Ontologies and Representations of Matter

Ernest Davis
Dept. of Computer Science
New York University
davise@cs.nyu.edu

Abstract

We carry out a comparative study of the expressive power of different ontologies of matter in terms of the ease with which simple physical knowledge can be represented. In particular, we consider five ontologies of models of matter: particle models, fields, two ontologies for continuous material, and a hybrid model. We evaluate these in terms of how easily eleven benchmark physical laws and scenarios can be represented.

1. Introduction

Physical matter can be thought of, and is thought of, in many different ways. Determining the true ontology of matter is a question for physicists; determining what ontologies are implicit in the way people think and talk is a question for cognitive scientists and linguists; determining the history of views of the true ontology is a question for historians of science. For those of us who are researching knowledge representations for intelligent systems, the question is, what ontology can be most easily and effectively used as the basis of a knowledge base that will be used for intelligent tasks?

In the long run the answer almost certainly involves using multiple, mutually inconsistent ontologies, with a meta-level structure to choose between them or combine them when necessary. However, at present, such a system would be very hard to design or validate. If we wish to address the representation of matter without taking on the difficulties of inconsistency resolution, the question becomes, what *single* model is most effective as a basis for reasoning?

This paper is a comparative study addressing this question. We consider five ontologies of matter: particle models, fields, two ontologies for continuous matter, and a hybrid model. We evaluate these in terms of how easily eleven benchmark physical laws and scenarios can be represented and in terms of inherent characteristics of the theories.

As each ontology is presented, previous uses of these ontologies in physical reasoning programs and KR theories will be discussed. There is also a large philosophical, linguistic, and KR literature on ontologies for objects, matter, and substance which is not particularly addressed to physical reasoning; some works of of particular interest here are (Bunt 1985; Galton 2004; Sider 2001).

2. Benchmarks

We will use the following eleven benchmarks:

Part-Whole: One piece of matter can be part of another.

Additivity of Mass: The mass of the union of disjoint pieces of matter is the sum of their individual masses.

Rigid solid objects: A solid object moves continuously and maintains a rigid shape. Two objects do not overlap.

Continuous motion of fluids: The shape of a body of fluid deforms continuously.

Chemical reactions. One combination of chemicals can transform into a different combination of chemicals. Chemical reactions obey the **law of definite proportion**.

Continuity at chemical reactions: In a chemical reaction, the products appear at exactly the same time and place as the reactants disappear.

Gas equilibrium: A gas in a stationary or slowly moving closed container quickly attains a state of equilibrium.

Gas laws: A body of gas in equilibrium in a closed container obeys the ideal gas law and the law of partial pressure.

Liquid in a cup: A body of liquid remains at rest at the bottom of a motionless cup. It can be carried without spilling if the cup is moved smoothly and slowly.

Availability of oxygen during combustion: If fuel is burned in air then (ordinarily) oxygen remains available at the point of combustion so that the process can continue.

Passivized surface of metals: In passivization, the surface atoms of a bar of pure metal combine with atmospheric oxygen to form the metallic oxide. The oxide is chemically inert and adheres to the metal, so the interior of the bar remains the pure metal. The problem is to represent a surface that is chemically different from the interior.

This collection of benchmarks, like most such collections, is largely arbitrary. These are problems that have come up in my research that raise striking ontological issues. I did not *seek them out* as examples of ontological issues, but I did *select* them here for that reason. There might well be a different, equally plausible, benchmark set that would lead to quite different conclusions.

3. Ontologies

We consider five ontologies: particles, fields, two ontologies for continuous matter, and a hybrid model.

3.1 Particle ontology

In the particle ontology, matter is composed of a finite collection of molecules; each molecule is a structure composed of atoms. Each atom is an instance of an element. Each molecule is an instance of a chemical, which is a particular arrangement of elements. Since our benchmarks do not involve subatomic particles or radioactive decay, we take atoms to be indivisible and eternal.

There are two variants to consider here. The first is the physically correct model, in which there are on the order of 10^{23} particles in any physical entity of perceptible or manipulable size. Obviously if this model is used to reason about reasonably-sized situations, then it will not be possible to reason about individual particles; in some way one must reason about large collections of particles. I am not aware of any AI program or KR theory that has actually used this.

