

Tracking student skills real-time through a continuous-variable dynamic Bayesian network

Hildo Bijl
Utrecht University of Applied Sciences
info@hildobijl.com

The field of Knowledge Tracing is focused on predicting the success rate of a student for a given skill. Modern methods like Deep Knowledge Tracing provide accurate estimates given enough data, but being based on neural networks they struggle to explain how these estimates are formed. More classical methods like dynamic Bayesian networks can do this, but they cannot give data on the accuracy of their estimates and often struggle to incorporate new observations in real-time due to their high computational load.

This paper presents a novel method, Performance Distribution Tracing (PDT), in which the distribution of the success rate is traced. It uses a dynamic Bayesian network with continuous random variables as nodes. By tracing the success rate distribution, there is always data available on the accuracy of any success rate estimation. In addition, it makes it possible to combine data from similar/related skills to come up with a more informed estimate of success rates. This makes it possible to accurately predict exercise success rates, even when an exercise requires a combination of different skills to solve. And through the use of the right conjugate priors, all distributions are available in analytical form, allowing efficient online updates upon new observations.

An application to an engineering physics course has shown that the PDT algorithm works well in real-life settings. Students mainly appreciate the overview given of their learning process as well as the corresponding practice recommendations.

Keywords: Education, Knowledge Tracing, Bayesian Networks, Data Mining

1. INTRODUCTION

When a teacher/tutor interacts with a student, the teacher tracks (often subconsciously) the student's progress. There is always an idea in the back of the teacher's mind, 'The student is good with these skills and less good with these other skills.' This tracking is subject to three important properties.

1. Success rate estimates are done real-time, taking into account student learning effects.
2. The teacher can clarify the origin and certainty/accuracy of said estimates.
3. The estimates take into account links between a large variety of skills.

The goal of this paper is to set up an algorithm that satisfies the above three requirements, effectively automating part of what the teacher does. The resulting success rate estimates can

then be used to (among others) provide recommendations to students on what skills to practice and which exercises are on the right level to do so, even further automating the tutoring process. This is then applied to an educational setting.

There is a long history behind skill tracking and digital tutoring systems. Most work relates to *Knowledge Components* (KCs), as described by [Koedinger et al. \(2012\)](#). A distinction can be made between *explicit KCs* (facts/principles) and *implicit KCs* (skills). Often skills are practiced and tested through exercises, and the fundamental question then is, ‘What is the chance that the student will do the next exercise correctly?’ This probability is known as the *success rate* and the study of tracking/estimating it is known as *Knowledge Tracing* (KT). Overviews are given by [Sapountzi et al. \(2019\)](#) and [Liu et al. \(2021\)](#) though a brief summary is provided below.

Classical work on KT revolves around *Bayesian Knowledge Tracing* (BKT) ([Corbett and Anderson, 1994](#)). This method directly estimates the success rate, constantly updating this estimate on each subsequent exercise execution using a hidden Markov model. Probabilities of the student slipping (failing despite having mastery) or guessing (succeeding while not having mastery) are taken into account ([Baker et al., 2008](#)) as well as potentially the exercise difficulty ([Pardos and Heffernan, 2011](#)). Extensions exist for further individualization towards students ([Yudelson et al., 2013](#); [Yudelson, 2021](#)).

Instead of using probabilistic models, an alternative set of methods uses logistic models. The fundamentals have been presented by [Cen et al. \(2006\)](#) in the *Learning Factors Analysis* (LFA) method. By using weight factors to account for the difficulty of a KC, the ease of learning of a KC and various other parameters, and by inserting the result into a logistic function, the success rate for a KC is estimated. Extensions/variants include *Performance Factor Analysis* (PFA) ([Pavlik Jr et al., 2009](#)) and *Knowledge Tracing Machines* (KTM) ([Vie and Kashima, 2019](#)).

The methods above mainly work for individual KCs, struggling to take into account links between KCs. This so-called *domain knowledge* can be taken into account in two ways: prespecified or learned by the algorithm. With the introduction and growing popularity of *Massive Open Online Courses* (MOOCs) ([Kaplan and Haenlein, 2016](#)) the amount of data that can be used for KT has increased significantly. As a result, more advanced methods like *Deep Knowledge Tracing* (DKT) have appeared, with fundamental work done by ([Piech et al., 2015](#)). In DKT recurrent neural networks are used to model students’ performance as they work through a course. Though this method requires a large amount of data, and is prone to overfitting ([Guo et al., 2021](#)), it does not require domain knowledge: the algorithm determines for itself which links between KCs exist. Possible extensions include taking into account the learning curve ([Yang et al., 2021](#)). A downside is that neural networks struggle to explain their estimates and their accuracy, while such accountability is becoming increasingly more important in machine learning applications ([Kim and Doshi-Velez, 2021](#)). Hence, to satisfy requirement 2, it is generally better to directly use existing domain knowledge, especially since such domain knowledge in education is generally readily available.

The most common way to take into account domain knowledge is through *Bayesian networks*. When a learning effect should also be considered, these networks are extended to *Dynamic Bayesian Networks* (DBNs) as outlined by [Levy and Mislevy \(2016\)](#). There are various example applications of DBNs to education ([Käser et al., 2017](#); [Reichenberg, 2018](#); [Cui et al., 2019](#); [Levy, 2019](#); [Sheidaei et al., 2022](#)). Though giving very promising results, many of these implementations require numerically complex evaluations, often through methods like Gibbs sampling. This prevents the real-time updating of success rate estimates (requirement 1). On

top of this, all considered DBNs model skill acquisition as binary variables: there either is or isn't mastery. A claim 'There is 40% chance of mastery' hence does not have any accuracy data attached to it. Do we simply not know much about the student and hence loosely guess '40%'? Or has an extensive set of data shown that the student successfully demonstrates a skill 40% of the time? It is left unclear, violating requirement 2.

In this paper a new method is proposed called *Performance Distribution Tracing* (PDT). In PDT the nodes of the DBN are not binary variables (mastery or not) whose chance is estimated, but continuous random variables (the success rates) whose distribution is constantly updated. By choosing the right conjugate priors, it is possible to update the success rate distributions analytically in real-time, satisfying requirement 1. And since the *distributions* of the success rates are known, it gives information on the certainty of the resulting success rate estimates (requirement 2). On top of that, knowing the success rate distributions also allows combining data from multiple sources – for instance multiple KCs/skills – facilitating the use of domain knowledge (requirement 3).

This paper is set up as follows. Chapter 2 studies tracing the success rate distribution of a student for a single skill. In Chapter 3 these methods are extended to take into account exercises involving multiple skills. This is further extended in Chapter 4, studying links between various skills and how these links can be taken into account to get more informed success rate estimates. The PDT method has been incorporated into a tutoring app that has been applied to a basic engineering course. Results of this are discussed in Chapter 5. The paper closes off with conclusions and recommendations in Chapter 6.

2. MODELING THE PROFICIENCY OF A SINGLE SKILL

Consider any single skill A (an implicit knowledge component) that a student might do. This skill A could be as basic as 'multiply two small numbers' or as advanced as 'perform a complex engineering mechanics calculation'. For this skill, we want to track the student's performance, incorporating new observations and making predictions on future executions.

2.1. MATHEMATICALLY DESCRIBING THE PROBABILITY OF SUCCESS

Key to modeling the student's proficiency at skill A is the *success rate*: the chance a that a student performs it successfully. Existing KT methods see a as a deterministic variable that must be found. However, this provides no data on the confidence with which a is known: is it a rough guess or a near-certain estimate?

This problem is solved if we describe the success rate by a random variable \underline{a} . (The underline notation denotes random variables.) In this case \underline{a} is not described by a single value, but by its distribution, expressed through its Probability Density Function (PDF) $f_{\underline{a}}(a)$. Examples of success rate PDFs are shown in Figure 1. Note that \underline{a} only takes values between 0 and 1 because it is a probability. We hence always have $f_{\underline{a}}(a) = 0$ for $a < 0$ or $a > 1$.

