UvA-DARE (Digital Academic Repository)

Multimodal Classification of Urban Micro-Events

Sukel, M.; Rudinac, S.; Worring, M.

DOI

[10.1145/3343031.3350967](https://doi.org/10.1145/3343031.3350967)

Publication date

2019

Document Version

Final published version

Published in

MM'19

License

Article 25fa Dutch Copyright Act

[Link to publication](#)

Citation for published version (APA):

Sukel, M., Rudinac, S., & Worring, M. (2019). Multimodal Classification of Urban Micro-Events. In *MM'19: proceedings of the 27th ACM Conference on Multimedia : October 21-25, 2019, Nice, France* (pp. 1455-1463). Association for Computing Machinery.
<https://doi.org/10.1145/3343031.3350967>

General rights

It is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), other than for strictly personal, individual use, unless the work is under an open content license (like Creative Commons).

Disclaimer/Complaints regulations

If you believe that digital publication of certain material infringes any of your rights or (privacy) interests, please let the Library know, stating your reasons. In case of a legitimate complaint, the Library will make the material inaccessible and/or remove it from the website. Please Ask the Library: <https://uba.uva.nl/en/contact>, or a letter to: Library of the University of Amsterdam, Secretariat, Singel 425, 1012 WP Amsterdam, The Netherlands. You will be contacted as soon as possible.

Multimodal Classification of Urban Micro-Events

Maarten Sukel
University of Amsterdam
Amsterdam, The Netherlands
m.m.sukel@uva.nl

Stevan Rudinac
University of Amsterdam
Amsterdam, The Netherlands
s.rudinac@uva.nl

Marcel Worring
University of Amsterdam
Amsterdam, The Netherlands
m.worring@uva.nl

ABSTRACT

In this paper we seek methods to effectively detect urban micro-events. Urban micro-events are events which occur in cities, have limited geographical coverage and typically affect only a small group of citizens. Because of their scale these events are difficult to identify in most data sources. However, by using citizen sensing to gather data, detecting them becomes feasible. The data gathered by citizen sensing is often multimodal and, as a consequence, the information required to detect urban micro-events is distributed over multiple modalities. This makes it essential to have a classifier capable of combining them. In this paper we explore several methods of creating such a classifier, including early, late and hybrid fusion as well as representation learning using multimodal graphs. We evaluate performance in terms of accurate classification of urban micro-events on a real world dataset obtained from a live citizen reporting system. We show that a multimodal approach yields higher performance than unimodal alternatives. Furthermore, we demonstrate that our hybrid combination of early and late fusion with multimodal embeddings outperforms our other fusion methods.

CCS CONCEPTS

• **Computing methodologies** → *Supervised learning; Machine learning algorithms; Boosting; Feature selection.*

KEYWORDS

urban micro-events, citizen as a sensor, multimodal classification, event detection, representation learning

ACM Reference Format:

Maarten Sukel, Stevan Rudinac, and Marcel Worring. 2019. Multimodal Classification of Urban Micro-Events. In *Proceedings of the 27th ACM Int'l Conf. on Multimedia (MM'19)*, Oct. 21–25, 2019, Nice, France. , 9 pages. <https://doi.org/10.1145/3343031.3350967>

1 INTRODUCTION

Cities are living organisms where numerous events take place at different geographical and temporal scales. Some of these events are macroscopic, involving a large geographical area and a large number of people. Other events occur in a very limited geographical area and have a smaller number of people involved. In this paper we focus on such small scale events, referring to them as *urban*

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

MM '19, October 21–25, 2019, Nice, France

© 2019 Association for Computing Machinery.

ACM ISBN 978-1-4503-6889-6/19/10...\$15.00

<https://doi.org/10.1145/3343031.3350967>

Figure 1: Example of how an urban micro-event has to be detected between multiple modalities.

micro-events. In particular we focus on the subclass of urban micro-events that need attention from the municipality. Examples of such urban micro-events are speeding boats, graffiti on walls, trash on the street, broken streetlights or a bicycle wreckage that needs to be removed (cf. Figure 1). In this paper we focus on automatically detecting those kind of urban micro-events.

Different types of events have been studied in various areas of multimedia. For example, events have been intensively studied in video analysis [1], but such events normally have a clear and consistent visual pattern and occur over multiple frames. The task we try to perform is different in nature because the text or image might not directly mention or show the related issue. In recent years it has been researched how to detect natural disaster related events [2] by combining social multimedia with satellite imagery and furthermore research has been done detecting if tweets related to real-world events are fake or real [4]. These tasks combine multimedia data to perform classification, which makes them relevant to our task at hand. However, the target events are very different from the urban micro-events we attempt to classify. For detecting urban micro-events the above methods give inspiration, but are not directly applicable.

Urban micro-events are difficult to find due to their small scale and the large variety of forms they may take. However, citizens are motivated to report possible issues, since they are likely to cause issues resulting in a negative impact on livability and thus citizens in the neighborhood. Using citizens to collect data about their surroundings is often called the *citizen as a sensor* paradigm [5] [36], and alternatively, participatory sensing or human-in-the-loop sensing. Such a paradigm can also be deployed to collect, analyze and mine information about events. Pervasiveness of smart phones with inbuilt high-quality sensors makes such collected information increasingly valuable.

