A Bayesian-Optimal Principle for Child-Friendly Adaptation in Learning Games*

Janne V. Kujala† Ulla Richardson Heikki Lyytinen

February 26, 2008

Abstract

Adaptation in learning games typically aims for a certain success rate assumedly ideal for learning. However, simple counterexamples show that any success rate by itself cannot guarantee learning results. Instead, we propose that the tasks that yield the most information about the student's skills would also facilitate learning. This approach naturally avoids exceedingly easy and exceedingly difficult tasks as their results are predictable and thus uninformative. However, as failures can lower motivation, we propose the more child-friendly objective of maximizing the information gain divided by the failure rate. We apply these principles to a model of idiosyncratic item responses.

Keywords: Bayesian adaptive estimation, mutual information, item response theory, adaptive learning game, educational game, E-learning, Hidden Markov Model, Fast Fourier Transform.

1 Introduction

In the recent years, there has been an increasing interest in computerized learning games (Ketamo, 2003; Manske and Conati, 2005; Wilson et al., 2006; Lyytinen et al., 2007). In these games, the driving principle for optimizing learning appears to be that if the number of mistakes is too high or if it takes the student too long to complete tasks, the learning material is too difficult, and if the number of mistakes is very low or it takes very little time for the student to complete tasks, the learning material is too easy (Ketamo, 2003).

*This research was supported by the Academy of Finland (grant number 121855) and by the European Commission's FP6, Marie Curie Excellence Grants (MCEXT-CE-2004-014203). The authors are grateful to Rauli Ruohonen for comments.

†Corresponding author. Address: Agora Center, University of Jyväskylä, P.O.Box 35, FI-40014 Jyväskylä, Finland. Email address: jvk@iki.fi. Fax: +358 14 2604400.

Thus, the adaptation logic typically follows some variation of the ad hoc algorithm that increases the difficulty of the learning tasks after correct answers and decreases it after incorrect answers. The algorithm is tuned so as to yield around 75%, 80%, or some such percentage of correct answers. While this simple logic may work reasonably well when the learning material has only one adaptable dimension of difficulty, it runs into trouble when the learning material is more complicated.

Wilson et al. (2006) model multiple dimensions of difficulty of the learning material by applying ad hoc formulas to update estimates of the student's success probability over the "learning space" with these dimensions. The adaptation logic aims at keeping the success probability around 75%. Although this is welcome multidimensional generalization, the algorithm is relatively complicated with many tuning constants and is not based on a statistical model; the only explicit principle underlying the adaptation logic is still a certain prescribed success rate (even though the algorithm may implicitly implement other principles).

An average success rate around certain numbers may be necessary for efficient learning, but it is certainly not sufficient; for example, alternating exceedingly easy and exceeding difficult tasks in 3:1 proportion, or always giving a multiple-choice task with three correct choices within four identical (unmarked) options, will yield a 75% success rate, but these tasks are obviously suboptimal for learning. To avoid such pathologies, the adaptation principles should be based on a full-blown statistical model of the student's skills as well as the task rather than on simple estimates of success probability or response times.

Principled Bayesian student models are already becoming prevalent in the field (e.g., Manske and Conati, 2005), but they have mostly been used for offering individualized support for the student rather than actively choosing the most effective learning material. An interesting approach by Stacey et al. (2003) suggests the use of cognitive conflict to facilitate learning, that is, choosing learning material that the student is likely to get wrong due to misconceptions. However, there is an obvious conflict with keeping the student motivated and therefore it is not immediately clear how this approach would generalize the simple principle of a certain average success rate.

In this paper, we propose general principles of adaptation that are applicable in any Bayesian student model. In the following sections, we first describe Bayesian adaptive estimation framework and propose our principles of adaptation as an extension of it. Then, we present an example model and apply the principles to it using some novel computational techniques. Finally, we evaluate the resulting algorithm using simulations and results from a pilot study with a real computerized learning game and real students. We conclude with a discussion of the philosophy behind our proposed principles.

2 Bayesian adaptive estimation

The ad hoc adaptation rules generally used in adaptive teaching systems are analogous to the staircase method of measurement (see Treutwein, 1995), in which a sensory detection threshold is estimated by increasing the stimulus intensity after each correctly detected stimulus and decreasing it after each undetected stimulus. The staircase method is not very efficient for measurement, often wasting several trials before reaching the overall scale of the threshold and then moving randomly around the optimal measurement points.

Bayesian adaptive estimation can be much more efficient, as it can efficiently use the information provided by all earlier trials even when the information is too scarce to obtain meaningful estimates of the parameters. The underlying Bayesian model (Watson and Pelli, 1983) can be represented graphically as

$$\begin{array}{ccccc} & \Theta & & & \\ \swarrow & \downarrow & \searrow & & \\ r_{x_1} & r_{x_2} & \ldots & r_{x_I} & \end{array}, \tag{1}$$

where r_{x_1}, r_{x_2}, $\ldots$, r_{x_I} are the trial results (with contents x_1, x_2, $\ldots$, x_I) having a known statistical dependence $p(r_x \mid \theta)$ on the unknown parameter Θ that is being estimated[1]. Assuming a prior distribution $p(\theta)$ for the unknowns Θ, Bayes' formula yields the posterior distribution

$$p(\theta \mid y) = \frac{p(y \mid \theta)p(\theta)}{p(y)} = \frac{\left[\prod_{i=1}^{I} p(r_{x_i} \mid \theta)\right] p(\theta)}{\int \left[\prod_{i=1}^{I} p(r_{x_i} \mid \theta)\right] p(\theta) d\theta}$$

of the unknowns Θ given the data $y = (r_{x_1}, \ldots, r_{x_I})$.

