Robustness Analyses for Repeated Mobility Surveys in Outdoor Advertising

Dirk Hecker[1], Christine Körner[2], Michael May[3]

[1, 2, 3] *Fraunhofer IAIS*

Schloss Birlinghoven, 53754 Sankt Augustin, Germany

[1]dirk.hecker@iais.fraunhofer.de

[2]christine.koerner@iais.fraunhofer.de

[3]michael.may@iais.fraunhofer.de

***Abstract*— A growing number of companies use mobility data in their day-to-day business. However, as the data grows older, new data has to be collected in order to keep applications up-to-date. Consequently, it is of great importance to know the impact that a different mobility sample may cause. This aspect of analysis has been largely neglected in mobility data mining research so far. In this paper we therefore analyze the robustness of performance measures with respect to a changed GPS sample in outdoor advertisement. The evaluation of outdoor advertising campaigns is a challenging application because it requires the evaluation of mobility data on a very fine spatial level. Thus, the application has a higher dependency on routes of individual test persons than classical mobility surveys. In our robustness analysis we apply bootstrapping and subsampling in order to measure the effect of a) a repeated mobility survey and b) a mobility survey of smaller size. We conduct our experiments on a real-world data set from Swiss outdoor advertising. Our results show that the effect is comparably small for a typical campaign and may be mitigated further by increasing the campaign size.**

***Keywords*— mobility mining, robustness analysis, sampling, GPS, outdoor advertising, bootstrap, standard error**

I. Introduction

During the past years the interest in the exploitation of mobility information has increased significantly. Algorithms have been developed e.g. for the clustering of trajectories, detection of relative motion patterns or anonymization of movement data. However, all these approaches are restricted to a given trajectory sample. As tracking technologies have become more easily available and less expensive over the past years, mobility surveys can be repeated over time. Therefore, the question arises how stable measurement results and derived variables are when mobility surveys are extended or replaced after some time. What happens to an application that relies on mobility data when measurements are repeated several years later in the same geographic region, with a different mobility sample? Can we estimate the robustness of results in advance? Can we learn from a given GPS-Survey which variance to expect due to a different, possibly smaller sample in the upcoming survey?

These are critical questions for many applications. In this paper we analyze how a substitution of an old GPS survey and a possible reduction of the sample size affect performance measures in outdoor advertisement. Our analysis is based on a real-world business application for the Swiss Poster Research Plus (SPR+). SPR+ conducted a GPS mobility survey with over 11.000 test persons for 7-10 days in major Swiss conurbations over the past years. The data is used to determine the number and (spatial) distribution of poster passages. The usage of GPS data has the advantage that poster performance can be differentiated with respect to the location of poster sites as well as the socio-demography and origin of target groups. In this paper we focus on the performance measure reach, which is one of the most important measures in the advertising branch. Reach describes the percentage of population that passes at least one poster of the campaign within a given period of time.

As the data grows older, new data has to be collected in order to keep performance measures up-to-date. However, as poster performance has a direct influence on the pricing of a site, the update process is a critical task. It is therefore of great importance to know in advance which impact a changed GPS sample may cause. In addition, as the first GPS survey has been comparably large, a reduction of the sample size may be necessary to decrease survey costs. Therefore, we are also interested in the influence that a reduced sample size has on poster performance.

In this paper we estimate the robustness of poster reach for a given sample size by bootstrapping and determine the impact of prospective smaller samples by subsampling with different sizes. In addition, we vary the size of the poster campaigns in our experiments because it is an important variable for the stability of performance measures.

We believe that this type of analysis is relevant for a number of applications where GPS samples have to be replaced after a given period of time as it helps to understand the flexibility and limitations of a given GPS sample.

The remaining of this paper is organized as follows. In the next section we discuss related work. Section III introduces our analysis approach. Section IV applies the approach to the Swiss mobility survey and Section VI concludes our work.

