

Urban Emotions – Geo-semantic Emotion Extraction from Technical Sensors, Human Sensors and Crowdsourced Data

Bernd Resch^{*,**,***}, Anja Summa^{**}, Günther Sagl^{*,**}, Peter Zeile^{****}, Jan-Philipp Exner^{****}

* Department of Geoinformatics, University of Salzburg, Austria

** Institute of Geography – GIScience, Heidelberg University, Germany

*** Center for Geographic Analysis, Harvard University, USA

**** Computergestützte Planungs- und Entwurfsmethoden (CPE), University of Kaiserslautern, Germany

Abstract How people in the city perceive their surroundings depends on a variety of dynamic and static context factors such as road traffic, the feeling of safety, urban architecture, etc. Such subjective and context-dependent perceptions can trigger different emotions, which enable additional insights into the spatial and temporal configuration of urban structures. This paper presents the Urban Emotions concept that proposes a human-centred approach for extracting contextual emotional information from human and technical sensors. The methodology proposed in this paper consists of four steps: 1) detecting emotions using wristband sensors, 2) “ground-truthing” these measurements using a People as Sensors location-based service, 3) extracting emotion information from crowdsourced data like Twitter, and 4) correlating the measured and extracted emotions. Finally, the emotion information is mapped and fed back into urban planning for decision support and for evaluating ongoing planning processes.

Keywords: People as Sensors, Urban Planning, VGI, Crowdsourcing, Emotion Detection

1. Introduction

The development of a digital city into an intelligent city offers new opportunities to capture spatial and temporal data in near real time. The enabling

driver is a continuous connection between the physical and the digital worlds by a variety of "sensors" - from calibrated measurement equipment to human sensors (Goodchild 2007). Recent developments use psychophysiological measurements in urban space, for instance, to map emotions (Zeile et al. 2009), or mobile phone data and social network data to assess collective human behaviour patterns (Sagl et al. 2012). These new data and information layers can provide additional insights into the development of both the physical and social structures of inherently complex and dynamic urban environments.

Yet, the city as a functional construct shall not only be seen as a place of technological infrastructure, financial transactions, a network of technical nodes, a geographical agglomeration area or a political landscape, but more as an actuated multi-dimensional conglomerate of heterogeneous processes, in which the citizens are the central component (Resch et al. 2012). This interaction between humans and urban space, i.e., where, when and in particular how people respond to urban processes, needs more attention from a quantitative point of view in order to derive more reliable results compared to current urban analysis approaches.

Urban Emotions aims to address this shortcoming by providing a human-centred approach for extracting contextual emotion information from technical sensor data (measurements from calibrated bio-sensors) and human sensor data (subjective observations by citizens). The results are used in the domain of urban planning for decision support and the evaluation of ongoing planning processes (Zeile et al. 2014). Like this, the realization of a Smart City is not only to be tackled from a technological viewpoint (as most previous research efforts did), but from a human-centred viewpoint that claims that a city requires "Smart Citizens" to be intelligent itself.

Figure 1 illustrates the general *Urban Emotions* concept that comprises four steps: 1) detecting emotions using wristband sensors, 2) "ground-truthing" these measurements using a smartphone-based People as Sensors location-based service (LBS) in near real time, 3) extracting emotion information from crowdsourced data like Twitter (detecting the type of emotion), and 4) correlating the measured and extracted emotions. Subsequently, the emotion information is mapped and fed back into urban planning processes. The paper at hand mainly covers the two modules of ground-truthing emotion information using the People as Sensors concept and the extraction of emotion information from text-based Volunteered Geographic Information (VGI). For this reason, we employ a graph-based semi-supervised learning (SSL) algorithm, i.e., we use a small set of labelled training data to assign a label to each Tweet - where the process of *labelling* stands for assigning an emotional category to each Tweet (s. section 4).

Thus in contrast to previous approaches, which have relied on methods from a single discipline like GIScience, computational linguistics (CL), sociology, or computer science (CS), we propose a trans-disciplinary method.



