Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Equitable carotid ultrasound screening for Indigenous Peoples through open artificial intelligence: a community-engaged proof-of-concept from the PAI Project (Brazil)

Américo Mota

americomota@gmail.com

State University of Bahia (UNEB)

Rodrigo Ramos

Federal University of the São Francisco Valley (UNIVASF)

Bruna Almeida

AMMO Institute / CDI

Isabella Diniz Souza

AMMO Institute / CDI

Jandir Nicacio

State University of Bahia (UNEB)

Vanessa Pereira

Federal University of the São Francisco Valley (UNIVASF)

Thiago Carvalho

AMMO Institute / CDI

Paula Maduro

University Hospital of the Federal University of the São Francisco Valley (HU-UNIVASF / EBSERH)

Paulo Schwingel

University of Pernambuco (UPE)

Carlos Souza

Federal University of the São Francisco Valley (UNIVASF)

Orlando Gomes

State University of Bahia (UNEB)

Ana Pinto

AMMO Institute / CDI

Caroline Chieco

AMMO Institute / CDI

Maria Santana

AMMO Institute / CDI

Alane Santos
Federal University of the São Francisco Valley (UNIVASF)
Anderson Armstrong
State University of Bahia (UNEB)

Research Article

Keywords: Health equity, Indigenous health, Artificial intelligence, Cardiovascular screening, Carotid ultrasound, Tele-ultrasound, Brazil, Truká, Fulni-ô, Underserved populations

Posted Date: May 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-9815931/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Table of Contents

Abstract....4
Background....6
Cardiovascular inequities in Indigenous Peoples....6
Artificial intelligence: opportunity and risk for health equity....7
The PAI Project as a platform for equitable AI development....9
Objective and scope....10
Methods....11
Study design and ethical framework....11
Scope of the present analysis and rationale for non-stratified design....13
Equity rationale for methodological choices....14
Image acquisition across the PAI program....15
Methodological arm 1 — Open AI for automated longitudinal cIMT measurement....17
Mask R-CNN architecture and training....18
Post-processing: comparative evaluation of eight techniques....19
Automated cIMT estimation....21
Cross-sectional segmentation: Viola-Jones cascade detector....21
Mask R-CNN applied to cross-sectional views....22
Mean radius computation....22
Methodological arm 2 — Open AI for cross-sectional radial wall-deformation tracking....23
Viola-Jones cascade detector....23
Mask R-CNN applied to cross-sectional views....24
Custom 24-direction radial sampling algorithm for mean-radius computation....25
Methodological integration of the two arms....25
Performance evaluation....26
Results....28
Results of methodological arm 1 — Longitudinal cIMT measurement....28
Results of methodological arm 2 — Cross-sectional radial wall-deformation tracking....31
Integrated outputs: from segmentation to clinically interpretable wall-deformation....32
Discussion....33
Algorithmic generalizability and the case for population-specific training....33
A staged research design: from population-level validation to subgroup analysis....34
Open methods, proprietary software, and structural barriers to access....36
Indigenous data sovereignty, community engagement, and authorship....37
Integrated pipeline as a model for equitable cardiovascular screening....37
Limitations and the boundary between proof-of-concept and clinical deployment....38
Implications for health systems serving Indigenous Peoples....39
Conclusion....40
Declarations....41
Ethics approval and consent to participate....41
Consent for publication....42
Availability of data and materials....42
Competing interests....43
Funding....43
Authors' contributions....43
Acknowledgements....44
AI Declaration....44
Figures and legends....45
Figure 1. Mask R-CNN segmentation of the carotid intima-media region and motivation for post-processing refinement (longitudinal view.)....45

Figure 2. Comparative evaluation of post-processing techniques for refinement of the Mask R-CNN longitudinal segmentation output.....46
Figure 3. Automated cIMT estimation pipeline and time-series output across the cardiac cycle.....47
Figure 4. Viola-Jones cascade detector applied to cross-sectional carotid ultrasound segmentation and tracking. 47
Figure 5. Mask R-CNN cross-sectional segmentation, 24-direction radial sampling algorithm, and time-series wall deformation output. 48
References 49

Equitable carotid ultrasound screening for Indigenous Peoples through open artificial intelligence: a community-engaged proof-of-concept from the PAI Project (Brazil)

Américo Mota[1,2,3,*]; Rodrigo Pereira Ramos[4]; Bruna Felix de Almeida[2,3]; Isabella Diniz de Souza Lins Bunge[2]; Jandir Mendonça Nicacio[1,3,4]; Vanessa Cardoso Pereira[4]; Thiago Augusto Cavalcante de Carvalho[2]; Paula Andreatta Maduro[3,4]; Paulo Adriano Schwingel[5]; Carlos Dornels Freire de Souza[4]; Orlando Vieira Gomes[1,3,4]; Ana Paula Pereira Rolim Coimbra Pinto[2]; Caroline Guimarães da Fonseca Chieco[2]; Maria Luíza Carvalho Santana[2]; Alane Mota dos Santos[4]; Anderson da Costa Armstrong[1,2,3,4]

[1] Graduate Program in Human Ecology and Socio-Environmental Management (PPGEcoH), State University of Bahia (UNEB), Juazeiro, Bahia, Brazil

[2] AMMO Institute / CDI, Petrolina, Pernambuco, Brazil

[3] University Hospital of the Federal University of the São Francisco Valley (HU-UNIVASF / EBSERH), Petrolina, Pernambuco, Brazil

[4] Federal University of the São Francisco Valley (UNIVASF), Petrolina, Pernambuco, Brazil

[5] University of Pernambuco (UPE), Petrolina Campus, Petrolina, Pernambuco, Brazil

* **Corresponding author:** Américo Mota, MD, MSc. ORCID: 0000-0003-0477-8330. Email: americomota@gmail.com

ORCID identifiers of the authors:

Américo A. M. Junior: 0000-0003-0477-8330

Rodrigo P. Ramos: 0000-0002-9243-3887

Bruna F. Almeida: 0000-0002-5259-4088

Isabella D. S. L. Bunge: 0009-0008-1175-6942

Jandir M. Nicacio: 0000-0001-9139-2414

Vanessa C. Pereira: 0000-0002-1030-9277

Thiago A. C. Carvalho: 0000-0002-2307-9300

Paula A. Maduro: 0000-0003-2174-2460

Paulo A. Schwingel: 0000-0002-2935-3403

Carlos D. F. Souza: 0000-0001-7995-1893

Orlando V. Gomes: 0000-0001-6324-7594

Ana P. P. R. C. Pinto: 0000-0002-1770-5619

Caroline G. F. Chieco: 0000-0002-7042-8438

Maria Luíza C. Santana: 0009-0007-8083-1910

Alane Mota dos Santos: 0000-0002-4195-3217

Anderson C. Armstrong: 0000-0003-3161-8922

Abstract

Background. Indigenous Peoples worldwide experience persistent inequities in cardiovascular health and limited access to noninvasive subclinical risk stratification. Artificial intelligence (AI) could democratize such access, but AI tools for carotid ultrasound have been developed almost exclusively in non-Indigenous, high-income populations, raising concerns about their external validity and the perpetuation of disparities. The Project of Atherosclerosis among Indigenous Populations (PAI Project) is a community-engaged initiative conducted with the Truká and Fulni-ô peoples of Northeast Brazil.

Methods. We report an integrated proof-of-concept combining two open AI pipelines: (i) automated longitudinal carotid intima-media thickness (cIMT) measurement using Mask R-CNN deep learning complemented by eight classical post-processing techniques, and (ii) automated cross-sectional radial wall deformation tracking using Viola-Jones cascade detection with Mask R-CNN segmentation and a custom 24-direction radial sampling algorithm. This Phase 2 study replaced the manual speckle tracking used in Phase 1 with end-to-end open automation. Images were obtained from Phase 1 (2016–2017, 999 adults, including Truká, Fulni-ô, and urban reference) and Phase 2 (2021–2023, 2,566 Truká and Fulni-ô adults), acquired using portable ultrasound suitable for field conditions. All methods used open-source frameworks, deliberately avoiding the use of proprietary software. The performance was evaluated using the Sørensen-Dice similarity coefficient and mean squared error (MSE).

Results. For longitudinal cIMT segmentation, Mask R-CNN achieved a Dice of 0.8954, accuracy of 95.24% (at Dice ≥ 0.80), and MSE of 17.91%. Among the eight post-processing techniques, Canny edge detection provided the most consistent boundary delineation, preserving the cyclic contraction–dilation pattern across the cardiac cycle. For cross-sectional radial wall deformation tracking, Mask R-CNN achieved a Dice of 0.8954 (95.24% accuracy), outperforming Viola-Jones (Dice 0.8938; 93.85%). The 24-direction radial sampling produced biologically plausible time series of the mean carotid radius across the cardiac cycle.

Conclusions. An open, non-proprietary AI pipeline combining longitudinal cIMT measurements and cross-sectional radial wall

deformation tracking is technically feasible in Indigenous communities in Northeast Brazil. This proof-of-concept study replaces operator-dependent manual analysis with locally validated automation, establishing a methodological and ethical foundation for the equitable deployment of cardiovascular AI in underserved populations.

Keywords: Health equity, Indigenous health, Artificial intelligence, Cardiovascular screening, Carotid ultrasound, Tele-ultrasound, Brazil, Truká, Fulni-ô, Underserved populations.

Background

Cardiovascular inequities in Indigenous Peoples

Cardiovascular diseases (CVDs) are the leading cause of mortality worldwide and contribute disproportionately to the burden of disease among Indigenous Peoples globally [1,2]. Indigenous populations are heterogeneous; however, available evidence consistently documents a higher prevalence of hypertension, diabetes, obesity, and premature cardiovascular mortality than in non-Indigenous populations in the same countries [3,4]. The Project of Atherosclerosis among Indigenous Populations (PAI), conducted with the Truká and Fulni-ô peoples in Northeast Brazil, has previously documented that greater urban influence is associated with a higher obesity prevalence, more hypertension, and rising cardiovascular mortality in these communities [5,6]. Complementary investigations within the PAI Project have further documented elevated cardiometabolic risk, altered cardiac biomarkers, chronic kidney disease burden, and adverse social determinants of health across these communities, with women and older adults being disproportionately affected [7–11].

Subclinical carotid artery alterations, measurable through ultrasound as increased intima-media thickness (cIMT) or reduced wall deformation, represent an early window for cardiovascular risk stratification, ideally before clinical events such as stroke or myocardial infarction [12,13]. However, the routine assessment of these markers depends on infrastructure that is structurally inaccessible in most Indigenous health districts, such as high-end ultrasound scanners, proprietary cardiovascular imaging software licensed at substantial costs per scanner-year, and specialized sonographers and cardiologists trained in advanced post-processing techniques such as speckle tracking [14]. The resulting reality is one in which diagnostic technologies that could identify subclinical cardiovascular disease and prevent stroke or premature death are concentrated in tertiary urban centers and are largely unavailable to Indigenous Peoples.