Citizen as a sensor creates valuable data sources for the detection of urban micro-events. An application of the citizen as a sensor paradigm frequently seen in cities around the world is a service

Figure 2: To detect urban micro-events in data gathered by using citizen sensing (1), features are extracted and additional contextual data is added from several additional data sources (2), unimodal and multimodal classifiers are created (3) and performance is evaluated on a real world dataset collected from a live citizen report system (4).

where residents can report issues in the public space, which creates data containing rich information about urban micro-events. For example, in most cities in the United States it is possible to call 311 for non-emergency service requests, report issues through web forms or contact the local government through social media channels such as Twitter or WhatsApp. With such a service cities are trying to get closer to the citizens by increasingly responding to urban service requests, making it important to process the request properly which in turn requires it to be given the right issue category in order to offer a timely and appropriate solution. These request often report and describe an urban micro-event. Since a large number of cities have such a citizen report system that could benefit from the accurate classification of urban micro-events, an effective solution to the problem could have a large potential for improving city livability.

The detection of urban micro-events in service requests is a challenging task due to the heterogeneous nature of these events, which range from anything related to nuisance in the public area caused by begging or boats with loud music to potholes and dangerous traffic situations (cf. Figure 2). Because the reporting citizen is not always familiar with the exact meaning of the class the wrong class could be selected which creates several problems. One of the main problems is that choosing the wrong class results in the issue not being sent to the right department, which in turn results in the issue not being solved or being solved with a delay. One of the solutions for this is letting experts decide the correct class. However this is a labor intensive exercise which is also likely to cause delays. The use of an automated classifier that performs better than both citizen and the expert would allow for a faster detection of urban micro-events and their underlying issues, and consequently lead to a timely issue resolution and a reduced negative impact on the city livability.

Reports that are made by citizens often consist of text, image, spatial and temporal data. Therefore it is essential to be able to combine these modalities. An example is a combination of text that reports trash on the street, and the associated image that shows the

type of trash. Complete understanding of the underlying issue thus requires information extracted from both text and visual content. Another example is a noise complaint, where the time and location of the report could be relevant for how the issue should be classified. The classifier in this case has to be capable of effectively extracting relevant features out of a wide range of data types, including visual, textual, spatial and temporal data.

For creating a multi-modal classifier we consider several fusion schemes: Early, late and hybrid fusion. In early fusion the features are first extracted from each modality and then combined in a joint representation. The approaches to early fusion range from simple concatenation to complex graph embeddings. Late fusion, on the other hand, combines different modalities by first performing classification on the unimodal feature vectors, and then using such obtained classification results as an input into the classifier combination technique. Both classic [19] and more recent studies [26] show that the optimal choice of classifier combining technique depends on the application. For example, research in video retrieval shows that late fusion tends to give slightly better performance, but where early fusion performs better the difference is more significant [34]. Since we are combining four modalities a hybrid fusion technique combining early and late fusion could result in even better performance. To arrive at an optimal method, in this paper we utilize a wide range of modalities and modality fusion techniques for an accurate classification of urban micro-events. The contributions of this paper can be summarized as follows:

- We present several fusion methods, both early, late and hybrid, to create a multi-modal classifier capable of detecting urban micro-events based on textual, visual, spatial and temporal data.
- We evaluate the resulting multi-modal learning method on a real life system deployed in Amsterdam for detecting urban micro-events.
- We elaborate on the usefulness of textual, visual, spatial and temporal features for the classification of urban micro-events by doing an extensive evaluation.

2 RELATED WORK

In this section we discuss related work. First we start with approaches to multimodal classification and from there we explore the citizen as a sensor paradigm and customer feedback systems. Finally we discuss approaches for event detection.

2.1 Multimodal classification

Multimodal classification has been a long-standing research topic in the multimedia community and over time a number of excellent solutions have been proposed. There are several general approaches, the most common being early and late fusion. Below we survey trends in both “schools”. Recent research efforts revolve around creating a joint item representation from the features of different modalities. For example, information extracted from the visual content and the text could be combined by generating a new ‘imagined’ vector as in [9] or by using a common subspace as in [40]. This is, however, not true for service requests since the image labels are often more subjective. For example a picture taken of a noise disturbance can capture many different scenes each with their own visual appearance. For the same reasons, using the text to determine attention as in [21] is not likely to work on urban micro-events. However, as shown in [43], using visual content to resolve ambiguities can significantly improve upon text-based event extraction results, even when the modalities are not well aligned. Multimodal classification requires extracting features from the visual content. Features can be extracted from images by using e.g. convolutional neural networks. These features can be visual concepts like a building or a bicycle, but often also more abstract visual features like the vector of elements defining the layer before the softmax as in [9]. A group of deep networks, referred to as residual networks or ResNet [12], achieved particularly good performance in various task such as image recognition or object detection. Such extracted features generally yield better results than the traditional alternatives, such as Bag of visual words, SIFT or HOG [18, 20], which is why we consider them a promising starting point for building multimodal representation of citizen reports. Graphs have also been deployed for embedding different modalities beyond text and visual content, and then the affinities between multimedia items were computed using e.g. random walks with restarts as in [33]. With the advent of deep learning, the general idea was revived in the approaches such as DeepWalk [29]. More recently, Grover and Leskovec proposed a node2vec framework for learning a low-dimensional representations of the nodes in a graph by optimizing a neighborhood preserving objective, allowing representation of complex networks [11]. We conjecture that this approach may be effective in encoding complex relations between different modalities in service requests, ranging from text and visual content to geolocation, time and weather conditions.

