Semantic Implicit learning

Albertyna Paciorek and John N. Williams

Department of Theoretical and Applied Linguistics, University of Cambridge

A slightly revised version of this manuscript appeared In P. Rebuschat (Ed.), *Implicit and Explicit Learning of Languages* (pp. 69-90): John Benjamins. 2015.

Abstract

Much previous research on implicit learning has examined form-based sequential regularities over letters and syllables. Recently, however, researchers have begun to examine implicit learning of systems in which the regularities are described at the level of meaning. We review existing work in this area, primarily from vision research and natural language. These studies suggest that meaning-based generalisations can be learned without intent and without awareness of what those generalisations are. In the case of language we review work on learning semantic constraints on determiner usage, and the acquisition of semantic preferences of verbs. We discuss outstanding issues: whether noticing of meaning, as well as form, is necessary, whether the effects reflect learning of new form-meaning connections as opposed to tuning of existing ones, and whether some semantic distinctions are more available to the implicit learning process than others.

Knowledge of word meaning is to a large extent intuitive. Native speakers are confident they know what a given word means, yet may have difficulty explaining the finer aspects of its meaning. This becomes particularly apparent in the case of understanding near synonyms. While most English speakers agree that in (1) and (2) either *big* or *large* is acceptable, such a simple interchange is not always an option. For example in (3) *large* appears to be a preferred alternative, whereas in (4) *big* generally sounds more natural.

(1) He entered into a room with shelves full of big/large leather-bound volumes.
(2) That shirt doesn't fit me. It's too big/large.

(3) DVDs are ideal for storing the large/*big amount of data associated with video.
(4) How big/*large are you around the waist?

The above examples illustrate implicit knowledge. It is the knowledge we are not aware that we have, yet which manifests itself in our behaviour in many areas of life, such as sports and games (e.g. knowing how to swim, ride a bicycle, play chess), music (e.g. dancing, playing an instrument), in creative processes (e.g. solving complex mathematical problems), in arts (e.g. having an intuitive sense of genuine art and fashion), even in social cognition (e.g. reading other people's intentions and mental states, understanding humour). A paradigm example of implicit knowledge is the command of one's native language. People instantly notice sentences violating grammar, even without being able to state the underlying principles. Similarly, a wrongly played note will disturb even non-musicians, who cannot competently justify why. For some people, even cooking is a demonstration of highly impressive knowledge which unfortunately often remains incommunicable, despite efforts. Implicit knowledge is therefore the tacit, non-verbalisable knowledge people may have, specific to a particular skill.

Within second language acquisition research Krashen (1981) proposed the distinction between unconscious acquisition and conscious learning, emphasising the acquisition of unconscious linguistic knowledge through implicit learning mechanisms. More recently though, a certain scepticism over this simplistic approach has come to the fore, most clearly expressed in the work of Schmidt (1990, 1995, 2001), and his seminal "noticing" hypothesis. Schmidt was primarily concerned with the roles that attention and awareness play in the learning process. He proposed "noticing" to be the minimal combination of attention and a low level awareness, and it is necessary and sufficient for converting input to intake (Schmidt, 1990). He proposed that stimuli outside focal attention and therefore outside awareness might activate preexisting memory representations subliminally, but that subliminal *learning* is impossible. Once 'noticed', a given aspect of language may be consciously analysed, compared to other things which were 'noticed', giving rise to a higher level of awareness – 'understanding' (which involves what is generally understood as ‘thinking’). Schmidt wrote sceptically about the possibility of forming generalisations and abstract rules without noticing at the level of understanding. This is the question which the research contained in this chapter addresses. Schmidt himself identified “unconscious induction and abstractness of the knowledge that results from learning” as “the most important issues involved in implicit learning studies” (Schmidt, 1995, p.35).

Schmidt provided evidence for his views from a case study of his own learning of Portuguese over a 6-month period in Brazil (Schmidt & Frota, 1986). He documented all aspects of the language he became aware of in a detailed diary, which also included recordings and transcriptions of his conversations with the natives. Schmidt noted a strong tendency for the things he noticed in the input to later appear in his productions, which led him to advocate that conscious ‘noticing’ should be crucial for learning. Additionally, the findings by (Leow, 1997, 2000) challenge the idea that there can be learning without awareness. His studies involved the learning of irregular (stem-changing) verbs in Spanish. Participants were solving a crossword puzzle during which they were required to report any thoughts coming to their mind. Interestingly, only those who commented on the changes of stem in the critical items later showed significant gains between pre-test and post-test. The study therefore suggests that there is no learning without ‘noticing’.

When one considers traditional artificial grammar learning research one can have sympathy with Schmidt’s scepticism about learning generalisations without awareness at the level of understanding. In such experiments people appear to learn fragmentary chunks rather than whole abstract systems, hence exposure to bigrams gives the same result as exposure to whole strings (Perruchet & Pacteau, 1990). Changing the letter set between training and test is a common way of assessing learning of abstract structure, but such effects that are obtained are traceable to knowledge of local patterns of repetition and doubling rather than the abstract structure of the whole grammar (Tunney & Altmann, 2001). No learning effects are obtained for bi-conditional grammars, where grammaticality judgements cannot be based on surface similarity to training examples (Johnstone & Shanks, 2001; Matthews, et al., 1989). Thus, there is little evidence for learning abstract regularities without awareness beyond acquisition of abstract repetition structures.

Research on invariant learning initially appeared to provide evidence of implicit learning of abstract rules that could be applied regardless of similarity to training items. Bright & Burton (1994) documented a study in which after exposure to clock faces showing times between 6-12 o'clock participants appeared to find new clocks displaying times within that range familiar, even if the test clocks were in a different analog/digital format, and crucially, even when no awareness of the regularity was reported in the final debriefing. Initially this was taken as evidence for abstraction of a rule. Subsequently, however, this conclusion was revised by Newell & Bright (2002) who, by carefully controlling the similarity of training and test items showed that, rather than extracting the "6 to 12" rule, participants simply judged test items with regards to how similar they were to training items stored in memory. For further examples and discussion of invariant learning see Berry & Cock (1998).

