



HAL
open science

Reasoning About Action and Change

Florence Dupin de Saint-Cyr, Andreas Herzig, Jérôme Lang, Pierre Marquis

► **To cite this version:**

Florence Dupin de Saint-Cyr, Andreas Herzig, Jérôme Lang, Pierre Marquis. Reasoning About Action and Change. Marquis, Pierre; Papini, Odile; Prade, Henri. A Guided Tour of Artificial Intelligence Research, 1 / 3, Springer International Publishing, pp.487-518, 2020, Knowledge Representation, Reasoning and Learning, 978-3-030-06163-0. 10.1007/978-3-030-06164-7_15 . hal-03015828

HAL Id: hal-03015828

<https://hal.science/hal-03015828v1>

Submitted on 26 Jun 2024

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.

Reasoning about Action and Change*

Florence Dupin de Saint-Cyr¹, Andreas Herzig¹, Jérôme Lang², Pierre Marquis³

¹ IRIT-CNRS, Université Paul Sabatier, Toulouse, France

² CNRS, Université Paris-Dauphine, PSL Research University, LAMSADE, Paris, France

³ CRIL-CNRS, Université d'Artois & Institut Universitaire de France, Lens, France

November 24, 2023

Abstract

This chapter presents the state of research concerning the formalisation of an agent reasoning about a dynamic system which can be partially observed and acted upon. We first define the basic concepts of the area: system states, ontic and epistemic actions, observations; then the basic reasoning processes: prediction, progression, regression, postdiction, filtering, abduction, and extrapolation. We then recall the classical action representation problems and show how these problems are solved in some standard frameworks. For space reasons, we focus on these major settings: the situation calculus, STRIPS and some propositional action languages, dynamic logic, and dynamic Bayesian networks. We finally address a special case of progression, namely belief update.

1 Introduction

In this chapter, we are interested in *formalizing the reasoning of a single agent* who can make *observations* on a *dynamic system* and considers *actions* to perform on it. Reasoning about action and change is among the first issues addressed within Artificial Intelligence (AI); especially, it was the subject of the seminal article by McCarthy and Hayes [1969]. Research in this area has been very productive until the late 1990s. Among other things, solutions to the various problems to be faced when dealing with action representation were put forward and a classification of action languages according to their expressive power was undertaken. Moreover, much progress towards the automatization of reasoning about action and change was made, for example through the design and the evaluation of algorithms implementing the reasoning processes of the main action languages and the investigation of the computational complexity of such processes.

The reasons why an agent may wish to act in order to modify the current state of a dynamic system or to learn more about it are numerous. For example, the goal can be to change the system into a configuration that the agent prefers over the actual one (such as moving a robot from a location to another one), or even into an optimal configuration. Alternatively, the objective can be to ensure that a certain property of the dynamic system is maintained, or that its successive states do not deviate too much from a normal path. The latter is for example the case when one wants to supervise and control a physical system, such as a furnace or that of a patient in an intensive care unit. Such scenarios involve concepts (state, action, observation, etc.) and processes connecting them (planning, prediction, explanation, etc.).

By ‘formalizing’, we first mean *modeling the concepts and processes* that are considered in such scenarios (the purpose is to define them rigorously from a mathematical point of view) and then *representing* them (that is, specifying how the information are encoded) and *automating* (designing algorithms suited to the processes under consideration). Note that there are typically two main reasons for modeling a dynamic system: *controlling* it (see Chapter 10 of Volume 2 about planning), and obtaining more information about it, for diagnosing it or supervising it (see Chapter 31 of this volume). Once modeled, the same concept can be associated with several representations. If the choice of a specific model typically depends on the available pieces of information and what one wants to do with them, the choice of a

*This is an accepted manuscript of the chapter Reasoning about Action and Change of the book A Guided Tour of Artificial Intelligence Research, 1 / 3, Springer International Publishing, pp.487-518, 2020

representation (suited to a given model) is based on other criteria, such as computational efficiency and succinctness.

2 Reasoning about Action: Models

2.1 Basic Concepts and the Corresponding Models

In this section, we define some mathematical notions corresponding to the key concepts considered in reasoning about action and change.

The *model* of a reasoning process on a *dynamic system* can be divided in two parts: the model of the system (with its own dynamics) and the model of the agent (including her knowledge about the system). Sandewall [1995] has developed a taxonomy of reasoning problems on dynamic systems, and the remainder of this chapter elaborates on it.

Throughout the chapter, we assume that *time* is discrete (which is a common assumption in Artificial Intelligence). The *horizon* of the process is the set \mathcal{H} of relevant steps for controlling and observing it. It can be finite ($\mathcal{H} = \{0, \dots, N\}$ with $N \in \mathbf{N}$) or infinite ($\mathcal{H} = \mathbf{N}$); a degenerate case of a finite horizon is when there is only one change step ($\mathcal{H} = \{0, 1\}$).

A *state* is the description of the system at a given time point. Unless stated otherwise, *the set of all states*, denoted by \mathcal{S} , will be assumed finite. A *state trajectory* is a sequence of elements of \mathcal{S} , indexed by elements of \mathcal{H} . The system states at the different time points of \mathcal{H} may only partially known by the agent.

The specification of a reasoning process on a dynamic system requires first the *beliefs* of the agent about the state of the system at different time points (including the initial time point) and about the general laws that govern the evolution of the system. Thus, we first have to choose a model for uncertain belief states. For space reasons, we will focus only on two uncertainty models in the following: the *binary model*, where belief states b are *non-empty subsets of \mathcal{S}* and the *Bayesian model*, where belief states b are *probability distributions on \mathcal{S}* .¹

Transitions from one state to another are triggered by *events*. These events generally change not only the state of the system, but also the beliefs of the agent. An *action* is a special event triggered by an agent. The agent has a model of each *action* available to her. The set of actions available to the agent is denoted by \mathcal{A} , and is supposed to be finite. The agent can also have a model for *exogenous events*, which are phenomena whose dynamics are similar to actions but which are not triggered by the agent. They are triggered by nature or possibly by other agents more or less well-identified (i.e., whose identity may be imperfectly known), and their occurrences are *a priori* not known by the agent. We distinguish between the *action type α* (defined very generally) from the action occurrence(s) at one or more time point(s): a given action may have no occurrence in an instance of a problem, or may have one or several occurrences. Actions have two types of effects: *ontic (physical) effects*, on the world, and *epistemic effects*, on the beliefs of the agent. Epistemic effects can either be caused by her projection of the physical effects of the performed action (for instance, if I know that the action “delete file F” has the effect that file F no longer appears on my computer, then, when I execute this action, the resulting belief state is such that I know that F no longer is on my computer) or from *observations* or any form of feedback (for instance, if after trying to turn the light on by flipping the switch, I observe that the light is off, then, in my new state of belief, I know that the bulb is broken or that the power is off).

Actions have generally two types of effects at once (as in the case of the “switch” action above). Some actions, referred to as *purely epistemic actions*, have only epistemic effects, and no effect on the state of the world; for example, measuring a temperature, or querying a database. Other actions, referred to as *purely ontic actions* have epistemic effects (it is hard to imagine actions without any epistemic effect, apart from the action “do nothing”), but these epistemic effects are the simple projection, by the agent, of what she knows about the ontic effects of the action (as for the action “delete file F” above). In other words, a purely ontic action gives *no feedback* to the agent: her belief state after the execution of such an action coincides with the belief state she could foresee before executing the action (“what you foresee is what you get”). Every action can be decomposed in a unique way into a purely ontic action and a purely epistemic action; without loss of generality, we can thus assume that each available action is either purely ontic or purely epistemic (and we will make such an assumption in the rest of the chapter, unless stated otherwise).

¹There are many other uncertainty models that should be mentioned but will not be, for space reasons - they include ordinal models, where belief states and action effects are modeled as pre-orders over \mathcal{S} , possibilistic models that are similar in spirit to them, non-Bayesian probabilistic models, where a belief state is a family of probability distributions, etc. (see Chapter 3 of this volume).

Let us start by describing *purely ontic actions*. The effects of a purely ontic action α are defined by a *transition system* between the states of the world, modeled as a *binary relation* R_α over \mathcal{S} .

The simplest case is when actions are *deterministic and always executable*. In this case, the transition system of α is a mapping R_α from \mathcal{S} to \mathcal{S} . An action has *conditional effects* if the resulting state after its execution depends on the state before its execution. For example, the action “switch off the light” may be regarded as deterministic and unconditional (if we assume that it always has the effect that the lamp is off after its execution). “Toggle the switch” can be considered as deterministic, and with conditional effects since its effects depend on the state (“on” or “off”) of the light before the execution of the action.

More generally, actions are not always executable: there can be states s such that $R_\alpha(s) = \emptyset$; actions can also be *non-deterministic*: there are states s such that $R_\alpha(s)$ contains more than one element. For example, the action “delete file F” is not executable if the file does not exist; in this case, the modeler will define the effect of the action only for states where the file exists, and executing the action will be forbidden in the other cases. Another model would make advantage of an action with conditional effects, where the transition associated to the action would be associated with the identity relation in situations where file F does not exist, and would lead to states where the file is deleted otherwise.

In the non-deterministic case, the transition model chosen depends on the nature of the uncertainty one wants to deal with; with each initial state is associated a belief state on the subsequent states. Note that choosing a deterministic or a non-deterministic model for a system may depend on the knowledge and the goals of the modeler: the action “turn the computer off” can be considered as non-deterministic for an agent who is not a computer scientist (since it may happen that after the execution of the action the computer is still on) but as deterministic, with conditional effects, for an expert in computer science (since this expert will be able to determine the cases where the computer stays on after being turned off). Modeling a dynamic system as a transition system between states amounts to making the implicit assumption that the system is Markovian.² Such an assumption can be made without loss of generality by considering more complex states (encoding state trajectories). For the sake of brevity, we will stick to the following two models: the binary non-deterministic model and the stochastic model.

In the binary non-deterministic model, the transition system of an action α is a mapping R_α from \mathcal{S} to $2^{\mathcal{S}}$ (or to $2^{\mathcal{S}} \setminus \{\emptyset\}$, when α is always executable). For example, if the system states are $\mathcal{S} = \{c_on, c_stand_by, c_off\}$ (representing the activity of a computer: “on”, “stand-by” or “off”) then the action of “shutting down the computer” may be modeled as $R_{shut_down}(c_on) = \{c_on, c_off\}$, $R_{shut_down}(c_stand_by) = R_{shut_down}(c_off) = \emptyset$ (meaning that one can “turn off the computer” only if it is “on”, and in this case, it is not sure that the “shut down” action succeeds). Note that if α is a purely epistemic action, then $R_\alpha(s) = \{s\}$ for all s .

In the stochastic model (here, the Bayesian model for uncertainty), R_α is a stochastic matrix, i.e., a family of probability distributions $p(.|s, \alpha)$ for $s \in \mathcal{S}$, where $p(s'|s, \alpha)$ is the probability to obtain the state s' after the execution of α in s . In this model, it is possible to specify how much the “shut down” action succeeds; thus, R_{shut_down} could be represented by $p(c_on | c_on, shut_down) = 0.1$, $p(c_off | c_on, shut_down) = 0.9$, $p(c_stand_by | c_on, shut_down) = 0$.

