Utilising Social Media to Identify the Cultural Loci of the City

Laurence Elsdon

Northumbria University

Contents

01 Introduction

Previous methods of research into the organisation of cities and patterns of user behaviour within them have typically required lengthy surveys and interviews with a variety of users for instance Kevin Lynch's Image of the City (Lynch, 1964) derived mental maps of the city by interviewing city residents, requiring them to draw maps, review photographs and walk the routes (Jiang, 2012). This was resource intensive and took a considerable time to produce what were only qualitative results.

New methods of research have been opened through advances in computing power and increasing availability of detailed databases of geospatial data which have enabled researchers to algorithmically derive quantitative results, for example, Jiang's (2012) quantitative image of the city.

Other researchers within the fields of space syntax, spatial cognition, spatial configuration and the broader fields of architecture, urban planning, and urban design have achieved similar feats using algorithms to quantitatively evaluate data sets to derive insights into human behaviour.

Within this paper it be will argued that without any prior knowledge the cultural loci of the city can be identified by analysing geolocated social media messages.

02 Background

Mapping Social Media

Aggregation and analysis of collective social actions has only become possible in recent years but has already been exploited by a number of researchers to identify patterns in human behaviour within downloaded databases of social activity.

Cranshaw et al. (2012) identified that people's activity patterns could be analysed in order to generate a new model of "municipal organizational units" (Cranshaw et al., 2012). These livehoods – unlike traditionally established neighbourhoods – are dynamic definitions that change, they represent how people really perceive a city today, not the outdated opinion of the original cartographers.

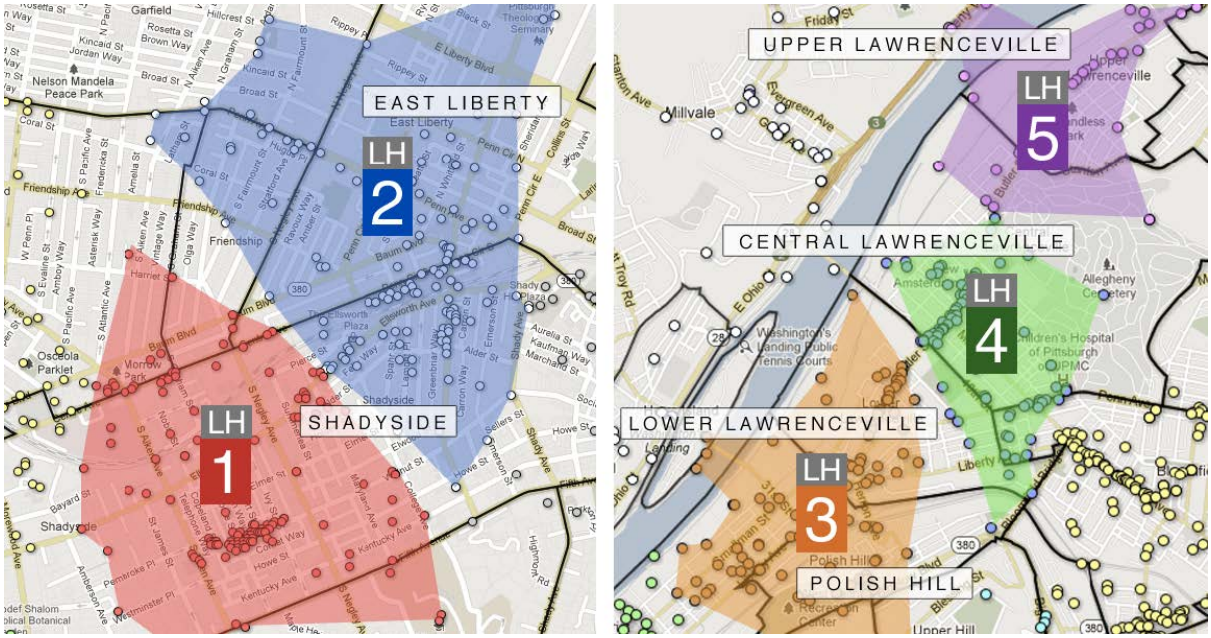

Figure 1 – Livehoods – The municipal borders (black) and Livehoods for Shadyside/East Liberty (Left) and Lawrenceville/Polish Hill (Right) (Cranshaw et al., 2012)

Cranshaw's livehoods are defined based upon clusters of *check-ins* on Foursquare – a location based social network. Foursquare's users opt to publically check-in to places so their friends and followers can see where they are and where they've been, this leaves a trail of geolocated flags. Cranshaw et al. downloaded a data set of all check-ins for the city of Pittsburgh then grouped the data into clusters which defined characteristic areas (see Figure 1 & Figure 2).

Figure 2 – Livehoods – The municipal borders (in black) and Livehoods for South Side Pittsburgh (Cranshaw et al., 2012)

Twitter

Twitter is a social media platform that enables anyone to create an account to read or post public messages. The content of these messages – or tweets as Twitter messages are known – vary greatly, often it is used by companies to advertise their products and engage with their users but it's also used by general people to discuss their daily activities. While the individual messages themselves are rarely of interest to academics the vast quantity of messages mean trends and patterns can be identified.

Figure 3 –Data for 6 January 2010 12Z, the map on the left shows postcode sectors where there is snow on the ground using Meteorological Office data. The map on the right shows the #uksnow counts from Twitter. (Grey et al., 2015)

In 2010 during a prolonged period of snowfall Twitter users began tweeting – the verb for sending a tweet – about the level of snow in their area. By collecting tweets that mentioned the first part of their

postcode and a rating from 0 to 10 for the level of snow Marsh (2011) was able to create an interface to map snowfall across the UK (see Figure 3).

Twitter API

Researchers at the Centre for Advanced Spatial Analysis (CASA) at UCL have developed a number of experiments and toolkits which exploit social media data (Grey et al., 2015) one of their early experiments – Tweet-o-Meter (CASA, no date) – involved measuring the number of messages posted to Twitter in eight cities every minute in real time (see Figure 4).

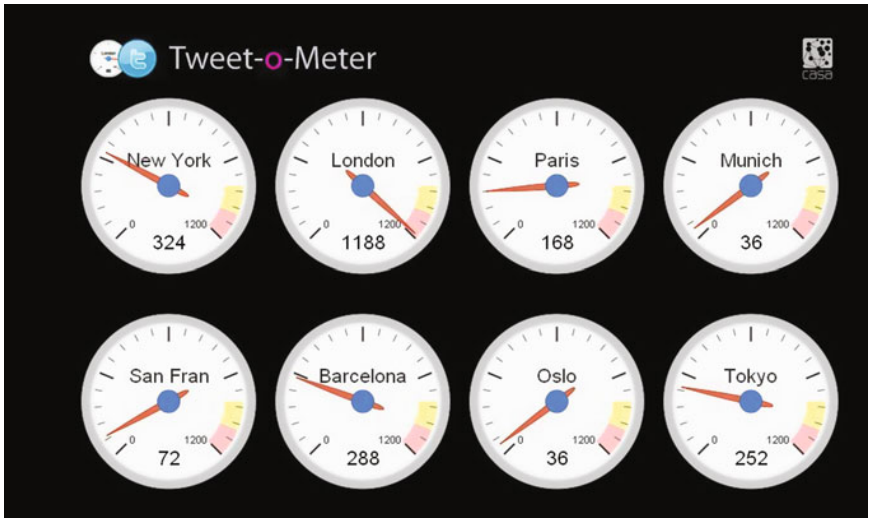

Figure 4 – Tweet-o-Meter real-time gauges for 8 cities. (Grey et al., 2015)

This is made possible by Twitter's public Application Programming Interface (API), which enables researchers to retrieve the latest tweets in a format suitable for computer applications to read. It also enables the researcher to restrict the returned data set. For example, you may only want tweets mentioning a certain topic, or in the case of the CASA's Tweet-o-meter research to have the results restricted to only geolocated messages less than 30km from the centre of London. Geolocation of tweets is optional but many users choose to enable the feature on their mobile phones, particularly when talking about places they've been. Twitter then use the phone's GPS to record the latitude and longitude.

Mapping Tweets

Chesire & Manley (2013) used the Twitter API to collect over ten million tweets geolocated within 30km of London over 3 years and then determined the language of each by running them through the Google Translate API. The resulting map (see Figure 5 & Figure 6) shows the cultural diversity of London.

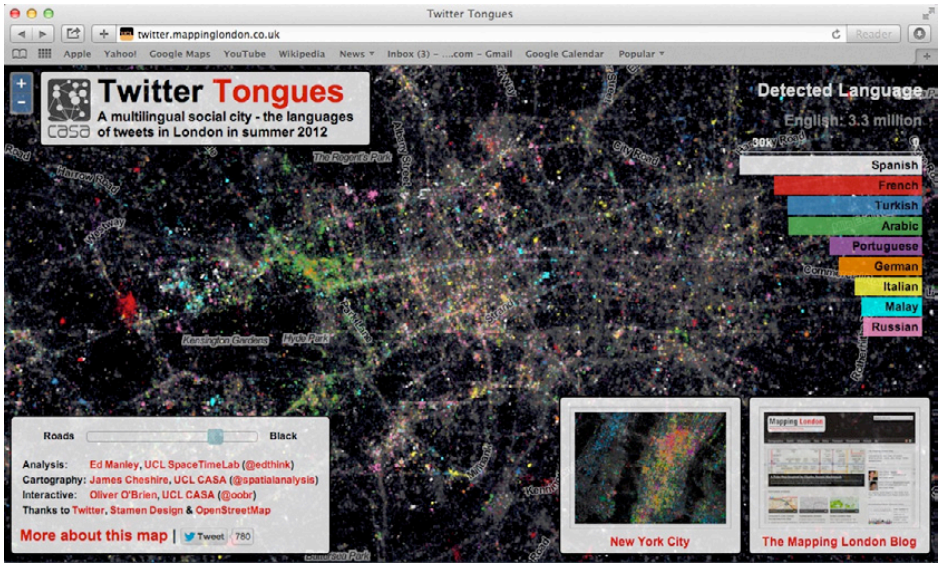

Figure 5 – New Media: The spatial distribution of tweets related to the distribution of ethnic groups in London from an analysis of language content of tweets (Batty & Hudson-Smith, 2014)

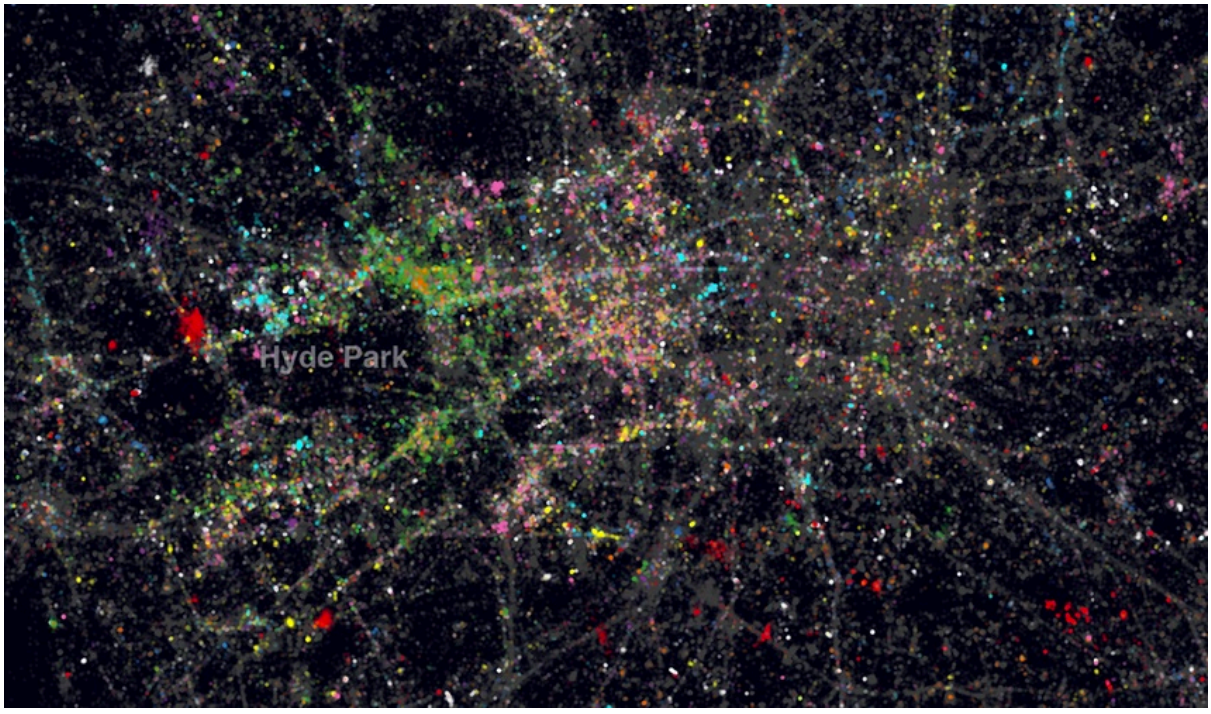

Figure 6 – The spatial density of tweets in London tweets by language where the grey foundation of the map is formed from the majority of tweets in English. Other nationalities represented by colours (Batty et al., 2013)

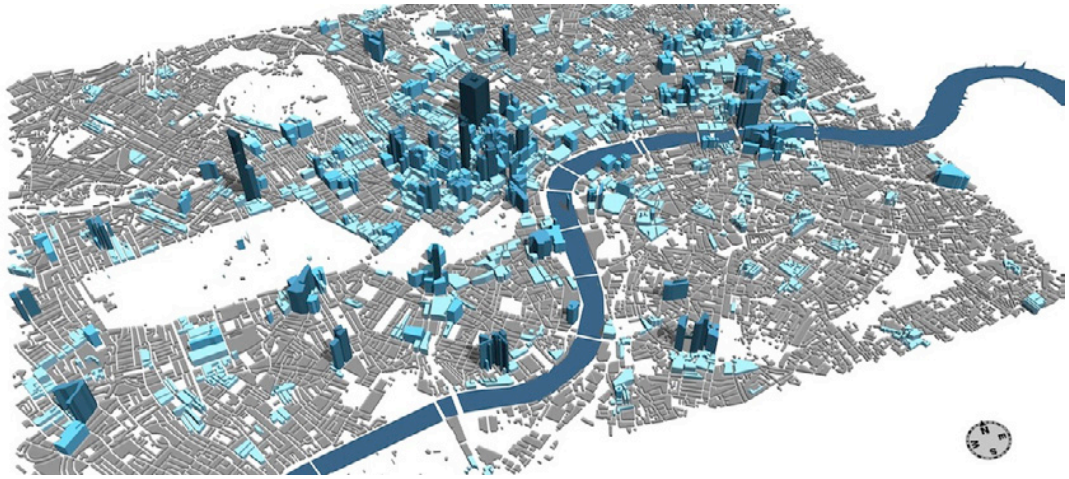

Figure 7 – New Media: The spatial distribution of tweets tagged to a 3d model of London, varying over 15 hours from 3500 tweets: buildings grow in proportion to the 'data' they generate

Hugel & Roumpani (2013) – of CASA, published in Batty & Hudson-Smith (2014) – present a method of associating tweets with specific buildings which they then visualise in 3d (see Figure 7) to show the volume of tweets across the city throughout a single day.

Culture

The chair's foreword for the World Cities Culture Report (World Cities Culture Forum, 2015) argues that culture offers an incredible return on investment:

> [Culture] generates billions in cash, is the reason tourists visit, breathes life into run-down parts of town, gives cities their unique character and bolsters international standing. [Culture] improves health and wellbeing, offers criminals a new perspective and builds the skills and confidence of future generations. It costs a fraction of other city budgets. And it makes life worth living.
>
> Simons (2015)

Culture is integral to urban planning and policy (World Cities Culture Forum, 2015, p. 8), plays an important role in the economic growth of cities which in turn drives culture and creativity (World Cities Culture Forum, 2015, p. 8) but "as budgets tighten it is often the thing that gets dropped" (Simons, 2015) because of the difficulty in quantifying its impact.

The intent of the World Culture Report is to measure the culture of its member cities and discuss their culture through interviews with local ambassadors. The report quantitatively ranks the world cities through a series of infographics (see Figure 8 to Figure 12) on the number of cultural institutions – theatres, museums, cinemas, restaurants and so on – each contain. While this is a good indicator of the level of culture within a city it does not measure the cultural engagement of the city's populace and does not investigate the cultural map of the city.

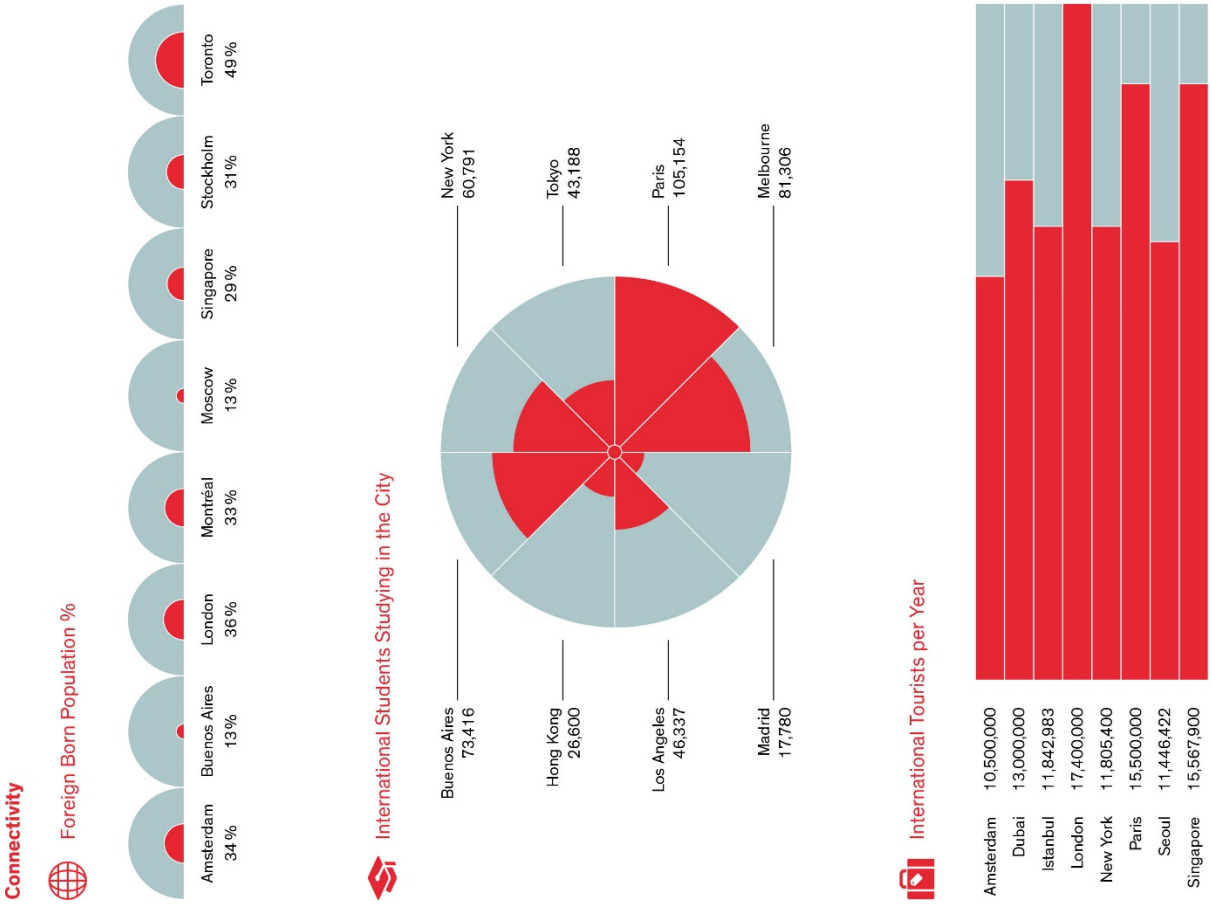

Figure 8 – Culture infographics (World Cities Culture Forum, 2015)

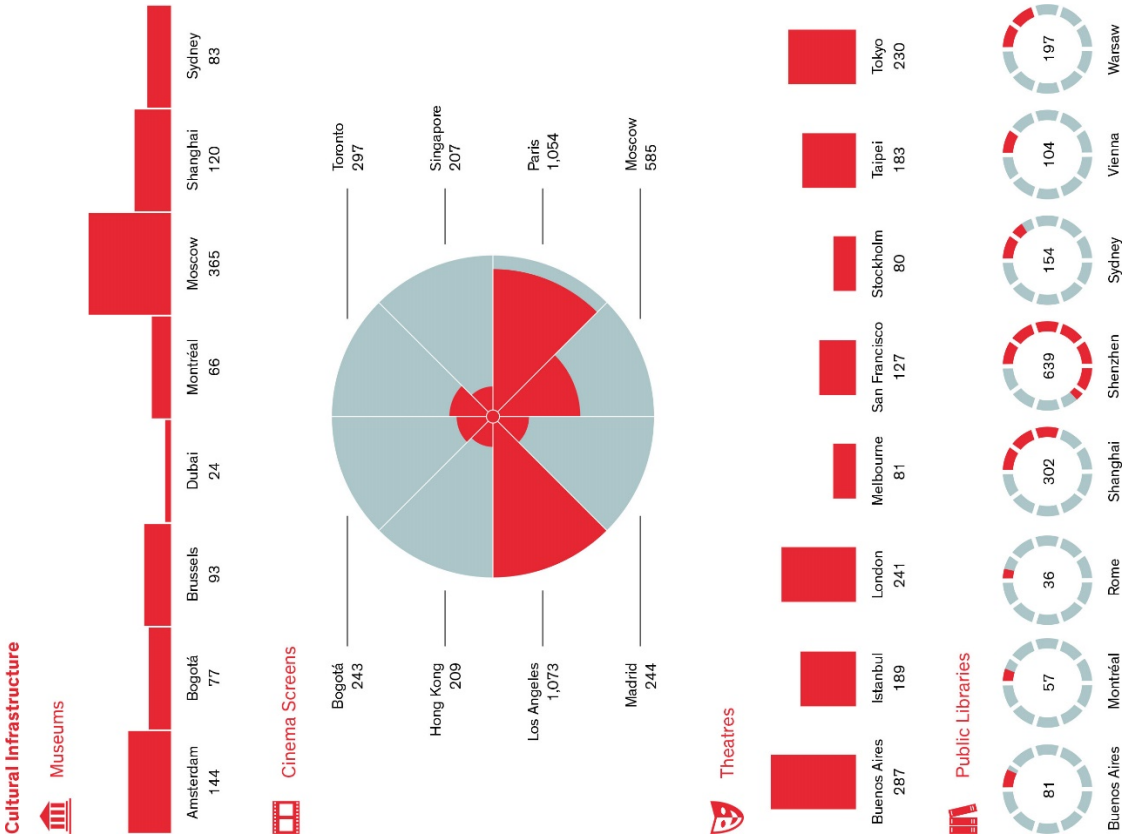

Figure 9 – Culture infographics (World Cities Culture Forum, 2015)

Figure 10 – Culture infographics (World Cities Culture Forum, 2015)

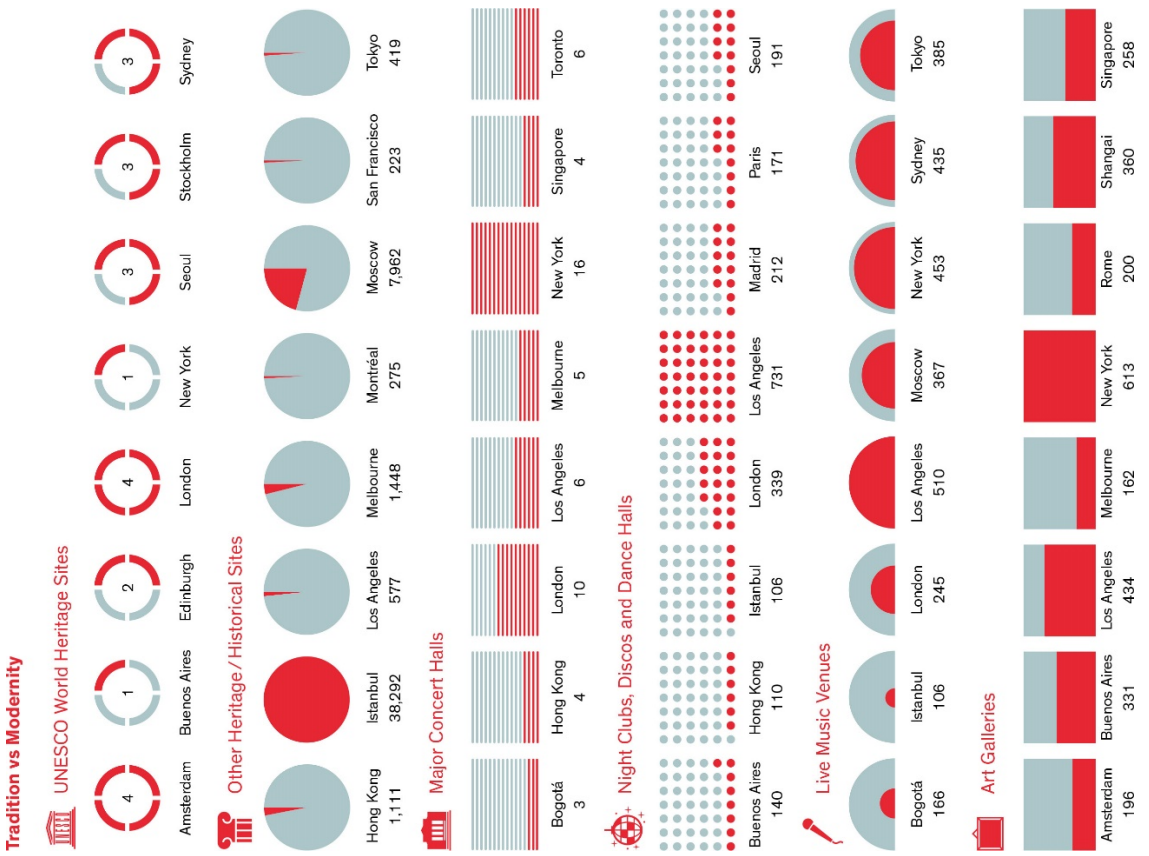

Figure 11 – Culture infographics (World Cities Culture Forum, 2015)

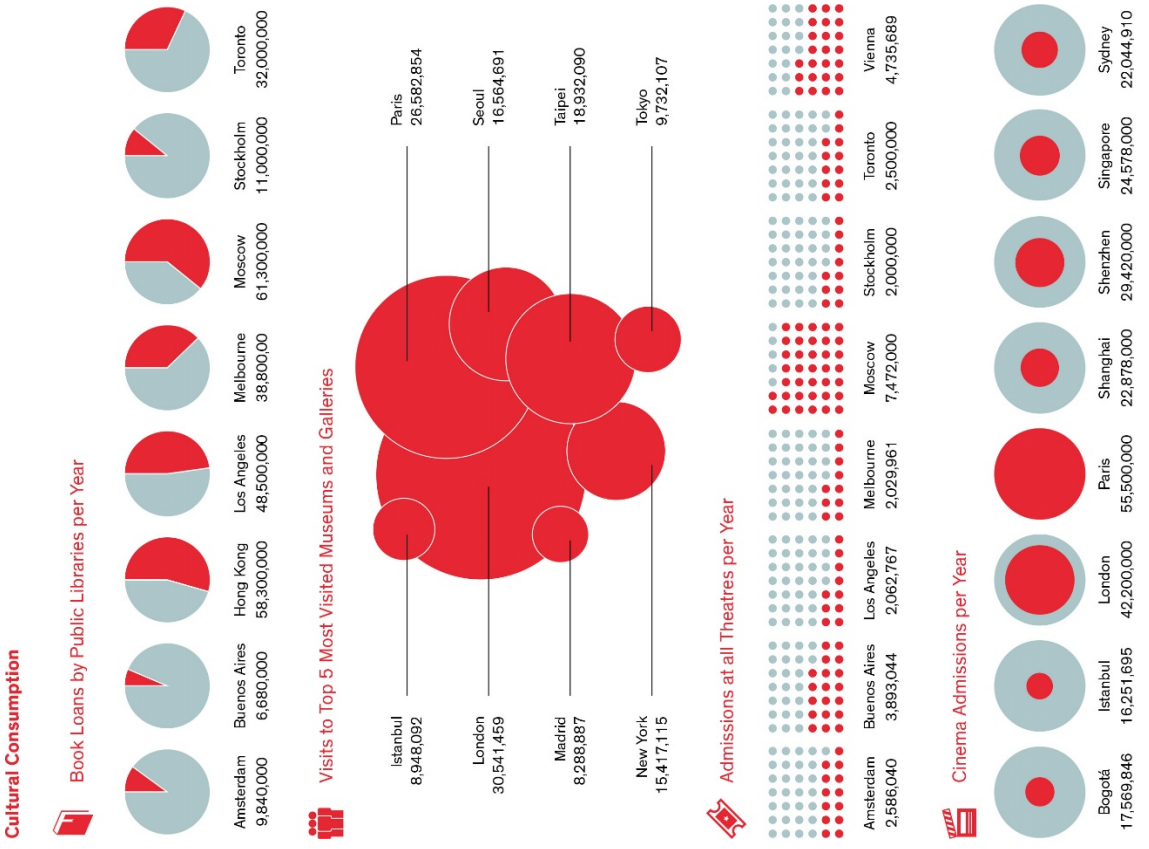

Figure 12 – Culture infographics (World Cities Culture Forum, 2015)

03 Proposal

Aims

To identify the cultural loci of the city by analysing geolocated social media messages.

Objectives

The overall objective will be to develop a data collection, analysis and visualisation application to complete the following objectives:

1. To retrieve the latest tweets that meet the following criteria:
    a. Are geolocated within a specific set of bounds – e.g. within 30km of the centre of London.
    b. Contain a cultural keyword – e.g. “museum”, “gallery” or “art”
2. To visualise the tweets on a map at their corresponding longitude and latitude
3. To group tweets on similar cultural topics visually
4. To analyse the resulting map and discuss whether it successfully shows the distribution of cultural loci within the city.
5. To ensure the validity of the results by making the application reusable by future researchers.

04 Methodology

In an experimental methodology this research will combine elements of Cranshaw et al.'s (2012) Livehoods research and Chesire & Manley's (2013) work on geolocating tweets within the city. The experimental methodology can be further broken down into two phases, build and evaluation.

Build Phase

Objectives 1, 2 and 3 constitute the build phase of the application, these will be accomplished through the development of an application. The User Experience (UX) wireframe illustrating how the user will interact with the application is shown in Figure 13.

Figure 13 – UX wireframe

The application will be formed of three pages, the first – criteria – accepts the user inputs, the city and the search terms. Once they hit compute the tweets will be searched and then returned on the second – schedule – which will show a table of results, the third screen – map – will show the visual distribution of geolocated tweets, with each marker coloured dependent upon its cultural topic.

Application Architecture

Similar to Cheshire & Manley's (2013) (Figure 5) map the application will be developed as a web application written in the JavaScript language, the application will take two parts a frontend and a backend (Figure 14).

The frontend serves as the user interface of the application, to aid development it will be based on AngularJS a popular framework developed by Google. The frontend runs within the user's browser and is the only part of the application they will see.

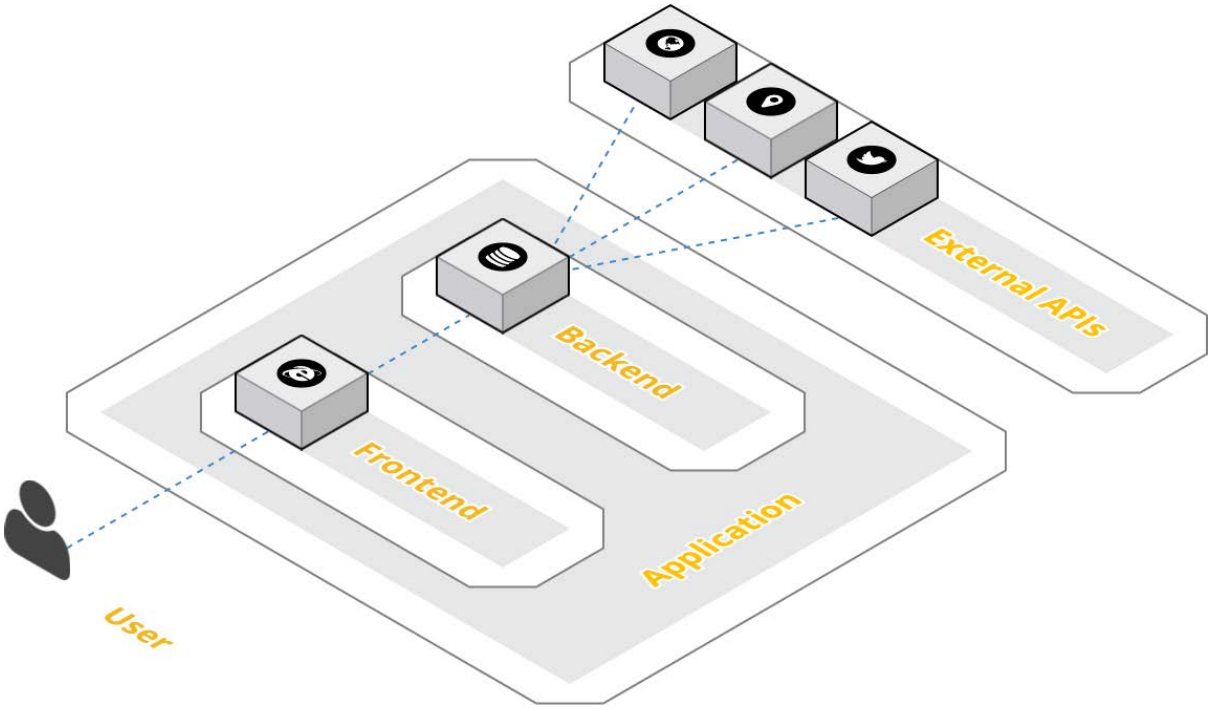

Figure 14 – Application architecture generated with CloudCraft

The backend is the hidden server element of the application, this receives requests from the frontend, makes API calls to the external endpoints, processes the responses and returns the results to the frontend application which visualises the data. The backend will be based on the Node.js platform.

Objective 1

The first input variable we require from the user is the name of the city. The name of the city will be translated through the Google Places API into longitude and latitude values. Next we require the search term or terms, which in combination with a search radius and longitude and latitude are the required parameters of the Twitter API call. Due to security limitations the API call cannot be made from the frontend – referred to as client side – it must be done using a backend server which acts as a middleman (Figure 14).

The frontend will send the request parameters to the backend server which will make the call to the Twitter search API endpoint which will return a maximum of 100 results and only results from the last 7 days. Therefore multiple asynchronously calls must be made in order to achieve a full set of results.

The application will request the first set of tweets, if there are exactly 100 then the next set will be requested, and so on until fewer than 100 results are returned, at which point all sets of tweets are

aggregated and sent to the frontend. Therefore, if there are 750 tweets for the search term “museum” then we will have to make 8 recursive API calls to obtain them all. Additionally, as we will accept multiple search parameters we will have to perform an unknown number of recursive calls for each of them. This is important to note because the API has a rate limit of 180 requests per 15 minute window (Twitter Developers, 2016).

Objective 2

Once objective 1 is completed, the frontend application should be in possession of a list of tweets, the objective will be to visualise them. First, the total counts will be quantitatively shown in a table then they will be visualised on a map.

The Google Maps JavaScript API provides methods to embed standard Google Maps within websites and allows the map to be styled and for markers to be placed. The map will be styled monochromatically showing only key rivers and roads to reduce visual clutter and give the geolocated markers more prominence

Objective 3

Similar to Chesire & Manley (2013) small coloured circles will be used to mark the geolocations of each tweet, clicking on the marker will dynamically open a popup revealing the related tweet. The colour of the marker will depend on the category of the search term to ensure for instance both “art” and “gallery” will be grouped together with the same colour.

Evaluation Phase

Objective 4

The fourth objective will analyse the results produced by the application to see whether they answer the aim of the research – to identify the cultural loci of the city by analysing geolocated social media messages. The theory is that within the resulting map there should be identifiable clusters of coloured markers. The clustering of markers – much like Cranshaw’s (2012) Livehoods – will represent cultural loci.

The analysis will be completed on a variety of cities in order to see the effectiveness of the algorithm across different city types.

Objective 5

Amaral et al. (2011) note that when reporting experimental research in computer science the reproducibility of the experiment is often forgotten this is, however, an essential part of academia. The application should therefore be published alongside the results so that an academic may attempt to complete the experiment themselves.

