

Stable Reasoning

Pedro Cabalar¹, David Pearce², and Agustín Valverde³

¹ University of Corunna, Spain, cabalar@udc.es

² Universidad Politécnica de Madrid, Spain, david.pearce@upm.es

³ Universidad de Málaga, Spain, a_valverde@ctima.uma.es

Abstract. We give an account of stable reasoning, a recent and novel approach to problem solving from a formal, logical point of view. We describe the underlying logic of stable reasoning and illustrate how it is used to model different domains and solve practical reasoning problems. We discuss some of the main differences with respect to reasoning in classical logic and we examine an ongoing research programme for the rational reconstruction of human knowledge that may be considered a successor to the logical empiricists' programme of the mid-20th Century.

1 INTRODUCTION

For most of the 20th Century logic as a scientific discipline was dominated by the paradigm of Frege and Hilbert and the aim of providing mathematics with a secure logical foundation. This challenging goal continued to dominate research programmes in logic long after Gödel's incompleteness theorems showed that the initial expectations were unreachable. Above all the Frege-Hilbert paradigm was and still is based on a standardised, classical conception of logic, on the axiomatic method as the basis for reconstructing mathematical knowledge and on the method of deduction as the central element in logical theorem proving.

The paradigm of Frege and Hilbert was challenged by various critics throughout its lifetime but only in recent years has it been eclipsed by more progressive research programmes within computational logic. A lively and searing critique of the Frege-Hilbert paradigm was published by the logician Carlo Cellucci in 1998 [6]. Cellucci calls the standard paradigm *mathematical logic* and discusses at length many of its significant features. Besides its emphasis on providing a secure foundation for mathematics, key features include the prominence of the axiomatic method, the idea of theories as embodying certain truths and theorems as logical deductions from these. Cellucci questions many of the assumptions of the paradigm, especially the manner in which it focuses on closed conceptual systems and problems of justification, while paying almost no attention to problems of (mathematical) discovery that may involve hypothetical reasoning, induction, abduction, analogy, heuristics and other methods. As a contrast to the axiomatic method Cellucci devotes much attention to describing and motivating

the analytic method in mathematical discovery and problem solving⁴ and justifying the importance of treating open systems and fallible reasoning.

While Cellucci doesn't offer a detailed description of modern *computational logic*, it is evident that some of the positive features of the analytic method can be found in computational approaches to logic and logic-based programming languages. He himself cites with approval six different features of Prolog that mirror aspects of the analytic method and reasoning with open systems (while at the same time noting that there are other important features of the analytic method that are not captured in Prolog). Since the publication of [6] computational logic has made many advances. One of them has been the elaboration of a new approach to logic-based programming known as *answer set programming* or ASP and the development of its underlying logical paradigm which we will call *stable reasoning*. Unlike Prolog, this is not a fully-fledged programming language, but rather a general approach to logic-based problem solving that can also be efficiently implemented in (answer set) solvers.

We cannot claim that stable reasoning currently satisfies all the requirements proposed by Cellucci for constituting a new and wholly adequate paradigm for logic. For one thing, the focus in [6] is on mathematics and mathematical problem-solving and discovery, while stable reasoning has a broad range of applications to many areas of inquiry. However, we do suggest that

1. it is distinctly different from the Frege-Hilbert paradigm of mathematic logic;
2. its approach to problem-solving has significant points in common with the analytic method;
3. it is able to embrace various aspects and methods of discovery;
4. it is able to deal with dynamical and open systems.

In this paper we will discuss stable reasoning, focussing mainly on items 1 and 2 above, while mentioning 3 and 4 briefly towards the end.

Stable reasoning is a recent and novel approach to problem solving from a formal, logical point of view. An important difference compared to reasoning based on classical logic is that stable reasoning can take account of default assumptions and conclusions that follow from them. It is not based on two-valued classical logic, since this does not allow for the distinction between certain truth and truth-by-default or the kind of truth that can be assumed when there is no evidence to the contrary. To account for defaults in the setting of classical logic usually special syntactic devices are employed or one has to distinguish between different kinds of inference, some defeasible others not. Stable reasoning does not require any special inference rules or syntactic devices because it is based on a many-valued logic where precisely the distinction between certain truth and truth-by-default can be made. There is just one basic kind of negation and one kind of inference.

There is another important difference compared to the axiomatic tradition of Frege and Hilbert that dominated the methodology of formal logic for much of the 20th Cen-

⁴ Closely associated with Plato and other classical scholars.

tury. Stable reasoning is closer to what might be termed a problem-solving approach to formal reasoning. In the axiomatic tradition a mathematical or empirical domain is formalised by introducing a language or vocabulary and a set of sentences (axioms) of the language intending to capture once and for all the entirety of knowledge governing that domain. Mathematical theorems are inferred through logical deduction from the axioms. Predictions or explanations from empirical theories are also deduced from their axioms once the initial conditions of a system are specified.

By contrast, stable reasoning is problem-driven. One wishes to find a solution or several possible solutions to a certain problem. One describes the problem domain by specifying the entities and relations that govern the domain. This description need not be complete nor need it capture the entirety of knowledge once and for all, but merely offer hypotheses sufficient to produce adequate answers or solutions. Some descriptions may indeed be robust and reusable, others may be fragile and of temporary validity/use. As in the classical, axiomatic case, answers will be produced once certain facts or initial conditions are specified. These answers or solutions are in general model-based. They are produced not by means of logical deduction or inference from axioms plus initial conditions, but rather by computing a model or state of affairs that embodies the solution in some obvious way.

