



Reality Sensing, Mining and Augmentation
for Mobile Citizen–Government Dialogue

FP7-288815

D2.2

Report on strategies of mobile sensing in eParticipation

Dissemination level:	Public
Contractual date of delivery:	Month 12, 2013/01/31
Actual date of delivery:	Month 18, 2013/07/11
Workpackage:	WP2 - Policy Modelling and Visualization for Mobile-aware Dialogue
Task:	T2.2 - Classification of citizen contributions into eGovernment responsibility
Type:	Report
Approval Status:	PMB Final Draft
Version:	1.1
Number of pages:	73
Filename:	D2.2-v1.1.docx
Abstract This deliverable reports on current and envisioned strategies for mobile sensing in eParticipation scenarios. It contains surveys on current algorithms for situational mobile sensing and available data sources that can be used to support the process of eParticipation. These tools are mapped to our three use case scenarios and it is shown how their integration is advantageous.	
The information in this document reflects only the author's views and the European Community is not liable for any use that may be made of the information contained therein. The information in this document is provided as is and no guarantee or warranty is given that the information is fit for any particular purpose. The user thereof uses the information at its sole risk and liability.	



This work was supported by the EU 7th Framework Programme under grant number IST-FP7-288815 in project Live+Gov (www.liveandgov.eu)

History

Version	Date	Reason	Revised by
0.1	2013/05/06	First draft	Matthias Thimm
0.2	2013/05/14	First draft Sec. 4.1.+4.3.	Matthias Thimm
0.3	2013/05/21	First draft Sec. 4.2.1	Laura Nittylä
0.4	2013/05/21	First draft Sec. 3	Heinrich Hartmann
0.5	2013/05/23	First draft Sec. 4.2.3	Maite De Arana Agiretxe
0.6	2013/05/29	First draft Sec. 4.2.2	Marjolein Terpstra Pieter Minnigh
0.7	2013/05/31	First draft Sec. 2	Laszlo Kovats
0.8	2013/06/10	Sec. 4.2.1, Sec. 5.1	Laura Nittylä
0.9	2013/06/11	Sec. 3 General elaboration of texts	Heinrich Hartmann Matthias Thimm
0.10	2013/06/15	Beta version	Matthias Thimm Pieter Minnigh Laszlo Kovats Jon Casasempere
0.11	2013/06/25	Incorporated further comments Editorial work	Matthias Thimm
0.12	2013/06/26	Sec. 4.2.3 and Sec. 5.3	Maite De Arana Agiretxe
0.13	2013/07/02	Incorporated further comments Editorial work	Matthias Thimm
0.14	2013/07/03	Minor revisions	Matthias Thimm
1.0	2013/07/08	Incorporated updates by partners Editorial work	Matthias Thimm
1.1	2013/07/10	Final version	Matthias Thimm

Author list

Organization	Name	Contact Information
UKob	Matthias Thimm	E-mail: thimm@uni-koblenz.de
UKob	Heinrich Hartmann	E-Mail: hartmann@uni-koblenz.de
MTS	Laura Nittylä	E-Mail: laura.niittyla@mattersoft.fi
BIZ	Maite De Arana Agiretxe	E-Mail: komunikazioa.tekniko1@biscaytik.eu
BIZ	Jon Casasempere	E-Mail: jon.casasempere@biscaytik.eu
YCT	Marjolein Terpstra	E-Mail: M.Terpstra@yucat.com
YCT	Pieter Minnigh	E-Mail: P.a.Minnigh@yucat.com
YCT	Barry Peet	E-Mail: <not assigned to Live+Gov anymore>
YCT	Jens Steensma	E-Mail: j.m.steensma@yucat.com
EuSoc	Laszlo Kovats	E-Mail: laszlo.kovats@eurosoc.de

Executive Summary

This deliverable deals with strategies on mobile sensing for eParticipation scenarios. We discuss the current status of incorporating mobile technology into the process of decision-making and describe the aims of the Live+Gov project to further improve this status. One of the novel approaches is to explore reality mining for the purpose of obtaining improved context information on the citizen and, thus, to use this information to provide a better eParticipation experience. We survey existing approaches to reality sensing and provide an overview on algorithms for activity recognition, image analysis, and text analysis (including topic detection and sentiment analysis).

We survey existing Open Governmental Data repositories and discuss how they can be used for enriching eParticipation approaches. We also present LISA, a working showcase that illustrates the use of Open Governmental Data by offering a visually appealing way to explore this data for ordinary citizens.

In the context of our three use cases, we investigate both, approaches to reality sensing and external data sources, and give mappings how those can be applied. We also report on first experiences with data mining approaches in the context of the Urban Maintenance use case.

Abbreviations and Acronyms

FT	Fourier Transformation
CSV	Comma-separated Values
DCT	Discrete Cosine Transformation
EM	Expectation Maximization
GPS	Global Positioning System
HSL	Helsinki Region Transport
HTML	Hypertext Markup Language
JSON	JavaScript Object Notation
LDA	Latent Dirichlet Allocation
LISA	Local Information, Search, and Aggregation
MAP	Maximum A Posteriori Estimation
MLE	Maximum Likelihood Estimation
pLSA	Probabilistic Latent Semantic Analysis
PDF	Portable Document Format
RBF	Radial Basis Function
RDF	Resource Description Framework
RSS	Really Simple Syndication
SQL	Structured Query Language
SVM	Support Vector Machine
TXT	Text file
UGC	User-Generated Content
VLAD	Vector of Locally Aggregated Descriptors
XML	eXtensible Markup Language
XLS	Microsoft Excel Sheet

Table of Contents

1	Introduction.....	8
2	Reality Mining in the context of Citizen Participation and Stakeholder Engagement.....	13
2.1	Reality Mining in the Live+Gov Context	13
2.2	Reality Mining and Transparency	14
2.3	Reality Mining and Public Participation	15
2.4	Reality Mining and Collaboration.....	15
3	Algorithms for situational mobile sensing.....	18
3.1	Knowledge Management and Discovery	18
3.2	Mobile Sensing and activity recognition	19
3.2.1	Activity Recognition	20
3.2.2	Mapping Sensors to Activities	24
3.3	Mobile Sensing and visual recognition	25
3.3.1	Approaches for Visual Recognition.....	27
3.3.2	Mapping algorithms to sensing needs	29
3.4	Mobile Sensing and Topic Detection	30
3.5	Summary	35
4	External data sources for eParticipation	36
4.1	External data sources for Live+Gov use cases.....	36
4.1.1	Mobility Use Case.....	36
4.1.2	Urban Maintenance Use Case	38
4.1.3	Urban Planning Use Case.....	43
4.2	Further Open Governmental Data repositories.....	43
4.3	A working show case: LISA.....	47
4.3.1	Development status and technical requirements.....	49
4.3.2	Evaluation	49
4.4	Summary	49
5	Reality Mining and External Data Sources in Live+Gov	50
5.1	Mobility use case	50
5.2	Urban Maintenance use case.....	52
5.3	Urban Planning use case.....	53
5.4	A Case Study on Analysing Issue Reports.....	55
6	Summary.....	64
7	References	65
A	Analyses Data Set	72
B	Patterns.....	73

List of Figures

Figure 1: The process of activity recognition	20
Figure 2: Framework overview for visual recognition	27
Figure 3: A graphical model explaining the process by which documents are generated.	32
Figure 4: A Language model that specifies how the documents are generated in the real world.	33
Figure 5: Bayesian inference reflecting the process for estimating the parameters of the language model specified in Figure 5	33
Figure 6: Graphical model that corresponds to the generative process of LDA	35
Figure 7: Screenshot of LISA	48
Figure 8: <i>Crisp-DM sub model</i>	56
Figure 9: <i>Knowledge base infrastructure</i>	61

List of Tables

Table 1: Approaches for Activity Recognition	21
Table 2 Mappings of sensors to activities	25
Table 3: Description of key-point detection algorithms (+++ represent the fastest and + the slowest).	27
Table 4: Description of feature extraction algorithms (+++ represent the fastest and + the slowest).	28
Table 5: Mapping of Algorithms to Sensing Needs	30
Table 6: Top terms under each topic discovered from a text collection (topic labels are subjectively chosen)	32
Table 7: External data sources for the Urban Maintenance Use Case	42

1 Introduction

eGovernance was always influenced by new developments in the sector of information and communication technologies. Back in the 1980s digitalization in administration processes have been deployed on big mainframe systems. Proceeding with personal computers and the world-wide web in the 1990s individual eParticipation was possible for the first time. In the same time the first mobile phones reached the mass market. The functions of these phones were limited to the essentials and were still no replacement for the desktop PC in solving complex tasks like eParticipation.

Around the turn of the millennium, the term Web 2.0 was coined to describe web sites where users can interact and collaborate with each other. Tools like blogs, wikis and social networks “provide for user-initiated information sources, exchanges, and dialogues about public issues that complement”, began to extent and “compete with, government sources and services” [97].

Now, in 2013 simple mobile phones continued to evolve to smartphones with various small apps and multiple sensors. Until this time the costs of information and communication technologies were permanently reduced and thus the barriers to eParticipation were also lowered. In the post PC era mobile devices with build-in camera, GPS and internet access have become mainstream. People have become accustomed to be online anytime and anywhere and to communicate independently from time and space. As argued in deliverable D2.1 *Conceptual documentation on issues, organization and stakeholder assessment*, this new possibility “meets the demands of the citizens in a modernising society on the one hand and offers advantages to the public administration on the other.” By more easily accessible ways of communication “citizens can participate more easily in public decision-making, can virtually attend consultation procedures and are not necessarily required to be present personally” and “most importantly, technical solutions provide political organisation to people without forcing them to attach to any form of political party. In other words, solutions for modern participation may aggregate the interests of non-organised citizens and bring them to a format that could later be utilised and worked with politically. Accordingly, they bring all citizens not represented by any political organisation back into the political process and therefore play an important part in legitimising public decision-making and the whole political system.”

A smartphone as input device does not only lower the threshold of participation but provides with its built-in sensors much more than that. Considering all the sensors it is possible to turn it into a unobtrusive, and inexpensive human activity recognition system. Textual input or images taken by a smartphone can easily be enriched with contextual information like the actual timestamp and the geo-location of the user. Beyond that the user’s motion can be logged and it can be attempted to predict his next actions. This type of metadata opens the door to further analysis and better ways to understanding the citizens’ needs. By mining this additional data, policy makers get closer to the participant and are able to serve their demands in a context sensitively manner.

Input provided by a special mobile app can even be enhanced with more high-level information. With data mining approaches, which will be developed within the Live+Gov

project, the sensor- and metadata can be combined to better recognize the contribution of each citizen. Thus it is possible to find out in what situation and environment an input was produced. An opinion, which was produced by a cyclist, is maybe different from that one a pedestrian or a car driver has on the same topic. By identifying such differences, policy makers can derive a better picture of the current opinions of the various groups.

At least as demanding as the technical aspect of mobile eParticipation is the connection between the citizen's input and the institutional responsibilities. Provided the "willingness of the public administration to transfer political power to the citizens giving up a certain extent of political control", "the degree of institutionalisation of political participation of citizens" was identified as critical issue in D2.1 *Conceptual documentation on issues, organization and stakeholder assessment*. Mobile participation should not be seen as a short-term marketing tool but continued in follow-up processes to influence plans and decisions. This "requires not only once-off investment but also on-going support by personnel" [98].

The Live+Gov solution is innovating in two particular respects: first, it is establishing a closer cooperation between the citizen and the public authorities by improving and facilitating Citizen Participation on its three dimensions transparency, public participation and collaboration, cf. D2.1 *Conceptual documentation on issues, organization and stakeholder assessment*. Accordingly, it is targeting the input legitimacy of the state. Second, the developed solution enhances the quality of public service delivery by supplying better information to the authorities. Accordingly, it is targeting the output legitimacy of the state at the same time.

The fundamental strategy of the Live+Gov project is to use mobile technology for achieving both ends. In particular, the sensing capacities of mobile devices (smartphones) capable of acquiring data about the actual context of the citizens and the mining possibilities bringing structure and recognising patterns in the collected data are being utilised by Live+Gov. Therefore, the related strategies for mobile sensing are threefold:

- First, the Live+Gov system is relying on passive sensing performed quietly by the mobile device, the immediate analysis and automatic service provision fitting to the actual need of the citizen.
- Second, the Live+Gov system is enabling active sensing initiated by the citizen and the immediate supply of information to the citizen responding to their actual demands.¹
- Third, the Live+Gov system stores the acquired data, performs background analyses and provides the public authorities (as well as the citizens) with high quality information about the state of the municipality. This is enabling them to respond to problems and issues of the citizens better.²

The three strategies are presented in greater detail below.

¹ The first and the second strategies are therefore working on the input legitimacy of the state.

² The third strategy is therefore working on the output legitimacy of the state.

Passive sensing – automated feedback

The first sensing strategy of the Live+Gov project is to establish an automated process of delivering valuable information to the citizen according to its actual context. This information loop starts off with sensing the context of the citizens, continues with analysing the incoming input, relating the input to specific information, and is ending with the immediate communication back to the citizen.

The sensor data, which is being collected relates basically to the geo-location of the mobile device (of the citizen) and its movement. In the context of this first strategy, the sensor data is collected quietly meaning that once the citizen activates the Live+Gov application their location and movement is tracked automatically. Accordingly, the citizen does not need to perform any other action on the mobile device.

The collected sensor data, in particular from GPS, compass, and gyroscope, are combined to a movement profile showing how the citizen is progressing through the municipality's area. This relates to the type of movement (walking, biking, driving by car, driving by public transport) and to the specific means of public transport (bus or tram and which lines respectively). The second server-side analysis relates to recognising the activity of the citizen, e.g. whether the person is commuting or pursuing a free time activity.

The derived information from the server-side analysis is utilised in two distinct ways: first, the actual location and activity of the citizen is related to available external data-sources in order to providing them with relevant background information. Second, the totality of the data provided by all mobile devices in the Live+Gov system are aggregated for identifying current problems and nuisances e.g. in the traffic system or with public transportation. This information is also communicated to the citizens and matches their actual needs. Hence, if for example there is a disturbance on a specific bus line, the citizen on this particular bus line receive a corresponding warning, which is augmented by additional indications where to obtain further background information (e.g. links to external websites or services for planning alternative routings). Accordingly, the Live+Gov system provides a “real-time” service assisting the citizens to travel through the traffic as efficiently as possible.

Active sensing - automated feedback

The second sensing strategy of the Live+Gov project is to provide on-demand information about public matters and confronting the citizen with participation tools in the very moment they wish being confronted with. Therefore, as soon as the citizen activates the Live+Gov application and performs a certain action and is, for example, pointing at a certain location with the mobile device's camera they are presented specific information related to this location.

The respective sensing data (camera view, geo-location data) are assembled, the exact location identified and related to stored information on the Live+Gov servers or to external data-sources. This information is then communicated back to the citizens. It can take the form of simple background information about a certain construction site or a specific maintenance level, can include “Augmented Reality” views as well as a poll asking the citizen for their opinion.³

³ Accordingly, Live+Gov understands the analysis of text, which is gathered by mobile devices as mobile sensing as well.

Data-mining and policy support

The data derived by mobile devices make up for a whole new source of valuable information about the municipality if collected over a longer period of time. The Live+Gov project covers three issue areas traffic, infrastructural planning, and maintenance, which are of high relevance for each municipality. The applications are, however, not restricted to these three.

Traffic. Mobile data capturing the movement profile of the citizen make up for a completely new method for describing the traffic flows in the municipality. This relates basically to the fact that they provide a comprehensive view from the citizens' perspective using the entirety of the public transport system and not one means in isolation. Effectively, they show how the different modes of transportation (walking, biking, driving a car, taking a bus, taking the tram, etc.) engage with one another. Accordingly, the system recognises the walking/biking/driving distances from the citizen's starting point (home) to the first means of transportation, recognises waiting and walking times between different facilities of the public transportation, recognises the distance from the last stop to the final destination. This aggregated information can serve as a sophisticated policy-support tool for public authorities for their infrastructural planning.

Maintenance. Data-analysis in the maintenance context is also highly relevant. Knowing who reports what as well as knowing where citizens spot particularly many issues and problems with the public infrastructure greatly improves the overview of the public authorities and helps identifying systematic problems of the city. In general, statistical analyses of citizens' characteristics (age, gender, neighbourhood) show the public authorities what the people think and what certain groups among the citizens consider important; statistical analyses of the reports (issues, areas) show where the municipalities are confronted with particular problems that need special attention.

Infrastructural planning. If a municipality initiates a poll via mobile devices e.g. as in the Urban Planning use-case, the voting results as well as the comments by the citizens represent a highly valuable source of information. Analysed for voter characteristics like age, gender and neighbourhood, the public authorities obtain a detailed picture about how the citizens evaluate a certain project and which groups may protest. Furthermore, such poll results could also be used in the planning phase of an infrastructural project assessing the usefulness and sheer necessity of the envisaged building or service, respectively.