The second models physical situations in terms of a few dozen or so particles (Gardin and Meltzer 1995; Johnston and Williams 2007). Physical behavior can be simulated by tracking particles and computing their interactions. The resemblance of this to the true particle theory is largely superficial; for most purposes, a system of 10^{23} particles resembles a field solution to a PDE more closely than a system of 30 particles. Collins and Forbus (1987) consider a different particle model of a liquid as a collection of small particles large enough to be characterized by thermodynamic properties but small enough to remain undivided in flow.

3.2 Field ontology

In the *field* ontology, a physical situation is characterized by a collection of spatially distributed scalar and vector fields. This can be done using either continuous or discrete models of space. A discrete model with a reasonably small number of spatio-temporal points (e.g. Funt 1995) can be used directly in a simulation; however, serious difficulties of discretization arise. The continuous model is commonly used in the physics literature; it is the model underlying the Eulerian formulation of fluid dynamics. It has occasionally been used in AI programs (Rajagopalan 1994).

In a pure field theory, all labels at a point or pixel must express purely local features, just as a label on an image pixel represents only the light intensity; labels that express global information, such as the name of an object, are disallowed. Thus there is no direct representation of the identity of "piece of matter" at different times or places. For instance, the swinging of a bat is not viewed as the movement of a persistent object from one place to another; rather, there are some points that are occupied by wood at the start, and at each later time, there are other points occupied by wood. There may be a vector field "Flow" which has a value at each point in space-time, but "Flow" is not associated with some "thing" moving from one place to another.

For solid objects this seems perverse. But imagine the case of pouring and mixing water back and forth between a number of cups. In that case, a question like, "Which part of the water in cup 3 at the end was in cup 17 at the beginning?" seems perverse; there is no meaningful answer, no way to know, and it makes no difference. Here, the Heraclitean view seems reasonable.

There are two variants of the field theory: one permits references to elements, the other prohibits it. For example, if the chemical composition at point P is pure water with density ρ, then one might say that the density of elemental oxygen is 0.89ρ and the density of elemental hydrogen is 0.11ρ or one might consider that concept invalid. Including elements increases the expressivity of the language, but seems unnatural in a non-atomic theory.

3.3 Continuous matter: Chunks

In this ontology and the next, matter is viewed as continuously distributed over space, and the theory tracks "pieces of matter" over time. The best known version of this is the *chunk* ontology. A chunk is a piece of matter that moves around over time. Archetypal examples are a solid object and a pool of fluid in a cup.

Chunks in this sense have been favored in knowledge representation work (Bennett et al. 2000; Davis 1993, 2008) and in certain forms of philosophical analysis (Needham 2002), because they are correspond to the way people talk about, and presumably think about, matter. Therefore, there is a good case to be made that a chunk-based understanding of the world is epistemically prior to a particle-based understanding; closer to direct experience and less a speculative theoretical construct. It is worth bearing in mind that as late as 1900 many physicists disbelieved in the reality of atoms.

Ideally, with these philosophical/cognitive objectives in mind, a chunk-based approach should be agnostic as to the small scale structure of matter and perhaps even of space and time, since ordinary human experience gives no direct evidence about these. This is one of the objectives of mereological theories; however, it is difficult to formulate a rich physical theory that adheres to this. Applications of chunk theory to physical reasoning (Davis 1993, 2008) have generally assumed a continuous model of space, time, and matter. An advantage of this is that it avoids the bumpiness of discrete theories; for example, an object described as a sphere can be an exact sphere, whereas in a particle model it is an approximation of a sphere, accurate to within a few Angstroms. A drawback is that it is susceptible to Zenonian paradoxes; e.g. a solid object can split into infinitely many pieces, of decreasing size (Davis 1992).

The precise definition of a chunk in the context of a rich physical domain is a delicate question; we will consider one definition here and a different one in section 3.5.

The first problem is to identify the class of chunks at some particular moment of time. Davis (1993, 2008) posits that chunks correspond to topologically regular regions. At any time T, any chunk occupies a regular region, and conversely for any regular region R, the piece of matter occupying R at T is a chunk. This, however, is insufficient in dealing with mixtures. Here we further posit that for any regular region R and chemical F, the matter that is in R and that is F constitutes a chunk; e.g. all the water vapor inside a closed jar. This solution somewhat undercuts the motivation of the chunk ontology, as such a chunk is not a perceptible, manipulable entity. Again, as in the field theory, one may choose either to include or to exclude elemental chunks (e.g. the chunk of all the elemental oxygen in a jar).

