

Aligning AI Research with the Needs of Clinical Coding Workflows: Eight Recommendations Based on US Data Analysis and Critical Review

Yidong Gan^{1,2}, Maciej Rybinski^{2,3}, Ben Hachey^{1,4}, Jonathan K. Kummerfeld¹

¹The University of Sydney ²CSIRO Data61 ³The University of Málaga ⁴Beamtree
yidong.gan@sydney.edu.au

Abstract

Clinical coding is crucial for healthcare billing and data analysis. Manual clinical coding is labour-intensive and error-prone, which has motivated research towards full automation of the process. However, our analysis, based on US English electronic health records and automated coding research using these records, shows that widely used evaluation methods are not aligned with real clinical contexts. For example, evaluations that focus on the top 50 most common codes are an oversimplification, as there are thousands of codes used in practice. This position paper aims to align AI coding research more closely with practical challenges of clinical coding. Based on our analysis, we offer eight specific recommendations, suggesting ways to improve current evaluation methods. Additionally, we propose new AI-based methods beyond automated coding, suggesting alternative approaches to assist clinical coders in their workflows.

1 Introduction

Clinical coding is a process that transforms clinical notes into a set of alphanumeric codes, which represent diagnoses and procedures during medical visits. This process is essential to tasks like hospital billing and disease prevalence studies. Manual clinical coding is labour-intensive and error-prone (Karimi et al., 2017; Li and Yu, 2020). To address these issues, automated coding has been widely explored. While many existing studies frame it as a multi-label classification task (Mullenbach et al., 2018; Vu et al., 2020; Huang et al., 2022; Li et al., 2023), no analysis has examined whether that matches the needs of clinicians.

This position paper critically reviews automated coding studies based on multi-label classification, focusing on those using the public US Medical Information Mart for Intensive Care (MIMIC) datasets (Goldberger et al., 2000; Johnson et al.,

2016, 2023). For consistency, we use the clinical term ‘code’ instead of ‘label’ throughout this paper. We focus on the MIMIC datasets as they are far larger than other public datasets. Specifically, MIMIC provides English electronic health records of acute and emergent inpatients, with the latest version, MIMIC-IV, including 331,675 patient admissions. The few other public datasets, such as those released by Pestian et al. (2007) and Miranda-Escalada et al. (2020), contain 978 and 1,000 admissions, respectively. This significant scale difference allows MIMIC to cover a broader range of diagnoses and procedures, making it more suitable for comprehensive analyses.

We show that current evaluations do not align with the needs of clinical contexts. In practice, coders must select from over a thousand codes and sequence them correctly (CMS and NCHS, 2024), yet many studies use smaller code sets and overlook code sequencing in evaluation. Additionally, the Area Under the Receiver Operating Characteristic Curve (AUC-ROC) is often the only threshold-independent metric reported, which is inappropriate given the imbalanced code distribution. Common human coding metrics like Exact Match Ratio and Jaccard Score are often omitted, making it difficult to measure the accuracy gap between automated and human coding. In other words, widely used evaluation strategies do not measure what truly matters for clinical application.

We then propose new methodologies to support clinical coders rather than automating the entire process. Given the limited effectiveness of automated coding (Edin et al., 2023), manual coding with software assistance remains prevalent, with code auditing essential in clinical workflows to reduce errors. Despite its significance, research on AI-assisted coding and auditing is limited. This underscores the need for future studies in these areas, which have the potential to yield practical solutions sooner than fully automated coding.

2 Clinical Coding Workflow

In this section, we present a typical clinical coding workflow in an inpatient setting, outline the individual coding processes, and provide an overview of the current state-of-the-art strategies and tools used in these processes. The workflow presented is specific to the US context; other countries may have similar or different processes.

The upper diagram in Figure 1 outlines a high-level, four-step clinical coding workflow for an inpatient episode. First, a patient is admitted to the hospital. During care, all relevant documentation (e.g., pathology reports) is entered into their Electronic Health Record (EHR). Upon discharge, the attending doctor writes a summary detailing the patient’s stay, including diagnosis and treatment. Clinical coders then assign International Classification of Diseases (ICD) codes based on the EHR. For reimbursement purposes, relevant ICD codes are grouped into a Diagnosis Related Group (DRG) code. DRG is a classification system that organises hospital cases into groups; each DRG has a specific payment rate based on the average resources needed to treat patients in that group. The lower diagram in Figure 1 gives additional details of the ICD coding task, illustrating the processes before and after coding episodes with ICD.

The implementation of ICD and DRG varies across countries, and the coding inputs can differ depending on healthcare settings (e.g., hospital policies). Previous AI coding studies often use discharge summaries as inputs but have also explored other sources, such as radiology reports (Pestian et al., 2007) and multilingual death certificates (Névéol et al., 2018). Given all these differences, it is important that AI coding research carefully considers the clinical coding workflow it is addressing. Of note, many references in the following subsection are corporate products, as previous research has largely focused on the automated coding approach in the clinical coding workflow.

2.1 Task Allocation

With many hospitals facing backlogs of uncoded cases (Alonso et al., 2020), optimising task allocation is crucial. The order in which cases are processed depends on hospital-specific business metrics, such as maximising profit and reducing backlog. Tools like Beamtree’s Q Coding platform (Beamtree, 2024) support task forecasting, rule-based distribution, and scheduling, enabling

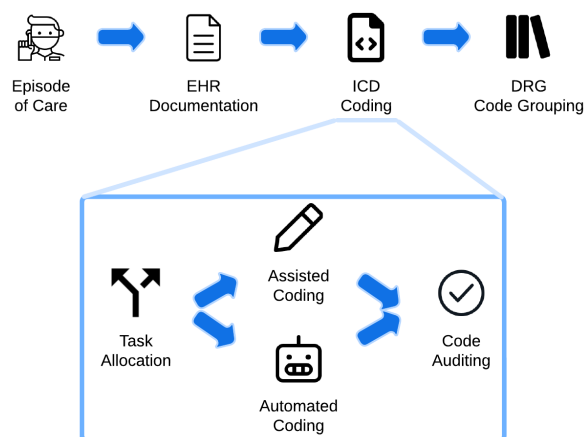


Figure 1: Typical clinical coding workflow for an inpatient episode of care in the US.

hospitals to address backlogs systematically and ensure critical cases are handled promptly.

After determining case priority, the next step is to allocate cases to the appropriate expert. In assisted coding (i.e., manual coding with computer software support), assigning cases based on coder specialty or experience effectively manages the complexity of coding tasks (Alonso et al., 2020; Beamtree, 2024). Likewise, with AI-based coding solutions, cases could be assigned to either assisted or automated coding pathways. This step is speculative and will be discussed further in Section 3.

2.2 Assisted Coding

In assisted coding, where coders use computer software to manually enter codes, each keystroke, word read, and thought incurs a cost. By minimising the manual effort involved, these tasks become more efficient. This promise of higher efficiency has motivated health tech companies to develop various assisted coding tools.