In general, the core mobile sensing strategies discussed in the three strategies above relate to utilising geo-location sensors, server-side-analyses algorithms, and the connection to external data-sources. In this deliverable, we provide a survey on methods for activity recognition, visual recognition, and text recognition using topic detection methods. Applying these methods using data obtained from mobile devices is beneficial for various scenarios in eParticipation. In our mapping of these technologies, we focus on application scenarios provided by the three use cases of Live+Gov, namely Mobility, Urban Maintenance, and Urban Planning. Besides the data gathered from mobile devices further data sources can enrich the eParticipation experience. Open Data initiatives such as Europe's Public Data⁴ provide information collected by municipalities and other governmental bodies on a wide variety of topics, e.g. census data, geo-location data, political data, traffic information,

⁴ <http://publicdata.eu>

criminality statistics, etc. While not long ago, this information was hidden from or not easily accessible by the citizen, Open Data initiatives promote free and easy access to this type of data. For eParticipation scenarios, this data is relevant as well as they can enrich both the analysis of data gathered from citizen as well as provide the citizen with useful information that may influence the way he/she wants to participate. Besides mobile technologies we also map these data sources to use case scenarios in eParticipation.

The remainder of this deliverable is organized as follows. In Section 2 we explore the general benefits of using mobile technologies, and specifically the area of *reality mining*, in eParticipation. In Section 3 we provide more technical details on different approaches to reality mining, such as activity recognition, visual recognition, and textual recognition. We continue in Section 4 with discussing external data sources. First, we report on specific data sources coming from our use case scenarios and, second, we give a general survey on Open Data repositories in Europe that can be exploited for both our use case scenarios and other eParticipation scenario. Third, we present LISA, a case study on how to use Open Data in order to bring benefit to the citizen. In Section 5 we bring mobile technologies and external data sources together and map them to the use case scenarios of Live+Gov. Finally, in Section 0 we conclude.

2 Reality Mining in the context of Citizen Participation and Stakeholder Engagement

Deliverable D2.1 has already sketched the large potential of mobile technology contributing to improved citizen participation and stakeholder engagement. The basic argument in this respect is that mobile devices can be used for identifying the actual context of the individual citizens and providing them with instantaneous information, which is fitting to their actual needs. Furthermore, gathered over a longer period of time, sensor data from mobile devices can be used as a valuable source of information improving the strategic mid-to-long-term planning of the public authorities. Accordingly, elaborating on the potential of improving “Reality Mining” through mobile technology and better recognising the context of the individual citizen represents the core innovation of the Live+Gov project.

Going more into the details, using mobile technology for better describing the context of the citizens has therefore two distinct effects – one on the citizens and one on public authorities. The citizen profits by having access to instantaneous information about certain issues in the municipality. This can take the form of real-time information about actual traffic conditions (e.g. an element of the mobility use-case) derived from the entirety of the delivered data from all users. Another option could be getting immediate information about infrastructural projects or being offered immediate participation or collaboration possibilities. Important is the fact that the citizen can access this service in the very moment they need or request it. Consequently, this lowers importantly the personal cost, effort and therefore the threshold for citizen participation and stakeholder engagement. Hence, this increases the probability that the citizens indeed participate.

The profits for the public authorities relate to the effectiveness and the efficiency of public service delivery. Harvesting the immediate input from the citizens’ mobile devices informs the authorities about the state of the municipality and about the attitudes and opinions of the citizens both in an unprecedented detail and speed: public authorities know more about the things to do in the municipality as well as about the priorities of the citizens and obtain both information almost immediately. Hence, they can better respond. Furthermore, utilising the technical systems necessary for raising this potential allows for far-reaching optimisation of the working processes within the administration. This can reduce the direct budgetary costs as well as the workload of the individual civil servant.

In the following, this section describes how “Reality Mining” is augmented by mobile technology and is implemented in the Live+Gov project. First, it defines the term “Reality Mining” and shows how this could be improved by sensing via mobile devices, the respective mining algorithms, and the combination with external data-sources. Doing this, it also outlines the components of the Reality Mining in the Live+Gov context, which are the mobile-side sensing techniques and the server side analysis.

2.1 Reality Mining in the Live+Gov Context

The reality-mining group of the MIT⁵ defines the term “reality mining” as “the collection of machine-sensed environmental data pertaining to social behaviour. This new paradigm of

⁵ realitymining.com

data mining makes possible the modelling of conversation text, proximity sensing, and temporal-spatial location through large communities of individuals. Mobile phones (and similarly innocuous devices) are used for data collection, opening social network analysis to new methods of empirical stochastic modelling” (realitymining.com).

The conception of “Reality Mining” in the context of the Live+Gov project is already very close to this definition. It will come up with solutions, which are facilitating the interpretation of “machine-sensed environmental data” (ibid.) via mobile devices. This allows for inferring the actual context and the activity of the individual citizen or groups of citizens. Additionally, it will utilise data-mining algorithms for recognising patterns in the social behaviour of the citizens enabling concrete and targeted municipal actions. However, Live+Gov’s understanding of “Reality Mining” goes a bit further. This is basically related to the fact that the fundamental aim of the project is not solely to derive information from the social reality of the citizens but from the public administration as well. In other words, the administration should be provided with better information about the reality of the citizens and the citizens should be provided with better information about the reality of the public authorities – the elected representatives and the public administration, respectively. This is represented by the element of transparency in the context of citizen participation and stakeholder engagement.

Consequently, Live+Gov uses sensors, which are integrated in mobile devices for capturing context data of the citizens and server-based data-mining algorithms for making sense of the abstract body of acquired data, giving them a substantial meaning and therefore recognising the activity of the citizens. Accordingly, the location and activity of the citizen is determining the most important information, which is delivered to them. In the Live+Gov context, this refers to how they could best participate in their actual context. Possible options could be to provide the relevant context related information (e.g. traffic information), presenting actual polls or surveys related to specific public questions that are probably interesting for the citizens, or offering certain forms of collaboration possibilities. Acquired over a longer period of time, the same information can be analysed for fundamental patterns in the municipality administration as recorded by the Live+Gov system, which is potentially revealing problems and issues that have not been detected before, e.g. sentiments of citizens. The public administration can then take these results into consideration when planning their mid-to-long-term policies (referred to in deliverable D2.1).

2.2 Reality Mining and Transparency

As it has already been discussed in deliverable D2.1, transparency measures need careful planning, well-conceived presentation and moderation. In particular, too much or complex information may overcharge the citizens (in his mobile context) and deter them from investing their time and effort for participatory purposes. This is especially relevant for transparency measures via mobile devices due to specific practical limitations. They relate to the smaller screen size and format, which is hampering the absorption of the presented information as well as the fact that the citizens are fundamentally in a mobile context and not particularly attuned to participation in a more complex manner.

Therefore, knowing about the context of the citizens and their actual activity (e.g. commuting or acquiring a particular kind of information via their mobile device) opens up the possibility of supplying specific and tailor-made information, which is fitting to their

actual needs and demands. This has the particular advantage that the citizen conceives the input as valuable, keeping their interest level for this issue high. Accordingly, the citizen may be motivated to gather more information at a later point in time or to participate more substantially. Therefore, it is of crucial importance for any kind of participation via mobile devices to deliver selected input as well as links to other sources of information or participation measures. If done properly, offering insights about public matters via mobile devices can function as a gateway leading the citizens to more advanced forms of participation (e.g. motivate them to participate in public deliberations or consultations).

2.3 Reality Mining and Public Participation

The potential of reality mining and mobile sensing for public participation arises from the possibility of confronting the citizen with a participatory process in the very moment in which they are most receptive and responsive to it. Accordingly, if a citizen is about to gather information about a certain issue (e.g. a construction site or the state of the public infrastructure) they can also be informed about the related participatory process. It may even be the case that the participatory process is partly or fully organised via the mobile device and can thus be executed immediately.

However, such participatory processes on mobile devices put high requirements on the presentation of the respective issue and the related alternatives or options. As transmitting substantial information via a mobile device is complicated due to practical limitations as well as the shorter attention span of the users (citizens) one needs to apply sophisticated techniques like for example “Augmented Reality” for presenting it. Visualising issues, options and alternatives makes it easier for the users to understand the topic and to take a meaningful decision. Therefore, in order to project an artificial model on the screen of a mobile device and showing the citizen what is planned on a specific construction site requires to connect the built-in sensing capacities of the mobile device like the camera, the compass, and instruments for identifying the geo-location (GPS).

2.4 Reality Mining and Collaboration

While reality mining and sensing technologies are importantly facilitating the transparency of public matters and are offering additional participation possibilities for the citizens, they are crucial for collaborative purposes between the municipal administration and the citizens. This relates essentially to the fact that collaboration is crucially facilitated by mobile technology. Traditional collaboration solutions via telephone, mail, or e-mail have been available in many municipalities but are necessarily implemented in rather complicated procedures.

After having realised an issue that they would like to report, citizens had to choose between rather cumbersome communication procedures. They could either make a phone call and explain their concern verbally and with all hampering misunderstandings or had to wait until they reach other means of communication like a computer or a piece of paper. In any case, the reporting required more effort and was often not even possible immediately. Therefore, citizens’ original urge to report was immediately weakened by the prospect of the large effort and/or by the time span that it would take to issue the report. Accordingly, the motivation of the citizen to collaborate was small, which had the effect that collaboration

did not take place very often – or at least as often to make a real difference for public maintenance or other areas of collaboration. Therefore, the facilitating momentum of mobile technology through a well-defined and easily understandable procedure and the possibility to make an immediate contribution is the central facilitating momentum for collaboration.

Hence, the citizen should be able to participate in the very moment they would like to do so. As soon as they want to report an issue (see: best practice example BuitenBeter) or feel the urge to become active, they can comply with their sentiment. The sensors of the citizens' mobile device facilitate this by allowing to take pictures of the issue to be reported as well as to record and transmit geo-location data. Citizens do not need to wait with their report until they reach other means of communication with the authorities and do not need to make ponderous explanations.

Therefore, knowing of the possibility to report these issues, citizens are beginning to view their municipality differently by paying attention to problems and flaws in the public infrastructure. This attitude is essential for the administration to get an adequate quantity of reports for making collaboration a real asset in the maintenance and planning of the municipality.

Another aspect is, however, even more important: effective and cost-efficient collaboration can only be implemented by utilising the sensing functionalities of mobile devices. The sheer possibility of sending pictures to the public administration, which are automatically augmented by geo-location information for the purpose of collaboration, adds the necessary level of detail to the report that is enabling the public employee to respond immediately from their remote desktop. In contrast, traditional communication systems like telephone, mail, or e-mail are far too inaccurate for clearly describing an issue. Consequently, reports by telephone or mail need to be double-checked by the local authorities. Accordingly, the administration has to send a person to the described location who adds the necessary details which, in turn, allows initiating the respective maintenance procedure. Accordingly, if being confronted with a high number of such reports or announcements the public authorities quickly reach the limit of their operating efficiency. Monitoring all the reports physically is far too costly and outweighs the positive effects of collaboration. Additionally, the process is way too slow for showing the citizen effectively how (well) the administration is working.

A further application of reality mining and mobile sensing in a collaborative way is tracking the movement of citizens and inferring their actual activity. Such activity recognition allows either for providing the users with tailor-made information that is relevant for them as well as deriving meaningful information for the administration about problems and issues in the public infrastructure.⁶ The Live+Gov project understands this as a form of collaboration because the citizen is explicitly allowing the Live+Gov system to record their movement profile by activating the respective functionality on their mobile device. However, inferring any meaningful information from the collected movement data requires additional server-based computing and sophisticated mining algorithms, which are correctly relating patterns in the data to a specific activity of the respective person.

⁶ This will be a fundamental element of the mobility use-case.

In the following this deliverable will go into greater detail of the different aspects of reality mining within the Live+Gov project. This relates mainly to situational mobile sensing, visual recognition and augmented reality as well as data-mining methodologies.

3 Algorithms for situational mobile sensing

The widespread adoption of smartphones and other sensor-capable devices and the availability of infrastructure capable of storing and processing sensor information have been the main enablers for the concept of Reality Mining. This new paradigm of data mining makes possible the modelling of conversation context, proximity sensing, and temporal-spatial location throughout large communities of individuals. Indeed, by collecting and aggregating the data streams generated from smart-phone installed sensors (e.g. GPS tracks, phone camera, location, and activity information) and combining them with user-generated content (UGC), we are able to gain insight into the dynamics of both individual and group behaviour [45]. What is particularly interesting is the fact that mobile and social network citizens can act as the living sensors of their city, its maintenance level, its mobility status, or even its infrastructure needs, by doing no more than their everyday activities.

In the context of Live+Gov we are interested in using the smartphone sensors to understand the citizen's context when interacting with his public administration. This is feasible both through passive collaboration strategies, where the data streams captured through the smartphone's hard sensors (i.e. accelerometer, GPS, etc) are transmitted, stored and processed to recognize the undertaken activity, as well as through active collaboration strategies, where the citizen sends visual or text input to inform his public administration about a certain issue, or suggestion. However, in order to sense the citizen's reality we need to employ a set of methods that are able to perform activity recognition, visual recognition and natural language understanding with topic detection. The goal of this section is to briefly survey the state-of-the-art algorithms used in these fields and discuss how they can be applied in settings for eParticipation. We start with giving some general comments on knowledge management and discovery in Section 3.1 and continue with approaches to mobile activity recognition in Section 3.2 . Afterwards, in Section 3.3 we survey approaches to visual recognition approaches and (very briefly) topic detection techniques in Section 3.4 . In Section 3.5 we conclude this section with a short summary.

3.1 Knowledge Management and Discovery

Al-khamayseh, Lawrence & Zmijewska [88] say that “over the past decade governments all over the world have been moving to providing services to their citizens via the web with varying degrees of success”. M-government is the extension of e-government to mobile frameworks. It also involves the use of mobile technologies in the provision of public sector services [89]. Kushchu & Kuscu [90] discuss some examples in their paper about law enforcement agencies in the USA that have adopted m-government technologies that lead to cost effective and efficient law enforcement programs by communicating among each other and access various information sources through the wireless network.

Dzhusupova, Ojo & Janowski [91] explain the need for knowledge management in government and explores the requirements of knowledge management for the public organizations”. In Singapore for example the government builds a knowledge management portal that enables employees to build capacity on each other's knowledge and therefore create better documentation of the procedure. In our research we discuss the value of gathering real time experiences from citizens in the public space to discover new knowledge from these real time experiences. Dzhusupova, Ojo & Janowski [91] say that “The capability

to gather and generate valuable knowledge from information is a defining feature of a modern state". Municipalities have to efficiently and effectively create, locate, capture, and share their organization's knowledge and expertise with the citizens to be able to gather and generate valuable knowledge. Only then can they improve the current government policies. This is why a knowledge base infrastructure is needed to automate and manage the process of gathering knowledge from citizens. Michael [92] states that a knowledge base infrastructure (he calls it Knowledge Management Architecture) "utilizes four primary resources, namely: repositories of explicit knowledge, refineries for accumulating, refining, managing, and distributing that knowledge, organization roles to execute and manage the refining process and Information technologies to support those repositories and processes".

When trying to extrapolate information from raw data we are talking about knowledge discovery. Data is useless unless information can be derived from it. On particular focus on approaches to knowledge discovery is *data mining*. "Data mining is the process of searching and analyzing data in order to find implicit, but potentially useful information. It involves selecting, exploring and modeling large amounts of data to uncover previously unknown patterns, and ultimately comprehensible information, from large databases" [93]. The requirement for forming a pattern is a minimum pair of two tags ($A \rightarrow B$), where an observation of A can predict an observation of B (e.g. Tree $\rightarrow$ Tree branches) [94]. One specific approach to conduct data mining is using the CRISP-DM methodology. The CRISP-DM methodology is described in terms of a hierarchical process model, consisting of sets of tasks described at four levels of abstraction (from general to specific): phase, generic task, specialized task and process instance [95]. CRISP-DM succeeds because it is soundly based on the practical, real-world experience of how people conduct data mining projects [95].

In the following, we take a closer look at knowledge discovery approaches for the purpose of recognizing activities and detecting concepts from images and text.

3.2 Mobile Sensing and activity recognition

The field of human activity recognition has been actively studied since digital sensors became cheaply available in the late 90ies [27], [28]. Modern smartphones offer a rich set of sensors, which are worn near the human body and allow performing activity recognition in an un-pervasive way. A variety of studies conducted since 2008 explore these new possibilities. Most recently (May 2013) Google integrated an activity recognition API into their Android mobile platform [40], [41]. This will most certainly lead to a vast increase in the usage of activity recognition in all kinds of application scenarios. Another on-going trend is the attachment of sensors to other clothing articles like shoes (Nike+iPod [42]), watches (Nike Sports Watch [44]) and glasses (Google Glas [43]). Data from these sensors will soon be available for data mining on smartphones and allow improvement of accuracy, and precision of current activity recognition methods.

The above mentioned efforts mainly focused on the recognition of activities like "walking" or "running" which are performed over a comparatively short interval of time. The task of extracting higher context information like "commuting to work" or "shopping" from sensor data, requires different methods and has attracted less attention (cf. [33], [34]). The Live+Gov project explores the use of activity recognition and context extraction in scenarios of eParticipation. In the following, we survey existing algorithms for activity recognition and

estimate their benefit for the different data mining tasks within the project, cf. D4.1. *Report on Live+Gov toolkit requirements and architecture*.

To automatically detect activities from a continuous stream of incoming sensor data a pipeline of signal processing and machine learning techniques is used. In the first step the stream of incoming data is divided into time windows of fixed size (typically 1-10 sec.) and for each window a set of features is computed. These features are filtering out relevant information from the raw signal. Typical features are standard deviation or modes of a Fourier transformation. In the initial training phase, test users collect annotated data samples for each of the required activities. We obtain a list of feature vectors together with annotations about the performed activity, which are stored together in a database. From this annotated sample data, data mining algorithms (e.g. decision trees, support vector machines) can derive activity classifiers of various sorts. The process is illustrated in Figure 1, cf. also the CRISP-DM model in Section 3.1 .




