

Prior Linguistic Knowledge Influences Implicit Language Learning

Janny H.C. Leung (jannyleung@hku.hk)

School of English, The University of Hong Kong,
735 Run Run Shaw Tower, Centennial Campus, The University of Hong Kong, Hong Kong

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

We report three experiments that explore the effect of prior linguistic knowledge on implicit language learning. Native speakers of English and native speakers of Cantonese participated in implicit learning (IL) experiments that involved different learning materials. In Experiment 1, both participant groups showed evidence of learning a mapping between articles and noun animacy. In Experiment 2, neither group showed learning of the mapping between articles and a linguistically anomalous concept (the number of capital letters in an English word or the number of strokes in a Chinese character). In Experiment 3, the Chinese group, but not the English group, showed evidence of learning a mapping between articles and a concept derived from the Chinese classifier system. It was concluded that first language knowledge affected implicit language learning, and that IL, at least when natural language learning is concerned, is not a completely unconstrained domain-general mechanism.

Keywords: implicit learning; form-meaning connections; vocabulary learning; second language acquisition; cross-linguistic influence

Introduction

Traditional implicit learning (IL) (as contrasted with explicit learning, EL) research has sought to minimise the effect of prior knowledge, either by using artificial grammars (e.g., Reber, 1967) or artificial event sequences (e.g., serial reaction time experiments; e.g., Nissen & Bullemer, 1987). Perruchet and Pacton (2006: 237) suggest that even arbitrary materials may not be “neutral” enough, for they may interact with related situational knowledge. It is typically assumed that domain general mechanisms underlie IL, but in many real life situations the learner may bring relevant domain-specific prior knowledge or dispositions, so the question arises as to whether, or how, these impact upon the IL process.

Natural language learning is a case in point – some believe that linguistic universals constrain both first and second language acquisition (Chomsky, 2006; Hawkins, 2004) and that in both cases the theoretically interesting learning processes operate at the implicit level. Even those who dispute the nature of such linguistic universals would accept that second language acquisition (SLA) is heavily influenced by first language (L1) knowledge, or L1-based processing strategies (MacWhinney, 2008; Ellis & Sagarra, 2011). Second language (L2) learners approach SLA with existing linguistic knowledge and habits they have gathered from their first language acquisition experience. Cross-linguistic influence is well documented in the SLA literature, much of which is concerned with identifying the ways in which elements from one language get incorporated into another, accounting for errors, contrastive analysis, and interaction of transfer effects with other factors. Ellis (1994/2001: 300) argues that no theory of SLA “can be considered complete” if it ignores the learner’s prior linguistic knowledge. In a similar vein, if IL is posited as an underlying mechanism of language acquisition, one must also consider whether and how the influences of prior linguistic experience on learning take place implicitly. In the SLA literature, cross-linguistic influences are sometimes thought of in terms of hypothesis testing and learner strategies (Corder, 1981; Tomasello & Herron, 1989), implying a certain degree of intention and awareness in the process. Although it is difficult to imagine that such influences involve only explicit processes, there does not seem to be empirical effort to demonstrate such influences operating at the implicit level during learning. Moreover, cross-linguistic influence is found to be subject to general constraints such as language proficiency, sociolinguistic factors, markedness, prototypicality, language distance and psychotypology, and developmental factors (Ellis, 1994/2001). The interaction of such constraints with domain general learning mechanisms begs for research.

Our earlier work has begun to show that learning processes supporting implicit language learning

effects are not completely unconstrained. Leung and Williams (2012) used an efficiency measure (Seger 1994) to measure learning of a semi-artificial article system comprising the pseudowords *gi*, *ul*, *ro*, *ne*. Participants were told that the use of the articles depended on the distance of the object described by the accompanying noun (*gi* and *ul* were used with an accompanying noun which referred to a near object; *ro* and *ne* were used if it referred to a far object). Unbeknownst to them, article use also depended on noun animacy (*gi* and *ro* were used with animate objects and *ul* and *ne* with inanimate objects). Objects were pictorially presented, along with an audio presentation of the corresponding noun phrase. The task for the participants was to respond as quickly as possible whether the noun referred to a living or a non-living thing¹. Learning was measured by an increase in reaction time when the hidden regularity was violated (Violation trials) compared with reaction time in grammatical trials (Control trials). After performing 272 training trials in the task, native speakers of English who claimed to be unaware of the hidden regularity (as revealed in a standardised verbal report) nevertheless slowed down significantly in animacy decisions when the correlation between article use and noun animacy was reversed. In contrast, in a second experiment article use correlated with the relative size of objects rather than their animacy. This time there was no significant change in relative size decision times when the mapping between article use and relative size of the objects described by the nouns was reversed. The findings suggest that some meanings are more amenable to the IL of form-meaning connections than others. One possible explanation for this is the availability of grammatical processes and representations based on participants' existing linguistic knowledge, but to probe into this issue, comparisons between participants with different first languages have to be made.

