# TEACHING CONSENSUS RULES (AND EXCEPTIONS) TO LLMS FOR TRUSTWORTHY MEDICAL REASONING

**Anonymous authors**Paper under double-blind review

000

001

002 003 004

010 011

012

013

014

016

017

018

019

021

025

026

027

028

029

031

034

037

040

041

042

043

044

046

047

048

051

052

#### **ABSTRACT**

Machine learning for early prediction in medicine has recently shown breakthrough performance, however, the focus on improving prediction accuracy has led to a neglect of faithful explanations that are required to gain the trust of medical practitioners. The goal of this paper is to teach LLMs to follow medical consensus guidelines step-by-step in their reasoning and prediction process. Since consensus guidelines are ubiquitous in medicine, instantiations of verbalized medical inference rules to electronic health records provide data for fine-tuning LLMs to learn consensus rules and possible exceptions thereof for many medical areas. Consensus rules also enable an automatic evaluation of the model's inference process regarding its derivation correctness (evaluating correct and faithful deduction of a conclusion from given premises) and value correctness (comparing predicted values against real-world measurements). We exemplify our work using the complex Sepsis-3 consensus definition. Our experiments show that small fine-tuned models outperform one-shot learning of considerably larger LLMs and of models that are trained on medical texts including consensus definitions under all evaluation metrics. Since fine-tuning on verbalized rule instantiations of a specific medical area yields nearly perfect derivation correctness for rules (and exceptions) on unseen patient data in that area, the bottleneck for early prediction is not out-of-distribution generalization, but the orthogonal problem of generalization into the future by forecasting sparsely and irregularly sampled clinical variables. We show that the latter results can be improved by integrating the output representations of a time series forecasting model with the LLM in a multimodal setup.

## 1 Introduction

Medical consensus definitions are guidelines stated by a representative group of experts on how to diagnose and treat a disease based on clinical evidence. From the perspective of logical inference, consensus guidelines include deductive and inductive inference rules. Deductive rules have the form of if-then relations where the conclusion is certain given the correctness of the premise, for example, in mapping thresholds on clinical measurements to step functions of diseases. If used for early prediction purposes (a.k.a. prognosis), inductive inference rules are required where the if-then relation between premise and conclusion is probabilistic. An illustrative example for a complex consensus guideline is the Sepsis-3 definition that identifies an organ dysfunction as an acute change in total SOFA score  $\geq 2$  points consequent to an infection (Singer et al., 2016; Seymour et al., 2016). The SOFA (Sepsis-Related Organ Failure Assessment) score is based on definitions for six organ system subscores, depending for their part on thresholded clinical variables observed during a 24h window ((Vincent et al., 1996), see Table 5 in Appendix A.1). These calculations comprise a logical rule system where time series forecasting (TSF) of clinical variables represents inductive rules, which are composed with deductive rules that calculate extrema over time, map clinical measurements onto step functions, and calculate changes of scores over time (see Figure 1).

The goal of this work is teach LLMs to follow the deductive and inductive rules of medical consensus guidelines step-by-step in their prediction and reasoning process<sup>1</sup>, in order to foster the trust of medical practitioners in the generated diagnosis or prognosis. We exemplify our work using the

<sup>&</sup>lt;sup>1</sup>In the following, in order to avoid an anthropomorphic attribution of "reasoning" capabilities to LLMs, we will speak of medical "rules" that are taught, and of an "inference process" that is being evaluated for the LLM.

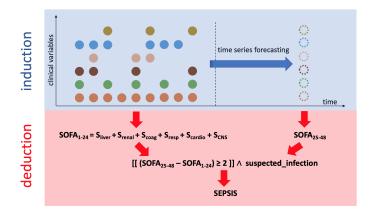


Figure 1: Inductive and deductive inference rules in the Sepsis-3 Consensus Definition (Singer et al., 2016; Seymour et al., 2016). Deductive rules calculate extrema over time, map thresholds onto step functions for SOFA scores, and calculate total SOFA and changes over time. Inductive rules involve time series forecasting of clinical variables 24 hours into the future.

complex Sepsis-3 consensus definition. Instantiations of verbalized medical inference rules to clinical patient data allow the model to learn the consensus rule and possible exceptions to the rule from patient examples in many medical areas (see Figure 2). Our work transfers learning of compositional inference tasks from mathematical problems to medical inference, with several advantages.

First, consensus definitions are ubiquitous in medicine, ranging from guidelines for mental disorders (American Psychiatric Association, 2013) to neurological (McDonald et al., 2001) and physiological diseases (KDIGO Acute Kidney Injury Work Group, 2012). For each area, verbalization of consensus rules can be done automatically by using templates that describe each step of an application of a consensus rule system to patient data. Furthermore, human curation can be integrated by annotating verbalized inference chains with possible exceptions and corrections to the rule. In contrast to prompting LLMs, supervised fine-tuning allows learning of rules and exceptions from example instantiations to patient data.

Second, the use of consensus rules enables an exact and automatic evaluation of the model's inference process against a trusted medical gold standard. We present an evaluation setup that differentiates the derivation correctness of a trained model — measuring a model's ability to learn to correctly deduce a conclusion from a given premise — from a model's value correctness — comparing the numerical value predicted in each inference step against real-world measurements. An evaluation of derivation correctness also can be seen to measure the faithfulness of the model's inference process in the sense of checking for an accurate representation of the reasoning process behind the model's prediction.

Our results show that small fine-tuned models (LLaMA 8B parameters) outperform one-shot learning of considerably larger LLMs (LLaMA 70B parameters) and LLMs that are trained on medical texts including the original consensus definitions (Me-LLaMA 8B parameters) under all evaluation metrics. In particular, fine-tuned LLMs show nearly perfect generalization to unseen patient data with respect to derivation correctness of rules and exceptions. This demonstrates that an adaptation of an LLM to verbalized rule instantiations for a specific medical area likely guarantees consistent results in that area, whereas training on the abstract definition texts themselves is not sufficient. We conjecture that generalization across consensus definitions in areas as different as psychiatry (based on interviews (American Psychiatric Association, 2013)), neurology (based partially on magnetic resonance imaging (McDonald et al., 2001)), or physiological diseases (based on vital signals and lab tests (KDIGO Acute Kidney Injury Work Group, 2012; Cederholm et al., 2019; Singer et al., 2016; Seymour et al., 2016)) is difficult to bridge without supervision. The bottleneck is thus not out-of-distribution generalization (Chu et al., 2025), but the orthogonal problem of generalization into the future, i.e., the inherent complexity of time series forecasting of sparsely and irregularly sampled clinical variables. The latter results can be improved by moving from a text-based encoding of clinical measurements to a multi-modal approach where the output representations of a time series forecasting model are fed as additional input to the LLM.

# 2 RELATED WORK

108

109 110

111

112

113

114

115

116

117

118

119

120

121

122

123

124

125

126 127

128

129

130

131

132

133

134

135

136

137

138

139

140

141

142

143

144

145 146

147

148

149

150

151

152

153

154

155

156

157

158

159

160

161

**Reasoning in LLMs** Recent evaluations of the "reasoning" process of LLMs have shown that they are able to correctly solve sub-tasks of multi-step inference tasks, but fail to compose them into a fully correct inference path, especially in extrapolation to more complex out-of-distribution data (Dziri et al., 2023; Zhang et al., 2023; Saparov & He, 2023; Yang et al., 2024; Mondorf & Plank, 2024). For compositional inference problems in mathematics or logics, fine-tuning on so-called "scratchpads" has been shown to yield consistently strong generalization (Nye et al., 2021; Hochlehnert et al., 2025). To our knowledge, our work is the first one to apply an automatic generation of scratchpads for fine-tuning and step-by-step evaluation to the inference process of medical LLMs. Recent works have addressed the aspect of faithfulness of chain-of-thought explanations of LLMs, asking whether they accurately represent the underlying inference process of the model (Jacovi & Goldberg, 2020). Faithfulness in chain-of-thought reasoning has been enforced by incorporating deterministic solvers into the inference process (Lyu et al., 2023; Xu et al., 2024). Evaluation of faithfulness has been done manually for LLMs with explicit (Turpin et al., 2023) and implicit biases (Arcuschin et al., 2025), or by defining automatic evaluation metrics of chain-of-thought inference processes (Lanham et al., 2023; Wang et al., 2025; Chen et al., 2025). The latter metrics are based on accuracy or confidence of the generated LLM outputs, and can thus be considered "reference-free" evaluation metrics, whereas a consensus rule can be used as a unique gold standard reference for evaluation.

**Medical Reasoning** Medical reasoning has been been characterized to include two contrasting mechanisms, termed analytic versus non-analytic strategies (Eva, 2005), causal versus example-based knowledge (Norman, 2005), or contrasting external clinical evidence with individual clinical expertise (Sackett et al., 1996). To our knowledge, no attempt has been made so far to differentiate these aspects in teaching and evaluating LLMs for healthcare. Instead, most evaluations are based on medical question-answering (Liu et al., 2023; Singhal et al., 2025), while it has been shown that the vast majority of questions in medical QA benchmarks can be answered by factual recall, not requiring multi-step inference at all (Thapa et al., 2025). While excelling in medical question-answering tasks, few-shot LLMs suffer from prompt brittleness (Hochlehnert et al., 2025; Reese et al., 2024) and fail in basic TSF tasks (Merrill et al., 2024; Tan et al., 2024) and lack adherence to consensus guidelines (Hager et al., 2024). This discrepancy calls for evaluations of medical LLMs that go beyond the final result and instead evaluate the inference process itself. Various approaches have been presented to teach LLMs medical inference rules, e.g., by pre-training or fine-tuning on medical texts including consensus guidelines (Chen et al., 2023; Xie et al., 2025), or by incorporating rule-based knowledge in knowledge graphs (Wu et al., 2025), reward models (Yun et al., 2025; Chen et al., 2024; Zhang et al., 2025), or by manual annotation by clinical experts (Liu et al., 2025). However, with the exception of the last work which relies on manual evaluation, all works are evaluated on standard medical question-answering benchmarks. Our work uses trusted consensus rules as training data and, instantiated to different patients, as gold standard references for evaluation, and obviates the need for manual evaluation.

Early Prediction of Sepsis Machine learning for early prediction of sepsis starts by labeling clinical data<sup>2</sup>, by applying a Sepsis consensus definition <sup>3</sup> to future clinical observations that are unseen to the model, and then training models to predict the label based on clinical measurements observed several hours before. Examples are the machine learning approaches taken by the 104 participants in the PhysioNet Challenge on "Early Prediction of Sepsis from Clinical Data" (Reyna et al., 2019), or the 21 approaches described in a recent overview (Moor et al., 2021). Machine learning for early prediction of sepsis includes a broad variety of learning algorithms, from linear learners to deep learning (see Reyna et al. (2019) and Moor et al. (2021)), up to most recent works using Transformers (Choi et al., 2024), graph neural networks (Yin et al., 2025), and pretrained LLMs (Li et al., 2024). The setup of predicting a label that is the outcome of a consensus definition can be abstractly viewed as the prediction of an effect that is caused by applying a consensus definition to future values of clinical variables. Another option is to directly predict the causes by forecasting clinical variables, and determine the effect by applying the consensus definition to the forecasted values (Staniek et al., 2024). Our approach is most similar to the latter, with the important difference that our approach not only learns rules, but also exceptions thereof from patient data.

