

# An Introduction to Entity Recommendation and Understanding

Hao Ma, Yan Ke  
Microsoft, One Microsoft Way, Redmond, WA 98052  
{haoma, yanke}@microsoft.com

## ABSTRACT

Entities and the Knowledge about the entities have become indispensable building blocks in modern search engines. This tutorial aims to present the current state of research in the emerging semantic search and recommendation field by studying how to help users effectively explore the knowledge base as well as answer their information needs. Many live applications and systems will be demonstrated throughout this tutorial. After the completion of the tutorial, the audience will have an introduction and overview of what is the emerging topics of entity recommendation and understanding. The audience will learn and be able to understand some current research work as well as industry practices using computational intelligence techniques in entity recommendation and understanding. One of the major goals of this tutorial is to help audience identify a few research directions that could have big impact in the near future.

## Categories and Subject Descriptors

H.3.3 [Information Search and Retrieval]: Search Process, Information Filtering

## Keywords

Entity Recommendation; Knowledge Graph; Question and Answering; Personalization; Semantic Search

## 1. TOPICS

The topics included in this tutorial are (but not limited to):

- **Introduction and Motivation**
  - Overview of the Tutorial
  - A Basic Introduction to Entities and Knowledge Bases
  - Challenges in Entity Recommendation and Understanding

- **Live Demonstration on Bing's Entity Experience**
  - Entity Recommendation and Ranking
  - Entity Recommendation Interpretation
  - Entity Recommendation Exploration
  - Entity Carousel Ranking (List Ranking)
  - Factoid Answers
  - Conversational Entity Question and Answering
  - Web Question and Answering
- **Entity Recommendation Techniques**
  - Commonalities and Differences with Traditional Recommender Systems
  - Architecture of a Real Entity Recommender System
  - Common Techniques in Entity Recommendation
  - Utilization of Query Logs and Context Information
  - Personalized Entity Recommender Systems
  - A New Three Way Personalized Entity Recommendation Technique
- **Other Related Techniques in Understanding Entity**
  - Entity Understanding in Text or Web Documents
  - Brief introduction on Entity Linking
  - Web Question and Answering Techniques
  - How These Techniques can Benefit Entity Recommendation Experience
- **Conclusions, Directions and Discussions**

## 2. TUTORIAL DESCRIPTION

Recent years have witnessed rapidly increasing interests on the research field of semantic search. Knowledge base powered entity search and recommendation experience has been widely adopted by major search engine companies. Although many work has been introduced in the field of traditional recommender systems or information retrieval, entity recommendation and search techniques differ significantly from them due to the introduction of knowledge base. The heterogeneity, semantic richness and large-scale nature of knowledge base make traditional approaches less effective.

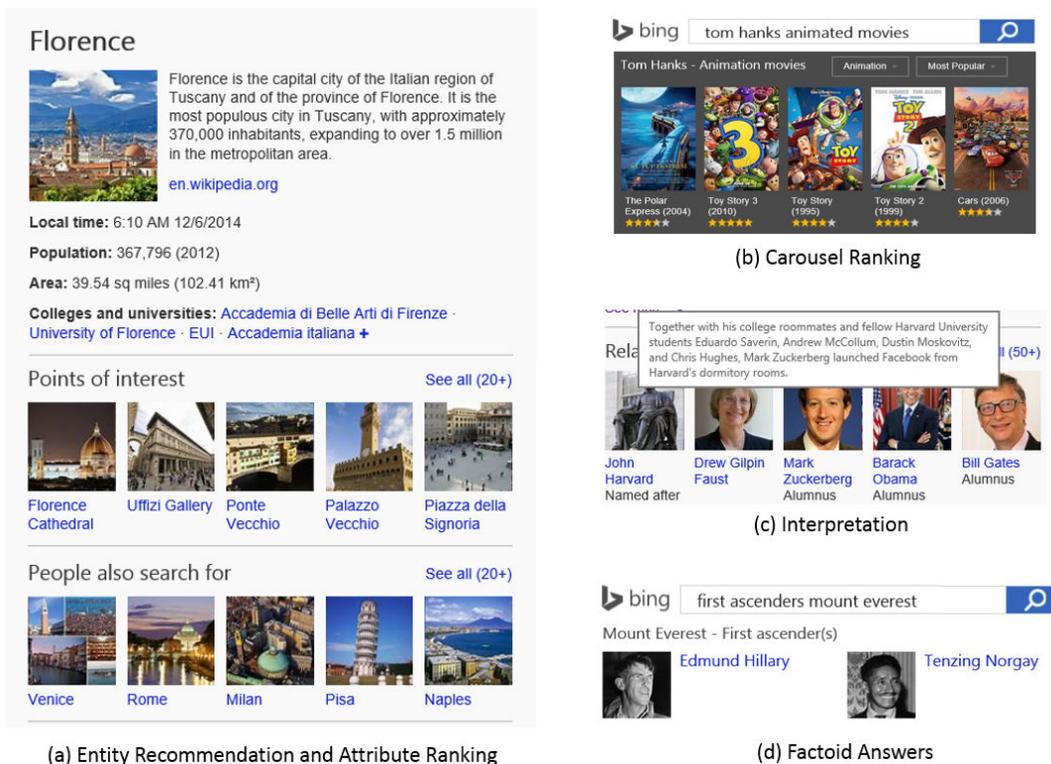


Figure 1: Screen Shots of Entity Recommendation and Understanding Live Systems

In this tutorial, we provide the first detailed introduction on how entity recommendation works and how various entity understanding techniques could further improve entity recommendation and search experience.

This tutorial consists of four major parts. In the first part, we give a brief introduction on entities and knowledge bases. We also show how we collect information from different data sources as well as how we infer users' interests on specific entities. In the second part, we demonstrate various entity recommendation and search applications we developed and productionized in Bing, including entity recommendation (Fig. 1(a)), natural language interpretation of recommendation (Fig. 1(c)), attribute ranking (Fig. 1(a)), carousel ranking (Fig. 1(b)), entity exploration, factoid answers (Fig. 1(d)), conversational question and answering, etc. The architectures, challenges, and corresponding solutions on these systems will also be briefly introduced in the second part of the tutorial. The third part will give a deep dive on the recommendation algorithms that are related to entity recommendation, including basic non-personalized recommendation algorithms as well as recommendation models that tailor related entities to an individual search user's unique taste and preference. The fourth part of this tutorial will focus on how to further improve semantic recommendation and search experience by employing other entity understanding techniques.

The tutorial will conclude by summarizing a whole area of exciting and dynamic research that is worthy of more detailed investigation for many years to come.

### 3. AUDIENCE

The target audience is both *researchers* as well as *industry practitioners*. Researchers in the field of semantic search, recommender systems, information retrieval and natural language processing will benefit the most as this will give them an exhaustive overview of the research in the direction of entity recommendation and understanding. Industry practitioners will also benefit from the this tutorial a lot since we will demonstrate many applications that we developed and shipped to Bing. We believe that the tutorial will give the newcomers/students a complete picture of the current work, introduce important research topics in this field, and inspire them to learn more.

### 4. PRE-REQUISITE

The tutorial will first give an overview of Entity Recommendation and Understanding and show some live demonstrations. This will be general knowledge and will not require any specific technical expertise. The tutorial will then move into more specific topics such as detailed recommendation algorithms. These topics will require more technical knowledge such as machine learning techniques, or some natural language processing preliminaries. However, the level will be aimed at college junior/senior level and should be easily followed for academic researchers and practitioners in the industry.