An Emotional Compass

Harvesting Geo-located Emotional States from User Generated Content on Social Networks and Using them to Create a Novel Experience of Cities.

Salvatore Iaconesi, Oriana Persico

ISIA Design, Florence, Italy
salvatore.iaconesi@artisopensource.net
oriana.persico@gmail.com

Abstract. This paper describes the design and implementation of a novel mobile interface under the form of an emotional compass. The interface has been created using the possibility to access large number of geo-located user generated contents across multiple social networks, which have then been processed using Natural Language Analysis techniques, to understand the emotions expressed by Internet users in urban environments. This possibility gives rise to the innovative interface for urban navigation presented in this paper, and opens up the opportunity to conceptualize and implement novel forms of services.

Keywords: User experience, urban navigation, social media, social networks, user generated content, location-based services, urban sensing, ubiquitous technologies, innovative interfaces, information visualization, emotional interfaces, natural interaction.

1 Introduction

"The map is not the territory." [18]

"The map is not the thing mapped." [4]

"The tale is the map that is the territory." [13]

"We say the map is different from the territory. But what is the territory? The territory never gets in at all. […] Always, the process of representation will filter it out so that the mental world is only maps of maps, ad infinitum." [3]

When we experience territories, we create stories. We model these stories using mental maps. These maps have seldom anything to do with what actually lies within the territories themselves. A mental map refers to one person's point of view perception of their own world, and is influenced by that person's culture, background, mood and emotional state, instantaneous goals and objectives. If we move along the streets

of my city in a rush, trying to find a certain type of shop or building, our experience will be different than the one we would have had if we were searching for something else.

Focus will change. We will see certain things and not notice other ones which we would have noticed otherwise. Some things we will notice because they are familiar, common, or because associate them to memories and narratives. Some will stand out because they react with some element of our culture or background. All this process continuously goes on as our feelings, emotions, objectives and daily activities change, creating the tactics according to which we traverse places and spaces, to do the things we do.

In the density of cities, this process happens for potentially millions of people at the same time. In his "the Image of the City" [20], Lynch described cities as complex time-based media, symphonies produced by millions of people at the same time in their polyphonic way of acting, moving, interpreting, perceiving and transforming the ambient around themselves: a massive, emergent, real-time, dissonant and randomly harmonic, work of time-based art with millions of authors that change all the time.

In this, our mental maps – the personal representations of the city which we build in our minds to navigate them to fulfill our needs and desires – live a complex life as our perception joins into the great performance of the city.

Dissonance is the essence of the city itself, and represents its complexity, density and opportunities for interaction.

Harmony represents affordances, the things which are recognized and shared by different cultures. Those elements of the perceptive landscape onto which we can agree upon, which we recognize and attribute compatible meanings, allowing us to collaborate, meet, do things together. For example, Haken and Portugali [15] have suggested a broad definition of landmarks to refer to any distinguished city elements that shape our mental map. Or as Appleyard [1], Golledge and Spector [14] who have conducted studies about the imageability of urban elements not because of their visual stimulus but because they possess some personal, historical, or cultural meaning.

These features found within our mental maps enable the possibility to design the affordances of places and spaces. We can use the understanding of what is consistently recognized and understood to design the elements of space/time which will be able to describe to people what is allowed or prohibited, suggested or advised against, possible or imaginable. Lynch's concepts of legibility and imageability are closely related to James J. Gibson's notion of affordances developed in his direct perception theory, according to which the objects of the environment can afford different activities to various individuals and contexts. And, again, in Haken and Portugali [15], all elements of a city afford remembering, as they shape in the mental maps in human minds.

In a further step in the direction of citizen activation, we can also imagine to make this type of understanding widely known and usable, to enable people to express

themselves (and their mental maps of how they perceive the world) more effectively and powerfully.

These possibilistic scenarios have become radically viable with the widespread of ubiquitous technologies. Nomadic devices (such as smartphones) and their applications we are able to merge our physical understanding of the world with the digital one: the two have, in fact, become so interweaved and interconnected as to form a new physicality, visuality and tactility which shape our everyday experiences of the world.

According to Mitchell's "City of Bits" [22], McCullough's Digital Ground [21], Zook's and Graham's DigiPlace [33] we are constantly immersed in emergent networks of interconnected data, information and knowledge which is produced by millions of different sources and subjects in the course of their daily lives. This data and information radically shapes the ways in which we have learned to work, learn, collaborate, relate, consume and perceive our environment.

If we are strolling in a park and we receive a notification of some sort on our smartphone, the natural environment could instantly transform into an ubiquitous, temporary office. If we want to make a decision about a certain thing we would like to purchase while in a shop, a quick look online will help define our opinion in ways that can be very powerful. If we receive a message on our smartphone, our mood could change for the rest of the day.

Situated and ubiquitous information is able to powerfully transform, in real-time, the ways in which we experience places, objects and services, by providing the wide accessibility of other people's stories, emotions, expectations and visions.

This scenario is the one we have tried to address in our research: the conceptualization, design and implementation of a tool for urban navigation, in which the emotional, narratives expressed by people while inhabiting and using urban places, spaces and objects become instantly and radically available, accessible and usable. We used this approach to define a novel vision on the opportunity to design new types of affordances for our cities.

