# Best *Faces* Forward: A Large-scale Study of People Search in the Enterprise

**Ido Guy, Sigalit Ur, Inbal Ronen**
IBM Research Lab
Haifa 31905, Israel
{ido, inbal, sigalit}@il.ibm.com

**Sara Weber, Tolga Oral**
IBM CIO's Office
Cambridge, MA 02142, USA
{sara_weber, tolga_oral}@us.ibm.com

**ABSTRACT**
This paper presents *Faces*, an application built to enable effective people search in the enterprise. We take advantage of the popularity *Faces* has gained within a globally distributed enterprise to provide an extensive analysis of how and why people search is used within the organization. Our study is primarily based on an analysis of the *Faces* query log over a period of more than four months, with over a million queries and tens of thousands of users. The analysis results are presented across four dimensions: queries, users, clicks, and actions, and lay the foundation for further advancement and research on the topic.

**Author Keywords:** People search, enterprise search, large scale

**ACM Classification Keywords:** H.5.3 Group and Organizational Interfaces: *Computer-supported cooperative work*

**General Terms:** Design, Experimentation, Human Factors

**INTRODUCTION**
Searching for other individuals is one of the most fundamental scenarios in an enterprise. As businesses become global and distributed, employees more often need to look for others in the organization and find out their job title, organizational unit, contact information, management chain, or office location. We define *people search* as any search in which the returned entities are people.

Despite its fundamentality, people search within the enterprise has received little attention in the literature. Most of the existing studies focus on the expertise location challenge, where the employee looks for another employee who is knowledgeable of a certain topic or field. However, searching by expertise is just one type of people search; other search criteria can include name, location, job role, email, phone number, or any combination of these.

Even outside the enterprise, studies on people search have been limited. Many works have studied web search engines, where the returned entities are web pages (e.g., [2,18]). There are a growing number of studies on vertical search engines, which specialize in a single domain of online content, such as books [11], scientific literature [12], blogs [15], or audiovisual content [8]. Although people search can be viewed as another type of vertical, it is only recently that a study provided the first comprehensive query log analysis of a commercial Dutch engine for people search [20]. The study showed that most people search is for celebrities, key players of bursty events, and friends or family members.

In the enterprise, people search has a rather different set of motivations. Employees may look up the details of an individual with whom they have a meeting or correspond with on email or instant messaging. They may explore the organizational unit or management chain of a complete stranger whose name they heard during a call. They may also look up a specific detail, such as the phone, email, or office location of a person they already know. In some cases, employees may only have partial information about the person they want to find: Alice whose last name starts with an 'H', Bob who works in Dublin, or someone whose last name is Johnson and works in the Research division.

In this paper, we present *Faces*, an application for people search in the enterprise, and analyze its use within IBM. In *Faces*, the results are presented and updated while the user is typing and fuzzy search handles misspelling. The user interface is designed to provide the most essential details of individuals and to allow easy navigation across the organizational chart. Scoring heuristics are used to bring the most relevant people as top results. The massive-scale backend pre-calculates and stores information to support fast response. Person data is kept in memory to speed scoring and display at runtime.

*Faces* has been rapidly adopted within our organization, gaining tens of thousands of users per month. In this work, we take advantage of its popularity to provide a large-scale evaluation of enterprise people search, relying on over four months of data, with over one million queries, and 35,000 distinct users. We also conducted a survey with 661 participants and interviewed 20 users. The goal of our analysis was to better understand the main scenarios and motivations that drive the search for other people in the enterprise. We wanted to gain insight on how employees perform their search, the people they search for, and how the *Faces* design supports this activity. To the best of our knowledge, this is the first study to provide a


*CHI'12,* May 5–10, 2012, Austin, Texas, USA.

comprehensive analysis of a people search engine's use within the enterprise.

The rest of the paper is organized as follows. We open with related work, followed by a description of *Faces* and its user interface. We briefly describe the scoring method and backend mechanisms, but those are not analyzed in this paper. The analysis section presents an extensive overview of how *Faces* is used to search for people in the enterprise, across four dimensions: queries, users, clicks, and actions. We conclude by discussing our findings and suggesting future directions.

## RELATED WORK

Searching people in the enterprise may be associated with the broader domain of *enterprise search*, typically referring to searching for content within the organization. Despite the advances in content search technology over the years, research shows that employees still spend a large amount of time searching for information [7]. Amitay et al. [1] describe a system for enterprise social search, which returns, in addition to documents, also people and tags. The people search part is approached and evaluated as an expertise location task.

*Expertise location* helps users find a person with knowledge or information about a certain technology, process, or domain. Typically performed within an organization, it is the only well-studied type of enterprise people search. Quite a few empirical field studies of enterprise expertise location systems have been conducted along the years. For example, Mcdonald and Ackerman [13] performed a field study within a medium-size software company and Reichling et al. [17] conducted a field study within a highly decentralized industrial organization. Yiman-Seid and Kobsa [21] identified two motives for seeking an expert: as a source of information and as someone who can perform a social or organizational role. Ehrlich and Shami [3] further enumerated four motives: getting answers to technical questions, finding people with specific skills, gaining awareness of "who is out there", and providing information.

Several studies suggested using other criteria, such as the organizational or social network, in addition to a topic, as query for expertise location. ReferralWeb [10] was one of the first systems to do so, allowing users to specify a search topic and a social criterion (e.g., people who are related by up to two degrees to John Doe). Expertise Recommender [14] filtered expert search results based on two elements of the user's network: organizational relationships and social relationships gathered through ethnographic methods. Smirnova et al. [19] proposed a model that ranks experts based on a linear combination of their knowledge on the queried topic and their contact time, estimated by their social distance from the user. All of the above still have the topic as the center of the query. In our work, the query is not constrained to include a topic and thus addresses a broader scope of people search than expertise location.

Our evaluation is primarily based on analysis of the *Faces* query log. *Query log analysis* is a common method for evaluating search engines and has been widely used in past work. One of the first studies [18] provided a large-scale query log analysis of the AltaVista web search engine, including popular query terms, query length, number of clicks per search, and more. Broder [2] classified queries in web search engines into three types: informational, navigational, and transactional. When Jansen and Spink [9] compared the query logs of nine web search engines, they found that entity search, including people search, has been on the rise.

