Available at
www.ElsevierComputerScience.com
POWERED BY SCIENCE DIRECT
Artificial Intelligence 153 (2004) 141–165
Artificial Intelligence
www.elsevier.com/locate/artint

An attempt to formalise a non-trivial benchmark problem in common sense reasoning

Murray Shanahan

Department of Electrical and Electronic Engineering, Imperial College London, Exhibition Road, London SW7 2BT, UK

Received 26 February 2002; received in revised form 9 March 2003

Abstract

Most logic-based AI research works at a meta-theoretical level, producing new logics and studying their properties. Little effort is made to show how these logics can be used to formalise object-level theories of common sense. In the spirit of Pat Hayes's Naive Physics Manifesto, the present paper supplies a formalisation of a non-trivial benchmark problem in common sense physical reasoning, namely how to crack an egg. The formalisation is based on the event calculus, a well-known formalism for reasoning about action. Along the way, a number of methodological issues are raised, such as the question of how the symbols deployed in the formalisation might be grounded through a robot's interaction with the world.

Keywords: Common sense reasoning; Naive physics; Cognitive robotics

1. Introduction

The "Naive Physics Manifesto" papers by Pat Hayes are well-known and widely cited [10,11]. Yet few researchers in AI have taken seriously the research programme he proposed.[1] Logic-based research in AI remains dominated by meta-theory. Papers typically present new logics, extend old ones, or study their meta-level properties. Examples that show how the logic in question is used are usually trivial and are frequently absent altogether. It seems to be taken for granted that, once we've got the right logic, using it will be straightforward. However, the provision of a principled axiomatic description of

E-mail address: m.shanahan@imperial.ac.uk (M. Shanahan).

[1] Notable exception include Davis [4,6], and the CYC project [13].

0004-3702/$ – see front matter
doi:10.1016/j.artint.2003.05.001

the common sense physical world is a difficult and complex task. An instructive analogy here is programming. Knowing that a device is capable of implementing any Turing-computable function is no help at all when faced with the task of writing a large C++ application. A whole battery of principles, techniques, and tricks must be acquired before a programmer can undertake a major project. Similarly, knowing that a logic has sufficient expressive power to represent a given problem domain is no help when it comes to actually constructing such a representation. What logic-based AI lacks is a rich enough body of representational principles, techniques, and tricks, analogous to the programmer's.

In this respect, the naive physics manifesto, with its emphasis on object-level axiomatisation, is still relevant today. Moreover, the problem of endowing a computer with a common sense understanding of the physical world is just as much of a stumbling block now as it was in the 1980s, and approaches to this problem based on logic as a representational formalism are still among the front-runners in the race to solve it. However, with two decades of hindsight, it's difficult—for the present author, at least—to accept all the tenets implicit in Hayes's original manifesto. First, it no longer seems plausible that a useable body of common sense knowledge about the physical world can be coded by hand. Second, the idea that researchers can make significant progress on the problem from their armchairs, that is to say without the "sanity check" of having to deploy their formalisations on a robot, looks ridiculous.

Why should it not be possible to devise useful formalisations without recourse to robotics? Of course, it would be foolhardy to argue this was impossible in principle. But in practise, the efforts made by logicians to represent an agent's interactions with the physical world tend to presuppose a set of predicate and function symbols at too high a level of abstraction, and rashly take it for granted that meaning can somehow be assigned to these symbols. In truth, the design of a set of bottom-level predicate and function symbols that can be *grounded* or *anchored* through sensors and actuators is critical to the success of the whole enterprise [3,9]. Unless formalisations are built on such foundations, they are nothing more than castles in the air.

This paper offers a formalisation of a non-trivial benchmark problem in common sense physical reasoning put forward by Ernie Davis, namely how to crack an egg and pour its contents into a bowl. In the light of the critical points above, it's natural to ask what benefit there is to this exercise. The answer is that the aim of the project is a better understanding of how to deploy formal logic as a medium for representing the everyday physical world, the sort of understanding that can only be acquired through practise at writing object-level theories. The aim is not a definitive object-level theory that will in itself be used in future research. Instead, we'll start to build up the sort of repertoire of principles, techniques, and tricks mentioned above. And throughout, we'll be concerned to address the issue of how the symbols used might be grounded through the sensory-motor apparatus of a robot.

Here is Davis's characterisation of the egg cracking problem [7].[2]

[2] This benchmark problem was originally submitted as a challenge to the participants of the Common Sense 98 workshop. The challenge was taken up by three members of the programme committee, including the present author. (See [15,19].)

> A cook is cracking a raw egg against a glass bowl. Properly performed, the impact of the egg against the edge of the bowl will crack the egg shell in half. Holding the egg over the bowl, the cook with then separate the two halves of the shell with his fingers, enlarging the crack, and the contents of the egg will fall gently into the bowl. The end result is that the entire contents of the egg will be in the bowl, with the yolk unbroken, and that the two halves of the shell are held in the cook's fingers.

Egg cracking is a worthy choice of example because it involves action and continuous change, motion, space and shape, materials, collisions and breakings apart, vessels, and gravity. To formalise it competently we have to address representational questions whose answers will carry over to numerous other problems. Indeed what makes this more than just an undergraduate logic exercise is the need to tackle these large issues in a principled way.

2. What constitutes a good formalisation?

Before embarking on the formalisation itself, let's try to set out some criteria by which it should be judged. How, in general, should an attempt to formalise a body of common sense knowledge to be evaluated? Consider what's wrong with the following naive representation of the egg cracking problem.

$$Initiates(CrackEgg, YolkInBowl, t) \leftarrow HoldsAt(BowlUnderEgg, t)$$

This formalisation has (at least) the following faults.

- It doesn't employ any knowledge that can be recycled for other problems. As already emphasised, we should expect a solution that appeals to general principles and techniques.
- Apart from their suggestive English names, no clue is given as to the meaning of *CrackEgg*, *YolkInBowl* and *BowlUnderEgg*. It's hard to imagine how the axiom could be put to any useful purpose, such as to program a robot to actually crack an egg. How could a robot recognise that the fluent *BowlUnderEgg* holds? How could the robot perform a *CrackEgg* action? We would like a solution constructed out of fluents which are more clearly related to the data obtainable from a robot's sensory apparatus and actions which are more plausibly executable by a robot.
- It doesn't allow for variations on the problem. In other words, it lacks *elaboration tolerance*, to use McCarthy's term [17]. Davis lists a number of elaborations of the egg cracking problem.

  > What happens if: The cook brings the egg to impact very quickly? Very slowly? The cook lays the egg in the bowl and exerts steady pressure with his hand? The cook, having cracked the egg, attempts to peel it off its contents like a hard-boiled egg? The bowl is made of looseleaf paper? of soft clay? The bowl is smaller than the egg? The bowl is upside down? The cook tries this procedure with a hard-boiled egg? With a coconut? With an M & M?

  We would like a solution in which elaborations such as these demand the least possible effort to make.

To summarise, we can judge a solution by (at least) the following three criteria: the degree to which it employs general-purpose, re-usable representational techniques, its usability in some (possibly imaginary) application, and its elaboration tolerance.

With the first criterion in mind, here's a rough survey of representational issues we'll have to confront to do the job properly. These are the areas where we should expect to develop reusable theories, or to be able to pick such theories off the shelf. To begin with, we require a formalism for reasoning about action, such as the event calculus or the situation calculus. The formalism needs to be able to handle continuous change, since we have to deal with the continuous motion of the egg and its contents. The topic of reasoning about action is one of the few that is already well-developed.

The act of hitting something against something else has to be formalised. A "start motion" action is involved, which initiates a period of continuous movement, which ends when the egg strikes the bowl. The treatment of this action should be general enough to deal with other sorts of agent-initiated continuous movement. On a more problem-specific level, the agent has to be holding the egg, the motion of the agent's arm must be in the right direction (towards the rim of the bowl), and the agent needs to strike with the right force.

Axioms for collision events will be required, and these should be as general-purpose as possible. The consequences of collision events depend on the fragility of the two objects involved, and the force of collision. In the case of an egg, the egg cracks but remains intact if the force of collision is just right.

Next, the act of prizing apart the two halves of the egg needs formalising. We need to represent the shape of the egg, the shape of the crack, and the shapes of the two halves that result when the egg is pulled apart. We need a way of representing objects coming into existence and ceasing to exist (i.e., the whole egg and the two halves of the egg). It will be harder to do all this in the context of a general theory, as the behaviour of a fractured egg while it's being prized apart is rather peculiar.

The pouring or dropping of the egg's contents into the bowl demands the representation of various facts about vessels and containment. We need to represent the fact that the egg is hollow, that its surface has no openings, and that it contains a liquid (or liquid-like substance). The behaviour of the yolk and egg-white as the egg halves are separated needs to be formalised. We might assume, for simplicity, that the egg's contents are a liquid that flows under the influence of gravity. We need to formalise the shape of the bowl, the fact that it's a vessel with an upward-facing aperture, and the fact that it lies below the egg. We need to represent the fact that liquid is retained in a vessel if the vessel is upright.

What follows is a crude formulation. Many of the issues mentioned above are addressed, but not all. Furthermore, numerous simplifying assumptions are made. Among the more gross idealisations it embodies is the assumption that the contents of the egg behave like water—so that it separates easily into two bodies of liquid in the two halves of the egg. It's also assumed that the egg breaks cleanly into two parts held upright. In reality, of course, when we prize open an egg-shell, generally with our thumbs at the bottom, the two halves remain touching, sometimes even joined, at the top. And because of the viscosity of the egg's contents, it initially remains inside the egg even though a crack has appeared through which a less viscous material would flow. Only when the opening has become wide enough does it all fall out.