Artificial intelligence: opportunity and risk for health equity

Artificial intelligence (AI), particularly deep learning approaches such as convolutional neural networks (CNNs) and region-based methods such as Mask R-CNN, have been increasingly applied to medical imaging segmentation tasks, including carotid ultrasound, with reported performance comparable to that of trained sonographers [15,16,17]. From an equity perspective, AI carries a dual potential: it can democratize access to high-quality diagnostic interpretation, reducing dependence on locally scarce specialists; however, it can equally reinforce existing inequities when AI tools deployed in underserved populations have been trained on data from non-representative cohorts [18–21].

A growing body of literature documents that medical AI systems trained on data from high-income, predominantly White or East Asian populations may perform systematically worse when applied

to other groups, perpetuating or amplifying healthcare disparities [18–21]. In the specific case of carotid ultrasound, published AI models have been trained almost exclusively on cohorts from China, Japan, the United States, and Western Europe [15,16,17,22]. To our knowledge, no published AI tool for carotid wall segmentation has been developed or validated in any Indigenous population worldwide, despite Indigenous Peoples bearing a substantial and growing burden of cardiovascular disease.

This gap has practical implications for the equity of AI-assisted cardiovascular screening that might be deployed in Indigenous territories. Without prior validation in the target population, externally developed AI tools could systematically misclassify Indigenous patients, with consequences ranging from missed diagnoses to inappropriate referrals to specialists. Therefore, validation in the population of interest is not an optional refinement but a precondition for ethical deployment.

The PAI Project as a platform for equitable AI development

The PAI program provides a uniquely positioned platform for developing and validating AI tools designed with and for Indigenous populations. Across Phase 1 (2016-2017) and Phase 2 (2021-2023, conducted under the broader scope of the PAI Project, locally referred to as PAI Project), the project has accumulated comprehensive cardiovascular imaging data from the Truká (more urbanized; Assunção Island, Cabrobó, Pernambuco) and Fulni-ô (less urbanized; Águas Belas, Pernambuco) communities, in addition to a complementary urban reference cohort from Petrolina in the São Francisco Valley region [5–11]. The project is conducted under the supervision of cardiology, vascular ultrasound, and radiology specialists with extensive regional experience, in partnership with Indigenous leadership, and includes a co-author of Truká ancestry,

reflecting a deliberate commitment to community-engaged research [7–11].

Within the broader PAI Project research program, our group has previously developed manual speckle tracking-based circumferential and radial strain analyses using proprietary cardiac software adapted for vascular use, applied to a regional urban reference cohort [14]. That work, corresponding to Phase 1 of this research program, demonstrated the clinical validity of carotid wall deformation analysis as a marker of subclinical cardiovascular disease, but depended on manual placement of six tracking points along the medial-adventitial layer of the vessel, an operator-dependent step requiring specialist training and proprietary software licensing. These constraints render manual speckle tracking fundamentally incompatible with the workflow of remote or itinerant Indigenous health teams. The present study reports Phase 2 of this research program: full replacement of the manual workflow by an end-to-end automated open AI pipeline that operates directly on raw DICOM input, requires no proprietary software, and is deployable on standard hardware available in resource-limited public health settings.

Objective and scope

In the present study (Phase 2 of the research program), we report an integrated proof-of-concept combining two complementary open AI pipelines developed within the PAI Project: (i) automated measurement of carotid intima-media thickness (cIMT) from longitudinal ultrasound images, based on Mask R-CNN deep learning and a panel of eight classical computer vision post-processing techniques; and (ii) automated tracking of radial wall-deformation from cross-sectional ultrasound images, based on Viola-Jones cascade detection combined with Mask R-CNN segmentation and a custom 24-direction radial sampling algorithm for mean radius

computation. We frame both pipelines as the protagonist components of a single research question — whether an open AI pipeline can support equitable, fully automated cardiovascular screening for Indigenous Peoples — rather than as separate technical contributions. The two pipelines address complementary anatomical questions: the longitudinal pipeline quantifies the thickness of the intima-media layer (a static structural marker of atherosclerosis), whereas the cross-sectional pipeline quantifies how the carotid wall deforms across the cardiac cycle (a dynamic functional marker of arterial stiffness).

Stratified analyses comparing AI-derived imaging features between the Truká, Fulni-ô, and urban reference groups and correlating these features with cardiometabolic and laboratory variables from the broader PAI cohort are part of a planned subsequent study and are not within the scope of the present manuscript. This staged research design reflects our view that population-level technical validation must precede subgroup-level epidemiological inferences.

Methods

Study design and ethical framework

This is a methodological proof-of-concept study based on a retrospective analysis of carotid ultrasound images obtained from the Project of Atherosclerosis among Indigenous Populations (PAI Project), Phase 1 of fieldwork (2016-2017) and Phase 2 of fieldwork (2021-2023). The present manuscript reports Phase 2 of the research program, in which the manual speckle-tracking workflow previously developed in Phase 1 [14] was replaced by an end-to-end automated Open AI pipeline. The PAI Project is a community-engaged research initiative conducted in partnership with the Truká and Fulni-ô Indigenous Peoples of Northeast Brazil. Research

priorities and protocols were developed in dialogue with community representatives, Indigenous health agents, and local leaders, following the broader principles of community-engaged Indigenous health research, as described elsewhere [7–11]. The reporting of this proof-of-concept study follows the EQUATOR Network guidance: the methodological description and performance reporting align with the Checklist for Artificial Intelligence in Medical Imaging (CLAIM), and the equity-focused framing of design, recruitment, and interpretation aligns with the STROBE-Equity reporting guideline. All research procedures complied with the principles of the Declaration of Helsinki (revised 2013) and applicable Brazilian regulations governing research with Indigenous Peoples, including Resolutions CNS 466/2012 and CNS 304/2000 and Operational Norm 001/2013. This work contributes to the United Nations Sustainable Development Goals (SDGs), in particular SDG 3 (Good Health and Well-being), SDG 10 (Reduced Inequalities), and SDG 17 (Partnerships for the Goals), by extending advanced cardiovascular imaging methods to historically underserved Indigenous communities through a community-engaged research framework.

The PAI Project was approved by the Brazilian National Research Ethics Commission (CONEP) under approval number 5.927.482, issued on March 7, 2023, with Certificate of Presentation for Ethical Appreciation (CAAE) number 60271422.2.0000.8807. This study was conducted in accordance with the principles of the Declaration of Helsinki (revised 2013). The study was approved by the Brazilian Resolution CNS nº 466/2012, Resolution CNS nº 510/2016, and Operational Norm nº 001/2013 of the National Health Council, which collectively govern research involving Indigenous Peoples in Brazil. Phase 1 of the PAI was previously approved by the same regulatory framework under a separate prior ethical authorization. The protocol was approved by the Indigenous leadership of the Truká and Fulni-ô communities. All participants or their legal

representatives provided written informed consent through a process culturally adapted to the local Indigenous contexts, with translation when needed.

In recognition of Indigenous data sovereignty, the present study follows the spirit of the CARE Principles for Indigenous Data Governance (Collective benefit, Authority to control, Responsibility, Ethics) [23]: the research is conducted in partnership with the communities, results are returned to community members and local health services, and the dataset is not transferred outside the research consortium without renewed community consent (see the Data Sharing Statement).

Scope of the present analysis and rationale for non-stratified design

The present analysis was conceived as a population-level technical validation: we evaluated whether the proposed Open AI pipeline performed adequately when trained and tested on carotid ultrasound images from the PAI cohort as a whole. We deliberately did not report performance metrics stratified by Indigenous community (Truká versus Fulni-ô) or sex in this manuscript. This design choice reflects two considerations: first, the sample size available for this proof-of-concept does not yet support reliable subgroup performance estimates, and underpowered fairness analyses may be more misleading than informative. Second, stratified subgroup comparisons, including the correlation between AI-derived imaging features and the cardiometabolic, laboratory, and clinical variables described in the broader PAI cohort [7–11], constitute the focus of a planned next-stage study, in which the validated pipeline will be applied across Truká, Fulni-ô, and urban reference participants with larger samples and structured fairness assessments. We consider this staged design to be more

methodologically transparent than reporting underpowered stratified metrics in the present proof-of-concept.

Equity rationale for methodological choices

Each methodological choice in this study was guided by an explicit equity rationale and technical considerations.

First, we deliberately avoided proprietary cardiovascular imaging software at all stages of the pipeline. Commercial speckle-tracking and cIMT measurement packages are typically licensed per scanner at costs that effectively exclude Indigenous health districts and other low-resource public health settings. Our pipeline was implemented using open-source frameworks (TensorFlow for deep learning, OpenCV-style classical computer vision routines) and standard algorithms described in the scientific literature (Mask R-CNN, Viola-Jones, Canny, Sobel, Laplacian, Otsu, Riddler-Calvard, adaptive thresholding). These methods are standard in the medical imaging research community; the equity contribution lies in assembling them into a deployable pipeline that does not require commercial licenses.

Second, the AI models were trained on data from the target population rather than being transferred uncritically from non-representative cohorts. By training directly on PAI images, we generated models whose validity was established in the population of interest, while still leveraging transfer learning from a generic visual foundation (COCO) to compensate for the limited training sample size.

Third, manual annotations that served as ground truth were performed by sonographers and trained reviewers familiar with the regional acquisition protocol and patient population, mitigating the risk that labeling biases imported from external annotators would

systematically distort the segmentation boundaries in this specific population.

Fourth, this proof-of-concept is deliberately framed within a longer research roadmap that includes community-engaged validation, results returned to participating communities, and progressive expansion, guided by Indigenous data sovereignty principles.

Image acquisition across the PAI program

Carotid ultrasound images were acquired across the two fieldwork phases of the PAI Project using two portable ultrasound platforms compatible with field conditions in Indigenous territories: a Philips CX 50 portable scanner (Philips Healthcare) and a Mindray MX7 portable system (Mindray Medical International). The use of portable equipment is consistent with the realistic operational profile of public health systems serving Indigenous and underserved populations, in which fieldwork in remote areas requires equipment that can be transported, set up, and operated outside conventional radiology. For each participant, ECG-synchronized video clips of the common carotid artery were obtained bilaterally in both the longitudinal and short-axis (cross-sectional) views. The images were exported in DICOM format and subsequently extracted to standard image formats (JPEG/PNG) for AI processing.

Acquisition supervision followed two distinct quality assurance protocols across the fieldwork phases. In Phase 1 of the fieldwork (2016-2017), image acquisition was performed by trained sonographers under direct on-site supervision by a multidisciplinary team of specialists, including radiologists, ultrasonographers, vascular ultrasound specialists, and cardiologists with subspecialty training in echocardiography. This supervisory model reflects the standard practice of academic medical imaging research and ensures the highest possible image quality at the cost

of substantial logistical and human resource investments per fieldwork campaign.

