

Chapter 15

An Emotional Compass: Emotions on Social Networks and a New Experience of Cities

Salvatore Iaconesi and Oriana Persico

15.1 Introduction

The map is not the territory (Korzybski 1933).

The map is not the thing mapped (Bell 1945).

The tale is the map that is the territory (Gaiman 2006).

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 (Bateson 1972).

When we experience territories, we create stories. We model these stories using mental maps, referring to one person's point of view perception of their own world, 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 our cultures, to memories, and narratives. 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.

S. Iaconesi (✉) · O. Persico

Department of Digital Design, ISIA School of Design, Florence, Italy

e-mail: salvatore.iaconesi@artisopensource.net

URL: <http://www.artisopensource.net>

O. Persico

e-mail: oriana.persico@gmail.com

URL: <http://www.artisopensource.net>

© Springer International Publishing AG 2018

V. Geroimenko (ed.), *Augmented Reality Art*, Springer Series on Cultural Computing, https://doi.org/10.1007/978-3-319-69932-5_15

277

In the density of cities, this process happens for potentially millions of people at the same time. In his “the Image of the City” (Lynch 1960), 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 (2003) have suggested a broad definition of landmarks to refer to any distinguished city elements that shape our mental map. Or as Appleyard (1969), Golledge and Spector (1978) 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.

We can imagine 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 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 (2003), 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 more effectively and powerfully.

These 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 forming a new physicality, visuality, and tactility which shape our everyday experiences of the world.

According to Mitchell’s “City of Bits” (1996), McCullough’s Digital Ground (2005), Zook’s and Graham’s DigiPlace (2007), 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, to design new types of affordances for our cities.

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

15.2 The Compass

The compass is a historically understood, ubiquitously known object dedicated to navigation and orientation: it finds the direction in which one wants to go.

Compasses 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 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 about the ways in which human beings express their emotions on social networks, and uses it to have insightful emotional experiences in the city?

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?

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

15.3 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 work on Biomapping (2004a) and Emotional Cartography (2004b). In the project, a rather large number of people have taken part in community mapping activities 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. A map is created which visualizes 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’) (2010), 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 Bismarck’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 is 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 (2013) 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 Web site.

Another example comes from an artwork titled *Consciousness of Streams* (2011). In the work, the artists have set up a series of devices or installations in several cities. 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 under the form of a "real-time interconnected emotional map of the planet" (Iaconesi and Persico 2012) showing a topography of human emotions, adjacencies, proximities, and distances which are not physical, but emotional.

Another relevant project is *Mappiness* (2012), 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 are 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, the 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.

An interesting project is "Testing, Testing!" (2011), 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.

Another project which we wish to highlight is the *Aleph of Emotions*, an experimental art project by Vigneshwara (2012): 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 are collected based on keywords that define certain emotions. The results are finally presented with an interactive object. 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—due to the keywords-based approach—and in the lacking sense of immersion in the information landscape.

In yet another example, the MONOLITT (2014) generative sculpture by Syver Lauritzsen and Eirik Haugen Murvold is a paint-emitting plinth that disperses color according to local mood on social media. The installation that quite literally paints the mood of the city using social media feeds as an input. The installation takes electronic signals and lets them manifest themselves in the physical world. Using sentiment analytics, the installation links tweets to corresponding colored paints in real time, feeding them out through the top of the sculpture, letting them flow into a procedurally generated three-dimensional painting. This is a form of environment-based Augmented Reality, much similar to something which could be done through projection mapping: a data phenomenon is taken and interpreted in terms of something that is “added” (or, as in this case, painted) to the architectural environment, allowing people to experience it and, thus, to experience the data phenomenon. While the emotional analysis technique used in this artwork seems far from optimal, the concept and implementation are very interesting in the fact that they produce a physical, tangible output that does not need people to use some technology to experience it. The emotional visualization is there, physical, painted onto the installation, and if properly labeled and explained (maybe through some information panel), is immediately experienceable from the users.

15.4 Concept and Methodology

Our goal was to create an Augmented Reality Compass on a smartphone showing the intensity of emotions in the directions around the user.

