Interactions Between Taxonomic Knowledge, Categorization, and Perception

Frédéric Gosselin, Annie Archambault, and Philippe G. Schyns

Department of Psychology

University of Glasgow

58 Hillhead St., Glasgow,

Scotland, UK, G12 8QB

Casual observers glancing at Figure 1 would have little difficulty identifying two faces. Faces comprise parts such as a *mouth*, two *eyes*, a *nose*, *eyebrows*, and so forth, organized in a typical spatial configuration. Even if the observer did not know the particular individuals, s/he should nevertheless identify the stimuli at a general level as male and female faces. Were the identity known, the stimuli might be categorized at a specific level as John and Mary. In yet another context, they might be categorized as neutral and happy faces.

It is clear that people can categorize objects at different levels of a taxonomy, but it is much less clear how these categorizations influence the perception of the object itself, if at all! Few would doubt that perception must influence categorization, but the opposite is still questioned (e.g. Pylyshyn, 1999). This chapter examines the hypothesis that people differently utilize and perceive the information of an object if they have learned to categorize it differently—e.g., at different levels of specificity. We consider the implications of the hypothesis in the domains of face, object and scene categorization and recognition, and examine how it forces us to reconsider the role of similarity in categorization.

Insert Figure 1 about here

Object recognition and categorization research are both concerned with the general question of "what is this object?" To recognize an object as a face or a computer is not very different from assigning it to the *face* or the *computer* category. The problem is to understand how information available in the visual array matches the information representing categories in memory. Both categorization and object recognition research therefore address the same fundamental issues of "what is the organization of information in memory?" "what is the information available in the input?" and "how do memory and input information interact to explain behavior?" (Schyns, 1998).

Despite these fundamental similarities, object recognition and categorization have evolved independently, without much dialogue. This could stem from a difference in focus: Categorization studies typically seek to unravel the rules governing the formation of categories (the idea that the visual attributes *feathers, wings, legs, beak, black*, but also functional attributes such as *fly, lay eggs, lives in the trees* represent a *crow* which is also a *bird*, an *animal* and a *living thing*), whereas researchers in recognition have mostly sought to understand the perceptual attributes of the recognition process (e.g., the visual properties of a crow such as its biological motion, size, ratio of its wings, color of its feathers, and so forth that authorize its initial recognition as a member of the *bird* category).

However, it has been recently suggested that the principles governing the formation of visual categories should be more tightly coupled with the perceptual aspects of recognition (Schyns, Goldstone & Thibaut, 1998). Instead of assuming that categorization operates on an already perceived input (see Pylyshyn, 1999, for discussions), it is proposed that categorization being executed can itself determine (at least partially) what the perception of the stimulus will be. More powerful, integrated theories of categorization, visual cognition and perception could emerge from these interactions. It is the purpose of this chapter to illustrate this point.

Diagnostic recognition: An overview.

It is worth remembering that in categorizing the visual input one seeks to obtain a close match between a category representation and a representation of the object in the input. This match between memory and input information is what we will call *a task* for the observer. We will here characterize a task as a strategy, as a means of directing an active search of information in the input for resolving a specific classification.

Generally speaking tasks or strategies are not rigid. Instead, different categorizations of an identical object tend to change the information requirements of the task at hand. For example, the information necessary to assign a visual event to the *Porsche, collie*, *sparrow*, *Mary*, or *New York* category might be comparatively more specific than that necessary for categorizing the same event as a *car*, *dog*, *bird*, *human face* or *city*. We will here examine the information required to place the input in a taxonomy, which is a sequence of progressively more inclusive categories such as *robin*, *bird*, *animal*. Within this hierarchy, we will concentrate on the so-called “perceptual” classifications (e.g., *robin* or *bird*), instead of the abstract functional classifications (e.g., *animal*) which tend to be less perceptual in nature. Henceforth, *task constraints* will denote the information that is required to place the visual input into one category of a hierarchy. Task constraints have traditionally been the main focus of categorization research, but they are an irreducible factor of *any* recognition task, and the first factor of the diagnostic recognition framework outlined here. Recognition can be viewed simply as the successful resolution of task constraints on a given input (Schyns, 1998).

The second factor governing diagnostic recognition is the *a priori* structure of perceptual information available to construct hierarchically organized categories. We group objects into perceptual categories because they "look alike" (see Ahn & Dennis, 2000; Hampton, 2000)–i.e., they share cues such as a similar silhouette or global shape, distinctive sets of parts similarly organized (e.g., *nose, mouth, eyes, ears, hair* and their *relationships*), typical surface properties (e.g., *smooth* vs. *discontinuous*, *symmetric* vs. *asymmetric*, and *textural*, *color* and *illumination cues*), or biological motion. Generally speaking, not all image cues are equally available; there are perceptual limitations on their extraction from the image. The structure and perceptual availability of object information has traditionally been the concern of perceptually-oriented object recognition researchers. However, perceptual cues are an irreducible factor of any object categorization, and the second factor of diagnostic recognition.

In the diagnostic recognition framework these two factors ---task constraints and persceptual cues--- interact: When the information required to assign an object to a category matches with input information, a subset of object cues become particularly useful (i.e., diagnostic) for the task at hand. Diagnosticity is the key component of recognition performance. However, perceptual limitations on the extraction of diagnostic cues should also affect performance. Thus, diagnostic recognition frames explanations of performance as interactions between cue diagnosticity and cue availability (Schyns, 1998). It is our view that the nature and the implications of these interactions has been largely neglected both in object recognition and in object categorization research. We will review some of these implications here, in part to demonstrate that these interactions are not just a theoretical construct, but instead lead to powerful and new phenomena which could force us to consider bridging what is traditionally independent: cognition and perception.

Specifically, this chapter is organized as follows. We first review recent evidence that constraints arising from a categorization task can both transiently modify the perception of common faces, objects and scenes, and more permanently affect the perception of expert object categories. While these results illustrate the phenomenology to be explained, it is difficult to derive specific predictions from them. To this end, we will present SLIP, a formalization of task constraints. In a nutshell, SLIP predicts (1) the successive steps of information acquisition prescribed in a categorization task, and (2) the perception that arises as a side-effect of (1).

Evidence that Categorization Changes the Perception of Objects

Studies in our laboratory have already provided evidence of flexible, task dependent perceptions. For example, in Schyns and Oliva's (1999) Experiment 1, three subject groups were instructed to each perform a different categorization (*male vs. female*, *expressive vs. non expressive*, *happy vs. angry vs. neutral*) on an identical set of *hybrid* faces presented tachystocopically on the computer monitor for 50 ms.. Hybrid faces associate a different perceptual content with a different spatial scale. For example, in the top face of Figure 1, the fine spatial scale (in fact, high spatial frequencies) represents a nonexpressive man. If you squint, blink, or step back from the picture you will perceive a smiling woman, represented at a coarse spatial scale (i.e. low spatial frequencies). We found that the categorization task modified the spatial scale that the observer used in a given hybrid face. Specifically, subjects in one task (expressive vs. nonexpressive judgments) would judge that the top picture was not expressive (using fine scale cues) whereas subjects in another task (which asks to identify the expression of the faces) saw the same face as happy (using coarse scale cues). These results reveal that scale perception is flexibly dependent on the information demands of different categorizations, an implication of diagnostic recognition (see also Oliva & Schyns, 1997, for other recent evidence). It also suggested that people can use different spaces composed of different visual cues (e.g., coarse blobs, vs. fine scale edges) spaces for the encoding of an identical stimulus.

Other evidence stems from the Change Detection literature (see Simons & Levin, 1997, for a review) which has often reported that people do not always perceive all features of a distal object. In a typical experiment, one feature of a real-world scene changes between successive presentations separated by a blank. For example, a particular object can change *location*, *color*, *texture*, or even disappear. Subjects tend to be blind to these straightforward changes, even though they that know there is a change. Rensink, O'Reagan and Clark (1997) showed that a change in a scene location of higher interest was detected faster than a change in a location of low interest, suggesting that visual information with little relevance for representing an object (Friedman, 1979) or a scene (Rensink et al., 1997), tends to be unnoticed (Simons & Levin, 1997; see also Dennett, 1991; Hochberg, 1982).

One of the authors recalls a particularly striking demonstration of change blindness (see the special issue of Visual Cognition for others) that appeared in the seventies on French television. An actor would stand on a Parisian sidewalk and ask a passerby for directions. During this conversation, two men holding a large mirror would walk between the actor and the passerby, in effect interrupting their conversation, and disrupting the situation. Behind the mirror was a second actor who would substitute for the first. The striking (and in this case funny) effect was that the passerby would resume giving directions when the mirror had gone without realizing that he was now talking to another, and often markedly different, person.

The evidence just reviewed suggests that the perception of an object is not exhaustive, but instead partial and contingent on a given task. Aside from these transient effects, expertise with objects is also known to give rise to sustained perceptual changes in restricted contexts. For example, experts categorizing X-rays (Christensen, Murry, Holland, Reynolds, Landay & Moore, 1981), dermatosis (Norman, Brooks, Coblentz & Babcock, 1992) or sexing chickens (Biederman & Shiffrar, 1987) can often perceive important differences that novices just do not see. Schyns and Murphy (1994; Schyns & Rodet, 1997) suggested that this could highlight a pervasive principle of category learning which creates new features for new object categorizations. In a typical experiment (see Schyns and Rodet, 1997), used three categories of unknown stimuli called "Martian cells" called *X*, *Y* and *XY*. Categories were defined by specific blobs common to all category members to which irrelevant blobs were added (to simulate various cell bodies). The main goal of Schyns and Rodet's experiment was to demonstrate that different categorization constraints could induce orthogonal perceptions of the defining *xy* component–i.e., perceptions of *xy* as an *x&y* feature conjunction, or as an *xy* unitary feature. One group of subjects was asked to learn *X* before *Y* before *XY* (*X->Y->XY*); the other group learned the same categories in a different order (*XY->X->Y*). Reliable classifications of *X*, *Y* and *XY* stimuli in the testing phase indicated, without any doubt, that all subjects saw and attended to the components *x* and *y*. *X-Y* cells were used to understand the perceptual analysis of *XY*. *X-Y* cells were *XY* exemplars in which the *x* and *y* components were not adjacent to each other. The reasoning was that subjects should categorize *X-Y* cells as *XY* members if they perceived and represented *XY* as conjunction of two individuated features. Results revealed that only one group (*X->Y->XY*) performed this categorization while the perception of *XY* in the other group prompted *X* or *Y* classifications of *X-Y*. In sum, orthogonal classifications of *X-Y*, when its component features were both clearly perceived and used in the experimental groups, suggested that different features were acquired to perceptually analyze and represent *XY* (see also Gauthier & Tarr, 1997, for evidence of perceptual expertise inducing enhanced sensitivity to the configurations of "face-like" abstract objects).

Effects of category expertise can also be found in simpler tasks, using much simpler stimuli, confirming the pervasiveness of the principles. For example, Goldstone (1994) trained humans to categorize squares varying in size or brightness. After prolonged training, subjects were tested in a same/different task. When a dimension had been relevant for categorization, same/different judgments along this dimension were more accurate than those of subjects for whom the dimension had been irrelevant, or of control subjects who had not undergone categorization training. The greatest acuity increase along the categorization-relevant dimension was found between those points that had served as the boundaries between the learned categories. This sensitization of the relevant dimension also extended to other values along that dimension. Dimensions that were irrelevant for categorization became less accurate than those of control subjects.

Using a related approach, Archambault and Schyns (1999) found that the perception of an entire object dimension (its texture, not just the trained textural values) could be enhanced when it became diagnostic via category training. In contrast to Goldstone (1994), the same/different task involved simultaneous, not sequential presentations of objects. That is, subjects saw two objects (e.g. computer-synthesized 3D fruits) appearing simultaneously on the screen and their goal was to judge whether the stimuli were physically identical. Even with this kind of stimulus display, and even though all subjects knew that the two stimuli could only differ in texture (i.e., their shape and color were identical), the subjects for whom this dimension was diagnostic during training perceived the difference with a significantly greater accuracy. Furthermore, they were better able to use this information in a subsequent categorization task.

In sum, the data reviewed so far suggest that the perception of faces, objects and scenes can either change transiently depending on the specific demands of a categorization task, or more permanently when expert categorizations require the creation of new distinct cues, or a new category distinction along a continuum. If an object can be differently perceived when the information requirements of its categorization change (either because people want to categorize this object at different levels of specificity, or because they become expert with it), it becomes crucial to predict the information requirements of a categorization to predict the visual information that must be perceived. SLIP was designed to this end.

SLIP: A Formal Model of Task Constraints

SLIP (for *Strategy Length & Internal Practicability*) is an ideal categorizer in the sense that it applies an optimal feature testing strategy to determine the category membership of an object in a taxonomy. To illustrate, SLIP would test that the stimulus presented in Figure 2 possesses a *cube* and a *wedge* to place it in the specific *PIM* category in the taxonomy of Figure 3a. It would verify that the same input possesses only a *cube* to assign it to the more general *LAR* category. SLIP uses one strategy (e.g., for the *PIM* category of Figure 3a it could be *Strat*(*X, PIM*) = [{“does *X* possess *cube?”*} & {“does *X* possess *wedge*"}]) to drive its perceptual systems to make an *ordered series* of feature tests (see Woodman & Luck, 1999, for a recent defense of serialism).

Insert Figure 2 about here

The behavior of SLIP is fully determined by two independent factors, the *length* and the *internal practicability* of strategies, themselves dependent on the organization of featural information in taxonomies. The length of a strategy is the minimal number of feature tests required for its completion. For example, two feature tests are required to identify a *PIM* in the taxonomy of Figure 3a. Thus, a strategy comprises sets of features and SLIP tests their presence, one set at a time, in a specific order. We suppose that response time is a function (typically linear) of the total number of features tested when SLIP executes a strategy. Since the higher level in Figure 3a comprises only one feature test whereas the lower level comprises two feature tests, SLIP predicts a slower access the lower level. A glance at the taxonomy of Figure 3a will indicate that the overlap of features between the categories of a taxonomy determine the length of categorization strategies (e.g. *cube* overlaps between several subordinate categories).

The second factor is the ease with which SLIP executes a particular test in a strategy. Figure 3b illustrates this idea. To categorize an object as a *LAR*, at the high level, the strategy is *Strat*(*X, LAR*) = [{“does *X* possess *macaroni?”*, “does *X* possess *cube*", "does *X* possess *pinched cube*"}]. That is, any one of these feature tests convey the same information about the category membership, and together they form an exhaustive set of *redundant* feature tests to access the category. A high level of redundancy in a strategy will elicit shorter categorization times. This occurs because the processes of SLIP are noisy and sometimes slip off the ideal track to test random object features. In general, slippage will increase the number of feature tests attempts and the time taken to reach a category decision. However, slippage to a diagnostic feature is more likely for categories with many redundant features than for those with fewer features. Thus, SLIP predicts a faster access to the more redundant higher level of Figure 3b.