Findings from invariant learning and artificial grammar learning might therefore lead to the conclusion that abstract regularities are not amenable to implicit learning. However, another, more recent, strand of research has begun to look at the formation of abstract semantic-based generalisations, what we might call 'semantic implicit learning'. These may be divided into two types. The first addresses the question of whether new associations between existing concepts can be formed implicitly, and the second whether connections between novel stimuli and (pre-existing) concepts can be established without awareness. Both of these types of learning are relevant for language acquisition.

Semantic implicit learning in non-language learning domains

Abstract sequences

An example study of the first kind, involving learning a thoroughly conceptual sequence is Goschke & Bolte (2007). Participants named pictures of objects, e.g. *table*, *shirt*, *nose*, *horse*, *trousers*, *ear*. What they did not know was that the pictures followed a repeating sequence based on their underlying category (furniture-clothing-body part-animal-clothing-body part). After some exposure people were faster to name new objects when the object sequence followed the category sequence than when it did not, thereby showing generalisation. The pattern remained significant when taking into account only those participants who had no awareness of the underlying regularity based on the follow up sequence reproduction and recognition tasks and in comparison with a control group. Learning of category sequences is most likely based on forming gradual associations between category representations which are repeatedly activated in succession. As a result, processing an exemplar from one category leads to preactivation of the next category, which speeds up naming of objects coherent with the preactivated category, and delays naming of those incoherent with it.

Brady & Oliva (2008) show a similar effect for pictures of scenes which followed a predictable sequence of categories (e.g. mountain-building-kitchen-forest), including transfer to word equivalents of the category sequences. Participants initially performed a simple task (repetition detection) on a stream of pictures. They were then tested on pairs of triplets of pictures, one of which followed the category sequence from the training, while the other did not. The participants' task was to choose which triplet they felt was more familiar. They were able to discriminate the two types above chance, even when the exact pictures had not appeared in the combinations presented but the categories to which they belonged had, and while reporting no awareness of the category sequences. These studies show that people can track sequential regularities at conceptual levels, rather than just the superficial level of letter identity, such as in artificial grammars (Reber, 1967) or spatial location (Lewicki, Czyzewska, & Hoffman, 1987). This may be quite beneficial in the daily interactions with the world since, as Brady and Oliva, suggest, it reduces information redundancy, which is ubiquitous in the daily organisation of objects and their locations. Thus, if we accept that stimuli can contact deep levels of conceptual representation, and that the same kinds of statistical learning mechanisms that operate over surface forms can operate here too, then we open the possibility for learning at levels of representation that are abstracted away from surface form.

Connecting novel stimuli and concepts

One of the early studies showing that people are sensitive to covariation at different levels of representation (form-level and a higher conceptual level) is Lewicki (1986). He presented participants with a series of photographs of short and long-haired people (visual forms) together with short auditory descriptions describing the people's characters (concepts). Unbeknown to the participants, the long-haired people were all 'kind', while short-haired people could be more accurately described as 'capable' although neither of these precise words appeared in the descriptions. The participants' subsequent forced-choice judgements (e.g., 'capable' – yes/no) of new pictured people, as well as their response latencies, indicated that they indeed learned this covariation. They did so without being able to report any dependencies in the stimuli. Because neither of the words ‘kind’ nor ‘capable’ appeared in the descriptions, participants must have extracted the relevant general notions about characters of the presented people and associated them with hair-length. This reinforces the notion that certain regularities based on semantic, conceptual generalisations are learnable without awareness.

Further interesting demonstrations of this phenomenon come from the area of visual perception and the derived attention paradigm. Lambert & Sumich (1996) showed that participants can become sensitive to the predictive relationship between the semantic category of words appearing on screen (living or non-living) and the position of a target dot located in a subsequent display, without awareness of the regularity throughout the experiment. Thus, novel stimuli (dots on a screen) become associated with abstract concepts. More recently, similar findings have been obtained using the contextual cuing paradigm. In Goujon, Didierjean, & Marmeche (2007) participants had to search for a target (a number 13 or 28) in a display containing other numbers as distractors. The time the participants took to indicate whether the target was on the left or right of the screen was recorded. What they didn’t know was that whether the target was on the left or right of the screen was predicted by whether the distractors were odd or even numbers. After some time performing the task, response times on such predictive trials were faster than on trials where the distracters were a mixture of odd and even numbers. This was the case even when the distractors were drawn from a different set of odd and even numbers from those encountered in training. In a follow-up test participants’ recognition of predictive and counterpredictive displays were at chance and they did not report any awareness of the regularity, so they appeared to have unconsciously abstracted the relationship between odd/even numbers and the target position. This, again, appears to be a case of learning an abstract relationship between a stimulus attribute (left or right of the screen) and an abstract conceptual distinction. Goujon, Didierjean, & Marmeche (2009) report a similar phenomenon in a situation in which target location is predicted by the semantic category of the distracter words (e.g. mammals predicted that the target would be on the left of the screen). Again, an effect was obtained even when the distracter words were drawn from a different set from those used in training.

In Goujon (2011) the category of the scene (e.g., a bedroom or a living room) predicted the location of a target. For example, for pictures of living rooms the target object were nearly always in the lower left corner, and search times were faster when this regularity was respected than when it was not, even for new examples of the scenes. Interestingly, though, the effect only arose when participants were asked to categorise a preview scene (even when it was different from the subsequent search scene) or could at least preview the search scene before the visual search trial, indicating that some processing enhancement may be necessary for such semantic implicit learning effects to occur.

The results of these experiments raise questions about the role of awareness in learning. Participants appear to learn an association between a stimulus attribute (e.g., length of hair or screen location) and a concept without awareness that there is any such association. This seems to contradict the view that awareness is always necessary for learning associations, even of the most basic kinds, such as between a tone and a puff of air in a classical conditioning experiment (Lovibond & Shanks, 2002). One might ask why there is this difference. Could it be that the situations studied in contextual cuing are more natural than those in the typical conditioning experiment? The contingency between a tone and a puff of air is an arbitrary invention of the experimenter, and there is no basis within the participant's prior knowledge for supposing a causal relationship between them. However, the point of contextual cuing experiments is to show how we can learn to orient our attention in the visual world through acquiring relationships between types of contexts and likely positions of interest, both at the low level of visual cues, and the deeper level of semantic categories (Chun, 2000). So one lesson from these experiments might be that when we begin to examine learning in more naturalistic contexts we might be surprised how little awareness is acquired to learn associations.

