

PriMock57: A Dataset Of Primary Care Mock Consultations

Alex Papadopoulos Korfiatis

Babylon

alex.papadopoulos¹

Radmila Sarac

radmila.sarac@gmail.com

Francesco Moramarco

Babylon, University of Aberdeen

francesco.moramarco¹

Aleksandar Savkov

Babylon

sasho.savkov¹

¹@babylonhealth.co.uk

Abstract

Recent advances in Automatic Speech Recognition (ASR) have made it possible to reliably produce automatic transcripts of clinician-patient conversations. However, access to clinical datasets is heavily restricted due to patient privacy, thus slowing down normal research practices. We detail the development of a public access, high quality dataset comprising of 57 mocked primary care consultations, including audio recordings, their manual utterance-level transcriptions, and the associated consultation notes. Our work illustrates how the dataset can be used as a benchmark for conversational medical ASR as well as consultation note generation from transcripts.

1 Introduction

The use of Automatic Speech Recognition (ASR) is widespread in the clinical domain but it is generally used to alleviate the administrative burden of clinical notes through dictation (Hodgson and Coiera, 2016; Kumah-Crystal et al., 2018).

However, the adoption of telemedicine, especially in primary care, generates vast quantities of clinical interaction recordings. Additionally, ASR models have become much more robust to applications in the clinical domain. In turn, this is beneficial for downstream Natural Language Processing (NLP) tasks, such as information extraction from clinical conversations (Selvaraj and Konam, 2021; Soltau et al., 2021) and automatic generation of consultation notes (Finley et al., 2018; Enarvi et al., 2020a; Quiroz et al., 2020; Molenaar et al., 2020).

Despite this being an active area of research it still lacks a commonly recognised ASR benchmark due to the sensitive nature of clinical conversations. Furthermore, as the datasets are not shared, research teams always need to invest time and resources into making their own private dataset. These limitations slow down progress in the field.

We release¹ a high quality public dataset of primary care consultation audio recordings, including manual transcriptions and associated consultation notes, which is the basis of our contributions:

1. a benchmark for ASR for primary care conversations;
2. a benchmark for automatic generation of consultation notes for primary care.

2 Related Work

Automated transcription of clinical consultations has attracted quite significant research interest; however, as mentioned above, there is no easily accessible common benchmark dataset in the style of Switchboard (Godfrey et al., 1992) or Fisher (Cieri et al., 2004), which are both non-medical conversational audio datasets. Because of this, comparing different approaches for clinical conversation ASR is challenging.

For example, Chiu et al. (2018) detail a dataset of $\approx 14,000$ hours of recorded and manually transcribed consultations that they use to train an end-to-end clinical conversation ASR model. Similarly, Kim (2020), Soltau et al. (2021) develop end-to-end ASR models for clinical conversations and Mani et al. (2020) train a sequence-to-sequence machine translation model to correct the errors of general-domain ASR engines; but they all use different, proprietary datasets. Johnson et al. (2014) and Kodish-Wachs et al. (2018) perform systematic reviews of the accuracy of a number of open-source and commercial ASR models for clinical conversation transcription; again, on proprietary datasets.

As for open-access datasets, He et al. (2020) compile and release two clinical dialogue datasets in Chinese and English, covering a wide range of clinical specialties. Ju et al. (2020) do the same for COVID-19 related clinical dialogue. These

¹<https://github.com/babylonhealth/primock57>

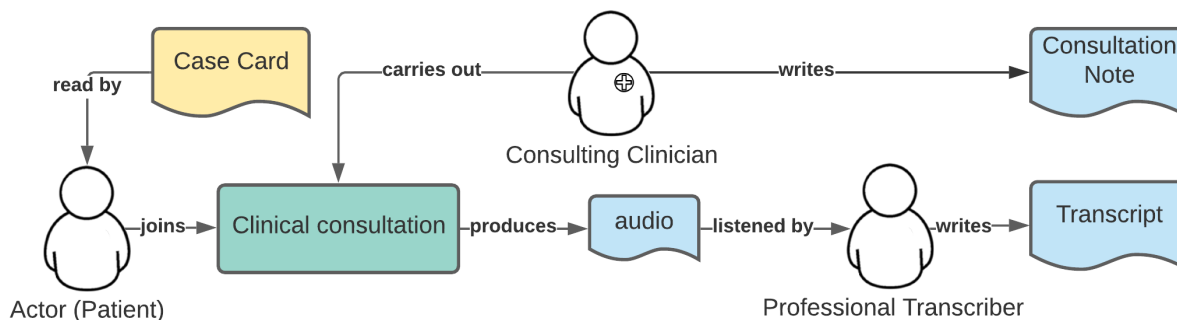


Figure 1: Overview of the data collection process. A mock patient, reading from a medical case card, has a consultation with a clinician which is recorded and transcribed. The resulting dataset includes the consultation audio recordings, notes and manual transcripts.

datasets are gathered from online clinical question answering sources; while they are relevant for clinical chatbot research, they are not representative of clinical interactions and do not include audio. Kazi et al. (2020) provide a dataset of audio recordings, automated transcripts and consultation notes for 70 mock psychiatric consultations — but no human transcripts.

