

Assessing the Quality of Home Detection from Mobile Phone Data for Official Statistics

Maarten Vanhoof¹, Fernando Reis², Thomas Ploetz³, and Zbigniew Smoreda⁴

Mobile phone data are an interesting new data source for official statistics. However, multiple problems and uncertainties need to be solved before these data can inform, support or even become an integral part of statistical production processes. In this article, we focus on arguably the most important problem hindering the application of mobile phone data in official statistics: detecting home locations. We argue that current efforts to detect home locations suffer from a blind deployment of criteria to define a place of residence and from limited validation possibilities. We support our argument by analysing the performance of five home detection algorithms (HDAs) that have been applied to a large, French, Call Detailed Record (CDR) data set (~ 18 million users, five months). Our results show that criteria choice in HDAs influences the detection of home locations for up to about 40% of users, that HDAs perform poorly when compared with a validation data set (resulting in 35°-gap), and that their performance is sensitive to the time period and the duration of observation. Based on our findings and experiences, we offer several recommendations for official statistics. If adopted, our recommendations would help ensure more reliable use of mobile phone data vis-à-vis official statistics.

Key words: Mobile phone data; home location; home detection algorithms; official statistics; big data.

1. Introduction

By now, big data has well and truly arrived and their potential as well as the challenges it poses for official statistics have become much more evident. Consequently, there has been a clear demand to invest in pilot projects that explore how big data can be integrated into official statistics (Eurostat 2014; Glasson et al. 2013).

From a practical perspective, pilot projects are useful not only to identify practical issues (e.g., legal issues, data management). They are also particularly useful when

¹ Open Lab, Urban Sciences Building, 1 Science Square, Science Central, Newcastle upon Tyne NE4 5TG, United Kingdom. Email: m.vanhoof1@newcastle.ac.uk

² Task Force Big Data, Eurostat, European Commission, Joseph Bech Building, 5, Rue Alphonse Weicker, L-2721 Luxembourg. Email: fernando.reis@ec.europa.eu

³ School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA, U.S.A., and Open Lab, Urban Sciences Building, 1 Science Square, Science Central, Newcastle upon Tyne NE4 5TG, United Kingdom. Email: thomas.ploetz@gatech.edu

⁴ Department SENSE, Orange Labs, Orange Gardens, 44 Avenue de la République, 92320 Châtillon, France. Email: zbigniew.smoreda@orange.com

Acknowledgments: The authors wish to thank Orange Labs France, especially the SENSE department, for making available the mobile phone data. We also like to thank INSEE, especially Stéphanie Combes, Marie-Pierre de Bellefon, Pauline Givord and Vincent Loonis for the construction of the validation dataset. Ultimately, we like to thank Morag Ottens and Andy Garbett for the help in preparin the article.

critically assessing the reliability of data sources and methodologies. It is reassuring that, regarding such assessments, [Karlberg et al. \(2015, 1\)](#) observe that: “There is a clear trend towards a more reflective approach, with an emphasis not only on producing high-quality statistics, but also on rendering explicit details on exactly how this is being achieved”. When it comes to big data the importance of providing explicit details is not to be underestimated as big data sources, typically, do not adhere to official statistics’ standards and principles – such as issues on coverage, representativity, quality, accuracy and precision ([Daas et al. 2015](#)) – and, consequently, neither do their methodologies.

In this article, we present a pilot study focusing on, arguably, the most important step for the application of mobile phone data in official statistics: identifying where someone lives, that is, detecting their home location. Current home detection methods for mobile phone data do not adhere to official statistics standards (or even to what could reasonably be expected from academic standards). We elaborate our argument by means of an extensive review of literature and an empirical analysis based on a large-scale, French, Call Detailed Records (CDR) data set. In doing so, we aim to show how current home detection practices came to be, how they are bound by limited validation possibilities and how they are sensitive to criteria choice or decision rule development. Given the lack of research on these problems, we argue that there is no clear framework on which to apprise the performance or the uncertainty of current home detection methods.

Our analysis evaluated the performance of five different home detection algorithms using a mobile phone data set from France. The case study allows us to reflect on the findings from a more practical point of view, whilst also contributing to our discussions and recommendations on the various uncertainties that underlie current home detection practices. We hope our contribution will help other researchers and practitioners to recognise the difficulties of integrating information on home locations sourced from mobile phone data into official statistics.

2. Mobile Phone Data, Official Statistics and the Role of Home Detection Methods

2.1. Mobile Phone Data and Official Statistics

Before looking at the methods used to identify home location, let us quickly consider how mobile phone data can be of interest for official statistics.

Over the last decade, the analysis of mobile phone data has grown into a mature research field with a wide array of applications that are being developed and applied ([Blondel et al. 2015](#)). One line of interest is that mobile phone data have the potential to capture temporal patterns of user presence ([Deville et al. 2014](#)), which could be used to estimate population density ([Ricciato et al. 2015](#)). In turn, these estimations could usefully support official statistics in developing countries ([Blondel et al. 2012](#); [de Montjoye et al. 2014](#)).

Another line of interest relates to the large-scale recording of mobility patterns. As mobile phones can capture individual mobility for millions of users, applications have been developed that estimate nationwide commuting figures ([Kung et al. 2014](#)), long-distance trips ([Janzen et al. 2016](#); [Janzen et al. 2018](#)), inbound tourism trips ([Raun et al. 2016](#)) and even domestic tourism trips ([Vanhoof et al., 2017b](#)).

These and similar developments have the potential to enhance official statistics in fields such as the delineation of urban areas (Vanhoof et al. 2017a), the understanding of migration patterns (Blumenstock 2012), or to complement tourism statistics (Ahas et al. 2008). They could even perform *nowcasting* of macro-economic and socio-economic aspects of populations (Baldacci et al. 2016; Marchetti et al. 2015; Giannotti et al. 2012; Pappalardo et al. 2016, and Vanhoof et al. 2018).

2.2. *The Role and Method of Home Location*

Common to many, if not all, mobile phone data research is the need to identify the home location of mobile phone users before proceeding to more advanced analysis. For example, knowing the place of residence is a prerequisite before analysing the amount of time spent at home and commuting patterns, which in turn fuel mobility and epidemiological models (Rubrichi et al. 2017). Besides its relevance within mobile phone data analysis, knowledge of home location also forms the crucial link between mobile phone data and other data sources, such as census data, making it a key enabler for the combination of information.

The method of pinpointing where someone lives consists of attributing a supposed home location to every single user in the database from the geographical metadata obtained from their mobile phone records. In practice, identifying a person's home means that a single cell tower is allocated as their home location. This allocation is based on the calling and movement patterns of each individual user. The spatial resolution of cell towers is used because most mobile phone data sets only have geographical data of the towers' positions. The assumption then is not that a user lives at that exact cell tower location, but rather somewhere in the area covered by the tower. It is remarkable that even though detecting the home location now forms a cornerstone of mobile phone research, home detection methods are often obscured in literature: details on their exact application, related uncertainties, perceived performance or even the validation processes are only rarely communicated.

In the following section, we show why current home detection practices are problematic. In an extended literature review, we show how, over time, methodologies for home detection have been simplified to single-step approaches using decision rules that are based on simple, *a priori* defined criteria of what defines a "home". Such methods are questionable because the possibilities to validate are limited, and there is a lack of knowledge on their sensitivity, specifically in respect to criteria choice. Our empirical work with a large, French, mobile phone data set exemplifies several of the problems we raise. It allows us to put the problems in a more practical context and outline their consequences in more detail.

3. **Identifying Homes from Large-Scale Location Traces**

Given the enormity of the data sets that capture geolocated traces of users, literature explains the automated methods developed for identifying the homes, or other meaningful places such as the workplace, of users. Here, it is necessary to distinguish between continuous location traces (e.g., GPS data) and noncontinuous location traces (e.g., mobile

phone data) where the latter do not provide a similar high-volume, high-resolution capture of location traces in time or space compared to the former.

As our main interest is to outline the deficits in the methods used for noncontinuous location traces, this section will start by reviewing the literature on automated home location.

3.1. Identifying Meaningful Places from Continuous Location Traces

The analysis of continuous location traces has been the focus of early developments in the automated identification of meaningful places. Related work typically used small-scale data sets, most commonly from continuous GPS traces, but also from Bluetooth, or Wi-Fi positioning (Wolf et al. 2001; Shen and Stopher 2014). The general methodology used to identify meaningful places from continuous location traces consists of a two-step approach.

In the first step, location traces are clustered in space (and sometimes in time) in order to detect *important places*. Techniques for clustering continuous location traces range from manual GIS analysis (Wolf, Guensler, and Bachman 2001; Gong et al. 2012) to automated, unsupervised analysis using, for example, k-means clustering (Ashbrook and Starner 2003), nonparametric Bayesian approaches (Nurmi and Bhattacharya 2008), or fingerprinting of the radio environment (Hightower et al. 2005).

