A Logic of Vision

Jaap van der Does and Michiel van Lambalgen[1]

Amsterdam, July, 1997

[1]This article is a completely rewritten version of Van Lambalgen and Van der Does 1996. The research is part of the PIONIER-project 'Reasoning with Uncertainty' sponsored by the Netherlands Organisation for Scientific Research (NWO) under grant PGS 22–262.

1

Introduction

> *Our mistake lies in supposing that things present themselves as they really are, names as they are written, people as photography and psychology give an unalterable picture of them. But in reality this is not at all what we ordinarily perceive. We see, we hear, we conceive the world in a lopsided fashion.*
>
> Marcel Proust, *La Fugitive*

Is there a logic of perception? And if there is, how is it related to the semantics and pragmatics of perception reports? These are the main questions we will address in the present essay. Inspired by Marr 1982, we shall argue that perception is approximative in many ways, and should hence allow for failure and correction. Consequently, the description of perception should also be retractable.

This position has been defended more often, but it still poses many interesting challenges. These are perhaps best exemplified by a passage from Proust's *La Fugitive*.

> I opened the Figaro. What a bore! The main article had the same title as the article which I had sent to the paper and which had not appeared. But not merely the same title... why, here were several words that were absolutely identical. This was really too bad. I must write and complain. But it wasn't merely a few words, it was the whole thing, and there was my signature... It was my article that had appeared at last! But my brain which, even at that period, had begin to show signs of age and to tire easily, continued for a moment longer to reason as though it had not understood that this was my article, like an old man who is obliged to complete a movement that he has begun even if it has become unnecessary, even if an unforeseen obstacle, in the face of which he ought at once to draw back, makes it dangerous. (Penguin Classics ed., p. 579)

When looking at the newspaper, the writer must have seen something; but what kind of object did he perceive? Whatever he saw approximated the real newspaper; yet, how do such approximations figure in the description of what he saw? Had he been less preoccupied, he would have noticed his article right away. Indeed, we normally do not expect gross misperception; so, what is the precise nature of this expectation?

In chapter 3 we try to clarify these issues by means of logico-semantical techniques. The logic proposed is abstracted from Marr's cognitive theory of vision, whence the name: logic of vision. Implicit in Marr's theory is a notion of partiality as lack of structure. For models which are partial in this sense, it is natural to represent reality as a 'regulative ideal'; i.e., a limit to which all understructured approximations converge. On the linguistic side, the approximative nature of perception is incorporated by means of so-called 'filtered' or 'conditional' quantification. This is a resource bounded form of quantification, which generalises the algebraic approach in Halmos 1952.

Chapter 4 shows how conditional quantification provides the semantical core of direct perception reports. On our view, the semantics of these reports should apply under both normal and abnormal circumstances, and may hence fail to be veridical (correct, successful); the truth of a perception report need not imply the truth of its complement. Normality, we hold, is rather a *pragmatic* issue, having to do with our expectation that direct perception reports are stable: truth relative to an approximation is preserved under refinement of information. We show that for 'positive' descriptions it is equivalent to assume veridicality, or to assume that the semantics of a perception report is stable. Consequently, we are able to formalise the pragmatics of direct perception reports by supplementing their semantical core with a defeasible rule stating the pragmatic expectation of stability. The combination of semantics and pragmatics implies veridicality for suitable descriptions. On the logical side we shall argue that the (in)validity of inferences depends on context. It is shown in particular how the logic of perception varies with the resources available for conditional quantification. The chapter ends with comparing partial objects in logic of vision with the pegs developed in Landman 1986.

Before presenting our theory, chapter 2 outlines its *desiderata*: What are the characteristics of perceived objects? Which (non)-inferences should the system be able to analyse? We prepare the ground for logic of vision by showing how simpler alternatives fail. Here we also discuss Grice's view on perception, whose division of labour between semantics and pragmatics is similar to ours (Grice 1961).

The last chapter, chapter 5, views logic of vision in the light of suggestions coming from psychology, linguistics, and philosophy. Firstly, we use neuropsychological studies of agnosia to indicate that the hierarchical model put forward here receives some support from experimental data. Secondly, we indicate how the model enables a study of the semantics of evidentiality; i.e., the linguistic means to indicate different sources of information. Finally, we sketch how logic of vision offers a fairly detailed formalisation of Husserl's theory of perception.

In the long run, our ambition is to understand the workings of natural discourses in terms of conditional quantification on rich models. The idea would be that a text changes the resources of conditional quantifiers as it goes along, and hence determines what concepts and objects are available at a certain locus. Here one could adopt instrumentalism, as once suggested by Lewis: 'In order to say what a meaning *is*, we may first ask what a meaning *does*, and then fid something that does that.' (Lewis 1972, 173). Instead, we think it is crucial to develop models with genuine predictive power based on psychological and philosophical insights. It is our hope that this will lead to a better understanding of our semantic capacities. This essay aims to be a first step in that direction.[1]

Acknowledgements

We are grateful for the comments and suggestions of Renate Bartsch, Johan van Benthem, Joan Bresnan, Martina Faller, Fritz Hamm, Anne Holzapfel, Rosalie Iemhoff, Uwe Mönnich, Dick Oehrle, Sussane Schüle, Keith Stenning, Frank Veltman, Marco Vervoort, and Frans Voorbraak. We also would like to thank the audiences of the DIP-colloquium (ILLC, University of Amsterdam), the Linguistics colloquium (Stanford University), and the Semantics colloquium (Seminar für Sprachwissenschaft, Universität Tübingen).

[1]The sections marked with a star ('*') are not required for an understanding of the body of the paper.

2

Direct perception

This article is concerned with direct visual perception and its description. Direct perception reports can take several forms; e.g., the 'simple' perception in (1), the naked infinitive form in (2), the gerundive form in (3), and the complementised form in (4).

(1) Jack saw Sharon.

(2) Jack saw Sharon wash her face.

(3) Jack saw Sharon washing her face.

(4) Jack saw that Sharon washed her face.

There are subtle semantic differences between the perception reports: sentence (1) describes the perception of objects, (2) and (3) that of scenes, while (4) gives the informational content of what is seen. Also, (1-3) report on what is seen directly, whereas (4) may state a conclusion inferred from what is actually perceived. The difference between (2) and (3) is aspectual; (2) concerns a finished action, (3) an ongoing one.

These differences manifest themselves in the restrictions placed on the main ingredients involved in the interpretation of perception reports; namely, (i) visual information, (ii) the semantic content of a report, and (iii) part of objective reality. Direct perception reports, in particular, require an immediate link between visual information and the semantic content of a report. Since fields of vision normally represent part of objective reality, this link may closely connect reality and semantic content, too; but a connection is not necessary.

From the fact that we consider visual perception, one should not conclude that generalisation to the other modalities is straightforward. For example, although it is possible to describe hearing and feeling by means of a naked infinitive report, as in (5) and (6), it seems impossible to use (7) and (8) for the description of smelling and tasting.

(5) Jack heard his mother call.

(6) Jack felt the earth shake.

(7) *Jack smelled the soup burn.

(8) *Jack tasted the water flow.

Perhaps these observations are connected with Viberg's lexical hierarchy of the five sense modalities (Viberg 1984), which is used to explain the lexicalisation of perception verbs.

(9) $\text{sight} > \text{hearing} > \text{touch} > \begin{cases} \text{smell} \\ \text{taste} \end{cases}$

The idea is that languages may have a perception verb extending to the right, e.g., describing touch, smell, and taste, while having separate means to describe the modalities higher up, e.g., hearing and seeing. The hierarchy suggests that the modalities which are connected to more recently developed parts of the brain take over from the modalities connected to the earlier parts in detecting the processes and events described by NI reports. The 'lower' modalities are then mainly used to detect states. If this is correct, the differences may be clarified by basing semantics for perception reports on the appropriate cognitive theories.

In what follows we first concentrate on the perception of objects, as described by (1). Next, this is generalised to the perception of scenes, as described by naked infinitive reports such as (2).

1 Perceiving objects

One of the most basic questions in developing the logic of direct perception reports is: what kind of objects are presumed in their interpretation and use? Our view on the matter is that they are 'threads'. Threads are a special kind of total object; they are the 'limits' of sequences of partial objects, which occur in various approximations to scenes; i.e., parts of the world in which objects stand in certain relations to each other (Barwise and Perry 1983, 185).

Although we work within a new framework, our approach can be seen as a considerable refinement of the treatment of perception reports in situation semantics. This will become clear in a critical overview of existing proposals, in which we develop our notion of object. Here, we use hallucinations and illusions to argue for a distinction between material and perceived objects. As a consequence, perception could fail to be veridical. Next, the more subtle issue of misperception shows that normally a perceived object approximates a material one (since it lacks most of its structure), and that despite change of approximations we are still able to keep track of one and the same material object. The partial nature of approximations is cashed out as a loss of certain logical principles. Finally, we use the insights obtained to define our notion of an object, and move on to the slightly more complicated case of scenes.

1.1 Simple objects

Realism in its most naive form holds that there is no difference between the objects perceived and material objects. Accordingly a naive realist is justified to hold that visual perception is a relation between two objects: a perceiver and a thing perceived. This corresponds to the meaning of transitive *see*, found in logic primers, as a relation on a domain of objects. For example, (10) means (11), or its mereological variant (12).

(10) Robert-Jan saw Tine

(11) $S(r,t)$

(12) $\exists t' \leq t[S(r,t')]$

If (10) is interpreted as (12), it is true iff Robert-Jan saw a part of Tine.

In our opinion, (11) and (12) are a far cry from capturing all niceties of (10). For one, they remain silent as to when the *see*-relation between objects obtains. This is in general a subtle matter where it is not always clear where the boundaries lie between semantic truth-conditions and pragmatic considerations of normality.

Suppose one knows (i) that the person standing in front of you is Tine, and (ii) that (part of) Tine occurs within Robert-Jan's field of vision. Given this much, does it make sense to report (10)? It depends. Plainly, Robert-Jan should see something for (10) to be true. Its truth even seems to require that Robert-Jan's perceptions approximate something which looks like a female. But there are abnormal circumstances (hallucinations, illusions, misperceptions) where this is less clear. So the question is: where does pragmatics take over from semantics?

Grice's work has made us sensitive to the often vague borderline between semantic and pragmatic inference. We should not decide too quickly to introduce a complicated semantics for certain parts of language, for some inferences might be explained on the basis of a fairly simple semantics supplemented by general pragmatic maxims governing the exchange of information. Grice's pragmatic analyses of the classical connectives are a case in point (cf. Grice 1989, chapter 2–4). Reasoning along these lines, one should consider whether the semantics of (10) might be the unlikely (11). Perhaps its consequence that Tine is part of what Robert-Jan saw is pragmatic, based on this simple semantic core. If so, existential generalisation implies that we should be able to use *any* property of a perceived object to give a true but perhaps pragmatically unacceptable description of it. But, is this so? For instance, looking out of your window, could you report (13), meaning (14)?[1]

(13) I see a globe turning around the sun.

(14) $\exists x[\text{globe-turing-around-the-sun}(x), S(i,x)]$

[1] Here and in what follows we use $\exists$ as a relation between two sets rather than as a property of sets; whence we write '$\exists x[\varphi, \psi]$' instead of: '$\exists x[\varphi \wedge \psi]$'. For this format, existential generalisation becomes: from $\varphi(t)$ and $\psi(t)$ infer $\exists x[\varphi, \psi]$.

Its apparent queerness should then be explained, say, by means of the maxims of quantity and of quality.[2] For example, the more general statement that you see a globe is certainly less informative then describing the hills, birds, and forests you happen to perceive from your window. Nevertheless, we think this pragmatic route to salvage the simple semantics of (13) is blocked: under the circumstances sketched (13) is false rather than true but odd. Apart from its pragmatic shortcomings, there is the strong semantical suggestion that one should at least see a globe rather than just a perhaps uncharacteristic part of it. A similar discrepancy between what is seen and how it is described may occur when seeing Tine (e.g., if you see the landscape of her skin using a macro-lens).

Although these remarks are far from conclusive, we take them as indicating that naive realism as formalised above is too simple. In chapter 4 we present a more complex semantics and a formalised pragmatics to deal with direct perception reports. Yet, the different ways in which the use of simple objects in naive realism fails gives a clear picture of what has to be accomplished by a more promising analysis. For this reason, we show that naive realism does not do justice to hallucination, misperception, and, perhaps most important of all, partial or incomplete perception.

1.1.1 Hallucination, illusion

Visions and hallucinations make particularly clear that a logic of perception needs more than just material objects; it also needs perceived, possibly non-material objects. So, for a short while let us give room to two persons who have described some particularly lively visions.

The Dutch mystic Hadewijch wrote her fourteen visions in the late thirteenth century. She received her seventh vision early one morning in Whitsuntide while singing matins in church. This is what she saw:

> Then He approached me, now in the form of the man that He was when He offered us for the first time His body, handsome and charming, with a countenance of rare beauty, and with the submissive demeanor of one who belongs totally to another. Then He gave himself to me in the customary manner of the sacrement, and then gave me to drink from the goblet: that tasted and seemed as usual. Then He approached close to me, took me completely in His arms, and pressed me to Him. All my limbs felt His to their total satisfaction, as my human heart desired. Likewise, I had the strength to bear this, but all too soon I began to lose sight of the so wonderfully handsome man, and I saw him fading and melting away till I could no longer perceive Him near me, and with me I could not distinguish Him from myself. At that moment I had the feeling that

[2] Recall that the maxim of quantity states: '1. Make you contribution as informative as is required (for the current purposes of exchange); 2. Do not make your contribution more informative than is required.' The second specific maxim falling under quality may also be applicable: 'Do not say that for which you lack adequate evidence.' (Grice 1989, 26–27).

> we were one together, without distinction. All this was real, visible, palpable and tastable...[3]

As religious ecstasy, fear combined with ambition may also have a liberating effect on one's senses. These are the reflections of a Thane of Glamis and Cawdor eager to become King of Scotland:

> Is this a dagger which I see before me,
> The handle toward my hand? Come, let me clutch thee:
> I have thee not, and yet I see thee still.
> Art thou not, fatal vision, sensible
> To feeling as to sight? Or art thou but
> A dagger of the mind, a false creation,
> Proceeding from the heat-oppressed brain?
> I see thee yet, in form as palpable
> As this which now I draw.
> (Shakespeare, *Macbeth*, act 2, scene 1, 32–40)

To account for visions and hallucinations a naive realist is hard-pressed. Some philosophers have even used such phenomena—or more mundane ones like sticks which look bent in water,—to argue that perceived objects are entirely mental. This argument from illusion begins with the observation that we perceive a material object from a certain perspective under certain physico-psychological conditions. (Self-conscious Macbeth furnishes a particularly apt example.) In fact the appearances of an object change with our perspectives and conditions to such an extent that the various properties it has over time are mutually incompatible with each other. (Hadewijch even experienced two objects merging to become one.) This would mean that the properties are not part of the material object itself. But if not material, they must be mental.

Like Ayer 1971, 188-9, and Hintikka 1969, among others, we don't think this reasoning is particularly convincing. The fact that the appearance of an object depends on a perceiver is not sufficient to conclude that perceived objects are entirely mental. According to Marr, the correct answer to the argument from illusion—using sticks in water—should be

> ...that usually our perceptual processing does run correctly (it delivers a true description of what is there), but although evolution has seen to it that our processing allows for many changes (like inconstant illumination), the pertubation due to the refraction of light by water is not one of them. And incidentally, although the example of the bent stick has been discussed since Aristotle, I have seen no philosophical inquiry into the nature of the perceptions of, for instance, the heron, which is a bird that feeds by pecking up fish first seen from above the water surface. For such birds the visual correction might be present.

[3]This translation is taken from the booklet '*De Materie; Louis Andriessen/Robert Wilson*'. Part two of this opera is based on Hadewijch's seventh vision.

However, illusions do justify a more modest position which distinguishes between total material objects and partial perceived objects. Then perception becomes a relation between a perceiver and possibly non-material objects perceived. To argue for this distinction, Anscombe uses an analogy with aiming (Anscombe 1981, 10, 17). A man is aiming at a dark patch in the foliage, but 'that dark patch against the foliage was in fact his father's hat with his father's head in it'. Normally one is allowed to say that the intentional object of his aiming is the dark patch, while the material object is his father's hat (or even just his father). But there is no such relation between the intentional and the material object in case the man is hallucinating. *Mutatis mutandis* the same is true for the perceived and the material object, whenever perception is normal and attentive.

1.1.2 Veridicality

Let us call the logic which results from using domains of quantification with both material and perceived objects 'adjusted realism'. In this logic there are at least two ways to generalise existentially, which correspond, for example, with having formal analogues of the contrast between (15) and (16).

(15)
$$\frac{\text{Hadewijch saw a man}}{\text{Some man-like object was seen by Hadewijch}}$$

(16)
$$\frac{\text{Macbeth saw a dagger}}{\text{Some material dagger was seen by Macbeth}}$$

In the literature, inferences like (16) have been called *veridicality*, *success*, or *correctness*. The contrast between (15) and (16) results from the fact that *a man* and *a dagger* may respectively quantify over perceived and over material objects. Accordingly, to adopt the distinction amounts to assuming, as we will do, that simple perception reports may fail to be veridical.

There is an interesting parallel between adjusted realism and Grice's analysis of perception reports, in that Grice's analysis is also based on a semantic core that may fail to be veridical (Grice 1961). Veridicality is rather ensured by pragmatic principles that govern normal conversation. In particular, he proposes to analyse such statements as (17) by means of propositions of the form (18).

(17) Jack saw Sharon.

(18) x looked P to Jack.

Here (18) is a 'sense-datum statement' describing the properties P of x as seen by Jack. One could object that statements like (18), other than (17), are often used when one doubts or denies that x has property P (Grice calls this *the D-or-D condition*). Although Grice admits that this use is most common, he warns against taking this as a matter of semantics. Instead he suggests the use should

> . . . be explained by reference to a general principle governing the use of language. Exactly what this principle is I am uncertain, but a *first shot* would be the

> following: 'One should not make a weaker statement rather than a stronger one unless there is good reason for so doing'.

This passage seems to contain the first formulation of a maxim of conversation.[4]

In many ways the semantics presented in chapter 3 can be seen as a formalisation of Grice's view (this is surprising since we take Marr 1982 as point of departure). We too will base our semantics on weak, possibly non-veridical statements, which correspond directly with those of form (18). Further, we also claim that veridicality arises from a defeasible pragmatic principle; namely: one expects that what is perceived will remain the case under refinement of perceptual information. But this pragmatic rule formalises an insight from cognitive psychology,[5] not from the theory of conversation.

Adjusted realism is not yet rich enough to cause the failure of principles other than veridicality which are due to the partiality of direct objects. It also raises the question of how material objects and perceived objects are related to each other. Misperception is most suitable to throw light on this matter.

1.1.3 Misperception

Misperception is more interesting than hallucination; not only does it enforce a distinction between material and perceived objects, it also makes one to consider the relation between them. It is especially important to account for the fact that misperceived objects, when corrected, could still be related to one and the same material object. An amusing example is furnished by the following children's story first heard at primary school.

(19) As always, poor Jack was short of money. But today Fortune was at his side. Looking from a window of his parent's penthouse just above the 17th floor, his sharp eyes saw a dime lying on Main Street. He rushed downstairs, checking every now and then whether some lucky bastard would find it before him. A miracle happened... At the 11th floor, it turned out to be a quarter, at the 5th floor even a dollar! How great his disappointment, when out of breath at ground floor he noticed to have chased a trash can.

It is worthwhile to observe that the story is about one coin-like object, despite the incompatible properties ascribed to it in the various stages of misperception.[6] To model this, adjusted realism would need several objects (the material and the perceived ones), but it would fail to clarify how they are related. On our view the

---

[4] The quote is from Grice 1961, section 3, which is not reproduced in Grice 1989, chapter 15. Cf. also the lucid overview of Grice's work in Neale 1992, 515–519.

[5] ... or from empiricism for that matter; cf. David Hume's use of *habits* in his Treatise and Inquiry.

[6] Another example comes from an attempt to save naive realism. A philosopher once held that things look to us, not as they appear but as they are. In fact, during a meeting with some colleagues he even argued they *only* look to us that way. One of the colleagues countered him by bringing in 'a glass vessel of water with a stick in it. "Do you mean to say," he asked, "that this stick does not look bent?" "No," said the other bravely: "It looks like a straight stick in water." So he took it out and it *was* bent.' (Anscombe 1981, 14) Again, talk is of one stick.

objects seen at the various stages are approximations of one and the same material trash can; in general, this is how perceived objects are presented in the logic defined in chapter 3.

Note further that the story is based on a particular use of the verb *to see*, which is quite common in narration. When interpreting perception reports one has to be careful in distinguishing the sources of information used to interpret them. In general, the information available to the speaker will be different from the information available to the perceiver. Sometimes it is clear that these sources are identical; e.g., in case of a first person report such as (13). In case of a third person report, when presented in isolation, one tends to identify the speaker as the source of perceptual evidence on which the report is based. But it is also possible to use a third person report when this source is the subject itself. Without such a device narration quickly becomes cumbersome. It is also used in such dialogues as the following between Pim, Jack, and Sharon.

P: (*to Jack*) 'I saw the sun rising.'

J: (*to Sharon*) 'Pim saw the sun rising.'

It would not do for Jack to weaken his statement to: *Pim thought he saw the sun rising*, or: *Pim saw something that looks like the sun rising*; for these statements voice a doubt which may be unjustified in a normal conversation.[7] There are languages that would not allow the indicated ambiguity, since they have more pronounced linguistic means to indicate the source of evidence. We come back to these more refined systems of evidentiality in chapter 5.

1.1.4 Partial perception

In the adjusted form of naive realism, seeing is identified as a relation between the perceiver and a possibly non-material perceived object. Such perceived objects are still much like the 'points' in the domains of first order models; i.e., objects which are 'accessible' independent of their properties. Although such idealised objects may suffice to determine the classical validities, it goes without saying that they are too idealised if it comes to modelling perception. Our next move should be to partialise the objects perceived.

In life it often happens that we perceive a person (or an object) only partially, but are nonetheless able to reason about this person and to incorporate him or her in our schemes. Such a situation has been eloquently described by Marcel Proust in *La Fugitive*:

> Had I been obliged to draw from memory a portrait of Mlle d'Eporcheville, to furnish a description of her, or even to recognise her in the street, I should have found it impossible. I had glimpsed her in profile, on the move, and she had struck me as being simple, pretty, tall and fair; I could not have said more. But

[7]This seems in conflict with Grice's view on perception statements. But note that Grice introduces the pragmatic principle in his analysis to allow for a less normal use.

all the reflexes of desire, of anxiety, of the mortal blow struck by the fear of not seeing her if my father took me away, all these things, associated with an image of which on the whole I knew nothing, and as to which it was enough that I knew it to be agreeable, already constituted a state of love. (Penguin Classics ed., p. 577-8)

There is a long tradition arguing for the view that perceived objects approximate material objects. In this century it reached a first, somewhat misty peak in the work of Edmund Husserl (cf. chapter 5), but it is also a common assumption in the more analytical tradition.

As approximations, perceived objects are partial, incomplete: they lack most of the structure of the corresponding material objects, and they may not decide for all properties whether or not they apply. Of course, partiality should reflect itself in the logic: classical principles tailored for total objects may fail to hold for partial ones. This logical phenomenon is perhaps first observed by Hintikka, when he discussed incomplete perceptual identification (Hintikka 1969, 164). For example, in classical logic the following reasoning is valid:

$$\frac{\text{Tine is a smiling woman} \quad \text{Robert-Jan saw Tine}}{\text{Robert-Jan saw a smiling woman}}$$

i.e., an instance of existential generalisation:

$$\frac{P(t) \quad S(r,t)}{\exists x[P(x), S(r,x)]}$$

with P any of Tine's properties. But as we have already indicated at the beginning of this chapter, there is a use of indefinite descriptions in perception reports where truth requires the properties described to be actually perceived. When asked: 'Did you see a smiling woman?', Robert-Jan might be justified in answering: 'No. Though I did see Tine, I didn't see her face.' It is on this use that existential generalisation may fail. In fact, due to the partiality of perceived objects, the logic should be rich enough to cause the failure of other inferences as well, such as (20).

(20)
$$\frac{\text{Macbeth saw a dagger}}{\text{Macbeth saw a dagger, and three boys or less than three boys}}$$

Here *three boys or less than three boys* denotes the union of the quantifier *at least three boys* and its complement *less than three boys*, which is hence the *true* quantifier holding of all properties. But as with the premise, there is a use of the conclusion where Macbeth has to see boys if it is to have a truth-value at all (cf. the observations in Barwise 1981 on the failure of classical logic when perceiving scenes).

Similarly, a quantificational analogue of *tertium non datur* would fail; (21) is invalid with 'fifty' meaning: at least fifty.

(21)
$$\frac{}{\text{Carol saw fifty leaves or less than fifty leaves on this spray}}$$

One reason for failure may be that Carol should see leaves in the first place. Anscombe notices another reason why (21) may not to be true:

> . . . there must be some definite number of leaves on a spray that I see, but there need not be some number such that I *see* that number of leaves on the spray. (Anscombe 1981, 3)

All in all, this means that direct perception may involve a relation between a perceiver and an intensional object. But what kind of intensional object could that be?

1.2 Partial objects

Let us recapitulate the characteristics of objects discerned thus far. Hallucinations and illusions were used to introduce a distinction between material and perceived objects; accordingly, perception could be non-veridical. Misperception indicated that under normal circumstances a perceived object approximates material objects, and that it may change its appearance as our perceptual information changes. The partial nature of approximations was cashed out as loss of certain logical principles (partial perception, adding 'empty' information, *tertium non datur*). We shall now sketch a notion of partial object that complies with these observations.

1.2.1 Partiality

The partial nature of perception and its description has many interdependent sources: imprecise visual information, vague conceptualisation, lack of knowledge, feeble reasoning, underdetermined meaning of parts of language. So let us delimit more precisely what kinds of partiality will be studied here.

In the literature partiality as due to *underdetermined* information and meaning has attracted much attention. As we shall show in section 4.4,[8] this notion of partiality also plays a promininent rôle in perception; we may not be able to discern objects from each other, so that the single object we sense might be any one from a collection of objects. For example, a dot seen at the horizon could be a boat, a floating lighthouse, an airplane, a u.f.o.,. . . The properties perceived do not fully determine an object, and objects which share these properties are indiscernible from each other.

In the example partiality is due to *lack of structure*, which is one the ways in which an object can be underdetermined. This kind of partiality is a crucial aspect of perception, and, we think, also of the semantics of perception reports. To continue the example, the dot seen may develop into a boat, which gradually unfolds into an intricate structure of hull, mast, cables and sails. Alternatively, at one of the stages one may have to retract one's identification as a boat in favour of, say, a floating lighthouse. To make this precise, one should describe what it means for a perceived object to lack the structure of a material object.

[8] The numeral '4.4' means: section 4 of chapter 4.

1.2.2 Partial objects

The previous discussion has made clear that a perceived object could change its properties and structure when perceptual information is altered; also, we should be able to keep track of an object despite such changes. Think of the information states in a set T as (at least) partially ordered by $\geq$ (i.e., $\geq$ is reflexive, transitive, anti-symmetric on T); and, relative to an information state, take a perceived object to be an approximation of an object. Then, a material object o can be identified with a function

$$o : T \longrightarrow \bigcup_{t \in T}\{\text{approximate objects at stage t}\}$$

In an information state $t \in T$ an object o is a underdetermined structure o_t of its parts. For example, a car may be appear as a dot, or as a tiny cylinder, or as a slightly larger cylinder with four small discs at its sides, or ... *ad infinitum*. On this view, if state s is at least as informative as state t ($s \geq t$), the approximation c_s of car c in s may have more structure than its approximation c_t in state t. Accordingly, $s \geq t$ should provide a structure preserving map h_{st} from state s into state t such that $h_{st}(c_s) = c_t$: the structure of c_s can be mapped unto that of c_t, but perhaps not conversely. In general, an object o is function as above such that:

$$h_{st}(o_s) = o_t$$

whenever $s \geq t$. This notion of object can be seen as a concrete instance of the abstract notion of a 'perceptually individuated object' in Hintikka 1969, 171–172; but in chapter 3 we shall show in detail how it can be developed in general topological terms. The objects are also related to the pegs in Landman 1986 (cf. section 4.4). For now we leave the discussion of perceived objects to discuss the slightly more complex issues concerning the perception of scenes.

2 Perceiving scenes

At the beginning of this chapter, we have recalled the several ways in which we report on what is seen. Now that we have concluded our overview on simple perception reports, we continue to discuss naked infinitive reports (whence: 'NI reports') like (2); repeated here as (22).

(22) Jack saw Sharon wash her face.

