



Artificial intelligence for oral health in low- and middle-income countries: mobile health, teledentistry, and frugal innovation

Inteligencia artificial para la salud oral en países de ingresos bajos y medios: salud móvil, teleodontología e innovación frugal

Ana María Chaves Cano¹  

¹ Fundación Universitaria Juan N. Corpas. Bogotá, Colombia.

Submitted: 29-07-2025 | Revised: 05-10-2025 | Accepted: 29-11-2025 | Published: 30-11-2025

Corresponding Author: Ana María Chaves Cano 

ABSTRACT

Artificial intelligence is being applied to oral health in low- and middle-income countries through three overlapping channels: mobile health, teledentistry, and frugal innovation. This review contrasts the conditions, uses, and evidence of these applications against those of high-income settings. In high-income countries, artificial intelligence in dentistry is dominated by image analysis of radiographs and intraoral scans, supported by rich datasets and established regulatory pathways; in low- and middle-income countries, the applications are oriented toward access, using smartphones as the capture device, teledentistry as the delivery channel, and low-cost models as the means, and they target the caries and oral cancer burden that the workforce shortage leaves unattended. The evidence is strongest for teledentistry as a route to access, moderate for smartphone-based screening, and weakest for the claim that artificial intelligence itself, as opposed to the mobile channel around it, changes oral health outcomes. Most reported performance comes from high-income datasets, and transfer to low-resource populations is largely unproven. The review concludes that the value of artificial intelligence in low- and middle-income countries currently lies less in the algorithm than in the integration of the algorithm into affordable, connected workflows, and that realizing that value depends on local data, local validation, and the infrastructure that the literature, so far, mostly takes for granted.

Keywords: artificial intelligence; oral health; developing countries; telemedicine; mobile health; frugal innovation.

RESUMEN

La inteligencia artificial se está aplicando a la salud oral en los países de ingresos bajos y medios a través de tres canales superpuestos: la salud móvil, la teleodontología y la innovación frugal. Esta revisión contrasta las condiciones, los usos y la evidencia de estas aplicaciones con los de los entornos de ingresos altos, donde la inteligencia artificial se orienta al análisis de imágenes, mientras que en los países de ingresos bajos y medios se orienta al acceso. La evidencia es más sólida para la teleodontología como vía de acceso, moderada para el cribado basado en teléfonos inteligentes y más débil para la afirmación de que la inteligencia artificial en sí cambia los resultados de salud oral. La mayor parte del rendimiento comunicado procede de conjuntos de datos de ingresos altos, y la transferencia a poblaciones de bajos recursos está sin demostrar. La revisión concluye que el valor de la inteligencia artificial reside hoy menos en el algoritmo que en su integración en flujos de trabajo asequibles y conectados, y que materializar ese valor depende de datos locales, validación local e infraestructura.

Palabras clave: inteligencia artificial; salud oral; países en desarrollo; telemedicina; salud móvil; innovación frugal.

INTRODUCTION

Artificial intelligence has arrived in oral health with a striking asymmetry. In high-income countries, the dominant applications are the automated analysis of radiographs, intraoral photographs, and cone-beam computed tomography, built on large annotated datasets, evaluated against clinical gold standards, and increasingly regulated as medical devices; in low- and middle-income countries, where the burden of untreated caries and oral cancer is highest and the dental workforce is thinnest, the applications take a different shape, oriented not to refining diagnosis for an already-served population but to extending reach to a population that currently receives almost no care.(1,2)

The contrast between these two ways of applying the same technology is the subject of this review, and it is informative precisely because the technology is the same while the conditions are not: the algorithm that reads a radiograph does not care whether the patient is in Seoul or in rural Bangladesh, but the infrastructure that produces the image, transmits it, and acts on its result differs entirely between the two settings.

The asymmetry is not a matter of intellectual lag but of constraint: a low- and middle-income setting cannot wait for the density of radiograph machines, specialists, and datasets that high-income artificial intelligence assumes, and it therefore reaches for the tools that are already present, the smartphone, the messaging network, and the community health worker. The scale of the unmet need gives the asymmetry its urgency: untreated caries in permanent teeth and severe periodontitis are among the most prevalent conditions on earth, and oral cancer is disproportionately a disease of South and Southeast Asia, where late diagnosis converts a curable lesion into a fatal one.(1,2)

A technology that could move even a fraction of this burden from late to early, from unexamined to examined, would be among the more cost-effective uses of artificial intelligence in health, and it is this possibility, rather than any particular algorithm, that motivates the work reviewed here. This review examines three channels through which artificial intelligence reaches oral health in low- and middle-income countries, namely mobile health, teledentistry, and frugal innovation, and contrasts each against its high-income counterpart, in the belief that the contrast, more than either side taken alone, reveals what the technology can and cannot yet do for the world's oral health.

The three channels are conceptually distinct but practically entangled. Mobile health uses the smartphone as a platform for screening, promotion, and self-care, turning a device that is already in people's pockets into an instrument of oral health. Teledentistry uses telecommunications to connect a patient or a generalist provider to a distant clinician, and artificial intelligence enters it as the assistant that analyzes the transmitted image before the clinician does.

