



News and the city: understanding online press consumption patterns through mobile data

Salvatore Vilella^{1,4*} , Daniela Paolotti⁴, Giancarlo Ruffo¹ and Leo Ferres^{2,3,4}

*Correspondence:

salvatore.vilella@unito.it

¹ Department of Computer Science,
University of Turin, Turin, Italy

⁴ ISI Foundation, Turin, Italy

Full list of author information is
available at the end of the article

Abstract

The always increasing mobile connectivity affects every aspect of our daily lives, including how and when we keep ourselves informed and consult news media. By studying a DPI (*deep packet inspection*) dataset, provided by one of the major Chilean telecommunication companies, we investigate how different cohorts of the population of Santiago De Chile consume news media content through their smartphones. We find that some socio-demographic attributes are highly associated to specific news media consumption patterns. In particular, education and age play a significant role in shaping the consumers behaviour even in the digital context, in agreement with a large body of literature on off-line media distribution channels.

Keywords: News consumption; Mobile data; Deep packet inspection; Urban; Geo-referenced analysis

1 Introduction

Internet, the World Wide Web and, more recently, the pervasiveness of mobile technologies have radically transformed the way individuals consume cultural content. One of the areas that was impacted the most by new forms of digital media is journalism: how newspapers are accessed, how news are consumed. The abundance and diversity of topics have forever been altered [1]. Gaining a better understanding of how different population groups are using and benefiting from on-line news services is now even more important. By leveraging mobile information, it is possible to map the variability of news consumption patterns of the population onto socio-demographic features such as the education level, age, income etc.

In this work, we want to explore the association between social inequality and news consumption in the context of digital news outlets. We focus on Chile, a country that in the past three decades has experienced an incredibly rapid economic growth along with a deep political transformation (from a dictatorship to a democracy). Exactly for these reasons, the case of Chile is particularly interesting. It is currently among the countries with larger social disparity and yet, according to the *Newzoo's Global Mobile Market Report*, the penetration of digital technologies remains the highest in Latin America. The

© The Author(s) 2020. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

expected effect is that easier access to digital news outlet should “democratize” the news consumption. Indeed, it has been found that newspapers and magazines have lost their elite character, diversifying in order to adapt to popular preferences throughout the 20th century while books remain a marker of ‘elite culture’, and while status is not an important determinant of magazine and newspaper reading, education and income are [2].

News consumption patterns have been increasingly studied in the last decade, especially now that digital media have started providing a huge variety of new platforms that facilitate—and, on the other hand, makes more complex—the fruition of news. Access to news outlets from mobile phones, for example, is now more fragmented and shallower with respect to traditional media [3]. Technology also drives the consumption of content in terms of continuity of access (“anywhere, anytime”) [4]. Previous literature [5] has also shown how individual choices with respect to news consumption can be influenced by the social status of the users. A strong and systematic association between status and newspaper readership has previously been established [6]. Influence of education, income and age on newspaper consumption and platform preferences has also been explored [7].

The role of education is particularly critical. Researchers have found education to be a strong predictor of media behavior [8]. It is also positively associated with general news exposure [9]. In some context, news consumption is more unequally distributed than income, with greater social inequality in online news consumption than in off line news consumption [10]. Within the particular context of Chile, news media have already been studied to outline, for example, their ownership structures or their political bias and to understand the effects of certain press manipulations by owners of the media outlet [11].

The contribution of this work is two-fold: first, we explore the association between socio-demographic attributes and news consumption in the Chilean context, in line with the literature we just described (that usually based the analyses in single-nation frameworks, like the US or Argentina) and our own previous work [11–14]. Second, we rely upon a unconventional and unique data source: mobile phone data. Traditionally, this topic has been investigated using surveys among various population groups ([15–18], just to mention a few). We believe that mobile technology can provide new insights that complement traditional methods of data collection. Indeed, many diverse research questions have been addressed using mobile communication data: to elaborate indices of the economic development of a region [19], to study traffic flows, urban human mobility, social mixing phenomena [20–25] as well as epidemic spreading on fine spatial and temporal scales [26]. We believe that such a data source could add interesting new insights also in the context of this research area, besides the fact that these methods are less expensive, less time consuming, and they scale efficiently.

In the following sections we present the results of a study on users accessing on-line news media in the city of Santiago de Chile (SCL) through mobile devices (using DPI (*deep packet inspection*) data. The focus is on the time window that spans from the 6th of July 2016 to the 2nd of August 2016. In order to characterize each area by the socio-demographic information of its residents, we used the 2017 census data (see Sect. 2.1.1). We find that some socio-demographic attributes are highly associated to specific news media consumption patterns. In particular, education and age play a significant role in shaping the consumers behaviour even in the digital context, in agreement with a large body of literature on off-line media distribution channels.

2 Methods

2.1 Datasets and data pre-processing

In this section, we present a description of the data used and of the procedures that we followed for data cleaning and pre-processing.

2.1.1 Census data

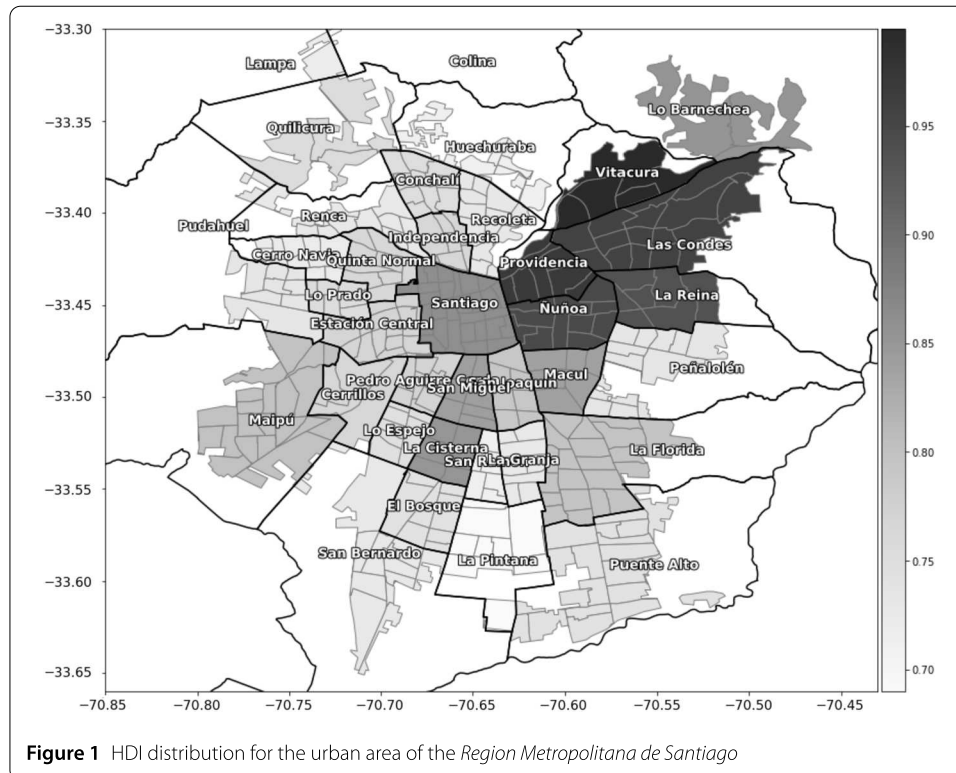
To characterize Chile socio-economically, we used the official 2017 Chilean census,^a surveying 17,574,003 people (51.1% females). Geo-politically, from largest to smallest, Chile consists of “regions”, the largest administrative division, followed by “comunas” (similar to counties in the United States), census districts, census zones, and finally blocks, the finest level of geographical granularity. The public dataset is available at the level of blocks. Here, we focus on the Metropolitan Region of Santiago de Chile (RM, for short). The RM consists of 52 “comunas” for a total of 7,112,808 individuals (51.3% females), accounting for about 40% of the total population of Chile. In this work we correlate census information with data from mobile web traffic records. Since minors (persons whose age is <18) cannot sign mobile phone subscriptions, we decided to be conservative and only work with adults in the census to have a better match between the sample populations of the mobile data and census data sets, ending up with 5,450,592 individuals in the census and 2,455,148 accesses by unique users in the general phone dataset.

