

Deep Learning Powered Question-Answering Framework for Organizations Digital Transformation

Nuno Ramos Carvalho
United Nations University (UNU-EGOV)
Rua Vila Flor 166,
4810-445 Guimarães
Portugal
ramos.de.carvalho@unu.edu

Luís Soares Barbosa
INESC TEC - University of Minho and United Nations
University (UNU-EGOV)
Rua Vila Flor 166, 4810-445 Guimarães
Portugal
barbosa@unu.edu

ABSTRACT

In the context of digital transformation by governments, the public sector and other organizations, many information is moving to digital platforms. Chatbots and similar question-answering systems are becoming popular to answer information queries, opposed to browsing online repositories or webpages. State-of-the-art approaches for these systems may be laborious to implement, hard to train and maintain, and also require a high level of expertise. This work explores the definition of a generic framework to systematically build question-answering systems. A sandbox implementation of this framework enables the deployment of turnkey systems, directly from already existing collections of documents. These systems can then be used to provide a question-answering system communication channel to enrich the organization digital presence.

CCS CONCEPTS

• **Applied computing** → **Computers in other domains** →
Computing in government → *E-government*

KEYWORDS

Deep learning, question-answering system, digital transformation

ACM Reference format:

N. Ramos Carvalho, L. Soares Barbosa. 2019. Deep Learning Powered Question-Answering Framework for Organizations Digital Transformation. In *Proceedings of the 12th International Conference on Theory and Practice of Electronic Governance (ICEGOV2019), Melbourne, VIC, Australia, April 3-5, 2019, 4 pages*.
<https://doi.org/10.1145/3326365.3326375>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

ICEGOV2019, April 3–5, 2019, Melbourne, VIC, Australia
© 2019 Association for Computing Machinery.
ACM ISBN 978-1-4503-6644-1/19/04...\$15.00
<https://doi.org/10.1145/3326365.3326375>

1. INTRODUCTION

Digital transformation of organizations encompasses an increased adoption of digital platforms, including online repositories of information available for searching and browsing. But with the increase of digitally available data and information, the effort of using it in an effective way has also increased [6].

Different communication channels, e.g. e-mail, social media, websites, mobile apps, smart devices, are being explored to enhance communication between organizations and the public. Chatbots are an approach that is currently increasing its' adoption [2]. This is mainly due to recent advances in the areas of: Machine Learning [7], Deep Learning [4], and Natural Language Processing [5], that allow the implementation of programs that are able to maintain a natural language conversation with an human, while addressing user queries or performing simple tasks. Some drawbacks of these approaches are: they may be laborious to implement, and specific models need to be trained on specific information for example, which requires a level of expertise and set of skills that may not be available in the IT departments of most public organizations.

The overarching goal of this work is to define a generic framework for implementing question-answering systems from a collection of documents, without additional processing, or model training. Please note that the term document here is used in an all-encompassing way, i.e. documents may refer to a Word or PDF document, or to the content of a web-page, or wiki page, a document simply refers to a well delimited piece of plain text. So, the problem being addressed by the framework is defined as: given a collection of documents and an information query, retrieve an answer to the query from some text in some document.

For example, one of the European Union (EU) websites (<https://europa.eu/european-union/topics/en>) provides a set of

documents on heterogeneous topics about the EU. An example information query that this information body is prone to address might be the expected cut in greenhouse emission for 2050, a possible question that captures this information query can be:

Q: “What is the cut in greenhouse gases for 2050?”

From the content of the text about the “Energy” topic, we find the following sentence: “An 80-95% cut in greenhouse gases compared with 1990 levels.”, which entails that the answer to the question can be “80-95%”. In a nutshell, the goal of the framework is to enable a system that can find this answer on itself from the collection of documents alone. Without any extra processing layer or other extra resources, like for example a knowledge base build from the documents collection.

Besides the framework definition, this work entails the definition of a sandbox implementation, i.e. an actual implementation of the functions defined in the framework that can be used in real experimentations.

