Seth, A. K.
2008 Post-decision wagering measures metacognitive content, not sensory consciousness. *Consciousness and Cognition*, 17, 981–983.

Seth, A., Dienes, Z., Cleeremans, A., Overgaard, M., & Pessoa, L.
2008 Measuring Consciousness: Relating Behavioural and Neurophysiological Approaches. *Trends in Cognitive Sciences*, 12, 314–321.

Shanks, D. R.
2005 Implicit learning. In K. Lamberts & R. Goldstone (Eds.), *Handbook of cognition* (pp. 202–220). London: Sage.

Shanks, D. R., & St. John, M. F.
1994 Characteristics of dissociable human learning systems. *Behavioral & Brain Sciences*, 17, 367–447.

Smoke, K. L.
1932 An objective study of concept formation. *Psychological Monographs*, 42, whole number 191.

Sun, R.
2002 *Duality of mind: A bottom-up approach toward cognition.* Erlbaum.

Tanaka, D., Kiyokawa, S., Yamada, A., Dienes, Z., Shigemasu, K.
2008 Role of Selective Attention in Artificial Grammar Learning. *Psychonomic Bulletin and Review*, 15, 1154–1159.

Tunney, R. J.
2007 The subjective experience of remembering in artificial grammar learning. *European Journal of Cognitive Psychology*, 19, 934–952.

Wan, L. L., Dienes, Z., & Fu, X. L.
2008 Intentional control based on familiarity in artificial grammar learning. *Consciousness and Cognition*, 17, 1209–1218.

Williams, J. N.
2004 Implicit learning of form-meaning connections. In J. Williams, B. VanPatten, S. Rott & M. Overstreet (Eds.), *Form Meaning Connections in Second Language Acquisition* (pp. 203–218). Mahwah, NJ: Lawrence Erlbaum Associates.

Williams, J. N.
2005 Learning without awareness. *SALL*, 27, 269–304.

Ziori, E., & Dienes, Z.
2006 Subjective measures of unconscious knowledge of concepts. *Mind & Society*, 5, 105–122.

Ziori, E. & Dienes, Z.
2008 How does prior knowledge affect implicit and explicit concept learning? *Quarterly Journal of Experimental Psychology*, 61, 601–624.

How implicit is statistical learning?

Phillip Hamrick and Patrick Rebuschat

Introduction

How learners extract knowledge from the environment is one of the fundamental questions in cognitive science. In this chapter, we will focus on two approaches that have gained in prominence over the past 15–20 years, namely *implicit learning* and *statistical learning* (see also Dienes, this volume; Misyak, Goldstein, & Christiansen, this volume). Implicit learning research began with Reber's (1967) early work and developed into one of the major paradigms in cognitive psychology (see Perruchet, 2008, for an overview). Statistical learning research was sparked by the work of Saffran and colleagues (Saffran, Aslin, & Newport, 1996) and now represents an important research strand in developmental psychology (see Gómez, 2007, for an overview).

Both approaches focus on how we acquire information from the environment and both rely heavily on the use of artificial grammars. In typical experiments, subjects are first exposed to stimuli generated by an artificial system and then tested to determine what they have learned. Given these and other similarities, Perruchet & Pacton (2006) suggested that implicit and statistical learning represent two approaches to a single phenomenon. Conway & Christiansen (2006) go as far as combining the two in name: *implicit statistical learning*.

Despite the considerable overlap between implicit and statistical learning research, there are several important differences. For example, one of the most distinctive features of statistical learning research is the careful manipulation of statistical information in the input. This aspect is generally absent in implicit learning studies. In addition, statistical learning research generally concentrates on how we acquire linguistic information, while implicit learning research focuses on information in general.[1] For this reason, statistical learning researchers tend to employ artificial systems that resemble

1. As one of our reviewers pointed out, there are also several studies on statistical learning in different modalities and with non-linguistic stimuli.

natural languages more closely (phrase-structure grammars instead of finite-state systems, and pseudowords instead of letter sequences).

Another important difference is that implicit learning researchers are generally concerned with the question of whether subjects acquire conscious (explicit) or unconscious (implicit) knowledge as a result of exposure. For this purpose, implicit learning studies usually contain measures of awareness.[2] This holds across implicit learning paradigms (i.e., artificial grammar learning, sequence learning, control of complex systems). In contrast, statistical learning studies do not typically feature any measures of awareness as part of the experimental design. This is, of course, partially explained by the fact that infants are unable to provide verbal reports, indicate confidence levels, or perform on fragment completion tasks. However, many of the experiments conducted within the statistical learning framework employ adults as subjects, which means that basic measures of awareness could be administered. Usually, lack of awareness is assumed but not empirically assessed (see Aslin, Saffran, & Newport, 1999, though see e.g., Dell, Reed, Adams, & Meyer, 2000; Warker & Dell, 2006). As such, it is unclear whether statistical learning typically results in conscious or unconscious knowledge. Given that language comprehension and production are thought to be based on implicit knowledge, it seems important to determine whether subjects in statistical learning research develop this type of knowledge.

