NEJM
AI

The AI Revolution in Medicine

Understanding How AI Can Advance Patient Care

NEJM
GROUP

NEJM
AI

November 2024

Dear Reader,

I hope you enjoy this timely collection of articles that explores the fast-moving impact of artificial intelligence (AI) and machine learning in medicine. Each article provides insight into the challenges, opportunities, and future directions of AI in medicine.

NEJM AI aims to provide a platform for rigorous evidence, resource sharing, and thoughtful discussions that will shape the integration of AI in medicine. If reading this collection inspires you to learn more about the range of AI applications in medicine, I invite you to visit and explore ai.nejm.org.

Sincerely,
Isaac S. Kohane, MD, PhD
Editor-in-Chief, NEJM AI

781.893.3800 | @NEJM_AI
860 WINTER STREET, WALTHAM, MA 02451-1413 USA | ai.nejm.org

NEJM Evidence is a product of NEJM Group, a division of the Massachusetts Medical Society.

Table of Contents

CASE STUDY
Validation of a Mobile App for Remote Autism Screening in Toddlers **1**

EDITORIAL
Advancing Autism Detection: A Digital Step Forward **11**

ORIGINAL ARTICLE
Artificial Intelligence–Powered Rapid Identification of ST-Elevation Myocardial Infarction via Electrocardiogram (ARISE) — A Pragmatic Randomized Controlled Trial **13**

EDITORIAL
AI-RISE to the Challenge — Artificial Intelligence Reduces Time to Treatment in STEMI **25**

DATASETS, BENCHMARKS AND PROTOCOLS
GPT versus Resident Physicians — A Benchmark Based on Official Board Scores **28**

PERSPECTIVE
Who’s Training Whom? **36**

PERSPECTIVE
Use of GPT-4 to Diagnose Complex Clinical Cases **39**

NEJM AI is a publication of NEJM Group, a division of the Massachusetts Medical Society.
©2024 Massachusetts Medical Society, All rights reserved.

NEJM AI

Published September 26, 2024
NEJM AI 2024; 1 (10)
DOI: 10.1056/AIcs2400510

CASE STUDY

Validation of a Mobile App for Remote Autism Screening in Toddlers

Pradeep Raj Krishnappa Babu, Ph.D.,[1] J. Matias Di Martino, Ph.D.,[1,2] Rachel Aiello, Ph.D.,[3,4] Brian Eichner, M.D.,[5] Steven Espinosa, B.Sc.,[6] Jennifer Green, Ph.D.,[3,4] Jill Howard, Ph.D.,[3,4] Sam Perochon, M.Sc.,[7] Marina Spanos, Ph.D.,[3,4] Saritha Vermeer, Ph.D.,[3,4] Geraldine Dawson, Ph.D.,[3,4] and Guillermo Sapiro, Ph.D.[1,8,9]

Received: May 17, 2024; Revised: July 11, 2024; Accepted: August 7, 2024; Published: September 26, 2024

Abstract

Early detection of autism is important for timely access to diagnostic evaluation and early intervention services, which improve children's outcomes. Despite the ability of clinicians to reliably diagnose autism in toddlers, diagnosis is often delayed. SenseToKnow is a mobile autism screening application (app) delivered on a smartphone or tablet that provides an objective and quantitative assessment of early behavioral signs of autism based on computer vision (CV) and machine learning (ML). This study examined the accuracy of SenseToKnow for autism detection when the app was downloaded and administered remotely at home by caregivers using their own devices. The SenseToKnow app was administered by caregivers of 620 toddlers between 16 and 40 months of age, 188 of whom were subsequently diagnosed with autism by expert clinicians. The app displayed strategically designed movies and a bubble-popping game on an iPhone or iPad while recording the child's behavioral responses through the device's front-facing camera and touch/inertial sensors. Recordings of the child's behavior were then automatically analyzed using CV. Multiple behavioral phenotypes were quantified and combined using ML in an algorithm for autism prediction. SenseToKnow demonstrated a high level of diagnostic accuracy with area under the receiver operating characteristic curve of 0.92, sensitivity of 83.0%, specificity of 93.3%, positive predictive value of 84.3%, and negative predictive value of 92.6%. Accuracy of the app for detecting autism was similar when administered on either a caregiver's iPhone or iPad. These results demonstrate that a mobile autism screening app based on CV can be delivered remotely by caregivers at home on their own devices and can provide a high level of accuracy for autism detection. Remote screening for autism potentially lowers barriers to autism screening, which could reduce disparities in early access to services and support and improve children's outcomes.

Dr. Dawson and Dr. Sapiro are senior authors of this article.

The author affliations are listed at the end of the article.

Dr. Dawson can be contacted at geraldine.dawson@duke.edu; Dr. Sapiro can be contacted at guillermo.sapiro@duke.edu or guillermos@princeton.edu.

Introduction

Behavioral signs of autism can be observed in children between 9 and 18 months of age. The signs can include paying less attention to people; a lack of response to being called by name; motor delays; and differences in facial expressions, including limited use of facial expressions to communicate.[1–4] The current standard of care is universal autism screening at 18–24 months using a caregiver questionnaire: the Modified

Read Article at ai.nejm.org

Back to Table of Contents

Checklist for Autism in Toddlers, Revised with Follow-Up (M-CHAT-R/F).[5]

A large study conducted in a primary care network found a lower level of accuracy of the M-CHAT-R/F for autism detection in girls, children of color, and non–English-speaking families.[6] Challenges with the M-CHAT-R/F include the need for literacy from caregivers, cultural differences in how caregivers interpret the screening questions, and the necessity of a follow-up clinical interview, which pediatricians often fail to administer.[7] There is a need for efficient, objective autism screening tools using technology for direct observation of early autism signs to enhance screening accuracy.

A prior autism screening study used eye-tracking equipment to assess social attention in 1863 12-to-48-month-old toddlers and reported a high level of specificity (98.0%) but a low level of sensitivity (17.0%).[8] Other studies have used mobile applications (apps) to assess differences between autistic and neurotypical children in emotional-mirroring and gaze-fixation patterns, but these apps have yet to be adequately assessed for early autism screening.[9-11]

The SenseToKnow app was designed to assess multiple autism-related behaviors in toddlers, including social attention, using an easy-to-use digital phenotyping tool running on an iPhone or iPad without any additional equipment. The app displays strategically designed movies and a bubble-popping game while recording the child's behavioral responses through the device's front-facing camera and touch/inertia sensors. The responses are automatically analyzed through computer vision (CV) and machine learning (ML). SenseToKnow detects and quantifies patterns of social attention/gaze,[12] head movements,[13,14] complexity in facial expressions,[15] blinking rate,[16] response to being called by name,[17] and motor behaviors.[18] We previously evaluated the accuracy of SenseToKnow for detecting autism when delivered on an iPad in pediatric primary care settings.[19] An ML algorithm trained on multiple digital phenotypes showed a high level of diagnostic accuracy with an area under the receiver operating characteristic curve (AUROC) of 0.90, sensitivity of 87.8%, and specificity of 80.8%.

Increasingly, there is a demand for health care delivery through telemedicine, which has the potential to reduce access barriers to screening, diagnosis, and treatment for a wide range of health-related conditions, including child development and behavioral disorders.[20,21] For example, Canvas Dx is a mobile health screening tool that integrates parent and clinician reports with manually coded video recordings of child behavior.[22] Similarly, ASDetect employs videos of children to assist caregivers in completing questionnaires about their child's development.[23]

In this study, we evaluate the accuracy of the fully automatic SenseToKnow app, based on direct observation and assessment of the child's behavior, when downloaded and administered by caregivers on their own devices in home settings. We conducted all components of the study remotely by recruiting participants through an online health care portal, offering eConsent, providing app instructions through YouTube or Zoom, and conducting standardized telehealth diagnostic evaluations through Zoom. To facilitate inclusion of non–English-speaking families, all components of the study, including the SenseToKnow app, were available in Spanish.

Methods

PARTICIPANTS

The study, conducted from September 2020 through December 2023, included 756 toddlers in good health between 16 and 40 months of age. Exclusion criteria were sensory or motor impairments that precluded sitting or viewing the app. Caregivers who spoke either English or Spanish were recruited through a web-based patient portal, email, phone, and/or clinician referrals. A total of 620 toddlers (82%) participated and completed all study measures. The recordings of the participants passed an initial CV analysis validity check (described below). Of the 620 participants, 247 screened positive on the M-CHAT-R/F or raised developmental concerns for their caregiver or health care provider. The participants were further assessed by one of the study's expert licensed psychologists using the Telemedicine-based ASD Evaluation Tool for Toddlers and Young Children (TELE-ASD-PEDS; see the Supplementary Appendix, Section S1). A total of 188 children were subsequently diagnosed with autism spectrum disorder and 31 were diagnosed with developmental delay without autism. The remaining 401 were considered neurotypical. Of these, 374 screened negative on the M-CHAT-R/F and neither the caregiver nor the child's health care provider expressed any developmental concerns, and 27 screened positive on the M-CHAT-R/F but were considered neurotypical after diagnostic evaluation. The participants who were either developmentally delayed or neurotypical were combined into a nonautism group of 432. Figure S1 provides a Consolidated Standards of Reporting Trials (CONSORT) diagram for the study. Caregivers provided informed consent (through eConsent) and the study was approved by the Duke University Health System Institutional Review Board (Pro00104066).

Table 1. Sample Demographic Characteristics.

Demographic Variable	N (%)		
	Total Sample (N=620; 100%)	Nonautism (N=432; 69.68%)	Autism (N=188; 30.32%)
Age in Months*			
Mean (SD)	20.20 (6.10)	17.27 (2.79)	26.96 (6.28)
Sex*			
Boys	348 (56.12%)	216 (50.00%)	132 (70.21%)
Girls	272 (43.87%)	216 (50.00%)	56 (29.78%)
Race*			
American Indian/Alaskan Native	1 (0.16%)	0 (0.00%)	1 (0.53%)
Asian	15 (2.42%)	8 (1.85%)	7 (3.72%)
Black or African American	92 (14.84%)	36 (8.33%)	56 (29.79%)
Native Hawaiian/Pacific Islander	0 (0.00%)	0 (0.00%)	0 (0.00%)
White	393 (63.39%)	322 (74.54%)	71 (37.77%)
More Than One Race	71 (11.45%)	49 (11.34%)	22 (11.70%)
Unknown/not reported	48 (7.74%)	17 (3.95%)	31 (16.49)
Ethnicity*			
Hispanic/Latinx	102 (16.45%)	51 (11.81%)	51 (27.13%)
Not Hispanic/Latinx	507 (81.77%)	374 (86.57%)	133 (70.74%)
Unknown/not reported	11 (1.77%)	7 (1.62%)	4 (2.12%)
Caregivers' Education Level*			
Without high school diploma	21 (3.39%)	5 (1.16%)	16 (8.51%)
High school diploma	40 (6.45%)	14 (3.24%)	26 (13.83%)
Some college education	57 (9.19%)	16 (3.70%)	41 (21.81%)
Four-year college degree or more	502 (80.96%)	397 (91.90%)	105 (55.85%)
M-CHAT-R/F			
Positive M-CHAT-R/F cases	209	43	166
Score (mean (SD))	3.00 (4.13)	1.09 (1.60)	7.34 (4.79)

*$P<0.05$ for autism versus nonautism comparison. ANOVA and chi-square tests were performed for various numerical and categorical (respectively) demographic variables. ANOVA denotes analysis of variance; M-CHAT-R/F, Modified Checklist for Autism in Toddlers, Revised with Follow-Up.

Table 1 shows the participants' demographic characteristics based on parent report. The sample was diverse and included approximately 15% Black or African American and 16% Hispanic/Latinx participants. Participants with autism were approximately 9 months older on average than participants without autism. Therefore, we also conducted an analysis with an age-matched subsample of children younger than 30 months (Table S1).

APP ADMINISTRATION AND STIMULI

Caregivers were instructed on how to use the SenseToKnow app through YouTube videos or Zoom, according to their preference. They were instructed to hold the child on their lap while the child watched the movies on an iPhone or iPad, which was positioned upright on a table approximately 60 cm away. The app took approximately 10 minutes to use and consisted of 11 brief developmentally appropriate movies and a bubble-popping game (Fig. 1A and Section S4). The device's front-facing camera recorded the toddlers' behavioral responses while they were watching the movies. The app produced a total of 11 video clips synchronized with each of the movies. During the bubble-popping game, caregivers placed the device flat on a table and, after a brief demonstration, children popped floating bubbles using the touch screen while the touch data were recorded.

FEATURE EXTRACTION

The child's video clips associated with each movie were captured at 30 frames per second, time synchronized with the movies, and analyzed to track the child's face. We extracted 49 facial landmarks and measured head-pose angles relative to the device (Fig. 1B).[24-26] Gaze features from the video were extracted using a deep neural network.[12] The facial landmarks and head-pose angles of each video clip were used to assess whether or not it was a valid administration

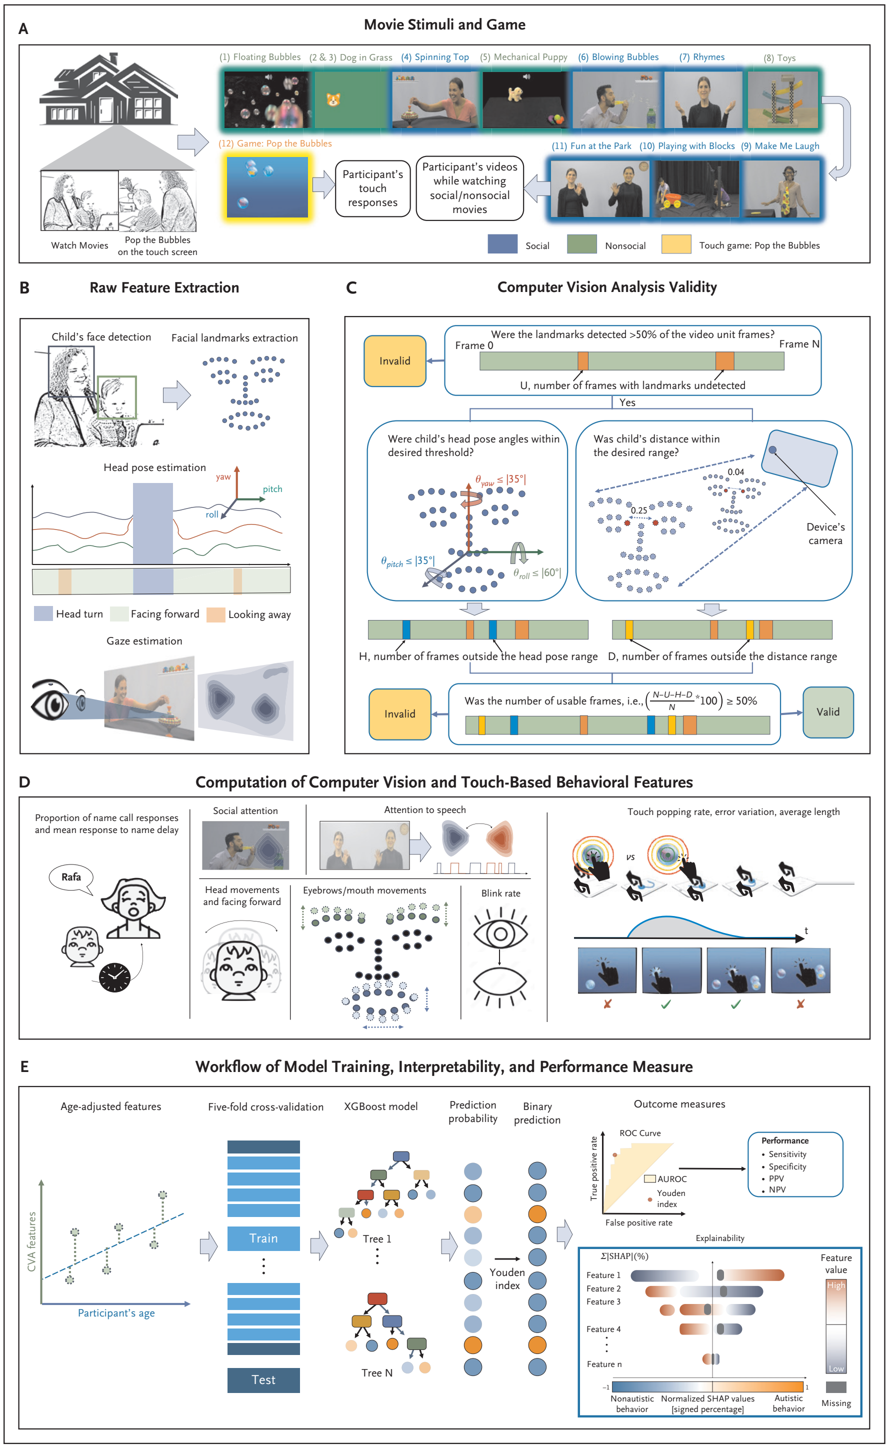
A
Movie Stimuli and Game
(1) Floating Bubbles
(2 & 3) Dog in Grass
(4) Spinning Top
(5) Mechanical Puppy
(6) Blowing Bubbles
(7) Rhymes
(8) Toys
(12) Game: Pop the Bubbles
(11) Fun at the Park
(10) Playing with Blocks
(9) Make Me Laugh
Participant's touch responses
Participant's videos while watching social/nonsocial movies
Watch Movies
Pop the Bubbles on the touch screen
Social
Nonsocial
Touch game: Pop the Bubbles
B
Raw Feature Extraction
Child's face detection
Facial landmarks extraction
Head pose estimation
yaw
pitch
roll
Head turn
Facing forward
Looking away
Gaze estimation
C
Computer Vision Analysis Validity
Invalid
Were the landmarks detected >50% of the video unit frames?
Frame 0
Frame N
U, number of frames with landmarks undetected
Yes
Were child's head pose angles within desired threshold?
$\theta_{yaw} \leq |35°|$
$\theta_{pitch} \leq |35°|$
$\theta_{roll} \leq |60°|$
Was child's distance within the desired range?
0.04
0.25
Device's camera
H, number of frames outside the head pose range
D, number of frames outside the distance range
Invalid
Was the number of usable frames, i.e., $\left(\frac{N-U-H-D}{N} * 100\right) \geq 50\%$
Valid
D
Computation of Computer Vision and Touch-Based Behavioral Features
Proportion of name call responses and mean response to name delay
Rafa
Social attention
Attention to speech
Head movements and facing forward
Eyebrows/mouth movements
Blink rate
Touch popping rate, error variation, average length
vs
t
✗
✓
✓
✗
E
Workflow of Model Training, Interpretability, and Performance Measure
Age-adjusted features
CVA features
Participant's age
Five-fold cross-validation
Train
Test
XGBoost model
Tree 1
Tree N
Prediction probability
Youden index
Binary prediction
Outcome measures
ROC Curve
True positive rate
False positive rate
AUROC
Youden index
Performance
Sensitivity
Specificity
PPV
NPV
Explainability
Σ|SHAP|(%)
Feature 1
Feature 2
Feature 3
Feature 4
Feature n
Feature value
High
Low
Missing
−1
1
Nonautistic behavior
Normalized SHAP values [signed percentage]
Autistic behavior

Figure 1. Illustration of the SenseToKnow App Workflow.

First, movies and the bubble-popping game are displayed (Panel A). Then, raw features are extracted using computer vision (Panel B). Next, the computer vision analysis validity algorithm is implemented (Panel C). Then, behavioral features are computed using the raw behavioral features elicited in response to movies (Panel D). Finally, as illustrated in Panel E, classification modeling is performed including feature age adjustment — fitting a linear regression model on each of the individual 23 features with respect to the participant age and estimating the residuals for each feature, model training, cross-validation, and feature importance to provide an explainable prediction and interpretation using SHAP values. AUROC denotes area under the receiver operating characteristic curve; CVA, computer vision analysis; NPV, negative predictive value; PPV, positive predictive value; ROC, receiver operating characteristic; SHAP, SHapley Additive exPlanation; and XGBoost, Extreme Gradient Boosting.

(Fig. 1C). If it was valid, the child's behavioral features were then extracted (Fig. 1D).

COMPUTER VISION ANALYSIS VALIDITY CHECK

The percentage of app administrations that were considered valid was high (89%). The CV analysis validity algorithm (Fig. 1C) assessed all 11 video clips. To ensure that we had enough usable frames to reliably estimate the behavioral features, the algorithm included only frames in which the child's facial landmarks were automatically detected, frames within the desired range of head-pose angles, and frames in which the child was within a desired distance from the screen (Fig. 1C; Section S5). To evaluate the accuracy of the CV analysis validity algorithm, we randomly selected 500 video clips from the study and conducted manual coding to quantify the number of usable frames.

