Mastery Learning Heuristics and Their Hidden Models

Shayan Doroudi^(✉)

University of California, Irvine, Irvine, CA 92697, USA
doroudis@uci.edu

Abstract. Mastery learning algorithms are used in many adaptive learning technologies to assess when a student has learned a particular concept or skill. To assess mastery, some technologies utilize data-driven models while others use simple heuristics. Prior work has suggested that heuristics may often perform comparably to model-based algorithms. But is there any reason we should expect these heuristics to be reasonable? In this paper, we show that two prominent mastery learning heuristics can be reinterpreted as model-based algorithms. In particular, we show that the N -Consecutive Correct in a Row heuristic and a simplified version of ALEKS' mastery learning heuristic are both optimal policies for variants of the Bayesian knowledge tracing model. By putting mastery learning heuristics on the same playing field as model-based algorithms, we can gain insights on their hidden assumptions about learning and why they might perform well in practice.

Keywords: Mastery learning · Heuristics · Learner models · N -CCR · Bayesian knowledge tracing · Change-point detection · CUSUM

1 Introduction

Mastery learning is an instructional technique popularized by Benjamin Bloom [7], but which at least dates back to progressive educational movements in the early twentieth century [17, 24]. The idea is to give students just the right amount of instruction or practice that they need in order to mastery a particular topic before moving them on to the next topic. Today, mastery learning underlies many adaptive learning technologies, including Khan Academy, Duolingo, ASSISTments, ALEKS, and cognitive tutors like MATHia. Each of these platforms are being used by thousands to millions of students yearly, and as such, the way they assess mastery can have real consequences for students. A lot of work has been invested in developing statistical techniques to infer models of student learning that make predictions about whether a student has learned a skill, which could in turn be used in mastery learning [5, 9, 12, 20, 25]. However, in practice, many state-of-the-art adaptive learning systems assess mastery in simple ways, which are seemingly not very “intelligent” [4]. For example, some platforms use simple heuristics to assess whether students have reached mastery, such as having students receive practice on a skill until they answer questions correctly three times

in a row. Even platforms that use models of student learning often have the model parameters manually set by system designers or domain experts [10, 22], rather than leveraging data-driven techniques.

One advantage of using simple heuristics is that they are interpretable and easy-to-convey to teachers, students, and other stakeholders. Moreover, they may seem intuitively reasonable. However, it is not clear whether these intuitions *in principle* align with (a) our understanding of how students learn or (b) inferences we can make about student mastery from the data.

In this paper, we present a means for better understanding mastery learning heuristics, by re-interpreting them as model-based algorithms. In particular, we show that the *N-Consecutive Correct in a Row (N-CCR)* heuristic used by ASSISTments and a simplified version of the mastery learning heuristic used by ALEKS are both optimal mastery learning policies for variants of the Bayesian knowledge tracing (BKT) model. By placing mastery learning heuristics in the same playing field as model-based mastery learning algorithms, we hope to better understand the theoretical assumptions about learning that mastery learning heuristics are making, and as such, help guide designers of adaptive learning systems to make more intentional decisions about what heuristics to use.

2 Background

In what follows, a *mastery learning policy* is any instructional policy that considers topics–skills, concepts, or knowledge components (KCs)—one at a time, and decides how many practice opportunities to give for the current topic before moving on to the next. In all of the mastery learning policies we consider, the decision will be made purely based on whether previous answers were correct or incorrect on the students’ first attempt on each question for the same KC. An *optimal* mastery learning policy under a model is one that gives the optimal amount of practice subject to some accuracy threshold (e.g., 95% confidence that the student has mastered the skill).

A popular approach to mastery learning that underlies cognitive tutors, such as MATHia, is to use the Bayesian knowledge tracing (BKT) model. The standard BKT model for a single KC is a two-state hidden Markov model that assumes that after receiving a practice opportunity t , the student is in one of two knowledge states: the learned state, where they know the KC ($K_t = 1$), or the unlearned state, where they do not know the KC ($K_t = 0$) [12]. When the student begins using the adaptive learning system, BKT assumes they start in the unlearned state with probability $P(L_0) = P(K_1 = 1)$. If a student is in the unlearned state, every time they attempt a practice opportunity and receive feedback, they have some fixed probability of learning the KC, $P(T) = P(K_{t+1} = 1 | K_t = 0)$. When the student is in the learned state, they are assumed to stay there forever (i.e., no forgetting). Every time the student is given a practice opportunity, we can see whether they answered the question correctly ($C_t = 1$) or incorrectly ($C_t = 0$). If the student is in the unlearned state, they will answer correctly with some probability of guessing $P(G)$ and

otherwise answer incorrectly. If the student is in the learned state, they will answer correctly unless they slip with probability $P(S)$. The BKT model for a single skill is thus fully described by four parameters: $P(L_0)$, $P(T)$, $P(G)$, and $P(S)$. When using the BKT model, one can continuously update the probability that the student has learned the KC so far ($P(K_t|C_1, C_2, \dots, C_{t-1})$). The optimal mastery learning policy for the BKT model continues to give practice opportunities to the student until this probability exceeds some threshold, typically 0.95 [11, 12].

