

Health-LLM: Large Language Models for Health Prediction via Wearable Sensor Data

Yubin Kim
MIT, Cambridge, MA

YBKIM95@MIT.EDU

Xuhai Xu
MIT, Cambridge, MA

XOXU@MIT.EDU

Daniel McDuff
Google Research, WA

DMCDUFF@GOOGLE.COM

Cynthia Breazeal
MIT, Cambridge, MA

CYNTHIAB@MIT.EDU

Hae Won Park
MIT, Cambridge, MA

HAEWON@MIT.EDU

Abstract

Large language models (LLMs) are capable of many natural language tasks, yet they are far from perfect. In health applications, grounding and interpreting domain-specific and non-linguistic data is crucial. This paper investigates the capacity of LLMs to make inferences about health based on contextual information (e.g. user demographics, health knowledge) and physiological data (e.g. resting heart rate, sleep minutes). We present a comprehensive evaluation of 12 state-of-the-art LLMs with prompting and fine-tuning techniques on four public health datasets (PMData, LifeSnaps, GLOBEM and AW_FB). Our experiments cover 10 consumer health prediction tasks in mental health, activity, metabolic, and sleep assessment. Our fine-tuned model, HealthAlpaca exhibits comparable performance to much larger models (GPT-3.5, GPT-4 and Gemini-Pro), achieving the best performance in **8 out of 10** tasks. Ablation studies highlight the effectiveness of context enhancement strategies. Notably, we observe that our context enhancement can yield up to **23.8%** improvement in performance. While constructing contextually rich prompts (combining user context, health knowledge and temporal information) exhibits synergistic improvement, the inclusion of health knowledge context in prompts significantly enhances overall performance.

Data and Code Availability In this study, we conduct experiments with 4 public datasets:

(1) PMData ¹ (Thambawita et al. (2020)), (2) LifeSnaps ² (Yfantidou et al. (2022)), (3) GLOBEM ³ (Xu et al. (2022)) and (4) AW_FB ⁴ (Fuller (2020)). Code is available at <https://github.com/mitmedialab/Health-LLM>

Institutional Review Board (IRB) This study has no human-subject research and only uses publicly available data.

1. Introduction

The performance of large language models (LLMs) OpenAI (2023b); Singhal et al. (2022); Nyberg et al. (2021) in diverse text generation and knowledge retrieval tasks presents wide-ranging opportunities Nori et al. (2023); Hegselmann et al. (2023); Gandhi et al. (2023); Wu et al. (2023d). However, in sensitive domains like healthcare, their true capabilities and limitations remain largely unexplored, especially when it comes to harnessing the diverse collection of *multi-modal, time-series* data generated by wearable sensors. Unlike static text, this data presents unique challenges for LLMs due to its high dimensionality, non-linear relationships, and continuous nature, requiring them to understand not only individual data points but also their dynamic patterns over time. Al-

1. <https://datasets.simula.no/pmdata/>
2. <https://github.com/Datalab-AUTH/LifeSnaps-EDA>
3. <https://the-globem.github.io/datasets/overview>
4. <https://dataverse.harvard.edu/dataset.xhtml?persistentId=doi:10.7910/DVN/ZS2Z2J>

though specialized medical-domain LLMs have shown promise in capturing domain knowledge [He et al. \(2023\)](#); [Singhal et al. \(2022, 2023\)](#); [Han et al. \(2023\)](#); [Toma et al. \(2023\)](#); [Thirunavukarasu et al. \(2023\)](#); [McDuff et al. \(2023\)](#), their application to consumer health tasks, which rely heavily on physiological (e.g. heart rate) and behavioral time-series data (e.g. daily steps), remains largely untested due to the challenges of grounding LLMs in non-linguistic data and the lack of standardized evaluation benchmarks.

In this paper, we propose **Health-LLM**, a framework tailored to the healthcare domain that aims to bridge the gap between pre-trained knowledge in current LLMs and consumer health problems. We conducted a comprehensive evaluation of 12 state-of-the-art LLMs that are publicly accessible: MedAlpaca ([Han et al., 2023](#)), PMC-Llama ([Wu et al., 2023b](#)), Llama 2 ([Hugo Touvron and Scialom, 2023](#)), BioMedGPT ([Zhang et al., 2024](#)), BioMistral ([Labrak et al., 2024](#)), Asclepius ([Kweon et al., 2023](#)), Clinical-Camel ([Toma et al., 2023](#)), Flan-T5 ([Chung et al., 2022](#)), Palmyra-Med ([Kamble and Alshikh, 2023](#)), GPT-3.5 ([OpenAI, 2023a](#)), GPT-4 ([OpenAI, 2023b](#)) and Gemini-Pro ([Gemini Team, 2023](#)). We cover 10 health prediction tasks across mental health, activity tracking, metabolism and sleep assessment. Our experiments include four steps: (i) zero-shot prompting, (ii) few-shot prompting along with chain-of-thoughts (CoT) and self-consistency (SC) prompting, (iii) instructional fine-tuning and (iv) ablation studies with context enhancement in *zero-shot* setting, where context enhancement refers to the strategic inclusion of additional information - 1) user profile, 2) health knowledge, 3) temporal context, and 4) a combination of these in the prompts for LLMs to improve their understanding in healthcare domain.

We found in (i) that zero-shot prompting shows comparable results to task-specific baseline models. Comparing (i) and (ii) shows that few-shot prompting with bigger LLMs, like GPT-3.5, GPT-4 and Gemini-Pro, can effectively ground numerical time-series data, resulting in significant improvements over zero-shot learning and fine-tuned models in some tasks. Through step (iii), our Alpaca-based fine-tuned model, namely *HealthAlpaca*, exhibits the best performance in 8 out of 10 tasks despite being a substantially smaller than GPT-3.5, GPT-4 and Gemini-Pro. In (iv), our ablation study indicates that the context enhancement strategy yields up to 23.8% performance improvement, emphasizing the importance of contextual information in the prompt for LLMs

in the healthcare domain. Finally, we present two case studies demonstrating the step-by-step reasoning process of representative LLMs in health prediction tasks, illustrating their ability to capture time-series data and offer personalized recommendations.

The contribution of our paper can be summarized as follows:

- We present Health-LLM, a framework that enables LLMs to adapt to health predictions by prompting/training via wearable sensor data.
- We combine four publicly available health datasets, curate ten novel consumer health tasks, and conduct evaluations with twelve state-of-the-art LLMs.
- We show the effectiveness of context enhancement strategies for Health-LLMs and release our fine-tuned model HealthAlpaca, as the first set of open-source LLMs targeted for consumer health prediction tasks.

2. Related Work

2.1. Wearable Sensor Data with LLMs

Wearable sensor technology has transformed personal health monitoring, enabling continuous tracking of vital physiological data such as heart rate variability and step counts [Perez et al. \(2019\)](#). These time-series data can be represented in different forms, including statistical summaries (e.g., mean, standard deviation), Fourier transforms, and more.

The representation of data before feeding it into the model can also vary. Data can be input as raw time-series text, where the textual information is directly used, or as encoder embeddings obtained from modality-specific ([Belyaeva et al., 2023](#)) and timestamp encoders ([Zhou et al., 2022, 2021](#)). Also, the integration of temporal information has proven to be effective in enhancing outcomes ([Wen et al., 2023](#)).

The synergy between wearable sensor data and advanced machine learning techniques holds promise in predicting diverse health outcomes, such as depression scores ([Mullick et al., 2022](#); [Englhardt et al., 2023](#)), early detection of atrial fibrillation ([Chen et al., 2022](#)), and monitoring stress levels ([Vos et al., 2023](#)). Additionally, this integration facilitates personalized monitoring in areas like nutrition ([Sempionatto et al., 2021](#); [Romero-Tapiador et al., 2023](#)) and stress management ([Tazarv et al., 2021](#)).

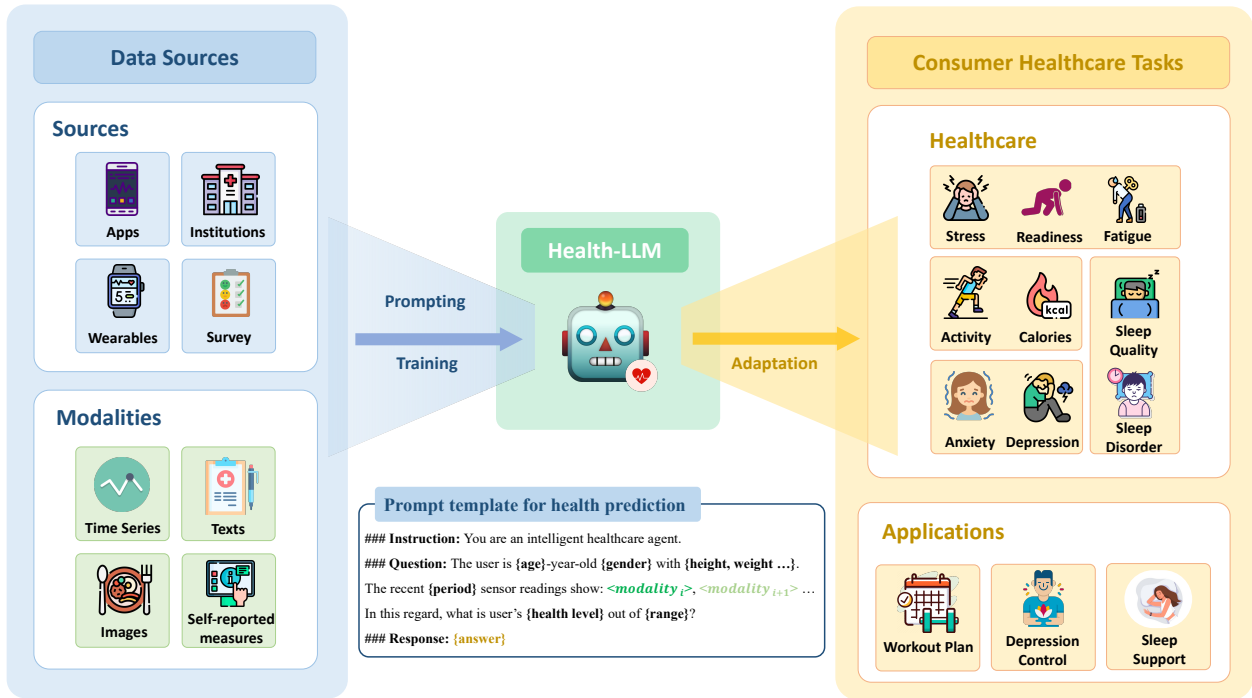


Figure 1: **Health-LLM**. We present a framework for evaluating LLM performance on a diverse set of health prediction tasks, training and prompting the models with multi-modal health data.

2.2. Health LLMs

The integration of LLMs in healthcare is a rapidly growing research field (Tu et al., 2024; Wang et al., 2023a; Liu et al., 2023; Han et al., 2023; Tang et al., 2023; Belyaeva et al., 2023). For instance, Singhal et al. (2022) demonstrated the efficacy of LLMs through a combination of base improvements (PaLM 2), medical domain fine-tuning, and ensemble refinement, outperforming benchmarks across various datasets. Additionally, Xu et al. (2023) explored LLM adaptation for diverse mental health tasks using online-text data, employing techniques like zero-shot, few-shot, and fine-tuning, while our work covers multi-modal time-series data collected from wearable sensors.

