Reality Mining in Sensor-based Mobile-driven Environments

Dennis G. Geisse[1], Iven John[2], Sebastian Kotstein[3]

Abstract: Reality Mining refers to an application of data mining, using sensor data to derive behavioral patterns in the real world. However, research in this field started a decade ago when technology was far behind today's state of the art. This paper discusses which requirements are now posed to applications in the context of Reality Mining. A survey has shown which sensors are available in state-of-the-art smartphones and usable to gather data for Reality Mining. As another contribution of this paper, a Reality Mining Application Architecture is proposed to facilitate the implementation of such applications. A proof of concept verifies the assumptions made on Reality Mining and the presented architecture.

Keywords: Reality Mining, application architecture, sensor-based environments, smartphone features, wearable sensors, Android applications, big data applications

1 Introduction

When Nathan Eagle and Alex Pentland began conducting their research in a field they named Reality Mining in 2004, their aim was to derive behavioral patterns from data collected by sensors in the real world [EP06]. As technology has evolved since then and big data technologies are increasingly adopted, questions arise how this evolution affects the initial understanding of Reality Mining. While Eagle and Pentland were restricted to few available sensors built into phones at that time, currently available smartphones feature much more sensor technology. There are even more measurements available when expanding the range with environmental sensors. How that improves the possibilities to capture reality, i.e. the movement of an individual in the real world, has not yet been evaluated.

How to use the potential of such sensor technology is also not yet explored in research. Insights of possible challenges for the acceptance of such applications as well as architectural approaches would be valuable for building Reality Mining applications. While there is existing research regarding the architecture of general big data applications [PP15, Kr14, Bi16], none of them considers the generation of data through remote sensors and its transfer to a central storage and analytics system. They also do not consider how to make data available for remote applications.

[1] Reutlingen University, Herman Hollerith Zentrum, dennis_gregor.geisse@student.reutlingen-university.de
[2] Reutlingen University, Herman Hollerith Zentrum, iven.john@student.reutlingen-university.de
[3] Reutlingen University, Herman Hollerith Zentrum, sebastian.kotstein@student.reutlingen-university.de

The contribution of this paper is an illustration of how smartphones have evolved in previous years and what impact this has on the idea of Reality Mining. In addition, an architecture for building Reality Mining applications is introduced and applied in a proof of concept.

2 Methodology and Focus of This Paper

To further focus the research of this paper, there are two research questions derived from the changes introduced by the smartphone evolution, examining what this means to Reality Mining and how these changes may be leveraged to support Reality Mining.

1. How does the variety of available sensors refine the idea of Reality Mining?

2. What are the core components required for a Reality Mining application capturing and processing data to achieve user benefit and how do they relate?

Before these research questions are answered, Section 3 aims to clarify the initial position on changes in smartphones since the first approach to Reality Mining was researched by Eagle and Pentland in 2004 [EP06]. The goal is to determine whether or not the technology advances – limited to features, sensors and technology that provide added value for Reality Mining – where substantial and how they evolved over time. This study contributes to providing the foundation of available data for a more refined Reality Mining approach and shows how the variety of sensors available in mobile devices in 2015 has increased manyfold. This provides opportunities for new use cases when looking into the subject Reality Mining.

Given research from the smartphone survey [GJK16] it is assumed that the sensor landscape has evolved since 2004. This leads to challenging Reality Mining in definition and execution. When applying well-spread as well as new technologies to the idea of Reality Mining, research question one will determine if and how Reality Mining evolves alongside these technologies. This includes classification criteria for Reality Mining applications.

Having determined how Reality Mining evolved alongside technologies, research question two applies this knowledge to identify common components in Reality Mining applications based on the aforementioned research. The idea is to have a common set of tools to base Reality Mining applications on. Any derived solution needs to be tested against a practical use case for verification purposes and to determine advantages, disadvantages and restrictions.

The research paper will be oriented along the research questions introduced above. As such it will explore each question in-depth before concluding the findings and introducing pointers on further research opportunities.