Epistemic effects of actions are expressed in terms of feedback. The actions that the agent decides to execute do not depend directly on the system state (which may be unknown to the agent) but on the agent’s beliefs (and in particular, on what has been observed earlier). Ideally, the current state and what the agent may observe coincide. In this case, the belief state of the agent is perfect, but this hypothesis reflects an ideal case and does not often hold. In order to define the epistemic effects of actions, an *observation space* Ω can be introduced in the model. This space, unless otherwise indicated, is supposed finite. The observations are the feedback given by the system and each observation at a given time point from the horizon is a projection of a state (not necessarily totally observed) of the system. The observations are called *reliable* if this projection corresponds to the *actual state* of the system (unreliable observations can arise from faulty sensors, for example).

Taking observations into account concerns two distinct stages: the off-line stage when the decision policy is generated, and the on-line stage when it is exploited (i.e., when the plan is executed). During the off-line stage, the agent who generates the policy takes advantage of her knowledge about the observations which could be made at the on-line stage. During the on-line phase, the actions which are triggered by the agent typically depend on the observations which are effectively made. Note that nothing prevents from having two distinct agents (one who computes a decision policy and another one who executes it).

Two assumptions corresponding to two extreme cases are commonly made: when the system is *fully observable*, the observation space is identical to the state space: when she generates a decision policy,

²A system is Markovian if the transition of the system to any given state depends only on the current state and not on the previous ones.

the agent knows that at the on-line stage, the actual state of the system will be known exactly at each time point; when the system is *non-observable*, the observation space is a singleton $\{o^*\}$, where o^* is a fictitious observation (the empty observation): the system gives no feedback.

When none of these two extreme assumptions hold, one faces the more general situation of *partial observability*, where observations and system states are constrained by an observation-state *correlation structure*, the definition of which varies with the uncertainty model under consideration. In general, each state s and each *epistemic action* α correspond to a *belief state* $\mathbf{O}_\alpha(s)$ over the space of observations, representing the *prior* beliefs about the observations obtained when action α is performed in state s . In the binary model for uncertainty, a family of sets $\mathbf{O}_\alpha(s)$ for $s \in \mathcal{S}$ is thus considered, where $\mathbf{O}_\alpha(s)$ is a *non-empty set of Ω* ; so $o \in \mathbf{O}_\alpha(s)$ reflects the fact that s is a state compatible with the observation o which results from the execution of α . If α is purely ontic, then $\mathbf{O}_\alpha(s) = \{o^*\}$ for all s . In the Bayesian model for uncertainty, the *feedback* is modeled by a probability distribution $p(\cdot|s, \alpha)$ on Ω where $p(o|s, \alpha)$ is the *probability to observe o when α is executed in s* .

In this section, we assumed that only one action at a time can be executed. In some problems, it is natural to perform several actions in a *concurrent* way. This requires to be able to define the effects of combinations of actions; for reaching this goal, the same models as previously considered can be used, viewing every possible combination of actions as a specific action. A typical example [Thielscher, 1995] is the one of a table that can be lifted by the right side or by the left side: the two actions performed in sequence and independently do not have the same effect as when they are executed simultaneously, especially when a glass of water is on the table!

2.2 Types of Reasoning and their Implementations

Reasoning on a dynamic system requires to take account of a time horizon, the *prior* beliefs on the system (general laws of the domain and action effects), the occurrences of actions at some time points, and the observations at given time points (it is a simplified model – see [Sandewall, 1995] for a more general one, where, in particular, the actions can have a duration). We are now going to approach some specific types of reasoning implying reasoning on a dynamic system, as well as their implementation by means of algorithms.

2.2.1 Prediction and Postdiction with Ontic Actions

Prediction (also called *projection*) consists in determining, according to one initial state of belief b and the description of a purely ontic action α , the new state of belief b' resulting from the application of α in b . The transformation of a state of belief into another one by an action is called *progression*; noted $b' = \text{prog}(b, \alpha)$. Of course, the formal definition of *prog* depends on the nature of the space of the beliefs (static and dynamic), thus it depends on the chosen representation of uncertainty. In the simplest case (that of classical planning) where belief states are perfect, actions are deterministic and always achievable, each $\text{prog}(\cdot, \alpha)$ is an application mapping a state to another one. In the binary nondeterministic model, a state s' is possible after the execution of α in the state of belief $b \subseteq \mathcal{S}$ if there exists a possible state s in the whole set of states corresponding to the initial belief b , such as s' is a possible result of α in s , i.e. $\text{prog}(b, \alpha) = \bigcup_{s \in b} R_\alpha(s)$. In the probabilistic model, the obvious choice is obtained by identifying the model of the process to a Markov chain: $\text{prog}(b, \alpha)$ is the probability distribution b' on \mathcal{S} defined by $b'(s') = \sum_{s \in \mathcal{S}} b(s) p(s'|s, \alpha)$ (where $p(\cdot|s, \alpha)$ is the probability distribution associated with R_α).

The second type of reasoning is *postdiction*. It consists in determining, according to one final state of belief b' and the description of a purely ontic action α which has just been carried out, the state of belief b before the action was done. This transformation of a state of belief into another, is sometimes also called *regression* or *weak regression*; noted $b = \text{reg}_w(b', \alpha)$. The weak regression corresponds to the progression by the reverse action of α (noted α^{-1}), which transition system $R_{\alpha^{-1}}$ is the reciprocal relation of the relation R_α ; thus it holds that $\text{reg}_w(b', \alpha) = \text{prog}(b', \alpha^{-1}) = \{s | R_\alpha(s) \cap b' \neq \emptyset\}$.

Postdiction must be distinguished from *goal regression*, also called *strong regression*, which is the reverse transformation of progression. It is defined only for the binary model³: given a belief state $b' \subseteq \mathcal{S}$ and a purely ontic action α , the aim is to find the belief state $b = \text{reg}_S(b', \alpha)$ such that $\text{prog}(b, \alpha) \subseteq b'$ and b is maximum for set inclusion; this belief state is the least informative state of belief (thus the least conjectural) which guarantees that the execution of α in it led to the goal b' .

³In the probabilistic model, there may not exist a unique probability distribution b on \mathcal{S} satisfying $b'(s') = \sum_{s \in \mathcal{S}} b(s) \cdot p(s'|s, \alpha)$, $b'(s')$ with $p(s'|s, \alpha)$ being known for all s, s' and α .

Let us notice that $reg_S(b', \alpha) \subseteq reg_w(b', \alpha)$ with the particular case $reg_S(b', \alpha) = reg_w(b', \alpha)$ when α is deterministic.

Progression and regression are two key processes of reasoning for *planning* (see Chapter 10 of Volume 2), which consists in determining the actions to carry out to make evolve the system as the agent wishes it (for example, get as close as possible to a reference trajectory in the case of the supervision, or to reach a goal state in the case of classical planning). On the other hand, postdiction has little interest for planning itself (because if b is obtained as a possible postdiction from b' with action α , it is not guaranteed that by carrying out the action α one would again obtain the state b' , while strong regression guarantees it by definition).

2.2.2 Prediction and Postdiction with Epistemic Actions

The progression of a belief state by an epistemic action depends on the nature of the reasoning process. In the case of a supervision process or a diagnosis, the agent reasons online and thus has access to all the observations coming from the actions *feedback* during its reasoning; thus it is enough to define the progression of a belief state by an observation, which is related to *belief revision* (see Chapter 14 of this volume). In the binary model, the progression of one belief state $b \subseteq \mathcal{S}$ by an observation o after having carried out the action α is $b \cap S(o)$, where $S(o) = \{s \mid o \in O_\alpha(s)\}$; while in the Bayesian model, the revision of b by o is the probability distribution $b(\cdot | S(O))$.

The *filtering* process consists in determining the new state of belief b' , given an initial belief state b , an action α , and an observation o resulting from the execution of α . In the binary model, this new belief state is simply $prog(b, \alpha) \cap S(o)$. In the Bayesian model, the probability distribution b' obtained after having carried out α and observing o is $b'(s') = \frac{p(o|s', \alpha) \cdot \sum_{s \in \mathcal{S}} b(s) \cdot p(s'|s, \alpha)}{\sum_{s'' \in \mathcal{S}} (p(o|s'', \alpha) \cdot \sum_{s \in \mathcal{S}} b(s) \cdot p(s''|s, \alpha))}$; it is the formula expressing the revision of the beliefs by the *feedback* in the partially observable Markov decision processes (see Chapter 10 of Volume 2).

In the case of a planning process, where the aim is to build an off line plan and to reason on its effects, the progression of a belief state by an epistemic action is in general not a unique belief state, but a set of such states (one for each possible observation, since the actual observation cannot be known off line). In the binary model, $prog(b, \alpha)$ is the set of belief states $\{b \cap S(o) \mid b \cap S(o) \neq \emptyset\}$ for o varying in Ω . For sake of shortness, we do not give details on regression by epistemic actions.

2.2.3 Event Abduction

The third type of reasoning is *event abduction*. It concerns reasoning on the event which took place between two successive time points t and $t + 1$, starting from the description of the possible events and from the description of the belief states at time t and $t + 1$ ⁴. If the event in question is exogenous, this reasoning is called *explanation*.

As for planning, progression and goal regression are two key processes for event abduction: in planning, one must choose the actions to be carried out to make the system evolve as desired starting from its current state; in event abduction, the objective is to determine which event α led the system to evolve as it did between t and $t + 1$ (even if this evolution was not desirable). In the binary model, to compute such α consists in searching among the possible events those satisfying $b' \subseteq prog(b, \alpha)$ (or equivalently $b \supseteq reg_F(b', \alpha)$).

2.2.4 Scenario Extrapolation

More generally, these types of reasoning, that we defined in a context where there is only one change stage (thus two time points), take place in situations where the horizon is unspecified, and where the input information is a complex *scenario* describing a partial trajectory of the system (at each time point, some information may be available about the occurrence of an action and/or an observation). In the typical case where no action was carried out and where the user wants to find the events (or more simply, the elementary changes) which occurred at each time point, the process is called *extrapolation*.

Another situation is when one seeks to recognize some trajectories among a set of reference trajectories in order to predict the events that will occur and/or the states that the system will reach; this process is called *scenario monitoring* or *scenario recognition*.

⁴A more complex abduction problem consists in reasoning not only on the event which took place, but also on the system states at time points t and $t + 1$, on which one wishes to obtain more precise beliefs.

The sequence of observations can also contain action occurrences [Dupin de Saint Cyr, 2008; Delgrande and Levesque, 2012; Hunter and Delgrande, 2015] (scenari are also called narratives or histories) and the two previous tasks of completing or recognizing some trajectories can be done in this more complex context. These tasks involve both prediction, postdiction and event abduction in situation that can be pervaded with uncertainty (fallible knowledge, erroneous perception, exogenous actions, and failed actions).