The second problem is cross-temporal identity. Does the identity of a chunk survive processes such as mixing, cutting, melting, vaporizing, or chemical reactions? Here, we will take a chunk to be an abstraction of a set of molecules; thus the chunk survives mixing, cutting, and phase changes, but not chemical reactions. A chunk, like a molecule, is extant over a certain lifespan; it comes into existence when all its component parts exist and ceases to exist when any part is annihilated (e.g. by a chemical reaction). For example, suppose that a pool of alcohol partially and gradually burns away from time $T1$ to $T2$. Then at any time T in between, the initial entire chunk is no longer extant, but any subchunk that has not yet been partially consumed is still extant.

3.4 Chunks and moleculoids

As we will discuss in section 4.3, it is very difficult to represent **continuity** in the language of chunks. One solution to this problem is to augment the theory of chunks with a continuous theory of particles, which we will call *atomoids* and *moleculoids*. These are analogous to atoms and molecules; each atomoid has an associated element and each moleculoid has an associated chemical. The difference is that, instead of being finite in number, there is a moleculoid and its associated atomoids at each of the uncountably many points of Euclidean space. In fact, there are (countably) infinitely many moleculoids of a given chemical F at each geometric point where the density of F is greater than zero.

This seems wildly far-fetched, even as a theoretical construct, and as far as I know has never before been studied. However, it does have four advantages. First, continuous motion of matter is expressed in the simple constraint that each atomoid moves continuously. Second, this is the implicit model in the Lagrangian formulation of fluid mechanics. The Lagrangian equations track the trajectory of particles at every point of space; so implicitly there *is* a particle at every point of space. Third, it supports a simple characterization of chunks; a chunk is just a set of moleculoids.

Fourth, it abstracts the awkward very small quantities of the true particle theory into zero quantities, which are mathematically more elegant. An atomoid has zero mass (in the same way that a geometric point has zero volume). The atoms in a molecule lie at the same geometric point.[1] A reaction between a collection of molecules can take place when all the molecules involved are literally at the same point. This is why we need infinitely many moleculoids of each chemical at each point; a chemical reaction may involve arbitrarily many instances of the molecule. For instance, we view any instance of the reaction $2H_2 + O_2 \rightarrow 2H_2O$ as involving the conversion of two moleculoids of hydrogen and one of oxygen into two moleculoids of water, all at a single time and point.

3.5 Hybrid ontology

The final ontology is a hybrid model combining particles, fields, and a variant of chunks. The theory views all three models as present simultaneously and coextensively; there are particles, continuous fields, and chunks of stuff in the same place at the same time. Note that this is a single consistent theory, axiomatized in a first-order language.

The behavior of each of these is governed partly by within-model axioms and partly by cross-model axioms. Each constraint is axiomatized in terms of the model where the axiomatization is simplest; the bridge axioms then carry over the consequences to the other models. For example, the constraint that a gas inside a container attains a state of uniform density is stated in terms of the continuous model of density. Combining this with the bridge axiom that the continuous density corresponds roughly to the number of nearby particles gives an implicit constraint on particle distribution. On the other hand, the constraint that matter moves continuously is stated in terms of particles. Combining this with the axiom that the region occupied by a chunk is close to the centers of its component atoms gives an implicit constraint on the shape deformations possible for a chunk.

The model is necessarily rather complex; see (Davis 2010) for a complete description. Here, we will discuss just the nature of chunks and the bridge axioms.

Chunks are defined differently here than in section 3.3. Rather than a fixed set of molecules, a chunk here will be a fluent whose value at each time is a set of molecules. For example, if a metal bar passivizes, it remains the same chunk, though the molecules on the surface change. ((Davis 1993) and (Bennett 2002) likewise consider chunks whose composition changes over time.) Not every fluent over sets of molecules is a chunk; nor is every set of molecules a possible value of a chunk. In fact, we do not give necessary and sufficient conditions for being a chunk. But the theory does enumerate some specific, important categories of chunks, called “basic chunks”: (1) A solid object; (2) A connected pool of liquid; (3) The body of gas within a closed container, and the open atmosphere. (4) For any chemical F and any chunk C in categories 1-3, the part of C that is F; for example, the water vapor inside a closed container. ((Davis 1993) and (Bennett 2002) similarly consider a solid object that maintains its identity over time despite being cut into or having small pieces chipped away.)

