Utilizing Web Analytics in the Context of Learning Analytics for Large-Scale Online Learning

Tobias Rohloff
Hasso Plattner Institute
Potsdam, Germany
tobias.rohloff@hpi.de

Sören Oldag
Hasso Plattner Institute
Potsdam, Germany
soeren.oldag@student.hpi.de

Jan Renz
Hasso Plattner Institute
Potsdam, Germany
jan.renz@hpi.de

Christoph Meinel
Hasso Plattner Institute
Potsdam, Germany
christoph.meinel@hpi.de

***Abstract*—Today, Web Analytics (WA) is commonly used to obtain key information about users and their behavior on websites. Besides, with the rise of online learning, Learning Analytics (LA) emerged as a separate research field for collecting and analyzing learners' interactions on online learning platforms. Although the foundation of both methods is similar, WA has not been profoundly used for LA purposes. However, especially large-scale online learning environments may benefit from WA as it is more sophisticated and well-established in comparison to LA. Therefore, this paper aims to examine to what extent WA can be utilized in this context, without compromising the learners' data privacy. For this purpose, Google Analytics was integrated into the Massive Open Online Course platform of the Hasso Plattner Institute as a proof of concept. It was tested with two deployments of the platform: openHPI and openSAP, where thousands of learners gain academic and industry knowledge about engineering education. Besides capturing behavioral data, the platforms' existing LA dashboards were extended by WA metrics. The evaluation of the integration showed that WA covers a large part of the relevant metrics and is particularly suitable for obtaining an overview of the platform's global activity, but reaches its limitations when it comes to learner-specific metrics.**

***Keywords*—Learning Analytics, Web Analytics, Learning Dashboards, MOOCs, Online Learning Environments**

I. Introduction

With the rise of the World Wide Web, the need of website operators to gather information about users and their behavior arose. To satisfy this demand, the field of Web Analytics (WA) emerged. Originally intended for e-commerce websites, it captures users' interactions and reveals valuable insights about the audience and their activity. Therefore, it has especially gained attention from business corporations, which utilize WA for decision-making processes. Consequently, WA rapidly evolved and is a common technique today that is widely used and no longer restricted to e-commerce websites only [1].

The development of the Internet also results in learning being revolutionized. While facilitated or blended learning takes place online only partly, Massive Open Online Courses (MOOCs) are usually held exclusively on the Internet. MOOCs enable anyone in participating in online courses similar to actual university courses. Therefore, they have to deal with a large scale of learners with different backgrounds from all over the world [2]. Besides the benefits and potential of these new forms of learning, they also bear new challenges. In a traditional classroom setup, instructors can observe their students face-to-face. However, in online courses teachers cannot directly watch students while learning. Therefore, monitoring systems are needed, which keep track of learners' progress and interactions. For this purpose, the field of Learning Analytics (LA) emerged for collecting and analyzing data of learners with the goal to support their learning process [2].

WA and LA are subtypes of the general field of analytics and are thus related to each other. Both methods gather and analyze data about users and their interactions on online platforms to understand the audience and their behavior. This data is eventually utilized to derive actions for optimizations. Even though the underlying objective differs, LA may benefit from integrating WA. While LA is a relatively new and active research field, WA is sophisticated and well-established meanwhile. Thus, by using it for analyzing the behavior of learners one could take advantage of its advanced development. Nevertheless, WA tools have not been profoundly used for this purpose, so far [3]. Therefore, this paper examines the research question:

RQ1 To what extend can Web Analytics methods be used in the context of Learning Analytics to gather insights in learning behavior and outcome on e-learning platforms?

Addressing this question involves considering sub-questions, which focus on different aspects of the main problem:

RQ1.a Can online learning activity be mapped to Web Analytics concepts?

RQ1.b How can different stakeholders of e-learning platforms be provided with Web Analytics insights?

RQ1.c Can Web Analytics methods improve the usefulness of Learning Analytics insights?

To answer these questions, a WA tool is integrated exemplarily into the white label MOOC platform of the Hasso Plattner Institute (HPI) [4] as a proof of concept. It was evaluated with two deployments of the platform: openHPI and openSAP, where thousands of learners gain academic and industry knowledge about engineering education. For testing the integration of WA, dashboards are used as a typical use case for analytics in general. The applicability of WA in the context of LA is evaluated by discussing potentials and limitations of WA by reference to the proof of concept. Besides, the usability of the revised dashboards is evaluated by conducting an expert survey.

II. Related Work

In general, WA is widely used on the Internet. However, there has not been much research in making use of WA capabilities for analyzing learner's behavior on e-learning platforms so far. Previous work related to this topic was still reviewed and is presented in this section.

Cooper [3] claims that the reasons for the missing utilization of WA tools in the e-learning context are mainly privacy concerns regarding collected activity data. As the majority of WA tools stores behavioral data on external servers, control over captured data is lost. Open-source alternatives, such as Matomo, enable operators to control the collected data. However, according to the authors tracking of these tools is usually less fine-grained. In general, WA does not meet all needs of LA as it does not cover all information of a learning process that might be useful.

Moissa et al. [5] developed a visualization tool for behavioral data collected in the e-learning environment AdaptWeb that uses Piwik (now Matomo) to capture and store analytics events. Besides the WA tool, the implemented application also retrieves data from the existing database of the platform. The tool provides 20 metrics by combining both data sources. However, the paper does not reveal, which metrics are based on Piwik and which are computed by querying the local database. In addition, evaluation and limitations of the use of WA in the e-learning context are not discussed as well.

Romanowski and Konak [2] integrated Google Analytics into the website of a course of the Penn State University to understand how students interact with it. For data collection, page tagging was used. Different pages and contents were compared in regard to the number of page views and the average time on page to discover which features of the website are most effective. The authors concluded that Google Analytics can gather enough data to understand learners' behavior, but should be combined with further log data of the platform itself to accomplish comprehensive analysis results.

Luo et al. [6] conducted a case study to ascertain potentials and limitations of utilizing Google Analytics for LA purposes in the context of advanced degree online programs. Activity of students of an online course of the Pennsylvania State University was captured using page tagging. For analysis, the researchers considered learner demographics, traffic metrics, efforts of learners, sequence of interactions with contents, and used technology. According to the authors, Google Analytics is well suited for providing an overview of learning processes on e-learning platforms. However, it can not be used to generate personalized learning reports. Therefore, they inferred that using Google Analytics alone might be to limiting.