05 Application

A live version of the application is available at:

http://tweetmap.studiole.uk

The application has only been tested in the latest version of Google Chrome (Version 50.0.2661.87 m), other browsers may not render correctly. If you have any issues with using the application or find it to be temporarily unavailable please contact the author directly for a demonstration.

The code for the application is split across two repositories – frontend and backend – available at:

https://github.com/StudioLE/TweetMapFrontend

https://github.com/StudioLE/TweetMapBackend

Figure 15 to Figure 20 show a series of screenshots illustrating the functionality.

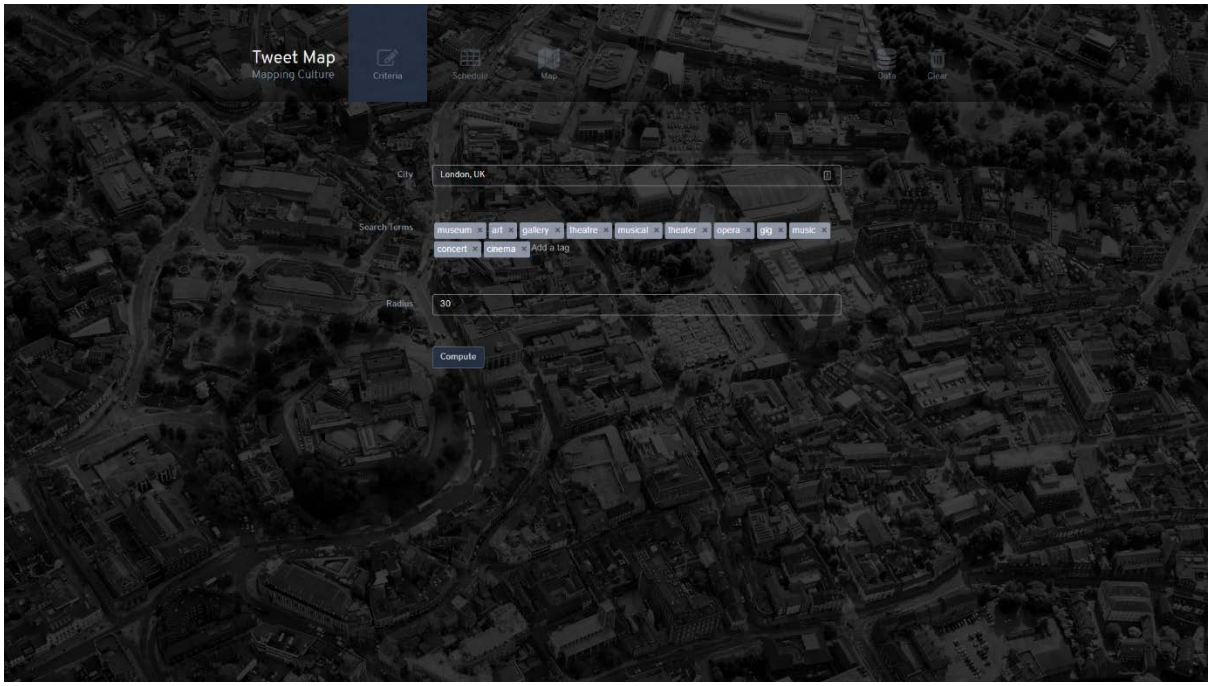

Figure 15 – Criteria page

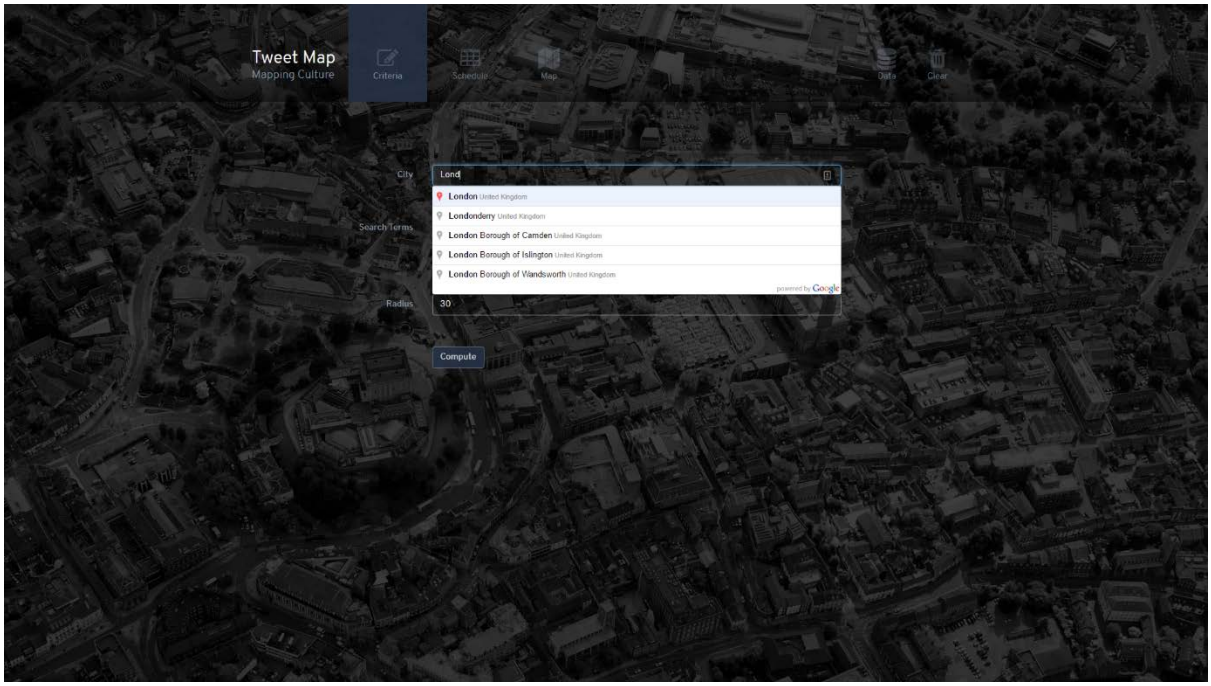

Figure 16 – Criteria page showing autocomplete dropdown

Figure 17 – Schedule page

Figure 18 – Schedule page showing options to change marker colour

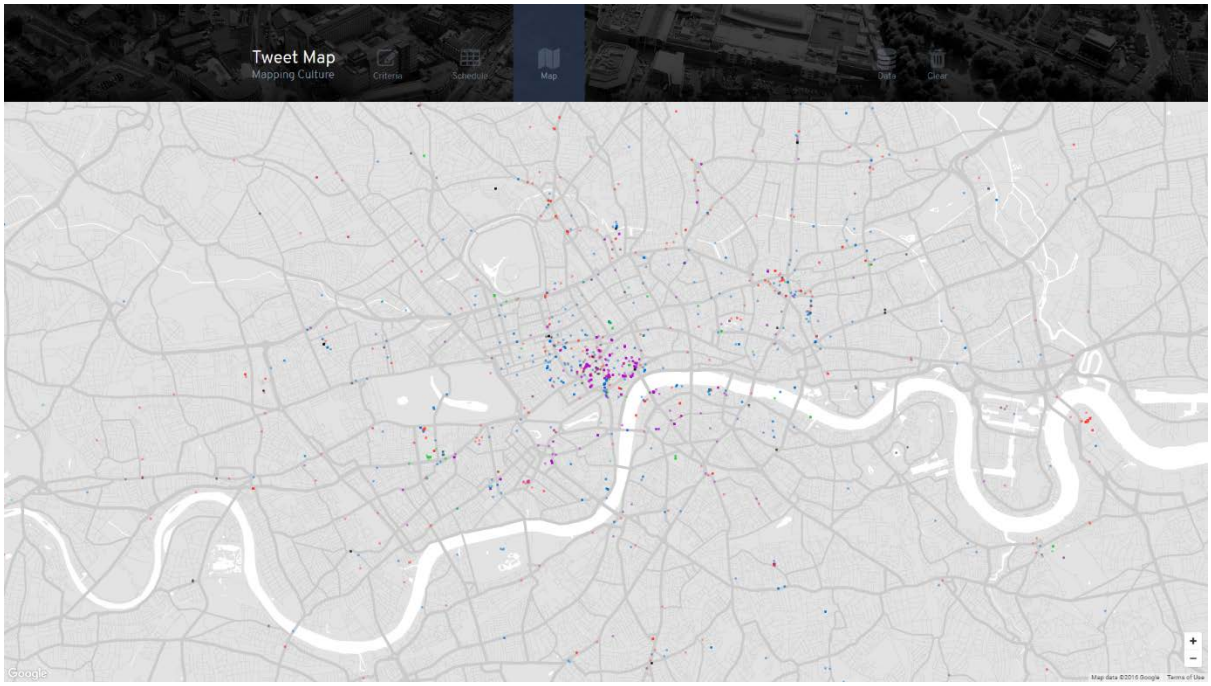

Figure 19 – Map page

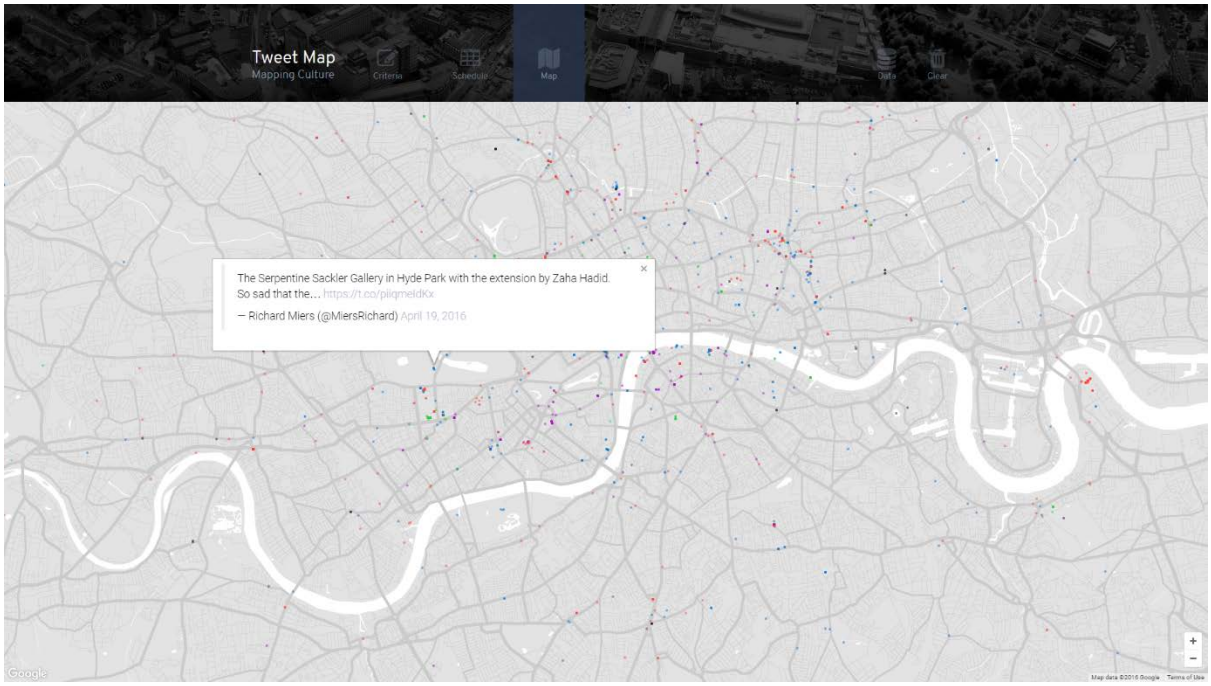

Figure 20 – Map page with Tweet Pop-up

06 Results

Following a series of experiments into suitable candidates five categories of cultural terms were defined to investigate the social culture of the city: museums; art; theatre; music and; cinema.

Category	Marker colour	Search terms
Museum	Green	“museum”
Art	Blue	“art”, “gallery”
Theatre	Purple	“theatre”, “theater”, “musical”, “opera”
Music	Red	“music”, “gig”, “concert”
Cinema	Black	“cinema”

Table 1 – Cultural category colours and search terms

Table 1 identifies the marker colours associated with each cultural category and the search terms used. These categories and terms were selected based on their inclusion within the World Cities Culture Report (World Cities Culture Forum, 2015). Both American and British spellings of theatre have been included to ensure the same methodology can be applied equally to British and American cities.

Three cities have been selected for analysis within the discussion of this document: London, Newcastle upon Tyne and Berlin, as each represent a unique city typology. Results for additional cities can be found within Appendix 1. All results were collected on Thursday, 21st April meaning all tweets are within the preceding seven day window.