In Phase 2 of the fieldwork (2021-2023), the operational model evolved toward a tele-ultrasound workflow designed to be sustainable in Indigenous health districts beyond the duration of dedicated research campaigns. Image acquisition was performed by medical students, medical residents, and general practitioners deployed in the field, with asynchronous remote quality assurance performed by experienced radiologists who accessed the acquired images through a Picture Archiving and Communication System (PACS). This decentralized acquisition model is operationally aligned with the actual workforce composition of the Brazilian Indigenous Health Subsystem (SESAI) and demonstrates that high-quality carotid ultrasound images can be collected by non-specialist practitioners under asynchronous expert oversight, which is a critical precondition for scalable cardiovascular screening in remote Indigenous territories.

The image set available for AI development spans both fieldwork phases of the PAI Project (Phase 1 of fieldwork: 999 adults aged 30–70 years from Truká, Fulni-ô, and an urban reference group; Phase 2 of fieldwork: 2,566 Truká and Fulni-ô adults aged ≥18 years) [5–11]. The training, validation, and test partitions used for the AI models in the present study were defined retrospectively from this broader image archive, prioritizing annotated frames of adequate technical quality.

Methodological arm 1 — Open AI for automated longitudinal cIMT measurement

The first methodological arm targets the automated measurement of carotid intima-media thickness (cIMT) using longitudinal B-mode ultrasound images. Accurate cIMT quantification is a recognized early marker of subclinical atherosclerosis [12] and has historically depended on manual or semi-automated software developed by ultrasound equipment manufacturers, requiring per-scanner licensing that effectively excludes Indigenous health districts. Our pipeline replaces this proprietary workflow with a fully open architecture based on Mask R-CNN deep learning for segmentation of the intima-media region, followed by automated post-processing through eight classical computer vision techniques whose mathematical foundations are openly described in the scientific literature [24–28].

Mask R-CNN architecture and training

Mask R-CNN is a region-based CNN architecture that extends Faster R-CNN by adding a branch for predicting binary segmentation masks on each Region of Interest in parallel with the classification and bounding box regression branches [24]. The architecture comprises three main components: (i) a CNN backbone (ResNet-based) that generates a feature map from the input image; (ii) a Region Proposal Network that scans the feature map and proposes candidate object regions; and (iii) three parallel heads applied to each proposed region after RoIAlign-based feature pooling. We chose Mask R-CNN over simpler alternatives because its architecture is fully described in the public literature, its reference implementations are open-source, and its performance on small medical imaging datasets is well documented [15,16,17,24].

The model was implemented in Python using the TensorFlow framework with transfer learning from a model pretrained on the COCO dataset and subsequent fine-tuning on our task-specific data. The longitudinal cIMT segmentation dataset comprised 114

training, 57 validation, and 42 test frames, which were manually annotated point-by-point using the VGG Image Annotator (VIA) software (Visual Geometry Group, University of Oxford). The training was conducted for 30 epochs, and the final model was selected based on the lowest validation error.

Post-processing: comparative evaluation of eight techniques

Visual inspection of the raw Mask R-CNN output revealed boundary uncertainty in a fraction of the test frames, with weakened intensity transitions, making the precise intima-media contour ambiguous. We evaluated eight classical computer-vision techniques for mask refinement. The mathematical foundation of each technique is summarized below, providing a transparent description that does not require access to the commercial documentation.

Simple thresholding applies a fixed threshold T to each pixel intensity I(x,y): pixels below T are set to zero, and the others are set to the maximum value, producing a binary image [25].

Adaptive mean thresholding computes a local threshold for each pixel as the mean intensity of a defined neighborhood window minus a constant C, which is useful when the image presents non-uniform illumination, which is a common situation in ultrasound.

Adaptive Gaussian thresholding is similar to adaptive mean thresholding; however, the local threshold is computed as a Gaussian-weighted sum of neighborhood intensities, with weights decreasing with distance from the central pixel, minus a constant C.

Otsu's method automatically determines an optimal global threshold by minimizing the intra-class variance between two pixel classes (object and background) and approximating the image histogram as a sum of two Gaussian distributions [26].

The Riddler-Calvard algorithm estimates the threshold through an iterative isodata procedure, with the threshold updated as the average of the means of the two classes defined by the current threshold until convergence. We applied Otsu and Riddler-Calvard sequentially as a combined refinement strategy.

The Sobel operator approximates the spatial gradient of the image intensity through two 3×3 convolution kernels (horizontal Gx and vertical Gy); the gradient magnitude $G = \sqrt{(Gx^2 + Gy^2)}$ is high at the edges.

The Laplacian filter computes the second spatial derivative of the image using a single isotropic 3×3 convolution kernel, emphasizing regions of rapid intensity change and responding to edges in all directions simultaneously.

The Canny edge detector is a multi-stage algorithm comprising (a) Gaussian smoothing for noise reduction, (b) gradient computation using Sobel kernels, (c) non-maximum suppression to thin edges, (d) double thresholding, and (e) edge tracking by hysteresis, which preserves weak edges only when connected to strong edges [27]. Canny is widely considered the reference standard for edge detection in noisy images, such as ultrasound images.

Automated cIMT estimation

Following post-processing, the refined intima-media contour was used to automatically estimate cIMT through a minimal bounding rectangle fitted to the segmented region, with the rectangle height (in pixels) interpreted as a proxy for the wall thickness. Pixel-to-millimeter conversion used the spatial calibration in the DICOM header. The procedure was repeated frame-by-frame, generating a

normalized time series of cIMT values showing the contraction-dilation cycle of the artery during the cardiac cycle.

Cross-sectional segmentation: Viola-Jones cascade detector

The Viola-Jones object detection framework combines integral image representation, Haar-like rectangular features, the AdaBoost learning algorithm, and a cascade structure of weak classifiers (C1, C2, …, Cn), in which sub-windows rejected at any stage are discarded and only those passing all stages are declared positive detections [28]. We selected the Viola-Jones algorithm for this proof-of-concept because of its computational efficiency, which makes it deployable on hardware available in resource-limited public health settings (including standard laptops without dedicated GPUs).

For cross-sectional carotid detection, the Viola-Jones detector was trained with 149 positive images (cropped carotid regions) and 596 negative images (background), with an additional 65 frames reserved for testing. Rectangular bounding box annotations were performed using VIA software. A circular mask inscribed within the detected bounding box was used for the initial cross-sectional segmentation.

Mask R-CNN applied to cross-sectional views

Two limitations of the Viola-Jones approach motivated the training of a second Mask R-CNN model for cross-sectional segmentation: (i) the circular mask did not perfectly delineate the carotid wall, particularly in vessels with irregular geometry, and (ii) small frame-to-frame variations in the bounding box generated artifactual fluctuations in the estimated radius. The cross-sectional Mask R-CNN was trained using 114 training, 57 validation, and 42 test

frames with point-by-point manual annotations of the carotid wall, following the protocol described for longitudinal segmentation.

Mean radius computation

Because the carotid cross-section is not strictly circular, the mean radius was computed as follows: the centroid of the segmented region was determined as the mean of the x and y pixel coordinates belonging to the mask; for every 15° around the centroid (24 radial directions), the Euclidean distance from the centroid to the corresponding wall pixel was computed; and the mean of these 24 distances yielded the mean radius for that frame. Repeating this computation across all frames produces a time series of mean radius values.

Methodological arm 2 — Open AI for cross-sectional radial wall-deformation tracking

The second methodological arm targets the automated tracking of carotid radial wall deformation from short-axis (cross-sectional) ultrasound video clips. While the longitudinal cIMT arm quantifies a static structural marker of atherosclerosis (wall thickness), this arm quantifies a dynamic functional marker: how the carotid wall expands during systole and recoils during diastole, providing information about arterial stiffness and elastic function, properties known to be altered by hypertension and other cardiovascular risk factors [13,14]. Phase 1 of this research program addressed the same functional question through manual speckle tracking with proprietary software, requiring the operator to place six tracking points along the medial-adventitial layer of the vessel [14]. Phase 2, described here, replaces the operator-dependent workflow with end-to-end open AI.

Viola-Jones cascade detector

The Viola-Jones object detection framework combines integral image representation, Haar-like rectangular features, the AdaBoost learning algorithm, and a cascade structure of weak classifiers (C1, C2, …, Cn), in which sub-windows rejected at any stage are discarded and only those passing all stages are declared positive detections [28]. We selected Viola-Jones for the initial cross-sectional detection stage of this pipeline because of its computational efficiency, which makes it deployable on hardware available in resource-limited public health settings, including standard laptops without dedicated GPUs.

For cross-sectional carotid detection, the Viola-Jones detector was trained with 149 positive images (cropped carotid regions) and 596 negative images (background), with an additional 65 frames reserved for testing. The training set was expanded relative to a preliminary version of this pipeline, with 39 additional training frames and 10 additional test frames incorporated to improve the detection robustness. Rectangular bounding box annotations were performed using VIA software (Figure 4C). A circular mask inscribed within the detected bounding box was used as the initial cross-sectional segmentation (Figure 4A, schematic; Figure 4B, applied to the carotid artery).

Mask R-CNN applied to cross-sectional views

Two limitations of the Viola-Jones-based segmentation motivated the training of an additional Mask R-CNN model dedicated to cross-sectional views: (i) the inscribed circular mask does not perfectly delineate the carotid wall, particularly in vessels with irregular geometry, and (ii) small frame-to-frame variations in the bounding box generate artifactual fluctuations in the estimated radius that are not biologically meaningful. The cross-sectional Mask R-CNN was

trained using 114 training, 57 validation, and 42 test frames, with point-by-point manual annotations of the carotid wall along the medial-adventitial interface (Figure 5A), performed in the VIA software, following the same protocol described for the longitudinal model. Training was conducted using transfer learning from a COCO-pretrained model over 30 epochs in the TensorFlow framework.

Custom 24-direction radial sampling algorithm for mean-radius computation

Because the carotid cross-section is not strictly circular, computing a single representative radius from the segmented mask requires a sampling strategy that captures the wall geometry in multiple directions. We implemented a custom 24-direction radial sampling algorithm (Figure 5B) with the following steps: (i) the centroid of the segmented region was determined as the mean of the x and y pixel coordinates belonging to the mask; (ii) for every 15° around the centroid (24 radial directions in total, spanning the full 360°), the Euclidean distance from the centroid to the corresponding wall pixel was computed; (iii) the mean of these 24 distances yielded the mean radius for that frame. Repeating this computation across all frames of an ECG-synchronized video clip produces a time series of mean carotid radius values, from which radial wall deformation across the cardiac cycle is derived. The 24-direction sampling provides a balance between angular resolution and computational efficiency, and it generalizes naturally to vessels with non-circular geometry, including the common carotid artery in older or hypertensive participants, where elliptical deformation is more frequent.