<sup>&</sup>lt;sup>2</sup>E.g., MIMIC-III (Johnson et al., 2016), MIMIC-IV (Johnson et al., 2023), or eICU (Pollard et al., 2018).

<sup>&</sup>lt;sup>3</sup>Sepsis-3 or Sepsis-2 (Levy et al., 2003; Dellinger et al., 2013) in earlier works.

## 3 TEACHING LLMS TO DIAGNOSE ACCORDING TO THE SEPSIS-3 DEFINITION

**Deductive Inference** Sepsis-3 contains several deductive inference steps: A first step is a calculation of "worst" values (minima or maxima) of clinical variables over 24 hours. The crucial deductive task in Sepsis-3 can be seen as a deductive syllogism consisting of rules (the major premises) that are instantiated to thresholds on clinical variables (the minor premises) that are mapped to step functions for six SOFA subscores, yielding scores from 0-4 (the conclusions). Each SOFA subscore corresponds to an organ system, namely the central nervous system, and the cardiovascuolar, respiratory, coagulation, liver, and renal organ systems. The rules themselves are abbreviated as  $S_{CNS}$ ,  $S_{cardio}$ ,  $S_{resp}$ ,  $S_{coag}$ ,  $S_{liver}$ , and  $S_{renal}$  in the lower part of Figure 1 and shown in the columns in Table 5 in Appendix A.1. The total SOFA score is calculated by summing these six subscores to a score ranging from 0-24, once for clinical measurements during the first 24 hours (SOFA $_{1-24}$ ), and for predicted measurements in the next 24 hours (SOFA $_{25-48}$ ). A further deductive inference step is the computation of a change  $\geq 2$  as

$$SOFA_{diff} := [(SOFA_{25:48} - SOFA_{1:24}) \ge 2],$$
 (1)

where [a] = 1 if a is true, 0 otherwise. Finally, a binary Sepsis label is assessed by combining the indicator function in Equation 1 with a binary indicator suspected\_infection<sup>4</sup>, yielding

$$SEPSIS := SOFA_{diff} \land suspected\_infection.$$
 (2)

Inductive Inference Sepsis-3 is laid out to detect a life-threatening organ dysfunction by an acute change in the total SOFA score. Measuring an increase in SOFA score over time requires TSF of future values of a set of clinical features F (see upper part of Figure 1), and an application of the SOFA definition to current and future clinical values. Formally, given a representation x of an input time series, an output vector of predicted clinical values  $\hat{y}_t \in \mathbb{R}^{|F|}$  is produced. The features used in our experiments are 131 clinical measurements extracted from the MIMIC-III database (Johnson et al., 2016). A full list is given in Appendix A.4. The TSF task in our experiments is defined to predict the next 24 hours from a history of preceding 24 hours. Dedicated models to perform the TSF task are described in Appendix A.2.

Teaching Verbalized Consensus Rules to Autoregressive LLMs The key requirement in teaching LLMs to generate a chain of inference rules that adheres to a certain definition is to deploy an autoregressive architecture where next-token prediction is based on a history of previously predicted tokens, and token-wise errors on a target inference chain can be backpropagated through the system. This procedure coincides with standard supervised fine-tuning, where an input question, a goldstandard verbalization of the inference process, and the answer, is provided to the model during training. An example for fine-tuning data including inference rules following the Sepsis-3 definition is shown in Figure 2, left column. The data starts with a verbalization of a sparse multivariate input time series of measurements for a 24 hour window, followed by an instruction to classify the patient as septic in the next 24 hours, given information about suspected infection<sup>5</sup>. The third text block in the fine-tuning data consists of a gold standard inference chain explaining the scores of the six SOFA systems from the corresponding clinical measurements, and computing their sum, given the measurements of the first 24 hour window. The fourth text block includes forecasts of the clinical variables relevant for SOFA for the next 24 hours, and the corresponding inference about the six SOFA subscores, and their sum. The last text block contains a gold standard inference about the Sepsis prediction and the responsible SOFA system, if applicable.

**Handling Exceptions to the Rules** Since consensus rules cannot capture every individual disease progression, it is important to consider the possibility of manual corrections to consensus-based inference. Imagine a feedback loop where a medical practitioner edits the output of an LLM trained on a consensus definition: While corrections and additions to the inference process can be handled by repeated fine-tuning on manually curated inference chains, exceptions to the consensus rule

<sup>&</sup>lt;sup>4</sup>Following Singer et al. (2016); Seymour et al. (2016), a suspected infection is defined as a combination of antibiotics treatment and blood cultures, starting within the first 24 hours after admission.

<sup>&</sup>lt;sup>5</sup>The onset of suspicion of infection is thus given prior to the onset of salient organ dysfunction in our approach, corresponding to labeling scheme H1 of Cohen et al. (2024), and consistent with other work on early prediction of sepsis where suspected infection is included in the list of input features (Nemati et al., 2018).

219

220

221

222

224

225

226

227

228

229

230

231

234

235

236

237

238

239

240

241

242

243

244

245

246

247

248249

251

252253

Patient is 75.0 years old and is male. Given all the information in this text, answer the question at the end. Here are the measurements: DBP at time -22.37: 49.0, SBP at time -22.37: 105.0, DBP at time -20.37: 52.0, GCS\_eye at time -20.37: 4.0, GCS\_motor at time -20.37: 6.0, GCS\_verbal at time -20.37: 1.0, SBP at time -20.37: 117.0, DBP at time -19.37: 56.0, FiO2 at time -19.37: 0.5, SBP at time -19.37: 127.0, DBP at time -18.37: 43.0,...

Now answer the following question: The doctors suspect an infection, based on this information and the other information in this text, will the patient be classified as septic tomorrow?

First we need to calculate the SOFA scores given the extracted values. The SOFA scores for the current time are the following: The minimum value of GCS\_eye is 4.0, GCS\_motor is 6.0 and GCS\_verbal is 1.0, this produces the sum 11.0 and means the CNS SOFA is 2. Because minimum MAP is 55.333, max Dopamine is 0, max Dobutamine is 0, max Epinephrine is 0 and max Norepinephrine is 0 with a patient weight of 62.8 kg, the cardiovascular SOFA is 1. Given that minimum PO2 is 100.0 and minimum FiO2 is 0.5 the calculated PAO2FIO2 is 200.0, this means the respiratory SOFA is 2. Because the minimum Platelet count is 310.0 the coagulation SOFA is 0. The maximum Bilirubin (Total) is 1 leading to a liver SOFA of 0. Because total Urine output is 1095.0 and maximum creatinine in the blood is 0.4 the renal SOFA is 0. To summarize: the patient has a total SOFA score of 5.

Now we need to calculate the SOFA scores with forecasted values. The SOFA scores in the future based on the forecasted values are the following: The minimum value of GCS\_eye will be 4.0, GCS\_motor will be 6.0 and GCS\_verbal will be 1.0, this produces the sum 11.0 and means the CNS SOFA will be 2. Because future minimum MAP will be 65.333, future max Dopamine will be 0, future max Dobutamine will be 0, future max Epinephrine will be 0 and future max Norepinephrine will be 0 with a patient weight of 62.8 kg, the cardiovascular SOFA will be 1. Given that minimum PO2 will be 100.0 and minimum FiO2 will be 0.5 the forecasted PAO2FIO2 will be 200.0, this means the respiratory SOFA will be 2. Because the Platelet count will be 310.0 the coagulation SOFA is going to be 0. The maximum Bilirubin (Total) will be 1 leading to a liver SOFA of 0. Because Urine output will be 150.0 and maximum creatinine in the blood will be 0.4 the renal SOFA will be 4. To summarize: the patient will have a future total SOFA score of 9.

This calculation means that the patient will likely experience a kidney failure since SOFA increased by 4. The patient will develop sepsis in the next 24 hours, because total SOFA increased by 4 and infection is suspected.

Patient is 75.0 years old and is male. Given all the information in this text, answer the question at the end. Here are the measurements: DBP at time -22.37: 49.0, SBP at time -22.37: 105.0, DBP at time -20.37: 52.0, GCS\_eye at time -20.37: 4.0, GCS\_motor at time -20.37: 6.0, GCS\_verbal at time -20.37: 1.0, SBP at time -20.37: 117.0, DBP at time -19.37: 56.0, FiO2 at time -19.37: 0.5, SBP at time -19.37: 127.0, DBP at time -18.37: 43.0,...

Now answer the following question: The patient has an existing precondition given by the ICD-10 code N18.9. The doctors suspect an infection, based on this information and the other information in this text, will the patient be classified as septic tomorrow?

First we need to calculate the SOFA scores given the extracted values. The SOFA scores for the current time are the following: The minimum value of GCS\_eye is 4.0, GCS\_motor is 6.0 and GCS\_verbal is 1.0, this produces the sum 11.0 and means the CNS SOFA is 2. Because minimum MAP is 55.333, max Dopamine is 0, max Dobutamine is 0, max Epinephrine is 0 and max Norepinephrine is 0 with a patient weight of 62.8 kg, the cardiovascular SOFA is 1. Given that minimum PO2 is 100.0 and minimum FiO2 is 0.5 the calculated PAO2FIO2 is 200.0, this means the respiratory SOFA is 2. Because the minimum Platelet count is 310.0 the coagulation SOFA is 0. The maximum Bilirubin (Total) is 1 leading to a liver SOFA of 0. Because total Urine output is 1095.0 and maximum creatinine in the blood is 0.4 the renal SOFA is 0. To summarize: the patient has a total SOFA score of 5.

Now we need to calculate the SOFA scores with forecasted values. The SOFA scores in the future based on the forecasted values are the following: The minimum value of GCS\_eye will be 4.0, GCS\_motor will be 6.0 and GCS\_verbal will be 1.0, this produces the sum 11.0 and means the CNS SOFA will be 2. Because future minimum MAP will be 65.333, future max Dopamine will be 0, future max Dobutamine will be 0, future max Epinephrine will be 0 and future max Norepinephrine will be 0 with a patient weight of 62.8 kg, the cardiovascular SOFA will be 1. Given that minimum PO2 will be 100.0 and minimum FiO2 will be 0.5 the forecasted PAO2FIO2 will be 200.0, this means the respiratory SOFA will be 2. Because the Platelet count will be 310.0 the coagulation SOFA is going to be 0. The maximum Bilirubin (Total) will be 1 leading to a liver SOFA of 0. Because Urine output will be 150.0 and maximum creatinine in the blood will be 0.4 the renal SOFA will be 4. To summarize: the patient will have a future total SOFA score of 5.

The patient will not develop sepsis in the next 24 hours, because total SOFA increased by 0 and infection is suspected.

Figure 2: Fine-tuning data including a verbalization of Sepsis-3 inference (left column) and inference under an exception due to medical preconditions (right column). Differences are shown in bold blue font. The general prompt is shown above the horizontal line, the gold standard answer below.

