Reasoning with misperception in the Features and Fluents Framework

Silvia Coradeschi
Department of Computer and Information Science
Linköping University, Sweden
E-Mail: silco@ida.liu.se

Abstract

In this paper we consider a way to deal with the problem of misperception when an agent is reasoning about dynamical systems as they are formalized and systematically studied in Sandewall's approach to reasoning about action and change. The presence of misperception can be detected in case it generates a contradiction. In this case a revision function resolves the inconsistency by constructing alternative consistent descriptions of the system. However, revision is in general expensive and therefore we define a delayed revision that postpones doing revision for some steps and we prove that it gives results similar to immediate revision. We also examine some forms of preferential revision that reduce the number of alternative descriptions of the system. We finally consider the relations between our work and Gärdenfors' approach to belief revision.

1 Introduction

This paper addresses the problem of dealing with misperception in inert and inhabited dynamical systems. A *dynamical system* is a system in which the state changes over time. State changes depend on earlier states of the system, and on the input that the system receives from outside. In a dynamical system we can have *agents* which influence the system states performing *actions* and that receive *observations* about the present state of the world. In this case the dynamical system is said to be *inhabited*.

In this paper, the case in which the observations that the agent receives from the world can be incorrect is considered. We call the presence of incorrect observations misperception. The agent detects the existence of incorrect observations by discovering contradictions in his knowledge, but in general, even if it is aware that something is incorrect, the agent can not precisely identify what is incorrect. We propose a solution where the agent, in case it finds a contradiction, construct several alternative descriptions of the world, each one containing a consistent subset of the observations received.

This work is inserted in Sandewall's *Features and Fluents* framework (Sandewall 1994), a systematic approach to the representation of knowledge and the reasoning about inert and inhabited dynamical systems. Sandewall defines an ontological and epistemological taxonomy which characterizes classes of dynamical systems with specific restrictions on them. The class that is mostly studied in (Sandewall 1994) is the class where the knowledge is explicit, accurate and correct, that is the knowledge is not partial, there are not misperceptions and illusions. We study the class of dynamical systems in the taxonomy in which there can be misperception in the form of erroneous observations. In (Coradeschi 1995) is explained in details how our work relate with Sandewall's approach to the representation of knowledge.

Our approach can be considered a coherence-based approach as defined in (Benferhat, Dubois, & Prade 1995), that is an approach where the aim is revising the knowledge base and restoring consistency. In particular it can be considered near Rescher and Mannor approach to reasoning with inconsistency (Rescher & Mannor 1970).

The problem of unreliable observations is related to the problem of belief revision and in particular to Gärdenfors' work (Gärdenfors 1988). In the last section we consider the relation with this work and we point out similarities and differences.

2 Scenario description

A scenario description is a description of a dynamical system and of how the states of the system have changed over time. Scenario descriptions are defined in (Sandewall 1994) for the class of dynamical systems in which it is assumed that our knowledge of the system is explicit, correct and accurate, that is the actions are exactly specified and the observations of the world are correct. It is also assumed that the system is *inert*, that is the state of the system does not change unless there is an action that changes it.

Definition 1 (scenario description (Sandewall 1994)) *A scenario description is a tuple* $\langle \mathcal{O}, Law, SCD, OBS \rangle$ *where* $\mathcal{O}$ *is a finite object domain,* Law *is a set of action laws,* SCD *is a schedule specifying a number of actions and their temporal order and* OBS *is a set of observations of the world at some point in time.*

The following is an example of a scenario description.

$$
\begin{array}{ll}
obs1 & [0]on(glass, table). \\
obs2 & [4]on(glass, shelf). \\
law1 & [t,s]Move(glass, table, shelf) \Rrightarrow \\
 & [t]on(glass, table) \rightarrow \\
 & [t,s]on(glass, shelf) {:=} T \wedge \\
 & [t,s]on(glass, table) {:=} F. \\
scd1 & [2,3]Move(glass, table, shelf).
\end{array}
\tag{1}
$$

