Delft University of Technology

Curriculum pacing

A new approach to discover instructional practices in classrooms

Patel, Nirmal; Sharma, Aditya; Sellman, Collin; Lomas, Derek

DOI

[10.1007/978-3-319-91464-0_38](https://doi.org/10.1007/978-3-319-91464-0_38)

Publication date

2018

Document Version

Submitted manuscript

Published in

Intelligent Tutoring Systems - 14th International Conference, ITS 2018, Proceedings

Citation (APA)

Patel, N., Sharma, A., Sellman, C., & Lomas, D. (2018). Curriculum pacing: A new approach to discover instructional practices in classrooms. In J. Vassileva, R. Nkambou, & R. Azevedo (Eds.), *Intelligent Tutoring Systems - 14th International Conference, ITS 2018, Proceedings* (pp. 345-351). (Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics); Vol. 10858 LNCS). Springer Science+Business Media. https://doi.org/10.1007/978-3-319-91464-0_38

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

Curriculum Pacing: A New Approach to Discover Instructional Practices in Classrooms

Nirmal Patel¹, Aditya Sharma¹, Collin Sellman², and Derek Lomas³

¹ Playpower Labs, India

{nirmal,aditya.sharma}@playpowerlabs.com

² Arizona State University, USA

collin.sellman@asu.edu

³ Delft University of Technology, Netherlands

j.d.lomas@tudelft.nl

Abstract. This paper examines the use of “pacing plots” to represent variations in student learning sequences within a digital curriculum. Pacing plots are an intuitive and flexible data visualizations that have a potential for revealing the diversity of blended classroom instructional models. By using curriculum pacing plots, we identified several common implementation patterns in real-world classrooms. After analyzing two years’ worth of data from over 150,000 students in a digital math curriculum, we found that a PCA and K-Means clustering approach was able to discover pedagogically relevant instructional practices.

Keywords: curriculum analytics, curriculum pacing, sequence mining, clustering, visualization

1 Introduction

New data instrumentation methods have enabled digital learning products to capture large amounts of fine-grained data describing student behavior (e.g., clickstream data.) However, it is still a challenge to transform big educational data into real-world benefits for students and teachers. It is common for large-scale analyses to rely on simple, aggregated metrics like average score, total events, or total usage time. While these metrics are essential, they do not make use of the temporal dynamics of student actions present in the data and actualize the potential for big data in digital learning [9]. Our study was motivated by the potential to determine correlations between temporal trajectories of learners and various outcome measures such as student and teacher engagement, implementation fidelity, learning gains, fluency etc.

Different methods have been used to analyze educational sequence data, including association rule mining, sequential pattern mining, and process mining [10]. As we expand our capacity to analyze sequence data, certain challenges arise. For instance, because of the diversity of usage behaviors, sequence mining techniques can produce overly complex and hard-to-interpret “spaghetti” models. Although these complex models can be used for predicting student actions

over time more precisely, they have little interpretability. In this case, rather than analyzing sequence data directly, we can use clustering methods to group similar sequences together and analyze them separately [3, 8].

Sequence data from learners provide new opportunities to make data-driven intelligent tutors. For example, clickstream data have been used to produce next-step recommendations [7], induce reinforcement learning teaching strategies [12], and build data-driven pedagogical models [2]. These studies show that it is possible for intelligent tutors to base their recommendations upon models of sequential (or temporal) data.

2 Curriculum Pacing

One temporal aspect of classroom activity is what educators refer to as *pacing*, which has been described as “the rate at which new instructional material is introduced to students” [1]. This definition of pacing is, however, somewhat limiting, as it is common for students to revisit learning materials to ensure mastery before moving on to more advanced concepts. Thus, we define pacing as “the progression through curriculum over time.” This definition attempts to include all aspects of instruction over time in the construct of pacing. Although we have analyzed pacing in digital classrooms, we note that a great deal of classroom activity cannot be represented by the clickstream data.

It is straightforward to build a representation of pacing by using timestamped logs of student use of different curricular activities within a digital learning system. In addition to timestamps, the only other requirement is that learning activities should have attached metadata that indicate where the activity falls within a linear curriculum (e.g., unit 1, unit 2, unit 3, etc.) The notion of a linear curriculum remains inherent and central to the construct of pacing, which tells us how a student goes through the curriculum over time. However, even in cases where the curriculum is not entirely linear, pacing plots can still be used. For instance, each activity can be tagged with the average time when it is used in the curriculum. This approach can capture the linearity of activities that are used in sequence.