Figure 1: Urban Emotions concept: 1) emotion sensing, 2) ground-truthing using People as Sensors, 3) extraction of emotion information from VGI, 4) correlating measured and extracted emotions; plus visualisation and enrichment of urban planning processes.

This paper is structured as follows: After this introduction, we provide a description of related work including a clear identification of research gaps (section 2). Then, we present our methodology for ground-truthing emotion measurements using a People as Sensors app (section 3), followed by a description of our proposed methodology for extracting emotion information from Twitter Tweets (section 4). Finally, we discuss our approach's integrability into urban planning processes (section 5), and end the paper with a number of key conclusions and future research avenues (section 6).

2. Related Work

For the scope of this paper, related work needs to be examined in three areas: using emotion measurements for urban applications (sub-section 2.1), emotion extraction from VGI (sub-section 2.2), and using emotion information in the field of urban planning (sub-section 2.3).

2.1. Measuring Emotions in the Urban Context

Salesses et al. (2013) describe an approach to extract information about citizens' perception of safety. They developed an online platform that allows people to compare and rate two randomly selected pictures showing different urban environments from a street-view perspective with respect to safety. This allows for a qualitative subjective assessment of a static situation (i.e., a picture). Yet, a continuous subjective perception and quantitative assessment with respect to the dynamic situational urban context, providing insights beyond a snapshot, is not addressed in the approach.

Gartner (2010) investigates the use of emotions to support way-finding tasks. The paper describes methods how emotions can be sensed and presents a conceptual framework for the use of emotion information in way-finding. However, no concrete field tests have been carried out, and no implementation and validation of the concept have been performed. The Urban Emotions approach goes one step further by applying emotion information in concrete real-world urban management and planning use cases. In addition, our approach performs a "ground-truth" for the emotional spike (s. section 3).

The approach by Klettner and Schmidt (2012) aims to cartographically visualise emotions. In other words, the use of emotion information is addressed from a purely cartographic presentation viewpoint, while no in-depth analysis and feedback to real-world processes is done.

2.2. Extracting Emotion Geo-Information from VGI

The research field of *Sentiment Analysis* only deals with a word's, sentence's, or document's polarity, i.e., whether it conveys a positive, negative, or neutral sentiment. Additionally, research has been done to determine the expressed sentiment's strength (Liu & Zhang, 2012). In the Urban Emotions approach, we use a more sophisticated emotion model as knowledge purely concerning a Tweet's polarity is not sufficient.

Emotion Detection from Tweets focuses on classifying Twitter posts according to a number of distinct emotions. The two approaches by Roberts et al. (2012) and Bollen et al. (2011) analyse the results of large events that cover the entire USA and influence Twitter traffic for one or several days. In doing so, singular small-scale variations of Twitter "traffic" might be overseen. For urban planning, these smaller events may be important as they affect smaller, local areas. Besides, both approaches lack the geographic component, which is essential to our approach as laid out in section 1. Also, these previous efforts neglect the possibilities that arise for emotion detection from emoticon analysis. Finally, both approaches propose insufficient methodologies that we aim to improve: Roberts et al. (2012) only work in a

supervised manner and thus require a large dataset to produce satisfying results. Bollen et al. (2011) do not evaluate their approach against a ground-truthed gold standard, but correlate their results manually to events that occurred at the same time the Tweets were sent. In our work, we overcome these methodological problems through using a semi-supervised learning approach, which can be applied to a dataset with few labelled, but numerous unlabelled instances (a smaller data set), and evaluated on a test set.

Another approach by Hauthal (2013) aims to detect emotions in VGI and to map emotional hot spots in a city. However, the approach works on a simple syntactical word-matching algorithm that is not able to cope with the complexity of unstructured text data like Twitter Tweets.

