4th International Conference on
Development and Learning and on Epigenetic Robotics
October 13-16, 2014. Palazzo Ducale, Genoa, Italy

Computational model for syntactic development: Identifying how children learn to generalize nouns and verbs for different languages

Yuji Kawai*, Yuji Oshima*†, Yuki Sasamoto*, Yukie Nagai*, and Minoru Asada*
*Graduate School of Engineering, Osaka University
2-1 Yamada-oka, Suita, Osaka, 565-0871 Japan
†NTT Software Innovation Center since April 2014
Email: {yuji.kawai, yukie, asada}@ams.eng.osaka-u.ac.jp

***Abstract*—By three years of age, children are supposed to start learning to understand syntactic structures, and at around five years of age, they are reported to be able to infer a syntactic category, such as a noun or a verb, for a novel word. Finding the syntactic cue enables them to infer a target directed by a novel word in visual stimuli. The study also found that their inference performances depended on their native languages. In this article, we propose a model to explain how children learn to generalize novel nouns and verbs in the Japanese, English, and Chinese languages. We use a Bayesian hidden Markov model (BHMM) to learn syntactic categories represented as hidden states in a BHMM. Here, an increase in the number of hidden states indicates the children's syntactic development. A model with a larger number of hidden states is able to infer a clearer syntactic category of a novel word, resulting in the correct choice of a category for the visual target. Syntactic categories that depend on input languages are acquired by BHMMs, and therefore result in different performances among the languages. We entered English-, Japanese-, or Chinese-corpus into the model and examined how the model inferred a correct target indicated by a novel word through the acquired syntactic categories. The results showed that the performances by our model are very similar to the children's performances. Further analysis of representations of hidden states clarified that the model acquires syntactic categories reflecting orders of words in English, suffixes in Japanese, and adverbs in Chinese.**

I. Introduction

There have been many observational and model studies on children's syntactic development because understanding how children acquire syntactic categories is one of the key steps to reveal mechanisms of language faculty. Children start to speak multi-word sentences at about twenty to twenty-four months of age. Their sentences seem to already follow some syntactic rules [1], but their syntactic categories, representing the role of a word in a sentence (e.g., a noun and a verb), are not yet mature [2], [3]. Syntactic categories are gradually acquired from thirty-six months [3]. However, the underlying mechanism for the development of syntactic categories has not been completely understood. One method to study children's syntactic development is to utilize the nouns and verbs generalization task, which examines whether a child can infer a syntactic category of a novel noun or verb and map it to a correct visual target (e.g., an object or an action) [3], [4]. A child observes the standard stimulus that a woman conducts an unfamiliar action with an unfamiliar object. An English-speaking child is given, at the same time of the standard stimulus, a novel word "dax" within one of the following sentences:

- Noun condition: This is a dax.
- Bare verb condition: Daxing.
- Verb with arguments condition: She is daxing it.

Then the two test stimuli are shown to the child. In the first one, only the object is different from the standard stimulus (Object-change: OC stimulus). In the second, the action differs, but the object is same as the standard stimulus (Action-change: AC stimulus). The child is asked to choose which stimulus is "dax." The correct choice is AC stimulus under the noun condition or OC stimulus under the verb conditions. Here, the child must choose a visual target indicated by a novel word according to a syntactic cue. For example, the child infers a syntactic category of a novel word, a noun or a verb, based on words before and after the novel word, and can generalize the novel word based on the knowledge that a noun and a verb correspond to an object and an action, respectively.

Imai et al. [3] carried out this test with English-, Japanese-, and Chinese-speaking children. The results indicated that three-year-olds were able to correctly generalize a novel noun, but failed to generalize a novel verb, irrespective of their native language. In contrast, the ability to generalize a novel verb in the case of five-year-olds depended on their native language. Japanese five-year-olds successfully generalize a novel verb in both verb conditions, because they can discriminate a verb from a noun through suffixes in Japanese. English- and Chinese-speaking children, on the other hand, cannot generalize a novel verb in the bare verb condition. The reasons are that a bare verb rarely appears in English, and that it cannot be discriminated from a noun in Chinese [3]. Children acquire syntactic categories that reflect their native language from three to five years of age. However, observational studies only can hardly explain a detailed structure of syntactic categories that enable children to generalize nouns and verbs during the developmental process.