Although the syntactic structure of naked infinitives is still a matter of debate, we take it to be 'NP see NP VP$_{ni}$', with the head verb in VP$_{ni}$ a naked infinitive. The presence of a VP, even if unconjugated, makes an NI complement almost propositional. Yet, there are several restrictions, absent for the complementised form, which indicate that the meaning of the complement is not just any proposition.

2.1 Non-stative complements

The restrictions on NI complements all have to do with the fact that 'stative' complement are often disallowed. This is true of simple complements:[9]

(23) a. *We saw the lamp stand on the table.
b. *We saw Jack be drunk.
c. *We saw Jack own/want a house.
d. *We saw Jack know a linguistic fact.

But also parts of speech that result in stative predicates, such as identity, negation, modals, and conditional forms, are doubtful if not plainly unacceptable.

(24) a. *Jack saw Sharon be (identical with) the woman in red.
b. ?Jack saw Sharon not smoke.
c. *Jack saw Sharon enter if Bill leave.
d. *Jack saw Sharon must leave Bill.
e. *Jack saw Sharon possibly leave.

Like all other theories we have not much more to offer here than to assume an unexplanatory taxonomy of events and propositions, and hold that NI complements denote something like the first. Our main contribution rather lies in providing a resource sensitive logic, which shows among other things how the semantics of different kinds of perception reports could be retractable. So let us now concentrate on the semantic phenomena within our compass.

2.2 Semantics

Ideally, NI reports are used to report on the perception of reality. Indeed, since the seminal Barwise 1981 most semantics formalise the idea that (22) is true iff Jack saw a scene in which Sharon washes her face.[10] A scene is a discernible part of the world in which objects stand in certain relations to each other (Barwise and Perry 1983, 185). Let σ be a partial function which assigns a unique scene to each perceiver, then (22) is true iff (25) is.

(25) $\sigma(j)$ verifies $W(s)$

The semantics in (25) makes simple NI reports highly factual; if (22) is true, there is part of the material world in which Sharon washes her face. Slight variants of the arguments that led to distinguish between material and perceived objects, can be used to argue for a similar distinction at the level of scenes. Interestingly, Barwise uses the distinction in his paper Barwise 1989a, 53, written at about the same time as Barwise and Perry 1983, to explain the workings of constraints. But also in his earlier work he was well aware of the fact that NI reports may fail to be veridical.

[9] The following data are from Akmajian 1977, Gee 1977, Higginbotham 1983, Mittwoch 1990. The observation that negation results in a stative predicate is in Verkuyl 1993 and earlier.

[10] See Barwise 1981 (also in Barwise 1989b), Higginbotham 1983, Kamp 1984, Asher and Bonevac 1987, Landman 1986, Asher and Bonevac 1989, Muskens 1989, van der Does 1991, Hendriks 1993, Koons 1996, among other people; also for a discussion of the logical principles.

At several places it is stressed that he restricts himself to normal usages of the reports, but once their truth conditions are unravelled

> ... we can back up and see what modifications would be necessary to accomodate nonveridical readings. (Barwise 1989b, 12)

In many ways our approach will be dual to the one suggested here. We prefer to start with an understanding of the semantical core of perception reports, which should be applicable in both normal and abnormal situations. Formalised pragmatic insights, added to the semantical kernel, should then predict what to expect under normal circumstances.

Generally speaking, the connection between reality on the one hand, and perception and semantics on the other is rather uncertain. It is this looseness which makes NI reports retractable over time; as in

(26) Hadewijch saw Christ approach to her, but later she realised it was Thomas.

When discussing the logical transparency of NI reports one often assumes an intimate relationship between semantic content and reality. As soon as one takes the uncertainty in this relationship seriously, most of the transparency is lost. To see this, let us revisit some well-known logical principles which have been considered to hold for the 'normal' usage of NI reports.

2.3 Partial perception

One of the most basic non-inferences concerns the interplay between perception of objects and NI reports; (28) does not follow from (27).

(27) Jack saw Sharon, and Sharon winked.

(28) Jack saw Sharon wink.

There is a twofold explanation of this fact: either Sharon's action is not within Jack's visual field, or it is too subtle to be discerned by him. The two possibilities combined identify a range of vision with a coarsened part of reality.

2.4 Veridicality

Veridicality is the principle which can be used to conclude (30) from (29).

(29) Jack saw Sharon was her face.

(30) Sharon washed her face.

Since the NI-complement is interpreted in Jack's field of vision, veridicality depends on what this field is taken to be. If it is part of the model for the conclusion (as in situation semantics), simple NI-complements will be veridical. But if Jack's perceptual field approximates this model (as in the semantics to be developed below), veridicality may fail even here. This would allow for Macbeth hallucinating, as for other NI complements that depend on the viewer's perspective. For example, it does not follow from (31) that Sharon walked left and right of Maria, as veridicality would have it.

(31) Jack saw Sharon walk left of Maria and Jill saw Maria walk left of Sharon.

simpy because 'to the left of' and 'to the right of' only make sense relative to a perspective which could be lacking for the conclusion.

Negative NI complements may block veridicality as well.

(32) Jack saw no girl wink.

(33) No girl winked.

But what about veridicality for numerical statements?

(34) Jack saw three girls wink.

(35) Three girls winked.

In case Jack's visual field approximates the model for the conclusion, our intuitions about numerals become feeble. Then the objects perceived are partial, and how does one count such objects? E.g., when repeatedly looking at the nocturnal sky, did the Babylonians see two stars (the Morning star and the Evening star), or just one (Venus)? Presumably, if the percepts of the stars are *un*stable over time there is no reason to expect veridicality; no matter how the numerals are taken. But what if they *are* stable under such refinement? Section 4 discusses the predictions of the present logic (cf. also Landman 1986, chapter 3).

2.5 Boolean connectives

In the literature there is also a discussion of inferences with the connectives 'and', 'or', and 'not'. We quickly review some of the main findings.

2.5.1 Conjunction.

On the whole, there is consensus concerning the equivalence of (36) and (37).

(36) Jack saw Sharon wink and Mary smile.

(37) Jack saw Sharon wink and Jack saw Mary smile.

It is indeed hard to find counterexamples to the transparency of 'to see' for conjunctive NI complements. Perhaps the following is one, since (39) does not appear to be a consequence of (38).

(38) Jack saw an ant walk nearby and Jack saw a beetle walk at a distance.

(39) Jack saw an ant walk nearby and a beetle walk at a distance.

The invalidity should be due to the impossibility to focus on scenes at different distances at the same time.

2.5.2 Disjunction.

For disjunction the question is whether (40) is equivalent to (41).

(40) Jack saw Sharon smile or stare.

(41) Jack saw Sharon smile or Jack saw Sharon stare.

There is a natural tendency to interpret the disjunction in (41) exclusively, which would block the inference from (40) to (41). But Grice has argued convincingly that this effect is pragmatic, not part of the semantics. For inclusive disjunction (40) and (41) are equivalent; and we will show how the exclusive interpretation may result from pragmatic reasoning.

2.5.3 Negation.

What about negative NI complements? Disregarding the fact that negative NI complements are often odd (cf. section 2.1), one normally takes (42) to imply (43) but not conversely.

(42) Jack saw Sharon not cry.

(43) Jack didn't see Sharon cry.

On the assumption that 'to see' denotes a relation between an object and a factual scene, we only have a non-trivial implication if there are negative facts of some sort. However, in some of us lives a Mr. X, who in a famous discussion on logical atomism with Bertrand Russell doubted the existence of such facts (Russell 1988, 215–16). Mr. X would maintain that the logical form of the premise has no negation within the scope of the perception operator; that is, the negation in the premise has scope over 'to see', so that (42) and (43) are equivalent.

Besides, there are those who read the premise as 'Jack saw Sharon refrain from crying', with 'to refrain from crying' the antonym of 'to cry' (Higginbotham 1983). Formally, this corresponds to introducing disjoint positive and negative extensions of a relation, much as in Feferman 1984. Then we do not have equivalence but we do stay within the realm of positive information, and keep Mr. X satisfied.

Denials in dialogues also ask for a special treatment of negative information in perception reports. Consider

'Did you see that hawk there?'

'I saw *something*, but it was not a hawk.'

It makes perfectly good sense to retract the last sentence by saying 'No, you're right, it *is* a hawk'.

2.6 Quantifiers

Intuitions similar to those concerning veridicality influence our judgments on quantificational behaviour. Apart from scope phenomena *pur sang*, there is the question whether the perceptual field described by an NI report alters the domains of quantification within the scope of 'to see'. In this way someone's visual field determines the extent in which quantifiers may be imported into or exported out of the scope of a perception verb.

It seems we have to following situation. If the perceptual field described by an NI report does not affect the domains of quantification, they may be moved freely into and out of the scope of 'to see'. Then, (44) is equivalent with (45), (46) with (47), and (48) with (49).

(44) Jack saw a girl swim.

(45) A girl is such that Jack saw her swim.

(46) Jack saw no girl swim.

(47) No girl is such that Jack saw her swim.

(48) Jack saw every girl leave his party.

(49) Every girl is such that Jack saw her leave his party.

But if NI reports are not neutral in this sense, all the pairs are independent of each other. Jack may be too confused to perceive a 'real' girl as a girl, and what he perceives to be a girl need not be one. As we shall see in chapter 4, logic of vision predicts more subtle distinctions.

This ends our discussion of the interaction between connectives, quantifiers, and the perception verb. We now give an informal sketch of the semantics developed in the remainder of our paper.

3 Sketch of the formal semantics

Logics for perception reports are often based on the assumption that the perceived objects are 'points' which cannot be refined any further; all partiality comes from their properties. In the previous sections we have seen many cases where it is more natural to assume that the objects themselves are partial. We therefore propose to see the infinitely precise points as arising in the limit of increasingly refined stages. In the formalisation each stage is a first order model. The properties of an object at one stage may fail to hold at more refined stages, although one normally expects no such variation. Despite such instability, it is possible to define an inverse limit, which we take to represent reality.

Reality is but one side of perception's coin, we also need linguistic means to describe it. Since the dynamics of retraction is crucial to the semantics of perception reports, this cannot just be a 'static' logic. Instead, we shall use so-called conditional quantification, which generalises Halmos 1952. This approach is crucially different from the generalisation of quantifiers in Mostowski 1957 and Lindström 1966. Conditional quantification offers a natural way to relativise quantification to varying measures of accuracy, and is hence well-suited for our purposes. For instance, if the resources of conditional quantifiers are rich enough, they can be used to mimick a stage in the above limit construction in the limit itself. Within this framework, perception can be analysed as consisting of an approximative core, which becomes veridical by our expectation that what is perceived will remain stable under refinement of visual information.

David Marr's theory of vision is the heuristic backdrop against which the formal theory is developed (Marr 1982). The next chapter starts with the essentials of his theory.

3

Vision, and a blurred view on logic

We have seen in section 2.2.4 that the principle of veridicality is an idealisation which does not allow for the retraction of perception reports. We believe that retraction is a very real phenomenon, and that any semantics for perception reports should account for this. Moreover, a semantics should also allow for partially perceived objects. This could possibly be achieved by introducing partial objects in the domain, but we favour a principled solution in which, roughly speaking, partial perception is the rule not the exception.

The semantics for perceptual expressions introduced here is characterised by the following features:

1) it is completely model theoretic in nature;
2) it tries to stay close to psychological models of perception;
3) it takes veridicality to be a defeasible principle which allows for the possibility to retract a perception report.

The reader might think there is a certain tension between 1) and 2), since typically the psychological models involve mathematical constructs such as Gaussians, Laplace operators etc., which one would not like to have in one's semantics. Indeed, it is incumbent on us to show that these psychological theories contain a model theoretic 'core' that is relevant to a semantics of perception. We believe that the two central notions here are 'inverse limit' and 'conditional quantification'; whether these indeed capture the semantically significant part of psychological modelling we must leave for the reader to judge. In any case, whatever the fate of this proposal, we are in agreement with Marr when he writes, criticising Gibson's 'realistic' approach

> The underlying point is that visual information processing is actually very complicated, and Gibson was not the only thinker who was misled by the apparent simplicity of the act of seeing. The whole tradition of philosophical inquiry seems not to have taken seriously enough the complexity of the information processing involved (Marr 1982, 30)

We maintain that, in order to explain the logic of perception, some of these complications have to be imported in the model theoretic machinery.

Inevitably, this chapter will be rather technical. We have tried to organise the presentation in such a manner that the main thrust of the argument can be followed also by those not willing to delve into the technicalities. Section 1 gives a rapid introduction to David Marr's theory of vision in so far as it is relevant to our concerns. In section 2 we extract from his work two model theoretic notions, that of a refining inverse system of models, and the inverse limit thereof. A model theoretic correlate of a third notion, that of a filter (in the sense of, say, an UV- or a Dolby-filter, not in the familiar logical sense), is studied in section 3. The three notions are linked in section 4. Next, section 5 revisits object recognition, which leads to a logical study of reliable information; see the preservation theorems in section 6.

1 David Marr on vision

We start by explaining the psychological motivation underlying the model theory. Here, we base ourselves on an abstract account of Marr's theory of vision (1982). Of course, basing one's semantics on an empirical theory brings with it the danger that the empirical theory is wrong; indeed, it has been claimed that Marr's views are 'almost completely wrong' (Mayhew, as quoted in Boden 1988, 74). Nevertheless, we hope to convince the reader that Marr's theory is extremely suggestive from a model theoretic point of view. In particular Marr's idea of a hierarchy of three dimensional models has a good model theoretic correlate; 'good' in the sense that the associated preservation and non-preservation theorems may shed some light on the logic of perception. The model theory is suffciently abstract to be compatible with other approaches based on the idea of a hierarchy of perceptual models, such as P.K. Allen's (1987). In chapter 5 we shall even suggest that the model provides a general semantics of information processing.

Marr's fundamental idea is that vision is in many ways approximate. Filtering takes place at many of the earlier levels of visual processing, leading up to the so called primal sketch; and, at the other end, the perception of 3-D objects and scenes takes place by means of a hierarchy of ever more refined, but never completely accurate models. Here we shall concentrate on the last stage of the visual process.

Seeing a 3-D object involves two processes: constructing an image from visual data, and matching the image to a catalogue of 3-D models, where the matching is based on some salient features derived from the image. At this point we can do no better than reproduce the following illustration from Marr 1982, 306.

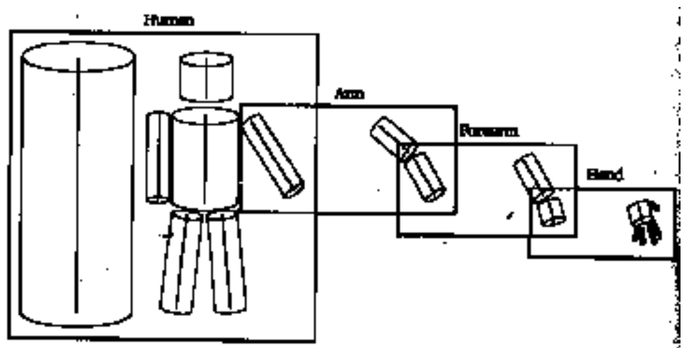

Refinement of an arm

What is depicted is an increasingly detailed series of models of an arm. Obviously, this series can be extended further by detailing the shape of the fingers, by replacing the cylinders by less rigid shapes, by decomposing other parts of the body, etc. Marr's point is that we recognise an object in the real world in terms of these 3-D models, and that we may often use a rather rough approximation to correctly identify the object.

> Modularity [...] allows the representation to be used more flexibly in response to the needs of the moment. For example, it is easy to construct a 3-D model description of just the arm of a human shape that could later be included in a new 3-D model description of the whole human shape. Conversely, a rough but usable description of the human shape need not include an elaborate arm description. Finally, this form of modular organisation allows one to trade off scope against detail. This simplifies the computational processes that derive and use the representation, because even though a complete 3-D model may be very elaborate, only one 3-D model has to be dealt with at any time, and individual 3-D models have a limited and manageable complexity (Marr 1982, 307).

There exists an interplay between the clues derived from an image and the matching process (cf. Marr 1982, 321): after a 3-D model has been selected (guided by the image), it can be used to search for additional clues in the image; in turn, these can be used (when necessary) to match the image to a more detailed 3-D model (cf. section 5). However, it may turn out to be impossible to find a more detailed 3-D model *of the kind we expected*. Indeed, like all computationally efficient heuristics, the use of such approximate models brings with it the possibility of error: what is identified as a real arm with respect to a given approximation may turn out to be something else (e.g., a wooden arm) when 'looking closer', i.e. with respect to a more refined approximation. (This point is not much emphasised in Marr 1982 though.) In any case a theory such as Marr's is well-suited to account for partial perception of an object: this is simply the matching of an object to a 3-D model without an expectation as to the direction in which the model can be refined. These observations suggest a formal semantics for visual reports in terms of approximate models and a stability condition. For instance, informally still, the expression:

(1) I see an arm

can be taken to mean the conjunction of (i) and (ii).

i) with the present approximation the object that I focus on is identified as an arm;
ii) I expect this to be the case for every more refined approximation.

That is, the arm reported on in (1) is viewed as a (possibly infinite) series of ever more accurate approximations; recognizing something as an arm means finding a matching 3-D model somewhere in this series. The stability condition says that we could also find less approximate models in this series, if we would care to submit the image to more elaborate processing. By contrast, if we say

(2) What I see looks like an arm

we imply only condition (i), not (ii).

2 Inverse limits

Consider again Marr's suggestive example of 3-D models of an arm. Viewed abstractly, what we see is a series of first order models, composed of objects and relations between them, together with a mapping specifying how an object occurring at one level is decomposed at the next. This situation can be represented abstractly by means of an *refining system* of first order models. The basic ingredient is the following:

Definition 1 Let T be a set directed by a partial order $\geq$; i.e., $\geq$ is reflexive, transitive, anti-symmetric, and for $s, t \in T$ there is $r \in T$ such that $r \geq s, t$. A *refining inverse system* (indexed by T) is a complex $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ with

i) for each $s \in T$, $\mathcal{M}_s$ is a model for signature $\mathcal{S}_s$;
ii) for any R in the union of the signatures there is $t \in T$ such that R is in $\mathcal{S}_s$ if $s \geq t$.
iii) for each $s, t \in T$ with $s \geq t$ there is $h_{st} : |\mathcal{M}_s| \longrightarrow |\mathcal{M}_t|$ with for each R in $\mathcal{S}_t$, and hence in $\mathcal{S}_s$

   (ρ) $\{\langle h_{st}(d_1), \ldots, h_{st}(d_n) \rangle : \langle d_1, \ldots, d_n \rangle \in [\![R]\!]_s\} \subseteq [\![R]\!]_t$;
iv) $h_{rr} = id_r$, and for $s \geq t \geq r$, $h_{sr} = h_{st} \circ h_{tr}$.

Here $|\mathcal{M}|$ denotes the domain of $\mathcal{M}$, and $[\![R]\!]_i$ the interpretation of R in $\mathcal{M}_i$. If no confusion is likely, we omit $[\![-]\!]$ and write, say, R_s instead of $[\![R]\!]_s$. The mappings h_{st} will be called *bonding* mappings. A refining inverse system is *total* if in addition the bonding mappings are surjective; we then have the usual concept of homomorphism between models. In the following, we shall often use 'refining system' instead of: refining inverse system.

Condition (ρ) is the reason for using the adjective 'refining'. Further, it is mainly for conceptual reasons that we allow the signatures of the models to vary; we may wish to say that a predicate is not yet applicable at a certain stage. For instance, in a rough approximation of 'human being' given by six appropriately positioned cylinders, representing torso, arms, legs and head, the predicate 'hand' is not applicable. Yet, once introduced a predicate should remain applicable. The above definitions can be simplified somewhat by working with a fixed signature for all models, and assuming that intuitively non-applicable predicates trivially apply to all objects in an approximation. But for present purposes we think this is an artifact which is better circumvented.

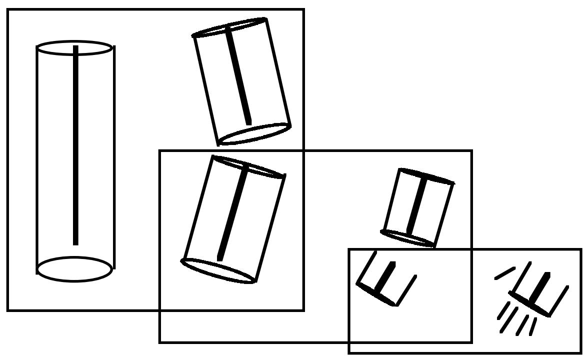

Arm – Forearm – Hand

We now show that the set of 3-D models (in the sense of Marr) can be given the structure of a refining system of models (in our sense): the condition of directedness then says that two refinements will themselves have a common refinement, as in the picture 'Arm-Forearm-Hand' after Marr 1982, 306.

Indeed, Marr's hierarchy of 3-D models can be obtained if we restrict ourselves to refining systems of *finite* models of a special kind. In the example, let $\mathcal{M}_t$ be a model $\langle D_t; A, \ldots \rangle$ with A a unary predicate for 'arm'. An arm $a \in A$ may become a composite $\{e, f\}$ in a more refined version $\mathcal{M}_s$ of $\mathcal{M}_t$ $(s \geq t)$, with e an upperarm joined to a forearm f. On this view, unanalysed objects and composite objects have different types, but this difference can be eliminated by identifying an unanalysed object d with the singleton $\{d\}$.

Definition 2 A *Marr model* is a finite first order model with domain $\wp^+(E) := \wp(E) - \{\emptyset\}$ and a 'part-of'-relation $\subseteq$, among other relations.

Inverse systems of Marr models allow for bondings between complexes which preserve $\subseteq$.

Definition 3 An *refining inverse system of Marr models* is a refining inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ such that

i) each $\mathcal{M}_s$ is a Marr model;

ii) each h_{st} is a bonding mapping with

a) $|X| \geq |h_{st}(X)|$;
b) $h_{st}(Y) \subseteq h_{st}(X)$ if $Y \subseteq X$ (*monotonicity*).

Condition (iia) ensures that the bonding mappings never assign a more refined object to a coarser one, while condition (iib) says that parts of more refined complexes are mapped onto parts of the corresponding coarser object. Write d° for the singleton $\{d\}$. To continue the example, $\mathcal{M}_s$ could be a model $\langle D; A, U, F, J, \ldots\rangle$, in which $\{e, j\}$ is an arm ($\{e, j\} \in A_s$) with e° an upperarm ($e^\circ \in U$) and f° a forearm ($f^\circ \in F$) joined to each other $\langle e^\circ, f^\circ\rangle \in J$. Clearly $\mathcal{M}_s$ represents the same situation in more detail than $\mathcal{M}_t$ via the bonding mapping h_{st} with $h_{st}(e^\circ) = h_{st}(j^\circ) = h_{st}(\{e, j\}) = a^\circ \in A_t$. But creating structure as in $h_{st}(e^\circ) = \{a, b\} \in D_t$ is disallowed.

A model in a refining system approximates reality, but reality itself is assumed to be infinitely precise. This means, e.g., that for a sequence of models

$$\ldots \mathcal{M}_{n+1} \xrightarrow{h_n} \mathcal{M}_n \longrightarrow \ldots \mathcal{M}_1 \xrightarrow{h_0} \mathcal{M}_0$$

with h_n the bonding mapping $h_{n+1,n}$, reality arises as an inverse limit $\mathcal{M}$ refining all models in this sequence

$$\mathcal{M} \xrightarrow{\pi_i} \mathcal{M}_i$$

via projections $\pi_i : |\mathcal{M}| \longrightarrow |\mathcal{M}_i|$. Topological results can be used to show that such inverse limits exist for *any* refining inverse systems of finite models (not just the linearly ordered ones).

Definition 4 Let $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ be a refining inverse system. Its *inverse limit*

$$\mathcal{M} := \varprojlim{}_{t\in T} \mathcal{M}_t$$

is defined as follows

i) the domain $|\mathcal{M}|$ consist of the *threads* in the product $\Pi_{t\in T}|\mathcal{M}_t|$; i.e., functions $\xi : T \longrightarrow \bigcup_{t\in T}|\mathcal{M}_t|$ satisfying: $\xi_t \in |\mathcal{M}_t|$, and $h_{st}(\xi_s) = \xi_t$ for $s \geq t$.
ii) the interpretation function $[\![-]\!]$ is given by: for each R there exists $t \in T$ such that

$$[\![R]\!](\xi^1, \ldots, \xi^n) \Leftrightarrow_{\text{def}} \forall s \geq t : [\![R]\!]_s(\xi^1_s, \ldots, \xi^n_s)$$

Again, when no confusion is likely, we omit $[\![-]\!]$.

The inverse limit $\mathcal{M}$ is a submodel of the direct product $\Pi_{t\in T}\mathcal{M}_t$ (Chang and Keisler 1990, 224); however, the domain of this submodel might be empty. Under the additional assumption that the $\mathcal{M}_s$ are finite this cannot be so.

Theorem 1 *Suppose $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ is a refining inverse system of finite models. Then $|\mathcal{M}|$ is non-empty.* 2

The proof uses the fact that the $|\mathcal{M}_s|$ are compact Hausdorff in the discrete topology, and that the bonding mappings are continuous in this setting (cf. Engelking 1989, 141). It can also be shown that the inverse limit comes with projections

$$\pi_t : \mathcal{M} \longrightarrow \mathcal{M}_t$$

for $t \in T$ defined by $\xi \mapsto \xi_t$. The projections satisfy

$$h_{st}(\pi_s(\xi)) = \pi_t(\xi)$$

for $s \geq t$ in T. If the bonding mappings h_{st} are surjective, π_s is surjective as well, so that $\mathcal{M}$ is a subdirect product of the $\mathcal{M}_s$ (Engelking 1989, 142).

Reality is taken to be an inverse limit of a refining inverse system. Inherent in this approach is a certain claim concerning the relationship between partial objects (the elements in an $|\mathcal{M}_s|$) and the material objects (the threads in the inverse limit). In particular, if the inverse limit is obtained from a total refining system (all bonding mappings are surjective), then each partial object is real in the sense that it is an approximation of a real object. But if a partial object may disappear under closer scrutiny, such as Macbeth's dagger, the inverse limit should be constructed from a refining system *simpliciter*.[1]

3 Conditional quantification

In the previous section we have seen how object recognition can be described by means of a hierarchy, formally a refining system, of 3-D models. Reality was viewed as the inverse limit of this hierarchy. In this section we study the inverse process: we start from the assumption that reality can be perceived only to varying degrees of approximation; perception is seen as the application of a suitable 'filter' to reality. The word 'filter' here should not be taken in its usual logical meaning (as a set of sets closed under intersection and supersets); it derives rather from physical analogues such as the Gaussian filters of Marr 1982, 54 *passim*. Their function is to blot out details which occur at some specified scale (hence at smaller scales). When applied to a picture, this type of filter introduces a blurring of the picture. Put differently, the effect is that pictures which differ only at scales smaller than the specified level, are perceived as identical. Hence, informally at least, there is a connection between filters and equivalence relations.

Logically speaking, a filter in the sense intended here is a new kind of generalised quantifier, which applied to a formula of n free variables, in general yields a new formula, also in n free variables. Hence this notion of generalised quantification differs from the more traditional Mostowski - Lindström generalised quantification, which does bind variables. It will be seen however, that the new notion of quantification is generalised in the sense that ordinary $\exists$ is a special case. Furthermore, the process of filtering is truly the inverse of taking an inverse limit: it can be shown that a suitable collection of filters applied to model $\mathcal{M}$ yields an refining system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ such that $\mathcal{M}$ is the inverse limit of $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ (for a more accurate formulation, cf. theorem 4 below).

To explain filtered quantification, let us return to the refining systems and inverse limits of the last section. An object $a \in \mathcal{M}_s$ can be identified with the collection of threads ξ such that $\xi_s = a$. Similarly, an assignment $g \in \mathcal{F}_s$ (the set of assignments

[1] For a similar construction in the logic of time, see Thomason 1989.

on $|\mathcal{M}_s|$) can be identified with the set of assignments $f \in \mathcal{F}$ (the set of assignments on $|\mathcal{M}|$) such that $f_s = g$; i.e., the assignments f such that for all variables y, $f(y)_s = g(y)$. In other words, an assignment $g \in \mathcal{F}_s$ can be seen as an equivalence class of assignments $f \in \mathcal{F}$ via an equivalence relation E_s which is defined as

$$E_s(f,h) \text{ iff for all variables } y : f(y)_s = h(y)_s$$

An equivalence relation E_s determines a generalised existential quantifier $\exists_s$ as follows

$$M, f \models \exists_s \varphi \Leftrightarrow_{\text{def}} \exists h \in \mathcal{F}[E_s(f,h) \;\&\; \mathcal{M}, h \models \varphi]$$

In the remainder of this section we shall study the quantifiers $\exists_s$ in greater detail, but for now we note the following. By the definition of the inverse limit, a predicate A applies to a thread ξ if for all 'scales' s, $A_s(\xi_s)$, where A_s denotes the interpretation of A on $\mathcal{M}_s$. The effect of $\exists_s$ applied to A is that we obtain a coarser predicate (coarser, since we have $A \subseteq \exists_s A$) defined on the inverse limit, where details at scales $t > s$ have been blotted out.