2.3. Using Emotion Information in Urban Planning

In the context of urban planning, only few research efforts deal with the question, how planners can make use of subjective feelings in urban surroundings and perceptions of citizens. Lynch's approaches of the "Image of the City" and "Mental Maps" suggest that citizens investigate a city during a walk and afterwards sketch a map of the investigated area out of their mind. Yet in this approach, test persons need excellent drawing skills to produce the maps. Sorin Matei et al. (2001) introduced a first digital approach of showing feelings in a digital map. This "Mental Maps Concept" visualises feelings not only on a digital map, but also on a three-dimensional model of the city for a better understanding of "geolocation of fear" in Los Angeles, USA. The art project "Biomapping" by Christian Nold was the first work that combined "emotional data" (physiological parameters like skin resistance levels) with GPS datasets (Nold 2009).

Despite such technical challenges, all experts in this field pointed out that citizens are the main and most important actor in urban planning processes. Thus, our approach focuses on using emotion information from a variety of sources and on feeding it back directly to urban planning.

3. Emotion Measurements: People as Sensors

As laid out in section 2, one of the major shortcomings of previous approaches is that emotions measured by technical sensors cannot be unambiguously correlated with a person's actual emotion (the type of emotion) and the cause why an emotion occurred (the context of the emotion). This is because currently available emotion sensors are only capable to detect a person's emotional spike, but not their causal trigger. Most current sensors observe a person's "additional heart rate", skin conductance, body temperature – the latter two variables are counter-rotating, i.e., in case an emotion-

al spike occurs, skin conductance rises and body temperature drops because the person emits cold sweat.

To account for this shortcoming, we designed and implemented a People as Sensors LBS, through which persons can enter the emotional category and the according context in case an emotional spike occurs. In the design of the app, we aimed to fulfil common guidelines for mobile applications. First, we followed the design principles: “make it direct”, “keep it lightweight”, “stay on the page”, “provide an invitation”, “use transitions”, and “react immediately”, as defined by Scott and Neil (2009). According to these principles, we decided to design a simple interface that offers users the possibility to first input the type of emotion, followed by a screen to input the context of the emotion. Both inputs immediately guide the user back to the start page. Thus, the interface is lightweight in terms of design elements, the number of “clicks” required, and the information provided to the user. To comply with the requirement of displaying an invitation, we notify the user when their input is required; i.e., when the wristband sensor measures an emotional spike, the user is requested to input the type and context of the emotion. Furthermore, we decided to greatly avoid strong colouring of the app and assigning colours to emotions in order to prevent emotional biases (Mohammad 2011). Figure 2 shows the application interface in three steps, 1) the input for the type of emotion, 2) the input for the context, and 3) the main page for submitting the data.

In order to comply with the requirement of immediate feedback to the user, we intend to show a map containing the emotions of a city right after a user has entered their information. Like this, the motivation for entering emotion information shall be kept high and users can immediately “compare” their own impressions with other persons’ perceptions.

In the next design phase, we also aim to improve the input modalities. One option that we are currently investigating is the use of a “colour wheel” (comparable to the RGB colour wheel) to enable the input of a combined emotion in a quasi-continuous emotion space instead of just a single emotion. Even though this does not comply with the idea of avoiding bias (s. above), it could be important to allow for inputting several emotions as emotional categories are mostly related to others and they do not appear as single isolated emotions. Furthermore, we investigate the possibility of allowing users to input the intensity of their emotion, which can be used for assessing the intensity’s correlation with the spikes’ amplitudes measured by the emotion sensor.



Figure 2: Three screens for entering the type of emotion, the context for the emotion and for sending off the information (left to right).

4. Emotion Extraction from VGI

As mentioned in section 1, previous approaches of extracting information, particularly emotion information, from crowdsourced data have relied on methods from a single discipline. The approach proposed in this paper, to our knowledge, is the first method that extracts information in a trans-disciplinary algorithm, using methods from computational linguistics and GIScience.

Previously, little research has been conducted on emotion extraction on Twitter. Consequently, few emotion-annotated Tweet datasets are available (e.g., Roberts et al. 2012). Yet, none of these corpora is sufficient for our purpose because they all lack the geographical component. Thus, our workflow involves the creation of a manually annotated emotion-tagged Twitter corpus, which will be employed for both training and evaluation purposes.