A crucial aspect of the reasoning about change approaches in artificial intelligence is that they assign a prominent role to *inertia* : by default, the system tends to remain static, and the changes other than those which are directly caused by action occurrences are rare, this is why one seeks to minimize them. This assumption is crucial if one wants to reason about action in presence of uncertainty without losing too much information. Very often, reasoning about change amounts to *minimize change* ; we will come back on this subject when we will approach the languages for reasoning about action. Indeed, according to the way actions are represented, there exist numerous ways of carrying out the progression of a state or the regression of a formula (encoding a set of states) by an action. Concerning the temporal or dynamic logic representations, progression and regression can be computed via some formula transformations (in particular, conjunction and forgetting). The use of change minimization principles is often proposed as a means to solve the frame problem (cf. Section 3.1), but it seems henceforth admitted that it is rather necessary to set up processes which remove the solutions containing abnormal changes (not caused by actions) than processes which minimize them.

This idea to focus on abnormal changes rather than on maximising inertia is well in accordance with the approaches that reason on a world under continuous change where the agent should adapt its action model when a surprise occurs (discrepancy between what is observed and what was expected). This kind of research is more related to the domain of planning in a dynamic world and particularly in the context of goal driven autonomy agents (GDA) that must reason about partially observable domains with a partial knowledge about available actions [Molineaux and Aha, 2014; Dannenhauer et al., 2016; Dannenhauer and Cox, 2017].

3 Reasoning about Action: Languages

3.1 Problems related to the Representation of Actions

In the majority of real problems, the system is naturally described by some variables, called *state variables*, that represent some features about some objects, etc. In this case, a state of the system corresponds to the description of a value for each one of the state variables, these values being able to change with the course of time. State variables are usually called *fluents*, a fluent describes a dynamic property of the system. Obviously, the number of possible states is exponential in the number of variables. The *explicit* description of the effects of the actions, which consists in specifying *in extenso* the functions R_α , becomes then unfeasible in practice, and is somewhat unnatural, because the user is obliged to describe the actions state by state. Similar considerations apply to the description of the correlations between states and observations and with the computation of the operations of progression, regression, etc.

However, there often exist much more economic and natural ways to represent the effects of the actions. For example, let us consider the action to “flip a switch” which causes the alternative “lighting on” or “off” of a bulb. If the representation of the problem requires to take into account the “on”/“off” states of 10 bulbs, then 1024 states of the system will have to be considered (all possible configurations of the 10 bulbs) in order to describe one flip action, whereas this action only causes the change of state of one particular bulb. To describe such action, one rather wants to be limited to indicate that it changes the state of this bulb and, implicitly, that it leaves the other bulbs in their current state.

Action languages were built precisely to this aim: obtaining representations of the effects of the actions which are both more economic (or more compact) and more natural. The problem of preventing the user from explicitly describing the fluents that are not modified by an action in the various possible contexts is known as *the frame problem* [McCarthy and Hayes, 1969]. It is indeed a problem involved in the choice of a representation of the actions (and not of a modeling problem, i.e. the problem does not rely on the choices of the fluents used to model the system but on the coding of actions in general).

In the same vein, one may face a problem that is the dual of the frame problem, known as *ramification problem* [Finger, 1987] which is solved when the action language makes it possible to avoid describing explicitly all the fluents that an action modifies, directly or indirectly, in the various possible contexts. Following up on the previous example, each flip of the switch causes the lighting on/off of the associated

bulb, then the room where the bulb is located becomes enlightened and consequently one can settle there for reading. This derived fact is a consequence of the execution of the action but it is not natural, when the action is described, to specify it directly: it results rather from a (static) law which expresses that when a room is lit, one can practice the reading there.

When one deal with action representation, the *qualification* problem [McCarthy, 1977] is also often evoked; this problem expresses the incapacity to describe all the pre-conditions that guarantee to obtain the “normal” effect of an action. To deal with this problem, it is first necessary to circumscribe the world with the individuals and the properties explicitly present in the representation; for example, flipping the switch when the associated bulb is off will cause the lighting of this one only if the conflict between Bordures and Syldaves did not cause the destruction of the electric line feeding the house. From our point of view, this problem is not intrinsic within the action representation, it occurs more primarily as soon as the modeling phase starts and simply reflects the difference existing between a situation of the physical world and a representation of this one, which necessarily abstracts it. However, in order to give the pre-conditions of an action, this restriction to the situations that have a representation in the language does not remove the need for reviewing all the situations in which the action is carried out normally. Solving the qualification problem means being able to state the “natural” pre-conditions of an action without having to describe explicitly the list of all the values of the fluents which allow the action to normally take place.

Once actions are represented, it is necessary to build algorithms allowing the computation of the basic operations (progression, regression, etc). The choice of an action language thus depends, on the one hand, of its more or less natural aspect, on the other hand, of its compactness (or space efficiency), and finally, of the complexity of the basic operations when the actions are represented in this language (its computational efficiency).

There exist many action languages which were developed and studied by the community. They can be gathered in several families, according to the nature of the mathematical objects that they use (propositional or first order logic formulas, temporal or dynamic logic formulas, Bayesian networks, state automats, etc). Giving an exhaustive panorama would be too long and little digest. We will thus only sketchily present the languages which received the most attention from the community, and which are sufficiently representative of the range of the existing languages. Each following sub-section approaches a particular language, or a family of languages, by briefly giving its specificities.

3.2 The Situation Calculus

From an historic perspective the *situation calculus* introduced by McCarthy and Hayes [1969] is the first formalism devoted to reasoning about actions. The definitions given by these authors enabled them to set the basic concepts (presented higher) on reasoning about change and action. The situation calculus is a typed first order logic language with equality, whose types are fluents, states (called situations), actions and objects. In order to simplify the presentation, here we only consider propositional fluents, which have one situation as single argument; we do not mention the objects of the world. Thus, $\neg P(S_0)$ express that the fluent P is false in the situation S_0 . S_0 denotes the state of the system at the initial time point of the horizon. For situations and actions we need both variables (denoted respectively s, \dots and x, \dots) and constants (denoted respectively S, \dots and A, \dots). The function *do* applies to a situation and an action and returns a situation. Thus, the formula $\neg P(S_0) \wedge P(do(A_1, S_0))$ expresses that P is false in S_0 and true in $do(A_1, S_0)$, i.e. in the situation obtained by applying the action A_1 in S_0 . The formula $\forall s \neg P(s)$ expresses that P is always false. The formula $\forall s ((\forall x \neg P(do(x, s)) \leftrightarrow x = A_0)$ expresses that A_0 is the single action which guarantees to make P false in any state where it is applied.

McCarthy and Hayes set a general representation framework enabling them to represent actions by their pre-conditions and their effects (represented by logical formulas). Many approaches were then proposed in order to characterize the “good” consequences of these formulas. Initially, all the authors bet on *change minimization* in order to restrict the set of models so that the properties resulting from the inertia principle can be deduced without having to mention them explicitly. This was accomplished thanks to a second order logic formula, and various circumscription policies were studied for this purpose (the reader can refer to [Moinard, 2000] for a review). McCarthy [1986] and then Hanks and McDermott [1986] used the circumscription of abnormality predicates (by considering that a fluent must persist unless otherwise explicitly indicated) within the framework of the situation calculus. However, there are some examples where circumscription does not give the expected result. One of most famous is the *Yale Shooting Problem* proposed by Hanks and McDermott: someone is alive in the initial situation, and one carry out successively the three actions “Load”, “Wait” then “Shoot”. The action “Shoot” is described

by the formula: $\forall s, (\text{loaded}(s) \rightarrow (\text{Abnormal}(\text{Alive}, \text{Shoot}, s) \wedge \neg \text{Alive}(\text{do}(\text{Shoot}, s))))^5$. The fact that, by default, the fluents are persisting is described by the second order logic formula $\forall f, s, a, ((f(s) \wedge \neg \text{Abnormal}(f, a, s) \rightarrow f(\text{do}(a, s)))^6$. The circumscription of the predicate *Abnormal* makes it possible to obtain a logical model in which the person is alive at the initial time point and dead (non alive) after the action “Shoot”. However, another model is possible: the one where the rifle unloaded itself during “Waiting” and the person is still alive after “Shoot”. Circumscribing the *Abnormal* predicate does not allow for preferring the first model to the second one because the two models have incomparable sets of abnormalities w.r.t. set inclusion (in the first model, it is “Alive” which is abnormal in the presence of the action “Shoot”; in the second one, it is “Loaded” which is abnormal w.r.t. “Wait”). Chronological ignorance, proposed by Shoham [1988] and consisting in preferring models where the changes occur the latest, allows one to obtain a satisfactory answer for this example. But this last *ad hoc* approach was challenged by other examples that it handles badly [Sandewall, 1995; Friedman and Halpern, 1994].

Another solution suggested by Lifschitz and Rabinov [1989] is to impose that all the fluents that are modified by an action are systematically non inert when this action is carried out. This idea is close to the solution, proposed by Castilho, Gasquet and Herzig [1999], to use a dependence relation between an action and the atoms on which it may act. The reader can refer to [Sandewall, 1995] for an excellent synthesis of all these works.

In short, approaches based on change minimization are based on *non-monotonic* logics and are very complex; they are not able to deduce all the intuitive consequences that are expected from a description of a set of actions and an initial situation.


The situation changed with the publication of what was called *Reiter’s solution* to the frame problem [Reiter, 1991]. Reiter suggests a *monotonic* solution based on *successor state axioms* (SSA). These axioms must be given for each fluent *P* (which is equivalent to an assumption of complete information about the conditions of change of truth value of a fluent) and they have the following form:

$$\forall s, x \quad (P(\text{do}(x, s)) \leftrightarrow \gamma_P(x, s))$$

where $\gamma_P(x, s)$ is a formula which does not contain the function symbol *do* and which can only contain S_0 as situation constant. Thus, the SSA for *P* describes the conditions under which *P* is true after an action has been performed, in function of what was true before.

Let us consider the Toggle-switch example [Lifschitz, 1990]:

In a room, the light is on only if both switches are up or both down. Initially, the switch *a* is up and the switch *b* in down, the light is thus off, someone toggles the switch *a*.



The fluents are U_a (“the switch *a* is up”) and U_b (“the switch *b* is up”). In this example, the SSA for fluent U_a can be written:

$$\forall s, x \quad U_a(\text{do}(x, s)) \leftrightarrow ((\neg U_a(s) \wedge x = T_a) \vee (U_a(s) \wedge x \neq T_a))$$

where T_a is the action to toggle the switch *a*, i.e., flip its position.

Reiter explains that using Successor States Axioms is a solution to the frame problem because one can reasonably expect the size of the set of SSA to be in the order of the number of fluents (which contrasts with the size of the explicit description of the frame axioms that would be in the order of the number of fluents set multiplied by the number of actions).

According to Reiter, quantification over actions is the key solution to the frame problem. As we will show in Section 3.4, the assumption of complete information about the conditions under which fluents change their truth value (translated by the \leftrightarrow in the SSA) allows Reiter to deal with the frame problem in a satisfactory way.

The presence of an SSA for each fluent allows for *regressing* formulas : atoms of the form $P(\text{do}(\alpha, \sigma))$ (where α and σ are terms built with variables, constants and the function *do*) are replaced by the right

⁵If the rifle is loaded in the situation *s* then the fluent “Alive” is abnormal (i.e., non persistent) when the action “Shoot” takes place in *s* and the person will not be alive any more in the resulting situation.

⁶If the fluent is not abnormal with respect to an action then it keeps its value after the execution of this action.

member of the SSA for P , by applying first the suitable substitution ; this process is reiterated until complete elimination of the function do . By construction, the formula thus obtained only relates on the initial state S_0 .