To illustrate the difference between the two orientations, consider the following well-known graph-theoretical problem.

Example 1 (Hamiltonian cycles). A *Hamiltonian cycle* of a graph G is a cyclic path that traverses each node in G exactly once. The same graph may have different Hamiltonian cycles or none at all. \square

Suppose we want to obtain the Hamiltonian cycles of graph G in Figure 1. To this aim we must decide:

- (a) how to represent the graph;
- (b) how to represent the cycles;
- (c) the type of logical reasoning task to obtain the desired solution.

Regarding (a), the obvious solution in Predicate Calculus is just using a pair of predicates, say $Node(x)$ and $Edge(x, y)$, to respectively describe the nodes and edges of G . Still, even at this elementary level, a first difference between the classical approach and stable reasoning already arises. In classical logic, if we just list the set of ground atoms:

$$\begin{aligned} &Node(0), \quad Node(1), \quad Node(2), \quad Node(3), \\ &Edge(0, 1), \quad Edge(1, 2), \quad Edge(1, 3), \quad Edge(2, 0), \\ &Edge(2, 3), \quad Edge(3, 2), \quad Edge(3, 0). \end{aligned} \tag{1}$$

this will leave free the possibility for other objects of becoming nodes or edges – we would have models where, for instance, $Edge(0, 2)$ could become true. In order to pre-




Fig. 1. A simple graph with a pair of Hamiltonian cycles.

cisely capture the nodes and edges we should actually use instead the pair of formulas⁵:

$$Node(x) \leftrightarrow x = 1 \vee x = 2 \vee x = 3 \vee x = 4 \quad (2)$$

$$Edge(x, y) \leftrightarrow (x = 0 \wedge y = 1) \vee (x = 1 \wedge y = 2) \vee (x = 1 \wedge y = 3) \vee (x = 2 \wedge y = 0) \vee (x = 2 \wedge y = 3) \vee (x = 3 \wedge y = 2) \vee (x = 3 \wedge y = 0) \quad (3)$$

In the stable reasoning approach, however, the set of facts (1) would suffice because predicates are subject to the so-called *Close World Assumption* (CWA). Informally speaking, this means that anything not explicitly stated will be false by default. We could, of course, still use the stronger version (2) \wedge (3) for stable reasoning, but the set of facts (1) is more flexible in the sense that new nodes or edges can be included by the simple *addition* of new formulas, rather than manipulating the existing ones (this feature is often referred as *elaboration tolerance* [20]). Predicates like $Node(x)$ or $Edge(x, y)$ whose extensions are defined by enumerating their lists of (true) atoms conform what is usually called the *extensional database* (a term inherited from Database theory) as opposed to those described by additional conditional formulas (or *rules* in Logic Programming) which receive the name of *intensional predicates*.

Another important feature about CWA is that, if necessary, we can remove this assumption in a selective way for any formula ϕ or, in particular, any atom $P(\bar{x})$. To this aim, it just suffices with adding an axiom like:

$$P(\bar{x}) \vee \neg P(\bar{x}) \quad (4)$$

Note that although (4) has the form of a classical tautology (the so-called *excluded middle axiom*), it is not a tautology for stable reasoning, since negation here has a different meaning. As said before, the effect of (4) is that CWA for predicate P is removed, so it behaves in a “classical” way. For instance, adding the axiom $Node(x) \vee \neg Node(x)$ would remove $Node$ to be false by default.

Let us move now to consider problem (b), that is, how to represent Hamiltonian cycles. As any of them visits all nodes in the graph, we obviously must refer to the set of edges involved in the cycle to differentiate one from each other. In classical

⁵ We assume that free variables, like x, y here, are implicitly universally quantified.

logic, we would typically represent sets of edges using some additional notation: assume, for instance, that standard set terms are allowed. Then, we would use some predicate, say $HamCycle(s)$, to represent that the set s of edges constitutes a Hamiltonian cycle for the graph, including a hypothetical set of axioms $\Gamma(HAM_G)$ that includes the graph description $(2) \wedge (3)$ and the meaning of predicate $HamCycle$. Under the axiomatic method, our problem (c) would then reduce to decide for which sets s of edges we can derive $\Gamma(HAM_G) \vdash HamCycle(s)$. For instance, in our example, we should conclude that this holds for $s = \{(0, 1), (1, 3), (3, 2), (2, 0)\}$ and for $s = \{(0, 1), (1, 2), (2, 3), (3, 0)\}$. It is worth to notice that as these two facts for predicate $HamCycle(s)$ are derived as theorems, they will be simultaneously true in *all* models of $\Gamma(HAM_G)$. In the general case, the extension of $HamCycle$ would contain the whole set of Hamiltonian cycles in *any* model of $\Gamma(HAM_G)$.

Under the stable reasoning approach, tasks (b) and (c) become model-oriented, rather than theorem-based. One of the main features of this methodology is that we identify each solution of the original problem with (a distinguished part of) each model of our logical representation. For instance, in our example, rather than collecting all cycles in a single predicate $HamCycle(s)$, our purpose is obtaining a *different model* per each possible Hamiltonian cycle of the original graph. Using this approach, we need not dealing with sets any more: we can just collect the edges *in* the current cycle (that is, the cycle represented inside the current model) as the extent of some predicate, call it $In(x, y)$. Thus, if we call again $\Gamma(HAM_G)$ to our logical representation of the problem, we are interested in obtaining the different models M such that $M \models \Gamma(HAM_G)$ so that each Hamiltonian cycle can be directly obtained by selecting some relevant information in each obtained model M . In our example, we should obtain a pair of models M_1 and M_2 such that:

$$M_1 \models In(0, 1) \wedge In(1, 3) \wedge In(3, 2) \wedge In(2, 0)$$

$$M_2 \models In(0, 1) \wedge In(1, 2) \wedge In(2, 3) \wedge In(3, 0)$$

and no other atoms for $In(x, y)$ hold in each model.