In a recent comprehensive evaluation of GPT-4 (Nori et al., 2023), the general-purpose model without fine-tuning, surpassed the USMLE passing score by over 20 points, outperforming earlier models like GPT-3.5 and medically fine-tuned models. This study showed GPT-4’s qualitative abilities in explaining medical reasoning, personalizing explanations, and crafting counterfactual scenarios.

3. Methods

3.1. Zero-shot Prompting

The aim of zero-shot prompting is to evaluate the capability of pre-trained knowledge in LLMs on health prediction tasks. To this end, we first design a basic prompt setup (*bs*) that summarizes the wearable sensor data into a paragraph. Then, we introduce a comprehensive zero-shot prompting along with four types of *context enhancements* introduced in Table 1 and Equation 1-2. 1) **User Context** (*uc*) provides user-specific information such as age, gender, weight, height, etc., which provides additional information that affects the understanding of health knowledge. 2) **Health Context** (*hc*) provides the definition and equation that controls certain health targets to inject new health knowledge into LLMs. 3) **Temporal Context** (*tc*) is adopted to test the importance of temporal aspects in time-series data. Instead of using aggregated statistics, we utilize the raw time-series sequence. Among different sets of temporal context representations, we empirically observe that using natural language string showed the best performance.

Table 1: Different Types of Contexts in Health Prompts.

Context	Prompt
Basic (<i>bs</i>)	The analysis of recent period: {14} days averaged sensor readings show: Steps: {812.0} steps, Burned Calories: {97.0} calories, Resting Heart Rate: {66.54} beats/min, Sleep Minutes: {487.0} minutes, Mood: {3} out of 5. In this regard, what is the predicted readiness score/level between 0 and 10?
User Context (<i>uc</i>)	Given the user’s profile as age: {23}-year-old, sex: {male} and height: {182} cm, the analysis of recent period: {14} days averaged sensor readings show: Steps: {812.0} steps, Burned Calories: {97.0} calories, Resting Heart Rate: {66.54} beats/min, Sleep Minutes: {487.0} minutes, Mood: {3} out of 5. In this regard, what is the predicted readiness score/level between 0 and 10?
Health Context (<i>hc</i>)	Readiness score is an indicator of how prepared our body is for physical activity. It is decided by activity, recent sleep and heart rate variability. The analysis of recent period: {14} days averaged sensor readings show: Steps: {812.0} steps, Burned Calories: {97.0} calories, Resting Heart Rate: {66.54} beats/min, Sleep Minutes: {487.0} minutes, Mood: {3} out of 5. In this regard, what is the predicted readiness score/level between 0 and 10?
Temporal Context (<i>tc</i>)	The analysis of recent period: {14} days sensor readings show: Steps: {"NaN, 991.0, ..., NaN"} steps, Burned Calories: {"NaN, 94.0 ..., NaN"} calories, Resting Heart Rate: {"69.32, 67.72, ..., 64.55"} beats/min, Sleep Minutes: {"534.0, 455.0, ..., 405.0"} minutes, [Mood]: 3 out of 5. In this regard, what is the predicted readiness score/level between 0 and 10?
All (<i>all</i>)	Readiness score is an indicator of how prepared our body is for physical activity. It is decided by activity, recent sleep and heart rate variability. The user is 23-year-old male with 182 cm. The analysis of recent period: {14} days sensor readings show: Steps: {"NaN, 991.0, ..., NaN"} steps, Burned Calories: {"NaN, 94.0 ..., NaN"} calories, Resting Heart Rate: {"69.32, 67.72, ..., 64.55"} beats/min, Sleep Minutes: {"534.0, 455.0, ..., 405.0"} minutes, [Mood]: 3 out of 5. In this regard, what is the predicted readiness score/level between 0 and 10?

$$\text{Prompt}_{ZS} = \text{Instruction}_{ZS} + \text{Context} + \text{Question} + \text{Output Format} \quad (1)$$

$$\text{Context} = C_{health}^* + C_{user}^* + TE^*(\text{TimeSeriesData}) \quad (2)$$

$$\text{Prompt}_{FS} = \text{Instruction}_{FS} + (\text{Context} + \text{Question, Answer})_N + \text{Prompt}_{ZS} \quad (3)$$

3.2. Few-shot Prompting

Few-shot prompting involves using a limited selection of demonstration examples within the prompts to facilitate in-context learning. In our case, we adopt 3-shot setting. These demonstrations are used exclusively within the prompts, while the model parameters remain static. This approach is akin to providing the model with a handful of case studies to help it grasp and apply healthcare domain knowledge effectively. In addition to few-shot prompting, we further enriched the prompting strategy by integrating Chain-of-Thoughts (CoT) (Wei et al., 2022) and Self-Consistency (SC) (Wang et al., 2022) prompting techniques. The incorporation of CoT prompting facilitates a more cohesive and contextually nuanced understanding, allowing the model to connect ideas seamlessly. Simultaneously, SC prompting contributed to refining the model’s responses by promoting internal coherence and logical consistency. Together, these methodologies synergistically formed a robust prompt, leveraging diverse prompting techniques to optimize the LLMs health understanding capabilities (Wang et al., 2023b). In this work,

we used $N = 5$ candidate reasoning paths for SC prompting.

3.3. Instruction Tuning

Instruction tuning is a technique where all parameters of a pre-trained model are further trained or fine-tuned on a target task. This process allows the model to adapt its pre-trained knowledge to the specificities of the new task, optimizing its performance. In the context of health prediction, fine-tuning allows the model to deeply understand physiological terminologies, mechanisms, and context, thereby enhancing its ability to generate accurate and contextually relevant responses.

Instead of fine-tuning all parameters, methods like LoRA (Hu et al., 2021) involves training a small proportion of parameters, by injecting trainable low-rank matrices into each layer of the pre-trained model. In the Health-LLM context, these Parameter Efficient Fine-tuning (PEFT) techniques enable the model to adapt to healthcare tasks while maintaining computational efficiency.

3.4. Temporal Encoding Methods

The methodologies for encoding time-series data into textual formats is crucial in performing health predictions with LLMs. In the context of our consumer health prediction task, we selected to use the **Natural Language String** method. This was due to its simplicity, effectiveness, and wide acceptance in the field, as evidenced by various prior studies (Gruver et al., 2023; Liu et al., 2023). This method’s inherent interpretability aligns seamlessly with our objective to develop models that are easily understood and manipulated by users. Furthermore, it’s capacity to handle missing values without imputation by adopting special symbols (e.g. NaN), making them more resilient to data corruption, is highly advantageous for health-related tasks where data may often be incomplete or irregular. Lastly, our decision was further influenced by the organization of physiological data in our datasets by specific time windows (daily, weekly, monthly), which are intuitively represented through natural language. This facilitates a clearer understanding and processing by LLMs, enhancing our model’s capability to make nuanced predictions.

4. Experiment

4.1. Datasets and Tasks

We consider four wearable sensor datasets that contained: (1) multi-modal physiological data, (2) user self-reported measures, (3) enough distinct time windows to evaluate over. Table 6 summarizes the dataset topic, tasks, metric to evaluate, size and text length and Table 2 presents the features used in the prompt for each task. For the train/test split, we selected 0.1 portion of the original set as the test set and randomly sampled the data from different participants as possible. The choice of ten tasks across four datasets were inspired by the functions provided by consumer health wearables (e.g. Fitbit, Apple Watch) and the previous works of LLMs in diverse applications (Liu et al., 2023; Wu et al., 2023c).

PMData Thambawita et al. (2020) A dataset of $n = 16$ participants (twelve men and three women, in the range of 25-60 years, with an average age of 34 years) during 5 months using Fitbit Versa 2 smartwatch wristbands (Objective Biometrics and Activity Data), Google Forms (Demographics, Food, Drinking, and Weight) and PMSys sports logging smartphone application (self-reported measures like

fatigue, mood, stress, etc) that combines conventional lifelogging data with sports-activity data. Tasks associated with this dataset include:

- **Stress (STRS)**: Estimation of an individual’s stress level based on physiological data and self-reported measures.
- **Readiness (READ)**: Assessment of an individual’s preparedness for physical activity/exercise.
- **Fatigue (FATG)**: Monitoring of signs of tiredness or exhaustion.
- **Sleep Quality (SQ)**: Assessment of factors including total sleep time, efficiency, frequency and duration of awakenings during the night.

LifeSnaps Yfantidou et al. (2022) A comprehensive, multi-modal dataset collected over a period more than 4 months by $n = 71$ participants (42 male and 29 female, half under 30 and half over 30). The annotations were collected from Fitbit Sense (Automatically Synced Data; sleep, heart rate, stress, etc) watch, SEMA3 Data (Ecological Momentary Assessments; context and mood, step goal, etc) and from the validated surveys (Self Reported Data; demographics, health, etc). Tasks associated with this dataset include:

- **Stress Resilience (SR)**: Assessment of an individual’s ability to recover from or adapt positively to stressors.
- **Sleep Disorder (SD)**: Identification of potential sleep disorders like insomnia or sleep apnea through analysis of recorded data.

GLOBEM Xu et al. (2022) A multi-year passive sensing datasets over 705 user-years and $n = 497$ participants’ (58.9% of females, 24.2% of immigrants, 38.2% of first-generations, and 9.1% of disability, and 53.9% of Asian and 31.9% of White) data collected from mobile (AWARE framework), wearable sensors (Fitbit Flex2 and Inspire 2) and survey data (Ecological Momentary Assessment). Tasks associated with this dataset include:

- **Depression (DEP)**: Use of machine learning algorithms that analyze patterns in user behavior and language use for depression detection.
- **Anxiety (ANX)**: Identification of anxiety often relies on behavioral markers such as irregular sleep patterns or heightened physiological responses like increased heart rate.

Table 2: **Consumer Health Tasks.** We define thirteen tasks from six datasets and classify them into four topics. * in the prompt indicates the optional contexts for the ablation.