Although conducted within the area of vision research and scene perception, it is possible to see analogies between the contextual cuing experiments and learning form-meaning connections in language. If we regard possible target locations as forms, then in contextual cuing what participants are learning is how the distribution of those forms is explained by the semantic properties of the context. But this is merely an analogy to word learning in natural language. Is there evidence that similar learning processes are at work when learning how semantic properties of contexts predict the distribution of word forms?

Grammatical form-meaning connections

Initial reports cast doubt on the possibility of learning form-meaning connections implicitly. DeKesyer (1995) utilised an artificial language with rich inflectional morphology for marking gender, number and object role. In an extensive study involving 20 learning sessions, participants were initially asked to indicate whether a given sentence correctly described a picture or not, e.g., *Bep-on warufk-at rip-us* (Worker-PL build-PL house-OBJ). The subsequent test involved a production task, where they had to describe pictures using the novel language. It turned out that they were only able to do so when it was possible to use stem-inflection combinations which had occurred in training. When items required novel stem-inflection combinations they performed at chance, indicating they did not learn the general, semantic properties of the inflectional morphemes. However, the sentence-picture training task that was used in this study did not appear to require participants to pay attention to the inflectional morphemes. Given the importance of noticing form for learning it is therefore not surprising that no learning of form-meaning connections was obtained.

More promising results come from the research by Williams (2005) Participants were first taught a system of four determiners: *gi, ro, ul* and *ne.* They were told that they had a similar role to the English definite article but they also encoded the distance between the speaker and the object. In the system, *ul* and *ne* were used for distant objects, whereas *gi* and *ro* were used for the near ones. They were not told, however, that the choice of the determiner also depended on the animacy of the following noun. Animate objects required either *gi* (near-animate) or *ul* (far-animate), while inanimates required *ro* (near-inanimate) or *ne* (far-inanimate). The participants were shown correct sentences (e.g. *I heard the sound of ul monkey in the tree; I knocked over ro cup and coffee spilled on my book*) and were trained on the distinction between far and near objects. While reading the sentences they were also asked to form mental images of the described situations. After training, they were given a test, where they had to choose between two possible determiners. In the test, familiar words appeared in new contexts, e.g. *The art collector went to Greece to collect ul / ne vases* (correct answer "*ne* vases"). Neither *ul* or *ne* had occurred with the noun *vase* during training. Those who claimed not to have been aware of the relevance of animacy to article usage during training or testing phases answered accurately 61% of the time in one experiment, and 64% in another, significantly above chance in both cases. It appeared that people learned an association between targets (the forms of the articles) and the animacy of the accompanying nouns without being aware of that association. Their attention was drawn independently to the form of the articles through the near/far decision, and to the meaning of the nouns through imagery, yet they remained unaware of the connection between the articles and noun meaning. Analogously to the contextual cuing experiments described above, learning generalised to new exemplars because it is supported by an abstract conceptual distinction.

Subsequent research in this area has explored different methods of assessing awareness of semantic regularities. After all, any demonstration of semantic implicit learning is only as good as the methods used to assess the role of conscious knowledge. Simply asking participants after the experiment if they had noticed a particular regularity is not the most sensitive test of awareness of course, and may fall foul of both the information criterion - test performance must reflect the same knowledge that is probed in the awareness test, and the sensitivity criterion - e.g., people may forget the criteria they were using during the test phase when asked afterwards (Shanks & St. John, 1994). One approach, advocated in SLA research by Leow and colleagues is to use think-aloud protocols. Hama & Leow (2010) report a conceptual replication of the Williams (2005) study in which participants were required to think aloud during the training and test phases. Interestingly they found no instances of participants commenting on the relevance of animacy to the articles during the training phase, confirming that this kind of task successfully prevents the learners from thinking about the hidden regularity. After the exclusion of all participants who mentioned animacy, or who formulated mal-rules during the test phase, there was no evidence of a bias towards selecting the article with the appropriate animacy value in the forced choice test task. However, it should be noted that this study introduced a number of radical procedural changes with respect to the original Williams (2005) study. Most significantly, test items were presented auditorily with a beep sound at the article position, followed by all four possible article choices (e.g., "I had to read the manual to find out how to adjust [beep] clock up on the wall." Choices: (a) gi clock, (b) ro clock, (c) ul clock, (d) ne clock). The task excises the noun phrases from the context, forces them to hold the context in memory, and clearly focuses the participants' attention on the distance variable in guiding article selection. In contrast, in Williams (2005), test items were presented visually, reducing the memory demands of the task, and only two choices of the same distance value were provided, neutralising this factor, and hence potentially allowing animacy biases to have more of an effect. It should also be noted that in Hama & Leow (2010) accuracy was below chance even for items that were repeated from training (compared to over 70% correct in the original study). With such poor memory for actual article-noun combinations it is not surprising that the relevant generalisation was not learned.

An alternative method for assessing awareness during judgment tasks is to require participants to indicate the basis of each judgement, indicating whether their decision was based on a guess, intuition, memory, or rule (Dienes, 2008; Dienes & Scott, 2005; Rebuschat, in press). Accuracy under 'guess' and 'intuition' is assumed to reflect the level of performance obtained when people do not consciously know on what knowledge their decision is based, even though in the case of intuition they may have moderate confidence that they are correct. Chen et al. (2011) repeated the original Williams (2005) study in Chinese, this time requiring participants to indicate the basis for each decision in the test phase. It was found that responses to generalisation items under guess and intuition were 55% correct, significantly above chance. Performance was only slightly higher for the memory- and rule-based responses (58%). None of the participants in this experiment reported awareness of the regularity in the post-experiment debriefing. In a second experiment, novel generalisation items were employed that had not appeared in training at all. Now accuracy for decisions based on guess and intuition was 58%, significantly above chance. This provides evidence for generalisation of the animacy rule beyond the nouns used in training.

