An Annotation System as an Abstraction Layer to Support Collaborative Knowledge Building

Isabela Chambers[1], Polyana Costa[1], Wallas Sousa[1], Rodrigo Costa[1], and Marcio Moreno[1]

IBM Research, Rio de Janeiro - RJ, Brazil
{ichambers,polyana.bezerra, wallas.sousa, rodrigo.costa}@ibm.com
{mmoreno}@br.ibm.com

Abstract. In this poster, we present an annotation system as an abstraction layer to enrich the collaborative knowledge creation and curation experiences by structuring data extracted from the exchanges between users, between users and AI services, and from users' input on content. It supports the definition of more meaningful relations between concepts and richer discussion processes among users, contributing to the expansion and evolution of knowledge bases that feed off the aforementioned structured data. It is also capable of yielding relevant results to semantic queries by which users can retrieve content and knowledge they contributed to creating. Our results show that users found this method of joint knowledge building to be useful and that it could optimize tasks, mainly because a) it allows access to fresh insights, correlations, and valuable knowledge exchange, and b) it supports data retrieval via semantic queries.

Keywords: Annotation Systems · Multimedia and Multimodal Retrieval · Hyperknowledge

1 Introduction

This poster[1] presents the ongoing work around the Hyperknowledge Annotation System (HAS), by describing a qualitative approach to understanding user needs following the speculative development of a proposed system. The work disclosed in this piece was elaborated in the context of the difficulties around interacting with knowledge bases to curate and enrich them. It can be a tiring and complex activity, especially for those who are not familiar with the field of knowledge engineering, which can be further complicated when there are multiple inputs from users from different backgrounds.

One of the many ways one may interact with such bases is by means of annotation systems. In general, these focus on one type of media (text, image, video or audio) [2] and allow users to collaborate by accessing annotations from other users, commenting on them, and curating them [3]. However, many and especially those directed at end-users , and not knowledge engineers do not support extracting abstract concepts from content fragments and their contexts, neither do they structure or store said data. Some of these systems do store the knowledge retrieved from the annotations on knowledge bases (such as triple stores) allowing queries over the saved content [2]. However, most are not friendly to users who are not in the habit of working with knowledge engineering, as they require direct manipulation of the knowledge bases and do not typically explore correlations in the annotated data to leverage knowledge structuring and allow for semantic queries.

Our proposed approach, the HAS, provides an abstraction layer that allows users to collaborate (with each other and with artificial intelligence services) when creating and curating knowledge to enrich knowledge bases. The system supports multimodal annotations over multimedia content segments so that annotators can use diverse types of content to create annotations, and the retrieval of information from the knowledge bases, which represents a reward for engaging in the activity in the first place.

It is all made possible by extracting and structuring concepts from annotations and their anchors (selected piece of content), as well as by offering suggestions through understanding annotators' discourse with the support of artificial intelligence (AI) algorithms. To structure the annotation content, the HAS uses its own conceptual model called Hyperknowledge [4], which allows semantic queries over the annotated data. We defined a use case scenario of research and development activities and sought to understand these users' pain points in their process of dealing with impressive amounts of data strewn across different types of content while creating, curating, organizing, storing, collaborating on, and retrieving data.

Our results showed that participants were able to make annotations, curate suggestions, understand how to collaborate, and make queries; understanding that in order to obtain results to their queries, the information needed to have been previously added to the base. When the sheer volume of information one operates with becomes ineffective to deal with in other methods such as keyword search, users stated that the HAS is a better alternative as it allows semantic queries that are a clear advantage to the process.

2 Background and Technical Aspects

The HAS systems design was first introduced by Moreno et al. [1], in which three main aspects of the system are defined: the multilayer architecture; the human-machine collaborative scope; and the effective integration of the annotation with the multimedia content via hyperknowledge, a knowledge representation model [4].