3 Evolution of Smartphone Features

The acquisition of data in Reality Mining applications is critical, which should not be neglected while planning such an application. Depending on the use case of the application, different kinds of data must be acquired. This requires a special capturing device, e.g. a smartphone or a smart watch having appropriate sensors. This device must be affordable for and available to a large target group of potential users, otherwise only a fraction of the target group would be able to access and use the application. Also, considering group prediction capabilities, the application requires data from many users – otherwise the quality would suffer. So all in all, two issues have to be considered: On the one hand, state-of-the-art smartphones and mobile devices must provide the required features (e.g. sensors) for capturing the data and on the other hand those features must be supported by many "popular" devices (i.e. devices used by a large target group) [GJK16]. Hence, a smartphone feature survey was conducted to identify and classify the state-of-the-art smartphone features which might be potentially relevant for Reality Mining applications. Furthermore, the survey illustrates the evolution of these identified features in the smartphones market – especially the support for each feature – between 2004 and 2015 [GJK16]. Presenting the whole approach and procedure of this survey would go beyond the scope of this paper. Nevertheless, this paper will provide the most important relevant facts and results of this survey.

The survey has identified 22 smartphone features relevant for Reality Mining applications which are capable of providing details about the user behavior and environmental conditions. Additionally, it has identified three more partially relevant features which are not content of further analysis [GJK16]. The identified features and their classifications are shown in Table 1.

Feature	Category	Feature	Category
Touchscreen	User Input	Finger Sensor	Human Sensor
Wifi	Communication	Gesture Sensor	
Bluetooth		Heart Rate Sensor	
NFC		Proximity Sensor	
Infrared		SpO2	
(GSM)		Microphone	Environmental Sensor
GPS	Localization & Movement	Camera	
GLONASS		Luxmeter	
(Beidou)		UV Sensor	
(Step Counter)		Magnetometer	
Accelerometer		Thermometer	
Gyroscope		Barometer	
		Humidity Sensor	

Tab. 1: Identified features and their classifications. Excluded features in brackets [GJK16].

The survey selected 143 smartphones released in Germany between January 2004 and May 2015. These smartphones have been analyzed regarding their support for the identified relevant features by determining whether a smartphone supports a feature. This results in a feature counter for each smartphone that was used to rank all analyzed devices regarding their overall capability for Reality Mining.

By analyzing the support of each feature per year, the survey reveals an interesting fact which applies on nearly every feature observed in this survey: The support of an introduced feature is commonly growing over the next years after its first occurrence which means that the vendors would rather add new additional features to a smartphone than replacing old features. For instance, state-of-the-art features of 2006 (e.g. Bluetooth, Wi-Fi or GPS) are still supported by smartphones released in 2015. But it has to be considered that technology revisions might have changed (e.g. upgrading from IEEE802.11g to IEEE802.11n), which is not considered the survey.

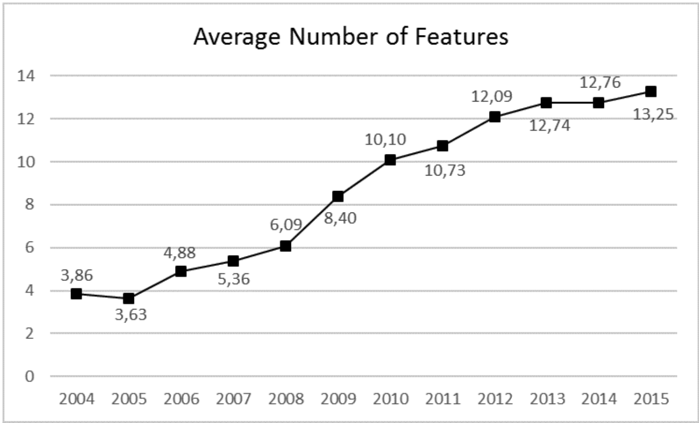

Fig. 1: Average number of features supported by a smartphone released per year [GJK16]

Figure 1 shows the average number of features of smartphones introduced in the corresponding years. As depicted, the average number of features has been increased from 3.86 features in 2004 to 13.25 in the year 2015.