Table 1
The language of the event calculus

Formula	Meaning
$Initiates(e, f, t)$	Fluent f starts to hold after action e occurs at time t
$Terminates(e, f, t)$	Fluent f ceases to hold after action e occurs at time t
$Releases(e, f, t)$	Fluent f ceases to be subject to the common sense law of inertia after action e occurs at time t
$Trajectory(f1, t, f2, d)$	Fluent $f2$ holds at time $t+d$ if fluent $f1$ starts to hold at t and continues to hold up to $t+d$
$Initially_P(f)$	Fluent f holds from time 0
$Initially_N(f)$	Fluent f does not hold from time 0
$Happens(e, t)$	Action or event e occurs at time t
$HoldsAt(f, t)$	Fluent f holds at time t
$Clipped(t1, f, t2)$	Fluent f is terminated between times $t1$ and $t2$
$Declipped(t1, f, t2)$	Fluent f is initiated between times $t1$ and $t2$

3. The event calculus

The event calculus will be adopted as a formalism for representing actions and their effects [24]. The egg cracking problem places considerable demands on an action formalism, and the full representational power of the event calculus will be required, in order to capture continuous change, actions with non-deterministic effects, and actions with indirect effects. The event calculus is based on first-order predicate calculus, extended with circumscription to overcome the frame problem. Table 1 presents the essentials of the language of the calculus, which includes sorts for fluents, actions (events), and time points. The sort of time points is assumed to be interpreted by the positive reals.

The basic event calculus axioms, including those for continuous change, are as follows. Throughout the paper, all variables in a formula are universally quantified, with maximum possible scope, unless stated otherwise.

$$HoldsAt(f, t) \leftarrow Initially_P(f) \wedge \neg Clipped(0, f, t) \tag{EC1}$$

$$\neg HoldsAt(f, t) \leftarrow Initially_N(f) \wedge \neg Declipped(0, f, t) \tag{EC2}$$

$$\begin{aligned} &HoldsAt(f, t2) \leftarrow \\ &\quad Happens(e, t1) \wedge Initiates(e, f, t1) \wedge t1 < t2 \wedge \neg Clipped(t1, f, t2) \end{aligned} \tag{EC3}$$

$$\begin{aligned} &\neg HoldsAt(f, t2) \leftarrow \\ &\quad Happens(e, t1) \wedge Terminates(e, f, t1) \wedge t1 < t2 \wedge \neg Declipped(t1, f, t2) \end{aligned} \tag{EC4}$$

$$\begin{aligned} &Clipped(t1, f, t2) \leftrightarrow \\ &\quad \exists e, t[Happens(e, t) \wedge t1 < t < t2 \wedge [Terminates(e, f, t) \vee Releases(e, f, t)]] \end{aligned} \tag{EC5}$$

$$\begin{aligned} &Declipped(t1, f, t2) \leftrightarrow \\ &\quad \exists e, t[Happens(e, t) \wedge t1 < t < t2 \wedge [Initiates(e, f, t) \vee Releases(e, f, t)]] \end{aligned} \tag{EC6}$$

$$\begin{aligned} &HoldsAt(f2, t2) \leftarrow \\ &\quad Happens(e, t1) \wedge Initiates(e, f1, t1) \wedge t1 < t2 \wedge \\ &\qquad t2 = t1 + d \wedge Trajectory(f1, t1, f2, d) \wedge \neg Clipped(t1, f1, t2) \end{aligned} \tag{EC7}$$

A consequence of Axioms (EC3) and (EC5) is that a fluent, once it has been initiated by an event, will hold over an interval which is open to the left and closed to the right. In other words, it doesn't start to hold until immediately *after* the event that initiates it, but it still holds at the time of the event that terminates it. Likewise, a consequence of Axioms (EC4) and (EC6) is that once a fluent has been terminated, it only starts not to hold immediately *after* the terminating event, but it still does not hold at the time of the next event that initiates it. This observation only applies to *inertial* fluents, that is to say those that are initiated and terminated by events (or actions). Continuously varying fluents, for example, are non-inertial, and fall under the control of Axiom (EC7).

To overcome the frame problem, we use circumscription to minimise the predicates *Happens*, *Initiates*, *Terminates*, and *Releases* [24]. If E is a domain description (*Initiates*, *Terminates*, and *Releases* formulae) and N is a narrative description ($\mathit{Initially}_P$, $\mathit{Initially}_N$, *Happens* and temporal ordering formulae), then we consider,

$$\mathit{CIRC}[\mathrm{N}; \mathit{Happens}] \wedge$$
$$\quad \mathit{CIRC}[\mathrm{E}; \mathit{Initiates}, \mathit{Terminates}, \mathit{Releases}] \wedge \mathrm{CEC}$$

where CEC is the conjunction of Axioms (EC1) to (EC7). State constraints—that is to say, formulae constraining the combinations of fluents allowed to hold at the same time—must be conjoined to CEC. A collection of uniqueness-of-names axioms for fluents and actions must also be conjoined to the above formula. These will be taken for granted in what follows, and omitted from the presentation. A detailed presentation of the event calculus, in a more tutorial form, may be found in [25].

4. Collisions and breakings apart

We begin the more domain specific formalisation with some general axioms describing the action of striking one thing against another. Table 2 summarises the actions and fluents used in the axioms that follow. Each of the terms in Table 2 will be taken as primitive. This assumption is unrealistic from a robotics perspective, as discussed below. Furthermore, it's less than satisfactory from a logical point-of-view. Consider the *Distance* fluent, for example. Although this is possibly the simplest concept featured in Table 2, there are still several ways to define the distance between two objects in more primitive terms [8, Chapter 7]. Notwithstanding this, we're forced to idealise ruthlessly in some areas in order to make any progress at all. In what follows, time and change are given the most detailed treatment, at the expense of space and shape.

Here are the effect axioms for the actions and fluents of Table 2. They say that an object moving towards another object will hit it eventually, and then stop. Needless to say, there are exceptions to this, such as when the first object passes right through the second. These exceptions are neglected here. Furthermore, the event calculus axioms will easily accommodate the possibility that the moving object is intercepted by a third object. But we won't explore this possibility here. It is assumed that a robot executing a *PropelTo* action doesn't let go of the object x, and holds on to it while the *MovingTo* fluent holds.

$$\mathit{Initiates}(\mathit{PropelTo}(x, y, v), \mathit{MovingTo}(x, y, v), t) \tag{E1.1}$$

Table 2
Collisions and continuous movements

Term	Sort	Meaning
$PropelTo(x, y, v)$	Action	Object or part x is propelled towards object or part y at velocity v
$MovingTo(x, y, v)$	Fluent	Object or part x is moving towards object or part y at velocity v
$Stop(x)$	Action	Stop moving object or part x
$Distance(x, y, z)$	Fluent	The distance between object or part x and object or part y is z
$CollidesWith(x, y, v)$	Action	Object or part x collides with object or part y at velocity v
$Fractured(x)$	Fluent	Object x is fractured

$$Releases(PropelTo(x, y, v), Distance(x, y, z), t) \tag{E1.2}$$

$$Terminates(Stop(x), MovingTo(x, y, v), t) \tag{E1.3}$$

$$Initiates(Stop(x), Distance(x, y, z), t) \leftarrow HoldsAt(Distance(x, y, z), t) \tag{E1.4}$$

$$\begin{aligned}&Trajectory(MovingTo(x, y, v), t, Distance(x, y, z2), d) \leftarrow\\ &\quad HoldsAt(Distance(x, y, z1), t) \wedge z2 = z1 - v \cdot d\end{aligned} \tag{E1.5}$$

$$\begin{aligned}&[HoldsAt(Distance(x, y, z1), t) \wedge\\ &\quad HoldsAt(Distance(z, y, z2), t)] \rightarrow z1 = z2\end{aligned} \tag{E1.6}$$

$$\begin{aligned}&Happens(CollidesWith(x, y, v), t) \leftarrow\\ &\quad HoldsAt(MovingTo(x, y, v), t) \wedge HoldsAt(Distance(x, y, 0), t)\end{aligned} \tag{N1.1}$$

This group of axioms exemplifies the style in which continuous change is formalised using the event calculus. Of particular note is the presence of the *Releases* axiom (E1.2). This axiom ensures that, at the time of a *PropelTo* action, *Distance* changes from an inertial fluent, under the influence of Axioms (EC1) to (EC4), to a non-inertial fluent which is no longer subject to the common sense law of inertia. While this fluent is non-inertial, its value is determined by Axiom (EC7) and the relevant domain-specific *Trajectory* formula, in this case (E1.5)[3]. Moreover, for every continuously varying non-inertial fluent, there is a corresponding inertial "second-order" fluent governing it. In this case, the second-order fluent that governs *Distance* is *MovingTo*. A *PropelTo* action not only makes *Distance* non-inertial, but also initiates *MovingTo* (E1.1). The period of continuous variation in *Distance* lasts while *MovingTo* holds, and is ended by a *Stop* action (E1.3). At the end of this period, *Distance* reverts to an inertial fluent, its value being set according to Axiom (E1.4).

