Invitation to submit an EoI to commercialize and deploy AI-based cervical cancer screening in LMICs

Frequently Asked Questions (FAQ)

This FAQ document consolidates questions from potential partners regarding deployment and scale-up of AI-assisted Automated Visual Evaluation (AVE) for cervical cancer screening in low- and middle-income countries (LMICs).

1. Technical Requirements

Q: What are the minimum device specifications?

A: The model was developed using images from Samsung J8, A21, and A23 phones. Our testing shows introducing photos from new phones leads to some performance degradation, but that relatively small training datasets (~500 images) were sufficient to improve model performance with images from different devices. Further analysis, and selection of a deployment approach that maintains acceptable performance, would be the responsibility of the respondent / eventual partners.

Q: Are camera add-ons or lighting equipment required?

A: We expect partners to determine this. We found during the development of the current model that if a phone's auto-focus succeeds in focusing on the cervix (and not the speculum, vaginal wall, etc.) and the cervix is illuminated, the image is almost always of sufficient quality both for human review and the AVE model.

Q: What qualifies a cervical image as "high quality"? This could include resolution, lighting requirements, color accuracy, or specific anatomical coverage.

A: Image should contain most or all of the cervix, the cervix should occupy at least 30% of the image, and the cervix should be illuminated and in focus. The AVE code package will include an image quality algorithm (and training pipeline for further development) that can help ensure high quality images are captured. For the core model, the cervix bounding box is resized to height x width of 384 x 512 (after rotating images if they have 'portrait' orientation as usually the case for phones). Because the cervix doesn't occupy the whole image, the original image should be at least about double this in each dimension. In the end, the main test of image quality is whether the quality negatively impacts the core models ability to accurately detect precancer+ vs. <precancer.

Q: Beyond a "positive or negative result", are there expectations for confidence scores, image regions of interest, or other diagnostic information to be displayed to the end user?

A: We do not have expectations or requirements in this regard. The version of the product used during research displayed only "positive" or "negative." We welcome respondents to propose what they believe will be a feasible and useful output to end users.

Q: Is offline functionality required? Should AI inference run locally on the device, in the cloud, or hybrid?

A: We invite respondents to determine and propose the most appropriate deployment model (offline, mobile edge, cloud-based, or hybrid), taking into account factors such as performance, connectivity in rural health facilities, data security, regulatory requirements, cost, and scalability.

Q: What turnaround time for results is expected?

A: During the patient's appointment, so within a few minutes. Some early market testing identified a turnaround time of <2 minutes as ideal.

Q: Are there specific accuracy benchmarks (sensitivity, specificity, PPV, NPV)?

A: We would expect a similar level of performance to that demonstrated in our 5-country observational study. More details on the model and performance to-date are available online:

- Internal validation of Automated Visual Evaluation (AVE) on smartphone images for cervical cancer screening in a prospective study in Zambia: https://pmc.ncbi.nlm.nih.gov/articles/PMC11176573/
- A Prospective 5-Country Observational Study of the Performance of a New Artificial Intelligence-Based Tool for Cervical Pre-Cancer Screening in the African Region: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4968226

2. Open Source, Data & Intellectual Property

Q: What exactly will be open-sourced and under what license?

A: The code is being open sourced by Global Health Labs under a permissive license, such as the MIT license. The training recipe, weights, and code will all be included.

Q: What rights will partners have to adapt and commercialize solutions? What conditions apply to derivatives of the open-sourced model?

A: Partners will have the right to adapt, deploy, and commercialize solutions built off the code. Partners will retain ownership of any intellectual property (IP) they generate, including derivatives of the open-sourced code, while ensuring compliance with any applicable licensing requirements of the original open-source components and access commitments.

Q: Who owns additional data collected by partners?

A: We expect partners would retain ownership of the data they collect. Negotiated arrangements with governments or collaborating entities may influence data ownership and governance responsibilities.

Q: What governance model for community contributions or updates to the open-source package are envisioned?

A: Respondents are encouraged to propose whether / what type of governance model they would find useful, and why.

Q: Are there restrictions on technology use in other regions?

A: We are not aware of any specific restrictions. We expect national regulations and cervical cancer treatment guidelines to inform use and adoption. Please note that the model was trained only with data from five African countries, so it is expected that additional data and studies may need to be conducted for use in other regions.