In contrast to the assumption of Cooper, several big MOOC platforms have integrated Google Analytics in their websites. A manual examination revealed that edX, Coursera, and Udacity have included the Google Analytics page-tagging snippet in their website. EdX specifies in their developer's guide that Google Analytics is used to track all page views and obtain metrics, such as referrers and search terms, used to find the website [7]. Consequently, WA is not used for improving the learning experience of users, but to measure and increase awareness of the platform. However, the other providers do not state their actual intentions and purposes for using WA.

It can be summarized that an integration of WA tools for LA purposes was done only in a basic scope so far. Related work is limited to collecting behavioral data using page tagging and analyzing a fundamental choice of different dimensions and metrics. Although privacy concerns of page tagging are discussed, other data collection methods were not considered in this context, yet. Some limitations of using WA in the context of e-learning were ascertained. Using WA alone might be too limiting to analyze learners' behavior in its entirety. Instead, it could be used in combination with additional LA capabilities to achieve comprehensible results.

Compared with the related work, this paper considers the full potential of WA by taking into account different tools, data collection methods, and analysis capabilities. Consequently, results could become more meaningful and universal. However, limitations identified by the presented papers might be valid for this approach as well.

III. Privacy Concerns

When analyzing user activity, a huge amount of data about users and their behaviors is collected and stored. Therefore, privacy laws have to be considered when integrating LA or WA into a website. Applicable regulations depend on the type of data that is processed. When collecting only anonymous data, information about individual users cannot be derived and their privacy is not affected. Consequently, data privacy laws are only relevant if collected data contains Personally Identifiable Information (PII). For this paper, utilization of WA tools is evaluated using the example of the HPI MOOC platform. Therefore, only applicable regulations were examined in the following. As the service is based in Germany, the European Union's General Data Protection Regulation (GDPR) is the law in force for governing processing of personal data.

Art. 4 GDPR defines personal data as "any information relating to an identified or identifiable natural person ('data subject')" and an identifiable natural person as "one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one or more factors specific to the [...] identity of that natural person". Collecting this kind of data is allowed only if any of the prerequisites listed in Art. 6 GDPR is fulfilled. Among others, this might be the explicit consent of the data subject or the necessity of data processing for purposes based on legitimate interests of the controller. As LA is exclusively used for improving the learning experience of users and optimizing the platform, it is considered as a legitimate interest. Therefore, collecting and processing behavioral data for these purposes is allowed and does not require an explicit consent of the learners. This also applies to the envisaged utilization of WA in context of this paper, where additional pseudonymization techniques are applied.

IV. Concept and Implementation

To evaluate the applicability of WA in context of LA and thus answer the main research question of this paper, a WA tool is integrated into the HPI MOOC platform as a proof of concept. Therefore, this section presents the concept and implementation for realization of this goal.

A. Choice of Web Analytics Service

There is a great number of different analytics suites available that can be used for analysis of learners' behavior. Even though this work aims to evaluate the utilization of WA tools for this purpose only, there are still many services to choose from. We decided to integrate only one of these as an example and representative for WA tools in general, as their core features are mainly the same. However, there are differences in regard to more specific and advanced analysis capabilities, processing limitations, and pricing models. We evaluated the proprietary tools Google Analytics and Adobe Analytics, and Matomo as an open source alternative.

Google Analytics is well-established as it is the most popular WA tool. Consequently, it is a paragon in its field and therefore well suited for examining the applicability of WA for LA purposes in general. It comes with a wide range of features, which enable evaluation of different aspects of WA. Even though Adobe Analytics still exceeds these analysis capabilities, the majority of additional features are not applicable in the context of e-learning. Furthermore, Adobe Analytics is highly complex and not as good documented as Google Analytics. Therefore, integration of it would be more complicated and costly. In contrast to the traditional, self-hosted setup of Matomo, Google Analytics and Adobe Analytics run on servers in the cloud. As a consequence, it does not have to be taken care of deployment and maintenance of the services. Furthermore, the corresponding machines are highly performant, which results in relative short response times even for more complex computations. Nevertheless, data privacy might be an issue when storing user activity data on external servers, especially when they are located outside the European Union (EU). All in all, Google Analytics is the best suited WA tool for the purpose of this work when the data privacy concerns are addressed, as it supports a broad range of functions, is easy to set up, and satisfies the needs and requirements for integration into the existing infrastructure and architecture of the platform. As the general concepts and main features of WA are the same for all related tools, the findings of this paper are for the most part also valid for the utilization of WA in context of LA in general.

B. Integration and Data Collection

This section presents a concept for collecting user interaction data on the HPI MOOC platform and transmitting it to Google Analytics as foundation for further analysis tasks.

1) Processing Pipeline: There are different data collection methods available in Google Analytics. The most common and easiest one is page tagging, which requires to insert a small JavaScript snippet provided by Google Analytics into each page. This snippet takes care of gathering needed data and sending it to the WA service. Although integration using this technique is simple and effortless, it comes with some issues. Page tagging slightly increases page loading times as another JavaScript file needs to be loaded and executed. Besides, the existing data collection procedure cannot be used as page tagging would incorporate a separate event tracking. Moreover, page tagging cannot be used properly in the native apps, where mobile SDKs would have to be utilized. This would cause code duplication and is vulnerable for inconsistencies between the different clients. Besides these technical issues, there are also privacy concerns in regard to page tagging as control over data that is sent to the service would be lost.

The HPI MOOC platform already has an analytics infrastructure [8], [9], which takes care of tracking and persisting certain user activities for LA purposes. In the platform's Service Oriented Architecture (SOA), the *lanalytics* service receives interaction events from any client and executes pipelines each representing an independent Extract, Transform, Load (ETL) process. Thanks to the flexible and extensible architecture, we integrated the data collection for Google Analytics into the existing service (Figure 1). Thus, a new pipeline was added for transforming interaction events according to the Google Analytics hit schema and emitting them via the Google Analytics Measurement Protocol. The pipeline consists of multiple steps: extraction, enrichment, pseudonymization, schema transformation and batching before transferring them.

