

Editorial Algorithms: Using Social Media to Discover and Report Local News

Raz Schwartz

Jacobs Institute
Cornell Tech
raz.schwartz@cornell.edu

Mor Naaman

Jacobs Institute
Cornell Tech
mor@jacobs.cornell.edu

Rannie Teodoro

Rutgers University School of
Communication and Information
rteodoro@rutgers.edu

Abstract

The role of algorithms in the detection, curation and broadcast of news is becoming increasingly prevalent. To better understand this role we developed CityBeat, a system that implements what we call “editorial algorithms” to find possible news events. This fully functional system collects real time geo-tagged information from social media, finds key stories, makes an editorial decision whether these events are of interest and eventually visualizes these stories on a big screen display. The system was designed in collaboration with journalists and tested at four New York City newsrooms. Our results show that while journalists were receptive to the idea of machine-generated news stories, the actual results of the system confirmed current concerns among journalists and researchers about the dangers of outsourcing news finding tasks to machines. This paper therefore exemplifies how news sourcing systems based on social media may favor specific type of news events, do not report results quickly enough, and cater to a biased population and interests.

Introduction

Machine-sourced news is a hot topic among newsrooms. Recent technological developments in machine learning algorithms together with the proliferation of information sources such as sensors, drones and social media, create new opportunities to revolutionize the practice of journalism (Pavlik 2013; Goldberg, Corcoran, and Picard 2013).

Big newsrooms and national publications are already on this wagon, employing data scientists, statisticians and data visualization designers to transform the way they find and report news (Hunter 2014). Local news outlets however, despite the constant demand to provide live and full coverage of local events, have limited access to resources such as time, money, or people (Stencel, Adair, and Kamalakanthan 2004). As a result, many local publications are looking for new ways to provide local news coverage as they cut down on print publications and move to online only formats (Farhi 2014).

In this work, we examine the opportunity and challenges in using social media to automatically produce local news. We developed CityBeat, a real time system that consists of an “editorial algorithm” that finds possible news events

based on geo-tagged social media and makes an editorial decision about the level of accuracy and interest of a news event.

We performed a multistage user-centered study that included a process of requirements gathering, design and iterations, system development, and finally deployment and evaluation. We chose to co-develop and test CityBeat with a group of journalists and newsrooms as they are “domain experts” in the field of local news. After the initial development, we conducted an iterative design process, and further tested via deployment and case studies in live newsroom settings, at four well-established news organizations in New York City.

This paper reports on the initial requirements gathering, provides the technical details of the developed CityBeat large screen ambient display as well as the deployments of CityBeat in the newsrooms and concludes with a discussion of the dangers and biases inherited in “Editorial Algorithms.”

Related Work

Researchers have long been studying the intersection of new information tools and journalistic practices (Park et al. 2011; Zubiaga, Ji, and Knight 2013; Schifferes et al. 2014). In this section we review previous research that examined social media and news production, event detection from social media, as well as tools and systems that utilize social media data to inform citizens and city stakeholders.

Journalists have always utilized various information sources to find stories and cover local news. From older sources like local informers to police scanners, journalistic sources recently expanded to include additional real-time information streams that provide clues into the live conditions of our cities. These new sources include various technologies like sensors, video feeds, and drones as well as crowdsourced services (Muthukumaraswamy 2010; Hermida 2010; Pavlik 2013).

The substantial adoption of social media platforms like Instagram and Twitter therefore introduces a new information source that promotes the developments of tools and techniques to study human activity. Previous works that studied social media data mostly focused at activities on a world or a national level (Sakaki, Okazaki, and Matsuo 2010; Weng and Lee 2011), or did not consider real time

aspects (Cranshaw et al. 2012; Cheng et al. 2011). In this realm, considerable work explored the use of social media data to source for local news and information. These works found that local information sourced from social media tends to be geographically-centered to where that information emanated (Yardi and boyd 2010) and that users communicate informational and helpful messages in times of crisis (Heverin and Zach 2012; Shaw et al. 2013).

Over the past few years researchers have developed a number of tools that use social media data to find sources and summarize breaking news events. SRSR (Diakopoulos, De Choudhury, and Naaman 2012) and TweetCred (Gupta and Kumaraguru 2012), for example, helps journalists find and assess sources in Twitter around breaking news events, while other tools like VOX (Diakopoulos, Naaman, and Kivran-Swaine 2010) help newsrooms to extract stories around large scale broadcast events. Other works used sensors as well as crowdsensing information systems, albeit sparse, to get real-time data about various areas of a city (Roitman et al. 2012).

Event detection on social media received considerable attention in recent years. Works in this field looked at various ways to utilize massive flows of information on platforms like Twitter and Flickr to find trending events as they happen (Becker, Naaman, and Gravano 2011; Walther and Kaiser 2013; Marcus et al. 2011; Aiello et al. 2013). Focusing on textual content, previous works detected events by grouping a few significant terms whose frequency of being used in geo-tagged Twitter content change significantly during a short period (Weiler et al. 2013; Graves, McDonald, and Goggins 2014). Other works examined events on social media and tried to understand the contributions of different users (De Choudhury, Diakopoulos, and Naaman 2012) and classified detected events into categories (Ritter et al. 2012).

