PERGAMON

Language & Communication 22 (2002) 211–229

LANGUAGE & COMMUNICATION

www.elsevier.com/locate/langcom

A computational approach to linguistic knowledge

Sandy Boucher*, Ian Gold

Department of Philosophy, Monash University, Wellington Road, Clayton, Victoria 3168, Australia

1. Introduction

The rejection of behaviorism in the 1950s and 1960s led to the view, due mainly to Noam Chomsky, that language must be studied by looking at the mind and not just at behavior. It is an understatement to say that Chomskyan linguistics dominates the field. Despite being the overwhelming majority view, it has not gone unchallenged, and the challenges have focused on different aspects of the theory. What is almost universally accepted, however, is Chomsky's view that understanding language demands a theory that posits mental states that represent rules of language. Call this claim, following Cowie (1999), *Representationalism* or (R). According to (R), "[e]xplaining language mastery and acquisition requires the postulation of contentful mental states and processes involving their manipulation" (Cowie, 1999, p. 154).

Although (R) is nothing more than the general assumption on which cognitive psychology is founded applied to the case of language, even it has had its detractors. Critics have argued that linguistic competence should not in fact be thought of as based on the possession of a body of linguistic knowledge but should be thought of, rather, as a kind of skill. This is an important challenge because one might be inclined to think that no recognizable form of Chomskyan linguistics could withstand the falsification of (R).

In this paper we attempt to show that in fact (R) could be false without doing much damage to Chomskyan linguistics at all. Indeed, it is possible that the Chomskyan position could be made *more* coherent by adopting the view we will

* Corresponding author. Present address: Department of Philosophy, University of Connecticut, Storrs CT 06269, USA.

E-mail addresses: sandycboucher@hotmail.com (S. Boucher), ian.gold@arts.monash.edu.au (I. Gold).

PII: S0271-5309(02)00003-4

sketch. Our claim, therefore, is that critics of (R) might be right, but that this does not obviously make them serious critics of the Chomskyan program.

2. Chomsky's program and its critics

As Cowie (1999) has recently argued, Chomskyan theory is not a monolith but is composed of a number of distinct claims. In addition to (R), she identifies four others: *Biological Boundedness,* the view that there are constraints on the kinds of thought that human beings can entertain; *Domain Specificity,* the view that learning a language requires that the learner's thoughts are constrained by principles that apply to the domain of language only; *Innatenesss,* the view that constraints on language are innate; and *Universal Grammar,* the view that the constraints and principles referred to in the principles of domain specificity are identical to the principles of Chomsky's Universal Grammar.

(R) is the least contentious of the Chomskyan claims. Of it Cowie says: "Chomsky's defeat of behaviorism was decisive in that it succeeded in establishing (R)" (p. 159); and "[d]ealing adequately with the phenomenon of language, Chomsky made clear, requires the admission of mental states, in particular representational states, into the ontology of psychological theory. It requires endorsement of (R)" (p. 162). For Chomsky, (R) implies that linguistic competence is best understood as an implicit grasp by the speaker of a language of the rules governing language use. The task of linguistics is the explication in theoretical terms of the grammar already known by the language user. He writes:

> The problem for the linguist... is to determine... the underlying system of rules that has been mastered by the speaker-hearer... Hence, in a technical sense, linguistic theory is mentalistic, since it is concerned with discovering a mental reality underlying actual behavior (Chomsky, 1965, p. 4)

The precise sense in which Chomsky thinks ordinary speakers have actual knowledge of a system of grammar is not entirely clear from his writings (see Stich, 1971; Devitt and Sterelny, 1989). Chomsky certainly wants to ground linguistic competence in *some* sort of knowledge, and in some places he appears to endorse the view that this knowledge is propositional, but he does not claim that a speaker possesses *explicit* propositional knowledge of syntactic rules, nor that a speaker is in any way conscious of her possession of them (see Fodor, 1981). Rather, the speaker's knowledge is "tacit," or "implicit": "Obviously every speaker of a language has mastered and internalized a generative grammar that expresses his knowledge of the language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them... " (Chomsky, quoted in Harman, 1967, p. 76).

This thesis has been challenged by a number of philosophers on the grounds that a speaker cannot be said to possess propositional knowledge of *any* kind about the syntactic rules governing the use of her language. Devitt and Sterelny (1989), Stich (1971), and others, have pointed out that it is difficult to see what this kind of propositional knowledge could have in common with the ordinary kind. Propositions that we can be said to know are usually taken to be propositions that we have reflected upon or are able to reflect upon, that we can understand when stated, and that we can be brought to acknowledge. Yet, as Stich (1971) argues,

> [p]eople—exempting a few linguists—have never been aware of the facts of linguistic theory; they are incapable of understanding them. And some, though competent speakers, are intellectually incapable of ever coming to understand them. If, nonetheless, people know the propositions of linguistic theory, it is surely an unusual case of knowledge. (p. 486)

Does the qualification that this knowledge is tacit help? On the face of it, it would appear not. For something to count as a piece of tacit, or implicit, propositional knowledge, there must still be a sense in which its possessor is *capable* of becoming consciously aware of it.[1] What we usually mean when we say that a person possesses tacit knowledge of some proposition is that although there is no ordinary circumstance in which they are likely to entertain it, they are nonetheless in a position to assent to it if given suitable prompting. However, an ordinary speaker is not in a position to assent to the propositions detailing such rules even if prompted. Adding to the mystery is the fact that it is unclear how such tacit knowledge of grammar could be appropriately integrated with other beliefs that we hold. These true beliefs would have to be of a special kind, "inferentially isolated" from our other beliefs (Devitt and Sterelny, 1999). It is thus hard to understand how tacit syntactic knowledge could be propositional, yet this appears to be Chomsky's view (see e.g. Chomsky, 1969).