The idea behind this particular group of axioms is that the robot will execute a *PropelTo* action followed by a *Stop* action at the time of the collision. But let's pause here to assess the plausibility of this choice of symbols. What exactly are $PropelTo(x, y, v)$ and $Stop(x)$ supposed to denote? The implicit assumption is that they denote actions that are, in some sense, primitive. That is to say, these actions don't require further breaking down into sub-actions, but might be issued as a direct command to a robot arm. Even if we set aside the

[3] Axiom (E1.5) assumes that the motion of the egg has a constant velocity. In reality, of course, it will unergo a short period of acceleration, and even after this its velocity is unlikely to be absolutely constant. But all that really counts here are the more qualitative properties of the trajectory that ensure the egg reaches its target. These could be captured using qualitative reasoning techniques, such as those axiomatised in [5].

whole issue of how the egg is to be grasped in the first place—an extremely complex matter in itself—the idea that *PropelTo* and *Stop* could be primitive actions seems flawed.

To see this, consider that perception and action must operate in a tight feedback loop to ensure that the grasped object meets its target and that the force behind the strike ceases at the very moment of impact. This requires a solution to the problem of visual object recognition of a generality that far exceeds the current state-of-the-art in machine vision, because the point of impact on the target object has to be visually identified regardless of the target's shape, when lighting conditions are poor, where shadows, highlights, and surface patterns abound, and when other objects clutter the scene. None of this would matter in the context of the present exercise in knowledge representation were it not for the fact that robust perception almost certainly depends on *expectation*, which in turn relies on a common sense understanding of the physical world [27].

Am I, then, suggesting that we've put the cart before the horse, that we should be tackling the topic of perception before that of common sense reasoning? Not at all. What we're talking about here are not horses and carts but—if you'll forgive the pun—chickens and eggs. To the question of which comes first, perception or common sense reasoning, the answer is neither. The recommendation is rather that we should be cautious about working on either topic in isolation. Moreover, as the champions of active perception have argued persuasively, perception itself is best studied in the context of action [1,2]. These arguments, taken in addition to those just advanced, entail that perception, action, and cognition form a trinity of subjects that shouldn't be separated.

In spite of this, there's much still to be learned if we proceed under the pretence that the *PropelTo* and *Stop* actions are a reasonable choice of primitive actions. The catalogue of issues that arises is both instructive and representative. The next clutch of formulae captures the effects of the *CollidesWith* action. It would be desirable to have a generic formalisation of the consequences of collisions, one in which the appearance of cracks was a special case. But for now we'll go for the easy option, which is to write specialised axioms for eggs and bowls. Table 3 describes the language features employed.

Here's the main axiom, which says that the collision of the egg shell with the bowl's rim brings about a fracture in the egg along its circumference if the velocity of the collision is just right. The formula $IsA(x, y)$ represents that object x is an instance of type y. The formula $Material(x, y)$ means that the whole of object x, including all its parts, is made of material y.

Table 3
Eggs and bowls

Formula/term	Sort	Meaning
$Shell(x)$	Object part	The shell of egg x
$Rim(x)$	Object part	The rim of bowl x
$Middle(x)$	Object part	The contents of egg x, including the egg-white and the yolk
Egg	Type of object	Eggs
Bowl	Type of object	Bowls
$Hard(x)$		Object or part x is hard
$Material(x, y)$		All of object or part x is made of material y
$JREC(v)$		Velocity v is "Just Right for Egg Cracking", in other words not so small that the egg remains intact nor so great that the egg smashes to pieces

$$Initiates(CollidesWith(Shell(x), y, v), Fractured(Shell(x)), t) \leftarrow IsA(x, Egg) \wedge Hard(y) \wedge JREC(v) \quad \text{(E1.7)}$$

The following group of axioms deals (very superficially) with materials, and with some problem-specific aspects of parts and wholes. The issue of parts and wholes is treated more fully shortly.

$$Hard(x) \leftarrow Material(x, Glass) \quad \text{(B1.1)}$$

$$Material(x, y) \leftarrow PartOf(x, z) \wedge Material(z, y) \quad \text{(B1.2)}$$

$$PartOf(Rim(x), x) \leftarrow IsA(x, Bowl) \quad \text{(B1.3)}$$

$$PartOf(Shell(x), x) \leftarrow IsA(x, Egg) \quad \text{(B1.4)}$$

Again, if we take a step back and ask how these symbols might be grounded in the sensory-motor activity of a robot, we expose serious shortcomings in the representation. Especially glaring is the absence of a proper treatment of shape and space in the context of which the terms $Shell(x)$ and $Rim(x)$ would have a more precise sense. A more thorough axiomatisation would flesh these concepts out in terms of spatial occupancy and shape. In [12], for example, we find an attempt to formally characterise the concept of a container in terms of enclosures and portals. Although Hayes's choice of ontological primitives, being at a lower level, is more apt than the present choice, they both share a major shortcoming. What we really want is a formal theory that will tie in to a vision system capable of recognising vessels. If tackled in full, this issue alone would fill several Ph.D theses. So once again, we'll content ourselves with a compromise in order to make a decent attempt at a total formalisation.

The next axiom describes the effect of splitting the cracked egg into two. Table 4 sets out the meanings of the language elements used. Some of these are general purpose language features for containment, and parts and wholes. The two shapes and the *SplitEgg* action are problem-specific.

In practise, a human egg-cracker usually prizes the fractured shell apart in such a way as to allow the whole contents to fall out as soon as the resulting fissure is wide enough. But we'll assume a more peculiar way of doing the job, which results in the contents of the egg being divided into two parts distributed between the two egg-shell halves. This will force us to confront the interesting issue of parts-and-wholes. In unreified shorthand, the effect

Table 4
Containment, parts and shapes

Formula/term	Sort	Meaning
$Contains(x, y)$	Fluent	Vessel x contains object y
$Shape(x, y)$		Object x has shape y
WholeShell	Shape	The shape of an intact egg shell
HalfShell	Shape	The shape of half an egg shell
$Comprises(x, y1, y2)$		Object x can be split into objects $y1$ and $y2$
$PartOf(x, y)$		Object x is a proper part of object y
$Half1(x)/Half2(x)$	Object part	The first/second half-shell of egg x
$SplitShell(x)$	Action	The shell of egg x is broken into two halves

we expect from a *SplitShell* action is this. Assuming the egg-shell is fractured in the right way, before splitting the egg we will have $\mathit{Contains}(\mathit{Shell}(x), \mathit{Middle}(x))$. After splitting the egg, we want, for some $x1$, $x2$, $y1$, and $y2$,

$$\begin{aligned}&\mathit{Contains}(x1, y1) \wedge \mathit{Contains}(x2, y2) \wedge\\&\quad \mathit{Shape}(x1, \mathit{HalfShell}) \wedge \mathit{Shape}(x2, \mathit{HalfShell}) \wedge\\&\qquad \mathit{Comprises}(\mathit{Shell}(x), x1, x2) \wedge \mathit{Comprises}(\mathit{Middle}(x), y1, y2)\end{aligned}$$

The following axiom does the job.

$$\begin{aligned}&\exists y1, y2[\mathit{Initiates}(\mathit{SplitShell}(x), \mathit{Contains}(\mathit{Half}1(x), y1), t) \wedge\\&\quad \mathit{Initiates}(\mathit{SplitShell}(x), \mathit{Contains}(\mathit{Half}2(x), y2), t) \wedge\\&\qquad \mathit{Comprises}(\mathit{Middle}(x), y1, y2)] \leftarrow\\&\qquad\quad \mathit{IsA}(x, \mathit{Egg}) \wedge \mathit{HoldsAt}(\mathit{Contains}(\mathit{Shell}(x), \mathit{Middle}(x)), t) \wedge\\&\qquad\qquad \mathit{HoldsAt}(\mathit{Fractured}(\mathit{Shell}(x)), t)\end{aligned} \tag{E1.8}$$

The purpose of the $\mathit{Half}1$ and $\mathit{Half}2$ functions is to permit the two halves resulting from a *SplitShell* action to be named. Logically speaking, they are acting as Skolem functions that substitute for existentially quantified variables in (E1.8). Strictly, these functions should include a time argument, as distinct *SplitShell* actions could result in different divisions of the same egg-shell. However, as we can safely assume that an egg-shell will never suffer more than one splitting in its lifetime, this parameter is omitted here. The following pair of axioms specify the shapes of the two egg-shell halves. Note that there's no need to specify the shape exactly.

$$\mathit{Shape}(\mathit{Shell}(x), \mathit{WholeShell}) \leftarrow \mathit{IsA}(x, \mathit{Egg}) \tag{B1.5}$$

$$\mathit{Shape}(\mathit{Half1}(x), \mathit{HalfShell}) \leftarrow \mathit{IsA}(x, \mathit{Egg}) \tag{B1.6}$$

$$\mathit{Shape}(\mathit{Half}2(x), \mathit{HalfShell}) \leftarrow \mathit{IsA}(x, \mathit{Egg}) \tag{B1.7}$$

The predicate *Comprises* is defined as follows.

$$\begin{aligned}&\mathit{Comprises}(x, y1, y2) \leftrightarrow\\&\quad \neg\exists y3[\mathit{PartOf}(y3, x) \wedge \forall y4[\mathit{PartOf}(y4, y3) \rightarrow\\&\qquad [\neg\mathit{PartOf}(y4, y1) \wedge \neg\mathit{PartOf}(y4, y2)]]] \wedge\\&\qquad\quad \neg\exists y3[\mathit{PartOf}(y3, y1) \wedge \mathit{PartOf}(y3, y2)]\end{aligned} \tag{B1.8}$$

In other words, x comprises $y1$ and $y2$ if there is no part of x that is entirely outside of both $y1$ and $y2$, and if $y1$ and $y2$ don't overlap. In particular, an egg-shell is comprised of its two halves.

$$\begin{aligned}&\mathit{Comprises}(\mathit{Shell}(x), \mathit{Half}1(x), \mathit{Half}2(x)) \leftarrow \mathit{IsA}(x, \mathit{Egg})\\&\quad \mathit{HoldsAt}(\mathit{Fractured}(\mathit{Shell}(x)), t)\end{aligned} \tag{B1.9}$$

The next axiom is required because, as a consequence of an egg's destruction by the *SplitShell* action, it can no longer be said to contain anything.

$$\mathit{Terminates}(\mathit{SplitShell}(x, x1, x2), \mathit{Contains}(\mathit{Shell}(x), y), t) \leftarrow \tag{E1.9}$$

Some remarks about existence and non-existence are in order here. Nowhere has the fact been represented that the original egg shell ceases to exist when the two halves are pulled apart. Instead, the consequences for the *Contains* fluent have been emphasised—a non-existent egg can't contain anything. Indeed, more generally, no physical fluent can hold for a non-existent object, and we could set down an axiom to this effect. A more principled approach to physical existence and non-existence is taken by Davis in his treatment of the effects of cutting an object in two [6], and also in [22], where an object is considered to exist if and only if it occupies some region of space.