Note that ordinary $\exists x$ is also determined by an equivalence relation, namely

$$f =_x g \text{ iff for all } y \text{ different from } x\text{: } f(y) = g(y)$$

hence quantifiers determined by equivalence relations are really 'generalised'. Another example is furnished again by inverse limits: instead of E_s we may choose a coarser equivalence relation R_s defined by

$$R_s(f,g) \text{ iff for all formulas } \varphi\, [\mathcal{M}_s, f_s \models \varphi \Leftrightarrow \mathcal{M}_s, g_s \models \varphi]$$

In general, R_s is coarser than E_s, since elements of $\mathcal{M}_s$ need not be distinguishable by formulas.

Before we present a number of further examples, we give an alternative construction of filtered quantifiers, which is closer to proof theory. Let $\mathcal{F}_s$ be the set of assignments for $|\mathcal{M}_s|$, and let $\mathcal{B}_s$ be the Boolean algebra over the subsets $\{f \in \mathcal{F}_s : \mathcal{M}_s, f \models \varphi\}$ of $\mathcal{F}_s$. The reader may verify that there exists the following relationship between R_s and $\mathcal{B}_s$:

$$R_s(f,g) \Leftrightarrow \forall C \in \mathcal{B}_s[f \in C \Leftrightarrow g \in C]$$

Let us write $\exists(\bullet|\mathcal{B}_s)$ for the quantifier determined by R_s, i.e. the quantifier which satisfies

$$f \in \exists(\varphi|\mathcal{B}_s) \Leftrightarrow_{\text{def}} \exists g[R_s(f,g) \;\&\; g \in \varphi]$$

The reason for this notation is that it will be advantageous for our purposes to take as fundamental the notion 'quantification with respect to a subalgebra' (or, more generally, with respect to a sublattice). Since this notion has strong similarities with the concept of conditional expectation from probability theory, we adopt the same notation, and dub this form of quantification 'conditional quantification'.

Observe that $\exists(\bullet|\mathcal{B}_s)$ satisfies the property, known as a Galois correspondence:

$$\forall \psi \in \mathcal{B}_s[\varphi \subseteq \psi \Leftrightarrow \exists(\varphi|\mathcal{B}_s) \subseteq \psi]$$

Indeed, the direction from right to left follows from the reflexivity of R_s, which entails $\varphi \subseteq \exists(\varphi|\mathcal{B}_s)$. For the converse, choose $f \in \exists(\varphi|\mathcal{B}_s)$; by definition this means that there exists g such that $R_s(g, f)$ and $g \in \varphi$. By hypothesis, $g \in \psi$, whence by the connection between R_s and $\mathcal{B}_s$, $f \in \psi$.

Definition 5 (*preliminary version*) Let $\mathcal{M}$ be a model, $\mathcal{F}$ the set of assignments on $\mathcal{M}$, $\mathcal{G}$ an algebra of subsets of $\mathcal{F}$. By a *quantifier conditional on* $\mathcal{G}$—notation: $\exists(\bullet|\mathcal{G})$—we mean a mapping which applied to a set $[\![\varphi]\!]_{\mathcal{M}} := \{f \in \mathcal{F} : \mathcal{M}, f \models \varphi\}$ yields an element of $\mathcal{G}$ such that

$$(*)\ \ \forall C \in \mathcal{G}[[\![\varphi]\!]_{\mathcal{M}} \subseteq C \Leftrightarrow \exists([\![\varphi]\!]_{\mathcal{M}}|\mathcal{G}) \subseteq C]$$

In the following we often write φ for $[\![\varphi]\!]_{\mathcal{M}}$, when the model is clear from context.

Roughly speaking, $\exists(\varphi|\mathcal{G})$ represents the best estimate of φ on the basis of the information available in $\mathcal{G}$; hence we require $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$. Note that a quantifier with these properties, when it exists, is unique. Indeed, the Galois correspondence (*) implies that $\exists(\varphi|\mathcal{G})$ must be defined as $\bigwedge\{C \in \mathcal{G}|\varphi \subseteq C\}$.

Definition 5 is not yet quite what we want, because there may be $\mathcal{G}$ and φ for which $\exists(\varphi|\mathcal{G}) \notin \mathcal{G}$; but it suffices for the following examples. Below we shall show how to modify the algebra so as to retain the Galois correspondence for conditional quantifiers.

Example 1: *Generalised quantification.* For the conditioning algebra, we take the algebra $\mathcal{G}_x$ consisting of sets $\{f \in \mathcal{F} : \mathcal{M}, f \models \varphi\}$ where x does not occur free in φ. Then $\exists(\varphi|\mathcal{G}_x) \subseteq \exists x\varphi$ from the Galois correspondence for $\exists(\bullet|\mathcal{G}_x)$. The rules for $\exists x$ yield the converse inclusion, since in this case $\exists(\varphi|\mathcal{G}_x) \in \mathcal{G}_x$! Hence $\exists(\bullet|\mathcal{G})$ is truly a generalised quantifier.

Example 2: *Blurring of individuals. In nuce*, the following example describes our approach to a semantics of perception. A statement like 'y sees φ' is rendered formally as: '$\exists x[\varphi(x) \wedge S(y,x)]$', where φ defines a unique x. Symbol S gives the denotation of the transitive verb 'to see', and it delimits the set of objects $\{d : Sad\}$ seen by a. The quantifier $\exists x$ ranges over completely accurate objects; in formal terms, the elements of the inverse limits constructed in the previous section. To accomodate actual perception, which always has finite precision, we replace the quantifier $\exists x$ by a filtered quantifier $\exists(\bullet|\mathcal{G})$, where $\mathcal{G}$ represents the degree of blurring. Here we present only a simple case.

Let $\mathcal{M} = \langle\{1,2,3,4,5\}, S, W, U\rangle$, with '$S(a,b)$' for: a sees b, 'W' for: West, and 'U' for: up. In particular, $S = \{\langle 2,1\rangle, \langle 2,3\rangle, \langle 4,5\rangle\}$, $W = \{1,5\}$, $U = \{1,2,3\}$.

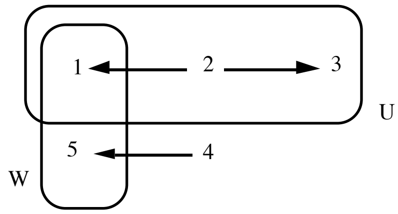

Up west

Put $\varphi(x,y) = S(y,x) \wedge W(x) \wedge U(y)$, then for f such that $f(y) = 2$:

$$\mathcal{M}, f \models \exists! x \varphi(x,y)$$

('y occupies the 'up' position and, y sees exactly one object looking to the West'); in fact:

$$\begin{aligned} &\{f : \mathcal{M}, f \models \exists x \varphi(x,y)\} \\ &\quad = \quad \{f : \mathcal{M}, f \models \exists(\varphi(x,y)|\mathcal{G}_x)\} \\ &\quad = \quad \{f : f(y) = 2\} \end{aligned}$$

The algebra $\mathcal{G}_x$ represents the case that the viewer y has complete information, both about its own position ('up') and the direction in which it is looking ('west'). Let us now vary this situation, for instance by depriving the viewer of the information that he is in the 'up' position. The (finite) algebra $\mathcal{G}$ on the set of assignments, corresponding to this state of affairs, is determined by the formula algebra generated by the set $\{S(y,x), W(x)\}$. We now have:

$$\begin{aligned} &\{f : \mathcal{M}, f \models \exists(\varphi(x,y)|\mathcal{G})\} \\ &\quad = \quad \bigwedge\{\psi \in \mathcal{G} : \varphi \leq \psi\} \\ &\quad = \quad \{f : \mathcal{M}, f \models S(y,x) \wedge W(x)\} \\ &\quad = \quad \{f : (f(y) = 2 \wedge f(x) = 1) \vee (f(y) = 4 \wedge f(x) = 5)\} \end{aligned}$$

Similarly, if y is so disoriented that he does not know whether he is looking East or West, we may describe his predicament by the algebra $\mathcal{H}$ generated by $\{S(y,x), U(y)\}$. In this case:

$$\begin{aligned} &\{f : \mathcal{M}, f \models \exists(\varphi(x,y)|\mathcal{H})\} \\ &\quad = \quad \bigwedge\{\psi \in \mathcal{H} : \varphi \leq \psi\} \\ &\quad = \quad \{f : \mathcal{M}, f \models S(y,x) \wedge U(x)\} \\ &\quad = \quad \{f : f(y) = 2 \vee (f(x) = 1 \wedge f(x) = 3)\} \end{aligned}$$

We trust the reader can subject y to a still more savage experiment.

In order to highlight some formal features of this example, a technical point: we adopt the following convention concerning *free variables*:

$$\mathrm{FV}(\exists(\varphi|\mathcal{G})) = \mathrm{FV}(\varphi) \cap \bigcup\{\mathrm{FV}(\psi)|\psi \in \mathcal{G}\}.$$

Notice that the number of free variables need not be reduced by conditional quantification! In the formula $\exists(\varphi(x,y)|\mathcal{G})$ x is neither fully bound nor fully free; whereas ordinary existential quantification $\exists x$ has the effect of abolishing all restrictions on x, a filtered quantifier need only liberalise restrictions on x to some extent, without fully abolishing them.

Conditioning algebras can be seen as representing information available about the possible values of variables. In the case of $\mathcal{G}_x$, we have no information about x, but complete information about variables other than x. We may now replace the algebra $\mathcal{G}_x$ by an arbitrary algebra, for instance a proper subalgebra $\mathcal{G}$ of $\mathcal{G}_x$. The meaning of this move is that we are prepared to consider situations in which we have less than complete information about variables other than x. This may seem strange on the usual view of variables, but it becomes less strange when variables are taken as observable quantities, which are measured by means of some device which has finite accuracy. If the variable is real-valued, this could mean, for instance, that a measurement only constrains the variable to lie in an interval. To put this in currently fashionable terms: we wish to introduce a resource-bounded form of quantification, where the resource is the degree of accuracy with which one can 'measure' variables. As indicated above, ordinary predicate logic has unbounded resources in this respect.

Notice that in the above example we also have: $\exists(\varphi(x,y)|\mathcal{G}) \in \mathcal{G}$, as required by the preliminary definition. The next example shows that there are particular conditioning algebras for which it may be impossible to satisfy this condition.

Example 3: *Truck through door?* We now replace the dichotomy up-down by a continuum of possibilities. A truck has to pass through a narrow door in a wall; a person behind the truck checks whether this is possible. In this case the viewer must choose his position rather accurately: too far left or right means that he will see only one side of the truck, if he comes too close the truck will occlude the door, but if he goes too far down he cannot accurately estimate the distance between truck and doorposts. See the picture 'Truck through door'.

Let $O(x)$ be the predicate 'x is a (sufficiently large) opening', $S(y,x)$ the relation 'y sees x' (where the viewer y is identified by the coordinates of his position in the plane) and let $A_n(y)$ be a countable collection of predicates denoting sets in the plane. We assume the following logical relations between these predicates:

1) $\forall n \forall x \forall y(O(x) \wedge S(y,x) \rightarrow A_n(y))$
2) $\forall n \forall y(A_{n+1}(y) \rightarrow A_n(y))$
3) $\forall n \exists y(\neg A_{n+1}(y) \wedge A_n(y))$
4) $\forall n \exists y(A_n(y) \wedge \forall x(O(x) \rightarrow \neg S(y,x)))$.

Condition 4) expresses that it is hard to find the exact position from which an opening can be accurately observed; each set A_n contains positions from which no opening is visible. This condition is not used in the proof to follow.

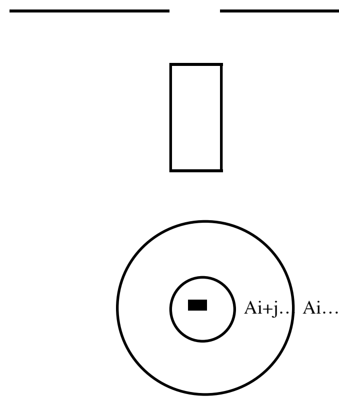

Truck through door

Let $\mathcal{M}$ be a model for (1)-(3), and suppose $\mathcal{G}$ is the Boolean algebra of subsets of $\mathcal{F}$, the set of assignments for $\mathcal{M}$, generated by

$$\{A_n(y) : n \in \omega\}$$

This algebra represents the situation that the viewer has no precise information about his location y; the only available information is in the form of the open sets $A_n(y)$. This has a twofold consequence:

a) $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \notin \mathcal{G}$,

b) $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ is not first order definable.

For (a) and (b), we observe that on $\mathcal{M}$,

$$\exists(O(x) \wedge S(y,x)|\mathcal{G}) \equiv \bigwedge_{n\in\omega} A_n(y)$$

To see this, first note that for all $\varphi \in \mathcal{G}$

(*) $\quad O(x) \wedge S(y,x) \subseteq \varphi \Leftrightarrow$ *there is* $k \in \omega[\varphi \equiv A_k(y)]$

From right to left this is immediate from (1). Conversely, assume $O(x) \wedge S(y,x) \subseteq \varphi$ for $\varphi \in \mathcal{G}$. Since φ is a Boolean combination of the A_n, it can be written in disjunctive normal form. By (1) conjunctions with negative occurrences of the A_n cancel out, so it follows from (2) that ψ is equal to some $A_k(y)$.

It is immediate from (*) and (1) that $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \equiv_{(*)} \bigwedge\{A_n(y) : O(x) \wedge S(y,x) \subseteq A_n(y)\} \equiv_{(1)} \bigwedge_{n\in\omega} A_n(y)$.

As to a), assume for a contradiction $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \in \mathcal{G}$. Then also $\bigwedge_{n\in\omega} A_n(y) \in \mathcal{G}$. Since $O(x) \wedge S(y,x) \subseteq \bigwedge_{n\in\omega} A_n(y)$, (*) implies $\bigwedge_{n\in\omega} A_n(y) \equiv A_k(y)$ for some k; contradicting (3).

As to b), assume $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ is first order definable, say by $\chi(x,y)$. Let $\mathcal{M}$ be a ω-saturated model for (1)–(3).[2] On $\mathcal{M}$, $\exists(O(x) \wedge S(y,x)|\mathcal{G}) \equiv \bigwedge_{n\in\omega} A_n(y)$,

[2]Recall that $\mathcal{M}$ is *ω-saturated* iff for each each type $\Gamma(x_1,\ldots,x_n)$ in the language for $\mathcal{M}$ with finitely many constants added: if $\mathrm{Th}(\mathcal{M}) \cup \Gamma(x_1,\ldots,x_n)$ is consistent, then $\Gamma(x_1,\ldots,x_n)$ is realisable in $\mathcal{M}$. The notion of *ω_1-saturation* is defined similary with 'at most countably many constants added' instead of: finitely many constants added.

as before. So the theory Γ given by

$$\mathrm{Th}(\mathcal{M}) \cup \{A_n(y) : n \in \omega\} \cup \{\neg\chi(x,y)\}$$

must be consistent; for otherwise

$$\mathrm{Th}(\mathcal{M}) \cup \{A_n(y) : n \in \omega\} \models \chi(x,y)$$

which by compactness of $\models$ and the previous observations implies

$$\mathrm{Th}(\mathcal{M}) \models A_k(y) \leftrightarrow \bigwedge_{n\in\omega} A_n(y)$$

for some $k \in \omega$. This is a contradiction, since $(1), (2), (3) \in \mathrm{Th}(\mathcal{M})$. Due to ω-saturation $\{A_n(y) : n \in \omega\} \cup \{\neg\chi(x,y)\}$ is realisable in $\mathcal{M}$: there are $c, d \in |\mathcal{M}|$ such that $\mathcal{M} \models \bigwedge_{n\in\omega} A_n(c) \wedge \neg\chi(c,d)$. This conflicts with $\chi(c,d) \equiv \bigwedge_{n\in\omega} A_n(c)$. We conclude that $\exists(O(x) \wedge S(y,x)|\mathcal{G})$ is not first order; in particular it is not equivalent to $\exists x(O(x) \wedge S(y,x))$.

We have seen that (a) and (b) are true. Clearly the trouble is that $\exists(\bullet|\mathcal{G})$ need not be in $\mathcal{G}$ because the required infima do not exist in $\mathcal{G}$. For the final definition of conditional quantification we therefore move to a slightly different structure.

3.0.1 Conditional quantifiers on frames

Let $\mathcal{N}$ be an arbitrary first order model, with set of assignments $\mathcal{F}$. We put a topology on $\mathcal{F}$ by specifying as basis for the topology sets of the form $\{f \in \mathcal{F} : \mathcal{N}, f \models \varphi\}$. These basic open sets are closed, too, since $\{f \in \mathcal{F} : \mathcal{N}, f \models \neg\varphi\}$ is also in the basis. Clearly the collection of closed sets is closed under finite unions and arbitrary intersections. This is the prime example of a $\vee, \bigwedge$-frame:

Definition 6 A lattice L is a $\vee, \bigwedge$-*frame* if it is closed under arbitrary meets, such that the following distributive law holds:

$$a \vee \bigwedge_{i\in I} b_i = \bigwedge_{i\in I}(a \vee b_i)$$

A *subframe* L' of L is a sublattice of L which is a frame, such that the meets in L', computed in L, coincide with those of L'.

The customary notion of frame, the complete Heyting algebra or $\wedge, \bigvee$-frame, is dual to our notion. However, since we shall be concerned with $\vee, \bigwedge$-frames only, we refer to these structures simply as frames. We aim to define quantification with respect to a subframe of a given frame.

Definition 7 Let **Form** be the frame generated by the first order definable sets of assignments, and let $\mathcal{G}$ be a subframe of **Form**. $\exists(\bullet|\mathcal{G})$, the *existential quantifier conditional on* $\mathcal{G}$, is the unique mapping **Form** $\longrightarrow \mathcal{G}$ satisfying

(*) For all $C \in \mathcal{G}$ $[\varphi \subseteq C$ iff $\exists(\varphi|\mathcal{G}) \subseteq C]$

Condition (*) implies that $\exists(\varphi|\mathcal{G})$ must be defined as $\bigwedge\{C \in \mathcal{G} | \varphi \subseteq C\}$. Since $\mathcal{G}$ is a frame, this meet exists in $\mathcal{G}$, hence indeed $\exists(\bullet|\mathcal{G}) :$ **Form** $\longrightarrow \mathcal{G}$. From this we also see that $\exists(\bullet|\mathcal{G})$ has the following properties

1) $\exists(\mathbf{0}|\mathcal{G}) = \mathbf{0}$, $\exists(\mathbf{1}|\mathcal{G}) = \mathbf{1}$;

2) $\varphi \subseteq \psi$ implies $\exists(\varphi|\mathcal{G}) \subseteq \exists(\psi|\mathcal{G})$ (monotonicity);
3) $\varphi \subseteq \exists(\varphi|\mathcal{G})$ (coarsening);
4) $\exists(\varphi \vee \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \vee \exists(\psi|\mathcal{G})$ (additivity);
5) $\exists(\varphi \wedge \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \psi$ where $\psi \in \mathcal{G}$ ('taking out what is known')

Note that (2) and (6) imply that $\exists(\bullet|\mathcal{G})$ is the identity on $\mathcal{G}$

$$\exists(\varphi|\mathcal{G}) = \exists(\mathbf{1} \wedge \varphi|\mathcal{G}) = \exists(\mathbf{1}|\mathcal{G}) \wedge \varphi = \mathbf{1} \wedge \varphi = \varphi$$

for $\varphi \in \mathcal{G}$. Thus, (6) is the analogue of the Frobenius property in logic

$$\exists(\varphi \wedge \exists(\psi|\mathcal{G})|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \exists(\psi|\mathcal{G}).$$

The relation between conditional quantifiers and filtered quantifiers defined by means of an equivalence relation is clarified by the following proposition.

Theorem 2 *Let R be an equivalence relation determined by an algebra $\mathcal{B}$ of first order definable sets of assignments of an ω_1-saturated model such that*

$R(f, g)$ *iff* $\forall C \in \mathcal{B}[f \in C \Leftrightarrow g \in C]$

and let the filtered quantifier $\exists$ be defined by

$f \in \exists\varphi$ *iff* $\exists g[R(f, g)\ \&\ g \in \varphi]$

Let $\mathcal{G}$ be the frame generated by $\mathcal{B}$, then $\exists\varphi = \exists(\varphi|\mathcal{G})$, for φ a first order formula.

Proofsketch One first shows by means of a compactness argument that $\exists\varphi \in \mathcal{G}$ (here one uses ω_1-saturation); this then yields $\exists(\varphi|\mathcal{G}) \subseteq \exists\varphi$ by Galois. For the converse direction, note that $\exists\varphi$ satisfies a Galois correspondence with respect to $\mathcal{B}$; since $\varphi \subseteq \exists(\varphi|\mathcal{G})$ and $\exists(\varphi|\mathcal{G})$ is an intersection of elements from $\mathcal{B}$, we have $\exists\varphi \subseteq \exists(\varphi|\mathcal{G})$. 2

Now that we have defined conditional quantifiers, we may extend our language with quantifiers '$\exists(\bullet|\mathcal{G})$', where $\mathcal{G}$ must be interpreted as a frame. This allows for formulas involving iterated conditional quantification. For reasons which will gradually become clear we are mainly interested in positive formulas.

Definition 8 A first order formula is *positive*, iff it is constructed using $\vee, \wedge, \exists, \forall$ and $\perp$. Henceforth, **Pos** denotes the frame generated by sets of assignments definable by positive formulas.

Definition 9 A formula in a first order language with conditional quantifiers is *positive* if it is constructed using $\vee, \wedge, \exists, \forall, \perp$ and $\exists(\bullet|\mathcal{G})$.

Lemma 3 *Let $\mathcal{G}$ be a subframe of* **Pos**. *For any positive formula φ (in the extended sense) $\exists(\varphi|\mathcal{G}) \in \mathcal{G}$.* 2

3.0.2 *Refining systems from filtering

In what follows we work with assignment spaces $\mathcal{F}_s := |\mathcal{M}_s|^{\mathrm{Var}}$ and $\mathcal{F} := |\mathcal{M}|^{\mathrm{Var}}$ on the models $\mathcal{M}_s$ and $\mathcal{M}$. Let $\mathcal{F}_s$ be given the topology generated by the clopen sets $[\![\varphi]\!]_{\mathcal{M}_s}$, this will be called the *formula topology* on $\mathcal{F}_s$. Since the models $\mathcal{M}_s$ are finite, they are compact Hausdorff in the discrete topology. Hence the product topology on $\mathcal{F}_s$ is compact Hausdorff as well, so that the formula topology on $\mathcal{F}_s$, as a subtopology of the product topology, is compact.

Define a projection $\mathbf{p}_s : \mathcal{F} \longrightarrow \mathcal{F}_s$ by: $\mathbf{p}_s(f)(x) := f(x)_s$. Sometimes we write '$f_s$' instead of: $\mathbf{p}_s(f)$. The projections $\mathbf{p}_s$ are surjective: Let $g \in \mathcal{F}_s$. For each x, choose ξ^x such that $g(x) = \pi_s(\xi^x)$. This is possible, because the projections $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ are surjective. Define: $f(x) := \xi^x$, then $\mathbf{p}_s(f) = g$.

Let $\mathbf{T}$ be the smallest topology on $\mathcal{F}$ which makes all the $\mathbf{p}_s$ continuous, when the $\mathcal{F}_s$ are given the formula topology. $\mathbf{T}$ is a subtopology of the product topology on $\mathcal{F}$. Since $|\mathcal{M}|$ is compact Hausdorff in the topology induced by the $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ (Engelking 1989, 141), so is $\mathcal{F}$; whence $\mathbf{T}$ is compact. $\mathbf{T}$ is the topology that is used in theorems and proofs.[3]

To motivate the following theorem, we return to the beginning of this chapter, where it was argued that perception should be viewed as applying a 'filter' to reality. Here we show that 'filtering' is in a sense inverse to taking limits of refining systems. Of necessity, the formulation of the result is somewhat sloppy; a full formulation and proof will be published elsewhere.

Definition 10 Let T be a directed set. A collection of algebras $\{\mathcal{B}_s : s \in T\}$ is a *net* if we have: $s \leq t$ implies $\mathcal{B}_s \subseteq \mathcal{B}_t$.

Theorem 4 *Let $\mathcal{M}$ be an ω_1-saturated model, $\mathcal{B}$ the algebra of assignments on $\mathcal{M}$. Let $\{\mathcal{B}_s : s \in T\}$ be a net of algebras such that $\bigcup \mathcal{B}_s = \mathcal{B}$. Then there exists an inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ such that*

1. *$\mathcal{M}$ is the inverse limit of $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$;*
2. *$\mathcal{M}_s, \mathbf{p}_s(f) \models R(x_1, \dots, x_n)$ iff $f \in \exists(R(x_1, \dots, x_n)|\mathcal{B}_s)$, for all relations $R(x_1, \dots, x_n)$.*

Condition (2) says that the interpretation of predicates on $\mathcal{M}_s$ is given by $\exists(\bullet|\mathcal{B}_s)$; i.e. $\mathcal{M}_s$ is the model that corresponds to the filter $\exists(\bullet|\mathcal{B}_s)$. This result is parallel to the observation in situation semantics that a first order model can always be viewed as the union of a directed set of partial submodels, but note the difference: the models $\mathcal{M}_s$ are not submodels of $\mathcal{M}$, but approximations to $\mathcal{M}$. This reflects the subtly different concept of partiality employed here.

4 *From refining systems to conditional quantifiers and back

This section is not required for an understanding of the body of the paper; it is included here because it justifies the logical form of perception reports adopted in the next section.

Given the material of the previous section, an obvious question arises. Suppose we start from an inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ with inverse limit $\mathcal{M}$. Let $\mathcal{B}_s$ be the

[3] There exists another topology on $\mathcal{F}$ determined by first order definable subsets of assignments on $\mathcal{M}$; cf. section 3.0.1. This topology will not be used.

following lattice:

$$\mathcal{B}_s := \{C \subseteq \mathcal{F} : C = \mathbf{p}_s^{-1}(\{g \in \mathcal{F}_s : \mathcal{M}_s, g \models \varphi\}) \text{ for some } \textit{positive } \varphi\}$$

Let $\mathcal{G}_s$ be the frame generated by $\mathcal{B}_s$. If A is any predicate, do we have:

$$\mathcal{M}_s, \mathbf{p}_s(f) \models A \Leftrightarrow f \in \exists(A|\mathcal{G}_s)?$$

This is almost true, as the next results show.

Theorem 5 *If φ is positive, $\varphi \subseteq \mathbf{p}_s^{-1}(\varphi_s)$.*

Proof. If φ is atomic, this is the truth definition of φ on $\mathcal{M}$. The cases $\vee$, $\wedge$ are trivial. Suppose $f \in \exists x\psi$, then there exists $g =_x f$ such that $g \in \psi$. By the inductive hypothesis, $g_s \in \psi_s$. But $g =_x f$ implies $g_s =_x f_s$, whence $f_s \in \exists x\psi$. For the case $f = \forall x\psi$ we need the surjectivity of the projection mappings. Indeed, let $f \in \forall x\psi$; we have to show that $f_s \in (\forall x\psi)_s$. Let $h \in \mathcal{F}_s$ be such that $h =_x f_s$. By surjectivity, there is $g \in \mathcal{F}$ such that for all $t \in T : g_t(y) = f_t(y)$ $(y \neq x)$ and $g_s(x) = h(x)$; then $g_s = h$ and $g \in \psi$, whence $h = g_s \in \psi_s$. 2

Corollary 6 *If φ is positive, $\exists(\varphi|\mathcal{G}_s) \subseteq \mathbf{p}_s^{-1}(\varphi_s)$.*

Proof. Since $\mathbf{p}_s^{-1}(\varphi_s) \in \mathcal{G}_s$, the statement follows from the Galois correspondence for $\exists(\bullet|\mathcal{G}_s)$. 2

This result is best possible in the sense that we cannot expect it to hold even for negations of atomic formulas: $\neg A$ may be false on $\mathcal{M}$ because it fails on a coordinate $t > s$, and in that case $\exists(\neg A|\mathcal{G}_s)$ will not be contained in $\mathbf{p}_s^{-1}(\neg A_s)$. However, the corollary can be strengthened in the sense that under an additional, harmless, condition, we have equality when φ is a predicate.

Definition 11 The pair $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ is *proper refining* with respect to a predicate A, if there does not exist a positive formula τ such that $A \subseteq \mathbf{p}_s^{-1}(\tau_s)$ and $\emptyset \neq \tau_s \subset A_s$.

