Semantic Implicit Learning in Language

Albertyna Paciorek (awp23@cam.ac.uk)
John N. Williams (jnw12@cam.ac.uk)

Department of Theoretical and Applied Linguistics, University of Cambridge
9 West Road, Cambridge CB3 9DP, United Kingdom

Abstract

Previous studies of semantic implicit learning in language have only examined learning grammatical form-meaning connections where learning could have been supported by prior linguistic knowledge. Also, these studies assessed awareness by verbal report, which is arguably not the most reliable measure. Here we target the domain of verb meaning, specifically semantic preferences of novel verbs (e.g. a novel verb takes abstract objects). Using a reaction time methodology we show that after exposure to correct verb-noun combinations, reaction times to incorrect combinations are slowed down even for participants who are unaware of the semantic regularity. This effect was also obtained even when the semantic regularity was irrelevant to the tasks being performed, suggesting that the semantic generalisation is learned and exerts its influence automatically, hence satisfying one criterion for implicitness. Combined with a lack of verbalisable knowledge in any participant these experiments provide strong evidence for semantic implicit learning in language.

Keywords: implicit learning; consciousness; form-meaning connections; vocabulary learning; verb learning, second language acquisition; automaticity.

Introduction

Most research on implicit learning has examined regularities at the level of form, be they in sequences of letters generated by artificial grammars, screen positions in repeating sequences, and in the domain of language, phonological (Dell et al., 2001) and orthographic (Pacton et al., 2001) patterns. This limits generalizability to other aspects of language learning where regularities might be conditioned by distinctions at the level of meaning, as opposed to form. Some research in visual perception has exposed semantic-based implicit learning, notably using the contextual cuing paradigm, where target locations are predicted from semantic properties of contexts, (Goujon, 2007). But can semantic implicit learning effects be obtained in the domain of language, especially in the adult language learner? Given arguments that even in children vocabulary acquisition requires declarative, explicit, memory

(Ellis, 1994) and shared attention (Bloom, 2000) one might suspect not. However, these arguments relate to learning referential meaning. Others have hypothesised that other aspects of word meaning, such as connotation and collocational behaviour, might be learned implicitly by the non-declarative system (Paradis, 2004). Here we test this proposal in the context of semantic preferences of verbs.

Previous research on semantic implicit learning in language has focused exclusively on article-noun agreement regularities (e.g. Williams, 2005; Leung & Williams, 2012). This work has demonstrated sensitivity to the semantic properties of nouns in learning about the distribution of articles in miniature semi-artificial languages, and has provided evidence of implicitness of knowledge through post-experiment verbal report. The present experiments extend this work in terms of generalizability to other aspects of language, and in terms of methodology. We will discuss methodology in Experiment 2, but with regard to generalizability, there was evidence in these earlier experiments that implicit learning was dependent upon prior knowledge of article agreement systems in other languages (Williams, 2005). There was also evidence that learning depended on the semantic regularity in question, since effects were obtained for agreement based on animacy, but not on relative size (Leung & Williams, 2012). Putting these two together one might argue that learning was dependent both on prior knowledge of the potential for article noun agreement, and on dispositions based on the "potentially encodable distinctions" that can be grammaticised in language (Bickerton, 1999). The question arises, therefore, whether similar effects could be obtained for an aspect of language that falls outside the realm of grammar, and is not so potentially affected by prior dispositions. This motivated the current investigation of learning the collocational behaviour, specifically semantic preferences, of novel verbs.

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, and are related to surface form, rather than paradigmatic relations regarding meaning (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, knowing that *pack* collocates with *dog, hounds, wolves* etc. while *swarm* with *bees, mosquitoes, bats*, naturally suggests other animals which would be appropriate in either set. It makes sense therefore, to predict that the existence of such semantically preferred sets of collocates involves generalisations at a level higher than form. The question is, can such semantic generalisations be learned implicitly?