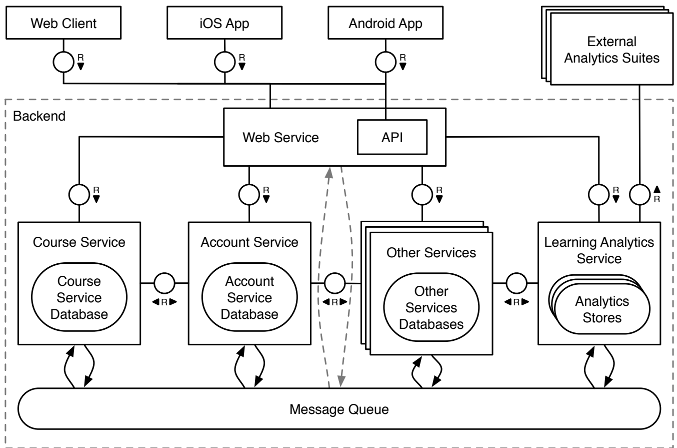

Fig. 1. Platform's Architecture with internal and external Analytics Stores

The asynchronous data collection has no impact on the performance of the website. Since tracking is already implemented in the clients, there is no need to adapt either the web client or mobile apps. Instead, all logic related to Google Analytics is encapsulated in the *lanalytics* service. Moreover, the basis of data stored locally in analytics stores and hits sent to the WA tool are the same, which prevents inconsistencies. As hits are constructed, pseudonymized and sent manually, we can completely decide, which interaction data is sent to Google Analytics. Thus, control over the data that is sent to third parties is regained.

2) Mapping Analytics Events to Hits: In order to be processed by Google Analytics, each analytics event needs to be transformed to a hit, which follows the schema defined by the Measurement Protocol and represents the underlying interaction in the best possible way. Therefore, depending on the event type and available context data appropriate parameters are specified manually based on the different analytics events. In general, each hit has a type indicating the kind of interaction it describes. Some parameters may be set only for specific types. The types of hits constructed by the mapping are limited to *pageview* and *event*. All events triggered when a user visits a certain page are mapped to *pageview* hits. Otherwise generic *event* hits are created. To ensure data privacy the SHA-256 hash of the user ID is used, which cannot be used by third parties to identify the user. Also the IP address and User-Agent are omitted, by sending empty payloads. The implementation of this mapping proves that online learning activity can be mapped to WA concepts, which satisfies RQ1.a. However, creating a generic mapping is virtually impossible as each e-learning platform and WA tool has different data schemas and capabilities. This situation could be improved by using a standardized format on the platform side, like xAPI.

3) Hit Batching and Emitting: The Measurement Protocol supports sending batches of hits inside a single Hypertext Transfer Protocol (HTTP) request. This feature is used in this context to lower the number of requests sent to Google Analytics, and thus increase performance. Batching and emitting of hits with an internal message queue prevents data loss and simplifies error handling. As requests are sent to an external service, connection errors are more likely to occur than it is the case when accessing local databases. One limitation of the Measurement Protocol is that it can process only hits that are not older than four hours [10]. To cope with this restriction during low activity times, a timeout less then four hours is assigned to each received hit. If this timeout expires before the hit was emitted, all outstanding hits are dispatched even though the maximum batch size it not reached, yet. For error handling, a message received by the consumer is acknowledged only if it was sent successfully to Google Analytics. Consequently, acknowledgement of messages is outstanding as long as the maximum batch size if not reached. If an error occurs while sending a batch of hits, the corresponding messages are negatively acknowledged. This results in the messages being requeued. Consequently, the consumer receives these messages again, and thus automatically retries sending them to Google Analytics. Besides, the mentioned data loss issue is prevented. The message queue is configured to be durable, i. e., store unacknowledged messages on disk [11]. This additional persistence layer makes sure, that no hit is getting lost even if the *lanalytics* service, the message broker, or even the machine running the service is shut down or crashes. The whole process is visualized in Figure 2.

4) Data Privacy: As discussed, data privacy laws need to be considered only when processing PII. Therefore, it must be determined first, whether PII is collected. As described before, the hashed user ID is sent to Google Analytics. The ID of a user is considered as a pseudonym of the person and thus as PII according to the GDPR. Even though it is sent as a hash, it is possible for anyone that has access to the platform's databases to identify a certain user given its hashed ID by re-computing all hashes. Nevertheless, Google does not have access to the platform's databases and consequently cannot identify single users. According to the GDPR, location data also belongs to personal data. As the geographical location is retrieved from the user's IP address, it is only a rough estimation of the city or the country of the user's location. Thus, it is not considered as PII. According to the GDPR, users have the right to receive a copy of collected data about themselves and can claim correction and erasure of this data. Google Analytics provides the possibility to download a file containing all collected data of a certain user. Besides, the entire personal data of a single user can be deleted either in the web frontend or via User

