Michael Ramscar

Research Statement

My research seeks to answer two questions: how do we as individuals learn to make meaningful sense of the world, and how do we share our understanding through language?

In what follows, I describe the theoretical framework my collaborators and I have developed to answer these core questions, and elaborate the conceptual principles and psychological mechanisms underlying this view. To illustrate the benefits of this approach, I detail the research I have undertaken and subsequently elaborate the many insights (and discoveries) it has produced. In particular, I describe how we have applied this framework to some well-known and longstanding puzzles, including: children's and adults' symbolic category learning, children's acquisition of number sense, and the distribution of perfect pitch in the general population. I conclude by reviewing other applications of this approach, and the particular conception of cognition it embodies.

Learning

My approach to language learning is rooted in the formal theories devised to explain the basic processes of animal learning. These mathematical theories specify the computational processes involved in learning [1], and have led to a significant growth in our understanding of the neurobiological mechanisms that allow humans and animals to understand their worlds [2, 3]. However, historically the application of learning theory to higher cognitive processes has been marked by misconceptions and misunderstandings [4]. A major strand of my research seeks to improve our understanding of the relationship between learning and higher-level cognition, and I have sought to use this work to make learning theory more comprehensible to cognitive scientists [5].

An aspect of learning theory that has been frequently misunderstood is the role that informativity plays in learning. Although learning models are usually referred to as "associative"—which might seem to suggest that learning works by simply noting where cues and events co-occur in the environment—this is not the case. Indeed, for over 40 years the process of "associative learning" has been formally conceived of in terms of **prediction** and **discrimination**: learning processes discriminate between predictive cues on the basis of the information they provide. What this means in practice is that learning is driven far more by predicted events that *don't occur*, than it is by positive pairings of cues and events.

This point, in particular, is not well understood in psychology. However, it can be clearly and simply illustrated via conditioning experiments on rats. If a rat is exposed to a series of tones followed by mild shocks, it will condition to the relationship between shocks and tones, responding fearfully to future tones. However, if tones that <u>don't</u> lead to expected shocks are added to the tone-shock pairings, the rat's conditioned responses will weaken in proportion to the increased 'background rate' of tones. Experiments such as these have revealed that the strength of the rat's conditioned response depends on how *informative* tones are about the shocks [6].

We have examined the way children use informativity in learning in a variety of ways. In one simple experiment, we trained children on novel word meanings while manipulating the 'background rates' of the objects paired with the labels that they learned. Children first saw two different novel objects 'A' and 'B' together, and heard them labeled ambiguously as a "DAX." Subsequently, 'B' was presented with a new object, 'C,' and another ambiguous label – "PID." This training was repeated, and then the children were presented with all three objects, and asked to identify either the "DAX," the "PID," or the "*WUG*" (which they hadn't heard before). Because 'B' occurs with both "PID" and "DAX," it has a higher background rate than 'A,' which makes 'A' more informative about "DAX" than 'B'. For the same reason, 'C's informativity about "PID" is greater than 'B's. So far, so obvious, perhaps, but importantly, B's higher background rate **also** makes it less informative about "WUG" than either 'A' or 'C' (which are both equally informative about "WUG").

From a purely informational perspective, 'A' is a DAX, 'C' is a PID, and 'A' or 'C' are WUGs. The two year olds we tested roundly agreed. However, while children match objects and labels based on informativity, undergraduates we tested did not. They agreed with the children about 'A' and 'C' but chose 'B' as the WUG. Further, when Developmental Psychologists were surveyed about their predictions for children's behavior in our task, they too thought 'B' was the WUG. Adults appear to reason—logically—that if B is not a DAX or a PID, it must be a WUG, whereas children appear to care more about informativity than logic [7]. (Given that the same object might be "Fido," "a dog," "a dumb mutt" or "pooch" depending upon the context, this strategy is not necessarily unwise).