The present study aims to consolidate and extend our earlier work by further exploring potential first language influences on IL of a semi-artificial grammatical system. We report below three experiments involving different learning materials and two language groups (native Chinese and native English speakers).

¹ Although the task drew attention to meaning, it did not draw attention to the form-meaning connection which was the target of learning.

Experiment 1

Objective

To test whether animacy, a conceptually salient feature, may be implicitly mapped onto articles by native speakers of English and Chinese.

Participants

Thirty native speakers of English from the University of Cambridge and 27 native speakers of Cantonese Chinese from the University of Hong Kong, with a varied second language background. All Hong Kong participants spoke English as an L2, and some also knew other second/foreign languages.

Materials and Procedure

All experimental materials were digitalized and presented with E-Prime software.

Participants were told that they would be introduced to a miniature article system from a language not known to them (Table 1). They were told that the articles² were used to encode the distance between the speaker and the object (*gi* and *ro* for near objects and *ul* and *ne* for far objects). Therefore *gi dog* may be read as 'the near dog', *ro table* as 'the near table', *ul mouse* as 'the far mouse', and *ne car* as 'the far car'. Participants could spend as much time as they needed to remember the mapping between the articles and the distance system. Participants were however not told that the use of these articles also depended on the animacy of the accompanying noun (*gi* and *ul* for animate objects and *ro* and *ne* for inanimate objects).

Table 1 The miniature article system in Experiment 1

Miniature article system	Participants not told		
	animate	inanimate	
Participants were told	near	<i>gi</i>	<i>ro</i>
	far	<i>ul</i>	<i>ne</i>

A total of 176 animate and inanimate nouns were used for the experiment, each appearing twice (once with each possible article). In the training phase, participants were exposed to visually presented noun phrases (article and noun combinations). While the same article system was used for all participants, the nouns were presented in each participant group's first

² Since the Chinese language has no article system, Chinese participants were simply told that these were words that were used before the noun. The Chinese participants were however familiar with the concept of articles from their L2.

language (i.e., English or Chinese; see table 2). For instance, an English subject may see *gi fox* and a Chinese subject may see *gi 狐狸* ('fox') on the screen. The task for the participants was first to make a decision about the animacy of the object (living or non-living, M or C key) described by the noun. The noun phrase then disappeared and was replaced by the prompt "N/F". Participants had to indicate the distance meaning of the article (M = near, K = far). Both decisions had to be made as quickly as possible and reaction time was recorded. Response buttons were configured such that the near/far buttons were in a logical arrangement - one above the other, which also helped reduce interference with the other decision which is in a horizontal arrangement; see figure 1 for visualisation. Reaction time for the animacy decision was measured from the onset of the noun phrase; reaction time for the near/far decision was measured from the onset of the N/F (near/far) prompt, which appeared immediately after the animacy response. Feedback was provided; if participants gave a wrong response, the display did not change. Eight practice trials were provided before the experiment started. A total of 204 grammatical trials were presented in the training phase, with equal numbers of trials presented with each article. These trials were divided into four blocks, although no division between blocks was apparent to the participants. The first block consisted of 84 trials, which were made up of 84 nouns used in combination with an equal number of one of the two grammatically possible articles for each noun. Block 2 contained 60 trials, comprising 60 new nouns, again in combination with an equal number of appropriate articles. In Block 3, 28 of these nouns were repeated with a correct article of opposite distance from Block 2 (e.g. if *gi pig* had occurred in Block 2, then *ul pig* occurred in Block 3). Within Blocks 1 to 3 the trials were divided into fixed groups of four, with each article occurring once. For each participant the order of trials within groups was randomised as was the order of groups. This procedure meant that no more than two successive trials would involve the same article. In Block 4 the remaining 32 nouns from Block 2 were repeated. Half of them occurred with an article of different distance from Block 2, but correct animacy (e.g., if *ul parrot* had occurred in Block 2 then *gi parrot* occurred in Block 4). These were Control trials. The other half of the Block 4 trials occurred with the article of opposite distance and animacy (e.g. if *ro tent* had occurred in Block 2, then *ul tent* occurred in Block 4). These were Violation trials. Control and Violation trials were randomly intermixed, and the nouns were rotated around conditions across subjects,

resulting in two presentation lists. Note that although nouns were repeated from Block 2 to Blocks 3 and 4, no article-noun combination was repeated throughout the experiment. After Block 4 participants filled out a questionnaire which probed awareness of the relevance of animacy to article usage, awareness of violations, and at what point in the experiment awareness of the regularity developed. Participants who reported no awareness were then encouraged to guess what factors (apart from near/far) determined when the articles were used. The whole experiment took about 45 minutes to complete.

