

Characterizing the Value of Information in Medical Notes

Chao-Chun Hsu¹, Shantanu Karnwal²,
Sendhil Mullainathan³, Ziad Obermeyer⁴, Chenhao Tan²

¹ University of Chicago, ² University of Colorado Boulder

³ Chicago Booth School of Business, ⁴ University of California, Berkeley

chaochunh@uchicago.edu

{shantanu.karnwal, chenhao.tan}@colorado.edu

sendhil@chicagobooth.edu, zobermeyer@berkeley.edu

Abstract

Machine learning models depend on the quality of input data. As electronic health records are widely adopted, the amount of data in health care is growing, along with complaints about the quality of medical notes. We use two prediction tasks, readmission prediction and in-hospital mortality prediction, to characterize the value of information in medical notes. We show that as a whole, medical notes only provide additional predictive power over structured information in readmission prediction. We further propose a probing framework to select parts of notes that enable more accurate predictions than using all notes, despite that the selected information leads to a distribution shift from the training data (“all notes”). Finally, we demonstrate that models trained on the selected *valuable* information achieve even better predictive performance, with only 6.8% of all the tokens for readmission prediction.

1 Introduction

As electronic health records (EHRs) are widely adopted in health care, medicine is increasingly an information science (Stead et al., 2011; Shortliffe, 2010; Krumholz, 2014): *Obtaining* and *analyzing* information is critical for the diagnosis, prognosis, treatment, and prevention of disease. Although EHRs may increase the accuracy of storing structured information (e.g., lab results), there are growing *complaints* about unstructured medical notes (henceforth “*notes*”) (Gawande, 2018; Payne et al., 2015; Hartzband et al., 2008).

These complaints can be grouped into two perspectives: consumption and production. On the one hand, information overload poses a critical challenge on the consumption side. That is, the sheer amount of information makes it difficult to glean meaningful information from EHRs, including notes (Weir and Nebeker, 2007).

On the other hand, from the perspective of production, for every hour spent on patient interaction, physicians have an added one-to-two hours finishing the progress notes and reviewing results among other things, without extra compensation (Patel et al., 2018). The additional work contributes to physician burnout, along with low-quality notes and even errors in the notes. Consequently, physicians tend to directly copy large volumes of patient data into notes, but may fail to record information *only* available through interaction with patients. For instance, they may miss the wheezing breath for the diagnosis of the chronic obstructive pulmonary disease, or fail to have engaging conversations for evaluating signs of depression (Zeng, 2016).

While the NLP community has focused on alleviating the challenges in analyzing information (e.g., information overload), we argue that it is equally important to help caregivers *obtain* and *record* valuable information in the first place. We aim to take a first step towards this direction by characterizing the value of information in medical notes computationally. In this work, we define *valuable information* as information that is useful for evaluating medical conditions and making medical decisions.

To do that, we first examine the value of notes *as a whole* conditioned on structured information. While narrative texts can *potentially* provide valuable information only accessible through physician-patient interaction, our analysis addresses the typical complaint that notes contain too many direct copies of structured information such as lab results. Therefore, a natural question is whether notes provide additional predictive power for medical decisions beyond structured information. By systematically studying two critical tasks, readmission prediction and in-hospital mortality prediction, we demonstrate that notes

are valuable for readmission predictions, but not useful for mortality prediction. Our results differ from previous studies demonstrating the effectiveness of notes in mortality prediction, partly because Ghassemi et al. (2014) use a limited set of structured variables and thus achieve limited predictive power with structured information alone.

We then develop a probing framework to evaluate the prediction performance of *parts* of notes selected by value functions. We hypothesize that not all components of notes are equally valuable and some parts of notes can provide stronger predictive power than the whole. We find that discharge summaries are especially predictive for readmission, while nursing notes are most valuable for mortality prediction. Furthermore, we leverage hypotheses from the medical literature to develop interpretable value functions to identify valuable sentences in notes. Similarity with prior notes turns out to be powerful: a mix of most and least similar sentences provide better performance than using all notes, despite containing only a fraction of tokens.

Building on these findings, we finally demonstrate the power of valuable information beyond the probing framework. We show that classification models trained on the selected valuable information alone provide even better predictive power than using all notes. In other words, our interpretable value functions can effectively filter noisy information in notes and lead to better models.

We hope that our work encourages future work in understanding the value of information and ultimately improving the quality of medical information obtained and recorded by caregivers, because information is after all created by *people*.

2 Our Predictive Framework

We investigate the value of notes through a predictive framework. We consider two prediction tasks using MIMIC-III:¹ readmission prediction and mortality prediction. For each task, we examine two questions: 1) does a model trained on both notes and structured information outperform the model with structured information alone? (§3) 2) using a model trained on all notes, are there interpretable ways to identify parts of notes that are more valuable than all notes? (§4)

¹MIMIC official website: <https://mimic.physionet.org/>.