The 2017 Census contains a great deal of information on several aspects of the socio-economic composition of the Chilean population. It has the issue that it is only an abridged version of the usual census surveys, since a more comprehensive one will be released in 2022. Given the particular structure of the questionnaire, we were able to perform a manual selection of the census features. Indeed, the vast majority of the questions in the survey revolve around a few topics (e.g., on top of the basic question “*Do you consider yourself as a member of a native minority?*” there are other additional questions to distinguish between the different native minorities). We thus decided to select our features only among the following basic socio-demographic information:

- Age;
- *Escolaridad* (years of formal education attained);
- Student status (whether an individual is still studying or not);
- Membership to a native population.

The census does not report on individuals’ income directly. Thus, information about the economic situation had to be inferred. To do so, we chose *escolaridad* as a proxy. In Chile, at least, there is a well-known strong correlation between formal education and income distribution and inequality [27, 28]. Moreover, the variable *escolaridad* weighs heavily in the calculation of the Human Development Index (HDI), which is a widely used indicator of wellness and quality of life, that in turn can be used to get a first, qualitative measure of the soundness of our results (a map of Santiago HDI distribution at the level of municipalities can be seen in Fig. 1).

Census data can be aggregated at the different levels already mentioned: region, comuna, districts, zones and blocks. As we go into more granularity, census information becomes less specific in order to avoid identification and preserve privacy of individuals. We have chosen the level of *district* as a good trade-off between granularity, privacy preservation and availability of information, see Fig. 2. Finally, each value was expressed as a percentage—or a mean, depending on the typology of the quantity—over the population of each district.



Other than the provisos made above, the rest of the census dataset was not modified or preprocessed in any other way.

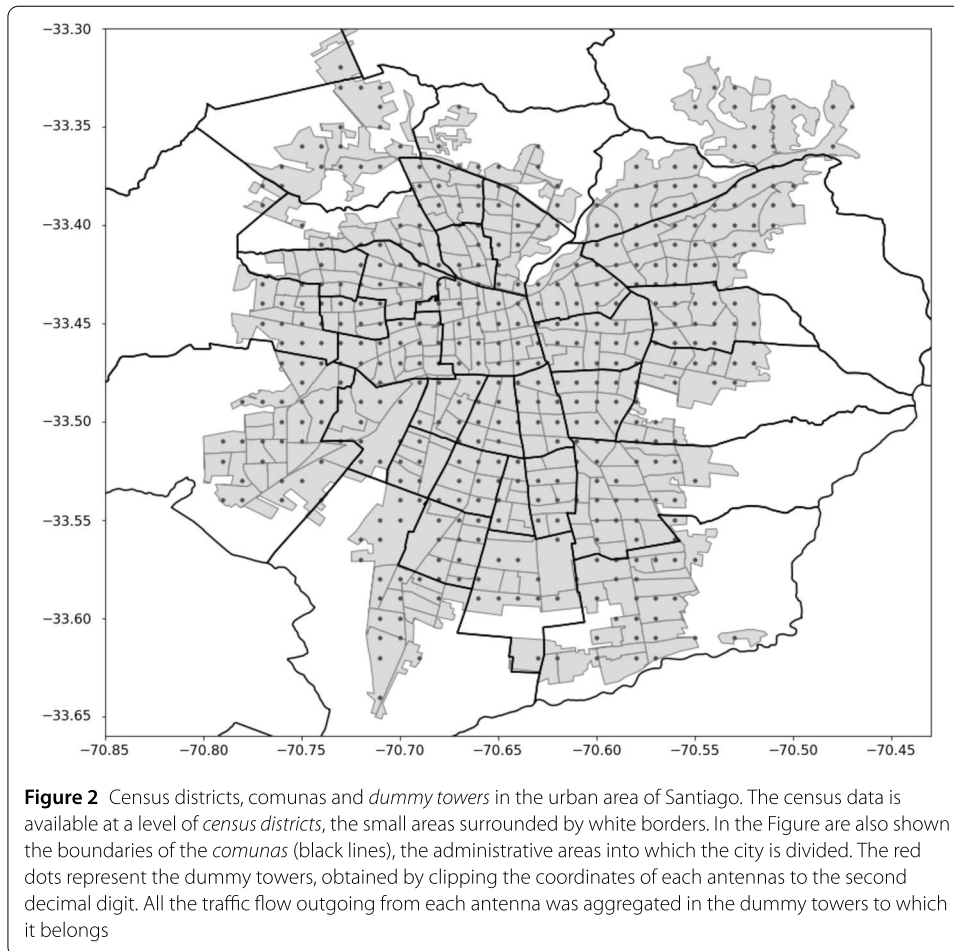
2.1.2 Building a mobile web connections data set

We obtained a DPI (Deep Packet Inspection) data set from one of the largest company in terms of mobile subscriptions market share (around 30%) in Chile. This data set is a record of Internet connections from smartphones or any device that contains a SIM card and is subscribed to the telco network. This dataset consists of almost a month (between the 6th of July 2016 and the 2nd of August 2016) of anonymized *events*. An event is defined as a connection of an individual device to an IP address through a cell tower (or antenna). In order to preserve privacy, information was aggregated by antenna and by hour, without any single user information, making it virtually impossible to de-anonymize news consumers.

Our DPI data set includes the *number of unique users* that connected from an antenna to an IP address at a certain hour (00, 01, 02, ..., 23), as in the following:

In the case of the small sample above, the first row of the raw dataset tells us that from antenna 00000000, on July 6, 2016, there were two unique (i.e. distinct) users visiting 200.12.26.117 at 11 in the morning.

The urban area of Santiago de Chile contains about 15,000 cell phone signal receivers, called *antennas*. Most outdoor antennas are placed on top of poles called cell phone *towers*. There are about 1700 towers in our area of interest, with an average of 10 antennas per tower. There exist also standalone indoor antennas (sometimes called “small cells”), usually attached to indoor walls of large spaces such as hospitals, malls and public buildings. We have information on the exact latitude and longitude of the different *towers* and indoor antennas, allowing us to geo-reference our analyses. Abusing terminology, we call



every unique geo-located cell phone receiver an “antenna”, this may be a tower with dozens of antennas on it, or just one single indoor antenna in a hospital. In high demand areas, such as the financial districts or industrial complexes, there are tens of towers with tens of antennas (plus indoor antennas) within a small area of a few tens of meters. Which antennas a user connects to depends on many factors, including distance, traffic load (many devices may be connected to that antenna already), power, azimuth, among others. If the user moves a few meters, this could mean triggering the network to assign her a new antenna every few minutes. For that reason, to make the results more stable, in this paper, we group together all the antennas within a 1.1 km radius, obtaining a grid of about 700 points. We call these “dummy” antennas (the red dots in Fig. 2). “Dummy” towers are the result of clipping the full latitude and longitude to two decimals, so that instead of having the “real” antennas at high-precision (and much finer level of granularity), we now have a “dummy” antenna at the center of the 1.1 km² square cell. These antennas should be taken as our “sensors”, the finest level of granularity we have access to, which will also be mapped to the census, as we will see later.

Our dataset is a part of the complete *deep packet inspection* dataset of the telco provider. This subset contains only requests to those IP addresses that belong to some kind of news media outlet. In order to navigate among these outlets, we used a curated list of news or-

Table 1 The first lines of the DPI dataset

	Antenna	Date	Hour	IP	Users
1	00000000	20160706	11	200.12.26.117	2
2	00000000	20160706	14	190.153.242.131	1
3	00000000	20160706	14	200.12.20.11	1
4	00000000	20160706	15	190.110.123.219	1
...

