

Chapter 2

The Structure of Ill-Structured Solutions: Boundary Objects and Heterogeneous Distributed Problem Solving

Susan Leigh Star

Abstract

The paper argues that the development of distributed artificial intelligence should be based on a social metaphor, rather than a psychological one. The Turing Test should be replaced by the "Durkheim Test," that is, systems should be tested with respect to their ability to meet community goals. Understanding community goals means analyzing the problem of due process in open systems. Due process means incorporating differing viewpoints for decision-making in a fair and flexible manner. It is the analog of the frame problem in artificial intelligence. From analyses of organizational problem solving in scientific communities, the paper derives the concept of boundary objects, and suggests that this concept would be an appropriate data structure for distributed artificial intelligence. Boundary objects are those objects that are plastic enough to be adaptable across multiple viewpoints, yet maintain continuity of identity. Four types of boundary object are identified: repositories, ideal types, terrain with coincident boundaries, and forms.

2.1 Introduction: Larger than Life and Twice as Natural

Artificial intelligence has long relied on natural and social metaphors in a variety of ways, ranging from a source of inspiration for design to attempts at modelling natural information processing.¹ Why?

The reasons have fallen roughly into two categories: attempts at intelligence and attempts at intelligibility. *Attempts at intelligence* have had as their long term goal the creation of a human or biological simulacrum—however that is defined—something that will pass the Turing Test. Metaphors have long been a way of bridging the enormous gap between the current capabilities of machines and the state of the art in computer science, and the complexity and sophistication of natural information processing systems. *Attempts at intelligibility* have had as their long term goal the production of something that will be usable and understandable by human intelligence. Metaphors used for these purposes point to the embeddedness of systems, user-friendliness, situated action, and so forth.

Yet in the metaphoric use of natural information processing, some important considerations become implicit. This especially includes understanding the relationship between the original source of metaphors and the final artifact. Some of the methodological debates in artificial intelligence reflect a deep uncertainty about the status of natural metaphors. Would a completely formal system allow them at all? If one is committed to a formal system, wherein does the fidelity to nature lie? Or do natural and artificial systems share formal properties to be discovered? ([Hall and Kibler 1985] review these issues.) Many of these concerns are being brought to light by research in distributed artificial intelligence. This is first because the original Turing goal could not be met by distributed work, and secondly because the social, not the psychological or the biological, appears to many researchers in the field both as an important metaphor and as part of the system.

2.2 From the Turing Test to the Durkheim Test

The original Turing Test [Turing 1950] involved a computer being able to mimic a woman well enough so that a human observer could not distinguish between a human male and a "female" computer. The test was predicated on a closed-universe model using "discrete state digital computers."

¹There are numerous descriptions of attempts to use such models; see e.g., [Ericsson and Simon 1979] for a review of sources of evidence on cognition.

rat
the
can
act
knc
nur
the
The
mit
the
as
for
dig

Later in
long as r
Turix
going ba
of artific
confusio
it also h
intellige:

- tes
- mi:
- co
- be
- lar
- th-

Criti
for some
world sy
informa:
ponents.

"The prediction which we are considering [of all future states] is, however, rather nearer to practicability than that considered by Laplace. The system of the 'universe as a whole' is such that quite small errors in the initial conditions can have an overwhelming effect at a later time . . . Even when we consider the actual physical machines instead of the idealized machines, reasonably accurate knowledge of the state at one moment yields reasonably accurate knowledge any number of steps later . . . Provided it could be carried out sufficiently quickly the digital computer could mimic the behavior of any discrete-state machine. The imitation game could then be played with the machine in question and the mimicking digital computer and the interrogator would be unable to distinguish them. Of course the digital computer must have an adequate storage capacity as well as working sufficiently fast. Moreover, it must be programmed afresh for each new machine which it is desired to mimic . . . This special property of digital computers . . . is described by saying that they are *universal machines*."

Later in the article Turing reiterates that these computers can meet any new situation, as long as they have enough storage capacity.

Turing's model is more than a quaint, outdated vision of what computers can do. By going back to the original source, some fundamental values (and value conflicts) in the field of artificial intelligence are revealed, and therein some of the reasons for the ambivalence and confusion about metaphors. Turing's test world is closed, as already pointed out above. But it also has the following properties that are being hotly contested by distributed artificial intelligence at this time:

- testing is done by individuals, not communities. There is no doubt in the tester's mind about what constitutes a valid result (in this case, stereotyped female behavior);
- computers, because they are programmable, are universal. Once a situation can be formally analyzed, it becomes amenable to understanding through this universal language;
- the only restriction on intelligence is lack of storage capacity (or processing power).

Critiques of these propositions have been coming from distributed artificial intelligence for some time. For example, Hewitt's open systems model posits that all nontrivial real-world systems are open. They include properties of the real world, including distributed information processing, asynchronous updates, arms' length relationships between components, negotiation, and continual evolution [Hewitt and DeJong 1983], [Hewitt 1986],

[Hewitt 1988]. These systems are open in several senses: there is no global temporal or spatial closure, and there is an absence of a central authority. Thus, rigid *a priori* protocols that will homogenize data and decision-making both beg the question of openness and limit the problem-solving capacity of the system in the real world. Flexibility and evolution are the central concerns.

