

Analysis of Student Behaviour in *Habitable Worlds* Using Continuous Representation Visualization

Zachary A. Pardos¹, Lev Horodyskyj²

Abstract

We introduce a novel approach to visualizing temporal clickstream behaviour in the context of a degree-satisfying online course, *Habitable Worlds*, offered through Arizona State University. The current practice for visualizing behaviour within a digital learning environment is to generate plots based on hand-engineered or coded features using domain knowledge. While this approach has been effective in relating behaviour to known phenomena, features crafted from domain knowledge are not likely well suited to making unfamiliar phenomena salient and thus can preclude discovery. We introduce a methodology for organically surfacing behavioural regularities from clickstream data, conducting an expert in-the-loop hyperparameter search, and identifying anticipated as well as newly discovered patterns of behaviour. While these visualization techniques have been used before in the broader machine-learning community to better understand neural networks and relationships between word vectors, we apply them to online behavioural learner data and go a step further, exploring the impact of the parameters of the model on producing tangible, non-trivial observations of behaviour that suggest pedagogical improvement to the course designers and instructors. The methodology introduced in this paper led to an improved understanding of passing and non-passing student behaviour in the course and is applicable to other datasets of clickstream activity where investigators and stakeholders wish to organically surface principal patterns of behaviour.

Notes for Practice

- Continuous-representation visualization can produce a high-level view of emergent student behaviour online without the need to define features or tag.
- Differential visualization of passing and non-passing student course behaviours can help identify deep- and shallow-learning strategies and provide instructors with essential information for modifying the curricula to discourage strategies associated with failure.
- Involving instructors in the tuning of the visualization and model parameters produces analyses with a desirable mixture of expected and unexpected, but explainable, patterns.
- Layering on additional data, such as when students create a discussion post, further contextualizes insight into student learning strategies from visualizations.

Keywords

Behaviour visualization, representations learning, feature engineering, dimensionality reduction, clickstream, online learning, skip-gram, t-SNE, *Habitable Worlds*, ASU, higher education.

Submitted: 16.09.2017 — **Accepted:** 03.04.2019 — **Published:** 15.04.2019

Corresponding author ¹ Email: zp@berkeley.edu Address: Graduate School of Education and School of Information, University of California, Berkeley, 2121 Berkeley Way, Suite 4232, Berkeley, CA 94720, USA ORCID: 0000-0002-6016-7051

² Address: Center for Education Through eXploration, School of Earth and Space Exploration, Arizona State University, 781 Terrace Mall, Tempe, AZ 85287, USA ORCID: 0000-0003-4606-4379

1. Introduction

A highly touted benefit of a completely online or otherwise digital course has been that a detailed record of the interactions of learners with course materials is kept that can then be mined for potential actionable insights. It is therefore no surprise that descriptive statistics of the interactions between learners and pedagogical materials were among the first type of information sought by online course instructors. Among their questions of curiosity were: How did students utilize the course resources? How did struggling students differ in their usage patterns compared to passing students? Where did students need help and what can the data provide that can help improve future offerings? While prior work has made headway, room for

methodological improvement remains. The current visualization practice of depicting graphs of common student transitions from one resource to another can produce unsatisfying answers to these questions because it is ultimately a descriptive statistic summarization of behaviour that can struggle to surface less-anticipated patterns that manifest over a longer period of time. Predictive models can learn complex behavioural patterns, but while they can at times convey enough corroborating evidence to convince researchers and educators that their predictions are robust, and even useful, communicating their learned patterns to provide enough novel and interpretable information to increase an instructor's existing understanding of their own course is a challenge. Linear representation learning approaches, detailed in the methodology section, strike a happy medium between straightforward descriptive approaches and the opacity of more complex machine learning models. As opposed to support vector machines and the cadre of neural network approaches labelled as "deep learning," representation learning algorithms (e.g., skip-grams) are in the class of simple linear feed-forward neural network models. While they are trained by optimizing a predictive outcome, it is the structure found in the model's learned parameters (i.e., the representation) that is of value. In this paper, we show that when used in combination with dimensionality reduction techniques, robust emergent patterns can be surfaced with a nuance not attainable with existing descriptive approaches.

2. Background

The practice of training a machine-learned classifier in predictive learning analytics has involved the customary process of feature engineering, a step of transforming data from its original form into a set of more abstract, handcrafted features using the domain knowledge of the researcher. This is often a process of aggregation, normalization, or combination of multiple attributes to extract essential descriptors from the original unstructured or semi-structured data (Freitag, 2000). The motivation for this process is often twofold: (i) to bring the prior knowledge of the researcher to bear on the problem and (ii) to transform the original data into a form that is syntactically compatible with the chosen classifier(s). This standard practice can be seen in learning analytics work on predicting student affect (Baker et al., 2012), drop-out (Boyer, Veeramachaneni, 2015), and question correctness (Stamper & Pardos, 2016), among many other applications of data mining in education (Koedinger et al., 2015). The premise of this approach is that by funnelling data through the prism of the researcher's domain knowledge, the engineered set of features will better represent the prediction task than the original untreated data. "Better," in this case, is defined as leading to better model fit or predictive generalization, the intuition being that there are certain hand-engineered transformations of the original attributes that closely relate to what is being predicted but would be difficult for a classifier to learn from the original data in the process of training. This approach has been quite effective at producing reasonably accurate classifiers. When working with smaller datasets, the feature-engineering process can be seen as a type of non-linear human-specified functional transformation between data and an intermediate representation which is closer to the target of prediction than the original data. In scenarios with smaller data, bringing the researcher's expert generalization to bear to create these representations often leads to more robust models than attempting to statistically learn them.

Visualization, too, can be seen as a manual process of feature engineering whereby the researcher brings their prior domain knowledge and hypotheses to bear on the transformation of the data from its original form. The difference between this process in the context of visualization and machine classification is that in visualization, this feature representation is being presented to a human learner, rather than a machine classifier, to understand and reason about. In the area of visualizing behaviour in an online course environment, which is the case study of this paper, common engineered features have been descriptive statistics of MOOC certification (Breslow et al., 2013), dwell time by resource category (Seaton et al., 2013), and counts of clickstream actions in quantized windows over time combined with summaries of forum activity (Crossley et al., 2016). Other work visualizes behaviour using conditional probabilities (e.g., Markov graphs), which express the most common transitions from one resource to another (Köck & Paramythis, 2011; Caprotti, 2017) or a sub-sequence of commonly occurring transitions called motifs (Davis et al., 2016). In the case of the descriptive approach, a bar chart or scatter plot is used, where the X-axis is typically a categorical (e.g., types of resources) or time (e.g., week 1 through 10) scale and the Y-axis is the attribute being summed (e.g., dwell time or number of actions). In Xu et al. (2014), for example, the Y-axis of their three scatter plots was the learners' course grade, with the X-axis being quiz clicks, lecture page views, and discussion page views. An overview of the past art on visualizing behaviours within massive open online courses (MOOCs) can be found in Emmons et al. (2017). The transition frequency approach can be seen as an attempt to disaggregate the descriptive approach and study the relationships between individual resources or resource types. In this approach, a graph is visualized, with nodes representing resources and edges representing transitions. The frequency of the transition can be expressed with the thickness or length of the edge. In the case of a report on the first four years of MITx MOOCs (Chuang & Ho, 2016), a prominent figure had each node representing a course and the size of the node depicting the number of enrolments.

In graphs, the angular orientation between two vertices with respect to a frame of reference is often irrelevant. Applications like Graphviz (Ellson et al., 2001) use variants of force-directed algorithms that optimize angles and orientations for visual appeal but do not necessarily encode any additional information. The only constraints are the set of edges and vertices and,

optionally, the length of those vertices. By visualizing machine-learned representations instead of hand-specified features, we hypothesize that patterns of greater novelty and significance can be revealed. Work has begun to bridge visualization with modelling, using regression with pre-hypothesized features and outcomes of the environment (Fratamico et al., 2017; Park et al., 2017), and using singular-value decomposition to study relationships between assessment constructs, or epistemic elements, based on their coded co-occurrences in discourse (Shaffer & Ruis, 2017). Bergner et al. (2014) discuss the remaining difficulties of visualizing variable-length learner clickstream data and the inadequacies of current methods. The representation approach introduced in this paper provides a method to address these difficulties.