Automatic consultation note generation and other long-form text summarisation tasks have rapidly developed due to recent advances in Natural Language Generation (NLG) architectures (Vaswani et al., 2017; Devlin et al., 2019). Several studies (Liu et al., 2019; MacAvaney et al., 2019; Zhang et al., 2020; Enarvi et al., 2020b; Joshi et al., 2020; Krishna et al., 2021; Chintagunta et al., 2021; Yim and Yetisgen-Yildiz, 2021; Moramarco et al., 2021; Zhang et al., 2021) use proprietary datasets of transcripts and notes to train NLG models end-to-end, and a number of them carry out automatic or human evaluations on their proprietary test sets. However, in a similar fashion to the ASR studies discussed above, most studies don’t publish these resources; hence, it is again prohibitively difficult to compare their proposed methods. Kazi et al. (2020) provide the only open access clinical dataset that could be used as a benchmark but it only contains psychiatric consultations, which is less applicable to primary care.

3 Dataset

The requirements for releasing a dataset containing Personal Health Information (PHI) are typically costly and involve collecting patient consent and/or de-identification, which is especially challenging with audio recordings. We built a mock consultation dataset as close as possible to the real conditions as a pragmatic alternative. The diagram in

Consultation type	Count
Otitis	2
Anaphylactic reaction	3
Cardiovascular	11
Dermatitis	4
Fever	4
Urinary tract infection	6
Upper respiratory infection	6
Asthma	2
Gastroenteritis	8
Mental health	3
Physical injury	2
Migraine	6

Table 1: A breakdown by consultation case card. The case card diagnoses were selected to be representative of common telemedicine presenting complaints.

Figure 1 shows an overview of the data collection process.

3.1 Mock consultation recordings

We employed 7 clinicians and 57 actors posing as patients from a range of ethnicities. The clinicians had experience with virtual consultations. Participation was optional and anyone could choose to withdraw at any time. Four of the clinicians were men and three were women; five of them had British English accent, and two of them Indian. The patient accent distribution is as follows: British English (47.4%), various European (31.6%), other English (10.5%), and other non-English (10.5%). The gender distribution was relatively even (52.6% women, 47.4% men); most participants were from 25 to 45 years old (see Figure A.1).

Each mock patient was given a case card that included background information (age, social history, family history of illnesses) as well as information about their presenting complaint, symptoms, condi-

Demographics (age, gender): 23 year old female
Presenting Complaint: Lower abdominal pain Duration of symptoms: 2 days
History, on open questioning: Have a terrible ache in my lower tummy and feeling hot and sweaty.
Symptoms and risk factors: There is some blood in the urine – pink colour Pain below belly button Feeling nauseated but no vomiting * * *

Table 2: An abridged example of a clinical case card for a Urinary Tract Infection. Mock patients were given a case card and asked to study it before consulting with the clinician. Full version available in the Appendix.

tions, and medications. The case cards were drawn from a pool of primary care conditions, representative of presenting complaints in UK primary care. For a breakdown of presenting complaints, see Table 1. An example case card is given in Table 2.

We recorded 57 mock consultations (8h38m6s in total) over 5 days, using proprietary telemedicine software that allowed us to export the individual clinician and patient audio channels.² In order to emulate real clinical practice, clinicians were using laptops while patients were using mobile phones in an office environment with background noise. Clinicians were asked to act as close as possible to their actual consultation sessions, including conforming to a consultation length of 10 minutes and writing a consultation note in the SOAP format (Pearce et al., 2016). The resulting mock consultations ranged between 3m48s and 14m18s, with an average consultation length of 9m5s.

3.2 Manual transcription

To transcribe the consultation recordings, we employed transcribers with experience in the clinical conversation domain, who were asked to:

1. Listen to the consultation audio recordings, in separate channels for clinicians and patients;
2. Identify the start and end points of individual utterances (continuous speech segments ending in a pause);

²Due to limitations of the software, audio was exported in compressed form (WebM encoder, Opus codec at a variable bitrate).