2.1 An Overview of MIMIC-III

MIMIC-III is a freely available medical database of de-identified patient records. This dataset includes basic information about patients such as admission details and demographics, which allows us to identify outcomes of interest such as mortality and readmission. It also contains detailed information that characterizes the patients' health history at the hospital, known as *events*, including laboratory events, charting events, and medical notes. The data derived from these events are elicited while patients are in the hospital. Our goal is to characterize the value of such elicited information, in particular, notes, through predictive experiments. Next, we break down the information into two categories: structured vs. unstructured.

Structured information. The structured information includes the numeric and categorical results of medical measurements and evaluations of patients. For example, in MIMIC-III, structured information includes status monitoring, e.g., respiration rate and blood glucose, and fluids that have been administered to or extracted from the patients.

Notes (unstructured texts). Caregivers, including nurses and physicians, record information based on their interaction with patients in notes. There are fifteen types of notes in MIMIC-III, including nursing notes and physician notes. Table 1 shows the number of notes in each type and their average length.

Not all admissions have notes from caregivers. After filtering patients under 18 and other invalid data (see details in the supplementary material), discharge summary appears in most admissions (96.7%); however, only 0.1% of admissions have consult notes. The most common types of notes include nursing,² radiology, ECG, and physician. There is also significant variation in length between different types of notes. For instance, discharge summary is more than 8 times as long as nursing notes.

Fig. 1 presents the total number of tokens in all types of notes within one admission. As discussed in the introduction, a huge amount of information (11,135 tokens on average) is generated in the form of unstructured texts for a patient in an admission. We hypothesize that not all of them are useful for medical purposes.

²We merge "Nursing" and "Nursing/other" in MIMIC-III.

CATEGORY	COUNT	%	LEN.
Nursing	506,528	73.0	241
Radiology	338,834	83.3	449
ECG	123,042	61.3	43
Physician	92,426	18.2	1369
Discharge summary	47,572	96.7	2195
Echo	34,064	45.8	464
Respiratory	32,798	8.1	205
Nutrition	7,971	6.4	602
General	7,710	6.4	290
Rehab Services	5,321	4.6	622
Social Work	2,294	2.8	446
Case Management	939	1.3	260
Pharmacy	97	0.1	512
Consult	78	0.1	1206

Table 1: Statistics of note events in MIMIC-III after data preprocessing. % denotes the proportion of admissions having this type of note. “LEN.” means the average length.




Figure 1: Distribution of token length of admissions. Average: 11,135 tokens. Median: 6,437 tokens.

2.2 Task Formulation & Data Representation

We consider the following two prediction tasks related to important medical decisions.

- **Readmission prediction.** We aim to predict whether a patient will be re-admitted to the hospital in 30 days after being discharged, given the information collected within one admission.
- **In-hospital mortality prediction.** We aim to predict whether a patient dies in the hospital within one admission. Following [Ghassemi et al. \(2014\)](#), we consider three time periods: 24 hours, 48 hours, and retrospective. The task is most difficult but most useful with only information from the first 24 hours. We thus focus on that time period in the main paper (see the supplementary material for 48 hours and retrospective results).

Formally, our data is a collection of time series with labels corresponding to each task, $\mathcal{D} = \{(E_i, y_i)\}_{i=1}^N$ where N is the number of admissions (instances). For each collection of time series $E = \{(h_t, \tau_t, x_t)\}_{t=1}^T$ of an admission, h_t represents the timestamp (e.g., $h_t = 4.5$ means 4.5

hours after admission) and $\tau_t \in \{0, 1\}$ captures the type of an event (0 indicates that the event contains structured variable and 1 indicates that the event is a note) and x_t stores the value of the corresponding event. Our goal is to predict label $y \in \{0, 1\}$: in readmission, y represents whether a patient was re-admitted within a month. In mortality prediction, y represents whether a patient died in this admission.³

As a result, we obtained a total of 37,798/33,930 unique patients and 46,968/42,271 admissions for readmission/mortality prediction (24 hours).

Representing structured information. As structured information is sparse over timestamps, we filter event types that occur less than 100,000 times (767 event types remaining).⁴ Following [Harutyunyan et al. \(2017\)](#), we represent the time series data of structured variables into a vector by extracting basic statistics of different time windows. Specifically, for events of structured variables, $E_i^{T=0} = \{(h_t, \tau_t, x_t) | \tau_t = 0\}_{t=1}^T$ where $x_t \in \mathbb{R}^d$, $d = 767$, we apply six statistical functions on seven sub-periods to generate $e_i \in \mathbb{R}^{d \times 7 \times 6}$ as the representation of structured variables. The six statistical functions are maximum, minimum, mean, standard deviation, skew, and number of measurements. The seven sub-periods are the entire time period, first (10%, 25%, 50%) of the time period, and last (10%, 25%, 50%) of the time period. We then impute missing values with the mean of training data and apply min-max normalization.