The most relevant related work is a very recent query log analysis of a Dutch people search engine on the Web [20]. As the authors state, "*it is the first time a query log analysis is performed on a people search engine*". The authors found that the percentage of one-query sessions (over 50%) is higher as compared to web document search, and that the click-through rate (17%) is much lower than for document search. Additionally, less than 4% of the queries included a keyword along with the person's name. Our result analysis and discussion further relate to that study and highlight the commonality and difference between web and enterprise people search.

## THE *FACES* APPLICATION

### Design Goals

*Faces* is a web application used to find people within an enterprise. Two applications existed in the enterprise before we deployed *Faces*: a user interface into the Corporate Directory (CD) [4] and an enterprise social network site [61] (SNS). A user profile in the SNS includes the employee's "friends" and tags applied by others, as well as recent activity in enterprise social media, such as blogs, wikis, bookmarks, and forums. *Faces* was designed to overcome many of the deficiencies found in these existing applications and included the following goals:

- Return as many results as possible, as fast as possible, and score them so that the most relevant people show up first
- Emphasize closer matches in the interface
- Search a mixture of user profile attributes; support partial matches and misspelling
- Allow quick navigation over the *organizational environment*: direct reports, peers, and managers
- Show people's faces

Fast response was of vital importance. We were introducing *Faces* into an enterprise that had existing, extensively-used applications for people search. It was essential to introduce significant improvements in functionality and performance to get users to switch. Our goal was to have results displayed within 100 milliseconds for each search.

The rest of this section describes *Faces*. We provide a comprehensive overview of the application being used to analyze enterprise people search. The backend, scoring

method, and design choices are not specifically evaluated in this work. Rather, we strive to provide a sense of how and why people search for people within the enterprise, relying on the large user-base accumulated by *Faces*.

### User Interface

The *Faces* application starts out with a simple interface that includes an empty search box and a prompt "*the best way to find employees. Period. Just type. No Waiting. Give it a shot!*" The user simply enters the information related to the person they are looking for. This may be a first or last name, or some other data about the person such as their job title, location, or a tag associated with them. Tags are retrieved from an enterprise people tagging application, which allows employees to annotate each other with descriptive terms [4]. As the user starts typing, *Faces* updates the results dynamically. A search is performed with each character typed, unless another one is typed less than 100 milliseconds thereafter (to avoid redundant searches when the user types very fast or pastes a whole string). If the user types another character while the search results to the previous string are still being processed, they are discarded. Otherwise, the top results are displayed to the user with both a picture ("face") and basic identification information, as shown in Figure 1. The display is continuously updated with results as the user types.

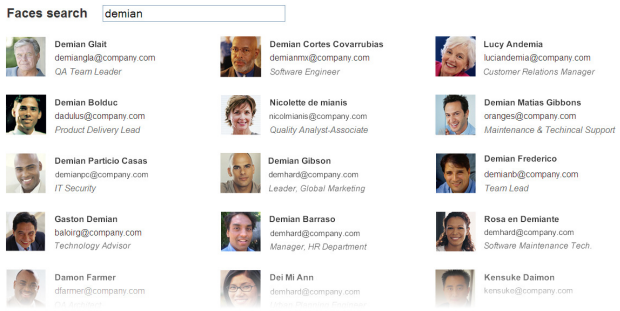


**Figure 1. Initial list of results.**

In the existing CD and SNS applications, searching is not performed until the full text is typed and submitted by the user. In contrast, *Faces* provides instant feedback while the user is typing. The user may often find the person they are looking for even before they finish typing the data.

As more text is entered into the search box, the confidence in the potential results increases. When the results become more distinctive, we present the top ones with a larger thumbnail to make them easier for the user to notice. Figure 2 depicts a case of two larger results. We only present larger results if they have a score that is at least 20% higher than the score of the subsequent result. The number of larger results is determined according to the gap in score of the subsequent results and does not exceed four.

*Faces* provides basic information for initial results of face, name, email, and job title (see Figure 1). More detailed results also include location (city and country), division,

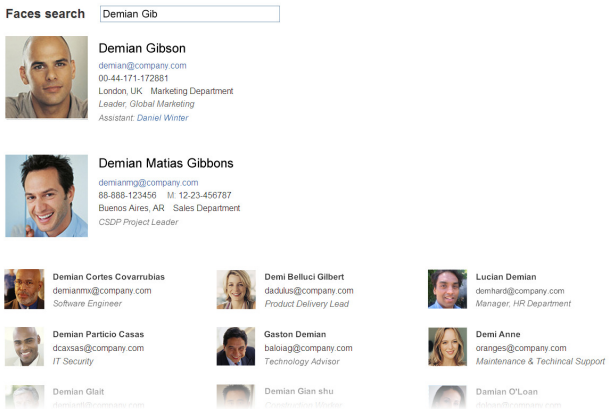


**Figure 2. Enlarged results when score is high.**

phone number, and a linked name to their assistant, if one exists (See Figure 2). When users click a result, a larger "lightbox" pops up with more information, as depicted in Figure 3. This information includes the *organizational environment*: management chain, peers, and direct reports (in case of a manager). It also contains a "more info" link, which replaces the organizational environment view with details such as the person's office location (building, floor, number), serial number, links to both the person's CD and SNS pages, as well as a "permalink" URL to this person's information within *Faces*. The user can click again to switch back to the organizational environment view.

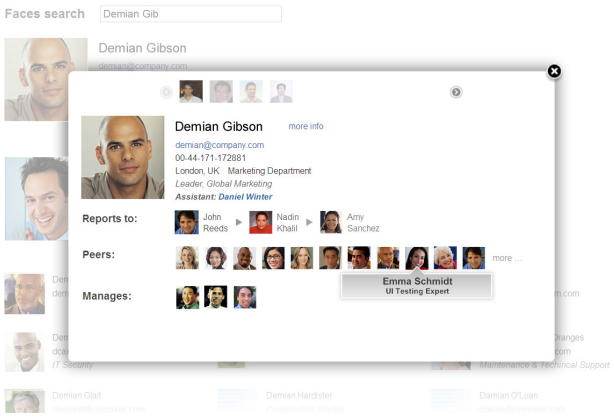


**Figure 3. Lightbox of a clicked person result.**

Upon hovering over a face in the organizational environment, a tooltip appears with the person's name and job title. In addition, users can easily browse up, down, and across the organization chain. Clicking a person's face displays their information on the lightbox, instead of the person currently presented. As people are selected from the organization environment, their pictures are kept at the top of the lightbox as breadcrumbs to allow jumping conveniently to any of them (see Figure 3).