Devitt and Sterelny (1989) have also suggested that Chomsky's propositional thesis harmonizes badly with a central notion of Chomskyan transformationalism, namely the idea that autonomous, self-contained "modular systems" in the mind are the real repositories of the relevant internalized syntactic information which is hard-wired into them. The very autonomy and separateness of modular systems rules out the possibility of any access to their contents by the processing function in the mind

[1] There is, however, a tradition (see Evans, 1981; Davies, 1987, 1989; Peacocke, 1986) of construing tacit knowledge in a more metaphorical way according to which a subject possesses tacit knowledge of a theory if she exhibits dispositions corresponding to the axioms of the theory (Davies, 1987). This shifts the focus from mental representation to observable behavior and in this way side-steps the question of the mental states that underlie the dispositions. Because the tacit knowledge approach is silent on the question of the cognitive bases of the dispositions at issue, it is quite compatible with the view that what underlies the behavioral dispositions are mental representations of the very axioms to which the dispositions correspond. Positing tacit knowledge, therefore, does not exclude the possibility of there being *explicit* knowledge as well; rather, the tacit knowledge theorist is merely adopting a behaviorist methodology according to which mental representations must be treated as theoretical posits to be inferred from behavioral data in the manner of methodological behaviorism. Despite the virtues of this approach, we do not adopt it here. The tacit knowledge theorist is attempting to produce a sympathetic combination of knowledge-that and knowledge-how in the sense that it is know-how (behavioral dispositions) that warrant the claim of tacit knowledge, but possessing tacit knowledge is compatible with possessing knowledge-that as the underlying basis of the know-how. Our concern, however, is not to combine the two kinds of knowledge but to propose a third way "between" them.

responsible for the formation of conscious beliefs. Therefore the syntactic abilities encoded in the modules cannot consist in the speaker's knowledge of the propositions expressing those abilities (Devitt and Sterelny, 1989; see also Fodor, 1983).

3. Semantic representationalism: Dummett

The central features of the debate concerning whether and how syntactic competence ought to be represented as propositional knowledge possessed by speakers have been reproduced at the semantic level of meaning, casting further doubt on the tenability of (R). What is at stake in this debate has become clear in the light of Michael Dummett's attempt to ground a version of anti-realism in an analysis of the nature of semantic competence.[2] In Dummett's (1973) opinion, "a theory of meaning is a theory of understanding" (p. 92). Elsewhere, (Dummett, 1993b) he says:

> What we are after is an account of the sort of understanding which a speaker of the language has... Its function is solely to present an analysis of the complex skill which constitutes mastery of a language, to display, *in terms of what he may be said to know,* just what it is that someone who possesses that mastery is able to do... (p. 37, emphasis ours).

Further, a "theory of meaning must do much more than simply analyse the way in which a sentence is determined as true, when it is true, in accordance with it composition: if has, among other things, to say what a speaker knows when he understands an expression of the language..." (Dummett, 1982, p. 105). "A theory of meaning will, then, represent the practical ability possessed by a speaker as consisting in his grasp of a set of propositions..." A speaker can be said to possess "implicit knowledge of those propositions by means of which we give a theoretical representation of that ability" (Dummett, 1993b, p. 36).

It thus appears that for Dummett semantics centrally involves the ascription of knowledge to speakers of a language. The problem of accounting for meaning is assumed to come down to the problem of what knowledge we take a language-user's understanding of a given sentence to consist in.[3] It would perhaps seem obvious that

[2] We are concerned here only with Dummett's account of semantic competence, not his argument for anti-realism.

[3] There is an exegetical question here as philosophers are far from united in interpreting Dummett's views in this way. Alexander George, for example, has objected to Devitt's characterization of Dummett's attitude to knowledge-ascription in semantics (see George, 1984). According to George, it is not the case that Dummett is disposed to be comfortable with the idea of the attribution of semantic knowledge to language users, given that it would have to be tacit/implicit knowledge, and thus not testable in the usual verbal way. He takes it that what Dummett is doing is proposing a solution to the difficulties raised by knowledge-ascription for the benefit of those for whom the basic idea of knowledge attribution is unproblematic. Dummett's proposal is that, given the implicit nature of this kind of semantic knowledge, the only way theories which appeal to such knowledge can have any explanatory force is by construing the knowledge in question as knowledge of the assertibility conditions of sentences manifestable in practical language use.

given that we accept that the meaning of a sentence is its truth-conditions, and given that we know the meaning of sentences, it follows that we know the truth-conditions of sentences in the propositional sense of knowledge-*that*, say, the truth-conditions of sentence *S* are *P*. Devitt (1984) asserts, however, that this is "not so much an argument as a play on words" (p. 270). For the argument to work, he claims, "[s]ome entity has to exist which is both the meaning of *S* and the truth conditions of *S*, such that if *X* is acquainted with the one, he is acquainted with the other..." (Devitt, 1984, p. 271). That this is the case is not at all obvious and requires an argument rarely to be found in the literature. Indeed, there are acceptable ways of interpreting the slogan that the meaning of a sentence is its truth conditions without positing the existence of either the meaning or the truth-conditions of a sentence. *Prima facie,* therefore, it can be the case that the meaning of a sentence is its truth conditions, without it being the case that knowing the meaning of a sentence consists in knowing *that* the truth conditions of the sentence are such-and-such.

Dummett's thesis faces other difficulties. Harman (1967), for example, has highlighted the problem that looms when—given the assumption that a speaker must, on the propositional account, be capable of representing to herself the truth-conditions of sentences—we ask in what does *this* ability consist and find ourselves facing a regress. Devitt and Sterelny (1989) have also raised the possibility, inconceivable on Dummett's view, that a speaker could be competent in all parts of a natural language excluding the part which counts as the semantic vocabulary for the language. And, as in the syntactic case, it would appear that the propositional semantic knowledge in question is unacceptably obscure in that it appears neither to (a) dispose a speaker to assent to the propositions expressing it when prompted, nor (b) stand in the kinds of internal relations to other beliefs that we would expect propositional knowledge to do.