Representing notes. For notes in an admission, we apply sentence and word tokenizers in the NLTK toolkit to each note ([Loper and Bird, 2002](#)). See §2.3 for details on how we use tokenized outcomes for different machine learning models.

2.3 Experimental Setup

Finally, we discuss the experimental setup and models that we explore in this work. Our code is available at <https://github.com/BoulderDS/value-of-medical-notes>.

Data split. Following the training and test split⁵ of patients in [Harutyunyan et al. \(2019\)](#), we use 85% of the patients for training and the rest 15% for testing. To generate the validation set, we first

³Discharge summaries are not used for mortality prediction since they may leak the label.

⁴We exclude structured information such as ICD-9 codes that might have been filled in retrospectively.

⁵Data split can be found here: <https://github.com/YerevaNN/mimic3-benchmarks>.

split 20% of the patients from training set and then collect the admissions under each patient to prevent information leaking for the same patient.

Models. We consider the following models.

- Logistic regression (LR). For notes, we use tf-idf representations. We simply concatenate structured variables with ℓ_2 -normalized tf-idf vector from notes to incorporate structured information. We use scikit-learn (Pedregosa et al., 2011) and apply ℓ_2 regularization to prevent overfitting. We search hyperparameters C in $\{2^x | x \in \mathbb{Z}, -11 \leq x \leq 0\}$.
- Deep averaging networks (DAN) (Iyyer et al., 2015). We use the average embedding of all tokens in the notes to represent the unstructured information, which can be considered a deep version of bag-of-words methods. Similar to logistic regression, we concatenate the structured variables with the average embedding of words in notes to incorporate structured information.
- GRU-D (Che et al., 2018). The key innovation of GRU-D is to account for missing data in EHRs. It imputes the missing value by considering all the information available so far, including how much time it elapses since the last observation and all the previous history. Similar to DAN, we use the average embedding of tokens to represent notes. See details of GRU-D in the supplementary material.

Although it is difficult to apply the family of BERT models to this dataset due to their input length limitation compared to the large number of tokens from all medical notes, we experiment with ClinicalBERT (Alsentzer et al., 2019) based on the selected valuable information in §4.

Evaluation Metrics ROC-AUC is often used in prior work on MIMIC-III (Ghassemi et al., 2014; Harutyunyan et al., 2017). However, when the number of negative instances is much larger than positive instances, the false positive rate in ROC-AUC becomes insensitive to the change of false positive instances. Therefore, area under precision-recall curve (PR-AUC) is considered more informative than ROC-AUC (Davis and Goadrich, 2006). In our experiments, the positive fraction is only 7% and 12% in readmission prediction and mortality prediction respectively.⁶ As

⁶To emulate class ratios in real data, we do not subsample to achieve balanced data.

precision is often critical in medical decisions, we also present precision at 1% and 5%.

3 Do Medical Notes Add Value over Structured Information?

Our first question is concerned with whether medical notes provide any additional predictive value over structured variables. To properly address this question, we need a strong baseline with structured information. Therefore, we include 767 types of structured variables to represent structured information (§2.2). Overall, our results are mixed for readmission prediction and in-hospital mortality prediction. We present results from GRU-D in the supplementary material because GRU-D results reveal similar trends and usually underperform logistic regression or DAN in our experiments.

Notes outperform structured variables in PR-AUC and ROC-AUC in readmission prediction (Fig. 2a-2d). For both logistic regression and DAN, notes are more predictive than structured information in readmission prediction based on PR-AUC and ROC-AUC. In fact, in most cases, structured variables provide little additional predictive power over notes (except PR-AUC with DAN). Interestingly, we observe mixed results for precision-based metrics. Structured information can outperform notes in identifying the patients that are most likely to be readmitted. For DAN, combining notes and structured information provides a significant boost in precision at 1% compared to one type of information alone, with an improvement of 16% and 13% in absolute precision over notes and structured variables respectively.

Structured information dominates notes in mortality prediction (Fig. 2e-2h). We observe marginally additional predictive value in mortality prediction by incorporating notes with structured information. In our experiments, the improvement is negligible across all metrics. This result differs from Ghassemi et al. (2014). We believe that the reason is that Ghassemi et al. (2014) only consider age, gender, and the SAPS II score in structured information, while our work considers substantially more structured variables. It is worth noting that logistic regression with our complete set of structured variables provides better performance than DAN and the absolute number in ROC (0.892) is better than the best number (0.79) in prior work (Che et al., 2018). The reason for the




Figure 2: Results of PR-AUC/ROC-AUC/Precision at 1%/Precision at 5% on logistic regression (LR)/deep averaging networks (DAN) models in readmission prediction and mortality prediction (24 hours). Notes are valuable for readmission predictions, but are marginally valuable in mortality prediction.

limited value of notes might be that mortality prediction is a relatively simple task where structured information provides unambiguous signals.