Figure 2: Average utterance length for clinician and patient as a function of conversation turns. The patient initially speaks more than the clinician but later in the consultation this trend is reversed.

3. Provide an accurate transcription of each of the utterances identified.

Thus we obtained a collection of start times, end times, and utterance-level transcriptions, important for the ASR evaluation described below.

Consultations have 92 conversation turns and 1,489 words on average; clinicians tend to speak more than patients (897 vs. 592 words per consultation) and take longer turns (19.3 vs 12.8 words per turn). Interestingly, patients tend to take longer turns than clinicians in the beginning of the consultation, where they presumably state their presenting complaint; turns are more balanced in the middle, and clinicians seem to take over during the diagnosis and management at the end (see Figure 2).

4 ASR Benchmark

We perform a baseline study of ASR for clinical conversations by passing the audio recordings of the mock consultations through commonly used open-source and commercial speech-to-text engines:

1. **Kaldi:** This is our baseline system, built using the Kaldi (Povey et al., 2011) speech recognition toolkit, running locally. It uses a pre-trained acoustic model from Zamia Speech³ and a 3-gram language model trained on a proprietary medical question answering dataset.
2. **NeMo QuartzNet & Conformer:** These systems use QuartzNet (Kriman et al., 2020) and Conformer (Gulati et al., 2020) ASR models, which we load using Nvidia’s NeMo toolkit.⁴

³<http://zamia-speech.org/asr/>

⁴<https://github.com/NVIDIA/NeMo>

ASR	WER								ECCA		
	Gender		Role		Accent			Pr	Re	F1	
	mean	stdev	M	F	Clinician	Patient	en-gb				other
GC STT	30.9 †	12.7	32.7	28.9	28.5	33.4	30.0	32.2	0.83	0.82	0.81
Azure STT	31.3 †	12.8	32.7	29.6	26.7	35.8	30.2	32.7	0.87	0.79	0.82
ATM	34.0‡	13.9	33.8	34.2	32.8	35.2	31.6	37.2	0.79	0.75	0.78
Kaldi	48.9	14.9	52.7	44.6	47.0	50.8	49.5	48.2	0.64	0.69	0.68
QuartzNet	46.4	15.5	48.4	44.1	48.1	44.7	46.6	46.1	0.67	0.49	0.56
Conformer	34.4‡	14.5	36.8	31.7	35.6	33.2	35.0	33.7	0.79	0.71	0.75

Table 3: Word Error Rate (WER) scores for a number of Speech-to-text engines, and Extracted Clinical Concepts Accuracy (ECCA) based on recognised clinical terms. The gender, role and accent breakdowns show how each factor affects the mean WER. † indicates lack of statistical significance between mean WER scores ($p = 0.097$); ‡ is weak significance ($p = 0.026$); all other scores are $p < 0.001$.

Both models are end-to-end and do not use a language model.

3. **Google Cloud Speech-to-text (GCSTT)**:⁵ a commercially available, general domain service. We use the *video* enhanced model which is only available for the *en-us* language.
4. **Amazon Transcribe Medical (ATM)**:⁶ a commercially available service, tailored specifically for medical use cases. There are models available for *clinical dictation* and *clinical conversation*; we use the conversation model with *speciality=Primary Care*.
5. **Azure Speech-to-text (ASTT)**:⁷ a commercially available, general domain service. We use the *Standard* model.

To test the accuracy of the above services, we first extract the audio for each individual utterance identified by our human transcribers. We then generate a transcript for the utterance using each of the ASR engines. We ensure consistency by performing the following post-processing steps on both human and automatic transcripts:

1. Remove disfluencies ("umm", "uhh", etc.). These are included in the reference transcripts, but often omitted in each STT service;
2. Replace numerals ("5", "9th", "1984") with written equivalents ("five", "ninth", "nineteen eighty-four") to ensure uniformity;
3. Remove all punctuation, collapse multiple spaces and convert to lowercase.

⁵<https://cloud.google.com/speech-to-text>

⁶<https://aws.amazon.com/transcribe/medical/>

⁷<https://azure.microsoft.com/en-us/services/cognitive-services/speech-to-text/>

Finally, we compute the Word Error Rate (WER) for each utterance using SCKT's *sc-lite*⁸ tool. The mean WER, including a breakdown by gender, role, and accent can be seen in Table 3. Even though both are general domain, Google and Azure together are the best performing models on our dataset ($p = 0.097$). Conformer performs surprisingly well, given that it is a character-level model evaluated on a word-level metric.