Insert Figure 3 about here

It is useful to contrast SLIP to another, and very different processing model of categorization: Lamberts' (1994) EGCM (Extended Generalized Context Model). SLIP is uses a task-dependant rule to represent a category, whereas EGCM represents a category with all its stored exemplars. SLIP accumulates features in tine, whereas EGCM compares them in parallel. SLIP categorizes and object whenever the collected features match the representation of a category in memory. In contrast, the probability that EGCM places an object in a category is a function of the distances (i.e., similarities), in a multi-dimensional psychological space, between this object and all the exemplars of memorized categories. Finally, SLIP has a single free parameter (i.e., the probability of a slip), whereas EGCM has several (i.e., one *inclusion rate* and one *utility value* for every dimensions).

In sum, SLIP predicts that an object should be categorized faster in category X than in category Y (1) if the length of the optimal strategy that identifies the object as X is smaller than the length of the optimal strategy that identifies the same object as Y and (2) if the optimal strategy associated with category X comprises more redundant attributes than that of category Y (see Gosselin & Schyns, 1999a and the Appendix for a formal presentation).

SLIP incorporates the two critical aspects of diagnostic recognition discussed earlier: task constraints and information availability. Task constraints correspond here to the different strategies associated with a vertical organization of categories–i.e., the idea that different categorization strategies can be applied to the same object. From these, we can ask different sorts of questions. For example, we can examine whether the two determinants of SLIP just discussed (i.e., strategy length and internal practicability) determine the speed of access to the different levels of a taxonomy. However, speed of access *per se* is not sufficiently powerful to reveal the locus of processing at which strategies influence categorization speed. Strategies could tap into the perception of the input, the memory retrieval of category information, the access to category names, the decision between different category names, or a combination of these factors. Nevertheless, the object categorization literature makes a number of specific predictions for the speed of access to the categories of a hierarchy. As a safeguard to the empirical validity of SLIP, we must verify that its principles make similar predictions. The section below will investigate these in detail. SLIP also allows a more precise study of the interactions between the specific requests of information associated with different categorizations and the perception of the input itself. We will examine these later in the chapter.

Effects of taxonomy on categorization: basic-level performance

It has long been established that speed of access varies with the level of categorization considered. In Rosch, Mervis, Gray, Johnson and Boyes-Braem's (1976, Experiment 7) seminal paper, participants were taught the name of 18 objects at three levels of categorization–the *subordinate* (e.g., *Levis, Macintosh*), *basic* (e.g., *pants, apple*) and *superordinate* (e.g., *clothes, fruit*)[1]. These objects belonged to one of six possible non-biological taxonomies: *musical instruments*, *fruit*, *tools*, *clothing*, *vehicles*, and *furniture*. In a verification task, subjects were shown a category name followed by a stimulus picture, and had to determine whether they matched. Categories at the basic-level were fastest to verify, and categories at the subordinate level slowest (see also Hoffmann & Ziessler, 1983; Jolicoeur, Gluck & Kosslyn, 1984; Murphy, 1991; Murphy & Smith, 1982; Murphy & Brownell, 1985; Tanaka & Taylor, 1991).

The basic level is superior in many other respects: (1) objects are named quicker at this level than at any other level of abstraction (Hoffmann & Ziessler, 1983; Jolicoeur, Gluck & Kosslyn, 1984; Murphy, 1991; Murphy & Smith, 1982; Murphy & Brownell, 1985; Rosch et al., 1976; Tanaka & Taylor, 1991); (2) objects are designated preferentially with their basic-level names (Berlin, 1992; Brown, 1958; Rosch et al., 1976; Tanaka & Taylor, 1991; Wisniewski & Murphy, 1989); (3) many

---

[1]It is worth pointing out that the usage of "basic level" is ambiguous. It can refer to the middle-level of a three-level hierarchy (with the level above called "superordinate" and the one bellow "subordinate"–e.g., Markman, 1989; Tanaka & Taylor, 1991), to an index of performance (the fastest level, or the one most often used to name things, and so forth–e.g., Corter and Gluck, 1992; Anderson, 1990, 1991), or to both the level of categorization and the index of performance (e.g., Rosch et al., 1976; Mervis & Crisafi, 1982). Henceforth, the *basic-levelness* of a category will denote a measure of performance. Whenever possible, we will refer to the levels of abstraction as the subordinate, basic, and superordinate. Otherwise, we will use a set of unambiguous level descriptors–e.g., low, middle and high. The subordinate-basic-superordinate trio has the advantage of having a phase known to most psychologists.

more features–especially shapes–are listed at the basic level than at the superordinate level, with only a slight increase at the subordinate level (Rosch et al., 1976; Tversky and Hemenway, 1984); (4) throughout development, basic level names are learned before those of other categorization levels (Anglin, 1977; Brown, 1958; Rosch et al., 1976; Horton & Markman, 1980; Markman, 1989; Markman and Hutchinson, 1984; Mervis and Crisafi, 1982); and (5) basic names tend to be shorter (Brown, 1956; Rosch et al., 1976). Convergence of these performance measures is crucial to establish a preferred categorization level, even though verification speed is the most commonly used.

To recapitulate, the behavior of a SLIP observer is fully determined by the *length* and *internal practicability* of the strategies it uses. We will first review experiments that involve only the internal practicability of SLIP. These include experiments with artificial (Murphy & Smith, 1982; Murphy, 1991; Gosselin & Schyns, 1999a) and natural taxonomies (Rosch et al., 1976; Tanaka & Taylor, 1992). Second, we will report experiments that varied the length of strategies, the second determinant of SLIP (Hoffmann & Ziessler, 1983; Gosselin & Schyns, 1998, 1999a). All used artificial taxonomies. Finally, we will discuss the only experiment that integrates the two determinants of SLIP (Gosselin & Schyns, 1999a).

Practicability determines speed of access

Faster access at an intermediate level. One of the most influential experiment on the basic-level is that of Murphy & Smith (1982, Experiment 1). It is influential because most subsequent basic level experiments have used the same procedure. Their participants were initially taught the artificial taxonomy represented at the top of Figure 4. In a later testing phase, they were shown a category name followed by a stimulus. Subjects' task was to verify as quickly as possible whether the name and stimulus matched. As shown in Figure 4, mid-level categories have the highest practicability. Table 1 illustrates that they were verified faster, and the high-level categories slowest. Using the same taxonomic organization of category attributes, Murphy (1991, Experiment 4, Simple) replicated these results. In fact, the highest practicability of the middle level is also responsible for its faster access in the artificial taxonomies of Mervis & Crasifi (1982) and Murphy (1991, Experiment 4, Enhanced) (see Figure 4 and Table 1).

Insert Figure 4 about here

Moreover five natural taxonomies had a greater redundancy at the intermediate level: Rosch et al. (1976, Experiment 7[2]), Tanaka & Taylor (1991, Novice[3]), and Johnson & Mervis (1997, advanced songbird expert, intermediate songbird expert, and novice[4]). For these, we assumed that the features subjects listed reflected their representations (see Rosch & Mervis, 1975). In addition, following Tversky and Hemenway (1984) and Tanaka and Taylor (1991) we assumed that one feature was never listed for two contrasting categories. Table 1 reveals that, as predicted in SLIP, basic-levelness was a direct function of average number of redundant attributes at each level.

Insert Table 1 about here

Faster access at the lower level. Murphy and Smith (1982, Experiment 3) added a unique set of attributes to Murphy and Smith's (1982) artificial tools at the lower-level. Figure 5 illustrates the abstract organization of the features. It also shows that the lower-level categories were more practicable because they had more redundant attributes. As predicted in SLIP, Table 1 reveals that these categories were accessed faster than categories at the other levels. Tanaka & Taylor's (1992, Expert[5]) is a variation on this theme: they used expertise to "add" redundant features at the lower level and thus speed up its access. They found that the basic and subordinate categories were equally fast and the superordinate

---

[2]In Rosch et al. (Rosch et al., 1976, Table 2, non biological taxonomies, raw tallies), the mean number of added redundant features was of 1.85, 5.55 and 3.5 for subordinate, basic, and superordinate, respectively.

[3]In Tanaka and Taylor's (1991, Novice) subjects listed approximately 8, 12, and 7 new redundant features for the superordinate, basic, and subordinate levels of categorization, respectively. The basic-level categories were the fastest, and the subordinate-level categories were the slowest (in Table 1 we give the mean RTs of bird novices and of dog novices).

[4]Johnson and Mervis's (1997, Experiment 1, Songbirds condition) used four-level natural taxonomies in a verification task. Their advanced songbird experts listed 1.75, 5, 6.02, and 3.75 for the superordinate, basic, subordinate and sub-subordinate levels, respectively. For the intermediate songbird experts, these numbers were 1, 4.87, 4.28, and 2.47, for the same levels. For the novices and the tropical freshwater fish experts, the numbers were 1.08, 2.47, 0.23, and 0.02.

[5]Their subjects listed approximately 8, 10, and 10 new features for the superordinate, basic, and subordinate levels of categorization, respectively (compare this with 8, 12, and 7 for the superordinate, basic, and subordinate levels in their Novice condition in the previous section).

categories the slowest (Table 1 gives the mean RTs of bird and dog experts). Gosselin & Schyns' (1999a, Experiment 2, LOW_FAST) inserted computer-synthesized four-geon chains similar to the one in Figure 2 in the two-level taxonomy illustrated in Figure 5. All strategies had length 1 but the high and low levels differed in practicability. Low-level strategies had greater practicability than high-level ones, and were verified faster (see Table 1).

Insert Figure 5 about here

Faster access at the higher level. In his Experiment 5, Murphy (1991) added a set of unique values to the high-level categorizations of Murphy and Smith's (1982) artificial tools. Figure 6 shows that this level becomes more practicable and Table 1 reveals that it was indeed accessed faster, as predicted in SLIP.

Insert Figure 6 about here

Gosselin & Schyns (1999a, Experiment 2, HIGH_FAST) inserted computer-synthesized four-geon (Biederman, 1987) chains (see Figure 2) in the two-level taxonomy illustrated in Figure 6 (see also Figure 3b). All strategies had length 1 but the high and low levels differed in practicability; high-level strategies had greater practicability than high-level ones, and Table 1 shows that they were verified faster.

Strategy Length determines speed of access.

In all the experiments reviewed so far, the length of the categorization strategies was held constant. Variations of strategy lengths were first tested in Hoffmann and Ziessler (1983, Hierarchy I). They used "PacMan ghosts" artificial objects organized in the top taxonomy of Figure 7. Strategy length was 1 at the high- and middle- levels, but 2 at the low-level. Participants accessed the high- and mid-levels categories equally fast, and were slower for low-level categories (see Table 1).

In Gosselin and Schyns (1998) participants learned the taxonomy of Figure 7 applied to artificially *textured* and *colored geons*. This taxonomy ascribes strategies of different lengths to the different levels of abstraction: length 1 for the high-level, length 2 for mid-level, and length 3 for low-level categories. Categorization was fastest at the higher level, and slowest at the lower level (see Table 1).

Gosselin and Schyns' (1999a, Experiment 1, HIGH_FAST and LOW_FAST) isolated strategy length using a set of computer-synthesized four-geon chains (see Figure 2) inserted in the two two-level taxonomies illustrated in Figure 7 (see also Figure 3a for HIGH_FAST condition). These taxonomies were designed to induce orthogonal patterns of categorization speed across conditions. Participants systematically verified length 1 strategies faster than length 2 strategies, irrespective of the considered level (low vs. high).

Insert Figure 7 about here

Interactions between Strategy Length and Internal Practicability determine speed of access.

We have so far shown that the two determinants of SLIP (strategy length and internal practicability) can independently determine a faster access to any level of a taxonomy. Gosselin and Schyns (1999a, Experiment 3) explored how these two factors could interact to determine performance. Of many possible interactions, they investigated three (EQUAL, SL_DOWN and IP_UP). In EQUAL, strategies at the high and low-levels had an equal length of 1 and the same constant practicability (see Figure 8). SLIP predicts that categorization speeds should be equal across levels, and this is what they found. In SL_DOWN, faster categorizations at the lower level were produced by augmenting the length of the strategies that access the high-level categories (see Figure 8). In the IP_UP scenario, the difference of strategy length just discussed was preserved, but the high-level became fastest because the practicability of the low level was decreased (see Figure 8). In sum, starting from an EQUAL access to two levels of a taxonomy, Gosselin and Schyns manipulated strategy length and internal practicability to modify the fastest level of a taxonomy. A change of strategy length in SL_DOWN produced faster categorizations at the low level. From this, a decrease in the internal practicability of the low level in IP_UP produced faster categorization at the high level (see Table 1).

Insert Figure 8 about here

Summary

SLIP correctly predicts most results of these experiments on the basic level (see Table 1). A Monte-Carlo study showed that the probability of this, or a better fit of the data being due to chance alone is smaller than .001–in any case, better than the fit obtained with other established basic-level measures (i.e., Jones', 1983, category feature-possession; Corter & Gluck's, 1991, category utility; Pothos & Chater's, 1998, compression measure; and Medin & Schaffer's, 1978, context model, modified by Estes, 1994). Note that none of these measures predicts verification latencies per se; they predict a superiority of a measure (i.e., greatest *utility* in the cases of category feature-possession, category utility, and compression measure; and maximum *within-category similarity and between-category dissimilarity* for the context model) that is correlated with verification latencies. To the extent that any model of categorization implements computational constraints (even if these are not well specified), the conclusion is that those of SLIP are closest to those underlying the speed of access to the categories of a taxonomy. At this stage, it is worth remembering that these two constraints are strictly of a categorical nature: in SLIP, speed of access is a direct function of the representation of object information in memory. In other words, the modeling of task constraints in SLIP accounts for the first aspect of diagnostic recognition. We now turn to the question of whether these strategies can modify the perception of visual object information.

Effects of Taxonomies on Perception

It is notoriously difficult to assess that the perception of an object has changed as a result of a change in categorization (see Pylyshyn, 1999; Schyns et al., 1998 for debates). SLIP, however, makes two testable perceptual predictions. A first prediction is that the features prescribed in a categorization strategy are only those sampled in the input. This implies that changing the features of a strategy (e.g., via the acquisition of conceptual expertise) could, to some extent, control the features that are (or not) seen in a given object. Below, we review studies (Archambault, O'Donnell & Schyns, 1999) that used *Change Blindness* to demonstrate that basic and subordinate categorization strategies can induce different perceptions of an identical object. A second perceptual prediction of SLIP is that features in a strategy are tested serially, in a specific order. Another study reviewed below (Gosselin an Schyns, 1999b) used a masking technique to determine the order of feature testing.

Representation-Driven Blindness

In Figure 9a, the mug could either be categorized as simply *a mug* at the general level of categorization, or as *Peter'*s *mug* at a more specific level. If a cylindrical shape and the presence of a handle would be diagnostic of *mug* more specific (and relatively independent) information (e.g., the specific texture on the mug) might be required to classify the input as *Peter's mug*. People could preferentially attend to and represent the visual properties required in their categorization strategies and differently perceive the same mug.