The practice of feature engineering is often an informationally lossy one. For one, this is because much of feature engineering is based on aggregation and summarization, and because features created by hand are limited to the scope of the researcher's intuition and domain knowledge. When an ample degree of domain knowledge is present, feature engineering can be an effective way of explicitly defining relevant relationships in the data that a statistical learning approach might struggle to identify. For domain areas in which there is little theory or expertise, such as many behavioural contexts, feature engineering may not be a viable option. In these areas, features may be engineered which effectively describe known phenomena but are ineffective at describing unfamiliar phenomena, thus inhibiting downstream discovery. The general paradigm of using connectionist models (i.e., neural networks) to generate these features is called representation learning (Bengio et al., 2013). There has been no field where the lossiness of feature engineering has been made more apparent than in computer vision. The dominant domain expert intuition for the task of classifying what object is featured in an image was to engineer a set of features which describe the edges present in an image and to then present this set of edge descriptions to a classifier for training. The accuracy results of this approach were eclipsed by convolutional neural networks (Krizhevsky et al., 2012), which automatically learn rich feature representations of the image from the original pixel data. We hypothesize that by applying representation learning to student behavioural data in an online course, as it has been applied to problem interaction sequences in math tutoring systems (Pardos & Dadu, 2017; Pardos et al., 2018) and course enrolment data (Pardos et al., 2019), the most significant features of behaviour can be revealed where they may never have been detected if first filtered through the prism of one's domain knowledge. The visualization of the learned representations in our approach is a visualization of the patterns or regularities of student behaviour learned by a model which combs through course clickstream, a kind of data hardly scrutable to an instructor in untreated form.

3. Methodology

This paper introduces a novel technical methodology involving representation learning and visualization as well as a qualitative methodology in which the visualization is tuned and interpreted in close collaboration with the course instructor.

The technical methodology began with a dataset containing the chronological sequences of interactions of students with an online for-credit course offered by Arizona State University. A representation learning model common in computational text, called a skip-gram (Mikolov, Yih, & Zweig, 2013), was applied to this behavioural sequence. This is a novel application of the algorithm because it is customarily applied to sequences of words, not behaviours. We hypothesize that the analogy to language representation will hold. This process learns a high-dimensional vector representation of every course element that students interacted with. These vectors were reduced to two dimensions for visualization using a non-linear dimensionality reduction technique called t-stochastic neighbourhood embedding (t-SNE), designed to visualize hidden layers in a neural network while retaining significant structure (Maaten & Hinton, 2008).

The tuning and interpretation methodology involved close communication with a subject-matter expert, the co-creator and instructor of the course. Since different hyperparameter values of the representation learning method can produce dramatically different vectors, and consequently dramatically different reduced dimensionality visualizations, different values for vector size and window size were used to produce 21 different visualizations of course behaviour that the course instructor was asked to then rate for the amount of novel information they contained. Additional guidance was given that favourable visualizations might be ones that contained both behavioural patterns the instructor was confident existed a priori and patterns that were not anticipated but still plausible. The purpose of the visualization was to surface information not already known by the expert. The rationale was that if the visualization depicted behavioural patterns known to be true by the expert, other aspects of the visualization might also be true but not yet known. After identifying useful parameters for the base models, visualizations of student discussion posts and differences between passing and failing students were designed in concert with the instructor. An interactive d3-based visualization was developed simultaneously with this research to enable the subject-matter expert to inspect the visualization by hovering the mouse cursor over plot points to reveal semantic meta-information about the element, such as the name of the question and which lesson it belonged to. A colouring feature was also made available whereby the data points could be coloured by lesson or any other categorical feature of the element. The features of the in house-created

d3 visualization¹ were very similar to features offered by a commercial software package called Tableau but could easily be linked via URL in shared web documents and is now open-sourced to the community.

3.1. Representation Learning with Skip-Grams

When applied to natural language, a skip-gram model will learn a vector representation of a word based on the many contexts in which it appears across a large corpus of text. The length of these contexts, or number of words used to the left and right of a word to learn its representation, is called the window size. The objective function of the model is to predict the probability of words appearing in context with a given input word. The probability of all words in the vocabulary must sum to one, and the error, or loss, of the model is calculated based on the probability values of the context words. The vector representation of the word is the output of the single hidden layer network that comprises the model. The length of this vector is referred to as the vector size, which is equal to the number of nodes specified in the hidden layer. When two words have similar contexts, they will likely have similar learned vector representations. Often, synonyms of words will have vectors located close to one another in this vector space, also referred to as an embedding. The novel intuition of our application of this to student course sequences is that instead of learning the structure of language by training on sequences of words, we are learning the structure of learner behaviour from sequences of page views. It was previously found that clickstream behaviours within MOOCs could be predicted using a recurrent neural network (RNN) with 70% accuracy, compared to the 45% accuracy provided by the expected path through the course when following the existing course structure (Tang et al., 2017). This work builds on the observation that patterns exist in learner clickstream behaviours. Instead of focusing on prediction, in this work we seek to scrutinize course component embeddings to qualitatively learn what these patterns are and what their relevance is to the pedagogical design of the course. While RNNs are useful for prediction, skip-gram models are better suited for interpretation because they are linear models which create representations embedded in a vector space, subject to arithmetic and scalar manipulation. These representations, in language, are evaluated based on the percentage of pre-enumerated semantic and syntactic relationships they encode (Mikolov, Chen, Corrado, & Dean, 2013). In our application, the subject-matter expert’s implicit knowledge of the course serves to validate and interpret plausible relationships not previously known, a methodological step not yet undertaken in computational linguistics.

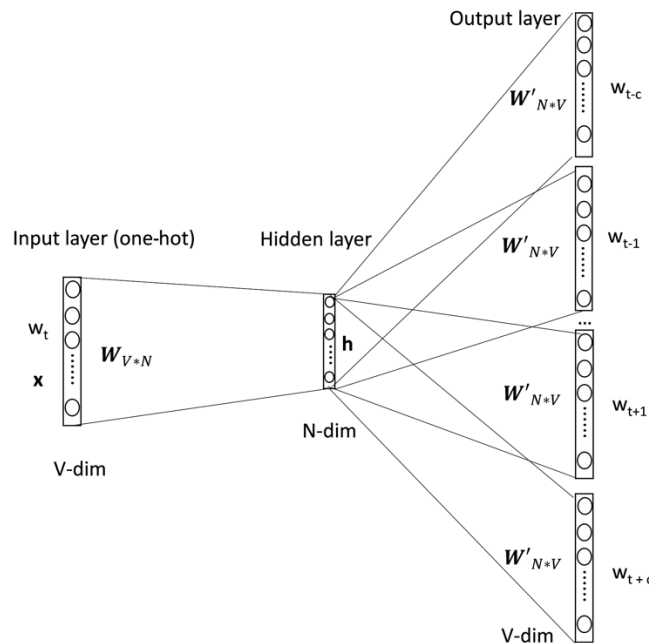


Figure 1 The skip-gram model architecture

Formally, a skip-gram is a simple feed-forward three-layer neural network with one input layer, one hidden layer, and one output layer (Figure 1). The input, in our context, is a one-hot (or dummy variable) representation of the course element, and the output can be described as a multi-hot representation of the specified number of elements in context. The number of elements in context is two times the window size, which is a hyperparameter of the model. The objective of the model is to

¹ <https://github.com/CAHLR/d3-scatterplot>

predict the elements in context given the input element. Since multiple elements are being predicted, the loss function, categorical cross-entropy, is calculated for each element in context. The number of weights trained in the model does not increase with the window size since the same weights are used to make predictions of every element in context. The second major hyperparameter is the size of the hidden layer, which is ultimately equivalent to the number of dimensions of the learned continuous course element vector. The continuous vector is the weights associated with the edges stemming from the one-hot position of the element to all the nodes in the hidden layer. In the case of natural language, the inputs are words in a vocabulary, processed by sweeping sequentially through the words in a large corpus, where the model objective is to predict the words in context given an input word. In the context of our online course interactions, the inputs are elements a student interacts with, and the model sweeps across chronological sequences of these elements. In our particular dataset of a course created in the Smart Sparrow course platform, the logged actions are students’ interactions with screens, or pages, in the course’s lessons, containing simulations, practice problems, and graded assessments. Specifically, the inputs are screen IDs, and the outputs being predicted are the screen IDs before and after the current screen ID accessed by students. No correctness information is used in this model because it is the student’s navigational behaviour that is the focus, not their performance.