No amount of increased storage capacity can bypass the problems posed by open systems. The structure of the original Turing Test, relying solely on a fixed repertoire of rules in order to mimic a range of behaviors, cannot accommodate this type of distributed system. The reasons are the same as for Hewitt's original critique: it could not analyze conflicting viewpoints within the system, and the fundamentally open nature of real world systems inevitably gives rise to such conflicts.

The conceptual struggle in distributed artificial intelligence has been with the tensions implied by the idea of a universal formal language and the inconsistency which arises from the distributed, open nature of the system itself. For example, [Durfee and Lesser 1987] propose the idea of partial global plans that dynamically model and incorporate the findings from distributed nodes of a system, maintaining the openness of the system but achieving coherence across nodes. [Cammarata *et al.* 1983] state that

"A main challenge to distributed problem solving is that the solutions which a distributed agent produces must not only be locally acceptable, achieving the assigned tasks, but also they must be interfaced correctly with the actions of other agents solving dependent tasks. The solutions must not only be reasonable with respect to the local task, they must be *globally coherent* and this global coherence must be achieved by *local computation alone*."

In response to this challenge, the metaphors in this line of work have gone from single humans or human psychology² to organizations, interactions, negotiation, blackboards, networks and communities. For example, [Fox 1981] discusses the "technology transfer" possible between human organizations and artificial intelligence systems; [Gasser 1987] calls for cooperation between distributed artificial intelligence and other fields of study concerned with coordinated action and distributed problem solving. I propose that this change in metaphoric base be recognized by replacing the vision of the Turing Test with a test adequate to meet the challenges of distributed open systems: the Durkheim Test.

Emile Durkheim (1858-1917) was a French sociologist who attempted to demonstrate the irreducible nature of what he called "social facts." For example, you could not understand differential suicide rates in different locations by simply saying that each case was

²I include network models of cognition here; I mean that the metaphors have moved away from individualist, black boxed models of single actors.

Why not
collective
behavior

patholog
terms of
the follo
facts pr
the codi
social e
The
This m
system.
never p
"When
each ne
Thus, t
Durkhe
It is als
order t
commu

So
of a sy
applied
pass m
of an c
design
sign an
blur. I
process
ipant s
for hur
system
of the

The
power.
closed.
ecolog
both t
forms

pathological; something was happening at the "system level" that did not reduce to the terms of lower levels. Social facts, he said, are thus *sui generis* (or irreducible). He proposed the following law: "The determining cause of a social fact should be sought among the social facts preceding it and not among the states of the individual consciousness," followed by the codicil: "The function of a social fact ought always to be sought in its relation to some social end" [Durkheim 1938].

The test of intelligence of a distributed open system is necessarily an ecological one. This means that it is *sui generis* at the social/system level, incorporating all parts of the system. Testing one node will not give reliable results; testing the whole open system is never possible (see e.g., [Lesser and Corkill 1981]. In the words of [Davis and Smith 1983], "When control is decentralized, no one node has a global view of all activities in the system; each node has a local view that includes information about only a subset of the tasks." Thus, the very concept of a test must change in order to deal with such systems. Following Durkheim, we can say that it would be communal, irreducible, distributed, and dynamic. It is also important to note that it cannot be applied solely after a design is complete. In order to understand the acceptance and use of a machine in and by a community, that community must be actively present as it evolves.

So the Durkheim test would be a real time design, acceptance, use and modification of a system by a community. Its intelligence would be the direct measure of usefulness applied to the work of the community; its ability to change and adapt, and to encompass multiple points of view while increasing communication across viewpoints or parts of an organization. Such a test also changes the position of metaphors with respect to design and use considerations. In an open, evolving system, the boundaries between design and use, between technology and user, between laboratory and workplace, necessarily blur. Neither is the organization of work something that can be added after the design process [Kling and Scacchi 1982]. [Chang 1987] develops a model of this he calls participant systems. Thus, social metaphors may remain sources of inspiration, or guidelines for human-computer interface. But if we are stringently to apply the principles of open systems to design, and account for differing viewpoints and evaluation criteria at every step of the way, social systems become deeply implicated at all times.

The futility of the Turing Test comes not from lack of storage capacity or processing power, but from a fundamental misunderstanding of the nature of computers and society as closed, centralized, and asocial. As that misunderstanding gets replaced by an open system, ecological, and political model of organizations, workplaces, and situations (which include both machines and human organization), the Turing Test will be replaced by different forms of evaluation. (For a discussion of this from both sociological and computer science

perspectives, see [Bendifallah *et al.* 1988].)

2.3 Due Process, the Frame Problem, and Scientific Communities

As noted above, the distributed and open nature of real systems gives rise to the existence of different viewpoints within the system. A viewpoint in this sense can occur at any level of organizational scale, from hardware to human organization. It can arise from, for example, asynchronous updates to a knowledge base, resulting in different ways of processing information at different nodes based on differences in a knowledge base. At higher levels, it can result from differences in the structure of tasks performed, different commitments, or different long or short term goals.