A third replication of the Williams (2005) study is reported by Rebuschat et al. (in press), again sticking closely to the original procedure but with the addition of source judgements in the test phase. This time, the majority of the participants (9 out of 13) showed some level of awareness of the animacy rule in the post-experiment questionnaire. Nevertheless, 31% of the responses made by all participants were attributed to guess and intuition, with an accuracy that was significantly above chance (67% and 75% respectively). The provision of the 'rule' option presumably encouraged participants to search for rules during the test phase, but a significant proportion of responses were still based on veridical unconscious knowledge.

One way of avoiding the complications associated with verbal report and subjective measures is to construct a test situation in which the test task does not directly tap the knowledge of interest. Ideally there should be no discernable test phase from the participant's point of view, making it less likely that they will enter a rule discovery mode because they think that their knowledge is being tested. Leung & Williams (2011, 2012) developed a reaction time procedure to do this. For example, Leung & Williams (2011) used a system in which articles before proper names signalled thematic roles (e.g. *gi* was associated with agents). The participants' task was merely to indicate on which side of picture the named person was located, for example in a scene in which a boy is pushing a girl into a pool. For a person who knew the system, on hearing "gi Tom" their attention would be oriented to the agent even before hearing the proper (male) name, and their responses would be predicted to be disturbed in the case where the wrong article was used in the context of an actor that turned out to have a different role (e.g. "ro Tom" where *ro* refers to patients). This effect was indeed obtained in reaction times, but only for those participants who claimed to be unaware of the correlation between articles and thematic roles in a post-experiment debriefing (80% of the participants). These experiments demonstrate how implicitly learned semantic regularities can influence online processing, as opposed to off-line judgements (Leung & Williams, 2012 used a similar method to show implicit learning of animacy constraints on article usage). Speeded responses are more likely (though not guaranteed) to tap implicit knowledge (Ellis, 2005), but when combined with verbal report this method provides relatively robust evidence for the acquisition of implicit semantic knowledge.

Learning constraints on article usage has become something of a test case for the principle of semantic implicit learning in language. However, it is a domain in which prior expectations about possible grammatical agreement relations between words might have an effect (and indeed, some evidence for the effect of such expectations deriving from prior language learning background was presented in

Williams, 2005). Is there evidence for semantic implicit learning outside of the grammatical domain, where such expectations would be less likely to have an effect?

Implicit learning of semantic preferences

A semantic preference can be understood as a particular type of collocation, where 'collocation' refers to higher than chance co-occurrence of two or more words. Collocates sound natural together and substituting one of them with a near-synonym results in a loss of naturalness for native speakers, for example in English it is better to say *fast car* and *fast food*, rather than **quick car* or **quick food*. Conversely, it is more natural to say *quick glance* and *quick meal* instead of **fast glance* or **fast meal*.

It has been traditionally proposed that collocations reflect syntagmatic relations between words, therefore related to their surface structure, rather than paradigmatic relations – regarding their meaning. "Meaning by collocation is an abstraction at the syntagmatic level and is not directly concerned with the conceptual or idea approach to the meaning of words. One of the meanings of 'night' is its collocability with 'dark'..." (Firth, 1957). However, syntagmatic regularities may not be the optimal, or the sole way of accounting for the existence and acquisition of semantic preferences. After all, new collocates can be freely generated, as long as they follow implicit assumptions regarding applicable semantic sets. For example, in English it is correct to say *a pack of dogs/hounds/wolves/coyotes/mules/rats/weasels* and *a swarm of bees/insects/mosquitoes/bats/ladybirds*. It sounds unnatural and even humorous to say **a pack of bats*, **a pack of insects* or **a swarm of dogs*. It is therefore plausible that in the process of learning, the speaker's knowledge of the meanings of these words gains the shape of, roughly: SWARM + [ANIMAL, FLYING, rather SMALL, most likely an INSECT], PACK + [ANIMAL, FOUR-LEGGED, most probably DOG-LIKE]. It makes sense therefore, to predict that the existence of such semantically preferred sets of collocates involves abstraction at a level higher than form. Intuitions about collocability form an important aspect of attaining a native-like command of a language.

Idiomatic expressions may be numerous in language, yet even more numerous are arrays of less fixed constructions. The examples are plentiful and may involve various degrees of abstraction around the set of collocates. In phrases such as *beyond* *belief/description/doubt/question/recognition/repair* the noun that follows the verb, although different in each case, always indicates some "verbal process", and the meaning of the whole may be roughly paraphrased as "too great to believe/describe/doubt...". Schmitt (2010) describes such phenomena as "formulae with open slots". The existence of these variable expressions, using Sinclair's (2004) terminology, poses a challenge for learners of both first and second language. Yet what the above examples show, is that having a generalised notion of what meaning is required in a given slot in a construction may potentially be very helpful in determining which words can and cannot be used. Native speakers develop such generalised notions incidentally, and possibly are not always aware of precisely what identifies legitimate collocates. Is it possible for second language learners to develop such notions implicitly as well?

In Paciorek (2012) and Paciorek & Williams (submitted) participants read sentences containing four novel verbs, e.g. *powter, mouten, gouble,* and *fonell*. For each sentence they had to indicate whether the novel verb meant, broadly, to increase or decrease something, and to also indicate how important the content was to them. Examples include (with correct response in brackets): *Nightingale worked tirelessly to improve public health and POWTER the status of nurses* (increase); *Start the day with a stewed apple. It is light on your digestion and will MOUTEN nutrients rather than emptying calories* (increase); *Avoid stimulants because they FONELL serotonin over time - this works the same as stress does* (decrease). Other sentences did not contain a novel verb, so only a required a judgment of importance, e.g. *In clinical trials, peptides derived from food proteins have shown an effect on blood pressure; Emergency treatment is indicated if potassium is very high, or if severe symptoms are present*. After this task there was a surprise test in which participants were presented with pairs of words and had to indicate if they remembered seeing them together in the same sentence in the previous part. They also had to rate their confidence in their decision. All words had occurred in training sentences, but not necessarily together, e.g. *MOUTEN nutrients* (correct answer yes), *FONELL proteins* (correct answer no), *POWTER potassium* (correct answer no). Unbeknownst to the participants there was a hidden regularity concerning thecollocates of the novel verbs – *POWTER* and *GOUBLE* (meaning increase and decrease respectively) went with abstract collocates, and *MOUTEN* and *FONELL* (meaning increase and decrease) went with concrete collocates. The critical factor in the test phase was that half of the 'no' items respected this regularity (e.g. *FONELL proteins*) and half of them violated it (e.g. *POWTER potassium*). The question was whether the probability of thinking that a new word combination had been seen before would be higher for pairs that respected the regularity than those that did not.