II. Related Work

A. Estimation of Performance Measures in Outdoor Advertising

Outdoor advertisement is one of the oldest advertising media and continues to play an important role in the advertisement industry. In 2008 the turnover was 684 million CHF (about 460 million Euros) in Switzerland. In recent years the market has changed rapidly. The change is predominately caused by two factors, namely the competition with other advertising media and the emergence of digital media. First, outdoor advertisement competes with other media including the classic television, radio and press as well as the modern online ads and direct mailing. In order to become incorporated by media planners in an advertisement mix, transparent measures are needed for the performance of a campaign. Typical measures are (1) the reach of a campaign, which is the percentage of persons within a target group defined by socio-demographic attributes that has had contact with a campaign in a certain time interval (often one week), and (2) the number of total contacts this group has had [1]. GPS technology has established itself as a new standard in Switzerland and Germany, greatly improving the possibilities of fine-grained media planning. Other countries, such as Austria and the UK, are currently preparing GPS studies, and it can be expected to become a worldwide standard. Fig. 1 shows the GPS trajectories of 1,956 test persons in Zurich, which we use for our experiments.

Fig. 1 GPS trajectories of test persons in Zurich

Given the trajectories of all test persons and the visibility areas of the poster panels, the resulting passages can be calculated by geographic intersection (Fig. 2).

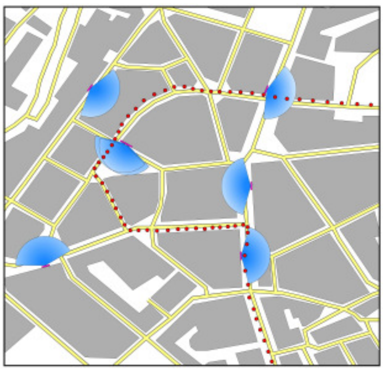

Fig. 2 Geographic intersection of poster sites and trajectories

Afterwards, the passages are weighted with individual contact factors of the poster sites, which results in the so called *visibility adjusted contacts* (VAC). The contact factors weight a passage according to several criteria, including speed, size, exposure time and angle of the panel. For example, if people pass a poster site by car they usually pass it faster than if they walk or cycle. The resulting VAC are then used to calculate poster performance measures.

Note that in this paper we perform our experiments based on poster passages instead of VAC because we would like to obtain a better understanding of the spatio-temporal influence of different trajectories in general.

A. Mobility Data Analysis

Mobility data analysis has been conducted for a long time in the area of travel surveys. In recent years also many data mining algorithms have been developed for mobility data. In this section we give a short description of both research direction and delineate our work from existing approaches.

The main focus of travel surveys is the analysis of general movement characteristics on a regional or national level. They are typically repeated after a specified period of time to monitor urban or national changes in mobile behavior. One example is the study Mobility in Germany [2] which is based on a computer assisted telephone interview (CATI) and records personal mobility for a single day. The study evaluates variables as average travelled kilometers per day, commuting behavior or utilized means of transportation. The analysis is conducted on coarse spatial level (e.g. nationwide, statewide) and therefore relies on a large and comparably stable mobility sample. In comparison, our application requires the evaluation of individually selected poster campaigns, which may range from a few poster sites to a few hundred posters. Therefore, the number of people that contribute poster passages is less than the complete sample. In addition, depending on the spatial distribution of the campaign, the composition of passing people changes. Thus, it can be expected that a changed GPS sample has a higher influence on our application than in classical mobility surveys.

In recent years trajectory data has drawn the attention of the data mining community. Algorithms have been developed for the clustering of (parts of) trajectories [3,4,5], detection of relative motion patterns [6,7] or sequential analysis of movement [8,9,10]. However, all these approaches focus on algorithmic developments and do not consider the sample perspective. All results are restricted to the one provided trajectory sample. However, it is unknown how results may change over time or how large samples are required to be in order to guarantee stable results under a changed mobility data set.

D. Bootstrap and the Estimation of Standard Errors

Most surveys use a data sample in order to study some quantity θ of a defined population. Besides the estimate $\hat{\theta}$ it is then important to know the accuracy of the estimate. The most commonly used accuracy measure is the standard error (*se*), which states the standard deviation of the estimate induced by the data sample. If we know the standard error, we can determine, for example, confidence intervals for the true value of θ. Although the standard error is a very simple measure for statistical accuracy, it has the disadvantage that for most quantities – with exception to the mean – it cannot be computed by a formula from the data sample [11]. One solution to this problem provides the bootstrap method as introduced by Efron in 1979 [12]. It offers a very easy, simulation-based way to estimate the standard error and other quantities from a data sample.