4 Refining Reality Mining

In 2004, Eagle and Pentland have demonstrated the potential of smartphones being aware of the behavior of their users [EP06]. They have equipped 100 test users with Nokia 6600 smartphones as wearable sensors for recognizing social patterns in the daily behavior of an individual user as well as of a whole group [EP06]. Compared to modern phones, those devices have a limited feature set, consisting of Bluetooth, Infrared, Global System for Mobile Communications (GSM), a Camera and a Microphone [GS15], but only Bluetooth and GSM have been used as sensors for their study. In the meantime, the average number of Reality Mining-relevant features has increased from 3.86 features in 2004 to 13.25 features in 2015 – for comparison: The Nokia 6600 has 4 features, considering that GSM has been excluded from the Smartphone Feature Survey [GJK16]. This leads to the question, whether the initial idea of Reality Mining by Eagle and Pentland [EP06] can be refined and extended with new technological or even social aspects considering the fact that the smartphone market has experienced a remarkable technological progress in the last decade with products like the first iPhone [Ke12] or the introduction of the Android operating system [An16].

To answer this question, the aspects of the study of Eagle and Pentland [EP06] have to be examined and discussed from today's perspective. Their study has two elementary objectives beside the demonstration of the ability to use smartphones as wearable sensors, namely the creation of a predictive classifier describing daily, weekly or even yearly patterns of an individual user as well as the identification of community structures, like friendships, work groups and organizational groups. These predictive and descriptive models have primarily relied on the location of the users as well as their proximity to other users and static known devices. Moreover, collected information about the phone status, call logs and use of applications improves the models. For detecting the location of a user both Bluetooth as a short-range radio frequency (RF) network for detecting static devices (e.g. Desktop Computers) within a range of 10 meters and GSM as a long-range RF technology being capable of tracking the Cell Tower IDs were used. Using these techniques together completes the data in case a user has no service (e.g. in the center of a building), but might be in range of a static device – or vice versa. The proximity of users is also based on Bluetooth, by discovering other Bluetooth devices (representing users) that are in range. To summarize, the study has returned descriptive data reflecting individual user behavior as well as the relationships between users.

From a technological point of view in 2015, they have only used the Bluetooth and GSM features of the *Communication* category [GJK16]. More accurate tracking technologies (category *Location & Movement*), e.g. Global Positioning System (GPS), have been considered but were not supported as built-in features in smartphones from 2004 [GJK16]. If Eagle and Pentland had had access to GPS capable smartphones in 2004, the location determination would have been more accurate and less complex. This leads to the conclusion that today's variety of additional features, when compared to 2004, enables new use cases which can rely on more accurate data. Nevertheless, if a feature is planned to be used in a Reality Mining application, its support on the smartphone market should be analyzed to avoid limitations on the target group [GJK16].

Additionally, social aspects might have changed – especially due to the fact that the sharing of personal and private data is a sensitive matter today. This has to be considered for any Reality Mining use case. The privacy aspect has already been considered by Eagle and Pentland, but their study is a proof of concept under lab conditions with the approval of all participants [EP06]. Today's applications have to respect the *Privacy* of their users, which is a legal and moral issue. But even if privacy is respected, the users might not be willing to share their data. However, users are more likely to provide their private data if they have a benefit by sharing them [QR15]. This research has adapted this statement into a generic approach which applies to every Reality Mining application.

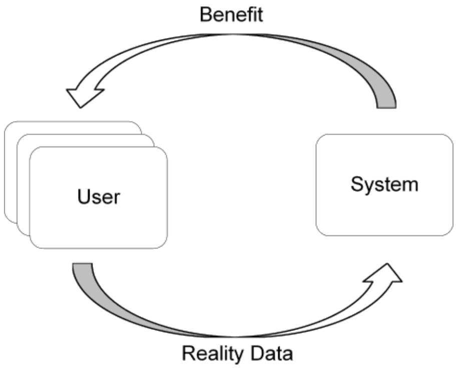

Fig. 2: Reality Mining Cycle

The Reality Mining Cycle, illustrated in Figure 2, exhibits the fact that Reality Mining applications have to provide a benefit to each individual user, otherwise the users would not use the application and would not share the data describing their behavior with the system. Furthermore, the *Goal Support* of an application can provide further incentive for using the application. In this context, support describes the degree of simplification and automation of a daily process of the user using the application whereas the added *User Benefit* is more of the objective of the application. Moreover, a Reality Mining application can be rated by its *Social Value*, which is the benefit for the community when using the application. These four dimension – *Privacy*, *User Benefit*, *Goal Support* and *Social Value* – introduced in this section were identified as the base criteria for today's Reality Mining applications.