Figure 1: The process of activity recognition

For the actual recognition task, the raw sensor data is again sliced into time windows and features are extracted. The resulting feature vector is supplied to the trained classifier, which outputs the annotation which, based on the presented training data, is most likely to describe the performed activity.

The extraction of higher context information, like “working” or “shopping” is an additional step, which takes the stream of classified activities as input to build models or hierarchies of higher activity patterns.

3.2.1 Activity Recognition

There are several important choices involved in the process of activity recognition: the actual activity recognition algorithm, the context extraction algorithm, and the mapping between available sensors and the activities that have to be recognised. In the following

sections, we will explain those choices in more detail and describe how they have been taken in the literature.

The most important features for activity recognition that can be extracted from sensor data can be grouped into two different types: statistical features and frequency-domain features. Statistical features include, mean value, standard deviation, bin distributions/histograms, mean absolute deviation, correlation, and entropy. They can be applied to a real-valued time series (e.g. accelerometer, magnetometer, etc.). However, the above mentioned features do not take into account the ordering of the incoming sensor samples. Frequency domain features can detect regularity patterns inherent in activity sensor data. The Fourier transform (FT), discrete cosine transform (DCT), wavelet analysis have been successfully applied in [31], [32], [30] for that purpose.

Furthermore, the most important machine learning algorithms for the task of activity recognition are Naive Bayes, decision trees and support vector machines. Other algorithms that are used for activity recognition are k-nearest neighbours, multilayer perceptron, fuzzy basis function, fuzzy interference system and regression models cf. [24].

Decision tree learning is the most popular method for activity recognition. It has been successfully applied in [3], [4], [5], [10], [14], [17], [18], [19], [20], [22], and [23]. A decision tree consists of a tree which has the classification attributes as leaves and at each internal node a comparison between an individual feature to a fixed value (e.g. $\text{mean}(S) > 4$). The classification of an incoming feature vector is obtained by following the tree from the root in directions (left or right) depending on the comparisons. A decision tree can be automatically derived from training data, e.g. using the C4.5 algorithm [76], which uses on the concept of information gain to generate the tree structure.

Naive Bayes is a probabilistic classification algorithm which has been used in [14], [15], [19], [20], and [23]. It is based on the assumption that individual features are mutually independent and only depend on the class attribute. Using the Bayesian conditional-probability formula it assigns to each incoming feature vector a probability that it belongs to a given class. Despite its simple structure the naive Bayes classifier shows a high classification performance in practice, even in cases where dependencies between the individual features are given.

Support vector machines separates the feature space into regions divided by hyperplanes in such a way that training samples of the same category belong to the same region and the distance between the hyperplanes to the training samples is maximized. Of course such a linear separation is not always possible, but non-linear embeddings (“Kernel trick”) and the introduction of slack variables allow the application of SVM methods also in those cases. Support vector machines have been proven to be an effective tool for activity recognition in [5], [12], [14], [20], and [22].

Table 1 summarizes the literature on activity recognition with smartphone sensors (the column “additional sensors” refers to additional sensors besides accelerometer).

Table 1: Approaches for Activity Recognition

#	Additional sensors	Activities	Algorithms
---	--------------------	------------	------------

[3]	Microphone, GPS	sitting, standing, walking, running	decision tree
[4]	GPS, Wifi	walking, not-walking, in vehicle	decision tree
[5]	Gyroscope	lying, sitting, walking, running, in bus	decision tree, support vector machine (SVM), others
[7]	-	sitting, walking, cycling	naive Bayes
[8]	-	sitting, walking, stairs, running	others
[9]	-	sitting, standing, lying, walking, stairs, cycling	others
[10]	GPS, Wifi, Microphone	sitting, walking, running, in vehicle	decision tree
[11]	-	sitting, standing, lying, walking, running, jumping	others
[12]	-	sitting, walking, running, cycling, stairs, in vehicle	SVM
[13]	-	walking, stairs, falling down	others
[14]	-	sitting, walking, running, cycling, in vehicle	decision tree, naive Bayes, SVM, others
[15]	camera intensity, proximity	sitting, standing, walking, running	naive Bayes, others
[16]	-	lying, sitting, standing, walking, stairs, running	others
[17]	GPS	sitting, walking, running, cycling, in vehicle	decision tree, others
[18]	-	sitting, standing, walking, running, stairs	decision tree, others
[19]	-	walking, not-walking	decision tree, naive Bayes, SVM
[20]	-	standing, walking, running, stairs	decision tree, naive Bayes, SVM
[21]	-	walking, running/jogging,	others

		cycling, jumping	
[22]	-	standing, walking, running, stairs	decision tree, SVM
[23]	-	walking, running, in bus, in metro	decision tree, naive Bayes, others

As an example, the study presented in [4] was conducted in Chicago with the aim to provide the user with precise information about public transportation. A two-step process was implemented to detect if a user is currently using means of public transportation. In the first step, using accelerometer and GPS signals, the distinction between “moving” and “standing still” was made using a decision tree. This classification task was achieved with a very high accuracy of 97%. If movement was detected, GPS and WiFi were used to get location information, which was compared to timetable and location information of public transport vehicles.

In [19] the distinction “walking/not walking” was sensed with a special emphasis on the position of the smart-phone. It does make a difference for the detection if the smartphone is placed in a trouser pocket, in the shirt’s top pocket or inside a coat. The study showed that it is possible to detect the smartphone position with high accuracy and use this information to enhance the initial recognition task.

Another important aspect of recognizing activities is the choice of training data. The sensor data from a walking child and an old man can be severely different. In practice, the training data is often generated by a small group of people (often the researchers themselves), which can lead to a bias of the trained classifier. However, this problem was studied in [7] in more depth, where an experiment was conducted with a group of 8 persons. For each person a personalized classifier for the activities “sitting”, “walking”, “running” and “biking” was trained (“within person model”) and another one on the basis of the training data from the seven other persons (“cross person model”). The cross person model showed with 97% only slightly lower precision than the within person model (99%). This shows that the classification of the above mentioned activities is very robust against the choice of training data.

Although activity recognition is able to provide useful information about the citizen’s reality, there are many cases where we need to lift our interpretation one level higher for understanding the citizen’s context. As already mentioned, the extraction of higher context information requires an additional step, which takes the stream of classified activities as input and builds models or hierarchies of higher activity patterns. Some of the most popular methods for extracting higher context information from sensor data include bag-of-words, topic models, hidden Markov models and conditional random fields.

Bag-of-words methods count the number of times $N(a)$ that an activity a is performed in a given time window of observation. The resulting activity frequencies $(N(a_1), \dots, N(a_n))$ are treated as a vector and linear algebra methods borrowed from text analysis community, like LSA and ESA, can be applied to cluster activity vectors and extract context information. In

[34] the authors refine this model by considering n-grams of activities as their vector space basis.

Topic models are another approach for context extraction which was successfully applied in [33] and [38]. For a sequence d of detected activities, the possible context c that the user might have is modelled as a probability distribution $p(c/d)$. For each activity a and context attribute c we can assign a probability $p(a/c)$ that a is performed in context c . Counting the number of times activity a occurs in d gives a third probability distribution $p(a/d)$ which is related to the other two. Now, Latent Dirichlet Allocation is a method to reverse this equation to estimate the context distribution $p(c/d)$ for the given series d , cf. Section 3.4.

Finally, Hidden Markov Models have been used in [35] to cluster observed events like “closing the door” or “walking down an aisle” into “high-level scenes” such as “visiting the supermarket”.

3.2.2 Mapping Sensors to Activities

The following sensors are available in modern smart phones (cf. D1.1): a) Accelerometer, b) Gyroscope, c) Magnetometer, d) GPS, e) Microphone, f) WiFi / Bluetooth.

The accelerometer sensor measures the acceleration of the device in the direction of the three axes in m/s^2 . Gravity affects the measurement of this sensor, so that a smartphone resting on a table will experience an acceleration of 9.81 m/s^2 in the direction the z-axis. It is possible to filter out the gravity signal in software (e.g. using low-pass filter). The accelerometer sensor is by far the most important sensor for activity recognition. All of the studies found in the literature are using the accelerometer sensor, most of them even exclusively. In our own experiments (cf. D1.1 *Sensor Data Application*), we have verified that, for those activities we aim to recognize within our use cases, activity recognition is possible based on the accelerometer data alone. The accelerometer sampling frequency varies from 10Hz to 100Hz in the literature. Current smart phones typically offer sampling frequencies up to 50Hz. The interesting finding of [29] is that there is no significant improvement in classification accuracy above 20Hz for ambulation activities.

The gyroscope measures the rotation velocity around the three axes of the device. In [5] this sensor was used for activity recognition, but it remained unclear how much impact this additional sensor had on classification results.

The magnetometer measures the magnetic field in micro-Tesla (T) and can be used as a compass to sense the pointing direction of the device with respect to the magnetic field of the earth. This sensor is useful in specialized mining scenarios like “bus line recognition” but was not used directly on activity recognition in the literature yet.

Other used sensors include GPS, WiFi, and microphone. With GPS the position and velocity of a user can be detected on the expense of high battery usage [4]. WiFi was also indirectly used for location detection in [10]. With the aid of the microphone the environment noise was classified to trigger the activation of GPS when moving outside. Also the proximity sensor was used in [15] only indirectly to judge the position of the smart phone.

Higher context extraction has not been studied widely in the literature. Available methods build upon activity recognition systems. Topic models seem to provide the best result for context categories like “commuting”, “lunch”, or “office work”.

In the requirements analysis of D4.1 *Report on Live+Gov toolkit requirements and architecture* we identified the following activity recognition attributes, which we have to classify. On the one hand we want to identify basic activities such as

- Ambulation: sitting, standing, walking, running
- Transportation: cycling, driving car, in public transport
- Means of public transport: bus, train, tram, subway, metro, ferry.

On the other hand, we want to identify high-level context such as commuting to work, shopping, sight seeing (tourist), lunch routine, dinner routine, working.

As we have seen above, the ambulation and transportation tasks can be reliably detected with data from the accelerometer, statistical features like mean and standard deviation, and a decision tree classifier. Table 2 gives an overview on the mapping of activities to the sensors that are needed to recognize them.

Category	Attribute	Sensors	Features	Mining
Ambulation	sitting	Mainly Accelerometer.	Mainly statistical: mean, variance, bin distribution.	decision trees, SVM, naive Bayes
	standing			
	walking			
	running			
Transportation	cycling	Also: Microphone, GPS, Wifi, Gyroscope	Also: Fourier transform	
	driving car			
	in public transport			
	- in bus			
	- in train	Mainly GPS.	Direct comparison of location to realtime GPS data of service vehicles	
	- in tram			
	- in subway			
	- in metro			
	- in ferry			
	If not available: Wifi, GSM, Acc., Compass,.			

Higher Context	Attribute	Extraction Algorithms
Routines	commuting to work	First: bag-of-words approach. Then: Topic models.
	shopping	
	sight seeing (tourist)	
	lunch routine	
	dinner routine	
	working	

Table 2 Mappings of sensors to activities

Besides those sensors that “implicitly” collect data for the purpose of activity recognition (as discussed so far), most smart phones also provide an optical sensor in the form of a camera. While the use of this sensor is a more deliberate action of the user it also allows, under many circumstances, a larger amount of usable information for assessing the context. We will take a closer look on visual recognition in the following subsection.

3.3 Mobile Sensing and visual recognition

High-end mobile phones have developed into capable computational devices equipped with high-quality displays, high-resolution digital cameras, and real-time hardware-accelerated 3D graphics. All these functionalities enabled a new class of visual recognition applications,

which use the phone camera to initiate search queries about objects in visual proximity to the user. Pointing with a camera provides a natural way of indicating one's interest and browsing information available at a particular location. Once the system recognizes the user's target it can provide further information or services [46], allowing to link between the physical and the digital world [47]. In this context, visual search has been extensively researched in recent years [48], integrating mobile augmented reality [49] and outdoor coordinate systems [46] with visual search technology.

Moreover, apart from augmented reality and visual search the visual recognition technologies have also been used for automatic metadata extraction and content understanding, as a pre-step to multimedia indexing and organization. For instance, ALIPR (Automatic Linguistic Indexing of Pictures - Real Time) [50] was one of the first attempts by researchers to incorporate visual content in search engines, and frameworks like the one presented in [51] were proposed for the efficient large-scale object retrieval in web-scale databases. More recent methods that make use of visual recognition technologies for indexing and organizing user contributed content, include [52] where the images contributed by social media users are clustered to automatically identify the most important places of interest in a city, and [53] where the visual and tag information carried by flickr images is effectively combined to facilitate their semantic indexing.

In the context of Live+Gov and the envisaged eParticipation scenarios, visual recognition will be considered both in its online version, which will basically serve as an enabling component of mobile augmented reality, as well as in its offline version, which will be part of the server-side mining, aiming among others to automatically extract the category of an issue accompanied by an image, as well as to identify the issues referring to the same real-world incident. In the following, we briefly explain the process of image recognition, we review the state-of-the-art methods in this field and eventually we map the most prominent of these algorithms to the analysis modes envisaged by the eParticipation scenarios of Live+Gov.

As in the case of activity recognition, visual recognition is essentially a problem of pattern recognition that is considered to consist of two parts, the image representation and the classification part. In recent literature, the use of local features combined with some variation of bag-of-visual words (BOW) [54] is the most popular approach for image representation. More specifically, points of interest are detected and for every key-point a set of features are extracted so as to be encoded into a single vector using the BOW approach. Thus, the process of image representation mainly consists of the: a) key-point detection, b) feature descriptor extraction, and c) encoding process. The next step is to identify the pattern followed by the extracted image representation. This is typically accomplished using a classification process that relies on machine learning. Figure 2 is a schematic diagram of the aforementioned pipeline. More details about the process of visual recognition can be found in [55].




Figure 2: Framework overview for visual recognition

3.3.1 Approaches for Visual Recognition

The literature related to visual recognition focuses mainly on the algorithms for image representation and pattern classification. In the following, we provide an excerpt from [55] that presents the state-of-the-art algorithms for: a) key-point detection, b) feature extraction, c) feature encoding, and d) classification. Our goal is to emphasize on the advantages and disadvantages of each method and justify the proposed mapping with the Live+Gov eParticipation scenarios.

Key-point detection describes the process of detecting well-defined positions in an image that can robustly characterize its content. They are usually points of high interest (e.g. corners) that are used to represent important aspects of the image. The differentiating factors among the existing algorithms are their invariance in image transformations, computational time and repeatability. Table 3 summarizes the state-of-the-art algorithms for key-point detection. It is easy to identify in this table the trade-off between the good properties of translation invariance and repeatability versus the computational cost.

Table 3: Description of key-point detection algorithms (+++ represent the fastest and + the slowest).

Algorithm	Speed	Translation invariance	Description
Dense	+++	✗	Key-points are selected on a grid using a step of a fixed number of pixels.
FAST [59]	+++	✗	Key-points are selected based on the intensities of the pixels around the examined pixel

GFTT [60]	++	✓	Pixels are ranked based on a quality measure that is extracted using eigen analysis
HARRIS [57]	++	✗	Detects high intensity variations of sliding windows around the pixel and defines a corner based on a score
MSER [58]	+	✓	Detects connected regions with little change across several intensity thresholdings of the image
SIFT [61]	++	✓	Locations at minima and maxima of a Difference of of Gaussian function are selected as key-points
STAR [62]	+	✓	Detects the extrema of the Gaussian operator's Laplacian across scale
SURF [63]	++	✓	Uses integral images and box filtering techniques
ORB [64]	++	✗	Augments the FAST detector with a pyramid scheme and the Harris corner measure

Feature extraction attempts to describe the surrounding environment of each key-point so that it captures characteristic and indicative information of its visual content. The differentiating factors in this case are the rotation and scale invariance of the algorithm, its computational and memory requirements. Table 4 is a summary of the state-of-the-art algorithms for feature extraction. Again it is easy to identify the trade-off between the good properties of an algorithm (i.e. scale and rotation invariance) and its computation cost.

Table 4: Description of feature extraction algorithms (+++ represent the fastest and + the slowest).

Algorithm	Speed	Size	Rotation invariance	Scale invariance	Description
BRIEF [65]	+++	32	✗	✗	Bit string description of an image patch, constructed by a set of binary tests
ORB [64]	+++	32	✓	✗	Steered version of BRIEF according to the orientation of key-points
SURF [63]	++	128	✓	✓	Uses box filtering techniques and integral images
SIFT [66]	+	128	✓	✓	A histogram of orientations is computed in a 16x16 area around the key-point

After applying the key-point detection and the feature description algorithms, each image is represented by a set of n-dimensional vectors. In order to transform these vectors into a single vector representation we need to apply a *feature encoding* algorithm. The most popular method, borrowed from text retrieval, is the bag of visual words that has been used with the following variations.

Hard assignment: The local feature descriptors of the image are matched with the visual words of the vocabulary. A histogram of the visual descriptors is populated by adding ones to the corresponding bins.

Soft (kernel codebook) assignment [67], [68]: In this case instead of assigning a descriptor to a single corresponding visual word we assign it to k bins in a soft manner. More specifically, for every descriptor we add a quantity q to the bins of the k top nearest visual words. This quantity q is the Gaussian kernel (Radial Basis Function) distance of the descriptor and the visual word.