Methodological integration of the two arms

The two arms of this proof-of-concept share a single technical foundation — Mask R-CNN deep learning trained on annotated carotid ultrasound images from the PAI Project — applied to two complementary anatomical views. The longitudinal arm targets the intima-media layer for cIMT measurement, and the cross-sectional arm targets the full vessel wall for radial deformation tracking. Together, they cover the two principal carotid imaging outputs used in cardiovascular risk stratification: a static structural marker (cIMT) and a dynamic functional marker (radial wall deformation), both obtained without any operator-dependent manual delineation and without any proprietary software dependency.

This integration represents the central methodological advance of the present study relative to Phase 1 of this research program [14], in which the same dynamic functional question was addressed through manual speckle tracking with proprietary software, requiring six operator-placed tracking points per participant. The fully automated Open AI pipeline described here is operationally compatible with the asynchronous tele-ultrasound workflow piloted in Phase 2 of the fieldwork, in which acquisition by medical students, residents, and general practitioners can be combined with remote AI-assisted analysis, which is a precondition for scalable cardiovascular screening in Indigenous territories.

Performance evaluation

The performance was evaluated for both methodological arms using the Sørensen-Dice similarity coefficient, which measures the agreement between algorithmic and manual segmentations:

$$Dice = 2 \times |A \cap M| / (|A| + |M|)$$

where A is the algorithmically segmented region and M is the manually annotated reference region,

with values ranging from 0 (no overlap) to 1 (perfect agreement). Segmentations with Dice ≥ 0.80 were considered correct for computing the accuracy rate, following the thresholds commonly adopted in the medical imaging literature [15]. The mean squared error (MSE) of the Dice coefficients across the test set was computed as follows:

$$MSE = (1/n) \times \Sigma\, |1 - Dice_i|^2$$

For the longitudinal cIMT arm (Arm 1), the performance was assessed using a held-out test set of 42 frames. For the cross-sectional radial wall-deformation arm (Arm 2), the performance was assessed for both the Viola-Jones detector (65 test frames) and cross-sectional Mask R-CNN model (42 test frames), enabling a direct comparison between the two architectures. Given the proof-of-concept nature of this study, the primary outputs were descriptive: mean Dice, accuracy at Dice ≥ 0.80, and MSE for each pipeline. Visual inspection of representative test cases and time-series plots of cIMT and mean carotid radius across the cardiac cycle complemented the quantitative metrics by demonstrating the biological plausibility of the AI outputs.

The sample sizes adopted for AI training and testing (114/57/42 frames for the longitudinal and cross-sectional Mask R-CNN models and 149/596 positive/negative images and 65 test frames for the Viola-Jones detector) were chosen to balance computational tractability with statistical robustness for proof-of-concept validation in an Indigenous research setting. This deliberate scale is

consistent with the staged research design used in this study. The present analysis prioritizes population-level technical feasibility as the precondition for the planned next-stage study, in which expanded image datasets from the broader PAI Project cohort, including stratified samples across Truká, Fulni-ô, and the urban reference group, will support multi-observer validation and structured fairness assessments.

Results

Results of methodological arm 1 — Longitudinal cIMT measurement

The Mask R-CNN model trained on annotated longitudinal frames successfully localized the intima-media region in the test frames that were not included in the training or validation sets. The model produced binary masks delineating the intima-media layer on the raw DICOM input (Figure 1A), with biologically plausible boundaries upon visual inspection in the majority of cases. In a fraction of the test frames, the raw model output showed weakened intensity transitions at the wall boundary, producing uncertainty in the precise contour position (Figure 1B). The eight post-processing techniques described in the Methods section were then applied to refine the masks (Figure 2; see Table 1 for qualitative ranking).

Among the eight post-processing techniques evaluated, the Canny edge detector produced the cleanest closed-contour delineation of the intima-media boundary (Figure 2E), followed by the combined Otsu + Riddler-Calvard thresholding (Figure 2B). Simple and adaptive thresholding methods (Figure 2A) tended to over-segment, including residual signals outside the intima-media region. The Sobel operator (Figure 2C) yielded sharper but discontinuous edges, whereas the Laplacian filter (Figure 2D) highlighted the region of

interest with notable noise. When these techniques were applied to refine the Mask R-CNN output prior to bounding rectangle-based cIMT estimation (Figure 3A), the time-series curves of the normalized wall thickness exhibited the expected contraction-dilation pattern across the cardiac cycle (Figure 3B). The Canny-derived curves preserved greater amplitude and clearer cyclic structure than the Otsu+Riddler-Calvard-derived curves, suggesting that more precise edge delineation preserves fine-grained temporal variation that would otherwise be attenuated by less specific binarization methods.

Table 1. Comparative qualitative ranking of post-processing techniques applied to Mask R-CNN longitudinal segmentation.

Post-processing technique	Qualitative performance on intima-media boundary delineation
Simple thresholding	Sensitive to non-uniform illumination; frequently over- or under-segments depending on threshold choice
Adaptive mean thresholding	Better handling of non-uniform illumination; tendency to over-segment with residual signal outside the intima-media
Adaptive Gaussian thresholding	Smoother local adaptation than mean variant; similar over-segmentation tendency
Otsu's method	Avoids manual threshold selection; performance depends on bimodal histogram assumption
Otsu + Riddler-Calvard	Improved refinement over Otsu alone; second-best overall

	performance in our evaluation
Sobel operator	Sharper edges but with discontinuities at wall boundary
Laplacian filter	Sensitive to edges in all directions but with substantial noise
Canny edge detector	Best overall: clean closed-contour delineation; preserves cyclic contraction-dilation pattern

Results of methodological arm 2 — Cross-sectional radial wall-deformation tracking

For cross-sectional carotid detection and segmentation, the Mask R-CNN model achieved a mean Dice coefficient of 0.8954, an accuracy of 95.24% (considering Dice ≥ 0.80), and an MSE of 17.91% on the held-out test set of 42 frames. The Viola-Jones cascade detector, trained on 149 positive and 596 negative images, achieved a mean Dice of 0.8938, an accuracy of 93.85%, and an MSE of 22.12% on the 65-frame test set (Table 2). Therefore, both methods produced segmentations within the range typically considered acceptable in the medical imaging literature [15,16,17,22], demonstrating that an open AI pipeline trained on images from the PAI Project — including Indigenous participants — attained performance comparable to that reported in much larger non-Indigenous cohorts published in high-income settings. Mask R-CNN outperformed Viola-

Jones across all three metrics, supporting the addition of the deep learning model as a refinement over the cascade detector alone.

Table 2. Performance metrics for cross-sectional carotid segmentation in the PAI Project cohort.

Method	Accuracy (Dice ≥ 0.80)	Mean Dice	MSE
Viola-Jones	93.85%	0.8938	22.12%
Mask R-CNN	95.24%	0.8954	17.91%

Using the 24-direction radial sampling algorithm described in the Methods section, time-series curves of the mean carotid radius were obtained from cross-sectional video clips for both Viola-Jones and Mask R-CNN segmentations. In the majority of the analyzed clips, both algorithms produced biologically plausible curves displaying systolic dilation and diastolic recoil consistent with the expected physiological pattern of arterial wall deformation (Figure 4D and 5C). In a minority of clips, tracking failures were observed, characterized by non-physiological spikes in the estimated radius, attributable to localized segmentation errors in individual frames. These failures were more frequent and severe with Viola-Jones (Figure 4E, panel d showing catastrophic collapse) than with Mask R-CNN (Figure 5D, panel d showing only an isolated spike), supporting the superiority of the deep learning approach for downstream wall deformation analysis, where the temporal continuity of the radius signal is essential.

Integrated outputs: from segmentation to clinically interpretable wall-deformation

Taken together, the two arms demonstrate the technical feasibility of an integrated, open, and locally deployable pipeline for cardiovascular screening in under-served regions. Arm 1 provides automated cIMT measurements with a performance comparable to that reported in the international literature. Arm 2 provides automated radial wall deformation tracking that captures the dynamic functional properties of the vessel, which were previously accessible only through manual speckle tracking with proprietary software [14]. The pipeline operates end-to-end on raw DICOM input, requires no operator-dependent manual delineation, and runs entirely on open-source software, removing the principal structural barriers to the deployment of advanced carotid imaging in Indigenous health districts and other low-resource public health settings.

Discussion

This integrated proof-of-concept study demonstrates the technical and clinical feasibility of an open, non-proprietary pipeline for cardiovascular screening in Indigenous communities in Northeast Brazil. The pipeline combines two methodological arms developed within the PAI program: deep learning-based automated segmentation of the carotid wall (Mask R-CNN, Viola-Jones, and a panel of classical computer vision techniques) and speckle tracking-based quantitative strain analysis via asynchronous tele-ultrasound. Beyond the technical results, this study contributes to ongoing debates on the equitable design and deployment of medical AI in underserved populations. We discuss the five implications below.

Algorithmic generalizability and the case for population-specific training

The principal equity-relevant contribution of this study is the demonstration that open AI methods, trained directly on data from Indigenous communities, attain segmentation performance comparable to that reported in much larger non-Indigenous cohorts published in high-income settings [15,16,17,22]. This finding addresses a recurring concern in the algorithmic fairness literature: medical AI models developed predominantly on non-representative cohorts may underperform when applied to populations excluded from the training data, perpetuating or amplifying healthcare disparities [18–21]. Our results suggest that, at least for carotid ultrasound segmentation, the technical adaptation of open methods to Indigenous populations is feasible without requiring datasets at the scale of high-income country cohorts, provided that transfer learning from generic visual foundations is used to compensate for limited sample sizes.

Our pipeline provides a methodological pathway for population-specific validation without dependence on commercial cardiovascular software packages, whose internal algorithms are opaque to external audits. The burden of proof should be reversed: rather than assuming that a model trained in one population generalizes to another, deployment in an Indigenous community should be preceded by a demonstration of validity in that community.

A staged research design: from population-level validation to subgroup analysis

We deliberately framed the present analysis as a population-level technical validation without stratified subgroup metrics. This design choice deserves explicit discussion because it represents a substantive methodological position rather than a limitation imposed by sample size alone.

In our view, premature subgroup stratification in proof-of-concept studies, particularly when sample sizes are insufficient, can produce confidence intervals so wide that any observed difference is uninterpretable and any observed similarity is unfalsifiable. Such underpowered analyses are more misleading than informative, particularly when they involve sensitive demographic categories. Therefore, in this manuscript, we prioritized a robust population-level validation of both AI pipelines as the precondition for the planned next-stage study, in which the validated pipelines will be applied across Truká, Fulni-ô, and urban reference participants with larger samples and structured fairness assessments. The next-stage study will specifically address (i) performance disparities between Indigenous communities and the urban reference for both cIMT measurement and radial wall-deformation tracking; (ii) sex-specific performance, particularly relevant given the cardiometabolic risk profile documented across the PAI Project [7–11]; and (iii) correlation of AI-derived imaging features with clinical and laboratory cardiometabolic variables from the broader PAI cohort.