4.1. Basic Workflow

The basic workflow for our approach is depicted in Figure 3. It starts with a set of unlabelled, raw Twitter posts along with their metadata (location, time, user ID, post ID, etc.), and results in a set of labelled Tweets (where

the labels are the emotion categories). The workflow comprises five distinct steps that are described below.




Figure 3: Workflow for extracting emotion information from twitter tweets.

The **first step** is to filter the Tweets by language because in our first use case, we use only English Tweets to respect our annotators' language skills and facilitate manual inspection. Furthermore, the reduction to English language simplifies data processing in our initial field tests due to a large number of Natural Language Processing (NLP) resources.

The **second step** is the annotation process, in which Tweets are classified manually by volunteers with respect to whether they contain one of the emotions of interest or are considered neutral. Like this, we create a “gold standard” for the evaluation later in the workflow. This gold standard is the basis for a quantitative evaluation of our results to prove the effectiveness of our approach. In other words, a quantitative evaluation of our results is necessary to prove that our approach can actually detect a particular emotion in a single Tweet. As mentioned in section 2, qualitative evaluation methods like manually correlating emotion detection results to large-scale events do not ensure this. For our research, we use the emotion model proposed by Ekman and Friesen (1971), which distinguishes between the six basic emotions joy, anger, fear, sadness, surprise, and disgust. The reasons for using this model are 1.) that it consists a solid, well-established emotion model, and 2.) that it is used in similar research, e.g., by Roberts et al. (2012), which ensures comparability of the datasets, and the results and proves that this theoretical model is applicable to Tweets.

The **third step** serves for selecting the subset of labelled Tweets that are employed as seeds for the subsequent graph-based semi-supervised learning (SSL) algorithm. Seed selection is a critical step in semi-supervised settings, as the seeds strongly influence the program's output as the program can only use labels that are contained in the seed set, which thus needs to contain all labels it is supposed to use. Additionally, the distribution of labels over the seeds influences the distribution of labels in the result. Thus, we are investigating two ways of seed selection, 1) to align the seed distribu-

tion with the actual emotion distribution in the entire dataset, or 2) to select uniformly distributed seeds.

The **fourth step** performs the actual emotion detection and classification, and is further divided into four sub-steps. Since this is the core part of our workflow, these steps are described in detail in sub-section 4.2 below.

The **fifth step** is the evaluation of the labelling process' results (s. sub-section 4.2) against our hand-annotated gold standard. The subjective nature of the task of identifying emotions necessitates the comparison of the computer's performance against that of human beings. A suitable evaluation measure must be applied, depending on the importance of precision ("positive predictive value") and recall ("sensitivity") or their ratio ("F1 score").

4.2. Similarity Computing and Label Propagation

As mentioned above, step four in our overall workflow is the central part of detecting emotion information. In this step, unlabelled Tweets are converted to labelled Tweets, i.e., each Tweet is assigned an emotional category. Figure 4 illustrates the workflow for the labelling procedure. The single steps are explained below.

