

Adaptive Robot Language Tutoring Based on Bayesian Knowledge Tracing and Predictive Decision-Making

Thorsten Schodde
CITEC, Bielefeld University
Bielefeld, Germany
tschodde@techfak.uni-
bielefeld.de

Kirsten Bergmann
CITEC, Bielefeld University
Bielefeld, Germany
kirsten.bergmann@uni-
bielefeld.de

Stefan Kopp
CITEC, Bielefeld University
Bielefeld, Germany
skopp@techfak.uni-
bielefeld.de

ABSTRACT

In this paper, we present an approach to adaptive language tutoring in child-robot interaction. The approach is based on a dynamic probabilistic model that represents the interrelations between the learner's skills, her observed behaviour in tutoring interaction, and the tutoring action taken by the system. Being implemented in a robot language tutor, the model enables the robot tutor to trace the learner's knowledge and to decide which skill to teach next and how to address it in a game-like tutoring interaction. Results of an evaluation study are discussed demonstrating how participants in the adaptive tutoring condition successfully learned foreign language words.

CCS Concepts

•Computing methodologies → Probabilistic reasoning; Cognitive robotics; •Applied computing → Interactive learning environments; •Human-centered computing → Empirical studies in HCI;

Keywords

Language tutoring; Education; Assistive robotics; Bayesian Knowledge Tracing; Decision making

1. INTRODUCTION

The use of robots for educational purposes has increasingly moved into focus in recent years. This is due to two major developments. First, robots became cheaper and more robust so that applications in everyday environments are now conceivable. In particular, technology has matured up to a point where intuitive interaction using natural language or gesture has become feasible. Second, the need for second language learning becomes increasingly important, and empirical evidence has demonstrated that learning with and from a physically present, interactive robot can be more effective than learning from classical on-screen media [14, 15, 20, 22]. In fact, recent research showed that performance

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HRI '17, March 06 - 09, 2017, Vienna, Austria

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4336-7/17/03...\$15.00

DOI: <http://dx.doi.org/10.1145/2909824.3020222>

can increase up to 50% [17]. It can, hence, be assumed that tutoring using social robots is qualitatively different from alternative digital tutoring technologies. Nowadays, first practical applications can be found, e.g. in nursery where toy robots teach the alphabet to kids in a very simple way. More generally, findings from a variety of settings seem to suggest that robots can help small children to develop in an educational setting [10, 18, 24, 27].

In the L2TOR project¹, we investigate in how far a social robot can support children at pre-school age with respect to second language learning. Learning a language is a very complex task. It involves not only acquiring vocabulary, but also learning prosodic features, syntactical structures, semantic meanings as well as situation-dependent language use. Yet, it has been argued that social robots can create the interactive environment and motivational experience needed to learn languages [19].

One of the most important aspects in tutoring is the robot's ability to keep track of the knowledge state, i.e. the learned and not-yet-learned skills, of the child interacting with it. This information is indispensable to enable a personalized tutoring interaction and to optimize the learning experience for the child [27]. The tutor has to structure the tutoring interaction, choose the skills to be trained, adjust the difficulty of the learning tasks appropriately and has to adapt its verbal and non-verbal behaviour.

The importance of personalized adjustments in the robot's behaviour has been substantiated in recent research showing that participants who received personalized lessons from a robot (based on heuristic skill assessment) outperformed others who received a non-personalized training [22]. Sub-optimal robot behaviour (e.g. too much, too distracting, mismatching or in other ways inappropriate) can even hamper learning [17]. In this paper we present an integrated approach for tracing the knowledge of the learner during a L2 learning interaction together with a strategic adaptation of tutoring actions.

In the following, we discuss related work in Section 2. In Section 3 an extension of Bayesian Knowledge Tracing is presented as well as a model to select the next tutoring actions based on the predicted effects they may have on the learner's knowledge state. This model has been implemented in a robot that provides language tutoring in a game-like fashion. Section 4 introduces the empirical basis for this scenario and observational studies on language tutoring in kindergarten. Section 5 presents an evaluation study carried out with this robot and Section 6 discusses the results.

¹<http://www.l2tor.eu>

2. RELATED WORK

Numerous studies have investigated the effects of social robots in tutoring scenarios. Empirical evidence demonstrates that learning with and from a physically present, interactive robot can be more effective than learning from classical on-screen media [14, 15, 20], and that robots can help children to develop in educational settings [10, 18, 24, 27]. However, at the same time, it is found that suboptimal behaviour of the robot can hamper learning [17]. Thus, a crucial ingredient for successful robot tutoring is the ability to provide personalized lessons [22] and to adapt in appropriate ways to the needs of the learner. The key question is when and how to adapt robot tutoring, according to which adaptation strategies, and based on what features of the state of learner or the tutoring interaction.