Archambault, O'Donnell & Schyns (1999) investigated this hypothesis. In a first experiment, one group (*MUG-computer*) learned the mugs at a general level and the computers at a specific level whereas the other group (*mug-COMPUTER*) learned the opposite assignment of category level to objects: mugs as specific and computers as general. This ensured that subjects learned an identical set of objects at different levels of categorization. It was expected that these learning differences would lead to different categorization strategies in the different groups.

A change detection task (see Simons & Levin, 1997) tested the visual encodings of the objects. The mugs and computers were inserted in a complex office scene (see the frames of Figure 9b). In a trial, two office photographs were successively presented, separated by a blank. Between the two frames, a mug could change (be replaced by a different mug) or disappear, a computer could change or disappear, or other office objects could disappear. All subjects (i.e., MUG-computer and COMPUTER-mug) were exposed to the same object changes and disappearances. Their task was to identify the difference between the two frames.

Archambault et al. (1999) found that subjects were "blind" (i.e., took longer to perceive) changes involving objects they learned at a general level, whereas other subjects were aware of the same changes when they learned these objects at a specific level. These different perceptions did not simply arise because subjects looked preferentially at the location of objects learned at a specific level because all subjects equally fast in perceiving the disappearances of all objects—i.e., those learned at a specific level *and* those learned at general level. In sum, these orthogonal perceptions of identical object changes when disappearances were detected equally fast isolated the effect of different categorization strategies on perceived object features.

A second experiment replicated these results with a within-subjects design that ruled out the objection that subjects preferentially scanned the image location where subordinate changes appeared. In their Experiment 1, *mug* and *computer* were orthogonally assigned to general and specific levels across groups. In their Experiment 2, subjects learned to categorize mugs and computers at *both* the general and specific levels. That is, subjects learned a subset of mugs and a subset of computers at a specific level, and the remaining mugs and computers were learned at a general level. Consequently, each image location now embodied either a general- or a specific-level object change. Attention to the location of a specific-level change now implied necessarily attention to the general-level object change also occurring in the same image location. Different perceptions of general- and specific-level object changes in these conditions confirmed that the nature of categorization did not affect selective attention to locations, but selective encodings of the objects present in these locations.

Insert Figure 9 about here

In SLIP, the features involved in a categorization strategy are those preferentially sampled from the input. In the previous experiments, and even though each frame of a sequence appeared for 5 s on the screen, giving ample time for careful visual encoding, subjects who knew the subordinate categorization of mugs and computers encoded properties that the others did not. This elicited different perceptions of identical objects and so the first perceptual prediction of SLIP is confirmed. We come back to its implications in the General Discussion.

Serial, ordered testing of features and perception

A second perceptual prediction of SLIP is that strategies specify an order of feature testing. If this order is respected, then the perceptual appearance of the stimulus could change. To illustrate, a *PIM* in the taxonomy of Figure 3a is represented either by *Strat*(*X*, *PIM*) = [{*wedge*} & {*cube*}], or by *Strat*(*X*, *PIM*) = [{*cube*} AND {*wedge*}]. These two strategies have equal speed of access, but the order in which the two features are tested differs. Why would one adopt the first or the second strategy? In the taxonomy of Figure 3a, one strategy (i.e., *Strat*(*X*, *PIM*) = [{*cube*} & {*wedge*}]) is more *robust* in categorization under time pressure. It is more robust because it is more likely to lead to a valid, if approximative, categorization of the input. We know that the input is at least a *LAR* if it has a *cube*. *Category robustness* is critical in everyday categorization because unseen features can be inferred from a partial categorization (e.g., Anderson, 1991).

Gosselin and Schyns (1999b) tested the prediction that people who categorized an identical scene could integrate their luminance and color in a reverse order. To this end, they synthesized four stimuli by combining two different luminance patterns (that we call *flat* and *hilly*) with two different chromatic patterns (called *grassy* and *sandy*). A learning procedure was devised to induce a different, two-level, taxonomic knowledge of these four stimuli in two subject groups (called LUMI and CHRO; see Figure 10).

Insert Figure 10 about here

At the general level, LUMI subjects learned to separate the four scenes into "flat" and "hilly", on the basis of luminance cues, whereas CHRO subjects learned to separate the same scenes into "grassy" and "sandy" on the basis of chromatic cues. At the specific level, LUMI and CHRO subjects all learned to categorize the stimuli as either "field" (the combination of *flat* and *grassy*), "desert" (*flat* and *sandy*), "mountain" (*hilly* and *grassy*) or "dune" (*hilly* and *sandy*). Note that the specific categorizations are strictly identical in the two groups. The conjunctive nature of the stimuli warrants that the input scene can only be recognized as, e.g. "field," when its flat luminance and its grassy chrominance are perceived and integrated.

This property can be used to ascertain whether subjects are more sensitive to the dimension defining the general than the specific level of their taxonomy, and therefore perceive the scenes according to their organization of knowledge. Suppose that the field picture is briefly presented on the screen, immediately followed by a mask. Subjects can make three errors at the lower level, depending on which information they misperceive; "dune" implies a misperception of both the flat luminance and the green chrominance of the field; "mountain" implies a misperception of only the flat luminance, whereas "desert" implies a misperception of only the green chrominance.

Subjects were only tested on the specific-level categorizations of the four scenes. Gosselin and Schyns predicted that the organization of luminance and chromatic information in the LUMI and CHRO taxonomies would determine different orders of feature testing, in turn leading to different perceptions of identical stimuli. They found that subjects placed in an identical condition of stimulation (e.g. seeing a field) and response (choosing between "field," "mountain," "desert" or "dune") produced opposite patterns of categorization errors (i.e. respond more often "desert" than "mountain" in LUMI, but "mountain" than "desert" in CHRO), revealing a differential sensitivity to luminance and chrominance in the groups (see Figure 10). A similar analysis and results applied to all four stimuli of the experiment.

These results have potentially far reaching implications. It is well established that luminance and color are two of the main dimensions of visual processing. If different categorization strategies applied to strictly identical scenes, in strictly identical conditions of response and stimulus presentation, can produce a different order of integration of luminance and chromatic cues, then this would constitute strong evidence that categorization strategies can determine perception. We come back to this point and its implications in the *General discussion*.

General discussion

Diagnostic recognition is a framework that seeks to establish a dialogue between categorization tasks and object perception. A categorization task specifies the information that must be obtained from the visual array to place the input object in a category. The framework proposes that these information demands can modify the perception of the object itself. We presented a formal model of task constraints called SLIP (standing for Strategy Length & Internal Practicability). In SLIP, task constraints are implemented as different categorization strategies. They specify a series of tests on the presence of visual properties in the input object. The visual properties tested in a strategy therefore depend on the categorization considered. Turning to the organization of categories in memory, we noted that they formed taxonomies of progressively more inclusive membership–e.g., a *Porsche* is a *car* is a *vehicle*. In a taxonomy, categories can share visual properties (e.g., the part *wheel* will be shared between *car*, *motorcycle*, *bus*, *plane*, and so forth). When a feature is shared between two categories, it does not by itself isolate one of the categories. Instead, a *combination* of features is sometimes necessary to identify the category. In SLIP, the number of features entering the combination determines the *length* of a strategy. Longer strategies are slower to verify and so SLIP predicts that categories with overlapping features are slower to verify. On the other hand, redundant features provide the same categorization information (to the enthusiast, a real-life Porsche 911 can be identified from many different *shape* features, including its *silhouette*, *characteristic hood*, *windscreen*, and so forth). That is, testing one, two or all redundant features of a category does not add any information to the membership decision. However, the feature pick-up process of SLIP is noisy–it slips off its ideal track–and the likelihood of randomly slipping to a diagnostic feature is higher when the category comprises many redundant features. Thus, more redundant categories are recognized faster. Numerical simulations of SLIP and experimental data revealed that it provides the closest fit to classical basic-level experiments data.

We argued earlier that a close fit to reaction time is not sufficiently powerful to inform the locus at which strategies affect behavior–i.e., perception, memory, or decision. A close behavioral fit nevertheless licenses further explorations of the perceptual assumptions of SLIP. We reviewed two of these. SLIP assumes that the observer will preferentially perceive the object features tested in a categorization strategy. Experiments on *Change Blindness* (see Archambault et al., 1999) demonstrated that subordinate categorizations induced the perception of the properties of an object that is basic categorization did not. SLIP also predicts that object properties are tested in a specific order. A scene categorization experiment revealed that subjects who had to integrate the same luminance and color cues could selectively respond more accurately to one of the information sources, if their taxonomy was based on this source (Gosselin & Schyns, 1999b). Together, the evidence presented suggest that categorization should be more closely bound to perception, if only because it influences it. We now turn to limitations of the framework and future research.

Limitations of Diagnostic Recogition.

Diagnostic recognition is a framework in which the information goals of object categorization tasks are considered before their perceptual representations. Although this is a good, generally recommended approach to theory construction (e.g., Marr, 1982), it nevertheless presents serious limitations for the study of object representations.

The reason is simply that thinking from task constraints to their perceptual representations could over-represent the considered information demands in the proposed representation. For example, if it were discovered that the information requirements of an object categorization were X, then it would be an easy step to assume that the representation of this object was effectively that X. But then, how would we know whether X represents the object, or the task itself?

Even though we acknowledge that diagnostic recognition is not designed to study issues of face, object, and scene representation, it is not clear to us that such a framework does exist yet, irrespective of what researchers want to believe. We do not yet have the means to study object representations, and when we think we do, we might only be studying representations of tasks. What diagnostic recognition does well is predict which visual information (irrespective of its actual representational format) is required to access a category and which information is not. On this basis, new issues can be explored that attempt to determine the visual information that is necessarily perceived, irrespective of task demands, and which information is strictly contingent on the categorization task. This could constitute a first step towards resolving the problem of object vs. task representations.

SLIP and reduction of uncertainty

The take-home message of this chapter is that categorization tasks can determine (at least in part) the perception of objects via the selective use of diagnostic information. It is therefore crucial to further examine the information requirements of different categorization tasks. Looking at the formal organization of visual information in taxonomies, we argued that categories are accessed via *overlapping* and/or *redundant* attributes. We showed that the particular sequence in which feature tests are performed is of critical importance for perception. We suggested that *category robustness* (i.e., the idea that feature tests are ordered so as to categorize the input as quickly as possible) could be a determinant of this order. There are other possibilities just as consistent with Gosselin & Schyns' (1999b) results, and it would be worth exploring them further. For example, features in a categorization strategy could be tested in an order that minimizes the uncertainty of the considered categorization. If the categorizer needs to judge whether or not the input is , e.g., a dog, he or she would first pick the features that lead to the greatest reduction of the uncertainty of the categorization. The categorizer would also optimize a measure of robustness, but on the information, not the category.

Unravelling features in the input

This proposal assumes that people use visual input features to categorize objects, but it does not provide the means to isolate what these input features are. It is notoriously difficult to ascertain what the features of an object are (Schyns et al., 1998). However, in the absence of an efficient empirical technique to unravel the object features actually *used* in different categorizations, it is impossible to know the distribution of information in the categories of a taxonomy (in terms of redundancy and overlap), and the approach proposed here disintegrates. It is therefore on our agenda to develop techniques to identify the visual features involved in different categorizations. For example, in Schyns and Oliva (1999), the use of hybrid stimuli similar to those of Figure 1 licensed the conclusion that different categorizations used different bandwidths of spatial information. Even though this is not sufficiently specific to know what features were used, the hybrid techniques narrows down the search for features to a subspace (a restricted bandwidth of spatial frequencies) of the original input space (the full spectrum). This is a useful starting point to start working on the features of natural images.

We are developing further an existing technique called reverse correlation (Ahumada & Lovell, 1971). This technique enables the extraction of the features diagnostic of a particular task. In a nutshell, the technique introduces noise in the visual input (enough to make several mistakes), and then adds together the images that lead to hits and false alarms while subtracting those that lead to correct rejections and misses. The result is a *classification image*, a template of the information *used* in the input to complete the task. Again, this technique will not reveal an actual breakdown of the original stimulus into its psychologically-relevant visual features (i.e., those of a strategy), but it will present a significant reduction of information in which a more precise search can be conducted.

Implications for similarity

From a logical standpoint, we know that there are an infinite number of ways in which one thing can be similar to another (e.g., stones A and B could be similar because they both weight less than 1000 pounds, than 1001 pounds, than 1002 pounds, …, than a 1000+n pounds), and thus it is impossible to ground categorization on similarity (Quine, 1977). Similarity is too unconstrained to be the basis of categorization. Additional constraints are thus needed to specify the respects in which objects are similar if similarity must be the basic mechanism of categorization. Diagnostic recognition provides such constraints: the respects of similarity are the dimensions that are diagnostic for the task at hand. We therefore believe that it is critical to understand the processes that set the appropriate dimensions of the space of stimulus encoding in different categorization tasks. A similar view, but applied only to the similarity of two objects is defended in Medin, Goldstone and Gentner (1993).

Concluding Remarks

Diagnostic recognition is a new approach within which to frame recognition and categorization problems. We have here presented SLIP, an implementation of the framework that places a central focus on the notion that information picking categorization strategies can partially determine perception. Even though the framework is still in its infancy, and much groundwork remains to be accomplished to arrive at any satisfactory theory of categorization, we believe that the framework offers a powerful and integrated approach to the understanding of face, object and scene categorization and perception.

Acknowledgements

Frédéric Gosselin was supported by scholarships from the Fonds pour la formation de Chercheurs et l'Aide à la Recherche (FCAR) and from the University of Glasgow during this research. This research reported here was funded by ESRC grants R000 222 700 and R000 237 412 to Philippe G. Schyns

Reference

Adelson, B. (1985). Comparing natural and abstract categories: a case study from computer science. *Cognitive Science, 9*, 417-430.

Ahn. W& Dennis, M. J. (2000). Dissociation between categorization and similarity judgment: differential effect of causal status on feature weights. In Hahn, U and Ramscar, M. (Eds.) *Similarity and categorization*. Oxford: Oxford University Press.

Ahumada Jr., A.J. & Lovell, J. (1971). Stimulus features in signal detection. *Journal of Acoustical Society of America*, *49*, 1751-1756.

Anderson, J. R. (1991). The adaptative nature of human categorization. *Psychological Review*, *98* (3), 409-429.

Anglin, J. M. (1977). *Word, object and conceptual development*. New York: Norton.

Archambault, A. & Schyns, P. G. (1999). Ceci n'est pas une orange Manuscript submitted for publication.

Archambault, A., O'Donnell, C. & Schyns, P. G. (1999). Blind to object changes: when learning the same object at different levels of categorization modifies its perception. *Psychological Science*, *10*, 249-255.

