low-level perception and high-level phenomena of stereotyping and prejudice that affect society as a whole.

References

1 Malpass, R.S. and Kravitz, J. (1969) Recognition for faces of own and other race. *J. Pers. Soc. Psychol.* 13, 330–334
2 Shapiro, P.N. and Penrod, S. (1986) Meta-analysis of facial identification studies. *Psychol. Bull.* 100, 139–156
3 Anthony, T. *et al.* (1992) Cross-racial facial identification: a social cognitive integration. *Pers. Soc. Psychol. B.* 18, 296–301
4 Furl, N. *et al.* (2002) Face recognition algorithms and the other-race effect: Computational mechanisms for a developmental contact hypothesis. *Cogn. Sci.* 26, 797–815
5 Sigala, N. and Logothetis, N.K. (2002) Visual categorization shapes feature selectivity in the primate temporal cortex. *Nature* 415, 318–320
6 Freedman, D.J. *et al.* (2001) Categorical representation of visual stimuli in the primate prefrontal cortex. *Science* 291, 312–316
7 Goldstone, R.L. (1998) Perceptual learning. *Annu. Rev. Psychol.* 49, 585–612
8 Chiroro, P. and Valentine, T. (1995) An investigation of the contact hypothesis of the own-race bias in face recognition. *Q. J. Exp. Psychol.* 48A, 879–894
9 Cross, J.F. *et al.* (1971) Sex, race, age, and beauty as factors in recognition of faces. *Percept. Psychophys.* 10, 393–396
10 Lavarkas, P.J. *et al.* (1976) A perspective on the recognition of other-race faces. *Percept. Psychophys.* 20, 475–481
11 Ng, W. and Lindsay, R.C.L. (1994) Cross-race facial recognition: failure of the contact hypothesis. *J. Cross. Cult. Psychol.* 25, 217–232
12 Swets, D.L. and Wang, J. (1996) Discriminant analysis and eigenspace partition tree for face and object recognition from views. *Proceedings of the Second International Conference on Automatic Face and Gesture Recognition*, pp. 192–197, IEEE Computer Society Press
13 Light, L.L. *et al.* (1979) Recognition memory for typical and unusual faces. *J. Exp. Psychol. Hum. Learn.* 5, 212–228
14 Levin, D.T. (2000) Race as a visual feature: using visual search and perceptual discrimination tasks to understand face categories and the cross-race recognition deficit. *J. Exp. Psychol. Gen.* 129, 559–574
15 Levin, D.T. (1996) Classifying faces by race: the structure of face categories. *J. Exp. Psychol. Learn. Mem. Cogn.* 22, 1364–1382

PII: S1364-6613(02)00036-0

An *ABSURDIST* model vindicates a venerable theory

Eric Dietrich

Philosophy Department, Binghamton University, Binghamton, NY 13902, USA

Can relations between concepts in one system be sufficient on their own for locating matching concepts in another system? This question in concept semantics has been contentious. Now, contrary to one important theory, the answer appears to be yes. Robert Goldstone's and Brian Rogosky's new neural-net program, ABSURDIST, has demonstrated what purely theoretical considerations suggested could not be accomplished, showing once again the power of experimental computer modeling.

'The moment of truth is a working program'. This well-known aphorism in artificial intelligence (AI) succinctly captures the fact that a working computer program can wreak havoc with, or provide compelling support for, a psychological or AI theory. Both of these are at work in Robert Goldstone's and Brian Rogosky's new neural net program, ABSURDIST, detailed in their recent, fascinating paper [1]. The theory jeopardized by ABSURDIST is Fodor's theory of concepts and conceptual semantics, a major consequence of which is that relations between concepts in one system are insufficient to establish correspondences with similar concepts in another system. The theory supported by ABSURDIST is the theory that relations between concepts constitute a large part of the meaning or semantics of individual concepts – even a necessary part. To appreciate fully the importance of ABSURDIST, some crucial background is needed.

Corresponding author: Eric Dietrich (dietrich@binghamton.edu).

A computer model confronts a theory

The relational semantics theory called the 'conceptual web theory' (or sometimes the 'conceptual role theory' or 'holism') says that the functional roles that concepts play in the larger cognitive system in which they exist help to determine their content. Fodor's theory, called by Goldstone and Rogosky the 'external grounding account of meaning' (or just 'externalism'), contrasts sharply with the conceptual web account (there are several different varieties of externalism, Fodor's is just the prominent one). On Fodor's view, the meaning of a concept is entirely determined by its causal, informational connection to the external environment in which the cognitive system lives.

The conceptual web theory has a long and distinguished history. It began to flourish in the last century with the work of the philosopher Quine, and really came into its own with the advent of semantic nets in AI. These were systems of interconnected nodes labeled with words, and the meanings of the nodes were determined by the connections, or arcs, between them [2,3]. But some theorists saw semantic webs as lacking any semantics, because the networks were not connected at all to the external environment. They likened such networks to meaningless symbols or strings connected to one another via arcs labeled with more meaningless symbols.

