has a cost function over the set of operations, i.e., $cost_i : OP_i \to \mathbb{N}$, where $cost_i(op)$ represents the cost of performing operation op.

Let S be a belief state for M_e and $\mathcal{O}$ an instant observation for M_e. Then, for each $1 \leq i \leq n$, we define the cost of the operations to be applied to obtain the knowledge base of the next time instant as:

$$Cost_i(S, \mathcal{O}) = \sum_{op(s) \in app_i^{next}(S, \mathcal{O})} cost_i(op)$$

Summing for all evolving contexts, we obtain the global cost of S given $\mathcal{O}$:

$$Cost(S, \mathcal{O}) = \sum_{i=1}^{n} Cost_i(S, \mathcal{O})$$

Now that we have defined a cost function over belief states, we can define a minimization function over possible equilibria of eMCS $M_e[K]$ for a fixed knowledge base configuration $K \in \mathbf{KB}_{M_e}$. Formally, given $\mathcal{O}$ an instant observation for M_e, we define the set of equilibria of $M_e[K]$ given $\mathcal{O}$ which minimize the global cost of the operations to be applied to obtain the knowledge base configuration of the next time instant as:

$$MinCost(K, \mathcal{O}) = \{S : S \text{ is an equilibrium of } M_e[K] \text{ given } \mathcal{O} \text{ and there is no equilibrium } S' \text{ of } M_e[K] \text{ given } \mathcal{O} \text{ such that } Cost(S', \mathcal{O}) < Cost(S, \mathcal{O})\}$$

Note that, instead of using a global cost, we could have also considered a more fine-grained criterion by comparing costs for each context individually, and define some order based on these comparisons. Also note that the particular case of taking $cost_i(op) = 1$ for every $i \in \{1, \ldots, n\}$ and every $op \in OP_i$, captures the scenario of minimizing the total number of operations to be applied.

The function $MinCost$ allows for the choice of those equilibria that are minimal with respect to the operations to be performed to the current knowledge base configuration in order to obtain the knowledge base configuration of the next time instant. Still, for each choice of an equilibrium S, we have to deal with the existence of several alternatives in the set $mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$. Our aim now is to discuss how we can apply some notion of minimal change that allows us to compare the elements in $mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$. The intuitive idea is to compare the distance between the current equilibria and the possible equilibria resulting from the elements in

[2] We use the notion of cost in an abstract sense, i.e., depending on the context, it may refer to, e.g., the computational cost of the operation, or its economic cost.

$mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$. Of course, given the possible heterogeneity of contexts in an eMCS, we cannot assume a global notion of distance between belief sets. Therefore, we assume that each evolving context has its own distance function between its beliefs sets. Formally, for each $1 \leq i \leq n$, we assume the existence of a distance function d_i, i.e., $d_i : \mathbf{BS}_i \times \mathbf{BS}_i \to \mathbb{R}$ satisfying for all $S_1, S_2, S_3 \in \mathbf{BS}_i$:

1. $d_i(S_1, S_2) \geq 0$
2. $d_i(S_1, S_2) = 0$ iff $S_1 = S_2$
3. $d_i(S_1, S_2) = d_i(S_2, S_1)$
4. $d_i(S_1, S_3) \leq d_i(S_1, S_2) + d_i(S_2, S_3)$

There are some alternatives to extend the distance function of each context to a distance function between belief states. In the following we present two natural choices. One is to consider the maximal distance between belief sets of each context. The other is to consider the average of distances between belief sets of each context. Formally, given S^1 and S^2 belief states of M_e we define two functions $\overline{d}_{\text{max}} : \mathbf{BS}_{M_e} \times \mathbf{BS}_{M_e} \to \mathbb{R}$ and $\overline{d}_{\text{avg}} : \mathbf{BS}_{M_e} \times \mathbf{BS}_{M_e} \to \mathbb{R}$ as follows:

$$\overline{d}_{\text{max}}(S^1, S^2) = Max\{d_i(S_i^1, S_i^2) \mid 1 \leq i \leq n\}$$
$$\overline{d}_{\text{avg}}(S^1, S^2) = \frac{\sum_{i=1}^{n} d_i(S_i^1, S_i^2)}{n}$$

We can prove that $\overline{d}_{\text{max}}$ and $\overline{d}_{\text{avg}}$ are distance functions between belief states.

