The Determinants of Basic-Level Performance

Frédéric Gosselin (GOSSELIF@PSY.GLA.AC.UK)
Philippe G. Schyns (PHILIPPE@PSY.GLA.AC.UK)
Department of Psychology, University of Glasgow
58 Hillhead St., Glasgow G12 8QB UK

Abstract

SLIP (*Strategy Length & Internal Practicability*) is a new model of basic-level performance that postulates two computational constraints on the basic-levelness of a category: the number of feature tests required to place the input in a category (*Strategy Length*) and the ease with which these tests are performed (*Internal Practicability*). This article reports three experiments that examined the validity of SLIP in two-level taxonomies of computer-synthesized artificial objects. Experiment 1 isolated *strategy length*, Experiment 2, *practicability*, and Experiment 3 explored the interactions of these factors. Whereas SLIP predicted the RT of these experiments, two established basic-level models of basic-level performance, Jones's (1983) *category feature-possession* and Corter and Gluck's (1991) *category utility*, did not.

What distinguishes your cellular phone, your fountain pen, your computer, your house, and other everyday objects of yours from those of your neighbors is often a combination of features. For example, to identify your pink Porsche 911 in a parking lot also comprising a pink Toyota Tercel and a lime Porsche 911, you must examine both the *color* and the *shape* of the cars. This is so because real-world things share many features. The hierarchical organization of categories is a direct consequence of this sharing of features.

In a seminal article, Rosch, Mervis, Gray, Johnson & Boyes-Braem (1976) distinguished between three natural levels of categorization hierarchy (or taxonomy), the subordinate (e.g., “Porsche 911”), basic (e.g., “car”) and superordinate (e.g., “vehicule”), from the most specific to the most general. Of all these levels, they showed that the basic was the best in many respects. People tend to: designate things by their basic-level names; list many more features at the basic level than at any other level; decide more rapidly that things are members of their basic categories than of any others; and so on.

SLIP (Strategy Length & Internal Practicability) is designed to model one of the most important index of basic-level performance: categorization speed. It postulates an ideal categorizer that performs the fewest possible number of features tests to classify things. Its name is derived from the fact that its attention slips off its ideal track once in a while.

Going back to the parking lot example, you had to check both the *color* and the *shape* of cars to find your pink Porsche 911. Fewer tests would have not lead to a definitive decision (because there were also a lime Porsche 911 and a pink Toyota Tercel). We call this optimal series of tests a *strategy*. Two aspects of strategies fully determine the response time of SLIP: their *length* and their *internal practicability*. Strategy length is simply the minimal number of tests required to complete a strategy. In the parking lot example, strategy length is equal to 2 (one test on *shape*; one test on *color*). The longer a strategy associated with a category, the more time it will take to categorize an object in this category. The second factor of SLIP–*internal practicability* (or *practicability*, for short)–is simply the ease which with a particular test in a strategy can be executed (e.g., the number of features that uniquely define this category). The greater the practicability of a category, the less time it will take to verify that an object belongs to this category. SLIP integrates strategy length and internal practicability to predict categorization time (see the Appendix for formal details).

In Gosselin and Schyns (1997, 1999) we demonstrated that the principles of SLIP better predict the results of 22 classic basic-level experiments (from Rosch et al., 1976; Murphy and Smith, 1982; Mervis & Crisafi, 1982; Hoffmann & Ziessler, 1983; Murphy, 1991; Lassaline, 1990; Tanaka & Taylor, 1991; Johnson & Mervis, 1997; and Gosselin & Schyns, 1998) than the leading models (e.g., Jones's, 1983, *category feature-possession* and Corter & Gluck's, 1991, *category utility*).

No matter how successful, these simulations are *a posteriori* accounts of data. The validity of SLIP would be better tested with a direct examination of strategy length and internal practicability[1]. In three experiments, we isolated the possible role of these factors on basic-levelness. Specifically, Experiment 1 isolated the effect of strategy length, Experiment 2, the effect of internal practicability, and Experiment 3 the interactions of the two factors.

Experiment 1

Experiment 1 examines the effect of *strategy length* on basic-levelness. Strategy length is the minimum number of required feature tests to perform a given categorization.