This suggests two things about how children learn words: first, that they go about it in a way that is qualitatively different from that of adults; and second, and more intriguingly, that error-driven learning mechanisms can offer much insight into the nature of children's statistical learning. Given that it is widely believed that such relatively simple mechanisms cannot possibly account for any aspect of verbal learning, this is a powerful finding.

To begin to cash out these insights, much of my recent research has focused on the question of how we can optimally structure information in a child's learning environment to speed or improve language learning (see e.g., my work on color and number learning, and overregularization, below). First, however, I should delve into the theoretical work that has made these practical applications possible.

How do symbols serve as abstractions?

Several years ago, I conducted a formal analysis of the way *linguistic symbols* are learned, to better understand the role they play in human thought.

This analysis begins with a straightforward observation from learning theory: that when learning a complex relationship in the environment, it is far better to learn from a **rich** set of cues to a **sparse** set of outcomes, than the reverse. So, for example, in the rat experiment, compare a rat trying to learn that a variety of different tones (cues) predict a shock (one outcome) versus a rat trying to learn that the same tone (one cue) predicts either a shock, a massage, or a food pellet (outcomes). The rat in the first scenario will quickly realize that any given tone it hears should prompt a particular reaction (e.g., cowering), while the rat in the second scenario will have no clear idea how to react, because the single cue-type will not allow for discrimination between possible outcomes (which could be wonderful or terrifying, depending).

This point about the role of cues in **coding** what gets learned is critical to the question of how we learn what *words* mean, because while objects and events in the environment are 'rich' in features, verbal labels are not. To get a feel for the contrast, compare the many varied features of a flesh-and-blood dog (color, size, hair, snout, bark, etc) with the acoustic features of '*dog*.' Outside of a relatively small set of phonemic category distinctions, people aren't good at discriminating within the sounds that language builds on. Thus in comparison to the world we perceive, which has innumerable features that discrimination learning reshapes into informative codes for an enormous variety of events, the sounds we use to communicate are almost featureless, providing a very barren medium for coding.

Because of this, word learning is a necessarily asymmetrical process: depending on the temporal ordering of learning (i.e., whether objects serve as cues and words as outcomes, or the reverse) very *different* patterns of learning and very different learned representations will result [5]. When **features predict labels (FL)**, learning is straightforward and informative. Because a number of cues are present at any time, cue competition occurs in learning. In this competition, more informative cues win out over less informative cues, which generate error, resulting in informative discrimination learning. However, when **labels predict features (LF)**, competition is essentially absent (a phoneme can't compete with itself), and this inhibits discrimination learning, markedly changing the nature of what gets learned. Computational simulations confirm this analysis, showing that learning produces different representations depending on the order in which features and labels are encountered, and that better discrimination learning occurs when objects predict labels [5].

We have shown that these informational constraints affect human symbolic learning. For example, if participants in category learning experiments are trained to predict labels from exemplars, they discriminate and categorize well, even when highly confusable categories are trained at markedly different frequencies. However, if participants are trained to predict exemplars from labels, they fail to learn to discriminate these confusable exemplars at all [5]. This is despite the fact that each participant sees exactly the same labels and exemplars for exactly the same amount of time in training; all that varies between conditions is the order in which labels and exemplars are presented.

These Feature-Label-Order (FLO) effects have been replicated in numerous studies that bridge both visual and auditory modalities. Additionally, an fMRI investigation into FLO-effects in collaboration with Sam McClure revealed greater **dorsal striatal activity** in FL-trained participants – which correlated with overall performance – and greater **PFC recruitment** in LF-trained participants – which correlated with poorer performance on low frequency items. These findings are highly compatible with what is known about the function of these areas: a wealth of evidence links the striatum to discrimination learning, while the PFC is known to be responsible for effortful, strategic thinking. The increased striatal activity we observed in the FL-trained participants is consistent with their superior discrimination learning, while the correlation between PFC activity and poorer discrimination in the LF-trained participants supports the idea that the unhelpful structure of information in training caused these participants to "think too much" about the task; it seems they consciously fixated on cues that were frequent and salient, but ultimately uninformative.