Let us consider now how to specify $\Gamma(HAM_G)$ under the stable reasoning approach. Typically, we would first consider models for all possible subsets of edges, and then include additional formulas that rule out the ones that do not correspond to Hamiltonian cycles. In a first attempt, we could include:

$$Edge(x, y) \rightarrow (In(x, y) \vee \neg In(x, y)) \quad (5)$$

$$In(x, y) \wedge In(x, z) \rightarrow y = z \quad (6)$$

$$In(y, x) \wedge In(z, x) \rightarrow y = z \quad (7)$$

together with a graph description, that is, an extensional database like (1).

As a first important remark, note that the consequent $In(x, y) \vee \neg In(x, y)$ in (5) has the form of an excluded middle formula like (4). As we explained before, this means that predicate $In(x, y)$ will not be subject to CWA, provided that x, y is a pair of nodes forming an edge. As a result, the effect of (5) is that we would have a model per each

possible subset of edges in the graph. Formula (6) (resp. (7)) specifies that we never pick two different outgoing (resp. incoming) edges for a given node x . These two formulas can also be represented by using \perp in the consequent (formulas like these receive the name of *constraints*):

$$\begin{aligned} In(x, y) \wedge In(x, z) \wedge y \neq z &\rightarrow \perp \\ In(y, x) \wedge In(z, x) \wedge y \neq z &\rightarrow \perp \end{aligned}$$

This first attempt (5)-(7), however, is not enough to capture Hamiltonian cycles yet. We can still get disconnected groups of edges like, for instance, $\{(0, 1), (2, 3), (3, 2)\}$ or even nodes that are not connected at all. To rule out these cases, we should further specify that any node can be *reached* from another one. Curiously, this property (graph reachability) is a well-known example of a problem that *cannot be represented* in classical First Order Logic. So, in fact, if we tried to represent Hamiltonian cycles using the classical axiomatic method, $\Gamma(HAM_G)$ should be represented in Second Order Logic. Under stable reasoning, however, we can easily capture reachability with an auxiliary predicate $Reach(x, y)$ defined with the pair of formulas:

$$In(x, y) \rightarrow Reach(x, y) \tag{8}$$

$$In(x, z) \wedge Reach(z, y) \rightarrow Reach(x, y) \tag{9}$$

Since $Reach(x, y)$ is subject to CWA, these two formulas behave as an inductive definition, as those frequently used in Mathematics. In other words, $Reach(x, y)$ is true if: (1) we have taken the edge $In(x, y)$; or (2) we took an edge to some z and we can reach y from that z , $Reach(z, y)$. Otherwise, $Reach(x, y)$ will be false, due to CWA. A predicate like this, whose extension is determined by some rules like those above in combination with CWA, receives the name of *intensional* predicate.

To complete our example, we would just further need to rule out models where we can find a pair of nodes that are not reachable one from each other. This can be done using the formula:

$$\neg \exists x, y (Node(x) \wedge Node(y) \wedge \neg Reach(x, y))$$

or its equivalent representation as a constraint:

$$Node(x) \wedge Node(y) \wedge \neg Reach(x, y) \rightarrow \perp \tag{10}$$

Again, it is worth to note that, while (10) is classically equivalent to:

$$Node(x) \wedge Node(y) \rightarrow Reach(x, y)$$

the latter has a quite different meaning under the stable reasoning approach – we would read it as a third option for defining $Reach(x, y)$ facts, and in this case, we would get that any pair of nodes would always be reachable, regardless the edges we had in the graph.




Fig. 2. Kinds of truth in stable reasoning.

Proposition 1. *Let $\Gamma(G)$ be the set of formulas (5)-(10) plus an extensional database for graph G like (1). Then, M is a stable model of theory $\Gamma(G)$ iff $\{(x, y) \mid M \models \text{In}(x, y)\}$ is a Hamiltonian cycle for G . \square*

2 A logic from first principles

Stable reasoning has very simple logical underpinnings. Its base logic can be constructed from first principles in a few easy steps and with only a few fundamental assumptions. The first idea is that formulas may be true or false, but that there are two kinds of truth: certain truth and truth-by-default. We can picture this as shown in the left hand side of the Figure 2.

Suppose that the disk represents the collection of true formulas; while outside the outer circle formulas are false. True formulas come in two kinds: those that are only weakly true or more particularly *true by default* lie within the unshaded, outer part of the disk, those lying within the shaded, inner circle are *certain*, or *true in the strong sense*. Truth in the weak or general sense covers both certain truth and truth by default, so the outer circle contains all the formulas of the disk, both inner and outer parts.

We can assign numbers to these different semantic properties, say 0 for falsity and 2 for certain truth. We assign the value 1 to formulas that lie within the outer but not the inner circle: they are only true in the weaker but not the stronger sense. The pictures looks like as shown in the right hand side of the figure 2.

Let us say that our base logic comprises, in the propositional case, the usual set of logical connectives $\{\wedge, \vee, \rightarrow, \neg\}$, standing for conjunction, disjunction, implication and negation. We may also make use of the falsum constant, \perp . There is an infinite set $Prop$ of propositional atoms from which formulas are constructed in the usual way, as in classical or intuitionistic logic. Let p, q, r, \dots stand for atomic propositions. As prescribed, the semantics is given by the three truth-values, 2, 1 and 0.