Thus, a computational model is suggested as it may reveal a detailed structure as to how children may acquire syn-

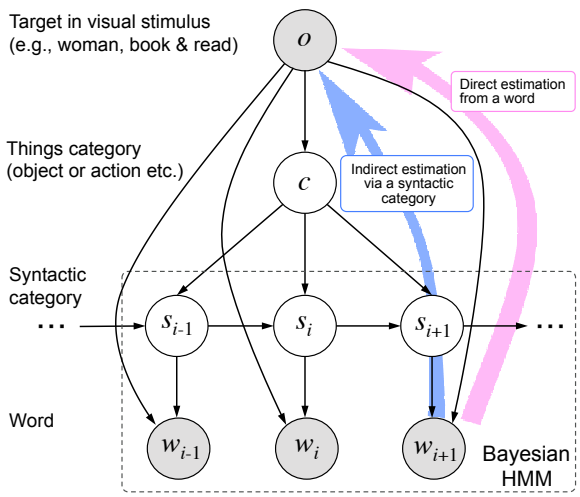

Fig. 1. A graphical model to infer a visual target indicated by a word based on a syntactic cue

tactic categories. An Elman network can represent syntactic categories as a result of predictive learning of sequences of words. Language learning with an artificial neural network is often compared to corresponding human language learning [5], [6]. Toyomura and Omori [7] proposed a model to infer a perceptual target indicated by a novel word based on the syntactic representations of an Elman network. The network is given simple three-word-sentences in English and then acquires syntactic categories. The syntactic categories are associated with perceptual targets that enable the model to infer a target indicated by a novel word based on a syntactic category. However, it is very difficult for a simple Elman network to learn a complex language structure, including word abbreviations and changes of word order as seen in Japanese; error minimization algorithm used by the neural network may neglect minority sentence patterns. This is why their model cannot explain the syntactic development and cross-lingistic difference reported by Imai et al. [3].

We propose a model that explains the developmental change in the syntactic understanding of three- to five-year-olds and its dependence on the language structure reported by Ima et al. [3]. We also describe an analysis of internal representations of the model that reveals a structure of syntactic categories. We used a Bayesian hidden Markov model (BHMM) [8] as the learner of syntactic categories. Our model represents syntactic development as an increase in the number of hidden states of a BHMM. The stochastic syntactic representation in our model enables it to cope with complex language rules, such as word abbreviations and changes of word order often observed in Japanese that are difficult for simple connectionist models to represent.

II. A Model for Syntactic Development

A. Overview

Fig. 1 illustrates the proposed graphical model. The BHMM at the bottom is given word sequences $\boldsymbol{w}$ and estimates the hidden state s_i of each word w_i as a syntactic category. A things category c (i.e., abstract and discrete perceptual information such as an action and an object) generates a series of syntactic categories $\boldsymbol{s}$, and c is stochastically associated with a target o in a visual stimulus. There are two pathways to estimate a probability of o from w_i: The first is the direct estimation $P(o|w_i)$ if w_i is known. The second is the indirect estimation $P(o|c)P(c|s_i)P(s_i|w_i)$ through s_i if w_i is unknown. The children inferred a *novel* target directed by a novel word in the experiment of Imai et al. [3]. The things category c is needed for the inference of o which has not been learned; the model cannot learn an association between novel o and any linguistic variables (w and s). The indirect pathway through $\boldsymbol{s}$ and c requires an accurate estimation of $P(s_i|w_i)$. If the number of hidden states S is small, then a hidden state represents multiple parts of speech; for example syntactic categories where nouns, verbs, and suffixes are mixed (see Fig. 2(a)), resulting in difficulty in an accurate estimation of c and o. A sufficiently large S enables s_i to represent a specific part-of-speech, and thus the model can correctly judge a directed target (see Fig. 2(b)). We consider this improvement in inference as syntactic development: The models with small and large Ss represent three-year-olds and five-year-olds, respectively.

B. Inference of a visual target from a word

The conditional probability of a visual target o under a word w_i in a sentence $\boldsymbol{w}$ is given by

$$P(o|\boldsymbol{w}, \boldsymbol{s}_{-i}) = \sum_{s_i} \sum_{c} P(o|c)P(c|s_i)P(s_i|\boldsymbol{w}, \boldsymbol{s}_{-i})P(o|w_i), \tag{1}$$

where $\boldsymbol{s}_{-i}$ represents $\boldsymbol{s}$ excluding s_i which is a hidden state of w_i. From Bayes' theorem Eq.(1) is written as

$$\begin{aligned} &P(o|\boldsymbol{w}, \boldsymbol{s}_{-i}) \\ &\propto \sum_{s_i} \sum_{c} P(c|o)P(s_i|c)P(c)P(s_i|\boldsymbol{w}, \boldsymbol{s}_{-i})P(w_i|o)P(o), \end{aligned} \tag{2}$$

where the first term on the right side represents a mapping from a directed target to its things category, which is given *a priori*. The second term is a concurrence relation between a syntactic category and a things category, which is given by

$$P(s_i|c = c_j) = \frac{n(s_i, c_j)}{n(c_j)}, \tag{3}$$

where $n(s_i, c_j)$ and $n(c_j)$ denote how many times s_i and c_j appear simultaneously, and c_j occurs in all sentences, respectively. The prior probability distributions of the third and sixth terms are uniform over the currently presented stimuli (e.g., objects and actions). The fourth term represents the inference of syntactic categories, which is obtained by a BHMM [8]. The fifth term denotes a correspondence between