254255256

257

258

259

260

261

262

263

264

265

266

267

268

might cause conflicts in machine learning. A possible example are medical preconditions with specific treatments, for example, dialysis in case of chronic kidney disease, which will alter clinical measurements and thus the kidney SOFA score. A medical practitioner might decide to disregarded the kidney SOFA in the calculation of total SOFA, leading to an exception to the Sepsis-3 definition. We test this hypothetical scenario by synthetic data. We synthesized fine-tuning data including exceptions to the rule by adding preconditions for five organ systems in form of ICD-10 codes. During data generation, each datapoint was randomly assigned a precondition, either one of the predefined types or the no-precondition option, with equal probability. An example for a verbalization is shown in Figure 2, right column. The format is the same as for inference according to the definition, except for the inclusion of a precondition in the second text block — in our example, the ICD code N18.9 indicating chronic kidney disease — implying an exception to the rule to disregard the SOFA score of the corresponding organ system in the fourth and fifth text blocks. We prepared the data such that some preconditions are seen during training and testing, while we also chose some codes from the same ICD class as out-of-distribution data that were only seen at test time. The codes and their distribution are shown in Tables 9 and 10 in Appendix A.6.

270271272

281

283 284

287288

289 290

291

292

293

295

296

297

298

299

300

301

302

303

304

305 306

307

308

309

310

311

312

313

314

315

316

317

318

319 320

321

322

323

Table 1: Derivation correctness of predicted values. Forced derivation correctness in brackets.

variable	one-shot	one-shot-70B	me-llama	fine-tuned	pipeline	multimodal
S <sub>CNS</sub>	0.613	0.651	0.508	1.000	1.000	1.000
$S_{cardio}$	0.485	0.634	0.382	0.998	0.998	0.975
$S_{resp}$	0.389	0.291	0.466	0.995	0.996	0.996
$S_{coag}$	0.622	0.615	0.611	0.966	0.966	0.970
$S_{liver}$	0.546	0.860	0.710	0.993	0.992	0.992
$S_{renal}$	0.502	0.675	0.483	1.000	0.998	1.000
$SOFA_{1:24}$	0.639	0.886	0.953	1.000	1.000	0.999
future S <sub>CNS</sub>	0.444 (0.817)	0.472 (0.839)	0.412 (0.782)	1.000 (1.000)	1.000 (1.000)	1.000 (1.000)
future S <sub>cardio</sub>	0.565 (0.660)	0.583 (0.668)	0.483 (0.655)	0.987 (0.997)	0.975 (0.997)	0.965 (0.997)
future S <sub>resp</sub>	0.508 (0.735)	0.285 (0.802)	0.456 (0.715)	0.995 (0.996)	0.996 (0.996)	0.995 (0.996)
future S <sub>coag</sub>	0.747 (0.800)	0.616 (0.787)	0.644 (0.800)	0.967 (1.000)	0.970 (0.999)	0.966 (1.000)
future S <sub>liver</sub>	0.529 (0.797)	0.843 (0.952)	0.703 (0.627)	0.992 (0.999)	0.992 (0.999)	0.992 (1.000)
future S <sub>renal</sub>	0.632 (0.689)	0.694 (0.714)	0.505 (0.695)	1.000 (1.000)	1.000 (1.000)	1.000 (1.000)
$SOFA_{25:48}$	0.568 (0.916)	0.860 (0.996)	0.935 (0.992)	1.000 (1.000)	0.999 (1.000)	1.000 (1.000)
$SOFA_{diff}$	0.799 (0.667)	0.831 (0.703)	0.947 (0.722)	1.000 (1.000)	1.000 (1.000)	1.000 (1.000)
SEPSIS	0.816 (0.728)	0.893 (0.738)	0.895 (0.593)	1.000 (1.000)	1.000 (1.000)	1.000 (1.000)

# 4 EXPERIMENTAL SETUP

**Data and Models** In our experiments, we use electronic health records (EHRs) from the MIMIC-III data (Johnson et al., 2016). After filtering for patients with an ICU stay of at least 24 hours with reported gender and age of at least 18 years, our dataset contained 44,858 ICU stays with 56 million data points. We split the data into partitions for fine-tuning (28,708), development (7,270), and testing (8,880). For computational reasons, we further subsampled 15,000 datapoints for fine-tuning, and 3,000 datapoints for development and testing, respectively. The resulting percentage of positive Sepsis cases in the test set was 7.33%. As features, we considered 131 clinical variables and the demographic variables gender and age (see Appendix A.4). This selection comprises all vital signs and laboratory values used in the PhysioNet challenge for early prediction of sepsis (Reyna et al., 2019), together with information on suspected infection as in Nemati et al. (2018). The time series of clinical observations in the data were split into a 24 hour observation window, followed by a 24 hour prediction window. During training and testing, we use a sliding window of 24 hours so that full admission days are given as input and output. Using the consensus definition given by Table 5 and Equation 2, SOFA scores and Sepsis label were calculated deterministically for the given data. Following this, features and SOFA scores were verbalized into texts as seen in Figure 2. In TSF, missing values are being carried forward from the previous day only.

The basic LLM used in our experiments is a pretrained Llama-3 model with 8B parameters (Grattafiori et al., 2024) <sup>6</sup>. We use the 8B LLM in **one-shot** mode using the prompt shown in Appendix A.5. Our fine-tuned model uses LORA adapters (Hu et al., 2022) on verbalizations of inference processes. Furthermore, we use a dedicated medical TSF forecaster, pre-trained on MIMIC-III to improve the inductive inference part of the LLM (see Appendix A.2). The TSF model is used in two ways: First, we extract predictions from the forecaster and augment the prompt with those predictions in a pipeline approach. In this approach, the Llama-3 model is still finetuned using LORA, but the forecaster is kept fixed. Second, we use a multimodal setup to connect the forecaster with the Llama-3 model. We extract the full decoder prediction from the forecaster. The output vectors are then fed into a connector MLP to produce vectors with Llama-3 embedding dimensionality. These time series "token embeddings" are then prepended to the actual text embeddings. The LORA adapters in the Llama-3 model, the connector MLP, and the TSF model get updated during the finetuning process. A list of hyperparameter settings chosen on the validation set (except the default setting for the Llama-3 8B model) is given in Appendix A.3. For further comparison, we use an 8B parameters LLM pretrained on medical texts, including publications and wikipedia texts on consensus guidelines (me-llama (Xie et al., 2025)<sup>7</sup>), and 70B pretrained Llama-3 model in one-shot mode (**one-shot-70B**<sup>8</sup>).

 $<sup>^{6} \</sup>texttt{meta-llama/Llama-3.1-8B-Instruct}$ 

<sup>7</sup>https://huggingface.co/YBXL/Med-LLaMA3-8B

 $<sup>^{8}</sup>$ https://huggingface.co/meta-llama/Llama-3.1-70B-Instruct

Table 2: Derivation correctness for in- and out-of-distribution exceptions to inference rules.

variable	% changes ID	% changes OOD	ID score	OOD score
SOFA <sub>1:24</sub>	0.404	0.424	1.000	1.000
SOFA <sub>25:48</sub>	0.433	0.452	1.000	1.000

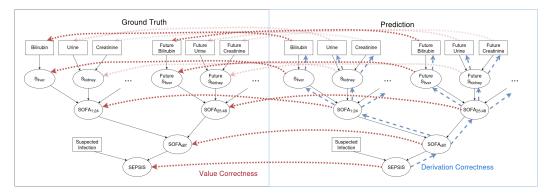


Figure 3: **Derivation correctness** (dashed blue arrows) checks for the predicted inference graph whether each child node (conclusion) follows from the parent node (premise) according to the consensus rule. **Value correctness** (dotted red arrows) maps each node in the predicted inference graph to its corresponding node in the ground truth graph, consisting of real-world clinical measurements for first and second 24 hours, and deterministic calculation of SOFA and SEPSIS on these values.

**Evaluation Metrics** We evaluate the generated output of LLMs in three different ways. First, we evaluate the LLM output according to **derivation correctness** (dashed blue arrows in Figure 3) which checks whether each child node (conclusion) in the predicted inference graph follows from the parent node (premise) according to the consensus rule. This evaluation measures the ability of the LLM to learn the deductive inference rules of the consensus definition. Furthermore, it measures faithfulness of the model's inference process in the sense of checking for an accurate representation of the reasoning process behind the model's prediction. In order to avoid trivial cases of derivation correctness, for example, inference of unchanged SOFA scores from unchanged future measurement predictions, we measure in addition forced derivation correctness. This is done by forcing the decoder to use the correct inference chain up to a last measurement token of a SOFA system as context, then continuing the decoding process until the specified target node. Furthermore, we conduct an evaluation according to value correctness (dotted red arrows in Figure 3) which compares the numerical value of each partial inference node in the prediction graph with a corresponding real-world output. Since errors earlier in the inference necessarily propagate through the inference chain, this evaluation measures the utility of the LLM output in practical applications. Comparison in all three cases is done by parsing the LLM output for relevant keywords, extracting numbers, and checking if the predicted number lies in a specific interval around the true value. A match is considered positive if the result deviates from the true value by 5%. In case of value correctness, we calculate a contingency table by comparing the application of the Sepsis-3 definition to a model involving forecasted clinical measurements with Sepsis-3 calculated for real-world clinical measurements in the second 24 hours block. From this, metrics such as accuracy, specificity, sensitivity and F1 are computed.

## 5 RESULTS AND DISCUSSION

**LLMs Shine in Deductive Inference According to Consensus Rules** Our first research question asks whether LLMs can learn the deductive inference rules underlying medical consensus definitions by fine-tuning on verbalized instantiations of the inference process to patient data. Table 1 shows that all fine-tuned models, from basic fine-tuning with text encoding to pipeline and multimodal approaches, achieve nearly perfect derivation correctness. This includes mapping clinical variables to SOFA step functions for the first 24 hours (rows 1-7), mapping forecasted clinical variables to future SOFA scores (rows 8-15), and computing a Sepsis label (row 16). In contrast, one-shot learning

Table 3: Value correctness of predicted clinical variables compared to real-world measurements. All fine-tuning improvements over one-shot learning are statistically significant according to an approximate randomization test.

