Invited talk

Automatically explaining health information

Emiel Krahmer
`e.j.krahmer@tilburguniversity.edu`

Tilburg center for Cognition and Communication (TiCC)

Abstract

Modern AI systems automatically learn from data using sophisticated statistical models. Explaining how these systems work and how they make their predictions therefore increasingly involves producing descriptions of how different probabilities are weighted and which uncertainties underlie these numbers. But what is the best way to (automatically) present such probabilistic explanations, do people actually understand them, and what is the potential impact of such information on people's wellbeing?

In this talk, I adress these questions in the context of systems that automatically generate personalised health information. The emergence of large national health registeries, such as the Dutch cancer registry, now make it possible to automatically generate descriptions of treatment options for new cancer patients based on data of comparable patients, including health and quality of life predictions following different treatments. I describe a series of studies, in which our team has investigated to what extent this information is currently provided to people, and under which conditions people actually want to have access to these kind of data-driven explanations. Additionally, we have studied whether there are different profiles in information needs, and what the best way is to provide probabilistic information and the associated undertainties to people.