4. Linguistic competence as know-how?

The opponents of the propositional, or knowledge-that, view about linguistic competence have urged the alternative proposal that linguistic competence is to be construed entirely as know-how. On this view, the ability to speak a language does not consist in the grasp of a set of propositions; speaking a language is something we learn to do on the basis of practical experience. It is a skill or ability akin to bike riding, playing tennis, knitting, and the like. Devitt and Sterelny (1999), for example, write:

> ...competence in a language does not consist in the speaker's semantic propositional knowledge of or representation of rules. It is a set of skills or abilities... It consists in the speaker being able to do things with a language, not in having thoughts about it. Understanding a language no more involves having propositional knowledge... than being able to ride a bicycle involves having propositional knowledge of a mechanical sort... (p. 187)

Richard Kirkham (1989) has actually attributed this view to Dummett. According to Kirkham, in the passages in which Dummett appears to his critics to be endorsing the propositional assumption, he is in fact endorsing the idea that the semantic knowledge possessed by speakers "consist[s] in abilities which can be *represented* with propositions" (Kirkham, 1989, p. 213–214). While Devitt (1984) takes it that Dummett's use of the words 'implicit' and 'tacit' in describing semantic knowledge are indicative of Dummett's concern to qualify and soften his propositional thesis, Kirkham insists that Dummett uses them to deny the propositional thesis altogether. Take a passage in Dummett such as this:

> A theory of meaning will, then, represent the practical ability possessed by a speaker as consisting in his grasp of a set of propositions... In general, it cannot be demanded of someone who has any given practical ability that he have more that an implicit knowledge of those propositions by means of which we give a theoretical representation of that ability. (Dummett, 1993b, p. 36)

It is not propositional knowledge possessed by the speaker that Dummett is describing, but rather non-propositional knowledge—know-how—which can be represented by "experts" as *if it were* propositional knowledge. Thus, Kirkham: "To have *implicit* knowledge of the proposition that *X* is to have a bit of non-propositional knowledge—an ability—which can be *represented* with the proposition that *X*" (Kirkham, 1989, p. 213).[4] Because the knowledge of the propositions expressing the speakers' semantic abilities is implicit knowledge, it cannot be propositional. Suppose, for example, that someone has the ability to type but is not capable of verbally specifying where on the keyboard certain letters are. This person has a practical ability which does not consist in propositional knowledge but which can be represented, or modeled, as propositional knowledge—as the grasp of a set of propositions. Kirkham (1989) says that Dummett would describe this "sort of epistemic relationship I have with these propositions as 'implicit knowledge,' meaning I do not really know *them* at all, but it is as though I did" (p. 212).

Many theorists acknowledge that while the proponents of the know-how view have raised powerful objections to the propositional thesis, their alternative of modeling competence as a practical skill seems equally unsatisfactory. For example, anti-Chomskyans are fond of highlighting the apparent discrepancies between the postulated propositional knowledge of language and ordinary cases of propositional knowledge, but we can just as easily recognize discrepancies between a person's ability to speak a language and other ordinary skills or abilities. Perhaps most significantly, language-learning is a skill acquired, apparently, without directed practice, and this distinguishes it from skills as ordinarily construed. Further, when a person has purely practical knowledge of how to perform an activity, typically the activity can be recognized or identified by the person prior to their acquiring the ability to engage in it. Yet there is no way in which the first language use can be identified in advance by a novice hoping to learn and become proficient at it. As

[4] The claim that this is Dummett's position finds support in Dummett (1993a), p. 131–134.

Bar-On (1996, p. 149) notes, "[m]ost of what we learn in learning a language concerns *what* to do, not simply how to do something we can already identify prior to having the ability." Consider riding a bicycle as a contrast case. In that case it is clear to the learner what the skill is that she is trying to acquire and how to go about trying to learn it. She has some basic physical skills that are required to make a start; she can get her body onto the bicycle, move her legs to get the pedals turning, and so on. It is not clear that anything like this is true in the case of language. A potential speaker has the auditory capacities to hear speech sounds, but how, on the know-how theory, does the child get a foothold in those sounds? It is well-known that the physical stimulus is continuous. Separation into linguistically relevant parts is a substantial task in itself (Cole and Jakimik, 1980), and this problem is prior to any consideration of categorising those sounds grammatically. Further, although it is natural to suppose, as Locke and the tradition following him did, that words for objects can easily he identified contextually by children, that does not appear to be the case (Gleitman, 1994). A fortiori, it is entirely unclear how syntax will be picked up and practiced by the child until the skill of sentence comprehension and production is fully developed. In the case of ordinary skills, the practitioner must be in a position to locate herself in the skill domain and have the basics to begin practice. In the case of language, it is very hard to see how this is possible. If language is a skill, it appears that much more has to be presupposed about the skill-user's initial capacities than in the case of ordinary skills. It is hard, therefore, to see what a theory of the acquisition of the know-how would look like. We will return to this point below in the discussion of Gibson's view of perception.

When the Chomskyans and Dummettians advanced the knowledge-that view, they recognized an important truth, namely that linguistic competence is qualitatively different from practical competence of the usual sort. This insight retains its validity despite the failure of the attempt to ground linguistic competence in propositional knowledge.[5] In what follows we describe a different approach which, if successful, promises to provide a detailed model for linguistic knowledge. This approach derives from computational theory, especially the computational theory of

[5] An attempt to develop a more nuanced approach to the question of speakers' knowledge has been made by Dorit Bar-On (1996) who aims to transcend the "simple-minded practical/propositional knowledge dichotomy" (p. 148). There are activities a person can be competent in where the competence appears to consist in a combination of practical and propositional knowledge—both know-how and knowledge-that. Competence in playing chess is, on the face of it, a practical ability. But it also involves *explicit* knowledge of propositions concerning the rules of the game and strategies for winning. These propositions can be consciously imparted to others until they have, on the basis of internalizing the set of rules expressed by those propositions, acquired the practical ability to play chess. This contrasts with other cases of non-propositional, practical knowledge, such as the knowledge of how to cycle. The latter does not usually involve conscious knowledge of the propositions detailing what to do at each stage of the cycling process. It is pure know-how (Bar-On, 1996). Bar-On suggests that linguistic competence plausibly consists in a combination of practical and propositional knowledge. Indeed, she suggests that this should be taken to be Dummett's view. She cites Dummett's (1991) conclusion "that we must see knowledge of language as lying between two extremes. At the one extreme lies *explicit theoretical knowledge*... At the other extreme lies 'simple practical knowledge of how to do something which has to be learned'" (Bar-On, 1996, p. 148).

vision. We begin by reviewing computational vision briefly and then indicating how it can be transposed to language. We note at the outset that our discussion is programmatic and that our interest is only in sketching a model. The details of the model remain to be developed.