Describing the PDF of a success rate \underline{a} is done through basis functions. Specifically, we use the probability density functions of the beta distribution (Johnson et al., 1995, chapter 25) as basis functions. We define the *basis functions* $g_{i,n}(a)$, for some order n , as

$$g_{i,n}(a) = (n + 1) \binom{n}{i} a^i (1 - a)^{n-i}. \quad (1)$$

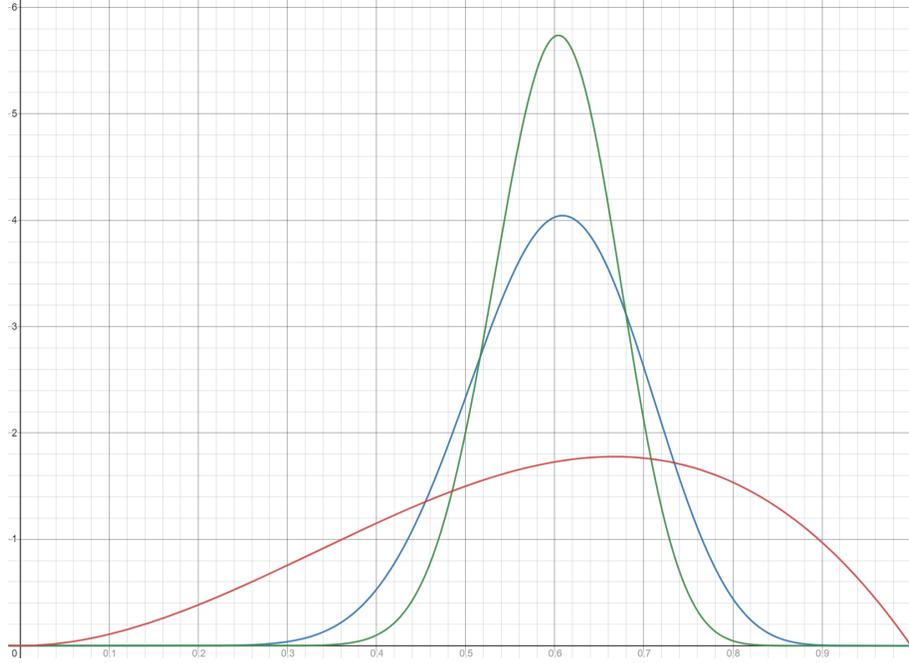


Figure 1: Three example distributions $f_{\underline{a}}(a)$ plotted from 0 to 1. They all have an expected value of 0.6, denoting a 60% estimated success rate, but their peakedness and hence their degree of certainty varies.

Note that the above function is the PDF of a beta distribution with parameters α and β satisfying $\alpha = i + 1$ and $\beta = n - i + 1$. Using the above function as basis function, we write the PDF of \underline{a} as

$$f_{\underline{a}}(a) = \sum_{i=0}^n c_i g_{i,n}(a) = \mathbf{c}^T \mathbf{g}_n(a). \quad (2)$$

Note that the *coefficient vector* \mathbf{c} is the vector with all *coefficients* c_i and identically the *basis function vector* $\mathbf{g}_n(a)$ is the vector function containing all basis functions $g_{i,n}(a)$. Obviously, not all functions can be described through this basis form, but by choosing our priors in the right way (as conjugate priors) the PDFs we encounter do fit this form. It effectively allows us to reduce a complete function to a set of coefficients.

There are some interesting consequences to describing the distribution of \underline{a} in this way. Because $f_{\underline{a}}(a)$ is a PDF, its integral must equal one. Hence,

$$\int_0^1 f_{\underline{a}}(a) da = \int_0^1 \left(\sum_{i=0}^n c_i g_{i,n}(a) \right) da = \sum_{i=0}^n c_i \left(\int_0^1 g_{i,n}(a) da \right) = \sum_{i=0}^n c_i = 1. \quad (3)$$

Note that we have used the fact that the integral over the beta distribution PDF $g_{i,n}(a)$ equals 1. The above shows that the sum of all the coefficients must always equal one, and that there are hence only n degrees of freedom in the $n + 1$ coefficients. If the coefficients add up to one, the coefficients are said to be *normalized*. If the coefficients are not normalized, then it is always possible to normalize them by dividing each coefficient by the sum of all coefficients. In practice this is often done, as well as ensure that all coefficients remain non-negative, to compensate for potential numerical inaccuracies in computers.

Given $f_{\underline{a}}(a)$ we can also calculate the expected probability of success. This is given by

$$E[\underline{a}] = \sum_{i=0}^n c_i \left(\int_0^1 a g_{i,n}(a) da \right) = \sum_{i=0}^n c_i \frac{i+1}{n+2}. \quad (4)$$

More generally, any moment $E[\underline{a}^m]$ can be calculated through

$$E[\underline{a}^m] = \sum_{i=0}^n c_i \frac{(n+1)!}{(n+m+1)!} \frac{(i+m)!}{i!}. \quad (5)$$

Instead of looking at the success rate \underline{a} , we can also study the *failure rate* $1 - \underline{a}$. The distribution of the failure rate, described by $f_{1-\underline{a}}(a)$, can be found by reversing the order of the coefficients c_i . That is,

$$f_{1-\underline{a}}(a) = \sum_{i=0}^n c_{n-i} g_{i,n}(a). \quad (6)$$

This concludes all the relevant properties of our basis function description. It is time to apply it.

2.2. UPDATING THE SKILL SUCCESS RATE DISTRIBUTION

Initially, when no data is present on a skill A , we start with the flat prior $f_{\underline{a}}(a) = 1$. This corresponds to a coefficient vector of $\mathbf{c} = [1]$. The order of this coefficient vector is hence 0: there is no information yet.

Next suppose that, after a student has done $k - 1$ exercises for some skill A , the performance distribution $f_{\underline{a}}(a|D_{k-1})$ is known. Here, D_{k-1} denotes all data related to the first $k - 1$ exercise executions. In other words, the coefficient vector \mathbf{c} is known up to this point. If the student then performs an exercise again (execution k), which is either a success or a failure, how is this data incorporated?

The posterior distribution is given through Bayes' law as

$$f_{\underline{a}}(a|D_k) = \frac{p(D_k|a, D_{k-1}) f_{\underline{a}}(a|D_{k-1})}{p(D_k|D_{k-1})}. \quad (7)$$

Note that, if the last exercise execution was successful, then $p(D_k|a, D_{k-1}) = a$, while on a failure it equals $p(D_k|a, D_{k-1}) = 1 - a$. Also note that $p(D_k|D_{k-1})$ is a constant. Assuming that afterwards we normalize our coefficient vector \mathbf{c} , which we always do, we can safely ignore it.

The posterior distribution $f_{\underline{a}}(a|D_k)$ described above is once more a performance distribution in basis form. That means we can describe it through a new set of coefficients. If the old distribution has order n , the new distribution can be shown to have order $n_* = n + 1$ with coefficients (for $0 \leq i \leq n_*$)

$$c_i^* = \begin{cases} i c_{i-1} & \text{on success,} \\ (n+1-i) c_i & \text{on failing,} \end{cases} \quad (8)$$

where afterwards we apply normalization to the coefficients to get their sum to equal one once more. (The full derivation of the above is omitted for reasons of brevity, but follows directly from (7).) Note that \mathbf{c}^* (the star superscript) denotes the coefficients of the updated distribution, incorporating the latest observation: this notation will be used more often in this paper. Through this update law, we can continuously incorporate more data, allowing the performance distribution to become more peaky over time.

2.3. INFERRING THE SUCCESS RATE FOR THE NEXT EXECUTION

The above has assumed that the actual value of the success rate \underline{a} is a constant: it may be unknown, but it does not vary over time. In reality this is not the case. After every exercise the student may have learned something new, or possibly misinterpreted something. The value of \underline{a} hence slowly shifts, upwards or downwards. This is known as the *learning effect* (Newell and Rosenbloom, 1993).

To account for the learning effect, we use (as is customary in KT) multiple random variables: we define \underline{a}_k as the probability of success of execution k of skill A . Initially we know (as our prior) the distribution $f_{\underline{a}_1}(a)$. After the first exercise, this will be updated using (8) as $f_{\underline{a}_1}(a|D_1)$. But what does this tell us about the probability of success \underline{a}_2 of the second execution?

To be able to say anything about \underline{a}_2 , we must first assume a joint prior between subsequent skill executions \underline{a}_k and \underline{a}_{k+1} : how similar are their success rates? Normally in KT, when the success rates \underline{a} are modeled as binary random variables, this is done through a simple probability table. Instead, because we have continuous random variables, we must set up a joint prior probability density function. A convenient conjugate joint prior is

$$f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1}) = \frac{\mathbf{g}_{n_s}^T(a_k) \mathbf{g}_{n_s}(a_{k+1})}{n_s + 1}, \quad (9)$$

where n_s is known as the *smoothing order*. (This name will be explained later.) The shape of this prior is shown in Figure 2. Note that the shape is symmetric with respect to \underline{a}_k and \underline{a}_{k+1} . It may be sensible to skew this prior to encourage \underline{a}_{k+1} to be (on average) larger than \underline{a}_k , since students generally get better at a skill with more practice instead of worse. However, for simplicity reasons this has not been done.