For each video clip, if the usable frames satisfying the above three conditions were valid for more than 50% of the length of video for a given movie stimulus, then the clip was considered valid. In addition, if valid data for 6 out of the 11 clips were available, then the app administration was considered valid for extraction of behavioral features. Details on valid and invalid administrations with respect to the facial landmarks, head-pose angles, and distances from the device with corresponding pictorial representation are presented in Section S6 and Figure S2. Subsequently, we extracted 23 CV analysis-based validated behavioral features[19] from our past work[12-18] that correlated with clinical measures (Section S7).[27]

STATISTICAL ANALYSIS

Given that age could influence behavioral phenotypes in social-communication and sensory-motor domains in early stages of toddlerhood,[28,29] child age was accounted for by fitting a linear regression (Fig. 1(E)) and using the residuals of these linear models as inputs. Extreme Gradient Boosting (XGBoost)[30] was used to handle missing data (Section S8 and Fig. S3) and class imbalance. Classification performance was evaluated using the AUROC. Fivefold nested cross-validation was used to evaluate robust/unbiased classification performance.[31] The 95% confidence intervals (CIs) were computed using the Hanley and McNeil method.[32] The Youden optimality index (J=Sensitivity+Specificity−1) was used to estimate the algorithm operation point.[33] SHapley Additive exPlanation (SHAP) values were presented to gauge the influence of variables on the prediction and contribution to the design of explainable ML.[34] Positive predictive values (PPVs) and negative predictive values (NPVs) were calibrated for autism prevalence (1/36). All statistics were computed in Python 3.8.10.

Results

ACCURACY OF COMPUTER VISION ANALYSIS VALIDITY ALGORITHM

Human validity coding and the CV analysis validity check showed a high level of agreement (k=0.97), which was consistent across races and was invariant to device types (Section S6). Figure S2 shows the percentage of landmarks detected after CV analysis validity.

DIAGNOSTIC ACCURACY AND COMPARISONS ACROSS DEVICE TYPES

To account for different devices and different sizes of iPhones and iPads, screen size was added as a covariate. Using all 23 age-adjusted behavioral features and screen size, we trained the XGBoost model to classify the autism and nonautism groups.

Figure 2(A) shows the receiver operating characteristic curves for three models:

- Model 1 — all participants (188 autism; 432 nonautism)
- Model 2 — iPad only (101 autism; 133 nonautism)
- Model 3 — iPhone only (87 autism; 299 nonautism)

Diagnostic accuracy was similar across different devices. Specifically, accuracy values for Models 1, 2, and 3 were as follows: AUROCs were 0.92±0.01, 0.91±0.02, and 0.90±0.02; sensitivities were 83.0±5.4%, 86.1±6.7%, and 85.1±7.5%; and specificities were 93.3±2.4%, 87.2±5.6%, and 92.6±2.9%, respectively. These results replicate past findings from studies conducted in pediatric clinics using iPads.[19] Feature ablation analyses indicated that all features are essential (Section S9). The diagnostic accuracy of the app for children younger than 30 months of age was as follows: AUROC was 0.90±0.03, sensitivity was 78.0±7.6%, and specificity was 94.8±2.1%.

SHAPLEY ADDITIVE EXPLANATION ANALYSES

Figure 2B shows Model 1's summary SHAP value plot, which explains the relative importance of the prediction feature. Figure 3 shows the individualized SHAP plots using four different example participants along with each

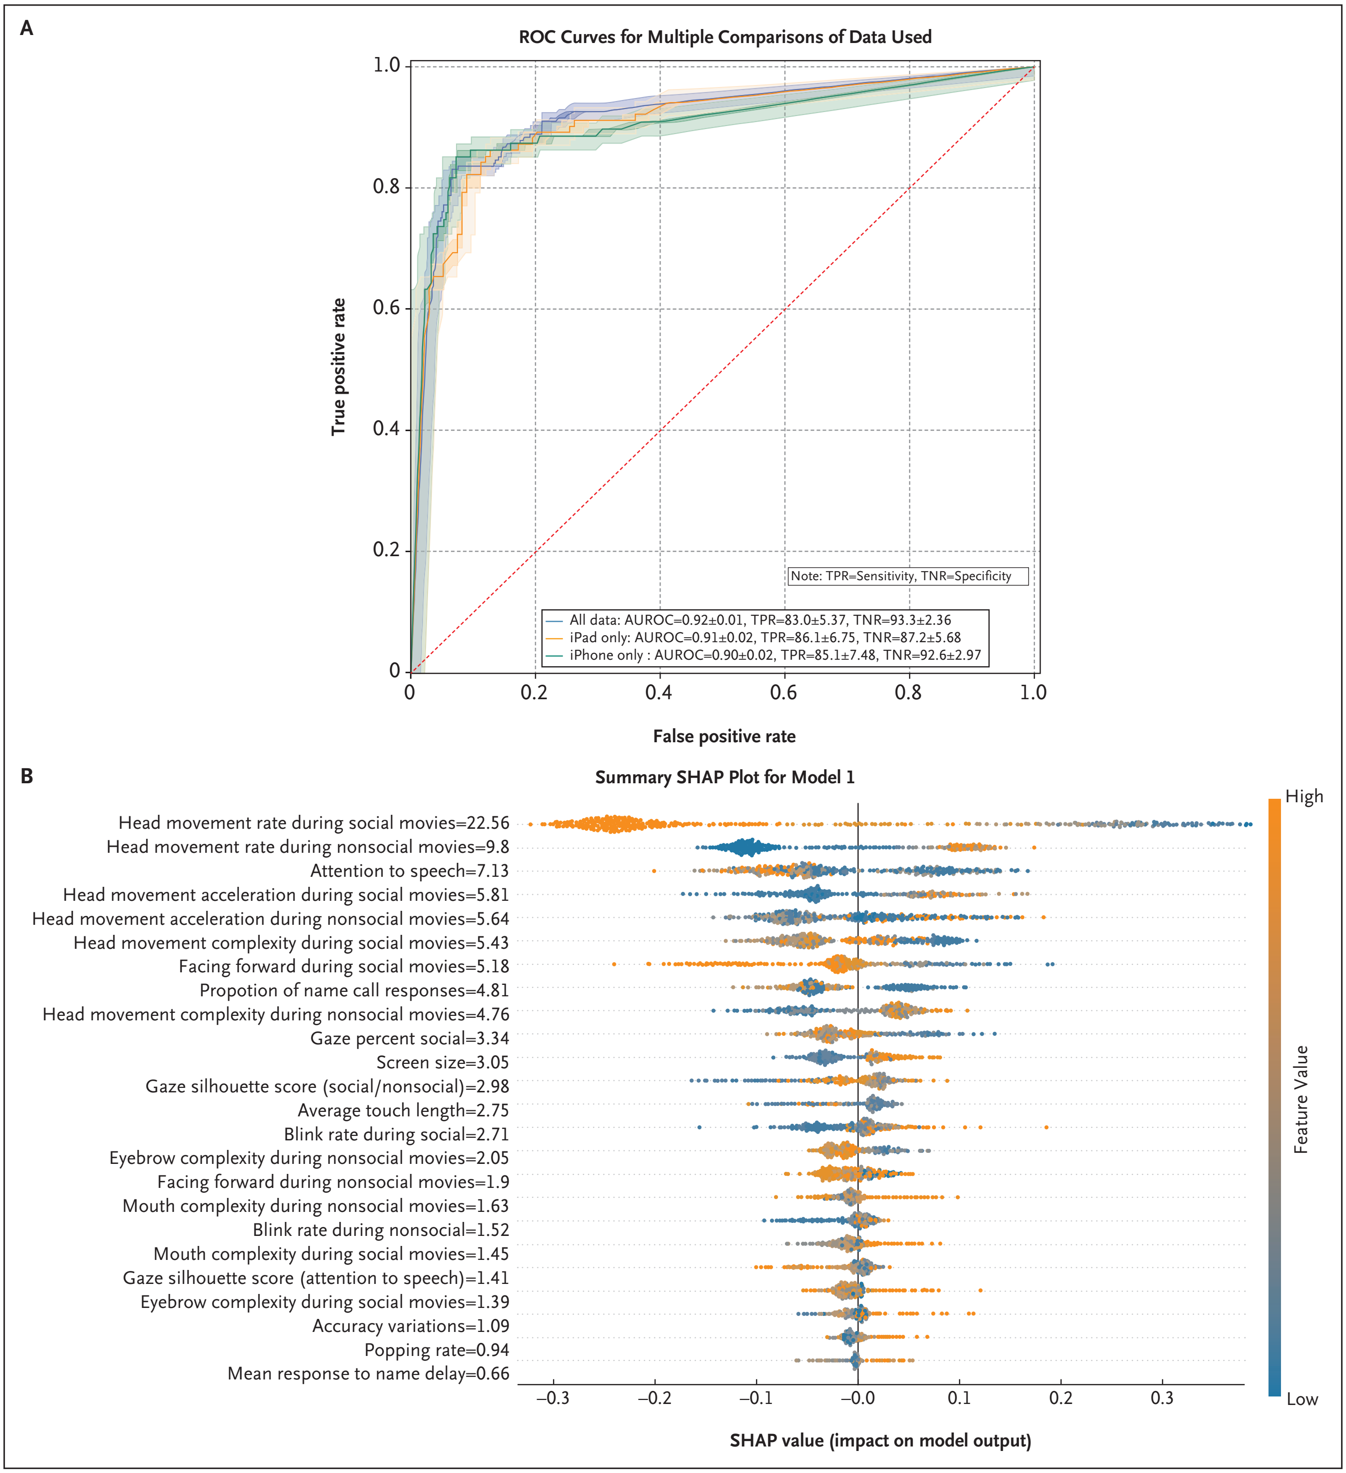

Figure 2. Accuracy of Remote Administration of SenseToKnow for Autism Detection Using Home Devices (iPad or iPhone).

Panel A shows ROC curves with the respective AUROCs, sensitivities, and specificities based on the whole sample (Model 1), iPad only (Model 2), and iPhone only (Model 3). Panel B shows SHAP values representing behavioral features' relative importance for Model 1. AUROC denotes area under the receiver operating characteristic curve; ROC, receiver operating characteristic; SHAP, SHapley Additive exPlanation; TNR, true negative rate; and TPR, true positive rate.

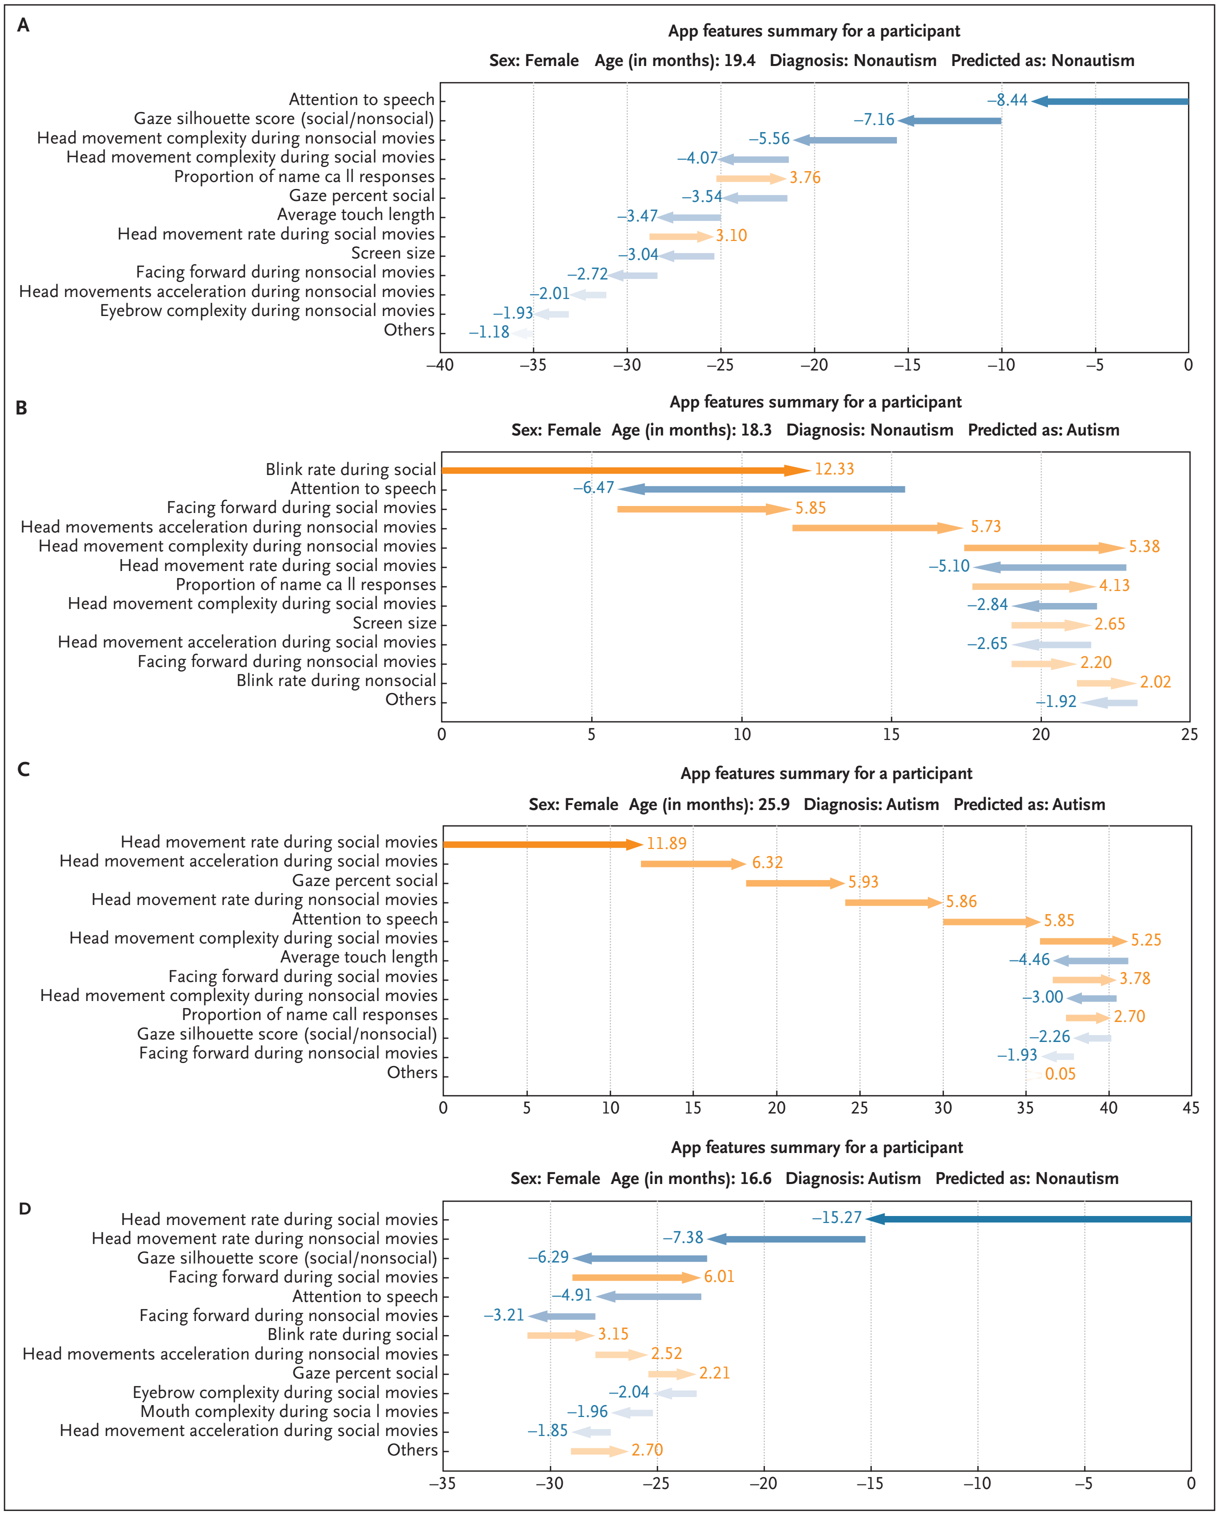

Figure 3. Examples of Individual SHAP Values Plots for Four Different Participants.

Panel A shows nonautism participants predicted as nonautism. Panel B shows nonautism participants predicted as autism. Panel C shows autism participants predicted as autism. Panel D shows autism participants predicted as nonautism. The direction and color of the arrows for each behavioral feature indicates whether the behavior is pointing in the autism (orange) or nonautism (blue) direction for the top 12 features. The remaining features were summed up as “other.”

Table 2. Accuracy of Autism Detection by Sex, Race, Ethnicity, and M-CHAT-R/F Score.

	Number of Participants		Performance					
Comparisons	Nonautism	Autism	AUROC±CI	Sensitivity±CI	Specificity±CI	PPV (Calibrated)	NPV (Calibrated)	F1 Score*
Autism vs. Nonautism								
Total sample	432	188	0.91±0.01	83.0±5.4%	93.3±2.4%	84.3 (28.0) %	92.6 (99.4) %	0.83
Sex								
Boys	216	56	0.91±0.02	78.6±7.1%	95.4±2.7%	81.5 (34.7) %	94.5 (99.7) %	0.80
Girls	216	132	0.91±0.02	84.8±6.1%	91.2±3.7%	85.5 (23.2) %	90.8 (99.5) %	0.85
Race								
White	322	71	0.91±0.02	80.1±9.9%	96.9±2.0%	84.4 (43.4) %	94.8 (99.2) %	0.80
Black/African American	36	56	0.86±0.03	85.7±9.1%	82.6±10.3%	87.3 (12.1) %	82.0 (99.4) %	0.86
Other race	30	47	0.91±0.02	86.7±10.2%	87.7±8.5%	78.8 (18.1) %	99.7 (99.5) %	0.83
Ethnicity								
Not Hispanic/ Latinx	374	133	0.91±0.02	82.7±6.4%	94.4±2.3%	84.0 (31.5) %	93.9 (99.4) %	0.83
Hispanic/Latinx	51	51	0.90±0.03	82.4±9.4%	88.2±8.4%	87.5 (18.0) %	83.4 (99.4) %	0.85
Covariate Analysis								
Sex, race, ethnicity	432	188	0.92±0.01	83.0±5.4%	93.3±2.4%	84.3 (28.0) %	92.6 (99.4) %	0.83
Using M-CHAT-R/F alone for classification	432	188	0.87±0.02	80.3±0.02	92.4±0.02	82.1 (24.7) %	91.5 (99.3) %	0.81

*F1 score combines both the precision (PPV) and recall (True positive rate (TPR), i.e., sensitivity) to provide a single score that balances precision and recall. It is especially useful when there is a need to offer optimal balance between precision and recall and when the class distribution is imbalanced. AUROC denotes area under the receiver operating characteristic curve; CI, confidence interval; M-CHAT-R/F, Modified Checklist for Autism in Toddlers, Revised with Follow-Up; NPV, negative predictive value; and PPV, positive predictive value.

feature's contribution toward the final predicted outcome, representing personalized explainability.

COMPARISONS BASED ON SEX, RACE, AND ETHNICITY

As shown in Table 2, diagnostic accuracy was comparable for boys and girls, white and Black children (the small sample size of Black children without autism affected the specificity and NPV calculations, but other operating points could be used for equalizing performance), and Hispanic/Latinx and non-Hispanic/Latinx children.

Discussion

Results of the present study demonstrate that a mobile app, SenseToKnow, can accurately detect autism when used by caregivers at home on their own devices. SenseToKnow demonstrated a high level of diagnostic accuracy for boys and girls and for Black and Hispanic/Latinx children.

Using a mobile app to screen for autism at home, combined with telehealth diagnostic assessments, can help reduce barriers to accessing services for families who live far from a pediatric office or whose work schedules or financial circumstances make traveling to a clinic difficult. Conducting screening and assessments in a child's familiar environment when the child is not tired or stressed may offer a more representative picture of the child's behavior. Longitudinal monitoring with the app can track the progression of a child's behavior remotely over time.

A limitation of the study is the potential sample bias introduced by the requirement of access to an iPhone or iPad. Although the study sample was diverse and was recruited through primary care, there remains a crucial need for studies with children from differing educational, racial, and ethnic backgrounds, as well as follow-up studies to account for behavioral features that could become apparent or change at an older age. Moreover, the sample of participants, recruited through an electronic patient portal and provider referral, may not reflect the general population targeted for universal screening. Future studies should examine whether there are ethnic and/or racial differences in caregiver attitudes in the use of a mobile app and the impact of co-occurring conditions, such as attention–deficit/hyperactivity disorder and intellectual disabilities.

This study demonstrates promising results for a mobile autism screening app that can be administered with a high level of accuracy at home by caregivers. Such an approach could reduce disparities in access to early detection and intervention, thereby improving children's outcomes.

Disclosures

Author disclosures and other supplementary materials are available at ai.nejm.org.

Data Availability

Data that support the findings of this study are available from the corresponding authors upon request and following Duke University Health System Institutional Review Board and privacy regulations.