Table 2 List of news outlets (or editorial groups) examined in the DPI data set

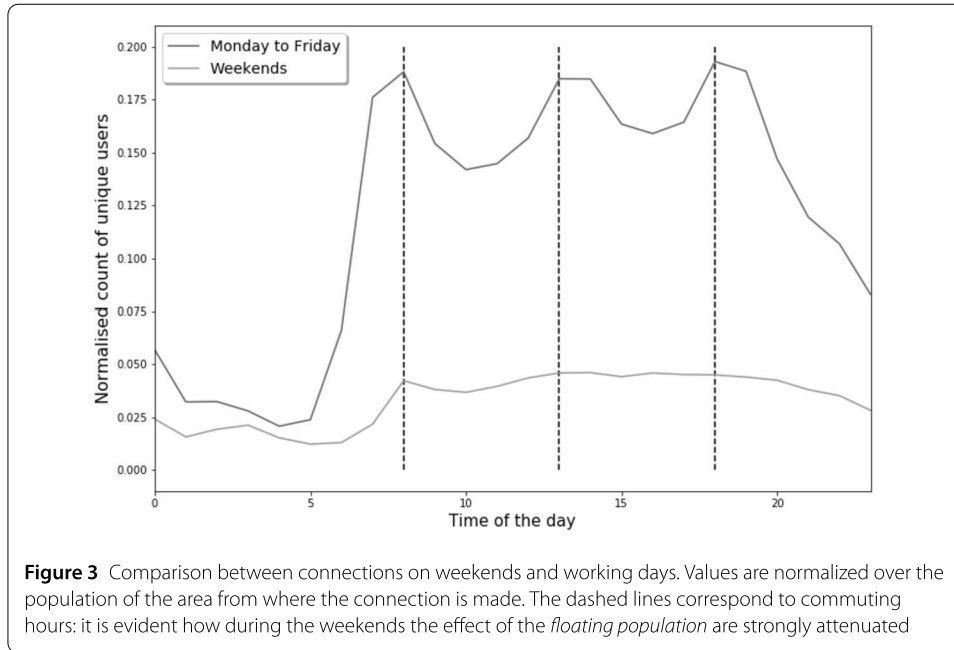
List of the news outlets examined	Number of unique users
BioBioChile	1,702,768
El Mercurio editorial group	815,152
Cooperativa	639,120
AdnRadioChile	465,149
The Clinic	337,048
Tele 13	188,409
Diario Financiero	103,450
Publmetro Chile	54,478

ganizations analyzed in a previous work [11]. These are around 400 news outlet accounts,^b twenty six outlets for which we knew their economic and political bias [12].

The process of identification of each IP address in order to associate the name of a website was not straightforward. There was often a many-to-one relationship between websites and IPs; for example, there are clusters of two or more websites that share the same IP. This is a critical issue, since in the dataset we can only see IP addresses, without any DNS reverse lookup. In any case, most of the websites that share the same IP belong to the same owner. This allows us to label each IP by its owner: this way we lose knowledge about the individual news outlet, that we are often unable to identify, but we still keep a satisfying amount of information by creating a unique matching between each IP and the editorial group (EG) that owns it. Thus, whenever needed, the EG was considered instead of the individual news media. Also, as shown in [11], the power structure in Chilean news media is strongly biased towards very few groups that share an identifiable editorial line. The only newspaper outside the above list is *The Clinic*, a Chilean satirical newspaper that is usually identified as *leftist*, which we added in order to cover the political spectrum as widely as possible, and use it as a baseline for informational extremes (“El Mercurio” would be right-conservative, and “The Clinic” would be left-liberal). The complete list of news outlets examined is shown in Table 2.

This selection is useful for various reasons: it allows us to remove some noise caused by the outliers, i.e. not-so-popular news media that display a very low number of connections, and, most importantly, it helps us focusing on a narrow but significant range of news outlets, for which we already have a quantitative knowledge of their political bias and of their popularity. Incidentally, with this filtering we were able to maintain in our analysis BioBioChile, El Mercurio and Cooperativa, that are the top 3 most visited outlets in our dataset, accounting for the vast majority of the volume of web traffic., thus not losing too much information.

After filtering by news outlet and considering only the urban area of Santiago de Chile, we have 4,313,964 events.



Finally, since we are crossing mobile data with census, which contains socio-demographic information about the *residents* of a certain area, we need to take into account the phenomenon of the *floating population*, i.e. people moving from one place to another, especially at commuting hours and during working days. This matter has been extensively studied, and there are study cases set up in Santiago that show exactly how the city is affected by the phenomenon (and how mobile data could be used to understand urban mobility, see for example [29]). To tackle this issue, we examined the temporal patterns of the connections, illustrated in Fig. 3. By comparing the trends between the weekends (Saturdays and Sundays) and the working days (Mondays to Fridays) we can easily notice the circadian rhythms of the city: the peak in connections starts when people wake up, continues when they commute to work, rises again at lunch time and finally when they go back home at 6 pm. This last peak is what differentiates the working days from the weekends: on Saturdays and Sundays, no afternoon peak can be observed, with a smooth decline of the connections towards night hours instead. The typical effects of the floating population phenomenon seem much attenuated during the weekends, and the volume of connections seems to remain almost constant throughout the whole day. The correlation between the number of unique users and the number of residents in each comuna is rather high, with a Pearson coefficient $P = 0.75$ ($p\text{-value} = 1.85e-07$). This means that the geographical distribution of users accessing news media websites from mobile is a good proxy for the geographical distribution of the population. There is a clear trade-off in our choice of considering only the connections at the weekends: on the one hand we reduce our data to almost one quarter of the initial volume, i.e. $\sim 1\text{M}$ events, but on the other hand we are able to address the issue of the *floating population*, at least partially reducing the noise caused by this phenomenon.

2.2 Ethical considerations

Working with mobile data could raise privacy concerns. Of course, DPI data was handed out by the mobile provider already anonymized and grouped per tower: data appeared as

shown in Table 1, displaying only the number of unique devices connected to a certain antenna, without any knowledge about the identity or the profile of the customer. Therefore, we work with highly aggregated data and, for good measure, the resolution was then made even coarser, as described in Sect. 2.1.2. The same applies to census data: even though they are already anonymized and publicly available, the research was not carried out at the finest possible granularity, in order to have a good balance between available amount of information and privacy preservation. Our concern in relation to privacy stems from the possibility of cross-referencing raw DPI data with census information which could have unforeseen consequences and result in potential de-anonymization.

No attempt was made to infer statistics about individuals: the nature of DPI itself prevents this from happening. All the results are *general* trends of *very large* groups of residents—between hundreds of thousands and millions of people.

2.3 Mapping census data to districts

All the connections to the websites are grouped by hour and by antenna: this provides with a good approximation of the users' position in the city. Our goal is to study the connections based on the census features of the areas from where they are originated. Thus, we want to assign to each census district (CD) a label that helps us distinguishing between different segments of census. To estimate the most appropriate number of groups (and, therefore, labels), we cluster the CDs based on the census features that we selected earlier. We use a k-means algorithm, running the procedure several times for $2 \leq k \leq 30$. Since we need a clear and easily interpretable grouping, we limit our choice to $3 \leq k \leq 8$: having only 2 census groups would not be enough for such an analysis, while more than 8 would be hard to interpret. The values are standardized, and the algorithm was initialized with a random seed. For each k we compute the Gini coefficient of the distribution of the population (Fig. 4): a low Gini represents a situation of order, in which the population is almost equally distributed across all the clusters, while a high value is retrieved for an heterogeneous distribution throughout the different groups.

The final value, $k = 5$, is chosen as it is the value that maximizes the Gini coefficient and, thus, the heterogeneity of the distribution of the population. This is based on the assumption that a very ordered situation, in which we have almost the same number of people in each cluster, would be non realistic.

