Chapter 9

VET analytics

Richard Lee Davis, Son Do-Lenh, Mina Shirvani Boroujeni, Ramtin Yazdanian

This book highlights a variety of digital technologies developed during the Dual-T project, from AR to social networking platforms. One of the features that these technologies had in common was that they generated new kinds of data. In some cases, these data were fine-grained traces of individual students' activities. In other cases, these data provided higher-level perspectives on entire VET institutions or groups of institutions. We believe that these data were an essential part of the digital technologies we were developing and we devoted considerable effort to figuring out ways to utilise these data in our work.

Most of these data could not be easily analysed or understood with standard statistical methods. Making sense of these data necessitated a turn to the world of learning analytics. Learning analytics is an interdisciplinary approach that brings together methods from machine learning to data mining to statistics to HCI. A widely-used definition of learning analytics is 'the measurement, collection, analysis and reporting of data about learners and their contexts, for purposes of understanding and optimising learning and the environments in which it occurs' (*https://tekri.athabascau.ca/analytics/*).

Learning analytics can be used to help meet a variety of objectives, including monitoring and analysing learners' activities in order to gain a better sense of what is happening in the classroom, making predictions about students' performance to identify and support those who need extra attention, assessing what students know and using this as a source of feedback to aid students in their studies, and personalising and adapting instructions to students' needs. Some of the methods used to meet these objectives include predictive modelling (e.g. regression), social network analysis, clustering, and data visualisation (readers interested in a more extensive treatment of the objectives and methods of learning analytics in the literature could look to Baker (2010), Chatti et al. (2012) or Clow (2013).)

This breadth of methods and objectives can be overwhelming, and a number of taxonomies have been proposed for categorising and organising learning analytics. One that we found to be both simple and useful was developed by Chatti et al. (2012). This 'reference model' suggests that learning analytics projects can be organised according to four questions: How? What? Why? Who? Any given project uses a specific set of methods (How?) on a dataset (What?) to achieve some objective or answer some question (Why?) posed by stakeholders (Who?). This reference model is helpful when trying to organise and make sense of the different types of analyses that fall under the umbrella of learning analytics. It is also useful in setting a boundary around the field of learning analytics. For example, if the answer to 'Why?' does not directly relate to student learning, it is not learning analytics. Likewise, if the answer to 'Who?' includes people from outside the formal educational system, such as managers in a company, it is not learning analytics.

Practically, this means that many of the ways we have used data in the Dual-T project do not technically fall under the umbrella of learning analytics. One of the points of this book (and the Dual-T project) is that the vocational education system, especially the dual vocational system, is uniquely different from other education systems. It has different objectives, stakeholders, and forms of assessment and evaluation, and it is distributed across multiple sites. Just as these differences need to be considered when designing and implementing digital technologies for VET, they must also be considered when developing analytics tools and methods. For this reason, we have coined a new term that describes our work: VET analytics. VET analytics is the measurement, collection, analysis and reporting of data from *the entire VET system* for the purposes of understanding and optimising *all aspects of vocational education and training*. This is a more expansive definition, which encompasses a larger group of stakeholders (e.g. apprentices, administrators, instructors, in-company trainers, employer associations) and a larger set of outcomes (e.g. communication patterns between stakeholders, feedback patterns related to learning documentation, changes in the types of skills valued by industry).

The VET system and the world of industry are deeply interconnected. Employer associations play an important role in shaping the curriculum so that students learn skills which are relevant to their industries, and industries take on some of the responsibility of training students through apprenticeship programmes. This means that a narrow focus on student learning in the classroom is not appropriate for VET. While this narrow focus might improve students' understanding in the classroom, it ignores the many connections to points outside of the classroom which are essential to VET. For example, this narrow focus ignores whether the things being learned in the classroom connect meaningfully to the students' apprenticeships and fails to consider whether the things students are learning are relevant to future employers. VET analytics takes all of these factors into account. Its focus is not only on improving student learning, but on optimising the entire VET system.

We illustrate VET analytics with three vignettes.