3 N -CCR as a Model-Based Algorithm

The N -Consecutive Correct in a Row (N -CCR) heuristic keeps giving students practice problems on a given topic or skill until the student answers the problems correctly N times correctly in a row. This heuristic is used in ASSISTments' Skill Builders exercises [6, 15] and was previously used by Khan Academy [13]. Recently, Khan Academy has switched to a more gamified way of implementing mastery learning, where students can go through a series of Mastery Levels, but to reach the Proficient Level students still have to get a certain number of problems correct in a row [16].

We now show that N -CCR can be viewed as the optimal mastery learning policy for certain BKT models. Note that when using the N -CCR heuristic, if a student gives any number of consecutive correct answers less than N followed by an incorrect answer, then they are back in the same "state" as though they had not given any correct answers. Now suppose that the true model of learning is a BKT model. Then, this must mean once we see $C_{t-1} = 0$, the student is identified as having been in state $K_{t-1} = 0$, or $P(K_{t-1} = 0|C_{t-1} = 0) = 1$. Using Bayes' rule, it can be shown that this implies that the probability of slipping $P(S)$ must be zero. By setting the values of $P(G)$, $P(T)$, and $P(L_0)$ appropriately, it can be easily seen that N -CCR is the optimal mastery learning policy for some BKT model where $P(S) = 0$ and for a given accuracy (e.g., 95%). We demonstrate this precisely in the online supplementary material¹. It is worth noting that the BKT model with $P(S) = 1$ was actually a well-studied model in the 1960s mathematical psychology community, known as the one-element model [8, 14].

This re-formulation of N -CCR can help explain why it may seem to perform well in practice. For example, Pelánek and Řihák [19] showed in simulation that, even if students learn according to BKT models, the N -CCR heuristic (with the optimal value of N) often performs almost as good as the optimal BKT mastery learning policy. This may seem surprising, but our findings indicate that if $P(S)$ is small, then the best N -CCR heuristic will correspond to the mastery learning policy for a BKT model that might be close to optimal.

4 TOW as a Model-Based Algorithm

The Tug-of-War (TOW) Heuristic is what we are calling a mastery learning heuristic that gives points to students for answering questions correctly, removes

¹ <https://bit.ly/aied2020-heuristics>.

points for answering questions incorrectly (while keeping the minimum number of points at zero), and keeps giving practice until the student achieves a certain number of points. We will use $TOW_{+i,-j,N}$ to designate the specific TOW heuristic where the student gets i points for a correct answer, loses j points for an incorrect answer and needs N points for mastery. ALEKS, a prominent adaptive learning system, implements mastery learning using the $TOW_{+1,-1,5}$ heuristic (or $TOW_{+1,-1,3}$ in some cases) with a few differences² [2]. Previously, ALEKS used the $TOW_{+2,-1,5}$ and $TOW_{+2,-1,3}$ heuristics [23].

To see how some TOW heuristics can be interpreted as model-based algorithms, consider a variant of the BKT model where we make no assumptions as to how or when the student learns a KC. That is, $P(T)$ need not be a fixed probability (e.g., it can increase over time), it need not be the same for all students, and it need not even be probabilistic. Since we do not make assumptions about how students are learning, to implement mastery learning here, we simply want to detect when there appears to have been a sudden increase in the probability of answering correctly (from $P(G)$ to $1 - P(S)$). This is known as change-point detection, which is a well studied problem in statistics and fields like quality control [1, 3, 18]. Specifically, we can use the Bernoulli CUSUM chart algorithm [21]. We describe this method and how it can be applied to mastery learning in the online supplementary material³. Given that we do not make assumptions about the probability of learning, we cannot make any statement about how confident we are that the student has learned the skill. Instead, we can set the parameters such that 95% of students who have not learned the skill would have had at least a certain number of practice opportunities (which we can choose) before we would mistakenly declare mastery. It can be shown that a variety of TOW heuristics, including all of the ones mentioned above can be implemented using the CUSUM algorithm with an appropriate choice of parameters.

5 Conclusion

Given that these heuristics may be more principled than they might appear at first sight, perhaps their use in adaptive learning systems is warranted, especially given that they are much easier to communicate to students than complex model-based policies. However, other considerations need to also be taken into account. For example, the N -CCR heuristic might be demotivating for students given that a single “slip” punishes the student. On the other hand, if we believe students can slip, then perhaps the N -CCR heuristic should not be used altogether! All in all, a more comprehensive understanding of mastery learning heuristics and their hidden models can hopefully help us ensure that adaptive learning systems perform mastery learning in productive ways.