2.1 Mutual information

For each prospective trial content x to be presented next, the information theoretic *mutual information*

$$I(R_x; \Theta \mid y) := \iint p(r_x, \theta \mid y) \log \frac{p(r_x, \theta \mid y)}{p(r_x \mid y)p(\theta \mid y)} dr_x \, d\theta$$

(given the previously observed results y) of the unknown variables Θ and the (random) result R_x of the next trial provides an elegant way of quantifying the trial's expected informativity in light of all prior information (Lindley, 1956; MacKay, 1992). The "greedy" measurement strategy of always presenting the content x yielding the highest value of $I(R_x; \Theta \mid y)$ works well in practice, even though it is not necessarily optimal

[1]Following standard notation, upper case letters denote random variables and lower case letters their values. However, we differ from the common notation $p(r \mid x,\ \theta)$ and instead use the notation $p(r_x \mid \theta)$ of (Kujala and Lukka, 2006), *indexing* a set of distinct result variables with the content x. This notation works better with the information-theoretic notation of the next section.

when the experiment consists of more than one trial. Variations of this general formulation have been applied successfully in several different fields, including psychometric measurement (Kontsevich and Tyler, 1999), machine learning (MacKay, 1992), student assesment (Liu, 2005), etc.

For dichotomous results $R_x \in \{0, 1\}$, the mutual information is conveniently computed (Kujala and Lukka, 2006) as

$$I(R_x; \Theta \mid y) = H(R_x \mid y) - H(R_x \mid \Theta, y), \tag{2}$$

where

$$\begin{aligned} H(R_x \mid y) &= h\left(\int p(R_x = 1 \mid \theta) p(\theta \mid y) d\theta\right), \\ H(R_x \mid \Theta, y) &= \int h(p(R_x = 1 \mid \theta)) p(\theta \mid y) d\theta, \end{aligned}$$

and $h(p) = -p \log p - (1 - p) \log(1 - p)$ denotes the entropy of a binary distribution with probabilities p and $1 - p$.

2.2 Adaptation in teaching systems

In measurement, the adaptation's only goal is to optimize the measurement efficiency. In a teaching system, there can be several goals, but the most important is the optimization of the learning result.

Adaptive Bayesian methodology can be used to quickly find a desired difficulty level for the learning material in a manner similar to a binary search. The general framework automatically allows more than one "dimension of difficulty" of the content to be adapted simultaneously, making it relatively easy to enforce a certain success rate after enough information has been accumulated. However, whether a certain success rate should be aimed for, or if there is some more important principle for optimal learning is open to question.

In fact, we emphasize the auxiliary goal of good measurement efficiency instead. Information of student's skills is valuable for analysis and feedback, and at least some measurement information is necessary for optimization of the learning effects anyway. The key hypothesis of our proposed approach is that the trial contents that are good for measurement, which is highly desirable in its own right, would also be among those that are the most efficient for learning.

Indeed, optimization of the expected information gain of each trial appears to avoid naturally trial contents that are extremely difficult or extremely easy as the model can predict the results of such trials to be correct or incorrect and so the actual answers would yield little new information. This alogrithm can be seen as a formalization of some aspects of the cognitive conflict approach pursued by Stacey et al. (2003): material that best differentiates between alternative hypotheses about student's conceptions is preferred.

However, there is one obvious flaw in pure optimization of the information gain as a general principle: the trial contents may be too difficult to keep children motivated, even though these contents might still be efficient for learning if the child would keep playing. In the following subsection, we propose a principled solution.

2.3 Child-friendly measurement principle

It is generally agreed that failures of a child may lower her motivation and therefore the more failures there are the less motived the child may become. In order to maintain motivation in a teaching system, we propose the following adjustment of the pure measurement goal: instead of simply optimizing the expected information gain of each trial, we optimize the expected gain divided by the expected cost measured as the estimated probability of an incorrect answer. In other words, we try to obtain as much measurement information as possible per each failure of the child. Thus, instead of maximizing $I(R_x; \Theta \mid y)$, the content x is chosen so as to maximize

$$\frac{I(R_x; \Theta \mid y)}{E(C_x \mid y)},$$

where the expected value of the cost $C_x = [R_x = 0]$ in light of the previously observed data y is given by

$$E(C_x \mid y) = \int p(R_x = 0 \mid \theta) p(\theta \mid y) d\theta.$$

The formal justification for this heuristic and conditions for its optimality as a strategy for optimizing the total gain per total cost ratio will be dealt with in another paper (Kujala, 2008).

On a more practical side, the reader may wonder what happens if the expected probability of failure is zero for some content. In that case, there would be no cost of presenting that content and if it was expected to give even the slightest amount of information, it would always be presented. However, a realistic model is bound to assume some probability of careless mistakes and therefore the expected cost should never be zero.

In a task where accuracy of answers is not an issue, one could instead define the cost C_x as the response time, assuming that the longer a student has to struggle with the response, the less motivated she might become. In that case, there would be no conflict with the measurement goal — the most information per time unit would also be ideal for measurement.

3 A model of idiosyncratic item responses

To test the proposed child-friendly adaptation principle, we formulated a Bayesian model of a student and a task used in the Literate (Ekapeli) learning game (Lyytinen et al., 2007). The task is described in Fig. 1.