For example, the formula

$$U_a(do(T_a, do(T_a, S_0)))$$

is first replaced by:

$$(\neg U_a(do(T_a, S_0)) \wedge T_a = T_a) \vee (U_a(do(T_a, S_0)) \wedge T_a \neq T_a)$$

which can be simplified into $\neg U_a(do(T_a, S_0))$. In a second step, this last formula is replaced by $\neg(\neg U_a(S_0) \wedge T_a = T_a) \vee (U_a(S_0) \wedge T_a \neq T_a)$ which can be simplified into $U_a(S_0)$. We have thus proven by regression that the switch is up after two executions of T_a if and only if it is up in the initial state S_0 .

In order to decide whether the application of the action α in the state S_0 leads to a state in which ψ holds, it is enough to decide if the formula $\phi(S_0) \rightarrow \psi(do(\alpha, S_0))$ is valid. The regression of $\psi(do(\alpha, S_0))$ results in a formula $\psi'(S_0)$. If the argument S_0 is eliminated, we obtain the propositional formula $\phi \rightarrow \psi'$, whose validity can be checked by using a suitable prover.

This solution was combined with epistemic logic [Scherl and Levesque, 2003], which gives a formalism close to dynamic logic, described in Section 3.4. Moreover the framework of the situation calculus with a Successor State Axioms has been recently used by [Batusov and Soutchanski, 2018] for causal ascription.

3.3 Propositional Action Languages

A weak point of the approaches based on the situation calculus is the difficulty of their algorithmic implementation. For this reason, researchers have also developed approaches based on propositional logic, which can benefit from off-the-shelf ASP or SAT solvers (see Chapters 4 and 5 of Volume 2).

In action languages based on propositional logic, action effects are represented by local rules specifying only the fluents that change, possibly together with the conditions under which they change. Let F be a finite set of fluents. The states of \mathcal{S} are the propositional interpretations over F , that is, $\mathcal{S} = 2^F$.

The most basic action language is arguably *STRIPS* [Fikes and Nilsson, 1971], where an action is represented by a precondition and its effects (see Chapter 10 of Volume 2), a precondition being a conjunction of literals and an effect being a consistent set of literals.

To encode the light switch example, one may take as set of fluents $F = \{U_a, U_b\}$, where U_a (resp. U_b) is true (resp. false) when switch a (resp. b) is on (resp. off). An action with conditional effects such as T_a ('switch a ') can be written as

$$(U_a \mapsto \neg U_a) \wedge (\neg U_a \mapsto U_a).$$

The right member l of each rule of form $c \mapsto l$ is a direct action effect, which applies if and only if the corresponding condition is satisfied in the state in which the action is performed. Thus, applying $c \mapsto l$ in state s leaves s unchanged if s does not satisfy c and enforces the truth of l otherwise, leaving other fluents unchanged. This applies to each rule⁷. Thus, applying T_a in state s leads to change the truth value of U_a in s , as we expect. Importantly, such an action description rule is not a classical logical formula, and in particular, \mapsto is not material implication. Indeed, a STRIPS action α can be seen as a *constraint* linking the state of the world *before* it is performed and the state of the world *after* it has been performed. In particular, $c \mapsto l$ is not equivalent to $\neg l \mapsto \neg c$.⁸

One of the limits of the STRIPS language is the impossibility to express static laws. These laws are however needed for the ramification problem to be dealt with. For instance, in the previous example, one may want to introduce a new fluent L expressing that "the light is on". With standard STRIPS, integrating this new fluent would require to modify all actions by specifying what happens to L . This solution is not reasonable when the number of fluents is large. A way to cope with this lack of expressiveness consists in encoding actions with a set of *basic* fluents on which the available actions act directly (U_a and U_b in the example); the fluents that are not basic are called *derived* fluents. Progression is first computed as in classical STRIPS, and then there is one additional step so as to take the static laws into account and make some inferences on derived fluents. Thus, to compute the progression of state s by an action, one starts

⁷A pathological case is when the conditions of rules leading to complementary literals are conjointly satisfied in s ; in such a case, the progression is undefined; this can reflect an error when specifying the representation of the action, or the fact that s is impossible (and in this case corresponds to an implicit static law).

⁸If they were equivalent, then the encoding of action "Shoot" by $Loaded \mapsto \neg Alive$ in the *Yale Shooting Problem* would be equivalent to $Alive \mapsto \neg Loaded$, meaning that shooting on a living person results on the gun being magically unloaded (and the person staying alive...)

by projecting s on the basic fluents; then one performs the progression of this projection, and finally the obtained state is completed using the static laws. In the switch example, one may take as static law

$$((U_a \wedge U_b) \vee (\neg U_a \wedge \neg U_b)) \leftrightarrow L$$

where U_a and U_b are basic and L is derived. The progression of state $\{U_a, \neg U_b, \neg L\}$ by action T_a is thus $\{\neg U_a, \neg U_b, L\}$.

There are four main problems with STRIPS : it does not allow for representing (a) non-determinism, (b) static causal relations between fluents (as discussed in the previous paragraph), (c) concurrent actions, and (d) epistemic actions. To cope with this lack of expressiveness, more sophisticated action languages have been developed, both in the planning community (with ADL [Pednault, 1989] and PDDL [Ghallab et al., 1998]) and in the knowledge representation community. We will now focus on the languages stemming from the latter community.

In the 70s and 80s, the knowledge representation community used to think of actions as simple rules linking action preconditions and action effects. Subsequently, some researchers suggested that prediction could be computed using *minimization of change*, so as to impose that, by default, fluents that are not concerned by the action should persist (these fluents, of course, do not need to be specified in action effects, so as to cope with the frame problem). Then, since the 90's, minimization of change was progressively replaced by the use of propositional languages based on *causal implication*. The solutions of [Reiter, 1991], [Lifschitz and Rabinov, 1989] and [Castilho et al., 1999] for solving problems occurring with minimization of change consist in expressing dependencies between an action and its effects. This very principle has been implemented in works using *causal implication* (see Chapter 9 of this volume), which is distinct from material implication since it is meant to express these dependencies.

Some approaches using causal implication make use of the situation calculus [Stein and Morgenstern, 1994], [Lin, 1995]. Others use the modality C [Geffner, 1990; Giordano et al., 1998; Turner, 1999] or equivalently, define a new connective \Rightarrow [Giunchiglia et al., 2004]. Yet others define influence relations between fluents [Thielscher, 1997]. The main feature of these approaches is that they distinguish the fact of being true from the reason for being true, and use this distinction for computing the expected effect of actions for prediction or planning.

We give now some details about the action language \mathcal{A} proposed by Gelfond and Lifschitz [1993]. In this language, an action is expressed by means of *conditional causal rules* of the form

$$\text{if } c \text{ then } \alpha \text{ causes } l,$$

where α is an action name, c a conjunction of literals (omitted when it is equivalent to \top), and l a literal. A set of causal rules defines a deterministic transition system between states. Thus, the action α defined by the causal rules

$$\text{if } p \wedge q \text{ then } \alpha \text{ causes } \neg p, \text{ if } \neg p \wedge q \text{ then } \alpha \text{ cause } p \text{ and } \alpha \text{ cause } q$$

corresponds to the transition system R_α defined by $R_\alpha(pq) = R_\alpha(\bar{p}q) = \bar{p}q$ and $R_\alpha(\bar{p}q) = R_\alpha(pq) = pq$. An action α described by such causal rules corresponds to a *propositional action theory* Σ_α , expressing α by means of propositional symbols F_t and F_{t+1} , with $F_t = \{f_t | f \in F\}$ and $F_{t+1} = \{f_{t+1} | f \in F\}$, where f_t represents fluent f at time t , that is, before action α has been performed, and f_{t+1} represents f at time $t+1$, after action α has been performed. The causal rules are translated into Σ_α according to the following principle: fluent f is true at $t+1$ if and only if one of these two conditions holds: (a) it was true at t and the state at t does not satisfy any precondition of a causal rule whose conclusion is $\neg f$, or (b) it was false at t and the state at t satisfies the precondition of a causal rule whose conclusion is f . One finds here the principle at work in the situation calculus, which we called 'Reiter's solution' in Section 3.2.

Formally, let $\Gamma(f)$ (respectively $\Gamma(\neg f)$) the disjunction of all preconditions of rules whose conclusion is f (respectively $\neg f$); then Σ_α is the conjunction of all the formulas

$$f_{t+1} \leftrightarrow \Gamma(f)_t \vee (f_t \wedge \neg \Gamma(\neg f)_t)$$

for $f \in F$. Thus, the action theory Σ_α corresponding to the action α previously described by its causal rules is

$$\Sigma_\alpha = (p_{t+1} \leftrightarrow ((\neg p_t \wedge q_t) \vee (p_t \wedge \neg(p_t \wedge q_t)))) \wedge (q_{t+1} \leftrightarrow \top),$$

which simplifies into

$$\Sigma_\alpha = q_{t+1} \wedge (p_{t+1} \leftrightarrow (p_t \leftrightarrow \neg q_t)).$$

An extension of language \mathcal{A} is language \mathcal{C} [Giunchiglia and Lifschitz, 1998], which allows for expressing executability conditions and static rules, independently of any action, such as

Outside \wedge \neg Umbrella \wedge Umbrella causes \neg Dry,

that are also taken into account in the corresponding action theory. For instance, consider action Go-out with a unique causal rule

Go-out causes Outside;

the corresponding action theory, taking into account the previous static rule, is

$$\Sigma_{\text{Go-out}} = \text{Outside}_{t+1} \wedge (\text{Umbrella}_{t+1} \leftrightarrow \text{Umbrella}_t) \wedge (\text{Rain}_{t+1} \leftrightarrow \text{Rain}_t) \\ \wedge (\text{Dry}_{t+1} \leftrightarrow \text{Dry}_t \wedge (\text{Umbrella}_t \vee \neg \text{Rain}_t)).$$

Non-determinism can be expressed in several different ways, explored independently in different papers:

- by *complex effects*, such as

$$\alpha \text{ causes } (p \leftrightarrow q),$$

a choice that is at the heart of belief update, cf. Section 4;

- by *disjunction of effects*, which are similar to nondeterministic union in dynamic logic, cf. Subsection 3.4), such as

Toss causes Heads or causes \neg Heads;

- by recursive causal rules, which is a more technical solution that we will not discuss here.

Some action languages (such as language \mathcal{C}) also have *concurrency*, whereas others have *epistemic actions*, thus enabling the distinction between facts and knowledge; thus, the action of testing whether the fluent f is true or false is represented by the causal rule

$$\alpha \text{ causes } \mathbf{K}f \text{ or causes } \mathbf{K}\neg f,$$

where \mathbf{K} is the knowledge modality of epistemic logic S5 (see in particular [Herzig et al., 2003]).

