The Exceptional and the Everyday: 144 Hours in Kiev

Lev Manovich
The Graduate Center
City University of New York
New York, NY, U.S.A.
manovich.lev@gmail.com

Mehrdad Yazdani
California Institute for
Telecommunication and Information
La Jolla, CA, U.S.A.
myazdani@ucsd.edu

Alise Tifentale
The Graduate Center
City University of New York
New York, NY, U.S.A.
atifentale@gc.cuny.edu

Jay Chow
Web Developer
Katana
San Diego, CA, U.S.A.
jaychow0702@gmail.com

***Abstract*—How can we use computational analysis and visualization of content and interactions on social media network to write histories? Traditionally, historical timelines of social and political upheavals give us only distant views of the events, and singular interpretation of a person constructing the timeline. However, using social media as our source, we can potentially present many thousands of individual views of the events. We can also include representation of the everyday life next to the accounts of the exceptional events. This paper explores these ideas using a particular case study – images shared by people in Kiev on Instagram during 2014 Ukranian Revolution. Using Instagram public API we collected 13208 geo-coded images shared by 6165 Instagram users in the central part of Kiev during February 17-22, 2014. We used open source and our own custom software tools to analyze the images along with upload dates and times, geo locations, and tags, and visualize them in different ways.**

Keywords—social media; photography; digital humanities, Instagram, social movements

I. INTRODUCTION

Over a few days in February 2014, a revolution took place in Ukraine. How was this exceptional event reflected in user-generated media shared on Instagram in Kiev? What can visual social media tell us about the experiences of thousands of people and life in their city during a social upheaval?

While other researchers and journalists have used quantitative analysis of social media during social protests and uprisings, this is the first project to focus on Instagram. Using Instagram public API (see *Constructing the Dataset* section for details), we collected all publicly available geo-tagged Instagram images shared during February 17-22, 2014 in the central part of Kiev. The area was centered on Independence Square, the key place of protests and confrontations with government forces. Our dataset contains 13,208 images shared by 6,165 Instagram users. The images were tagged with 5,845 unique tags.

We used open source and our own custom software tools to analyze the images along with upload dates and times, geo locations, and tags, and visualize them in different ways. To compare the patterns during the days of the Revolution with a more normal period, we also downloaded over 400,000 images shared in the same area of Kiev between February 24 and May 14, 2014.

When news outlets report on events such as uprisings, wars or revolutions, they typically show us only the events taking place in particular areas – demonstrations, clashes with police, etc. – as opposed to everything else that is going on outside these areas. This makes sense since news is there to inform us about things that have global political and economic importance (for example, who will govern Ukraine next and how this will play into European and global geopolitics). As a result, when events of this scale are taking place, in media reports they stand in for the whole city or a country. Nothing else is visible at these times. For example, if we are to look at news images of Kiev published during Maidan (the name used to refer to the oppositional movement and 2014 Revolution) events, the whole city reduced to what was taking place on Independence Square.

But if we consider all Instagram photos shared during the same days in central area of Kiev that includes the Square, a very different picture emerges. Outside of the most intense days and the areas of protests, you would not even know that something political was happening. People post selfies and other photos of their lives. They dress up getting ready to go out, and take photos of cultural events. The images of Maidan clashes, political slogans, and burned cars and buildings appear right next to everything else. Most people continue their lives and post their "likes" as on any other day. The exceptional co-exists with the everyday. We saw this in the collected images, and this was our motivation to begin this project.

In presenting this analysis, we have no intention of downplaying the importance of 2014 Ukrainian Revolution, the heroism of people who made it happen, and the work of everybody else who supported Maidan movement. Nor are we saying that most people in Kiev who shared images on

978-1-4799-5666-1/14/$31.00 ©2014 IEEE

Instagram during the Revolution had no interest in politics, just because they posted photos of something else. Intentions and interests of a person can't be guessed from a single image. And the same image can mean different things depending on what else is next to it. What we do want to do is to explore the things that are usually left out from the brief news reports – the everyday, and its relationship with the exceptional, as it is reflected and staged on Instagram.

For journalists, social media is a window into what actually took place. Thus it becomes important to identify who posted what and who among them was a real participant. The web users who can be verified are treated as additional sources of the news. While we fully respect this approach and understand its practical usefulness, in this project we approach social media as its own reality, separate from the "real" reality on the ground.

13,208 geo-tagged images from one area shared on Instagram over six days with their 21,468 tags in three languages (and hundreds of thousands of words used in descriptions) paint their own fascinating picture. This picture is not a "photo" of social reality. Instead, it can be compared to a modern painting. It contains some references to the world outside, but it is not its realistic copy. Our goal was to "see" this picture dispersed between all images, tags, time stamps, and geo-locations. To do this, we explored all this data in as many ways as we could and then selected what we feel were most interesting "views" of this picture.

