ETH
Eidgenössische Technische Hochschule Zürich
Swiss Federal Institute of Technology Zurich

Distributed Computing

Prof. R. Wattenhofer

Small-Data Recommender: How to Learn Quickly

Netflix, Spotify, Amazon. What do these companies have in common? They all have massive user-numbers and therefore can employ state-of-the art machine learning techniques to build recommender-systems to provide their users with more content.

So, the moral of the story is: Only if you have a lot of data can you build a good system that will appeal to many users. Chicken-or-Egg problem, anyone? While it is indisputable that the recommender systems of aforementioned companies work very well, it is not clear that we cannot do well without a lot of data. How do we learn as much as possible about a user without having many users and collecting tons of data? How much can we learn, what are the best ways to learn and how do these "small data" methods compare to the Big Data ones? We already have some ideas on how to approach this problem, but we would like to hear your thoughts as well. Depending on your interests, this thesis can take many different directions.

If this sounds interesting to you, do not hesitate to contact us so that we can have a chat and exchange ideas.

You will: Learn about state-of-the-art recommender systems. Develop and implement your own algorithms that will help us to "learn quickly". Evaluate your system against state-of-the-art recommender systems. Learn how to read scientific papers. Meet with your supervisors on a weekly basis to discuss progress and open questions. And so on.

Interested? Please contact us for more details!

Contacts

- Gino Brunner: gino.brunner@tik.ee.ethz.ch, ETZ G63