Berlin, B. (1992). *Ethnobiological Classification: Principles of Categorization of Plants and Animals in Traditional Societies*. New Jersey: Princeton University Press.

Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, *94*, 115-147.

Biederman, I. & Shiffrar, M. M. (1987). Sexing day-old chicks: a case study and expert systems analysis of a difficult perceptual-learning task. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, *13*, 640-645.

Brown, R. (1958). How shall a thing be called? *Psychological Review, 65*, 14-21.

Bruner, J. S., Goodnow, J. J. & Austin, G. A. (1956). *A Study of Thinking*. New York: Wiley.

Cantor, N. & Mischel, W. (1979). Prototypes in person perception. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (pp. 3-52). New York: Academic Press.

Corter, J. E. & Gluck, M. A. (1992). Explaining basic categories: Features predictability and information. *Psychological Bulletin, 111*, 291-303.

Christensen, E. E., Murry, R. C., Holland, K., Reynolds, J., Landay, M. J., & Moore, J. G. (1981). The effect of search time on perception. *Radiology*, *138*, 361-365.

Dennett, D. C. (1991). *Consciousness Explained.* Boston: Little, Brown and Company.

Friedman, A. (1979). Framing pictures: The role of knowledge in automatized encoding and memory for gist. *Journal of Experimental Psychology: General*, *3*, 316-355.

Gauthier, I., & Tarr, M. J. (1997). Becoming a 'Greeble' expert: exploring the face recognition mechanism. *Vision Research*, *37*, 1673-1682.

Goldstone, R. L. (1994). Influences of categorization on perceptual discrimination. *Journal of Experimental Psychology: General*, *123*, 178-200.

Gosselin, F. & Schyns, P. G. (1997). Debunking the basic level. *Proceedings of the nineteenth annual conference of the cognitive science society* (pp.277-282). New Jersey: Lawrence Erlbaum.

Gosselin, F. & Schyns, P. G. (1998). The contingency of parts in object concept. *Proceedings of the twentieth annual conference of the cognitive science society* (p. 1222). New Jersey: Lawrence Erlbaum.

Gosselin, F. & Schyns, P. G. (1999a). *Why do we SLIP to the basic-level? Computational constraints and their implementation.* Manuscript submitted for publication.

Gosselin, F. & Schyns, P. G. (1999b). *New look at the interactions between between knowledge and perception.* Manuscript submitted for publication.

Hampton, J. A. (1999). The role of similarity in natural categorization. In M. Ramscar, U. Hahn, E. Cambouropolos, & H. Pain (Eds.) *Similarity and categorization.* Cambridge: Cambridge University Press.

Hochberg, J. (1982). How big is a stimulus. In J. Beck (Ed.). *Organization and Representation in Perception,* 191-217. Lawrence Erlbaum, NJ.

Hoffmann, J., & Ziessler, C. (1983). Objectidentifikation in kunstlichen begriffshierarchien [Object identification in artificial concept hierarchies]. *Zeitschrift fur Psychologie, 194*, 135-167.

Horton, M. S., & Markman, E. M. (1980). Developmental differences in the acquisition of basic and superordinate categories. *Child Development, 51*, 708-719.

Johnson, K. E., & Mervis, C. B. (1997). Effects of varying levels of expertise on the basic level of categorization. *Journal of Experimental Psychology: General, 126*, 248-277.

Jolicoeur, P., Gluck, M., & Kosslyn, S. M. (1984). Pictures and names: Making the connection. *Cognitive Psychology, 19*, 31-53.

Jones, G. V. (1983). Identifying basic categories. *Psychological Bulletin*, *94*, 423-428.

Lamberts, K. (1994). Flexible tuning of similarity in exemplar-based categorisation. *Journal of Experimental Psychology: Learning, Memory & Cognition, 20*, 1003-1021.

Mervis, C. B., & Crisafi, M. A. (1982). Order of acquisition of subordinate-, basic-, and superordinate-level categories. *Child Development, 53*, 258-266.

Malt, B. C. (1995). Category coherence in cross-cultural perspective. *Cognitive Psychology, 29*, 85-148.

Markman, E. M. (1989). *Categorization and Naming in Children: Problems of Induction.* Cambridge, Massachusetts: MIT Press.

Markman, E. M., & Hutchinson (1984). Children's sensitivity to constraints on word meanings: Taxonomic vs. thematic relations. *Cognitive Psychology, 16*, 1-27.

Medin, D. L., Goldstone, R. L. & Gentner, D. (1993). Respects for similarity. *Psychological Review*, *100*, 254-278.

Medin, D. L., & Schaffer, M. M. (1978). Context theory of classification learning. *Psychological Review, 85*, 207-238.

Morris, M. W., & Murphy, G. L. (1990). Converging operations on a basic level in event taxonomies. *Memory & Cognition, 18*, 407-418.Murphy, G. L., & Brownell, H. H. (1985). Category differentiation in object recognition: Typicality constraints on the basic category advantage. *Journal of Experimental Psychology: Learning, Memory and Cognition, 11*, 70-84.

Murphy, G. L. (1991). Parts in objects concepts: Experiments with artificial categories. *Memory & Cognition*, *19*, 423-438.

Murphy, G. L., & Lassaline, M. E. (1997). Hierarchical structure in concepts and the basic level of categorization. In K. Lamberts, D. R. Shanks, et al. (Eds.) *Knowledge, Concepts and Categories: Studies in Cognition* (pp. 93-131). Cambridge, Massachussett: MIT Press.

Murphy, G. L. & Smith, E. E. (1982). Basic level superiority in picture categorization. *Journal of Verbal Learning and Verbal Behavior*, *21*, 1-20.

Newport, E. L. and Bellugi, U. (1978). Linguistic expression of category in a visual-gestural language: A flower is a flower is a flower. In E. Rosch & B. B. Lloyd (Eds.), *Semantic factors in cognition* (pp. 137-168). Hillsdale, NJ: Erlbaum.

Norman, G. R., Brooks, L. R., Coblentz, C. L., & Babcock, C. J. (1992). The correlation of feature identification and category judgments in diagnostic radiology. *Memory & Cognition*, *20*, 344-355.

Oliva, A. & Schyns, P. G. (1997). Coarse blobs and fine edges? Evidence that information diagnosticity changes the perception of complex visual stimuli. *Cognitive Psychology*, *34*, 72-107.

Pothos, E. M., & Chater, N. (1998). Rational categories. In *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 848-853). New Jersey: Lawrence Erlbaum Associates, Publishers.

Quine, W. V. O. (1977). Natural kinds. In S. P. Swartz (Ed.), *Naming, necessity, and natural kinds*. Ithaca, New York: Cornell University Press.

Rensink, R. A., O'Regan, J. K. & Clark, J. J. (1997). To see or not to see: The need for attention to perceive changes on scenes. *Psychological Science*, *8*, 368-373.

Rifkin, A. (1985). Evidence for a basic level in event taxonomies. *Memory & Cognition, 13*, 538-556.

Rosch, E. & Mervis, C. B. (1975). Family ressemblances: Studies in the internal structure of categories. *Cognitive Psychology*, *7*, 573-605.

Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M. & Boyes-Braem, P. (1976). Basic objects in natural categories. *Cognitive Psychology, 8*, 382-352.

Schyns, P. G. (1998). Diagnostic recognition: Task constraints, objects information, and their interactions. *Cognition, 67*, 147-179.

Schyns, P. G. & Murphy, G. L. (1994). The ontology of part representation in object concepts. *The Psychology of Learning and Motivation*, *31*, 305-349.

Schyns, P. G. & Oliva, A. (1999). Dr. Angry and Mr. Smile: When categorization flexibly modifies the perception of faces in rapid visual presentations. *Cognition*, *69*, 243-265.

Schyns, P. G. & Rodet, (1997). Categorization creates functional features. *Journal of Experimental Psychology: Learning, Memory & Cognition, 23, 1-16.*

Schyns, P. G., Goldstone, R. L., & Thibaut, J. P. (1998). The development of features in object concepts. *Behavioral and Brain Sciences*, *21*, 1-53.

Shaver, P., Schwarz, J. Kirson, D. and O'Connor, D. (1987). Emotion knowledge: further explorations of prototype approach. *Journal of Personality and Social Psychology*, *52*, 1061-1086.

Simons, D. J. & Levin, D. T. (1997). Change Blindness. *Trends in Cognitive Sciences*, *1*, 261-267.

Smith, E. E. & Medin, D. L. (1981). *Categories and Concepts*. Cambridge, Massachusetts: Harvard University Press.

Tanaka, J. W. & Taylor, M. (1991). Object categories and expertise: Is the basic level in the eye of the beholder? *Cognitive Psychology*, *23*, 457-482.

Tarr, M. J., Bülthoff, H. H., Zabinski, M. & Blanz, V. (1997). To what extent do unique parts influence recognition across changes in viewpoint? *Psychological Science*, *8*, 282-289.

Thibaut, J. P. & Schyns, P. G. (1995). The development of feature spaces for similarity and categorization. *Psychologica Belgica*, *35*, 167-185.

Tversky, B. & Hemenway, K. (1984). Objects, parts and categories. *Journal of Experimental Psychology: General*, *113* (2), 169-191.

Wisniewski, E. J., & Murphy, G. L. (1989). Superordinate and basic category names in discourse: a textual analysis. *Discourse Processing, 12*, 245-261.

Woodman, G. F. & Luck, S. J. (1999). Electrophysiological measurement of rapid shifts of attention during visual search. *Science, 400*, 867-869.

Appendix: a formal implementation of SLIP

A strategy to categorize an object at a given level of abstraction is defined here as a list of features. Typically, some of these features are unique to this category and some overlap with the defining features of other categories. An optimal strategy is the shortest series of tests on the features defining the category. We posit that SLIP categorizers always use optimal strategies. We call redundant features, or set of redundant features, the collection of features which, individually, provide exactly the same information as to the category membership of objects. In other words, testing one, two, or more redundant features does not provide more information.

Formally, we will say that a strategy is a series of sets of redundant features. It has succeeded whenever all sets of redundant features have been completed *in a specific order*. And a set of redundant features is completed as soon as a test on the presence of one of its redundant features has been performed.

This usually happens after a succession of misses. The probability of having t-1 successive misses is given by $\left(1-\pi_j\right)^{t-1}$ where π_j–when redundancy of sets of features and the number of possible configurations that these can take in objects are taken into account–is equal to $C_j\left(1-S\right)+C_jSR_j$ that is, the practicability of set of redundant features j or the probability that it will be completed after a single attempt. S is the probability of a random slip (it was arbitrarily set to .5 throughout the simulations), and C_j is the probability that the target features will be in the expected configuration (1 / number of configurations). Thus the first term of π_j is the probability that the SLIP categorizer will guess the feature configuration correctly and that it will not slip. R_j is the probability that a random slip will result in a diagnostic test ([cardinality of j] / [number of features in objects]). The second term of π_j is the probability that the categorizer will slip, but that it will guess the correct configuration and will perform a diagnostic feature test.

The probability of a hit is 1 minus the probability of a miss. Thus, the probability that the set of redundant features j will be completed after t trials is

$$\left(1-\pi_j\right)^{t-1}\pi_j,$$

and the probability that a strategy of length n will have succeeded after t trials in a certain configuration of hits and misses is

$$\prod_{j=1}^{n}\left(1-\pi_j\right)^{\square}\pi_j,$$

where $\square$ is a function of j (it will remain unspecified) which gives the number of misses for the jth set of redundant features for that particular configuration. Usually, more than one such configuration exist. In fact, the number of possible configurations is easy to compute. The last hit necessarily happens at the tth trial; the n-1 other hits, however, can happen anywhere in the t-1 trials left, *in order*. Therefore, the number of possible configurations is the number of *combinations* of t-1 items taken n-1 by n-1 that is,

$$\square = \square\begin{matrix} t \square 1 \\ n \square 1 \end{matrix}\square = \frac{(t \square 1)!}{(t \square n)!(n \square 1)!}.$$

We can now give the global shape of the probability that a strategy of lengths n will succeed after t trials:

$$\mathop{\square}_{i=1}^{\square} \mathop{\square}_{j=1}^{n} \left(1 \square \square_j\right)^{\square} \square_j ,$$

where $\square$ is a function of i and j that specifies the number of misses for the jth set of redundant features for the ith configuration of hits and misses. We call this the Response Time Function (RTF). We still have to specify $\square$. We will establish a connection between this function and multinomial expansions. The multinome $\left(a_1 + a_2 + \ldots + a_n\right)^{t \square n}$ expands into $\square$ different terms, and the sum of the n exponents of each term is equal to t-n. It follows that $\square$ gives the jth exponent of the ith term in this multinomial expansion.

As a global measure of basic-levelness, we use *t_mean*, the mean number of tests required to complete a strategy. When internal practicability is constant within a strategy (this is true for all experiments reported in this chapter), the RTF is a Pascal density function and, thus, *t_mean* is equal to $n/\square$.