We have decided to start from the idea of a Compass.

2 The Compass

The compass is an historically understood, ubiquitously known object dedicated to navigation and orientation: it finds the direction in which one wants to go. It usually is a navigational instrument that shows a directions in a frame of reference that is stationary relative to the face of the earth.

The frame of reference defines the a number of cardinal directions, with a rotating indicator, pointing out the direction towards which the user is facing.

They are very easy to use (or, at least, to understand how they work) and are capable of providing direct, immediately accessible insights about the information they convey.

Different cultures and civilizations have used compasses for very different reasons, such as in the case of the Qibla compass, which is used by Muslims to show the direction to Mecca for prayers, or the Feng Shui compass, through which one is able to understand how to better orient houses' furniture and elements to obtain optimal energies.

The Feng Shui example is of particular relevance for the objectives our research. In its construction, the cardinal points are matched with an overwhelming amount of other information: over 40 concentric circles of writing and detail used to define the Bagua of your home, the ways in which energy flows. In the Feng Shui compass, the cardinal directions are combined with information coming from entirely different domains, and this combination gives rise to a completely different concept of orientation.

This is the idea that we wanted to explore in our research.

Is it possible to use the ubiquitous infoscape (the informational landscape) which is constantly produced by human beings on social networks to design novel forms of urban navigation? Novel ways of experiencing places? New ways for making decisions, for relating to one another, for consuming, for expressing and understanding emotions?

We wanted to design a new type of compass, and we wanted to use it as a way to also design a methodology to explore, conceptualize and implement new forms of orientation.

We started from the idea of emotions.

How is an emotional compass made?

How do you create a compass which harvests in real-time as much data, information and knowledge as possible about the ways in which human beings express their emotions on social networks? How to use this information to orient users to have insightful emotional experiences?

Is it possible to identify “emotional landmarks” – those places/spaces where, at a specific or recurring time, a certain emotion is expressed powerfully and abundantly – ?

If they do exist: do emotional landmarks change over time? Do they change according to the culture you are observing? To language? To the time of day, week, month or year? To the specific topic your compass is observing?

Is it possible to design a methodology that can be used to allow for the creation of urban sensing compasses of this kind? To allow for the creation of services and objects which can sense emotional expressions on social networks on multiple topics, to create a sense of emotional direction? Could this be done for emergency scenarios

such as revolts, riots or natural disasters? For art and tourism? For city planning and safety? For entertainment and consumption?

These, among many others, were the main questions which we asked ourselves in our research.

We started from trying to validate a a simple idea: an emotional compass, using public, real-time user generated information on social networks to tell us in which direction to go to arrive where people express more powerfully the basic emotion we selected.

3 Previous Work

Abundant work exist which explores the idea of emotionally mapping cities and to propose forms of navigation that go beyond classical way-finding.

For example, Christian Nold's fundamental work on Biomapping [23] and Emotional Cartography [24], which is a set of methodologies and tools for visualising people's reactions to the external world. In the project, a rather large number of people (about 2000) have taken part in community mapping projects in over 25 cities across the globe. In structured workshops, participants re-explore their local area with the use of a device which records the wearer's Galvanic Skin Response (GSR), which is a simple indicator of emotional arousal, in conjunction with their geographical location. In this way, a map is created which visualises points of high and low arousal. Nold's work can be considered to be a seminal one in exploring how devices can capture location-based emotional states, and make them accessible through maps and other means. In our research we wanted to focus more on more complex possibilities to interpret human emotions, coming from the usage of language, and on the possibility to not only record emotions, but to turn them into active, searchable, usable, knowledge which anyone could generate and access.

Another example, the Fuehlometer ('feel-o-meter') [31], was produced by german artists Wilhelmer, Von Bismarck, and Maus in the form of a public face, an interactive art installation that reflects the mood of the city via a large smiley face sculpture. It was installed atop a lighthouse in Lindau, Germany. A digital camera along the lake captured the faces of passersby, which were then analyzed by a computer program and classified as either happy, sad, or indifferent. The cumulative results determine the expression of the sculpture, whose mouth and eyes shift accordingly via a system of automated motors. Von Bismark's thoughts on the artwork are particularly interesting in this case: “we wanted people to start considering if they want people to read their emotions, and if they want to know others’ emotions; if they want to be private or they want to be public. That’s what it comes to in the end—what is private, and what is public?” The artwork itself provided us with precious guidelines about what we set forth to achieve: an immediately readable and understandable service. Yet the techniques it used proved to be very limited in terms of the possibility for interpretation of human emotions, and for the production of usable knowledge out of them,

including considerations on people's cultures, behaviors and relations in their interactions in the city.

Using a different approach, the City of Vilnius [8] has found a way to track emotions on its territory using a social tool that gauges the average residents' level of happiness. Residents submit their overall level of happiness for each given day using their smartphones, or by scanning a barcode on the post advertising the initiative dubbed the "Happiness Barometer." Votes are later totaled to determine the overall happiness level of the town – displayed on a large urban screen and on the website. While this is a potentially interesting strategy to use by a public administration to both capture and communicate the emotional states of its citizens, it also constitutes a very limited approach, for its hackability and extreme synthesis, reducing "happiness" to a single percentage.

