Real User Monitoring (RUM)

A Comparative Survey

Niklas Lorber, Maria Seiser, and Kilian Weisl

706.041 Information Architecture and Web Usability 3VU WS 2024/2025
Graz University of Technology

13 Dec 2024

Abstract

Real User Monitoring (RUM) is an essential technology for understanding user behaviour, optimising website performance, and enhancing user experience. However, selecting the right RUM tool can be overwhelming, due to the vast differences in features and pricing. This survey presents a comprehensive overview of the current landscape of RUM tools, categorising them into commercial and non-commercial solutions. It looks in detail at three state-of-the-art tools, Hotjar, Heap, and Clarity, describing their integration approaches, unique functionality, and practical applications. Examining aspects such as feature richness, usability, cost, and performance impact, the survey offers actionable guidance for selecting a RUM solution that balances feature relevance, ease of implementation, and cost effectiveness.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

Real User Monitoring (RUM) collects and measures real user interactions to observe user behaviour, detect potential issues, and assess user satisfaction. The supply of tools that offer RUM solutions is vast, but nevertheless heterogeneous. This survey explains what RUM is, what features it entails, and how different tools perform when implemented on a website with real user traffic.

1.1 Real User Monitoring (RUM)

RUM monitors and takes measurements from real user interactions. Typically, RUM applications follow the pattern illustrated in Figure 1.1. RUM tools are usually integrated into a web application by adding some JavaScript Code. For the actual users who visit the website the pages remain unchanged, only the various measurements are now also sent to an aggregator for collection and analysis. Table 1.1 collects some of the terminology associated with RUM.

1.1.1 Real vs. Synthetic Measurements

Contrary to popular belief, RUM is not only about how users interact with a website or some performance measurements. Real User Monitoring, first and foremost, emphasises that a real user provides the measurements. Typical data collected include page load time, latency, or Nielsen ratings. Synthetic Measurements are the opposite of RUM. Real users do not generate them. Instead, interactions are simulated by a computer, often elsewhere on the network using tools such as Application Performance Monitoring (APM) [Mastin 2016, Chapter 1].

1.1.2 Active vs. Passive Measurements

Active Measurements are controlled experiments that try to influence a site's performance and adaptability by artificially introducing stress or more traffic to the system. With active measurements, RUM can run controlled experiments on hypothetical scenarios to observe what problems might occur. Passive Measurements are just observational studies conducted on a website's existing traffic. With passive measurements, RUM can detect what problems are happening on the site right now, it can detect errors in real-time.

1.2 Benefits and Limitations of RUM

RUM tools offer significant benefits by providing precise and accurate metrics, ensuring a clear understanding of system performance from the user's view [Mastin 2016, Chapter 2]. These measurements, taken directly from the point of consumption, deliver highly relevant and actionable insights. RUM enables its users to observe real-time alerts for errors and track issues as soon as they arise [Croll and Power

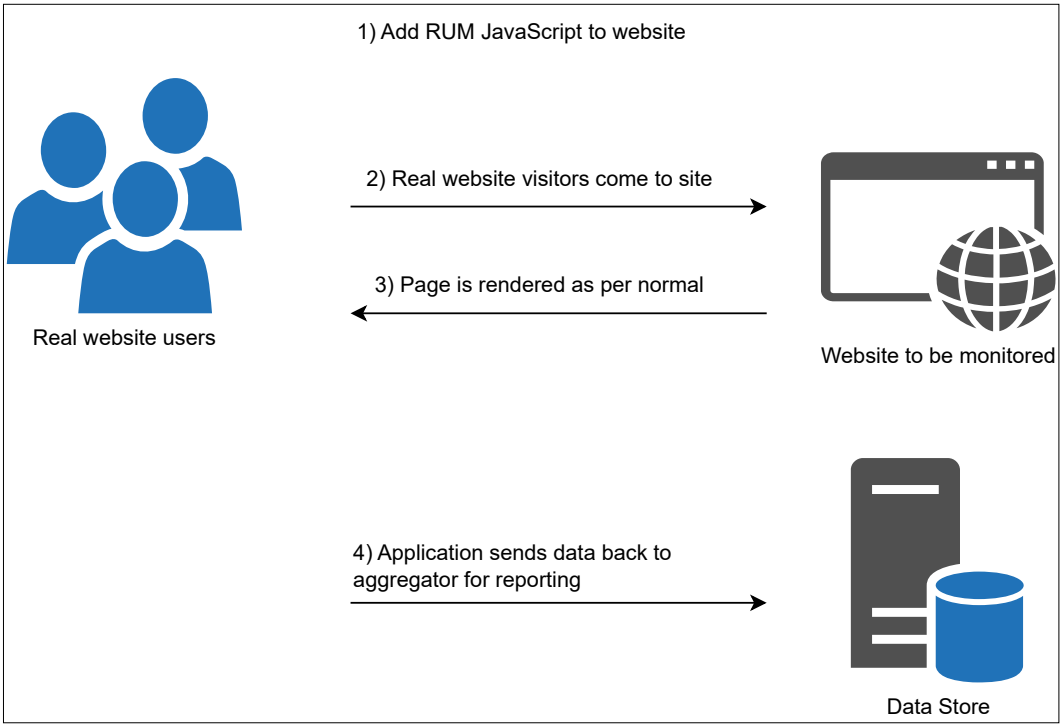

Figure 1.1: The typical flow of a RUM tool comprises adding JavaScript to the target website, having real users visit the site, rendering the page normally, and collecting various measurements in a data store. [Adapted by the authors from Figure 1-2 of Mastin [2016].]

Term	Definition
DEM	Digital Experience Monitoring
RUM	Real User Monitoring (or Measurement)
Synthetic Measurement	Measurement with synthetic user interactions.
Active Measurement	Measurement taken within a controlled experiment.
Passive Measurement	Measurement taken based on observation.
APM	Application Performance Monitoring
LCP	Largest Contentful Paint
INP	Interaction to Next Paint
CLS	Cumulative Layout Shift
CSAT	Customer Satisfaction
RTUT	Real-Time User Tracking
PLT	Page Loading Time

Table 1.1: Common terminology in RUM.

2009]. Additionally, the substantial volume of data generated offers a robust foundation for in-depth analysis, driving informed decision-making and continuous optimisation of the observed application.

Despite its benefits, RUM also has some limitations. Implementing RUM to make meaningful observations requires a substantial number of users to visit the target website, which generates a significant amount of data that must be stored somewhere. Integrating RUM tools into the website can sometimes also lead to page-speed reductions that influence the user experience. Monitoring and storing the data of users visiting a website needs to be handled carefully. Global and national laws regarding data privacy need to be taken into account.

Chapter 2

Canonical Features of RUM Tools

The variety of features available across existing tools and inconsistencies in their definition made establishing a canonical set of features necessary to ensure a clearer understanding of RUM and facilitate better comparison across the investigated tools. The features were chosen because they are widely prevalent, recognized as essential by most tools, or represent novel and innovative functionality that adds unique value. To describe the features, it is necessary to distinguish between the end user, who visits the website under investigation, and the user of a RUM tool who investigates a website. Here, the term "user" will be used for the former and "investigator" for the latter.

2.1 Real Time User Tracking (RTUT)

If a tool supports real-time user tracking, it can follow the live actions of a particular user on the investigated website. This functionality is implemented either by an event log or graphically.

In an events-only RTUT approach, a user can only be followed live by a textual description of events. This approach makes it impossible to follow the live session from the actual user's perspective, who can see the website under investigation in their browser. The investigator can only see textual descriptions of what a user does on the website. There is either a live view for all events happening on the investigated website or a live view for a specific user, typically uniquely identified by an ID. Examples of these events could be:

- Click on `i.fas.fa-shopping-cart`.
- View `/shop/products/xyz/`.
- Submit on `form.checkout.woocommerce-checkout`.

In a graphically supported RTUT approach, the investigator can see the website as the user sees it. RUM tools which support this kind of RTUT can also show the site at the exact size of a user's device. As a user interacts with the site, the investigator can see what a user does, including scrolling through pages, clicking links, and filling out forms in real time. To protect sensitive and personal information, the tool automatically detects form input fields which are not monitored and these are hidden from the investigator.