The bridge axioms enforce consistency between the distribution of matter as described in terms of fields, molecules, and chunks. For instance, one axiom asserts that if chemical F has positive density at point P, then some molecule of F is near P. (Density is a continuous field.) Another asserts that the mass of chunk C is equal to the integral of the density over a region that approximates the extent of C.

DeCuyper et al. (1995) proposed a somewhat similar hybrid model, combining a chunk model, a small-number particle model, and a model of fields over discrete space. However, the details of this were never worked out.

4. The Evaluation

We now go through the cross-product of ontologies times benchmark examples and discuss how well each ontology deals with each example. Nearly all of the benchmark problems *can* in principle be represented in any of the ontologies; the question is how *easy* and *natural* is that representation

---

[1] This will of course interfere with a theory of the geometric structure of molecules.

and how *useful* is it for supporting commonsense, qualitative inference. The first two criteria are obviously subjective. The third is less so, but it can be determined reliably only after a task domain has been specified, and an inference engine has been built based on each of the ontologies. Also the meaning of a program "using" an ontology is somewhat vague. With these caveats, we proceed.

4.1 Particle model

Part-whole is the subset relation on sets of molecules. **Additivity of mass** follows from the axiom that the mass of a body of matter is the sum of the masses of the atoms.

The elegant representation of **chemical reactions** is one of the strongest features of the particle model; a chemical reaction is the rearrangement of atoms as a new set of molecules. The law of **definite proportion** follows directly.

Continuity of both kinds is the statement that each atom move continuously. **Rigid motion** is the constraint that every molecule in a solid object is rigidly connected to its neighbors. **Non-overlap** is easily expressed if the molecules in a solid object are viewed as packed tightly.

The form of **passivization** follows from the facts that molecules can combine chemically only if they are very close together, and that a solid object cannot be interpenetrated by other matter. Hence, the oxygen in the air can reach the surface atoms of the metal and can react with them, but cannot reach the interior atoms.

However, the remaining representational problems are difficult. In the true molecular theory, **liquid dynamics** are very complicated even in cases that seem mesoscopically simple. A cup of water at rest seems to be just sitting there with a horizontal top; but actually the molecules are in constant thermal motion, constrained by complicated Van der Waals forces; molecules condense from, and evaporate into, the air; and the "horizontal" top is randomly bumpy. Calculating the dynamics of water in a moving cup is much harder.

Some AI systems have used models with small numbers of particles for calculating liquid dynamics (Gardin and Meltzer 1995; Johnston and Williams 2007), but I am not convinced that these reliably give realistic results; for instance, they seem to permit molecules of liquid in a cup to be stacked in a pyramid, like cannon balls. (In the true theory, this cannot occur because of the thermal motion; but these models do not generally incorporate thermal motion.)

The **gas laws** and the **availability of oxygen** rely on probabilistic, statistical facts that are very complex to state and to reason about in a formal theory. For instance, the fact that oxygen in the air remains available for chemical reactions is the statement that *with high probability* a molecule of oxygen will be *very close*, *very soon*. The assertion that the density of a gas is uniformly ρ throughout a container is actually the statement that *with high probability* the total mass of the gas in a *small* region R around any point at any given time is *approximately* $\rho \cdot \text{Volume}(R)$. How these could be handled in a model with few particles is not at all clear.

4.2 Field Ontology

We will first consider a first-order language over the domain of geometric points and time instant; then a language that includes regions and time instants; then a language of four-dimensional regions of space-time.

First-order language of points and instants

We assume a language of points that includes the sorts points, distance, and vectors, with basic operations. The language also contains primitives denoting fields, such as the Boolean field $\text{SolidInterior}(P, T)$, which is true if point P is in the interior of a solid object at time T, or the vector field $\text{Flow}(P, T)$ which is the rate of flow at point P and time T.

Almost none of the benchmarks can be directly represented in this language, because as stated they refer to extended regions and objects, such as "a solid object" or "a pool of liquid". However, what can be done is that the physical laws underlying the inference can be represented in the language of points and then some further external mechanism used to translate between the language of points and the language of extended matter. This would seem wholly inadequate, except that something very similar is *always* necessary in comparing representational systems. To determine whether a problem formulated in one language $\mathcal{L}$ can be solved in theory $\mathcal{T}$ expressed in language $\mathcal{M}$, it is necessary to translate the problem into $\mathcal{M}$; and that translation process is external to $\mathcal{T}$.