Another example comes from an artwork titled Consciousness of Streams [2]. In the work the artists have set up a series of devices or installations in several cities, as well as making available a web application that could be used from anywhere. Using these devices and applications users were able to contribute their geographic location, emotional state, as well as an image of their face or sound recording. The resulting information is constantly visible online and do not come under the form of statistics or maps, but, rather, under the form of a "real-time interconnected emotional map of the planet" [16] showing a novel geography whose objective is not to show buildings, streets and other elements of the landscape, but a topography of human emotions, showing adjacencies, proximities and distances which are not physical, but emotional.

Another relevant project is Mappiness [19], part of a research project at the London School of Economics. This mobile app and online system actively notifies users once a day, asking how they're feeling. The data gets sent back along with users' approximate geographical location and a noise-level measure, as recorded from the phone's microphone. In this way users can learn interesting information about their emotions – which they see charted inside the application – and the operator can learn more about the ways in which people's happiness is affected by their local environment — air pollution, noise, green spaces, and so on. This is an interesting mechanism, but also one that lacks the possibility to sense the natural emergence of emotions as linked to urban daily life, in people's language and expressions, as it relegates users' interactions to strictly encoded forms.

An interesting project is "Testing, Testing!" [9], an experiment developed by Colin Ellard and Charles Montgomery, and conducted in New York, Berlin, and Mumbai. By inviting participants to walk through the urban terrain, and measuring the effects of environment on their bodies and minds, Ellard aimed to collect data in real, living urban environments. That data would then be available for application within urban planning and design to enhance urban comfort, increase functionality, and keep city dwellers' stress to acceptable levels.

The last project which we wish to highlight among the many that we have analyzed during our research, is the Aleph of Emotions [29], an experimental art project by Mithru Vigneshwara: a camera-like interface allows users to point along a particular

direction, focus to a place along that direction, and click to view a representation of emotions in that place. The intention is to explore and find patterns in human emotions with relation to space and time. Data is collected based on keywords that define certain emotions. The results are finally presented with an interactive object. A custom software was then written to collect tweets that contain these keywords. We felt, to a certain degree, this project to be really close to what we wanted to achieve. The major limitations which we have identified in its conception lie in the impossibility to comprehend human emotions in significant ways (as described in the following sections, keyword misuse, irony, the possibility to handle multiple languages and the lack of context awareness all constitute enormous limits if one's objective is to gain a level of understanding that is as wide as possible), and in the lacking sense of immersion in the information landscape: when you hold the device, it does not show information around you, but the one coming from distant cities.

4 Concept and Methodology

Our goal was to create:

- a compass (i.e.: a tool that shows the user a direction in a set frame, as simple to use as a compass) available on a smartphone showing the direction(s) toward the emotion selected by the user;
- a system, to be used by the compass to provide the direction, which is able to harvest all user generated content in the area on major social networks, and process them using Natural Language Analysis and geo-referencing techniques to infer the emotion (if any) expressed by the messages and their geographical location.

For this we broke down the activity into different domains:

- the system to harvest messages from major social networks in real time;
- the geo-referencing/geo-coding techniques;
- the Natural Language Processing techniques;
- interface design and interactive information visualization.

4.1 A System to Harvest Messages from Major Social Networks in Real Time.

There are many different techniques and technologies using which a system of this kind can be implemented.

The main issues we were faced with during the design and implementation process were the following:

- legal issues;
- technical issues.

Starting from the legal issues: users and developers wishing to use the features of major social networks have to abide to the rules dictated in the providers' Terms of Service (ToS). These are very complex legal documents which state what can be done and what is prohibited, also establishing various forms of liability for all parts involved. Different providers have different ToS, which can vary substantially in describing the ways in which you can and cannot use the information generated on their services.

In the specifics, most of our focus was oriented towards the ToS documents offered for developers.

Most social networks offer Application Programming Interfaces (API) of some sort, which developers can use to build their own applications by interacting with the social network's ecosystem (users, communities, content, etcetera).

These APIs offer an incredible opportunity for service designers and developers, as they permit accessing a vast amount of data about people's expressions and positions, the topics they discuss and the relations which they maintain, allowing for the creation of a variety of useful services.

APIs usage is constrained by the ToS, which limits the degree to which any developer or company is able to capture, process, use and visualize information coming from social network operators.

Limits are mainly imposed on:

- ownership of the data, which is hybrid to various degrees, as some operators choose to leave it to the users, but claim the exclusive right to use it, while others claim ownership of the data, with the opportunity for users to “have it back”;
- number of queries, which are also limited, and cannot go beyond set quantities or sizes per day; costly agreements are almost always possible to overcome these kinds of limitations;
- storage of the information, which is generally forbidden apart from the users, who are granted the right to store their own data;
- processing of the harvested information, according to which different limitations exist in the possibility to form aggregates, statistics and elaborations on the data coming from social networks; these kinds of limitations are usually really difficult to define and enforce, as is the act of actually recognizing a developer breaking them (how can you know what I've done with the data after I extracted it using the APIs?);
- visualization and branding, according to which the single information element (for example, a Twit) must be shown on any interface it appears on according to a precise set of guidelines, which include the visibility of the branding, the functionalities associated with it (i.e.: the retweet button).