The base WER metric treats all words in a transcript as equally important; this may be less desirable in the clinical domain, where the correct transcription of specific clinical terms is expected to be more important. To test this, we use a proprietary clinical information extraction engine based on fuzzy string matching, linking to SNOMED-CT (Donnelly et al., 2006). We extract medical concepts from each utterance in both reference and hypothesis transcripts, then compare the concepts extracted to estimate accuracy based on clinical terminology (ECCA in Table 3). The results mostly match the WER comparisons; the medical-domain Amazon model does not seem to perform better.

5 Consultation Note Generation Benchmark

The consultation transcripts and corresponding notes (see example in Table 4) are intended as a parallel dataset to evaluate methods for automatically generating primary care consultation notes. We propose a benchmark for this task by evaluating a number of baseline approaches and reporting common automatic metric scores on our dataset. The approaches considered include:

⁸<https://github.com/usnistgov/SCKT>

Transcript		Note
Clinician	So, um, tell me what's been going on. You've been saying there's a problem with your hearing. Is that right?	History: Hx of difficulty hearing left ear for 6 weeks with tinnitus and slight nausea/ dizziness. One previous similar episode in the past- resolved spontaneously. No discharge/fever/itchiness/pain Doesn't use cotton wool buds No Pmhx of note Ex: Looks well, not in pain. Imp: need to exclude impacted wax in ear canal first Pln: for face to face GP appointment in 5 days to examine ear If any problems in interim to ring us back Pt happy with and understands plan
Patient	Yeah, so I just feel I can't really hear as well as I used to, like my hearing is kind of deteriorating in some way.	
Clinician	Right, OK. How long has this been going on for?	
Patient	Uh about six weeks.	
Clinician	Six weeks, OK. Um, and before that have you had any hearing problem at all?	
Patient	Um I had something maybe, about a year ago, but it only lasted a couple of days, it wasn't anything as long as this.	
Clinician	Right, OK, OK. And, um, in this six week period, have you had anything else happen? Have you had any other ear symptoms at all?	
Patient	Um, I occasionally get like a ringing in my left ear, uh just on the one side and um there's actually been a few times when I felt kind of a bit sick or a bit dizzy as well.	

Table 4: Snippet of a mock consultation transcript and the corresponding note, written by the consulting clinician.

Model	R1	R2	RL	B
BART-CNN	0.17	0.02	0.10	0.80
BERT-ext	0.21	0.03	0.10	0.78
Random	0.19	0.02	0.09	0.78
BART-finet	0.31	0.08	0.17	0.81

Table 5: Average common metrics scores of different models on the 57 consultations. R1 through L represent Rouge F1 scores for unigrams, bigrams, and longest-common-subsequence. B represents non-rescaled BERTScore; score range is between 0.7 to 0.9, so differences are less pronounced.

BART-CNN: a neural sequence-to-sequence summariser based on the BART model (Lewis et al., 2020) and fine-tuned on the Daily-mail/CNN dataset (Nallapati et al., 2016);

BERT-ext: a general-purpose extractive summariser based on Bert embeddings (Miller, 2019);

Random: a baseline that extracts 15 random sentences from the transcript and collates them to form a note;

BART-finet: a BART-CNN model further fine-tuned on a proprietary dataset of 8,000 real transcripts and consultation notes.

We evaluate the models on our dataset and report common summarisation metrics scores: Rouge-1,

-2 & -L (Lin, 2004) which compute the F-score across ngrams between generated and human notes; and BERTScore (Zhang et al., 2019), which computes the similarity between BERT embeddings of the notes.

The results can be seen in Table 5: the fine-tuned BART model scores highest with all metrics, while *BART-CNN* and *BERT-ext* fail to outperform the *Random* baseline model. This highlights the differences between consultation note generation and general-purpose summarisation.

A more detailed evaluation of this task can be found in Moramarco et al. (2022); example notes can be found in Appendix Table A.3.

6 Conclusion

We present a dataset of 57 high quality mocked consultation audio recordings, their manually aligned and diarised transcripts, and consultation notes. By publishing this dataset, we hope to offer a benchmark for future studies in both ASR for clinical conversations and Consultation Note Generation for the primary care domain.