In sum, we find that note contributes valuable information over structured variables in readmission prediction, but almost no additional value in mortality prediction. Note that ROC-AUC tends to be insensitive to different models and information. We thus use PR-AUC in the rest of the work to discuss the value of selected information.

4 Finding Needles in a Haystack: Probing for the Valuable Information


The key goal of this work is to identify valuable components within notes, as we hypothesize that not all information in notes is valuable for medical decisions, as measured by the predictive power.

To identify valuable components, we leverage an existing machine learning model (e.g., models in Fig. 2) and hypothesize that the test performance is better if we only use the “valuable” components. Formally, assume that we trained a model using all notes, f_{all} . S_i denotes sentences in all notes ($E_i^{T=1} = \{(h_t, \tau_t, x_t) | \tau_t = 1\}_{t=1}^T$) for an admission in the test set. We would like to find a subset of sentences $s_i \subset S_i$ so that $f_{\text{all}}(s_i)$ provides more accurate predictions than $f_{\text{all}}(S_i)$. Note that s_i by definition entails a distribution shift from the data that f_{all} is trained on (S_i), because s_i is much shorter than S_i .

The challenge lies in developing interpretable ways to identify valuable content. We first compare the value of different types of notes in §4.1, which can be seen as trivial value functions based on type of note, and then propose interpretable value functions to zoom in on the content of notes (§4.2). Finally, we show that these valuable components not only provide accurate predictions with a model trained with all the notes, but also allow us to learn a model with better predictive power than that trained with all the notes (§4.3). In other words, we can effectively remove the noise by focusing on the valuable components.

4.1 Discharge Summaries, Nursing Notes, and Physician Notes are Valuable

To answer our first question, we compare the effectiveness of different types of notes within the top five most common categories: nursing, radiology, ECG, physician, and discharge summary. An important challenge lies in the fact that not every admission produces all types of notes. Therefore, we conduct pairwise comparison that ensures an admission has both types of note. Specifically, for each pair of note types (t_1, t_2), we choose admissions with both two types of note and make predictions using s_{t_1} and s_{t_2} respectively, where s_t refers to all the sentences in notes of type t . Each cell in Fig. 3 indicates $\text{PR-AUC}(f_{\text{all}}(s_{t_{\text{row}}}), y) - \text{PR-AUC}(f_{\text{all}}(s_{t_{\text{column}}}), y)$



(a) Readmission prediction.

(b) Mortality prediction (24 hrs).

Figure 3: Pairwise comparisons between different types of note with logistic regression (each cell shows $\text{PR-AUC}(f_{\text{all}}(s_{t_{\text{row}}}), y) - \text{PR-AUC}(f_{\text{all}}(s_{t_{\text{column}}}), y)$). To account for the differences in length, we subsample two types of note under comparison to be the same length and report the average values of 10 samples. Discharge summaries dominate all other types of notes in readmission prediction, while nursing notes are most useful for mortality prediction.

with LR (see the supplementary material for DAN results, which are similar to LR). For instance, the top right cell in Fig. 3a shows the performance difference between using only nursing notes and using only discharge summaries for admissions with both nursing notes and discharge summaries. The negative value suggests that nursing notes provide less accurate predictions (hence less valuable information) than discharge summaries in readmission prediction. Note that due to significant variance in length between types of note, we subsample $s_{t_{\text{row}}}$ and $s_{t_{\text{column}}}$ to be the same length in these experiments.

Discharge summaries dominate other types of notes in readmission prediction. Visually, most of the dark values in Fig. 3a are associated with discharge summaries. This makes sense because discharge summaries provide a holistic view of the entire admission and are likely most helpful for predicting future readmission. Among the other four types of notes, nursing notes are the second most valuable. In comparison, physician notes, radiology reports, and ECG reports are less valuable.

Nursing notes and physician notes are more valuable for mortality prediction. For mortality prediction, nursing notes provide the best predictive power. ECG reports always have the worst results. Recall that we subsample each type of notes to the same length. Hence, the lack of value in ECG reports cannot be attributed to its short length.

In summary, templated notes such as radiology reports and ECG reports are less valuable for predictive tasks in medical decisions. While physician notes are the central subject in prior work (Weir and Nebeker, 2007), nursing notes are as important for medical purposes given that there are

many more nursing notes and they record patient information frequently.

4.2 Identifying Valuable Chunks of Notes

Next, we zoom into sentences within notes to find out which sentences are more valuable, i.e., providing better predictive power using the model trained with all notes. We choose content from discharge summaries for readmission prediction because they are the most valuable. For mortality prediction, we select content from the last physician note since they play a similar role as discharge summaries. To select valuable sentences from S_i , we propose various value functions V , and for each V , we choose the sentences in S_i that score the highest using V to construct $s_i^V \subset S_i$. These value functions are our main subject of interest. We consider the following value functions.