Frugal innovation is the design philosophy of building for low-resource settings, and in this context it means the low-cost, often smartphone-based, sometimes open-source systems that are engineered to work where a radiographic machine or a specialist is unavailable. Each of these channels has a high-income analogue: mobile health in high-income countries means apps for behavior change and monitoring, teledentistry means remote consultation in well-resourced systems, and frugal innovation is largely absent because the constraint it addresses does not exist.(3,4)

The comparison therefore exposes not just what low- and middle-income countries do differently, but what the high-income literature takes for granted, and it is in the gaps, the missing datasets, the missing external validations, the missing referral pathways, that the real difference between the two settings is visible. A reader approaching this topic from the high-income artificial-intelligence literature may find the low- and middle-income applications technically modest, and in a narrow sense they are, but the modesty is the point: the constraint is not the capacity of the model but the reach of the system, and a model that runs on a phone in a village clinic, however simple, is doing work that a more sophisticated model in a teaching hospital cannot.

The review organizes its synthesis around two contrasts that cut across all three channels. The first is the contrast between low- and middle-income and high-income settings, in disease burden, workforce, infrastructure, and data, and it is this contrast that makes the low- and middle-income applications intelligible as more than poorer copies of high-income ones.

The second is the contrast between promise and demonstrated impact, between what a pilot demonstrates is possible and what a study shows is achieved, because the literature on artificial intelligence in global health is rich in the former and, so far, sparse in the latter.(5,6) The two contrasts are related: the promise-versus-impact gap is widest precisely in the settings where the infrastructure contrast is sharpest, because a technology that is promising in a well-resourced pilot faces its hardest test in the low-resource environment where it is most needed, and the gap between those two conditions is where the literature is thinnest.

The purpose is therefore not to catalogue every application but to show where the evidence is strong, where it is aspirational, and what would need to be true for the promise to become practice.

METHODS

This is a narrative review, and its methods are documented in the manner that the SANRA guidance recommends for reviews that synthesize rather than exhaust. Sources were identified by structured searches of PubMed, OpenAlex, and Crossref using terms for artificial intelligence, mobile health, teledentistry, frugal innovation, and oral health, restricted for the core literature to the years 2013 through 2025 and supplemented by systematic reviews and foundational papers without a lower date limit.

The searches were run during the review period, and the selection of sources was guided by relevance to the three channels and to the two contrasts, rather than by a prespecified exhaustive protocol. Where a statement of effectiveness or accuracy is made, the underlying study design is stated, because a proof-of-concept and a randomized trial carry different weight, and the synthesis distinguishes them throughout.

The comparability of the settings is addressed explicitly rather than assumed: where a low- and middle-income and a high-income result are compared, the comparison notes the differences in device, population, and outcome definition that could account for the difference, so that apparent gaps in performance are not mistaken for genuine deficits in the technology. The review does not perform a meta-analysis, and its conclusions are qualitative, organized by contrast rather than by a pooled estimate.

DEVELOPMENT

Teledentistry: the access channel with the strongest evidence

Teledentistry is the oldest of the three channels and the one with the most substantial evidence base, and its core finding is that remote examination is broadly concordant with in-person examination for the conditions that matter most, caries and oral mucosal lesions.^(7,8) The diagnostic-accuracy literature is specific enough to be useful: systematic reviews of teledentistry for caries diagnosis with extra-oral photography report agreement with clinical examination that is moderate to high, and a diagnostic meta-analysis of teledentistry for oral lesions reaches a similar conclusion, that remote assessment is a credible triage instrument even when it is not a replacement for the specialist's eye.^(9,10)

The same pattern holds at the level of the individual study, with smartphone telediagnosis of caries in geriatric patients achieving accuracy in the range of clinical examination, and this is the kind of evidence that supports deployment in settings where a clinician is not otherwise available.⁽¹¹⁾

Systematic reviews have repeatedly found that teledentistry improves access, particularly for rural and underserved populations, and that patient and provider acceptance is high, while the evidence for its effect on hard clinical outcomes remains thinner than the evidence for its effect on access.⁽¹²⁻¹⁴⁾

The qualitative literature adds the mechanism that the quantitative reviews only imply: teledentistry works as an enabler to care by removing the travel and the waiting that keep underserved patients away, and its value is experienced as much in the referral it produces as in the diagnosis it reaches.⁽¹⁵⁾

The role of artificial intelligence within teledentistry is more recent and less settled: where AI is added, it is typically to analyze the photograph transmitted by a smartphone or a mid-level provider before a clinician reviews it, and the available studies report diagnostic performance that approaches, but does not clearly exceed, a trained clinician reading the same image.^(16,17)

The contrast with the high-income literature is instructive: in high-income settings, teledentistry is a convenience and a triage tool, whereas in low- and middle-income settings it is a substitute for a clinician who is not otherwise reachable, and this difference changes what counts as success, from agreement with a reference standard to the number of patients who would otherwise receive no assessment at all.^(18,19)

The WhatsApp- and photograph-based telediagnosis studies of oral lesions illustrate the same point from the other side: they report high concordance with specialist diagnosis, but their

value lies in the referral they trigger, not in the image quality they achieve.(20,21)

The orthodontic monitoring case adds a further nuance, because it is the one teledentistry application where the remote workflow demonstrably reduces the number of in-person visits without loss of outcome, and it does so because the endpoint, tooth movement, is objectively measurable from photographs, which is a property that caries and oral cancer screening, with their softer endpoints, do not fully share.(22,23)

The lesson is that teledentistry's evidence is strongest where the outcome is measurable and the workflow is image-based, and weaker where the outcome is behavioral or the workflow requires a physical examination, a distinction that any policymaker deciding where to invest should apply directly.