Vector of Locally Aggregated Descriptors (VLAD) [69]: In this case, firstly we assign each descriptor to its closest visual word. Then for each visual word a vector is calculated by accumulating all the differences of the assigned descriptors with the visual word. Finally these vectors are concatenated into a single vector representation.

Finally, after representing each image with a single feature vector, a model that learns the correspondence between image labels and features needs to be trained. One way to accomplish this is by using probabilistic methods, which try to find the joint distribution between labels and features (e.g. Bayesian Network [77]) or the conditional distribution of the labels given the features (e.g. conditional random fields [78]). There are also the tree decision algorithms [79], which attempt to map observations about an example to conclusions about the examples true class. Random forests [80] is an example of such an algorithm which constructs a number of random decision trees in a controlled way in order to obtain better predictive performance and generalization. Neural networks [81] are inspired by the structure and functionalities of human neural networks and attempt to capture the structure in the unknown joint probability distribution between observed variables. Finally, there are the algorithms that attempt to split the feature space so that the different classes are separated. Logistic regression and Support Vector Machines (SVMs) [56] are the most popular of this category.

As in the previous cases, the classification schemes are characterized by the performance vs. computational cost trade-off. If we take for example SVMs, which is the most widely used machine learning scheme in visual recognition, the computational cost for building a model and testing a new image, is considerably different when using an RBF kernel, or a linear one. Thus, depending on the underlying requirements we may favour different configurations.

3.3.2 Mapping algorithms to sensing needs

As already mentioned, the role of visual recognition in the context of Live+Gov is to either facilitate the image-based functionality of mobile augmented reality, or the server-side data mining through automatic metadata extraction, or clustering of similar items. Thus, there is a clear need both for an online algorithmic configuration that will balance the trade-off in favour of speed, and an off-line algorithmic configuration that will balance the trade-off in favour of recognition performance. Based on the extensive study that we have performed in [55] for gaining detailed insights into this trade-off, we plan to adopt the mapping proposed in Table 5.

Table 5: Mapping of Algorithms to Sensing Needs

Mode	Functionality	Attribute	Key-point detection	Feature extraction	Encoding	Classificat.
online	Augmented-reality	Image-marker recognition	surf	surf	vlad	SVM (linear kernel)
		Object-class recognition				
offline	Server-side data mining	Automatic metadata extraction (e.g. issue category)	fast	sift	vlad	SVM (RBF kernel)
		Clustering of similar issues				

3.4 Mobile Sensing and Topic Detection

The explosion of social media and especially forums, wikis, blogs and micro-blogs has created unprecedented opportunities for citizens to voice their opinions, but has created serious bottlenecks when it comes to making sense of these opinions. Policy-makers and citizens do not yet have an effective way to make sense of this mass conversation and interact meaningfully with thousands of others. As a result of this paradox, the public debate in social media is characterized by short-termism and auto-referentiality. At the same time, the sheer amount of raw data is also an opportunity to better make sense of opinions. In this direction, opinion mining [70] can be defined as a sub-discipline of computational linguistics that focuses on extracting people's opinion from the web data. More specifically, given a piece of text, opinion-mining systems analyse: a) which part is opinion expressing, b) who wrote the opinion, c) what is being commented. Similarly, sentiment analysis is about determining the subjectivity, polarity (positive or negative) and polarity strength (weakly positive, mildly positive, strongly positive, etc.) of a piece of text.

Both opinion mining and sentiment analysis cannot be considered as a new research theme. Automated methods for content analysis have been increasingly used, and have increased at least 6 folds from 1980 to 2002 [71]. These methods are typically based in long established computer science disciplines, such as Natural Language Processing, Text Mining, Machine Learning and Artificial Intelligence, Automated Content Analysis, and Voting Advise Applications. However, according to [70], since 2001 we have seen a growing awareness of the related problems and opportunities, accompanied by numerous papers published on the

subject. What is new today is the sheer increase in the quantity of unstructured data, mainly due to the adoption of social media (wikis, blogs, forums, comments) that are available for machine learning algorithms to be trained on. Social media content by nature reflects opinions and sentiments, while traditional content analysis tended to focus on identifying topics [72]. As such, it deals with more complex natural language problems.

Driven by the increased data availability and the need to analyse more complex concepts, there has been a shift from semantic-based analysis towards greater use of statistics and visualisation. A number of tools exist in the market that are primarily geared towards analyzing customers' feedback about products and services, and therefore skewed towards sentiment analysis that detects positive/negative feelings by interpreting natural language. Among these tools we may identify *twitrratr*⁷ that simply analyses terms based on a pre-defined glossary, *wordle*⁸ that relies on wordclouds and provides an appealing design solution that can serve as an entry level in the opinion mining market, and *uservice.com*⁹ that relies on crowdsourcing by allowing the users to submit feedback and rank other people ideas. Apart from these simple and free applications, there is a flourishing market of enterprise-level software for opinion mining, which much more advanced features that has been also used to detect hostile or negative communication in governmental context [73].

Technically, these tools rely on machine learning with regards to identifying and classifying relevant comments, through a combination of latent semantic analysis, support vector machines, "bag of words" and Semantic Orientation [74], [75]. Most tools on the market rely on a combination of machine and human analysis, typically using machines to augment human capacity to classify, code and label comments. Automated analysis, on the other hand, is usually based on a combination of semantic and statistical analysis, which due to the sheer increase in the quantity of available datasets is becoming more and more important. Some of the latest trends in research include: improving the accuracy of algorithms for opinion detection, reduction of human effort needed to analyze content, semantic analysis through lexicon/corpus of words with known sentiment for sentiment classification, identification of policy opinionated material to be analysed, computer-generated reference corpuses in political/governance field, visual mapping of bipolar opinion, identification of highly rated experts.

In the following we have a closer look on the field of topic detection and mining. Techniques and methods from topic detection and mining have the potential for automatically organizing, understanding, searching and summarizing large collections of data. We can use those to uncover the latent topical patterns that dominate a collection as for example uncover trends in issue reporting in the Urban Maintenance Use case. However, the focus of Live+Gov lies in the application of methods for activity and visual recognition, rather than topic detection and mining. Therefore, we will discuss topic detection and mining techniques only briefly.

⁷ <http://twitrratr.com/>

⁸ <http://wordle.com>

⁹ <https://www.uservice.com/>

Table 6: Top terms under each topic discovered from a text collection (topic labels are subjectively chosen)

Topic: Weather	Topic: Semantic Web	Topic: Economy
cold	web	credit
snow	semantic	money
warm	language	market
weather	RDF	business
winter	knowledge	rate
morning	schema	economy
degree	service	home

In probabilistic topic modelling a ‘topic’ is seen as a multinomial distribution over a vocabulary that assigns high probability to a set of words that tend to appear in the similar documents. A qualitatively ‘better topic’ is that in which words that have a high probability are semantically related to each other and a human subject is able to say that “these words are about X”, where X can be any domain for example, weather, semantic web, economy etc. as show in Table 6. A “Topic Model” is seen as a model of the generative process by which documents are created and captures the word co-occurrence patterns in a document corpus to produce semantically coherent topics. It shows how the observations can be generated by realization of the random variables while traveling along the edges of the directed graph. As the model captures the casual process of observed data generation, therefore, such models are referred to as generative models. To explain how the documents are generated in the real world, we specify a language model where each data point corresponds to the words in the document with topics as latent variables. Figure 4 represents this model in the form of a graph. In the graph, z represent the latent variable (topic) of the model, w represent the observed variable (words), while α , β and θ are model parameters. Given a particular word, our goal is to find a posterior distribution over topics that explain the data best.




Figure 3: A graphical model explaining the process by which documents are generated.

The generative process that corresponds to the document generation example above is shown in Figure 5. The task of the Bayesian inference is to invert generative process as

shown in Figure 6 and find parameter values that explain the observed data best. From the observed words in a set of document, we would like to find which language model is most likely to have generated the set of documents. This involves inferring the probability distribution over words associated with each topic, the distribution over topics for each document, and often the topic responsible for generating each word.




Figure 4: A Language model that specifies how the documents are generated in the real world.




Figure 5: Bayesian inference reflecting the process for estimating the parameters of the language model specified in Figure 5

A variety of statistical models have been proposed for topic-based analysis and modeling text documents. To name few of them are unigram model, mixture of unigram model [88], latent semantic analysis (pLSA) [86] and Latent Dirichlet Allocation (LDA) [83]. pLSA and LDA are the two widely used models for topic discovery in text data.

Unigram and Mixture of Unigram Model

Under the unigram language model, the words of every document are drawn from a multinomial distribution Θ . The unigram model uses strong independence assumption that words are drawn independently from a multinomial distribution and throws away all conditioning context, and estimates each term independently. This is,

$$\mathcal{P}(w_{1:n}) = \prod_{i=1}^n P(w_i|\theta)$$

Then with the approach of and [88] the generative process corresponds to the following steps:

1. For each word in the document

- Draw a topic $z \sim \theta_z$
- Draw the word from topic specific distribution $w \sim \theta_z$

The document probability is given by,

$$\mathcal{P}(w) = \sum_z P(z) \prod_{n=1}^N P(w_n|z)$$

Probabilistic Latent Semantic Analysis

The assumption in mixture model that each document is generated by one topic is relaxed by probabilistic latent semantic analysis (pLSA). In pLSA each document is generated by the activation of multiple topics, and each topic is modeled as multinomial distributions over words and is given by,

$$\mathcal{P}(w, d) = P(d) \sum_z P(w_n|z) P(z|d)$$

However, pLSA model uses a distribution indexed by training documents, which means the number of parameters being estimated in pLSA grow linearly with the number of training documents. The parameters for a k-topic pLSA model are k multinomial distributions of size **V** and **M** mixtures over the **k** hidden topics. This gives **kV + kM** parameters and therefore linear growth in **M**. The linear growth in parameters suggests that the model is prone to overfitting in many practical applications.

Latent Dirichlet Allocation

Latent Dirichlet Allocation (LDA) overcomes the problems of pLSA by using the Dirichlet distribution to model the distribution of the topics for each document.




Figure 6: Graphical model that corresponds to the generative process of LDA

LDA (Figure 7) is a Bayesian network that generates a document using a mixture of topics. In its generative process, for each document d , a multinomial distribution Θ over topics is randomly sampled from a Dirichlet distribution with parameter α , and then to generate each word, a topic z is chosen from this topic distribution, and a word w is generated by randomly sampling from a topic-specific multinomial distribution Φ . The robustness of the model is greatly enhanced by integrating out uncertainty about the per-document topic distribution Θ .

Parameter Estimation

Different approaches have been used for parameter estimation in the topic-based probabilistic models. These approaches include Maximum likelihood estimation (MLE), Maximum a posteriori estimation (MAP) and Bayesian estimation. Expectation-maximization (EM) [85] is used to find the direct estimates of model parameters for MLE and MAP approaches. Alternatively, variational EM [83], expectation propagation [87], Gibbs Sampling [84] algorithms provide approximate inference of the model parameters in Bayesian estimation. Blei [83] suggested to use approximate methods where parameters Θ and Φ can be integrated out because explicit estimate methods suffer from problem of local maxima in topic models.

3.5 Summary

In this section we discussed different technical approaches for obtaining information from citizens that can be used to improve eParticipation approaches. We gave a short introduction into general data mining approaches and presented surveys on activity recognition methods, visual recognition methods, and text extraction methods such as topic detection. For Live+Gov, the data fed into these approaches all originate from the citizens' mobile devices, either via implicit sensing using mobile sensors, photos, or textual input. With these approaches we have the possibility to make sense out of a large amount of data collected and, thus, in better understanding the citizens' situation we can improve our targeted eParticipation scenarios.

4 External data sources for eParticipation

The information gathered using sensors of mobile devices already provide important information for assessing the context of citizens and, thus, already enrich the eParticipation process. Other important resources for enriching the eParticipation process are open governmental data repositories which are discussed in this chapter. In particular, in Section 4.1 we give a description of important data repositories that will be used to realize our three use cases on Mobility, Urban Maintenance, and Urban Planning. We continue in Section 4.2 with briefly reviewing some further open governmental data repositories that can be used to incorporate publicly available information on municipalities and other official governmental data into the process of eParticipation. In Section 4.3, we describe a working application that has been developed for Live+Gov as a showcase to illustrate the benefits of providing open governmental information to citizens in a visually appealing way. We conclude with a summary in Section 4.4

4.1 External data sources for Live+Gov use cases

Political decision-making is often criticised for its lack of transparency or for referencing ambiguous figures that cannot be verified easily. However, governmental institutions collect usually large amounts of data, which are of interest to the general public, but often only a small amount of datasets is made publicly available. The Open Data and eGovernment movement wants to address these issues by promoting open access to these kinds of datasets. This is to improve transparency and legitimacy of complex processes instantiated by policy decision-makers in the government or local authorities. Furthermore, these processes should be accompanied and supported by an active citizenship, i.e. ensure participation and the possibility for feedback at any time during the process in order to provide for better decision making in the future.

The Open Data domain comprises mainly three actors: data providers, policy decision-makers and the public. A Data Provider publishes data, preferably in a raw, unprocessed format, that is accessible to everyone. Note that data providers include private parties as well as public parties, though public or governmental agencies are the most common ones. Decision Makers need to push Open Data to increase transparency and to legitimate their proposals and decision. Active citizens or the public are furthermore encouraged to participate in the decision processes. They are asked for feedback and thus require a transparent processing and presentation of the data in a way that is commonly comprehensible.

This section gives a list of specific data repositories for our three use cases.

4.1.1 Mobility Use Case

In the mobility use case, most external data sources focus on public transportation information, which will be used to allow an in-depth analysis of both citizen preferences and needs, as well as to detect issues in the field of public transport. The three most important data sources all focus on HSL public transport environment (HSL - Helsinki Region Transport), providing information about routing, vehicle location, and detected disruptions in operations.

The Public Transport Registry (JORE) in Helsinki provides information about public transport routes and stop locations. The most interesting data for the mobility use case can be found from three files:

- *lines.dat* includes information about which lines are running on which transport means (tram, bus, metro, ferry)
- *pysakki.dat* includes stop information
- *reittimuoto.dat* including route geometry information for lines

All files are downloadable from HSL's FTP-interface (requires login). This data provides a solid base for being able to detect lines automatically from smart phone sensors, as well as help to find the most affected areas in the foreseen jam detection module.

Consortium partner Mattersoft provides data from HSL Live - a public transport information system providing location data from trams and a few buses operating in the HSL area. The location of trams are registered on a once-per-minute basis using GPS and transmitted to a central server. The current location is then compared to the schedule and based on the difference, a real time forecast to following stops are given. The real-time data of tram locations can be received from HSL Live. The description of the data interface can be found from <http://developer.reittipas.fi/pages/en/other-apis.php> (HSL Live push interface).

HSL provides official disruption information for public transport users. Disruption info provides information on all exceptions leading to delays of over 15 minutes in public transport, involving trams, buses, metro and ferries. The reported exceptions are all based on observations by passengers, drivers, or authorities, and they all need to be confirmed by HSL. Thus, this makes the information become available with a certain delay. This is one reason why including social disruption messaging from passenger to passenger is interesting to be tested for the HSL authorities as these two messaging channels support each other.

The information can be accessed using simple HTTP GET interfaces. Currently the information is available at the HSL website and displays on stops. Citizens can also obtain current information by SMS or e-mail (requires registration). More information about the disruption info can be found at <http://www.poikkeusinfo.fi/pinfo/poikkeusinfo?lang=3>.

Other additional data for the mobility use case is also possible to be included in order to provide a more comprehensive understanding of current status of transportation and reasons behind unexpected events. For the city of Helsinki, where the use case will be set up, a large number of different types of regional open data are available, see also Section 4.2 on Finland. There, the most profitable provider for open data is *Helsinki region infoshare* (www.hri.fi), containing over 1000 datasets. Most useful types of data that could be considered to be included in the use case at some point would of following types:

- weather
- recent traffic related maintenance activities
- public events (concerts, sports, etc.) with massive participation of citizens
- municipality information (areas)
- demographic information (population density, age structure of citizens)

Incorporation of the above types will be evaluated at a later point in time.

4.1.2 Urban Maintenance Use Case

The Dutch Government is stimulating transparency within the National Government, Provinces, and Municipalities by encouraging the publication of open data sources.

The National Government has a specific website, dedicated to bundle all these datasets: <https://data.overheid.nl/>. This website currently contains 5712 dataset, which are available to the public. Many of these datasets are published by municipalities and are municipality specific. On most topics, a bundled and standardized dataset for the entire Netherlands is – as far as can be seen – not available.

Another web site that publishes open datasets in the Netherlands is <http://opendatanederland.org/>. The homepage text describes why the initiative for this website is started:

We are delighted that such a large amount of open data of the Netherlands is available. Often these open datasets are mentioned on the web site of the publishing organization. Unfortunately, these web sites are not known to a broad public. OpenDataNederland.org is solving this by describing all Dutch open datasets in one catalogue and making it possible to search through this catalogue. (Translated from: <http://opendatanederland.org/2013-05-28>)

OpenDataNederland currently contains 591 open datasets. Datasets on OpenDataNederland can be searched by keyword or filtered by Organization that published the data, License type, (file) Format, Category, Region, Content (audio, visual, geographic, numeric, textual, video, happening). It is possible to show the available datasets on a map of the Netherlands. In this map, the municipality where the open dataset is published can be selected, and a list with sets from that municipality is shown.