These legal limits are different across different providers and also change quite frequently and arbitrarily, forcing companies and researchers to constantly adapt and maintain their applications: if your application is perfectly ToS compliant and working today, it might not be so tomorrow.

On top of that it must be said that the issue of expectation for publicness is also a large presence for what concerns the legal side of things. Just as it happens when we go to malls and shopping centers, we perceive them to be public spaces and, thus, we conform to what we have learned to be our rights and acceptable behaviors in public spaces. But this is not the case as different sets of rules apply in these spaces affecting anything from privacy, freedom of expression and basic rights. This is a issue which is rising in importance and relevance, and also in the awareness of people and organizations, and it is too broad to cover here. Yet it must be said that we have often clashed with it, for example in trying to harvest all the user expressions on their feelings towards public policies enacted by governments and administrations.

That said, in our research we have had to access the constant consultancy of a group of specialized lawyers to understand what we were allowed to do with the information harvested from the different social network operators, and we have managed to design a replicable model which includes clusters of technical rules which transform the legal specifications into technical and technological ones, and which we have been able to successfully use in these kinds of scenarios over the past three years.

Some limitations exist on the purely technical side, too.

In the first instance, the APIs allow for limited degrees of freedom in the querying and interaction with the databases of operators: not all of the information is made available and limitations on how developers are able to formulate the queries also exist.

Furthermore APIs frequently change, forcing development teams to constantly maintain and adapt the source code of the applications.

Once in a while, entire sets of features and possibilities disappear or change in form or availability, forcing designers and developers to go back to the drawing board and re-think or re-frame their services.

It can be said that the ideas of access and of interoperability are currently not among the priorities of social networking service providers.

City	From Date	To Date	N. of UGC
London	Jan. 1st 2011	Feb. 1st 2011	5143500
Rome	Oct 15th 2011	Oct. 16th 2011	91538

City	From Date	To Date	N. of UGC
Turin	Aug. 1st 2011	Sept. 20th 2011	240982
Berlin	Jan. 4th 2012	Jan. 20th 2012	1699240
Hong Kong	May 1st 2012	Jul. 1st 2012	5732487
Cairo	Jul. 27th 2013	Sept. 2nd 2013	3466388

Table 1: Number of UGC harvested from social networks in different experiments.

We resolved most of these issues adopting a radically modular approach. Interoperable connectors have been designed and created to take into account the different scenarios with the different operators, and to abstract the main service logic from their implementation details. And providing us with the possibility to limit the damages whenever ToS or regulations changed on the operators' side.

This part of the activity has revealed to be a truly fundamental one, as we have actually developed a service layer which implements an easily maintainable abstraction and interoperability among different social network providers, and we're thinking to dedicating to it a separate research effort, to design the ways in which it could be offered as a service or as a novel source of real-time Open Data.

We have been successfully able to use these technologies and techniques for some time now, and, at the time of writing we have been able to perform several experiments whose results, in terms of the number of captured User Generated Contents (UGC) over time, are listed in Table 1.

The systems we used in some of the experiments were quite small (for example standard desktop computers): by scaling up the infrastructures we have verified that the number of messages captured increases dramatically. For example, in a small local experiment in Rome, we used a setup in which Twitter queries fired off each 5 seconds, to limit the load of the computer used, ad to adhere to the limits in API usage which were allowed in the free option. When we augmented the rate (to 1 query per second) and made a commercial agreement with Twitter, we were able to dramatically increase the number of harvested messages: from about 9000 per hour, to about 20000 per hour.

4.2 Geo-referencing/Geo-coding Techniques and Named Places

A number of different possibilities exist in trying to attribute a geographical context to UGC:

- users employ the features offered by social networks for geo-referencing their own messages (either using the GPS on their smartphone, or providing additional information);

- users include in the message information which can lead to finding out a location that they are talking from or about;
- users may use none of the previous possibilities, but include an indication of their geographical position (either current or by default) in their profiles;
- users do none of the above: in this case it is not possible to gather the user's location.

The third case has a low level of reliability. For a number of reasons, users may lie about their current or "home" location. For example, they commonly choose their favorite city, or a 'cool' city, or a totally fictional location: on the popular social network Foursquare we currently reside in Mordor (taken from Tolkien's "Lord of the Rings"), which we have placed, using the standard features offered by the system, a few meters away from our lab.

For these reasons, in our research we tend not to use these kinds of location specification (the "home" location or the current location as specified in the user's profile).

The first case is also very easy to deal with: a geographical location (often paired with extensive sets of meta-data, such as in the case of Facebook and Foursquare) is explicitly provided in the message, and thus we are able to use it.

From the analysis of the results of our experiments, it turns out that the geolocation features offered by social networks are not very commonly used. This varies from service to service, from region to region, and across contexts. But it is really not used a lot. From what we have seen in our experiments, the most common user behavior is to either turn on the location sharing features when they download the applications to their smartphone, or forget about it.