The present study seeks to address this gap. Specifically, we investigated whether the knowledge acquired in a typical statistical learning experiment is conscious, unconscious, or both. We argue that measures of awareness can provide a richer understanding of the cognitive mechanisms involved in statistical learning and of the nature of the resulting knowledge. Before describing our experiment, though, it is important to review how the conscious and unconscious status of knowledge can be assessed.

2. In the research tradition started by Reber (1967), the use of the term *implicit* is generally restricted to those situations where subjects have acquired unconscious knowledge under incidental learning conditions. If incidental exposure in an experiment results in conscious knowledge, e.g. when subjects were able to figure out the rule system despite not having been told about its existence, the learning process is usually only characterized as being *incidental* and not as implicit. The same applies for those experiments that do not include a measure of awareness. The term *explicit learning* is usually applied to learning scenarios in which subjects are instructed to actively look for patterns, i.e. learning is intentional, a process which tends to result conscious knowledge.

Measuring implicit and explicit knowledge of language

The question of whether knowledge acquired during incidental learning experiments is actually "implicit" is controversial.[3] Proposals for measuring awareness include verbal reports, direct and indirect tests, and subjective measures (see Dienes & Seth, 2010; Rebuschat, submitted, for overviews).

Verbal reports A common way of measuring awareness is to prompt subjects to verbalize anything they might have noticed while doing the experiment (e.g., Reber, 1967). Knowledge is considered to be unconscious if subjects perform above chance despite being unable to verbalize the knowledge that underlies their performance. The view that knowledge is unconscious when subjects are unable to verbalize the knowledge they have acquired has been criticized for a variety of reasons (see Perruchet, 2008). One problem is that subjects may only be able to verbalize knowledge after a long exposure period. Another problem is that verbal reports constitute a relatively insensitive and incomplete measure of awareness. For example, subjects may fail to verbalize knowledge because low-confidence knowledge retrieval may be difficult.

Direct and indirect tests Several authors have advocated the contrastive use of direct and indirect tests as a more exhaustive measure of awareness (e.g., Reingold & Merikle, 1988). Direct and indirect tests are not, per se, measures of awareness, but actually measures of learning. Generally, the performance on two tasks is compared. The *direct test* is a measure that explicitly instructs subjects to make use of their knowledge (e.g., a free generation task). The task encourages subjects to access all relevant conscious knowledge in order to perform. The *indirect test* assesses subjects' performance without instructing them to use their acquired knowledge (e.g., serial reaction time task). Knowledge is assumed to be unconscious if an indirect test clearly indicates a learning effect, even though a direct test shows no evidence of learning.

In the case of sequence learning, for example, Jiménez, Méndez, & Cleeremans (1996) used the serial reaction time (SRT; Nissen & Bullemer, 1987) task as an indirect measure and a generation task as a direct measure. Subjects performed on the two tasks successively. In the SRT task, subjects

3. As one of our reviewers pointed out, this might well be reason why statistical learning researchers have avoided the implicit/explicit distinction in the first place.

saw a stimulus appear at one of several locations on a computer screen and were asked to press as fast and accurately as possible on the corresponding key. Unbeknownst to subjects, the sequence of successive stimuli was determined by an artificial grammar. In the generation task, subjects were asked to predict the location of the next stimulus by pressing the corresponding key. Jiménez et al. (1996) found that subjects had clearly learned to exploit the regularities inherent in the stimulus environment. More importantly, they also found that some knowledge about the sequential structure of the material was exclusively expressed in the indirect task (SRT), but not in the relatively similar direct task (generation). Their results suggest that this knowledge was unconscious.

Direct and indirect measures have been criticized for a number of reasons, perhaps most often because the tests lack exclusivity. That is, they may not solely measure what they are supposed to. As Reingold & Merikle (1988) argue, direct tests are inadequate measures of conscious knowledge because they may be contaminated by unconscious knowledge. When direct tests indicate greater than zero sensitivity, it is unclear whether performance is driven exclusively by conscious knowledge or not. Thus, any approach based on these measures runs the risk of underestimating the influence of unconscious knowledge.

Subjective measures Dienes (2004, 2008, this volume) has advocated the use of subjective measures in order to assess whether the knowledge acquired during Artificial Grammar Learning (AGL) tasks is conscious or unconscious. One way of dissociating conscious and unconscious processes is to collect confidence ratings (e.g., Dienes, Altmann, Kwan, & Goode, 1995). In AGL, for example, subjects can be asked to report, for each grammaticality judgment, how confident they were in their decision. Dienes et al. (1995) suggested two ways in which confidence rating data could serve as an index of unconscious knowledge. Firstly, knowledge can be considered unconscious if subjects believe to be guessing when their classification performance is, in fact, significantly above chance. Dienes et al. called this the *guessing criterion*. Secondly, knowledge is unconscious if subjects' confidence is unrelated to their accuracy. This criterion was labeled *zero correlation criterion* by Dienes et al. Several studies have shown that performance on standard AGL tasks can result in unconscious knowledge according to these criteria (e.g., Dienes et al., 1995).