5. A different approach: computational theory

The elements of our account are as follows:

(1) We focus on a language mechanism or module rather than on the speaker of the language. The account is thus meant to apply at the sub-personal level of cognitive theory.
(2) The account says that linguistic competence is to be analyzed in terms of the computational capacities of this module that have developed in the course of evolution. They are not learned by the speaker.
(3) The computational capacities of the module are to be understood in part in terms of the *natural constraints* on the set of computations embodied by the language module. These constraints are not *represented* in the language module. Rather, they are a propositional expression of the facts about the external world—in this case, the linguistic environment—that have to be true if the module is to function successfully.
(4) Linguistic competence is not know-how because language is not learned the way skills are. But neither is it based on knowledge-that because the propositions that represent at least the rules of language are not represented by the speaker himself but only by linguistic theory.

In the next sections we discuss computational theory as it appears in the theory of vision and then apply it to the case of language by developing each of the claims above in some detail.

5.1. Background

Computational theory is most developed in the domain of vision science due to work done by a number of theorists (many at MIT), most famously David Marr and his co-workers (see Marr, 1982; see also Horn, 1986). A clue that computational vision may provide a model for a third way in the theory of language is that there is a sense in which computational vision is a third way between the two traditional approaches to vision theory. Since Helmholtz (1866/1965; see also Gregory, 1998), it had been natural to think of perception as a process of *unconscious inference*. According to this view, the visual system is presented with a stimulus that is ambiguous, and it has to infer the properties of the environment from that stimulus. For example, in a natural environment things move. Suppose there is a motion stimulus in the lower left quadrant of one's visual field and in the upper right quadrant. Do these stimuli represent two moving objects or a single one? Since the stimulus itself is ambiguous, the perceiver must make an inference in order to answer the question.

Helmholtz thought that the inference had the following rough form. First one represents a general rule or law about the relation of the external environment to some feature of the stimulation effect on the observer. So, for example, one might represent the relation between the way in which moving objects in the environment bring about particular characteristic patterns on the retina that are different from the patterns brought about by disconnected stimuli. The second premise in the argument would be a representation of the particular set of patterns occurring on the retina over a particular interval of time. The conclusion would then have as its content a representation of the motion of an object that brought about the changing patterns on the retina. In this way one uses general information about the relation of the world to the perceiver to work backwards to the properties of the world that caused a particular effect on the perceiver. This general strategy is typical of the way vision scientists of various styles have approached vision. Representing the external environment is accomplished by working backwards form the properties of the retinal stimulation to the properties of the environment that cause that stimulation. For this reason, the problem to be solved in vision is often referred to as the *inverse problem*, or the *problem of inverse projection* (see Aloimonos and Rosenfeld, 1991).

Later theorists shared with Helmholtz the basic outlook on the inverse problem, but some differed on the question of the origin of the perceiver's knowledge of the general laws governing the relations between features of the environment and patterns of retinal stimulation. As an empiricist, Helmholtz believed that these general patterns of relation must be learned by the perceiver in the course of early experience. In contrast, theorists such as Rock (1983) and Shepard (1994) are rationalists who hold that these laws come to be represented innately by visual systems in the course of evolution (see Kubovy and Epstein, in press).

The second approach to vision that is relevant to computational theory is the *ecological approach* due to J.J. Gibson (1979). As against the traditional theory of unconscious inference, Gibson argued that the input to vision is not merely a brute stimulus but carries information. That is to say, the view that the stimulus is intrinsically ambiguous and has to be disambiguated is, on Gibson's view, false. The information needed to perceive the environment is actually in the stimulus, often in the form of higher-order variables, and available to be "picked up" by the perceiver. Perception is thus *direct,* on Gibson's view; it is not mediated by an inferential intermediary.

A familiar case is the perception of size. As an object recedes from a viewer, the image it creates on her retina decreases in area; a small object close up and a large object far away may occupy exactly the same area on the retina. The visual system must, therefore, be able to distinguish these cases. This is a typical problem of vision in which ambiguity has to be eliminated, and the traditional view held that only an inferential process could accomplish this. Gibson's view, in contrast, is that there is information in the environment that can do the work for the visual system (see Gibson, 1950; Sedgwick, 1983). For example, one source of information about size is texture. A second source of information is more specific and follows from what is known as the *horizon principle*. According to this principle, the ratio of an object perceived to be below the line of the horizon and above the line of the horizon carries information about the object's size. Objects with the same ratio are perceived to have the same height. These cases, according to Gibson, are typical. General principles do not have to be employed to see the properties of the environment. Information about the environment is present in the stimulus itself.

At the beginning of Marr's (1982) classic work, *Vision*, he assesses Gibson's position both sympathetically and critically. It is a crucial feature of the computational view that there is information contained in the stimulus. In this respect, Gibson is an important influence on computational vision. However, what Gibson fails to address, according to Marr, is the crucial question of how exactly we "pick up" this information. According to Gibson, light hits the eye and the information is made available to vision. But how?