From what we have been able to understand, the most location aware social networks are Foursquare and Instagram, with respectively 92% and 30% of the messages which have a location attached to them. Then comes Twitter, with 10–15%, according to time and context. Then Facebook: if we exclude the posts related to events (which have a location attached to them), the percentage drops to about 4%, and comes almost completely from messages generated using the mobile applications. These results are based on the messages we have collected over time in our experiments, and vary a lot across time and context. For example, many more messages with a location are generated on holidays and in times of vacation, and in the case of special events, such as the riots and revolts in Cairo, Egypt, during 2013. In this last case, for example, Twitter messages with a location specified rises up to as much as 18%.

The second case in the list are more complex and interesting. They take place when users do not use the platforms' features to include their location in the message, but, rather, mention the location which they're talking from or about in the text of the message itself.

First of all, it is important to try to understand whether the mention of a geographical location in a message is indicating that the message was produced in that location,

or if it was talking about it: according to the service which one wants to implement, these two possibilities may completely change the relevancy of the message.

That said, we have tried to formulate a working procedure with which to try and add location information to these kinds of messages.

We:

- built databases of Named Places for the various cities, including landmarks, street names, venues, restaurants, bars, shopping centers, and more, by combining the information coming from
  - publicly available data sets such as the ones available accompanying public cartography sets (for example for Italy we have used the named places provided by ISTAT, Italy's National Statistics Institute available at [17]);
  - the list of named places contained in the OpenStreetMap databases, for example as described in [25] and [26];
  - the list of named places provided by social networks themselves, which allow using their APIs to discover the locations used by users in writing their messages, for example on Facebook [10] or Foursquare [12];
  - lists of relevant words and phrases, such as event names or landmarks;
- used the text representation in various forms of the named places in a series of phrase templates to try to understand it the user writing the message was in the place, going to the place, leaving the place, or talking about the place;
  - for example, the template "*going to [named place]*" would identify the action of going, while "*never been in [named place]*" would identify the action of talking about a place;
  - templates have currently been composed in 29 different languages, for a total of more than 20000 different templates;
- each template was assigned a degree of confidence, evaluating the level of certainty according to which the sentence could be said to identify the intended information;
  - for example: "I'm going to [named place]" has a relevance of 1 (100%), while the "[named place]" taken by itself has a relevance of .2 (20%) as it might be a false match (imagine a bar with the same name of a famous landmark, for example);
  - a threshold was established; if the sum of relevance degrees for templates matched to sentences was above the threshold, the information about content location was kept, else it was thrown away. Currently the threshold we use for this is of 90% .

In the application we have, thus, chosen to gather geo-location information through explicit use of the location-based features of the services and, should they not have been provided, by combining them the results of the named places analysis.

By applying these rules we have been able to bring up the percentages of geo-located messages quite successfully. For example on Twitter , we brought it up from 15% to around 27% (judging from a series of sample-based statistics we produced on the messages that had not been directly geo-coded by the user). To our knowledge, the rest of the messages were not dealing with a specific location, or were not intended to deal with it.

Natural Language Processing and Artificial Intelligence to recognize emotions and topics in text

There is an extensive amount of research about the possibility to automatically interpret text to understand the emotion expressed by the writer, either on social networks or on more general texts.

We approached the possibility to recognize emotions by identifying in text the co-occurrence of words or symbols that have explicit a□ective meaning. As suggested in by Ortony et al. [27] we must separate the ways in which we handle words that directly refer to emotional states (e.g.: fear, joy) from the ones which only indirectly reference them, based on the context (e.g.: "killer" can refer to an assassin or to a "killer application"): each has different ways and metrics for evaluation.

For this, we have used the classification found in the WordNet [11] extension called WordNet Affect [32].

The approach we used was based on the implementation of a variation of the Latent Semantic Analysis (LSA). LSA yields a vector space model that allows for a homogeneous representation (and hence comparison) of words, word sets, sentences and texts. According to Berry [6], each document can be represented in the LSA space by summing up the normalized LSA vectors of all the terms contained in it. Thus a synset in WordNet (and even all the words labeled with a particular emotion) can also be represented in this way. In this space an emotion can be represented at least in three ways: (i) the vector of the specific word denoting the emotion (e.g. “anger”), (ii) the vector representing the synset of the emotion (e.g. {anger, choler, ire}), and (iii) the vector of all the words in the synsets labeled with the emotion.

This procedure is well-documented and used, for example in the way shown in [28], which we adopted for the details of the technique.

We adapted the technique found in [31] to handle multiple languages by using the meta-data provided by social networks to understand in which language messages were written in (and performing a best-effort analysis on those cases in which the meta-data seemed to be wrong due to the high number of non-existing words in a certain language), and using a mixture of the widely available WordNet translations and some which we produced during the research for specific use cases.

An annotation system was created on the databases to tag texts with the relevant emotions (as, within the same message, multiple emotions can be expressed).

We also tried to deal with the wide presence of irony, jokes and other forms of literary expression which are difficult to interpret automatically. To do this, we have followed the suggestions described in [7] and [5] with varying results.