London

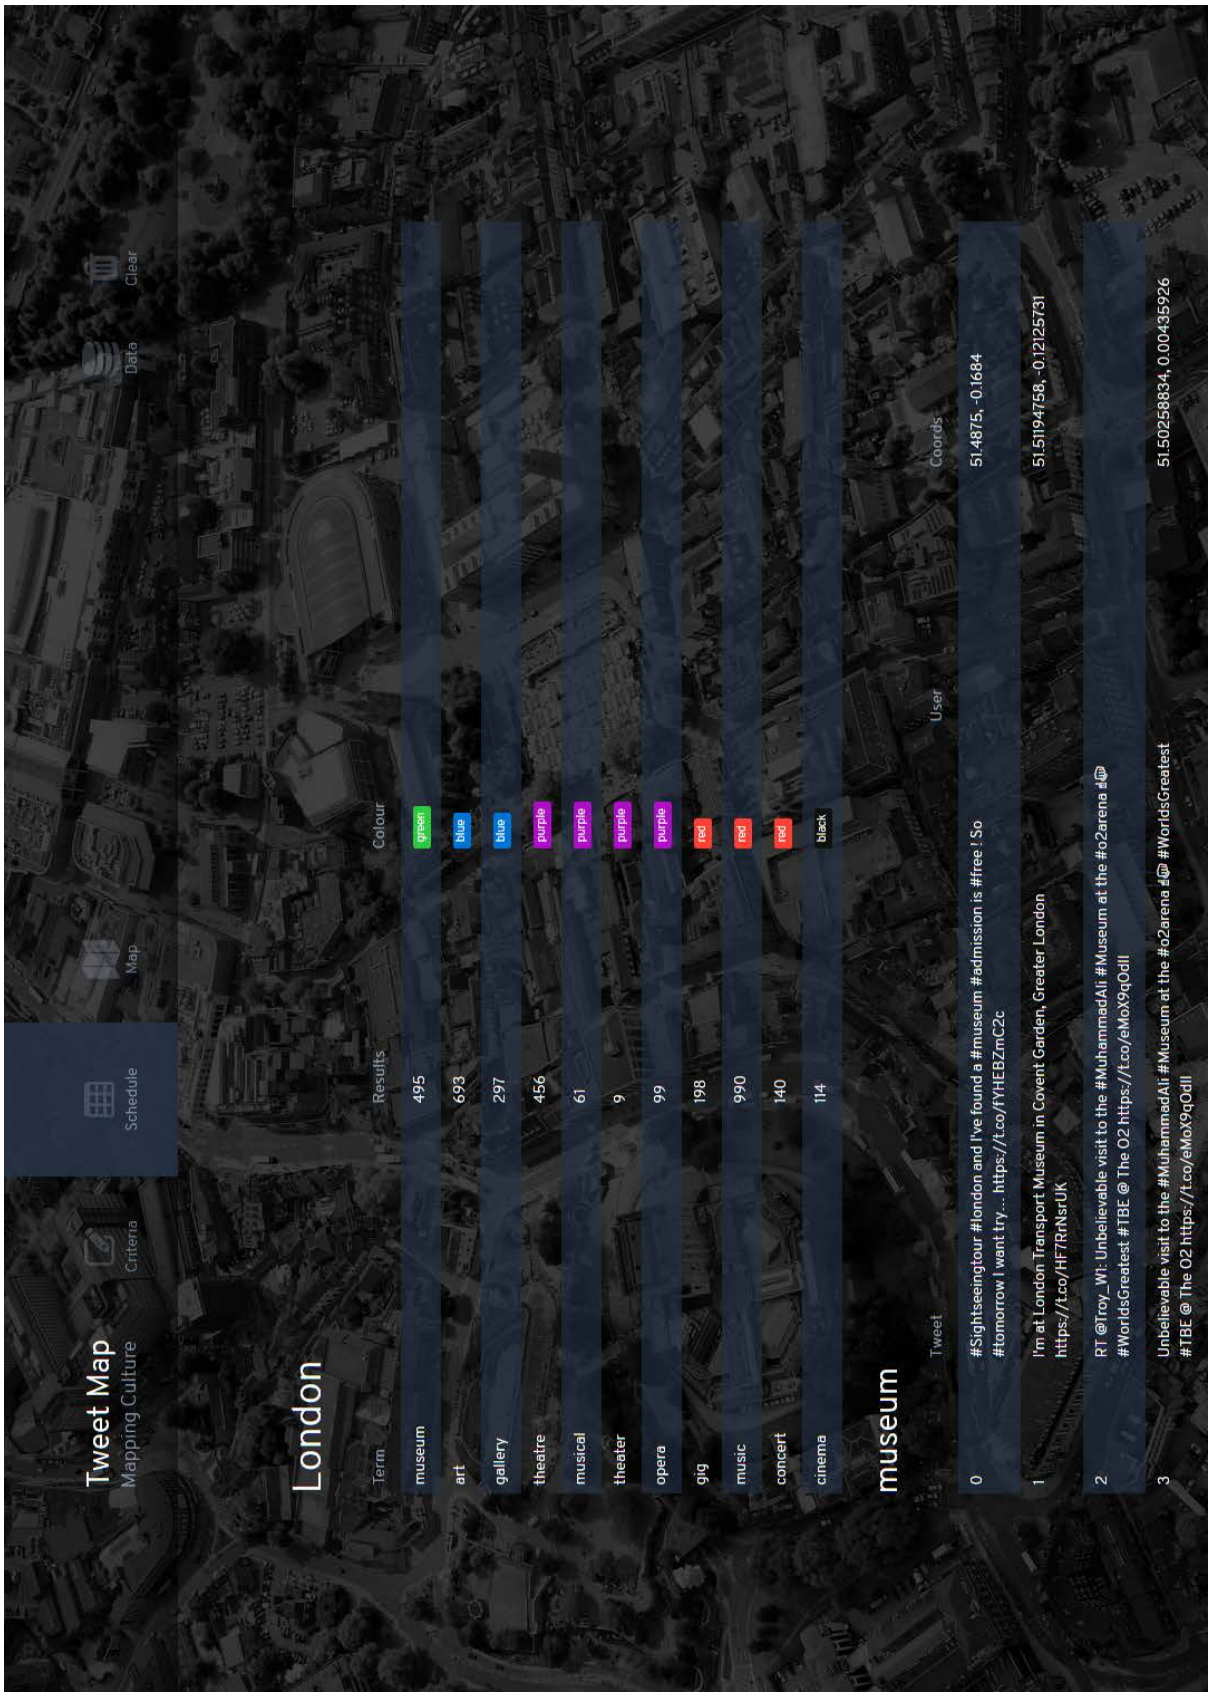

Figure 21 – London results table

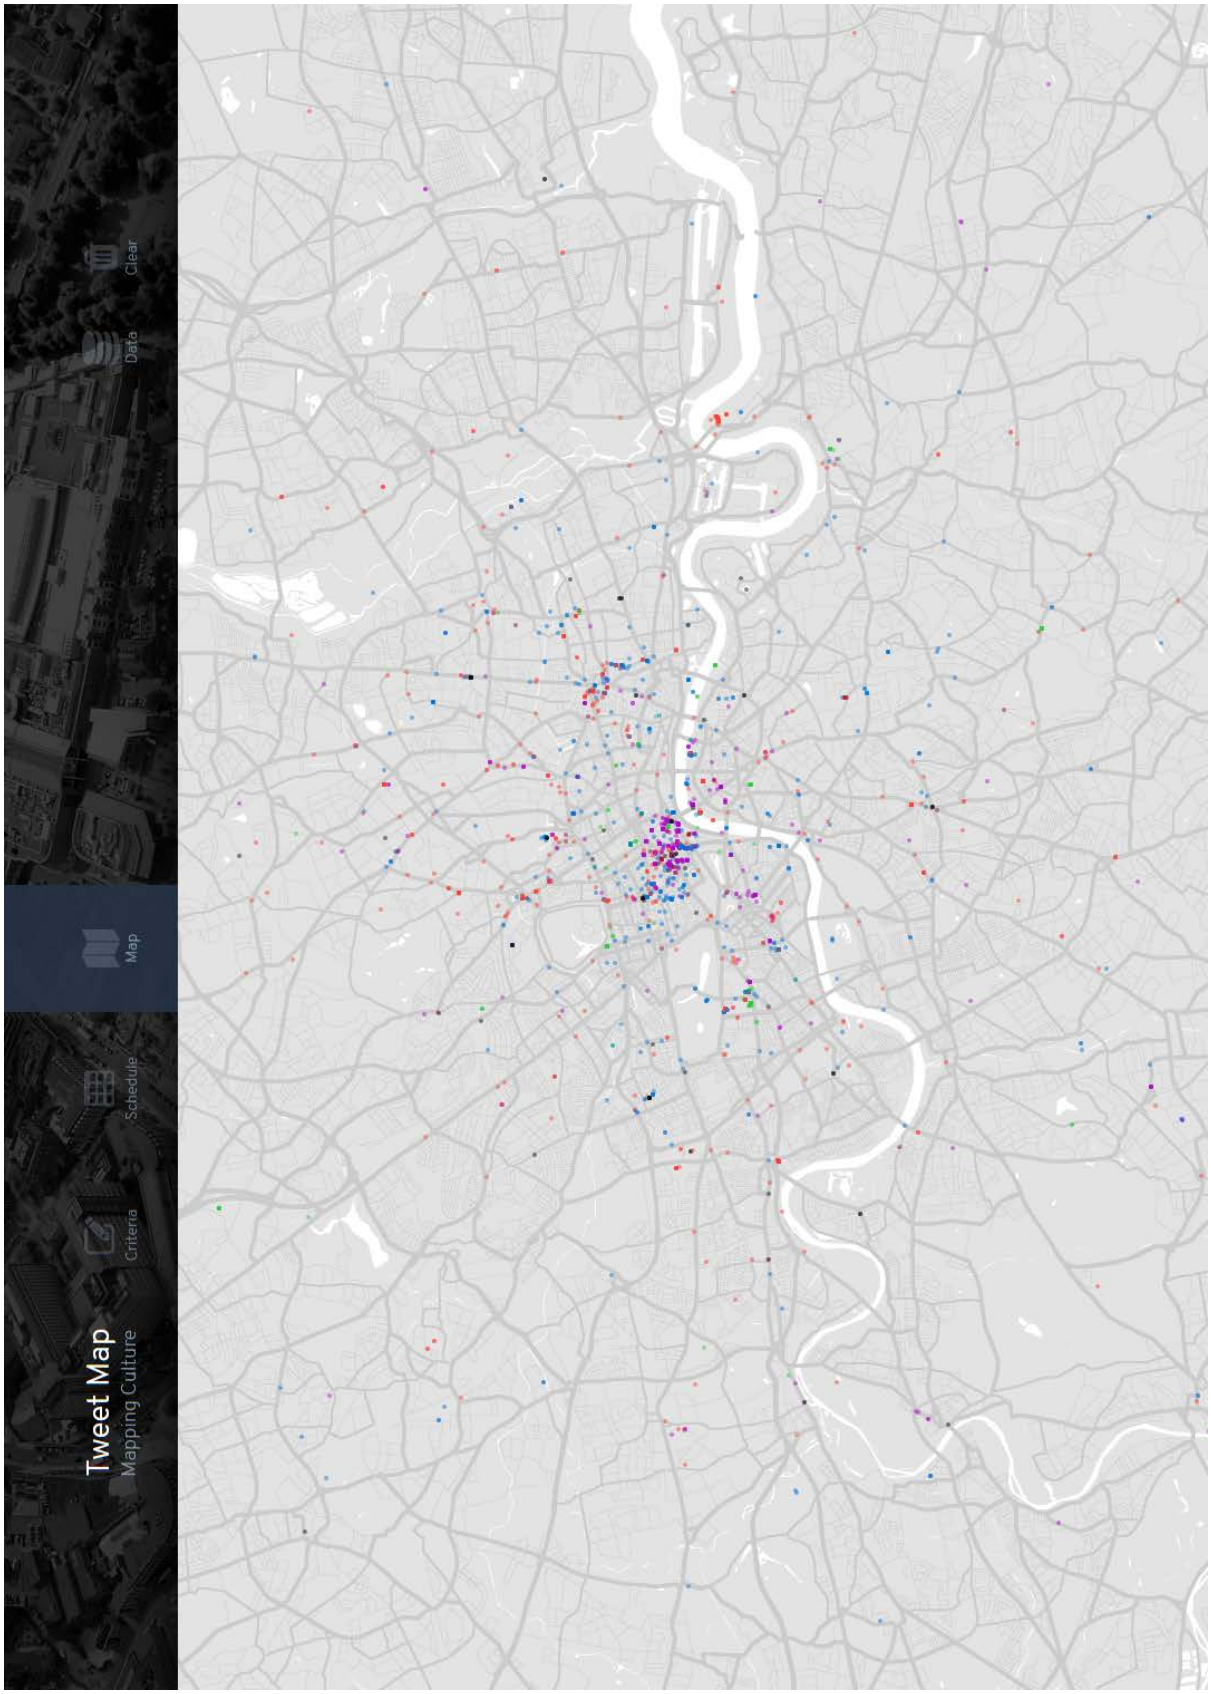

Figure 22 – London tweet culture map at zoom level 12

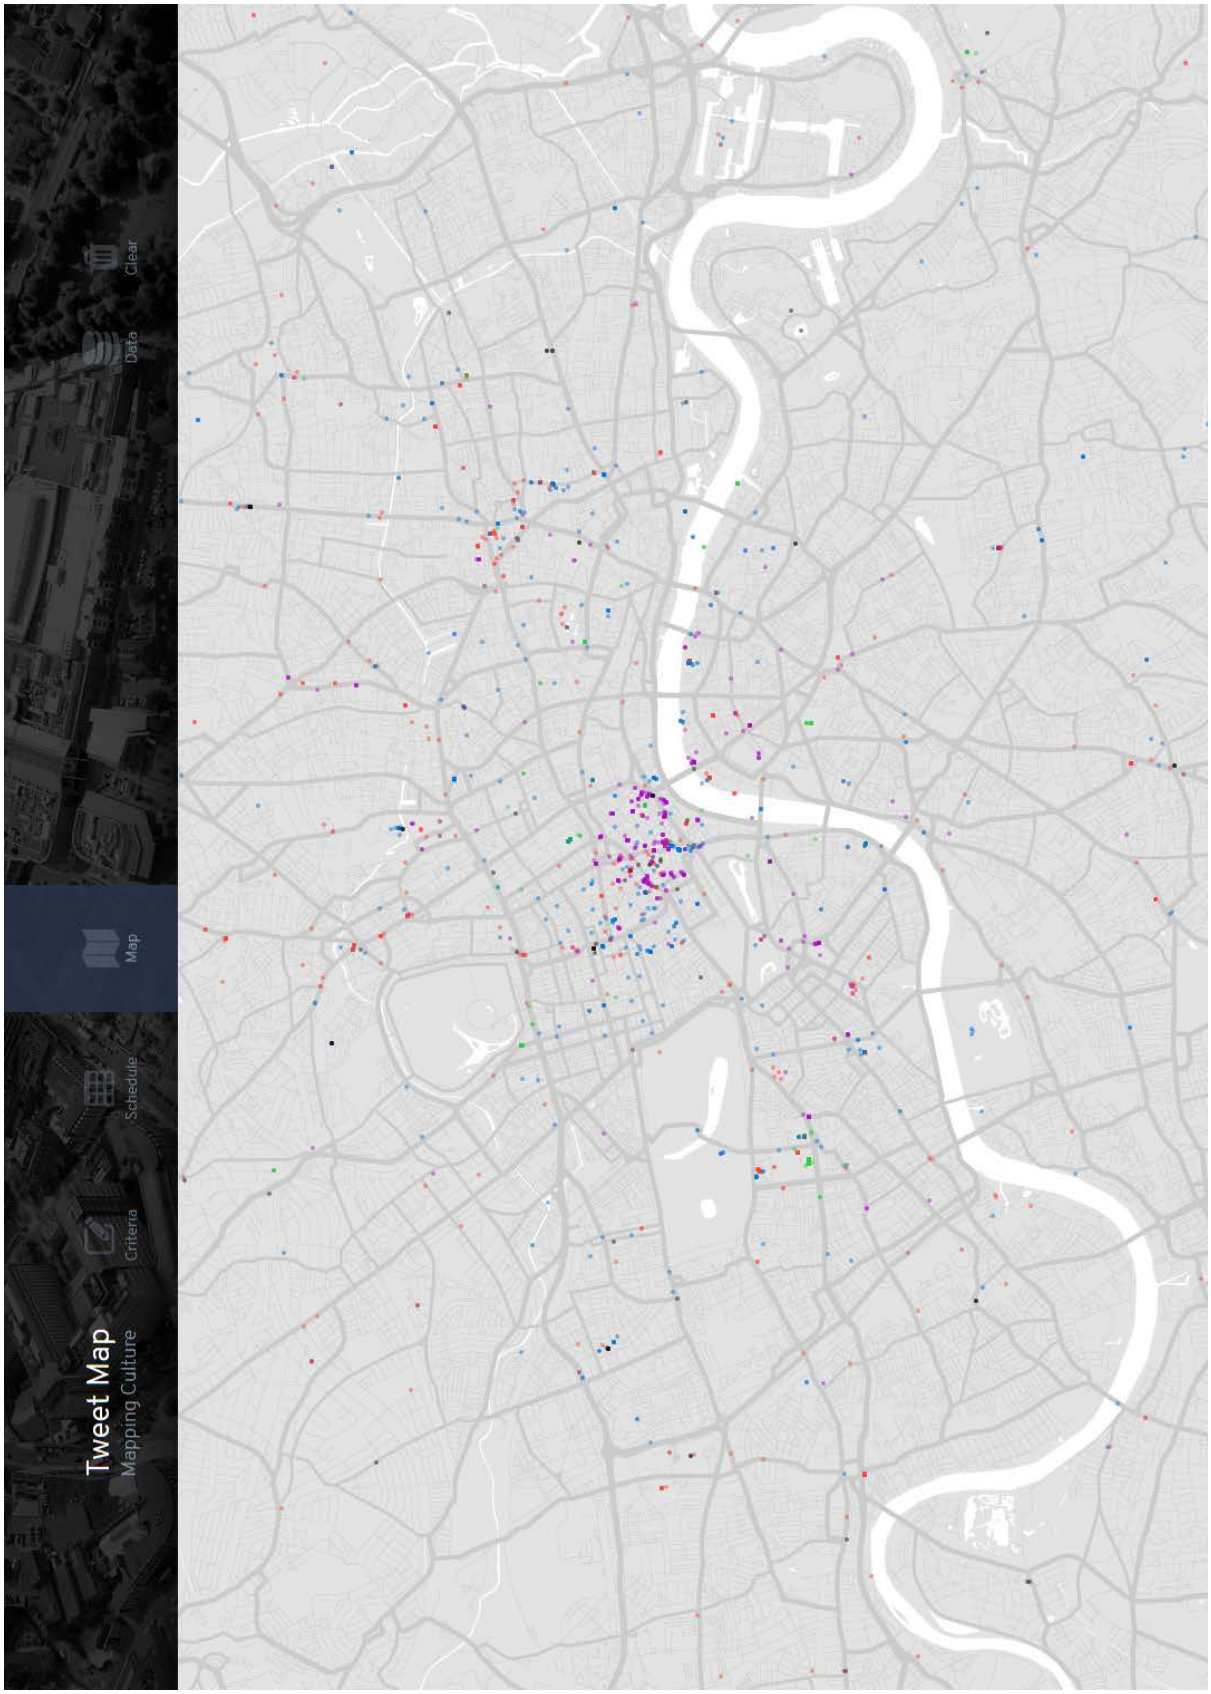

Figure 23 – London tweet culture map at zoom level 13

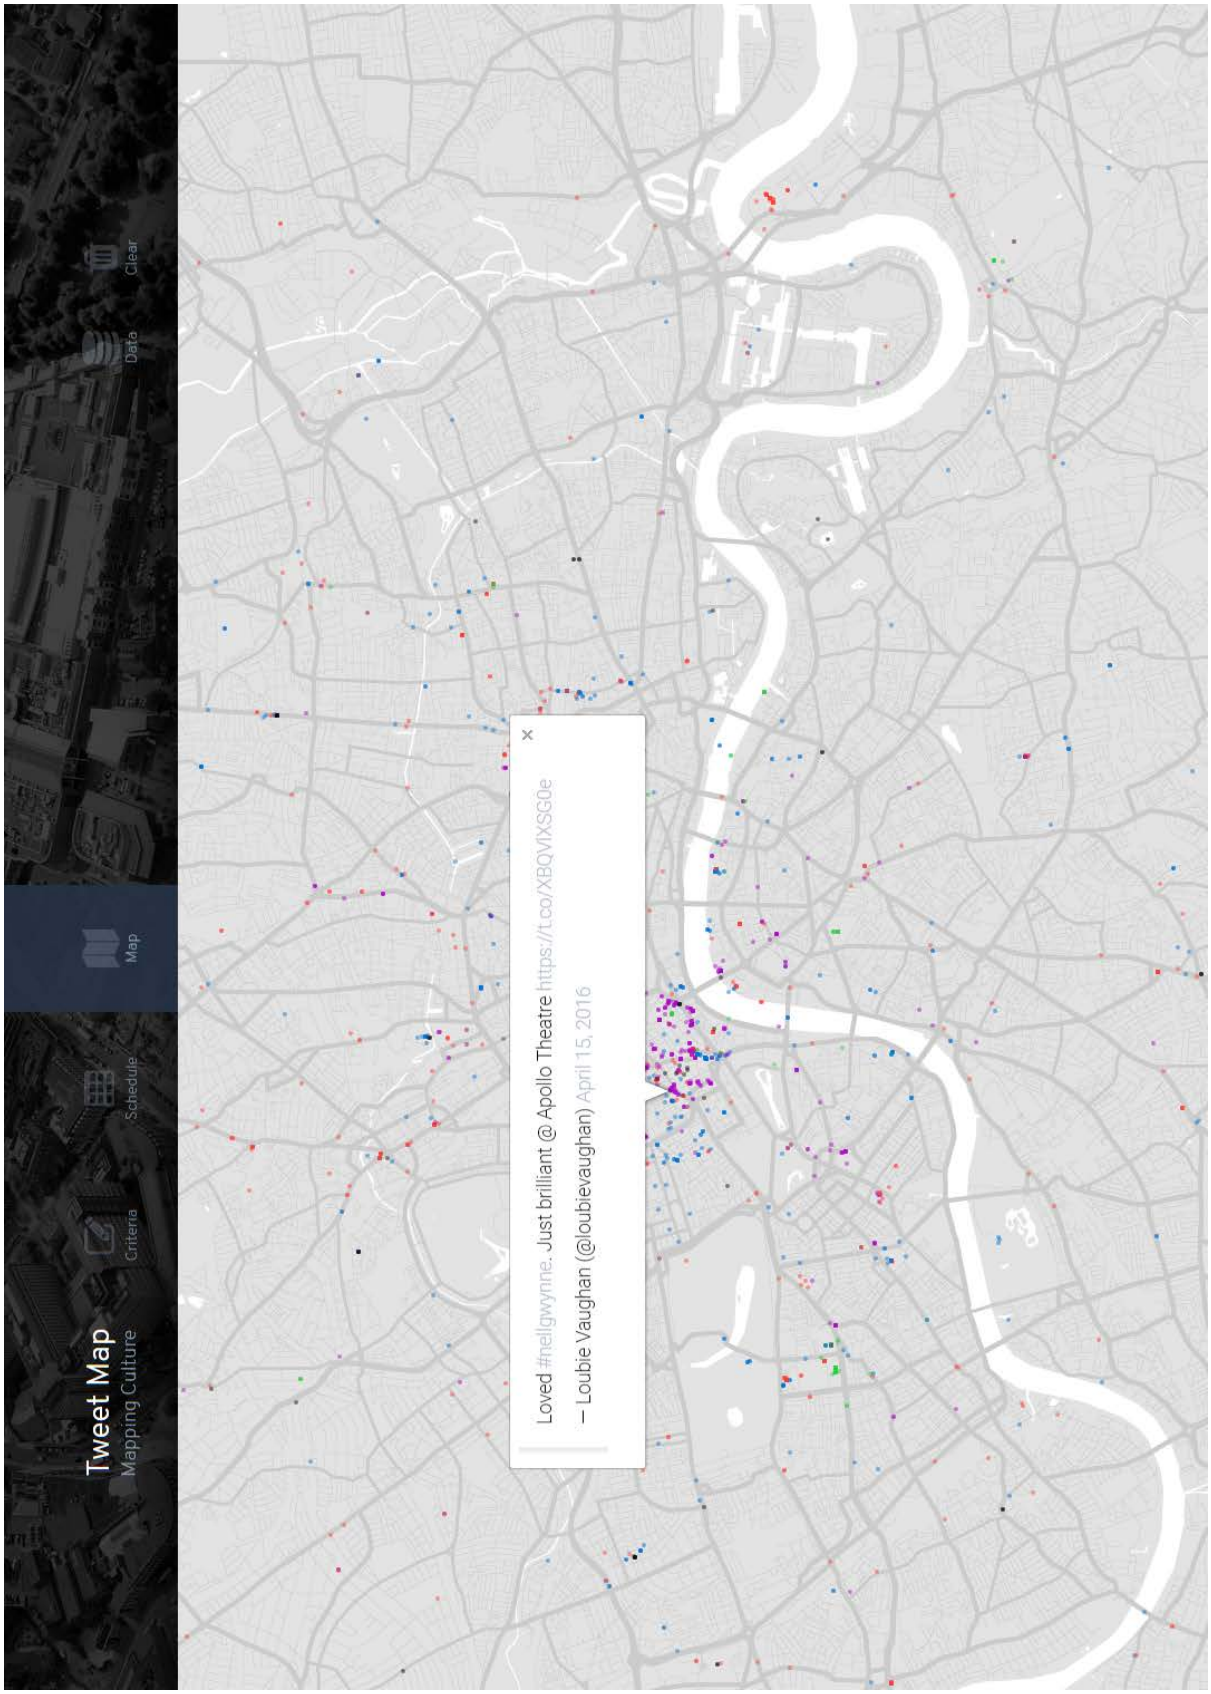

Figure 24 – West End, London

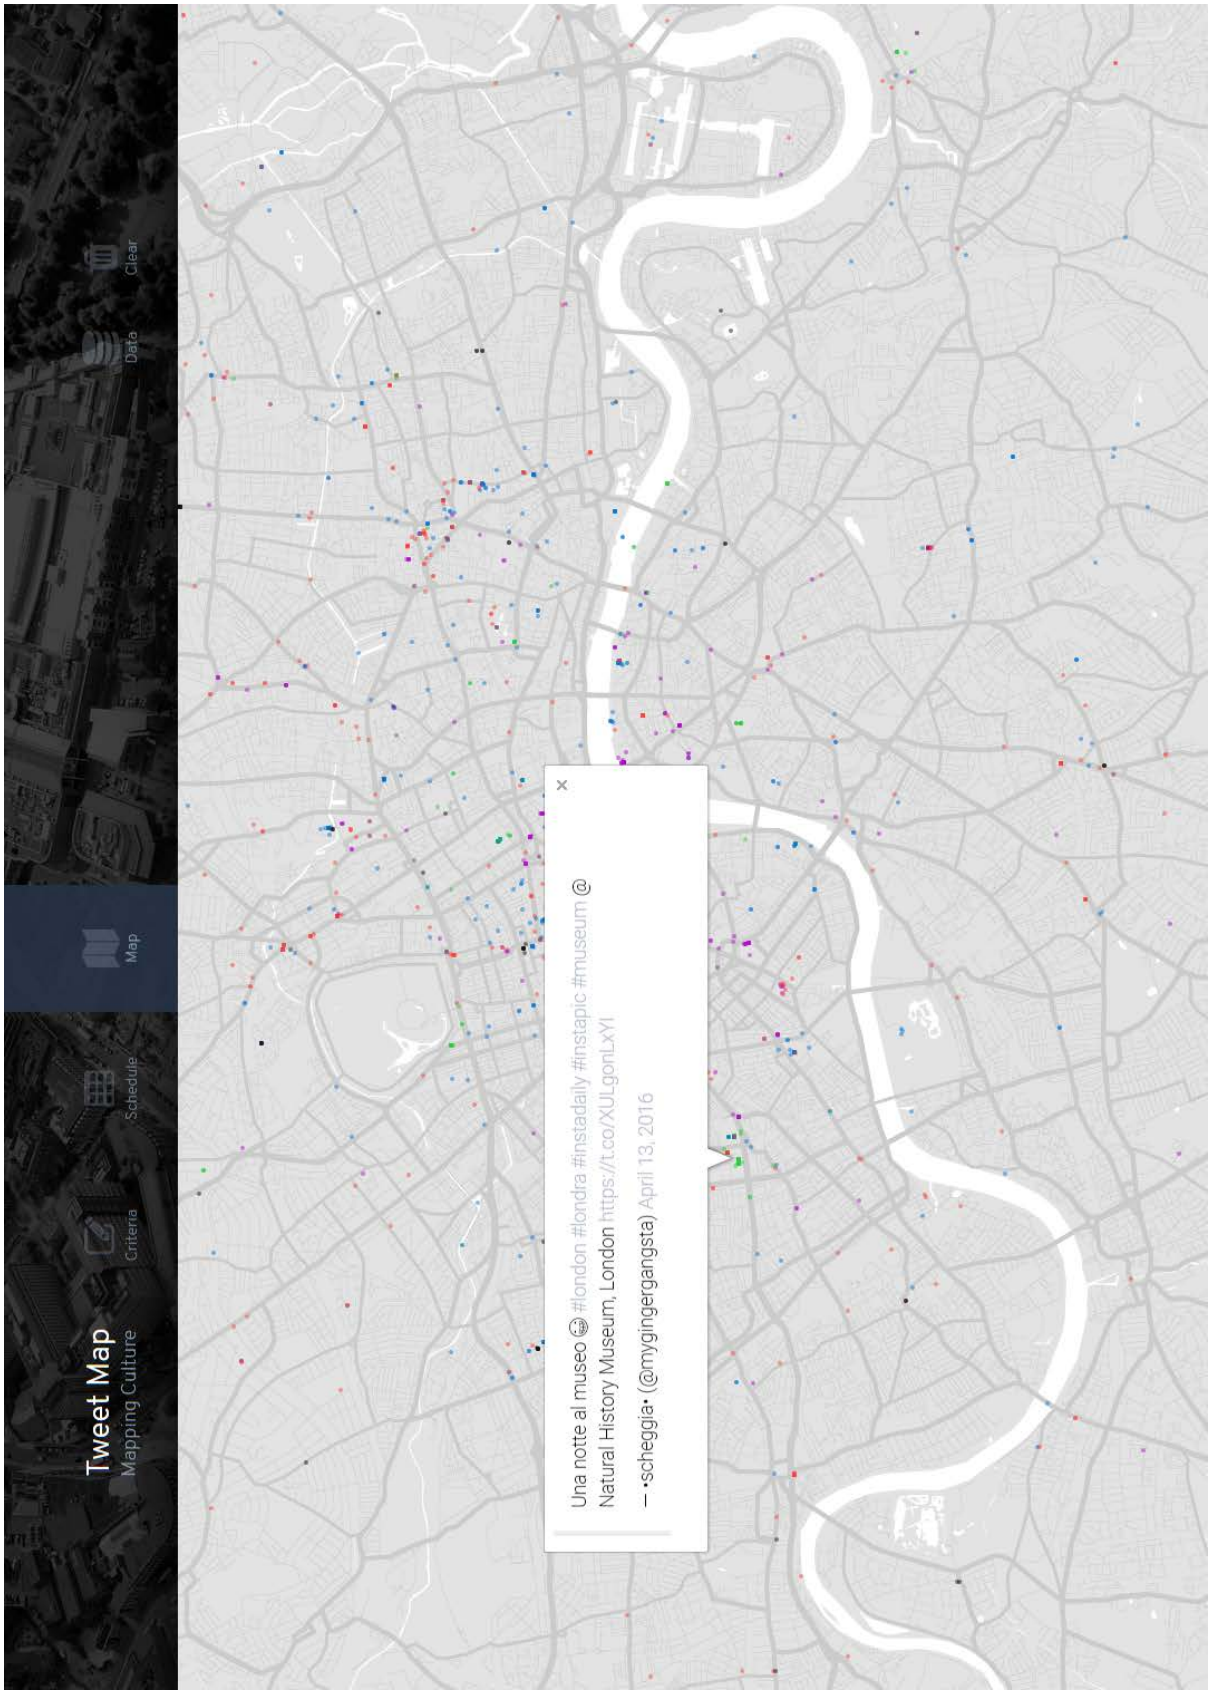

Figure 25 –Natural History Museum, London

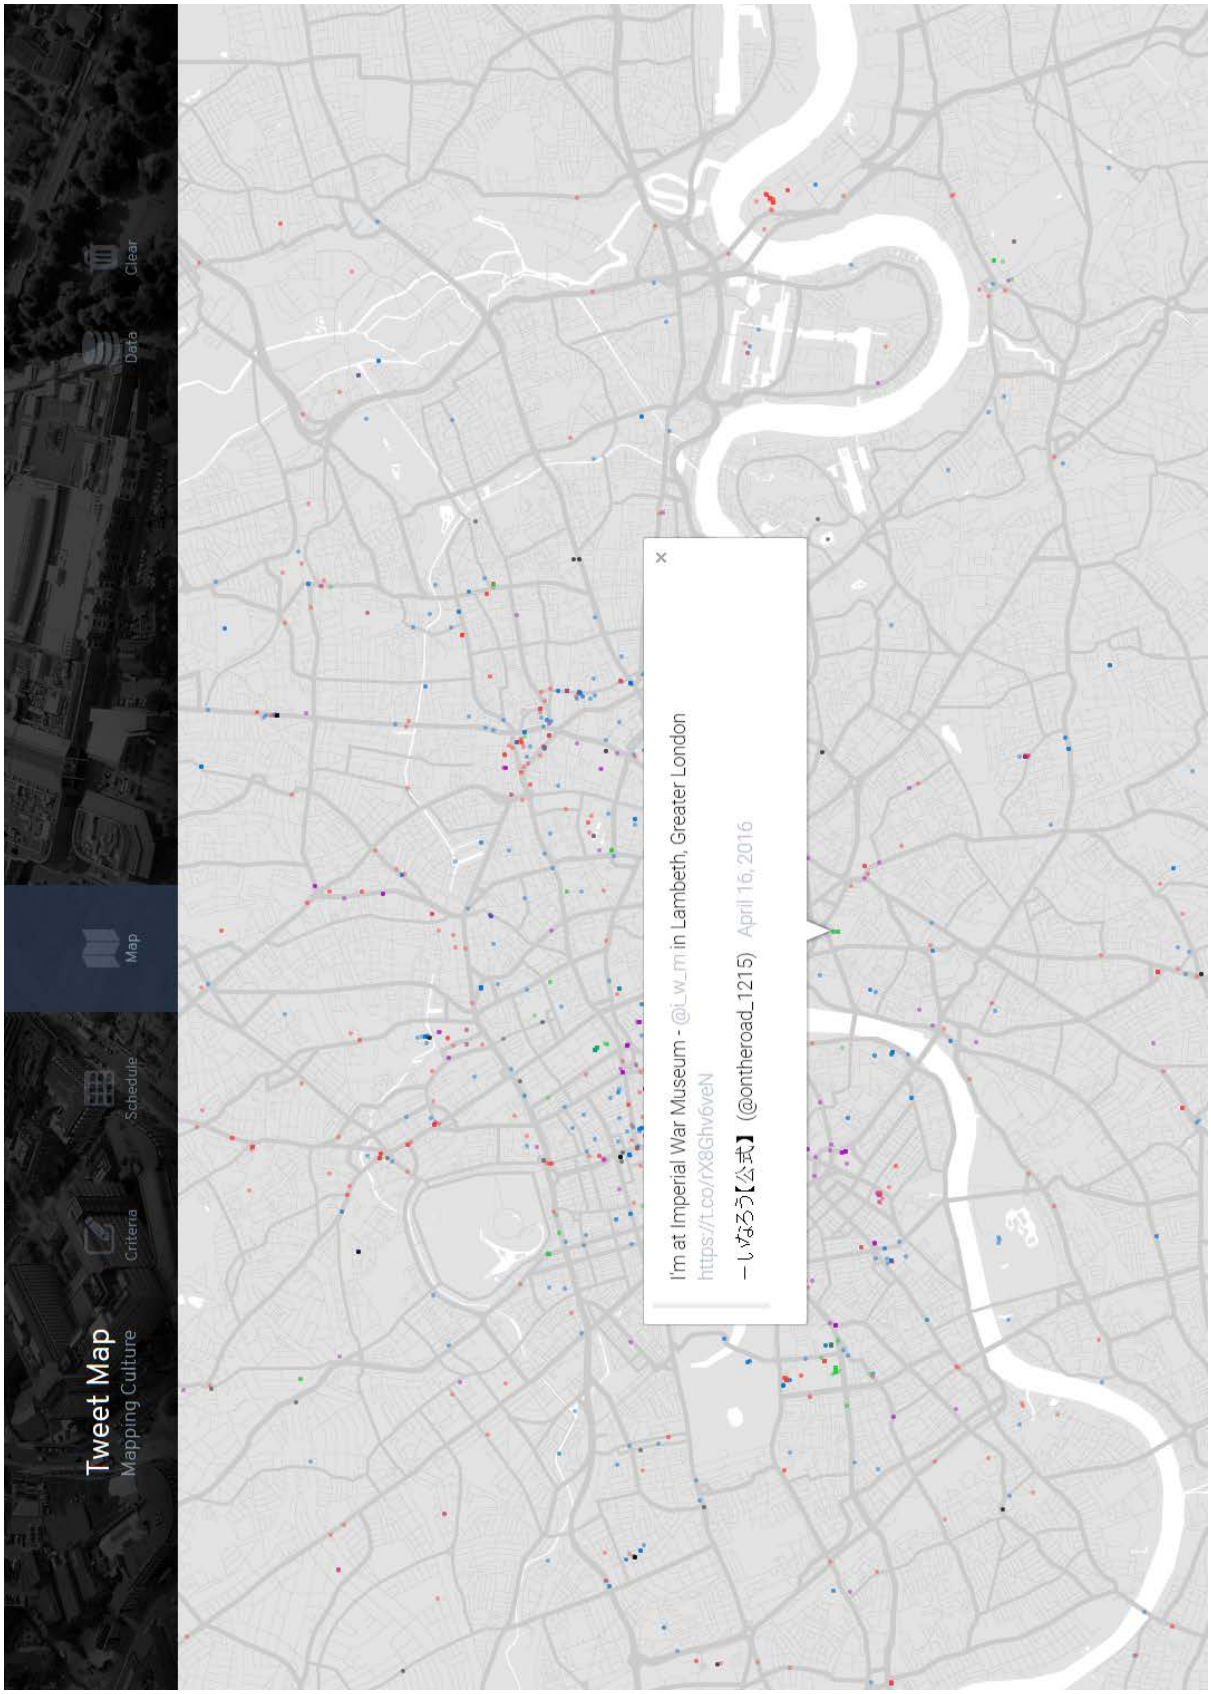

Figure 26 – Imperial War Museum, London

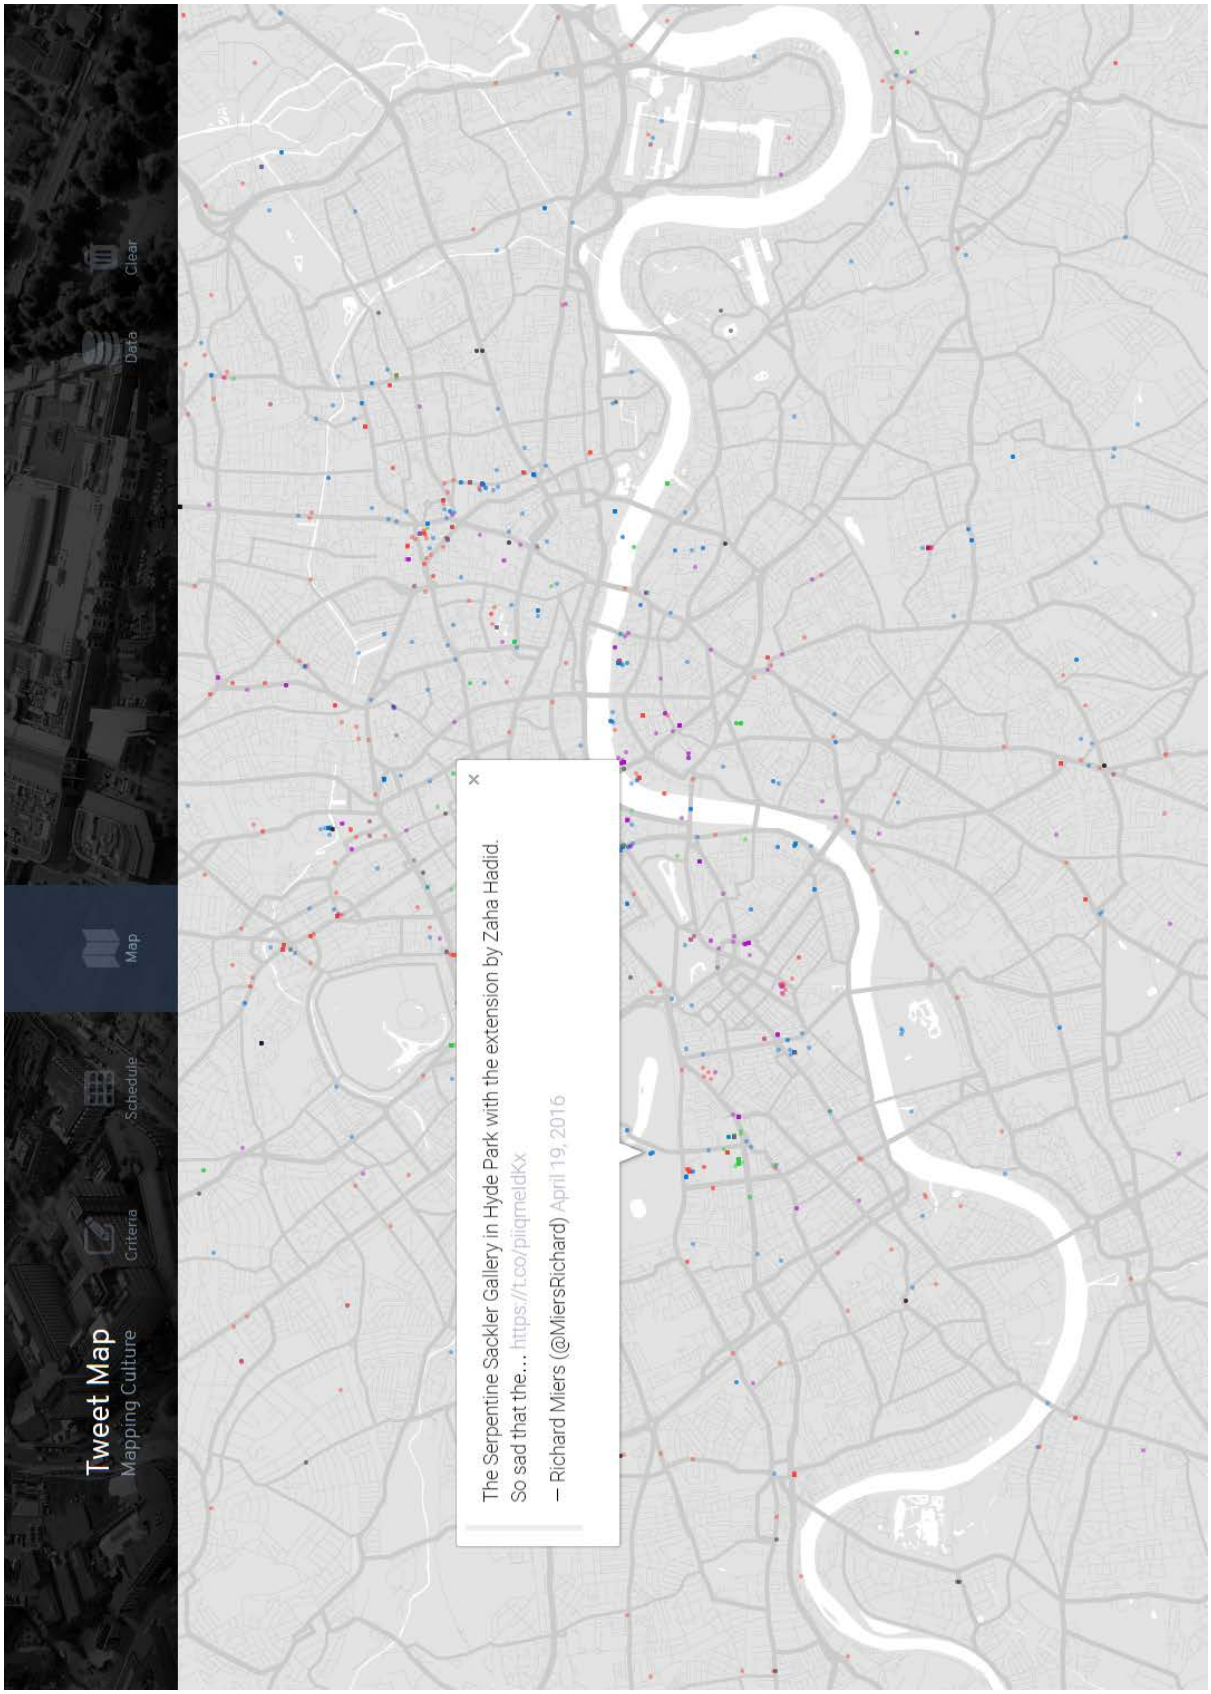

Figure 27 – Serpentine Gallery, London

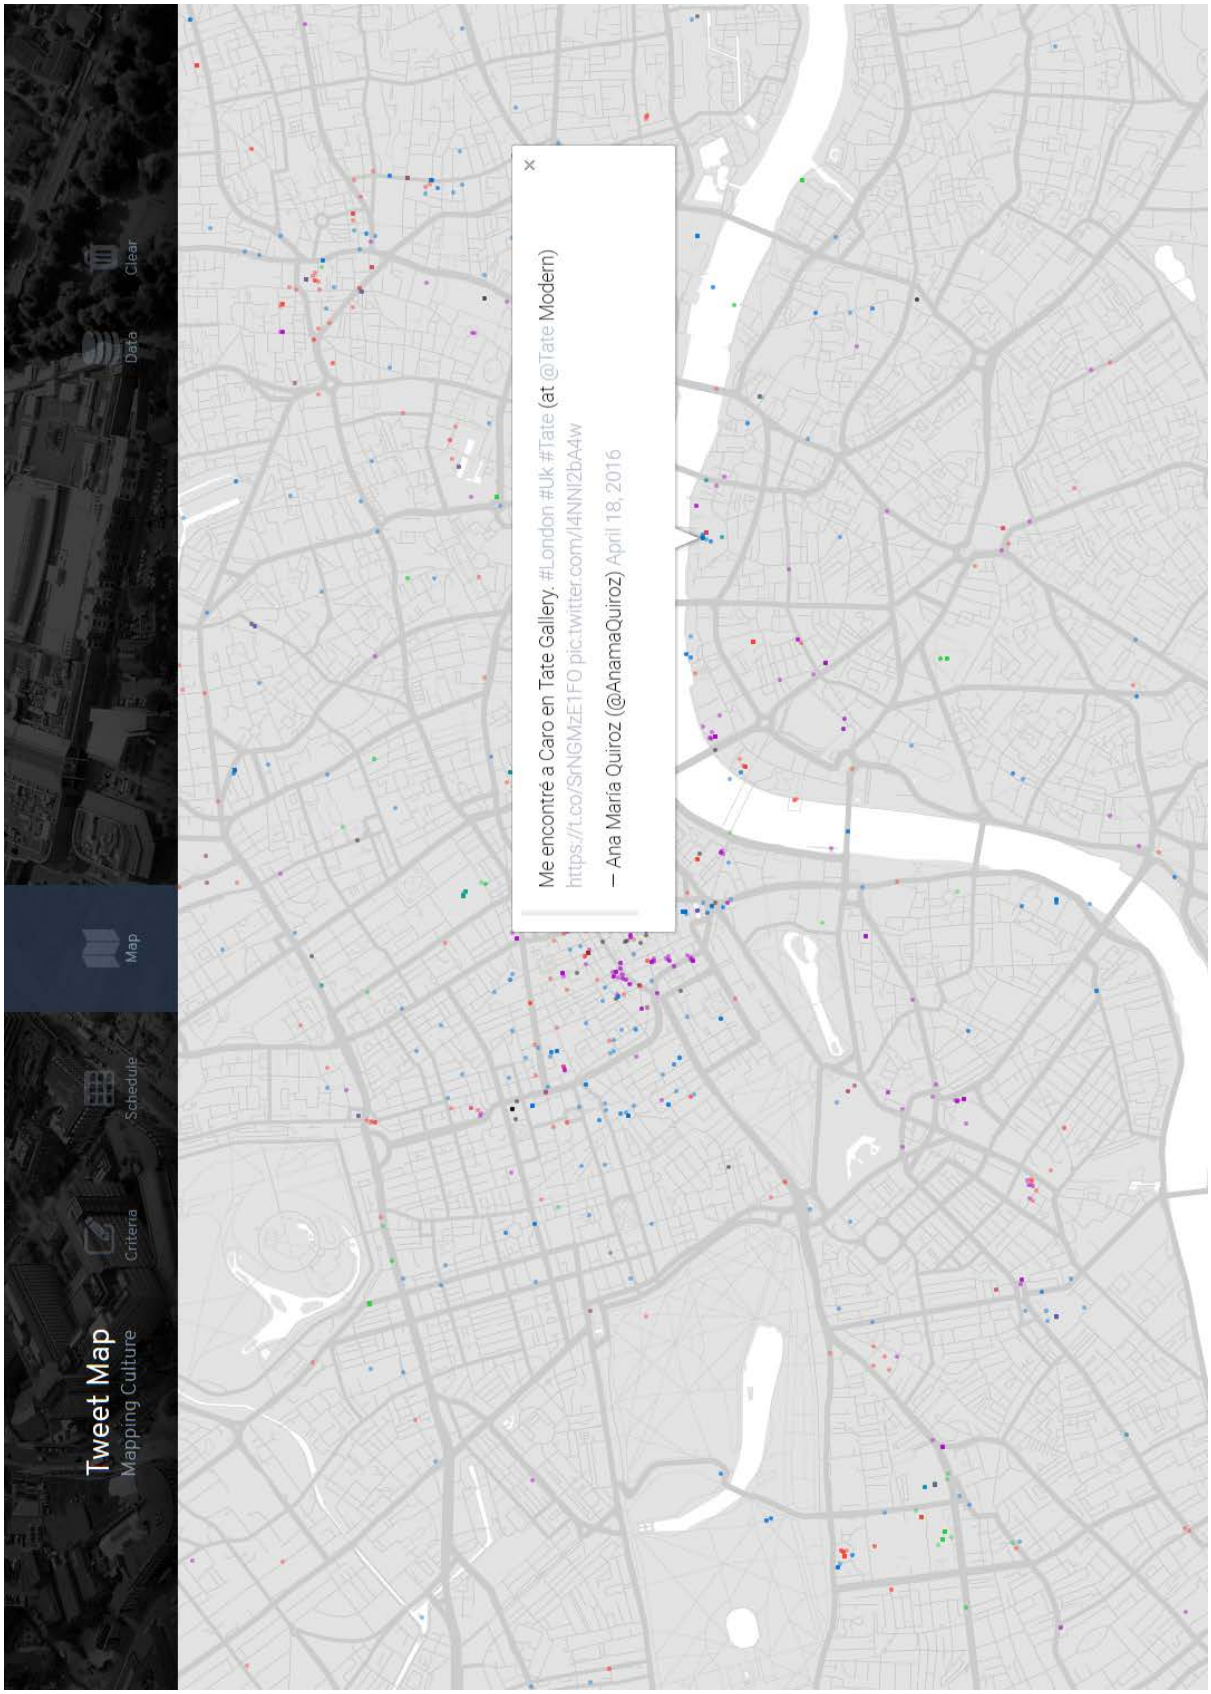
Tweet Map
Mapping Culture
Criteria
Schedule
Map
Data
Clear
Me encontré a Caro en Tate Gallery. #London #Uk #Tate (at @Tate Modern) https://t.co/SrNGMzE1FO pic.twitter.com/l4NNI2bA4w
— Ana María Quiroz (@AnamaQuiroz) April 18, 2016