This section has shown that the changes in the mobile technology landscape driven by a variety of new features broaden possibilities for use cases, but also bring challenges. New possibilities and emerging social aspects refine the original idea and focus of Reality Mining, answering the first research question.

5 Reality Mining Application Architecture

Building applications in the context of Reality Mining poses challenges regarding the applications architecture [CZ14]. As the Reality Mining cycle introduced in Figure 2 suggests, such applications always have some kind of data exchange between the users and the Reality Mining system at its core. In the following discussion on how a Reality Mining application can be architected, the Reality Mining system will be called the *backend*.

Reality Mining applications using the stated possibilities of today's smartphones can be considered big data applications, since smartphones are widely adopted and their variety of sensors may produce large amounts of data in several formats [CZ14]. Hence, an architecture for big data applications will serve as the starting point from where a Reality Mining architecture is built. With their research on common implementations of big data applications, Pääkkönen and Pakkala have proposed a reference architecture to realize big data systems [PP15], which, however, does not take into account where data originates and how it is transferred to the backend. Figure 3 shows a schematic architecture where elements of the Pääkkönen and Pakkala reference architecture [PP15] are used to build the backend. In comparison to the original architecture, it is extended with elements necessary to integrate remote sensors such as smartphones. The following description highlights components of the resulting architecture in *italics*.

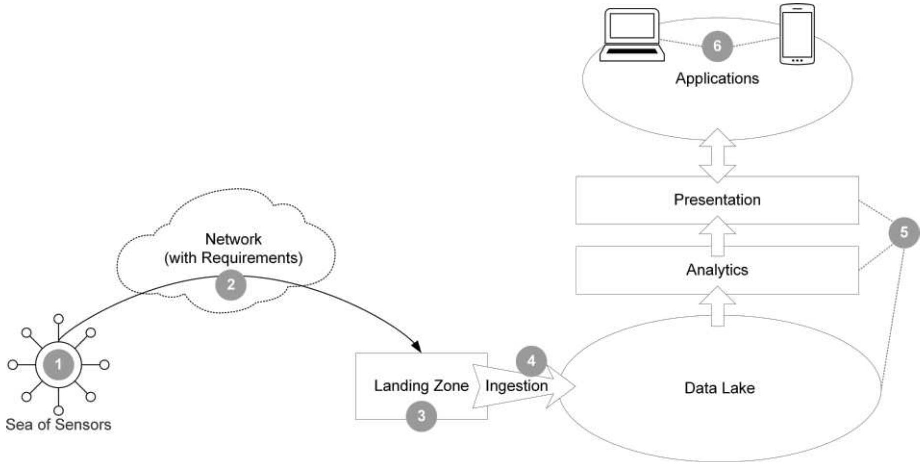

Fig. 3: Extended Big Data Reference Architecture

As the reference architecture proposes, the backend is comprised of elements for data extraction, loading, processing and analysis as well as interfacing [PP15]. In particular, *Ingestion* (4) represents the step of moving data into a central storage, the *Data Lake*, from where it can be loaded for further purposes. In this architecture there is no restriction as to what type of data can be ingested, because nowadays there are possibilities to process all kinds of data whether structured, semi-structured or unstructured [Kr14, PP15]. Based on this data, *Analytics* refers to steps of data processing in terms of both data preparation and data analysis. Further, *Presentation* allows for steps to create content that can be presented to the outside, i.e. to *Applications* (6) where users interact with the backend (5) through arbitrary interfaces.

To allow Reality Mining applications, this architecture needs to be extended by elements (1), (2) and (3). The *Sea of Sensors*, i.e. all sensors that are contributing to the Reality Mining application, are a crucial component in such applications and therefore mandatory. The *Network* is necessary to transfer data from its remote origin (1). Depending on the application there can be several requirements to the *Network*, e.g. regarding latency or bandwidth. Data is received by the *Landing Zone*, which then serves as a data source from where it can be ingested (4) into the backend (5).