Aside from the collection of datasets, OpenDataNederland has bundled 36 Open Data tools, tools that can be used to edit or visualize open datasets, as it is explained on the website: 'often it is unclear how a dataset can be used practically'.

For the Urban Maintenance use case apart from internal data sources – within Yucat, gathered with BuitenBeter – external data sources are also relevant for the use case. Table 7 takes a closer look on some of the open data repositories that are investigated with respect to their possible advantages in the Urban Maintenance use case.

Type	Examples of this type of data	Source	Relation to Urban Maintenance use-case (examples)	Additional information (free, static/dynamic feed?)
Municipal information – municipality web sites and/or	- Areas- neighborhoods - contact information - mayor - ...	Wikipedia http://www.wikipedia.org/ Official national information on multiple topics: http://www.overheid.nl/	Being used	Free Updated Free
Municipal boundaries	Shape-file data	Open-street maps based shapes	Being used	Free Updated if changes arise
Demographic information	- population density - education levels - safety - age - ...	Statistics Netherlands* http://www.cbs.nl/en-GB	To provide more information about municipality	Free Updated periodically
Weather information	- Real-time weather information in XML feed - Historical weather data, etc.	Real-time weather information: http://xml.buienradar.nl/ Historical data: http://www.knmi.nl/index_en.html	For specific problems that could arise (e.g. storm damage)	Free Dynamic Free

				Static
Information about local festivities in the municipality		National calendar? Individual websites of municipalities?	For specific categories to report (e.g. specific litter as a consequence of these festivities or damage to items in public space)	No explicit source found
(Seasonal) Health issue information	- hay fever/allergy risk - air pollution levels - oak processionary risk	Hay fever/allergies: https://www.lumc.nl/con/1070/85683/105795/?setlanguage=English&setcountry=en (Dutch) http://www.pollennieuws.nl/ (Dutch) Air pollution: http://www.lml.rivm.nl/data/verwachting/fijnstof.html (showing expected air pollution levels per municipality - Dutch) http://www.nsl-monitoring.nl/monitoring-nsl/inleiding/ (showing air pollution around roads, highly detailed and files also exportable to CSV and shapefiles - Dutch)	For specific problems that could arise (e.g. oak processionary, smog/air pollution, hay fever /allergies)	Free Dynamic
Maintenance quality norms (standards)	National norms for quality of public space	http://www.crow.nl/nl/Publicaties/publicatiedetail.aspx?code=557 (Dutch)	For information on how public space should look like (e.g. height of the grass)	Not free Static

Maintenance quality levels (specific actual location quality level 'policy')	Municipality-specific norms for quality of public space	Individual websites of municipalities	For information on how public space is maintained (e.g. frequency of mowing the grass)	Free Static?
Municipality budgeting?	Budgets on public space by municipality	Individual websites of municipalities	Information on public budgets, e.g. per maintenance issue	Free Changes every year
Maintenance 'news' / plans	- information about maintenance			
Zoning plans and participation possibilities, strategic spatial plans	- zoning plans - strategic spatial plans	http://ruimtelijkeplannen.nl/web-roo/roo/bestemmingsplannen?tabFilter=JURIDISCH# (zoning plans for municipalities and strategic spatial plans - Dutch)	Information on current and future policies per area	Free Dynamic?
Major Bicycle routes	- bicycle streets (major bicycle routes)	Individual websites of municipalities. For example, for Utrecht: http://www.utrecht.nl/images/dso/infraprojecten/fiets/fietsroutes.html (Dutch, interactive map) No map for entire Netherlands at once	Geographical representation of major bicycle routes	Free Dynamic?
Design of public space	- location of bins - location of benches	https://data.overheid.nl/ ; http://opendatanederland.org	Geographical representation of objects in public space	Free Static?

	- location of street lanterns - location of utilities (hospitals, movie theaters, GP's, etc.) - etc.			
Waste dumps	- location of waste dumps	https://data.overheid.nl/ ; http://opendatanederland.org	Citizens can provide creative ideas for the reuse of waste dumps	Free Static?
Other open data	- route for sprinkling salt when it's icy - location of sports venues - location of fire stations, police stations, etc. - etc.	https://data.overheid.nl/ ; http://opendatanederland.org		

Table 7: External data sources for the Urban Maintenance Use Case

4.1.3 Urban Planning Use Case

For the Urban Planning Use Case no specific data repositories are considered to be taken into account other than the general data repositories discussed in the next section. This data repositories can enrich the Urban Planning Use Case by e.g. providing municipal budget information when displaying planned projects.

4.2 Further Open Governmental Data repositories

In the following we provide a thorough survey on further existing (European) portals for publishing open governmental data, together with supported data formats. These data repositories will be used to augment the data used in eParticipation scenarios by taking information such as census data, geographical information, political information, etc. into account.

European-wide data repositories

- Europe's Public Data (<http://publicdata.eu>)
Formats: XLS, CSV, PDF, XML, HTML, RDF, TXT
- European Union Open Data Portal (<http://open-data.europa.eu>)
Formats: XML, HTML, XLS, PDF, SQL, RDF

Balearic Islands

- Dades Obertes (<http://www.caib.es/caibdatafront/>)
Formats: CSV, XML, RDF, XML, HTML, RSS, PDF

Belgium

- Data.gov.be (<http://data.gov.be>)
Formats: XML, XLS, HTML
- Public Data Belgium (<http://publicdata.belgium.be>)
Formats: misc.
- Open Data Antwerp (<http://opendata.antwerpen.be/datasets>)
Formats: misc.

Austria

- data.gv.at (<http://data.gv.at/>)
Formats: PDF, TXT, XLS, HTML, CSV, JSON, GML, KML, RSS
- Open Government Data Austria (<http://gov.opendata.at/site/>)
Formats: CSV, XLS, HTML, TXT, SQL, RDF+XML
- Open Government Data Niederösterreich (<http://data.noe.gv.at/>)
Formats: CSV, ZIP-Downloads
- Open Government Data Graz (<http://data.graz.gv.at/>)
Formats: CSV, WFS
- Linz Open Data (<http://data.linz.gv.at/>)
Formats: CSV, PDF, XML, TXT
- Open Government Data Vienna (<http://data.wien.gv.at/>)
Formats: CSV, KML, GeoRSS, GML

- Open Government Data Tyrol (<http://data.tirol.gv.at/>)
Formats: CSV, SHP

Denmark

- Datakataloget (<http://data.digitaliser.dk>)
Format: PDF, HTML, TIFF, RDF, SPARQL, XML, CSV, XLS(X)
- Open Data Aarhus (<http://www.odaa.dk>)
Formats: PDF, XML, HTML, TXT, DOCX, XLSX, DWG

Germany

- GovData (<https://www.govdata.de>)
Formats: PDF, HTML, JS
- Berlin Open Data (<http://daten.berlin.de>)
Formats: XLS(X), CSV, HTML, JSON, RSS
- Open-Government-Data-Portal Rhineland-Palatinate (<http://www.daten.rlp.de>)
Formats: CSV, GeoTIFF, XLS, XML, PDF, HTML
- Open Data Baden-Württemberg (<http://opendata.service-bw.de>)
Formats: CSV, XLS, PDF
- Open Data Portal Bayern (<http://opendata.bayern.de>)
Formats: misc.
- Offenes Parlament (<http://offenesparlament.de>)
Formats: TXT
- RIS Munich (http://www.ris-muenchen.de/RII2/RII/ris_startseite.jsp)
Formats: PDF
- Offenes Köln (<http://offeneskoeln.de>)
Formats: CSV, SQL, PDF, JPG, TIFF

Finland

- Helsinki Region Share (<http://www.hri.fi>)
Formats: XML, JSON, XLS(X), (Geo)TIFF, CSV, SHP, DWG
- Laatusuomen Kattava Kartoitus
(https://www.suomi.fi/suomifi/tyohuone/yhteiset_palvelut/avoin_data)
Formats: misc.
- Open Data in Jyväskylä (<http://data.jyvaskyla.fi>)
Formats: XLS(X), RSS, Excel
- Open Data Tampere (<http://tampere.fi/avoindata>)
Formats: CSV, XLS

France

- Data.gouv.fr (<http://www.data.gouv.fr/>)
Formats: CSV, XLS(X), XML, XSD, TXT, DOC
- ParisData (<http://opendata.paris.fr/opendata/jsp/site/Portal.jsp>)
Formats: CSV, XLS, TIFF
- Grand Lyon Data Catalog (<http://catalogue.data.grandlyon.com/>)
Formats: misc.

- Open Data Le Mans (<http://www.lemans.fr/page.do?t=2&uuid=16CB26C7-550EA533-5AE8381B-D7A64AF8>)
Formats: CSV, PDF, KML
- Open Data Bordeaux (<http://opendata.bordeaux.fr/>)
Formats: misc.
- Open Data La Rochelle (<http://www.opendata.larochelle.fr/>)
Formats: misc.
- Open Data Toulouse (<http://data.grandtoulouse.fr/>)
Formats: CSV, PDF, JPG, XML
- Open Data Montpellier (<http://opendata.montpelliernumerique.fr/>)
Formats: CSV, XLS, TXT, XML, RDF, DOC, PDF, ODT, PPT

Greece

- geodata.gov.gr (<http://geodata.gov.gr/geodata/>)
Formats: misc.

United Kingdom

- data.gov.uk (<http://data.gov.uk/>)
Formats: CSV, XLS, HTML, PDF, RDF
- London Data Portal (<http://data.london.gov.uk/>)
Formats: CSV
- DataGM (<http://www.datagm.org.uk/>)
Formats: CSV, XML, XLS(X)
- Open Data Birmingham (<http://www.birmingham.gov.uk/open-data>)
Formats: JSON, KML
- Open Data Warwickshire (<http://opendata.warwickshire.gov.uk/datasets/>)
Formats: XLS, CSV, RSS, PDF, KML

Ireland

- Dubl:nked (<http://www.dublinked.ie/>)
Formats: XML, JSON
- Fingal Open Data (<http://data.fingal.ie/>)
Formats: CSV, XML, KML

Italy

- OpenDataHub Italia (<http://www.opendatahub.it/>)
Formats: CSV, XLS, XML, HTML, JSON
- City of Rome Open Data (<http://dati.comune.roma.it/>)
Formats: CSV, JSON
- City of Milan Open Data (<http://dati.comune.milano.it/>)
Formats: CSV
- Provincia Lucca Data catalog (<http://opendata.provincia.lucca.it/>)
Formats: RDF, JSON, XML, XLS(X), CSV
- DATI.Piemonte.it (<http://www.dati.piemonte.it/>)
Formats: misc.

- Open Data Comune di Bologna (<http://dati.comune.bologna.it/>)
Formats: CSV, KML, SHP

Netherlands

- Netherlands Public Data Catalogue (<http://data.overheid.nl/>)
Formats: CSV
- Open Data Nederland (<http://opendatanederland.org/>)
Formats: XLS, ODS, XML, DOC, HTML, RDF
- Amsterdam O+S (<http://www.os.amsterdam.nl/>)
Formats: PDF
- Rotterdam Open Datastore (<http://www.rotterdamopendata.nl/>)
Formats: misc.
- Centraal Bureau voor Statistiek (<http://cbs.nl/>)
Formats: misc.
- Rijksoverheid Open data (<http://www.rijksoverheid.nl/opendata>)
Formats: JSON, XML

Norway

- data.norge.no (<http://data.norge.no/>)
Formats: CSV

Spain

- OpenDataBCN (<http://www.bcn.cat/opendata/>)
Formats: CSV, RDF, XLS, XML, PDF
- Open Data Catalog Navarra (<http://opendata.navarra.es/>)
Formats: CSV, JSON, ODS, RSS, SHP, XLS, XML
- Catalonia Open Data Catalog (<http://dadesobertes.gencat.cat/>)
Formats: XLS, CSV, XML, RDF, (Geo)RSS, KML
- Open Data Portal Galicia (<http://abertos.xunta.es/>)
Formats: XML, CSV, RSS, HTML, XLS, JSON, PDF

Sweden

- Öppna data (<http://oppnadata.se/>)
Formats: PDF, XLS, PPT, XML, DOC, JSON
- Open Stockholm (<http://open.stockholm.se/>)
Formats: XML, WMS, JSON
- OpenUmea (<http://www.openumea.se/>)
Formats: ODS, CSV, XML

Switzerland

- City of Zurich (<http://data.stadt-zuerich.ch>) & -\\
Formats: misc.

Slovakia

- Data.gov.sk (<http://data.gov.sk/>)
Formats: CSV, XSD, XML, HTML

Please note the above list is not exhaustive but already shows that opening data for public use is gaining a lot of attention. The data provided by these repositories can be included in our target eParticipation scenarios for e.g. providing better context information for citizens providing issue reports.

4.3 A working show case: LISA

In order to illustrate the usage of open governmental data and raise awareness to the Open Data Movement, we developed LISA (Local Information, Search, and Aggregation) as a best practice showcase.

From a citizen's perspective it is often hard to figure out if and where governmental data regarding the topic of interest is publicly available. Consider a person, let's call him Tom, who is moving to a new city and would like to find the best place for renting an apartment. Tom is 30 years old. Since he does not own a car, he wants to live close to the city center to have everything within walking distance. Therefore, he also needs a good connection with the public transportation to get to work. Moreover, he also sets value on living in an area with mostly younger people and recreational areas. All data, which is needed for such an evaluation, can basically be found on the Internet. However, the data is spread over various data sources and offered in different (proprietary) formats. Raw statistics and customized charts about demographics, infrastructure, public transportation, and geographic points of interest need to be aggregated manually. Hence, an automatic integration and aggregation of Open Data is necessary in order make all the data more accessible to the public.

Open data is usually published using either standards like RDF or using propriety data formats like Excel sheets or PDF files. In particular, in application areas that make use of governmental data aggregation, reasoning, and visualization of data in heterogeneous formats is mandatory.

LISA is a web-based application that utilizes data from different sources in order to assess the attractiveness of some location for a particular user. For example, the attractiveness of a city is a very subjective assessment that is based on different criteria depending on the age and life circumstances. Often it is difficult to bring own ideas in accordance with local conditions. Also, there are many different factors that can play a role in assessing a place. Consider the following toy use cases:

1. The Schmidt family moves to Berlin for employment reasons. But in what district should the new home be located? Father Herbert wants to reach the working place quickly, Mother Petra is looking for good shopping facilities and close proximity to a park, son Max still has to go to kindergarten and daughter Clara to school and wants to use sports facilities.
2. Volker Mueller wants to open a steak house in Bremen. But where should the location of the restaurant be? Factors such as competitors, target market, infrastructure, rents, etc. must be considered to find the optimal location.
3. Mrs Friedrich from the Office for Social Work in Wiesbaden wants to make the city more family-friendly. But where are already enough kindergartens, playgrounds and sports facilities, and where has to be build something new? Where do families live without enough playgrounds in the neighbourhood?

As the above toy use cases are concerned with issues related to attractive spots in a city, they can easily be seen in the context of e.g. the Urban Maintenance use case of Live+Gov.

Functionality

Above we gave three examples, for which LISA offers a solution. LISA will determine the attractiveness factor of a region based on various weighted data sources and displays it in an aggregated and integrated form clearly to the user. This is possible because more and more authorities provide urban, rural and federal data sources, such as location-based data and statistics to the general public. This allows citizens to receive interesting information about their cities and individual neighbourhoods. Unfortunately, so far in many applications lack a combined integration and aggregation of relevant data sources so that users are presented with information selected from multiple data sources, integrated and aggregated, and user preferences are prioritized accordingly.




Figure 7: Screenshot of LISA

LISA allows the user to select individual important indicators, prioritize and store them in a user profile. Then, an individual location-based attractiveness factor is calculated and clearly displayed on a map. The user can see at a glance the streets, neighbourhoods or regions that are interesting to him because they are marked accordingly. This allows arbitrary regions to be analysed quickly and easily using various individual criteria.

- Integration of local and open data
- Clear visualization, for example with a map
- Individual selection criteria, for example attractiveness as a place of residence, ideal location for business or indicators for urban planners
- Closer to citizens through transparent management decisions and publicly available data

- Direct feedback to the management of a city

4.3.1 Development status and technical requirements

LISA is realized in a prototypic way as a browser-based JavaScript application. The realization uses standard and open web technologies, like HTML5, JavaScript, OpenStreetMap¹⁰ and HTML5-Canvas. LISA is able to use freely available data sets from the public sector. All points of interest, e.g. location of schools, parks, pharmacies etc. (points of interest) used in LISA are derived from OpenStreetMap. On-going work is concerned with a more flexible way of incorporating new and heterogeneous data sets [1].

LISA is designed as a web-app, which provides an intuitive interface for exploration and map-based visualization of arbitrary data. The underlying framework is based on standard web technologies, like HTML5 and JavaScript and uses common data exchange standards like JSON, therefore LISA can easily be extended to further data sources or other places.

4.3.2 Evaluation

LISA won the first prize in der category “Apps single developer” of the open data challenge Apps4Deutschland¹¹ and is also honoured additionally with the “Public Brain Award” by Vodafone.

4.4 Summary

We presented a survey on open governmental data repositories that can be used for enriching the eParticipation process in the context of Live+Gov. We also presented LISA as a show case that demonstrates the usefulness of open governmental data.