Structural knowledge and judgment knowledge. One criticism that can be leveled at the use of confidence ratings concerns the type of knowledge that is assessed by this measure. Consider the case of natural language

acquisition (Dienes, 2008). Language acquisition is often considered a prime example of implicit learning. All cognitively unimpaired adults are able to discern grammatical sentences of their native language from ungrammatical ones, even though they are unable to report the underlying rule system. However, if asked how confident they are in their grammaticality decisions, most native speakers will report high confidence levels, as in: "John bought an apple in the supermarket" is a grammatical sentence and I am 100% confident in my decision, but I do not know what the rules are or why I am right." Since in these cases accuracy and confidence will be highly correlated, does this mean that language acquisition is not an implicit learning process after all? Probably not. Dienes (2008; Dienes & Scott, 2005) proposed a convincing explanation for this phenomenon, based on Rosenthal's (2005) Higher-Order Thought Theory.

Dienes suggested that, when subjects are exposed to letter sequences in an AGL experiment, they learn about the structure of the sequences. This *structural knowledge* can consist, for example, of knowledge of associations, whole exemplars, knowledge of fragments or knowledge of rules (e.g., "A letter sequence can start with an M or a V.") In the testing phase, subjects use their structural knowledge to construct a different type of knowledge, namely whether the test items shared the same underlying structure as the training items (e.g., "MRVXX has the same structure as the training sequences."). Dienes labeled this *judgment knowledge*. Both forms of knowledge can be conscious or unconscious. For example, a structural representation such as "An R can be repeated several times." is only conscious if it is explicitly represented, i.e., if there is a higher-order thought such as "I {know/think/believe, etc.} that an R can repeated several times." Likewise, judgment knowledge is only conscious if there is a corresponding higher-order thought (e.g., "I {know/think/believe, etc.} that MRVXX has the same structure as the training sequences.") The guessing and the zero correlation criteria measure the conscious or unconscious status of judgment knowledge, not structural knowledge.

Dienes & Scott (2005) assume that conscious structural knowledge leads to conscious judgment knowledge. However, if structural knowledge is unconscious, judgment knowledge could still be either conscious or unconscious. This explains why, in the case of natural language, people can be very confident in their grammaticality decisions without knowing why. Here, structural (linguistic) knowledge is unconscious while (metalinguistic) judgment knowledge is conscious. The phenomenology in this case is that of intuition, i.e. knowing that a judgment is correct but not knowing why. If, on the other hand, structural and judgment knowledge are unconscious, the phenomenology is that of guessing. In both cases the structural knowledge acquired during training is unconscious. Dienes and Scott proposed that in AGL experiments the conscious status of both structural and judgment knowledge can be assessed concurrently by adding source attributions to the confidence ratings in the testing phase. That is, in addition to asking subjects how confident they were in their grammaticality judgments, one also prompts them to report the basis (or source) of their judgments.

In the next section, we will describe an experiment that applied the subjective measures developed by Dienes (2004, 2008, this volume; Dienes & Scott, 2005) to an established paradigm in statistical learning research, namely cross-situational word learning.

Method

The following experiment had two objectives. The first objective was to illustrate how subjective measures can be applied to the investigation of statistical learning. The second objective was to determine what type of knowledge (conscious or unconscious) subjects acquire in a typical statistical (word) learning paradigm. To our knowledge, no statistical learning experiment has empirically assessed whether subjects acquire conscious or unconscious knowledge as a result of exposure. The experiment below adds subjective measures of awareness to the cross-situational word learning paradigm (e.g., Kachergis et al., 2010; Yu & Smith, 2007) to address this gap.[4]

Participants

Thirty native speakers of English (19 women and 11 men, mean age = 19.3) were randomly assigned to incidental or intentional learning conditions. There were no significant differences between the two groups in terms of age or language background, *ps* > .05.

4. A recent study by Kachergis et al. (2010) compared cross-situational word learning under incidental and intentional learning conditions. That is, the emphasis was on whether or not intention to learn played a role in word learning. However, the study did not including measures of awareness to assess the conscious or unconscious status of the acquired knowledge.

Stimuli

An artificial lexicon consisting of 27 auditory pseudowords was created for this experiment. All pseudowords were bisyllabic, stressed on the first syllable, and obeyed English phonotactics. The pseudowords were read aloud by a female native speaker of English, digitally recorded and subsequently edited by means of sound processing software (Audacity, version 1.2.4). Each pseudoword was then matched with one or more black-and-white drawings from the International Picture-Naming Project website (Szekely et al., 2004).