Proposition 1. *The functions $\overline{d}_{max}$ and $\overline{d}_{avg}$ defined above are both distance functions, i.e., satisfy the axioms 1) - 4).*

We now study how we can use one of these distance functions between belief states to compare the possible alternatives in the sets $mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$, for each $1 \leq i \leq n$. Recall that the intuitive idea is to minimize the distance between the current belief state S and the possible equilibria that each element of $mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$ can give rise to. We explore here two options, which differ on whether the minimization is global or local. The idea of global minimization is to choose only those knowledge base configurations $\langle k_1, \ldots, k_n \rangle \in \mathbf{KB}_{M_e}$ with $k_i \in mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$, which guarantee minimal distance between the original belief state S and the possible equilibria of the obtained eMCS. The idea of local minimization is to consider all possible tuples $\langle k_1, \ldots, k_n \rangle$ with $k_i \in mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$, and only apply minimization for each such choice, i.e., for each such knowledge base configuration we only allow equilibria with minimal distance from the original belief state.

We first consider the case of pruning those tuples $\langle k_1, \ldots, k_n \rangle$ such that $k_i \in mng_i(app_i^{next}(S, \mathcal{O}), kb_i)$, which do not guarantee minimal change with respect to the original belief state. We start by defining an auxiliary function. Let S be a belief state for M_e, $K = \langle k_1, \ldots, k_n \rangle \in \mathbf{KB}_{M_e}$ a knowledge base configuration for M_e, and $\mathcal{O} = \langle o_1, \ldots, o_\ell \rangle$ an instant observation for M_e. Then we define the set of knowledge base configurations that are obtained from K given the belief state S and the instant observation $\mathcal{O}$ as:

$$NextKB(S, \mathcal{O}, \langle k_1, \ldots, k_n \rangle) = \{\langle k'_1, \ldots, k'_n \rangle \in \mathbf{KB}_{M_e} : \text{ for each } 1 \leq i \leq n \text{ we have that } k'_i \in mng_i(app_i^{next}(S, \mathcal{O}), k_i)\}$$

For each choice $\overline{d}$ of a distance function between belief states, we define the set of knowledge base configurations that minimize the distance to the original belief state. Let S be a belief state for M_e, $K = \langle k_1, \ldots, k_n \rangle \in \mathbf{KB}_{M_e}$ a knowledge base configuration for M_e, and $\mathcal{O}^j$ and $\mathcal{O}^{j+1}$ instant observations for M_e.

$$MinNext(S, \mathcal{O}^j, \mathcal{O}^{j+1}, K) = \{(S', K') : K' \in NextKB(S, \mathcal{O}^j, K) \text{ and } S' \in MinCost(M_e[K'], \mathcal{O}^{j+1}) \text{ s.t. there is no } K'' \in NextKB(S, \mathcal{O}^j, K) \text{ and no } S'' \in MinCost(M_e[K''], \mathcal{O}^{j+1}) \text{ with } \overline{d}(S, S'') < \overline{d}(S, S')\}.$$

Note that $MinNext$ applies minimization over all possible equilibria resulting from every element of $NextKB(S, \mathcal{O}^j, K)$. Using $MinNext$, we can now define a minimal change criterion to be applied to evolving equilibria of M_e.

Definition 7. *Let $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$ be an eMCS, Obs $= \langle \mathcal{O}^1, \ldots, O^m \rangle$ an observation sequence for M_e, and let $S_e = \langle S^1, \ldots, S^s \rangle$ be an evolving equilibrium of M_e given Obs. We assume that $\langle K^1, \ldots, K^s \rangle$, with $K^j = \langle k_1^j, \ldots, k_n^j \rangle$, is the sequence of knowledge base configurations associated with S_e as in Definition 6. Then, S_e satisfies the* strong minimal change criterion *for M_e given Obs if, for each $1 \leq j \leq s$, the following conditions are satisfied:*

- $S^j \in MinCost(M_e[K^j], \mathcal{O}^j)$
- $(S^{j+1}, K^{j+1}) \in MinNext(S^j, \mathcal{O}^j, \mathcal{O}^{j+1}, K^j)$

We call this minimal change criterion the *strong* minimal change criterion because it applies minimization over all possible equilibria resulting from every possible knowledge base configuration in $NextKB(S, \mathcal{O}^j, K)$.