References

Bharath Chintagunta, Namit Katariya, Xavier Amatriain, and Anitha Kannan. 2021. Medically aware gpt-

- 3 as a data generator for medical dialogue summarization. In *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*, pages 66–76.
- Chung-Cheng Chiu, Anshuman Tripathi, Kat Chou, Chris Co, Navdeep Jaitly, Diana Jaunzeikare, Anjuli Kannan, Patrick Nguyen, Hasim Sak, Ananth Sankar, Justin Jesada Tansuwan, Nathan Wan, Yonghui Wu, and Frank Zhang. 2018. [Speech recognition for medical conversations](#).
- Christopher Cieri, David Miller, and Kevin Walker. 2004. The Fisher corpus: A resource for the next generations of speech-to-text.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT (1)*.
- Kevin Donnelly et al. 2006. Snomed-ct: The advanced terminology and coding system for ehealth. *Studies in health technology and informatics*, 121:279.
- Seppo Enarvi, Marilisa Amoia, Miguel Del-Agua Teba, Brian Delaney, Frank Diehl, Stefan Hahn, Kristina Harris, Liam McGrath, Yue Pan, Joel Pinto, Luca Rubini, Miguel Ruiz, Gagandeep Singh, Fabian Stemmer, Weiyi Sun, Paul Vozila, Thomas Lin, and Ranjani Ramamurthy. 2020a. [Generating Medical Reports from Patient-Doctor Conversations Using Sequence-to-Sequence Models](#). In *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*, pages 22–30, Online. Association for Computational Linguistics.
- Seppo Enarvi, Marilisa Amoia, Miguel Del-Agua Teba, Brian Delaney, Frank Diehl, Stefan Hahn, Kristina Harris, Liam McGrath, Yue Pan, Joel Pinto, et al. 2020b. [Generating medical reports from patient-doctor conversations using sequence-to-sequence models](#). In *Proceedings of the first workshop on natural language processing for medical conversations*, pages 22–30.
- Gregory Finley, Erik Edwards, Amanda Robinson, Michael Brenndoerfer, Najmeh Sadoughi, James Fone, Nico Axtmann, Mark Miller, and David Suendermann-Oeft. 2018. [An automated medical scribe for documenting clinical encounters](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Demonstrations*, pages 11–15, New Orleans, Louisiana. Association for Computational Linguistics.
- John J. Godfrey, Edward C. Holliman, and Jane McDaniel. 1992. SWITCHBOARD: telephone speech corpus for research and development. In *Proceedings of the 1992 IEEE international conference on Acoustics, speech and signal processing - Volume 1, ICASSP'92*, pages 517–520, USA. IEEE Computer Society.
- Anmol Gulati, James Qin, Chung-Cheng Chiu, Niki Parmar, Yu Zhang, Jiahui Yu, Wei Han, Shibo Wang, Zhengdong Zhang, Yonghui Wu, and Ruoming Pang. 2020. [Conformer: Convolution-augmented transformer for speech recognition](#).
- Xuehai He, Shu Chen, Zeqian Ju, Xiangyu Dong, Hongchao Fang, Sicheng Wang, Yue Yang, Jiaqi Zeng, Ruisi Zhang, Ruoyu Zhang, Meng Zhou, Penghui Zhu, and Pengtao Xie. 2020. [MedDialog: Two Large-scale Medical Dialogue Datasets](#). *arXiv:2004.03329 [cs, stat]*. ArXiv: 2004.03329.
- Tobias Hodgson and Enrico Coiera. 2016. [Risks and benefits of speech recognition for clinical documentation: a systematic review](#). *Journal of the American Medical Informatics Association : JAMIA*, 23(e1):e169–e179.
- Maree Johnson, Samuel Lapkin, Vanessa Long, Paula Sanchez, Hanna Suominen, Jim Basilakis, and Linda Dawson. 2014. [A systematic review of speech recognition technology in health care](#). *BMC Medical Informatics and Decision Making*, 14:94.
- Anirudh Joshi, Namit Katariya, Xavier Amatriain, and Anitha Kannan. 2020. [Dr. summarize: Global summarization of medical dialogue by exploiting local structures](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pages 3755–3763.
- Zeqian Ju, Subrato Chakravorty, Xuehai He, Shu Chen, Xingyi Yang, and Pengtao Xie. 2020. [Coviddialog: Medical dialogue datasets about covid-19](#). <https://github.com/UCSD-AI4H/COVID-Dialogue>.
- Nazmul Kazi, Matt Kuntz, Upulee Kanewala, and Indika Kahanda. 2020. [Dataset for automated medical transcription](#).
- Suyoun Kim. 2020. [End-to-End Speech Recognition on Conversations](#). thesis, Carnegie Mellon University.
- Jodi Kodish-Wachs, Emin Agassi, Patrick Kenny, and J. Marc Overhage. 2018. [A systematic comparison of contemporary automatic speech recognition engines for conversational clinical speech](#). *AMIA Annual Symposium Proceedings*, 2018:683–689.
- Samuel Kriman, Stanislav Beliaev, Boris Ginsburg, Jocelyn Huang, Oleksii Kuchaiev, Vitaly Lavrukhin, Ryan Leary, Jason Li, and Yang Zhang. 2020. [Quartznet: Deep automatic speech recognition with 1d time-channel separable convolutions](#). In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6124–6128. IEEE.
- Kundan Krishna, Sapan Khosla, Jeffrey Bigham, and Zachary C. Lipton. 2021. [Generating SOAP notes from doctor-patient conversations using modular summarization techniques](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint*

- Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4958–4972, Online. Association for Computational Linguistics.
- Yaa A. Kumah-Crystal, Claude J. Pirtle, Harrison M. Whyte, Edward S. Goode, Shilo H. Anders, and Christoph U. Lehmann. 2018. [Electronic Health Record Interactions through Voice: A Review](#). *Applied Clinical Informatics*, 09(3):541–552. Publisher: Georg Thieme Verlag KG.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Zhengyuan Liu, Angela Ng, Sheldon Lee, Ai Ti Aw, and Nancy F Chen. 2019. Topic-aware pointer-generator networks for summarizing spoken conversations. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 814–821. IEEE.
- Sean MacAvaney, Sajad Sotudeh, Arman Cohan, Nazli Goharian, Ish Talati, and Ross W Filice. 2019. Ontology-aware clinical abstractive summarization. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 1013–1016.
- Anirudh Mani, Shruti Palaskar, and Sandeep Konam. 2020. [Towards Understanding ASR Error Correction for Medical Conversations](#). In *Proceedings of the First Workshop on Natural Language Processing for Medical Conversations*, pages 7–11, Online. Association for Computational Linguistics.
- Derek Miller. 2019. Leveraging bert for extractive text summarization on lectures. *arXiv preprint arXiv:1906.04165*.
- Sabine Molenaar, Lientje Maas, Verónica Burriel, Fabiano Dalpiaz, and Sjaak Brinkkemper. 2020. [Medical Dialogue Summarization for Automated Reporting in Healthcare](#). *Advanced Information Systems Engineering Workshops*, 382:76–88.
- Francesco Moramarco, Alex Papadopoulos Korfiatis, Aleksandar Savkov, and Ehud Reiter. 2021. A preliminary study on evaluating consultation notes with post-editing. In *Proceedings of the Workshop on Human Evaluation of NLP Systems (HumEval)*, pages 62–68.
- Francesco Moramarco, Alex Papadopoulos Korfiatis, Mark Perera, Damir Juric, Jack Flann, Ehud Reiter, Anya Belz, and Aleksandar Savkov. 2022. (in press): Human evaluation and correlation with automatic metrics in consultation note generation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics*.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290.
- Patricia F Pearce, Laurie Anne Ferguson, Gwen S George, and Cynthia A Langford. 2016. The essential soap note in an ehr age. *The Nurse Practitioner*, 41(2):29–36.
- Daniel Povey, Arnab Ghoshal, Gilles Boulianne, Lukas Burget, Ondrej Glembek, Nagendra Goel, Mirko Hannemann, Petr Motlicek, Yanmin Qian, Petr Schwarz, et al. 2011. The kaldi speech recognition toolkit. In *IEEE 2011 workshop on automatic speech recognition and understanding*, CONF. IEEE Signal Processing Society.
- Juan C Quiroz, Liliana Laranjo, Ahmet Baki Kocaballi, Agustina Briatore, Shlomo Berkovsky, Dana Reza-zadegan, and Enrico Coiera. 2020. [Identifying relevant information in medical conversations to summarize a clinician-patient encounter](#). *Health Informatics Journal*, 26(4):2906–2914. Publisher: SAGE Publications Ltd.
- Sai P. Selvaraj and Sandeep Konam. 2021. [Medication Regimen Extraction from Medical Conversations](#). In Arash Shaban-Nejad, Martin Michalowski, and David L. Buckeridge, editors, *Explainable AI in Healthcare and Medicine: Building a Culture of Transparency and Accountability*, Studies in Computational Intelligence, pages 195–209. Springer International Publishing, Cham.
- Hagen Soltau, Mingqiu Wang, Izhak Shafran, and Laurent El Shafey. 2021. [Understanding Medical Conversations: Rich Transcription, Confidence Scores & Information Extraction](#). *arXiv:2104.02219 [cs]*. ArXiv: 2104.02219.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008.
- Wen-wai Yim and Meliha Yetisgen-Yildiz. 2021. Towards automating medical scribing: Clinic visit dialogue2note sentence alignment and snippet summarization. In *Proceedings of the Second Workshop on Natural Language Processing for Medical Conversations*, pages 10–20.
- Longxiang Zhang, Renato Negrinho, Arindam Ghosh, Vasudevan Jagannathan, Hamid Reza Hassanzadeh, Thomas Schaaf, and Matthew R Gormley. 2021.