The logic of this task rests on the false memory paradigm (Deese, 1959; Roediger & Mcdermott, 1995). For example, people erroneously report that they previously encountered a word like *mug* if the study phase contains highly related words (e.g., *cup, beer, bowl, coffee*). This procedure has been used before in the implicit learning literature to examine invariant learning in the "Past Midnight" studies (Bright & Burton, 1994), and in a study of word order pattern learning (Cleary & Langley, 2007). The advantage of the paradigm is that participants are concerned simply to judge their memory for items, not whether those items respect a regularity or not. Any influence of this regularity on memory performance would be indirect and automatic, and in defiance of the participants' attempts to judge the strength of the memory traces for events.

Over a series of experiments, Paciorek (2012) consistently found the predicted false memory effect – participants were more likely to erroneously think that they had seen new word pairs that respected the semantic preference rules than those that did not. Crucially, this was the case even amongst participants who, in a post-experiment forced choice task with think-aloud, appeared to have no knowledge of the semantic preference rules. This is an impressive result because the training task only required the participants to consider the potential increase/decrease dimension of the novel verb meaning, as inferred from the context, and did not encourage particular attention to the collocate. Thus, the experiment demonstrates semantic implicit learning in a highly natural situation in which participants encounter novel words in text, attempt to discern some aspect of their meaning, and at the same time acquire implicit knowledge of semantic preferences.

Further analyses of the confidence data suggested that in this paradigm, even if participants do have conscious knowledge they do not necessarily use it strategically in the memory task. The signature of conscious knowledge in this task would be a relatively high confidence in rejecting new word pairs that violate the preference rule (because if a pair violates the rule it is unlikely to have been encountered in training). However, evidence of this was only found for the participants with the clearest understanding of the system. Participants classed as only partially aware did not show any evidence that higher confidence in rejecting new ill-formed pairs was associated with a larger learning effect. This even applied to participants who had received prior instruction in the system, but showed imperfect recall of it in the post-test. It appeared that even participants with some conscious knowledge of the system did not use this knowledge strategically in the memory task. Clearly it is critical to distinguish having some reportable knowledge in a post-experiment questionnaire or task, and actually using that knowledge strategically to influence performance during the test phase (recall also the Rebuschat et al, in press, study where participants with reportable knowledge also produced above-chance judgements based on guess and intuition). When the test task is an indirect measure of the relevant knowledge, as here, it is quite possible that participants who have conscious knowledge do not use it strategically to perform the task. After all, this requires a level of reasoning that goes beyond the immediate task requirements. In such situations then, post-experiment verbal report is likely to over-estimate the extent to which test performance is actually influenced by conscious knowledge. Excluding all participants who show even partial awareness, as in the Paciorek studies, is therefore a conservative procedure.

Emerging issues in semantic implicit learning of language

In the above we have summarised recent research on implicit learning in language. Most of this work has been concerned to refine the measurement of awareness and to devise procedures that make contamination from conscious knowledge less and less likely. But there are many theoretical issues that need to be resolved before we can claim to have a full understanding of this phenomenon.

Awareness of form and meaning in learning form-meaning connections

According to the noticing hypothesis, awareness of form is critical to learning: “the objects of attention and noticing are elements of the surface structure of utterances in the input, instances of language, rather than any abstract rules or principles of which such instances may be exemplars” (Schmidt, 2001). Thus, in order to learn about form, it is necessary to pay attention to, and be aware of, form. Cognitive psychological research strongly supports the view that attention is a prerequisite for learning (Jiménez & Méndez, 1999; Logan & Etherton, 1994; Toro, Sinnett, & Soto-Faraco, 2005). In order for associations between stimuli to be learned it is necessary that those stimuli are attended, or more subtly, that the relevant dimensions of those stimuli are attended (see Williams, in press, for a review).

None of the demonstrations of implicit semantic learning contradict this aspect of the noticing hypothesis. In all cases the participants attend to, and are required by the task to be aware of, the relevant forms in order to make responses. The issue arises in relation to awareness of meaning, since Schmidt also states that "In order to acquire vocabulary one must attend to both word form (pronunciation, spelling) and to whatever clues are available in input that can lead to identification of meaning. In order to acquire pragmatics, one must attend to both the linguistic form of utterances and the relevant social and contextual features with which they are associated. In order to acquire morphology (both derivational and inflectional), one must attend to both the forms of morphemes and their meanings, and in order to acquire syntax one must attend to the order of words and the meanings they are associated with" (Schmidt, 2001). The question raised by the research reviewed above is whether it really is necessary to notice the relevant meaning in order for learning to occur.

In some cases, noticing of meaning does indeed seem to occur, even though participants remain unaware of its relevance. In two of Goujon et al's (2011) experiments the participants were required to categorise a scene before searching for the target, thereby encouraging them to notice the scenes as instances of the relevant categories (kitchens, living rooms, etc). In Williams & Leung (2012) learning effects in reaction time were only obtained when the participants were required to indicate whether the named object (e.g. "gi dog") was living or non-living. In these cases awareness of the relevant meaning is required by the task. Yet most of the participants still remained unaware of the systematic association between form and meaning, showing no awareness of the "abstract rules or principles of which such instances may be exemplars", or awareness at the "level of understanding" (Schmidt, 2001).

In other cases we don't know if the participants were aware of the relevant meaning or not. In contextual cuing experiments it is possible that the participants were aware that displays tended to contain mostly either odd or even numbers (Goujon, 2007) or words of certain categories (Goujon, 2009). This information was not elicited in the debriefing. But in the case of the Williams (2005) procedure involving reading sentences and forming images, Hama & Leow (2011) found no mentions of animacy in the think-aloud protocols from the training task. And it seems unlikely that in Paciorek's (2012) false memory experiments the participants noticed whether the object of a novel verb in a sentence was abstract or concrete.