Topic	Dataset	Task	Metric	Prompt	Target
MHealth	PMData	Stress Prediction	MAE ↓	{Target} refers to [Health Knowledge]*. Given the [User Info]*, and [Period] sequence of Steps: [Steps], Calories Burn: [Calories], Resting Heart Rate: [RHR], Sleep Duration: [SleepMinutes], Mood: [Mood]. What will my stress level be?	[Stress]
	LifeSnaps	Stress Resilience Prediction	MAE ↓	[Target] refers to [Health Knowledge]*. Given the [User Info]* and following [Period] sequence of data, predict the Stress Resilience Index. Stress Score: [StressScore], Positive Affect Score: [PosAffectScore], Negative Affect Score: [NegAffectScore], Lightly Active Minutes: [Duration], Moderately Active Minutes: [Duration], Very Active Minutes: [Duration], Sleep Efficiency: [SleepEfficiency], Sleep Deep Ratio: [SleepDeepRatio], Sleep Light Ratio: [SleepLightRatio], Sleep REM Ratio: [SleepREMRatio].	[Stress Resilience]
	GLOBEM	Estimate of PHQ4 Score	MAE ↓ MAE ↓	[Target] refers to [Health Knowledge]*. Steps during last [Period] sequence of maximum, minimum, average, median, standard deviation daily step count were [ListOfSteps] respectively. Sleep during last [Period] sequence of sleep efficiency, duration the user stayed in bed after waking up, duration the user spent to fall asleep, duration the user stayed awake but still in bed, duration the user spent to fall asleep are [ListOfDurations] in average. In this regard, what would be [Target]?"	[PHQ]
Activity	PMData	Readiness Prediction Fatigue Prediction	MAE ↓ Accuracy ↑	[Target] refers to [Health Knowledge]*. Steps: [Steps], Burned Calories: [Calories], Resting Heart Rate: [RHR], SleepMinutes: [Duration], Mood: [Mood]. What will my readiness level be?	[Readiness] [Fatigue]
	AW_FB	Activity Recognition	Accuracy ↑	[Target] refers to [Health Knowledge]*. Predict the activity type among [ListOfActivities] given the following information [User Info]*, Steps: [Steps], Burned Calories: [Calories], Heart Rate: [HR]."	[Activity]
Metabolic	AW_FB	Calorie Burn Estimate	MAE ↓	[Target] refers to [Health Knowledge]*. Predict the burned calories given the following information. [User Info]*, Steps: [Steps], Heart Rate: [HR].	[Calories]
Sleep	PMData	Sleep Quality Prediction	MAE ↓	[Target] refers to [Health Knowledge]*. Steps: [Steps], Burned Calories: [Calories], Resting Heart Rate: [RHR], SleepMinutes: [Duration], Mood: [Mood]. What will my sleep quality level be?	[SQ]
	LifeSnaps	Sleep Disorder Prediction	Accuracy ↑	[Target] refers to [Health Knowledge]*. Given the following data, predict whether there exists sleep disorder (1) or not (0). Sleep Duration: [Duration], Minutes Awake: [Duration], Sleep Efficiency: [Efficiency], Sleep Deep Ratio: [SleepDeepRatio], Sleep Wake Ratio: [SleepWakeRatio], Sleep Light Ratio: [SleepLightRatio], Sleep REM Ratio: [SleepREMRatio], RMSSD: [RMSSD], SPO2: [SPO2], Full Sleep Breathing Rate: [BreathingRate], BPM: [BPM], Resting Hour: [Duration].	[Sleep Disorder]

AW_FB Fuller (2020) A dataset examining the accuracy of consumer wearable devices (GENEActiv, Apple Watch Series 2 and Fitbit Charge HR2) collected from $n = 49$ participants (26 females and 23 males) over 104 hours of activity logs in a lab based protocol. Tasks associated with this dataset include:

- **Calorie Burn (CAL):** Estimation of the amount of energy expended by an individual during physical activities.
- **Activity (ACT):** Identification of the types of physical activities based on sensor data.

4.2. Models

We experimented with multiple LLMs with different sizes, pre-training targets, and availability.

- **MedAlpaca (7B, 13B) Han et al. (2023):** An advanced LLM, fine-tuned specifically for medical

question-answering. Built upon the foundations of Alpaca, it utilizes a diverse array of medical texts.

- **PMC-Llama (13B) Wu et al. (2023a):** A specialized open-source LLM, incorporating knowledge from 4.8M biomedical papers and 30K medical textbooks.
- **Asclepius (7B) Kweon et al. (2023):** A publicly shareable clinical LLM developed to circumvent privacy regulations by training on synthetic clinical notes extracted from biomedical literature.
- **ClinicalCamel (70B) Toma et al. (2023):** An open LLM, fine-tuned on the LLaMA-2 70B architecture using QLoRA.
- **Flan-T5 (3B) Chung et al. (2022):** An instruction fine-tuned version of T5 or Text-to-Text Transfer Transformer Language Model.

Table 3: **Performance Evaluation of LLMs on Health Prediction Tasks.** **STRS**: Stress, **READ**: Readiness, **FATG**: Fatigue, **SQ**: Sleep Quality, **SR**: Stress Resilience, **SD**: Sleep Disorder, **ANX**: Anxiety, **DEP**: Depression, **ACT**: Activity, **CAL**: Calories. “-” denotes the failure cases due to token size limit or unreasonable responses. “N/A” denotes the case where the prediction is not reported or cannot be conducted. For each column (task), the best result is **bolded**, and the second best is underlined. CoT denotes the chain-of-thoughts and SC denotes the self-consistency prompting. For each task, arrows in the parenthesis indicate the desired direction of improvement. \uparrow indicates higher values are better for accuracy, while \downarrow indicates lower values are better for mean absolute error.

Dataset	PMData				LifeSnaps		GLOBEM		AW_FB		
	Model	STRS (\downarrow)	READ (\downarrow)	FATG (\uparrow)	SQ (\downarrow)	SR (\downarrow)	SD (\uparrow)	ANX (\downarrow)	DEP (\downarrow)	ACT (\uparrow)	CAL (\downarrow)
Zero-shot	MedAlpaca	0.76 \pm 0.1	2.18 \pm 0.1	46.8 \pm 11	0.68 \pm 0.0	1.17 \pm 0.1	40.3 \pm 1.6	1.23 \pm 0.3	0.89 \pm 0.1	21.7 \pm 4.4	35.0 \pm 6.0
	PMC-Llama	1.33 \pm 0.0	4.83 \pm 1.2	0.00 \pm 0.0	2.25 \pm 0.0	1.21 \pm 0.8	41.7 \pm 8.3	2.33 \pm 0.5	2.23 \pm 0.3	-	43.4 \pm 1.2
	Asclepius	0.43 \pm 0.0	1.44 \pm 0.3	27.3 \pm 10	0.45 \pm 0.1	-	-	0.82 \pm 0.2	1.10 \pm 0.0	-	28.9 \pm 11
	ClinicalCamel	0.40 \pm 0.1	2.11 \pm 0.1	58.1 \pm 3.2	0.37 \pm 0.1	1.35 \pm 0.1	88.3 \pm 1.6	0.97 \pm 0.0	0.79 \pm 0.1	16.3 \pm 1.2	43.4 \pm 5.3
	Flan-T5	0.36 \pm 0.0	2.48 \pm 0.2	56.5 \pm 1.6	0.56 \pm 0.1	2.20 \pm 0.3	57.1 \pm 3.5	2.84 \pm 0.0	2.89 \pm 0.0	23.4 \pm 0.0	66.0 \pm 0.4
	Palmyra-Med	0.83 \pm 0.1	2.11 \pm 0.1	43.5 \pm 15	0.44 \pm 0.0	1.03 \pm 0.0	3.13 \pm 0.0	2.07 \pm 0.1	1.99 \pm 0.2	29.7 \pm 1.1	75.3 \pm 5.5
	Llama 2	0.57 \pm 0.2	2.86 \pm 0.4	41.2 \pm 13	0.89 \pm 0.3	-	-	1.19 \pm 0.1	1.23 \pm 0.3	-	-
	BioMedGPT	0.37 \pm 0.0	2.49 \pm 0.2	61.2 \pm 3.3	0.41 \pm 0.1	0.77 \pm 0.2	-	0.95 \pm 0.2	0.85 \pm 0.1	12.2 \pm 3.1	-
	BioMistral	0.55 \pm 0.1	2.12 \pm 0.2	56.6 \pm 3.1	0.45 \pm 0.0	1.59 \pm 0.5	90.0 \pm 0.0	0.90 \pm 0.1	-	18.4 \pm 1.1	41.0 \pm 4.3
	GPT-3.5	-	2.38 \pm 0.1	70.8 \pm 4.2	0.87 \pm 0.0	1.21 \pm 0.3	19.0 \pm 10	-	-	13.8 \pm 6.2	36.4 \pm 0.0
Few-shot	GPT-4	-	2.22 \pm 0.1	72.2 \pm 2.0	0.73 \pm 0.1	1.23 \pm 0.1	-	-	22.6 \pm 2.6	75.2 \pm 8.2	
	Gemini-Pro	0.79 \pm 0.0	1.69 \pm 0.1	34.0 \pm 9.8	0.78 \pm 0.1	2.67 \pm 1.1	84.6 \pm 3.1	1.03 \pm 0.1	0.95 \pm 0.1	17.7 \pm 0.0	31.4 \pm 1.9
	MedAlpaca	0.78 \pm 0.1	1.94 \pm 0.2	36.2 \pm 12	0.69 \pm 0.1	0.94 \pm 0.2	49.6 \pm 11	0.97 \pm 0.2	0.56 \pm 0.1	19.3 \pm 8.1	36.7 \pm 5.6
	+ CoT	0.83 \pm 0.1	2.61 \pm 0.0	9.68 \pm 0.0	0.58 \pm 0.0	1.13 \pm 0.3	33.3 \pm 13	1.67 \pm 0.5	1.26 \pm 0.3	18.1 \pm 0.6	36.1 \pm 0.5
	+ CoT-SC	0.50 \pm 0.2	2.08 \pm 0.1	22.6 \pm 6.5	0.51 \pm 0.0	0.99 \pm 0.2	46.7 \pm 2.7	1.32 \pm 0.1	0.58 \pm 0.1	19.2 \pm 0.3	35.3 \pm 0.0
	GPT-3.5	0.94 \pm 0.1	1.62 \pm 0.1	<u>73.9</u> \pm 16	0.79 \pm 0.1	0.80 \pm 0.0	58.5 \pm 4.4	1.98 \pm 0.1	0.68 \pm 0.2	26.3 \pm 1.1	26.6 \pm 5.5
	+ CoT	0.70 \pm 0.1	1.76 \pm 0.2	77.0 \pm 4.7	0.73 \pm 0.2	0.69 \pm 0.1	51.7 \pm 1.7	1.08 \pm 0.1	0.59 \pm 0.1	23.9 \pm 0.5	38.9 \pm 0.8
	+ CoT-SC	0.57 \pm 0.1	1.46 \pm 0.3	43.3 \pm 12	0.49 \pm 0.0	-	83.0 \pm 8.0	0.85 \pm 0.3	0.56 \pm 0.1	29.8 \pm 3.5	37.7 \pm 0.1
	GPT-4	0.76 \pm 0.1	1.64 \pm 0.3	61.3 \pm 4.8	0.60 \pm 0.1	0.45 \pm 0.0	73.4 \pm 2.3	1.11 \pm 0.3	0.60 \pm 0.1	15.4 \pm 2.3	24.0 \pm 1.4
	+ CoT	0.52 \pm 0.1	1.84 \pm 0.2	73.9 \pm 6.7	0.64 \pm 0.0	0.41 \pm 0.2	71.8 \pm 19	0.90 \pm 0.1	0.67 \pm 0.1	21.9 \pm 5.0	33.0 \pm 2.7
+ CoT-SC	0.33 \pm 0.1	0.86 \pm 0.4	46.0 \pm 4.0	0.48 \pm 0.0	<u>0.36</u> \pm 0.1	80.5 \pm 3.3	<u>0.56</u> \pm 0.0	<u>0.41</u> \pm 0.0	22.5 \pm 1.3	31.1 \pm 0.5	
Fine-tune	Gemini-Pro	1.10 \pm 0.0	2.20 \pm 0.2	24.8 \pm 6.5	0.80 \pm 0.1	1.18 \pm 0.1	71.8 \pm 3.6	1.30 \pm 0.1	1.05 \pm 0.1	15.0 \pm 3.3	37.2 \pm 2.2
	+ CoT	1.02 \pm 0.1	2.24 \pm 0.1	23.9 \pm 3.8	0.83 \pm 0.1	1.34 \pm 0.2	71.8 \pm 10	0.95 \pm 0.1	0.86 \pm 0.1	14.4 \pm 2.5	38.8 \pm 5.8
	+ CoT-SC	1.07 \pm 0.0	2.13 \pm 0.1	28.8 \pm 6.5	0.78 \pm 0.1	1.23 \pm 0.1	75.6 \pm 1.8	1.12 \pm 0.1	0.90 \pm 0.0	19.6 \pm 0.0	36.3 \pm 9.6
	HealthAlpaca-lora-7b	0.53 \pm 0.0	1.40 \pm 0.1	50.0 \pm 13	0.58 \pm 0.1	0.62 \pm 0.1	61.7 \pm 4.7	0.62 \pm 0.0	0.51 \pm 0.1	27.4 \pm 5.0	43.6 \pm 0.1
HealthAlpaca-7b	<u>0.31</u> \pm 0.2	1.32 \pm 0.2	70.7 \pm 2.9	<u>0.35</u> \pm 0.1	0.62 \pm 0.1	72.1 \pm 8.9	0.46 \pm 0.1	0.49 \pm 0.2	<u>41.7</u> \pm 1.1	31.5 \pm 6.4	
HealthAlpaca-lora-13b	0.34 \pm 0.0	1.56 \pm 0.0	54.8 \pm 9.7	0.39 \pm 0.1	0.70 \pm 0.0	<u>92.0</u> \pm 4.0	1.04 \pm 0.1	0.67 \pm 0.1	29.0 \pm 0.0	39.6 \pm 1.4	
HealthAlpaca-13b	0.21 \pm 0.0	<u>1.08</u> \pm 0.2	61.2 \pm 3.4	0.14 \pm 0.0	0.32 \pm 0.1	93.9 \pm 3.1	0.95 \pm 0.2	0.24 \pm 0.0	51.0 \pm 3.5	28.5 \pm 5.6	
Baseline	Majority	N/A	N/A	20.0	N/A	N/A	50.0	N/A	N/A	14.3	N/A
	MLP	1.56 \pm 0.3	3.70 \pm 0.2	27.1 \pm 11	1.61 \pm 0.3	1.20 \pm 0.1	23.1 \pm 5.7	1.39 \pm 0.7	1.45 \pm 0.5	13.2 \pm 5.6	44.3 \pm 0.7
	SVM	0.45 \pm 0.0	1.45 \pm 0.0	58.8 \pm 0.0	0.49 \pm 0.0	0.53 \pm 0.0	10.1 \pm 0.0	0.63 \pm 0.0	0.73 \pm 0.0	17.6 \pm 0.0	33.7 \pm 0.0
	RandomForest	0.57 \pm 0.0	1.58 \pm 0.0	61.8 \pm 2.4	0.52 \pm 0.0	0.41 \pm 0.0	9.00 \pm 0.6	1.31 \pm 0.1	0.82 \pm 0.0	32.3 \pm 1.4	25.4 \pm 0.5
	BERT	0.65 \pm 0.2	2.91 \pm 0.3	51.7 \pm 14	0.70 \pm 0.2	0.62 \pm 0.0	53.8 \pm 0.0	0.67 \pm 0.0	0.67 \pm 0.0	19.0 \pm 3.5	41.1 \pm 0.1
	BioMed-RoBERTa	0.39 \pm 0.0	1.31 \pm 0.0	46.9 \pm 0.0	0.47 \pm 0.0	2.01 \pm 0.0	-	2.93 \pm 0.0	3.21 \pm 0.0	-	44.7 \pm 0.0