The simultaneous existence of multiple viewpoints and the need for solutions which are coherent across divergent viewpoints is a driving consideration in distributed artificial intelligence. [Hewitt 1986] and [Gerson 1987] have discussed aspects of this as the problem of *due process*: a legal phrase that refers to collecting evidence and following fair trial procedures. The due process problem in either a computer or human organization is this: in combining or collecting evidence from different viewpoints (or heterogeneous nodes), how do you decide that sufficient, reliable and fair amounts of evidence have been collected? Who, or what, does the reconciling, according to what set of rules?

[Davis 1980] notes that cooperation is necessary in order to resolve this class of problems, but that many researchers who came to distributed processing via attempts to synthesize networked machines see cooperation as a form of compromise "between potentially conflicting views and desires at the level of system design and configuration." The two motivations he suggests for cooperation are insolubility by a single node and compatibility (joining of forces).

The interdependence suggested by these motivations would seem to work against pluralism of viewpoints. How can two entities (or objects or nodes) with two different and irreconcilable epistemologies cooperate? If understanding is necessary for cooperation, as is widely stated in the distributed artificial intelligence literature, what is the nature of an understanding that can cooperate across viewpoints?

There is a fundamental similarity between these concerns about cooperation, i.e., the due process problem, and the frame problem in artificial intelligence. The frame problem, as [Hayes 1987] notes, "arises when the reasoner is thinking about a changing, dynamic world, one with actions and events in it ... it only becomes an annoyance when one tries to describe a world of the sort that people, animals and robots inhabit." It is a problem,

he states,
temporal

Spatial
incompati
and space
on viewpo
depending
problem [

From t
lem are fi
and chang
cess means
making pr
bases. It is
patibility.
process m
[Hewitt 19

The fra
formation
moving th
out having
implosion."
it is an inte
objects: ac
robot is for
geneous, ex
closed worl

The rec
mischaracte
actual cont
as a reconc
of actions
because env
adjudicatin

³Sociologis
inertia is stru
open systems.

he states, not in computation, but in representation; it occurs in the presence of spatial or temporal change.

Spatial or temporal change is significant in this regard because of the epistemological incompatibilities that such change may bring about. As an actor moves through time and space, new information or new axiomatic requirements evolve (or devolve, depending on viewpoint), thus shifting the assumptive frame. Which axioms to retain or change, depending on which things can be taken for granted (or not) is at the heart of the frame problem [Pylyshyn 1987].

From the viewpoint of open systems, the problems of due process and the frame problem are figure-ground to one another. In the problem of due process, viewpoints evolve and change with new information and new situational constraints. The concept of due process means evaluating and synthesizing potentially incompatible viewpoints in the decision-making process: adducing evidence. The problem is one of drawing on different evidentiary bases. It is the differences in situation and viewpoint that make for epistemological incompatibility. In open systems, the lack of a sovereign arbiter means that questions of due process must be solved by negotiation, rules and procedures, case precedents, etc. (see [Hewitt 1988]).

The frame problem arose in the context of dealing with moving actors, absorbing information in a fashion that threatens the stability of their axiomatic structure. A robot, moving through novel open space, must find a robust way to deal with that novelty without having to add so many new axioms that it becomes bogged down in a "combinatorial implosion." But the problem is *not* really one of moving *through* neutral territory: in fact, it is an interactional problem. *Environment* really means a series of interactions with other objects: actors, events, and new kinds of ordered actions. In other words, the moving robot is forced to evaluate a series of interactions by picking and choosing from the heterogeneous, evolving, potentially incompatible viewpoints of other actors outside its original (closed) world.

The reconciliation between multiple viewpoints in the frame problem has thus been mischaracterized as a single actor problem. In fact, viewed temporally, and taking the actual content of the changing environments into account, the frame problem can be seen as a reconciliation between old and new experience in the same actor through a series of actions in open, distributed space.³ The content of this experience is interactional because environments are a set of new actors and events. Solving the frame problem means adjudicating decisions about which evidence is important for which circumstances, and

³Sociologists discuss this as the problem of continuity of identity. (See [Strauss 1969].) The problem of inertia is structurally similar to the track-record heuristic described by [Hewitt 1986] in his discussion of open systems.

which can be taken for granted. The continuity of the robot's actions relies on a set of metarules that are structurally identical to the due process problem: What data does it take from which viewpoint? What is kept and what is discarded (thus the many discussions of relevance and inertia in the frame problem literature)? How can a decision be reached that incorporates both novelty and sufficient closure for action?

Human actors routinely solve both the frame problem and the due process problem. They do so in a variety of ways, as noted both in the social science literature and in the frame problem literature, and in variably democratic ways. In the remainder of this paper I present one class of strategies employed by two scientific communities I have studied in some detail.

The studies began as an exploration of the scientific community metaphor in a long collaboration with Carl Hewitt. We analyzed issues that arose in the context of artificial intelligence research by looking at how human communities resolved them. These included issues such as due process [Gerson 1987], the resolution of conflict in a distributed community [Star 1989a], triangulation of evidence from domains with incompatible goals [Star 1986], resolution of local uncertainty into global certainty [Star 1985], local constraints on representing complex information [Star 1983], and the management of anomalous information [Star and Gerson 1987].