Category	Example trial: Target	Example trial: Choice-set
1	a	$\boxed{\text{a}}$ $\boxed{\text{s}}$
10	l	$\boxed{\text{m}}$ $\boxed{\text{n}}$ $\boxed{\text{l}}$ $\boxed{\text{r}}$
26	out	$\boxed{\text{oot}}$ $\boxed{\text{out}}$ $\boxed{\text{ouk}}$ $\boxed{\text{oup}}$
94	kattarmo	$\boxed{\text{kattarmo}}$ $\boxed{\text{kattairmo}}$ $\boxed{\text{katarmo}}$

Figure 1: In the pilot game implementing our proposed child-friendly adaptation principle, all possible learning tasks are divided into 94 categories of increasing difficulty, ranging from phonetically distinct letters to long, minimally distinct (in the Finnish language context) pseudowords. The target is presented auditorily and the subject's task is to choose the corresponding graphemic representation from the visually presented choices. The number of choices can vary within one category. After each trial, brief feedback is given to facilitate learning. As the Finnish writing system is practically 100% consistent (each letter has its own phoneme and each phoneme its own letter), learning the core knowledge of reading by this kind of drilling should be possible.

3.1 The model

The learning material is divided into categories $k = 1, \ldots, K$, where all content within one category is assumed to be equally difficult. The unknown difficulty of each category is denoted by Θ_k and the unknown overall skill of the student is denoted by B (upper case β). The result $R_{k,n}$ of an n-choice trial with content from category k is assumed to be distributed as

$$p(R_{k,n} = 1 \mid \theta_k, \beta) = \psi_n(\beta - \theta_k),$$

where $R_{k,n} = 1$ denotes success and $R_{k,n} = 0$ failure.

The link function ψ_n is given by

$$\psi_n(x) = \frac{1}{n} + \frac{n}{n-1} \cdot \frac{1-\delta}{1 + \exp(c \log_2(n/2) - x)},$$

where $\delta = 0.05$ is the probability of an incorrect answer due to careless mistakes and the lower asymptote is given by $1/n$, the probability of guessing correctly (see Fig. 2). Additionally, the parameter $c = 2$ shifts the threshold slightly to the right for large numbers of choices, modeling the fact that a large number of choices may overwhelm a student whose ability only just meets the difficulty of the task. The parameters δ and c could also be estimated, but for now, we have used constant values.

In our application, there are $K = 94$ categories, approximately ordered by difficulty, but we do expect some idiosyncratic deviations from this ordering. Therefore, we assume the prior distributions

$$\Theta_k \sim N(\mu_k, \sigma_k^2)$$

for the difficulties, where $\mu_k = 15 + 2k$ and $\sigma_k = 5$ for $k = 1, \ldots, 94$ (see Fig. 3). The unknown skill B is assumed to be distributed uniformly on $[0, 220]$. The attainable

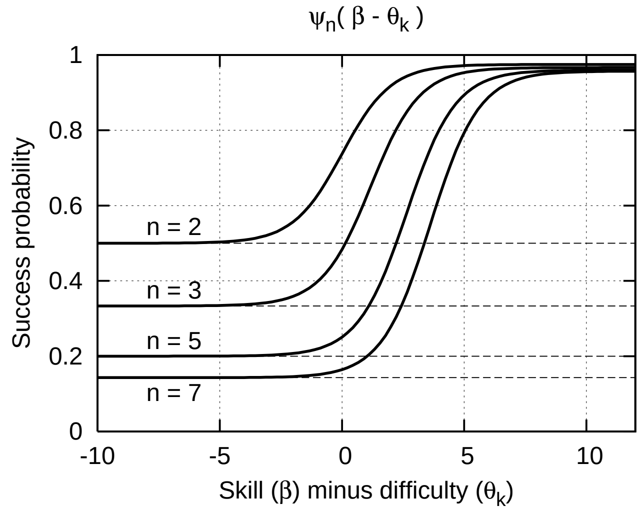

Figure 2: The success probability of an n-choice trial of category k.

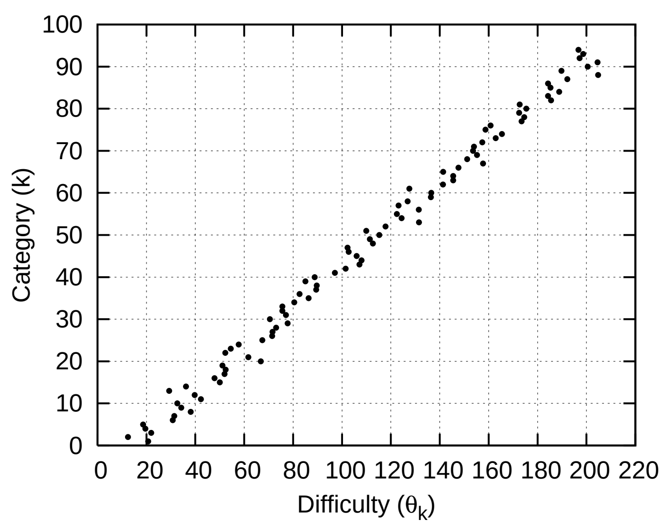

Figure 3: A realization drawn from the assumed prior distribution of catergory difficulties, illustrating the type of idiosyncratic deviations that the model anticipates from the ideal ordering of the categories.

values of the number of choices $n \in \{2, \ldots, 7\}$ vary between the categories; in some categories, there are only two-choice minimum pair trials, while in other categories, a larger selection of distractor items may be available. All distractors are assumed to be equally discriminable from the target so that we can ignore the specific answer and just model its correctness.