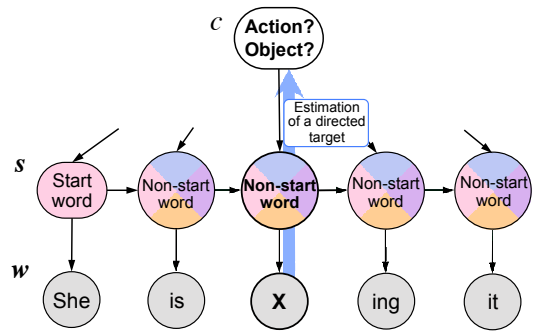

(a) Estimation of a probability of a directed target through a non-differentiated syntactic category

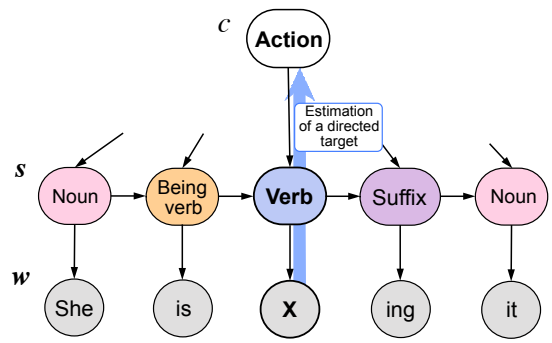

(b) Estimation of a probability of a directed target through a well-differentiated syntactic category

Fig. 2. Syntactic development by the differentiation of syntactic categories

directed targets and words, i.e., vocabularies, which is given as

$$P(w_i|o = o_t) = \frac{n(w_i, o_t)}{n(o_t)}, \tag{4}$$

where, $n(w_i, o_t)$ and $n(o_t)$ denote how many times w_i and o_t co-occur, and o_t is observed in learning data, respectively. This distribution is uniform if w_i is novel.

In the learning phase, the model given learning corpus $\boldsymbol{w}$ estimates a syntactic category s_i for each word w_i, that is, it approximately computes the right side fourth term in Eq. (2) by Gibbs sampling method [8]. The number of hidden states S is set by the experimenter. Then, Eqs. (3) and (4) are calculated. In the testing phase, the model is given a novel word X, and estimates a conditional probability of o given X from Eq. (2).

III. Experiment

A. Experimental setting

We created artificial corpora reflecting English, Japanese, or Chinese-language structure as learning data. Each corpus consists of nouns, verbs, adjectives, and so on, and their suffixes are dissociated from words (e.g., verb / ing in English). The grammatical structures in each corpus and their proportions are determined based on real corpora that consisted sentences spoken by caregivers to their two to five year old children [9]–[11] as well as on the basis of a study analyzing child-directed-speech [12] (see Appendix for more details). Abbreviations of a subject and an object word are consequently fewer in the English and Chinese corpora than in the Japanese one. The percentage of single sentences consisting "Verb-ing" is very small (0.2 %), especially in the English corpus. On the other hand, the Japanese corpus has many more abbreviations than the other two corpora.

The labels corresponding to nouns, verbs, and adjectives in sentences are also entered into the model as directed targets. When the sentence "She is read ing a red book." is selected, for example, the directed targets "girl", "read", "red", and "book" are also entered into the model. The model learns 10,000 sets of sentences and directed targets in this experiment. A novel word X that is not included by any corpora is given to the model after the learning. The form of the appearance of X is borrowed from that of Imai et al. [3].

English

- Noun condition: This is a X.
- Bare verb condition: X ing.
- Verb with arguments condition: She is X ing it.

Japanese

- Noun condition: X ga (nominal particle) aru (exist).
- Bare verb condition: X teiru (progressive).
- Verb with arguments condition: Oneesan (girl) ga (nominal particle) nanika (something) wo (accusative particle) X teiru (progressive).

Chinese

- Noun condition: Nali (there) you (exist) ge (classifier) X.
- Bare verb condition: X.
- Verb with arguments condition: Ayi (girl) zai (progressive) X yi (one) ge (classifier) dongxi (thing) ne (mode marking particle).