By setting $k = 5$ and inspecting the average values of the features for each group (see Table 3), we observe an almost hierarchical relationship among the different clusters. This is particularly true for the *escolaridad* variable, which correlates very well with the income, as mentioned above. The clusters were simply named *K1-K2-K3-K4-K5*, with K1 being the wealthiest area, and all the others following in an ordinal fashion. In light of this hierarchical relationship, we will denote these groups—or clusters—also as census *levels*.

The resulting map, shown in Fig. 5, resembles the distribution of the HDI shown in Fig. 1. In particular, the clustering procedure captures very well the presence of a very rich and highly segregated area (cluster K1) in the north-eastern part of the urban area, while the other census groups are distributed more heterogeneously throughout the city.

2.3.1 Study of the global spatial autocorrelation

As further evidence of the segregation that emerges from the k-means algorithm, we dig deeper into the spatial distribution of the census features. To do this, we measure Moran's

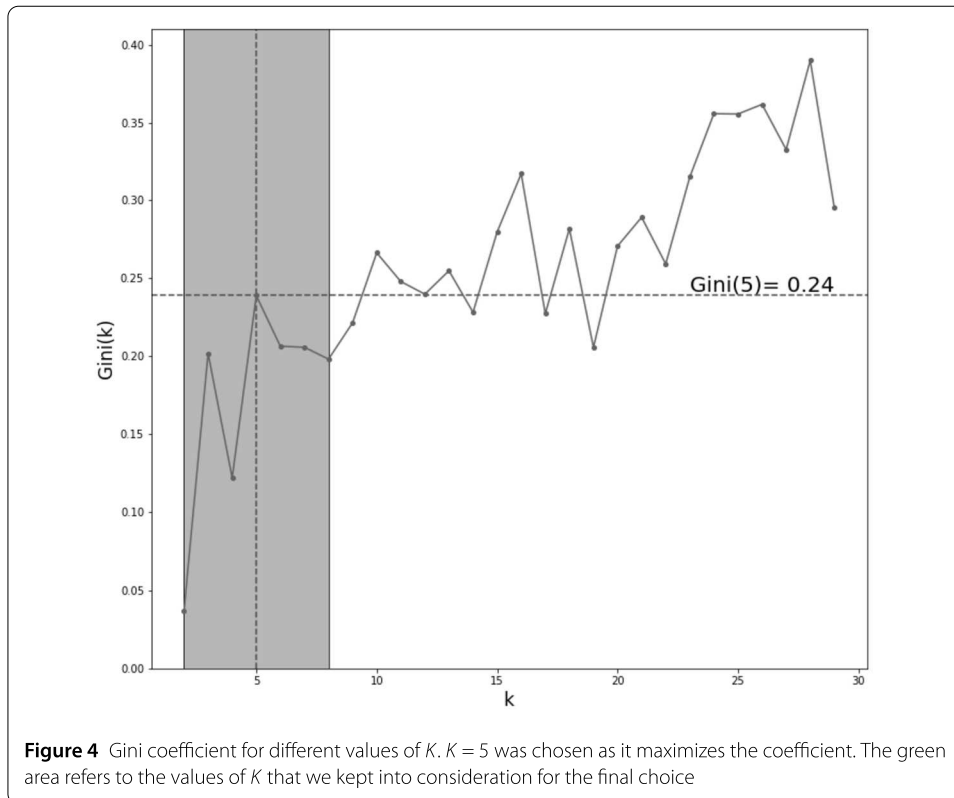


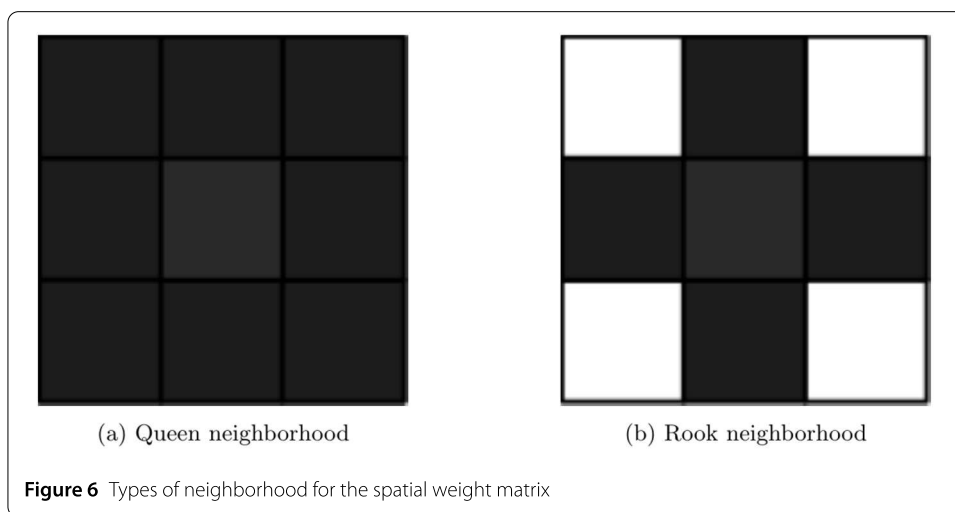
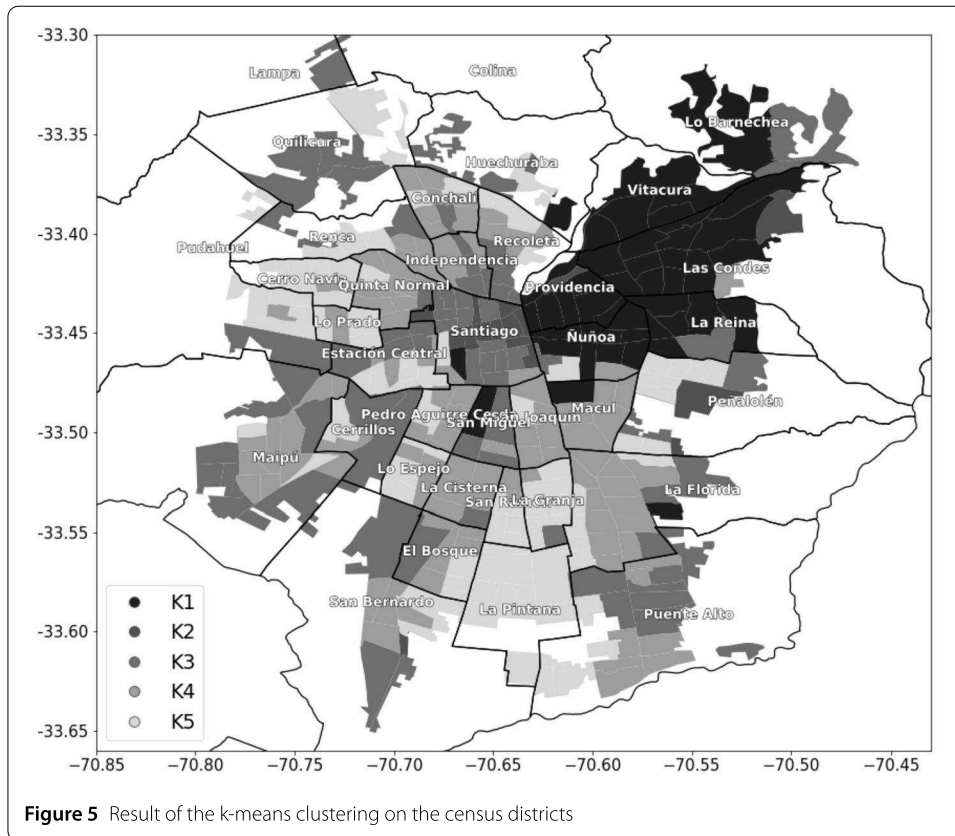
Table 3 Insights of the clusters. In these table are shown the mean values of the features for each cluster