In a second step, the important places identified are then annotated as *meaningful places* (such as home, work, recreation area). Annotation can be done either through interpretation, for example by expert judgment, by surveying the user that produced the traces, or through automation, mainly by means of time-space heuristics (Nurmi and Bhattacharya 2008).

3.2. Identifying Meaningful Places from Noncontinuous Traces

In contrast to the above-mentioned continuous location traces, the use of noncontinuous location traces has recently become very popular. Examples of activities that produce noncontinuous location traces in large-scale data sets are mobile phone usage, credit card transactions, or check-ins through location-based services (e.g., Foursquare) and online social networks (e.g., Twitter).

The identification of meaningful places from noncontinuous location traces poses substantial challenges, most notably due to the less frequent observations and the larger spatial resolution in which observations are captured (e.g., mobile phone data are only captured at the location of the cell tower used). However, these challenges are outweighed by the presumed advantages associated with the larger coverage, in terms of users, timespan and spatial extent of the data sources (Järv, Ahas, and Witlox 2014; Kung et al. 2014).

The following analysis will focus on one example of how to identify one meaningful place – the location of a user’s home, using one prominent example of noncontinuous traces: Call Detailed Record (CDR) data.

CDR data are mobile phone data captured by the network operator every time a user makes or receives a text or call (hence the noncontinuous tracing). Note that the methods

and problems described in the following sections are not limited to CDR data, but are relevant for all data sets covering noncontinuous location traces.

3.2.1. Two-Step Approaches for Noncontinuous Traces

As with the two-step approaches for continuous traces, initial methods to detect home locations from CDR data also clustered location traces into important places before annotating them as meaningful places. For example, in Isaacman et al. (2011) individual traces from CDR data are clustered using Hartigan's leader algorithm. Clusters are then annotated into meaningful places by means of a logistic regression model that is trained on data from 18 persons for which ground truth was available. Next, and for each user, the cluster with the highest score on the logistic regression model is chosen to be the presumed home area.

3.2.2. Single-Step Approaches for Noncontinuous Traces

However, two-step approaches for noncontinuous location traces quickly gave way to single-step approaches that are now widely deployed in literature (Calabrese et al. 2014; Calabrese et al. 2011; Kung et al. 2014; Phithakitnukoon et al. 2012). The difference between two-step and single-step methods is that the latter skips the clustering into important locations and thus acts directly on individual cell towers instead of groups of cell towers.

One of the reasons for switching to single-step approaches is that the standard clustering methods used in the two-step approaches make it difficult to construct consistent spatial traces when combined with noncontinuous location traces. Nevertheless, the main drawback of this switch to a single-step approach is that the spatial pattern of the location traces is largely neglected, as only single cell tower annotation is targeted. This increases the uncertainty of fixing home location, because single events at individual cell towers may be sufficient to undermine the method.

In practical terms, detecting a home in a single-step approach is done by using a decision rule that is based on an *a priori* definition of home – the *home criterion* as we call it – in order to produce a list of one or several cell towers that could be the home location. A standard example of a home criterion for the case of CDR data is “home is where calls are made during the night”. The problem with single-step approaches is that such decision rules are being applied as *heuristics*, meaning that one general rule is applied to the location traces of all users even though a different set of decision rules could potentially lead to better results.

In terms of identifying home location, applying heuristics implies that meaningful places (like the home) can be described similarly for all users in the data set, regardless of the user's characteristics as observed in their movements and calling patterns. It seems logical that the imposition of this assumption can only be done when a proper evaluation and validation of their movements has been carried out, or when clear evidence exists for the use of a specific criterion or decision rule. For this reason, the following paragraphs will discuss how to define decision rules for one-step home detection methods and which criteria to use.

3.3. Defining Decision Rules for Single-Step Home Identification

3.3.1. Simple Decision Rules for Single-Step Home Detection

The core challenge for single-step home detection is in defining a decision rule that is simultaneously capable of i) distinguishing between different important places, and ii) annotating the correct home location. Most research employs simple decision rules that are either based on information from official statistics or rely on precedents found in literature.

When examining the existing decision rules in research literature, the most popular are: time-based limitations for the night (“home is the location that has the most activity between x p.m. and y a.m.”), time-based aggregations (‘home is where the most distinct days, or weekend-days are spent), and spatial groupings (‘home is the location with the most activity in a spatial radius of x km around it), (Calabrese et al. 2011; Phithakkitnukoon et al. 2012; Frias-Martinez and Virseda 2012; Kung et al. 2014; Tizzoni et al. 2014). One example, using time-interval statistics from a Boston data set drawn from the American Time Use Survey (Calabrese et al. 2011), uses the highest distinct number of observations between 6 p.m. and 8 a.m. to derive home locations.

Almost all studies using simple decision rules rely on census data. They depend either on specific surveys and questionnaires to define the criteria deployed, (Calabrese et al. 2011) or, for high-level validation, on aggregated population density data (Phithakkitnukoon et al. 2012) or commuting Figures (Kung et al. 2014).

3.3.2. Complex Decision Rules for Single-Step Home Detection

A few studies have elaborated more complex decision rules for home detection. The seminal work of Ahas et al. (2010), for example, uses a tree-based approach that combines a set of criteria including distinct days of activities on a cell tower, the starting times of calls, deviations of starting time of calls, durations of calls, and this all for a training set of 14 people for which the ground truth was known. The decision rules, as defined by the classification tree, were consequently deployed to all users in an Estonian data set (as heuristics in other words), raising the question of how representative a training set of 14 people could possibly be for a large population.

The problem of small training sets was overcome in Frias-Martinez et al. (2010), who used a training set of 5,000 users to construct a complex decision rule for home detection. Deploying a Genetic Algorithm technique, they focus on finding the best combination of temporal criteria to denote home locations in an emerging economy. Their best performance is a correct prediction of around 70% for a subset of 50% of the users. Users were filtered on the basis of having at least a 20% difference in the percentage of total calls between the first and second eligible cell tower. The complex decision rule they use to obtain this result is to select the cell tower logging the most activity during the nights of Friday, Saturday, Sunday, Monday and Tuesday from 5:15 p.m. to 8:30 a.m.

The individual ground truth data in Frias-Martinez et al. (2010) are retrieved from users’ contracts with the provider. This data is not available in most countries due to legal obligations to anonymise users or bans on linking individual information to CDR data.

As a consequence, Csáji et al. (2013), tried to derive a temporal decision rule, but this time without a training/validation data set at individual level. Applying an unsupervised

k-means algorithm to the temporal activity patterns of frequently used cell towers in Portugal, they found clusters that are interpretable as temporal patterns typically relating to presence at home, at work, or as not interpretable at all. Consequently, their decision rule to detect home locations was based on these temporal patterns interpreted as home presence. Compared to [Frias-Martinez et al. \(2010\)](#), one of the drawbacks of their approach is that they did not construct their criteria based on individual observations. This raises the question as to the degree to which such criteria are realistic for different subsets of users.

In a way, the subset representativity problem persists for all single-step approaches, regardless of whether their decision rules are defined in a complex or simple way. If the same decision rules is applied to all phone users, careful investigation into the effect at individual level, or at population subset level should be carried out, in order to know the degree to which generalisation favours or disfavors subsets of users. In other words, if decision rules are applied generically, indepth validation of the single-step approaches is important.

3.4. *Validating Large-Scale Home Detection Methods*

The use of a particular decision rule, whether derived from a census, borrowed from literature or defined by training sets, is often based on comparing population counts from mobile phone data with census data. However, such high-level validation does not offer a direct evaluation of performance at individual user level, nor does it allow for comparison between cases. In fact, assessing the performance of different decision rules by comparing the resultant population counts with census data is, strictly speaking, a rather limited alternative solely justified by the absence of individual level validation data.

The absence of validation data at individual level is a common problem in published research, and is therefore often taken for granted. However, the absence of validation data has several consequences. First and foremost, it impedes the creation of evaluation metrics that can assess the performance of home detection at individual level. Such an individual level evaluation could allow us to better understand the workings of different decision rules on a specific data set and user subsets, which in turn could enable a comparison between different decision rules, data sets, users and areas.

Secondly, the absence of validation data at individual level is implicitly why single-step approaches apply decision rules as heuristics. In the absence of individual level validation data, it is impossible to understand which decision rules works best for any individual user. Consequently, case-adjusted, adaptive algorithms cannot be developed. This implicitly forces researchers and practitioners to adhere to a one-size-fits-all solution in order to be clear and consistent.

It is worth noting that, currently, high-level validation is still assumed to be a good solution in the absence of individual level validation data. In particular, two observations stand out.

Firstly, census data is often used for high-level validation. For example, comparisons for small geographical areas can be made between the counts of home locations identified from mobile phone data and the aggregated counts of peoples' residential locations obtained from censuses. This is a very opportunistic, if not naïve, validation attempt as

census data has never specifically been gathered to serve this purpose and little or no information exists on how, for example, different spatial delineations or the distorted market shares of mobile phone operators could influence this kind of validation.