Color label learning

This analysis of symbolic learning offers a solution to a puzzle first noted by Darwin: why do children struggle to learn color words?

Our answer to this question begins with the ubiquity of color. In order to dissociate colors from inappropriate labels, children need <u>informative</u> error. However, when it comes to color, children get—literally—too much information. At any given moment, children are exposed to most colors, making it extremely hard for them to discern any systematic covariance between labels and hues: when "red" is heard, red, yellow, brown, green, blue etc., will be almost certainly be present, and the same will be true for "yellow" or "green" or "blue," etc. The error from children's expectations in these circumstances in uninformative.

To get around this problem, children's attention needs to be directed to particular aspects of visual scenes as labels are heard, so that the process of discerning the informative cues to labels can be simplified (a ball can be simply red or blue, even if the world never is). Language can provide children with this direction, as long as they already know the names of other things, such as objects, and as long as language is structured appropriately. To explore these ideas, we conducted a training experiment with two-year-olds that manipulated the order in which color words were used in English. What we found was that children who heard object names predict color labels (hearing "this ball is blue" when shown a blue ball, **FL**) consistently sorted novel objects above chance levels, whereas those who heard color labels predict object names ("this is a blue ball," **LF**) did not [5].

In recent studies I have applied this technique for analyzing information structure to other developmental questions, such as the way children learn numbers. A formal analysis of the problem of number set identification showed that the limits on set size identification (our ability to "see" a set of, say, three objects; usually called "subitization") arise naturally in a discrimination-learning model. This is because of two factors: 1) the increased confusability of larger sets (a set of *two* objects also contains sets of one object, and a set of *three* objects also contains a set of two, and sets of one object, etc), and 2) the way sets are distributed in the learning environment (set salience appears to follow a power law distribution, with one being the most frequent and salient set). These factors make it easy to learn to discriminate smaller sets, but not larger ones.

As well as naturally explaining how subitization "capacity limits" arise out of experience, the model makes a novel prediction: Because our training deliberately changes the background rates of cues that aren't informative to number words, training children **feature to label (FL)** on the sets 2, 4 and 6 should also improve their ability to discriminate 3, 5 and 7! This striking prediction was confirmed in tests [8]. Given that the age at which children master the subitization of small sets is an important predictor of their later mathematical abilities, I am currently working on developing educational applications of this successful training paradigm.

"No representation without taxation"

The work described above shows how symbolic learning is, quite literally, a process of abstraction. Learning to discriminate between categories increases the value of informative cues at the expense of cues that are uninformative. This process distorts the completeness of what is learned by dissociating non-discriminating features, in much the same way that one might eliminate a great deal of detail to convey the gist of a paper in an abstract. In studies, I have explored the consequences of this process on the kind of representations participants form in learning. I found that when participants were required to categorize objects (discriminating at the category level of abstraction), they performed better when **FL**-trained. However, when asked to re-identify the exact exemplars they were trained on (discriminating at the exemplar level of abstraction), they performed better when **LF**-trained [9]. This is because while FL-trained participants learned to ignore features that are uninformative (as far as classification goes), these same features would have helped identify test exemplars on the second task. In other words, the principles that benefited the classification task were detrimental to the identification task. I have suggested that the mechanisms of human learning obey a basic principle: *no representation without taxation*, and from a learning perspective, both costs and benefits arise from learning to conceptualize the environment.

I have subsequently applied this approach to the question of "perfect" or "absolute pitch" – the ability to name notes based on absolute frequency information. It has long been known that 1) perfect pitch is typically quite rare among the general population and 2) that possessors perform less well than non-possessors on tests of *relative* pitch. The question is why? To begin to answer this, we used an FL-LF design to train participants to either acquire better within-category representations (LF) – which we predicted would be better for absolute pitch learning – or better between-category representations (FL) – which we predicted would be better for relative pitch learning. The results of our experiment bore this out: While LF trained participants were better at perceptually discriminating trained notes from near-identical lures (showing that they had a less distorted representation of what they had heard), they performed proportionally less well at discriminating octave transpositions of the trained notes from lures (showing they had learned the pitch categories less well). A similar pattern of results was obtained in tasks matching tones to labels. Crucially, FL trained participants showed a loss of sensitivity to within-category discriminations as a result of learning pitch category representations. In other words, their representations of pitch had become more *symbolic.* [10]