This extended big data architecture as depicted in Figure 3 features all necessary elements for Reality Mining applications. However, it is constructed strictly sticking to the big data reference architecture proposed by Pääkkönen and Pakkala. Thus, it is more of a big data architecture with Reality Mining capabilities than a Reality Mining architecture utilizing existing big data research. The aim of this paper, however, is to optimize it for Reality Mining. Weakening the dependency from the big data reference architecture leads to a similar yet more mobile-centric approach, the Reality Mining Application Architecture shown in Figure 4.

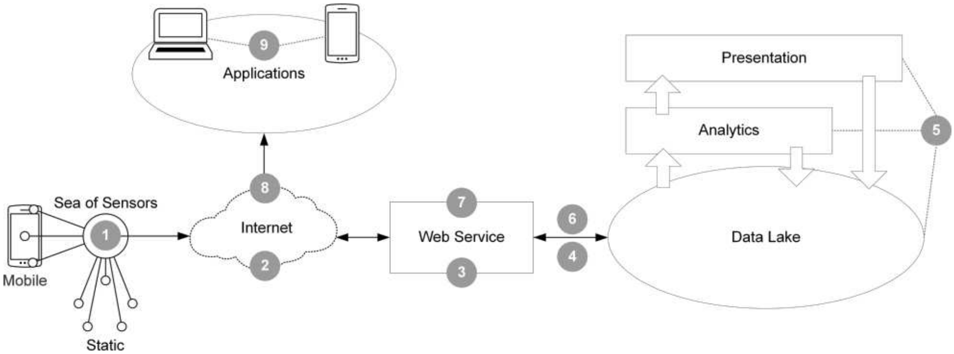

Fig. 4: Reality Mining Application Architecture

In the proposed Reality Mining Application Architecture, a *Web Service* replaces both the *Landing Zone* and the *Ingestion* introduced in the extended big data reference architecture approach, since Reality Mining data most likely arrives at the backend using the *Internet* as communication medium. Still, there is no restriction on the type of data ingested. This *Web Service* receives data and covers extraction into the *Data Lake*. In addition, communication is bi-directional such that the *Web Service* can also retrieve data from the *Data Lake*.

Similar to the backend in the extended big data reference architecture, in the presented architecture *Analytics* refers to data processing in terms of both data preparation and data analysis. *Presentation* covers steps to create content that can be presented to the outside. Additionally, the data flow is extended such that both *Analytics* and *Presentation* are able to add data, i.e. analytics results, to the *Data Lake*. This is necessary, since *Applications* are not bound to the backend but are loosely coupled to it through the *Internet* and the *Web Service*. This way, results can be retrieved from any *Application* without coupling it to the *Presentation* element, enabling mobility as required in Reality Mining applications. To add more precision on possible data origins, the Reality Mining Application Architecture states personal mobile sensors carried by individuals as well as static environmental sensors, which are placed on pre-defined locations.

Utilizing the proposed architecture, a Reality Mining application would collect data through a *Sea of Sensors* (1) and transmit it (2) to a server (3), from where it is ingested (4) into the backend (5). In case processing in the backend leads to output, this data can be retrieved (6) by the server and sent (7) via Internet (8) to an application displaying it (9). These are the components required for a Reality Mining application and their relation and dependency to each other, answering the second research question. How this architecture applies to a real use case is shown in Section 6.

6 Proof of Concept

After developing the Reality Mining Application Architecture presented in Section 5, the next step is to test it by applying it to a practical example application meeting the Reality Mining criteria identified in Section 4.

The use case for the proof of concept (or prototype) in this research study is a carpooling recommendation system. The idea is to provide meaningful advice on who to carpool with while requiring as little manual interaction from the ("carpoolees") as possible. To this end, the proposed solution is to track the trip history of the end user automatically in the background on a personal smartphone and match trips from all users for potential overlaps in the backend. These overlaps will be calculated and presented to the user as recommendations, without the need for user interaction until the recommendation is presented and may be acted upon. Trips in the sense of this use case are user locations in known regular intervals over a given amount of time between two locations – essentially a regular time-location mapping table between two geographical points.

