


# Comparative Experimentation of Accuracy Metrics in Automated Medical Reporting: The Case of Otitis Consultations

Wouter Faber<sup>\*1</sup> , Renske Eline Bootsma<sup>\*1</sup> , Tom Huibers<sup>2</sup>, Sandra van Dulmen<sup>3</sup>   
and Sjaak Brinkkemper<sup>1</sup> 

<sup>1</sup>Department of Information and Computing Sciences, Utrecht University, Utrecht, The Netherlands

<sup>2</sup>Verticai, Utrecht, The Netherlands

<sup>3</sup>Nivel: Netherlands institute for health services research, Utrecht, The Netherlands

**Keywords:** Automated Medical Reporting, Accuracy Metric, SOAP Reporting, Composite Accuracy Score, GPT.

**Abstract:** Generative Artificial Intelligence (AI) can be used to automatically generate medical reports based on transcripts of medical consultations. The aim is to reduce the administrative burden that healthcare professionals face. The accuracy of the generated reports needs to be established to ensure their correctness and usefulness. There are several metrics for measuring the accuracy of AI generated reports, but little work has been done towards the application of these metrics in medical reporting. A comparative experimentation of 10 accuracy metrics has been performed on AI generated medical reports against their corresponding General Practitioner's (GP) medical reports concerning Otitis consultations. The number of missing, incorrect, and additional statements of the generated reports have been correlated with the metric scores. In addition, we introduce and define a Composite Accuracy Score which produces a single score for comparing the metrics within the field of automated medical reporting. Findings show that based on the correlation study and the Composite Accuracy Score, the ROUGE-L and Word Mover's Distance metrics are the preferred metrics, which is not in line with previous work. These findings help determine the accuracy of an AI generated medical report, which aids the development of systems that generate medical reports for GPs to reduce the administrative burden.


## 1 INTRODUCTION


"Administrative burden is real, widespread and has serious consequences (Heuer, 2022)."


Although the Electronic Health Record (EHR) has its benefits, the consequences regarding time and effort are increasingly noticed by medical personnel (Olivares Bøgeskov and Grimshaw-Aagaard, 2019; Moy et al., 2021). In addition to less direct patient care (Lavander et al., 2016), documentation sometimes shifts to after hours: studies of (Anderson et al., 2020) and (Saag et al., 2019) show that physicians spend hours on documentation at home. Working after hours, a poor work/life balance, stress, and using Health Information Technology such as EHRs are associated with


less work-life satisfaction and the risk of professional burnout (Shanafelt et al., 2016; Robertson et al., 2017; Gardner et al., 2018; Hauer et al., 2018).

The notation used for documentation of General Practitioner (GP) consultations is the Subjective, Objective, Assessment and Plan (SOAP) notation, which has been widely used for clinical notation, and dates back to 1968 (Weed, 1968; Heun et al., 1998; Sapkota et al., 2022). Based on a consultation, the SOAP note has to be written by the clinician, to update the EHR. To reduce the administrative burden, generative Artificial Intelligence (AI) can be used to summarize transcripts of medical consultations into automated reports, which is also the purpose of the Care2Report program to which this study belongs (Molenaar et al., 2020; Maas et al., 2020; Kwint et al., 2023; Wegstapel et al., 2023). Care2Report (C2R) aims at automated medical reporting based on multimodal recording of a consultation and the generation and uploading of the report in the electronic medical record system (Brinkkemper, 2022). However, for these reports to be useful, the accuracy of the generated report has to

<sup>a</sup>  <https://orcid.org/0009-0001-8412-9009>

<sup>b</sup>  <https://orcid.org/0009-0008-0667-2517>

<sup>c</sup>  <https://orcid.org/0000-0002-1651-7544>

<sup>d</sup>  <https://orcid.org/0000-0002-2977-8911>

\* These authors contributed equally to this work

be determined. Several metrics exist to compare the accuracy of generated text (Moramarco et al., 2022), but the application in the Dutch medical domain has not been researched.

Therefore, this paper proposes research towards metrics in the Dutch medical domain, resulting in the following research question:

**RQ.** *What Is the Preferred Metric for Measuring the Difference Between an Automatically Generated Medical Report and a General Practitioner’s Report?*

This study contributes to the field of AI generated medical reports by providing a case-level look and adds to the larger field of Natural Language Generation (NLG). Furthermore, this work has societal relevance by providing the preferred accuracy measure for AI-generated Dutch medical reports. Being able to identify the accuracy of a medical report, research towards the generation of the reports can be extended, to ensure a high accuracy. Namely, since reports play a crucial role in patient care, diagnosis, and treatment decisions it is vital that generated reports are correct and complete. Reports with high accuracy would prevent the medical staff from spending a lot of time writing the report themselves or correcting the gener-

ated reports, which reduces the administrative burden.

First, a literature review will be presented in section 2. The method and findings will be discussed in section 3 and section 4 respectively, after which these will be discussed in section 5. Conclusions will be drawn and directions for future work will be given in section 6.

## 2 RELATED WORK

Different accuracy metrics for NLG exist, which can be compared in various ways.

