

Technical Disclosure Commons

Defensive Publications Series

October 2019

Page Navigation For Virtual Assistants Based On A Knowledge Graph

Thierry Coppey

Victor Carbune

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Coppey, Thierry and Carbune, Victor, "Page Navigation For Virtual Assistants Based On A Knowledge Graph", Technical Disclosure Commons, (October 25, 2019)

https://www.tdcommons.org/dpubs_series/2605



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Page Navigation For Virtual Assistants Based On A Knowledge Graph

ABSTRACT

A virtual assistant is designed to respond to user queries and enables users to navigate web pages and/or applications. However, full-text pages, e.g., pages of online encyclopedias or long-form journals, are not easy to navigate, especially via a voice interaction. Text is linear, includes reference points provided by authors, e.g., title, summary, etc., and, in raw form, is unsuitable for navigation via voice. Per the techniques of this disclosure, a text page or corpus is transformed, e.g., by the use of relation-extraction technique, into a distilled knowledge representation, e.g., a knowledge graph that includes relates the concepts on the text page with other concepts. A virtual assistant surfaces the distilled contents to the user in a manner that is more natural than a linear readout of the page. Navigation to relevant information is speeded up, and opportunities are offered to the user to explore related content.

KEYWORDS

- Knowledge database
- Knowledge graph
- Text search
- Natural language processing (NLP)
- Entity mining
- Relation extraction
- Open information extraction
- Knowledge distillation
- Page navigation

BACKGROUND

A virtual assistant is designed to respond to user queries and enables users to navigate web pages and/or applications. However, full-text pages, e.g., pages of online encyclopedias or long-form journals, are not easy to navigate, especially via a voice interaction. Text is linear, includes reference points provided by authors, e.g., title, summary, etc., and, in raw form, is unsuitable for navigation via voice. The development of text-navigation aids, e.g., links to concepts, sectioning of text, etc. for such content takes manual effort.

Knowledge databases, which are a fixed form of structured information, do not readily identify relationships across articles; even text search within databases generally requires exact matches. Hyperlinks enable the cross-referencing of articles but are somewhat inflexible as they need to be explicitly encoded in the text.

DESCRIPTION

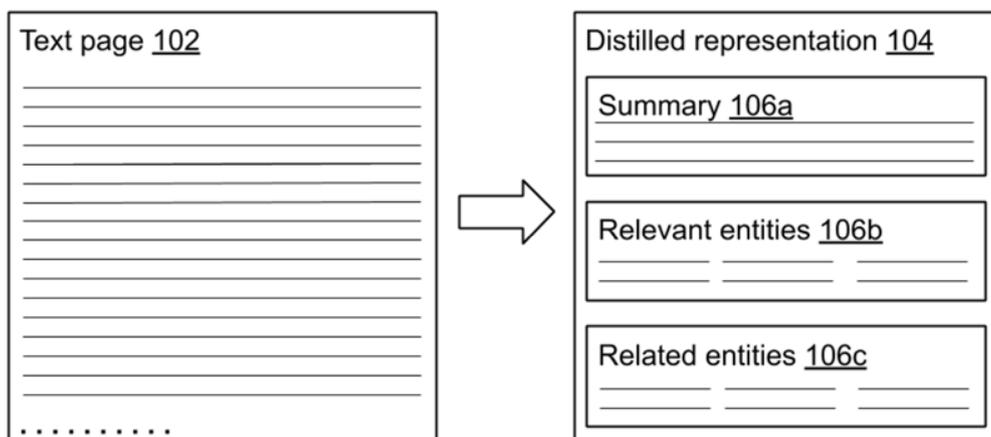


Fig. 1: Transformation from text to distilled knowledge representation

Per the techniques of this disclosure, illustrated in Fig. 1, a text page or corpus (102) is automatically transformed into a distilled knowledge representation (104) using relation-extraction techniques and a knowledge graph that is an overlay on knowledge databases that

reference common concepts across pages. Relation extraction (or entity mining) is a technique to create or enrich a knowledge graph using a knowledge database. The distilled knowledge representation can include a summary (106a), relevant entity/entities of the page (106b), a list of their related entities and relations (106c), etc.

- *Example:* An encyclopedia page on a film actress can include the following distilled knowledge representation. Summary: “American actress”; relevant entities: filmography, spouse, co-stars, etc.; related entities: “appearing in a movie,” “starred in a series,” “married,” “divorced,” etc.
- *Example:* An encyclopedia entry on a castle can have following distilled knowledge representation. Summary: “Former seat of X kingdom”; relevant entities: “King X, home, 1900-1915,” etc.; related entities: “King X, Y, daughter,” etc.

Navigation of the page content is enabled using one or more of the following interaction modes. (Examples refer to an encyclopedia page for an actress.)

- *Explain an entity*, e.g., go to a page or a snippet that states the summary, to gather more details about the entity.
- *Move to another entity*, e.g., from the actress page, move to a co-star.
- *Explore relations with another entity*, e.g., determine relations that connect the two entities. For the page on an actress, an example of a leading query would be “have they appeared together in a movie?”
- *Expand a relationship*, e.g., “appearing in a movie” can expand to the filmography of the actress.

- *Explain a relationship/entity*, e.g., from related-entity “married” is extracted a snippet of the page that explains the marital status, possibly restricting to a single entity based on user selection.
- *Find other pages with similar relationships*, e.g., when the user is on the "actress plus marriage” page, they might evince interest in the consistency with “spouse plus marriage” page, or similar pages.

The techniques of this disclosure enable a virtual assistant to surface facts from a page in a structured and summarized manner that is more natural than a linear readout of the page. Rather than manually sectioning the page or inserting hyperlinks in the text, the page is auto-navigated through the concepts identified in the page. Navigation by the user to the relevant information is thereby speeded up and more opportunities are offered to the user to explore related content.

The user can navigate content via a voice interface, where the virtual assistant speaks the navigation options and responds to queries for summaries, relevant links, relationships, etc. The navigation using the virtual assistant can also be via a displayed interface, e.g., tap-based navigation, such that when a user opens a page, relevant facts from the page, e.g., summaries, relevant links, relationships, etc. are visually presented.

To improve the latency of the response of the assistant, navigation can happen on-device once relevant facts from a page, e.g., summaries, relevant links, relationships, etc. are downloaded. The navigation can handle multimedia, e.g., pictures, sounds, videos (with some processing for interpretation by humans and machines), as well as text, which is denser and easier to process, but has a relatively narrow audience.

CONCLUSION

Per the techniques of this disclosure, a text page or corpus is transformed, e.g., by the use of relation-extraction technique, into a distilled knowledge representation, e.g., a knowledge graph that includes relates the concepts on the text page with other concepts. A virtual assistant surfaces the distilled contents to the user in a manner that is more natural than a linear readout of the page. Navigation to relevant information is speeded up, and opportunities are offered to the user to explore related content.