\wedge	0	1	2
0	0	0	0
1	0	1	1
2	0	1	2


\vee	0	1	2
0	0	1	2
1	1	1	2
2	2	2	2

\rightarrow	0	1	2
0	2	2	2
1	0	2	2
2	0	1	2


	\neg
0	2
1	0
2	0

Fig. 3. Truth tables for the connectives.

In the picture on the right p is certain, q and s are true by default and r and t are false. The values for complex formulas are assigned according to the usual meaning of connectives. In the picture $p \wedge q$ takes the value 1, since only p is certain, while q is not. On the other hand the fact that p is certain is sufficient for assigning $p \vee q$ and even $p \vee r$ the value 2. Since r is false, also $p \wedge r$ and $q \wedge r$ must be false, while $q \vee r$ takes the value 1.



Following this reasoning we can build the complete truth tables for conjunction and disjunction (see Figure 3); they are similar to those of other well-known logics.



Some care has to be taken with the values of formulas of form $A \rightarrow B$. In the picture on the left we would expect both $p \rightarrow r$ and $q \rightarrow r$ to be false since in both cases we have a true antecedent and false consequent. Notice that although q is not certain, we cannot put $q \rightarrow r$ in the inner circle since that would make it strongly true and yet false at the same time. Since r is false, evidently both $r \rightarrow q$, $r \rightarrow p$ and $r \rightarrow t$ are true, and even in the strong sense. Since p is certain, it is also true in the general sense and therefore $p \rightarrow q$ must be at least weakly true. But it is not certain, since q is not. By contrast, since p is certain, $q \rightarrow p$ is also. Lastly, consider the

case of two formulas q, s that lie only in the outer circle. Since neither atom is certain, but neither is false, the implications $q \rightarrow s$ and $s \rightarrow q$ can safely be placed in the inner circle of certainty.

The complete truth matrix for implication is given in the Figure 3. The case of negation is straightforward. Negating a false proposition such as r produces a certain proposition $\neg r$; while negating either kind of true formula produces a false one: in the picture both $\neg p$ and $\neg q$ will be false.

Equivalently we can regard negation to be definable, by $\neg p := p \rightarrow \perp$, where \perp takes the constant value 0. Then the table for negation is derivable from that for implication.

Our base logic is by no means new. It was first introduced by Heyting in his study of intuitionistic logic [11]. Shortly after it reappeared in Gödel's paper [10] showing that intuitionistic logic is not tabular (finite-valued). Heyting provided the truth matrices while subsequently Łukasiewicz [13] studied the logic in greater depth and gave the first axiomatisation based on the axioms and rules of the intuitionistic calculus extended by the addition of a weaker form of Peirce's Law, viz:

$$(\neg\alpha \rightarrow \beta) \rightarrow (((\beta \rightarrow \alpha) \rightarrow \beta) \rightarrow \beta). \quad (11)$$

Subsequently, the logic was studied by Smetanich [18] and Umezawa [19] who gave an alternative axiom to that of Łukasiewicz:

$$\alpha \vee \neg\beta \vee (\alpha \rightarrow \beta)$$

The completeness of this system was then proved by Hosoi [12].

Although Heyting made use of the logic as a technical device, it was clear from the start that it was of interest not only from a purely formal point of view. Łukasiewicz made a detailed comparison with his own 3-valued logic and Heyting already provided a natural interpretation of the third truth-value, claiming that it applies to a correct proposition that cannot be false but whose correctness cannot be proved. If we re-phrase this in terms of truth, then the interpretation 'true-by-default' is a natural one since it conveys the idea of a judgement that is accepted although not formally derivable in the system at hand.⁶

Our logic is known under a variety of names, we prefer to call it the logic of *here-and-there*, in symbols **HT**, whose etymology is easily explained by looking at an alternative way to describe the semantics. It is well-known that intuitionistic logic is complete for possible world models, that is triples of the form $\langle W, \leq, I \rangle$, where W is a non-empty set (of points, states or worlds), \leq is a partial ordering on W , and I is an interpretation function assigning a set of (verified) atoms to each $w \in W$, such that $I(x) \subseteq I(y)$ whenever $x \leq y$. The logic **HT** is also complete for a class of such models, but of an especially simple form: W comprises just two points, say h ('here') and t ('there'), with $h \leq t$. It follows that we can represent a model simply as an ordered pair $\langle H, T \rangle$ of sets of atoms, where $I(h) = H$ and $I(t) = T$. Evidently we always have $H \subseteq T$. These are then *here-and-there* models. Applying the usual semantics to evaluate formulas at worlds one can easily verify the correspondence to our earlier truth-tables. Note that in a here-and-there model $\langle H, T \rangle$, H represents the certain atoms having value 2, while T represents the non-false atoms (of value 1 or 2). Those in $T \setminus H$ are true-by-default, ie take the value 1. The complement of T in $Prop$ is the set of atoms that are false in the model, ie those corresponding to the value 0 in our truth tables. We adopt the usual convention and say that a formula φ is satisfied or holds in a model $\langle H, T \rangle$, in symbols $\langle H, T \rangle \models \varphi$, if it is strongly true or certain in the model, i.e. is satisfied at the world h .

⁶ Strictly speaking we should say that the base logic gives a monotonic approximation to reasoning by default. There is an additional aspect to defaults that emerges in the nonmonotonic extension of the logic.