As the next proposition shows, if the pair $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ is proper refining with respect to a predicate A, then $A_s = \{g \in \mathcal{F}_s : \mathcal{M}_s, g \models A\}$ represents our best estimate of A at stage s.

Corollary 7 *If the pair $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ is proper refining with respect to a non-empty predicate A, then $\exists(A|\mathcal{G}_s) = \mathbf{p}_s^{-1}(A_s)$.*

Proof. Since A is non-empty, we may assume $\exists(A|\mathcal{G}_s) \neq \mathbf{0}$. Suppose $\exists(A|\mathcal{G}_s)$ is strictly contained in $\mathbf{p}_s^{-1}(A_s)$. Since $\exists(A|\mathcal{G}_s) = \bigcap\{C \in \mathcal{B}_s : A \subseteq C\} \subset \mathbf{p}_s^{-1}(A_s)$, by compactness there must be $C \in \mathcal{B}_s$ such that $\exists(A|\mathcal{G}_s) \subseteq C \subset \mathbf{p}_s^{-1}(A_s)$. There is clopen $C' \subseteq \mathcal{F}_s$ such that $C = \mathbf{p}_s^{-1}(C')$. Note that C' can be defined by means of a positive formula. We then have $C' = \mathbf{p}_s"(\mathbf{p}_s^{-1}(C')) \subset A_s$ and $A \subseteq C = \mathbf{p}_s^{-1}(C')$. This contradicts the fact that $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ is proper refining with respect to A.
2

We can always transform a refining inverse system $\langle\mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ into a new inverse system $\langle\mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ with the same inverse limit $\mathcal{M}$, such that all pairs

$\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ are proper refining with respect to all predicates. Indeed, for any predicate A and any $s \in T$, put

$$A'_s := \bigcap\{\tau_s : \tau \text{ positive}, \emptyset \neq \tau_s \subseteq A_s, A \subseteq \mathbf{p}_s^{-1}(\tau_s)\}$$

This gives us a new refining inverse system of models $\mathcal{M}'_s$, with the same bonding mappings as before. (The condition that the τ are positive is essential in showing that the bonding mappings are homomorphisms with respect to the new predicates.) It then easily follows that $\mathcal{M}$ is also the inverse limit of the $\mathcal{M}'_s$, for we have: if $\mathcal{M}, f \models A$, then

$$\forall s \forall \tau (\tau \text{ positive}, \emptyset \neq \tau_s \subseteq A_s, A \subseteq \mathbf{p}_s^{-1}(\tau_s) \Rightarrow f \in \mathbf{p}_s^{-1}(\tau_s))$$

which implies $\mathcal{M}, f \models A'$. The converse is trivial. The new inverse system has the property that, at each stage s, the interpretation of A on $\mathcal{M}_s$ represents our best estimate of A at stage s.

Lastly, we extend Corollary 7 to formulas of the form $\exists x_1 \ldots \exists x_n \theta$, θ a conjunction of predicates.

Theorem 8 *Let φ be a formula of the form $\exists x \theta$, where θ is a conjuntion of atomic formulas; we assume θ is non-empty on $\mathcal{M}$. Assume furthermore that $\pi_s : \mathcal{M} \longrightarrow \mathcal{M}_s$ is proper refining w.r.t. θ. If x is a variable, let $\mathcal{B}_{s,x}$ be the subalgebra of $\mathcal{B}_s$ determined by formulas in which x does not occur free, and let $\mathcal{G}_{s,x}$ be the frame generated by $\mathcal{B}_{s,x}$. Then we have*

$$\mathcal{M}_s, g \models \exists x \theta \Leftrightarrow \forall f \in \mathcal{F}[f_s = g \Rightarrow \mathcal{M}, f \models \exists(\theta|\mathcal{G}_{s,x})]$$

Proof. Suppose $\mathcal{M}_s, g \models \exists x \theta$, and let f satisfy $f_s = g$; we have to show $f \in \exists(\theta|\mathcal{G}_{s,x})$. There exists $g' =_x g$ such that $\mathcal{M}_s, g' \models \theta$. By surjectivity of $\mathbf{p}_s$ there exists f' such that $f' =_x f$. By corollary 7, we must have $f' \in \exists(\theta|\mathcal{G}_s)$, hence $f' \in \exists(\theta|\mathcal{G}_{s,x})$. A little reflection shows that $\exists(\theta|\mathcal{G}_{s,x})$ is $=_x$-invariant, whence $f \in \exists(\theta|\mathcal{G}_{s,x})$. Conversely, suppose that for f with $f_s = g$ we have $\mathcal{M}, f \models \exists(\theta|\mathcal{G}_s, x)$. By theorem 5, $\theta \subseteq \exists x \theta \subseteq \mathbf{p}_s^{-1}((\exists x \theta)_s)$. Since $\mathbf{p}_s^{-1}((\exists x \theta)_s) \in \mathcal{G}_{s,x}$, by the Galois property $\exists(\theta|\mathcal{G}_{s,x}) \subseteq \mathbf{p}_s^{-1}((\exists x \theta)_s)$, in particular $\mathcal{M}_s, g \models \exists x \theta$. 2

Corollary 9 *Assuming the conditions of the theorem, we also have*

$$\mathcal{M}_s, g \models \exists x \theta \Leftrightarrow \forall f \in \mathcal{F}[f_s = g \Rightarrow \mathcal{M}, f \models \exists x \exists(\theta|\mathcal{G}_s)]$$

The extension to formulas of the form $\exists x_1 \ldots \exists x_n \theta$ is straightforward. This result explains the particular formalisation of perception reports adopted in the next section.

5 Object recognition revisited

We briefly return to Marr's theory of object recognition involving three-dimensional models, to explain the use of conditional quantifiers in the analysis of perception reports. We quote Marr *in extenso* because his description is very suggestive.

> Recognition involves two things: a collection of stored 3-D model descriptions, and various indexes into the collection that allow a newly derived description to be associated with a description in the collection. We shall refer to the above collection along with its indexing as the catalogue of 3-D models... [T]hree access paths into the catalogue appear to be particularly useful. They are the specificity index, the adjunct index, and the parent index. [...] A newly derived 3-D model may be related to a model in the catalogue by starting at the top of the hierarchy and working down the levels through models whose shape specifications are consistent with the new model's until a level of specificity has been reached that corresponds to the precision of the information in the new model. Once a 3-D model for a shape has been selected from the catalogue, its adjunct relations provide access to 3-D models for its components based on their locations, orientations, and relative sizes. This gives us another access path to the models in the catalogue, which we call the adjunct index... [T]he adjunct index provides useful defaults for the shapes of the components of a shape prior to the derivation of 3-D models for them from the image. [...] The third access path that we consider important is the inverse of the second, and we shall call it the parent index of a 3-D model. When a component of a shape is recognised, it can provide information about what the whole shape is likely to be. [...] It is important to note that the adjunct and the parent indexes play a role secondary to that of the specificity index, upon which our notion of recognition rests... [T]heir purpose is primarily to provide contextual constraints that support the derivation process... (Marr 1982, pp. 318–21)

The picture that emerges from this description is the following. We perceive an object x; the resulting retinal image is processed, by compressing information and imposing an object-centered coordinate system, until we arrive at a 3-D image. To determine which object this is an image of, one indexes into the catalogue of 3-D models, in our notation the refining inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$. One traces the beginning of a thread through these models, until the finite resolution of the image leads to different possible continuations of the initial segment of the thread. The net result can be described as 'object x has property A with resolution R'. The resolution of the image corresponds to an equivalence relation on the side of the catalogue of 3-D models. We see that object recognition essentially involves approximation: we start with a real object which has an infinite amount of detail and we end up with a symbolic description, meanwhile filtering out very many of those details. This process of filtering is summarised in the conditionally quantified formula '$\exists(A(x)|\mathcal{G})$', where $\mathcal{G}$ is the frame corresponding to the equivalence relation R (cf. theorem 2).

Marr proposes this model for the recognition of objects by means of shapes, on the grounds that only a modular organisation of this type allows one to cope with partial information of various sorts (finite accuracy, observation of a part not a whole, sufficiency of a rough description for the task at hand) in a computationally efficient manner. For example, it allows one to do the recognition process iteratively, by successive cycles of obtaining information from the image, matching this information against the catalogue of 3-D models, deriving a predication from the model chosen, now checking the predication against the image, perhaps analysed in finer detail. This process will be illustrated below. The very abstract model proposed here imposes this type of organisation for any property whatsoever, not just shape. One of course loses all connection with computational considerations in this way, but there are some advantages to be gained. Firstly, one does not seem to need a detailed computational model to account for the logic of perception reports. Secondly, the abstract model allows one to prove metamathematical results which bear on the process of object recognition, involving the two other indices discussed by Marr.

Consider the adjunct index. A concrete instance of its use is the following: if we have identified a shape as a human body, then we expect it to have two arms, even though we may not yet have observed them. That is, we consider a projection ξ_t of the thread ξ, the object underlying the perceived shape, in a model $\mathcal{M}_t$ in which there are no arms. The constraint 'human beings have a left and a right arm' is expressed by means of a sentence of the form $\forall x(\varphi \rightarrow \exists y \exists z \psi)$, where φ, ψ are quantifier free, and is assumed to be true of some more refined approximation $\mathcal{M}_s$, where $s \geq t$. I.e., $\mathcal{M}_s$ is an approximation in which there are arms. We now expect the property φ, observed in $\mathcal{M}_t$, to hold in $\mathcal{M}_s$ as well (how this expectation should be formalised is discussed in chapter 4), and we conclude $\exists y \exists z \psi$. The said constraint will be true in all sufficiently refined approximations, i.e. in all $s \geq t$, hence in all those $\mathcal{M}_s$ we will have $\exists y \exists z \psi$. The following natural question then arises: suppose that indeed for all $s \geq t$, $\mathcal{M}_s \models \exists y \exists z \psi$, does this imply that $\exists y \exists z \psi$ is true in reality, i.e. on the inverse limit $\mathcal{M}$ of the refining system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$? More generally, given a constraint of the form $\forall(\theta \rightarrow \tau)$, used in the manner above, for which τ can we be sure that $\exists t \forall s \geq t \mathcal{M}_s \models \tau$ implies $\mathcal{M} \models \tau$? This question is solved in section 6, on preservation theorems for inverse limits.

The same question applies to the antecedens of a constraint. Object recognition often proceeds iteratively (cf. Marr 1982, 313): one extracts information θ from an image, say at resolution t; one then looks for constraints which contain θ in their antecedent, constraints which typically apply only at levels $s \geq t$; this yields a disjunction of properties $\psi_1 \vee \ldots \vee \psi_n$ which is then checked against the image, ideally yielding one more piece of information in addition to θ, although the outcome may also be a rejection of θ. This procedure works only when the information θ extracted from the image is reliable, when it tells us something about reality. The logical question that immediately comes to mind is then the following: which information θ *can* be reliable, i.e. has the property that $\exists t \forall s \geq t \mathcal{M}_s \models \theta$ implies $\mathcal{M} \models \theta$? Put differently: how should one choose a representation?

Hence the general question is: what is the logical form of statements suitable for expressing information extractable from an image, and by implication, what is the logical form of the constraints which relate those pieces of information? This is of course an abstract version of one of Marr's main questions, namely, what can one derive from the available information?

6 *Preservation theorems

To motivate the material in this section, let us consider the following question. By the definition of inverse limit, the interpretation of a predicate A on $\mathcal{M}$ can be approximated by its interpretations on the models $\mathcal{M}_t$, in the sense that

$$\exists s \forall \xi (\xi \in A \text{ iff } \forall t > s(\xi_t \in A_t).$$

Topologically speaking, A determines a closed subset on $\mathcal{F}$. This has as a consequence that, if B is another predicate disjoint from A on $\mathcal{M}$, there will be a 'finite' stage t such that $A_t \cap B_t = \emptyset$, so that if some object ξ is an A but not a B we will know this after 'finite time': we do not need complete information about ξ. In principle, one could also allow predicates which are defined on the inverse limit only, not on the approximating models; if C such a predicate, one would need an infinite amount of information to decide whether $\xi \in C$. For a situation where this would be natural, cf. section 5.1.[4] But even if we do not allow this, there still arises the question which properties $\varphi(x)$ behave as predicates do, i.e., which properties $\varphi(x)$ determine closed sets. Below we give a complete answer by means of several preservation theorems.

Recall that an inverse system is *total and refining* if the bonding mappings are homomorphisms (hence by definition surjective), and *refining* if the bonding mappings h_{st} just satisfy condition (ρ) of section 2. The first concept reflects the case where all objects in an approximation are supposed to correspond to real objects, whereas the second concept allows for purported objects to vanish under closer scrutiny. A formula is *positive* if it is equivalent to a formula containing only $\exists, \forall, \wedge, \vee$ and $\perp$; a formula is *positive primitive* if it is equivalent to a formula containing only $\exists, \wedge, \vee$ and $\perp$.

A sentence θ is *preserved* by an inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ if whenever $\mathcal{M}_s \models \theta$ and the inverse limit $\mathcal{M}$ is nonempty, then $\mathcal{M} \models \theta$. To simplify the fomulation of definitions and theorems we assume that all structures involved are models for the same signature. The results easily transfer to the case where for any predicate A, there is $t \in T$ such that for $s \geq t$, A is interpreted on $\mathcal{M}_s$.

Theorem 10 (1) *Let $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ be a total and refining inverse system of finite models. Then any sentence of the form*

$$\forall x_1, \ldots, x_n (\varphi \to \psi),$$

[4] One might say that in this situation, conditional quantifiers come into their own: they allow one to estimate the extension of a predicate on a model $\mathcal{M}_t$ where the predicate is not interpreted.

where φ and ψ are positive formulas, is preserved by $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$.

(2) *Let $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ be a refining inverse system of finite models. Then any sentence of the form $\forall x_1, \ldots, x_n(\varphi \to \psi)$, where φ and ψ are positive primitive formulas, is preserved by $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$.*

The proof is easy; its main ingredient is the following

Lemma 11 (1) *Let $\langle \mathcal{M}_s, h_{st}\rangle_{s,t\in T}$ be a total and refining inverse system of finite models. Let $\varphi(x)$ be a positive formula. Then we have $\mathcal{M} \models \varphi(\xi)$ iff for all t, $\mathcal{M}_t \models \varphi(\xi_t)$.*

(2) *Analogously for refining inverse systems and positive primitive formulas.*

As we saw above, finiteness of the $\mathcal{M}_s$ ensures that the topology on the inverse limit is compact; this is essential if we want to allow existential quantification in ψ. Note that parts (1) and (2) of the lemma imply *inter alia* that the inverse limit is nonempty.[5] The proof of the lemma tells us something about *negations of positive formulas* as well: not only $\mathcal{M} \not\models \varphi(\xi)$ iff for some t, $\mathcal{M}_t \not\models \varphi(\xi_t)$, but actually if $\mathcal{M}_t \not\models \varphi(\xi_t)$ then for $s \geq t$, $\mathcal{M}_s \not\models \varphi(\xi_s)$. Here are examples of preserved formulas.

i) The class of preserved sentences, although defined by means of positive formulas, includes sentences involving negation: e.g.,

$$\forall x(\neg\varphi \to \psi)$$

is equivalent to $\forall x(\varphi \vee \psi)$;

$$\forall x(\varphi \to \neg\psi)$$

is equivalent to $\forall x(\varphi \wedge \psi \to \bot)$; and

$$\exists x \neg\varphi$$

is equivalent to $\forall x(\varphi \to \bot)$. Here, φ and ψ are positive. On the other hand, a simple formula such as $\exists x(A(x) \wedge \neg B(x))$ is not preserved. At any particular stage we can be uncertain whether an A is in fact also a B, even though in the limit there is no longer uncertainty.

ii) Universal Horn formulas: $\forall x_1, \ldots, x_n(\varphi \to \psi)$, where φ is a conjunction of predicates, and ψ is a predicate; these formulas express the presence of a law.

iii) Formulas which express 'there are at most n x such that φ', where φ is positive, since these formulas are of the form

$$\forall x, x_1, \ldots, x_n(\varphi \to x = x_1 \vee \ldots \vee x = x_n).$$

iv) Let us present the fragment preserved in a form familiar to linguists. As before, think of $\exists$ and $\forall$ as binary determiners

$$\exists x_1, \ldots, x_n(\varphi, \psi),\ \forall x_1, \ldots, x_n(\varphi, \psi),$$

then the fragment isolated in part (1) of the theorem, the *stable* formulas, are defined inductively as follows:

[5] The proof that disjunctions are preserved upwards contains a subtlety: it uses the fact that we may take a cofinal subset T' of T and still obtain an inverse limit isomorphic to $\mathcal{M}$.

a. Atomic formulas are stable.
b. If φ, ψ are positive, then $\exists x_1, \ldots, x_n(\varphi, \psi)$, $\forall x_1, \ldots, x_n(\varphi, \psi)$ are stable.
c. Stable formulas are closed under $\vee$ and $\wedge$.

Note that $\exists x_1, \ldots, x_n(\varphi, \psi)$ and $\forall x_1, \ldots, x_n(\top, \psi)$ are positive (where $\top$ is a tautology); in general, however, $\forall x_1, \ldots, x_n(\varphi, \psi)$ is not positive, so we cannot sharpen (iv,b) to: 'if φ, ψ are stable, then $\exists x_1, \ldots, x_n(\varphi, \psi)$ and $\forall x_1, \ldots, x_n(\varphi, \psi)$ are stable'.

To show that stable formulas are closed under $\vee$, observe that

$$\forall x_1, \ldots, x_n(\varphi, \psi) \vee \forall x_1, \ldots, x_n(\theta, \tau)$$

is equivalent to $\forall x_1, \ldots, x_n(\varphi \wedge \theta, \psi \vee \tau)$, which is stable if φ, ψ, θ and τ are positive.

Referring to part (2) of the theorem one can define an analogous concept of stable primitive formulas, in which clause (3) is modified to: 'if φ, ψ are positive primitive, then ...'.

Let us now return to the motivating example at the beginning of this section. One can show the following:

Theorem 12 *The following are equivalent:*

i) *For any inverse limit $\mathcal{M}$ of a total and refining inverse system, $\varphi(x)$ determines a closed set $\{f : \mathcal{M}, f \models \varphi\} \subseteq \mathcal{F}$*

ii) *φ is positive.*

PROOFSKETCH That (ii) implies (i) follows from Lemma 11. The converse uses Lyndon's theorem (cf. Chang and Keisler 1990, theorem 5.2.13). Analogous results hold for inverse limits of refining inverse systems. 2

An iterated use of Lyndon's theorem allows one to prove a stronger result. Note that a conspicuous example of a non-stable formula is the formula which expresses 'there are at least n x such that φ', i.e., $\exists x, x_1, \ldots, x_n(\neg(x = x_1) \wedge \ldots \wedge \neg(x = x_n) \wedge \varphi(x))$, and indeed it can be shown that such formulas need not be preserved. This is why statements such as 'a human being has two arms', in order to be preserved, have to be formulated positively, for example as 'a human being has a left and a right arm'. That this is no accident is shown by the next theorem.

Theorem 13 (1) *Suppose θ is preserved by all total and refining inverse systems. Then θ is equivalent to a conjunction of sentences of the form*

$$\forall x_1, \ldots, x_n(\varphi \to \psi),$$

where φ and ψ are positive formulas.

(2) *Suppose θ is preserved by all refining inverse systems. Then θ is equivalent to a conjunction of sentences of the form*

$$\forall x_1, \ldots, x_n(\varphi \to \psi),$$

where φ and ψ are positive primitive formulas.

The proof of this result is too complicated to be given here. What the theorems taken together show, is that the essence of preservation on inverse limits is given by Lemma 11; namely that positive (primitive) formulas are preserved upward.

4

Perception, or: non-monotonicity, the hard way

In the analytical philosophical tradition, and elsewhere, direct perception has been thought of as a hallmark of certainty; sense data were presented as a solid foundation on which to erect knowledge. In logical semantics remnants of this position can be found when direct perception reports are taken to be irretractable if true. This position is rather problematic. Philosophy of science has made clear that a form of perception which is epistemologically or theoretically neutral is hard to defend. The empirical basis of knowledge is open to revision, just like knowledge itself.

Logic of vision is based on the idea that we only have direct perceptual access to fallible approximations of material objects, not to the objects themselves. Material objects are more like Kantian 'things in themselves'. On this view the properties of a perceived object are in principle hypothetical; it may be impossible to decide whether they are due to material objects, to our cognitive abilities, or to intermediaries. Normally one disregards the hypothetical nature of perceived objects; but this is largely a pragmatic issue, based on the fact that their properties are often stable under change of perceptual information. It is this stability which makes it natural to assume that there are material objects bearing them. Stability also induces expectations as to how an object will appear when perspectives are altered. Nevertheless, properties may be erroneously ascribed; hypotheses expected to be true may still be refutable, and hence they should be treated as such.

This view on perception, which section 5.3 discusses in more detail, should have consequences for the logical semantics of direct perception reports. The semantics developed in section 2 is based on fields of vision as approximations to a part of reality, and relative to an approximation perception reports need not be veridical. As for other kinds of inference, the (in)validity of veridicality depends (i) on the perceptual resources available for conditional quantification, and (ii) on the pragmatic expectation that perception reports are stable under refinement of approximations. The results of section 3.6 can be used to show that if this expectation is justified for positive perception reports, veridicality follows. Such expectations can be pitched too high, though, and should therefore correspond to defeasible rules; as in the formalised pragmatics of section 3.

This chapter ends with a study of pegs in logic of vision. But first we concentrate on the logical form of perception reports. The inferences of chapter 2 are discussed as we go along.

1 The logical form of perception reports

The present set up sustains a natural distinction between directly and indirectly perceived objects. A *directly* perceived object is represented by a free variable, which can only be 'measured' with finite precision; the degrees of precision being given by conditional quantifiers $\exists(\bullet|\mathcal{G})$. Instead, '*indirectly*' perceived objects, such as the doctor in (1), are represented by means of bound variables.

(1) Jack saw Sharon phone a doctor.

Some examples should clarify the distinction further.

The direct perception report (2) is rendered formally as (3) (cf. the propositions of form (18), chapter 2, in Grice's analysis of perception reports).

(2) I see this arm.

(3) $\exists(A(x)|\mathcal{G}_i)$

Here, the demonstrative 'this' is interpreted as the variable x, which receives its value from a contextually given assignment. So, x corresponds with the arm perceived directly, and the frame $\mathcal{G}_i$ with the filter of the subject's perceptual field.

There are two ways to interpret proper names. Firstly, one may interpret (4) with a demonstrative element as: 'this object I see is Sharon'. Formally this corresponds to having a free variable in the representation, as in (5).

(4) I see Sharon.

(5) $\exists(x = s|\mathcal{G}_i)$

The constant s, for Sharon, denotes a thread in the inverse limit. Secondly, one may read (4) indefinitely as: 'an object I see is Sharon', as in (6).

(6) $\exists x \exists(x = s|\mathcal{G}_i)$

Theorem **??** shows that for rich enough frames (6) can be written as the open (7), *provided* the algebra has no information about the variable x.

(7) $\exists(x = s|\mathcal{G}_{i,x})$

In what follows, we shall often use the open form. Here are some further examples:

(8) 'I see an arm' becomes: $\exists x \exists(A(x)|\mathcal{G}_i)$; and

(9) 'I see Russell wink' becomes: $\exists(x = r \wedge W(x)|\mathcal{G}_i)$ or $\exists x \exists(x = r \wedge W(x)|\mathcal{G}_i)$.

An example of a report with an indirectly perceived object, as in (1), is (10) meaning (11).

(10) I saw Jack walk towards an ant hill.

(11) $\exists(x = j \wedge \exists y[Ah(y), W(x,y)]|\mathcal{G}_i)$

As did Hintikka 1969, the examples show that the semantics for perception reports describing objects is similar to those describing scenes; mainly the descriptive content within the scope of 'to see' varies. Below we discuss ways of interpreting arbitrary quantified noun phrases in perception statements.

It should be noticed that in all the examples the perception operator has an open statement within its scope; some of the objects described must be seen directly. This is crucial, for a frame $\mathcal{G}$ acts trivially on 'closed' truths or falsities; $\exists(\mathbf{1}|\mathcal{G}) = \mathbf{1}$ and $\exists(\mathbf{0}|\mathcal{G}) = \mathbf{0}$. From this we also see that substitution may fail

$$x = t \ \& \ \exists(\varphi(x)|\mathcal{G}) \not\Rightarrow \exists(\varphi(t)|\mathcal{G})$$

with t a constant or a variable. Section 5.2, on evidentials, indicates how to generalise the distinction between directly and indirectly perceived *objects* to direct and indirect perception *reports*.

2 Possibly non-veridical perception reports

The semantics of direct perception reports will be based on the idea that a field of vision approximates part of reality. Reality itself appears as a 'regulative ideal'; i.e., an inverse limit of first order models. So, the most important issues left for a semantics of possibly non-veridical NI reports are to specify the rôles of approximation and partiality. The relation between these concepts is given by the insight that a perceptual field does two things: (i) it determines a part of reality corresponding to what is seen, and (ii) it determines the filter through which this part is seen. Given the discussion in chapter 3, it is clear that an approximation can be modelled by a suitable conditional quantifier on the inverse limit. So it remains to import the notion of partiality. We do so in a discussion of the (non-)inference called partial perception.

2.0.3 Partial perception

Perhaps the first requirement for an analysis of 'see' is that it should be able to cope with reports of the following kind, employing NPs, infinite VPs, and PPs:

$$\frac{\text{Whitehead saw Russell} \quad \text{Russell winked}}{\text{Whitehead saw a winking man}}$$

$$\frac{\text{Whitehead saw Russell} \quad \text{Russell winked}}{\text{Whitehead saw Russell wink}}$$

$$\frac{\text{Whitehead saw Russell} \quad \text{Russell had his shirt unbuttoned}}{\text{Whitehead saw Russell with his shirt unbuttoned}}$$

The common feature of these examples of invalid inferences is that the sentence 'Whitehead saw Russell' refers to a stage of approximation s, whereas the factual statements are true in an approximation t possibly different from s. In particular, t may be incomparable to s—e.g., if Whitehead saw Russell from behind—or t may be a refinement of s but still the distance between Whitehead and Russell could have been too large for Whitehead to actually see, e.g., that Russell had his shirt unbuttoned.

Up until now we have formalised the possibly non-veridical semantics of 'w *saw* φ', φ positive, as $\exists(\varphi|\mathcal{H})$, with $\mathcal{H}$ a suitable frame assigned to perceiver j. This is still too crude, though, for it forces us to declare the above inferences valid; by monotonicity we must have

$$\exists(x = r|\mathcal{G}_s) \wedge \forall x(x = r \rightarrow W(x)) \rightarrow \exists(x = r \wedge W(x)|\mathcal{G}_s),$$

and analogously for the other examples.

However, recall that conditional quantifiers express *consistency*:

$$\exists(x = r \wedge W(x)|\mathcal{G}_s)$$

says that it is consistent with the information available to Whitehead that Russell winks. This would also be the case, for instance, if $\exists(x = r \wedge W(x)|\mathcal{G}_s)$ equals the full set of assignments $\mathcal{F}$, i.e., when no nontrivial element of $\mathcal{G}_s$ dominates $x = r \wedge W(x)$. Clearly, one would not call this 'seeing'. Accordingly, to accomodate partial perception the definition of the conditional quantifier has to be modified slightly, in such a manner that it is defined only when it represents nontrivial information.

Definition 1 A *pseudolattice* L is a partially ordered set in which finite *non-empty* meets and joins exist. (Hence L need not have top or bottom.) A pseudolattice L is an *evidential* $\vee, \bigwedge$*-frame* if it is closed under arbitrary non-empty meets, such that the following distributive law holds:

$$a \vee \bigwedge_I b_i = \bigwedge_I (a \vee b_i).$$

In this definition the meet $a \wedge b$ of elements a, b is non-empty, and whence in L, even if $a \wedge b = \mathbf{0}$. We now define quantification with respect to an evidential frame.

Definition 2 Let $\mathcal{G}$ be an evidential frame. $\exists(\bullet|\mathcal{G})$, the *existential quantifier conditional on* $\mathcal{G}$, is the unique mapping **Form** $\longrightarrow \mathcal{G}$ satisfying

(*) for all $C \in \mathcal{G}$: $\varphi \subseteq C$ if and only if $\exists(\varphi|\mathcal{G}) \subseteq C$.

As before, (*) implies that $\exists(\varphi|\mathcal{G})$ must be defined as $\bigwedge\{C \in \mathcal{G}|\varphi \subseteq C\}$, *when the meet exists in* $\mathcal{G}$. In general, however, $\exists(\bullet|\mathcal{G})$ will be a partial map: **Form** $\longrightarrow \mathcal{G}$.

