BOUTIQUE ULTRASOUND · METHODOLOGY WHITEPAPER · V2 · 2026

Ultrasound Vision™: from sound waves to a recognisable face

The imaging physics, the display problem, and how our AI recognises a baby's real face through it — the methodology behind 8K Ultrasound, anchored by Identity Lock.

METHODOLOGY & LITERATURE REVIEW COMPILED BY **MEETLAOMA** · BOUTIQUE ULTRASOUND

BoutiqueUltrasound.com

8K Ultrasound · Ultrasound Vision · Identity Lock

ABSTRACT

What this paper covers

An **8K Ultrasound** portrait turns a parent's real 3D/4D scan into a photo-realistic newborn image. Doing it faithfully is hard for one reason: an ultrasound is not a photograph — it is a map of *sound reflections*, degraded by the very physics that creates it. This paper works from first principles: (1) how an ultrasound image is actually formed, (2) why that image is hard to read and how imaging tries to improve it — contrast, shadows, and clarity, and (3) how our **Ultrasound Vision** pipeline recognises a real face through those constraints, anchored by **Identity Lock**. Each step is grounded in the imaging-physics and peer-reviewed deep-learning literature. We are explicit about what the science does and does not support.

PART 1 · THE PHYSICS

How an ultrasound image is formed

Medical ultrasound is a **pulse-echo** system. A piezoelectric transducer emits a short high-frequency sound pulse (typically 2–15 MHz) into the body, then listens for the echoes that return [P1].

Echoes come from acoustic-impedance boundaries

Every tissue has an **acoustic impedance** Z = density (ρ) × speed of sound (c). An echo is produced wherever Z *changes* — at the boundary between two tissues — and the strength of that echo is proportional to the *difference* in impedance across the boundary [P1]. So an ultrasound image is fundamentally a map of *interfaces*, not of surfaces or colour.

Depth from time, brightness from amplitude

The machine assumes a single average sound speed (1540 m/s) and converts each echo's round-trip **time of flight** into a depth. Each returning echo is then drawn as a dot whose **brightness** encodes its amplitude — bright (hyperechoic) for strong reflectors, dim (hypoechoic) and black (anechoic, e.g. fluid) for weak ones. Sweeping the beam across a plane builds the 2-D **B-mode** (brightness-mode) grayscale image; stacking planes builds the 3-D volume a keepsake studio renders [P1].

> **The key takeaway:** the face in a 3D scan is reconstructed from millions of impedance echoes — a grayscale surface of reflections, carrying *no* colour and physically distorted by how sound travels. Everything downstream has to respect that.

PART 2 · THE DISPLAY PROBLEM

Contrast, shadows, and clarity — and how imaging fights for them

Because the raw echo data spans a vast amplitude range and is corrupted by the body, sonographers and scanners spend most of their effort *improving the displayed image*. Three battles matter for a face.

(a) Black–white contrast — making structure readable

Echo amplitudes span a huge range, so they are compressed into the visible grayscale via **gain** and **dynamic range**. Narrowing the dynamic range raises contrast — low-level "haze" is suppressed while bright structures stay — at the cost of subtle mid-tone detail [P3]. Too much gain washes the face out; too little buries it. The right contrast is what separates the curve of a cheek from the noise around it.

(b) The dark side — attenuation and acoustic shadows

Sound **attenuates** as it travels, so deeper echoes return weaker; without correction the far side of the face would fade to black. **Time-Gain Compensation (TGC)** amplifies later (deeper) echoes so that similar tissue looks similarly bright regardless of depth [P2]. Separately, a strong reflector (a hand, bone, the cord) blocks sound and casts an **acoustic shadow** — a dark region that is *not* empty, just hidden. Distinguishing "dark because hidden" from "dark because absent" is one of the hardest reads in obstetric ultrasound.

(c) Clarity — resolution, speckle, and the trade-offs

Sharpness has two axes: **axial** resolution (along the beam, set by pulse length / frequency) and **lateral** resolution (across the beam, set by beam width and focus). Higher frequency sharpens both but penetrates less — a fixed trade-off. Modern scanners push clarity with **harmonic imaging** (imaging the second-harmonic signal generated in tissue, which raises contrast and spatial resolution and cuts artifacts) and **spatial compounding** (combining views from several angles to suppress speckle — at some cost to lateral resolution) [P3].

And over all of it sits **speckle**: the grainy texture is not random static but a coherent *interference pattern* from sub-resolution scatterers [3]. It simultaneously carries tissue texture and obscures fine edges — which is exactly why naive smoothing destroys the face along with the noise.

Why this matters for a portrait: contrast, shadow, and speckle are not cosmetic — they determine whether a nostril, an eyelid, or a lip edge even *exists* in the data. A generic photo-AI sees none of this; it treats the scan as a picture and fills the gaps by guessing.

PART 3 · THE CORE

How Ultrasound Vision recognises a face through the physics

Ultrasound Vision is built on a single principle: **recognise the real signal, separate it from the physics artifacts, anchor the baby's identity, and only then render.** Each stage corresponds to a problem from Parts 1–2 and mirrors a technique validated in the peer-reviewed deep-learning literature. We run these inside a frontier image model; the citations ground each *stage*, not a claim to reproduce those exact systems.