The remainder of this paper is organized as follows: Section 2 describes the introduced framework in more detail including how to address the two different identified problems; Section 3 discusses the sandbox implementation of the framework; Section 4 introduces some related work; and, Section 5 ends the paper with some closing remarks, and ongoing work trends to address identified challenges.

2. THE FRAMEWORK

The generic goal of the framework is to provide a set of function definitions that when combined together, provide a question-answering system for a collection of documents. The top-most function is function q , that simply takes as input a question Q and provides an answer A 2:

$$q : Q \rightarrow A$$

Both question Q and answer A are strings of text. Using this approach, the problem of finding an answer to a specific question is divided in two distinct and independent steps:

- 1) From a list of documents and a question, retrieve a context, i.e. the document that is prone to contain the answer to the question;
- 2) And, from a context and a question, retrieve an answer to the question in natural language from the context.

The next sections discuss these steps in detail. The underlying functions act as placeholders in the generic framework, which means that any suitable existing techniques or approaches can instantiate them. Each sub-section also introduces the current technique, or approach, used in our proof-of-concept tool to implement this function in the sandbox framework.

2.1. FROM DOCUMENTS TO CONTEXT

Step 1) defined in the previous section is addressed by function f and is defined as: given a collection of documents $[D]$ and a

question Q , find the document (context) C that is prone to contain the answer to question Q :

$$f : [D] \times Q \rightarrow C$$

The intuition behind this approach is that the document that is more related to the question (semantically) has the higher probability to contain the context (snippet) of text that is able to answer the question.

2.2. FROM CONTEXT TO ANSWER

The result of applying function f defined in the previous section, is a document. Which is used in this step as a context, i.e. a snippet of text. The goal of Step 2) is to retrieve an answer to the question from the context. This is handled by function g whose goal is: given a context C (a document), and a question Q , retrieve an answer A for question Q .

$$g : C \times Q \rightarrow A$$

The intuition behind this step is that the document retrieved by function f from the collection, is the document more semantically related to the information query, and hence is the most prone to contain the snippet of text that answers the question.

Finally, function q can be defined as applying function g , after computing a context using function f :

$$q = g (f [D] Q) Q$$

Which means that the only requirements for the framework are a definition for functions f and g . The next section describes a possible instantiation for these functions that allow the implementation of a sandbox framework.

3. SANDBOX IMPLEMENTATION

The goal of the sandbox implementation is to provide a generic proof-of-concept realization of the previously described framework. There are two major placeholders in the framework: function f and function g , that compute the context and the answer to the question respectively.

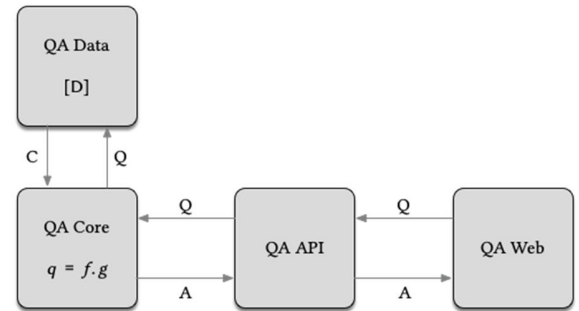


Figure 1: Question-answering framework generic architecture.

² The notation used for the formulas follows the Haskell functional programming language notation [18].

Figure 1 illustrates the sandbox implementation architecture with a high level of abstraction: QA Data is the collection of documents (there can be more than one collection), QA Core implements the core functions (f and g , and the query q function by composition of the functions f and g), QA API provides a REST interface (via HTTP) for the core library; and, QA Web implements a web interface to provide a user-friendly interaction with the core functions. The arrows connect the different components and illustrate the flow of the question from the web interface to the core functions, and the answer from the core functions back to the user via the web interface.

3.1. FUNCTION f

In the sandbox implementation of the framework, Latent Semantic Analysis (LSA) [3] is used as a placeholder for function f . In a nutshell, LSA is a technique for analyzing relations between words in documents, often used to measure the distance between a set of words (an information query for example), and a set of documents, to rank which documents are more related to the query. Its' often used in search engines to rank web pages related with an information query. This is a natural first approach for function f , i.e. find the document that is more semantically related to the question, the underlying rationale is that this document is more prone to contain the snippet of text that answers the question, than the semantically less related documents.