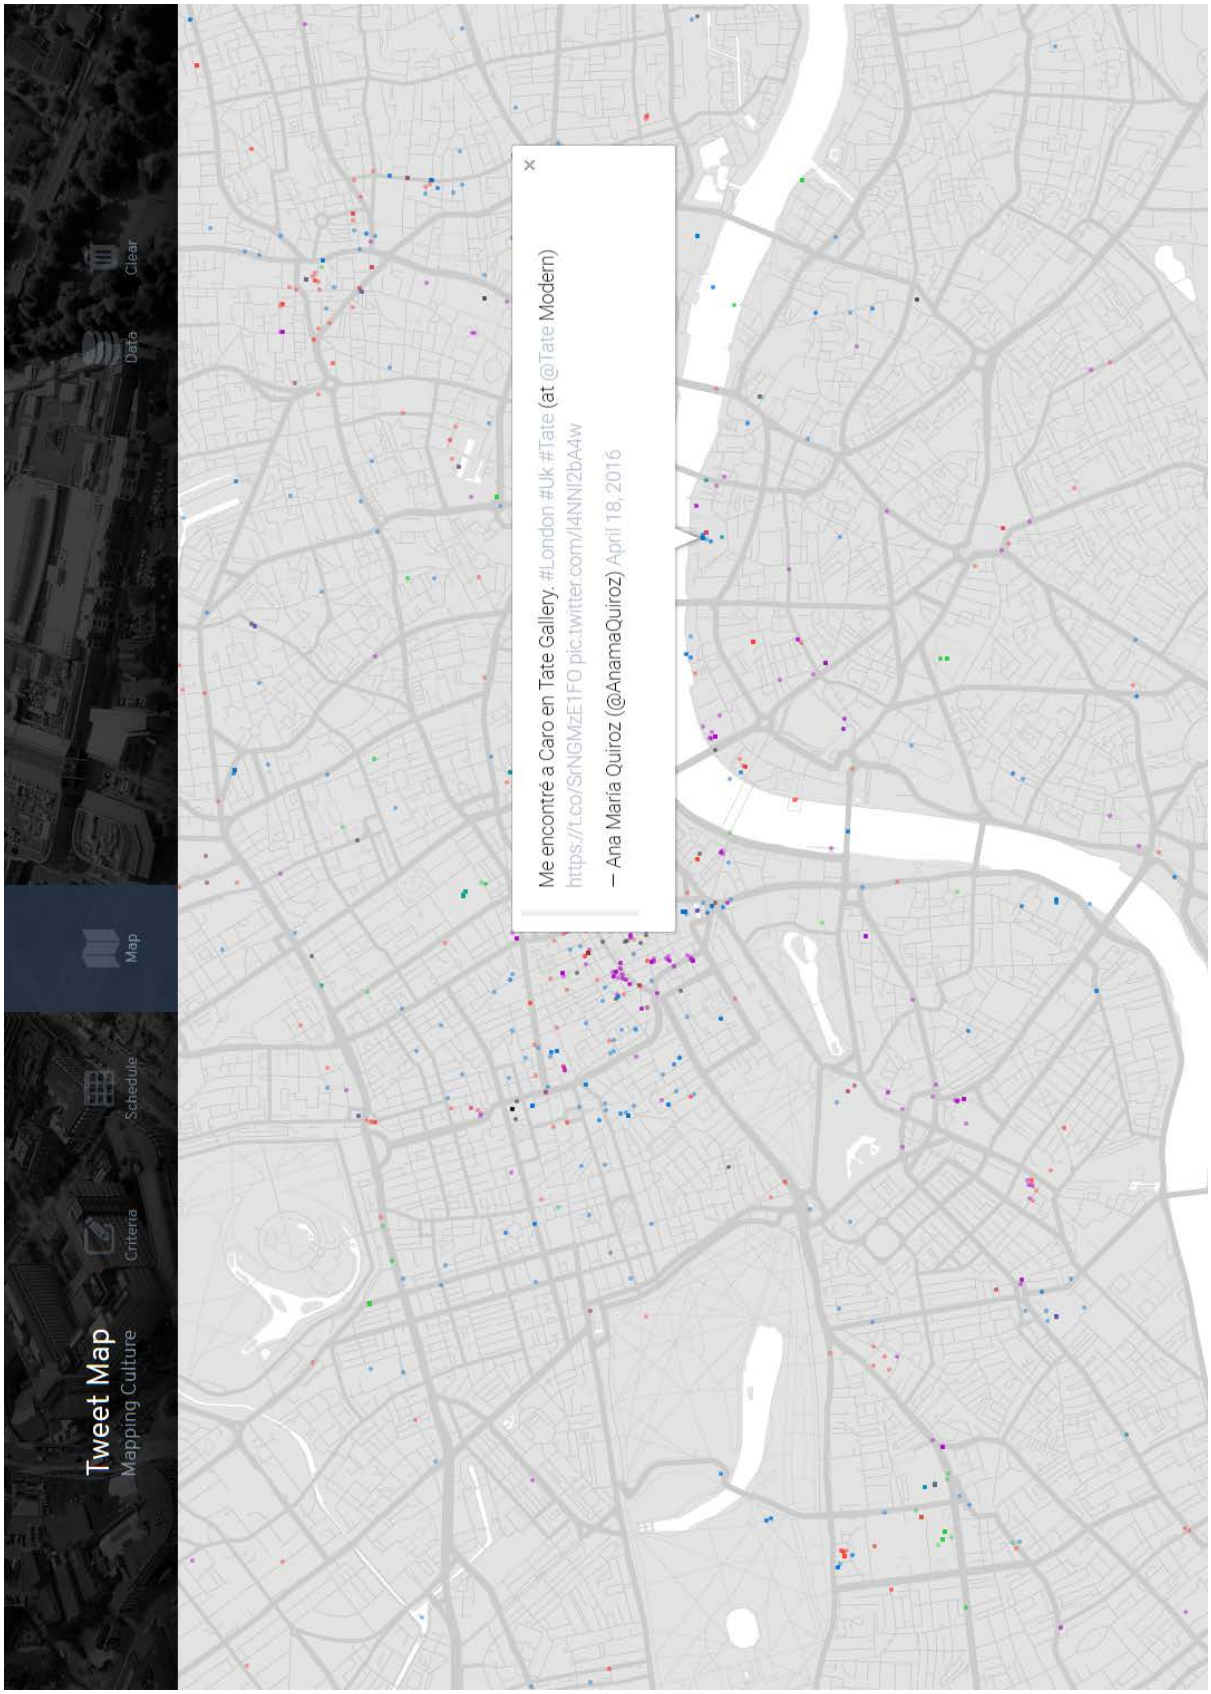

Figure 28 – Tate Modern, London

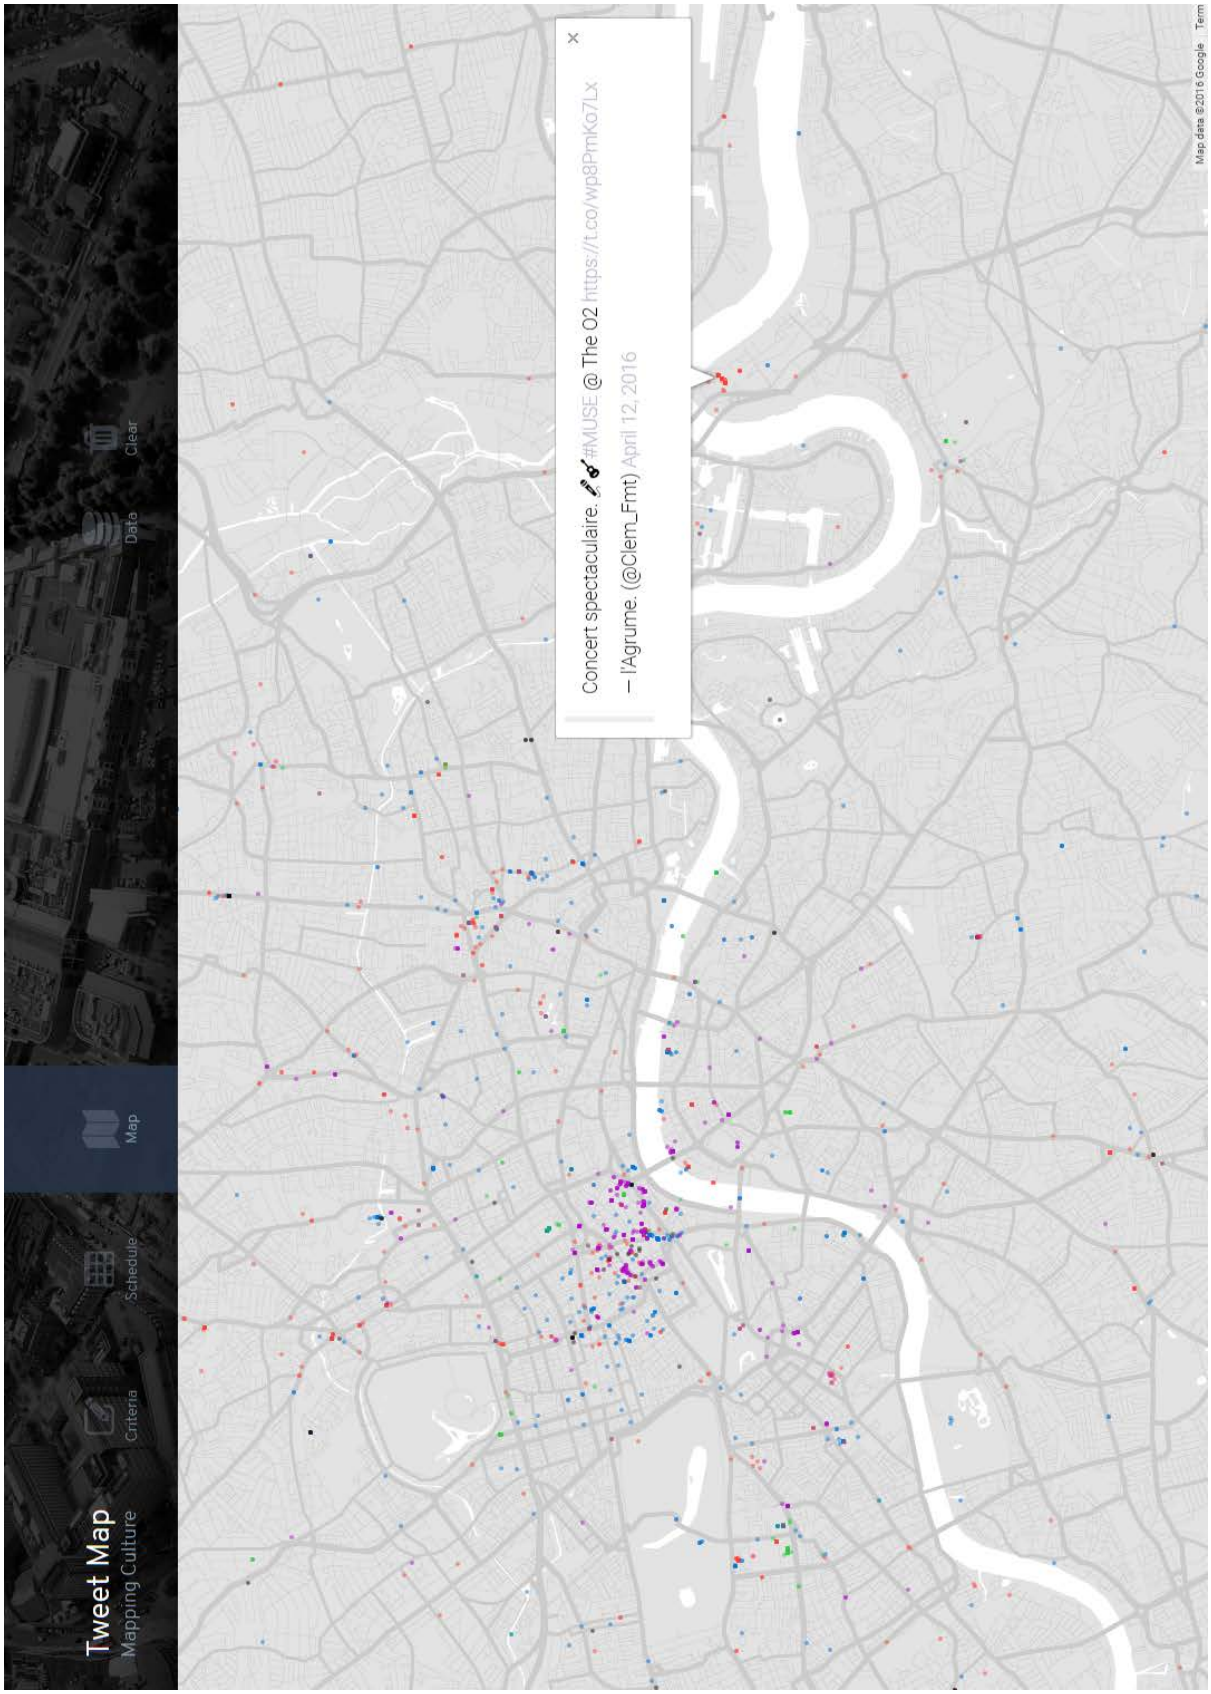

Figure 29 – The O2, London

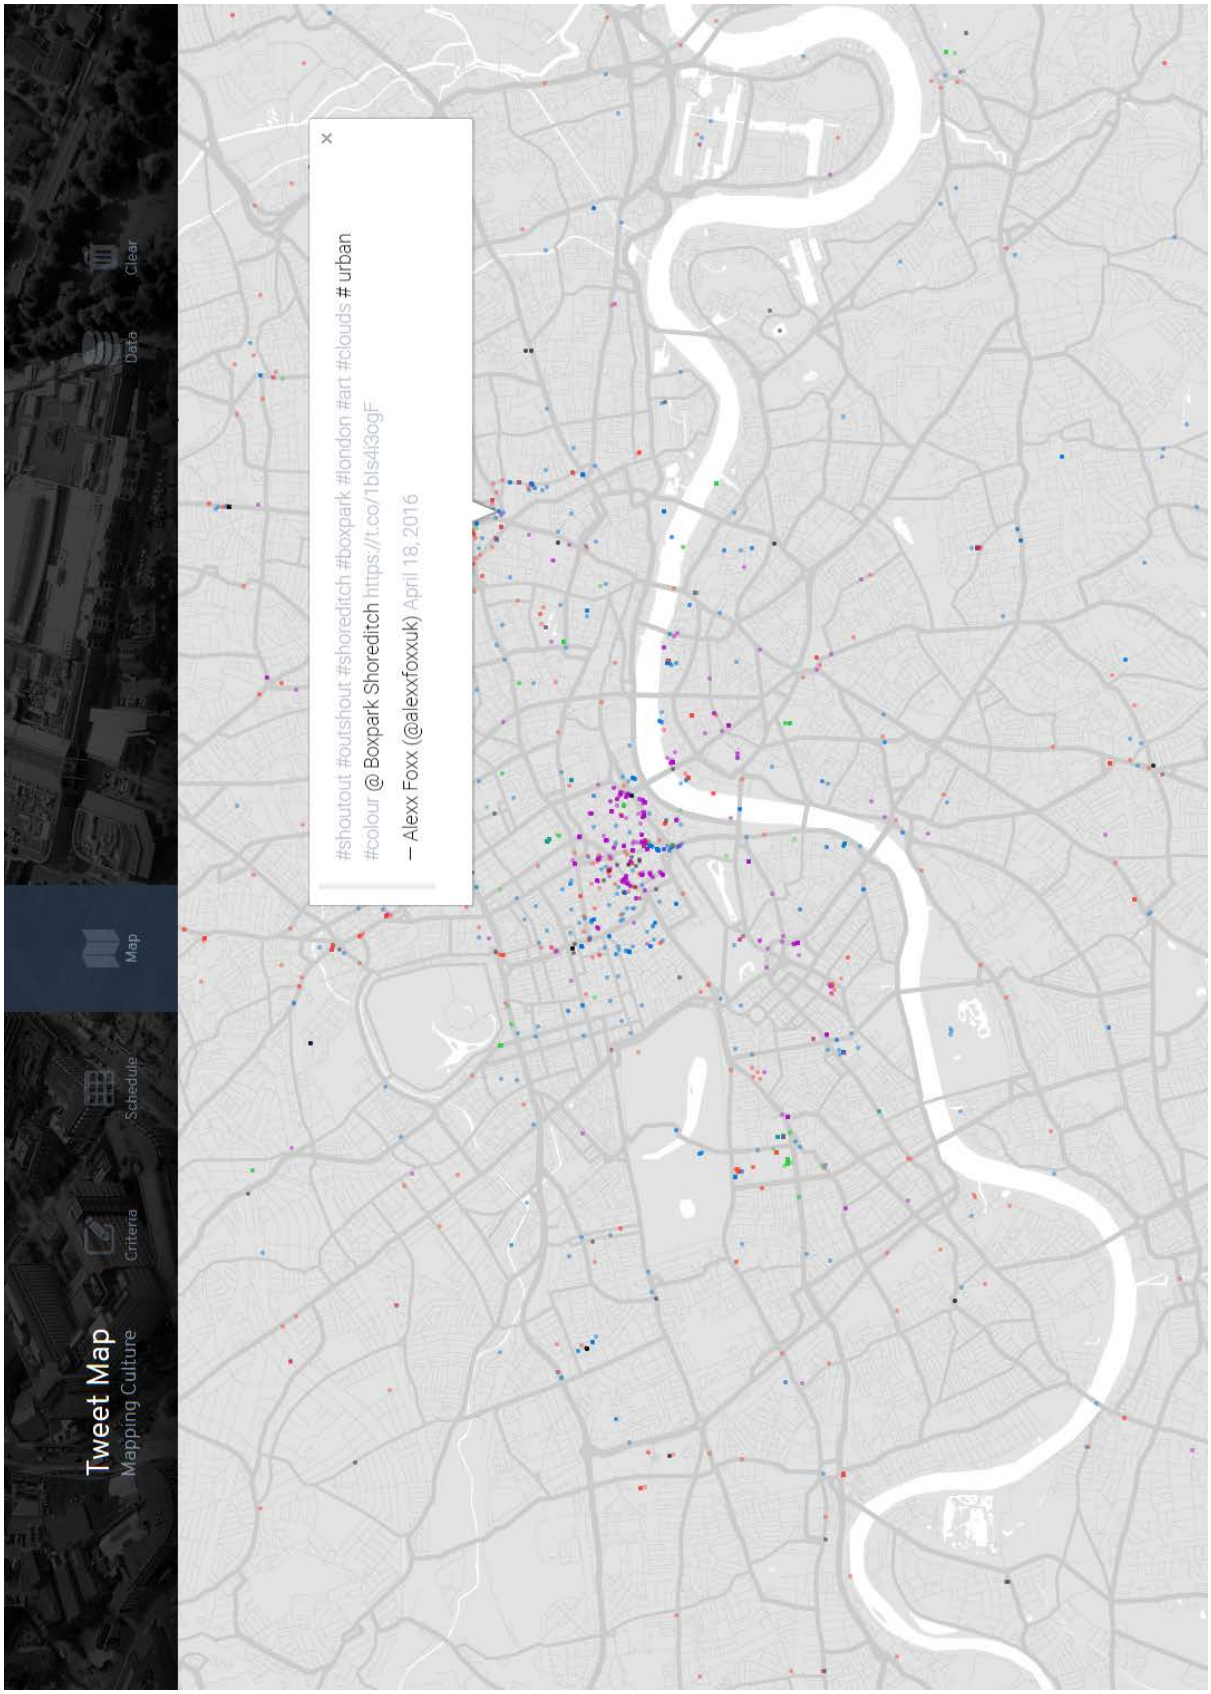

Figure 30 – Shoreditch, London

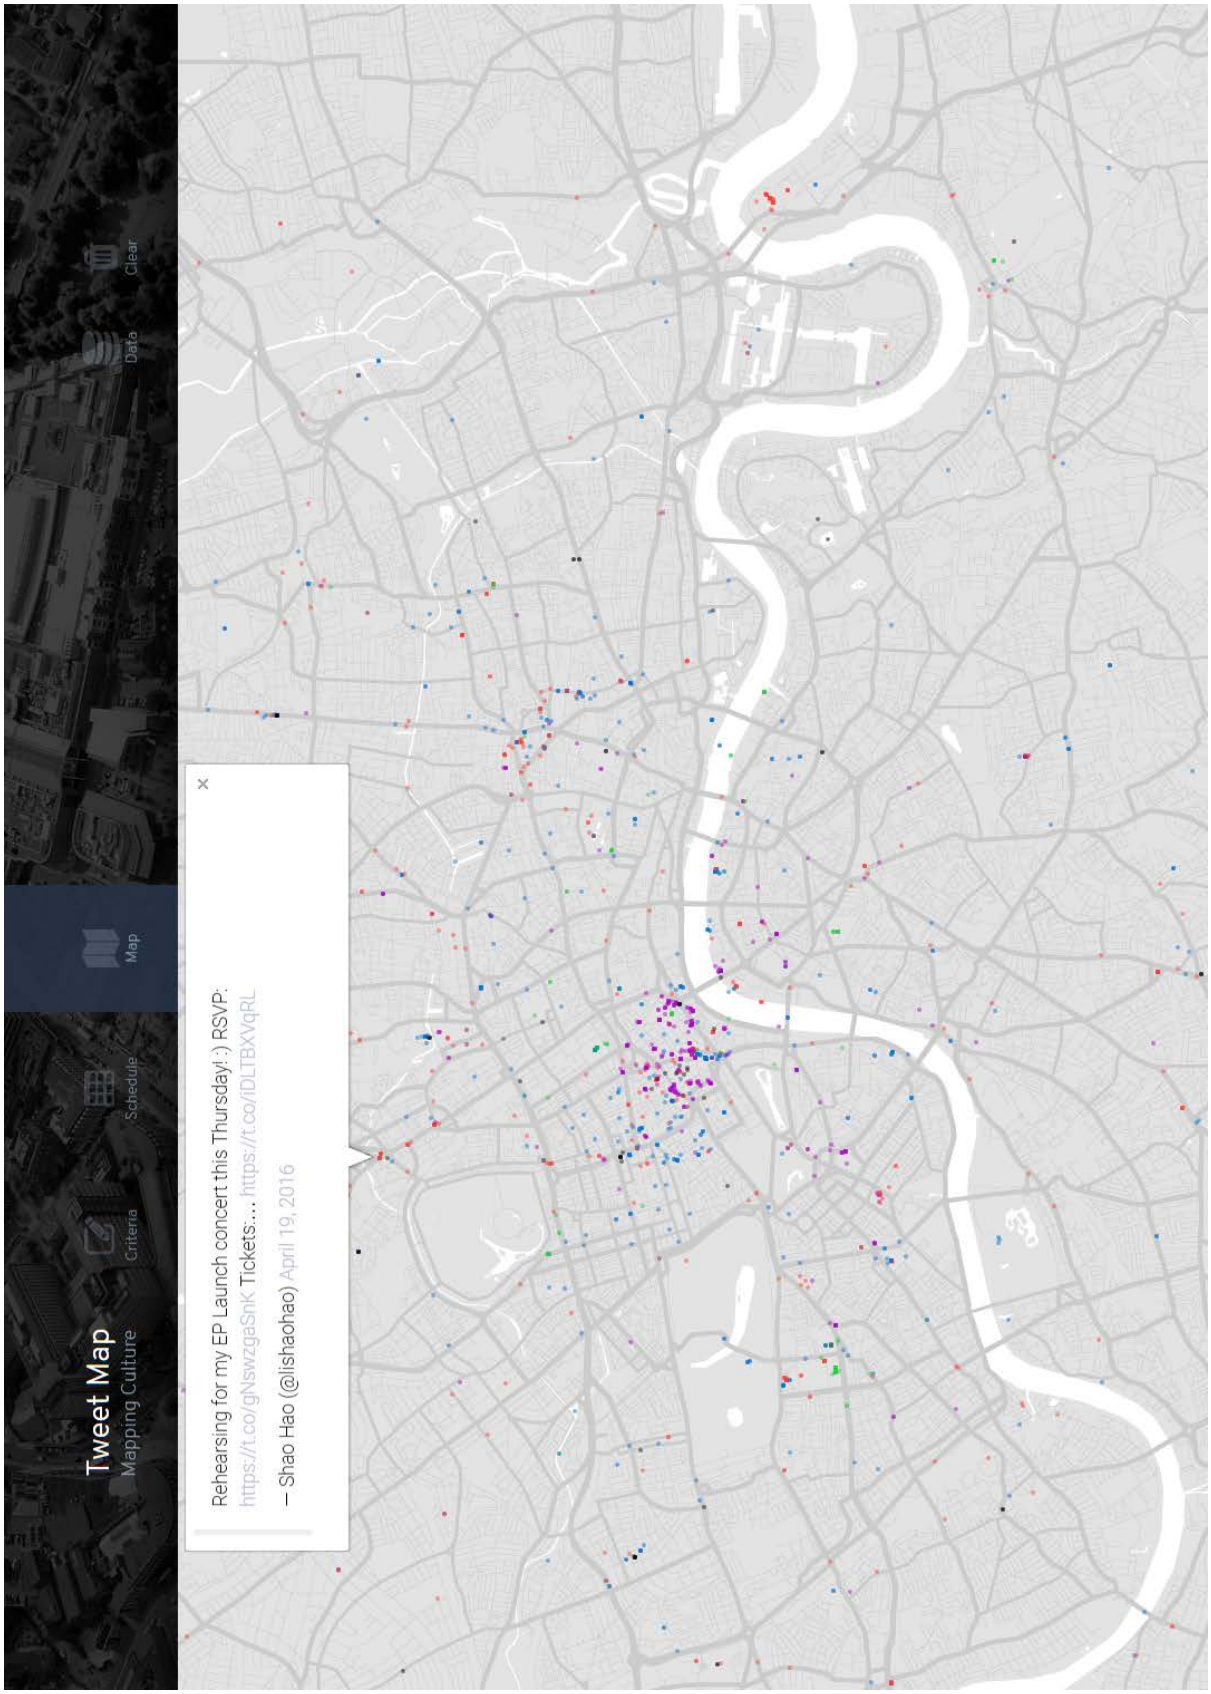

Figure 31 – Camden Town, London

Newcastle upon Tyne

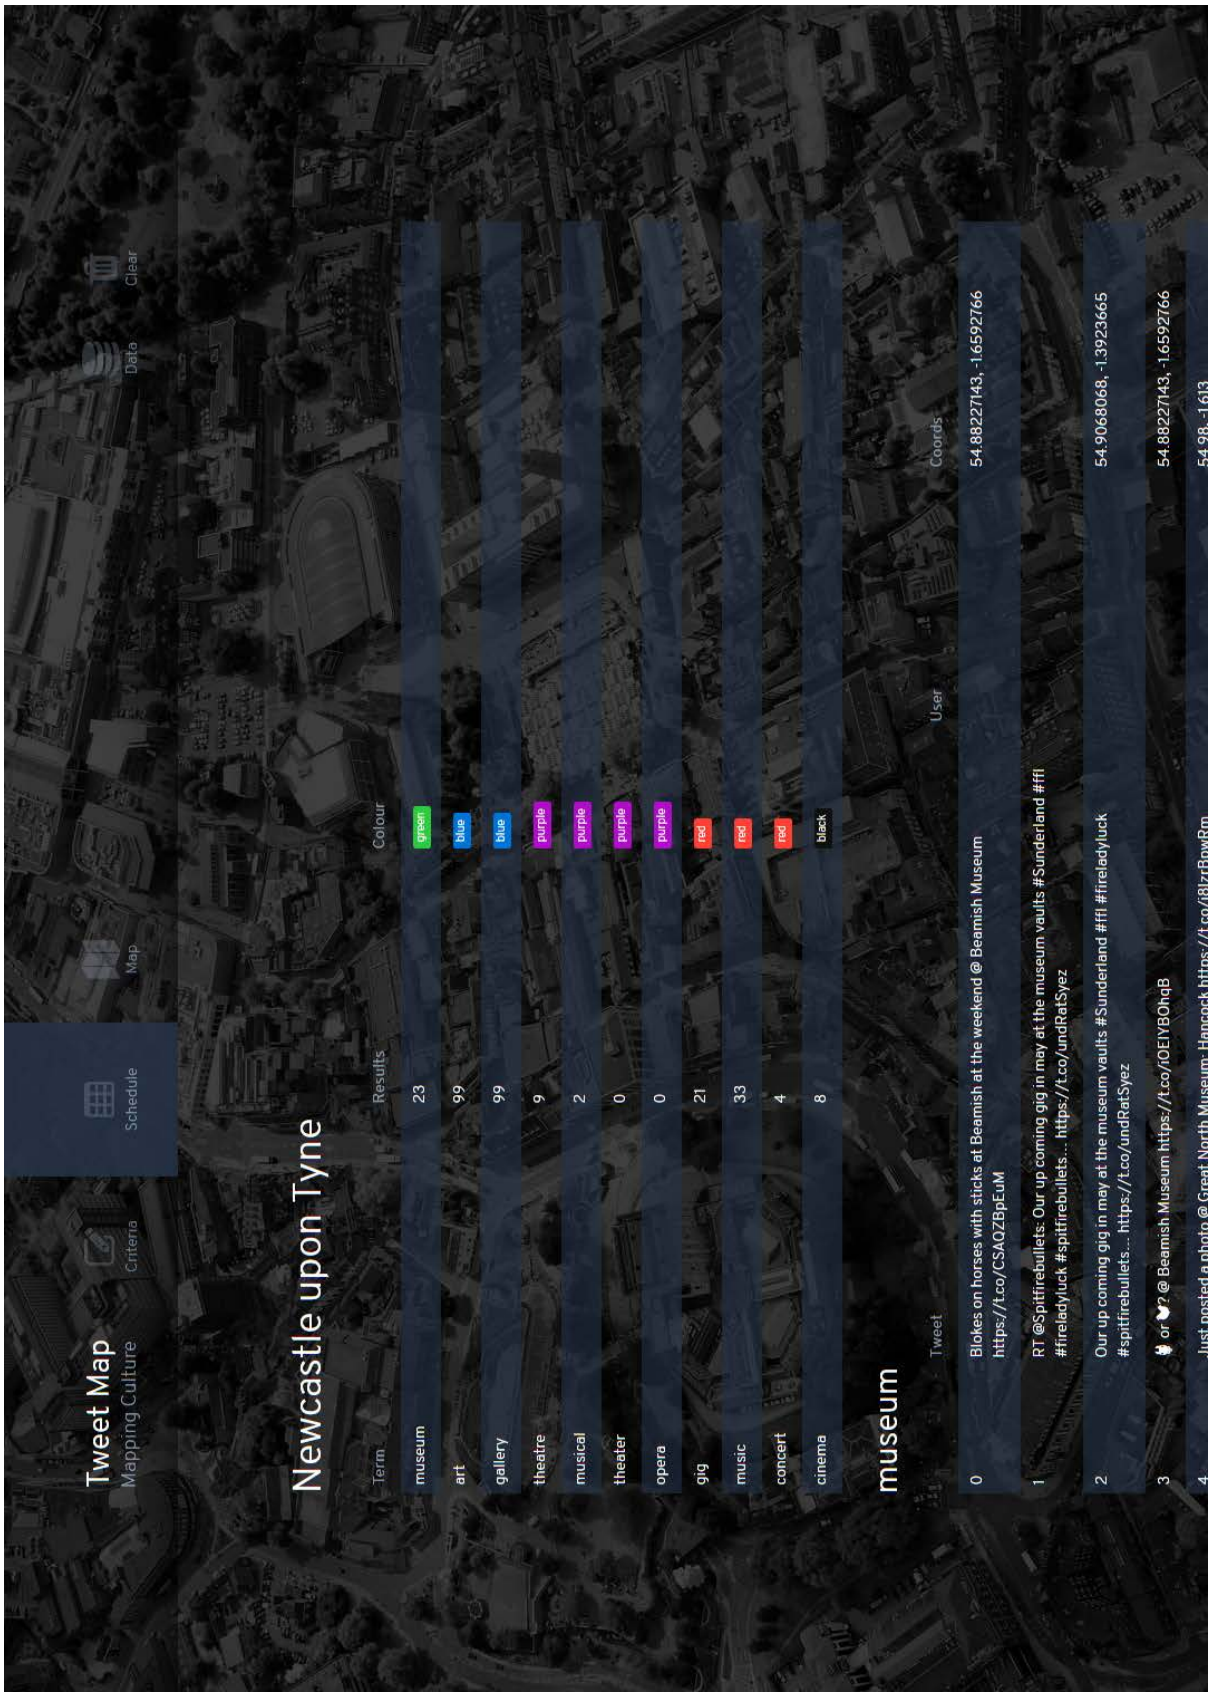

Figure 32 – Newcastle upon Tyne results table

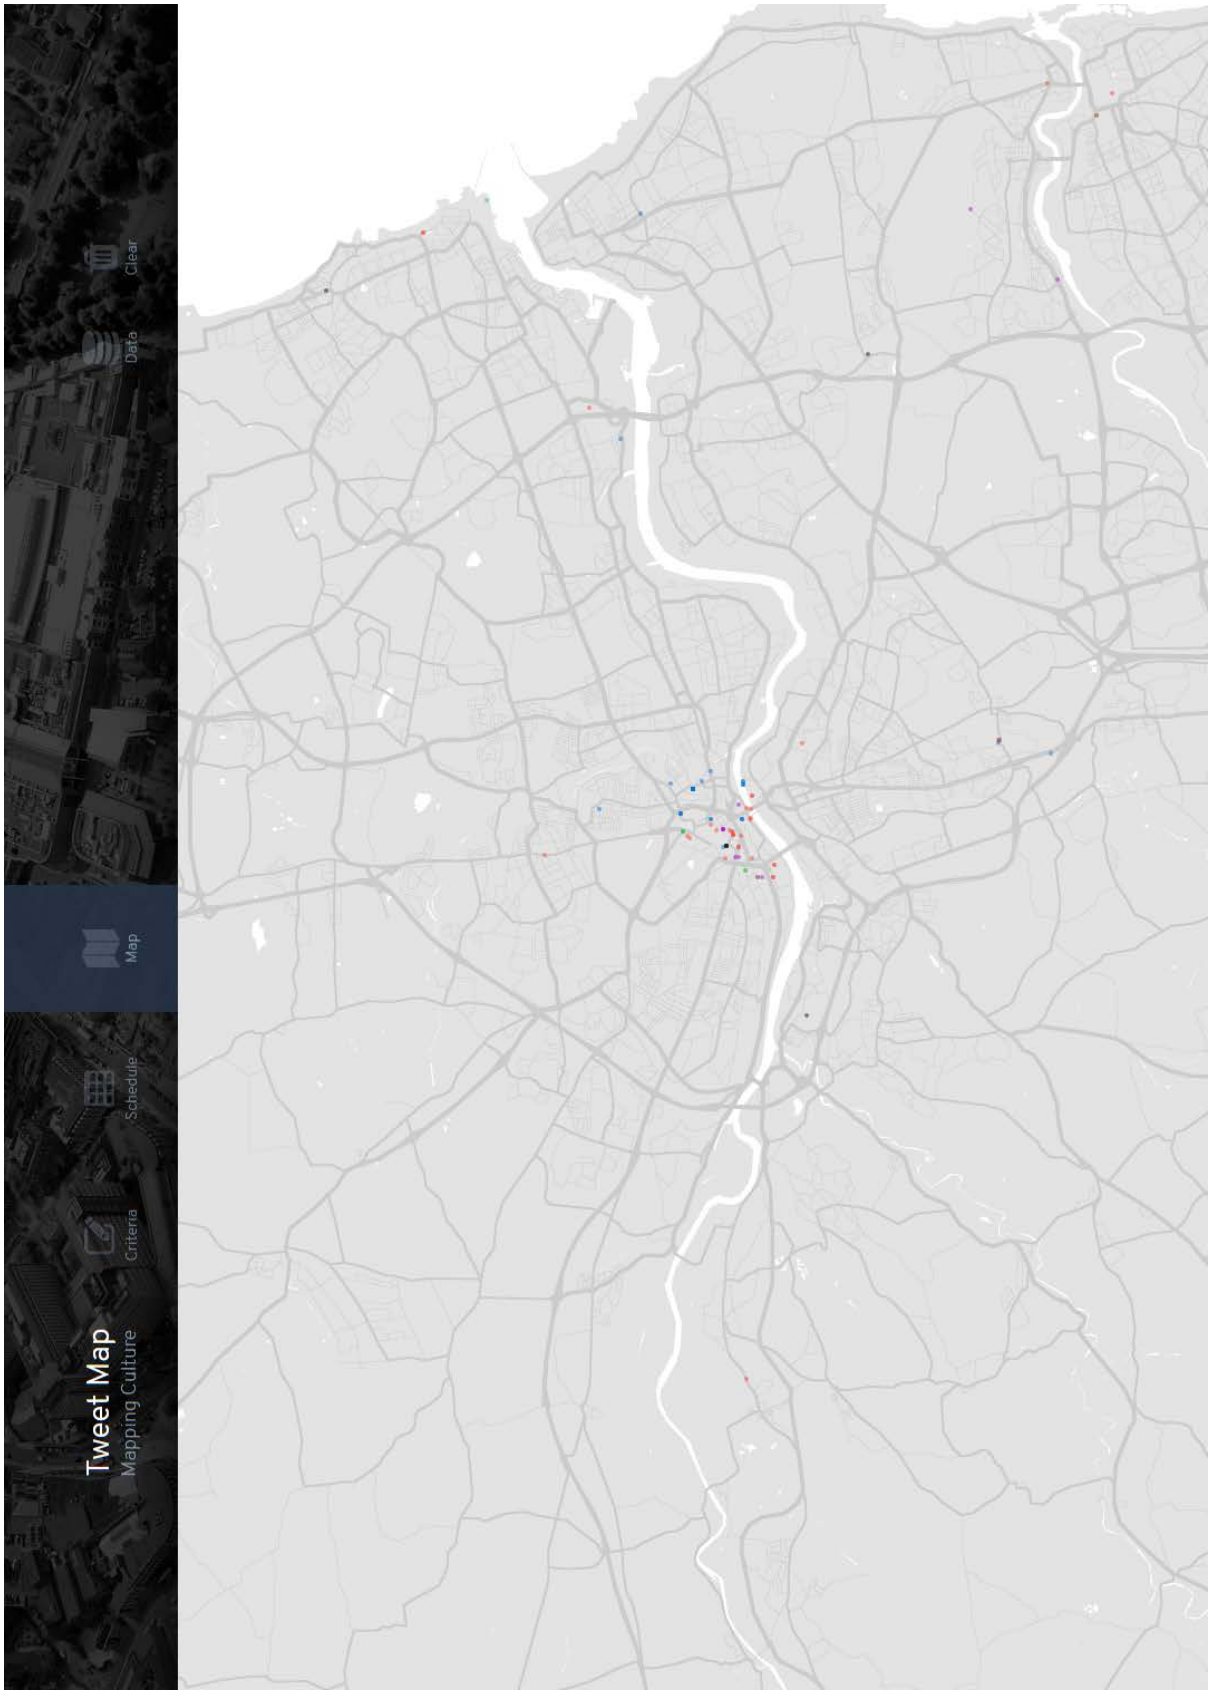

Figure 33 – Newcastle upon Tyne tweet culture map at zoom level 12

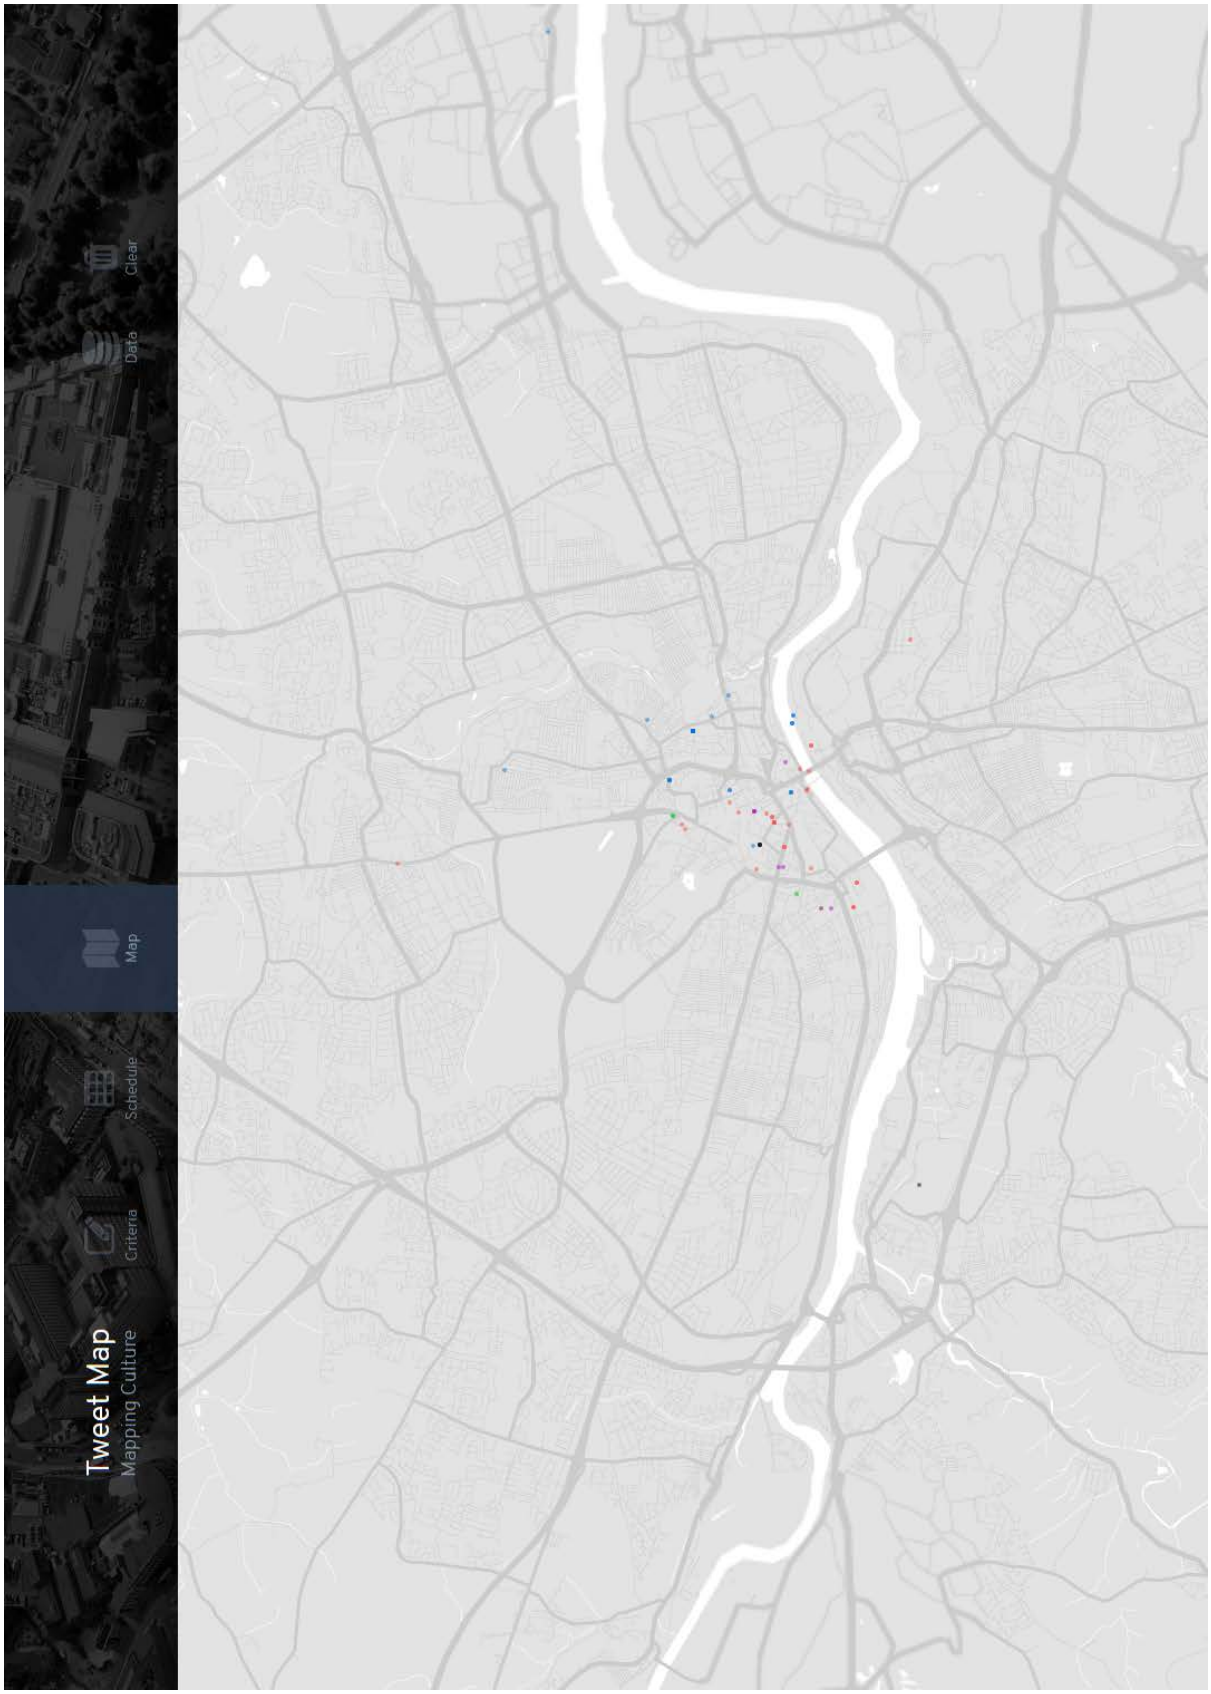

Figure 34 – Newcastle upon Tyne tweet culture map at zoom level 13

Figure 35 – Great North Museum, Newcastle upon Tyne

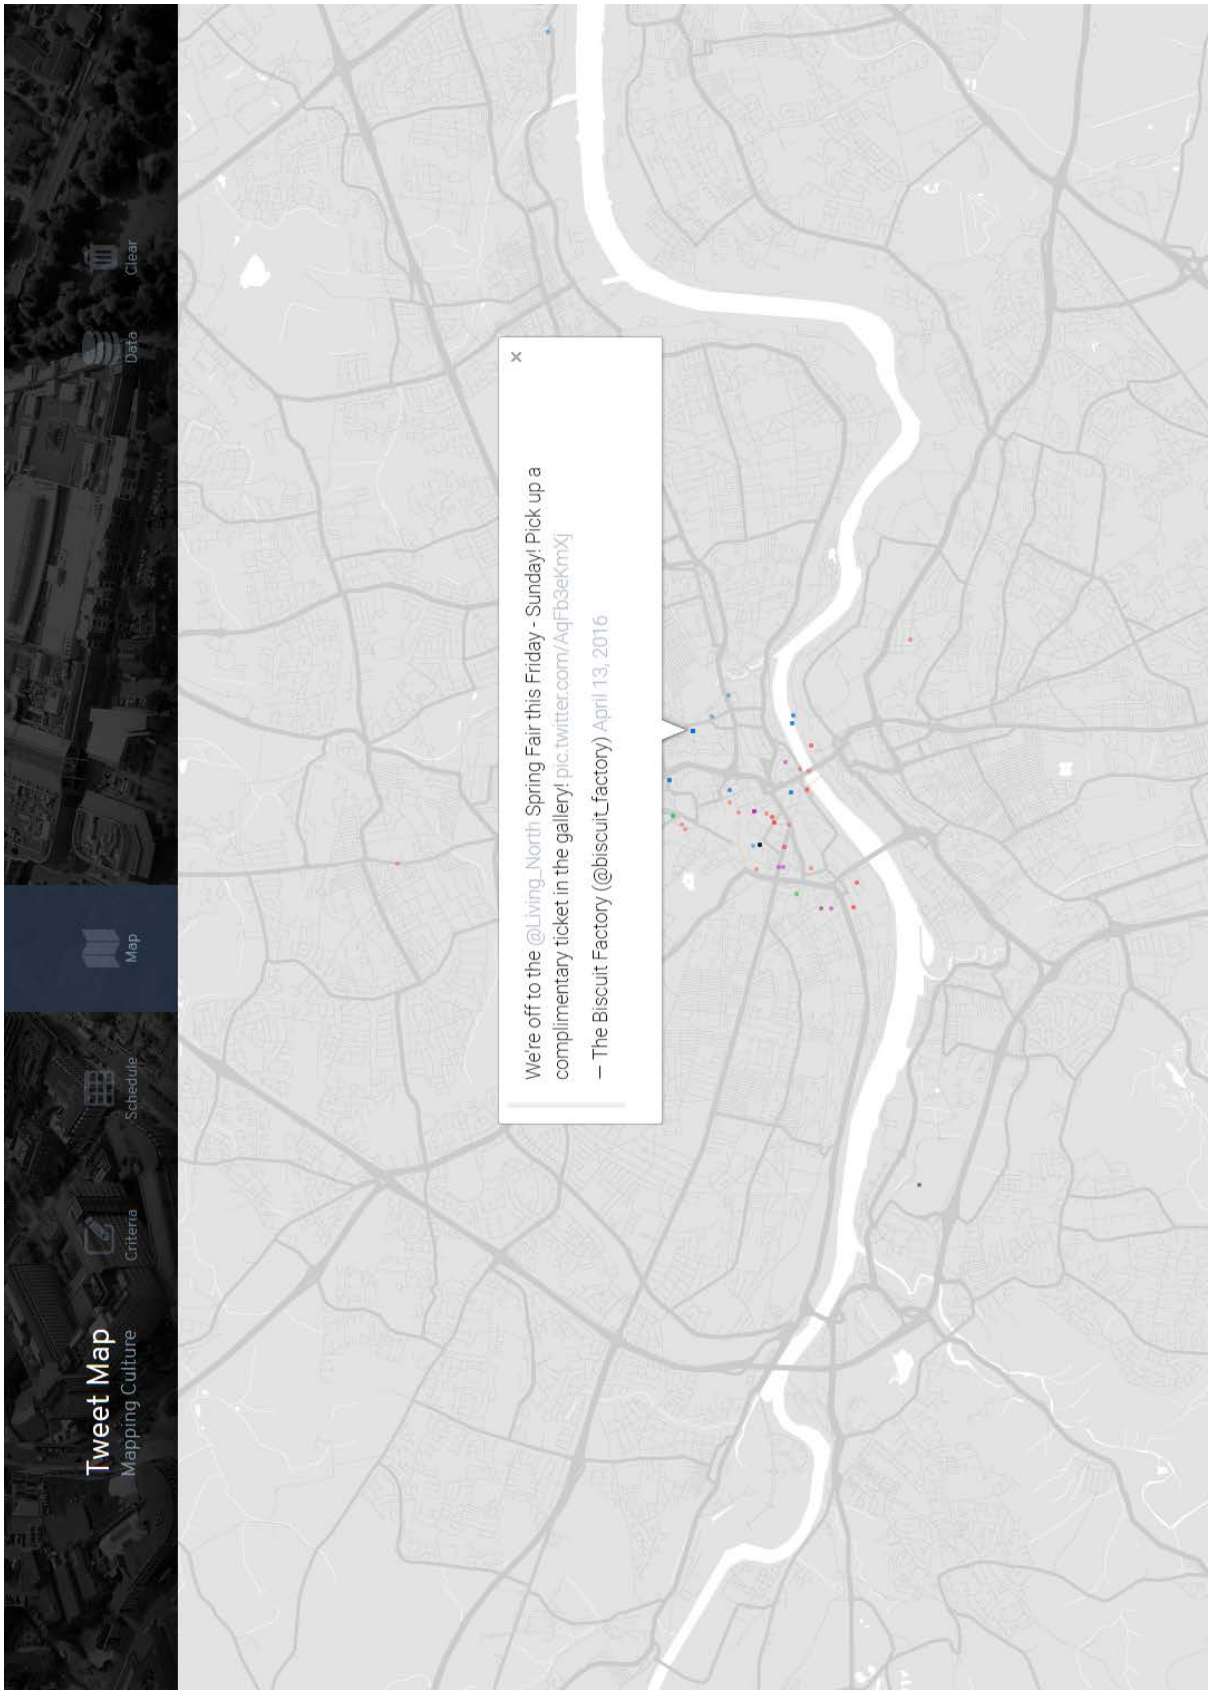

Figure 36 – The Biscuit Factory, Newcastle upon Tyne

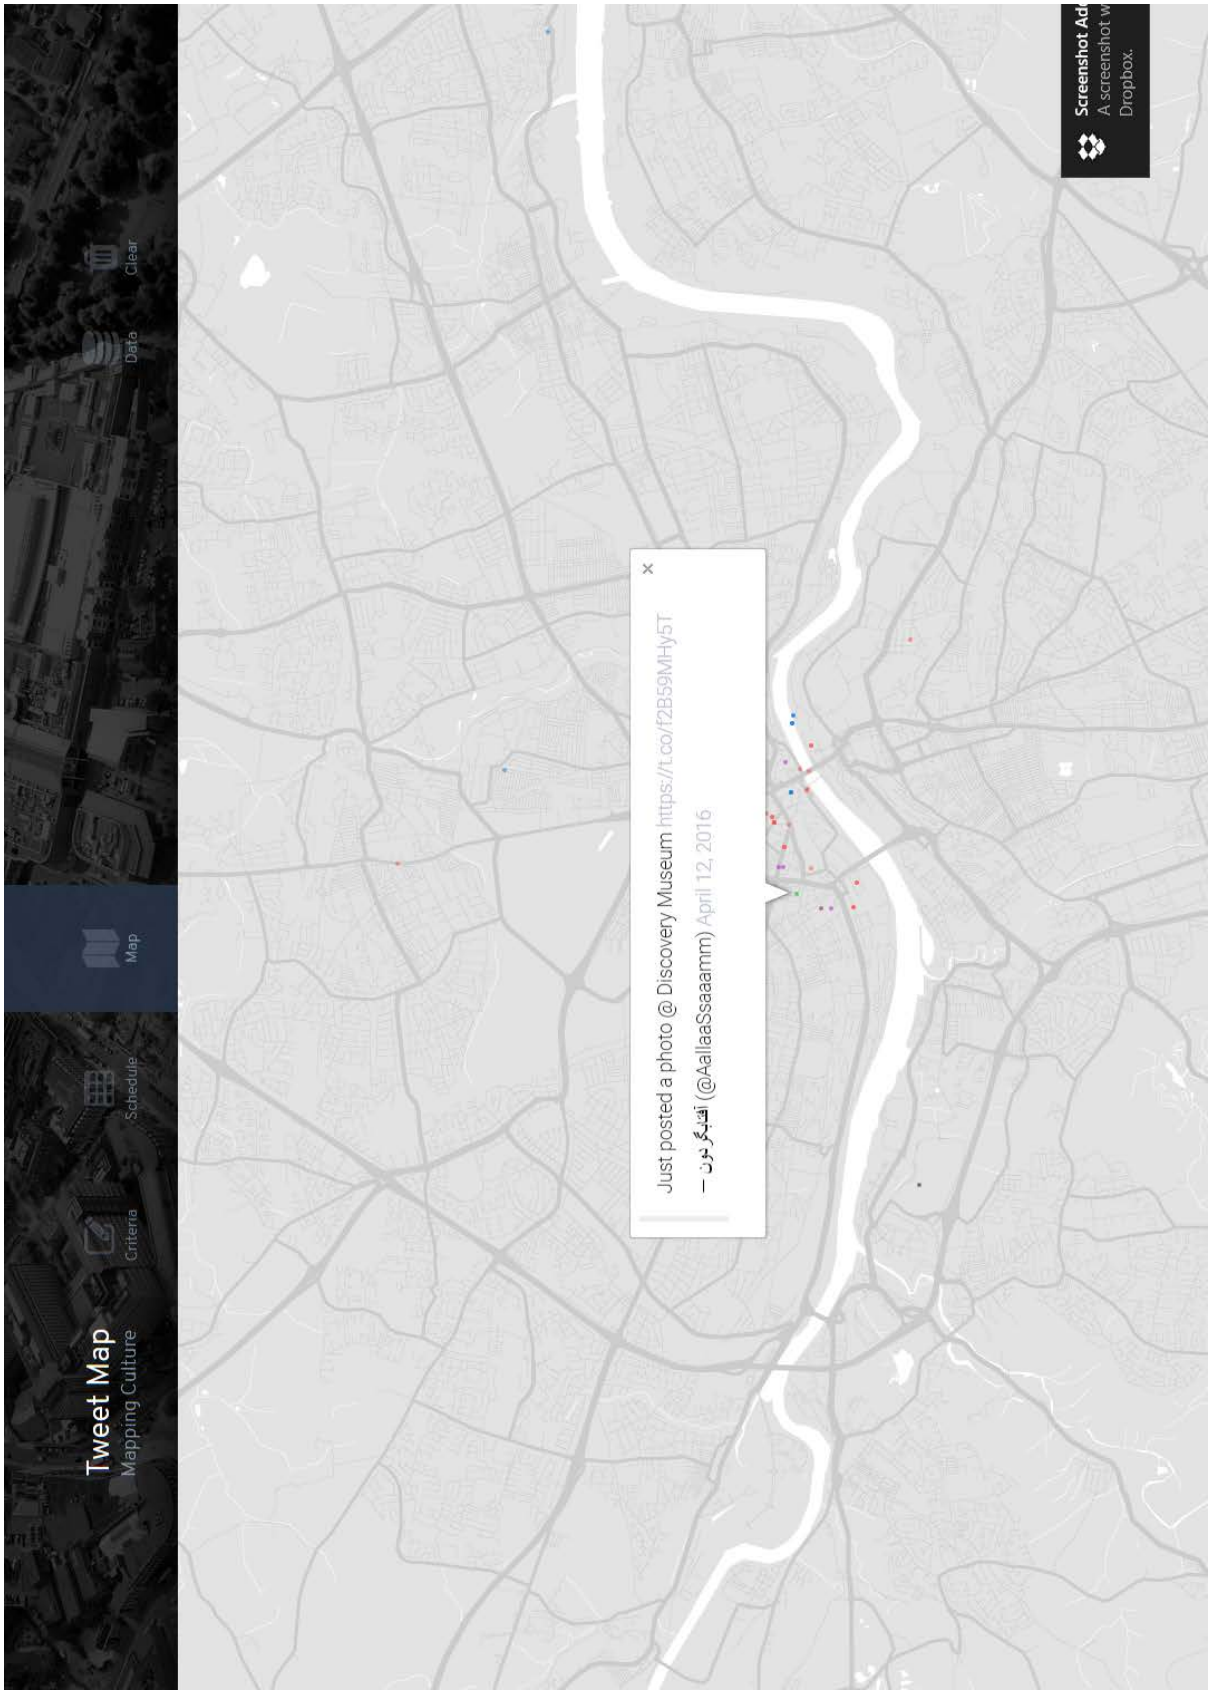

Figure 37 – Discovery Museum, Newcastle upon Tyne

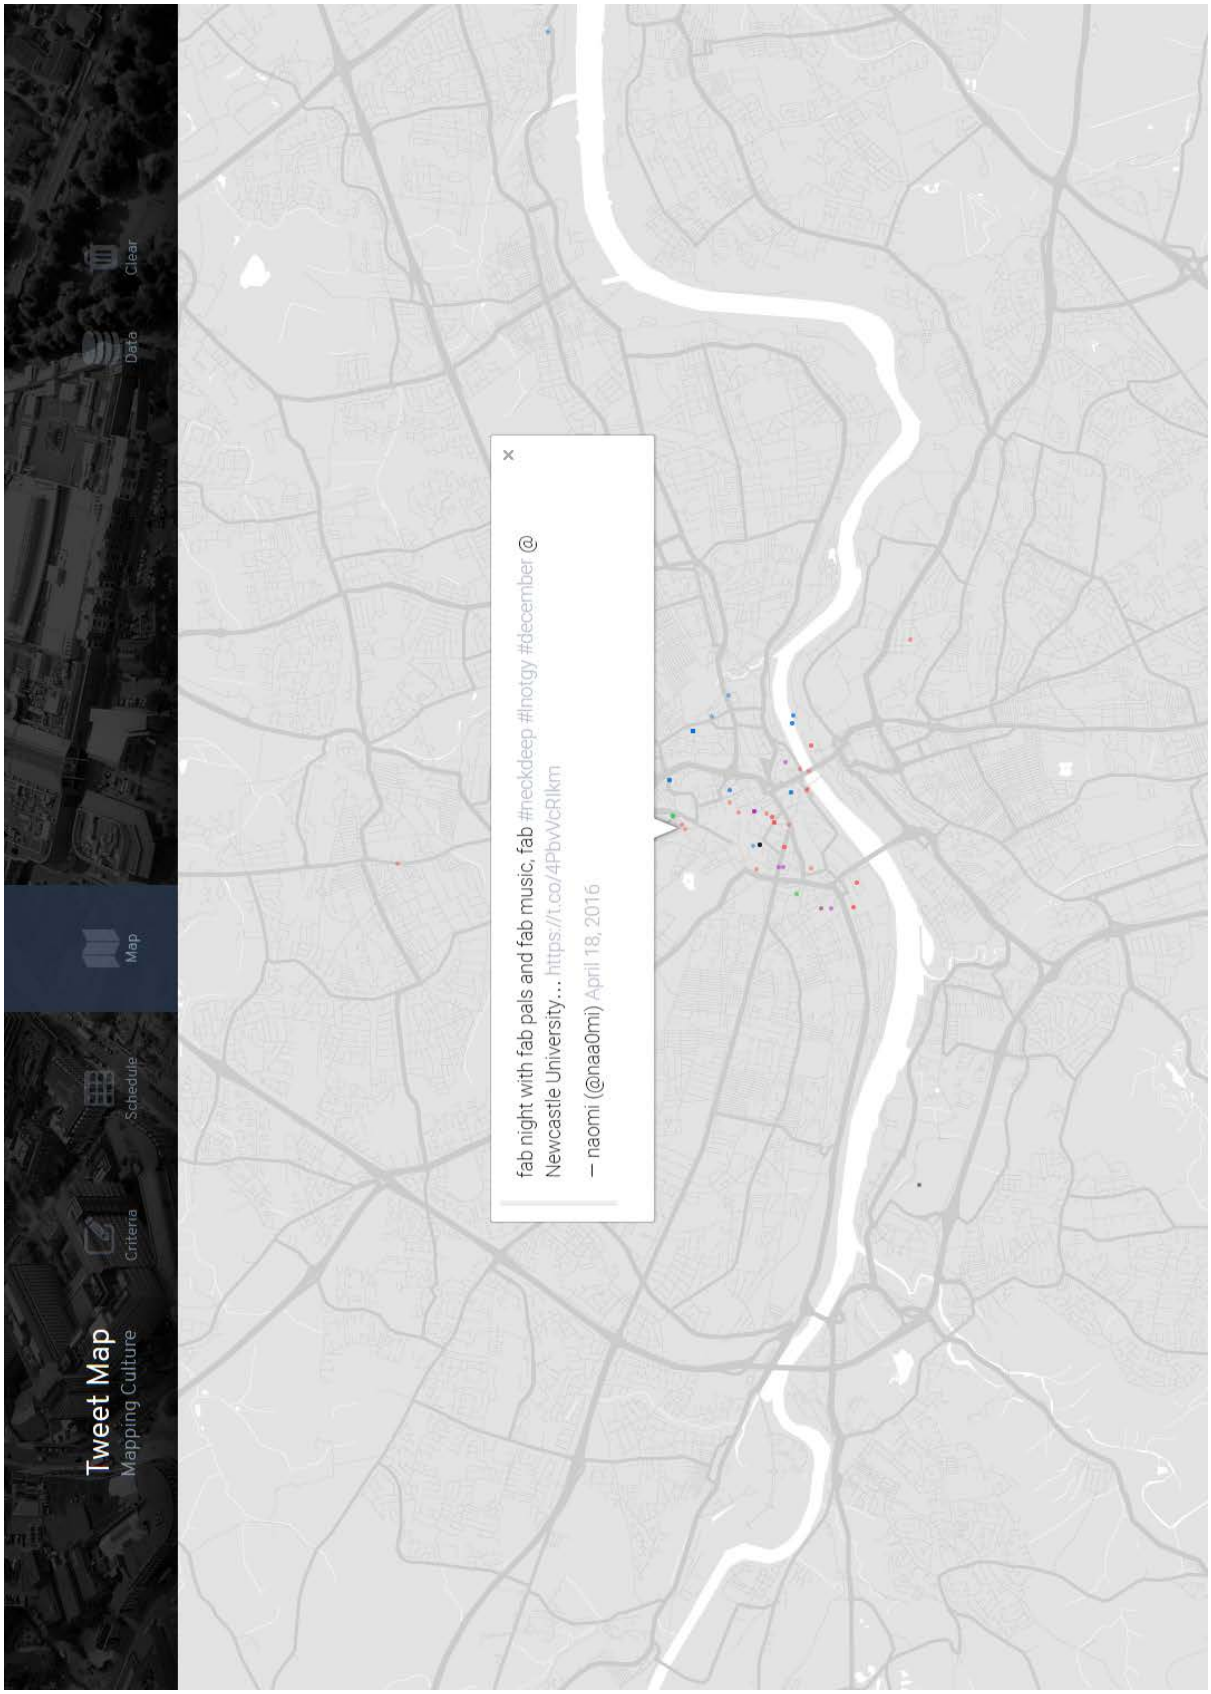

Figure 38 – Neck Deep at Newcastle University

Berlin

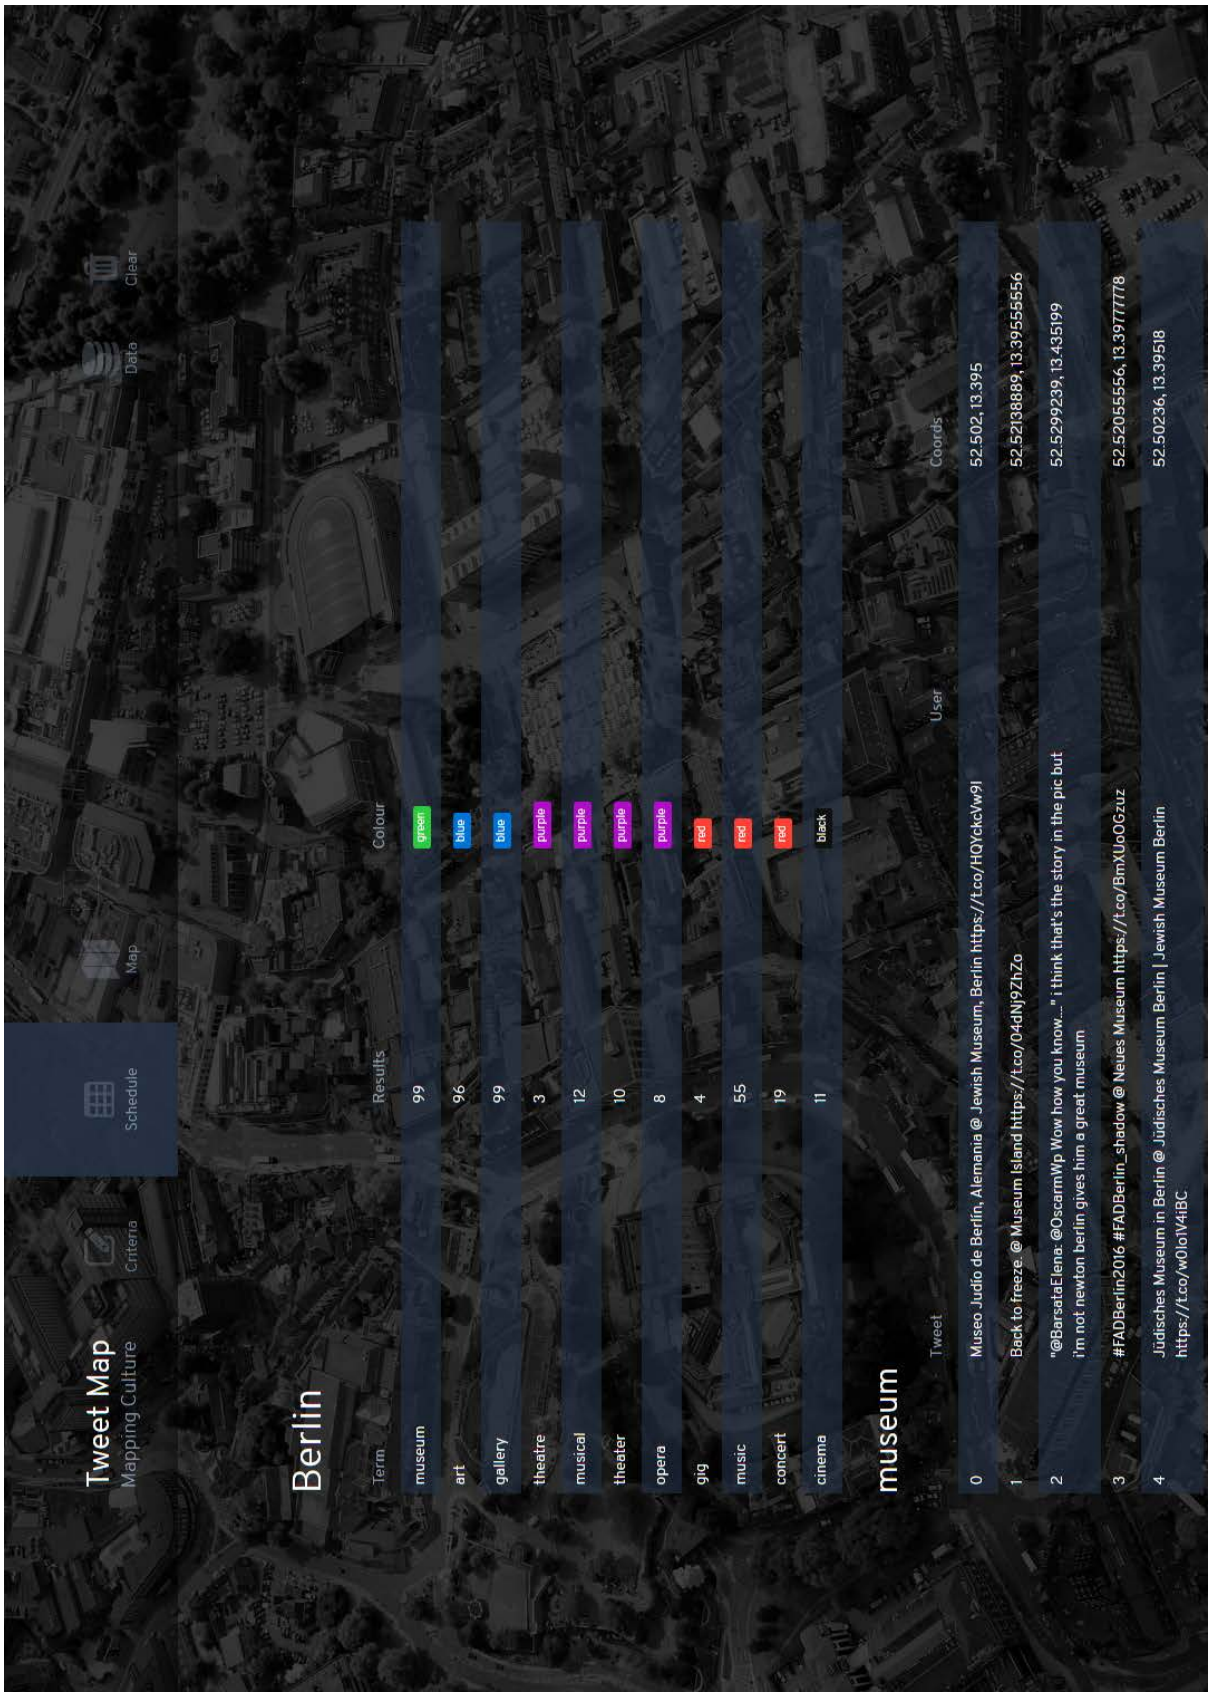

Figure 39 – Berlin results table

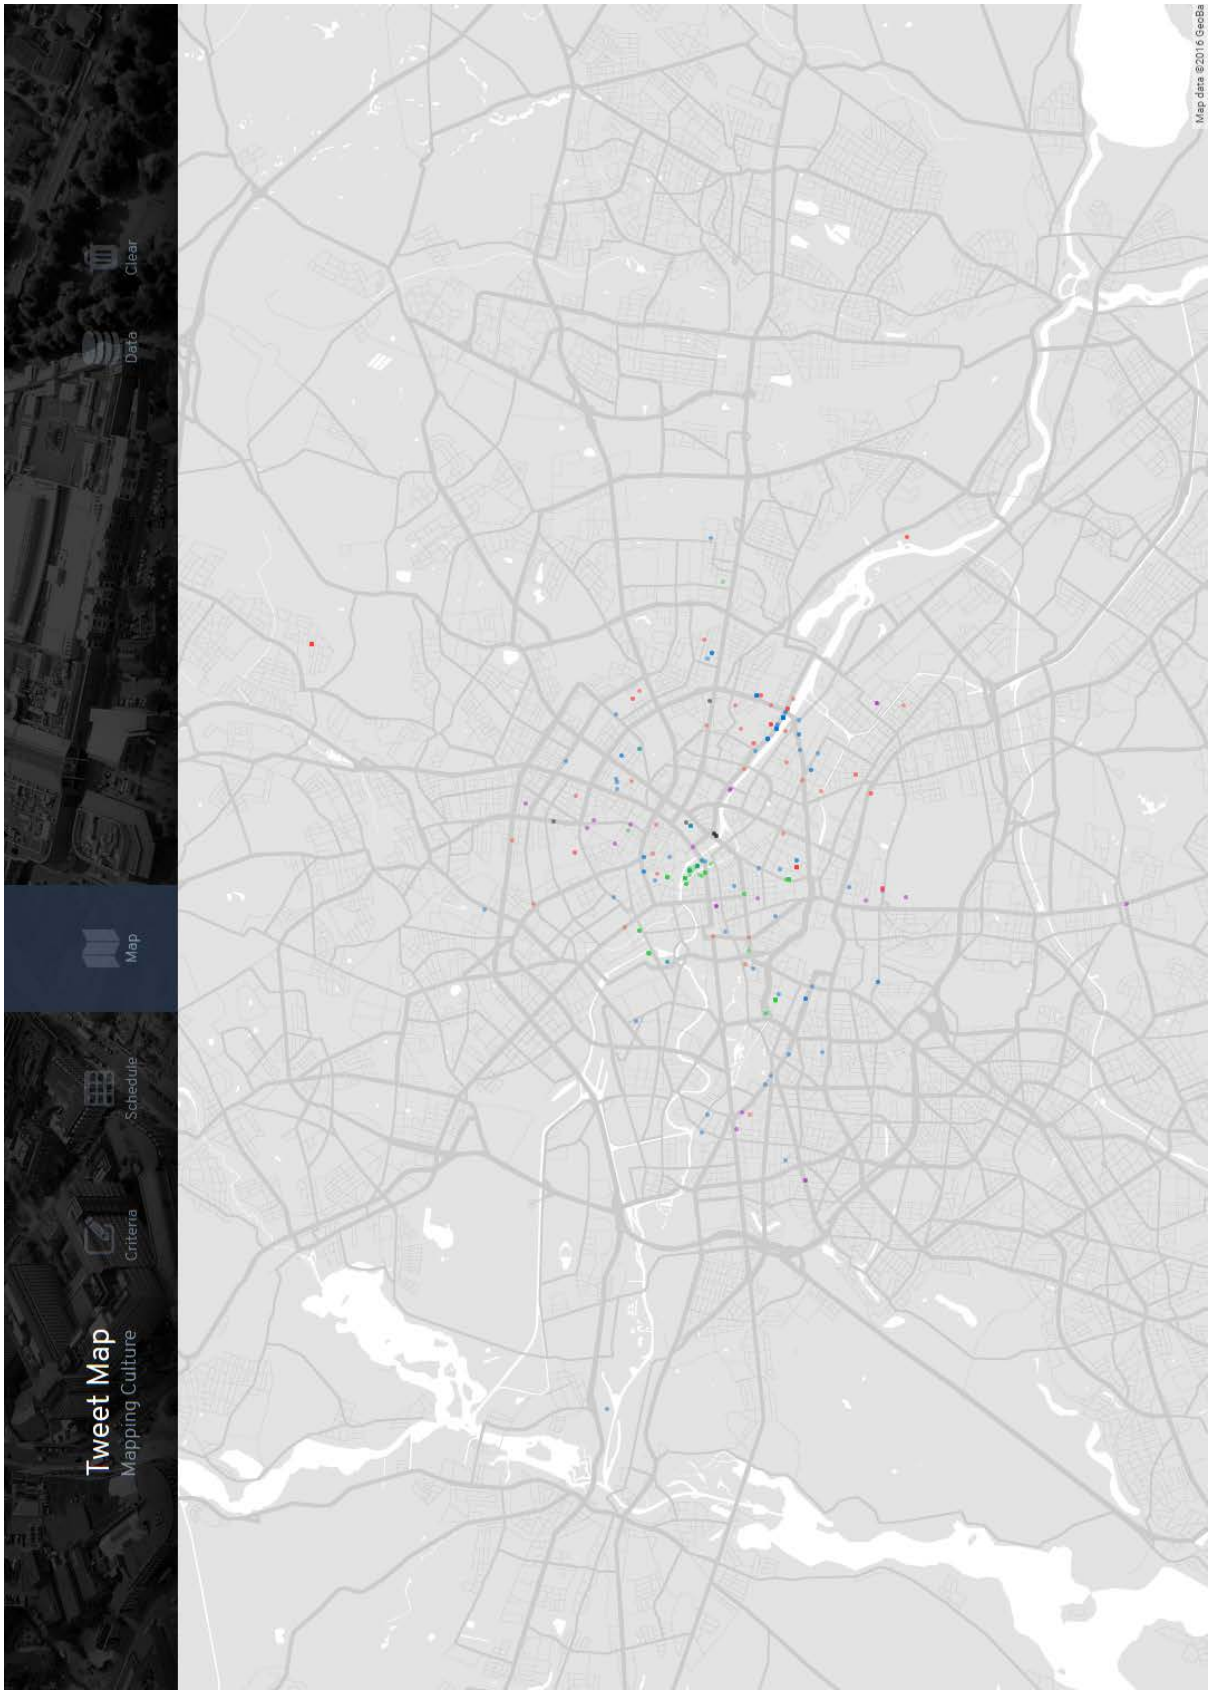

Figure 40 – Berlin tweet culture map at zoom level 12

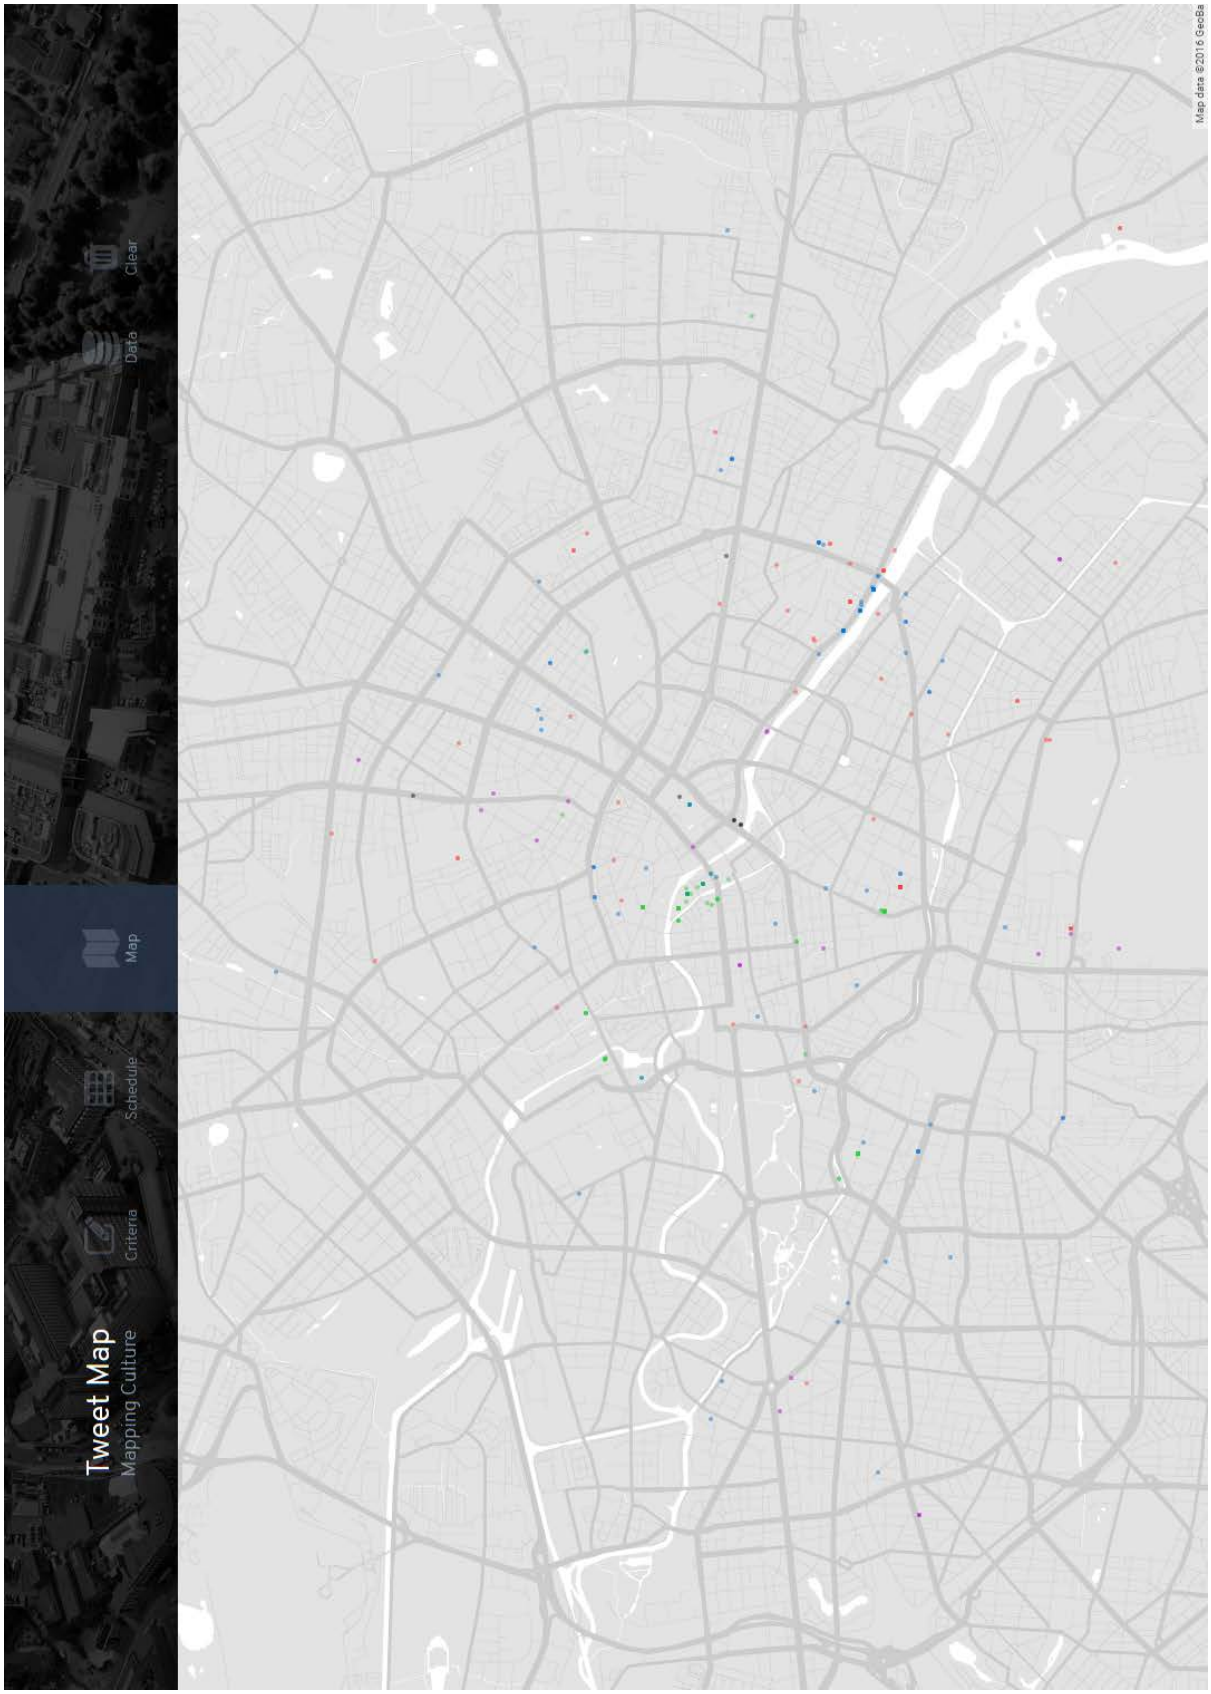

Figure 41 – Berlin tweet culture map at zoom level 13

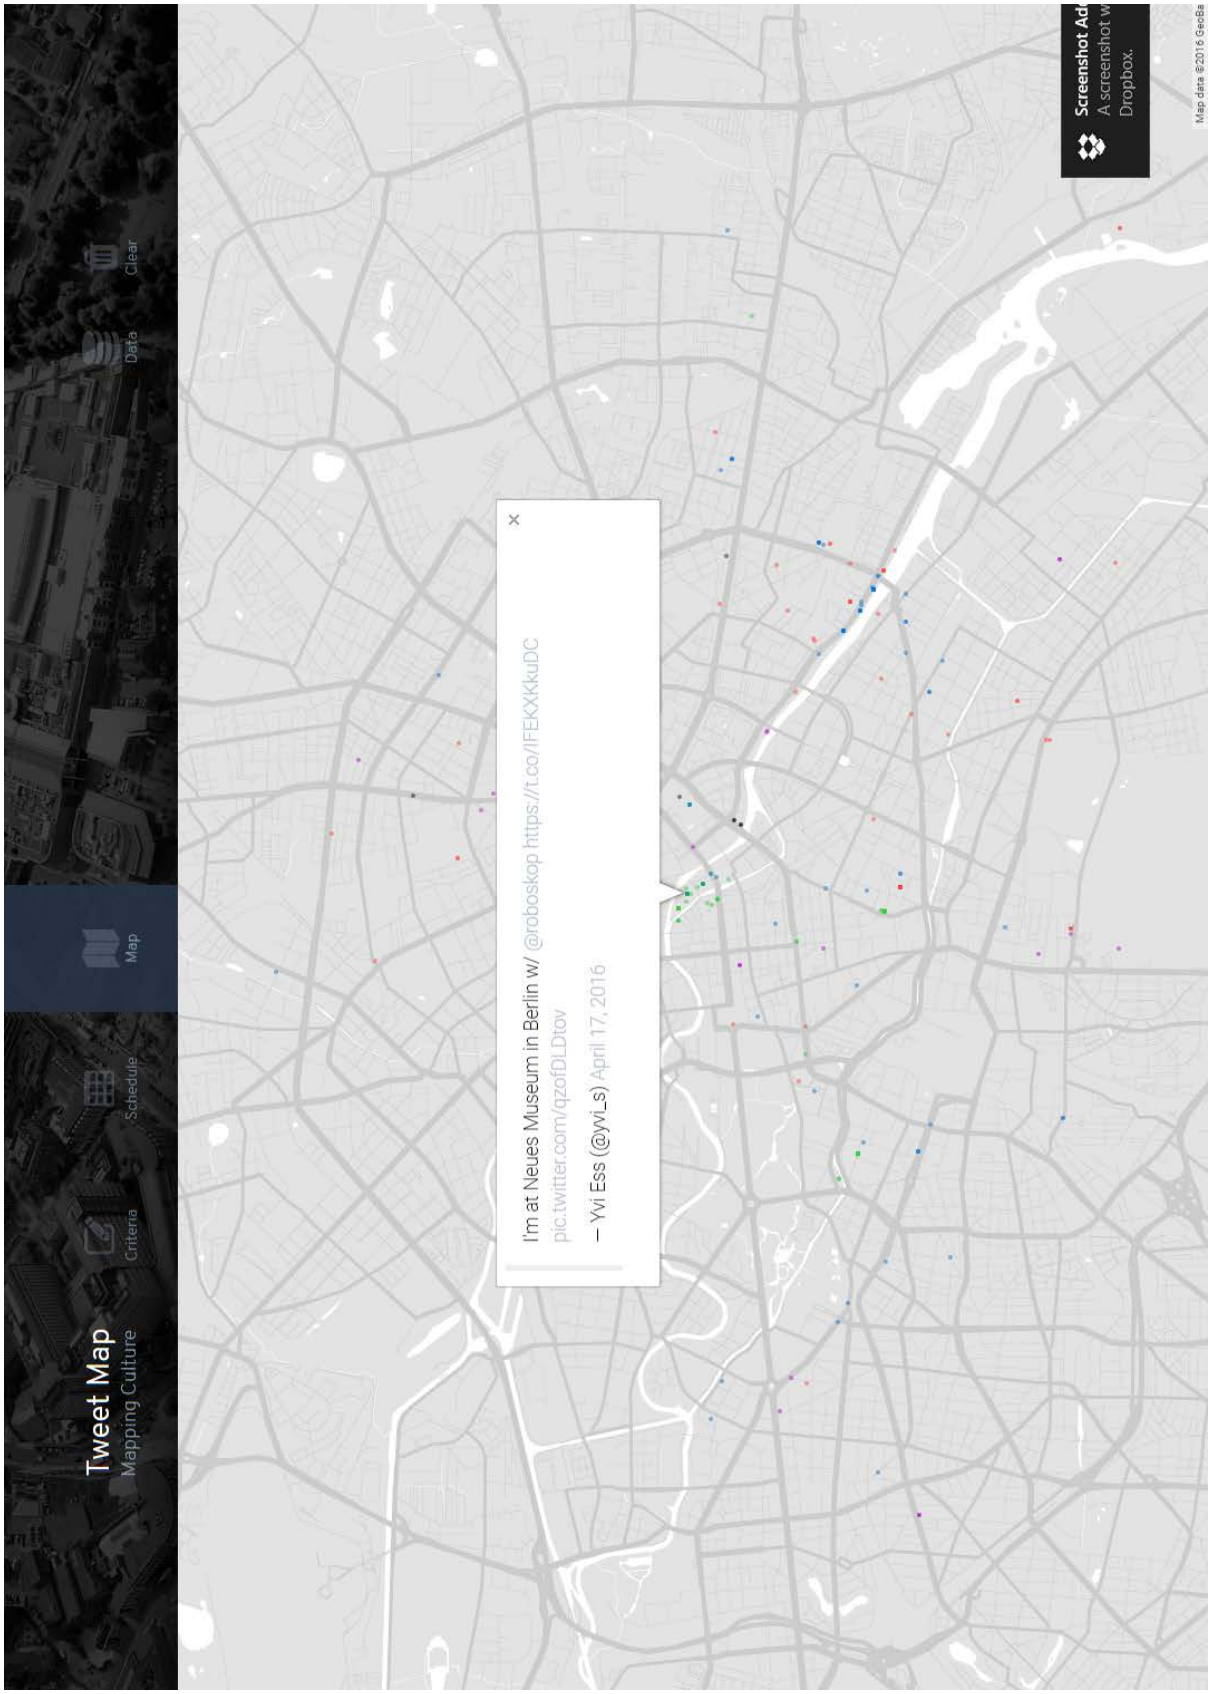

Figure 42 – Neues Museum, Berlin

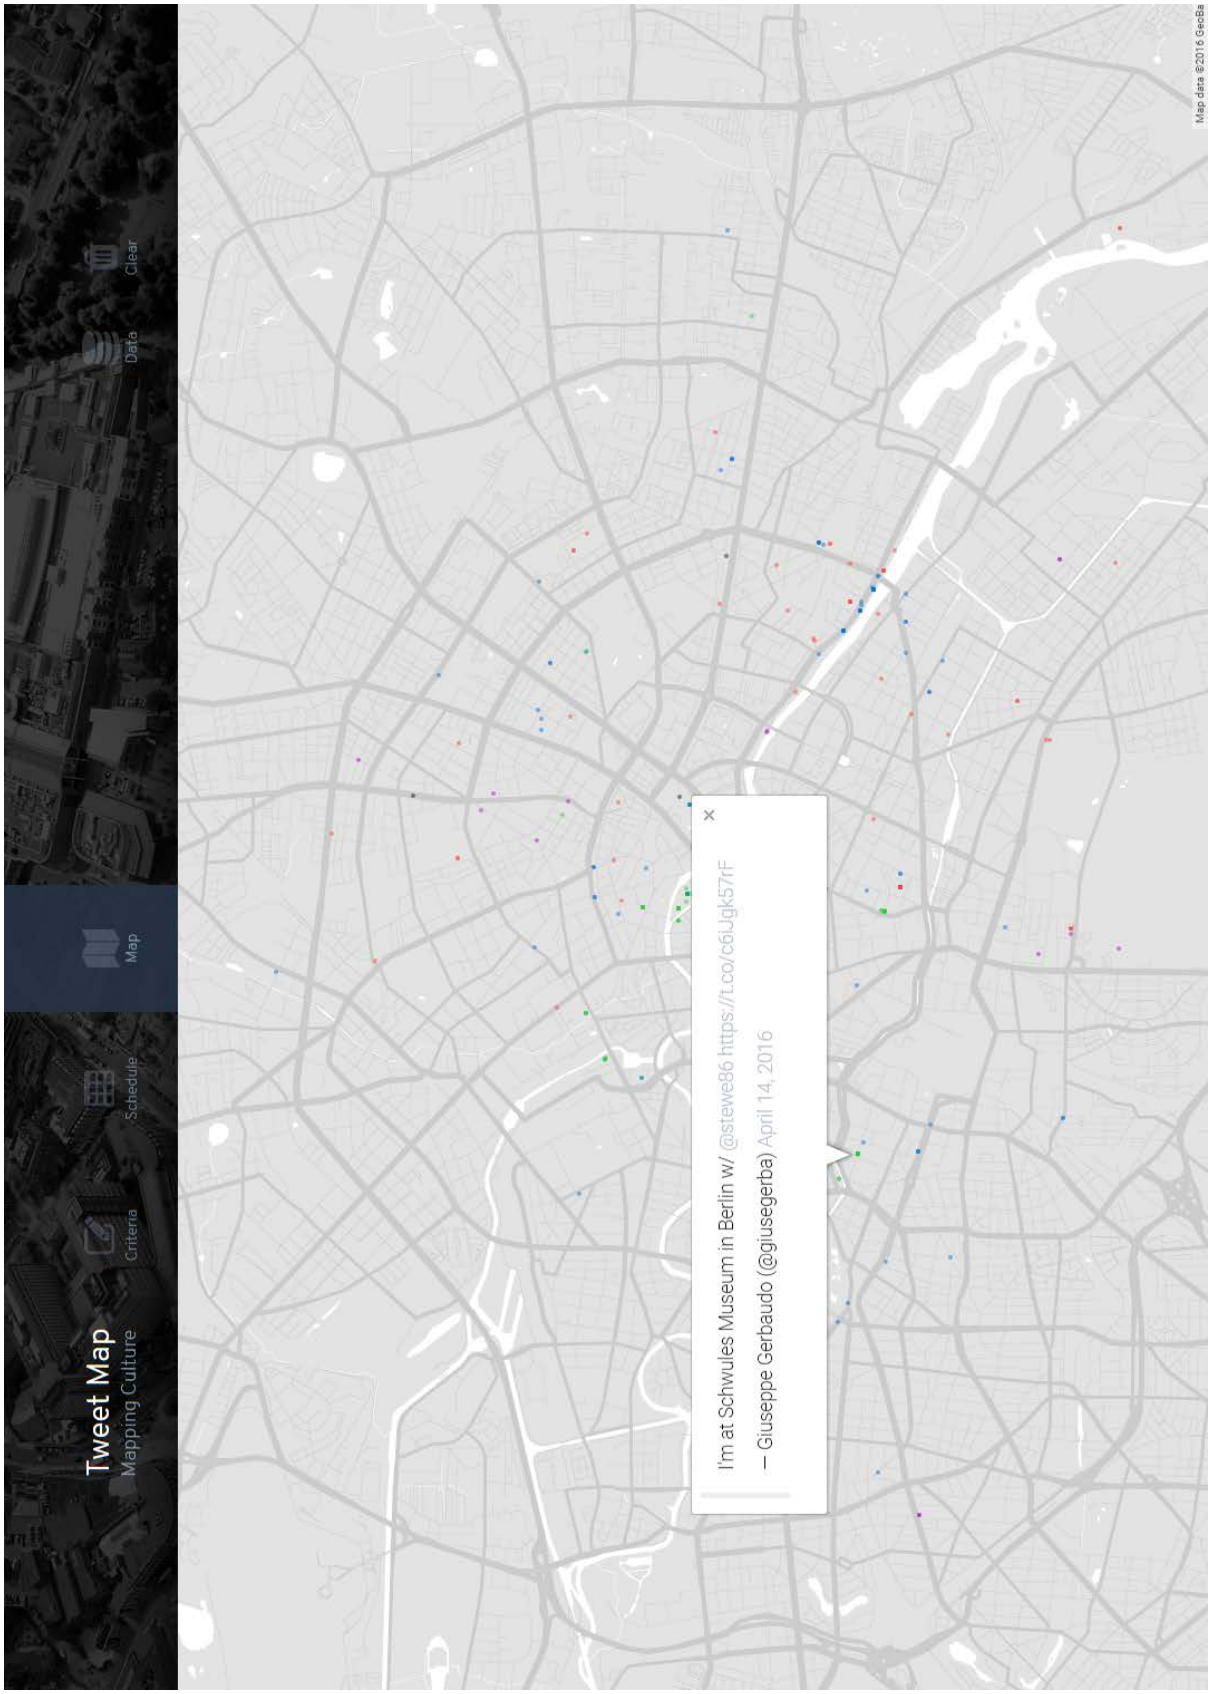

Figure 43 – Schwules Museum, Berlin

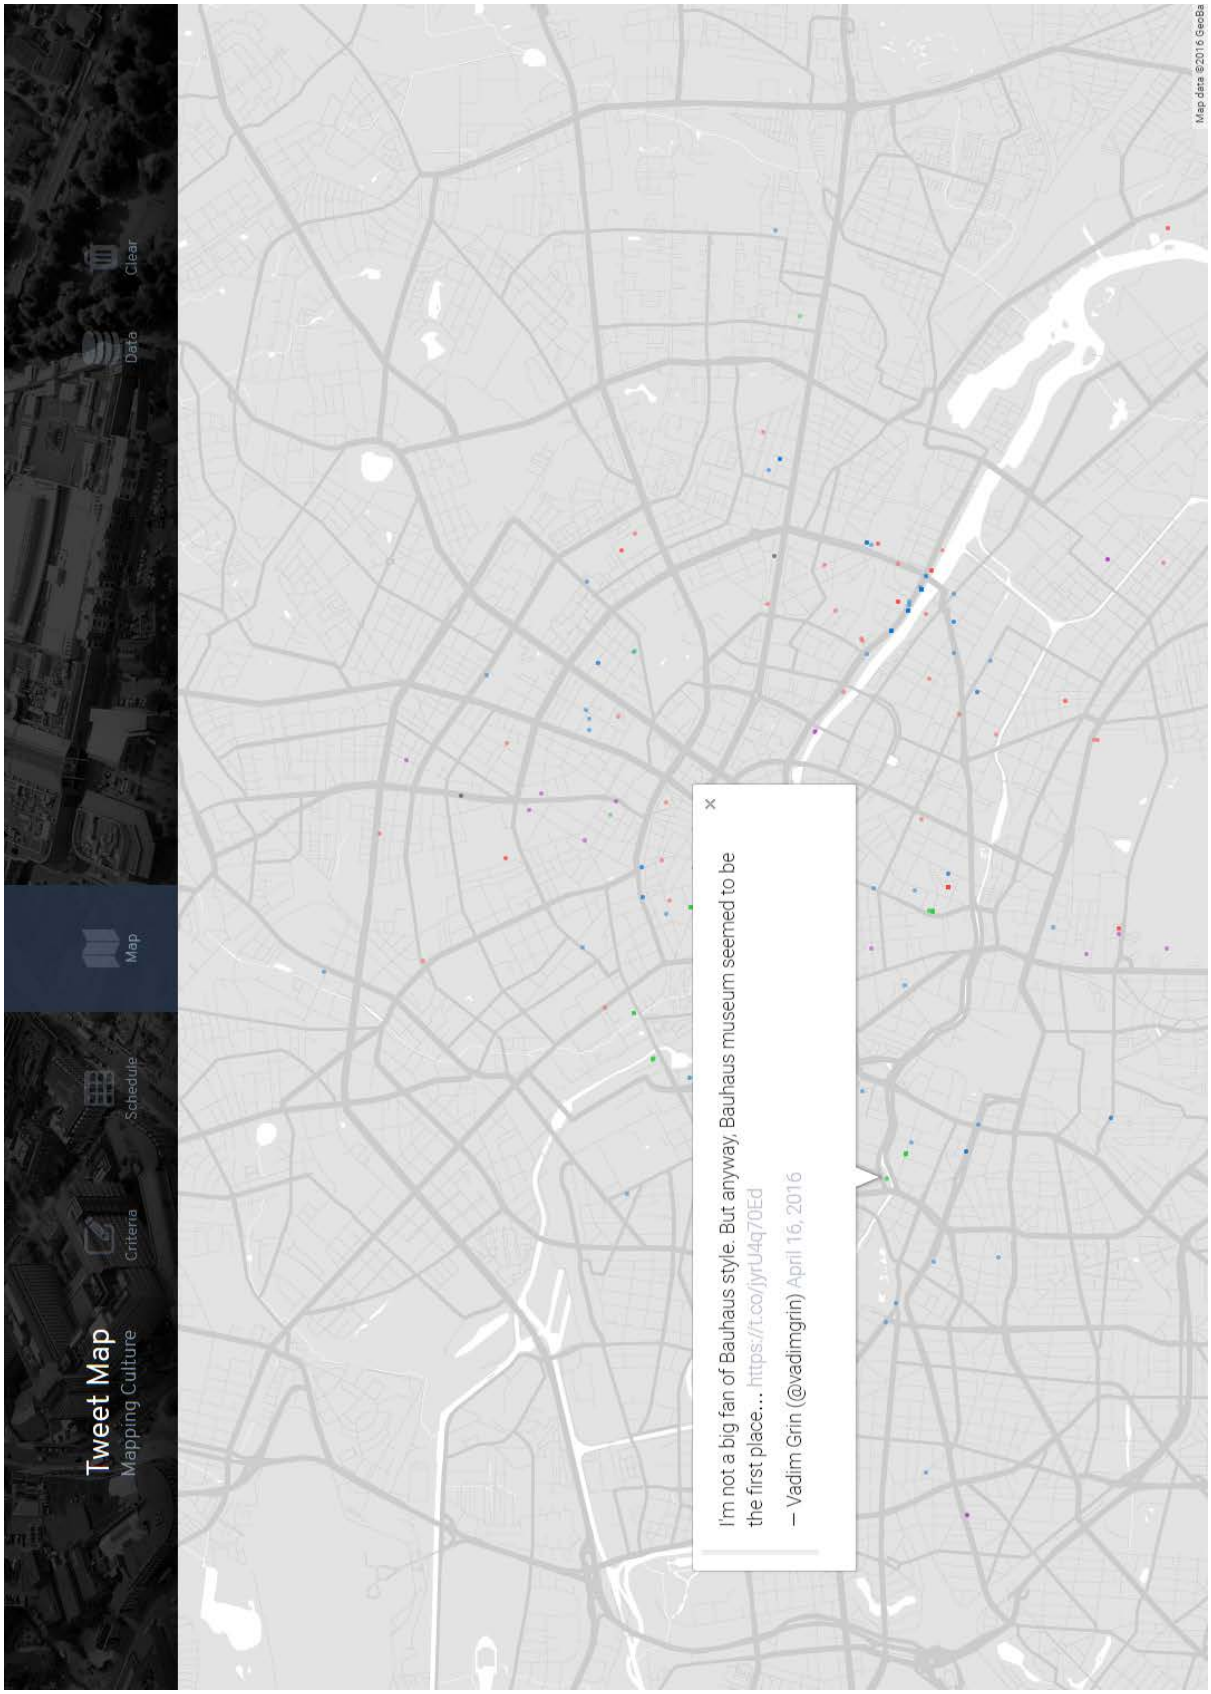

Figure 44 – Bauhaus Museum, Berlin

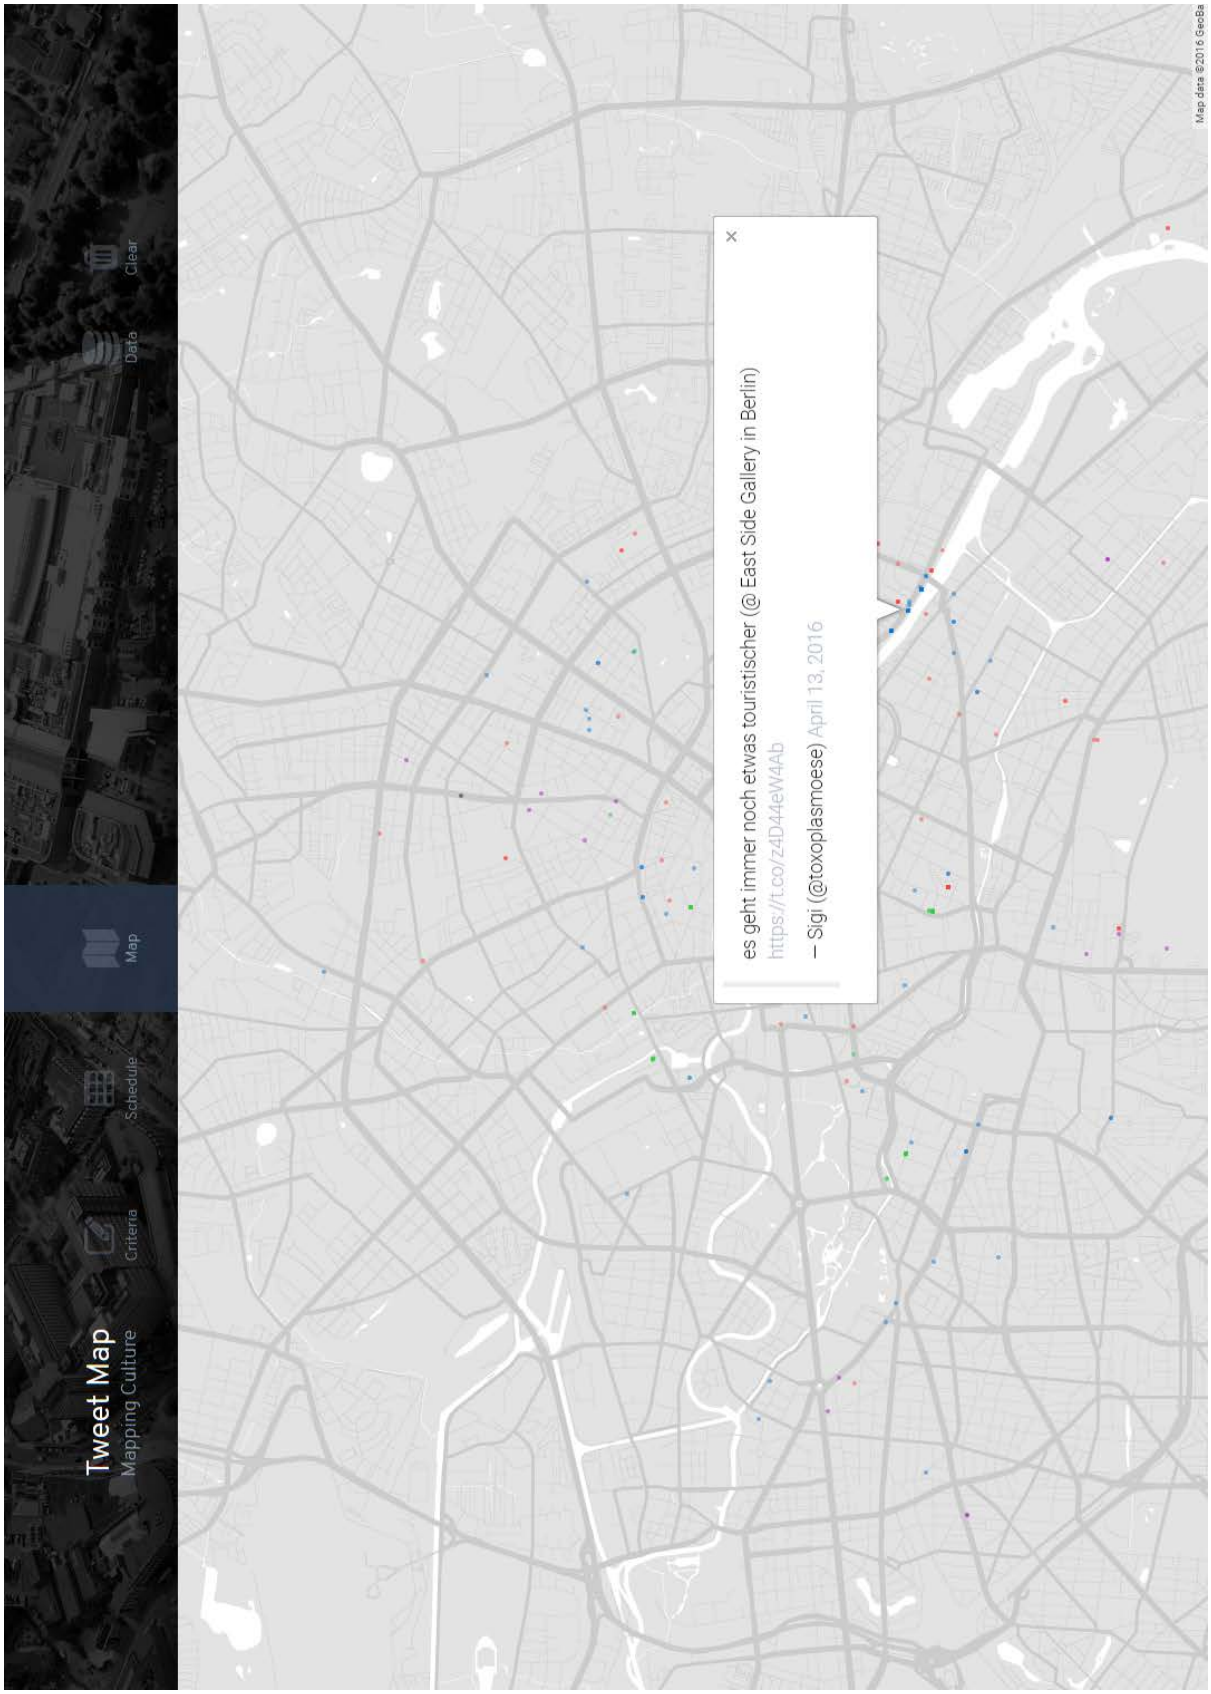

Figure 45 – East Side Gallery, Berlin

07 Discussion

London

The aim of this research was to identify the cultural loci of the city by analysing geolocated social media messages. The application that was developed within the methodology has mapped all tweets from the last 7 days that contain the selected cultural keywords. Even from a brief glance at the results in Figure 22 and Figure 23 it is evident that this has been achieved. Most notably there is a significant clustering of purple markers – indicating theatres – spread across the West End of London (Figure 24), just as would be expected.

To a lesser extent we can see clusters of museums – green markers – in the South Kensington area at the Science Museum and Natural History Museum (Figure 25), south of the river there are clusters of museum markers at the Imperial War Museum (Figure 26).