Existing tools include features such as searching and navigating codes with integrated guidelines, which help users quickly find relevant codes and follow best practices (Beamtree, 2024; 3M, 2024d). AI-suggested codes and DRG grouping helps in streamlining the coding process (Beamtree, 2024; 3M, 2024b,d). Customisable rules allow users to automate coding based on criteria drawn from their expertise (3M, 2024b). Online audits against organisational policies ensure quality control (3M, 2024a). Evidence linking for assigned codes facilitates efficient edits and reviews (Goins, 2024; 3M, 2024c). Collectively, these features contribute to a more efficient and cost-effective coding workflow.

2.3 Automated Coding

Automated coding refers to the process of assigning accurate diagnostic and procedural ICD codes without human intervention. It is closely related to other real-world problems, such as tagging in social networks (Coope et al., 2019) and indexing biomedical literature (Krithara et al., 2023), where a piece of text is categorised using multiple labels. Automated coding can address coding of episodes where the accuracy of an AI system is higher than human performance or is otherwise sufficient given operational cost-benefit considerations. CodeAssist (3M, 2024c), for example, is a commercial system used in many hospitals, with its main feature being automated coding.

Automated coding has also been the focus of AI coding research to date, with most studies relying solely on discharge summaries as model inputs. Mullenbach et al. (2018) propose the Convolutional Attention for Multi-Label Classification (CAML) model. CAML combines convolutional neural networks (CNNs) with per-code attention to focus on text sections relevant to each ICD code. This attention mechanism enhances interpretability by highlighting the text that contributes to the model's decisions. Li and Yu (2020) extend this framework with the Multi-Filter Residual Convolutional Neural Network (MultiResCNN), which uses multiple filter layers and residual connections for better feature extraction. On the other hand, Vu et al. (2020) introduce the Label Attention Model (LAAT), using a bidirectional long short-term memory (LSTM) to capture clinical context in text. LAAT uses a distinct per-code attention mechanism compared to CAML, incorporating additional steps in attention weight calculations and a hierarchical joint learning strategy, which leads to better predictions for rare ICD codes. Huang et al. (2022) adapt LAAT's attention mechanism for pretrained language models in ICD coding (PLM-ICD). PLM-ICD also incorporates domain-specific pretraining and segment pooling, addressing challenges such as the domain mismatch between pretraining and clinical text and the large code space.

Recent benchmarking by Edin et al. (2023) compares CAML, MultiResCNN, LAAT, PLM-ICD, and a few other models. PLM-ICD consistently achieves best results on MIMIC-III and MIMIC-IV datasets. However, all models, including PLM-ICD, struggle with rare codes, which is a persistent issue in automated coding. Notably, document

length had minimal effect on model performance, with little difference when truncating documents from 4,000 to 2,500 words. In Section 3, we demonstrate that prior studies fail to measure key factors critical for real-world applications. Our analysis reveals realistic upper bounds of existing state-of-the-art models compared to human performance.

2.4 Code Auditing

Clinical coding is complex, and even human coders often make mistakes (Burns et al., 2012). In the US, coding errors and quality improvement efforts cost an estimated \$25 billion annually (Xie and Xing, 2018). Even more concerning, some errors may be treated as fraud, leading to legal liability (Rudman et al., 2009). To address this, many asynchronous auditing tools have been developed. For example, 3M (2024a) provides offline tools for batch auditing, referencing documents and codes used in patient claims to ensure compliant coding, and integrated denial tracking for managing coding quality. Beamtree (2024) offers tools for audits against quality indicators, dynamic code sequencing, and code combination validation. In general, these tools aim to enhance coding quality and reduce error-related costs in both time and revenue.

3 Data and Automation: Analysis and Recommendations

In this section, we (1) identify key shortcomings in the evaluation methodologies widely used in automated coding studies, (2) analyse the widely used MIMIC datasets, and (3) offer corresponding recommendations. Many of these shortcomings are not limited to studies using the MIMIC datasets, but are also prevalent in other methodologies proposed for automated coding shared tasks involving different datasets (Pestian et al., 2007; Miranda-Escalada et al., 2020).

3.1 Evaluations Using the Top 50 Codes Do Not Reflect Real Effectiveness

Table 1 shows that many of the existing studies evaluate their methods using the 50 most frequent codes¹. When considering the application of models in a real-world healthcare environment, this evaluation strategy is sub-optimal (Liu et al., 2021); it focuses solely on the most frequent codes, failing

¹This is not intended as a systematic literature review. Instead, the aim is to demonstrate that many studies rely on the top 50 codes to validate their methodologies.

	Count	References
Reported only top 50 results	8	(Shi et al., 2017; Wang et al., 2018; Teng et al., 2020; Feucht et al., 2021; Sun et al., 2021; Liu et al., 2022; Michalopoulos et al., 2022; Liu et al., 2023)
Reported only full results	2	(Wiegrefe et al., 2019; Kim and Ganapathi, 2021)
Reported one or more top k results, excluding full results	3	(Prakash et al., 2017; Xie and Xing, 2018; Zhang et al., 2020)
Reported both top 50 and full results, but used top 50 results to emphasise contribution	4	(Li and Yu, 2020; Li et al., 2021; Shi et al., 2021; Yang et al., 2022)
Reported both top 50 and full results without emphasising top 50 results	8	(Mullenbach et al., 2018; Xie et al., 2019; Vu et al., 2020; Cao et al., 2020; Liu et al., 2021; Zhou et al., 2021; Yuan et al., 2022; Li et al., 2023)

Table 1: Overview of Automated Coding Studies and Their Evaluation Strategy.

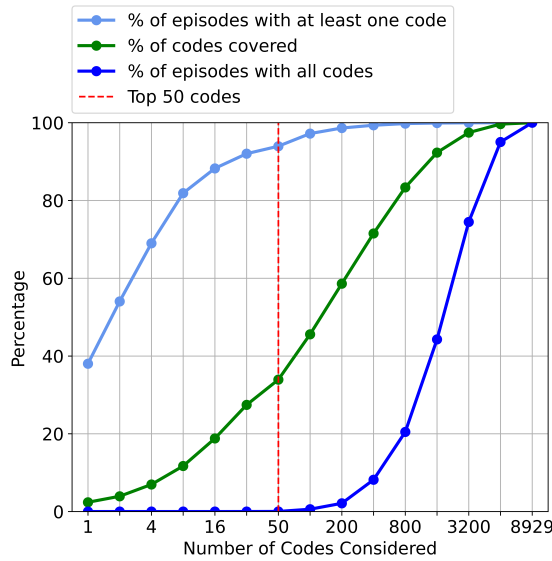


Figure 2: Code coverage in MIMIC-III.