Four directed targets corresponding to a "girl" (known), an adjective (known), a novel noun, and a novel verb are also given to the model. The model estimates the probability of a target o from X by Eq. (2), and chooses between an object and an action in a random manner if the model's estimation supports a "girl" or the target corresponding to adjectives. This assumption is consistent with the children's forced-choices between OC and AC stimuli even if their inference for a target of a novel word is other than an object or an action in the nouns and verbs generalization task [3], [4]. We conducted experiments, in which the model estimated a conditional probability of o given X under each sentence condition, twenty times with different initial BHMM parameters. In order to replicate the experimental results reported by Imai et al. [3], a number of the hidden states S in a three-year-old model and a five-year-old model were set to three and from five to seven, respectively. The one sample t-test (two-tailed) evaluated correctness of the estimation, which is indicated by the facts that the noun condition estimates were significantly smaller than the value for chance level (0.5) and the verb condition estimates were significantly larger than the value

for chance level (0.5). We, furthermore, analyzed structures of part-of-speech in the acquired representations of hidden states (i.e., syntactic categories). We tagged each word in the learned corpora with a part-of-speech, and calculated the sum of $P(s_i|\boldsymbol{w}, \boldsymbol{s}_{-i})$ in Eq. (2) in each part-of-speech. Rates of the summed probabilities in each hidden state were obtained as representations of part-of-speech in the hidden states.

B. Results

Fig. 3 depicts proportions that a novel word is inferred as an action under each language condition: (a) English-, (b) Japanese-, or (c) Chinese-language. The small solid and empty circles in Fig. 3 denote the children's results reported by Imai et al. [3]. These results demonstrate that our model provided estimates similar to those of the children studied by Imai et al. [3]. The noun generalizations were successful in all conditions (all ps $< .05$). Novel verb generalizations were successful in some conditions with a large S depending on the language. The English model, with seven hidden states successfully generalized a verb with arguments ($p < .01$), while it failed to generalize a bare verb ($p > .05$). This is due to the fact that the English corpus has very few bare verb sentences: "Verb ing." The Japanese-speaking five-year-old model ($S = 6$), correctly judged an action in both verb conditions (both ps $< .01$). This is because Japanese suffixes distinguish verbs from nouns. Results for the Chinese model were similar to the English one because Chinese bare verbs cannot discriminate between verbs and nouns. The different S values among languages were set in the five-year-old models in order to reproduce the children's performance in the nouns and verbs generalization task [3]. However, all language models with the same number of hidden states (e.g., $S = 7$) indicated trends similar to the results shown in Fig. 3.

Figs. 4–6 depict typical representations of parts-of-speech in all syntactic categories. In the English model (Fig. 4), the syntactic categories with $S = 3$ represented start words (nouns) in sentences (ID 1 in Fig.4(a)) and non-start words (ID 2 in Fig.4(a)). These categories enable the model to generalize a noun but not to generalize a verb. When S was large the model acquired well-differentiated syntactic categories. The categories of nouns were separated into start words, i.e., subjects (ID 2 in Fig. 4(b)), and other nouns (ID 1 in Fig.4(b)), suggesting that English syntactic categories depend on the word order. The English model when given a novel bare verb classified the start word as a noun, and consequently the verb was incorrectly inferred to be an object. In the Japanese model (Fig. 5), verb generalization failed because the hidden state representing nouns confused syntactic categories of verbs and adjectives. The model could generalize a noun because there was the relatively higher ratio of nouns in the syntactic category (ID1 in Fig.5(a)) than other parts-of-speech, and the syntactic categories in some trials could acquire differentiated representations of nouns. The representation of each part-of-speech was differentiated when $S = 6$, which lead to successful verb generalization. Note that suffixes were separated into auxiliary verbs (ID 4 in Fig. 5(b)) and endings of adjectives