Mobile health: promotion, screening, and the limits of the app

Mobile health in oral health spans two quite different uses, and the contrast between them is sharp. The first use is promotion and self-care, where the smartphone delivers reminders, education, and behavior-change support, and the evidence, from systematic reviews, is that such tools are feasible and acceptable but that their effect on clinical outcomes such as plaque or caries is modest and inconsistently demonstrated.(3,24)

The promotion literature is large in volume and small in effect: apps can remind, educate, and nudge, but translating those nudges into a measurable reduction in disease requires the sustained behavior change that a reminder alone does not reliably produce, and the reviews conclude that the promise of the channel currently outruns its demonstrated effect. The randomized trials that do exist, for example of smartphone reminders to improve oral hygiene during orthodontic treatment, show at most a transient or modest benefit, which is the expected result for an intervention that changes behavior through a channel the patient can ignore as easily as they can read it.(3)

This does not make promotion worthless, it remains the most scalable way to reach a population with prevention messages, but it does mean that promotion should be judged by reach and engagement rather than by a reduction in caries that the channel is not well placed to deliver alone. The second use is screening, where the smartphone camera captures an intraoral image that is then analyzed, increasingly by an artificial-intelligence model, and here the reported performance is more striking, with studies of caries and lesion detection on smartphone photographs reporting accuracy in the range of clinical examination.(25-27)

The screening use has a specific attraction for low- and middle-income settings that promotion does not share: a screening image can be captured by a non-specialist, a parent, or even the patient, and analyzed centrally, which decouples the detection of disease from the presence of a clinician, and this is precisely the decoupling that a workforce-shortage setting needs. The maternal teledentistry studies make this concrete, showing that a parent's smartphone photograph, analyzed remotely, can carry enough information to flag early childhood caries, and the tele-detection and referral protocols under trial extend the same logic to a population level.(26,28) The contrast between the two uses matters because they have different evidentiary requirements: a promotion app needs a trial showing behavior or outcome change, and the trials are few, whereas a screening app needs a diagnostic-accuracy study, and the studies are many but predominantly small and single-site.(28,29)

The further contrast, between mobile health in low- and middle-income and in high-income settings, is that the same app is a convenience in one and a lifeline in the other, and the external-validation studies that would establish whether a model built on high-income images works on low-income populations are exactly the studies that are missing.(30,31)

The external-validation pilot that does exist, for an AI gingivitis-detection tool among older adults in daycare centers, is a model of the kind of evidence the field needs, small, local, and reported with its limitations, and it is the exception rather than the rule.(29)

The mobile channel also reaches beyond detection into treatment support and behavior change, and the applications that combine a screening function with a motivational or educational function are the ones that most closely resemble the high-income app literature, with the same mixed evidence for their effect on hard outcomes, which suggests that the smartphone's greatest contribution in low-resource settings is as a capture and delivery device rather than as a stand-alone behavior-change instrument.(3,24)

Frugal artificial intelligence: the oral cancer and caries frontier

Frugal innovation is where the low- and middle-income specificity is most explicit, because the systems are built for the constraint itself. The strongest line of work concerns oral cancer

and potentially malignant disorders, which are concentrated in South and Southeast Asia and are diagnosed late precisely because screening is scarce; here, point-of-care, smartphone-based imaging with neural-network classification has been demonstrated in proof-of-concept and pilot studies that report encouraging accuracy, and in some cases the systems have been deployed in surveillance programs.(32-35)

The oral-cancer frugal literature has a coherence that the caries literature lacks, because the disease is concentrated, the consequences of missing it are severe, and the alternative, a specialist examination that the patient will not receive, is absent, so even a modest-accuracy triage tool has a clear place in the pathway.(36)

The Thai surveillance platform, built for remote oral-cancer self-examination and screening, and the multimodal mobile imaging systems under development both show the direction: the phone captures, the model flags, and a human decides, and the model's job is to reduce the number of patients who reach a specialist too late, not to make the final diagnosis.(35,37) Low-cost fluorescence and multimodal mobile platforms extend the same idea to treatment guidance and risk-based management, and the studies share a common logic: use a device that is already present, a smartphone, and a model that runs on it or near it, to bring specialist-grade triage to a community that has no specialist.(37,38)

For caries, the frugal direction is newer and thinner, with smartphone-based detection and classification reported mostly at proof-of-concept, and the evidence base does not yet reach the standard set by the teledentistry literature.(30,39) The caries frugal work is nonetheless strategically important, because caries is by far the more common condition and the one for which a screening tool would have the largest reach, but its challenge is that a caries screening tool must discriminate not just disease from health but the subset of lesions that warrant intervention from those that can be watched, and this finer discrimination is exactly what the proof-of-concept studies have not yet demonstrated.(31)

The contrast with high-income settings is again the point: high-income oral-cancer AI is built on large, curated datasets and aims to match or beat a specialist, while the frugal systems aim to do something a specialist is not present to do, and this difference in aim must be kept in mind when comparing their reported accuracies.(5,36) A comparison that treats a frugal triage model and a high-income diagnostic model as the same task will conclude that the frugal model is worse, when the fairer conclusion is that they are solving different problems under different constraints, and neither can be substituted for the other.