Secondly, it is noteworthy that no studies have used high-level validation to compare the performance of different decision rules. Nor are there studies that evaluate the sensitivity of high-level validation to criteria choice. This absence is probably because high-level validation is not informative enough to properly understand the differences between criteria, decision rules, and their performances. Given this, we are far from obtaining a consensus on which criteria are best, or on how to construct optimal decision rules. In fact, we are far from understanding the strengths and weaknesses of different home detection methods altogether. Given this, we should question the degree to which high-level validation contributes to the development and trustworthiness of home detection.

3.5. Current Deficits of Home Detection Using Noncontinuous Location Traces

In conclusion, we find a clear framework is missing to allow us to understand the performance, uncertainty and sensitivity of the criteria choice or decision rule development, especially at individual level, when using noncontinuous location traces to detect home location. Despite their widespread use, no clear reasoning exists as to why single-step approaches should be chosen over two-step approaches. Nor does a consensus exist on which criteria should be used, or how optimal decision rules for a given data set should be defined.

Similarly, it is striking that no work investigates the sensitivities of single-step approaches to criteria choice. Additionally, we find that the validation of large-scale home detection methods is severely limited because of the absence of ground truth data at individual level. As a result, current assessments of home detection methods are based on high-level validation, but the trustworthiness and exact contribution of this practice is rather dubious.

In summary, our findings indicate that the current methods to identify users' home locations for official statistics are rather questionable. We illustrate some of the aforementioned problems by means of a case study for identifying home locations using French CDR data.

4. Investigating Home Detection Algorithms for French CDR Data

To explore the application of single-step home detection methods on a French CDR data set, we start by constructing five home detection algorithms that incorporate different popular home criteria in simple decision rules. We apply these algorithms to the French data set, perform high-level validation, and investigate sensitivity to criteria choice. This allows us to demonstrate some of the aforementioned problems in an applied context.

4.1. The French CDR Data Set

CDR data are the most widely-used examples of mobile phone data in research. CDR data are passively gathered by operators for billing and maintenance purposes and are collected

every time a mobile phone user makes or receives a text or a call. Apart from technical metadata on the workings of the network, CDR data contain information on the time, the location (the cell tower used), as well as the caller and the call receiver.

For our analysis, we use an anonymised CDR data set from the mobile phone carrier Orange. The data covers the mobile phone usage of ~ 18 million users on the Orange network in France during a period of 154 consecutive days in 2007 (13 May 2007 to 14 October 2007). At that time, mobile phone penetration was estimated at 86% (ARCEP 2008). Given a population of 63.9 million inhabitants during the observed period (counted as the average of monthly estimates between May and October 2007 as obtained from the INSEE Website: www.insee.fr), this data set covers about 32.8% of all French mobile phone users and 28.6% of the total population.

The Orange France 2007 CDR data set is one of the largest CDR data sets available in terms of population-wide coverage and has been extensively studied before (Grauwin et al. 2017; Sobolevsky et al. 2013; Deville et al. 2014). It is the latest CDR data set available for France that allows for long-term, temporal continuous tracking of mobile phone users. Access to more recent data sets is limited by The French Data Protection Agency (CNIL), which is anticipating the EU General Data Protection Regulation and does not allow individual traces for periods of more than 24 hours to be collected, before being irreversibly recoded.

Some of the typical characteristics of CDR data sets that pose substantial challenges for their automated analysis are the temporal sparsity in observations and the spatially uneven distribution of the areas covered. The former results in only a few records per user per day. For example: for an arbitrary day in the French data set (Thursday, 1 October 2007), the median number of records per user was four, relating to only two different locations. Such statistics are representative for CDR based studies and can be deemed rather high compared to other large-scale noncontinuous datasets like credit-card transactions or Flickr photos (Bojic et al. 2015). The latter is the result of a demand-driven, nonuniform distribution of cell tower locations (higher densities of cell towers are found in more densely populated areas, such as cities or coastlines), meaning that the spatial accuracy of the dataset is restricted to the network's spatial resolution.

On the other hand, it is very attractive to have the possibility of researching the large-scale CDR data sets at population level, without users needing to share their locations. This increases the feasibility of automated applications such as home location. In addition, continuous data collection allows us to observe over extended periods, which in turn enables complex analysis and lessens any influence emanating from singular events and/or nonroutine behaviour.

4.2. *Applying Five HDAs to the French CDR Data*

4.2.1. *Constructing Five HDAs with Simple Decision Rules Based on Popular Home Criteria*

To perform home detection, we construct five basic Home Detection Algorithms (HDAs). Each incorporates one or two popular home criteria that are applied by means of simple decision rules. In order to select criteria, we took into account literature that dealt with

single-step approaches (e.g., [Ahas et al. 2010](#); [Isaacman et al. 2011](#); [Calabrese et al. 2011](#); [Tizzoni et al. 2014](#); [Chen et al. 2014](#); [Phithakkitnukoon et al. 2012](#); [Csáji et al. 2013](#); [Kung et al. 2014](#)). We also used distilled criteria that were sometimes used independently (e.g., [Tizzoni et al. 2014](#)), sometimes combined (e.g., [Ahas et al. 2010](#)), sometimes within simple decision rules (e.g., [Phithakkitnukoon et al. 2012](#)), and sometimes within complex decision rules (e.g., [Csáji et al. 2013](#) and [Frias-Martinez 2010](#)).

The HDAs we construct use the decision rules that ‘home’ is in the area of the cell tower where:

1. The majority of both outgoing and incoming calls and texts were made (amount of activities criterion),
2. The maximum number of distinct days with phone activities – both outgoing and incoming calls and texts – was observed (amount of distinct days criterion),
3. Most phone activities were recorded during 7 p.m. and 9 a.m. (time constraints criterion),
4. Most phone activities were recorded, implementing a spatial perimeter of 1,000 meters around a cell tower that aggregates all activities within (space constraints criterion) and
5. The combination of 3) and 4), thus most phone activities recorded during 7 p.m. and 9 a.m. and implementing a spatial perimeter of 1,000 meter (time constraints and space constraint criterion).

Note that throughout this article, we will estimate cell tower areas by means of the Voronoi tessellation of the cell tower network. The use of Voronoi polygons to describe the spatial patterns of cell tower coverage has disadvantages. Although widely used in literature, Voronoi polygons are a simplification of the actual capacity of cell towers to cover areas. In reality, capacity is dependent on factors such as humidity, urban environment, elevation of the cell tower, and orientation. Theoretically, developing estimation models for the coverage of cell towers should be possible, but such models need extensive field surveys for validation, surveys that are expensive and thus rarely available. Therefore, there exists an unobservable measurement error when using Voronoi polygons and most findings (including ours) are dependent on the assumption that this error has an insignificant impact.

Note also that [Bojic et al. \(2015\)](#) uses similar HDAs when assessing and comparing home detection methods for a credit card transaction and Flickr data set. This shows that the relevance of these algorithms goes beyond the case of CDR data and also serves other data sets with non-continuous location traces.

4.2.2. Applying Five HDAs to the French CDR Dataset

We apply all five HDAs to the Orange France 2007 CDR data set to detect the cell tower that covers the presumed home location (L1) for all users during all months in the data set (May to October). Besides the L1 cell tower, we gather information about the second (L2) and the third (L3) most plausible cell tower to cover the home location following the particular decision rule applied.

[Table 1](#) shows the total number of times each HDA could detect an L1, L2 or L3 cell tower based on the CDR data of ~ 18 million users and when applied to each month in the

Table 1. Number of times (in millions) for which an L1, L2 or L3 cell tower could be detected from an individual user's CDR data by the various Home Detection Algorithms (HDAs) when applied per month in the data set. Percentages are column-wise and with respect to the number of LI detections.

Number of users (in million) with	Amount of activities (algorithm 1)	Amount of distinct days (algorithm 2)	Time restraints (amount of act.) (algorithm 3)	Space restraints (amount of act.) (algorithm 4)	Time and Space restraints (algorithm 5)
Detected L1	109.4 (100%)	109.4 (100%)	98.4 (100%)	109.4 (100%)	98.4 (100%)
Detected L2	102.2 (93.5%)	102.2 (93.5%)	78.0 (81.3%)	102.0 (92.8%)	78.4 (79.6%)
Detected L3	96.1 (87.9%)	96.1 (87.9%)	65.0 (66.1%)	66.1 (86.6%)	62.3 (63.3%)

data set. Given the availability of six different months (mid-May to mid-October), non-restrictive algorithms (such as algorithm 1 and 2) will be capable of detecting an L1 cell tower for about 109.4 million users ($\sim 18\,000\,000 \cdot 6$). Restrictive algorithms, such as the time-constraining algorithm 3, have fewer users for which a presumed home cell tower (L1) can be detected. The reason is that some users might not have made or received calls or texts during the restricted timeframe, so no CDR records exist and therefore the algorithm cannot identify an L1 cell tower.