This is an important finding, since it essentially paints a probabilistic picture of symbols themselves: symbols can be seen as being more or less "symbolic" (and abstract) depending on the degree to which their underlying representations have been shaped by learning. This work provides a tractable method for specifying what "symbols" are, and for studying how they develop. In this vein, I am currently extending this work to an investigation of phonological category learning, and the development of linguistic symbols.

Language as an informative, predictive process

My work on the formal properties of symbolic learning has led me to develop a theory of communication based on a process of mutual prediction. On this view, when talking, speakers use their cultural and experiential knowledge to generate the utterances they believe are most likely to bring about changes in listeners' beliefs or behavior. At the same time, listeners, far from being passive decoders of tokens of meaning, use broadly the same processes to build up their understanding of what is being said. Listeners use both learned semantic cues and distributional information in the speech stream, in order to predict the behavior and intentions of speakers. Successful communication thus relies on shared prior knowledge that enables mutual predictability, and collaboration between speaker and listener to bring about understanding [5].

This view of language differs markedly from traditional approaches. Most theories of language are based on the idea that there is a "sender" and a "receiver" of tokens of meaning. A speaker sends a listener a message in the form of words, which the listener then decodes. This natural, folk-psychological view of language suffers from numerous problems, which are usually side-stepped in contemporary linguistic and psychological theories. For example, language is usually seen as referential, such that the relationship between words and their meanings (or the things in the world they represent) is treated as bidirectional, even though philosophers such as Wittgenstein and Quine have noted numerous deep problems in this. I have sought to formalize these analyses by considering reference from the point of view of formal theories of learning, and information and coding [1].

My work on symbolic learning has shown how the idea of theory of reverse abstraction that lies at the heart of referential theories is incompatible with both learning and information theory, and the behavior of adults and children in experiments. Why? In short: the idea that a word can 'convey' a meaning makes about as much sense as the idea that an abstract can 'covey' detailed information about the results and method sections of a paper it summarizes. It would be better to say that, given an abstract, we can make guesses about the results and methods sections – thereby making a kind of prediction about the kind of information they might contain. If we are expert on the topic at hand, then the likelihood that our predictions are more accurate, or even substantially correct, will increase. However, given no more than an abstract, we can do no more than make predictions. This is because the process of abstraction involves *discarding* information that cannot be recovered from an abstract representation. Since symbols are abstractions, it follows similarly that symbolic meanings can only be inferred. This is the basis for the idea that language must be a predictive process.

Predicting sequences when understanding

Predictive approaches to language have long been dismissed as implausible because the sequential probabilities speakers must acquire are argued to be unlearnable. These arguments rest on two claims that I believe are extremely weak: 1) that there is no generalization in prediction, and 2) that predictions are only made over words.

Regarding the first claim, Daniel Yarlett and I have developed a detailed cognitive model of cue generalization in predicting word sequences [11; 12], in which predictions about unobserved sequences get smoothed by blending in probability estimates from similar sequences that have been experienced. For instance, knowledge about *p(mat|cat)* is used in estimating the probability of *p(mat|dog)* if *mat* has never been heard after *dog*. In simulations trained on data that children might reasonably be expected to encounter, our model matches or exceeds the performance of state of the art language engineering models [12]. The model is able to successfully combat data-sparsity because it generalizes across neighborhoods of distributionally similar words.