- **Longest sentences.** Intuitively, longer sentences may contain valuable information. Hence, we use $V_{\text{longest}}(s) = \text{length}(s)$, where length gives the number of tokens.
- **Sentences with highest fractions of medical terms.** Medical terms are critical for communicating medical information. We develop a value function based on the fraction of medical terms in a sentence. Empirically, we observe that fraction alone tends to choose very short sentences, we thus use $V_{\text{frac}}(s) = \frac{\text{medical}(s)}{\text{length}(s)} * \sqrt{\text{length}(s)}$, where the medical terms come from OpenMed-Spel (Robinson, 2014) and MTH-Med-Spel-Chek (Narayanaswamy, 2014)⁷.
- **Similarity with previous notes.** A significant complaint about notes is the prevalence of copy-pasting. We thus develop a value function based

⁷<https://github.com/Glutanimate>.




Figure 4: Performance of the selected information based on different value functions using the logistic regression (LR) model trained on all notes. Despite the distribution shift (selected content is much shorter than the training data, i.e., all notes), the selected information outperforms using all notes with either LR or DAN.




Figure 5: Performance comparison for discharge summaries of different lengths in readmission prediction. Selecting valuable information is most useful for the fourth quartile, the longest discharge summaries.

on similarity with previous notes. As discharge summaries are the final note within an admission, we compute the max tf-idf similarity of a sentence with all previous notes.⁸ Specifically, we define $V_{\text{dissimilar}}(s) = -\max_{x_k \in X} \text{cossim}(s, x_k)$, where X refers to all previous notes: we find the most similar previous note to the sentence of interest and flip the sign to estimate dissimilarity.

Although we hypothesize dissimilar sentences are more valuable due to copy-pasting concerns (i.e., novelty), sentences may also be repeatedly emphasized in notes because they convey critical information. We thus also flip $V_{\text{dissimilar}}$ to choose the most similar sentences (V_{similar}) and use V_{mix} to select half of the most similar and half of the most dissimilar ones. Similarly, we apply these value functions on the last physician note to select valuable content for mortality prediction.

- **Important Section.** Finally, physicians do not treat every section in notes equally themselves, and spend more time on reading the “Impression and Plan” section than other sections (Brown et al., 2014). We use whether a sentence is in this section as our final value function. This only applies to physician notes.

In practice, sentences in medical notes can be very long. To be fair across different value func-

⁸We also consider average and normalizations to account for the effect of length. We tried to use previous sentences as basic units as well. The results are similar and see the supplementary material for details.

tions, we truncate the selected sentences to use the same number of tokens with each value function (see the implementation details in the supplementary material).

Parts of notes can outperform the whole. Fig. 4 shows the test performance of using different value functions to select a fixed percentage of tokens in the discharge summary or the last physician note, compared to using all notes. The underlying model is the corresponding logistic regression model. We also show the performance of using all notes with DAN as a benchmark.

Some value functions are able to select valuable information that outperforms using all notes with either logistic regression or DAN. Interestingly, we find that selected valuable information generally performs better based on the LR model, which seems more robust to distribution shifts than DAN (recall that selected valuable information is much shorter than the expected test set using all notes).

In readmission prediction, medical terms are fairly effective early on, outperforming using all notes with LR, using only 20% of the discharge summary. As we include more tokens, a mix of similar and dissimilar sentences becomes more valuable and is eventually comparable with DAN using 45% of the discharge summary. Table 2 presents an example of sentences selected from different value functions in readmission prediction using logistic regression.

In mortality prediction, the advantage of selected valuable information is even more salient.

Value Function	Prob.	Selected Sentences
similar	0.189	Congestive Heart Failure - Systolic and [**Month/Day/Year**] Failure - most recent echo on [**2123-9-3**] with EF 40 % 3. Valvular Disease - Moderate Aortic Stenosis - mild - moderate aortic regurgitation - mild - moderate mitral regurgitation 4. Coronary Artery Disease - [**2122-11-16**] - s/p BMS to OM2, D1, Left circumflex in [**2122-11-16**] for unstable angina and TWI in V2 - 4 - [**2123-5-24**] - NSTEMI s/p cardiac cath
dissimilar	0.291	No thrills, lifts. BNP elevated though decreased from prior. Please take tiotropium bromide (Spiriva) inhalation twice a day 2. Mom[**Name(NI) 6474**] 50 mcg / Actuation Spray , Non - Aerosol Sig : Two (2) spray Nasal twice a day. Troponins negative x3 sets. PO Q24H (every 24 hours). No S3 or S4. Resp were unlabored, no accessory muscle use. Occupation: general surgeon in [**Location (un) 4551**]. Abd: Soft, NTND. EtOH: 1 glass of wine or alcoholic drink /week. + [**4-16**] word dyspnea
mix-similar	0.620	Congestive Heart Failure - Systolic and [**Month/Day/Year**] Failure - most recent echo on [**2123-9-3**] with EF 40% 3. Valvular Disease - Moderate Aortic Stenosis - mild - moderate aortic regurgitation - mild - moderate mitral regurgitation 4. Coronary Artery Disease <u>No thrills, lifts. BNP elevated though decreased from prior. Please take tiotropium bromide (Spiriva) inhalation twice a day 2. Mom[**Name (NI) 6474**] 50 mcg / Actuation Spray, Non-Aerosol Sig: Two (2) spray Nasal twice a day. Troponins negative x3 sets. PO Q24H</u>