Our work builds on previous locally focused projects such as Whoo.ly (Hu, Farnham, and Monroy-Hernández 2013) that automatically summarizes hyperlocal information about local events, top users and places. Whoo.ly is meant to provide neighborhood based events list by identifying and grouping trending features in Twitter posts. In CityBeat we focus on a different type of data (Instagram-focused), methods, and visualization goals. In addition there are several commercial tools such as Dataminr (<http://dataminr.com>) and Banjo (<http://ban.jo>) that offer similar service. These services however focus on either on detecting breaking news on a global level (Dataminr) or use human editors to filter detected events.

These algorithms and tools, however, are viewed by journalists, media critics and theorists as "black boxes." As these black boxes control our information consumption we must understand their inherited biases and mistakes (Diakopoulos 2015). In particular, "filtering" decisions of information streams relates to the notion of the Filter Bubble, in which algorithms are presenting information that people already agree with and therefore amplifies biases (Pariser 2011). Following this line of thought, our work provides a real life live system that exemplifies these biases and examines the ways in which Editorial Algorithms detect, curate, and present local news and the implications of these black boxes.

Creating CityBeat

The development of CityBeat consisted of multiple stages, including gathering initial requirements for development, where we performed participants interviewed. Using the major take-aways from the interviews as a guide, we developed the initial design and implementation of CityBeat, and we describe the rationale and technical process for doing that in the next section. Finally, we deployed CityBeat in multiple New York City newsrooms; the next section also reports on lessons learned from the deployment.

To gather initial requirements needed for the development of a hyperlocal social media data aggregation news tool, we conducted a set of semi-structured interviews with journalists, as well as local government employees, administrators, and policy makers from various organizations. During the interviews, we asked the participants about their experiences in gathering city-level insights from social media platforms and any other software or methods. More specifically, we focused on better understanding our interviewees' current tools and practices as well as their unfulfilled needs and requirements. In addition to getting a sense of the existing software and methods used by our participants, we presented them with a set of mock tools and data visualizations to exemplify possible ways to explore and organize real-time geo-tagged data. We used these examples to inquire about their usefulness to our participants' current workflows.

In the beginning of 2013 we conducted a series of semi-structured interviews with 12 participants from various organizations, including *The New York Times*, *The Daily Beast*, and Newark's and Jersey City's mayoral offices (initial results are reported in more length in (Schwartz, Naaman, and Matni 2013)). The interviews were conducted via phone or Skype, and lasted 45 minutes on average. No compensation was provided. Through thematic analysis we examined the interviews for recurring patterns and developed a qualitative coding scheme to group our participants' statements. Based on these results, we identified themes that encompass our interviewees' interactions with city-level social media data. We report these results briefly here, focusing on those that informed the CityBeat system.

The interviews and analysis exposed three key categories of tasks for social media city data: Tracking, Analysis, and Synthesis. Each of these categories portrays a different set of tasks involved in handling city-level social media information. *Tracking* activities include keeping track of real-time local social media to discover new events and happenings; *analysis* describes deep investigation into specific interests, events and topics; and *synthesis* refers to extracting meaningful longer-term conclusions from aggregated data across various social media platforms.

Indeed, extracting news-worthy or news-related events in real-time was the most common social media practice or need, mentioned by eleven out of twelve participants. These participants reported that due to the ubiquity of social media data and the availability of real-time information, they look at this data as an important resource for local information. Tracking activities included detecting events in real-time, monitoring ongoing activity, with the intention of instantly reacting to special scenarios.

Social media was not an exclusive source of data for our participants, but it was becoming integral to their routine. Our interviewees described various other tactics and tools that help them track activities and identify events. For example, five journalists listed sources such as police scanners, news wires and online tools like Google Trends, while two local government employees mentioned local TV channels and news websites. From a social media perspective, eleven of our interviewees considered social media platforms as an integral part of their sources to the discovery of local information. As Lisa (journalist) puts it:

“In the old news room you would sit with your police scanner and you would listen to what the police are talking about and know if something’s going on. Now I ‘listen’ to social media and I watch the progress of how something is happening and that’s very valuable.”

Out of our twelve interviews, ten participants emphasized focusing on data that is unusual. Our participants emphasized the immense value they could gain from data that displays what is “out of the norm.” As Chris (journalist) explains:

“News begins with what’s unexpected or what’s unusual just monitoring is not useful. I mean, it’s useful in the collection of stuff, but seeing what’s being monitored right now is not necessarily useful. What’s useful is what’s different than what’s expected, what’s unusual.”

In this way, finding out first-hand, real-time information about places that display above than average activity levels was highly valuable. As Joe (journalist) emphasizes:

“Saying that most tweets are coming from the Times Square area is not a surprise to me. I don’t really care. But if it was something that was saying this rises above the norm, so that’s a thing that I care about. So, if there is something that’s happening and there’s a frequency that’s aberrant. . . to me that’s a hot-spot. It’s not about the volume.”

As Joe notes, seeing the spread and volume of social media data across the city is not particularly useful. On the other hand, comparing how today’s data differs from those of previous days might prove to be significantly useful. Indeed, two interviewees who work in local government organizations noted that detecting abnormality in the data could help them provide better and quicker service to local communities. As Mike (mayor’s office employee) explains:

“Event detection is super important, so there are certain [events] we would really want to know if people mentioned. And these triggers are super important. So, if someone mentions gunshots, we really want to know that. . . That’s a needle in a haystack that’s really important to us. . . Like, is this thing trending that normally isn’t? Or is this hub suddenly overwhelmed? That would be really important to us because we would then send more police there.”

