BioGen: Automated Biography Generation

Heer Ambavi*, Ayush Garg*, Ayush Garg*, Nitiksha*, Mridul Sharma*, Rohit Sharma*
Jayesh Choudhari, Mayank Singh*
Department of Computer Science and Engineering, Indian Institute of Technology Gandhinagar, India
singh.mayank@iitgn.ac.in

ABSTRACT

A biography of a person is the detailed description of several life events including his education, work, relationships and death. Wikipedia, the free web-based encyclopedia, consists of millions of manually curated biographies of eminent politicians, film and sports personalities, etc. However, manual curation efforts, even though efficient, suffers from significant delays. In this work, we propose an automatic biography generation framework BioGen. BioGen generates a short collection of biographical sentences clustered into multiple events of life. Evaluation results show that biographies generated by BioGen are significantly closer to manually written biographies in Wikipedia. A working model of this framework is available at *nlpbiogen.herokuapp.com/home/*

CCS CONCEPTS

•**Artificial intelligence** →**Natural language processing**;

KEYWORDS

Biography generation; English Wikipedia; Summarization

1 INTRODUCTION

As Internet technology continues to thrive, a large number of documents are continuously generated and published online. Online newspapers, for instance, publish articles describing important facts or life events of well-known personalities. However, these documents being highly unstructured and noisy, contain both meaningful biographical facts, as well as unrelated information to describe the person, for example, opinions, discussions, etc. Thus, extracting meaningful biographical sentences from a large pool of unstructured text is a challenging problem. Although humans manage to filter the desired information, manual inspection does not scale to very large document collections.

1.1 Textual Biographies

A textual biography can be represented as a series of facts and events that make up a person's life. Different types of biographical facts include aspects related to personal characteristics like date and place of birth and death, education, career, occupation, affiliation, and relationships. Overall, a general biography generation process can be described in three major steps: (i) identifying biographical sentences, (ii) classifying biographical sentences into different life-events, and (iii) relevancy-based ranking of sentences in each life-event class. Along with the biographical information, a Wikipedia profile of a person also consists of a consistently-formatted table on the top right-hand side of the page. This table or box is known

as an *infobox*, and it contains some important facts related to the person.

1.2 Machine Learning in Biography Generation

Majority of past literature focuses on Machine Learning techniques for information extraction. Zhou et al. [10] trained biographical and non-biographical sentence classifiers to categorize sentences. They also employ a Naive Bayes model with n-gram features to classify sentences into ten classes such as bio, fame factor, personality, etc. This work looks similar to ours, but this requires a good amount of human effort. Biadsy et al. [3] proposed summarization techniques to extract important information from multiple sentences. Liu et al. [7] also use multi-document summarization. For identifying salient information, the paragraphs are ranked and ordered using various extractive summarization techniques. However, both these systems ([3] and [7]) do not focus on sectionizing the biography. The works by Filatova et al. [5] and Barzilay et al. [2] focus on specific tasks such as identifying occupation related important events and sentence ordering techniques respectively. One of the recent works in generating sentences is by Lebret et al. [6]. They use *concept to text generation* approach to generate only single/first sentence using the fact tables present in Wikipedia.

1.3 Our Contribution

In this paper, we address the task of automatically extracting biographical facts from textual documents published on the Web. We pose this problem in the extractive summarization framework and propose a two-stage extractive strategy. In the first stage, sentences are classified into biographical facts or not. In the following stage, we classify biographical sentences into several life-event categories. Along with the biography generation task, we also propose a method to generate *Infobox* which is a consistently-formatted table mentioning some important facts and events related to a person. We experimented with several ML models and achieve significantly high F-scores.

Outline: Section 2 describes the datasets that are used to train the models. Section 3 describes the components of our system in more detail. Section 4 describes our experiments and results. Section 5 draws the final conclusion of our work.

2 DATASETS

The current work requires large textual biography datasets. Also, in order to categorize between biographical and non-biographical sentences, we leverage a non-biographical news dataset. Following is the descriptions of the dataset used.