### 2.1 Accuracy Metrics

Over the years, different evaluation metrics for measuring the accuracy of NLG systems have been developed, such as BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), and METEOR (Banerjee and Lavie, 2005). All these metrics compare a generated text with a reference text. In recent years, a number of studies have provided an overview of these metrics and divided them into different categories (Sai et al., 2020; Celikyilmaz et al., 2020; Fabbri et al., 2020; Moramarco et al., 2022). Our study adopts most

Table 1: Overview of existing accuracy metrics for Natural Language Generation.

Category	Metrics	Property	Common use <sup>a</sup>
Edit distance	Levenshtein	Cosine similarity	MT, IC, SR, SUM, DG & RG
	WER	% of insert, delete, and replace	SR
	MER	Proportion word matches errors	SR
	WIL	Proportion of word information lost	SR
Embedding	ROUGE-WE	ROUGE + word embeddings	SUM
	Skipthoughts	Vector based similarity	MT
	VectorExtrema	Vector based similarity	MT
	GreedyMatching	Cosine similarity of embeddings	RG
	USE	Sentence level embeddings	MT
	WMD	EMD <sup>b</sup> on words	IC & SUM
	BertScore	Similarity with context embeddings	DG
MoverScore	Context embeddings + EMD <sup>b</sup>	DG	
Text overlap	Precision	% relevant of all text	MT, IC, SR, SUM, DG, QG, & RG
	Recall	% relevant of all relevant	MT, IC, SR, SUM, DG, QG, & RG
	F-Score	Precision and recall	MT, IC, SR, SUM, DG, QG, & RG
	BLEU	<i>n</i> -gram precision	MT, IC, DG, QG, & RG
	ROUGE-n	<i>n</i> -gram recall	SUM & DG
	ROUGE-L	Longest common subsequence	SUM & DG
	METEOR	<i>n</i> -gram with synonym match	MT, IC, & DG
	CHRF	<i>n</i> -gram F-score	MT

<sup>a</sup>Abbreviations for the subfield, as introduced by (Celikyilmaz et al., 2020). **MT**: Machine Translation, **IC**: Image Captioning, **SR**: Speech Recognition, **SUM**: Summarization, **DG**: Document or Story Generation, Visual-Story Generation, **QG**: Question Generation, **RG**: Dialog Response Generation.

<sup>b</sup>EMD = Earth Mover’s Distance.

of the metrics and categories of (Moramarco et al., 2022), which was inspired by the categories stated by (Sai et al., 2020) and (Celikyilmaz et al., 2020). Metrics that are not specifically developed for summarization are also included, to ensure that the study does not become too narrow. (Moramarco et al., 2022) also introduce a new metric with its own category, the Stanza+Snomed metric, which is not included in this current study since there is no other known work or use of this metric. The remaining three groups of metrics are:

- **Edit distance metrics** count how many characters or words are required to convert the output of the system into the reference text. They include Levenshtein (Levenshtein et al., 1966), Word Error Rate (WER) (Su et al., 1992), Match Error Rate (MER) (Morris et al., 2004), and Word Information Lost (WIL) (Morris et al., 2004).
- **Embedding metrics** use encode units of text and pre-trained models to compute cosine similarity to find a similarity between the units. For this, they use word-level, byte-level, and sentence-level embeddings. The metrics include: ROUGE-WE (Ng and Abrecht, 2015), Skipthoughts (Kiros et al., 2015), VectorExtrema (Forgues et al., 2014), GreedyMatching (Sharma et al., 2017), Universal Sentence Encoder (USE) (Cer et al., 2018), Word Mover’s Distance (WMD) (Kusner et al., 2015), BertScore (Zhang et al., 2019), and MoverScore (Zhao et al., 2019).
- **Text overlap metrics** rely on string matching, and measure the amount of characters, words, or  $n$ -grams that match between the generated text and the reference. These metrics include BLEU (Papineni et al., 2002), ROUGE (Lin, 2004), METEOR (Banerjee and Lavie, 2005), and Character  $n$ -gram F-score (CHRF) (Popović, 2015). The F-measure, based on precision and recall (Marongsi et al., 2019), also falls under this category.

Table 1 provides an overview of all the metrics, along with their category, property, and common use. This table extends data from (Moramarco et al., 2022) and (Celikyilmaz et al., 2020). Information that was not available in these studies was derived from the original papers of the metrics.

## 2.2 Comparison of Metrics

In order to perform an evaluation of the accuracy metrics, a reference accuracy score is necessary to compare the calculated scores. There are different ways to determine these reference scores, which could involve a human evaluation. One could simply ask hu-

man evaluators to compare the generated text with a reference text and rate the general factual accuracy on a scale of 1 to 5 (Goodrich et al., 2019). However, this is a very broad measure that is heavily influenced by subjectivity. Alternatively, other studies used different dimensions to compare generated texts, such as *Adequacy*, *Coherence*, *Fluency*, *Consistency*, and *Relevance* (Fabbri et al., 2020; Kryściński et al., 2019; Turian et al., 2003). (Moramarco et al., 2022), use *Omissions*, *Incorrect statements*, and *Post-edit times* to evaluate automatically generated medical reports.

The ratings of the human evaluators can be compared with the results of the metrics, using correlation measures such as the Spearman, Pearson, or Kendall’s  $\tau$  coefficient (Goodrich et al., 2019; Moramarco et al., 2022; Fabbri et al., 2020).