Now we need an axiom relating parts and containment. The following axiom seems adequate on first examination.

$$\mathit{HoldsAt}(\mathit{Contains}(x,y),t) \leftrightarrow \forall z[\mathit{PartOf}(z,y) \rightarrow \mathit{HoldsAt}(\mathit{Contains}(x,z),t)]$$

But it turns out that a more useful axiom can be formulated.

$$\mathit{HoldsAt}(\mathit{Contains}(x,y),t) \leftrightarrow \neg\exists z1[\mathit{PartOf}(z1,y) \wedge \forall z2[\mathit{PartOf}(z2,z1) \rightarrow \neg\mathit{HoldsAt}(\mathit{Contains}(x,z2),t)]] \quad \text{(B1.10)}$$

This axiom insists that no part of an object can be wholly outside a vessel that contains it. In this form, the axiom interacts more straightforwardly with the definition of *Comprises*. Given the right formalisation of the part-whole relation, the second of these axioms will follow from the first. But no attempt will be made here to present and justify a fully-fledged theory of parts-and-wholes. (A detailed formal investigation of this subject can be found in [21]). Instead, we'll adopt just three axioms. The first two state that the *PartOf* relation is irreflexive and transitive.

$$\neg\mathit{PartOf}(x,x) \quad \text{(B1.11)}$$

$$\mathit{PartOf}(x,y) \wedge \mathit{PartOf}(y,z) \rightarrow \mathit{PartOf}(x,z) \quad \text{(B1.12)}$$

The third axiom we require insists that all parts have sub-parts, reflecting an assumption that space is infinitely divisible. (Metaphysically, this rules out atomic objects, although it still allows for an object that cannot physically be split.) Without this axiom, an object with a part with no sub-parts would be possible, and such an object would be contained by everything according to Axiom (B1.10).)

$$\mathit{PartOf}(x,y) \rightarrow \exists z\mathit{PartOf}(z,x) \quad \text{(B1.13)}$$

Getting back to the specifics of the benchmark problem, the egg-cracker needs to be able to move the egg over the bowl. Let's cheat a little here, and provide an action that does the job directly and instantaneously, ignoring the continuous motion involved (Table 5).

Table 5
Moving things above things

Term	Sort	Meaning
$\mathit{MoveAbove}(x,y)$	Action	Object x is moved above object y
$\mathit{Beneath}(x,y)$	Fluent	Object y is above object x

$$\mathit{Initiates}(\mathit{MoveAbove}(x, y), \mathit{Beneath}(y, x), t) \quad \text{(E1.10)}$$

Now we have two state constraints to ensure that the *Beneath* fluent respects common sense intuitions about parts and wholes and containment. All of an object's parts are above everything the object itself is above, and an object is above everything its container is above.

$$\mathit{HoldsAt}(\mathit{Beneath}(x, y), t) \leftarrow \mathit{PartOf}(y, z) \wedge \mathit{HoldsAt}(\mathit{Beneath}(x, z), t) \quad \text{(B1.14)}$$

$$\mathit{HoldsAt}(\mathit{Beneath}(x, y), t) \leftarrow \quad \text{(B1.15)}$$
$$\quad \mathit{HoldsAt}(\mathit{Contains}(z, x), t) \wedge \mathit{HoldsAt}(\mathit{Beneath}(z, y), t)$$

Two further axioms, with the arguments in different orders, are required but not shown.

5. Liquids and vessels

The final group of axioms describes the behaviour of liquids and vessels. Table 6 introduces some of the new language features required.

First we'll augment the effects of the *SplitShell* action to reflect the assumption that it is performed in such a way as to ensure that the resulting egg halves are held upright. (If carried out by a human, this means that the egg-cracker's thumbs are on top of the egg while it is pulled apart.)

$$\mathit{Terminates}(\mathit{SplitShell}(x), \mathit{Angle}(\mathit{Half}1(x), 0), t) \quad \text{(E2.1)}$$

$$\mathit{Terminates}(\mathit{SplitShell}(x), \mathit{Angle}(\mathit{Half}2(x), 0), t) \quad \text{(E2.2)}$$

Now, we can move from the problem-specific terms of Axioms (E1.4) to (E2.2) back into the generic realm with the following pair of background axioms.

$$\mathit{IsA}(x, \mathit{OpenVessel}) \leftarrow \mathit{Shape}(x, \mathit{HalfShell}) \quad \text{(B2.1)}$$

$$\mathit{IsA}(x, \mathit{ClosedVessel}) \leftarrow \mathit{Shape}(x, \mathit{WholeShell}) \quad \text{(B2.2)}$$

Of course, we also have the following.

$$\mathit{IsA}(x, \mathit{OpenVessel}) \leftarrow \mathit{IsA}(x, \mathit{Bowl}) \quad \text{(B2.3)}$$

A vessel isn't open if it's closed.

$$\neg \mathit{IsA}(x, \mathit{OpenVessel}) \leftarrow \mathit{IsA}(x, \mathit{ClosedVessel}) \quad \text{(B2.4)}$$

Table 6
Pouring liquid from a vessel

Term	Sort	Meaning
OpenVessel	Type of object	Vessels with an opening
ClosedVessel	Type of object	Sealed vessels
$\mathit{Tilt}(x)$	Action	Vessel x starts to tilt
$\mathit{Tilting}(x)$	Fluent	Vessel x is tilting
$\mathit{Angle}(x, h)$	Fluent	Vessel x is tilted at angle h from horizontal
$\mathit{Upturned}(x)$	Fluent	The contents of vessel x can fall out
$\mathit{Spill}(x, y)$	Action	The contents x of vessel y start to fall out

If you tilt an open vessel, eventually its contents start to fall out. This effect is captured through the $Angle(x, h)$ fluent, whereby the angle h increases continuously while the second-order fluent *Tilting* holds. When h reaches a critical point in its trajectory, the *Upturned* fluent starts to hold, and a *Spill* event is triggered. We'll assume the sort of "angles" only permits values between 0 and 2π radians, and that functions returning an angle are taken modulus 2π.

$$\exists h1[h1 > 0 \wedge [HoldsAt(Upturned(x), t) \leftrightarrow HoldsAt(Angle(x, h2), t) \wedge h2 \geqslant h1]] \quad \text{(E2.3)}$$

$$Initiates(Tilt(x), Tilting(x), t) \leftarrow \neg HoldsAt(Upturned(x), t) \quad \text{(E2.4)}$$

$$Releases(Tilt(x), Angle(x, h), t) \quad \text{(E2.5)}$$

$$Terminates(Stop(x), Tilting(x), t) \quad \text{(E2.6)}$$

$$Initiates(Stop(x), Angle(x, h), t) \leftarrow HoldsAt(Angle(x, h), t) \quad \text{(E2.7)}$$

$$\exists w[Trajectory(Tilting(x), t, Angle(x, h2), d) \leftarrow HoldsAt(Angle(x, h1), t) \wedge h2 = h1 + w \cdot d] \quad \text{(E2.8)}$$

$$[HoldsAt(Angle(x, h1), t) \wedge HoldsAt(Angle(z, h2), t)] \rightarrow h1 = h2 \quad \text{(E2.9)}$$

$$Happens(Spill(x, y), t) \leftarrow HoldsAt(Upturned(y), t) \wedge HoldsAt(Tilting(y), t) \wedge HoldsAt(Contains(y, x), t) \quad \text{(N2.1)}$$

$$Terminates(Spill(x, y), Contains(y, x), t) \quad \text{(E2.10)}$$

Table 7 presents the remaining language features introduced in this section. A *Spill* event initiates an interval over which the fluent *Falling* holds. During this interval, the fluent $Distance(x, y, z)$ undergoes continuous variation, where z is the distance between falling object x and an object y underneath it. The constant G denotes the rate of acceleration due to gravity.