To control for memorability, all possible stimuli (pseudowords and object images) were normed using two memory recognition tasks. Twelve undergraduates who were not involved in the main experiment participated in the memory recognition tasks. The procedure for both tasks was identical, with the only difference being whether or not pseudowords or pictures were being used. For the pseudoword memory task, we divided our pool of pseudowords (n = 68) into set A (n = 34) and set B (n = 34). In the memorization phase, half the participants were then instructed to memorize set A, while the other half was instructed to memorize set B. In the test phase, all subjects performed on a recognition task: They were presented with the complete pool (n = 68) of pseudowords and asked to indicate which items they had previously encountered. The procedure for the picture memory task was the same, except that the stimuli in this version were images of objects (n = 68). All twelve participants were given both pseudoword and picture memory tasks. The order of the tasks was counterbalanced across participants.

The analysis of the two tasks showed that the mean recognition rate was 3.52 for pseudowords and 3.76 for the pictures of objects. That is, on average 3.52 people recalled pseudowords correctly at test while 3.76 people recalled the pictures correctly. We decided to use items that had recognition rates between 3 and 4 (the closest to average) for our experiment, and discarded stimuli with memory rates of 1, 2, 5, and 6. This was because we neither wanted to use stimuli that were too easy to remember nor stimuli that were too difficult to retain. The selected stimuli were then randomly paired to create the artificial vocabulary.

The lexicon was divided into 12 target items and 15 fillers.[5] All filler items were unambiguous and only occurred once each in the input during

5. The reason for this manipulation is that we were also interested in determining the role of frequency in the development of implicit and explicit knowledge. This data will be reported elsewhere, however (Hamrick & Rebuschat, under review).

Table 1. Ambiguous and Unambiguous Target Items and Their Referents

Pseudoword	Referents (Co-Occurrence Frequency)
dobez	backpack (6), arrow (4), bathtub (2)
paylig	wheelchair (6), towel (4), bandage (2)
femod	bench (6), thumb (4), bridge (2)
whoma	comb (6), crib (4), fan (2)
houger	elephant (6), glass (4), pear (2)
jillug	ladder (6), leaf (4), mixer (2)
keemuth	mop (6)
nengee	panda (6)
zomthos	radio (4)
loga	stethoscope (4)
shrama	robot (2)
thueek	tank (2)

the exposure phase. The target items were subdivided into six lexically ambiguous pseudowords (one word, three matching referents) and six lexically unambiguous pseudowords (one word, one matching referent). All target words were manipulated in terms of their word-referent co-occurrence frequencies. Some words co-occurred with their matching referents six times, other words co-occurred with their appropriate referents four times, and others co-occurred with their appropriate referents twice. For example, the pseudoword *houger* occurred 12 times: Six times with an elephant, four times with a glass, and two times with a pear.

Procedure

The experiment was presented by means of a PC with a 15.6 inch screen using Microsoft Power Point 2007. Instructions were displayed in black text (Arial font sizes 20–24) on a white background. Pseudowords were played through headphones. The experiment consisted of an exposure phase and a testing phase. The testing phase was the same for both groups. The groups differed in how they interacted with the 57 exposure trials.

Exposure phase In the exposure phase, subjects in both conditions were presented with the same 57 trials. In each trial, two images were displayed on the screen at the same time, one on the left, the other on the right side

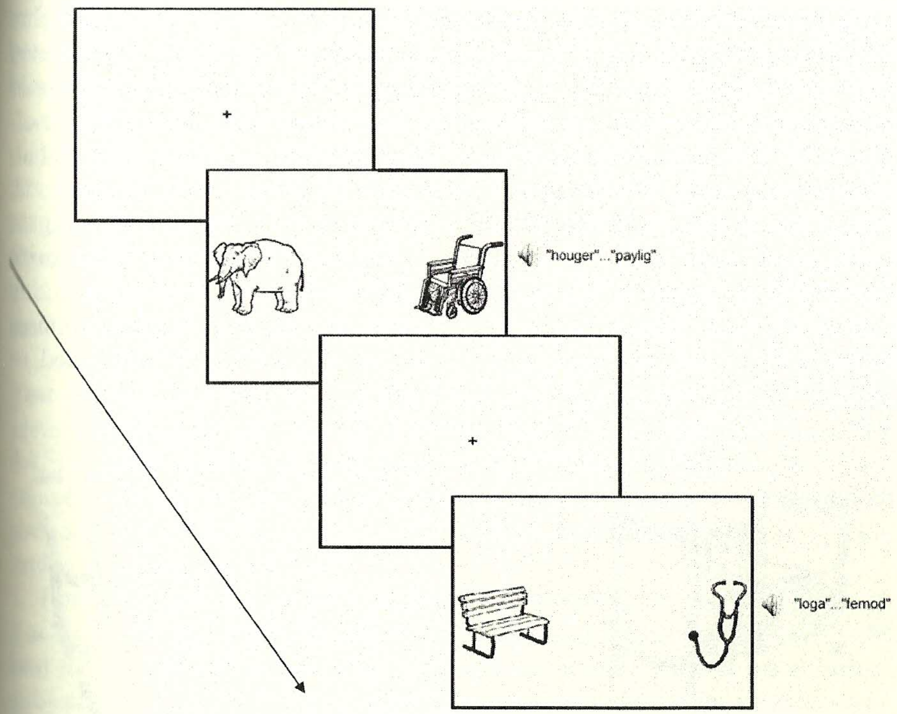

Figure 1. Sample screenshot from the exposure phase. Participants in both experimental conditions were presented with the same 57 trials. In each trial, a fixation cross was first displayed for two seconds. This was followed by the concurrent presentation of two pictures and two spoken pseudowords. Importantly, the presentation order of the pseudword was not related to the location of the image on the screen.