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.

15.4.1 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 both legal and technical.

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), which are very complex legal documents.

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 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 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;
- number of interrogations over time;
- storage of the captured information;
- processing of the harvested information;
- visualization and branding.

These legal limits are different across different providers and also change quite frequently and arbitrarily.

Furthermore, it must be said that the issue of expectation for publicness also represents a very important legal aspect. 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. We have often clashed with this kind of issue, for example in trying to harvest all the user expressions on their feelings toward public policies enacted by governments and administrations.

That said, with the help of legal consultants, we have managed to design a replicable model which includes clusters of 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 3 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.

Table 15.1 Number of UGC harvested from social networks in different experiments

City	From date	To date	No. of UGC
London	Jan. 1st 2011	Feb. 1st 2011	5,143,500
Rome	Oct. 15th 2011	Oct. 16th 2011	91,538
Turin	Aug. 1st 2011	Sept. 20th 2011	240,982
Berlin	Jan. 4th 2012	Jan. 20th 2012	1,699,240
Hong Kong	May 1st 2012	Jul. 1st 2012	5,732,487
Cairo	Jul. 27th 2013	Sept. 2nd 2013	3,466,388

We resolved most of these issues adopting a radically modular approach, using interoperable connectors 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. Table 15.1 shows the amount of data which we were able to capture using these methods in various occasions and experiments, across different cities.

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 are 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.

15.4.2 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 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 "the 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 do not 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 messages. Thus, we are able to use it.

From the analysis of the results of our experiments, the geo-location features offered by social networks are not very commonly used. The most common user behavior is to either turn on the location sharing features when they download the applications to their smartphone, or to forget about them.

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 is more complex and interesting. It takes place when users do not use the platforms' features to include their location in the message, but, rather, mention the location which they are 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: these two possibilities may completely change the relevancy of the message.

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 (for example for Italy we have used the named places provided by ISTAT, Italy's National Statistics Institute [2013](#));
 - the list of named places contained in the OpenStreetMap databases, for example as described in OpenStreetMap ([2013a](#), [b](#));

- 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 (2013) or Foursquare (2013);
- 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 if 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 20,000 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 0.2 (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.

15.4.3 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 affective meaning. As suggested in by Ortony et al. (1987), 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 (Fellbaum 1998) extension called WordNet Affect (Strapparava and Valitutti 2004).

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 (1992), 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, cholera, 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 Strapparava and Mihalcea (2008), which we adopted for the details of the technique.

We adapted the technique found in Iaconesi and Persico (2012) to handle multiple languages by using the meta-data provided by social networks to understand in which language messages were written in 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). For example, Fig. 15.4 shows the results of a full week of emotional harvesting in the city of Rome, for the emotion “trust.”

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 Carvalho et al. (2009) and in Bermingham and Smeaton (2010) with varying results.

15.4.4 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.

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.

In the first instance, we tried to use a standard AR canon, with the information being displayed on top of the live camera feed. In this layout, an arrow constantly showed the "forward" direction on the and was color coded as to indicate the level of danger for the current direction: from a vivid green showing a lack of evidence about violence, to a full red denoting the presence of many violence-related messages in that direction (see Fig. 15.1).

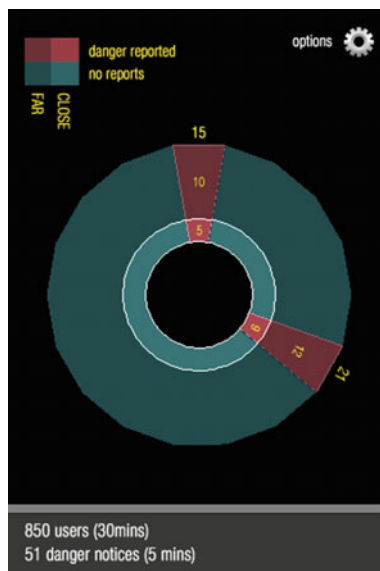
While this configuration of the interface turned out to be really usable and accessible, it did not satisfy us for its readability. The information shown was extremely synthetic, bringing down the complexity of the available information to a single color. Of course, this was in line with the idea of implementing a compass, showing information about the direction in which the user is facing. But we felt that the trade-off represented by losing all of the missing information (such as the declination of the different emotions involved in determining the output, or the possibility to show the messages taken into account to make the decision) was too steep.