The present experiments used four novel verbs, *powter, mouten, gouble*, and *conell*. The participants were exposed to these verbs in verb phrases containing a direct object noun. Their task required them to think about whether the verb conveyed an 'increase' or 'decrease' meaning, either as inferred from the context (Experiment 1) or as they had learned before the experiment (Experiment 2). What they were not told was that *powter* and *gouble* took abstract collocates, whereas *mouten* and *conell* took concrete collocates (Table 1). For example, correct verb phrases would be *powter the significance, gouble the power, mouten the calcium, conell the chocolate*. We tested whether participants would learn the semantic preferences of these novel verbs implicitly using two techniques. Experiment 1 embedded the verb phrases in sentence contexts, and required participants to make an explicit concreteness decision on the nouns, and indicate whether the verb meant increase or decrease at the end of the sentence. The prediction was that after exposure to many correctly formed trials concreteness decision times and/or increase/decrease response times would be faster to new verb-noun combinations that respect the rule than to combinations that do not (e.g. *powter* with a concrete collocate such as *compost*). These will be referred to as the New Grammatical (NG) and New Ungrammatical (NU) conditions respectively. If learning is implicit then this effect would be obtained even for participants who evinced no awareness of the relevance of concreteness to the collocational behavior of the verbs, as assessed by a post-test.

In Experiment 2 participants saw only the verb phrases, but this time they had to decide whether the noun conveyed positive or negative connotations (a decision that is subjective and irrelevant to the hidden regularity). This was followed by the increase/decrease decision, as before. Any effect here would arguably provide stronger evidence for implicit learning, and would speak to the automaticity of the semantic activation underlying the effect.

Table 1. The novel verbs used in the experiments

		Participants not told	
		Abstract collocate	Concrete collocate
Participants told to infer from context	increase	*powter*	*mouten*
	decrease	*gouble*	*conell*

Experiment 1

Participants

40 students of the University of Cambridge participated in the experiment. 17 were native speakers of English. All of the nonnative participants had achieved at least IELTS 7.5.

Materials and Procedure

A total of 80 sentences were created, 20 for each novel verb, in which the verb conveyed either an 'increase' or 'decrease' meaning with respect to the object. For procedural reasons the word order was scrambled so that the verb phrase occurred at the beginning of the sentence. Examples are shown in Table 2.

Table 2. Example sentences from Experiment 1

POWTER the prestige of wealthy families, artists can.
GOUBLE the role of nuclear weapons, Obama stresses the need to.
MOUTEN the nutrients you need, make sure you.
CONELL the histamine stores, the sweating helps to.

The experiment comprised two blocks of trials, although the participants were not aware of any division between them. In the first block there were 44 training trials in which each novel verb occurred equally often. The collocates occurred with both increase and decrease verbs (e.g. *POWTER the*

prestige and *GOUBLE the prestige* occurred, but in different sentence contexts). Block 2 contained 32 critical trials in which the novel verbs occurred in new sentences and with *new* collocates not encountered in Block 1. Half of these items respected the semantic preference rule ("new grammatical", NG, condition), and half violated it, for example by pairing a concrete noun with POWTER ("new ungrammatical", NU, condition). The new collocates in the critical test items were chosen so as to be roughly synonymous with an object noun that occurred in the training phase (e.g. *POWTER the importance* occurred in training, and *POWTER the significance* occurred in test). Each object noun appeared only once in the critical trials. Assignment of items to conditions was rotated around two presentation lists so as to control for item effects. Block 2 also contained an additional 44 grammatical sentences so as to provide more reinforcement for the rule. These were sentences repeated from Block 1. The order of trials in Blocks 1 and 2 was independently randomized for each subject.

The sentences were presented word by word in the centre of the screen. First the noun, e.g. POWTER, was presented in capital letters for 600 msecs followed by the article *the* for 600 msecs. This was followed by a noun, e.g. *prestige*, in red lower case letters. The participants were instructed to indicate as quickly as possible whether this noun referred to an abstract or a concrete object by pressing the left or right buttons on a response box. If they made an error the word remained on the screen until they pressed the correct button. Upon a correct response the display changed to a recall prompt '___ the ___' and they had to recall the noun phrase aloud, i.e. say "powter the prestige". They then pressed a response button and the remainder of the sentence appeared at a rate of 600 msecs per word, e.g. *of wealthy families artists can*. At the end of the sentence the prompt +/- appeared on the screen and the participants had to indicate as quickly as possible whether they thought the verb conveyed a broadly increase or decrease meaning by pressing the right or left buttons on the response box.

In order to assess awareness, at the end of the experiment the participants were presented with 8 new sentences from which the verb had been removed. They were asked to indicate which of the four novel verbs they thought should be used in that context and to think aloud as they made their decision. Any participant who referred to the abstract/concrete distinction, or similar, was classified as 'aware', regardless of their actual performance on the post-test.

