Journal of Universal Computer Science, vol. 25, no. 12 (2019), 1507-1530
submitted: 11/4/19, accepted: 31/10/19, appeared: 1/12/19 © J.UCS

Improving Teacher Game Learning Analytics Dashboards through ad-hoc Development

Antonio Calvo-Morata
(Complutense University of Madrid, Madrid, Spain
acmorata@ucm.es)

Cristina Alonso-Fernández
(Complutense University of Madrid, Madrid, Spain
calonsofernandez@ucm.es)

Iván J. Pérez-Colado
(Complutense University of Madrid, Madrid, Spain
ivanjper@ucm.es)

Manuel Freire
(Complutense University of Madrid, Madrid, Spain
manuel.freire@fdi.ucm.es)

Iván Martínez-Ortiz
(Complutense University of Madrid, Madrid, Spain
imartinez@fdi.ucm.es)

Baltasar Fernández-Manjón
(Complutense University of Madrid, Madrid, Spain
balta@fdi.ucm.es)

Abstract: Using games for education can increase the motivation and engagement of students and provide a more authentic learning environment where students can learn, test and apply new knowledge. However, the actual (serious) game application in schools is still limited, partly because teachers consider their use as a complex process. To increase game adoption, the integration of Game Learning Analytics (GLA) can provide teachers a thorough insight into the knowledge acquired by their students and usually presented through a visual dashboard. Although it is possible to provide a useful general metrics and a prefab dashboard, it may not fully cover teachers' expectations. In this paper, we study the ad-hoc adaptation of generic dashboards to increase their effectiveness through three case-studies. In these experiences, we adapt dashboards for teachers to include detailed information for more-focused analysis. With the positive results obtained from these scenarios, we have identified a methodological process to create ad-hoc GLA dashboards and extracted some lessons learned for dashboard development: simple but useful dashboards can provide a higher added value for stakeholders compared with more complex dashboards; teachers and game developers should be involved in dashboard design for better results; and, if possible, ad-hoc developed dashboards should be used as they have proved to be more effective than generic dashboards.

Keywords: Learning Analytics, Serious Games, Dashboards, xAPI, Game-Based Learning
Categories: K.3, L.5.1

1 Introduction

Games characteristics such as their engaging and motivating nature make games especially adequate for education [Connolly, Boyle, MacArthur, Hainey and Boyle 2012, Denis and Jouvelot 2005]. Despite these probed advantages, the adoption of serious games is still poor in formal education, partly due to a lack of standards for development, validation, and deployment in schools [Loh, Sheng and Ifenthaler 2015]. We consider that collecting and analyzing student gameplay is one of the keys to increase serious game adoption in schools. Collected analytics data can provide insight and improve all steps in the process, and the resulting insights can themselves constitute an important selling point supporting the use of games in the classrooms. In educational settings, Learning Analytics (LA) is used to provide insight into learners' actions to improve some aspects of the learning process and contexts. When applied to games, Game Learning Analytics (GLA) focuses on information gathered from players via in-game interactions.

Fig. 1 represents a generic Game Learning Analytics pipeline, focusing on two main stakeholders: students playing a serious game, whose information is tracked, stored and analyzed within the Analytics System; and teachers supervising the game session. Analyses and visualizations (embedded in a dashboard) provide information for teachers. Other analytics stakeholders, such as students themselves, game developers, and academic officials, can be presented with their own dashboards.

The analysis of data has not only risen dramatically in the field of games [Bauckhage, Drachen and Thurau 2015]: it is a key trend across many fields and domains, characterized by collecting and analyzing large amounts of data and, therefore, sometimes referred to as *big data*. For instance, business analytics has been identified as one of the key technology trends in the 2010s [Chen, Chiang and Storey 2012]. Following the analysis of data, the results are typically presented visually, for different purposes and target audiences. The visual analytics field also has a wide range

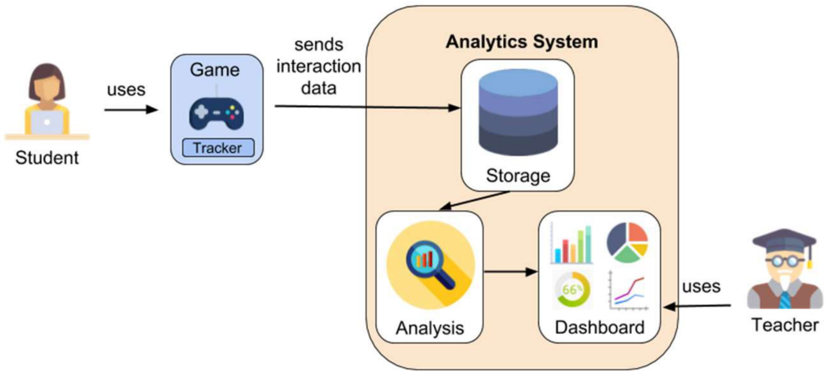

Figure 1: Game Learning Analytics pipeline: a student plays a game that includes a tracker component. The tracker sends interaction data to the Analytics System for storage, analysis and visualization. Teachers, among other stakeholders, interact with dashboards to gain insights into student gameplay and educational outcomes.

of possible application areas, including information synthesis, the discovery of unexpected results, and performing and communicating assessments in calls for action [Keim et al. 2008]. It is, therefore, no exception that, also for serious games, dashboards are the usual way to communicate information to stakeholders [Shoukry, Göbel and Steinmetz 2014]. They can display important metrics and provide a visual overview of other information while allowing filtering and limited query capabilities to gather more in-depth data. Besides, data visualization via dashboards is the most commonly used method in Game Analytics [Drachen, El-Nasr and Canossa 2013] to communicate insights and provide reusable reports for different stakeholders: visualizations can provide in a simple way a variety of information, such as heatmaps to display places where level crashes occurred.

It may appear that once a serious game implements tracking and event-reporting standards, it should be simple just to add Game Learning Analytics (GLA) to provide pedagogical insights that teachers can use in their classrooms. However, even though generic visualizations (and dashboards) provide a means to obtain information from commonly-present traces generated from GLA data, establishing the relationship between the visualization (or the event) and the consequence (e.g., the knowledge acquired) is not an easy task. When designing these generic visualizations and dashboards, pedagogy experts must help to determine how the Learning Goals (LGs) of a game can be identified in such simple visualizations built from standardized traces. They must participate in the creation of the Learning Analytics Model (LAM) [I. Perez-Colado, Alonso-Fernandez, Freire, Martinez-Ortiz and Fernandez-Manjon 2018] for the generic dashboard, as their knowledge helps to identify the suitable LGs that every game must-have for teachers to be able to evaluate and assess later.

These generic visualizations present relevant data such as times and scores, which can be compared against baselines to determine student performance. In a class, teachers can see whether whole tasks are working as intended or diagnose problems. How to display this information is far from standardized: a recent literature review on analytics in serious games [Alonso-Fernández, Calvo-Morata, Freire, Martínez-Ortiz and Fernández-Manjón 2019] found no consensus for the display and visualization of performance metrics for serious games. We have the same experience as we have found no standard patterns for this process, relying instead on different ad-hoc methods to visualize specific outcomes and general progress.

While generic dashboards might be enough to satisfy the basics, ad-hoc, game-specific dashboards allow a focused and higher level of introspection and understanding of the gameplay and its conclusions, as they are developed for a specific game and its stakeholders.

The stakeholders in GLA include students, teachers, game developers, administrators, and researchers, among others. This paper focuses on teachers, as they oversee the actual educational environments and are, in our opinion, the first stakeholder that needs to be considered to improve analytics for games in education: student dashboards are also important, but it is teachers that decide if games will be used in their classrooms. More narrowly, the last step of *Fig. 1* depicts the teacher using the dashboard, which requires teacher dashboards to be understandable, as described in [Schwendimann et al. 2017].