A user who is not identified is prompted to do so by a string that appears at the bottom of each page "*Want better results? Tell us who you are!*" After clicking the string,

users can identify by searching for themselves and clicking on the correct result. A persistent cookie is used to recognize returning users who have previously identified.

### Scoring

Scoring brings the most relevant results to the top of the list. As the user starts typing, the backend system gathers all people that have a profile field that matches the search term(s). Scoring is then performed on this set of matches by calculating a cumulative score that is a product of the importance of the matching field and the strength of the match. Profile fields are grouped into three categories, in decreasing order of importance: 1) first name, last name, and email; 2) job description, location, and tag; 3) middle name and organizational unit. The strength of the match designates how well the search term (token) matches the field's text; it has three possible values for: exact match, prefix (field text starts with the search term), and substring (field text contains the search term, but does not start with it). In case of multiple matches of a token to a field, the one that yields the highest value, taking into account both field importance and match type, is considered for scoring.

On top of the basic scoring, *Faces* applies a personalization boost when the user has identified. Personalization boost is added when (in descending order of strength): the person is in the searcher's network or vice versa; the person is in the searcher's management chain or vice versa; they share the same work location, organizational unit, or country.

If less than 250 results match the query, *Faces* performs a fuzzy search to catch phonetic misspellings and extend the set of results. For fuzzy search, the Metaphone 3 software (http://www.amorphics.com), which extends the Metaphone algorithm for phonetic encoding [16], is used. Up to one million additional results are fetched based on misspelling alternatives. The match strength value for these results is normalized according to the phonetic resemblance.

### Backend

The *Faces* backend is built to support very fast response to people search queries and enable the dynamic update of results. Project Voldemort (http://project-voldemort.com), a distributed key-value storage system, serves as the main data storage mechanism. Four main data structures are stored through Voldemort: 1) a mapping of person IDs to person entities with all relevant profile fields; 2) a mapping of person IDs to images. Images are pre-scaled to all different desired sizes; 3) a mapping of every possible substring of an employee's profile field value to person IDs of the relevant employees, who have at least one field whose value contains the string. Based on this data structure, *Faces* determines person IDs that are relevant to the query. For multiple-token queries, intersection is performed across tokens to determine the relevant person IDs; 4) for handling fuzzy search, *Faces* maps phonetically-close strings, resolved by Metaphone 3, and well-known nicknames to all names within the corporate directory.

The system uses Apache Hadoop's map-reduce paradigm (http://hadoop.apache.org) to distribute the burden of pre-computation and load the data into Voldemort. Data is retrieved from multiples sources, including the corporate directory (for most person fields and images), the people tagging application (for tags), and a social network aggregation system, called SONAR [5] (for social network information, used for personalization). SONAR calculates a weighed list of a person's familiar people in the organization, taking into account relationships as reflected across various enterprise systems, including the explicit enterprise SNS, the organizational chart, databases of projects and patents, and enterprise social media (wikis, blogs, forums). SONAR has shown to effectively produce the list of a person's familiar people in the enterprise [5].

The *Faces* runtime is implemented principally using Java web application servers. Person information is loaded into memory from Voldemort (data structure 1 above) at server start-up. In addition, mappings of all substrings of length 1, 2, and 3 to their matching person IDs (data structure 3) are loaded into memory to allow speedy response to the query's first few characters. Scoring is performed at runtime as explained in the previous section.

## ANALYSIS

### Setup

Our evaluation is based primarily on query log analysis. The *Faces* query log documents every query string sent to the server, along with its respective timestamp and the user's IP address and ID if they are identified. For each query, the log records the interface actions taken by the user, such as clicking on results or navigating the organizational environment. We analyzed the logs recorded from Sept. 19, 2010 to Feb. 7, 2011 (142 days overall).

We also conducted a short user survey to cover several aspects that could not be inferred from the logs, such as the most important piece of information about a person found or the use of copy-paste versus manual typing. The survey also prompted for general free-text comments. We sent the survey to the top 2000 *Faces* users and received 661 responses. Participants originated from 45 countries, spanning the different divisions within our organization. Furthermore, we interviewed 20 *Faces* users to get an in-depth understanding of why and how they use *Faces*. The interviewees originated from 13 countries, spanning different usage levels of *Faces*. The interviews were conducted by phone and lasted half an hour each.

### General

When asked about the most compelling features of *Faces*, our survey participants and interviewees noted the dynamic display of results and their high relevance, fuzzy search support, simplicity of the interface, and speedy performance. Many said they found *Faces* to be the most useful intranet application. Easy, cool, snappy, clean, practical, convenient, useful, handy, intuitive, and (most

commonly) fast, were among the popular adjectives used to describe it. One participant wrote: "*This is a fantastic application and I have come to depend on it for my day-to-day business function. The ability to search using various fields is very useful*" and another stated "*Together with email and calendar, this is my most used internal service*."

Some people mentioned that they keep a browser tab open with *Faces* on it, for "*fast access to people*". They use *Faces* to find a person in a matter of seconds and typically do not spend extra time on exploration. Some did mention that serendipity may occur when someone else in the display catches their eye.