## 2.3 SOEP Reporting for GPs

In the Netherlands, the SOEP convention is used by GPs for medical reporting, which is the Dutch alternative to SOAP (Van der Werf, 1996). Subjective (S) represents the state, medical history and symptoms of the patient. Objective (O) contains measurable data obtained from past records or results of the examination and Evaluation (E) (or Assessment (A)) offers the opportunity to note the assessment of the health problem and diagnosis. Finally, Plan (P) contains the consultant’s plan for the patient. Close attention should be paid to the division between the symptoms and signs (subjective descriptions and objective findings) since this is a common pitfall while writing SOEP notes (Podder et al., 2022; Seo et al., 2016).

## 3 RESEARCH METHOD

An overview of the research method of our study is shown in Figure 1, which will be explained in the subsections. The blue outline shows our research focus.

### 3.1 Materials

The C2R program provides data from seven transcripts of medical consultations between GPs and their patients concerning ear infections, namely *Otitis Externa* ( $n = 4$ ) and *Otitis Media Acuta* ( $n = 3$ ) (Maas et al., 2020). These transcripts are derived from video recordings, for which both the patients and GPs provided informed consent. The recordings were made as part of a study by Nivel (Netherlands institute for health services research) and Radboudumc

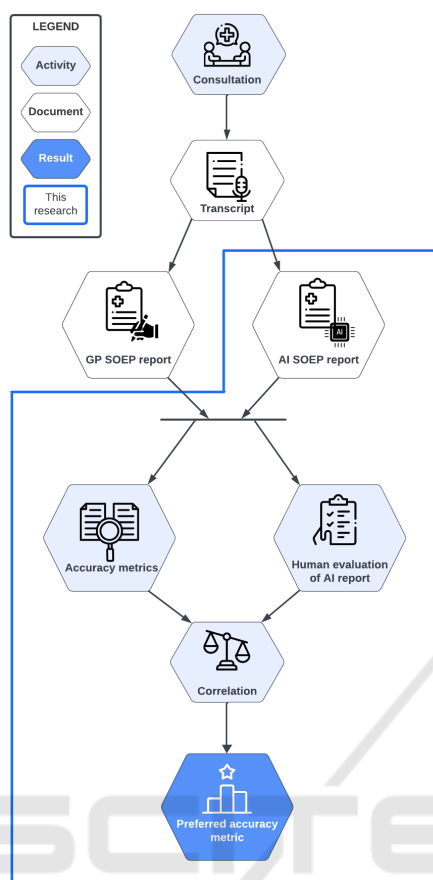


Figure 1: Research Method Diagram, showing the method along with its input and intended output.

to improve GP communication (Houwen et al., 2017; Meijers et al., 2019).

Based on the transcripts, GPs wrote a SOEP report (referred to as GP report), which is considered the ground truth for this study. These GPs did not perform the consultation but wrote the report solely based on the transcripts. Furthermore, software of the C2R program which runs on GPT 4.0 was used. The temperature was set to 0 to limit the diversity of the generated text. Based on the formulated prompt and transcript, the GPT generates a SOEP report (referred to as AI report) (Maas et al., 2020).

### 3.2 Pre-Study

Upon first inspection of the GP reports, it was noticed that abbreviations such as "pcm" meaning "paracetamol" are frequently used. To gain more insights into the experience and preferences of medical staff regarding the formulation of SOEP reports, a pre-study was conducted among Dutch medical staff ( $n = 5$ ; 1 physiotherapist, 1 paediatrician, 1 junior doc-

tor / medical director, 1 nurse and 1 nursing student). The participants were asked about their experience with (SOEP) medical reporting, important factors of SOEP, the use of abbreviations, and general feedback on the Care2Report program. All participants indicated to have knowledge of SOEP reporting, and have experience with writing medical reports, using SOEP or similar methods. Distinguishing between Subjective and Objective information was indicated to be a common mistake, which is in line with research (Podder et al., 2022; Seo et al., 2016). In addition, the notation of the Evaluation is important since this is "the essence of the consult", but is sometimes not filled in completely. Regarding abbreviations, the participants were divided. Some of them indicated that they preferred using abbreviations, to enable faster reading, but discouraged the use of difficult abbreviations. The other participants indicated always favouring written terms since this improves readability. All participants favoured using written terms when multiple staff members (from different backgrounds) were involved, for example when it came to patient transfer, with the exception of general abbreviations.

In general, the medical staff agreed that an AI report would be "a great solution" that "saves time, which enables more consultation time". In addition, two of the medical staff indicated writing reports after the consult due to time limits, which can cause a loss of information. This insight is in line with previous research (Olivares Bøgeskov and Grimshaw-Aagaard, 2019; Moy et al., 2021; Lavander et al., 2016; Anderson et al., 2020; Saag et al., 2019).

### 3.3 Prompt for Report Generation

The GPT software does not have the knowledge or capability to use medical abbreviations like GPs use in their SOEP report. This can result in the metrics falsely identifying differences between written terms and their abbreviations. That, in combination with the fact that using written terms was preferred by half of the medical staff of the pre-study and since the reports will be read by staff members from different disciplines, led to the decision to change abbreviations in the GP's report to the full expression.