variable	one-shot	one-shot-70B	me-llama	fine-tuned	pipeline	multimodal
GCS-eye	0.830	0.924	0.931	0.998	0.997	0.998
GCS-motor	0.928	0.939	0.904	0.998	0.998	0.998
GCS-verbal	0.953	0.963	0.916	0.999	0.998	0.999
MAP	0.049	0.275	0.068	0.993	0.976	0.994
Dopamine	0.870	0.951	0.913	0.972	0.964	0.968
Dobutamine	0.901	0.980	0.948	0.993	0.991	0.992
Epinephrine	0.890	0.967	0.934	0.979	0.980	0.980
Norepinephrine	0.802	0.876	0.846	0.909	0.895	0.906
Weight	0.906	0.809	0.876	0.999	0.999	0.999
PaO2/FiO2	0.262	0.540	0.124	0.998	0.997	0.997
PaO2	0.537	0.553	0.400	0.998	0.997	0.997
FiO2	0.877	0.714	0.848	1.000	1.000	1.000
Platelet	0.915	0.963	0.808	0.997	0.996	0.996
Bilirubin	0.301	0.281	0.311	0.999	0.998	0.999
Urine	0.114	0.245	0.114	0.966	0.761	0.880
Creatinine	0.896	0.968	0.790	0.998	0.998	0.998
future GCS-eye	0.684	0.701	0.653	0.712	0.743	0.789
future GCS-motor	0.777	0.800	0.759	0.811	0.832	0.869
future GCS-verbal	0.758	0.838	0.791	0.848	0.871	0.895
future MAP	0.076	0.225	0.151	0.274	0.289	0.294
future Dopamine	0.969	0.964	0.928	0.978	0.987	0.988
future Dobutamine	0.984	0.983	0.953	0.994	0.993	0.993
future Epinephrine	0.979	0.978	0.944	0.989	0.989	0.989
future Norepinephrine	0.906	0.888	0.865	0.913	0.914	0.913
future Weight	0.906	0.809	0.876	0.915	0.916	0.914
future PaO2/FiO2	0.082	0.168	0.068	0.565	0.584	0.596
future PaO2	0.139	0.187	0.174	0.616	0.659	0.715
future FiO2	0.685	0.549	0.738	0.821	0.840	0.867
future Platelet	0.216	0.261	0.266	0.364	0.383	0.396
future Bilirubin	0.159	0.169	0.201	0.811	0.832	0.866
future Urine	0.058	0.090	0.084	0.133	0.152	0.183
future Creatinine	0.299	0.305	0.344	0.395	0.424	0.455

does not perform nearly as good, even for models that include abstract consensus definitions in their training data (me-llama) or use significantly larger pre-training data (one-shot-70B). Forced derivation correctness (in brackets in Table 1) shows that predictions using correct histories are very similar to using predicted histories. Again, all one-shot learners underperform compared to fine-tuning.

**LLMs Can Learn Exceptions to Consensus Rules** The next question we ask is whether LLMs can learn exceptions to the inference rules that might conflict with the consensus definition. In our hypothetical scenario, we synthesize examples where ICD codes indicate a medical precondition related to an organ system, with the exception to the rule implying that the SOFA score for the respective organ system should be disregarded in the calculation of total SOFA and Sepsis. Table 2 shows that the addition of preconditions to the premises caused changes in 40-45% of the cases, however, the derivation correctness of total SOFA scores was 100% without losing performance on examples without preconditions. This shows that fine-tuning can indeed handle exceptions to the rule, opening the doors to possible extensions of our work to include human feedback in the loop.

The Bottleneck Lies in Inductive Inference, not Out-of-Distribution Generalization Our final evaluation compares the results of partial inference steps to real-world clinical measurements and to SOFA and Sepsis scores derived from these. As shown in the first half of Table 3, the numerical values that the LLM decoder has to extract from time series encodings are nearly perfect for fine-tuning approaches. The second half of Table 3 reveals the bottleneck of teaching medical inference rules to LLMs: Correctness of predicting values of 131 clinical variables 24 hours into the future drops from

Table 4: Value correctness of predicted SOFA and SEPSIS scores compared to derivations from real-world clinical measurements. All fine-tuning improvements over one-shot learning are significant.

variable	one-shot	one-shot-70B	me-llama	fine-tuned	pipeline	multimodal
S <sub>CNS</sub>	0.565	0.620	0.476	0.998	0.997	0.997
$S_{cardio}$	0.406	0.577	0.359	0.985	0.971	0.987
$S_{resp}$	0.262	0.397	0.183	0.996	0.994	0.996
$S_{coag}$	0.625	0.640	0.593	0.999	0.999	0.999
$S_{liver}$	0.478	0.812	0.591	1.000	0.999	1.000
$S_{renal}$	0.532	0.687	0.530	0.997	0.992	0.996
$SOFA_{1:24}$	0.132	0.228	0.095	0.979	0.958	0.981
future S <sub>CNS</sub>	0.516	0.543	0.475	0.701	0.732	0.760
future S <sub>cardio</sub>	0.439	0.525	0.406	0.700	0.738	0.745
future S <sub>resp</sub>	0.228	0.330	0.175	0.746	0.754	0.780
future S <sub>coag</sub>	0.635	0.624	0.589	0.839	0.823	0.834
future S <sub>liver</sub>	0.478	0.812	0.591	0.945	0.945	0.945
future S <sub>renal</sub>	0.477	0.585	0.467	0.553	0.451	0.473
$SOFA_{25:48}$	0.117	0.159	0.106	0.203	0.220	0.243
$SOFA_{diff}$	0.124	0.122	0.120	0.151	0.147	0.163
SEPSIS Accuracy	0.788	0.849	0.860	0.868	0.873	0.886
SEPSIS Specificity	0.823	0.899	0.910	0.936	0.920	0.922
SEPSIS Sensitivity	0.350	0.209	0.227	0.263	0.336	0.386
SEPSIS F1	0.195	0.169	0.192	0.254	0.272	0.309

the high nineties to values in the tens and twenties, especially for sparsely and irregularly observed lab values. The same is true for value correctness of predicted SOFA scores and Sepsis labels: As shown in Table 4, the numerical values produced by mapping clinical measurements in the first 24 hours to SOFA scores are nearly perfect, however, mapping values of forecasted clinical variables to step functions severely impacts correctness, especially on metrics like F1 that are sensitive to imbalanced data. One-shot learning is impacted even more, while fine-tuning results can be improved by moving from a text-based encoding to coupling dedicated TSF encoders with the LLM.

**Discussion** Leveraging consensus rules that exist for many medical areas, consistent diagnoses for unseen patients and faithful medical inference even including exceptions to the rules can be achieved by fine-tuning LLMs on instantiations of the respective consensus rule system to patient data. In contrast, training on abstract expositions of the rule system or on very large general-domain data is insufficient. Since an adaptation of an LLM to consensus rules for a specific medical area yields consistent results and faithful medical inference in that area, the bottleneck in consensus-based medical prediction lies the orthogonal problem of generalization into the future, i.e., in the inherent complexity of forecasting of sparsely and irregularly sampled clinical variables for early prediction. This problem hinders medical prognosis, and needs to be addressed by improved TSF.

#### 6 Conclusion

LLMs have been shown to excel at various health-related tasks, however, diagnostic AI has not realized the potential to actually reduce the cognitive demand and associated diagnostic errors of clinicians, due to a focus on predictive accuracy of final diagnostic labels (Adler-Milstein et al., 2021). Furthermore, instead of defining diagnostic excellence by perfect labeling based on fully established clinical criteria, Angus & Bindman (2022) advocate to learn patterns of clinical data that are yet unrecognized but still predictive of sepsis. Our approach can be seen as a first step towards these goals by putting a focus on teaching LLMs inference rules that support clinicians, and to learn exceptions to the rules that can capture new clinical patterns. Our experiments show that LLMs perfectly generalize deductive inference to unseen patients, while the bottleneck lies in inductive inference. Future work shall address the TSF bottleneck of early prediction models, and elevate our work from simulating exceptions to the rule to learning them from real-world clinical data. Furthermore, we will investigate the potential of task association learning (Cai et al., 2025) applied to related consensus definitions.

#### REPRODUCIBILITY STATEMENT

This work is based on publicly available data and open-source code. Code to reproduce our experiments is available as supplementary material and will be made public upon acceptance of the paper.

## REFERENCES

- Julia Adler-Milstein, Jonathan H. Chen, and Gurpreet Dhaliwal. Next-generation artificial intelligence for diagnosis: From predicting diagnostic labels to "wayfinding". *JAMA*, 326(24):2467–2468, 2021. URL https://doi.org/10.1001/jama.2021.22396.
- American Psychiatric Association (ed.). *Diagnostic and Statistical Manual of Mental Disorders*. Fifth edition, 2013. URL https://doi/book/10.1176/appi.books.9780890425787.
- Derek C. Angus and Andrew B. Bindman. Achieving diagnostic excellence for sepsis. *JAMA*, 327 (2):117–118, 2022. URL https://doi.org/10.1001/jama.2021.23916.
- Iván Arcuschin, Jett Janiak, Robert Krzyzanowski, Senthooran Rajamanoharan, Neel Nanda, and Arthur Conmy. Chain-of-thought reasoning in the wild is not always faithful. In *ICRL Workshop on Reasoning and Planning for Large Language Models*, 2025. URL https://openreview.net/forum?id=L8094Whth0.
- Ziyang Cai, Nayoung Lee, Avi Schwarzschild, Samet Oymak, and Dimitris Papailiopoulos. Extrapolation by association: Length generalization transfer in transformers. *arXiv*, abs/2506.09251, 2025. URL https://arxiv.org/abs/2506.09251.
- T. Cederholm, G.L. Jensen, M.I.T.D. Correia, M.C. Gonzalez, R. Fukushima, T. Higashiguchi, G. Baptista, R. Barazzoni, R. Blaauw, A. Coats, A. Crivelli, D.C. Evans, L. Gramlich, V. Fuchs-Tarlovsky, H. Keller, L. Llido, A. Malone, K.M. Mogensen, J.E. Morley, M. Muscaritoli, I. Nyulasi, M. Pirlich, V. Pisprasert, M.A.E. de van der Schueren, S. Siltharm, P. Singer, K. Tappenden, N. Velasco, D. Waitzberg, P. Yamwong, J. Yu, A. Van Gossum, C. Compher, Gordon L. Jensen, Compher Charlene, Tommy Cederholm, Andre Van Gossum, Maria Isabel T.D. Correia, M. Cristina Gonzalez, Ryoji Fukushima, Takashi Higashiguchi, G. Baptista, R. Barazzoni, R. Blaauw, A. Coats, A. Crivelli, D.C. Evans, L. Gramlich, V. Fuchs, H. Keller, L. Llido, A. Malone, K.M. Mogensen, J.E. Morley, M. Muscaritoli, I. Nyulasi, M. Pirlich, V. Pisprasert, M.A.E. de van der Schueren, S. Siltharm, P. Singer, K. Tappenden, N. Velasco, D. Waitzberg, P. Yamwong, and J. Yu. GLIM criteria for the diagnosis of malnutrition a consensus report from the global clinical nutrition community. *Clinical Nutrition*, 38(1):1–9, 2019. URL https://doi.org/10.1016/j.clnu.2018.08.002.
- Junying Chen, Zhenyang Cai, Ke Ji, Xidong Wang, Wanlong Liu, Rongsheng Wang, Jianye Hou, and Benyou Wang. Huatuogpt-o1, towards medical complex reasoning with LLMs. *arXiv*, abs/2412.18925, 2024. URL https://arxiv.org/abs/2412.18925.
- Yanda Chen, Joe Benton, Ansh Radhakrishnan, Jonathan Uesato, Carson Denison, John Schulman, Arushi Somani, Peter Hase, Misha Wagner, Fabien Roger, Vlad Mikulik, Samuel R. Bowman, Jan Leike, Jared Kaplan, and Ethan Perez. Reasoning models don't always say what they think. *arXiv*, abs/2505.05410, 2025. URL https://arxiv.org/abs/2505.05410.
- Zeming Chen, Alejandro Hernández Cano, Angelika Romanou, Antoine Bonnet, Kyle Matoba, Francesco Salvi, Matteo Pagliardini, Simin Fan, Andreas Köpf, Amirkeivan Mohtashami, Alexandre Sallinen, Alireza Sakhaeirad, Vinitra Swamy, Igor Krawczuk, Deniz Bayazit, Axel Marmet, Syrielle Montariol, Mary-Anne Hartley, Martin Jaggi, and Antoine Bosselut. Meditron-70b: Scaling medical pretraining for large language models. *arXiv*, abs/2311.16079, 2023. URL https://arxiv.org/abs/2311.16079.
- Jee-Woo Choi, Minuk Yang, Jae-Woo Kim, Yoon Mi Shin, Yong-Goo Shin, and Seung Park. Prognostic prediction of sepsis patient using transformer with skip connected token for tabular data. *Artificial Intelligence in Medicine*, 149:102804, 2024. URL https://doi.org/10.1016/j.artmed.2024.102804.