TREC-RCV1 [6]: This Reuters news corpus consists of ~8.5 million news titles, links and timestamps collected between Jan 2007 and Aug 2016. Dataset was used for training a 2-class classifier in

---

the first step of the biography generation process (see Section 3.1). All the sentences in this dataset are labeled as *non-biographical*.

`WikiBio` [1]: This dataset consists of ~730K biographical pages from English Wikipedia. For each article, the dataset consists of only the first paragraph. This dataset was used for training the 2-class classifier as mentioned in Section 3.1. All the sentences are labeled as *biographical* sentences.

`BigWikiBio`: We curated this dataset by crawling English Wikipedia articles. It consists of ~6M Wikipedia biographies. This dataset was used to train the 6-class classifier (see Section 3.2).

3 METHODOLOGY

The biography generation process involves multi-stage extractive subtasks. In this section, we describe these stages in detail. Along with the biographical information, a Wikipedia page also consists of an *infobox*. For the sake of completeness, we also describe an automatic approach to generate infoboxes similar to the one present on the Wikipedia page.

3.1 Identifying Biographical Sentences

A textual document that describes an event or some news related to a person contains a large number of non-biographical sentences as compared to biographical sentences. In the first stage, sentences were categorized into the above two categories.

3.1.1 Data Pre-processing. Given a text document, we partition it into a set of sentences. Next, we enrich extracted sentences by performing standard NLP tasks like special character removal, spell check, etc.

3.1.2 Sentence representation and classification. Each sentence is converted into a fixed-length `TF-IDF` vector representation. We consider sentences available in the TREC dataset as *non-biographical* sentences. Whereas sentences in `WikiBio` dataset are considered as *biographical*. We experiment with several machine learning models like Logistic Regression, Decision Trees, Naive Bayes, etc., to perform binary classification. Since the Logistic Regression model performed best (evaluation scores described in Section 4), we leverage its classified results for next stages.

3.1.3 Filtering False Positives. Our classifier resulted in some false positives — sentences that are non-biographical but were classified as biographical. We, further, filter out these cases by employing standard *Named Entity Recognition* technique. We consider only those sentences that contain at least one of the three named entities: (i) *PERSON*, (ii) *PLACE* or (iii) *ORGANIZATION*. In the next stage, we classify biographical sentences into several life-event categories.

3.2 Classifying Biographical Sentences

The categorized biographical sentences are further classified into six life-event classes namely *Education, Career, Life, Awards, Special Notes, and Death*. We leverage section information in `BigWikiBio` dataset (described in Section 2) to construct a mapping between sentences and life-events. We label each sentence in the Wikipedia page with its corresponding section heading and further map it to a broad life-event class. For example, sentences with sections headings as *College, High School, Early life and education, Education, etc.* are labeled as *Education* class, *Politics, Music career,*

Career, Works, Publications, Research, etc. are labeled as *Career* class, *Honors, Awards, Recognition, Championships, Achievements, Accomplishments, etc.* are labeled as *Honours/Awards* class, *Honors, Awards, Recognition, Championships, Achievements, Accomplishments, etc.* are labeled as *Honours/Awards* class, *Notes, Legacy, Personal, Gallery, Influences, Other, Controversies, etc.* are labeled as *Special Notes* class, and *Death, Death and Legacy, Later life, and Death, etc.* are labeled as *Death* class.

Next, we leverage the Logistic Regression model to perform this multi-class classification task. We construct similar fixed-length `TF-IDF` vector representation as described in Section 3.1. The classification results into clusters of similar sentences representing a single life-event of person.

3.3 Summarization

A single life-event cluster might contain hundreds of biographical sentences. We, therefore, rank the most important sentences by leveraging graph ranking algorithm[1] *Text Rank* [8]. For a given person, we apply *Text Rank* on each of the obtained six clusters. The ranking imparts flexibility in experimenting with multiple length values of the generated biography.

3.4 Generating Infobox

The infobox is a well-formatted table which gives a short and concise description of important facts related to the person. We use the following set facts in our proposed *Infobox* of a queried person.