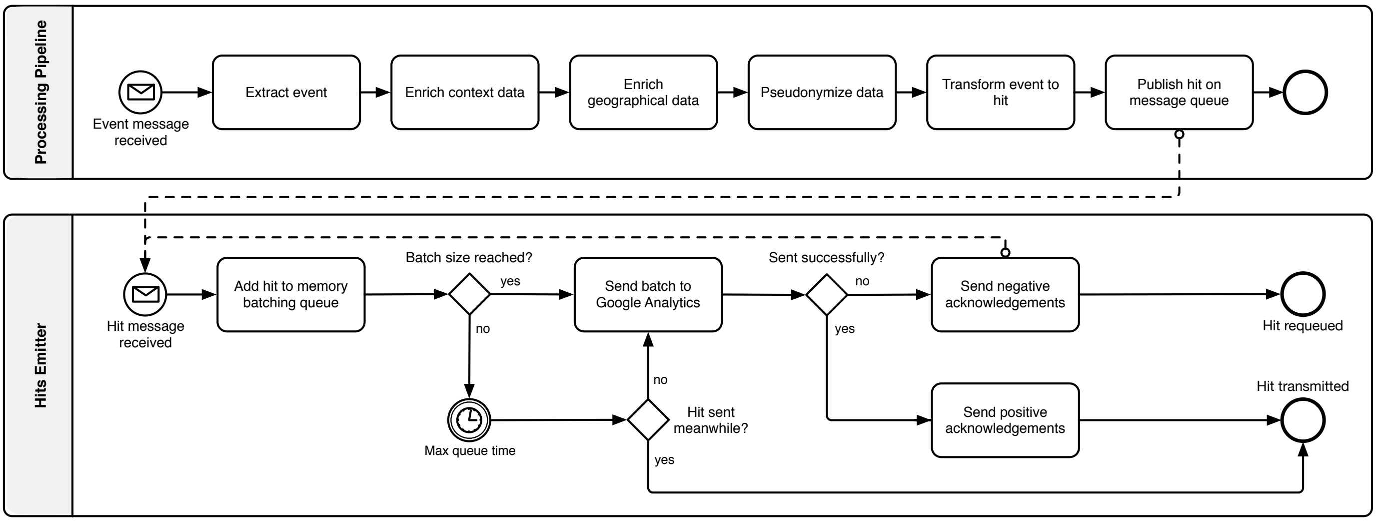

Fig. 2. Process for Interaction Data Collection, Processing, Batching and Emitting to Google Analytics in the Platform's Lanalytics Service

Deletion API. On request of a user, the person in authority can take care of providing the copy of data or deleting the data of the submitter. Nevertheless, hits once sent to Google Analytics cannot be modified anymore. Therefore, when a user requests correction of data, it can only be deleted to ensure correctness. Google is an US-based company and collected data is stored on servers in the US as well. Therefore, special regulations might apply, because the US is a third country from the perspective of the EU, in which the GDPR is not in force. However, Google has a Privacy Shield certificate, which causes the level of data privacy of the company to be classified as appropriate to the GDPR. Thus, an explicit approval of the user is not necessary.

C. Reporting of Analysis Results

After user interaction data has been collected and analyzed, results need to be reported to the stakeholders of the platform. There are currently several ways how LA insights are provided to the stakeholders in the HPI MOOC platform. For this paper, we decided to utilize dashboards as a typical analytics use case for two different objectives. First, RQ1.b is examined by integrating WA metrics retrieved from Google Analytics into the dashboards. Second, the main research question is approached by combining newly acquired WA insights and existing LA metrics and thus demonstrating to what extent WA can be utilized in the context of e-learning and which parts of the dashboards require LA-specific processing. For these purposes, different stakeholders need to be considered. In general, the platform has four types of stakeholders that differ in their needs in regard to LA:

Platform owners are interested in the overall performance of the platform across all courses.

Teaching teams are in charge of specific courses in which they inform and support the learners.

Learners are mainly interested in their own learning progress in courses they are enrolled in.

Researchers might use any kind of LA data depending on the research question they are examining.

Currently, there are two existing types of LA dashboards on the HPI MOOC platform: a global one mainly intended for platform owners and a course-specific one for the use by the teaching teams. The usefulness of WA insights in the e-learning context (RQ1.c) should be evaluated by comparing the usefulness of the existing dashboards and the extended ones incorporating WA metrics. Consequently, the focus is set on platform owners and teaching teams for this study.

1) Retrieval of Google Analytics Metrics: Independently from the actual approach for implementing the dashboards it is required to programmatically retrieve analysis results from Google Analytics. Therefore, two Application Programming Interfaces (APIs) are provided, which are be utilized for this purpose. The Reporting API enables retrieval of preprocessed and aggregated reports defined by a certain query. The Realtime API enables retrieval of realtime data. However, this API can return only a small range of basic dimensions and metrics. Google Analytics reports can be used to obtain or derive certain LA metrics, which can subsequently be integrated into the platform and visualized in the dashboards. For this purpose, the metrics provided by the *lanalytics* services are extended by metrics querying Google Analytics. Currently, metrics about the general activity, enrollments, learner progress, forum activity, device usage, geographical attributes and demographical attributes are visualized within the dashboards.

Instead of replacing existing metrics, this work focuses on integrating new WA metrics into the platform, which are relevant in the e-learning and MOOC context. A typical WA topic that is not well represented in the existing metrics are sessions. The reason for this is that computation of session-related metrics on raw event data is expensive. However, analyzing sessions of learners can help to understand, how often and how long users are learning on the platform. Consequently, several metrics are implemented querying appropriate Google Analytics dimensions and metrics, like the average session duration and days since last session.

Another metric that can be easily obtained from Google Analytics is the number of active users at a certain point in time, as it is also common for WA. Therefore, a metric is implemented that returns the number of active learners for each day and hour of a given date range. Besides, another metric is added aggregating this data by calculating the average number of active users for the hours of each day of week. While the first one can be used to obtain the actual activity of the past, the second metric gives on overview about typical weekdays and daytimes learners are accessing the platform.

Analyzing how users are navigating through a course can help to identify problems of its structure and contents. Unfortunately, it is also an expensive task when working with raw event data only. However, analysis of navigation paths is also a common WA task, which is why two corresponding metrics are integrated. The first one identifies exit items, i. e., items being regularly the last ones within a session and consequently might cause session exits. A high exit rate could indicate that the content is too complex or incomprehensible causing frustration of learners, which results in session exists. The other metric computes the percentage of backjumps for learning items, i. e., the proportion of page views originating from any succeeding item regarding the structure of the course. If during the progress of a course many learners return to a certain previous item again, this might indicate that prior knowledge being taught in this item was not understood well by a large part of learners.

2) Limitations: Although Google Analytics can be used to retrieve several metrics that are useful for operators and teaching teams of a MOOC platform, there are some limitations in regard to the kind of data that can be obtained. These limitations are also relevant for answering RQ1.a as they reveal mismatches between LA and WA concepts. In general, the web frontend of Google Analytics is used by the majority of customers, which is why Google mainly focuses on implementation of this component. As a result, a few information can only be extracted from the frontend, but not via API. Usually, WA is used to analyze behavior of the entire user base or certain user segments. Therefore, the API does return not any data about individual users [12], [13]. As a result, learner-specific metrics can not be implemented using Google Analytics. Due to the typical purpose of WA tools, Google Analytics comes with advanced e-commerce analysis capabilities. Among other features, this includes measuring conversions of predefined goals and analyzing the shopping behavior visualized by a funnel that shows at which stages users abandon the buying process. In context of this paper, the utilization of these e-commerce features in the MOOC context is examined. For example, the progress of a MOOC can be compared with the purchase of a product. Following this idea, different shopping stages can be mapped to actions concerning a MOOC and vice-versa as shown in Table I.