The experiment was run using Superlab software and a Cedrus response box.

Results and Discussion

Out of 40 participants, 13 revealed at least fragmentary explicit knowledge in the post-experiment debriefing, and were classified as 'aware'. The remaining 27 participants were classified as 'unaware'. For each participant, response times that were more than 2.5 standard deviations above the mean response time over the 32 critical trials were winsorized (i.e. replaced with the next highest value). Additionally, in cases where an error was made on the first (abstract/concrete) decision, the response on the second decision was removed from the analysis. This was because participants were likely to have been distracted on the subsequent increase/decrease decision by just having had to correct themselves.

An initial analysis of the data revealed that for both the aware and unaware groups there were no differences in either reaction time or error rate between the NG and NU conditions on the first, concrete/abstract, decision. This was rather surprising because we expected that if the regularity had been learned the verb would set up an expectation of a certain type of collocate. However, a significant NU-NG reaction time difference was obtained on the second, increase/decrease, decision. After excluding two unaware participants with excessively long response times (of 3231 msecs and 3359 msecs) the remaining 25 unaware participants had a mean response time of 1101 msecs in the NG condition and 1332 msecs in the NU condition, $F(1,23) = 5.11$, $p < 0.05$, $\eta2 = 0.20$. There was no difference in error rates (25.2% and 23.4% respectively). As for the aware participants, despite showing a large numerical difference between the NG and NU conditions on the increase/decrease decision (1478 msecs and 1611 msecs respectively) this difference was not significant, $F(1,11) = 1.70$, $p = 0.23$. Neither was there an effect in errors (22.5% and 23.9 % respectively).

A post hoc analysis was carried out to check for potential differences in performance between the native and non-native English-speaking participants. An ANOVA revealed no interaction between grammaticality and native/non-native speaker status, $F < 1.0$.

This experiment provides evidence of implicit learning of a semantic preference rule. The fact that learning effects are apparent on the decisions involving purely the indication of whether a

particular verb meant to 'increase' or 'decrease' is particularly compelling, since the decision was being made with reference to the meaning of the verb, not the collocate, and knowledge of the semantic preference rule does not directly inform this decision. This effect actually provides stronger evidence of the use of implicit knowledge than if it had occurred on the concreteness decision (a point which we will elaborate below).

Having said this, it is not clear how the effect of grammaticality on the increase/decrease decision arises. One possibility is that the mismatch between the verb and noun in the NU condition somehow disturbed the process of deriving the increase/decrease meaning from the verb and its context. It may also have caused confusion about the identity of the verb (since the collocate would have suggested other verb possibilities than the one that occurred). However, there is an alternative explanation that cannot be ruled out at this stage. The effect may not reflect learning of a semantic regularity at all, but rather associations between novel verbs and patterns of button presses (e.g. POWTER was associated with successive responses on the right-hand button). Experiment 2 was designed to rule out this possibility, as well as creating conditions under which awareness of the hidden regularity was much less likely to occur, and under which any effect of knowledge on behavior was more likely to reflect automatic, as opposed to controlled, behavior.

Even if we suppose that the effect obtained in Experiment 1 was semantic in origin, the question remains as to the nature of the generalization that was formed. Although the noun collocates in the critical test items were different from those that occurred in training they were in fact roughly synonymous with a noun that had occurred in the training phase. This means that it is hard to defend the claim that what was learned was a correlation between verbs and the abstract/concrete distinction as such. Rather the effect could have reflected the similarity between individual nouns in training and test. In order to address this issue, the noun collocates in Experiment 2 were changed so as to represent a more heterogeneous set of abstract and concrete nouns, and no noun in the test phase was a synonym of a noun in the training phase. Learning over these items would be more likely to reflect abstraction of a broad concreteness distinction.

Experiment 2

This experiment employed a reaction time methodology similar to Experiment 1. Two main changes were made. First, a simplified procedure was employed in which only verb phrases were presented. Participants were informed about the increase/decrease meanings of the verbs prior to the experiment. The second change was that instead of making a concreteness decision on the collocates participants now had to indicate whether the collocate had positive, negative, or neutral connotation. For example, *chocolate* and *holidays* would be expected to receive positive judgments, whereas *horror* would be expected to receive a negative judgment. Participants were informed that the choices were subjective and that there was no correct answer. Crucially, the semantic preference rule was exactly the same as in Experiment 1; *powter* and *gouble* went with abstract nouns and *conell* and *mouten* with concrete nouns. Given that no systematic alignment between connotative meaning judgments and concreteness is expected this means that there will be no systematic relationship between the button pressed on the first decision and the second increase/decrease decision. Thus, learning is unlikely to be based on associations between nonsense verbs and patterns of button presses. This also means that any influence of noun concreteness on the second decision must reflect automatic activation of this aspect of meaning, rather than explicit retrieval as part of the task (as was the case in Experiment 1).