In the first vignette, we describe how learning analytics was the key to unlocking the pedagogical value of the TinkerLamp. The story of the TinkerLamp has been told in rich detail in Chapter 6, but there the emphasis was on the co-development of the technology with VET stakeholders and the importance of integrating classroom orchestration tools with the platform. Here, we focus more on the TinkerBoard, the data visualisation dashboard which was the key to unlocking the pedagogical value of the TinkerLamp system in the classroom. This vignette illustrates the value of 'closing the loop', that is, using the data captured by digital technologies as a source of feedback which helps to ensure that the technologies are being used effectively. Closing the loop turns out to be of central importance in all three vignettes, and the importance of closing the loop is one of the key takeaways of this chapter.

In the second vignette, we talk about Mina Shirvani Boroujeni's analysis of communication patterns between instructors, in-company trainers and students on the Realto platform. This analysis identified dysfunctional communication patterns between the different stakeholders in the VET system, in particular how instructors and in-company trainers were not reliably responding to students' requests for feedback. The more we investigated this issue the more we realised that it was a widespread problem. It was also present on e-DAP and LearnDoc, online platforms we had developed for chefs and bakers respectively. Together with Christian Gianneti, we attempted to address this problem by setting up workshops for in-company trainers to show them the value of providing timely feedback to apprentices on the e-DAP platform (see Chapter 3) or by introducing them to Realto in the workshops that all three stakeholders attended together (see Chapter 4).

The third and final vignette focuses on the skills that students learn during their training. The VET system is designed to provide students with skill sets that prepare them to be effective and productive employees. In Switzerland, these skillsets are formalised in training plans which are updated every five years in consultation with regional and national industry organisations. This is a top-down approach to keeping the VET curriculum in line with the needs of industry. However, we realised that new forms of data had made it possible to pursue a bottom-up approach. Over the past decade, large numbers of online job ads had been collected across a variety of professions, including most VET professions. We extracted the skills from the plain text of each job advertisement and used this to track the rise and fall of in demand skills in the labour market. We then used this data to build a predictive model of which skills were most likely to emerge in the coming years. Our work on this project is ongoing, but we believe that when this information is used to close the loop and help design future curricula, it will make the VET system more responsive to the needs of the labour market.

The TinkerBoard Story

In Chapter 6 we introduced the TinkerLamp, a tangible, an AR system for teaching apprentice logisticians about storage optimisation. There was an aspect of that story that we only touched on briefly, but which deserves more attention for what it can teach us about VET analytics. This is the introduction of the TinkerBoard to the TinkerLamp system.

The TinkerBoard is a data visualisation dashboard designed by Son Do-Lenh which is meant to be permanently projected on a wall for the duration of the TinkerLamp activity. It is an interactive system which collects, processes and displays information while also providing the teacher with the ability to manipulate the display of information and control the TinkerLamps. In the following paragraphs, we re-introduce the TinkerBoard by telling the story of how Jacques Kurzo used it in his teaching. For a more complete overview of the TinkerBoard, see Chapter 6.

Jacques had split the class up into four groups and assigned each group to work with a TinkerLamp. After a brief introduction, the students began to create and test different warehouse layouts. Jacques walked around the classroom, answering questions and offering help to the different groups. Once the groups had created and tested five warehouse layouts, Jacques hit a button on the TinkerBoard to pause all the lamps. Freezing the simulations helped to bring the students' attention to the front of the class, where Jacques was standing next to the TinkerBoard projection (see Figure 6-9 in Chapter 6).

In front of the class, Jacques selected one layout from each group and added them to the ComparisonZone of the TinkerBoard. Presented under each layout were statistics about the gross area, the gross storage area, the net storage area, the number of shelves, the degree of surface utilisation and the average amount of time it took to move a palette to or from the shelves. Jacques led a discussion about each of the different designs, pointing out their strengths and weaknesses in terms of these statistics. This allowed him to clarify the meanings of the different terms, many of which were new to the students, and helped him explain the trade-offs between maximising utilisation of space and minimising the average time per palette by contrasting different groups' designs. The TinkerBoard made it easier to ground the explanation of the optimisation problem in the students' actual designs, and by contrasting different examples, Jacques was able to point out design strategies and heuristics that students could use to better optimise their layouts.

Recall that the TinkerBoard was developed to solve a problem: students in VET classrooms were not making connections between the TinkerLamp activity and the concepts the activity was designed to teach. To determine whether the TinkerBoard had solved this problem, we designed an experiment with four classrooms taught by two different teachers. Since we have already explained the details of this experiment in Chapter 6, we will only provide a quick overview here.