Interface Design and Interactive Information Visualization

Given the intensive preparation phase, the information was, at this point, ready to be visualized and the interaction designed. We chose a very minimal layout, to allow the user to focus on the interaction mechanism, providing little-to-none additional detail beyond the emotional compass.

The interface development followed a two-phase sequence. First was designed a rough interface to understand the accessibility and usability of this kind of tool. The design was created in occasion of our Rome based tests, following a city wide riot which had happened the previous year, and of which we had been able to capture the social network activity.

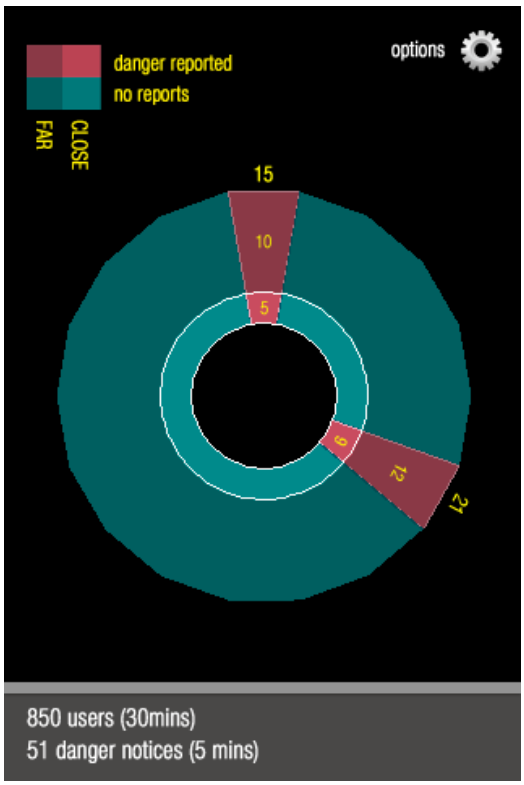

Figure 1: The first interface: a riot in Rome.

In this first scenario a mobile application was designed that would poll the database for new updates, which came under the form of a list of basic emotions and their intensity in the various directions, relative to the user's current position.

The information was drawn on screen using a radial diagram using basic trigonometry, while the on-board magnetic compass and accelerometer controlled the diagram's rotation, to keep track of the user's heading and the device orientation (see Figure 1).

The focus in this interface was to highlight the potentially dangerous scenarios, so that users would be able to avoid going in their directions. For this the default setup was pre-configured highlighting emotions of fear and grief, followed by anger and sadness. The user was able to use the settings button on the interface to choose from a drop-down (a scroll-wheel, on most smartphones) to choose from the other available emotions, so that the experience and goal of the experience could be personalized. The second iteration of the interface was more general purpose (Figure 2).

In this new form, a the color coded emotions would surround the white center, radially indicating the intensities of the emotions as they emerged around the user.

The result was a multi-compass, with each color showing an emotion, its thickness around the center indicating its intensity in the relative direction. In the picture, the color purple, indicating boredom, is thicker in the upper right and lower left, showing that the emotion has been recently manifested on social networks to the front-right of the user, and to his back-left.

A pull-up menu can be dragged up by the user to toggle on/off the various layers, also obtaining a visual legend for the meaning of the colors. From the same menu, cursor sliders can be used to configure the sensibility of the emotional compass: in distance, from 100 meters to 1 kilometer (e.g.: if you choose 500 meters, only the emotions generated within a 500 meters radius will be taken into account); and in time, from 5 minutes to 1 month (e.g.: if you choose 2 days, only the emotions expressed during the past 2 days will be used).

The transformation of the emotional color blobs around the center take place using smooth, interpolated transitions, both to give the user a clear vision of what is changing, and to achieve a “blobby”, organic look, which is able to visually communicate a situation in constant evolution.

Whenever the user reaches a location in which a certain emotion has recently been expressed with particular strength, the background starts pulsating in the color of the corresponding emotion: an emotional landmark has been reached.

5 User Experience of the Artwork

The artwork is currently available as a prototype application for iOS and Android smartphones. It will be available on major stores as soon as the final beta-testing stag-

es are complete (estimated late January 2014) and interested parties can request beta access by contacting the authors.

Throughout the interface design process we performed regular walks in the city which we observed on social networks to better understand how the application would transform our perception of the city.

The experience itself can be compared to the one of Rhabdomancy. While walking amidst the spaces of the city while using the compass, the ordinary way-finding reference items become less important. The color-coded intensity indicators for the various emotions provide the sensation of being able to access a geiger counter, or some sort of field intensity measurement device, showing the directions in which a certain emotion is stronger.