Generating dashboards starts with tracking interaction data, which must be done with care to guarantee privacy. Interaction data can then be displayed in different visualizations, for example displaying previously-identified KPIs (Key Performance Indicators), the choice of which will be different from other stakeholders. For example, school administrators will be interested in comparing summaries of the performance of each class; while teachers will want insights into individual students as compared to their classmates. Data collection and analysis should be as transparent as possible. The final step, the visualization of the information, is where game learning analytics can finally provide value. In this sense, when used by teachers, dashboards need to be evaluated as pedagogical tools, taking into account their goals, affect and motivation, and usability [Jivet, Scheffel, Drachsler and Specht 2018].

The process of collecting, analyzing, and displaying data from in-game interactions to yield useful teacher dashboards comprises several steps, each of them beset by possible issues:

a. Data collection: Collected interaction data cannot be easily shared unless a collection standard is being followed. Once standardized, privacy issues need to be addressed. Furthermore, what data should be captured depends on the games, and game developers are understandably more interested in designing games – rather than selecting what to send and then, on top, having to perform anonymization and sending it according to specific standards. While data collection is not an issue that is specific to teacher dashboards, decisions made at this step (particularly what is collected and how it is anonymized) greatly influence dashboard outcomes.
b. Low teacher expectations: Teachers are often new to analytics dashboards, and do not know what to expect. In our experience, when asked what they expect to see, teachers described only basic information, such as times of completion, difficulty, results in terms of counts of right and wrong answers, or the number of attempts, that if possible, should be displayed using simple visualizations. Also, teachers assume that analytics will only be available after the intervention, and do not expect to receive any information while students are playing.
c. Dashboard design: The design space of possible dashboards is vast, and designing useful visualizations requires both pedagogical knowledge and game-specific information. Teachers are generally not experts in dashboard design (see point *b* above); and are unwilling to make significant investments in dashboard design upfront, before the game is even available.
d. Changing dashboard requirements: Teachers will often request additional visualizations for their dashboards after the game has been played (see points *b* and *c*). Fulfilling these may be costly or even impossible (for example, if the requisite data was not originally collected; see point *a*) – unless the whole system has been designed to allow the necessary flexibility.
e. Beyond stand-alone games: Teachers may want to use games as parts of larger courses, which may, in turn, be games as well. In these cases, dashboard granularity needs to be configurable, allowing the game to be analyzed not only by itself but also as a part of a (generally hierarchic) set of activities.

Nevertheless, a recent literature review [Jivet, Scheffel, Drachsler and Specht 2017] regarding the use of learning analytics dashboards in education points out several flaws and improvements for them, including that dashboards usually lack a secondary goal beyond awareness and reflection, and that focusing on each student's needs instead of using the common technique of comparing each student with their peers may be a better pedagogical technique to promote motivation. Although default out-of-the-box visualizations are helpful, the real powerful insight is obtained when creating custom visualizations [Drachen et al. 2013].

Section 2 of this paper describes how our architecture, and specifically our teacher dashboards, tackle the above issues; and details how dashboards can be designed working with designers and teachers. Section 3 describes three specific use cases where we have developed teacher dashboards. Section 4 describes the lessons learned in teacher dashboard design. Finally, Section 5 summarizes our conclusions and Section 6 outlines future work.

2 Addressing the teacher dashboards issues

As mentioned before, the process of collecting, analyzing, and displaying data from game interactions on dashboards useful to the teacher is a complex task. Next, we will talk about the main decisions at the level of the analytical system, data collection and the process to create the dashboards that will show the analyzed data.

For data collection (issue *a* in the previous section), we use the Experience API for Serious Games Profile (xAPI-SG) as a standard collection and archival format. xAPI-SG defines a set of interactions that are usual in serious games, as detailed in [Serrano-Laguna et al. 2017]. We provide an easy to use library that greatly simplifies adding analytics to serious games, isolating game developers from the details of the standard. To avoid privacy issues, we use pseudonymous tokens for students. These tokens are unique strings of 4 characters created at the server when the game is deployed and provided by teachers to their students, who will then use them to access the game. We store no backup of the token-to-student correspondence table. We then rely exclusively on the tokens to identify students across play sessions, using them also to display information in the visualizations. Only teachers can, if they choose, keep the correspondence between tokens and actual students, which they can use to trace the identity of students in the dashboard.

By using pseudonymous tokens, we guarantee that if data is stolen or lost, only the owners of the pseudonym-to-name conversion table will be able to identify the students who played. Not only that, but we open the possibility of sharing the pseudonymous data with game-developers or researchers, which can benefit from real gameplay data to improve the lifecycle of games by identifying issues and interesting patterns. Given the increased awareness of privacy risks (as reflected in the EU's GDPR), the use of pseudonyms is not the only precaution that has been taken. Our databases are secured to only allow internal connections, and APIs are secured through API gateways. Additionally, if students want their data to be deleted, we allow teachers to permanently delete data for specific students (again, identified via a pseudonymous token) from the system to guarantee their right to be forgotten.

When conducting experiments in schools, we require the school to sign a consent form that describes the type of data collection to be performed and acknowledging that the data is collected under pseudonyms that we will not be able to track back to specific students. Additionally, at the beginning of each gameplay session within an experiment, students are informed that their gameplay data is being collected for research purposes and can opt-out from this data collection.

From our work in two EU H2020 Projects, we have developed a complete architecture to track, collect, store, analyze and display the data collected from serious games in a systematized way [Alonso-Fernandez, Calvo, Freire, Martinez-Ortiz and Fernandez-Manjon 2017, I. Perez-Colado et al. 2018]. A simplification of this architecture is depicted in *Fig. 2*. To understand the workflow, first, the student starts to play the game. Inside the game, there is a component (a tracker) that handles student authentication with the analytics server, and which continuously sends information about what is happening inside the game. To ease integration of the tracker into games, it is available for multiple frameworks and programming languages, including Unity, JavaScript, and Java; and implements the xAPI-SG standard. With a single line of code, developers can link an event inside the game with a trace that will be sent to the analytics backend. Once the analytics backend receives the trace, it is enqueued in Apache Kafka to be consumed by an Apache Storm streaming analysis. This analysis performs the needed calculations and stores both the results and the raw traces into ElasticSearch, which provides storage and querying, for later use by Kibana, which can generate visualizations from ElasticSearch. Teachers can then visualize real-time results through the Analytics Frontend, which builds dashboards from Kibana visualizations, to assess their students.

The tracker communicates game events to the backend by sending traces in xAPI format, which are then categorized and further analyzed. The generic analysis performs simple aggregations, such as counting correct/incorrect answers or calculating the highest score obtained in a specific part of the game. However, some analyses can be more complex, for example by pre-calculating averages or time-clustering the results. An example of the possibilities of an analysis is the Multi-Level Analysis [I. J. Perez-Colado, Rotaru, Freire-Moran, Martinez-Ortiz and Fernandez-Manjon 2018], in which a hierarchy of analyses is provided, and traces are filtered, modified, and re-enqueued into the system to provide, for example, the progress and completion status of each

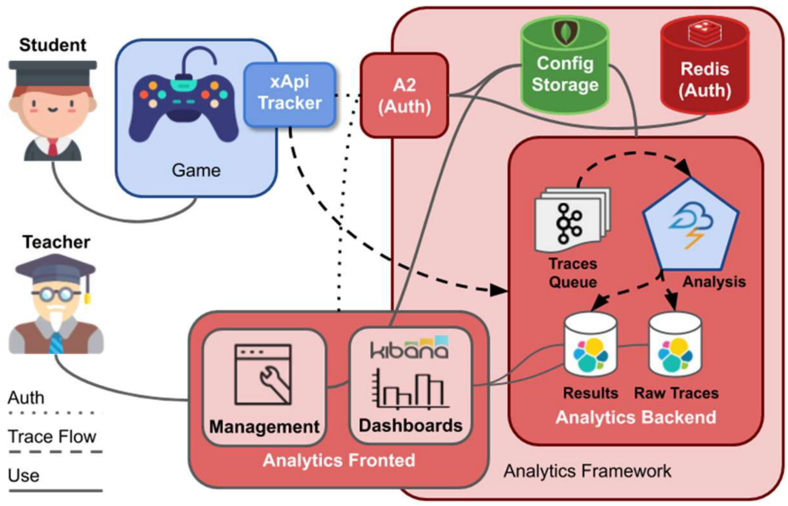

Figure 2: Game Learning Analytics pipeline for our analytics framework. A more specific explanation of what is described in Fig. 1. including technologies used.

node in an activity tree, where nodes can be games that include other games inside (see issue *e* in the previous section).