- **Name:** Name of the queried person.
- **Date of Birth & Date of Death:** We use regular expressions to extract the date, depending on the context phrases such as 'born on', 'birth', etc.
- **Place of birth:** We use a similar methodology as above. We leverage part-of-speech (POS) tagging and Named Entity Recognition[2] to identify the place of the birth.
- **Awards:** We extract award information by leveraging a list of all the awards available at: *Wikipedia List of Awards page* [3]. We, next, use standard string matching to identify an award name in the biographical sentences.
- **Education & Career:** Here also, we leverage education information (degree, courses, etc.) present at official government sites like `data.gov.in` and `usa.gov`. Similarly, career-related information was obtained using the Wikipedia page *list of Occupations*.[4]

As an additional feature, we also present a profile image of the person by performing an image-based Google search query. This additional feature enriches textual biography with visual aspects similar to a Wikipedia profile of a person.

4 EXPERIMENTS AND RESULTS

In this section, we present evaluation results of two sub-tasks: (i) *Biography Generation* and (ii) *Infobox Generation*.

---

[1] We use Gensim implementation [9].
[2] We use the Stanford NER library[4].
[3] https://en.wikipedia.org/wiki/List_of_awards
[4] https://en.wikipedia.org/wiki/Lists_of_occupations

4.1 Tasks and Evaluation Measures:

The evaluation metrics for the above tasks are as follows:

4.1.1 Biography Generation. We compare biographies generated by `BioGen` against corresponding Wikipedia page. Note that, biography generation task is similar to document summarization. We, therefore, use *ROUGE* score to evaluate our generated biographies. In the current paper, we three variations of ROUGE, *ROUGE-1*, *ROUGE-2*, and *ROUGE-L* scores.

4.1.2 Infobox Generation. To evaluate the quality of the infobox generated by `BioGen` we define a score for each field in the infobox. Let I_{BG} and I_W be the infobox generated by `BioGen` and the infobox present on Wikipedia page respectively. Let $f_{BG} \in I_{BG} = \{f_{BG}^{(c)}\}$ be the set of characteristics present in the field recovered by `BioGen` and $f_W \in I_W = \{f_W^{(c)}\}$ be the corresponding field set in the infobox present on the Wikipedia page. Then, the score for each field $f_{BG} \in I_{BG}$ is defined as: $S(f_{BG}) = \frac{\sum_{c \in f_{BG}} [f_{BG}^{(c)} \in f_W]}{|f_W|}$ And the total score for the generated infobox I_{BG} is given by the average over the present fields given as: $S(I_{BG}) = \frac{\sum_{f_{BG} \in I_{BG}} S(f_{BG})}{|I_{BG}|}$ As the Infobox is generated containing the only specific set of fields which are Name, Date-Of-Birth, Place-Of-Birth and Death, Awards, Education, and Career; the score is calculated only corresponding to those fields.

4.2 Results

4.2.1 Biography Generation Accuracy. Extracting information from an arbitrary webpage is a challenging task. We, therefore, leverage three web resources to construct a source document set for a given person query. The resources are *Ducksters*[5], *IMDB*[6], and *Zoomboola*[7]. Ducksters is an educational site covering subjects such as history, science, geography, math, and biographies. IMDB is the world's most popular and authoritative source for movie, TV and celebrity content. Zoomboola is a news website. We experiment with randomly selected 150 biographies belonging to various domains such as Academics, Politics, Literature, Sports, Film Industry, etc.

Table 1 describes the *ROUGE* scores with the increasing number of sources that were used to generate the biography. We observe that *ROUGE* score increases with an increase in the number of sources. I.e., the similarity of biography generated by `BioGen` with that of the Wikipedia page increases as the number of sources increase. This also demonstrates the fact that a Wikipedia page is composed of multiple references.

	One Source			Two Sources			Three Sources		
	F	P	R	F	P	R	F	P	R
Rouge-1	0.19	0.13	0.29	0.22	0.15	0.31	0.26	0.41	0.21
Rouge-2	0.05	0.03	0.10	0.06	0.04	0.10	0.10	0.17	0.08
Rouge-L	0.16	0.15	0.33	0.16	0.15	0.33	0.21	0.38	0.19

Table 1: Rouge Score (F1-Score, Precision, Recall) of generated text after summarization when the input document is taken from one, two and three sources.