Some of the students worked with the TinkerLamp+TinkerBoard while others worked only with the TinkerLamp. When we evaluated the students' conceptual understanding and their ability to solve a warehouse layout problem, we found that the students who worked with the TinkerBoard scored higher on both (see Table 6-1 in Chapter 6).

Why tell this story in the context of a chapter on VET analytics? **Because it emphasises the value and importance of data**. The data being captured and generated by the TinkerLamps were the key to unlocking their potential to help students learn. Until the introduction of the TinkerBoard, data were being collected but not being used during the classroom activity. The TinkerBoard transformed this raw, dormant data into informative, interactive visualisations that were used by teachers and students during the classroom activity.

One of the main explanations for why the TinkerBoard supported learning in this way is that the TinkerBoard 'closed the loop'. In other words, the TinkerBoard provided a feedback mechanism that both the teacher and the students could use to monitor and control the classroom activity. At various points during the classroom activity, the teacher checked the LayoutHistory panel of the TinkerBoard to view all of the groups' real-time activities and used this information to identify groups who needed special attention (Figure 9-1). The students also regularly consulted the TinkerBoard during the activity to compare their work to those of the other groups and used this information to modify and improve their own layouts (Figure 9-2). Both the teacher and the students used the information displayed on the TinkerBoard to keep the activity on track and to ensure that the TinkerLamps were being used optimally.

Figure 9-1 · The teacher consulted the TinkerBoard during the classroom activity to see the students' work simultaneously.

Figure 9-2 · Students regularly consulted the TinkerBoard during the classroom activity.

The TinkerBoard also helped the teacher connect the TinkerLamp activity to the discussion about optimisation. The TinkerBoard made it easy to use the students' layouts as examples when explaining the trade-offs between maximising utilisation of space and minimising the average time per palette. This helped ground the concepts from the discussion in the students' experiences, transforming it into a 'time for telling'.

We designed the TinkerLamp system to enable rich interactions before realising that the TinkerLamp data were actually a central part of the technology to be pedagogically exploited.

When digital tools are brought into classrooms, there are many ways that things can go wrong, and a few ways that things can go right. The added complexity can make it challenging for the teacher to monitor whether the students are using the tools as intended, and harder for students to see what their peers are doing. In the case of the TinkerLamp, the teachers and students could easily see the current warehouse layout of each group thanks to the tangible shelves, but the information contained in the projected simulation was harder to see, and there was no way to see the history of what different groups had tried. However, all of this information was contained in the TinkerLamp data. The TinkerBoard transformed this data into a public, interactive visualisation that restored everyone's ability to monitor and make sense of what was happening in the classroom, and this restored ability is what helped ensure that the activity remained on track and that students learned.

This story illustrates an important point that is especially relevant for the designers of learning technologies: Teachers and students need help managing the complexity that digital tools introduce, and the data that these tools generate is key to this. The data should not merely be treated as something to be analysed later by researchers or developers. Instead, the data should be used to show the teacher and students what is happening, what has happened and what might happen. When the data is used in this way, it acts as a feedback mechanism that teachers and students can use to ensure that they are using the tools correctly. In this chapter, we call this way of using data 'closing the loop', and we will see the importance of this idea in each of the stories.

The Story of Social Network Analysis

Realto was developed to make it easier to capture, share and manipulate experiences occurring in the different sites of the dual VET system (internship workplace and classroom). A core mechanism of Realto is the ability of teachers in the classroom and in-company trainers in the workplace to view, evaluate and give feedback on these experiences. When all is working well, these interactions allow information to flow more easily between the different sites, which makes it possible for teachers and in-company trainers to align the topics they teach (see Chapter 2 on the Erfahrraum model for more on this).

All of these interactions between stakeholders are preserved in the data collected by Realto. Every action, such as a student's request for feedback, an in-company trainer leaving feedback or someone liking or commenting on a post, is saved by Realto in its log files. We realised that this data could provide us with a way to model the entire Realto network. To do this, Mina Shirvani Boroujeni created a network analysis module within Realto (Boroujeni, 2018). This module automatically constructed graphs, where each stakeholder was represented as a node and the edges between the nodes represented different kinds of interactions. Figure 9-3 shows an example, a sub-network with nodes for apprentices, in-company trainers and teachers, and the directed edges representing communication between stakeholders.