3.2 Bayesian estimation using convolutions

Given any independent prior distributions $p_B(\beta)$ and $p_k(\theta_k)$ for the unknown parameters, the joint prior is

$$p(\beta, \theta) = p_B(\beta) \prod_k p_k(\theta_k).$$

Suppose we have observed the sequence $y := (k_1, n_1, r_1), \ldots, (k_I, n_I, r_I)$ of trial results. Then, the likelihood of the parameter values β and $\theta = (\theta_1, \ldots, \theta_K)$ given the data y is

$$p(y \mid \beta, \theta) = \prod_k f_k(\beta - \theta_k),$$

where

$$f_k(x_k) := \prod_{i=1}^{I} \begin{cases} \psi_{n_i}(x_k), & r_i = 1,\ k_i = k, \\ 1 - \psi_{n_i}(x_k), & r_i = 0,\ k_i = k, \\ 1, & k_i \neq k \end{cases}$$

is the likelihood of the random variable $X_k := B - \Theta_k$. It follows that the posterior is given by

$$p(\beta, \theta \mid y) = \frac{p_B(\beta) \prod_k f_k(\beta - \theta_k) p_k(\theta_k)}{\int p_B(\beta) \prod_k \left[\int f_k(\beta - \theta_k) p_k(\theta_k) d\theta_k\right] d\beta}. \tag{3}$$

Using the convoulution operator defined as

$$(f * g)(y) := \int f(y - x) g(x) dx,$$

the inner integrals in the denominator of (3) can be written as $(f_k * p_k)(\beta)$. It turns out that we can integrate out any variables in the posterior using convolutions. Thus, the marginal posterior of B is

$$p(\beta \mid y) = \int p(\beta, \theta \mid y) d\theta = \frac{p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta}{\int p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta},$$

and similarly, the marginal posterior of Θ_l is

$$\begin{aligned} p(\theta_l \mid y) &= \frac{\int \left[\prod_{k \neq l} f_k * p_k\right](\beta) f_l(\beta - \theta_l) p_l(\theta_l) d\beta}{\int p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta} \\ &= \frac{p_l(\theta_l) \left(\overleftarrow{f_l} * \left[\prod_{k \neq l} f_k * p_k\right]\right)(\theta_l)}{\int p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta}, \end{aligned}$$

where we denote $\overleftarrow{f}(x) = f(-x)$.

3.3 Adaptive trial placement

As the trial result $R_{k,n}$ depends on the unknown variables only through X_k, we have

$$I(R_{k,n}; B, \Theta \mid y) = I(R_{k,n}; X_k \mid y).$$

This is analogous to the well-known fact that $I(Y;\Theta) = I(T(Y);\Theta)$ if T is a sufficient statistic of the data Y depending on an unknown parameter Θ (Cover and Thomas, 1991, p. 37), although in this context X_k as a function of (Θ, B) is the sufficient statistic, that is, the roles of the parameters and the data are reversed here (which is fine as the mutual information is symmetric). Thus, we only need to find the posterior distribution of X_k given y to be able to compute the expected information gain of observing $R_{k,n}$ next.

Changing variables to (B, X), the prior is

$$p(\beta, x) = p_B(\beta) \prod_k p_k(\beta - x_k),$$

and the likelihood is $p(y \mid \beta, x) = \prod_k f_k(x_k)$, yielding the posterior

$$p(\beta, x \mid y) = \frac{p_B(\beta) \prod_k f_k(x_k) p_k(\beta - x_k)}{\int p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta}.$$

Thus, the marginal posterior of x_l is given by

$$p(x_l \mid y) = \frac{f_l(x_l) \left(\overleftarrow{p_l} * \left[p_B \prod_{k \neq l} f_k * p_k \right] \right)(x_l)}{\int p_B(\beta) \left[\prod_k f_k * p_k\right](\beta) d\beta}.$$

Note the "dualism" of the prior and likelihood between the two parameterizations, which is apparent in the expressions of $p(x_l \mid y)$ and $p(\theta_l \mid y)$.

As per (2), the expected information gain of observing $R_{k,n}$ next can be computed as

$$I(R_{k,n}; X_k \mid y) = H(R_{k,n} \mid y) - H(R_{k,n} \mid X_k, y),$$

where

$$\begin{aligned}
H(R_{k,n} \mid y) &= h\left(\int p(R_{k,n} = 1 \mid x_k) p(x_k \mid y) dx_k \right) \\
&= h\left(\int \psi_n(x_k) p(x_k \mid y) dx_k \right), \\
H(R_{k,n} \mid x_k, y) &= \int h(p(R_{k,n} = 1 \mid x_k)) p(x_k \mid y) dx_k \\
&= \int h(\psi_n(x_k)) p(x_k \mid y) dx_k,
\end{aligned}$$

and the expected cost, the probability of a failure, is given by

$$E(C_x \mid y) = p(R_{k,n} = 0 \mid y) = 1 - \int \psi_n(x_k) p(x_k \mid y) dx_k.$$

The resulting adaptation tree for the first 6 trials is shown in Fig. 4 for both the child-friendly and the pure information maximization strategies.

