A context-aware recommendation system for mobile devices

Jaehun Lee,[1] Taeho Hwang,[1] Jungho Park,[1] Yunsu Lee,[1] Boris Motik[2] and Ian Horrocks[2]

[1] Samsung Research and [2] University of Oxford

Recommendation systems are one of the most successful and widespread applications of AI, and are used in numerous domains, including retail and commerce (Amazon), and video and music streaming (Netflix, Spotify). The increasing popularity of smart phones and personal assistants makes it critical to provide accurate recommendations that are customised for specific users and their current contexts. There has been some work in this direction, but context is typically limited to location, time and other sensor-based data, and client-server architectures typically require relevant data to be transferred to a server [2]. We aim at a much deeper understanding of the user's context via analysis of the rich data that is typically stored on a mobile device including, e.g., contacts, calendar, e-mail and SMS. At the same time, we address privacy concerns by computing recommendations on-device, and without transferring user data to a server.

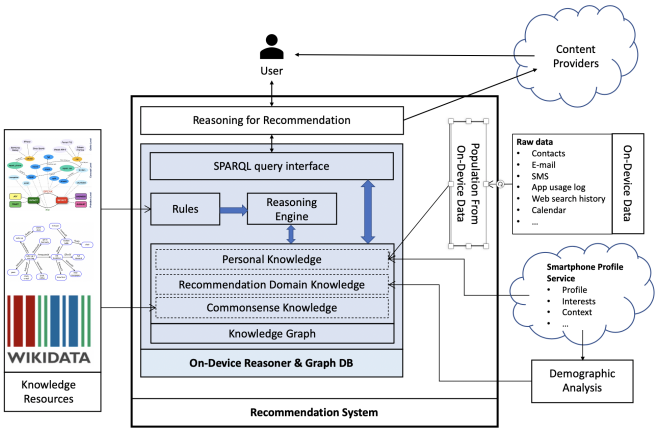

Fig. 1. On-device reasoning-based recommendation system.

System Architecture An on-device system brings new technical challenges, in particular constrained computational resources and limited amounts of available data. We address these challenges by building our system around an on-device version of the RDFox high-performance reasoner and graph database from Oxford Semantic Technologies.[1] Our reasoning-based approach is economical with computational resources and does not require large amounts of training data; it also mitigates the cold-start problem, and supports explanation for recommendations.

Fig. 1 illustrates the high level architecture of our system. The reasoner operates on an RDF knowledge graph, which is populated from a variety of sources. We use a template-based technique to extract contextual knowledge from user data [3], and demographic analysis to extract domain knowledge from user profile data stored on the

[1] `https://www.oxfordsemantic.tech/`

server. Background and "common-sense" knowledge is also needed to understand the semantics of the information being gathered, and this is derived from sources such as *Wikidata*, *ConceptNet* [5] and *SenticNet* [1]. The system uses rules to integrate the various sources of information and to derive an understanding of the user, their context and their preferences; the resulting enriched knowledge graph is accessed via a SPARQL interface and used to generate appropriate requests to content providers. RDFox reasons by materialising all the triples implied by the data and rules, which allows for fast query answering, but it also supports incremental materialisation, which allows for real-time responses even with rapidly changing contextual data [4].

Experimental Validation We have evaluated the system both w.r.t. performance of the on-device reasoner and quality of recommendations. We tested performance on a Samsung Galaxy Note 9 with 6GB RAM, an Exynos 9810 8 core CPU and running Android 9. We used two LUBM benchmark datasets with approximately 10^5 and 10^6 triples respectively. Performance was more than satisfactory: materialisation took less than 1s even for the larger dataset, most of the benchmark queries were answered in less than 1ms and no query took longer than 116ms. Moreover, RDFox could incrementally update the materialisation in only a few milliseconds when the data changed.

We tested the subjective quality of recommendations via a user study with a group of 30 Samsung employees of diverse age, gender, and experience. Users rated the recommendations from our system and from a conventional recommendation system based on an alternating Least Square (ALS) algorithm, and we compared the Discounted Cumulative Gain (DCG)10, and the non-Discounted Cumulative Gain (nDCG). The results showed that the average rating for recommendations from our context-aware system was 73% higher than for the control ALS system.

Outlook We plan to use the system in a new service running on Samsung smartphones that support a voice assistant. The recommendation service will encompass music, video, articles, applications, and contextual greetings, and will deliver an improved user experience by providing personalised and context-sensitive recommendations. We also plan to evolve the system to a client-server hybrid architecture in which the on-device reasoner continuously refines its knowledge of the user and sends requests to the server to obtain relevant domain and commonsense knowledge. This will enable us to more effectively manage the trade-off between requirements for comprehensive knowledge and limited memory usage on the client device.

References

1. Cambria, E., Speer, R., Havasi, C., Hussain, A.: Senticnet: A publicly available semantic resource for opinion mining. In: 2010 AAAI Fall Symposium Series (2010)
2. Colombo-Mendoza, L.O., Valencia-García, R., Alor-Hernández, G., Bellavista, P.: Special issue on context-aware mobile recommender systems. Pervasive Mob. Comput. **38** (2017)
3. Li, Y., Yang, Z., Guo, Y., Chen, X., Agarwal, Y., Hong, J.I.: Automated extraction of personal knowledge from smartphone push notifications. In: IEEE Int. Conf. on Big Data (2018)
4. Motik, B., Nenov, Y., Piro, R., Horrocks, I.: Incremental update of datalog materialisation: the backward/forward algorithm. In: Proc. of AAAI 15. pp. 1560–1568. AAAI Press (2015)
5. Speer, R., Chin, J., Havasi, C.: Conceptnet 5.5: An open multilingual graph of general knowledge. In: Proc. of the 31st Nat. Conf. on Artificial Intelligence (AAAI 17) (2017)