The following proposition states the desirable property that the existence of an equilibrium guarantees the existence of an equilibrium satisfying the strong minimal change criterion. We should note that this is not a trivial statement since we are combining minimization of two different elements: the cost of the operations and the distance between belief states. This proposition in fact follows from their careful combination in the definition of $MinNext$.

Proposition 2. *Let Obs $= \langle \mathcal{O}^1, \ldots, O^m \rangle$ be an observation sequence for M_e. If M_e has an evolving equilibrium of size s given Obs, then at least one evolving equilibrium of size s given Obs satisfies the strong minimal change criterion.*

Note that in the definition of the strong minimal change criterion, the knowledge base configurations $K \in NextKB(S^j, \mathcal{O}^j, K^j)$, for which the corresponding possible equilibria are not at a minimal distance from S^j, are not considered. However, there could be situations in which this minimization criterion is too strong. For example, it may well be that all possible knowledge base configurations in $NextKB(S^j, \mathcal{O}^j, K^j)$ are important, and we do not want to disregard any of them. In that case, we can relax the minimization condition by applying minimization individually for each knowledge base configuration in $NextKB(S^j, \mathcal{O}^j, K^j)$. The idea is that, for each fixed $K \in NextKB(S^j, \mathcal{O}^j, K^j)$ we choose only those equilibria of $M_e[K]$ which minimize the distance to S^j.

Formally, let S be a belief state for M_e, $K \in \mathbf{KB}_{M_e}$ a knowledge base configuration for M_e, and $\mathcal{O}$ an instant observation for M_e. For each distance function $\overline{d}$ between belief states, we can define the following set:

$$\begin{aligned} MinDist(S, \mathcal{O}, K) = \{ & S' : S' \in MinCost(M_e[K], \mathcal{O}) \text{ and} \\ & \text{there is no } S'' \in MinCost(M_e[K], \mathcal{O}) \\ & \text{such that } \overline{d}(S, S'') < \overline{d}(S, S') \} \end{aligned}$$

Using this more relaxed notion of minimization we can define an alternative weaker minimal change criterion to be applied to evolving equilibria of an eMCS.

Definition 8. *Let $M_e = \langle C_1, \ldots, C_n, O_1, \ldots, O_\ell \rangle$ be an eMCS, Obs $= \langle \mathcal{O}^1, \ldots, O^m \rangle$ an observation sequence for M_e, and $S_e = \langle S^1, \ldots, S^s \rangle$ an evolving equilibrium of M_e given Obs. We assume that $\langle K^1, \ldots, K^s \rangle$, with $K^j = \langle k_1^j, \ldots, k_n^j \rangle$, is the sequence of knowledge base configurations associated with S_e as in Definition 6. Then, S_e satisfies the* weak minimal change criterion *of M_e given Obs, if for each $1 \leq j \leq s$ the following conditions are satisfied:*

- $S^j \in MinCost(M_e[K^j], \mathcal{O}^j)$
- $S^{j+1} \in MinDist(S^j, K^{j+1}, \mathcal{O}^{j+1})$

We can now prove that the existence of an evolving equilibrium implies the existence of an equilibrium satisfying the weak minimal change criterion. Again note that the careful combination of the two minimizations – cost and distance – in the definition of $MinDist$ is fundamental to obtain the following result.

Proposition 3. *Let Obs $= \langle \mathcal{O}^1, \ldots, O^m \rangle$ be an observation sequence for M_e. If M_e has an evolving equilibrium of size s given Obs, then at least one evolving equilibrium of size s of M_e given Obs satisfies the weak minimal change criterion.*

We can now prove that the strong minimal change criterion is, in fact, stronger than the weak minimal change criterion.

Proposition 4. *Let M_e be an eMCS, Obs $= \langle \mathcal{O}^1, \ldots, O^m \rangle$ an observation sequence for M_e, and $S_e = \langle S^1, \ldots, S^s \rangle$ an evolving equilibrium of M_e given Obs. If S_e satisfies the strong minimal change criterion of M_e given Obs, then S_e satisfies the weak minimal change criterion of M_e given Obs.*

4 Related and Future Work

In this paper we have studied the notion of minimal change in the context of the dynamic framework of eMCSs [19]. We have presented and discussed some alternative definitions of minimal change criteria for evolving equilibria of an eMCS.