Art markers – blue – are significantly more spread across the capital, this could be due to the diverse nature of the word art which could appear within social media messages, it could just be through the prevalence of art across the city or additional research may be able to conclude that art is heavily ingrained in our culture so is discussed a great deal through social media and not necessarily in relation to a place. Three notable clusters of art have been identified at the Serpentine Gallery (Figure 27), Tate Modern (Figure 28) and around Shoreditch (Figure 30).

Similar to art, music is mentioned abundantly across London the most notable musical clusters are on the Greenwich peninsular at The O2 (Figure 29), around Shoreditch and in Camden Town (Figure 31).

Newcastle upon Tyne

The same methodology has been applied to Newcastle upon Tyne with the results shown in Figure 33 and Figure 34. Due to Newcastle's significantly lower population – 289,835 compared to 8,538,689 (ONS, 2015) – and the lower number of tourists there are significantly fewer tweets, this makes it very difficult to identify clusters which can be defined as cultural loci. It can still be noted that a few tweets appear at the Great North Museum (Figure 35) and the Biscuit Factory (Figure 36), however there is insufficient data to define these as cultural loci.

Berlin

Berlin was chosen for two reasons, primarily, it represents a foreign city where English is not the native language. Secondly, the urban planning of Berlin has led to the creation of Museumsinsel – the Museum Island – which should be strongly identifiable as a cultural loci within the results.

The exact same search terms have been applied to Berlin which has a population 41% the size of London – 3.502 million (UNdata, 2015) – while the tweet density is significantly less than London there are still a few identifiable clusters (Figure 40 and Figure 41).

First, as expected the Museum Island (Figure 42) is lit up with green markers, clearly identifying a cultural loci, another smaller museum cluster is visible to the West around the Schwules Museums (Figure 43). On the topic of art the East Side Gallery is a strongly representative cultural loci along Mühlenstraße (Figure 45).