of the monitor. The two images were displayed for six seconds. While the images were on display, two pseudowords were played once. For example, subjects might see an image of a panda on the left and an image of a glass on the right, while hearing first the pseudword *houger*, followed by the pseudoword *femod*. Importantly, the presentation order of the pseudword was not related to the location of the image on the screen. That is, each word could refer either to the image on the left or to the image on the right. The only way for participants to learn the artificial vocabulary was to keep track of the pseudoword-object co-occurrences across trials. The order of trials was randomized for each participant.

Subjects in the *intentional learning condition* (n = 15) were told that they were participating in a word-learning experiment and were instructed to "learn the meanings of the words". They were also told that they would be tested afterwards. In contrast, subjects in the *incidental learning condition* (n = 15) were not informed about the true purpose of the experiment, nor did they know that they would be tested after the exposure phase. Instead, they were told that the objective of the study was to investigate how people with different language experience perceive and categorize objects. Their task during the exposure phase was to indicate how many objects on each slide were animate. There were three possible responses (zero, one, or two animate objects) and participants were instructed to

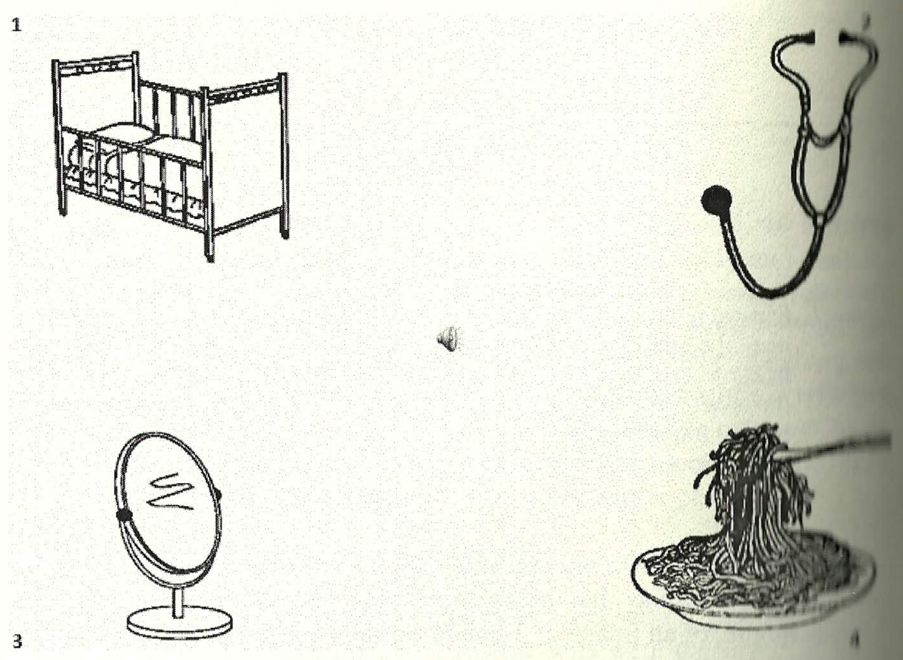

Figure 2. Sample screenshot of the four-alternative forced-choice (4AFC) picture matching task. The 4AFC task consisted of thirty trials. In each trial, participants were presented with four pictures, one in each corner of the screen, and a spoken pseudoword. Their task was to select the appropriate referent as quickly and accurately as possible. In addition, subjects were also asked to report how confident they were in their decision and what the basis of their decision was.

enter 0, 1, or 2 on their keypads. This task was made more difficult by the presence of pictures that were not easily classifiable as animate or inanimate (e.g., a thumb, a leaf). They were informed that they would have to do the task while hearing "nonsense" words through their headphones.

In sum, all experimental subjects were exposed to the same 57 trials. The key difference between subjects in the intentional group and subjects in the incidental group is how they interacted with the stimuli. Subjects in the former group were instructed to learn the meanings of words, whereas subjects in the latter group were asked to perform on an irrelevant task and to treat the auditory pseudowords as a distraction.

Test phase After the exposure phase, all participants were asked to complete a four-alternative forced-choice (4AFC) picture matching task. The 4AFC task consisted of thirty trials. In each trial, participants were presented with four pictures, one in each corner of the screen, and a spoken pseudoword. Their task was to select the appropriate referent as quickly and accurately as possible.