As there are currently no tools that provide this type of event detection, our participants developed various

workaround methods that involved keyword searches in real-time using Twitter clients such as Hootsuite and TweetDeck and link services like bit.ly. Five interviewees mentioned using existing tools to perform searches based on specific textual strings. For example, Joe (police detective) used names of local gangs to follow tweets about their activity while Lisa (journalist) chose TweetDeck saved searches to display a real-time time line of keywords like “murder” and “shooting.”

“I always have 6, 7, 8 searches running in my TweetDeck and my interns run searches on the bit.ly site. We are looking for people who are talking about a homicide via one of our search terms then we’ll be able to connect to them, we’ll reach out to them. Sometimes we’re able to get an ID of who the victim is early because people have been talking about it on social media and that comes out in our TweetDeck search.”

The use of keywords search provides our participants a way to browse a sub-stream of the information that focuses on a certain issue. As our participants noted, although this method does offer some degree of access to local information, it is limited to a small number of keywords. Moreover, using keywords search retrieves results that are not necessarily related to a certain geographic location. As Chris (journalist) notes:

“What you really want to do is you want to filter it by something that you’re actually paying attention to. To be able to filter, like if you we’re saying, ‘Ok, show me deviations for the norm in a certain area where shooting is involved’. That would be interesting.”

Four of our interviewees tried to overcome this difficulty by using aggregation tools such as GeoFeedia (<http://geofeedia.com/>) that provide search of social media data based on geographical area. However, these participants were not satisfied with these services, as these were not providing local alerts in real-time, and in many cases provided results that were not relevant to their interests.

Finally, seven participants mentioned social media data as source to monitor ongoing activity in the city after an event was detected. In this way, our interviewees follow the live stream of local tweets, pictures, checkins and video clips as they are shared over various social media platforms. Three interviewees mentioned the notion of civic reporting and how people that take pictures on their smartphones can sometimes get a better coverage than that of a reporter.

Initial Design and Feature Takeaways

Some key design takeaways informed the CityBeat system development based on the interviews. In particular, we focused on three key themes mentioned by the participants of our initial requirement gathering stage: detecting hyperlocal news events from raw data, creating an aggregated representation and exposing other deviations from the norm.

As mentioned above, many of our participants were interested in identifying different types of local news events in real time. This type of event detection can play a crucial role in real-time tracking of a city and provide invaluable infor-

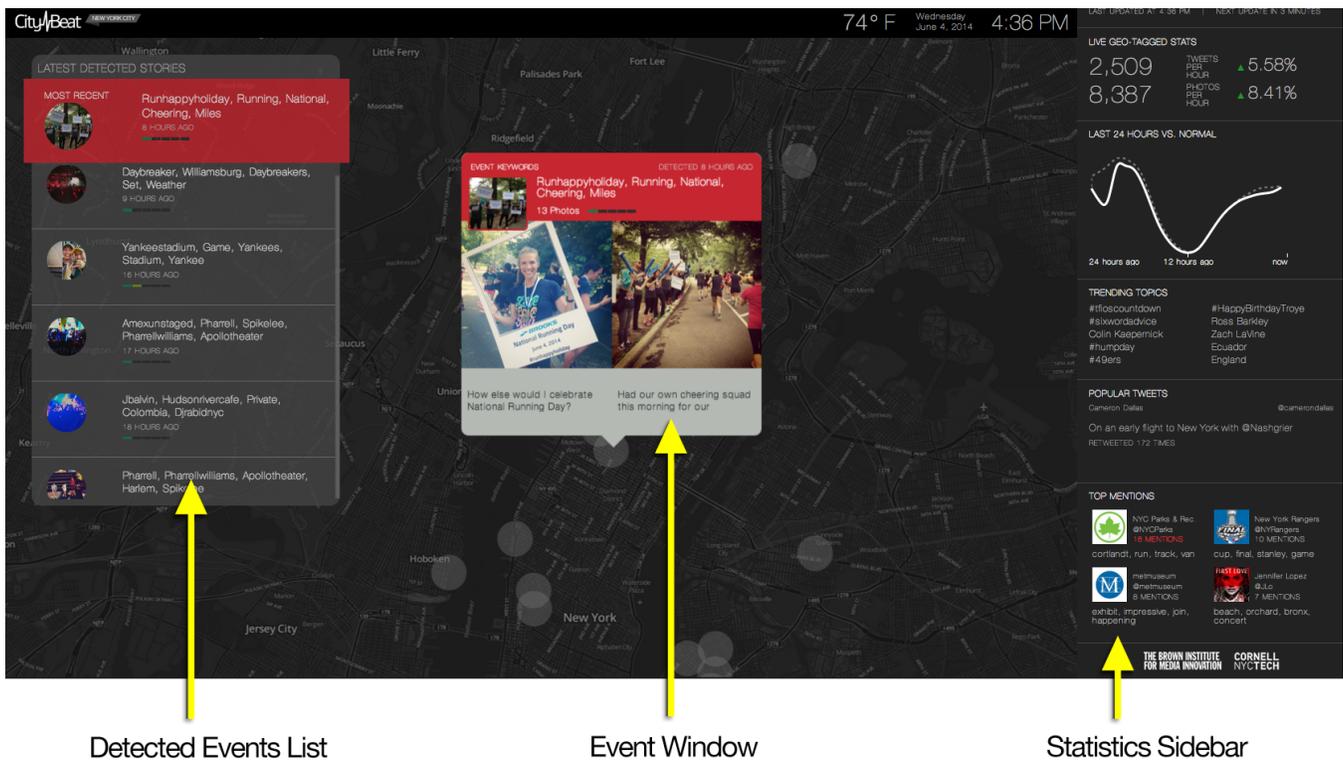


Figure 1: Screenshot of CityBeat interface showing the Detected Events List, Event Window and the Statistics Sidebar