These 'externalists' claimed further that because such networks lacked a semantics, they were useless in understanding the nature of thought. This strong claim followed from the idea that the *essential* feature of thoughts (i.e. of active concepts) is that they have meaning, a semantics. No semantics; no thinking. AI researchers countered that, even if this were true, it wouldn't follow that semantic nets were useless, for they at least modeled how connected webs of concepts could be used in everything from language processing to planning. AI researchers and other 'networkers' also claimed that a kind of semantics was provided, in the minimal sense that the relations making up the networks constrained any information that was received by the network from the outside world. There the matter stalled.

Fodor's entrance into the fray at this time had a major impact. He argued that the relations governing semantic networks could be *no* part of a semantics, not even a constraining part, and that the conceptual web account jeopardized cognitive psychology [4]. First, Fodor pointed out that, in natural systems, no two networks would be identical, as no two networks would connect up their concepts using the exact same arcs in the exact same structure. Then, because a node's meaning is determined by its connections throughout the entire web, it follows that no node in one network is going to mean the same as a node with the same label in another network. From this Fodor concluded that no two humans could think the same thoughts, hence no two humans (or any other cognizer) could have the same concepts. Therefore, on the network theory of meaning, translation between conceptual networks is impossible, which means that cognitive psychology is impossible because psychology's job is to express generalizations about human thinking and that requires expressing generalizations across different humans and their conceptual schemes. (For example, my concept of rock climbing might differ from yours in that my concept is connected to my taking a 100-foot fall while solo climbing Lankin Dome in Wyoming in the late summer of 1975, whereas yours is not. Hence you and I cannot think the same thought about rock climbing, even if the thought is expressed in identical words.) For Fodor, then, only an external account of conceptual semantics allows for identical concepts across individuals [4].

Now, you might object, 'But identity isn't required for picking out the same concepts across different individuals. All that is required is *similarity*.' Fodor has a technical argument that the notion of similarity cannot be made to do the job, which Goldstone and Rogosky explain admirably. (Put simply, it either presupposes the notion of identity or generates an infinite regress.)

The seminal importance of ABSURDIST is now easily stated. ABSURDIST shows how non-identical, but merely similar, concepts in two different systems can be placed in correspondence, thereby enabling a translation of conceptual meaning between the two systems. (ABSURDIST in fact stands for Aligning Between Systems Using Relations Derived Inside Systems for Translation.) Goldstone and Rogosky have therefore shown that none of the dire consequences that Fodor and other externalists see for network theories of meaning follow from embracing a network theory of meaning. And they have shown this in the best way possible: they have actually demonstrated in a working model that relations between concepts in one system are *sufficient* to establish correspondences with similar concepts in another system – precisely what Fodor denies.

The ABSURDIST model

As a simple example, suppose in system A that $D(\mathrm{q,r}) = 7$ (i.e. the distance between node q and node r is 7), $D(\mathrm{r,s}) = 2$, and $D(\mathrm{q,s}) = 9$, and in system B, $D(\mathrm{z,y}) = 3$, $D(\mathrm{y,x}) = 7$, and $D(\mathrm{z,x}) = 10$. ABSURDIST can place these two sets of nodes in correspondence, and finds that node q corresponds to node x, r to y, and s to z (to see this pictorially yourself, write them out in graphical notion using nodes and arcs).

Of course, Goldstone and Rogosky have made some simplifying assumptions in building ABSURDIST. Relying on distances between concepts as the sole relation between them is one such assumption. These assumptions naturally provide some wiggle room for Fodor and other externalists. They can argue that although ABSURDIST works in its domain, it is too constrained, too impoverished, to show that conceptual webs are psychologically real and useful, and that translation between real webs is possible relying solely on internal relations. This objection comes down to a long-term bet on the future. Goldstone and Rogosky are betting that the principles explored in ABSURDIST scale up, that they point to a way to translate between the conceptual webs of real human beings. Fodor is betting that they won't, and for fundamentally a priori reasons. The smart bet would seem to be on ABSURDIST. It is quite robust, handling different kinds of noise and other permutations and alterations. Moreover, ABSURDIST is not restricted to internal conceptual relations as the only way to determine meaning. Almost every cognitive scientist (with the exception of Fodor) thinks that meaning is determined by *both* internal relations and relations with external information coming into the system. In other words, a correct theory of conceptual meaning must combine conceptual web meaning and external meaning. ABSURDIST allows for just this kind of combination by allowing the activation of the nodes to be seeded with many different activations thereby representing externally derived information.

The ABSURDIST future

We do not yet have programs or machines that think. But, thankfully, we do have programs that help us in our difficult theorizing about how to build such machines, by robustly testing our theories. In spite of its name, ABSURDIST leads to the quite sensible conclusion that relations between concepts are quite important to their semantics. But with that settled, many new questions arise.