The GPT was given a Dutch prompt, of which the translated text is given in Listing 1. The formulation of the prompt was based on existing research within the C2R program (line 1, 2, 3, 4, 5, 7, 9), and has been adapted to incorporate the input of the medical staff (line 3, 4, 5, 6, 8) and literature (line 3, 4, 5). Mainly, the division between symptoms and signs (Subjective and Objective) and the definition of the Evaluation category have been added.

```

1 Write a medical s.o.e.p report based
  on a conversation between a gp and
  a patient and use short and
  concise sentences.
2 Report in the categories of
  subjective, objective, evaluation,
  and plan.
3 Make sure that for subjective the
  description of the complaints of
  the patient is noted.
4 Also, possible pain medication which
  is used by the patient and the
  information that emerges from the
  anamnesis may be noted here.
5 At objective, the observation of the
  symptoms by the gp and the results
  of the physical examination must
  be noted.
6 The evaluation contains the judgement
  of the examination and the
  diagnosis.
7 The treatment plan must be clear from
  the plan.
8 Make sure that the medical terms,
  such as the name of the medication
  are noted.
9 The content of the report must be
  derived from the given transcript.

```

Listing 1: Prompt used as input for the GPT.

### 3.4 Metric Selection and Execution

For this study, a spread of metrics between categories (see subsection 2.1) was chosen. More popular or common metrics were preferred due to frequent application and public availability. The following 10 metrics are part of the selection: *Levenshtein*, *WER*, *ROUGE-1*, *ROUGE-2*, *ROUGE-L*, *BLEU*, *F-Measure*, *METEOR*, *BertScore*, *WMD*.

Each accuracy metric is applied to the AI report with the GP report as reference. Five of the metrics could be run via an online application and the F-Measure was calculated using an R function. In addition, the embedding metrics, *BertScore* and *WMD*, required running Python code. For these metrics, Dutch embeddings were used. *BertScore* supports more than 100 languages via multilingual BERT, including Dutch, and for *WMD* the "dutch-word-embeddings" Word2Vec model was used (Nieuwenhuijse, 2018). *METEOR* uses *n*-gram comparison with synonym match. At the time of writing, there is no alternative for Dutch texts. Therefore, *METEOR* will mostly rely on *n*-gram comparison.

### 3.5 Human Evaluation

Concurrently, the AI reports are compared with the GP reports by the first authors, i.e., the human evaluation. This is inspired by the work of (Moramarco et al., 2022). This method of evaluation is adopted because it is a domain-specific method that includes the accuracy of the report and provides insight into the amount of work needed by the GP as well. For each AI report, seven aspects will be counted, as can be seen in Table 2.

Table 2: Human evaluation aspects along with their descriptions and abbreviations.

Aspect	Description	Abr.
Missing	Missing in AI report	MIS
Incorrect	Incorrect in AI report	INC
Added On-topic	Not in GP report, on-topic	ADD <sub>ON</sub>
Added Off-topic	Not in GP report, off-topic	ADD <sub>OFF</sub>
Post-edit time	Time (s) to correct AI report	PET
Nr. of characters	Nr. of characters in AI report	NRC
Word length	Avg. word length in AI report	WLE

Firstly, the number of *Missing statements* and *Incorrect statements*, which include wrongly stated information. Next, the *Additional statements*, which are divided into *On-topic* and *Off-topic*. *Added On-topic* statements contain information that is not present in the GP report but relates to the content, e.g. "There is no pus visible, but there is blood leaking from the ear". *Added Off-topic* statements contain information that is not present in the GP report and does not relate to the content, e.g. "The patient called in sick for work". In addition, the *Number of characters* and the number of words will be counted, to calculate the *Word length*. An independent samples *t*-test will be performed on the *Number of characters* and the *Word length* between the AI report and the GP report to gain more insight into a potential difference in report length. Lastly, the *Post-edit time* describes the time it takes to correct the AI report, i.e., adding *Missing*, changing *Incorrect* and removing *Additional* statements. This is interesting to consider since the goal of AI reporting is to reduce the time spent on reporting by GPs.

After performing the human evaluation, Pearson correlation coefficients are calculated between the aspects of the human evaluation and every metric, excluding the *Word length*, and *Number of characters*. In theory, the stronger the negative correlation, the more effective the metric is in the domain of medical reporting. Namely, an AI report ideally has low missing statements, low incorrect and low additional statements.

To compare the metrics, a single *Composite Ac-*

curacy Score (CAS) is calculated for each metric. For this, the correlations per *Missing*, *Incorrect* and *Added* statements with the metric are normalised on a scale from 0 to 1, where 0 is the lowest (negative) correlation and 1 is the highest. Based on the normalised correlations with the *Missing (MIS)*, *Incorrect (INC)*, *Added On-topic (ADD<sub>ON</sub>)*, and *Added Off-topic (ADD<sub>OFF</sub>)* statements, the *Composite Accuracy Score (CAS)* is calculated using Formula 1.

$$CAS = \frac{MIS + INC + ADD_{OFF} + 0.5 \times ADD_{ON}}{3.5} \quad (1)$$

Every score has a weight of 1.0, except the *Added On-topic* statements, which have been attributed a weight of 0.5 since their presence in the AI report is deemed less severe than the other aspects. The *Post-edit time* is not part of the *Composite Accuracy Score* because it is dependent on the other aspects of the human evaluation.