2.1 Approaches to Adaptive Tutoring

In the realm of Intelligent Tutoring System (ITS), dedicated *pedagogical modules* are employed for planning an optimal path through the curriculum by using an internal model of the learner’s present knowledge state (cf. [8]). Cakmak and Lopes [3], for example, proposed a teaching algorithm that selects the most informative demonstrations for the learner. This learning agent makes use of Inverse Reinforcement Learning (IRL) to reduce the learner’s hypothesis space of possible reward functions as fast as possible. In an evaluation, the authors showed that a learner trained with non-optimal selected expert demonstrations require significantly more demonstrations to achieve a similar performance as the optimally taught learner. This system, however, is designed for a sequential decision task in which no uncertainty about the learner’s knowledge/skill exists. This assumption does not hold for the domain of L2 learning where the learner’s current state of knowledge can, at best, be inferred from observed behaviour. Another important limitation of this approach is a lack of flexibility as no adaptation towards students’ individual needs is considered.

Addressing especially the issue of adaptation towards students’ individual needs, Partial Observable Markov Decision Processes (POMDPs) have been employed as basis for the pedagogical module of an ITS. Rafferty et al. [25], for instance, proposed different algorithms for planning an action-policy based on a POMDP and compared these against two different random and a maximum information gain (MIG) choice. They showed that even a simple action-policy based on a POMDP can achieve a significant faster skill learning than choosing actions randomly. But compared to the simple MIG algorithm, no significant difference was observed. Only with increasing skill space the MIG algorithm seems not to be sufficient anymore. A likely explanation for this finding is that the knowledge tracing model is insufficient. In addition, finding a good policy based on a POMDP is often computationally intractable.

Clement et al. [4] compared two algorithms choosing the next skill and action in a tutoring interaction against a lesson given by a human expert. Both algorithms based on prior knowledge, e.g. the impact of actions on the learning gain or the difficulty of different types of tasks, which had been annotated by experts beforehand. The algorithms differed with regard to the adaptation method and the amount of additional knowledge stored besides the prior. The authors showed that even if the ITS does not make use of an internal model to store beliefs about the child’s knowledge

state regarding a specific skill, the use of their algorithm can lead to a higher learning gain compared to an expert lesson. Furthermore their second proposed algorithm, which additionally stores information about the knowledge state of the child, performed even better. Clement et al. concluded that extending their system with a more complex model for tracing the knowledge state of a student might lead to a higher learning.

An often criticized issue in this line of research is the lack of an effective *knowledge-tracing* method in the pedagogical module of an ITS that could be profitable for the learning interaction, e.g. by increasing the students’ learning gain. Hence, we review research on knowledge tracing methods in the following.

2.2 Knowledge Tracing

Knowledge tracing aims to model learners’ mastery of the knowledge being tutored. An often used approach is *Bayesian Knowledge Tracing* (BKT). BKT is a specific type of *Dynamic Bayesian Networks* (DBN), or more precisely, of Hidden Markov Models consisting of observed and latent variables. The latent variables represent the ‘skills’ and are classically assumed to be binary. That is, a skill is represented to be mastered or not. Generally, separate BKT networks are used for each skill to be learned [5]. Belief update is based on the observation of an answer to a given task testing a specific skill. The observed answer is binary too. Further, BKT models have two types of parameters: The emission probability and the transition probability. The emission probabilities are given by the ‘guess probability’ $p(guess)$, the probability of answering correctly without knowing the skill, and the ‘slip probability’ $p(slip)$ of answering wrongly although knowing the skill. In contrast, the transition probabilities are given by $p(t)$, the skill transition from unknown to known, and $p(f)$ the probability of forgetting a previously known skill. Often $p(f)$ is assumed to be zero.

Spaulding et al. [29] recently adopted BKT to trace the language-reading skill of children in robot-based language tutoring. They proposed the ‘Affective BKT model’, which is characterized by two further observable variables called ‘smile’ and ‘engagement’ to take into account the affective state of the child. This model structure allows emotions to influence the belief-state of each skill as they are included in every belief-update. The authors showed that the affective state of the children can be successfully integrated into BKT and that this approach outperforms traditional models for tracing the knowledge state in learning situations [29].

Another modification of BKT was published by Käser et al. [16]. Instead of using a dedicated BKT for every skill, they defined one comprehensive DBN to trace the knowledge on all skills to be learned. This enables to trace the knowledge on each skill individually and, in addition, to represent and reason with skill inter-dependencies. This allows for searching some kind of order in which skills may be learned best. The authors could demonstrate that this more detailed model outperforms other traditional models of knowledge tracing, including the normal BKT, with regard to the accuracy of the skill belief [16].

Finally, Gordon et al. [11] recently presented a so-called ‘active learner model’ to trace the word-reading skills of small children. This model does not work on the basis of BKT but employs a simple distance metric to approximate the conditional probability $p(w_2|w_1)$ of whether the child

can read a word w_2 if it already knows the word w_1 . Their evaluation showed that their system is able to adapt to users of different age and to trace their reading knowledge up to a certain extent [11].

In this paper we present an expandable model based on BKT for knowledge tracing that, in contrast to the systems reviewed above [16, 29], allows for the simulation of actions and decision-making in teaching interactions.

3. ADAPTIVE LANGUAGE TUTORING

As a basis for our approach to adaptive language tutoring, we adopt the Bayesian Knowledge Tracing model [5] which has been successfully employed in other work and was shown to be easily extensible. However, we modify and extend the BKT model in order to enable predictive decision-making based on the represented beliefs about the learner’s knowledge state. In this section, we first introduce our version of BKT and then present the approach for decision-making.