6.1 Reality Mining Classification and Technological Feasibility

Section 4 introduced four common dimension criteria to measure Reality Mining applications by their *User Benefit*, *Privacy*, *Social Value* and *Goal Support*. The proposed carpooling application measures against the identified criteria as follows:

The most important criteria for a Reality Mining application is the eventual added *Benefit* for the end user. Carpooling allows the user to take trips without accommodating for schedule changes – maximizing efficiency while reducing stress and costs for the both parties. The added user benefit usually comes with an impact to *Privacy*, which needs to be measured and kept to a minimum. In the carpooling use case this means only enabling the collection of data if absolutely necessary and making sure that the captured data is handled as anonymously and securely as possible. For this, the prototype determines the performed activity of the user without storing this information and only enables location tracking when the user is in a vehicle, i.e. the user is driving. Additionally, to accommodate for false positives, the user will get the possibility to optionally verify detected trips, with denied trips not getting stored. The application also ensures that user data will not be exposed to other users with the only exception of the time and place to carpool in addition to user-managed contact data for this purpose. The dimension *Social Value* can be evaluated by the promotion of carpooling and the resulting environment-friendly cost and emission savings. As an added social component, carpooling allows to meet people with similar routines and trips and thus promotes social interactions. *Goal Support* is another important measure for Reality Mining applications. For the presented use case it means that the user will not be burdened with the additional time and effort requirements that come with finding a person to carpool with. The application automatically determines overlaps in trips and recommends them to the user – optionally without any user interaction up to the decision to accept a carpooling recommendation.

Using Table 1 of Section 3, the following paragraph will be dedicated to determining the feasibility of implementing the use case with the sensors available in today's personal mobile devices. Touchscreens are readily available in smartphones nowadays. These will allow for intuitive interaction with the end user (accepting/declining recommendations, approving trips). Wi-Fi and radios supporting new generations of GSM fulfill the prerequisite for exchanging location and recommendation data. The location data is determined with Wi-Fi, GPS/GLONASS/BeiDou as well as cell tower triangulation. In conjunction with the aforementioned technologies accelerometer, gyroscope and proximity sensors can provide data to more accurately determine activity performed by the user. Therefore, the variety of sensors available enables the collection of data needed to calculate accurate carpooling recommendations.

6.2 Using the Reality Mining Application Architecture

After determining the Reality Mining aspects and technical feasibility of the proof of concept, the next step is to adapt the Reality Mining Application Architecture introduced in Section 5. Figure 5 shows the frontend implementation architecture while Figure 6 details the backend architecture.

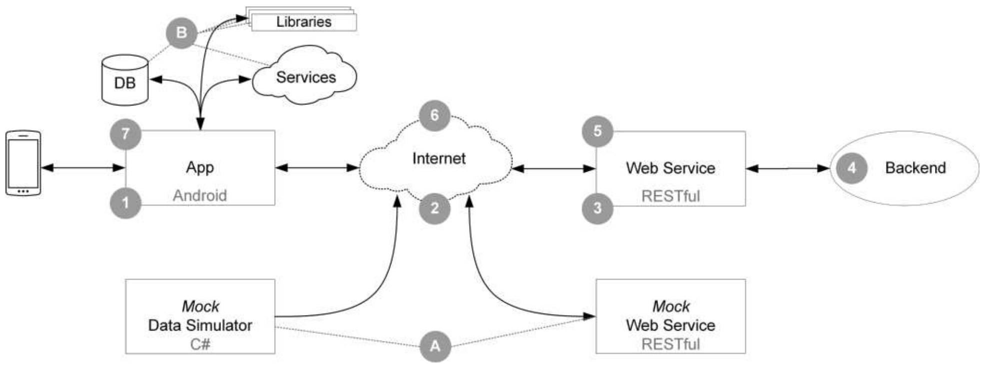

Fig. 5: Proof of Concept Frontend Architecture

Within the frontend, the location tracking part of the prototype is implemented as a smartphone *app* for the Android platform (1) [An16]. For this, the *app* heavily leverages platform features, *cloud services* as well as several *third party libraries* (B). Using the *Internet* as communication medium (2), trip data is send to a *RESTful web service* (3) [RR07] which will make the data accessible for the backend (4). The data mining and calculation process is further discussed later in this chapter. The results of this process (recommendations for carpooling) will then be available for the *web service* (5) to be requested (poll mechanism) on-demand (7) – again using the *Internet* (6) as transport medium.

