

Team Zoom @ AutoMin 2023: Utilizing Topic Segmentation And LLM Data Augmentation For Long-Form Meeting Summarization

Felix Schneider

Zoom Video Communications
felix.schneider@zoom.us

Marco Turchi

Zoom Video Communications
marco.turchi@zoom.us

Abstract

This paper describes Zoom’s submission to the Second Shared Task on Automatic Minuting at INLG 2023. We participated in Task A: generating abstractive summaries of meetings. Our final submission was a transformer model utilizing data from a similar domain and data augmentation by large language models, as well as content-based segmentation. The model produces summaries covering meeting topics and next steps and performs comparably to a large language model at a fraction of the cost. We also find that re-summarizing the summaries with the same model allows for an alternative, shorter summary.

1 Introduction

We participated in the 2023 workshop on automatic minuting at INLG 2023 (Ghosal et al., 2023). We chose to focus our efforts on the technical meetings in Task A, as it most closely aligns to a real-world use case of typical meetings. Meeting summarization is a challenging task, for both technical and semantic reasons:

- Summarizing spoken meetings by necessity relies on a transcription, which can be noisy, even if done by humans. Direct speech summarization is still in its infancy (Wang, 2022).
- There is a large amount of variability in what can be considered a good summary, more so than in other text-to-text tasks like machine translation. Human judgement is often necessary, as automatic, reference-based metrics can be insufficient.
- Summarization is a very hard task even for humans. As a result, very little training data is available, and that which is available can usually not be combined because of different styles.

- The meeting transcripts are longer than the maximum input length of many currently available models (over 8000 words in AutoMin).

On the other hand, meetings often follow an agenda, with specific topic being discussed sequentially. Our model exploits this structure by breaking the meeting down into smaller chunks by topic. We also address the lack of data by generating artificial labels using a large language model (LLM).

2 Related Work

Submissions to the previous AutoMin Workshop (Ghosal et al., 2021) used different approaches, including extractive summaries (Sharma et al., 2021; Williams and Haddow, 2021) and abstractive summaries using pre-trained language models. The teams using neural abstractive models used different approaches to dealing with the long input length in the dataset, including truncating (Garg and Singh, 2021), chunking by length (Shinde et al., 2021; Mahajan et al., 2021) and topic segmentation (Yamaguchi et al., 2021). We found the summary quality to be heavily dependent on the quality of the segmentation and found that purely length-based segmentation leads to worse summaries than topic-based segmentation. Because of the difficulties pointed out above, all of the best-performing systems in the previous AutoMin did not make use of the training data at all.

3 Model Architecture

Our main summarization model is derived from `bart-large` (Lewis et al., 2019). Because this model was only pretrained with 1024 position encodings, sequences longer than 1024 tokens cannot be processed by this model. While we have experimented with long-input models such as Longformer (Beltagy et al., 2020) and Big Bird (Zaheer

et al., 2020), we found those models not to perform better overall.

In order to summarize a meeting, which is generally much longer than 1024 tokens, we use Content Vector Segmentation (CVS) (Alemi and Ginsparg, 2015) to break the meeting down into smaller segments. CVS uses a dynamic programming algorithm to find a globally optimal segmentation as determined by a scoring function. It operates on sentence embeddings, which we generate with a pre-trained Sentence-Bert (Reimers and Gurevych, 2019). We observed that introducing more splits would almost always increase the score, which would almost always result in the meeting being split into the maximum number of segments. In order to allow for a variable number of segments per meeting, we introduced a CVS score improvement threshold of 1.2. Only solutions which improve upon the current best by at least the threshold are considered an improvement. The desired segment length was tuned with the 1024-token limit of Bart in mind, although the algorithm may still produce segments longer than that, in which case they are truncated.

4 Processing

The organizers provided a training set of 84 meetings annotated with summaries (some multiple times) by human annotators (Nedoluzhko et al., 2022). However, we found this training data unsuitable for the purposes of training our model. The annotation style is inconsistent across annotators and while we could apply the CVS segmentation to the source, we would need to create an alignment between source segments and summary sentences. We therefore opted to create pseudo-references by using a large language model.