Regarding low teacher expectations and dashboard design (issues *b* and *c* in the previous section), since we cannot expect teachers to provide detailed lists of what should be analyzed and how it should be displayed, we have developed a default, generic dashboard, which does not require any setup and can display basic data for any game that sends valid xAPI-SG interaction data. For example, since xAPI-SG has a specific vocabulary to indicate that a student has made a choice, and whether the game considers the choice to be correct or not, the default dashboard can easily display counts of correct/incorrect student answers. The use of a default, generic dashboard immediately provides value to teachers and provides a useful base to elicit requirements for more complex game-specific dashboards. On the other hand, game-specific dashboards can provide much better insights, since, for example, keeping track of a game's specific goals in a dashboard where each goal is represented by a specific visualization is much more intuitive than having to rely on one-size-fits-all generic dashboards.

Additionally, since our dashboards are updated in near real-time, with delays of few seconds between receiving interaction data (in xAPI-SG format) and displaying updated visualizations, we also include a simple alerts-and-warnings mechanism that can be configured to notify teachers of possible issues as they arise. We consider alerts to be higher-priority than warnings, but the underlying mechanism is the same; and, asides from increasing the situational awareness of teachers, its existence reminds teachers that the use of analytics is not limited to presenting post-mortem information on playthroughs; and that their role during gameplay sessions need not be limited to proctoring.

To provide the necessary dashboard flexibility (issue *d* above), we are not limited to the default dashboard, and allow game developers (presumably with teacher feedback) to create customized game-specific analyses and visualizations for their games. Since our dashboards are built on top of the Kibana and ElasticSearch open-source projects, dashboard creation is developer-friendly, although not recommended for non-programmers. If need be, custom analyses and dashboards can re-evaluate old data, allowing dashboards to be updated to display existing information in new ways. This architecture allows for requirements to evolve as teachers and game designers refine their understanding of how students play and learn with a serious game, or when dashboard usability issues are identified. Subsections 3.1 and 3.2 contain two case-studies of such custom game-specific dashboards.

Finally, regarding issue *e*, games are sometimes part of more complex course structures. For example, a game may contain several mini-games, each of which can merit its specific dashboards. However, it still makes sense to provide a global dashboard to monitor progress across all the minigames, and to be able to switch granularity on demand. Indeed, this functionality was a key piece of the H2020 BEACONING EU Project. The resulting multi-level dashboards provide insights on player progress at any level of the activity hierarchy, which could encompass both serious games and more traditional classroom activities.

Building dashboard visualizations for teachers requires collaboration between both game designers/developers and teachers. Developers are closest to the game's requirements and characteristics and in charge of designing the interactions that will yield observable data and report it to the analytics server. Subsection 2.1 describes this process in greater detail. On the other hand, their efforts will be in vain if the resulting dashboards are not useful to teachers. Subsection 2.2 describes teacher involvement in the teacher dashboard creation process. Note that communication between teachers and developers should occur as early as possible, since the design of the game and its choice of what to report when determines much of what can be achieved with the subsequent analysis and presentation in the form of dashboards. Several iterations may be needed before a good teacher dashboard is created.

In addition to these two roles, pedagogical experts are necessary to ensure that the game is pedagogically sound, and the game's analytics dashboard is an important tool to ensure this soundness. Pedagogical experts should be involved in the design of the game, and part of this design includes coming up with good ways to measure and assess learning – which should be embodied in the dashboard. After the game is finalized, pedagogical experts can be valuable to help teachers to build dashboards that "bring up the learning", by emphasizing the links between in-game actions and actual pedagogical outcomes.

2.1 Working with game designers and developers to create dashboards

Game designers and developers choose the game's mechanics and contents and are in the best place to map this knowledge of its internal workings into meaningful interactions to be reported to the analytics server. Therefore, their involvement is essential when designing a dashboard's individual visualizations, to ensure that the specific data that they rely on is correctly collected, analyzed and presented. This process can be divided into the following steps:

1. Dashboard goal, and choice of general layout and contents: The first step is to define a draft of the desired dashboard including the visualizations that we would like to show. This includes the general layout and contents of visualizations, encompassing both the information to be shown and the desired visual mapping. For instance, possible visualization contents could be "total player activity over time, as a line chart", "number of correct and incorrect answers per player, using stacked bars", or "average score per player and players who have completed the game, as a sorted scatterplot". Teacher involvement (see steps 2 and 3 in subsection 2.2) is important here.
2. Choose the data to fill the dashboards: Once the information to be shown in the dashboard is clearly defined, the next step is to define the specific variables that will be used to report this information. These report variables (encoded into different sections of xAPI statements) will be fed from data from the game. This step encompasses both the choice of xAPI encoding and that of choosing the game data from which it will be assembled. For instance, to show the number of correct and incorrect answers per player, the game would have to send, for each answer and player, whether the response is correct or not. To isolate game developers from as many xAPI details as possible, we provide easy-to-use libraries in several programming languages. Note that developers can choose to send additional analytics information, even if it does not make it into the final dashboards: this provides flexibility until dashboard design is finalized since it allows different dashboard ideas to be tested out with existing, previously-unused data. Indeed, they are encouraged to at least report

on all major xAPI-SG events; this is enough to feed the default analysis and dashboard (described above, and addressing issues *b* and *c*), and provides a good baseline on top of which game-specific improvements can be made.

3. Define how the data reported from the game will be analyzed. Building on the previous example, to show the number of incorrect and correct answers per player, given xAPI statements from different players stating whether each of their answers was correct or not, the analysis would have to aggregate results for each player and store it for later display.
4. Create the dashboard: following the previous steps, a new version of the dashboard can be created, which can be tested with actual game data, either by playing it anew or by replaying previously-collected data. Continuing with the previous example, we could choose to depict the number of correct and incorrect answers per player with a stacked bar chart where each bar corresponds to a player, and contains two smaller bars on top of the other, representing correct and incorrect answers, and encoded as green and red. The height of each bar would reflect answer counts. The resulting dashboard would then be evaluated by teachers, and possibly refined – see step 5 in the next section.

2.2 Working with teachers to create dashboards

Early teacher involvement is critical to the success of a teacher dashboard. While game designers or researchers may think that a given dashboard prototype is highly informative and useful, it is easy to forget that most of the teachers that will use the dashboards will not be experts in the game, and will not be familiarized with high-density visualizations favored by visualization experts; involving some non-designer teachers early during dashboard design helps it to be much more usable.

Beyond usability, teacher involvement allows dashboards to support teachers in the activities that surround their specific use of the game in the classroom. The process has the following steps:

1. Show the game to the teacher: Teachers should play and understand the game, its goals, how it is played, and the relationship between game mechanics and learning outcomes. Early teacher input will allow visualizations to be focused on their intended use of the game, making it more likely that the resulting dashboards will be both usable and informative. Without playing or understating the game, teachers cannot build an opinion on what they would like to see.
2. Define teacher goals: Once teachers have played, we can ask them to describe what they want to learn from how their students play. For example, teachers may be interested in knowing what parts of the game students are playing, whether they have problems in certain places or with certain questions, or the order in which they are choosing to solve the game's challenges.
3. Create a draft dashboard: given a set of goals, we can draft a possible dashboard that attempts to satisfy them. Each dashboard visualization will require different analyses, and data to feed them, which is the focus of the next step.
4. Ensure availability of underlying data: For each visualization, a suitable analysis must be available to produce the data which will be displayed, and

this analysis, in turn, requires the game to send the necessary interaction data to feed it.