In the next design iteration, the information was, then, drawn on screen using a radial diagram, 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 Fig. 15.2).

Fig. 15.1 First iteration of the design: the arrow shows the predominant emotion in the current direction



Fig. 15.2 Second iteration of the design: a rotating radial diagram highlights the danger zones



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 third iteration of the interface was more general purpose (Fig. 15.3).


In this new form, 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 m to 1 km (e.g., if you choose 500 m, only the emotions generated within a 500 m radius will be taken into account); and in time, from 5 min 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

Fig. 15.3 Third iteration of the design: the emotional multi-compass



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.

15.5 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 stages 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.

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 emotions 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 than 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.

Using the application to follow multiple emotions at the same time has proven to be somewhat hard: 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 get a really accurate sense of the compresence 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 (Fig. 15.4).

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.

Since its creation, we have tested and used the Emotional Compass a number of times, at festivals, conferences, events, workshops, and even in public scavenger hunts whose purpose was to navigate the city in a different way, and to try to




Fig. 15.4 An example of emotions captured in a city: “trust” in Rome, during the days/times of a full week

establish readable connections between the emotions detected on social networks and the things that were physically happening in the city.

All times, this was complex matter. As it is understandable through common sense and through detailed research (Barberà and Rivero 2014), what happens on social networks relates only partially to what can be understood by physically traversing the city.

This is due to multiple reasons: the different geographies of online cities, in which people express “for,” “from,” and “about” the locations of the city in a variety of ways, only some of which require them to be actually present (and, thus, physically “readable”) in the location itself.

Different forms of expression also characterize such activities, in which the measure and proportions of the different emotions expressed are determined sometimes and to some degree by the architecture and interfaces of the social networks themselves (Stieglitz and Dang-Xuan 2013), where on some social networks some forms of emotional expression appear to be definitely more frequent than others.

And, more in general, the form of the public, private and intimate spheres change on the different social networks, ranging from the Habermas-type “private-publics” of social applications such as WhatsApp, to complex hybrids such as Facebook, in which multiple forms of public, private, and intimate spheres coexist and penetrate each other, depending from users’ complex privacy settings and ways of publishing (for example, a certain user may publish a post to be visible only to some of their connections).

These limitations also turn out to be opportunities, as they provide us with tools to materialize, through Augmented Reality, other, different, physically unforeseen layers of reality to the world, improving the ways in which general reality may be read, understood, and experienced.

15.6 Conclusions

We have found this research path to be rewarding for its implications in terms of the possible artworks and 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 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.

References

- Appleyard D. Why buildings are known. *Environ Behav.* 1969;1:131–56.
- Art is Open Source. Consciousness of streams. <http://www.artisopensource.net/2011/01/12/cos-consciousness-of-streams/>. Accessed 1 Sept 2013, and <http://cos.artisopensource.net/> (2011). Accessed 20 Sept 2017.
- Barberà P, Rivero G. Understanding the political representativeness of Twitter users. *Soc Sci Comput Rev.* 2014.
- Bateson G. Form, substance and difference in steps to an ecology of mind. San Francisco: Chandler Publishing Company; 1972.
- Bell ET. Numerology: the magic of numbers. New York: United Book Guild; 1945.
- Bermingham A, Smeaton AF. Classifying sentiment in microblogs: is brevity an advantage? In: *CIKM'10 Proceedings of the 19th ACM international conference on information and knowledge management*, Toronto; 2010. p. 1833–6.
- Berry M. Large-scale sparse singular value computations. *Int J Supercomput Appl.* 1992;6(1): 13–49.