The structural contrast: burden, workforce, and data

The three channels are all responses to a structural imbalance that is easy to state and hard to close: the burden of oral disease and the shortage of the workforce that could address it are concentrated in the same places, low- and middle-income countries, and the data and infrastructure that modern artificial intelligence depends on are concentrated elsewhere.(1,2,40)The geographic maldistribution of the dental workforce, the so-called dental deserts, is itself a well-documented phenomenon that predates artificial intelligence and that no amount of algorithmic sophistication will by itself reverse, because a model cannot substitute for a clinic, a referral, or a course of treatment.(40)

The workforce shortage is the binding constraint, because artificial intelligence, whatever its accuracy, requires an image to analyze and a clinician to act on the result, and where neither is available the algorithm solves only part of the problem.(41,42) In fragile and conflict-affected states the constraints are sharper still, because the oral health system that a digital tool would plug into may be absent or disrupted, and the review literature on such settings treats infrastructure as the precondition that the technology literature mostly assumes.(41)

The data shortage is the second constraint, and it is the one that most directly limits the transfer of high-income models to low-resource populations, because a model trained on high-income images may not perform on images of different populations, disease presentations, or acquisition devices.(6)

The reviews of artificial intelligence developed in low- and middle-income countries themselves note that the literature is growing but that it is dominated by small, single-centre studies, that the datasets are in most cases not public, and that the evidence for real-world deployment, as opposed to proof of concept, is the exception.(5,6)

The practical consequence is that a policymaker cannot yet rely on an off-the-shelf model validated elsewhere, and must instead plan for local data collection and local validation as part of any adoption, which is a cost the literature does not account for.

Barriers and enablers across the three channels

Across the three channels, the barriers and enablers are remarkably consistent, and they matter because they determine whether a promising pilot becomes a durable service. The enablers are the spread of smartphones, the reach of messaging networks, and the existence of a cadre of community or mid-level health workers who can capture images and act on referrals, and these are precisely the assets that the frugal and mobile designs are built to exploit, which is why the designs concentrate on the capture device and the channel rather than on hardware that a low-resource clinic cannot afford.(3,34)

The barriers are connectivity, power, data, and the absence of a referral destination, and they are the mirror image of the enablers, the same infrastructure that makes the tool possible is, where unreliable, what makes it fail, so the enablers and the barriers are not separate lists but two descriptions of the same dependence on infrastructure.(42)The provider-side barriers are behavioral as much as technical, because a tool that no clinician is trained or willing to use will not be adopted however accurate it is, and the awareness-and-attitude surveys show that readiness is uneven even where the technology is available, which makes training and change management as much a part of implementation as the software itself.(43)

The most consequential enabler, and the one the literature mentions least, is a functioning referral pathway, because a detection without a destination is a result without a consequence, and the systems that report the most durable outcomes are those, like the orthodontic monitoring tools, that are embedded in an existing clinical workflow rather than bolted onto a disconnected one.(22,23)

This is why the pathway, and not the model, is the unit that a funder or a ministry should evaluate, because the model can be swapped, but a broken pathway breaks every model that passes through it, and the history of digital health is full of technically sound tools that failed for want of a destination.

The synthesis across the three channels is therefore that the technology is the easy part, and the infrastructure, the workforce, and the pathway are the hard part, which is the opposite of how the literature, with its emphasis on model accuracy, frames the problem. A condensed view of the three channels and their evidence is given in (Table 1).

Table 1. Synthesis of the three innovation channels.

Channel	Primary use	Target condition	Evidence strength	Main limitation
Teledentistry	Remote diagnosis and referral	Caries, oral lesions	Strong for access and agreement	Thin on clinical outcomes
Mobile health	Promotion and smartphone screening	Caries, gingivitis	Moderate for screening	Few outcome trials
Frugal artificial intelligence	Point-of-care smartphone AI	Oral cancer, caries	Proof of concept	No external validation

Note: The evidence column reflects the maturity and strength of the evidence for each channel's effect on access and on clinical outcomes. Source: Summary of the reviewed literature.

DISCUSSION

The differences documented in the development section have explanations that are contextual rather than technical, and distinguishing the two is the main work of this discussion. The finding that teledentistry has the strongest evidence while frugal artificial intelligence has the weakest is explained by maturity and by the difficulty of the target: teledentistry is a mature, low-technology intervention whose effect on access is easy to measure, whereas frugal artificial intelligence is a young intervention aimed at a hard problem, oral cancer detection, whose clinical endpoint is distant and whose validation is demanding.(7,32)

The same logic explains the internal ordering within each channel: within teledentistry, the access outcomes are the best documented because access is measurable by counting visits and referrals, while the diagnostic outcomes are harder because they require a reference standard and a blinded comparison, and the clinical outcomes are hardest of all because they require follow-up and, ideally, randomization, which is why the evidence thins as one moves from access to diagnosis to outcome.(44)

The finding that reported accuracies are lower in low-resource settings than in high-income settings is explained, at least in part, by the data, the devices, and the populations, and it would be an error to read a lower reported accuracy as evidence that the technology is unsuited, just as it would be an error to read a high pilot accuracy as evidence that a system is deployable.(6)

The two errors are mirror images of the same over-reading, and the discipline of holding them apart is what separates a careful reading of this literature from an impressionistic one. Association and causation must be kept apart here: the observation that high-income countries report higher accuracies is associated with richer data and stricter evaluation, and does not by itself establish that a given model would fail in a low-resource setting, but the absence of external validation means the question is unanswered rather than answered in either direction.