To test the psychological plausibility of this idea, I conducted a training study with adults with Melody Dye. We found that inserting a low frequency word like *samovar* into sentences with words that were distributionally similar to *kettle* resulted in measurable changes in both priming between *kettle* and *samovar*, and in semantic similarity judgments between them. Since participants were never actually exposed to *kettle* in training, these results support the idea that linguistic experience does indeed generalize across distributional neighborhoods [12]. This is an important finding because the claim that people could never learn the set of "parameters" needed to drive a predictive model of language *hinges* on the assumption that the generalizations that our participants made don't happen!

In conjunction with this line of work, Colin Bannard and I have examined the *kinds* of representations that people learn predictive parameters for. In a series of experiments, we have found that learning of probabilistic information about language is not restricted to words and phonemes. Both children and adults demonstrate sensitivity to manipulations of overall *sequence* frequency, in which word and bigram frequencies are held constant, revealing that they acquire detailed knowledge of the probabilities of multi-word sequences [13, see also 14]. Recently, we have extended this kind of modeling to other questions, showing how simple predictive models can be used to explain the effects of metaphoric priming [15], and how a simple distributional model can not only provide a single mechanistic account of the many claimed differences in people's perception of regular and irregular plurals in compounds, but can also be used to identify systematic differences among classes of nouns that are usually grouped together as "regular"[16].

Taken together, these findings provide strong support for a *usage-based* approach to language. From this perspective, "linguistic rules" are descriptions of regularities that are observed in the distributions of languages—they are theoretical constructs, rather than psychologically real structures that *guide* linguistic behavior. Linguistic patterns emerge out of the behavior of a community of speakers, and the role of a theory of language is to explain how speakers learn from the distribution of language they are exposed to, how they use this knowledge of language to communicate, and how the demands of communication shape linguistic distributions. The approach I have adopted in my work, which treats language as a predictive process, and communication as the reduction of uncertainty between interlocutors, fits naturally with both learning and information theory, and I believe that it is uniquely well equipped to begin to meet these explanatory demands.

Predicting sequences when reading

The approaches pioneered in my lab have begun to be adopted by other researchers. In a particularly exciting development, Harald Baayen and his collaborators have adapted our predictive approach to develop a simple discriminative model of reading time responses [17]. Because the model treats reading as a process of predicting symbolic meanings (which in turn predict symbolic sounds), it does not – as might be expected – assume separate representations for inflections or inflectional paradigms. Yet despite its simplicity, it can still successfully simulate the complex paradigmatic effects that characterize Serbian case inflection. Just as remarkably, in simulations of English data, frequency effects for complex words and phrases emerge in the model *without* the presence of whole-word or whole-phrase representations; meanwhile, family size effects that emerge across simple words, derived words, and compounds can be simulated without derived words or compounds being represented as such. In addition, the model replicates the finding that, on average, words with more productive affixes elicit longer response latencies, and predicts that productive affixes afford faster response latencies for new words. (The model has also been successfully applied to lexical decision data in Dutch [18]).

In short, this reading model succeeds in accounting for more data more accurately than any other model I know of, while making fewer representational and processing assumptions. The remarkable coverage of the model, given its equally remarkable simplicity, and its amenability to formal specification, offers a powerful demonstration of the broad potential of the approach described above.

Sequencing and learning in language

It is clear that children learning a language have to learn more than just how sounds map to meanings. They have to learn to identify patterns of regularity in the sounds themselves. In our models, learning is treated as a process of iterative discrimination, in which the cues to individual linguistic regularities are slowly refined, along with those regularities themselves. Accordingly, since words will generally be encountered in linguistic contexts, this context, as well as words themselves, serve as cues to other words. To explore the dynamics of this kind of contextual learning, we have used our predictive learning model as the basis for investigating adults' poor performance in learning grammatical gender during second language acquisition.