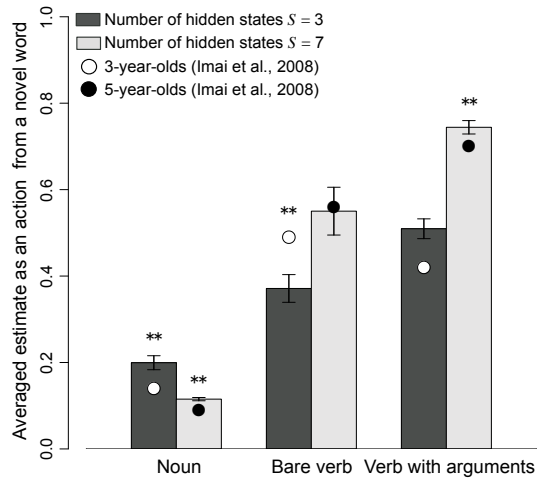

(a) English input

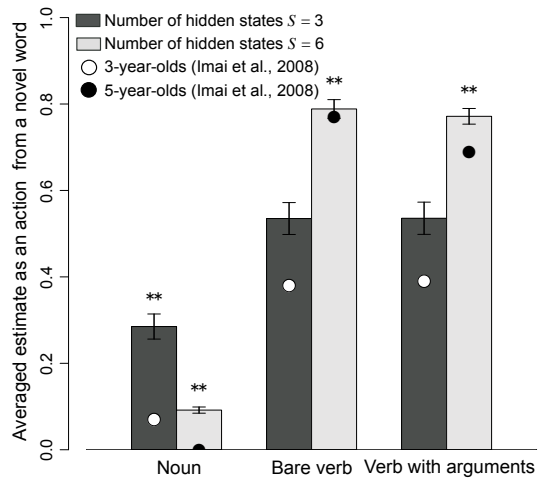

(b) Japanese input

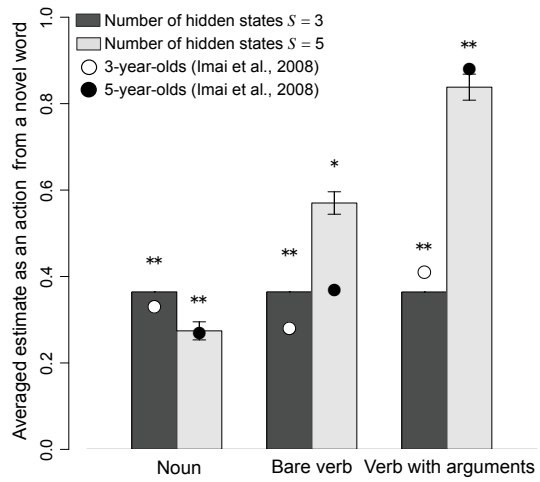

(c) Chinese input

Fig. 3. Estimation of a perceptual target indicated by a novel word. The vertical bars represent a standard error. The stars denote significant differences (*: $p < .05$, **: $p < .01$) between model estimation and chance (0.5). The points indicate the results of choices for children reported by Imai et al. [3]