- **Palmyra-Med** (20B) [Kamble and Alshikh \(2023\)](#): An LLM fine-tuned on a custom medical dataset, demonstrating superior performance on medical knowledge datasets like PubMedQA and MedQA.
- **BioMedGPT** (7B) [Zhang et al. \(2024\)](#): An open multi-modal generative pre-trained transformer (GPT) for biomedicine, which closes the gap between natural language modality and diverse biomedical data modalities via large generative language models.
- **BioMistral** (7B) [Labrak et al. \(2024\)](#): A Mistral-based further pre-trained open source model (BioMistral 7B DARE) suited for the medical domains.
- **Llama 2** (7B) [Hugo Touvron and Scialom \(2023\)](#): Meta AI’s open LLM, which possess features of including a longer context length, improved training on a larger dataset, and fine-tuning via RLHF for greater performance.
- **GPT-3.5** (175B) [OpenAI \(2023a\)](#): A variant of GPT-3.5 (`gpt-3.5-turbo-instruct`) from OpenAI, specifically fine-tuned to provide direct answers or complete text rather than simulating conversations.
- **GPT-4** [OpenAI \(2023b\)](#): A state-of-the-art autoregressive language model from OpenAI (`gpt-4`) exhibiting remarkable capabilities in various NLP tasks including translation, question answering, and text generation without task-specific fine-tuning.
- **Gemini-Pro** [Gemini Team \(2023\)](#): Google DeepMind’s versatile model optimized for diverse tasks across text, code, images, and audio. It offers a balance between capability and efficiency, suitable for both high-performance applications and on-device

Table 4: **Balanced Cross-Dataset Performance Summary of MedAlpaca-7b Finetuning on Single Dataset.** Numbers indicate the results of the model finetuned and tested on the same dataset. The bottom four rows are related Alpaca versions for reference. Green/Red color marks the ones with better/worse cross-dataset performance compared to the zero-shot version MedAlpaca. Box indicate the results of the case where the model fine-tuned and tested on the same dataset. For each task, arrow in the parenthesis indicate the desired direction of improvement. \uparrow indicates higher values are better for accuracy, while \downarrow indicates lower values are better for mean absolute error.

Test Dataset	PMDData				LifeSnaps		GLOBEM		AW_FB	
Finetune Dataset	STRS (\downarrow)	READ (\downarrow)	FATG (\uparrow)	SQ (\downarrow)	SR (\downarrow)	SD (\uparrow)	ANX (\downarrow)	DEP (\downarrow)	ACT (\uparrow)	CAL (\downarrow)
PMDData	0.38	0.94	71.4	0.90	2.02	72.1	2.01	1.76	-	-
LifeSnaps	0.71	3.26	-	-	0.45	80.0	0.86	0.62	25.0	34.8
GLOBEM	0.72	4.62	23.8	1.25	0.72	81.2	0.81	0.29	-	33.9
AW_FB	0.40	1.94	-	1.37	0.94	93.3	1.91	1.57	43.5	33.2
Reference										
MedAlpaca_ZS	0.76	2.18	46.8	0.68	1.17	40.3	1.23	0.89	21.7	35.0
MedAlpaca_FS*	0.50	1.94	36.2	0.51	0.94	49.6	0.97	0.56	19.3	35.3
MedAlpaca_T	0.53	1.38	58.1	0.50	0.68	73.3	0.98	0.81	16.1	6.52
HealthAlpaca	0.31	<u>1.32</u>	<u>70.7</u>	0.35	<u>0.62</u>	72.1	0.46	<u>0.49</u>	<u>41.7</u>	<u>31.5</u>

ZS: Zero Shot Prompting

FS*: Best Performance among Few Shot Prompting

T: Task-specific Fine-tuning

tasks, making advanced AI accessible to developers and enterprises alike.

- **Baseline:** We also compare LLMs against a set of baseline techniques, including simple majority class (classification only), traditional ML models (MLP, SVM, RandomForest), and traditional pre-trained language models (BERT, BioMed-RoBERTa, and PALM_{p-tuned}). Note that these baseline models are all task-specific and thus need to be trained on data from target tasks.

5. Results and Discussion

Table 3 (w/ MAE and Accuracy metric) and Table 7 (w/ MAPE and F1-score metric) in Appendix shows the performance of twelve off-the-shelf LLMs prompted/trained with a set of prompting/fine-tuning techniques and six task-specific baseline models across ten health prediction tasks and four health datasets.

5.1. Zero-shot and Few-shot Performance

In the *zero-shot* setting, there was no clear superiority among the models and surprisingly, relatively smaller models (number of parameters less than 100B) showed better performance (GPT series models and Gemini-Pro took proportion of only 20%

among best and second best performers) than GPT series models and Gemini-Pro. This may suggest that the effectiveness of a model in the health domain is not solely dependent on its size but could also be significantly influenced by its pre-training knowledge in a zero-shot setting. Asclepius for example, was pre-trained on 158k high-quality synthetic clinical dataset, performed on par with GPT-3.5-turbo showing task-specific knowledge can potentially outweigh the benefits of sheer model size [Kweon et al. \(2023\)](#).

In the few-shot setting however, GPT-3.5 and GPT-4 showed improvement with *few-shot* prompting along with CoT and SC in most of the tasks, while such improvement is not significantly observed with smaller models. Specifically, CoT-SC greatly enhanced model’s performance by showing the best performance with CoT-SC in 9 out of 10 tasks indicating the benefit of structured prompting in understanding and predicting behavior. This finding is aligned with prior findings from other non-health tasks in [Tay et al. \(2022\)](#); [Suzgun et al. \(2022\)](#). This indicates that large LLMs have a stronger capability of quickly learning from examples for health tasks.

5.2. Finetuning Performance

Across all categories, our fully *fine-tuned* model, HealthAlpaca, shows the best performance in 8 out of 10 consumer health tasks (Table 7). HealthAlpaca achieves comparable or better performance than GPT-series models and Gemini-Pro which has a two-magnitude larger number of parameters. HealthAlpaca-lora, a model *fine-tuned* with a parameter-efficient finetuning technique LoRA, also shows a performance boost over the larger models in almost all tasks. These results suggest that LLMs can be effectively tuned for tasks with multi-modal time-series wearable data.

Also, as predicted, HealthAlpaca-13b showed better performance than HealthAlpaca-7b and parameter-efficiently fine-tuned models in most of the tasks as it involved more number of parameters and updates in those larger number of parameters.

5.3. Generalization of Fine-tuned Models across Datasets

To assess the generalization capabilities of our *fine-tuned* models, we conducted dataset-specific fine-tuning. We then compared model performance with reference results from task-specific and multi-dataset *fine-tuned* models. The results presented in Table 4 indicate that not surprisingly, MedAlpaca, when fine-tuned on the same dataset with the target task mostly shows the best performance compared to when fine-tuned on other datasets. We also observed that while dataset-specific fine-tuning often failed to predict tasks from other datasets, the multi-dataset *fine-tuned* HealthAlpaca, exhibited reasonable generalization performance across tasks.

Moreover, we made an intriguing observation regarding the synergistic effect between certain datasets. In a few cases, such as AW_FB \rightarrow STRS, and LifeSnaps \rightarrow ANX, the performance of the *fine-tuned* models surpassed that of the zero-shot and dataset-specific fine-tuning approaches. These findings suggest that fine-tuning on a single dataset can provide health knowledge to a certain extent and thereby improve overall generalization results. However, such improvement is not consistently observed across all tasks, and it depends on the overlapping content across datasets.