Progression and *regression* can be applied directly in these languages. A belief state, in the binary uncertainty model, is a nonempty set of states, and can thus be represented by a consistent propositional formula. Progression and regression map a consistent formula and an action to a formula (which is always consistent in the case of progression and weak regression). The progression of formula φ by action α consists first in taking the conjunction of φ_t (expressing that φ is true before the action) and Σ_α , and then in forgetting in $\varphi_t \wedge \Sigma_\alpha$ all variables f_t , i.e., in deriving the strongest logical consequence of $\varphi_t \wedge \Sigma_\alpha$ independent of the variables f_t (see for instance [Lang et al., 2003]). Weak regression is computed similarly: the weak regression of ψ by α is the result of forgetting the variables f_{t+1} in $\psi_{t+1} \wedge \Sigma_\alpha$. The strong regression of ψ by α is obtained by computing the minimal conditions guaranteeing that the application of α will lead to a state satisfying ψ . Thus, in the previous example, the progression of Dry \wedge Umbrella by Go-out is (up to logical equivalence)

$$\text{Outside} \wedge \text{Umbrella} \wedge \text{Dry},$$

and the progression of \neg Umbrella by Go-out is

$$\text{Outside} \wedge \neg \text{Umbrella} \wedge (\text{Rain} \rightarrow \neg \text{Dry}),$$

whereas the weak regression of Dry \wedge Rain by Go-out is

$$\text{Umbrella} \wedge \text{Rain} \wedge \text{Dry}.$$

3.4 Dynamic Logic

There are other possible ways of representing actions and dealing with the corresponding problems. *Dynamic logic* is a formalism initially known in theoretical computer science for reasoning about program execution. In addition to Boolean operators, its language contains *modal operators* of the form $[\alpha]$, where α is a program. The combination of such an operator with a formula results in a formula of the form $[\alpha]\phi$, read ' ϕ is true after every execution of α '. Instead of a program, one may assume that α is an event or an action. For instance, the action of toggling switch a can be described by the two effect laws $(\neg U_a \rightarrow [T_a]U_a)$ and $(U_a \rightarrow [T_a]\neg U_a)$.

In the context of dynamic logic, an important aspect of reasoning about actions that was dealt with first in [Herzig and Varzinczak, 2007] concerns the consistency of a domain description. It has been shown that for expressive action languages, beyond logical consistency, a good domain description should be modular, in the sense that effect laws describing the actions should not allow for deriving




Figure 1: The DAG of a dynamic Bayesian network

new static laws. For instance, the static laws $P_1 \rightarrow [A]Q$, $P_2 \rightarrow [A]\neg Q$ and $\neg[A]\perp$ together imply the static law $\neg(P_1 \wedge P_2)$; if this law is not deducible from the other static laws and only them, then these effect laws should be considered problematic.

Unlike in situation calculus, states are not explicit in dynamic logic. Although dynamic logic does not allow either for quantifying over actions, which is a key feature of Reiter's solution for the frame problem, it has been shown that this solution can be implemented in dynamic logic for the rather general case of explicit SSAs [van Ditmarsch et al., 2011]. In such SSAs, x must be the only action variable of $\gamma_P(x, s)$ and if an action constant A does not appear in $\gamma_P(x, s)$ then $\gamma_P(A, s)$ should not be equivalent to $P(s)$. These conditions are natural for a system satisfying inertia. An example of SSA not satisfying them would be $\forall s(\forall x P(do(x, s))) \leftrightarrow \neg P(s)$, which means that P is changed in every state (thus P is a non-inert fluent). Note that the formula $\gamma_{U_a}(x, s)$ in our example from paragraph 3.2 satisfies these conditions. In order to translate these SSAs in dynamic logic, one introduces *assignment actions* of the form $P := \phi$; such an assignment describes an action where P takes the truth value that ϕ had in the previous state. This allows for associating with each action constant A the following set of assignments:

$$\sigma_{SSA}(A) = \{P := simp(\gamma_P(A)) \mid P \text{ appears in } \gamma_P(x)\}$$

where $simp(\gamma_P(A))$ is obtained from $\gamma_P(x)$ by eliminating argument s , substituting x by A and simplifying the equalities. In our example from paragraph 3.2, after substituting x by T_a we obtain:

$$\sigma_{SSA}(T_a) = \{U_a := (\neg U_a \wedge T_a = T_a) \vee (U_a \wedge T_a \neq T_a)\}$$

which can then be simplified into

$$\sigma_{SSA}(T_a) = \{U_a := \neg U_a\}.$$

Each occurrence of an abstract action symbol A is replaced by the corresponding assignment. As shown in [van Ditmarsch et al., 2011], this constitutes a solution (in Reiter's sense) to the frame problem. Thus Reiter's solution is transferred to dynamic logic, without any need to quantify over actions. It is also shown that Reiter's solution can be combined with epistemic logic, thus bridging it with epistemic-dynamic logics (see Chapter 2 of this volume).

A recent work in dynamic logic is to investigate epistemic extensions that are suitable for conformant planning [Li et al., 2017] and more generally for multiagent epistemic planning [Aucher and Bolander, 2013; Bolander et al., 2015; Cooper et al., 2016]. An overview paper about combinations of logics of action with logics of knowledge and belief is [Herzig, 2015].

3.5 Dynamic Bayesian Networks

A *dynamic Bayesian network* [Dean and Kanazawa, 1989] is a Bayesian network (see Chapter 8 of Volume 2) in which the variables exist in as many copies as there are time points: for any fluent f and time step t there is a fluent f_t . For each time step t , there exists a Bayesian network linking the variables corresponding to t . Moreover, between these 'instantaneous' networks, the only allowed edges are those that are directed from past to future. The temporal directed acyclic graph (DAG) given on Figure 1, equipped with probabilities for each variable, at each time step, conditionally on the values of the parents of the variable, constitutes a dynamic Bayesian network.

If the system is Markovian, in order to determine completely the behavior of the system it is enough to know the probability distribution for x_t and the conditional probability distribution for x_{t+1} given x_t . The Markovian assumption can reasonably be made for many classes of systems. A Markovian

temporal DAG cannot admit an arc linking variables distant from more than one time step: by deleting the edge between x_{t-1} and x_{t+1} the diagram on the example below becomes Markovian, and a description restricted to time steps t and $t + 1$ suffices.

The truth value of a fluent f at a given time step can depend on its value at earlier time steps $t - \Delta$, which translates in probabilistic terms into the following “survival equation”:

$$Pr(f_t) = Pr(f_t | f_{t-\Delta}) \cdot Pr(f_{t-\Delta}) + Pr(f_t | \neg f_{t-\Delta}) \cdot Pr(\neg f_{t-\Delta}).$$

The conditional probability $Pr(f_t | f_{t-\Delta})$ is called *survival function*. The survival function represents the tendency of propositions to persist given all events that can make them false. A classical survival function is: $Pr(f_t | f_{t-\Delta}) = \exp^{-\lambda \cdot \Delta}$, which indicates that the probability that f persists decreases, from the last time step where f was observed to hold, at an exponential speed determined by λ .

If one has some information about events that can affect the truth value of the fluent, then the survival equation no longer fits. Generally, the probability that a proposition f is true in t is a function of:

- the probability $Pr(f_{t-\Delta})$ that it is true at $t - \Delta$
- the probability $Pr(do(f_t))$ of the occurrence of an event that makes f true at t
- and the probability $Pr(do(\neg f_t))$ of the occurrence of an event that makes f false at t .

From the standpoint of expressiveness, the interest of this class of probabilistic approaches for reasoning about change is that it allows for expressing numerical uncertainty on beliefs, observations, and causal laws (see also [Hanks and McDermott, 1994] and by [Pearl, 1988]). On the other hand, a problem is that it requires the specification of many prior probabilities, even if it is not always necessary to ‘solve’ the whole probabilistic network to determine the probability of a proposition: one can instead focus on a few key time steps (and key propositions). Note that the use of a dynamic possibilistic network allows one to reason without knowing precisely these probabilities [Heni et al., 2007].

4 Reasoning about Change: Update

Update is a research domain at the intersection of reasoning about actions (whence its presence in this chapter) and belief change (see Chapter 14 of this volume). It is a process for integrating into a belief base a modification of the state of the system that is explicitly specified by a propositional formula. More precisely, given a belief base K and a propositional formula α , the update of K by α is the progression of K by an action, or alternatively by an exogenous event whose occurrence is known and whose effect is α . Belief update is opposed to belief revision where a new piece of information about the system is integrated into a belief base about that system, under the hypothesis that the latter did not evolve. The distinction was clarified by Katsuno and Mendelzon [1991], although update was studied before [Keller and Winslett, 1985; Winslett, 1988], partially by scholars from the database community (see Chapter 3 of Volume 3).

The distinction between revision and update deserves to be clarified a bit more here. If the new piece of information (‘the input’) completes our beliefs about the world then it is not the world that has evolved, but only the agent’s beliefs about the world. (This may be due to the questioning of an erroneous information about the world or a new piece of information about the characteristics of the world.) In that case we have to perform a *revision*.⁹ Such a revision amounts to a simple addition (also called *expansion*) when the input is consistent with the beliefs; however, in case of inconsistency revision selects some beliefs that have to be rejected in order to restore consistence (see Chapter 14 of this volume). If the input characterizes an explicit evolution of the world (i.e., is the effect of an action or an exogenous event) then we speak of an *update*. The updated belief base describes the world after its evolution. The update therefore corresponds to a *progression*.

The difference can be illustrated by the following example [Morreau, 1992]. Suppose there is a basket containing either an apple or a banana. If we learn that in fact it does not contain bananas then our beliefs have to be revised and we deduce that the basket contains an apple. However, if we learn that the world has evolved in a way such that there is no banana in the basket any more (e.g. because somebody has

⁹As shown in [Friedman and Halpern, 1999], revision remains relevant even if the initial belief state and the new formula do not refer to the same time point, as long as there is a syntactical distinction (via some time-stamping) between a fluent at a time point and the same fluent at another time point: what matters for revision is not that the world is static, but that the propositions that are used to describe the world are static. This also explains that belief extrapolation also corresponds to a revision process [Dupin de Saint-Cyr and Lang, 2011].

performed the action of taking the banana out of the basket if it was there) then we have to update our beliefs, i.e., that now the basket is either empty or contains an apple.