As a simple example, consider how to evaluate an implication $p \rightarrow q$ in a model $\langle H, T \rangle$, for atomic p, q . By the possible world semantics, this formula is true at t if $q \in T$ or $p \notin T$. In terms of truth-values it means that $p \rightarrow q$ takes a value of at least 1 if $q \in T$ no matter what value is given to p . However in the case that $p \notin T$, then also $p \notin H$. This means that $p \rightarrow q$ is true also at h and so the value of $p \rightarrow q$ is then 2.

Much is known about **HT** and its properties. It is one of 7 superintuitionistic (SI) logics having the interpolation property, [14], and it is also the strongest SI-logic properly contained in classical logic, all other SI-logics being properly contained in it.

3 Basic stable reasoning

The core logic of stable reasoning is based on **HT** but is actually a non-monotonic extension of it. It can be characterised in terms of a preferential entailment relation in the sense of Shoham (1988) defined on here-and-there models. The preference relation can be explained by considering some features of defaults.

When we build up a partial description of the world or of our problem domain we specify a set of formulas that we suppose, at least hypothetically, to be true in the certain sense. Call this our ‘theory’. When we then consider the three-valued **HT**-models of this theory (assuming it is consistent), they in turn generally admit formulas that are true only in the weak or default sense. We typically have models $\langle H, T \rangle$ where $H \subset T$ and then some formulas will be true only in the default sense, for instance all the propositions in $T \setminus H$. Equally, by consistency, we will also have models of the form $\langle T, T \rangle$ where all truths are certain. Our preference condition is to select just those models $\langle T, T \rangle$ where all truths are certain *and our theory does not admit any model with uncertainty whose true atoms are exactly T* ; in other words there is no model of our theory of the form $\langle H, T \rangle$, where $H \subset T$. In this sense our theory justifies the choice of T by not accepting any model where T forms the set of truths in the general sense and some of them are only true in the weak sense.

So here is our definition in full. A model $\langle H, T \rangle$ of a theory \mathcal{T} is said to be an *equilibrium* model of \mathcal{T} , if (i) $H = T$ and (ii) for any H such that $H \subset T$, $\langle H, T \rangle \not\models \mathcal{T}$. The term equilibrium model derives from [15], and the ensuing logic associated with this is equilibrium logic. However there is complete agreement between equilibrium models and the stable models of logic programs as defined by Gelfond and Lifschitz (1988). That is to say, if the formulas of a theory have precisely the shape of rules of logic programs, then a set of atoms T is a stable model of the theory if and only if $\langle T, T \rangle$ is an equilibrium model of it.⁷ For this reason equilibrium logic can serve as a

⁷ In logical notation the rules of a (disjunctive) logic program have the form

$$b_1 \wedge \dots \wedge b_m \wedge \neg b_{m+1} \wedge \dots \wedge \neg b_n \rightarrow a_1 \vee a_2 \vee \dots \vee a_k$$

where the a_i, b_j are atoms. If $k = 1$ everywhere the program is said to be *normal*.

foundation for stable reasoning.⁸ Though they are defined differently, we often use the terms stable model, answer set and equilibrium model inter-changeably.

4 Implementation of stable reasoning

Just as classical deduction is implemented in automated provers such as Prover9 or SAT-solvers, so stable reasoning has been implemented in various systems, collectively known as *answer set solvers*. These have obtained a fairly high degree of efficiency and sophistication and can be used to model and solve real-world problems in domains such as software verification, security and configuration management, model checking, agent technologies, constraint satisfaction, reasoning for the semantic web, software synthesis from specifications, knowledge representation, data and information integration, planning and diagnosis. The inputs to answer set solvers are called *answer set programs*, the branch of computer science and programming dealing with these is called *answer set programming*, or ASP for short. A closely related domain is that of DATALOG and deductive database systems. Indeed data and information management is one of the principal application areas of ASP, and one that currently enjoys some commercial success.

While they are successful in supporting real-world problem solving, ASP systems do not implement deduction in a formal logic. They are not designed to deduce theorems or prove the correctness of logical inferences, even though the solutions they compute can be precisely understood in terms of formal models of a logical system of deduction. This correspondence to logic however was a more recent discovery and did not directly guide the initial ASP implementations. While in traditional logic programming, systems of logical inference came first and computer programming applications came after, in ASP this order was reversed. The theory preceded the computer implementations that in turn preceded the logic. However, nowadays the logic exerts a growing influence on the development and comprehension of new systems, especially extensions of the initial ASP family of languages. It is also fundamental to understanding stable reasoning from a foundational point of view.

5 Practical examples

Let us now consider some additional examples of typical commonsense reasoning tasks that can be easily represented under the stable reasoning approach. We also show in some cases how these examples can be implemented using existing solvers. For instance, Figure 4 shows a possible implementation of the Hamiltonian cycles problem

⁸ Since the late 1990s the stable model semantics has been systematically extended to embrace wider classes of formulas, more recently including arbitrary propositional and even first-order theories. All the widely accepted extensions have coincided with equilibrium logic.

for graph in Figure 1, using the input language of the DLV solver⁹. The column in the right shows the correspondence to each formula in $\Gamma(HAM_G)$. The translation of our example theory into DLV language follows some standard syntactic conventions from Logic Programming and, in this case, is quite straightforward. Variables begin with uppercase letters and predicates with lowercase letters. Implications are reversed, so that $\alpha :- \beta$ stands for $\beta \rightarrow \alpha$. When $\alpha = \perp$ it is just omitted. Finally, a comma represents a conjunction, the symbol ‘ \vee ’ represents a disjunction and all formulas are ended by a full stop.

```

node(0) . node(1) . node(2) . node(3) .
edge(0,1) . edge(1,2) . edge(1,3) .           (1)
edge(2,0) . edge(2,3) . edge(3,2) . edge(3,0) .

in(X,Y) v out(X,Y) :- edge(X,Y) .           (5)

:- in(X,Y) , in(X,Z) , Y!=Z .                (6)
:- in(Y,X) , in(Z,X) , Y!=Z .                (7)

reach(X,Y) :- in(X,Y) .                      (8)
reach(X,Y) :- in(X,Z) , reach(Z,Y) .         (9)

:- not reached(X,Y) , node(X) , node(Y) .     (10)

```

Fig. 4. Hamiltonian cycle representation $\Gamma(HAM_G)$ for the graph in Figure 1 written for DLV.