5.4. Importance of different contexts in Prompt Designing for Healthcare LLMs

In Figure 2, we present the summarized effect of context enhancement strategies in *zero-shot* setting across LLMs and datasets. Overall, adding contexts can significantly improve the model performance. Among four types of context information, adding health context shows the biggest performance boost in average. More interestingly, our results show varied effectiveness on different LLMs and datasets. From an LLM perspective (shown in Figure 2b), Palmyra-Med benefited the most from the enhancement, where it shows up to 44.58% improvement when adding temporal context (*tc*). Palmyra-Med’s significant improvement can be attributed to its fine-tuning with two medical datasets: PubMedQA, a resource focused on question-answering in the medical domain, and MedQA, which provides a broad range of medical questions and answers. By integrating temporal context, Palmyra-Med leverages its enhanced medical literature comprehension and temporal pattern recognition, leading to substantial improvements in predictive accuracy by effectively interpreting medical conditions’ progression and patterns over time. However, *tc* does not help improve the performance for large-size LLMs like GPT-3.5 or GPT-4. This is probably because they already possess the capability to understand the statistical traits in the time-series data as shown in the case study in Figure 3 (e.g. average, max, variation). From the dataset perspective, our enhancement strategies boost performance in all datasets, and LifeSnaps shows the best improvement up to 34.46%. GLOBEM and AW_FB showed significant performance improvement when applying *tc* and *all* context enhancements. This might be because of temporal aspects in the dataset. For instance, GLOBEM provides 7- and 14-days statistics for each feature, which highlights the temporal dimension.

5.5. Importance of Training Size in Fine-tuning Performance

To understand the required amount of data for effective fine-tuning, we conduct experiments on fine-tuning Alpaca with varying down-sampled training sizes: 5%, 15%, 25%, 50%, and 100% of the original dataset, and evaluate health prediction performances with three seeds. Figure 4 in Appendix shows the overall results. With 15% of the original dataset, the fine-tuned model already outperforms the *zero-shot* performance on all 13 tasks. This number provides

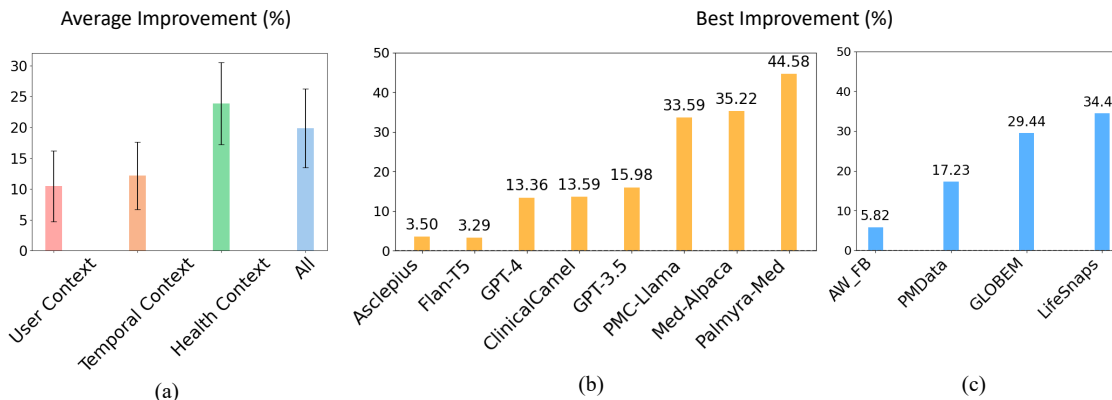


Figure 2: (a): Average Performance Improvement over basic (*bs*) across contexts. (b): Best Performance Improvement across LLMs. (c): Best Performance Improvement across Datasets. Note that few models (Llama 2, Gemini-Pro, BioMedGPT and BioMistral) were excluded in this experiment due to the prioritization of models based on integration timelines.

guidance for fine-tuning when computing resources are limited. As expected, we observe an increasing trend in performance with more training data.

5.6. Case Study of LLM’s Capability on Health Reasoning

In addition to health predictions, it is critical to see whether the health reasoning are also valid especially in healthcare domain. Here, we show two representative cases; 1) READ: Readiness Score Prediction and 2) SD: Sleep Disorder Prediction across four LLMs including our fine-tuned model HealthAlpaca.

5.6.1. LLM’S CAPABILITY TO UNDERSTAND THE TIME-SERIES DATA

In analyzing time-series data, LLMs show distinct approaches. HealthAlpaca for example, shows a conservative interpretation of data like steps and mood scores to suggest moderate readiness levels (see Figure 3). It emphasizes stability, suggesting suitability for moderate activities. GPT-4 on the other hand, identifies days with high activity to suggest greater readiness, highlighting its focus on peak performance.

5.6.2. FALSE AND HALLUCINATED REASONING FROM LLMs

False reasoning in LLMs occurs when data is misinterpreted or generic standards are misapplied. In readiness evaluations (Figure 3), such reasoning can lead to overestimation or underestimation of physical readiness. For instance, HealthAlpaca might not fully

value high-activity days, while GPT-4 could underestimate the need for consistent activity. In case of GPT-3.5, there was a misinterpretation of the question, particularly around the provided resting heart rate information. Gemini-Pro, on the other hand, makes reasonable predictions but misinterprets that the average calorie burns are in the range of recommended daily caloric burn.

6. Conclusion

In this paper, we present the first comprehensive evaluation of twelve off-the-shelf LLMs across ten consumer health prediction tasks (binary, multi-class classification, and regression) spanning four public health datasets. Our experiments encompass a variety of prompting and fine-tuning techniques. The results reveal several interesting findings. First, our context enhancement strategy boosts the performance across all datasets and LLMs, particularly emphasizing the importance of incorporating health knowledge context in prompts. More importantly, our fine-tuned model, HealthAlpaca, demonstrates the best performance in 8 out of 10 tasks, outperforming much larger models such as GPT-4 and Gemini-Pro, even when these are equipped with few-shot prompting. Additionally, we conducted a case study on selected examples to highlight the LLMs’ reasoning capabilities and limitations regarding false and hallucinated reasoning in health predictions. Ethical concerns regarding privacy and bias still remains, necessitating further investigation before real-world deployment.

Limitation and Future Work This study’s reliance on self-reported health data limits its clinical applicability Zack et al. (2023) and raises ethical considerations, particularly regarding data validity and user communication. Additionally, the “black-box” nature of LLMs complicates the assessment of their clinical validity. To address these issues, future work will focus on: 1) conducting evaluations with clinically diagnosed datasets in collaboration with healthcare professionals to enhance clinical relevance; 2) ensuring ethical and regulatory compliance, particularly in how health-related predictions are communicated to users; 3) improving LLMs’ explainability to facilitate understanding of their decision-making processes, thereby aiding in the accurate interpretation of health predictions; and 4) incorporating privacy-preserving technologies like federated learning Raemini (2023) to protect sensitive health information.

Acknowledgments

We thank Yoon Kim and Rosalind Picard at MIT, Vivek Natarajan and Ming-Zher Poh at Google for their revisions, feedback, and support.

References

- Anastasiya Belyaeva, Justin Cosentino, Farhad Hormozdiari, Krish Eswaran, Shravya Shetty, Greg Corrado, Andrew Carroll, Cory Y. McLean, and Nicholas A. Furlotte. Multimodal llms for health grounded in individual-specific data, 2023.
- Yunfan Chen, Chong Zhang, Chengyu Liu, Yiming Wang, and Xiangkui Wan. Atrial fibrillation detection using a feedforward neural network. *Journal of Medical and Biological Engineering*, 42(1): 63–73, 2022.
- Hyung Won Chung, ... Le Hou, and Jason Wei. Scaling instruction-finetuned language models, 2022.
- Zachary Englhardt, Chengqian Ma, Margaret E Morris, Xuhai Xu, Chun-Cheng Chang, Lianhui Qin, Xin Liu, Shwetak Patel, Vikram Iyer, et al. From classification to clinical insights: Towards analyzing and reasoning about mobile and behavioral health data with large language models. *arXiv preprint arXiv:2311.13063*, 2023.
- Daniel Fuller. Replication Data for: Using machine learning methods to predict physical activity types with Apple Watch and Fitbit data using indirect calorimetry as the criterion., 2020. URL <https://doi.org/10.7910/DVN/ZS2Z2J>.
- Kanishk Gandhi, Jan-Philipp Fränken, Tobias Gerstenberg, and Noah D. Goodman. Understanding social reasoning in language models with language models, 2023.
- Sebastian Borgeaud ... Oriol Vinyals Gemini Team, Rohan Anil. Gemini: A family of highly capable multimodal models, 2023.
- Nate Gruver, Marc Finzi, Shikai Qiu, and Andrew Gordon Wilson. Large language models are zero-shot time series forecasters, 2023.
- Tianyu Han, Lisa C Adams, Jens-Michalis Papaioannou, Paul Grundmann, Tom Oberhauser, Alexander Löser, Daniel Truhn, and Keno K Bressen. Medalpaca—an open-source collection of medical conversational ai models and training data. *arXiv preprint arXiv:2304.08247*, 2023.
- Kai He, Rui Mao, Qika Lin, Yucheng Ruan, Xiang Lan, Mengling Feng, and Erik Cambria. A survey of large language models for healthcare: from data, technology, and applications to accountability and ethics. *arXiv preprint arXiv:2310.05694*, 2023.
- Stefan Hegselmann, Alejandro Buendia, Hunter Lang, Monica Agrawal, Xiaoyi Jiang, and David Sontag. Tabllm: Few-shot classification of tabular data with large language models, 2023.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models, 2021.
- ... Hugo Touvron and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models, 2023.
- Ming Jin, Shiyu Wang, Lintao Ma, Zhixuan Chu, James Y. Zhang, Xiaoming Shi, Pin-Yu Chen, Yuxuan Liang, Yuan-Fang Li, Shirui Pan, and Qingsong Wen. Time-llm: Time series forecasting by reprogramming large language models, 2023.
- Kiran Kamble and Waseem Alshikh. Palmyra-med: Instruction-based fine-tuning of llms enhancing medical domain performance, 07 2023.
- Sunjun Kweon, Junu Kim, Jiyouon Kim, Sujeong Im, Eunbyeol Cho, Seongsu Bae, Jungwoo Oh, Gyubok