Among the most common usage scenarios people mentioned were searching for individuals who send mail to their inbox, appear in their calendar meetings, and participate in chats and phone calls. One participant wrote: "*I often use Faces when I get an email from someone I don't know or when people I do not recognize are copied*" and another noted: "*Faces fast performance allows me to look for someone while they are calling me, just by quickly typing in their phone number*." One interviewee said "*When I go to meetings abroad, I look up the people I'm going to be meeting beforehand*" and another told us: "*I have once set up a call with someone I was already working with, and he suggested that I add more participants to the conference, so I wanted to know their role, location, and org chart relation to him. I found out that one was his manager, one was his employee, and one worked in his lab on this topic."* Hearing someone's name during a meeting was also mentioned as a common use case. One interviewee said: "*I often hear a person's name or nickname in a meeting, do not get the full name, but try to find them with the first name and keywords from the context of the discussion.*" Other interesting scenarios described were: "*I use Faces when I hear about an organizational change and want to better understand it*" and "*I recommend it to every newcomer in my team and it is key to their integration*."

We next provide an in-depth analysis of the use of *Faces* across queries, users, clicks, and actions.

### Queries

Since *Faces* updates results as the user types, each character can lead to a new query of the backend. However, in our analysis we were interested in the final string of characters created once the user stops typing. We therefore merged successive query log entries by the same user that had an edit distance of at most 3 between their respective query strings. The *edit distance* between two strings is the number of operations required to transform one to another; we specifically used the Levenshtein distance definition. This merging method is required since users sometimes delete and/or enter characters in the middle of the query string (commonly to fix a typo). The value of 3 is needed since not every character edit is logged (due to fast typing). The rest of the section refers solely to the merged queries.

Our dataset included a total of 1,119,121 queries. There were 668,084 unique query strings (59.7%), a much higher portion than the 31.7% reported in [20]. This can be attributed to the fact that users do not need to complete the whole string or correct spelling mistakes due to the dynamic result updates and the support for partial matching. Of the queries, 310,587 (27.75%) ended with a click on a result; we refer to these as *clicked queries*. We note that the fact that a click was not made, does not mean the user did not find the desired result, since many of the details appear inline. Weerkamp et al. [20] reported a lower click ratio of 17% and mentioned that web search engines have reported a substantially higher ratio, between 50% and 87%.

Figure 4 shows the distribution of number of characters per query for clicked queries. The most common query length was 6 characters (11.3% of the queries), implying a low effort to get the desired result. The average was 8.94 (median: 7, max: 126). In our survey, we asked participants whether they use copy-paste of text for their queries rather than typing the text themselves. 3.2% indicated they always use copy-paste, 12% chose 'often', and 17.1% indicated they use it for about half of their queries. The majority (51.4%) selected 'sometimes, but not too often' and 16.3% chose 'never'. We conclude that while copy-paste is not the prevalent way for querying, it is used from time to time by most users. Hence, the number of characters for manually-typed queries is likely to be even lower. One interviewee said: "*I often copy names or email addresses from email or chat messages I get or from calendar meetings*."

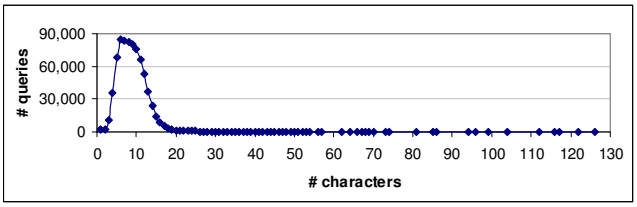


**Figure 4. Distribution of number of characters per query.**

Inspecting the clicked queries, the most common number of tokens per query was 2 with 53.3%. 41.3% contained one token, 4.9% had 3 tokens, and the rest (0.5%) had more. 7.28% of the queries included a keyword, i.e., a token that is neither a name nor contact information (email, phone number, etc.) This portion is higher than the 3.9% reported for the commercial engine [20]. It could be explained by the fact that in the organization, people are associated with more attributes, such as the organizational unit or the job description. The most popular keyword was 'manager' (appeared in 678 queries), followed by 'sales' (481), 'project' (278), 'business' (264), 'software' (242), 'marketing' (214), and 'testing' (206).

The upper rows of Table 1 show the distribution of number of characters per token (for clicked queries). The most common token lengths were 4, 5, and 6 characters (median is 5). 4% of the tokens had 9 or more characters. The lower

row of the table indicates, per each token length, the percentage of tokens that exactly match their corresponding field of the resulting person. Match percentage is low when token length is 1 or 2 characters and gradually increases as token length grows. For tokens with 9 or more characters it decreases slightly, probably since the likelihood for misspelling grows for lengthy tokens. Overall, 35.6% of the tokens did not have an exact match, indicating the support for partial matching and fuzzy search is indeed of need. One of our interviewees made an interesting comment in this context: "*Working in a global environment exposes us to names we have hard time to spell or even pronounce. Faces breaks the barrier of culture within the organization by allowing us to easily locate international colleagues.*"

| Length | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | ≥9 |
|---|---|---|---|---|---|---|---|---|---|
| Dist | 5.6 | 7.4 | 12.9 | 17.9 | 20.0 | 16.2 | 10.7 | 5.3 | 4.0 |
| Match | 8.2 | 18.8 | 42.0 | 66.1 | 78.5 | 81.0 | 82.9 | 84.7 | 80.3 |

**Table 1. Character length distribution (in %) and respective exact match percentage across all tokens of clicked queries.**

Separately inspecting the tokens of one, two, and three-token queries (denoted 1TQ, 2TQ, and 3TQ, respectively), the first token of the query was typically longer (median 6 for 1TQ and 5 for 2TQ and 3TQ) and with a higher exact match (73.6% for 1TQ, 85.5% for 2TQ, and 82.3% for 3TQ). 1- or 2-character strings were very uncommon (less than 3%) for the first token. By contrast, 27.2% of the second tokens in 2TQ were of 1 or 2 characters. The median length was 4 and only 38.1% exactly match their respective field. Similarly, the second and third tokens of 3TQ were of median 3 and 4, and with exact match of 51.5% and 40.6%, respectively. These numbers imply that users often provided one full token and only a subset for the second and third. One interviewee noted: "*I usually need to provide a name with another hint of character or two before I find the result*" and another said: "*I often don't remember correctly long family names, so it's a blessing Faces can suggest results after typing the first few letters.*"

We next examine the types of query tokens to understand what person fields were used as search criteria. For each clicked query, we matched each of its tokens with a field of the resulting person. Matching was done even when the token did not equal the field, based on substrings and phonetic resemblance, using the same mechanism used by *Faces* at runtime. Overall, we inspected the following field types: first name (FN), last name (LN), middle name (MN), email (eM), job title (Job), location (Loc), tag associated with the user (Tag), and organizational unit (Org).