2.2 Session Replay

Session replay refers to the ability of a RUM tool to show the user's interaction during completed and, therefore, non-live sessions. Similarly to RTUT, the investigator can see which elements the user interacts with on the website.

2.3 Error Tracking

All modern websites use JavaScript to create a more interactive experience for users. Since JavaScript runs client-side in a user's web browser, JavaScript errors are usually not reported back to the server and can therefore be overlooked. This can lead to unwanted or undefined behaviour on a website.

Since RUM tools can monitor a complete user session, JavaScript errors can also be caught and reported by the tool. This way, a website investigator can see which errors happened. Combined with RTUT or session replay, it is possible to recreate the complete environment in which the error occurred. Therefore, providing solutions and fixing errors on the website under investigation is easier.

2.4 Real Performance Measurements

Real performance measurements are standardised metrics for essential aspects of a specific user's perceived speed. Since loading time is crucial for a successful website, investigators are interested in how their application or website performs under different circumstances. Users may have different loading times depending on internet speed, location, and the website server's location. Google recommends three performance metrics known as the Core Web Vitals [Google 2024a], which can be measured with free tools, such as Google PageSpeed Insights [Google 2024b]:

- *Largest Contentful Paint (LCP)*: LCP measures when the main content is loaded: "LCP reports the render time of the largest image, text block, or video visible in the viewport, relative to when the user first navigated to the page." [Walton and Pollard 2024]. This includes the Time-To-First-Byte (TTFB), i.e. any unloading time from a previous page, redirect time, and connection set-up time. Since JavaScript frameworks and browsers can already show parts of a website before it is fully downloaded, the full loading time of a website is no longer a meaningful performance metric when it comes to the perceived speed of a website.

- *Interaction to Next Paint (INP)*: INP is used to measure the responsiveness of a website: "INP observes the latency of all interactions a user has made with the page and reports a single value, which shows all (or nearly all) interactions beneath. A low INP means the page was consistently able to respond quickly to all – or the vast majority – of user interactions." [Wagner 2024]

- *Cumulative Layout Shift (CLS)*: Many websites today use JavaScript to add content to the user's screen after the page has fully loaded. This may negatively affect a user's perceived experience. CLS measures how often users experience unexpected layout shifts [Mihajlija and Walton 2023].

2.5 User Paths

This feature is handy when it comes to the navigation on a website. More often than not, a user's first-visited page is not the home page. The intended navigation flow of a website may differ entirely from the actual navigation of a user. As an investigator of a website, it is helpful to see which page a user first visits and how the website's navigation is used. The user path enables investigators to see which pages have been visited by a user in which order.

2.6 User Segmentation

User segmentation allows an investigator filter data according to specific criteria or behavior. Some examples of user segmentation include:

- Users from country X.

- Mobile/Tablet/Desktop browser sessions.

- Users who visited page Y and clicked button Z.

2.7 User Surveys

An investigator might be interested in eliciting feedback from a user at the end of a particular workflow, for example after the checkout process of an e-commerce website. RUM tools implementing this feature can create surveys that appear to a particular user after a particular event. They collect user feedback and report data to the investigator. This is especially useful when combined with behavioral user segmentation.

2.8 User Satisfaction Metrics

User satisfaction metrics are used to derive assumptions about the user experience of a website, for example:

- *Rage Clicks*: When a user clicks repeatedly on the same button or element [Hotjar 2024e].
- *U-Turns*: When a user returns directly to a previous URL within a few seconds [Hotjar 2024a].
- *Exit Rate*: Measures the percentage of users who leave a specific page or website [Contentsquare 2024].
- *Bounce Rate*: Determines the percentage of users who visit a page and directly leave the site again without visiting any other pages [Wikipedia 2024a].

RUM tools can observe the complete interaction of a user with a website and can thus calculate these metrics.

Chapter 3

Overview of RUM Tools

Various RUM tools are available, both commercial and non-commercial, and closed-source and open-source. In an initial discovery phase, a number of tools were investigated for each category to determine whether they might be suitable candidates for more detailed exploration and trial implementation on a real website. At this stage, the tools were compared using published documentation and reviews. Criteria included available features (see Chapter 2), cost, usability, integration options, and scalability. For each tool, its key features, target audience, and unique selling points were assessed.

3.1 Commercial Tools

A wide variety of commercial RUM tools are available. They are typically characterised by an extensive range of functions and intricate and complicated payment plans. Despite their expensive and sometimes overloaded nature, these tools are impress with their simplicity to use and comprehensive support. They are targeted at enterprises with the budget to invest in high-quality solutions. Six commercial tools were considered in the discovery phase of the survey: Hotjar, Heap, Calibre, Sematext, Dynatrace, and UXCam.

3.1.1 Hotjar

Hotjar is a user-friendly tool designed for businesses of all sizes [Hotjar 2024d]. It excels at insights into user behaviour through features such as heatmaps and session recordings, and user feedback tools like surveys and questionnaires, making it a fit for teams focused on improving user experience and website usability. Hotjar was one of the first tools focused on RUM and is one of the largest players in the industry. The key characteristics of Hotjar are:

- *Key Features*: Engagement heatmaps, session recordings, surveys, and user feedback collection.
- *Target Audience*: Small to medium businesses, product teams, and UX designers.
- *Unique Selling Points*: Easy onboarding, intuitive interface, transparent pricing, and highly customisable surveys and questionnaires.

3.1.2 Heap

Heap is a RUM tool that automates data collection and analytics for user interactions [Heap 2024d]. It enables application providers to capture every user action on their website or application without requiring manual tracking. It also allows for real-time event tracking. The key characteristics of Heap are:

- *Key Features*: Automatic event tracking, funnel analysis, user segmentation, and advanced behavioural analytics.
- *Target Audience*: Product teams, marketing teams, and data-driven businesses.

- *Unique Selling Points*: Automatic tracking of all user interactions, intuitive interface, rich documentation, and customer support.

3.1.3 Calibre

Calibre is a RUM tool focusing on website performance and optimising user experience [Calibre 2024]. It provides detailed insights into page load times, Core Web Vitals, and other performance metrics out of the box, making it an excellent choice for developers and performance enthusiasts. The key characteristics of Calibre are:

- *Key Features*: Core Web Vitals monitoring, real-time performance tracking, and synthetic testing.
- *Target Audience*: Developers, web performance teams, and digital marketers.
- *Unique Selling Points*: Out-of-the-box Core Web Vitals and customisable alerts for performance surveillance.

3.1.4 Sematext

Sematext is an all-in-one RUM and performance monitoring tool [Sematext 2024], designed for organisations seeking infrastructure and application performance visibility combined with RUM. It combines real user monitoring with log management and synthetic monitoring for a more complete approach. The key characteristics of Sematext are:

- *Key Features*: End-to-end RUM, log management, infrastructure monitoring, synthetic testing, and performance monitoring.
- *Target Audience*: IT teams, DevOps, system architects, and enterprises.
- *Unique Selling Points*: All-in-one platform for performance monitoring, log management, RUM, and infrastructure monitoring.

3.1.5 Dynatrace

Dynatrace is a comprehensive RUM tool designed for enterprise-level monitoring and performance optimisation [Dynatrace 2024]. It offers AI-driven insights, full-stack monitoring, and real-time analytics to help organisations manage complex digital ecosystems. The key characteristics of Dynatrace are:

- *Key Features*: AI-powered anomaly detection, RUM, synthetic monitoring, server-side monitoring, and application security features.
- *Target Audience*: Enterprises, IT teams, and performance optimisation specialists.
- *Unique Selling Points*: Scalability, deep and sophisticated integration capabilities, and automation capabilities for problem tracking and resolution.