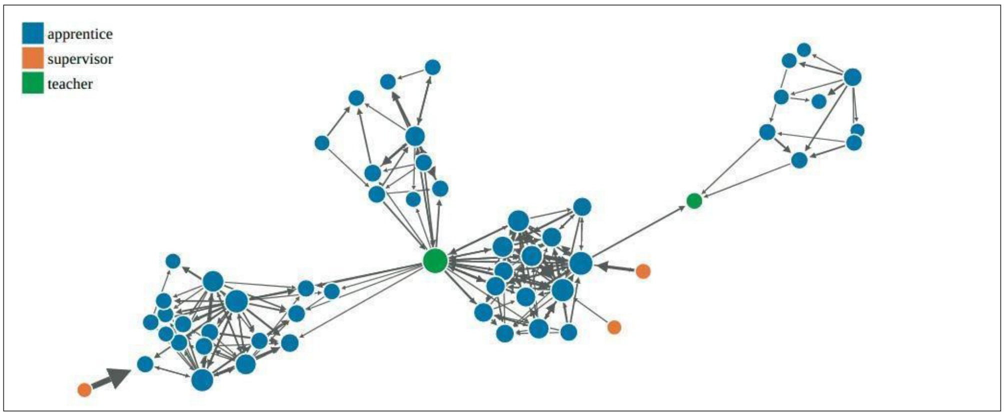

Figure 9-3 · Example of Realto sub-network of florist teachers, in-company trainers and apprentices

Was Realto actually helping to make connections between stakeholders and bridging the different sites of the VET system? To answer this question, we analysed four sub-networks containing connections between (1) apprentices, (2) teachers and apprentices, (3) in-company trainers and apprentices and (4) apprentices connected to both a teacher and an in-company trainer. We found that 80% of the apprentices using Realto had a connection with another apprentice, 47% of the apprentices were connected with a teacher, 26% of the apprentices had connected to an in-company trainer, and only 1% of the apprentices had connections with both a teacher and an in-company trainer (see the column apprentices count in Table 9-1 in the appendix).

Although these percentages suggest that many of the apprentices using Realto were not interacting with their teachers or in-company trainers, they do not capture the whole story. For example, interactions between teachers, students and in-company trainers that occurred offline were not observable in the Realto data. We were aware of many connections of this type, which were made outside of Realto and were not captured in these statistics.

However, there was a statistic that we took seriously. This was binary reciprocity. Binary reciprocity is a number between 0 and 1, which captures the proportion of two-way connections in a graph. If an apprentice posted their learning documentation and asked a teacher or an in-company trainer for feedback, that would be picked up as a one-way connection. If the teacher or the in-company trainer then provided feedback to the student, the request plus the response would be picked up as a two-way connection.

In our analysis, we found that binary reciprocity was below 0.5 in all cases (Table 9-2). This was a concerning finding: in all cases, communication between stakeholders was more likely to be one-way than two-way. This means it was more common to reach out to someone and not get a response than to get a response. This was particularly concerning in the context of learning documentation. In the Swiss VET system, apprentices are tasked with documenting their learning experiences and connecting them to a list of competencies. One of the requirements for receiving a diploma is to submit this learning documentation. Getting the learning documentation right is tricky for apprentices, and feedback from their in-company trainers and instructors is of enormous help. In fact, in-company trainers are required by federal regulations to give feedback on this learning documentation at least once a semester. Our findings about binary reciprocity on Realto suggest that apprentices are seeking out this feedback, but that instructors and in-company trainers often fail to provide it within Realto.

Figure 9-4 shows graphical representations of these communication networks. One-way connections are represented with a single-headed arrow, and two-way connections are represented with double-headed arrows. The infrequency of double-headed arrows provides a visualisation of the infrequency of reciprocal communication on Realto.