The reader may wish to check that a quantifier $\exists(\bullet|\mathcal{G})$ conditional on an evidential frame satisfies the following subset of the properties listed after definition 7, *provided* $\exists(\bullet|\mathcal{G})$ is defined for the relevant formulas.

1) $\varphi \subseteq \psi$ implies $\exists(\varphi|\mathcal{G}) \subseteq \exists(\psi|\mathcal{G})$ (partial monotonicity);
2) $\varphi \subseteq \exists(\varphi|\mathcal{G})$ (partial coarsening);
3) $\exists(\varphi \vee \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \vee \exists(\psi|\mathcal{G})$ (partial additivity);
4) $\exists(\varphi \wedge \psi|\mathcal{G}) = \exists(\varphi|\mathcal{G}) \wedge \psi$ where $\psi \in \mathcal{G}$ ('taking out what is known').

Let $\mathcal{G}$ be an evidential subframe of **Form**. $\exists(\mathbf{0}|\mathcal{G})$ will be defined if $\mathcal{G}$ is non-empty, but it could be different from $\mathbf{0}$, which may not be in $\mathcal{G}$.

In order to model the nonvalidity of the above inferences, it is elegant, although not strictly necessary, to import one more form of partiality into the conditional quantifiers. Suppose $\mathcal{M}_s$ represents Whitehead's approximation of the world, which contains two individuals: a_s and r_s, approximations at stage s of a and r, respectively. Suppose furthermore that we have only one predicate, W for 'wink'. The evidential frame $\mathcal{G}_s$ could then be taken to be generated by

$$\{\emptyset\} \cup \{\{f : \mathcal{M}_s, f_s \models W(x)\} : x \text{ a variable}\}.$$

Now suppose that, although both a and r may actually be winking, Whitehead is only in a position to see a winking. This can be modelled by switching to a different evidential frame: introduce a predicate W' such that $\mathcal{M}_s \models W'(a_s)$, $W'(r_s)$ undefined, and let $\mathcal{G}'_s$ be defined as $\mathcal{G}_s$, except that we use W' instead of W. If we now compute $\exists(W(x) \wedge x = r|\mathcal{G}'_s)$, we see that there is no element of $\mathcal{G}'_s$ lying above $\{f : \mathcal{M}_s, f_s \models W(x) \wedge x = r\}$, whence $\exists(W(x) \wedge x = r|\mathcal{G}'_s)$ is undefined, as desired.

It is important to observe that we need not actually change anything to the signature of $\mathcal{M}_s$; we may just take $\mathcal{G}'_s$ to consist of suitable subsets of elements of $\mathcal{G}_s$ and leave $\mathcal{M}_s$ unchanged. Hence, unlike in situation semantics, partiality is not introduced for the first order language; it resides in the conditionally quantified formulas. It would be possible to set up a four-valued partial logic for formulas in the language with conditional quantifiers, but we shall refrain from doing so here.[1]

The reader may have noticed that the frame $\mathcal{G}'_s$ is not a subframe of $\mathcal{G}_s$, since the interpretation of the predicate has been changed. So, whereas $\mathcal{G}_s$ derives from the algebra $\mathcal{B}_s$ determined by $\mathcal{M}_s$, $\mathcal{G}'_s$ has no immediate relation to $\mathcal{M}_s$. If we combine this with the results in the sections 3.3 and 3.4 we get the following picture. For suitable choices of frames, conditional quantifiers can recapture the approximating models $\mathcal{M}_s$; but conditional quantification is *more general* in the sense that it allows one to formalise several forms of partiality simultaneously, using different kinds of frames.

We conceive of the relationship between approximating models and filters in the following manner. The basic idea is that perception must be viewed as some form of filtering of reality $\mathcal{M}$. Part of the filtering consists of the inevitable blurring imposed upon us by our perceptual apparatus; this is formally captured by the inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ which has $\mathcal{M}$ as inverse limit. Additional filtering occurs because of restricted perceptual fields and the effect of perspective. This is not part of the perceptual apparatus, neither is it a property of reality; it arises as a consequence of 'being in the world', hence it is put in the filters only. Very roughly

[1] In the following we sometimes use Blamey's double-barrelled notion of consequence (Blamey 1983); partial consequence preserves truth from premise to conclusion and preserves falsity in the other direction.

speaking, the $\mathcal{M}_s$ are concerned with the possible structuring of experience, i.e., concepts, whereas the filters $\exists(\bullet|\mathcal{G}_s)$ relate to actual experience, i.e., perception.

Now we are in the position to study the formal semantics of direct perception reports in more detail.

2.0.4 Semantics

In chapter 2 we have seen examples indicating that someone's range of vision may come with different granularities; cf. (38). This would mean that a visual field consists of several approximations to possibly different submodels. In what follows, however, we often assume that each viewer comes with a unique range of vision; i.e., a single approximation to one partial submodel. We may therefore continue to interpret perception reports by means of formulas $\exists(\bullet|\mathcal{H})$, with φ a *positive* formula and $\mathcal{H}$ the frame assigned to perceiver j. The subtle issue of negation is discussed in section 3.0.11.

In chapter 2 we have showed that the semantics for perception reports of form 'a sees NP' is closely related to that of form 'a sees NP VP$_{ni}$'; all inferences for simple perception reports concerning objects have analogues within the domain of NI-reports. For this reason we shall mainly discuss inferences for this slightly more general setting. Here one could proceed by introducing a formal language, showing how its formulas are interpreted in terms of quantifiers $\exists(\bullet|\mathcal{G})$. For present purposes not much is gained by such precision, so we interpret the natural language examples directly into logic of vision; i.e., first order logic with quantifiers $\exists(\bullet|\mathcal{G})$ added (and inverse limits as models). To begin with, we restrict ourselves to possibly non-veridical uses; the veridical variants are discussed in section 3. The general strategy will be to show how the validity of inferences depend on the available perceptual resources, and compare this with informal linguistic judgments.

2.0.5 Retractability of perception reports

Recall the children's story on the apparent magic of dimes (i.e. (19) in section 1.1). In this story, one and the same object is described as it occurs in different perceptions. Its perceived properties vary, due to the change of granularity in Jack's visual field at the varying stages. Formally this corresponds to conditionally quantifying over assignments f in the inverse limit, such that at consecutive stages $s_1, s_2, \ldots, s_n$, $f(x)$ is perceived to be a dime: $f \in \exists(D(x)|\mathcal{G}^j_{s_1})$, a quarter: $f \in \exists(Q(x)|\mathcal{G}^j_{s_2})$, … a trash can: $f \in \exists(T(x)|\mathcal{G}^j_{s_n})$. This accounts for the consistency of (19) despite the fact that in reality D, Q and T are mutually incompatible.

2.0.6 Disjunction

What are the logical relationships between (12) and (13)?

(12) Jack saw Sharon smile or stare.

(13) Jack saw Sharon smile or Jack saw Sharon stare.

In chapter 2 we have already argued, following Grice, that the exclusive use of disjunction in (13) is best seen as a pragmatic effect, based on an inclusive interpretation. Due to the distributive law of frames, we have for inclusive disjunction:

$$\begin{array}{rl} & \exists(\varphi \vee \psi|\mathcal{H}) \\ = & \bigwedge\{C \in \mathcal{H} : \varphi \vee \psi \subseteq C\} \\ = & \bigwedge\{C' \in \mathcal{H} : \varphi \subseteq C'\} \vee \bigwedge\{C'' \in \mathcal{H} : \psi \subseteq C''\} \\ = & \exists(\varphi|\mathcal{H}) \vee \exists(\psi|\mathcal{H}) \end{array}$$

provided that $\exists(\bullet|\mathcal{H})$ exist for the relevant formulas. In section 3.0.11 it is shown how logic of vision obtains the exclusive use from pragmatic reasoning.

At this point we would like to emphasise that the logic of perception reports is *not* classical. Since evidential frames may lack a top element, a 'vacuous' statement such as $\exists(\varphi \vee \neg\varphi|\mathcal{H})$, which is equivalent to the positive $\exists(\top|\mathcal{H})$, could result in an undefined statement; *tertium non datur* is not part of the internal logic of 'to see', and similarly for other classical principles (cf. the earlier remarks on substitution).

2.0.7 Conjunction

Consider the following scenario. Jack walks about in his garden and sees an ant and a beetle crawl and fiddle. Since the insects fascinate him, he gets down on his knees to determine their species. First question: what are the species? Second question: relative to the perception operator, does the logic of conjunction alter as Jack's position changes? We have just seen that the answer to the second question is 'no' for inclusive disjunction, but conjunction is more sensitive in this respect; and logic of vision provides a flexible instrument to study the alternatives.

The general case In situation semantics, (14) and (15) are equivalent because conjunction has wide scope over the perception verb.

(14) Jack saw this ant crawl and fiddle.

(15) Jack saw this ant crawl and Jack saw this ant fiddle.

Indeed, (14) is true, iff Jack's scene verifies that this ant crawls and fiddles, iff, by definition, Jack's scene verifies that this ant crawls *and* Jack's scene verifies that this ant fiddles (cf. Barwise 1981, Kamp 1984, among others). Logic of vision predicts that (14) implies (15) in case of nontrivial perception. This is a consequence of partial monotonicity:

$$\frac{\varphi \subseteq \psi \quad \exists(\varphi|\mathcal{H})}{\exists(\psi|\mathcal{H})}$$

if $\exists(\varphi|\mathcal{H})$ and $\exists(\psi|\mathcal{H})$ exist. Therefore, conjunction satisfies the inequality (16) whenever the perception statements are defined.

(16) $\exists(\varphi \wedge \psi|\mathcal{H}) \subseteq \exists(\varphi|\mathcal{H}) \wedge \exists(\psi|\mathcal{H})$

This conforms to linguistic judgement; a linguist may even judge the sentences equivalent, especially if the inference is presented in isolation. However, the present logic suggests that this judgement could be based on background assumptions concerning perceptual information (i.e., the available frame) which are not always justified.

Logic of vision allows for poor perceptual resources; as a consequence there are counterexamples for the inference from (15) to (14). Suppose that at a certain point during Jack's walk he is unable to distinguish crawlers and fiddling things from each other, for Jack they are just moving objects. Under such circumstances, one could construct the following counterexample (cf. section 3.3, example 2, 'blurring of individuals'). Choose a model $\mathcal{M}$ in which:

i) there are crawlers and fiddling things;
ii) the crawlers and the fiddling things all move: $\forall y[C(y) \vee F(y), M(y)]$;
iii) no crawler fiddles: $\neg\exists y[C(y), F(y)]$ (equivalently: $C(y) \wedge F(y) = \mathbf{0}$);
iv) the object y that Jack perceives is a moving ant: $A(y) \wedge M(y)$.

Let Jack's visual information consist of the frame $\mathcal{H} := \{\mathbf{0}, M(y)\}$. This means that Jack cannot distinguish the crawlers and the fiddling things: $\exists(C(y)|\mathcal{H}) = M(y) = \exists(F(y)|\mathcal{H})$. Now (15), with 'this ant' wide scope, means (17).

(17) $A(y) \wedge \exists(C(y)|\mathcal{H}) \wedge \exists(F(y)|\mathcal{H})$

This is true, because $A(y) \wedge M(y)$ is. But (14) means (18).

(18) $A(y) \wedge \exists(C(y) \wedge F(y)|\mathcal{H})$

This is false because $\exists(C(y) \wedge F(y)|\mathcal{H}) = \exists(\mathbf{0}|\mathcal{H}) = \mathbf{0}$.[2]

Should the counterexample be considered an artifact of the logic, or can we make sense of it linguistically? A first reaction might be: it shows that conjunction always has wide scope, just as in situation semantics. But perhaps a more subtle judgment is possible, too, which would make the counterexample fit 'the facts'.

A speaker of (14) and (15) should know that the perception verb in an NI-report may influence the interpretation of the verbs within it scope. Hence, he will be careful in phrasing his report, for the conjunction of two verbs could have a different interpretation than the verbs taken separately, whence the perception verb may act differently on them as well. This is what happens in the counterexample: Jack has precise information on the empty property, but only approximate information on each of the verbs. We must ask the reader to judge whether this distinction is viable or not, but we think it is a virtue of logic of vision that it allows one to draw it.

Richer frames The counterexample may evoke yet another response: the verb 'to see' may have scope over conjunction, but its interpretation only makes sense for more realistic perceptual resources; that is, we should search for sensible constraints on frames, and study the logic in case they are satisfied.[3] We shall not consider such constraints here, but we do indicate how richer frames could make a difference.

A frame could be 'rich' in the sense that it has fully precise information for one

[2] There are slightly more involved counterexamples which do not use the artificial $\mathbf{0}$.

[3] The constraints could correspond, e.g, to the constraints put on the interpretation of 2-D images as 3-D scenes; Brady 1993 has a lucid introduction.

of the conjuncts. Then, (15) implies (14). For instance, $C(x) \in \mathcal{H}$ captures the fact that Jack accurately discerns crawlers, so

(19) $\exists(C(y)|\mathcal{H}) \wedge \exists(F(y)|\mathcal{H}) = \exists(C(y) \wedge F(y)|\mathcal{H})$

by 'taking out what is known'.

More interesting rich frames are discussed in section 3.3, where it is shown that conditional quantification may be equivalent to filtered quantification (given in terms of an equivalence relation). For instance, if Jack's frame $\mathcal{G}_s$ sustains the equivalence (20), $\mathcal{M}$ an inverse limit, he is somehow able to discern all individuals within its visual field (but not necessarily all their properties).[4]

(20) $\mathcal{M}, f \models \exists(\varphi|\mathcal{G}_s) \Leftrightarrow \exists g[f_s = g_s \;\&\; \mathcal{M}, g \models \varphi]$

Assume, e.g, that (21) and (22) are true of Jack's perceptual field. Given a rich $\mathcal{G}$, this field can then also be described by (23).

(21) Jack saw this ant carry a crumble.

(22) Jack saw this ant walk towards the ant hill.

(23) Jack saw this ant carry a crumble and walk towards the ant hill.

To show that in this case (21) and (22) imply (23), we reason as follows. Let (24) model the conjunction of (21) and (22).

(24) $\mathcal{M}, f \models \exists_s(A(x) \wedge Ca(x,y) \wedge Cr(y)) \wedge \exists_s(A(x) \wedge W(x,z) \wedge Ah(z))$

So 'this ant' has narrow scope and the indefinites have a specific reading (non-specific readings are possible as well). Suppose (24) is true. Then there are g, h with $g =_s f$ and $h =_s f$ such that

$$\mathcal{M}, g \models A(x) \wedge Ca(x,y) \wedge Cr(y) \;\&\; \mathcal{M}, h \models A(x) \wedge W(x,z) \wedge Ah(z).$$

Define f^* by: $f^*_s(y) := g_s(y)$, $f^*_s(z) := h_s(z)$, and $f^*(u) := f(u)$ elsewhere. Then, $f_s = f^*_s$, and $\mathcal{M}, f^* \models A(x) \wedge Ca(x,y) \wedge Cr(y) \wedge W(x,z) \wedge Ah(z)$. So (25) is true, as required.

(25) $\mathcal{M}, f \models \exists_s(A(x) \wedge Ca(x,y) \wedge Cr(y) \wedge W(x,z) \wedge Ah(z))$

Evaluating the above two examples in terms of constraints on frames, one might expect that realistic constraints will identify them as 'extremes' of some sort.

Different frames Let us return to Jack, when he was on his knees to look more closely at the ant and the beetle. In this position we could have reported the conjunction of (26) and (27) (cf. (38)).

(26) Jack saw this ant walk nearby.

(27) Jack saw that beetle walk at a distance.

We judged this conjunction unequivalent with (39)—here repeated as (28)—because Jack may not be able to focus on scenes at different distances at the same time.

(28) Jack saw this ant walk nearby and that beetle walk at a distance.

[4] Recall that $f_s = g_s$ means: $f_s(x) = g_s(x)$, for all variables x.

The example suggests another reason why the import of conjunction could fail; namely: in case of wide scope conjunction each conjunct may bring its own filter. To sketch the issues at stake, consider the 'poor' singleton frames $\mathcal{H}_a := \{\text{ant-nearby}(x)\}$ and $\mathcal{H}_b := \{\text{beetle-at-a-distance}(y)\}$. The first filter only has information on ants nearby, the second only on beetles at a distance; whence the conjunction of (26) and (27) could be true, for (29) is equivalent to (30).

(29) $\exists(\text{ant-nearby}(x)|\mathcal{H}_a) \wedge \exists(\text{beetle-at-a-distance}(y)|\mathcal{H}_b)$

(30) $\text{ant-nearby}(x) \wedge \text{beetle-at-a-distance}(y)$

On the other hand, the use of $\mathcal{H}_a$ with this conjunction within its scope would make (28) true for the wrong reason.

(31) $\exists(\text{ant-nearby}(x) \wedge \text{beetle-at-a-distance}(y)|\mathcal{H}_a)$

The filter eliminates the unseen part from the description, which results in

(32) $\text{ant-nearby}(x)$.

Instead, one would expect (28) to be undefined in this case; after all Jack does not see the beetle. For the frames $\mathcal{H}_a$ and $\mathcal{H}_b$ this means one must use (33).

(33) $\exists(\text{ant-nearby}(x)|\mathcal{H}_a) \wedge \exists(\text{beetle-at-a-distance}(y)|\mathcal{H}_a)$

The formula (33) is undefined because $\exists(B(y)|\mathcal{H}_a)$ is the top element $\bigwedge\emptyset \in \mathbf{Form}$, which is not in $\mathcal{H}_a$.

Perhaps this logical observation confirms the linguistic idea that conjunction should have wide scope over the perceptual verb. But apart from the fact that we doubt this to be true, it would surely be more interesting to arrive at this conclusion on the basis of more realistic frames.

This ends our discussion of conjunction. The next topic is quantificational inference, which begins with some options for the interpretation of quantifiers.

2.0.8 Quantifiers

A sentence such as 'Jacks sees $[_{\text{np}}\text{DET N}]$ VP_{ni}', with DET denoting a Lindström quantifier D, can be interpreted in at least three ways (narrow scope readings of D are discussed in section 3.0.10):[5]

T1 $Dx[\varphi, \exists(\psi|\mathcal{H})]$

T2 $Dx[\exists(\varphi|\mathcal{H}), \exists(\psi|\mathcal{H})]$

T3 $Dx[\exists(\varphi|\mathcal{H}), \exists(\varphi \wedge \psi|\mathcal{H})]$

[5] Let $\mathcal{M}$ be a model and $\mathcal{F}$ its space of assignments. Recall from section 3.3 that formulas φ are treated as subsets $[\![\varphi]\!]_{\mathcal{M}}$ of $\mathcal{F}$. So, to make sense of the formulas to follow, we have to 'lift' the determiner D as a relation between subsets of $|\mathcal{M}|$ to a map Dx, x a variable, from subsets of $\mathcal{F}$ onto a subset of $\mathcal{F}$. This can be done as follows:

$$f \in Dx[\mathbf{A}, \mathbf{B}] \text{ iff } D(\{d \in |\mathcal{M}| : f^x_d \in \mathbf{A}\}, \{d' \in |\mathcal{M}| : f^x_{d'} \in \mathbf{B}\})$$

for all $\mathbf{A}, \mathbf{B} \subseteq \mathcal{F}$. As a consequence

$$Dx[[\![\varphi]\!]_{\mathcal{M}}, [\![\psi]\!]_{\mathcal{M}}] = [\![Dx[\varphi, \psi]]\!]_{\mathcal{M}},$$

since: $\mathcal{M}, f^x_d \models \chi$ iff $f^x_d \in [\![\chi]\!]_{\mathcal{M}}$.

T2 and T3 restrict the domain of quantification by Jacks's field of vision—since the first argument is so restricted—while T1 does not. T3 incorporates a further restriction by repeating φ in D's second argument within the scope of the conditional quantifier.[6] Whence T2 and T3 are different, for by conservativity of D we respectively have (34) and (35).[7]

(34) $Dx[\exists(\varphi|\mathcal{H}), \exists(\varphi|\mathcal{H}) \wedge \exists(\psi|\mathcal{H})]$

(35) $Dx[\exists(\varphi|\mathcal{H}), \exists(\varphi|\mathcal{H}) \wedge \exists(\varphi \wedge \psi|\mathcal{H})]$

The previous discussion on conjunction shows that (34) and (35) are *not* equivalent, but that T3 and (35) *are* (again, provided $\exists(\bullet|\mathcal{G})$ defined). Does this mean that we have to opt for T3? Not necessarily, for the semantic differences between T2 and T3 are subtle. To see this, consider (36).

(36) Jack saw a dead man walk.

For simplicity, assume $\mathcal{G}_s$ is rich enough to make (37) and (38) equivalent (cf. section 3.3).

(37) $\mathcal{M}, f \models \exists(\varphi|\mathcal{G}_s)$

(38) $\exists g[E_s(f,g) \;\&\; \mathcal{M}, g \models \varphi]$

As before, E_s is the equivalence relation between assignments with: $E_s(f,g)$ iff $f_s = g_s$. So given an assignment f, (36) means: there is an object $f(x)$ which Jack's approximation cannot discern from: (T2) a material dead man and a material walker, or: (T3) a material dead man and a material dead man walking. Both T2 and T3 provide for a dead man walking *as far as Jack's field of vision is concerned*, but T3 has the supplementary 'transcendental' requirement that this is due to a material dead man walking.

Since we see no principled way to choose among T2 and T3, we opt for T3. It should be clear how the arguments to follow must be adapted for the simpler T2.

Consider the following pair of sentences, illustrating in- and exportation of a determiner DET:

(39) Jack saw DET children swim.

(40) DET children are such that Jack saw them swim.

Here, DET denotes a Lindström quantifier. In chapter 2 we have remarked that for DET equal to 'a', 'all', or 'no' (39) and (40) should be independent of each other, if one drops veridicality. On the other hand, the usual translation into first order logic, or even generalised quantifier logic, does give equivalence (cf. van der Does 1991). We shall now translate these sentences in the framework of conditional quantification, to see under what conditions the implications are predicted to hold.

[6] This trick is also used to embed generalised quantification into DRT (Kamp and Reyle 1993, ch. 4); we shall see shortly that it also helps logic of vision to get the correct interpretation of simple anaphoric elements.

[7] Conservativity is the constraint: $D(X,Y) \Leftrightarrow D(X, X \cap Y)$, satisfied by almost all natural language determiners.

We are interested in a reading which makes (39) different from (40). Thus, a translation by means of T1 is not what we are after, for then both (39) and (40) are interpreted as (41).

(41) $Dx[C(x), \exists(S(x)|\mathcal{H})]$

Instead, T3 reads (39) as (42), which may not be equivalent to (41).

(42) $Dx[\exists(C(x)|\mathcal{H}), \exists(C(x) \wedge S(x)|\mathcal{H})]$

As in case of conjunction, logic of vision does not give a single answer concerning the logical relationship between (41) and (42). Instead, it predicts how the (non-)inferences depend on the available perceptual resources. To show this, we follow Jack to a seaside resort, where he is a lifeguard in his spare time. At first Jack is sitting in his observation post, from which he can see all children and recognise them as such. After a short while a child gets in danger. So Jack runs to the seashore, where he does not see all children but does recognises the children that he sees, among which the unhappy one. A sudden wave almost drowns Jack, but his colleagues manage to rescue him (and the child). As soon as they have dragged Jack into a helicopter, Jack has all children within eyeshot again, but the water in his eyes precludes him to identify them precisely.

Let us go through this story once more to see what happens to Jack's logic of vision as he goes along.

Case (a) 'Jack recognises a child when he sees one, and he can see all children.' Formally this means: $C(x) \in \mathcal{H}$. Since $\exists(\bullet|\mathcal{H})$ is the identity on the elements of $\mathcal{H}$, $C(x) \in \mathcal{H}$ is a veridicality assumption for the first argument. Indeed, in this case (41) is equivalent with (42) for each determiner D:

$$\begin{aligned} & Dx[C(x), \exists(S(x)|\mathcal{H})] \\ = \; & Dx[C(x), C(x) \wedge \exists(S(x)|\mathcal{H})] \\ = \; & Dx[C(x), \exists(C(x) \wedge S(x)|\mathcal{H})] \\ = \; & Dx[\exists(C(x)|\mathcal{H}), \exists(C(x) \wedge S(x)|\mathcal{H})] \end{aligned}$$

The first step is by conservativity, the second step by 'taking in what is known', and the third step from $\exists(C(x)|\mathcal{H}) = C(x)$.

Case (b) 'Jack recognises a child when he sees one, but he might not see all children'; i.e., $C(j,x) \in \mathcal{H}$, where $C(j,x) \subseteq C(x)$ restricts the children to those in Jack's perceptual field (as in situation semantics). Let D be a $\uparrow$MON$\uparrow$ determiner (e.g., 'a' or 'at least n').[8] Then logic of vision predicts that (41) implies (42), but not conversely. To show the implication, first obtain

$$Dx[C(j,x), \exists(S(x)|\mathcal{H})] = Dx[\exists(C(j,x)|\mathcal{H}), \exists(C(j,x) \wedge S(x)|\mathcal{H})].$$

[8] Recall that

- $[\![D]\!]$ is MON$\uparrow$, iff: $[\![D]\!](X,Y)$ and $Y \subseteq Z$ implies $[\![D]\!](X,Z)$;
- $[\![D]\!]$ is MON$\downarrow$, iff: $[\![D]\!](X,Z)$ and $Y \subseteq Z$ implies $[\![D]\!](X,Y)$;

for all $X, Y, Z \subseteq E$. Left monotonicity is defined analogously.

as in case (a). By weak monotonicity and the fact that D is $\uparrow$MON$\uparrow$.

$$Dx[C(j,x), \exists(S(x)|\mathcal{H})] \subseteq Dx[\exists(C(x)|\mathcal{H}), \exists(C(x) \wedge S(x)|\mathcal{H})].$$

It is worthwhile to notice that $C(x)$ is imported into the scope of $\exists(\bullet|\mathcal{H})$ within D's second argument. This may seem surprising; for suppose Jack saw a girl instead of children. Then all C's in the above argument are changed to G, and (40) becomes (43) (similarly for (39)).

(43) A girl is such that Jack saw her swim.

But couldn't it be the case that there is a girl whom Jack perceives as swimming, without actually being aware that it is *she* who swims? No, because the anaphor *her* in (43) is taken to imply that the girl is imported in Jack's visual field; since he correctly identifies children, the conclusion follows.

At this point we should highlight that logic of vision offers different predictions than those of situation semantics. Situation semantics assumes veridicality at the lowest level: $\textsc{see}(j,R) \subseteq R$, for each relation R. Instead the filtered, possibly non-veridical properties used here, may become coarser: $R \subseteq \exists(R|\mathcal{H})$. This means, e.g., that in case of monotone determiners the inferences declared valid by the two semantics could be each other's opposite. Nevertheless, for non-veridical perception the predictions of the present logic accord with our informal semantic judgments; here as well as in the cases to follow. Further, in case of veridical perception the judgments of the logics are similar: both indicate that only positive formulas are veridical.

We now show that the converse fails for numerals $\mathbf{n}$ ($n \geq 1$) on the 'at least' and the 'exactly' use.[9] It may be assumed that $C(j,x)$ is properly contained in $C(x)$ (for otherwise the converse follows from (a)). Let Jack's perceptual resources be the frame $\mathcal{H}$ generated by $\{C(j,x), H(x), S(x)\}$ (H for: human). Assume

i) all children are human: $\forall x[C(x), H(x)]$;

ii) all swimmers are children: $\forall x[S(x), C(x)]$;

iii) there are n humans who swim: $\mathbf{n}\, x[H(x), S(x)]$;

iv) Jack sees a child: $\exists x.C(j,x)$;

v) no children seen by Jack swim: $\neg\exists x[C(j,x), S(x)]$.

Due to (ii–v), there is a child who does swim and there is one who does not, so: $\exists(C(x)|\mathcal{H}) = H(x)$. And due to (iii), $\exists(C(x) \wedge S(x)|\mathcal{H}) = H(x) \wedge S(x)$. Therefore, (42) is true:

$$\mathbf{n}\, x[\exists(C(x)|\mathcal{H}), \exists(C(x) \wedge S(x)|\mathcal{H})] \equiv \mathbf{n}\, x[H(x), S(x)].$$

[9] The argument can be adapted to all logical MON$\uparrow$ determiners on finite domains. Flum has observed that on finite domains MON$\uparrow$ determiners can be written as D^f with $f : \omega \longrightarrow \omega$ so that $f(n) \leq n$, and: $D^f_E(X,Y)$ iff $|X \cap Y| \geq f(|E|)$ (Flum 1985). We may assume that $\uparrow$MON$\uparrow$ D is non-trivial, for otherwise the converse does hold. So, there is E with $f(|E|) \geq 1$, on which a countermodel can be constructed as indicated.

