

Post-edit Analysis of Collective Biography Generation

Bo Han Will Radford Anaïs Cadilhac Art Harol Andrew Chisholm Ben Hachey
Hugo.ai

58-62 Kippax Street
Surry Hills, Australia

{bhan,wradford,acadilhac,aharol,achisholm,bhachey}@hugo.ai

ABSTRACT

Text generation is increasingly common but often requires manual post-editing where high precision is critical to end users. However, manual editing is expensive so we want to ensure this effort is focused on high-value tasks. And we want to maintain stylistic consistency, a particular challenge in crowd settings. We present a case study, analysing human post-editing in the context of a template-based biography generation system. An edit flow visualisation combined with manual characterisation of edits helps identify and prioritise work for improving end-to-end efficiency and accuracy.

Keywords

Text Generation, Collective Intelligence, Evaluation

1. INTRODUCTION

Natural language generation applications often use human post-editing to ensure quality of final output. Understanding the kind and amount of manual effort is critical to prioritising how generation can be improved and what editorial guidelines should be implemented. Recent work also suggests that some post-editing can be automated, with substantial improvements in a machine translation shared task (+5.5 BLEU score) [2].

Post-edit analysis also has a long history in the evaluation of text generation tasks, in particular machine translation [1, 3]. The closest work to ours [4] is a post-edit evaluation of 2,728 pairs of system-edited texts from a system for generating weather forecasts, reporting a range of edits including individual style preferences as well as corrections. We build on this work to perform a post-edit analysis complementary to existing evaluation metrics like ROUGE and BLEU. We introduce a visualisation that helps identify and prioritise improvements, and present a case study using data from a commercial biography generation system, reporting a detailed analysis of edit flow and characteristic edit actions.

2. SYSTEM OVERVIEW

Our case study evaluates an collective biography generation system that is part of a larger commercial tool for finding and summarising information about a person on the web. Input comprises facts and events derived from social and news sources (e.g., name, affiliations, education, skills, investments). Facts and events are obtained from a search tool, which guides internal crowd users through the search process to select a range of relevant and diverse sources. The core generation capability uses 26 templates to produce an average of 2.8 suggested sentences per biography, e.g.:

Suggested: *Philip has been the Co-Founder at High Fidelity, Inc. since January 2013 and an Investor at Milk, Sunglass.io, Akili Interactive, Crowdfunder and more, and specializes in virtual worlds, start-ups and software development. He received a BS degree in Physics from The University of California, San Diego in 1992. Philip frequently tweets the hashtags #vr, #highfidelity and #avatars. Philip recently interacted with @id_aa_carmack and @inc on Twitter. Philip's favorite movie is Meditate and Destroy. Philip is a member of Facebook groups "Virtual Blogging", "San Francisco Bay Area Free Yoga and Meditation Events" and "BVHS Class of 1986".*

Internal crowd users select, edit and add suggested sentences, e.g.:

Selected: *Philip has been the Co-Founder at High Fidelity, Inc. since January 2013 and an Investor at Milk, Sunglass.io, Akili Interactive, Crowdfunder and more, and specializes in virtual worlds, start-ups and software development. Edited: He 'received' → 'holds' a BS degree in Physics from The University of California, San Diego in 1992. Selected: Philip frequently tweets the hashtags #vr, #highfidelity and #avatars. Not Selected: Philip recently interacted with @id_aa_carmack and @inc on Twitter. Not Selected: Philip's favorite movie is Meditate and Destroy. Philip is a member of Facebook groups "Virtual Blogging", "San Francisco Bay Area Free Yoga and Meditation Events" and "BVHS Class of 1986". New: He is married to Yvette Forte Rosedale. People: In 2007, Philip was listed in Time Magazine's 100 Most Influential People in The World.*

The result is a concise and informative text for end users, e.g.:

Final: *Philip has been the Co-Founder at High Fidelity, Inc. since January 2013 and an Investor at Milk, Sunglass.io, Akili Interactive, Crowdfunder and more, and specializes in virtual worlds, start-ups and software development. He received a BS degree in Physics from The University of California, San Diego in 1992. Philip frequently tweets the hashtags #vr, #highfidelity and #avatars. He is married to Yvette Forte Rosedale. In 2007, Philip was listed in Time Magazine's 100 Most Influential People in The World.*

The goal of the analysis here is twofold: (1) identify areas for improving system coverage and precision to focus human effort on high-value cognitively demanding tasks and improve end-to-end efficiency; (2) identify strategies for improving consistency and quality of biographies produced from noisy web data. The analysis uses 10,320 final biographies for which we also have the original suggested sentences from the generation system. This allows a unique and detailed characterisation of collective biography generation, tracking the editorial process from automatically generated sentences to final sentences after human post-editing.

©2017 International World Wide Web Conference Committee (IW3C2), published under Creative Commons CC BY 4.0 License. WWW'17 Companion, April 3–7, 2017, Perth, Australia. ACM 978-1-4503-4914-7/17/04.

<http://dx.doi.org/10.1145/3041021.3054264>



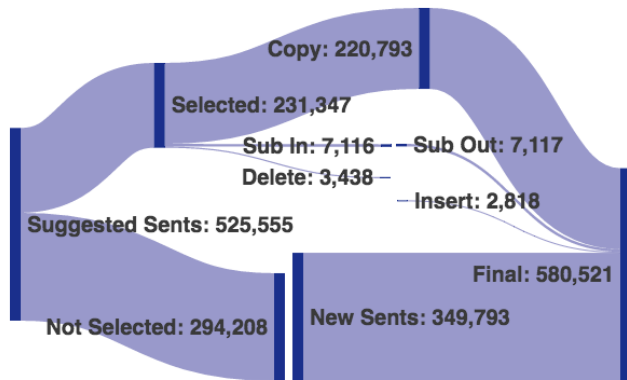


Figure 1: Flow of tokens through the post-editing process.