3.1 Bayesian Knowledge Tracing

The traditional approach to BKT uses only one latent variable S to represent the skill belief and one observable variable O for the user’s answer. This suffices to represent if a skill is mastered or not, and how probable it would be that the user will answer correctly. Also, this information can be used to choose the next skill to learn, e.g. the skill which has the lowest belief probability of having been mastered. However, this model does not include information about how a skill can be addressed for teaching. In consequence, there is no possibility to take possible actions and their influence on the update of skill beliefs into account. We thus add a decision node A for actions to the Bayesian network (see Figure 1). This node not only influences the possible observation but also the belief update in the next time step. Further, we use a latent variable S that can attain six discrete values for each skill, corresponding to six bins for the belief state (0%, 20%, 40%, 60%, 80%, 100%). This allows for a more detailed model of the impact of tutoring actions on the possible observations and skills. Moreover, it becomes possible to better quantify the robot’s uncertainty about the learner’s skill.

With these changes, especially the conditional probability table $p(O^t|S^t)$ and the additional influence of the action A^t on the observable (now $p(O^t|S^t, A^t)$), the classical BKT update function, which was based on simple assumptions about guessing $p(guess)$ and slipping $p(slip)$ during the answer process, cannot be applied anymore. Instead, we apply a normal Bayesian update rule for the conditioning of skill beliefs including a transition probability $p(S_i^{t+1}|s_k, O^t, A^t)$ where s_k identifies a bin of the skill S_i^t . As a simplification we substitute this probability with $p(S^{t+1}|s_k)$:

$$\begin{aligned} p(S_i^{t+1}) &:= p(S_i^{t+1}|O^t, A^t) \\ &= \sum_{s_k \in S_i^t} [p(S_i^{t+1}|s_k, O^t, A^t) \cdot p(s_k|O^t, A^t)] \\ &\approx \sum_{s_k \in S_i^t} \left[\frac{p(O^t|s_k, A^t) \cdot p(A^t|s_k) \cdot p(s_k)}{p(O^t, A^t)} \cdot p(S_i^{t+1}|s_k) \right] \end{aligned}$$

3.2 Predictive Decision-Making

The extended BKT model is used to decide which tutoring action the robot should take next. At first, the skill to




Figure 1: Dynamic Bayesian Network for BKT: The action node A^t predicts the observation O^t and influences the belief update of S^t for the next time step $t + 1$.

address with the next tutoring action is chosen. For this, the Kullback-Leibler divergence (KLD) between the current skill belief and the desired skill belief is used, the latter being a maximally certain belief in a maximally high skill of the learner:

$$next_skill = \underset{\forall S_i^t \in \mathbb{S}}{\operatorname{argmin}} [\alpha(S_i^t) \cdot KLD(p(S_i^t), p(S_{opt}))]$$

\mathbb{S} represents the set of all skills that can be addressed, which consists of all words to be taught to the user. $p(S_{opt})$ is the desired belief for each skill, which means 99.999% of probability mass in the last bin (100%). The factor $\alpha(S_i^t)$ has been added for each skill to regulate the skill occurrence frequency. It is decreased each time the skill is addressed, and it is increased if another skill is being practised. In this way, the skill-selection algorithm takes care of the maximization of each skill belief as well as the balancing of all skills.

After the skill has been chosen, the next step is to decide with which tutoring action this should be done. Here, we consider abstract tasks as tutoring actions. These tasks will have to be mapped onto concrete exercises or pedagogical acts at a later stage in the robot control architecture (see Section 4). For simplicity, we distinguish between tutoring actions according to the difficulty (easy, medium or hard) of the task that addresses the corresponding skill. Finding the best action a_l for a given skill S_i^t is thus a minimization problem of the following form:

$$next_action = \underset{\forall a_l \in A^t}{\operatorname{argmin}} [\alpha(a_l) \cdot KLD(p(S_i^{t+1}), p(S_{opt}))]$$

where

$$\begin{aligned} p(S_i^{t+1}) &:= p(S_i^{t+1}|a_l) \\ &= \sum_{s_k \in S_i^t} \sum_{o_j \in O^t} p(S_i^{t+1}|o_j, s_k, a_l) \cdot p(o_j, s_k|a_l) \\ &\approx \sum_{s_k \in S_i^t} p(s_k|a_l) \sum_{o_j \in O^t} p(S_i^{t+1}|s_k) \cdot p(o_j|s_k, a_l) \end{aligned}$$

with

$$p(o_j, s_k|a_l) = p(o_j|s_k, a_l) \cdot p(s_k|a_l)$$

Here, $p(S_i^{t+1})$ could be seen as predicting the effect of applying the current action a_l to the skill S_i , where we again

substitute the transition probability $p(S_i^{t+1}|o_j, s_k, a_t)$ with $p(S_i^{t+1}|s_k)$ regarding simplicity. In addition, here again the skill belief is compared with $p(S_{opt})$ which represents the desired tutor belief state for each skill. The factor $\alpha(a_t)$ provides a more detailed selection of the “best” action. This way, the model will select an easy task if the skill is believed to be low, a hard task if it is high, and medium in-between. The goal of this strategy is to create a feeling of flow which can lead to better learning results [2, 7, 12]. Thus, it strives not to overburden the learner with too difficult tasks nor to bore him with too easy tasks, both of which may lead to frustration and thus hamper the learning [9, 13].