541

542

543

544

546

547

548

549 550

551

552

553

554

558

559

561

562

563 564

565

566

567

568

569

570

571

572

573

574

575

576

577

578

579

580

581

582

583

584

585

588

589

592

Tianzhe Chu, Yuexiang Zhai, Jihan Yang, Shengbang Tong, Saining Xie, Dale Schuurmans, Quoc V Le, Sergey Levine, and Yi Ma. SFT memorizes, RL generalizes: A comparative study of foundation model post-training. In *Forty-second International Conference on Machine Learning (ICML)*, 2025. URL https://openreview.net/forum?id=dYur3yabMj.

Samuel N. Cohen, James Foster, Peter Foster, Hang Lou, Terry Lyons, Sam Morley, James Morrill, Hao Ni, Edward Palmer, Bo Wang, Yue Wu, Lingyi Yang, and Weixin Yang. Subtle variation in sepsis-III definitions markely influences predictive performance within and across methods. *scientific reports*, 14(1920):1–10, 2024. URL https://doi.org/10.1038/s41598-024-51989-6.

R.P. Dellinger, M.M. Levy, A. Rhodes, D. Annane, H. Gerlach, S.M.Opal, J.E. Sevransky, C.L. Sprung, I.S. Douglas, /R. Jaeschke ant T.M. Osborn, M.E. Nunnally, S.R. Townsend, K. Reinhart, R.M. Kleinpell, D.C. Angus, C.S. Deutschman, F.R. Machado, G.D. Rubenfeld, S.A. Webb, R.J. Beale, J.L. Vincent, and R. Moreno. Surviving sepsis campaign: international guidelines for management of severe sepsis and septic shock: 2012. *Critical Care Medicine*, 41(2):580–637, 2013. URL http://dx.doi.org/10.1097/CCM.0b013e31827e83af.

Nouha Dziri, Ximing Lu, Melanie Sclar, Xiang Lorraine Li, Liwei Jiang, Bill Yuchen Lin, Sean Welleck, Peter West, Chandra Bhagavatula, Ronan Le Bras, Jena D. Hwang, Soumya Sanyal, Xiang Ren, Allyson Ettinger, Zaid Harchaoui, and Yejin Choi. Faith and fate: Limits of transformers on compositionality. In *Thirty-seventh Conference on Neural Information Processing Systems* (NeurIPS), 2023. URL https://openreview.net/forum?id=Fkckkr3ya8.

Kevin W. Eva. What every teacher needs to know about clinical reasoning. *Med Educ*, 39(1):98–106, 2005. URL https://doi.org/10.1111/j.1365-2929.2004.01972.x.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, Amy Yang, Angela Fan, Anirudh Goyal, Anthony Hartshorn, Aobo Yang, Archi Mitra, Archie Sravankumar, Artem Koreney, Arthur Hinsvark, Arun Rao, Aston Zhang, Aurelien Rodriguez, Austen Gregerson, Ava Spataru, Baptiste Roziere, Bethany Biron, Binh Tang, Bobbie Chern, Charlotte Caucheteux, Chaya Nayak, Chloe Bi, Chris Marra, Chris McConnell, Christian Keller, Christophe Touret, Chunyang Wu, Corinne Wong, Cristian Canton Ferrer, Cyrus Nikolaidis, Damien Allonsius, Daniel Song, Danielle Pintz, Danny Livshits, Danny Wyatt, David Esiobu, Dhruv Choudhary, Dhruv Mahajan, Diego Garcia-Olano, Diego Perino, Dieuwke Hupkes, Egor Lakomkin, Ehab AlBadawy, Elina Lobanova, Emily Dinan, Eric Michael Smith, Filip Radenovic, Francisco Guzmán, Frank Zhang, Gabriel Synnaeve, Gabrielle Lee, Georgia Lewis Anderson, Govind Thattai, Graeme Nail, Gregoire Mialon, Guan Pang, Guillem Cucurell, Hailey Nguyen, Hannah Korevaar, Hu Xu, Hugo Touvron, Iliyan Zarov, Imanol Arrieta Ibarra, Isabel Kloumann, Ishan Misra, Ivan Evtimov, Jack Zhang, Jade Copet, Jaewon Lee, Jan Geffert, Jana Vranes, Jason Park, Jay Mahadeokar, Jeet Shah, Jelmer van der Linde, Jennifer Billock, Jenny Hong, Jenya Lee, Jeremy Fu, Jianfeng Chi, Jianyu Huang, Jiawen Liu, Jie Wang, Jiecao Yu, Joanna Bitton, Joe Spisak, Jongsoo Park, Joseph Rocca, Joshua Johnstun, Joshua Saxe, Junteng Jia, Kalvan Vasuden Alwala, Karthik Prasad, Kartikeya Upasani, Kate Plawiak, Ke Li, Kenneth Heafield, Kevin Stone, Khalid El-Arini, Krithika Iyer, Kshitiz Malik, Kuenley Chiu, Kunal Bhalla, Kushal Lakhotia, Lauren Rantala-Yeary, Laurens van der Maaten, Lawrence Chen, Liang Tan, Liz Jenkins, Louis Martin, Lovish Madaan, Lubo Malo, Lukas Blecher, Lukas Landzaat, Luke de Oliveira, Madeline Muzzi, Mahesh Pasupuleti, Mannat Singh, Manohar Paluri, Marcin Kardas, Maria Tsimpoukelli, Mathew Oldham, Mathieu Rita, Maya Pavlova, Melanie Kambadur, Mike Lewis, Min Si, Mitesh Kumar Singh, Mona Hassan, Naman Goyal, Narjes Torabi, Nikolay Bashlykov, Nikolay Bogoychev, Niladri Chatterji, Ning Zhang, Olivier Duchenne, Onur Çelebi, Patrick Alrassy, Pengchuan Zhang, Pengwei Li, Petar Vasic, Peter Weng, Prajjwal Bhargava, Pratik Dubal, Praveen Krishnan, Punit Singh Koura, Puxin Xu, Qing He, Qingxiao Dong, Ragavan Srinivasan, Raj Ganapathy, Ramon Calderer, Ricardo Silveira Cabral, Robert Stojnic, Roberta Raileanu, Rohan Maheswari, Rohit Girdhar, Rohit Patel, Romain Sauvestre, Ronnie Polidoro, Roshan Sumbaly, Ross Taylor, Ruan Silva, Rui Hou, Rui Wang, Saghar Hosseini, Sahana Chennabasappa, Sanjay Singh, Sean Bell, Seohyun Sonia Kim, Sergey Edunov, Shaoliang Nie, Sharan Narang, Sharath Raparthy, Sheng Shen, Shengye Wan, Shruti Bhosale, Shun Zhang, Simon Vandenhende, Soumya Batra, Spencer Whitman, Sten Sootla, Stephane Collot, Suchin Gururangan, Sydney Borodinsky, Tamar Herman, Tara Fowler, Tarek Sheasha,