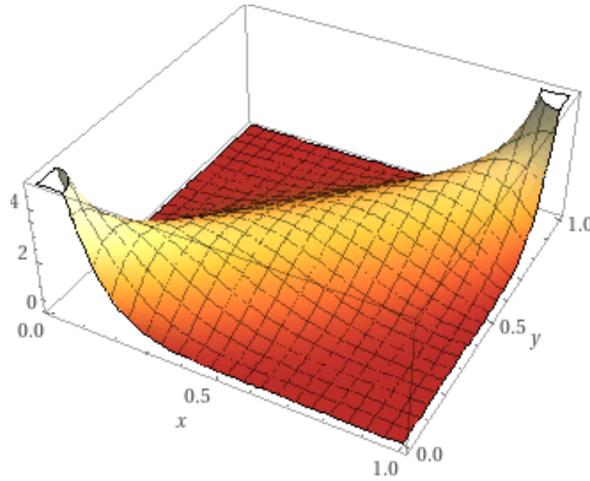


Figure 2: The joint prior for two subsequent skill success rates \underline{a}_k and \underline{a}_{k+1} , for smoothing order $n_s = 10$. Note that, given a value of \underline{a}_k , the value of \underline{a}_{k+1} will very likely be similar. Higher smoothing orders n_s give a more peaked ridge at the diagonal line, and hence assume a stronger correlation between \underline{a}_k and \underline{a}_{k+1} .

The above prior can be used to infer the distribution of \underline{a}_{k+1} from the distribution of \underline{a}_k . For this, we must first use the definition of the conditional distribution to write

$$f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1}|D_k) = f_{\underline{a}_{k+1}}(a_{k+1}|a_k, D_k) f_{\underline{a}_k}(a_k|D_k). \quad (10)$$

Note that, given \underline{a}_k , the data D_k does not add any information to \underline{a}_{k+1} , so it can be omitted. This turns $f_{\underline{a}_{k+1}}(a_{k+1}|a_k, D_k)$ into $f_{\underline{a}_{k+1}}(a_{k+1}|a_k)$. By applying the conditional distribution in reverse, and by subsequently marginalizing over \underline{a}_k , we find

$$f_{\underline{a}_{k+1}}(a_{k+1}|D_k) = \int_0^1 \frac{f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1}) f_{\underline{a}_k}(a_k|D_k)}{f_{\underline{a}_k}(a_k)} da_k. \quad (11)$$

The resulting distribution is once more one in basis form. We can hence solve the integral and derive the resulting coefficients $c_{k+1,i}$. For the mathematically curious, the full derivation is given in Appendix A. We only present the final outcome here as

$$c_{k+1,i} = \sum_{j=0}^{n_s} \binom{i+j}{i} \binom{n_* + n_s - i - j}{n_* - j} c_{k,j}^*, \quad (12)$$

where afterwards the coefficients must be normalized as usual. Note that the coefficients $c_{k,j}^*$ describe $f_{\underline{a}_k}(a_k|D_k)$ (the success rate of execution k given the data on the result of execution k) while $c_{k+1,i}$ describe $f_{\underline{a}_{k+1}}(a_{k+1}|D_k)$ (the success rate of execution $k+1$ without knowing the result of execution $k+1$). Also note that the order of this new set of coefficients $c_{k+1,i}$ is always the smoothing order n_s .

When applying the joint prior in this way, the distribution of \underline{a}_{k+1} is always slightly flatter (less peaked) than the distribution of \underline{a}_k . As a result, this step is known as the *smoothing step*. By applying this smoothing, we effectively incorporate the uncertainty of the learning effect. The smaller n_s is, the more smoothing is applied, while a large value of n_s leaves the distribution nearly unaffected.

The only unanswered question is how to choose n_s . In practice this is done based on various settings. The full method is discussed in Appendix B but it can be summarized in a few simple rules. If a student has practiced the skill a lot, then the learning effect is smaller so n_s is chosen to be larger to have less smoothing. Similarly, if the time since the last skill execution is large, the student may have forgotten a lot, so n_s is chosen to be smaller to have more smoothing. In practice, because the smoothing step is time-dependent, the coefficients that are stored in the database for each student are always the coefficients c_k^* after incorporating new observations yet prior to smoothing.

3. HANDLING EXERCISES COMBINING MULTIPLE SKILLS

So far we have considered exercises in which only a single skill needed to be applied. In practice, many exercises require a variety of skills. How can we incorporate results of these exercises into the corresponding skill success rate distributions? And how do we use this to predict exercise success rates?

3.1. MATHEMATICALLY DEFINING THE EXERCISE SET-UP

The first step is to describe the exercise. Consider an exercise X requiring skills A and B . The exact way this exercise is set up can be described through various operators.

- The and-operator: if a successful execution of the exercise requires a successful execution of both A and B , then we write $X = \text{and}(A, B)$. The chance \underline{x} that a student does both A and B correctly is $\underline{x} = \underline{ab}$.

- The **or**-operator: if an exercise can be solved by either doing A or B , with each method being suitable, then we write $X = \text{or}(A, B)$. In this case $\underline{x} = 1 - (1 - \underline{a})(1 - \underline{b}) = \underline{a} + \underline{b} - \underline{ab}$. Note that $\underline{x} \geq \max(\underline{a}, \underline{b})$ because a student can use the other method/skill to check his/her answer.

In theory an exercise can have any kind of set-up. For instance, we may have $X = \text{and}(A, \text{or}(A, B))$, in which case $x(\underline{a}, \underline{b}) = \underline{a}^2 + \underline{ab} - \underline{a}^2\underline{b}$. The first expression ‘ $\text{and}(A, \text{or}(A, B))$ ’ is known as the *exercise set-up* while the second expression ‘ $\underline{a}^2 + \underline{ab} - \underline{a}^2\underline{b}$ ’ is the corresponding *probability polynomial*. Note that any set-up can be readily turned into a probability polynomial.

3.2. UPDATING THE SKILL SUCCESS RATE DISTRIBUTIONS

Suppose that a student does exercise X correctly, or alternatively he/she fails it. This then provides another data point D_X that can be used to update the distributions of both \underline{a} and \underline{b} . We consider how to update the distribution of \underline{a} . It works identically for \underline{b} , as well as any other potentially involved skill.

Our goal is to find the posterior $f_{\underline{a}}(a|D_X)$. (All previously known data D_A and D_B on skills A and B is also taken into account, but this is left out of the notation for brevity.) By marginalizing the joint distribution $f_{\underline{a}, \underline{b}}(a, b|D_X)$ with respect to \underline{b} , followed up with Bayes’ law, we find

$$f_{\underline{a}}(a|D_X) = \int_0^1 f_{\underline{a}, \underline{b}}(a, b|D_X) db = \int_0^1 \frac{p(D_X|a, b)f_{\underline{a}, \underline{b}}(a, b)}{p(D_X)} db. \quad (13)$$

Note that $p(D_X)$ is a constant, not depending on \underline{a} , so we can leave it out if we normalize coefficients afterwards. Also note that $p(D_X|a, b)$ follows from the probability polynomial $x(a, b)$ as

$$p(D_X|a, b) = \begin{cases} x(a, b) & \text{on success,} \\ 1 - x(a, b) & \text{on failing.} \end{cases} \quad (14)$$

We generally also assume independence of \underline{a} and \underline{b} . In this case we can write $f_{\underline{a}, \underline{b}}(a, b)$ as $f_{\underline{a}}(a)f_{\underline{b}}(b)$ and hence get

$$f_{\underline{a}}(a|D_X) \sim \left(\int_0^1 p(D_X|a, b)f_{\underline{b}}(b) db \right) f_{\underline{a}}(a) = h_{\underline{a}}(a)f_{\underline{a}}(a), \quad (15)$$

where we have defined $h_{\underline{a}}(a)$ as the part between brackets in the above equation. Note that $h_{\underline{a}}(a)$ is merely, for a given a , the expected value of (on success) the probability polynomial $x(a, \underline{b})$ or (upon failure) the inverse $1 - x(a, \underline{b})$, with respect to all other random variables.

$h_{\underline{a}}(a)$ is a polynomial that only depends on a . This means we can write it as

$$h_{\underline{a}}(a) = \sum_{i=0}^{n_p} k_i a^i, \quad (16)$$

for some polynomial order n_p and some set of constants k_0, \dots, k_{n_p} . The exact value of these constants can be found from the probability polynomial with the help of (5).