4. ROBOT LANGUAGE TUTORING

The adaptive model as described in the previous section has been brought to application in a child-robot second language (L2) tutoring game on the basis of empirical data from adult-child language tutoring interactions.

4.1 Empirical Basis

To design a tutoring interaction that matches children’s needs, we decided to design the interaction on an empirical basis of language tutoring data. We collected video recordings of language tutoring games as they take place in kindergartens. Given that 1:1 interactions of educator and child can hardly be realized in kindergartens, the games typically involve one educator and a small group of children. Data of four language tutoring games have been collected: reading a picture book together with children in an interactive manner; card game “I spy with my little eye”; card game “I’m giving you a present”; and a rhyming game. The children were between four and six years of age. The data collected comprises round about 681 min of video data. These video data have been transcribed and annotated with regard to the following categories:

- **Dialogue acts:** Utterances are classified due to the underlying intention based on the DAMSL annotation scheme [6].
- **Children’s mistakes:** Types of language errors the children made, e.g. wrong plural form, missing articles, wrong syntax, etc.
- **Educator’s speech repair:** Pedagogical acts used to correct the errors, e.g. reformulation, corrected repetition, etc.
- **Nonverbal behaviour:** Nods, smiles, gestures etc. used by the educators.

On the basis of these annotations, we identified some overall patterns to inform the detailed design of the robot’s behaviour. These fall basically into two categories, (i) overall interaction structure and (ii) feedback behaviour by the educators.

4.1.1 Overall Interaction Structure

A common pattern in all language tutoring games under investigation was the following basic course of actions:

1. **Opening:** Marks the beginning of the interaction and should mitigate the children’s timidity as well as it should motivate the child.

2. **Game Setup:** This step is used to prepare the game by explaining the task and clarify the necessary terms.
3. **Test run:** A test run of the game is conducted to test whether the instructions have been understood and to practice the game flow.
4. **Game:** Here, the main interaction game takes place. Every move is accompanied by the educator’s feedback and motivations to continue.
5. **Closing:** Marks the end of the learning interaction. Additionally, it is used to ensure motivation for future interactions by acknowledge the participation, joint singing a goodbye song and an outlook on what’s going to happen next time.

4.1.2 Educator’s Feedback Behaviour

In addition, we analysed the educators’ behaviour when providing children with feedback. An important and common pattern is that language errors are almost never corrected explicitly. Instead, feedback is always provided in a positive way, falling into one of the following categories with the percentage of their occurrence given in squared brackets: (i) **praising the child** for a correct utterance whereby praise is often combined with a repetition of the correct word [13%] (ii) **implicit correction** in case of an error made by the child: repetition of the word as if correct (e.g. correct pronunciation, with article, plural form, etc.) [54%], (iii) **correct recasting of a sentence**, e.g. after syntax errors [32%], (iv) **moving on to next task**, e.g. when the child’s message is unclear due to incomprehensible pronunciation [1%]. All kinds of educators’ feedback behaviour is typically accompanied by looking at the child, smiling and nodding.

4.2 Game Setup

We have chosen the game “I spy with my little eye...” as a paradigm for our child-robot language tutoring game. The robot – in the role of a tutor, assisting the child in learning novel L2 vocabularies – is acting as ‘the spy’. The child-robot setting is further enriched with a tablet PC on which objects are displayed (see Figure 2). In addition, the tablet’s touch-screen displays three buttons to enable further user input in terms of ‘yes’ and ‘no’ answers as well as the option to let the robot repeat its previous statement.

A basic move of the game is structured as follows: It starts with a set of objects being displayed on the tablet screen and the robot saying “I spy with my little eye, something that is ...”, followed by a foreign language word that refers to a property of one of the items on the screen. The child’s task is now to respond by selecting the object referred to via touch input on the tablet. The robot’s feedback behaviour in response to a correct or false answer is realized on the basis of our empirical data (see Section 4.1.2). That is, the robot responds to correct answers by praising the learner as well as repeating the L2 word and the corresponding L1 translation. In case of a false guess by the child, the robot explains the correct meaning of the to-be-learned word one more time. In addition, the wrongly chosen object as well as the actually correct object are both displayed on the tablet screen and the child is asked to select the correct object. The overall game structure is framed by the other elements making up typical language tutoring games in adult-child interaction (see Section 4.1.1).

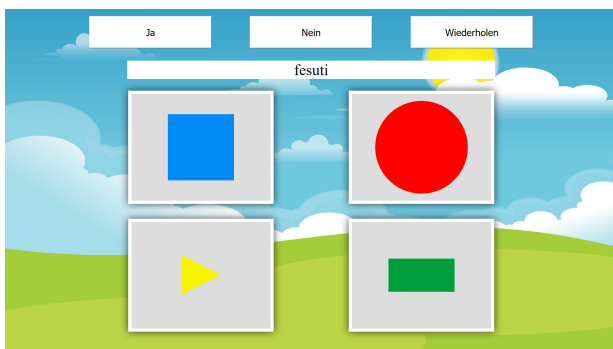


Figure 2: Experimental setup (left) with a participant sitting in front of a tablet displaying the graphical user interface (right). The robot *Nao* stands next to the tablet slightly rotated towards the user.