- Lee, Jong Hak Moon, Seng Chan You, Seungjin Baek, Chang Hoon Han, Yoon Bin Jung, Yohan Jo, and Edward Choi. Publicly shareable clinical large language model built on synthetic clinical notes, 2023.
- Yanis Labrak, Adrien Bazoge, Emmanuel Morin, Pierre-Antoine Gourraud, Mickael Rouvier, and Richard Dufour. Biomistral: A collection of open-source pretrained large language models for medical domains, 2024.
- 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 preprint arXiv:2305.15525*, 2023.
- Daniel McDuff, Mike Schaekermann, Tao Tu, Anil Palepu, Amy Wang, Jake Garrison, Karan Singhal, Yash Sharma, Shekoofeh Azizi, Kavita Kulkarini, et al. Towards accurate differential diagnosis with large language models. *arXiv preprint arXiv:2312.00164*, 2023.
- Tahsin Mullick, Ana Radovic, Sam Shaaban, and Afshaneh Doryab. Predicting depression in adolescents using mobile and wearable sensors: Multimodal machine learning-based exploratory study. *JMIR Form Res*, 6(6):e35807, Jun 2022. ISSN 2561-326X. doi: 10.2196/35807. URL <https://formative.jmir.org/2022/6/e35807>.
- Harsha Nori, Nicholas King, Scott Mayer McKinney, Dean Carignan, and Eric Horvitz. Capabilities of gpt-4 on medical challenge problems, 2023.
- Erik P. Nyberg, Ann E. Nicholson, Kevin B. Korb, Michael Wybrow, Ingrid Zukerman, Steven Mascaro, Shreshth Thakur, Abraham Oshni Alvandi, Jeff Riley, Ross Pearson, Shane Morris, Matthieu Herrmann, A.K.M. Azad, Fergus Bolger, Ulrike Hahn, and David Lagnado. Bard: A structured technique for group elicitation of bayesian networks to support analytic reasoning. *Risk Analysis*, 42(6):1155–1178, June 2021. ISSN 1539-6924. doi: 10.1111/risa.13759. URL <http://dx.doi.org/10.1111/risa.13759>.
- OpenAI. gpt-3.5-turbo-instruct, 2023a. URL <https://platform.openai.com/docs/models/gpt-3-5>.
- OpenAI. Gpt-4 technical report, 2023b.
- Marco V. Perez, Kenneth W. Mahaffey, Haley Hedlin, John S. Rumsfeld, Ariadna Garcia, Todd Ferris, Vidhya Balasubramanian, Andrea M. Russo, Amol Rajmane, Lauren Cheung, Grace Hung, Justin Lee, Peter Kowey, Nisha Talati, Divya Nag, Santosh E. Gummidipundi, Alexis Beatty, Melanie True Hills, Sumbul Desai, Christopher B. Granger, Manisha Desai, and Mintu P. Turakhia. Large-scale assessment of a smartwatch to identify atrial fibrillation. *New England Journal of Medicine*, 381(20):1909–1917, 2019. doi: 10.1056/NEJMoa1901183. URL <https://doi.org/10.1056/NEJMoa1901183>. PMID: 31722151.
- Mohammad Raeini. Privacy-preserving large language models (ppllms). Available at SSRN 4512071, 2023.
- Sergio Romero-Tapiador, Blanca Lacruz-Pleguezuelos, Ruben Tolosana, Gala Freixer, Roberto Daza, Cristina M Fernández-Díaz, Elena Aguilar-Aguilar, Jorge Fernández-Cabezas, Silvia Cruz-Gil, Susana Molina, et al. Ai4fooddb: a database for personalized e-health nutrition and lifestyle through wearable devices and artificial intelligence. *Database*, 2023:baad049, 2023.
- Juliane R Sempionatto, Victor Ruiz-Valdepenas Montiel, Eva Vargas, Hazhir Teymourian, and Joseph Wang. Wearable and mobile sensors for personalized nutrition. *ACS sensors*, 6(5):1745–1760, 2021.
- Karan Singhal, Shekoofeh Azizi, Tao Tu, S Sara Mahdavi, Jason Wei, Hyung Won Chung, Nathan Scales, Ajay Tanwani, Heather Cole-Lewis, Stephen Pfohl, et al. Large language models encode clinical knowledge. *arXiv preprint arXiv:2212.13138*, 2022.
- Karan Singhal, Tao Tu, JuraJ Gottweis, Rory Sayres, Ellery Wulczyn, Le Hou, Kevin Clark, Stephen Pfohl, Heather Cole-Lewis, Darlene Neal, et al. Towards expert-level medical question answering with large language models. *arXiv preprint arXiv:2305.09617*, 2023.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, et al. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*, 2022.

- Xiangru Tang, Anni Zou, Zhuosheng Zhang, Yilun Zhao, Xingyao Zhang, Arman Cohan, and Mark Gerstein. Medagents: Large language models as collaborators for zero-shot medical reasoning, 2023.
- Yi Tay, Mostafa Dehghani, Vinh Q Tran, Xavier Garcia, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Neil Houlsby, and Donald Metzler. Unifying language learning paradigms. *arXiv preprint arXiv:2205.05131*, 2022.
- Ali Tazarv, Sina Labbaf, Stephanie M Reich, Nikil Dutt, Amir M Rahmani, and Marco Levorato. Personalized stress monitoring using wearable sensors in everyday settings. In *2021 43rd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 7332–7335. IEEE, 2021.
- Vajira Thambawita, Steven Alexander Hicks, Hanna Borgli, Håkon Kvale Stensland, Debesh Jha, Martin Kristoffer Svensen, Svein-Arne Pettersen, Dag Johansen, Håvard Dagenborg Johansen, Susann Dahl Pettersen, Simon Nordvang, Sigurd Pedersen, Anders Gjerdrum, Tor-Morten Grønli, Per Morten Fredriksen, Ragnhild Eg, Kjeld Hansen, Siri Fagernes, Christine Claudi, Andreas Bjørn-Hansen, Duc Tien Dang Nguyen, Tomas Kupka, Hugo Lewi Hammer, Ramesh Jain, Michael Alexander Riegler, and Pål Halvorsen. Pmdata: A sports logging dataset. In *Proceedings of the 11th ACM Multimedia Systems Conference, MMSys '20*, page 231–236, New York, NY, USA, 2020. Association for Computing Machinery. doi: 10.1145/3339825.3394926.
- Arun James Thirunavukarasu, Darren Shu Jeng Ting, Kabilan Elangovan, Laura Gutierrez, Ting Fang Tan, and Daniel Shu Wei Ting. Large language models in medicine. *Nature medicine*, 29(8):1930–1940, 2023.
- Augustin Toma, Patrick R. Lawler, Jimmy Ba, Rahul G. Krishnan, Barry B. Rubin, and Bo Wang. Clinical camel: An open expert-level medical language model with dialogue-based knowledge encoding, 2023.
- Tao Tu, Anil Palepu, Mike Schaekermann, Khaled Saab, Jan Freyberg, Ryutaro Tanno, Amy Wang, Brenna Li, Mohamed Amin, Nenad Tomasev, Shekoofeh Azizi, Karan Singhal, Yong Cheng, Le Hou, Albert Webson, Kavita Kulkarni, S Sara Mahdavi, Christopher Sementurs, Juraj Gottweis, Joelle Barral, Katherine Chou, Greg S Corrado, Yossi Matias, Alan Karthikesalingam, and Vivek Natarajan. Towards conversational diagnostic ai, 2024.
- Gideon Vos, Kelly Trinh, Zoltan Sarnyai, and Mostafa Rahimi Azghadi. Ensemble machine learning model trained on a new synthesized dataset generalizes well for stress prediction using wearable devices. *Journal of Biomedical Informatics*, 148:104556, 2023. ISSN 1532-0464. doi: <https://doi.org/10.1016/j.jbi.2023.104556>. URL <https://www.sciencedirect.com/science/article/pii/S1532046423002770>.
- Hanyin Wang, Chufan Gao, Christopher Dantona, Bryan Hull, and Jimeng Sun. Drg-llama : Tuning llama model to predict diagnosis-related group for hospitalized patients, 2023a.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*, 2022.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models, 2023b.
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Fei Xia, Ed Chi, Quoc V Le, Denny Zhou, et al. Chain-of-thought prompting elicits reasoning in large language models. *Advances in Neural Information Processing Systems*, 35:24824–24837, 2022.
- Qingsong Wen, Tian Zhou, Chaoli Zhang, Weiqi Chen, Ziqing Ma, Junchi Yan, and Liang Sun. Transformers in time series: A survey, 2023.
- Chaoyi Wu, Weixiong Lin, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-llama: Towards building open-source language models for medicine, 2023a.
- Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-llama: Further fine-tuning llama on medical papers. *arXiv preprint arXiv:2304.14454*, 2023b.

- Ruolan Wu, Chun Yu, Xiaole Pan, Yujia Liu, Ningning Zhang, Yue Fu, Yuhan Wang, Zhi Zheng, Li Chen, Qiaolei Jiang, et al. Mindshift: Leveraging large language models for mental-states-based problematic smartphone use intervention. *arXiv preprint arXiv:2309.16639*, 2023c.
- Shijie Wu, Ozan Irsoy, Steven Lu, Vadim Dabravolski, Mark Dredze, Sebastian Gehrmann, Prabhajan Kambadur, David Rosenberg, and Gideon Mann. Bloomberggpt: A large language model for finance, 2023d.
- Xuhai Xu, Han Zhang, Yasaman Sefidgar, Yiyi Ren, Xin Liu, Woosuk Seo, Jennifer Brown, Kevin Kuehn, Mike Merrill, Paula Nurius, et al. Globem dataset: Multi-year datasets for longitudinal human behavior modeling generalization. *Advances in Neural Information Processing Systems*, 35: 24655–24692, 2022.
- Xuhai Xu, Bingsheng Yao, Yuanzhe Dong, Saadia Gabriel, Hong Yu, James Hendler, Marzyeh Ghassemi, Anind K. Dey, and Dakuo Wang. Mental-llm: Leveraging large language models for mental health prediction via online text data, 2023.
- Sofia Yfantidou, Christina Karagianni, Stefanos Efstathiou, Athena Vakali, Joao Palotti, Dimitrios Panteleimon Giakatos, Thomas Marchioro, Andrei Kazlouski, Elena Ferrari, and Šarūnas Girdzijauskas. Lifesnaps, a 4-month multi-modal dataset capturing unobtrusive snapshots of our lives in the wild. *Scientific Data*, 9(1):663, 2022.
- Travis Zack, Eric Lehman, Mirac Suzgun, Jorge A Rodriguez, Leo Anthony Celi, Judy Gichoya, Dan Jurafsky, Peter Szolovits, David W Bates, Rajal-Elie E Abdunour, et al. Coding inequity: Assessing gpt-4’s potential for perpetuating racial and gender biases in healthcare. *medRxiv*, pages 2023–07, 2023.
- Kai Zhang, Jun Yu, Eashan Adhikarla, Rong Zhou, Zhiling Yan, Yixin Liu, Zhengliang Liu, Lifang He, Brian Davison, Xiang Li, Hui Ren, Sunyang Fu, James Zou, Wei Liu, Jing Huang, Chen Chen, Yuyin Zhou, Tianming Liu, Xun Chen, Yong Chen, Quanzheng Li, Hongfang Liu, and Lichao Sun. Biomedgpt: A unified and generalist biomedical generative pre-trained transformer for vision, language, and multimodal tasks, 2024.
- Haoyi Zhou, Shanghang Zhang, Jieqi Peng, Shuai Zhang, Jianxin Li, Hui Xiong, and Wancai Zhang. Informer: Beyond efficient transformer for long sequence time-series forecasting, 2021.
- Tian Zhou, Ziqing Ma, Qingsong Wen, Xue Wang, Liang Sun, and Rong Jin. Fedformer: Frequency enhanced decomposed transformer for long-term series forecasting, 2022.