Cluster	Mean age	Avg years of schooling	% of students	% of people of indigenous ethnicity
K1	46.25 ± 0.27	16.91 ± 0.11	0.154 ± 0.003	0.047 ± 0.002
K2	38.78 ± 0.58	16.50 ± 0.17	0.179 ± 0.006	0.075 ± 0.004
K3	42.05 ± 0.21	14.65 ± 0.12	0.138 ± 0.002	0.096 ± 0.002
K4	46.36 ± 0.13	14.30 ± 0.09	0.122 ± 0.001	0.097 ± 0.001
K5	44.62 ± 0.18	12.86 ± 0.08	0.106 ± 0.001	0.131 ± 0.002

Index [30] on census data and test the global spatial autocorrelation of the features against a random null model. Moran’s I for a variable y measured over n spatial units is defined as

$$I = \frac{n}{S_0} \frac{\sum_i \sum_j z_i w_{ij} z_j}{\sum_i z_i^2},$$

where $w_{i,j}$ are the elements of a spatial weight matrix W , $S_0 = \sum_i \sum_j w_{ij}$ and $z_i = y_i - \bar{y}$ is the deviation of y from its mean value in the spatial unit i . The matrix W is essential in spatial autocorrelation analysis, since it provides the model with a measure of spatial contiguity. The definition of spatial contiguity is usually specified as a *neighborhood* relationship between spatial units. Since we are dealing with census districts, namely areas of the city, we decided to use the Queen neighborhood (Fig. 6). We believe that using a Rook neighborhood in a context of urban spatial analysis would be too limiting because, for each computation, we would not be considering the influence of the areas in the corners of the spatial unit under examination. Thus, we defined as *contiguous* all those areas



that surround the zone we are observing, considering sufficient that they only meet at corner vertexes [31]: all those areas interact, communicate and most likely influence one another.

We measured the global spatial autocorrelation for the census since we need to corroborate any evidence of clustering resulting from the application of the k-means algorithm. The Moran’s I is an univariate measure of correlation, hence we measured it for every census feature. All the results were then compared with the Moran’s I calculated for the case of a completely random spatial distribution of the features.

In order to have a better understanding of the spatial autocorrelation and have a *more local* view, we examined the Moran scatterplot of the data points for each feature. The analysis of the Moran scatterplot was first introduced by Anselin [32], and it is based on the interpretation of the Moran's Index as a coefficient of an ordinary least squares (OLS) regression of the *spatially lagged* variable on the variable itself:

$$\sum_j w_{ij}z_j = a + Iz_i,$$

where $\sum_j w_{ij}z_j$ is the *spatial lag* of the variable z (which in turn is y expressed as deviation from the mean). Hence the spatial lag is a weighted average of the variable z over the neighbours of i . Note that the interpretation of I as a coefficient of an OLS regression is valid for any statistic that can be expressed as a ratio between a quadratic form and the sum of the squares [32], which is exactly how the Moran's I is defined.

Plotting the spatial lag of z as a function of z allows us to have a local view of the spatial autocorrelation. The plot is divided into four quadrants, going from the 1st quadrant of high values of both z and the spatial lag (*high-high correlation*) to the 4th quadrant of high-low correlation, passing through the so-called low-high and low-low correlation. A point being far up in the first quadrant of the Moran scatterplot of the *average schooling year* feature, represents an area that has a very high mean value of schooling, and is surrounded by neighbours with similarly high values.

As we can see in Fig. 7, there is strong evidence of segregation for almost all the census features, in particular for the *extreme* census segments K1 and K5 which, depending on the feature, are usually located way up near the bisector of the first and third quadrants. The census segments in between, instead, display a far more mixed situation.

3 Results and discussion

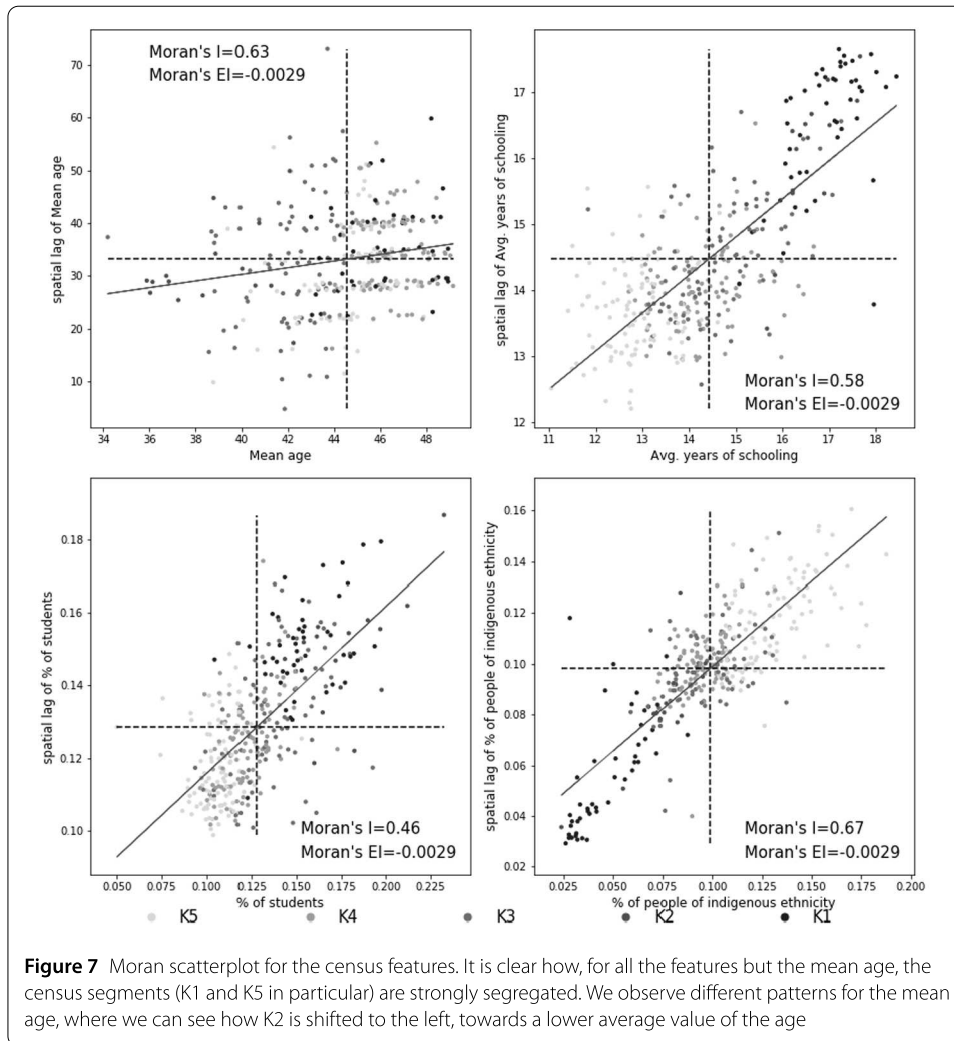
3.1 Geo-referenced analysis of the DPI dataset

The labeling of the census zones based on socio-demographic features finally enables us to proceed with a geo-referenced analysis of the connections towards the selected news outlets websites. The goal is to find differences in the consumption of news media content by areas of the city that have different socio-demographical attributes.

A list of the news outlets can be found in Table 2: wherever we encountered the issues described in Sect. 2.1.2, we grouped the news outlets by owner and considered the resulting group as an individual entity. This is the case of what happened with all the outlets belonging to the *El Mercurio* editorial group, which we included in our final list as a single entry.

In Fig. 9 we plot, for each News Outlet (NO), the number of unique users in each CD, normalized over the number of residents of the cluster the CD belongs to. Figure 10 shows the same quantity considering all the NOs together, i.e. the total amount of users connecting to news outlets websites during the examined weekends of July and August 2016.