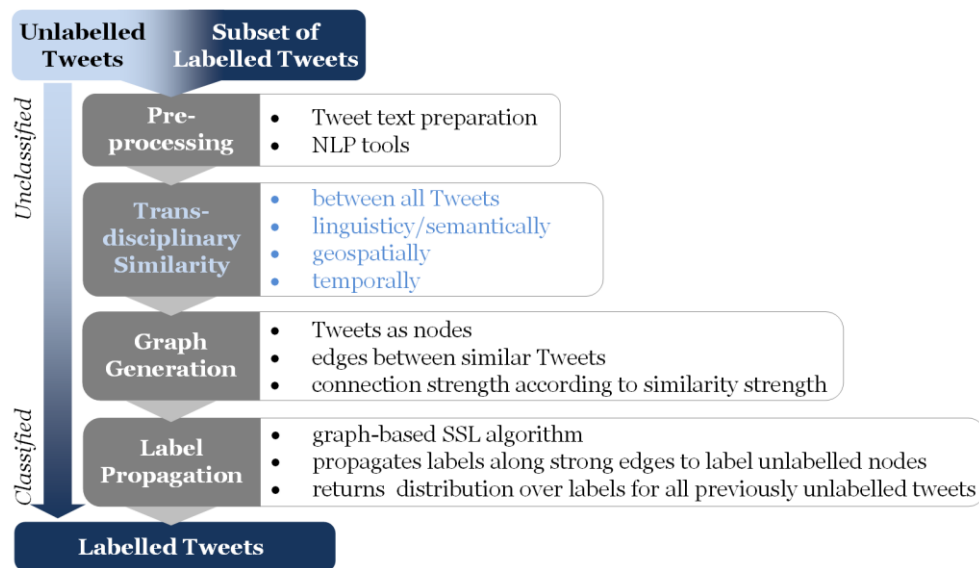


Figure 4: Workflow for labelling Tweets in a trans-disciplinary approach to similarity.

First, the Tweet text is **pre-processed** using basic text processing algorithms. This process includes enriching the text with different kinds of se-

mantic information to reveal the text's hidden structure. We apply the part-of-speech (POS) tagger presented by Owoputi et al. (2013) to detect emoticons, negations, and other specific POS. Pre-processing in our case also includes constraining possible emotion labels for unlabelled Tweets by comparing their content to the Affective Norms for English Words (ANEW) word list (Bradley and Lang 1999).

Second, the **similarity** between all Tweets, both seeds and unlabelled, is computed. This, along with the Tweets themselves, constitutes the input for the graph constructed in the next sub-step. We compute similarity according to three factors: *linguistic and semantic content*, *temporal distance*, and *geo-spatial distance*. We argue that Tweets, which are alike with respect to certain aspects, such as having emotional words, emoticons, and other linguistic features in common, will likely also carry the same emotion.

The main novelty of our approach is the combination of the concepts of similarity, distance and clusters that exist in both the domains of GIScience and computational linguistics, into a single model. Concretely in accordance with Takhteyev et al. (2012), we apply Waldo Tobler's First Law of Geography (Tobler 1970) to our research by assigning a higher similarity score to Tweets, which are closer in geographic space. In consequence, our combined model considers Tweets to be "similar" if the distance to each other is small in semantic space, in geographic space and in the time domain.

As for the CL dimension of the method, different degrees of similarity can be defined (Agirre et al. 2012): *completely equivalent*, as they mean the same thing; *mostly equivalent*, but some minor details differ; *roughly equivalent*, but some important information differs; *not equivalent*, but share some details; *on different topics*. However, this definition does not exactly hit the core of our interpretation of similarity because we do not look for Tweets with a similar meaning (i.e., textual content), but those containing the same emotion, for which the semantic similarity is only an incomplete proxy. Thus, our future research will include finding features that are common for one emotion, while uncommon for the other ones.

Third, a **graph is constructed** consisting of the seeds and unlabelled Tweets as nodes, as well as their respective similarity scores as edges between them. Consequently, the higher the similarity score for two Tweets is, the stronger the edge between them is. Naturally, Tweets that are not found to be similar at all have no edge between them.

Fourth, **graph-based label propagation algorithm** Modified Adsorption (MAD), as presented by Talukdar and Crammer (2009), is applied to

the graph. MAD takes a small set of labelled nodes and a large set of unlabelled ones as input along with the similarity scores for all pairs of nodes. From this initial position, the MAD algorithm iteratively propagates the labels along the strongest edges (those edges with highest similarity) to the unlabelled nodes. This results in a probability distribution over all labels for each previously unlabelled node.

We consider MAD best suited for our purpose because it improves the original Label Propagation (Zhu and Ghahramani 2002) and Adsorption (Baluja et al. 2008) algorithms in performance for small sets of seeds and preserves a more variable distribution of labels over the result than Label Propagation. Furthermore, MAD's implementation is easily usable and adaptable as it is publicly available and stores all information like nodes and edge values in text files. Concerning the choice for the family of graph-based semi-supervised learning algorithms, we consider a graph to be the native environment to combine CL and GI because the concept of graphs is well established in both disciplines. Additionally, semi-supervised learning is the only way to receive labelled output from a small set of labelled instances, i.e., the annotation process is less laborious while still guaranteeing high performance of the learning algorithm.