TABLE I
POSSIBLE MAPPING OF E-COMMERCE STEPS TO MOOC ACTIONS

E-Commerce Step		MOOC Action
Click on a certain product	→	Click on a certain course
View product details	→	View course details
Add product to cart	→	Enroll for course
Remove product from cart	→	Unenroll for course
Several checkout steps	→	Visiting learning items of course
Complete purchase	→	Pass exam and complete course

Based on this mapping, additional hits could be sent to Google Analytics containing the corresponding e-commerce parameters. As a result, respective analysis capabilities could also be utilized. For example, completion rates of courses could be calculated using conversions and corresponding funnels could be analyzed. This would make it possible to identify sections or items of a course that cause learners to abandon the course. The main problem of this idea is the fact that the e-commerce metrics are based on single sessions and cannot be calculated across multiple sessions of the same user [14]. This contradicts the general concept of MOOCs as an entire course can usually not be completed within a single session. Instead, a course typically runs over several weeks and is elaborated by a user in multiple sessions. However, the course progress of an user would not be considered in its entirety, but as several independent attempts to complete the course. Therefore, the e-commerce concept of Google Analytics can only be applied to MOOCs partly.

3) Requirements and Issues of existing Dashboards: The HPI MOOC platform already provides a global and course dashboard. However, several issues with these existing solutions have been identified by conducting expert interviews with relevant stakeholders. Six employees of openSAP were interviewed about usage scenarios of LA reporting capabilities of the platform, especially the global and course dashboard. The participants hold different occupational roles in context of openSAP. Three of them are in charge of certain courses as members of the corresponding teaching teams while two hold the role of the platform owner. The last person has experience in both roles. Despite the small number of interviewees, the gained insights are highly relevant. The respondents are experts in their field of duty, who work with the dashboards on a daily basis. Besides, openSAP is a big and professional MOOC platform with roughly $3,000,000$ enrollments in about 270 course (July 2018). Consequently, the views and opinions of the interviewees are considerable in this context.

The employees were asked for what purposes they use LA data in their daily work and how they utilize the dashboards for accomplishment of these tasks. In this context, special attention was paid to identifying parts of the dashboards that are essential and those that are not used at all by the individual persons. Additionally, the interviewees were asked for technical and conceptual issues as well as suggestions for improvements for the existing solutions. One problem that was mentioned by all is the performance of the pages. Especially when loading the course dashboard, it takes a lot of time until the page is eventually shown in the browser. In addition, the dashboards are usually visited frequently, which reinforces the issue. The reason for these long loading times is that the page is not rendered until all required LA data is loaded and visualized metrics and statistics are calculated. Three of the interviewees charged that the dashboards were cluttered. The pages contain many different visualizations and as explained previously, the majority of them is not relevant for all stakeholders. As a consequence, users might scroll over a number of components until they reach the visualization they were actually looking for. Especially long tables, such as referrer or social share statistics take up a lot of space, but are used only by certain users.

4) Revised Dashboard Concept: To meet the needs of the stakeholders and solve issues of the existing solutions, the entire concept of the dashboards is revised from different perspectives. On a structural level the general goal is to clean up the existing dashboards to simplify the access to metrics and statistics. The actual objective of dashboards in general is to visualize complex data in a simple way to provide a quick overview about a certain topic, in this case the performance of the platform or specific courses.

In the interviews it became clear that the existing dashboards contain a great number of different visualizations, whereas the majority of them is not relevant for all users. The new concept focuses on visualizations that are relevant for the majority of the users while providing possibilities to obtain extensive information on demand. As Key Performance Indicators (KPIs) are highly relevant for all interviewees, the corresponding sections of both dashboards are retained. In contrast, detailed visualizations built for special purposes are moved to separate pages, referred as *statistic pages* in the following. However, the information of the moved parts should still be represented in the dashboards. Therefore, for each statistic page there is a component in the dashboard visualizing the underlying data at a higher level, which takes up less space and is also easier to understand. At the same time, it serves as a link to the corresponding statistic page. For example, the list of social networks the courses have been shared in is moved to such a separate page. Along with this, the total number of course shares is added as KPI to the dashboards. In this way, users receive an overview about the performance of the entire platform or a specific course and can follow the links in case they are interested in more detailed information.

In terms of content, global and course dashboards are assimilated to each other. Both pages start with showing relevant KPIs and visualize the geographical locations, age distribution, client usage, and top referrers below. The course dashboard additionally shows the development of enrollments, forum activity, and helpdesk tickets over time. In addition, the age distribution illustrated as a bar chart is added to the dashboards as suggested by interviewees.

From a technical perspective all pages are rendered server-sided, but required analytics data is retrieved asynchronous in the client now. This approach has the advantage that the initial page is loaded quickly in the browser and the users already see the structure of the page while needed data is loaded in the background. To retrieve all data that is visualized in the dashboard, multiple API requests to different endpoints are necessary. These requests can be sent and processed in parallel. As a result, data is shown on the page as soon as it is received. Therefore, metrics that can be calculated more quickly are already visible in the user interface while more expensive operations are still running.

5) Integration of Web Analytics Metrics: In addition to revision of the existing concept, the new Google Analytics metrics are integrated into the dashboards. This is done by adding new visualizations and creating new statistic pages. The average session duration is added to the KPI section of the global dashboard. Besides, a heat map shows the average number of learners per hour and day of week on both dashboards. These two components additionally link to a new activity statistic page, which is available for both the global and course context as well. This page shows histograms of session durations and the number of days between two subsequent sessions. The corresponding metrics group values to buckets, which also ensures clarity and understandability of the visualizations. When viewing the statistic page in course context, these bar charts additionally show the platform average for each bucket, which makes it possible to compare the activity of a course with the average of all courses. Next to these two visualizations another heat map shows the number of active users for each day and hour in the last two months (global) or the course time frame. This visualization is similar to the heat map of the existing course dashboard that shows the temporal activity of users.