595

596

597

598

600

601

602

603

604

605

606

607

608

610

611

612

613

614

615

616

617

618

619

620

621

622

623

624

625

626

627

628

629

630

631

632

633

634

635

636

637

638

639

640

641

642

644

645

646

647

Thomas Georgiou, Thomas Scialom, Tobias Speckbacher, Todor Mihaylov, Tong Xiao, Ujjwal Karn, Vedanuj Goswami, Vibhor Gupta, Vignesh Ramanathan, Viktor Kerkez, Vincent Gonguet, Virginie Do, Vish Vogeti, Vítor Albiero, Vladan Petrovic, Weiwei Chu, Wenhan Xiong, Wenyin Fu, Whitney Meers, Xavier Martinet, Xiaodong Wang, Xiaofang Wang, Xiaoqing Ellen Tan, Xide Xia, Xinfeng Xie, Xuchao Jia, Xuewei Wang, Yaelle Goldschlag, Yashesh Gaur, Yasmine Babaei, Yi Wen, Yiwen Song, Yuchen Zhang, Yue Li, Yuning Mao, Zacharie Delpierre Coudert, Zheng Yan, Zhengxing Chen, Zoe Papakipos, Aaditya Singh, Aayushi Srivastava, Abha Jain, Adam Kelsey, Adam Shajnfeld, Adithya Gangidi, Adolfo Victoria, Ahuva Goldstand, Ajay Menon, Ajay Sharma, Alex Boesenberg, Alexei Baevski, Allie Feinstein, Amanda Kallet, Amit Sangani, Amos Teo, Anam Yunus, Andrei Lupu, Andres Alvarado, Andrew Caples, Andrew Gu, Andrew Ho, Andrew Poulton, Andrew Ryan, Ankit Ramchandani, Annie Dong, Annie Franco, Anuj Goyal, Aparajita Saraf, Arkabandhu Chowdhury, Ashley Gabriel, Ashwin Bharambe, Assaf Eisenman, Azadeh Yazdan, Beau James, Ben Maurer, Benjamin Leonhardi, Bernie Huang, Beth Loyd, Beto De Paola, Bhargavi Paranjape, Bing Liu, Bo Wu, Boyu Ni, Braden Hancock, Bram Wasti, Brandon Spence, Brani Stojkovic, Brian Gamido, Britt Montalvo, Carl Parker, Carly Burton, Catalina Mejia, Ce Liu, Changhan Wang, Changkyu Kim, Chao Zhou, Chester Hu, Ching-Hsiang Chu, Chris Cai, Chris Tindal, Christoph Feichtenhofer, Cynthia Gao, Damon Civin, Dana Beaty, Daniel Kreymer, Daniel Li, David Adkins, David Xu, Davide Testuggine, Delia David, Devi Parikh, Diana Liskovich, Didem Foss, Dingkang Wang, Duc Le, Dustin Holland, Edward Dowling, Eissa Jamil, Elaine Montgomery, Eleonora Presani, Emily Hahn, Emily Wood, Eric-Tuan Le, Erik Brinkman, Esteban Arcaute, Evan Dunbar, Evan Smothers, Fei Sun, Felix Kreuk, Feng Tian, Filippos Kokkinos, Firat Ozgenel, Francesco Caggioni, Frank Kanayet, Frank Seide, Gabriela Medina Florez, Gabriella Schwarz, Gada Badeer, Georgia Swee, Gil Halpern, Grant Herman, Grigory Sizov, Guangyi, Zhang, Guna Lakshminarayanan, Hakan Inan, Hamid Shojanazeri, Han Zou, Hannah Wang, Hanwen Zha, Haroun Habeeb, Harrison Rudolph, Helen Suk, Henry Aspegren, Hunter Goldman, Hongyuan Zhan, Ibrahim Damlaj, Igor Molybog, Igor Tufanov, Ilias Leontiadis, Irina-Elena Veliche, Itai Gat, Jake Weissman, James Geboski, James Kohli, Janice Lam, Japhet Asher, Jean-Baptiste Gaya, Jeff Marcus, Jeff Tang, Jennifer Chan, Jenny Zhen, Jeremy Reizenstein, Jeremy Teboul, Jessica Zhong, Jian Jin, Jingyi Yang, Joe Cummings, Jon Carvill, Jon Shepard, Jonathan McPhie, Jonathan Torres, Josh Ginsburg, Junjie Wang, Kai Wu, Kam Hou U, Karan Saxena, Kartikay Khandelwal, Katayoun Zand, Kathy Matosich, Kaushik Veeraraghavan, Kelly Michelena, Keqian Li, Kiran Jagadeesh, Kun Huang, Kunal Chawla, Kyle Huang, Lailin Chen, Lakshya Garg, Lavender A, Leandro Silva, Lee Bell, Lei Zhang, Liangpeng Guo, Licheng Yu, Liron Moshkovich, Luca Wehrstedt, Madian Khabsa, Manav Avalani, Manish Bhatt, Martynas Mankus, Matan Hasson, Matthew Lennie, Matthias Reso, Maxim Groshev, Maxim Naumov, Maya Lathi, Meghan Keneally, Miao Liu, Michael L. Seltzer, Michael Valko, Michelle Restrepo, Mihir Patel, Mik Vyatskov, Mikayel Samvelyan, Mike Clark, Mike Macey, Mike Wang, Miquel Jubert Hermoso, Mo Metanat, Mohammad Rastegari, Munish Bansal, Nandhini Santhanam, Natascha Parks, Natasha White, Navyata Bawa, Nayan Singhal, Nick Egebo, Nicolas Usunier, Nikhil Mehta, Nikolay Pavlovich Laptev, Ning Dong, Norman Cheng, Oleg Chernoguz, Olivia Hart, Omkar Salpekar, Ozlem Kalinli, Parkin Kent, Parth Parekh, Paul Saab, Pavan Balaji, Pedro Rittner, Philip Bontrager, Pierre Roux, Piotr Dollar, Polina Zvyagina, Prashant Ratanchandani, Pritish Yuvraj, Qian Liang, Rachad Alao, Rachel Rodriguez, Rafi Ayub, Raghotham Murthy, Raghu Nayani, Rahul Mitra, Rangaprabhu Parthasarathy, Raymond Li, Rebekkah Hogan, Robin Battey, Rocky Wang, Russ Howes, Ruty Rinott, Sachin Mehta, Sachin Siby, Sai Jayesh Bondu, Samyak Datta, Sara Chugh, Sara Hunt, Sargun Dhillon, Sasha Sidorov, Satadru Pan, Saurabh Mahajan, Saurabh Verma, Seiji Yamamoto, Sharadh Ramaswamy, Shaun Lindsay, Shaun Lindsay, Sheng Feng, Shenghao Lin, Shengxin Cindy Zha, Shishir Patil, Shiva Shankar, Shuqiang Zhang, Shuqiang Zhang, Sinong Wang, Sneha Agarwal, Soji Sajuvigbe, Soumith Chintala, Stephanie Max, Stephen Chen, Steve Kehoe, Steve Satterfield, Sudarshan Govindaprasad, Sumit Gupta, Summer Deng, Sungmin Cho, Sunny Virk, Suraj Subramanian, Sy Choudhury, Sydney Goldman, Tal Remez, Tamar Glaser, Tamara Best, Thilo Koehler, Thomas Robinson, Tianhe Li, Tianjun Zhang, Tim Matthews, Timothy Chou, Tzook Shaked, Varun Vontimitta, Victoria Ajayi, Victoria Montanez, Vijai Mohan, Vinay Satish Kumar, Vishal Mangla, Vlad Ionescu, Vlad Poenaru, Vlad Tiberiu Mihailescu, Vladimir Ivanov, Wei Li, Wenchen Wang, Wenwen Jiang, Wes Bouaziz, Will Constable, Xiaocheng Tang, Xiaojian Wu, Xiaolan Wang, Xilun Wu, Xinbo Gao, Yaniv Kleinman, Yanjun Chen, Ye Hu, Ye Jia, Ye Qi, Yenda Li, Yilin Zhang, Ying Zhang, Yossi Adi, Youngjin Nam, Yu, Wang, Yu Zhao, Yuchen Hao, Yundi Qian, Yunlu Li, Yuzi He, Zach Rait, Zachary DeVito, Zef Rosnbrick, Zhaoduo Wen, Zhenyu Yang, Zhiwei Zhao, and Zhiyu Ma. The Llama 3 herd of models. arXiv, abs/2407.21783, 2024.

```
URL https://arxiv.org/abs/2407.21783.
```

- Paul Hager, Friederike Jungmann, Robbie Holland, Kunal Bhagat, Inga Hubrecht, Manuel Knauer, Jakob Vielhauer, Marcus Makowski, Rickmer Braren, Georgios Kaissis, and Daniel Rueckert. Evaluation and mitigation of the limitations of large language models in clinical decision-making. *Nature Medicine*, 30, 2024. URL https://doi.org/10.1038/s41591-024-03097-1.
- Andreas Hochlehnert, Hardik Bhatnagar, Vishaal Udandarao, Samuel Albanie, Ameya Prabhu, and Matthias Bethge. A sober look at progress in language model reasoning: Pitfalls and paths to reproducibility. In *Second Conference on Language Modeling (COLM)*, 2025. URL https://openreview.net/forum?id=90UrTTxp50.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations (ICLR)*, virtual, 2022. URL https://openreview.net/forum?id=nZeVKeeFYf9.
- Alon Jacovi and Yoav Goldberg. Towards faithfully interpretable NLP systems: How should we define and evaluate faithfulness? In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL)*, Online, 2020. URL https://aclanthology.org/2020.acl-main.386/.
- Alistair E. W. Johnson, Lucas Bulgarelli, Lu Shen, Alvin Gayles, Ayad Shammout, Steven Horng, Tom J. Pollard, Sicheng Hao, Benjamin Moody, Brian Gow, Li-wei H. Lehman, Leo A. Celi, and Roger G. Mark. MIMIC-IV, a freely accessible electronic health record dataset. *Scientific Data*, 10(1), 2023. URL https://doi.org/10.1038/s41597-022-01899-x.
- Alistair E.W. 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. MIMIC-III, a freely accessible critical care database. *Scientific Data*, 3(1):160035, 2016. URL https://doi.org/10.1038/sdata.2016.35.
- KDIGO Acute Kidney Injury Work Group. Kidney disease: Improving global outcomes (KDIGO) Cinical practice guideline for acute kidney injury. *Kidney international supplements*, 2(1):1–138, 2012. URL doi:10.1038/kisup.2012.1.
- Tamera Lanham, Anna Chen, Ansh Radhakrishnan, Benoit Steiner, Carson Denison, Danny Hernandez, Dustin Li, Esin Durmus, Evan Hubinger, Jackson Kernion, Kamilė Lukošiūtė, Karina Nguyen, Newton Cheng, Nicholas Joseph, Nicholas Schiefer, Oliver Rausch, Robin Larson, Sam McCandlish, Sandipan Kundu, Saurav Kadavath, Shannon Yang, Thomas Henighan, Timothy Maxwell, Timothy Telleen-Lawton, Tristan Hume, Zac Hatfield-Dodds, Jared Kaplan, Jan Brauner, Samuel R. Bowman, and Ethan Perez. Measuring faithfulness in chain-of-thought reasoning. *arXiv*, abs/2307.13702, 2023. URL https://arxiv.org/abs/2307.13702.
- Mitchell M. Levy, Mitchell P. Fink, John C. Marshall, Edward Abraham, Derek Angus, Deborah Cook, Jonathan Cohen, Steven M. Opal, Jean-Louis Vincent, and Graham Ramsay. 2001 SCCM/ESICM/ACCP/ATS/SIS International Sepsis Definitions Conference. *Critical Care Medicine*, 31(04):1250–1256, 2003. URL https://doi.org/10.1097/01.ccm.0000050454.01978.3b.
- Qiang Li, Hanbo Ma, Dan Song, Yunpeng Bai, Lina Zhao, and Keliang Xie. Early prediction of sepsis using chatgpt-generated summaries and structured data. *Multimedia Tools and Applications*, pp. 1–23, 2024. URL https://doi.org/10.1007/s11042-024-18378-7.
- Xiaohong Liu, Hao Liu, Guoxing Yang, Zeyu Jiang, Shuguang Cui, Zhaoze Zhang, Huan Wang, Liyuan Tao, Yongchang Sun, Zhu Song, Tianpei Hong, Jin Yang, Tianrun Gao, Jiangjiang Zhang, Xiaohu Li, Jing Zhang, Ye Sang, Zhao Yang, Kanmin Xue, Song Wu, Ping Zhang, Jian Yang, Chunli Song, and Guangyu Wang. A generalist medical language model for disease diagnosis assistance. *Nature Medicine*, 31:932–942, 2025. URL https://doi.org/10.1038/s41591-024-03416-6.