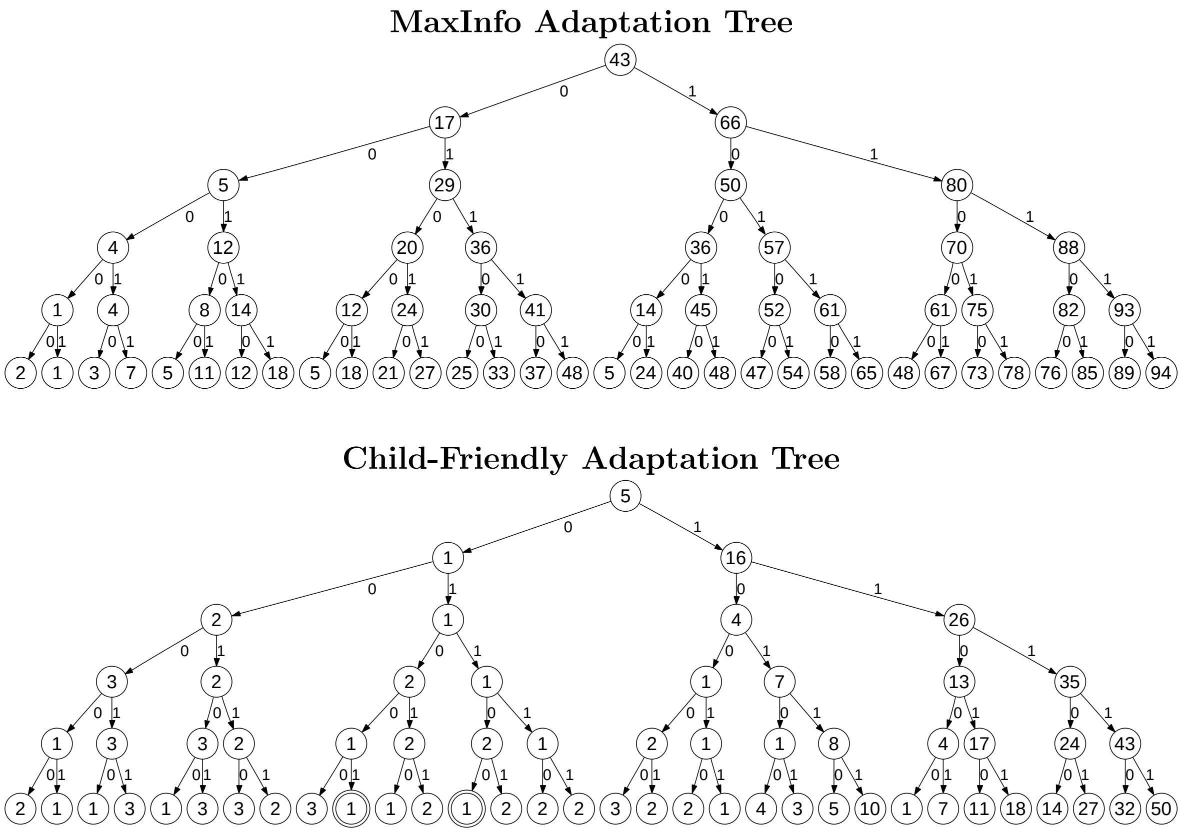

Figure 4: Six first layers of the adaptation trees of the pure information maximization principle and the child-friendly variation applied to the model. The edges between the nodes indicate correct (1) and incorrect (0) answers and each node shows the category number of the trial to be presented next. The pure information maximization works similar to a binary search in the first few steps. The child-friendly variation, however, tests a category much closer to the easier end at each step and thereby yields a higher proportion of correct answers. In these first layers, the number of choices for each trial is always the maximum possible for that category and is not indicated. As the estimate of the skill level gets more accurate, the child-friendly variation begins to present trials with lower than the maximum number of choices. In the child-friendly tree, the two nodes marked with ◎ have the exact same data (but in a different order) on the paths leading to them and so the subtrees rooted at these nodes are identical.

3.4 Practical implementation using FFT

If we discretize β and θ_k using a uniformly pitched sampling, the convolutions can be efficiently computed using the Fast Fourier Transform (Cooley and Tukey, 1965). By precomputing the functions $a_l = \prod_{k<l} f_k * p_k$ and $b_l = \prod_{k>l} f_k * p_k$ for all $l = 1, \ldots, K$ at each trial, the expressions $\prod_{k \neq l} f_k * p_k = a_l b_l$ can be evaluated quickly for all different $l = 1, \ldots, K$. The whole Bayesian updating, computation of marginal distributions, and evaluation of the expected information gain for all possible k and n can be done in $\mathcal{O}(KNM \log M)$ time and $\mathcal{O}(KM)$ space, where K is the number of categories, N is the number of different numbers of choices in a trial and M is the number of discretized values of the β and θ_k variables.

In practice, the computations took only about a second per trial (with $K = 94$, $M = 6$, and $M = 1024$) which is a huge speedup over the naive implementation with a $(K+1)$-dimensional parameter-hypercube, which with its $\mathcal{O}(KNM^{K+1})$ running time per trial would become impractical already at around $K = 3$ with present computers.

This idea is similar to the speedups presented in (Kujala and Lukka, 2006).

3.5 Dynamic changes

In Bayesian adaptive estimation, the underlying model has traditionally been static, that is, it is assumed that the measured variables do not change over time. This is of course an inappropriate assumption for a learning game, but can be assumed to hold within one session (around 35 trials) of the game. Thus, we assume a static model within one session but anticipate changes in the true values of the parameters between sessions.