An important assumption in ABSURDIST is that the notion of distance in one conceptual system is identical to the notion of distance in another system. This is because distance is based on generic similarity. A key question is whether this assumption can be relaxed; that is, can the ABSURDIST algorithm be extended to handle relations between nodes other than distances so defined? One way to get at this question is to ask whether the ABSURDIST algorithm can be extended to symbolic AI, where the relations between nodes are characterized by labeled arcs rather than distance measures? One method for addressing this question is to use abstraction. As the non-distance relations in different conceptual systems need not be even conceptually the same, the way to compare them is to abstract them, leaving out or packing away certain information until a relational match is found. Then, perhaps, a symbolic variant of the ABSURDIST algorithm could be applied. There is a tantalizing but surely not absurd future for conceptual webs.

References

1 Goldstone, R. and Rogosky, B. (2002) Using relations within conceptual systems to translate across conceptual systems. *Cognition* 84, 295–320
2 Shapiro, S. (1979) The SNePS Semantic Network Processing System. In *Associative Networks: Representation and Use of Knowledge by Computers* (Findler, N., ed.), pp. 179–203, Academic Press
3 Rapaport, W.J. (2002) Holism, conceptual-role semantics, and syntactic semantics. *Minds Machines* 12, 3–59
4 Fodor, J. (1998) *Concepts: Where Cognitive Science Went Wrong*, Oxford University Press

doi:10.1016/S1364-6613(02)00042-6

Letters

Human spatial representation derived from a honeybee compass

Martin Giurfa[1] and Randolf Menzel[2]

[1]Centre de Recherches sur la Cognition Animale, Université Paul-Sabatier, CNRS, 118 Route de Narbonne, 31062 Toulouse cedex 4, France
[2]AG Neurobiologie, Institut für Biologie, Freie Universität Berlin, Königin-Luise-Str. 28/30, D-14195 Berlin, Germany

Wang and Spelke [1] reviewed different navigation strategies of animals and their underlying spatial representations and concluded that humans, like insects, apply predominantly egocentric ‘primitive navigational systems’, such as path integration, beaconing and view-dependent recognition of landmarks. Despite their arguing in favour of common spatial representations between humans and non-human animals, Wang and Spelke concluded that the uniqueness of humans resides in their capacity to overcome such primitive navigational systems to construct ‘true geocentric maps of the environment’ [1]. We believe that this reasoning assumes a simplistic view of animal navigation in general, and of insect navigation in particular. We focus on insect navigation, because Wang and Spelke’s arguments were based in part on it, and because it may not be purely egocentric and primitive.

The traditional method for studying insect navigation is to train a bee or an ant along a route, and then release it at a novel site and assess the influence of specific information (e.g. sky compass information or landmarks) on its subsequent steering course. Under such conditions, insects use path integration and visual landmarks en route and between their goals [2]. Much richer navigation strategies have been found in honeybees using different training procedures [3]. Furthermore, the use of new techniques such as harmonic radar allows recording for the first time of complete flight paths of bees in their natural environment [4].

In experiments of Menzel *et al.* [3], bees foraged at a feeder that was regularly rotated around the hive, within its close vicinity. Thus they could not establish a route memory, but were guided by a spatial memory acquired by latent learning during their exploratory flights. Bees trained in this way and released at different novel sites around the hive could return from any location within the distance of their orientation flights without reference to a beacon or a landmark constellation [3]. Bees trained to a single location following conventional route training first applied the navigation vector corresponding to the route learned when released, and only later referred to the ‘general landscape memory’. Thus, bees access the general landscape memory only when they do not have a route memory, or have used it but have not yet arrived at the goal [3].

Furthermore, route-trained bees carrying a transponder enabling harmonic radar to locate them were captured and released at a novel site, either when leaving the feeder to return to the hive, or when arriving empty at the feeder [5]. Both groups of bees first flew their usual vector when released at a novel site, but then headed towards the hive after a phase of circling flights. Again, beacon orientation and navigation according to landscape features were excluded. Most importantly, bees sometimes also decided to fly back to the feeder first rather than directly to the hive. These and additional experiments can be explained by assuming that during their orientation flights, bees learned different locations in their surroundings and attached to the landmarks characterizing these locations local vectors pointing towards the hive. Equipped with such a hive-centered map, bees would be able to perform exactly as they did in the earlier study [3]. However, the fact that bees foraging at a distant and constant feeder could decide to fly first back to the feeder rather than directly to the hive [5] indicates a form of spatial memory in which some geometrical relationships between defined points in space are preserved, in agreement with Tolman’s seminal paper [6].

Wang and Spelke suggest that human navigation must interpreted in the light of egocentric and dynamically updated spatial representations that are common to non-human animals. We agree with this conclusion. But we also think that ignoring the richness of animal spatial representations constitutes an error. The fact that insects can apply navigational strategies more complex than

Corresponding author: Martin Giurfa (giurfa@cict.fr).