

Figure 6: Expertise location using WCI: web app/user interaction for the construction of a query.

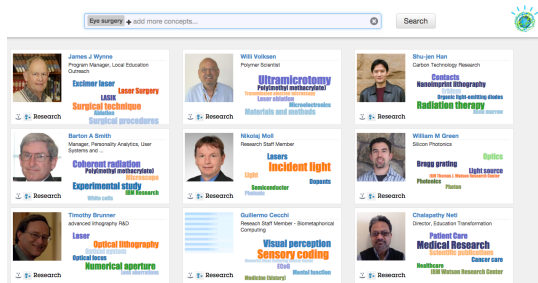


Figure 7: Expertise location using WCI: search results.

3.3 Expertise location with WCI

A basic approach to building an expert location system using WCI is to create a corpus that contains one document per person, describing the expertise of the person. For a large fraction of different expertise, it is generally possible to directly use at least part of the direct output from people. Due to the internal model for concept relationships WCI is effective at "filling the gaps," which results in deriving a picture of someone's expertise even if the description is not as complete. In the example shown in figures 6 and 7 we illustrate an expertise locator for IBM researchers. For this particular corpus, the use of the direct product of the expert's work as a document is particularly efficient (in this case the research publications). The actual data set consists of a profile page each researcher is maintaining together with the list of publications of each researcher.

The first problem to solve when interacting with the system is to help the end user construct a usable query in an intuitive way. The biggest help comes here from the fact that the concepts used in the system are identifiable by their English name (or in certain instances, short description). However not every English word or construct is a valid concept in the system therefore the chosen approach is to provide an intuitive auto-complete widget that proposes a list of the closest concept names together with a short abstract for the concept in order to enable disambiguation on sight for the end user (a task the average end user happens to be extremely good at). This is illustrated in Figure 6 where a query on quantum computing is being auto completed.

An example results page for the query eye surgery is shown in Figure 7. Note the explanation tags on this query: thanks to both WCI and the richness of the dataset highly related terms, potentially unknown or vaguely familiar to the user, show up. Experience shows that a common human reaction

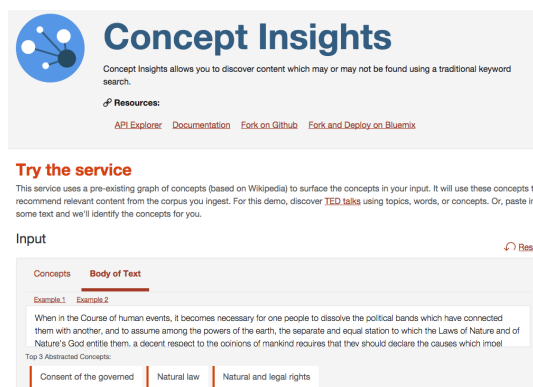


Figure 8: Content recommendation using WCI: constructing a query from textual user input.

to this is to proceed and explore the meaning of the terms (each one is clickable and links to Wikipedia).

3.4 Content Recommendation with WCI

An additional example is content recommendation. In this case we selected TED talks as a rich data set encompassing a large span of human interests. Recommendations can be driven by directly engaged user input or gathered textual evidence. In Figure 8 we show the content recommendation application gathering concepts from an excerpt of text. The most relevant three concepts are shown in the bottom left of the figure. We do not show the resulting recommendations due to lack of space.

4. CONCLUSIONS

We introduced Watson Concept Insights, an information retrieval system designed to foster end-user exploration of semantically and linguistically sparse corpora of documents. The system uses as key features mentions to concepts familiar to humans. The demonstration practically illustrates these characteristics focusing on both developer and end-user experience.

5. REFERENCES

- [1] R. T. Fielding. *Architectural styles and the design of network-based software architectures*. PhD thesis, University of California, Irvine, 2000.
- [2] I. Land, S. Huettinger, P. Hoehner, J. B. Huber, et al. Bounds on information combining. *Information Theory, IEEE Transactions on*, 51(2):612–619, 2005.
- [3] R. Mihalcea and A. Csomai. Wikify!: linking documents to encyclopedic knowledge. In *Proceedings of the sixteenth ACM conference on Conference on information and knowledge management*, pages 233–242. ACM, 2007.
- [4] D. Namiot and M. Sneps-Sneppe. On micro-services architecture. *International Journal of Open Information Technologies*, 2(9):24–27, 2014.
- [5] L. Page, S. Brin, R. Motwani, and T. Winograd. The pagerank citation ranking: bringing order to the web. 1999.
- [6] L. Soares, M. Franceschini, and L. Lastras. Watson concept insights api explorer, 2015. [Online; accessed 22-December-2015].