Summary

Clusters of geolocated tweets have been observed which identify cultural loci in both London and Berlin, however, the density of tweets within Newcastle upon Tyne was not sufficient to identify any clusters.

Limitations

Although the methodology and the developed application has been successful in identifying cultural loci it has a number of limitations in its current form. The most significant limitation is from the restriction of the Twitter API to only data from the last 7 days.

The 7 day window is unlikely to reflect the cultural nature of the city throughout the year, it also means the data can be skewed by the presence or absence of cultural events. For example, Mardi Gras or Notting Hill Carnival are significant cultural events within a city but are unlikely to be reflected within the 7 day sample, and if they are included they will overwhelm other potential cultural loci.

The 7 day window also makes it difficult to analyse smaller cities as can be seen with the example of Newcastle. In order to obtain a suitable data sample the city would need to be sampled over a longer period, perhaps a month or quarter.

It must be noted that the demographic of Twitter users mean there is a bias within the data (Figure 46) 37% of 18-29 year olds use Twitter but only 12% of 50-64 and 10% of 65+ (Spruce Social, 2015) therefore the cultural loci identified are more likely to be representative of younger generations.

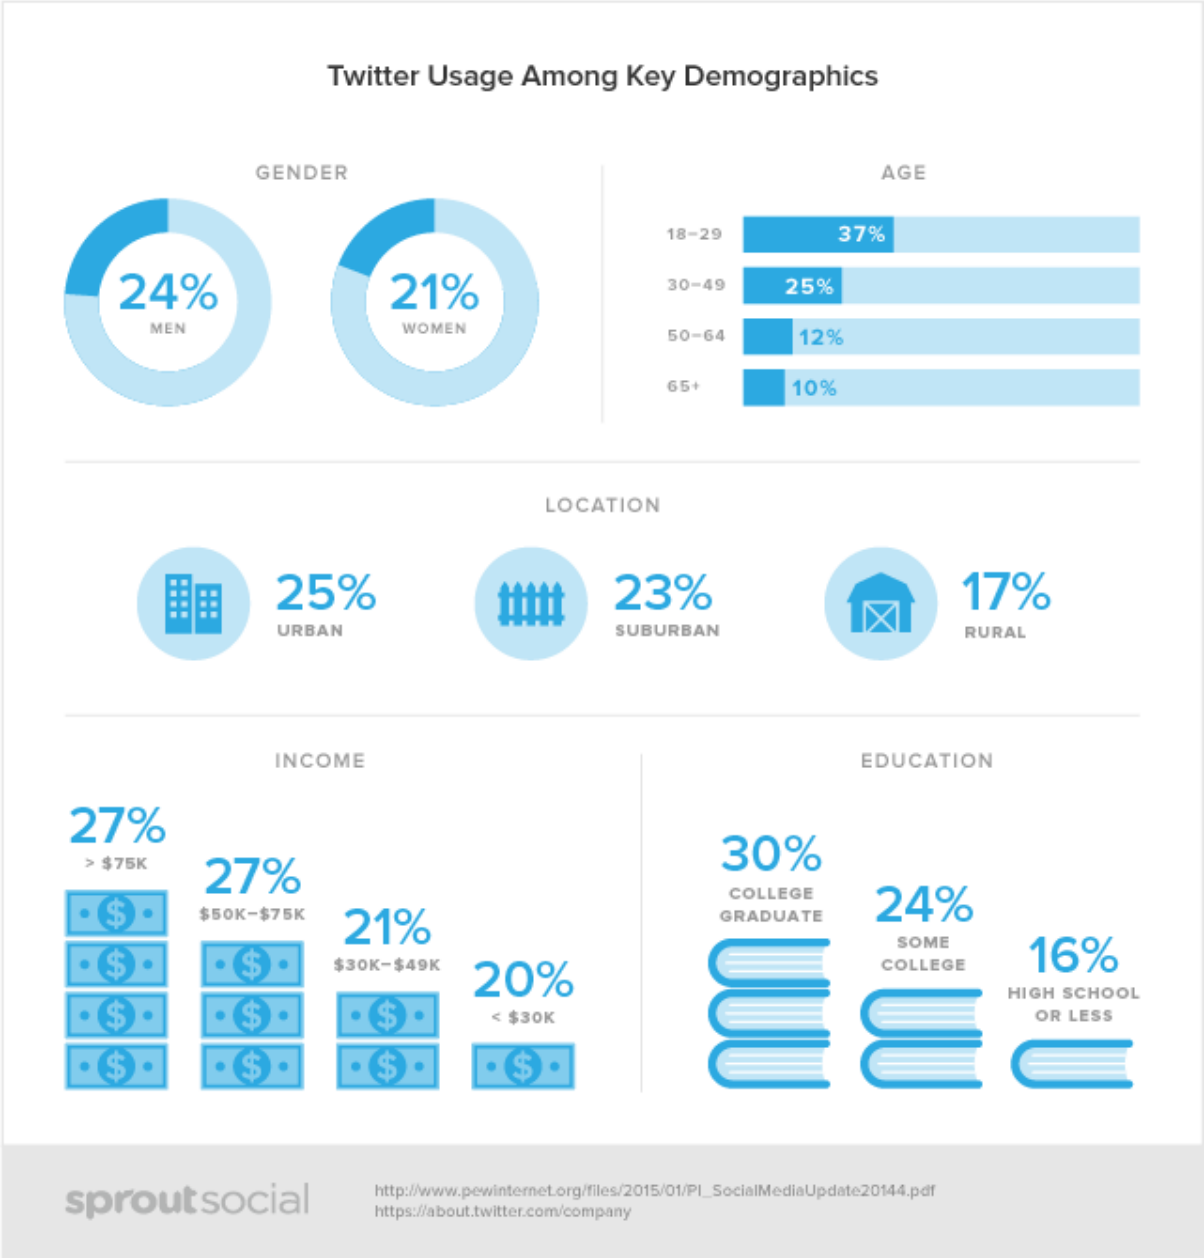

Figure 46 – Twitter usage among key demographics (Sprout Social, 2015)

The list of cultural categories was kept to a minimum in order for cluster to be visually identifiable, this means a number of potential cultural terms had to be removed. These could have included cultural topics such as libraries, sport, cuisine, parks, gardens, festivals and celebrations. Food and drink was considered as a category represented by the search terms: “club”, “bar”, “café”, “restaurant”, and “pub”,

however the search produced too many results which over dominated the legibility of other results (Figure 47 in Appendix 1).