The only newly introduced feature with respect to the first order theory $\Gamma(HAM_G)$ is the use of an auxiliary predicate $Out(x, y)$ to replace the formula $\neg In(x, y)$ since DLV does not accept negation to the left of $:-$ operator. This replacement is correct provided that $Out(x, y)$ is not used elsewhere. Otherwise, to fully replace the original formula we should include the pair of rules:

```

out(X,Y) :- not in(X,Y) , edge(X,Y) .
:- in(X,Y) , out(X,Y) , edge(X,Y) .

```

which are logically equivalent to:

$$Edge(x, y) \rightarrow (Out(x, y) \leftrightarrow \neg In(x, y))$$

A fundamental property of stable reasoning we have not exploited yet in the previous example is the use of default negation for non-monotonic reasoning (NMR). To illustrate this concept, consider the following classical example in the NMR literature.

⁹ <http://www.dlvsystem.com>

We want to capture the default “birds typically fly” and the exception “penguins are birds but do not fly.” These two assertions can be respectively encoded as:

$$Bird(x) \wedge \neg CannotFly(x) \rightarrow Flies(x) \quad (12)$$

$$Penguin(x) \rightarrow Bird(x) \wedge CannotFly(x) \quad (13)$$

Suppose we have theory Γ' containing the two formulas above plus the ground atom $Bird(Tweety)$. As there is no evidence that *Tweety* is a penguin, we derive the literal $\neg Penguin(Tweety)$ by CWA. This falsifies the antecedent of (13), and so no evidence on $CannotFly(Tweety)$ can be obtained from (13). In fact, no evidence on $CannotFly(x)$ can be obtained from (12) either. This is because implication in stable reasoning acquires a kind of directionality. In particular, (12) should be read as a *definition rule* for $Flies(x)$, saying that it will hold when x is a bird for which we have no evidence on $CannotFly(x)$. As a result, we conclude $\neg CannotFly(Tweety)$ by CWA and then, $Flies(Tweety)$ from (12).

This simple example became a challenge for NMR approaches because a wrong predicate minimisation policy could easily lead to models where $\neg Flies(Tweety)$ was decided first and then $CannotFly(Tweety)$ was derived by applying (12) as the classically equivalent formula:

$$Bird(x) \wedge \neg Flies(x) \rightarrow CannotFly(x)$$

In stable reasoning, such an equivalence does not hold. In fact, the formula above has a quite different reading from (12): it states that a bird x cannot fly if we cannot find any evidence for $Flies(x)$.

To complete the example, consider now the extended theory

$$\Gamma'' = \Gamma' \cup \{Penguin(Tweety)\}.$$

Since we do obtain now evidence on $CannotFly(Tweety)$ from (13), we lose the justification for $Flies(Tweety)$ we obtained before from (12). As a result, $Flies(Tweety)$ is not derived any more. This shows the non-monotonic nature of stable inference, since the addition of a new formula, $Penguin(Tweety)$, has made a previous conclusion $Flies(Tweety)$ to be retracted.

The representation of this example in DLV is shown below:

```
flies(X) :- bird(X), not cannotfly(X).
bird(X)  :- penguin(X).
cannotfly(X) :- penguin(X).
bird(tweety).
penguin(tweety).
```

6 Stable reasoning, open systems and the analytic method

In this final section we look briefly at one of the directions in which stable reasoning has developed into a programme for the logical reconstruction of human knowledge in a practical setting. Then we return to some of the conditions suggested by Cellucci [6] that should be fulfilled by a logical paradigm to replace the axiomatic tradition of mathematical logic. We consider briefly some of the ways in which stable reasoning conforms to these requirements.

6.1 Gelfond's programme

The rise of the axiomatic method and formal reasoning based on classical logic in the 20th Century was closely linked to two philosophical schools based in central Europe: the Lwów-Warsaw School of Logic and Philosophy and the logical empiricist movement of the Vienna Circle. Particularly the latter school endorsed a programme of rational reconstruction of scientific and other forms of knowledge as well as adopting the idea of explicating philosophically important concepts typically by formalising them within a logico-mathematical system. The limits of the logical empiricist programme and the growing criticism it faced from the 1960s onwards are well-known and well-documented. Especially vulnerable was the ideal to reconstruct (ultimately all of) scientific knowledge in formal languages governed by the classical laws of deduction. A pivotal point of this criticism was the claim that logic and experience are not sufficient to explain the rationality of science, catalogue its methods and reconstruct the knowledge it generates.

With hindsight there is an irony to this criticism and to the ultimate downfall of the empiricist programme. It came about just at a moment in time when logic was about to undergo a radical change and a major shift in its boundaries. This change was led by Montague in natural language processing, by Hintikka in the study of propositional attitudes of knowledge and belief, by McCarthy in artificial intelligence and commonsense reasoning, by Simon in learning and scientific discovery, and by a range of applications in computer science and programming.