Closely related to eMCSs is the framework of reactive Multi-Context Systems (rMCSs) [5, 14, 8] inasmuch as both aim at extending mMCSs to cope with dynamic observations. The key difference between them is that the operator $next$ of eMCSs allows for a clear separation between persistent and non-persistent effects, and the specification of transitions based on the current state.

Another framework closely related to eMCSs is that of evolving logic programs EVOLP [2] which deals with updates of generalized logic programs, and the two frameworks of reactive ASP, one implemented as a solver *oclingo* [15] and one described in [5]. Whereas EVOLP employs an update predicate that is similar in spirit to the $next$ predicate of eMCSs, it does not deal with heterogeneous knowledge, neither do both versions of Reactive ASP. Moreover, no notion of minimal change is studied for these frameworks.

This work raises several interesting paths for future research. Immediate future work includes the study of more global approaches to the minimization of costs of operations, namely by considering the global cost of an evolving equilibrium, instead of minimizing costs at each time instant. A topic worth investigating is how to perform AGM-style belief revision at the (semantic) level of the equilibria, as in Wang et al [39], though necessarily different since knowledge is not incorporated in the contexts. Also interesting is to study a paraconsistent version of eMCSs, grounded on the work in [24] on paraconsistent semantics for hybrid knowledge bases. Another important issue open for future work is a more fine-grained characterization of updating bridge rules (and knowledge bases) as studied in [18] in light of the encountered difficulties when updating rules [34, 35, 37] and the combination of updates over various formalisms [35, 36]. Also, as already outlined in [21, 27], we can consider generalized notions of minimal and grounded equilibria [6] for eMCSs to avoid, e.g., self-supporting cycles introduced by bridge rules, or the use of preferences to deal with several evolving equilibria an eMCS can have for the same observation sequence. Also interesting is to apply the ideas in this paper to study the dynamics of frameworks closely related to MCSs, such as those in [28, 20, 26, 23].

Finally, and in line with the very motivation set out in the introduction, we believe that the research in MCSs – including eMCSs with the different notions of minimal change – provides a blue-print on how to represent and reason with heterogeneous dynamic knowledge bases which could (should) be used by developers of practical agent-oriented programming languages, such as JASON [4], 2APL [11], or GOAL [22], in their quest for providing users and programmers with greater expressiveness and flexibility in terms of the knowledge representation and reasoning facilities provided by such languages. To this end, an application scenario that could provide interesting and rich examples would be that of norm-aware multi-agent systems [10, 29, 9, 16, 38, 31].

Acknowledgments We would like to thank the referees for their comments, which helped improve this paper. R. Gonçalves, M. Knorr and J. Leite were partially supported by FCT under project ERRO (PTDC/EIA-CCO/121823/2010) and under strategic project NOVA LINCS (PEst/UID/CEC/04516/2013). R. Gonçalves was partially supported by FCT grant SFRH/BPD/100906/2014 and M. Knorr was partially supported by FCT grant SFRH/BPD/86970/2012.