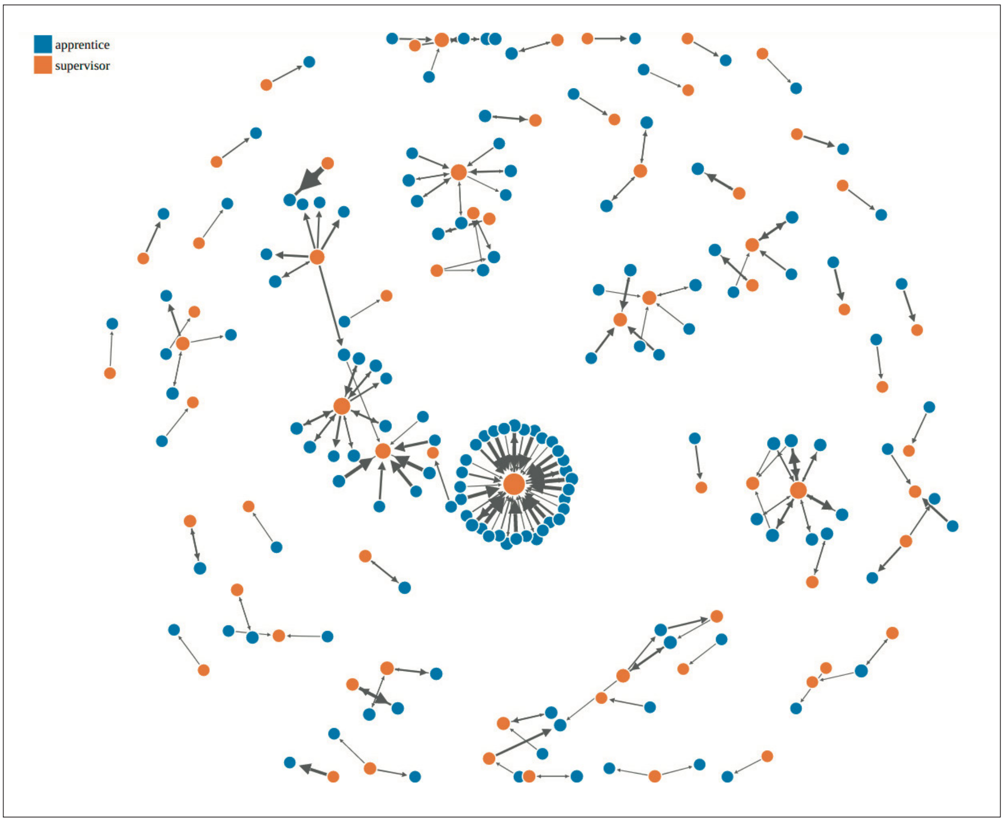

Figure 9-4 · In-company trainer-apprentice communication sub-network. The direction of the arrows indicates the direction of the interaction (e.g., from apprentice to trainer), and the boldness of the arrows indicates the numbers of interactions between two people.

This finding was upsetting, but it was not the end of the story.

A similar situation was playing out with bakers and chefs. Each of these professions has its own learning documentation platform which they use instead of Realto (bakers use a platform called LearnDoc and chefs, a platform called e-DAP). Both were created as part of the Dual-T project and predate Realto by several years. The bakers and chefs continue to use these platforms instead of Realto because they were designed specifically for their professions and because they satisfy their needs and work properly.

While neither LearnDoc nor e-DAP have a network analysis module like the one built into Realto, it was nevertheless possible to directly analyse the log files to investigate what kinds of feedback the in-company trainers of the chefs and bakers were giving. As with Realto, we found that many of the in-company trainers were giving little to no feedback to their apprentices on the platform. And as with Realto, this was concerning because soliciting and receiving feedback from in-company trainers was one of the main reasons that the e-DAP platform was developed.

A potential solution to this problem was found by Christian Gianneti, a former chef and in-company trainer and now a full-time teacher. In 2019, Christian invited the in-company trainers to a workshop where he demonstrated how e-DAP could be used to provide feedback and emphasised the added value of providing feedback on the platform.

Continued monitoring of the log files showed that two things began to change around this time. First, both the number of feedback requests from students and the number of responses from in-company trainers increased dramatically, reaching their highest levels in nearly a decade. Second, the average amount of time between an apprentice's request for feedback and the response from an in-company trainer decreased to its lowest point on record (see Figure 3-8 in Chapter 3).