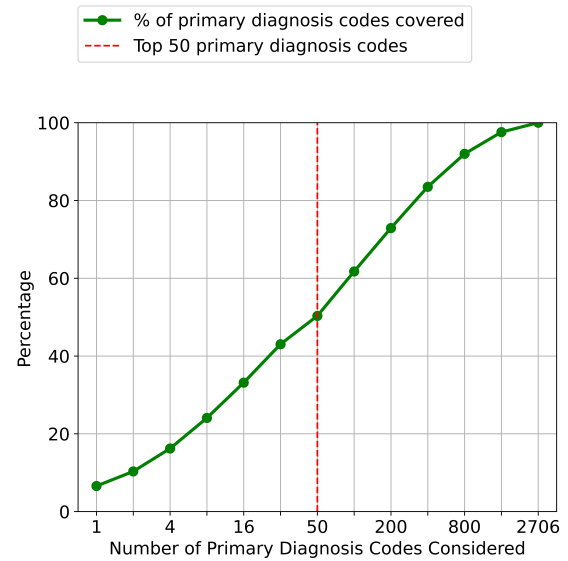


Figure 3: Code coverage for the first primary diagnosis code in an episode only.

to cover all diagnoses and procedures encountered in clinical practice.

We demonstrate the coverage issue by showing the distribution of codes in MIMIC-III. The green line in Figure 2 shows that the top 50 codes comprise only 33.92% of total code occurrences. We further measure how well the episodes² in MIMIC-III are covered by the top 50 codes. Specifically, we calculate the percentage of episodes where all assigned codes are within the top 50. The dark blue line shows that none of the episodes are fully covered, meaning every episode has some codes outside the top 50. Even when we increase this to include the top 800 codes, the coverage rate remains very low, at 20.48%. In other words, even with the top 800 codes, we miss some coding information in about 80% of episodes. If we use a more generous measure, calculating the percent-

²A single episode usually has many codes.

age of episodes that include at least one of the top 50 codes, the light blue line shows a coverage rate of 93.95%. If we preprocess the MIMIC-III episodes to retain only the first primary diagnosis code, which typically has the highest clinical and billing priority, the issue of low coverage rates persists. We observe an 8% higher coverage rate with the top 50 primary diagnosis codes (see Figure 3), yet the improvement is still marginal in terms of overall coverage.

The low coverage rate leads to a generalisation issue, where the ranking of different models shows low correlation between the top 50 and full code settings. According to Mullenbach et al. (2018), in the top 50 setting, CNN outperforms their proposed CAML model in terms of Micro F1, Macro F1, and Precision@5. However, CAML performs better in the full code setting. This contradiction is also observed in a reproducibility study (Edin

et al., 2023). If we rank the six reproduced models by macro AUC-ROC scores from lowest to highest, the order in the top 50 setting is Bi-GRU, CAML, CNN, MultiResCNN, LAAT, PLM-ICD, whereas in the full code setting, it is CNN, Bi-GRU, CAML, LAAT, MultiResCNN, PLM-ICD. **Therefore, we suggest that future studies compare models cautiously using top 50 results and prioritise full code setting improvements.**

3.2 Global Thresholds Do Not Account for Variable Error Cost

In many existing studies, a threshold is defined as the minimum confidence level required to assign a specific code, and a fixed threshold of 0.5 is used without justification (Mullenbach et al., 2018; Wang et al., 2018; Li and Yu, 2020; Vu et al., 2020). This is important because the standard evaluation includes the F1 score, a threshold-dependent classification metric. Edin et al. (2023) found that not tuning the threshold significantly degraded the performance of most models, so they fine-tuned a single threshold by maximising the model’s micro F1 score on the validation set. This does not account for the inherent differences between codes, such as prior probabilities and misclassification costs. This leads to an interesting direction of research: adapting dynamic thresholds (Wu et al., 2019; Alotaibi and Flach, 2021) for classification-based automated coding methods. **Additionally, this highlights the importance of threshold-independent metrics, such as Area Under the Receiver Operating Characteristic Curve (AUC-ROC) and Area Under the Precision-Recall Curve (AUC-PR), as they provide comprehensive assessment of automated coding models across various thresholds.**

3.3 AUC-ROC Score Is Not Ideal for Imbalanced Datasets

AUC-ROC is a binary classification metric. In multi-label classification, a common approach is to convert the task into a one-vs-all setting, where one code (class) is treated as the positive code and all remaining codes as the negative code. Then, the AUC-ROC scores for each code are computed individually. AUC-PR offers a different perspective by focusing on the relationship between precision and recall. One way to calculate AUC-PR is by using Average Precision (AP), which represents the average of precision values at different recall levels as the threshold varies. Since multiple codes can be predicted for each episode, we can calculate

the mean AP for all possible codes, known as the Mean Average Precision (MAP). In other words, the AUC-PR for each code can be calculated using AP, while MAP extends this by averaging the AP values across all predicted codes.

In an imbalanced dataset like MIMIC, the dominance of the negative code can result in a misleadingly high AUC-ROC score. A study by Edin et al. (2023) on three MIMIC splits shows that the SOTA automated coding model, PLM-ICD, consistently achieves macro AUC-ROC scores greater than 95%, indicating it is very effective at scoring relevant codes higher than irrelevant ones. By only looking at this metric, one might infer that this is a robust model. However, PLM-ICD’s MAP remains below 70% across all three splits, indicating that when it predicts a certain code given various threshold values, it is often incorrect. Previous work (Mullenbach et al., 2018; Liu et al., 2021; Yuan et al., 2022; Huang et al., 2022) reported only the AUC-ROC score, which can be misleading, as the model’s precision trade-off is not well reflected. **Thus, we recommend reporting AUC-PR as well as AUC-ROC.**

3.4 Automated Coding Evaluation Should Match Human Coding Evaluation

Existing studies aim for automated coding, which assumes that the model acts as an independent coder. However, their evaluations do not include the common accuracy metrics that are used to measure human performance.

The term ‘accuracy’ can be confusing due to its many possible implementations. Different studies have reported human accuracy in clinical coding, but their definitions of accuracy are inconsistent (Burns et al., 2012). In this position paper, we will examine two implementations of accuracy. Instance accuracy, or Exact Match Rate (EMR), measures the percentage of instances (i.e., episodes of care or medical cases) where the predicted code set exactly matches the true code set. Code accuracy, or Jaccard Score, is defined as the ratio of the size intersection of a predicted code set and a target code set to the size of their union. Instance accuracy is stricter, as it requires a perfect match for every code in an instance.