Intuitively, at least, it seems plausible that in these cases the relevant semantic information was in itself implicit in the participant's understanding of the context. As such it was not "noticed", but it was active enough in working memory to participate in associative learning. In fact, to borrow a term from discourse processing, one might argue that the relevant information is in "implicit focus" (Sanford & Garrod, 1981).[1] Linguistic contexts make available a large amount of semantic information that remains implicit in our understanding of events, and which is available to guide inferencing, and, we would argue, participate in learning.

The question remains why in some cases it is necessary to force participants to notice the relevant meaning. In fact, in Goujon's (2011) experiment, overt categorisation of the scenes did not turn out to be necessary. Learning could be obtained just by allowing the participants a preview of the scene before the target appeared. All that may be required then is time for sufficient semantic processing of the context. Leung & Williams (2012) represents the only language study so far not to embed the learning targets in complete sentences. It is possible that in the absence of a sentence context there is insufficiently strong activation of the relevant semantic information for learning to occur. As one might expect, semantic implicit learning depends on there being a sufficient depth of semantic processing, but this does not require that the relevant meanings are noticed.

Learning new form-meaning connections versus tuning

All of the studies of semantic implicit learning in language described here share one important methodological feature – participants are encouraged to map novel word forms onto pre-specified meanings, e.g. to 'near' and 'far' (Williams, 2005; Leung & Williams, 2012), or increase/decrease (Paciorek, 2012). The implicit learning effects

---

[1] To make this concept clear, consider the case of bridging inferences. When we read "Mary dressed the baby" we implicitly know that clothes were put on the baby, and hence have no trouble reading the continuation "The clothes were made of pink wool". Strictly speaking 'the clothes' should be problematic because the definite description requires an antecedent, and none is explicitly mentioned in the previous sentence. But 'dressed' entails clothing, and this information is active in working memory, or in Sanford & Garrod's (1981) terms, in "implicit focus", and hence is available as an antecedent.

relate to the acquisition of further semantic constraints on the use of the words, over and above those consciously entertained by the participants. We regard this as an important aspect of naturalistic vocabulary acquisition, whereby an initial rough hypothesis about the meaning of a word becomes refined through experience. It is this process of refinement, or “tuning”, that we claim is implicit.

The question remains, therefore, whether the creation of novel form-meaning connections can occur implicitly. One common line of argumentation against this comes from the fact that vocabulary acquisition in amnesic patients is grossly impaired (Glisky, Schacter, & Tulving, 1986), suggesting that vocabulary learning depends on declarative, memory systems, especially the hippocampus. However, this does not in itself mean that vocabulary learning requires conscious awareness of the relevant form-meaning connections. The hippocampus is important for forming relational representations (Eichenbaum, Otto, & Cohen, 1994), and hence is critical even for forms of implicit learning such as contextual cuing (Chun & Phelps, 1999; Park, et al., 2004), and, plausibly, learning form-meaning connections. In the intact brain it is perfectly possible for learning to be dependent on the declarative system and yet still be implicit (Ullman, 2005).

A stronger argument comes from research on joint attention during vocabulary learning in infants (Bloom, 2000). Children will only learn that a word uttered by an adult refers to an object if the child is attending to the object, and the child is aware that the adult is also attending to that same object. So, for example, words will not be learned as labels for objects if they just happen to be played over a loudspeaker in the moment that the child is attending to the object. Vocabulary learning is not a simple matter of associative learning, but depends upon assumptions about the intentions of speakers.

But even if an initial form-meaning connection depends on conscious attention to both form and a hypothesised meaning, this does not mean that further tuning of that meaning requires awareness. Indeed, Paradis (2004) proposed a terminological distinction between “vocabulary” and “lexicon”, where ‘vocabulary’ refers to the referential meaning of words, supported by the declarative memory system, and “lexicon” refers to, for example, grammatical constraints on word usage, and collocational behaviour. He proposes that these aspects are supported by the procedural system, and hence acquired and represented implicitly. Whilst we would not want to assume that the kind of “lexical” knowledge being acquired in the experiments described here is represented in the procedural system, we would argue that it is implicit.

It actually remains an open empirical question whether associations between entirely novel forms and meanings can be laid down unconsciously in memory in the absence of awareness, and thereby participate in some way in the development of referential meaning. The work on semantic implicit learning in language reviewed here does not address this issue. Rather, these studies show that after seeding the word learning process with some explicit knowledge of meaning, further, subtle, biases in usage can be acquired unconsciously.

Constraints on semantic implicit learning

Statistical and associative learning mechanisms are often thought to be unconstrained in the sense that any and all patterns of association in the environment are learnable. However, there are already indications that not all conceptual distinctions are equally amenable to implicit learning. Leung & Williams (2012) found implicit learning of a correlation between novel articles and noun animacy, but there was no effect when article usage was determined by the relative size of an object (for example, *gi* would be used for the smaller of two objects on the screen, regardless of their absolute size, hence with a fish in the context of an armchair, but with a paperclip in the context of a fish). Chen et al. (2011) obtained a similar contrast between learning of correlations between determiners and animacy and no learning of a correlation with size relative to a dog. The concept of relative size is not one that is typically encoded by grammatical morphemes, so these null results raise the question of whether semantic implicit learning in this domain is constrained by knowledge of what Bickerton (2001) referred to as "potentially encodable distinctions", a list of semantic distinctions that are encoded in the world's languages (such as animacy and thematic role, amongst others). Clearly far more work is required to test this idea, and to distinguish the effect of linguistic universals from the general cognitive availability of different conceptual distinctions. For example, the pattern of breakdown of conceptual knowledge after brain damage often reflects gross categorical distinctions between living and non-living things, or abstract and concrete concepts, suggesting that these distinctions are fundamental to conceptual representation in the brain (Warrington & Shallice, 1984). One might also wonder whether the extent to which a semantic distinction is encoded in one's L1 influences the ease with which it can be implicitly acquired in the L2. These are all questions that await further investigation as we seek to understand how the implicit learning mechanism is constrained by linguistic knowledge.