An analogous effect was seen with Realto-based on the way the training workshops were organised (see Chapter 4). The amount of learning documentation filled in by the students was considerably higher and more sustained when all the stakeholders participated in the same training session in contrast to the case when training was done separately. As we saw in Chapter 3, apprentices using the e-DAP platform in 2020 reported higher feelings of connectivity between the school and the workplace than apprentices who did not use the platform. Although we cannot say for sure, the positive changes in feedback given by in-company trainers likely contributed to these apprentices' feelings of connectivity.

This story illustrates two points. The first is one we have already encountered, which is the importance of using data to close the loop. Analytics are essential if the goal is to successfully integrate digital technologies into VET. Unless there is a good reason, the data produced when stakeholders use digital technologies should be leveraged to ensure that the tools are being used optimally. The second point is about the value of VET analytics. In this account, we showed how analysing data from the school, the workplace and the intercompany courses uncovered serious problems, and pointed towards ways to address these problems once uncovered. A purely learning analytics approach would most likely have missed these problems since it would not have considered data from outside the school. The VET system is fundamentally different from other educational systems, and this story shows the importance of using a specially tailored analytics approach to better understand what is happening in all parts of the system.

The Story of Emerging Skills

Our third and final story is about the potential of VET analytics to inform and globally shape the entire VET system. The primary purpose of the VET system is to provide students with skills and expertise that will enable them to take on a skilled labour job immediately upon graduation. There are two sides to this equation. The industries that require skilled workers have job openings that they hope graduates of the VET system will be equipped to fill while the VET system trains students with specifics skills that they hope will be in demand once they graduate. Keeping this equation balanced requires both industry and the VET system to coordinate with one another. In Switzerland, this coordination occurs every five years when the heads of the VET system meet with the regional and national industry organisations. During this meeting, skillsets are formalised into training plans that will determine curricula across the VET system.

This is a top-down approach to steering the VET curriculum. Stakeholders with experience and influence make informed predictions about the future needs of industry and adjust the VET curriculum to meet those needs. This is a valid approach, one that has been working for decades, but it has shortcomings that should not be overlooked.

One of the more obvious shortcomings is that the five-year update cycle may be too slow. This is especially true today, with the AI-powered fourth industrial revolution in full swing. The dramatic increase in automation and the rise of new technologies has rapidly transformed the skillsets required for a wide variety of jobs and industries. Highly valued skills have fallen out of favour while new skills have risen to prominence. These changes can occur rapidly, creating a situation where students in the VET system are stuck learning outdated skills or fail to learn skills that are now in high demand. One solution to this problem would be to increase the frequency with which the VET curriculum is updated, but there is still a problem that this approach would not solve.

That problem is predicting the future. The industrial landscape is complex, and it is not so easy to predict which skills will be valuable in years to come. For example, few would have predicted that bankers would need to understand peer-to-peer networking, cryptographic hashing or public ledgers, but these are precisely the concepts involved in something which has profoundly altered the landscape of banking over the past decade: Bitcoin and other forms of cryptocurrency. This is the kind of change that few people would have seen coming, regardless of their level of expertise or influence.

One possible way around these shortcomings is to introduce a bottom-up approach into the process. This kind of approach would also provide the VET system with information about future skills, but this information would be directly derived from the labour market. Ramtin Yazdanian explored the feasibility of this approach in his PhD (Yazdanian, 2021). His approach was to train predictive models on labour market data to predict emerging skills. There is a lot to unpack in that sentence, so let us take it term by term.

First, what is a predictive model? Predictive modelling is a method that uses statistical models trained on historical data to predict future outcomes. A simple example of a predictive model is a line fit to data. Maybe you want to try and predict what your favourite basketball team's final score will be while you are sitting through commercials at halftime. You could plot the score on the y-axis and the time on the x-axis, and then fit a line to the data. Extend the line out to the end of the fourth quarter, and it will sit directly on a prediction of the final score. While this method is simple, it is not usually accurate. There are more sophisticated methods that can produce better predictions, but these often require large amounts of data to be effective.

Ramtin was able to use these more sophisticated methods because he found an enormous, and largely untapped, source of labour market data: job ads. For the past 10-15 years, companies have been posting job advertisements on websites such as monster.com and linkedin.com. These sites possess huge archives of historical job ad data, and these archives contain information on the dynamics of skills in the labour market. Ramtin's hypothesis was that predictive models trained to learn these dynamics should be able to predict the rise and fall of skills in the labour market.