But (41) is false: $\mathbf{n}\,x[C(j,x), \exists(S(x)|\mathcal{H})] \equiv \mathbf{n}\,x[C(f,x), S(x)]$, and no child seen by Jack swims.

What about determiners other than $\uparrow$MON$\uparrow$ numerals? Since $\downarrow$MON$\downarrow$ determiners, such as 'no' and 'at most n', are negations of $\uparrow$MON$\uparrow$ D, they sustain the converse (non)-implications valid for the $\uparrow$MON$\uparrow$ determiners. In a sense this argument is too quick, though, since it ignores the problematic aspects of negation; this issue is studied in section 3.0.11.

It is also immediate that for a $\downarrow$MON$\uparrow$ D such as 'every', we may infer (41) from (42):

$$Dx[\exists(C(x)|\mathcal{H}), \exists(C(x) \wedge S(x)|\mathcal{H})] \subseteq Dx[C(j,x), \exists(S(x)|\mathcal{H})],$$

for $C(j,x) \subseteq \exists(C(x)|\mathcal{H})$ and $\exists(C(x) \wedge S(x)|\mathcal{H}) \subseteq \exists(S(x)|\mathcal{H})$. We leave it to the reader to find a counterexample for the converse (cf. also footnote 9).

Case (c) 'Jack can see all children, but he cannot identify them precisely.' We show for $\uparrow$MON$\uparrow$ numerals that in this case (41) and (42) are independent of each other, as we think they should be.

For a start we show that one of Jack's colleagues in the helicopter cannot infer (42) from (41). Given Jack's poor shape there may be at least two children swimming within his eyesight, but Jack need not perceive such children. To show this formally, choose a model $\mathcal{M}$ and a frame $\mathcal{H}$ such that

i) 'There are n children': $\mathbf{n}\,x.C(y)$ and 'there are swimmers': $\exists y.S(y)$;
ii) 'No child swims': $\neg\exists[C(y), S(y)]$;
iii) 'Jack only discern humans and the empty property': $\mathcal{H} := \{H(y), \mathbf{0}\}$;
iv) 'All children and swimmers are within Jack's field of vision': $\forall y[C(y) \vee S(y), H(y)]$.

Nota bene, Jack's visual resources are too poor to discern children: $C(x) \notin \mathcal{H}$! Under these unfortunate circumstances: $\exists(S(y)|\mathcal{H}) = H(y)$, and hence (41) is true. But $C(y) \wedge S(y) = \mathbf{0}$, whence $\exists(C(y) \wedge S(y)|\mathcal{H}) = \mathbf{0}$, so that (42) is false.

The direction from (42) to (41) should be invalid, too, because what appears to be a child to Jack need not actually be one. This is a point where it is crucial to have evidential frames rather than just frames; when $D \equiv \exists$ the inference *is* valid for each $\mathcal{H}$ with $\mathbf{0} \in \mathcal{H}$. Suppose $\exists[C(y), \exists(S(y)|\mathcal{H})]$ is false. Since $S(y) \subseteq \exists(S(y)|\mathcal{H})$, $\exists[C(y), S(y)]$ is false as well. Whence $\exists(C(y) \wedge S(y)|\mathcal{H}) = \mathbf{0}$, so that $\exists y[\exists(C(y)|\mathcal{H}), \exists(C(y) \wedge S(y)|\mathcal{H})]$ is false as before.

The inference is invalid for evidential $\mathcal{H}$. Let Jack's perceptual frame be generated from the properties F, for: having fun, and W, for: woman, which include swimming and children, respectively. $\mathcal{H}$ is the frame $\{W(y) \wedge F(y), W(y), F(y), W(y) \vee F(y)\}$. Assume

i) there are n women having fun;
ii) there are children but none of them has fun; so $\exists(C(y)|\mathcal{H}) = W(y)$;
iii) no children swim (for it is fun); so $\exists(C(y) \wedge S(y)|\mathcal{H}) = W(y) \wedge F(y)$; and

iv) there is a non-female swimmer; so $\exists(S(y)|\mathcal{H}) = F(y)$.

All in all this makes (42) is true, because:

$$\mathbf{n}\, y[\exists(C(y)|\mathcal{H}), \exists(C(y) \wedge S(y)|\mathcal{H})] \equiv \mathbf{n}\, y[W(y), F(y)]$$

But (41) is false: $\mathbf{n}\, y[C(j,y), \exists(S(y)|\mathcal{H})] \equiv \mathbf{n}\, y[C(f,y), F(y)]$. The inference remains invalid when Jack is able to discern the swimmers precisely (use the same argument, but assume that 'swimming' and 'having fun' coincide: $\forall x[S(x) \leftrightarrow F(c)]$).

For 'every' the situation is of course different. First, the direction from (42) to (41) *does* follow, given the fact that it is $\downarrow$MON$\uparrow$. But the converse still fails, as the reader may wish to check for himself.

This section has shown that non-veridical perception reports already uses some of the intricacies of the logic of vision, but the situation becomes even more interesting in case of veridical perception.

3 Veridical perception reports

Recall that we took the expression 'I see an arm' to mean the conjunction of (i) and (ii).

i) 'with the present approximation the object that I focus on is identified as an arm',

ii) 'I expect this to be the case for every more refined approximation'.

The second condition is evidently non-monotonic: more precise information may contradict the expectation expressed in (ii). As such, (ii) does not yet express that our perception will be veridical: for this we would need a result which says that if φ is true in every approximation, then φ is true ('in reality', i.e., on the inverse limit). Combined with such a result, (ii) yields *veridicality*.

The purpose of this section is to give a precise formulation to veridicality conceived of as a defeasible rule. We first discuss the standard format for defaults, and then present a slightly deviant form, more suitable for applications to perception. This is then applied to study veridicality inferences.[10]

3.0.9 Pragmatic inference from default rules

In Reiter's version of default logic (Reiter 1980), a *default* is a rule of the form

$$\alpha : \beta_1, \ldots, \beta_n / \omega$$

where α is the *prerequisite* of the rule, $\beta_1, \ldots, \beta_n$ are its *justifications*, and ω is its *consequent*. The customary interpretation of the rule is: 'if α has been derived and $\beta_1, \ldots, \beta_n$ are consistent with what has been derived, conclude ω'.

A *normal* default is one in which there is a single justification which is identical to the consequent; this is the kind of default of interest to us. Normally, defaults

[10] The interpretation of default rules owes much to conversations with Frans Voorbraak.

are used to express rules with exceptions, such as 'Birds fly', formalised as

$$B(a) : F(a)/F(a),$$

for every constant a. A default theory consists of a set of facts and a set of default rules. The facts ('Tweety is a bird') are taken to be specific and reliable information, and the defaults represent general information. In our case the situation is slightly different: the specific information consists of perceptual judgments, which are approximate, hence defeasible. The default rules should express the expectation that the judgment will continue to be true in more refined approximations. In this slightly different situation, the intended interpretation of defaults also undergoes a subtle change.

For the following discussion we assume given an inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ with inverse limit $\mathcal{M}$ and a corresponding system of quantifiers $\exists(\bullet|\mathcal{G}_s)$ (cf. section 3.3.0.2). If 's' stands for the present approximation, the statement (i)

> with the present approximation, x, the object that I focus on, is identified as an arm

is formalised as $\exists(A(x)|\mathcal{G}_s)$. In Reiter's format, statement (ii) would then be expressed by the default rule

$$\exists(A(x)|\mathcal{G}_s) : \exists(A(x)|\mathcal{G}_t)/\exists(A(x)|\mathcal{G}_t) \text{ (for } t > s),$$

which says that if evidence at stage s supports $A(x)$, and if it is still consistent to assume $A(x)$ at stage t, then assume $A(x)$ at stage t.

However, a different, 'evidential', interpretation better fits the perceptual situation. The discussion of Marr's theory of object recognition in 3.5 strongly suggests that the stage s, given by (i), should represent the *maximal available information*: we move downward in the hierarchy of 3-D models until the finite resolution of the image leads to a branching. But in that case we cannot have any evidence against $\exists(A(x)|\mathcal{G}_t)$ at stage s. For simplicity assume that $\mathbf{0}$ is an element of the frame $\mathcal{G}_s$. Suppose the evidence at stage s is summarised by $\exists(\varphi|\mathcal{G}_s) \neq \mathbf{0}$, so that we have $\exists(\varphi|\mathcal{G}_s) \subseteq \exists(A(x)|\mathcal{G}_s)$. Suppose furthermore that $\exists(\varphi|\mathcal{G}_s) \cap \exists(A(x)|\mathcal{G}_t) = \mathbf{0}$. Then, $\exists(\exists(\varphi|\mathcal{G}_s) \cap \exists(A(x)|\mathcal{G}_t)|\mathcal{G}_s) = \mathbf{0}$; so $\exists(\varphi|\mathcal{G}_s) \cap \exists(\exists(A(x)|\mathcal{G}_t)|\mathcal{G}_s) = \mathbf{0}$; whence $\exists(\varphi|\mathcal{G}_s) \cap \exists(A(x)|\mathcal{G}_s) = \mathbf{0}$, a contradiction.

Hence we shall take a default to be a rule of the form

$$\exists(\varphi|\mathcal{G}_s)/\exists(\varphi|\mathcal{G}_t)$$

with φ a *positive* formula. This rule should be interpreted as: 'if I have observed φ at stage s, and s represents the maximum available accuracy, then I may assume that I will observe φ at stage t.[11]

Having introduced the expectation inherent in every perceptual experience in

[11] The difference between Reiter's interpretation and ours is that we take the consistency of the justification to be relative to a stage s, whereas in Reiter's case it refers to an extension of the default theory. Here, we shall forego a discussion of the possible notions of extensions applicable in this context; see Voorbraak 1997.

the form of a default rule, we return to the principle of veridicality, which would sanction inferences of the form

$$\frac{\text{I see an arm}}{\text{There is an arm}}$$

or formally

$$\frac{\exists x \exists (A(x)|\mathcal{G}_s)}{\exists x A(x)}$$

To see the connection between the default rules and veridicality, assume that we have an assignment f such that for all t, $f \in \exists(A(x)|\mathcal{G}_t)$; does it then follow that $\mathcal{M}, f \models A(x)$? For positive formulas, the answer is 'yes':[12]

Theorem 1 *Assume the bonding mappings are surjective. For* positive φ*,*

$$\bigcap_{s \in T} \exists(\varphi|\mathcal{G}_s) = \varphi.$$

This result fails already for negations of atomic formulas! 2

A related result, with 'positive' replaced by 'positive primitive', holds if the bonding mappings are not necessarily surjective. We conjecture that the identity $\bigcap_{s \in T} \exists(\varphi|\mathcal{G}_s) = \varphi$ *implies* that φ is positive, but we have not been able to prove this. A compactness argument now gives

Corollary 2 *Assume the bonding mappings are surjective and let φ be positive. If for all $s \in T$, $\mathcal{M} \models \exists x_1 \ldots \exists x_n \exists(\varphi|\mathcal{G}_s)$, then $\mathcal{M} \models \exists x_1 \ldots \exists x_n \varphi$.* 2

We go on to show how these insights can be utilised in the study of veridical perception reports.

3.0.10 Veridical perception reports

Veridicality as a defeasible pragmatic inference is a subtle issue, as should become clear from the discussion of numerals and negation.[13]

Monotone increasing numerals In section 2.0.8 some inferences with numerals failed because the percepts of Jack as a lifeguard were treated as possibly non-veridical. But what if his percepts are 'rounded off' by adding the expectation that they are stable? In particular, if Jack saw at least two chiddren swim, and his identification of children is presumed to remain correct ($\exists(C(x)|G_t) \subseteq \exists(C(x)|G_s)$

[12] The proof essentially uses the fact that positive formulas determine closed sets of assignments; this was established in section 3.6.

[13] The import of conjunction into the scope of the perception verb can also be formulated as a defeasible principle:

$$\frac{\exists(\varphi|\mathcal{G}_s) \wedge \exists(\psi|\mathcal{G}_s)}{\exists(\varphi \wedge \psi|\mathcal{G}_s)} \ \textsc{conj}$$

If we are justified in expecting that each (positive) conjunct is stable under refinement of information, then conjunction can be imported into the scope of 'to see'. This is immediate from Theorem 1.

for $t \geq s$ by defeasible assumption), does it then follow that at least two children swam? It depends. To show what is at stake we consider the numeral 'two'.[14]

The familiar first order definition of 'two' contains a negative element; it states: there are objects *un*equal to each other such that.... This negative element blocks veridicality for the definition, since Theorem 1 cannot apply. However, it is defensible that $x \neq y$ is *inferred* from a irreflexive relation D (for: different). If this view is taken seriously, we get veridicality for $\uparrow$MON$\uparrow$ numerals in the following sense (we state $n = 2$).[15]

$$\frac{\exists x \exists y[\exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y)|\mathcal{G}_s)]}{\exists x \exists y[x \neq y \wedge \varphi(x) \wedge \varphi(y)]} \uparrow\text{NUM}\uparrow$$

Assume $\mathcal{M}, f \models \exists x \exists y[\exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y)|\mathcal{G}_s)]$. Then there is $g =_{x,y} f$ such that $\mathcal{M}, g \models \exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y)|\mathcal{G}_s)$. Choose such a g. Veridicality gives: $\mathcal{M}, g \models \exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y)|\mathcal{G}_t)$, for all stages t. Whence, for all stages t: $\mathcal{M}, f \models \exists x \exists y \exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y)|\mathcal{G}_t)$. By Corollary 2: $\mathcal{M}, f \models \exists x \exists y[D(x,y) \wedge \varphi(x) \wedge \varphi(y)]$. Further, the fact that D is irreflexive is preserved: since $\forall x[Dxx, \perp]$ is true in all stages, it is true in $\mathcal{M}$ as well. We conclude: $\mathcal{M}, f \models \exists x \exists y[x \neq y \wedge \varphi(x) \wedge \varphi(y)]$, as required.

Observe that the same argument works for numerals within the scope of 'see'. Since $\exists(\bullet|\mathcal{G}_s)$ is only defined for positive formulas, these numerals should be definable within positive first order logic (perhaps with irreflexive R added).[16] An example of such a numeral is in (44) formalised by (45).

(44) Jack saw Sharon phone two lifeguards.

(45) $\exists(x = s \wedge \exists y, z[D(y,z) \wedge L(x) \wedge L(y) \wedge P(x,y) \wedge P(x,z)]|\mathcal{G}_s)$

In the descriptive part of (45) all relations occur positive; hence they are monotone increasing. This means that after deriving $x = s \wedge \exists y, z[D(y,z) \wedge L(x) \wedge L(y) \wedge P(x,y) \wedge P(x,z)]$ using veridicality, the irreflexive D can be replaced by $\neq$, as required. By Lyndon's theorem, which states that a predicate is monotone increasing iff it occurs positive, the same is true for all positive formulas.

A check of the above argument reveals moreover that if one carefully discerns between the first and the second argument of the numerals—φ and ψ below—veridicality can be used to derive (47) from (46) (cf. the counterexamples in case (c) above).

(46) $\exists x \exists y \exists(D(x,y) \wedge \varphi(x) \wedge \varphi(y) \wedge \psi(x) \wedge \psi(y)|\mathcal{G}_s)$

(47) $\exists x \exists y[x \neq y \wedge \varphi(x) \wedge \varphi(y) \wedge \exists(\psi(x) \wedge \psi(y)|\mathcal{G}_s)]$

[14] We now use the first order defitions of this numeral, not the generalised quantifier version. Note the subtle difference between $\mathbf{2}\, x\exists(\varphi \wedge \psi|\mathcal{H})$ and $\exists x, y[x \neq y \wedge \exists(\varphi(x) \wedge \psi(x)|\mathcal{G}_s) \wedge \exists(\varphi(y) \wedge \psi(y)|\mathcal{G}_s)$, because $\exists(\bullet|\mathcal{G}_s)$ is sensitive to the variables chosen!

[15] Due to the defeasible rule CONJ stated in footnote 13, the relative scope of $\wedge$ and $\exists(\bullet|\mathcal{G}_s)$ in the premise is immaterial.

[16] Question: which determiners have this property? Only the positive Boolean combinations of quantifiers 'at least n'? Another option would be to see whether the result for positive formulas can be extended to include $\uparrow$MON$\uparrow$ determiners.

Thus we obtain an analogue of the informal idea that export of quantifiers is related to veridicality.

Monotone decreasing numerals Recall that in situation semantics veridicality was 'built in' at the atomic level: $\textsc{see}(j, R) \subseteq R$, while logic of vision often allows coarsening of a predicate: $R \subseteq \exists(R|\mathcal{G}_s)$. This difference may show up in the predictions concerning veridicality. In particular, numeric NI-complements with $\downarrow$MON$\downarrow$ numerals, such as 'at most n', are not veridical in situation semantics; but depending on the available resources they may be veridical here. For example, (48) is formalised as (49). So if (48) is true, and hence $\exists(P(x) \wedge P(y)|\mathcal{G}_s)$ in (49) is defined, then (50) meaning (51) is true as well, by $\downarrow$MON$\uparrow$ of $\forall$.

(48) Lonely Jack saw at most one pebble.

(49) $\forall x \forall y[\exists(P(x) \wedge P(y)|\mathcal{G}_s), x = y]$

(50) There is at most one pebble.

(51) $\forall x \forall y[P(x) \wedge P(y), x = y]$

But isn't this result worrying? No, it just indicates that in this case Jack has accurate perceptual information on *all* pebbles; whence the result. As soon as his resources get poorer in this respect, the inference may fail. For example, if we restrict the pebbles to those seen by Jack ($P(j, x) \subseteq P(x)$, as in situation semantics), the inference no longer holds. It is also blocked if Jack lacks information about pebbles: then the conclusion (51) could be false, and the premise (49) undefined.

We must leave it to the reader to ponder over non-monotonic numerals (e.g., 'just ten').

3.0.11 Negation

In perception reports, negation may occur in several forms, sometimes with veridical import, sometimes without. Here we give a formal discussion of the examples introduced in chapter 2, and we add some more.

First a formal point. Thus far we have only considered quantifiers $\exists(\bullet|\mathcal{G})$ applied to positive formulas. This will remain the case; for each R there may be several antonymic relations $\sim R$ disjoint from R, which are positive approximations of R's complement. As we shall see shortly, the approximations may vary for several reasons; that is why we allow more than one antonym of a relation. The only 'negations' which may occur inside $\exists(\bullet|\mathcal{G})$ are $\sim R$. In this way, all formulas of interest remain positive, and the machinery developed so far applies.

Since Jack's talents as a birdwatcher are limited, (52) should not imply (53).

(52) Jack saw no hawk fly.

does not imply

(53) No hawk flew.

Formally, this can be seen by adapting the second argument under Case (c) for quantifiers, in section 2. But it is also instructive to see why this is so at an informal level.

Let $\exists(\bullet|\mathcal{H})$ be the conditional quantifier associated with Jack's perceptual field, incorporating both a filter and a domain restriction. Now (52) may be formalised by (54)

(54) $\neg\exists x[\exists(H(x)|\mathcal{H}), \exists(H(x) \wedge F(x)|\mathcal{H})]$,

if we read 'hawk' as being in the scope of 'saw'. On the other hand, (53) is rendered formally as (55)

(55) $\neg\exists x[H(x), F(x)]$.

To derive (53) from (52), one could try to argue as follows: assume $H(x)$ and $F(x)$, then we have $\exists(H(x)|\mathcal{H})$ and $\exists(F(x)|\mathcal{H})$, whence also: $\exists x[H(x), F(x)]$ implies $\exists x[\exists(H(x)|\mathcal{H}), \exists(F(x)|\mathcal{H})]$. However, the first step fails because $\exists(\bullet|\mathcal{H})$ is conditional on an evidential frame, whence not necessarily, say, $F(x) \subseteq \exists(F(x)|\mathcal{H})$, because the latter need not exist. A similar argument applies if 'hawk' is taken to be outside the scope of 'saw', i.e., when (52) is represented by

(56) $\neg\exists x[H(x), \exists(F(x)|\mathcal{H})]$,

(similarly if $H(x)$ is restricted by Jack's field of vision); cf. case (c) above for formal details.

Sometimes, however, a perception report involving negation does carry with it the implication of veridicality. This appears to be the case when the negative statement is actually derived from positive information. Consider Jack's musing:

(57) I see a bird. It is not a hawk.

Here one is tempted to infer non-monotonically that the perceived bird is not a hawk. The reason is that we actually perceive the bird to possess a property which is an antonym ($\sim H$) of being a hawk (H): $\sim H$ can be taken to be a collection of birds which are either black, or yellow, or 10 cms long, or...

To formalise this example, one could proceed as follows: we treat the anaphor 'it' in the manner of DPL, so that the last occurrence of x is bound by the quantifier $\exists x$:

(58) $\exists x \exists(B(x)|\mathcal{G}) \wedge \exists(\sim H(x)|\mathcal{G})$;

now apply veridicality to the second conjunct to obtain $\sim H(x)$. This conclusion may fail for two reasons:

i) one may have $\exists(\sim H(x)|\mathcal{G}) \cap \exists(H(x)|\mathcal{G}) \neq \emptyset$.

ii) $\sim H$ was taken too large.

A good example of situation (ii) is furnished by a type of buzzard, *buteo rufinus*, whose colour shows two phases: a common light phase and a rare chocolate brown, almost black, phase. Not knowing that the latter phase exists, i.e., taking $\sim H$ too large, may easily lead to misidentification.

This way of reasoning can also be applied to the dialogue of Sharon and Jack, from chapter 2.

S: 'Did you see that hawk there?'

J: 'I saw *something*, but it was not a hawk...
No, you are right, it is a hawk after all.'

This example can be analysed either by assuming that the conclusion '... it was not a hawk' is based on positive information P implying $\sim H$ such that (unbekownst to us) $\exists(P(x)|\mathcal{G}) \cap \exists(H(x)|\mathcal{G}) \neq \emptyset$, or by arguing that the positive predicate P includes too many features.

For another use of antonyms, recall sentences (59) and (60) from chapter 2:

(59) Jack saw this hawk not fly.

(60) Jack didn't see this hawk fly.

Once more, 'not fly' is interpreted as one of the antonyms of 'fly'. It was noted that one normally takes (59) to imply (60), but not conversely. The reason that the converse implication fails is presumably that Jack may not be in a position to see the hawk fly or not fly, e.g., he may be looking at the ground. Can we reproduce these intuitions? Formally, these sentences would be rendered thus:

(61) $\exists(H(x) \wedge \sim F(x)|\mathcal{G})$

(62) $\neg\exists(H(x) \wedge F(x)|\mathcal{G})$.

Now in principle $\exists(H(x) \wedge \sim F(x)|\mathcal{G})$ and $\exists(H(x) \wedge F(x)|\mathcal{G})$ may overlap, so that neither of (61) and (62) implies the other. However, we almost get that $\exists(H(x) \wedge \sim F(x)|\mathcal{G})$ implies $\neg\exists(H(x) \wedge F(x)|\mathcal{G})$, in the following sense. Being positive, both formulas $H(x) \wedge \sim F(x)$ and $H(x) \wedge F(x)$ determine closed sets of assignments on $\mathcal{M}$; these sets are disjoint. Consider the set of quantifiers $\exists(\bullet|\mathcal{G}_s)$; by theorem 1, $\exists(H(x) \wedge \sim F(x)|\mathcal{G}_s)$ converges to $H(x) \wedge \sim F(x)$, and $\exists(H(x) \wedge F(x)|\mathcal{G}_s)$ converges to $H(x) \wedge F(x)$. By compactness, there will be a stage s such that $\exists(H(x) \wedge \sim F(x)|\mathcal{G}_s)$ and $\exists(H(x) \wedge F(x)|\mathcal{G}_s)$ are disjoint. For such s we have indeed that $\exists(H(x) \wedge \sim F(x)|\mathcal{G})$ implies $\neg\exists(H(x) \wedge F(x)|\mathcal{G})$.

Mutatis mutandis the same reasoning can be used to establish the defeasible rule

$$\frac{\exists(\varphi \vee \psi|\mathcal{G}_s)}{\exists t > s : \exists(\varphi|\mathcal{G}_t) \uplus \exists(\psi|\mathcal{G}_t)} \text{ EX-DISJ}$$

with $\uplus$ denoting exclusive disjunction. The rule EX-DISJ says that an inclusive disjunction will become exclusive from a certain stage onwards. Again, this is defeasible rule is a consequence of veridicality.

3.1 Concluding remark

At first sight it might seem that the logic of vision is rather weak, since it declares so many inferences invalid. We view the matter differently: in general the inferences are invalid, but given suitable pragmatic constraints they become valid; the logic of vision provides us with a precise instrument to isolate the necessary additional assumptions. This is of course a general feature of resource bounded logics: the underlying logics are weak, but stronger principles can be obtained by strengthening the resource.

This chapter ends with a formal comparison of partial objects in logic of vision with the so-called pegs of Landman 1986.

4 Partial objects

According to Landman (1986, 97) the main problem in semantics is to explain the intersubjective character of language use: how can it be that we often assume to speak about the same objects—not just about our private experiences,—while disregarding the fact that some of the objects we take to be different might actually collapse into one object? To shed light on this puzzling aspect of our semantic milieu, it is crucial to come to grips with partial objects (or 'variables' in the traditional sense). Here, we study some notions of partial object as they can be found in logic of vision. This prepares the ground for studying pegs within this framework, which gives yet another notion.

In what follows, we fix a total inverse system $\langle \mathcal{M}_s, h_{st} \rangle_{s,t \in T}$ that is proper refining with respect to all atomic sentences; $\mathcal{M}$ is the inverse limit of this system. Accordingly we have (i) for each $d \in |\mathcal{M}_s|$ there is $\xi \in \mathcal{M}$ such that $\xi_s = d$ (by totality); and (ii) for each conjunction of atomic formulas α: $\mathbf{p}_s^{-1}(\alpha) = \exists(\alpha|\mathcal{G}_s)$, with $\mathcal{G}_s$ the frame generated from the lattice of positive formulas $\mathcal{B}_s$ (by proper refinement). Cf. section 3.4.

4.1 Some kinds of partial object in logic of vision

In logic of vision, the threads $\xi \in \mathcal{M}$ are total objects, but there are several ways in which an object can be partial. We name a few alternatives, together with their partial ordering:

a) sets of threads (called: *blurred* objects), ordered by inclusion;
b) the proper tails of a thread, ordered 'by being a tail of';
c) the elements in $\bigcup_{t \in T} |\mathcal{M}_s|$, ordered by: $a \leq b$, iff $h_{st}(b) = a$ for some h_{st}.

In chapter 3 we have mainly worked with a particular kind of blurred object. It will appear that a version of this notion is indeed well suited to logic of vision.

The tails in (b) and the elements in (c) determine the same partial objects. By totality of the refining system, we may think of each $d \in \mathcal{M}_s$ as ξ_s for some thread $\xi \in \mathcal{M}$, and d fixes the tail of all threads running through it; $\xi_s = \xi'_s$ iff (by the definition of thread) for all $t \leq s$: $\xi_t = \xi'_t$. Conversely, each proper tail τ begins in a certain $s \in T$; so τ can be identified with the element τ_s. Thus one sees that the map from tails to elements is a bijection preserving the partial order 'is a tail of': the elements and the tails are interchangeable. This is not true for elements–hence: tails—and blurred objects.

Each $\xi_s \in |\mathcal{M}_s|$ determines the blurred object $\mathrm{bo}^\xi(s)$ with

$$\mathrm{bo}^\xi(s) := \{\xi' \in |\mathcal{M}| : \xi_s = \xi'_s\} = \{\xi' \in |\mathcal{M}| : \forall t \leq s(\xi'_t = \xi_t)\}.$$

This map preserves the partial order of blurred objects, as follows: if $s \geq t$ then $\mathrm{bo}^\xi(s) \subseteq \mathrm{bo}^\xi(t)$. For let $\xi' \in \mathrm{bo}^\xi(s)$ and $t \geq s$. Then: $\xi'_t = h_{st}(\xi'_s) = h_{st}(\xi_s) = \xi_t$. So $\xi' \in \mathrm{bo}^\xi(t)$. As to the converse, in general the map from elements to blurred objects is not an isomorphism. A blurred object could be any set of threads, perhaps totally unrelated to each other; in particular the tails of its threads may differ.