6) Summary: This section demonstrated how WA insights can be provided to stakeholders in a user-friendly way (RQ1.b). For this purpose, the structural and technical concept of the existing dashboards was revised first, since several issues in regard to their usability have been identified. Afterwards, the existing LA capabilities were enhanced by extending the dashboards by WA insights obtained from Google Analytics. Thereby, no differences between reporting of both kinds of insights were realized. The use of WA does not simplify visualizing LA data. Instead, it comes with the same challenges and problems. Communicating such complex data is always a difficult task and interpretations usually require expert knowledge. Thus, this issue is not changed by utilizing WA methods.

V. Evaluation

This section focuses on evaluating different aspects of this work. First, the usability of the revised and extended LA dashboards is evaluated on basis of a conducted survey. This includes answering RQ1.c by examining the usefulness of the implemented WA metrics. Afterwards, an answer to the main research question is given by discussing the applicability of WA in context of LA based on the implemented proof of concept and answered sub-questions.

A. Usability of Learning Analytics Dashboards

This section presents an evaluation of the usability of the implemented LA dashboards. By comparing the existing and revised dashboards in regard to usefulness, ease of use, and satisfaction, it is examined whether the usability could be improved by adding new WA metrics and refactoring the underlying technical concept.

1) Methodology: To achieve the goals of this evaluation, a survey was conducted addressing platform owners and teaching team members of openHPI and openSAP as they are the target group of the according dashboards. While platform owners were instructed to consider the global dashboards, teaching teams should compare the course dashboards. As the audience is really specific and thus small, only 11 respondents could be acquired for answering the questionnaire. However, the importance of participants' views and positions are still highly relevant as they are experts in their fields, who utilize LA insights in their daily work. Consequently, results of this survey are still meaningful despite the small number of respondents. The participants were asked to express their agreement with the following ten statements separately for the existing and the revised dashboard:

Q01 The dashboard helps me to monitor the activity of the platform or my courses.
Q02 The dashboard facilitates access to relevant metrics.
Q03 The dashboard meets my needs.
Q04 I regularly use the dashboard for my work.
Q05 The dashboard is easy to use.
Q06 The dashboard is understandable.
Q07 The dashboard loads fast.
Q08 The dashboard is clear and tidy.
Q09 The dashboard works the way I would expect.
Q10 I like to use the dashboard.

A symmetric Likert scale with the following four levels and corresponding scores was utilized for giving answers:

- Strongly agree (3)
- Somewhat agree (2)
- Somewhat disagree (1)
- Strongly disagree (0)

For evaluating the significance of differences, a Wilcoxon test was performed based on the answers' scores for each question. Additionally, effect sizes were computed with Cohen's d. These statistics can be found in Table II. The participants were also approached for qualitative feedback by means of free-text questions. This includes asking for suggestions for improvements. While the first two questions are considered in the following, the mentioned improvement suggestions are discussed later as they are part of the future work.

TABLE II
DESCRIPTIVE AND INFERENTIAL STATISTICS FOR USABILITY BEFORE AND AFTER REVISION OF THE DASHBOARDS

	Existing Dashboards		Revised Dashboards		Wilcoxon	
Q	Mean	Std.Dev.	Mean	Std.Dev.	p-value	Cohen's d
01	2.1818	0.6030	2.2727	0.6467	0.6547	0.1454
02	2.0909	0.7006	2.4545	0.8202	0.1573	0.4767
03	1.6364	0.6742	1.9091	0.7006	0.1797	0.3967
04	2.4545	0.6876	2.5455	0.5222	0.5637	0.1489
05	1.9091	0.9439	2.3636	0.8090	0.1025	0.5170
06	1.3636	0.6742	2.1818	0.8739	0.0235	1.0484
07	1.0909	0.5394	2.0000	0.7746	0.0152	1.3621
08	1.2727	0.6467	2.3636	0.8090	0.0097	1.4895
09	1.8182	0.6030	2.1818	0.7508	0.0455	0.5340
10	1.6364	0.5045	2.2727	0.7862	0.0196	0.9633

The survey results were analyzed in regard to two aspects of this work. First, the usefulness of WA insights in context of LA is evaluated for answering RQ1.c. Afterwards, the ease of use and satisfaction of the revised concept are assessed based on the answers of the participants to examine whether the goals of the revision could be accomplished.

2) Usefulness of Web Analytics Insights: From a content-related perspective, integrating the implemented WA metrics did not result in a significant difference of the usefulness ($p < 0.05$). The new metrics are not highly relevant for the majority of participants (Q02) and could neither help them in monitoring the activity on the platform or in courses (Q01), nor increase the satisfaction of their needs (Q03) in a significant extent. Furthermore, the additional insights did not lead to respondents planing to use the dashboards more often in a remarkably scope (Q04). The performance of the dashboards (Q07) could be significantly increased ($p = 0.0152$) with a large effect ($d = 1.3621$). This was achieved partly by replacing certain existing metrics with appropriate WA metrics, which can be retrieved faster. Consequently, the integration of WA contributed to this improvement. Nevertheless, the performance of the local analytics stores might also be boosted for example by upgrading the underlying hardware. However, this would cause additional costs for purchase and maintenance. Therefore, especially for a non-commercial project, such as openHPI, making use of the provided cloud infrastructure of WA tools, which is in case of Google Analytics even free of charge, is a meaningful decision to improve the performance of complex queries. The majority of the underlying metrics could also be covered by using WA methods. Consequently, WA could provide further insights, whose usefulness is already proven. Nevertheless, the effect of these insights cannot be measured in this context as there is no basis of comparison. Nevertheless, e-learning platforms that do not have such sophisticated LA capabilities as the HPI MOOC platform could benefit from integrating these WA metrics.