5. Evaluate the dashboard with teachers: The draft dashboard can now be implemented and shown to teachers to evaluate its usability and usefulness. Based on teacher feedback, visualizations may be added or removed from the dashboard, prompting a new prototype/provision/develop/test cycle.

Note that the dashboard development processes for game designers and developers, on the one hand, and for teachers, on the other, are interrelated: steps 2 and 3 above feed into designer/developer step 1 (choice of dashboard layout and contents); and step 5 follows designer/developer step 4 (dashboard implementation). Furthermore, step 4 above can only be fulfilled with developer collaboration, since only developers will be aware of the exact data being sent, and how it is being analyzed. Finally, steps 2-5 are meant to be iterated. See *Fig. 3* for an illustration.

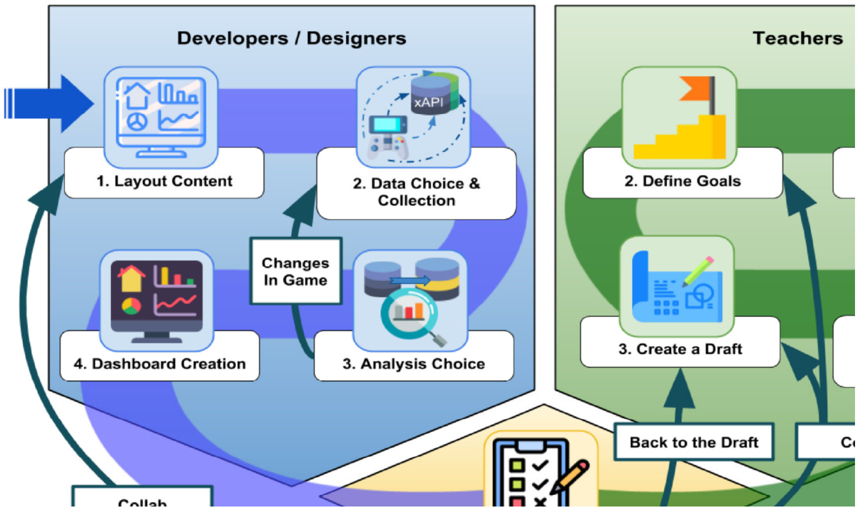

Figure 3: Dashboard creation process.

3 Use cases

We have used teacher dashboards and game learning analytics in multiple experiments, with goals that included the validation of the underlying analytics architecture, improving the serious games that we were analyzing, and evaluating game use and performance. In this section, we summarize three experiences in the creation and use of game learning analytics and teacher dashboards.

Note that all the dashboards are designed to be visualized in a web-browser in a traditional monitor, with a mouse used for interaction. Using the dashboards on a mobile device may result in poor performance since the frontend relies on CPU-intensive Kibana to display the real-time visualizations.

3.1 A custom dashboard for a serious game on cyberbullying awareness

Conectado [Calvo-Morata et al. 2018] is a serious game to raise awareness against cyberbullying by placing players in the role of a transfer student that suffers bullying and cyberbullying after arriving at a new school. Players experience the life of this student in first person, during each of 5 in-game days, while being exposed to feelings of helplessness and increased (in-game) social isolation. The game keeps track of the level of the friendship of the story's protagonist with each classmate. For instance, the variable called *friendship risk* indicates, based on the player's choices, the risk of being bullied in the game. The risk is encoded as an integer between 0 and 100, where higher values correspond to higher risk, that is, worse social standing. There are indicators of risk for each character as well. The decisions that players make during the game, including whether players decide to ask for help to being bullied to the parents or the teacher or not, also determine the ending.

Some of the game-dependent visualizations developed for *Conectado* can be seen in *Figure 4*. From left to right, and from top to bottom:

a. Average friendship risk: this general metric describes whether the average of the class has low, medium or high friendship risk (shown in green, yellow and red, respectively). In general, the lower the risk, the better, since this means that the students are keeping good relations with their in-game classmates by choosing good dialogue options.
b. The number of players per game day: this bar chart provides a vision of the progress of players in the game. As there are five days, teachers can see in real-

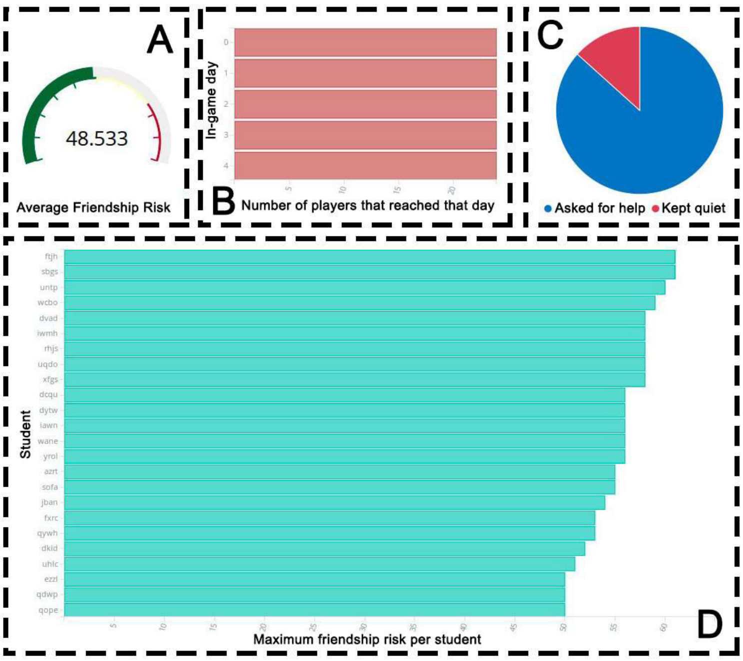

Figure 4: Some of the game-dependent visualizations of the game Conectado.

time how many students have played through each in-game day and help students who are too far behind their classmates.

c. The number of players that have taken each possible action that determines the ending: this pie chart compares the number of players who have decided to complain about bullying (have asked for help) to the in-game parents and teachers vs. those that have decided to remain silent. Note that communicating that the harassment is happening is the first step to solve it .
d. Maximum friendship risk per student: this visualization provides a more detailed view that complements the general metric provided in visualization (a) and allows teachers to quickly identify which students are doing best and worst in the game. Since the visualization is sorted, it also provides an overview of the distribution of risk scores throughout the class. This visualization is key for a quick response of the teacher in order to help the proper students, as students with higher risk need more help.

These visualizations have been designed trying to cover some of the information usually required by teachers: progress (b), decisions taken (c) and specific metrics, both general (a) and per student (d).

By default, our dashboards display individual visualizations in two rows, with each visualization of the same shape and size. As seen in *Figure 4*, this can be modified to realize almost any design, by providing control over size, shape, and location of individual visualizations.

In the *Conectado* case, and as a part of our experiments in schools, we also asked teachers for feedback on the utility and usability of their dashboards while the videogame was being played in their classrooms. Teacher feedback resulted in a more complex dashboard, found in *Fig. 5*. This dashboard shows, in real-time, what players are doing, the percentage that has completed each phase, and their answers in key points in the story:

e. The last trace sent: displays, in a list, for each player, the time at which that player last interacted with the game; together with the player's general progress in the game. This allows teachers to quickly see which users are inactive or are further behind, in order to take action and notify inactive users that they must play.
f. Games started vs games completed: allows teachers to assess the general progress of all players. If all players have finished, then the pie chart will display two equally-large portions, one purple (games started) and another green (games completed). The ratio between both portions reflects the difference between these numbers. Ideally, the lesson should end with the same number of started and completed.
g. Scenes started: displays, for all players, the game scenes that they have visited in their games, stacked from oldest (bottom) to newest (top). This allows teachers to spot trends and understand not only how far along they are, but also how they reached that point. This visualization is more useful after the lesson has ended, as it allows teachers to see the exact sequence of scenes played by each student, and might relate a bad performance with a set of scenes.