A further contextual factor is the workforce itself: in a high-income setting an artificial-intelligence system competes with an available specialist, and must therefore clear a high bar of non-inferiority, whereas in a low-resource setting it competes with nothing, and therefore clears the bar by existing, which is why the two literatures report such different framing of the same technology even when their accuracies are similar.(5)

This difference in the comparison group is one of the least appreciated reasons that the two bodies of evidence read so differently, and any synthesis that ignores it will misjudge the technology in both directions at once.

The three channels also interact in ways that the literature treats separately, and this separation is itself a finding. A frugal smartphone system for oral cancer is a mobile-health intervention, a teledentistry intervention, and an artificial-intelligence intervention at once, and its success depends on all three layers working together, the device, the transmission, and the model.(16,34)

This layered character means that the three channels are best understood not as alternatives but as a stack, in which mobile health supplies the device, teledentistry supplies the delivery, and artificial intelligence supplies the analysis, and the value of the whole is limited by whichever layer is weakest in a given setting, a fact that explains why the same system can succeed in one region and stall in another.

Studies that isolate one layer, the accuracy of the model, say, without the connectivity and the referral pathway, measure something necessary but not sufficient, and the evidence-gap maps of teledentistry make this point by showing that the literature is richest on diagnostic agreement and poorest on the downstream outcomes that actually constitute care.(18,44)

A synthetic view therefore suggests that the value of artificial intelligence in low- and middle-income oral health currently lies less in the algorithm than in the workflow into which the algorithm is embedded, and that an algorithm of moderate accuracy embedded in a working referral pathway is worth more than a state-of-the-art algorithm stranded in a pilot.

This is not an argument against improving the algorithms, but a reminder that the binding constraints in this setting are, in order, the clinician who is absent, the connectivity that is unreliable, and the model that is imperfect, and that the field's attention has been concentrated on the third while the first two remain the harder problems. What artificial intelligence specifically adds, beyond the mobile and telehealth channel that already delivers images, is scalability of interpretation: a clinician can review a handful of transmitted photographs, but only a model can screen the volume of images that a population-level program would generate, and it is at that scale, where the human is the bottleneck, that the algorithm becomes genuinely indispensable rather than merely helpful.(16,27)

The appropriate measure of the algorithm's value in a low-resource setting is therefore not its accuracy in isolation but the throughput it enables per scarce clinician-hour, and this is a metric that almost none of the reviewed studies report, which is a further sign of how far the field's framing has drifted from the setting's needs. The acceptance side of the same gap is better documented than the outcome side: surveys of dental practitioners show awareness of teledentistry but uneven adoption, and patient-satisfaction reviews report positive experience with remote care, which means the human substrate for adoption exists even where the clinical-outcome evidence does not.(43,45)

The orthodontic use case is the exception that proves the rule, because remote monitoring of treatment is one of the few teledentistry applications with direct evidence of reduced visits and maintained outcomes, and it works because the outcome is measurable and the workflow is already image-based.(22,23)

A synthetic view therefore suggests that the value of artificial intelligence in low- and middle-income oral health currently lies less in the algorithm than in the workflow into which the algorithm is embedded, and that an algorithm of moderate accuracy embedded in a working referral pathway is worth more than a state-of-the-art algorithm stranded in a pilot.

The contrast between promise and demonstrated impact is the most consistent theme across all three channels, and it deserves to be stated bluntly. The literature is rich in proof-of-concept demonstrations that a smartphone and a model can detect caries or oral cancer, and it is far poorer in the prospective, population-based evaluations that would show whether those detections change care.(5,6)

The reasons are partly generic, the difficulty and cost of field trials, and partly specific, the fragmented regulatory and funding landscape for digital health in low- and middle-income countries, where the institutions that would commission or demand such evidence are themselves under-resourced.

The consequence is that a policymaker reading this literature can be confident that the tools are possible, and cannot yet be confident that they work where they are needed, and the gap between those two kinds of confidence is the gap the field must close. It is also worth stating what would count as closure: not a larger accuracy study but a demonstration that a detection, once made, reliably produces a referral, and that a referral, once made, reliably produces care, because those two links are where most of the real-world value, and most of the real-world failure, will be found.

The conditions for closing it are concrete and are the same across the three channels: local data, because population and device shift are real; local validation, because external validation from a different population is the minimum standard; and infrastructure, because connectivity, power, and a referral pathway are the substrate without which the algorithm has nothing to act on.(41,42)

A research agenda that follows from these conditions would prioritize, in order, the prospective evaluation of the full pathway rather than the model alone, the release of locally collected datasets, and the head-to-head comparison of low-cost systems against the current standard of care, which in a low-resource setting is often no screening at all rather than a specialist examination. The field does not lack demonstrations; it lacks the evaluations that would convert a demonstration into a recommendation, and that conversion is where the next decade of work lies.

This review has limitations that follow from its method. It is a narrative review, not a systematic one, and its selection of sources is guided by relevance rather than by an exhaustive protocol, so it may omit relevant work, particularly work published outside the indexed, English-language literature in which these applications are increasingly reported, and the true volume of activity in low- and middle-income countries is likely understated here. It does not appraise the risk of bias of the included studies with a formal tool, and its synthesis of accuracy is qualitative, so the statements about performance should be read as indicative rather than as pooled estimates.