Table 2 shows the *ROUGE* score when we do not use the *Summarization* step of the biography generation process in `Biogen`. We can see that the Recall is better in this case, and this is because *Summarization* step filters out some text from the biography, and thus the overlap with the Wikipedia page decreases. Table 3 shows

	1-Source	3-Sources
Rouge-1	0.3869	0.5977
Rouge-2	0.2071	0.2961
Rouge-L	0.3618	0.5698

Table 2: Rouge scores (Recall values) without summarization with different number of sources.

Amitabh Bachchan's (Bollywood Film Industry Actor) biography generated using `BioGen` as an example. If we compare the generated biography to Amitabh Bachan's Wikipedia page[8], it turns out that the sentences which are highlighted in *italics* are very much related to the class in which they have been classified, which says that `BioGen` does a fairly good job in not only generating relevant sentences but also placing them in the proper sections. Also, we can see that the row corresponding to the field *Death* is empty. I.e. `BioGen` did not extract any information related to Amitabh Bachan's death, which is also true as per Wikipedia, as there is no information about his death. It is also important to note that `BioGen` did not add any arbitrary information in that field.

We experimented `BioGen` with one more parameter, which is the length of the generated biography. Figure 4 shows the change in the *ROUGE* score with the change in the length of the biographies generated. Again, we can see that as the length increases the recall increases, but the precision decreases. Which is evident, because as we add in more and more content in the biography, we cover more and more information present on Wikipedia resulting in the increase of recall. However, as the length increases, the amount of new information that we get goes on reducing, which is seen by the decrease in the precision value. Table 3 shows a sample biography
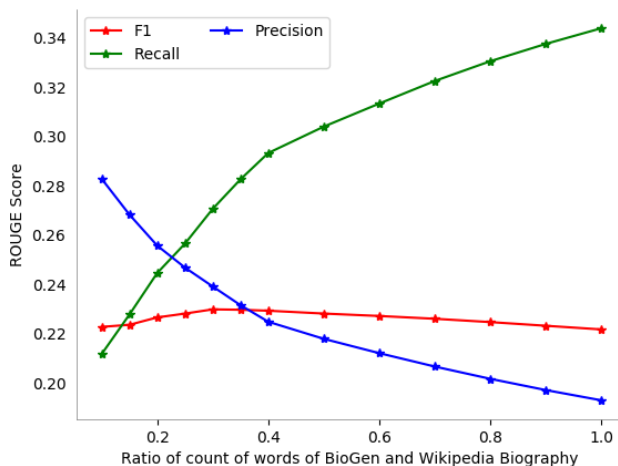

Figure 4: Change in *ROUGE* score with changing ratio of lengths of `BioGen` generated and Wikipedia biographies.

Career	*Bachchan's career moved into fifth gear after Ramesh Sippy's Sholay (1975).*The movies he made with Manmohan Desai (Amar Akbar Anthony, Naseeb, Mard) were immensely successful but towards the *latter half of the 1980s his career went into a downspin.*However, the importance of being Amitabh Bachchan is not limited to his career, although he reinvented himself and experimented with his roles and acted in many successful films.
Life	*Amitabh Bachchan was born on October 11, 1942 in Allahabad. He is the son of late poet Harivansh Rai Bachchan and Teji Bachchan.Son of well known poet Harivansh Rai Bachchan and Teji Bachchan.He has a brother named Ajitabh.*He got his break in Bollywood after a letter of introduction from the then Prime Minister Mrs. Indira Gandhi, as he was a friend of her son, Rajiv Gandhi.*He married Jaya Bhaduri, an accomplished actress in her own right, and they had two children, Shweta and Abhishek.*His son, Abhishek, is also an actor by his own rights. On November 16, 2011, he became a Dada (Paternal Grandfather) when Aishwarya gave birth to a daughter in Mumbai Hospital.*He is already a Nana (maternal grandfather) to Navya Naveli and Agastye - Shweta's children.*After completing his education from Sherwood College, Nainital, and Kirori Mal College, Delhi University, he moved to Calcutta to work for shipping firm Shaw and Wallace.
Awards	*In 1984, he was honored by the Indian government with the Padma Shri Award for his outstanding contribution to the Hindi film industry.France's highest civilian honour, the Knight of the Legion of Honour, was conferred upon him by the French Government in 2007, for his "exceptional career in the world of cinema and beyond"*
Death	
Rejected	Amitabh was in Goa during the last weekend to be one of the speaker at the THINK festival where he was honoured like he is being honoured in any and every place he makes his presence. At the very outset Bachchan was humble enough to let all those in the audience know that he held De Niro as one of his major sources of inspiration, was once forced to clear immigration in his hotel in Cairo, because his Egyptian fans became overly enthusiastic at the airport. Image caption