The vast majority of the tokens referred to first names and last names, as can be seen in the upper rows of Table 2. Out of the keyword fields, Job was most common (1.7%), followed by Location (1.3%). Employee ID (0.3%) and phone number (0.13%) appeared in a small portion of the tokens and are not shown on the table. The lower row of Table 2 shows the percentage of tokens of that type for which there was an exact match. Names had a higher match percentage than keywords. First name had a particularly high portion of exact matches. Our participants indicated they occasionally copy and paste a name or an email to search *Faces*. This may contribute to the high percentage of exact matches for these fields.

| Type | FN | LN | MN | eM | Job | Loc | Tag | Org |
|---|---|---|---|---|---|---|---|---|
| Dist | 47.1 | 42.6 | 4.0 | 1.6 | 1.7 | 1.3 | 0.8 | 0.5 |
| Match | 78.6 | 54.8 | 55.3 | 49.8 | 18.6 | 27.3 | 46.0 | 15.6 |

**Table 2. Token type distribution (in %) and respective exact match percentage across all tokens of clicked queries.**

Table 3 depicts the breakdown of type distribution by tokens in queries of length 1, 2, and 3. For 1TQ, most tokens were evenly distributed between first name and last name. For 2TQ, first name was predominant as the first token, while last name was prevalent as the second. Email was mostly used as a single token, since it is a unique identifier. Keywords were used more often as part of multiple-token queries, especially in three token queries. Job and Tag were spread evenly across the tokens, while Location and Org tended to appear in the second or third, as a means to further narrow down the search results, as one interviewee noted: "*If I don't have the person's exact name, or can't spell it, I try to add hints like location or division".*

| Type | FN | LN | MN | eM | Job | Loc | Tag | Org |
|---|---|---|---|---|---|---|---|---|
| 1TQ 1st | **44.7** | **46.1** | 1.2 | 4.6 | 1.1 | 0.4 | 0.5 | 0.2 |
| 2TQ 1st | **83.7** | 12.3 | 0.6 | 0.3 | 1.5 | 0.5 | 0.7 | 0.3 |
| 2TQ 2nd | 14.8 | **73.9** | 5.4 | 0.8 | 1.6 | 2.0 | 0.7 | 0.4 |
| 3TQ 1st | **77.3** | 10.7 | 1.0 | 0.2 | 5.3 | 1.5 | 2.4 | 1.5 |
| 3TQ 2nd | 20.5 | 17.0 | **45.3** | 1.6 | 6.3 | 3.5 | 2.4 | 3.0 |
| 3TQ 3rd | 7.0 | **62.7** | 12.1 | 1.4 | 5.2 | 7.0 | 2.1 | 2.5 |

**Table 3. Token type distribution (in %) for queries of length 1, 2, and 3.**

### Users

For our study period, 265,940 people used *Faces*, according to distinct IP addresses. There were 36,663 distinct identified users, who originated from 82 countries. The U.S. was by far the most common (21.9%), followed by India (9.9%), Brazil (8.9%), Germany (7.2%), UK (6.6%), and Japan (5.1%). Users spanned the different divisions in our organizations: 41.9% were from Services, 23.4% Sales, 9.3% Operations, 8.9% Software, 8.1% Headquarters, 6.7% Systems, and 1.7% Research. The portion of users from Services and Sales is higher than their relative division sizes. Indeed, they typically need to work with a wider set of individuals and teams, compared to, say, software developers or system engineers.

Figure 5 (left plot) shows the distribution of the number of queries for identified users, which follows a power law with slope $\alpha$=*-1.93*. The average number of queries per user, for identified users, was 16.92 (stdev: 30.74, median: 5, max: 864). This is a substantially higher utilization, compared to the commercial engine, where only 21.65% of the users

issued more than one query over a similar analysis period [20]. Our numbers may be even higher in practice, as users may have not been identified for all queries they issued.

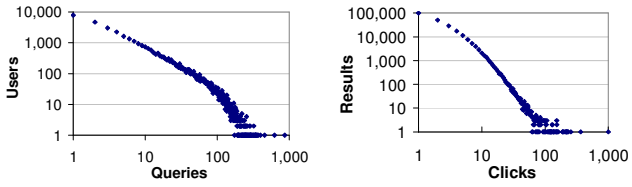


**Figure 5. Distribution of (Left:) queries over users and (Right:) result frequencies.**

Figure 6 shows the percentage of managers within the top *Faces* users. Top usage is determined according to the number of queries issued by a user along the time period. The corresponding number of queries by the last person on the list is shown in brackets. For example, the top 100 include users with 209 queries or more. The "All Users" bar shows the percentage of managers for all identified users, while the "All IBM" bar shows the percentage of managers within the entire organization, for reference. Clearly, the percentage of managers is higher among heavier *Faces* users (except for an anomaly between the top 100 and top 200). It gradually drops from 25.7% for the top 200 to 16% for the top 10,000, and 14.5% for all 36,663 users, which is still higher than the general percentage within IBM (12.9%). Overall, it is evident that managers use *Faces* more frequently. This can be explained by the fact that managers typically need to communicate with a large numbers of people, receive more email, and participate in more meetings, making people search more fundamental to their job.

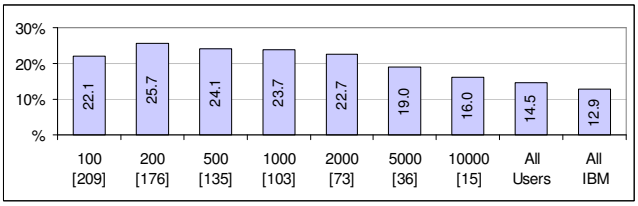


**Figure 6. Percentage of managers within top *Faces* users.**

## Clicks

The overall number of distinct employees who were clicked within *Faces* over the analysis period, either as a result or through the organizational environment, was 274,390. The overall number of employees clicked at least once as a direct result of a query was 243,354. The distribution of these employees according to the number of times they were clicked as a result is depicted in Figure 5 (right plot). It follows a power law with slope $\alpha$=*-2.49*. The average rank of a clicked result was 1.82 for identified users and 2.34 for non-identified users, indicating that the personalization of results indeed improves ranking. Even without personalization, the relevant person is typically very high on the list of results once the user clicks.