2.2 Citizen as a sensor

The use of citizen as a sensor brings multiple challenges [5]. Typically, domain experts are tasked with determining the category of the issue described by the report, but their time is limited and costly. The range of reported issues is also broad, requiring knowledge of multiple domains and location specific information. Another challenge is that citizens have no way of discerning truth from

falsehood. The same applies for analyzing the reports, citizens that are interested in skewing information can create bias in the data. The use of machine learning for processing reports also has shortcomings since it requires training data and feedback. However, having humans perform the classification is resource intensive task and, in addition, they may become tired over time, possibly missing or misinterpreting signals.

Customer feedback data is a rich source for the creation of classifiers. For example textual classification of customer feedback chats is discussed in [28]. Being able to predict customer satisfaction based on textual data would certainly be a nice addition to our approach, but our primary goal is detection of micro events in an urban environment. In [22] methods to detect classes (i.e. comment, request, bug, complaint, meaningless, and undetermined) in customer service data are compared and the best scoring method is a bidirectional LSTM+CNN. Similarly, [23] explores understanding of the same classes in multilingual customer feedback. The data used has a heterogeneous nature, consisting of multiple languages. These categories are different from the urban micro-events we seek to classify, however the type of textual input data and output have a similar nature so the best scoring methods might also apply for the classification of urban micro events.

Several designs of citizen feedback systems are compared in [27] and several important factors of how such a system impacts the interaction between citizens and municipalities are discussed. The study identifies the way in which category selection works as one of the most important factors determining effectiveness of the system.

Finally, as stressed by Tang et al., utilizing the full potential of citizen reports requires effective methods for their analysis and categorization, based on the increasingly heterogeneous multimedia data they contain [37]. In this paper we present several solutions to the category selection problem.

2.3 Event detection

Event detection has been a long-standing interest of the research community. Examples are plentiful and range from classic works on detection of events in news articles using text retrieval and clustering techniques [42] to multimedia event detection in video [1]. Methods for finding events in news and social data streams are particularly well researched [6, 14, 30, 39, 41], however the scope, such as festivals and international incidents, is of a different geographical scale than the urban micro-events we attempt to classify. Still the fine-grained characterization of events from social data streams might prove a useful technique for our purpose. Social data is also used to detect emergency situations [17] and combined with satellite images to detect natural disasters [2]. In addition, it is also used to determine whether real-world events are fake or real, by fusing multiple modalities (social, text and visual) [4, 16].

When working with social data it is possible to use additional information such as URLs, users and hash tags as in [8] and [24]. For example, [24] utilize hashtags and user handles for pooling tweets and extracting latent topics of a higher quality. Similarly, in [13] textual and geographical data are combined to detect different type of users in geographical and textual social media data. Due to privacy and ethical concerns, no information about users is collected by the citizen reporting system discussed in this paper. However

when using metadata, such as time and geolocation, inspiration can be drawn from the above-mentioned related work. Methods for detecting generic micro events using multi-modal techniques are described in [15]. These micro events are different from the urban micro-events we seek. They are defined as transient occurrences within larger events.

Most related work in event detection seeks for events of a different nature, having a different geographical coverage and affecting a different number of citizens. However, the additional use of multimedia content is likely to also be beneficial to the classification of urban micro-events due to also having the relevant information in multiple modalities. While most approaches on event detection center on social media data or news articles, the main focus of this paper is detection of urban micro-events in a real world data set containing extremely heterogeneous multimedia data. Our objective is therefore aligned with the recent efforts of the multimedia community towards rethinking the very concept of event in the age of multimedia data that goes beyond simple text and visual modalities [32].

3 APPROACH

In this section we describe our approach to classifying urban micro-events, which consists of the following steps: (3.1) Discovering features for different modalities, (3.2) Creating unimodal classifiers using these features (3.3) Creating multimodal fusion techniques combining these classifiers.

3.1 Features

We propose several methods for extracting features out of text, image, geo and temporal data. The source code that shows how the features are extracted on a sample dataset can be found in the linked repository [35].

3.1.1 Textual. Since the textual description of an urban micro-event often has a large quantity of information, we will explore several methods of representing the data and evaluate what method works best for detecting urban micro-events. Motivated by the recent research in the information retrieval community, which shows that for some categories word embeddings work better, while for the others traditional vector space models still yield a better performance [38], we decided to evaluate both TF-IDF [31] and word2vec [25] representations. Our initial experiments with different variants of Latent Dirichlet Allocation [3] yielded unsatisfying performance, mostly due to the short length of the reports and a varying quality of conversational language used in them, which is why we do not report on them in this paper.