We consider this staged approach to be more transparent than reporting underpowered stratified metrics in the present proof-of-concept, and we note that this methodological position is consistent with emerging guidance on responsible AI development in healthcare [18]. The next stage of this study will also extend the integrated pipeline to additional validation tasks, including the agreement between AI-derived cIMT and reference manual measurements by experienced sonographers and the correlation between AI-derived radial wall-deformation parameters and Phase-1 manual speckle-tracking outputs, where image pairs are available.

Open methods, proprietary software, and structural barriers to access

The second equity-relevant contribution concerns the structural barriers introduced by proprietary cardiovascular imaging software. Specialized cIMT and speckle tracking packages, typically licensed at substantial costs per scanner year, effectively exclude public health systems serving Indigenous and other underserved populations. The reliance on such software in the academic literature creates a downstream barrier to equitable deployment; even when a research result obtained with proprietary software has clinical merit, its translation to underserved settings requires either licensing that the local health system cannot afford or independent re-implementation that the local technical workforce may lack capacity to perform [14,29].

By assembling a pipeline based entirely on open-source frameworks and standard algorithms openly described in the scientific literature — Mask R-CNN, Viola-Jones, Canny, Sobel, Laplacian, Otsu, Riddler-Calvard, and adaptive thresholding — we demonstrate that the technical components for AI-assisted carotid assessment are freely accessible to the research and clinical community. The question is no longer whether algorithms exist, but whether the institutional, training, and infrastructural conditions for their deployment in Indigenous health districts are in place. In our view, this shift from technical to institutional framing is where future equity-oriented investments should be directed.

Indigenous data sovereignty, community engagement, and authorship

The third implication concerns the governance of AI development in Indigenous populations. The collection, storage, and use of medical imaging data from Indigenous communities raises issues that go beyond standard biomedical research ethics, encompassing the principles of Indigenous Data Sovereignty articulated in the CARE Principles for Indigenous Data Governance [23]. In the present

study, we followed these principles by conducting research in partnership with the Truká and Fulni-ô communities, with the contribution of a co-author of Truká ancestry (Dr. Orlando Vieira Gomes), whose perspectives helped inform the interpretation of findings, returning results through accessible channels, and not transferring the dataset outside the research consortium without renewed community consent.

Integrated pipeline as a model for equitable cardiovascular screening

A central methodological contribution of this study is the demonstration that two complementary AI pipelines, longitudinal cIMT measurement and cross-sectional radial wall-deformation tracking, can be developed as an integrated framework for cardiovascular screening in resource-limited settings. Both pipelines share a single technical foundation (Mask R-CNN deep learning trained on annotated carotid ultrasound images from the PAI Project), applied to two complementary anatomical views and addressing two complementary clinical questions: a static structural marker of atherosclerosis (intima-media thickness) and a dynamic functional marker of arterial stiffness (radial wall deformation across the cardiac cycle).

This integrated proof-of-concept study builds upon and extends Phase 1 of our research program [14], in which manual speckle tracking with proprietary software was used to demonstrate the clinical validity of carotid wall deformation analysis as a marker of subclinical cardiovascular disease in a regional cohort. While Phase 1 established what to measure, Phase 2, described here, establishes how to measure it in a manner compatible with the operational reality of Indigenous health districts: without manual delineation, proprietary software licensing, or specialist on-site supervision. The transition from a manual to an automated workflow is not merely a

technical refinement; it represents the difference between cardiovascular screening that is possible in tertiary urban centers and cardiovascular screening that is scalable to remote Indigenous territories.

Limitations and the boundary between proof-of-concept and clinical deployment

We must be explicit about the limitations of this study and the boundaries they impose on its interpretation. First, the AI outputs were not directly compared with cIMT or wall deformation values measured by clinical specialists using standard commercial software, which is required to establish diagnostic agreement. Second, the implementation code and trained model weights are not currently publicly available; future stages of the project will address this in coordination with Indigenous Data Sovereignty principles. Third, images were acquired with two different portable ultrasound platforms (Philips CX 50; Mindray MX7), which is a strength in terms of operational realism but a potential source of inter-equipment variability that has not been formally quantified in the present analysis.

Taken together, these limitations mean that the present study should not be read as proposing a clinical tool that is ready for deployment in Indigenous health districts. Rather, it establishes a methodological and ethical foundation for community-engaged validation studies, in which open pipelines can be progressively expanded, tested, and refined in partnership with Indigenous communities. Complementary cardiometabolic, clinical, and laboratory data from the broader PAI Project cohort [7–11] will serve as the basis for correlating AI-derived imaging features in the next stage of the study.

Implications for health systems serving Indigenous Peoples

If the technical feasibility demonstrated here is confirmed in larger, multisite validation studies, the implications for health systems serving Indigenous Peoples are potentially substantial. Tele-ultrasound, which combines image acquisition by mid-level health workers in Indigenous territories with asynchronous remote interpretation supported by AI, offers a viable path for systematic subclinical cardiovascular screening in populations currently excluded from such services. Phase 2 of fieldwork in the PAI Project demonstrated the operational viability of this model: image acquisition by medical students, residents, and general practitioners, with asynchronous remote quality assurance by experienced radiologists through a PACS system, produced images of adequate quality for advanced AI-assisted analysis. The remaining barriers to scalable deployment are predominantly institutional: training, equipment procurement and maintenance, integration into existing Indigenous health information systems (SESAI), and financing sustainability. AI is a necessary but insufficient component of this larger puzzle.

Conclusion

This integrated proof-of-concept demonstrates that two complementary open AI pipelines — automated longitudinal cIMT measurement based on Mask R-CNN deep learning with classical post-processing and automated cross-sectional radial wall deformation tracking based on Viola-Jones cascade detection combined with Mask R-CNN segmentation and 24-direction radial sampling — can be developed and validated in carotid ultrasound images from Indigenous communities in Northeast Brazil. Both pipelines achieved performance comparable to studies conducted in

much larger non-Indigenous cohorts published in high-income settings while operating entirely on open-source software and standard hardware deployable in resource-limited public health settings. The present work represents Phase 2 of a longer research program: it replaces the manual speckle-tracking workflow developed in Phase 1 [14] with end-to-end automation, removing the operator-dependent and proprietary software bottlenecks that previously limited the deployment of carotid wall deformation analysis in remote Indigenous territories. By anchoring this technical development in a community-engaged research program guided by Indigenous data sovereignty principles and the participation of a co-author of Truká ancestry, we contribute to a methodological and ethical foundation for the equitable deployment of cardiovascular AI in underserved populations. The validated pipelines will be applied in a planned next-stage study to correlate AI-derived imaging features with the rich cardiometabolic and laboratory data from the broader PAI Project cohort with stratified analyses across the Truká, Fulni-ô, and urban reference groups. Realizing the broader potential of this line of research requires moving beyond technical proof-of-concept toward community-engaged validation, capacity building, and sustainable health-system integration.

Declarations

Ethics approval and consent to participate

The PAI Project was approved by the Brazilian National Research Ethics Commission (CONEP) under approval number 5.927.482, issued on March 7, 2023 (CAAE: 60271422.2.0000.8807). Ethical approval was obtained following Brazilian Resolutions CNS nº 466/2012 and nº 510/2016 and Operational Norm nº 001/2013 of the National Health Council. Phase 1 of the PAI received prior ethical

approval within the same regulatory framework. The protocol was approved by the Indigenous leadership of the Truká and Fulni-ô communities. All participants or their legal representatives provided written informed consent through a process culturally adapted to the local Indigenous contexts.

Consent for publication

Not applicable. No individually identifiable participant data are presented in this study.

Availability of data and materials

Individual participant data and AI model weights are not publicly available. This restriction is based on Brazilian national legislation and ethical regulations protecting the collective rights, privacy, and cultural integrity of Indigenous Peoples. Data were collected under specific agreements with the Truká and Fulni-ô leadership, prioritizing Indigenous data sovereignty and autonomy. Where journal verification of findings is required, the authors may facilitate a review by a third-party auditor appointed by the journal, provided that such access is previously approved by the participating Indigenous communities, the National Research Ethics Commission (CONEP), and the relevant regulatory bodies.

Competing interests

The authors declare no competing interests.

Funding

This work was supported by the National Council for Scientific and Technological Development (CNPq) and institutional resources from

the Federal University of the São Francisco Valley (UNIVASF), State University of Bahia (UNEB), University of Pernambuco (UPE), and University Hospital of UNIVASF/EBSERH. The AMMO Institute/CDI provided the technical infrastructure for AI pipeline development. Funding sources had no role in the study design, data collection, data analysis, data interpretation, or decision to submit this manuscript for publication.

Authors' contributions

AAMJ conceived the integrated proof-of-concept combining the longitudinal cIMT and cross-sectional radial deformation AI arms, participated in all stages of both arms, performed image analysis, and drafted the manuscript as the first and corresponding author. RPR led the technical development and validation of the Mask R-CNN and Viola-Jones AI pipelines. BFA, IDSLB, TACC, APPRCP, CGFC, MLCS, and AM performed manual image annotation and contributed to the image analysis across both arms. JMN contributed to drafting and providing clinical supervision. The VCP led community engagement with the Truká and Fulni-ô leadership. PAM contributed to the clinical and ultrasonographic imaging supervision. PAS and CDFS contributed to the statistical and methodological analyses. OVG contributed perspectives informed by her Truká ancestry and helped strengthen community linkage. ACA served as the senior author, conceived the broader PAI Project, secured institutional support, and supervised the project as the principal investigator. All authors had full access to the data, critically reviewed the manuscript, approved the final version, and had the final responsibility for the decision to submit it for publication.

Acknowledgements

We are deeply grateful to the Truká and Fulni-ô peoples for their partnership, trust, and permission to conduct this study on their territories. We thank the Indigenous field research teams from both communities, whose local knowledge, language mediation, and door-to-door outreach were indispensable to the PAI program. We honor the Truká ancestry of co-author Dr. Gomes, whose contributions helped ensure that Indigenous perspectives were incorporated into the interpretation of the findings. The PAI program relies on the voluntary efforts of students and health professionals, to whom we are deeply indebted.

AI Declaration

As native Portuguese speakers, the authors used artificial intelligence-based tools to prepare this manuscript in standard scientific English. The use of these tools was limited to translation, grammatical and linguistic correction, and structural and stylistic refinement of the English text, consistent with the conventions of scientific writing. All AI-assisted edits were reviewed, revised, and approved by the authors, who took full responsibility for the accuracy, integrity, and final content of the manuscript.