3. VISUALISING EDIT FLOW

Sentence alignment We first align suggested with final sentences using a minimum Levenshtein ratio of 0.8, defined as $(l_{(a+b)} - d(a, b)) / l_{(a+b)}$ where $l_{(a+b)}$ is the combined length of the input strings and $d(a, b)$ is the Levenshtein edit distance between the input strings. Suggested biographies have a total 28,842 sentences and 525,555 tokens, while final biographies have a total of 33,238 sentences and 580,521 tokens. So, final biographies tend to have more sentences (3.2 versus 2.8), but each sentence is shorter on average (17.5 versus 18.2).

Edit flow Figure 1 visualises the flow of tokens from suggestions through to the final biography using a Sankey diagram. The total number of tokens in suggested sentences breaks down into two categories: 231,347 (44%) are selected by users and 220,793 (56%) are not selected. Looking at tokens in final sentences: 349,793 (60%) are new sentences written by human editors. This suggests that work on improving end-to-end efficiency should prioritise extending fact collection and generation to provide more useful suggestions (more detail in Section 4). Trimming unused suggestions is less valuable as ignoring these is relatively easy.

Token alignment We refine the edit flow analysis by identifying phrase substitutions within the selected sentences aligned in the first step. We leverage the standard diff library in Python to collect token-level edit operations. This provides more detailed phrase-level substitution, deletion and insertion statistics corresponding to the middle of the edit flow diagram. Interestingly, we find that the vast majority (95%) of tokens in aligned sentences are copied directly from selected to final sentences. Phrase edits collectively represent only 5% of tokens in selected sentences, suggesting that generation obtains high accuracy where facts and templates are available.

4. CHARACTERISING EDITS

Selected From token alignments in selected sentences, we observe a total of 1,833 substitution, 621 deletion and 468 insertion phrase pairs and 4,568, 1,343 and 1,284 instances respectively. Overall, selection edits tend to result in shorter sentences with an average of 1.8 tokens before editing compared to 1.6 post-editing. Table 1 lists some of the most common edit actions. We observe that the top edits are generally grammatical (e.g., ‘the Engineer’ \mapsto ‘an Engineer’) or stylistic (e.g., ‘received a BFA’ \mapsto ‘holds a BFA’). Other edits address capitalisation consistency or errors from collected facts (e.g., ‘The University’ \mapsto ‘the University’, ‘microsoft office’ \mapsto ‘Microsoft Office’).

Rank	Edit	# Instances
1	‘of’ \mapsto ‘at’	320 (4.4%)
2	‘’ \mapsto ‘the’	272 (3.8%)
3	‘’ \mapsto ‘also’	233 (3.2%)
4	‘a’ \mapsto ‘the’	166 (2.3%)
5	‘with’ \mapsto ‘in’	130 (1.8%)
19	‘The’ \mapsto ‘the’	39 (0.5%)
20	‘has been’ \mapsto ‘is’	34 (0.5%)
27	‘has’ \mapsto ‘have’	25 (0.3%)
30	‘the’ \mapsto ‘an’	23 (0.3%)
38	‘received’ \mapsto ‘holds’	20 (0.3%)

Table 1: Example edit actions.

Not Selected Ignoring the sentence alignment threshold, we find that 67% of suggested sentences that were not selected have some overlap with sentences in the final biography. The remaining 33% were not directly useful for their biographies. By inspection of a random 100 of these, we observe that most talk about hobbies (e.g., interests, travel, books), social media interactions, or a monolingual person’s sole language.

New Sentences Looking from the other side, we find that 53% of new sentences have some overlap with suggestions. These comprise sentences where editors deleted/introduced many tokens, or split/merged suggested sentences. The remaining 47% of new sentences are due primarily to coverage issues in search or fact collection feeding into generation.

5. DISCUSSION

We introduced an edit flow visualisation for analysis of human post-editing of automatically generated text. This helps to identify and prioritise work for improving end-to-end efficiency. It is encouraging that most new sentences include suggested content, but these sentences still represent significant human cost. This suggests that end-to-end efficiency can be improved most by expanding generation templates to provide editors more style and content options. The next priority is then to automate post-edits where possible, learnt from data here or external resources. Another issue discovered during analysis is the lack of clear editorial policy, with users sometimes making conflicting edits. Editorial style policy is challenging to implement with crowd-sourced editing, but we expect that gradual introduction of high-precision rules should actively guide editors toward a more consistent style. In current work, we are implementing the changes suggested above, developing a virtuous circle to improve both efficiency and quality over time.

6. REFERENCES

- [1] W. Aziz, S. Castilho, and L. Specia. PET: a tool for post-editing and assessing machine translation. In *LREC*, pages 3982–3987, 2012.
- [2] O. Bojar, R. Chatterjee, C. Federmann, Y. Graham, B. Haddow, M. Huck, A. Jimeno Yepes, P. Koehn, V. Logacheva, C. Monz, M. Negri, A. Neveol, M. Neves, M. Popel, M. Post, R. Rubino, C. Scarton, L. Specia, M. Turchi, K. Verspoor, and M. Zampieri. Findings of the 2016 conference on machine translation. In *WMT*, pages 131–198, 2016.
- [3] C. Scarton. Discourse and document-level information for evaluating language output tasks. In *NAACL Student Research Workshop*, pages 118–125, 2015.
- [4] S. G. Sripada, E. Reiter, and L. Hawizy. Evaluating an NLG system using post-editing. In *IJCAI*, pages 1700–1701, 2005.