- Xin Liu, Daniel McDuff, Geza Kovacs, Isaac Galatzer-Levy, Jacob Sunshine, Jiening Zhan, Ming-Zher Poh, Shun Liao, Paolo Di Achille, and Shwetak Patel. Large language models are few-shot health learners. *arXiv*, abs/2305.15525, 2023. URL https://doi.org/10.48550/arXiv.2305.15525.
  - Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. Faithful chain-of-thought reasoning. In *Proceedings of the 13th International Joint Conference on Natural Language Processing and the 3rd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics (IJCNLP)*, Nusa Dua, Bali, 2023. URL 10.18653/v1/2023.ijcnlp-main.20.
  - W. I. McDonald, A. Compston, G. Edan, D. Goodkin, H. P. Hartung, F. D. Lublin, H. F. McFarland, D. W. Paty, C. H. Polman, S. C. Reingold, M. Sandberg-Wollheim, W. Sibley, A. Thompson, S. van den Noort, B. Y. Weinshenker, and J. S. Wolinsky. Recommended diagnostic criteria for multiple sclerosis: guidelines from the international panel on the diagnosis of multiple sclerosis. *Annals of Neurology*, 50(1):121–127, 2001. doi: https://doi.org/10.1002/ana.1032.
  - Mike A Merrill, Mingtian Tan, Vinayak Gupta, Thomas Hartvigsen, and Tim Althoff. Language models still struggle to zero-shot reason about time series. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, Miami, Florida, USA, 2024. URL 10.18653/v1/2024.findings-emnlp.201.
  - Philipp Mondorf and Barbara Plank. Beyond accuracy: Evaluating the reasoning behavior of large language models a survey. In *First Conference on Language Modeling (COLM)*, 2024. URL https://openreview.net/forum?id=Lmjgl2n11u.
  - Michael Moor, Bastian Rieck, Max Horn, Catherine Jutzeler, and Karsten Borgwardt. Early prediction of sepsis in the ICU using machine learning: A systematic review. *Frontiers in Medicine*, 8, 2021. URL https://doi.org/10.3389/fmed.2021.607952.
  - Shamim Nemati, Andre Holder, Fereshteh Razmi, Matthew D. Stanley, Gari D. Clifford, and Timothy G. Buchman. An interpretable machine learning model for accurate prediction of sepsis in the ICU. *Critical Care Medicine*, 46(4):547–553, 2018. URL http://dx.doi.org/10.1097/ccm.00000000000002936.
  - Geoffrey Norman. Research in clinical reasoning: past history and current trends. *Med Educ*, 39(4): 418–27, 2005. URL https://doi.org/10.1111/j.1365-2929.2005.02127.x.
  - Maxwell I. Nye, Anders Johan Andreassen, Guy Gur-Ari, Henryk Michalewski, Jacob Austin, David Bieber, David Dohan, Aitor Lewkowycz, Maarten Bosma, David Luan, Charles Sutton, and Augustus Odena. Show your work: Scratchpads for intermediate computation with language models. *arXiv*, abs/2112.00114, 2021. URL https://arxiv.org/abs/2112.00114.
  - Tom J. Pollard, Alistair E. W. Johnson, Jesse D. Raffa, Leo A. Celi, Roger G. Mark, and Omar Badawi. The eICU collaborative research database, a freely available multi-center database for critical care research. *Scientific Data*, 5(180178), 2018. URL https://doi.org/10.1038/sdata.2018.178.
  - Justin T. Reese, Daniel Danis, J. Harry Caufield, Tudor Groza, Elena Casiraghi, Giorgio Valentini, Christopher J. Mungall, and Peter N. Robinson. On the limitations of large language models in clinical diagnosis. *medRxiv*, 2024. URL https://doi.org/10.1101/2023.07.13.23292613.
- Matthew A. Reyna, Christopher S. Josef, Russell Jeter, Supreeth P.B. Shashikumar, M. Brandon Westover, Shamim Nemati, Gari D. Clifford, and Ashish Sharma. Early prediction of sepsis from clinical data: The physionet/computing in cardiology challenge 2019. *Critical Care Medicine*, 48 (2):210–217, 2019. URL https://doi.org/10.1097/CCM.000000000004145.
- David L Sackett, William M C Rosenberg, J A Muir Gray, R Brian Haynes, and W Scott Richardson. Evidence based medicine: what it is and what it isn't. *BMJ*, 312(7023):71–72, 1996. URL https://www.bmj.com/content/312/7023/71.

- Abulhair Saparov and He He. Language models are greedy reasoners: A systematic formal analysis of chain-of-thought. In *The Eleventh International Conference on Learning Representations (ICLR)*, 2023. URL https://openreview.net/forum?id=qFVVBzXxR2V.
- Christopher W. Seymour, Vincent X. Liu, Theodore J. Iwashyna, Frank M. Brunkhorst, Thomas D. Rea, André Scherag, Gordon Rubenfeld, Jeremy M. Kahn, Manu Shankar-Hari, Mervyn Singer, Clifford S. Deutschman, Gabriel J. Escobar, and Derek C. Angus. Assessment of clinical criteria for sepsis for the third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA*, 315(8):762–774, 2016. URL https://doi.org/10.1001/jama.2016.0288.
- Mervyn Singer, Clifford S. Deutschman, and Christopher Warren Seymour. The third international consensus definitions for sepsis and septic shock (Sepsis-3). *JAMA*, 315(8):801–810, 2016. URL 10.1001/jama.2016.0287.
- Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R. Pfohl, Heather Cole-Lewis, Darlene Neal, Qazi Mamunur Rashid, Mike Schaekermann, Amy Wang, Dev Dash, Jonathan H. Chen, Nigam H. Shah, Sami Lachgar, Philip Andrew Mansfield, Sushant Prakash, Bradley Green, Ewa Dominowska, Blaise Agüera y Arcas, Nenad Tomašev, Yun Liu, Renee Wong, Christopher Semturs, S. Sara Mahdavi, Joelle K. Barral, Dale R. Webster, Greg S. Corrado, Yossi Matias, Shekoofeh Azizi, Alan Karthikesalingam, and Vivek Natarajan. Toward expert-level medical question answering with large language models. *Nature Medicine*, 2025. URL https://doi.org/10.1038/s41591-024-03423-7.
- Michael Staniek, Marius Fracarolli, Michael Hagmann, and Stefan Riezler. Early prediction of causes (not effects) in healthcare by long-term clinical time series forecasting. *Proceedings of Machine Learning Research: Machine Learning for Healthcare*, 252:1–29, 2024. URL https://arxiv.org/abs/2408.03816.
- Mingtian Tan, Mike A Merrill, Vinayak Gupta, Tim Althoff, and Thomas Hartvigsen. Are language models actually useful for time series forecasting? In *The Thirty-eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024. URL https://openreview.net/forum?id=DV15UbHCY1.
- Rahul Thapa, Qingyang Wu, Kevin Wu, Harrison Zhang, Angela Zhang, Eric Wu, Haotian Ye, Suhana Bedi, Nevin Aresh, Joseph Boen, Shriya Reddy, Ben Athiwaratkun, Shuaiwen Leon Song, and James Zou. Disentangling reasoning and knowledge in medical large language models. *arXiv*, abs/2505.11462, 2025. URL https://arxiv.org/abs/2505.11462.
- Miles Turpin, Julian Michael, Ethan Perez, and Samuel R. Bowman. Language models don't always say what they think: Unfaithful explanations in chain-of-thought prompting. In *Thirty-seventh Conference on Neural Information Processing Systems (NeurIPS)*, 2023. URL https://openreview.net/forum?id=bzs4uPLXvi.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS)*, Long Beach, CA, 2017. URL https://proceedings.neurips.cc/paper\_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf.
- J.L. Vincent, R. Moreno, J. Takala, S. Willatts, A. De Mendonça, H. Bruining, C. Reinhart, P. Suter, and L. Thijs. The SOFA (Sepsis-related Organ Failure Assessment) score to describe organ dysfunction/failure. *Intensive Care Medicine*, 22(7):707–710, 1996. URL https://doi.org/10.1007/BF01709751.
- Zezhong Wang, Xingshan Zeng, Weiwen Liu, Yufei Wang, Liangyou Li, Yasheng Wang, Lifeng Shang, Xin Jiang, Qun Liu, and Kam-Fai Wong. Chain-of-probe: Examining the necessity and accuracy of CoT step-by-step. In *Findings of the Association for Computational Linguistics:* NAACL 2025, Albuquerque, New Mexico, 2025. URL https://aclanthology.org/2025.findings-naacl.140/.
- Juncheng Wu, Wenlong Deng, Xingxuan Li, Sheng Liu, Taomian Mi, Yifan Peng, Ziyang Xu, Yi Liu, Hyunjin Cho, Chang-In Choi, Yihan Cao, Hui Ren, Xiang Li, Xiaoxiao Li, and Yuyin Zhou.

- Medreason: Eliciting factual medical reasoning steps in llms via knowledge graphs. *arXiv*, 2025. URL https://arxiv.org/abs/2504.00993.
  - Qianqian Xie, Qingyu Chen, Aokun Chen, Cheng Peng, Yan Hu, Fongci Lin, Xueqing Peng, Jimin Huang, Jeffrey Zhang, Vipina Keloth, Xinyu Zhou, Lingfei Qian, Huan He, Dennis Shung, Lucila Ohno-Machado, Yonghui Wu, Hua Xu, and Jiang Bian. Medical foundation large language models for comprehensive text analysis and beyond. *npj Digital Medicine*, 8(1), 2025. URL https://doi.org/10.1038/s41746-025-01533-1.
  - Jundong Xu, Hao Fei, Liangming Pan, Qian Liu, Mong-Li Lee, and Wynne Hsu. Faithful logical reasoning via symbolic chain-of-thought. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (ACL)*, Bangkok, Thailand, 2024. URL 10.18653/v1/2024.acl-long.720.
  - Haotong Yang, Yi Hu, Shijia Kang, Zhouchen Lin, and Muhan Zhang. Number cookbook: Number understanding of language models and how to improve it. *arXiv*, abs/2411.03766, 2024. URL https://arxiv.org/abs/2411.03766.
  - Changchang Yin, Shihan Fu, Bingsheng Yao, Thai-Hoang Pham, Weidan Cao, Dakuo Wang, Jeffrey Caterino, and Ping Zhang. Sepsiscalc: Integrating clinical calculators into early sepsis prediction via dynamic temporal graph construction. *arXiv*, 2025. URL https://arxiv.org/abs/2501.00190.
  - Jaehoon Yun, Jiwoong Sohn, Jungwoo Park, Hyunjae Kim, Xiangru Tang, Yanjun Shao, Yonghoe Koo, Minhyeok Ko, Qingyu Chen, Mark Gerstein, Michael Moor, and Jaewoo Kang. Medprm: Medical reasoning models with stepwise, guideline-verified process rewards, 2025. URL https://arxiv.org/abs/2506.11474.
  - Honghua Zhang, Liunian Harold Li, Tao Meng, Kai-Wei Chang, and Guy Van den Broeck. On the paradox of learning to reason from data. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 3365–3373, 2023. URL https://doi.org/10.24963/ijcai.2023/375.
  - Sheng Zhang, Qianchu Liu, Guanghui Qin, Tristan Naumann, and Hoifung Poon. Med-RLVR: Emerging medical reasoning from a 3B base model via reinforcement learning. *arXiv*, abs/2502.19655, 2025. URL https://arxiv.org/abs/2502.19655.

# A APPENDIX

## A.1 SOFA SCORE

Table 5: SOFA score for six organ systems, calculated by thresholding corresponding clinical measurements (Vincent et al., 1996). In our setting, we had to recalculate MAP (mean arterial pressure) from SBP and DBP (systolic and diastolic blood pressure), the Horowitz coefficient from PaO2 and FiO2, and had no knowledge about the kind of mechanical ventilation. If no value for calculation in a SOFA subsystem was available, we took a value of 0. Abbreviations: CNS = Central nervous system; GCS = Glasgow Coma Scale; MV = mechanically ventilated including CPAP; MAP = mean arterial pressure, UO = Urine output.