First, the architecture is composed of four layers, each defining a level of abstraction: layout structure layer; syntactic layer; semantic layer; and pragmatic layer. The layout structure layer supports information extraction in a document by identifying semantically related structure (e.g. bullets and headers). The syntactic layer handles the grammatical structure of sentences (e.g. identifying a substructure in a sentence as its subject). The semantic layer is in charge of specifying content meaning (e.g. of a given word or concept). Lastly, the pragmatic layer provides support to annotation at a natural language level (e.g. manually annotating a concept).

The human-machine collaborative scope relates automated annotations from AI services with users annotation. It captures part of the contextual information from a users annotation and provides an automatically generated annotation. The users annotation can be done on and in a range of media types, and to support that, the HAS establishes contracts between a media type and the appropriate AI services.

Fig. 1. shows an example of the aforementioned annotation interaction on an image, in which the white rectangle is the users annotation anchor referring to the concept of a player and the red rectangle refers the AI services output which identifies the player as Neymar. The media node $Image\_A$ represents the image you see to the left (Neymar playing). It contains one more anchor besides the default λ anchor. The anchor *anchor_1* can be linked with connectors of type *depicts* to nodes of type *instance* (*sprint_17*) of concept ($Move$). In this example, *sprint_17* is an instance of the class $Move$ and is linked to an anchor of $Image$. For the ontology in question, a $Move$ (such as *sprint_17*) is executed by a $Player$, which, in this case, is the instance $Neymar$. Finally, the facts are inside a context called $Match\_3$, but the nodes $Neymar$, *Player*, and *Move*, are in dashed lines, which indicates they are being reused. In other words, reusing allows entities that belong to different contexts to be linked without having to define them once more. How to proceed with the definition of the entities in contexts is up to their application. To structure the annotations and store them on the knowledge base, the conceptual model behind HAS - Hyperknowledge - uses domain-based ontologies. In this particular example, the chosen domain was soccer [5], but any other use case scenario could have been used, if an ontology that represents it was given.

3 User Tests

The methods we used to test user interactions and assess the system's value to them were small-scale and qualitative in nature, but enough to drive investigations into our two main questions: *(a)* Do people understand and perform according to their role of calibrating the AI algorithms and enriching the knowledge base? *(b)* Do users perceive the advantage in contributing to the system in order to reap the benefits of knowledge retrieval via semantic queries?

We interviewed scientists of different backgrounds, all of which engaged in research and development activities that require a lot of information to consume, analyze, share, and build upon, about which they often needed to retrieve specific information such as the temperature used in a specific experimental setting. In this case, digital or physical notes have to be associated with a digital

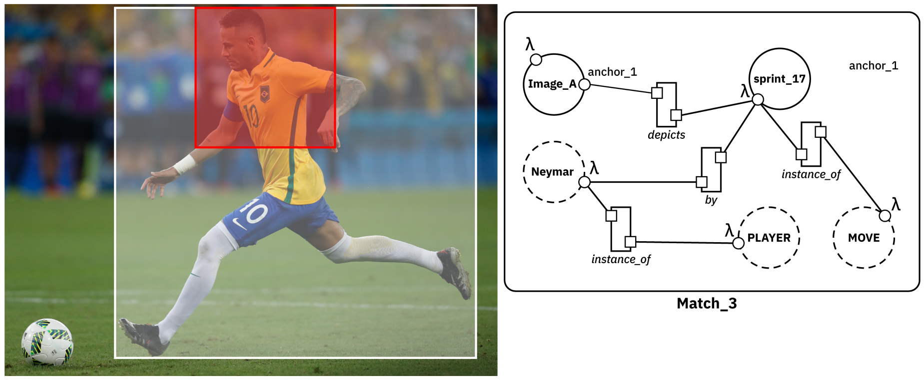

Fig. 1. Annotated image (left, soccer player) via the HAS; and the hyperknowledge model (right, graph) generated from the annotation.

image file. To find such information as the temperature used in experimental settings that had returned a particular type of result, they would have to parse physical documents in a binder; or, in the case of digital documents, type in the appropriate keywords in a document finder or software, and then look for the specific data among all the results that the keyword search returns. Saving a digital picture of the physical note does not help search for words that might be in the content itself, and not the file's metadata. In that sense, it would be of great benefit to them if they could use different media types to directly annotate on multimedia content.