In automated coding, even if a case contains a single error, human coders must re-code or correct the errors. This means the effort and cost associated with reviewing the entire clinical note are not

	MIMIC-III Clean	MIMIC-IV ICD-9	MIMIC-IV ICD-10
Three digits accuracy	52.84 \pm 0.34	55.22 \pm 0.19	51.17 \pm 0.22
Four digits accuracy	46.21 \pm 0.33	49.28 \pm 0.19	44.97 \pm 0.22
Full code accuracy	44.01 \pm 0.33	46.75 \pm 0.18	42.05 \pm 0.22

Table 2: The code accuracy of PLM-ICD calculated using Jaccard Score, with values spread within the 95% confidence interval determined using the Z-score.

mitigated, highlighting the need for measuring instance accuracy. On the other hand, code accuracy is useful because not all codes are equally important. When assigning a DRG code, the primary diagnosis code is usually the first and most important determinant factor (IHACPA, 2023). Given this, when evaluating coding accuracy, it is beneficial to allow for partial matches, acknowledging that capturing overlap in codes can still be valuable.

We recommend that future studies include both of these accuracy metrics to better demonstrate the performance gap between AI and human coding.

3.5 Low Automation Accuracy Suggests Subset-Specific Automation

When considering instance accuracy, a study in the UK that included 50 episodes of care showed that human accuracy is 54%. In a UK hospital setting, Abdulla et al. (2020) reported an average of 67.5% accuracy each month over four months. In contrast, Among the ‘MIMIC-III clean’, ‘MIMIC-IV ICD-9’, and ‘MIMIC-IV ICD-10’ splits, none of the six advanced models replicated by Edin et al. (2023) achieved an instance accuracy greater than 1.1%.

When considering code accuracy, the median human performance in the UK was 83.2%, with large variance among thirty-two studies (50-98%) (Burns et al., 2012). However, the definition of accuracy is inconsistent across the explored studies; some defined inaccurate coding as inaccurate three digit coding, while the majority defined it as inaccurate four digit coding. In fact, many inaccuracies occur at the four digit level (Burns et al., 2012) instead of the three digit level. We select PLM-ICD, the best model according to Edin et al. (2023), and report its code accuracy with respect to three-digit, four-digit, and full code levels across three MIMIC splits. For each split, we train a single PLM-ICD instance on the full code prediction task, adjusting the accuracy measure across different digits to account for varying evaluation complexities. Table 2 shows that PLM-ICD’s three digit accuracy, the most generous evaluation measure, is much higher than the

other two in all splits. If we compare this result against human accuracy, PLM-ICD is only half as good as an average human at best, indicating there still exists a notable gap to reach full automation.

Instead of full automation, we could consider the task allocation part of clinical coding workflows. A recent study in radiology (Agarwal et al., 2023) suggests that the best approach for combining human expertise with AI is to delegate cases to either AI or humans, rather than having AI augment human decisions. In other words, automating a subset of tractable episodes may be a promising direction for AI coding. The main challenge, however, is the choice of the subset. Figure 4 shows that in MIMIC-IV, approximately only 1% of episodes contain one unique ICD-10-CM three-digit code, while more than half include at least six. This suggests we cannot simply choose the subset based on a single disease or symptom. **We encourage future research to investigate the selection of tractable subsets of care and estimate realistic upper bounds for these subsets.**

3.6 MIMIC Episodes Are Challenging to Fully Automate

The MIMIC cohort consists intensive care unit (ICU) and emergent inpatients, who often present complex conditions requiring multiple diagnoses (see Figure 4) and treatments (Alonso et al., 2020). More details on the MIMIC cohort can be found in Appendix B. Campbell and Giadresco (2020) noted that, compared to inpatient coders, outpatient coders are more concerned that assisted coding will replace their role. One possible reason for this is that outpatient episodes often involve less complex conditions (Alonso et al., 2020). This suggests that outpatient episodes, which are not included in MIMIC, may be better automation candidates.

A common problem in the MIMIC datasets is the imbalanced code distribution, where less than half of the full codes occur at least 10 times, except in the MIMIC-IV ICD-9 collection (see Appendix C for more details). MIMIC-III’s coverage of the ICD-9-CM code space is relatively low, rep-

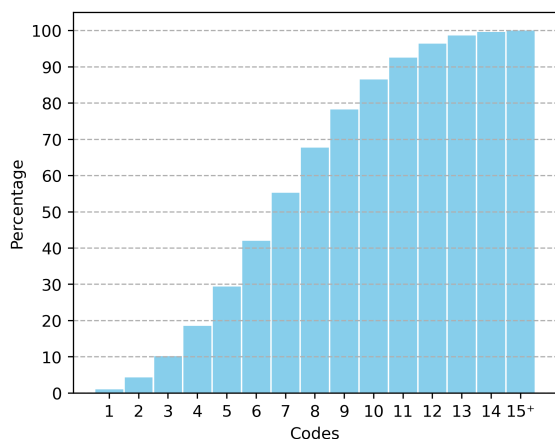


Figure 4: Distribution of MIMIC-IV episodes by the number of unique three-digit ICD-10-CM codes (e.g., G47 for ‘sleep disorders’). About 4% of the episodes have at most two of these three-digit codes.

representing only 50.16% of the possible 17,800 full codes (Wikipedia, 2024). This limitation is worse in MIMIC-IV’s ICD-10-CM/PCS data, which includes 18.78% of the possible 139,000 full codes (Wikipedia, 2024). In other words, the evaluations are omitting a large proportion of ICD codes that may be easier to automate but are missing or rare in MIMIC due to the nature of the dataset.

Large public datasets that cover a broader range of care types are currently lacking. Therefore, developing such datasets would be invaluable for automated coding.

Additionally, we suggest exploring the use of MIMIC for AI coding research in other key parts of the workflow: task allocation, assisted coding, and code auditing. We will discuss this further in Section 4.

3.7 Code Sequence Matters

It is not sufficient to consider only the correctness of the assigned codes; their sequence is also important. For instance, certain conditions have both an underlying etiology and manifest across multiple body systems. The official ICD-10-CM American guideline requires the underlying condition to be sequenced before the manifestation (CMS and NCHS, 2024). Similarly, the official Australian ICD-10-AM guideline requires the anaesthetic codes to be sequenced immediately following the procedure code to which they relate (IHPA, 2022). Therefore, it is evident that clinical coding requires the modeling of both code sequence and dependency. However, the sequence of the target codes is ne-

glected by existing work despite this information being provided in the MIMIC datasets. In fact, both sequence and dependency issues are inherent in some multi-label classification tasks. Many corresponding solutions have been developed (Read et al., 2011; Alvares-Cherman et al., 2012; Yeh et al., 2017; Yang and Liu, 2019). **We recommend that future studies include code sequences in their evaluation to better estimate the real impact on workflow.**

4 New Workflow-Inspired Methodologies

In Section 3, we explored limitations of the current formulation of the task as a multi-label classification problem. Now, we propose alternatives to integrate AI into clinical coding workflows as recommendation systems or asynchronous (offline) auditing assistants. In these cases, information retrieval metrics such as Precision@k, Recall@k, and Coverage Error are more appropriate.

4.1 New Assisted Coding Methodologies

We consider three types of tasks to augment the manual coding process: (1) a sequential task, where the system predicts one code at a time and receives human feedback after each prediction; (2) a recall task, where coders start with a large set of possible codes in mind and select from a set of system suggested codes; and (3) a structural task, where coders take a top-down approach, and instead of predicting complete codes, the system only predicts partial codes.