For example, when we compare the number of times algorithm 1 (all activities) was capable of detecting an L1 compared to algorithm 3 (only nighttime activities), we can derive that up to 10% (98.4/109.4) of the users did not have mobile phone activities during the night. This made it impossible for the time-constraint HDAs to detect a cell tower presumably covering the home location. It is also interesting to note that, depending on decision rule of the algorithm, between 79.6 and 93.5% and between 62.3 and 87.9% of users have an L2 or L3 cell tower that could also be nominated as the home location cell tower, as they only varied by a slight degree compared to the L1 (or L2) cell tower(s). In other words, the decision rules applied do not overly discriminate between the eligibility of different cell towers to be the presumed home location. This raises the question of whether the French data set would not have benefited from a two-step approach.

4.3. Comparison of HDAs at Individual Level

One intriguing question is whether, for the same individual user, different HDAs would detect different home locations (L1 cell towers). We assess to which degree two different algorithms detect similar home locations for all individual users in the data set by evaluating the Simple Matching Coefficient (SMC) (Bojic et al. 2015):

$$\%SMC(\text{algorithm}_A, \text{algorithm}_B) = 100 \cdot \frac{\sum_{i=1}^N \delta(\text{Home}_{A,i}, \text{Home}_{B,i})}{N} \quad (1)$$

where $i = 1, \dots, N$ denotes the N users analysed, and $\delta(\text{Home}_{A,i}, \text{Home}_{B,i})$ is the Kronecker delta which is equal to 1 when the home detected by algorithm A for the i -th user is identical to the home detected by algorithm B for the same user. The Kronecker delta becomes 0 otherwise. Values of per cent SMC thus range between 0 and 100 and can be interpreted as the percentage of individual cases for which both algorithms detected the same home locations. When calculating SMC values, we omit all cases where one of the algorithms failed to detect a home location (e.g., when no observations were left after implementing a time constraint).

Figure 1 shows the SMC values for all pair combinations of HDAs during the different months in the data set. In general, pair accordance ranges between 61.5% and 96.4% of the detected homes, resulting in discordance rates between about 40% and 4%. In absolute numbers, this means that different decision rules predict different homes for between 6.8 and 0.6 million users. The patterns of (dis)similarities between HDAs are rather clear. Algorithms that incorporate time-constraints (algorithms 3 and 5) have a high degree of variance with algorithms that count the amount of activities (algorithm 1), distinct days (algorithm 2), or perform spatial groupings (algorithm 4), all of which show rather high degrees of pair accordance. The different results for the time-constraints algorithms might

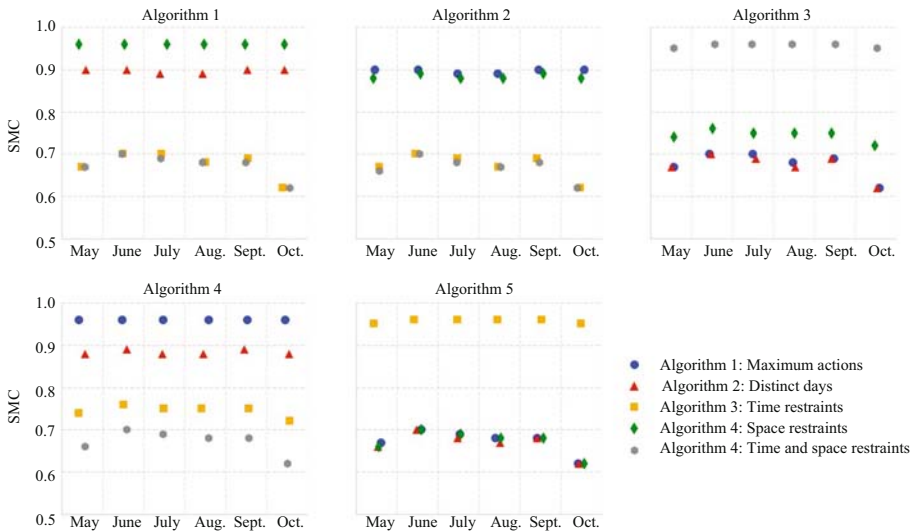


Fig. 1. SMC values for all pair combinations of HDAs, for each month in the data set. SMC values express the ratio of users for which two HDAs detect the same home.

stem from sparser observations or different movement patterns during the night, but exact reasons are unknown.

4.4. High Level Validation of Home Detection Algorithms

Given that different HDAs give different results for a considerable share of all individual users, the question becomes which decision rule should be preferred. As discussed previously, no consensus exists in literature on which decision rule(s) are best. This is partly because of the absence of comparative studies, but mainly because of the lack of proper validation data at individual level. In our case too, individual-level ground truth data was not available and so our assessment is at high-level, comparing census figures with population counts produced by HDAs.

4.4.1. National Statistics Validation Data Set

In contrast to related works, our high-level validation is based on a unique validation data set that was created in collaboration with the French National Statistics Institute (INSEE). To construct the validation data set, the Public Finances Directorate General (DGFiP) collected individual (or household) home locations from revenue declarations, housing taxes and the directory of taxable individuals. It then aggregated this information into population counts at the resolution of the Orange cell tower network (see also Figure 3a). In other words, an estimation of the population numbers, based on census data, for the geographical areas created by the Voronoi polygons of the Orange cell towers was produced and made available to the research project under a non-disclosure agreement.

It is a huge advantage to have access to a validation data set that has the same spatial resolution as the mobile phone network. It avoids the spatial translation of statistical zones

to the cell tower Voronoi areas, which is complicated and prone to errors (Frias-Martinez et al. 2010), given the spatially uneven distribution of cell towers.

Unfortunately, the individual (or household) home locations used to construct the validation data set could only be made available for the year 2010. However, for reasons explained in the previous paragraph we do opt to use this validation data set with its temporal mismatch (the mobile phone data set covers 2007) over the low resolution, publicly available census data that are updated every year. Since we only use the validation data set for relative comparisons between HDAs (i.e., no absolute validation is attempted), the assumption we introduce concerning this temporal mismatch is that relative population patterns do not change drastically within three years.

4.4.2. Validation of HDA Results at Cell Tower Level

To compare results from HDAs with the proposed validation data set, we evaluate the degree of similarity in population counts attributed to all cell tower areas. Note that we do not target an absolute assessment of similarity, as this is impossible given the unknown spatial distribution of the 28.7% sample of Orange users and the differences in times of collection between the CDR data set (2007) and our validation data (2010). Instead, we compare general patterns of estimated populations by means of vector comparison.

In our case, a first vector denotes the estimated population by one HDA for all cell tower areas and is compared to a second vector that describes the validation population count for exactly the same cell tower areas. Both vectors thus have an equal length ($n = 18\,273$ the amount of cell towers in the Orange network). To quantify the similarities and differences between both vectors, we use a standard Cosine Similarity Metric (CSM). According to Ye (2011, 91): “The cosine similarity is a classic measure used in information retrieval and is the most widely reported measure of vector similarity”, and it is based on the angle between two vectors described by its cosine:

$$\cos(\vec{x}, \vec{y}) = \frac{\sum_{i=1}^n x_i * y_i}{\sqrt{\sum_{i=1}^n x_i^2} * \sqrt{\sum_{i=1}^n y_i^2}} \quad (2)$$

where x_i and y_i are components of vectors \vec{x} and \vec{y} respectively and n is the total number of cell tower areas.

Values of the cosine will range between -1 and 1 . A value of 1 indicates the highest similarity in orientation (the angle between \vec{x} and \vec{y} is zero degrees), 0 indicates the lowest similarity in orientation (the angle between \vec{x} vector and \vec{y} vector is 90 or -90 degrees) and -1 indicates an opposite orientation (the angle between \vec{x} and \vec{y} is 180 degrees). Deriving the angle between two vectors and expressing it in degrees ($^\circ$) consequently gives us the CSM value we want:

$$\text{CSM}(\vec{x}, \vec{y}) = \left| \cos^{-1} \left(\frac{\sum_{i=1}^n x_i * y_i}{\sqrt{\sum_{i=1}^n x_i^2} * \sqrt{\sum_{i=1}^n y_i^2}} \right) * \frac{180}{\pi} \right| \quad (3)$$

A CSM value of 0° denotes the highest possible similarity, 90° indicates the lowest similarity in orientation whereas 180° degrees refers to an opposite orientation.

4.4.3. Validation with Census Data: CSM

Figure 2 shows the calculated CSM values for all HDAs and for different months. The distinct days that the algorithm performs best in replicating the population pattern of the validation data set, followed by the number of activities and the time-constrained number of activities. The HDAs that involve grouping in space perform worst, even though the applied perimeter (1 kilometre) in reality does not correspond to a substantial distance. It is worth noting that the performance of all HDAs range between 34° and 38° . This is substantially different from the intended 0° , which would signify a perfect match with the validation set. In other words, a ‘gap’ of about 35 degrees exists when using the CSM measure. This is indicative for the limited performance of our HDAs and raises the question of whether there is a structural limitation on the performance of single-step HDAs when applied to the French data set or to CDR data in general.