In a skip-gram, the vector representation of an input screen is defined as $v_{w_i} = W^T \delta(w_i)$, where W^T is the left-side weight matrix in Figure 1, indexed by the one-hot of the input screen, $\delta(w_i)$.

A softmax layer, common to classification tasks, is used to produce a probability distribution over screens to predict screens in context: $p(w_o | w_i) = \exp(W' \delta(w_o) v_{w_i}) / \sum_{j=1}^V \exp(W' \delta(w_j) v_{w_i})$.

The probability of a given output screen, w_o , in the vocabulary, is the exponential normalization defined by the exponentiation of the input screen’s vector, v_{w_i} , multiplied by the output screen’s vector, $W' \delta(w_o)$, divided by the sum of all screens’ exponentiation of their output vector multiplied by the input vector. An output vector is the multiplication of the right-side weight matrix, W' , by a one-hot of the output screen, w_o .

The cross-entropy loss (i.e., log loss) across all students’ sequence of screens, which is back-propagated, is

$$C = - \sum_{s \in S} \frac{1}{T} \sum_{t=1}^T \sum_{-c \leq i \leq c, i \neq 0} \log p(w_{t+i} | w_t),$$

where, for each student, s , the average loss is calculated over the input screens at each time slice, t . The loss for a single time slice of a student is the sum of the log of the model’s probability of observing the screens within a time slice window size, c , of the current time slice. While the model is trained to minimize error in predicting the screens in context, the intended extract from the model after training is not its predictions of screens but rather the learned representations of the screens in the form of the weight vectors associated with each screen. These weight vectors, a product of the single hidden layer of the model, are the automatic featurization of the screen. Screens which have similar contexts become mapped to vectors of similar magnitude and direction to minimize the loss.

4. Case Study: *Habitable Worlds*

Habitable Worlds is an introductory level for-credit online lab science course developed by Dr. Ariel Anbar and Dr. Lev Horodyskyj. The course was developed from 2011 to 2015 in Smart Sparrow’s Adaptive eLearning Platform (AeLP), a Microsoft Powerpoint-esque development environment that gives instructors full control over content, layout, adaptivity, and learning pathways. The platform collects data while students interact with the educational content, and data are reported to the instructor via a dashboard and through downloadable CSV sheets should the instructor desire to conduct more detailed analyses on a particular activity. Because of these capabilities, continual refinement of *Habitable Worlds* through data analyses, student feedback, and instructor intuition is possible. A full account of the pedagogy and design philosophy of the course can be found in Horodyskyj et al. (2018). The course reached a stable version by the fall semester of 2015, with minimal subsequent changes. Despite improvements in course content through changes based on descriptive statistics like time-on-question, number of attempts, and the frequency of certain incorrect answers, general behavioural insights were difficult to come by due to the sheer amount of available data and the lack of algorithms for parsing the data into usable information.

Course content is divided into three types of activities: *training*, *applications*, and *project*. Most content in the course is linear, with occasional adaptive pathways to remediate misconceptions or bypass mastered content. *Training* activities introduce students to new concepts and give them the opportunity to explore and experiment without penalty, usually with surveys, short written texts, images, short videos (< 8 minutes), simulators, equations, problem sets, and virtual field trips (Figure 2). Points are accumulated as students pass milestones. Unlimited attempts are allowed; however, students cannot progress until they correctly complete the required activity. *Applications* are equivalent to quizzes in that students are expected to demonstrate competency on them after learning in the associated training activities, with points deducted for incorrect

answers. The *project*, essentially a final exam, is a comprehensive activity requiring students to locate rare habitable worlds in a field of 500 randomly generated stars. Course content is released on a weekly basis during the course’s 7.5-week deployment period, and students have a week to complete newly released exercises. Training activities remain open and accessible for full credit for the entire term. *Applications* close off week by week, and the project opens on the second week of the course and remains available for the rest of the term (Figure 3).

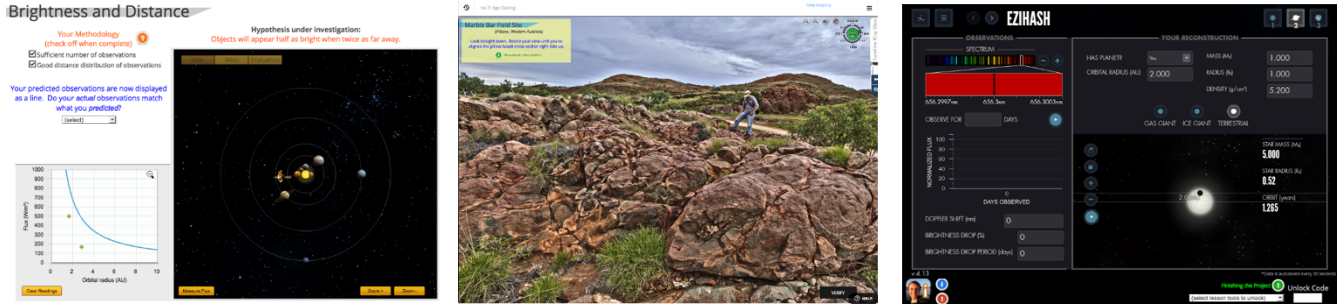


Figure 2: (left) An experimental activity with a simulator involving a hypothesis, a check on the execution of a student’s methodology, and an evaluation of their hypothesis. (centre) An observational activity at a virtual field trip. Here, students are instructed to rotate their view to observe basaltic rocks in the field in a particular orientation. (right) The course project, where students are required to find rare habitable worlds in a field of 500 randomly generated stars.

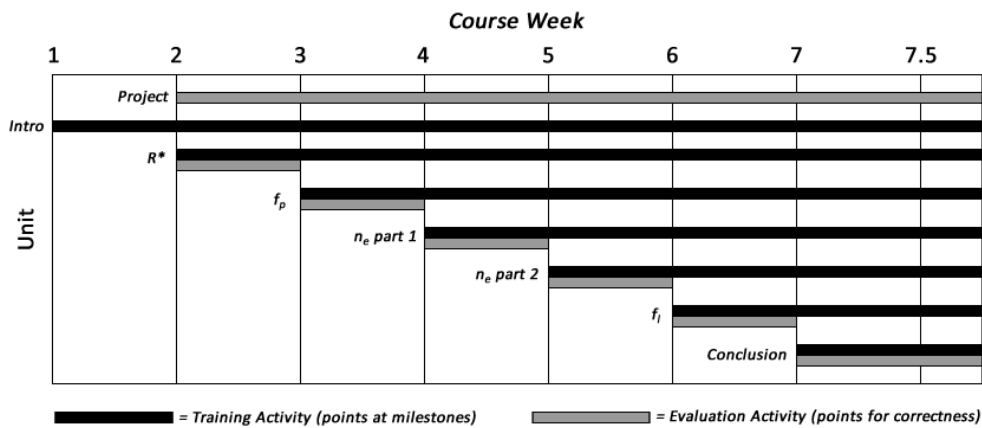


Figure 3: Schedule of activities for Habitable Worlds. Training activities open week by week and remain open throughout the term. Applications open for one-week windows. The project opens in week 2 and remains open throughout the term. Unit titles (R^* = stars, f_p = planets, n_e = Earth-like planets, f_l = life) are based on the Drake equation, a common astrobiology construct for exploring the question of life in the universe.

The intended completion strategy for *Habitable Worlds* is to alternate between a training activity and its paired application, while accessing the project on a weekly basis as units are completed. However, students are not restricted to the prescribed path. Based on aggregate data generated by the AeLP, student behaviour observed on a paired discussion board, and general intuition, it was hypothesized that students who failed the course were taking non-optimal pathways; however, the nature of those pathways had been opaque to the course instructors to this point.