3. Funding & Support

Q: What financial, technical, and regulatory support is available? What forms of market-shaping support (e.g., volume guarantees, advance purchase commitments) are available?

A: Please review the call for an example list. Respondents are invited to specify the type and nature of support they may require to deploy and scale the solution, and why such support is needed. This information will be used to refine the eventual scope of support, contingent on prioritized need, compliance with target product characteristics, and available funding.

Q: Are matching funds or co-investments required?

A: We have not yet determined or specified any such requirement, but co-investment is encouraged because funding available is unlikely to cover the full regulatory approval process and entirety of go-to-market activities.

Q: Are there ROI, equity, or revenue model expectations?

A: At this stage, there are no specific ROI, equity, or revenue model expectations. Respondents are encouraged to propose pricing models and describe their approach to ensuring affordability, sustainability, and scalability in LMIC contexts. More detailed pricing information will be requested at the RFP stage.

4. Regulatory and Quality

Q: Which regulatory clearances are required (FDA, CE, national)? Are separate approvals needed for each country? Is ISO 13485 certification mandatory?

A: We invite respondents to propose the regulatory clearance pathway they believe is most appropriate. Proposals should take into account our overarching goal of maximizing access and ensuring scalability across LMIC markets and explain how their approach to regulatory would enable these objectives. We anticipate AI-enabled tools would need to be included in the national cervical cancer screening guidelines of each country of deployment. Device and software must meet appropriate

international quality standards; depending on final product this could include IEC 62304, ISO 13485, and/or ISO 14971.

Q: Will hardware add-ons require separate approvals?

A: We anticipate hardware add-ons may require independent regulatory clearance, depending on their intended medical use and jurisdiction. This would be the responsibility of respondents to eventually determine as part of their plans for productization and regulatory engagement.

Q: Are additional clinical trials required beyond existing studies?

A: This will depend on the regulatory pathway chosen and pursued by respondents.

Q: Will CHAI/Unitaid provide direct support in regulatory submissions?

A: CHAI's contribution may include sharing regulatory learnings, facilitating connections with relevant stakeholders, and providing feedback on proposed strategies, with the aim of accelerating progress toward approval and market access. Responsibility for preparing, submitting, and securing regulatory approvals will remain with the respondent.

5. Deployment & Operations

Q: What is the roadmap for rollout and scaling across LMICs?

A: By end of 2027 the expectation is to have the AVE model deployed and productised, with introduction underway in multiple LMICs. The intermediary milestones and phased rollouts may be co-created between RfP winners, CHAI, and Unitaid. Respondents are invited to propose such roadmaps.

Q: Which countries and sites are prioritized, and why?

A: Current CHAI/Unitaid focal countries for AVE introduction include Kenya, Malawi, Nigeria, Rwanda, Zambia, Zimbabwe. Respondents do not need to prioritize all focal countries, and may also prioritize other LMICs.

Q: Are there minimum deployment scale requirements?

A: There is no set minimum deployment scale.

Q: Should urban or rural areas be prioritized?

A: Respondents are encouraged to submit proposals that they believe will maximize access and usability of the solution long-term.

Q: Who is responsible for training, equipment maintenance, and user support? Who covers costs of cloud infrastructure and maintenance?

A: Respondents / eventual partners would be responsible for these items.

Q: What are the requirements for multilingual support and accessibility?

A: The more specificity respondents can provide regarding plans to ensure usability and feasibility across LMICs, the better.

6. Integration & Interoperability

Q: Should AVE be a standalone screening solution, or can it be integrated with other workflows? What level of integration or interoperability with existing digital health systems is expected? Is compliance with HL7 FHIR, DICOM, or WHO SMART guidelines required?

A: The goal is to maximize use of the tool, so respondents are encouraged to design for integration with existing digital health platforms and health information systems wherever possible. Proposals should address interoperability, alignment with user workflows, and minimizing additional burden on end users. To further support broad access and scalability, proposals should also outline plans for intellectual property management, including options for sub-licensing or open-source configurations where feasible.

Q: Is centralization of post-deployment data mandatory?

A: No.

7. Training Data

Q: What is the size of the training dataset that will be shared with partners?

A: Potentially up to 32,000 images linked to de-identified participant clinical information.

Q: How have the images been captured to date, and should partner solutions be compatible with similar capture conditions?