We propose various assisted coding systems to address the three tasks. While these systems provide a great starting point, they are not exhaustive. The proposed assisted coding systems leverage both human and AI strengths, but their implementation is likely not straightforward. Multiple user designs can be implemented for the same system, leading to varying performance results. Additional challenges include the time needed for humans to adapt to collaborating with AI and the risk of over-trust or under-trust in AI suggestions (Agarwal et al., 2023). Therefore, we recommend including user studies or field tests when evaluating assisted coding systems, measuring not only accuracy and efficiency but also user satisfaction and trust in the system, and user-AI effectiveness.

4.1.1 The Sequential Task

The goal of this task is to transform the multi-label classification problem into a simpler multi-

classification problem. Instead of predicting all codes in a single step, the objective is to assign codes sequentially like real coding practice. We can implement this in several ways:

1. Chain a group of classifiers sequentially, similar to the classifier chain proposed by [Read et al. \(2011\)](#) for related problems in other domains. Each classifier makes a binary choice about assigning a code and receives user feedback, where the feedback is then used as part of the input for the next classifier.
2. Train a single multi-class classifier that predicts one code at a time, receives user feedback, and stops when a predefined termination criterion is met.
3. Treat all ICD codes as vocabulary and use a seq2seq model to predict the codes, prompting the user for feedback after each decoding step.

All three designs add sequential information as part of the input to tackle the code sequencing and dependency issue. Relevant evaluation metrics include Precision@k, the number of steps required to achieve full recall, and the model's convergence rate with human feedback. Importantly, while it is true that human-in-the-loop is not absolutely mandatory in the suggested designs, the inclusion of human input offers significant benefits. These include higher accuracy over long run, real-time error correction, and improved user trust (critical for clinical applications) as users contribute to the model's learning process.

4.1.2 The Recall Task

This task can be framed as a multiple-choice problem, where the objective is to maximise relevant choices while minimising the total number of choices. We propose two designs:

1. A system where all predicted codes are presented to a human expert.
2. A system where high-confidence predicted codes are assigned automatically, whereas low-confidence predicted codes are presented to a human expert.

Both designs can still be approached using multi-label classification; however, the evaluation metrics and optimisation methods should differ. Typically, multi-label classification uses a loss function that

penalises any incorrect outputs. In this task, it is acceptable if some of the model's top confident codes are incorrect, as long as the correct options are included. This approach involves optimising for Recall@k, ensuring high recall with a low value of k. In the second design, the ranking of positive codes becomes important as well. A key challenge in this task is designing effective choice minimisation strategies, which may include grouping by ontology or confidence interval. For example, grouping codes by related ICD categories can reduce the number of options a coder needs to review, thereby speeding up the coding process.

4.2 The Structural Task

The structural task we propose is inspired by [Nguyen et al. \(2023\)](#), but their setup does not incorporate human input. The objective of their work is to use a two-stage decoder that first predicts the parent (i.e., the first three digits) codes and then the child (i.e., the digits following the first three) codes. Their model's parent prediction on 'MIMIC-III Full' achieves micro and macro F1 scores of 29.1% and 69.0% respectively, outperforming the overall micro and macro F1 scores of 10.5% and 58.4% by a considerable margin. This confirms that parent prediction is much simpler for AI models. Based on this, we could design systems where the second stage benefit from human input:

1. The system delegates the challenging task of predicting child codes entirely to humans, while its parent code predictions are used to augment human decision-making.
2. The system could predict a set of parent codes, and upon a human selecting a parent code, it would use that input to predict the relevant child codes. This approach leverages human expertise to refine the AI system's broader categorisations.

In addition to standard classification metrics, explainability metrics (e.g., matching annotated evidence spans) are useful for evaluating the first design, as the system output is intended to inform human decision-making. Hierarchical evaluation metrics (e.g., [Falis et al., 2021](#)) are very useful as they account for the hierarchical code structure, avoiding equal penalisation for all mispredictions. Overall, the proposed designs are more robust in handling rare codes, as parent codes (i.e., the broader

categories) have more training examples. By incorporating a human-in-the-loop approach, the system effectively reduces the likelihood of error propagation from the first stage, resulting in improved reliability and efficiency.

4.3 New Code Auditing Methodologies

Clinical coding is challenging and humans often make errors (Searle et al., 2020; Cheng et al., 2023). If a model achieves high Precision@k, it could be integrated as an offline auditor that starts up immediately after human coders finish coding an episode. For example, if a model's Precision@1 score is 95%, then its most confident code is correct 95% of the time. This is beneficial in addressing under-coding. For instance, the model can flag a high-confidence code that is missing in the episode, prompting the coder for review. Such offline designs ensure AI interventions do not disrupt the human's standard coding workflow, improving coding quality and reducing the back-and-forth communication between coders and auditors. When evaluating these models, considering specific metrics such as the decrease in under-coding incidents and the improvement in billing accuracy post-AI review will provide clear insights into the AI's efficacy in improving the workflow.

5 Summary of Recommendations

1. The top 50 codes have low episode coverage, limiting the generalisation of evaluation results to the full code setting, which is more practical. We recommend against validating methodologies solely with the top 50 codes.
2. Clinical coding involves variable error costs, different thresholds might be optimal for different codes. We recommend reporting more threshold-independent scores for a comprehensive evaluation.
3. In imbalanced datasets like MIMIC, a high AUC-ROC score does not necessarily indicate a high precision score due to the dominance of the negative code (class). We recommend reporting both AUC-PR and AUC-ROC.
4. Existing studies on automated coding often exclude common human accuracy metrics like EMR and Jaccard Score. We recommend including these metrics to demonstrate the practical effectiveness of automated solutions.

5. Our findings indicate that current SOTA models do not achieve human accuracy on MIMIC datasets. We recommend more AI coding research on deliberate task allocation, focusing on a manageable subset of episodes.
6. MIMIC includes acute inpatients, which are challenging even for human coders. We recommend creating datasets of other care types for automated coding research.
7. MIMIC episodes are challenging to automate. We recommend using MIMIC datasets for AI coding research in other parts of clinical coding workflows, such as developing systems for code suggestion and coding audit assistance, as described in our proposed methodologies.
8. Code sequence is part of real-world coding practice. We recommend including it in future evaluations to provide a more comprehensive assessment of system effectiveness.

6 Conclusion

Most studies have approached clinical coding as a traditional multi-label classification task. In this position paper, we conduct a critical review and data analysis of these studies and the MIMIC dataset. We show key shortcomings in existing evaluation methodologies, which fail to align with clinical contexts. We offer eight recommendations for aligning research with the needs of clinical coding workflows. Additionally, we introduce new AI-based methodologies beyond automated coding, proposing alternative research directions with a practical impact on clinical coding workflows.