3.1.6 UXCam

UXCam is a RUM tool specialising in mobile application analytics [UXCam 2024]. It provides deep insights into user behaviour on mobile platforms, enabling app developers to identify friction points and improve the user experience. The key characteristics of UXCam are:

- *Key Features*: Session recordings, heatmaps for mobile gestures, funnel analysis, and crash analytics.
- *Target Audience*: Mobile app developers, product managers, and UX designers.
- *Unique Selling Points*: Focused on mobile app analytics, detailed visualisation of user interactions, and actionable insights to optimise mobile experiences.

3.1.7 Discussion

A plethora of commercial tools for RUM are available, with probably more to come as the need for insightful data analysis and investigation only accelerates the trend for RUM tools in a commercial context. Serious businesses need robust and reliable commercial solutions to understand their users and data to form comprehensive and impact full decisions for their companies' success.

Using a commercial solution for RUM offers the following benefits:

- Comprehensive support and frequent updates.
- Advanced analytics and robust reporting.
- Comprehensive and extensive feature set.
- Frequent updates and improvements.
- Reliability and service level agreement.

When deciding to use a commercial solution for RUM, the following disadvantages should be considered:

- High cost and complex payment plans.
- Potential vendor lock-in.
- Data ownership concerns.
- Limited trial periods.

3.2 Non-Commercial Tools

Far fewer non-commercial RUM tools are available than commercial tools, and their functionality, documentation, and support are often more limited. These tools appeal to smaller projects and companies with smaller budgets. Three non-commercial tools were considered in the discovery phase of the survey: Clarity, OpenObserve, and Basic RUM.

3.2.1 Clarity

Clarity is a free-to-use, non-commercial, closed-source RUM tool by Microsoft [Microsoft 2024], which provides insights into website user behaviour. It offers tools for session recordings, heatmaps, and advanced analytics to help webmasters and UX designers understand and improve user interactions. The key characteristics of Clarity are:

- *Key Features*: Heatmaps, session recordings, click and scroll tracking, out-of-box web vitals, and traffic analytics.
- *Target Audience*: Small businesses, website owners, non-profit organisations, startups, and UX designers seeking to improve website usability.
- *Unique Selling Points*: Completely free to use, no traffic caps, seamless integration with other Microsoft services, and GDPR-compliant data collection.

3.2.2 OpenObserve

OpenObserve is an open-source, cloud-native observability platform [OpenObserve 2024], which handles logs, metrics, traces, and basic RUM functionality. It offers a unified solution for monitoring and analysing system performance and user interactions. The key characteristics of OpenObserve are:

- *Key Features*: Logs, metrics, traces, basic RUM support, and cost-efficient scaling architecture.

- *Target Audience*: Developers, DevOps teams, enterprises, and open-source enthusiasts seeking customisable monitoring tools.
- *Unique Selling Points*: Open-source and free to use, community-driven development, real-time observability, scalable architecture, and an intuitive GUI for data visualisation.

3.2.3 Basic RUM

Basic RUM is another open-source RUM tool [Basic RUM 2024], designed to help enthusiasts, consultants, and developers visualise and analyse trends in performance data. It integrates Boomerang JS for data collection, ClickHouse DB for storage, and Grafana for visualisation. However, extensive technical knowledge is required to properly set up and maintain the system in production environments. The key characteristics of Basic RUM are:

- *Key Features*: Collection of performance metrics, efficient data storage, customisable dashboards, sophisticated monitoring of Core Web Vitals, and cost-effective scalability.
- *Target Audience*: Web performance enthusiasts, consultants, operations teams, front-end developers, and organisations seeking an open-source RUM solution focusing on expandability.
- *Unique Selling Points*: Completely open-source, capable of handling high-traffic websites, flexibility to customise and extend functionalities, and community-driven development.

3.2.4 Discussion

There are far fewer non-commercial tools than commercial tools. However, these should not be neglected, as they still have many use cases depending on the project and requirements. Especially when entering the broad field of RUM, starting with a tool like Clarity can make sense to keep costs low. However, the support and ongoing operation of complex tools such as basic RUM can sometimes be underestimated. In addition, hosting should not be overlooked for self-hosted open-source projects, as these tools generate a large amount of traffic and data. The dependence on the community behind the open-source projects should also be kept in mind, particularly if maintenance and development cannot be handled in-house. For some companies, however, it may be essential to use self-hosted, open-source systems, especially if data sovereignty is to remain within the company.

Using a non-commercial solution for RUM offers the following benefits:

- Free to use.
- Focus on essential core features.
- Ideal for startups, small businesses, or educational purposes.
- Potential for expandability (open-source).
- Data sovereignty (self-hosted).
- Vendor-independence.

When deciding to use a non-commercial solutions for RUM, the following disadvantages should be considered:

- Limited functionality compared to commercial tools.
- Lack of dedicated customer support.
- Inconsistent updates and maintenance, dependent on community contributions.
- Higher initial setup complexity.

- Potential security risks.
- Requires significant in-house expertise for customisation and optimisation.
- Scaling infrastructure for high traffic may require additional resources and manpower.

3.3 Selected Tools

As can be seen, many different RUM tools are available, each with unique characteristics and strengths. Particularly in commercial tools, there are plenty of alternatives, each trying to position itself in the relevant niche with unique features.

Three RUM tools were chosen for more detailed comparison by integrating them into a real website. Hotjar and Heap were chosen from the commercial realm, due to their popularity and the fact that they can be tested using a test license while still providing most of the relevant features. Clarity was chosen for the non-commercial realm, as no immense infrastructure and setup costs are needed compared to the open-source solutions. However, the non-commercial open-source tools and the other alternatives should not be neglected if one wants to seek a reliable RUM tool for a project.

Chapter 4

Comparison of Three RUM Tools

The three RUM tools selected in the discovery phase described in Chapter 3, namely Hotjar, Heap, and Clarity, were then investigated in more detail by installing them and trying them out with a real working e-commerce web site The canonical feature set established in Chapter 2 was used to assess and compare the three tools. This chapter provides a comprehensive overview of the setup process for each tool, detailing the steps involved in configuring them for use, highlights the unique features of each tool, and provides an analysis of their strengths and weaknesses.

4.1 Hotjar

Hotjar is one of the pioneering RUM solutions, launched in 2014, and acquired by Contestsquare in 2021. It is used by over 1,306,323 websites in over 180 countries [Hotjar 2024b]. The focus of Hotjar is user-centred design and analysing not only the interactions but also the user experience, which makes it unique among the countless tools that exist for monitoring users. Hotjar offers nearly all of the essential features associated with RUM and provides additional exciting components for conducting user surveys. Hotjar is a commercial tool that operates on a subscription-based model. However, it does offer a free trial period of 14 days, which made its implementation possible and allowed in-depth observation of its features for this survey.

4.1.1 Setup

The setup of Hotjar is well-documented and easy to carry out. The implementation requires two simple steps, after which the tool will start monitoring the target website and collecting data from users visiting the page. The onboarding for the tool's main functionality provides an intuitive step-by-step guide to new users. This yields a fast understanding of the tool and serves as a low threshold for new users to utilise all of the features available to them. In particular:

- *Integration*: Hotjar generates a unique JavaScript code snippet called a "Tag"; an example is shown in Figure 4.1. The Tag code must be placed in the `<head>` section of every page one wants to track. After integrating the JavaScript code on the website, users must verify the site they want to observe by entering the URL and clicking a button. Hotjar is thenactive and running.

- *Onboarding*: To explore all of the features, Hotjar provides an extensive step-by-step guide to their users. The guide provides a link to each feature, instructions on how to collaborate with colleagues, and how to integrate other tools like Google Analytics into Hotjar.