This represents a scenario description where the the set of action laws *Law* contains the formula labeled $law1$. This formula specifies exactly the action of moving a glass from the table to the shelf. If the glass is on the table when the action is performed, then the glass will be on the shelf at the end of the action. If the glass is not on the table, then the action will have no effect. The observation $obs1$ says that at the time point 0 the glass is on the table. In the schedule there is an action of moving the glass from the table to the shelf and the observation $obs2$ confirms that the glass is on the shelf at the time point 4. The action instances in the schedule are expanded according to the action laws. For example the action $[2,3]Move(glass, table, shelf)$ is expanded to $[2]on(glass, table) \rightarrow [2,3]on(glass, shelf){:=}T \wedge [2,3]on(glass, table){:=}F$.

$[t,s]on(glass, shelf){:=}T$ and $[t,s]on(glass, table){:=}F$ are reassignment formulae stating that the feature "on(glass,shelf)" will get the value true,"T", and the feature "on(glass,table)" will get the value false, "F", somewhere in the interval $(t,s]$.

In general the information that is contained in the scenario is incomplete, but we can deduce some of the missing information considering the assumption of inertia. The set of models of a scenario description is the subset of the classical models of the formulas in the scenario that satisfy the inertia principle. How these models are defined is specified in detail in (Sandewall 1994) .

For instance in the example above, all models have in common that the glass is on the table from time point 0 to time point 2 and it is on the shelf from time point 3 on:

$$\begin{array}{l} [0,2]on(glass, table) \\ [3,\infty)on(glass, shelf) \wedge \neg on(glass, table) \end{array} \qquad (2)$$

3 Hypothetic scenario description, narrative and revision

One of the assumption in the previous section is that all knowledge that the agent has is correct, including knowledge obtained by observations. In case of misperception this assumption is relaxed.

The problem in case it is possible to have misperception is that we cannot detect in general if a specific observation is incorrect. Otherwise, the solution would be just to eliminate it. The presence of an incorrect observations can be detected only if the observation contradicts some other observation or the rest of the scenario. However, we cannot in general detect which observation is incorrect.

In particular, there are three cases in which we can detect if there are incorrect observations. The first case is when there is a conflict between the observations. This happens if there are contradictory observations at the same time point, or if observations at different time points conflict because of the principle of inertia. The latter happens when new observations indicate changes in the world that are not justified by any action in the schedule. As it is assumed that only actions can change the world, this conflict of observations is ascribed to the presence of incorrect observations, even if there is no possibility in general to know which of the observations is incorrect. For instance in the scenario (3) the glass is on the table at time point 0 and it is not on the table at time point 2 without any action that can justify this change.

$$\begin{array}{ll} law1 & [t,s]Move(glass, table, shelf) \Rrightarrow \\ & \quad [t]on(glass, table) \rightarrow \\ & \quad [t,s]on(glass, shelf){:=}T \\ & \quad \wedge[t,s]on(glass, table){:=}F. \\ obs1 & [0]on(glass, table). \\ obs2 & [2]\neg on(glass, table). \\ scd1 & [2,3]Move(glass, table, shelf). \end{array} \qquad (3)$$

The second case is when the initial conditions of the action are in accordance with the observations but the observations do not agree with the effects that the actions are supposed to have given these conditions. This means, as we assume that the actions perform as they are supposed to do, that there are incorrect observations either relating to the initial conditions or to the effect of the action. In the following example, scenario (4), the condition $[2]on(glass, shelf)$ is satisfied. In fact, $obs1$ says that the glass is on the table at 0 and there are no actions that move it between 0 and 2, but the action does not seem to have the effect that it is supposed to have. Instead, $obs2$ says that the glass is not on the shelf. Thus, either obs1 or obs2 is incorrect.