Appendix A. Implementation Details

We fine-tune our Health-LLMs on 4 NVIDIA A100 80GB GPUs with a batch size of 128 with different number of epochs for the purpose of fine-tuning. We based on MedAlpaca-7b and -13b models and conducted instruction fine-tuning with cross entropy as the loss function, we backpropagate and update model parameters in 3 and 5 epochs (7b and 13b respectively), with Adam optimizer and a learning rate as $2e^{-5}$ (cosine learning rate scheduler and warmup steps of 100). It took about 3.4 hours for 5 epochs of training with the default training setting. For more training details, we follow the default parameters in the original code repository of MedAlpaca⁵. We set default decoding method as sampling and also use other decoding methods such as controlling the temperature in sampling to generate different reasoning paths. The codes and fine-tuned models will be made publicly available upon the release of the camera-ready version of this paper.

For the zero-shot and few-shot prompting, we utilized open-source models from huggingface with four Nvidia A6000 GPUs and used OpenAI and GenAI APIs for the closed-source models. We conducted oversampling for each dataset to resolve class imbalance issue and merged the train sets together. For the task-specific and multi-dataset used for fine-tuning, we will upload the codes to generate these datasets in the code repository by camera-ready version.

Appendix B. Temporal Encoding Methods

B.0.1. Natural Language String

Following the approach presented in Gruver et al. (2023), we also transform raw time-series data into a language string. This method, known for its simplicity and effectiveness, converts sequential data into an understandable format for both humans and AI models. For instance, a 14-day sensor reading for a consumer health prediction task can be transformed into daily readings of steps, and resting heart rate ({Nan, 991.0, ..., Nan}).

B.0.2. Modality-specific Encoding

In line with the methodology proposed in Belyaeva et al. (2023), this method utilizes pre-trained encoders for each modality to embed non-text data

modalities including time-series data into the same latent space as the language tokens. Though this method may provide nuanced data representations, it introduces complexity and computational overhead that may not always yield performance gains.

B.0.3. Statistical Summary

Similar to the methods used in Jin et al. (2023), this method encapsulates time-series data into statistical summaries such as mean, std, median, etc. While this approach reduces the volume of data to be processed, it may risk overlooking critical temporal patterns inherent in the raw data.

Appendix C. Additional Experiments

Fine-tuned Dataset	Tested on LifeSnaps	
	STRS (MAE \pm SD)	CAL (MAE \pm SD)
PMDData	0.44 \pm 0.1	-
AW_FB	-	44.2 \pm 7.2

Table 5: Summary of additional experiments showing the performance of HealthAlpaca-7b on the LifeSnaps dataset for stress level (STRS) and calorie burn (CAL) prediction.

Additional experiments were conducted to address concerns regarding potential dataset overlap between fine-tuning and evaluation datasets, potentially favoring HealthAlpaca unfairly. Specifically, we focused on 1) stress level (STRS) prediction and 2) calorie burn (CAL) prediction using LifeSnaps dataset labels, which were fine-tuned on PMData and AW_FB datasets. These experiments were designed to showcase HealthAlpaca’s generalizability and robustness by entirely separating the training and evaluation datasets. Results demonstrated promising performance with a mean absolute error (MAE) of 0.44 \pm 0.1 for stress level prediction and MAE of 44.2 \pm 7.2 for calorie burn prediction. Furthermore, we plan to extend these findings by incorporating additional experimental results with HealthAlpaca-7b-lora, HealthAlpaca-13b-lora, and HealthAlpaca-13b models. Additionally, future work will include more cross-dataset experiments using the GLOBEM dataset to further evaluate HealthAlpaca’s performance across diverse datasets.

5. <https://github.com/kbressemed/medAlpaca>

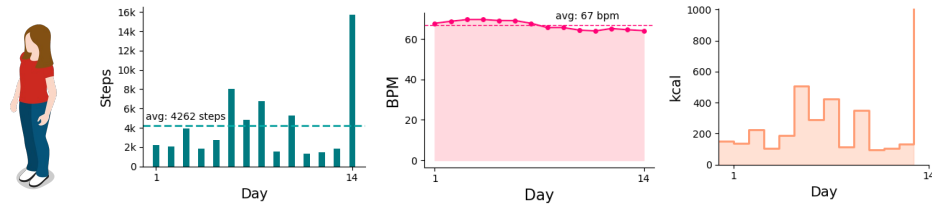
Table 6: **Consumer Health Tasks.** We curate ten tasks with four datasets, across four domains. For the Dataset Size, we indicate the number of QA samples and the label distribution in the parentheses.

Domain	Dataset	Task	Metric	Dataset Size	Text Length (Token)
MHealth	PMData	Stress Prediction	MAE ↓	1784 (0: 1, 1: 21, 2: 315, 3: 833, 4: 240, 5: 8)	202 ± 42
	LifeSnaps	Stress Resilience Prediction	MAE ↓	150 (0-1: 30, 1-2: 67, 2-3: 35, 3-4: 13, 4-5: 5)	202 ± 42
	GLOBEM	Estimate of PHQ4 Depression Estimate of PHQ4 Anxiety	MAE ↓	1735 (0: 748, 1: 555, 2: 296, 3: 108, 4: 28) 1735 (0: 1093, 1: 372, 2: 194, 3: 63, 4: 13)	216 ± 0
Activity	PMData	Readiness Prediction Fatigue Prediction	MAE ↓ Accuracy ↑	1420 (0-1: 32, 2-3: 217, 4-5: 628, 6-7: 423, 8-10: 120) 1418 (1: 53, 2: 410, 3: 845, 4: 101, 5: 9)	202 ± 42
	AW_FB	Activity Recognition	Accuracy ↑	652 (0: 151, 1: 93, 2: 87, 3: 93, 4: 100, 5: 128)	119 ± 0
Metabolic	AW_FB	Calorie Burn Estimate	MAE ↓	652 (0-50: 387, 51-100: 265)	87 ± 1
Sleep	PMData	Sleep Quality Prediction	MAE ↓	1418 (1: 12, 2: 295, 3: 743, 4: 351, 5: 17)	202 ± 42
	LifeSnaps	Sleep Disorder Prediction	Accuracy ↑	148 (0: 40, 1: 108)	303 ± 0

Table 7: **Performance Evaluation of LLMs on Health Prediction Tasks.** **STRS:** Stress, **READ:** Readiness, **FATG:** Fatigue, **SQ:** Sleep Quality, **SR:** Stress Resilience, **SD:** Sleep Disorder, **ANX:** Anxiety, **DEP:** Depression, **ACT:** Activity, **CAL:** Calories. “-” denotes the failure cases due to token size limit or unreasonable responses. “N/A” denotes the case where the prediction is not reported or cannot be conducted. For each column (task), the best result is **bolded**, and the second best is underlined. CoT denotes the chain-of-thoughts and SC denotes the self-consistency prompting. For each task, arrows in the parenthesis indicate the desired direction of improvement. ↑ indicates higher values are better for macro F1-score, while ↓ indicates lower values are better for mean absolute percentage error (MAPE).