Table 4 shows the top 10 most popular resulting employees, in terms of *click counts* (referring to clicks on search results only, excluding clicks from the organizational environment) and distinct users who clicked. Per each person result, the table also shows the level in the organizational hierarchy (i.e., the number of managers up the organizational chain) and the number of people managed, either directly or indirectly (i.e., the number of direct or indirect children in the organizational tree). The company's CEO is by far the most popular, while the rest on the list are also very senior executives, either senior vice presidents (SVP) or general managers (GM). All have a high level in the organizational hierarchy and a large number of people managed. It should be noted that these senior executives do not lead the list of *Faces* frequent users and, in fact, most did not use *Faces* at all during the inspected period.

| Role | Level | #managed | Count | Users |
|---|---|---|---|---|
| CEO | 0 | 156,301 | 1149 | 502 |
| SVP, Software & Systems | 1 | 48,544 | 365 | 190 |
| SVP, Sales & Marketing | 1 | 15,573 | 256 | 158 |
| GM, Growth Markets | 0 | 108,101 | 231 | 108 |
| GM, India / South Asia | 0 | 111,336 | 230 | 77 |
| GM, Germany | 0 | 22,503 | 221 | 126 |
| SVP, Services | 1 | 72,778 | 220 | 116 |
| GM, Brazil | 3 | 17,662 | 217 | 104 |
| GM, Dev. & Manufacturing | 3 | 19,131 | 215 | 87 |
| SVP, Software Solutions | 2 | 6,368 | 211 | 114 |

**Table 4. Ten most popular results.**

Figure 7 displays the average level and number of people managed for the most popular *Faces* results. Popularity was determined by click count; the number in brackets is the corresponding click count for the last person on the list. For example, the top 100 most popular results had 116 clicks or more. The trends are generally very clear: popular results tend to be of senior people who are high up the hierarchy and manage many employees. The average level of the top 100 results is 2.6 and gradually goes down the hierarchy to 4.8 for the top 10,000, which is still much higher than the organizational average at 6.2. The average number of people managed goes down from over 11,000 for the top 100 (median: 2094) to 266 for the top 10,000 (median: 8). It should also be noted that the percentage of managers among the top results is very high: 93.6% for the top 100 and 60.7% for the top 10,000.

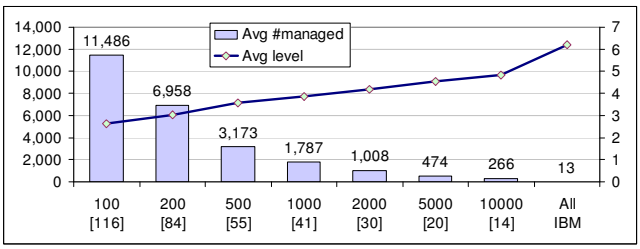


**Figure 7. Average level and number of people managed for top *Faces* results. The right Y axis refers to the level value.**

We next analyze the relationship of users to the people they searched for. We first examine how many of a person's searches fell among her top 150 familiar people, as calculated by SONAR. Overall, only 7.5% of the searches were of people among the user's top 150 familiar people. In our survey, a somewhat different picture was received. The vast majority of participants (74.5%) chose 'about equally strangers and people I know' in response to a question about whom they search for more often. 15.9% chose 'mostly strangers', 8.8% chose 'mostly people I know', 0.8% chose 'only strangers' and 0 chose 'only people I know'. Comments made by our interviewees shed more light on the reason for this gap. One said "*I know many people across the company, so most people I search for I vaguely know or remember from before*" and another mentioned: "*Many of the people I don't personally know, but rather know of […] for example, people who were on meetings with me or whom I previously searched*". Apparently, the 150 people returned by SONAR represent the employee's stronger ties within the organization, but typically the number of people they are familiar with, or are aware of, within the organization, may be much higher.

While users mostly searched for strangers or faint ties, we found that these often belonged to the same country and division. 74.7% of the searches were for people in the same country. In 30% of the cases, these were people from the same office location. Additionally, 56.9% of the searches were within the same organizational division.

While searching for strong ties was rare, self searching was rather common. 43.8% of the identified users searched for themselves at least once. Of these, 65% searched for themselves exactly once, while 6.9% did so 5 times or more. The most ardent self searcher did this 336 times. Overall, 31,202 queries were self-directed (2.79% of all queries), 18,390 of which ended with a click on self. Interestingly, this is a substantially higher click ratio than for all queries (58.94% vs. 27.75%).

In our survey, we asked participants about the single most important piece of information, once they find the person they are searching for. Figure 8 shows the distribution of answers. Contact information – phone number (24.5%) and email address (20%) – is at the top of the list. The organizational environment is also among the most important, with 20.1%, as demonstrated by this comment: "*I look at the peers to see if I know someone on the team they are on.*" And also: "*I often examine the management chain to understand to which organization they belong and who is the executive in charge.*" Job description was selected by 11% ("*I first look at the job title to understand their role, which often explains why they were mentioned to me or copied on an email I got*"), while the person's full name was selected by 10% ("*If I did not remember the exact name, this is the first thing I typically want to know.*")

Photo was selected by 7.7% of the participants. One interviewee told us: "*I am a very visual beast, I remember people by picture rather than by name […] then complete the original sense I get when I see the face by seeing the other attributes.*" Location and organizational unit were selected by very few participants as most important fields, and are perceived by most as complementary information. One participant wrote: "*I often look up a person's photo who works in the same building when talking to them the first time, then walk to their location and find them by comparing the photo to real faces.*"