	CNS	Cardiovascular	Respiratory	Coagulation	n Liver	Renal
Score	GCS	MAP or vasopressors	PaO2/FiO2 (mmHg)	Platelets $(\times 10^3/\mu l)$	Bilirubin (mg/dl)	Creatinine (mg/dl) or UO
+0	15	$MAP \ge 70 \text{ mmHg}$	≥ 400	≥ 150	< 1.2	< 1.2
+1	13-14	MAP < 70  mmHg	< 400	< 150	1.2-1.9	1.2-1.9
+2	10-12	dopamine $\leq 5 \mu \text{g/kg/min}$	< 300	< 100	2.0-5.9	2.0-3.4
		OR dobutamine (any dose)				
+3	6-9	dopamine $> 5 \mu g/kg/min$	< 200	< 50	6.0-11.9	3.5-4.9
		OR epinephrine $\leq 0.1 \mu\text{g/kg/min}$	AND MV			OR < 500  ml/day
		OR norepinephrine $\leq 0.1$				·
		μg/kg/min				
+4	< 6	dopamine $> 15 \mu g/kg/min$	< 100	< 20	> 12.0	> 5.0
		OR epinephrine $> 0.1 \mu\text{g/kg/min}$	AND MV			OR < 200  ml/day
		OR norepinephrine > 0.1				·
		$\mu$ g/kg/min				

# A.2 TIME SERIES FORECASTING

Our TSF model uses the implementation of Staniek et al. (2024) which is based on a Transformer encoder-decoder architecture (Vaswani et al., 2017). First, sparse multivariate input time series are represented as quadruplets  $S = \{(f_i, t_i, v_i, n_i)\}_{i=1}^n$ , where  $f_i \in F$  is a clinical variable identifier,  $t_i \in \mathbb{R}_{\geq 0}$  is a time index,  $v_i \in \mathbb{R}$  the observed value of  $f_i$  at  $t_i$ , and  $n_i$  the unique stay identifier. Then the quadruplets for a 24 hour time series are encoded into a dense representation x where every timestep is a vector of feature values representing one hour. We construct this vector by choosing the first observed value during the represented hour for each feature. If no value was observed, we impute zero which corresponds to the mean value due to standardization of the data. Additionally, a mask indicating whether a value was imputed is generated and appended to the vector. For TSF, we use a Transformer model with an autoregressive iterative multistep (IMS) decoder that generates an output vector  $\hat{y}_t \in \mathbb{R}^{|F|}$ . The predicted output  $\hat{y}_t$  is a function of the history  $\hat{y}_{< t}$  of predicted timesteps until time t, the encoded input x, and the model parameters  $\theta$ :  $\hat{y}_t = f_{\theta}(\hat{y}_{< t}, x)$ . To perform long-term TSF using the autoregressive setup, the outputs  $\hat{y}_t$  from each time step  $t = 1, \ldots, T$  are concatenated. The complete model is trained with masked mean squared error (MSE).

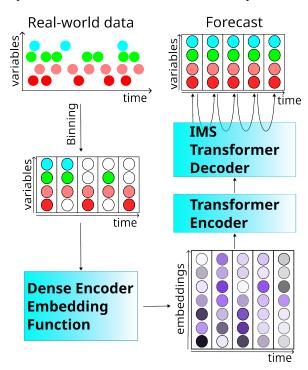


Figure 4: Time Series Forecasting using a dense encoder and iterative multistep decoder architecture.

# A.3 HYPERPARAMETERS AND COMPUTE INFRASTRUCTURE

Table 6: Hyperparameter settings of the time series forecaster.

Hyperparameter	value
Embedding Size	512
Hidden size encoder	512
Hidden size decoder	512
# Encoder layers	2
# Decoder layers	1
learning rate	0.0005
attention heads encoder	8
attention heads decoder	1
dropout	0.05
max epochs	100
patience (early stopping)	6
Random Seed	unixtime variation
# GPUs	1
Training all clock time	5 hours
GPU type	Nvidia GTX 1080 Ti

Table 7: Hyperparameter settings of the multi-modal and LLM finetuning architectures.

Hyperparameter	value
optimizer	adam
learning rate	0.000002
max epochs	10
connector first layer	131x2096
connector second layer	2096x4096
lora-r	16
lora-alpha	16
lora-dropout	0.1
lora-target-modules	all-linear
# GPUs	1
Training all clock time	40 hours
GPU type	NVIDIA A100
Starting model	meta-llama/Llama-3.1-8B-Instruct

## A.4 CLINICAL FEATURES

Table 8: Feature list for MIMIC-III: Besides the following 131 dynamic variables, only age and gender were extracted. The 15 variables marked with an asterisk are directly used for calculating the SOFA score.

1031	SOFA score.			
1032	ALP	Epinephrine*	LDH	Packed RBC
1033	ALT	Famotidine	Lactate	Pantoprazole
1034	AST	Fentanyl	Lactated Ringers	Phosphate
	Albumin	FiO2*	Levofloxacin	Piggyback
1035	Albumin 25%	Fiber	Lorazepam	Piperacillin
1036	Albumin 5%	Free Water	Lymphocytes	Platelet Count*
1037	Amiodarone	Fresh Frozen Plasma	Lymphocytes (Absolute)	Potassium
1038	Anion Gap	Furosemide	MBP	Pre-admission Intake
1039	BUN	GCS_eye*	MCH	Pre-admission Output
1040	Base Excess	GCS_motor*	MCHC	Propofol
	Basophils	GCS_verbal*	MCV	RBC
1041	Bicarbonate	GT Flush	Magnesium	RDW
1042	Bilirubin (Direct)	Gastric	Magnesium Sulfate (Bolus)	RR
1043	Bilirubin (Indirect)	Gastric Meds	Magnesium Sulphate	Residual
1044	Bilirubin (Total)*	Glucose (Blood)	Mechanically ventilated	SBP*
1045	CRR	Glucose (Serum)	Metoprolol	SG Urine
1046	Calcium Free	Glucose (Whole Blood)	Midazolam	Sodium
	Calcium Gluconate	HR	Milrinone	Solution
1047	Calcium Total	Half Normal Saline	Monocytes	Sterile Water
1048	Cefazolin	Hct	Morphine Sulfate	Stool
1049	Chest Tube	Heparin	Neosynephrine	TPN
050	Chloride	Hgb	Neutrophils	Temperature
1051	Colloid	Hydralazine	Nitroglycerine	Total CO2
1052	Creatinine Blood*	Hydromorphone	Nitroprusside	Ultrafiltrate
	Creatinine Urine	INR	Norepinephrine*	Urine*
1053	D5W	Insulin Humalog	Normal Saline	Vancomycin
1054	DBP*	Insulin NPH	O2 Saturation	Vasopressin
1055	Dextrose Other	Insulin Regular	OR/PACU Crystalloid	WBC
1056	Dobutamine*	Insulin largine	PCO2	Weight
057	Dopamine*	Intubated	PO intake	pH Blood
	EBL	Jackson-Pratt	PaO2*	pH Urine
1058	Emesis	KCl	PT	
1059	Eoisinophils	KCl (Bolus)	PTT	

1080 A.5 ONE-SHOT PROMPT 1082 Please answer in the following style, where current time values are the correct summary values of the given measurements according to the Sepsis-3 definition. The future values should be your best guesses on how the values develop. The SOFA scores should be calculated 1083 according to the Sepsis-3 definition. Only answer like in the given example. Here is the example: 1084 Patient is 63.0 years old and is male. Given all the information in this text, answer the question at the end. Here are the measurements: DBP at time -23.68: 36.0, SBP at time -23.68: 71.0... Now answer the following question in the given format: 1086 Doctors suspect an infection, based on this information and the other information in this text, will the patient be classified as septic tomorrow? First we need to calculate the SOFA scores given the extracted values. The SOFA scores for the current time are the following: 1087 The minimum value of GCS\_eye is 4.0, GCS\_motor is 6.0 and GCS\_verbal is 5.0, this produces the sum 15.0 and means the CNS SOFA is 0. 1088 Because minimum MAP is 43.333, max Dopamine is 0, max Dobutamine is 0, max Epinephrine is 0 and max Norepinephrine is 0 with a patient weight of 80 kg, the cardiovascular SOFA is 1. 1089 Given that minimum PO2 is 141.0 and minimum FiO2 is 1 the calculated PAO2FIO2 is 141.0, this means the respiratory SOFA is 3. 1090 Because the minimum Platelet count is 235.0 the coagulation SOFA is 0. The maximum Bilirubin (Total) is 1.8 leading to a liver SOFA of 1. 1091 Because total Urine output is 1585.0 and maximum creatinine in the blood is 1.4 the renal SOFA is 1. 1092 To summarize: the patient has a total SOFA score of 6. Now we need to calculate the SOFA scores with forecasted values. The SOFA scores in the future based on the forecasted values are the 1093 1094 The minimum value of GCS\_eye will be 4.0, GCS\_motor will be 6.0 and GCS\_verbal will be 5.0, this produces the sum 15.0 and means the CNS SOFA will be 0. 1095 Because future minimum MAP will be 55.667, future max Dopamine will be 0, future max Dobutamine will be 0, future max Epinephrine will be 0 and future max Norepinephrine will be 0 with a patient weight of 80 kg, the cardiovascular SOFA will be 1. Given that minimum PO2 will be 141.0 and minimum FiO2 will be 1 the forecasted PAO2FIO2 will be 141.0, this means the respiratory SOFA will be 3. 1098 Because the Platelet count will be 295.0 the coagulation SOFA is going to be 0. The maximum Bilirubin (Total) will be 1.8 leading to a liver SOFA of 1. 1099 Because Urine output will be 1635.0 and maximum creatinine in the blood will be 1.1 the renal SOFA will be 0. 1100 To summarize: the patient will have a future total SOFA score of 5. The patient will not develop sepsis in the next 24 hours, because total SOFA increased only by -1 and infection is suspected. 1101 The example is now finished. Say "The patient will develop sepsis" in the last sentence if the criteria are met (if total SOFA changed by 2 and 1102 infection is suspected). Patient is 76.0 years old and is female. Given all the information in this text, answer the question at the end. 1103 Here are the measurements: DBP at time -22.97: 56.0, GCS\_eye at time -22.97: 4.0, GCS\_motor at time -22.97: 6.0... 1104 Now answer the following question in the given format: The doctors don't suspect an infection, based on this information and the other information in this text, will the patient be classified as septic 1105 tomorrow? 1106 1107 1108 1109 1110 1111 1112 1113 1114 1115 1116 1117 1118 1119 1120 1121 1122 1123 1124 1125 1126 1127 1128 1129 1130

## A.6 SETUP FOR LEARNING FROM PRECONDITIONS

Table 9: Medical preconditions for five organ systems indicated by ICD-10 codes. In-distribution (ID) data were seen, out-of-distribution (OOD) data were unseen during fine-tuning.

	lung	kidney	coagulation	liver	cardiovascular
ID	J40, J41, J42	N18.9, N28	D68.4, D68.5	K70.0, K70.41	I50.0, I50.9
OOD	J44.9	N19	D68.6	K70.3	I50.1

Table 10: ICD-10 Code frequency for in- and out-of-distribution testsets.

- To Code frequen	•	
ICD-10 Code	# ID Test	# OOD Test
J40	263	194
J41	239	173
J42	253	194
J44.9	-	166
N18.9	266	172
N19	-	204
N28	281	151
D68.4	242	171
D68.5	244	162
D68.6	-	176
I50.0	264	185
I50.1	-	193
I50.9	229	162
K70.0	230	165
K70.3	-	162
K70.41	215	182
No-Preexisting	274	189