Before analyzing these plots, it is worth going back to Table 3. The grouping into 5 clusters portrays a very clear division between the first two census segments and the remaining three. Indeed, groups K1 and K2 display a very high average value of schooling years (> 16), with group K2 being composed of much younger individuals and a higher percentage of students. Groups K3, K4 and K5 instead are more comparable with each other, displaying



a progressively lower value of schooling years and percentage of students, and a higher value of people of indigenous ethnicity. As pointed out in Sect. 2.3.1, these three groups are spatially mixed, while K1 and K2 are segregated in the north-eastern and, to a lower extent, in the central part of the urban area of the *Region Metropolitana de Santiago*.

In light of this grouping, we can inspect the results shown in the plots. From Fig. 10 we can gather the first, clear evidence: group K2 completely dominate the chart. This means that the majority of the connections are originated in those areas that are characterized by a high percentage of young and highly educated residents. As said, for each cluster the connections have been normalized over the number of people in it, meaning that we are looking at a measure of the *activity* of a cluster with respect to its size. This dominance of the group K2 and, in smaller measure, of K1, is confirmed also by computing the Pearson correlation between the number of connections and the number of residents in each cluster. We get $P = -0.92$ (p-value = 0.029), a very high anti-correlation, meaning that—in proportion to their size—the smallest clusters show the highest activity. Indeed, the smallest clusters are K2 (with only 255,699 people) and K1 (682,053 inhabitants), while the remaining three all have more than 1M residents each.

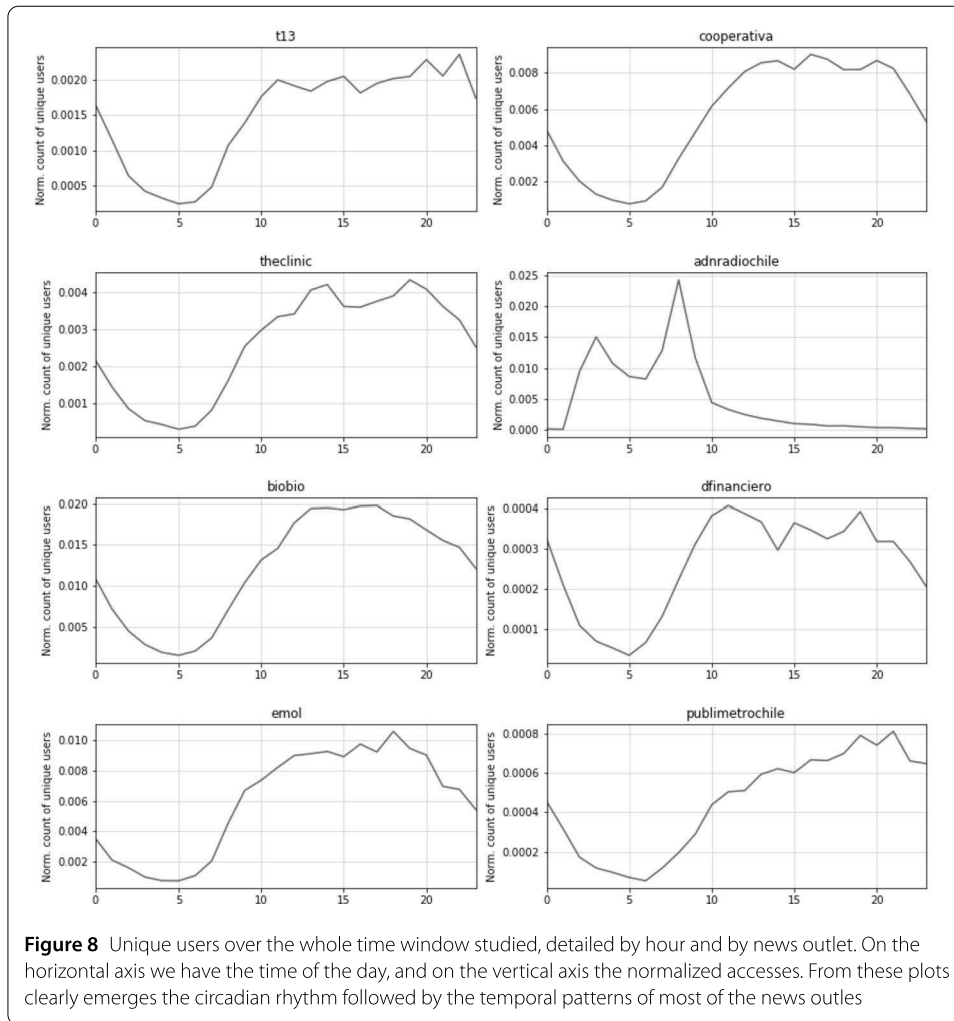
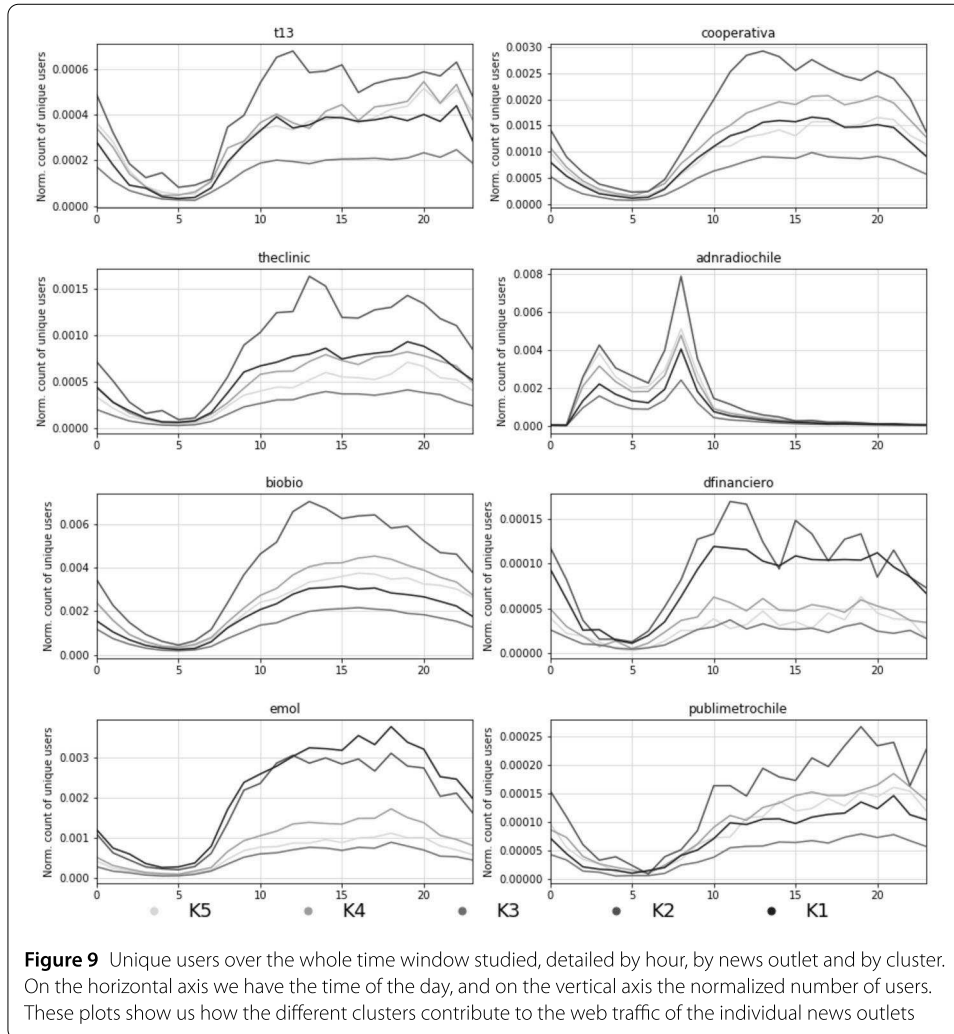


Figure 8 Unique users over the whole time window studied, detailed by hour and by news outlet. On the horizontal axis we have the time of the day, and on the vertical axis the normalized accesses. From these plots clearly emerges the circadian rhythm followed by the temporal patterns of most of the news outlets

We can also analyze the same trends for the individual news media websites, shown in Fig. 8 and disaggregated by clusters in Fig. 9. From Fig. 9 we can see that the areas belonging to group K2 are again on top in the majority of the plots. The only exception is for the websites belonging to *El Mercurio*, a right wing editorial group, which is most accessed by those areas belonging to the K1 segment. Similar numbers between K1 and K2 can be found also for the website of *Diario Financiero*, a financial newspaper. As for the other census groups no clear pattern can be identified, but still it seems pretty clear that, both in Figs. 9 and 10, group K3 is always the least active. In general, segments K3, K4 and K5 display an activity comparable to that of groups K1 and K2 only for those news media, like *Biobio*, *Cooperativa* or *Tele13*, that can be classified as containers of *generic* news and information, unlike *Diario Financiero*, which is highly specialised on financial matters, or *El Mercurio*, which is openly politically oriented.