Our visual model of curriculum pacing aims to present student activity in the curriculum over time. This model has two dimensions: X dimension representing time (e.g., number of weeks) and Y dimension representing distance through the curriculum (e.g., unit 1, unit 2, etc.) Figure 1 presents three examples of pacing plots. Pacing plots examples shown in Figure 1 are just an instance of a broader visual design space. We explored various values for different design factors of the plot, including **Data** (single student, all students in a class,) **X-axis** (# of weeks, calendar data,) **Y-axis** (lesson number, average week used,) **Plot type** (scatter plot, heatmap,) and **Fills** (usage, score, percentile.) These variations were used to help identify different classroom implementation models across a school district. The addition of score or percentile information can help to indicate where students or entire classrooms have struggled in a curriculum.

Fig. 1: Examples of pacing plots. Each plot corresponds to a unique trace that one or more students left in the digital learning system. X-axis represents week of usage, and Y-axis represents the digital curriculum chapter numbers. Cells of the plot indicate whether student accessed the curriculum chapter in the given week or not.

We also found it useful to combine or average multiple plots to produce an aggregated pacing plot; e.g., to represent an entire district’s “typical” pacing.

3 Meaningful Characteristic Variations

Classrooms have a great deal of diversity, and this variation is bound to appear in the pacing plots. During our initial review of pacing plots, we identified several characteristic variations in the form of the plots. These variations appeared to represent different classroom conditions and implementation approaches.

The most notable of all the patterns that appeared in these plots were vertical lines indicating that students are revisiting previous topics (e.g., “icicles”) and horizontal lines indicating that students are practicing the same material over and over again (e.g., “ruts”.) We identified several patterns that were likely to appear in classrooms: **Lockstep pacing**, a pattern where all students used the same material, represented as a tight pacing line; **Flexible pacing**, a pattern where there was variation in the material used, represented as a fuzzier pacing line; **Icicles**, a graph feature that appeared to occur when classes would engage in review; **Cram-to-complete**, the common tendency for an accelerated pace at the end of the school year; and **Glaciers** representing students entering at a later time during the year and catching up.

These variations were identified in a relatively small subset of data (< 100 classrooms.) To explore these variations at a larger scale, we surmised that a clustering method might be helpful for automating the identification and quantification of these pacing patterns. Effective clustering should be capable of revealing the patterns already identified and also, potentially, capable of revealing new patterns. We had two hypotheses for our study:

- **H1:** Clustering will identify previously known curriculum pacing variations that are meaningful to experts.
- **H2:** Clustering will identify previously unknown curriculum pacing variations that are meaningful to experts.

4 Participants and Method

For our analysis, we used anonymized data from a large-scale online math curriculum that has been used by more than 150,000 students across the United States. The curriculum is divided into topics and subtopics; within each subtopic, there are a variety of activities such as videos, scaffolded practice quizzes, formative assessments, and homework assignments. The program is designed to go from the start to end in a linear fashion. Teachers assign resources to students, and students turn them in after completion. Using clickstream data from this program, we extracted pacing plots for individual students. Plots that were the same were merged. This deduplication reduced the size of the dataset by approximately 20% ($N = 121,502$). This produced a dataset with unique instructional patterns as data points, instead of students. This meant that we clustered different usage patterns, giving all the patterns same weight regardless of the difference in their frequencies.

To cluster pacing plots, we used K-Means clustering. Clustering high dimensional data can lead to the curse of dimensionality [5], where a large number of clusters is needed to discover meaningful patterns. Indeed, each of our pacing plots had 5824 features (52 weeks x 112 digital book chapters,) so we reduced the dimensionality of these data points using PCA as a preprocessing step. We used 212 principal components that explained 50% of the variance in data.

5 Results

We ran K-Means clustering on 121,502 unique instructional sequences using their lower dimensional representations. We chose a total of 50 clusters ($K = 50$), which produced a model capturing 26.5% of the variation between the clusters. As the number of clusters was large, many clusters exhibited similar patterns, but this allowed us to find both unexpected and nuanced patterns. A large number of patterns did not have any characteristic variations in them. We found smaller clusters that identified both known and unknown characteristic variations in student learning sequences. We also found that many small clusters had little difference between them. Only a fraction of clusters are shown here.