3.2. FUNCTION g

To find the answer to the question in the context, the sandbox implementation of the framework uses Bi-Directional Attention Flow (BiDAF) [9], a widely used Deep Learning comprehension baseline model, that achieves a state-of-the-art accuracy on the SQuAD dataset [8]. This model is trained to find an answer to a question in a (small) passage of natural text, the models' final prediction is the starting token (word), and the ending token (word) that answers the question, this means that the answer is the sequence of words (string) between these two tokens. The main advantage of similar approaches is that they do not require a pre-processing stage to build a knowledge base, triples collection of facts about the text, etc. the only requirement is the collection of documents.

The sandbox implementation of the framework is available at: <http://qa.egov.uminho.pt/>, where a web interface can be used to query the collection of documents mentioned earlier concerning EU topics.

4. RELATED WORK

Open domain question-answering systems are mainly concerned with retrieving answers to questions from passages of text, or documents. These systems are often explored to enable communication channels, like chatbots or virtual. These are currently being mainly used by private organizations, mainly to implement solutions with increases effectiveness at lower costs, some examples include: Apple Siri, Microsoft Cortana, Amazon Alexa, or Google Assistant. But local and central public

organizations around the world are also successfully exploring the use of virtual assistants approaches to enable public services, while reducing costs [12]. For example, the Los Angeles CHIP chatbot has achieved a 50% reduction in the e-mail volume that needs to be handled by city staff, by answering a sub-set of citizens questions [13].

Modern approaches for implementing state of the art question-answering systems use artificial neural networks of different architectures. Feng et al explore a framework based on Convolutional Neural Networks (CNN) to find answers to questions [14]; Yao et all use a stack of recurrent networks, an encoder-decoder architecture with an attention layer, to model a conversation [15]; Vinyals and Le discuss a sequence to sequence framework to implement a simple language model to train a conversational engine [16]. The currently available literature seems to show a trend for implementing question-answering systems, of moving from complex pipelines of tools and techniques, to end-to-end systems based on neural networks [17].

5. CONCLUSION

Question-answering systems, chatbots, and similar artefacts are currently being used by organizations to provide channels of communications with the public. In order to achieve state-of-the-art results, these can be laborious and/or complex to implement. This work describes ongoing research being undertaken to define a generic framework that allows the implementation of these systems.

Figure 2 illustrates the front web interface after placing some questions to the system using the sandbox framework implementation.

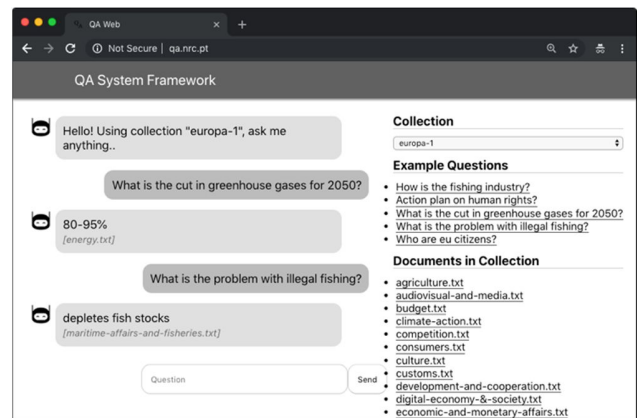


Figure 2: Web interface screenshot while using the sandbox implementation to query the EU collection of documents.

Some challenges that the framework still needs to address include: taking context in consideration between questions, or questions that depend on the previous question/answer; provide a pattern-match based approach for handling questions not related with the documents collection (e.g. greetings, "how are

you?”). Explore different implementations of functions f and g in the framework, and provide a scored rank of possible answers, is also in order.