Participants

46 students of the University of Cambridge participated in the experiment. Three participants were excluded due to problems with the task. Of the remaining group of 43 participants, 22 were native speakers of English. All of the nonnative participants had achieved at least IELTS 7.5.

Materials and Procedure

The experimental design was identical to Experiment 1. The nouns for the training and test items were changed so as to comprise a more heterogeneous set of abstract and concrete nouns. The broadened category of abstract nouns included ones as different as *happiness, wisdom, impact, understanding*. The category of concrete nouns was similarly broadened to include, for example, *chocolate, luggage, metal and paper*. Only verb phrases were presented.

Participants were first informed about the increase/decrease meanings of the novel verbs. The

simplified procedure in Experiment 2 comprised the following sequence of events. First the verb was presented in capital letters for 600 msecs. This was followed by the word *the* for 600 msecs, followed by the noun in red. The participants made their decision about the connotation of the noun within an allocated time of 3 secs. Responses were entered on a millisecond accurate keyboard where M indicated 'positive', Z 'negative', and space bar 'neutral'. Upon making any response the noun was replaced by the 'inc/dec' prompt and the participant indicated whether the verb meant increase or decrease by pressing M and Z respectively. After every two stimuli participants received prompts that required them to recall out loud one of the phrases they had just seen. The prompt revealed either the first part of the phrase, for example "MOUTEN the _____", or the second part: "_____ the prestige", and participants were asked to pronounce the complete phrase. The memory task was to encourage full attention to the materials and the data were not analysed. All stimuli were presented in the centre of the screen. The experiment was followed by a post-test similar to that in Experiment 1.

Results and Discussion

None of the participants appeared to have any awareness of the correlation between the novel verbs and the concreteness of the noun (whereas in Experiment 1 32.5% of the participants were classed as aware). The data were treated in the same way as in Experiment 1. Two of the native participants were excluded on the basis of excessively long second decision response times. As before there were no reaction time effects on the first decision. This time there were no effects on the second, increase/decrease, decision either. Response times in the NG and NU conditions were 616 msecs and 627 msecs respectively, and error rates were 6.1% and 6.7% (note that reaction times were much faster than in Experiment 1, presumably because the decision was made immediately after the noun, rather than being delayed until the end of a sentence). When the data for the native English-speaking and non-native groups were compared an interesting pattern emerged. For the natives, reaction times in the NG and NU conditions were 556 msecs and 612 msecs, with error rates of 7.2% and 5.9%. In contrast for the non-natives the reaction times in the NG and NU conditions were 690 msecs and 655 msecs, with error rates of 5.1% and 7.4%. An ANOVA was performed on the reaction time data with Group (native or non-native) and Presentation List as between-subjects factors, and Condition (NG vs NU) as a within-subjects factor. The interaction between Group and

Condition was significant, $F(1,37) = 8.74$, $p < 0.01$, $\eta2 = 0.19$. Follow-up ANOVAs showed that the learning effect was only significant for the native speaker group $F(1,17) = 9.13$, $p < 0.01$, $\eta2 = 0.34$. There were no significant effects on errors.

Experiment 2 provides stronger evidence for semantic implicit learning than Experiment 1 because the learned generalisation was unrelated to the tasks being performed. Knowledge of the correlation between verbs and noun concreteness was unrelated both to the decision being made on the noun (connotative meaning) and to the increase/decrease decision. It has been argued that implicit knowledge should exert its influence on behaviour in an automatic rather than controlled way (Cleeremans & Jimenez, 2002), and that the strongest test of implicit knowledge is to be obtained in situations where knowledge has an effect on performance even though it is irrelevant to the task at hand (Tzelgov, 1997; Vinter & Perruchet, 1999). The present experiment seems to satisfy those criteria. Furthermore, on this occasion none of the participants demonstrated awareness of the semantic regularity in the post-test, which in itself suggests that the relevant knowledge was well below the level of awareness.