The contrast between low- and middle-income and high-income settings is itself a simplification, because both categories are internally heterogeneous, and a middle-income country with a developed dental sector differs from a low-income country with almost none, a distinction the review acknowledges but does not fully resolve.

The framing around artificial intelligence also risks overstating its role: many of the benefits attributed to artificial intelligence in this literature are really benefits of the mobile and telehealth channels that surround it, and a review that insisted on isolating the algorithm's contribution would find a smaller body of direct evidence, which is itself a finding the review returns to throughout. These limitations bound the precision of the conclusions without changing their direction.

CONCLUSIONS

Artificial intelligence for oral health in low- and middle-income countries is real, growing, and oriented toward access rather than toward the refinement of diagnosis that dominates high-income settings. Teledentistry has the strongest evidence as a route to access, mobile health spans promotion and screening with uneven evidence, and frugal artificial intelligence is the most specific but least validated, concentrated on oral cancer and caries at the point of care.

The reported accuracy of these systems is generally lower than in high-income settings, and this reflects data, devices, and evaluation as much as technology, but the transfer of high-income models to low-resource populations is largely unproven.

The value of artificial intelligence in this context lies in its integration into affordable, connected workflows, not in the algorithm alone, and realizing that value requires local data, local validation, and the infrastructure that connects a detection to a referral and a referral to care.

The field has demonstrated that the tools are possible; it has not yet demonstrated that they work where they are needed, and closing that gap, rather than another increment of model accuracy, is the task that would make the promise real.

The path to that demonstration is not mysterious, and it is the same across the three channels: build on the channels that already work, teledentistry and the smartphone, validate the models on local data, and evaluate the whole pathway, from detection to referral to care, rather than the algorithm in isolation. If the field does those things, the next review of this topic will be able to say what this one cannot, which is whether artificial intelligence, integrated into the mobile and telehealth infrastructure that low- and middle-income countries already possess, actually changes the course of oral disease for the people who currently carry its burden.

