

# Social Networks and Discovery in the Enterprise (SaND)

Inbal Ronen, Elad Shahar, Sigalit Ur, Erel Uziel, Sivan Yogev, Naama Zwerdling,  
David Carmel, Ido Guy, Nadav Har'El, Shila Ofek-Koifman  
IBM Research Lab, Haifa 31905, Israel  
{inbal,elads,sigalit,erelu,sivany,naamaz,carmel,ido,nyh,shila}@il.ibm.com

**Categories and Subject Descriptors:** H.3.3 [Information storage and retrieval]: Information search and retrieval

**General Terms:** Algorithms

**Keywords:** social search, enterprise search, social network

## Introduction

Traditional information discovery methods are based on content: documents, terms, and the relationships between them. In Web 2.0, people come into play as they create documents and tags in many forms. Personalized search, social graphs, content and people recommendation, are some of the tasks that can take advantage of this newly formed ecosystem.

The Social Networks and Discovery (*SaND*) platform is an aggregation tool for information discovery and analysis over social data gathered from Web 2.0 applications in the enterprise. It leverages complex relationships between content and people as surfaced through the social applications to unleash the value of information. Its integrated index supports combining content-based analysis and people-based analysis over a rich data foundation. Enterprise social data is easily modeled and ingested into *SaND*, and can be further combined with data from external social applications.

This demo will present three main functions provided by *SaND*:

**Social search:** *SaND* supports search over the social data using a unified approach [1] in which all system entities (documents, people, tags) are searchable and retrievable (See Figure 1). The search UI enables the searcher to get a wider view on the query topic through results from all entity types, while uncovering the relationships between the on-screen entities.

**Entity recommendation:** *SaND* can be utilized to recommend people and content for the searcher (Figure 2 shows the “Do You Know” widget for people recommendation). People are recommended according to their “social similarity” relations with the searcher, e.g. organizational and friending ties, similar tagging activity and more. Similarly, content that is related to people that are “socially related” to the searcher is recommended as valuable interesting data.

**Personalization:** Search results are personalized by considering the relations of retrieved entities with the searcher. Entities are ranked according to their relevance to the query as well as according to their relationship strength with the searcher.

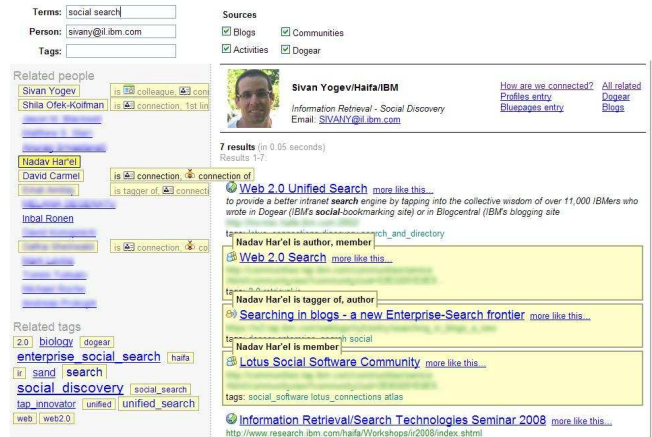


Figure 1: A screenshot of *SaND*'s search results



Figure 2: *SaND*'s “Do you know” widget

## How does it work?

The cornerstone of the *SaND* platform is the ability to aggregate information from different sources for any entity while preserving its relations with other entities. Through its aggregation model, *SaND* supports queries over any entity in the system, be it a textual term, a person or a tag, retrieving a ranked list of entities related to that entity. In particular, the social network of a person based on a flexible set of relationships can be obtained [2].

## 1. REFERENCES

- [1] E. Amitay, D. Carmel, N. Har'El, S. Ofek-Koifman, A. Soffer, S. Yogev, and N. Golbandi. Social search and discovery using a unified approach. In *HyperText*. ACM, 2009.
- [2] I. Guy, M. Jacovi, E. Shahar, N. Meshulam, V. Soroka, and S. Farrell. Harvesting with SONAR: the value of aggregating social network information. In *CHI*. ACM, 2008.