$$Initiates(Spill(x, y), Falling(x, z), t) \leftarrow HoldsAt(Beneath(z, y), t) \quad \text{(E2.11)}$$

$$Releases(Spill(x, y), Distance(x, z, w), t) \leftarrow HoldsAt(Beneath(z, y), t) \quad \text{(E2.12)}$$

$$Terminates(Fill(x, y), Falling(y, x), t) \quad \text{(E2.13)}$$

Table 7
Liquids and vessels

Term	Sort	Meaning
$Falling(x, y)$	Fluent	The contents x of vessel y is falling out
$Fill(x, y)$	Action	Vessel x starts to fill with object y
$Filling(x, y)$	Fluent	Vessel x is filling with object y
$Capacity(x)$		The capacity of vessel x
$Volume(x)$		The volume of object x
$Opening(x)$	Region	The region of empty space that comprises the opening of vessel x
$PContains(x, y, k)$	Fluent	Vessel x contains a portion of object y, corresponding to k times its volume, where k ranges from 0 to 1
$Overflows(x)$	Action	Vessel x overflows
$GotAll(x, y)$	Action	Vessel x contains all of object y

$$Initiates(Fill(x, y), Distance(x, y, 0), t) \tag{E2.14}$$

$$Trajectory(Falling(x, y), t, Distance(x, y, z2), d) \leftarrow \\ HoldsAt(Distance(x, y, z1), t) \wedge z2 = z1 - 0.5 \cdot G \cdot d^2 \tag{E2.15}$$

If a falling body encounters the opening of a vessel, it starts to fill that vessel. This effect is captured by the triggering of a *Fill* event when the distance between the falling body and the opening becomes zero. A $Fill(x, y)$ event initiates a period of continuous change during which the proportion of object y contained by vessel x gradually increases, as represented by the *PContains* fluent.

$$Happens(Fill(x, y), t) \leftarrow \\ HoldsAt(Distance(y, x, 0), t) \wedge HoldsAt(Falling(y, Opening(x)), t) \wedge \\ IsA(x, OpenVessel) \wedge \neg HoldsAt(Upturned(x), t) \tag{N2.2}$$

$$Initiates(Fill(x, y), Filling(x, y), t) \tag{E2.16}$$

$$Releases(Fill(x, y), PContains(x, y, k), t) \tag{E2.17}$$

$$Releases(Fill(x, y), Contains(x, y), t) \tag{E2.18}$$

$$\exists r[Trajectory(Filling(x, y), t, PContains(x, y, k), d) \leftarrow k = r \cdot d] \tag{E2.19}$$

$$[HoldsAt(PContains(x, y, k1), t) \wedge \\ HoldsAt(PContains(z, y, k2), t)] \rightarrow k1 = k2 \tag{E2.20}$$

If the vessel continues to fill, eventually either an *Overflows* event or a *GotAll* event occurs. An $Overflows(x)$ event occurs if the combined volume of the proportions of the two objects that have fallen into vessel x exceeds the capacity of x. Ideally, the conditions that trigger an $Overflows(x)$ event would take account of arbitrarily many sources of flow into x, whereas Axiom (N2.3) below assumes that exactly two objects are falling into x. A more comprehensive treatment of this issue can be found in [18].

$$Happens(Overflows(x), t) \leftarrow \\ HoldsAt(PContains(x, y1, k1), t) \wedge \\ HoldsAt(PContains(x, y2, k2), t) \wedge \\ \neg\exists y3[PartOf(y3, y1) \wedge PartOf(y3, y2)] \wedge \\ Capacity(x) = k1 \cdot Volume(y1) + k2 \cdot Volume(y2) \tag{N2.3}$$

$$Initiates(Overflows(x), PContains(x, y, k), t) \leftarrow \\ HoldsAt(PContains(x, y, k), t) \tag{E2.21}$$

$$Terminates(Overflows(x), Filling(x, y), t) \tag{E2.22}$$

For Axiom (N2.3) to work properly when the two parts of an egg are recombined in a bowl, volume must be additive.

$$Comprises(x, y, z) \rightarrow Volume(x) = Volume(y) + Volume(z) \tag{B2.5}$$

A $GotAll(x, y)$ event occurs if the whole of object y ends up in vessel x. Note that, if the capacity of a vessel is equal to the volume of the object(s) falling into it, a *GotAll* even and an *Overflows* event will occur simultaneously.

$$Happens(GotAll(x, y), t) \leftarrow HoldsAt(PContains(x, y, 1), t) \wedge HoldsAt(Filling(x, y), t) \quad \text{(N2.4)}$$

$$Initiates(GotAll(x, y), Contains(x, y), t) \quad \text{(E2.23)}$$

$$Terminates(GotAll(x, y), Filling(x, y), t) \quad \text{(E2.24)}$$

The axioms of this section are intended to capture some aspects of the common sense physics of spilling, falling, filling, and overflowing.[4] Their adequacy for the benchmark problem at hand is demonstrated in the next section.

6. A cracking narrative

The aim of this section is to show that, taken together, the axioms of Sections 3–5 are sufficient to represent the narrative of events set out in Davis's description of the egg-cracking benchmark. To begin with, we have an egg and a glass bowl.

$$IsA(Egg0, Egg) \quad \text{(B3.1)}$$

$$IsA(Bowl0, Bowl) \quad \text{(B3.2)}$$

$$Material(Bowl0, Glass) \quad \text{(B3.3)}$$

$$Capacity(Bowl0) \geqslant Volume(Middle(Egg0)) \quad \text{(B3.4)}$$

$$Initially_N(Upturned(Bowl0)) \quad \text{(B3.5)}$$

$$Initially_P(Distance(Shell(Egg0), Rim(Bowl0), Z)) \quad \text{(B3.6)}$$

$$Z > 0 \quad \text{(B3.7)}$$

$$Initially_P(Contains(Shell(Egg0), Middle(Egg0))) \quad \text{(B3.8)}$$

First, the egg is hit against the rim of the bowl. The robot propels the side of the egg towards the rim of the bowl, and stops as soon as the collision occurs.

$$Happens(PropelTo(Shell(Egg0), Rim(Bowl0), V), T0) \quad \text{(N3.1)}$$

$$JREC(V) \quad \text{(N3.2)}$$

$$Happens(Stop(Shell(Egg0)), t1) \wedge t1 < T1 \leftarrow T0 < t1 \wedge HoldsAt(Distance(Shell(Egg0), Rim(Bowl0), 0), t1) \wedge \neg\exists t2[HoldsAt(Distance(Shell(Egg0), Rim(Bowl0), 0), t2) \wedge T0 < t2 < t1] \quad \text{(N3.3)}$$

Next the robot moves the fractured egg above the bowl.

$$Happens(MoveAbove(Egg0, Opening(Bowl0)), T1) \quad \text{(N3.4)}$$

$$T0 < T1 \quad \text{(N3.5)}$$

Now the robot prizes apart the two halves of the egg.

[4] A comparison with Hayes's approach to liquids and vessels would be valuable [12].

$$Happens(SplitShell(Egg0), T2) \quad \text{(N3.6)}$$

$$T1 < T2 \quad \text{(N3.7)}$$

Finally, the robot tilts the two egg halves until their contents spill out.

$$Happens(Tilt(Half1(Egg0)), T3) \quad \text{(N3.8)}$$

$$Happens(Tilt(Half2(Egg0)), T3) \quad \text{(N3.9)}$$

$$T2 < T3 \quad \text{(N3.10)}$$

$$\begin{aligned} &Happens(Stop(Half1(Egg0)), t1) \leftarrow \\ &\quad T3 < t1 \wedge HoldsAt(Upturned(Half1(Egg0)), t1) \wedge \\ &\quad\quad \neg\exists t2[HoldsAt(Upturned(Half1(Egg0)), t2) \wedge T3 < t2 < t1] \end{aligned} \quad \text{(N3.11)}$$

$$\begin{aligned} &Happens(Stop(Half2(Egg0)), t1) \leftarrow \\ &\quad T3 < t1 \wedge HoldsAt(Upturned(Half2(Egg0)), t1) \wedge \\ &\quad\quad \neg\exists t2[HoldsAt(Upturned(Half2(Egg0)), t2) \wedge T3 < t2 < t1] \end{aligned} \quad \text{(N3.12)}$$

Proposition 6.1. *Let* N *be the conjunction of all the above axioms numbered* (N...), *let* E *be the conjunction of all the above axioms numbered* (E...), *and let* B *be the conjunction of all the above axioms numbered* (B...). *Finally let* U *be the conjunction of a set of uniqueness-of-names axioms for actions and fluents. Then we have,*

$$\begin{aligned} &CIRC[\mathrm{N}; Happens] \wedge \\ &\quad CIRC[\mathrm{E}; Initiates, Terminates, Releases] \wedge CEC \wedge \mathrm{B} \wedge \mathrm{U} \models \\ &\quad\quad \exists t[t > T3 \wedge HoldsAt(Contains(Bowl0, Middle(Egg0)), t)] \end{aligned}$$

In other words, if an agent carries out the actions in the above narrative, the contents of the egg will end up in the bowl.