Table 2. Sample animate and inanimate nouns used in Experiment 1

gi/ul		ro/ne	
English version	Chinese version	English version	Chinese version
<i>fox</i>	狐狸	<i>piano</i>	鋼琴
<i>buffalo</i>	水牛	<i>microscope</i>	顯微鏡
<i>gorilla</i>	大猩猩	<i>telescope</i>	望遠鏡
<i>seal</i>	海豹	<i>kettle</i>	水壺

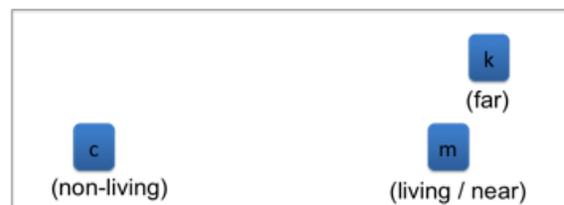


Figure 1. Configuration of response buttons on millisecond accurate keyboard/response box.

Results

In the post-experiment debriefing, two levels of unawareness were assessed: a failure to report knowledge of the hidden regularity, and a failure to suggest any relevance of the target concept to article use when prompted to guess. Data from participants who failed to report knowledge are reported below; data from participants who failed to guess were found to show a pattern consistent with the summarized findings and are not reported here.

Outlying response times, at cutoff limits at ± 2.5 standard deviations from each subject's mean in the Control and Violation blocks respectively were removed. In addition, data were excluded from any participant whose mean response time for the first decision over the two critical conditions was more than 2.5 standard deviations from the group mean.

Considerable variability in response times was observed, possibly because participants varied in how they distributed processing time over the two decisions and so inordinately slow participants were excluded on the basis that they may have been approaching the task in a different way from the majority.

Evidence of learning is based on the difference in reaction time across Control and Violation trials, as tested by a paired-sample t-test. Only data for the first (in this case animacy) decision are reported (no effects in any experiment were obtained for the second, near/far, decision).

Twenty native English and 20 native Chinese did not report relevant knowledge of the regularity in response to the first debriefing question (awareness rates 33% and 26% respectively). One Chinese participant was excluded due to slow overall animacy decision response times. Response times to make animacy decisions were significantly slower for Violation than Control trials even for participants who showed no awareness of the relevance of animacy, or related concepts, to article usage (see fig. 2, ‘English animacy’ and ‘Chinese animacy’), indicating that the animacy rule has been learned implicitly. The combination of alphabetical articles with characters is orthographically odd for the Chinese, but this did not seem to have affected their learning.

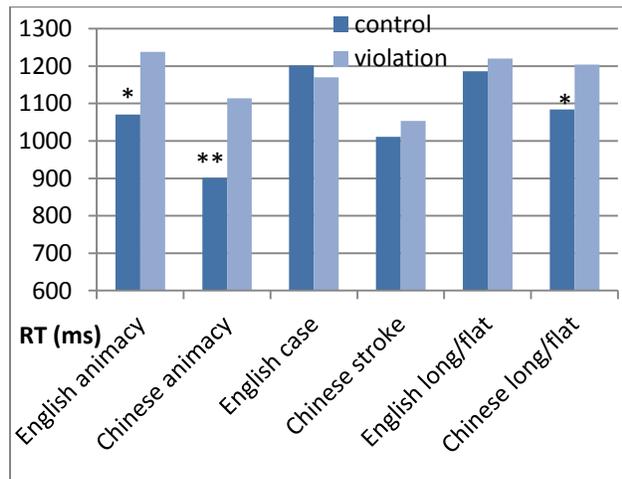


Figure 2. Reaction times (RT) in milliseconds (ms) for unaware participants in all three experiments, * $p < 0.05$, ** $p < 0.01$ for difference between control and violation conditions

The question that follows is whether anyone with any language background will learn any regularity. A

linguistically unnatural learning target is adopted in the following experiment to investigate this.