Both sensor data gathered from the mobile device of citizens and open governmental data provide the opportunity to obtain a broader understanding of the situation of the citizen and the situation of the governmental body. In the next section, we continue with a discussion on strategies how these components can be incorporated within the Live+Gov system and, specifically, how they can enrich our three use case examples.

¹⁰ <http://www.openstreetmap.org/>

¹¹ <http://www.apps4deutschland.de>

5 Reality Mining and External Data Sources in Live+Gov

This chapter relates the information presented so far on current best practices of reality mining technology and external data sources with our use case scenarios. For that we present both planned work but also already implemented approaches. Our discussion follows the three use cases of Live+Gov. We discuss the Mobility use case in Section 5.1, the Urban Maintenance use case in Section 5.2, and the Urban Planning use case in Section 5.3. In Section 5.4, we conclude with a case study we performed on using data mining techniques for analysing issue reports.

5.1 Mobility use case

Citizen participation in the mobility use case focuses on two parts: providing travel information and issue reporting, both depend heavily on mobile technology and mobile sensing. To be able to get this type of information a set of mobile sensing techniques are required to define where and how citizens travel. With the mobile application it is possible to collect data and later aggregate specific information on the server side about citizen movement profiles, tendencies, and issues related to these.

Tracing citizens' movements is the core functionality for the travel information collection. With the help of accelerometer data of the mobile device, it can be automatically sensed when citizens are using different means of transportation; are they walking, waiting or in a bus, cf. Section 3.2. To form a comprehensive picture of the journey, a data analysis needs also to be done. In the analysis, the location data is compared to public transport data, so that it can then be defined which service lines are used during the journey. By collecting this type of data and with the help of several data analysis and aggregation methods over a period of time, travel tendencies are detected and movement profiles of various users can be formed. This information can then be used to help public transport planners to find inconveniences on services. Besides techniques for mobile activity recognition (see Section 3.2 general data mining and knowledge discovery techniques have to be applied.

In particular, in the mobility field trial there are three major types of activity recognition through advanced algorithms for situational sensing of the travel information. All three types are linked to the next one in the order of presentation. The first problem focuses on recognizing activities of the citizen, whether they are walking, running, in a vehicle etc. The recognition is based on the citizen movement activities, constantly monitored and detected by the accelerometer sensors of the mobile device. Approaches for this problem are used as basis for the detection of usage of public transport travel chains. With this information it is possible to detect if a user constantly needs to have several different activities (e.g. walk, take bus, walk, take train, walk and take bus) in a row before reaching a destination, which would indicate of a poor service level between two areas. The second problem is about detecting the service lines when citizens are using public transport. Once detected that the citizen is using public transport, there is a need to define which line they are on. Based on the location data (based on either GPS-data or positioning via WiFi) of both, the citizen and a vehicle, this information can be extracted. Alternatively, for the vehicles, if no exact location can be defined, the assumed location of a vehicle can be defined based on routing and timetable information. It is important to detect the specific line for two purposes: in order to detect preferred routes and in order to be able to provide personalized messages about

issues regarding the specific line. The third major problem is for traffic jam detection. Based on vehicle locations, the algorithm compares the status of each tracked vehicle regarding their schedule and once a jam is detected, an alert to the citizens who are detected by the second algorithm to be likely to travel through the congested area are alerted.

Solutions to these problems provide a whole new insight on citizen behavior, patterns and preferences as their activities can be understood based on completely new information and better services can then be provided. From the reports many different types of information can be seen; where the citizens travel, what transport means they use, how direct routes they travel, how long the journey takes and how much they need to wait. Services can then be improved when reports show significant lack on service between certain areas or within a certain time range. This means that new service lines are created or routing is changed for better coverage between specific areas. Another possibility is that schedules are modified to decrease waiting times if long interchanges are constantly detected in journeys between specific areas. Also, when gaining knowledge of citizens' tendencies on their travel, they can be provided with more personalized information than before. Most of the personalized information relies on either other citizens' reports or external data sources, from where only the relevant information for each user is filtered. Real-time information about their current journey and possible expectations is generally considered welcome and not disturbing. By receiving more personalized information, citizens will have more confidence in using public transportation services and also increasing the satisfaction of service level.

By reporting issues it is possible for the citizens to state their opinions, report problems and give feedback on transport related issues in different categories. Providing categories for reports improves the possibility to later direct the report to the responsible service provider. Reported issues will include GPS-data of the location of the issue and by using the camera of the mobile device, the user can attach an image describing the issue in more detail. A textual description of the issue provides explanation and personal opinions about the issue. By categorizing different issues (using topic detection mechanisms as described in Section 3.4 , the reports can be filtered and redirected more efficiently to the responsible service provider and image recognition with the attached images can help even better to categorize issues even more efficiently, cf. Section 3.3 . Issue-reporting is an important feature in increasing citizen participation and improving services; authorities receive direct feedback from the users and are able to react immediately when needed. When noticing improvement on services, that is done based on the given reports, citizen satisfaction increases. When informing the users about actions taken based on the reported issues, direct interaction between authorities and citizens increase. In this, the messaging feature and textual input is of great importance.

Public participation and collaboration on the other hand includes also direct communication with citizens in order to improve services, which lead to the possibility to intervene in underlying issues more efficiently. Most of the public participation and collaboration is based on manual features – issue reporting and alerts regarding jams, delays and other exceptions. This not only provides authorities information about the traffic situation, but also helps other citizens to re-consider their plans. Also automatic features, such as travel tendency detection is a strong part of the collaboration and participation based on the fact

that by using the application, citizens are willingly providing authorities valuable information to help their planning processes. This information provision is also a major actor in the improvement of mid-to-long term planning when learning about citizen preferences, areas of low service level and problematic areas can be done on a more stronger and reliable basis.

For the mobility use case, the external data sources include mainly data of the public transport tracking and routing, cf. Section 4.1 . The external data sources in all problems described above are used to determine passenger preferences and jams in a more detailed and area-specific way. HSL Live offers real time vehicle data, which are used to detect the vehicles users are on, by comparing the user location to the vehicle location. Also jam detection is heavily based on the HSL Live data; tracked vehicles' locations are constantly compared to the scheduled location and, once several vehicles are detected to be increasingly delayed, a jam is defined for that area. Approaches for both problems are developed mainly based on the vehicle data and therefore the presence of real-time data is valuable for the use case.

More static data of the public transport is available for the use case from Jore-register, which includes information about routing and timetables. Jore-data is important for detecting the lines passengers use. Detecting is based on comparing the passenger location to line information. Jore-data is also used for determining stops affected in jam detection-module. Disruption information and open data sources are used to provide additional information to citizens, thus helping to form a comprehensive reporting channel of the system and also to gain more understanding of prevalent conditions to support planning processes when looking into reports.

5.2 Urban Maintenance use case

In the Urban Maintenance use case the main medium of communication is the issue report. The current version of the *BuitenBeter* mobile application allows citizens to take a picture of a real-world issue they would like to report. It also allows the citizen to annotate the picture with a comment. Besides raising awareness to particular issues and triggering a process for repair, a municipality can further exploit this information by analyzing it and extract patterns that might be useful in obtaining a larger understanding of the underlying situation. In a first case study, tag annotations on pictures of issue reports have been analyzed to detect co-occurrence of issues, see Section 5.4 for more details on our approach.

In general, both pictures and comments of issue reports can be analyzed using the methods described in Section 3.3 and 3.4 respectively. Furthermore, employing activity recognition approaches (Section 3.2) to detect the activity of the citizen reporting an issue, the issue report can be enriched by including this information as well. Further analysis of this information on the server-side will also yield a better understanding of the particular citizen's context and the general situation of the citizens in the municipality.

General reality mining methods such as activity recognition methods also allow the implicit submission of issue reports, i.e. the submission of issue reports without the explicit

involvement of the citizen himself. For example, using sensor-mining techniques on accelerator data can detect potholes in streets, simply by monitoring the citizen while driving. This passive sensing (see also Section 1) does not require the citizen to actively report issues and, therefore, may lower the threshold for participation of many citizens. Similarly, using visual recognition techniques the system can detect e.g. broken lamps when making videos or taking pictures.

Visual recognition techniques are particularly useful when analyzing pictures from issue reports. Pictures of similar issues can be used as the input for learning approaches in order to obtain a classifier for future issue reports. In the current version of *BuitenBeter* pictures of issues have to be manually inspected and annotated with a description about what issue this picture really is. With the above classifier the daily-work of the municipality can be facilitated as the exact issue is detected automatically.

A statistical analysis of the citizens' demographic information helps to improve policies of the municipality, in particular, when it comes to specific policies of specific areas. By analyzing e.g. age, education, and location distributions of the citizens posting a particular issue, correlations between demographic information and issues can be detected. This might raise awareness and lead to a change in the assumed quality expectations of e.g. public parks or streets, depending on the location. Demographic information on the submitters of reports can also help the municipality to advertise participation for specific user groups that are currently underrepresented in the participation process.

Taking both, information of external data sources about e.g. weather conditions and specific reports on weather-related issues into account, also allows the training of predictors so that further issues can even be prevented by reinforcing certain infrastructures. In general, preventing future issues instead of fixing reported ones is a key benefit for extracting patterns and other useful information from obtained data.

5.3 Urban Planning use case

Citizen participation in public life is increasingly gaining relevance and mobile applications open a new channel for this, where location and recognition of activity or daily routine can give extra information. Through the mobiles applications that will be offered through the Urban Planning use case local entities will be able to increase transparency in their decision making processes. Information will be offered to the citizens about actual plans of the municipality, including technical and economical details related to the issue, or anything that can be of interest at a certain moment in time.

The possibility of tracking the location of the mobile phone and thus offer information based on this location, with the augmented reality layer enhancing the user's experience and perception of the world around them, personalizes the way a citizen is reached directly on the handheld personal device. At the same time, feedback is also personalized and more details about the citizen participating can be used in the system, without interfering with privacy and personal data protection. This also means that the information is direct feedback from the citizens, and can serve as an input for the decision making process. Moreover, this can make a global impact in the community, as there are measurable results of the public

opinion about issues. Direct communication through these mobile applications allows a new wave of open government solutions that allow participation and collaboration, and at the same time administrations offer transparency. However in this case, most of the information is user input that is provided by the user either the first time the application is used or given for specific planning issues.

The aim of these strategies is to develop methodologies for a sustainable urbanization in terms of social, environmental and economic sustainability, enhancing public facilities and services, contributing to municipal integration (urban-rural), enhancing city safety, creating civic capital and promoting citizenship, and at the same time learn about citizens preferences. The benefit and reason for using mobile phones for this eParticipation initiative is not so much the challenges of technology involved in terms of activity recognition and data mining, but more related to the opportunities that augmented reality and visual recognition techniques (see Section 3.3 offer for presentation of the scenery).

In the Urban Planning use case, sensing means both hard sensing, related to sensors such as the GPS and compass for location detection and the augmented reality view, but also reality sensing, which has a much broader interpretation and is the major advance of this use case as it creates valuable collective information about the perception of the world that surrounds citizens and how they feel about it, cf. Section 2. This reality sensing and extracting intelligible information from people is a challenge that can totally change public life and make citizens and administrations better understand each other and have a more positive impression of each other.

The first way in which hard sensing can be used for the use case is related to location and orientation detection, extracted from GPS sensors or mobile triangulation techniques and compass, that are necessary for the augmented reality layer. Extracting activities through activity recognition approaches (see Section 3.2 will enhance the report that a citizen has about a planned project, as it will give context. For example, a person who votes about the use of a bridge, on his way to work, has interest in this issue relating to commuting to work. However if it is done so when sight-seeing, the interest of this issue would be related to tourism of the area. Therefore the results can be analyzed in their context so they generate greater impact.

The second way in which this location can be used is to segment the decisions or votes provided by users based on the location with respect to a facility or issue of urban planning, related to the augmented reality view. This can give insights in cases where there is dependence on the side from which something is looked at, or different features on sides. The information provided by citizens can be considered reality sensing feedback, and this can also be studied in order to segment the results depending on different factors, such as age, gender, education level, etc.

For the Urban Planning use case external data comes principally from the local council. Information about related plans, or zoning plans can be external data of interest to give more insights on the issues. Maps of the area with location of points of interest can also complement the issues planned. Information extracted from these sources will be used to

give users context information for better understanding and to make the application useful for them as they will receive insights about the plan. Then specific information for the plan will need to also be shared by the council in order to customize each issue.

For the analysing data obtained from the reports of citizens, a specific knowledge-structure is envisioned. The system will provide a knowledge base, which will be used to collect information about participation of citizens and their profile. For data mining procedures the data will be anonymized. This system will keep different types of information. First, it will store the input from citizens, the opinion issued on the issues raised. This information shall contain the direct input delivered by people. The system will also store information about age, gender, and the address of citizens. This information will also actively provided by users. Second, the system will also provide information obtained through mobile sensors without direct citizen participation, like the type of transport used, if the citizen is walking, time, etc. This type of information must be treated in compliance with very specific legal requirements, since the data mining personal data have legal implications that must be taken into account. By analyzing the information, it is possible to know the opinion of citizens considering different profiles such as age or gender. For example, combining the provenance information about the socioeconomic data with the demographic data of the municipality, it is possible to estimate the opinions from different socio-economic profiles of the municipality.

The main objective of this system is getting a multidimensional view of the data, using different dimensions to analyse it, and crossing the information with other data sources owned by the municipality, like socioeconomic mapping, habits, etc. to get a clear vision of the input of citizens and their context.

5.4 A Case Study on Analysing Issue Reports

Use case partner Yucat (YCT) is currently providing a mobile application (*BuitenBeter*), which citizens can use to report “real time experiences” of the public space. Real time experiences can be different things namely, problems, emotions or suggestions concerning the public space. To give a example of these experiences; a problem could be “dog feces on the street” the emotion with it could be frustration (towards the owner or government for not cleaning it if it has been reported multiple times) and a suggestion could be something to improve the public space, for instance a dog bin that contains bags and room to deposit. The real time experiences are sent as reports containing information about the location, date, problem type and an image. The image can contain more information about the real time experience. To extrapolate this information from the image the process of image tagging can be used to add descriptive tags to the images [96]. A tag is a textual concept that humans can understand (e.g., Car, Tree and House) that is added to describe information that is contained in the image.

In the following, a report on our first experiences in using reality-sensing techniques in the Urban Maintenance use case is given. For that, the following research question is investigated: *“Can a knowledge base infrastructure be used to automate and manage the discovery of implicit, but potentially useful information derived from a mobile government application, to improve government policies?”* It discusses the translation process from real-time experiences into explicit knowledge that can be used to improve government policies

on the public space. We researched how reality-mining methods can be implemented and effectively used with the data obtained from m-government mobile applications. Based on this research a prototype knowledge base infrastructure will be set-up to manage and translate the real time experiences from citizens to further improve government policies.

5.4.1 Research Methodology

The data used are the reports and images sent to the municipalities. These images and reports are obtained by citizens with smartphones who have BuitenBeter running on their smartphone and reported real time experiences from anywhere within the municipality in a period of 5 months.

The reports and images are stored in databases. The data used is a sample set of the database for a period of 5 months in text format. Personal information about the citizens was not included in this text dump, only the information about the report. Apart from the reports, also copies of all the images corresponding with these reports are included in the sample set. From the total of 1002 reports a sample of 300 reports were taken to represent the population (Margin of error = 0.05, Confidence level = 0.95 and response distribution = 0.5). The reports selected had all fields filled in and had a valid image (some reports had images but the image was corrupted, or blacked out). Specific attributes used by the mobile application were removed. All calculations were done on this sample. All reports were grouped by district. This was done to compare problems between the different districts of the municipality.

Image Tagging

Images obtained from the mobile application are unstructured and have to be structured with the help of descriptive tags. The process of image tagging was done manually by experts. The images were viewed and the contents were summarized into keywords that best described what the contents of the image were. For all images 6 descriptive tags (e.g., Street, House, Tree, Trash) were added. Tagging provides us with a way of structuring the images without adding structure in the way that they are stored. To prevent inconsistencies or redundant synonyms of tags a set of rules were set for the process of adding the descriptive tags. These rules consist of the use of singular or plural. A rule to determine when a object should be included in the 6 descriptive tags based on the distance and relevance to the main problem type.

Data Mining Approach

The CRISP-DM model is used to create and test an implementable and reusable data mining process, because the objectives were quite clear from the start and the data did not contain a lot of errors for data mining. A subset of this model that focuses more on data preparation and data mining is used.



Figure 8: Crisp-DM sub model

Figure 8 shows a version of the CRISP-DM reference model, which presents us with the iterative cycle, and different phases that are used when adopting CRISP-DM. The data was prepared by first analyzing the descriptive tags that were added by manual image tagging. Tags that were too closely related were reformatted to more generic tags, because no prior knowledge was known about correct tags to be used and meaningful patterns. First, an exploratory data mining process was done by writing a JAVA program that will find all possible tag pairs in the sample¹². The images used in our sample have 6 descriptive tags. The JAVA program reads all records and takes these 6 tags and calculates all possible pairs (pairs of two and three) that can be made with the combination of these 6 tags. The pairs are stored in a list. The program then writes the results to a file containing all pairs that were found and their frequencies in the sample. The results were a huge set of pairs. Second, With these patterns a more explanatory data mining process was done by looking at the frequencies of the tags found and doing more analytical tests to find which had significant relationships to the 14 main problem categories.