For what follows, let us recast the above observations in terms of the assignment space $\mathcal{F}$. Assume the thread $\xi \in |\mathcal{M}|$ to be named by the constant $\underline{\xi}$,[17] and choose an $x \in \mathrm{VAR}$. The notion corresponding to $\mathrm{bo}^\xi(s) \subseteq |\mathcal{M}|$ is $\exists(x = \underline{\xi}|\mathcal{G}_s) \subseteq \mathcal{F}$.

$$\exists(x = \underline{\xi}|\mathcal{G}_s) = \mathbf{p}_s^{-1}(x = \underline{\xi}) = \{f \in \mathcal{F} : f_s(x) = \xi_s\}.$$

The relation between $\mathrm{bo}^\xi(\bullet)$ and $\exists(x = \underline{\xi}|\mathcal{G}_\bullet)$ is given by

$$f(x) \in \mathrm{bo}^\xi(t) \Leftrightarrow f \in \exists(x = \underline{\xi}|\mathcal{G}_t)$$

for all $t \in T$. Thus: $\mathrm{bo}^\xi(s) \subseteq \mathrm{bo}^\xi(t)$ iff $\exists(x = \underline{\xi}|\mathcal{G}_s) \subseteq \exists(x = \underline{\xi}|\mathcal{G}_t)$, which means that there is no principal difference between using $\mathrm{bo}^\xi(\bullet)$, or using $\exists(x = \underline{\xi}|\mathcal{G}_\bullet)$ for a certain x.

4.2 Pegs in logics of vision

Until now we have considered partial objects without paying attention to the facts holding true of them. The notion of a peg, developed by Landman within data semantics,[18] is based on the alternative view, which holds that such 'bare' partial objects are best seen as derived from more basic 'clothed' ones.

Pegs are objects that may vary their guises—the facts true of them—with the information state in which they occur. Three intuitions concerning pegs are prominent. Firstly, they are partial objects, in that they may leave undecided whether or not a fact is true of it. Secondly, relative to information states pegs may approximate each other; in particular, they may approximate total, fully determined pegs. Thirdly, due to the partiality of their guises, they may be indiscernible from each other. It is assumed moreover that pegs are consistent: in each information state the propositions true of it are compatible.

What kind of formal object would comply with this circumscription? In what follows pegs are modelled by functions from information states to consistent sets of propositions.[19] The propositions are thought of as partially describing the total objects a peg may develop into if more information becomes available.

Landman introduces pegs within a fairly abstract set up, but for present purposes a concrete set-theoretic version of his ideas suffice. To do justice to all aspects of pegs one should use a partial set theory. We refrain from doing so here, but it

[17] It is understood that $[\![\underline{\xi}]\!]_s = \xi_s$.

[18] Data semantics is first introduced in Veltman 1981.

[19] In Landman 1986, 113–132, pegs are not defined formally, but we take it that the definition given stays close to his intentions.

should be clear how this feature could be added. As we go along, it will become apparent that pegs in logic of vision have several interesting properties, which are absent in the general theory.

Let T be a set of information states, and P a set of propositions in a so-called information system (Landman 1986, 115; see also below).

Definition 3 A (concrete) *peg* is a function from information states to consistent sets of propositions:

i) $o : T \longrightarrow \wp\wp(P)$;

ii) $o(t)$ is a proper filter, for all $t \in T$.[20]

That $o(t)$ is a proper filter corresponds to the idea that peg o should appear consistent in information state t. Observe, though, that this requirement gives *finite* consistency; for infinite $o(t)$ it could still be the case that $\bigcap o(t) = \emptyset$.[21]

When using pegs to model directly perceived partial objects, one may wonder whether consistency formulated in terms of filters is appropriate. For example, a directly perceived fact may imply a fact which cannot be observed, but filters in a sufficiently rich information state *are* closed under such implications.[22] Physics teaches us that each human is a near empty cloud of elementary particles; yet, there is a sense in which the following inference is invalid

$$\frac{\text{Jack saw Sharon}}{\text{Jack saw a near empty cloud of particles}}$$

Be this as it may, it could still be so that directly perceived objects are filters once restricted to observable facts. To decide whether or not this is so is a subtle matter. Is it possible to directly observe green without directly observing colouredness? Can we see a shade without seeing lack of light? Fortunately we may sidestep the issue, since we shall see shortly that there is a variant of pegs for which the problem does not arise.

Pegs come naturally in logic of vision. Each stage $\mathcal{M}_s$ induces an information system $\langle P_s, \wedge, \perp, \sqsubseteq \rangle$, as follows:

i) $P_s := \{\exists(\varphi|\mathcal{G}_s) : \varphi$ a conjunction of atomic formulas$\}$;

ii) $\wedge$ is intersection; $\perp$ is the empty set, and $\sqsubseteq$ is inclusion.

Recall that for $\exists(\varphi|\mathcal{G}_s) \in P_s$: $\exists(\varphi|\mathcal{G}_s) = \mathbf{p}_s^{-1}(\varphi) = \{f \in \mathcal{F} : \mathcal{M}_s, f_s \models \varphi\}$, by Theorem 8. So conjunction is well-defined: $\exists(\varphi|\mathcal{G}_s) \cap \exists(\psi|\mathcal{G}_s) = \mathbf{p}_s^{-1}(\varphi) \cap \mathbf{p}_s^{-1}(\psi) = \mathbf{p}_s^{-1}(\varphi \wedge \psi) = \exists(\varphi \wedge \psi|\mathcal{G}_s) \in P_s$ (this is why we use conjunctions of atomic formulas).

[20] A proper filter F is a set of sets such that: $\emptyset \notin F$, and: $X, Y \in F$ iff $X \cap Y \in F$.

[21] This is one of the instances where the observations in this section are only interesting for the infinite case.

[22] Filters are monotone: $X \in F$ and $X \subseteq Y$ implies $Y \in F$. Indeed, the issue is related to the problem of omniscience for propositional attitudes.

Accordingly, definition 3 makes a peg a function that maps each $t \in T$ onto a proper filter $o(t) \subseteq \wp(P_s)$.

In logic of vision pegs are not just finitely consistent, they are consistent *simpliciter*.

Proposition 3 *Let o be a peg. For each $t \in T$: $\bigcap o(t) \neq \emptyset$,*

PROOF. Each $X \in o(t)$ can be written as $\exists(\varphi|\mathcal{G}_s)$, with φ a conjunction of atomic formulas. Thus we know from section 3.6 that each $X \in o(t)$ is a closed set in $\mathcal{F}$. Since $o(t)$ is a proper filter, it has the finite intersection property. By compactness, $\bigcap o(t) \neq \emptyset$. 2

Proposition 3 connects pegs with the kind of blurred objects that have been used in the previous sections. It also shows that the present framework disregards pegs whose consistency in an information state can only be detected by 'infinite means'. This is a reasonable assumption. It even suggests to redefine pegs so that $\bigcap o(t) \neq \emptyset$, for each $t \in T$ (i.e., a change of (ii) in definition 3). In this way one circumvents the discussion whether a peg in an information state should be a filter or not.

Intuitively, we think of $o(t)$ as the facts true of object o in t. To make this precise one has to require there be a variable x occurring free in the defining conjunction of the $X \in o(t)$. For this reason we write o^x from now on.

Which total objects does $o^x(t)$ approximate? One answer is: it approximates each thread in the set $\{f(x) : f \in \bigcap o^x(t)\}$. Since $\exists(x = \underline{f(x)}|\mathcal{G}_t)$ gives an approximation of $f(x)$ in t, it natural to ask: can the set $\{f(x) : f \in \bigcap o^x(t)\}$ be obtained as the union of these blurred objects? The answer to this question shows once more that logic of vision is sensitive to variables. Let $\mathrm{FV}(o(t))$ be the union of all free variables in the definitions of the elements $X \in o(t)$. It is not difficult to check that

(*) $\bigcap o^x(t) = \bigcup_{f \in \bigcap o(t)} \bigcap_{x \in \mathrm{FV}(o(t))} \exists(x = \underline{f(x)}|\mathcal{G}_t)$

Without information about all free variables of $o(t)$, one only has that $o(t)$ *sharpens* the disjunction of the partial descriptions $\exists(x = \underline{f(x)}|\mathcal{G}_t)$:

$$\bigcap o^x(t) \subseteq \bigcup_{f \in \bigcap o^x(t)} \exists(x = \underline{f(x)}|\mathcal{G}_t)$$

One would have equality in this case, if the elements of $o^x(t)$ are defined by formulas with only x free.

Pegs also allow for another notion of total object, different from threads. Assume that the set of information states T is closed under suprema of chains: if $C \subseteq T$ is a chain, then $\bigvee C \in T$.[23] Call a peg *polished* iff $o(\bigvee C) := \bigcup_{s \in C} o(s)$ for each chain C. Polished pegs are monotone: if $s \geq t$, then $\{s,t\}$ is chain, so $o(s) = o(s) \cup o(t) \supseteq o(t)$. For a polished peg o^x and a maximal chain C, it is natural to view $o^x(\bigvee C)$ as a total object to which o^x converges; for $o^x(t) \subseteq o(\bigvee C)$, $t \in C$, and due to maximality of C the peg o^x cannot develop into a more refined object with this property.

This notion of total object is clearly different from total objects as threads, since

[23] A *chain* $C \subseteq T$ in partial order T has: $s \geq t$ or $t \geq s$, for all $s, t \in T$. A chain is *maximal* iff there there is no chain in T properly extending it.

in general $\bigcap o^x(\bigvee C)$ need not be equivalent to $x = \underline{\xi}$ for a thread ξ.[24] There are pegs o for which $o(\bigvee C)$ does yield such a 'singleton'. Consider, for example, the peg o^ξ, ξ a thread, with $o^\xi(t)$ the filter generated by $\exists(x = \underline{\xi}|\mathcal{G}_t)$. To see that this yields a 'singleton,' first notice that the blurred object $\mathrm{bo}^\xi(\bullet)$ converges to ξ, since

$$\bigcap\nolimits_{t\in T}\mathrm{bo}^\xi(t) = \{\xi\}.$$

On the assignment space $\mathcal{F}$ the analogue of this singleton is the proposition $x = \underline{\xi}$. But o converges to $x = \underline{\xi}$, because

$$\bigcap\nolimits_{t\in T}\exists(x = \underline{\xi}|\mathcal{G}_t) = (x = \underline{\xi})$$

by Theorem 1.[25]

4.3 Indiscernibility

A basic trait of pegs is that the may be indistinguishable from each other within a certain information state. But what does that mean? Adapting Landman's definition to logic of vision, one gets:

Definition 4 Let o and o' be pegs and $t \in T$ an information state. One says that $o(t)$ and $o'(t)$ are *discernible* from each other—$o(t) \not\equiv_1 o'(t)$—iff (i) or (ii).

i) $\exists X \in P_t[X \in o(t)\ \&\ \exists Y \in o'(t)[X \cap Y = \emptyset]]$;

ii) $\exists X \in P_t[\exists Y \in o(t)[X \cap Y = \emptyset]\ \&\ X \in o'(t)]$.

And $o(t)$ and $o'(t)$ are *in*discernible—$o(t) \equiv_1 o'(t)$—iff they are not discernible from each other.

There is also a notion of discernibility—$\not\equiv_2$—which is weaker than $\not\equiv_1$ but more in line with the tenet of logic of vision. The relation $\not\equiv_2$ requires $\exists Z \in o(t)[Z \subseteq X]$ instead of $X \in o(t)$, etc.; so the discerning property X must be *consistent with* a property of the pegs, they need not *have* the property. For the points we want to make there is not much difference between the two notions.

Proposition 4 *Let o and o' be pegs and let $t \in T$ be an information state.*

$$o(t) \not\equiv_i o'(t) \Leftrightarrow \bigcap o(t) \cap \bigcap o'(t) = \emptyset$$

with $i \in \{1, 2\}$.

Proof. First assume, $o(t) \not\equiv_1 o'(t)$. Then, say, $\exists X \in P_t[X \in o(t)\ \&\ \exists Y \in o'(t)[X \cap Y = \emptyset]]$ (the other possibility is similar). Since $\forall X \in o(t)\forall Y \in o'(t)[\bigcap o(t) \cap \bigcap o'(t) \subseteq X \cap Y]$, it follows that $\bigcap o(t) \cap \bigcap o'(t) = \emptyset$.

Conversely, assume $o(t) \equiv_1 o'(t)$. Then $\forall X \in o(t)\forall Y \in o'(t)[X \cap Y \neq \emptyset]$. This means that $o(t) \cup o'(t)$ is a set of closed sets with the finite intersection property

[24] Note in passing that as a consequence the following two defeasible assumptions are equivalent for polished pegs o (cf. section 3).

M_i There is maximal chain C with $i \in C$;

T_i There exists a total object τ such that $o(i) \sqsubseteq \tau$.

This means that for the present notion defeasible expectations of (perceived) pegs converging to total objects are available.

[25] One could also use chains here, provided they are cofinal in T; see footnote 5.

(because $o(t)$ and $o'(t)$ already are such). By compactness $\bigcap(o(t) \cup o'(t)) \neq \emptyset$. But $\bigcap(o(t) \cup o'(t)) \subseteq \bigcap o(t) \cap \bigcap o'(t)$, so that $\bigcap o(t) \cap \bigcap o'(t) \neq \emptyset$ as well. Almost the same argument works for $\not\equiv_2$. 2

What is still lacking is that pegs become distinguishable through information growth (Landman 1986, 127). For this reason, Landman introduces a notion of indiscernibility involving properties a peg *must* or *may* have. In logic of vision one could do the same. Proposition 5 indicates what can be expected along these lines.

Proposition 5 *Let o and o' be polished pegs. For each chain $C \subseteq T$:*

$$o(\bigvee C) \equiv_i o'(\bigvee C) \Leftrightarrow \forall s \in C : o(s) \equiv_i o'(s)$$

with $i \in \{1, 2\}$.

Proof. First assume $\forall s \in C : o(s) \equiv_i o'(s)$. By monotonicity of o and o' on C both $\{\bigcap o(s)\}_{s \in C}$ and $\{\bigcap o'(s)\}_{s \in C}$ are decreasing sequences of closed sets. Since by proposition 4, $\bigcap o(s) \cap \bigcap o'(s) \neq \emptyset$, it follows that $\{\bigcap o(s) \cap \bigcap o'(s)\}_{s \in C}$ is a set of closed sets with the finite intersection property. By compactness, $\bigcap_{s \in C}(\bigcap o(s) \cap \bigcap o'(s)) \neq \emptyset$. Since $\bigcap o(\bigvee C) = \bigcap_{t \in C} \bigcap o(t)$, it follows that $\bigcap o(\bigvee C) \cap \bigcap o'(\bigvee C) \neq \emptyset$. So, $o(\bigvee C) \equiv_i o'(\bigvee C)$, as required.

To converse is even simpler. Assume $o(\bigvee C) \equiv_i o'(\bigvee C)$. Then $\bigcap o(\bigvee C) \cap \bigcap o'(\bigvee C) \neq \emptyset$ by proposition 4. It follows that $\bigcap_{s \in C}(\bigcap o(s) \cap \bigcap o'(s)) \neq \emptyset$; whence, again by proposition 4, $\forall s \in C : o(s) \equiv_i o'(s)$. 2

Proposition 5 says that two pegs are distinguishable in the limit of a chain, iff they are already distinguishable at a certain stage in that chain (and conversely). It is open to debate whether this consequence of compactness is desirable or not. If one thinks of $o(\bigvee C)$ and $o'(\bigvee C)$ as 'regulative ideals' approximated by $o(s)$ and $o'(s)$, $s \in C$, there seems to be no metaphysical reason why their differences should show up at a 'finite' stage. On the other hand, the fact that these differences are absent at a more mundane level invites a parsimonious mind to neglect them altogether; only those objects should be discerned which can be distinguished with 'finite' means.

The propositions 4 and 5 show that important features of pegs are reducible to properties of the blurred objects we have been using thus far. We leave it to the reader to judge whether this reduction makes pegs superfluous for logic of vision. Section 5.3 discusses the philosophical aspect of partial objects in more detail.

5

Postlude

In the body of the paper we have been mainly concerned with Marr's approach to 3-D vision and the logic of perception reports. This section tries to buttress the proposed model, in particular the use of inverse limits and the attendant notion of partial object, by comparing it with suggestions put forward in the psychological, the linguistic, and the philosophical literature.

1 Marr, agnosia and hierarchical models

In his 'Artificial Intelligence–a personal view' (Marr 1990, first published 1977; cf. also Marr 1982, p. 357 *passim*), Marr argues that his hierarchy of 3-D models, which can be indexed in several ways (cf. section 3.5) may contain features which apply more generally:

1. The perception of an event or object must include the simultaneous computation of several different descriptions of it, that capture diverse aspects of the use, purpose, or circumstances of the event or object.
2. That the various descriptions described in 1. include coarse versions as well as fine ones. These coarse descriptions are a vital link in choosing the appropriate overall scenarios . . . and in establishing correctly the roles played by the objects and actions that caused those scenarios to be chosen. (Marr 1990, p. 140).

Marr gives the example of the following pair of sentences:

(A) The fly buzzed irritatingly against the window pane.

(B) John picked up the newspaper.

The juxtaposition of (A) and (B) strongly suggests that John picked up the newspaper to squash the fly. This process can be seen as the move from a partial model $\mathcal{M}_t$ (satisfying (A)) and a partial model $\mathcal{M}_s$ (satisfying (B)) to a more refined model $\mathcal{M}_r$ ($r \geq s, t$) where the newspaper has the additional property of being capable to squash insects against a brittle surface. In the model $\mathcal{M}_r$, the newspaper might still be represented as a single object; however, sentence (B) might have continued '. . . and sat down to read' so one should simultaneously consider a refinement $\mathcal{M}_v$

$(v \geq s, t)$, where the newspaper has the property of being reading material, and where it is decomposed into a collection of articles. So the general features of the model introduced in section 3 are also present here: bonding mappings which connect objects in coarse and finer descriptions, and predicates which appear from a certain stage onward.

Now it turns out that in the area of cognitive neuropsychology several experiments have been performed which point toward the presence of inverse system like structures in the human semantic system. These experiments are concerned with the phenomenon of *visual agnosia,* a failure to identify meaningful pictures (or objects) that cannot be accounted for in terms of defective perceptual analysis. One of the first extensive experimental investigations of patients suffering from this syndrome can be found in Warrington 1975. Typically, such a study proceeds via *picture naming*, where a subject is presented with a picture for which he has to provide a name. This process is thought to require access to three different representations: to stored structural knowledge about objects, to semantic knowledge and to a stored phonological description, and thus provides an avenue for obtaining insight in the interactions between these representations. The theoretical interpretation of the experimental findings is still controversial (see the special issue of *Cognitive Neuropsychology* **5** (1) 1988, so we shall only give a brief description of some of the literature relevant to the proposed model; the references will provide further details and, occasionally, opposing views.

In Warrington's 1975 experiments, subjects were presented with a picture and had to answer yes/no questions such as: is it an animal? is it a bird? is it dangerous? is it English? is it larger than a telephone directory? She found that the error probability for the first type of question was low, whereas the subjects' answers to the latter type of question appeared to be completely random. In a further test for object recognition, she obtained the interesting result that subjects tended to choose a superordinate category (e.g. 'animal' for 'cow'). In particular, objects from categories consisting of many exemplars only differentiated by detail presented great difficulties: subjects could recognise a flower, but not which particular flower; they could differentiate between fruits and vegetables, but had trouble identifying which particular one. Two types of incorrect response frequently occurred: choice of a (correct) general superordinate category, or the incorrect choice of an exemplar from the same category ('dog' for 'cat') (Warrington 1975, p. 642).

Warrington herself interpreted these findings as being evidence for a hierarchical structure of the semantic system: on being presented with a picture, a subject first activates the most general category to which the object belongs, and then moves down the hierarchy. Visual agnosia was taken to be an impairment of this *semantic* system. Later critics (e.g. Rapp and Caramazza 1989, p. 270) argued that patients suffering from visual agnosia can extract only a limited amount of information from the visual array; indeed it seems reasonable to assume that in order to identify an object precisely, one must extract all the visual features that distinguish the object from others in its category. Now clearly one needs less information to determine whether an object is an animal, than to determine whether it is larger than a telephone directory, hence it is not surprising that patients did not perform above chance level on the latter task.

Such a view can be modelled quite well in the framework proposed here, with its two intertwined notions of partiality: the visual system provides the thread, up to some approximation, and the semantic system checks whether a predicate is applicable, and if so, whether it does apply, to the partially given thread. One could even use the distinction between inverse system and inverse limit here: one might say, somewhat hyperbolically, that one needs an infinite amount of information, i.e. precise knowledge, of an object to determine whether it is larger than a telephone directory; such a predicate could be taken to live only on the inverse limit, not on its approximations.

It should be said though, that Warrington and co-workers still believe that this view, agnosia as a consequence of restricted visual information, is too simple (cf. Shallice 1988).

Recall that picture naming requires access to three different types of representation: structural knowledge about objects, semantic knowledge and phonological descriptions. Humphreys, Riddoch and Quinlan 1988 studied the process of picture naming with the aim of obtaining more detailed information about the interaction of these levels. One may entertain (at least) two different theories on the exact nature of accessing these representations:

1) The process is *discrete* in the sense that information is only transmitted to the next stage after the construction of the representation has been finished; for example, the structural description of a picture or an object must be finished to the extent that no other description remains activated, before it is passed on to the semantic level.
2) On the other hand, the process could be a *cascade* in the sense that semantic information about a picture can be activated prior to the completion of the structural description of the object.

There exists a clear model theoretic distinction between the two views: on the first view, a semantic system is best represented as an ordinary first order model, with predicates applicable to objects whose structural description is completed, so, one might as well say, to unstructured objects; whereas on the second view, predicates should also be applicable when the structural description is not yet completed, so that it becomes important to keep track of the stages of structural description of an object. The latter option is more like an inverse system of first order models. Hence we view the experiment to be described as a rough indication of which type of semantic organisation is to be preferred.

Humphreys et al. 1988 proposed to decide between the discrete and the cascade theories in the following manner. The discrete theory predicts that structural similarity or dissimilarity between pictures will have no influence on the probability that the subject will come up with the correct name, since name giving starts only after the structural description has been completed, even when this takes a relatively long time (as in the case of structurally similar pictures). The cascade theory, on the other hand, predicts that structural similarity between pictures must have an influence on the probability of a correct answer: before the structural description has stabilised, there is ample time for interaction between semantic and structural description.

More precisely, Humphreys et al. studied the interaction between picture name frequency and structural similarity of pictures. Name frequency (the frequency with which a name occurs in print) is thought to affect the access to a picture's phonological representation, hence should be conditionally independent of structural similarity (given the semantic representation). The experimental results showed that there is little effect of name frequency in the case of structurally similar pictures (whose descriptions take a fairly long time period to access), but a large effect in the case of structurally dissimilar pictures (which are relatively easy to access). This result is what the cascade theory would predict: since name information is made available during the completion of the structural description, name frequency, which pertains to the phonological representation, has no effect. A further interesting result was, that in the case of structurally similar pictures, the reaction times for naming correlated strongly with the degree of structural similarity, and not so for the case of structurally dissimilar pictures. This seems to show that there must be a relatively high degree of structural similarity before it starts affecting naming performance. One explanation for this phenomenon is that in the case of structurally similar pictures a superordinate, 'generic' structural description is activated, corresponding to a category name (say 'bird'; here the authors refer to Marr's hierarchy of 3-D models.), which in turn activates descriptions of many exemplars belonging to the category, thus further slowing down the process of name-giving.

The upshot of all this seems to be that, in the context of semantics it is sensible to distinguish two notions of partiality, one pertaining to objects and one to predicates, as has been done above.

2 Interpreting evidentials

The semantics of evidentials is a second empirical field where conditional quantification on inverse limits seems promising.[1] This section begins with a quick overview of what evidentials are, and then indicates how logic of vision could be used to provide semantics.

[1] Good introductions to the subject are Anderson 1986, and especially Willett 1988, which gives an overview of the articles in Chafe and Nichols 1986, among other things.

2.0.1 Evidentials

An example of an evidential in English is the phrase ‘I hear’ in (1) (with small capitals indicating stress).

(1) I hear Johan has won the PRIZE

The main claim of (1) is that Johan has won the prize, and the differently tensed ‘I hear’ indicates that this claim is based on hearsay. By contrast, ‘heard’ in (2) is *not* an evidential.

(2) I HEARD that Johan has won the prize.

For (2) claims that the speaker heard something; the perception verb is part of the main predicate, it does not just indicate the source of information. Before discussing further examples of evidentials, also from languages other than English, let us first delimit more precisely what evidentials are and what they are not.

Whenever a speaker makes a factual claim, it is based on a source of information, such as perception, the reports of others, or an (inductive) inference. Evidentials are the linguistic means to indicate these sources. Thus, they are among the epistemic modalities, but there is overlap into the areas of tense and aspect as well (Willett 1988, pp. 51–55). Despite vagueness, the borderlines between evidentails and other parts of speech are clear enough to phrase a working definition.

Definition 1 A *true (gramaticised) evidential* shows the kind of justification a speaker has for a factual claim, in such a way that

i) it is a specification added to a claim about something else, not the main predication;

ii) it indicates the source of evidence as its primary meaning, not just as a contextual implication.

Morphologically, evidentials are inflections, clitics, or other free syntactic elements, not compounds or derivational forms. Cf. Anderson 1986, pp. 274–275, and Willett 1988, p. 84.

Since there are many sources of information, the questions arise (i) what *are* the evidential contrasts that occur in language, and (ii) *how* are they marked. According to Willett, the primary distinction for evidentials is whether the speaker's information is based on *direct* or on *indirect* evidence. Further distinctions can be found in the following table (Willett 1988, p. 56).

$$\text{Types of Evidence}\begin{cases} \text{Direct} \quad - \ \textbf{Attested} \begin{cases}\text{Visual}\\ \text{Auditory}\\ \text{Other sensory}\end{cases} \\ \text{Indirect} \begin{cases} \textbf{Reported} \begin{cases}\text{Hearsay}\\ \text{Folklore}\end{cases} \\ \textbf{Inferring} \begin{cases}\text{Results}\\ \text{Reasoning}\end{cases}\end{cases}\end{cases}$$

The types of evidence can be indicated with varying detail and in several ways. For instance, Himalayan languages such as Chepang, Sherpa and Tibetan have developed verb suffix systems for this purpose. Sherpa, in particular, uses the suffix *-no(k)* in the present tense to indicate direct evidence, and in the past and future tense to indicates hearsay. Sherpa also has a suffix for unspecified direct evidence in the past tense, and one for unspecified indirect evidence in the nonpast. There are also languages with fixed particles to express evidentiality, such as the Burmese-Lolon language Akha. (cf. Willett 1988, p. 76 and Appendix B.) We give a more detailed example of evidentaility in Japanese (cf. Horie 1993, pp. 9–10, or Holzapfel 1997).

In Japanese there is overlap between evidentiality and tense, in that the present tense indicates directness (this is a feature of several languages). This phenomenon can be found in the distinction between describing one's own sensations and those of others. In the first person (3) the direct form 'atui', for: *be hot*, is acceptable, while it is unacceptable in the third person (4).[2]

(3) *Watasi wa atui.*
I T.M be hot.
'I am hot.' (self's sensation)

(4) *Kare wa *atui*
He T.M be hot.
'He is hot.' (other's sensation)

Indeed, third person descriptions are only grammatical for 'objectivised' predicates like 'feel hot' or 'appears to be hot', as in (5).

(5) *Kare wa atui yoo da.*
He T.M. hot appear COP.
'He appears to be hot.' (other's sensation)

We shall now indicate how logic of vision can be used to develop a formal semantics of evidentiality.

2.0.2 The semantics of evidentials

Logic of vision seems well-suited to study the interpretation of evidential expressions. As a step in this direction, we first propose a formal semantics for the Japanese evidentials *yoo da* and *rasi*; evidentiality in English is considered next.

Japanese uses nonpast tense to indicate directness; as for example in (6), again from Horie 1993, p. 10.

(6) *Hanako wa byooki da.*
Hanako T.M sickness COP.
'Hanako is sick.'
EVID: fact (known)

[2] In the following 'T.M.' is short for: topic marker, and 'COP' for : copula.

In this case the semantics could just be: $S(h)$, interpreted on the inverse limit.[3] The evidential *yoo da* signals that the source of evidence is direct perception (cf. also Aoki 1986, 229).

(7) *Hanako wa byooki no yoo da.*
Hanako T.M sickness GEN appear COP.
'Hanako looks sick.'
EVID: judgment based on direct visual information

Hence we propose the interpretation: $\exists(x = h \wedge S(x)|\mathcal{G}^{\mathrm{SP}})$. By contrast, the semantics of indirect forms is more complex (but still available). The form 'rasii' indicates that that the claim is inferred from visible information[4].

(8) *Hare teiru rasii.*
clear result appear.
'It seems to be clear.'
EVID: judgment based on indirect visual information

Kasiola (as quoted by Aoki 1986, p. 231) points out that (8) is unacceptable when the speaker is looking at the sky; but (8) could be used if it is based on perceiving, e.g., the brightness of a room. Let φ describe what is perceived directly. Then this case involves the following inference (with ψ_t the interpretation of ψ at stage t, and the variable e running over situations):

$$\frac{\exists(x = e \wedge \varphi(x)|\mathcal{G}_s^{\mathrm{SP}}) \quad \forall e[\varphi(e), C(x)]_t \quad (t \geq s)}{C(e)_t}$$

This rule follows from combining Veridicality for the first premise with monotonicity for $\exists(\bullet|\mathcal{G}_s^{\mathrm{SP}})$.