mation for reporters, local government officials and individuals. Events can also be associated with different categories such as entertainment, crime, or breaking news. As various users would want to pay attention to different events, some tools could be specific for a certain type of event or category, while others could learn the users' preferences, or allow the users to manually set their preferences. In CityBeat, we explored the capabilities (and, as we report below, the limitations) of hyperlocal event detection from social media.

The historical social media data of a city can provide a crucial baseline for the discovery of patterns. The ability to compare current volumes and spread of social media data across the city with past statistics can offer a more detailed account to explain deviations in the data. In our design of CityBeat, we included a number of indicators beyond event detection that reflect the difference between historic models and real-time data.

Hyperlocal data aggregation tools can help monitor patterns of people's activity in the city by providing a real-time aggregated representation of the data. Creating a live representation of the city from real-time data requires special attention to the spatio-temporal characteristics of the data, for example, plotting the data over a city map. In addition, tracking live data requires a significant amount of attention. To overcome this difficulty, an aggregation tool could include mechanisms to draw attention to breaking or significant occurrences, or draw the users' attention to any abnormal activity that they set as relevant for them. To this end, we designed CityBeat as an ambient interface with a map-based

component, using visual cues to draw attention to deviations and key elements that require attention from the user. Next, we describe the CityBeat system and its deployment for case studies in several New York City newsrooms.

System Development

The development of CityBeat included several designs and iteration stages based on the requirements detailed above, addressing the need for (1) real-time event detection, (2) aggregate representation, and (3) incorporation of historical data versus live data. In the development process, we worked with a group of journalists as co-developers who provided ongoing feedback on the design on the system described here.

CityBeat is implemented as a web-based ambient visualization, as shown in Figure 1, meant to be shown on large displays in newsrooms. The visualization, based on data from Instagram and Twitter has three main components: the Detected Events List, Event Window, and the Statistics Sidebar.

The Detected Events List, as seen in Figure 1, is based on a social media event detection algorithm using Instagram data, and is visualized as an auto-scrolling list of all events discovered from the data in the past 24 hours. Each detected event on the list is noted on the background map by a pulsing circle placed in a specific geographical location. As the map moves from one location to another at the center of the screen, an Event Window presents a summary of the currently in-focus event including keywords, time of detection,

and relevant photos, all automatically computed by our system. Although meant as an ambient display with limited interaction, one interaction that the system did allow is clicking on an Event Window to go to a more detailed event page where additional event information is presented.

As mentioned above, CityBeat focused on detection and presentation of hyperlocal events. To detect events in a stream of city-level social media data, we devised and implemented an algorithm that is geared to extract hyperlocal events (Xia et al. 2014). While previous work mostly focused on event detection at the global or national level (Chen and Roy 2009; Mathioudakis and Koudas 2010), CityBeat focuses on events that are occurring in a small region, e.g. a street corner or a venue. These events can range from music concerts and exhibitions to emergencies like fires and car accidents.

Performing robust hyperlocal event detection is challenging, given the noise and scale of social media data, where the sparse signal for any location can be easily overwhelmed by irrelevant content. Furthermore, as finding out and following breaking events in real-time is highly important for journalists, our system has to focus not only on robust and precise detection but also on efficient detection in real time (Sakaki, Okazaki, and Matsuo 2010). While the full details of our event detection algorithm can be found (Xia et al. 2014), an overview is provided here for completeness.

The input to the algorithm is the city-wide stream of geo-tagged Instagram data. Instagram is a popular photo sharing service that allows users to add location information to their photo; publicly-available photos with location coordinates can be retrieved using an API. Recent work had shown that Instagram “attention” is more geographically focused than the other sources of unstructured social media Twitter like Twitter (Flatow et al. 2015), and is thus more appropriate for such a task. We collect all geo-tagged photos shared from New York City, averaging about 80,000 items per day (as of December 2014). The system divides the data to sub-regions and models the time series data for each sub-region. An online alert engine compares the region’s real-time data with the time series model for that region. If the engine discovers a significant deviation, a candidate event is created, containing all the photos that contributed to the abnormal signal. In the following step, the candidate event is represented using a vector of spatial, lexical/topical and historical features based on the photos’ data. A trained classifier determines whether the candidate event is an actual event, based on these features.

A final step uses real-time crowdsourcing to improve the classification precision by removing false positives. Candidate events are instantly sent to Amazon Mechanical Turk workers which are then asked to cast their judgment about the existence of an event, the type of event, as well as help curate the relevant photos. In this way detected events are further analyzed to find representative photos and flag important keywords for an event. This crowdsourced judgment, had proven to be problematic. In many instances Amazon Mechanical Turk workers would get confused by the number of different photos that appeared and would classify actual events as noise.