Table 2: Example of selected sentences (5% of tokens) by different value functions from discharge summaries for readmission prediction. This patient was **readmitted** to the hospital in 30 days after discharge. Underlined sentences in *mix-similar* function come from dissimilar sentences. “Prob.” shows the output probability of readmission with the LR model trained on all notes given selected sentences.

Consistent with Brown et al. (2014), “assessment and plan” is indeed more valuable than the whole note. It alone outperforms both LR and DAN with all notes. Different from readmission prediction, sentences dissimilar to previous notes are most effective. The reason might be that dissimilar sentences give novel developments in the patient that relate to the impending death. As structured information dominates notes in this task, selected information adds little value to structured information (see the supplementary material).

The effectiveness of value functions varies across lengths. To further understand the effectiveness of value functions, we break down Fig. 4a based on the length of discharge summaries. Intuitively, it would be harder to select valuable information for short summaries, and Fig. 5a confirms this hypothesis. In all the other quartiles, a value function is able to select sentences that outperform both LR and DAN using all notes. The medical terms are most effective in the second and third quartiles. In the fourth quartile (i.e., the longest discharge summaries), dissimilar content is very helpful, which likely includes novel perspectives synthesized in discharge summaries. These observations resonate with our earlier discussion that dissimilar content contribute novel information.

4.3 Leveraging Valuable Information

Building on the above observations, we leverage the selected valuable information to train models

based on only valuable information. Fig. 6 shows the performance of these models on readmission prediction.⁹ Here we include DAN with note-level attention (“DAN-Att”) as a model-driven oracle weighted selection approach, although it does not lead to interpretable value functions that can inform caregivers during note-taking.

First, models trained only using discharge summaries (“last note”) improves the performance over using all notes by 41% (0.219 vs. 0.155), and outperform DAN and DAN-att as well. Using medical terms and all types of similarity methods, we can outperform using all notes with models only trained on 20% tokens of discharge summaries, that is, 6.8% of all notes. Compared to Fig. 4a, by focusing exclusively on these selected 20% of tokens, the model trained with selected dissimilar sentences outperforms logistic regression by 24.3% (0.194 vs. 0.156), DAN by 8.2% (0.194 vs. 0.178), and DAN-Att by 2% (0.194 vs. 0.190). We also experiment with ClinicalBERT with a fixed number of tokens (see the supplementary material). ClinicalBERT provides comparable performance with logistic regression, and demonstrates similar qualitative trends.

Recall that medical notes dominate structured information for readmission prediction. It follows that our best performance with selected valuable information all outperform the best performance

⁹This is hard to operationalize for mortality prediction since not all admissions have physician notes.




Figure 6: Performance of trained models with selected valuable information (20% of discharge summaries).

obtained in §3.

5 Related Work

We summarize additional related work into the following three areas.

Value of medical notes. Prior work shows that some important phenotypic characteristics can only be inferred from text reports (Shivade et al., 2014). For example, Escudié et al. (2017) observed that 92.5% of information regarding autoimmune thyroiditis is only presented in text. Despite the potential valuable information in medical notes, prior work also points out the redundancy in EHRs. Cohen et al. (2013) proposed methods to reduce redundant content for the same patient with a summarization-like fingerprinting algorithm, and show improvements in topic modeling. We also discuss the problem of redundancy in notes, but provide a different perspective by probing what type of information is more valuable than others using our framework.

NLP for medical notes. The NLP community has worked extensively on medical notes to alleviate information overload, ranging from summarization (McInerney et al., 2020; Liang et al., 2019; Alsentzer and Kim, 2018) to information extraction (Wiegrefe et al., 2019; Zheng et al., 2014; Wang et al., 2018). For instance, information extraction aims to automatically extract valuable information from existing medical notes. While our operationalization seems similar, our ultimate goal is to facilitate information solicitation so that medical notes contain more valuable information.

Recently, generating medical notes has attracted substantial interest that might help caregivers record information (Liu et al., 2018; Krishna et al., 2020), although they do not take into account the value of information.