Interestingly, the performance of all HDAs is rather similar. Especially in their temporal patterns, where lower CSM values for June and September, and higher values for May, July, August and October are observed. A possible explanation for the high SMC values for May and October is the limited number of available days for these months in the data set (18 and 14 days respectively). This indicates that data should be collected for a certain duration for the HDA to perform properly.

The highest CSM values are observed during summer (July and August). All algorithms are sensitive to this temporal change, most likely because of the changing spatial behaviour of users who go on holiday (see also Deville et al. 2014; Vanhoof et al., 2017b). Time-limited criteria are more sensitive to temporal changes, which raises questions about their widespread adaptation in literature. In addition, it is interesting to note that differences between each algorithm are smaller than the differences of each algorithm

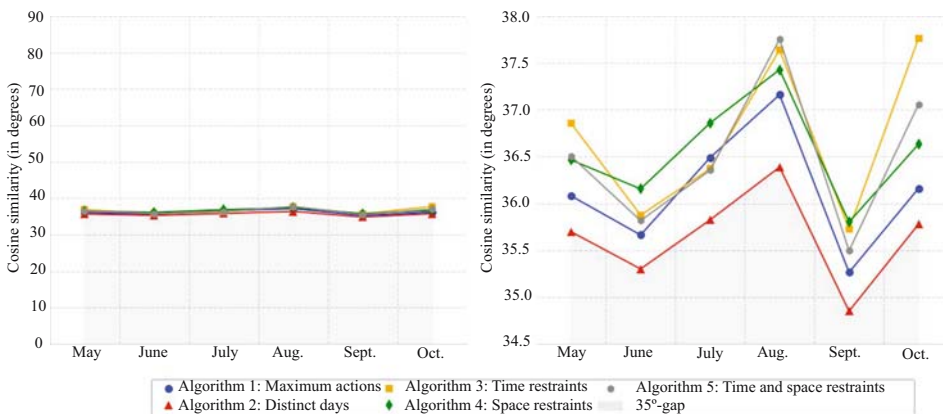


Fig. 2. CSM values (in degrees) of the comparison with ground truth data, for all HDAs applied to all months in the dataset. CSM values were calculated at cell tower level. The 35° gap is denoted as the difference between the best performing HDA and the expected CSM of 0° in the case of a perfect match between population counts from home location and the validation data set.

over time. Future analysis of HDA performance should therefore take into account the time period.

4.4.4. Spatial Patterns of Population Count

Although the CSM values for all HDAs are within a rather small range, it is important to realise that small differences in CSM values can imply major differences in the related spatial patterns of population counts.

Figures 3c and 3d for instance, show the spatial patterns of population counts obtained by the number of activity algorithms for June and August respectively. The difference in CSM values between both is a mere 1.08° , but their spatial pattern, as emphasised by the Getis-Ord G_i^* statistic (Getis and Ord 1992), is rather different. This statistic shows statistically significant clusters of high (hotspots) or low (coldspots) population counts. In August, for instance, the detected hotspots illustrate clear clusters of high numbers of home locations near sea and mountain areas. This is in contrast to an expected spatial pattern, where high clusters of population counts are found near cities and in urban areas, as can be seen from the spatial pattern of the validation data set in Figure 3b.

The spatial pattern of the differences between the validation data sets and detected homes in June and August are given in Figures 3e and 3f and visualise this contrast. Note that in Figure 3, the centre of Paris is often denoted as a coldspot because of the high density of cell towers, so each tower has a lower number of users, resulting in apparent coldspots. This effect is also visible in other city centres where cell tower density is high.

5. Discussion

5.1. Differences at Individual Level and the Absence of Ground Truth Data

Our results showed high discordance rates between different HDAs (ranging from 4% to 40% of the individual mobile phone users). This finding challenges the use of single-step home detection approaches for the French CDR data set when done without fully justifying the home location criteria used and the decision rules involved in the HDAs. As we argued, such justification is currently absent in the aforementioned literature, mainly because of the absence of ground truth data at individual level. Our case study clarifies how the absence of individual ground truth data necessitates a heuristic application of decision rules in current home detection methods. By this, we mean that one decision rule is applied to all users in the data set, regardless of the nature of their CDR traces. Obviously, the better approach would be to have non-generic algorithms that could flexibly select decision rules (and validation) based on the characteristics of individual user traces. Such a solution, however, would require large training samples (individual ground truth) to learn how to switch between different decision rules. As yet, these are not available.

5.2. Sensitivity of Performance Considering Time and Decision Rule Choice

Performing high-level validation on five HDAs, by comparing population counts with census data, unveiled rather poor performances (CSM values between 34° and 38°), and a clear sensitivity to the chosen time period. In fact, for the French data set, defining a time

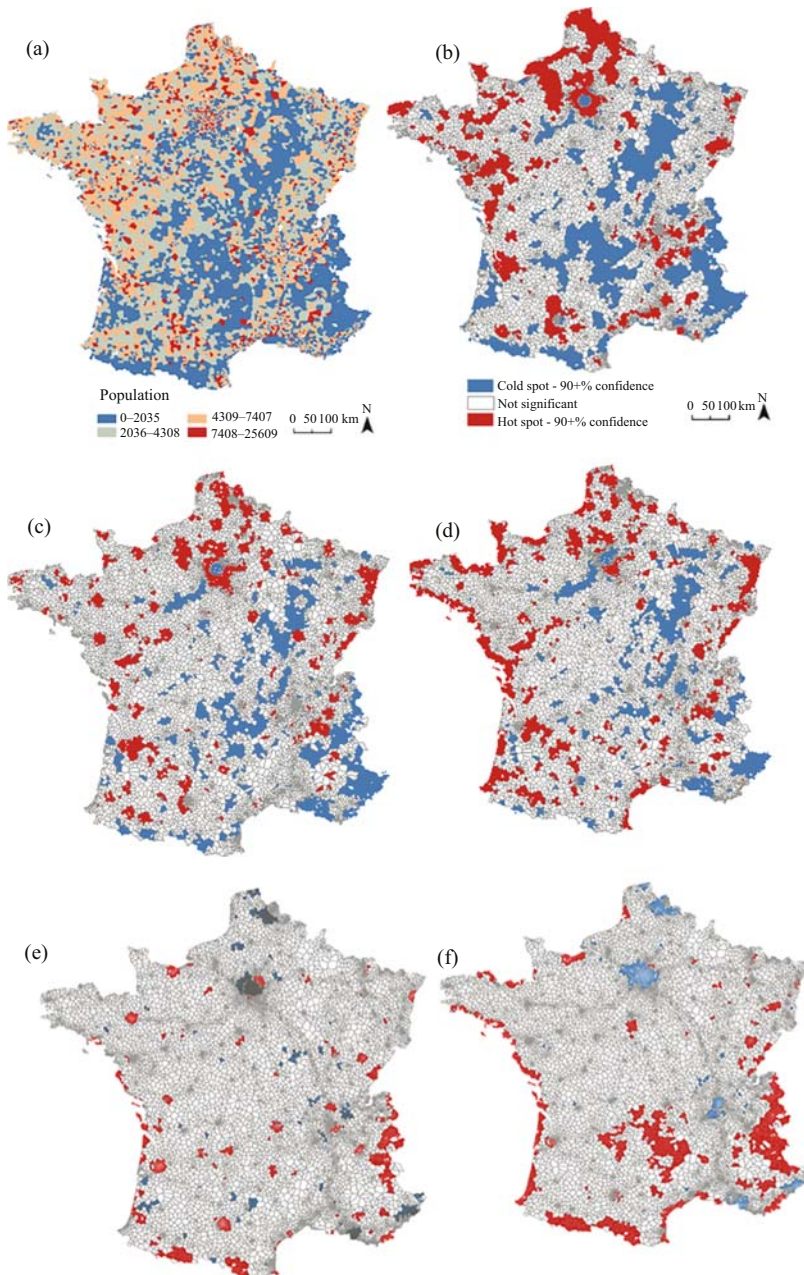


Fig. 3. Population counts of the validation data set: (a) Hotspots (red) and coldspots (blue) as defined by the 90+% interval of the Getis-Ord G_i^* statistic for the population counts of the validation data set (b); for the number of detected homes using the amount of activities algorithm in June (c); for the number of detected homes using the amount of activities algorithm in August (d); for the log(ratio) between the amount of activities algorithm in June and the population counts of the validation data set (e); and for the log(ratio) between the amount of activities algorithm in August and the population counts of the validation data set (f). All maps are compiled from Voronoi tessellation of the Orange cell towers. Figures b, c, d, e, and f share the legend of Figure 2.

period to carry out home detection seems as important as criteria choice. As was illustrated by the spatial pattern of population counts in [Figure 3d](#), and by increasing CSM values in July and August in [Figure 2](#), summer periods should be avoided when running HDAs. Additionally, shorter observation periods (like May and October in our case) also seem to influence the performance of HDAs.