Experiment 2

Objective

To test whether a linguistically anomalous, form-based, distinction may be implicitly mapped onto articles by native speakers of English and Chinese.

Participants

Twenty three native speakers of English from the University of Cambridge and 29 native speakers of Chinese from the University of Hong Kong, with a varied second language background, participated in the study.

Materials and Procedure

The experiment shares a similar design with Experiment 1. The animacy system was now replaced by a non-semantic distinction for each participant group. In the English version, the hidden regularity was that the choice of article depended on whether the word had one capital letter or two. In the Chinese version, the hidden regularity was whether the first character in a two-character noun has more strokes than the second and vice versa, with a strokes difference between the characters being big enough that no counting would be necessary (see Table 3 below).

Table 3. Sample nouns used in Experiment 2

gi/ro		ul/ne	
English version (case)	Chinese version (stroke)	English version (case)	Chinese version (stroke)
foX	天鵝	goRilla	獅子
piAno	月餅	TelesCope	剪刀
buffaLo	牙醫	sEA1	學生
miCroscope	天橋	kEttLE	瀑布

The design and procedure were the same as Experiment 1 except that here participants had to indicate, as their first decision, whether the noun contained one or two capital letters (M and Z keys respectively) or whether there were more strokes in the second or first character (M and Z keys respectively). The second decision indicated whether the article meant ‘near’ (M) or ‘far’ (K), as before.

Results

Twenty native English and 28 native Chinese did not report relevant knowledge of the regularity in response to the first debriefing question (awareness rates 13% and 3% respectively). One slow English unaware participant was excluded.

In both language groups, the reaction times across Control and Violation trials for unaware participants were not significantly different (see fig. 2), indicating that implicit language learning in this domain does not occur for a linguistically unnatural, and non-semantic, distinction. The next experiment examines whether learning occurs for any linguistically natural semantic distinction, regardless of whether it is reflected in the first language.

Experiment 3

Objective

To test whether a concept derived from the Chinese classifier system, which is not grammaticised in English, may be implicitly mapped onto articles by native speakers of English and Chinese.

Participants

Twenty seven native speakers of English from the University of Cambridge and 32 native speakers of Chinese from the University of Hong Kong, with a varied second language background, participated in the study.

Materials and Procedure

The learning target in this experiment is derived from a shape distinction in the Chinese classifier system. The Chinese classifier 張 (*zoeng1* in Cantonese) is generally used with thin flat objects (e.g., sheets of paper, photos, blankets), and the counter 條 (*tiu4* in Cantonese) is generally used with long thin objects (e.g., rivers, straws, ties)³. Both classifiers frequently occur in daily usage.

The same design as the above experiments was adopted, except that the long/flat distinction became the hidden regularity governing article use in this experiment. The same article system was used with the two participant groups; nouns were presented in the first language of the participants. For example, one may find items such as *gi shoelace* and *ro envelope* in the English version, and their equivalents in the Chinese version (e.g., *gi 鞋帶* and *ro 信封*).

³ In both cases exceptions exist. For instance, *zoeng1* is also associated with furniture items. Such exceptional items were not included in the experiment.

The design and procedure were the same as in previous experiments except that here participants had to indicate, as their first decision, whether the noun referred to an object that was long (M key) or flat (Z key). The second decision indicated whether the article meant 'near' (M) or 'far' (K), as before.

Results

Twenty four native English and 26 native Chinese did not report relevant knowledge of the regularity in response to the first debriefing question (awareness rates 11% and 19% respectively). One slow English and two slow Chinese unaware participants were excluded. A significant slowdown in Violation trials when compared with the Control trials was obtained among the unaware Chinese participants but not the unaware English participants (see fig. 2), suggesting that implicit language learning is sensitive to prior linguistic knowledge.

Discussion

One might imagine that in this paradigm participants merely learn associations between articles and patterns of keystrokes, e.g., that 'gi' is associated with the sequence M (living) – M (near), or 'ne' with Z (non-living) – K (far). Control trials respect these patterns, but Violation trials break them, e.g. 'gi' would occur with the sequence Z (non-living) – M (near). However, if this were the case, then the nature of the categorization being performed on the noun should have made no difference whatsoever. The fact that it did suggests that the learning effects were due to learning associations between the articles and conceptual categories.

It is important to note that in all of the experiments the relevant 'hidden' conceptual distinction had to be attended and computed in order to perform the task. But this did not guarantee that the association with the articles would be learned. Some equate statistical learning with IL (e.g., Conway & Christiansen, 2006), but statistical computations should not be sensitive to the nature of the data (Perruchet & Pacton, 2006).