Table 3: Amitabh Bachchan's biography generated by `BioGen`.

generated for a famous Bollywood actor Amitabh Bachchan. It can be seen that our model does fairly well in extracting important sentences and classifies those into the six life-event classes. Also, as the actor is still alive, there should not be any 'death' related event, which our model outputs correctly.

4.2.2 Infobox Generation Accuracy. Table 5 shows an example of an Infobox (for Amitabh Bachchan) generated using `BioGen`. This Infobox recovers almost all the information present on the original Wikipedia page[9], and achieves a score of $S(I_{BG}) = 0.84$.

Name	Amitabh Bachchan
POB	Allahabad
DOB	1942-10-11
Education	Delhi University
Career	Actor, Artist, Assistant, Producer
Awards	Padma Vibhushan, Padma Bhushan, Padma Shri, Government of India

Table 5: An example of an Infobox generated using `BioGen`.

5 CONCLUSION AND FUTURE WORK

In this work, we proposed a system that generates a biography of a person, given the name and reference documents as input. However, we aim to build a system that just takes the name as an input and generates a biography by extracting information from the web. We would also like to enhance our system by incorporating coreference resolution so that it takes care of identifying the sentences related to the entity we are interested in. Right now, the system works well for extractive summarization, but we feel that adding in the abstractive form would help in creating better biographies. One of the other enhancements to check would be neural network based models for the classification and sentence generation task.

REFERENCES

[1] Massih R. Amini, Nicolas Usunier, and Cyril Goutte. 2009. Learning from Multiple Partially Observed Views -an Application to Multilingual Text Categorization. In *Proceedings of the 22Nd International Conference on Neural Information Processing Systems (NIPS'09).* Curran Associates Inc., USA, 28–36. http://dl.acm.org/citation.cfm?id=2984093.2984097

[2] Regina Barzilay, Noemie Elhadad, and Kathleen R. McKeown. 2001. Sentence Ordering in Multidocument Summarization. In *Proceedings of the First International Conference on Human Language Technology Research (HLT '01).* Association for Computational Linguistics, Stroudsburg, PA, USA, 1–7. DOI: http://dx.doi.org/10.3115/1072133.1072217

[3] Fadi Biadsy, Julia Hirschberg, and Elena Filatova. 2008. An unsupervised approach to biography production using wikipedia. *Proceedings of ACL-08: HLT* (2008), 807–815.

[4] Steven Bird, Ewan Klein, and Edward Loper. 2009. *Natural Language Processing with Python* (1st ed.). O'Reilly Media, Inc.

[5] Elena Filatova and John Prager. 2005. Tell Me What You Do and I'Ll Tell You What You Are: Learning Occupation-related Activities for Biographies. In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing (HLT '05).* Association for Computational Linguistics, Stroudsburg, PA, USA, 113–120. DOI: http://dx.doi.org/10.3115/1220575.1220590

[6] Rémi Lebret, David Grangier, and Michael Auli. 2016. Neural text generation from structured data with application to the biography domain. *arXiv preprint arXiv:1603.07771* (2016).

[7] Peter J Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198* (2018).

[8] Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing.*

[9] Radim Řehůřek and Petr Sojka. 2010. Software Framework for Topic Modelling with Large Corpora. In *Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks.* ELRA, Valletta, Malta, 45–50. http://is.muni.cz/publication/884893/en.

[10] Liang Zhou, Miruna Ticrea, and Eduard Hovy. 2005. Multi-document biography summarization. *arXiv preprint cs/0501078* (2005).

[9]https://en.wikipedia.org/wiki/Amitabh_Bachchan