7 Limitations

Our analysis focuses exclusively on automated coding studies that use the US MIMIC datasets. Therefore, some findings, such as the coverage issue of the top 50 common codes, may not generalise to datasets with patient cohorts different from MIMIC. Likewise, workflow-related discussions may not be universally applicable. This limitation, however, arises from the scarcity of public clinical coding datasets, which is also why MIMIC datasets are widely used in AI coding research. One of our eight recommendations is to create new public datasets for other types of care, which would benefit the entire AI coding community and enable more comprehensive evaluations.

Acknowledgements

This material is partially supported by the Australian Research Council through a Discovery Early Career Researcher Award and the Commonwealth Scientific and Industrial Research Organisation (CSIRO). We also thank the anonymous reviewers for their constructive feedback on our submission.

References

- 3M. 2024a. 3m™ 360 encompass™ audit expert system. https://www.3m.com/3M/en_US/health-information-systems-us/improve-revenue-cycle/coding/audit-compliance/audit-expert/. [Online; accessed 29-July-2024].
- 3M. 2024b. 3m™ 360 encompass™ professional system. https://www.3m.com/3M/en_US/health-information-systems-us/improve-revenue-cycle/coding/professional/360-encompass-professional-system/. [Online; accessed 29-July-2024].
- 3M. 2024c. 3m™ codeassist™ system. https://www.3m.com/3M/en_US/health-information-systems-us/improve-revenue-cycle/coding/professional/code-assist/. [Online; accessed 29-July-2024].
- 3M. 2024d. 3m™ codefinder™ software. https://www.3m.com.au/3M/en_AU/health-information-systems-au/improve-revenue-cycle/coding/facility/codefinder/. [Online; accessed 29-July-2024].
- Suha Abdulla, Natalie Simon, Kelvin Woodhams, Carla Hayman, Mohamed Oumar, Lucy Rose Howroyd, and Gulshan Cindy Sethi. 2020. Improving the quality of clinical coding and payments through student doctor–coder collaboration in a tertiary haematology department. *BMJ Open Quality*, 9:e000723.
- Nikhil Agarwal, Alex Moehring, Pranav Rajpurkar, and Tobias Salz. 2023. Combining human expertise with artificial intelligence: Experimental evidence from radiology. Technical report, National Bureau of Economic Research.
- Vera Alonso, João Vasco Santos, Marta Pinto, Joana Ferreira, Isabel Lema, Fernando Lopes, and Alberto Freitas. 2020. Problems and barriers during the process of clinical coding: a focus group study of coders’ perceptions. *Journal of medical systems*, 44:1–8.
- Reem Alotaibi and Peter Flach. 2021. Multi-label thresholding for cost-sensitive classification. *Neurocomputing*, 436:232–247.
- Everton Alvares-Cherman, Jean Metz, and Maria Carolina Monard. 2012. Incorporating label dependency into the binary relevance framework for multi-label classification. *Expert Systems with Applications*, 39(2):1647–1655.
- Beamtree. 2024. Q coding platform™. <https://beamtree.com.au/our-solutions/coding-platform/>. [Online; accessed 29-July-2024].
- Elaine M Burns, E Rigby, R Mamidanna, A Bottle, P Aylin, P Ziprin, and OD Faiz. 2012. Systematic review of discharge coding accuracy. *Journal of public health*, 34(1):138–148.
- Sharon Campbell and Katrina Giadresco. 2020. Computer-assisted clinical coding: A narrative review of the literature on its benefits, limitations, implementation and impact on clinical coding professionals. *Health Information Management Journal*, 49(1):5–18.
- Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Shengping Liu, and Weifeng Chong. 2020. HyperCore: Hyperbolic and co-graph representation for automatic ICD coding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3105–3114, Online. Association for Computational Linguistics.
- V Chandrabalan. 2023. comorbidipy. <https://github.com/vvcb/comorbidipy>. [Online; accessed 30-July-2024].
- Hua Cheng, Rana Jafari, April Russell, Russell Klopfer, Edmond Lu, Benjamin Striner, and Matthew Gormley. 2023. MDACE: MIMIC documents annotated with code evidence. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7534–7550, Toronto, Canada. Association for Computational Linguistics.
- CMS and NCHS. 2024. Icd-10-cm official guidelines for coding and reporting. <https://www.cms.gov/files/document/fy-2024-icd-10-cm-coding-guidelines-updated-02/01/2024.pdf>.
- Sam Coope, Yoram Bachrach, Andrej Žukov-Gregorič, José Rodriguez, Bogdan Maksak, Conan McMurtie, and Mahyar Bordbar. 2019. A neural architecture for multi-label text classification. In *Intelligent Systems and Applications: Proceedings of the 2018 Intelligent Systems Conference (IntelliSys) Volume 1*, pages 676–691. Springer.
- Joakim Edin, Alexander Junge, Jakob D Havtorn, Lasse Borgholt, Maria Maistro, Tuukka Ruotsalo, and Lars Maaløe. 2023. Automated medical coding on mimic-iii and mimic-iv: A critical review and replicability study. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pages 2572–2582.