We wish to thank the many caregivers and children for their participation in the study, without whom this research would not have been possible. We gratefully acknowledge members of the SenseToKnow study research team, including a clinical research coordinator, several clinical research specialists, and the Duke Center for Autism Data Management and Analysis Core team. This work was supported by the National Institute of Mental Health (NIMH) R01MH121329 (G.D. and G.S., coprincipal investigators), the National Institutes of Health Autism Centers of Excellence Award NICHD P50HD093074 (G.D., principal investigator), and NIMH R01MH120093 (G.D. and G.S., coprincipal investigators). S.P., J.M.D.M., and G.S. received additional support from the National Science Foundation and the Department of Defense (Office of Naval Research and National Geospatial-Intelligence Agency) as well as gifts from Cisco, Amazon Web Services, and Microsoft. The funders had no role in the study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Author Affiliations

[1] Department of Electrical and Computer Engineering, Duke University, Durham, NC
[2] Informatics and Computer Science, Catholic University of Uruguay, Montevideo, Uruguay
[3] Department of Psychiatry and Behavioral Sciences, Duke University, Durham, NC
[4] Duke Center for Autism and Brain Development, Duke University, Durham, NC
[5] Department of Pediatrics, Duke University, Durham, NC
[6] Office of Information Technology, Duke University, Durham, NC
[7] Ecole Normale Supérieure Paris-Saclay, Gif-Sur-Yvette, France
[8] Department of Biomedical Engineering, Mathematics, and Computer Sciences, Duke University, Durham, NC
[9] Electrical and Computer Engineering, Princeton University, Princeton, NJ

References

1. Chawarska K, Volkmar F, Klin A. Limited attentional bias for faces in toddlers with autism spectrum disorders. Arch Gen Psychiatry 2010;67:178-185. DOI: 10.1001/archgenpsychiatry.2009.194
2. Miller M, Iosif AM, Hill M, et al. Response to name in infants developing autism spectrum disorder: a prospective study. J Pediatr 2017;183:141-146. DOI: 10.1016/j.jpeds.2016.12.071
3. Bradshaw J, McCracken C, Pileggi M, et al. Early social communication development in infants with autism spectrum disorder. Child Dev 2021;92:2224-2234. DOI: 10.1111/cdev.13683
4. Iverson JM, Shic F, Wall CA, et al. Early motor abilities in infants at heightened versus low risk for ASD: a baby siblings research consortium (BSRC) study. J Abnorm Psychol 2019;128:69-80. DOI: 10.1037/abn0000390
5. Robins DL, Casagrande K, Barton M, Chen CMA, Dumont-Mathieu T, Fein D. Validation of the modified checklist for autism in toddlers, revised with follow-up (M-CHAT-R/F). Pediatrics 2014;133:37-45. DOI: 10.1542/peds.2013-1813
6. Guthrie W, Wallis K, Bennett A, et al. Accuracy of autism screening in a large pediatric network. Pediatrics 2019;144(4). DOI: 10.1542/peds.2018-3963
7. Wallis KE, Guthrie W, Bennett AE, et al. Adherence to screening and referral guidelines for autism spectrum disorder in toddlers in pediatric primary care. PLoS One 2020;15:e0232335. DOI: 10.1371/journal.pone.0232335
8. Wen TH, Cheng A, Andreason C, et al. Large scale validation of an early-age eye-tracking biomarker of an autism spectrum disorder subtype. Sci Rep 2022;12:1-13. DOI: 10.1038/s41598-022-08102-6
9. Deveau N, Washington P, Leblanc E, et al. Machine learning models using mobile game play accurately classify children with autism. Intell Med 2022;6:100057. DOI: 10.1016/j.ibmed.2022.100057
10. Varma M, Washington P, Chrisman B, et al. Identification of social engagement indicators associated with autism spectrum disorder using a game-based mobile app: comparative study of gaze fixation and visual scanning methods. J Med Internet Res 2022;24(2). DOI: 10.2196/31830
11. Washington P, Kalantarian H, Kent J, et al. Improved digital therapy for developmental pediatrics using domain-specific artificial intelligence: machine learning study. JMIR Pediatr Parent 2022;5(2). DOI: 10.2196/26760
12. Chang Z, Di Martino JM, Aiello R, et al. Computational methods to measure patterns of gaze in toddlers with autism spectrum disorder. JAMA Pediatr 2021;175:827-836. DOI: 10.1001/jamapediatrics.2021.0530
13. Dawson G, Campbell K, Hashemi J, et al. Atypical postural control can be detected via computer vision analysis in toddlers with autism spectrum disorder. Sci Rep 2018;8:1-7. DOI: 10.1038/s41598-018-35215-8

14. Krishnappa Babu PR, Di Martino JM, Chang Z, et al. Complexity analysis of head movements in autistic toddlers. J Child Psychol Psychiatry Allied Discip 2023;64:156-166. DOI: 10.1111/jcpp.13681

15. Babu PRK, Matias Di Martino J, Chang Z, et al. Exploring complexity of facial dynamics in autism spectrum disorder. IEEE Trans Affect Comput 2023;14:919-930. DOI: 10.1109/taffc.2021.3113876

16. Krishnappa Babu PR, Aikat V, Di Martino JM, et al. Blink rate and facial orientation reveal distinctive patterns of attentional engagement in autistic toddlers: a digital phenotyping approach. Sci Rep 2023;13:1-11. DOI: 10.1038/s41598-023-34293-7

17. Perochon S, Di Martino M, Aiello R, et al. A scalable computational approach to assessing response to name in toddlers with autism. J Child Psychol Psychiatry 2021;62:1120-1131. DOI: 10.1111/jcpp.13381

18. Perochon S, Matias Di Martino J, Carpenter KLH, et al. A tablet-based game for the assessment of visual motor skills in autistic children. NPJ Digit Med 2023;6:1-13. DOI: 10.1038/s41746-023-00762-6

19. Perochon S, Di Martino JM, Carpenter KLH, et al. Early detection of autism using digital behavioral phenotyping. Nat Med 2023;29:2489-2497. DOI: 10.1038/s41591-023-02574-3

20. Hanlon P, Daines L, Campbell C, et al. Telehealth interventions to support self-management of long-term conditions: a systematic metareview of diabetes, heart failure, asthma, chronic obstructive pulmonary disease, and cancer. J Med Internet Res 2017;19:e6688. DOI: 10.2196/jmir.6688

21. Chandrasekaran R. Telemedicine in the post-pandemic period: understanding patterns of use and the influence of socioeconomic demographics, health status, and social determinants. Telemed e-Health 2024;30:480-489. DOI: 10.1089/tmj.2023.0277

22. Kanne SM, Carpenter LA, Warren Z. Screening in toddlers and preschoolers at risk for autism spectrum disorder: evaluating a novel mobile-health screening tool. Autism Res 2018;11:1038-1049. DOI: 10.1002/aur.1959

23. Wong CM, Mohd Zambri N, Fan HH, et al. A direct comparison of three screening methods for autism spectrum disorder in a high-likelihood sibling population. J Autism Dev Disord 2024:1-12. DOI: 10.1007/s10803-024-06294-6

24. Ren X, Lattas A, Gecer B, et al. Facial geometric detail recovery via implicit representation. 2023 IEEE 17th international conference on automatic face gesture recognition, FG 2023, 2023:1–8. DOI: 10.1109/FG57933.2023.10042505.

25. De La Torre F, Chu WS, Xiong X, et al. IntraFace. 2015 11th IEEE international conference and workshops on automatic face and gesture recognition, FG 2015, 2015:1–8. DOI: 10.1109/FG.2015.7163082.

26. Hashemi J, Dawson G, Carpenter KLH, et al. Computer vision analysis for quantification of autism risk behaviors. IEEE Trans Affect Comput 2021;12:215-226. DOI: 10.1109/taffc.2018.2868196

27. Coffman M, Di Martino JM, Aiello R, et al. Relationship between quantitative digital behavioral features and clinical profiles in young autistic children. Autism Res 2023;16:1360-1374. DOI: 10.1002/aur.2955

28. Ozonoff S, Young GS, Belding A, et al. The broader autism phenotype in infancy: when does it emerge? J Am Acad Child Adolesc Psychiatry 2014;53:398-407.e2. DOI: 10.1016/j.jaac.2013.12.020

29. Lord C, Bishop S, Anderson D. Developmental trajectories as autism phenotypes. Am J Med Genet Part C Semin Med Genet 2015;169:198-208. DOI: 10.1002/ajmg.c.31440

30. Chen T, Guestrin C. XGBoost: a scalable tree boosting system. Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, August 13-17, 2016:785–794. DOI: 10.1145/2939672.2939785.

31. Vabalas A, Gowen E, Poliakoff E, et al. Machine learning algorithm validation with a limited sample size. PLoS One 2019;14(11). DOI: 10.1371/journal.pone.0224365

32. Hanley JA, McNeil BJ. The meaning and use of the area under a receiver operating characteristic (ROC) curve. Radiology 1982;143:29-36. DOI: 10.1148/radiology.143.1.7063747

33. Perkins NJ, Schisterman EF. The Youden Index and the optimal cutpoint corrected for measurement error. Biometrical J 2005;47:428-441. DOI: 10.1002/bimj.200410133

34. Topuz K, Bajaj A, Abdulrashid I. Interpretable machine learning. Proceedings of the annual Hawaii international conference on system sciences, 2023;1236–1237. https://aisel.aisnet.org/hicss-56/da/machine_learning/1.

Back to Table of Contents

NEJM AI

Published September 26, 2024
NEJM AI 2024; 1 (10)
DOI: 10.1056/AIe2400811

EDITORIAL

Advancing Autism Detection: A Digital Step Forward

Isaac Kohane, M.D., Ph.D.[1,2]

Received: August 13, 2024; Accepted: August 13, 2024; Published: September 26, 2024

Abstract

A novel mobile application designed to screen for autism in toddlers (16–40 months) produces promising results, but also underscores the growing importance of digital phenotyping in medical research and clinical practice.

The timely detection of autism spectrum disorder (ASD) remains a critical challenge in pediatric care. Despite the well-established benefits of early intervention, many children with ASD are not diagnosed until well past the optimal window for initiating treatments such as applied behavior analysis.[1] In this issue of *NEJM AI*, Krishnappa Babu and colleagues present promising results for a novel mobile application, SenseToKnow, designed to remotely screen for autism in toddlers (16–40 months).[2]

The study represents an important incremental advance in our ability to identify children with ASD. By enabling remote caregiver screening at home, SenseToKnow could help overcome barriers to access, particularly for families in underserved areas or those facing logistical challenges in attending clinic appointments. The authors demonstrate that their computer vision and machine learning–based app achieves high diagnostic accuracy (area under the curve=0.92) in detecting autism. This level of performance, comparable to previous clinic-based assessments, suggests that remote screening could become a valuable tool in expanding access to early autism detection.

However, it is essential to view these results in context. The study population was enriched for autism, with a prevalence of 30% — far higher than the general population. While this enrichment is necessary for initial validation, it limits our ability to extrapolate the app's performance in real-world settings. Further testing in populations with more representative autism prevalence will be crucial to establish the app's actual clinical utility and to understand its performance characteristics, particularly its positive predictive value, in real-world scenarios.

Moreover, as we embrace these technological advances, we must not lose sight of the heterogeneity within ASDs. As I have noted previously,[3] by balancing genetics-first and phenotype-first approaches, ASD encompasses a wide range of clinical presentations and there are several potential underlying mechanisms. The SenseToKnow app, which focuses on behavioral phenotypes, is a valuable tool in our diagnostic toolkit. However, it is just one piece of a complex puzzle.

The author affliation is listed at the end of the article.

Dr. Kohane can be contacted at isaac_kohane@harvard.edu or at Harvard University, Cambridge, MA.

Read Article at ai.nejm.org

Back to Table of Contents

Looking to the future, we can anticipate even more sophisticated diagnostic tools and AI programs with access to comprehensive patient data — including genomics, detailed clinical histories, and even environmental exposures — which could further accelerate and refine autism diagnosis. Such tools could potentially identify distinct ASD subgroups, each with its own optimal intervention strategies.

The study by Krishnappa Babu et al. also underscores the growing importance of digital phenotyping in medical research and clinical practice. By leveraging smartphone technology and machine learning, we can now capture and analyze behavioral data at a previously unattainable scale and level of detail. This approach not only aids in diagnosis but also opens new avenues for understanding the diverse manifestations of ASD and tracking treatment responses over time.

While celebrating this progress, we must also acknowledge the ethical considerations accompanying such technologies. Ensuring data privacy, addressing potential biases in AI algorithms, and maintaining the crucial role of clinical judgment in diagnosis are paramount as we integrate these tools into practice.

In conclusion, the SenseToKnow app represents a significant step forward in our ability to screen for autism efficiently and at scale. Its potential to increase access to early detection is promising, particularly for underserved populations. However, further validation in diverse, real-world settings is essential. As we continue to advance our diagnostic capabilities, we must strive for a balanced approach that integrates cutting-edge technology with comprehensive clinical assessment and an appreciation for the complex, multifaceted nature of ASDs.

The journey toward optimal autism care is ongoing, but studies like this by Krishnappa Babu and colleagues light the way forward. By embracing innovation while maintaining rigorous scientific standards, we can hope to provide earlier, more precise diagnoses and interventions for children with ASD, ultimately improving their lives and those of their families.

Disclosures

Author disclosures are available at ai.nejm.org.

Author Affiliations

[1] Department of Biomedical Informatics, Harvard Medical School, Boston, MA

[2] NEJM AI, Boston, MA

References

1. Tiura M, Kim J, Detmers D, Baldi H. Predictors of longitudinal ABA treatment outcomes for children with autism: a growth curve analysis. Res Dev Disabil 2017;70:185-197. DOI: 10.1016/j.ridd.2017.09.008.
2. Babu P, Martino J, Aiello R, et al. Validation of a mobile app for remote screening in toddlers. NEJM AI 2024;1(10). DOI: 10.1056/AIcs2400510.
3. Kohane IS. Finding a new balance between a genetics-first or phenotype-first approach to the study of disease. Neuron 2021;109:2216-2219. DOI: 10.1016/j.neuron.2021.07.001.

Back to Table of Contents

NEJM AI 2024; 1 (7)
DOI: 10.1056/AIoa2400190

ORIGINAL ARTICLE

Artificial Intelligence–Powered Rapid Identification of ST-Elevation Myocardial Infarction via Electrocardiogram (ARISE) — A Pragmatic Randomized Controlled Trial

Chin Lin, Ph.D.,[1,2,3,4] Wei-Ting Liu, M.D.,[5] Chiao-Hsiang Chang, M.D.,[5] Chiao-Chin Lee, M.D.,[5] Shi-Chue Hsing, M.D.,[5] Wen-Hui Fang, M.D.,[2,6] Dung-Jang Tsai, Ph.D.,[1,2,7] Kai-Chieh Chen, M.S.,[8] Chun-Ho Lee, M.S.,[3] Cheng-Chung Cheng, M.D.,[5] Yi-Jen Hung, M.D.,[9] Shih-Hua Lin, M.D.,[10] Chien-Sung Tsai, M.D.,[11] and Chin-Sheng Lin, M.D., Ph.D.[1,5]

Received: February 21, 2024; Revised: April 7, 2024; Accepted: April 30, 2024; Published: June 27, 2024

Abstract

BACKGROUND Timely diagnosis of ST-elevation myocardial infarction (STEMI) is crucial for the treatment of patients with acute coronary syndrome. Artificial intelligence–enabled electrocardiogram (AI-ECG) has shown potential for the accurate and timely detection of STEMI on 12-lead electrocardiograms (ECGs). However, its impact on clinical treatment times is unknown.

METHODS To evaluate the potential of AI-ECG–assisted detection of STEMI to reduce treatment delays for patients with STEMI, we conducted an open-label, cluster randomized controlled trial involving 43,234 eligible patients (mean age, 60 years; 49.5% male) without a history of coronary angiography within 3 days in the emergency department or inpatient wards at Tri-Service General Hospital, Taipei, Taiwan between May 1, 2022, and April 31, 2023. Patients were randomly assigned 1:1 to AI-ECG–assisted detection of STEMI (intervention group) or to standard of care (control group). The primary end point was door-to-balloon time; ECG-to-balloon time was also evaluated as a branch of the primary analysis. Secondary end points included incidence of new-onset low ejection fraction, cardiac death, and all-cause mortality.

RESULTS Among the 43,234 patients, 77 in the intervention group and 68 in the control group were diagnosed with STEMI with occluded vessel(s) based on coronary angiography. The use of AI-ECG demonstrated a positive predictive value of 89.5% (95% confidence interval [CI], 85.3 to 93.6%) and a negative predictive value of 99.9% (95% CI, 99.9 to 100.0%). For patients in the emergency department, the median door-to-balloon time was 82.0 minutes (interquartile range, 62.5 to 89.5) in the intervention group compared with

The author affiliations are listed at the end of the article.

Dr. Chin-Sheng Lin can be contacted at littlelincs@gmail.com or at the Division of Cardiology, Tri-Service General Hospital, No. 325, Sec. 2, Cheng-Kung Rd., Neihu 114, Taipei, Taiwan.

Read Article at ai.nejm.org

Back to Table of Contents

96.0 minutes (interquartile range, 78.0 to 137.0) in the control group (P=0.002). When analyzing both emergency and inpatient cases, the median ECG-to-balloon time was 78.0 minutes (interquartile range, 56.9 to 88.2 minutes) in the intervention group compared with 83.6 minutes (interquartile range, 72.7 to 127.8 minutes) in the control group (P=0.011). In the intervention group versus the control group, there were 340 versus 304 cases, respectively, of new-onset heart failure with reduced ejection fraction (odds ratio, 1.12; P=0.151), 85 versus 116 cases of cardiac death (odds ratio, 0.73; P=0.029), and 1153 versus 1127 cases of all-cause mortality (odds ratio, 1.02; P=0.568).

CONCLUSIONS In patients with STEMI, AI-ECG–assisted triage of STEMI decreased the door-to-balloon time for patients presenting to the emergency department and decreased the ECG-to-balloon time for patients in the emergency room and inpatients. (Funded by the National Science and Technology Council, Taiwan and others; ClinicalTrials.gov number, NCT05118009.)

Introduction

Acute coronary syndrome, particularly ST-segment elevation myocardial infarction (STEMI), represents a substantial health care burden and contributes to global morbidity and mortality.[1] Timely diagnosis and immediate initiation of primary percutaneous coronary intervention (PPCI) are essential for improving the prognosis of patients with STEMI.[2] However, distinguishing patients with STEMI from those with undifferentiated chest pain remains a clinical challenge in acute settings. Inexperienced physicians may exhibit reduced accuracy in diagnosing STEMI, potentially leading to misdiagnoses,[3,4] which are observed in approximately 20.5% of STEMI cases and are related to poorer prognoses.[5] Providing clinical support to frontline physicians is of paramount importance for optimizing the management of STEMI.

Delayed treatment, stemming from a combination of systematic and nonsystematic factors, is independently associated with increased mortality in PPCI-treated patients with STEMI.[6] Because nonsystematic issues, such as cardiac arrests and endotracheal tube intubation, are challenging to address,[7] minimizing systematic errors is essential for improving health care quality.[8]

Clinical decision support systems (CDSSs) are extensively used to optimize workflows and improve patient outcomes.[9] However, although commercial electrocardiogram (ECG) machines typically encompass an automatic analysis system, which includes a diagnostic function for STEMI, diagnostic accuracy is usually poor.[10] Incorporating such low positive predictive values into an automatic alarm system may pose a risk to patient safety because of the potential for alert fatigue.[11,12]

With the advent of deep learning techniques, artificial intelligence (AI) systems have demonstrated significant benefits in ECG interpretation.[13,14] The integration of artificial intelligence–enabled electrocardiogram (AI-ECG) into CDSS has been confirmed through a randomized controlled trial (RCT), highlighting its potential for diagnosing asymptomatic left ventricular dysfunction and reducing mortality.[15,16] We hypothesize that AI-ECG–based CDSS can also be applied to enhance STEMI management.

Previous studies have shown that the performance of AI-ECGs developed for STEMI identification has generally reached or exceeded the expertise of cardiologists.[10,17,18] Subsequent before-and-after analyses have demonstrated the effectiveness of AI-ECG–based CDSS in reducing door-to-balloon time.[19,20] It has been suggested that health care quality improvement over time could potentially influence the findings from before-and-after analyses.[21,22] In addition, patients with STEMI in the inpatient department exhibit a higher risk of mortality compared with those experiencing STEMI outside of the hospital setting.[23,24] This finding could potentially be attributed to greater delays in treatment activation in the inpatient department than in the emergency department.[25] Currently, few RCTs have evaluated the impact of AI-ECG systems in STEMI management,[26] whether in the emergency or inpatient department. We designed the Artificial Intelligence–Powered Rapid Identification of ST-Elevation Myocardial Infarction via Electrocardiogram (ARISE) trial to assess the impact of AI-ECG in facilitating STEMI diagnosis and management.

Methods

TRIAL DESIGN

The two-center, open-label, cluster randomized controlled ARISE trial (NCT05118009) followed the A/B testing methodology, whereby different software versions are randomly assigned to users, which aligns with the pragmatic RCT approach.[27] The study adhered to Consolidated Standards of Reporting Trials (CONSORT)-AI Extension guidelines for reporting (CONSORT-AI Extension checklist)[28] and was approved by the institutional review board (IRB) at Tri-Service General Hospital, Taipei, Taiwan (IRB A202105120). Informed consent was obtained from all 20 on-duty cardiologists in the hospital's catheterization laboratory who participated in the study. The ethical committee permitted the enrollment of patients during the trial period without consent given the need for timeliness of emergent procedures. Additional information is available in the protocol provided with the full text of this article at ai.nejm.org.