Before we continue, we must first put $h_{\underline{a}}(a)$ in an alternate form. We want to write $h_{\underline{a}}(a)$ as

$$h_{\underline{a}}(a) = \sum_{i=0}^{n_p} k'_i \binom{n_p}{i} a^i (1 - a)^{n_p - i}, \quad (17)$$

for another set of constants k'_0, \dots, k'_{n_p} . The link between these two sets of constants is

$$\binom{n_p}{n_p - i} k'_i = \sum_{j=0}^i \binom{n_p - j}{n_p - i} k'_j. \quad (18)$$

Thanks to this alternate way of writing $h_{\underline{a}}(a)$, we can now update the coefficients of $f_{\underline{a}}(a)$. If this PDF used to have order n , with coefficients c_i for $0 \leq i \leq n$, then the posterior PDF $f_{\underline{a}}(a|D_X)$ will be of order $n_* = n + n_p$. Its coefficients c_i^* , with $0 \leq i \leq n_*$, can be found through

$$c_i^* = \sum_{j=\max(0, i-n)}^{\min(n_p, i)} \binom{i}{j} \binom{n_* - i}{n_p - j} c_{i-j} k'_j, \quad (19)$$

where afterwards coefficients must once more be normalized. Note that this method is a generalization of the update law of (8). It allows us to update the distributions of any skills A , B , and potentially more, based on exercises with any set-up involving these skills.

In practice this update law has some intuitive consequences. Consider the situation where a student has mostly mastered skill A and is still struggling with skill B . If he/she then tries an exercise with set-up (A, B) and fails, then the probability theory inherent in the above method automatically searches for ‘blame’. In this case, the fault most likely lies in skill B , so that skill is more strongly penalized, while the success rate for skill A is only slightly adjusted downwards. This behavior is exactly what seems sensible, showing that the system works in the desired way.

3.3. INFERRING THE SUCCESS RATE FOR THE EXERCISE

When considering which exercise to present to the student, it is useful to know how likely the student is to successfully complete each exercise. Given an exercise X with known set-up and probability polynomial, we want to know the success rate \underline{x} , or more simply put the expected value $E[\underline{x}]$ of it.

The expected value $E[\underline{x}]$ is actually straightforward to find. If we assume that \underline{a} and \underline{b} are independent, then the expected value of the probability polynomial $E[x(\underline{a}, \underline{b})]$ follows directly from the application of (5).

In some cases it is also useful to know the complete distribution of \underline{x} . This tells us (among others) how certain we are of the estimated success rate $E[\underline{x}]$. Finding the distribution of \underline{x} is rather involved, mostly because it generally *cannot* be described through our basis functions. However, we can define a nearly identical random variable $\hat{\underline{x}}$ whose distribution *can* be described through basis functions.

We want the exercise success rate \underline{x} and this new random variable $\hat{\underline{x}}$ to have as similar distributions as possible. To capture this similarity, we define the prior $f_{\hat{\underline{x}}, \underline{x}}(\hat{\underline{x}}, \underline{x})$, identically to (9), as

$$f_{\hat{\underline{x}}, \underline{x}}(\hat{\underline{x}}, \underline{x}) = \frac{\mathbf{g}_{n_i}^T(\hat{\underline{x}}) \mathbf{g}_{n_i}(\underline{x})}{n_i + 1}, \quad (20)$$

for some *inference smoothing order* n_i . In practice, the order n_i for this application is often chosen to be on the lower side ($n_i \approx 10$) for practical reasons we will soon see.

We want to find the posterior distribution $f_{\hat{x}}(\hat{x}|D_{A,B})$ given all data $D_{A,B}$ on skills A and B . Identically to how we found (11) we can find

$$f_{\hat{x}}(\hat{x}|D_{A,B}) = \int_0^1 \frac{f_{\hat{x},x}(\hat{x},x)f_x(x|D_{A,B})}{f_x(x)} dx. \quad (21)$$

This integral can subsequently be solved and rewritten into the basis function form. This gives us a set of coefficients $c_{x,i}$ describing the posterior distribution of \hat{x} . These coefficients satisfy

$$c_{x,i} = \int_0^1 \int_0^1 g_{i,n_i}(x(a,b))f_a(a|D_A)f_b(b|D_B) da db, \quad (22)$$

where subsequent normalization must be applied, as usual. Alternatively, we can say that

$$\mathbf{c}_x = \mathbb{E}[\mathbf{g}_{\mathbf{n}_i}(x(\underline{a}, \underline{b}))]. \quad (23)$$

So the coefficients $c_{x,i}$ are simply the expected values of the basis functions g_{i,n_i} when inserting the probability polynomial into these basis functions. This idea works identically if more than two skills A and B are involved, and it is known as the *inference step* of the algorithm.

The remaining question is how to calculate these coefficients from the coefficients of \underline{a} and \underline{b} . Note here that each basis function g_{i,n_i} is also a polynomial function. Inserting a polynomial into a polynomial function once more gives a polynomial. If g_{i,n_i} does not have too large powers (that is, we keep n_i reasonably small) this can be expanded through a binomial expansion. Subsequently, the expected value can be found through (5).

Finally, it might be interesting to discuss the intuitive view of this new variable \hat{x} : how can we interpret it? Some may argue that x , which follows directly from \underline{a} and \underline{b} through the probability polynomial $x(\underline{a}, \underline{b})$, is *not* the most accurate way to describe whether the student will successfully complete exercise X . After all, when doing exercise X , the student must also recognize the steps to take, which has nothing to do with A and B individually, but does affect the success rate for exercise X . As a result, we must define a *true success rate* \hat{x} , which takes the estimated success rate x based on the skills A and B , and adds some uncertainty (read: smoothing of the distribution) on top of this.

4. DESCRIBING RELATIONS BETWEEN SKILLS

The strength of the PDT algorithm lies in its ability to link related skills. A course generally does not consist of a collection of unrelated skills. The skills build up on each other in an often complex way. This can be modeled and incorporated. The complete method may seem convoluted at first, but the overview at the end (Figure 3) should clarify this.

4.1. DEFINING THE SKILL SET-UP

A *composite skill* is a skill made up of multiple smaller *subskills*. For instance, to calculate $2 + 3 \cdot 4$ you must first figure out the order of operations, then apply multiplication and finally apply addition. In this simple example, we could say that skill S (evaluating basic expressions) has a set-up of $\text{and}(A, B, C)$, with A , B and C the aforementioned subskills.

Skills have a set-up similar to exercises, but while exercises always have a *deterministic set-up*, only using and and or operations, skills may have a *non-deterministic set-up*. In the above example, instead of using addition, we may perhaps require subtraction half the times, resulting in a varying set-up. We therefore add the following non-deterministic operations.

- The pick-operator: from a list of subskills, we select only one. For instance, if a skill first requires either A or B (but always only one of them) followed by C , then we write $S = \text{and}(\text{pick}(A, B), C)$. To find the probability polynomial, we may reduce $\text{pick}(A, B)$ to $\frac{1}{2}\underline{a} + \frac{1}{2}\underline{b}$. Extensions exist where multiple skills from a list are selected and/or weights are applied upon selection.
- The part-operator: if a skill only requires a skill A in a part p of the cases, always followed by a skill B , then we may write $S = \text{and}(\text{part}(A, p), B)$. Alternatively, if sometimes we can apply either A or B , and sometimes only B , we may write $S = \text{or}(\text{part}(A, p), B)$.

For the part-operator, the probability polynomial depends on the surrounding operator. On a surrounding and-operator we reduce $\text{part}(A, p)$ to $1 - p(1 - \underline{a})$, while on a surrounding or-operator $\text{part}(A, p)$ becomes $p\underline{a}$.

With these extra operators, it is still always possible to turn a skill set-up into a probability polynomial. As a result, using the merging methods from Section 3.3, we can always predict the success rate of a skill S based on data from its subskills, just like for an exercise X .

4.2. MERGING OBSERVATIONS ON SKILLS AND SUBSKILLS

Consider the situation where a composite skill S has subskills A and B , and where data is available on *all* these skills.

- First the student practices subskills A and B , giving data $D_{A,B}$. Using the methods from Section 3.3 – specifically the coefficients from (23) – we can hence infer $f_{\underline{s}}(s|D_{A,B})$.
- Subsequently the student practices skill S directly, giving data D_S . Using *only* this data, we can also find a distribution $f_{\underline{s}}(s|D_S)$. This is done through (8) or more generally with (19).

How can this data then be ‘merged’ into $f_{\underline{s}}(s|D_{A,B}, D_S)$? To answer this question, we must apply Bayes’ law,

$$f_{\underline{s}}(s|D_{A,B}, D_S) = \frac{p(D_{A,B}, D_S|s)f_{\underline{s}}(s)}{p(D_{A,B}, D_S)}. \quad (24)$$

Assuming that the observation sets $D_{A,B}$ and D_S are conditionally independent given \underline{s} , we may reduce $p(D_{A,B}, D_S|s)$ to $p(D_{A,B}|s)p(D_S|s)$. Applying Bayes’ law in the reverse direction subsequently gives

$$f_{\underline{s}}(s|D_{A,B}, D_S) = \frac{p(D_{A,B})p(D_S)}{p(D_{A,B}, D_S)} \frac{f_{\underline{s}}(s|D_{A,B})f_{\underline{s}}(s|D_S)}{f_{\underline{s}}(s)}. \quad (25)$$

Note that all terms apart from $f_{\underline{s}}(s|D_{A,B})$ and $f_{\underline{s}}(s|D_S)$ are constants. We can therefore find the posterior distribution of \underline{s} simply by multiplying these separately inferred distributions and normalizing the result.