4.1. Data

We used the time-stamped interactions (Table 1) of 778 anonymized students from two offerings of *Habitable Worlds*, fall 2015 and spring 2016. A skip-gram model² was used to learn continuous representations of course materials from 1.4 million temporal interactions of students with 1,644 pages of the course, called screens, within 67 different lessons. This model had several tuneable hyperparameters that changed the learned representations and thus the resultant visualizations. These were the context window sizes and the number of nodes in the hidden layer (vector size). Information on students’ final grades and

² Our research code utilizes the python genism implementation of word2vec.

contributions to the discussion board were also included in the dataset.

Table 1. Fields used to construct the dataset (left). Example sequences of screenIDs (right). These sequences, specified one row per student, constituted the dataset used to train the skip-gram model. Application (graded) screens are in bold.

Field	Description	userID	Example sequence of screenIDs
userID	Unique identifier for each student	Sam	s:1 s:2 s:3 s:2 s:3 s:4 s:5 s:6 s:7 ...
screenID	Unique identifier for each “screen” in a lesson	Erica	s:1 s:6 s:5 s:6 s:3 s:6 s:7 s:15 s:14 s:15 s:12 s:15 s:12 s:15 s:16 s:21 ...
interactionID	Unique chronologically ordered identifier for an event recorded by the AeLP (typically triggered when students click a “Check” button on the screen)	Paulo	s:1 s:2 s:3 s:4 s:5 s:6 s:7 s:8 s:9 s:10 s:11 s:12 s:13 s:14 s:15 s:7 s:8 s:9 s:10 s:11 s:12 s:13 s:14 s:15 s:15 s:15 s:16 ...

4.2. Parameter Tuning and Visualization Evaluation Results

There is no established rule of thumb as to the appropriate hyperparameters for skip-grams to use to bring out visually salient patterns in a dataset. As a result, part of our methodology was to find good values using a limited range of hyperparameters to create a set of 21 representations and respective t-SNE visualizations which were rated for their usefulness by the co-creator of the course (second author) on a five-point scale (Figure 4) in similar fashion to Géryk (2015).

As mentioned earlier, the guidance given to the expert rater was to favour visualizations that both confirmed existing intuitions and provided additional insight. After initially scanning all visuals to identify common patterns, three confirmatory patterns quickly emerged: (i) a grouping of week 2 quizzes — expected due to observations showing students racing to finish expiring activities before a hard deadline, which first happens in week 2; (ii) the beginnings of new units connecting to the endings of old units — expected based on observations of high-performing students beginning new activities as soon as they were released; and (iii) the splitting of a particular week 5 activity — expected, as the activity is excessively long. Visualizations were rated based on the presence of these three patterns. A visualization would receive extra points if, in addition to possessing the three patterns, the particular plot highlighted novel but explainable patterns.

Analyzing the results in Figure 4 from an algorithmic standpoint, the hyperparameters with smaller vector size generally lead to simpler, more singular and linearly connected data points, representing the underlying intended sequencing of the course. There is an analogy here to principal component analysis. When only allowed to represent course elements with two continuous values, these values capture the most predominant structure, which is the courseware sequencing, followed by many students. Large vector sizes can bring out second- and third-order patterns of significance, which will be looked at more closely in the next section. When the vector size was two, window sizes of one and three produced more linear, connected plots than a window size of five. With a smaller window with which to learn representations, patterns consisting of many behaviours cannot be easily considered and the representation takes on a form more common when constructing transition plots based only on the frequency of transition from one element to the next, effectively a context window of one. At higher vector sizes, a higher context window size struck a balance between providing a paucity of sequence patterns and retaining identifiable and anticipated patterns of behaviour.

4.3. Results

In summary, the highest-rated visualization was produced by a skip-gram with a window size of five and a vector size of 12. Other high vector sizes, paired with high window sizes, also yielded visualizations with useful information. Low vector sizes yielded excessively sequential plots, while low window sizes yielded plots with almost no connections between screens.

With a model of satisfactory representational quality in hand, the following sections further interrogate this informational artifact to investigate more pointed questions regarding student behaviour and its connection to the course’s pedagogy. Each iteration of the analysis adds an extra layer of complexity in order to answer research questions of increasing specificity.

4.3.1. Iteration 1 — Inspecting the Highest-Rated Visualization

The annotated visualization, most highly rated from the parameter tuning, is shown in Figure 5, with the prescribed course sequence marked with black arrows and the beginnings of new units starred. Essentially, the plot shows that students tend to approach a course that was designed to be linear in a mostly linear fashion. The visualization reveals behaviours that were previously intuited from aggregate data and discussion board observations. Although helpful in visualizing and confirming student behaviours that had been previously assumed, the skip-gram revealed only limited additional information for course improvements, mostly related to which lessons ought to be split due to being too lengthy.