This paper continues our previous investigations where we visualize and interpret patterns in visual social media such as 2.3 million of Instagram images shared in 13 global cities [1].All visualizations that appear in this paper are also available in high resolution at the project web site http://www.the-everyday.net/.

II. SIX DAYS IN KIEV

To make the following discussion understandable, we first briefly summarize the events of the 2014 Ukrainian Revolution:

TABLE I.

Date	Summary of the events
2/17/2014	The night before the 18th, oppositional parties called for all concerned citizens to take a part in "peace offensive."
2/18/2014	In the morning, tens of thousands of demonstrators were attacked by police. After a day of fighting in a few areas in the city, police launched the attack on Independence Square at 8pm. But by the midnight, 20,000 people still remained on the square.
2/19/2014	Government closed the metro and blocked main roads. According to one report, 30,000 were now at the square, preparing for more confrontations.
2/20/2014	Another day of fighting between the protesters and the police and Berkut (special government forces). During these days, 103 protesters and 13 police were killed.
2/21/2014	An agreement between the protesters and the president calling for constitutional reform and new elections was reached. Soon thereafter President Yanukovych and most of his ministers fled the city.
2/22/2014	Former President Yulia Tymoshenko was released from prison and she addressed over 100,000 people on the square. By February 23, transition towards new temporary government was underway.

[a] This brief summary was compiled from: http://en.wikipedia.org/wiki/2014_Ukrainian_revolution; http://en.wikipedia.org/wiki/List_of_people_killed_during_Euromaidan. The numbers of people at Maidan Square are estimates by different journalists and agencies.

III. WHY INSTAGRAM?

The word *euromaidan* (the name for the popular movement in Ukraine which led to the 2014 Revolution) was first used as a hashtag on Twitter [2]. Protesters and their supporters actively employed Twitter and Facebook to organize the gatherings and demonstrations and communicate news to the outside world. The protesters and other opposition parties made most political announcements on Facebook, which had approximately 3 million users in Ukraine in February 2014 [3]. "EuroMaydan" page became the most-liked Ukrainian page on Facebook (with 304,590 likes as of 10/04/2014). Even more people in Ukraine were using VKontakte - the largest social network in Europe with English, Russian and Ukrainian as official languages.

In comparison, Instagram had significantly less users in Ukraine. According to the current Alexa.com data, VKontakte is third top web site in Ukraine, right after Google's international page and Ukrainian page. Facebook is no. 6, Twitter is no. 2, and Instagram is no. 56 [4]. However, despite its smaller size in Ukraine, Instagram gives us a unique view into the events in Kiev during 2014 Ukrainian revolution. As a global network organized around photography, it offers a different picture than other social networks: a visual account of the life in a city, the desires and imaginations of its people (or at least, people in their 20s), and their actions and thoughts during important social and political events. (The data on demographics of Instagram users around the world is not available. Our research on Instagram selfies in New York, Berlin, Sao Paulo, Bangkok and Moscow showed that at least among people posting selfies in these cities, the largest numbers were in their mid-20s [5].)

Examining images and data we collected, it appears that unlike Facebook and Twitter, Instagram was not used systematically for communication by protesters, oppositional parties or the government. Our image set is not dominated by a few power users posting disproportional numbers of images. And the images themselves are quite varied - we do not see a few images repeating endlessly. Therefore, although we will never claim that Instagram picture of the Ukrainian Revolution days is "objective," it is at least representative of interests and experiences of significant numbers of people. Thus, its relatively small user base in Ukraine in February 2014 makes it more (as opposed to less) useful for research.

IV. CONSTRUCTING THE DATASET

Using Instagram API, we downloaded all available geo-tagged Instagram images and video publically shared in the central area of Kiev between 02/02/2014 and 05/15/2014. The collection area is a rectangle centered on Independence Square: 3.9 miles by 6.2 miles (6.3 km by 9.9 km). The dataset contains 463,989 media files (%3.6 are video and the rest are still images).

We focused our analysis on the days of 2014 Ukrainian revolution: February 18-21. In order to better understand Instagram patterns during these exceptional days, we also included one day before and one day after. Consequently, our final period for the project is February 17-February 22, 2014. During this six-day period, 6,165 Instagram users shared 13,208 images in the central part of Kiev. This media has 5,845 unique tags; the total number of tags is 21,468.