Just as for revision, there does not exist a unique update operator that would suit all applications. It is therefore interesting to define criteria which determine which of these operators are ‘rational’, these criteria can be written under the form of rationality postulates. Paralleling Alchourrón, Gärdenfors and Makinson’s postulates characterizing ‘rational’ revision operators (the so-called AGM postulates) [1985], Winslett [1990] was the first to define postulates for update operators. These postulates inspired Katsuno and Mendelzon [1991] who defined a new set of postulates. Similarly to the AGM postulates (see Chapter 14 of this volume), they are related to the existence of a set of preorder relations about the states of the system, where with each state there is associated a preorder. In [Katsuno and Mendelzon, 1991], the authors implement an idea that had been put forward by Winslett: in order to update a belief base one may update each of the models of the base independently. Katsuno and Mendelzon’s contribution is the idea that each model has to be updated towards the ‘closest’ models (in the sense of the preorder associated to the original model). According to Katsuno and Mendelzon, an update operator is a function \diamond which given a formula \mathcal{K} representing beliefs about the world and a formula ϕ representing the information about the evolution of the world, returns a new formula $\mathcal{K} \diamond \phi$. The postulates they propose in order to characterize the ‘rational’ operators \diamond are the following:

- U1 $\mathcal{K} \diamond \phi$ implies ϕ .
- U2 If \mathcal{K} implies ϕ then $(\mathcal{K} \diamond \phi)$ is equivalent to \mathcal{K} .
- U3 If \mathcal{K} and ϕ are satisfiable then $\mathcal{K} \diamond \phi$ is satisfiable.
- U4 If \mathcal{K}_1 is equivalent to \mathcal{K}_2 and ϕ_1 is equivalent to ϕ_2 then $\mathcal{K}_1 \diamond \phi_1$ is equivalent to $\mathcal{K}_2 \diamond \phi_2$.
- U5 $(\mathcal{K} \diamond \phi) \wedge \psi$ implies $\mathcal{K} \diamond (\phi \wedge \psi)$.
- U6 If $\mathcal{K} \diamond \phi_1$ implies ϕ_2 and $\mathcal{K} \diamond \phi_2$ implies ϕ_1 then $\mathcal{K} \diamond \phi_1$ is equivalent to $\mathcal{K} \diamond \phi_2$.
- U7 If \mathcal{K} is complete then $(\mathcal{K} \diamond \phi_1) \wedge (\mathcal{K} \diamond \phi_2)$ implies $\mathcal{K} \diamond (\phi_1 \vee \phi_2)$.
- U8 $(\mathcal{K}_1 \vee \mathcal{K}_2) \diamond \phi$ is equivalent to $(\mathcal{K}_1 \diamond \phi) \vee (\mathcal{K}_2 \diamond \phi)$.
- U9 If \mathcal{K} is complete and $(\mathcal{K} \diamond \phi_1) \wedge \phi_2$ is satisfiable then $\mathcal{K} \diamond (\phi_1 \wedge \phi_2)$ implies $(\mathcal{K} \diamond \phi_1) \wedge \phi_2$.

U1 stipulates that ϕ is a piece of information describing the world after its evolution (this is one of Winslett’s postulates). U2 says that if ϕ was already true in all states of the system before the update then the system does not evolve. This is the postulate requiring that inertia has always to be preferred to spontaneous evolution. It is however not always desirable because it forbids the existence of transitory states, i.e., states within which the system may not stay because it immediately evolves towards other states. U3 expresses that a consistent representation of the system and of its evolution can always be updated in a consistent way (which was also one of Winslett’s postulates). However, systems may exist where there is no transition between two states: for example, when $\phi = \text{dead}$ and $\alpha = \text{alive}$ then one may wish the update to fail. So this postulate is not always desirable. U8 (also one of Winslett’s postulates) means that the update is defined as a progression operator. U9 is a restriction of the converse of U5. We refer to [Dubois et al., 1995; Herzig and Rifi, 1999] for a more detailed critique of these postulates.

The following representation theorem relates these postulates to the existence of a set of preorder relations between states of the world:

Theorem 1 (Katsuno, Mendelzon) \diamond satisfies U1, U2, U3, U4, U5, U8, U9¹⁰ if and only if for every $\omega \in \Omega$ there is a total preorder \leq_ω such that

- (1) $\forall \omega' \in \Omega, \quad \omega <_\omega \omega'$ (\leq_ω is “faithful”);
- (2) $\text{Mod}(\mathcal{K} \diamond \phi) = \bigcup_{\omega \models \mathcal{K}} \{\omega' \models \phi \text{ such that } \forall \omega'' \models \phi, \omega' \leq_\omega \omega''\}$.

Item (1) means that for each model ω of \mathcal{K} , the models of the update of ω by ϕ are the models of ϕ that are closest to ω w.r.t. \leq_ω , and (2) means that the set of models of the update of \mathcal{K} by ϕ is the union of the sets of models resulting from the update of each model of \mathcal{K} by ϕ (which follows directly from postulate U8).

Numerous update operators were proposed in the literature. Thanks to the above theorem they can be defined by associating with each state a faithful total preorder relation between states. In practice, such a set of preorders is a way of *minimizing change*. For example, Winslett defined a relation \leq_ω^{PMA} between states that she called ‘Possible Models Approach’, abbreviated PMA. It is based on the function $\text{diff}_{\text{PMA}}(\omega_1, \omega_2)$ (the set of variables whose value differs between the two states ω_1 and ω_2): $\omega_1 \leq_\omega^{\text{PMA}} \omega_2 \Leftrightarrow_{\text{def}} \text{diff}_{\text{PMA}}(\omega_1, \omega) \subseteq \text{diff}_{\text{PMA}}(\omega_2, \omega)$. This relation is faithful and therefore defines an update operator. The corresponding update operator can also be characterized in terms of independence (the logical consequences of \mathcal{K} that are independent of ϕ persist) [Marquis, 1994]. In the examples of

¹⁰If U6 and U7 are used instead of U9 then the theorem gives us a faithful preorder that is only partial.

[Morreau, 1992], the initial beliefs are $\mathcal{K} = (banana \wedge \neg apple) \vee (apple \wedge \neg banana)$, so there are two models $\omega_1 = \{\neg apple, banana\}$ and $\omega_2 = \{apple, \neg banana\}$. When the agent then learns (φ) that somebody took the banana if it was there (update by $\neg banana$), the states of the system representing the information φ are ω_2 and $\omega_3 = \{\neg apple, \neg banana\}$. The updated base $\mathcal{K} \diamond_{\text{PMA}} \varphi$ can be computed by taking the union, for all models ω of \mathcal{K} , of the models of φ that are closest to ω . Here, the model of φ that is closest to ω_1 for the relation $\leq_{\omega}^{\text{PMA}}$ is ω_3 ; the model of φ closest to ω_2 is ω_2 itself (because $\leq_{\omega}^{\text{PMA}}$ is faithful). So the set of models of $\mathcal{K} \diamond_{\text{PMA}} \varphi$ is $\{\omega_2, \omega_3\}$. This means that after the update there is either an apple in the basket or the basket is empty.

The PMA relation has been refined by assigning priorities to some fluents [Winslett, 1988], which allows for handling fluents that do not persist in the same way. Other update operators go for increased expressiveness, e.g. the one proposed by Cordier and Siegel [1995] which allows for more or less priority transition constraints. These constraints take the form of pairs of formulas (φ, ψ) and are satisfied by a pair of models (ω, ω') when ω satisfies φ and ω' satisfies ψ . Then ω' is considered closer to ω than ω'' if the transition (ω, ω') violates less priority constraints than the transition (ω, ω'') .

Updates *à la* Katsuno and Mendelzon (and in particular Winslett’s PMA [1988] but also Forbus’ operator [Forbus, 1989]) are built on minimization of change. However, minimization of change is not always desirable for updates. In particular, Herzig and Rifi [1999] have shown that the approaches building on minimization of change do not allow updates by disjunctions; more formally, an update operator satisfying the Katsuno-Mendelzon postulates cannot handle disjunctions correctly, the culprit being postulate U5.

For that reason, several scholars studied update operators that are not built on minimization. They in particular studied a family of update operators that is based on the concept of *dependence*. Such updates of a belief base β by a formula α consists in first forgetting in β “all information concerning α ” (leaving the truth values of the variables that are not concerned by the update unchanged), and then adding α to the result. It remains to work out what “all information concerning α ” means. Such a kind of relevance is induced by a dependence relation between formulas: α concerns β if and only if β depends on α . This approach is general because the notion of dependence between formulas can vary.

Most of the dependence-based approaches to update consider that the dependence relation is expressed first between formulas and propositional variables, and can then be extended to a dependence relation between formulas: α and β are dependent if and only if there is at least one variable on which α and β are dependent. Examples of such update operators can be found in [Herzig, 1996], [Doherty et al., 1998] et [Herzig and Rifi, 1999]. This principle allows for remediating several counterintuitive aspects of minimization-based approaches and moreover is generally of lower computational complexity. A slight drawback is however that it is too little conservative: too much information of the initial base is forgotten. This can be counterbalanced by replacing the dependence relation between formulas and variables by a dependence relation between formulas and literals [Herzig et al., 2013].

As an update by a formula α can be viewed as a progression by a particular action whose effect is α (“to make α true”) (see a discussion in [Lang, 2007]), it makes sense to situate update w.r.t. propositional action languages. We start by observing that STRIPS is a particular case of both formalisms, corresponding to an update by conjunctions of literals. Axiom U8—which requires that the update of a set of models is the union of the update of the individual models—is exactly the definition of the progression of a belief state by an action. The two paradigms however differ in the variety of available actions: on the one hand, update offers the possibility of taking into account disjunctive effects (representing a unique but imperfectly known effect) and more generally effects consisting of arbitrary propositional formulas. On the other hand, action languages allow for conditional effects such as

(if Heads then flip-coin causes \neg Tails, if \neg Heads then flip-coin causes Tails),

nondeterministic effects such as

Toss-coin causes Heads or causes \neg Heads,

concurrent effects such as

if G and D are actions consisting in lifting the left and the right side of a table and if a glass of water is on the table then

G causes Spilled, D causes Spilled, G concurrently with D cause \top

as well as static causal rules allowing for ramifications. Approaches aiming at unifying the potentialities of various approaches are not numerous. Some update approaches take ramification into account by resorting to integrity constraints [Doherty et al., 1998] or allow for nondeterministic updates [Brewka and Hertzberg, 1993], or conditional or concurrent updates [Herzig et al., 2001]. However, an embedding of

Winslett's and Forbus' update operator and of Dalal's revision operator into dynamic logic was recently provided in [Herzig, 2014].

Update corresponding to the progression of an action, there exists a generalization (rightly called *generalized update*) that enables both revision and event abduction [Boutilier, 1998]. Generalized update allows for example for handling the following scenario: an agent wakes up in the morning and believes that the lawn is dry just as it was when she went to sleep. She subsequently observes that the lawn is wet, which first of all leads to a revision of her beliefs, then to the abduction of an event (it rained), and finally to an update by the effects of the event (the road is wet, too). Belief extrapolation [Dupin de Saint-Cyr and Lang, 2011] and other related formalisms such as [Berger et al., 1999] only handle the abduction of events. Finally, update can be viewed as an ordinal form of Lewis's *imaging* operator [Dubois and Prade, 1993] as well as the predictive phase of the Kalman filter [Cossart and Tessier, 1999; Benferhat et al., 2000].

Goldszmidt and Pearl [1992] were also interested by accounting for revision and update at the same time. They reason about a set of default rules and are of the causal kind, from which they deduce an order on the pairs of states of the world that they are filtering according to the input. If the last operator is a revision by φ then the pairs of states where the final state satisfies φ see their plausibility increased. In the case of an update by φ , one has to perform a revision by the dummy action $do(\varphi)$.

The contributions of Winslett, and of Katsuno and Mendelzon, are important for two reasons. First of all, they established a clear distinction between revising and updating a belief base. Second, they elaborated a set of postulates guaranteeing that a rational update is related to the existence of a set of preorders between the possible states of the system. Katsuno and Mendelzon, and Winslett, implicitly opted for the particular case where the fluents are by default inertial (they only change if an action or an event occurs that changes them). There is a further implicit hypothesis embodied by postulate U3 (Winslett's MB4): asserting that any update can be performed means that the input is always consistent with the possible evolution of the world.