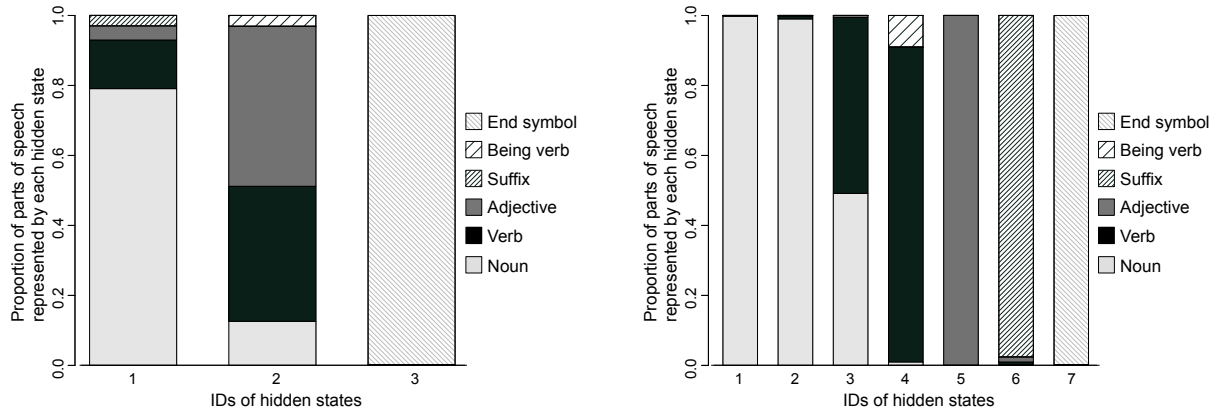

(a) A number of hidden states is 3

(b) A number of hidden states is 7

Fig. 4. A typical representation of syntactic categories in English input

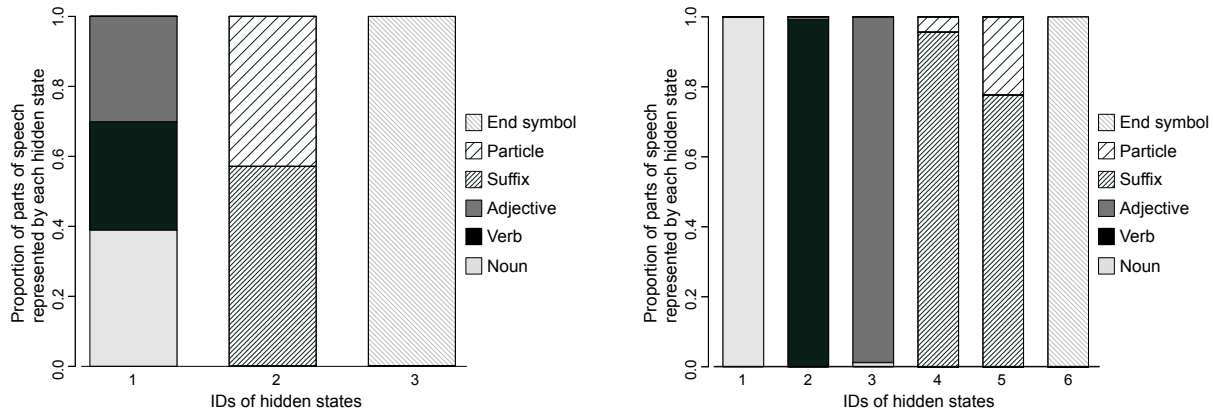

(a) A number of hidden states is 3

(b) A number of hidden states is 6

Fig. 5. A typical representation of syntactic categories in Japanese input

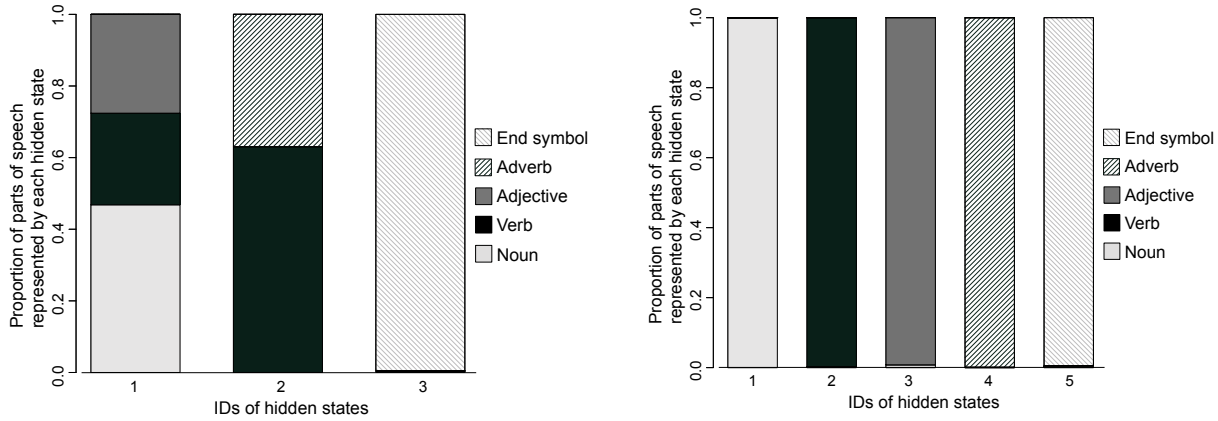

(a) A number of hidden states is 3

(b) A number of hidden states is 5

Fig. 6. A typical representation of syntactic categories in Chinese input

(ID 5 in Fig. 5(b)). This structure of the syntactic categories for suffixes enabled the Japanese model to distinguish a verb from a noun and thus to infer a novel verb as an action in the bare verb condition. In the Chinese model (Fig. 6), parts-of-speech were differentiated as S increased. Chinese, however, does not have suffixes like Japanese, resulting in failure to generalize bare verbs. In contrast to the other languages, the Chinese syntactic category (ID 4 in Fig. 6(b)) represented progressive words (adverbs). A progressive word is an important syntactic cue in the Chinese language, whereas only Chinese bare verbs do not have distinctive information between nouns and verbs.

IV. Discussion

We proposed a model that infers a hidden state (a syntactic category) of a novel word by using a BHMM and judges a visual target indicated by a novel word based on a syntactic category. We represented children's syntactic development as an increase in the number of hidden states of a BHMM. Our simulations showed that the estimations of a target by the proposed model were consistent with the children's inferences reported by Imai et al. [3]. We note that the model explains the dependence on a native-language in the bare verb generalization by five-year-olds. Further analysis of the typical representations of the hidden states clarified structures of syntactic categories underlying the developmental phenomena in the nouns and verbs generalization task. English-, Japanese-, and Chinese-language structures cause syntactic categories to depend on word orders, suffixes, and adverbs (progressives), respectively. These syntactic categories that reflect the language structure produce different performances among the input languages in the verbs generalization task.

These results suggest that syntactic understanding ability from three to five years of age is realized through simple statistical learning of sequential rules, similar to how a BHMM learns. The development of this ability is due to an improvement in the representative capacity of hidden states, that is, a differentiation process of representations of syntactic categories. However, in our model S is set in advance, and the mechanism behind increases in S remains unidentified. Findings that children and their caregivers' syntax of verbs become similar [13] indicate that language input to children is one of the triggers of their syntactic development. What specific language input leads to children's syntactic development is as yet obscure. Recently, a non-parametric method that can automatically optimize S of a BHMM was proposed [14]. We plan to investigate the characteristics of language input contributing to an increase of S with such a method.