Proof. The proof is contained in Appendix A. □

7. Elaborations and limitations

Let's return to the elaborations suggested by Davis and listed in Section 2. Some of these can be straightforwardly accommodated by the above formalisation, and some of them are tricky. What happens if the cook brings the egg to impact very quickly? Or very slowly? Very little extra machinery is required to deal with a slow impact. With the current axioms, if the impact is not "just right", the egg will be unaffected. Instead of (N3.2), we have the following, where $TSEC(v)$ means velocity v is "Too Slow for Egg Cracking" and $TFEC(v)$ means v is "Too Fast for Egg Cracking".

$$TSEC(V)$$

The following axioms guarantee that the different velocity ranges are mutually exclusive.

$$v1 > v2 \leftarrow TSEC(v1) \wedge JREC(v2)$$

$$v1 > v2 \leftarrow JREC(v1) \wedge TFEC(v2)$$

Some extra effect axioms partially take care of the case where the impact is too fast. In addition to (E1.7) we can include the following.

$$\mathit{Initiates}(\mathit{CollidesWith}(\mathit{Shell}(x), y, v), \mathit{Shattered}(y), t) \leftarrow$$
$$\mathit{IsA}(x, \mathit{Egg}) \wedge \mathit{Hard}(y) \wedge \mathit{TFEC}(v)$$

$$\mathit{Terminates}(e, \mathit{Contains}(x, y), t) \leftarrow \mathit{Initiates}(e, \mathit{Shattered}(x), t)$$

However, this elaboration draws further attention to one of the formalisation's major limitations: there is no explicit notion of spatial occupancy or location. The chief consequence of shattering the egg, namely that the shell is no longer capable of holding onto its contents, should be a common sense consequence of the shapes of the egg fragments that result from the impact, and a suitable collection of axioms relating shape and containment. Instead, this fact has been written down as a special case.

The point here is that the medium through which the egg is moving on its way to the rim of the bowl and through which the yolk later falls is the same as that through which the shape of the bowl curves enabling it to retain liquid, namely space. The relationship between the bowl and the egg at the moment of impact is one of extreme spatial proximity. The effect of prizing apart the halves of the egg-shell is one of increasing spatial separation. Yet nowhere in the formalisation is there any indication of the common spatial foundation to the terms *CollidesWith, MovingTo, Contains, Falling, Filling, OpenVessel, ClosedVessel*, and so on. Axiomatising a foundational common sense theory of space, however, is extremely difficult. (See [20] for one attempt.)

Let's move on to some of Davis's other elaborations. What happens if the bowl is made of soft clay? Then, instead of Axiom (B3.3), we have the following.

$$\mathit{Material}(\mathit{Bowl}0, \mathit{Clay})$$

$$\neg \mathit{Hard}(x) \leftarrow \mathit{Material}(x, \mathit{Clay})$$

Now the impact of the egg against the bowl will have no effect. Further axioms would be required to capture the distorting effect of the impact on the shape of the bowl in this case.

What if the bowl is smaller than the egg? To begin with, we have the following axiom instead of (B3.4).

$$\mathit{Volume}(\mathit{Middle}(\mathit{Egg}0)) > \mathit{Capacity}(\mathit{Bowl}0)$$

From Axiom (N2.3), this leads to an $\mathit{Overflows}(\mathit{Bowl}0)$ event, which terminates the *Filling* fluents for both halves of the egg's contents. This prevents a *GotAll* event from occurring for the contents of at least one of the egg-halves. More precisely, there will be three classes of models with respect to the time T of the *Overflows* event. Let $Y1$ and $Y2$ denote the contents of $\mathit{Half}1(\mathit{Egg}0)$ and $\mathit{Half}2(\mathit{Egg}0)$, respectively. In the first class of model, $\mathit{PContains}(\mathit{Bowl}0, Y1, 1)$ holds at time T because a $\mathit{GotAll}(\mathit{Bowl}0, Y1)$ event occurs before T, while $\mathit{PContains}(\mathit{Bowl}0, Y2, k)$ holds at T for some $k < 1$. In the second class of model, $\mathit{PContains}(\mathit{Bowl}0, Y2, 1)$ holds at time T because a $\mathit{GotAll}(\mathit{Bowl}0, Y2)$ event occurs before T, while $\mathit{PContains}(\mathit{Bowl}0, Y1, k)$ holds at T for some $k < 1$. In the third class of model, no *GotAll* events occur, $\mathit{PContains}(\mathit{Bowl}0, Y1, k1)$ holds at time T for some $k1 < 1$, and $\mathit{PContains}(\mathit{Bowl}0, Y2, k2)$ holds at T for some $k2 < 1$. In none of these classes of model can two *GotAll* events occur. Therefore the inference that the bowl eventually contains the whole of the egg's contents is blocked.

The current set of axioms neglects to tell us exactly what happens instead. However, the following axiom can be introduced to complement Axioms (E2.21) and (E2.22).[5]

$$\mathit{Initiates}(\mathit{Overflows}(x), \mathit{Overflowing}(x, y), t) \leftarrow \mathit{HoldsAt}(\mathit{Filling}(x, y), t)$$

Further axioms in the same style as those already written down can be used to describe how the situation evolves while the $\mathit{Overflowing}(x, y)$ fluent holds. Specifically, they would describe the body of liquid y flowing down the side of the vessel x until a surface is reached, and then spreading outwards. To capture the resulting mess in logic might seem a challenge. But thanks to the versatility of the existential quantifier, the challenge can no doubt be met.

Finally, let's briefly consider a class of elaborations not in Davis's list. The formalisation of this paper is quite tolerant to the introduction of interfering events. For example, the cook could drop the egg before it strikes the bowl, or while the contents are pouring out. Similarly, another cook with an irritating sense of humour might remove the bowl at any time. The techniques used here for representing continuous change ensure that knowledge of such extra events can be easily absorbed.

8. Concluding remarks

In Section 1, it was claimed that the exercise of formalising a non-trivial benchmark in common sense reasoning would help to build up a repertoire of knowledge representation principles, techniques, and tricks for later deployment, possibly in a robotics context. To what extent have we succeeded? What principles have been uncovered, and what techniques and tricks have been devised? The topic given the most thorough treatment in the paper is reasoning about action, and this has highlighted the advanced state of development in this area. In particular, the formalisation demonstrates that sophisticated kinds of common sense reasoning about continuous change, including motion, can be captured with the event calculus. Moreover, the set of axioms presented here conforms to a pattern that can be used to represent continuous motion in other domains, such as mobile robotics [23].

This claim is reinforced by the fact that continuous motion has been formulated in a way that can accommodate concurrent actions (such as two simultaneous Tilt events), and triggered events (such as a collision or an overflow). In addition, the formalisation is robust in the presence of certain kinds of incompleteness. Under the right circumstances, precise knowledge of an object's trajectory is not required to license the conclusion that it will collide with another object. This has been achieved through the use of existential quantification in both *Trajectory* formulae, such as (E2.19), and in formulae characterising fluents that are dependent on *Trajectory* formulae, such as (E2.3). These formulae are written in a style that can be mapped onto other domains. Furthermore, proofs that invoke

[5] Ideally, this axiom should have further conditions to block the initiation of the *Overflowing* fluent if a simultaneous *GotAll* event occurs.

these formulae, such as that in Appendix A, exhibit patterns that will recur in reasoning about continuous change in other domains.

In contrast to reasoning about action, the more complex area of common sense reasoning about space and shape is less mature. A proper theory of spatial occupancy would dramatically improve the present formalisation, and this is the subject of ongoing work. Especially striking is the fact that there is no formal link in the present theory between the intuitively connected concepts of motion, parts-and-wholes, and containment. However, the careful incorporation of predicates such as *Distance, Contains,* and *PartOf* has provided a potential interface to a future deep theory of spatial occupancy. And although the theory of parts and containment deployed here is far from fully worked out, it is sufficient to suggest that a full theory could be glued onto an event calculus formulation of continuous change without much difficulty.

The formalisation given in this paper exhibits a good degree of elaboration tolerance, and therefore meets one of the main design criteria set out in Section 2. Many of Davis's suggested elaborations, as well as others, can be accommodated with minimal revision to the axioms. This elaboration tolerance is partly inherited from the ability of the underlying event calculus formalism to absorb extra effect axioms and extra event occurrences, and is partly due to the high level of physical detail in the formalisation.

The egg-cracking benchmark scenario has also been tackled by Lifschitz [15] and Morgenstern [19], so a few words of comparison are in order. For reasoning about action, Lifschitz uses a causal logic based on the work of McCain and Turner [16]. Lifschitz's paper convincingly demonstrates the applicability of this logic to a complex scenario like egg-cracking, but doesn't attempt a deep formalisation of shape, space, or continuous change. These issues are tackled more fully in Morgenstern's work, which merits a closer comparison with the present paper. The comparison is facilitated by the fact that Morgenstern's formalisation, like the present paper, uses the event calculus to represent and reason about action and continuous change, and deploys exactly the same set of predicates and axioms as presented here in Section 3.

Morgenstern's formalisation addresses the representation of shape, and in particular of containment, in a more principled way than the present work. An object's shape is considered as a region of space, and the concepts of a closed and open container are defined in these terms. In addition to straightforward non-porous containers, Morgenstern's paper tackles the concept of a leaking container, a class of object not considered in this paper. On the other hand, the present formalisation can handle simultaneous pourings, which, though not a realistic feature of the egg-cracking scenario, do present an interesting challenge. Moreover, the present paper offers a treatment of parts-and-wholes, which allows an object (such as the egg's contents) to be broken up and to come back together again.[6] Both papers address the issue of continuous change, although Morgenstern's attention is confined to falling and pouring, while the present work also deals with the continuous motion involved in propelling the egg towards the bowl and with the continuous rotation involved in tilting the broken egg.

[6] This is a prerequisite for representing what Morgenstern calls the "Humpty Dumpty theorem", namely that an egg-shell, once broken, cannot be put back together again.