In Figure 2a, we see a group of instructional patterns where students accessed the same material with minimal variation (**Lockstep pacing**.) Figure 2b shows a “fuzzier” pacing line, indicating that students who followed these patterns did different things in the same week (**Flexible pacing**.) Figure 2c captures patterns where, at the end of the year, previous chapters of the textbook were reviewed (**Icicles**,) and Figure 2d shows patterns where materials were covered quickly at the end of the year (**Cram-to-complete**.) We did not find any clusters

Fig. 2: Visualizations of clusters from resulting analysis. Each plot is a combination of all data points in the cluster. Many of the patterns in the dataset had minimal usage, and meaningful clusters were smaller in size.

representing Glaciers, potentially due to our use of relative time in the pacing plots. Together, these clusters provide evidence that partially supports **H1**, that clustering will identify known curriculum pacing variations.

Figure 2e shows a group of unexpected instructional patterns: most of the curriculum was covered but in an unusual sequence. Some later topics were covered first, and some earlier topics were visited later. Figure 2f shows a group of patterns where the material was covered in bursts. These findings provide some evidence in the support of **H2**, that clustering will identify previously unknown pacing variations.

6 Discussion and Conclusion

Our approach shows how digital curriculum pacing plots might provide insight into variations in classroom instruction. Although we found evidence in support of our hypotheses, this preliminary work is limited. While clustering helped identify novel behavioral patterns, we have not evaluated their meaningfulness with, for instance, teachers who participated in those classrooms. We also suspect that the clustering techniques were unable to identify certain variations that

were common in our classroom-level analyses. For instance, in our initial observations, we characterized different classrooms as “lockstep” or “personalized,” based on the presence of “icicles”; these icicles appear to be cases where teachers were helping students by assigning them prerequisite skills. Why were these not observed in our clustering attempts? This may be because personalization behaviors were observed at a classroom level, whereas our analysis was focused on individual students. Alternatively, the dimensionality reduction approach that we took for clustering might be washing out finer details of the plots. This can be investigated in future analyses.

We expect that variations in pacing will predict variations in student outcomes [4, 11]. In our future work, we will investigate the addition of student performance data in the pacing plots, which can illustrate nuances in instructional success at every step of the curriculum. These future approaches are intended to help understand and support classroom instructional practices at scale. One future possibility is that these models could support teacher-facing adaptive recommendation systems [6] that could help teachers learn from the aggregated decision-making of thousands of other teachers.

References

1. Barr, R., Dreeben, R.: *How Schools Work*, p. 33 (1991)
2. Chi, M., VanLehn, K., Litman, D.: Do micro-level tutorial decisions matter: Applying reinforcement learning to induce pedagogical tutorial tactics. In: *Proceedings of the 10th International Conference on ITS*. pp. 224–234. Springer (2010)
3. Hansen, C., Hansen, C., Hjuler, N., Alstrup, S., Lioma, C.: Sequence modelling for analysing student interaction with educational systems. In: *Proceedings of the 10th International Conference on Educational Data Mining* (2017)
4. Hoadley, U.: Time to learn: Pacing and the external framing of teachers’ work. *Journal of education for teaching: international research and pedagogy* 29(3), 265–277 (2003)
5. Keogh, E., Mueen, A.: Curse of dimensionality. In: *Encyclopedia of machine learning*, pp. 257–258. Springer (2011)
6. Koedinger, K.R., Brunskill, E., Baker, R.S., McLaughlin, E.A., Stamper, J.: New potentials for data-driven intelligent tutoring system development and optimization. *AI Magazine* 34(3), 27–41 (2013)
7. Pardos, Z.A., Tang, S., Davis, D., Le, C.V.: Enabling real-time adaptivity in moocs with a personalized next-step recommendation framework. In: *Proceedings of the 4th ACM Conference on Learning@Scale*. pp. 23–32. ACM (2017)
8. Patel, N., Sellman, C., Lomas, D.: Mining frequent learning pathways from a large educational dataset. In: *Proceedings of the 3rd International Workshop on Graph Educational Data Mining*. pp. 27–30 (2017)
9. Reich, J.: Rebooting mooc research. *Science* 347(6217), 34–35 (2015)
10. Romero, C., Ventura, S., Pechenizkiy, M., Baker, R.S.: *Handbook of educational data mining*. CRC press (2010)
11. Smith, J., Smith, B., Bryk, A.: *Setting the pace: Opportunities to learn in chicago public elementary schools*. Tech. rep. (1998)
12. Zhou, G., Wang, J., Lynch, C.F., Chi, M.: Towards closing the loop: Bridging machine-induced pedagogical policies to learning theories. In: *Proceedings of the 10th International Conference on Educational Data Mining*. pp. 112–119 (2017)