References

1. Alberti, M., Knorr, M., Gomes, A.S., Leite, J., Gonçalves, R., Slota, M.: Normative systems require hybrid knowledge bases. In: van der Hoek, W., Padgham, L., Conitzer, V., Winikoff, M. (eds.) Procs. of AAMAS. pp. 1425–1426. IFAAMAS (2012)
2. Alferes, J., Brogi, A., Leite, J., Pereira, L.: Evolving logic programs. In: Flesca, S., Greco, S., Leone, N., Ianni, G. (eds.) Procs. of JELIA. LNCS, vol. 2424, pp. 50–61. Springer (2002)
3. Benerecetti, M., Giunchiglia, F., Serafini, L.: Model checking multiagent systems. J. Log. Comput. 8(3), 401–423 (1998)
4. Bordini, R.H., Hübner, J.F., Wooldridge, M.: Programming Multi-Agent Systems in AgentSpeak Using Jason (Wiley Series in Agent Technology). John Wiley & Sons (2007)
5. Brewka, G.: Towards reactive multi-context systems. In: Cabalar, P., Son, T.C. (eds.) Procs. of LPNMR. LNCS, vol. 8148, pp. 1–10. Springer (2013)
6. Brewka, G., Eiter, T.: Equilibria in heterogeneous nonmonotonic multi-context systems. In: Procs. of AAAI. pp. 385–390. AAAI Press (2007)
7. Brewka, G., Eiter, T., Fink, M., Weinzierl, A.: Managed multi-context systems. In: Walsh, T. (ed.) Procs. of IJCAI. pp. 786–791. IJCAI/AAAI (2011)
8. Brewka, G., Ellmauthaler, S., Pührer, J.: Multi-context systems for reactive reasoning in dynamic environments. In: Schaub, T., Friedrich, G., O'Sullivan, B. (eds.) Procs. of ECAI. Frontiers in Artificial Intelligence and Applications, vol. 263, pp. 159–164. IOS Press (2014)
9. Cardoso, H.L., Oliveira, E.C.: A context-based institutional normative environment. In: Hübner, J.F., Matson, E.T., Boissier, O., Dignum, V. (eds.) COIN Revised Selected Papers. LNCS, vol. 5428, pp. 140–155. Springer (2009)
10. Criado, N., Argente, E., Botti, V.J.: THOMAS: an agent platform for supporting normative multi-agent systems. J. Log. Comput. 23(2), 309–333 (2013)
11. Dastani, M.: 2APL: a practical agent programming language. Journal of Autonomous Agents and Multi-Agent Systems 16(3), 214–248 (2008)
12. Dastani, M., Hindriks, K.V., Novák, P., Tinnemeier, N.A.M.: Combining multiple knowledge representation technologies into agent programming languages. In: Baldoni, M., Son, T.C., van Riemsdijk, M.B., Winikoff, M. (eds.) DALT, Revised Selected and Invited Papers. LNCS, vol. 5397, pp. 60–74. Springer (2008)
13. Dragoni, A., Giorgini, P., Serafini, L.: Mental states recognition from communication. J. Log. Comput. 12(1), 119–136 (2002)
14. Ellmauthaler, S.: Generalizing multi-context systems for reactive stream reasoning applications. In: Procs. of ICCSW. OASICS, vol. 35, pp. 19–26. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, Germany (2013)
15. Gebser, M., Grote, T., Kaminski, R., Schaub, T.: Reactive answer set programming. In: Delgrande, J.P., Faber, W. (eds.) Procs. of LPNMR. LNCS, vol. 6645, pp. 54–66. Springer (2011)
16. Gerard, S.N., Singh, M.P.: Evolving protocols and agents in multiagent systems. In: Gini, M.L., Shehory, O., Ito, T., Jonker, C.M. (eds.) Procs. of AAMAS. pp. 997–1004. IFAAMAS (2013)
17. Giunchiglia, F., Serafini, L.: Multilanguage hierarchical logics or: How we can do without modal logics. Artif. Intell. 65(1), 29–70 (1994)
18. Gonçalves, R., Knorr, M., Leite, J.: Evolving bridge rules in evolving multi-context systems. In: Bulling, N., van der Torre, L.W.N., Villata, S., Jamroga, W., Vasconcelos, W.W. (eds.) Procs. of CLIMA. LNCS, vol. 8624, pp. 52–69. Springer (2014)
19. Gonçalves, R., Knorr, M., Leite, J.: Evolving multi-context systems. In: Schaub, T., Friedrich, G., O'Sullivan, B. (eds.) Procs. of ECAI. Frontiers in Artificial Intelligence and Applications, vol. 263, pp. 375–380. IOS Press (2014)