Predictive tasks with EHRs. Readmission prediction and mortality prediction are important tasks that have been examined in a battery of stud-

ies (Johnson et al., 2017; Ghassemi et al., 2014; Purushotham et al., 2018; Rajkomar et al., 2018). In MIMIC-III, to the best of our knowledge, we have experimented with the most extensive structured variables and as a result, achieved better performance even with simple models. Other critical tasks include predicting diagnosis codes (Ford et al., 2016) and length of stay (Rajkomar et al., 2018). We expect information in medical notes to be valued differently in these tasks as well.

6 Conclusion

Our results confirm the value of medical notes, especially for readmission prediction. We further demonstrate that parts can outperform the whole. For instance, selected sentences from discharge summaries can better predict future readmission than using all notes and structured variables. Our work can be viewed as the reverse direction of adversarial NLP (Wallace et al., 2019): instead of generating triggers that fool NLP models, we identify valuable information in texts towards enabling humans to generate valuable texts.

Beyond confirming intuitions that “assessment and plan” in physician notes is valuable, our work highlights the importance of nursing notes. Our results also suggest that a possible strategy to improve the value of medical notes is to help caregivers efficiently provide novel content while highlighting important prior information (mixed similarity). Substantial future work is required to achieve the long-term goal of improving the note-taking process by nudging caregivers towards obtaining and recording valuable information.

In general, the issue of effective information solicitation has been understudied by the NLP community. In addition to model advances, we need to develop human-centered approaches to collect data of better quality from people. As Hartzband et al. (2008) argued, “as medicine incorporates new technology, its focus should remain on interaction between the sick and healer.” We hope that our study will encourage studies to understand the interaction process and the note-taking process, beyond understanding the resulting information as a given. After all, people are at the center of data.

Acknowledgments. We thank helpful comments from anonymous reviewers. We thank the MIMIC team for continually providing invaluable datasets for the research community.

References

- Emily Alsentzer and Anne Kim. 2018. Extractive summarization of ehr discharge notes. *arXiv preprint arXiv:1810.12085*.
- Emily Alsentzer, John Murphy, William Boag, Wei-Hung Weng, Di Jin, Tristan Naumann, and Matthew McDermott. 2019. [Publicly available clinical BERT embeddings](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 72–78, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- PJ Brown, JL Marquard, B Amster, M Romoser, J FridERICI, S Goff, and D Fisher. 2014. What do physicians read (and ignore) in electronic progress notes? *Applied clinical informatics*, 5(02):430–444.
- Zhengping Che, Sanjay Purushotham, Kyunghyun Cho, David Sontag, and Yan Liu. 2018. Recurrent neural networks for multivariate time series with missing values. *Scientific reports*, 8(1):6085.
- Raphael Cohen, Michael Elhadad, and Noémie Elhadad. 2013. Redundancy in electronic health record corpora: analysis, impact on text mining performance and mitigation strategies. *BMC bioinformatics*, 14(1):10.
- Jesse Davis and Mark Goadrich. 2006. The relationship between precision-recall and roc curves. In *Proceedings of the 23rd international conference on Machine learning*, pages 233–240.
- Jean-Baptiste Escudié, Bastien Rance, Georgia Malamut, Sherine Khater, Anita Burgun, Christophe Cellier, and Anne-Sophie Jannot. 2017. A novel data-driven workflow combining literature and electronic health records to estimate comorbidities burden for a specific disease: a case study on autoimmune comorbidities in patients with celiac disease. *BMC medical informatics and decision making*, 17(1):140.
- Elizabeth Ford, John A Carroll, Helen E Smith, Donia Scott, and Jackie A Cassell. 2016. Extracting information from the text of electronic medical records to improve case detection: a systematic review. *Journal of the American Medical Informatics Association*, 23(5):1007–1015.
- Atul Gawande. 2018. [Why doctors hate their computers](#).
- Marzyeh Ghassemi, Tristan Naumann, Finale Doshi-Velez, Nicole Brimmer, Rohit Joshi, Anna Rumshisky, and Peter Szolovits. 2014. Unfolding physiological state: Mortality modelling in intensive care units. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 75–84. ACM.
- Pamela Hartzband, Jerome Groopman, et al. 2008. Off the record—avoiding the pitfalls of going electronic. *New England Journal of Medicine*, 358(16):1656–1657.
- Hrayr Harutyunyan, Hrant Khachatrian, David C Kale, and Aram Galstyan. 2017. Multitask learning and benchmarking with clinical time series data. *arXiv preprint arXiv:1703.07771*.
- Hrayr Harutyunyan, Hrant Khachatrian, David C. Kale, Greg Ver Steeg, and Aram Galstyan. 2019. [Multitask learning and benchmarking with clinical time series data](#). *Scientific Data*, 6(1).
- Mohit Iyyer, Varun Manjunatha, Jordan Boyd-Graber, and Hal Daumé III. 2015. Deep unordered composition rivals syntactic methods for text classification. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1681–1691.
- Alistair EW Johnson, Tom J Pollard, and Roger G Mark. 2017. Reproducibility in critical care: a mortality prediction case study. In *Machine Learning for Healthcare Conference*, pages 361–376.
- Kundan Krishna, Sopan Khosla, Jeffrey P. Bigham, and Zachary C. Lipton. 2020. [Generating soap notes from doctor-patient conversations](#).
- Harlan M Krumholz. 2014. Big data and new knowledge in medicine: the thinking, training, and tools needed for a learning health system. *Health Affairs*, 33(7):1163–1170.
- Jennifer Liang, Ching-Huei Tsou, and Ananya Poddar. 2019. [A novel system for extractive clinical note summarization using EHR data](#). In *Proceedings of the 2nd Clinical Natural Language Processing Workshop*, pages 46–54, Minneapolis, Minnesota, USA. Association for Computational Linguistics.
- Jingshu Liu, Zachariah Zhang, and Narges Razavian. 2018. Deep ehr: Chronic disease prediction using medical notes. *arXiv preprint arXiv:1808.04928*.
- Edward Loper and Steven Bird. 2002. Nltk: The natural language toolkit. In *In Proceedings of the ACL Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics*. Philadelphia: Association for Computational Linguistics.
- Denis Jered McInerney, Borna Dabiri, Anne-Sophie Touret, Geoffrey Young, Jan-Willem van de Meent, and Byron C Wallace. 2020. Query-focused ehr summarization to aid imaging diagnosis. *arXiv preprint arXiv:2004.04645*.
- Rajasekharan Narayanaswamy. 2014. [Mth-med-spelchek of mt-herald](#).
- Rikinkumar S Patel, Ramya Bachu, Archana Adikey, Meryem Malik, and Mansi Shah. 2018. Factors related to physician burnout and its consequences: a review. *Behavioral Sciences*, 8(11):98.