Experiment 1 is set up as a category verification task of two two-level taxonomies of artificial objects (see Figure 2). The taxonomies are designed to induce orthogonal patterns of categorization speed across conditions. In the HIGH_FAST taxonomy, overlap of geons between categories defines a shorter strategy at the higher than at the lower level. In the LOW_FAST taxonomy, different geon arrangements yield the reverse situation—i.e. longer strategies at the top than at the bottom level. SLIP predicts that shorter strategies are completed faster than longer strategies, irrespective of categorization levels.

[1] Furthermore, it must be noted that our data set of 22 published experiments is itself biased to mid-then-high-then-low level. Any model that predicts this RT sequence will be 58% right, irrespective of the actual hierarchy.

Hence, on the basis of only strategy length, SLIP predicts orthogonal categorization performance across taxonomies.

Method

Participants

Twenty University of Glasgow students with normal or corrected vision were paid to participate in the experiment.

Stimuli

Stimuli were computer-synthesized chains of four geons (see Figure 1) similar to those used in Tarr, Bülthoff, Zabinski and Blaz (1997) in the context of object recognition. We designed stimuli with a three-dimensional modeling software on a Macintosh computer.

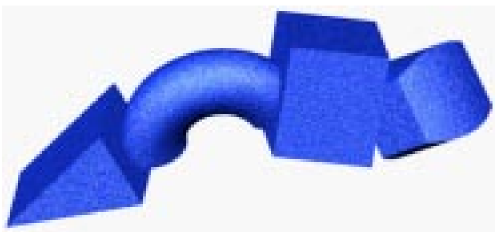

Figure 1. A four-geon chain used in Experiment 1.

Five geons defined the categories of the HIGH_FAST taxonomy. One different geon defined each one of three high-level categories. Each one of six possible low-level categories was further specified by one of the two remaining geons. Figure 2 illustrates this taxonomy.

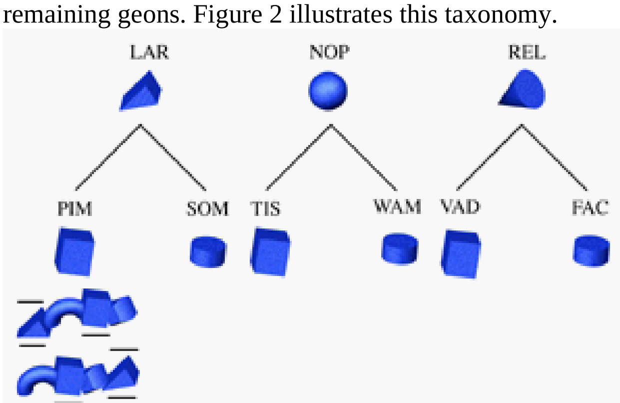

Figure 2. The HIGH_FAST taxonomy of Experiment 1. The geons specify the defining information of each category. The bottom geon chains are the two PIM exemplars (they are also LAR exemplars) used in the experiment.

In Figure 2, strategy length equals 1 for the higher-level categories. A length 1 strategy means that only one feature needs to be tested (the feature defining each high-level category) to determine the membership of the objects at this level. Strategy length equals 2 at the lower-levels, because these categorizations require two feature tests. The longer strategies arise from the overlap of features across lower-level categories. Shortly put, in the taxonomy of Figure 2, lower-level categorizations require conjunctions of features to be tested. This difference between strategy lengths across the levels of a taxonomy is the backbone of Experiment 1. To create the actual experimental stimuli, we further added two geons that served as *fillers* to obtain six four-geon objects. Fillers were identical across objects and so could not be used to categorize them. We created two exemplars per low-level category by changing the location of the diagnostic geons in the chain (see Figure 2, the bottom geon chains).

Nine geons defined the LOW_FAST taxonomy. A unique combination of two geons (sampled from a set of three) defined each one of three top-level categories (see Figure 3). High-level strategies had length 2 because a two-geon conjunction had to be tested. A unique diagnostic geon further specified the bottom-level categories. Bottom-level categories had length 1 strategies because a single feature test determined membership. Figure 3 shows the LOW_FAST taxonomy. We added one filler to generate six four-geon chains. From these, we created two exemplars per category (see Figure 3).