4.3 Technical Realization

We employed the Nao robot² for our language tutoring game. It is standing in a bit more than 90 degrees rotated, to the right of the participant. In addition a Microsoft Surface Pro 4³ tablet PC is used to catch the user input and to display the graphical user interface realized via a HTML website. For the implementation of the interaction and dialogue structure, the state-chart based dialogue-manager IrisTK has been used [28]. NAOqi⁴ has been applied as middleware between the robot, the graphical user interface, the dialogue manager, and our developed adaptive tutoring model. NAOqi is shipped with each Nao robot and allows to communicate via a simple event system between various programming-languages (Python, Java, C++, JScript).

5. EVALUATION STUDY

To assess the effects of our adaptive model on L2 word learning, we set up an evaluation study based on the language tutoring game described in the previous section. The major objective behind this study was to evaluate the effects of the adaptive model on learners' performance. We used the Nao robot to deliver all task information and direct feedback to the learner. This enables us to test the model within the desired final setting, including the effects of a robot's presence in the tutoring interaction. Given that children show a high degree of inter-individual variation and might further need child-specific adaptations of, for instance, synthesized speech to enable them to understand what the robot says, we decided to conduct this first study with adult learners.

We employed a between-subjects design with a manipulation of training type: Participants learned L2 vocabulary items either with the fully adaptive model, or in a random control condition. In the adaptive condition, the skill to be taught and the action to address the skill were chosen by the model as described in Section 3. In our language tutoring game, skill relates to the foreign language words and action refers to the specific task used in the game (target word, objects displayed). The difficulty of the actions/tasks in this study were implemented by using less or more distractor objects that were shown together with the correct

object on the screen. For instance, an easy task consisted of two distractor objects, whereas a hard task had four distractors. Distractors were chosen with respect to the skill beliefs of the user, with the set of objects mainly consisting of items for which the L2 words were still/mostly unknown by the learner.

As shown by Craig et al. [7], better learning performance is to be expected when learners have to expend the right amount of cognitive effort (i.e. not too hard or too easy tasks). Accordingly, while learning with our model in the adaptive condition, no hard tasks are shown until the system believes the user to have basic knowledge on all skills. Then, the system will increase task difficulty (as determined by the adaptive tutoring model) by adding distractor objects. Note, however, that at a certain point the user will know too many skills/words so that finding a distractor set (i.e. task difficulty) that cannot be sorted out by exclusion becomes impossible. In the control condition, all skills are taught in a random order and always with 'medium' task difficulty.

Participants' performance was assessed with two measures: (1) we tracked learners' response behaviour over the course of the training to investigate the progress of learning, (2) we conducted a post-test on the taught vocabulary in the form of both L1-to-L2 translations and L2-to-L1 translations to assess participants' state of knowledge subsequent to the intervention.

5.1 Materials

The training materials for the study comprised German-Vimmi word pairs. Vimmi is an artificial language created for experimental purposes [23] that aims to avoid associations with other known words or languages. The Vimmi items are created according to Italian phonotactic rules. Ten items have been chosen: four colour terms, four shape-encoding terms and two terms describing size (see Table 1).

5.2 Procedure

Upon entering the lab, participants were randomly assigned to one of the two experimental conditions. They were informed that they take part in an experiment on foreign language learning and were asked to sign an informed consent form. They also filled out a questionnaire that covered personal information like age and nationality as well as a personal estimation of language learning skills in general and memorization ability for L2 vocabulary.

²<https://www.aldebaran.com/en/cool-robots/nao>

³<https://www.microsoft.com/surface/en-gb/devices/surface-pro-4>

⁴<http://doc.aldebaran.com/2-1/naoqi/>

N	German	Vimmi	English translation
1	blau	bati	blue
2	grün	uteli	green
3	gelb	dirube	yellow
4	rot	fesuti	red
5	rund	beropuga	round
6	dreieckig	pewo	triangular
7	quadratisch	tanedila	square
8	rechteckig	paltra	rectangular
9	klein	kiale	small
10	groß	ilado	big

Table 1: The 10 words from Vimmi to be learned in the evaluation study with its corresponding translation in German as well as English for comprehension purposes.

Next, a list of the to-be-learned Vimmi items were presented to the participants for 30 seconds. This was to enable participants to practice the items right from the first game interaction on. Then, the learning interaction with the Nao robot began. After introducing itself, the robot explained the “I spy with my little eye”-game and started a test-run with the participants. Once this test run was finished and the participants agreed that (s)he understood the game, the main interaction consisting of a total of 30 trials (game moves) began. Each trial addressed one vocabulary item as described in Section 4.2. That is, the robot asked for one of the objects displayed on the tablet screen, whereby the question was in L1 (German) for the most part, except for the referring, to-be-learned word in L2 (Vimmi). After 30 trials, the game was finished, the Nao robot thanked the participants and said goodbye.