The specifc dynamic formulation we propose is a Hidden Markov Model (HMM), that is, it is assumed that the student has a certain state θ_t (corresponding to both β and θ in the present model) that can change over time, but so that given the current state θ_t, the following random state Θ_{t+1} is independent of the past states $\theta_{t-1}, \theta_{t-2}, \ldots$ (i.e., there is no memory of the past states other than what is encoded in the present state). This measurement model can be represented graphically as

$$\begin{array}{ccccccccc} \Theta_0 & \rightarrow & \Theta_1 & \rightarrow & \Theta_2 & \rightarrow & \Theta_3 & \rightarrow & \ldots \\ & & \downarrow & & \downarrow & & \downarrow & & \\ & & y_1 & & y_2 & & y_3 & & \end{array}, \tag{4}$$

where y_t indicates the observed data for each session t. In this general formulation, we allow Θ_t to have any structure, the only assumption being the Markov (no-memory) property. Therefore, we use the term HMM to describe the general formulation although specific models of this formulation may be better described as Dynamic Bayesian Networks (DBNs) as some authors reserve the term HMM only for models with a single (discrete) variable.

In the implementation of our specific model, at the end of each session, we throw away any posterior interdependencies between B and the Θ_k, that is, we replace their joint posterior with the product of its marginals. This slight deviation from the ideal Bayesian model is needed to maintain the simple form (3) of the posterior for the next session, where we assume that the true difficulties Θ_k of the categories and the skill B of the student may have changed from those of the previous session according to certain transition densities $T_k(\theta'_k \mid \theta_k)$ and $T_B(\beta' \mid \beta)$. This yields the following priors for the next session:

$$p'_k(\theta'_k) := \int T_k(\theta'_k \mid \theta_k) p_k(\theta_k \mid y) d\theta_k,$$
$$p'_B(\beta') := \int T_B(\beta' \mid \beta) p_B(\beta \mid y) d\beta.$$

The specific transition densities we use in the pilot version of the game are defined so that

$$B' \mid \beta \sim \begin{cases} N_{[0,220]}(\beta, \sigma^2), & \text{with probability } 0.8, \\ \text{Uniform}[0, 220], & \text{with probability } 0.2, \end{cases}$$

(where $N_{[0,220]}$ denotes a normal distribution such that any value falling below 0 or above 220 is mirrored back in, i.e., transitions going outside the range bounce back) and

$$\Theta'_k \mid \theta_k \sim N_{[0,220]}(s_k(\theta_k - \mu_k) + \mu_k, \sigma^2),$$

where $\sigma = 3$ determines the expected magnitude of the transitions and the constant

$$s_k := \left(1 - \frac{\sigma^2}{\sigma_k^2}\right)^{1/2}$$

is chosen so that without new observations from this category, the distribution of Θ_k eventually converges back to its prior. Thus, in addition to small changes to the skill and difficulties, we expect that the student's skill may occasionally jump to any value. That way the model can keep up if the student happens to learn (or forget) outside the game. These transition densities can be efficiently applied by using the FFT convolution.

This complete description of our implementation of the model serves to show that the inherent computational complexity of the proposed principles can be overcome in a practical way. There is a lot more that could be said about the rationale and details of these design choices, but it is beyond the scope of the present paper.

3.6 Evaluation

As both variations of the algorithm look only one trial ahead (greedy optimization) they are only approximations to the truly optimal strategy under their assumptions. Therefore, we compared the true measurement efficiencies of both algorithms under various session lengths. The results shown in Fig. 5 confirm that the child-friendly variation indeed yields better estimation efficiency per the number of mistakes of the child while the pure information maximization strategy is more efficient per the number of trials.