When comparing criteria, it is clear that the space constraints criterion is outperformed by all others. The main logic behind grouping close locations together (in this case, within a 1 km perimeter) is to avoid frequent handovers between close cell towers. However, on a large scale, such a precaution, seems to have a negative impact. Furthermore, the extreme volatility of the performance of the time constraints criteria is remarkable. Clearly, this criterion is not able to cope with (changing) user behaviour during the summer months, resulting in the worst performance obtained.

5.3. The 35°-gap in High-Level Validation

The most telling result of our analysis, is that all tested HDAs have CSM values which are still far off from the intended 0°, as can be observed in [Figure 2](#). Additionally, it is remarkable that CSM values for all HDAs occur in the same, rather small, range of CSM values (even though a small difference in CSM can induce a rather profound change in spatial patterns). The 35°-gap observed is indicative for the (current) limits of single-step approaches based on simple decision rules, at least at cell tower level (as aggregation to higher levels might diminish the gap considerably).

The 35°-gap also adds to the discussion of high-level validation. As mentioned before, the absence of individual-level validation is hindering a clear understanding of why the performance between algorithms may differ. Playing devil's advocate, one couldn't care less about individual correctness, as long as the statistical performance at nationwide level is sufficient. However, given the considerable differences between census and mobile phone home location data at cell tower level, as revealed by both the 35°-gap, in [Figure 2](#), and the clear differences in spatial patterns in [Figure 3](#), it seems inevitable that investigations at individual or subset level will need to be undertaken to improve insights into the workings of HDA and, ultimately, the performance of home detection methods in general.

It is clear that the 35°-gap requires further exploration so as to understand its constituent parts. We consider, at least, the following elements to be of importance:

- *Distorted local market shares*: Local market shares for individual mobile phone operators can be highly volatile and are often unknown. This causes a lot of uncertainty when it comes to high-level validation with census data, as the percentage of the population that the operator actually captures in different regions is unclear. Unknown local market shares therefore impede both validation techniques that perform in pairs and/or absolute comparisons between population estimations and ground truth figures collected by nationwide censuses. They also most probably hinder validation techniques that are based on relative differences (like the CSM metrics).
- *Diversity of mobile phone use*: Differences in mobile phone use between users and/or regions can structurally influence the validation of single-step HDAs. When concentrated, differences in mobile phone use influence high-level validation in the

same way as distorted market shares would. Additionally, it is clear that mobile phone usage changes rapidly over time. It can be argued that the use of phones for professional or private purposes was different in 2007 than it is today. Unfortunately, such usage contexts are not available in CDR data. Neither can they be easily derived since, in general, privacy regulations ban the linking of CDR data and customer databases that gather for example, billing addresses or type of payment information. In other words, traces in CDR data will be of a different nature at different times because of differences in mobile phone usage. This implies that information on mobile phone usage is necessary to understand the effect on home detection performance.

- *Differing definitions of home*: Differences between the definition of home in census data and the definition of home by HDAs may cause structural discordance when validating the latter by the former. Even though official statistical practices have a tradition of distinguishing between different definitions of home, such as ‘usual resident population’ and ‘second home population’, it remains unclear to what degree mobile phone data is capable of capturing such concepts of home and to what degree different decision rules would favour the detection of different types of homes.
- *Technical aspects of the data collection and methodology*: Research has paid wide attention to the technical aspects of mobile phone data, especially when it comes to the estimation of cell tower areas and their translation into statistical areas (Ricciato et al. 2015). In our case, we avoided the translation problem by constructing a validation data set at cell tower level, but for many cases this is not an option. Estimation of cell tower areas was done by Voronoi polygons, which introduces errors at a local scale, but could also structurally influence high-level validation if, for example, areas of cell towers in high population density locations were consistently underestimated. Here too, the effects on high-level validation practices are currently unknown, but we expect them to be minor compared to previous points.

For all the points raised above, no quantification of their effect(s) has yet been explored. Additionally, it is worth noting that some of these points may become more or less relevant over time due to, factors such as technical advancements or regulations. The EU General Data Protection Directive 2018, for example, will probably make it harder to work at the individual level (individual mobile phone use, different types of home definition). This will make high-level validation techniques more relevant, and thus increase the need to have proper knowledge of local market shares. For this reason, it is difficult to assess the relevance of the findings given in this article. However, we strongly suggest that all are the topics of future research and consideration.

6. Recommendations

Throughout this work, we have given suggestions concerning the use of HDAs for mobile phone data. In summary, we believe we can compile these suggestions into a set of three recommendations, which are relevant at different levels.

1. **Individual level**: Currently, the biggest problem in ensuring the reliable use of HDAs for mobile phone data (and, in extension, other similar data sources like

location based services or geotagged online social networks) is the absence of ground truth data at the individual level. We strongly recommend the collection of ground truth data linking mobile phone usage, the related CDR data, and movement patterns of individual users. Even if collected for only small samples of users, this step is essential for giving proper estimations of error and performance of HDAs at individual level. It also would help in understanding the differences between decision rules, and plausibly allow for the development of non-generic HDAs that could switch between decision rules based on the characteristics of individual traces, instead of being generically applied to all users. Additionally, the availability of individual level ground truth could shed light on the structural effects that currently obscure high-level validation practices, such as the changing usage of mobile phones and the differences between declared homes and lived-in homes as captured by census and mobile phone data respectively.

2. **National level:** Apart from ensuring nationally representative sampling of individual level ground truth data, we believe it to be important either to understand local market shares of single operators, or to collect mobile phone data from all operators in the territory. Without this information, high-level validation of population estimations at nationwide level will remain flawed, making it impossible to describe correctly the performance of HDAs at a larger geographical scale. In addition, resolving the local market share issue is a crucial step in the investigation of the (spatial) representativity of available mobile phone data sets, as unknown market shares at local level impede the analysis of subset populations in data sets.
3. **International level:** Finally, we believe that one of the key components to ensure reliable use of mobile phone data in official statistics is the opportunity to test ideas and methodologies on different data sets, which contain differing populations and cover various time periods. This is not necessarily a matter of testing for uniformity. On the contrary, it is a matter of understanding the limits of current methodologies, assessing the true potential for applications and anticipating the wider challenges posed by fast-evolving technology usage and deployment. All of these factors are necessary to ensure the future applicability of mobile phone data sources in official statistics.

As we reflect on the direction that further investigation should take, together with the feasibility of carrying out the recommendations proposed, we realize that this is a larger intervention than any single researcher, research group, national statistics office or even operator can be expected to take. Therefore, it is encouraging to see that collaborations are being formed to address different parts of the problem.

In France, for example, a collaboration between the operator Orange and the national statistics office INSEE is investigating different aspects of the high-level validation of home detection practices, such as translating Voronoi polygons into existing statistical grids ([Sakarovitch et al. in preparation.](#)).

On a European scale, the ESSnet Big Data project has been organising the exchange of best practices for the integration of mobile phone data (and multiple other big data sources) in official statistics. Its goal is directly in line with the recommendations previously described (especially recommendation 3: international level), facilitating the

uniformity of quality and methodologies for the use of big data sources in European official statistics (ESSnet Big Data 2018).

As a last example, the Open Algorithm project (OPAL) is a collaboration between operators, academia, and institutional partners who are building a platform to allow the use of *Open Algorithms* on mobile phone data sets from different operators (OPAL 2018). The idea is that users can launch a predefined set of algorithms (such as home detection algorithms), which are then run behind the firewalls of the operators before returning the aggregated results back to the user. Although the project is currently still in its test phase (with pilots in Senegal and Colombia), hopes are that it could facilitate cooperation between different operators in sharing basic statistical information from their data sets (as captured by the predefined set of algorithms). If all of a country's mobile phone operators would engage in this form of cooperation, the problem of dealing with a distorted market share, for example, would be solved.

Hence, the bottom line is that although the home location problem is mainly a methodological one, the paths to address the problem are much more complex. They require the combination of collaborative, technical, methodological, institutional and strategic actions. Optimistically, we believe that official statistics offices are in a good position to (continue to) play a prominent role, because of their organisational structure, methodological knowledge and recognised institutional role within a country.

7. Conclusion

Big data sources in general, and mobile phone data in particular, create intriguing new opportunities and challenges for official statistics. Because of this, there has been a clear call for exploratory pilot projects to be carried out, as well as a trend towards critical investigation and transparency of methodologies to produce high-quality statistics. This article adhered to both of these calls in its analysis of home detection practices for non-continuous location traces, focusing mainly on mobile phone data.

Based on a critical review of literature, we discussed how existing methods to identify home locations using non-continuous location traces mainly consist of single-step approaches that deploy simple decision rules and use high-level validation only.