Marr's objection can be expressed more clearly in the following way. Consider again the case of depth from texture. It is quite true that texture signals depth and can therefore provide information about size, *but how does one come to see a surface as having an ever finer texture to begin with?* Before a perceiver can have this percept, the visual stimulus has to be interpreted correctly as indicating a surface with this texture. Only then can the information about texture be made available to the perceiver. But it appears that the problem of ambiguity arises here again because it is unclear how information about texture is picked up. The stimulus from the surface could, it seems, be as easily interpreted as indicating a large number of objects of ever smaller sizes stacked one on top of the other. In order to see the scene as one of changing texture, depth itself must be perceived. But seeing depth was the problem to be solved! Similarly, once one can see a scene as having objects and a horizon ratio, that information can be used to see size correctly, but how does the visual system disambiguate the visual stimulus sufficiently to produce a coherent percept of this scene in the first place?

Recall that in our discussion of the know-how view above, we noted that a prima facie problem with that view is that it is hard to see how a potential practitioner of the skill of language could get a foothold in the skill domain that would be sufficient to begin to practice and become proficient in the skill. Marr's point here is a closely related one. Information pickup may be possible once the perceiver is located in a visual scene that is disambiguated. For example, if the visual system can achieve figure-ground segregation, then the horizon principle is of considerable use. But the horizon principle is useless in the absence of a coherent visual scene. Marr's critique of Gibson is that Gibson has ignored this all-important first-step in visual perception.

5.2. Computational vision

Marr's answer to the question of how information is picked up in vision embraces Gibson's view that the visual stimulus is information-bearing. His problem—the problem of computational vision generally—is to explain how this information is extracted by the visual system. This is a problem because the visual stimulus does indeed seem to be ambiguous. This ambiguity can be made clear in the following way. Think of the input to the visual system at a single moment as a pattern of bright and dark spots on the two-dimensional grid constituted by differing activations of the carpet of photoreceptors on the retina. Over time, the visual system will have a succession of such patterns to work with and nothing else. That is, the "film" made up of the set of frames of two-dimensional patterns of light and dark spots is all that vision has to work with to extract information about three-dimensional shape, texture, motion, color, and so on (see Ullman, 1979 for this picture applied to the problem of motion).

Notice that the difficulty posed for vision is that there are *too many* possible answers to the question of what the patterns of light and dark spots signify. Indeed, there are indefinitely many such answers. Is the dark spot in one stimulus frame part of an edge or part of a dark surface? Does the changing pattern of light and dark from one frame to the next represent a moving stimulus or a change in the conditions of illumination? And so on. A well-known problem of this type in the domain of motion perception is the *correspondence problem*. Consider the series of frames in Fig. 1. Suppose we know that the stimuli in the frames represent dots on a moving surface. Before the visual system can determine the nature of the motion of the object, it must decide which dot in frame 1 corresponds to which dot in frame 2. Different patterns of correspondence will produce different percepts of movement, and nothing in the stimulus itself gives any indication of how the correspondence problem should be solved. Until the number of possibilities is restricted, the information in the stimulus cannot be extracted. The correspondence problem is one instance of the inverse problem described above, and visual problems are typically of the same form. The inverse problem in general, therefore, is insoluble as stated.

Problems which are insoluble because there are too *many* possible solutions are said to he *ill-posed* (Poggio and Koch, 1985), and they can only be solved by making assumptions about the problem that restrict the domain of solutions—ideally to a single option. A natural solution to the correspondence problem, for example, would be provided by making a *nearest neighbor* assumption according to which each point in a frame corresponds to the point in the succeeding frame which is nearest to it in two-dimensional space (Ullman, 1979). Whether or not the visual system actually makes use of this assumption to solve the correspondence problem,[6] the solution makes clear what sort of principle is necessary for visual problems to be solved. The visual system must make assumptions that restrict the domain of possible interpretations of the stimulus. With this restriction in place, a solution becomes

[6] It almost certainly doesn't. What determines how the correspondence problem is solved is which elements are perceived to have the greatest "affinity," or similarity, and there are a number of dimensions of similarity that may affect visual phenomenology. The nearest neighbor principle is supposed to represent the solution to the correspondence problem when no other dimensions of affinity are available. There is evidence, however, that the nearest neighbor principle is not sufficient to explain visual behavior even in this case. It can be shown that under certain conditions the visual system opts for correspondences over greater spatial distances if those correspondences issue in information about the three-dimensional structure of moving objects (Gold et al., unpublished MS).

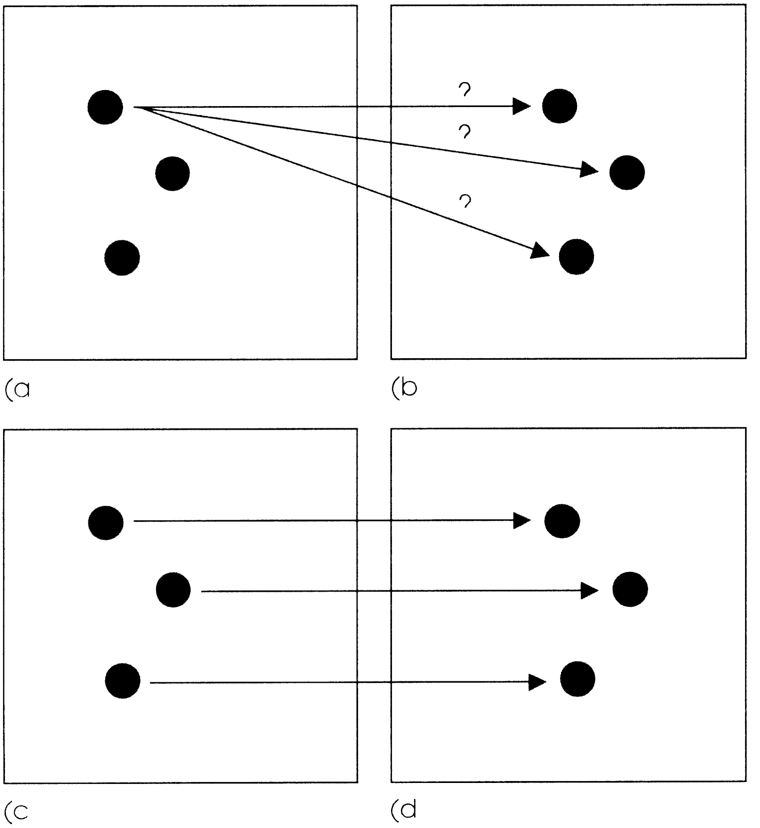

Fig. 1. A series of frames of a "film" of a moving object. It is ambiguous which dots in (a) correspond to those in (b). In (c) and (d) the "nearest neighbour" principle is applied to resolve the ambiguity.

possible. Man called these assumptions *natural constraints*. They are constraints because they reduce the number of possible solutions of an ill-posed visual problem; they are natural because, on Marr's view, many of them express general facts about the natural environment being perceived. For example, the nearest neighbor principle is a natural constraint because it embodies a fact about the kinematics of solid objects (see Shepard, 1994).