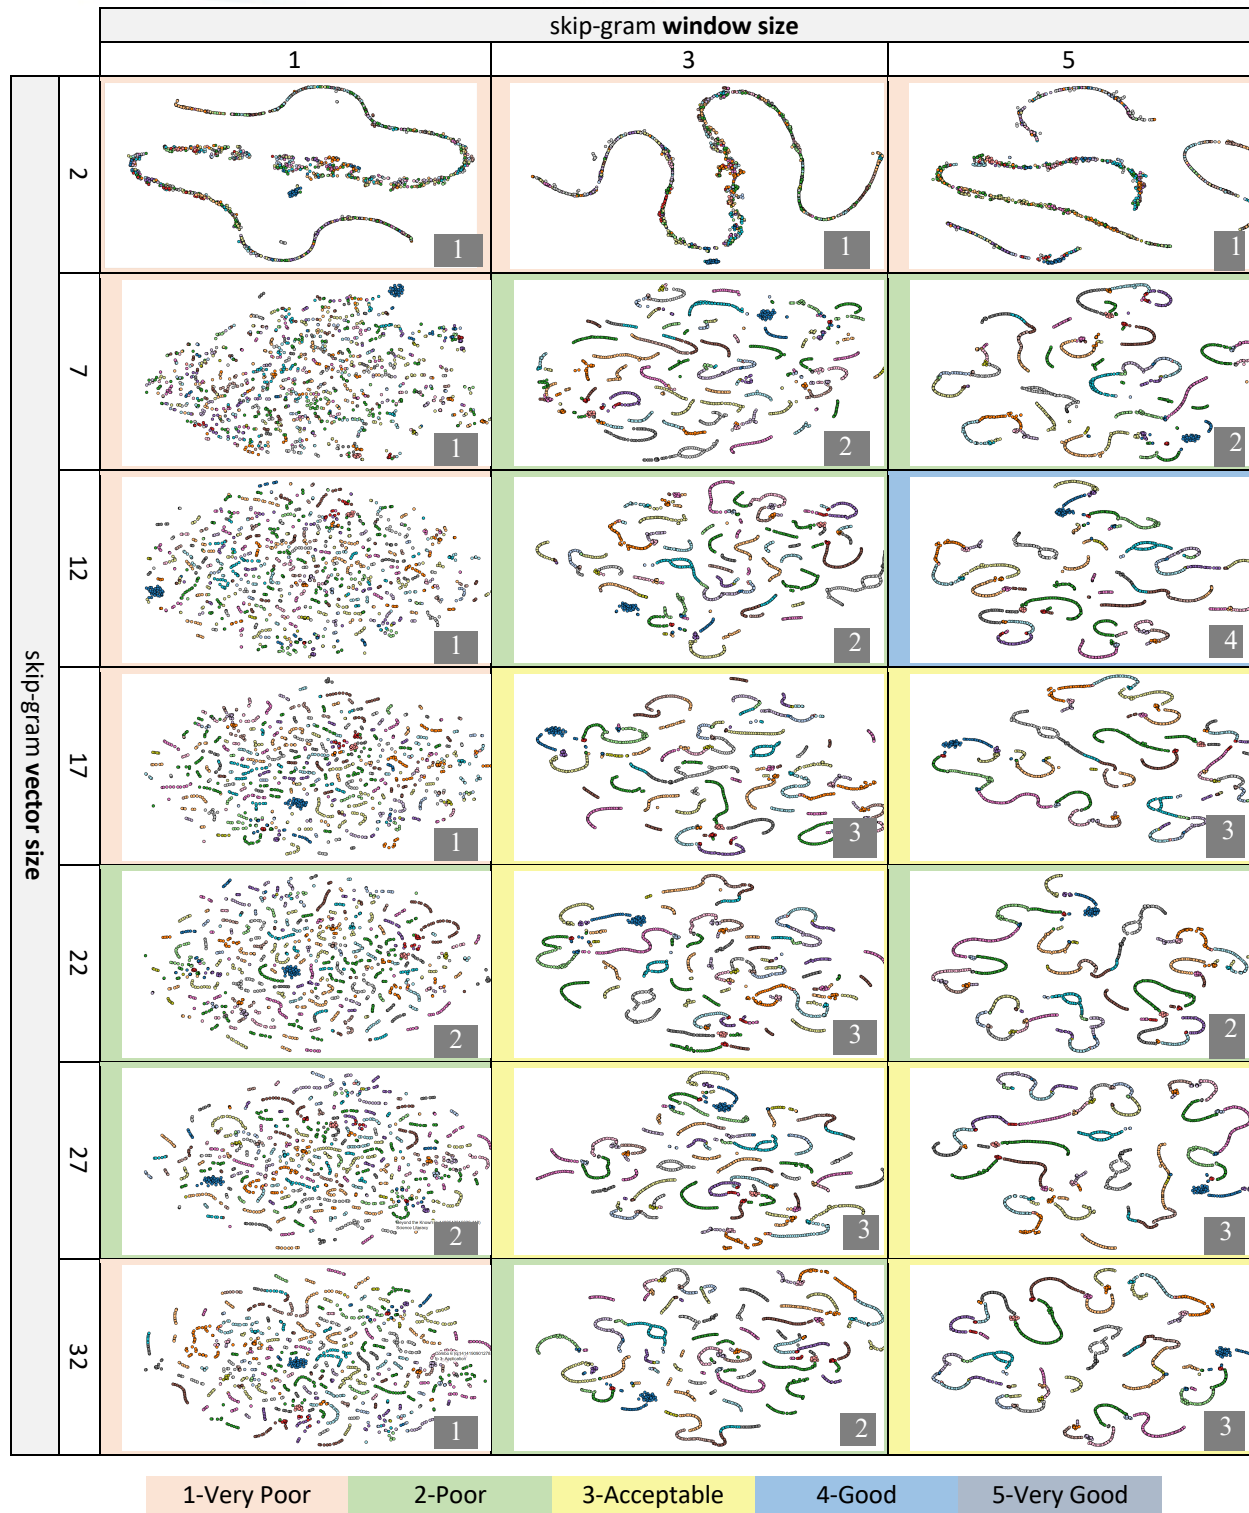


Figure 4: Skip-gram t-SNE visualizations with a variety of window sizes (1–5) and vector sizes (2–32) with instructor usefulness ratings shown in the lower right corner. Each of the 21 scatter plots depicts the complete set of elements in the course and their relationship to one another. Colours of plot points represent the lesson within the course to which each point belongs. Useful plots were those which depicted hypothesized and plausibly explainable relationships between course elements.

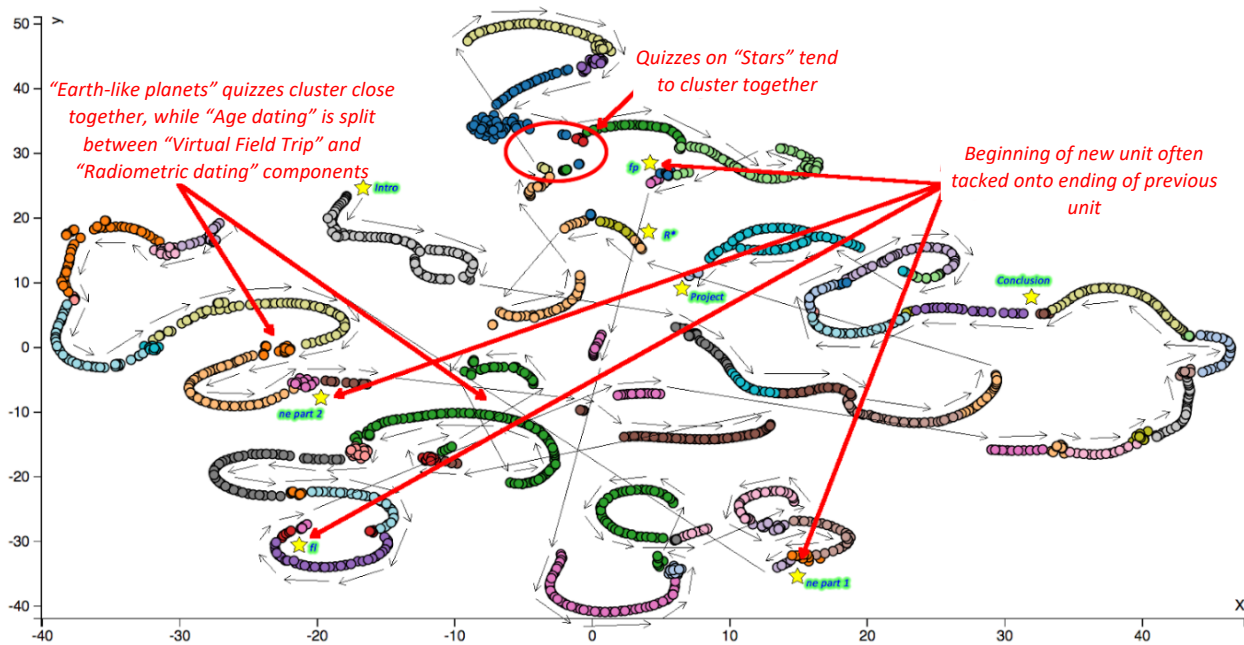


Figure 5: Highest-rated visualization of course elements from parameter-tuning stage with instructor annotations. Axes are abstract and represent dimensionality reduced vector values. Grey arrows indicate the sequence of content in the course, stars indicate the beginnings of a new unit, and red notes indicate patterns of interest. Each lesson is represented by a different colour. The typical pattern consists of a long linear sequence representing the training activity followed by a more clustered quiz activity. Course material progresses from *Intro* to R^* to f_p to n_e (part 1) to n_e (part 2) to f_i to *Conclusion*.

4.4. Iteration 2 — Differentiating Behaviours Exhibited by Passing and Failing Students

Failure in courses is often attributed to student difficulties with content or disengagement. Instructors often have little insight into why a student has failed a course, aside from observing that a student has either stopped attending class or performed poorly on an assignment. Analytical data from learning management systems and MOOCs have provided new windows into student behaviour, particularly when interacting with digital content. Previous informal analyses of *Habitable Worlds* offerings showed differences in engagement between passing, failing, and withdrawing students, in addition to level of content mastery and “attendance” based on whether course content was accessed. Passing students fully engaged with a majority of the content, while students who withdrew engaged with little, if any, content. Failing students, surprisingly, showed persistence in the course, often engaging with content week after week, though not completing it successfully.

The second iteration of the skip-gram featured the incorporation of pass/no pass information, creating two of each course element. If a student had passed the course, each element in the student’s sequence would be prefixed with a “p” (e.g., “p-s:12”); the element would be prefixed with an “n” if the student had not passed. A single model was trained to represent these elements along with a single dimensionality reduction, but since no student sequence consisted of both pass and no-pass elements, the two sets of representations do not substantively share the same space because they never appear in each other’s context windows. This single projection of passing and non-passing student course representations can be seen in Figure 6. Although content in *Habitable Worlds* is linear, individual training (10–60 screens), application (~5 screens), and project (~5 screens) lessons can be accessed in any order, if desired, during any given week. The intended course pathway is to complete a training activity followed by its associated application activity, repeated for each topic (between four and six for each week). Once the unit is complete, the associated component of the project can be completed. Release of units of the course is time-gated, but lessons within the units, once released, are not gated.

The introductory unit of the course, completed in week 1, does not have any associated applications, only training activities. Paired applications begin with week 2 content. For both passing and failing students, week 2 (R^*) applications cluster together, indicating that students are accessing them close together in time. This may indicate that students are racing the deadline and attempting to complete the hard-deadline content (training activities have no deadline). Passing students seem to realize that this is not an optimal strategy and switch to a more sequential approach in subsequent units, where they complete training activities followed by applications as they appear in the sequence of the unit, a pattern consistent with self-regulated learning.