The AutoMin data has been deidentified, meaning all names of persons, organizations and projects are replaced by “PERSON[N]”, “ORGANIZATION[N]” and “PROJECT[N]”, respectively. We used models that had not been trained on such deidentified data, so we converted to pseudonyms instead: We replaced the “PERSON[N]” tags by arbitrarily chosen English names, respecting the gender information present in the training set (for the test set, no such information is available, so we used male names for all participants). For projects and organizations, we randomly generated three- and four-letter acronyms, respectively, making sure not to use ones that represent well-known organiza-

tions or projects. Note that we made no attempt to determine the actual identities of the persons in the meetings (as required in the usage agreement), we simply replaced the generic tags by randomly chosen names. Before submitting the results, we restored the original deidentification.

For the LLM, we used `text-davinci-003` from OpenAI, which is from the GPT-3.5 series.¹ We used two priming examples from the training set (see Appendix A), which we annotated by hand to demonstrate the style that we wanted. We then sampled with a temperature of 0.7.

We found that the outputs are quite sensitive to the choice of priming example. When choosing two examples that are about a similar topic, something similar to overfitting occurs and the generated summaries will always be about the same topic as the examples, regardless of source content. The priming examples allow us to influence the summary length, style and grammar to a certain extent, but not fully. For example, we would have preferred summaries to be completely in the simple present tense (as the priming examples), but in many generated summaries, the first sentence was in the present progressive, even when giving more explicit instructions to the LLM. We reviewed model outputs on the AutoMin 2021 development set.

We generated one pseudo-reference for each segment of the training set for a total of 704 examples. We then fine-tuned a `bart-large` model that was previously trained on XSum (Narayan et al., 2018) and SamSum (Gliwa et al., 2019)² for 5 epochs using Adam with a learning rate of 0.0002. We applied the model using beam search with beam size 4. The resulting model performs similarly to GPT-3 while being much smaller.

The final summary for a meeting is the concatenation of all segment summaries. We also provided an alternative, shorter summary by applying the same Bart model as above to the concatenated summaries, producing a summary of summaries (SoS).

An overall similar architecture was proposed by (Asi et al., 2022), which was developed contemporaneously to ours. Also, (Shinde et al., 2021) also used a segment-then-summarize approach in AutoMin 2021, although they did not use content-based segmentation, LLM pseudo-labels, or a sum-

¹<https://platform.openai.com/docs/models/gpt-3-5>

²`lidiya/bart-large-xsum-samsum` from HuggingFace Hub

Model	Dev R-1	Test R-1	Test2 R-1	Test23 R-1	Test23 Adequacy	Test23 Gramm.	Test23 Fluency	Test23 Relevance
Bart	20	20	37	41	4.6	4.8	4.7	4.5
Bart SoS	29	22	28	29	3.3	5.0	3.6	4.7
GPT-3.5	22	21	40	41	3.8	5.0	4.0	4.5
GPT-4				44	4.6	5.0	4.8	5.0
Next best				40	4.1	4.9	4.3	4.6
Reference					4.0	4.9	4.3	4.6

Table 1: Metrics of submitted systems and other entries in the shared task. Human scores are out of 5. Gramm. is Grammaticality.

mary of summaries.

5 Summary Style

In choosing our priming examples for the LLM, we made a deliberate decision not to follow the summary style from the AutoMin data. While this may cause our results to be unfairly rated by automatic metrics, we did not feel that the style of summary from the data was ideal for two reasons: First, full sentences in simple present tense aligns the data more closely to the SamSum pretraining task, which we hypothesized to help with task transfer, especially given the small training set. Second, the summaries from AutoMin often omit attribution, i.e. they do not state who gave a particular piece of information. We feel that it is important for a summarization model to always attribute statements in its output. Unless it is a very sophisticated system, the model has no way of distinguishing when a statement in the source represents a fact and when it represents an opinion. Without attribution, a statement in the summary becomes a factual statement by the summarization model, which may not reflect reality. Furthermore, if the summary states an opinion as fact, it could be misunderstood as the model actually holding that opinion. We therefore choose to alter the summary style to always include attribution and use full sentences in the simple present tense.

6 Results and Discussion

We submitted three models for evaluation: Our Bart model, the same model with summary of summaries and GPT-3.5. Note that we included the latter only to get a reference human evaluation, we do not claim it as our work.

Table 1 shows our results, as well as the official evaluations of GPT-4 and the “Kmjec” system.³

³At the time of submission, we were not informed what

Human evaluation was performed on the four axes of “Adequacy”, “Grammaticality”, “Fluency” and “Relevance” (Ghosal et al., 2023). When averaging the four human metrics, ours was the second-highest rated submission behind GPT-4. Our long summaries outperform GPT-3.5 and compare on par with the human reference. While the long summaries compared on par with GPT-4 on Adequacy, they rate lower on Relevance, so the annotators would have preferred shorter, more focused summaries. Predictably, the short (summary-of-summary) summaries rated higher on relevance, but lost a lot of adequacy, so perhaps the annotators found them too brief.