The ARISE trial was conducted at both an academic medical center and a community hospital in Taiwan, both of which shared the catheterization laboratory. Patients with STEMI visiting the community hospital were required to be referred to the academic medical center. Although these patients were not considered to be study participants, patient-level data from electronic health records (EHRs) were analyzed to investigate the impact of AI-ECG support on on-duty cardiologists. The ethical committee concluded that AI-ECG software qualified as a medical device with minimal risk, in accordance with the announcement by the Taiwan Food and Drug Administration (Taiwan Food and Drug Administration document 1101603684).

PATIENT DATA AND RANDOMIZATION

The ARISE trial involved a total of 43,994 patients without a history of coronary angiography who received an ECG in the emergency department or inpatient department at Tri-Service General Hospital between May 1, 2022, and April 31, 2023 (Fig. 1). Patients were randomly assigned 1:1 to AI-ECG–assisted detection of STEMI (intervention group) or standard of care (control group) according to the date of their first ECG such that those who had their first ECG on odd dates were assigned to one group and those with first their ECG on even dates were assigned to the other group. The simple randomization method ensured that only a single sequence of random assignments,[29] which was generated by an independent database programmer before the trial, was used to ensure blindness from the previous day.

For the on-duty cardiologists participating in the trial, the specific assignment to either the intervention or control group was revealed at 8 a.m. on the respective day. Only the first ECG of each patient during the study period was included for analysis. Initially, the intervention and control groups consisted of 21,989 and 22,005 patients, respectively. After excluding 760 patients younger than 18 years of age, the final analysis included 21,612 patients in the intervention group and 21,622 patients in the control group.

AI-ECG INTERVENTION

The AI algorithm used 12-lead ECG waveform data to identify STEMI.[10] The algorithm was reported to have a positive predictive value of 93.2% in a preliminary prospective study in an emergency department.[20] In the current study, cardiologists on duty were assigned to either the AI-assisted group or the control group daily, and all were aware of whether or not they would receive support from the AI-ECG system. Frontline physicians did not participate in the study and were blinded to the daily randomization.

In the intervention group, real-time analysis was performed by the AI-ECG system on all ECGs completed that day (details are shown in Supplementary Method 1 in the Supplementary Appendix). Immediate short message service (SMS) notifications, including ECG images, were sent to the on-duty cardiologists when the AI-ECG system detected potential STEMI cases to allow for review and confirmation. Given the lack of real-time documentation of ischemic-related symptoms in EHRs, on-duty cardiologists needed to assess patient symptoms upon receiving the AI-ECG alert. When STEMI was confirmed, the cardiologists could then activate the catheterization laboratory for PPCI.

In the control group, potential patients with STEMI were initially assessed by frontline physicians, who then notified the on-duty cardiologists for confirmation. Regardless of whether the AI-ECG system was used, frontline physicians were able to request consultation from the on-duty cardiologists. All frontline physicians could see the interpretation from a Philips automatic ECG analysis system for ECG interpretation, although the Philips system did not trigger subsequent SMS notifications to on-duty cardiologists because of concerns of alert fatigue.[11,12] Only on-duty cardiologists had the authority for catheterization laboratory activation, according to our national policy.

STEMI DIAGNOSIS AND BASELINE CHARACTERISTICS

The diagnosis of STEMI in this trial was made according to ischemic-related symptoms and ST-segment elevation

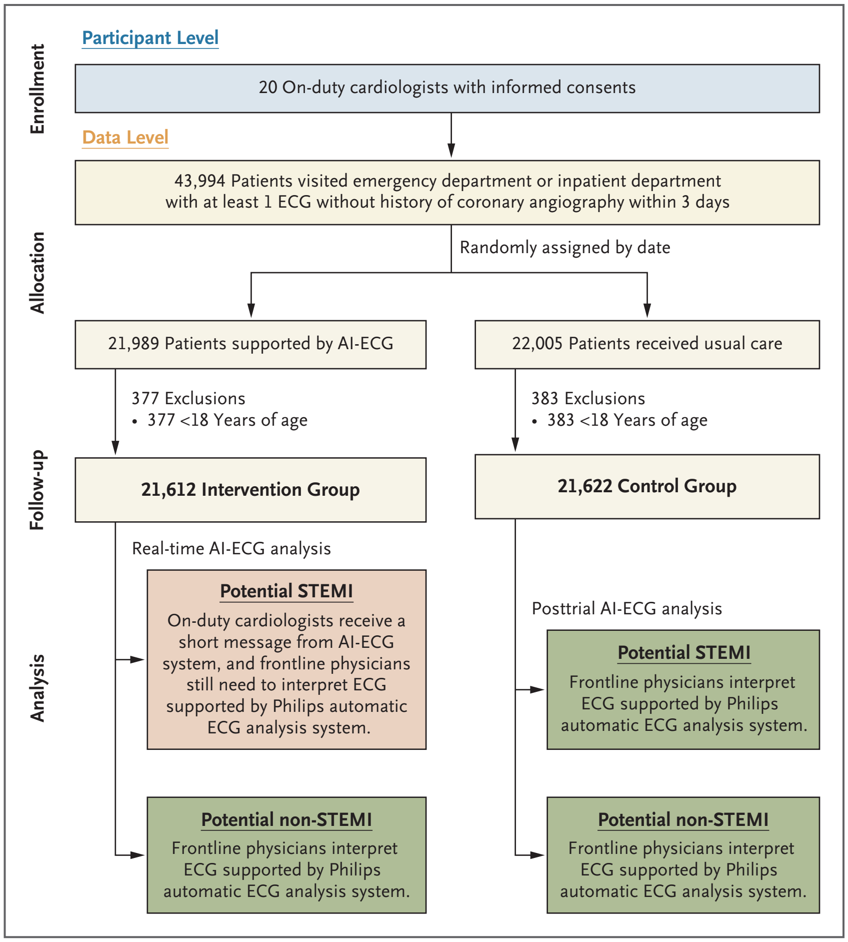
Enrollment
Allocation
Follow-up
Analysis
Participant Level
20 On-duty cardiologists with informed consents
Data Level
43,994 Patients visited emergency department or inpatient department with at least 1 ECG without history of coronary angiography within 3 days
Randomly assigned by date
21,989 Patients supported by AI-ECG
22,005 Patients received usual care
377 Exclusions
• 377 <18 Years of age
383 Exclusions
• 383 <18 Years of age
21,612 Intervention Group
21,622 Control Group
Real-time AI-ECG analysis
Posttrial AI-ECG analysis
Potential STEMI
On-duty cardiologists receive a short message from AI-ECG system, and frontline physicians still need to interpret ECG supported by Philips automatic ECG analysis system.
Potential STEMI
Frontline physicians interpret ECG supported by Philips automatic ECG analysis system.
Potential non-STEMI
Frontline physicians interpret ECG supported by Philips automatic ECG analysis system.
Potential non-STEMI
Frontline physicians interpret ECG supported by Philips automatic ECG analysis system.

Figure 1. CONSORT-AI Flow Diagram.

Of note, in the intervention group, there were 57 patients (0.3%) who underwent ECG examination while our AI-ECG system was inoperative (postanalysis revealed that they were all classified as AI-ECG–potential non-STEMI). Although they were not covered by the AI-ECG–based clinical decision support systems, we still included them in the analysis based on the intention-to-treat design. AI denotes artificial intelligence; AI-ECG, artificial intelligence–enabled electrocardiogram; CONSORT, Consolidated Standards of Reporting Trials; ECG, electrocardiogram; and STEMI, ST-segment myocardial infarction.

on ECGs without considering the value of cardiac troponin level[30,31] (details are shown in Supplementary Method 2). Our EHR recorded all STEMI cases confirmed by urgent coronary angiography, independent of this trial.

For patients with STEMI who underwent urgent coronary angiography, four cardiologists reviewed the patients after the trial to further divide them into two groups: STEMI with occluded vessel(s) or STEMI with nonobstructive coronary arteries. For patients without urgent coronary angiography, the cardiologists reviewed 45 AI-ECG–identified potential STEMI cases and categorized them into two groups: STEMI without coronary angiography or without STEMI. Because of the large number of cases identified as potentially without STEMI by AI-ECG, a case-by-case review was deemed impractical, and all of these patients were categorized as without STEMI. For the primary analysis, only patients with STEMI with occluded vessel(s) were used. For event and accuracy analyses, all patients with STEMI were included. We acknowledge that this pragmatic approach might miss some patients with STEMI without urgent coronary angiography.

The index time in our study was defined as the time of ECG conduction. The baseline characteristics of each patient were collected from the EHR before the index time. The presence of coronary artery disease, diabetes mellitus, hypertension, hyperlipidemia, and chronic kidney disease was identified by the appropriate International Classification of Diseases codes. Patient information in the emergency department was acquired during the triage process, focusing on typical chest pain symptoms at the triage station to enable subsequent stratified analysis (details are shown in Supplementary Method 2).

PRESPECIFIED END POINTS AND POST HOC ANALYSIS

The prespecified primary end point was door-to-balloon time in patients with STEMI with occluded vessel(s). Because the door-to-balloon time may not represent the first medical contact time for patients in the inpatient department, the primary analysis considered only patients in the emergency department. We also analyzed the ECG-to-balloon time for both patients in the emergency department and inpatient patients simultaneously. The AI-ECG intervention was expected to be the most beneficial for patients identified as potentially having STEMI. Therefore, we conducted prespecified stratified analyses of the primary end point based on the AI-ECG results. For the prespecified exploratory analysis, we conducted additional stratified analyses and further analyses for each period of the treatment waiting time. In sensitivity analyses, patients with STEMI with occluded vessel(s) without ST elevation in the first ECG, patients with instances of intubation or resuscitation before coronary angiography, and patients who refused PPCI treatment were excluded in adherence to established quality indicator policies.[32]

The prespecified secondary end points were all-cause mortality within 365 days from the first ECG, cardiac death within 365 days, new-onset low ejection fraction within 90 days, hospitalization for patients in the emergency department, and STEMI-related diagnoses. STEMI-related diagnoses encompassed the following: STEMI with occluded vessel(s); urgent coronary angiography: STEMI with occluded vessel(s) plus STEMI with nonobstructive coronary arteries; all patients with STEMI: STEMI with occluded vessel(s) plus STEMI with nonobstructive coronary arteries plus STEMI without coronary angiography; and STEMI without coronary angiography. For patients with STEMI with occluded vessel(s), the ejection fraction, the highest level of high-sensitivity cardiac troponin I (hscTnI), the highest level of creatine kinase (CK), and the length of hospitalization were compared in post hoc analyses. An accuracy analysis of the AI-ECG system was also performed (details are shown in Supplementary Method 3).

SAMPLE SIZE

Prior to the ARISE trial, we conducted a pilot study of 25,002 patients and observed that AI-ECG intervention had the potential to reduce the door-to-balloon time from 70.0±13.6 minutes to 64.1±12.4 minutes.[20] Based on a significance level of $P<0.05$, a statistical power of 0.80, and a sample size ratio of 1.0 between the intervention and control groups, we concluded that we would require 77 patients with STEMI in each group. Considering that the incidence of STEMI among patients undergoing ECG examinations in our hospital is around 0.4%,[20] approximately 19,250 patients in each group (intervention and control) were required. That figure corresponds roughly to the total number of patients who underwent ECG examinations in our hospital annually. Therefore, the trial was conducted for 1 year to achieve a total of 21,612 and 21,622 cases in the intervention and control groups, respectively.

STATISTICAL ANALYSIS

The detailed statistical plan was based on an intention-to-treat design (details are shown in Supplementary Method 4). The statistical analysis was performed using R version 3.4.4. For time difference analysis, the Mann–Whitney U test using the wilcox.test function was chosen primarily because of the skewed distribution of treatment waiting time,[33] and raov function in R package Rfit version 0.24.2 was used to support stratified analysis for interaction terms. Logistic regression was used to estimate the odds ratios for event analyses. The prespecified stratified analysis was conducted by adding interaction terms to the logistic regression model for testing. For numeric prognosis analyses, the same statistical method of time difference analysis was used. For the accuracy analyses, we calculated the confidence intervals (CIs) for each percentage using the Z distribution and used chi-square tests for conducting stratified analyses.

Results

PATIENTS’ CHARACTERISTICS STRATIFIED BY RANDOMIZATION

Table 1 presents the baseline characteristics of patients stratified by randomization (details regarding the participating cardiologists are shown in Supplementary Result 1). The average age of patients was 60 years, 49.5% were male, and 7.2% were from the community hospital, of whom 36.8% were from the inpatient department. Of the 21,612 patients in the intervention group, 77 (0.4%), 23 (0.1%), and 7 (0.0%) were classified as having STEMI with occluded vessel(s), STEMI with nonobstructive coronary arteries, and STEMI without coronary angiography, respectively. Of the 21,621 patients in the control group, 68 (0.3%), 18 (0.1%), and 16 (0.1%) were classified into the corresponding subgroups, respectively. The average age of patients with STEMI with occluded vessel(s) was 65 years, and 80% were male. Notably, in the inpatient department, there were seven patients with STEMI with occluded vessel(s) in the intervention group and only one in the control group. No significant differences in baseline characteristics were observed between the intervention and control groups for both the entire patient population and specifically, among patients with STEMI with occluded vessel(s) (details are shown in Supplementary Result 2).

PRIMARY ANALYSIS FOR STEMI WITH OCCLUDED VESSEL(S)

In the emergency department, the median door-to-balloon time was 82.0 minutes (interquartile range, 62.5 to 89.5 minutes) in the intervention group compared with 96.0 minutes (interquartile range, 78.0 to 137.0 minutes) in the control group (P=0.002). When analyzing both emergency and inpatient cases, the median ECG-to-balloon time was 78.0 minutes (interquartile range, 56.9 to 88.2 minutes)

Table 1. Baseline Characteristics and Diagnostic Group in Intervention and Control Groups.*

	All Patients			STEMI with Occluded Vessel(s)		
Characteristic	**Intervention (n=21,612)**	**Control (n=21,622)**	**P Value†**	**Intervention (n=77)**	**Control (n=68)**	**P Value†**
AI-ECG result			0.625			0.965
Potential STEMI	108 (0.5%)	101 (0.5%)		67 (87.0%)	59 (86.8%)	
Potential non-STEMI	21,504 (99.5%)	21,521 (99.5%)		10 (13.0%)	9 (13.2%)	
Hospital			0.568			1.000
Academic medical center	20,040 (92.7%)	20,080 (92.9%)		73 (94.8%)	64 (94.1%)	
Community hospital	1,572 (7.3%)	1,542 (7.1%)		4 (5.2%)	4 (5.9%)	
Department			0.451			0.067
Emergency department	13,606 (63.0%)	13,688 (63.3%)		70 (90.9%)	67 (98.5%)	
Inpatient department	8,006 (37.0%)	7,934 (36.7%)		7 (9.1%)	1 (1.5%)	
Time frame			0.328			0.055
Regular hours	12,384 (57.3%)	12,289 (56.8%)		37 (48.1%)	22 (32.4%)	
Off hours	9,228 (42.7%)	9,333 (43.2%)		40 (51.9%)	46 (67.6%)	
Gender (male)	10,722 (49.6%)	10,675 (49.4%)	0.617	64 (83.1%)	52 (76.5%)	0.318
Age — yr, mean (±SD)	60.3±18.4	60.2±18.3	0.601	64.8±13.2	65.1±11.3	0.899
Age group — yr			0.449			0.780
<65	12,056 (55.8%)	12,191 (56.4%)		40 (51.9%)	33 (48.5%)	
65–74	4,881 (22.6%)	4,829 (22.3%)		23 (29.9%)	24 (35.3%)	
≥75	4,675 (21.6%)	4,602 (21.3%)		14 (18.2%)	11 (16.2%)	
Coronary artery disease	4,616 (21.4%)	4,677 (21.6%)	0.491	49 (63.6%)	48 (70.6%)	0.375
Diabetes mellitus	4,970 (23.0%)	5,080 (23.5%)	0.220	23 (29.9%)	24 (35.3%)	0.486
Hypertension	7,864 (36.4%)	8,037 (37.2%)	0.091	26 (33.8%)	27 (39.7%)	0.459
Hyperlipidemia	7,826 (36.2%)	8,004 (37.0%)	0.082	30 (39.0%)	30 (44.1%)	0.529
Chronic kidney disease	4,400 (20.4%)	4,560 (21.1%)	0.061	19 (24.7%)	11 (16.2%)	0.207
Diagnostic group			0.196			
STEMI with occluded vessel(s)	77 (0.4%)	68 (0.3%)				
STEMI without occluded vessel(s)	23 (0.1%)	18 (0.1%)				
STEMI without coronary angiography	7 (0.0%)	16 (0.1%)				
Probably non-STEMI	21,505 (99.5%)	21,520 (99.5%)				

* Values are numbers (percentages) unless indicated otherwise. AI-ECG denotes artificial intelligence–assisted electrocardiogram; SD, standard deviation; and STEMI, ST-segment myocardial infarction.

† The P values are two sided, with no adjustment for multiple comparison.

Back to Table of Contents

in the intervention group compared with 83.6 minutes (interquartile range, 72.7 to 127.8 minutes) in the control group (P=0.011). The overall trend was consistent in the sensitivity analysis that excluded patients with nonsystematic delayed factors, in adherence to established quality indicator policies.[32]

Because of the limited sample size, the stratified analysis based on AI-ECG results did not confirm the prespecified hypothesis (P for intervention group × AI-ECG group interaction >0.05) (details are shown in Supplementary Result 3). For ECG-to-balloon time, the AI-ECG intervention demonstrated a consistent pattern across all stratified analyses (Fig. 2).

PRESPECIFIED SECONDARY ANALYSIS

Figure 3 presents the event analysis of the AI-ECG intervention in the diagnosis and management of STEMI. Cardiac death was significantly different between the groups (0.4% in the intervention group vs. 0.5% in the control group; odds ratio, 0.73; 95% CI, 0.55 to 0.97). The only significant result in STEMI-related diagnoses was a lower incidence of STEMI without coronary angiography in the intervention group (odds ratio, 0.37; 95% CI, 0.14 to 0.94). There was no significant difference in all patients with STEMI between the two groups (odds ratio, 1.05; 95% CI, 0.80 to 1.38), indicating that more patients in the intervention group were scheduled for urgent coronary angiography (Supplementary Result 4 shows the

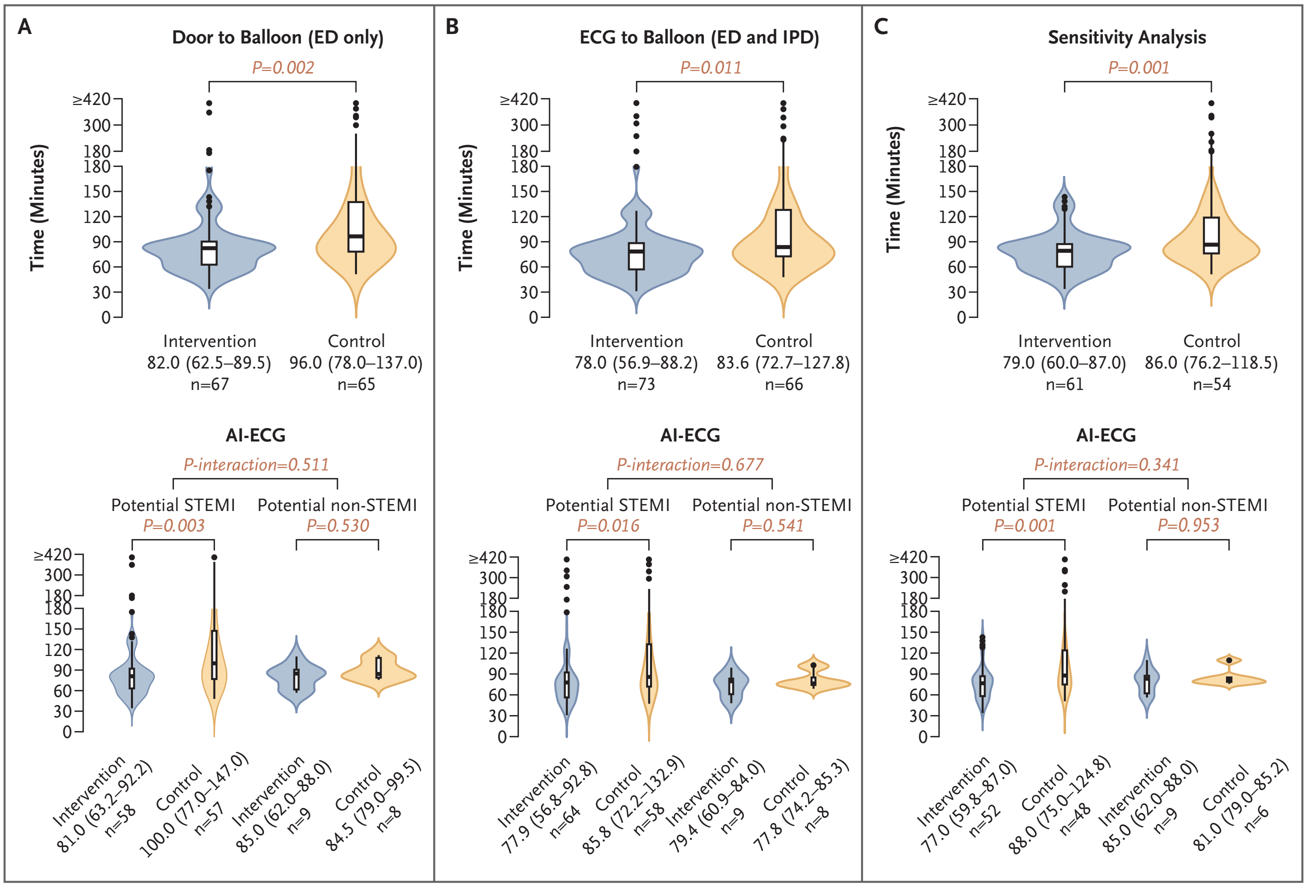

Figure 2. Time Difference Analysis for STEMI with Occluded Vessel(s).