As expected ([24, 33, 34]), the circadian rhythm of the population can still be seen from these figures, even though to a lower extent than the same analysis performed on the weekdays (Fig. 3), as explained in Sect. 2.1.2. The decline in connections towards night hours is smooth, and the trends are very similar across all the census segments.



3.1.1 Addressing the limitations

This method presents some limitations. First of all, we acknowledge a possible bias in our dataset: while census is representative of the whole population, mobile data accounts only for the customers of the service provider, which can be biased towards a certain segment of the population depending on multiple factors (typologies of the proposed contracts, marketing strategies, etc.). Nonetheless, bearing in mind that this provider alone owns more than the 37% of the marketshare, a correlation check between our dataset and the number of residents in each census zone gave us satisfactory results, as pointed out in Sect. 2.1.2. Another influence on the results could come from the diverse penetration rate of WiFi technologies among the various census groups. One of our basic assumptions is that people, during weekends, access the web mostly from their houses or from nearby locations. This means that, especially in their houses but also in some public places, people could be surfing the web via WiFi instead of using their own mobile data. This, again, strongly depends on the specific country and on its telco market, and it could also be related to the fact that—in general—WiFi is more popular among those individuals that have, for example, a bank account (sometimes needed even just to get a broadband contract in the first place) or at least a certain yearly income. Thus, we acknowledge that these census

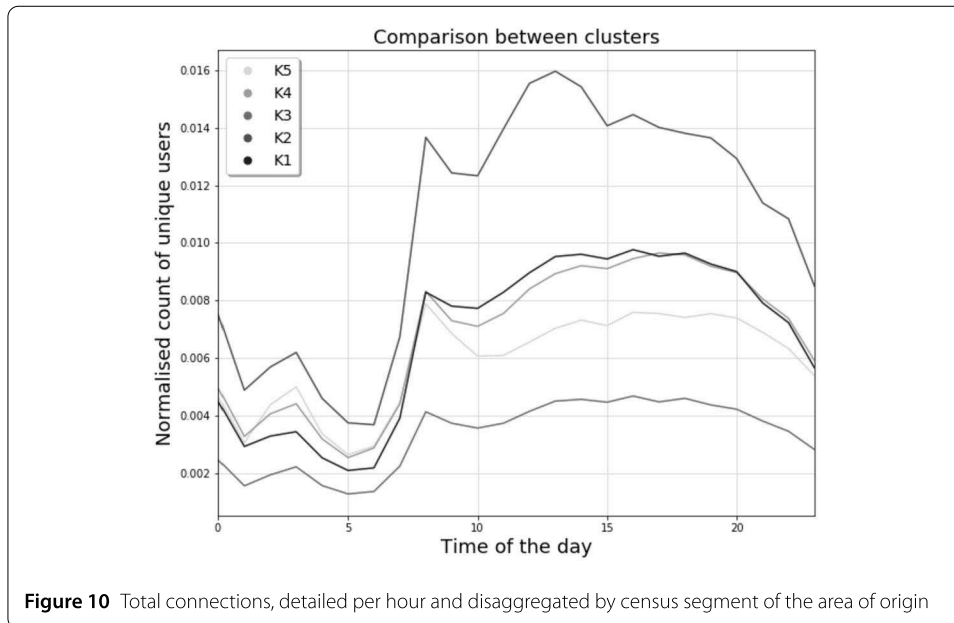


Figure 10 Total connections, detailed per hour and disaggregated by census segment of the area of origin

groups, in particular K1 and K2 segments according to our clustering algorithm, could be slightly under-represented in our dataset. Nonetheless, the positive correlation between the number of unique accesses and the residents in each municipality is a reassuring evidence with respect to this matter. Moreover, for the particular case of Chile, the effects of this issue should be limited: according to official data (Subtel), nearly 80% of all Internet access is made from mobile phones [35].

A further limitation of this study is data availability. We are under strict non-publication policies from the telco, and the data cannot leave the telco's premises. This is to be expected given the novelty of the data sets we are working with, and their yet-to-be-understood privacy implications (this notwithstanding, please see the "Data Availability" section below)

4 Conclusions

The main purpose of this work is to understand whether individuals living in different areas of the city of Santiago de Chile display different behaviors in terms of access to digital news media content with respect to socio-demographic attributes. In order to do so, we studied a record of a month (July–August 2016) of anonymized and geo-referenced accesses to several websites of Chilean news media through the cellphone network of a major telco provider. To reduce the effects of the *floating population* phenomenon, we analyzed only a portion of the connections, narrowing the window to the weekends only.

By applying a k-means algorithm on official Chilean census data, in order to group the population into 5 different census levels, we find that the wealthiest areas of the city, K1 and K2, are very similar in terms of *years of schooling*, with K2 members being—on average—much younger than K1. The other three clusters (K3–K4–K5) are indeed pretty far from the first two in terms of average education level, and far more spatially mixed: K1 and K2 are mostly grouped in the north-eastern quadrants of the city.

The news access patterns tell us that segment K2 is overall the most active group. This means that the most highly educated and youngest are also the most informed, with re-

spect to older people with similar education level. If we break down the number of access by individual media outlet, we obtain a more composite picture (Fig. 3). K2 remains the most comprehensively informed with respect to the diversity of media outlets; on the other hand the ranking of the other groups can vary depending on the flavour of the news media.

These findings confirm empirically that socio-demographic features have an influence on the consumption of news. Our results suggest that the consumption of news media content does not always necessarily increase with the education level of the user. While it is true that the highly educated are those who access news media websites the most, the opposite cannot be inferred: indeed group K3, which includes all those areas of the city whose residents are averagely educated (see Table 3), is the cluster that displays the lowest activity.

The predominance of group K2 also tells us that age plays a key role in this context. Young generations are likely more comfortable in using mobile devices in their daily lives, and this has been also explored in the specific field of news consumption [36].

In summary, although we found clear correlations between the consumption of online news media content and the socio-demographics of users, this relationship seems to be non-trivial. While the highly educated are the most eager users of mobile news media outlets, there are important clues that a group of well educated people, that we could identify as middle-class, show a lack of interest in accessing news media via mobile. This could have a significant impact on the civic engagement of a relevant share of the population, as well as on the economical and political life of a Country, not mentioning the media market itself.

There are positive correlations between the other two features—*student status* and *percentage of people of native ethnicity*. Therefore, a more detailed analysis of the ties between these features and the consumption patterns of the individual news outlets—namely, the habits and preferences of students and native minorities—could be of much interest in the context of studying inequalities in the accesses between different cohorts of the population, and it is left as future work.

We believe that this work could be useful to shed some light on how digital platforms can contribute to the already complex interplay between socio-demographic characteristics of the population and the news consumption behaviour. Additional studies with more comprehensive datasets in other social contexts are needed, to assess the full impact of digital technologies on this topic, as well as to understand if these findings are highly Country-specific or if there is a shared worldwide-trend.