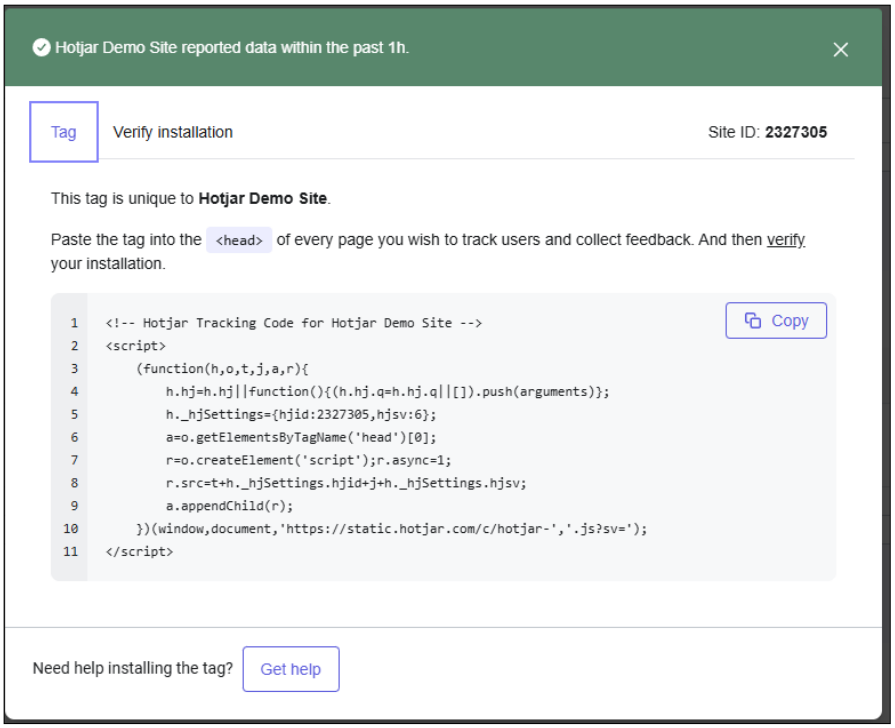

Figure 4.1: Hotjar: An example of the JavaScript Tag code snippet. [Screenshot taken by Maria Seiser]

4.1.2 Feature Set

Hotjar provides its users with extensive features and covers nearly all of the canonical features described in Chapter 2. Compared to Heap and Clarity, Hotjar offers additional unique features for monitoring the user experience and behaviour.

4.1.2.1 Real-Time User Tracking (RTUT)

Although Hotjar offers a wide range of features for real user monitoring, it currently does not support real-time tracking of user interactions. Following a live user is the only feature that cannot be achieved by using this tool.

4.1.2.2 Session Replay

The session replay feature of Hotjar can be found under the tab Recordings. There, all user sessions are listed alongside some metadata automatically extracted from the sessions. The table of user session recordings in Hotjar can be sorted by the following categories:

- *Relevance*: A score that predicts how relevant a session might be based on the user's activity.
- *Date*: Day and time of the recorded session.
- *User*: A user id to identify the same user's session.
- *Country*: The country where the user is currently located.
- *Action #*: The number of actions the user took.
- *Pages #*: The number of pages visited during the session.
- *Duration*: Amount of time the user spends on the website.
- *Device Type*: The type of device (desktop, mobile phone, tablet) the user had.

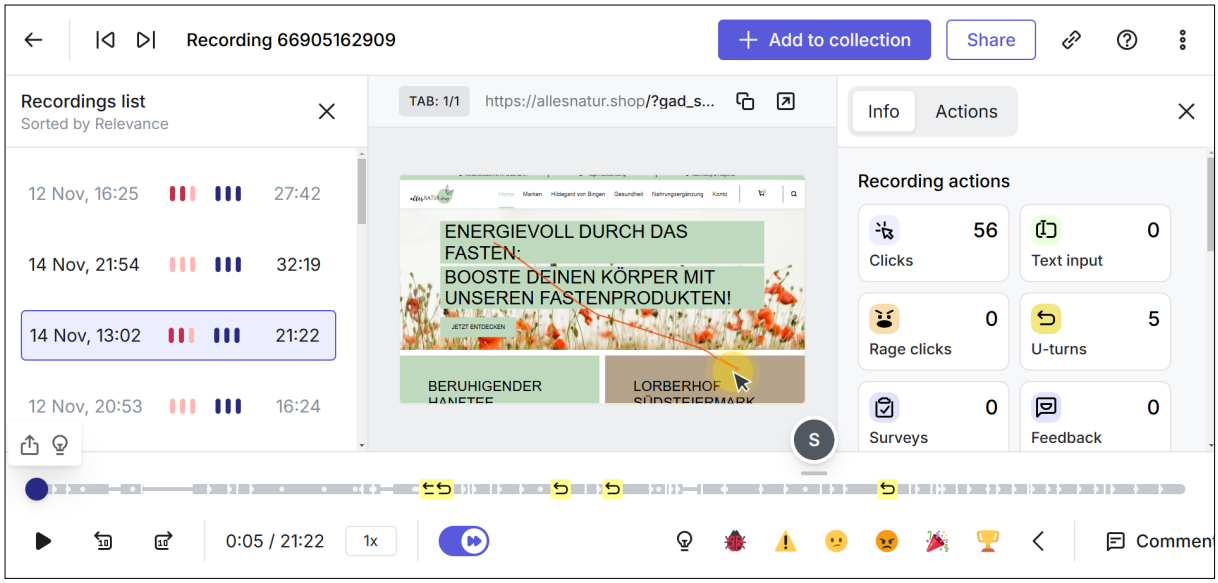

Figure 4.2: Hotjar: Replaying a user session. [Screenshot taken by Maria Seiser]

- *Browser*: The browser type used.
- *Operating System*: The operating system of the device.
- *Landing Page*: The first page the user visited.
- *Exit Page*: The last page the user visited.

Additionally, Hotjar also marks entries that have been identified as highlights by a team member. As shown in Figure 4.2, the page to view individual recordings is packed with features to investigate the session further. It features a list of timestamps for relevant actions, different categories for actions that occurred or can occur, a feature to collaborate with colleagues and mark actions within a session as highlights or errors, and an error log.

4.1.2.3 Error Tracking

Hotjar extracts substantial metadata from the individual user sessions. The sessions record the user interaction and if a JavaScript code on the webpage throws an error. This is extremely helpful for RUM use cases where debugging and accurate error tracking is essential.

4.1.2.4 Real Performance Measurements

Hotjar does not monitor performance measurements itself, but they can be included by integrating tools like Google Analytics [Google 2024a].

4.1.2.5 User Paths

The tool allows investigators to explore user navigation and behaviour. Using Funnels, paths can be analysed and compared by looking at drop-off rates and conversion times. Investigators can filter for pages visited, but also for individual components interacted with, and then analyse the sessions where those interactions happened.

4.1.2.6 User Segmentation

Hotjar provides extensive options for filtering and grouping users by various categories, such as new users, returning users, mobile users, etc. Features like Funnels allow for filtering based on behaviour, such as users who clicked a specific button or visited a specific page after visiting the main page.

4.1.2.7 User Surveys

In contrast to Heap and Clarity, Hotjar is the only tool observed in this study that offers user survey integration. Alongside the option to create individual surveys, Hotjar offers the following user survey templates:

- *Feedback Button*: A simple button on the side that opens up a feedback form with a 1 to 5 scale for the user to rate their experience.
- *NPS Survey*: A feedback form based on the market research metric Net Promoter Score [Wikipedia 2024b] to measure long-term loyalty.
- *Exit Intent Survey*: A feedback form that captures feedback from users before they leave the website. This can help in reducing the bounce rate.

4.1.2.8 User Satisfaction Metrics

Hotjar offers bounce rate, and rage click metrics. It also calculates a frustration score and an engagement score for each user session based on the user's actions and errors.

4.1.3 Privacy and GDPR

As a European company, Hotjar is subject to the General Data Protection Regulation (GDPR) and can be used by European companies without issue. The following steps must be followed to collect user data with Hotjar:

- Include information about the usage of Hotjar on the website in the Terms of Services or Privacy Policy.
- Sign a Data Processing Agreement (DPA).

4.1.4 Advantages and Disadvantages

Hotjar is an exceptional tool when looking at the available features. It is intuitive to use and does not require specific skills. It does, however, fall short in terms of providing specific features for performance analytics.