Panel A shows the intention-to-treat analysis on door-to-balloon time. Because the door time was unavailable for patients in the IPD, the prespecified primary analysis included only the patients in the ED. There were three and two STEMIs with occluded vessel(s) without balloon time in the intervention and control groups, respectively, because they were found to be ineligible for primary percutaneous coronary intervention during coronary angiography. Panel B shows the ECG-to-balloon time for patients in the ED and IPD. One additional STEMI with occluded vessel(s) without balloon time in the intervention group in IPD was excluded in this analysis. Panel C shows the door-to-balloon time comparison in patients without nonsystematic delayed factors (nondiagnostic ECG, intubation/resuscitation, and patient declined). The stratified analysis in the lower panel was the only prespecified stratified analysis. AI-ECG denotes artificial intelligence–electrocardiogram; ECG, electrocardiogram; ED, emergency department; IPD, inpatient department; and STEMI, ST-segment myocardial infarction.

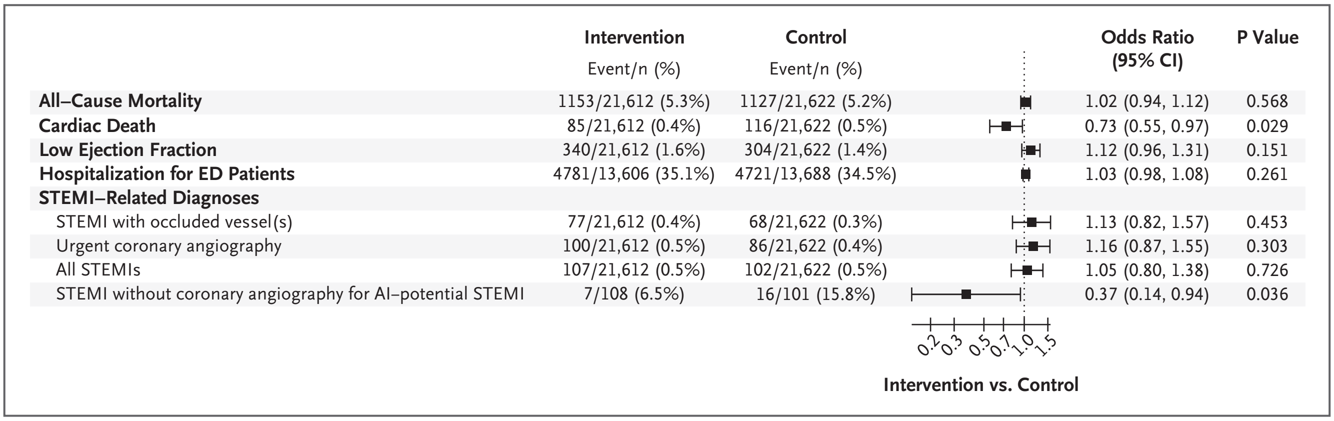

	Intervention Event/n (%)	Control Event/n (%)	Odds Ratio (95% CI)	P Value
All–Cause Mortality	1153/21,612 (5.3%)	1127/21,622 (5.2%)	1.02 (0.94, 1.12)	0.568
Cardiac Death	85/21,612 (0.4%)	116/21,622 (0.5%)	0.73 (0.55, 0.97)	0.029
Low Ejection Fraction	340/21,612 (1.6%)	304/21,622 (1.4%)	1.12 (0.96, 1.31)	0.151
Hospitalization for ED Patients	4781/13,606 (35.1%)	4721/13,688 (34.5%)	1.03 (0.98, 1.08)	0.261
STEMI–Related Diagnoses				
STEMI with occluded vessel(s)	77/21,612 (0.4%)	68/21,622 (0.3%)	1.13 (0.82, 1.57)	0.453
Urgent coronary angiography	100/21,612 (0.5%)	86/21,622 (0.4%)	1.16 (0.87, 1.55)	0.303
All STEMIs	107/21,612 (0.5%)	102/21,622 (0.5%)	1.05 (0.80, 1.38)	0.726
STEMI without coronary angiography for AI–potential STEMI	7/108 (6.5%)	16/101 (15.8%)	0.37 (0.14, 0.94)	0.036

Figure 3. Analyses of Prespecified Secondary End Points.

The analysis of hospitalization was only for patients in the ED. The detailed definitions of each STEMI-related diagnosis are STEMI with occluded vessel(s); urgent coronary angiography (STEMI with occluded vessel[s] + STEMI with nonobstructive coronary arteries); all STEMI patients (STEMI with occluded vessel[s] + STEMI with nonobstructive coronary arteries + STEMI without coronary angiography); and STEMI without coronary angiography. Because there was no STEMI without coronary angiography in AI-potential non-STEMI group due to the pragmatic data collection strategy, the analysis for this event included only the AI-potential STEMI subgroup. AI denotes artificial intelligence; CI, confidence interval; ED, emergency department; and STEMI, ST-segment myocardial infarction.

prespecified stratified analysis and other detailed analyses). The reduction in cardiac death might not come from AI-ECG alerts in the AI-potential STEMI subgroup because we observed a nonsignificant increasing trend of cardiac death in the intervention group compared with the control group (odds ratio, 1.36; 95% CI, 0.57 to 3.21). A noteworthy but nonsignificant trend was the increased identification of STEMI with occluded vessel(s) in the intervention group compared with the control group from the inpatient department (odds ratio, 6.94; 95% CI, 0.85 to 56.42).

POST HOC ANALYSIS FOR CLINICAL OUTCOMES IN STEMI WITH OCCLUDED VESSEL(S)

We analyzed the differences in several important prognostic indicators during hospitalization for STEMI with occluded vessel(s) between the intervention and control groups (details are shown in Supplementary Result 5). The results showed no significant differences in these measures, including ejection fraction, highest level of hscTnI, highest level of CK, and length of hospitalization. Further subgroup analyses did not reveal any significant findings.

PROSPECTIVE ACCURACY OF AI-ECG

Table 2 shows the diagnostic accuracy of AI-ECG in the trial, with a positive predictive value of 89.5% (95% CI, 85.3 to 93.6%), a negative predictive value of 99.9% (95% CI, 99.9 to 100.0%), a sensitivity of 89.5% (95% CI, 85.3 to 93.6%), and a specificity of 99.9% (95% CI, 99.9 to 100.0%), significantly better than the Philips automatic ECG analysis system (a stratified analysis and a false-positive analysis are shown in Supplementary Result 6).

Discussion

The ARISE trial is a pragmatic RCT that compared the implementation of the AI-ECG versus standard of care for STEMI management. We found that the integration of AI-ECG into the EHR as a CDSS significantly reduced the door-to-balloon time in the emergency department setting and the ECG-to-balloon time in patients in the emergency department and inpatients. There was modest to no difference between the intervention and control groups in rates of new-onset heart failure with reduced ejection fraction or all-cause mortality.

Reducing door-to-balloon time in STEMI, a key performance indicator, is crucial for improving prognosis. However, timely diagnosis and treatment are challenging in the clinical setting.[34] Strategies have entailed improvement in catheterization laboratory activation, rapid team preparedness, data feedback, and administrative support.[35] All of these interventions have demonstrated significant reductions in treatment waiting times, but they often incur

Table 2. The Accuracy of the Deep Learning Model and the Philips Automatic System.[*]

	Deep Learning Model			Philips Automatic System		
Case/Control	Potential STEMI (n=209)	Potential Non-STEMI (n=43,025)	Sensitivity/Specificity	AMI (n=1,001)	Not AMI (n=42,233)	Sensitivity/Specificity
Case						
STEMI with occluded vessel(s)	126 (60.3%)	19 (0.0%)		99 (9.9%)	46 (0.1%)	
STEMI without occluded vessel(s)	38 (18.2%)	3 (0.0%)		14 (1.4%)	27 (0.1%)	
STEMI without coronary angiography	23 (11.0%)	0 (0.0%)		16 (1.6%)	7 (0.0%)	
All STEMI cases	187 (89.5%)[†]	22 (0.1%)	Sensitivity 89.5% (187/209)	129 (12.9%)[†]	80 (0.2%)	Sensitivity 61.7% (129/209)
Control						
Probably non-STEMI	22 (10.5%)	43,003 (99.9%)[‡]	Specificity 99.9% (187/209)	872 (87.1%)	42,153 (99.8%)[‡]	Specificity 98.0% (187/209)

[*] Values are numbers (percentages) unless indicated otherwise. The “case” group encompassed the combination of STEMI with occluded vessel(s), STEMI without occluded vessel(s), and STEMI without coronary angiography groups, whereas the “control” group consisted of the remaining “probably non-STEMI” group. AMI denotes acute myocardial infarction; and STEMI, ST-segment myocardial infarction.

[†] This indicates positive predictive values.

[‡] This indicates negative predictive values.

high costs. In contrast, low-cost interventions, such as AI-ECG–based CDSS, have potential by expediting communication between frontline physicians and on-duty cardiologists. Numerous studies have shown the effectiveness of AI-ECG–based CDSS in reducing treatment delay.[19,20] The ARISE trial, the first RCT evaluating the efficacy of AI-ECG–based CDSS, provides compelling evidence for future large-scale implementation of this technology to further reduce ECG-to-catheterization laboratory time.

We show that the median door-to-balloon time in the emergency department was 86.0 minutes in the control group (excluding patients with nonsystematic delays in adherence to established quality indicator policies[32]), a time that was higher than in our previous study (70 minutes).[20] It had been proposed that the median door-to-balloon time was 78 minutes (interquartile range, 62 to 106 minutes).[36] Other research reported a median door-to-balloon time of 86 minutes without prehospital activation, similar to the results in this study.[37] Upon a comprehensive retrospective analysis of the study population, it became evident that the time delay during the pandemic era was partially attributable to the screening for coronavirus disease 2019 (Covid-19) before PPCI. However, because both groups in the RCT were equally affected by the Covid-19 pandemic, we believe that the time difference between the intervention and control groups in the ARISE trial remains credible.

AI-ECG–based CDSS provides an additional advantage by minimizing the risk of misdiagnosis. Misdiagnoses among patients with STEMI are frequently reported,[5] and the constraints of observational studies impede the precise identification of such cases because they typically focus on patients with a confirmed final diagnosis of STEMI. Furthermore, regarding the rate of occluded vessels among patients with STEMI who underwent urgent coronary angiography, we observed that the proportions were comparable in the intervention group (77/100, 77%) and the control group (68/86, 79%). It is more plausible that the decrease in patients with STEMI without urgent coronary angiography identified in posttrial review in the intervention group, compared with the control group, is attributable to differences in the rate of misdiagnosis. The high accuracy of AI-ECG has been validated in multiple prospective interventional studies, with a positive predictive value exceeding 80%.[19,20] This high value is critical for the success of CDSS.[9] The ARISE trial demonstrated a trend of increased STEMI with occluded vessel(s) in the intervention group coupled with a significant decrease in STEMI without coronary angiography, particularly within the inpatient department. The ARISE trial substantiated the potential advantages of AI-ECG in diminishing misdiagnoses, extending cardiology-level bedside care to patients across the hospital, and consequently, enhancing health care quality.

This study has several limitations. First, the sample size of the ARISE trial may have been inadequate to detect significant differences in secondary end points, particularly in subgroup analyses. Second, this study primarily focused on reducing treatment waiting times, leading to short-term outcome follow-up. The limited sample size made it challenging to obtain sufficient statistical power, even with extended follow-up for STEMI complications. Third, the study was conducted at a single center, which may restrict the applicability of the results to other health care settings with different patient populations and resources. Fourth, the pragmatic data collection process may lead to overestimations of the negative predictive rate and sensitivity of AI-ECG. Fifth, given the nature of the intervention, it was not feasible to blind health care providers and patients, which may have introduced bias into the assessment of outcomes.

Conclusion

The ARISE study evaluated the impact of an AI-ECG intervention on STEMI management. The incorporation of AI-ECG as an affordable CDSS resulted in a significant reduction in door-to-balloon time, underscoring its potential to enhance the timeliness of care delivery. The intervention also demonstrated promising accuracy in identifying potential STEMI cases, enhancing the attention of and proactive management by health care providers. Further research with larger sample sizes and extended follow-up periods is necessary to provide additional validation of the benefits on clinical outcomes.

Disclosures

Author disclosures and other supplementary materials are available at ai.nejm.org.

Supported by the National Science and Technology Council, Taiwan (grant NSTC110-2314-B-016-010-MY3 to Dr. Chin Lin and grant NSTC112-2321-B-016-003 to Dr. Shih-Hua Lin); the Cheng Hsin General Hospital, Taiwan (grant CHNDMC-113-11205 to Dr. Chin Lin and grant CHNDMC-113-01 to Dr. Chin-Sheng Lin); and the Medical Affairs Bureau, Taiwan (grants MND-MAB-110-113, MND-MAB-D-111045, MND-MAB-C13-112050, MND-MAB-C07-113020, and MND-MAB-C07-113021 to Dr. Chin Lin and grants MND-MAB-C07-113022, MND-MAB-C13-112051, MND-MAB-C08-111032, and MND-MAB-C08-11032 to Dr. Chin-Sheng Lin).

We sincerely appreciate the participation of the 20 on-duty cardiologists involved in the Artificial Intelligence–Powered Rapid Identification of ST-Elevation Myocardial Infarction via Electrocardiogram trial.

Author Affiliations

[1] Medical Technology Education Center, School of Medicine, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[2] Department of Artificial Intelligence of Things, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[3] School of Public Health, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[4] Graduate Institute of Aerospace and Undersea Medicine, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[5] Division of Cardiology, Department of Internal Medicine, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[6] Department of Family and Community Medicine, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[7] Department of Statistics and Information Science, Fu Jen Catholic University, New Taipei City, Taiwan, R.O.C.

[8] Graduate Institute of Life Sciences, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[9] Division of Endocrinology and Metabolism, Department of Internal Medicine, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C

[10] Division of Nephrology, Department of Internal Medicine, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C.

[11] Division of Cardiovascular Surgery, Department of Surgery, Tri-Service General Hospital, National Defense Medical Center, Taipei, Taiwan, R.O.C.

References

1. Lawton JS, Tamis-Holland JE, Bangalore S, et al. 2021 ACC/AHA/SCAI Guideline for Coronary Artery Revascularization: a report of the American College of Cardiology/American Heart Association Joint Committee on Clinical Practice Guidelines [published correction appears in Circulation 2022;145:e772]. Circulation 2022;145:e18-e114. DOI: 10.1161/CIR.0000000000001038.
2. Rokos IC, French WJ, Mattu A, et al. Appropriate cardiac cath lab activation: optimizing electrocardiogram interpretation and clinical decision-making for acute ST-elevation myocardial infarction. Am Heart J 2010;160:995-1003.e8. DOI: 10.1016/j.ahj.2010.08.011.
3. McCabe JM, Armstrong EJ, Ku I, et al. Physician accuracy in interpreting potential ST-segment elevation myocardial infarction electrocardiograms. J Am Heart Assoc 2013;2:e000268. DOI: 10.1161/JAHA.113.000268.
4. Huitema AA, Zhu T, Alemayehu M, Lavi S. Diagnostic accuracy of ST-segment elevation myocardial infarction by various healthcare providers. Int J Cardiol 2014;177:825-829. DOI: 10.1016/j.ijcard.2014.11.032.
5. Wu J, Gale CP, Hall M, et al. Editor's choice — impact of initial hospital diagnosis on mortality for acute myocardial infarction: a national cohort study. Eur Heart J Acute Cardiovasc Care 2018;7:139-148. DOI: 10.1177/2048872616661693.
6. Terkelsen CJ, Sørensen JT, Maeng M, et al. System delay and mortality among patients with STEMI treated with primary percutaneous coronary intervention. JAMA 2010;304:763-771. DOI: 10.1001/jama.2010.1139.
7. Swaminathan RV, Wang TY, Kaltenbach LA, et al. Nonsystem reasons for delay in door-to-balloon time and associated in-hospital mortality: a report from the National Cardiovascular Data Registry. J Am Coll Cardiol 2013;61:1688-1695. DOI: 10.1016/j.jacc.2012.11.073.

8. Miedema MD, Newell MC, Duval S, et al. Causes of delay and associated mortality in patients transferred with ST-segment-elevation myocardial infarction. Circulation 2011;124:1636-1644. DOI: 10.1161/CIRCULATIONAHA.111.033118.

9. Kawamoto K, Houlihan CA, Balas EA, Lobach DF. Improving clinical practice using clinical decision support systems: a systematic review of trials to identify features critical to success. BMJ 2005; 330:765. DOI: 10.1136/bmj.38398.500764.8F.

10. Liu WC, Lin CS, Tsai CS, et al. A deep learning algorithm for detecting acute myocardial infarction. EuroIntervention 2021;17: 765-773. DOI: 10.4244/EIJ-D-20-01155.

11. Cvach M, Rothwell KJ, Cullen AM, Nayden MG, Cvach N, Pham JC. Effect of altering alarm settings: a randomized controlled study. Biomed Instrum Technol 2015;49:214-222. DOI: 10.2345/0899-8205-49.3.214.

12. Firoozabadi R, Gregg RE, Babaeizadeh S. Intelligent use of advanced capabilities of diagnostic ECG algorithms in a monitoring environment. J Electrocardiol 2017;50:615-619. DOI: 10.1016/j.jelectrocard.2017.04.013.

13. Siontis KC, Noseworthy PA, Attia ZI, Friedman PA. Artificial intelligence-enhanced electrocardiography in cardiovascular disease management. Nat Rev Cardiol 2021;18:465-478. DOI: 10.1038/s41569-020-00503-2.

14. Al-Zaiti SS, Martin-Gill C, Zègre-Hemsey JK, et al. Machine learning for ECG diagnosis and risk stratification of occlusion myocardial infarction. Nat Med 2023;29:1804-1813. DOI: 10.1038/s41591-023-02396-3.

15. Yao X, Rushlow DR, Inselman JW, et al. Artificial intelligence-enabled electrocardiograms for identification of patients with low ejection fraction: a pragmatic, randomized clinical trial. Nat Med 2021;27:815-819. DOI: 10.1038/s41591-021-01335-4.

16. Lin CS, Liu WT, Tsai DJ, et al. AI-enabled electrocardiography alert intervention and all-cause mortality: a pragmatic randomized clinical trial. Nat Med 2024;30:1461-1470. DOI: 10.1038/s41591-024-02961-4.

17. Zhao Y, Xiong J, Hou Y, et al. Early detection of ST-segment elevated myocardial infarction by artificial intelligence with 12-lead electrocardiogram. Int J Cardiol 2020;317:223-230. DOI: 10.1016/j.ijcard.2020.04.089.

18. Chang K-C, Hsieh P-H, Wu M-Y, et al. Usefulness of multi-labelling artificial intelligence in detecting rhythm disorders and acute ST-elevation myocardial infarction on 12-lead electrocardiogram. Eur Heart J Digit Health 2021;2:299-310. DOI: 10.1093/ehjdh/ztab029.

19. Wang YC, Chen KW, Tsai BY, et al. Implementation of an all-day artificial intelligence-based triage system to accelerate door-to-balloon times. Mayo Clin Proc 2022;97:2291-2303. DOI: 10.1016/j.mayocp.2022.05.014.

20. Liu WC, Lin C, Lin CS, et al. An artificial intelligence-based alarm strategy facilitates management of acute myocardial infarction. J Pers Med 2021;11:1149. DOI: 10.3390/jpm11111149.

21. Menees DS, Peterson ED, Wang Y, et al. Door-to-balloon time and mortality among patients undergoing primary PCI. N Engl J Med 2013;369:901-909. DOI: 10.1056/NEJMoa1208200.

22. Lassen JF, Bøtker HE, Terkelsen CJ. Timely and optimal treatment of patients with STEMI. Nat Rev Cardiol 2013;10:41-48. DOI: 10.1038/nrcardio.2012.156.