4.3. Data Requirements

In accordance with our research goals, we formulated several requirements for the data sets that can be processed using our proposed method. First, Tweets have to be geolocated. However, this criterion is not sufficient for our purpose because we emphasise the connection between VGI analysis and Urban Planning. Therefore, knowing the location of Tweets is not enough, but our approach also requires that many of them actually concern the location where they were created.

Second, to reduce the amount of data to be annotated, the Tweet set needs to contain as many “emotional” posts as possible. Thus, for testing our approach, we decided to use Tweets in the area of New York City (WGS84 BBOX -74.08,40.64,-73.89,40.88) during the time of the New York Fashion Week 2014 (6-13 February 2014), and in the area of Boston (WGS84 BBOX -71.21,42.29,-70.95,42.45) during the Boston Marathon event 2013 (15 April 2013). Both of these events attracted comprehensive media attention, including Twitter use, to cause emotional reactions. More, the Tweets were greatly produced in-situ (we used geo-referenced Tweets from within the cities, not those talking about the cities), thus representing these particular events well enough. Preliminary results show that we are able to detect and classify emotions in a semi-supervised manner and identify spatio-temporal

clusters with the help of CL methods. Detailed results will be presented in a separate follow-up publication.

5. Integration of Emotion Information into Urban Planning Processes

One of the main future tasks for urban and spatial planning is to deal with the impacts and issues of digital spatial data and in the context of big data. Intensified collaboration between the research fields of GIScience and urban planning is essential in the future, particularly in the area of real-world data collected by LBS. At the same time, creative urban scientific approaches have to be developed. The laboratory “urban space” is the test area for developing, adopting and perhaps discarding new planning concepts.

Ideally, the integration in the planning processes takes place during formal consideration, similar to other aspects of “public interest”. In an informal planning process, the integration of “urban emotions” can be another channel or aspect in participation processes.

Especially the planning disciplines are in danger of not being able to create adequate new visions or models to react accordingly to influences by these new developments. To enhance “traditional” Planning Support Systems (PSS) like the in the way as it is described by Batty (2014) and Geertman (2002), in which all the known methods and digital techniques are pooled like an early “mash-up” for digital planning with the developments of “wearables”, the “quantified self-movement”, human and technical sensors for near real-time data gathering opens up a new dimension in urban planning. For instance, the extraction of georeferenced emotions could be used to identify areas where the citizens’ well-being is not optimal and where urban planning actions are necessary. The scientific potential of these trends is not exploited in all its varieties and the resulting possibilities for innovative urban analysis and simulation needs to be evaluated in a wide understanding of a “Science of Cities” (Batty 2014).

All these developments urge planners to intensify their scientific efforts in the context of the ecosystem between humans, sensors, the city, and data acquisition, to gather a new quality of urban information and to detect unknown urban patterns. Without understanding these special interactions and relations between all entities, the development of a “smart city” will take place on a technical level and not on an urban planning level. Yet, this does not mean that planners should be increasingly technocratic, but that they are “lawyers” of social aspects in the smart city movement. In addition, the interdisciplinary character of planning has to be one of the main devel-

opments of self-conception urban in planning, especially in context of the knowledge society (Streich 2011). Hence, technology aspects constitute a main part of the urban planning's interdisciplinary nature. Only the combination of expertise in technology (which datasets and information can be gathered?), government (how does local administration work?) and sociology/societal skills (which are the societal impacts?) can be a firm basis to develop cities which are really "smart" and improve citizens' lives (Exner 2014). Consequently, future urban planners will need to be far more interdisciplinary in their composition than the equivalent groups producing definitions for the Web only.