V. CONSTRUCTING THE PICTURE

What can we see in the world if we only use social media content such as Instagram photos and their metadata (descriptions, tags, locations)? Analysis and visualization of large samples of social media can provide an alternative to summaries of the events presented by historians, individual journalists, or groups of writers (e.g., collaboratively authored Wikipedia articles). This is especially true for the visual summaries of the events. Instead of only a few views we can now have thousands or event millions of separate views.

Of course, often these are only fragments and bits too short to articulate a full statement - but often they are not. And with images, the results can be particularly interesting, since even a single image can contain much more information than many Twitter posts put together.

While media outlets also personalize their reports by interviewing some of the participants and then including parts from these interviews into the reports, this is not the same thing. The diversity of perspectives by tens of thousands of participants can be much larger than that of only a few who were interviewed.

Combination of computer data analysis and visualization can help us to juxtapose these perspectives. We can find commonalities and differences, and discover typical as well as unique perspectives. But we have to remember that as any other visual media, data visualization is not neutral. By organizing images and data in particular way, we can tease this or that pattern. Some patterns may be given too much attention, while others may remain hidden.

For example, organizing all 13,208 images shared by 6,165 Instagram users strictly by their date and time creates visualization where the exceptional and the everyday are dramatically juxtaposed (figures 1 and 2). At the same moment as one person shares a photo of the demonstrations, another person is posting her portrait, and yet another is posting a photo from a party the night before. Such a "film" created by projecting thousands of images taken by thousands of people over a large city area onto a single linear time dimension creates a picture of extreme fragmentation which is perhaps even more intense than the modernist collages of cities created hundred years earlier. But it is important to remember that this particular picture is not "native" or "natural" to Instagram in general, or the use of Instagram by people in Kiev during Maidan effect. It is the result of our systematic decision to organize Instagram images shared by thousands of individuals in a particular way for the visualization.

VI. THE EXCEPTIONAL AND THE EVERYDAY

Having discussed our motivations and goals, the construction of the dataset, and the construction of its "views" though computer analysis and visualization, we will now present some of these views and their interpretations. They are selected from hundreds of graphs and visualizations we generated while working with the data.

A. Flow

We start with a visualization of all 13,208 images shared by 6,165 Instagram (Fig. 1).

Fig. 1. A visualization showing 13,208 images shared on Instagram central part of Kiev during February 17-22, 2014. The images are organized by shared date/time (top to bottom, left to right).

Six light-to-dark "waves" correspond to the six days (lighter images during the day, darker images at night). From this bird's eye view, we don't see any obvious reflections of the exceptional events that took place during this period. It seems as though the Revolution never took place.

Fig. 2 is the close-up view of the visualization, with the addition of time stamp right above each image. We also added a dark background and a space around every image. Each row shows all photos shared within a short period during the night and morning of February 18. Looking at this scale, we now see photos of the events at Independence Square (fires and crowds of people in first and second row) next to the photos showing other subjects.

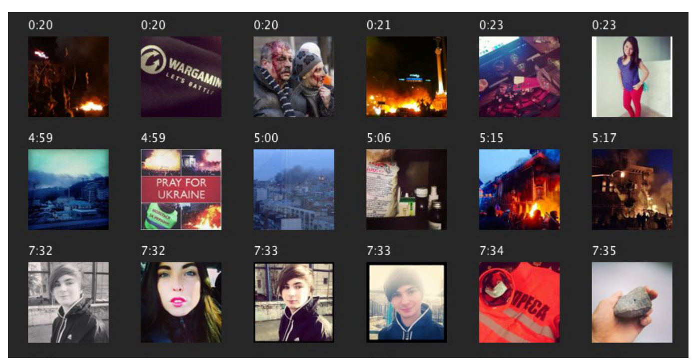

Fig. 2. The close-up view of the visualization in Fig. 1, with the addition of time stamp right above each image.

B. Time

Now we will visualize the data as a graph, plotting the number of shared images over time (Fig 3). Grey part of the graph shows all shared images between February 17 and 22. Red part corresponds to images shared only around Independence Square (1,900 images, or %14 of the total).

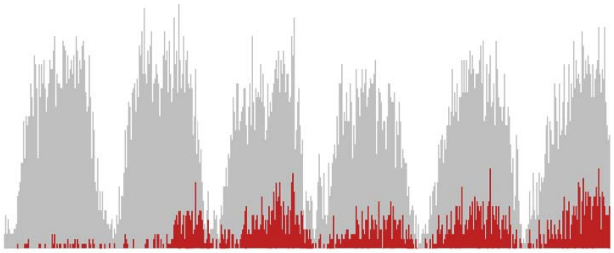

Fig. 3. The numbers of Instagram images shared in the central part of Kiev during February 17-22 (left to right). Each "mountain" corresponds to one 24-hour period. Grey: all shared images. Red: images shared around Independence Square.