After some years, with the development of the open systems model and the evolution of our own social science work, the "metaphor gap" seems to be closing.⁴ The status of the social/community metaphor in the face of real-world systems embedded in organizations has shifted as the boundaries of "computer," "system," and "actors" are perceived as being larger and wider than Turing's closed world model. Because advances in both artificial intelligence and social science call for the development of new *ecological units of analysis, methods, and concepts*, both the *content and role of metaphors have shifted*.

The concept of boundary objects as presented below thus is simultaneously metaphor, model, and high-level requirement for a distributed artificial intelligence system. The more seriously one takes the ecological unit of analysis in such studies, the more central human problem-solving organization becomes to design—not simply at the traditional level of human-computer interface, but at the level of understanding the limits and possibilities of a form of artificial intelligence [Star 1989b].

⁴Another factor may contribute to closing the gap. The metaphor, as a source of inspiration, models, or design specifications, works both ways: artificial intelligence is also a metaphor for sociological research (see [Star 1989a] for a discussion of this process)!

2.4

[Kornfeld
source of
distributed
have arms
tendency of ar
The inform
different k
contingenc

Scienti
continually
station gi
piecemeal
years.

Scienti
informatic
toward co
by locatio
fluid [Lato
are consta

Yet wi
ings. They
hold up w
nated Hew
how is rob
entific con
coherent t.

Any sc
meets thos
that prese
the frame
of scientifi
as it is ac
words of e
[Wimsatt

2.4 The Scientific Community and Open Systems

[Kornfeld and C. Hewitt 1981] proposed that the scientific community be taken as a good source of metaphors for open systems work. Because real world information systems are distributed and decentralized, they evolve continuously, embody different viewpoints, and have arms-length relationships between actors requiring negotiation. The internal consistency of an open system cannot be assured, due to its very character as open and evolving. The information in an open system is thus heterogeneous, that is, different locales have different knowledge sources, viewpoints, and means of accomplishing tasks based on local contingencies and constraints.

Scientific workplaces are open systems in Hewitt's sense of the term. New information is continually being added asynchronously to the situation. There is no central "broadcasting" station giving out information simultaneously to scientists. Rather, information is carried piecemeal from site to site (when it is carried at all), with lags of days, months, or even years.

Scientific work is distributed in this way. Thus, there is no guarantee that the same information reaches participants at any time, nor that people are working in the same way toward common goals. People's definitions of their situations are fluid and differ sharply by location; the boundaries of a locality or workplace are simultaneously permeable and fluid [Latour 1988]. Scientific theory-building is deeply heterogeneous: different viewpoints are constantly being adduced and reconciled.

Yet within what may sound like near chaos, scientists manage to produce robust findings. They are able to create smooth-working procedures and descriptions of nature that hold up well enough in various situations. Their ability to do so was what originally fascinated Hewitt and Kornfeld. In the absence of a central authority or standardized protocol, how is robustness of findings (and decision-making) achieved? The answer from the scientific community is complex and twofold: they create objects that are both plastic and coherent through a collective course of action.

Any scientific workplace can thus be described in two ways: by the set of actions that meets those local contingencies that constantly buffet investigators, or the set of actions that preserves continuity of information in spite of local contingencies (due process and the frame problem simultaneously). Understanding this requires a different appreciation of scientific theories than that traditionally put forward by philosophers. Scientific truth *as it is actually created* is not a point-by-point closed logical creation. Rather, in the words of ecologist Richard Levins, "our truth is the intersection of independent lies" (in [Wimsatt 1980]). Each actor, site, or node of a scientific community has a viewpoint, a

partial truth consisting of beliefs, local practices, local constraints, and resources—none of which are fully verifiable across all sites. *The aggregation of those viewpoints is the source of the robustness of science.*

2.5 Heterogeneous Problem Solving and Boundary Objects

In the face of the heterogeneity produced by local constraints and divergent viewpoints, how do communities of scientists reconcile evidence from different sources? The problem is an old one in social science; indeed, one could say it reflects the core problem of sociology. One major concern of early sociologists, such as Robert Park and Georg Simmel, was to describe interaction between participants from groups (or worlds) with very different "definitions of the situation." This concern gave rise to a series of case studies of ethnicities, work groups, and subcultures now grouped loosely under the rubric "Chicago school sociology." Everett Hughes, a leader of this group, argued for an ecological approach to understanding the participation of heterogeneous groups within a workplace, neighborhood, or region. By this he meant that the different perspectives, or viewpoints, of the participants need to be understood in a *sui generis* fashion, not simply as a compilation of individual instances, and as situated action.

Some findings from our studies of scientists of potential interest to distributed artificial intelligence are that scientists

1. cooperate without having good models of each other's work;
2. successfully work together while employing different units of analysis, methods of aggregating data, and different abstractions of data;
3. cooperate while having different goals, time horizons, and audiences to satisfy.