How then does the visual system extract information from the stimulus? On the computational view, vision makes use of assumptions in the form of natural constraints to deal with the different dimensions of visual perception (motion, shape, etc.). Each stimulus is then disambiguated by interpreting it in accordance with these constraints to yield unambiguous information about the object or scene. The application of a natural constraint to an ill-posed visual problem is thus much like Helmholtz's application of a general principle to an ambiguous stimulus. However, while Helmholtz took the process to be one performed unconsciously by a perceiver, the computational view takes it to be accomplished sub-doxastically by a visual mechanism. The process is thus computational rather than inference properly so called.

5.3. Natural constraints

What does it mean to say, however, that the visual system "makes assumptions" in the form of natural constraints? One possibility is that these assumptions are propositional in form and are explicitly represented within the visual system though they will not be accessible at the doxastic level to the perceiver herself.

This is the Helmholtzian line, but computational vision does not adopt it. The computational solution is quite different and represents one of the central contributions of computational vision to modern psychology and philosophy of mind. According to computational vision, the visual system does not represent natural constraints at all. Rather, these constraints are descriptions of the *environment* and have to do with the way in which the visual system evolved. In order to see what this means and why it is relevant, imagine the evolution of the visual system of a particular species. The species makes its living in a particular niche, and some of the properties of this niche will be usefully represented visually. The color of ripe fruit against the background of green leaves, for example, is thought to have been an important feature of primate visual evolution (Osorio and Vorobyev, 1996). In this niche, therefore, the ability to detect color will constitute an adaptive advantage. Now imagine two hypothetical animals in this environment. Suppose both of these animals have genetic mutations that alter the properties of their visual systems. In the first animal, this change causes it to be selectively sensitive to different wavelengths and thereby to be able to detect surface color. The second animal exhibits an alteration that allows it to detect the polarization of light. Because the former change is adaptive in this environment but the latter is not (though it may be in other environments, such as that of the bee), the first animal, let us suppose, survives, and the second does not. In this way color vision establishes itself in the species. Notice, however, that the development of color vision makes no reference to any assumptions made by the visual system or to any unconscious or implicit knowledge possessed by the animal. The alteration in vision may be nothing more than the evolution of a new cone pigment in the animal's retina. What makes color vision possible is a random change in the animal's visual repertoire in an environment that possesses a particular property the detection of which is adaptive.

When the vision theorist comes to consider the algorithm embodied in primate vision for detecting color, she may hypothesize that the algorithm is one that works *as if* it had a goal and that goal is to distinguish ripe fruit against green backgrounds. This may lead to the hypothesis that one of the ways in which the visual system functions is as if it assumed that wavelength contrast signaled a difference between distinct objects in a scene. This assumption, in turn, will contribute to the visual system's ability to disambiguate the stimulus. It would be a mistake, however, to infer from this that the visual system itself represents that assumption. The constraint is a description of the environment that shaped vision and not an assumption of the visual system itself. It thus represents a concern of vision *theory* rather than of the perceiver himself. Computational vision thus takes the scientist out of the brain and puts him back into the lab where he belongs.

Given that the assumptions captured by natural constraints are not represented anywhere outside of vision theory, one must characterise *successful* computational strategies in a somewhat roundabout way. For example, to say that motion perception makes use of the natural constraint expressed by the nearest neighbor principle means that if the natural environment in which vision occurs actually satisfies the kinematic principle expressed by the principle, then the computation carried out by the visual system will typically issue in a veridical representation of the motion of objects. If the environment is not correctly described by the constraint, then the percepts produced by the visual system are likely to be inaccurate. If one further assumes that evolution governs the way in which vision operates, then one can suppose that there will be pressures on vision to develop computations that in fact operate under natural constraints that correctly describe the environment. One can thus explain how vision produces veridical (or at least adaptive) percepts without having to *represent* natural constraints at all. As Kubovy and Epstein (in press) put it, according to the Helmholtzian style of explanation, "assumptions and knowledge are mental contents that are active in the perception process." Natural constraints, in contrast, "are neither (1) lodged in the mind nor (2) are they active constituents in the perceptual process. They are the conditions which the world must satisfy if the computational algorithms are to go through" (p. 2).

Computational vision thus takes visual perception to be explained, in general terms, with reference to the capacities of a visual mechanism that has developed (both phylogenetically and ontogenetically) so as to function successfully just in case the environment satisfies the natural constraints articulated by computational theory. Nowhere outside of vision theory itself are the principles satisfied articulated or represented. The ability of a perceiver to see, then, is reduced to, or explained by, the capacities of a visual module or mechanism the successful function of which is described in terms of the satisfaction of the natural constraints identified by vision theory but which does not represent those constraints either explicitly or implicitly.

5.4. Overview of the debate

We have seen three approaches to the inverse problem in vision, and a brief overview will highlight the relevant issues. Because the stimulus is ambiguous, almost all vision theorists have thought that perceivers must have some sort of assumptions or prior knowledge in order to achieve a coherent visual representation. Gibson, in contrast, dealt with the problem by denying that there was a problem to begin with. The inverse problem, on his view, can be seen to be a pseudo-problem once one recognizes that there is information in the ambient array that is sufficient for vision. The light stimulus is not ambiguous, and the problem of vision is not ill-posed. One does not, therefore, have to posit either assumptions or natural constraints to solve it.