3) Ease of Use and Satisfaction of Revised Concept: Besides examining the usefulness of the implemented WA metrics, the survey also aims to ascertain the ease of use and satisfaction of the revised dashboard concept. In addition to the improved performance, whose effect was already determined previously, the redesign led to a significant increase ($p = 0.0235$) of the understandability (Q06) with a proven large effect ($d = 1.0484$). Besides, a highly significant difference ($p = 0.0097$) was found in regard to the clearness (Q08) with a large practical effect as well ($d = 1.4895$). In terms of satisfaction of participants, a significant improvement ($p = 0.0455$) with an intermediate effect ($d = 0.5340$) has been measured in the extent the dashboards work as expected by the participants (Q09). Additionally, the respondents also prefer working with the revised version (Q10) as a significant difference ($p = 0.0196$) of respective answers with a large effect ($d = 0.9633$) has been ascertained as well. However, the revision had no significant impact ($p = 0.1025$) on the simplicity of the dashboards (Q05). In addition to the quantitative evaluation, the participants were asked to mention aspects of the revised concept they like the most. The majority brought up the improved clarity caused by moving detailed statistics to separate pages and enabling the possibility to drill-down. Besides, also the fact that data for each chart is loaded independently was well received. Other mentioned aspects are the availability of new KPIs and charts, the improved visualizations, and the increased number of LA insights.

B. Applicability of Web Analytics

As shown in the previous sections, WA tools can be successfully integrated into a MOOC platform and thus utilized for LA purposes. A large part of relevant metrics can be retrieved using WA capabilities. However, there are some limitations, which is why this method can not be used exclusively in this context. This section examines to what extend WA tools can be utilized to gather insights in learning behavior on e-learning platforms. A great number of LA metrics correspond or can be mapped to WA metrics. For example, the session duration provides information about how long users are learning in one piece and page views indicate how often and in which order learning item are visited. Besides, the number of active users and characteristics of the audience, such as temporal access patterns, used clients and devices, and geographical origins, are relevant for both fields. Thanks to generic event tracking, which is supported by the majority of WA tools, any type of interaction can be tracked and therefore analyzed. As a result, especially LA KPIs can be calculated easily using WA as they usually just count the occurrences of a certain event type in a given date range.

In general, LA focuses on optimizing the learning experience on online platforms. This also includes improving the user experience as it is done by WA. However, the intention differs between both fields. While WA aims to help businesses in decision-making processes and is intended to increase revenue, users are in the center of LA as they should be supported while learning. This also includes encouraging individual learners for example by identifying users at risk that are likely to dropout soon and might therefore need special assistance. This is not an use case of WA and consequently corresponding tools do typically not support analyzing behavior of single users. Instead, only metrics regarding user segments (e. g. mobile users or users of a certain country) or the entire user base can be accessed.

In WA, sessions are the central element for analyses. The amount of metrics that are calculated across subsequent sessions of the same user are very limited in the majority of existing tools. For instance, conversion rates and e-commerce metrics in Google Analytics are restricted to sessions. In contrast, a learning process usually extends over a long period of time. This fact also applies to online learning. For example, the majority of MOOCs has a length of six weeks, in which contents are typically published gradually. Consequently, learning takes place in a great number of sessions and perhaps on multiple devices. To analyze the entire learning process of users within a course, activity data has to be considered across sessions. For these purposes, WA can not be utilized with its current set of features.

To sum up, certain aspects of LA can also be accomplished by utilizing WA. However, the aptitude of a metric for being retrieved using WA strongly depends on the type of stakeholder it is intended for. As explained before, openHPI has four different stakeholders: *platform owners*, *teaching teams*, *learners*, and *researchers*. While researchers are interested in any kind of LA data, the needs of the other three differs. For platform owners metrics concerning the overall performance of the platform are relevant, which is why they most closely correspond to the typical user of WA. Consequently, the majority of metrics relevant for this role can be queried using WA as well. Especially highly aggregated metrics, such as KPIs, can be easily obtained this way, but are nevertheless essential for platform owners. When it comes to teaching teams, WA can be used only partly. As long as information about the general activity and progress of a course should be gathered, it still works fine. However, it quickly reaches its limitations when trying to retrieve metrics about smaller groups of users sharing certain characteristics or even single users. As teaching teams are responsible for supervising and supporting learners of a course, WA can help them only with a fraction of their duties. Finally, the method can be used only in a minor extent to provide individual learners with LA data. Metrics about single users, which are most important in this context, can usually not be obtained. Nevertheless, there are some use cases where WA is helpful. For example, information about the average performance of learners can help individuals to reflect on their own performance. All in all, WA mainly meets the needs of platform owners, but can assist teaching teams only in some cases. For individual learners there are only rare cases where WA data is relevant.

While this evaluation was done based on the proof of concept integration of Google Analytics into the HPI MOOC platform, the key findings also apply to WA and e-learning platforms in general. The revealed potentials and limitations of WA are not specific to the HPI MOOC platform. This is mainly because characteristics of learning concepts and processes as well as the needs of stakeholders in regard to their use of LA insights are similar among e-learning platforms.

Besides, core concepts and features are the same for the majority of WA tools. Differences exist only in regard to more specific and advanced analysis capabilities, processing limitations, and pricing models. As this evaluation considers the general concepts of WA instead of concrete features of Google Analytics, the findings also hold for other tools and are thus valid for the field of WA in general.

VI. Conclusion

So far, WA has not been profoundly used to analyze learners' behavior on e-learning platforms. However, LA may benefit from this sophisticated and well-established method. Therefore, the goal of this paper was to examine how WA can be utilized for LA purposes and what limitations it has in this context (RQ1).

To answer this question, Google Analytics was integrated into the HPI MOOC platform as a proof of concept to evaluate the applicability of WA tools in the context of large-scale online learning. For this purpose, the platform's *lanalytics* service was extended by another processing pipeline that transforms captured interaction data according to the schema defined by Google Analytics and sends it to the WA service. For this transformation a mapping was developed that models learning activity as WA hits (RQ1.a). Instead of using the typical data collection technique of page tagging, the Measurement Protocol is utilized to transmit hits from the platform's backend. Consequently, the solution took advantage of the existing event tracking engine resulting in consistency between the local and external analytics stores. Besides, this approach reinforces data privacy as the amount of data sent to third parties is selected manually.

Based on the captured data, it was considered how gathered insights can be provided to stakeholders of the platform so that they can make use of it (RQ1.b). Thereby, LA dashboards were chosen for evaluation purposes as they are a typical use case of analytics in general. This paper focused on dashboards intended for platform owners and teaching teams as these were already existing, which enabled comparing the usefulness of WA metrics (RQ1.c). Requirements of the stakeholders and issues of the existing solutions were identified by conducting expert interviews. Based on these insights, the concept of the dashboards was revised to improve their usability. Afterwards, they were extended by newly implemented WA metrics, which were integrated into the existing architecture of the *lanalytics* service. In doing so, no differences between visualizing WA and LA insights were recognized. Instead, both types of metrics were handled exactly the same as the dashboards do not distinguish different data sources in this context.