If we look at the grey part of the graph, all the days are quite similar – few images in early morning, more in the afternoon, and then gradual decrease in the evening. The last day in our period (February 22) is somewhat different probably because it is Saturday, so people are waking up later and going out later as well. Again, looking only at the volume of all shared images, you would not know that a revolution took place.

Filtering the data to select only images shared around Independence Square (red area) tells a different story. There is almost nothing on the first day before the events start (2/17). The next day there are confrontations between the protestors and government forces in a few areas (2/18). Police stages a big attack on Independence Square at 8 pm, and it continues for hours. In the graph, we see a big jump already around 6pm. The next day there is no fighting on the square, but more people are arriving on the square in preparation for the next fight, and many also coming to see and take photos with their phones (2/19). This high level of activity continues for the rest of our selected period.

C. Space

Having separated time into two streams – everything shared in the central part of Kiev and the part of the dataset containing images shared around the area of Independence Square – we made the exceptional clearly visible. We also revealed its temporal shape (red part in Fig. 3).

We will now switch from time to space, and map the locations of images. Figure 4 shows two maps. The first displays locations of images shared during the day before confrontations start (February 17). The second shows locations of images shared during next five days, i.e. February 18-22. (Note that since our collection was limited to a rectangle centered on Independence Square, there are no points outside of this rectangle).

Fig. 4. The locations of images shared during February 17 (top) and February 18-22 (bottom).

There is no spike in activity on Independence Square on the 17th, but over next five days the exceptional – the fights taking place around the square, and the massive meetings there afterwards – clearly stand out.

We can "slice" these map across a single coordinate (latitude or longitude). The graphs in fig. 5 show volume of shared images (vertical axis) against latitude coordinate (horizontal axis) for February 17 and February 18-22. The story told by these maps and slices is similar to what we saw in time plot in Fig 3. – but now the exceptional stands out more dramatically. On the top graph in Fig. 6 the area of Independence is just a part of overall busy activity in the city center. But during the next five days it completely dominates all other areas (the bottom graph).

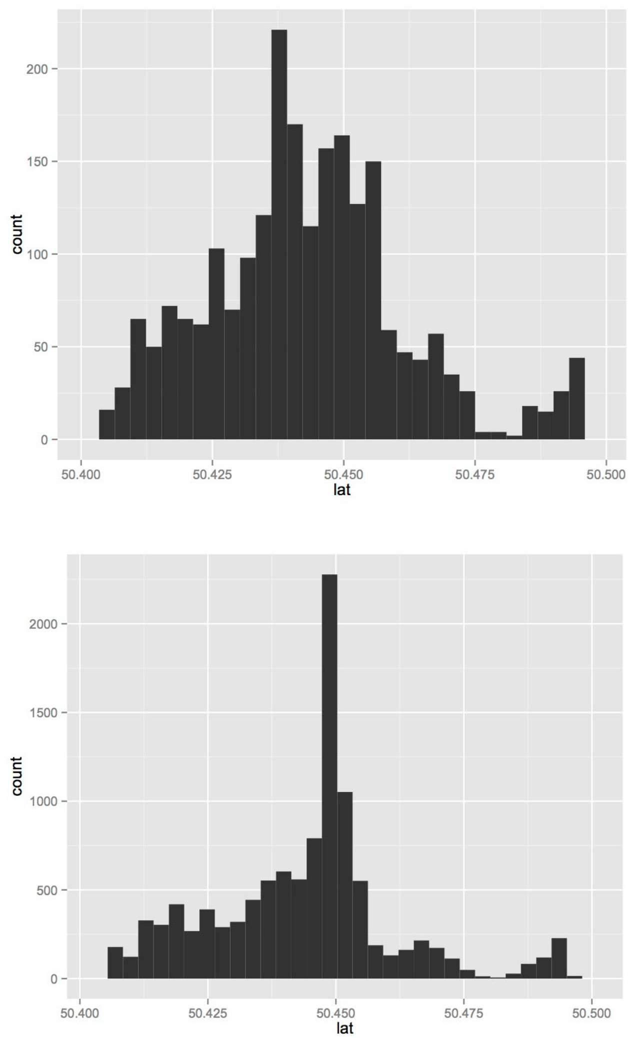

Fig. 5. Volumes of shared images (Y) against their latitude coordinate (X). Top: February 17. Bottom: February 18-22.