Hotjar has the following advantages:

+ *Easy Setup*: The onboarding and installation process is well structured, intuitive, and supported by step-by-step guides.
+ *Easy to Use*: Features are mostly self-explanatory and do not require coding skills or a deeper understanding of the technology.
+ *Transparent Pricing*: Clearly defined prices for different subscriptions and what they entail can be found on their website [Hotjar 2024c].
+ *User Surveys and Interviews*: Hotjar is the only tool featuring templates for User Surveys and Interviews.
+ *Customer Support*: The tool offers many guides and blogs on various topics, a help centre to find support, and an integrated chat with customer service on the home page.

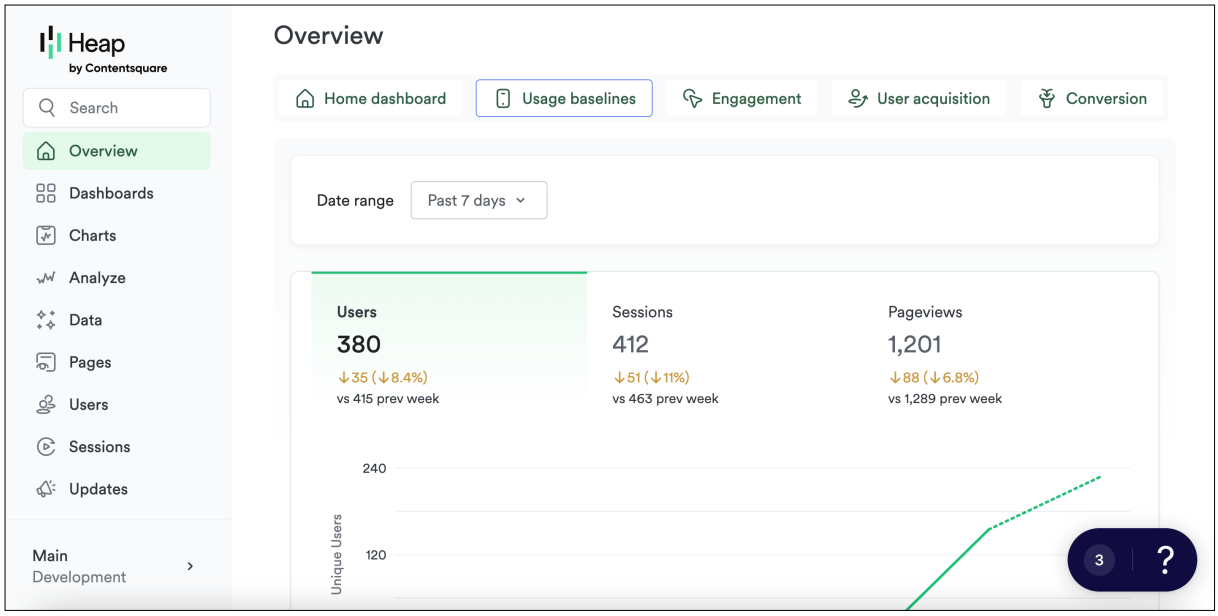

Figure 4.3: Heap: Dashboard. [Screenshot taken by Kilian Weisl]

Hotjar has the following disadvantages:

- *Commercial Tool*: Although Hotjar offers a Basic version free of charge, its features are minimal, and only a maximum of 35 user sessions per day are tracked.
- *Scaled Pricing*: Hotjar subscription plans become increasingly expensive, depending on how many users visit the site daily.
- *Resource Consuming*: The installation of Hotjar led to a noticeable performance drop. The bundle size of the tool is around 400 KB.

4.2 Heap

Heap was founded in 2013 and acquired by Contentsquare in 2023 [Heap 2024d]. Heap is a product analytics platform focusing on easy-to-use insights to illuminate customers' key digital behaviours and pinpoint quantitative and qualitative insights. Over 10,000 companies worldwide use Heap's cross-channel and cross-device analytics to improve adoption and visualize customer engagement. Heap is a software-as-a-service (SaaS) service that charges customers based on a subscription-based model. For this survey, the free 14-day trial was used, which offers a pro-version (not premium-version), and afterward changed to the free plan, which allows up to 10,000 monthly sessions [Contentsquare 2023]. The Heap Dashboard is shown in Figure 4.3.

4.2.1 Setup

Heap has a rich documentation service called Heap University [Heap 2024c], where it is possible to learn how to use the software in general and also best practices in Heap. Additionally, it offers live training and events. The developer portal focuses on technical topics, such as setting up Heap on a website and documentation about the Application Programming Interface (API) and Software Development Kit (SDK) [Heap 2024a]. In particular:

- *Account Creation*: Before integrating Heap, one must create an account on their website and choose a subscription model.
- *Integration*: To integrate Heap in a website, a web developer needs to add a specific script element

to the <head> section of each web page to be tracked, including the unique App-ID to identify the Heap account.

- *Additional Changes*: If a Content-Security-Policy is active on the website under investigation, the web developer must also add some entries to it to load certain files that Heap needs [Heap 2024a].

4.2.2 Feature Set

Based on the canonical feature set that was defined in Chapter 2, Heap does not support as many features as Hotjar or Clarity.

4.2.2.1 Real-Time User Tracking (RTUT)

Heap groups user interactions into different types of events, for example:

- Click on `i.fas.fa-shopping-cart`
- View `/shop/products/xyz/`
- Submit on `form.checkout.woocommerce-checkout`

In Heap, it is possible to follow events in real time, but only in textual form in an event log. Unlike Clarity, Heap is not able to show the website as the user perceives it.

4.2.2.2 Session Replay

Heap provides session replays only for premium users, therefore this feature was not tested as part of this survey.

4.2.2.3 Error Tracking

Heap provides error tracking only for premium users, therefore this feature was not tested as part of this survey.

4.2.2.4 Real Performance Measurements

Heap does not provide real performance measurements for its customers.

4.2.2.5 User Paths

To find out, which users visited specific pages, Heap provides a section called View Events on the Sessions page, shown in Figure 4.4. Another way to see standard pages users visit is the so-called Journey feature, which can be found on the Analyze page. With numerous filters, this provides insights into what percentages of users followed a specific path on the website. A very common behaviour in Heap is the creation of charts, which can be saved and used throughout the platform to directly see essential data important for a specific use case.

4.2.2.6 User Segmentation

Heap provides a page called Users, where it is possible to filter users based on an extensive set of filters. These filters can be activated and combined as needed and always consist of a property, an operator, and a value. Properties can be segments, behavioural, active usage, defined user, and standard user properties. Depending on the property and operators, the value can vary. When it comes to operators, Heap provides standard operators (e.g. “contains”, “does not contain”, “is defined”, etc.) but also more sophisticated ones (e.g. “wildcard matches”, “wildcard does not match”, etc.). This gives Heap users extensive power over searching and retrieving data for a specific use case.

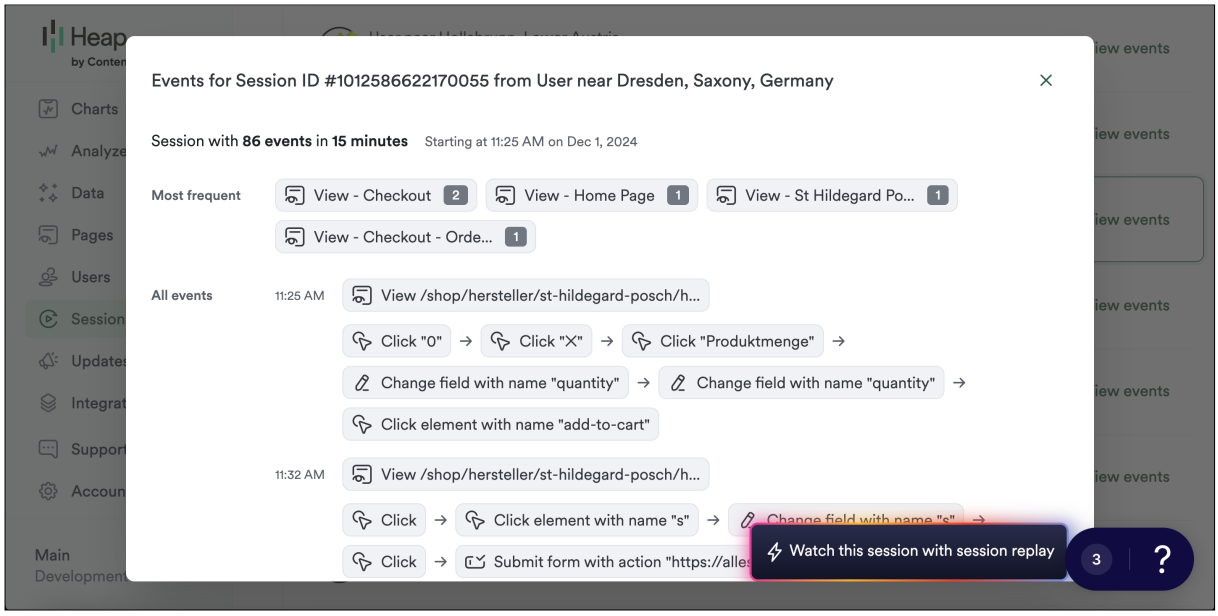

Figure 4.4: Heap: Sessions Page. [Screenshot taken by Kilian Weisl]

4.2.2.7 User Surveys

Heap does not provide user surveys.