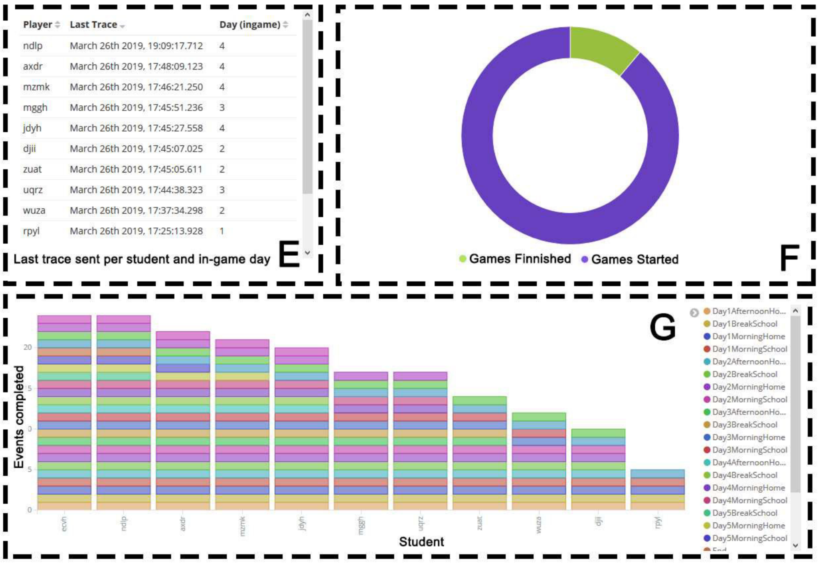

Figure 5: Game-specific visualizations added to the teacher dashboard from Fig. 3, based on teacher feedback.

We have developed the Conectado dashboard while also developing the game itself, including also game design and its pedagogical design. Concentrating all these roles greatly simplifies dashboard development since the points from sections 2.1.1 and 2.2.1 that refer to the game developer and teacher involvement can when handled by the same team that has also built the analytics system, require drastically less communication than what would be needed in less ideal circumstances. We later conducted several experiments to validate the dashboard's usability and usefulness.

We carried out experiments both to validate Conectado and its analytics dashboard., sending traces to the analytics framework described in Section 2. Details on the Conectado validation experiments are available in [Calvo-Morata et al. 2018], including participants, methodology, and results. Dashboard-wise, the main steps were as follows:

1. Participating teachers supervised the gameplay session in their classes using the dashboard, and used it to assess their students and understand what was happening in the class while using it. For example, they could identify inactive students using the "last traces sent" visualization (Fig. 5.E).
2. Once that all students finished the game, teachers were asked to fill a technology acceptance model (TAM) survey, which we used to evaluate the dashboard. The survey contained items on:
    - Previous technology-related experience
    - Previous analytics-related experience
    - Usability (how easy-to-use it is)
    - Usefulness (increased control individually and by group, supports assessment, etc.)

- ○ Intention to use

Validation dashboard experiments took place in two different schools of Spain with a total of 5 different teachers. Each teacher lead multiple game validation sessions. The results of these experiments show that most teachers were satisfied with the dashboard and its usability and usefulness, therefore validating its design. Results also confirmed that teachers that had previous experience in dashboard or analytics usage quickly understood and took advantage of the dashboard, reporting a positive experience in their feedback; while the teachers that reported little or no experience with dashboards and analytics often experienced problems during their first steps and understood individual visualizations worse. Teachers also provided feedback on UI-related issues such as font size and the use of clearer titles to increase usability (this is already apparent in *Fig. 5*). Finally, most of the teachers pointed out that, if available, they would gladly use these tools and dashboards with more games to support student assessment.

3.2 A custom dashboard and analysis for a serious game on workplace interaction

This dashboard, where again both analysis and visualizations were developed ad-hoc, was built for a game centered on workplace interactions developed within the EU H2020 RAGE project. In the game, named *Watercooler* [Hollins, Humphreys, Yuan, Sleightholme and Kickmeier-Rust 2017], the player must work as an office assistant in a simulated game studio to improve the teamwork between employees. Its goal is that, through improving, enabling, prompting and challenging the attitudes, values and social skills, the player leads the company to success. Information regarding game validation experiments is available in [Hollins et al. 2017, TUGraz 2018, 2019a].

The game design included a requirement to use the Thomas–Kilmann Conflict Mode Instrument (TKI) [Thomas and Kilmann 2008] to measure and display responses to the different conflict situations that the player is exposed to while working as a team leader in a simulated game development company. The TKI is based on two dimensions of behavior, *assertiveness,* and *cooperativeness*; and defines five different approaches based on the balance between both dimensions: *competing*, *accommodating*, *avoiding*, *collaborating* and *compromising*; see element (e) in *Fig. 6*. Specific analysis and visualization were developed to display the TKI categorization for each player. Additionally, certain situations allowed the player to exhibit, or avert, certain types of biases (for instance, based on gender, race, or fashion sense). Finally, the game allowed players to track office morale, productivity (in terms of shipped games), and awards for quality.

Fig. 6 displays the seven visualizations developed for this dashboard:

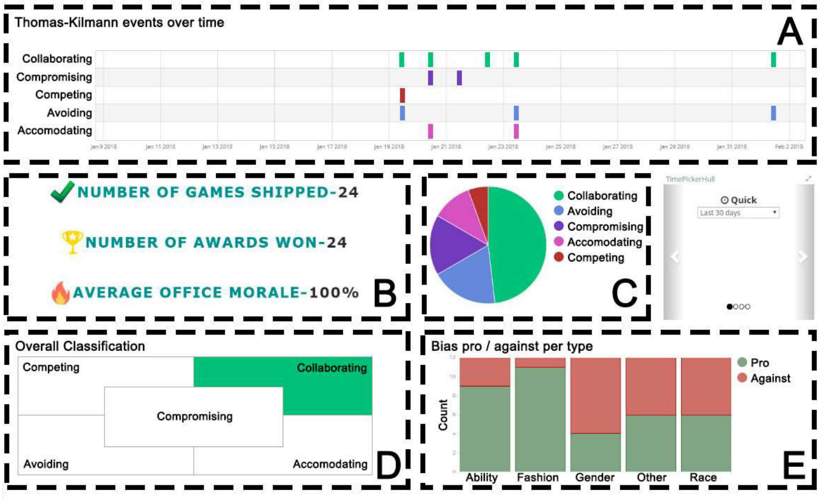

Figure 6: Some of the game-specific visualizations developed for the Thomas-Kilmann Conflict Mode Instrument.

a. Thomas-Kilmann classification of a specific student's answers over time. The swim lanes receive and sort the type of the answers received, providing an overview of progress and showing, at a glance, whether the game is working properly.
b. Three badges that provide a running score: Games shipped, a measure of team productivity; awards won, a measure of team quality; and Office morale, a percentage indicating the degree to which the player's in-game coworkers are happy with the player's choices. These scores help to maintain student engagement by allowing them to gauge how well they are doing compared to their classmates.
c. Overall Thomas-Kilmann classification for the player, displaying the category that has appeared more times from the five pre-set categories. This visualization is part of the standard TKI, and should be interpreted as having 2 axes. The horizontal axis corresponds to *cooperativeness*, while the vertical axis represents *assertiveness*. It is split into 5 areas, one for each TKI category. When the visualization of "a" gets too complex, this visualization quickly identifies the category.
d. Pie chart displaying the distribution of answers according to the 5 TKI categories. In addition to "c", this helps to identify if the distance between categories is close or far, making more significative the result displayed in c.
e. Bar chart displaying, for each bias, the ratio of responses where it was averted (green) or exhibited (red). Note that both counts are relevant, as most possible choices to the in-game conversations did not offer the opportunity to either avert or exhibit a bias.

Our role in this scenario was limited to dashboard developers, as we had minimal input regarding game knowledge or learning goals. Communication between dashboard development and game/learning design took place through email and videoconference.