Stable reasoning has spawned a successor to the logical empiricist programme of rational reconstruction. It is a research programme that aims to reconstruct some of the most basic forms of human knowledge and to exploit this knowledge for practical problem solving. While logical empiricist efforts were largely theoretically oriented, this programme deals with a mix of theory and practice. It combines scientific and engineering knowledge of real systems (...) with practical human skills and abilities and commonsense reasoning. It deals with both static and dynamic domains. The scientist most closely identified with this programme is Michael Gelfond, also a co-founder of stable reasoning itself. His programme combines the physicalist language of engineering and physical systems with epistemic notions such as belief, agency and action. The new programme of rational reconstruction is much less self-conscious than its predecessor.

sor. The latter formed part of a manifesto with a clear philosophical, and sometimes political, message. It was stated and re-stated many times. The new programme is scarcely articulated and hardly known. Nevertheless its goals and methodology are largely clear, if sometimes buried in technical articles and lectures.¹⁰

Gelfond's programme for representing and reasoning about knowledge has two main objectives. First it aims to achieve an understanding of "basic commonsense notions we use to think about the world: beliefs, knowledge, defaults, causality, intentions, probability, etc., and to learn how one ought to reason about them." Secondly it aims "to understand how to build software components of agents – entities which observe and act upon an environment and direct its activity towards achieving goals." These goals shape the criteria used to evaluate and select languages for Knowledge Representation (KR). In particular he endorses four main adequacy criteria:

1. Clarity: the logical vocabulary should have a clear and intuitive meaning.
2. Elegance: the corresponding mathematics should be simple and elegant.
3. Expressiveness: the KR language should suggest systematic and elaboration tolerant representations of a broad class of phenomena of natural language, including belief, knowledge, defaults, causality and others.
4. Relevance: a large number of interesting computational problems should be reducible to reasoning about theories formulated in this language.

It is interesting to compare these criteria with the requirements that Carnap proposes for the adequate explication of concepts [5]. Clearly criterion 2 is close to Carnap's requirement for the task of concept explication that "the explicatum should be as *simple* as possible." Criterion 4, on the other hand, and to a somewhat lesser extent 3, are close to Carnap's idea that an explicatum is to be a *fruitful* concept, "that is, useful for the formulation of many universal statements." Although Elaboration Tolerance is a more modern idea, one may suppose that it would also feature among the properties of fruitfulness. Carnap's suggestion that the "explicatum is to be *similar to the explicandum*" in many cases of usage does not appear explicitly in Gelfond's list. However since Gelfond's aim is to reconstruct commonsense knowledge and practical reasoning one can assume that Carnap's requirement is one he would also endorse and is somehow implicit in his programme.¹¹

It is also revealing to consider some adequacy criteria for KR languages that Gelfond *rejects*. Among these are two that in the past were often considered sacrosanct in the AI community of Knowledge Representation. One is the idea that any KR language should be *supra-classical* ie. extend first-order classical logic. The other is the requirement of *efficiency* understood in a computational sense. We have seen already that the underlying logic of stable reasoning is not classical, and nor does it become so when additional features and functionalities are provided. On the other hand, the expressive-

¹⁰ But see specially [9] for an overview.

¹¹ All quotations above are from [5], Introduction.

ness of the basic language – a positive feature – also results in it being less efficient computationally than some other languages.¹²

6.2 Open conceptual systems

Let us return to the idea of open conceptual systems that we mentioned in the introduction. According to Cellucci there are many similarities between open conceptual systems and the notion of open physical systems. Here is a summary of some of the basic features of the former taken from [6] (pp. 313–315).

1. Open conceptual systems take account of the manner in which the solution of a problem is to be arrived at.
2. Unlike in the axiomatic approach there is no unifying idea that serves as a foundation once and for all.
3. The unifying impulse for problem solving is data driven rather than reductionist.
4. Open conceptual systems tackle the problem to be solved directly, from first principles. There is no *a priori* set of concepts and principles that precede the problem formulation.
5. The rules of the game are not given at the beginning and fixed once and for all, but may be introduced and changed during the course of the game.
6. Open systems are dialogical, since partially given information may be extended through interaction with other systems.
7. The rules of the system give only a partial and dynamically changing representation of knowledge.

Besides these characteristics, Cellucci emphasises the ampliative nature of logical inference, as well as the need to deal with global inconsistencies and incoherences. There are many other features of open systems and the analytic method discussed in [6] and a detailed examination would take us beyond the scope of this paper. Likewise, here we have described only some core features of stable reasoning and ASP. Many other features emerge in the practical development of systems and their application to problem solving. We conclude with a shortlist of some of the characteristics that may bear on Cellucci's challenge to develop an alternative logical paradigm.

- Stable reasoning in its basic form already deals with weak and strong exceptions to defaults. However to deal with contradictions that arise indirectly as consequences of default conclusions, an extension of ASP with consistency-restoring rules (CR-Prolog) was developed and applied by Balduccini and Gelfond [1]. This is essentially an abductive mechanism.

¹² Generally speaking the complexity of stable reasoning lies at the second level of the polynomial hierarchy. Nevertheless answer set solvers are relatively efficient in being able to deal with quite large amounts of data. Notice that from a representational point of view answer set programs are very efficient in being able to encode complex problems in a concise manner.