As the proof of concept heavily relies on user-generated trip data and to enable parallel development of the core two segments, frontend and backend, both a *data simulation application* for mock trip data as well as a *mock web service* are implemented (A) and used to test frontend and backend separately.

Mapping the components of the proof of concept to the core components of the Reality Mining Application Architecture introduced in Section 5 and depicted in Figure 4, the smartphone acts as both mobile sensor and presentation layer for the *application*. The Internet and a RESTful web service are implemented as stated above and can be identified directly as they are named analogously. The backend is implemented as shown in Figure 6.

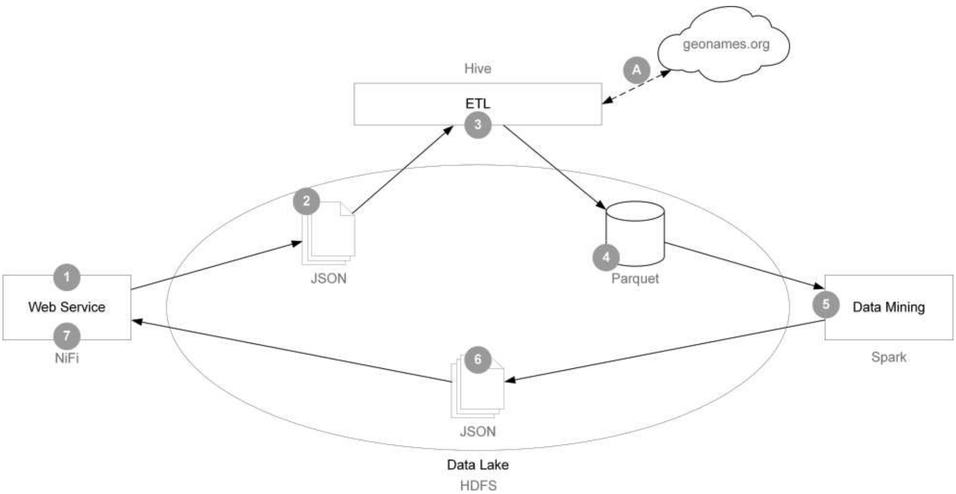

Fig. 6: Proof of Concept Backend Architecture

Data arriving at the *web service* (1) is ingested to the *Data Lake* and stored in semi-structured JSON documents (2). Periodically, an *ETL* workflow (3) loads this data and applies some data preparation, including enrichment through data from third party services (A) [LK15]. As a result of the ETL process, data is now in a structured format (4) and again stored in the Data Lake. Utilizing this structured and enriched data, a *data mining* algorithm (5) determines possible carpooling matches and creates recommendations that are put into the *Data Lake* in JSON format (6). The web service can now retrieve carpool recommendations on user request (7).

As shown in Figure 6, the backend is implemented with technologies from within the Hadoop ecosystem, which provides a broad range of technologies to handle data that is not necessarily structured, like JSON in this case [Kr14, Ap16a]. The web service is leveraging the various functionality of NiFi, a tool to handle data flow in an effective manner [Ap16b]. As data warehousing tool, Hive provides the possibilities to handle the ETL processes efficiently in conjunction with the Hadoop Distributed File System HDFS [Ap16a], which implements the Data Lake. In contrast to other data warehouses, Hive is able to provide relational models directly on original data in HDFS – supporting the concept of "schema later" where data can be stored first without knowing its schema [Kr14]. Additionally, these software packages are working in highly scalable distributed environments what makes them suitable for projects from small proof of concepts to large applications in production [Kr14, Ap16a, Bi16]. This also holds for Spark, a tool for distributed processing that implements the data mining functionality in this proof of concept [Ap16c].

Reflecting the elements within the proposed Reality Mining Application Architecture on this proof of concept, the ETL process serves as analytics implementation. The data mining algorithm represents both the analytics and the presentation layer: Applying a statistical model to derive patterns (*Analytics*) before creating output that is to be passed to the application and that happens to be the most valuable information in this particular use case (*Presentation*).