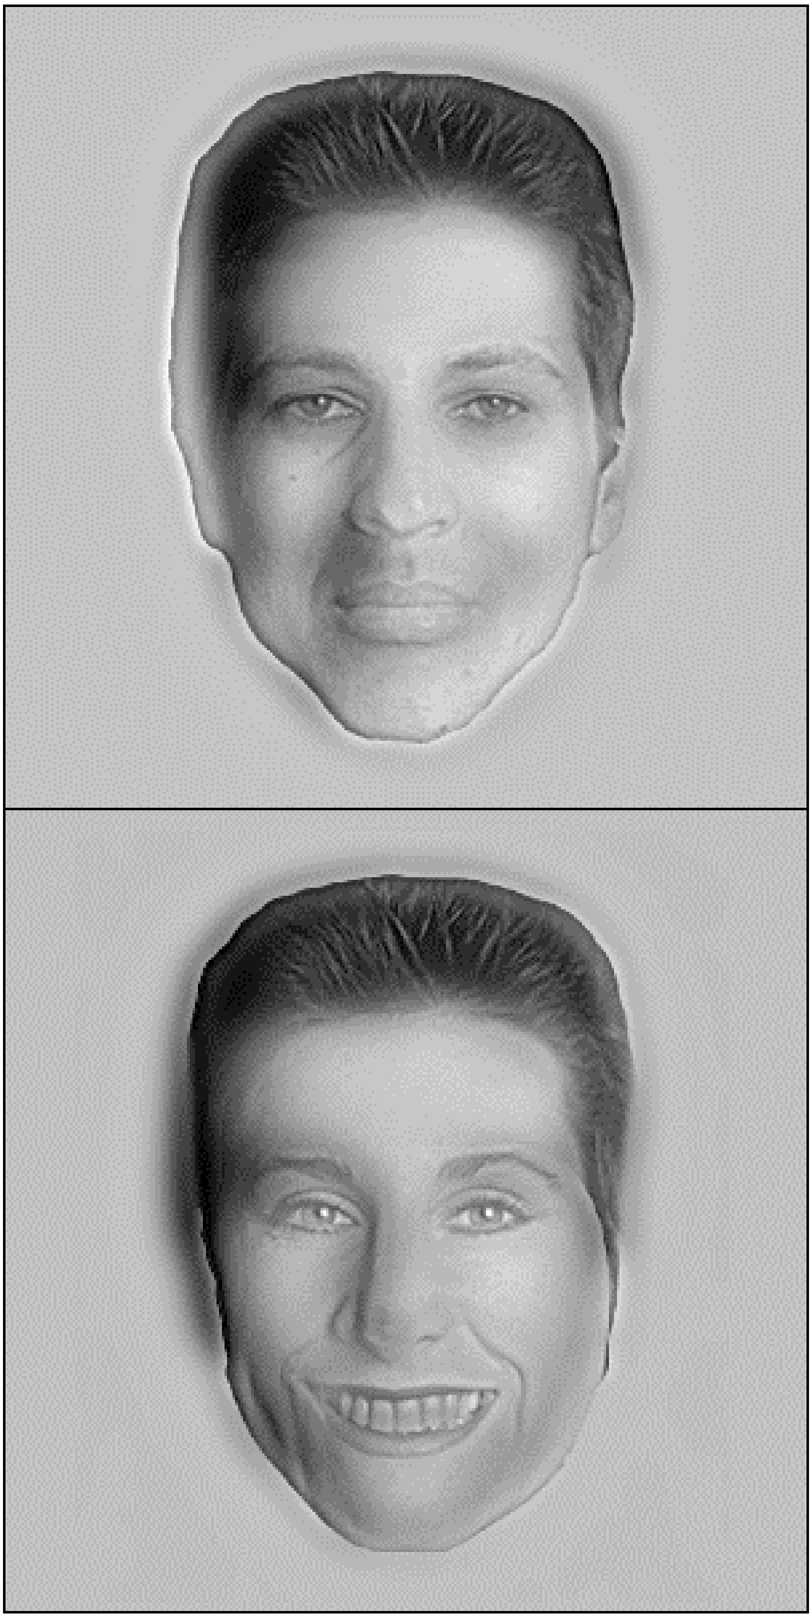

Figure 1. The top picture is a neutral male and the bottom picture a smiling female. If you squint, blink, or step away from the pictures, the opposite situation should appear, namely the smiling female at the top and the neutral male at the bottom. This occurs because squinting and blinking modify the

relative availability of high spatial frequencies, which represent the initial perception of the pictures, whereas low spatial frequencies represent the other one

.

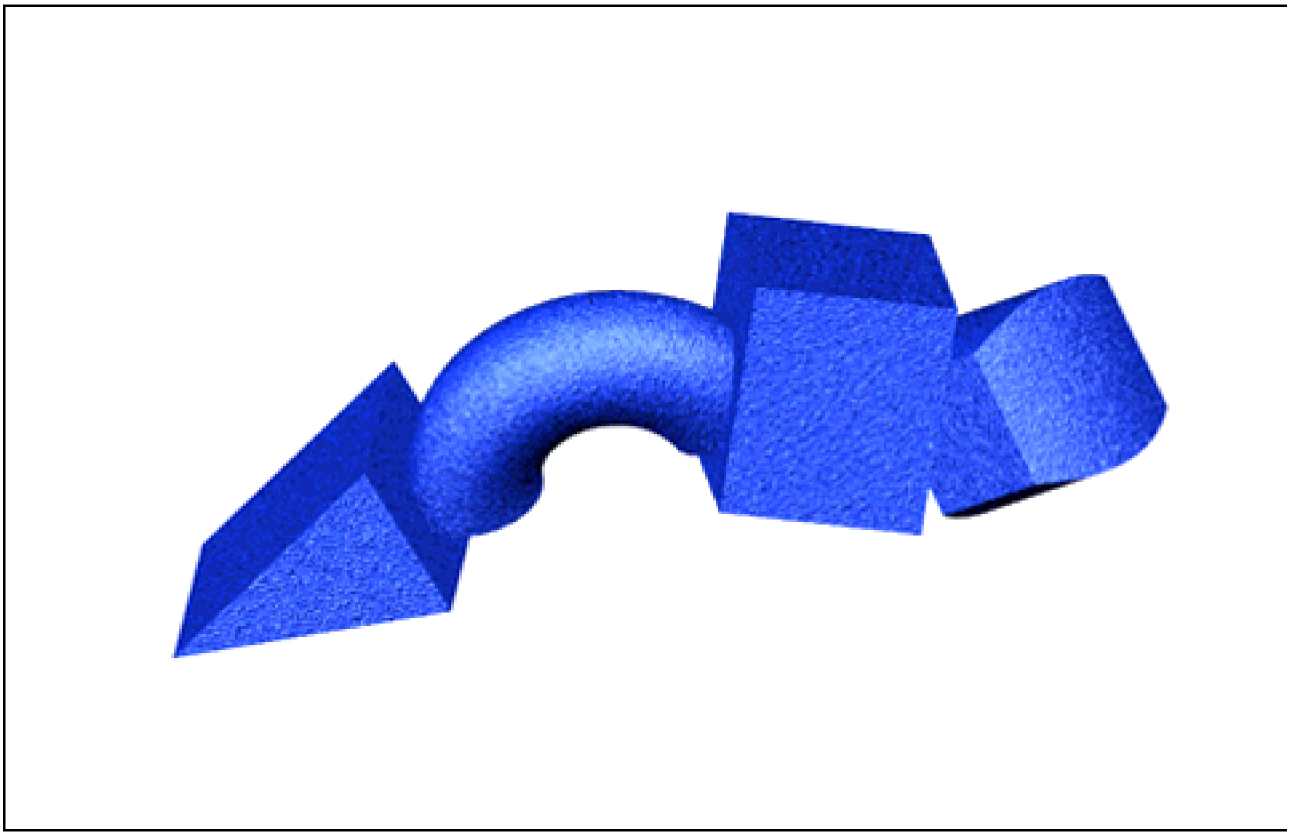

Figure 2. An artificial object used in Gosselin & Schyns, 1999a, Experiments 1 and 2.

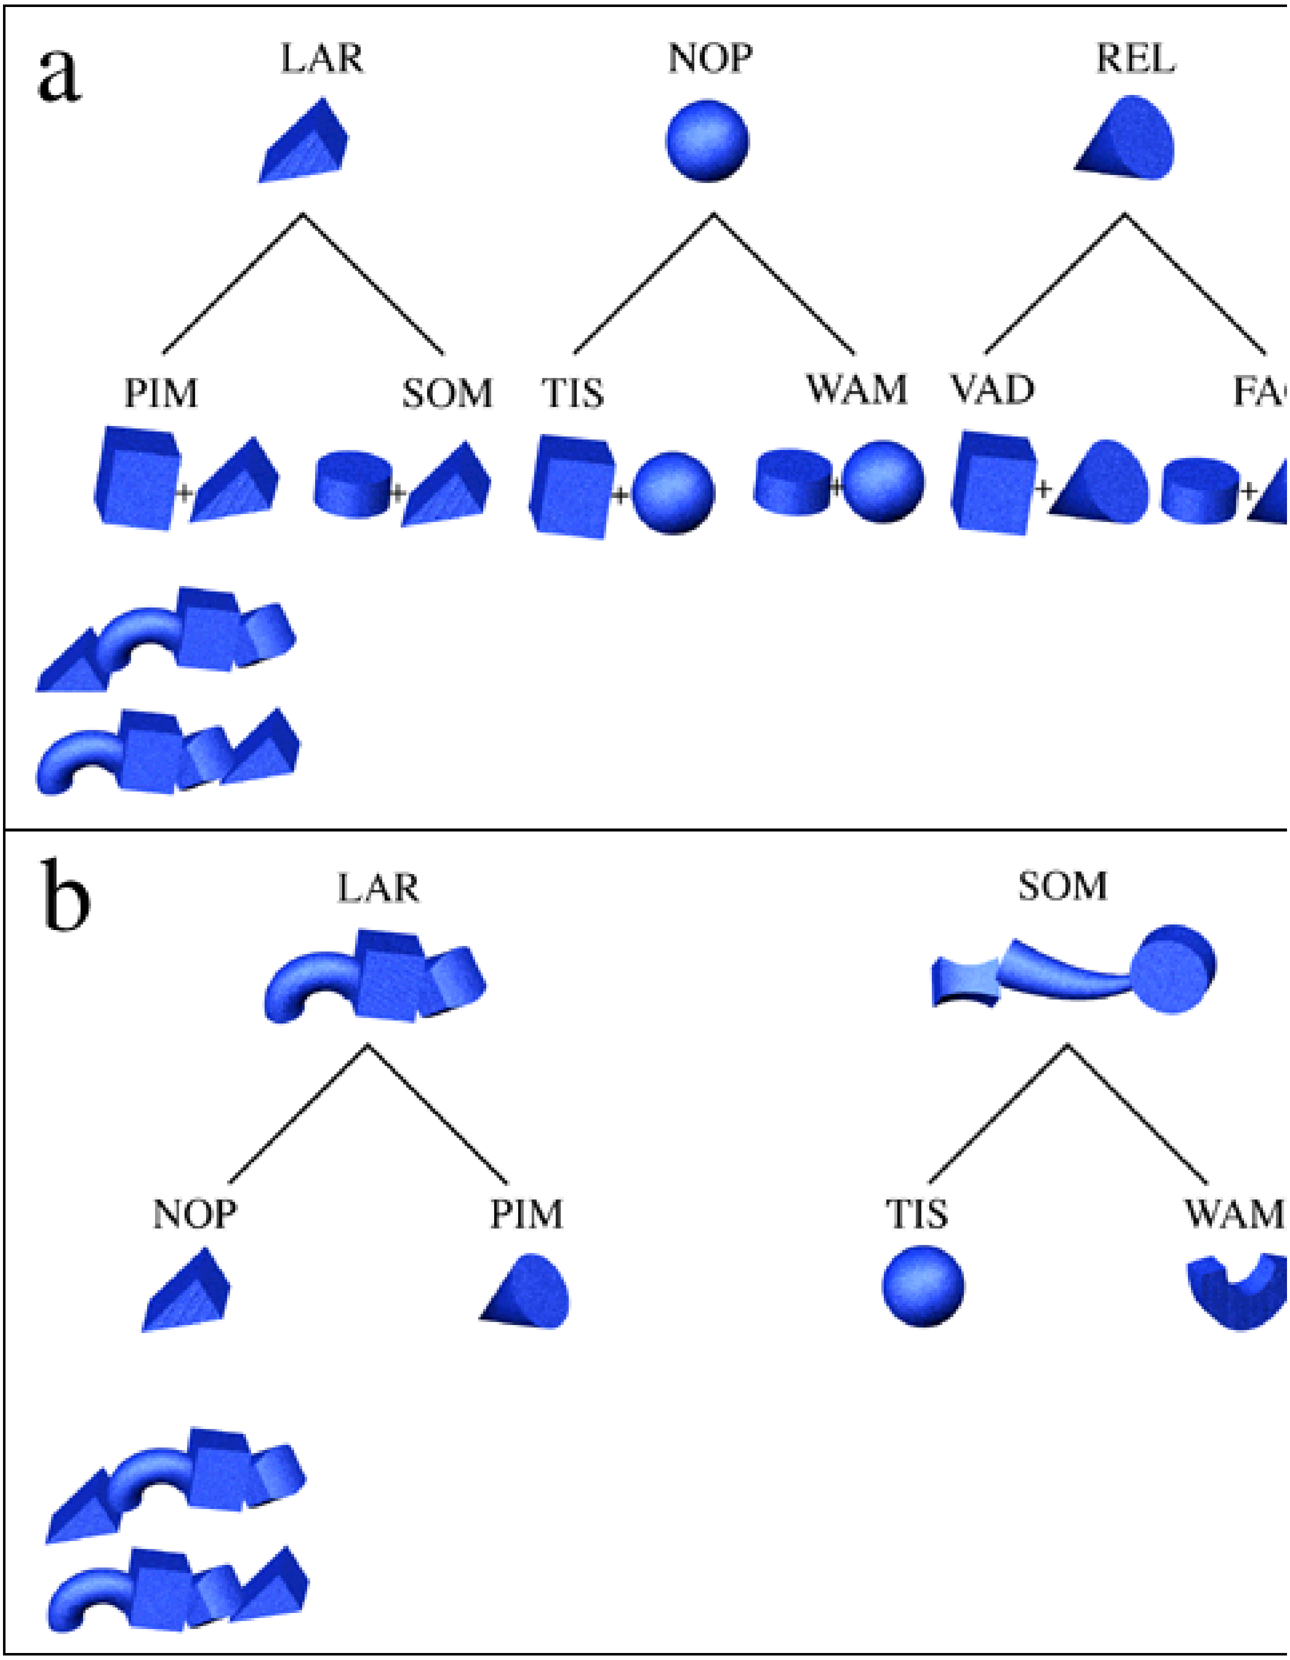
a
LAR
NOP
REL
PIM
SOM
TIS
WAM
VAD
b
LAR
SOM
NOP
PIM
TIS
WAM

Figure 3. (a) The HIGH_FAST taxonomy of Gosselin & Schyns, 1999a, Experiment 1. This illustrates a variation of strategy length. (b) The HIGH_FAST taxonomy of Gosselin & Schyns, 1999a, Experiment 2. This illustrates a change of internal practicability.