They do so by creating objects that serve much the same function as a blackboard in a distributed artificial intelligence system. I call these *boundary objects*, and they are a major method of solving heterogeneous problems. Boundary objects are objects that are both plastic enough to adapt to local needs and constraints of the several parties employing them, yet robust enough to maintain a common identity across sites. They are weakly structured in common use, and become strongly structured in individual-site use.

Like the blackboard, a boundary object "sits in the middle" of a group of actors with divergent viewpoints. Crucially, however, there are different types of boundary objects depending on the characteristics of the heterogeneous information being joined to create

them.
joining
paper p
to the c
humans

2.6

In study
The ana
rated ra
of a cor
This gr
tors, at
colleagu
ley (Sta
amateur
conserva

Wha
intellige
different
Space pr
are summ

First.
clinical r
theoretic
(an insta
alizations
events su
theory to
rected to
about the
from the e
nally, for
side effect
make their

one of source
lary
s, how
1 is an
r. One
scribe
itions
work
y."
nding
n. By
to be
ances,
ificial
ds of
in a
or
both
oying
eakly
with
jects
reate

them. The combination of different time horizons produces one kind of boundary object; joining concrete and abstract representations of the same data produces another. Thus, this paper presents not just one blackboard, but a system of blackboards structured according to the dynamic, open-systems requirements of a community (including both machines and humans).

2.6 Types of Boundary Objects

In studying scientists, I identified heterogeneous subgroups within the scientific workplace. The analysis of boundary objects presented here draws on two case studies that incorporated radically different viewpoints in the conduct of work. First, I conducted a study of a community of neurophysiologists at the end of the nineteenth century in England. This group included both clinical and basic researchers, as well as hospital administrators, attendants, experimental animals, journalists and patients [Star 1989a]. Second, my colleagues and I conducted a study of a zoological museum from 1900-1940 at Berkeley [Star and Griesmer 1989], [Gerson 1987]. This group included professional biologists, amateur collectors, university administrators, animals, and local trappers, farmers, and conservationists.

What is interesting about these studies from the point of view of distributed artificial intelligence is that the structure and attributes of the information brought in from the different participants were distributed and heterogeneous, yet were successfully reconciled. Space prohibits a detailed discussion of all the differences in viewpoint, but two salient ones are summarized below.

First, in comparing clinical and basic research evidence, the following differences obtain: clinical research operates with a much shorter time horizon (cure the patient, not find the theoretical generalization) than basic research; the case is the unit of analysis for clinicians (an instance-based form of explanation), whereas for basic researchers it is analytic generalizations about classes of events. In clinical research, attention is directed toward concrete events such as symptoms, treatments and patient trajectories. Diagnosis draws on medical theory to validate concrete observations of this nature. In basic research, attention is directed toward analytic generalizations such as refinements to others' theories, statements about the applicability of an experiment to a larger body of knowledge. Work proceeds from the experimental situation and is directed outwards toward a body of knowledge. Finally, for the clinician, interruptions to work come in the form of complications, which are side effects to be dealt with locally and discarded from the evidentiary body (they never make their way into publication of the cases). Interruptions to work for the basic researcher

come in the form of anomalies which must be accounted for in the body of evidence, either by controlling them or introducing them into the findings.

Second, in the world of the natural history museum, one primary source of comparison is between amateur and professional biologists. There are some similar differences as between clinicians and basic researchers. For the amateur collector of specimens, the specimen itself is the unit of analysis—a dead bird or a bone found in a specific location. Collecting, like clinical work, is the art of dealing on an instance-by-instance basis with examples and local contingencies. For the professional biologist, on the other hand, the specimens collected by amateurs form a part of an abstract generalization about ecology, evolution, or the distribution of species. The particular bug or beetle is not as important as what it represents. Furthermore, the work organization is highly distributed, ranging from the museum in Berkeley to various collecting expeditions throughout the state of California.

In analyzing these types of heterogeneity, I found four types of boundary objects created by the participants. The following is not an exhaustive list by any means. These are only analytic distinctions, in the sense that we are really dealing here with *systems* of boundary objects that are themselves heterogeneous.

2.6.1 Repositories

These are ordered piles of objects that are indexed in a standardized fashion. Repositories are built to deal with problems of heterogeneity caused by differences in unit of analysis. An example of a repository is a library or museum. They have the advantage of modularity.




Figure 2.1: Boundary object: repositories

2.6.2 Ide

This is an ob... details of any... However, it is... a means of co... parties. Exam... described no... a means of co... degree of abstr... in the deletio... adaptability.

Condit
Result
and gi
Advan

2.6.3 To

These are co... They arise i... tributed ove... each site car... area with th... of coinciden... ject for worl...

⁹See [Wims

2.6.2 Ideal Type or Platonic Object

This is an object such as a map or atlas which in fact does not accurately describe the details of any one locality. It is abstracted from all domains, and may be fairly vague. However, it is adaptable to a local site precisely because it is fairly vague; it serves as a means of communicating and cooperating symbolically—a sufficient road map for all parties. Examples of platonic objects are the early atlases of the brain, which in fact described no brain, which incorporated both clinical and basic data, and which served as a means of communicating across both worlds. Platonic objects arise with differences in degree of abstraction such as those that obtain in the clinical/basic distinction. They result in the deletion of local contingencies from the common object, and have the advantage of adaptability.