Conclusion

One of the aims of this chapter has been to draw out the connections between research on semantic implicit learning within general cognitive psychology, primarily vision research, and language acquisition. There are domain-general learning processes at work here that are able to abstract over events to form meaning-based generalisations, processes that, it is claimed, operate without conscious intent, and without awareness of the resulting knowledge. What is also noteworthy is the ease with which generalisation to new instances of language is obtained. In learning form-level regularities, such as in letter sequences generated by artificial grammars, changes in surface form always result in reduced learning effects, what Pacton et al. (2001) refer to as the "transfer decrement". In contrast, when semantic regularities are employed, generalisation to new forms is readily obtained, a reflection of the fact that the regularity is defined at a level that is independent of form. Thus semantic implicit learning may be a particularly important mechanism for explaining the productivity of language.

Clearly much of the work to date has also been concerned to establish that knowledge of semantic-based regularities can be acquired without awareness of what those regularities are. Such questions will always be difficult to answer convincingly simply because of the inherent difficulty of proving the non-existence of something (in this case, conscious awareness). We would argue, however, that the primary issue is the ability to acquire knowledge spontaneously without employing consciously directed explicit learning strategies. In fact current research on "statistical learning" appeals to the same kinds of simple associative learning mechanisms that are assumed to underlie implicit learning but without being concerned with the learners' awareness of what is learned (Misyak, Goldstein, & Christiansen, 2012). The difference is that in statistical learning the nature of the assumed computations (e.g. tracking transition probabilities) is clearly unconscious and hence the issue of contamination from conscious learning strategies hardly arises, even though these might well result in restructuring of conscious experience (Perruchet & Vinter, 1998). However, in semantic implicit learning research the target systems are relatively simple, and in principle amenable to explicit learning. The emphasis on the unconscious nature knowledge is a strategy for ensuring that the learning process was implicit. Implicitness of knowledge guarantees implicitness of acquisition (except in the case of automatisation of explicit knowledge through massive amounts of practice, which was not the case in the studies discussed here). But implicitness of acquisition does not guarantee implicitness of knowledge, since it is intuitively plausible that unconscious knowledge somehow emerges into awareness, directly or not, through processes of "insight", processes which are in themselves highly interesting, but under-researched (see Runger & Frensch, 2008, for discussion in the context of sequence learning). Ruling out all cases of conscious knowledge, whether gained through explicit learning or spontaneous insight, is a conservative strategy for focussing on the workings of the implicit learning process. It is primarily in relation to this kind of learning mechanism that questions such as the importance of linguistic constraints, and the availability of different conceptual distinctions, become theoretically interesting, and lead us into a potentially deeper understanding of the way that simple learning mechanisms might interact with complex representational structures. The methods that have been developed in this line of research should be seen as tools for investigating the nature of the unconscious learning mechanisms, and not ends in themselves for establishing the existence of unconscious knowledge.

References

Bickerton, D. (2001). Okay for content words, but what about functional items? Commentary on Bloom: How children learn the meanings of words. *Behavioral and Brain Sciences, 24*, 1104-1105.

Bloom, P. (2000). *How Children Learn the Meanings of Words*: MIT Press.

Brady, T. F., & Oliva, A. (2008). Statistical learning using real-world scenes: Extracting categorical regularities without conscious intent. *Psychological Science, 19*, 678-685.

Chen, W. W., Guo, X. Y., Tang, J. H., Zhu, L., Yang, Z. L., & Dienes, Z. (2011). Unconscious structural knowledge of form-meaning connections. *Consciousness and Cognition, 20*, 1751-1760.

Chun, M. M. (2000). Contextual cueing of visual attention. *Trends in Cognitive Sciences, 4*, 170-178.

Chun, M. M., & Phelps, E. A. (1999). Memory deficits for implicit contextual information in amnesic subjects with hippocampal damage. *Nature Neuroscience, 2*, 844-847.

Cleary, A. M., & Langley, M. M. (2007). Retention of the structure underlying sentences. *Language and Cognitive Processes, 22*, 614-628.

Deese, J. (1959). On the prediction of occurrence of particular verbal intrusions in immediate recall. *Journal of Experimental Psychology, 58*, 17-22.

DeKeyser, R. M. (1995). Learning second language grammar rules: An experiment with a miniature linguistic system. *Studies in Second Language Acquisition, 17*, 379-410.

Dienes, Z. (2008). Subjective measures of unconscious knowledge. *Progress in Brain Research, 168*, 49 - 64.

Dienes, Z., & Scott, R. (2005). Measuring unconscious knowledge: distinguishing structural knowledge and judgment knowledge. *Psychological Research, 69*, 338-351.

Eichenbaum, H., Otto, T., & Cohen, N. J. (1994). Two functional components of the hippocampal memory system. *Behavioral and Brain Sciences, 17*, 449-518.

Ellis, R. (2005). Measuring implicit and explicit knowledge of a second language: A psychometric study. *Studies in Second Language Acquisition, 27*, 141-172.

Firth, J. R. (1957). Modes of meaning *Papers in Linguistics 1934-1951* (pp. 190-215). Oxford: Oxford University Press.

Glisky, E. L., Schacter, D. L., & Tulving, E. (1986). Computer learning by memory-impaired patients: Acquisition and retention of complex knowledge. *Neuropsychologia, 24*, 313-328.

Goschke, T., & Bolte, A. (2007). Implicit learning of semantic category sequences: Response-independent acquisition of abstract sequential regularities. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 33*, 394-406.

Goujon, A. (2011). Categorical implicit learning in real-world scenes: Evidence from contextual cueing. *Quarterly Journal of Experimental Psychology, 64*, 920-941.

Goujon, A., Didierjean, A., & Marmeche, E. (2007). Contextual cueing based on specific and categorical properties of the environment. *Visual Cognition, 15*, 257-275.

Goujon, A., Didierjean, A., & Marmeche, E. (2009). Semantic contextual cuing and visual attention. *Journal of Experimental Psychology: Human Perception and Performance, 35*, 50-71.

Hama, M., & Leow, R. P. (2010). Learning without awareness revisited. *Studies in Second Language Acquisition, 32*, 465-491.

Jiménez, L., & Méndez, C. (1999). Which attention is needed for implicit sequence learning? *Journal of Experimental Psychology: Learning, Memory, and Cognition, 25*, 236-259.