The **part-whole** relation and **additivity of mass** actually become vacuous on this view; their entire content gets pushed into the translation mechanism.

Gas equilibrium is characterized in the constraint that the fields pressure, density, and temperature are constant at all interior points of a gas. The **ideal gas laws** are state constraints asserting arithmetic relations over pressure, density, and temperature. **Availability of oxygen** is the constraint that, under specific conditions, the density of oxygen remains positive. A **passivized surface** is a feature of the boundary points of the metal.

The other benchmarks are difficult, using either discrete or continuous space-time. Essentially the only useful language for constraints are difference equations in the discrete case, and PDE's in the continuous case. The physics involved in the benchmarks can be expressed in these languages, though not easily; but deriving qualitative conclusions from these formulations involves very difficult analysis (necessarily in the meta-theory, since the qualitative constraints cannot even be formulated in the object language).

Moreover, the representations of these physical laws require the use of flow fields. These are undesirable, first because they are rather abstract, second because they are highly unstable as regards qualitative boundary conditions, further complicating the integration process. I discuss below only the continuous case; the discrete case is similar.

Continuous motion of fluids without chemical reactions is equivalent to conservation of the mass of each chemical. This can be expressed in the flow PDE $\partial \rho_q / \partial t = \vec{\nabla} \cdot \vec{F}_q$ where ρ_q is the density of chemical q and $\vec{F}_q$ is the flow density of chemical q. There is a further constraint that, if $\rho_q(x) = 0$ throughout a neighborhood of point x, then $\vec{F}_q(0) = \vec{0}$ (if it isn't there it can't flow).

Chemical reactions are characterized in terms of the diminution of the density of the reactants and the increase of the density of the products. In a theory with elemental density, the law of **definite proportion** is expressed in terms of state constraints relating the density of a chemical to the density of its constituent elements (e.g. in pure water, the elemental density of hydrogen is 0.11 times the density of water) and **continuity at reactions** is local conservation of mass of each element, as above. In a theory without elemental density, the law of definite proportion amounts to a constraint on the rates of the consumption of the reactants and the production of the products. Continuity is expressed in the equation $\partial\rho_q/\partial t = \vec{\nabla} \cdot \vec{F}_q + \sum_w \alpha_w \beta_{wq}$ where w ranges over reactions, α_w is a field giving the rate of reaction w at a given point and time, and β_{wq} is the fractional rate of production/consumption of chemical q by w. (α_w is in units of gm/sec·cm^3. β_{wq} is dimensionless.)

Rigid motion of a solid object can be characterized by a set of PDEs over the Flow field, restricted to points where the Boolean field SolidInterior holds: if $\vec{F}$ is the flow field and $\hat{u}, \hat{v}$ are orthogonal then $\partial F_u/\partial v - \partial F_v/\partial u = 0$. Since the PDE's do not hold at the boundaries of a solid object, two abutting objects can move independently. The fact that a solid object cannot split and that two solid objects cannot merge are also consequences of these constraints.

Deducing the possibility of the **rest state of fluid in a motionless cup** involves showing that it is a solution of the incompressible fluid equation; deducing that the fluid will attain this state involves showing that it is the only solution in an equation that includes energy dissipation. Similarly inferring that gas in a closed container **attains equilibrium** requires stating and (at the meta-level) solving the heat equation and the gas dynamic equation. The PDE analysis of **fluid in a moving cup** is extremely difficult.

First order language of regions and instants

Mereologists prefer a language of regions to a language of points. **Part-whole** and **additivity of mass** are now meaningful; they can be expressed in a mereological terms.

Rigid motion can be defined over short time intervals as follows: Region R is defined as *a solid-object region* at time T if it is a maximal connected region where SolidInterior is true. If R is a solid-object region at T and $T1 < T < T2$, then $R, T, T1, T2$ constitutes a *short rigid motion* if, at any TP between $T1$ and $T2$ there is a solid-object region RP which overlaps more than half of R and is congruent to R. (Note here that one is not allowed to characterize an entire function from time to regions.)