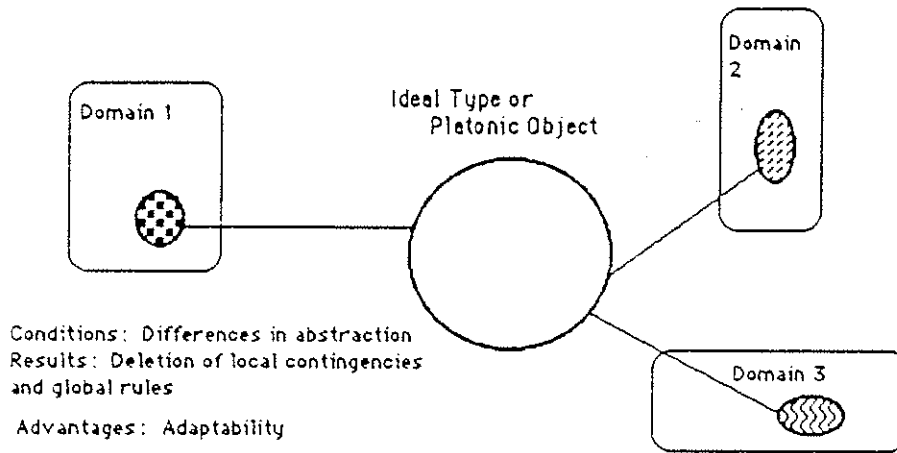


Figure 2.2: Boundary object: platonic object

2.6.3 Terrain with Coincident Boundaries

These are common objects which have the same boundaries but different internal contents.⁵ They arise in the presence of different means of aggregating data and when work is distributed over a large-scale geographic area. The result of such an object is that work in each site can be conducted autonomously, but cooperating parties can work on the same area with the same referent. The advantage is the resolution of different goals. An example of coincident boundaries is the creation of the state of California itself as a boundary object for workers at the museum. The maps of California created by the amateur collectors

⁵See [Wimsatt 1980] for a fuller discussion of these issues.

and the conservationists resembled traditional roadmaps familiar to us all, and emphasized campsites, trails, and places to collect. The maps created by the professional biologists, however, shared the same outline of the state (with the same geopolitical boundaries), but were filled in with a highly abstract, ecologically-based series of shaded areas representing "life zones," an ecological concept.




Figure 2.3: Boundary object: terrain with coincident boundaries

2.6.4 Forms and Labels

These are boundary objects devised as methods of common communication across dispersed work groups. Both in neurophysiology and in biology, work took place at highly distributed sites, conducted by a number of different people. When amateur collectors obtained an animal, they were provided with a standardized form to fill out. Similarly, in the hospital, night attendants were given forms on which to record data about patients' symptoms of epileptic fits in a standardized fashion; this information was later transmitted to a larger data base compiled by the clinical researchers attempting to create theories of brain and nervous system function. The results of this type of boundary object are standardized indexes and what Latour would call "immutable mobiles" (objects that can be transported over long distance and convey unchanging information). The advantages of such objects are that local uncertainties (for instance, in the collecting of animals or in the observation of epileptic seizures) are deleted. Labels and forms may or may not come to be part of repositories.

Condi
 Resu
 stanc
 Adva

2.7 Su

What are the
 of boundary
 of the sort c
 his terminol
 combinatori
 representati
 respects loca
 a logical Esp
 should search
 above produ
 they are ill-s
 functional an

The prob
 require a de
 constraints a
 on cooperati
 problem solv
 inhabit differ
 should be use

Future dir

Forms and Labels

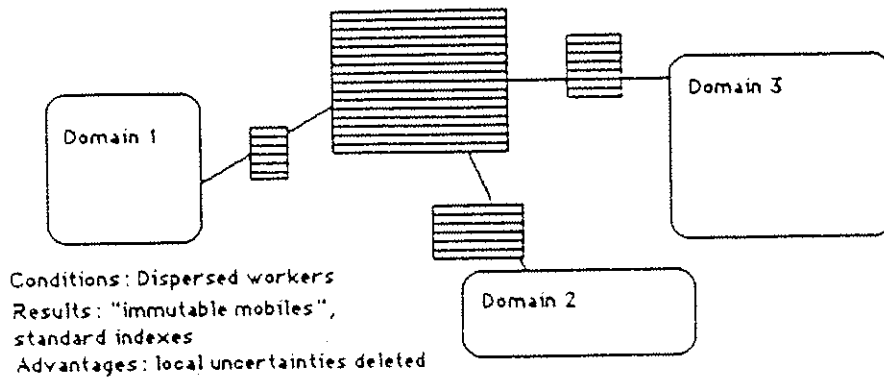


Figure 2.4: Boundary object: forms/labels

2.7 Summary and Conclusions

What are the implications for distributed artificial intelligence of understanding the creation of boundary objects by scientists? First, boundary objects provide a powerful abstraction of the sort called for by [Chandrasekaran 1981] to organize blackboards. They are, to use his terminology, neither committee nor hierarchy. They bypass the sort of problems of combinatorial implosion feared by Kornfeld and also bypass hierarchical delegation and representation. Unlike Turing's universal computer, the creation of boundary objects both respects local contingencies and allows for cross-site translation. Instead of a search for a logical Esperanto, already proved impossible in a distributed open systems context, we should search for an analysis of such objects. Problem-solving in the contexts described above produces workable solutions that are not, in Simon's terms, well-structured. Rather, they are ill-structured: they are inconsistent, ambiguous, and often illogical. Yet, they are functional and serve to solve many tough problems in distributed artificial intelligence.