- Matúš Falis, Hang Dong, Alexandra Birch, and Beatrice Alex. 2021. [CoPHE: A count-preserving hierarchical evaluation metric in large-scale multi-label text classification](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 907–912, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Malte Feucht, Zhiliang Wu, Sophia Althammer, and Volker Tresp. 2021. [Description-based label attention classifier for explainable ICD-9 classification](#). In *Proceedings of the Seventh Workshop on Noisy User-generated Text (W-NUT 2021)*, pages 62–66, Online. Association for Computational Linguistics.
- Goinvo. 2024. Natural language processing (nlp) software for hospital coding. <https://www.goinvo.com/work/3m-coderyte/>. [Online; accessed 29-July-2024].
- Ary L Goldberger, Luis AN Amaral, Leon Glass, Jeffrey M Hausdorff, Plamen Ch Ivanov, Roger G Mark, Joseph E Mietus, George B Moody, Chung-Kang Peng, and H Eugene Stanley. 2000. Physiobank, physiotoolkit, and physionet: components of a new research resource for complex physiologic signals. *circulation*, 101(23):e215–e220.
- Chao-Wei Huang, Shang-Chi Tsai, and Yun-Nung Chen. 2022. [PLM-ICD: Automatic ICD coding with pre-trained language models](#). In *Proceedings of the 4th Clinical Natural Language Processing Workshop*, pages 10–20, Seattle, WA. Association for Computational Linguistics.
- IHACPA. 2023. Australian refined diagnosis related groups version 11.0. <https://www.ihacpa.gov.au/sites/default/files/2023-03/AR-DRG%20V11.0%20Final%20Report.PDF>. [Online; accessed 30-July-2024].
- IHPA. 2022. National coding advice: Coding rules and faqs for icd-10-am/achi/acs eleventh edition. <https://www.ihacpa.gov.au/sites/default/files/2022-08/National%20Coding%20Advice%20-%20Coding%20Rules%20and%20FAQs%20for%20Eleventh%20Edition.pdf>. [Online; accessed 14-June-2024].
- Alistair EW Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, et al. 2023. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific data*, 10(1):1.
- Alistair EW Johnson, Tom J Pollard, Lu Shen, Li-wei H Lehman, Mengling Feng, Mohammad Ghassemi, Benjamin Moody, Peter Szolovits, Leo Anthony Celi, and Roger G Mark. 2016. MIMIC-III, a freely accessible critical care database. *Scientific data*, 3(1):1–9.
- Sarvnaz Karimi, Xiang Dai, Hamed Hassanzadeh, and Anthony Nguyen. 2017. [Automatic diagnosis coding of radiology reports: A comparison of deep learning and conventional classification methods](#). In *BioNLP 2017*, pages 328–332, Vancouver, Canada. Association for Computational Linguistics.
- Byung-Hak Kim and Varun Ganapathi. 2021. Read, attend, and code: Pushing the limits of medical codes prediction from clinical notes by machines. In *Machine Learning for Healthcare Conference*, pages 196–208. PMLR.
- Anastasia Krithara, James G Mork, Anastasios Nentidis, and Georgios Paliouras. 2023. The road from manual to automatic semantic indexing of biomedical literature: a 10 years journey. *Frontiers in Research Metrics and Analytics*, 8:1250930.
- Fei Li and Hong Yu. 2020. ICD coding from clinical text using multi-filter residual convolutional neural network. In *proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 8180–8187.
- Xingwang Li, Yijia Zhang, Faiz ul Islam, Deshi Dong, Hao Wei, and Mingyu Lu. 2021. IJan: medical code prediction via joint learning attention networks and denoising mechanism. *BMC bioinformatics*, 22:1–21.
- Xinhang Li, Xiangyu Zhao, Yong Zhang, and Chunxiao Xing. 2023. Towards automatic ICD coding via knowledge enhanced multi-task learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*, pages 1238–1248.
- Junping Liu, Shichen Yang, Tao Peng, Xinrong Hu, and Qiang Zhu. 2023. ChatICD: Prompt learning for few-shot ICD coding through ChatGPT. In *2023 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*, pages 4360–4367. IEEE.
- Yang Liu, Hua Cheng, Russell Klopfer, Matthew R. Gormley, and Thomas Schaaf. 2021. [Effective convolutional attention network for multi-label clinical document classification](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 5941–5953, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zichen Liu, Xuyuan Liu, Yanlong Wen, Guoqing Zhao, Fen Xia, and Xiaojie Yuan. 2022. [TreeMAN: Tree-enhanced multimodal attention network for ICD coding](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 3054–3063, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- George Michalopoulos, Michal Malyska, Nicola Sahar, Alexander Wong, and Helen Chen. 2022. [ICDBig-Bird: A contextual embedding model for ICD code classification](#). In *Proceedings of the 21st Workshop on Biomedical Language Processing*, pages 330–336, Dublin, Ireland. Association for Computational Linguistics.

- Antonio Miranda-Escalada, Aitor Gonzalez-Agirre, Jordi Armengol-Estapé, and Martin Krallinger. 2020. Overview of automatic clinical coding: Annotations, guidelines, and solutions for non-english clinical cases at codiesp track of clef ehealth 2020. *CLEF (Working Notes)*, 2020.
- James Mullenbach, Sarah Wiegrefe, Jon Duke, Jimeng Sun, and Jacob Eisenstein. 2018. [Explainable prediction of medical codes from clinical text](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1101–1111, New Orleans, Louisiana. Association for Computational Linguistics.
- Aurélie Névoul, Aude Robert, Francesco Grippi, Claire Morgand, Chiara Orsi, Laszlo Pelikan, Lionel Ramadier, Grégoire Rey, and Pierre Zweigenbaum. 2018. Clef ehealth 2018 multilingual information extraction task overview: Icd10 coding of death certificates in french, hungarian and italian. In *CLEF (Working Notes)*, pages 1–18. CEUR-WS.
- Thanh-Tung Nguyen, Viktor Schlegel, Abhinav Ramesh Kashyap, and Stefan Winkler. 2023. [A two-stage decoder for efficient ICD coding](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 4658–4665, Toronto, Canada. Association for Computational Linguistics.
- John P. Pestian, Chris Brew, Pawel Matykiewicz, DJ Hovermale, Neil Johnson, K. Bretonnel Cohen, and Wlodzislaw Duch. 2007. [A shared task involving multi-label classification of clinical free text](#). In *Biological, translational, and clinical language processing*, pages 97–104, Prague, Czech Republic. Association for Computational Linguistics.
- Aaditya Prakash, Siyuan Zhao, Sadid Hasan, Vivek Datla, Kathy Lee, Ashequl Qadir, Joey Liu, and Oladimeji Farri. 2017. Condensed memory networks for clinical diagnostic inferencing. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31.
- Hude Quan, Vijaya Sundararajan, Patricia Halfon, Andrew Fong, Bernard Burnand, Jean-Christophe Luthi, L Duncan Saunders, Cynthia A Beck, Thomas E Feasby, and William A Ghali. 2005. Coding algorithms for defining comorbidities in icd-9-cm and icd-10 administrative data. *Medical care*, 43:1130–1139.
- Jesse Read, Bernhard Pfahringer, Geoff Holmes, and Eibe Frank. 2011. Classifier chains for multi-label classification. *Machine learning*, 85:333–359.
- William J Rudman, John S Eberhardt, William Pierce, and Susan Hart-Hester. 2009. Healthcare fraud and abuse. *Perspectives in Health Information Management/AHIMA, American Health Information Management Association*, 6(Fall).
- Thomas Searle, Zina Ibrahim, and Richard Dobson. 2020. [Experimental evaluation and development of a silver-standard for the MIMIC-III clinical coding dataset](#). In *Proceedings of the 19th SIGBioMed Workshop on Biomedical Language Processing*, pages 76–85, Online. Association for Computational Linguistics.
- Haoran Shi, Pengtao Xie, Zhiting Hu, Ming Zhang, and Eric P Xing. 2017. Towards automated icd coding using deep learning. *arXiv preprint arXiv:1711.04075*.
- Wei Shi, Jiewen Wu, Xiwen Yang, Nancy Chen, Ivan Ho Mien, Jung-Jae Kim, and Pavitra Krishnaswamy. 2021. [Analyzing code embeddings for coding clinical narratives](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 4665–4672, Online. Association for Computational Linguistics.
- Wei Sun, Shaoxiong Ji, Erik Cambria, and Pekka Martinen. 2021. Multitask recalibrated aggregation network for medical code prediction. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 367–383. Springer.
- Fei Teng, Zheng Ma, Jie Chen, Ming Xiao, and Lufei Huang. 2020. Automatic medical code assignment via deep learning approach for intelligent healthcare. *IEEE journal of biomedical and health informatics*, 24(9):2506–2515.
- Carl Van Walraven, Peter C Austin, Alison Jennings, Hude Quan, and Alan J Forster. 2009. A modification of the elixhauser comorbidity measures into a point system for hospital death using administrative data. *Medical care*, 47:626–633.
- Thanh Vu, Dat Quoc Nguyen, and Anthony Nguyen. 2020. [A label attention model for icd coding from clinical text](#). In *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 3335–3341. International Joint Conferences on Artificial Intelligence Organization. Main track.
- Guoyin Wang, Chunyuan Li, Wenlin Wang, Yizhe Zhang, Dinghan Shen, Xinyuan Zhang, Ricardo Henao, and Lawrence Carin. 2018. [Joint embedding of words and labels for text classification](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2321–2331, Melbourne, Australia. Association for Computational Linguistics.
- 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.
- Wikipedia. 2024. ICD-10 — Wikipedia, the free encyclopedia. <https://en.wikipedia.org/wiki/ICD-10>. [Online; accessed 11-June-2024].