There are clear differences in the way that adults and children approach the task of language learning: For example, unlike children, adults already have a language, as well as folk theories about how language works, and a greater proportion of their initial learning may take place in a classroom, as opposed to in context. If adults learn the meanings of words out of context, our model suggests that this should have the effect of blocking later learning of associations between linguistic contexts and words. To test this idea, Inbal Arnon and I trained participants on an artificial language, utilizing a classic blocking design: half of our participants were trained on blocks of noun-meaning mappings prior to encountering the nouns in sentences with their determiners, whereas for the other half this order of training was reversed. Consistent with our blocking hypothesis, we found that the participants who first encountered the nouns in sentences learned to map them to their gendered determiners far better than those who first trained on the nouns alone. Moreover, the sentence-first group also learned the noun-meaning mappings better, suggesting that grammatical gender markers may serve to carry information by narrowing the scope of subsequent lexical predictions [19].

In related work, my collaborators and I have found that a Feature-Label-Order (FLO) asymmetry may explain why many languages are biased to add inflections to the *end* of words. A corpus analysis of English confirmed that suffixes are more informative than prefixes about the grammatical category of root-words. Additionally, an artificial language learning task revealed that suffixes (which are predicted by root-words, i.e., FL-learning) were learned significantly more accurately than prefixes (which predict root-words, i.e., LF-learning) [17].

Context, discrimination, and a surprising prediction about over-regularized plurals

The predictive framework described here has has lead to a very fruitful reanalysis of the reasons behind children's tendency to over-regularize plurals (i.e. *mouses*), a much debated question in recent debates about the nature of language. In our predictive model, over-regularization arises initially because the semantic cue of 'plurality in general' predicts both regular and irregular forms [21, 22]. (For example, if a child has not discriminated between 'rats' as a label to multiple rats, and rats as a label to multiple things, such as 'mice'). This causes interference and over-regularization, because the high frequency of regular plurals in English causes their representations to strengthen more rapidly than those of irregular forms. In the model, this interference resolves itself naturally in time, as a result of error-driven learning. Until 'plurality in general' has been unlearned as a cue to pluralization, the occurrence of regular plurals leads to irregulars being expected, resulting in negative learning when regular plurals are encountered instead. In time, this negative learning dissociates irregulars from the more general plural features that lead to interference and over-regularization.

The model makes very clear predictions: that children can stop over-regularizing simply by being exposed to a normal distribution of plurals, and that they will do so without the need for explicit feedback [22]. Moreover, it predicts a point in learning where exposing children to regular plurals **alone** will *reduce* over-regularization. Studies have provided strong support for these predictions [22, 23], further demonstrating children's sensitivity to information in learning.

The architecture of learning as an adaptation for culture

A key question that all theories of language must answer is this: what are the origins of linguistic conventions? Symbolic communication is, in its essence, conventional. Given a symbol, a social animal needs to be able to infer and understand (and often, to do) the *appropriate thing* in the *appropriate context*. For this to happen, "symbolic values," must be conventionalized and internalized. Sharon Thompson-Schill and I have proposed that part of the answer to this question lies in the slow pattern of development of control processes in prefrontal cortex (PFC) in childhood [24, 25]. PFC functions enable adults (and animals) to filter their behavior and attention, and thus direct their learning. A great deal of evidence – behavioral, computational and neurobiological – supports the idea that children barely select between, or filter competing responses at all prior to their fourth year, and that adult levels of PFC function develop very slowly, over the course of childhood. We suggest that since PFC function provides a mechanism for active filtering, it will result in more "individualized" patterns of learning, whereas the absence of PFC function (and filtering) will result in learning that is more environmentally determined and conventionalized. In learning the appropriate cues to symbols, unsupervised cue competition will tend to produce very similar patterns of learning if learners are exposed to similar distributions of environmental cues and symbols. From this perspective, the delayed pattern of PFC development can be seen as an adaptation for cultural and linguistic convention learning.

Context and control in learning

Laboratory studies routinely reveal that children under the age of four are inflexible thinkers. Yet this poses a puzzle: why in the normal course of events is this inflexibility so hard to detect? Why do young children appear to be perfectly capable of switching responses and matching their behavior to context once they are outside the lab? Recently, we have explored the idea that **contextual learning** may be the answer to this: while toddlers are (in general) quite good at matching their behavior to context once they have learned to do so, they fail laboratory tasks because the novelty of the tasks requires active response selection, an ability which they have not yet developed.