D. *What is #euromaidan?*

Our images have seven different Maidan tags, spelling the word differently in English, Russian and Ukrainian. These tags are: #майдан, #maidan, #euromaidan, #євромайдан, #евромайдан, #euromaydan, #Euromaidan. Overall, 1,340 images have at least one of these tags (%10 of the total). Figures 6 and 7 show the visualization of these 1,340 images organized by upload date and time (left to right, top to bottom).

Fig. 6. 1340 images which have Maidan tags organized by time (left to right, top to bottom).

Zooming into this visualization we find that while most of the images with Maidan tags are indeed directly related to Maidan events, some are not (fig. 7). (We also visualized the images that do not have any Maidan tags, and discovered that occasionally they show the events on the square. And when we visualized the images shared around the square, we found that some of them have no obvious relation to Maidan events.) So we cannot exclusively rely on tags to predict the subjects of images. For a computer scientist concerned with detecting social upheavals in social media, this finding would indicate a problem that needs to be solved. Such images may be considered as "noise," only to be removed from the "signal." But for us, they are the real "finding." They show that the everyday and the exceptional do not simply "co-exists" side-by-side (as presented in Figures 1-2). Instead, they "leak" into each other, so to speak.

As we discovered, this "leakage" has its own patterns. Out of 1,340 images that have one or more Maidan tags, %30 also have some other tag(s). Images with only Maidan tags typically indeed refer to Maidan events. But if it also has other tags, there is a chance that it shows another subject (such as the selfie in upper left corner of the close-up in Fig. 7).

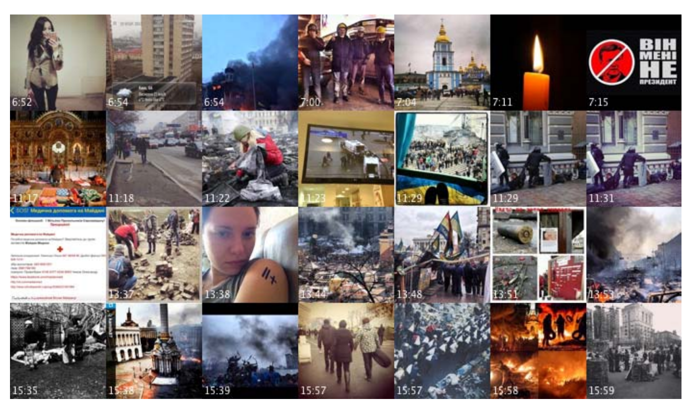

Fig. 7. A close-up of the visualization in Fig. 6.

E. *"Speaking the images"*

When a person assigns multiple tags to a single image, this image literally multiplies - because it will now show up in searches for any of these tags. For example, let's say an image has #euromaidan tag and also #майдан tag (the word "maidan" spelled in Russian). If you use Instagram to search for these tags, this image will show up in both results.

When a person applies multiple Maidan tag versions to the same image, every tag potentially says something else. "Maidan" spelled in English, in Ukrainian, or in Russian is not the same thing. So the reason behind using more than one tag for the same word is not only to have wider dissemination of images, or to address different linguistic communities. Assigning a new tag to an image is like saying something else about this image.

Our next visualization (Fig. 7) is based on this idea. Like the visualization in Fig 6, it shows all images that have one or more Maidan tags, but now every image is repeated for each of its Maidan tags. For example, if an image has #euromaidan and #майдан tags, it is repeated twice. As a result, 1,340 images turn into 2,917.

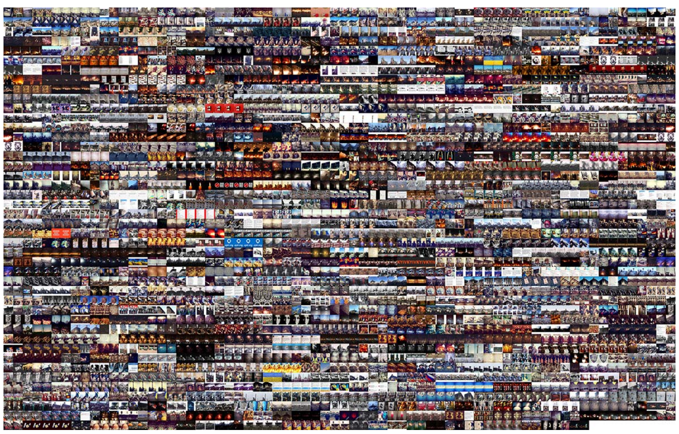

Figure 8. 1340 unique images with Maidan tags. Each image is repeated for each of its tags. The images are organized by date and time (left to right, top to bottom).

Figure 9. A close-up of the visualization in Fig. 8.