We argued that, given the absence of ground truth data at individual level, i) it is unclear why one-step approaches are preferred over two-step approaches that are typically used for continuous location traces; ii) no consensus in literature exists on which criteria are best to deploy when creating decision rules for home detection methods, nor has work been done to investigate the sensitivity of the results to these decision rules and criteria; and iii) the trustworthiness of high-level validation and its added value to the home detection practices are questionable at best.

By deploying five algorithms with simple decision rules to a large French CDR data set, we demonstrated several of the problems. At individual level, we found home detection methods to be rather sensitive to criteria choice, with pair comparison of different home detection algorithms resulting in different identified homes for up to 40% of users. When looking at high-level validation, we found that five different home detection algorithms performed in a similar range (34°–38°) with a similar sensitivity to the time period and the duration for which the mobile phone data was collected. Even though we found that the

sensitivity to time and the differences between different HDA algorithms does not seem large when expressed in CSM values, we showed how small changes to CSM values influence substantive and nationwide changes in the spatial patterns of population counts.

Our most noteworthy finding is the magnitude of the mismatch (the 35°-gap) between population counts constructed from mobile phone-based data on home location and a validation data set based on census data. This large mismatch is indicative of the severity of the home location problem and challenges the validity of single-step approaches in literature. In our discussion, we listed several elements that plausibly effect this mismatch but go unnoticed when only high-level validation is undertaken. We believe that these (structural) elements, such as unknown market shares and differences in mobile phone usage, need further investigation if ever home detection methodologies are to comply with official statistics' standards.

Finally, we compiled our findings, insights, and experiences into a set of specific recommendations, ranging from the collection of individual ground truth data to the testing of methods on multiple data sets. Given the nature of these recommendations and the tasks at hand, we think that it is unlikely that individual researchers, research groups, national statistics offices, or even mobile phone operators can, or will, invest in them. Therefore, we call on and support any ongoing, collaborative actions that tackle these problems, while recognising the prominent role official statistics can (continue to) play in this area.

8. References

- Ahas, R., A. Aasa, A. Roose, Ü. Mark, and S. Silm. 2008. "Evaluating Passive Mobile Positioning Data for Tourism Surveys: An Estonian Case Study." *Tourism Management* 29(3): 469–486. Doi: <http://dx.doi.org/10.1016/j.tourman.2007.05.014>.
- Ahas, R., S. Silm, O. Järvi, E. Saluveer, and M. Tiru. 2010. "Using Mobile Positioning Data to Model Locations Meaningful to Users of Mobile Phones." *Journal of Urban Technology* 17(1): 3–27. Doi: <http://dx.doi.org/10.1080/10630731003597306>.
- ARCEP. 2008. "Le Suivi Des Indicateurs Mobiles – Les Chiffres Au 31 Décembre 2007 (Publication Le 4 Février 2008)." Available at: <http://www.arcep.fr/index.php?id=9545> (Last accessed February 2018).
- Ashbrook, D. and T. Starner. 2003. "Using GPS to Learn Significant Locations and Predict Movement across Multiple Users." *Personal and Ubiquitous Computing* 7(5): 275–286. Doi: <http://dx.doi.org/10.1007/s00779-003-0240-0>.
- Baldacci, E., D. Buono, G. Kapetanios, S. Krische, M. Marcellino, G. Mazzi, and F. Papailias. 2016. "Big Data and Macroeconomic Nowcasting: From Data Access to Modelling." Luxembourg: Eurostat. Doi: <http://dx.doi.org/10.2785/360587>.
- Blondel, V.D., M. Esch, C. Chan, F. Clerot, P. Deville, E. Huens, F. Morlot, Z. Smoreda, and C. Ziemlicki. 2012. "Data for Development: The D4D Challenge on Mobile Phone Data." *arXiv:1210.0137*. Available at: <http://arxiv.org/abs/1210.0137> (Last accessed February 2018).
- Blondel, V.D., A. Decuyper, and G. Krings. 2015. "A Survey of Results on Mobile Phone Datasets Analysis." *EPJ Data Science* 4(10): 1–57. Doi: <http://dx.doi.org/10.1140/epjds/s13688-015-0046-0>.

- Blumenstock, J.E. 2012. “Inferring Patterns of Internal Migration from Mobile Phone Call Records: Evidence from Rwanda.” *Information Technology for Development* 18(2): 107–125. Doi: <http://dx.doi.org/10.1080/02681102.2011.643209>.
- Bojic, I., E. Massaro, A. Belyi, S. Sobolevsky, and C. Ratti. 2015. “Choosing the Right Home Location Definition Method for the given Dataset.” In *Social Informatics – 7th International Conference, SocInfo 2015, Beijing, China, December 9–12, 2015, Proceedings*, edited by Tie-Yan Liu, Christie Napa Scollon, and Wenwu Zhu, 9471: 194–208. Beijing: Springer. Doi: http://dx.doi.org/10.1007/978-3-319-27433-1_14.
- Calabrese, F., L. Ferrari, and V.D. Blondel. 2014. “Urban Sensing Using Mobile Phone Network Data: A Survey of Research.” *ACM Computing Surveys* 47(2): 1–20. Doi: <http://dx.doi.org/10.1145/2655691>.
- Calabrese, F., G. Di Lorenzo, L. Liu, and C. Ratti. 2011. “Estimating Origin-Destination Flows Using Mobile Phone Location Data.” *IEEE Pervasive Computing* 10(4): 36–44. Doi: <http://dx.doi.org/10.1109/MPRV.2011.41>.
- Chen, C., L. Bian, and J. Ma. 2014. “From Traces to Trajectories: How Well Can We Guess Activity Locations from Mobile Phone Traces?” *Transportation Research Part C: Emerging Technologies* 46: 326–337. Doi: <http://dx.doi.org/10.1016/j.trc.2014.07.001>.
- Csáji, B.C., A. Browet, V.A. Traag, J.C. Delvenne, E. Huens, P. Van Dooren, Z. Smoreda, and V.D. Blondel. 2013. “Exploring the Mobility of Mobile Phone Users.” *Physica A: Statistical Mechanics and Its Applications* 392(6): 1459–1473. Doi: <http://dx.doi.org/10.1016/j.physa.2012.11.040>.
- Daas, P.J.H., M.J. Puts, B. Buelens, and P.A.M. van den Hurk. 2015. “Big Data as a Source for Official Statistics.” *Journal of Official Statistics* 31(2): 249–262. Doi: <http://dx.doi.org/10.1515/JOS-2015-0016>.
- de Montjoye, Y.-A., Z. Smoreda, R. Trinquart, C. Ziemlicki, and V.D. Blondel. 2014. “D4D-Senegal: The Second Mobile Phone Data for Development Challenge.” *arXiv:1407.4885*. Available at: <http://arxiv.org/abs/1407.4885> (Last accessed February 2018).
- Deville, P., C. Linard, S. Martin, M. Gilbert, F.R. Stevens, A.E. Gaughan, V.D. Blondel, and A.J. Tatem. 2014. “Dynamic Population Mapping Using Mobile Phone Data.” *Proceedings of the National Academy of Sciences* 111(45): 15888–15893. Doi: <http://dx.doi.org/10.1073/pnas.1408439111>.
- Eurostat. 2014. “ESS Big Data Action Plan and Roadmap 1.0. Approved by the 22nd Meeting of the European Statistical System Committee.” Available at: https://ec.europa.eu/eurostat/cros/content/ess-Big-DataAction-Plan-and-Roadmap-10_en (Last accessed February 2018).
- ESSnet Big Data. 2018. “ESSnet Big Data.” Available at: https://webgate.ec.europa.eu/fpfis/mwikis/essnetbigdata/index.php/ESSnet_Big_Data (Last accessed February 2018).
- Frias-Martinez, V. and J. Virseda. 2012. “On the Relationship between Socio-Economic Factors and Cell Phone Usage.” In *Proceedings of the Fifth International Conference on Information and Communication Technologies and Development*, March 12–15, 2015. 76–84. New York, NY: ACM Press. Doi: <http://dx.doi.org/10.1145/2160673.2160684>.
- Frias-Martinez, V., V. Jesus, A. Rubio, and E. Frias-Martinez. 2010. “Towards Large Scale Technology Impact Analyses.” In *Proceedings of the 4th ACM/IEEE*