3.1.2 Visual. Similar to [18], we will use the 2048-dimensional output of the last layer before softmax of the ResNet50 model [12], pre-trained on ImageNet [10] as the visual features.

Another possibility is using the output of the network containing confidences for ImageNet classes, but our visual data is too different from ImageNet, which might result in poor classification performance. For example, bicycles are likely correctly detected, but the garbage containers and garbage bags will not. By using the hidden representation the features are more general. However, for

Figure 3: Examples of how the density and proximity of geographical object types is captured.

the creation of the graph embedding the output of the network will be used to reduce the number of nodes.

3.1.3 Geo objects. For creating a fingerprint capturing the geographical location of reports, a reference database with geographical data that describes the environment will be used. Using this we can create new features such as the distance to the closest container, the mean average distance of five double flowered chestnut trees and the number of residential buildings within n meters. In our case the data about geo objects and their location consist of 552.999 geo objects extracted from <https://maps.amsterdam.nl/>. This results in 1856 features describing the proximity and density of objects in the environment which are defined as follows

- Proximity is represented by taking the distance to the closest geo object of each type, but also by taking the mean of the closest five, ten and hundred objects per type.
- Density is represented by counting the occurrences of a geo object type within 25, 50, 100 and 200 meter.

Using available historical geographical data of urban micro-events we will create a historical profile of the area by using the same proximity and density features as for the geographical objects, this creates another 472 geo features on 57 different types of urban micro-events

3.1.4 Temporal and weather features. Temporal data is one-hot encoded information such as the month of the year, the weekday and the hour of that day. Furthermore historical weather data per hour is added, for example temperature, wind speed, snow and rain. This results in 37 temporal features and 18 weather features.

3.2 Unimodal classifiers

The best performing textual classifier has been implemented in a system that routes citizen reported urban micro events to the correct department of the City of Amsterdam. This allowed for the evaluation of this classifier in a real world setting and to gather multimodal data for the multimodal experiment.

The task of the implemented classifier is identifying the correct class using textual data. Whenever a prediction is made with a high enough probability the classifier automatically routes the customer service requests to the corresponding department. When the probability of the prediction is not higher than a predefined threshold experts will perform the task of the classification. After the classification has been made, by the classifier or by the expert, it is possible

Figure 4: A schematic example of the graph representation.

that a mistake is made. Whenever a department receives a customer service request that is not classified correctly they will assign it to the correct class. This allows evaluation of both the classifier and experts performing the same task. In addition customer evaluation is done by asking the initial reporter to rate the entire process.

The implemented textual classifier was used to gather data of performance in a real world setting and the results can be seen in Table 1. While the textual classifier scored 89% accuracy in a real world setting, experts performing the same task using all modalities for the classification score 91% accuracy.

When looking at a survey done for 2768 customers reports, we find that reports that have been automatically classified receive a customer satisfaction of 3,2/5. The manually classified reports receive a 2,9/5. This leads us to believe that a decrease in time needed to resolve the issue likely had a positive effect on customer satisfaction. Since the textual classifier is outperformed by an expert that can see location, time and the added image as well it is likely performance can increase by adding more modalities to the classifier. Improved classification performance in combination with much lower processing time than in case of manual annotation seems to be a promising path to an increased customer satisfaction.

3.3 Multimodal Fusion

To assess what kind of fusion method works best for the multimodal classification of urban micro-events we will consider a number of different methods, starting with a simple early fusion, followed by the creation of a graph embedding. After this we will discuss a late fusion and a hybrid fusion method of creating a classifier.

The classifications for all types of fusion are done with XGBoost [7] and a Logistic Regression to be able to evaluate performance of the classifiers on different modalities and fusion techniques.

3.3.1 Early fusion. For every possible combination of feature sets a classifier is created. Here the features are simply concatenated and used as input for the classifier. An example of an early fusion is:

Visual||Textual||Geo||Time||Weather

3.3.2 Graph embedding. To represent different modalities in one representation it is possible to place them in a graph. To evaluate this method a graph has been created for classification purposes. Let

$G = (V, E)$ be our undirected weighted graph with the set of nodes V and the set of edges E . We choose to use an undirected graph since the relations between the nodes and edges are symmetric.

The following nodes are added:

- Report nodes $R = \{r_1, r_2, \dots, r_n\}$ For every report in the dataset a node is created.
- Geo objects. $G = \{g_1, g_2, \dots, g_n\}$ For all reports the closest geo objects are added.
- Geo location. $L = \{l_1, l_2, \dots, l_n\}$ For all report the coordinates are added as a node.
- Visual concepts. $V = \{v_1, v_2, \dots, v_n\}$ For all images the top two visual concepts are added to the graph.
- Words. $T = \{t_1, t_2, \dots, t_n\}$ The top 5000 words in the corpus are added to the graph.
- Time. $H = \{h_1, h_2, \dots, h_{24}\}$ and $W = \{w_1, w_2, \dots, w_7\}$ For every weekday and hour a node is created.