Let’s say that $f_{\underline{s}}(s|D_{A,B})$ has order $n_{A,B}$ and coefficients $c_i^{A,B}$, and identically $f_{\underline{s}}(s|D_S)$ has order n_S and coefficients c_i^S . The resulting posterior distribution $f_{\underline{s}}(s|D_{A,B}, D_S)$ then has order $n_m = n_{A,B} + n_S$ and coefficients c_i^m satisfying

$$c_i^m = \sum_{j=\max(0, i-n_{A,B})}^{\min(n_S, i)} \binom{i}{j} \binom{n_m - i}{n_S - j} c_{i-j}^{A,B} c_j^S, \quad (26)$$

where the coefficients must be normalized afterwards. Note the similarity of the above inference law with the update law (19).

The above method effectively shows how to merge two distributions of $f_{\underline{s}}(s)$, based on different, separate and conditionally independent data, together into one distribution. As a result the above method is known as *merging distributions*, which is also why we use a subscript/superscript m here.

4.3. INCORPORATING CORRELATIONS BETWEEN SKILLS

Next to parent/child dependencies between skills, we in practice also encounter correlated skills. For instance, the skill ‘expand brackets with numbers’ with example exercise ‘expand $(2+3) \cdot 4$ ’ is generally strongly correlated with the skill ‘expand brackets with letters’ with example exercise ‘expand $(a+b) \cdot c$ ’. Note that these skills are similar but not identical: many a mathematics teacher can confirm that students who have mastered the first skill often still struggle with the second one. Nevertheless, someone’s performance on the first skill *does* give a little bit of information about the expected performance on the second skill.

To model this, we consider a skill S with a correlated skill R . We write their success rates as \underline{s} and \underline{r} , respectively. We now want to use data D_R on skill R to infer the distribution of \underline{s} . Once more, the first step is to define a joint prior between \underline{r} and \underline{s} . Identically to (9) we assume

$$f_{\underline{r}, \underline{s}}(r, s) = \frac{\mathbf{g}_{n_c}^T(r) \mathbf{g}_{n_c}(s)}{n_c + 1}, \quad (27)$$

for some *correlation smoothing order* n_c . The stronger the correlation, the higher n_c must be. In practice correlations are often weak, so n_c does not go above a value of 10.

Once the joint prior is defined, we notice that the problem we have here is exactly the same as that from Section 2.3. The distribution $f_{\underline{s}}(s|D_R)$ is hence simply the smoothed version of $f_{\underline{r}}(r|D_R)$, using smoothing order n_c . The coefficients follow from (12).

Once $f_{\underline{s}}(s|D_R)$ has been determined, we can merge it into $f_{\underline{s}}(s|D_S)$ using the methods from Section 4.2; specifically through (26). This gives us the posterior $f_{\underline{s}}(s|D_R, D_S)$. After all, the problem is identical to when we tried to derive the posterior distribution of \underline{s} based on data from subskills. Additionally, in case data from subskills is *also* present, we can merge that in too. The order of merging does not matter.

A more complex problem appears when there are groups of correlated skills. Imagine there are two skills Q and R that are both correlated with skill S . You could treat all these correlations individually and separately merge $f_{\underline{s}}(s|D_Q)$ and $f_{\underline{s}}(s|D_R)$ into $f_{\underline{s}}(s|D_S)$. However, this neglects correlations between skills Q and R which are most likely present. To take this into account, we can define a joint prior like

$$f_{\underline{q}, \underline{r}, \underline{s}}(q, r, s) = \frac{1}{n_c + 1} \sum_{i=0}^{n_c} g_{i, n_c}(q) g_{i, n_c}(r) g_{i, n_c}(s). \quad (28)$$

By applying Bayes’ law back and forth a few times, similar to what we did for (25), we can find that the joint posterior distribution given data on Q and R is proportional to

$$f_{\underline{q}, \underline{r}, \underline{s}}(q, r, s|D_Q, D_R) \sim f_{\underline{q}}(q|D_Q) f_{\underline{r}}(r|D_R) f_{\underline{q}, \underline{r}, \underline{s}}(q, r, s). \quad (29)$$

Marginalizing over \underline{q} and \underline{r} then results in the posterior distribution $f_{\underline{s}}(s|D_{Q,R})$. The lengthy mathematics are omitted for reasons of brevity, but the final procedure is as follows.

- Start with the coefficients c_i^q and c_i^r describing $f_q(q|D_Q)$ and $f_r(r|D_R)$. (Ensure that the learning effect and time decay are already taken into account.)
- Smooth these distributions separately with (12) giving coefficients $c_i^{q'}$ and $c_i^{r'}$ describing $f_{\underline{s}}(s|D_Q)$ and $f_{\underline{s}}(s|D_R)$.
- Multiply the coefficients $c_i^{q'}$ and $c_i^{r'}$ element-wise as

$$c_i^{s'} = c_i^{q'} \cdot c_i^{r'} \quad (30)$$

to find the coefficients $c_i^{s'}$ describing $f_{\underline{s}}(s|D_Q, D_R)$.

- Merge the distributions $f_{\underline{s}}(s|D_S)$ (described by c_i^s) and $f_{\underline{s}}(s|D_Q, D_R)$ (described by $c_i^{s'}$) using (26). The resulting coefficients describe $f_{\underline{s}}(s|D_Q, D_R, D_S)$ as intended.

This allows us to even take into account groups of correlated skills. In practice this is hardly ever necessary: correlated duos do occur, but triplets or quadruplets are extremely rare and often indicate an inconveniently chosen course set-up.

4.4. OVERVIEW OF PERFORMANCE DISTRIBUTION TRACING STEPS

We have so far seen how, given data on a skill S , its subskills A, B, \dots and its correlated skills Q, R, \dots , the posterior distribution of \underline{s} is determined. Let's create an overview. The entire procedure is summarized in Figure 3.

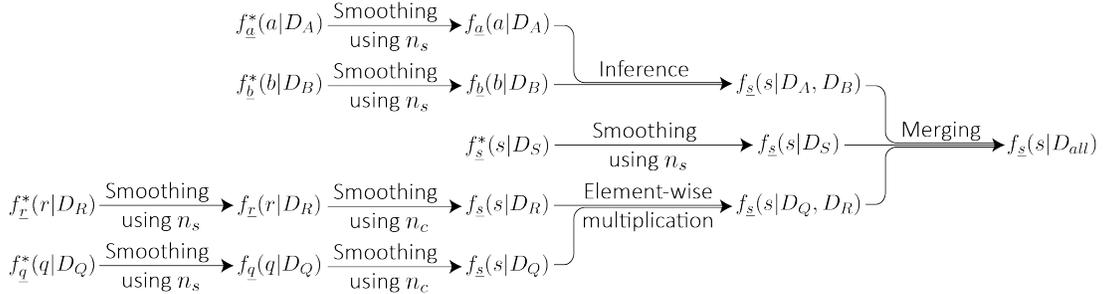


Figure 3: Overview of how to incorporate data from subskills and correlated skills. The star distributions – stored in the database – denote the distributions *after* incorporating a new observation from an exercise, but *before* applying any smoothing for practice and time decay. For smoothing use (12), for inference use (23) and for merging use (26).

Effectively, the coefficients stored in the database are always the ones describing a distribution of \underline{s} *after* data on an exercise execution has been incorporated, but *before* smoothing has been performed to take into account the learning effect (practice decay) as well as time decay. So before anything is done, the stored coefficients c^* are first always smoothed using (12) with smoothing order n_s . Note that, as explained in Appendix B, the value of n_s depends (among others) on the time passed since the last exercise, so it is likely to have a different value for every skill.

If a skill is a composite skill, and hence has subskills, an inference step is applied using the skill set-up. This is done through (23). Afterwards distributions are merged using (26).

Similarly, if a skill is correlated with other skills, then these other skills are first smoothed with order n_c using (12). Here n_c depends on the strength of the correlation. If there is more than

one correlated skill, the resulting coefficients for each of these correlated skills are multiplied element-wise using (30). This distribution is then once more merged in using (26). The merging order is irrelevant.

Through this method, we can always find the most informed distribution of \underline{s} , incorporating as much data as is possible/sensible. And, because we always use distributions, the certainty/accuracy of said data is always correctly taken into account. In theory, if correlated skills Q, R, \dots also have subskills, or if subskills A, B, \dots also have correlated skills or further subskills, these can be taken into account as well. In practice, due to the repetitive smoothing, the tiny bit of extra information resulting from this is generally not considered worth the effort. Hence the plan of Figure 3 suffices.