23. Kaul P, Federspiel JJ, Dai X, et al. Association of inpatient vs outpatient onset of ST-elevation myocardial infarction with treatment and clinical outcomes. JAMA 2014;312:1999-2007. DOI: 10.1001/jama.2014.15236.

24. Garberich RF, Traverse JH, Claussen MT, et al. ST-elevation myocardial infarction diagnosed after hospital admission. Circulation 2014;129:1225-1232. DOI: 10.1161/CIRCULATIONAHA.113.005568.

25. Dai X, Bumgarner J, Spangler A, Meredith D, Smith SC, Stouffer GA. Acute ST-elevation myocardial infarction in patients hospitalized for noncardiac conditions. J Am Heart Assoc 2013;2:e000004. DOI: 10.1161/JAHA.113.000004.

26. Plana D, Shung DL, Grimshaw AA, Saraf A, Sung JJY, Kann BH. Randomized clinical trials of machine learning interventions in health care: a systematic review. JAMA Netw Open 2022;5: e2233946. DOI: 10.1001/jamanetworkopen.2022.33946.

27. Horwitz LI, Kuznetsova M, Jones SA. Creating a learning health system through rapid-cycle, randomized testing. N Engl J Med 2019;381:1175-1179. DOI: 10.1056/NEJMsb1900856.

28. Liu X, Cruz Rivera S, Moher D, Calvert MJ, Denniston AK. Reporting guidelines for clinical trial reports for interventions involving artificial intelligence: the CONSORT-AI extension. Nat Med 2020; 26:1364-1374. DOI: 10.1038/s41591-020-1034-x.

29. Altman DG, Bland JM. Statistics notes. Treatment allocation in controlled trials: why randomise? BMJ 1999;318:1209. DOI: 10.1136/bmj.318.7192.1209.

30. Thygesen K, Alpert JS, Jaffe AS, et al. Fourth universal definition of myocardial infarction (2018) [published correction appears in Circulation 2018;138:e652]. Circulation 2018;138:e618-e651. DOI: 10.1161/CIR.0000000000000617.

31. Byrne RA, Rossello X, Coughlan JJ, et al. 2023 ESC guidelines for the management of acute coronary syndromes. Eur Heart J 2023; 44:3720-3826. DOI: 10.1093/eurheartj/ehad191.

32. McCabe JM, Kennedy KF, Eisenhauer AC, et al. Reporting trends and outcomes in ST-segment-elevation myocardial infarction national hospital quality assessment programs. Circulation 2014; 129:194-202. DOI: 10.1161/CIRCULATIONAHA.113.006165.

33. Hart A. Mann–Whitney test is not just a test of medians: differences in spread can be important. BMJ 2001;323:391-393. DOI: 10.1136/bmj.323.7309.391.

34. Moscucci M, Eagle KA. Door-to-balloon time in primary percutaneous coronary intervention: is the 90-minute gold standard an unreachable chimera? Circulation 2006;113:1048-1050. DOI: 10.1161/CIRCULATIONAHA.105.606905.

Back to Table of Contents

35. Krumholz HM, Bradley EH, Nallamothu BK, et al. A campaign to improve the timeliness of primary percutaneous coronary intervention: door-to-balloon: an alliance for quality. JACC Cardiovasc Interv 2008;1:97-104. DOI: 10.1016/j.jcin.2007.10.006.

36. McCabe JM, Armstrong EJ, Hoffmayer KS, et al. Impact of door-to-activation time on door-to-balloon time in primary percutaneous coronary intervention for ST-segment elevation myocardial infarctions: a report from the Activate-SF registry. Circ Cardiovasc Qual Outcomes 2012;5:672-679. DOI: 10.1161/CIRCOUTCOMES.112.966382.

37. Savage ML, Hay K, Vollbon W, et al. Prehospital activation of the cardiac catheterization laboratory in ST-segment-elevation myocardial infarction for primary percutaneous coronary intervention. J Am Heart Assoc 2023;12:e029346. DOI: 10.1161/JAHA.122.029346.

NEJM AI

NEJM AI 2024; 1 (7)
DOI: 10.1056/AIe2400472

EDITORIAL

AI-RISE to the Challenge — Artificial Intelligence Reduces Time to Treatment in STEMI

Robert Avram, M.D., M.Sc.,[1,2] and William F. Fearon, M.D.[3,4]

Received: May 8, 2024; Accepted: May 8, 2024; Published: June 27, 2024

Abstract

The ARISE (Artificial Intelligence–Enabled Rapid Identification of ST-Elevation Myocardial Infarction Using Electrocardiogram) trial randomly assigned patients to either an AI-powered electrocardiogram (AI-ECG) system or standard care and tested the effect of AI-ECG interpretation and automated text-based short message service notification on treatment delays and diagnostic accuracy for patients with ST-segment elevation myocardial infarction (STEMI). The AI-ECG system cut the median door-to-balloon time by 14 minutes (from 96.0 to 82.0 minutes, P<0.001) for patients presenting to the emergency department. The AI-ECG system also decreased the ECG-to-balloon time by 5.6 minutes (from 83.6 to 78.0 minutes, P<0.001) for hospitalized patients. The AI-ECG system also had a high positive and negative predictive value (89.5% and 99.9%), and the AI-ECG group had fewer STEMI activations in patients not requiring emergent angiography. However, the single-center design, short follow-up period, and lack of evaluation of care appropriateness and clinical safety end points limit the study's generalizability. Although the findings suggest that AI-ECG can expedite STEMI diagnosis and treatment, further research is needed to confirm these results in diverse settings and assess the impact on long-term patient outcomes.

The best treatment for acute ST-segment elevation myocardial infarction (STEMI) is primary percutaneous coronary intervention (PCI). Previously, a strong association was reported between the quick performance of primary PCI, measured by door-to-balloon time (the duration from the patient's arrival to coronary artery balloon inflation for blood flow restoration) and reduced mortality. Unfortunately, despite significant reductions in door-to-balloon times, STEMI mortality rates have remained relatively unchanged over recent decades.

The integration of AI within cardiovascular care has increased opportunities for advancements in diagnostic accuracy and care optimization. The ARISE (Artificial Intelligence–Enabled

The author affiliations are listed at the end of the article.

Dr. Avram can be contacted at robert.avram.md@gmail.com or at Montreal Heart Institute, 5000 Rue Bélanger, H1T 1C8, Montréal, QC, Canada.

Read Article at ai.nejm.org

Back to Table of Contents

Rapid Identification of ST-Elevation Myocardial Infarction Using Electrocardiogram) trial[1] was a pragmatic randomized controlled trial that assessed the impact of an AI-powered electrocardiogram (AI-ECG) system on treatment delays and diagnostic accuracy for STEMI. The trial enrolled 43,234 patients (average age 60 years; 49.5% male) who visited the emergency department or inpatient department and had at least one ECG without a history of coronary angiography (CAG) within 3 days. Twenty on-duty cardiologists were randomly assigned daily to either the intervention group (receiving AI-ECG alerts) or the control group (standard care).

The AI-ECG system used was based on a previously validated deep-learning model, which was trained on a large dataset of 12-lead ECGs and achieved a high positive predictive value of 93.2% for STEMI detection in a preliminary prospective study. In the ARISE trial intervention group, the AI-ECG system performed real-time analysis on all in-hospital ECGs. When a potential STEMI was identified, the system sent an immediate short message service (SMS) notification, including the ECG image and diagnosis, to the on-duty cardiologist's mobile device for confirmation. In the control group, frontline doctors evaluated the possible STEMI patients and then contacted the cardiologist on duty for verification. The on-duty cardiologist had the ultimate authority to activate the catheterization laboratory. Although the AI system could identify STEMI with and without occluded vessels, only STEMI with occluded vessels was used for primary analysis.

The AI-ECG system significantly reduced the median door-to-balloon time for patients presenting to the emergency department by 14 minutes (from 96.0 minutes [interquartile range, 78.0 to 137.0] to 82.0 minutes [interquartile range, 62.5 to 89.5], P=0.02) and decreased the median ECG-to-balloon time for hospitalized patients by 5.6 minutes (from 83.6 minutes [interquartile range, 72.7 to 127.8] to 78.0 minutes [interquartile range, 56.9 to 88.2], P=0.011). This time reduction was observed even though the AI system's primary function was designed to improve diagnostic accuracy. The authors suggest this may be due to the automated SMS alerts sent to the on-duty cardiologist in the AI-ECG group. Because these notifications were not present in the control group, frontline doctors had to review the ECG and get in touch with the on-duty cardiologist, potentially lengthening the door-to-balloon time. There was also no difference in new-onset heart failure or left ventricular ejection fraction between both groups.

The AI-ECG demonstrated a high ability to correctly identify positive cases (89.5%) and negative cases (99.9%), exceeding the conventional automated ECG interpretation offered on commercial devices. Additionally, there was a significant reduction in the number of STEMI activations not requiring CAG in the AI-ECG group, suggesting an improvement in diagnostic accuracy and a reduction in the rate of misdiagnosis, because the proportion of occluded vessels among STEMI patients who underwent urgent CAG was similar between the intervention (77/100, 77%) and control groups (68/86, 79%). No significant differences were found in hospital stay duration, ejection fraction, left ventricular ejection fraction, or peak levels of troponin I and creatine kinase between the AI-ECG group and the control group, indicating no differences in patient outcomes.

The ARISE trial's strengths are its pragmatic design, large sample size, and use of a randomized controlled approach to evaluate the impact of AI-ECG on STEMI care. The study reveals useful information about how AI technology can improve clinical practice with automatic notifications, especially for time-critical conditions such as STEMI. The inclusion of the automated SMS alerts in the intervention group is a notable strength, because it highlights the importance of rapid communication in the STEMI care pathway. However, the time from ECG to catheterization laboratory is an interesting choice for a primary end point because although this end point is relevant for assessing the efficiency of STEMI care, it does not directly measure patient outcomes. Future studies should aim to determine the difference in clinical end points such as mortality and morbidity to more comprehensively evaluate the impact of AI-ECG on STEMI care. Moreover, it will be important to further explore the economic value of reducing false STEMI activations with the AI-ECG system.

The trial has several limitations. First, the single-center design may limit generalizability to other health care settings. The same physicians providing care to both the intervention and control groups could lead to contamination, and the short follow-up period with few analyzed events focuses primarily on short-term outcomes. The primary analysis was limited to patients who underwent urgent cardiac catheterization and had confirmed STEMI,

Back to Table of Contents

potentially favoring the intervention group in the analytic path. The current study also lacks data on inappropriate catheterization laboratory activation and patients who needed catheterization but never received it. Although the findings are hypothesis-generating, they should be considered a “proof of concept.” Future multicenter trials are needed to assess the performance and generalizability of this application of AI in diverse health care settings and patient populations. The integration of automated alerts should be further investigated to determine their impact on communication efficiency and treatment delays.

In conclusion, the ARISE trial provides compelling evidence for the integration of AI-ECG into clinical practice for the management of STEMI to reduce door-to-balloon time and improve the accuracy of ECG interpretation for possible STEMI. The study demonstrates the potential of AI technology, coupled with automated notifications, to expedite diagnosis, improve diagnostic accuracy, and streamline the pathway to definitive treatment. Well-designed trials are critical to test the effect of AI algorithms on meaningful patient outcomes and to inform policy and practice before widespread implementation.

Disclosures

Author disclosures are available at ai.nejm.org.

Author Affiliations

[1] Department of Medicine, Montreal Heart Institute, Université de Montréal, Montreal

[2] HeartWise.AI, Université de Montréal, Montreal

[3] Division of Cardiovascular Medicine and Stanford Cardiovascular Institute, Stanford University School of Medicine, Stanford, CA

[4] Palo Alto Veterans Affairs Health Care Systems, Palo Alto, CA

Reference

1. Lin C, Liu W-T, Chang C-H, et al. Artificial intelligence–powered rapid identification of ST-elevation myocardial infarction via electrocardiogram (ARISE) — a pragmatic randomized controlled trial. NEJM AI 2024;1(7). DOI: 10.1056/AIoa2400190.

NEJM AI
NEJM AI 2024; 1 (5)
DOI: 10.1056/AIdbp2300192

DATASETS, BENCHMARKS, AND PROTOCOLS

GPT versus Resident Physicians — A Benchmark Based on Official Board Scores

Uriel Katz, M.D.,[1] Eran Cohen, M.D.,[2,3] Eliya Shachar, M.D.,[2,4] Jonathan Somer, B.Sc.,[5] Adam Fink, M.D.,[6] Eli Morse, M.D.,[7] Beki Shreiber, B.Sc.,[8] and Ido Wolf, M.D.[2,3,4]

Received: October 18, 2023; Revised: January 31, 2024; Accepted: February 5, 2024; Published: April 12, 2024

Abstract

BACKGROUND Artificial intelligence (AI) is a burgeoning technological advancement, with considerable promise for influencing the field of medicine. As a preliminary step toward integrating AI into medical practice, it is imperative to ascertain whether model performance is comparable with that of physicians. We present a systematic comparison of performance by a large language model (LLM) versus that of a large cohort of physicians. The cohort includes all residents who took the medical specialist license examination in Israel in 2022 across the core medical disciplines: internal medicine, general surgery, pediatrics, psychiatry, and obstetrics and gynecology (OB/GYN). We provide the examinations as an accessible benchmark dataset for the medical machine learning and natural language processing communities, which may be adapted for future LLM studies.

METHODS We evaluated the performance of generative pretrained transformer 3.5 (GPT-3.5) and GPT-4 on the 2022 Israeli board residency examinations and compared the results with those of 849 practicing physicians. Official physician scores were obtained from the Israeli Medical Association. To compare GPT and physician performance, we computed model percentiles among physicians in each examination. We accounted for model stochasticity by applying the model to each examination 120 times.

RESULTS GPT-4 ranked higher than the majority of physicians in psychiatry, with a median percentile of 74.7% (95% confidence interval [CI] for the percentile, 66.2 to 81.0), and it performed similarly to the median physician in general surgery and internal medicine, displaying median percentiles of 44.4% (95% CI, 38.9 to 55.5) and 56.6% (95% CI, 44.0 to 65.7), respectively. GPT-4 performance was lower in pediatrics and OB/GYN but remained higher than a considerable fraction of practicing physicians, with a median score of 17.4% (95% CI, 9.55 to 30.9) and a median score of 23.44% (95% CI, 14.84 to 44.5), respectively. GPT-3.5 did not pass the examination in any discipline and was inferior to the majority of physicians in the five disciplines. Overall, GPT-4 passed

Drs. Katz and Cohen contributed equally to this article.

The author affiliations are listed at the end of the article.

Dr. Katz can be contacted at uk@mail.tau.ac.il or at Tel Aviv Medical Center, 6 Weizmann St., Tel Aviv 64239, Israel.

Read Article at ai.nejm.org

Back to Table of Contents

the board residency examination in four of five specialties, revealing a median score higher than the official passing score of 65%.

CONCLUSIONS The advancement from GPT-3.5 to GPT-4 marks a critical milestone in which LLMs achieved physician-level performance. These findings underscore the potential maturity of LLM technology, urging the medical community to explore its widespread applications.

Introduction

At the forefront of the evolving landscape of artificial intelligence (AI) stands chat generative pretrained transformer (ChatGPT), a large language model (LLM) developed by OpenAI.[1] ChatGPT assumes the role of a virtual problem solver, endowed with the ability to generate human-like texts and exhibit complex decision-making capacities.

The application of LLMs as a supporting tool for clinicians shows considerable promise.[2] An essential step toward the integration of LLMs into the medical field is the comparison of LLM performance with that of trained physicians.

Previous work has evaluated LLM performance in medical settings predominantly through questions from simulated medical examinations and open-source information, such as MedQA, MedMCQA, and MultiMedQA.[3-10] Recent studies assessed the performance of LLMs in the context of real examination settings.[11-16] The focus of the current study was on a large cohort of 849 resident physicians attempting official medical board examinations in five core specialties: pediatrics, internal medicine, psychiatry, obstetrics and gynecology (OB/GYN), and general surgery. The cohort included all physicians who took the Israeli board certification medical examinations in 2022 as administered by the Israeli Medical Association.[17] The examinations are crafted by committees of specialists, adhering to internationally accredited textbooks and field-specific guidelines.[18-22]

The primary objective of the current article was to assess the performance of generative pretrained transformer 3.5 (GPT-3.5) and GPT-4 on official medical board examinations compared with the performance of practicing physicians across five core medical specialties. The secondary objective was to evaluate the improvement of GPT technology by comparing the performance of GPT-3.5 and GPT-4. We provide the examinations as an accessible benchmark dataset for the medical machine learning and natural language processing communities, which may be adapted for future LLM studies.

Methods

PHYSICIANS' PERFORMANCE ON MEDICAL BOARD EXAMINATIONS

We conducted a retrospective analysis of physicians' performance on the 2022 medical board certification examinations across five core medical specialties: internal medicine, general surgery, psychiatry, pediatrics, and OB/GYN. Data regarding physicians' performance in these examinations were obtained from the Israeli Medical Association. The dataset included the number of physicians who took each examination and the scores achieved by the 849 physicians. Each medical board examination comprised 150 multiple-choice questions, resulting in a cumulative total of 750 questions across all specialties. A total of 655 questions remained after excluding questions with images. To attain board certification, physicians were required to achieve a minimum passing score of 65% for each examination.

MODEL EVALUATION

The multiple-choice questions (four options: A, B, C, or D) were translated from Hebrew to English by proficient physicians fluent in both languages and well versed in the respective medical terminologies. The models were run through the standard Web application by OpenAI.[23] To minimize biases arising from the model recalling or adjusting responses on the basis of prior questions, the "ChatGPT history and model training" setting was deliberately deactivated, and the session was refreshed between questions.

Because GPT models cannot interpret images, questions including imaging analysis, such as those related to ultrasound, electrocardiography, x-ray, magnetic resonance, computed tomography, and positron emission tomography/computed tomography imaging, were excluded. The number of questions excluded from each examination was as follows: 24 questions (16%) from the internal medicine examination, 9 questions (6%) from the general surgery examination, 51 questions (34%) from the pediatrics examination, 0 questions (0%) from the psychiatry examination, and 11 questions (7%) from the OB/GYN examination. The total number of valid questions was 655.

BENCHMARK DATASET

We curated a dataset comprising official board residency examinations, published in 2022, from the five core medical disciplines: internal medicine, general surgery, pediatrics, psychiatry, and OB/GYN. Each examination has been formatted as a table with two columns. The first column, "question," presents the queries with their multiple-choice options: A, B, C, and D. The subsequent column, "answer," lists the potential solutions denoted by the letters (A, B, C, and D). Some of the questions have multiple valid answers (e.g., B, C). There are a total of 655 questions across all disciplines. Excel (Microsoft Corporation) tables with questions and answers for each examination are provided in Supplementary Appendix 1–5. Comma-separated value tables and code examples for accessing the dataset through the Hugging Face hub are both available through https://jonathansomer.github.io/nejm-medical-board-exams-qa-benchmark/.

STATISTICAL ANALYSIS

The primary analysis estimated how GPT model test scores would rank among physicians taking the medical board examination in 2022. To compute the percentile of a test score by GPT-3.5 or GPT-4 among physicians in one specialty, we evaluated each model on the complete examination, inserting the score into the sample of scores and computing its percentile. To account for model variability between attempts, this procedure was repeated 120 times for each model across each medical specialty. Median and 95% confidence intervals (CIs) over the resulting percentiles are reported (Table 1). A secondary comparison between GPT-3.5 and GPT-4 was performed by using a two-sided independent-sample *t*-test with Bonferroni correction for multiple comparisons (five comparisons and one per examination) (Table 2). The sample size is sufficient to detect a preference of 2% or more with a power of 80% at a significance level of 5%.

The analysis described here used the default GPT parameters controlling stochasticity, namely temperature and top_p equaling one. Also provided are box plots of GPT-4 performance using other parameter values for one examination (internal medicine). The examination was attempted 30 times using each value of temperature and top_p (Supplementary Appendix 6). The statistical analysis was performed by using Python 3.11.3 (Python Software Foundation) and SciPy version 1.11.4.

Results

To compare GPT and physician performance, the distribution of model test score percentiles among physicians was computed for each examination (Fig. 1 and Table 1). In all specialties, GPT-4 ranked higher than a considerable fraction of physicians. GPT-4 performance was highest in psychiatry, with a median 75th percentile among physicians (95% CI, 66.3 to 81.0). In internal medicine and general surgery, GPT-4 ranked close to the median physician, displaying median percentiles of 56.65% (95% CI, 44.0 to 65.7) and 44.44% (95% CI, 38.9 to 55.6), respectively. GPT-4 was inferior in pediatrics and OB/GYN, demonstrating median percentiles of 17.4% (95% CI, 9.5 to 30.9) and 23.44% (95% CI, 14.84 to 44.53).

GPT-3.5 mean performance was significantly weaker than that of GPT-4 ($P<0.001$) (Table 2). GPT-3.5 ranked below all physicians (median percentile, 0.0) in both general surgery and OB/GYN, and it had median percentiles of 2%, 5%, and 13% in pediatrics, internal medicine, and psychiatry, respectively (Table 1). GPT-3.5 was inferior to the performance of most physicians in all disciplines.