And the following edges are created:

- From the reports edges are created to the two closest geo objects, the weight of the edge is the distance.
- The reports are linked to the top two visual objects with the highest probability, with that probability as the weight of the edge.
- For all the words in the text of the report an edge is created, with TF-IDF as the weight of the edge.
- For the weekday and hour of the report a binary weighted edge is added, and for all the neighboring weekdays and hours also an edge is created.
- An edge from the report to the geographical location is made. The geographical location is linked to the two closest geographical locations.

In Figure 4 a schematic overview is given of the connection between edges and nodes. After the creation of the graph node2vec [11] was used to create a 256-dimensional representation for every node. The node2vec framework learns low-dimensional representations for all the nodes in the graph. This is done by optimizing a neighborhood preserving objective by simulating random walks. The reports representation will then be used as a feature for classification and evaluation purposes.

3.3.3 Late fusion. Stacking will be used as ensemble learning technique to combine information out of two or more models into a new model. This is done by using the probabilistic output of classifiers as input for a new classifier. An example of a late fusion is:

ProbVisual||ProbTextual||ProbGeo||ProbTime||ProbWeather

3.3.4 Hybrid. The hybrid fusion classifier is a combination of the early fusion and the late fusion classifier. For some feature sets the late fusion probabilistic output is used. For other feature sets the original features will be as in the early fusion. These features are combined as input for the classifier. For every possible combinations of early and late fusion features a classifier is created and the predictions are evaluated on the test set to determine what the optimal method of hybrid fusion is on this specific dataset. An example of a hybrid fusion is:

Text||ProbVisual||ProbGeo||ProbTime||ProbWeather

Figure 5: Performance of issue level ($n = 57$) classification using logistic regression, for all none fusion classifiers and the best performing fusion classifier, with and without text. The results are reported in terms of F1-score and for every class the normalized support is shown. For improved readability only the 40 classes with the highest support are shown.

4 EXPERIMENTAL SETUP

In this section we describe the evaluation criteria and the data we use for the experiments.

4.1 Evaluation criteria

For evaluation weighted F1 is used since the smaller classes are also important to classify correctly and we seek for a balance of recall and precision. To allow for comparison evaluation will be done with the same train/test split 80/20 for all experiments.

4.2 Data

The data used is a set of citizen reports from the City of Amsterdam. For the experiments multiple different subsets have been used:

- 523,651 reports with textual information and their corresponding issue class have been used for creating a textual classifier.
- Of those reports 29,408 reports also had visual data.
- 9,362 of the reports with visual data had their label corrected by a domain expert.

The multimodal experiments will be done on the subset of 29,408 reports that has visual data. The reports can be grouped in eight main issue classes. The largest group is garbage and most of these reports are about bulky waste. Other issues in this group are litter, full garbage containers, broken garbage containers and construction waste. The second largest group is about anything related to roads, traffic and furniture. This group consists of issues about the maintenance of roads, traffic signs, clogging drains, slippery roads, broken streetlights, issues with playgrounds and dangerous traffic situations. The third largest group is about disturbance in the public space, ranging from bicycle wrecks, illegal parking, objects blocking the sidewalk, noise nuisance and dog poo. The group of green and water is about any green that needs maintenance or quay wall that needs to be repaired. Animals is about disturbances from rats, wasps or pigeons. Disturbances by people, business or boats consists of noise nuisance, smell disturbances or e.g. speeding boats. Note that these are weak labels due to them being collected from a real world system, making it possible that an urban micro-event is labelled incorrectly.

Table 1: Customer satisfaction and performance (Accuracy with eight classes) of domain experts and a textual classifier.

	Customer satisfaction	Resolved in same issue class
Domain expert	2.9/5	91%
Textual classifier	3.2/5	89%

Table 2: Top textual and non textual results for best hybrid fusion experiments and no fusion on the main classes ($n = 8$). Logistic regression (LR) and XGBoost (XGB) used for classification. Weighted F1-score is used for evaluation.

Classifier	Features	F1
LR	text, graph, prob_time, prob_image	0.882
XGB	geo, image, graph, geo_hist, prob_text	0.875
LR	text	0.865
XGB	text	0.851
LR	graph	0.844
XGB	graph	0.810
XGB	prob_image, prob_geo_hist, prob_weather, time, geo	0.737
LR	image, geo_hist, prob_geo	0.730
LR	image	0.710
XGB	image	0.706
XGB	geo_hist	0.508
XGB	geo	0.505
LR	geo	0.481
LR	geo_hist	0.464
XGB	weather	0.400
XGB	time	0.391
LR	time	0.380
LR	weather	0.380

Table 3: Evaluation of several textual classification methods trained on 418.920 samples and evaluated on 104.731 samples. The task was detecting which of seven main classes the report belongs to.

Classifier	F1 macro avg	F1 micro avg
TF-IDF + LR	0.79	0.87
Bidirectional CNN+LSTM	0.72	0.88
W2V on reports + LR	0.73	0.83
W2V 160 combined [22] + LR	0.60	0.72

5 EXPERIMENTAL RESULTS

In this section several experiments and their results are discussed, to answer the following research questions.