Since the launch of the system, CityBeat detected various events such as conferences, music concerts, outdoor festivals, gallery openings, sports events as well as emergencies like fires. For example, during New York City Fashion Week (September 5 - 12, 2013), CityBeat detected a series of fashion shows and events that took place throughout the city.

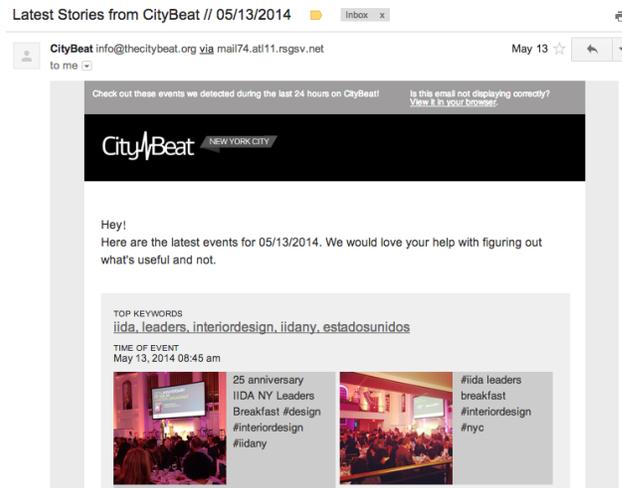


Figure 2: Example of the CityBeat Daily Digest Email

The Statistics Sidebar seen in Figure 1 provides aggregate representation of social media with additional insights into trends of live data and historical data. The sidebar includes the number of tweets and photos posted in the last hour in New York together with the percentage of change from the previous hour, the top Twitter accounts which were mentioned the most during the past hour in Twitter items geo-tagged to New York City, and the five most retweeted tweets that were geo-tagged in New York. We also show the city’s trending topics extracted from Twitter API. These elements intend to give the viewers a different quick idea of the “social media state” of the city.

The Statistics Sidebar also includes a plot of the time series of photos and tweets volume for the city during the past 24 hours. The data from the past 24 hours are visualized by a solid colored curve, while the dashed curve represents the predicted time series (expected, “normal” levels). These two curves, representing the city’s historical data vs. live data, can provide a quick visual indication when some unusual city-wide activity is occurring.

Feedback on the initial implementation led to further design iterations, including the addition of visual elements (e.g., colors, background) that draw the users’ attention to critical or important information in the interface. For example, red font colors call out attention to Twitter accounts that were mentioned much more than normal levels in the Statistics Sidebar.

Deployment and Evaluation

CityBeat was installed and deployed for an extended case study in four New York City-based newsrooms, including *Buzzfeed*, *Gothamist*, *The New York Times*, and *New York*

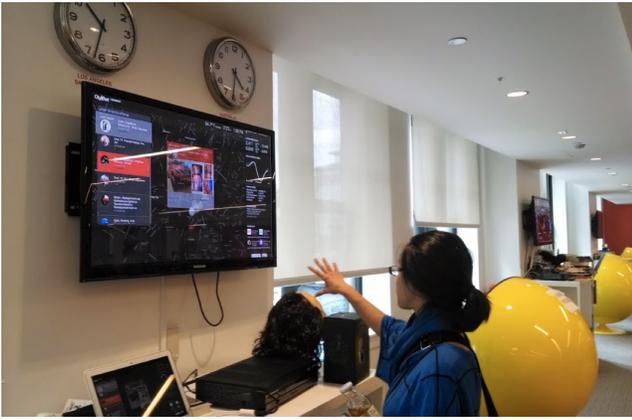


Figure 3: Deployment of CityBeat at the BuzzFeed Newsroom

World. In each organization (other than the Times) we installed CityBeat on a large-screen monitor inside the newsroom for a test period of at least two weeks (see Figure 3). In this setup, CityBeat was visible to editors and journalists, and they could look at the display throughout their working day. Additionally, individual representative of each news organization received the CityBeat daily digest (see Figure 2). The Digest was sent in the form of a daily morning email, listing the latest events identified and retrieved by CityBeat. The email format mirrored the CityBeat visualization, including the events, and for each the top mentioned keywords, the time of the event, and two representative Instagram photos. Overall, approximately twenty working journalists were directly exposed to CityBeat on a daily basis.

To assess the use of CityBeat in actual newsroom settings we conducted semi-structured interviews with six journalists who were exposed to the CityBeat display and received daily emails of the CityBeat digest. Interviews lasted between 17 to 56 minutes. Two of the researchers conducted a thematic analysis of the responses from the interviews. Questions inquired about the overall utility of CityBeat, its content, and its quality as an informative ambient display. We chose to qualitatively evaluate the tool using case studies in normal work context and settings, avoiding the usability testing trap (Olsen 2007). This kind of evaluation was used in previous work as a way to better understand the importance of a system outside of a lab evaluation setting (Wang et al. 2011; Diakopoulos, De Choudhury, and Naaman 2012).

Findings

Based on our participants' responses, some key themes emerged which we discuss and elaborate on next in the context of the system's goals.

The Popularity Bias Our participants reported that CityBeat was heavily biased towards major (albeit local) events. Specifically, three participants mentioned that the news events that CityBeat found provided access only to what a large group of New Yorkers found interesting or, even more critically, worthy of sharing (on Instagram).