$$\begin{array}{ll} law1 & [t,s]Move(glass, table, shelf) \Rrightarrow \\ & \quad [t]on(glass, table) \rightarrow \\ & \quad [t,s]on(glass, shelf){:=}T\wedge \\ & \quad [t,s]on(glass, table){:=}F. \\ obs1 & [0]on(glass, table). \\ scd1 & [2,3]Move(glass, table, shelf). \\ obs2 & [4]\neg on(glass, shelf). \end{array} \qquad (4)$$

The third case is the case when one observation $obsj$ conflicts with the scenario in itself. The scenario description (5) is an example of an observation, $obs1$, which conflicts with the scenario in itself. In fact, if the glass is on the table at time point 4, it was on the table also at time point 2 as a consequence of inertia and the action of moving it on the shelf should have been successful. However this implies that the glass should not be on the table anymore at time point 3 and, as there are no other actions, it should not be on the table at time point 4 either.

$$\begin{array}{ll} law1 & [t,s]Move(glass, table, shelf) \Rrightarrow \\ & \quad [t]on(glass, table) \rightarrow \\ & \quad [t,s]on(glass, shelf){:=}T\wedge \\ & \quad [t,s]on(glass, table){:=}F. \\ scd1 & [2,3]Move(glass, table, shelf). \\ obs1 & [4]on(glass, table). \end{array} \qquad (5)$$

In case of observations that conflict with the schedule we can say that these observations are surely incorrect, as we assume that the actions have always the effect that are supposed to have, and we can eliminate them.

In the other two cases we know that there are incorrect observations, but we do not know which observations are incorrect. We then consider several sets of observations that represent alternative descriptions of the world. This is the idea behind the definition of hypothetic scenario description.

Definition 2 (hypothetic scenario description) *A hypothetic scenario description is a tuple: $\langle \mathcal{O}$, Law, $SCD, \{OBS_1, \ldots, OBS_n\}\rangle$ where $\mathcal{O}$ is an object domain, Law is a set of action law, SCD is a schedule specifying a number of actions and their temporal order, and each OBS_i is a set of hypothetic observations of the world at some points in time. The observations are represented as atomic formulae.*

A hypothetic scenario description is correct if has models for every set of observations.

Definition 3 (correct hypothetic scenario description) *A hypothetic scenario description $\langle \mathcal{O}$, Law, SCD, $\{OBS_1, ..., OBS_n\}\rangle$ is correct if for all OBS_k, the set of models of $\langle \mathcal{O}, Law, SCD, OBS_k\rangle$ is not empty.*

Now we can define the concept of narrative. The idea behind this definition is that to each time point corresponds a hypothetic scenario description in which we have the observations that we have got until that time point. At each time point t_i we get a set of new observations that we call $NewObs_i$ and that we add to our sets of observations. The sets of observations in the hypothetic scenario description at t_i are included in the union of one of the set of observations at t_{i-1} and $NewObs_i$. The inclusion can be proper if among the new observations there are observations that contradict each other, or contradict previous observations, or contradict the rest of the scenario.

Definition 4 (narrative) *A narrative is a mapping from $\{t_0, \ldots, t_k\}$ to hypothetic scenario descriptions which have $\mathcal{O}$, Law and SCD in common and which satisfy the following condition: if t_i maps to $\langle \mathcal{O}$, Law, SCD, $\{OBS_1^i, ..., OBS_n^i\}\rangle$ and t_{i+1} maps to $\langle \mathcal{O}$, Law, SCD, $\{OBS_1^{i+1}, \ldots, OBS_m^{i+1}\}\rangle$ and $NewObs_{i+1}$ is the set of new observations at time t_{i+1}, then for all $j \in \{1, \ldots m\}$ there exists an $s \in \{1, \ldots n\}$ such that $OBS_j^{i+1} \subseteq OBS_s^i \cup NewObs_{i+1}$.*

A narrative can be correct, completely correct or faulty.

Definition 5 (correct narrative) *A correct narrative is a narrative in which the hypothetic scenario description that corresponds to the last time point t_k of the narrative is correct.*

We call *faulty narrative* a narrative that is not correct.