Recently, belief change (including belief update) within the framework of fragments of propositional logic has gained attention. A propositional fragment simply is a subset of a propositional language which has some valuable properties (typically, from the computational side) but is not fully expressive w.r.t. propositional logic (some propositional formulas do not have any equivalent representation in the fragment). For instance, the Horn CNF fragment is the set of CNF formulas where each clause is Horn, i.e., it contains at most one positive literal. It is well-known that the satisfiability of any Horn CNF formula can be decided in linear time but that some propositional formulas (e.g. the clause $a \vee b$) cannot be turned into equivalent Horn CNF ones. Other fragments which are often considered are the Krom one (the set of all CNF formulas where each clause is binary) and the affine fragment (the set of all conjunctions of exclusive-or clauses), and each of them offers the same tractability property as the Horn one w.r.t. the satisfiability issue and the same limitations as to expressiveness. In order to preserve the tractability benefits, when a belief base from a given fragment has to be updated, it is expected that the updated base belongs to the same fragment. However, update operators satisfying all the Katsuno-Mendelzon postulates (especially Winslett's PMA and Forbus' operator) do not ensure this property. This calls for a notion of refinement of an update operator for a given fragment, which warrants that the result of any update is in the fragment when the initial base is in the fragment as well. Any refined operator is required to approximate the behavior of the operator considered at start (especially, leading to the same updated base as it when this base fits in the fragment), the price to be paid being the loss of some rationality postulates. A constructive approach to such refinements of update operators has been introduced in [Creignou et al., 2015], and the Katsuno-Mendelzon postulates satisfied by the refined operators identified as well.

Several tentatives were also made to extend update to more expressive frameworks:

- *ASP*: Slota and Leite [2010] adapt Katsuno & Mendelzon's postulates to logic programs [Slota and Leite, 2010]. They also define update operators for hybrid belief bases [Slota et al., 2011; Slota, 2012]. Such hybrid bases are made up of an ontology component, expressed in the language of the description logic ALCIO (ALC with inverse and nominals), and a rule component, expressed in the language answer-set programming under the stable semantics (see Chapter 4 of Volume 2). Update operators are studied in particular for the strong equivalence semantics for ASP as well as for hybrid belief bases [Motik and Rosati, 2010].
- *Belief states*: Lang, Marquis and Williams [2001] define update operators over epistemic states. In addition to beliefs, such epistemic states, represented by orders on worlds, allow for expressing the relative plausibility of beliefs. The authors extend the class of dependence-based update operators

to epistemic states. Baral and Zhang [2005] generalize update so as to distinguish between facts and knowledge, as in the epistemic logic S5. Their process is called *knowledge update* and allows one to account for the effects of epistemic actions by an *updating* the epistemic formulas describing the agent’s beliefs. Such a framework takes the viewpoint of a modeler agent O who reasons about the belief state of another agent ag. For example, the update of an S5 model by $K_{ag}\varphi$ means that O updates her beliefs about ag’s beliefs; the mental state of ag is seen by O as part of the external world, and the update by $K_{ag}\varphi$ corresponds to an action whose effect is to make $K_{ag}\varphi$ true (for example, the action of telling ag that φ is true).

- *Description logics* (see Chapter 6 of this volume): Liu et al. [2011] update assertions of an “ABox” (that is, the factual component of the belief base). They highlight an expressiveness problem that arises in that framework: sometimes the expected result of an update cannot be encoded by an assertion of the basic description logic ALC. For example, the assertion

$$\text{mary} : \text{Person} \sqcap \exists \text{child_of} . \text{Person} \sqcap \forall \text{child_of} . (\text{Person} \sqcap \text{Happy})$$

expresses that every child of Mary is happy. If one updates the ABox containing this information by the fact that Peter becomes unhappy, i.e., by the assertion $\text{peter} : \text{Person} \sqcap \neg \text{Happy}$, then one has to take two possible cases into account: the case where Peter is among Mary’s children and the case where he is not. Intuitively, the principle of minimal change requires that the result of the update is on the one hand the new piece of information (Peter is unhappy) and on the other hand the fact that every child of Mary either has the property of being happy or has the property of being Peter. The latter assertion (i.e., Mary’s children are either happy or called Peter) cannot be expressed in ALC: it requires an extension by object names. The extended logic (called ALCO) allows for writing

$$\text{mary} : \text{Person} \sqcap \exists \text{child_of} . \text{Person} \sqcap \forall \text{child_of} . (\text{Person} \sqcap (\text{Happy} \sqcup \{\text{pierre}\})).$$

It turns out that almost all description logic have similar expressiveness problems. If we allow for object names as concepts as done in ALCO, this wipes out the distinction between ABox assertions (which are about particular objects) and TBox concept inclusions which should not be about particular objects. This is unsatisfactory because the distinction is one of the very basic ideas of description logics.

- *Action descriptions*: Eiter et al. [Eiter et al., 2010] define a framework for minimal change of action descriptions that they call “action description updates”.
- *Abstract argumentation*: researchers in that domain (see Chapter 13 of this volume) are also interested by update and more generally change operators. An abstract argumentation system is represented by a graph whose vertices are arguments and whose edges are attacks between arguments. Some authors [Boella et al., 2009; Cayrol et al., 2010; Liao et al., 2011; Booth et al., 2013; Coste-Marquis et al., 2013] are interested by the impact of a change (by adding / withdrawing arguments or attacks) on such systems. Baumann and Brewka [2010; 2012] introduced the notion of *enforcement*, which is very similar to the notion of update (cf. [Bisquert et al., 2013; Dupin de Saint Cyr et al., 2016]) because the idea is to minimally modify an argumentation system in a way such that it satisfies a given goal (usually expressed in terms of arguments that should be accepted). They define a preference relation between argumentation systems, which is similar to preference relations between models in classical update.

5 Conclusion

Reasoning about action and change is one of the oldest topics in Artificial Intelligence. Since 1995, the topic is the subject of a biennial workshop *International Workshop on Nonmonotonic Reasoning, Action and Change (NRAC)* held in conjunction with the IJCAI (*International Joint Conference on Artificial Intelligence*) conference.

Several periods followed, during which the researchers were interested in conceiving several formal settings for modeling the important tasks related to reasoning about action: STRIPS first, then approaches based on minimal change, and then approaches based on *successor state axioms*. To this variety of formal settings corresponds a variety of languages for representing and reasoning about action: propositional logic, situation calculus, dynamic logic, graphical models (among others).

Reasoning about action and change has close connections with other areas of Artificial Intelligence, including non-monotonic reasoning, belief change, reasoning under uncertainty, planning (and in particular Markov decision processes) ; it also has links with control theory (more precisely, Kalman filtering and discrete event systems).

References

- Alchourrón, C., Gärdenfors, P., and Makinson, D. (1985). On the logic of theory change : partial meet contraction and revision functions. *J. of Symbolic Logic*, 50:510–530.
- Aucher, G. and Bolander, T. (2013). Undecidability in epistemic planning. In Rossi, F., editor, *IJCAI 2013, Proceedings of the 23rd International Joint Conference on Artificial Intelligence, Beijing, China, August 3-9, 2013*, pages 27–33. IJCAI/AAAI.
- Baral, C. and Zhang, Y. (2005). Knowledge updates: Semantics and complexity issues. *Artif. Intell.*, 164(1-2):209–243.
- Batusov, V. and Soutchanski, M. (2018). Situation calculus semantics for actual causality. In *Proceedings of the Thirty-Second AAAI Conference on Artificial Intelligence, New Orleans, Louisiana, USA, February 2-7, 2018*.
- Baumann, R. (2012). What Does it Take to Enforce an Argument? Minimal Change in Abstract Argumentation. In *Proc. European Conf. on Artificial Intelligence (ECAI'12)*, pages 127–132.
- Baumann, R. and Brewka, G. (2010). Expanding Argumentation Frameworks: Enforcing and Monotonicity Results. In *Proc. Int. Conf. on Computational Models of Argument (COMMA'10)*, pages 75–86.
- Benferhat, S., Dubois, D., and Prade, H. (2000). Kalman-like filtering in a possibilistic setting. In *Proc. European Conf. on Artificial Intelligence (ECAI'00)*, pages 8–12.
- Berger, S., Lehmann, D. J., and Schlechta, K. (1999). Preferred history semantics for iterated updates. *J. of Logic and Computation*, 9(6):817–833.
- Bisquert, P., Cayrol, C., Dupin de Saint Cyr, F., and Lagasquie-Schiex, M.-C. (2013). Enforcement in Argumentation is a kind of Update. In Liu, W., Subrahmanian, V., and Wijsen, J., editors, *International Conference on Scalable Uncertainty Management (SUM), Washington DC, USA, 16/09/2013-18/09/2013*.
- Boella, G., Kaci, S., and van der Torre, L. (2009). Dynamics in argumentation with single extensions: Attack refinement and the grounded extension. In *Proc. Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS'09)*, pages 1213–1214.
- Bolander, T., Jensen, M. H., and Schwarzenrüber, F. (2015). Complexity results in epistemic planning. In Yang, Q. and Wooldridge, M., editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 2791–2797. AAAI Press.
- Booth, R., Kaci, S., Rienstra, T., and van der Torre, L. (2013). A logical theory about dynamics in abstract argumentation. In Liu, W., Subrahmanian, V., and Wijsen, J., editors, *International Conference on Scalable Uncertainty Management (SUM), Washington DC, USA, 16/09/2013-18/09/2013*.
- Boutilier, C. (1998). A unified model of qualitative belief change: A dynamical systems perspective. *Artificial Intelligence*, 98(1-2):281–316.
- Brewka, G. and Hertzberg, J. (1993). How to do things with worlds: On formalizing actions and plans. *J. of Logic and Computation*, 3(5):517–532.
- Castilho, M., Gasquet, O., and Herzig, A. (1999). Formalizing action and change in modal logic I: The frame problem. *J. of Logic and Computation*, 9(5):701–735.
- Cayrol, C., Dupin de Saint-Cyr, F., and Lagasquie-Schiex, M.-C. (2010). Change in abstract argumentation frameworks: Adding an argument. *J. of Artificial Intelligence Research*, 38:49–84.