Dataset	PMData				LifeSnaps		GLOBEM		AW_FB		
	Model	STRS (↓)	READ (↓)	FATG (↑)	SQ (↓)	SR (↓)	SD (↑)	ANX (↓)	DEP (↓)	ACT (↑)	CAL (↓)
Zero-shot	MedAlpaca	54.0 ± 0.0	49.2 ± 0.0	0.26 ± 0.0	38.3 ± 0.0	87.1 ± 0.0	0.22 ± 0.0	66.7 ± 0.0	72.2 ± 0.0	0.13 ± 0.0	77.1 ± 0.0
	PMC-Llama	57.6 ± 1.8	81.6 ± 14	0.00 ± 0.0	37.7 ± 8.8	55.5 ± 10	0.25 ± 0.1	75.8 ± 7.8	79.3 ± 6.5	-	62.2 ± 11
	Asclepius	31.3 ± 0.0	51.9 ± 0.0	0.07 ± 0.0	13.2 ± 0.0	-	-	47.6 ± 0.0	<u>25.3 ± 0.0</u>	-	72.5 ± 0.0
	ClinicalCamel	12.8 ± 0.0	25.5 ± 0.0	0.16 ± 0.0	13.3 ± 0.0	77.4 ± 0.0	0.48 ± 0.0	<u>33.0 ± 0.0</u>	25.7 ± 1.3	0.10 ± 0.0	120 ± 21
	Flan-T5	16.1 ± 0.0	37.9 ± 0.0	0.24 ± 0.0	11.6 ± 0.0	41.0 ± 0.0	0.38 ± 0.1	83.5 ± 0.0	93.6 ± 7.4	0.13 ± 0.0	94.1 ± 6.9
	Palmyra-Med	56.0 ± 0.0	20.6 ± 0.4	0.25 ± 0.2	13.1 ± 2.1	23.4 ± 3.7	0.47 ± 0.0	85.1 ± 17	89.8 ± 3.9	0.17 ± 0.1	131 ± 14
	Llama 2	18.0 ± 0.0	80.4 ± 0.0	0.18 ± 0.0	34.7 ± 0.0	-	-	74.3 ± 0.0	74.5 ± 0.0	-	-
	BioMedGPT	72.9 ± 0.0	56.3 ± 3.4	0.07 ± 0.0	97.9 ± 0.0	28.2 ± 0.0	-	33.5 ± 0.0	-	0.16 ± 0.1	-
	BioMistral	19.6 ± 0.0	<u>23.8 ± 0.0</u>	0.27 ± 0.0	12.4 ± 0.0	73.7 ± 0.0	0.67 ± 0.2	33.3 ± 0.0	-	0.12 ± 0.0	92.1 ± 11
	GPT-3.5	-	38.4 ± 10	0.12 ± 0.0	24.6 ± 3.7	96.6 ± 0.0	0.23 ± 0.1	-	-	0.13 ± 0.1	93.3 ± 12
GPT-4	-	41.2 ± 4.0	0.26 ± 0.1	23.5 ± 1.4	55.0 ± 14	-	-	-	0.18 ± 0.0	107 ± 23	
Gemini-Pro	24.5 ± 5.6	25.0 ± 3.5	0.19 ± 0.0	24.9 ± 3.5	57.1 ± 8.0	0.63 ± 0.0	47.2 ± 9.8	43.1 ± 8.8	0.11 ± 0.0	116 ± 18	
Few-shot	MedAlpaca	38.5 ± 4.1	46.5 ± 7.6	0.25 ± 0.1	19.2 ± 2.3	77.9 ± 8.5	0.39 ± 0.1	52.3 ± 5.4	72.4 ± 5.6	0.19 ± 0.0	91.1 ± 4.5
	+ CoT	26.0 ± 5.4	43.2 ± 6.4	0.18 ± 0.1	18.1 ± 3.8	66.1 ± 5.9	0.31 ± 0.1	38.7 ± 2.3	56.7 ± 3.5	0.18 ± 0.0	88.8 ± 10
	+ CoT-SC	23.8 ± 2.8	41.1 ± 5.5	0.19 ± 0.0	18.0 ± 0.9	41.9 ± 5.6	0.47 ± 0.2	40.1 ± 3.4	66.4 ± 5.4	0.15 ± 0.0	82.3 ± 6.7
	GPT-3.5	54.1 ± 4.2	36.0 ± 1.5	0.25 ± 0.0	30.2 ± 1.6	79.1 ± 19	0.43 ± 0.1	64.1 ± 3.6	51.1 ± 11	0.10 ± 0.0	84.3 ± 31
	+ CoT	47.3 ± 3.6	41.0 ± 2.4	0.20 ± 0.1	32.5 ± 2.2	65.3 ± 4.2	0.34 ± 0.0	46.4 ± 4.3	39.2 ± 6.9	0.11 ± 0.0	71.1 ± 33
	+ CoT-SC	50.3 ± 3.9	51.9 ± 2.5	0.22 ± 0.0	34.1 ± 2.8	65.4 ± 2.2	0.37 ± 0.1	47.4 ± 3.3	34.4 ± 5.8	0.12 ± 0.0	65.2 ± 24
	GPT-4	20.0 ± 2.6	24.0 ± 4.8	0.27 ± 0.0	14.8 ± 0.8	57.3 ± 1.8	0.54 ± 0.0	51.3 ± 16	59.2 ± 7.3	0.14 ± 0.0	85.7 ± 19
	+ CoT	16.8 ± 3.9	35.7 ± 1.7	0.23 ± 0.1	19.9 ± 1.5	51.2 ± 2.8	0.61 ± 0.0	37.0 ± 1.1	53.2 ± 8.6	0.16 ± 0.0	72.4 ± 16
	+ CoT-SC	16.4 ± 5.8	36.7 ± 1.8	0.23 ± 0.1	20.7 ± 1.5	49.5 ± 3.8	0.61 ± 0.0	46.1 ± 5.9	48.5 ± 6.9	0.17 ± 0.0	<u>64.8 ± 17</u>
	Gemini-Pro	35.4 ± 4.7	41.9 ± 2.8	0.17 ± 0.1	26.2 ± 4.7	54.1 ± 3.2	0.46 ± 0.0	58.9 ± 4.9	55.4 ± 2.7	0.12 ± 0.0	163 ± 23
+ CoT	30.3 ± 2.9	50.4 ± 5.1	0.17 ± 0.0	33.7 ± 2.2	54.0 ± 6.8	0.44 ± 0.0	54.9 ± 1.3	49.1 ± 8.1	0.10 ± 0.0	181 ± 19	
+ CoT-SC	27.1 ± 3.0	47.2 ± 1.6	0.20 ± 0.0	29.1 ± 3.3	55.7 ± 3.0	0.43 ± 0.1	50.6 ± 7.2	50.3 ± 6.0	0.11 ± 0.1	128 ± 16	
Fine-tune	HealthAlpaca-lora-7b	13.9 ± 1.4	37.3 ± 1.4	0.19 ± 0.0	17.2 ± 1.4	38.4 ± 3.8	0.47 ± 0.1	35.9 ± 0.0	34.4 ± 1.0	0.17 ± 0.1	96.9 ± 3.9
	HealthAlpaca-7b	13.1 ± 0.0	33.3 ± 4.7	0.27 ± 0.1	16.7 ± 1.6	36.1 ± 8.8	0.37 ± 0.0	33.2 ± 0.0	26.5 ± 2.2	<u>0.29 ± 0.0</u>	73.5 ± 18
	HealthAlpaca-lora-13b	<u>12.8 ± 0.0</u>	55.9 ± 16	0.45 ± 0.0	<u>6.53 ± 1.7</u>	<u>21.4 ± 2.2</u>	0.78 ± 0.1	29.9 ± 2.4	29.2 ± 2.0	0.21 ± 0.1	75.6 ± 9.4
	HealthAlpaca-13b	11.5 ± 0.0	26.4 ± 0.0	0.58 ± 0.0	2.86 ± 0.0	14.0 ± 0.0	0.91 ± 0.1	21.1 ± 0.0	12.1 ± 1.2	0.31 ± 0.1	73.0 ± 3.2
Baseline	Majority	N/A	N/A	0.20	N/A	N/A	0.50	N/A	N/A	0.14	N/A
	MLP	39.4 ± 3.4	47.5 ± 1.4	0.16 ± 0.0	39.8 ± 2.5	43.8 ± 9.2	<u>0.79 ± 0.1</u>	-	-	0.23 ± 0.0	180 ± 1.3
	SVM	16.3 ± 0.0	42.4 ± 0.0	0.25 ± 0.0	14.5 ± 0.0	22.7 ± 0.0	0.49 ± 0.0	51.7 ± 0.0	56.6 ± 0.0	0.11 ± 0.0	212 ± 11
	RandomForest	15.5 ± 0.7	50.0 ± 0.7	0.34 ± 0.0	15.3 ± 0.8	24.3 ± 0.4	0.48 ± 0.0	102 ± 5.0	69.3 ± 3.1	0.23 ± 0.0	184 ± 16
	BioMed-RoBERTa	46.4 ± 11	59.3 ± 5.6	0.24 ± 0.0	43.9 ± 8.8	23.9 ± 0.9	0.36 ± 0.0	39.4 ± 1.3	37.4 ± 2.6	0.15 ± 0.0	80.3 ± 2.4

Prompt:



Input Prompt | Ground Truth (Readiness Score): 5

The user is **38-year-old female** with **179 cm**. The analysis of recent **2-weeks** sensor readings show:
[Steps]: [2184.0, 2091.0, ..., 15754.0] steps
[Burned Calories]: [153.0, 138.0, ..., 1005.0] calories
[Resting Heart Rate]: [67.7, 68.7, ..., 64.1] beats/min
[Mood]: 3 out of 5
 Please return the readiness score between 1 and 10 and step-by-step reasoning.

Answer: The user's **readiness score is 5**.
 This indicates an **average level of readiness**, suggesting that the user may be able to handle a moderate level of physical activity but not intense exercise.
 The user's **average steps per week over the past two weeks is around 4000, which is lower than the recommended daily count of 10,000**. The user's **resting heart rate seems to be stable and within a normal range**, which is good.
 The user's **mood is rated as 3 out of 5, indicating an average mood state** but this marks room for improvement.

HealthAlpaca | Answer: 5

1. Steps: The user had a wide range of daily steps from a **minimum of around 1,323 to a maximum of 15,754 during the last 2 weeks**. This indicates inconsistency in her physical activities. However, some days show a high amount of steps which is a good sign.
 2. Burned Calories: Similarly to steps, the user also has a **varied amount of burned calories from 97 to 1,005 kcal**. Higher amounts of burned calories indicate a higher level of physical activity.
 3. Resting Heart Rate (RHR): The provided **RHR values are within the normal range (60-100 bpm for adults)**. A lower RHR generally indicates a more efficient heart function and better cardiovascular fitness.
 4. Mood: The given **mood score is average**, suggesting moderate levels of stress or emotional well-being.
 Considering the above points, the **readiness score could be around 7 out of 10**.

GPT-4 | Answer: 7

Since we don't have the actual resting heart rate values, we can't calculate the exact readiness score.
 However, based on the available information, we can conclude that the **readiness score will be 5**.

GPT-3.5 | Answer: 5

1. Steps: **Average steps per day: 4546 steps**. This is within the recommended range of 7000-10000 steps per day for adults.
 2. Burned Calories: **Average calories burned per day: 317 kcal. This is within the recommended range of 1800-2400 kcal per day for women.**
 3. Resting Heart Rate: **Average resting heart rate: 67 bpm**. This is within the normal range of 60-100 bpm for adults.
 4. Mood: The user's **mood is 3 out of 5, which is neutral**.
 Overall, the user's health metrics are within normal ranges, and her mood is neutral. Therefore, her **readiness score is 7 out of 10**.

Gemini-Pro | Answer: 7

Figure 3: A Case Study on Readiness Score Prediction (READ) from PMData dataset. Here, we display the responses from 1) our fine-tuned model, HealthAlpaca, 2) GPT-3.5, 3) GPT-4 and 4) Gemini-Pro. **Green Bolded texts** highlights the valid reasoning and **Red Bolded texts** highlights the false or irrelevant reasoning to the input.

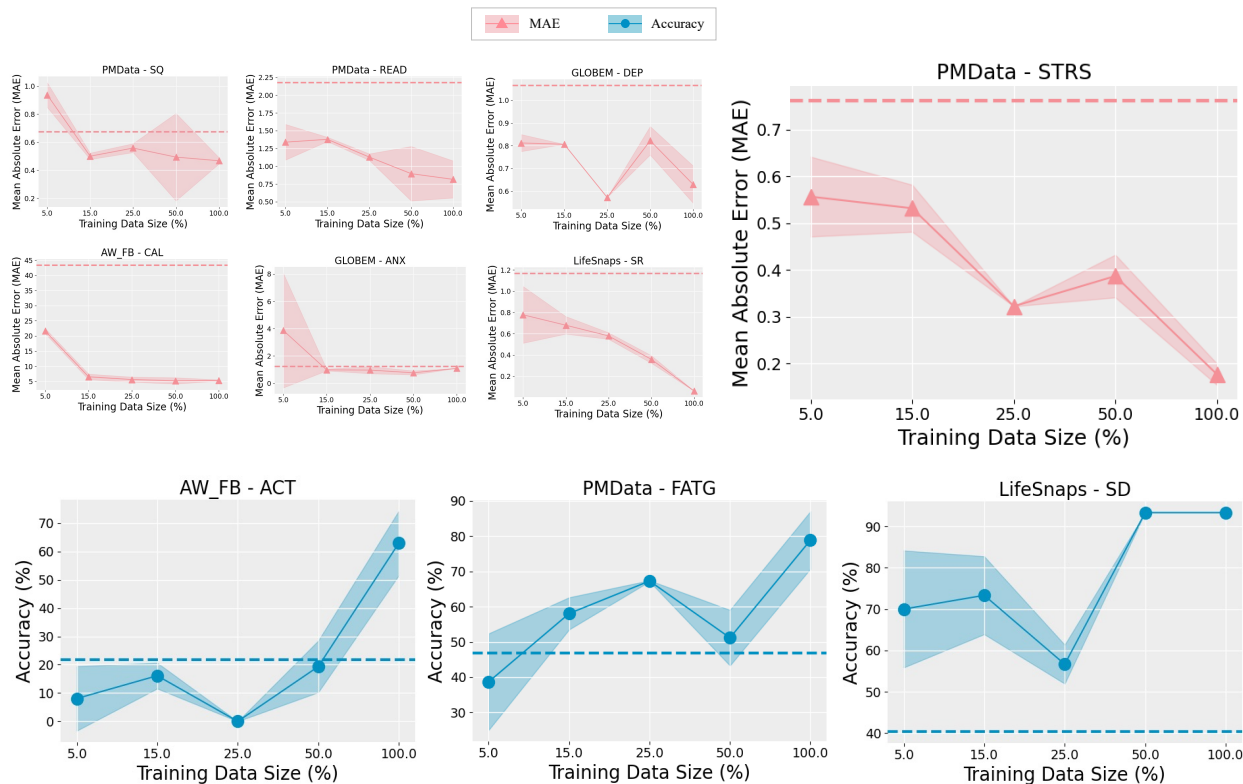


Figure 4: **Health Prediction Performance of Fully *fine-tuned* MedAlpaca with Different Training Sizes.** The instruction fine-tuning is conducted across ten tasks across four datasets. The solid lines represents the *fully fine-tuned* model’s performance whereas the dashed lines represents the *zero-shot* performance of MedAlpaca which serves as baselines. Note that the color indicates the metrics used to evaluate the tasks.