The finding that IL effects are sensitive to the nature of the concepts involved points to an interaction between the domain general learning mechanism and linguistic knowledge, which, according to many linguists, is domain specific. Semantic IL in natural language is constrained by the availability of conceptual distinctions to grammaticisation, which varies cross-linguistically.

Where no evidence of learning was obtained, it is possible that measurable learning would develop over

time. This is to say that linguistically unnatural or unfamiliar semantic categories may not be unlearnable. However, the variable amount of learning obtained after equivalent exposure shows that implicit language learning is sensitive to prior linguistic knowledge. The present study thus extends the existing literature and is congruent with studies which show that cross cultural processing biases also apply to unconscious knowledge (such as a preference for local versus global perspectives in Kiyokawa et al., 2012), and can help explain SLA studies which find L1 semantic structures in L2 processing (e.g., Jiang 2004).

We provide evidence that cross-linguistic influences may take place implicitly, and caution against a presumption that L1 transfer is based on hypothesis testing or learner strategy. It remains unclear to what degree such influences take place implicitly, or explicitly, or both implicitly and explicitly, in different SLA settings, and it seems likely that individual differences exist. A better understanding of the mechanism underlying cross-linguistic influences has obvious theoretical and pedagogical implications. Theoretically it sheds light on debates in language acquisition on domain specificity and linguistic universals; pedagogically it informs teaching/learning methodologies that aim to promote or discourage different kinds of influences.

But the kind of cross-linguistic influence that we have demonstrated may have gone beyond simple transfer. In our experiments, English participants were only sensitive to animacy – a fundamental conceptual distinction, even though it is only subtly marked in English, and there is no article-noun agreement. Chinese participants were sensitive to all semantic (but not non-semantic) distinctions, presumably through experience of their classifier system, but generalizing to novel distinctions.

Although many assume that first language acquisition is essentially implicit, it is only recently that research has shown possibilities of adult SLA taking place implicitly. Apart from further exploring the mechanisms of crosslinguistic influence, which are pertinent to SLA research, future research into implicit language learning may also inform IL research as to its nature and constraints.

Acknowledgments

This project (code RES-000-22-3030) has been generously supported by a joint research scheme under Research Grants Council, Hong Kong, and Economic and Social Research Council, United Kingdom.

References

- Conway, C. M. and Christiansen, M. H. (2006) Statistical Learning Within and Between Modalities: Pitting Abstract Against Stimulus-Specific Representations. *Psychological Science*, 17:10, 905-912.
- Chomsky, N. (2006) *Language and Mind*. Cambridge University Press.
- Corder, P. (1981) *Error analysis and interlanguage*. Oxford: Oxford University Press.
- Ellis, N. C. & Sagarra, N. (2011) Learned attention in adult language acquisition: A replication and generalization study and meta-analysis. *Studies in Second Language Acquisition*, 33: 4, 589-624.
- Ellis, R. (1994/2001) *The Study of Second Language Acquisition*. Oxford, U.K.: Oxford University Press.
- Hawkins, John. A. (2004) *Efficiency and complexity in grammars*. Oxford, U.K.: Oxford University Press.
- Jiang, N. Semantic Transfer and Its Implications for Vocabulary Teaching in a Second Language. *The Modern Language Journal*, 88:3, 416-432.
- Kiyokawa, S., Dienes, Z., Tanaka, D., Yamada, A. and Crowe, L. (2012) Cross cultural differences in unconscious knowledge. *Cognition*, 124: 16-24.
- Leung, J. H. C. and Williams, J. N. (2012) Constraints on Implicit Learning of Grammatical Form-Meaning Connections. *Language Learning*, 62: 2, 634-662.
- MacWhinney, B. (2008) A Unified Model. In N. Ellis & P. Robinson (Eds.) *Handbook of cognitive linguistics and second language acquisition*. Lawrence Erlbaum Press.
- Nissen, M.J., & Bullemer, P. (1987) Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19:1, 1-32.
- Perruchet, P. and Pacton, S. (2006) Implicit learning and statistical learning: one phenomenon, two approaches. *Trends in Cognitive Science*, 10:5, 233-238.
- Reber, A. S. (1967). Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6, 855-863.
- Tomasello, M. and Herron, C. (1989) Feedback for Language Transfer Errors. *Studies in Second Language Acquisition*, 11, 386-395.
- Williams, J. N. (2005). Learning without awareness. *Studies in Second Language Acquisition*, 27, 269-304.