To test this idea we used an FL / LF training manipulation to teach children color- and shape-naming games prior to having them participate in the Dimensional Change Card Sort Task (DCCS). In the DCCS, children play a game in which they sort items by color, and then in a new game, switch to sorting by shape. Notably, three year olds typically fail the switch trials. However, a simulation predicted that appropriate (FL) training on *naming* would facilitate contextual learning, enabling three-year olds to successfully pass the DCCS. In line with this prediction, we found that over 70% of children given FL-training in naming the cards' dimensions later successfully completed the DCCS switch trials. By contrast, less than 40% of LF-trained children, and less than 20% of an untrained control group managed to switch their sorting behavior appropriately [26].

As well as providing an answer to the puzzle above, this line of work allows for further exploration of the role of context in cognitive processing (an important part of the language as prediction hypothesis). The results also provide for a more nuanced view of the cognitive demands of "cognitive control" tasks, suggesting that the degree to which a task invokes active cognitive filtering is very much dependent on experience.

Prediction and language in context

The view of language described here is very different to many other contemporary theories, which adhere to a view of language built around 'senders' and 'receivers' of messages, which encode meanings. Our treatment of linguistic exchange as a process of mutual prediction is consistent with the idea that people learn how words are *used* in communication, rather than learning discrete word meanings. Theoretically, this approach is consistent with the ideas of major linguistic philosophers such as Quine and Wittgenstein. It is also consonant with a wealth of findings in psycholinguistics that have revealed the ubiquity of prediction in language processing. From a methodological perspective, this approach offers traction when it comes to understanding the learning, processing and nature of language. An equally attractive aspect of this approach is practical: it has enabled us to build models that successfully predict hitherto unforeseen phenomena in language learning, and to formally design training experiments that have succeeded in teaching color and number word discrimination at a much earlier age than had been thought possible.

A final practical benefit of this work lies in its potential to change our everyday understanding of communication and communicative behavior. The idea that linguistic exchange is a collaborative, predictive process grounded in learning effectively broadens the scope of what counts as "linguistic behavior." It paints a picture in which things like gesture, tone, and the many other devices we use to reduce uncertainty in communication are not brushed aside as 'pragmatics,' but are instead understood as integral to any communicative effort. In this, it makes sense of the *active* role a listener has in conversation, emphasizing the importance of the nods and grunts and expressions that a listener uses to provide feedback on her ongoing efforts to predict a speaker. Equally—and perhaps most importantly—it emphasizes the need for the speaker to attend to this feedback, so that he can shape a message that actually informs his listener. I believe that the ideas summarized here will not only improve our scientific understanding of communication, but will improve the way we communicate with one another as well. To this end, I am currently at work on a book that I hope will succeed in expressing these ideas to a much wider audience.

Publications and Manuscripts Cited

[1] Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effective- ness of reinforcement and non-reinforcement. In Black & Prokasy (Eds.), *Classical conditioning. II. Current research and theory* (pp. 64 99). New York: Appleton-Century-Crofts.

[2] Waelti, P., Dickinson, A. and Schultz, W. (2001): Dopamine responses comply with basic assumptions of formal learning theory. *Nature* 412: 43-48

[3] P. N. Tobler, C. D. Fiorillo, and W. Schultz (2005) Adaptive Coding of Reward Value by Dopamine Neurons. *Science* 307, 1642-1645

[4] Rescorla, R.A. (1988) Pavlovian conditioning: It's not what you think it is. *American Psychologist*, 43, 151-160.

[5] Ramscar, M., Yarlett, D., Dye, M. Denny, K.G., & Thorpe. K.L. (2010). The Effects of Feature-Label-Order and their implications for symbolic learning. *Cognitive Science*. 34(6), 909-957.

Ramscar, M. (2010) Computing Machinery and Understanding. *Cognitive Science*. 34(6), 966-971.