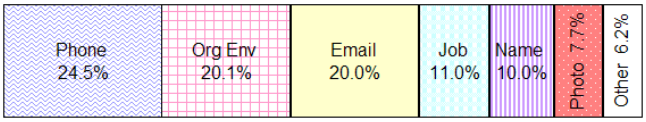


**Figure 8. Most important information based on survey results.**

Some of the participants mentioned that the selection depends on the search scenario. "*If I know the name, I will be looking for email or phone, but if I know the [job] title, I want the name*" wrote one participant and another commented: "*When I look for someone I know, I'm usually after their contact info, and when I look for strangers, I'm after all the other info: organizational structure, location, division, photo, etc.*"

### Actions

The overall number of actions within *Faces* along the time period was 2,487,518. Figure 9 details the different action types and their distribution of occurrence (percentage out of the total number of actions). Clicks on query results constituted 30.4% of all actions. Clicks on the organizational environment (peers, managers, and direct reports) were even higher at 34.9% overall. Most common were clicks on peers (16.8%), followed by managers (11.5%) and then reports (6.6%). It is interesting that peers were explored more often than the management chain; this could be due to fact that a person typically has more peers than managers. The lower portion of direct report clicks is expected since this option exists only for managers.

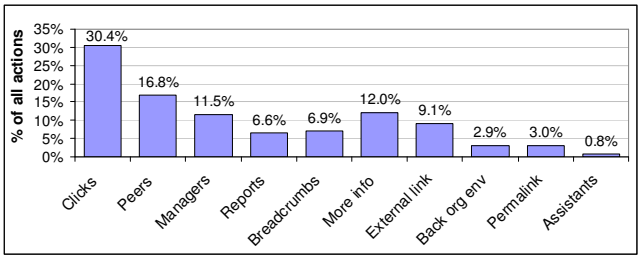


**Figure 9. Action distribution by type.**

The breadcrumbs within the organizational environment view were clicked in 6.9% of all cases, indicating that users often navigated back and forth across the organizational environment of the searched person. Overall, including the breadcrumbs, nearly 41% of the actions relate to exploration of the organizational environment, telling us that easy navigation across the organizational environment

is indeed a highly important feature for enterprise people searching. Clicking on "more info" to see additional details about the searched person was another common action (12%), leading to clicks on external links, either to the CD or the SNS (9.1% in total). Going back to organizational environment view ("back org env") was less frequent (2.9%).

The average number of actions per clicked query was 5.55 (stdev: 10.66, median: 3, max: 881). 28.44% of the clicked queries had only one action (the click itself), while 21.56% had 5 or more. Overall, we observed a high level of diverse activity beyond the click itself, especially around exploration of the organizational environment.

We also compared how managers use *Faces* versus regular employees. Table 5 summarizes the comparison across different aspects related to queries, clicks, and actions. All differences were found to be statistically significant. First, the average number of searches by managers was higher than by employees, again indicating that managers use *Faces* more frequently. A higher percentage of the managers searched for themselves. A comment by an interviewee, who manages a department of a few hundred people, provides insight for this difference: "*I often need someone within my own organization, so I start by searching my own name and then dig into my team [through the organizational environment] to find them*". Indeed, while 14.9% of the managers searched themselves 10 times or more, only 8.1% of the employees did so. Keyword tokens were used more often by employees. The only type of keyword used more frequently by managers was organizational unit, supposedly due to better organizational orientation.

Managers searched for people who were higher up the hierarchy (i.e., average level of a clicked result has smaller value than for employees). The difference was not very large, as employees were also looking for people who are high up the organizational chain. The average percentage of results among the top 150 familiar people was higher for managers than for employees, presumably since managers search more often within their own hierarchy. Finally, managers carried out more actions per clicked result than employees, indicating that they conduct more exploration of the organizational environment than regular employees.

| | Searches | Self | Keywords | Level | Familiar | Actions |
|---|---|---|---|---|---|---|
| M | 20.82 | 48.4% | 6.6% | 5.47 | 10.3% | 5.84 |
| E | 16.26 | 43.0% | 7.4% | 5.76 | 7.0% | 5.50 |

**Table 5. Manager (M) vs. Employee (E) usage of *Faces*.**

## DISCUSSION AND FUTURE DIRECTIONS

Our analysis indicates that enterprise people search, as implemented by *Faces*, is a very important tool for the workforce in a large enterprise. Employees typically use it on a frequent basis to quickly find the details of another employee based on partial information they have. The number of queries per user is much higher than reported for the commercial engine [20], indicating that people search in the enterprise is a more fundamental action. The information sought for is diverse: users pointed out phone number, email, job title, photo, location, division, and organizational environment as valuable information pieces in different scenarios. For example, phone and email are essential when looking to contact a known person; job title, photo, and organizational environment are important when looking up the details of a stranger or a distant tie. The search criteria also vary accordingly, although the vast majority of searches are done by name. However, if name information is not available or has been forgotten, other attributes, such as location, job title, or a tag, come to aid.

Most employees used *Faces* to search for individuals they barely know or complete strangers, often within the same country or division. The overall user-base of *Faces* was very broad, both in terms of unique users and unique results, while managers were found to use it more often. Even more prominently, the popular searches were mostly for managers and senior executives, as opposed to 'regular' employees. This suggests that boosting results by hierarchy level or number of people managed can be effective.

In their study of the logs for the commercial people search engine, Weerkamp et al. [20] found that 35% of the sessions are "repetitive sessions", where the user repeatedly made small changes in a name to correct its spelling. One of the most prominent features of *Faces* is its dynamic result update and support for partial matching. Users indicated these features are of great value, allowing them to find a person even if they remember just a few letters of the name, cannot figure out the correct spelling, or misheard it. This was said to be especially important for long or foreign names. Our results indicate that *Faces* queries are indeed typically short (median of 7 characters). Over 35% of the tokens did not precisely match a field of the clicked result, indicating the practical application of these features.