In testing the HAS, they were instructed to simulate uploading a file (which, in this case, was an image), annotating on it, reviewing the suggestions made by the AI (in that case, IBM's Watson Image Recognition), saving that annotation, and then going over that annotation's details and properties in order to contribute to it via replies and further annotations. Finally, we gave them 2 minutes to freely pose the system with queries that they would like to be able to make, and they came up with twenty of them, a few of which were: show videos of test 4; show highlighted points of interest in a content; how did the colorimetric response of a given indicator vary over time. All of the queries they wished to make were feasably supported by the HAS, provided that the relevant data was present in the base and appropriately structured; that parameters for properties such as "colorimetric response" were defined; and that query inputs were adapted to one of the supported query languages, such as SPARQL [1].

4 Results

We were able to successfully answer the research questions posed beforehand in the following manner:

(a) Do people understand and perform according to their role of calibrating the AI algorithms and enriching the knowledge base? We were able to conclude

through observations and direct user quotes ("you have to keep in mind that all that you might want to ask it depends on what has been annotated") that users did indeed understand where inputs came from (themselves, AI suggestions, and mutual feedback/curation between users and between users and AI); their roles in providing these inputs and curating them; that the enrichment of the knowledge base depended on that; and that the possibility of querying did as well.

(b) Do users perceive the advantage in contributing to the system in order to reap the benefits of knowledge retrieval via semantic queries? We reached the conclusion that indeed they do. That same quote we highlighted as part of the answer to question above encapsulates a fundamental factor to answering this one. If they could understand and accept that they had to make an effort in order to be able to make queries, it is clear that that is something they wish to able to do. They realized that the queries afforded by the HAS could significantly optimize their process, and so they represent an advantage to which the effort required to contribute to the base did not seem disproportionate (and neither, as stated before, more laborious than their current process), especially as they greatly reduced the number of tasks required to get access to specific knowledge during a research and development project, in a way that contributes to reducing cognitive exhaustion.

Furthermore, users stated that working alongside other users and an AI annotator's suggestions in the HAS provided them with fresh insights into relationships between concepts, which helped them establish other correlations they might not have thought of, and which brought them new ideas to, in turn, bring to discussions, and to further enrich their knowledge building process (and even if they're not in direct contact with it, the corresponding knowledge base).

References

1. Moreno, M., Santos, W., Costa, R., Cerqueira, R.: Supporting Knowledge Creation through HAS: The Hyperknowledge Annotation System. In: IEEE International Symposium on Multimedia, 2018.
2. Takis, J., Islam, A. S., Lange, C., Auer, S.: Crowdsourced Semantic Annotation of Scientific Publications and Tabular Data in PDF. In: International Conference on Semantic Systems, 2015.
3. Stenetorp, P., Pyysalo, S., Topi, G., Ohta, T., Ananiadou, S., Tsujii, J.: BRAT: a Web-based Tool for NLP-Assisted Text Annotation. In: Conference of the European Chapter of the Association for Computational Linguistics, 2012.
4. Moreno, M., Brando, R., Cerqueira, R.: Extending Hypermedia Conceptual Models to Support Hyperknowledge Specifications. In: IEEE International Symposium on Multimedia, 2016.
5. Moreno, M., Santos, W., Santos, R., Ramos, I., Cerqueira, R.: Supporting Soccer Analytics through HyperKnowledge Specifications. In: 2019 Second International Conference on Artificial Intelligence for Industries (AI4I) (pp. 13-16). IEEE, 2019.