Leveraging pretrained models for automatic summarization of doctor-patient conversations. *arXiv preprint arXiv:2109.12174*.

Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.

Yuhao Zhang, Derek Merck, Emily Tsai, Christopher D Manning, and Curtis Langlotz. 2020. Optimizing the factual correctness of a summary: A study of summarizing radiology reports. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5108–5120.

Appendix




Figure A.1: Accent and age group distributions for patients in the 57 mock consultations.

Demographics (age, gender):
23 year old female
Presenting Complaint:
Lower abdominal pain
Duration of symptoms: 2 days
History, on open questioning:
Have a terrible ache in my lower tummy and feeling hot and sweaty.
Symptoms and risk factors:
There is some blood in the urine – pink colour
Pain below belly button
Feeling nauseated but no vomiting
Going to the toilet a little more often but drinking lots of fluids
No urine urgency or pain when passing urine.
Was constipated until 1 week ago but that has cleared up now
Had sexual intercourse 4 days ago
No new sexual partner since last STI screen 6 months ago
No vaginal discharge
Has Implanon contraceptive implant for 1 year
No change in vaginal bleeding
No loin pain
Activities of daily living: No problems performing daily activities
Family history: nil
Past Medical History: nil
Drug History: Implanon
Allergies: Amoxicillin

Table A.1: Example clinical case card for a Urinary Tract Infection. Mock patients were given a case card and asked to study it before consulting with the clinician.

Human Transcription	Google Speech-to-text
<p>Doctor: Hello?</p> <p>Patient: Hello. Can you hear me well?</p> <p>Doctor: Uh uh yes. I think. It's a bit better. It's a bit, it's a bit, it's not very clear. But let's continue anyway.</p> <p>Patient: OK.</p> <p>Doctor: Uh, OK. Let's start again. So how can I help you sir?</p> <p>Patient: Yes. So, it's been a few days now. I have like a sore, and a red skin. It's kind of, it's really itchy, and it's like super annoying. So I'd like to find something quick to solve it.</p> <p>Doctor: OK. No, no problem. I'm happy to help. Um whereabouts in your skin is it affected?</p> <p>Patient: Uh, mostly like my chest, my, my hands, my arms. Like, like really, it's it's super annoying. Like it's itching a lot, like all the time. And I can't even sleep at night. I really need something quickly to, to solve it. Because even at work I, I can, when I'm in a meeting and I have to, like uh think about my work, I can't focus, I can't actually focus on my work. It's really annoying because I can't actually think about, uh, what I have to say. I'm always like, uh, disturbed by this disease.</p> <p style="text-align: center;">* * *</p> <p>Doctor: OK. OK. So it's something for you to think about. you can get different types of antihistamines. I can give you something a little bit stronger today as well. Um, something like Fexofenadine, which I can give to you today. It's definitely worth trying, and it's not going to do you any harm.</p> <p>Patient: OK.</p> <p>Doctor: Um but I think using the steroids and the emollients, um on a regular basis Uh over the next week to ten days, should hopefully control your symptoms. But do come back and see me next week, if things don't get better.</p> <p>Patient: That sounds good.</p> <p>Doctor: OK? Um do you have any questions for me?</p> <p>Patient: Uh, no that's it. Thank you very much. Bye. Thank you as well. Bye.</p>	<p>Doctor: Hello.</p> <p>Patient: Hello, can you hear me wet?</p> <p>Doctor: Yes, I think it's a bit better. It's a bit. It's a bit. It's not very clear. But let's continue. Anyway,</p> <p>Patient: Okay.</p> <p>Doctor: okay, let's talk again. So, how can I help you, sir?</p> <p>Patient: Yes, so it's been a few days now. I have like a sore and the Redskin it's kind of it's really itchy and it's like super annoying.</p> <p>Doctor: Okay.</p> <p>Patient: So I'd like to find something quick to serve it.</p> <p>Doctor: No, no problem. Happy to help whereabouts of your skin is affected.</p> <p>Patient: Mostly like my chest my my hands my arms like agree. It's super annoying like it's itching a lot like all the time and I can't even sleep at night. Like I really need something quickly to study because even at work I like when I'm in the meeting and I have to like think about my work Focus like actually focus on my work. It's</p> <p>Doctor: Yeah.</p> <p>Patient: really annoying because I can actually think about what happened say, I'm always like disturbed by this disease.</p> <p style="text-align: center;">* * *</p> <p>Doctor: It didn't okay. So something for you to think about a you can get different types of and system means I can give you something Little Bit Stronger today as well</p> <p>Patient: Okay.</p> <p>Doctor: something like fix the penalty in which I can give to you today. It's definitely worth trying it's not gonna do you any harm but I say anything using the steroids and the emollients on a regular basis over the next week to 10 days should hopefully care control your symptoms, but do come back and see me next week if things don't get better.</p> <p>Patient: That sounds good.</p> <p>Doctor: Okay any questions for me?</p> <p>Patient: And now that's it.</p> <p>Doctor: Okay. Well, I wish you all the best.</p> <p>Patient: Thank you very much.</p> <p>Doctor: Hope you have a good day.</p> <p>Patient: Bye-bye.</p>