Failing students, however, make this switch much later in the term. This results in a hub-and-spoke pattern for each unit on the visualization, where applications cluster together, with training activities radiating from them. This indicates that students may be attempting to complete the applications first, only proceeding to the training activities to find the answers to the applications.

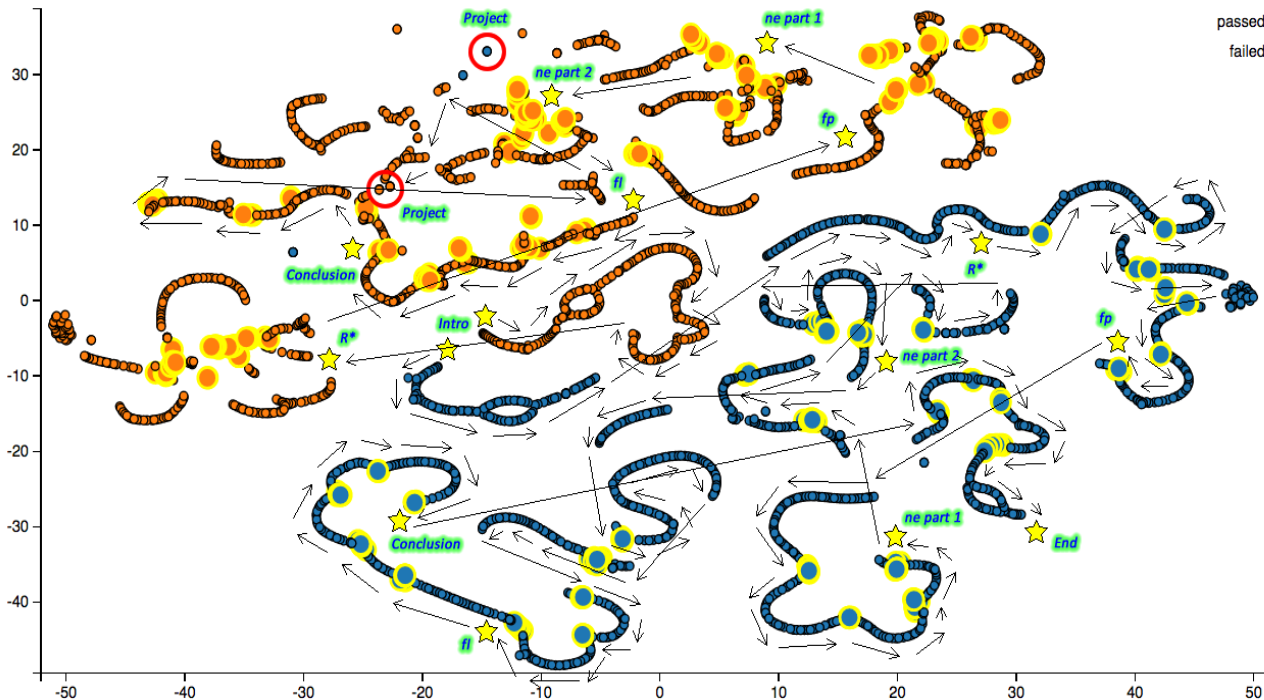


Figure 6: Behaviour of failing (orange) and passing (blue) students. Screens associated with training activities are outlined in black, quiz/application screens are highlighted in bright yellow, and project screens are circled in red. Passing students show a linear progression through content, with quizzes/applications appearing sequentially after their associated training activities. Failing students show a hub-and-spoke pattern for most units, with training activities converging onto a cluster of quiz/application activities for that unit. The project clusters far from activities for passing students (indicating random access during the semester) and close to the concluding unit for failing students.

In addition, there are differences between passing and failing students in the visualization with respect to the final project. The project is released in week 2 and is a fairly complex endeavour, requiring students to utilize skills learned in almost every unit of the course and assembling those skills into a methodology to find a handful of habitable worlds in a field of hundreds of stars. The optimal strategy is to engage with the project early and complete components of the project as the concepts are learned. For passing students, the project clusters quite far away from the rest of the course, indicating that they are accessing the project throughout the course, hence the lack of association with any particular week’s activities. For failing students, however, the project clusters very close to the Conclusion unit, indicating that failing students do not engage with the project until all other course material has been completed.

Overall, this visualization, where students are differentiated based on their grade, revealed that although both groups of students take the precarious strategy early in the course, passing students subsequently adopt a more optimal strategy while failing students do not, continuing to struggle week after week. These results are consistent with previous work showing that successful students in active-learning settings utilize deep-learning strategies while struggling students in the same context suffer because they are adopting shallow-learning strategies (Gašević et al., 2017). This is, however, inconsistent with previous work indicating that students move toward less effective learning strategies over time (Jovanović et al., 2017). This may be a function of course design, where shallow learning strategies (such as guessing, relying heavily on preconceptions, or skimming video and reading content) are consistently frustrated by dead-end alternative pathways or loops designed to rectify misconceptions. This results in a course where the “path of least resistance” requires the continuous application of deep-learning strategies, because shallow-learning strategies result in a significant increase in work for a student to complete.

4.5. Iteration 3 — Discussion Forums

Habitable Worlds is paired with a discussion help forum on the Piazza platform. Students who are stuck or require a more detailed explanation of any course concept can post their questions to the forum, where either an instructor or a fellow student

can reply. Average response time is 5–10 minutes during most of the day. Instructors typically offer assistance for training activities only, while fellow students are allowed to offer assistance on training, applications, and the project. Because most of the evaluations are randomly generated using equations and algorithms in the AeLP, students can exchange techniques for solving problems but not answers, as no two students have the same problem sets. This results in a collaborative environment in which students and instructors work together to learn and master skills and concepts, mimicking a real scientific environment.

Forum posts, but not views, were also part of our time-stamped data and were added to the skip-gram input sequences to understand how passing and failing students utilized the discussion board (Figure 7). These representations were also learned with a single model, but, for this analysis, the projection (i.e., dimensionality reduction) for each pass/no-pass group was generated separately. With few exceptions, passing students’ forum postings do not tend to cluster with any particular training/application combination. This suggests that passing students are posting on the discussion board only when they need help, and this differs from student to student, resulting in clustering of posts away from any particular lessons.




Figure 7. Behaviour of failing (orange) and passing (blue) students. Screens associated with training activities are outlined in black, quiz/application screens are highlighted in bright yellow, project screens are circled in red, and forum discussion posts are grey-filled circles. (top) Failed student forum behaviour, showing some correlations between forum posts and application activities and little association between forum posts and project work. (bottom) Passed student forum behaviour showing little association between forum activity and application activities, but strong association between forum posts and project work.

Failing students’ forum postings tend to have a slightly stronger association with applications than for passing students, although this is not universal. This indicates that when running into problems, failing students reach out for assistance, a desired behaviour. However, because this tends to happen during the *application* activities rather than the training activities, this indicates that these students have likely not mastered the content in the training activities, and so when asked to apply their knowledge, they struggle. The reason for the lack of content mastery is unclear. Piazza records the number of posts and threads “read” by students but does not timestamp this particular activity. Many students have reported that they do not post on the

board often because they start activities late and find that when they become stuck, their questions have already been answered on the forum days earlier. Training activities cannot always be randomized because they are teaching conceptual frameworks, and these frameworks often have single explanations (e.g., there are only so many ways to teach “higher temperatures = more evaporation”). If a student is using the discussion board as a crutch to complete difficult training content by doing what previous students have reported as resulting in success without understanding *why* it led to success, it is likely that they are not internalizing the concepts of the course. Hence, when they reach the application, which creates randomized activities based on the underlying concepts, it is likely that they will continue to struggle despite reaching out for help because they never gained a deeper understanding of the underlying concept that is being tested in the application.

There is a notable reversal of this behaviour for the course project. The project is a summative assessment, equivalent to a final exam. Students are presented with 500 randomized stars and are required to find rare habitable worlds, synthesizing techniques learned in the class, and reinforced in applications, into a search strategy. Students can develop and exchange strategies and check each others’ work, but they cannot trade answers because each star set is unique to each student. A collaborative environment is not necessary to successfully complete the project and is purely optional. Passing students’ project pages cluster very closely with forum posts, indicating that these students are taking advantage of the collaborative environment available to them. This matches observations where instructors have seen collaborative groups forming and developing strategies together. For failing students, project pages do not show a close correlation with discussion posts. This clustering pattern of project pages suggests that failing students are *not* collaborating on an activity for which collaboration is most beneficial. They may be skimming information from other discussions for strategies but do not seem to be engaged in the process of synthesizing strategies themselves.