5.4.2 Results

For a better analysis of potential problem areas within the municipality, first, the reports have been categorized by district. The municipality was divided into 7 districts (A, B, C, D, E, F and G). This makes it easier to find differences or relationships between potential problem areas within the municipality. There are 14 main problem categories: (1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13 and 14). With 234 reports district B had the most reports. For all districts except district B, the main problem category Other was most reported. The first analyses of the data set can be seen in Appendix A.

Patterns

A sample of 300 reports was selected to be used for the pattern finding. From this sample there were 6 reports that had images that were either blacked out or corrupt, these were removed. The end sample used contained 294 reports. From this sample there were 283 unique descriptive tags. Each report has 6 descriptive tags that can be any combination of the 283 unique descriptive tags. The top 20 descriptive tags with the highest frequencies are:

- Tree,
- Car,
- House,
- Street,
- Grass,
- Pavement,

¹² In this first approach we use a simplified variant of the Apriori-algorithm for mining association rules, cf. [99].

- Bicycle road,
- Building,
- Trash,
- Tree branches,
- Bush,
- Post,
- Path,
- Pavement tile,
- Dangerous,
- Sand,
- Broken,
- Light Post,
- Wheat, and
- Nuisance.

These are descriptive tags that were taken from the images. Indicating that citizens are experiencing things in the public space around these objects. The descriptive tags add extra information to the reported main problem category. The pairs were counted for combinations found between these descriptive tags and the main problem categories using the JAVA program for data mining, see Appendix B. Patterns between these tags and the 14 main problem categories revealed relationships between the location of problems reported by looking at the descriptive tags from the images. The main problem category (1) had 51 reports. From these reports 47 had a pair with the tag Tree, 30 with the tag Tree branches and 12 with bush. These patterns are expected to be found, because they have strong relation to the concept of the problem category (1) and confirm that citizens are reporting the right real-time experiences for this main problem category. Interesting to see is that more than half of the reports had patterns with the tag House with 15 pairs and the tag Building with 11 pairs. Indicating that these reports happen near populated areas. The main problem category (6) with 13 reports on the other hand had 10 pairs with the tags Tree, 3 pairs with the tag Grass, 6 pairs with the tag Bush and 3 pairs with the tag Sand but no patterns were found with the tags building or House. Indicating that this problem does not happen in populated areas, but in a more remote area with trees and bushes.

Patterns found confirm real time experiences from the citizens and add more context information of the surroundings of the experience. The top 40 tags found is a good starting set for further analyses of the data set. Two pairs were more frequent but have many relationships between problem categories and make it harder to mine meaningful information from these patterns. Three and four pairs form more meaningful information in relation to the main problem category but the frequencies are lower and are less significant in the data set. A larger sample with more than 6 descriptive tags can find more meaningful patterns. A challenge here is defining which patterns are interesting. These patterns can have some significance in their relationship with the main problem categories, but are of no

real interest to the municipality. To solve this problem a set of requirements for specific aggregated information needs should be predefined by the municipality. The requirement should be more in a vocabulary term than in terms of attributes. This will form a sort of template that municipalities can use to match the patterns found with specific requirements.

Analysis of the mobile application process

The analysis of BuitenBeter's current process showed that the acquisition of information, the real time experiences in the form of a digital report happens from an external source namely a mobile application used by citizens in the public space. The report consists of information about the Location, Date, Citizen information, Municipality, specific application related information and an image. The report information and the image are sent separately. The reports undergo a refining process of indexing (where they get a unique report ID) and labeling (images are renamed and given the same report ID). The reports are stored in databases on servers at YCT. The images are stored separately in a file-system. The repositories are only accessible internally in YCT. All citizens receive an e-mail of the report. Municipalities that are using YCT's mobile application receive this report directly into their current system of handling reports. YCT has a direct connection to the municipality's system. This direct connection can give feedback to the citizens about the status of the report that they made. The municipalities that are not using YCT's mobile application receive an e-mail of the report and have to add it manually into their own system.

Knowledge base infrastructure

The current process was expanded with the image tagging and data mining process to create a new knowledge base infrastructure that will automate and manage the process of adding descriptive tags to images and using data mining to find patterns that can improve government policies, see Figure 9. It shows the supportive role of the Knowledge base infrastructure to the process of data gathering of citizen BuitenBeter reports (including images) to knowledge discovery, which results in patterns for governmental usage. The process to come from sheer data to knowledge is following the Crips-DM principles of understanding, preparation, mining, evaluation, process steps that are supported by the infrastructure. The current process of gathering information from the citizens was not changed. The acquisition and refining processes were added in the new knowledge base infrastructure and the image tagging and data mining process were integrated to work on top of this current process. The first step in the image tagging process is adding descriptive tags from information of the reports. This is done by using an application that can read and write tags to images. This step is done to enable fast searching and categorizing of the images for internal use by YCT using image organizing software. The images will be stored in the same file-system that the current process uses for storing images. YCT can choose to keep backups of the original images that do not contain any descriptive tags depending on storage limitations/requirements. Tags from the images are stored in a separate database that only contains a reference to the report ID and the corresponding tags. The data mining process can use this database to access and store tags used for pattern finding. The second step in the image tagging process, which is called automatic image tagging will add more

contextual descriptive tags to the images. These tags will be used for finding patterns that are not apparent by only looking at the report information. YCT will have to use third party services to implement the automatic image tagging. This requires sending the images to online services where descriptive tags will be added to the images. The tagged images are send back and are stored in the same file-system.

Finding patterns requires the use of different data mining techniques. The CRISP-DM model offers a model for implementing a data mining process. For this experiment a sub-set of this model has been used, but the knowledge base infrastructure should use all steps of this model to create a reliable, reusable and implementable data mining process that can be used to find patterns with. The data mining process is a iterative cycle that will prepare samples taken from the tags database and use these samples to look for patterns. The patterns found are stored in a knowledge repository. Data mining can be exploratory or explanatory. By storing meaningful and frequent patterns the process can look for specific patterns. This is more explanatory and will not focus on new knowledge but more on confirming current knowledge. Exploratory data mining will focus on finding new patterns and does not rely on previous knowledge. The results may or may not be of any use to the municipality. Depending on the requirements of the municipality the process can switch between these goals.

The selection phase is where the knowledge and requirements of the municipalities are combined and the patterns found from the data mining process are evaluated and new knowledge can be extrapolated.

Using knowledge and requirements from the municipalities new patterns can be formulated and are added to the repository and old patterns that are no longer required can be removed. Government policies that are formulated in the selection phase are evaluated by municipalities in the feedback phase. This feedback can result in new government policies and the improvement of current policies.




Figure 9: Knowledge base infrastructure

5.4.3 Conclusions

The mobile application of BuitenBeter is used by municipalities in the Netherlands as a mobile platform for their citizens. Citizens are using this mobile platform to send real time experiences from the public space to their municipalities. These real time experiences are sent in the form of reports to the municipality containing information about the problem and an image. Images sent with these reports can contain more information about the context of the reported problem. This information in the image is unstructured and to be useful must be structured. By adding descriptive tags of the contents of the image, structured information is derived from the image.

Descriptive tags added by image tagging can contain more information about a real time experience. The descriptive tags add more information about the context of the problem than the reported problem type. A sample of 294 reports was used to find patterns from these descriptive tags in relation to the main problem types. The data mining process was based on a modified CRISP-DM model. A JAVA program was written to find two and three pairs in the sample data set. The number of possible combination between the descriptive tags was too much to find real significant patterns. The patterns found did confirm the real time experiences from the citizens. For the problem category *Tree nuisances* the most frequent patterns were with *Tree*, *Tree branches* and *Bush* confirming the environment related to that problem. The problem category *Other* which had the most report's send had patterns with the tags *Tree*, *Car*, *House*, *Street*, *Pavement* and *Post* indicate that there is much overlapping with the other main problem categories in where these problems happen, but more descriptive tags are needed to really identify the problem and create new possible categories. The combination of image tagging and data mining revealed some interesting patterns but the sample used was too small to find real meaningful patterns. Combining the processes of image tagging and data mining with YCT's current process for the mobile application resulted in a new knowledge base infrastructure. This knowledge base infrastructure is an automated process that manages the discovery of implicit, but potentially useful information by gathering real time experiences of the public space from citizens using mobile applications.

The knowledge base infrastructure described is based on Michael's knowledge management architecture. The knowledge base infrastructure contributes to the practices of implementing this knowledge management architecture and further explores new architectures that contribute to knowledge management. The research done can be used to introduce new m-government programs aimed at better understanding and reacting to problems in the public space. As the example (construction of a knowledge-base for the South African Department of Arts, Culture, Science and Technology) stated in the literature review, this research contributes to the fact that Government using Knowledge discovery can use this to further improve Government policies.

Due to time constraints and the complexity of setting up a reliable automatic image tagging solution, the image tagging was done manually with a small sample of 294 reports to find patterns as a proto-type experiment. The images were opened one by one and descriptive tags were added. There was no real prior knowledge on what tags were best to use. Some rules were implemented to reduce the risks of inconsistent and redundant tags. To further improve the findings, the process should expand the number of descriptive tags used. The automatic image tagging process should be implemented with custom requirements to extrapolate domain specific descriptive tags from the images. Some images were close ups of certain problems and did not contain a lot of content (e.g. a close up of a broken pavement tile only shows the tile and no extra contextual information about the surroundings). For these images it was hard to add descriptive tags, because of the lack of content. To require the citizen to take a specific format of images is not feasible, this means more research has to be done on how to better deal with these type of close up images and the (automatic) recognition of their content.

The knowledge base infrastructure presented in this research adds new value for this mobile application. By implementing this knowledge base infrastructure new possibilities arise for

the mobile application in the light of policy modeling. This would include possibilities for an integrative system where more insights are received from the process of finding patterns in the reports send.

Future research should be done on fully implementing the knowledge base infrastructure. The processes of image tagging and data mining were only were done manually and using a small sample. The process of implementing a fully automated process that can be integrated into a current process needs further in depth research on automatic image tagging and creating a CRISP-DM based data mining process.

6 Summary

This deliverable reported on strategies regarding mobile sensing for eParticipation scenarios. We discussed the current status of incorporating mobile technology into the process of decision-making and described the aims of the Live+Gov project to further improve this status. In particular, we explored reality mining for the purpose of obtaining improved context information on the citizen and, thus, to use this information to provide a better eParticipation experience. We surveyed existing approaches to reality sensing and provided an overview on algorithms for activity recognition, image analysis, topic detection, and sentiment analysis.

We also surveyed existing Open Governmental Data repositories and discuss how they can be used for enriching eParticipation approaches. We presented LISA, a working showcase that illustrates the use of Open Governmental Data by offering a visually appealing way to explore this data for ordinary citizens.

In the context of our three use cases, we investigated those approaches to reality sensing and external data sources, and gave mappings how those can be applied. We also reported on first experiences with data mining approaches in the context of the Urban Maintenance use case.

7 References

- [1] Görlitz, Olaf; Gröner, Gerd; Hachenberg, Christian; Scheglmann, Stefan; Thimm, Matthias (2012): The Open Data Workflow: Towards an Open Data Middleware. In: W3C Workshop USING OPEN DATA: policy modeling, citizen empowerment, data journalism.
- [2] Y.-J. Hong, I.-J. Kim, S. C. Ahn and H.-G. Kim, „Mobile health monitoring system based on activity recognition using accelerometer,“ *Simulation Modelling Practice and Theory* 18 (4), pp. 446-455, 2010.
- [3] E. Miluzzo, N. D. Lane, K. Fodor, R. Peterson, H. Lu, M. Musolesi, S. B. Eisenman, X. Zheng and A. T. Campbell, „Sensing meets mobile social networks: the design, implementation and evaluation of the CenceMe application,“ *SenSys '08 Proceedings of the 6th ACM conference on Embedded network sensor systems* , pp. 337-350, 2008.
- [4] A. Thiagarajan, J. Biagioni, T. Gerlich and J. Eriksson, „Cooperative transit tracking using smart-phones,“ *SenSys '10 Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, pp. 85-98, 2010.
- [5] Z. Prekopcsák, S. Soha, T. Henk and C. Gáspár-Papanek, „Activity Recognition for Personal Time Management,“ *Aml '09 Proceedings of the European Conference on Ambient Intelligence*, pp. 55-59, 2009.
- [6] Y. Maruno, K. Cho, Y. Okamoto, H. Setoguchi and K. Ikeda, „An Online Human Activity Recognizer for Mobile Phones with Accelerometer,“ *ICONIP'11 Proceedings of the 18th international conference on Neural Information Processing - Volume Part II*, pp. 358-365, 2011.
- [7] T. S. Saponas, J. Lester, J. Froehlich, J. Fogarty and J. Landay, „iLearn on the iPhone: Real-Time Human Activity Classification on Commodity Mobile Phones,“ 2008.
- [8] A. M. Khan, Y.-K. Lee, S. Y. Lee and T.-S. Kim, „Human Activity Recognition via An Accelerometer-Enabled-Smartphone Using Kernel Discriminant Analysis,“ *The 5th International Conference on Future Information Technology (FutureTech2010)*, pp. 1-6, 2010.
- [9] M. Berchtold, M. Budde, D. Gordon, H. R. Schmidtke and M. Beigl, „ActiServ: Activity Recognition Service for mobile phones,“ *2010 International Symposium on Wearable Computers (ISWC)*, pp. 1-8, 2010.
- [10] Y. Wang, J. Lin, M. Annavaram, Q. A. Jacobson, J. Hong and B. Krishnamachari, „A framework of energy efficient mobile sensing for automatic user state recognition,“ *MobiSys '09 Proceedings of the 7th international conference on Mobile systems, applications, and services*, pp. 179-192, 2009.
- [11] A. Henpraserttae, S. Thiemjarus and M. Marukatat, „Accurate Activity Recognition using a Mobile Phone regardless of Device Orientation and Location,“ *2011 International Conference on Body Sensor Networks (BSN)*, pp. 41 - 46, 2011.
- [12] L. Sun, D. Zhang, B. Li, B. Guo and S. Li, „Activity recognition on an accelerometer embedded mobile phone with varying positions and orientations,“ *Proceedings of*

- the 7th international conference on Ubiquitous intelligence and computing, pp. 548-562, 2010.
- [13] T. Brezmes, J.-L. Gorricho and J. Cotrina, „Activity Recognition from Accelerometer Data on a Mobile Phone,” IWANN '09 Proceedings of the 10th International Work-Conference on Artificial Neural Networks: Part II: Distributed Computing, Artificial Intelligence, Bioinformatics, Soft Computing, and Ambient Assisted Living, pp. 796-799, 2009.
- [14] J. Yang, H. Lu, Z. Liu and P. P. Boda, „Physical Activity Recognition with Mobile Phones: Challenges, Methods, and Applications,” Multimedia Interaction and Intelligent User Interfaces, pp. 185-213, 2010.
- [15] L. Grokop, A. Sarah, C. Brunner, V. Narayanan and S. Nanda, „Activity and device position recognition in mobile devices,” UbiComp '11 Proceedings of the 13th international conference on Ubiquitous computing, pp. 591-592, 2011.
- [16] A. M. Khan, „Human Activity Recognition Using A Single Tri-axial Accelerometer,” Kyung Hee University, Seoul, 2011.
- [17] S. Reddy, J. Burke, D. Estrin, M. Hansen and M. Srivastava, „Determining transportation mode on mobile phones,” ISWC '08 Proceedings of the 2008 12th IEEE International Symposium on Wearable Computers, Nr. 25-28, 2008.
- [18] J. R. Kwapisz, G. M. Weiss and S. A. Moore, „Activity Recognition using Cell Phone Accelerometer,” SIGKDD Explorations, pp. 74-82, 2010.
- [19] D. Sugimori, T. Iwamoto and M. Matsumoto, „A Study About Identification of Pedestrian by Using 3-axis Accelerometer,” 17th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications, pp. 134-137, 2011.
- [20] N. Ravi, N. Dandekar, P. Mysore and M. L. Littman, „Activity Recognition from Accelerometer Data,” IAAI'05 Proceedings of the 17th conference on Innovative applications of artificial intelligence - Volume 3, Nr. 3, pp. 1541-1546, 2005.
- [21] J. Frank, S. Mannor and D. Precup, „Activity Recognition with Mobile Phones,” ECML PKDD'11 Proceedings of the 2011 European conference on Machine learning and knowledge discovery in databases - Volume Part III, pp. 630-633, 2011.
- [22] Z. Zhao, Y. Chen, J. Liu, Z. Shen and M. Liu, „Cross-People Mobile-Phone Based Activity Recognition,” IJCAI'11 Proceedings of the Twenty-Second international joint conference on Artificial Intelligence - Volume III, pp. 2545-2550, 2011.
- [23] Y.-C. Yang, T. Toida and C.-M. Hong, „Transportation prediction using build-in triaxial accelerometer in cell phone,” 2010.
- [24] Lara, O.; Labrador, M., "A Survey on Human Activity Recognition using Wearable Sensors," Communications Surveys & Tutorials, IEEE , vol.PP, no.99, pp.1,18, 0
- [25] doi: 10.1109/SURV.2012.110112.00192
- [26] J. Yin, Q. Yang, and J. Pan, “Sensor-based abnormal human-activity detection,” IEEE Transactions on Knowledge and Data Engineering, vol. 20, no. 8, pp. 1082–1090, 2008.