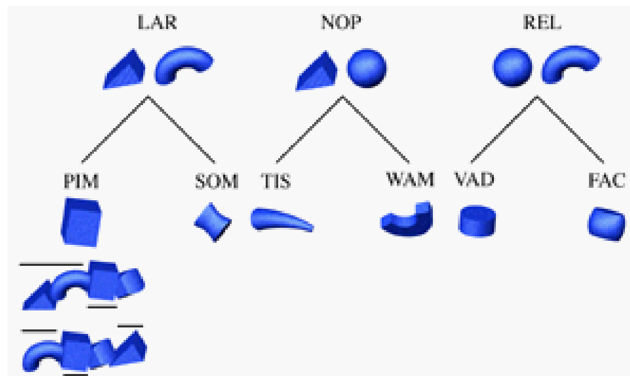

Figure 3. The LOW_FAST taxonomy of Experiment 1. The geons specify the defining information of each category. The bottom geon chains the two PIM exemplars (they are also LAR exemplars) used in the experiment.

Procedure

The procedure followed closely that of Murphy (1991). In a learning phase, participants were evenly split between the learning of the HIGH_FAST and LOW_FAST taxonomies. We instructed participants to learn the names and the defining geon(s) of nine categories. Participants saw their taxonomy on a sheet of paper (see Figures 2 and 3); this learning phase was not constrained in time.

We tested participants' knowledge of the taxonomy by asking them to list the features associated with each category name. Learning criterion was to list twice in a row, without any mistake, the defining features of each category. Corrective feedback was provided.

When subjects knew the taxonomy, a category verification task measured categorization time at each level. Each trial began with the presentation of a category name. Subjects could then press the "continue" computer keyboard button to see the list of all learned definitions on the screen (each definition corresponded to a set of geons per category). Participants had to identify the list associated with the previously shown category name. This insured that subjects accessed the representation of this category. After a 200 ms delay, an object appeared on the screen. Subjects had to decide–as fast as they possibly could–whether or not the named category and the object matched by pressing the "yes" or "no" computer keyboard key. We recorded response latencies. Note that low-level categories are more numerous than high-level categories. We normalized the number of positive and negative trials with the constraint of equating the number of trials per level.

Results and discussion

We performed the analysis of RTs on the correct positive trials that were within two standard deviations from the means. Table 1 reports the mean RTs at the low- and high-levels for the two taxonomies tested (see Observations in Table 1).

Table 1. Mean RTs for the Positive Trials of All Experiments As Well As Predictions of SLIP, Category Feature-Possession, and Category Utility (Erroneous Predictions Are Shaded).

		Level	
	Model	Low	High
Exp. 1, HIGH_FAST	Oservation	1,256	896
	Possession	2	3
	Utility	.195	.222
	SLIP	6.4	3.2
Exp. 1, LOW_FAST	Observation	948	1,240
	Possession	1	3
	Utility	.25	.333
	SLIP	3.2	6.4
Exp. 2, HIGH_FAST	Observation	788	660
	Possession	1	3
	Utility	.375	.500
	SLIP	3.2	2.286
Exp. 2, LOW_FAST	Observation	740	774
	Possession	3	1
	Utility	.624	.500
	SLIP	2.286	3.2
Exp. 3, EQUAL	Observation	672	680
	Possession	1	5
	Utility	.176	.260
	SLIP	1.714	1.714
Exp. 3, SL_DOWN	Observation	920	1,058
	Possession	1	5
	Utility	.250	.333
	SLIP	1.714	3.429
Exp. 3, IP_UP	Observation	928	775
	Possession	1	5
	Utility	.250	.333
	SLIP	6.857	3.429

A two-way (GROUP x STATEGY LENGTH) ANOVA of the RTs with repeated measures on one factor (STRATEGY LENGTH) revealed a main effect of STRATEGY LENGTH, $F(1, 18) = 77.08$, $p < .0001$, (mean length 1 strategies = 922 ms verification time; mean length 2 strategies = 1248 ms verification time), meaning that participants systematically verified length 1 strategies faster than length 2 strategies, irrespective of the considered level (low vs. high). Neither the interaction between GROUP and STRATEGY LENGTH, $F(1, 18) = .84$, ns, nor the main GROUP effect, $F(1, 18) = .02$, ns, were significant. The error rate was low overall and was not correlated with RT ($r = -.17$, ns), ruling out a speed-accuracy trade-off.

Remember that SLIP predicts that length 1 strategies should be completed faster than length 2 strategies, irrespective of categorization level (see SLIP in Table 1 for numerical predictions with $S = .25$). The data reported here confirms that strategy length, rather than categorization level, determines participants RTs.