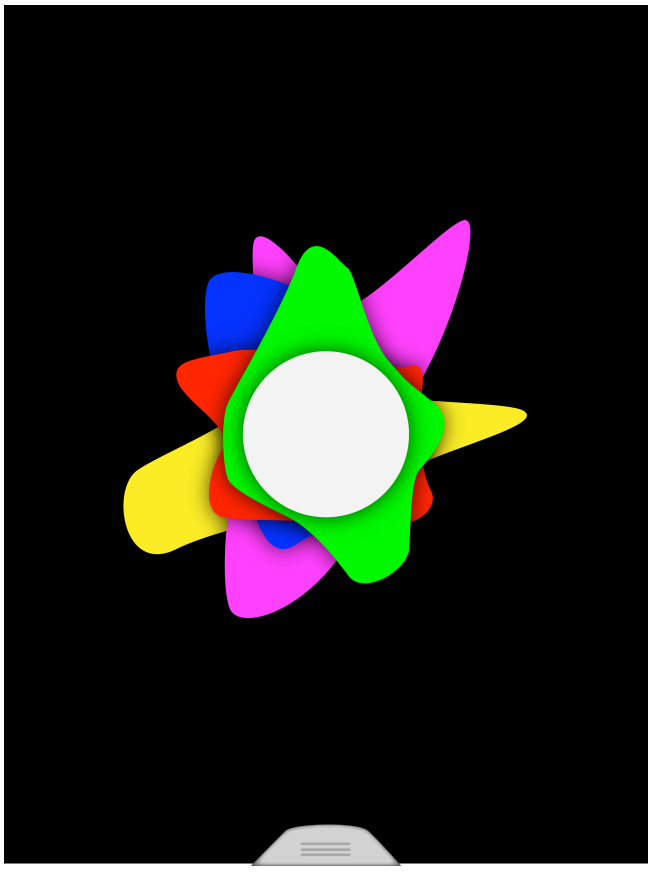

Figure 2: the second interface.

The impossibility to access street and topography based directions, for example, is strange at times sometimes. On the other hand, it gives the exact perception of being able to access a different kind of geography: one that is based on the intensity of emo-

tions in a certain place, rather than its name or street number. It is definitely the perception of an energy field, of a radiation. As an example, while following the peak level of a certain emotion, we were faced with a wall, or a building or block that was standing in our way. In this kind of situation, the system did not provide any clue about the fact that the peak itself was to be found inside the obstacle (for example in the building) or beyond it. As we tried to go around the building we would be able to gain better understanding: if the peak reversed its direction once we were around it, it would clearly mean that the peak emotion was inside the building; if it kept on pointing in the same direction, it meant that the peak intensity was beyond the obstacle.

A similar effect could be achieved by acting on the slider which regulates the sensibility in terms of distance. Once faced with an obstacle it was possible to act on the slider to lower the senseable distance. By doing this it was sufficiently clear that if the peak disappeared at when the slider was lowered to the point of being nearer that the obstacle's perpendicular thickness, it would mean that the emotional peak was to be found within it.

Identifying emotional peaks in closed spaces proved to be quite a challenge: the lack of GPS coverage in closed spaces allows to easily identify the buildings in which a certain emotional peak can be found, but not to continue to search within them.

While using the application has proven to be somewhat hard to follow multiple emotions at the same time: with the different peak indicators all being independent, it has come out to be much easier just to follow one main emotion, and to eventually check the other emotional levels once arrived at a certain location

The addition of sounds has also proven to be extremely useful. A different drone-based sound loop of specific tones and texture was associated to each basic emotion, and its volume was connected to the instantaneous intensity of the emotion at the current user location. By wearing headphones users gets a really accurate sense of the com-presence of the emotions in the place they are currently in, also being able to momentarily switch off the various emotions/tones to associate each tone to the relative emotion. Creating sounds which have a drone-like, constant tone, but with evolving texture has been proven to give the best effects: users can create a generative song by walking around, depending on how social networks users expressed in that location.

Also, the pairing of the sounds with the indicators, with specific focus on the color-coded on-screen alert which appears when an emotional peak is reached, has proven to be really effective, with the alert matching the maximum volume of the relative sound: when users heard these kinds of high volumes, they consistently checked the application display to see if the alert appeared. This also allowed users to use the compass from their pockets, navigating the city by following volume augmentations, and pulling the smartphone out only when the volume would be high, to check the visual confirmation that the emotional peak had been reached.

6 Conclusions

We have found this research path to be rewarding for its implications in terms of the possible services that could be designed by using the proposed methodology, and of the possibility to observe and experience urban environments in truly innovative ways. We can imagine services highlighting the sense of security, of enjoyment or satisfaction, with enormous potentials for tourism, real-estate, entertainment, events and for public administrations wishing to discover and expose the ways in which people feel in the city.

On the other side, using these kinds of techniques, we are now able to understand cities better, in how people live their daily lives across cultures, languages, occupations and interests. For example, by simply filtering the meta data about language, we would be able to know the emotions of people in the city coming from different countries and cultures. We could see how they move around the city, we could compare them and the emotions they express, finding the ways in which they feel the same, or differently, at the different times of the days and weeks. We could use this information to better understand our cities, providing ways to empower multicultural ecosystems to form in more harmonious ways. The concept of the emotional landmark has proven to be very interesting. Which are the places in which different cultures more powerfully express a certain emotion, in different times of the day? How can we use this information? How can we design a city for emotions? These and more will be the questions which we will try to answer in the next phases of our research, together with the idea of opening up the process, promoting the accessibility and interoperability of this novel source of real-time, emergent Open Data that we have helped to shape: publicly expressed human emotions.