Continuity of fluid motion is expressed without using flow fields as follows: Let MassIn(Q, R) be the mass of chemical Q in region R. Then

1. For any Q, R, MassIn(Q, R) is a continuous function of time. (This can be expressed in a first-order language of time instants using the standard ϵ, δ definition.)
2. For any chemical Q, regions R, RO, and times $T1 < T2$, if NTPP(R, RO) and MassIn($Q, RO - R$)=0 throughout $[T1, T2]$ then MassIn(Q, R) is constant throughout $[T1, T2]$. (Flow into or out of R can occur only if matter goes through $RO - R$.) NTPP is non-tangential partial part (Randell, Cui, and Cohn 1992).
3. If Mass(Q, R)=0 at time T then for any RI such that NTPP(RI, R) there exist $T1 < T < T2$ such that Mass(Q, RI)=0 throughout $[T1, T2]$.

Continuity at chemical reactions requires either extending conditions (1-3) to elements or adding the chemical consumption and production of Q in R to conditions (2,3).

Density is uniform distributed over a region R just if the mass of matter in any subregion $R1 \subset R$ is proportional to the volume of $R1$. Uniform temperature and pressure can be defined using functions MaxTemp(R), MinTemp(R), MaxPress(R), MinPress(R). The density is greater than zero throughout R (except on a set of measure zero) if the mass of any subregion is positive. The **gas laws** and **availability of oxygen** are now quite straightforward. Note that the cavity of a closed container is a simple topological predicate.

Liquid in a motionless cup is easily characterized; a static pool of liquid has a boundary consisting of boundaries of solid objects together with a horizontal top. Representing the motion of **liquid in a moving cup** over a short time interval is analogous to rigid motion.

Passivized surface is problematic but not hopeless. Generally mereological theories require regions to be of full dimensionality, and so would askance at physical characteristics that hold only at the two-dimensional surface; but it may be possible to develop a consistent theory that allows it (Galton 2007). Another possibility is to posit that oxygen may seep into the metal to a very thin depth.

First order language of 4D regions

Extending the language of regions and instances include 4D regions of space-time or region-valued fluents (Hayes 1985; Muller 1998) greatly simplifies the representation of **solid objects** and **liquids in a moving cup** as these can be identified with continuous fluents; e.g. a solid object corresponds to a continuous region-based fluent that is always a solid-object region.

4.3 Continuous matter-based theories

The formulations of most of these example in the language of chunks is almost the same as in the theory of 4D regions, but simpler in that the cross-temporal identity relation is built-in rather than derived from continuity constraints. **Part-whole, additivity of mass, gas laws, availability of oxygen** and **liquid in a motionless cup** are essentially identical to the 4D region representation. **Rigid motion** and **liquid in a moving cup** are much the same, but simpler because of the inherent cross-temporal identity. **Passivized surface** encounters the same difficulties as in the region-based representation and has the same possible solutions.

Continuity is not so easy. **Fluid continuity** ought to be expressible purely as the statement that each chunk moves continuously; i.e. the region occupied by each chunk is a continuous function of time. But what it means for a function from time to spatial regions to be “continuous” is debatable. There are a number of different plausible topologies over the domain of spatial regions (Galton 2000) each of which gives a different meaning for "continuity", and it is not clear what the best choice is. (Davis 2008) posits a complex pair of constraints: every chunk moves continuously relative to the volume of the symmetric difference metric, and every chunk contains a subchunk that moves continuously relative to the Hausdorff distance.

Unless one admits chunks of elements, which would vitiate the philosophical justification of using chunks, **continuity at chemical reactions** is much harder. I believe that it can be formulated using the primitive Transforms($C1, C2$) asserting that $C2$ is composed from $C1$ via a chemical transformation. The formulation is too long to be presented here; it is neither easy, natural, nor usable.

As discussed above, if the ontology is extended to include moleculoids and atomoids, then continuity become simply the continuous motion of each atomoid.

4.4 Hybrid theory

All these benchmarks are easy in the hybrid theory. **Part-whole, additivity of mass** and both types of **continuity** are described in terms of particles. **Passivized surface of metals** is described in terms of the outer level of molecules. Motions of **rigid solid objects** and of **liquids** are described in terms of chunks. **Gas laws** and **availability of oxygen** are described in terms of fields.

The difficulties associated with this theory are in formulating the bridge axioms and in establishing that the general theory is consistent with well-posed problem specifications.

4.5 Summary

Two of the ontologies – the theory of chunks with particloids and the hybrid theory – support quite straightforward representations of all the benchmark problems. With all the other ontologies, the representation of some of the benchmark problems is either logically complex, indirect, or awkward. The theory of chunks with particloids has the drawback that particloids are a strange construct; the hybrid theory has the drawback that the bridge axioms are complex and verifying their consistency is difficult.