Language was another limitation, Berlin was included as a sample of a non-native English speaking population however there are many English speaking Germans and there will also be a significant English speaking tourist population visiting.

08 Conclusion

Within this research paper a methodology to identify the cultural loci of the city by analysing geolocated social media messages has been identified, an application has been developed and the results analysed. It can be concluded that it is possible to identify cultural loci within cities, this has been demonstrated in the analysis of London and Berlin but Newcastle did not have a sufficiently large density of tweets across a 7 day window.

There is therefore potential for future research on this topic by developing an application which collects tweets across a larger period of time than just 7 days. This could be achieved by using Twitter's Streaming API which collects tweets as they are sent which can then be stored in a database and analysed at a later date. Using this method it would also be possible to collect tweets over an entire year, so the cultural loci can be better represented year round, by comparing differences in clusters month to month it would be possible to identify temporary loci such as festivals and carnivals.

To more effectively apply this avenue of research to non-English speaking countries future iterations of the application could automate the translation of search terms to the native language. Finally, this research has included only 11 search terms across 5 categories therefore, there is a significant amount of culture that has not be mapped, future research could explore any number of topics from the potentially endless list of cultural terms.

This research will have significant influence on my future academic studies in the realms of urban design, most notably it will assist with my selected investigation on the topic of graphing connections within the city. It will also aid future architectural masterplanning exercises which require an analysis of the dynamic clusters of social or cultural loci across the city.

09 Bibliography

Amaral, J. N., Buro, M., Elio, R., Hoover, J., Nikolaidis, I., Salavatipour, M., Stewart, L., & Wong, K. (2011). About Computing Science Research Methodology.

Batty, M., & Hudson-Smith, A. (2014). Visual Analytics for Urban Design.Urban Design, (132).

Batty, M., Gray, S., Hudson-Smith, A., Milton, R. W., O'Brien, O., & Roumpani, F. (2013). Visualising spatial and social media.

Batty, M., Hudson-Smith, A., Hugel, S., & Roumpani, F. (2015). Visualising Data for Smart Cities. Handbook of Research on Social, Economic, and Environmental Sustainability in the Development of Smart Cities, 339.

Carmona, M. (Ed.). (2014). Explorations in Urban Design: An Urban Design Research Primer. Ashgate Publishing, Ltd..

CASA (2012) Twitter Tongues. Available at: http://twitter.mappinglondon.co.uk/ (Accessed: 19th April 2016)

CASA (No date) Tweet-o-Meter. Available at: http://www.casa.ucl.ac.uk/tom/ (Accessed: 19th April 2016)

Cheshire, J. And Manley, E. (2013) The Twitter Map Of London, Guardian Newpaper, Tuesday 26-03-13, 28-29

CNN Travel (2012) World's most cultured cities. Available at: http://travel.cnn.com/explorations/life/worlds-best-cities-culture-vultures-832691/ (Accessed: 19th April 2016)

Gray, S., Milton, R., & Hudson-Smith, A. (2015). Advances in Crowdsourcing: Surveys, Social Media and Geospatial Analysis: Towards a Big Data Toolkit. In Advances in Crowdsourcing (pp. 163-179). Springer International Publishing.

Howe, J. (2006). The rise of crowdsourcing. Wired magazine, 14(6), 1-4.

Hudson-Smith, A. (2014). 11 smart cities, social networks and the Internet of things. Explorations in Urban Design: An Urban Design Research Primer, 123.

Hudson-Smith, A. (2014). Tracking, Tagging and Scanning the City.Architectural Design, 84(1), 40-47.

Hugel, S. & Roumpani, F. (2013) TweetCity: Building London using Real time feeds and CityEngine. Available at: http://en-topia.blogspot.co.uk/2013/07/tweetcity-re-populating-london.html

Liu, S. B., & Palen, L. (2010). The new cartographers: Crisis map mashups and the emergence of neogeographic practice. Cartography and Geographic Information Science, 37(1), 69-90.

Milton, R., Anand, S., Batty, M., Crooks, A., Hudson-Smith, A., Jackson, M., & Morley, J. (2010). Data mash-ups and the future of mapping.

ONS (2015) Population Estimates for UK, England and Wales, Scotland and Northern Ireland, Mid-2014. 25 June 2015. Available at: http://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/datasets/populationestimatesforukenglandandwalesscotlandandnorthernireland (Accessed: 22nd April 2016)

Pointer, G. (2007) "The UK's major urban areas" (PDF). statistics.gov.uk. (Accessed: 22nd April 2016)

Sprout Social (2015) Social Media Demographics for Marketers. Available at: http://sproutsocial.com/insights/new-social-media-demographics/ (Accessed: 22nd April 2016)

Twitter Developers (2016) GET search/tweets. Available at: https://dev.twitter.com/rest/reference/get/search/tweets (Accessed: 19th April 2016)

UNdata (2015) Population. Available at: https://www.google.co.uk/publicdata/explore?ds=z5567oe244g0ot_&met_y=population&idim=city_proper:015520:015780:015160&hl=en&dl=en (Accessed: 22nd April 2016)

World Cities Culture Forum (2015) World Cities Culture Report 2015. Available at: http://www.worldcitiesscultureforum.com/publications/world-cities-culture-report-2015 (Accessed: 19th April 2016)

World Cities Culture Forum (2016) The World Cities Culture Report 2015. Available at: http://www.worldcitiesculturereport.com/ (Accessed: 19th April 2016)