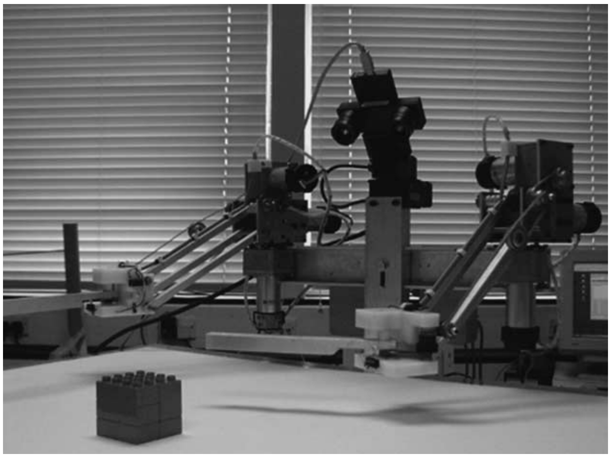

Fig. 1. LUDWIG the humanoid robot.

Finally, this exercise has highlighted certain important methodological issues. In particular, the question arises of the appropriateness of an ontology that has been devised without the requirement that its lowest-level predicate and function symbols are grounded through interaction with the world. The obvious way to deflect this sort of objection is to marry the project of formalising common sense to robotics [23,26]. This is the approach being taken by ongoing work in cognitive robotics at Imperial College, where an upper-torso humanoid robot has been constructed for fundamental research on perception and spatial reasoning (Fig. 1). At present, the main thrust of this work is in the area of vision, and a logic-based theory has been developed that casts visual perception as a form of abduction [27]. Perhaps, in the not too distant future, cognitive robotics research in this style will lead to the sort of deep understanding of common sense reasoning that is surely required to fulfill the long-term ambitions of the field of AI.

Acknowledgements

This work was partly supported by EPSRC grants GR/L20023 “Cognitive Robotics”, GR/N13104 “Cognitive Robotics II” and GR/S02891/01 “Spatial Reasoning and Perception in a Humanoid Robot”. Thanks to Leora Morgenstern, Erik Meuller, Ernie Davis, and the two anonymous referees.

Appendix A. Proof of Proposition 6.1

Let $S = \mathit{Shell}(\mathit{Egg}0)$ and $R = \mathit{Rim}(\mathit{Bowl}0)$. From Proposition 2 in [14], it follows that, since all the *Happens* formulae in N are in Horn clause form, the circumscription of N entails the predicate completion of *Happens*. From this, we derive,

$$\mathit{Happens}(\mathit{PropelTo}(S, R, v), t) \leftrightarrow t = T0 \wedge v = V \tag{A.1}$$

Similarly, from the circumscription of E, we obtain the predicate completions of *Terminates* and *Releases*. (In fact, some formulae in E, such as (E1.8), aren't in Horn clause form. But this can be fixed by replacing the existential quantifiers with Skolem functions.) These completions entail the following.

$$\neg\exists e[\mathit{Terminates}(e, \mathit{Distance}(S, R, z), t)] \tag{A.2}$$

$$\mathit{Releases}(e, \mathit{Distance}(S, R, z), t) \rightarrow \exists v[e = \mathit{PropelTo}(S, R, v)] \tag{A.3}$$

From (B3.6), (A.1)–(A.3), and (EC1), we get,

$$\mathit{HoldsAt}(\mathit{Distance}(S, R, Z), t) \leftarrow 0 \leqslant t \leqslant T0 \tag{A.4}$$

In other words, the distance from the side of the egg to the rim of the bowl is Z up to and including the time of the *PropelTo* action.

Now, from (EC7), (N3.1), (E1.1), (E1.5) and (A.4), it follows that,

$$\begin{aligned} &\mathit{HoldsAt}(\mathit{Distance}(S, R, z), T0 + d) \leftarrow \\ &\quad d > 0 \wedge z = Z - V \cdot d \wedge \neg \mathit{Clipped}(T0, \mathit{MovingTo}(S, R, V), T0 + d) \end{aligned} \tag{A.5}$$

In other words, after the *PropelTo* action, the distance from the side of the egg to the rim of the bowl is a linear function of the elapsed time since the *PropelTo* action, until the first event that clips the *MovingTo* fluent. The next step is to establish the time of that first clipping event.

From the predicate completion of *Happens* we have,

$$\begin{aligned} &\mathit{Happens}(\mathit{Stop}(S), t1) \leftrightarrow \\ &\quad T0 < t1 < T1 \wedge \mathit{HoldsAt}(\mathit{Distance}(S, R, 0), t1) \wedge \\ &\qquad \neg\exists t2[\mathit{HoldsAt}(\mathit{Distance}(S, R, 0), t2) \wedge t2 < t1] \end{aligned} \tag{A.6}$$

Similarly, from the predicate completions of *Terminates* and *Releases*, we have the following.

$$\mathit{Terminates}(e, \mathit{MovingTo}(x, y, v), t) \rightarrow e = \mathit{Stop}(x) \tag{A.7}$$

$$\neg\exists e[\mathit{Releases}(e, \mathit{MovingTo}(x, y, v), t)] \tag{A.8}$$

Let $D = Z/V$. We need to show that the distance between the egg and the bowl reaches zero at time $T0 + D$. First, we prove that no *Stop*(S) event occurs before $T0 + D$, in other words that,

$$\neg\exists t[\mathit{Happens}(\mathit{Stop}(S), t) \wedge t < T0 + D] \tag{A.9}$$

To see that (A.9) is true, assume the first *Stop*(S) event occurs at some time $t1$ before $T0 + D$, where $t1 = T0 + d$ for some $d < D$. Given (A.6), it follows from this assumption that the distance between the egg and the bowl at time $t1$ is zero. Now, from (A.4), we know this distance is Z up to and including time $T0$. Since $Z > 0$, $t1$ must therefore be after $T0$. So, from (A.2), we know that the distance between the egg and the bowl at time $t1$ is $Z - V \cdot d$. (The fluent *MovingTo*(S, R, V) is not clipped before $t1$ by assumption.) We know $d < D$, and therefore $d < Z/V$. If $d < Z/V$, then $Z - V \cdot d > 0$. In other words, the distance from the egg to the bowl at $t1$ is greater than zero, which is a contradiction. Therefore no *Stop*(S) event can occur before time $T0 + D$.

Now, from (A.7)–(A.9) and (EC5), we obtain,

$$\neg \mathit{Clipped}(T0, \mathit{MovingTo}(S, R, V), T0 + D) \tag{A.10}$$

From this plus (A.5), we get,

$$\mathit{HoldsAt}(\mathit{Distance}(S, R, z), T0 + d) \leftarrow 0 < d \leqslant T0 + D \wedge z = Z - V \cdot d \tag{A.11}$$

and, more specifically,

$$\mathit{HoldsAt}(\mathit{Distance}(S, R, 0), T0 + D) \tag{A.12}$$

From (EC3), (N3.1), (E1.1) and (A.10), we also have,

$$\mathit{HoldsAt}(\mathit{MovingTo}(S, R, V), T0 + D) \tag{A.13}$$

From (N3.3) and (A.11), we get,

$$\mathit{Happens}(\mathit{Stop}(S), T0 + D).$$

From (A.12), (A.13) and (N1.1), we get,

$$\mathit{Happens}(\mathit{CollidesWith}(S, R, V), T0 + D)$$

A similar argument to that for (A.9) above shows that no *CollidesWith* event occurs before $T0 + D$. In the rest of the proof, analogous non-occurrences will be assumed without argument.

Now, given the material properties of the egg and the bowl ((B1.1) to (B1.4) and (B3.1) to (B3.3)), it can be seen from Axioms (E1.7) and (N3.2) that the *CollidesWith* event at $T0 + D$ initiates the fluent $\mathit{Fractured}(S)$. From the event calculus axioms and the circumscriptions of N and E, it can be shown that this fluent then persists indefinitely. In other words,

$$\mathit{HoldsAt}(\mathit{Fractured}(S), t) \leftarrow t > T0 + D \tag{A.14}$$

The next event to occur is the robot's *MoveAbove* action at time $T1$ (N3.4). From (E1.10), this event initiates the fluent $\mathit{Beneath}(\mathit{Opening}(\mathit{Bowl}0), \mathit{Egg}0)$. From the event calculus axioms and the circumscriptions of N and E, it can be shown that this fluent then persists indefinitely. In other words,

$$\mathit{HoldsAt}(\mathit{Beneath}(\mathit{Opening}(\mathit{Bowl}0), \mathit{Egg}0), t) \leftarrow t > T1 \tag{A.15}$$

After the *MoveAbove* action, the next event to occur is the robot's *SplitShell* action at time $T2$. From (A.14), we know that the egg-shell is fractured at $T2$. It can also be proved, from (B3.8), that $\mathit{Shell}(\mathit{Egg}0)$ contains $\mathit{Middle}(\mathit{Egg}0)$ at $T2$. Therefore, from (E1.8), this *SplitShell* action will initiate an interval during which the contents of the egg is distributed between the two resulting half-shells. Since no event occurs between $T2$ and the two simultaneous *Tilt* events at time $T3$, we know this interval lasts at least until $T3$. More precisely, we have,

$$\begin{aligned} &\exists y1, y2[\mathit{HoldsAt}(\mathit{Contains}(\mathit{Half}1(\mathit{Egg}0), y1), t) \wedge \\ &\quad \mathit{HoldsAt}(\mathit{Contains}(\mathit{Half}2(\mathit{Egg}0), y2), t) \wedge \\ &\quad\quad \mathit{Comprises}(\mathit{Middle}(\mathit{Egg}0), y1, y2) \leftarrow T2 < t \leqslant T3] \end{aligned} \tag{A.16}$$