5. Mechanisms: A General Difficulty

There is a general difficulty in reasoning about chemical reactions using any of these representations. The most obvious form of reasoning to be done is to reason that, under appropriate circumstances, a chemical reaction will take place. For a chemical reaction to take place, there has to be some physical and chemical mechanism; for example, for an iron bar to rust through, the rust must flake off or open up (if this can be prevented, the iron bar will passivize, like aluminum). But these mechanisms are in fact complicated and not always well understood, and certainly the naive reasoner or student can understand the general characteristics of the reaction without understanding the mechanism. After all, the whole point of a qualitative theory is that you can use the knowledge that you have without being required to know everything. I think the problem must be that the fundamental reasoning task is not being correctly characterized; but I have not found a solution.

Acknowledgements

Thanks to Leora Morgenstern and Marsha Berger for helpful suggestions. This research was supported in part by NSF grant IIS-0534809.

References

Bunt, H. 1985. The formal representation of (quasi) continuous concepts. In Hobbs, J. and Moore, R. (eds.) *Formal Theories of the Commonsense World.* Norwood, N.J.: Ablex Pubs.

Bennett B. et al. 2000. Describing Rigid Body Motions in a Qualitative Theory of Spatial Regions. *AAAI-00*, 503-509.

Bennett, B. 2002. Physical objects, identity, and vagueness. *KR-02,* 395-406.

Collins, J. and Forbus, K. 1987. Reasoning about Fluids via Molecular Collections, *AAAI-87,* 590-595.

Davis, E. 1992. Infinite Loops in Finite Time: Some Observations. *KR-92.* 47-58.

Davis, E. 1993. The Kinematics of Cutting Solid Objects. *Annals of Math. and AI,* 9:253-305.

Davis, E. 2008. Pouring Liquids. *Artificial Intelligence,* 173:1540-1578.

Davis, E. 2010. A hybrid model of matter for commonsense reasoning about chemical processes. In preparation.

DeCuyper, J., Keymeulen, D., and Steels, L. 1995. A Hybrid Architecture for Modelling Liquid Behavior, in (Glasgow, Chandrasekaran, and Narayanan 1995) 731-751.

Funt, B.V. 1995. Problem-solving with diagrammatic representations. In (Glasgow, Chandrasekaran, and Narayanan 1995) 33-68.

Galton, A. 2000. *Qualitative Spatial Change.* Oxford: Oxford U. Press.

Galton, A. 2004. Fields and objects in space, time and space-time. *Spatial Cognition and Computation,* **4**(1) 39-67.

Galton, A. 2007. On the paradoxical nature of surfaces: Ontology at the physics/geometry interface. *The Monist,* **90:** 379-390.

Gardin, F. and Meltzer, B. 1995. Analogical Representations of Naive Physics. In (Glasgow, Chandrasekaran, and Narayanan 1995) 669-689.

Glasgow, J., Chandrasekaran, B. and Narayanan, N.H, (eds.) *Diagrammatic Reasoning,* 669-689. Cambridge: MIT Press.

Hayes, P. 1985. Naive Physics 1: Ontology for Liquids, In Hobbs, J. and Moore, R. (eds.) *Formal Theories of the Commonsense World.* Norwood, N.J.: Ablex Pubs.

Johnston, B. and Williams, M.A. 2007. A generic framework for approximate simulation in commonsense reasoning systems. *Commonsense-07,* 71-76.

Needham, P. 2002. Chemical Substances and Intensive Properties. In J. Earley Sr. (ed.) *Chemical Explanation: Characteristics, Development, Autonomy*, Annals of the New York Academy of Sciences, 988:99-113.

Muller, P. 1998. A Qualitative Theory of Motion Based on Spatio-Temporal Primitives. *KR-98,* pp. 131-141.

Rajagopalan, R. 1994. A Model for Integrated Qualitative Spatial and Dynamic Reasoning about Physical Systems. *AAAI-94,* 1411-1417.

Randell, D.A., Cui, Z. & Cohn, A.G.. (1992). A spatial logic based on regions and connection. *KR-92,* 165-176.

Sider, T. 2001. *Four-dimensionalism: A ontology of persistence and time.* Oxford: Oxford U. Press.