Johnstone, T., & Shanks, D. R. (2001). Abstractionist and processing accounts of implicit learning. *Cognitive Psychology, 42*, 61-112.

Krashen, S. (1981). *Second language acquisition and second language learning*. London: Pergamon.

Leow, R. P. (1997). Attention, awareness, and foreign language behavior. *Language Learning, 47*, 467-505.

Leow, R. P. (2000). A study of the role of awareness in foreign language behavior. *Studies in Second Language Acquisition, 22*, 557-584.

Leung, J. H. C., & Williams, J. N. (2011). The implicit learning of mappings between forms and contextually derived meanings. *Studies in Second Language Acquisition, 33*, 33-55.

Leung, J. H. C., & Williams, J. N. (2012). Constraints on implicit Learning of grammatical form-meaning connections. *Language Learning, 62*, 634-662.

Lewicki, P. (1986). Processing information about covariations that cannot be articulated. *Journal of Experimental Psychology: Learning Memory and Cognition, 12*, 135-146.

Lewicki, P., Czyzewska, M., & Hoffman, H. (1987). Unconscious acquisition of complex procedural knowledge. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 13*, 523-530.

Logan, G. D., & Etherton, J. L. (1994). What is learned during automatization? The role of attention in constructing an instance. *Journal of Experimental Psychology: Learning, Memory and Cognition, 20*, 1022-1050.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes, 28*, 3-26.

Matthews, R. C., Buss, R. R., Stanley, W. B., Blanchard-Fields, F., Cho, J.-R., & Druhan, B. (1989). The role of implicit and explicit processes in learning from examples: A synergistic effect. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 15*, 1083-1100.

Misyak, J. B., Goldstein, M. H., & Christiansen, M. H. (2012). Statistical-sequential learning and development. In P. Rebuschat & J. N. Williams (Eds.), *Statistical Learning and Language Acquisition* (pp. 13-54). Boston: Walter de Gruyter.

Paciorek, A. W. (2012). *Implicit Learning of Semantic Preferences.* University of Cambridge, Cambridge.

Paciorek, A. W., & Williams, J. N. (submitted). Implicit learning of selectional preferences of verbs.

Pacton, S., Perruchet, P., Fayol, M., & Cleeremans, A. (2001). Implicit learning out of the lab: The case of orthographic regularities. *Journal of Experimental Psychology: General, 130*, 401-426.

Paradis, M. (2004). *A Neurolinguistic Theory of Bilingualism*. Amsterdam/Philadelphia: John Benjamins.

Park, H., Quinlan, J., Thornton, E., & Reder, L. M. (2004). The effect of midazolam on visual search: Implications for understanding amnesia. *Proceedings of the National Academy of Sciences, 101*, 17879–17883.

Perruchet, P., & Pacteau, C. (1990). Synthetic grammar learning: Implicit rule abstraction or explicit fragmentary knowledge? *Journal of Experimental Psychology: General, 119*, 264-275.

Perruchet, P., & Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language, 39*, 246-263.

Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior, 6*, 855-863.

Rebuschat, P. (in press). Measuring awareness in implicit learning research: A review. *Language Learning*.

Rebuschat, P., Hamrick, P., Sachs, R., Riestenberg, K., & Ziegler, N. (in press). Implicit and explicit knowledge of form-meaning connections: Evidence from subjective measures of awareness *Festschrift for Richard Schmidt*.

Roediger, H. L., & Mcdermott, K. B. (1995). Creating false memories: Remembering words not presented in lists. *Journal of Experimental Psychology: Learning Memory and Cognition, 21*, 803-814.

Runger, D., & Frensch, P. A. (2008). How incidental sequence learning creates reportable knowledge: The role of unexpected events. *Journal of Experimental Psychology: Learning, Memory, and Cognition, 34*, 1011-1026.

Sanford, A. J., & Garrod, S. C. (1981). *Understanding written language: Explorations in comprehension beyond the sentence*. Chichester: Wiley.

Schmidt, R. (1990). The role of consciousness in second language learning. *Applied Linguistics, 11*, 129-158.

Schmidt, R. (1995). Consciousness and foreign language learning: A tutorial on the role of attention and awareness in learning. In R. Schmidt (Ed.), *Attention and Awareness in Foreign Language Learning* (pp. 1-63). Honolulu: Hawai'i: University of Hawai'i Press.

Schmidt, R. (2001). Attention. In P. Robinson (Ed.), *Cognition and Second Language Instruction* (pp. 3-32). Cambridge: Cambridge University Press.

Schmidt, R., & Frota, S. N. (1986). Developing basic conversational ability in a second language: A case study of an adult learner of Portuguese. In R. R. Day (Ed.), *Talking to learn: Conversation in second language acquisition* (pp. 237-326): Rowley, MA: Newbury House.

Schmitt, N. (2010). *Researching Vocabulary: A Vocabulary Research Manual*: Palgrave Press.

Shanks, D. R., & St. John, M. (1994). Characteristics of dissociable human learning systems. *Behavioral and Brain Sciences, 17*, 367-447.

Sinclair, J. M. (2004). *Trust the Text: Lexis, Corpus, Discourse*. London: Routlege.

Toro, J. M., Sinnett, S., & Soto-Faraco, S. (2005). Speech segmentation by statistical learning depends on attention. *Cognition, 97*, B25-B34.

Tunney, R. J., & Altmann, G. T. M. (2001). Two modes of transfer in artificial grammar learning. *Journal of Experimental Psychology: Learning Memory and Cognition, 27*, 614-639.

Ullman, M. T. (2005). A cognitive neuroscience perspective on second language acquisition: The Declarative/Procedural Model. In C. Sanz (Ed.), *Mind and Context in Adult Second Language Acquisition: Methods, Theory, and Practice* (pp. 141-178). Washington, DC: Georgetown University Press.

Warrington, E. K., & Shallice, T. (1984). Category specific semantic impairments. *Brain, 107*, 829-854.

Williams, J. N. (2005). Learning without awareness. *Studies in Second Language Acquisition, 27*, 269-304.

Williams, J. N. (in press). Attention, awareness, and noticing in language processing and learning *Festschrift for Richard Schmidt*.