For each trial, the screen contained one correct referent and three foils. Each picture was numbered (1 through 4) and participants indicated the best match by writing down their answers on an answer sheet. In addition, subjects were also asked to report how confident they were in their decision and what the basis of their decision was. Subjects were asked to place their confidence on a continuous scale, ranging from 50% (*complete guess*) to 100% (*complete certainty*). We emphasized that subjects should only use "50%" when they believed to be truly guessing, i.e., they might as well have flipped a coin. In the case of the source attributions, there were three response options: *guess*, *intuition*, and *memory*. The guess category indicated that subjects believed the classification decision to be based on a true guess. The intuition category indicated that they were somewhat confident in their decision but did not know why it was right. The memory category indicated that the judgment was based on the recollection of pseudoword-referent mappings from the exposure phase. All participants were provided with these definitions before starting the testing phase.

At the end of the test phase, all subjects completed a debriefing questionnaire, which asked them to report if they had learned any of the pseudoword-referent mappings during exposure, whether or not they had used any specific learning strategies and, if so, what kind of strategies.

Results

Performance on the 4AFC task served as the measure of learning. Awareness was measured by means of confidence ratings and source attributions.

Four-alternative forced-choice task

The analysis of the 4AFC task showed that the incidental group classified 44.4% ($SD = 7.5\%$) of the test items correctly and the intentional group 73.3% ($SD = 10.7\%$). Both the incidental group, $t(14) = 9.99$, $p < .05$, and the intentional group, $t(14) = 17.53$, $p < .05$, performed significantly above chance (25%), which indicates that there was a clear learning effect for both groups. Further analysis showed that the difference between the two groups was significant, $t(28) = 8.559$, $p < .001$, i.e., the learning effect was greater under intentional learning conditions.

Confidence ratings

The average confidence level was 61.3% ($SD = 7.2\%$) in the incidental group and 80.6% ($SD = 6.3\%$) in the intentional group. The difference was significant, $t(28) = 7.79$, $p < .05$. Further analysis showed that accuracy and confidence were significantly correlated in the intentional group, $r = .77$, $p < .05$, but not in the incidental group, $r = .45$, $p > .05$. When intentional learners were confident in their decision, they tended to be accurate. This suggests that subjects in the intentional group had acquired conscious judgment knowledge: These participants were partially aware that they had acquired some knowledge during the exposure phase. In contrast, subjects in the incidental group were not aware of having acquired knowledge, despite the fact that their performance on the 4AFC task clearly indicates that they did. The zero correlation criterion was thus met in the case of the incidental group.

We then analyzed all classification decisions for which subjects gave a 50% rating, i.e., they believed to have guessed when deciding on the appropriate referent for the pseudoword. When subjects in the incidental group gave a confidence rating of 50%, their classification performance was 34.0% ($SD = 47.5\%$), which was significantly above chance, $t(140) = 2.26$, $p < .05$. In the case of the intentional group, when subjects gave a confidence rating of 50% their classification performance was 45.0% ($SD = 50.3\%$), also significantly above chance, $t(39) = 2.51$, $p < .05$. That is, the guessing criterion for unconscious judgment knowledge was satisfied in both groups. Subjects in both conditions had acquired at least some unconscious judgment knowledge.

The confidence ratings thus indicate that the incidental group was largely unaware of having acquired knowledge during the exposure phrase. In the case of the intentional group, subjects were clearly aware of having acquired knowledge (see correlation between confidence and accuracy), though some of their judgment knowledge did remain unconscious (as indicated by guessing criterion).

Source attributions

In terms of proportion, the incidental group most frequently believed their classification decisions to be based on a guess or intuition (86% of judgments). The memory category was selected least frequently (only 14% of all judgments). That is, when performing on the 4AFC task, subjects in the incidental group generally based their decisions on the more implicit categories. In the case of the intentional group, the memory category was selected most frequently (61% of judgments), followed by guessing and intuition. In terms of accuracy, the analysis showed that the incidental group scored highest when reporting that their classification was based on memory, followed by the intuition and guess categories. The same pattern was observed in the intentional group, i.e., these subjects were most accurate when attributing their classification decision to memory. They were, however, considerably more accurate, performing close to 90% accuracy.

Further analyses revealed significant effects of source attribution in both the incidental group, $F(2,16) = 8.247$, $p < .05$, and the intentional group, $F(2,22) = 5.49$, $p < .05$. In the case of the incidental group, the difference between decisions based on guessing and decisions based on intuition was significant, $p < .05$, as was the difference between decisions based on guessing and those based on memory, $p < .05$. In the case of the intentional group, the differences between decisions based on guessing and intuition, guessing and memory, and intuition and memory were all significant, $p < .05$.

Interestingly, subjects in both groups performed significantly above chance across categories, irrespectively of whether they attributed their decision to guessing, intuition, or memory. The guessing criterion was therefore satisfied in both groups: When subjects gave a confidence rating of 50%, indicating that they were forced to guess the right answer in the 4AFC task, their actual classification performance suggests that they had acquired the knowledge to make that decision. This suggests that subjects

Table 2. Accuracy and proportions (%) across source attributions

		Guess	Intuition	Memory
Incidental	Accuracy	35.8*	48.5**	61.4**
	Proportion	44.2	41.7	14.1
Intentional	Accuracy	54.2**	61.9**	88.9**
	Proportion	23.2	27.9	48.9

Significance from chance (25%): $*p < .01$, $**p < .001$.

in both groups acquired at least some unconscious structural knowledge. Table 2 shows the classification performance for the different attributions.