It is worthwhile to observe that similar evidentials are available in English; in particular, (9a-c) appear to have the same semantics as (6–8).

(9) a. It's a fact that Jack is SICK.
b. I see Jack look SICK.
c. I see that the sun is RISING.

Stress is added to preclude what is now an evidential to be the main predication. In fact, English, like other languages, has a rich class of phrases to express evidentiality (see Chafe 1986 for an overview). We discuss the phrases 'may', 'must', 'oddly enough' besides perceptual evidentials, and propose a semantics in logic of vision.

'May' The modals 'may' and 'might' concern the reliability of the information source as in (10) (cf. also Veltman 1996, among others).

(10) a. I see a man. He may be John.
b. $\exists x \exists(M(x)|\mathcal{G}_s^{\mathrm{SP}}) \wedge \exists t \geq s.\exists(M(x) \wedge x = j|\mathcal{G}_t^{\mathrm{SP}})$

[3]To stress the epistemic character one might prefer: for all $t \in T$: $S(h)_t$, with $S(h)_t$ the interpretation of '$S(h)$' in stage $\mathcal{M}_t$ of a refining system. The preservation theorems in section 6 show that for the present case this is equivalent to $S(h)$ in the inverse limit.

[4]'Rasi' can also be used to indicate hearsay.

Here and in what follows the quantifier $\exists x$ is taken as in DPL, so that it binds the variable x in the second conjunct. The anaphor sentence of (10) says that the information available to the speaker is consistent with the man introduced by the antecedent sentence being John; but this consistency could show up at a later stage. As in section 4, the gender of the pronoun 'he' imports the fact that John is a man into the scope of $\exists(\bullet|\mathcal{G}_s^{\mathrm{SP}})$.

'Must' The evidential auxiliary 'must' signals an inductive inference. For instance, the antecedent sentence of (11) combined with background information allows the speaker to conclude that he sees John.

(11) I see a man. He must be John.

This inference has the following form:

$$\frac{\exists x.\exists(M(x)|\mathcal{G}_s^{\mathrm{SP}}) \quad \forall x[\beta(x), M(x) \rightarrow x = j]_s}{\exists(M(x) \wedge x = j|\mathcal{G}_s^{\mathrm{SP}})}$$

Again this rule is based on the monotonicity of $\exists(\bullet|\mathcal{G}_s^{\mathrm{SP}})$, which can be employed as soon as the background information expressed by β holds for the x. We trust the reader knows how to adapt this rule if he thinks different stages are involved.

'Oddly enough' Evidentials could also indicate that the facts fall short of expectations, as 'oddly enough' in (12).

(12) I saw a man. Oddly enough it turned out to be John.

In (12) the antecedent sentence combined with already available information induces the expectation that the man the speaker saw was not John. The anaphor sentence states the surprise in discovering that this expectation can be countered. More formally, in a suitable context the antecedent sentence gives (i) and defeasible (ii):

i) $\exists x.\exists(M(x)|\mathcal{G}_s^{\mathrm{SP}})$;

ii) $\forall t \geq s : \exists(M(x) \wedge \mathit{non\text{-}id}(x, j)|\mathcal{G}_t^{\mathrm{SP}})$.

Here '*non-id*' is a positive approximation of '$\neq$' (cf. sections 3.0.10 and 3.0.11). Information available at one of the later stages $t > s$ implies that $x = j$, so that the defeasible assumption that x was not John should be withdrawn, oddly enough.

Perceptual evidentials Finally, let us consider perceptual evidentials at work in a perceptual claim; e.g. (13).

(13) a. I observed that Jack saw Sharon DANCE.

b. $\exists(\exists(x = s \wedge D(x)|\mathcal{G}_s^{\mathrm{J}})|\mathcal{G}_s^{\mathrm{SP}})$

The semantics of (13) asks for iterating the perception operator. Such iterations are already available in logic of vision. But it might be more realistic (and interesting) to assign to each perceiver PR a refining inverse system $\mathcal{S}^{\mathrm{PR}} := \langle \mathcal{M}_s^{\mathrm{PR}}, h_{st}^{\mathrm{PR}} \rangle_{s,t} \in T$, all converging to the *same* inverse limit $\mathcal{M}$. The operator $\exists(\bullet|\mathcal{G}^{\mathrm{PR}})$ and its defeasible expectations are then interpreted in terms of $\mathcal{S}^{\mathrm{PR}}$, as before. This approach would be a first step towards a full-fledged theory of intensionality, where alternative worlds are considered as well. At the technical side it leads to the question: when do two refining inverse system induce isomorphic inverse limits?

This concludes our preliminary thoughts on the semantics of evidentials. It is clear, or so we think, that this is a promising area of research where logic of vision can be applied. It would be particularly interesting to provide semantics for evidentials in languages other than English and Japanese. To name but two candidates out of many, we think of (i) the Californian Indian Languages, such as Northern Pomo, which use verbal suffixes to indicate roughly the same evidential distinctions as available in Japanese (cf. O'Connor 1987, pp. 46, 289-291), or (ii) the Burmese-Lolon language Akha which has expectation as one of its evidential parameters (Thurgood 1986, Willett 1988, pp. 78-79, 83).

3 Husserl on perception

As a third example, we show that the model introduced in section 3 has similarities to Edmund Husserl's description of perception, in particular with the notion of partiality introduced by him. Here, we shall present a very brief outline of Husserl's theory with some representative quotations, and we compare the ingredients of the formal model with Husserl's informal suggestions.[5]

3.1 Husserl's notion of partiality

One of Husserl's main concerns was the persisting identity of perceived objects, in the face of the fact that all perception is partial. An object is always viewed from a certain 'perspective' (*Abschattung*), which dictates the kind of questions that can be asked about the object, leaving many other possible questions undecided. What accounts for the identity of the object, when we move from one perspective to another? Husserl sought the solution in the peculiar way in which our partial perceptual knowledge is organised.

> *Das Wahrnehmen ist . . . ein Gemisch von wirklicher Darstellung, die das Dargestellte in der Weise originaler Darstellung anschaulich macht, und leeren Indizieren, das auf mögliche neue Wahrnehmungen verweist . . . eben diejenigen, in denen sich derselbe Gegenstand von immer neuen Seiten zeigen würde* (Husserl 1966, p. 5)
>
> Perception is a mixture of actual representation, which makes the object represented intuitive in the manner of an orginal representation, and empty indexing, which points toward new possible perceptions . . . namely those in which the object would show new aspects of itself.

Hence perception is the very opposite of merely receiving sense impressions; somehow, in any given perceptual experience possibilities for future experiences are included. Indeed, this is the fundamental structure of all knowledge:

[5]The translations have been prepared using Cairns 1973.

> *So ist eine Fundamentalstruktur des Weltbewußtseins . . . die Struktur der Bekanntheit und Unbekanntheit mit der ihr zugehörigen durchgängigen Relativität und der ebenso durchgängigen relativen Unterscheidung von unbestimmter Allgemeinheit und bestimmt er Besonderheit.* (Husserl 1972, p. 33)
>
> Hence the fundamental structure of our consciousness of the world. . . is a structure of acquaintedness and unacquaintedness with its attendant ubiquitous relativity, and the likewise ubiquitous relative distinction between indeterminate generality and determinate particularity.

Partiality in the sense of Husserl is therefore a positive concept; it denotes not only lack of knowledge, but also ways of filling up the lacunae. The essential notion here is that of a 'horizon', one meaning of which is introduced in the following quotation:

> *So hat jede Erfahrung von einen einzelnen Ding ihren Innenhorizont; und 'Horizont' bedeutet hierbei die wesensmäßig zu jeder Erfahrung gehörige und von ihr untrennbare Induktion in jeder Erfahrung selbst. [. . .] Diese urspüngliche 'Induktion' oder Antizipation erweist sich als . . . hinausmeinend nicht nur in der Weise eines Antizipieren von Bestimmungen, die als sich herausstellende jetzt erwartet werden, sondern auch nach anderer Seite hinausmeinend über dieses Ding selbst mit allen seinen antizipierten Möglichkeiten künftiger Weiterbestimmung, hinausmeinend auf die anderen mit ihm zugleich, wenn auch zunächst bloß im Hintergrund bewußte Objekte. Das heißt, jedes erfahrene Ding hat nicht nur einen Innenhorizont, sondern es hat auch einen offen endlosen Außenhorizont von Mitobjekten . . .* (Husserl 1972, p. 28)
>
> Hence every experience of a single object comes with an inner horizon; 'horizon' here points to the induction which inseparably belongs to every experience. . . . This original 'induction' or anticipation can be seen as projection; not only as anticipation of determinations whose occurrence is now expected, but also projecting beyond the object itself, with its anticipated possibilities for further determination, projecting toward other objects, of which we are initially conscious only as present in the background. In other words, to every object is attached an inner horizon and an infinite outer horizon of simultaneously present objects . . .

The inner horizon thus comprises the questions, with their expected answers, that can be posed about the perceived object itself; the outer horizon is concerned with questions and expected answers about the relation of the perceived object to other objects. The essential phrase here is '. . . the induction which inseparably belongs to every experience': the horizon is always present whenever we experience something. The horizon thus allows us to experience the identity of an object in different circumstances:

> *Jedes Object, jeder Gegenstand überhaupt (auch jeder immanente) bezeichnet eine Regelstruktur. . . Als sein vorgestelltes, wie immer Bewußtes bezeichnet es*

> *sofort eine universale Regel möglichen sonstigen Bewußtsein von demselben...* (Husserl 1977, p. 55)
>
> Any 'objective' object, any object whatsoever (even an immanent one) points to a structure... that is governed by a rule. As something the ego objectivates, something of which he is conscious of in any manner, the object indicates forthwith a universal rule governing possible other consciousness of it as identical...

One could phrase this distinction between inner horizon and outer horizon in logical terms as that between properties of an object, and relations of the object to other objects. Husserl introduces other notions of horizon as well. For instance in the *Cartesianische Meditationen* (Husserl 1977), the horizon is viewed temporally, as the pattern of recollections and expectations of past and future perceptual experiences. In the *Analysen zur passiven Synthesis* (Husserl 1966), the horizon is again viewed 'spatially', but now the inner horizon refers to already perceived properties of an object, which can be further determined, whereas the outer horizon refers to properties which can be perceived from different perspectives on that object. Although partiality is positive in the sense that it contains within it possibilities to reduce it, still complete knowledge (perceptual or otherwise) of an object is unattainable.

> *So gehört zu jeder äußeren Wahrnehmung eine im Unendlichen liegende Idee, die Idee des voll bestimmten Gegenstandes .... Ich sprach von einer im Unendlichen liegenden, also unerreichbaren Idee, denn daß es eine Wahrnehmung geben könnte (als einen abgeschlossenen Prozeß kontinuierlich ineinander übergehender Erscheinungsverläufe), die eine absolute Kenntnis des Gegenstandes schüfe ... das ist durch die Wesensstruktur der Wahrnehmung selbst ausgeschlossen; denn evidenterweise ist die Möglichkeit eines plus ultra prinzipiell nie ausgeschlossen.* (Husserl 1966, p. 20–21)
>
> Hence to every perception belongs an infinitely removed idea, namely the idea of a completely determined object ... Deliberately I called this an infinitely removed, hence unattainable idea, because the structure of perception itself excludes the possibility that there could be a perception (as a finished process of continuous flow of impressions) which gives absolute knowledge of the object; for obviously the possibility of a *plus ultra* is in principle never excluded.

However, this does not necessarily lead to scepticism. Knowledge is often used as a basis for acting, and we may act reliably using partial information only.

> *Das thematische Interesse, das in Wahrnehmung sich auslebt, ist in unserem wissenschaftlichen Leben von praktischen Interessen geleitet, und das beruhigt sich, wenn gewisse für das jeweilige Interesse optimale Erscheinungen gewonnen*

sind, in denen das Ding soviel von sich selbst zeigt, als dieses praktische Interesse fordert. (Husserl 1966, p. 23)

The thematic interest present in perception is guided by pragmatic interests, and these can be satisfied when certain optimal (relative to the present interest) impressions have been obtained in which the object shows itself just so much as the pragmatic interest demands.

3.2 Husserl and a logic of vision

It is fairly easy to find parallels between Husserl's theory of perception and the formal constructions introduced in this article.

Perspectives can be seen as (finite) first order models. For any two such models $\mathcal{M}$ and $\mathcal{N}$, if they are perspectives of the same situation, then there will be a mapping $h : \mathcal{M} \longrightarrow \mathcal{N}$ (say) which identifies the objects in $\mathcal{M}$ with objects in $\mathcal{N}$. Hence we are led to a notion of inverse system as discussed in section 3. The horizon determines the possible questions that can be asked; this corresponds to the signatures of the various approximating models. The expected answers to the questions suggested by the horizon were formalised by means of the default rule

$$\frac{\exists(\varphi|\mathcal{G}_s)}{\exists(\varphi|\mathcal{G}_t)}\ t \geq s$$

at least for the simple case where we expect a feature to persist upon 'looking closer'. The default rule, i.e. the second component of the analysis of 'see' above, captures Husserl's ideas about '...the induction which inseparably belongs to every experience'. However, much more general expectations can be treated in this framework. Consider e.g. the following adaptation of the rabbit/duck example: we perceive a silhouette against the sky which can be interpreted as either a duck or a rabbit. Now it is a 'law of nature' that ducks have feathers and rabbits have fur, so upon coming closer we should be able to decide whether what we perceive is rabbit or duck. Formally this situation can be represented as follows. At stage s our perception ('this is a duck' or 'this is a rabbit') is given by $\exists(D(x)|\mathcal{G}_s)$ or by $\exists(R(x)|\mathcal{G}_s)$, perhaps alternatingly. The variable x is used to represent the demonstrative.

There is a stage $r \geq s$ such that for all $t \geq r$ we have $\mathcal{M}_t \models \forall x(\mathrm{D}(x) \rightarrow \mathrm{Fe}(x)) \equiv \forall x(\mathrm{R}(x) \rightarrow \mathrm{Fu}(x))$. Note that this statement need not hold at $\mathcal{M}_s$, for example because the predicate Fe is not interpreted in $\mathcal{M}_s$. The fact that we are concerned with a *law* is rendered formally by the condition that the universal Horn sentences involved hold at *all* $t \geq r$; it expresses that we have not found a counterexample. If the perception is given by $\exists(D(x)|\mathcal{G}_s)$, then we expect $\exists(D(x)|G_t)$, and this triggers $\exists(\mathrm{Fe}(x)|G_t)$; similarly for the other case.

Note that the preservation of universal Horn formulas allows us to satisfy a consistency requirement: from $\exists(D(x)|\mathcal{G}_s)$ we may deduce non-monotonically $D(x)$ and apply the law $\forall x(D(x) \longrightarrow \mathrm{Fe}(x))$ in the inverse limit to obtain $\mathrm{Fe}(x)$; or we may proceed as above and deduce non-monotonically $\mathrm{Fe}(x)$ from $\exists(\mathrm{Fe}(x)|\mathcal{G}_t)$.

If reality-as-perceived is modelled by an refining system of approximations or perspectives, and reality-as-it-is by the inverse limit of this refining system, the unattainability of complete knowledge about an object, as emphasised by Husserl, can be represented formally by allowing only filtered quantifiers on the inverse limit, no ordinary first order quantifiers. Since $\exists$ and $\forall$ are special cases of filtered quantifiers, we have to say precisely what we mean here. The existential quantifier $\exists x$ is determined by the equivalence relation $=_x$; to implement impossibility of precise knowledge we could for instance allow only quantifiers determined by equivalence relations of the form $=_x \ \& \ R$, where R has sufficiently large classes. E.g. in the cases considered here, the R-equivalence classes could be closed uncountable sets.

Lastly, let us consider Husserl's observation that for practical purposes, less-than-complete information may suffice. This is connected with preservation properties of disjunctions. Acting on an observation usually presupposes that we decide which disjunct of a disjunction is true. In practice, these disjuncts will often be mutually disjoint. In that case, however, a compactness argument shows that there will be some stage $s \in T$ (a finite stage, loosely speaking) at which it is decided which disjunct obtains; there is no need to evaluate the disjunction on the inverse limit itself. This is a consequence of the compactness of the inverse limit (as a topological space), hence of the fact that we started from an inverse system of finite models.[6]

There is one aspect of Husserl's informal theory for which we haven't yet introduced a formal correlate. At several points in *Ideen I*, Husserl emphasises that our notion of the persisting identity of objects is inextricably bound up with the 'continuity' of approximations:

> *Dieses Kontinuum bestimmt sich näher als allseitig unendliches, in allen seinen Phasen aus Erscheinungen desselben bestimmbaren X bestehend, derart zusammenhängend geordnet… daß jede beliebige Linie desselben in der stetigen Durchlaufung einen einstimmigen Erscheinungszusammenhang ergibt… in welchen das eine und selbe immerfort gegebene X sich kontinuierlich-einstimmig 'näher' und niemals 'anders' bestimmt.* (Husserl 1950, p. 351)
>
> This continuum can be characterised as being infinite in all directions, and it consists in all its stages of impressions of the same determinate X, ordered and connected in such a way… that every line through this ordering yields a consistent sequence of impressions… in which the same given X is continuously and consistently determined ever more accurately, and never differently.

What can be gleaned from this passage is that the stages of approximation, here the index set T, should form an ordered continuum, in such a manner that every 'line',

[6] Predicates A are closed on the inverse limit, hence so are the $\exists(A|\mathcal{G}_s)$. We know that $\lim_{s\in T} \exists(A|\mathcal{G}_s) = A$. Now suppose A and B are disjoint predicates, then $\lim_{s\in T} \exists(A|\mathcal{G}_s) = A$ is disjoint from $\lim_{s\in T} \exists(B|\mathcal{G}_s) = A$. By compactness there will be $s \in T$ such that $\exists(A|\mathcal{G}_s)$ and $\exists(B|\mathcal{G}_s)$ are disjoint.

i.e. maximal linearly ordered set, should completely determine an object (formally, a thread). If T is a partial order, the latter condition forces T to be directed. It is an interesting exercise to put a topology on T and (the union of) the universes of the models $\mathcal{M}_t$ so that threads become continuous functions. This can be done, but at present it is not clear that this move changes anything to the logic, so we leave the matter here, except for noting that Marr makes the same point about the importance of continuity (Marr 1982, p. 355).

Bibliography

Akmajian, A. 1977. The Complement Structure of Perception Verbs in an Autonomous Syntax Framework. In *Formal Syntax*, ed. Culicover et al. 427–460. New York: Academic Press.

Allen, P. K. 1987. *Robotic Object Recognition through Vision and Touch*. Dordrecht: Kluwer Academic Publishers.

Anderson, L. B. 1986. Evidentials, paths of change, and mental maps: typologically regular asymmetries. In *Evidentiality: the linguistic coding of epistemology*, ed. W. Chafe and J. Nichols. Advances in Discourse Processes, Vol. XX, Chap. 17, 273–312. Norwood, New Jersey: Ablex Publishing Corporation.

Anscombe, G. E. M. 1981. *The Intensionality of Sensation: A Grammatical Feature*. Oxford: Basil Blackwell. This chapter published in 1965.

Aoki, H. 1986. Evidentials in Japanese. In *Evidentiality: the linguistic coding of epistemology*, ed. W. Chafe and J. Nichols. Advances in Discourse Processes, Vol. XX, Chap. 14, 223–238. Norwood, New Jersey: Ablex Publishing Corporation.

Asher, N., and D. Bonevac. 1989. Determiners and Resource Situations. *Linguistics and Philosophy* 10:567–596.

Ayer, A. J. 1971. *Language, Truth and Logic*. Pelican Books. Harmondsworth, Middlesex, Engeland: Penguin Books. First published: 1936, second edition: 1946, fifth Pinguin reprint: 1978.

Barwise, J. 1981. Scenes and other Situations. *Journal of Philosophy* LXXVIII:369–397.

Barwise, J. 1989a. Logic and information. In *The Situation in Logic*. Lecture Notes, Vol. 17. Stanford, California: CSLI Publications. Distributed by Cambridge University Press.

Barwise, J. 1989b. *The Situation in Logic*. Lecture Notes, Vol. 17. Stanford, California: CSLI Publications. Distributed by Cambridge University Press.

Barwise, J., and J. Perry. 1983. *Situations and Attitudes*. Cambridge, Massachusetts: The M.I.T. Press.

Blamey, S. 1983. Partial Logic. In *Handbook of Philosophical Logic*, ed. D. Gabbay and F. Geunthner. 1–70. Dordrecht: D. Reidel Publishing Company.

Boden, M. A. 1988. *Computer Models of Mind*. Problems in Behavioural Sciences, Vol. 4. Cambridge: Cambridge University Press.

Brady, M. 1993. Computational Vision. In *The Simulation of Human Intelligence*, ed. D. Broadbent. Chap. 5, 121–150. Oxford: Blackwell.

Cairns, D. 1973. *Guide for translating Husserl*. The Hague: Martinus Nijhoff.

Caramazza, A., A. E. Hillis, B. C. Rapp, and C. Romani. 1990. The multiple semantics hypothesis: multiple confusion? *Cognitive Neuropsychology* 7:161–189.

Chafe, W. 1986. Evidentiality in English converstation and academic writing. In *Evidentiality: the linguistic coding of epistemology*, ed. W. Chafe and J. Nichols. Advances in Discourse Processes, Vol. XX, Chap. 16, 261–272. Norwood, New Jersey: Ablex Publishing Corporation.

Chafe, W., and J. Nichols. 1986. *Evidentiality: the linguistic coding of epistemology*. Advances in Discourse Processes, Vol. XX. Norwood, New Jersey: Ablex Publishing Corporation.

Chang, C. C., and H. J. Keisler. 1990. *Model Theory*. Studies in Logic and the Foundations of Mathematics, Vol. 73. Amsterdam: North-Holland. Third edition. First edition: 1973, second edition 1977.

Engelking, R. 1989. *General Topology*. Sigma Series in Pure Mathematics, Vol. 6. Berlin: Heldermann. Revised and completed edition.

Feferman, S. 1984. Towards Useful Type-Free Theories I. In *Truth and the Liar Paradox*, ed. R.L. Martin. 237–287. Oxford: Oxford University Press.

Flum, J. 1985. Characterizing Logics. In *Model-Theoretic Logics*, ed. J. Barwise and S. Feferman. Chap. III, 77–120. New York: Springer-Verlag.

Gee, J. 1977. Comments on the Paper by Akmajian. In *Formal Syntax*, ed. Culicover et al. 461–481. New York: Academic Press.

Grice, H. P. 1961. The Causal Theory of Perception. *Proceedings of the Aristotelian Society* 35 (suppl.):121–152. Revised version published as chapter 15 in Grice 1989.

Grice, H. P. 1989. *Studies in the Ways of Words*. Cambridge, Massachusetts: Harvard University Press.

Halmos, P. R. 1952. Algebraic Logic, I: Monadic Boolean Algebras. *Compositio Mathematica* 12:217–249.

Hendriks, H. 1993. *Studied Flexibility. Categories and Types in Syntax and Semantics.* Doctoral dissertation, ILLC, University of Amsterdam., Amsterdam. ILLC Dissertation Series 1993–5.

Higginbotham, J. 1983. The Logic of Perceptual Reports: An Extensional Alternative to Situation Semantics. *Journal of Philosophy* LXXX:100–127.

Hintikka, K. J. J. 1969. On the logic of perception. In *Models for Modalities*. Synthese Library, Vol. 23. Dordrecht: D. Reidel Publishing Company.

Holzapfel, A. 1997. Zur Evidentialität Im Japanischen. Research report. Eberhard-Karls-Universität Tübingen: Seminar für Sprachwissenschaft.

Horie, K. 1993. *A cross-linguistic study of perception and cognition; Verb complements: a cognitive perspective*. Doctoral dissertation, University of Southern California, Los Angeles, August.

Humphreys, G. W., M. J. Riddoch, and P. T. Quinlan. 1988. Cascade processes in picture identification. *Cognitive Neuropsychology* 6:67–103.

Husserl, E. 1950. *Ideen zu einer reinen Phänomenologie und phänomenologische Philosophie*. Husserliana, Vol. Band I. The Hague: Martinus Nijhoff.

Husserl, E. 1966. *Analysen zur passiven Synthesis*. Husserliana, Vol. Band XI. The Hague: Martinus Nijhoff.

Husserl, E. 1972. *Erfahrung und Urteil*. Philosophische Bibliothek, Vol. Band 280. Hamburg: Felix Meiner.

Husserl, E. 1977. *Cartesianische Meditationen*. Philosophische Bibliothek, Vol. Band 291. Hamburg: Felix Meiner.

Kamp, H. 1984. A Scenic Tour through the Land of Naked Infinitives. Unpublished manuscript, Univeristy of Texas, Austin.

Kamp, H., and U. Reyle. 1993. *From Discourse to Logic*. Studies in Linguistics and Philosophy, Vol. 42. Dordrecht: Kluwer Academic Publishers. Student edition.

Koons, R. C. 1996. Information, Representation and the Possibility of Error. In *Logic, Language, and Computation*, ed. J. Seligman and D. Westerståhl. Chap. 22, 333–345. Lecture Notes. Stanford, California: CSLI Publications.

Landman, F. 1986. *Towards a Theory of Information*. Doctoral dissertation, Department of Philosophy, University of Amsterdam, Amsterdam. Also published with Foris.

Lewis, D. 1972. General Semantics. In *Semantics of Natural Language*, ed. D Davidson and G. Harman. 169–218. Dordrecht: D. Reidel Publishing Company.

Lindström, P. 1966. First-order Predicate Logic with Generalized Quantifiers. *Theoria* 32:186–195.

Marr, D. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* San Fransisco: W.H. Freeman.

Mittwoch, A. 1990. On the Distribution of Bare Infinitive Complements in English. *Journal of Linguistics* 26:103–131.

Mostowski, A. 1957. On a Generalization of Quantifiers. *Fundamenta Mathematicae* 44:12–36.

Muskens, R. 1989. *Meaning and Partiality*. Doctoral dissertation, University of Amsterdam, Amsterdam.

Neale, S. 1992. Paul Grice and the Philosophy of Language. *Linguistics and Philosophy* 15:509–559. Review article of Grice 1989.

O'Connor, M. C. 1987. *Topic in Northern Pomo Grammar*. Doctoral dissertation, University of California, Berkeley,, Berkeley, March.

Rapp, B. C., and A. Caramazza. 1989. General to specific access to word meaning: a claim reexamined. *Cognitive Neuropsychology* 6:251–272.

Reiter, R. 1980. A logic for default reasoning. *Artificial Intelligence* 13:81–132.

Russell, B. 1988. The Philosophy of Logical Atomism. In *Logic and Knowledge*, ed. Robert Charles Marsh. 177–281. London: Unwin Hyman Ltd. First published in 1956.

Shallice, T. 1988. Specialisation within the semantic system. *Cognitive Neuropsychology* 5:133–142.

Thomason, S. K. 1989. Free Construction of Time from Events. *Journal of Philosophical Logic* 18:43–67.

Thurgood, G. 1986. The nature and origine of the Akha evidentails system. In *Evidentiality: the linguistic coding of epistemology*, ed. W. Chafe and J. Nichols. Advances in Discourse Processes, Vol. XX, Chap. 13, 214–222. Norwood, New Jersey: Ablex Publishing Corporation.

van der Does, J. 1991. A Generalized Quantifier Logic for Naked Infinitives. *Linguistics and Philosophy* 14:241–294.

Veltman, F. 1981. Data Semantics. In *Formal Methods in the Study of Language*, ed. J. Groenendijk, T. Janssen, and M. Stokhof. Mathematical Centre Tracts, Vol. 35. Amsterdam: Mathematical Centre.

Veltman, F. 1996. Defaults in Update Semantics. *Journal of Philosophical Logic* 25:221–261.

Verkuyl, H. 1993. *A Theory of Aspectuality. The interaction between Temporal and Atemporal Structure*. Cambridge: Cambridge University Press.

Viberg, Å. 1984. The verbs of perception: a typological study. In *Explanations for language universals*, ed. Butterworth, Comrie, and Dahl. Berlin: Mouton De Gruyter.

Voorbraak, F. 1997. A nonmonotonic observation logic. Unpublished Manuscript, ILLC, University of Amsterdam.

Warrington, E. K. 1975. The selective impairment of semantic memory. *Quarterly Journal of Experimental Psychology* 27:635–657.

Willett, T. 1988. A Cross-linguistic survey of the grammaticization of evidentiality. *Studies in Language* 12(1):51–97.