The debate between Helmholtz and Gibson mirrors the debate between the knowledge-that and know-how camps. Like Gibson, the know-how theory takes the linguistic environment to be sufficiently accessible to the potential practitioner that language can be acquired with nothing more than the basic perceptual and motor skills. The knowledge-that position, in contrast, sees the linguistic environment, like the visual environment, as inaccessible in the absence of prior assumptions or knowledge.

We have seen that the computational approach to visual perception accepts the Helmholtzian claim that the visual stimulus is ambiguous and cannot be disambiguated in the absence of any constraint. However, it rejects the idea that these constraints must be embodied as mental representations in the perceiver. Like Gibson, computational vision holds that the perceiver is in a position to pick up information—as long as he has a visual mechanism that has been shaped by evolution to function as if it makes assumptions about what the environment is like.

If we are correct in suggesting that there is an isomorphism between the Helmholtz–Gibson debate and the debate about linguistic knowledge, then a computational approach to language seems promising. We turn to this suggestion next.

6. Language: a proposal

We have already noted Cowie's (1999) observation that Chomskyan theory is not monolithic but rather composed of at least five distinct theses of varying degrees of generality and for which, on Cowie's view, there are differing degrees of evidential support. The one of interest to us is (R), the claim that linguistic competence requires contentful mental states representing rules of language. As we have said, this proposition is the most general of the Chomskyan theses and, as the core tenet of cognitive science, the one that is most widely accepted.

Cowie suggests that confusion about Chomskyan linguistics arises because it is often thought that these different theses have to be accepted or rejected as a package, and she argues that one can pick and choose among them. The kind of objection we rehearsed in Section 2 advocates rejecting (R) with respect to language, and one can see how it might be natural to think that in rejecting (R) one must also reject the other Chomskyan theses. This is because the other theses seem to qualify the kinds of representations posited by (R). If one rejects the claim that there are representations to begin with, then the rest of Chomskyan theory seems to be dispensed with at the same time.

Our claim is that this is not the case. One can reject (R) and retain the remaining Chomskyan theses if one replaces (R) with a principle based on the notion of natural constraints developed above. Consider, therefore, a principle expressing the natural constraints view:

> (NC) Explaining language mastery and acquisition requires the postulation of a set of computational mechanisms characterized in part by a set of natural constraints which describe features of the linguistic environment.

On the model of visual constraints, we take linguistic natural constraints to be a description of important features of the linguistic environment that restrict the number of possible solutions to ill-posed linguistic problems. This restriction permits the hearer or speaker to produce a single solution to problems of linguistic ambiguity that are consistent with the grammar of their language.

What might a natural constraints view look like? First, we will assume that, like vision, language is subserved by a module or a set of modules each of which is responsible for some aspect of linguistic ability. For the purposes of illustration, we will focus on knowledge of grammar, and we will suppose (contrary to fact) that there is a single module subserving syntactic comprehension (see, e.g. Friedmann, 2001). Now consider a child developing within a particular linguistic community. The child is exposed to sound patterns that he must eventually be able to interpret in accordance with syntactic rules and that will form the basis of his own syntactically well-formed utterances. The task for the module is, therefore, to function in such a way as to be capable of forming representations of the utterances that accord with the principles of the grammar of the language. The knowledge-that theorist (the defender of R) holds that that ability to interpret utterances is the learned or innate representation of the syntactic structures of his language. The knowledge-how theorist takes the child's ability to be syntactic know-how (whatever exactly that amounts to) manifested in his capacity to understand utterances.

The problem here is much like that faced by vision. The comprehension module faces the task of finding an interpretation of the incoming auditory stimulus that is veridical in the sense that it accords with the syntactic rules of the language. Because a very large number of interpretations are possible, the module must be able to restrict the interpretations so as to make a solution possible. On a computational view, this restriction is achieved by having the module function in accordance with a set of natural constraints. A computational approach thus does away with (R) but does not do away with the idea that linguistic ability requires built-in interpretive mechanisms. The approach will hold (1) that adequate syntactic production by a speaker is to be analyzed in terms of the capacities of the relevant module and not in terms of a property of the speaker herself; (2) that 'adequacy' is to be explicated as performance by the module in accordance with a set of constraints or principles that reduces the number of syntactic interpretations of the linguistic stimulus to make it comprehensible (presumably some of the principles of the grammar itself); and (3) that these principles are not represented in any form by the module itself but only by the linguist's theory. The computational approach thus makes use of the familiar distinction between following a rule and acting in accordance with a rule. The module must satisfy the constraints necessary to interpret the linguistic stimulus but because these rules are not represented by the module, it should not be said to obey them. It rather operates in accordance with the rules as articulated by the theory of the grammar of the language.

6.1. Chomskyan theory and (NC)

We return, finally, to Chomskyan theory. (R) can be detached from the theory. What does the theory look like if we replace it with (NC)? In our view, (NC) is not only an acceptable option, it is one to be preferred on the grounds that it makes Chomskyan linguistics more coherent. Recall Cowie's (1999) taxonomy of Chomskyan principles other than (R):

Biological Boundedness (BB): There are constraints on the kinds of thought that human beings can entertain;
Domain Specificity (DS): Learning a language requires that the learner's thoughts are constrained by principles that apply to the domain of language only;
Innateness (I): Constraints on language are innate; and
Universal Grammar: The constraints and principles referred to in the principles of domain specificity are identical to the principles of Chomsky's Universal Grammar.