- Cooper, M. C., Herzig, A., Maffre, F., Maris, F., and Régnier, P. (2016). A simple account of multi-agent epistemic planning. In Kaminka, G. A., Fox, M., Bouquet, P., Hüllermeier, E., Dignum, V., Dignum, F., and van Harmelen, F., editors, *ECAI 2016 - 22nd European Conference on Artificial Intelligence, 29 August-2 September 2016, The Hague, The Netherlands - Including Prestigious Applications of Artificial Intelligence (PAIS 2016)*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 193–201. IOS Press.
- Cordier, M.-O. and Siegel, P. (1995). Prioritized transitions for updates. In *Proc. European Conf. on Symbolic and Qualitative Aspects of Reasoning under Uncertainty (ECSQARU'95)*, pages 142–150.
- Cossart, C. and Tessier, C. (1999). Filtering vs. revision and update: Let us debate! In *Proc. European Conf. on Symbolic and Qualitative Aspects of Reasoning under Uncertainty (ECSQARU'99)*, pages 116–127.
- Coste-Marquis, S., Konieczny, S., Mailly, J.-G., and Marquis, P. (2013). On the revision of argumentation systems: Minimal change of arguments status. *Proc. TAFE*.
- Creignou, N., Ktari, R., and Papini, O. (2015). Belief update within propositional fragments. In *Proc. European Conference on Symbolic and Quantitative Approaches to Reasoning with Uncertainty (ECSQARU'15)*, pages 165–174.
- Dannenhauer, D., Munoz-Avila, H., and Cox, M. T. (2016). Informed expectations to guide gda agents in partially observable environments. In *IJCAI*, pages 2493–2499.
- Dannenhauer, Z. A. and Cox, M. T. (2017). Rationale-based visual planning monitors for cognitive systems. In *Proceedings of the Thirtieth International Florida Artificial Intelligence Research Society Conference*, pages 182–185. AAAI Publications.
- Dean, T. and Kanazawa, K. (1989). A model for reasoning about persistence and causation. *Computational Intelligence*, 5(2):142–150.
- Delgrande, J. P. and Levesque, H. J. (2012). Belief revision with sensing and fallible actions. In *KR*.
- Doherty, P., Lukaszewicz, W., and Madalinska-Bugaj, E. (1998). The PMA and relativizing minimal change for action update. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 258–269.
- Dubois, D., Dupin de Saint-Cyr, F., and Prade, H. (1995). Update postulates without inertia. In *Proc. European Conf. on Symbolic and Qualitative Aspects of Reasoning under Uncertainty (ECSQARU'95)*, pages 162–170.
- Dubois, D. and Prade, H. (1993). Belief revision and updates in numerical formalisms: An overview, with new results for the possibilistic framework. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'93)*, pages 620–625.
- Dupin de Saint Cyr, F. (2008). Scenario Update Applied to Causal Reasoning. In *International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 188–197, <http://www.aaai.org/Press/press.php>. AAAI Press.
- Dupin de Saint Cyr, F., Bisquert, P., Cayrol, C., and Lagasquie-Schiex, M.-C. (2016). Argumentation update in YALLA (Yet Another Logic Language for Argumentation). *International Journal of Approximate Reasoning*, 75:57–92.
- Dupin de Saint-Cyr, F. and Lang, J. (2011). Belief extrapolation (or how to reason about observations and unpredicted change). *Artificial Intelligence*, 175(2):760–790.
- Eiter, T., Erdem, E., Fink, M., and Senko, J. (2010). Updating action domain descriptions. *Artif. Intell.*, 174(15):1172–1221.
- Fikes, R. and Nilsson, N. (1971). Strips : A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 2:189–208.
- Finger, J. (1987). *Exploiting constraints in design synthesis*. PhD thesis, Stanford University, Stanford, CA.

- Forbus, K. (1989). Introducing actions into qualitative simulation. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'89)*, pages 1273–1278.
- Friedman, N. and Halpern, J. (1994). A knowledge based framework for belief change part II : revision and update. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'94)*, pages 190–201.
- Friedman, N. and Halpern, J. Y. (1999). Modeling belief in dynamic systems, part II: Revision and update. *J. of Artificial Intelligence Research*, 10:117–167.
- Geffner, H. (1990). Causal theories for nonmonotonic reasoning. In *Proc. National Conf. on Artificial Intelligence (AAAI'90)*, pages 524–530.
- Gelfond, M. and Lifschitz, V. (1993). Representing action and change by logic programs. *J. of Logic Programming*, 17:301–321.
- Ghallab, M., Howe, A., Knoblock, C., McDermott, D., Ram, A., Veloso, M., Weld, C., and Wilkins, D. (1998). The planning domain definition language. Technical report, AIPS-98 Planning Competition.
- Giordano, L., Martelli, A., and Schwind, C. (1998). Dealing with concurrent actions in modal action logics. In *Proc. European Conf. on Artificial Intelligence (ECAI'98)*, pages 537–541.
- Giunchiglia, E., Lee, J., Lifschitz, V., McCain, N., and Turner, H. (2004). Nonmonotonic causal theories. *Artificial Intelligence*, 153:49–104.
- Giunchiglia, E. and Lifschitz, V. (1998). An action language based on causal explanation: Preliminary report. In *Proc. National Conf. on Artificial Intelligence (AAAI'98)*, pages 623–630.
- Goldszmidt, M. and Pearl, J. (1992). Rank-based systems: A simple approach to belief revision, belief update, and reasoning about evidence and actions. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'92)*, pages 661–672.
- Hanks, S. and McDermott, D. (1986). Default reasoning, nonmonotonic logics, and the frame problem. In *Proc. National Conf. on Artificial Intelligence (AAAI'86)*, pages 328–333.
- Hanks, S. and McDermott, D. (1994). Modelling and uncertain world i : symbolic and probabilistic reasoning about change. *Artificial Intelligence*, 66:1–55.
- Heni, A., Ben Amor, N., Benferhat, S., and Alimi, A. (2007). Dynamic possibilistic networks: Representation and exact inference. In *Proc. IEEE Int. Conf. on Computational Intelligence for Measurement Systems and Applications (CIMSAS'07)*, pages 1–8.
- Herzig, A. (1996). The PMA revisited. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'96)*, pages 40–50.
- Herzig, A. (2014). Belief change operations: A short history of nearly everything, told in dynamic logic of propositional assignments. In Baral, C., Giacomo, G. D., and Eiter, T., editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Fourteenth International Conference, KR 2014, Vienna, Austria, July 20-24, 2014*. AAAI Press.
- Herzig, A. (2015). Logics of knowledge and action: critical analysis and challenges. *Autonomous Agents and Multi-Agent Systems*, 29(5):719–753.
- Herzig, A., Lang, J., and Marquis, P. (2003). Action representation and partially observable planning using epistemic logic. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'03)*, pages 1067–1072.
- Herzig, A., Lang, J., and Marquis, P. (2013). Propositional update operators based on formula/literal dependence. *ACM Trans. Comput. Log.*, 14(3):24:1–24:31.
- Herzig, A., Lang, J., Marquis, P., and Polacsek, T. (2001). Updates, actions, and planning. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'01)*, pages 119–124.
- Herzig, A. and Rifi, O. (1999). Propositional belief base update and minimal change. *Artificial Intelligence*, 115(1):107–138.

- Herzig, A. and Varzinczak, I. J. (2007). Metatheory of actions: beyond consistency. *Artificial Intelligence*, 171:951–984.
- Hunter, A. and Delgrande, J. (2015). Belief change with uncertain action histories. *Journal of Artificial Intelligence Research*, 53:779–824.
- Katsuno, H. and Mendelzon, A. (1991). On the difference between updating a knowledge base and revising it. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 387–394.
- Keller, A. and Winslett, M. (1985). On the use of an extended relational model to handle changing incomplete information. In *IEEE Trans. on Software Engineering*, volume SE-11:7, pages 620–633.
- Lang, J. (2007). Belief Update Revisited. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'07)*, pages 2517–2522.
- Lang, J., Liberatore, P., and Marquis, P. (2003). Propositional independence: Formula-variable independence and forgetting. *J. of Artificial Intelligence Research*, 18:391–443.
- Lang, J., Marquis, P., and Williams, M.-A. (2001). Updating epistemic states. In *Proc. Australian Joint Conf. on Artificial Intelligence (AI'01)*, pages 297–308.
- Li, Y., Yu, Q., and Wang, Y. (2017). More for free: a dynamic epistemic framework for conformant planning over transition systems. *J. Log. Comput.*, 27(8):2383–2410.
- Liao, B., Jin, L., and Koons, R. C. (2011). Dynamics of argumentation systems: A division-based method. *Artificial Intelligence*, 175(11):1790 – 1814.
- Lifschitz, V. (1990). Frames in the space of situations. *Artificial Intelligence*, 46:365–376.
- Lifschitz, V. and Rabinov, A. (1989). Things that change by themselves. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'89)*, pages 864–867.
- Lin, F. (1995). Embracing causality in specifying the indirect effects of actions. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pages 1985–1991.
- Liu, H., Lutz, C., Milicic, M., and Wolter, F. (2011). Foundations of instance level updates in expressive description logics. *Artificial Intelligence*, 175(18):2170–2197.
- Marquis, P. (1994). Possible models approach via independency. In *Proc. European Conf. on Artificial Intelligence (ECAI'94)*, pages 336–340.
- McCarthy, J. (1977). Epistemological problems of artificial intelligence. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'77)*, pages 1038–1044.
- McCarthy, J. (1986). Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence*, 28(1):1038–1044.
- McCarthy, J. and Hayes, P. (1969). Some philosophical problems from the standpoint of artificial intelligence. In *Machine Intelligence*, volume 4, pages 463–502.
- Moinard, Y. (2000). Note about cardinality-based circumscription. *Artificial Intelligence*, 119(1-2):259–273.
- Molineaux, M. and Aha, D. W. (2014). Learning unknown event models. In *AAAI*, pages 395–401.
- Morreau, M. (1992). Planning from first principles. In *Belief revision*, pages 204–219. Cambridge University Press.
- Motik, B. and Rosati, R. (2010). Reconciling description logics and rules. *J. of the Association for Computing Machinery*, 57(5).
- Pearl, J. (1988). Embracing causality in formal reasoning. *Artificial Intelligence*, 35:259–271.

- Pednault, E. (1989). Adl: Exploring the middle ground between strips and the situation calculus. In *Proc. Int. Conf. on Principles of Knowledge Representation and Reasoning (KR'89)*, pages 324–332.
- Reiter, R. (1991). The frame problem in the situation calculus: A simple solution (sometimes) and a completeness result for goal regression. In *Artificial Intelligence and Mathematical Theory of Computation: Papers in Honor of John McCarthy*, pages 359–380. Academic Press.
- Sandewall, E. (1995). *Features and Fluents*. Oxford University Press.
- Scherl, R. and Levesque, H. J. (2003). The frame problem and knowledge producing actions. *Artificial Intelligence*, 144(1-2).
- Shoham, Y. (1988). *Reasoning about Change - Time and Causation from the Standpoint of Artificial Intelligence*. MIT Press.
- Slota, M. (2012). *Updates of Hybrid Knowledge bases*. PhD thesis, Universidade Nove de Lisboa.
- Slota, M. and Leite, J. (2010). On semantic update operators for answer-set programs. In *Proc. European Conf. on Artificial Intelligence*, pages 957–962.
- Slota, M., Leite, J., and Swift, T. (2011). Splitting and updating hybrid knowledge bases. *Theory and Practice of Logic Programming*, 11(4-5):801–819.
- Stein, L. and Morgenstern, L. (1994). Motivated action theory: A formal theory of causal reasoning. *Artificial Intelligence*, 71(1):1–42.
- Thielscher, M. (1995). The logic of dynamic systems. In *Proc. Int. Joint Conf. on Artificial Intelligence (IJCAI'95)*, pages 1956–1962.
- Thielscher, M. (1997). Ramification and causality. *Artificial Intelligence*, 89(1-2):317–364.
- Turner, H. (1999). A logic of universal causation. *Artificial Intelligence*, 113(1-2):87–123.
- van Ditmarsch, H., Herzig, A., and de Lima, T. (2011). From Situation Calculus to Dynamic Epistemic Logic. *J. of Logic and Computation*, 21(2):179–204.
- Winslett, M. (1988). Reasoning about action using a possible models approach. In *Proc. National Conf. on Artificial Intelligence (AAAI'88)*, pages 89–93.
- Winslett, M. (1990). *Updating Logical Databases*. Cambridge University Press.