Definition 6 (completely correct narrative) *A completely correct narrative is a narrative in which all the hypothetic scenario descriptions of the narrative are correct.*

Given a faulty narrative, a revision is a function that transforms it into a correct narrative. We first give the auxiliary definition of maximal consistent subset and then the definition of revision.

Definition 7 (maximal consistent subset) *A set of observations OBS' is a maximal consistent subset of a set of observations OBS relative to a scenario $\langle \mathcal{O}$, Law, SCD, $\emptyset\rangle$ if $OBS' \subseteq OBS$, the set of models of $\langle \mathcal{O}$, Law, SCD, $OBS'\rangle$ is not empty, and for any OBS'' such that $OBS' \subset OBS'' \subseteq OBS$, the set of models of $\langle \mathcal{O}$, Law, SCD, $OBS''\rangle$ is empty.*

Definition 8 (revision) *Let $\langle \mathcal{O}$, Law, SCD, $\{OBS_1, ..., OBS_n\}\rangle$ be the last step of a faulty narrative N. A revision is a function Rev(N) which maps the faulty narrative N into the narrative which is equal to N, except for the last scenario, which has this form: $\langle \mathcal{O}$, Law, SCD, $\{OBS_1', \ldots, OBS_m'\}\rangle$ where $\{OBS_1', ..., OBS_m'\}$ contains for every element in $\{OBS_1, ..., OBS_n\}$ exactly the maximal consistent subsets of this element relative to $\langle \mathcal{O}$, Law, SCD, $\emptyset\rangle$.*

Theorem 1 *Rev(N) always returns a correct narrative and this narrative is unique.*

The proof of the theorem is in (Coradeschi 1995).

A narrative can be constructed as follow:

- First we have the narrative that only contains the mapping from time point t_0 to the hypothetic scenario description $\langle \mathcal{O}, Law, SCD, \{\emptyset\}\rangle$.
- If we have a narrative whose last time point is t_k, in order to add another mapping element $t_{k+1} \rightarrow \langle \mathcal{O}, Law, SCD, \{OBS_1^{k+1}, \ldots, OBS_n^{k+1}\}\rangle$ we add the new observations that we get at t_{k+1} to every $OBS_1^i, ..., OBS_n^i$ and if we get a faulty narrative we do a revision.

Constructing a narrative in this way, we obtain a completely correct narrative. In section 5 we consider the case in which delayed revision is done. In this case, the narrative that is obtained is correct, but not completely correct.

4 Models of hypothetic scenario description

We define as models of the hypothetic scenario description the union of the models that we get in the scenario descriptions which contain just one of the sets of observations of the hypothetic scenario description. Formally:

Definition 9 (Models of hypothetic scenario description) *Let Mod(T) denote the set of models of a scenario description T. If $S = \langle \mathcal{O}$, Law, SCD, $\{OBS_1, \ldots, OBS_n\}\rangle$ is a hypothetic scenario description, then the set of models of S, Modh(S), is:*

$$Modh(S) = \bigcup_{1 \leq i \leq n} Mod(\langle \mathcal{O}, Law, SCD, OBS_i\rangle)$$

We can suppose without losing generality that the set of observations at t_0 is empty. In this case the models of the hypothetic scenario description at t_0 in a narrative are all the models that satisfy the schedule and which preserve inertia. Adding new observations in the successive time points results in some of the models being thrown away, but doing revision some of the models that have been thrown away at previous time points are made alive again.

Theorem 2 *Let N be a narrative where the set of observations at t_0 is empty.*

(a) The models of the hypothetic scenario description at t_0 are all the models that satisfy the schedule and which preserve inertia.

(b) If at t_{i+1} we do not do revision, the models of the hypothetic scenario description at t_{i+1} are a subset of the models of the hypothetic scenario description at t_i.

(c) If at t_{i+1} we do revision, each of the models of the hypothetic scenario description that we obtain after the revision was also model of at least one previous scenario description.