4.2.2.8 User Satisfaction Metrics

As part of the free pro version trial, Heap does provide satisfaction metrics (e.g. rage clicks). Due to the extensive filtering methods, it is also possible to find out which paths a user took. The Analyze page provides a tab called Retention, which can be used to additionally check for u-turns, exit rates, or bounce rates.

4.2.3 Privacy and GDPR

Heap passes responsibility regarding GDPR and data privacy to its customers, as stated on their privacy page: "To ensure we don't collect any special classes of personal information, it is the customer's responsibility to not send any sensitive personal data into Heap." [Heap 2024b]

However, they provide several features to minimise the risk of sending unwanted or sensitive information to Heap, including:

- *Target Text Auto-Capture*: This is a toggle implemented in Heap, which ensures that sensitive information included in elements of web page is not captured.
- *DisableTextCapture API*: An alternative to Target Text Auto-Capture.
- *Heap Redact*: An HTML attribute that can be added to elements to prevent the data inside the element being sent to Heap.

The complete list of privacy features Heap provides can be found in [Heap 2024b].

4.2.4 Advantages and Disadvantages

Heap is easy-to-install, given some technical knowledge about web development. The interface is intuitive and visually appealing. Heap provides rich documentation, usage guides, courses, and is especially beginner-friendly. However, its pricing is not as transparent as other investigated tools, and the JavaScript bundle size is also not lightweight compared to Clarity.

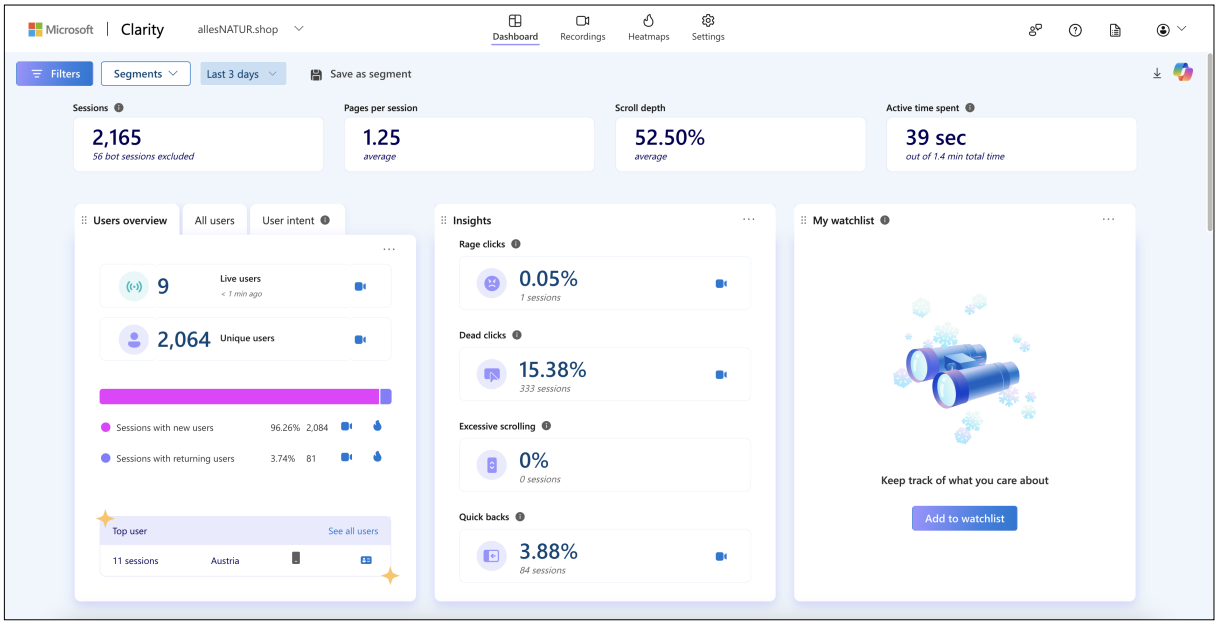

Figure 4.5: Clarity: Dashboard. [Screenshot taken by Niklas Lorber.]

Heap has the following advantages:

+ *Visually Appealing*: The graphical user interface is clean and easy to navigate.
+ *Intuitive to Use*: Feature names and navigation speak for themselves, so it is easy-to-use.
+ *Documentation and Learning Support*: Heap University [Heap 2024c] provides extensive material, courses, and live events to improve knowledge about the platform.
+ *Customer Support*: Heap provides extensive support for its customers, which is especially helpful when problems or questions arise.

Heap has the following disadvantages:

- *Commercial Tool*: Heap is not free and charges users based on website traffic.
- *Intransparent Pricing*: Heap does not provide information about exact prices on their pricing page. One must be signed-up already in order to get an estimate of the price depending on the chosen plan, depending on the website traffic.
- *Onboarding Complexity*: Compared to other tools, Heap does not provide an easy-to-install plug-in (e.g., like Hotjar does in WordPress).
- *Resource Consuming*: Heap has a JavaScript bundle size of around 307 KB.

4.3 Clarity

Microsoft launched Clarity, a free-to-use, closed-source RUM tool, on 28 Oct 2020 [Microsoft 2024]. Their goal was to build a broadly available tool containing the most essential features of an analytic RUM tool. The tool started with three integral features: Insights Dashboard, Heatmaps, and Session Replay [Dicks 2020]. At the time of writing, no further significant features have been introduced since then. The Clarity Dashboard is shown in Figure 4.5.

4.3.1 Setup

Clarity's setup process is not as straightforward and documented as the commercial tools, but it is still feasible for users with varying technical levels of expertise. In particular:

- *Integration*: Integrating Clarity involves adding a JavaScript snippet to the header section of every page to be tracked. This is not explained as well as Hotjar and Heap. However, no further coding skills are required, if one has access to the full source code of the website to be investigated.
- *Dashboard Initialisation*: After integration, users can access a dedicated dashboard that automatically begins tracking user interactions. No further setup is required for tracking. However, the application needs a few minutes to initialise the data so that it works accordingly.
- *Ease of Use*: The tool's functionality and interface are not as well explained as in the commercial tools. However, it is still fairly easy to familiarise oneself with the tool's capabilities.

4.3.2 Feature Set

Clarity provides a well-rounded feature set focusing on core RUM functionalities. While it may not match the extensive offerings of commercial tools, it excels in delivering essential capabilities to track and analyse user interactions.

4.3.2.1 Real-Time User Tracking (RTUT)

Clarity supports an entirely out-of-the-box, functioning, real-time user tracking feature. It aggregates the user's events and reports them to the application in real time. The session can then be followed in a graphical recording-like format, like a live stream. However, during this survey, there were some issues with the reliability of this feature. The real-time nature of the session was also occasionally in doubt.