Our model explains children's developmental change in noun and verb generalization by syntactic mechanism i.e., the differentiation of syntactic categories. Imai et al. [3] found that three-year-olds could successfully generalize a novel noun but failed to generalize a novel verb regardless of their native language. Our model suggests an explanation of these results: Nouns have more syntactic cues to specialize a syntactic category, e.g., a starting word in a sentence, than verbs, irrespective of language structures. In contrast, Imai et al. [3] suggested that the reason for the phenomena is that it is more difficult to parse actions than objects from the environment [4], [15]: The perceptual boundaries of objects are clearer than those of actions. The effect of perceptual development on performance in the generalization task is, however, not clear. Waxman et al. [16] reported that three-year-olds can generalize a novel verb if a known word, rather than a pronoun, is used as an object word in the same task as Imai et al. [3], [4]. Therefore, they concluded that the verb generalization failure by three-year-olds is caused by difficulty with syntactic understanding rather than perceptual immaturity [16]. Our model can reproduce children's performance development in the nouns and verbs generalization task without perceptual development, supporting the Waxman et al. [16]'s position that the developmental change is attributed to syntactic development. The ability to parse or code perceptual targets is a very important factor in vocabulary learning. The perceptual space in our model is constant and discrete. It would be interesting to investigate the role of differentiation of this perceptual space in syntactic development and generalization of a novel word in future research.

V. Conclusion

We proposed a model that can explain children's developmental change and the dependence on their native language in the nouns and verbs generalization task reported by Imai et al. [3]. Our model suggests two possibilities: the first is differentiation of children's syntactic categories underlying the developmental phenomena in the task, and the second is that acquisition of syntactic categories depending on language structure leads to the different performance outcomes for the task among languages. Our analysis of the model's hidden states clarified that syntactic categories are acquired based on the order of words in English, suffixes in Japanese, and adverbs (progressives) in Chinese. It is noted that a single model acquired these structures of syntactic categories in multiple languages. A further study is planned to verify the hypothesis deduced from the results in our current study.

Appendix

The proposed model learned artificial corpora reflecting linguistic characteristics of English, Japanese or Chinese. Figs. 7 show transition rules between word categories for sentence productions of each language. All sentences start at BOS (beginning of sentence) and then finish at EOS (end of sentence). The blue and orange decimals in the figures mean transition probabilities of word categories and probabilities that a category (i.e., a light blue circle) is dropped off, respectively. There is a half chance that an adjective is arranged before a noun in all corpora. Words are selected out of candidates in word categories in a random manner. Table I indicates numbers of the candidates. In sentences with a verb, there are three tenses: present, past and progressive forms, which are equally selected.

The transition probabilities are designed so that syntactic characteristics of the artificial corpora match with those of

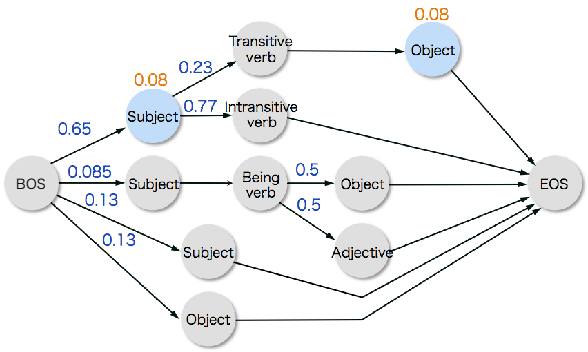

(a) English rule

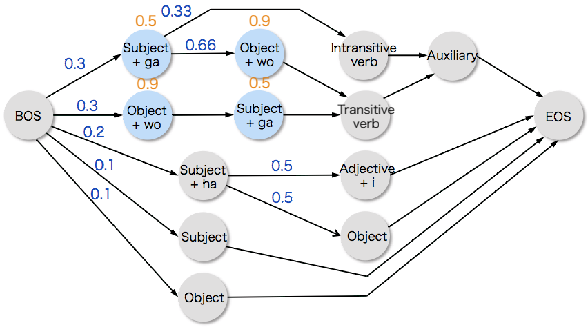

(b) Japanese rule

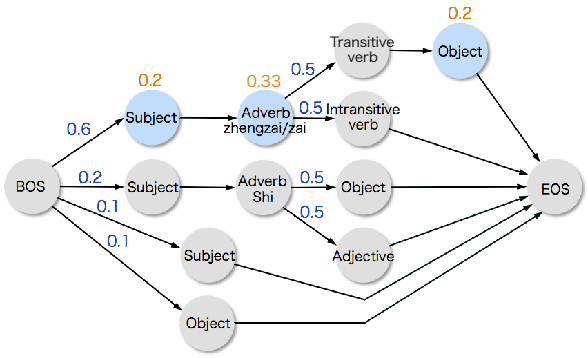

(c) Chinese rule

Fig. 7. Generation rules for (a) English, (b) Japanese and (c) Chinese corpora. BOS and EOS stand for beginning and end of a sentence, respectively. Blue decimals above arrows are transition probabilities. Orange decimals above blue circles are probabilities that the item is dropped off. An adjective is placed before a noun (a subject and object word) with the probability of one half in all rules. An indefinite article is placed before a countable noun (or an adjective) with the probability of one half in the English rule.