F. *Clusters*

The words in the tags and the number of times the tags with with the same concept are applied to the same images often can tells what images represent, but they also can be misleading. But even if the tags and image content match perfectly, they usually can't completely describe everything an image represents, and how it represents it (composition, focus, etc.). If it is true that "an image equals a thousand words," we certainly are out of luck - no single Instagram image in our dataset has as many tags.

Let's take an Instagram photo of a person as an example. The tags can tell us about the gender of a person (#girl, #guy), her or his mood (#happy), and perhaps location (#beach) - but not about the composition, or colors, or all other visual dimensions.

For our next stage of the analysis, we will bypass language. Instead, we will use computer algorithms to automatically separate images into groups based only on visual similarity. Each group contains some images that have something in common. We use digital image analysis to extract visual characteristics of images, and disregard tags, locations and date and time information.

Using cluster analysis, we divided all images into 60 clusters. A few of these clusters consist from mostly Maidan related images; most others are not. These latter reveal the types of "everyday" in Kiev as it presented on Instagram: double portraits, objects against a white background, city views with light sky and darker lower parts, etc.

Fig. 10 show three examples of our clusters. The first cluster contains mostly city views with light sky and darker lower part. The images in the second cluster show a single subject in the center framed by a light background. These examples illustrate how a large image collection may contain groups of visually similar images that are not visible through direct examination of a collection. While we can notice separate instances of types of images shown in these clusters, it is much harder to see how many such instances a collection contains.

Clustering image by their visual characteristics allows us to understand how representations of the exeptional and the everydays on Instagram are related in yet another way. Each cluster shown in Fig. 10 contains mostly images of the everyday. But because of the similary in composition, the clusters also "catch" a few Maidan related images. In the first cluster, these are the landscapes manipulated to only contain blue and yellow colors of Ukraine flag. In the second cluster, this is a political text situated in front of a candle (second to last row, first and second image from the left).