The automatic metrics reveal a change in annotation style over the subsequent test sets: The references get progressively longer from dev to test to test2 (and presumably to test23). Where the average reference in dev has 110 words, test has 163 and test2 418 words. The training set is in the middle with 260 words on average. It seems therefore unwise to evaluate models on the dev portion of the data, as it does not reflect the training or testing summaries.

Our generated summaries on average 616 words for the concatenated segment summaries and 118 words for the summary-of-summary, which explains why the SoS scores better on the dev set than the long summaries, but not on test2, where the length aligns more closely with the long summaries. The ability of the same model to perform the summary of summaries is somewhat surprising, given that this task does not match its most recent finetuning domains (the closest would be XSum). It could indicate that the source domain (prose or dialog) is only of secondary importance if the expected output style does match the training. Or it could be that because the inputs are the model’s

team made this submission.

own output, their internal representations are more “familiar” to it, allowing it to understand them more easily than other prose texts. We leave this question for future experiments.

The summary of summaries is often very extractive, but without the lead bias typical of news summaries. Because sentences from different segments often immediately follow one another, it could create a false impression of connection. The human evaluation will show whether this is a serious issue.

7 Conclusion

We presented our system for the automatic meeting summarization task. Our results represent a significant improvement over our previous systems, and we made several valuable observations:

- Using content-based segmentation significantly improves downstream results as opposed to using length-based segmentation.
- We can leverage Large Language Models to generate pseudo-references in a controllable style using minimal annotation effort.
- The same model is able to re-summarize its own summaries, adding the option for a shorter summary variant.

We would like to thank the organizers and look forward to the next installation of the shared task.

References

- Alexander A Alemi and Paul Ginsparg. 2015. Text segmentation based on semantic word embeddings. *arXiv preprint arXiv:1503.05543*.
- Abdelkadir Asi, Song Wang, Roy Eisenstadt, Dean Geckt, Yarin Kuper, Yi Mao, and Royi Ronen. 2022. An end-to-end dialogue summarization system for sales calls. *arXiv preprint arXiv:2204.12951*.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Amitesh Garg and Muskaan Singh. 2021. [Team Symantelytical @ AutoMin 2021: Generating Readable Minutes with GPT-2 and BERT-based Automatic Minuting Approach](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 65–70.
- Tirthankar Ghosal, Ondřej Bojar, Marie Hledíková, Tom Kocmi, and Anja Nedoluzhko. 2023. Overview of the second shared task on automatic minuting (automin) at intl 2023. In *Proceedings of the 16th International Conference on Natural Language Generation: Generation Challenges*. Association for Computational Linguistics.
- Tirthankar Ghosal, Ondřej Bojar, Muskaan Singh, and Anja Nedoluzhko. 2021. [Overview of the First Shared Task on Automatic Minuting \(AutoMin\) at Interspeech 2021](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 1–25.
- Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. Samsun corpus: A human-annotated dialogue dataset for abstractive summarization. *arXiv preprint arXiv:1911.12237*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Parth Mahajan, Muskaan Singh, and Harpreet Singh. 2021. [Team AutoMinuters @ AutoMin 2021: Leveraging state-of-the-art Text Summarization model to Generate Minutes using Transfer Learning](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 34–40.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Anna Nedoluzhko, Muskaan Singh, Marie Hledíková, Tirthankar Ghosal, and Ondřej Bojar. 2022. ELITR Minuting Corpus: A novel dataset for automatic minuting from multi-party meetings in English and Czech. In *Proceedings of the 13th International Conference on Language Resources and Evaluation (LREC-2022)*, Marseille, France. European Language Resources Association (ELRA). In print.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*.
- Umang Sharma, Muskaan Singh, and Harpreet Singh. 2021. [Team The Turing TESTament @ AutoMin 2021: A Pipeline based Approach to Generate Meeting Minutes Using TOPSIS](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 71–77.
- Kartik Shinde, Nidhir Bhavsar, Aakash Bhatnagar, and Tirthankar Ghosal. 2021. [Team abc@ automin 2021: Generating readable minutes with a bart-based automatic minuting approach](#). *Proceedings of the First Shared Task on Automatic Minuting at Interspeech*, pages 1–8.
- Jun Wang. 2022. Essumm: Extractive speech summarization from untranscribed meeting. *arXiv preprint arXiv:2209.06913*.