The proof of the theorem is in the technical report (Coradeschi 1995).

5 Delayed revision

Since it can be expensive to check consistency and to do revision, we consider the possibility of delaying the revision for some time points. If instead of checking whether a narrative is faulty we would go on developing the narrative, just adding the new observations at each time point and doing the revision at time point k, does then the revision at time point k give the same result as we would have obtained at k if we had done the revision every time it was necessary?

The following theorem proves that we obtain the same scenario description at time point k.

Theorem 3 *Let N be a narrative. If the revision is skipped between n additions of observations in N the same hypothetic scenario description is obtained as a result of the revision after the n additions as if the revision would have been done every time it was necessary.*

The proof of the theorem is in (Coradeschi 1995).

6 Preferential revision

When we do a revision of a narrative we can get several sets of observations. Instead of taking all the sets, we can make a choice and take just some of them. This is the idea behind the definition of preferential revision.

Definition 10 (preferential revision) *Let N be a narrative, N' the narrative obtained doing the revision of N and $\langle \mathcal{O}$, Law, SCD, $\{OBS_1,...,OBS_n\}\rangle$ the last hypothetic scenario description of N'. A preferential revision returns a narrative equal to N' except for the last hypothetic scenario description which has the form $\langle \mathcal{O}$, Law, SCD, $\{OBS_1',...,OBS_m'\}\rangle$, where $\{OBS_1',...,OBS_m'\} \subseteq \{OBS_1,...,OBS_n\}$.*

The criterion for choosing some sets of observations instead of others depends mostly on the application we are dealing with. We can anyway give some general criteria of choice. In the rest of the section we examine four of these criteria.

6.1 Preference for the last or the first observations

The first case we consider is the case in which we prefer the last observations that have been received. A way to prefer the last observations is to take, when we do the revision, just the maximal subsets that contain all the new observations.

Definition 11 (Last observation preferential revision) *Let N be a narrative, NewObs be the new observations that we have added at the last time point, N' the narrative obtained doing the revision of N and $\langle \mathcal{O}$, Law, SCD, $\{OBS_1,...,OBS_n\}\rangle$ the last hypothetic scenario description of N'. A last observation preferential revision returns a narrative equal to N' except for the last hypothetic scenario description which has the form $\langle \mathcal{O}$, Law, SCD, $\{OBS_1',\ldots,OBS_m'\}\rangle$ where $\{OBS_1',...,OBS_m'\}$ contains all the OBS_i in $\{OBS_1,...,OBS_n\}$ such that $NewObs \subseteq OBS_i$.*

This preferential revision has the advantage that the choice of the preferred subset is very easy: one only needs to check a set inclusion. A disadvantage is that we must assume that the set of the new observations is consistent. This can be reasonable, but of course it is a limitation.

Another thing that we can notice is that we do not consider the order in which the observations had arrived in the successive revisions.

Let us now consider the preference of the first observations. In this case, we choose the maximal subsets that contain all the old observations. The formal definition is quite similar to the previous one:

Definition 12 (First observation preferential revision) *Let N be a narrative, N' the narrative obtained doing the revision of N and $\langle \mathcal{O}$, Law, SCD, $\{OBS_1,...,OBS_n\}\rangle$ the last hypothetic scenario description, corresponding to the time point t_i, of N'. A first observation preferential revision returns a narrative equal to N' except for the last hypothetic scenario description which has the form $\langle \mathcal{O}$, Law, SCD, $\{OBS_1',\ldots,OBS_m{}'\}\rangle$ where $\{OBS_1',\ldots,OBS_m'\}$ contains all the OBS_i in $\{OBS_1,\ldots,OBS_n\}$ such that there exists a set of observations OBS_j^{i-1} in the hypothetic scenario description at t_{i-1}: $OBS_j^{i-1} \subseteq OBS_i$.*

In this case we include without checking its consistence a OBS_j^{i-1} in every set that we obtain doing the revision, but this does not give any problems because the OBS_j^{i-1} are all consistent. We still have the same considerations as before about the fact that we do not retain the order of arrival of the observations in the successive revisions.