Experiment 2

Practicability refers to the ease with which the features identify a category at any level of a taxonomy. A category has high practicability if many of its defining features are uniquely diagnostic of this category (or if the features occupy few positions across exemplars). It will have low practicability if only one feature defines the category (or if the features can occupy many positions across exemplars). Practicability has so far been the only factor under study in basic-level experiments (see Gosselin & Schyns, 1997). Never has it been shown, however, that practicability could affect the basic-levelness of all categorization levels.

Experiment 2 isolates practicability in a two-level taxonomy using objects similar to those of Experiment 1. All strategies had length 1 but the high and low levels differed in practicability. In the HIGH_FAST condition, high-level strategies had greater practicability than low-level strategies. The opposite applied to the LOW_FAST condition, with low-level strategies having higher practicability. SLIP predicts that categories with higher practicability will be verified faster, irrespective of their level in the taxonomy.

Method

Participants

Twenty students from University of Glasgow with normal or corrected vision were paid to participate in the experiment.

Stimuli

Stimuli were similar to those of Experiment 1: four-geon chains synthesized with a three-dimensional modeling software on a Macintosh computer.

The HIGH_FAST condition used 10 diagnostic geons. Three different geons defined each one of two high-level categories; one different geon further defined each low-level category (see Figure 4). We generated two exemplars per category by changing the location (either rightmost or leftmost of the chain) of the three geons

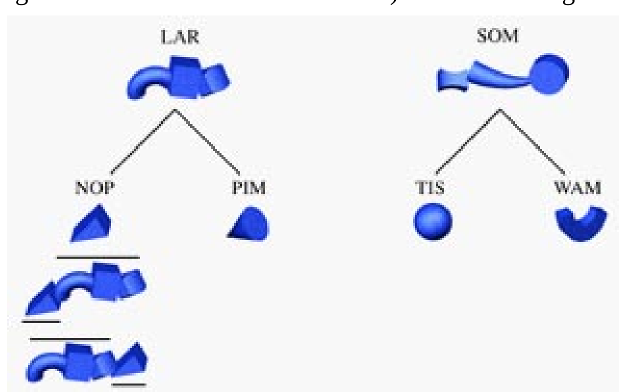

defining the high-level categories (see Figure 4).

Figure 4. The HIGH_FAST taxonomy of Experiment 2. The geons specify the defining information of each category. The bottom geon chains are the two NOP exemplars (they are also LAR exemplars) used in the experiment.

The LOW_FAST condition involved fourteen diagnostic geons. A single diagnostic geon defined each one of two high-level categories, and three different geons further defined each one of four low-level categories (see Figure 5). As before, we created two category exemplars by changing the location (either far right or far left of the object) of the triplets defining the low-level categories (see Figure 5).

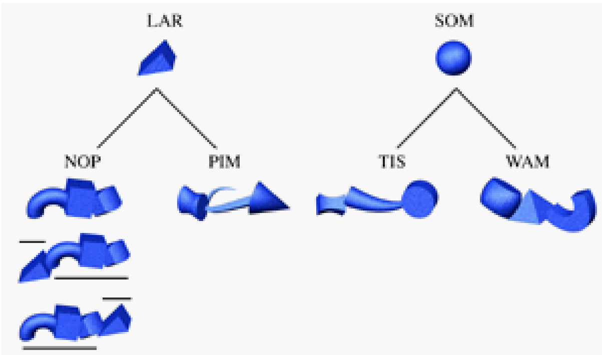

Figure 5. The LOW_FAST taxonomy of Experiment 2. The geons specify the defining information of each category. The bottom geon chains are the two NOP exemplars (they are also LAR exemplars) used in the experiment.

Practicability is greater for high-level categories in the HIGH_FAST condition and for the low-level categories in the LOW_FAST condition because more unique features are associated with the top- and bottom-level categories, respectively. SLIP predicts a faster verification performance for categories with higher practicability (high in HIGH_FAST and low in LOW_FAST) irrespective of the level of the taxonomy considered.

Procedure

The procedure followed in all respects that of Experiment 1: Participants were randomly assigned to the HIGH_FAST and LOW_FAST conditions. They were taught their respective taxonomy before being measured on the categorization speeds of its levels. Each one of 280 trials consisted in the initial presentation of a category name followed by an object. Participants had to decide as fast as they possibly could whether the two matched and we recorded response latencies.

Results and discussion

We analyzed only the correct positive trials RTs within two standard deviations from the means. Table 1 shows the mean RTs at the low and high-levels for the HIGH_FAST and for the LOW_FAST taxonomies.