- What method of textual classification works best for the classification of urban micro-events?
- What visual, geo, time and textual features can be used for the detection of urban micro-events?

Table 4: Top textual and non textual results for best hybrid fusion experiments and no fusion on the issue level ($n = 57$). Logistic regression (LR) and XGBoost (XGB) used for classification. Weighted F1-score is used for evaluation.

Classifier	Features	F1
LR	text, graph, prob_geo, prob_time, prob_image	0.700
XGB	prob_text, prob_time, image	0.699
LR	text	0.675
XGB	text	0.655
LR	graph	0.625
XGB	graph	0.573
XGB	geo_hist, prob_time, geo, prob_image, weather ...	0.474
LR	image	0.446
XGB	image	0.427
XGB	geo_hist	0.258
XGB	geo	0.255
LR	geo	0.213
LR	geo_hist	0.202
XGB	weather	0.129
XGB	time	0.120
LR	time	0.107
LR	weather	0.091

- What type of fusion is effective for the classification of urban micro-events?

The source code of experiments can be found in the linked repository [35].

5.1 Textual classifier

For the textual classifier evaluation results are shown in Table 3. The best performing model with regard to F1 macro measure was TF-IDF with a logistic regression. CNN+LSTM performed best when looking at overall performance, but failed detecting some smaller, yet important classes.

5.2 Multimodal classifier

The best performing multimodal classifier was created by using a Logistic regression for classification and a hybrid fusion method, for both class level classification and issue level classification, the results can be seen in Tables 2 and 4. The following features have been used by the best performing classifier: textual, late fusion geographical objects, late fusion time, late fusion visual and the graph embedding. The unimodal classifiers using time, weather or geo features yielded a better performance when using XGBoost, but the multimodal classifiers and the unimodal classifiers using visual and text data performed better when using logistic regression.

The results show that it is possible to improve on the textual classifier. For the class level classification adding visual information, time information and the graph embedding to the text with a hybrid fusion method increased performance from .865 to .882. When comparing this increase with the experts annotations as seen in Table 1, it is possibility that on a larger dataset the multimodal

Table 5: Several examples of classification with example features.

	Example 1	Example 2	Example 3
Image
Visual features	mountain_bike = 0.81 bicycle-built-for-two = 0.03	ash_can = 0.62 turnstile = 0.02	chainlink_fence = 0.17 ashcan = 0.10
Text	They are still here	There is garbage next to the container	Graffiti
Geo	bicycles_wreck_within_100m = 6 bicycles_wreck_mean_5_closest = 0.04 km	garbage_container_within_200m = 7 hist_garbage_nearest = 0.01 km	hist_graffiti_within_100m = 3 hist_graffiti_nearest = 0.01 km
Time	2018-12-25 09.05	2018-11-22 15.30	2018-12-22 15.39
Label	Bicycles wreckage	Bulky waste	Graffiti
Prediction	Bicycles wreckage	Litter	Graffiti
True class	Bicycles wreckage	Home garbage	Graffiti

classifier will outperform domain experts performing the same task while increasing customer satisfaction.

For the issue level classification adding visual, geo, time information and the graph embedding to the textual data leads to an increase in performance from 0.675 to 0.700. This is a small increase but it has to be taken into consideration that the data used partly has weak labels generated by using the textual data, which creates a bias for the performance of the textual classifier. Even with this bias the performance has increased, making the results promising. Also the data used has a large class unbalance as can be seen in Figure 5, making the size of the available training data a potential limitation. This makes it likely that when more data is available performance will increase even further.

In Figure 5 the F1-score can be seen per class and per feature set. For most classes the multimodal classifier improves on performance over the textual baseline. An example of classes that improve from the multimodal classifier are "Sunken boats", "Floating garbage" and "Graffiti". These are all classes that have a clear visual hint, like a boat, water or graffiti on the image. Several example classifications are shown in Table 5.

A few classes do not get classified better when using the multimodal classifier. Examples of these classes are "loud boats", "nuisance from music" and "plastic container being full." The Loud boats textual classifier scores perfectly in evaluation, and the addition of other modalities decreases performance. The cause of this could be the partial use of weakly supervised data and the absence of relevant information in the other modalities.

Finally we observe that for almost all classes with sufficient class size the multimodal classifier works better than the textual classifier.

6 CONCLUSION

In this paper we investigated the potential for automatically classifying urban micro events based on heterogeneous information describing them in citizen reports, which ranges from text and image to metadata about event geolocation, time and weather. We further deploy a number of approaches for fusing information extracted from different modalities, including traditional early and late fusion, but also novel representation learning on graphs and hybrid fusion. Finally, we investigate contribution of individual modalities to the overall classification performance. The experiments were conducted on a real-world dataset collected from a live citizen reporting system. Our main conclusion is that a multimodal classifier yields a higher performance than the unimodal alternatives. Text appears to be the single most important modality, which is expected since the citizens are usually careful when describing the issues. In addition, our experiments show that the representation learning on graphs is effective in embedding heterogeneous information extracted from all different modalities into a compact, but discriminative representation. Indeed, a hybrid fusion of such created representation with different modalities associated with the citizen reports emerges as the overall best performing classification approach. In our future work we will further investigate the undoubtedly large potential of multimodal graph embeddings and the possibilities of incorporating different fusion mechanisms into representation learning.