With respect to editing, a metric can be considered preferred if it has a low *Composite Accuracy Score* as well as a strong negative correlation with the *Post-edit time*. If a metric fulfils these requirements, it is an adequate tool to measure the accuracy of the report itself as well as the administrative burden.

## 4 FINDINGS

Performing the method resulted in measured human evaluation aspects, correlations and the *Composite Accuracy Scores*.

### 4.1 Human Evaluation

As mentioned in subsection 3.5, the seven human evaluation aspects were counted for each AI report, of which the first five can be seen in Table 3. Often, the error in question was *Missing* in the AI report. *Incorrect* statements were statements which were wrongly stated or wrongly defined as being said by the GP. Of the statements that were not in the GP report (*Added*), the distinction between *On-topic* and *Off-topic* was less direct. Mainly, *On-topic* statements contained additional information regarding the medical history, complaints or treatment. Statements regarding other topics then discussed in the GP report and explanations to the patient were classified as *Off-topic* because these would not be of any relevance to the SOEP report, written by the GP.

The *Number of characters* and the number of words were used to calculate the *Word length* for both the GP report and AI report. The results of the independent samples *t*-test show that the AI reports are significantly longer in terms of characters ( $1199.29 \pm$

Table 3: Human evaluation aspects per AI report. The averages of the aspects are rounded to whole numbers.

Human Evaluation Aspects	R1	R2	R3	R4	R5	R6	R7	Average
Missing statements	12	5	8	7	9	8	7	8
Incorrect statements	2	1	1	2	1	5	1	2
Added statements - On-topic	6	9	5	7	8	3	7	6
Added statements - Off-topic	5	5	0	2	1	5	0	3
Post-edit time (s)	378	196	170	213	193	186	169	215

Table 4: Pearson correlation between human evaluation aspects and metrics, along with the *Composite Accuracy Score*.

Metric	Miss.	Incorr.	Additional		CAS	PET
			On-topic	Off-topic		
<i>Levenshtein</i>	0.122	-0.178	-0.011	-0.798	<b>0.229</b>	-0.320
<i>WER</i>	0.673	-0.042	-0.409	-0.315	0.434	0.385
<i>BertScore</i>	-0.272	0.126	0.319	0.759	0.618	0.329
<i>WMD</i>	-0.564	-0.168	0.381	-0.289	<b>0.241</b>	<b>-0.591</b>
<i>ROUGE-1</i>	-0.063	0.123	-0.394	-0.634	0.284	<b>-0.483</b>
<i>ROUGE-2</i>	-0.153	0.131	-0.021	-0.259	0.401	-0.201
<i>ROUGE-L</i>	-0.233	-0.109	-0.056	-0.597	<b>0.209</b>	<b>-0.461</b>
<i>BLEU</i>	0.109	-0.013	0.002	-0.462	0.364	-0.258
<i>F-Measure</i>	0.698	0.119	-0.434	-0.339	0.501	0.333
<i>METEOR</i>	0.315	0.467	-0.220	0.045	0.677	0.103

The negative correlations are indicated by different intensities of orange, and the positive correlations are indicated by different intensities of blue. The three lowest Composite Accuracy Scores and PET correlations are in **bold**.

197) than the GP report ( $410.71 \pm 94.32$ ),  $t(12) = -9.520$ ,  $p < 0.001$ . The words used in the AI report ( $6.04 \pm 0.33$ ) are significantly shorter than in the GP report ( $7.62 \pm 0.29$ ),  $t(12) = 9.555$ ,  $p < 0.001$ .

## 4.2 Correlation Between Metrics and Human Evaluation Aspects

Using the metric scores and the human evaluation aspects, the mutual correlation has been calculated. In contrast to the other metrics, for the edit distance metrics (Levenshtein and WER) a low score equals a good accuracy. To enable easier comparison between the metrics, the correlations of the edit distance metrics have been inverted by multiplying with -1. All correlations between the metrics and human evaluation aspects are shown in Table 4. Ideally, these correlations are strongly negative since the medical report should be concise and contain all, and only, relevant information. The three strongest negative correlations with the *Post-edit time* and the three lowest *Composite Accuracy Scores* have been indicated in bold in Table 4

## 5 DISCUSSION

Based on the findings in section 4, notable observations were found.

**Human Evaluation.** The results of the human evaluation (Table 3) show that all AI reports contain on average 9 *Added* statements, compared to the GP report. Consequently, the AI reports are longer than the GP reports. Most *Added* statements are *On-topic*. Despite the additional information and length of the AI report, each report misses on average 8 statements.

**Added Statements.** The first noticeable correlation in Table 4 appears between more than half of the metrics correlating moderately ( $-0.3 < r < -0.5$ ) or strongly ( $r < -0.5$ ) negatively with *Added Off-topic* statements. Interestingly, only three metrics have a moderate negative correlation with *Added On-topic* statements. This can be explained by the *On-topic* statements adding extra, relevant, information to the content of the SOEP report. Even though these statements are added, the metrics could define this as relevant information thus not having a strong negative correlation.