A two-way (GROUP x PRACTICABILITY) ANOVA on the RTs with repeated measures on one factor (PRACTICABILITY) revealed a main effect of practicability, $F(1, 18) = 16.83$, $p = .001$ (mean verification time = 700 ms for high practicability strategies; 781 ms for low practicability strategies). Out of 20 participants, only three did not respond faster to the greater practicability categories. Neither the GROUP x PRACTICABILITY interaction, $F(1, 18) = 5.53$, ns, nor the main GROUP effect, $F(1, 18) = .06$, ns, were significant. The error rate was low overall and was not correlated with RT ($r = .05$, ns), ruling out a speed-accuracy trade-off.

In sum, SLIP predicted that greater practicability strategies should yield faster categorization decisions (see SLIP in Table 1 for numerical predictions with $S = .25$). The results of Experiment 2 reveal that this factor determined RTs at different categorization levels.

Experiment 3

Experiments 1 and 2 respectively revealed that strategy length and internal practicability–the two computational determinants of SLIP–can contribute independently to faster categorizations at any level of a taxonomy. Experiment 3 further explores how these two factors interact to determine performance.

There are many possible interactions to investigate and we will not investigate them all. Instead, we have selected to examine three scenarios that selectively change the fastest categorization level by selectively modifying either strategy length or internal practicability.

In the EQUAL scenario, strategies at the high and low-levels have an equal length of 1 and a constant practicability. SLIP predicts that in these circumstances, categorization speeds should be equal across levels. EQUAL is our baseline condition. In the SL_DOWN scenario, practicability is constant across levels, but whereas low-level strategies have length 1, high-level strategies have length 2. SLIP predicts faster categorizations at the lower level. The IP_UP scenario preserves the difference in strategy lengths, but it changes the fastest categorizations to the higher level by decreasing the practicability of the low level.

Together, EQUAL, SL_DOWN and IP_UP illustrate how the faster categorization level can go up and down a taxonomy by changing strategy length or the internal practicability, the two factors of SLIP.

Method

Participants

Thirty students from University of Glasgow with normal or corrected vision were paid to participate in the experiment.

Stimuli

Stimuli were similar to those of Experiments 1 and 2: geon chains designed with a 3D-object modeling software.

Nine diagnostic geons entered the composition of categories in the EQUAL, SL_DOWN and IP_UP conditions. In EQUAL, one geon defined each one of the nine categories of the taxonomy (see Figure 6). We added four fillers to each defining geon to form a total of six six-geon chains. We placed the geons defining the high-level categories at the far left of the chains, and those defining the low-level categories at the far right (see Figure 6).

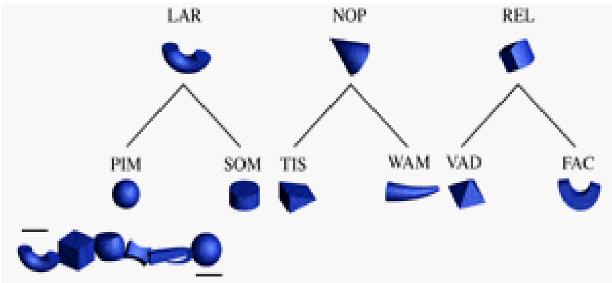

Figure 6. This illustrates the EQUAL taxonomy of Experiment 3. The geons specify the defining information of each category. The bottom geon chain is the PIM exemplar (it is also a LAR exemplar) used in the experiment.

In SL_DOWN, a unique combination of two of the nine geons defined each top-level category. The addition of one different geon further defined each lower-level category. (SL_DOWN employed the Experiment 1, LOW_FAST taxonomy but with a different set of geons). We produced six six-geon chains by adding three fillers. We placed the geon pairs defining the high-level categories at the far left of the chains, and those defining the low-level categories at the far right.

These chains also served to construct the exemplars of condition IP_UP. Here, we generated four exemplars per category by changing only the location in the chain of the single geon defining the low-level categories (one of the four rightmost positions in the six-geon chains).

Procedure

The procedure was almost identical to that of experiments 1 and 2. Participants were randomly assigned to one of three conditions (EQUAL, SL_DOWN, and IP_UP). Following a learning of their taxonomy, they did 240 verification trials. Each trial consisted in the presentation of a category name followed by an object. Participants had to decide whether these matched and we measured response latencies.