At the beginning of the development, dashboard developers received a hand-drawn mockup of the dashboard. After several prototypes and developing the requested game-specific dashboard visualizations, we arrived at the dashboard displayed in Fig. 6.

The usability and usefulness of the improved dashboards was then evaluated by the stakeholders by a reviewing process where they evaluated the whole Game Analytics Suite [TUGraz 2019b]. Some of the evaluation aspects included: requirements fulfillment, usefulness, or impact on the effectiveness and added value. The review from the stakeholders showed overall good results, emphasizing in its usefulness speeding up the game development. Reviewers noted that using it is a non-frustrating experience, being easy to use. Finally, the results also pointed out that stakeholders were satisfied with the product and it provides added value for the game.

3.3 A custom dashboard with parametrized analysis for a serious game on formal specification

The last presented scenario is another ad-hoc dashboard created for the integration between the FormalZ [Prasetya et al. 2019] game and the analytics framework, as part of the IMPRESS Erasmus+ EU project. This game falls within the genre of "tower defense", where players seek to protect a base by placing defensive towers that must destroy wave after wave of enemies launched against it, with the waves increasing in difficulty as the players upgrade their defenses. In FormalZ, each level corresponds to an electronic circuit in where the circuit chips must transform the values of incoming particles into certain output values. However, the chips are not working properly, so the student must specify preconditions and postconditions that must hold before the particle is allowed into the chip, and on the transformed particles once they exit. Allowing invalid particles into the chip, or generating invalid outputs, resulted in loss of in-game "lives". These conditions are specified in natural language, which the student must transform into logic asserts to be interpreted by towers that destroy unwanted incoming particles.

To assess the students, game designers and analytics developers identified several indicators to track during gameplay, including money, the number of towers, and the number of remaining lives. These values were reported by the game to the analytics server in real-time, using xAPI-SG statements. Each wave initialization includes details about the conditions used, how close students are to the solution (proximity), and how much time the student took to write them. The last and most important event contains the player's choices of preconditions and postconditions. Every time the student decides to apply a new condition, an interaction event is reported to the server, mentioning how long the student took to write it. Time-dependent visualizations display all lines as starting at the start of the game session of each player. While not all players started their sessions at the exact same time, using the start of each session as zero makes comparisons between time-lines much easier to perform.

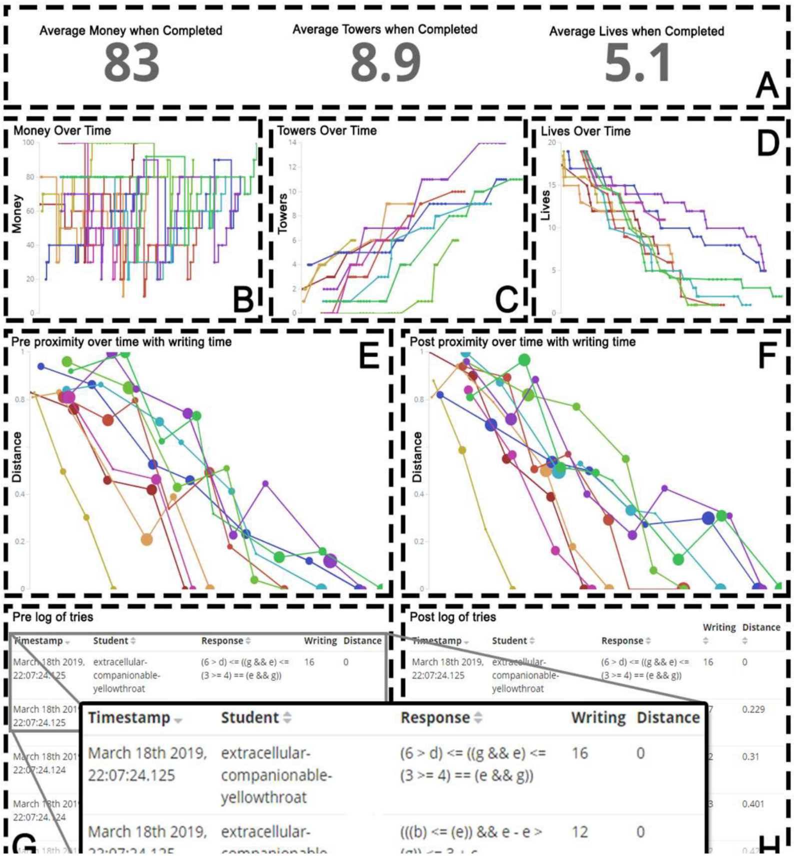

Figure 7: Dashboard created for FormalZ game as part of the Erasmus+ project Impress

Based on teacher requirements, we developed the dashboard displayed in *Fig. 7*, which contains 8 visualizations:

a. Averages when the level is completed: Money, Towers, and Lives. These three metrics work not only as control values to check performance compared with design values but also as badges to compare with other groups to improve the engagement.
b. Evolution of money over time, per student. This complex visualization may help to identify good and bad gameplay by the amount of money obtained. However, it is much more helpful filtered by a student as it provides a very good backtracking tool.
c. Number of towers over time, per student. Similar to b, but it provides other types of conclusions, e.g., a student that passes the game with a low number of towers, used it more efficiently.

d. Number of lives over time, per student. Like b and c, but in this case, the higher the lives, the better the performance. If a student loses lives quickly, something must be wrong, and assistance is needed immediately.
e. Distance between correct and player-selected precondition over time per student: The closer to zero, the closer to the correct answer. Additionally, dot size is used to represent the actual time the students spent writing the conditions, with bigger dots corresponding to longer time intervals. The filtered-by-student view is much more useful, as in b. Ideally, the closer and faster the student gets to zero distance, the better the performance.
f. Distance between correct and player-selected postcondition over time, per student: as in (e), but for the postcondition.
g. Detailed log of the preconditions as written by students, including the distance, writing time, timestamp, student, and the exact response given. This visualization is key to evaluate as it shows the exact answer given by the student. Also, it is useful to backtrack, as in this visualization, the evolution of the precondition can be seen.
h. Exact log of the post conditions given by the students, like (g) but for postconditions.

The dashboard can be used not only to assess the students but also to assess the difficulty of individual tower defense levels. For example, if too many towers need to be used, or students keep running out of budget or lives, despite having already shown knowledge of the level, this would indicate that the level needs to be revised. Conversely, having all students breeze through a level would indicate that that level lacks significant challenge and that students are probably being bored by the exercise.

While developing this dashboard, we also performed full analytics integration: we explained the possibilities of the xAPI-SG standard to the game's developers, shared a list of example visualizations, and created and discussed mockups while meeting with the game team and evaluators. After several rounds, the dashboard reached its final version shown in *Fig. 6.* We also created a gameplay-emulator to populate functional dashboard mockups with example data – indeed, the results displayed in *Fig. 6* were populated using this emulator, as the main experiments have not taken place yet. As of this writing, the dashboard and game are still in validation, but the early results and feedback received by both game developers and teachers are very promising.

4 Lessons learnt in teacher dashboard design

Each of the case-studies described in the previous section has led to improvements in the analytics platform and teacher dashboards. Additionally, we have tested the default teacher dashboard (not displayed above), including the alerts and warnings, in experiments with over 600 students and 150 teachers. Based on our experience, we make the following recommendations to analytics dashboard authors:

- Consider the specific requirements of teachers when designing visualizations, and be aware that the teachers may not be aware of these requirements until they start to see the possibilities and limitations of the actual dashboards.
- Start with small, simple dashboards, and let the teachers propose additional visualizations before including them. An overly-complex dashboard will lead to information overload and decrease usability and teacher engagement.

- Ensure that all visualizations include helpful, descriptive labels; and include mechanisms that allow teachers to see more detailed explanations on demand.
- Allowing teachers to move and reorganize dashboard visualizations on demand also leads to greater use; additionally, we have observed that teachers appreciate aesthetic use of colors, and having beautiful visualizations drives engagement.