As stated in [11] the standard error of a data sample can be obtained by bootstrap as follows. Let $x = (x_1, x_2, \ldots, x_n)$ denote a data sample and $x^{*i} = (x^*_1, x^*_2, \ldots, x^*_n)$ denote a *bootstrap sample* which is generated by n times randomly sampling with repetition from the original data sample x. We repeat the sampling r times in order to obtain a number of independent bootstrap samples $x^{*1}, x^{*2}, \ldots, x^{*r}$. From each sample we calculate the so-called *bootstrap replication* $\hat{\theta}^{*i}$ with $i = 1..r$ of our quantity of interest. The estimate of the standard error then results from the standard deviation of the bootstrap replications

$$s\hat{e} = \sqrt{\frac{\sum_{i=1}^{r}(\hat{\theta}^{*i} - \hat{\theta}^{*(\cdot)})^2}{r-1}} \qquad \text{with} \quad \hat{\theta}^{*(\cdot)} = \frac{\sum_{i=1}^{r}\hat{\theta}^{*i}}{r}.$$

If the number r of replications approaches infinity, the estimate $s\hat{e}$ approaches the ideal bootstrap estimate for the given sample. Typically, the number of replications lies between 25 and 200.

III. Method Overview

In this section we give an overview of our approach to analyze the robustness in repeated mobility surveys. We explain the workflow and introduce all involved components. We rely on the well-understood bootstrapping and subsampling method, and apply them systematically in the spatial trajectory context.

The aim of our work is twofold. First, we would like to estimate the variance in poster performance for a changed GPS sample of the same size. Second, we would like to estimate the variability that a reduced GPS sample causes in order to determine appropriate sample sizes for future surveys. In both cases our statistic of interest is the reach of poster campaigns. Naturally, reach varies with the specific campaign under consideration. However, the reach of randomly selected poster campaigns of a given size in a city is comparably stable. Note that an increasing number of posters in a campaign has a stabilizing effect on performance measures, because local effects are more often allowed to cancel each other out. We conduct our experiments therefore for varying sizes of poster campaigns. For each campaign size we randomly select several campaigns and average the observed effects.

In order to measure the effect of a changed GPS sample of the same size we set up a bootstrapping scenario. As described in Section II, bootstrapping allows to estimate the standard error of arbitrary statistics and therefore fits our application. In our scenario we repeatedly create bootstrap samples based on the complete GPS data set and evaluate the reach of selected campaigns. From each application of the bootstrap we obtain the standard error for a specific campaign. This error is then averaged over all campaigns of the same size. The details of our experiment setup are shown in Alg. 1.

Alg. 1: Bootstrap on Full Sample

```
Input:
   =  set of campaign sizes S_c = {10,20,....,100}
   =  set of test persons Pers, set of poster locations Loc
      and set of poster passages Pass
   =  # bootstrap repetitions r_b, # campaign repetitions r_c
Output:
   =  sê = (sê_10, sê_20, ..., sê_100) vector with estimates of
      standard error for campaign sizes S_c
Method:
for (s ∈ S_c) { // iterate over campaign sizes
  for (j = 1..r_c) { // iterate over r_c campaigns per size
    C = sample (Loc, s)  // sample campaign
    for (i = 1..r_b) { // calc. bootstrap replications
      θ̂^{*i} = calcBootstrapReplication(Pers, C, Pass)
    }
    sê_j = std(θ̂^{*1}, θ̂^{*2}, ..., θ̂^{*r_b})  // calc. stand. error
  }
  sê_s = avg(sê_j | j = 1..r_c) // average se per camp. size
}
sê = (sê_10, sê_20, ..., sê_100)
```

When we analyse the effect of a reduced GPS sample, we can distinguish two situations. The first situation observes solely the effect of a smaller sample size. The second situation monitors also the effect of a newly drawn sample of smaller size. We represent both situations in our second experiment. In addition to the evaluation of different campaign sizes as in the previous experiment, we now vary also the size of the GPS sample. Similar to the sampling of campaigns, we draw several subsamples for each GPS sample size to compensate sampling effects.

In order to measure the effect of smaller sample sizes only, we calculate the root mean squared error (RMSE) for each GPS subsample when compared to the performance measure of the full GPS sample. Naturally, we expect that the error increases with decreasing sample and campaign size. An interesting question hereby is the relationship and the strength of both quantities.

In order to measure the combined effect of changed samples and small sample sizes, we again apply a bootstrap schema. Hereby, we first reduce the GPS sample size and then apply bootstrap on the reduced set. Again, we draw several GPS subsamples of the same size in order to reduce random effects. Alg. 2 shows the details of both experiments.