[6] Rescorla, R. A. (1968). Probability of shock in the presence and absence of CS in fear conditioning. *Journal of Comparative and Physiological Psychology*, 66, 1-5

[7] Ramscar, M., Dye, M., & Klein, J. Children Value Informativity Over Logic in Word Learning. *In submission at Science*.

[8] Ramscar, M., Dye, M., Popick, H. & O'Donnell-McCarthy, F. (*Under review)* How children learn to value numbers: Information structure and the acquisition of numerical understanding. Second round review *Science*.

[9] Ramscar, M. & Dye, M. (2009b). No representation without taxation: The costs and benefits of learning to conceptualize the environment. *Proceedings of the Analogy 09 Conference*, Sofia, Bulgaria.

[10] Ramscar, M., Suh, E. & Dye, M. (*Under review)* For the price of a song: How pitch category learning comes at a cost to absolute frequency representations. Second round review *Science*.

[11] Yarlett, D (2008) Language Learning Through Similarity-Based Generalization, PhD Thesis, Stanford University.

[12] Yarlett, D., Ramscar, M. & Dye, M. (*Under review*) Language Learning Through Similarity-Based Generalization. Second round review *Cognitive Science*.

[13] Bannard, C. (2006) Acquiring Phrasal Lexicons from Corpora. PhD Thesis, University of Edinburgh.

[14] Bannard, C. & Matthews, D. E. (2008) Stored Word Sequences in Language Learning: The Effect of Familarity on Children's Repetition of Four-Word Combinations, *Psychological Science*, 19, 241-248.

[15] Ramscar, M.J.A., Matlock, T., & Dye, M.W. (2010) Running down the clock: the role of expectation in our understanding of time and motion. *Language and Cognitive Processes*, 25(5), 589-615.

[16] Ramscar, M.J.A., & Dye, M.W. (*in press*) Learning Language from the Input: Why level ordering can't explain noun-compounding. *Cognitive Psychology*.

[17] Baayen, R. H., Milin, P., Filipovic Durdjevic, D., Hendrix, P., & Marelli, M. (*Under review*). An amorphous model for morphological processing in visual comprehension based on naive discriminative learning. *Psychological Review*.

[18] Baayen, R. H. (*in press*). Assessing the processing consequences of segment reduction in Dutch with naive discriminative learning. *Lingue & Linguaggio*.

[19] Arnon, I., & Ramscar, M. (2009) Granularity and the acquisition of grammatical gender: How order-of-acquisition affects what gets learned. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Amsterdam, Netherlands. (Longer article in second round review at *Cognition*.)

[20] St. Clair, M., Monahan, P., & Ramscar, M. (2009) Relationships between Language Structure and Language Learning: The Suffixing Preference and Grammatical Categorization. *Cognitive Science*. 33(7), 1317-1329.

[21] Ramscar, M.(2002) The role of meaning in inflection: Why the past tense does not require a rule. *Cognitive Psychology*, 45(2), 45-94.

[22] Ramscar, M. & Yarlett, D. (2007) Linguistic self-correction in the absence of feedback: A new approach to the logical problem of language acquisition. *Cognitive Science,* 31, 927-960.

[23] Ramscar, M., & Dye, M. (2009) Expectation and negative evidence in language learning: the curious absence of mouses in adult speech. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*, Amsterdam, Netherlands. A manuscript describing an extended version of this work is under review at *Cognition*.

[24] Ramscar, M., & Gitcho, N. (2007). Developmental change and the nature of learning in childhood. *Trends In Cognitive Science, 11*(7), 274-279.

[25] Thompson-Schill, S. L., Ramscar, M., & Evangelia, C. G., (2009) Cognition without control: When a little frontal lobe goes a long way. *Current Directions in Psychological Science*. 18(5).

[26] Ramscar, M., Dye, M., Witten, J., & Klein, J. (2009) Two routes to cognitive flexibility: Learning and response conflict resolution in the dimensional change card sort task. *Proceedings of the 31st Meeting of the Cognitive Science Society*, Amsterdam, Netherlands.