Subsequent to the interaction with the robot, participants’ learning performance was assessed with a post-test. In an interview with the experimenter, they had to translate the ten to-be-learned vocabulary items from German to Vimmi and likewise from Vimmi to German (both in randomized order). The whole interaction and the vocabulary-post-test at the end of the study were recorded with an external camera. Also the system decisions taken during the interaction and the probability distributions for each updated skill belief were logged.

5.3 Participants

A total of 40 participants (20 per condition) with an average age of 24.13 ($SD = 3.82$) took part in this study (16 males and 24 females). All participants had very good command of the German language and normal or corrected sight. All of them were paid or received credits for their participation.

5.4 Results

5.4.1 Learning Progress During Training

In order to assess the learners’ progress during training, we compared the number of correct responses addressing the initial quarter of the tutoring game (first seven items) against the final quarter (last seven items). When an item occurred repeatedly within the initial quarter, the first occurrence has been taken into account. When an item oc-

	Adaptive (A)		Control (C)		A, C	
	M	SD	M	SD	M	SD
F7	3.75	1.37	4.00	1.17	3.88	1.27
L7	6.90	0.31	5.15	1.69	6.03	1.49
F7, L7	5.33	0.69	4.58	1.12		

Table 2: Means (M) and standard deviations (SD) of correct answers for the initial quarter of the training interaction (first seven items – F7) and the final quarter (last seven items – L7) in each condition, as well as the inter-model (A, C) and intra-model (F7, L7) means and standard deviations.




Figure 3: Mean numbers of correct answers at the beginning (first 7) and end (last 7) of the interaction in the different conditions.

curred repeatedly within the final quarter, the last occurrence has been considered.

A mixed-design ANOVA with training phase (initial, final) as a within-subjects factor and training type (adaptive-model-based, control) as between-subjects factor has been conducted. Results are summarized in Table 2 and Figure 3. Not surprisingly, there was a main effect of training phase at a significant level ($F(1, 38) = 66.85, p < .001, \eta^2 = .64$): Learners’ performance was significantly better in the final phase as compared to the initial phase. In the first quarter of training, participants achieved a mean of 3.88 ($SD = 1.27$) correct responses, whereas in the final quarter, a mean of 6.03 ($SD = 1.49$) items was selected correctly. More interestingly, there was also a main effect of training type ($F(1, 38) = 6.52, p = .02, \eta^2 = .15$) such that participants who learned in the adaptive condition had a higher score of correct answers ($M = 5.33, SD = .69$) as compared to learners in the control condition with an average of $M = 4.58$ ($SD = 1.12$) correct answers. Finally, the interaction between training phase and training type was also significant ($F(1, 38) = 14.46, p = .001, \eta^2 = .28$) indicating that the benefit of adaptive-model-based training develops over time (see Figure 3). While participants’ response behaviour in the first quarter of training was similar across conditions, a benefit of training with the adaptive model became evident in the final quarter. At this stage of training, participants in the adaptive model condition achieved a mean of $M = 6.9$ ($SD = .31$) correct responses, whereas participants in the control condition achieved a mean of $M = 5.15$ ($SD = 1.69$) correct responses.

	Adaptive (A)		Control (C)	
	M	SD	M	SD
German-to-Vimmi	3.95	2.56	3.35	1.98
Vimmi-to-German	7.05	2.56	6.85	2.48

Table 3: Results of both post-tests (German-to-Vimmi and Vimmi-to-German): Means (M) and standard deviation (SD) of correct answers grouped by the experimental conditions.




Figure 4: Participant-wise amount of correct answers grouped by the different conditions for the German-to-Vimmi post-test.

5.4.2 Post-Test

Participants’ learning performance subsequent to the intervention has been measured with two translation tests (L2-to-L1 and L1-to-L2). Results are summarized in Table 3. Paired-samples t-tests were conducted to compare the number of correctly recalled words after training with the adaptive model as compared to training in the control condition. For the German-to-Vimmi translation, there was no significant main effect ($T(38) = .25, p = .80$). Participants who trained with the adaptive-model recalled a mean of 3.95 ($SD = 2.56$) out of ten words correctly, while participants in the control condition recalled a mean of 3.35 ($SD = 1.98$) words. Likewise, there was no significant main effect ($T(38) = .83, p = .41$) for the Vimmi-to-German translation task. Participants’ performance after learning with the adaptive model amounted to a mean of 7.05 ($SD = 2.56$) correct items, participants’ performance in the control condition to a mean of 6.85 ($SD = 2.48$) correct items.

Although no main effect of training type emerged in the post-test, some details might nevertheless be worth mentioning. In the German-to-Vimmi post-test, a maximum of ten correct responses was achieved by participants in the adaptive-model condition, whereas the maximum on the control condition were six correct answers. Moreover, there were two participants in the control condition who did not manage to perform any German-to-Vimmi translation correctly. In the adaptive-model condition, all participants achieved at least one correct response (see Figure 4).

6. CONCLUSION

In this paper we have presented a novel approach to personalize language tutoring in human-robot interaction. This adaptive tutoring is enabled through a model of how tutors

mentalize about learners – by keeping track of their knowledge state and by selecting the next tutoring actions based on their likely effects on the learner. This is realized via an extended model that combines Bayesian Knowledge Tracing (of the learned) with tutoring actions (of the tutor) in one causal probabilistic model. This allows, for selecting skills and actions based on notions of optimality – here the desired learner’s knowledge state as well as optimal task difficulty – to achieve this for a given skill. This model has been implemented into a robot language tutoring game and tested in a first evaluation study.