This implementation of a Reality Mining application shows the possibilities that arise with today's smartphones and their sensor technology. In addition, the proof of concept verifies the applicability of the proposed Reality Mining Application Architecture.

7 Conclusion

This research paper covered Reality Mining in a sensor-based mobile-driven environment. Initially defined research questions were consecutively answered. The concept of being *sensor-based* is defined in Section 3 using the findings of a smartphone feature survey. These findings also allowed for a refined Reality Mining approach, answering the first research question by the introduction of some concepts to define modern Reality Mining applications. *Mobile-driven* was reflected proposing an architecture approach to provide a common ground for developing Reality Mining applications and hence answering the second research question. This approach was then verified using a proof of concept, also detailing advantages and restrictions of the proposed architecture.

The architecture proposed in Section 5 and explored in Section 6 by implementing a proof of concept proved to cover the requirements for Reality Mining applications. Advantages covered by the architecture include modularity, loose coupling, scalability as well as extensibility.

While the proof of concept implemented a solution for a more individual-driven use case, the architecture allows to extend the input systems. In this way, it is possible to e.g. incorporate static environmental systems or even complex systems comprised of sensors and applications for data enrichment to drive implementations of use cases focusing on group behavior.

Looking forward, future research may evaluate software that is available to implement the individual parts of the proposed architecture. Similar to what Pääkkönen and Pakkala presented for big data in particular [PP15], this could lead to a selection of software that is suited to tackle the different parts of the Reality Mining Application Architecture. Additionally, the evolution of smartphones probably will continue to make new kinds of sensors available. To enable Reality Mining to exploit all features of smartphones, this evolution needs to be accompanied by research in the field to maintain a state of the art.

References

[An16] Android, http://www.android.com/, accessed: 2016-04-10.

[Ap16a] Apache Hadoop, https://hadoop.apache.org/, accessed: 2016-04-10.

[Ap16b] Apache NiFi, https://nifi.apache.org/, accessed: 2016-04-10.

[Ap16c] Apache Spark, https://spark.apache.org/, accessed: 2016-04-10.

[Bi16] Bilal, M. et al.: Big data architecture for construction waste analytics (CWA): A conceptual framework. Journal of Building Engineering 6, pp. 144–156, 2016.

[CZ14] Chen, C.L.P.; Zhang, C.: Data-intensive applications, challenges, techniques and technologies: a survey on big data. Information Sciences 275, pp. 314–347, 2014.

[EP06] Eagle, N.; Pentland, A.: Reality mining: sensing complex social systems. Personal and Ubiquitous Computing 10(4), pp. 255–268, 2006.

[GJK16] Geisse, D. G.; John, I.; Kotstein, S.: A Survey on Smartphone Features Relevant for Reality Mining. In (Hertweck, D.; Decker, C. Hrsg.): Digital Enterprise Computing (DEC 2016), Böblingen, Germany, June 14-15, 2016.

[GS15] GSMArena.com, www.gsmarena.com/, accessed: 2015-07-07.

[Ke12] Kelly, H.: 5 ways the iPhone changed our lives, http://edition.cnn.com/2012/06/28/tech/mobile/iphone-5-years-anniversary/, accessed: 2016-04-09.

[Kr14] Kromer, M.: Modern Hybrid Big Data Warehouse Architectures. Business Intelligence Journal 19(4), pp. 48–55, 2014.

[LK15] Lee, J.-G.; Kang, M.: Geospatial Big Data: Challenges and Opportunities. Visions on Big Data, Big Data Research 2(2), pp. 74–81, 2015.

[PP15] Pääkkönen, P.; Pakkala, D.: Reference Architecture and Classification of Technologies, Products and Services for Big Data Systems. Big Data Research 2(4), pp. 166–186, 2015.

[QR15] Quint, M.; Rogers, D.: What Is the Future of Data Sharing? Consumer Mindsets and the power of brands. Columbia Business School, 2015.

[RR07] Richardson, L.; Ruby, S.: RESTful web services. O'Reilly, Beijing, Köln, 2007