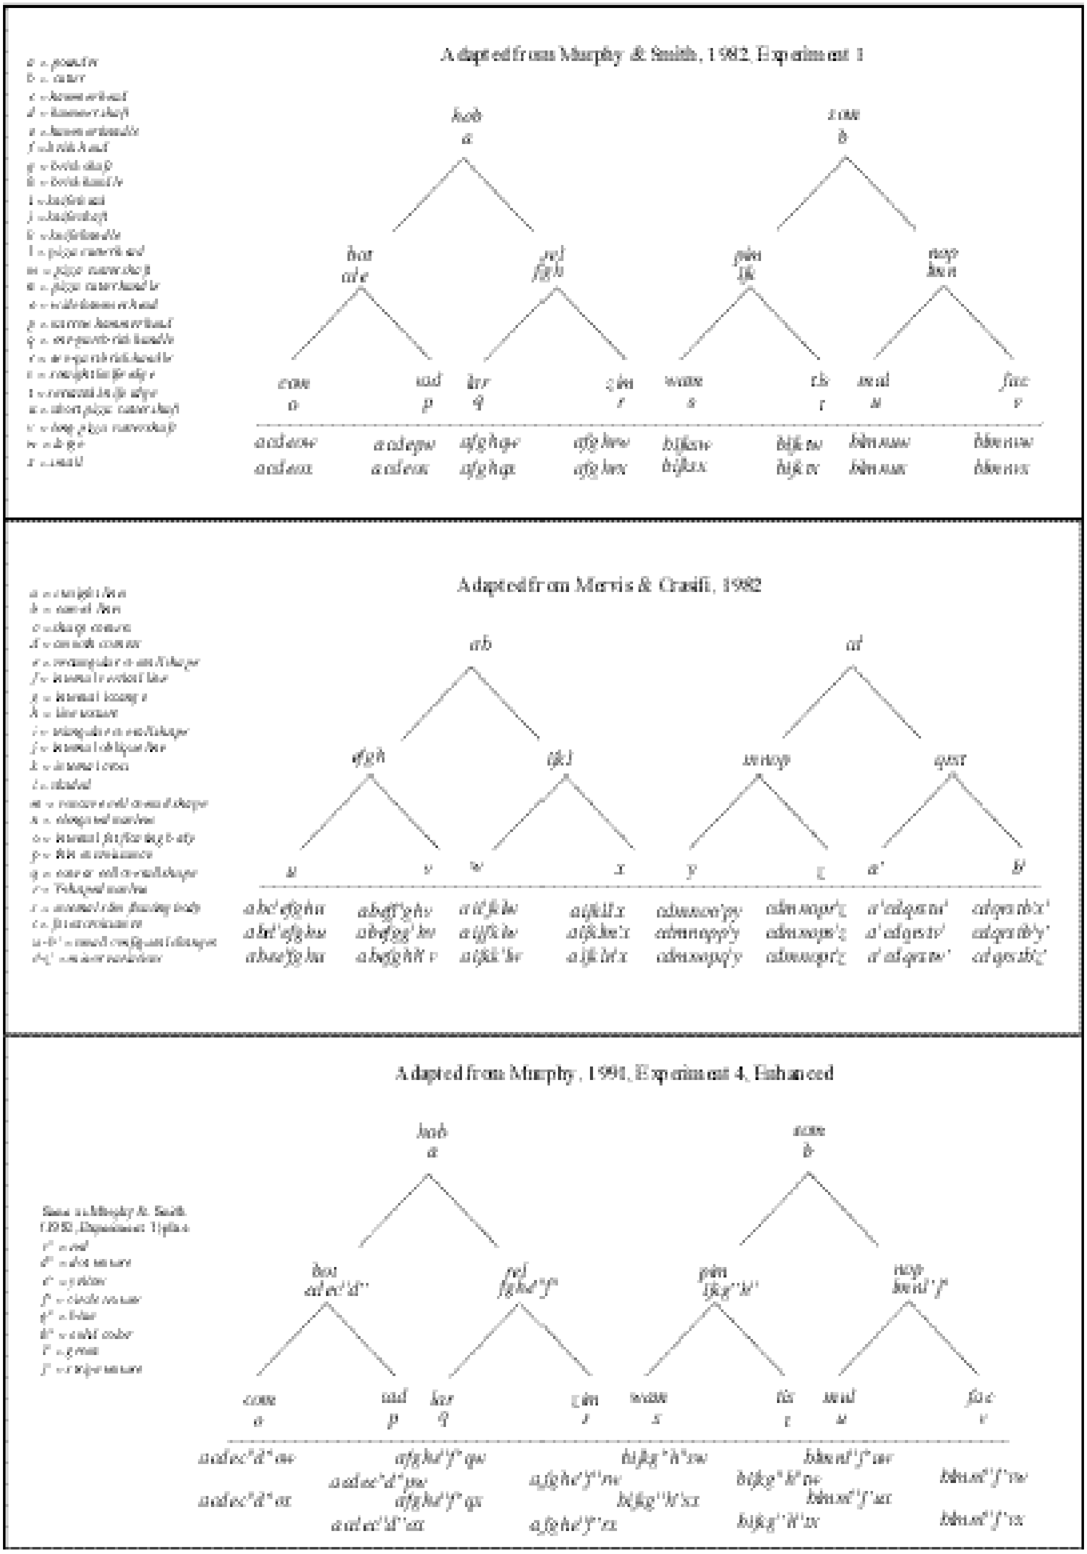

Figure 4. Taxonomies of all experiments with varying redundancy that exhibited an advantage for an intermediate level of categorization. Underneath the category names, we provide the optimal strategies fed

to SLIP. The feature constitution of all exemplar is giving underneath each taxonomy. An index for these abstract features is also provided.

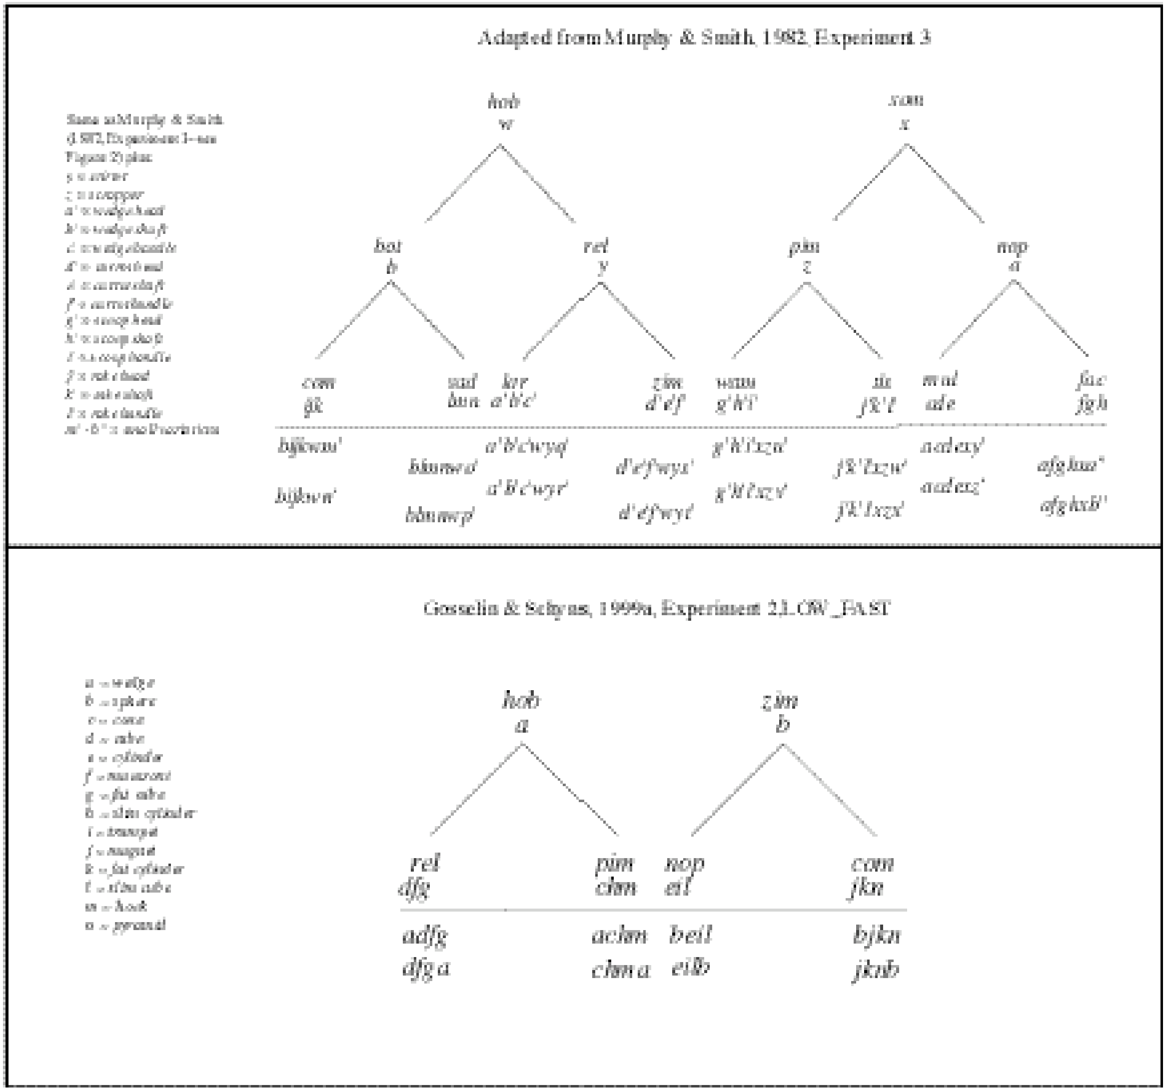

Figure 5. Taxonomies with varying redundancy that exhibited an advantage at the lower level of categorization. Underneath the category names, we provide the optimal strategies fed to SLIP. At the bottom of the taxonomy, the abstract feature constitution of all exemplars is given. An index for these abstract features is provided left of the taxonomy. The feature constitution of all exemplar is giving underneath the taxonomy.

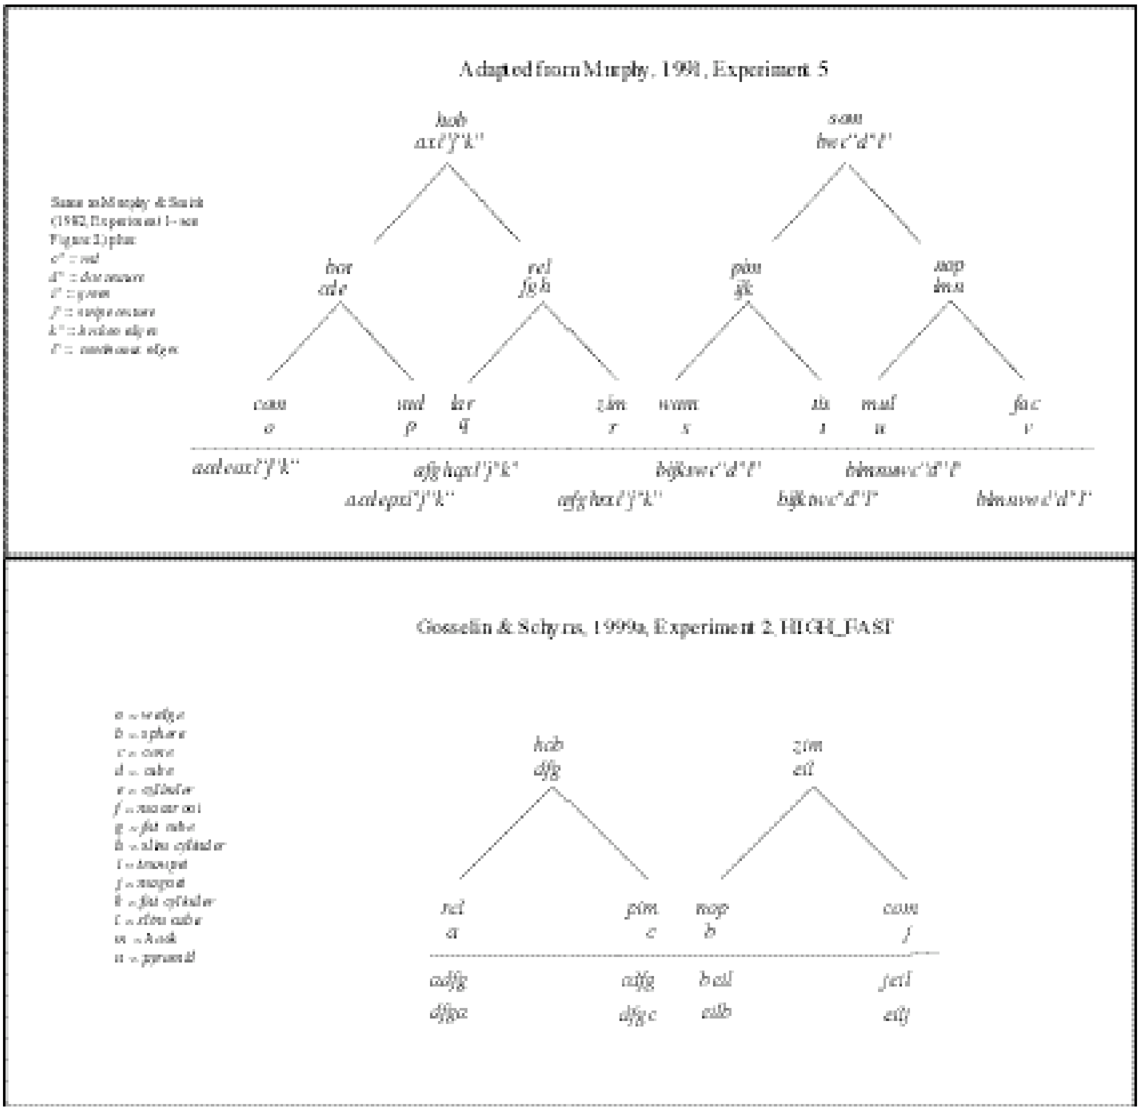

Figure 6. Taxonomies that exhibited an advantage at the higher level of categorization. Underneath the category names, we provide the optimal strategies fed to SLIP. At the bottom of the taxonomy, the abstract feature constitution of all exemplars is given. An index for these abstract features is provided left of the taxonomy. The feature constitution of all exemplar is giving underneath the taxonomy.