To make visualizations more understandable, we have completely reworked their names and labels, with a focus on clarity; and include customizable pop-up descriptions to explain the more complex visualizations. For instance, in *Fig. 6* (g), the visualization of biases exhibited/averted requires, at first, a significant description to understand what is being displayed; but once understood, there should be no further need for bringing up the description, and it should therefore not receive permanent display space. Additionally, as can be seen in *Figs. 4-7*, we now allow dashboards to combine visualizations of different sizes so that more visually complex visualizations can be rendered in larger areas. Additionally, while in the initial teacher dashboards the positions of individual visualizations were not fixed, we have now made positioning entirely controllable. This has considerable advantages when comparing the dashboards for different experiments based on the same game, for example. On a related note, we also allow teachers to reorder and hide visualizations; this allows them to configure dashboards to display only what they find most useful, and facilitates comparisons in visualizations that share common axes.

The display of alerts and warnings has also been improved, by making better use of display area; for instance, showing triggered alerts and warnings directly in the general view if they are not too many, or showing only the most recent otherwise. *Fig. 8* depicts these improvements: in the original version, teachers must click on a student's name or access token to see the detailed alerts and warning that the student has triggered. In the updated version, teachers see details for each student directly on the main view, avoiding an extra interaction to display details.

Making a dashboard more understandable also makes it easier for teachers to reason on the underlying information, and to take actions based on these decisions, which is the goal of any analytics system: to induce new meaning or change behavior [Verbert, Duval, Klerkx, Govaerts and Santos 2013]. Therefore, usability is an important first step towards actionable feedback. We are also exploring other avenues to provide recommended actions for teachers. For example, we are considering the use of alerts to highlight statistical deviations from a baseline. This would first require sufficient baseline data to be gathered; for example, we can take all completion times from a validation run, and use these times to identify students who take significantly longer (say, one standard deviation) than their colleagues to finish. Since this analysis can be performed regardless of the game, it can be rolled into the default alerts system, benefitting all future users of the analytics system at essentially no increased cost for users.

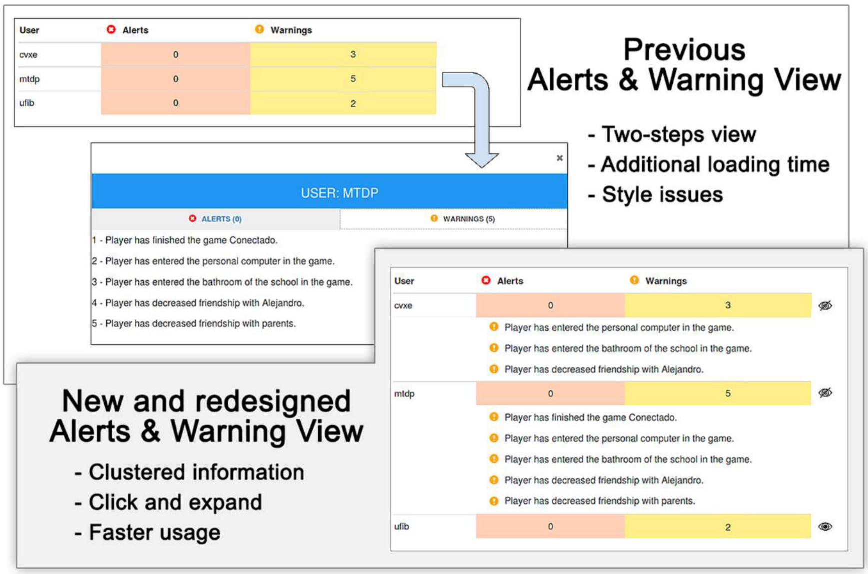

Figure 8: Previous (top) and updated (bottom) alerts & warnings view. The previous display only displayed counts of alerts and warnings; teachers had to click on student names to view the actual alerts and warnings for those students. The updated version does not require this context switch.

The entire process of creating visualizations, analyzing, and collecting game data requires a complex and flexible infrastructure underneath. This combination of complexity and flexibility is potentially fragile, and can give rise to a variety of problems: it may not support the volume of interaction traces sent by games; or receive types of data that break specific analyses, causing run-time errors that leave visualizations without data to display; or authors may remove a game for which data is still being received; among a host of possible problems. It is therefore important to continue working on simplifying the infrastructure itself and increasing the system's resiliency while streamlining the process of creating and using dashboards. The goal here is to allow this type of infrastructure to be deployed as easy-to-manage stand-alone servers that allow non-expert users to take full advantage of reliable but powerful learning analytics.

5 Conclusions

Teachers are key to increasing the adoption of serious games by schools, and Game Learning Analytics should, therefore, prioritize their specific needs. We consider that teachers' requirements should determine what information is to be collected and analyzed, to be later displayed on dashboards that are easy to understand for an average teacher. Dashboards should help teachers to make informed decisions not only after the games are played, but also while the game is ongoing and teacher interventions can still help players make the most of their sessions.

In this paper, we have identified several issues with teacher dashboards, including privacy and data collection, low teacher expectations regarding the outputs of the dashboards, lack of initial input when creating initial dashboards vs. late dashboard design requirements, and the use of dashboards for non-standalone games; and we have described how we have met these challenges by using simple anonymization via tokens and the xAPI-SG standard, a default set of visualizations that provides teachers with quick and easy-to-understand information to act on the previous contexts, support for custom-built dashboards (we present three case-studies), a flexible alerts and warnings system, and hierarchical dashboards. We have also provided a detailed step-by-step process for teacher dashboard authoring in collaboration with game designers and developers; if possible, we recommend involving both stakeholders in the process, as they provide complementary inputs from their different points of view: developers and designers from their perspective closer to the game; and teachers closer to the actual use of dashboards.

In our experiments using these dashboards, we have identified recommendations for dashboard authors, together with several improvements to our dashboards that make them more understandable and usable. Many of these improvements are already implemented, and others will be added on subsequent iterations, to be tested and validated using the *Conectado* serious game and other games from RAGE and BEACONING H2020 projects.

6 Future Work

While we have done significant work to develop the analytics framework that collects, analyzes, and displays data from different games, we still have much to do if we wish to allow non-specialists to create good dashboards by themselves. Improving the process of dashboard creation will require two key steps: first, we will have to add simple editors to configure both visualizations and, much more difficult, the associated analyses from where they will receive their data. Second, we must provide a quick method to test whether these analyses and visualizations are performing as expected since non-specialists are more likely to use a trial-and-error approach to development due to their lack of development expertise. One approach is to allow the upload of small amounts of xAPI-SG or CSV-formatted interaction traces to test with, even if these have not been collected by running a real game (as illustrated in Section 3.3). The platform would process, analyze and show visualizations of these, as usual, allowing dashboards to be tinkered with before applying them to an actual game. In the same sense, decoupling sets of traces from dashboards should allow trying out a new dashboard on data from an existing game, and keeping it if the teachers consider it better than the old version; or even allowing them to choose the dashboard that they feel most comfortable with.

In addition to this, we will continue to test and evaluate the platform and its dashboards in further experiments with new games.

Acknowledgements

This work has been partially funded by Regional Government of Madrid (eMadrid P2018/TCS4307), by the Ministry of Education (TIN2017-89238-R), by the European Commission (RAGE H2020-ICT-2014-1-644187, BEACONING H2020-ICT-2015-687676, Erasmus+ IMPRESS 2017-1-NL01-KA203-035259) and by the Telefónica-Complutense Chair on Digital Education and Serious Games.