Philip Williams and Barry Haddow. 2021. [Team UEDIN @ AutoMin 2021: Creating Minutes by Learning to Filter an Extracted Summary](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 78–81.

Atsuki Yamaguchi, Gaku Morio, Hiroaki Ozaki, Ken-ichi Yokote, and Kenji Nagamatsu. 2021. [Team Hitachi @ AutoMin 2021: Reference-free Automatic Minuting Pipeline with Argument Structure Construction over Topic-based Summarization](#). In *Proc. First Shared Task on Automatic Minuting at Interspeech 2021*, pages 41–48.

Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in neural information processing systems*, 33:17283–17297.

A Priming Examples

Note that these include the pseudonyms described in section 4.

Transcript:

Ernie: Let me copy it.

Ernie: So here in the tap box I'm I'll be where is the message box, yeah.

Ernie: Send messages.

Ernie: Tap box in this, how do I send message?

Ernie: In in here?

Ernie: Ok 'll I'll paste in the GSR document.

Annabel: There's the side panel, you can use.

Annabel: If you click side panel maybe, there you could send a message.

Ernie: I can see only the list of the persons' It has to be somewhere.

Annabel: If you click side panel, it's

Ernie: Aha, ok, I got it yeah.

Ernie: So I send the link there and if you click there if you click it.

Ernie: And password, user name password is editor editor.

Ernie: I'll try that too.

Ernie: So it will once you once you're able to view the subtitling platform.

Annabel: Ok,

Ernie: So there are like lot of languages.

Annabel: Ok.

Ernie: And you might want to like unselect most of the languages from the left side.

Ernie: From from the right side.

Annabel: Ok, so now I see them all and I would have to click on each of them to unselect them?

Ernie: Yeah yeah.

Annabel: Ok. Can I left Oh.

Annabel: Well I think instead of Romanian it's showing Russian.

Ernie: I, yeah, already switch last night I think.

Ernie: I think Russian and Romanian they are like swapped.

Ernie: So can you choose Russian?

Annabel: Yeah, no, mhm, Russian looks more like Czech or Slovak I Oh ok, there are thing they are like Polish is Romanian.

Annabel: So I think they're just mhm, the UI is just it should be.

Ernie: Yeah, ok.

Ernie: So that's slight bug backward that means.

Ernie: So 'll fix it finding it now.

Ernie: So right now the Polish is the real Romanian right?

Annabel: Yes.

Ernie: Ok, ok.

Ernie: So we need to keep up in in order to view the subtitles.

Ernie: So think Jan will might want to take over from here.

Annabel: I think now it stopped.

Annabel: The thing.

Ernie: Yeah, because we aren't talking so.

Ernie: So can you comment anything about the subtitles, the Romanian subtitles?

Ernie: I mean if you look at the real English ASR and compare how how worse is it?

Ernie: Or is it?

Annabel: Uh, so I think it does have a little bit of trouble telling when new turn started.

Annabel: Or maybe just has lag.

Annabel: And it does have some some non sense in there to be honest.

Ernie: Uh huh.

Annabel: But I don't know if it's because of the quality of the how how we're speaking.

Annabel: Um, so I I couldn't say what we were talking about by reading, um.

Ernie: Uh, ok, yeah please.

Annabel: guess I would also, I'm curious, if you have some mhm like what kind of quality feedback would you like me to give.

Annabel: Maybe I could look over the transcript later on and give it some annotations or do you want just a general feel to it?

Summary:

Ernie shows Annabel the subtitling platform. There is a bug in the UI where the languages are swapped. Annabel is not sure about the quality of the Romanian subtitles.

###

Transcript:

Ayden: Okay so Lina, I have -

Aurelia: Lina?

Ayden: I have few questions

Ayden: So, let us just be on the same page regarding the annotation.

Ayden: So first we have the ASR right?

Ayden: From the -

Lina: Yes.

Ayden: Video meetings.

Ayden: Right?

Lina: Yes.

Ayden: Then, then we create the reference summary.

Lina: Then we correct the ASR.

Ayden: Okay.

Ayden: And these are the files that you can find in the data with MAN.

Ayden: M-A-N.

Ayden: This is manual.

Ayden: So the files named: first date, then, eh no, no.