6.2 Preference for observations in which we have different belief

To define this kind of preferential revision we first introduce the new concept of *belief observation*, that is, a pair *(observation, number)* where the number represents our belief in the observation.

Definition 13 (belief observation/scenario description/narrative) *A belief observation is a pair (obs, n) where obs is one observations and n is a number that represents the belief in the observation. A belief scenario description is an hypothetic scenario description containing belief observations instead of observations. A belief narrative is a narrative containing belief scenario description instead of hypothetic scenario description.*

The definitions of faulty and correct belief narrative are analogous to faulty and correct narrative, with belief scenario description instead of hypothetic scenario description.

Definition 14 (get-val) *Get-val is a function that takes a belief observation (obs, n) and give back n.*

We can now define a *belief preferential revision*; in this case we prefer the sets of observations for which the sum of the beliefs of the observations is maximal.

Definition 15 (Belief preferential revision) *Let N be a belief narrative, N' the belief narrative obtained doing the revision of N and $\langle \mathcal{O}$, Law, SCD, $\{BOBS_1,...,BOBS_n\}\rangle$ the last hypothetic scenario description of N'. A belief preferential revision returns a narrative equal to N' except for the last hypothetic scenario description which has the form $\langle \mathcal{O}$, Law, SCD, $\{BOBS'_1,\ldots,BOBS'_m\}\rangle$ where $\{BOBS'_1,...,BOBS'_m\}$ contains all the $BOBS_i$ in $\{BOBS_1,...,BOBS_n\}$ such that $\sum$ get-val(obs_i, n) is maximal (no other $BOBS_j$ gives a higher value) and where the $(obs_i,\ n)$ are the belief observations of $BOBS_i$.*

With the aid of the belief preferential revision we can define different versions of preference of the last and the first observations. In the first case we can just put as belief value the time of arrival of the observations. In the second case we can put as belief value the time of arrival changed of sign or we can put the time of arrival as it is, but instead of preferring the sets of observations for which the sum of the belief of the observations is maximal, we prefer the sets for which it is minimal.

With this new definition we do not have the problems that we had with the previous definitions of first and last observation preference revision. In fact we do not need to assume that the new observations are consistent as the new observations are checked in the same way as the other observations and the order of arrival is preserved also in the successive revisions.

Another property that can be interesting is that if we are in the case of last observation preference and we have to choose between a smaller set of more recent observations and a larger set of earlier observations, it can happen that we choose the second set of observations.

6.3 Context-dependent belief

In the previous section we have assumed that a belief level is directly associated with each observations. However in many realistic applications the belief in observations can be context-dependent. If there is not so much light we may be more prepared to accept that we make a mistake in an observation, than if it is clear daylight.

Also, the observations can be differently affected by context factors, for example the observations that concern the sight can be affected by the amount of light, while the observations that concern the hearing can be affected by the amount of noise in the environment.

To formalize context-dependent belief we can divide our observations into two categories: *context observations*, that is observations that give information about the context, for example: it is daylight, there is music, and *non-context observation*. Each context observation changes the value of belief of the subset of non-context observations which are affected by it.

Definition 16 (context function) *A context function is a function that takes a set of context observations and a non context dependent observation (obs, n) and gives back a pair (obs, n'), where n' is the belief value of the observation modified according to the context observations.*

Definition 17 (context belief narrative) *A context belief narrative is a belief narrative where at each time point t_i, instead of adding $NewObs_i$ at each $BOBS_j$, the set $NewObs'_i$ is added, which is obtained by applying the context function to all the (obs, n) in $NewObs_i$, with the context observations in $BOBS_j \cup NewObs_i$ as set of context observations.*

An algorithm for calculating the context function is the following:

- we decide the observations that will be context observations;
- we associate to each context observation the set of observations that are affected by it and how much the fact that this observation is present increases or decreases the belief value of the observations (we suppose that the context observation affects all the observations in the same way);
- the context function takes a set of context observations C and a pair (obs, n), checks if obs is in the set associated with a context observation in C and, if that is the case, changes the value n accordingly. The observation obs can be in the set associated with several context observations. The value that the context function gives back is the one after all the changes.