The problems of instantiating descriptions in distributed systems [Pattison *et al.* 1987] require a device similar to the creation of boundary objects for accounting for shifting constraints and organizational structures. [Durfee *et al.* 1987] suggest a system that relies on cooperation and plan-based nodes that arrive at locally complete solutions for distributed problem solving. Again, the notion that systems of actors create common objects that inhabit different nodes in different fashions, and are thus locally complete but still common, should be useful here.

Future directions for research on these questions would include the following:

1. expanding the taxonomy of boundary objects and refining the conceptions of the types of information used in their construction;
2. examining the impact of combinations of boundary objects, and beginning to develop a notion of systems of such objects;
3. examining the problem of scaling up or applying an ecological, human/machine analysis to what is called multigrained systems in [Gasser *et al.* 1986].

The "Durkheim Test" referred to in the beginning of this paper is important in evaluating the construction and use of boundary objects. That is, the construction of such objects is a community phenomenon, requiring at least two sets of actors with different viewpoints. Analysis of the use of such an object at only one point in the system, or apart from its relationship to other nodes, will produce a systematic reductionist bias of the sort described by [Wimsatt 1980]. Heuristics used in such a fashion will reflect the neglect of the *sui generis* nature of the system. Furthermore, if the ecological unit of analysis recommended here and elsewhere in artificial intelligence is adopted, it should be noted that human designers, users, and modifiers of the computer systems involved will *make* boundary objects out of the information systems at every stage of the information processing trajectory.

Acknowledgements

Conversations with Geof Bowker, Lee Erman, Les Gasser, James Griesemer, Carl Hewitt, Rob Kling, Steve Saunders, Randy Trigg and Karen Wieckert were very helpful in formulating the ideas expressed here.

References

- [Bendifallah *et al.* 1988] S. Bendifallah, F. Blanchard, A. Cambrosio, J. Fujimura, L. Gasser, E. M. Gerson, A. Henderson, C. Hewitt, W. Scacchi, S. L. Star, L. Suchman, and R. Trigg, "The Unnamable: A White Paper on Socio-Computational Systems," unpublished draft manuscript available from Les Gasser, Department of Computer Science, University of Southern California, Los Angeles, California, 1988.
- [Cammarata *et al.* 1983] S. Cammarata, D. McArthur, and R. Steeb, Strategies of cooperation in distributed problem solving. *Proceedings IJCAI-83*, Karlsruhe, West Germany, August 1983, pp. 767-770.
- [Chandrasekaran 1981] B. Chandrasekaran, "Natural and Social System Metaphors for Distributed Problem Solving: Introduction to the Issue," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-11, No. 1, 1981, pp. 1-5.
- [Chang 1987] E. Chang, "Participant Systems," in M. N. Huhns, ed., *Distributed Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA, 1987, pp. 311-339.

[Da
[Dav
[Dur
[Dur
[Dur
[Eric
[Fox
[Gasse
[Gasse
[Gerso
[Hall a:
[Hayes
[Hewitt
[Hewitt
[Hewitt
[Kling ar
Te