ALG. 2: BOOTSTRAP AND SUBSAMPLING FOR VARYING NUMBERS OF GPS PERSONS

Input:

= set of campaign sizes $S_c = \{10,20,\ldots,100\}$

= set of person subsample sizes in percent $S_s = \{2.5, 5.0, \ldots, 97.5\}$

= set of test persons *Pers,* set of poster locations *Loc* and set of poster passages *Pass*

= # bootstrap repetitions r_b, # campaign repetitions r_c and # subsample repetitions r_s

Output:

= $rmse = (rmse_{ts})$ and $\hat{se} = (\hat{se}_{ts})$ with $s \in S_c$ and $t \in S_s$ matrix with estimates of RMSE and standard error for campaign sizes S_c and subsample sizes S_s

Method:

for $(s \in S_c)$ { // iterate over campaign sizes
  for $(t \in S_s)$ { // iterate over person subsample sizes
    for $(j = 1..r_c)$ { // iterate over r_c campaigns per size
      $C = sample\ (Loc,\ s)$ // sample campaign
      for $(k = 1..r_s)$ { // iterate over r_s persons groups per subsample size
        D = $sample\ (Pers,\ t)$ // sample pers. subsample
        for (i = 1..r_b) { // calc. bootstrap replication
          $\hat{\theta}^{*i} = calcBootstrapReplication(D, C, Pass)$
        }
        $\hat{se}_{jk} = std(\hat{\theta}^{*1}, \hat{\theta}^{*2}, \ldots, \hat{\theta}^{*r_b})$ // calc. stand. error
        $err_{jk} = calcReach(D, C, Pass) - calcReach(Pers, C, Pass)$ // calc. error
      }
      $rmse_j = \sqrt{\sum_{k=1}^{r_s} err_{jk}^2} \Big/ r_s$ // calc. RMSE over subsamp.
    }
    $\hat{se}_{ts} = avg(\hat{se}_{jk} \mid j = 1..r_c, k = 1..r_s)$ // average results
    $rmse_{ts} = avg(rmse_j \mid j = 1..r_c)$ // average results
  }
}
$\hat{se} = (\hat{se}_{ts})$
$rmse = (rmse_{ts})$

IV. EXPERIMENTS

A. Application Data and Experiment Setup

In this section we apply our approach to a large real-world dataset in Switzerland. The GPS-study captures the mobility of about 11,000 test persons residing in 12 Swiss conurbations. The number of test persons per region depends on the size of each of the surveyed conurbations. The survey period of each test person lasts seven days. In addition to mobility data the empirical study contains information about the poster sites (55,000 in Switzerland). Besides geographic coordinates, a visibility area for each panel is defined from within which the poster can be seen. Given the trajectories of an individual and the visibility area of a poster panel, all resulting passages can be calculated by geographic intersection and the performance measures can be derived. For this paper we restrict our results to the conurbation of Zurich with 1,956 test persons and 10,093 poster sides. In total, the test persons generate 2,071,124 poster passages within 7 days.

B. Bootstrap on Full GPS Sample

In the first experiment we evaluate the effect of a changed GPS sample for different campaign sizes. We hereby apply bootstrap with 30 repetitions in order to determine the standard error. We varied the campaign size between 10 and 100 posters by an increment of 10, and averaged results over 15 different campaigns per size. Table I and Fig. 3 show the average reach and average standard error as computed by bootstrap for the different campaign sizes. Clearly, the reach increases with the number of posters in the campaign. The standard error, however, decreases with increasing campaign size. This behavior is expected, because larger campaigns allow to cancel out local effects more often.