REFERENCES

1. Peres MA, Macpherson LMD, Weyant RJ, Daly B, Venturelli R, Mathur MR, et al. Oral diseases: a global public health challenge. *The Lancet*. 2019;394(10194):249-260. doi: 10.1016/s0140-6736(19)31146-8
2. Watt RG, Daly B, Allison P, Macpherson LMD, Venturelli R, Listl S, et al. Ending the neglect of global oral health: time for radical action. *The Lancet*. 2019;394(10194):261-272. doi: 10.1016/s0140-6736(19)31133-x
3. Pascadopoli M, Zampetti P, Nardi MG, Pellegrini M, Scribante A. Smartphone Applications in Dentistry: A Scoping Review. *Dentistry Journal*. 2023;11(10):243. doi: 10.3390/dj11100243
4. Surdu A, Foia CI, Luchian I, Trifan D, Budala DG, Scutariu MM, et al. Telemedicine and Digital Tools in Dentistry: Enhancing Diagnosis and Remote Patient Care. *Medicina*. 2025;61(5):826. doi: 10.3390/medicina61050826
5. Reddy CL, Mitra S, Meara JG, Atun R, Afshar S. Artificial Intelligence and its role in surgical care in low-income and middle-income countries. *The Lancet Digital Health*. 2019;1(8):e384-e386. doi: 10.1016/s2589-7500(19)30200-6
6. Umer F, Adnan S, Lal A. Research and application of artificial intelligence in dentistry from lower-middle income countries – a scoping review. *BMC Oral Health*. 2024;24(1):220. doi: 10.1186/s12903-024-03970-y
7. Gurgel-Juarez N, Torres-Pereira C, Haddad AE, Sheehy L, Finestone H, Mallet K, et al. Accuracy and effectiveness of teledentistry: a systematic review of systematic reviews. *Evid Based Dent*. 2022;1-8. doi: 10.1038/s41432-022-0257-8
8. Estai M, Kanagasingam Y, Tennant M, Bunt S. A systematic review of the research evidence for the benefits of teledentistry. *J Telemed Telecare*. 2018;24(3):147-156. doi: 10.1177/1357633x16689433
9. Kargozar S, Jadidfard MP. Teledentistry accuracy for caries diagnosis: a systematic review of in-vivo studies using extra-oral photography methods. *BMC Oral Health*. 2024;24(1):828. doi: 10.1186/s12903-024-04564-4
10. Uhrin E, Domokos Z, Czumbel LM, Kóti T, Hegyi P, Hermann P, et al. Teledentistry: A Future Solution in the Diagnosis of Oral Lesions: Diagnostic Meta-Analysis and Systematic Review. *Telemedicine and e-Health*. 2023;29(11):1591-1600. doi: 10.1089/tmj.2022.0426
11. Pandey P, Jasrasaria N, Bains R, Singh A, Manar M, Kumar A. The Efficacy of Dental Caries Telediagnosis Using Smartphone: A Diagnostic Study in Geriatric Patients. *Cureus*. 2023;15(1):e33256. doi: 10.7759/cureus.33256
12. Alavi S, Gholami M, Malik L, Matti R, Oktai S, Al-Najafi F, et al. Assessment of teledentistry in improving access to dental care: a systematic review. *Australian Dental Journal*. 2025;70(1):4-41. doi: 10.1111/adj.13045
13. Al-Buhaisi D, Karami S, Gomaa N. The role of teledentistry in improving oral health outcomes and access to dental care: An umbrella review. *J of Oral Rehabilitation*. 2024;51(11):2375-2389. doi: 10.1111/joor.13836
14. Emami E, Harnagea H, Shrivastava R, Ahmadi M, Giraudeau N. Patient satisfaction with e-oral health care in rural and remote settings: a systematic review. *Syst Rev*. 2022;11(1):234. doi: 10.1186/s13643-022-02103-2
15. Irving M, Stewart R, Spallek H, Blinkhorn A. Using teledentistry in clinical practice as an enabler to improve access to clinical care: A qualitative systematic review. *J Telemed Telecare*. 2018;24(3):129-146. doi: 10.1177/1357633x16686776
16. Department of Orthodontics, Faculty of Dentistry, Jamia Millia Islamia, New Delhi, India, Batra P, Tagra H, HCL Technologies Ltd, New Delhi, India, Katyál S, Department of Orthodontics, Faculty of Dentistry, Jamia Millia Islamia, New Delhi, India. Artificial Intelligence in Teledentistry. *Discoveries*. 2022;10(3):e153. doi: 10.15190/d.2022.12
17. Wayne State University, Detroit, Michigan, Kaushik R, Rapaka R, University of Connecticut, Storrs, Connecticut. A Patient-Centered Perspectives and Future Directions in AI-powered Teledentistry. *Discoveries*. 2024;12(4):e199. doi: 10.15190/d.2024.18
18. Kengne Talla P, Allison P, Bussi  res A, Rodrigues A, Bergeron F, Giraudeau N, et al. Teledentistry for Improving Access To, and Quality of Oral Health Care: Overview of Systematic Reviews and Meta-Analyses. *J Med Internet Res*. 2025;27:e65211-e65211. doi: 10.2196/65211
19. Kall  s MS, Krause DTW, Childs CA, Marchini L, Busse AL. Teledentistry for Older Community-Dwelling Patients: A Scope Review of Concordance Between Virtual and Clinical Examination. *Special Care in Dentistry*. 2025;45(5):e70101. doi: 10.1111/scd.70101
20. Nguyen J, Yang S, Melnikova A, Abouaki M, Lin K, Takesh T, et al. Novel Approach to Improving Specialist Access in Underserved Populations with Suspicious Oral Lesions. *Current Oncology*. 2023;30(1):1046-1053. doi: 10.3390/curroncol30010080
21. Israel MS, Freire NDA, Rodrigues BTG, Dias YM, Martins MD, de Sousa CAM, et al. Remote Consultations for Monitoring Actinic Cheilitis: A Preliminary Non-Randomized Clinical Trial. *Oral Diseases*. 2025;31(12):3316-3322. doi: 10.1111/odi.70010
22. Sangalli L, Alessandri-Bonetti A, Dalessandri D. Effectiveness of dental monitoring system in orthodontics: A systematic review. *J Orthod*. 2024;51(1):28-40. doi: 10.1177/14653125231178040
23. TORRES DKB, SANTOS MCCD, NORMANDO D. Is teledentistry effective to monitor the evolution of orthodontic treatment? A systematic review and meta-analysis. *Dental Press J Orthod*. 2023;28(4):e2322195. doi: 10.1590/2177-6709.28.4.e2322195.oar
24. Fern  ndez CE, Maturana CA, Coloma SI, Carrasco-Labra A, Giacaman RA. Teledentistry and mHealth for Promotion and Prevention of Oral Health: A Systematic Review and Meta-analysis. *J Dent Res*. 2021;100(9):914-927. doi: 10.1177/00220345211003828
25. Estai M, Kanagasingam Y, Mehdizadeh M, Vignarajan J, Norman R, Huang B, et al. Mobile photographic screening for dental caries in children: Diagnostic performance compared to unaided visual dental examination. *J Public Health Dent*. 2022;82(2):166-175. doi: 10.1111/jphd.12443
26. Iranparvar P, Ghorbani Z, Zandieh S, Hartoonian S, Shirazi M. Evaluating Maternal Accuracy in Smartphone-Based Tele-Dentistry for Early Childhood Caries Detection. *Clinical & Exp Dental Res*. 2025;11(5):e70231. doi: 10.1002/cre2.70231