Figures and legends

Figure 1. Mask R-CNN segmentation of the carotid intima-media region and motivation for post-processing refinement (longitudinal view).

(A) Original B-mode ultrasound frame extracted from a DICOM file showing the common carotid artery in a longitudinal view (left) and the binary segmentation mask of the intima-media region (right) generated by the Mask R-CNN model after transfer learning from a

COCO-pretrained backbone and fine-tuning on annotated images from the PAI Project cohort. (B) Regions of boundary uncertainty in the raw Mask R-CNN output, highlighted by red rectangles at the lateral extremities of the segmented mask, where the gradient of the pixel intensity weakens and generates ambiguity about the precise contour of the wall. These regions of uncertainty motivated the comparative evaluation of eight classical computer-vision post-processing techniques, as shown in Figure 2.

Figure 2. Comparative evaluation of post-processing techniques for refinement of the Mask R-CNN longitudinal segmentation output.

(A) Adaptive thresholding applied to the Mask R-CNN segmentation output, with mean (left) and Gaussian (right) variants; the local threshold is computed within a defined neighborhood window, compensating for non-uniform illumination but tending to over-segment. (B) Otsu thresholding combined with the Riddler-Calvard iterative isodata algorithm produced a consolidated continuous binary mask with substantially less residual background signal than adaptive thresholding. (C) Sobel operator, emphasizing edges through the spatial gradient computed via two 3×3 convolution kernels (horizontal and vertical), with discontinuous boundary delineation. (D) Laplacian filter: computing the second spatial derivative through a single isotropic 3×3 convolution kernel with low overall contrast and noise sensitivity. (E) Canny edge detector, a multi-stage method combining Gaussian smoothing, gradient computation, non-maximum suppression, double thresholding, and edge tracking by hysteresis, produces the cleanest, thinnest, and most continuous closed-contour delineation of the intima-media region among the eight techniques evaluated.

Figure 3. Automated cIMT estimation pipeline and time-series output across the cardiac cycle.

(A) Final step of the automated longitudinal pipeline: the refined intima-media contour obtained through Canny edge detection (Figure 2E) is fitted with a minimum bounding rectangle (in red); the height of this rectangle, in pixels, is converted to millimeters using the spatial calibration encoded in the DICOM file header, providing the automated cIMT value for the analyzed frame. (B) Time-series curves of normalized cIMT across the cardiac cycle derived frame-by-frame from the same longitudinal video clip after two different post-processing strategies: Otsu thresholding combined with Riddler-Calvard (left) and Canny edge detection (right). The Canny-derived curves preserved greater amplitude and clearer cyclic contraction-dilation structure, supporting the interpretation that more precise edge delineation captures fine-grained temporal variation that would otherwise be attenuated.

Figure 4. Viola-Jones cascade detector applied to cross-sectional carotid ultrasound segmentation and tracking.

(A) Schematic diagram of the cascade of classifiers in the Viola-Jones detection framework: image sub-regions are evaluated sequentially by classifiers C1, C2, …, Cn; sub-regions failing any classifier (F) are discarded, and those passing all (V) are declared detected. (B) Cross-sectional B-mode ultrasound image of the common carotid artery overlaid with the Viola-Jones detector output: red rectangular bounding box and green circular mask inscribed within the bounding box. (C) Rectangular manual annotation of the carotid position using the VGG Image Annotator (VIA) software, generating ground truth labels for the training and validation of the Viola-Jones detector. (D) Frame-by-frame tracking of the carotid artery in a

cross-sectional video clip (2×3 grid of consecutive frames), each overlaid with the cyan circular mask returned by the cascade detector; subtle variations in mask size and position across frames reflect systolic dilation and diastolic recoil of the arterial wall. (E) Time-series curves of the carotid radius (in pixels) versus time (in seconds) computed by the Viola-Jones-based segmenter. Panels a, b, and c illustrate successful tracking with physiologically plausible cyclic patterns, whereas panel d illustrates a catastrophic tracking failure with the collapse of the estimated radius to zero and physiologically implausible negative values.

Figure 5. Mask R-CNN cross-sectional segmentation, 24-direction radial sampling algorithm, and time-series wall deformation output.

(A) Point-by-point manual annotation of the carotid wall in a cross-sectional ultrasound frame using the VIA software with the polygon tool, producing pixel-precise ground-truth masks for training and validation of the cross-sectional Mask R-CNN model. The annotation conforms closely to the medial-adventitial interface and accommodates a non-circular vessel geometry. (B) Illustration of the custom 24-direction radial sampling algorithm: after Mask R-CNN segmentation, the centroid of the segmented region is determined as the mean of x and y coordinates of wall pixels; from the centroid, 24 radial directions are defined at angular intervals of 15° (24 × 15° = 360°); the wall pixel intersected by each radial direction is identified (white markers); the Euclidean distance from the centroid to each wall pixel is computed; the mean of these 24 distances yields the mean radius for that frame. The red overlay highlights the 15° angular interval between two consecutive radial directions of the disc. (C) Frame-by-frame tracking of the carotid artery using the cross-sectional Mask R-CNN model (2×3 grid of consecutive frames); the masks conform to the irregular geometry of the vessel and

capture subtle elliptical deformation across the cardiac cycle. (D) Time-series curves of carotid radius (in pixels) versus time computed by the Mask R-CNN-based segmenter; panels a, b, and c illustrate successful tracking with physiologically plausible cyclic patterns; panel d illustrates a less severe tracking failure than that observed with Viola-Jones (Figure 4E), manifesting as an isolated spike rather than a catastrophic collapse.

References

1. Pandian JD, Gall SL, Kate MP, Silva GS, Akinyemi RO, Ovbiagele BI, et al. Prevention of stroke: a global perspective. Lancet. 2018;392(10154):1269–1278.

2. Anderson I, Robson B, Connolly M, Al-Yaman F, Bjertness E, King A, et al. Indigenous and tribal peoples' health (The Lancet–Lowitja Institute Global Collaboration): a population study. Lancet. 2016;388(10040):131–157.

3. King M, Smith A, Gracey M. Indigenous health part 2: the underlying causes of the health gap. Lancet. 2009;374(9683):76–85.

4. Cardoso AM, Tavares IDN, Werneck GL. Indigenous health in Brazil: from vulnerable to protagonists. Lancet. 2022;400(10368):2011–2014.

5. Armstrong ADC, Ladeia AMT, Marques J, Armstrong DMFO, Silva AML, Morais Junior JC, et al. Urbanization is Associated with Increased Trends in Cardiovascular Mortality Among Indigenous Populations: the PAI Study. Arq Bras Cardiol. 2018;110(3):240–245.

6. Armstrong ADC, Souza CDF, Santos JM, Carmo RF, Armstrong DMFO, Pereira VC, et al. Urbanization and cardiovascular health among Indigenous groups in Brazil. Commun Med (Lond). 2023;3(1):17.

7. Nicacio JM, Souza CDF, Gomes OV, Souza BV, Lima JAC, Carmo RF, et al. Cardiac Biomarkers in a Brazilian Indigenous Population Exposed to Arboviruses: A Cross-Sectional Study. Viruses. 2024;16(12):1902.

8. Gomes OV, Souza CDF, Nicacio JM, Carmo RF, Pereira VC, Barral-Netto M, et al. Epidemiology of chronic kidney disease in older Indigenous peoples of Brazil: findings from a cross-sectional survey. Aging Clin Exp Res. 2023;35(10):2201–2209.

9. Pereira VC, Coelho DLLC, Santos JM, Armstrong DMFO, Patriota PVAM, Lima JAC, et al. Traditional pipe smoking (xanduca) and respiratory function in the Fulni-ô indigenous people, Brazil: Project of Atherosclerosis among Indigenous Populations (PAI) study. J Bras Pneumol. 2022;48(2):e20210468.

10. Carmo TR, Santos LG, Farias AVSG, Cavalcanti MPL, Dias GS, Nunes BEBR, et al. Can urbanisation influence alcohol consumption by Indigenous groups? A brief analysis of Brazilian data. Drug Alcohol Rev. 2022;41(4):890–894.

11. Koike BDV, Valões RMP, Cazal C, Pereira VC, Fraga CAC, Carmo RF, et al. Oral health of an indigenous population in northeastern Brazil: a cross-sectional study of the Fulni-ô ethnic group. Sao Paulo Med J. 2024;142(1):e2022355.

12. Øygarden H. Carotid intima-media thickness and prediction of cardiovascular disease. J Am Heart Assoc. 2017;6(1):e005313.

13. Yang WI, Kang KW, Lee HY, Shim CY, Hong GR, Chung N, et al. Incremental predictive value of carotid arterial strain in patients with stroke. Int J Cardiovasc Imaging. 2018;34(6):893–902.

14. Mota A, Santos AM, Salvioni NCP, Vasconcelos HD, Ramos RP, Venkatesh BA, et al. Tele-ultrasound in community with limited resources: The association of hypertension on reduced carotid artery circumferential strain. JODI. 2023;1(1):e2023001.

15. Biswas M, Kuppili V, Araki T, Edla DR, Godia EC, Saba L, et al. Deep learning strategy for accurate carotid intima-media thickness measurement: an ultrasound study on Japanese diabetic cohort. Comput Biol Med. 2018;98:100–117.

16. Jain PK, Sharma N, Roy S. Hybrid Deep Learning Models for Segmentation of Atherosclerotic Plaque in B-mode Carotid Ultrasound Image. In: Sharma S, Subudhi B, Sahu UK, editors. Intelligent Control, Robotics, and Industrial Automation. Singapore: Springer Nature Singapore; 2023. p. 807–819.

17. Jamthikar A, Gupta D, Johri AM, Mantella LE, Saba L, Suri JS. A machine learning framework for risk prediction of multi-label cardiovascular events based on focused carotid plaque B-Mode ultrasound: A Canadian study. Comput Biol Med. 2022;140:105102.

18. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science. 2019;366(6464):447–453.

19. Char DS, Shah NH, Magnus D. Implementing Machine Learning in Health Care — Addressing Ethical Challenges. N Engl J Med. 2018;378(11):981–983.

20. Larrazabal AJ, Nieto N, Peterson V, Milone DH, Ferrante E. Gender imbalance in medical imaging datasets produces biased classifiers for computer-aided diagnosis. Proc Natl Acad Sci U S A. 2020;117(23):12592–12594.

21. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential Biases in Machine Learning Algorithms Using Electronic Health Record Data. JAMA Intern Med. 2018;178(11):1544–1547.

22. Song J, Zou L, Li Y, et al. Combining artificial intelligence assisted image segmentation and ultrasound based radiomics for the prediction of carotid plaque stability. BMC Med Imaging. 2025;25(1):89.

