CUSTOMER EXPERIENCE

Search Relevance

Search is the single biggest facilitator of product discovery for a customer purchasing on Amazon. The concept of relevance is not easily captured from data, is multi-dimensional and ever evolving. For example, customers would sometimes like to find fast deliverable products and sometimes they look for products that are of the highest quality. At Amazon, we constantly endeavor to improve the overall search experience by improvising existing and formulating new ML techniques to fully understand customer's intent on search. We are working towards areas of spell error correction, leveraging phonetics, and supporting transliterated and colloquial usages to improve the way we analyze and address customer queries. We build product category specific models for search relevance and take a multi-objective optimization-based approach to capture the various aspects of relevance a customer may imply.

Question-Answering

Product detail pages often overwhelm customers with a wealth of data, making discoverability of relevant information a challenge, especially on mobiles. At IML, we are building automated conversational agents to address this problem. With such agents, customers do not have to read through thousands of words on a product detail page, across product description, technical specifications and product reviews. They can ask a question in their own words, without worrying about knowing specific terms or technical jargon. Building such a system poses an exciting set of challenges, such as understanding the user intent, semantic matching of user questions and answers on the detail page, requirement of high precision, lack of existing training datasets, mixed-language questions and natural-language responses. To address these challenges, we are building conversational systems using state-of-the-art Deep Neural Network frameworks, leveraging techniques like LSTMs, CNNs, attention, phrasal-embeddings and reinforcement learning.

Size recommendations: Appropriate sizing is a primary friction point for customers shopping online for apparel and shoes. Brands have different sizing conventions - size 6 in Nike may be different from a Reebok's size 6, and physically measuring every product is difficult. Survey based mechanisms can sometimes be insightful but have limited penetration rates amongst customers. Based on customer purchase and returns data, we have built Bayesian and non-Bayesian size recommendation algorithms to recommend product sizes to customers. Bayesian algorithms are particularly useful to deal with issues arising from noise and data sparsity. Some of this work has also been accepted at ACM Recsys 2017, and WWW-2018.