Results and discussion

We performed the analysis of RTs on the positive, correct trials that were within two standard deviations from the means. Table 1 shows the mean RTs.

A two-way (GROUP x LEVEL) ANOVA with repeated measures on LEVEL revealed a significant interaction between GROUP and LEVEL, $F(2, 27) = 11.85$, $p < .001$, simple main effects of GROUP(SL_DOWN) by LEVEL, $F(1, 27) = 10.58$, $p = .003$, GROUP(IP_UP) by LEVEL, $F(1, 27) = 13.09$, $p = .001$, and GROUP(EQUAL) by LEVEL, $F(1, 27) = .04$, ns. The error rate was low overall and was positively correlated with RT ($r = .31$, $p < .05$), ruling out a speed-accuracy trade-off.

SLIP predicted all the results observed in Experiment 3 (see SLIP in Table 1 for numerical predictions with $S = .25$). Participants categorized equally fast at both levels in EQUAL. Increasing the strategy length of the higher level in SL_DOWN induced faster categorizations of the lower level. Diminishing practicability at the lower level then made the high level faster. Thus, the two computational factors of SLIP predicted speed of categorization in taxonomies.

General Discussion

SLIP (Strategy Length & Internal Practicability) is a new model of basic-level performance. Three verification experiments tested the two computational determinants of the model: strategy length and internal practicability. In Experiment 1, strategy length was shown to decide basic-levelness. In Experiment 2, practicability was shown to be a second determinant of basic-level performance. In Experiment 3, interactions between strategy length and internal practicability in SLIP predicted the observed RTs.

SLIP performance can be compared to that of two well-established measures of basic-levelness, *category feature-possession* (Jones, 1983) and *category utility* (Corter & Gluck, 1990). The predictions of the models are given in Table 1. The scores of both category utility and category feature-possession should be inversely proportional to RTs; SLIP's scores should be directly proportional to RTs. The best predictor is SLIP with seven correct RT patterns out of seven, followed by category feature-possession with a hit rate of 5/7 (the mistakes have been shaded in Table 1), and trailed by category utility with a 4/7 hit rate.

It is instructive to decompose these scores into strategy length and internal practicability scores. For the conditions testing only practicability (Experiment 2, HIGH_FAST and LOW_FAST, and Experiment 3, EQUAL and IP_UP), category feature-possession and category utility both predict 3/4 of all RT patterns. We have demonstrated elsewhere (Gosselin & Schyns, 1999) that these models are biased to faster responses at higher levels.

Interestingly, for the conditions testing only strategy length (Experiment 1, HIGH_FAST and LOW_FAST, and Experiment 3, EQUAL and SL_DOWN) category feature-possession and category utility only predict 4/8 and 2/8 of the RTs, respectively. (Note that Experiment 3, EQUAL, is included in the break-down into strategy length and internal practicability; it is an extreme case of both.) This confirms the argument that category feature-possession and category utility neglect strategy length as a specific factor of basic level performance (Gosselin & Schyns, 1997). This is a serious problem because attributes do overlap between categories in the real-world, and so strategy length is an important factor of categorization performance outside the laboratory.

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.

Acknowledgements

The first author 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 was funded by ESRC grant R000237901 to the second author.

Reference

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

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. (1999). *Why do we SLIP to the basic-level? Computational constraints and their implementation.* Manuscript submitted for publication.

Hoffmann, J. and Ziessler, C. (1983). Objectidentifikation in kunstlichen begriffshierarchien. *Zeitschrift fur psychologie*, *194*, 135-167.

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

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

Lassaline, M. E. (1990). *The basic level in hierarchical classification.* Unpublished master's thesis, University of Illinois.

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

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

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

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.

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.

Appendix

A category is defined by 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 $\left(1-\psi_j\right)^{t-1}$ where ψ_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 $C_j(1-S) + C_jSR_j$ 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 ψ_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 ψ_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 simply 1 minus the probability of a miss. Thus, the probability that the set of redundant features *j* will be completed after *t* trials is

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

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

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

where ϕ is a function of *j* (it will remain unspecified) which gives the number of misses for the *j*th 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,

$$\lambda = \binom{t-1}{n-1} = \frac{(t-1)!}{(t-n)!(n-1)!}.$$

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

$$\sum_{i=1}^{\lambda}\prod_{j=1}^{n}\left(1-\psi_j\right)^{\omega}\psi_j,$$

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