Acknowledgements

Daniela Paolotti and Salvatore Vilella acknowledge support from the Lagrange Project of the Institute for Scientific Interchange Foundation (ISI Foundation) funded by Fondazione Cassa di Risparmio di Torino (Fondazione CRT). LF acknowledges financial support from Movistar—Telefónica Chile, the Chilean government initiative CORFO 13CEE2-21592 (2013-21592-1-INNOVA_PRODUCION2013-21592-1), Conicyt PAI Networks (REDES170151) “Geo-Temporal factors in disease spreading and prevention in Chile”, and Project PLU180009, Tenth “Fondo de Estudios sobre el Pluralismo en el Sistema Informativo Nacional”, 2018, “Geo-Temporal Access to Chilean news outlets using digitales traces”.

Availability of data and materials

2017 Chilean Census data is available at <https://www.censo2017.cl>. DPI dataset was provided by the mobile telecommunication company, and therefore is not publicly available. However, if any researcher would like to access/work on the data, the telco's offices are open to collaboration through the authors of this paper. Please send an email to lferres@udd.cl.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

Conceptualization: LF, DP, GR, SV; Data curation: SV, LF; Formal analysis and Methodology: SV, DP, GR, LF; Writing: SV, LF, DP, GR. All authors read and approved the final manuscript.

Author details

¹Department of Computer Science, University of Turin, Turin, Italy. ²Data Science Institute, Universidad del Desarrollo, Santiago, Chile. ³Telefónica R&D, Santiago, Chile. ⁴ISI Foundation, Turin, Italy.

Endnotes

^a Released by the *Instituto Nacional de Estadística (INE)* See <https://www.censo2017.cl>.

^b <https://bit.ly/2Ukpbly>.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 22 July 2019 Accepted: 21 April 2020 Published online: 29 April 2020

References

1. Alejandro J (2010) Journalism in the age of social media. Reuters Institute Fellowship Paper
2. Torche F (2007) Social status and cultural consumption: the case of reading in Chile. *Poetics* 35(2–3):70–92
3. Molyneux L (2018) Mobile news consumption: a habit of snacking. *Dig Journal* 6(5):634–650
4. Boczkowski PJ, Mitchelstein E, Matassi M (2018) “News comes across when I’m in a moment of leisure”: understanding the practices of incidental news consumption on social media. *New Media Soc* 20(10):3523–3539
5. Lindell J (2018) Distinction recapped: digital news repertoires in the class structure. *New Media Soc* 20(8):3029–3049
6. Chan TW, Goldthorpe JH (2007) Social status and newspaper readership. *Am J Sociol* 112(4):1095–1134
7. Anderson B (2018) Influence of education, income and age on newspaper use and platform preference. *Elon J Undergrad Res Commun* 9(1):108–114
8. Self CC (1988) Perceived task of news report as a predictor of media choice. *Journal Q* 65(1):119–125
9. Poindexter PM, McCombs ME (2001) Revisiting the civic duty to keep informed in the new media environment. *Journal Mass Commun Q* 78(1):113–126
10. Kalogeropoulos AK, Nielsen RK (2017) Social inequalities in news consumption
11. Bahamonde J, Bollen J, Elejalde E, Ferres L, Poblete B (2018) Power structure in Chilean news media. *PLoS ONE* 13(6):e0197150
12. Elejalde E, Ferres L, Herder E (2018) On the nature of real and perceived bias in the mainstream media. *PLoS ONE* 13(3):e0193765
13. Elejalde E, Ferres L, Herder E, Bollen J (2018) Quantifying the ecological diversity and health of online news. *J Comput Sci* 27:218–226. <https://doi.org/10.1016/j.jocs.2018.06.003>
14. Elejalde E, Ferres L, Schifanella R (2019) Understanding news outlets’ audience-targeting patterns. *EPJ Data Sci* 8(1):16. <https://doi.org/10.1140/epjds/s13688-019-0194-8>
15. Rodrigues UM, Paradies Y (2018) News consumption habits of culturally diverse Australians in the digital era: implications for intercultural relations. *J Intercult Commun Res* 47(1):38–51
16. Edgerly S, Thorson K, Thorson E, Vraga EK, Bode L (2018) Do parents still model news consumption? Socializing news use among adolescents in a multi-device world. *New Media Soc* 20(4):1263–1281
17. Fletcher R, Nielsen RK (2017) Paying for online news: a comparative analysis of six countries. *Dig Journal* 5(9):1173–1191
18. Matthijs K, De Coninck D, Debrael M, d’Haenens L, De Cock R (2019) Unpacking attitudes on immigrants and refugees: a focus on household composition and news media consumption. *Media Commun* 7(1):43–55
19. Mao H, Shuai X, Ahn Y-Y, Bollen J (2015) Quantifying socio-economic indicators in developing countries from mobile phone communication data: applications to Côte d’Ivoire. *EPJ Data Sci* 4(1):15
20. Gonzalez MC, Hidalgo CA, Barabasi A-L (2008) Understanding individual human mobility patterns. *Nature* 453(7196):779–782
21. Calabrese F, Colonna M, Lovisolio P, Parata D, Ratti C (2011) Real-time urban monitoring using cell phones: a case study in Rome. *IEEE Trans Intell Transp Syst* 12(1):141–151
22. Calabrese F, Ferrari L, Blondel VD (2015) Urban sensing using mobile phone network data: a survey of research. *ACM Comput Surv* 47(2):25
23. Iqbal MS, Choudhury CF, Wang P, González MC (2014) Development of origin–destination matrices using mobile phone call data. *Transp Res, Part C, Emerg Technol* 40:63–74
24. Graells-Garrido E, Ferres L, Caro D, Bravo L (2017) The effect of Pokémon Go on the pulse of the city: a natural experiment. *EPJ Data Sci* 6(1):23
25. Beiró MG, Bravo L, Caro D, Cattuto C, Ferres L, Graells-Garrido E (2018) Shopping mall attraction and social mixing at a city scale. *EPJ Data Sci* 7(1):28
26. Wesolowski A, Buckee CO, Engø-Monsen K, Metcalf CJE (2016) Connecting mobility to infectious diseases: the promise and limits of mobile phone data. *J Infect Dis* 214(Suppl_4):S414–S420
27. De Gregorio J, Lee J-W (2002) Education and income inequality: new evidence from cross-country data. *Rev Income Wealth* 48(3):395–416
28. Jerrim J, Macmillan L (2015) Income inequality, intergenerational mobility, and the Great Gatsby Curve: is education the key? *Soc Forces* 94(2):505–533
29. Graells-Garrido E, Saez-Trumper D (2016) A day of your days: estimating individual daily journeys using mobile data to understand urban flow. In: Proceedings of the second international conference on IoT in urban space, pp 1–7
30. Moran PA (1950) Notes on continuous stochastic phenomena. *Biometrika* 37(1/2):17–23
31. O’Sullivan D, Unwin D (2014) Geographic information analysis. Wiley, Hoboken

32. Anselin L (1993) The Moran scatterplot as an ESDA tool to assess local instability in spatial association. Regional Research Institute, West Virginia University, Morgantown, WV
33. De Nadai M, Staiano J, Larcher R, Sebe N, Quercia D, Lepri B (2016) The death and life of great Italian cities: a mobile phone data perspective. In: Proceedings of the 25th international conference on world wide web, pp 413–423
34. Graells-Garrido E, Peredo O, García J (2016) Sensing urban patterns with antenna mappings: the case of Santiago, Chile. *Sensors* 16(7):1098
35. 2017 digital news report, Reuters Institute for the Study of Journalism and University of Oxford. <http://www.digitalnewsreport.org/survey/2017/chile-2017/#fn-6155-1>. Accessed: 2019-02-26
36. Chan M (2015) Examining the influences of news use patterns, motivations, and age cohort on mobile news use: the case of Hong Kong. *Mob Media Commun* 3(2):179–195

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