- International Conference on Information and Communication Technologies and Development – ICTD '10, December 13–16, 2010. 1–10. New York, NY: ACM Press. Doi: <http://dx.doi.org/10.1145/2369220.2369230>.
- Getis, A. and J.K. Ord. 1992. “The Analysis of Spatial Association by Use of Distance Statistics.” *Geographical Analysis* 24(3): 189–206. Doi: <http://dx.doi.org/10.1111/j.1538-4632.1992.tb00261.x>.
- Giannotti, F., D. Pedreschi, A. Pentland, P. Lukowicz, D. Kossmann, J. Crowley, and D. Helbing. 2012. “A Planetary Nervous System for Social Mining and Collective Awareness.” *The European Physical Journal Special Topics* 214(1): 49–75. Doi: <http://dx.doi.org/10.1140/epjst/e2012-01688-9>.
- Glasson, M., J. Trepanier, V. Patruno, P. Daas, M. Skaliotis, and A. Khan. 2013. “What Does ‘Big Data’ Mean for Official Statistics?” Available at: <http://www1.unece.org/stat/platform/pages/viewpage.action?pageId=77170614> (Last accessed February 2018).
- Gong, H., C. Chen, E. Bialostozky, and C.T. Lawson. 2012. “A GPS/GIS Method for Travel Mode Detection in New York City.” *Computers, Environment and Urban Systems* 36(2): 131–139. Doi: <http://dx.doi.org/10.1016/j.compenvurbsys.2011.05.003>.
- Grauwin, S., M. Szell, S. Sobolevsky, P. Hövel, F. Simini, M. Vanhoof, Z. Smoreda, A.-L. Barabasi, and C. Ratti. 2017. “Identifying and Modelling the Structural Discontinuities of Human Interactions.” *Scientific Reports* 7: 46677. Doi: <http://dx.doi.org/10.1038/srep46677>.
- Hightower, J., S. Consolvo, A. LaMarca, I. Smith, and J. Hughes. 2005. “Learning and Recognizing the Places We Go.” In *UbiComp 2005: Ubiquitous Computing*, edited by M. Beigl, S. Intille, J. Rekimoto, and H. Tokuda, 159–176. Berlin, Heidelberg: Springer Berlin Heidelberg. Doi: <http://dx.doi.org/10.1007/11551201>.
- Isaacman, S., R. Becker, R. Caceres, S. Kobourov, M. Martonosi, J. Rowland, and A. Varshavsky. 2011. “Identifying Important Places in People’s Lives from Cellular Network Data.” In *Pervasive Computing: Pervasive 2011*, edited by K. Lyons, J. Hightower, and E.M. Huang, 133–151. Berlin, Heidelberg: Springer Berlin Heidelberg. Doi: http://dx.doi.org/10.1007/978-3-642-21726-5_9.
- Janzen, M., M. Vanhoof, and K.W. Axhausen. 2016. “Estimating Long-Distance Travel Demand with Mobile Phone Billing Data.” In Proceedings of the 16th Swiss Transport Research Conference (STRC 2016), May 18–20, 2016. Available at: http://www.strc.ch/conferences/2016/Janzen_EtAl.pdf (Last accessed February 2018).
- Janzen, M., M. Vanhoof, Z. Smoreda, and K.W. Axhausen. 2018. “Closer to the Total? Long-Distance Travel of French Mobile Phone Users.” *Travel Behaviour and Society* 11: 31–42. Doi: <http://dx.doi.org/10.1016/j.tbs.2017.21.001>.
- Järv, O., R. Ahas, and F. Witlox. 2014. “Understanding Monthly Variability in Human Activity Spaces: A Twelve Month Study Using Mobile Phone Call Detail Records.” *Transportation Research Part C: Emerging Technologies* 38: 122–135. Doi: <http://dx.doi.org/10.1016/j.trc.2013.11.003>.
- Karlberg, M., S. Biffignandi, P.J.H. Daas, A. Holmberg, B. Hulliger, P. Jacques, R. Lehtonen, R.T. Münnich, N. Shlomo, R. Silberman, and I. Stoop. 2015. “Preface.” *Journal of Official Statistics* 31(2): 149–153. Doi: <http://dx.doi.org/10.1515/jos-2015-0011>.

- Kung, K.S., K. Greco, S. Sobolevsky, and C. Ratti. 2014. “Exploring Universal Patterns in Human Home-Work Commuting from Mobile Phone Data.” *PLoS ONE* 9(6): e96180. Doi: <http://dx.doi.org/10.1371/journal.pone.0096180>.
- Marchetti, S., C. Giusti, M. Pratesi, N. Salvati, F. Giannotti, D. Pedreschi, S. Rinzivillo, L. Pappalardo, and L. Gabrielli. 2015. “Small Area Model-Based Estimators Using Big Data Sources.” *Journal of Official Statistics* 31(2): 263–281. Doi: <http://dx.doi.org/10.1515/JOS-2015-0017>.
- Nurmi, P. and S. Bhattacharya. 2008. “Identifying Meaningful Places: The Non-Parametric Way.” In *Pervasive Computing*, edited by J. Indulska, D. Patterson, J. Rodden, and M. Ott, 111–127. Berlin: Springer Berlin.
- OPAL. 2018. “The OPAL project.” Available at: <http://www.opalproject.org/> (Last accessed February 2018).
- Pappalardo, L., M. Vanhoof, L. Gabrielli, Z. Smoreda, D. Pedreschi, and F. Giannotti. 2016. “An Analytical Framework to Nowcast Well-Being Using Mobile Phone Data.” *International Journal of Data Science and Analytics* 2(1–2): 75–92. Doi: <http://dx.doi.org/10.1007/s41060-016-0013-2>.
- Phithakkitnukoon, S., Z. Smoreda, and P. Olivier. 2012. “Socio-Geography of Human Mobility: A Study Using Longitudinal Mobile Phone Data.” *PloS One* 7(6): e39253. Doi: <http://dx.doi.org/10.1371/journal.pone.0039253>.
- Raun, J., R. Ahas, and M. Tiru. 2016. “Measuring Tourism Destinations Using Mobile Tracking Data.” *Tourism Management* 57: 202–212. Doi: <http://dx.doi.org/10.1016/j.tourman.2016.06.006>.
- Ricciato, F., P. Widhalm, M. Craglia, and F. Pantisano. 2015. “Estimating Population Density Distribution from Network-Based Mobile Phone Data.” Luxembourg: Publications Office of the European Union. Doi: <http://dx.doi.org/10.2788/162414>.
- Rubrichi, S., Z. Smoreda, and M. Musolesi. 2017. “A Comparison of Spatial-Based Targeted Disease Containment Strategies Using Mobile Phone Data.” *arXiv:1210.0137*. Available at <https://arxiv.org/pdf/1706.00690.pdf> (Last accessed February 2018).
- Sakarovitch, B., P. Givord, M.-P. de Bellefon, and M. Vanhoof. In Preparation. “Allô t’es où ? Estimer la population résidente à partir de données de téléphonie mobile, une première exploration.” *Economie et Statistique/Economics and Statistics*. (Preprint available upon request to authors).
- Shen, L. and P.R. Stopher. 2014. “Review of GPS Travel Survey and GPS Data-Processing Methods.” *Transport Reviews* 34(3): 316–334. Doi: <http://dx.doi.org/10.1080/01441647.2014.903530>.
- Sobolevsky, S., M. Szell, R. Campari, T. Couronné, Z. Smoreda, and C. Ratti. 2013. “Delineating Geographical Regions with Networks of Human Interactions in an Extensive Set of Countries.” *PloS One* 8(12): e81707. Doi: <http://dx.doi.org/10.1371/journal.pone.0081707>.
- Tizzoni, M., P. Bajardi, A. Decuyper, G.K.K. King, C.M. Schneider, V. Blondel, Z. Smoreda, M.C. Gonzalez, and V. Colizza. 2014. “On the Use of Human Mobility Proxies for Modeling Epidemics.” *PLoS Computational Biology* 10(7): e1003716. Doi: <http://dx.doi.org/10.1371/journal.pcbi.1003716>.
- Vanhoof, M., S. Combes, and M.-P. de Bellefon. 2017a. “Mining Mobile Phone Data to Detect Urban Areas.” In *Proceedings of the Conference of the Italian Statistical Society*

- (*SIS*), edited by A. Petrucci and R. Verde, 1005–1012. Florence: Firenze University Press. ISBN (online) 978-88-6453-521-0.
- Vanhoof, M., L. Hendrickx, A. Puusaar, G. Verstraeten, T. Ploetz, and Z. Smoreda. 2017b. “Exploring the Use of Mobile Phones during Domestic Tourism Trips.” *Netcom* 31(3/4): 335–372.
- Vanhoof, M., W. Schoors, A. Van Rompaey, T. Ploetz, and Z. Smoreda. 2018. “Correcting Mobility Entropy for Regional Comparison of Individual Movement Patterns.” *Journal of Urban Technology* 25(2): 27–61. Doi: <http://dx.doi.org/10.1080/10630732.2018.1450593>.
- Wolf, J., R. Guensler, and W. Bachman. 2001. “Elimination of the Travel Diary: Experiment to Derive Trip Purpose from GPS Travel Data.” *Transportation Research Record* 1768: 125–134. Doi: <http://dx.doi.org/10.3141/1768-15>.
- Ye, J. 2011. “Cosine Similarity Measures for Intuitionistic Fuzzy Sets and their Applications.” *Mathematical and Computer Modelling* 53(1–2): 91–97. Doi: <http://dx.doi.org/10.1016/j.mcm.2010.07.022>.

Received June 2017

Revised March 2018

Accepted April 2018