Once the PDT algorithm has determined, for every relevant skill S , the most well-informed performance distribution $f_{\underline{s}}(s|D_{all})$, this information can then be applied. It can be used to inform a student of his/her progress, to recommend skills to practice, or to select exercises to try out. And once a new data point becomes available – the student tries and succeeds at or fails a certain exercise – update law (19) can be used to directly update all relevant stored coefficients.

Note that, when using update law (19), the coefficients c_i are the coefficients of $f_{\underline{s}}(s|D_S)$ *without* any additional information, while the polynomial coefficients k'_j do have *all* available data taken into account, also from subskills and correlated skills. This may seem strange: why store the distributions $f_{\underline{s}}(s|D_S)$ and not the distributions $f_{\underline{s}}(s|D_{all})$? The reason here is practical: if we did, then when a user practices a subskill A , we would have to update *all* its parent skills. For very fundamental subskills like solving a linear equation, there may be countless parent skills, possibly also including future skills that have not been added to the system yet. As a result, storing coefficients for distributions of \underline{s} based on *only* observations directly related to skill S is practically much more sensible.

5. PRACTICAL APPLICATION

Because the PDT is fundamentally different from other somewhat similar algorithms, it is difficult to make a direct comparison. Existing data sets often do not have sufficient domain knowledge documented, which is required by PDT. Instead, to test the PDT algorithm, it has been built into an existing tutoring web-app called Step-Wise for some limited tests. We will study how this app works, what students thought of the success rate estimates, and how the recommendations made by the app turned out.

5.1. GENERAL FUNCTIONING OF THE STEP-WISE APP

The Step-Wise app is a practice support app. A university-level Thermodynamics course has been split up into a large learning tree of 35 skills, many with complex interdependencies. The students were taught about the theory in regular lectures. They were then instructed to practice using the Step-Wise app. For each of the 35 skills, a set of randomly generated exercises has been added, effectively resulting in an infinite amount of practice material for the students. An example Step-Wise exercise is shown in Figure 4.

When using the Step-Wise app, students are continuously kept up-to-date on their estimated performance for every skill. This is done through the *skill indicator*: a small colored sphere, as shown in Figure 4. For a skill S , the color of the sphere would be determined by the expected success rate $E[s]$: red for 20%, gradually shifting to green at 80% and even turning blue on

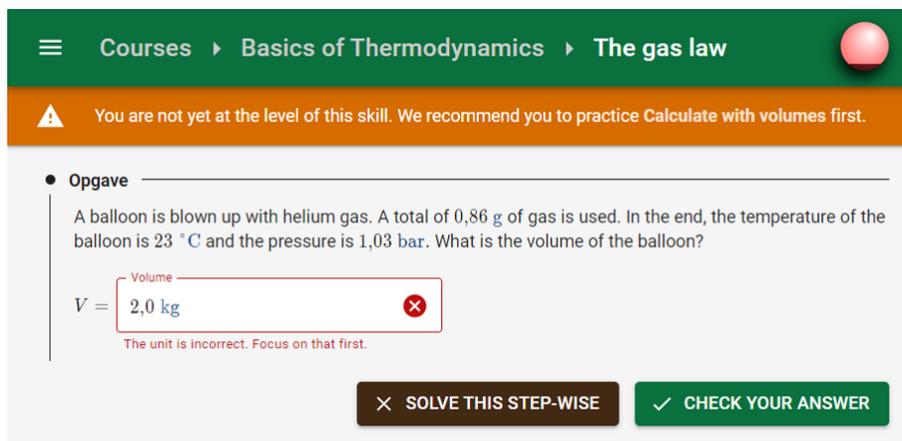


Figure 4: A screenshot of an exercise from the Step-Wise app. Numbers are randomly generated. Every exercise is connected to a skill, and the success rate estimate is continuously shown as a colored sphere on the top right. By hovering over the sphere, the student sees an explanation of what this score is and is based on. Exercises are programmed to provide individualized feedback.

higher percentages. In addition, if the estimate was rather uncertain – the distribution of \underline{s} would be flat – this color would be partly grayed out. The more certain the PDT algorithm gets of its estimates, the brighter the colors become.

Each exercise in the Step-Wise app also has an extra feature. If the student cannot figure it out, he/she can choose the option ‘Solve this Step-Wise’. This splits the exercise up into steps, each representing a subskill needed to solve the exercise. The student can then do these steps one by one. Students appreciate this approach, as it helps them to better understand the content matter. Also, because they have earlier proven mastery for each of these subskills, they are certain they will always eventually be able to solve the exercise. For the PDT algorithm this Step-Wise approach proves very useful too, as it allows it to pinpoint in which subskill the student is deficient. The Step-Wise app can then recommend the student to practice this subskill some more separately.

A numerical analysis of the results of applying the Step-Wise app could not be provided. After all, the sample group (60 students) was too small, the regular variation between students was too large, and even if there were significant changes then, due to the corona crisis, these could not necessarily be attributed to this intervention. However, through a study of how students used the app, and detailed interviews held with said students, conclusions can still be drawn.

5.2. EXPERIENCE WITH SUCCESS RATE ESTIMATES

Students are generally happy with the continuous success rate estimates provided by the app. Observations can be grouped into three categories.

- In advance, when starting to practice a new skill, the PDT algorithm can use data on subskills to estimate how proficient a student would be on that skill. One student for instance said, ‘I know when I’m going to study a difficult skill, so I can mentally prepare. That helped.’
- During practice, the success rate indicators mainly motivated students to work harder. ‘I really wanted the indicator to be bright green. I worked deep into the night to get

there.’ This was often not without any frustration. ‘On a silly calculation error the app already counts it as a fail, dropping my carefully accumulated score. On an exam I do get partial points for a correct method!’ This actually shows that the PDT algorithm works as intended. After all, it only counts successes and failures and not partial successes. A positive side-effect is that the students started checking their own work for errors more often, resulting in fewer calculation errors down the line.

- After practice, students were once more happy with the skill level indicator. ‘It is very reassuring to have an app tell me that I have practiced enough.’

All in all, providing students with skill level estimates has shown to increase motivation, increase the time spent practicing and help students become more consistent in their work.

In general there were no reported cases where the PDT algorithm gave illogical skill level estimates. As a result, it can be concluded that the PDT works as intended. In addition, the data on the certainty of the estimates, as provided to the students in the coloring scheme, had a mild increase on the motivation of the students. They wanted to get the indicator to be bright green. We can conclude that this extra information was also useful, albeit only mildly.

5.3. RECOMMENDATIONS OF SKILLS AND EXERCISES

Based on the skill level estimates, students are also given recommendations on what skills to practice. If students practice a skill they have already mastered, they would be advised, ‘Try [this more complex] skill instead.’ Or similarly, if students would practice a skill for which a prerequisite was not mastered, they would be advised, ‘You may want to try practicing [this more basic skill] first.’ (See Figure 4 for an example of the latter.) In the thermodynamics course, the students initially liked these recommendations since it structured their learning process, and they eagerly and obediently followed them. This did reduce significantly over time. Often at some point a student would skip a skill. ‘I think I already mastered it, but the app doesn’t believe it, because I keep making small calculation errors.’ Afterwards, the app would continuously send the student back to this skipped skill, much to the chagrin of said student. This caused the recommendations to be ignored more often. Small tweaks in the user-friendliness of the recommendation system are hence still necessary.

The Step-Wise app also includes an intelligent exercise selection method, roughly similar to (Wu et al., 2020). Effectively, every skill has three or four exercises connected to it, each with randomly chosen parameters. Though they are similar, each exercise does have minor variations, like a different step somewhere along the way or simply an entire extra step. As a result, when we use the inference step of the PDT algorithm to estimate the success rate of all these exercises, we do get different values. The Step-Wise app was programmed to take this into account. Exercises were still chosen randomly, but if an exercise had an estimated success rate close to 50% then it would be much more likely to be chosen than an exercise with a 20% (too hard) or 80% (too easy) success rate. For the teachers this add-on was positive: they noticed the students starting with the easier exercises and later on getting harder ones. However, for the students this functionality caused a more monotonous practice session. ‘I keep getting the balloon exercise. Give me something else for a change!’ Adding more exercises for each skill will mostly circumvent this problem.

6. CONCLUSIONS AND RECOMMENDATIONS

This paper has presented the Performance Distribution Tracing algorithm. For any skill S , this algorithm continuously traces the distribution of the success rate \underline{s} of this skill, directly applying updates when new observations become available. This is done while taking into account the learning effect and practice decay. If there is an exercise X requiring the usage of multiple skills, in any possible set-up, then the success rate \underline{x} of this exercise can also be inferred, and observations from this exercise can be taken into account for all related skills. In addition, if domain knowledge is present, in terms of links between skills, then this knowledge can be taken into account, resulting in more informed estimates of the success rates of all respective skills.