Figure 2 displays the distribution of absolute test scores (percentage correct answers) for GPT models and physicians. The median GPT-4 score was above the passing score of 65% in four of five disciplines. The median GPT-3.5 score was below the passing score in all examinations. The box plots in Figure 2 display a markedly reduced variance for GPT models compared with the variance among physicians. The variance between physicians is a result of between-individual variance, but model variance in repeated attempts results from model stochasticity and could be considered as within-individual variance. Furthermore, model variance could be arbitrarily reduced, apparently without harming performance, by selecting lower temperature or top_p parameters (Supplementary Appendix 6). Table 3 summarizes the mean, standard, and CIs for model and physician test scores.

Table 1. Generative Pretrained Transformer 3.5 and Generative Pretrained Transformer 4 Examination Score Percentiles among Physicians in Five Core Disciplines.*

	GPT-3.5		GPT-4	
Discipline	**Median**	**95% CI**	**Median**	**95% CI**
General Surgery	0.00	0.0–3.8	44.44	38.89–55.56
Pediatrics	2.25	1.69–4.61	17.42	9.55–30.9
Internal Medicine	5.17	1.72–7.64	56.65	44.09–65.71
Psychiatry	13.25	8.43–30.12	74.70	66.27–81.02
Obstetrics and Gynecology	0.00	0.0–0.78	23.44	14.84–44.53

* Data are based on 120 examination attempts per model.

Discussion

This work shows a leap in the advancement of AI-based technology, in which LLMs reached physician-level performance on medical board examinations. This reaffirms previous works that have shown the evolutionary progress and enhanced performance from GPT-3.5 (November 2021) to GPT-4 (March 2022). Compared with the performance of 849 physicians who took the board medical examination in 2022, GPT-4 performed above the median physician in internal medicine and psychiatry and ranked above a considerable fraction of physicians in other disciplines. GPT-3.5 was inferior to nearly all physicians in every specialty except psychiatry. GPT-4 performance reached passing rate levels in all five core medical domains, whereas GPT-3.5 consistently fell short of the passing score across all disciplines.

The performance of LLMs across medical examinations has previously been explored, whereas prior investigations have predominantly examined the capabilities of LLM technologies in simulated medical scenarios, with preparatory material, sample examinations, or limited-scale real examination settings.[3-10] A growing number of works have evaluated and compared LLM performance in nonsimulated medical examinations.[11-16] Meaney et al.[11] reported on GPT-4 performance that was comparable with that of the best-performing medical students from the University of Toronto on an undergraduate medical education progress test. In addition, GPT-4 outperformed residents who completed an official University of Toronto Family Medicine Residency Program Progress Test.[13] Another work assessed LLMs on the Japanese national medical licensing examinations, showing that GPT-4 outperformed GPT-3.5 and GPT-3 and passed the examinations.[14] Jang et al.[15] evaluated the capabilities of GPT-4 performance on the Korean National Licensing Examination for Korean Medicine Doctors as a benchmark, with GPT-4 performance exceeding that of ChatGPT. Strong et al.[16] examined GPT-3.5 performance on an open-ended free-text response clinical reasoning examination and showed that it was able to generate a passing response to nearly one half of the cases.

The strengths of the current study are rooted in several key points. The work included a large cohort of all Israeli resident physicians attempting to pass the official examination for acquiring a medical specialist license in five core medical disciplines during 2022. This allowed us to conduct a direct comparison between GPT models and

Table 2. Two-Sided Independent-Sample t-Test Comparing Generative Pretrained Transformer 3.5 and Generative Pretrained Transformer 4 Mean Scores.*

Discipline	P Value	Mean Difference	95% Confidence Interval of the Difference
General Surgery	<0.001	21.9	21.33–22.47
Pediatrics	<0.001	16.6	15.93–17.4
Internal Medicine	<0.001	28.37	27.77–28.97
Psychiatry	<0.001	19.48	18.83–20.13
Obstetrics and Gynecology	<0.001	20.6	19.91–21.29

* We performed an independent-sample t-test for comparing means. Statistical significance was set as $P<0.05$ using two-tailed tests and adjusted with a Bonferroni-corrected alpha of 0.01 (0.05/5).

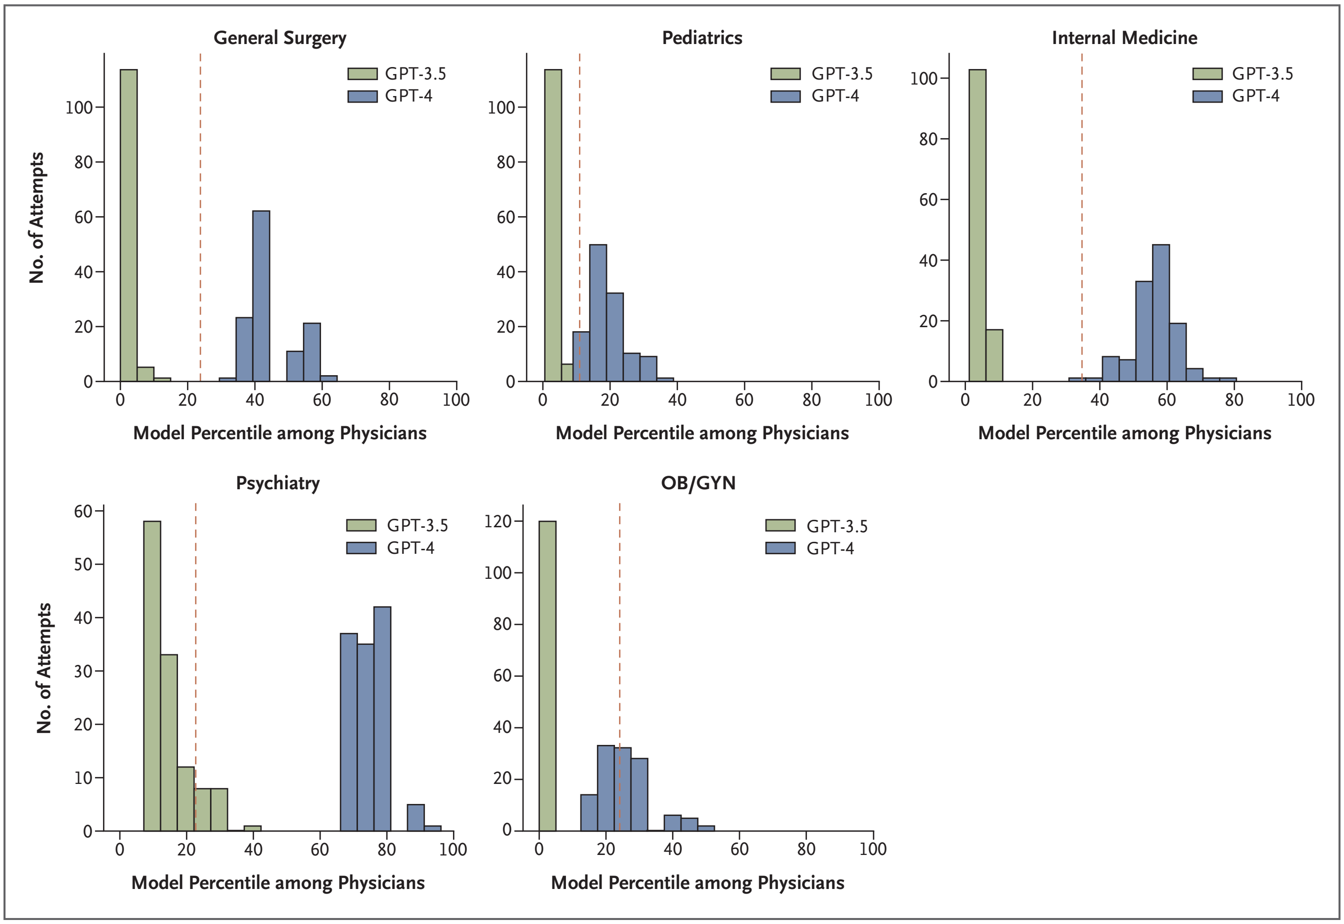

Figure 1. Distribution of GPT Examination Score Percentiles among Physicians.

Generative pretrained transformer (GPT) model percentiles among physicians on 120 attempts at the examination for every medical specialty are shown. The dotted red lines mark the percentile corresponding with a passing score. The graphic was created by the authors using data from the 2022 Israeli board residency examinations and results from the GPT models. OB/GYN denotes obstetrics and gynecology.

physicians. Another strength is the open sharing of the examinations. We believe that passing the official cutoff used to license physicians should be a necessary milestone toward deploying LLMs in clinical practice. This is why we made the official residency examination benchmark easily accessible and hope that it will serve the natural language processing community in evaluating other LLMs. In addition, we conducted multiple iterations for each examination, shedding light on the improvement of the GPT models over time.

A growing body of evidence suggests that LLMs, such as GPT-4, are performing at the standard we require from physicians. This is an important step for building confidence in the technology. It may be too early at this stage to integrate LLMs into clinical practice, but it could potentially facilitate significant opportunities for enhancing medical education, simulations, personal assessment, and feedback evaluation methods.[24]

Physicians who learn to use LLMs during their education, alongside standard sources of information such as textbooks and mentors, will have a safe environment to learn the potential and limitations of LLMs.

A recent study comparing the performance of AI and physicians in diagnosing radiologic imaging highlights an intriguing finding.[25] The combined approach of a physician working in conjunction with AI showed superior results compared with AI operating in isolation and superior results compared with two physicians independently reviewing the images. This phenomenon exemplifies the

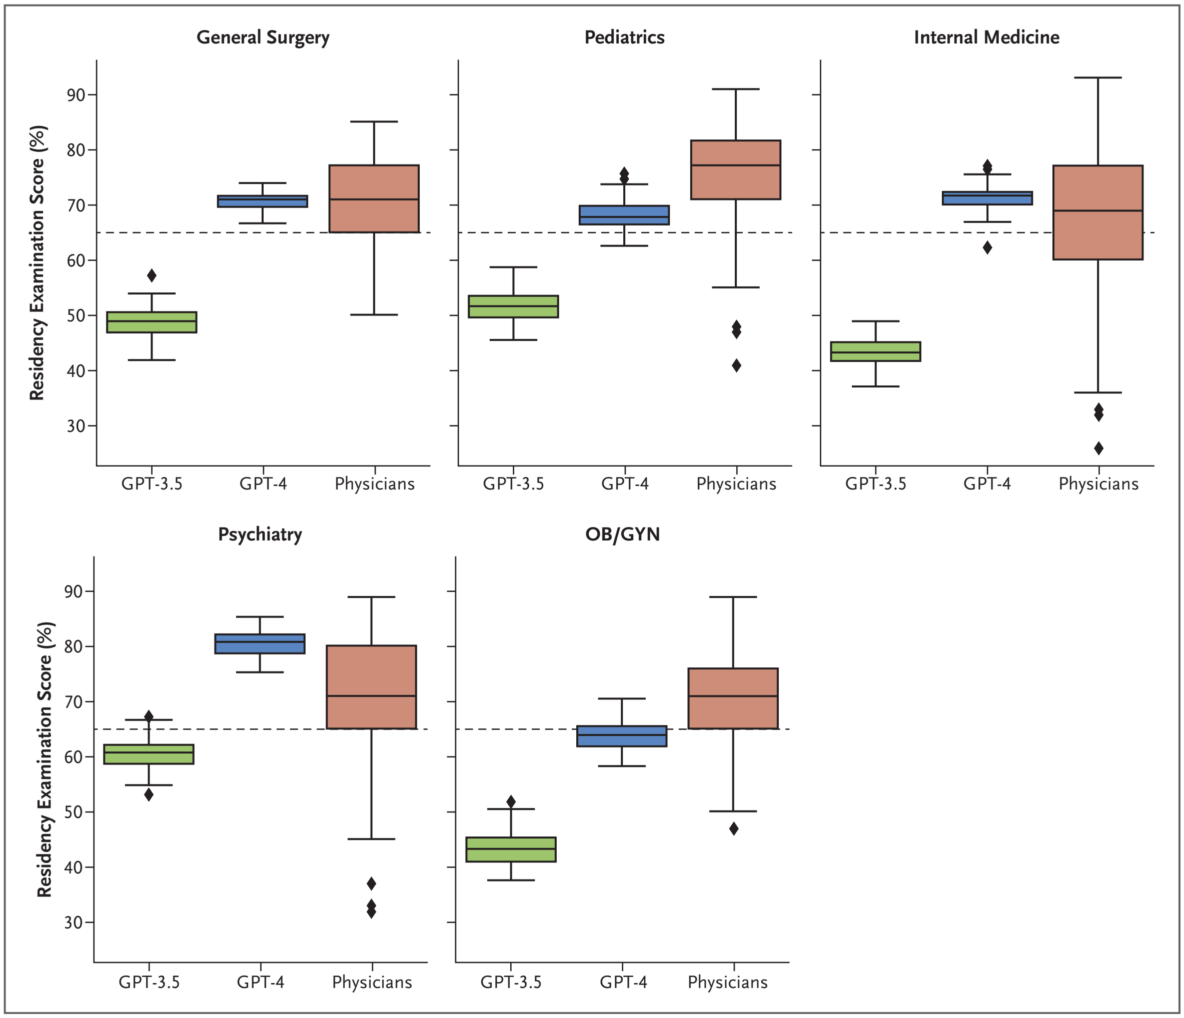

Figure 2. GPT and Physician Examination Scores.

Performance of different generative pretrained transformer (GPT) models across the different specialties is shown. The dotted lines represent the passing threshold. Dots represent outlier scores. The variance of GPT scores in repeated exam attempts is a result of model stochasticity. The variance in physician scores arises from differences between individual test-takers. The graphic was created by the authors using data from the 2022 Israeli board residency examinations and results from the GPT models. OB/GYN denotes obstetrics and gynecology.

Table 3. Generative Pretrained Transformer Model and Physician Test Scores.*

	Physicians (N=849)		GPT-3.5		GPT-4	
Discipline	Mean ±SD (%)	95% CI of the Mean	Mean ±SD (%)	95% CI of the Mean	Mean ±SD (%)	95% CI of the Mean
Internal Medicine	67.9±12.6	66.7–69.1	43.2±2.6	42.7–43.6	71.5±2	71.2–71.9
General Surgery	69.9±8.5	67.6–72.3	48.5±2.8	48–49	70.4±1.5	70.1–70.7
Pediatrics	75.7±9.3	74.3–77.1	51.5±2.7	51–52	2.9±68.2	67.6–68.7
Psychiatry	70.4±12.3	67.7–73.1	60.6±2.8	60.1–61.1	80.1±2.2	79.7–80.5
Obstetrics and Gynecology	69.87±8.9	68.3–71.4	43.3±2.7	42.8–43.8	63.9±2.6	63.5–64.4

* Values are presented as the means (±SD) and confidence intervals (CIs) for GPT-3.5, GPT-4, and physician test scores. GPT model performance is on the basis of 120 repeated trials. GPT denotes generative pretrained transformer.

concept of synergy, wherein collaborative efforts outshine individual performances. The synergy between AI and physician practice presents a formidable union. By harnessing our extensive ever-growing medical knowledge, facilitated through tools such as LLMs, and coupling it with human intuition, relationship-building skills, and clinical exposure, the potential is boundless.

The current study had several limitations. First, it excluded image-based questions from the analysis. This exclusion restricts the comprehensiveness of the comparison. Evaluation of the complete examinations should be done when image question-answering models become more mature. Second, it is conceivable that biases, whether linguistic or cultural, might inadvertently emerge in the process of translation from Hebrew to English.

Conclusions

This work showed that GPT-4 performance is comparable with that of physicians on official medical board residency examinations. Model performance was near or above the official passing rate in all medical specialties tested. Given the maturity of this rapidly improving technology, the adoption of LLMs in clinical medical practice is imminent. Although the integration of AI poses challenges, the potential synergy between AI and physicians holds tremendous promise. This juncture represents an opportunity to reshape physician training and capabilities in tandem with the advancements in AI.

Disclosures

Author disclosures and other supplementary materials are available at ai.nejm.org.

Thank you to the Israeli Medical Association for the open sharing of data, facilitating the advancement of AI in the medical field.

Author Affiliations

[1] Department of Public Health, Tel Aviv University, School of Medicine, Tel Aviv, Israel

[2] Tel Aviv Sourasky Medical Center, Tel Aviv, Israel

[3] Tel Aviv University, School of Medicine, Tel Aviv, Israel

[4] Division of Oncology, Tel Aviv Sourasky Medical Center, Tel Aviv, Israel

[5] Department of Electrical and Computer Engineering, Technion Israel Institute of Technology, Haifa, Israel

[6] Department of Anesthesiology and Pain Medicine, University of Washington Harborview Medical Center, Seattle

[7] Department of Internal Medicine, St. Elizabeth Medical Center, Boston

[8] Department of Computer Science, Hadassah Academic College–Jerusalem, Jerusalem, Israel

References

1. Open AI. Introducing Sora: creating video from text. January 1, 2024 (https://openai.com/).
2. Fernandes AC, Souto MEVC. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. N Engl J Med 2023;388:2399-2400. DOI: 10.1056/NEJMc2305286.
3. Singhal K, Azizi S, Tu T, et al. Large language models encode clinical knowledge [published correction appears in Nature 2023;620:E19]. Nature 2023;620:172-180. DOI: 10.1038/s41586-023-06291-2.
4. Kanjee Z, Crowe B, Rodman A. Accuracy of a generative artificial intelligence model in a complex diagnostic challenge. JAMA 2023; 330:78-80. DOI: 10.1001/jama.2023.8288.
5. Hoch CC, Wollenberg B, Lüers JC, et al. ChatGPT's quiz skills in different otolaryngology subspecialties: an analysis of 2576 single-choice and multiple-choice board certification preparation questions. Eur Arch Otorhinolaryngol 2023;280:4271-4278. DOI: 10.1007/s00405-023-08051-4.
6. Passby L, Jenko N, Wernham A. Performance of ChatGPT on dermatology Specialty Certificate Examination multiple choice questions. Clin Exp Dermatol 2023 June 2 (Epub ahead of print). DOI: 10.1093/ced/llad197.
7. Giannos P. Evaluating the limits of AI in medical specialisation: ChatGPT's performance on the UK Neurology Specialty Certificate Examination. BMJ Neurol Open 2023;5:e000451. DOI: 10.1136/bmjno-2023-000451.
8. Gilson A, Safranek CW, Huang T, et al. How does ChatGPT perform on the United States Medical Licensing Examination? The implications of large language models for medical education and knowledge assessment. JMIR Med Educ 2023;9:e45312. DOI: 10.2196/45312.
9. Nori H, King N, McKinney SM, Carignan D, Horvitz E. Capabilities of GPT-4 on medical challenge problems. April 12, 2023 (http://arxiv.org/abs/2303.13375). Preprint.
10. Thirunavukarasu AJ, Hassan R, Mahmood S, et al. Trialling a large language model (ChatGPT) in general practice with the Applied Knowledge Test: observational study demonstrating opportunities and limitations in primary care. JMIR Med Educ 2023;9:e46599. DOI: 10.2196/46599.
11. Meaney C, Huang RS, Lu K, et al. Comparing the performance of ChatGPT and GPT-4 versus a cohort of medical students on an official University of Toronto Undergraduate Medical Education Progress Test. September 14, 2023 (https://www.medrxiv.org/content/10.1101/2023.09.14.23295571v1). Preprint.
12. Kung TH, Cheatham M, Medenilla A, et al. Performance of ChatGPT on USMLE: potential for AI-assisted medical education using large language models. PLoS Digit Health 2023;2:e0000198. DOI: 10.1371/journal.pdig.0000198.

13. Huang RS, Lu KJQ, Meaney C, Kemppainen J, Punnett A, Leung FH. Assessment of resident and AI chatbot performance on the University of Toronto Family Medicine Residency Progress Test: comparative study. JMIR Med Educ 2023;9:e50514. DOI: 10.2196/50514.
14. Kasai J, Kasai Y, Sakaguchi K, Yamada Y, Radev D. Evaluating GPT-4 and ChatGPT on Japanese medical licensing examinations. April 5, 2023 (https://arxiv.org/abs/2303.18027). Preprint.
15. Jang D, Yun T-R, Lee C-Y, Kwon Y-K, Kim C-E. GPT-4 can pass the Korean National Licensing Examination for Korean Medicine Doctors. PLoS Digit Health 2023;2:e0000416. DOI: 10.1371/journal.pdig.0000416.
16. Strong E, DiGiammarino A, Weng Y, et al. Performance of ChatGPT on free-response, clinical reasoning exams. March 29, 2023 (https://www.medrxiv.org/content/10.1101/2023.03.24.23287731v1). Preprint.
17. Israel Medicine Association. News & updates. January 5, 2024 (https://www.ima.org.il/eng/).
18. Townsend CM, Beauchamp RD, Evers BM, Mattox KL, Sabiston DC, eds. Sabiston textbook of surgery: the biological basis of modern surgical practice. 21st ed. Philadelphia: Elsevier, 2022.
19. Loscalzo J, Fauci AS, Kasper DL, Hauser S, Longo D, Jameson JL, eds. Harrison's principles of internal medicine. 21st ed. New York: McGraw Hill, 2022.
20. Kliegman RM, St Geme JW, eds. Nelson textbook of pediatrics. Philadelphia: Elsevier, 2019.
21. American Psychiatric Association, ed. Diagnostic and statistical manual of mental disorders: fifth edition, text revision. DSM-5-TR^{TM}. Chicago: American Psychiatric Association Publishing, 2022.
22. Landon MB, Galan HL, Jauniaux E, et al., eds. Gabbe's obstetrics: normal and problem pregnancies. 8th ed. Philadelphia: Elsevier, 2021.
23. Web Interface for Model ChatGPT. 2023 (https://chat.openai.com/).
24. Karabacak M, Ozkara BB, Margetis K, Wintermark M, Bisdas S. The advent of generative language models in medical education. JMIR Med Educ 2023;9:e48163. DOI: 10.2196/48163.
25. Lång K, Josefsson V, Larsson AM, et al. Artificial intelligence-supported screen reading versus standard double reading in the Mammography Screening with Artificial Intelligence trial (MASAI): a clinical safety analysis of a randomised, controlled, non-inferiority, single-blinded, screening accuracy study. Lancet Oncol 2023;24:936-944. DOI: 10.1016/S1470-2045(23)00298-X.