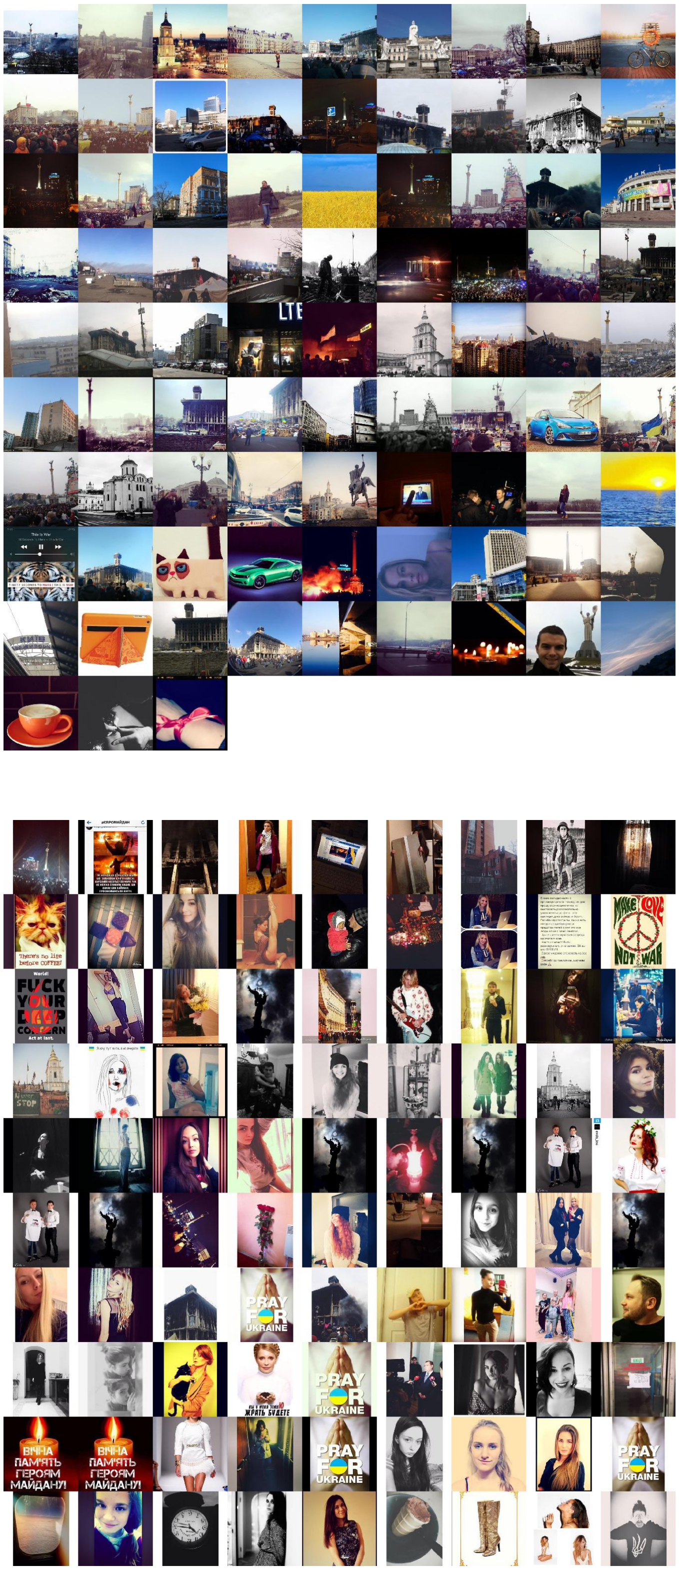

Fig. 9. Two among 60 clusters identified by cluster analysis of our image set.

G. Language

Instagram users assigned 5845 unique tags to images they shared in the larger central area of KIev during February 17-22. If we analyze these tags, would this tell us the same story as the previous graphs and visualizations, or a different one?

Fig. 10 shows top 10 tags for each of the six days, sorted from top to bottom in order of their frequency.

2/17	2/18	2/19	2/20	2/21	2/22
kiev	kiev	kiev	kiev	kiev	kiev
ukraine	ukraine	euromaidan	ukraine	euromaidan	euromaidan
love	киев	ukraine	euromaidan	ukraine	ukraine
follow	euromaidan	євромайдан	євромайдан	євромайдан	євромайдан
followme	євромайдан	евромайдан	евромайдан	евромайдан	евромайдан
instagood	евромайдан	kyiv	revolution	kyiv	киев
me	kyiv	киев	киев	revolution	revolution
киев	revolution	revolution	kyiv	киев	kyiv
like	love	майдан	майдан	Ukraine	followme
photooftheday	followme	Ukraine	love	love	instagood

Fig. 10. Top ten tags for February 17-22. We used blue for the geography tags, grey for the "everyday" tags, and "red" for the tags related to the revoltuon (Maidan tags variations and #revolution).

On the first day (February 17) before confrontations start, we see typical Instagram pattern common to numerous places around the world. The geographic tags which identify the location appear first, followed by "universal" Instagram tags not specific to any location: #love, #follow, followme, #instagood, #me, #photooftheday. These tags are among top 15 tags on Instagram around the world every day (To see the top 100 Instagram tags on any given day, consult http://websta.me/hot/.)

As confrontations begin, the words "EuroMaidan," "Maidan" and "revolution" immediately jump to the top. On February 18, only two universal Instagram tags still make it to the bottom of the top ten ranking (#love and #followme). February 19 is truly exceptional: no generic Instagram tags are in the list. And on February 20 and 21, only single generic tag #love appears in the bottom. Finally, on the 22nd, as the fighting is over and the revolution has succeeded, two generic tags slip back in (#followme and #instagood).

Overall, we get a perfect arc. Before the exeptional events, "the Instagram everyday" dominates the list (February 17th). then it disappears completely (19th), and after that gradually starts coming back (20th-22nd). The exceptional local events push out the universal everyday - but only for a short time.

The pattern is remarkably clear, but it raises a question. The proportion of images shared around Indepedence Square is only %14 of all 13,208 images shared in the larger central part of Kiev. But when we consider the tags of all these images, Maidan tags dominate during the revolution days. Why? One reasonable explanation is that people who did not come to Maidan were still concerned with the events, and they used Maidan tags to show their support of (or other attitudes towards) the revolution.

VII. Conclusion

Humans are always looking for signals standing out against noise. But as modern society developed techniques to generate progressively more data, this became particularly important.

Modern "news" pick out what is important for us to know. Claude Shannon defined information as the amount of unpredictability in a message. Flickr pioneered the use "interestingness" to filter the photos [6]. Google Analytics "monitors your website's traffic to detect significant statistical variations, and then automatically generates alerts, or *Intelligence Events*, when those variations occur" [7]. These are only two examples of the key technology of "data society" - data mining – that uses automatic computational techniques "with the intention of uncovering hidden patterns in large data sets" [8].

The metaphor of mining suggests that you recover what is valuable and discard the rest. But what if we reverse the procedure?