Verbal reports

Analysis of the verbal reports showed that learners in the intentional condition became aware of many pseudoword-referent pairs and were able to name a few. When prompted for strategies, the most commonly reported strategies were repeating the pseudowords, making a link between pseudowords and prior knowledge (e.g., "that sounded like something in French"), and hypothesis testing. In contrast, subjects in the incidental group reported deliberately trying to block out the pseudowords. Indeed, many interpreted the pseudowords to be a distraction and consequently tried to ignore them.

Discussion

The results of the present experiment show that subjects can use statistical information to learn new words, which is consistent with previous research (e.g., Yu & Smith, 2007). Moreover, statistical word learning can take place under both intentional and incidental learning conditions (see also Kachergis et al., 2010). Subjects learn to associate pseudowords with their appropriate referents even when they are instructed to perform on an irrelevant task (indicating animacy) and disregard the auditory pseudowords. However, the learning effect is greater under intentional learning conditions. Instructing subjects to "learn the meanings of the words" resulted in higher accuracy in the 4AFC task.

The analysis of the confidence ratings and the source attributions showed that the learning condition plays an important role in the type of

knowledge that subjects acquire. Under incidental learning conditions, subjects developed primarily implicit knowledge. These subjects were not aware of having acquired knowledge, reported most of their decisions to be based on guessing or on intuition, and performed significantly above chance in the discrimination task even when they believed to be guessing. In contrast, subjects in the intentional learning condition acquired primarily conscious knowledge. These subjects were aware of having acquired knowledge and tended to be highly accurate when reporting high levels of confidence. Moreover, these subjects also acquired some unconscious knowledge. These findings are consistent with implicit learning research, where it is often found that subjects develop unconscious knowledge even when explicitly instructed to discover rule structure (e.g, Guo et al., in press; Rebuschat, 2008, Experiment 6; Rebuschat & Williams, 2009).

Our data suggest that incidental statistical learning is more likely to result in implicit knowledge. They also suggest that instructing subjects to learn the meanings of words prompts them to use strategies resulting in explicit knowledge. Interestingly, subjects using explicit strategies still acquired some unconscious knowledge. This supports the view that implicit learning may proceed in parallel with explicit learning (cf. Guo et al., in press). Thus, statistical word learning may result in both implicit and explicit knowledge, and learning conditions appear to influence the extent to which each type of knowledge develops.

The experiment illustrates the usefulness of including measures of awareness when researching statistical learning. Depending on the learning conditions of the experiment, subjects may rely on general learning mechanisms, explicit strategies, or both, and the use of subjective measures provides a method for determining the contributions of different learning conditions. Likewise, it may also be the case that different statistical regularities may differentially promote implicit or explicit knowledge (Perruchet & Pacton, 2006). Measures of awareness would allow us to empirically verify such a hypothesis.

In summary, the present study shows that statistical word learning can result in both implicit and explicit knowledge; however, the amount and quality of each kind of knowledge was influenced by learning condition. Thus, subjective measures of awareness may provide researchers with a richer characterization of the knowledge acquired in statistical learning experiments. We argue that the present paper brings to light some important considerations for further investigations of the relationship between implicit learning and statistical learning.

Acknowledgments

The authors would like to thank Luke Amoroso, Natalie Brito, Katie Jeong-Eun Kim, Julie Lake, Kaitlyn Tagarelli and three anonymous reviewers for their helpful feedback. Special thanks to Jennifer Johnstone and Michelle To for their support.

References

Aslin, R. N., Saffran, J. R., & Newport, E. L.
1999 Statistical learning in linguistic and non-linguistic domains. In B. MacWhinney (Ed.), *Emergentist approaches to language*. Hillsdale, NJ: Lawrence Erlbaum.

Conway, C. M., & Christiansen, M. H.
2006 Statistical learning within and between modalities. *Psychological Science*, 17, 905–912.

Dell, G. S., Reed, K. D., Adams, D. R., Meyer, A. S.
2000 Speech errors, phonotactic constraints, and implicit learning: A study of the role of experience in language production. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26, 1355–1367.

Dienes, Z.
2004 Assumptions of subjective measures of unconscious mental states: Higher order thoughts and bias. *Journal of Consciousness Studies*, 11(9), 25–45.

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

Dienes, Z.
this volume Conscious versus unconscious learning of structure. In P. Rebuschat & J. N. Williams (Eds.) *Statistical learning and language acquisition*. Berlin: Mouton de Gruyter.

