Demo: Mobile Contextual Advertising Platform based on Tiny Text Intelligence

Jung-Hyun Lee, So-Young Jun, So-Jung Park, Kang-Min Kim, SangKeun Lee
Department of Computer Science and Engineering
Korea University, Seoul, Republic of Korea
{jhbslpd, syzzang, sojungpark, kangmin89, yalphy}@korea.ac.kr

ABSTRACT

In-app advertising has become a significant source of revenue for mobile app. In order to improve the effectiveness of in-app ads, most ad networks focus on targeting a user based on the user's personal information collected from their ad library inside mobile apps and the global knowledge built from big data on ad servers. However, sharing user's sensitive information with the ad servers may raise privacy concerns. As opposed to targeting users, mobile contextual advertising seeks to target the app page a user is viewing.

In this demo, we present a novel mobile contextual advertising platform, called *MoCA*, which is designed to improve the semantic relevance of in-app ads in a stand-alone, privacy-protecting manner on mobile devices. MoCA understands the semantics of app page and ads, and then matches semantically relevant ads to the page inside mobile devices. To the best of our knowledge, this is the first work to implement the mobile contextual advertising platform based on the semantic approach without resort to ad servers.

The technical challenge is to perform the classification and searching task of the semantic approach, which needs high computational resources, inside mobile devices with limited hardware resources. To this end, MoCA utilizes the *tiny text intelligence* [1] that is suitable for the mobile environment. It consists of the light-weight classifier and the similarity model on the compressed knowledge base which requires only 34.3MB storage. The tiny text intelligence has been utilized and its effectiveness has already been proven for several intelligent services [2].

Figure 1 illustrates the system architecture of MoCA. In the background process, MoCA first pre-collects ad candidates from ad servers. In the real-time process, it scrapes the dynamic content of the app page viewed by the user through its ad library, and then captures the semantics of the content and matches semantically relevant ads from ad candidates by utilizing the tiny text intelligence.

We evaluate the relevance of the ads shown by MoCA using a real-world dataset consisting of app pages crawled from popular apps in the Google Play Store. We observe that MoCA achieves the relevance performance of as much as 93.1% compared to the semantic approach in the server environment. We also observe that the average runtime and energy consumption of the real-time process on 1,000 ad candidates is 0.965 seconds and 0.352 *mAh* per ad request, which is equivalent to 2.73 seconds of music play. These results show that the overhead of MoCA is not considerable.

We implement a MoCA prototype and a test app using its ad library as independent applications on Android OS. The test app is a simple news aggregation app that shows a list of recent news articles dynamically crawled from the RSS feed of Wall Street Journal. During the demo session, users will be able to read the recent news articles along with relevant ads dynamically provided by MoCA. The demo video is also available here: https://youtu.be/G2O0j7ucYCk.

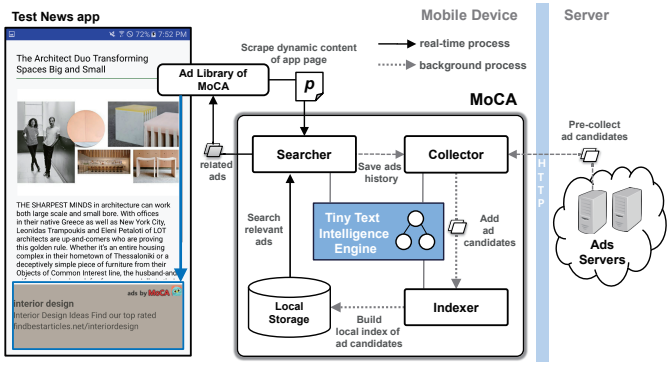

Figure 1: System architecture of MoCA

Acknowledgment

This research was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT and Future Planning (number 2015R1A2A1A15052701).

1. REFERENCES

[1] J. Ha, J.-H. Lee, and S. Lee. EPE: An embedded personalization engine for mobile users. *IEEE Internet Computing*, 18(1):30–37, January 2014.

[2] W.-J. Ryu, J.-H. Lee, K.-M. Kim, and S. Lee. meCurate: Personalized curation service using a tiny text intelligence. In *Proc. WWW '17 Companion*, pages 269–272, 2017.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MobiSys'17, June 19-23, 2017, Niagara Falls, NY, USA.

© 2017 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-4928-4/17/06.
DOI: http://dx.doi.org/10.1145/3081333.3089326