4.6. Habitable Worlds Case Study Discussion and Conclusions

The visualization methodology corroborated a variety of existing intuitions held by the instructor from interacting with students via discussion boards, by e-mail, and in person over the years that *Habitable Worlds* has run. The pathway visualization supported the expectation that students followed a mostly linear path through the course. The relative length of lesson paths, on a cursory level, helped to identify content that could be split into smaller pieces to benefit learners. More significantly, the analyses revealed substantial navigational differences between passing and failing students. The interpretation of differences depicted by the visualization was that passing students take the quizzes in tandem with the practice material while the failing students go directly to the quizzes and then seek out the answers in the learning material. This information can prove essential to formative course development.

Habitable Worlds was designed to be approached in whatever way a student finds most comfortable. A student can pace their work across an entire week or complete it all in a short burst. Students can work alone or together, either online or in person. The course project can be completed throughout the entire semester or all at the very end. This design was intentional to allow maximum flexibility for students, many of whom are non-traditional and have significantly more responsibilities than a typical on-campus student. However, for many failing students, this design may be detrimental as it depends on a certain amount of self-discipline and awareness of personal limitations, a skill that is often not well developed in novice learners. The representation learning visualizations confirmed an existing assumption about failing students’ unsophisticated strategy toward engaging with the course project but also revealed that this strategy was applied to the rest of the course content and likely the usage of the discussion board as well. Students who fall into this category may require a more structured approach or alternative incentives to adopt successful study strategies for completing the course. Future versions of the course will focus on building better supports, such as an early module on successful course strategies, and course structure enhancements, such as adding flexible deadlines and enforced or incentivized content ordering. These planned modifications are expected to better support those who are struggling not just with the subject matter of the course but also with learning to learn in an online environment.

5. CONTRIBUTIONS AND DISCUSSION

We have introduced a methodology around the nascent field of representation learning visualization and its interpretation and applied it to the general topic of summarizing the temporal behaviours of learners. The intention of this approach was to require as few assumptions about behaviour as possible, instead allowing prominent features of learner navigation to surface organically with the aid of careful model and visualization tuning that took place between the researcher and the practitioner closest to the domain. By allowing a model to take the place of an expert in creating the abstract featurization of behaviour, novel insights could be made by the expert in his interpretation of the visualization.

This approach can be seen as a type of human-computer interaction (HCI) or, more specifically, a type of human-AI interaction, where (1) an AI learns a set of representations of the world (via skip-gram) that allow it to make predictions of behaviour, (2) a human expert builds confidence in the AI’s epistemic validity by selecting the representation that exhibits information known by the expert (i.e., model selection), and (3) observations are made by the expert from the representation

visualization that support novel insights about behavioural patterns. Like any other expert, an AI can be wrong, and corroborating support for an interpretation should be gathered before a pattern is considered a bona fide generalizable phenomenon. Nevertheless, a machine-learned representation can serve as a unique perspective on behaviour that can prove to be an indispensable companion source of information in design.

6. FUTURE WORK

The machine-learned representations in this work, with the addition of assessment constructs (Pardos & Dadu, 2017), might be comparable to an epistemic frame, which Shaffer (2006) describes as akin to the relationship between discursive practices (e.g., forms of interaction) and structures of knowledge. Future work could investigate the nature of cognitive assessment information encoded by a learned representation versus a more explicit expert representation (Shaffer & Ruis, 2017). Although screen sequences and discussion posts were utilized in this study, there is a multitude of additional data that can be used in future iterations of the work applied to courseware or other contexts in which temporal learner data are collected. These additional data include problem answer text and correctness information, granular activity involving within-page interactive widgets (e.g., simulations), and interactions and communications with peers and instructional staff outside the discussion board.

An obstacle to adoption of our method by a broader audience of instructors, without researcher support, is the rather careful process by which the parameters of the visualization were initially tuned. Since this was a collaborative research endeavour, guided by the educational theory-informed goals of the instructor (Hillaire, Rappolt-Schlichtmann, & Ducharme, 2016), it was justified for this laborious effort to be undertaken by a member of the instructional staff; however, this should not be a prerequisite, in practice, for a typical instructor to take on before interfacing with these analytics. The technique needs to be applied to more courses and more settings that vary by the type of students, course material, and platform to ascertain whether common useful settings emerge in general contexts which would reduce or eliminate the upfront manual tuning effort. It is also an open question whether a threshold level of analytical data fluency or professional development is required for a domain expert to independently begin to make sense of these visualizations.

Acknowledgements

We thank Smart Sparrow and the Inspark Teaching Network for facilitating this work and providing anonymized student data. This research was conducted in accordance with the Arizona State University institutional review board (Study #3679).

Declaration of conflicting interest

The authors declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

Funding

The publication of this article received financial support from the Bill & Melinda Gates Foundation and the National Science Foundation (Award #1225741).

References

- Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8), 1798–1828. <http://doi.org/10.1109/TPAMI.2013.50>
- Bergner, Y., Shu, Z., & von Davier, A. (2014). Visualization and confirmatory clustering of sequence data from a simulation-based assessment task. In Z. Pardos & J. Stamper (Eds.). *Proceedings of the 7th International Conference on Educational Data Mining (EDM '14)*, 4–7 July 2014, London, UK (pp. 177–184). International Educational Data Mining Society. http://educationaldatamining.org/EDM2014/uploads/procs2014/long%20papers/177_EDM-2014-Full.pdf
- Boyer, S., & Veeramachaneni, K. (2015). Transfer learning for predictive models in massive open online courses. In C. Conati, N. Heffernan, A. Mitrovic, & M. F. Verdejo (Eds.). *Proceedings of the 17th International Conference on Artificial Intelligence in Education (AIED '15)*, 22–26 June 2015, Madrid, Spain (pp. 54–63). Springer. https://doi.org/10.1007/978-3-319-19773-9_6
- Breslow, L., Pritchard, D. E., DeBoer, J., Stump, G. S., Ho, A. D., & Seaton, D. T. (2013). Studying learning in the worldwide classroom: Research into edX's first MOOC. *Research & Practice in Assessment*, 8, 13–25. <https://www.rpajournal.com/dev/wp-content/uploads/2013/05/SF2.pdf>
- Caprotti, O. (2017). Shapes of educational data in an online calculus course. *Journal of Learning Analytics*, 4(2), 76–90. <http://doi.org/10.18608/jla.2017.42.8>