27. Adnan N, Faizan Ahmed SM, Das JK, Aijaz S, Sukhia RH, Hoodbhoy Z, et al. Developing an AI-based application for caries index detection on intraoral photographs. *Sci Rep*. 2024;14(1):26752. doi: 10.1038/s41598-024-78184-x
28. Aly NM, El Kashlan MK, Giraudeau N, El Tantawi M. A Tele-detection and referral pathways model for early childhood caries control- a protocol for a randomized factorial study: the TRACE study. *BMC Oral Health*. 2024;24(1):934. doi: 10.1186/s12903-024-04706-8
29. Chau RCW, Cheng ACC, Mao K, Thu KM, Ling Z, Tew IM, et al. External Validation of an AI mHealth Tool for Gingivitis Detection among Older Adults at Daycare Centers: A Pilot Study. *International Dental Journal*. 2025;75(3):1970-1978. doi: 10.1016/j.identj.2025.01.008
30. Boy AF, Akhyar A, Arif TY, Syahrial S. Artificial intelligence for dental caries detection: A mixup, fine-tuning, and quantization approach on the MobileNetV2 model. *J Conserv Dent Endod*. 2025;28(8):764-771. doi: 10.4103/jcde.jcde_362_25
31. Kaur A, Jyoti D, Sharma A, Yelam D, Goyal R, Nath A. Deep caries detection using deep learning: from dataset acquisition to detection. *Clin Oral Invest*. 2024;28(12):677. doi: 10.1007/s00784-024-06068-5
32. Uthoff RD, Song B, Sunny S, Patrick S, Suresh A, Kolur T, et al. Point-of-care, smartphone-based, dual-modality, dual-view, oral cancer screening device with neural network classification for low-resource communities. *PLoS ONE*. 2018;13(12):e0207493. doi: 10.1371/journal.pone.0207493
33. Song B, Sunny S, Li S, Gurushanth K, Mendonca P, Mukhia N, et al. Mobile-based oral cancer classification for point-of-care screening. *J Biomed Opt*. 2021;26(06):1-10. doi: 10.1117/1.jbo.26.6.065003
34. Wongsapai M, Wudtijureepun K, Tamdee D, Suthachai T, Wuttisarnwattana P, Suebnukarn S. Oral Cancer Surveillance in Remote Areas Using a Smartphone Digital Platform. In: *Studies in Health Technology and Informatics*. IOS Press; 2025. p. 1898-1899. doi: 10.3233/shti251270
35. Wongsapai M, Jaidee E, Thanathornwong B, Suebnukarn S, Wuttisarnwattana P. RISKOCA: A Smartphone-Based Digital Platform for Oral Cancer Self-Examination. In: *Studies in Health Technology and Informatics*. IOS Press; 2024. p. 164-165. doi: 10.3233/shti240369
36. Haj-Hosseini N, Lindblad J, Hasséus B, Kumar VV, Subramaniam N, Hirsch JM. Early Detection of Oral Potentially Malignant Disorders: A Review on Prospective Screening Methods with Regard to Global Challenges. *J Maxillofac Oral Surg*. 2024;23(1):23-32. doi: 10.1007/s12663-022-01710-9
37. Mitbander R, Brenes D, Coole JB, Kortum A, Vohra IS, Carns J, et al. Development and Evaluation of an Automated Multimodal Mobile Detection of Oral Cancer Imaging System to Aid in Risk-Based Management of Oral Mucosal Lesions. *Cancer Prev Res (Phila)*. 2025;18(4):197-207. doi: 10.1158/1940-6207.capr-24-0253
38. Siddiqui SA, Siddiqui S, Hussain MAB, Khan S, Liu H, Akhtar K, et al. Clinical evaluation of a mobile, low-cost system for fluorescence guided photodynamic therapy of early oral cancer in India. *Photodiagnosis and Photodynamic Therapy*. 2022;38:102843. doi: 10.1016/j.pdpdt.2022.102843
39. Nayak PP, Shetty V, S S, Zacharias L, Gore I. AI-driven approaches in the management of early childhood caries: A path toward global oral health. *Journal of Oral Biology and Craniofacial Research*. 2025;15(5):1134-1140. doi: 10.1016/j.jobcr.2025.07.022
40. Edwards JW, Bogale B, Gallagher JE. Tackling geographic barriers to primary dental care (dental deserts): a systematic review. *Br Dent J*. 2025. doi: 10.1038/s41415-025-8487-8
41. Bogale B, Scambler S, Mohd Khairuddin AN, Gallagher JE. Oral health system strengthening in fragile and conflict-affected states: A systematic review. *J Glob Health*. 2024;14:04132. doi: 10.7189/jogh-14-04132
42. Kengne Talla P, Inquimbert C, Dawson A, Zidarov D, Bergeron F, Chandad F. Barriers and Enablers to Implementing Teledentistry From the Perspective of Dental Health Care Professionals: Protocol for a Systematic Quantitative, Qualitative, and Mixed Studies Review. *JMIR Res Protoc*. 2023;12:e44218. doi: 10.2196/44218
43. Lin GSS, Koh SH, Ter KZ, Lim CW, Sultana S, Tan WW. Awareness, Knowledge, Attitude, and Practice of Teledentistry among Dental Practitioners during COVID-19: A Systematic Review and Meta-Analysis. *Medicina*. 2022;58(1):130. doi: 10.3390/medicina58010130
44. Nascimento da Silva Mulder J, Ramos Pinto M, Aníbal I, Dornellas AP, Garrido D, Huanca C, et al. Teledentistry Applied to Health and Education Outcomes: Evidence Gap Map. *J Med Internet Res*. 2024;26:e60590. doi: 10.2196/60590
45. Ingle NA, AlOraini RA, Radhi FM, AlSaydalani GK, Mansour AM, AlHarkan AF, et al. Impact of AI-Driven Technologies on Patient Satisfaction in General Dentistry: A Systematic Review and Meta-Analysis. *J Pharm Bioallied Sci*. 2025;17(Suppl 2):S1297-S1300. doi: 10.4103/jpbs.jpbs_1920_24

FUNDING

No financing

CONFLICT OF INTEREST

None

AUTHORSHIP CONTRIBUTIONS

Drafting – original draft: Ana María Chaves Cano.
Writing–review and editing: Ana María Chaves Cano.