TABLE I
AVERAGE REACH AND STANDARD ERROR FOR VARYING CAMPAIGN SIZES

camp. size	10	20	30	40	50	60	70	80	90	100
avg. reach	38.8	58.7	67.0	74.9	79.3	83.3	84.3	86.5	89.2	90.1
avg. se	1.5	1.4	1.4	1.3	1.2	1.1	1.1	1.0	0.9	0.9

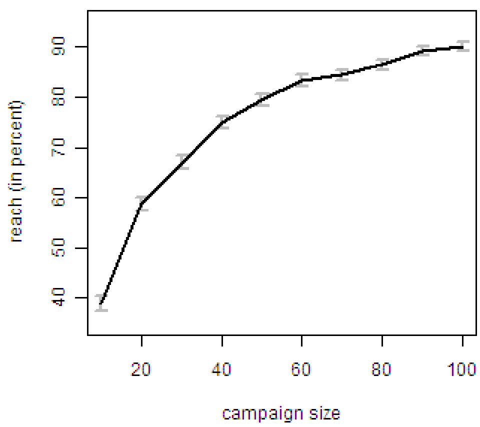

Fig. 3 Average reach and standard error for varying campaign sizes

Fig. 3 shows that the standard error for all campaign sizes is comparably small. Fig. 4 shows again the standard error as obtained by bootstrapping, however magnified to show the differences in detail. The figure also contains a linear regression function which has been derived from the data. It illustrates the linear relationship between the standard error of a changed GPS sample and the size of poster campaigns. In practice, this means that the number of posters in a campaign has a direct influence on the variability of performance measures.

In Zurich a typical campaign contains about 50 posters. The average standard error for this size has a value of only about 1.5% of the average reach, which is an acceptable variation in practice. For the standard case we thus conclude that the complete exchange of the Zurich GPS sample does not lead to large changes in performance measurements for randomly distributed campaigns.

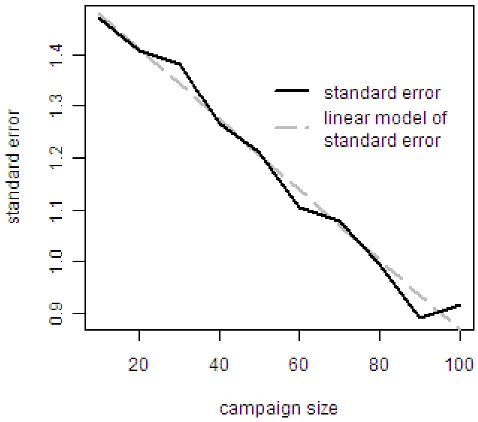

Fig. 4 Average standard error for varying campaign sizes

C. *Subsampling and Bootstrap with Reduced Size*

In the second experiment we evaluate the effected that a smaller sample size has on the variability of performance measures. We again varied the campaign size between 10 and 100 posters by an increment of 10 and averaged results over 15 different campaigns per size. In addition, we formed groups between 2.5 and 97.5 percent of the sample size by an increment of 2.5 percent. For each group size we drew 15 different samples of test persons in order to avoid random effects. In total, our results are thus averages over 225 experiment runs per selected campaign and subsample size.

In a first step we evaluated the error induced only by a decreased GPS sample size and calculated the root mean squared error (RMSE) when evaluating reach based on a GPS subsample instead of the full sample. The results for a subset of subsample sizes are depicted in Table II, and Fig. 5 visualizes the results for all subsample and campaign sizes.

The RMSE is highest for a small campaign and subsample size while it is lowest for a large number of posters and test persons. Again, we can see a linear dependency between the campaign size and the RMSE. However, for decreasing sample sizes the RMSE increases exponentially. Until a sample size of about 40 percent the RMSE curve is still even and afterwards increases rapidly. In practice, this means that we may decrease the GPS sample to about 50 percent of its current size and still obtain reliable results of average reach. However, the current evaluation considers only performance measures calculated for the total population. If we are interested in the reach of a certain socio-demographic group, the sample size will decrease additionally. Therefore, rather than applying the statistics to the total sample number, they must be applied to the smallest socio-demographic group that is of interest for evaluation. In the case of unevenly distributed characteristics, our experiments suggest to implement a stratified sampling process, so that the area with harmful RMSE can be avoided.

TABLE II
AVERAGE ROOT MEAN SQUARED ERROR FOR REDUCED GPS SAMPLES

	camp. size	10	20	30	40	50	60	70	80	90	100
subsample size	2.5%	8.7	8.2	8.7	7.9	6.9	6.9	5.4	6.1	5.7	4.9
	10.0%	3.9	4.1	4.2	3.5	3.8	3.1	3.0	2.6	2.5	2.5
	25.0%	2.3	2.3	2.3	2.1	2.0	1.8	1.7	1.8	1.4	1.5
	50.0%	1.5	1.4	1.2	1.2	1.1	1.0	1.0	1.1	0.8	0.9
	75.0%	1.0	1.0	0.9	0.9	0.9	0.9	0.8	0.9	0.7	0.7
	90.0%	0.7	0.7	0.6	0.7	0.7	0.6	0.6	0.7	0.5	0.5
	97.5%	0.7	0.6	0.5	0.6	0.6	0.5	0.5	0.7	0.4	0.4