However, simply predicting the rise and fall of skills was not the main outcome of interest. Recall that the goal of this work was to predict which skills would be in demand when apprentices were leaving the VET system and entering the labour market. Ramtin called these emerging skills and defined them as 'previously low-demand skills that have recently experienced a surge in hiring demand'. Successfully predicting emerging skills from historical data is exceptionally challenging because they are largely indistinguishable from other low-demand skills that never emerge.

As a proof-of-concept, Ramtin developed a predictive model to identify emerging skills in the information technology sector. He began with a dataset of job ads, where each ad was represented by the list of skills it contained. He used this data to create a time series of each skill's popularity in the labour market and extracted hundreds of features from these time series. These features include summary statistics (e.g. mean, various quantiles, variance), linear trends, measures of non-linearity and spikes and FFT coefficients. The most informative features were identified using a combination of methods and brought together into a single logistic regression model. (The full technical details of this model can be found in Yazdanian et al., 2022.)

Ramtin found that predicting emerging skills was, in fact, possible. His best model was able to consistently beat a number of strong baselines, showing that job ad data contain enough information to predict many of the skills that will emerge in the future.

Although this had worked on data from the IT sector, it was not obvious that it would also work in VET domains. The IT sector moves at a breakneck pace, with new programming languages, technologies, frameworks, and platforms being introduced every day. Needless to say, a VET domain such as masonry does not experience such rapid change.

So what happened when Ramtin applied his method to job ad data from VET industries? The results were not as clear. He used his method on job ads from two VET sectors: logistics and healthcare. In some cases, his model was able to produce useful predictions, but in other cases, his model failed. One of the most likely reasons is that this method, which was developed on data from the IT labour market, is ill-suited to the VET labour market. We saw a similar issue in the story about social network analysis. There, Mina's methods, which she developed using MOOC data, did not work when applied to data from the VET domain. In her case, it required developing new methods that were specifically tailored to VET. A similar approach is probably needed in this case as well. Rather than dismiss Ramtin's results as a failure, we see them as reminding us, yet again, that VET is a unique domain that requires its own special methods and approaches. With persistence and luck, we (or someone else) will find a way to use job ad data from the VET labour market to help predict the skills that students will need once they graduate.

What is the moral of this story? There are two. First, this story illustrates why VET analytics must have a wider scope than learning analytics, which is solely concerned with optimising student learning. Optimising student learning is important. However, in the VET system it is also important to make sure that the skills that students are learning will be in demand when they graduate. In other words, the curriculum also needs to be optimised. A learning analytics approach would optimise student learning while not concerning itself with the content in the curriculum. Over time, this approach would produce apprentices with excellent, but obsolete, skills. The wider scope of VET analytics avoids this by respecting the complexity and interconnectedness of the VET system.

The second moral of this story is one that you may be sick of by now. This is the importance of using data to close the loop and ensure that things are working as intended. In this account, the loop is at the most zoomed out, macro level, and the technology in question is the entire VET system. If we are ever able to use the data generated by the VET system (labour market data) to help make predictions about which skills are likely to become important in the future, this will close one of the many loops in the VET system. This, in turn, will help make VET more agile, ensuring that apprentices graduating from the system have more relevant skills.

So what?

If there are only three things that you take away from this chapter, let it be these:

First takeaway: VET analytics is a novel approach that embraces the complexity of the VET system. The Story of Emerging Skills shows why a learning analytics approach that focuses on student learning while ignoring other aspects of the VET system is inappropriate. The different parts of the VET system have deep connections that tie them together. The things that students learn in the classroom are determined by a curriculum that was designed to meet the future needs of companies. Optimising for student learning without also optimising the curriculum makes no sense within the VET context. VET analytics is built on an understanding of the complexity and interconnectedness of the VET system, and recognises the need to consider multiple stakeholders, contexts and outcomes.

The Story of Social Network Analysis also illustrates this idea. To see why, we must provide a bit more background. Mina pioneered her methods on data from a MOOC (Boroujeni, 2017). However, she was unable to directly apply these methods to the Realto data due to differences between the structure of MOOCs and that of the VET system. She needed to modify her methods to account for the differences in stakeholders, forms of social interaction and communication patterns. Again, this illustrates the unsuitability of existing analytics approaches and the need for a new approach that understands and embraces the complexity of VET.