20. Gonçalves, R., Alferes, J.: Parametrized logic programming. In: Janhunen, T., Niemelä, I. (eds.) Procs. of JELIA. LNCS, vol. 6341, pp. 182–194. Springer (2010)
21. Gonçalves, R., Knorr, M., Leite, J.: Towards efficient evolving multi-context systems (preliminary report). In: Ellmauthaler, S., Pührer, J. (eds.) Procs. of ReactKnow (2014)
22. Hindriks, K.V.: Programming rational agents in GOAL. In: El Fallah Seghrouchni, A., Dix, J., Dastani, M., Bordini, R.H. (eds.) Multi-Agent Programming:, pp. 119–157. Springer US (2009)
23. Ivanov, V., Knorr, M., Leite, J.: A query tool for $\mathcal{EL}$ with non-monotonic rules. In: Alani, H., Kagal, L., Fokoue, A., Groth, P.T., Biemann, C., Parreira, J., Aroyo, L., Noy, N.F., Welty, C., Janowicz, K. (eds.) Procs. of ISWC. LNCS, vol. 8218, pp. 216–231. Springer (2013)
24. Kaminski, T., Knorr, M., Leite, J.: Efficient paraconsistent reasoning with ontologies and rules. In: Procs. of IJCAI. IJCAI/AAAI (2015)
25. Klapiscak, T., Bordini, R.H.: JASDL: A practical programming approach combining agent and semantic web technologies. In: Baldoni, M., Son, T.C., van Riemsdijk, M.B., Winikoff, M. (eds.) DALT 2008, Revised Selected and Invited Papers. LNCS, vol. 5397, pp. 91–110. Springer (2008)
26. Knorr, M., Alferes, J., Hitzler, P.: Local closed world reasoning with description logics under the well-founded semantics. Artif. Intell. 175(9-10), 1528–1554 (2011)
27. Knorr, M., Gonçalves, R., Leite, J.: On efficient evolving multi-context systems. In: Pham, D.N., Park, S. (eds.) Procs. of PRICAI. LNCS, vol. 8862, pp. 284–296. Springer (2014)
28. Knorr, M., Slota, M., Leite, J., Homola, M.: What if no hybrid reasoner is available? hybrid MKNF in multi-context systems. J. Log. Comput. 24(6), 1279–1311 (2014)
29. Meneguzzi, F., Rodrigues, O., Oren, N., Vasconcelos, W.W., Luck, M.: BDI reasoning with normative considerations. Eng. Appl. of AI 43, 127–146 (2015)
30. Moreira, A.F., Vieira, R., Bordini, R.H., Hübner, J.F.: Agent-oriented programming with underlying ontological reasoning. In: Baldoni, M., Endriss, U., Omicini, A., Torroni, P. (eds.) DALT, Selected and Revised Papers. LNCS, vol. 3904, pp. 155–170. Springer (2005)
31. Panagiotidi, S., Álvarez-Napagao, S., Vázquez-Salceda, J.: Towards the norm-aware agent: Bridging the gap between deontic specifications and practical mechanisms for norm monitoring and norm-aware planning. In: Balke, T., Dignum, F., van Riemsdijk, M.B., Chopra, A.K. (eds.) COIN Revised Selected Papers. LNCS, vol. 8386, pp. 346–363. Springer (2013)
32. Roelofsen, F., Serafini, L.: Minimal and absent information in contexts. In: Kaelbling, L., Saffiotti, A. (eds.) Procs. of IJCAI. pp. 558–563. Professional Book Center (2005)
33. Sabater, J., Sierra, C., Parsons, S., Jennings, N.R.: Engineering executable agents using multi-context systems. J. Log. Comput. 12(3), 413–442 (2002)
34. Slota, M., Leite, J.: On semantic update operators for answer-set programs. In: Coelho, H., Studer, R., Wooldridge, M. (eds.) Procs. of ECAI. Frontiers in Artificial Intelligence and Applications, vol. 215, pp. 957–962. IOS Press (2010)
35. Slota, M., Leite, J.: Robust equivalence models for semantic updates of answer-set programs. In: Brewka, G., Eiter, T., McIlraith, S.A. (eds.) Procs. of KR. AAAI Press (2012)
36. Slota, M., Leite, J.: A unifying perspective on knowledge updates. In: del Cerro, L., Herzig, A., Mengin, J. (eds.) Procs. of JELIA. LNCS, vol. 7519, pp. 372–384. Springer (2012)
37. Slota, M., Leite, J.: The rise and fall of semantic rule updates based on se-models. TPLP 14(6), 869–907 (2014)
38. Vasconcelos, W.W., Kollingbaum, M.J., Norman, T.J.: Normative conflict resolution in multi-agent systems. Autonomous Agents and Multi-Agent Systems 19(2), 124–152 (2009)
39. Wang, Y., Zhuang, Z., Wang, K.: Belief change in nonmonotonic multi-context systems. In: Cabalar, P., Son, T.C. (eds.) Procs. of LPNMR. LNCS, vol. 8148, pp. 543–555. Springer (2013)