4.3.2.2 Session Replay

Clarity excels at session replay functionality, allowing investigators to review user interactions post-session in a graphical way. The most important capabilities of Clarity's session replay are:

- Clicks, scrolls, and mouse movements are fully visualised.
- Automatic anonymisation of sensitive data (e.g., input fields).
- Device and viewport detection, enabling accurate reproduction of user experiences across mobile, tablet, and desktop devices.
- Milestone integration, allowing to jump to predefined events more efficiently.

Clarity's session replay is extremely convincing. Thanks to its simplicity, even non-technical users can easily gain insights from the sessions. Furthermore, the operation is intuitive, even new users can quickly find their way around.

4.3.2.3 Error Tracking

Clarity also comes with integrated error tracking. The way Clarity handles error tracking is tightly coupled to the session replay feature. As described in Section 2.3, this feature is used to identify potential JavaScript issues. Clarity constantly monitors the website for upcoming JavaScript errors and classifies them accordingly. Due to the sophisticated filtering system, the investigator can later find and react to the desired problematic sessions with JavaScript errors.

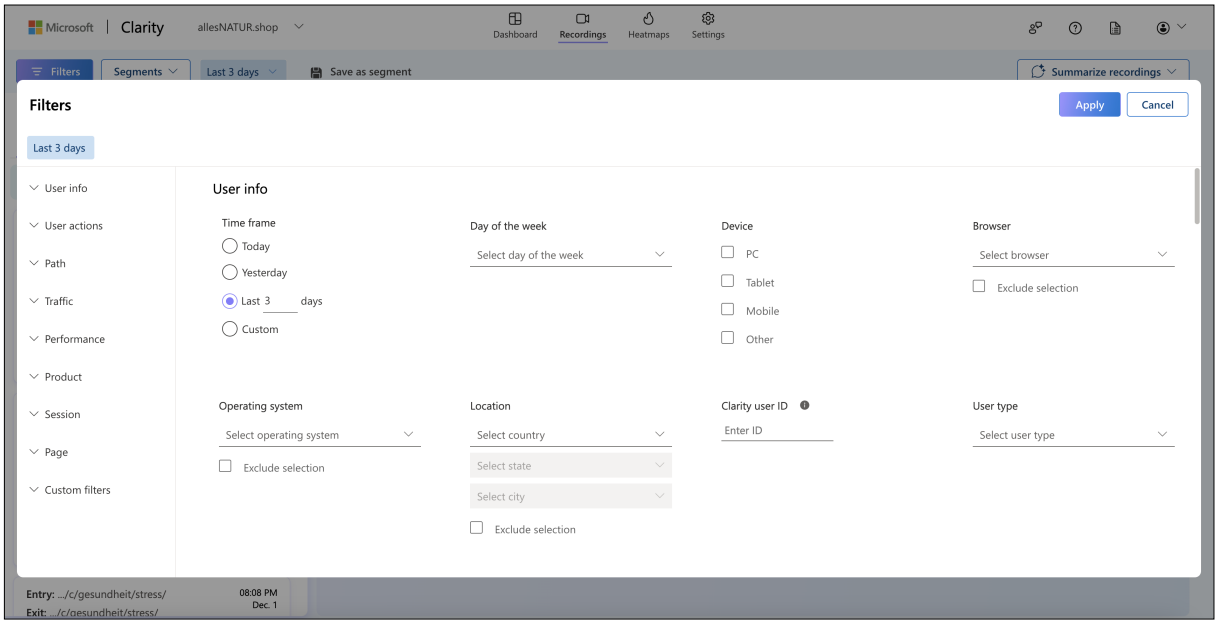

Figure 4.6: Clarity: Filtering and segmentation options are available to filter for device type, time of visit, operating system, and many more. [Screenshot taken by Niklas Lorber.]

4.3.2.4 Real Performance Measurements

Clarity provides basic performance measurements, focusing primarily on aggregated data like page load times and user interaction metrics. However, it does support the Core Web Vitals without the need to integrate an additional third-party plugin.

4.3.2.5 User Paths

The tool allows investigators to analyse user navigation paths through its Insights Dashboard. Investigators can identify popular entry and exit points, helping optimise the website's structure and usability. However, user paths are not as granular or customisable as those in commercial solutions.

4.3.2.6 User Segmentation

Clarity excels at basic user segmentation. It allows filtering sessions by device type, country, browser, and more, allowing the building of creative and sophisticated persistent segments, as shown in Figure 4.6. However, it lacks more advanced segmentation techniques (e.g., specific clicks or workflows).

4.3.2.7 User Surveys

Clarity does not include user surveys or questionnaires, limiting its ability to gather direct user feedback or satisfaction with specific workflows.

4.3.2.8 User Satisfaction Metrics

Clarity comes with some user satisfaction metrics:

- *Rage Clicks*: Users repeatedly click or tap on the same small area, indicating frustration.
- *Excessive Scrolling*: Users scroll more than usual, often indicating confusion.
- *Dead Clicks*: Users expected clickable events without further action, possibly leading to dissatisfaction.

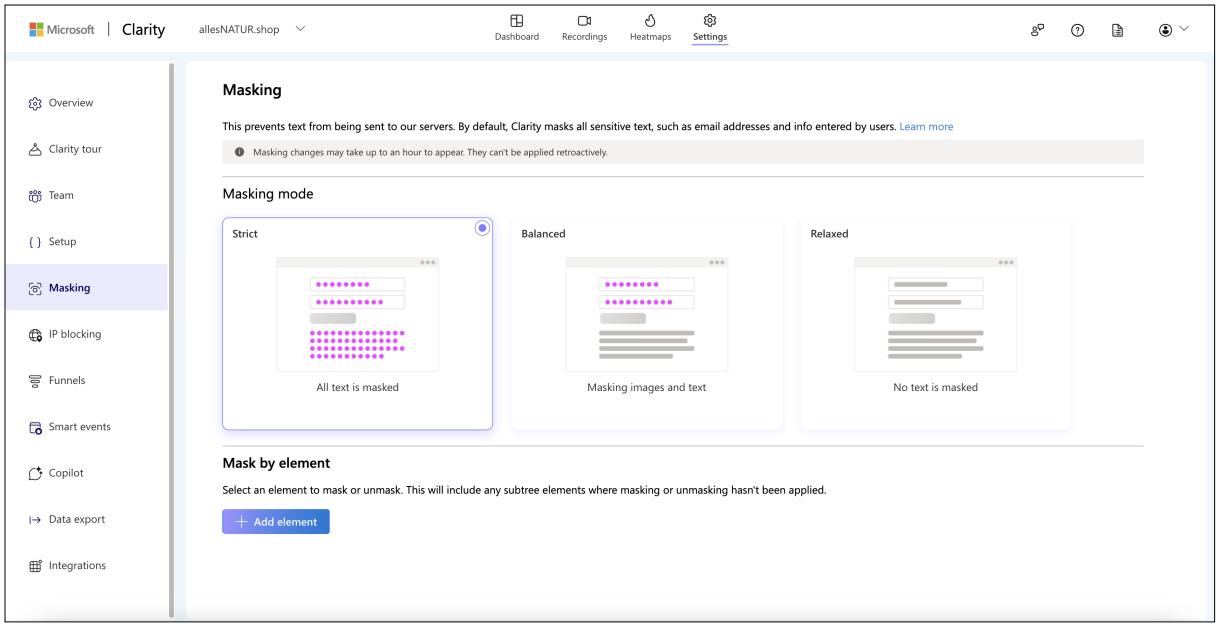

Figure 4.7: Clarity: Different modes of masking user input on the monitored website. [Screenshot taken by Niklas Lorber.]

4.3.3 Privacy and GDPR

Microsoft Clarity is DPF-certified, so its use in the European Union is trouble-free. However, some aspects should be considered to use Microsoft Clarity in compliance with the GDPR:

- Mask input fields, as illustrated in Figure 4.7.
- Obtain cookie consent.
- Enter into a Data Processing Agreement (DPA).
- Customise privacy policy.

Once the points listed above have been completed, nothing should legally stand in the way of operating Microsoft Clarity [eRecht24 2024].