Appendix 1 – Additional Results

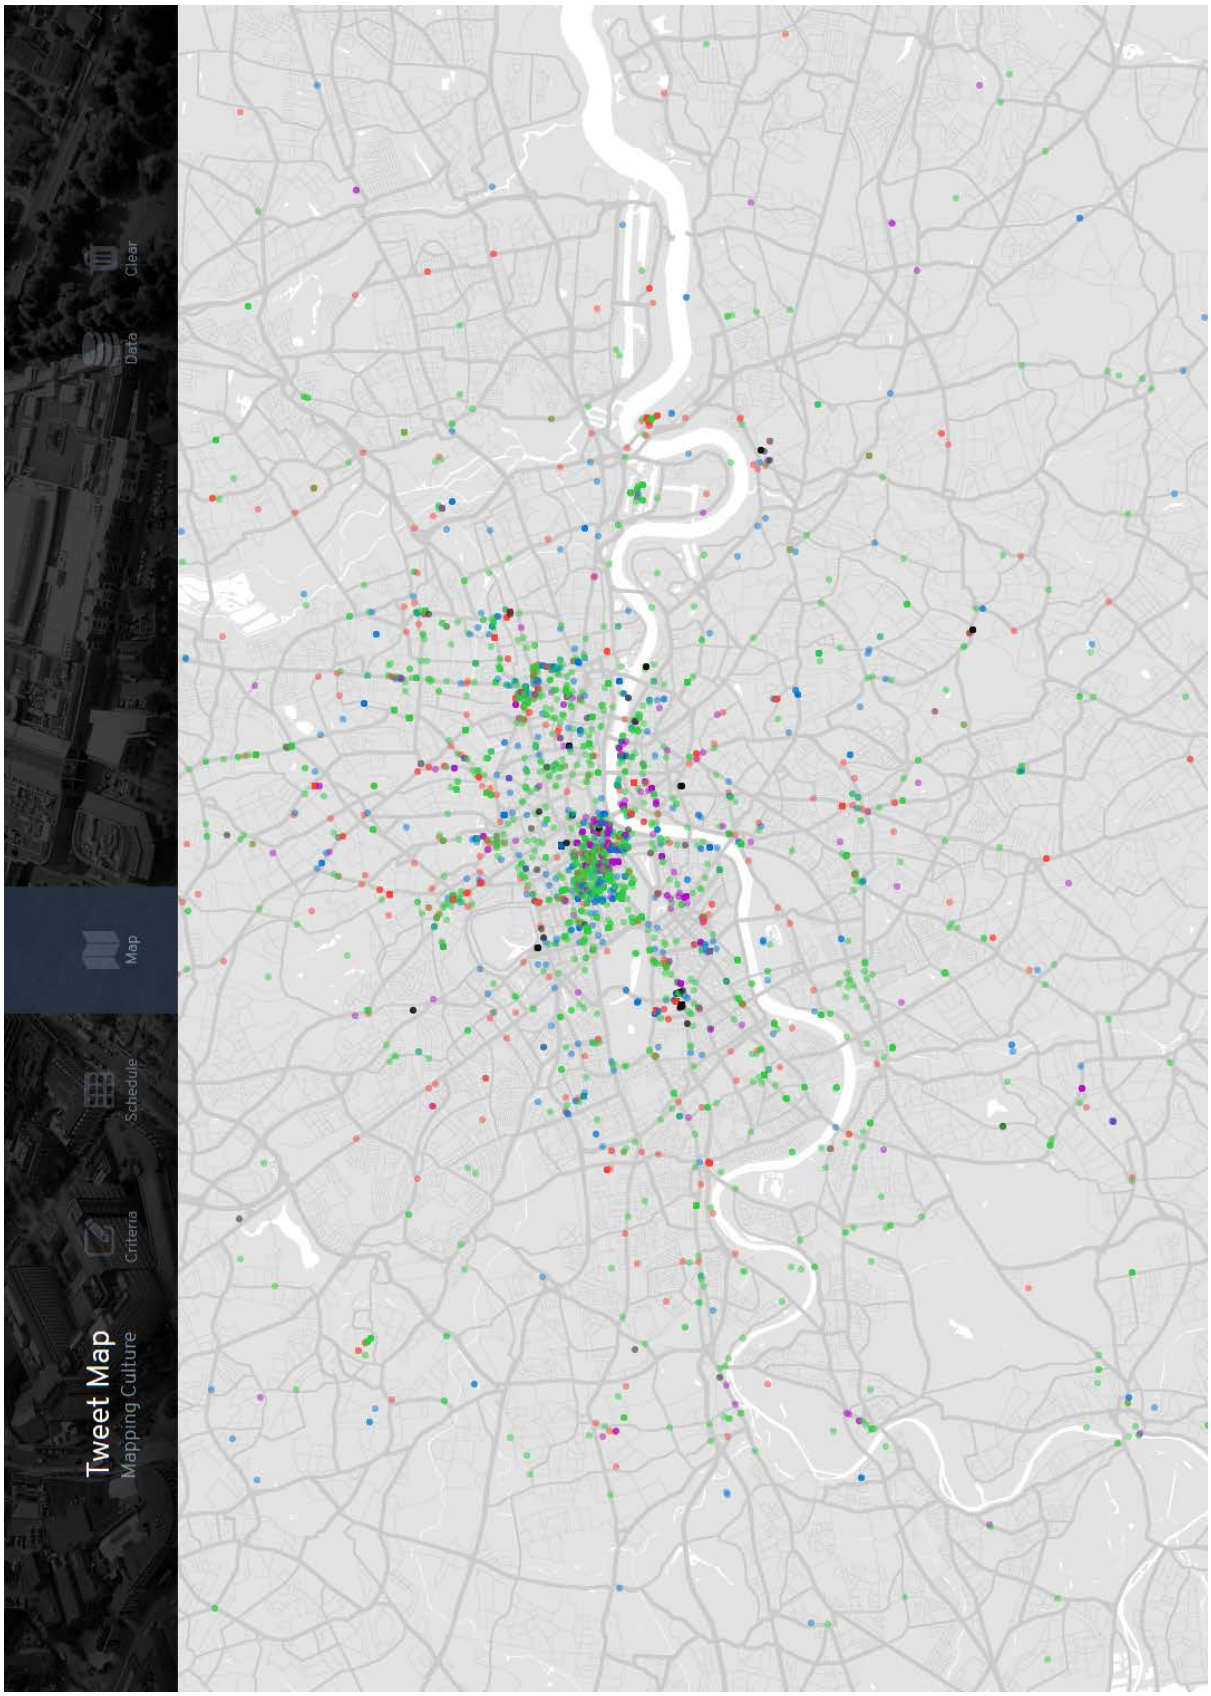

Figure 47 – London tweet culture map with restaurants in green at zoom level 12

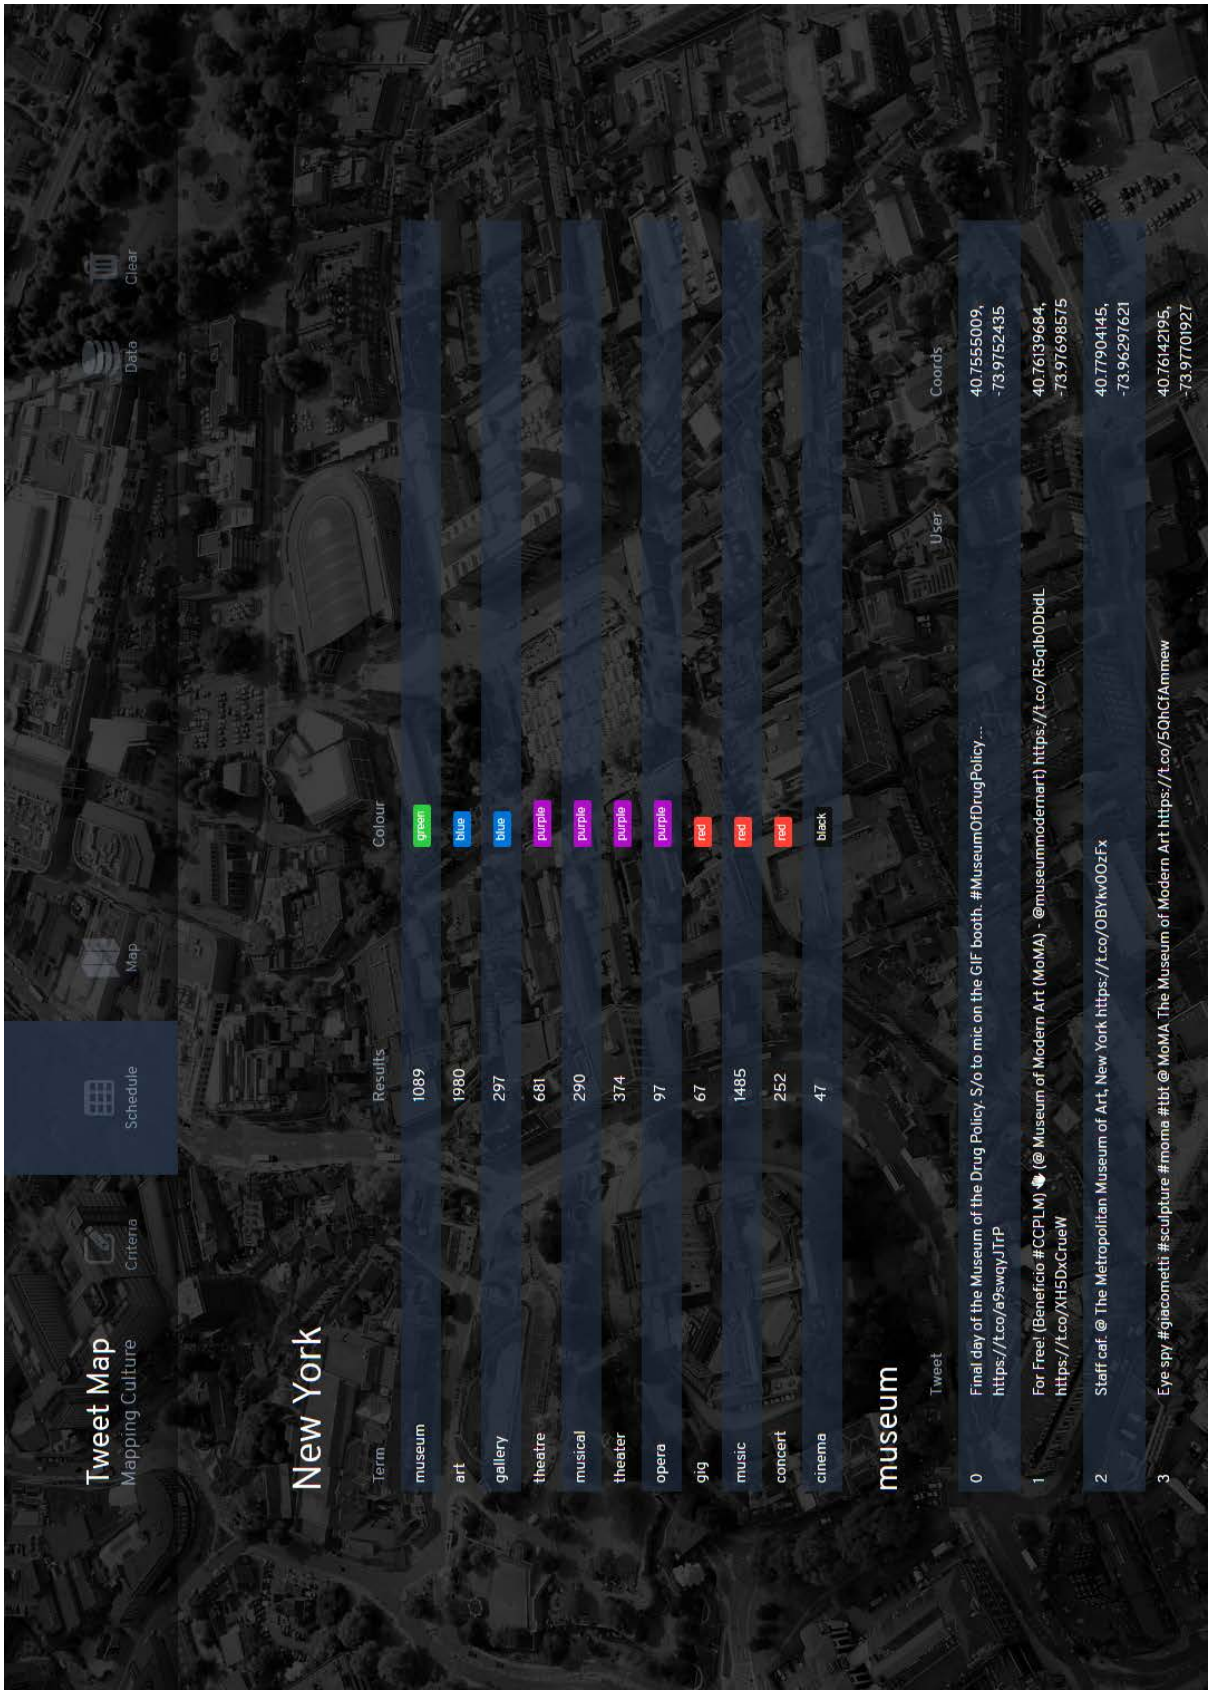

Figure 48 – New York results table

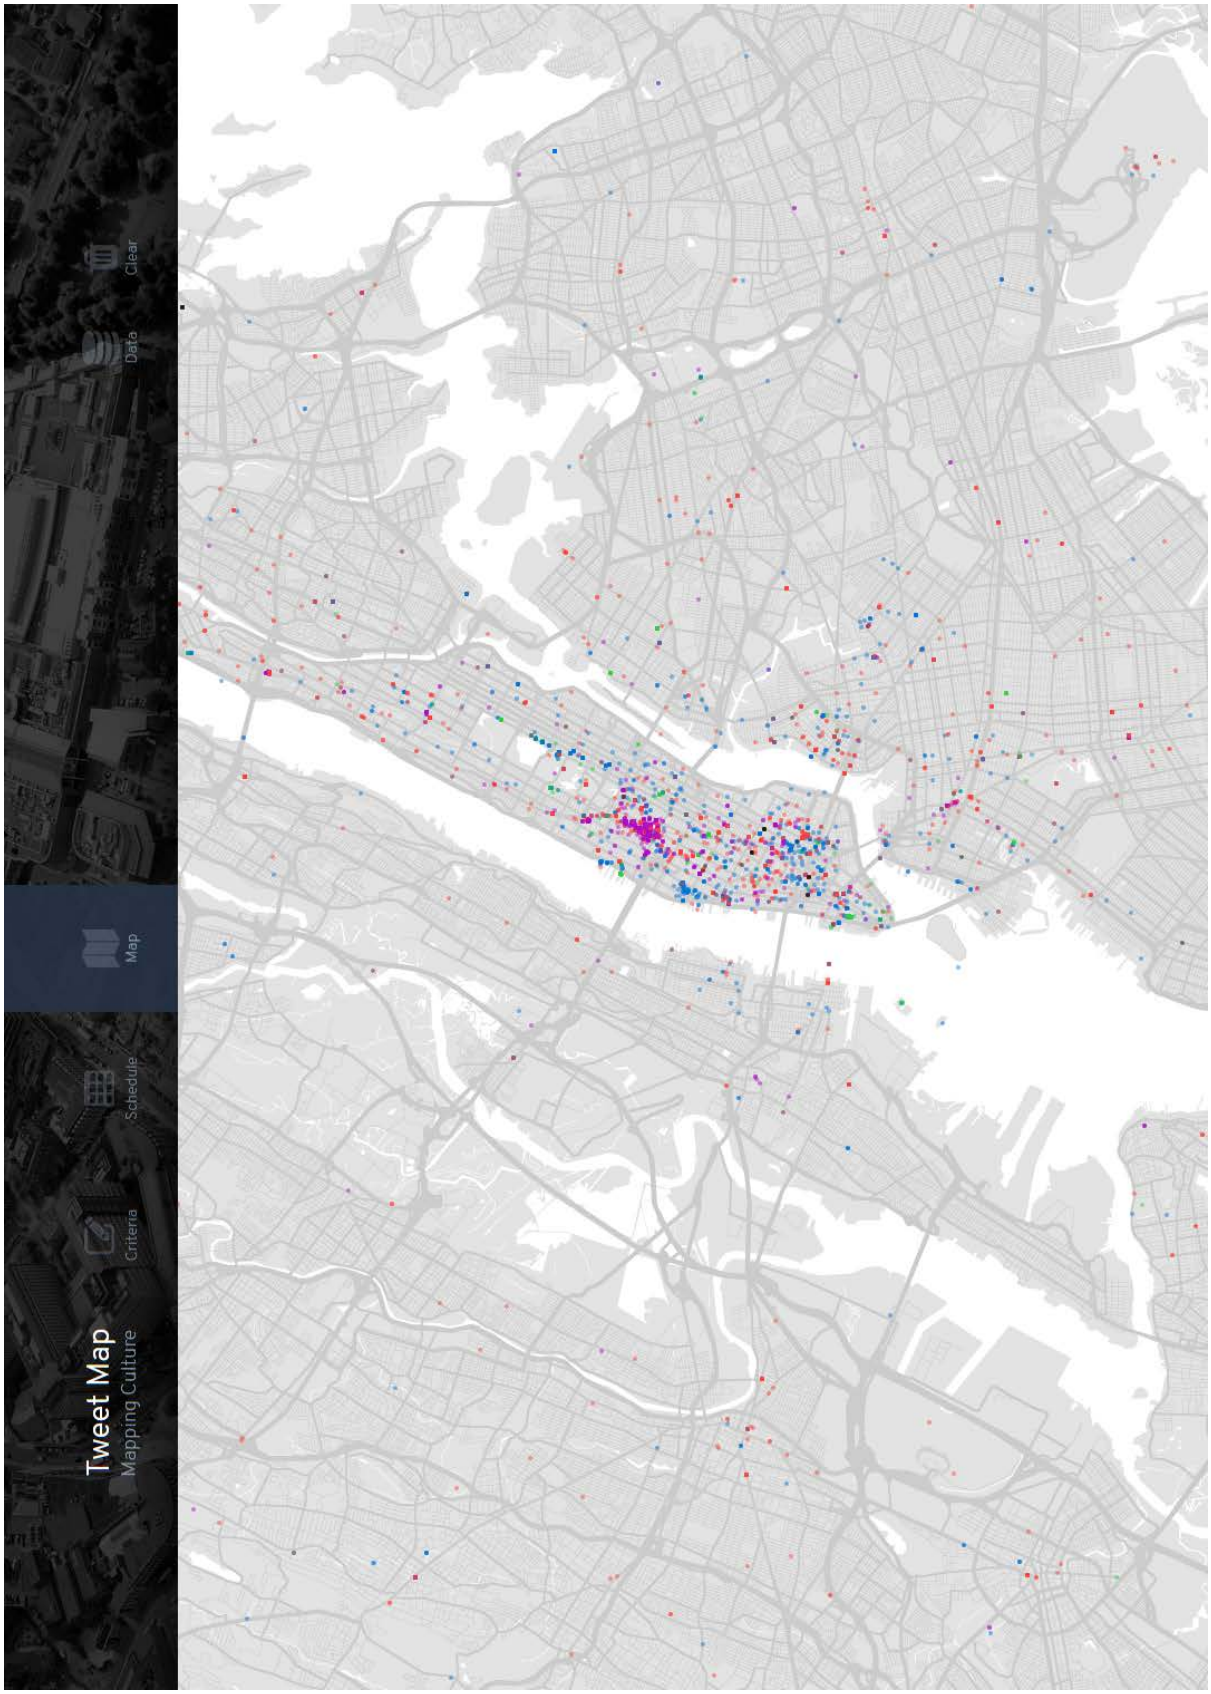

Figure 49 – New York tweet culture map at zoom level 12

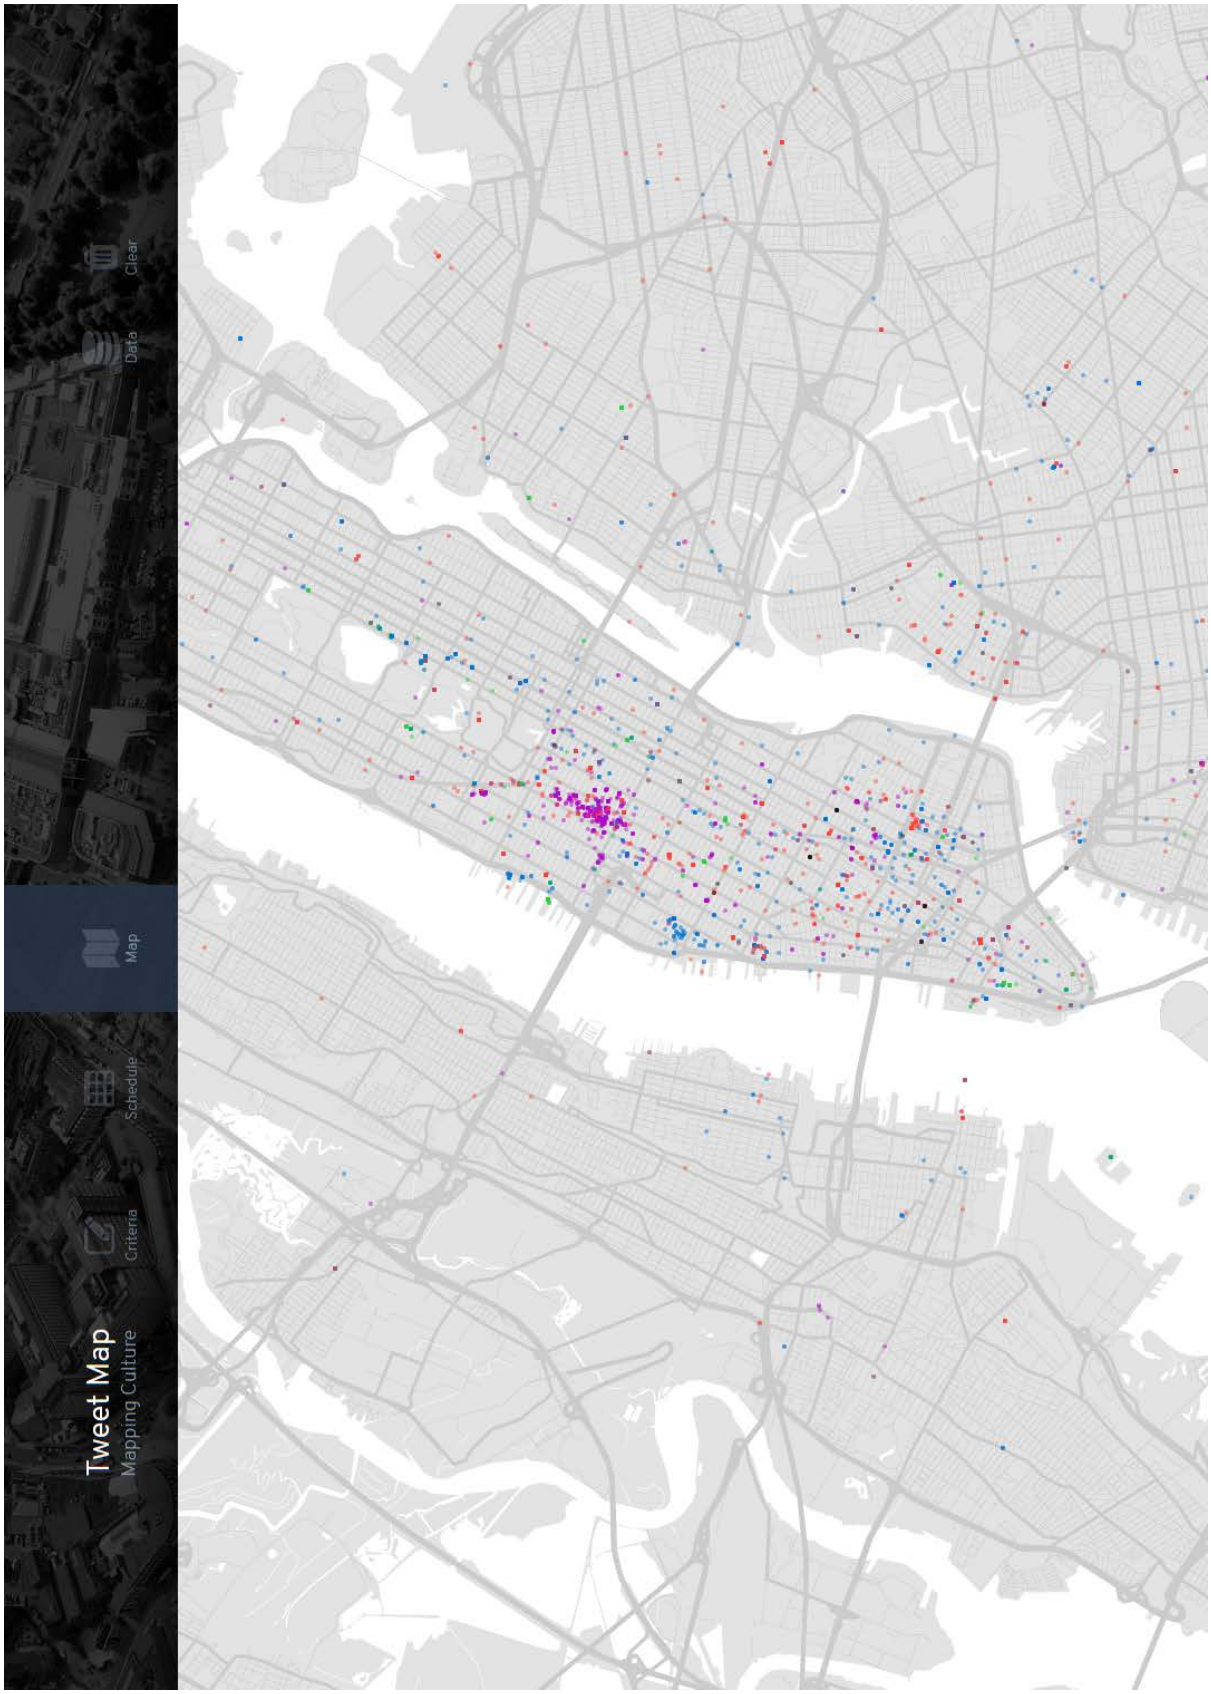

Figure 50 – New York tweet culture map at zoom level 13

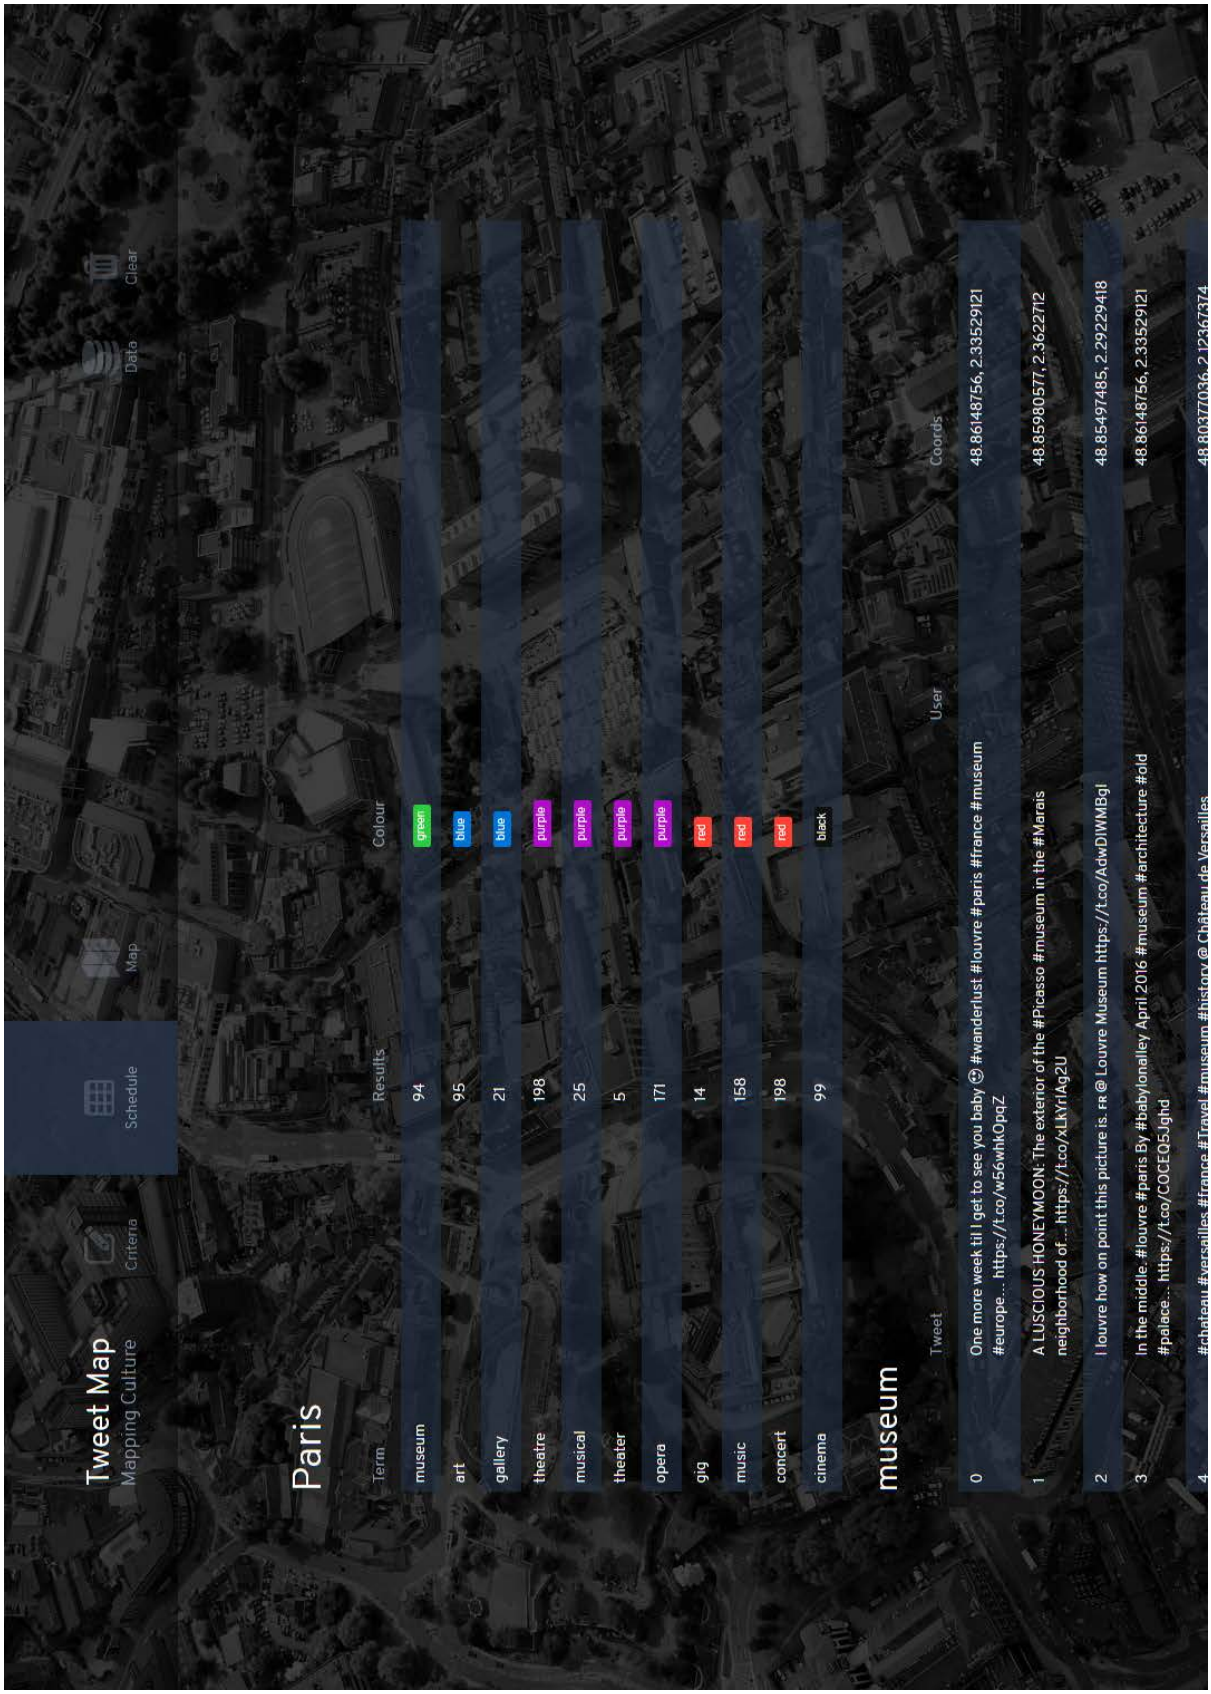

Figure 51 – Paris results table

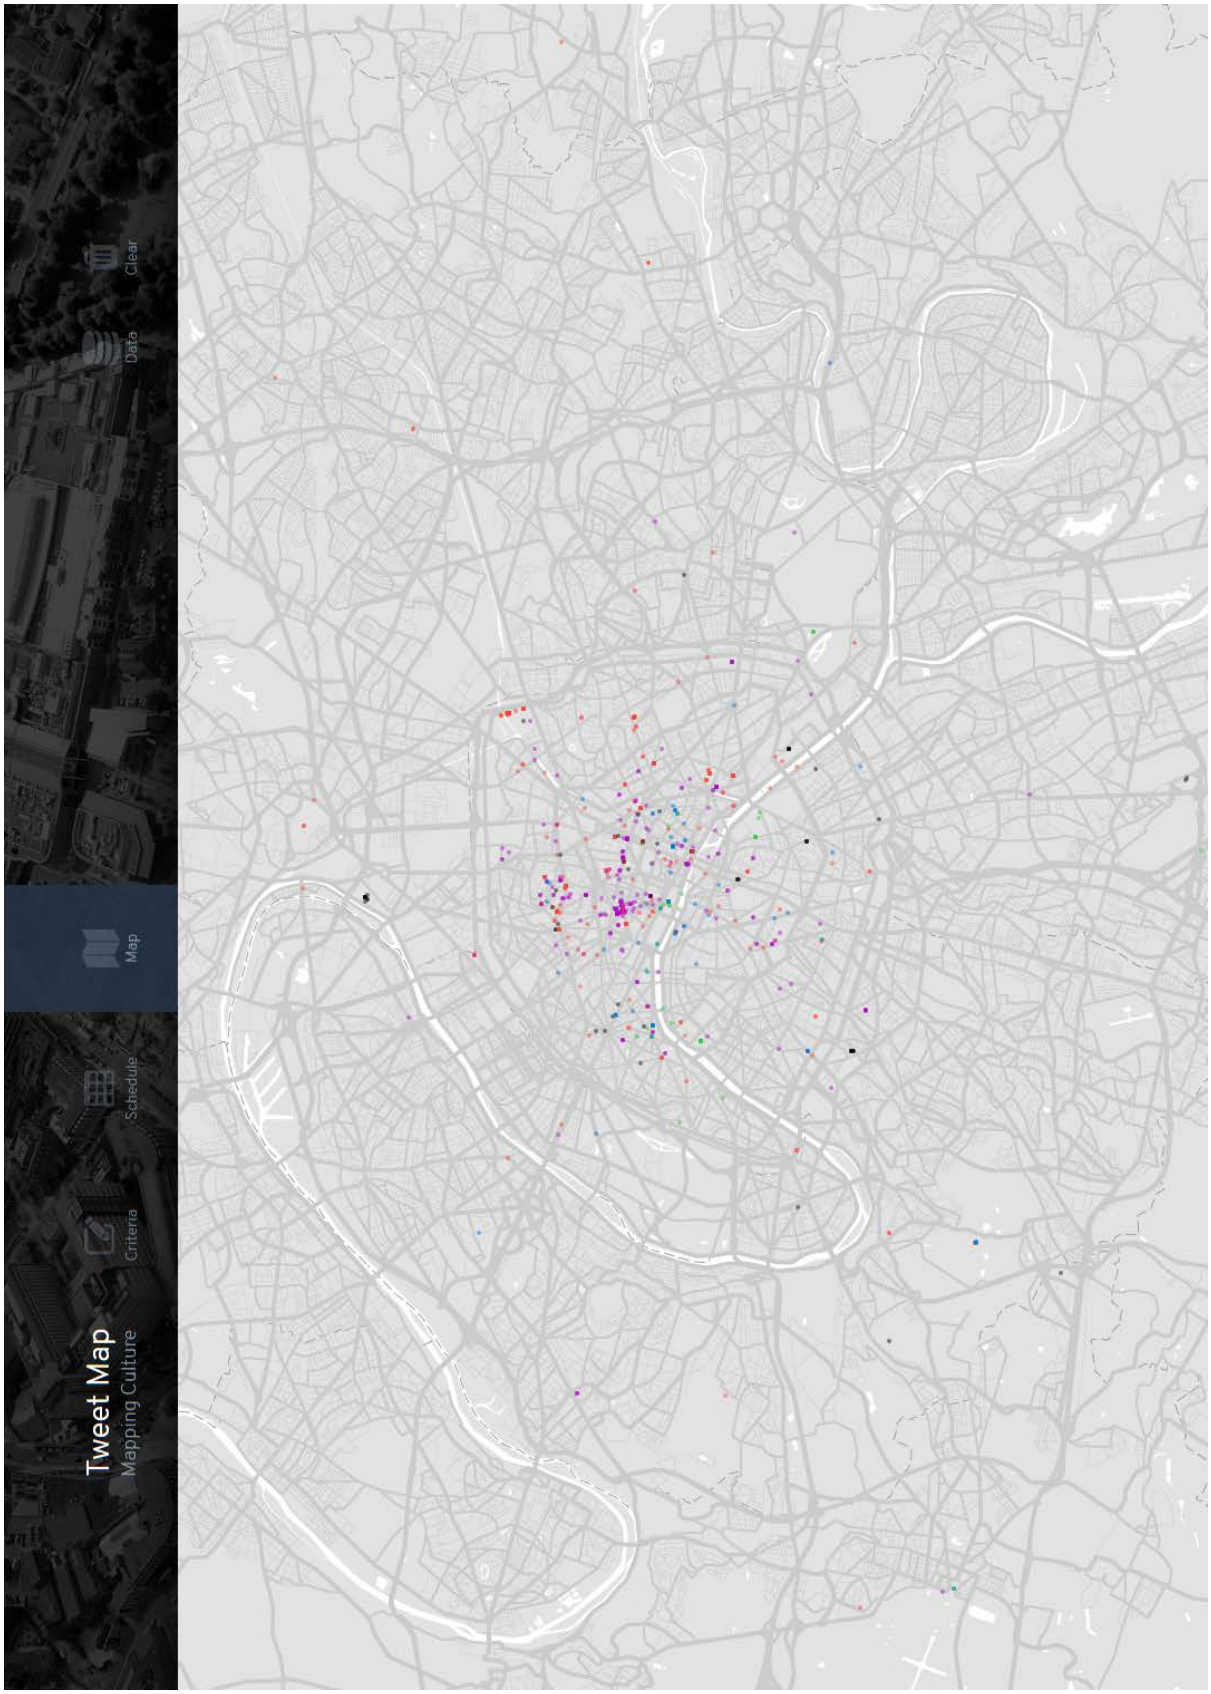

Figure 52 – Paris tweet culture map at zoom level 12

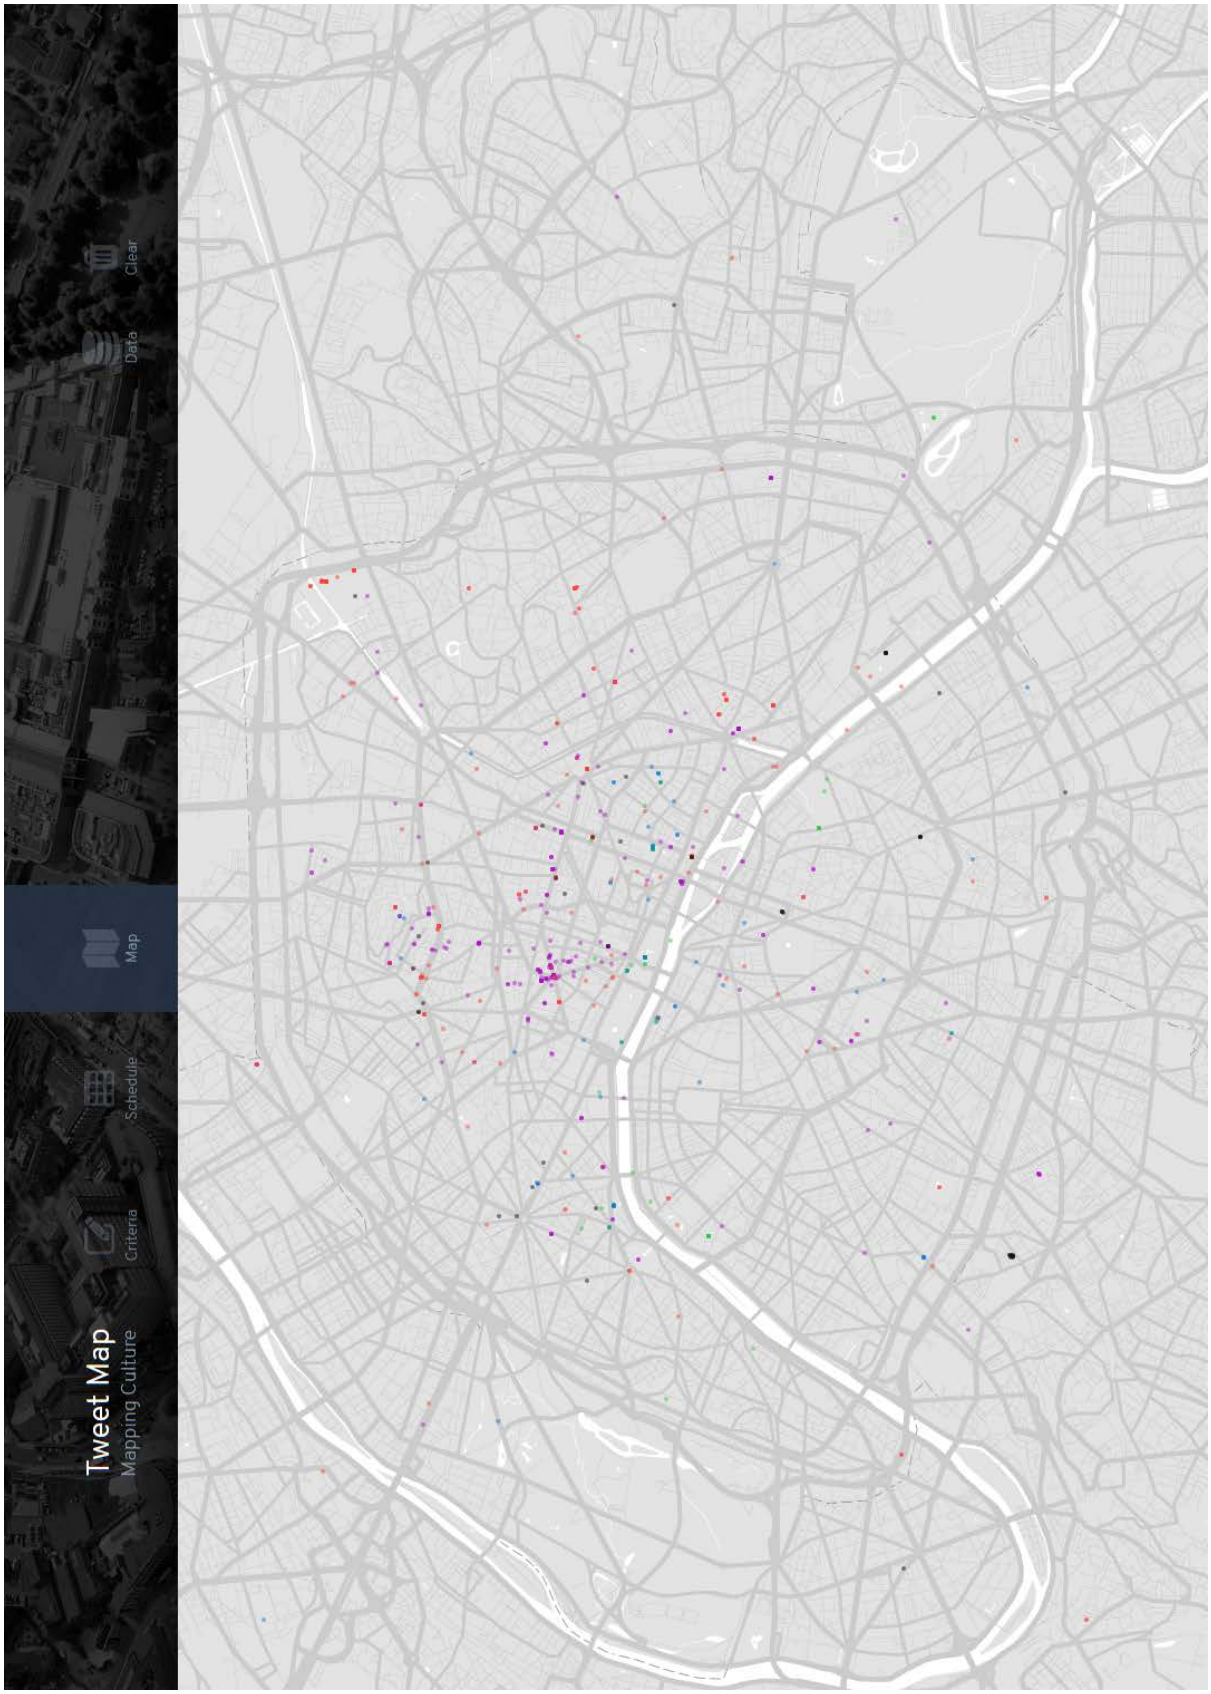

Figure 53 – Paris tweet culture map at zoom level 13

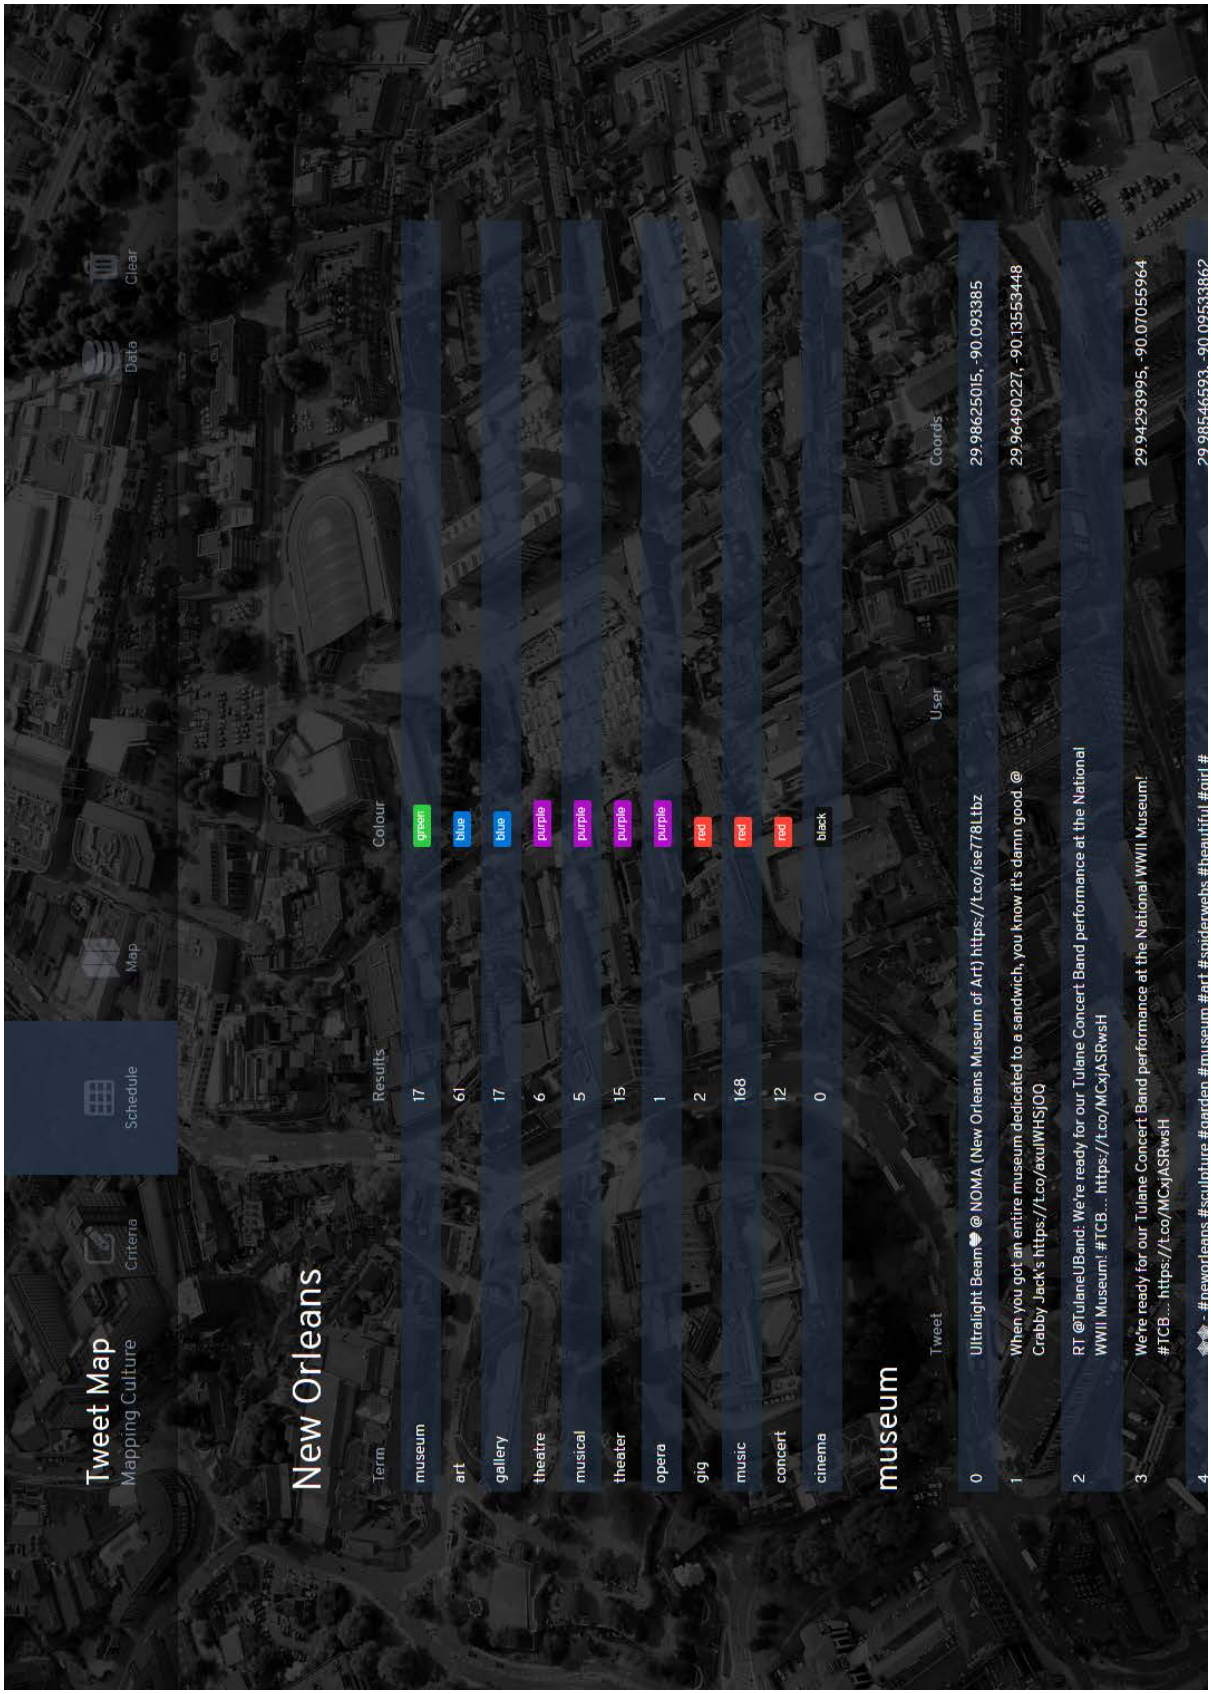

Figure 54 – New Orleans results table

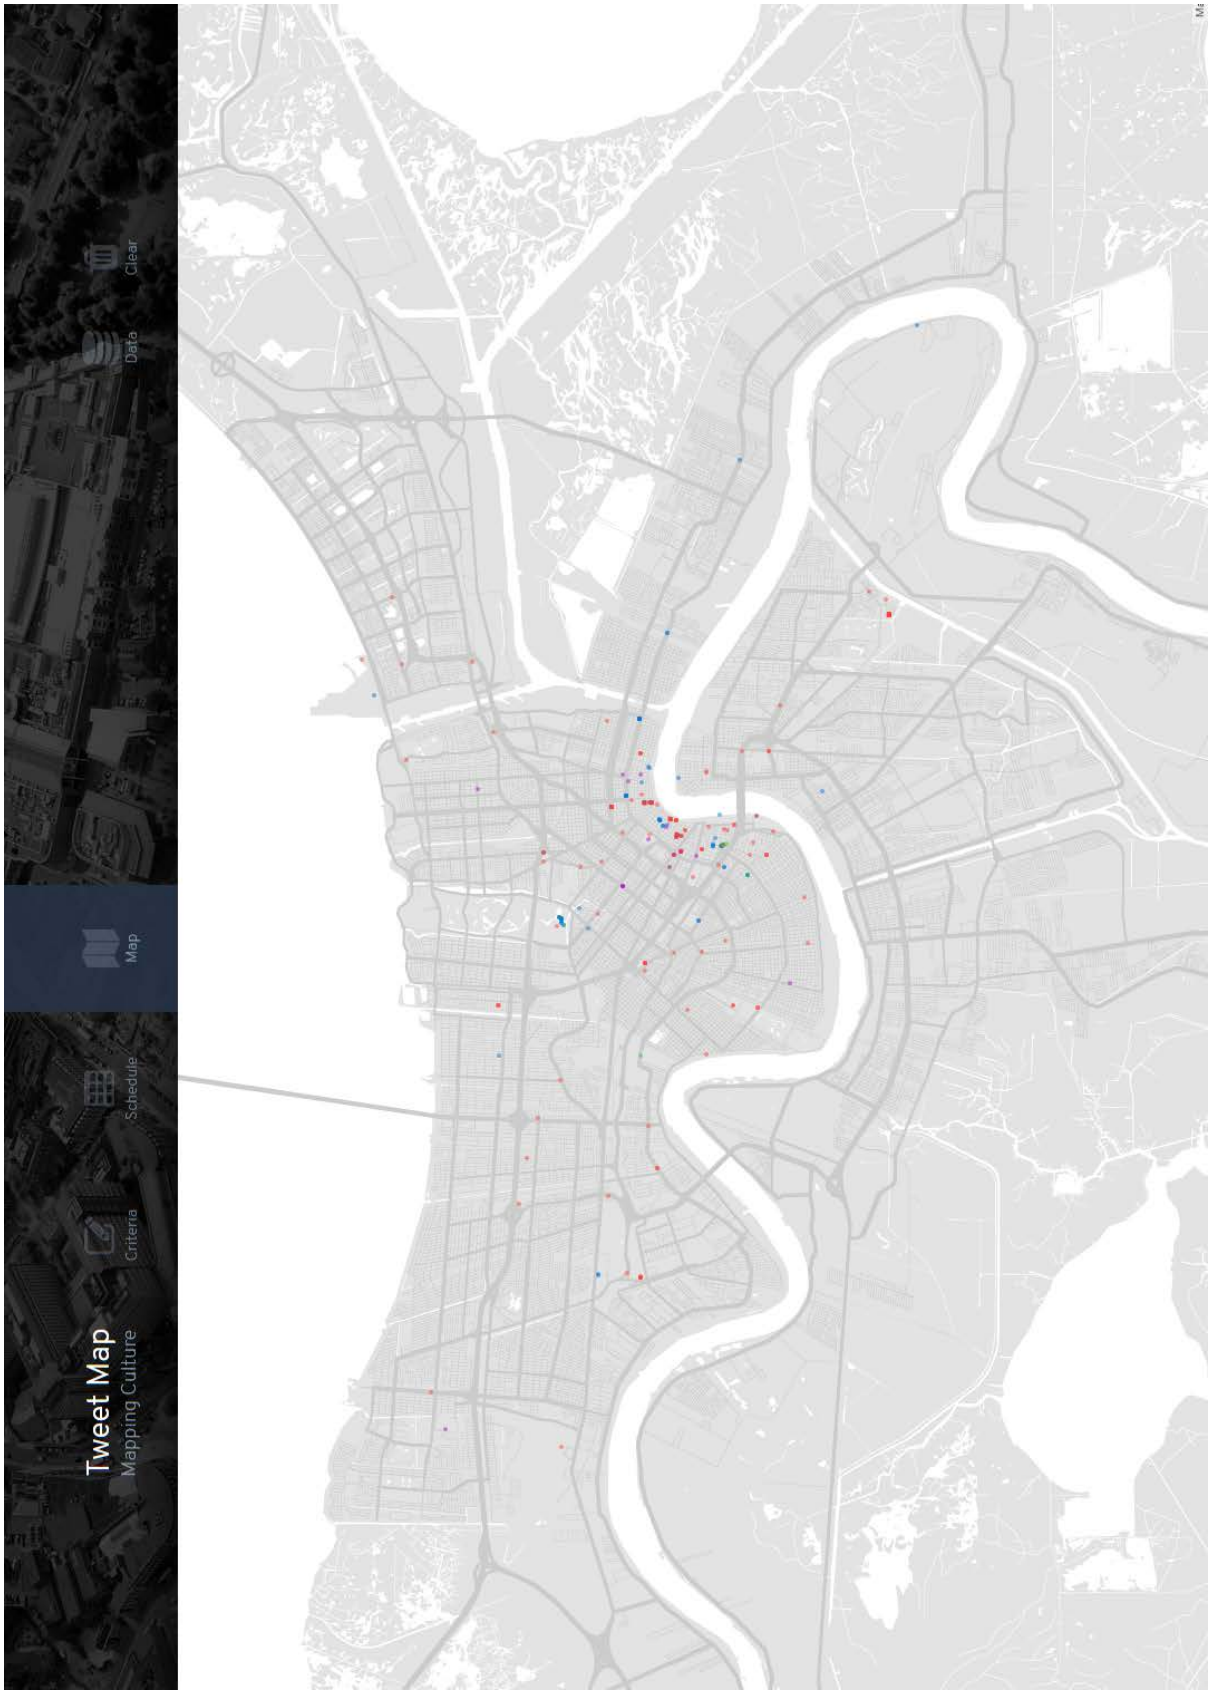

Figure 55 – New Orleans tweet culture map at zoom level 12

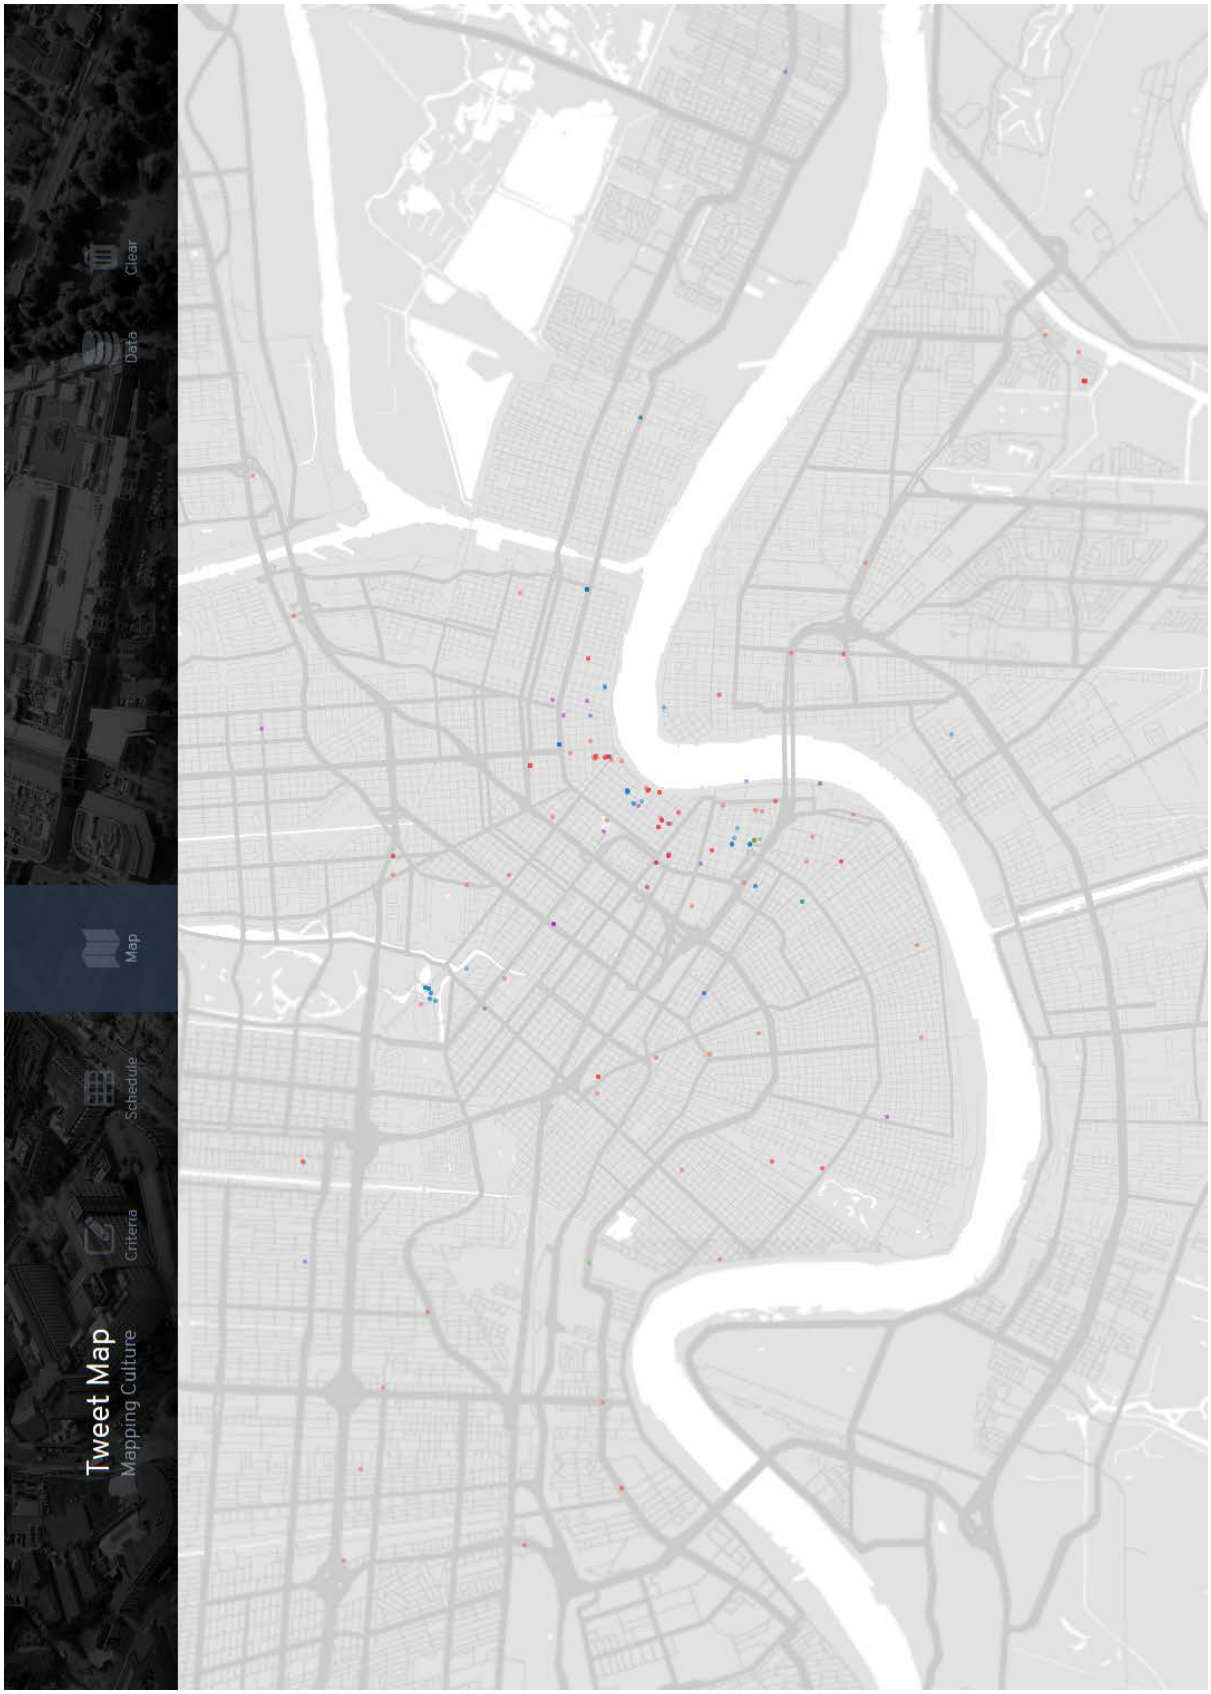

Figure 56 – New Orleans tweet culture map at zoom level 13