- Jiawei Wu, Wenhan Xiong, and William Yang Wang. 2019. [Learning to learn and predict: A meta-learning approach for multi-label classification](#). 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 4354–4364, Hong Kong, China. Association for Computational Linguistics.
- Pengtao Xie and Eric Xing. 2018. [A neural architecture for automated ICD coding](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1066–1076, Melbourne, Australia. Association for Computational Linguistics.
- Xiancheng Xie, Yun Xiong, Philip S Yu, and Yangyong Zhu. 2019. Ehr coding with multi-scale feature attention and structured knowledge graph propagation. In *Proceedings of the 28th ACM international conference on information and knowledge management*.
- Zhenyu Yang and Guojing Liu. 2019. Hierarchical sequence-to-sequence model for multi-label text classification. *IEEE Access*, 7:153012–153020.
- Zhichao Yang, Shufan Wang, Bhanu Pratap Singh Rawat, Avijit Mitra, and Hong Yu. 2022. [Knowledge injected prompt based fine-tuning for multi-label few-shot ICD coding](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Chih-Kuan Yeh, Wei-Chieh Wu, Wei-Jen Ko, and Yu-Chiang Frank Wang. 2017. Learning deep latent space for multi-label classification. In *Proceedings of the AAAI conference on artificial intelligence*, volume 31.
- Zheng Yuan, Chuanqi Tan, and Songfang Huang. 2022. [Code synonyms do matter: Multiple synonyms matching network for automatic ICD coding](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 808–814, Dublin, Ireland. Association for Computational Linguistics.
- Zachariah Zhang, Jingshu Liu, and Narges Razavian. 2020. [BERT-XML: Large scale automated ICD coding using BERT pretraining](#). In *Proceedings of the 3rd Clinical Natural Language Processing Workshop*, pages 24–34, Online. Association for Computational Linguistics.
- Tong Zhou, Pengfei Cao, Yubo Chen, Kang Liu, Jun Zhao, Kun Niu, Weifeng Chong, and Shengping Liu. 2021. [Automatic ICD coding via interactive shared representation networks with self-distillation mechanism](#). 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 5948–5957, Online. Association for Computational Linguistics.

A AI Coding Task Definition

Previous studies have approached clinical coding as a fine-grained multi-label classification task. The objective is to create a model that maps an input text (usually a discharge summary) to a set of labels (ICD codes). The term ‘fine-grained’ indicates that the code space is extremely large, covering thousands of codes or more. Both code frequency and the number of associated codes per instance (an episode or medical case) is highly imbalanced. The code space is presented as a hierarchy, where codes exhibit parent-child relationships.

B Descriptive Statistics of the MIMIC Cohort

An overview of the patient cohorts in MIMIC-III/IV is shown in Table 3. The preprocess code is built upon the work of [Edin et al. \(2023\)](#). Admissions without assigned ICD codes are removed. Age is calculated based on the patients’ age at the time of admission. The Elixhauser index is calculated using ([Chandrabalan, 2023](#)), with [Quan et al. \(2005\)](#)’s mapping and [Van Walraven et al. \(2009\)](#)’s score weighting, and is adjusted for patient age. LOS refers to the length of stay. Some admissions contain missing or invalid data, such as missing LOS values or cases where the discharge time is earlier than the admission time. Note that many ICU LOS values are missing in MIMIC-IV. These admissions are excluded from corresponding statistical calculations, and the number of included admissions, N , is provided accordingly.

C Overview of the MIMIC datasets

Table 4 shows descriptive statistics of the ICD coding data in MIMIC-III/IV. In the ICD-10 collection of MIMIC-IV, the median document length is 1,492 words with an interquartile range (IQR) of 1,147–1,931, and the median number of full codes per summary is 15, with an IQR of 10–21. These are comparable across the other two collections. The varying number of codes in most summaries reflects considerable differences in the complexity of medical issues documented, aligning with the fluctuating Elixhauser index presented in Table 3.

	ICD-9		ICD-10
	MIMIC-III	MIMIC-IV	MIMIC-IV
Years collected	2001–2012	2008–2019	2008–2019
Number of admissions	52,722	209,359	122,316
Number of patients	41,126	97,727	65,685
Age: Median (IQR)	63 (49–77)	61 (48–74)	61 (48–72)
Elixhauser index: Median (IQR)	7 (2–14)	6 (2–12)	7 (2–15)
LOS: Median (IQR) [days]	7.18 (4.26–12.78) <i>N=52,671</i>	3.28 (1.79–5.96) <i>N=209,306</i>	3.88 (2.02–6.96) <i>N=122,297</i>
ICU LOS: Median (IQR) [days]	2.43 (1.30–5.29) <i>N=51,938</i>	1.95 (1.09–3.84) <i>N=38,717</i>	2.14 (1.18–4.30) <i>N=26,606</i>

Table 3: Overview of the MIMIC-III v1.4 and MIMIC-IV v2.2 patient cohort.

	ICD-9		ICD-10
	MIMIC-III	MIMIC-IV	MIMIC-IV
Words per document: Median (IQR)	1,375 (965–1,900)	1,320 (997–1,715)	1,492 (1,147–1,931)
Full codes per document: Median (IQR)	14 (10–20)	12 (8–17)	15 (10–21)
Number of unique three-digit codes	1,606	1,712	2,239
Number of unique four-digit codes	6,120	7,337	11,254
Number of unique full codes	8,929	11,331	26,098
Full code coverage [%]	50.16	63.66	18.78
Full code freq ≥ 10 [%]	41.23	54.28	30.43

Table 4: Overview of the MIMIC-III v1.4 and MIMIC-IV v2.2 ICD coding data. The preprocess code is built upon the work of [Edin et al. \(2023\)](#). Only admissions without assigned ICD codes are removed.