6. Discussion and Conclusion

Emotion information of how people perceive their surroundings in the city can build a vital base for innovative urban planning. Thus, *Urban Emotions* provides a human-centred approach for extracting contextual emotional information from technical and human sensor data. The methodology used in our approach comprises four steps: 1) detecting emotions using wristband sensors, 2) "ground-truthing" these measurements using a smartphone-based People as Sensors LBS in near real time, 3) extracting emotion information from crowdsourced data like Twitter, and 4) correlating the measured and extracted emotions. The results are used in the domain of urban planning for decision support and the evaluation of ongoing planning processes.

The uniqueness of Urban Emotions is fourfold: First, the concept improves previous research in that it proposes a **trans-disciplinary approach** combining methods from GIScience, computational linguistics and urban sociology by merging the concepts of **semantic, geographic and temporal distance**, and semantic, geographic and temporal clusters. Second, Urban Emotions provides the first application for "**ground-truthing**" emotions in near real time in an urban context using the concept of "People as Sensors". Third, unlike other research efforts, our approach offers direct **feedback to real-world processes** in urban management and planning, and will help to detect previously unseen urban patterns. Finally, the *Urban Emotions* approach is generic so that it is **usable in other areas** like public health, traffic analysis and management, public safety, tourism, etc.

One clear limitation of the methodology for extracting emotion information from VGI is the current reliance on Twitter Tweets, assuming that Tweets are written in-situ, i.e., the posts concern the location and time at which they are published. Yet, it has been stated by Hahmann and Burkhardt (2013) that this simplification cannot be assumed to reflect reality.

Apart from this shortcoming, future research includes the optimisation of the People as Sensors user interface, the integration of other base data sources (different sensors and various crowdsourced data repositories), and the creation of a clear set of guidelines for using emotion information in urban planning.

Acknowledgements

The authors would like to express their gratitude to the German Research Foundation (DFG – Deutsche Forschungsgemeinschaft) for supporting the project “Urban Emotions”, reference number ZE 1018/1-1 and RE 3612/1-1. This research has been supported by the Klaus Tschira Stiftung gGmbH.

References

- Agirre E, Diab M, Cer D, Gonzalez-Agirre A (2012) Semeval-2012 task 6: a pilot on semantic textual similarity. In: Proceedings of the First Joint Conference on Lexical and Computational Semantics, June 7-8, Montreal, Canada:385-393
- Batty, M. (2014). Can it happen again? Planning support, Lee's Requiem and the rise of the smart cities movement. *Environment and Planning B: Planning and Design*, 41(3), 388–391
- Baluja S, Seth R, Sivakumar D, Jing Y, Yagnik J, Kumar S, Ravichandran D, Aly M (2008) Video suggestion and discovery for youtube: taking random walks through the view graph. In: Proceedings of the 17th ACM International World Wide Web Conference, 21-25 April 2008, Beijing, China:895-904
- Bollen J, Mao H, Pepe A (2011) Modeling public mood and emotion: twitter sentiment and socio-economic phenomena. In: Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media (ICWSM), 17-21 July 2011, Barcelona, Spain
- Bradley, M M, Lang, P J (1999) Affective norms for English words (ANEW): Instruction manual and affective ratings. Technical Report C-1, The Center for Research in Psychophysiology, University of Florida: 1-45
- Ekman P, Friesen, W V (1971) Constants across cultures in the face and emotion. *Journal of personality and social psychology*, 17(2):124-129
- Exner, J-P (2014) Smart Cities & Smart Planning. In: M.Schrenk et al. (eds.) Proceedings of RealCORP 2014, 21-23 May 2014, Vienna, Austria:603-610
- Gartner, G (2010) Emotional response to space as an additional concept of supporting way-finding in ubiquitous cartography. In: Mapping Different Geographies, Springer Berlin Heidelberg, Germany:67-73