- Carvalho P, Sarmento L, Silva MJ, 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, Hong Kong; 2009. p. 53–6.
- City of Vilnius. The happy barometer. <http://happybarometer.com/> (2013). Accessed 20 Sept 2017.
- Ellard C. Testing, Testing! <http://www.bmwguggenheimlab.org/where-is-the-lab/mumbai-lab/mumbai-lab-city-projects/testing-testing-mumbai> (2011). Accessed 20 Sept 2017.
- Facebook. Obtaining a list of places in a geographical area. <https://developers.facebook.com/docs/reference/api/search/> (2013). Accessed 1 Sept 2013.
- Fellbaum C. WordNet. An electronic lexical database. Cambridge, MA: MIT Press; 1998.
- Foursquare. Obtaining a list of places in a geographical area. <https://developer.foursquare.com/docs/venues/search> (2013). Accessed 20 Sept 2017.
- Gaiman N. Fragile things. New York: HarperCollins; 2006.
- Golledge RJ, Spector A. Comprehending the urban environment: theory and practice. *Geogr Anal.* 1978;10:403–26.
- Haken H, Portugali J. The face of the city is its information. *J Environ Psychol.* 2003;23(4): 385–408.
- Iaconesi S, Persico O. ConnectiCity: real-time observation and interaction for cities using information harvested from social networks. *Int J Art Cult Design Technol (IJACDT).* 2012;2 (2):14–29.
- ISTAT. Territorial Data Sets, including named places. <http://sitis.istat.it/sitis/html/>, <http://www.istat.it/it/prodotti/banche-dati>, <http://www.istat.it/it/archivio/44523> (2013). Accessed 20 Sept 2017.
- Korzybski A. A Non-Aristotelian system and its necessity for rigour in mathematics and physics, American association for the advancement of science 1931. In: Conference proceedings. Reprinted in *Science and Sanity*, New Orleans; 1933. p. 747–61.
- Lauritzsen S, Murvold EH. MONOLITT. <https://www.youtube.com/watch?v=8a7BNgDfRXo> (2014) Accessed 20 Sept 2017.
- London School of Economics. Mappiness. <http://www.mappiness.org.uk/> (2012). Accessed 20 Sept 2017.
- Lynch K. The image of the city. Cambridge, MA: MIT Press; 1960.
- McCullough M. Digital ground: architecture, pervasive computing, and environmental knowing. Cambridge, MA: MIT Press; 2005.
- Mitchell WJ. City of bits: space, place, and the Infobahn. Cambridge, MA: MIT Press; 1996.
- Nold C. Biomapping. <http://biomapping.net/> (2004a). Accessed 20 Sept 2017.
- Nold C. Emotional cartography. <http://emotionalcartography.net/> (2004b). Accessed 20 Sept 2017.
- OpenStreetMap. Key:place. <http://wiki.openstreetmap.org/wiki/Key:place> (2013a). Accessed 20 Sept 2017.
- OpenStreetMap. Map features. http://wiki.openstreetmap.org/wiki/Map_Features (2013b). Accessed 20 Sept 2017.
- Ortony A, Clore GL, Foss MA. The psychological foundations of the affective lexicon. *J Pers Soc Psychol.* 1987;53:751–66.
- Stieglitz S, Dang-Xuan L. Emotions and information diffusion in social media—sentiment of microblogs and sharing behavior. *J Manage Inf Syst.* 2013;29(4).
- Strapparava C, Mihalcea R. Learning to identify emotions in text. In: SAC'08 Proceedings of the 2008 ACM symposium on Applied computing, Fortaleza; 2008. p. 1556–60.

- Strapparava C, Valitutti A. WordNet-affect: an affective extension of WordNet. In: Proceedings of the 4th international conference on language resources and evaluation (LREC 2004), Lisbon; 2004. p. 1083–6.
- Vigneshwara M. Aleph of emotions. <http://www.mithru.com/projects/aleph.html> (2012). Accessed 20 Sept 2017.
- Wilhelmer R, Von Bismarck J, Maus B. Fuehlometer. <http://richardwilhelmer.com/projects/fuhl-ometer> (2010). Accessed 20 Sept 2017.
- Zook MA, Graham M. Mapping DigiPlace: geocoded internet data and the representation of place. *Environ Plan B Plan Design*. 2007;34(3):466–82.