Definition 18 (context belief preferential revision) *The definition of context belief preferential revision is the same as the definition of belief preference, substituting context belief narrative for belief narrative.*

7 Relations with Gärdenfors' approach to belief revision

In his book 'Knowledge in flux' (Gardenfors 1988), Gärdenfors considers the case in which people are forced to revise their beliefs, that is, they have a new belief and if they want to accept this new belief they must give up some of their old beliefs. The problem is deciding which beliefs to give up. Unfortunately, there are not in general any logical reasons for choosing one belief other another. However one can consider a *criterion of informational economy* that demands that as few beliefs as possible should be given up in order to accept the new belief. In some sense the change should be a *minimal change*. Gärdenfors defines a number of postulates that have the aim to delimit the meaning of "minimal change". He assumes that with the postulates and with certain *non logical* factors it is possible to obtain a unique revision of a set of beliefs.

Our approach is in some ways similar: we revise a particular kind of beliefs, that is, the observations of the world, and we respect the principle of minimal change. In fact, we take the maximal sets of observations that are correct for our narrative, but we do not have general non logical factors giving us a unique revised set of observations. In general, it is not possible, given a set of observations, to decide which ones are incorrect. Therefore, our revision gives back several alternative sets of observations. Non-logical factors can give a preference in the observations, and this is the case that we have examined with preference revision, but in general preference revision do not always yield a unique set of observations. It can just choose some of the sets.

In Gärdenfors' work it is assumed that the revised set of beliefs should contain the new belief A. The belief A is not compared with the old beliefs that should be thrown away to accept A, A is privileged in every case. In our case, the new observations are not in general privileged and in the definition of revision without preference we keep both the new and the old observations; we just put the contradictory observations in different sets and we consider these sets as alternative descriptions of the world.

The postulates of Gärdenfors apply in general to our revision with some adaptation due mostly to the differences outlined in the previous paragraphs. In (Coradeschi 1995) the applicability of Gärdenfors' postulates to our revision is considered in detail.

Gärdenfors' postulates are in general applicable also to preferential revision. In particular in last observation preference revision the new observations are always preferred, and this corresponds to Gärdenfors' assumption that the new belief is always accepted.

8 Conclusion

This paper presents a way to revise knowledge in a dynamical systems in case incorrect observations generate a contradiction. We introduce a revision function, a delayed revision that postpones the revision of some steps, and a preferential revision that prefers some sets of observations over the other.

We also want to mention that revision and preferential revisions have been implemented in Prolog.

We intend to extend this work considering the case in which the presence of unreliable observations influences the schedule, for example introducing informative actions, and we will in particular consider planning with unreliable observations.

9 Acknowledgments

I would like to thank Lars Karlsson, Erik Sandewall, Maria Simi, Patrick Lambrix and Pavlos Peppas for comments on earlier drafts of this paper.

10 References

Benferhat, S., Dubois, D., and Prade, A. 1995. How to infer from inconsistent belief without revising? In Proceedings the 14th International Joint Conference on Artificial Intelligence, 1449–1455. Montreal, Canada: International Joint Conferences on Artificial Intelligence, Inc.

Coradeschi, S. 1995. Reasoning with unreliable observations in the features and fluents framework, Technical Report, LITH-IDA-R-95-44. Dept. of Computer and Information Science, Linköping Univ.

Gärdenfors, G. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* MIT Press.

Rescher, N. & Mannor, R. 1970. On inference from inconsistent premises. In *Theory and Decision*, 1:179–219.

Sandewall, E. 1994. *Features and Fluents.* Oxford Press.