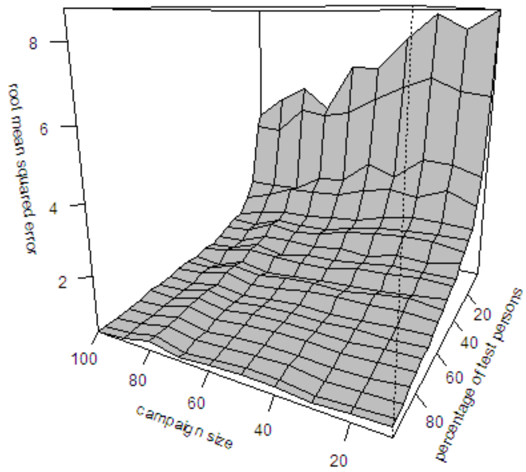

Fig. 5 Average root mean squared error for reduced GPS samples

So far, our experiment accounts only for variability due to smaller sample sizes of a given GPS sample. However, it does not consider variation due to a different and smaller GPS sample. We therefore added a bootstrap scenario to the subsampling experiment. More detailed, for each combination of campaign and subsample size (and iteration of randomly drawn set of posters and persons of the given sizes) we performed bootstrap with 30 repetitions. The obtained results are depicted in Table III and Fig. 6. Note that the surface in Fig. 6 is smoother as in Fig. 5 because each experimental run contains 30 bootstrap repetitions.

The trend in the results is similar to our previous observations. The standard error reacts indirect proportional to the subsample and campaign sizes. One difference, however, is the error height. For a subsample size of more than 25 percent, the standard error obtained by bootstrap lies above the RMSE error of subsampling. In these cases, the actual exchange of the GPS sample introduces additional variance. For subsample sizes below 25 percent, the values of both experiments are about equal, which means that the subsamples are nearly independent of the original sample. In general, the differences between the errors of both experiments are not high. Therefore, we can conclude that the simultaneous reduction and exchange of the GPS sample may be carried out in practice without large implications for the price calculation of campaigns.

Both experiments show another interesting result. For small numbers of test persons the impact of a large campaign on the

RMSE or standard error is higher than for large numbers of test persons. This is a welcome effect in practice, because it means that in case of small sample sizes larger campaigns may reduce the error to a reasonable size.

TABLE III
AVERAGE STANDARD ERROR FOR NEWLY DRAWN REDUCED GPS SAMPLES

	camp. size	10	20	30	40	50	60	70	80	90	100
subsample size	2.5%	8.2	8.5	8.1	7.3	6.9	6.2	6.1	5.8	5.0	4.7
	10.0%	4.0	4.2	4.0	3.7	3.5	3.2	3.0	2.8	2.7	2.5
	25.0%	2.5	2.6	2.5	2.3	2.2	2.0	1.9	1.9	1.7	1.6
	50.0%	1.8	1.9	1.7	1.6	1.6	1.4	1.4	1.3	1.2	1.1
	75.0%	1.5	1.5	1.5	1.3	1.3	1.2	1.1	1.1	1.0	0.9
	90.0%	1.3	1.4	1.3	1.2	1.2	1.1	1.0	1.0	0.9	0.9
	97.5%	1.3	1.3	1.3	1.2	1.1	1.0	1.0	1.0	0.9	0.8