- Thomas H Payne, Sarah Corley, Theresa A Cullen, Tejal K Gandhi, Linda Harrington, Gilad J Kuperman, John E Mattison, David P McCallie, Clement J McDonald, Paul C Tang, et al. 2015. Report of the amia ehr-2020 task force on the status and future direction of ehrs. *Journal of the American Medical Informatics Association*, 22(5):1102–1110.
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830.
- Sanjay Purushotham, Chuizheng Meng, Zhengping Che, and Yan Liu. 2018. Benchmarking deep learning models on large healthcare datasets. *Journal of biomedical informatics*, 83:112–134.
- Alvin Rajkomar, Eyal Oren, Kai Chen, Andrew M Dai, Nissan Hajaj, Michaela Hardt, Peter J Liu, Xiaobing Liu, Jake Marcus, Mimi Sun, et al. 2018. Scalable and accurate deep learning with electronic health records. *NPJ Digital Medicine*, 1(1):18.
- R. Robinson. 2014. [Openmedspel of e-medtools \(version 2.0.0\)](#).
- Chaitanya Shivade, Preethi Raghavan, Eric Fosler-Lussier, Peter J Embi, Noemie Elhadad, Stephen B Johnson, and Albert M Lai. 2014. A review of approaches to identifying patient phenotype cohorts using electronic health records. *Journal of the American Medical Informatics Association*, 21(2):221–230.
- Edward H Shortliffe. 2010. Biomedical informatics in the education of physicians. *JAMA*, 304(11):1227–1228.
- William W Stead, John R Searle, Henry E Fessler, Jack W Smith, and Edward H Shortliffe. 2011. Biomedical informatics: changing what physicians need to know and how they learn. *Academic Medicine*, 86(4):429–434.
- Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2019. Universal adversarial triggers for nlp. In *EMNLP*.
- Yanshan Wang, Liwei Wang, Majid Rastegar-Mojarad, Sungrim Moon, Feichen Shen, Naveed Afzal, Sijia Liu, Yuqun Zeng, Saeed Mehrabi, Sunghwan Sohn, et al. 2018. Clinical information extraction applications: a literature review. *Journal of biomedical informatics*, 77:34–49.
- Charlene R Weir and Jonathan R Nebeker. 2007. Critical issues in an electronic documentation system. In *AMIA Annual Symposium Proceedings*, volume 2007, page 786. American Medical Informatics Association.
- Sarah Wiegrefe, Edward Choi, Sherry Yan, Jimeng Sun, and Jacob Eisenstein. 2019. [Clinical concept extraction for document-level coding](#). In *Proceedings of the 18th BioNLP Workshop and Shared Task*, pages 261–272, Florence, Italy. Association for Computational Linguistics.
- M Zeng. 2016. Opinion: When the doctor must choose between her patients and her notes. *Accessed January*, 19.
- Jin Guang Zheng, Daniel Howsmon, Boliang Zhang, Juergen Hahn, Deborah McGuinness, James Hendler, and Heng Ji. 2014. Entity linking for biomedical literature. In *Proceedings of the ACM 8th International Workshop on Data and Text Mining in Bioinformatics*, pages 3–4.