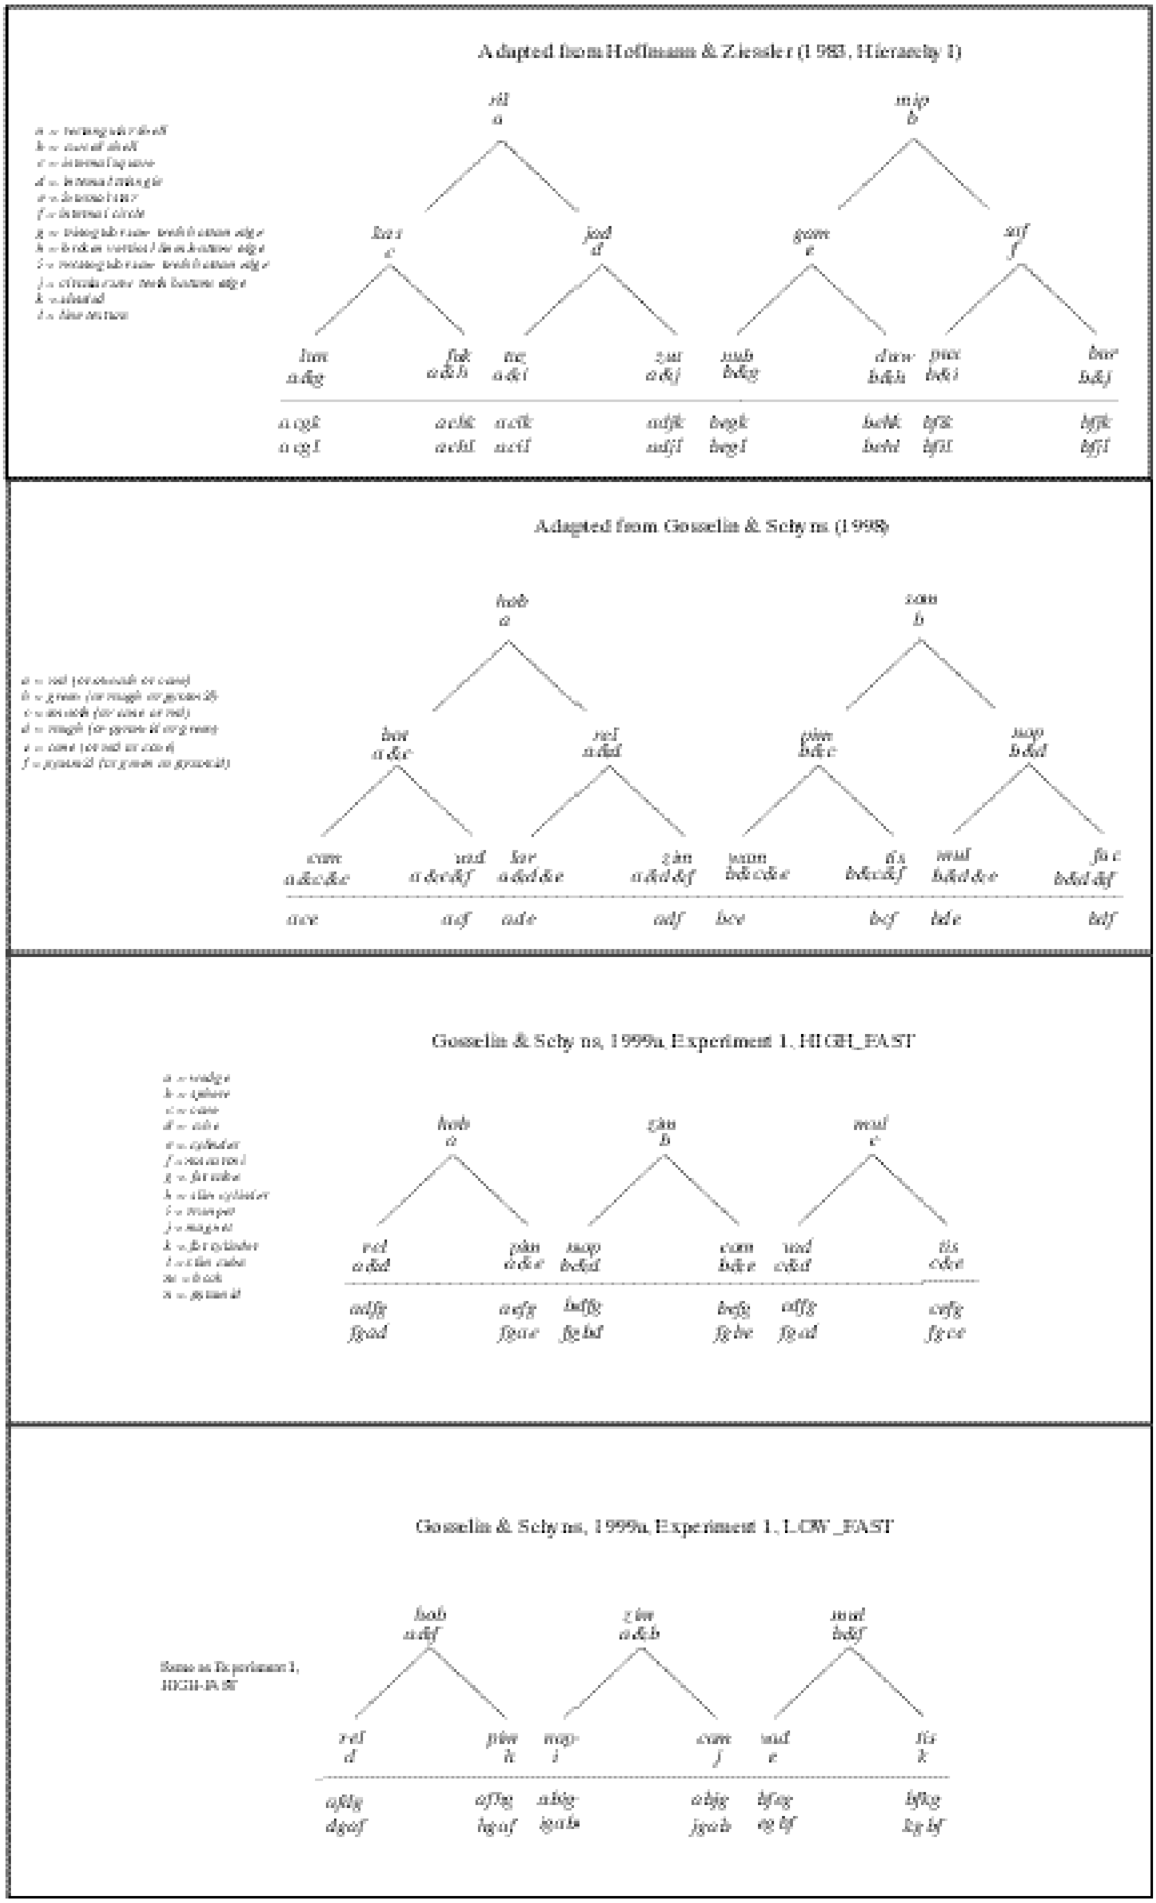
Adapted from Hoffmann & Ziessler (1983, Hierarchy I)
Adapted from Gosselin & Schyns (1998)
Gosselin & Schyns, 1999a, Experiment 1, HIGH_FAST
Gosselin & Schyns, 1999a, Experiment 1, LOW_FAST
Same as Experiment 1, HIGH-FAST

Figure 7. Abstract taxonomies of all experiments with varying strategy length. Underneath the category names, we provide the optimal strategies fed to SLIP. The feature constitution of all exemplar is giving underneath each taxonomy. An index for these abstract features is also provided.

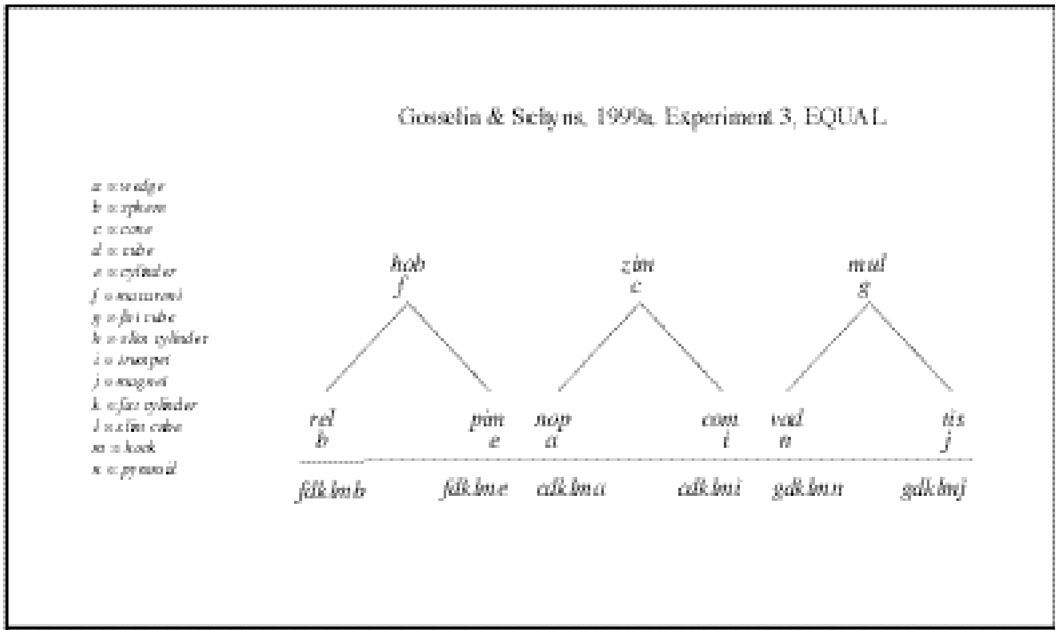
Gosselin & Schyns, 1999a, Experiment 3, EQUAL
a = wedge
b = sphere
c = cone
d = cube
e = cylinder
f = macaroni
g = flat cube
h = slim cylinder
i = trumpet
j = magnet
k = fat cylinder
l = slim cube
m = hook
n = pyramid
hob
f
zim
c
mul
g
rel
b
pim
e
nop
a
com
i
vad
n
tis
j
fdklmb
fdklme
cdklma
cdklmi
gdklmn
gdklmj

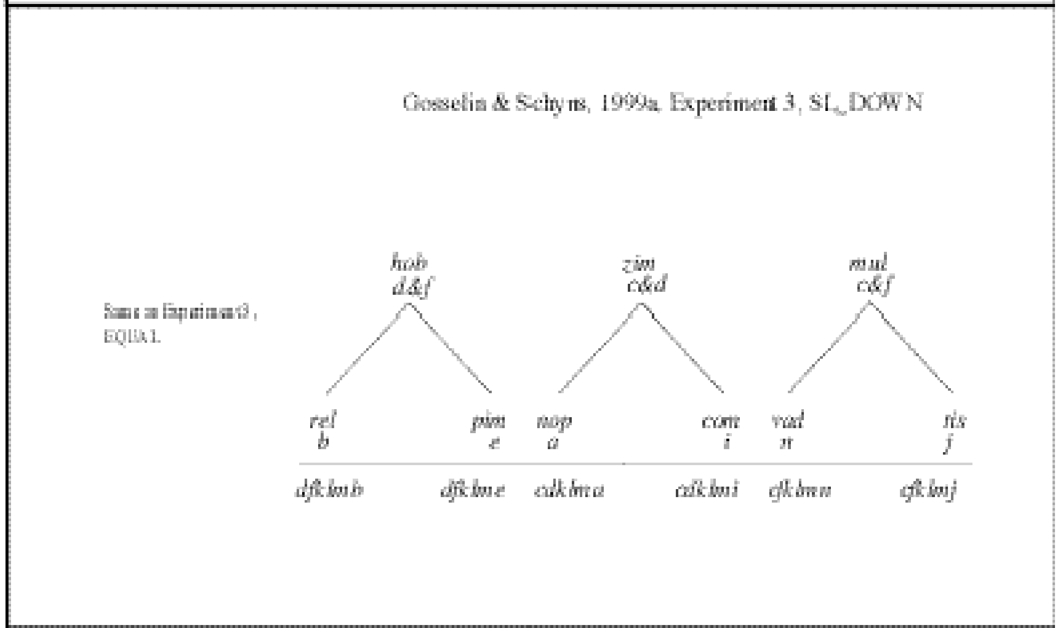
Gosselin & Schyns, 1999a, Experiment 3, SL_DOWN
Same as Experiment 3, EQUAL
hob
d&f
zim
c&d
mul
c&f
rel
b
pim
e
nop
a
com
i
vad
n
tis
j
dfklmb
dfklme
cdklma
cdklmi
cfklmn
cfklmj

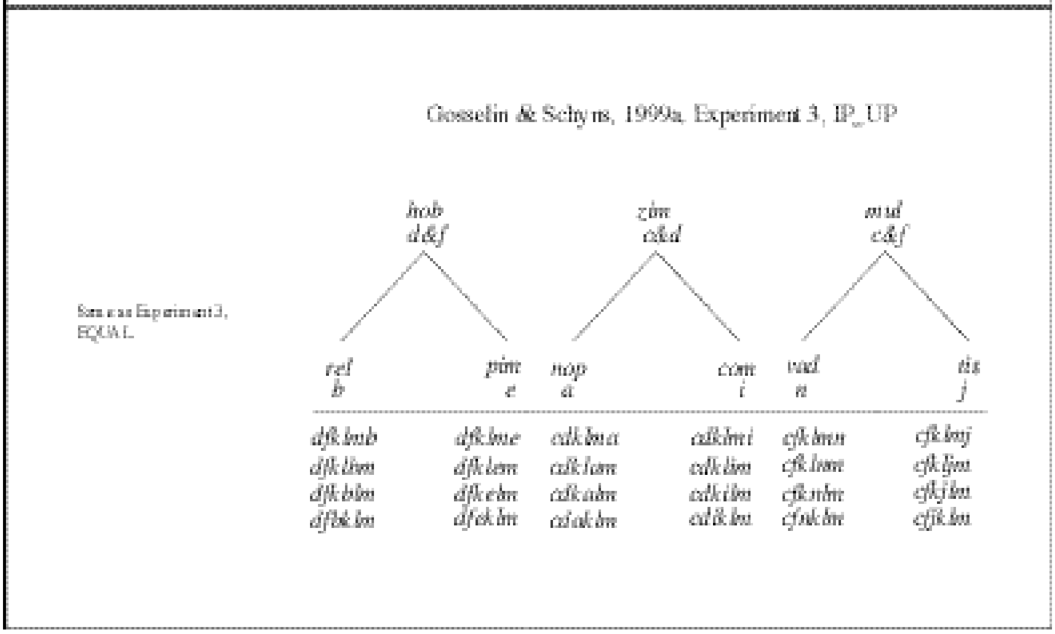
Gosselin & Schyns, 1999a, Experiment 3, IP_UP
Same as Experiment 3, EQUAL
hob
d&f
zim
c&d
mul
c&f
rel
b
pim
e
nop
a
com
i
vad
n
tis
j
dfklmb
dfklbm
dfkblm
dfbklm
dfklme
dfklem
dfkelm
dfeklm
cdklma
cdklam
cdkalm
cdaklm
cdklmi
cdklim
cdkilm
cdiklm
cfklmn
cfklnm
cfknlm
cfnklm
cfklmj
cfkljm
cfkjlm
cfjklm

Figure 8. Taxonomies of Experiment 3, EQUAL, SL_DOWN, and IP_UP (from top to bottom). Strategy length and internal practicability interacts here. Underneath the category names, we provide the optimal strategies fed to SLIP. The feature constitution of all exemplar is giving underneath each taxonomy. An index for these abstract features is also provided.