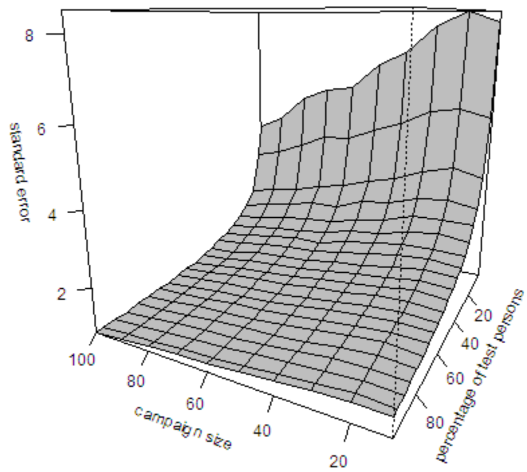

Fig. 6 Average standard error for newly drawn reduced GPS samples

V. CONCLUSIONS

For a number of companies mobility data has become a critical data source. However, as the data grows older, new data has to be collected in order to keep the applications up-to-date. Consequently, it is of great importance to know the impact that a different sample has on the application. In this paper we analyze the effect of a changed GPS samples for outdoor advertising, where mobility data is used to calculate performance measures and poster prices. Outdoor advertising is a very challenging application because it requires the evaluation of mobility data on a very fine spatial level. In comparison, classical mobility surveys focus on movement characteristics on regional or national level and are therefore less influenced by changes in the data sample. Current work on mobility mining algorithms has also not considered stability of results with respect to the data sample so far. In our robustness analysis we apply bootstrapping and subsampling in order to measure the effect of a) a repeated mobility survey and b) a mobility survey of smaller size. Our results show that the induced standard error is comparably small. In addition, the standard error shows a linear relationship to the size of poster campaigns. When we decreased the size of the mobility sample we detected an exponential relationship with the standard error. However, the increase in standard error remained comparably even until we reached a GPS subsample size of about 40 percent. We may therefore conclude that the simultaneous exchange of the GPS sample and reasonable reduction in size may be carried out in practice without large implications for the price calculation of campaigns.

So far our presented evaluations consider only performance measures calculated for the total population. In future work we intend experiments in a stratified sampling process. The total sample number of test persons should be applied to the smallest socio-demographic group that is of interest in media planning.

We believe that in future this type of analysis is relevant for a number of applications in mobility mining where GPS samples have to be replaced after a given period of time as it helps to understand the flexibility and limitations of a given GPS sample.

REFERENCES

[1] J. Z. Sissors and R. B. Baron, Advertising Media Planning, Chapters 4-5, McGraw-Hill, 2002.

[2] BMVBS – Bundesministerium für Verkehr, Bau und Stadtentwicklung Mobilität in Deutschland, Abschlussbericht (Mobility in Germany 2008, final report). Available: http://www.mobilitaet-in-deutschland.de, 2008.

[3] N. Pelekis, I. Kopanakis, I. Ntoutsi, G. Marketos, G. Andrienko and Y. Theodoridis, Similarity search in trajectory databases, In: Proc. of the 14th IEEE International Symposium on Temporal Representation and Reasoning (TIME 2007). IEEE Computer Society Press, pp 129-140, 2007.

[4] S. Rinzivillo, D. Pedreschi, M. Nanni, F. Giannotti, N. Andrienko and G. Andrienko, Visually driven analysis of movement data by progressive clustering. In: Information Visualization 7(3):225-239, 2008.

[5] M. Nanni and D. Pedreschi, Time-focused density-based clustering of trajectories of moving objects. In: Journal of Intelligent Information Systems (JIIS), 27(3):267-289, Special Issue on Mining Spatio-Temporal Data, 2006.

[6] J. Gudmundsson, M. Kreveld and B. Speckmann, Efficient detection of patterns in 2D trajectories of moving points. In: Geoinformatica 11(2):195-215, 2007.

[7] P. Laube and S. Imfeld, Analyzing relative motion within groups of trackable moving point objects. In: Proc. of the 2nd International Conference on Geographic Information Science (GIScience'02). Springer, pp 132–144, 2002.

[8] Y. Zheng, L. Zhang, X. Xie and W. Ma, Mining Interesting locations and travel sequences from GPS Trajectories. In: Proc. of the 18th International World Wide Web Conference (WWW'09). ACM, pp 791-800, 2009.

[9] F. Giannotti, M. Nanni, D. Pedreschi and F. Pinelli, Trajectory pattern mining. In: Proc. of the 13th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD'07). ACM, pp 330-339, 2007.

[10] Y. Yang and M. Hu, TrajPattern: mining sequential patterns from imprecise trajectories of mobile objects. In: Proc. of 10th International Conference on Extending Database Technology. Springer, pp 664-681, 2006.

[11] B. Efron and R. J. Tibshirani, An Introduction to the Bootstrap, Chapman & Hall, 1993.

[12] B. Efron, Bootstrap Methods: Another Look at the Jackknife, In: The Annals of Statistics 7 (1), 1–26, 1979.