² In the ALEKS Learning Sequence, if the student gets two consecutive answers correct in a row, the student gets an extra point. Also, if the student revises an incorrect answer on their second attempt, they subsequently gain a point, netting zero points.

³ <https://bit.ly/aied2020-heuristics>.

References

1. Adams, R.P., MacKay, D.J.: Bayesian online changepoint detection. arXiv preprint [arXiv:0710.3742](https://arxiv.org/abs/0710.3742) (2007)
2. ALEKS: New ALEKS student module: reference guide (2016). https://www.aleks.com/resources/New_Student_Module_Ref_Guide.pdf
3. Aroian, L.A., Levene, H.: The effectiveness of quality control charts. *J. Am. Stat. Assoc.* **45**(252), 520–529 (1950)
4. Baker, R.S.: Stupid tutoring systems, intelligent humans. *Int. J. Artif. Intell. Educ.* **26**(2), 600–614 (2016)
5. Baker, R.S.J.D., et al.: Contextual Slip and prediction of student performance after use of an intelligent tutor. In: De Bra, P., Kobsa, A., Chin, D. (eds.) UMAP 2010. LNCS, vol. 6075, pp. 52–63. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-13470-8_7
6. Beck, J.E., Gong, Y.: Wheel-spinning: students who fail to master a skill. In: Lane, H.C., Yacef, K., Mostow, J., Pavlik, P. (eds.) AIED 2013. LNCS (LNAI), vol. 7926, pp. 431–440. Springer, Heidelberg (2013). https://doi.org/10.1007/978-3-642-39112-5_44
7. Bloom, B.S.: Learning for mastery. *Eval. Comment* **1**(2), 1–5 (1968)
8. Bower, G.H.: Application of a model to paired-associate learning. *Psychometrika* **26**(3), 255–280 (1961)
9. Cen, H.: Generalized learning factors analysis: improving cognitive models with machine learning. Ph.D. thesis (2009)
10. Cen, H., Koedinger, K.R., Junker, B.: Is over practice necessary?-Improving learning efficiency with the cognitive tutor through educational data mining. *Front. Artif. Intell. Appl.* **158**, 511 (2007)
11. Corbett, A.: Cognitive mastery learning in the act programming tutor. In: AAAI Tech. rep. SS-00-01 (2000)
12. Corbett, A.T., Anderson, J.R.: Knowledge tracing: modeling the acquisition of procedural knowledge. *User Model. User-Adapt. Interac.* **4**(4), 253–278 (1995)
13. Hu, D.: How Khan Academy is using machine learning to assess student mastery (2011). <http://david-hu.com/2011/11/02/how-khan-academy-is-using-machine-learning-to-assess-student-mastery.html>
14. Karush, W., Dear, R.: Optimal strategy for item presentation in a learning process. *Manag. Sci.* **13**(11), 773–785 (1967)
15. Kelly, K.M., Wang, Y., Thompson, T., Heffernan, N.T.: Defining mastery: knowledge tracing versus n-consecutive correct responses. In: Proceedings of the 8th International Conference on Educational Data Mining (2015)
16. Khan Academy: what are course and unit mastery? (2020). <https://khanacademy.zendesk.com/hc/en-us/articles/115002552631-What-are-Course-and-Unit-Mastery->
17. Morrison, H.C.: The Practice of Teaching in the Secondary School. University of Chicago Press, Chicago (1926)
18. Page, E.S.: Continuous inspection schemes. *Biometrika* **41**(1/2), 100–115 (1954)
19. Pelánek, R., Řihák, J.: Experimental analysis of mastery learning criteria. In: Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization, pp. 156–163 (2017)
20. Piech, C., et al.: Deep knowledge tracing. In: Advances in Neural Information Processing Systems, pp. 505–513 (2015)

21. Reynolds Jr., M.R., Stoumbos, Z.G.: A CUSUM chart for monitoring a proportion when inspecting continuously. *J. Qual. Technol.* **31**(1), 87–108 (1999)
22. Salden, R.J., Aleven, V., Schwonke, R., Renkl, A.: The expertise reversal effect and worked examples in tutored problem solving. *Instr. Sci.* **38**(3), 289–307 (2010)
23. Uzun, H.: Personal communication (2020)
24. Washburne, C.W.: The individual system in Winnetka. *Elementary Sch. J.* **21**(1), 52–68 (1920)
25. Zhang, J., Shi, X., King, I., Yeung, D.Y.: Dynamic key-value memory networks for knowledge tracing. In: *Proceedings of the 26th International Conference on World Wide Web*, pp. 765–774 (2017)