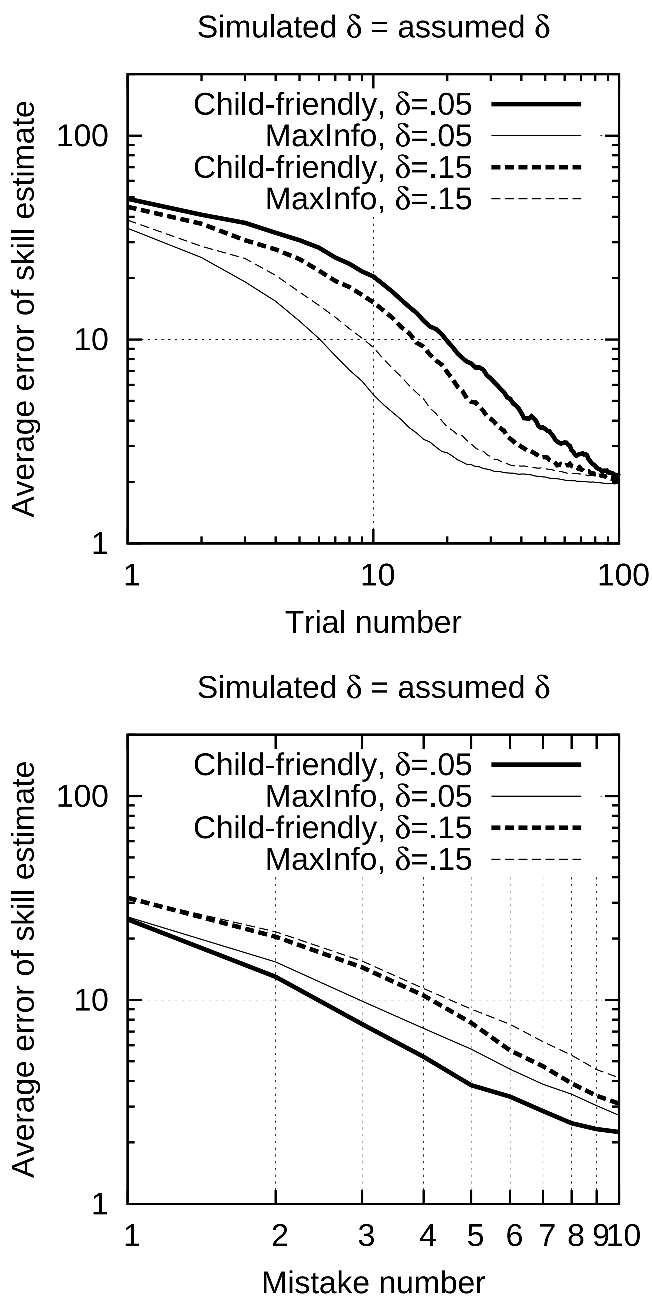

Figure 5: Accuracy of the final skill estimate computed as the posterior mean based on the data from a session (hypothetically) terminated after the nth trial (top) or the nth mistake (bottom). Average over (4×) 1000 simulation runs. As expected, the pure information maximization strategy yields the best estimation efficiency per number of trials whereas the child-friendly variation yields the best efficiency per the number of mistakes.

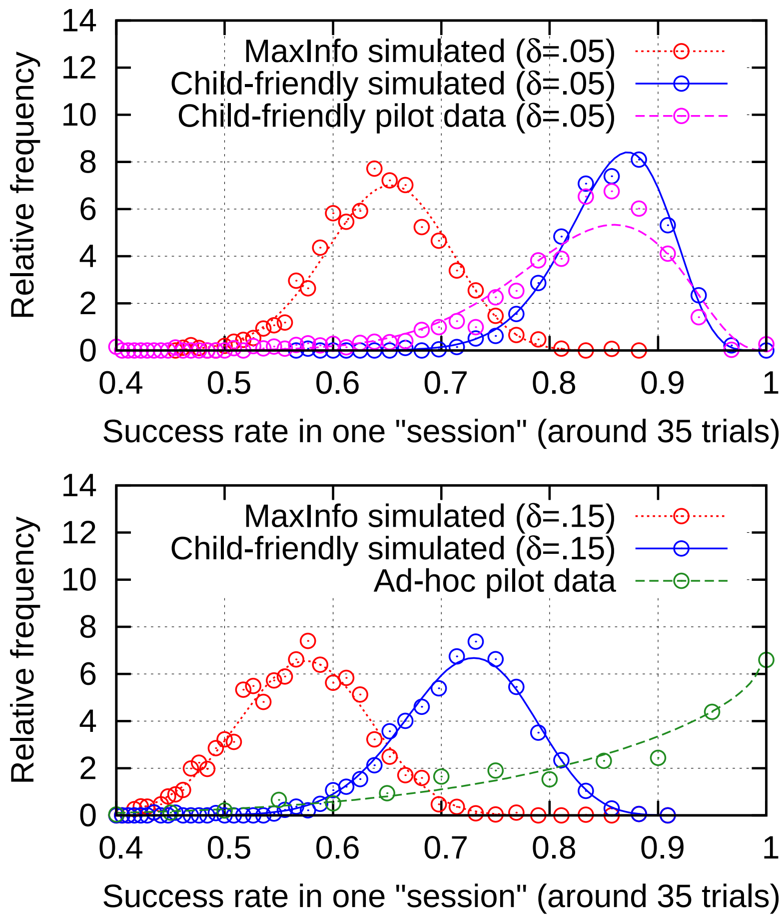

Figure 6: Comparison of the success rates yielded by the two variations of the algorithm. The child-friendly variation appears to yield reasonable success rates just by following its implicit logic. The slight deviation of the real data results from the simulated ones is mostly due to a few students whose skills were apparently far below the simulated range. The difference between the upper and lower panels indicates that the assumed value of δ has a significant effect on the resulting success rate. However, for a given assumed value of δ, varying the simulated true value of δ between .05 and .15 had almost no effect on the distributions (around $\pm.01$ shift of the modes). The lower panel also shows the results from an ad-hoc adaptation algorithm implemented in another pilot version of the game (see text).

In the pilot game, each session terminates after 30 correct answers[2], which usually happens at around 35 trials, but the player is free to continue for as many sessions as she wishes. We simulated (4×) 100 players for up to 10 sessions following this session termination rule and drawing the true values of the skill and difficulties from the assumed priors and transition densities. Additionally, the simulation was stopped early if the skill of the simulated player drifted outside the range $[\mu_1, \mu_{94}]$. Figure 6 depicts the resulting distribution of average success rates in one session. The real data also shown there are from 28 real players totaling approximately the same number of sessions as the simulated data. The smooth lines are beta densities fitted to the histogram data. The results for just the first simulated session of each player were practically identical to those of the full data so the exact type of true dynamics is apparently not important for the resulting success rate under these conditions. However, as mentioned in the figure caption, the boundary effects at the upper and lower range of skills are significant for the resulting success rate.