An application to a university Thermodynamics course with 60 students has shown that the PDT algorithm works in practice. The data it provides, including data on the uncertainty of its estimates, has proven to be helpful at automatically coaching students in their learning process, and it increases the motivation of students to practice longer and more extensively. Due to limited numbers, no concrete numeric information about the accuracy of the algorithm could be obtained.

An extension of the PDT algorithm may concern individualization: can the algorithm be tuned to a specific student, resulting in better updates to give more accurate estimates? Possibly the smoothing order n_s or even the prior $f_{\underline{s}}(s)$ can be adjusted on-the-go, although the latter would require a renewed derivation of all equations in the PDT algorithm. Or possibly a non-symmetrical smoothing step can be used, to better take into account the learning effect.

Another possible extension includes taking into account the cooperation of multiple students. If two students work together, and the distributions of their individual success rates \underline{s}_1 and \underline{s}_2 for some skill S are known, how can we predict the success rate $\underline{s}_{1,2}$ of the two students working together? And if they together successfully execute an exercise, how can we update their individual skill distributions?

A. MATHEMATICAL DERIVATION OF COEFFICIENT EQUATIONS

This appendix has as goal to show how an expression with PDFs like (11) can be turned into an expression for coefficients like (12). We will go through the full derivation, first isolating coefficients and then solving the integral inside their expression.

A.1. ISOLATING THE COEFFICIENTS

Our first task is isolating the coefficients. We start at (11), for easy reading repeated as

$$f_{\underline{a}_{k+1}}(a_{k+1}|D_k) = \int_0^1 \frac{f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1}) f_{\underline{a}_k}(a_k|D_k)}{f_{\underline{a}_k}(a_k)} da_k. \quad (31)$$

In this expression we have the joint prior $f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1})$ defined by (9), or written in its sum notation as

$$f_{\underline{a}_k, \underline{a}_{k+1}}(a_k, a_{k+1}) = \frac{1}{n_s + 1} \sum_{i=0}^{n_s} g_{i, n_s}(a_k) g_{i, n_s}(a_{k+1}). \quad (32)$$

Similarly, the posterior distribution of \underline{a}_k , described by $f_{\underline{a}_k}(a_k|D_k)$, follows from (2) as

$$f_{\underline{a}_k}(a_k|D_k) = \sum_{i=0}^{n_*} c_{k,i}^* g_{i, n_*}(a_k), \quad (33)$$

where $c_{k,i}^*$ are the coefficients describing \underline{a}_k . These coefficients are known from (8) or alternatively (19). Finally, we have the prior $f_{\underline{a}_k}(a_k)$ which equals 1 and can hence be ignored.

Inserting all the above expressions into (31) turns it into

$$\int_0^1 \frac{1}{n_s + 1} \left(\sum_{i=0}^{n_s} g_{i,n_s}(a_k) g_{i,n_s}(a_{k+1}) \right) \left(\sum_{i=0}^{n_*} c_{k,i}^* g_{i,n_*}(a_k) \right) da_k. \quad (34)$$

We can rearrange the above to write it as

$$\sum_{i=0}^{n_s} \left(\frac{1}{n_s + 1} \int_0^1 \left(\sum_{j=0}^{n_*} c_{k,j}^* g_{j,n_*}(a_k) \right) g_{i,n_s}(a_k) da_k \right) g_{i,n_s}(a_{k+1}). \quad (35)$$

Comparing this with the standard basis function notation (2), we can directly see that the coefficients $c_{k+1,i}$ describing $f_{\underline{a}_{k+1}}(a_{k+1}|D_k)$ equal

$$c_{k+1,i} = \frac{1}{n_s + 1} \int_0^1 \sum_{j=0}^{n_*} c_{k,j}^* g_{j,n_*}(a_k) g_{i,n_s}(a_k) da_k. \quad (36)$$

The coefficients have hence been isolated. What remains is solving the above integral.

A.2. SOLVING THE INTEGRAL

Consider (36). Note that, if we normalize the coefficients $c_{k+1,i}$ afterwards, we can safely ignore constant multiplications: factors not depending on i . Hence $\frac{1}{n_s+1}$ drops out. If we expand the basis functions using their definition (1), and ignore the constant factors $(n+1)$ within, we can write the above expression for $c_{k+1,i}$ as

$$c_{k+1,i} = \int_0^1 \sum_{j=0}^{n_*} c_{k,j}^* \binom{n_*}{j} a_k^j (1-a_k)^{n_*-j} \binom{n_s}{i} a_k^i (1-a_k)^{n_s-i} da_k. \quad (37)$$

Rearranging the above turns it into

$$c_{k+1,i} = \sum_{j=0}^{n_*} c_{k,j}^* \binom{n_s}{i} \binom{n_*}{j} \int_0^1 a_k^{i+j} (1-a_k)^{n_*+n_s-i-j} da_k. \quad (38)$$

Note that in general, for any integer values of i and n with $0 \leq i \leq n$, it holds that

$$\int_0^1 g_{i,n}(x) dx = \int_0^1 (n+1) \binom{n}{i} x^i (1-x)^{n-i} dx = 1. \quad (39)$$

After all, integrating over the PDF of a beta distribution must result in one. Using the above relation, we can solve the integral from (38). Ignoring constant multiplications as well, we find

$$c_{k+1,i} = \sum_{j=0}^{n_*} \frac{\binom{n_s}{i} \binom{n_*}{j}}{\binom{n_*+n_s}{i+j}} c_{k,j}^*. \quad (40)$$

By applying the definition of the binomial, the above can be expanded into

$$c_{k+1,i} = \sum_{j=0}^{n_*} \frac{\frac{n_s!}{i!(n_s-i)!} \frac{n_*!}{j!(n_*-j)!}}{\frac{(n_*+n_s)!}{(i+j)!(n_*+n_s-i-j)!}} c_{k,j}^* = \sum_{j=0}^{n_*} \frac{\frac{(i+j)!}{i!j!} \frac{(n_*+n_s-i-j)!}{(n_*+n_s)!(i+j)!}}{\frac{(n_*+n_s)!}{n_*!n_s!}} c_{k,j}^*. \quad (41)$$

The denominator here equals $\binom{n_*+n_s}{n_*}$, which is a constant, so it can be ignored. We remain with

$$c_{k+1,i} = \sum_{j=0}^{n_s} \binom{i+j}{i} \binom{n_*+n_s-i-j}{n_*-j} c_{k,j}^*, \quad (42)$$

which is the final result (12) that we wanted to get. Note that in theory it is possible to take into account all constants in the derivation of $c_{k+1,i}$ too, but in practice it is much easier to ignore constants and simply normalize coefficients afterwards.

B. CHOOSING AND APPLYING THE SMOOTHING ORDER n_s

In Section 2.3 it was discussed that, to incorporate the learning effect and practice decay, the distribution of the success rate \underline{a} would be smoothed after every exercise. This is done with some smoothing order n_s . The question remains: how should n_s be chosen? This is not a mathematical matter, but one of settings and preferences.

B.1. LINKS BETWEEN ORIGINAL AND SMOOTHED DISTRIBUTIONS

Consider a skill A . Let's write the success rate *before* smoothing as \underline{a}_k and the success rate *after* smoothing as \underline{a}_{k+1} . In this case, based on (12), a link can be determined between their expected values $E[\underline{a}_k]$ and $E[\underline{a}_{k+1}]$. To be precise, this link satisfies

$$\left(E[\underline{a}_{k+1}] - \frac{1}{2} \right) = \frac{n_s}{n_s + 2} \left(E[\underline{a}_k] - \frac{1}{2} \right). \quad (43)$$

In words, the smoothed success rate \underline{a}_{k+1} always has a mean closer to 1/2 than the original success rate \underline{a}_k . To be precise, the reduction factor that is applied is

$$r = \frac{n_s}{n_s + 2}. \quad (44)$$

We call r the *decay ratio*. If it equals 1 there is no smoothing, while a decay ratio of 0 means everything is forgotten. There are now two questions: which decay ratio r is appropriate? And how do we turn it into a smoothing order n_s ?

B.2. CHOOSING A DECAY RATIO

Smoothing should take into account both the learning effect and time decay. For time decay an exponential decay seems appropriate. Something like

$$r = (1/2)^{t/t_{1/2}}, \quad (45)$$

where t is the time since the last exercise and $t_{1/2}$ is the time after which the student has 'lost' half of its skills. The latter is generally set to a year.