Ayden: First the word" transcript" then M-A-and - it means that manual it means that manual transcript and it means that it is manually corrected ASR.

Ayden: So the ASR outputs are mostly very bad.

Ayden: So you can't use them for, to understand what the meeting was about.

Ayden: Right.

Lina: So if you look at the output.

Lina: If you look at the ASR and the file name which contains ASR in the name you can just see that it is senseless.

Lina: And that these MAN are corrected, yes.

Ayden: Okay.

Ayden: So we have that ASR first, eh, from the, from the

minute, eh, from the meetings.
 Ayden: Then we correct the ASR.
 Ayden: That is a manually corrected ASR.
 Ayden: Ad then, eh, the reference summaries are created from the manually corrected ASR.
 Lina: Yes.
 Ayden: So who creates this reference?
 Lina: So the annotators worked in such way that they have got, eh, they had a video or audio file and they corrected the transcript according to that.
 Lina: And then they created this summary.
 Aurelia: So Lina we have, always we have 2 annotators for, annotating each minute?
 Lina: Not always.
 Lina: Eh, not always.
 Lina: This is reflected in the table you have access to.
 Lina: So if you look at the table there is a colon, eh twice transcript or twice minutes.
 Lina: And if there is "1" in the colon it means that it is really doubled.
 Lina: That we have double annotation.
 Lina: And if it doesn't have "1" if it has "0" it means that we don't have -
 Aurelia: I'm sorry.
 Aurelia: Which, which table are you referring to Lina?
 Lina: OPDI minuting annotation.
 Lina: I, I have, I gave the link here to the chat.
 Lina: Eh, if you look at the OFM chat today.
 Aurelia: Yes, yes.
 Lina: And click on the table, you will be inside.
 Aurelia: Okay.
 Aurelia: Thank you.
 Ayden: Okay, so, erm, then we have the reference summary, which we refer to as the gold standard summary that are generated by human annotators.
 Ayden: Right?
 Lina: Uuf, yes.
 Lina: Well, so why I'm so kind of unsure.
 Lina: Because annotators are humans, mostly students, not always from the computer linguistic fields so creating the summary is quite subjective and sophisticated task and the summaries that are created by different annotators, when I look at them with my eyes, are very often not very similar.
 Lina: So they are gold annotations.
 Lina: But if two people make the same they are never the same.
 Ayden: Okay.
 Ayden: So for each manually corrected ASR.
 Ayden: So sometimes we have 1 reference summary and sometimes we have more.
 Ayden: Is that correct?
 Lina: Yes.
 Ayden: Okay, okay.
 Lina: So they are, eh, so we are getting new and new every day.
 Lina: So I work with more than 10 annotators and they gave me back the text.
 Lina: So this is the working, eh, flow.
 Ayden: Okay.

Summary:
 Ayden and Aurelia have some questions about the annotation process. Lina explains that the automatic ASR transcripts are first manually corrected. Then, one or more human annotators create reference summaries. If there is more than one annotation, they will often be quite different.

B Sample

Summary of meeting_en_dev_008:

B.1 Bart

PERSON7 and PERSON5 are discussing the Overleaf interface and the structure of the deliverable documents. PERSON7 is trying to fix the program in Overleaf to use it as a plain text editor, but he is not sure if it will work. PERSON5 explains that they need to include the ORGANIZATION1 presentation platform and the ORGANIZATION7 presentation platform in the documents. PERSON1 joins the conversation and asks if PERSON3's deliverable is working. PERSON3 confirms that it is.

PERSON7 and PERSON5 discuss the structure of the deliverable for the ORGANIZATION4 demo paper. They discuss the need to describe the internals of the project and the need for a reviewer for the publishing platform deliverable. PERSON7 suggests that the ORGANIZATION7 platform should be included and suggests that PERSON2, the intermediary with ORGANIZATION7, should talk to them about it. They also discuss the ORGANIZATION3 sheet and the continuous reporting. PERSON5 will email all of the reviewers and explain that they are late with the preparation.

PERSON7 and PERSON5 are discussing the ELG document and the progress to the template. PERSON7 suggests that they should have their full content ready by Wednesday and that PERSON1 should have some documentation ready by the end of the week. They agree that it is feasible for PERSON1 to submit the document, as it is a demonstrator and not a report. They also agree that the content should include a picture of the user interface and a description of how the information flow works.