Figure 6 also shows the success rate that resulted from an ad hoc adaptation logic implemented in another pilot version of the game. That version used roughly the same categories of content, but presented only trials of one category within one session. Within each session, the difficulty was controlled by adapting the number of choices according to the ad hoc rules discussed in the introduction. The game advanced to the next category only after a certain low percentage of mistakes was reached in the current category. The early categories were too easy for many children, which yielded the large probability mass near the 100% success rate. Also, for some children some categories were exceptionally difficult within their position in the ordering, which resulted in the long left tail of the distribution as the children struggled through these categories. The Bayesian principles avoid all these problems by anticipating deviations from the difficulty order of the categories and by freely moving between the categories in each session. However, at this point we do not have an objective comparison of the induced learning effects between the different adaptation principles.

4 Discussion

Learning is a cumulative process and so the specific effects of a single learning trial are difficult to assess. Even if such effects could be accurately modeled, their interactions would likely be complicated. Therefore, we do not model any explicit effect of the presentation of a certain content, that is, the true state of the student in the future is not assumed to explictly depend on the contents of the past trials. We do anticipate the possibility of changes in the state of the student, but we do not expect that exposures to certain content would automatically trigger certain changes. Instead, we have adopted a less explicit but more easily quantifiable view on the induced learning effects.

Analogous to the fact that measurements of a physical system can alter its state,

[2] This termination rule may seem a bit odd, but it has little significance in the Bayesian framework and therefore its rationale is not important here.

measurements of student's skills can alter her state. We assume that conducting the measurements where there is the most uncertainty about the student's skills would yield the best chances of the student learning. For example, if the prerequisites of a certain task have not been learned, there is little uncertainty about the task's result and so its measurement would be useless; similarly, if a certain task is already mastered, there is little uncertainty about its prerequisites and so their measurement would be useless. This reasoning leads to our proposed principle of optimizing the measurement efficiency as a means of optimizing the learning results.

Apart from suitable learning material, another important factor for learning is motivation — too many failures of a student should be avoided. Assuming that the chances of learning are proportional to the amount of measurement information given by the results of the learning tasks, and inversely proportional to the number failures of the student, the mathematically optimal teaching strategy is to optimize the information gain divided by the number of failures.

This formulation exemplifies the philosophy behind our approach: instead of enforcing ad hoc constraints on content selection, we aim at modeling the phenomena behind the constraints so that the mathematical solution to the optimization problem yields the desired behavior without any post hoc adjustments.

Our proposed principle of child-friendly adaptation is general, applicable in any Bayesian model having the notion of a success or failure of the student in each learning task. The example model we presented illustrates that the proposed principle can be expected to work generally in a sensible way without any tuning parameters, simply by following its implicit internal logic. It should be noted, however, that if the student model does not adequately capture the structure of the learning material, then the proposed principle cannot be expected to work very well. For example, if there is no prior information about the relative difficulties of different tasks, then the model has no way of knowing which ones might be suitable for the student without trying all of them in an arbitrary order until one is found. Thus, even though the principles themselves are completely general, they do not obviate the need to carefully develop an appropriate student model for the specific problem.

References

Cooley, J. W. and Tukey, J. W. "An Algorithm for the Machine Calculation of Complex Fourier Series." *Mathematics of Computation*, 19(90):297–301 (1965).

Cover, T. M. and Thomas, J. A. *Elements of Information Theory*. John Wiley (1991).

Ketamo, H. "An Adaptive Geometry Game for Handheld Devices." *Educational Technology & Society*, 6(1):83–95 (2003).

Kontsevich, L. L. and Tyler, C. W. "Bayesian adaptive estimation of psychometric slope and threshold." *Vision Research*, 39(16):2729–2737 (1999).

Kujala, J. V. “Bayesian adaptive estimation under a random cost of observation associated with each observable variable.” University of Jyväskylä, Department of Mathematics and Statistics, Preprint 364 (2008).

Kujala, J. V. and Lukka, T. J. “Bayesian Adaptive Estimation: The Next Dimension.” *Journal of Mathematical Psychology*, 50(4):369–389 (2006).

Lindley, D. V. “On a Measure of the Information Provided by an Experiment.” *The Annals of Mathematical Statistics*, 27(4):986–1005 (1956).

Liu, C.-L. “Using Mutual Information for Adaptive Item Comparison and Student Assessment.” *Educational Technology & Society*, 8(4):100–119 (2005).

Lyytinen, H., Ronimus, M., Alanko, A., Poikkeus, A.-M., and Taanila, M. “Early identification of dyslexia and the use of computer game-based practice to support reading acquisition.” *Nordic Psychology*, 59(2):109–126 (2007).

MacKay, D. J. “Information-based objective functions for active data selection.” *Neural Computation*, 4(4):590–604 (1992).

Manske, M. and Conati, C. “Modelling Learning in Educational Games.” In *Proceedings of the 12th International Conference on AI in Education (AIED 05), Amsterdam, July 19–23* (2005).

Stacey, K., Sonenberg, E., Nicholson, A., Boneh, T., and Steinle, V. “A Teaching Model Exploiting Cognitive Conflict Driven by a Bayesian Network.” In *Proceedings of the 9th International Conference on User Modeling*, 352–362. Johnstown, PA, USA (2003).

Treutwein, B. “Adaptive psychophysical procedures.” *Vision Research*, 35(17):2503–2522 (1995).

Watson, A. B. and Pelli, D. G. “QUEST: A Bayesian adaptive psychometric method.” *Perception & Psychophysics*, 33(2):113–120 (1983).

Wilson, A. J., Dehaene, S., Pinel, P., Revkin, S. K., Cohen, L., and Cohen, D. “Principles underlying the design of “The Number Race”, an adaptive computer game for remediation of dyscalculia.” *Behavioral and Brain Functions*, 2(19) (2006).