This popularity bias was considered as a distortion to the concept of traditional local news which focuses on small, exclusive small-scale stories. Our interviewees pointed out that the emphasis on critical volume (even if not excessively large) in the event detection resulted in the identification of only large-scale events and filtered out the small-scale events that might provide leads for the stories they wanted. As Barry (journalist) explains:

“Let’s say there’s a fire that happens in Brooklyn we want that coverage but then there’s also a pillow fight that happens in Brooklyn. There’s more people taking photos of that pillow fight so the fire is no longer the biggest event and the fire gets pushed back... maybe the fire would mean more to us than the pillow fight or the parade or the concert... the news that we want is not being marked.”

This finding suggests that while the volume of social media data can be a solid indicator of popular interest in, and attendance of, a local event, the system could benefit from classifying and ranking events based on type (e.g., breaking/urgent or not, planned vs. unplanned, small vs. large). Indeed, CityBeat was considered as a biased tool that provides a distorted bird’s eye view of what the most popular events in the city.

Timeliness A common critique of CityBeat was its lack of *real* breaking news events, as well as a lag in the response time. As a news reporting tool, CityBeat failed to display photos and textual content within minutes of an event our participants regarded as news. In the cases when CityBeat *did* detect and displayed relevant breaking news events, it did not do so quickly enough after the event broke. For example, in reference to a Harlem Fire where a water main collapsed into a gas line (Sanchez 2014), Mark (journalist) noted:

“[The Harlem Fire] did show up, but it was half an hour later... at that point we’re not using Instagram.”

Indeed, due to the process and method of detection used, which relied on a certain volume of social media content, news events would only show up between 15 and 30 minutes after the event was first detected. This time lag became crucial in times of urgent, breaking events like the aforementioned building collapse incident.

Not So Breaking News In addition to the desire for timely information, journalists consistently critiqued the “news-worthiness” of the aggregated social media content emerging out of CityBeat. Despite the hope for social media content to reveal stories of the city from an authentic grassroots perspective, journalists sought content that was unplanned and not already on their radar by traditional means such as press releases.

When asked about the ratio of planned events versus unplanned events they observed on CityBeat, all of our interviewees felt the vast majority of the events detected were either planned or public relations events such as conferences, gallery openings, and music concerts. These planned events together with most of the other unplanned events were usu-

ally not interesting from a news perspective. As John (data journalist) explained:

“It’s a great tool for finding social events happening around the city. That’s what appears most often is concerts, parades, that pillow fight...”

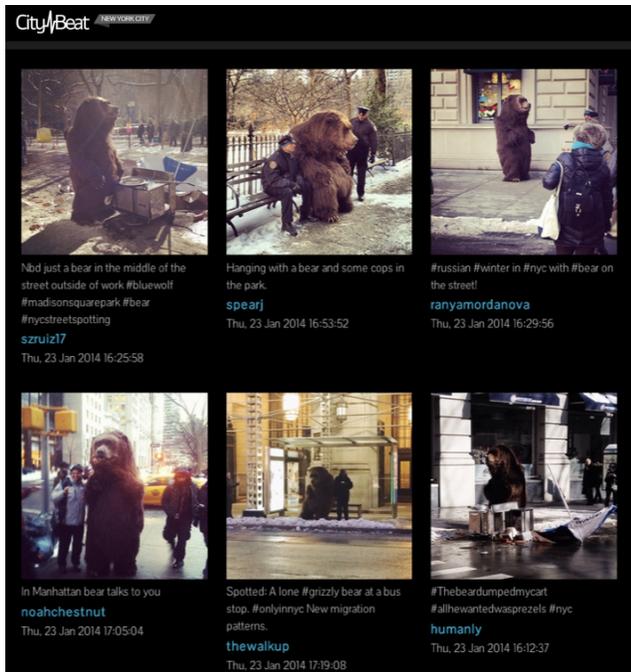


Figure 4: CityBeat detected event: A Person in a bear custom walking around the streets of NYC

Even unplanned events that appeared on CityBeat were not always interesting to the newsrooms. Another editor, Andrea, referencing a person wearing a bear custom walking around the streets of NYC (Figure 4):

“We don’t care about these types of events. Most of these are promotional – run by PR companies anyway. That’s just not the kind of stuff we cover.”

This clear need to parse out the breaking news from the noise exposed the challenges journalists frequently face when trying to use social media tools to pull out relevant stories and information. These tasks become especially difficult for our algorithm in a city like New York, where much of the social media activity revolves around tourist attractions, planned events, and PR events.

Unpacking the Black Box

We use CityBeat here to help think about a future where local news is sourced, curated and presented by machines. We identified various biases and insights while developing and deploying CityBeat as a mockup version of that future. These insights, summarized below, can help better inform the development of future tools and promote critical discussion about the role of Editorial Algorithms in our daily lives.

We discuss the differences between an algorithmically detected event and a newsworthy event, explore the tensions related to crowdsourcing news judgment and emphasize the importance of deployment and evaluation of tools in a real life setting.

Algorithmic Event vs. Newsworthy Event

One of the main challenges while developing CityBeat was in translating into algorithm the notion of a newsworthy event. The type of events that were exposed in the social media data were very different from what was required by the journalists. More specifically, while a cluster of photos and tweets that all shared the same characteristics or semantics was considered an event by CityBeat, in many (most) instances journalists perceived it as noise, unworthy of a news story. The journalists we worked with contended that what citizens determine as “sharable” or “interesting” is very much different than the type of stories news journalists find “print worthy” (we are using the word “print” quite lightly here, of course).