In his 1973 text *Species of Spaces and Other Pieces*, French writer Georges Perec noted how the news only talk about the exceptional, but never the everyday: "What speaks to us, seemingly, is always the big event, the untoward, the extra-ordinary: the front-page splash, the banner headlines. Railway trains only begin to exist when they are derailed, and the more passengers that are killed, the more the trains exist… How should we take account of, question, describe what happens every day and recurs everyday: the banal, the quotidian, the obvious, the common, the ordinary, the infra-ordinary, the background noise, the habitual?" [9]

During a weekend in October 1974, Perec set out to realize his idea of systematically capturing *infra-ordinary*. Over three days, he described what he saw from a window of a café in Place Saint-Sulpice: buses, cars, people passing by, ordering coffee, the pigeons, and so on. The results were published as a short book *An Attempt to Exhaust a Place in Paris.*

Similar to Perec, in our project we use a rectangle as a frame, and only take into account what this frame captures. For Perec, it was a window in a café; for us, it's the rectangle on the map defined by longitude and latitude coordinates passed to Instagram API. But since we substituted a single human point of view by the social media network, we can stretch our frame, to capture over much larger area. And this what we did in our project. This allowed us to observe both infra-ordinary and the extra-ordinary, and reconstruct some of the ways in which they interact.

VIII. CONCLUSION

"What can we see in the world if we only use social media content such as Instagram photos and their metadata?" In this paper we investigated this question using the images shared during 2014 Ukrainian revolution as our case study. In the news reports of this exceptional event, the whole city was reduced to a single square where the confrontations before the protesters and goverment were taken place. Instagram gives us a very different picture. The images of the Revolution appear right next to many other subjects. Using a number of visualization techniques and different parts of the data (images, tags, dates, etc.), we explored the various ways in which the representations of the exceptional and the ordinery interact with each other.

IX. APPENDIX: CLUSTER ANALYSIS DETAILS

To find clusters of similar images, we use the k-means clustering algorithm provided by the R statistical programming environment. All images are 150-by-150 pixels (the standard Instagram thumbnail dimension). We extract images features in Python using the scikit-image library [10]. From this library, we use the Histogram of Oriented Gradients (HoG) features [11] with 8 gradient orientations and non-overlapping 16-by-16 window cells. In addition, we also use the raw grayscale pixel values of the images as features, as well as the latitude and longitude of image locations. Since these feature have different scales and units, we re-scale them to have zero mean and unity standard deviation.

In k-means clustering, the number of clusters must be specified a priori. We experimented with several values for k, ranging from 10 to 100. For each value of k, we use 10 random starts to ensure that we find clusters that converge to the same cluster centers. When the number of clusters is large, there are redundant clusters that can be merged based on our evaluation of similarity. To evaluate the found clusters, we use the ImageMontage [12] tool to visualize all images in each cluster (see examples in Fig. 10). The primary advantage of using a large number of clusters is that non-linear cluster boundaries can be approximated with k-means. After a number of experiments, we decided that k=60 provides best results, and this is the number we used in this paper.

REFERENCES

[1] N. Hochman and L. Manovich, "Zooming into an Instagram city: Reading the local through social media." First Monday, 6/1/2013.

[2] http://en.wikipedia.org/wiki/Euromaidan#Name_history

[3] J. Dickinson, "The Revolution Will Be Live-streamed: Social Media and the Maidan," presented at SOYUZ Symposium, Havighurst Center, 3/1/2014. Unpublished.

[4] http://www.alexa.com/topsites/countries/UA, retrieved 10/04/2014.

[5] L. Manovich, M. Stefaner, M. Yazdani. D. Baur, D. Goddemeyer, A. Tifentale, N. Hochman, J. Chow. "Selfiecity." New York, February 2014. www.selfiecity.net

[6] Flickr, Interestingness Ranking of Media Objects, U.S. Patent #US008732175, Filed February 8, 2006. http://pdfpiw.uspto.gov/.piw?PageNum=0&docid=08732175&IDKey=B8DCE9513A3F

[7] https://support.google.com/analytics/answer/1320491?hl=en&ref_topic=1032994, accessed 11/8/ 2014

[8] http://en.wikipedia.org/wiki/Data_mining

[9] G. Perec, Species of Spaces and Other Pieces, transl. J. Sturrock. London; New York: Penguin Books, 1997. First published in 1973 in *Cause Commune*.

[10] van der Walt S, Schönberger JL, Nunez-Iglesias J, Boulogne F, Warner JD, Yager N, Gouillart E, Yu T, the scikit-image contributors, "scikit-image: image processing in Python," *PeerJ* 2:e453 http://dx.doi.org/10.7717/peerj.453, 2014.

[11] N. Dalal B. Triggs, B., "Histograms of Oriented Gradients for Human Detection," IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2005, San Diego, CA, USA.

[12] https://github.com/culturevis/imagemontage