Table A.2: An example of a human transcript and a Google Speech-to-text transcript for one of the mock consultations. Both are partial.

Human Note	Hx: 1 week history of spontaneous elbow swelling left. Not painful. No trauma. No FH of rheumatological disease- NB pt says he has been old he has OA previously by doctors- ? need to confirm this Works in a desk job Not happened before Otherwise well- PMHx: nil of note FH: nil of note DH: not on any medication, allergic to peanuts SH: exercises regularly, active Ex: looks well, not in pain. Mild erythema and minimal swelling (if any) around olecranon process left elbow Imp: possible bursitis Plan: for NSAIDs- usual advice re SE For rheum bloods: esr, crp, fbc, rheum factor and urate Review thereafter in person/ via video To contact us back in interim if any deterioration/concerns- pt warned re symptoms of septic arthritis.
BART-CNN	Doctor Deen Mirza from GP at Hand sees John Smith. John says he has a weird swelling on his left elbow. He also says he is allergic to peanuts. Deen takes a look at John's elbow to see if there is anything wrong with it. Do you have any other illnesses at all?
BERT-ext	Before we start your appointment, could you please tell me your first name and your date of birth. And I was born on the fifth of April, , nineteen seventy three. But it's just, just a bit, a bit weird, to see that. , and , , in terms of your job, do you do anything physical? so you know you said you think you've got , , osteoarthritis. and, do you have any other illnesses at all? , I run regularly, like two, three times a week. , what I think we should do is, I think you should be on some anti-inflammatory medication, in the, in the first instance. And, there'll be instructions within that pack, about where to go to get those blood tests done. and , your, your joint doesn't look like that. However, if your, the elbow was to become very red, very painful, , and the redness was to spread or become , you know more intense. That would require more immediate assessment, more immediate treatment. do you, do you think it's something dangerous? Like something, like could I die from that, or is it, is it No. that's four hundred milligrams, two times a day. Maybe within a , actually you know, the follow-up appointment doesn't have to be face-to-face, if it's more convenient for you do, to do it over the phone, we can do that over the phone, , over video. We can do that as well, that's, that's your call.
Random	Sure. No, no I haven't noticed that before. OK, OK, great. Yes, a few years ago. do you, do you think it's something dangerous? Fantastic. But you contact us, , after you've had the blood test done, and we can review things then, OK. OK. OK, yeah that sounds good. OK. -. , yeah, no, I'm, think I'm healthy. . So, , this, this is not the case right now. I run regularly, like two, three times a week. don't need to worry. All right then, OK. , take care then.
BART-finet	You have a problem with your left elbow. 1 week ago noticed a weird swelling on the left elbow. Not painful at all, but slightly warm, slightly warm. No pain, no swelling, no fluid in the elbow. No injury. No previous history of this. No injury to the elbow. NKDA. SH: Mobile and active, exercise 2-3 times a week, running. Osteoarthritis of the elbow. You should start the treatment you have been prescribed. You should begin the treatment prescribed as we discussed. You may want to take some ibuprofen or paracetamol in addition to any prescribed medication.

Table A.3: Examples of a human written note and automatically generated notes with the four baseline models.