23. Carroll SR, Garba I, Figueroa-Rodríguez OL, Holbrook J, Lovett R, Materechera S, et al. The CARE Principles for Indigenous Data Governance. Data Sci J. 2020;19:43.

24. He K, Gkioxari G, Dollár P, Girshick R. Mask R-CNN. In: Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV). Venice: IEEE; 2017. p. 2980–2988.

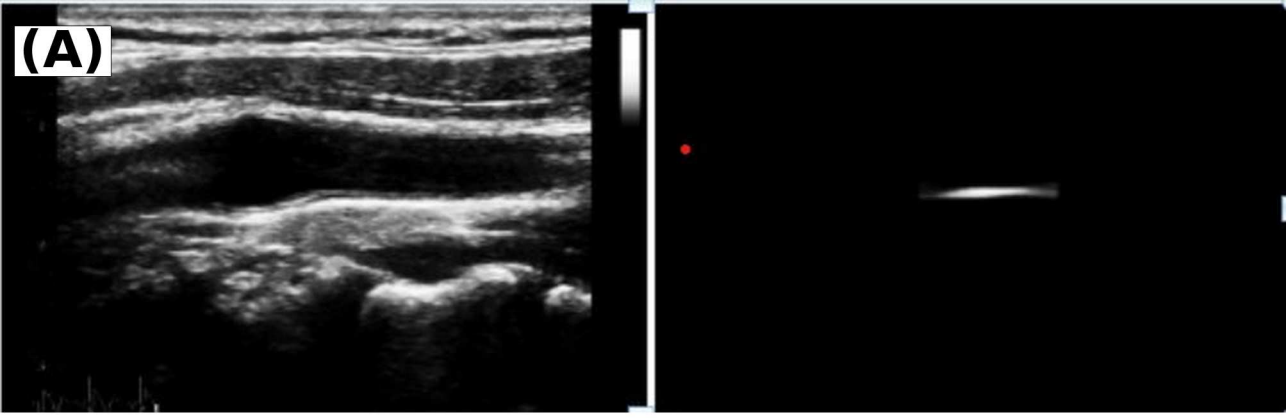
25. Gonzalez RC, Woods RE. Digital image processing. 4th ed. New York: Pearson; 2018.

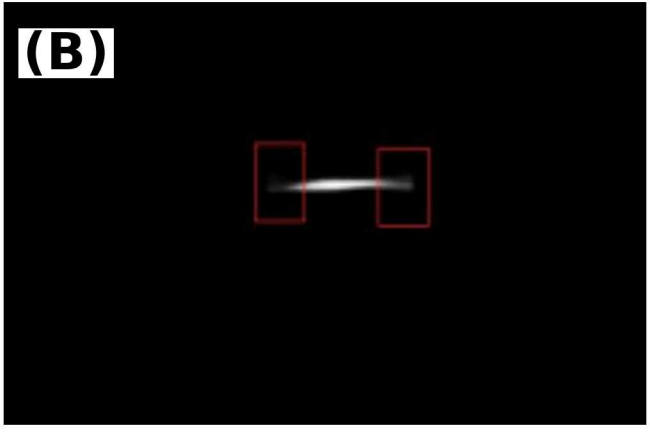
26. Otsu N. A Threshold Selection Method from Gray-Level Histograms. IEEE Trans Syst Man Cybern. 1979;9(1):62–66.

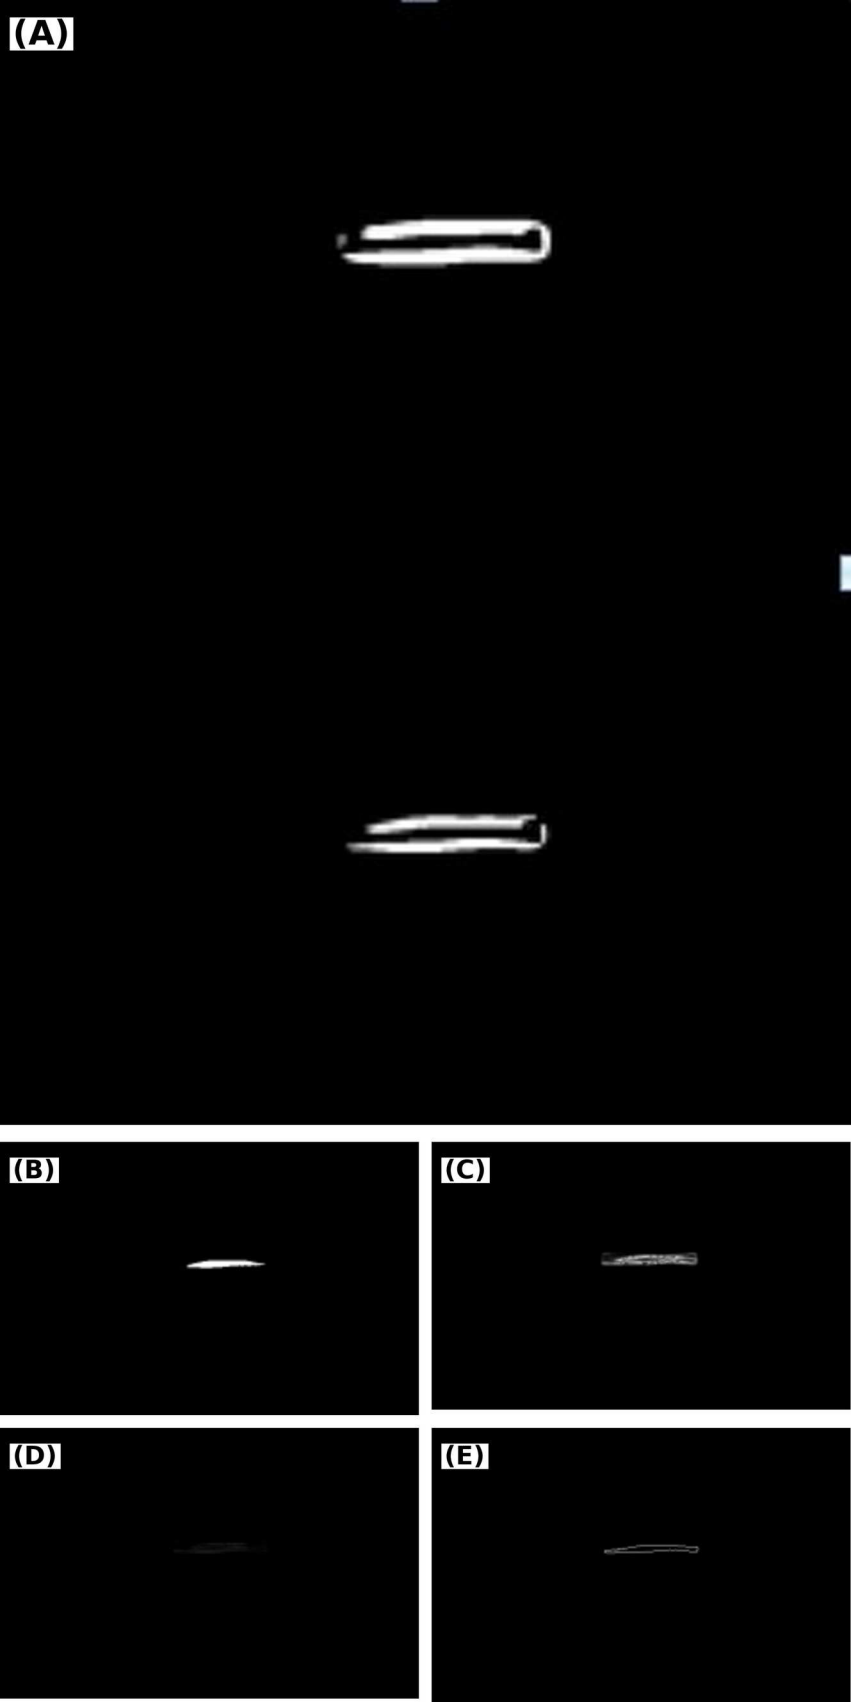
27. Canny J. A Computational Approach to Edge Detection. IEEE Trans Pattern Anal Mach Intell. 1986;PAMI-8(6):679–698.

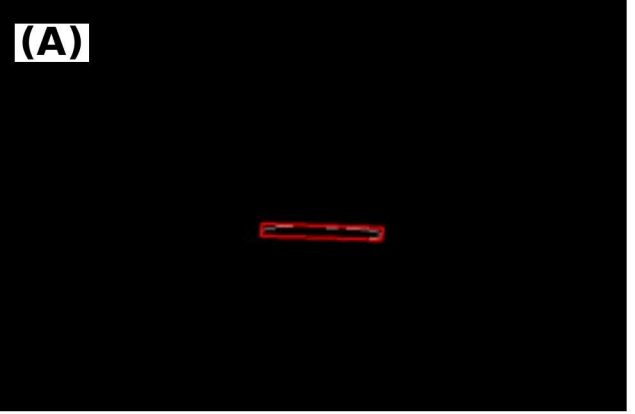
28. Viola P, Jones M. Rapid object detection using a boosted cascade of simple features. In: Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2001). Kauai (HI): IEEE Computer Society; 2001. p. I-511–I-518.

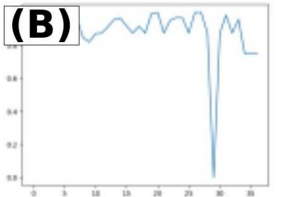
29. Wahl B, Cossy-Gantner A, Germann S, Schwalbe NR. Artificial intelligence (AI) and global health: how can AI contribute to health in resource-poor settings? BMJ Glob Health. 2018;3(4):e000798.