Let $Y1$ and $Y2$ be the contents of $Half1(Egg0)$ and $Half2(Egg0)$, respectively, as characterised in (A.16). Now let's consider the aftermath of the two *Tilt* actions that occur at $T3$. From (E2.1) to (E2.3), it can be shown that,

$$\neg HoldsAt(Upturned(Half1(Egg0)), T3)$$

$$\neg HoldsAt(Upturned(Half2(Egg0)), T3)$$

Therefore the two *Tilt* actions initiate $Tilting(Half1(Egg0))$ and $Tilting(Half2(Egg0))$, respectively. Now, the combination of (E2.8) and (E2.3) guarantees the existence of a smallest delay $D1$ and an angle H such that,

$$Trajectory(Tilting(Half1(Egg0)), T3, Angle(Half1(Egg0), H), D1) \quad \text{(A.17)}$$

where

$$HoldsAt(Upturned(Half1(Egg0)), t) \leftarrow HoldsAt(Angle(Half1(Egg0), H), t) \quad \text{(A.18)}$$

Using the same method as for the proof of (A.9), it can be shown that no event affecting $Half1(Egg0)$ occurs between $T3$ and $T3 + D1$. Hence we have,

$$\neg Clipped(T3, Tilting(Half1(Egg0)), T3 + D1)$$

and thus, from (EC7), (A.17) and (A.18),

$$HoldsAt(Upturned(Half1(Egg0)), T3 + D1)$$

Similarly, we obtain, for some delay $D2$,

$$HoldsAt(Upturned(Half2(Egg0)), T3 + D2)$$

From (N2.1), this entails two *Spill* events, at times $T3 + D1$ and $T3 + D2$, respectively. Let's focus on just one.

$$Happens(Spill(Y1, Half1(Egg0)), T3 + D1)$$

Since, from (A.15), we have $Beneath(Opening(Bowl0), Egg0)$ at time $T3 + D1$, this event initiates $Falling(Y1, Opening(Bowl0))$, given (E2.11), (B1.14) and (B1.15). It can now be shown from (E2.15) and (EC7), using the same method as for the proofs of previous event occurrences, that there exists a time $T4 > T3 + D1$ such that,

$$HoldsAt(Distance(Y1, Opening(Bowl0), 0), T4)$$

Therefore, from (N2.2), we have,

$$Happens(Fill(Bowl0, Y1), T4)$$

This initiates $Filling(Bowl0, Y1)$, from (E2.16). Given (B3.4) and (B2.5), it can be proved that no *Overflows* event can occur to terminate the filling via (N2.3). Therefore, from (E2.19), (EC7) and (N2.4), it can be shown that,

$$Happens(GotAll(Bowl0, Y1), T5)$$

for some time $T5 > T4$. From (E2.23), this initiates an interval during which *Bowl*0 contains $Y1$. Since there are no further events that could terminate this interval, we obtain,

$$\mathit{HoldsAt}(\mathit{Contains}(\mathit{Bowl}0, Y1), t) \leftarrow t > T5 \tag{A.19}$$

Concurrently with the story of $Y1$, the remainder of the egg's contents, $Y2$, falls from $\mathit{Half}2(\mathit{Egg}0)$ into *Bowl*0, yielding,

$$\mathit{HoldsAt}(\mathit{Contains}(\mathit{Bowl}0, Y2), t) \leftarrow t > T6 \tag{A.20}$$

for some $T6 > T3 + D2$.

Now, from (A.16), we know that $\mathit{Middle}(\mathit{Egg}0)$ comprises $Y1$ and $Y2$, and from (A.19) and (A.20) we know that both $Y1$ and $Y2$ are contained in the bowl at any time after both $T5$ and $T6$. With these three lemmas, we can show that $\mathit{Middle}(\mathit{Egg}0)$ is contained in the bowl at such a time.

First, note that, since $\mathit{Middle}(\mathit{Egg}0)$ comprises $Y1$ and $Y2$, there can be no part of $\mathit{Middle}(\mathit{Egg}0)$ that is wholly separate from $Y1$ and wholly separate from $Y2$, from (B1.8). In other words, every part of $\mathit{Middle}(\mathit{Egg}0)$ overlaps with either $Y1$ or $Y2$. Now suppose there exists a part P of $\mathit{Middle}(\mathit{Egg}0)$ that is not contained by *Bowl*0. From (B1.10), this entails the existence of some part Q of P that is wholly outside the container. By the transitivity of *PartOf*, Q must also be a part of $\mathit{Middle}(\mathit{Egg}0)$. But if Q is wholly outside the container, then it cannot overlap with either $Y1$ or $Y2$, both of which are contained by the bowl, so it cannot be a part of $\mathit{Middle}(\mathit{Egg}0)$, which is a contradiction. Therefore every part of $\mathit{Middle}(\mathit{Egg}0)$ is contained by *Bowl*0. From (B1.10), this entails that $\mathit{Middle}(\mathit{Egg}0)$ itself is contained by *Bowl*0. In other words, we have,

$$\mathit{HoldsAt}(\mathit{Contains}(\mathit{Bowl}0, \mathit{Middle}(\mathit{Egg}0)), t) \leftarrow t > T5 \wedge t > T6$$

from which the main theorem follows directly. □

References

[1] J. Aloimonos, I. Weiss, A. Bandyopadhyay, Active vision, in: Proc. 1st International Conference on Computer Vision, London, 1987, pp. 35–54.

[2] D.H. Ballard, Animate vision, Artificial Intelligence 48 (1991) 57–86.

[3] S. Coradeschi, A. Saffioti, Anchoring symbols to sensor data, in: Proceedings AAAI-2000, Austin, TX, 2000, pp. 129–235.

[4] E. Davis, Representations of Commonsense Knowledge, Morgan Kaufmann, San Mateo, CA, 1990.

[5] E. Davis, Axiomatizing qualitative process theory, in: Proc. 3rd Internat. Conf. on Principles of Knowledge Representation and Reasoning (KR-92), Cambridge, MA, 1992, pp. 177–188

[6] E. Davis, The kinematics of cutting solid objects, Ann. Math. Artificial Intelligence 9 (1993) 253–305.

[7] E. Davis, Physical reasoning: Cooking problem, The Common Sense 98 Problem Page, http://www.dcs.qmw.ac.uk/conferences/CS98/CS98Problems.html.

[8] A. Galton, Qualitative Spatial Change, Oxford University Press, Oxford, 2000.

[9] S. Harnad, The symbol grounding problem, Phys. D 42 (1990) 335–346.

[10] P.J. Hayes, The naive physics manifesto, in: D. Michie (Ed.), Expert Systems in the Microelectronic Age, Edinburgh University Press, Edinburgh, 1979.

[11] P.J. Hayes, The second naive physics manifesto, in: J.R. Hobbs, R.C. Moore (Eds.), Formal Theories of the Commonsense World, Ablex, 1985, pp. 1–36.

[12] P.J. Hayes, Naive physics 1: Ontology for liquids, in: J.R. Hobbs, R.C. Moore (Eds.), Formal Theories of the Commonsense World, Ablex, Norwood, NJ, 1985, pp. 71–107.

[13] D.B. Lenat, R.V. Guha, Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project, Addison-Wesley, Reading, MA, 1990.

[14] V. Lifschitz, Circumscription, in: D.M. Gabbay, C.J. Hogger, J.A. Robinson (Eds.), The Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning, Oxford University Press, Oxford, 1994, pp. 297–352.

[15] V. Lifschitz, Cracking an egg: An exercise in commonsense reasoning, Unpublished manuscript, 1998.

[16] N. McCain, H. Turner, Causal theories of action and change, in: Proceedings AAAI-97, Providence, RI, 1997, pp. 460–465.

[17] J. McCarthy, Mathematical logic in artificial intelligence, Daedalus (1988) 297–311.

[18] R.S. Miller, M.P. Shanahan, Reasoning about discontinuities in the event calculus, in: Proc. 5th Internat. Conf. on Principles of Knowledge Representation and Reasoning (KR-96), Trento, Italy, 1996, pp. 63–74.

[19] L. Morgenstern, Mid-sized axiomatizations of commonsense problems: A case study in egg cracking, Studia Logica 67 (2001) 333–384.

[20] D.A. Randell, Z. Cui, A.G. Cohn, A spatial logic based on regions and connection, in: Proc. 3rd Internat. Conf. on Principles of Knowledge Representation and Reasoning (KR-92), Cambridge, MA, 1992, pp. 165–176.

[21] P. Simons, Parts: A Study in Ontology, Oxford University Press, Oxford, 1987.

[22] M.P. Shanahan, Default reasoning about spatial occupancy, Artificial Intelligence 74 (1995) 147–163.

[23] M.P. Shanahan, Robotics and the common sense informatic situation, in: Proceedings ECAI-96, Budapest, 1996, pp. 684–688.

[24] M.P. Shanahan, Solving the Frame Problem: A Mathematical Investigation of the Common Sense Law of Inertia, MIT Press, Cambridge, MA, 1997.

[25] M.P. Shanahan, The event calculus explained, in: M.J. Wooldridge, M. Veloso (Eds.), Artificial Intelligence Today, in: Lecture Notes in Artificial Intelligence, vol. 1600, Springer, Berlin, 1999, pp. 409–430.

[26] M.P. Shanahan, Reinventing Shakey, in: J. Minker (Ed.), Logic-Based Artificial Intelligence, Kluwer Academic, Dordrecht, 2000, pp. 233–253.

[27] M.P. Shanahan, A logical account of perception incorporating feedback and expectation, in: Proc. 8th Internat. Conf. on Principles of Knowledge Representation and Reasoning (KR-2002), Toulouse, France, 2002, pp. 3–13.