**Comparison with (Moramarco et al., 2022).** Except for the WMD metric, none of the metrics

strongly correlate negatively with *Missing* statements and *Post-edit time*, which is not in line with the findings of (Moramarco et al., 2022). In their findings, METEOR and BLEU scored good on detecting *Missing* statements and Levenshtein and METEOR rank highly on the *Post-edit time*. Additionally, none of the metrics moderately or strongly correlate negatively with *Incorrect* statements, which is also not in line with the findings of (Moramarco et al., 2022), where ROUGE scored good on identifying these statements.

**Opposite of Preferred Correlations.** The WER and F-Measure strongly correlate ( $r > 0.5$ ) *Missing* statements with better accuracy, and BertScore correlates a high number of *Off-topic* statements with better accuracy. These results indicate exactly the opposite of what is preferred and therefore seem to be less suitable for the evaluation of automatically generated reports.

**Post-Edit Time.** Six metrics have a negative correlation with the *Post-edit time*. WMD, ROUGE-1, and ROUGE-L have the strongest negative correlation, meaning that they are the preferred metrics concerning the correlation with *Post-edit time*.

**Composite Accuracy Score.** When looking at the CAS, The WER, BertScore, F-Measure and METEOR metrics score high ( $> 0.5$ ), indicating that these metrics are not suitable for the current application. The high CAS of the BertScore is remarkable since this metric performed as one of the best in the study of (Moramarco et al., 2022). The high CAS of METEOR could be explained due to the fact that the used transcripts are Dutch, which is an unsupported language by the metric. Therefore, it cannot use synonym matching, which is the added benefit of the METEOR metric compared to other text overlap metrics. There is no consensus within the categories of edit distance, embedded and text overlap metrics. Consequently, no conclusions can be drawn regarding preferred performing categories.

**Preferred Metrics.** Based on the CAS and the *Post-edit time* correlations, ROUGE-L and WMD are the preferred metrics, since they are in the top 3 for both. The WMD scores slightly worse in terms of CAS, which can be explained by the fact that it has a positive correlation (0.381) with the *Added On-topic* statements, whereas ROUGE-L has only negative correlations with the human evaluation metrics. However, WMD scores better than ROUGE-L when looking at the *Post-edit time*. (Moramarco et al., 2022)

found that Levenshtein, BertScore, and METEOR are the most suitable metrics, which does not correspond with the findings of our work.

## 6 CONCLUSION

AI generated medical reports could provide support for medical staff. These reports should be as accurate as possible, to limit the time needed by the medical staff making corrections. To determine the accuracy of a text, metrics can be used. This research investigated the performance of 10 accuracy metrics by calculating the correlation between the metric score and the following human evaluation aspects: *Missing* statements, *Incorrect* statements, *Additional* statements and *Post-edit time*.

For each metric, the *Composite Accuracy Score* has been calculated, indicating its performance. Based on the *CAS* and the correlation with the *Post-edit time*, the *ROUGE-L* and *Word Mover's Distance (WMD)* metrics are preferred to use in the context of medical reporting, answering the research question:

*What Is the Preferred Metric for Measuring the Difference Between an Automatically Generated Medical Report and a General Practitioner's Report?*

Based on the results, we see that there is a diversity among the applications of the different metrics. Both strong positive and negative correlations with the human evaluation aspects are found, which can be explained by the different methods used by the metrics. The preferred method depends heavily on the context of use and which aspect is deemed more important. Therefore, no unambiguous answer can be given. However, we created the *CAS* score based on our context of use, identifying the preferred metrics in the context of medical reporting.

### 6.1 Limitations

The outcome of this study is not in line with previous research, which could be due to the limitations. There are three main limitations to this study.

Firstly, the data set used for running the accuracy metrics consists of just seven AI reports. Additionally, the transcripts used are all from GP consultations on *Otitis Externa* and *Otitis Media Acuta*, making the data limited in its medical diversity. These factors make it difficult to draw general conclusions on accuracy metrics that work for all AI generated medical reports.

Adding to that, the GP reports were written solely based on the transcripts, and not by the GP who per-

formed the consultation, which is not standard practice.

Lastly, the human evaluation was performed by researchers who have no prior experience in writing medical reports. Even though medical staff was consulted in this study, it would be preferred if the human evaluation was done by people with medical expertise. That way, those with a deeper comprehension of what should be included in a report could handle the more challenging cases of evaluating the generated statements' relevance.

### 6.2 Future Work

The main limitations should be addressed in future work. Mainly, the study should be repeated with more medical reporting, on other pathologies. Besides, it would improve the quality of the study if the human evaluation were executed by healthcare professionals. Furthermore, the current AI reports result in low accuracy scores for each metric. Therefore, it would be beneficial if further research was done into optimizing the prompt formulation, resulting in more accurate AI reports. Additionally, the human evaluation of this study does not take wrongly classified statements to the SOEP categories into account, which could be adopted in future work. Finally, the use of abbreviations in generated reports could be further explored, since this was taken out of the equation for this study.

## ACKNOWLEDGEMENTS

Our thanks go to the medical staff who helped us with the pre-study. In addition, the icons of Flaticon.com enabled us to create Figure 1. Finally, many thanks go to Bakkenist for the support of this research project.