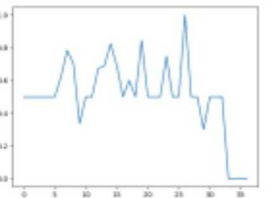

(A)

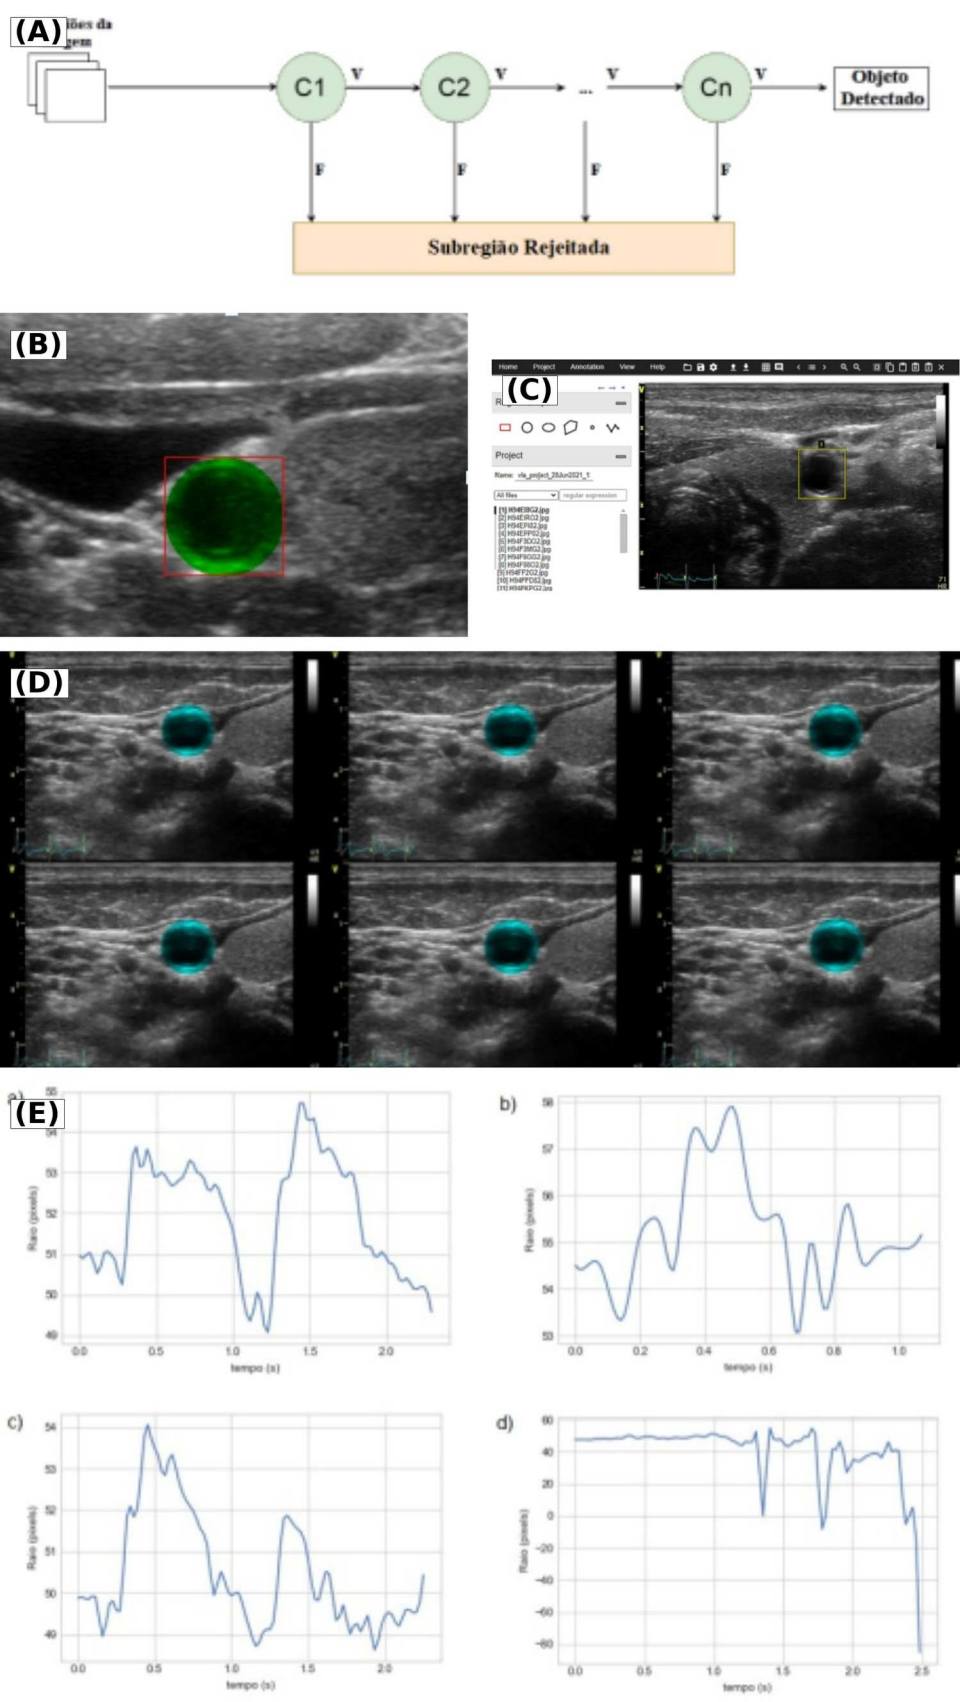

(B)

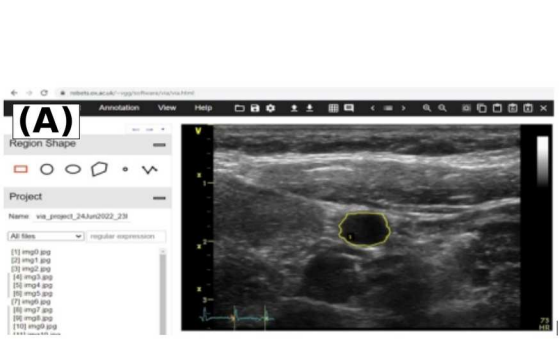

(A)
(B)
(C)
(D)
(E)

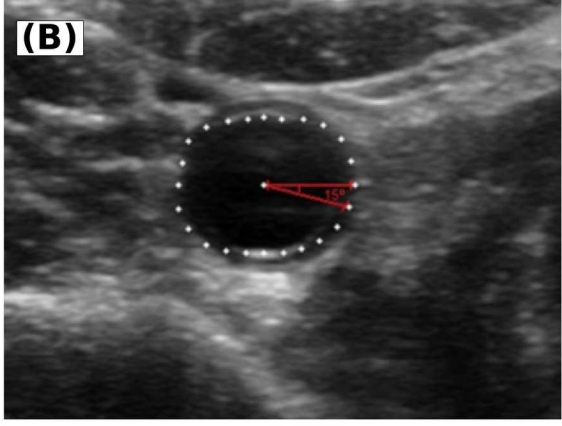

(A)

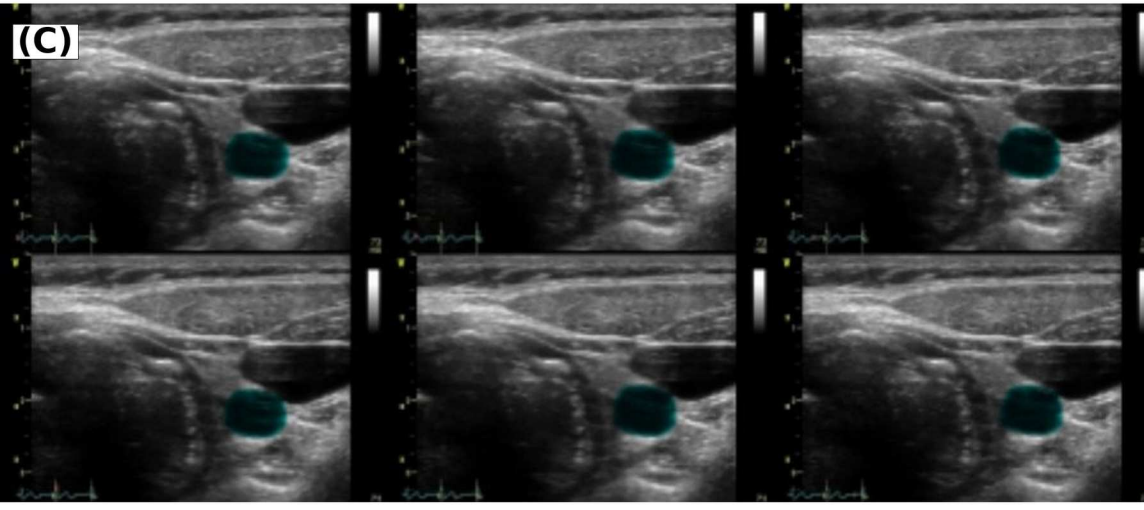

(B)

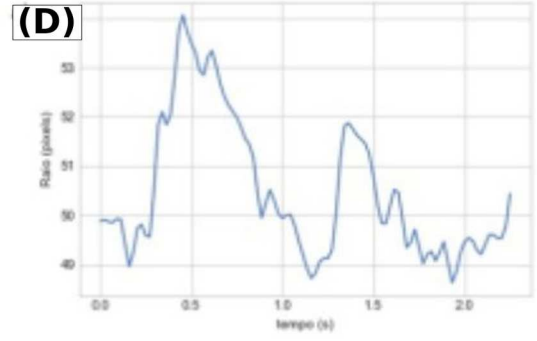


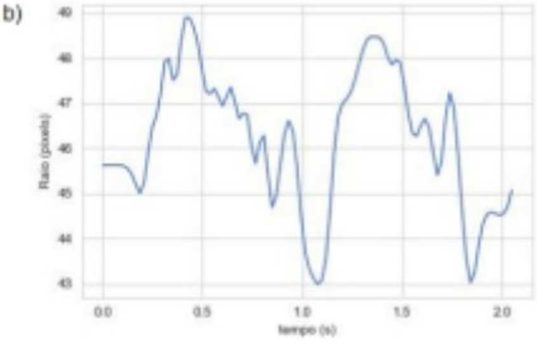

(A)
C1
C2
...
Cn
V
F
Objeto Detectado
Subregião Rejeitada
(B)
(C)
Home
Project
Annotation
View
Help
Project
(D)
(E)
a)
b)
c)
d)
Raio (pixels)
tempo (s)

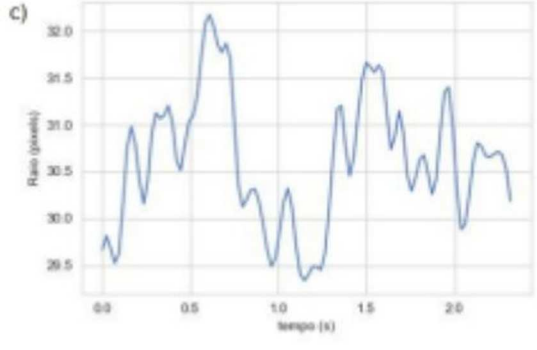

(A)
Annotation
View
Help
Region Shape
Project

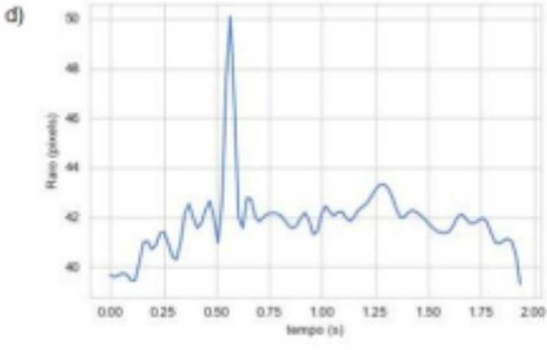

(B)
150


(C)


(D)
Raio (pixels)
tempo (s)


b)
Raio (pixels)
tempo (s)


c)
Raio (pixels)
tempo (s)


d)
Raio (pixels)
tempo (s)