A: Images were captured by a nurse conducting a cervical exam with acetic acid in a primary or secondary health facility in African settings. These conditions are consistent with the target use case for the product and so partner solutions should be compatible with such conditions.

Q: What ground truth mechanisms were used for data annotation (e.g., classification, localization)? Will annotated datasets include lesion classification, localization, or only binary outcomes?

A: The model was trained based on the histopathology outcome. The main outcome is a binary classification: Precancer+ vs. <Precancer, where Precancer is defined as CIN2-3. We also have HPV

result and VIA result for each patient image. Women who were negative on three screenings (HPV, VIA, and AVE) were assumed to be reference negative and thus no histopathology result is available. We have annotations on the location of the cervix within the image, and annotations for image quality (manual review of whether the cervix is in focus).

8. Commercialization & Pricing

Q: What is the target price range for devices/services in LMICs? What pricing models are acceptable (per scan, subscription, bundled, licensing)?

A: There is not yet a target price for the device, but a commitment to make the device accessible at affordable price and of acceptable quality in LMICs to public sector (as a priority). We invite respondents to share initial thinking on feasible pricing and approaches to reaching maximum scale. Respondents are invited to propose the most appropriate commercialization approach for the AVE solution (e.g., per scan, per device, subscription model, or other), considering factors such as affordability, sustainability, and scalability which should also be included in the response.

Q: Should government procurement or private market be prioritized?

A: We invite respondents to propose the most appropriate approach to maximize access, sustainability, and scalability of the solution. As our priority is to detect as many cases of precancer as possible, and most health care in LMICs is provided through the public sector, we anticipate successful proposals will include (but need not be limited to) a strategy to scale through the public sector.

Q: Are regional exclusive contracts allowed?

A: As the core model is being open sourced, we do not see the value in regional exclusivity.

Q: What KPIs will be used to evaluate commercial success?

A: This has not yet been defined, but our focus is on maximizing access, sustainability, and scalability of the solution.

9. Risk & Legal

Q: Who holds liability in cases of malpractice or device failure?

A: Respondents would be responsible for their product once they have deployed it.

Q: What are the protocols for recalls, regulatory changes, or downtime?

A: This would be determined between the respondent and the regulatory body through which they have secured approvals.

10. Application, Eligibility, & Selection Process

Q: What are the main evaluation criteria in the RfP?

A: The RfP and its evaluation criteria are still in development and will focus on analyzing the feasibility, scalability, and usability of the respondent's proposal.

Q: Will multiple AI tools and approaches be supported, or is convergence around a single model expected?

A: Our goal is to maximize accessibility and usage of AI solutions for cervical cancer screening to prevent as many cancer cases and deaths as possible. We are open to supporting multiple partners and approaches, depending on the proposals we receive. While our financial and technical support is directed toward the model developed by Global Health Labs, our broader objective is to ensure access to low-cost, effective cervical cancer screening. We do not intend to force the market to converge around a single model, but rather to contribute to a diverse and sustainable ecosystem of solutions.

Q: Is LMIC experience mandatory?

A: This is not a strict requirement, but respondents should provide evidence for their capacity to ensure product accessibility, scale, and usability in LMICs.

Q: Who is eligible to apply (e.g., developers, device manufacturers, academics, governments)? Is consortium formation or collaboration with NGOs/governments allowed?

A: Yes. We expect respondents could include software developers, medical device manufacturers, academics, governments, and AI-based solution entrepreneurs and innovators, and more. Consortiums and collaborations are eligible.

Q: Do we need to have a registered AI product to be eligible to respond to this call?

A: No, this is not a requirement.

Q: Will partnerships with local suppliers, health systems, or funders be facilitated? Are there existing networks and engagement mechanisms already?

A: Yes - our ambition is to contribute to a diverse and sustainable ecosystem for scaling AI-based tools for cervical cancer screening, which we expect would require robust engagement and partnership with local entities.

Q: Is the project limited to Africa, or open to other LMICs?

A: The project is open to other LMICs with a focus on introduction in the target countries in Africa.

Q: What is the expected RfP release date and selection timeline?

A: We expect to release the RfP by end-2025 and select partner(s) in Q1 2026.

Q: Is there a required template for applications?

A: There is no specified template to respond to this call. Templates to respond to the RfP will be shared at a later date.