ACKNOWLEDGMENTS

This project has received funding from the City of Amsterdam.

REFERENCES

- [1] George Awad, Asad Butt, Keith Curtis, Yooyoung Lee, Jonathan Fiscus, Afzal Godil, David Joy, Andrew Delgado, Alan F. Smeaton, Yvette Graham, Wessel Kraaij, Georges Quénot, Joao Magalhaes, David Semedo, and Saverio Blasi. 2018. TRECVID 2018: Benchmarking Video Activity Detection, Video Captioning and Matching, Video Storytelling Linking and Video Search. In *Proceedings of TRECVID 2018*. NIST, USA.
- [2] Benjamin Bischke, Patrick Helber, Christian Schulze, Venkat Srinivasan, Andreas Dengel, and Damian Borth. 2017. The Multimedia Satellite Task at MediaEval 2017. In *MediaEval*.
- [3] David M. Blei. 2012. Probabilistic Topic Models. *Commun. ACM* 55, 4 (April 2012), 77–84. <https://doi.org/10.1145/2133806.2133826>
- [4] Christina Boididou, Symeon Papadopoulos, Yiannis Kompatsiaris, Steve Schiffrer, and Nic Newman. 2014. Challenges of computational verification in social multimedia. In *Proceedings of the 23rd International Conference on World Wide Web*. ACM, 743–748.
- [5] Maged N Kamel Boulos, Bernd Resch, David N Crowley, John G Breslin, Gunho Sohn, Russ Burtner, William A Pike, Eduardo Jezewski, and Kuo-Yu Slayer Chuang. 2011. Crowdsourcing, citizen sensing and sensor web technologies for public and environmental health surveillance and crisis management: trends, OGC standards and application examples. *International journal of health geographics* 10, 1 (2011), 67.
- [6] Longbiao Chen, Jérémie Jakubowicz, Dingqi Yang, Daqing Zhang, and Gang Pan. 2017. Fine-Grained Urban Event Detection and Characterization Based on Tensor Cofactorization. *IEEE Trans. Human-Machine Systems* 47, 3 (2017), 380–391.
- [7] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. ACM, 785–794.
- [8] Tao Chen, Xiangnan He, and Min-Yen Kan. 2016. Context-aware image tweet modelling and recommendation. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 1018–1027.
- [9] Guillem Collell, Ted Zhang, and Marie-Francine Moens. 2017. Imagined Visual Representations as Multimodal Embeddings. In *AAAI*. 4378–4384.
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. IEEE, 248–255.
- [11] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In *Proceedings of the 22nd ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 855–864.
- [12] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [13] Tianran Hu. 2018. Decode Human Life from Social Media. In *2018 ACM Multimedia Conference*. ACM, 820–824.
- [14] Yuqian Huang, Yue Li, and Jie Shan. 2018. Spatial-Temporal Event Detection from Geo-Tagged Tweets. *ISPRS International Journal of Geo-Information* 7, 4 (2018), 150.
- [15] Kasthuri Jayarajah and Archan Misra. 2016. Can instagram posts help characterize urban micro-events?. In *Information Fusion (FUSION), 2016 19th International Conference on*. IEEE, 130–137.
- [16] Zhiwei Jin, Juan Cao, Han Guo, Yongdong Zhang, and Jiebo Luo. 2017. Multimodal fusion with recurrent neural networks for rumor detection on microblogs. In *Proceedings of the 25th ACM international conference on Multimedia*. ACM, 795–816.
- [17] Diptesh Kanojia, Vishwajeet Kumar, and Krithi Ramamritham. 2016. Civique: Using Social Media to Detect Urban Emergencies. *arXiv preprint arXiv:1610.04377* (2016).
- [18] Douwe Kiela and Léon Bottou. 2014. Learning image embeddings using convolutional neural networks for improved multi-modal semantics. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 36–45.
- [19] Josef Kittler, Mohamad Hater, and Robert PW Duin. 1996. Combining classifiers. In *Proceedings of 13th international conference on pattern recognition*, Vol. 2. IEEE, 897–901.
- [20] Hongzhi Li, Joseph G Ellis, Heng Ji, and Shih-Fu Chang. 2016. Event specific multimodal pattern mining for knowledge base construction. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 821–830.
- [21] Jingjing Li, Lei Zhu, Zi Huang, Ke Lu, and Jidong Zhao. 2018. I read, I saw, I tell: Texts Assisted Fine-Grained Visual Classification. In *2018 ACM Multimedia Conference*. ACM, 663–671.
- [22] Shuying Lin, Huosheng Xie, Liang-Chih Yu, and K Robert Lai. 2017. SentiNLP at IJCNLP-2017 Task 4: Customer Feedback Analysis Using a Bi-LSTM-CNN Model. *Proceedings of the IJCNLP 2017, Shared Tasks* (2017), 149–154.
- [23] Chao-Hong Liu, Declan Groves, Akira Hayakawa, Alberto Poncelas, and Qun Liu. 2018. Understanding Meanings in Multilingual Customer Feedback. *arXiv preprint arXiv:1806.01694* (2018).