- Geertman, S (2002) Participatory planning and GIS: a PSS to bridge the gap. *Environment and Planning B: Planning and Design*, 29:21–35
- Goodchild M F (2007) Citizens as sensors: the world of volunteered geography. *GeoJournal* 69(4):211–221
- Hauthal, E (2013) Detection, analysis and visualisation of georeferenced emotions. In: *Proceedings of the 26th International Cartographic Conference (ICC 2013)*, Dresden, Germany, 25-30 August 2013
- Klettner S, Schmidt M (2012) Visualisierung von Emotionen im Raum. In: *Kartographisches Denken*, Springer, Vienna, Austria:398-399
- Liu B, Zhang L (2012) A survey of opinion mining and sentiment analysis. In: *Mining Text Data*, Springer US:415-463
- Mohammad, S M (2011) Even the abstract have colour: consensus in word–colour associations. In: *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics*, 19-24 June 2011, Portland, USA:368-373
- Nold, C (2009) *Emotional Cartography - technologies of the self.*, ISBN 978-0-9557623-1-4S. <http://emotionalcartography.net/EmotionalCartography.pdf>, 2009. (14 May 2014)
- Owoputi O, O'Connor B, Dyer C, Gimpel K, Schneider N, Smith N A (2013) In: *Proceedings of the Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL HLT 2013)*, 9-15 June 2013, Atlanta, USA:380-390
- Resch, B, Britter, R, Ratti, C (2012) Live Urbanism - Towards the Senseable City and Beyond. In: *Pardalos, P, Rassia, S (Eds.) (2012) Sustainable Architectural Design: Impacts on Health*:175-184
- Resch B (2013) People as Sensors and Collective Sensing - contextual observations complementing geo-sensor network measurements. In: *Progress in Location-Based Services*, Springer Berlin Heidelberg:391-406
- Roberts K, Roach M A, Johnson J, Guthrie J, Harabagiu S M (2012) EmpaTweet: annotating and detecting emotions on twitter. *Proceedings of the International Conference on Language Resources and Evaluation (LREC)*, 21-27 May 2012, Istanbul, Turkey:3806-3813
- Sagl G, Resch B, Hawelka B, Beinat E (2012) From social sensor data to collective human behaviour patterns: analysing and visualising spatio-temporal dynamics in urban environments. In: *GI-Forum 2012: Geovisualization, Society and Learning*, Wichmann Verlag, Berlin:54-63
- Salesses P, Schechtner K, Hidalgo C A (2013) The collaborative image of the city: mapping the inequality of urban perception. *PLoS ONE*, 8(7): e68400

- Scott B, Neil T (2009) Designing web interfaces: principles and patterns for rich interactions. Sebastopol, USA: O'Reilly Media
- Streich, B (2011) Stadtplanung in der Wissensgesellschaft – Ein Handbuch (2. Edition), Springer Berlin Heidelberg
- Takhteyev Y, Gruzd A, Wellman B (2012) Geography of Twitter networks. *Social Networks*, 34(1):73-81
- Talukdar P P, Crammer, K (2009) New regularized algorithms for transductive learning. In: *Machine Learning and Knowledge Discovery in Databases*, Springer Berlin Heidelberg: 442-457
- Tobler, W R (1970) A computer movie simulating urban growth in the detroit region. *economic geography*:234-240
- Zeile, P, Resch, B, Exner, J-P, Sagl, G, Summa, A (2014) Urban Emotions - Kontextuelle Emotionsinformationen für die Räumliche Planung auf Basis von Echtzeit-Humansensorik und Crowdsourcing-Ansätzen. In: Strobl, J, Blaschke, T, Griesebner, G (Eds.) (2014) *Angewandte Geoinformatik 2014*, Wichmann Verlag, Heidelberg:664-669
- Zeile, P, Höffken, S, Papastefanou, G (2009) Mapping people? – The measurement of physiological data in city areas and the potential benefit for urban planning. In: *Proceedings of RealCORP 2009*, 22-25 April 2009, Vienna, Austria:341-352
- Zhu X, Ghahramani Z (2002) Learning from labeled and unlabeled data with label propagation. Technical Report CMU-CALD-02-107, Carnegie Mellon University, Pittsburgh, USA