-
- [27] O. X. Schmilch, B. Witzschel, M. Cantor, E. Kahl, R. Mehmke, and C. Runge, "Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring," *Computers in Human Behavior*, vol. 15, no. 5, pp. 571–583, 1999
 - [28] F. Foerster, M. Smeja, and J. Fahrenberg, "Detection of posture and motion by accelerometry: a validation study in ambulatory monitoring," *Computers in Human Behavior*, vol. 15, no. 5, pp. 571–583, 1999.
 - [29] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher, "Activity recognition and monitoring using multiple sensors on different body positions," in *International Workshop on Wearable and Implantable Body Sensor Networks*, (Washington, DC, USA), IEEE Computer Society, 2006
 - [30] [0.26] L. Bao and S. S. Intille, "Activity recognition from user-annotated acceleration data," in *Pervasive*, pp. 1–17, 2004
 - [31] Y.-P. Chen, J.-Y. Yang, S.-N. Liou, Gwo-Yun=Lee, and J.-S. Wang, "Online classifier construction algorithm for human activity detection using a tri-axial accelerometer," *Applied Mathematics and Computation*, vol. 205, no. 2, pp. 849–860, 2008.
 - [32] K. Altun and B. Barshan, "Human activity recognition using inertial/magnetic sensor units," in *Human Behavior Understanding*, *Lecture Notes in Computer Science*, pp. 38–51, Springer Berlin / Heidelberg, 2010
 - [33] D. T. G. Huynh, "Human activity recognition with wearable sensors", Dissertation at TU Darmstadt, 2008.
 - [34] R. Hamid, S. Maddi, A. Johnson, A. Bobick, and C. Isbell I. Essa. Unsupervised discovery and characterization of activities from event-streams. In *UAI*, 2005.
 - [35] B. Clarkson and A. Pentland. Unsupervised clustering of
 - [36] Ambulatory audio and video. In *ICASSP*, 1999.
 - [37] E. Kim, S. Helal, D. Cook, "Human Activity Recognition and Pattern Discovery", *IEEE Pervasive Comput.* 2010; 9(1): 48. doi:10.1104/MPRV.2010.7
 - [38] T. Huynh, M. Fritz, and B. Schiele. "Discovery of Activity Patterns using Topic Models," *Proceedings of the Tenth International Conference on Ubiquitous Computing*, pp. 10-19, 2008.
 - [39] U. Maurer, A. Smailagic, D. P. Siewiorek, and M. Deisher, "Activity recognition and monitoring using multiple sensors on different body positions," in *International Workshop on Wearable and Implantable Body Sensor Networks*, (Washington, DC, USA), IEEE Computer Society, 2006.
 - [40] Android Activity Recognition API. Accessed at 2013-06-08. <http://developer.android.com/training/location/activity-recognition.html>
 - [41] Google IO 2013. Keynote wrapup. Accessed at 2013-06-08. <http://android-developers.blogspot.de/2013/05/android-at-google-io-2013-keynote-wrapup.html>
 - [42] Nike + iPod. Accessed 2013-06-08. <http://www.apple.com/ipod/nike/>
 - [43] Google Glas. Accessed 2013-06-08. <http://www.google.de/glass/start/>
-

-
- [44] Nike Sports Watch. Accessed 2013-06-08.
http://nikeplus.nike.com/plus/products/sport_watch/
 - [45] A. Thiagarajan, J. Biagioni, T. Gerlich, and J. Eriksson, "Cooperative transit tracking using smart-phones," in Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 85–98.
 - [46] Gabriel Takacs, Vijay Chandrasekhar, Natasha Gelfand, Yingen Xiong, Wei-Chao Chen, Thanos Bimpigiannis, Radek Grzeszczuk, Kari Pulli, and Bernd Girod. 2008. Outdoors augmented reality on mobile phone using loxel-based visual feature organization. In Proceeding of the 1st ACM international conference on Multimedia information retrieval (MIR '08). ACM, New York, NY, USA, 427-434.
 - [47] Xu Liu, Jonathan Hull, Jamey Graham, Jorge Moraleda, Timothee Bailloeul. 2010. Mobile Visual Search, Linking Printed Documents to Digital Media. In CVPR 2010 Demonstrations.
 - [48] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," Proc. of the IEEE Conference on Computer Vision and Pattern Recognition, 2007.
 - [49] D. Wagner, G. Reitmayr, A. Mulloni, T. Drummond, and D. Schmalstieg, "Pose tracking from natural features on mobile phones," Proc. of the 7th IEEE/ACM Int. Symp. on Mixed and Augmented Reality (Sept. 15 - 18, 2008), pp. 125-134.
 - [50] J. Li and J. Z. Wang, "Real-time Computerized Annotation of Pictures", Proceedings of the ACM Multimedia Conference, pp. 911-920, ACM, Santa Barbara, CA, 2006.
 - [51] J. Philbin, O. Chum, M. Isard, J. Sivic, and A. Zisserman, "Object retrieval with large vocabularies and fast spatial matching," in Proc. IEEE Conf. Computer Vision and Pattern Recognition, 2007.
 - [52] S. Papadopoulos, C. Zigkolis, Y. Kompatsiaris, A. Vakali. "Cluster-based Landmark and Event Detection on Tagged Photo Collections". In IEEE Multimedia Magazine 18(1), pp. 52-63, 2011
 - [53] S. Nikolopoulos, S. Zafeiriou, I. Patras and I. Kompatsiaris, "High-order pLSA for indexing tagged images", Signal Processing Elsevier, Special Issue on Indexing of Large-Scale Multimedia Signals, Volume 93, Issue 8, August 2013
 - [54] J. Sivic and A. Zisserman. 2003. Video Google: A Text Retrieval Approach to Object Matching in Videos. In Proceedings of the Ninth IEEE International Conference on Computer Vision - Volume 2 (ICCV '03), Vol. 2. IEEE Computer Society, Washington, DC, USA,
 - [55] D3.1, "Platform and prototype application for augmented reality", Live+Gov consortium, June 2013.
 - [56] C. Cortes and V. Vapnik, "Support-vector networks in Machine Learning", pp. 273–297, 1995.
 - [57] C. Harris and M. Stephens, "A combined corner and edge detector," in Proc. 4th Alvey Vision Conf., 1988.
-


-
- [58] J. Matas, O. Chum, M. Urban, and T. Pajdla, “Robust wide baseline stereo from maximally stable extremal regions,” in Proc. British Machine Vision Conf. (BMVC), Cardiff, Wales, Sept. 2002.
 - [59] E. Rosten and T. Drummond, “Machine learning for high speed corner detection,” in Proc. European Conf. Computer Vision (ECCV), Graz, Austria, May 2006.
 - [60] J. Shi and C. Tomasi, “Good features to track”, In IEEE Conference on Computer Vision and Pattern Recognition, pages 593400, June 1994.
 - [61] D. Lowe, “Object recognition from local scale-invariant features,” in Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR), Los Alamitos, CA, Aug. 1999.
 - [62] M. Agrawal, K. Konolige, and M. R. Blas. CenSurE: Center surround extremas for realtime feature detection and matching. In Proceedings of the European Conference on Computer Vision (ECCV), 2008.
 - [63] H. Bay, A. Ess, T. Tuytelaars, and L. V. Gool, “Speeded-up robust feature,” Comput. Vis. Image Understand., vol. 110, no. 3, pp. 346–359, 2008.
 - [64] E. Rublee, V. Rabaud, K. Konolige, and G. Bradski. ORB: An efficient alternative to SIFT or SURF. In International Conference on Computer Vision, Barcelona, 2011.
 - [65] M. Calonder, V. Lepetit, C. Strecha, and P. Fua. 2010. BRIEF: binary robust independent elementary features. In Proceedings of the 11th European conference on Computer vision: Part IV (ECCV'10), K. Daniilidis, P. Maragos, and N. Paragios (Eds.). Springer-Verlag, Berlin, Heidelberg, 778-792, 2010.
 - [66] D. Lowe, “Distinctive image features from scale-invariant keypoints,” Int. J. Comput. Vis., vol. 60, no. 2, pp. 91–110, 2004.
 - [67] J. C. van Gemert, J. M. Geusebroek, C. J. Veenman, and A. W. M. Smeulders. Kernel codebooks for scene categorization. In Proc. ECCV, 2008
 - [68] J. C. van Gemert, C. J. Veenman, A. W. M. Smeulders, J. M. Geuse-broek, Visual word ambiguity, IEEE Transactions on Pattern Analysis and Machine Intelligence 32 (7) (2010) 1271–1283.
 - [69] H. Jegou, M. Douze, C. Schmid, and P. Pérez, “Aggregating local descriptors into a compact image representation,” in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, 2010
 - [70] Pang, B., Lee, L., (2008) Opinion Mining and Sentiment Analysis. In Foundations and Trends in Information Retrieval, 2 (1-2) Hanover, MA, USA: Now Publishers Inc. pp. 1-135
 - [71] Neuendorf, K. A. (2002), The Content Analysis Guidebook, Cleveland State University, USA, 2002
 - [72] Pang, B., Lee, L., and Vaithyanathan, S., (2002) Thumbs up? Sentiment classification using machine learning techniques, Proceedings of EMNLP, pp. 79-86, 2002
 - [73] Abbasi, A., Chen, H., and Salem, A. (2008), Sentiment analysis in multiple languages: Feature selection for opinion classification in Web forums. ACM Trans. Inf. Syst. 26, 3, Article 12 (June 2008)
-

-
- [74] Palau. R. M. and Moens, M. F. (2009), Argumentation mining: the detection, classification and structure of arguments in text, pp. 98-107 in Proceedings of the 12th International Conference on Artificial Intelligence and Law.
 - [75] Vinodhin, G., Chandrasekar, RM., (2012) Sentiment Analysis and Opinion Mining: A Survey, International Journal of Advanced Research in Computer Science and Software Engineering, Volume 2, Issue 6, June, 2012
 - [76] Quinlan, J. R. C4.5: Programs for Machine Learning. Morgan Kaufmann Publishers, 1993
 - [77] Domingos, P., Pazzani, M.J.: On the optimality of the simple bayesian classifier under zero-one loss. Machine Learning 29(2-3), 103–130 (1997)
 - [78] Xuming He, Richard S. Zemel, and Miguel Á. Carreira-Perpiñán. 2004. Multiscale conditional random fields for image labeling. In Proceedings of the 2004 IEEE computer society conference on Computer vision and pattern recognition (CVPR'04). IEEE Computer Society, Washington, DC, USA, 695-703.
 - [79] Breiman, L., Friedman, J., Olshen, R., Stone, C.: Classification and Regression Trees. Wadsworth and Brooks, Monterey, CA (1984)
 - [80] Leo Breiman. 2001. Random Forests. Mach. Learn. 45, 1 (October 2001), 5-32. DOI=10.1023/A:1010933404324 <http://dx.doi.org/10.1023/A:1010933404324>
 - [81] Egmont-Petersen, M., de Ridder, D., Handels, H.: Image processing with neural networks – a review. Pattern Recognition 35(10), 2279–2301 (2002), doi:10.1016/S0031-3203(01)00178-9


Hosmer, David W., and Stanley Lemeshow. Applied logistic regression. New York: Wiley, 2000
 - [82] Anderson, J. Q., & Rainie, L. (2012). *What is the potential future influence of Big Data by 2020?* Pew Research Center and the Imagining the Internet Center at Elon University.
 - [83] Blei, D. M., Ng, A., & Jordan, M. I. (2003). Latent dirichlet allocation. *The Journal of Machine Learning Research*, 993-1022 .
 - [84] Griffiths, T. L., & Steyvers, M. (2004). Finding Scientific Topics. *Colloquium on Mapping Knowledge Domains* (pp. 5228--5235). Natl Acad Sci & Engrn, Arnold & Mable Beckman Ctr, Irvine, CA: NATIONAL ACADEMY OF SCIENCES .
 - [85] Hofmann, T. (1999). Probabilistic latent semantic indexing. *ACM SIGIR conference on Research and development in information retrieval* (pp. 50-57). Berkeley, California, United States: ACM.
 - [86] Hofmann, T. (2001). Unsupervised Learning by Probabilistic Latent Semantic Analysis. *Machine Learning*, 177-196.
 - [87] Minka, T., & Lafferty, J. (2002). Expectation-Propagation for the Generative Aspect Model. *Conference on Uncertainty in Artificial Intelligence* (pp. 352-359). Morgan Kaufmann.
-

-
- [88] Nigam, K., McCallum, A. K., Thrun, S., & Mitchell, T. (2000). Text Classification from Labeled and Unlabeled Documents using EM. *Machine Learning*, 103-134. Alkhamayseh, S., Lawrence, E., & Zmijewska, A. (2007). *Towards Understanding Success Factors in Interactive Mobile Government*. Sydney: University of Technology.
- [89] Carroll, J. (2006). „What’s in It for Me?":Taking M-Government to the People. *eValues*, (pp. 1-13). Bled, Slovenia.
- [90] Kushchu, I., & Halid Kusc, M. (2003). From E-government to M-government: Facing the Inevitable. *3rd European Conference on e-Government*, (pp. 1-13).
- [91] Dzhusupova, Z., Ojo, A., & Janowski, T. (2009). *ORGANIZING AND MANAGING KNOWLEDGE FOR E-GOVERNMENT: Issues, Practices and Challenges*. Macao: Centre for Electronic Governance, United Nations University.
- [92] Michael, Z. H. (1999). Managing Codified Knowledge. *Sloan Management Review* , 45-58.
- [93] Shaw, M. J., Subramaniam, C., Tan, G. W., & Welge, M. E. (2001). Knowledge management and data mining for marketing. *Decision Support Systems* (pp. 127–137). Elsevier.
- [94] Silberschatz, A., & Tuzhilin, A. (1996). What Makes Patterns Interesting in Knowledge Discovery Systems. *Knowledge and Data Engineering, IEEE Transactions on Knowledge and Data Engineering* , 970 - 974.
- [95] IBM. (2010). *CRISP-DM 1.0 Step-by-step data mining guide*. New York, United States of America: IBM.
- [96] Renn, M., van Beusekom, J., Keyzers, D., & Breuel, T. M. (2010). Automatic Image Tagging using Community-Driven Online Image Databases. *Lecture Notes in Computer Science* (pp. 112-126). SpringerLink.
- [97] Sharon S. Dawes. The evolution and continuing challenges of e-governance. *Public Administration Review*, 2008.
- [98] Steffen Albrecht, Niels Köhlrausch, Herbert Kubicek, Barbara Lippa, Oliver Maerker, Matthias Trnel, Volker Vorwerk, Hilmar Westholm, and Christian Wiedwald. eParticipation - Electronic Participation of Citizens and the Business Community in eGovernment. *Study on Behalf of the Federal Ministry of the Interior, Division IT 1*, 2008.
- [99] Rakesh Agrawal and Ramakrishnan Srikant Fast algorithms for mining association rules in large databases. Proceedings of the 20th International Conference on Very Large Data Bases, VLDB, pages 487-499, Santiago, Chile, September 1994.
-

A Analyses Data Set



Report location spread in relation to municipality center



Main problem category frequency per district

B Patterns

Main Category	Tree	Car	Home	Street	Grass	Street pavement	Bicycle path	Building	Waste	Tree branches	Bush	Street pole	Pad	Sidewalks	Dangerous	Sand	Broken	Lamppost	Nuisance
1	22	14	11	13	16	14	6	11	32	1	7	3	9	2	2	8	3	0	6
3	5	7	8	8	4	4	4	0	3	0	2	4	7	23	5	4	7	0	0
4	7	1	1	8	6	5	12	0	0	1	2	4	1	4	6	1	5	0	0
5	5	4	7	2	1	2	1	3	0	0	3	1	3	0	0	1	0	1	0
6	3	5	3	3	1	2	0	5	0	1	5	1	1	0	0	0	2	9	0
7	2	0	0	0	1	1	0	0	0	1	0	0	0	0	0	0	0	0	0
9	0	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0	0	0	0
10	10	1	0	1	3	1	2	0	2	4	6	0	0	0	0	3	0	0	2
12	47	17	15	7	15	6	14	11	0	30	12	0	0	0	9	0	0	0	8
13	0	1	0	0	0	0	0	1	0	0	1	0	0	0	0	0	0	1	0
14	43	27	23	23	17	20	14	19	7	7	7	19	13	5	9	10	7	7	6

Table 1. Top 20 tag frequency pair with main problem categories. (The sample did not have all main problem categories).

SC1	SC2	FREQUENCY
Tree	Branches	27
Sidewalk	Loose	12
Tree	Car	11
Tree	Grass	11
Tree	Procession caterpillars	11
Grass	Waste	10
Tree	Bush	9
Tree	Bicycle path	9
Tree	Dangerous	9
House	Tree	9

Table 2. Top 10, two pair patterns found using the 294 samples

SC1	SC2	SC3	FREQUENCY
Tree	Bush	Branches	5
Tree	Branches	Bicycle path	5
Tree	Branches	Knocked down	5
Tree	Branches	Dangerous	5
Tree	Car	Branches	5
Tree	Branches	Lie down	5
Sidewalk	Loose	Dangerous	5
Tree	Bush	Procession caterpillars	5
Tree	Branches	Nuisance	4
Tree	Grass	Branches	4

Table 3. Top 10, three pair patterns found using the 294 samples

All tags (294 sample)	FREQUENCY
Tree	144
Car	78
House	68
Street	65
Grass	64
Street pavement	55
Bicycle path	53
Building	51
Waste	50
Branches	45

Table 4. Top 10, tags found using the 294 samples