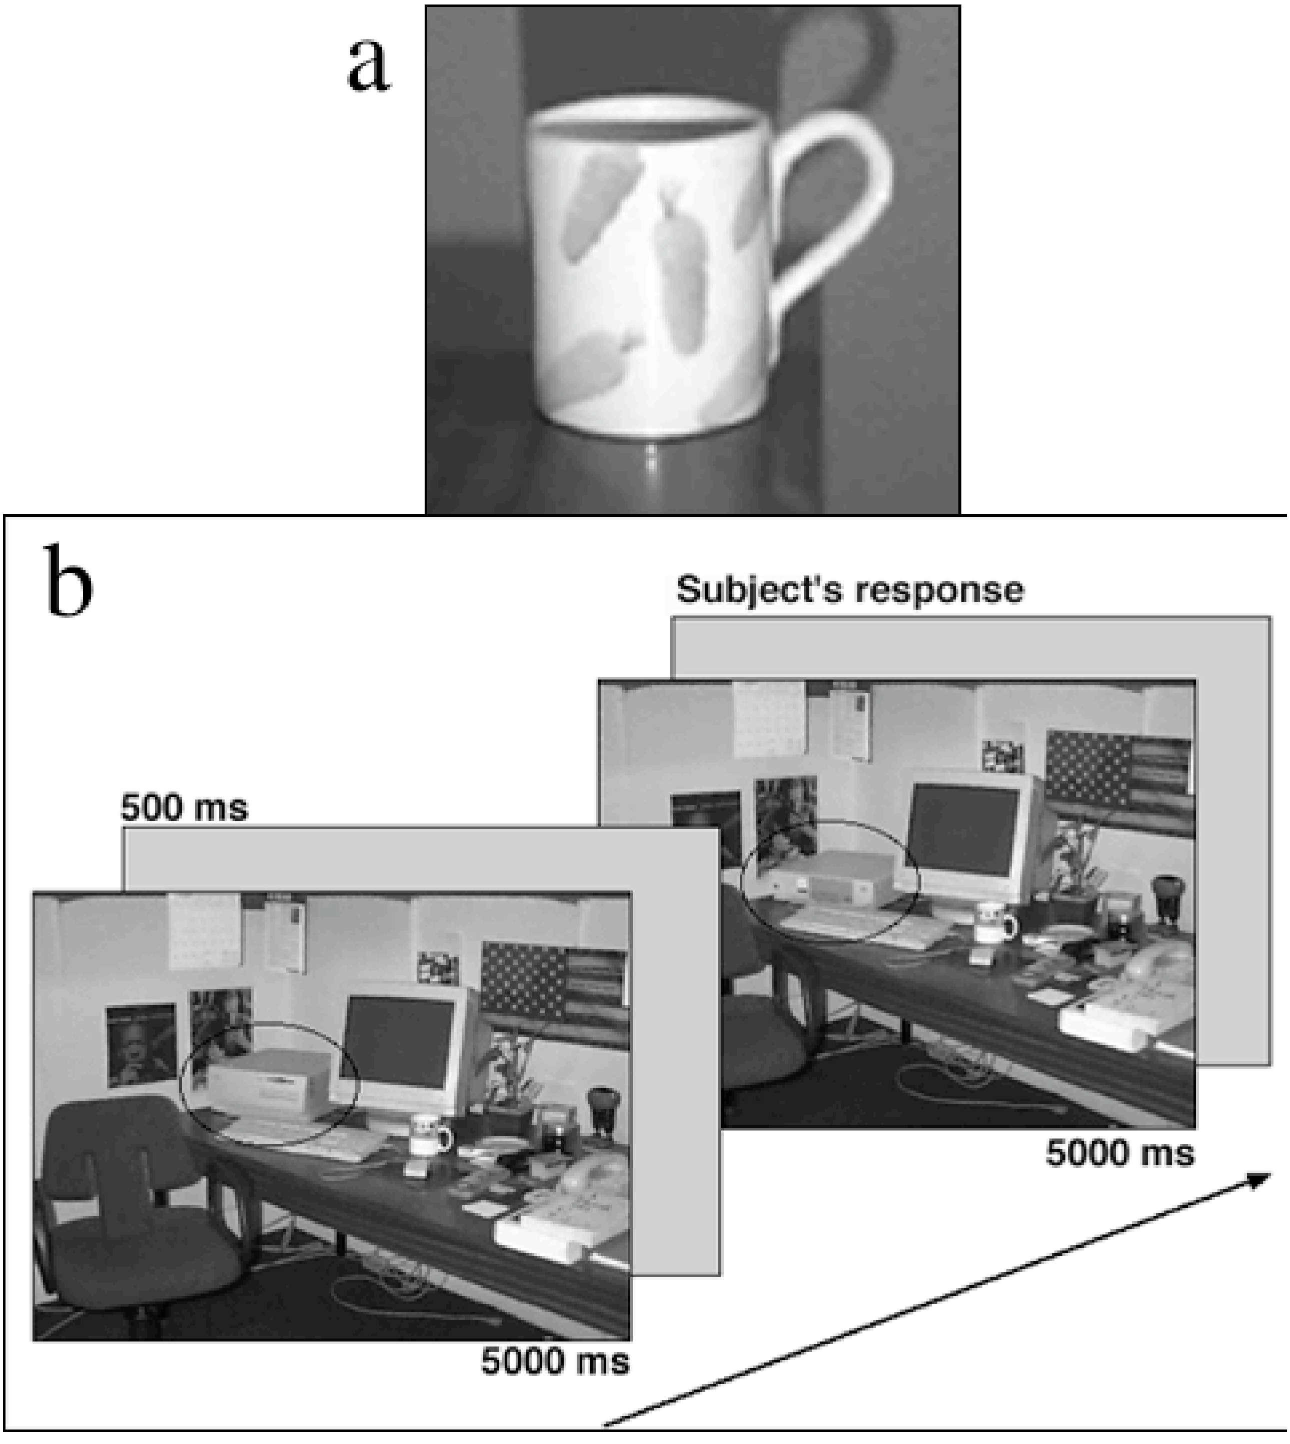

Figure 9. (a) Object known as *Peter's mug* to half the subjects and as a *mug* to the others. (b) This illustrates a trial of the change detection task–a computer change. Each frame of a two-frame sequence was presented for 5 sec, separated by a 500 ms blank. The two-frame sequence was repeated until subjects perceived the change and correctly identified it. The number of repetitions was used as an independent measure of change perception.

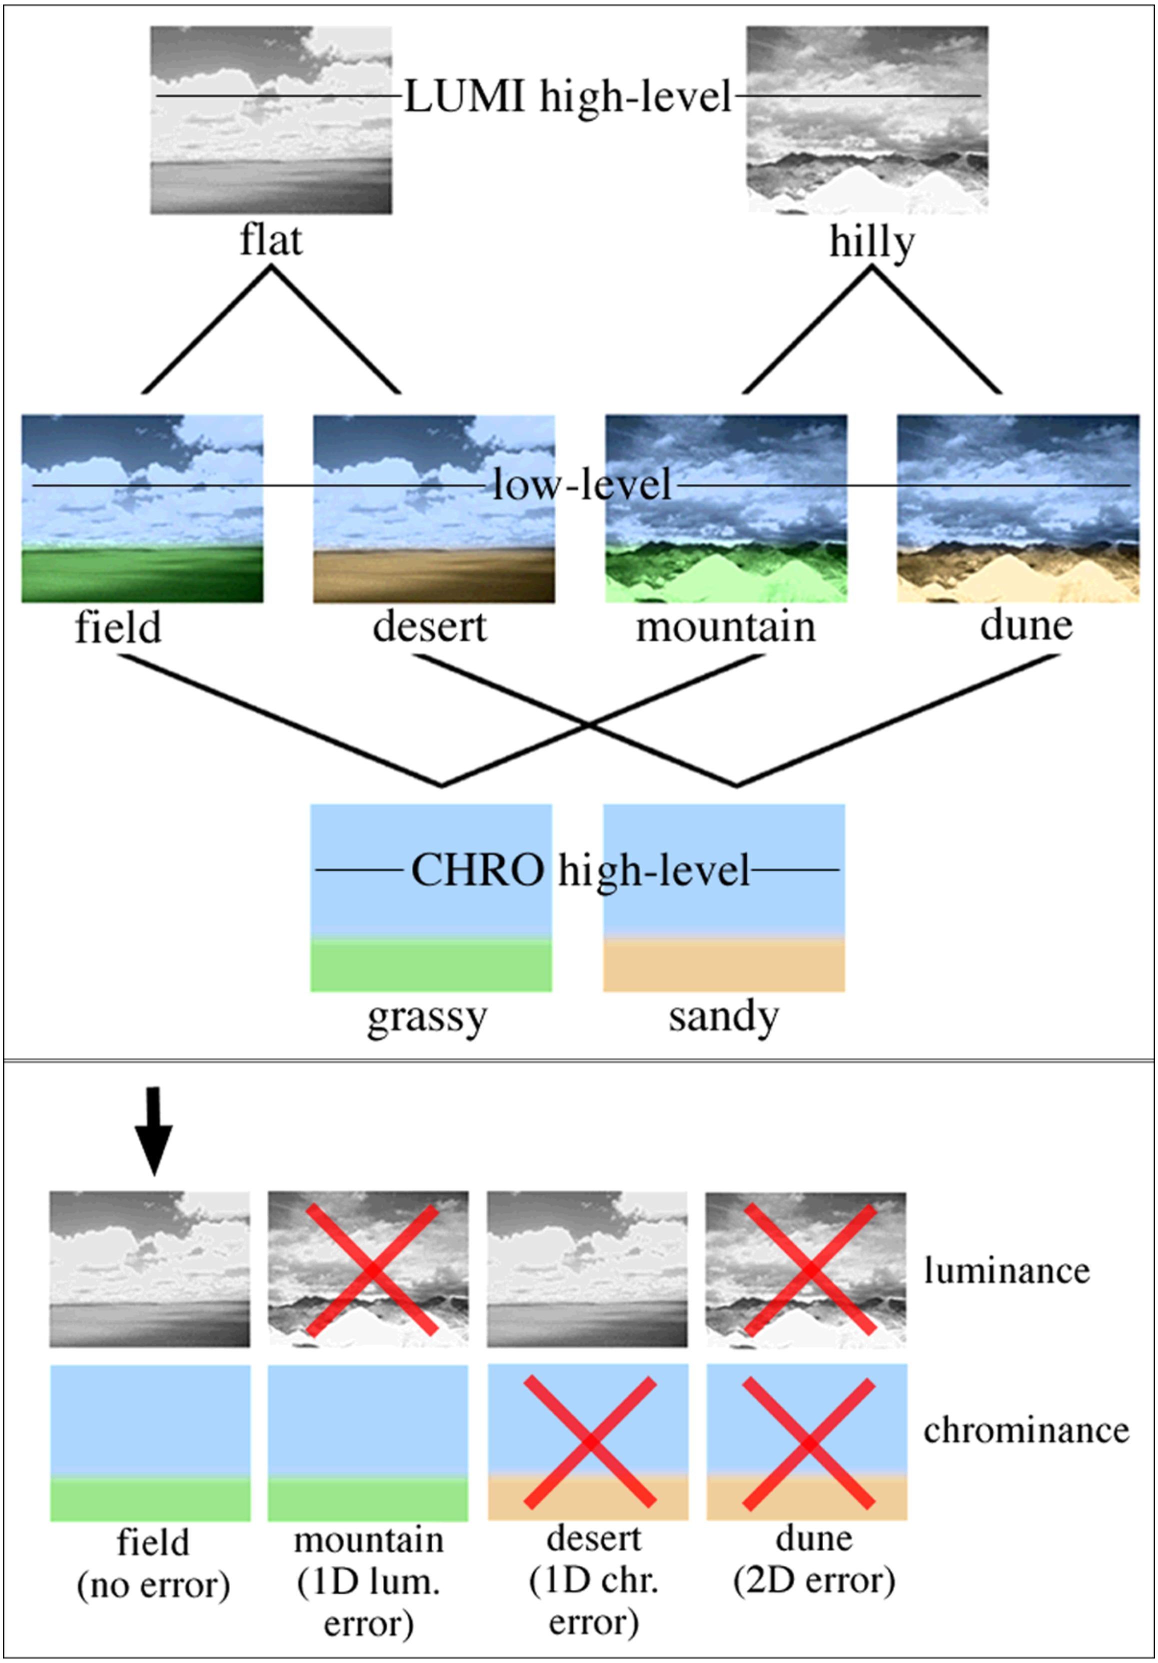

Figure 10. A the top, the four scenes used in this experiment and the corresponding low-level category names learned by all participants ("field", "mountain", "desert", and "dune"), sandwiched by the two high-

level categorizations ("flat" and "hilly") LUMI subjects learned, and those ("grassy" and "sandy") CHRO subjects learned.

Table 1.

Mean observed correct positive trials verification reaction times for various category structures and SLIP's predictions (with $S = .5$; see Appendix). For each experiment, we summed the absolute values of the differences between the predicted and actual categorization level ranks (see the rightmost column). The smaller this distance measure, the better the fit.

		Level				
Source		H - 3	H - 2	H - 1	H (Highest)	Dis
Murphy & Smith, Exp. 1	Observation		723 ms	678 ms	879 ms	
	SLIP		1.667	1.25	1.667	
Murphy, Exp. 4, Simple	Observation		862 ms	811 ms	980 ms	
	SLIP		1.667	1.25	1.667	
Mervis & Crasifi	Observation		3rd	1st	2nd	
	SLIP		1.75	1.273	1.556	
Murphy, Exp. 4, Enhanced	Observation		1,132 ms	854 ms	955 ms	
	SLIP		1.75	1.167	1.75	
Rosch et al., Exp. 7	Observation		659 ms	535 ms	591 ms	
	SLIP		1.714	1.333	1.5	
Tanaka & Taylor, Novice	Observation		778 ms	678 ms	746 ms	
	SLIP		1.588	1.385	1.543	
Johnson & Mervis, Songbird Novice	Observation	~2100 ms	~1950 ms	~1600 ms	~1900 ms	
	SLIP	1.990	1.886	1.212	1.557	
Johnson & Mervis, Songbird Intermediate	Observation	~1725 ms	~1600 ms	~1550 ms	~1800 ms	
	SLIP	1.673	1.493	1.443	1.853	
Johnson & Mervis, Songbird Expert	Observation	~1600 ms	~1625 ms	~1500 ms	~1750 ms	
	SLIP	1.630	1.466	1.535	1.808	
Murphy & Smith, Exp. 3, Size	Observation		574 ms	882 ms	666 ms	

	SLIP		1.25	1.667	1.667
Gosselin & Schyns, 1999a, Exp. 2, LOW_FAST	Observation			740 ms	774 ms
	SLIP			2.286	3.2
Tanaka & Taylor, Expert	Observation		622 ms	623 ms	729 ms
	SLIP		1.474	1.474	1.556
Murphy, Exp. 5	Observation		1,072 ms	881 ms	854 ms
	SLIP		1.8	1.5	1.286
Gosselin & Schyns, 1999a, Exp. 2, HIGH_FAST	Observation			788 ms	661 ms
	SLIP			3.2	2.286
Hoffmann & Ziessler, Exp. 1	Observation		~ 700 ms	~ 500 ms	~ 500 ms
	SLIP		3.2	1.6	1.6
Gosselin & Schyns, 1998, Overall	Observation		1,184 ms	1,012 ms	819 ms
	SLIP		4.5	3	1.5
Gosselin & Schyns, 1999a, Exp. 1, HIGH_FAST	Observation			1256 ms	896 ms
	SLIP			6.4	3.2
Gosselin & Schyns, 1999a, Exp. 1, LOW_FAST	Observation			948 ms	1240 ms
	SLIP			3.2	6.4
Gosselin & Schyns, 1999a, Exp. 3, EQUAL	Observation			672 ms	680 ms
	SLIP			1.714	1.714
Gosselin & Schyns, 1999a, Exp. 3, SL_DOWN	Observation			920 ms	1058 ms
	SLIP			1.714	3.429
Gosselin & Schyns, 1999a, Exp. 3, IP_UP	Observation			928 ms	775 ms
	SLIP			6.857	3.429