Even while using *Faces* for short, focused transactions, users took advantage of the interface to extensively explore the organizational environment by moving along and across the organizational chart. Overall, more than half of the clicks originated from the organizational environment view rather than the search results. This combination of "search and browse" behavior bears some similarity to web search. These findings support the centrality of organizational navigation in *Faces* design. One of the suggestions for improving *Faces* referred to searching by social network: "*I would like to be able to find Collean who works with Ajamu Smith*" wrote one participant. The org chart navigation capability provides partial support for this already, as another interviewee explained: "*I sometimes make my way to the person I need by actually searching for someone I know who is close in the hierarchy*". Showing the user's

network beyond the formal organizational hierarchy can further support this kind of navigation.

The most common feature request was integration with other business tools, such as email and instant messaging: "*I would like to be able to instantly open a chat window with the person I found*." A few participants suggested adding instant messaging availability as a search criterion (e.g., search for online people only). Another common comment referred to better enablement for mobile devices. "*Faces would make a perfect mobile app, as I need to find people on the go*", wrote one participant. Another repeating remark was that tags do not cover all expertise areas, so adding more keywords based on the employee's content can be of value. Other interesting ideas included: "*I would like queries to be saved, so I can access my history of searches*" and "*Showing Twitter ID or LinkedIn ID would be handy*."

Throughout our analysis we pointed out the similarities and differences between *Faces* and the commercial people search engine [20]: in both, one-query sessions are popular and search by name is dominant, but in *Faces* keywords are more common; click ratio is low in both engines, compared to web search engines, however, it is higher in *Faces*; the percentage of unique query strings is expectedly higher in *Faces* as well as the average number of queries per user. Obviously, some of these stem from the different designs of the tools. Overall, however, both people search engines have unique characteristics that distinguish them from other search engines, e.g., the prevalence of one-query sessions, low click ratio, and massive use of names for querying.

This study lays the foundation for further in-depth exploration of enterprise people search. Comparing search characteristics among job roles (e.g., sales people vs. developers) can be one direction of interest. Further examination of self searches, to examine to what degree they are followed by exploration of the organizational environment, can also be of value. Inspecting more closely the level of user repeated searches is also important and can support the request for saving search history, as expressed by one participant. Finally, the volume of searches for an individual can serve as an indication of her influence or reputation in the enterprise. It would be interesting to perform an in-depth exploration of who appears high on the list of popular searches, especially when those individuals do not hold senior positions (e.g., we found that 16 of the top 200 searched people were not managers).

Our analysis is based on experimentation in a single enterprise and is thus affected by its unique characteristics, such as the organizational structure and distribution of employees across countries and divisions. We believe, however, that our organization is a good representative of a large global enterprise. Given the importance of enterprise people search, we call for future research to study the topic further and conduct experiments in other organizations.

## REFERENCES


1. Amitay, E., Carmel, D., Har'el, N., Soffer, A., Golbandi, N., Ofek-Koifman, S., & Yogev, S. 2009. Social search and discovery using a unified approach. *Proc. HT '09*, 199-208.
2. Broder, A. 2002. A taxonomy of web search. *SIGIR Forum* 36, 2 (September 2002), 3-10.
3. Ehrlich, K. & Shami, N.S. 2008. Searching for expertise. *Proc. CHI '08*, 1093-1096.
4. Farrell, S., Lau T., Nusser S., Wilcox E., & Muller, M. 2007. Socially augmenting employee profiles with people-tagging. *Proc. UIST '07*, 91-100.
5. Guy, I., Jacovi, M., Shahar, E., Meshulam, N., Soroka, V., & Farrell, S. 2008. Harvesting with SONAR: the value of aggregating social network information. *Proc. CHI '08*, 1017-1026.
6. Guy, I., Ur, S., Ronen, I., Perer, A., & Jacovi, M. 2011. Do you want to know?: recommending strangers in the enterprise. *Proc. CSCW '11*, 285-294.
7. Hawking, D. 2004. Challenges in enterprise search. *Proc. ADC '04*, 15-24.
8. Huurnink B., Hollink L., van den Heuvel W., & de Rijke, M. 2010. Search behavior of media professionals at an audiovisual archive: a transaction log analysis. *JASIST* 61, 6 (June 2010), 1180-1197.
9. Jansen, B.J. & Spink, A. 2006. How are we searching the world wide web?: a comparison of nine search engine transaction logs. *Inf. Process. Manage.* 42, 1 (Jan. 2006), 248-263.
10. Kautz, H., Selman, B., & Shah, M. 1997. Referral Web: combining social networks and collaborative filtering. *Commun. ACM* 40, 3 (Mar. 1997), 63-65.
11. Kazai, G. & Doucet, A. 2008. Overview of the INEX 2007 Book Search Track (BookSearch'07). *SIGIR Forum* 42, 2–15.
12. Lawrence S., Bollacker K., & Giles, C.L. 1999. Indexing and retrieval of scientific literature. *Proc. CIKM '99*, 139-146.
13. McDonald, D.W. & Ackerman M.S. 1998. Just talk to me: a field study of expertise location. *Proc. CSCW '98*, 315-324.
14. McDonald, D. W. & Ackerman, M. S. 2000. Expertise recommender: a flexible recommendation system and architecture. *Proc. CSCW '00*, 231-240.
15. Mishne, G. & de Rijke, M. 2006. A study of blog search. *Proc. ECIR '06*, 289–301.
16. Philips, L. 2000. The double metaphone search algorithm. *C/C++ Users J.* 18, 6 (June 2000), 38-43.
17. Reichling, T., Veith, M., & Wulf, V. 2007. Expert recommender: designing for a network organization. *JCSCW* 16, 4-5 (October 2007), 431-465.
18. Silverstein, C., Marais, H., Henzinger, M., & Moricz, M. 1999. Analysis of a very large web search engine query log. *SIGIR Forum* 33, 1 (Sept.1999), 6-12.
19. Smirnova, E. & Balog, K. 2011. A user-oriented model for expert finding. *Proc. ECIR '11*, 580-592.
20. Weerkamp, W., Berendsen, R., Kovachev, B., Meij, E., Balog, K., & de Rijke, M. 2011. People searching for people: analysis of a people search engine log. *Proc. SIGIR '11*, 45-54.
21. Yiman-Seid, D. & Kobsa, A. 2003. Expert-finding systems for organizations: Problem and domain analysis and the DEMOIR approach. *JOCEC* 13, 1 (2003), 1-24.