- Chuang, I., & Ho, A. D. (2016). HarvardX and MITx: Four years of open online courses (Fall 2012–Summer 2016). <http://doi.org/10.2139/ssrn.2889436>.
- Crossley, S., Paquette, L., Dascalu, M., McNamara, D. S., & Baker, R. S. (2016). Combining click-stream data with NLP tools to better understand MOOC completion. In S. Dawson, H. Drachler & C. Rosé (Eds.). *Proceedings of the 6th International Conference on Learning Analytics & Knowledge (LAK '16)*, 25–29 April 2016, Edinburgh, UK (pp. 6–14). New York: ACM. <https://doi.org/10.1145/2883851.2883931>
- Davis, D., Chen, G., Hauff, C., & Houben, G. J. (2016). Gauging MOOC learners' adherence to the designed learning path. In E. Snow & J. Rowe (Eds.). *Proceedings of the 9th International Conference on Educational Data Mining (EDM2016)*, 29 June–2 July 2016, Raleigh, NC, USA (pp. 54–61). International Educational Data Mining Society. http://www.educationaldatamining.org/EDM2016/proceedings/paper_63.pdf
- de Baker, R. S. J., Gowda, S. M., Wixon, M., Kalka, J., Wagner, A. Z., Salvi, A., Aleven, V., Kusbit, G., Ocumpaugh, J., & Rossi, L. (2012). Towards sensor-free affect detection in cognitive tutor algebra. In K. Yacef, O. Zaïane, H. Hershkovitz, M. Yudelson, & J. Stamper (Eds.). *Proceedings of the 5th International Conference on Educational Data Mining (EDM '12)*, 19–21 June 2012, Chania, Greece (pp. 126–133). International Educational Data Mining Society. <https://files.eric.ed.gov/fulltext/ED537205.pdf>
- Ellson J., Gansner E., Koutsofios L., North S.C., Woodhull G. (2002). Graphviz — Open source graph drawing tools. In P. Mutzel, M. Jünger, & S. Leipert (Eds.). *Graph Drawing (GD 2001)*. *Lecture Notes in Computer Science*, 2265. Berlin, Heidelberg: Springer. https://doi.org/10.1007/3-540-45848-4_57
- Emmons, S. R., Light, R. P., & Börner, K. (2017). MOOC visual analytics: Empowering students, teachers, researchers, and platform developers of massively open online courses. *Journal of the Association for Information Science and Technology*, 68(10), 2350–2363. <http://doi.org/10.1002/asi.23852>
- Fratamico, L., Perez, S., & Roll, I. (2017). A visual approach towards knowledge engineering and understanding how students learn in complex environments. In C. Thille & J. Reich (Eds.). *Proceedings of the 4th ACM Conference on Learning @ Scale (L@S 2017)*, 20–21 April 2017, Cambridge, MA, USA (pp. 13–22). New York: ACM. <https://doi.org/10.1145/3051457.3051468>
- Freitag, D. (2000). Machine learning for information extraction in informal domains. *Machine Learning*, 39(2–3), 169–202. <https://doi.org/10.1023/A:1007601113994>
- Gašević, D., Jovanović, J., Pardo, A., & Dawson, S. (2017). Detecting learning strategies with analytics: Links with self-reported measures and academic performance. *Journal of Learning Analytics*, 4(2), 113–128. <https://doi.org/10.18608/jla.2017.42.10>
- Géryk, J. (2015). Using visual analytics tool for improving data comprehension. In C. Romero & M. Pechenizkiy (Eds.). *Proceedings of the 8th International Conference on Educational Data Mining (EDM '15)*, 26–29 June 2015, Madrid, Spain (pp. 327–334). International Educational Data Mining Society. <https://files.eric.ed.gov/fulltext/ED560521.pdf>
- Hillaire, G., Rappolt-Schlichtmann, G., & Ducharme, K. (2016). Prototyping visual learning analytics guided by an educational theory informed goal. *Journal of Learning Analytics*, 3(3), 115–142. <https://doi.org/10.18608/jla.2016.33.7>
- Horodyskyj, L. B., Mead, C., Belinson, Z., Buxner, S., Semken, S., & Anbar, A. D. (2018). Habitable Worlds: Delivering on the promises of online education. *Astrobiology*, 18(1), 86–99. <https://doi.org/10.1089/ast.2016.1550>
- Jovanović, J., Gašević, D., Dawson, S., Pardo, A., & Mirriahi, N. (2017). Learning analytics to unveil learning strategies in a flipped classroom. *The Internet and Higher Education*, 33, 74–85. <https://psycnet.apa.org/doi/10.1016/j.iheduc.2017.02.001>
- Köck, M., & Paramythis, A. (2011). Activity sequence modelling and dynamic clustering for personalized e-learning. *User Modeling and User-Adapted Interaction*, 21(1), 51–97. <https://doi.org/10.1007/s11257-010-9087-z>
- Koedinger, K. R., D'Mello, S., McLaughlin, E. A., Pardos, Z. A., & Rosé, C. P. (2015). Data mining and education. *Wiley Interdisciplinary Reviews: Cognitive Science*, 6(4), 333–353. <https://doi.org/10.1002/wcs.1350>
- Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Advances in Neural Information Processing Systems*, 25, 1097–1105. <https://doi.org/10.1145/3065386>
- Maaten, L. V. D., & Hinton, G. (2008). Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(Nov), 2579–2605. <http://www.jmlr.org/papers/volume9/vandermaaten08a/vandermaaten08a.pdf>
- Mikolov, T., Chen, K., Corrado, G., & Dean, J. (2013). Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.
- Mikolov, T., Yih, W., & Zweig, G. (2013). Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 9–14 June 2013, Atlanta, GA, USA (pp. 746–751). Stroudsburg, PA, USA: The Association

- for Computational Linguistics. <http://aclweb.org/anthology/N13-1090>
- Pardos, Z. A., Dadu, A. (2017). Imputing KCs with representations of problem content and context. In F. Cena & M. Desmarais (Eds.). *Proceedings of the 25th Conference on User Modeling, Adaptation and Personalization (UMAP '17)*, 9–12 July 2017, Bratislava, Slovakia (pp. 148–155). New York: ACM. <https://doi.org/10.1145/3079628.3079689>
- Pardos, Z. A., Farrar, S., Kolb, J., Peh, G. X., & Lee, J. H. (2018). Distributed representation of misconceptions. In J. Kay & R. Luckin (Eds.). *Proceedings of the 13th International Conference of the Learning Sciences (ICLS '18)*, 23–27 June 2018, London, UK (3 volumes, pp. 1791–1798). International Society of the Learning Sciences. https://www.researchgate.net/publication/325672563_Distributed_Representation_of_Misconceptions
- Pardos, Z. A., Fan, Z., Jiang, W. (2019). Connectionist recommendation in the wild: On the utility and scrutability of neural networks for personalized course guidance. *User Modeling and User-Adapted Interaction*. <https://doi.org/10.1007/s11257-019-09218-7>
- Park, J., Denaro, K., Rodriguez, F., Smyth, P., & Warschauer, M. (2017). Detecting changes in student behavior from clickstream data. In X. Ochoa, I. Molenaar, & S. Dawson (Eds.). *Proceedings of the 7th International Conference on Learning Analytics & Knowledge (LAK '17)*, 13–17 March 2017, Vancouver, BC, Canada (pp. 21–30). New York: ACM. <http://doi.org/10.1145/3027385.3027430>
- Ribeiro, L. F., Saverese, P. H., & Figueiredo, D. R. (2017). struc2vec: Learning node representations from structural identity. In R. Kumar, T. Eliassi-Rad, R. J. Bayardo, & C. Elkan (Eds.). *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '17)*, 13–17 August 2017, Halifax, NS, Canada (pp. 385–394). New York: ACM. <https://doi.org/10.1145/3097983.3098061>
- Seaton, D. T., Bergner, Y., Chuang, I., Mitros, P., & Pritchard, D. E. (2013). Who does what in a massive open online course? *Communications of the ACM*, 57(4), 58–65. <https://doi.org/10.1145/2500876>
- Shaffer, D. W. (2006). Epistemic frames for epistemic games. *Computers & Education*, 46(3), 223–234. <https://doi.org/10.1016/j.compedu.2005.11.003>
- Shaffer, D. W., & Ruis, A. R. (2017). Epistemic network analysis: A worked example of theory-based learning analytics. In C. Lang, G. Siemens, A. Wise, & D. Gašević (Eds.). *Handbook of Learning Analytics*, 1st edition (pp. 175–187). Alberta, Canada: Society for Learning Analytics Research (SoLAR). <https://doi.org/10.18608/hla17.015>
- Stamper, J., & Pardos, Z. A. (2016). The 2010 KDD Cup Competition Dataset: Engaging the machine learning community in predictive learning analytics. *Journal of Learning Analytics*, 3(2), 312–316. <https://doi.org/10.18608/jla.2016.32.16>
- Tang, S., Peterson, J., & Pardos, Z. (2017). Predictive modelling of student behaviour using granular large-scale action data. In L. C. Lang, G. Siemens, A. Wise, & D. Gašević (Eds.). *Handbook of Learning Analytics*, 1st edition (pp. 223–233). Alberta, Canada: Society for Learning Analytics Research (SoLAR). <https://doi.org/10.18608/hla17.019>
- Xu, Z., Goldwasser, D., Bederson, B. B., & Lin, J. (2014). Visual analytics of MOOCs at Maryland. In A. Fox, M. T. H. Chi, & M. Hearst (Eds.). *Proceedings of the 1st ACM Conference on Learning @ Scale (L@S 2014)*, 4–5 March 2014, Atlanta, GA, USA (pp. 195–196). New York: ACM. <http://doi.org/10.1145/2556325.2567878>