The analysis of participants’ response behaviour over the course of training has clearly shown that participants learned the L2 words during the human-robot interaction. Importantly, they learned more successfully with our adaptive model as compared to a randomized training. That is, the repeated trials addressing still unknown items as chosen by the adaptive model (until the belief state about these words equalled that of known items) outperformed the tutoring of the same material (same number of trials and items) but in randomized order. In the post-test, however, there was no significant difference across experimental conditions, despite a trend towards better performance in the adaptive model conditions over the controls.

Different explanations may account for this inconsistent finding. One potential explanation could be that the way how responses were prompted was not identical in training sessions and post-test. In the training sessions, participants saw pictures reflecting the meaning of the to-be-learned words whereas in the post-test they just received a linguistic cue in form of a word they had to translate. It might be that repeated trials as they were particularly supported for difficult-to-remember items by the adaptive model, led to stronger associations between linguistic and imagistic materials. This might have caused a stronger decline of correct responses for participants who trained with the adaptive model as opposed to those in the control condition. An alternative explanation could be that test results measured immediately after the training session are subject to strong inter-individual differences among learners. This is the reason why studies on vocabulary learning usually range over repeated sessions spread over several days (cf. 1). A typical pattern is that significant results emerge after two or three sessions/days and/or in the long-term (measured several weeks after training took place). So it might well be that further training sessions or delayed tests might result in a post-test performance that matches the picture of the during-session performance.

One might argue that the performance of our adaptive model is comparable to the vocabulary learning technique of *spaced repetition* as implemented, for instance, in the Leitner system [21]. In this system flashcards are sorted into groups according to how well the learner knows each one. Learners try to recall items written on a flashcard. If they succeed, the card is sent to the next group. If they fail, the card is sent back to the first group. Each succeeding group has a longer period of time before the learner is required to revisit the cards. This way all items, that are hard to remember for the learner will be repeated more often. In contrast to such spaced repetition systems, our model is more flexible as it can vary the difficulty of the tasks by providing more or less distractor items. In addition, we plan a more comprehensive action space of the model to account for motivating actions

where necessary or adaptations in the robot’s verbal or non-verbal behaviour.

Overall, results from the evaluation study are, at least, in parts very promising: learners’ performance during training was significantly improved by personalized robot tutoring based on the adaptive model. Nevertheless, the fact that this positive effect did not hold in the post-test, *inter alia*, marks a starting point for further refinements of the model: Training stimuli should be designed such that they match the way language learners need to apply them best possible. That is, when the aim is to enable people to translate words from one language to another, training stimuli should provide cues for this process of mapping linguistic materials on each other. Moreover, a further study with more learning sessions (e.g. over several days as common in many vocabulary studies) should be conducted. Regarding the model itself, we plan to incorporate skill-interdependencies as well as to take the affective user state into account, too. Both kind of extensions have been shown to improve learning [16, 29]. Additionally, the model can (and is meant to) provide a basis for exploiting the full potential of an embodied tutoring agent. Regarding this, we plan to advance the model such that the robot’s verbal and non-verbal communicative behaviour is adapted to the learner’s state of knowledge and progress. Specifically, we aim to enable dynamic adaption of (i) embodied behaviour such as iconic gesture use to be known to support vocabulary acquisition as a function of individual differences across children (cf. [26]); (ii) the robot’s synthetic voice to enhance comprehensibility and prosodic focusing of content when needed; and (iii) the robot’s socio-emotional behaviour depending on the learners’ current level of motivation or engagement. Further, as the long-term goal of our work is to enable robot-supported language learning for pre-school children, another important goal is to make children-specific adaptations to the language game and test it in child-robot interaction studies.

7. ACKNOWLEDGEMENTS

This work was supported by the L2TOR (www.l2tor.eu) project supported by the EU Horizon 2020 Program, grant number: 688014, and by the Cluster of Excellence Cognitive Interaction Technology ‘CITEC’ (EXC 277) at Bielefeld University, funded by the German Research Foundation (DFG).