For practice decay, we could introduce another factor. But instead, slightly more intuitively, we use the concept of *equivalent time*. We say that practicing one more exercise is equivalent to a time t_e (set as two months) of not practicing. We call t_e the *equivalent inactive time*. This gives a decay ratio of

$$r = (1/2)^{(t+t_e)/t_{1/2}}. \quad (46)$$

However, we still go one step further. The first time a student does an exercise, the learning effect is still strong. It's the student's first time, so he/she is expected to learn a lot from it. However, after ten exercises or so, this effect is a lot smaller, and exercise fifty is unlikely to still have much of a learning effect at all. To incorporate that, we reduce the equivalent inactive time t_e based on the amount of practice a student has had. If the student has already practiced the skill n times, then we define

$$t_e = t_{e,0} (1/2)^{n/n_{1/2}}. \quad (47)$$

Here $t_{e,0}$ is the *initial equivalent inactive time*, set to two months, and $n_{1/2}$ is the number of times a student must practice a skill to get half the learning effect. Generally we set $n_{1/2} = 8$, although ideally this is larger for smaller (more simple) skills and smaller for larger (more complex) skills.

Using the above equations, we can always calculate an appropriate decay ratio r . It only needs to be applied.

B.3. APPLYING A DECAY RATIO

To apply a given decay ratio r , we could simply choose a smoothing order. From (44) you would expect it to equal

$$n_s = \frac{2r}{1-r}. \quad (48)$$

However, the smoothing order n_s must be an integer, which complicates matters. A rough solution would be to round n_s to the nearest integer value and apply that, but this results in discontinuities. Perhaps, from one day to the next, a student's score suddenly jumps. That is not desirable and may result in confusion/frustration.

A cleaner and more continuous solution would be to write $r = r_1 \cdot r_2 \cdot r_3 \cdot \dots$, for which each individual decay subratio r_i does result in an integer smoothing order $n_{s,i} = \frac{2r_i}{1-r_i}$. We can then apply smoothing multiple times. Choosing the subratios r_1, r_2, \dots is done by applying the following steps for $i = 1, 2, \dots$

- Choose the smoothing order $n_{s,i} = \lceil \frac{2r}{1-r} \rceil$.
- Pick $r_i = n_{s,i} / (n_{s,i} + 2)$ accordingly.
- Update the remaining decay ratio r using $r \leftarrow r / r_i$.
- If $n_{s,i} > n_{s,max}$ then stop: ignore r_i and further.

In practice, for numerical/practical reasons, usually $n_{s,max}$ is set around 100 or 120 or so. This also ensures that we often get at most three subratios r_1, r_2, r_3 , or in very rare cases four.

It must be noted here that the subratios r_1, r_2, \dots are automatically sorted largest to smallest, meaning their corresponding smoothing orders $n_{s,1}, n_{s,2}, \dots$ are sorted smallest to largest. Mathematically the order in which these smoothing orders are applied does not matter, but in practice it is wisest to apply them largest to smallest, ending with $n_{s,1}$. This ensures that the final coefficient vector \mathbf{c} we wind up with has order $n_{s,1}$. This reduces the number of coefficients we need to work with later on, slightly speeding up further computations.

REFERENCES

- BAKER, R. S. J. D., CORBETT, A. T., AND ALEVEN, V. 2008. More accurate student modeling through contextual estimation of slip and guess probabilities in bayesian knowledge tracing. In *Intelligent Tutoring Systems*, B. P. Woolf, E. Aïmeur, R. Nkambou, and S. Lajoie, Eds. Springer Berlin Heidelberg, Berlin, Heidelberg, 406–415.
- CEN, H., KOEDINGER, K., AND JUNKER, B. 2006. Learning factors analysis – a general method for cognitive model evaluation and improvement. In *Intelligent Tutoring Systems*, M. Ikeda, K. D. Ashley, and T.-W. Chan, Eds. Springer Berlin Heidelberg, Berlin, Heidelberg, 164–175.
- CORBETT, A. T. AND ANDERSON, J. R. 1994. Knowledge tracing: Modeling the acquisition of procedural knowledge. *User Modeling and User-Adapted Interaction* 4, 253–278.
- CUI, Y., CHU, M.-W., AND CHEN, F. 2019. Analyzing student process data in game-based assessments with bayesian knowledge tracing and dynamic bayesian networks. *Journal of Educational Data Mining* 11, 1 (06), 80–100.
- GUO, X., HUANG, Z., GAO, J., SHANG, M., SHU, M., AND SUN, J. 2021. *Enhancing Knowledge Tracing via Adversarial Training*. Association for Computing Machinery, New York, NY, USA, 367–375.
- JOHNSON, N. L., KOTZ, S., AND BALAKRISHNAN, N. 1995. *Continuous Univariate Distributions*, 2 ed. Vol. 2. Wiley.
- KAPLAN, A. M. AND HAENLEIN, M. 2016. Higher education and the digital revolution: About MOOCs, SPOCs, social media, and the Cookie Monster. *Business Horizons* 59, 4, 441–450.
- KIM, B. AND DOSHI-VELEZ, F. 2021. Machine learning techniques for accountability. *AI Magazine* 42, 1 (Apr), 47–52.
- KOEDINGER, K., CORBETT, A., AND PERFETTI, C. 2012. The knowledge-learning-instruction (KLI) framework: Toward bridging the science-practice chasm to enhance robust student learning. *Cognitive science* 36, 757–98.
- KÄSER, T., KLINGLER, S., SCHWING, A. G., AND GROSS, M. 2017. Dynamic bayesian networks for student modeling. *IEEE Transactions on Learning Technologies* 10, 4, 450–462.
- LEVY, R. 2019. Dynamic bayesian network modeling of game-based diagnostic assessments. *Multivariate Behavioral Research* 54, 6, 771–794.
- LEVY, R. AND MISLEVY, R. J. 2016. *Bayesian Psychometric Modeling*. Chapman & Hall.
- LIU, Q., SHEN, S., HUANG, Z., CHEN, E., AND ZHENG, Y. 2021. A survey of knowledge tracing.
- NEWELL, A. AND ROSENBLOOM, P. 1993. Mechanisms of skill acquisition and the law of practice. *Cognitive Skills and Their Acquisition Vol. 1*.
- PARDOS, Z. A. AND HEFFERNAN, N. T. 2011. Kt-idem: Introducing item difficulty to the knowledge tracing model. In *User Modeling, Adaption and Personalization - 19th International Conference, UMAP 2011, Girona, Spain, July 11-15, 2011. Proceedings*, J. A. Konstan, R. Conejo, J. L. Marzo, and N. Oliver, Eds. Lecture Notes in Computer Science, vol. 6787. Springer, 243–254.
- PAVLIK JR, P., CEN, H., AND KOEDINGER, K. 2009. Performance factors analysis - a new alternative to knowledge tracing. In *Frontiers in Artificial Intelligence and Applications. Frontiers in Artificial Intelligence and Applications* 200, 531–538.
- PIECH, C., BASSEN, J., HUANG, J., GANGULI, S., SAHAMI, M., GUIBAS, L. J., AND SOHL-DICKSTEIN, J. 2015. Deep knowledge tracing. In *Advances in Neural Information Processing Systems*, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, Eds. Vol. 28. Curran Associates, Inc.

- REICHENBERG, R. E. 2018. Dynamic bayesian networks in educational measurement: Reviewing and advancing the state of the field. *Applied Measurement in Education* 31, 4 (10), 335–350.
- SAPOUNTZI, A., BHULAI, S., CORNELISZ, I., AND KLAVEREN, C. 2019. Dynamic knowledge tracing models for large-scale adaptive learning environments. *International Journal on Advances in Intelligent Systems* 12, 93 – 110.
- SHEIDAEI, A., FOROUSHANI, A. R., GOHARI, K., AND ZERAATI, H. 2022. A novel dynamic bayesian network approach for data mining and survival data analysis. *BMC Medical Informatics and Decision Making* 22, 1 (09).
- VIE, J.-J. AND KASHIMA, H. 2019. Knowledge tracing machines: Factorization machines for knowledge tracing. *Proceedings of the AAAI Conference on Artificial Intelligence* 33, 750–757.
- WU, Z., LI, M., TANG, Y., AND LIANG, Q. 2020. Exercise recommendation based on knowledge concept prediction. *Knowledge-Based Systems* 210, 106481.
- YANG, S., ZHU, M., AND LU, X. 2021. Deep knowledge tracing with learning curves.
- YUDELSON, M. 2021. *Individualization of Bayesian Knowledge Tracing Through Elo-infusion*. Springer, 412–416.
- YUDELSON, M. V., KOEDINGER, K. R., AND GORDON, G. J. 2013. Individualized bayesian knowledge tracing models. In *Artificial Intelligence in Education*, H. C. Lane, K. Yacef, J. Mostow, and P. Pavlik, Eds. Springer Berlin Heidelberg, Berlin, Heidelberg, 171–180.