7 References

1. Appleyard, D. Why buildings are known, Environment and Behavior 1 (1969), 131–156.
2. Art is Open Source, Consciousness of Streams. http://www.artisopensource.net/2011/01/12/cos-consciousness-of-streams/ (last accessed 01/09/2013) and http://cos.artisopensource.net/ (last accessed 01/09/2013)
3. Bateson, G. Form, Substance and Difference in Steps to an Ecology of Mind, Chandler Publishing Company (1972).
4. Bell, E.T. in Numerology: The Magic of Numbers, Williams and Wilkins (1933).
5. Bermingham, A., Smeaton, A.F. Classifying sentiment in microblogs: is brevity an advantage? in CIKM '10 Proceedings of the 19th ACM international conference on Information and knowledge management (2010), 1833-1836.
6. Berry, M. Large-scale sparse singular value computations. International Journal of Supercomputer Applications, 6(1) 1992, 13–49.
7. Carvalho, P., Sarmento, L., Silva, M.J., de Oliveira, E. Clues for detecting irony in user-generated contents: oh...!! it's "so easy" ;-) in TSA '09 Proceedings, CIKM workshop on Topic-sentiment analysis for mass opinion (2009), 53–56.
8. City of Vilnius, the Happy Barometer. http://happybarometer.com/ (last accessed 01/09/2013)

9. Ellard, C. Testing, Testing! http://www.bmwguggenheimlab.org/where-is-the-lab/mumbai-lab/mumbai-lab-city-projects/testing-testing-mumbai (last accessed 01/09/2013)
10. Facebook, obtaining a list of places in a geographical area. https://developers.facebook.com/docs/reference/api/search/ (last accessed 01/09/2013)
11. Fellbaum, C. WordNet. An Electronic Lexical Database. The MIT Press (1998).
12. Foursquare, obtaining a list of places in a geographical area. https://developer.foursquare.com/docs/venues/search (last accessed 01/09/2013)
13. Gaiman, N. in Fragile Things, HarperCollins (2006).
14. Golledge, R.J., Spector A. Comprehending the urban environment: Theory and practice, Geographical Analysis, 10 (1978), 403–426.
15. Haken, H., Portugali, J. The face of the city is its information, Journal of Environmental Psychology, 23(4) (2003), 385-408.
16. Iaconesi, S., Persico, O. ConnectiCity: Real-Time Observation and Interaction for Cities Using Information Harvested from Social Networks in International Journal of Art, Culture and Design Technologies (IJACDT) Volume 2, Issue 2 (2012), 14–29.
17. ISTAT, Territorial Data Sets, including named places. http://sitis.istat.it/sitis/html/ (last accessed 01/09/2013), http://www.istat.it/it/prodotti/banche-dati (last accessed 01/09/2013), http://www.istat.it/it/archivio/44523 (last accessed 01/09/2013)
18. Korzybski, A. A Non-Aristotelian System and its Necessity for Rigour in Mathematics and Physics in American Association for the Advancement of Science 1931, Conference Proceedings. Reprinted in Science and Sanity (1933), 747–61.
19. London School of Economics, Mappiness. http://www.mappiness.org.uk/ (last accessed 01/09/2013)
20. Lynch, K. The Image of the City, MIT Press (1960).
21. McCullough, M. Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing, MIT Press (2005).
22. Mitchell, W.J. City of Bits: Space, Place, and the Infobahn, MIT Press (1996).
23. Nold, C. Biomapping. http://biomapping.net/ (last accessed 01/09/2013)
24. Nold, C. Emotional Cartography. http://emotionalcartography.net/ (last accessed 01/09/2013)
25. OpenStreetMap, Key:place. http://wiki.openstreetmap.org/wiki/Key:place (last accessed 01/09/2013)
26. OpenStreetMap, Map Features. http://wiki.openstreetmap.org/wiki/Map_Features (last accessed 01/09/2013)
27. Ortony, A., Clore, G.L., Foss, M.A. The psychological foundations of the a□ective lexicon in Journal of Personality and Social Psychology (1987), 751–766.
28. Strapparava, C., Mihalcea, R. Learning to Identify Emotions in Text in SAC'08 (2008).
29. Vigneshwara, M. Aleph of Emotions. http://www.mithru.com/projects/aleph.html (last accessed 01/09/2013)
30. Warner, D. Compasses and Coils: The Instrument Business of Edward S. Ritchie, Rittenhouse, Vol. 9, No. 1 (1994), pp. 1-24.
31. Wilhelmer, R., Von Bismarck, J., Maus, B. Fuehlometer. http://richardwilhelmer.com/projects/fuhl-o-meter (last accessed 01/09/2013)
32. Carlo Strapparava and Alessandro Valitutti. WordNet-Affect: an Affective Extension of WordNet, in Proceedings of the 4th International Conference on Language Resources and Evaluation (LREC 2004), Lisbon, May 2004, pp. 1083-1086.
33. Zook, M.A., Graham, M. Mapping DigiPlace: geocoded Internet data and the representation of place in Environment and Planning B: Planning and Design 34(3) (2007) 466–482.