4.3.4 Advantages and Disadvantages

Clarity is a particularly convincing tool, because it is free and lightweight, and yet still offers a wide range of functionality. Clarity's RTUT and sophisticated filtering were especially convincing.

Clarity has the following advantages:

+ *Non-Commercial*: Free-to-use, but not open-source.
+ *Functionality*: Covers almost all canonical RUM features.
+ *Lightweight*: Very lightweight, with a 70 KB total bundle size.
+ *AI Chatbot Support*: Includes an AI-powered chatbot to assist with the dashboard.

Clarity has the following disadvantages:

- *Onboarding Complexity*: Requires technical knowledge to set up.
- *Aesthetics*: Not as visually appealing as other tools.
- *Lack of Support*: No dedicated support is available.

Chapter 5

RUM Tool Recommendation

Three RUM tools which were investigated in detail for this survey: Hotjar [Hotjar 2024d], Heap [Heap 2024d], and Clarity [Microsoft 2024]. They were integrated and tested in the context of a live e-commerce website. Table 5.1 summarises the features supported by each of the three tools. Since each of the RUM tools impacts a website's performance, it is not advisable to install multiple tools in parallel.

The following aspects should be considered when choosing a RUM tool:

- *Resource Consumption*: The size of the included JavaScript bundle affects loading time. Clarity excels in this aspect.
- *Ease of Use*: How much expert-knowledge is necessary to use a tool? Heap excels in this aspect.
- *Features*: How many features are implemented by a tool. Hotjar excels in this aspect.
- *Budget Constraints*: How much it costs to use a tool. Clarity excels in this aspect.

The choice of tool also depends on the needs of the specific project. For instance, if user surveys are a fundamental requirement, only Hotjar can be used. On the other hand, if performance is critical, Clarity 4.3 is the best option. However, based on the observations from this comparative survey, Clarity is recommended as a good first choice for most projects:

- *Cost Effective*: Completely free to use, ideal for low budget projects.
- *Performant*: Very lightweight, with only 70 KB bundle size.
- *Feature Rich*: Covers all investigated features, except for user surveys.
- *Innovative Tools*: Includes an AI chatbot for enhanced user interaction and data analysis.

Feature	Hotjar	Heap	Clarity
Real Time Tracking:	✗	Events Only	✓
Session Replay:	✓	Premium Version	✓
Error Tracking:	✓	Premium Version	✓
Real Performance Measurements:	GA Integration	✗	✓
User Paths:	✓	✓	✓
User Segmentation:	✓	✓	✓
User Surveys:	✓	✗	✗
User Satisfaction Metrics:	✓	✓	✓

Table 5.1: Comparison of features between Hotjar, Heap, and Clarity. A green tick indicates a feature is supported. A red cross indicates it is not supported. Otherwise, a textual description is provided. GA stands for Google Analytics.

Chapter 6

Concluding Remarks

This survey examined the diverse landscape of Real User Monitoring (RUM) tools. After an initial discovery phase looking at six commercial and three non-commercial tools, three tools were selected for more datailed evaluation: Hotjar, Heap, and Clarity. These were then investigated by integrating them into a live ecommerce website. The survey identified each tool's strengths and weaknesses, helping to provide a holistic guide for selecting an appropriate RUM tool based on specific project requirements.

While commercial tools like Hotjar and Heap provide extensive functionality, especially in their paid programs, including questionnaires and user surveys, the non-commercial tool Clarity offers robust core features at no cost, making it ideal for projects with limited monetary resources. In terms of performance, Clarity also excels here with its lightweight implementation, whereas Heap focuses more on user-friendliness as well as extensive documentation, making it suitable for non-technical users. Overall, Clarity is recommended as a good first choice RUM tool.

Bibliography

Basic RUM [2024]. *Basic RUM: Open-Source Real User Monitoring*. 2024. https://basicrum.com/ (cited on page 12).

Calibre [2024]. *Calibre: Monitor Core Web Vitals and Page Speed with Synthetic and Real-User Data*. 2024. https://calibreapp.com/ (cited on page 10).

Contentsquare [2023]. *Acquisition of Heap*. 2023. https://contentsquare.com/blog/contentsquare-completes-acquisition-heap/ (cited on page 19).

Contentsquare [2024]. *Use Cases for Filtering Recordings*. 2024. https://contentsquare.com/guides/customer-satisfaction/metrics/ (cited on page 7).

Croll, Alistair and Sean Power [2009]. *Complete Web Monitoring*. O'Reilly, 12 Jun 2009. ISBN 0596551355 (cited on page 1).

Dicks, Aaron [2020]. *Microsoft launches Clarity web analytics: free heatmaps and session recording*. 2020. https://impressiondigital.com/blog/microsoft-clarity-free-heatmaps-recordings/ (cited on page 22).

Dynatrace [2024]. *Dynatrace: Real User Monitoring*. 2024. https://dynatrace.com/platform/real-user-monitoring/ (cited on page 10).

eRecht24 [2024]. *Datenschutzerklärung für Microsoft Clarity*. 2024. https://e-recht24.de/dsg/12700-microsoft-clarity.html (cited on page 25).

Google [2024a]. *Core Web Vitals*. 2024. https://web.dev/articles/vitals (cited on pages 6, 17).

Google [2024b]. *PageSpeed Insights*. 10 May 2024. https://pagespeed.web.dev/ (cited on page 6).

Heap [2024a]. *Heap Dev Hub*. 2024. https://developers.heap.io/docs/overview (cited on pages 19–20).

Heap [2024b]. *Heap Privacy Features*. 2024. https://help.heap.io/data-privacy/privacy-features/using-heap-to-comply-with-data-privacy-legislation/ (cited on page 21).

Heap [2024c]. *Heap University*. 2024. https://learn.heap.io/ (cited on pages 19, 22).

Heap [2024d]. *Heap: Product Analytics for Digital Experiences*. 2024. https://heap.io/ (cited on pages 9, 19, 27).

Hotjar [2024a]. *7 Hotjar Features You May Not Know About*. 2024. https://hotjar.com/blog/features-you-may-not-know-about (cited on page 7).

Hotjar [2024b]. *About Hotjar*. 2024. https://hotjar.com/about-us/ (cited on page 15).

Hotjar [2024c]. *Hotjar Pricing*. 2024. https://hotjar.com/pricing/ (cited on page 18).

Hotjar [2024d]. *Hotjar: Behavior Analytics and User Feedback Tools*. 2024. https://hotjar.com/ (cited on pages 9, 27).

Hotjar [2024e]. *How to Use Rage Clicks to Improve User Experience*. 2024. https://hotjar.com/blog/rage-clicks/ (cited on page 7).

Mastin, Pete [2016]. *Real User Measurements*. O'Reilly, Oct 2016. ISBN 1491944056 (cited on pages 1–2).

Microsoft [2024]. *Clarity: Free Heatmaps and Session Recordings*. 2024. https://clarity.microsoft.com/ (cited on pages 11, 22, 27).

Mihajlija, Milica and Philip Walton [2023]. *Cumulative Layout Shift (CLS)*. 2023. https://web.dev/articles/cls (cited on page 6).

OpenObserve [2024]. *OpenObserve: Open Source Observability Platform*. 2024. https://github.com/openobserve/openobserve (cited on page 11).

Sematext [2024]. *Sematext: IT System Monitoring Tools for DevOps*. 2024. https://sematext.com/ (cited on page 10).

UXCam [2024]. *UXCam: Mobile App Analytics To Understand User Behavior*. 2024. https://uxcam.com/ (cited on page 10).

Wagner, Jeremy [2024]. *Interaction to Next Paint (INP)*. 2024. https://web.dev/articles/inp (cited on page 6).

Walton, Philip and Barry Pollard [2024]. *Largest Contentful Paint (LCP)*. 2024. https://web.dev/articles/lcp (cited on page 6).

Wikipedia [2024a]. *Bounce Rate*. 2024. https://wikipedia.org/wiki/Bounce_rate (cited on page 7).

Wikipedia [2024b]. *Net Promoter Score*. 2024. https://wikipedia.org/wiki/Net_promoter_score (cited on page 18).