This work is included in a broader project aiming at reducing administrative burden [1, 10], which is a major concern for governments in many countries. Typically, this encompasses initiatives around the re-engineering of administrative processes, including the provision of communication channels between citizens and public offices, where question-answering systems (e.g. chatbots) can play an important role, by providing a low-cost 24x7 service, that is able to instantly provide answers to simple information queries.

ACKNOWLEDGMENTS

This work is financed by the ERDF - European Regional Development Fund through the Operational Programme for Competitiveness and Internationalisation - COMPETE 2020 Programme and by National Funds through the Portuguese funding agency, FCT - Fundação para a Ciência e a Tecnologia, within project POCI-01-0145-FEDER-029946 (DaVinci).

REFERENCES

- [1] Carvalho, N.R. and Barbosa, L.S. 2018. Transforming Legal Documents for Visualization and Analysis. Proceedings of the 11th International Conference on Theory and Practice of Electronic Governance - ICEGOV '18 (2018), 23–26.
- [2] Drift SurveyMonkey Audience, S. and myclever 2018. The 2018 State of Chatbots Report.
- [3] Dumais, S.T. 2004. Latent semantic analysis. Annual Review of Information Science and Technology. 38, 1 (2004), 188–230. DOI:<https://doi.org/10.1002/aris.1440380105>.
- [4] Goodfellow, I. et al. 2016. Deep Learning. MIT Press.
- [5] Jurafsky, D. and Martin, J.H. 2009. Speech and Language Processing: An Introduction to Natural Language Processing, Computational Linguistics, and Speech Recognition. Speech and Language Processing An Introduction to Natural Language Processing Computational Linguistics and Speech Recognition. 21, (2009), 0–934. DOI:<https://doi.org/10.1162/089120100750105975>.
- [6] Keim Daniel, K.J. and Mansmann, G. rey E. and F. 2010. Mastering the Information Age Solving Problems with Visual Analytics.
- [7] Michalski, R.S. et al. 2013. Machine learning: An artificial intelligence approach. Springer Science & Business Media.
- [8] Rajpurkar, P. et al. 2016. SQuAD: 100,000+ Questions for Machine Comprehension of Text. (Jun. 2016).
- [9] Seo, M. et al. 2016. Bidirectional Attention Flow for Machine Comprehension. (Nov. 2016).
- [10] Veiga, L. et al. 2016. Digital Government and Administrative Burden Reduction. Proceedings of the 9th International Conference on Theory and Practice of Electronic Governance. (2016), 323–326. DOI:<https://doi.org/10.1145/2910019.2910107>.
- [11] Veiga, L. et al. 2016. Digital Government and Administrative Burden Reduction. Proceedings of the 9th International Conference on Theory and Practice of Electronic Governance. (2016), 323–326. DOI:<https://doi.org/10.1145/2910019.2910107>.
- [12] Bermejo Razquin, J., 2018. A friendly approach to Open Government: Chatbots.
- [13] Azure Government Chatbot Reduces Email By 50%. 2017. Retrieved from: <https://cfocussoftware.com/government-chatbots/azure-government-chatbot-reduces-email-50-percent/>.
- [14] Feng, M., Xiang, B., Glass, M.R., Wang, L. and Zhou, B., 2015, December. Applying deep learning to answer selection: A study and an open task. In Automatic Speech Recognition and Understanding (ASRU), 2015 IEEE Workshop on (pp. 813-820). IEEE.
- [15] Yao, K., Zweig, G. and Peng, B., 2015. Attention with intention for a neural network conversation model. arXiv preprint arXiv:1510.08565.
- [16] Vinyals, O. and Le, Q., 2015. A neural conversational model. arXiv preprint arXiv:1506.05869.
- [17] Sun, H., Dhingra, B., Zaheer, M., Mazaitis, K., Salakhutdinov, R. and Cohen, W., 2018. Open Domain Question Answering Using Early Fusion

- of Knowledge Bases and Text. In Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (pp. 4231-4242).
- [18] Hall, C. and O'Donnell, J., 2000. Introduction to Haskell. In Discrete Mathematics Using a Computer (pp. 1-33). Springer, London..