Experiment 2 also shows that the effects obtained in Experiment 1 were not due to learning associations between the novel verbs and patterns of keystrokes. This was because the nouns in the abstract and concrete categories would have elicited a range of 'positive' (M), 'negative' (Z) and 'neutral' (space bar) responses. Thus the effects must have had a semantic origin. Furthermore the learning effect in natives was obtained over sets of nouns that were more heterogeneous than in Experiment 1, and the critical test nouns were not synonyms of any nouns in training. Therefore the learning effect must have been supported by a broad generalisation, which we assume is essentially based on the abstract versus concrete distinction.

The fact that in this experiment learning was not obtained for non-native speakers is perhaps not surprising. There is a wealth of evidence that the mapping from second language words to meaning is less automatic than from first language words. For example, automatic semantic priming from second language words can only be obtained at very high levels of proficiency (Basnight-Brown & Altarriba, 2007). Thus, in a situation in which a decision is made about one aspect of meaning it is not surprising that the non-natives in Experiment 2 did not activate other aspects of meaning with sufficient strength to produce a learning effect.

Conclusion

The present experiments demonstrate that semantic implicit learning of linguistic regularities can be obtained outside of the realm of grammar, and can extend into learning about verb meaning. However, this is not to say that all aspects of word meaning have the potential to be learned implicitly. Paradis (2004) has proposed that learning referential meaning requires the declarative system, and hence presumably depends upon awareness of the connection between form and meaning. But other aspects of meaning, such as semantic preferences, could be acquired implicitly. Although it is not clear whether the kind of learning demonstrated here actually depends on the operation of the procedural system, as hypothesised by Paradis, the present results do show that this aspect of verb meaning is amenable to implicit learning.

It should also be stressed that the present experiments demonstrate learning semantic preferences in a situation in which some aspect of the meaning of the verbs (i.e., their increase or decrease meaning) is already explicitly known (as in Experiment 2) or being intentionally inferred from context (Experiment 1). Thus, we regard these learning effects as essentially reflecting the process of 'tuning' an already-established meaning of which the participants are aware. Whilst this tuning process undoubtedly forms an important part of word learning through usage, it has to be distinguished from the process of actually forming new form-meaning connections from scratch.

Finally, the fact that there was no semantic implicit learning effect in Experiment 2 for non-native speakers provides a cautionary note in relation to the role this process may play in second language acquisition. The implication is that semantic regularities will be most effectively learned when attention is drawn to the relevant aspects of meaning by the task. Otherwise semantic implicit learning effects may be limited unless semantic access from known words in the context is highly automatic. This does not mean, though, that learners have to be aware of the actual underlying regularities. Even if in some cases it may be necessary to be aware of both form and meaning, it is not necessary to be aware of the semantic generalisations that they license.

References

Basnight-Brown, D. M., & Altarriba, J. (2007). Differences in semantic and translation priming across languages. *Memory & Cognition, 35*, 953-965.

Bickerton, D. (1999). Creole languages, the language bioprogram hypothesis, and language acquisition. In W. C. Ritchie & T. K. Bhatia (Eds.), *Handbook of child language acquisition*: Academic Press.

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

Cleeremans, A., & Jiménez, L. (2002). Implicit learning and consciousness: A graded, dynamic perspective. In R. M. French & A. Cleeremans (Eds.), *Implicit Learning and Consciousness* (pp. 1-40). Hove: Psychology Press.

Dell, G. S., Reed, K. D., Adams, D. R., & Meyer, A. S. (2000). Speech errors, phonotactic constraints, and implicit learning. *Journal of Experimental Psychology-Learning Memory and Cognition, 26*, 1355-1367.

Ellis, N. C. (1994). Vocabulary acquisition: The explicit ins and outs of explicit cognitive mediation. In N. C. Ellis (Ed.), *Implicit and Explicit Learning of Languages* (pp. 211-282). London: Academic Press.

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

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

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

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.

Tzelgov, J. (1997). Specifying the relations between automaticity and consciousness: A theoretical note. *Consciousness and Cognition, 6*, 441-451.

Vinter, A., & Perruchet, P. (1999). Isolating Unconscious Influences: The Neutral Parameter Procedure. *Quarterly Journal of Experimental Psychology: Section A, 52*, 857-875.

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