References

[Alonso-Fernández, Calvo-Morata, Freire, Martínez-Ortiz and Fernández-Manjón 2019] Alonso-Fernández, C., Calvo-Morata, A., Freire, M., Martínez-Ortiz, I., Fernández-Manjón, B.: 'Applications of data science to game learning analytics data: A systematic literature review'; Computers and Education, Vol. 141 (2019). https://doi.org/10.1016/j.compedu.2019.103612

[Alonso-Fernandez, Calvo, Freire, Martinez-Ortiz and Fernandez-Manjon 2017] Alonso-Fernandez, C., Calvo, A., Freire, M., Martinez-Ortiz, I., Fernandez-Manjon, B.: 'Systematizing game learning analytics for serious games'; In 2017 IEEE Global Engineering Education Conference (EDUCON). IEEE (2017), pp. 1111–1118. https://doi.org/10.1109/EDUCON.2017.7942988

[Bauckhage, Drachen and Thurau 2015] Bauckhage, C., Drachen, A., Thurau, C.: 'The Age of Analytics'; IEEE Transactions on Computational Intelligence and AI in Games, Vol. 7, No. 3 (2015), pp. 205–206. https://doi.org/10.1109/tciaig.2015.2467166

[Calvo-Morata, Rotaru, Alonso-Fernandez, Freire, Martinez-Ortiz and Fernandez-Manjon 2018] Calvo-Morata, A., Rotaru, D. C., Alonso-Fernandez, C., Freire, M., Martinez-Ortiz, I., Fernandez-Manjon, B.: 'Validation of a Cyberbullying Serious Game Using Game Analytics'; IEEE Transactions on Learning Technologies (2018), pp. 1–12. https://doi.org/10.1109/TLT.2018.2879354

[Chen, Chiang and Storey 2012] Chen, H., Chiang, R. H. L., Storey, V. C.: 'Business Intelligence and Analytics: From Big Data to Big Impact'; MIS Quarterly, Vol. 36, No. 4 (2012), pp. 1165–1188.

[Connolly, Boyle, MacArthur, Hainey and Boyle 2012] Connolly, T. M., Boyle, E. A., MacArthur, E., Hainey, T., Boyle, J. M.: 'A systematic literature review of empirical evidence on computer games and serious games'; Computers & Education, Vol. 59, No. 2 (2012), pp. 661–686. https://doi.org/10.1016/j.compedu.2012.03.004

[Denis and Jouvelot 2005] Denis, G., Jouvelot, P.: 'Motivation-driven educational game design: applying best practices to music education'; In Proceedings of the 2005 ACM SIGCHI International Conference on Advances in computer entertainment technology (2005), pp. 462–465.

[Drachen, El-Nasr and Canossa 2013] Drachen, A., El-Nasr, M. S., Canossa, A.: 'Game Analytics: Maximizing the Value of Player Data'; (M. Seif El-Nasr, A. Drachen & A. Canossa, Eds.). London: Springer London (2013). https://doi.org/10.1007/978-1-4471-4769-5

[Hollins, Humphreys, Yuan, Sleightholme and Kickmeier-Rust 2017] Hollins, P., Humphreys, S., Yuan, L., Sleightholme, G., Kickmeier-Rust, M.: 'The 'Water Cooler' Game'; In Proceedings of the 11th European Conference on Games Based Learning, ECGBL 2017 (2017).

[Jivet, Scheffel, Drachsler and Specht 2017] Jivet, I., Scheffel, M., Drachsler, H., Specht, M.: 'Awareness Is Not Enough: Pitfalls of Learning Analytics Dashboards in the Educational Practice'; (2017), pp. 82–96. https://doi.org/10.1007/978-3-319-66610-5_7

[Jivet, Scheffel, Drachsler and Specht 2018] Jivet, I., Scheffel, M., Drachsler, H., Specht, M.: 'License to evaluate: Preparing learning analytics dashboards for the educational practice'; In Learning Analytics & Knowledge Conference (2018).

[Keim, Andrienko, Fekete, Görg, Kohlhammer and Melançon 2008] Keim, D., Andrienko, G., Fekete, J.-D., Görg, C., Kohlhammer, J., Melançon, G.: 'Visual Analytics: Definition, Process and Challenges'; In Information Visualization (2008), pp. 154–175. https://doi.org/10.1007/978-3-540-70956-5

[Loh, Sheng and Ifenthaler 2015] Loh, C. S., Sheng, Y., Ifenthaler, D.: 'Serious Games Analytics'; (C. S. Loh, Y. Sheng & D. Ifenthaler, Eds.). Cham: Springer International Publishing (2015). https://doi.org/10.1007/978-3-319-05834-4

[I. Perez-Colado, Alonso-Fernandez, Freire, Martinez-Ortiz and Fernandez-Manjon 2018] Perez-Colado, I., Alonso-Fernandez, C., Freire, M., Martinez-Ortiz, I., Fernandez-Manjon, B.: 'Game learning analytics is not informagic!'; In 2018 IEEE Global Engineering Education Conference (EDUCON). IEEE (2018), pp. 1729–1737. https://doi.org/10.1109/EDUCON.2018.8363443

[I. J. Perez-Colado, Rotaru, Freire-Moran, Martinez-Ortiz and Fernandez-Manjon 2018] Perez-Colado, I. J., Rotaru, D. C., Freire-Moran, M., Martinez-Ortiz, I., Fernandez-Manjon, B.: 'Multi-level game learning analytics for serious games'; 2018 10th International Conference on Virtual Worlds and Games for Serious Applications, VS-Games 2018 - Proceedings (2018). https://doi.org/10.1109/VS-Games.2018.8493435

[Prasetya, Leek, Melkonian, Tusscher, van Bergen, Everink, et al. 2019] Prasetya, I. S. W. B., Leek, C. Q. H. D., Melkonian, O., Tusscher, J. ten, van Bergen, J., Everink, J. M., et al.: 'Having fun in learning formal specifications'; In Proceedings of the 41st International Conference on Software Engineering: Software Engineering Education and Training (2019), pp. 192–196.

[Schwendimann, Rodriguez-Triana, Vozniuk, Prieto, Boroujeni, Holzer, et al. 2017] Schwendimann, B. A., Rodriguez-Triana, M. J., Vozniuk, A., Prieto, L. P., Boroujeni, M. S., Holzer, A., et al.: 'Perceiving learning at a glance: A systematic literature review of learning dashboard research'; IEEE Transactions on Learning Technologies (2017), pp. 30–41. https://doi.org/10.1109/TLT.2016.2599522

[Serrano-Laguna, Martínez-Ortiz, Haag, Regan, Johnson and Fernández-Manjón 2017] Serrano-Laguna, Á., Martínez-Ortiz, I., Haag, J., Regan, D., Johnson, A., Fernández-Manjón, B.: 'Applying standards to systematize learning analytics in serious games'; Computer Standards & Interfaces, Vol. 50, No. September 2016 (2017), pp. 116–123.

https://doi.org/10.1016/j.csi.2016.09.014

[Shoukry, Göbel and Steinmetz 2014] Shoukry, L., Göbel, S., Steinmetz, R.: 'Learning Analytics and Serious Games: Trends and Considerations'; In Proceedings of the 2014 ACM International Workshop on Serious Games - SeriousGames '14. New York, New York, USA: ACM Press (2014), pp. 21–26. https://doi.org/10.1145/2656719.2656729

[Thomas and Kilmann 2008] Thomas, K. W., Kilmann, R. H.: 'Thomas-Kilmann Conflict Mode'; TKI Profile and Interpretive Report (2008), pp. 1–11.

[TUGraz 2018] TUGraz: 'RAGE - Data Evaluation Watercooler pilot 1'; (2018). https://doi.org/10.5281/zenodo.2577071

[TUGraz 2019a] TUGraz: 'RAGE - Data Evaluation Watercooler pilot 2'; (2019a). https://doi.org/10.5281/ZENODO.2564932

[TUGraz 2019b] TUGraz: 'RAGE - Data Final Evaluation RAGE Components'; (2019b). https://doi.org/10.5281/zenodo.2576902

[Verbert, Duval, Klerkx, Govaerts and Santos 2013] Verbert, K., Duval, E., Klerkx, J., Govaerts, S., Santos, J. L.: 'Learning Analytics Dashboard Applications'; American Behavioral Scientist, Vol. 57, No. 10 (2013), pp. 1500–1509. https://doi.org/10.1177/0002764213479363