- Another feature of Gelfond’s programme has been the aim to reason about the degrees of belief of a rational agent. This has led to a system (P-log) that combines logical and probabilistic reasoning based on ASP [2].
- Although in its basic form stable reasoning is highly declarative, when ASP is used in practice the problem representation takes account of the way in which the solver will successfully and economically reach a solution. Features such as cardinality and integrity constraints and more generally *aggregates* are employed to direct the computational mechanism and possibly enhance efficiency [8].
- Basic ASP already deals with some problems of temporal projection. However, to deal with a wider range of problems for dynamically changing domains, a temporal version of equilibrium logic and ASP has been developed and studied, following [4].
- A central property of open systems is the necessity to interact with other systems. Scholars have developed different logical semantics based on ASP that facilitate this interaction. In particular, logic program rules may contain concepts whose meanings are partially determined by external data sources such as knowledge bases or ontologies [7, 16]. This gives rise to hybrid theories that mix different reasoning systems (e.g. monotonic and non-monotonic). Equilibrium logic can be applied to give a simple and uniform treatment of such theories [3].
- To deal with problem solving in a dynamically evolving setting, it is important to consider the problem of updating knowledge in light of new data and knowledge discovery. This has led to the study of theory and program updates in the framework of ASP [17].

7 Acknowledgements

Work on this paper has been partially supported by the projects MERLOT (TIN 2013-42149-P) and FOREST. It is a great pleasure to dedicate this paper to our friend Luis Fariñas del Cerro who has given us generous encouragement and support over a long period. After devoting many years to the advancement of logic and its applications, Luis has recently become also a major contributor to the programme of stable reasoning. We look forward to many more years of intense and fruitful cooperation.

References

1. M. Balduccini and M. Gelfond. Logic Programs with Consistency-Restoring Rules. In P.Doherty, J. McCarthy, and M-A. Williams, eds, *International Symposium on Logical Formalization of Commonsense Reasoning*, AAAI 2003 Spring Symposium Series, pp. 9-18, 2003.
2. C. Baral, M. Gelfond, and N. Rushton. Probabilistic Reasoning with Answer Sets. In *Proceedings of LPNMR-7*, 2004.
3. J. de Bruijn, D. Pearce, A. Polleres and A. Valverde. A semantical framework for hybrid knowledge bases, *Knowledge and Information Systems* (2010) 25: 81–104

4. P. Cabalar, G. Vega, Temporal equilibrium logic: A first approach, in: *Proc. 11th EUROCAST*, LNCS 4739, Springer, 2007, pp. 241–248.
5. R. Carnap. *The Logical Foundations of Probability*, University of Chicago Press, 1950.
6. C. Cellucci, *Le Ragioni della Logica*, Laterza, Roma-Bari, 1998.
7. T. Eiter, G. Ianni, R. Schindlauer, and H. Tompits. A uniform integration of higher-order reasoning and external evaluations in answer-set programming. In: *IJCAI 2005*, pp 90?96.
8. M Gebser, R Kaminski, B Kaufmann, and T Schaub. *Answer set solving in practice*. Synthesis Lectures on Artificial Intelligence and Machine Learning 6 (3), pp. 1-238, 2012.
9. M. Gelfond, Personal Perspective on the Development of Logic Programming Based KR Languages (2011). Available online at <http://www.depts.ttu.edu/cs/research/krlab/papers.php>
10. K. Gödel. Zum intuitionistischen aussagenkalkül. *Anzeiger der Akademie der Wissenschaften Wien, mathematisch, naturwissenschaftliche Klasse*, 69:65–66, 1932.
11. A. Heyting. Die formalen regeln der intuitionistischen logik. *Sitzungsberichte der Preussischen Akademie der Wissenschaften, Physikalisch-mathematische Klasse*, pages 42–56, 1930.
12. T. Hosoi. The axiomatization of the intermediate propositional systems S_n of Gödel. *Journal of the Faculty of Science of the University of Tokio*, 13:183–187, 1966.
13. J. Łukasiewicz. Die Logik und das Grundlagenproblem, in *Les Entretiens de Zürich sur les Fondaments et la Méthode des Sciences Mathématiques* 6-9, 12 (1938), Zürich, 1941, 82-100.
14. L. Maksimova. Craig’s interpolation theorem and amalgamable varieties, *Doklady Akademii Nauk SSSR*, 237, no. 6, (1977), 1281-1284 (translated as *Soviet Math. Doklady*).
15. D. Pearce. A new logical characterisation of stable models and answer sets. In *Non-Monotonic Extensions of Logic Programming, NMELP 96*, LNCS 1216, pages 57–70. Springer, 1997.
16. R. Rosati. DL+log: tight integration of description logics and disjunctive datalog. In: *Proceedings KR 2006*, pp. 68?78.
17. M. Slota and J. Leite. On semantic update operators for answer-set programs. In Coelho, H.; Studer, R.; and Wooldridge, M., eds., *Proceedings of the 19th European Conference on Artificial Intelligence (ECAI 2010)*, volume 215 of *Frontiers in Artificial Intelligence and Applications*, pp. 957?962. IOS Press, 2010.
18. Ya. S. Smetanich. On Completeness of a Propositional Calculus with an additional Operation of One Variable (in Russian). *Trudy Moscovskogo matrmaticheskogo obshchestva*, 9 (1960), 357-372.
19. T. Umezawa. On intermediate many-valued logics. *Journal of the Mathematical Society of Japan* 11(2) (1959).
20. John McCarthy, ‘Elaboration Tolerance’, in *Proc. of the 4th Symposium on Logical Formalizations of Commonsense Reasoning (Common Sense 98)*, pp. 198–217. London, UK (1998).