Moreover, there is no one definition of a newsworthy event. Each newsroom has different news interests and reporters who cover the different news stories of the city. In this way, the events that more traditional newsrooms like *The New York Times* found interesting are different from those that are interesting to newer newsrooms such as *Buzzfeed* or cultural media outlets such as *TimeOut New York*. In this context, an additional challenge to be addressed is how to identify events that match conceptions and interests of editors and journalists in different newsrooms and media venues.

Demography and Geography Matter

It is not a surprise that geo-tagged social media data is heavily biased towards highly populated and affluent urban areas, and is far from uniformly distributed even in urban settings (Hecht and Stephens 2014; ?). Moreover, detecting hyper-local events based on volume presents an additional challenge for under-represented areas due to the scarcity of the data. As a result, the detected events are usually situated in very dense areas of the city and therefore show a skewed image of the type of events and activities that are taking place in other areas of the city that are not as well-represented, e.g. due to socioeconomic factors. Working with social media data to produce news has the potential to sustain and amplify these biases.

Crowdsourcing News Judgment

The implementation of crowd-sourced news judgment as part of CityBeat was highly controversial. Journalists objected to the use of Amazon Mechanical Turk workers due to the workers’ lack of “journalistic intuition.” This lack of understanding what can be considered as news, would lead to the wrong classification of events and the annotation of certain public relations or marketing events as important. As part of the integration of the crowdsourced layer to CityBeat, the journalists insisted on getting access to the full set

of results as annotated by the crowdsourced workers. By incorporating and displaying the full judgments of the crowdsourced workers, the journalists were more inclined to accept the results and use them as another way to signal the significance of a certain event.

Lab Evaluation vs. Evaluation Out In The Wild

There is a considerable difference between the feedbacks received while evaluating the tool in a lab setting vs. the feedback from participants after using it for a test period in a live newsroom setting. In our initial lab evaluation we would mostly receive positive feedback about the design and form factor of the tool. Our participants were positive and excited about the possibilities this tool will provide them.

Due to the live nature of the tool, it was difficult to judge if the events that we detected were useful as the type of events changes throughout the days of the weeks and the hours of the day. Moreover, not everyday was there a breaking news event that would demonstrate the way the system works. As a result, the feedback we received was lacking and tended to “confirmation bias” (tendency of people to confirm the validity of the tool). Only by testing the tool in a real newsroom setting over a certain period of time, we received a genuine feedback that listed the pitfalls and biases of the tool. Without having our participants incorporate the tool into their daily work environment and have them make sense and extract insights from it, we would have not be able to get this kind of valuable feedback.

Conclusions

In CityBeat, we experimented with outsourcing local news to machines. To do so, we built a tool that utilizes information from various publicly available social streams (Instagram and Twitter) and employs algorithmic judgment to find and decide what is local news. The goal of this paper was to showcase through a live, fully functional system the inherited biases and dangers that are embalmed into the design of Editorial Algorithms.

CityBeat should be considered as an example and a reference point for future projects and ongoing efforts to build new journalistic tools. We emphasize the need for translation of critical thinking and journalistic concerns around automatic news publishing, into a more functional (and aware) system that can showcase and formalize these issues into a visible tool. We hope that our experience and system provides an important reference to the critical discussion about the growing role of news algorithms in our lives.

Acknowledgements

This work is supported in part by the National Science Foundation CAREER award 1446374.

References

Aiello, L. M.; Petkos, G.; Martin, C.; Corney, D.; Papadopoulos, S.; Skraba, R.; Goker, A.; Kompatsiaris, I.; and Jaimes, A. 2013. Sensing Trending Topics in Twitter. *IEEE Transactions on Multimedia* 15(6):1268–1282.

Becker, H.; Naaman, M.; and Gravano, L. 2011. Beyond trending topics: Real-world event identification on Twitter. In *Proceedings of ICWSM '11*.

Chen, L., and Roy, A. 2009. Event detection from flickr data through wavelet-based spatial analysis. *Proceeding of CIKM '09*.

Cheng, Z.; Caverlee, J.; Lee, K.; and Sui, D. 2011. Exploring Millions of Footprints in Location Sharing Services. *Proceedings of ICWSM '11*.

Cranshaw, J.; Schwartz, R.; Hong, J.; and Sadeh, N. 2012. The livelihoods project: Utilizing social media to understand the dynamics of a city. *Proceedings of ICWSM '12*.

De Choudhury, M.; Diakopoulos, N.; and Naaman, M. 2012. Unfolding the event landscape on twitter. *Proceedings of CSCW '12*.

Diakopoulos, N.; De Choudhury, M.; and Naaman, M. 2012. Finding and assessing social media information sources in the context of journalism. *Proceedings of CHI '12*.

Diakopoulos, N.; Naaman, M.; and Kivran-Swaine, F. 2010. Diamonds in the rough: Social media visual analytics for journalistic inquiry. *2010 IEEE Symposium on Visual Analytics Science and Technology*.

Diakopoulos, N. 2015. Algorithmic accountability: Journalistic investigation of computational power structures. *Digital Journalism*.

Farhi, P. 2014. Charting the years-long decline of local news reporting. *The Washington Post*.