## REFERENCES

- Anderson, J., Leubner, J., and Brown, S. (2020). Ehr overtime: an analysis of time spent after hours by family physicians. *Family Medicine*, 52(2):135–137.
- Banerjee, S. and Lavie, A. (2005). Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72.
- Brinkkemper, S. (2022). Reducing the administrative burden in healthcare: Speech and action recognition for automated medical reporting. In Buchmann, R. A., Silaghi, G. C., Bufnea, D., Niculescu, V., Czibula, G., Barry, C., Lang, M., Linger, H., and Schneider,



- C., editors, *Proceedings of the Information Systems Development: Artificial Intelligence for Information Systems Development and Operations*, Cluj-Napoca, Romania: Babeş-Bolyai University.
- Celikyilmaz, A., Clark, E., and Gao, J. (2020). Evaluation of text generation: A survey. *CoRR*, abs/2006.14799. arXiv.
- Cer, D., Yang, Y., Kong, S.-y., Hua, N., Limtiaco, N., John, R. S., Constant, N., Guajardo-Cespedes, M., Yuan, S., Tar, C., et al. (2018). Universal sentence encoder. *arXiv preprint arXiv:1803.11175*.
- Fabbri, A. R., Kryściński, W., McCann, B., Xiong, C., Socher, R., and Radev, D. (2020). Summeval: Re-evaluating summarization evaluation. *Transactions of the Association for Computational Linguistics*, 9:391–409.
- Forgues, G., Pineau, J., Larchevêque, J.-M., and Tremblay, R. (2014). Bootstrapping dialog systems with word embeddings. In *Nips, modern machine learning and natural language processing workshop*, volume 2, page 168.
- Gardner, R. L., Cooper, E., Haskell, J., Harris, D. A., Poplau, S., Kroth, P. J., and Linzer, M. (2018). Physician stress and burnout: the impact of health information technology. *Journal of the American Medical Informatics Association*, 26(2):106–114.
- Goodrich, B., Rao, V., Liu, P. J., and Saleh, M. (2019). Assessing the factual accuracy of generated text. *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 166–175.
- Hauer, A., Waukau, H., and Welch, P. (2018). Physician burnout in wisconsin: An alarming trend affecting physician wellness. *Wmj*, 117(5):194–200.
- Heuer, A. J. (2022). More evidence that the healthcare administrative burden is real, widespread and has serious consequences comment on” perceived burden due to registrations for quality monitoring and improvement in hospitals: A mixed methods study”. *International Journal of Health Policy and Management*, 11(4):536.
- Heun, L., Brandau, D. T., Chi, X., Wang, P., and Kangas, J. (1998). Validation of computer-mediated open-ended standardized patient assessments. *International Journal of Medical Informatics*, 50(1):235–241.
- Houwen, J., Lucassen, P. L., Stappers, H. W., Assendelft, W. J., van Dulmen, S., and Olde Hartman, T. C. (2017). Improving gp communication in consultations on medically unexplained symptoms: a qualitative interview study with patients in primary care. *British Journal of General Practice*, 67(663):e716–e723.
- Kiros, R., Zhu, Y., Salakhutdinov, R. R., Zemel, R., Urtasun, R., Torralba, A., and Fidler, S. (2015). Skip-thought vectors. *Advances in neural information processing systems*, 28.
- Kryściński, W., Keskar, N. S., McCann, B., Xiong, C., and Socher, R. (2019). Neural text summarization: A critical evaluation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 540–551.
- Kusner, M., Sun, Y., Kolkin, N., and Weinberger, K. (2015). From word embeddings to document distances. In *International conference on machine learning*, pages 957–966. PMLR.
- Kwint, E., Zoet, A., Labunets, K., and Brinkkemper, S. (2023). How different elements of audio affect the word error rate of transcripts in automated medical reporting. *Proceedings of BIOSTEC*, 5:179–187.
- Lavander, P., Meriläinen, M., and Turkki, L. (2016). Working time use and division of labour among nurses and health-care workers in hospitals—a systematic review. *Journal of Nursing Management*, 24(8):1027–1040.
- Levenshtein, V. I. et al. (1966). Binary codes capable of correcting deletions, insertions, and reversals. In *Soviet physics doklady*, volume 10, pages 707–710. Soviet Union.
- Lin, C.-Y. (2004). Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Maas, L., Geurtsen, M., Nouwt, F., Schouten, S., Van De Water, R., Van Dulmen, S., Dalpiaz, F., Van Deemter, K., and Brinkkemper, S. (2020). The care2report system: Automated medical reporting as an integrated solution to reduce administrative burden in healthcare. In *HICSS*, pages 1–10.
- Maroengsit, W., Piyakulpinyo, T., Phonyiam, K., Pongnumkul, S., Chaovalit, P., and Theeramunkong, T. (2019). A survey on evaluation methods for chatbots. In *Proceedings of the 2019 7th International conference on information and education technology*, pages 111–119.
- Meijers, M. C., Noordman, J., Spreeuwenberg, P., Olde Hartman, T. C., and van Dulmen, S. (2019). Shared decision-making in general practice: an observational study comparing 2007 with 2015. *Family practice*, 36(3):357–364.
- Molenaar, S., Maas, L., Burriel, V., Dalpiaz, F., and Brinkkemper, S. (2020). Medical dialogue summarization for automated reporting in healthcare. In Dupuy-Chessa, S. and Proper, H. A., editors, *Advanced Information Systems Engineering Workshops*, pages 76–88, Cham. Springer International Publishing.
- Moramarcio, F., Korfiatis, A. P., Perera, M., Juric, D., Flann, J., Reiter, E., Savkov, A., and Belz, A. (2022). Human evaluation and correlation with automatic metrics in consultation note generation. In *ACL 2022: 60th Annual Meeting of the Association for Computational Linguistics*, pages 5739–5754. Association for Computational Linguistics.
- Morris, A. C., Maier, V., and Green, P. (2004). From wer and ril to mer and wil: improved evaluation measures for connected speech recognition. In *Eighth International Conference on Spoken Language Processing*.
- Moy, A. J., Schwartz, J. M., Chen, R., Sadri, S., Lucas, E., Cato, K. D., and Rossetti, S. C. (2021). Measurement of clinical documentation burden among physicians and nurses using electronic health records: a scoping review. *Journal of the American Medical Informatics Association*, 28(5):998–1008.
- Ng, J. P. and Abrecht, V. (2015). Better summarization evaluation with word embeddings for rouge. In *Proceed-*