- [24] Rishabh Mehrotra, Scott Sanner, Wray Buntine, and Lexing Xie. 2013. Improving lda topic models for microblogs via tweet pooling and automatic labeling. In *Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval*. ACM, 889–892.
- [25] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. Efficient Estimation of Word Representations in Vector Space. *CoRR abs/1301.3781* (2013). <http://dblp.uni-trier.de/db/journals/corr/corr1301.html#abs-1301-3781>
- [26] Mohamed Mohandes, Mohamed Deriche, and Salihu O Aliyu. 2018. Classifiers combination techniques: A comprehensive review. *IEEE Access* 6 (2018), 19626–19639.
- [27] Dietmar Offenhuber. 2014. Infrastructure legibility - a comparative analysis of open311-based citizen feedback systems. *Cambridge Journal of Regions, Economy and Society* (2014), rsu001.
- [28] Kunwoo Park, Jaewoo Kim, Jaram Park, Meeyoung Cha, Jiin Nam, Seunghyun Yoon, and Eunhee Rhim. 2015. Mining the Minds of Customers from Online Chat Logs. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management (CIKM '15)*. ACM, New York, NY, USA, 1879–1882. <https://doi.org/10.1145/2806416.2806621>
- [29] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. DeepWalk: Online Learning of Social Representations. In *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '14)*. ACM, New York, NY, USA, 701–710. <https://doi.org/10.1145/2623330.2623732>
- [30] Shengsheng Qian, Tianzhu Zhang, and Changsheng Xu. 2016. Multi-modal multi-view topic-opinion mining for social event analysis. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 2–11.
- [31] Juan Ramos et al. 2003. Using tf-idf to determine word relevance in document queries. In *Proceedings of the first instructional conference on machine learning*, Vol. 242. 133–142.
- [32] Stevan Rudinac, Tat-Seng Chua, Nicolas Diaz-Ferreira, Gerald Friedland, Tatjana Gornostaja, Benoit Huet, Rianne Kaptein, Kristin Linden, Marie-Francine Moens, Jaakko Peltonen, et al. 2018. Rethinking summarization and storytelling for modern social multimedia. In *International Conference on Multimedia Modeling*. Springer, 632–644.
- [33] Stevan Rudinac, Alan Hanjalic, and Martha Larson. 2013. Generating visual summaries of geographic areas using community-contributed images. *IEEE Transactions on Multimedia* 15, 4 (2013), 921–932.
- [34] Cees GM Snoek, Marcel Worring, and Arnold WM Smeulders. 2005. Early versus late fusion in semantic video analysis. In *Proceedings of the 13th annual ACM international conference on Multimedia*. ACM, 399–402.
- [35] M. Sukel. 2019. Multimodal classification of urban micro-events. <https://github.com/maartensukel/multimodal-classification-of-urban-micro-events>.
- [36] Mengfan Tang, Siripen Pongpaichet, and Ramesh Jain. 2016. Research challenges in developing multimedia systems for managing emergency situations. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 938–947.
- [37] Mengfan Tang, Siripen Pongpaichet, and Ramesh Jain. 2016. Research Challenges in Developing Multimedia Systems for Managing Emergency Situations. In *Proceedings of the 24th ACM International Conference on Multimedia (MM '16)*. ACM, New York, NY, USA, 938–947. <https://doi.org/10.1145/2964284.2976761>
- [38] Christophe Van Gysel, Maarten de Rijke, and Marcel Worring. 2016. Unsupervised, Efficient and Semantic Expertise Retrieval. In *Proceedings of the 25th International Conference on World Wide Web (WWW '16)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 1069–1079. <https://doi.org/10.1145/2872427.2882974>
- [39] Maximilian Walther and Michael Kaisser. 2013. Geo-spatial event detection in the twitter stream. In *European Conference on Information Retrieval*. Springer, 356–367.
- [40] Bokun Wang, Yang Yang, Xing Xu, Alan Hanjalic, and Heng Tao Shen. 2017. Adversarial cross-modal retrieval. In *Proceedings of the 25th ACM international conference on Multimedia*. ACM, 154–162.
- [41] Xiaoshan Yang, Tianzhu Zhang, and Changsheng Xu. 2016. Abnormal event discovery in user generated photos. In *Proceedings of the 24th ACM international conference on Multimedia*. ACM, 47–51.
- [42] Yiming Yang, Tom Pierce, and Jaime Carbonell. 1998. A study of retrospective and on-line event detection. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*. ACM, 28–36.
- [43] Tongtao Zhang, Spencer Whitehead, Hanwang Zhang, Hongzhi Li, Joseph Ellis, Lifu Huang, Wei Liu, Heng Ji, and Shih-Fu Chang. 2017. Improving event extraction via multimodal integration. In *Proceedings of the 25th ACM international conference on Multimedia*. ACM, 270–278.