Flatow, D.; Naaman, M.; Xie, K. E.; Volkovich, Y.; and Kanza, Y. 2015. On the accuracy of hyper-local geotagging of social media content. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining, WSDM '15*, 127–136. New York, NY, USA: ACM.

Goldberg, D.; Corcoran, M.; and Picard, R. G. 2013. Remotely piloted aircraft systems & journalism: opportunities and challenges of drones in news gathering. *Reuters Institute for the Study of Journalism. University of Oxford*.

Graves, I.; McDonald, N.; and Goggins, S. P. 2014. Sifting signal from noise: A new perspective on the meaning of tweets about the big game. *new media & society*.

Gupta, A., and Kumaraguru, P. 2012. Credibility ranking of tweets during high impact events. *Proceedings of PSOSM '12*.

Hecht, B., and Stephens, M. 2014. A Tale of Cities: Urban Biases in Volunteered Geographic Information. *Proceedings of ICWSM '14*.

Hermida, A. 2010. Twittering The News. *Journalism Practice* 4(3):297–308.

Heverin, T., and Zach, L. 2012. Use of microblogging for collective sense-making during violent crises: A study of three campus shootings. *J. Am. Soc. Inf. Sci.* 63(1):34–47.

Hu, Y.; Farnham, S. D.; and Monroy-Hernández, A. 2013. Who.ly. *Proceedings of CHI '13*.

Hunter, I. 2014. New york times chief data scientist chris wiggins on the way we create and consume content now. *Fast Company*.

- Marcus, A.; Bernstein, M. S.; Badar, O.; Karger, D. R.; Madden, S.; and Miller, R. C. 2011. Twitinfo:Aggregating and Visualizing Microblogs for Event Exploration. *Proceedings of CHI '11*.
- Mathioudakis, M., and Koudas, N. 2010. TwitterMonitor. *Proceedings of SIGMOD '10*.
- Muthukumaraswamy, K. 2010. When The Media Meet Crowds Of Wisdom. *Journalism Practice* 4(1):48–65.
- Olsen, D. R. 2007. Evaluating user interface systems research. *Proceedings of UIST '07*.
- Pariser, E. 2011. *The Filter Bubble: What the Internet Is Hiding from You*. Penguin Group.
- Park, S.; Ko, M.; Kim, J.; Liu, Y.; and Song, J. 2011. The politics of comments. *Proceedings of CSCW '11*.
- Pavlik, J. V. 2013. *Journalism and new media*. Columbia University Press.
- Ritter, A.; Mausam; Etzioni, O.; and Clark, S. 2012. Open domain event extraction from twitter. *Proceedings of KDD '12*.
- Roitman, H.; Mamou, J.; Mehta, S.; Satt, A.; and Subramaniam, L. 2012. Harnessing the crowds for smart city sensing. *Proceedings of CrowdSens '12*.
- Sakaki, T.; Okazaki, M.; and Matsuo, Y. 2010. Earthquake shakes Twitter users. *Proceedings of WWW '10*.
- Sanchez, R. 2014. Nine still missing after manhattan explosion leaves at least 4 dead, 63 hurt. <http://www.cnn.com/2014/03/12/us/manhattan-building-explosion>.
- Schifferes, S.; Newman, N.; Thurman, N.; Corney, D.; Göker, A.; and Martin, C. 2014. Identifying and Verifying News through Social Media. *Digital Journalism* 1–13.
- Schwartz, R.; Naaman, M.; and Matni, Z. 2013. Making sense of cities using social media: Requirements for hyper-local data aggregation tools. In *Proceedings of the WCMCW at ICWSM '13*.
- Shaw, F.; Burgess, J.; Crawford, K.; and Bruns, A. 2013. Sharing news, making sense, saying thanks: Patterns of talk on Twitter during the Queensland floods. *Australian Journal of Communication* 40(1).
- Stencel, M.; Adair, B.; and Kamalakanthan, P. 2004. The goat must be fed: Why digital tools are missing in most newsrooms. *A report of the Duke Reporters Lab*.
- Walther, M., and Kaisser, M. 2013. Geo-spatial Event Detection in the Twitter Stream. *Lecture Notes in Computer Science* 356–367.
- Wang, T. D.; Wongsuphasawat, K.; Plaisant, C.; and Shneiderman, B. 2011. Extracting Insights from Electronic Health Records: Case Studies, a Visual Analytics Process Model, and Design Recommendations. *Journal of Medical Systems* 35(5):1135–1152.
- Weiler, A.; Scholl, M. H.; Wanner, F.; and Rohrdantz, C. 2013. Event identification for local areas using social media streaming data. *Proceedings of the ACM SIGMOD Workshop DBSocial '13*.
- Weng, J., and Lee, B.-S. 2011. Event Detection in Twitter. In *Proceedings of ICWSM '11*.
- Xia, C.; Schwartz, R.; Xie, K.; Krebs, A.; Langdon, A.; Ting, J.; and Naaman, M. 2014. Citybeat: real-time social media visualization of hyper-local city data. In *Proceedings of the companion publication of the 23rd international conference on World wide web companion*.
- Yardi, S., and boyd, d. 2010. Tweeting from the Town Square: Measuring Geographic Local Networks. *Proceedings of ICWSM '10*.
- Zubiaga, A.; Ji, H.; and Knight, K. 2013. Curating and contextualizing Twitter stories to assist with social newsgathering. *Proceedings of IUI '13*.