- ings of the 2015 Conference on Empirical Methods in Natural Language Processing, pages 1925–1930.
- Nieuwenhuijse, A. (2018). Coosto - Dutch Word Embeddings. <https://github.com/coosto/dutch-word-embeddings>, Accessed 2023-10-31.
- Olivares Bøgeskov, B. and Grimshaw-Aagaard, S. L. S. (2019). Essential task or meaningless burden? nurses' perceptions of the value of documentation. *Nordic Journal of Nursing Research*, 39(1):9–19.
- Papineni, K., Roukos, S., Ward, T., and Zhu, W.-J. (2002). Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th annual meeting of the Association for Computational Linguistics*, pages 311–318.
- Podder, V., Lew, V., and Ghassemzadeh, S. (2022). Soap records. In: StatPearls Publishing, <https://www.ncbi.nlm.nih.gov/books/NBK482263/>, Accessed 2023-10-20.
- Popović, M. (2015). chrF: character n-gram f-score for automatic mt evaluation. In *Proceedings of the tenth workshop on statistical machine translation*, pages 392–395.
- Robertson, S. L., Robinson, M. D., and Reid, A. (2017). Electronic health record effects on work-life balance and burnout within the i3 population collaborative. *Journal of graduate medical education*, 9(4):479–484.
- Saag, H. S., Shah, K., Jones, S. A., Testa, P. A., and Horwitz, L. I. (2019). Pajama time: working after work in the electronic health record. *Journal of general internal medicine*, 34:1695–1696.
- Sai, A. B., Mohankumar, A. K., and Khapra, M. M. (2020). A survey of evaluation metrics used for nlg systems. *ACM Computing Surveys*, 55.
- Sapkota, B., Shrestha, R., and Giri, S. (2022). Community pharmacy-based soap notes documentation. *Medicine*, 101(30).
- Seo, J.-H., Kong, H.-H., Im, S.-J., Roh, H., Kim, D.-K., Bae, H.-o., and Oh, Y.-R. (2016). A pilot study on the evaluation of medical student documentation: assessment of soap notes. *Korean journal of medical education*, 28(2):237.
- Shanafelt, T. D., Dyrbye, L. N., Sinsky, C., Hasan, O., Satele, D., Sloan, J., and West, C. P. (2016). Relationship between clerical burden and characteristics of the electronic environment with physician burnout and professional satisfaction. *Mayo Clinic Proceedings*, 91(7):836–848.
- Sharma, S., El Asri, L., Schulz, H., and Zumer, J. (2017). Relevance of unsupervised metrics in task-oriented dialogue for evaluating natural language generation. *CoRR*, abs/1706.09799.
- Su, K.-Y., Wu, M.-W., and Chang, J.-S. (1992). A new quantitative quality measure for machine translation systems. In *COLING 1992 Volume 2: The 14th International Conference on Computational Linguistics*.
- Turian, J. P., Shen, L., and Melamed, I. D. (2003). Evaluation of machine translation and its evaluation.
- Van der Werf, G. T. (1996). Probleemlijst, soep en icpc. *Huisarts Wet*, 39:265–70.
- Weed, L. (1968). Medical records that guide and teach. *New England Journal of Medicine*, 278(11):593–600. PMID: 5637758.
- Wegstapel, J., den Hartog, T., Mick Sneekes, B. S., van der Scheer-Horst, E., van Dulmen, S., and Brinkkemper, S. (2023). Automated identification of yellow flags and their signal terms in physiotherapeutic consultation transcripts. In *Proceedings of the 16th International Joint Conference on Biomedical Engineering Systems and Technologies, BIOSTEC*, volume 5, pages 530–537.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., and Artzi, Y. (2019). BERTScore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Zhao, W., Peyrard, M., Liu, F., Gao, Y., Meyer, C. M., and Eger, S. (2019). Moverscore: Text generation evaluating with contextualized embeddings and earth mover distance. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.