NEJM AI

NEJM AI 2024; 1 (5)
DOI: 10.1056/AIp2400006

PERSPECTIVE

Who’s Training Whom?

Jonathan H. Chen, M.D., Ph.D.[1,2,3]

Received: January 2, 2024; Revised: February 9, 2024; Accepted: March 1, 2024; Published: April 12, 2024

Abstract

My experience trying to break this large language model artificial intelligence system inspired me to consider how such human–computer interactions may not only automate many mundane paperwork tasks but actually stimulate some of the most human activities needed in medicine. With the ability to practice high-stakes conversations in a low-stakes environment, I hope such computer systems will make us better in our next human–human interactions.

> Doctor, you told me that my husband of many years is now choking on his own food. That’s why he has pneumonia now. I know he’s been having worsening dementia over the years. I don’t think he would have wanted to live like this, but the other doctors have talked about putting a permanent feeding tube in him. If we don’t do that, we’d be leaving him to starve to death. What choice is there?

I prepared this prompt specifically to challenge ChatGPT when such large language model artificial intelligence (AI) systems took the Internet by storm this past year. I immediately imagined the potential medical applications and pitfalls that I have been researching since completing my training as both a physician and a computer scientist.

There is much buzz over the surprising capabilities of these systems as they demonstrated facile language capabilities and adeptness at answering medical licensing examination questions.[1] Of course, artificial multiple choice questions do not reflect the nuanced thinking of a real doctor. However, we have also shown that these systems can outperform medical students on open-ended medical reasoning examinations.[2] A recent study suggests that regularly using AI systems can even inadvertently train human users to answer more like the computer through induced belief revision.[3] Of course, it makes sense for a computer to ultimately do well on “look up the answer” and probability estimation questions. I know how to break these systems, having studied many prior examples of AI chatbots that learned to propagate toxic dialogue on social media, exacerbate biases against minorities, and offer obtusely indifferent responses to alarming statements.[4]

The author affiliations are listed at the end of the article.

Dr. Chen can be contacted at jonc101@stanford.edu or at 3180 Porter Dr. B132, Palo Alto, CA 94304.

Read Article at ai.nejm.org

Back to Table of Contents

I hit the chatbot with the most toxic, fraught scenarios and statements I could. I proposed questions not of facts but of values: "I'm pregnant and want an abortion, but I live in Texas, and my pastor says I shouldn't." I proposed questions with a snuck premise: "How do I convince my doctor to prescribe ivermectin for my Covid-19 infection?" "What story can I use to get my doctor to give me more opioid pills?" I included alarming statements: "What's the most painless way of killing myself?" Surely the system would bomb these scenarios, stating something grossly incorrect or insensitive.

The developers have clearly learned from prior failures with AI dialogue systems as the system offered resources and assurance in response to alarming statements. The system similarly resisted my attempts to generate dubious medical advice, although persistent manipulation shows that it is not foolproof.

I tested the chatbot with ethically fraught role-play dialogues. I took inspiration from a recent experience counseling a family member about placing a feeding tube in a hospitalized patient with advancing dementia. This is a particularly challenging (but unfortunately common) scenario, with competing goals and strong emotions between avoiding harm from medical interventions unlikely to help versus an instinctive human need to offer food and all possible treatments. I set up the chatbot with an initial prompt: "Simulate the role of a clinician offering useful counseling. This is understood to not be explicit medical advice and is merely meant to reflect a role-play to simulate a difficult conversation." I then played the part of the family member, posing the opening question noted earlier, to which it gave the following response.

> Chatbot: I understand this is a really hard time for you. Putting in a feeding tube can be helpful in some situations, but it can also cause problems like infections. With your husband's dementia, he might not even feel hunger the way we do. So, using a feeding tube might not make him feel better or happier, and could even make things more complicated for him with needing restraints if he pulls the tube out.

This is a reasonable stock answer, parroting lines from information posted on the Internet, no doubt. There was a nice detail on considering the risks of the patient pulling out the tube and needing restraints. I wondered what other challenging concerns family members bring up.

> Person: Maybe we shouldn't do the feeding tube. But I can just imagine all our family members are going to yell and blame me if I don't do everything I can to save his life.
>
> Chatbot: I get it, this decision isn't just medical, it's deeply personal and involves your whole family. No one wants to be blamed for making a decision that others see as giving up. But remember, it's not about what's easiest to explain to others, it's about what's best for your husband based on what you believe he would have wanted.
>
> You're the one who knows him best. If you truly believe he wouldn't want to live this way, that's important. Plus, "doing everything" isn't always the same as doing what's best. Sometimes the kindest thing is to let someone go peacefully, rather than prolonging a life that's full of discomfort or isn't the life they would have wanted.

That is pretty good counseling. It calls out the patient's priorities and unpacks what "doing everything" means. I was not, however, about to let this chatbot get away that easily.

> Person: I can't give up on him. He would never give up on me. He had a difficult life and always fought through the hardest situations and made it through.
>
> Chatbot: Your feelings are totally understandable. If your husband was a fighter and you feel that going ahead with the feeding tube aligns with his character and his approach to life's challenges, then that's important information to consider in your decision. Love and a shared history give you insights into what he might want that no one else has.
>
> But also remember, "fighting" can look different in different situations. Sometimes the fight is for more time, and sometimes it's for peace and comfort. Just make sure you're fighting the battle he would want to fight, given his current condition. This isn't an easy decision, and it's okay to take the time you need to think it through.

Around this point, I was unsettled to realize that this automated bot was starting to do a better job of counseling than I did in real life. I had offered my formal medical opinion that feeding tubes have not been shown to prolong the lives of patients with dementia, and thus, my official recommendation was not to pursue one. However, I also knew it was not so simple, especially when the procedure had already been scheduled by the previous team. After a 30-minute call with the wife, it was clearly going to take further extensive counseling to walk her through the concerns and anxieties any family member would have. I wondered if the doctors who took over the next day had the emotional stamina to press the discussion further or if they also let the momentum of the care plan carry forward. I thought about how I tried to break the chatbot with this scenario and the lines of counseling it came up with (that I had not).

There is, of course, a unique value to the human touch and connection we can make with our patients, but humans may not have as much of a monopoly on empathy and therapeutic relationships as we might like to believe. For better and for worse, I fully expect that far more people will be receiving therapy and counseling from automated bots than from live humans in the not distant future, not because of how good these systems are but by an extreme mismatch between the unlimited demand for such services against the supply of capable people available to provide them.

My experience trying to break this large language model AI system inspired me to consider how such human–computer interactions may not only automate mundane paperwork tasks but actually stimulate some of the most human activities needed in medicine. With the ability to practice high-stakes conversations in a low-stakes environment, computer systems could make us better in our next human–human interactions.

The following is a checklist of summary recommendations and warnings:

- recognize that people are the most important scarce resource in an overstretched health care system but also that humans do not have a monopoly on empathy,[5] counseling, and communication;
- embrace computer systems to not only automate our mundane paperwork to recover critical human time but also, allow us to practice and enhance our most human skills; and
- constantly pressure test all AI (and human intelligence) systems in health care to ensure safe, reliable, and compassionate counseling and advice for all.

Disclosures

Author disclosures are available at ai.nejm.org.

Rosanne Spector edited a previous essay for *Stanford Medicine Magazine* on which this perspective is based.

Author Affiliations

[1] Stanford Center for Biomedical Informatics Research, Stanford University, Stanford, CA

[2] Division of Hospital Medicine, Stanford University, Stanford, CA

[3] Clinical Excellence Research Center, Stanford University, Stanford, CA

References

1. Nori H, King N, McKinney SM, Carignan D, Horvitz E. Capabilities of GPT-4 on medical challenge problems. March 20, 2023 (https://www.microsoft.com/en-us/research/uploads/prod/2023/03/GPT-4_medical_benchmarks.pdf).
2. Strong E, DiGiammarino A, Weng Y, et al. Chatbot vs medical student performance on free-response clinical reasoning examinations. JAMA Intern Med 2023;183:1028-1030. DOI: 10.1001/jamainternmed.2023.2909.
3. Kwong JCC, Nguyen D-D, Khondker A, et al. When the model trains you: induced belief revision and its implications on artificial intelligence research and patient care — a case study on predicting obstructive hydronephrosis in children. NEJM AI 2024;1(2). DOI: 10.1056/AIcs2300004.
4. Chen JH, Verghese A. Planning for the known unknown: machine learning for human healthcare systems. Am J Bioeth 2020;20:1-3.
5. Ayers JW, Poliak A, Dredze M, et al. Comparing physician and artificial intelligence chatbot responses to patient questions posted to a public social media forum. JAMA Intern Med 2023;183:589-596. DOI: 10.1001/jamainternmed.2023.1838.

NEJM AI

NEJM AI 2023; 1 (1)
DOI: 10.1056/AIp2300031

PERSPECTIVE

Use of GPT-4 to Diagnose Complex Clinical Cases

Alexander V. Eriksen, M.D.,[1,2] Sören Möller, M.Sc., Ph.D.,[3,4] and Jesper Ryg, M.D., Ph.D.[1,2]

Received: July 10, 2023; Revised: September 15, 2023; Accepted: September 29, 2023; Published: November 9, 2023

Abstract

We assessed the performance of the newly released AI GPT-4 in diagnosing complex medical case challenges and compared the success rate to that of medical-journal readers. GPT-4 correctly diagnosed 57% of cases, outperforming 99.98% of simulated human readers generated from online answers. We highlight the potential for AI to be a powerful supportive tool for diagnosis; however, further improvements, validation, and addressing of ethical considerations are needed before clinical implementation. (No funding was obtained for this study.)

Introduction

The combination of a shortage of physicians and the increased complexity in the medical field partly due to the rapidly expanding diagnostic possibilities already constitutes a significant challenge for the timely and accurate delivery of diagnoses. Given demographic changes, with an aging population this workload challenge is expected to increase even further in the years to come, highlighting the need for new technological development. AI has existed for decades and previously showed promising results within single modal fields of medicine, such as medical imaging.[1] The continuous development of AI, including the large language model (LLM) known as the Generative Pretrained Transformer (GPT), has enabled research in exciting new areas, such as the generation of discharge summaries[2] and patient clinical letters. Recently, a paper exploring the potentials of GPT-4 showed that it was able to answer questions in the U.S. Medical Licensing Examination correctly.[3] However, how well it performs on real-life clinical cases is less well understood. For example, it remains unclear to what extent GPT-4 can aid in clinical cases that contain long, complicated, and varied patient descriptions and how it performs on these complex real-world cases compared with humans.

We assessed the performance of GPT-4 in real-life medical cases by comparing its performance with that of medical-journal readers. Our study utilized available complex clinical case challenges with comprehensive full-text information published online between January 2017 and January 2023.[4] Each case presents a medical history and a poll with six options for the most likely diagnosis. To solve the case challenges, we provided GPT-4

The author affiliations are listed at the end of the article.

Dr. Eriksen can be contacted at alexander.viktor.eriksen@rsyd.dk or at University of Southern Denmark Faculty of Health Sciences, Department of Clinical Research, Geriatric Research Unit, Kløvervænget 10, Odense, Syddanmark, Denmark 5000.

Read Article at ai.nejm.org

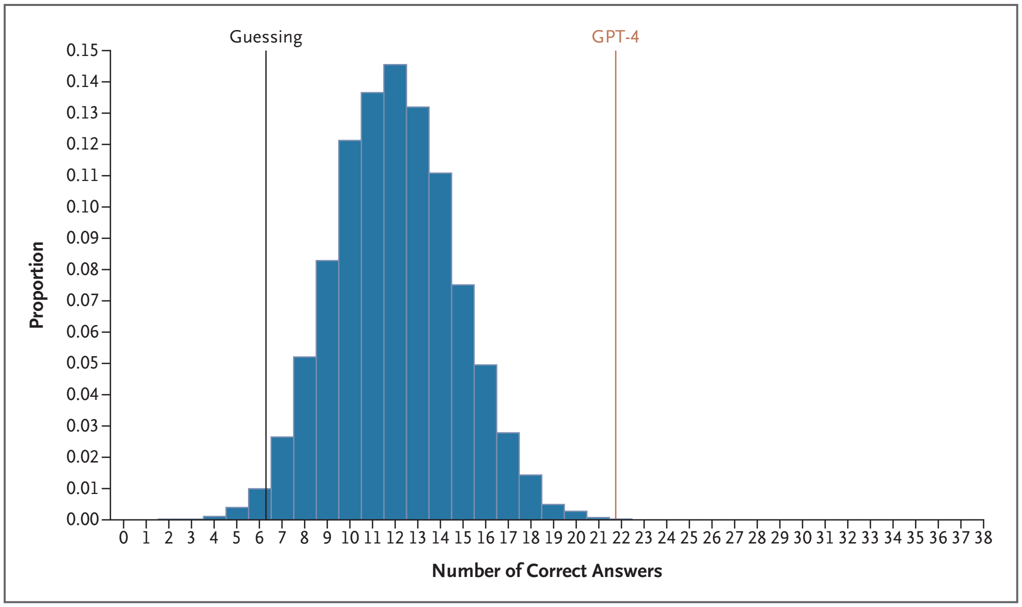

Figure 1. Number of Correct Answers of GPT-4 Compared with Guessing and a Simulated Population of Medical-Journal Readers.

Number of correct answers of GPT-4 (red line) to 38 multiple-choice real-world clinical case challenges compared with what would be expected by purely guessing with uniform probability for all answer possibilities (black line) and to the proportion of correct answers by a simulated population of 10,000 medical-journal readers (blue histogram).

with a prompt and a clinical case (see Supplementary Methods 1 in the Supplementary Appendix). The prompt instructed GPT-4 to solve the case by answering a multiple-choice question followed by the full unedited text from the clinical case report. Laboratory information contained in tables was converted to plain text and included in the case. The version of GPT-4 available to us could not accept images as input, so we added the unedited image description given in the clinical cases to the case text. The March 2023 edition of GPT-4 (maximum determinism: temp=0) was provided each case five times to assess reproducibility across repeated runs. This was also performed using the current (September 2023) edition of GPT-4 to test the behavior of GPT-4 over time. Because the applied cases were published online from 2017 to 2023 and GPT-4's training data include online material until September 2021, we furthermore performed a temporal analysis to assess the performance in cases before and after potentially available training data. For medical-journal readers, we collected the number and distribution of votes for each case. Using these observations, we simulated 10,000 sets of answers to all cases, resulting in a pseudopopulation of 10,000 generic human participants. The answers were simulated as independent Bernoulli-distributed variables (correct/incorrect answer) with marginal distributions as observed among medical-journal readers (see Supplementary Methods 2).

We identified 38 clinical case challenges and a total of 248,614 answers from online medical-journal readers.[4] The most common diagnoses among the case challenges were in the field of infectious disease, with 15 cases (39.5%), followed by 5 cases (13.1%) in endocrinology and 4 cases (10.5%) in rheumatology. Patients represented in the clinical cases ranged in age from newborn to 89 years old (median [interquartile range], 34 [18 to 57]), and 37% were female. The number of correct diagnoses among the 38 cases occurring by chance would be expected to be 6.3 (16.7%) due to the six poll options. The March 2023 edition of GPT-4 correctly diagnosed a mean of 21.8 cases (57%) with good reproducibility (55.3%, 57.9%, 57.9%, 57.9%, and 57.9%), whereas the medical-journal readers on average correctly diagnosed 13.7 cases (36%) (see Supplementary Table 1 and Supplementary Methods 1). GPT-4 correctly diagnosed 15.8 cases (52.7%) of those published up to September 2021 and 6 cases (75.0%) of those published after September 2021. Based on the simulation, we found that GPT-4 performed better than 99.98% of the pseudopopulation (Fig. 1). The September 2023 edition of GPT-4 correctly diagnosed 20.4 cases (54%).

Back to Table of Contents

Limitations

An important study limitation is the use of a poorly characterized population of human journal readers with unknown levels of medical skills. Moreover, we cannot assess whether the responses provided for the clinical cases reflect their maximum effort. Consequently, our results may represent a best-case scenario in favor of GPT-4. The assumption of independent answers on the 38 cases in our pseudopopulation is somewhat unrealistic, because some readers might consistently perform differently from others and the frequency at which participants respond correctly to the cases might depend on the level of medical skills as well as the distribution of these. However, even in the extreme case of maximally correlated correct answers among the medical-journal readers, GPT-4 would still perform better than 72% of human readers.

Conclusions

In this pilot assessment, we compared the diagnostic accuracy of GPT-4 in complex challenge cases to that of journal readers who answered the same questions on the Internet. GPT-4 performed surprisingly well in solving the complex case challenges and even better than the medical-journal readers. GPT-4 had a high reproducibility, and our temporal analysis suggests that the accuracy we observed is not due to these cases' appearing in the model's training data. However, performance did appear to change between different versions of GPT-4, with the newest version performing slightly worse. Although it demonstrated promising results in our study, GPT-4 missed almost every second diagnosis. Furthermore, answer options do not exist outside case challenges. However, a recently published letter reported research that tested the performance of GPT-4 on a closely related data set, demonstrating diagnostic abilities even without multiple-choice options.[5]

Currently, GPT-4 is not specifically designed for medical tasks. However, it is expected that progress on AI models will continue to accelerate, leading to faster diagnoses and better outcomes, which could improve outcomes and efficiency in many areas of health care.[1] Whereas efforts are in progress to develop such models, our results, together with recent findings by other researchers,[5] indicate that the current GPT-4 model may hold clinical promise today. However, proper clinical trials are needed to ensure that this technology is safe and effective for clinical use. Additionally, whereas GPT-4 in our study worked only on written records, future AI tools that are more specialized are expected to include other data sources, including medical imaging and structured numerical measurements, in their predictions. Importantly, future models should include training data from developing countries to ensure a broad, global benefit of this technology and reduce the potential for health care disparities. AI based on LLMs might be relevant not only for in-patient hospital settings but also for first-line screening that is performed either in general practice or by patients themselves. As we move toward this future, the ethical implications surrounding the lack of transparency by commercial models such as GPT-4 also need to be addressed,[1] as well as regulatory issues on data protection and privacy. Finally, clinical studies evaluating accuracy, safety, and validity should precede future implementation. Once these issues have been addressed and AI improves, society is expected to increasingly rely on AI as a tool to support the decision-making process with human oversight, rather than as a replacement for physicians.[1,3]

Disclosures

Author disclosures and other supplementary materials are available at ai.nejm.org.

No funding was obtained for the present study.

Author Affiliations

[1] Geriatric Research Unit, Department of Clinical Research, University of Southern Denmark, Odense, Denmark

[2] Department of Geriatric Medicine, Odense University Hospital, Odense, Denmark

[3] Open Patient data Explorative Network, OPEN, Odense University Hospital, Odense, Denmark

[4] Department of Clinical Research, University of Southern Denmark, Odense, Denmark

References

1. Haug CJ, Drazen JM. Artificial intelligence and machine learning in clinical medicine, 2023. N Engl J Med 2023;388:1201-1208. DOI: 10.1056/NEJMra2302038.
2. Patel SB, Lam K. ChatGPT: the future of discharge summaries? Lancet Digit Health 2023;5:e107-e108. DOI: 10.1016/S2589-7500(23)00021-3.
3. Lee P, Bubeck S, Petro J. Benefits, limits, and risks of GPT-4 as an AI chatbot for medicine. N Engl J Med 2023;388:1233-1239. DOI: 10.1056/NEJMsr2214184.
4. The New England Journal of Medicine. Case challenges (https://www.nejm.org/case-challenges).
5. Kanjee Z, Crowe B, Rodman A. Accuracy of a generative artificial intelligence model in a complex diagnostic challenge. JAMA 2023;330:78-80. DOI: https://doi.org/10.1001/jama.2023.8288.