The TinkerBoard story showed that methods from learning analytics still have a central place in VET analytics. This story showed how introducing a data visualisation dashboard into the classroom helped to ensure that the TinkerLamp was having a positive impact on student learning. This was a straightforward, effective use of learning analytics in a VET classroom. This story is a reminder that VET analytics has all the tools of learning analytics at its disposal. The other stories show us why it is important to go beyond these tools if VET analytics is to effectively optimise the entire VET system.

Second takeaway: The data is an essential part of the technology, not a side effect. The second takeaway is '*Pay attention to the data!*' We are aiming this mainly at the designers of learning technologies. The stories in this chapter show that ignoring produced data fails to tap in the potential of learning technologies. In the case of the TinkerLamp, hiding the interaction data from users was one of the main reasons why students were not learning. And in the case of Realto, only once Mina dug into the data did it become clear that the problem of in-company trainers and instructors failing to give feedback was more widespread and serious than we had realised.

Our takeaway from these stories is that the data being generated and logged by learning technologies is not a side effect but an essential part of the technology. Designers should recognise this and ensure that users can easily access and understand the data. Providing easy ways of accessing the data, filtering it and visualising will likely help ensure that the technology is being used effectively.

We applied this principle to the entire VET system in the Story of Emerging Skills. In that story, the users were the administrators and industrial organisations who decide on the VET curriculum every five years, and the data they were not using was job ad data containing information about skills in the labour market. Our research showed that this data contained information that could be used to predict emerging skills, but we have yet to provide the users with a way of using this information in their decision-making. If the other stories are any indication, once this data is brought into the decision-making process, it will help ensure that the 'technology' of the entire VET system works as intended.

Third takeaway: To improve the VET system, VET analytics must 'close the loop'. Our final takeaway is about what should be done with the results of VET analytics. Any results or insights should be used to 'close the loop'. In other words, they should be used as a source of feedback that can help the VET system improve itself. If the results of these analyses only end up in reports or papers, they will have minimal impact on the VET system. These insights need to be fed back into the system that generated the data so the system can keep itself on track.

In the TinkerLamp story, information extracted from the interaction data was visualised on the TinkerBoard dashboard. Instructors and students alike regularly consulted the dashboard to better understand what was happening in the classroom and instructors used it to connect the TinkerLamp activity to the concepts in the lesson. Without the TinkerBoard, it was more difficult for the instructors to identify students who had gone off track or for students to monitor their own activities in relation to their peers. The TinkerBoard provided a source of feedback: Only when it was introduced into the classroom did the TinkerLamp activity result in significant learning gains.

Closing the loop does not need to be a high-tech exercise. In the case of Realto, the information that instructors and in-company trainers were not providing enough feedback to students was used to close the loop in the form of a workshop. Christian showed teachers and in-company trainers how to see requests for feedback and how to respond to that feedback and explained the importance of providing feedback through the e-DAP system. Following the workshop, there was a notable increase in feedback given and a drop in the amount of time it took an in-company trainer or teacher to provide feedback. An analogous effect was present when manipulating the setting of the workshops (see Chapter 4).

We have yet to learn what will happen when we use the data on emerging skills to close the loop on the VET curriculum. We hope that it will help to make the VET system more agile, make the VET curriculum more relevant and help equip students with the skills and knowledge that will prepare them for the future of industry.

Appendix

Classrooms in Bulle		**Classrooms in Yverdon**	
No of Students	Condition	No of Students	Condition
15	No TinkerBoard	15	No TinkerBoard
17	TinkerBoard	16	TinkerBoard

Table 9-1 · The setup for the classroom study designed to evaluate the TinkerBoard

Sub-network type	**Apprentices count (%)**	**Teachers count (%)**	**Trainers count (%)**	**Binary reciprocity**
Apprentices	458 (80%)	NA	NA	0.44
Teachers to Apprentices	266 (47%)	54 (52%)	NA	0.21
In-company trainers to Apprentices	147 (26%)	NA	68 (86%)	0.49
Teachers to In-company trainers to Apprentices	53 (1%)	22 (21%)	15 (19%)	0.28

Table 9-2 · User distribution and reciprocity of Realto sub-networks