PERSON1 and PERSON5 are trying to work of Overleaf outside of it. PERSON7 explains that they need to create a special password for Overleaf, but it should work. He also explains that PERSON6 has a problem with cloning, but if they clone it and there is a make file, it will compile.

PERSON7 and PERSON1 are discussing the deliverable for the ORGANIZATION4 demo paper. PERSON5 suggests that PERSON1 should use the structure suggested in the Overleaf document and provide a screenshot and a paragraph to describe how the platform behaves. PERSON7 suggests using the second strategy and asking ORGANIZATION7 to review his understanding of the platform. He also suggests creating a preview file in Git at the six one preview makes. PERSON1 then tries to make the demo paper, but it fails.

PERSON7 and PERSON5 are discussing the make and compiler. PERSON5 has just cloned the Git repo and is trying to figure out how to commit. PERSON1 joins the conversation and explains that it should be making Git commit from Git stash, and then he will push it.

PERSON1 and PERSON5 are trying to figure out how to install a package manager. PERSON1 explains that it is either the standard packages that PERSON5 has or the Tex Live manager. He suggests using the YAM install tex live dash last page or the B T.

PERSON1 and PERSON5 are discussing the Overleaf project. PERSON1 is trying to find the Paragraph view text, but it is not working. PERSON5 explains that it is still the wrong package, and she will check the make file for the target names. She also mentions that it may be easier to work online without it, and PERSON1 agrees to work on it over night.

B.2 Bart (summary of summaries)

PERSON7 and PERSON5 are discussing the structure of the deliverable documents for the ORGANIZATION4 demo paper. They discuss the need to include the ORGANIZATION1 presentation platform and the ORGANIZATION7 presentation

platform in the documents and the need for a reviewer for the publishing platform deliverable. They agree that PERSON1 should use the structure suggested in the Overleaf document and provide a screenshot and a paragraph to describe how the platform behaves. They also discuss the ORGANIZATION3 sheet and the continuous reporting. PERSON1 agrees to work on the demo paper over night.

B.3 text-davinci-003

PERSON7 and PERSON5 are discussing the structure of their deliverable documents. They are using the Overleaf platform to collaborate on the project, but it is not working properly. PERSON5 suggests the documents should be published on the ORGANIZATION1 presentation platform. PERSON1 joins the conversation and has trouble accessing the chat window. PERSON7 sends the link again and PERSON1 is able to access it. PERSON7 and PERSON5 discuss the requirements and structure of the deliverable for the PROJECT1 project. They need to include a screenshot and description of the ORGANIZATION7 platform, and PERSON5 will be responsible for providing most of the content. They also need to find an internal reviewer for the document. PERSON5 will also be preparing documents to summarize the information and it is important to have the same structure for all three platforms.

PERSON7, PERSON5, and PERSON1 discuss the ELG document they need to submit for the EAMT workshop. PERSON7 is emailing colleagues from ORGANIZATION6 and ORGANIZATION5 to get feedback on the deliverable. They need to have the content ready by Wednesday, and the document should be two pages with pictures to demonstrate the functionalities. The paper is due on February 21st.

PERSON1 and PERSON5 are discussing how to use Overleaf for a project. PERSON7 explains that they need to use the menu button in the top left corner and click on Git. They will need to create credentials in order to use the command Git clone. PERSON7 also mentions that they can use a make file if they clone it and that it should work.

PERSON7 and PERSON5 discuss the structure of the deliverable and what should be included. PERSON5 suggests that she write the deliverable right away and PERSON7 agrees. PERSON1 has a problem with the ORGANIZATION4 demo paper, but PERSON7 offers to modify the makefile and commits with Overleaf Git. PERSON7 disconnects and PERSON1 and PERSON5 start a conversation. PERSON1 introduces himself and PERSON5 asks him about how to compile the code. PERSON1 explains that she should type 'make' and it will do the Git commands internally. PERSON5 mentions that it prompts her for a file name and PERSON1 suggests that it is probably related to the last page style.

PERSON1 and PERSON5 are discussing how PERSON5 can install a package manager. PERSON1 suggests TLMGR or a standard package, and suggests inserting the package name, "tex live dash". PERSON1 also suggests installing the "last page" package.

PERSON1 and PERSON5 are discussing how to use the Overleaf platform to create a paper. PERSON1 has sent a link to a Wikibooks page. PERSON5 is installing the dependencies and PERSON1 is trying to find the Paragraph view text. They both agree that Overleaf is convenient and that PERSON1 should be the one to dictate what is included in the document. Finally, they thank each other for their work.