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

Dienes, Z., & Seth, A.
2010 The conscious and the unconscious. In G. Koob, M. Le Moal, & R. F. Thompson (Eds), *Encyclopedia of Behavioral Neuroscience*, vol. 1, (pp. 322–327). Oxford: Academic Press.

Dienes, Z., Altmann, G., Kwan, L., & Goode, A.
1995 Unconscious knowledge of artificial grammars is applied strategically. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 21(5), 1322–1338.

Gómez, R. L.
2007 Statistical learning in infant language development. In M. G. Gaskell (Ed.), *The Oxford Handbook of Psycholinguistics*. Oxford: Oxford University Press.

Guo, X., Zheng, L., Zhu, L., Yang, Z., Chen, C., Zhang, L., Ma, W., & Dienes, Z.
in press Acquisition of conscious and unconscious knowledge of semantic prosody. *Consciousness and Cognition.*

Hamrick, P., & Rebuschat, P.
under review Measuring frequency effects on the development of implicit and explicit lexical knowledge. *Proceedings of the 2012 Georgetown University Roundtable on Languages and Linguistics Conference.* Washington, D.C.

Jiménez, L., Méndez, C., & Cleeremans, A.
1996 Comparing direct and indirect measures of implicit learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 22, 948–969.

Kachergis, G., Yu, C., & Shiffrin, R. M.
2010 Cross-situational statistical learning: Implicit or intentional? *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Misyak, J. B., Goldstein, M. H., & Christiansen, M. H.
this volume Statistical-sequential learning in development. In P. Rebuschat & J. N. Williams (Eds.) *Statistical learning and language acquisition*. Berlin: Mouton de Gruyter.

Nissen, M. J., & Bullemer, P.
1987 Attentional requirements of learning: Evidence from performance measures. *Cognitive Psychology*, 19, 1–32.

Perruchet, P.
2008 Implicit learning. In J. Byrne (Ed.), *Learning and memory: A comprehensive reference* (Vol. 2, pp. 597–621). Oxford: Elsevier.

Perruchet, P., & Pacton, S.
2006 Implicit learning and statistical learning: One phenomenon, two approaches. *Trends in Cognitive Sciences*, 10(5), 233–238.

Reber, A. S.
1967 Implicit learning of artificial grammars. *Journal of Verbal Learning and Verbal Behavior*, 6, 317–327.

Rebuschat, P.
2008 *Implicit learning of natural language syntax*. Unpublished PhD dissertation, University of Cambridge.

Rebuschat, P.
submitted Measuring implicit and explicit knowledge in second language research: A review.

Rebuschat, P., & Williams, J. N.
2009 Implicit learning of word order. *Proceedings of the 31st Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society.

Reingold, E. M., & Merikle, P. M.
1988 Using direct and indirect measures to study perception without awareness. *Perception & Psychophysics*, 44, 563–575.
Rosenthal, D. M.
2005 *Consciousness and mind.* Oxford: Oxford University Press.
Saffran, J. R., Aslin, R. N., & Newport, E. L.
1996 Statistical learning by 8-month-old infants. *Science*, 274, 1926–1928.
Szekely, A., Szekely, A., Jacobsen, T., D'Amico, S., Devescovi, A., Andonova, E., Herron, D., . . . Bates, E.
2004 A new on-line resource for psycholinguistic studies. *Journal of Memory and Language*, 51, 247–250.
Warker, J. A., Dell, G. S.
2006 Speech errors reflect newly learned phonotactic constraints. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 32, 387–398.
Yu, C., & Smith, L. B.
2007 Rapid word learning under uncertainty via cross-situational statistics. *Psychological Science*, 18, 414–420.

What Bayesian modelling can tell us about statistical learning: What it requires and why it works

Amy Perfors and Daniel J. Navarro

Introduction

The purpose of this chapter is to address some issues regarding the *why* and *what* of statistical learning, with a particular focus on Bayesian computational modelling and language acquisition. We begin with a brief introduction to Bayesian modelling, contrasting it with the other primary computational approach to statistical learning (connectionist modelling), and demonstrating how it clarifies some common confusions about what statistical learning is and is not. The chapter is structured around a series of questions: What is statistical learning? What data does statistical learning operate on? What knowledge does learner acquire from the data? What assumptions do learners make about the data? What prior knowledge does the learner possess? Finally, why does statistical learning work? Each of these is a big topic in itself, so we aim only to provide a general introduction to them, covering some but not all of the issues involved.

What is statistical learning?

Statistical learning encompasses a wide variety of learning situations in which the knowledge acquired by the learner is highly dependent on the statistical structure of the data that they are given. From an empirical perspective, researchers are interested in finding out whether and to what extent people are sensitive to statistical structure (e.g., the frequencies of different events) when learning from data. From a formal perspective, we aim to describe the abstract principles and processes that are necessary to explain how the learner might acquire knowledge based on statistical input. Statistical learning can be distinguished from learning that relies solely on deterministic rules, such as the subset principle (Berwick, 1986) or learning that requires a certain type of input before acting, like "trigger" learning (Gibson & Wexler, 1994).