

USING DEEP LEARNING TO PREDICT DEMOGRAPHICS FROM MOBILE PHONE METADATA

Bjarke Felbo

Technical University of Denmark
Copenhagen, Denmark
bfelbo@dtu.dk

Pål Sundsøy

Telenor Research
Fornebu, Norway
pal-roe.sundsoy@telenor.com

Alex ‘Sandy’ Pentland

Massachusetts Institute of Technology
Cambridge, MA
sandy@media.mit.edu

Sune Lehmann

Technical University of Denmark
Copenhagen, Denmark
sljo@dtu.dk

Yves-Alexandre de Montjoye*

Massachusetts Institute of Technology
Cambridge, MA
yvesalexandre@demontjoye.com

ABSTRACT

Mobile phone metadata are increasingly used to study human behavior at large-scale. There has recently been a growing interest in predicting demographic information from metadata. Previous approaches relied on hand-engineered features. We here apply, for the first time, deep learning methods to mobile phone metadata using a convolutional network. Our method provides high accuracy on both age and gender prediction. These results show great potential for deep learning approaches for prediction tasks using standard mobile phone metadata.

1 INTRODUCTION

Our mobile phones produce metadata every time we send or received a text or a phone call. These metadata – recording who calls or texts who, for how long, and from where – provide a detailed view of human behavior including mobility at large-scale. This data has great potential for good but often lacks basic demographic information, which is why there has recently been a growing interest in predicting demographic information, such as age and gender, from mobile phone metadata. Previous approaches relied on standard machine learning algorithms and hand-engineered features (Sarraute et al., 2014; Frias-Martinez et al., 2010).

Convolutional networks (ConvNets) have recently systematically outperformed existing approaches in analyses of large-scale image datasets (Krizhevsky et al., 2012; Simonyan & Zisserman, 2014). We show in this work how a ConvNet can be used to predict demographic information such as age and gender from standard mobile phone metadata.

2 TEMPORAL REPRESENTATION

We focus on using the temporal information contained in mobile phone metadata. We represent the data as 8 matrices summarizing mobile phone usage on a given week with hours of the day on the x-axis and the weekdays on the y-axis (see Figure 1). These 8 matrices are combined into a tensor (less formally, a 3-dimensional matrix) with a separate ‘channel’ for each of the 8 variables of interest. These 8 channels are the number of unique contacts, calls, texts and the total duration of calls for respectively incoming and outgoing interactions. We model each week separately.

*Corresponding author: yvesalexandre@demontjoye.com




Figure 1: The mean number of outgoing calls, one of the channel our matrix representation uses, when averaged across the population.

3 CONVNET ARCHITECTURE

Choosing the right architecture of the ConvNet for the task at hand is crucial to reaching high accuracy for the prediction task. We use a series of 5 horizontal conv. layers followed by a vertical conv. filter and 2 dense layers (see Table 1) to capture increasingly abstract patterns.

Table 1: Architecture for the ConvNet. The leaky ReLU activation layers are not shown for brevity.

Layer Name	Conv. Filter Size
<i>Input</i>	-
<i>conv₁</i>	4x1
<i>conv₂</i>	4x1
<i>conv₃</i>	4x1
<i>conv₄</i>	4x1
<i>conv₅</i>	12x1
<i>conv₆</i>	1x7
<i>dense₇</i>	-
<i>dense₈</i>	-
<i>softmax₉</i>	-

The ConvNet performs a prediction of the demographic attribute on each week of data using the softmax layer. These predictions are traditionally averaged to generate a single prediction for the user. We find that using the ConvNet as a feature extractor with a support vector machine (SVM) for prediction increases accuracy by 1 to 3 percentage points (see Table 2).

4 RESULTS

We demonstrate the effectiveness of our method on gender and age prediction. Our results are based on anonymized call detail records (CDRs) for 150,000 people in a Western European country for 15 weeks. Table 2 shows that our method achieves a high accuracy on both age and gender prediction.

Table 2: Accuracy of classifiers when predicting age and gender.

	Age	Gender
Random	35.7%	56.3%
ConvNet	60.7%	78.3%
ConvNet-SVM	63.1%	79.7%

5 CONCLUSION

We here applied, for the first time, deep learning algorithms to prediction tasks using mobile phone metadata. Our method provides a high accuracy on both age and gender prediction. These results show the great potential of deep learning techniques for prediction tasks using mobile phone metadata.

ACKNOWLEDGMENTS

This work was supported in part by the Army Research Laboratory under Cooperative Agreement Number W911NF-09-2-0053. Sune Lehmann acknowledges support from the Villum Foundation and the Danish Council for Independent Research. Yves-Alexandre de Montjoye was partially supported by a grant from the Media Lab and Wallonie-Bruxelles International.

REFERENCES

- Vanessa Frias-Martinez, Enrique Frias-Martinez, and Nuria Oliver. A gender-centric analysis of calling behavior in a developing economy using call detail records. In *AAAI Spring Symposium: Artificial Intelligence for Development*, 2010.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Carlos Sarraute, Pablo Blanc, and Javier Burrioni. A study of age and gender seen through mobile phone usage patterns in mexico. In *Advances in Social Networks Analysis and Mining (ASONAM), 2014 IEEE/ACM*, pp. 836–843. IEEE, 2014.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.