- [Davis 1980] R. Davis, "Report on the workshop on distributed artificial intelligence," *SIGART Newsletter*, 1980, p. 73.
- [Davis and Smith 1983] R. Davis and R. G. Smith, "Negotiation as a Metaphor for Distributed Problem Solving," *Artificial Intelligence* Vol. 20, 1983, pp. 63-109.
- [Durfee and Lesser 1987] E. Durfee and V. Lesser, "Using Partial Global Plans to Coordinate Distributed Problem Solvers," *Proceedings IJCAI-87*, Milan, Italy, August 1987, pp. 875-883.
- [Durfee et al. 1987] E. Durfee, V. Lesser, and D. Corkill, "Cooperation through Communication in a Distributed Problem Solving Network," in M. N. Huhns, ed., *Distributed Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA, 1987, pp. 29-58.
- [Durkheim 1938] Emile Durkheim, *The Rules of Sociological Method*, Free Press, New York, 1938.
- [Ericsson and Simon 1979] K. A. Ericsson and H. A. Simon, "Sources of evidence on cognition: An historical overview," C.I.P. Working Paper No. 406, Carnegie Mellon Dept. of Psychology, October 1979.
- [Fox 1981] M. Fox, "An Organizational View of Distributed Systems," *IEEE Transactions on Systems, Man and Cybernetics*, Vol. SMC-11, No. 1, 1981, pp. 70-80.
- [Gasser 1987] L. Gasser, "Distribution and Coordination of Tasks among Intelligent Agents," *First Scandinavian Conference on Artificial Intelligence*, Trumsoe, Norway, March, 1987.
- [Gasser et al. 1986] L. Gasser, C. Braganza, and N. Herman, "MACE: A Flexible Testbed for Distributed AI Research," Technical Report CRI 87-01, Computer Research Institute, University of Southern California, 1986.
- [Gerson 1987] E. M. Gerson, "Audiences and Allies: The Transformation of American Zoology, 1880-1930," *Society for the History, Philosophy and Sociology of Biology*, Blackburg, VA, 1987.
- [Hall and Kibler 1985] R. P. Hall and D. F. Kibler, "Differing Methodological Perspectives in Artificial Intelligence Research," *The AI Magazine*, Fall 1985, pp. 166-178.
- [Hayes 1987] P. J. Hayes, "What the frame problem is and isn't," in Z. W. Pylyshyn, ed., *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Ablex, NY, 1987, pp. 123-137.
- [Hewitt 1986] C. Hewitt, "Offices Are Open Systems," *ACM Trans. on Office Information Systems*, Vol. 4, 1986, pp. 271-287.
- [Hewitt 1988] C. Hewitt, "Organizational Knowledge Processing," presented at *8th AAAI Conference on Distributed Artificial Intelligence*, Lake Arrowhead, California, May 1988.
- [Hewitt and DeJong 1983] C. Hewitt and P. DeJong, "Analyzing the Roles of Descriptions and Actions in Open Systems," *Proceedings of the AAAI*, 1983, pp. 162-167.
- [Kling and Scacchi 1982] R. Kling and W. Scacchi, "The Web of Computing: Computer Technology as Social Organization," *Advances in Computers*, Vol. 21, 1982, pp. 1-90.

- [Kornfeld and C. Hewitt 1981] W. Kornfeld and C. Hewitt, "The Scientific Community Metaphor," *IEEE Transactions on Systems, Man and Cybernetics*, SMC-11, No. 1, 1981, pp. 24-33.
- [Latour 1988] B. Latour, *Science in Action*, Harvard University Press, Cambridge, MA, 1988.
- [Lesser and Corkill 1981] V. Lesser and D. Corkill, "Functionally Accurate, Cooperative Distributed Systems," *IEEE Transactions on Systems, Man and Cybernetics*, SMC-11, No. 1, 1981, pp. 81-96.
- [Pattison *et al.* 1987] H. E. Pattison, D. Corkill and V. Lesser, "Instantiating Descriptions of Organizational Structures," in M. N. Huhns, ed., *Distributed Artificial Intelligence*, Morgan Kaufmann, Los Altos, CA, 1987, pp. 59-96.
- [Pylyshyn 1987] Z. Pylyshyn, ed., *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Ablex, Norwood, NJ, 1987.
- [Star 1989a] S. L. Star, *Regions of the Mind: Brain Research and the Quest for Scientific Certainty*, Stanford University Press, Stanford, CA, 1989.
- [Star 1989b] S. L. Star, "Human beings as material for artificial intelligence: Or, what computer science can't do," presented to *American Philosophical Association*, Berkeley, CA, March 1989.
- [Star 1986] S. L. Star, "Triangulating Clinical and Basic Research: British Localizationists, 1870-1906," *History of Science*, Vol. 24, 1986, pp. 29-48.
- [Star 1985] S. L. Star, "Scientific Work and Uncertainty," *Social Studies of Science*, Vol. 15, 1985, pp. 391-427.
- [Star 1983] S. L. Star, "Simplification in Scientific Work: An Example from Neuroscience Research," *Social Studies of Science*, Vol. 13, 1983, pp. 205-228.
- [Star and Gerson 1987] S. L. Star and E. Gerson, "The Management and Dynamics of Anomalies in Scientific Work," *Sociological Quarterly*, Vol. 28, 1987, pp. 147-169.
- [Star and Griesmer 1989] S. L. Star and J. R. Griesmer, "Institutional Ecology, 'Translations,' and Coherence: Amateurs and Professionals in Berkeley's Museum of Vertebrate Zoology, 1907-1939," *Social Studies of Science*, August 1989.
- [Strauss 1969] A. Strauss, *Mirrors and Masks: The Search for Identity*, The Sociology Press, San Francisco, 1969.
- [Turing 1950] Alan Turing, "Computing Machinery and Intelligence," *Mind*, Vol. 59, 1950, pp. 433-460. Reprinted in E. Feigenbaum and J. Feldman, eds., McGraw-Hill, NY, 1963, pp. 11-35.
- [Wimsatt 1980] W. C. Wimsatt, "Reductionist Research Strategies and Their Biases in the Units of Selection Controversy," in T. Nickles, ed., *Scientific Discoveries: Case Studies*, D. Reidel, Dordrecht, 1980.

Susan Leigh Star
 Department of Information and Computer Science
 University of California
 Irvine, CA 92717

Cha
 Rep
 Org
 Dis

Les G
 Lieb

Abst

This pa
 for distr
 a partic
 about o
 set of q
 differen
 are deve
 implem

3.1

In the
 beyond
 ple wor
 long an
 spite th
 nized.
 of hum