The evaluation of the usability of the revised and extended dashboards showed that the newly implemented WA metrics do not have a statistically significant impact on the usefulness of the dashboards. However, a large part of the existing LA capabilities, which are indeed proven to be useful for the stakeholders, could also be realized by using WA methods. Consequently, WA can still provide useful insights in the context of LA. In addition, it can contribute to an increase of the general performance by making use of the cloud infrastructure of WA tools (RQ1.c). Besides, the revision of the dashboards had a significant impact on the ease of use and satisfaction. Especially the clearness could be improved on a large scale as shown by the quantitative, but also qualitative evaluation. Consequently, the intentions and goals of the revision have been accomplished.

To finally answer the main research question (RQ1) of this paper, the proof of concept integration was evaluated in regard to the applicability of WA for gathering LA insights. It was shown that WA can indeed be used to retrieve a large part of metrics relevant in context of LA. However, its applicability highly depends on the type of stakeholder the corresponding metrics are intended for. The needs of platform owners of e-learning platforms and websites in general do not differ much. Consequently, the majority of insights relevant for this role can be retrieved using WA. Especially KPIs are essential for this type of stakeholder and can easily be obtained from WA tools. Nevertheless, when it comes to teaching teams, the technique can be utilized only to a limited extent. While WA can provide an overview of the general performance of a course, it reaches its limitations when considering learner-specific metrics since WA is not designed for retrieving user-level information. Consequently, it is also not suitable for providing LA data to individual students. In this context, WA might be used only to support self-reflection by providing information about the average performance on the platform as a point of reference. Besides, more advanced features of WA tools, such as e-commerce analysis, can not be utilized in context of LA due to a mismatch of concepts.

Additionally, there is still room for further extensions and improvements. So far, the integration of Google Analytics is deep-seated in the platform's Service Oriented Architecture. Nevertheless, also other e-learning platforms could benefit from integrating WA capabilities in a similar way. For this purpose, it would be necessary to generalize the implementation, e. g., by providing a mapping interface from the xAPI standard to the Google Analytics Measurement Protocol. The dashboards revised in the context of this paper are intended for platform owners and teaching teams. However, there are also other stakeholders for whom LA data can add value. Individual learners can directly benefit from LA as well. There is much research going on concerning student-facing dashboards and how they can be utilized to improve the learning experience in MOOCs. As already mentioned, at the moment there is no LA dashboard specifically for learners implemented in the platform. The existing tracking and reporting capabilities could be used as a foundation for such a feature. While WA is suited only in a minor extent for this purpose, the local analytics stores can be used to obtain relevant metrics for learner dashboards. In this way, students' self-reflection, awareness, and self-assessment can be encouraged.

It can be concluded that e-learning platforms can benefit from utilizing WA for improving the learning experience. By integrating Google Analytics into the HPI MOOC platform, the existing LA capabilities could be extended by new, valuable insights in learners' behavior. Thereby, WA works well for obtaining an overview about the general activity on the platform or within single courses. However, for receiving insights in the behavior individual learners, WA is not applicable. For these purposes, LA-specific methods need to be utilized. This might change with further development of WA and corresponding tools.

References

[1] E. Brynjolfsson, L. M. Hitt, and H. H. Kim, "Strength in Numbers: How Does Data-Driven Decisionmaking Affect Firm Performance?" *SSRN eLibrary*, 2011. [Online]. Available: http://ssrn.com/paper=1819486

[2] B. Romanowski and A. Konak, "Using Google Analytics to Improve the Course Website of a Database Course," 2016. [Online]. Available: https://www.hofstra.edu/pdf/academics/colleges/seas/asee-fall-2016/asee-midatlantic-f2016-konak.pdf

[3] A. Cooper, "Learning Analytics Interoperability - The Big Picture In Brief," *Learning Analytics Community Exchange*, 2014. [Online]. Available: http://laceproject.eu/publications/briefing-01.pdf

[4] C. Meinel and C. Willems, "openHPI: the MOOC offer at Hasso Plattner Institute," Tech. Rep., 2013. [Online]. Available: https://publishup.uni-potsdam.de/frontdoor/index/index/docId/6548

[5] B. Moissa, L. S. de Carvalho, and I. Gasparini, "A Web Analytics and Visualization Tool to Understand Students' Behavior in an Adaptive E-Learning System," in *Learning and Collaboration Technologies. Designing and Developing Novel Learning Experiences*, 2014.

[6] H. Luo, S. Rocco, and C. Schaad, "Using Google Analytics to Understand Online Learning: A Case Study of a Graduate-Level Online Course," in *International Conference of Educational Innovation through Technology (EITT)*, 2015.

[7] edX Inc., "Open edX Developer's Guide - Analytics." [Online]. Available: http://edx.readthedocs.io/projects/edx-developer-guide/en/latest/analytics.html

[8] J. Renz, G. Navarro-Suarez, R. Sathi, T. Staubitz, and C. Meinel, "Enabling Schema Agnostic Learning Analytics in a Service-Oriented MOOC Platform," in *Proceedings of ACM Learning at Scale Conference (L@S2016)*. ACM, 2016.

[9] T. Rohloff, M. Bothe, J. Renz, and C. Meinel, "Towards a Better Understanding of Mobile Learning in MOOCs," in *Learning with MOOCs 2018 (LWMOOCS V)*. IEEE, 2018.

[10] Google Inc., "Measurement Protocol Reference." [Online]. Available: https://developers.google.com/analytics/devguides/collection/protocol/v1/reference

[11] Pivotal Software Inc., "RabbitMQ: Reliability Guide." [Online]. Available: https://www.rabbitmq.com/reliability.html

[12] Google Inc., "User Explorer." [Online]. Available: https://support.google.com/analytics/answer/6339208

[13] ——, "Dimensions & Metrics Explorer." [Online]. Available: https://developers.google.com/analytics/devguides/reporting/core/dimsmets

[14] ——, "Enhanced e-commerce reports." [Online]. Available: https://support.google.com/analytics/answer/6014872