8. REFERENCES

- [1] K. Bergmann and M. Macedonia. *A Virtual Agent as Vocabulary Trainer: Iconic Gestures Help to Improve Learners’ Memory Performance*, pages 139–148. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.
- [2] D. E. Berlyne. Conflict, arousal, and curiosity. 1960.
- [3] M. Cakmak and M. Lopes. Algorithmic and human teaching of sequential decision tasks. In *AAAI Conference on Artificial Intelligence (AAAI-12)*, 2012.
- [4] B. Clement, D. Roy, P.-Y. Oudeyer, and M. Lopes. Multi-armed bandits for intelligent tutoring systems. *arXiv preprint arXiv:1310.3174*, 2013.
- [5] A. T. Corbett and J. R. Anderson. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User modeling and user-adapted interaction*, 4(4):253–278, 1994.
- [6] M. G. Core and J. Allen. Coding dialogs with the damsl annotation scheme. In *AAAI fall symposium on communicative action in humans and machines*, volume 56. Boston, MA, 1997.
- [7] S. Craig, A. Graesser, J. Sullins, and B. Gholson. Affect and learning: An exploratory look into the role of affect in learning with autotutor. *Journal of Educational Media*, 29(3):241–250, 2004.
- [8] C. Dede. A review and synthesis of recent research in intelligent computer-assisted instruction. *International Journal of Man-Machine Studies*, 24(4):329–353, 1986.
- [9] S. Engeser and F. Rheinberg. Flow, performance and moderators of challenge-skill balance. *Motivation and Emotion*, 32(3):158–172, 2008.
- [10] M. Fridin. Storytelling by a kindergarten social assistive robot: A tool for constructive learning in preschool education. *Comput. Educ.*, 70:53–64, Jan 2014.
- [11] G. Gordon and C. Breazeal. Bayesian active learning-based robot tutor for children’s word-reading skills. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI’15*, pages 1343–1349. AAAI Press, 2015.
- [12] J. Gottlieb, P.-Y. Oudeyer, M. Lopes, and A. Baranes. Information-seeking, curiosity, and attention: computational and neural mechanisms. *Trends in cognitive sciences*, 17(11):585–593, 2013.
- [13] M. J. Habgood and S. E. Ainsworth. Motivating children to learn effectively: Exploring the value of intrinsic integration in educational games. *The Journal of the Learning Sciences*, 20(2):169–206, 2011.
- [14] J. Han, M. Jo, S. Park, and S. Kim. The educational use of home robots for children. In *ROMAN 2005. IEEE International Workshop on Robot and Human Interactive Communication, 2005.*, pages 378–383, Aug 2005.
- [15] E. ja Hyun, S. yeon Kim, S. Jang, and S. Park. Comparative study of effects of language instruction program using intelligence robot and multimedia on linguistic ability of young children. In *RO-MAN 2008 - The 17th IEEE International Symposium on Robot and Human Interactive Communication*, pages 187–192, Aug 2008.
- [16] T. Käser, S. Klingler, A. G. Schwing, and M. Gross. *Beyond Knowledge Tracing: Modeling Skill Topologies with Bayesian Networks*, pages 188–198. Springer International Publishing, Cham, 2014.
- [17] J. Kennedy, P. Baxter, and T. Belpaeme. The robot who tried too hard: Social behaviour of a robot tutor can negatively affect child learning. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction, HRI ’15*, pages 67–74, New York, NY, USA, 2015. ACM.
- [18] J. Kory and C. Breazeal. Storytelling with robots: Learning companions for preschool children’s language development. In *The 23rd IEEE International Symposium on Robot and Human Interactive Communication*, pages 643–648, Aug 2014.
- [19] J. Kory Westlund, G. Gordon, S. Spaulding, J. Lee, L. Plummer, M. Martinez, M. Das, and C. Breazeal. Learning a second language with a socially assistive robot. In *The 1st International Conference on Social Robots in Therapy and Education*, 2015.

- [20] H. Kose-Bagci, E. Ferrari, K. Dautenhahn, D. S. Syrdal, and C. L. Nehaniv. Effects of embodiment and gestures on social interaction in drumming games with a humanoid robot. *Advanced Robotics*, 23(14):1951–1996, 2009.
- [21] S. Leitner. So lernt man lernen: Der weg zum erfolg [learning to learn: The road to success]. *Freiburg: Herder*, 1972.
- [22] D. Leyzberg, S. Spaulding, M. Toneva, and B. Scassellati. The physical presence of a robot tutor increases cognitive learning gains. In *CogSci*. Citeseer, 2012.
- [23] M. Macedonia, K. Müller, and A. D. Friederici. Neural correlates of high performance in foreign language vocabulary learning. *Mind, Brain, and Education*, 4(3):125–134, 2010.
- [24] J. R. Movellan, M. Eckhardt, M. Virnes, and A. Rodriguez. Sociable robot improves toddler vocabulary skills. In *Human-Robot Interaction (HRI), 2009 4th ACM/IEEE International Conference on*, pages 307–308, March 2009.
- [25] A. N. Rafferty, E. Brunskill, T. L. Griffiths, and P. Shafto. Faster teaching via pomdp planning. *Cognitive Science*, 2015.
- [26] M. L. Rowe, R. D. Silverman, and B. E. Mullan. The role of pictures and gestures as nonverbal aids in preschoolers’ word learning in a novel language. *Contemporary Educational Psychology*, 38(2):109–117, 2013.
- [27] M. Saerbeck, T. Schut, C. Bartneck, and M. D. Janse. Expressive robots in education: Varying the degree of social supportive behavior of a robotic tutor. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’10, pages 1613–1622, New York, NY, USA, 2010. ACM.
- [28] G. Skantze and S. Al Moubayed. Iristk: A statechart-based toolkit for multi-party face-to-face interaction. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction*, ICMI ’12, pages 69–76, New York, NY, USA, 2012. ACM.
- [29] S. Spaulding, G. Gordon, and C. Breazeal. Affect-aware student models for robot tutors. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems*, AAMAS ’16, pages 864–872, Richland, SC, 2016. International Foundation for Autonomous Agents and Multiagent Systems.