If (NC) is true, then one expects that at least some of the constraints that are relevant to language-learning wifi be domain-specific and thus that some version of (DS) will be true. Further, (NC) is set against the background view that evolution has exerted selective pressure on the language system to function so as to solve problems that the speaker faces in the linguistic environment. It follows that the properties of the language system that satisfy those constraints are innate features of that system. That is, if we assume (NC) and the appropriate background assumptions, then (I). Finally, if (NC) is true, then there are indeed constraints on the kinds of "thoughts" that can be entertained. (NC) is based on the idea that in order to solve ill-posed problems, one must rule out incorrect interpretations. Because these interpretations are ruled out by the operation of the linguistic system, they can be construed as linguistic states that are out of bounds—as linguistic "thoughts" that one cannot have. Indeed, these thoughts *must* be ruled out if language is to be learned at all.

It seems, therefore, that while nothing of Chomskyan theory follows from (R), a good deal of it follows from—or is at least in the spirit of—(NC). Since coherence may be thought a theoretical virtue, (NC) is to be preferred over (R). To the extent that it constitutes a third way between the knowledge-that and the knowledge-how accounts of language, it may also be preferable on broader philosophical grounds as well.

Acknowledgements

We are grateful to Max Coltheart and Natalie Stoljar for many helpful comments on earlier drafts of this paper.

References

Aloimonos, Y., Rosenfeld, A., 1991. Computer vision. Science 1249–1254.
Bar-On, D., 1996. Anti-realism and speaker knowledge. Synthese 106, 139–166.
Chomsky, N., 1965. Aspects of a Theory of Syntax. MIT Press, Cambridge.

Chomsky, N., 1969. Linguistics and philosophy. In: Hook, S. (Ed.), Language and Philosophy: A Symposium. New York.
Cole, R., Jakimik, J., 1980. In: Cole, R. (Ed.), Perception and Production of Fluent Speech. Erlbaum, Hillsdale, pp. 133–163.
Cowie, F., 1999. What's Within. Oxford University Press, New York.
Davies, M., 1987. Tacit knowledge and semantic theory: can a five per cent difference matter? Mind 96, 441–462.
Davies, M., 1989. Tacit knowledge and subdoxastic states. In: George, A. (Ed.), Reflections on Chomsky. Basil Blackwell, Oxford.
Devitt, M., 1984. Realism and Truth. Oxford.
Devitt, M., Sterelny, K., 1989. Linguistics: what's wrong with 'the right view'. In: Tomberlin, J.E. (Ed.), Philosophical Perspectives 3: Philosophy of Mind and Action Theory. Atascadero, Ridgeview.
Devitt, M., Sterelny, K., 1999. Language and Reality: An Introduction to the Philosophy of Language. Blackwell, Oxford.
Dummett, M., 1973. Frege: Philosophy of Language. Duckworth, London.
Dummett, M., 1982. Realism. Synthese 52, 55–112.
Dummett, M., 1991. The Logical Basis of Metaphysics. Harvard University Press, Cambridge.
Dummett, M., 1993a. Language and truth. In: The Seas of Language. Oxford University Press, Oxford.
Dummett, M., 1993b. What is a theory of meaning (II). In: The Seas of Language. Oxford University Press, Oxford.
Evans, G., 1981. Semantic theory and tacit knowledge. In: Holtzman, S., Leich, C. (Eds.), Wittgenstein: To Follow a Rule. Routledge and Kegan Paul, London.
Fodor, J., 1981. Introduction: some notes about what linguistics is talking about. In: Block, N. (Ed.), Readings in Philosophy of Psychology, vol. 2. Harvard University Press, Cambridge.
Fodor, J., 1983. Modularity of Mind. MIT Press, Cambridge.
Friedmann, N., 2001. Agrammatism and the psychological reality of the syntactic tree. Psycholinguist Research 30, 71–90.
George, A., 1984. On Devitt on Dummett. Journal of Philosophy 516–527.
Gibson, J.J., 1950. The Perception of the Visual World. Houghton Mifflin, Boston.
Gibson, J.J., 1979. The Ecological Approach to Visual Perception. Houghton Mifflin, Boston.
Gleitman, L.R., 1994. Words, words, words... . Philosophical Transactions of the Royal Society of London, B Biological Science 346, 71–77.
Gold, I.J., Anderson, K.C., Siegel, R.M. A violation of minimal mapping in structure-from-motion. Unpublished MS.
Gregory, R., 1998. Eye and Brain, fifth ed. Oxford University Press, Tokyo.
Harman, G.H., 1967. Psychological aspects of the theory of syntax. Journal of Philosophy 64, 75–87.
Helmholtz, H.L.F.v., 1866/1965. Hermann Ludwig Ferdinand von Helmholtz (1821–1894) on empiricism in perception. In: Hermstein, R.J., Boring, E.G. (Eds.), A Source Book in the History of Psychology. Harvard University Press, Cambridge.
Horn, B.K.P., 1986. Robot Vision. MIT Press, Cambridge.
Kirkham, R.L., 1989. What Dummett says about truth and linguistic competence. Mind 390, 207–224.
Kubovy, M., Epstein, W., in press Internalization: a metaphor we can live without. Behavioral 1 and Brain Sciences 24(3).
Marr, D., 1982. Vision. Freeman Publishers, San Francisco.
Osorio, D., Vorobyev, M., 1996. Colour vision as an adaptation to frugivory in primates. Proceedings of the Royal Society of London B 263 (1370), 593–599.
Peacocke, C., 1986. Explanation in computational psychology: language, perception and level 1.5. Mind and Language 1, 101–123.
Poggio, T., Koch, C., 1985. Ill-posed problems in early vision: from computational theory to analogue networks. Proceedings of the Royal Society of London B 226, 303–323.
Rock, I., 1983. The Logic of Perception. MIT Press, Cambridge.
Sedgwick, H.A., 1983. Environment-centered representation of spatial layout: available visual information from texture and perspective. In: Rosenthal, A., Beck, J. (Eds.), Human and Machine Vision. Academic Press, New York.

Shepard, R.N., 1994. Perceptual-cognitive universals as reflections of the world. Psychonomic Bulletin and Review 1, 2–28.
Stich, S.P., 1971. What every speaker knows. Philosophical Review 80, 476–496.
Ullman, S., 1979. The Interpretation of Visual Motion. MIT Press, Cambridge.