TABLE I
THE NUMBERS OF CANDIDATES OF WORD CATEGORIES IN EACH CORPUS.

	English	Japanese	Chinese
Subject	41	45	56
Object	37	45	56
Adjective	35	34	34
Transitives	20	21	20
Intransitives	14	15	14
Article	2	–	–
Auxiliary	–	3	–
Adverb (zhengzai/zai)	–	–	2

real corpora of child-directed speech [9]–[12]. We calculated percentages of occurrences of predicative verbs, subjects, objects and single sentences consisting a noun in the real corpora, and made percentages of the artificial corpora correspond with the real ones. The percentages of single utterances consisting "Verb-ing" in the artificial corpus is equal to the real one [9]. Furthermore, the artificial English corpus has the same ratios of three construction types: fragments (sentences without subject and predicate), copulas (sentences in which the verb is some form of to be) and sentences with both a subject and a single predicative intransitive as those of real child-directed utterances reported by [12].

ACKNOWLEDGMENT

The authors gratefully acknowledge the advice by Prof. Mutsumi Imai of Keio University. This work was supported by JSPS Grants-in-Aid for Specially Promoted Research (24000012) and Grants-in-Aid for JSPS Fellows (13J00756).

REFERENCES

[1] M. D. S. Braine and M. Bowerman, "Children's first word combinations," *Monographs of the society for research in child development*, vol. 41, no. 1, pp. 1–104, 1976.
[2] R. Olguin and M. Tomasello, "Twenty-five-month-old children do not have a grammatical category of verb," *Cognitive Development*, vol. 8, no. 3, pp. 245–272, 1993.
[3] M. Imai, L. Li, E. Haryu, H. Okada, K. Hirsh-Pasek, R. M. Golinkoff, and J. Shigematsu, "Novel noun and verb learning in chinese-, english-, and japanese-speaking children," *Child Development*, vol. 79, no. 4, pp. 979–1000, 2008.
[4] M. Imai, E. Haryu, and H. Okada, "Mapping novel nouns and verbs onto dynamic action events: Are verb meanings easier to learn than noun meanings for japanese children?" *Child Development*, vol. 76, no. 2, pp. 340–355, 2005.
[5] J. Elman, "Finding structure in time," *Cognitive science*, vol. 14, no. 2, pp. 179–211, 1990.
[6] M. Morifuji and T. Inui, "A connectionist model of vocabulary acquisition based on development of grammatical knowledge," in *Conference on Neuro-Computing and Evolving Intelligence*, 2004.
[7] A. Toyomura and T. Omori, "A computational model for taxonomy-based word learning inspired by infant developmental word acquisition," *IEICE transactions on information and systems*, vol. 88, no. 10, pp. 2389–2398, 2005.
[8] S. Goldwater and T. Griffiths, "A fully bayesian approach to unsupervised part-of-speech tagging," in *Annual meeting-association for computational linguistics*, vol. 45, no. 1, 2007, p. 744.
[9] A. Henry, *English Belfast Corpus*. PA: TalkBank, 2004, ISBN 1-59642-037-5.
[10] S. Miyata, *MiiPro Nanami Corpus*. PA: TalkBank, 2013, ISBN 1-59642-473-7.
[11] T. Tardif, *Chinese Beijing2 Corpus*. PA: TalkBank, 2007, ISBN 1-59642-287-4.
[12] T. Cameron-Faulkner, E. Lieven, and M. Tomasello, "A construction based analysis of child directed speech," *Cognitive Science*, vol. 27, no. 6, pp. 843–873, 2003.
[13] S. Choi, "Early development of verb structures and caregiver input in korean: Two case studies," *International Journal of Bilingualism*, vol. 3, no. 2-3, pp. 241–265, 1999.
[14] Y. W. Teh, M. I. Jordan, M. J. Beal, and D. M. Blei, "Hierarchical dirichlet processes," *Journal of the american statistical association*, vol. 101, no. 476, pp. 1566–1581, 2006.
[15] D. Gentner, *Why nouns are learned before verbs: Linguistic relativity versus natural partitioning*, S. A. Kuczaj, Ed. Erlbaum, Hillsdale, 1982.
[16] S. R. Waxman, J. L. Lidz, I. E. Braun, and T. Lavin, "Twenty four-month-old infants interpretations of novel verbs and nouns in dynamic scenes," *Cognitive Psychology*, vol. 59, no. 1, pp. 67–95, 2009.