01 Quality gate

Reject scans too degraded (noise, movement, field-of-view, occlusion) to support a faithful portrait — limits documented across the fetal-reconstruction literature. GROUNDED IN: ALOMAR 2021/22 [1]; SIVERA 2024 [2]

02 Speckle recognition & correction

Because speckle is tissue-dependent and approximately multiplicative, generic de-noisers fail; ultrasound-specific self-supervised methods separate texture from true structure. GROUNDED IN: SPECKLE2SELF, MEDICAL IMAGE ANALYSIS 2025 [3]

03 Shadow / occlusion recognition

We estimate where the image is dark because it is *hidden* (acoustic shadow, limb, cord) versus genuinely absent — so a hand over the mouth is never hallucinated into a smile. GROUNDED IN: MENG ET AL., IEEE TMI 2019 [4]

04 Contrast & clarity enhancement

An enhancement stage plays a role analogous to the scanner's TGC and dynamic-range controls — pulling facial structure out of low-contrast and attenuated regions before identity is read. GROUNDED IN: IMAGE-OPTIMIZATION PRINCIPLES [P2][P3]; SUPER-RESOLUTION [8]

05 Landmark recognition

We locate the baby's real eyes, nose and mouth on the source scan — the same fetal-landmark problem studied at MICCAI and in 3DFETUS — and let the studio confirm them before rendering. GROUNDED IN: 3DFETUS 2025 [5]; XU ET AL., MICCAI 2020 [6]

06 Identity Lock (core method)

The portrait is anchored to those landmarks — to *her* face, not an imagined one. Identity preservation is handled as its own module, a principle established for generative face models; multiple scans of the same baby (Identity Pack) tighten the lock. GROUNDED IN: IDENTITY-AWARE CYCLEGAN [7]

07 Geometry correction

Because scanners assume a fixed 1540 m/s sound speed, real tissue introduces measurable geometric distortion (a large part of why a nose reads wider on a scan than in life); we correct for it. GROUNDED IN: BLAND ET AL. 2015 [9]

08 Heritage colouring

Grayscale ultrasound carries no colour. Skin tone and hair come from parent-provided details as an artistic choice — never a prediction of the baby's true colouring. HONESTY BOUNDARY — NO ACADEMIC BASIS FOR COLOUR INFERENCE

09 Render → up to 8K

Only now do we render on a frontier image model and upscale toward 8K print resolution — the same GAN super-resolution approach validated to improve fetal-ultrasound quality. GROUNDED IN: REAL-ESRGAN ON FETAL US, SCI REP 2025 [8]

This recognise-then-render sequence — not a one-step filter — is what a generic "AI baby" tool skips, and why its outputs drift toward generic doll-faces while ours hold the baby's real structure.

PART 4 · HONESTY

The limits we will not cross

No colour recovery. Every credible result in the fetal-imaging literature is grayscale *shape* only. A scan holds no skin tone, hair, or eye colour — colour is an artistic, parent-guided choice, not a prediction.

No validated "match %". There is no rigorous published figure for how closely a scan-based portrait matches the later baby; cohorts are tiny and validation is geometric. Anyone quoting a match percentage is guessing — we won't.

A likeness, not a forecast. What the scan genuinely carries is facial shape and proportion — which is real, and why parents recognise the nose or lips at birth.

Positioning: ultrasound computer-vision and fetal-face recognition exist in clinical and academic work (e.g. automatic fetal-face navigation on commercial scanners). Ultrasound Vision is, to our knowledge, the first to bring this methodology to consumer keepsake 8K portraits and explain it transparently — not a claim to have invented ultrasound vision itself.

REFERENCES — IMAGING PHYSICS

1. [P1] AAPM/RSNA Physics Tutorial for Residents: Topics in US (RadioGraphics); Radiopaedia, "Acoustic impedance" & "B-mode ultrasound."
2. [P2] "Time Gain Compensation" — ScienceDirect Topics; Wikipedia, Time gain compensation.
3. [P3] Tissue harmonic imaging, spatial compounding & speckle-reduction imaging — Radiology Key, "Advancements in ultrasound"; Resolution and Speckle Reduction in Cardiac Imaging (PMC8034817).

REFERENCES — DEEP LEARNING

1. Alomar et al. "Reconstruction of the 3D fetal face from ultrasound." SCITEPRESS 2021 / Comput. Methods Programs Biomed. 2022.
2. Sivera et al. "Fetal face shape analysis from prenatal 3D ultrasound images." Scientific Reports, 2024.
3. Li, Navab, Jiang. "Speckle2Self: self-supervised ultrasound despeckling." Medical Image Analysis, 2025.
4. Meng et al. "Weakly Supervised Estimation of Shadow Confidence Maps in Fetal Ultrasound." IEEE Trans. Medical Imaging, 2019.
5. Alomar et al. "3DFETUS: standardising fetal facial planes in 3D ultrasound." arXiv:2511.10412, 2025.
6. Xu et al. "Region Proposal Network with Graph Prior… Landmark Detection in 3D Ultrasound." MICCAI, 2020.
7. Huang et al. "Face Translation using Identity-aware CycleGAN." arXiv:1712.00971.
8. Rahman et al. "Enhancing fetal ultrasound image quality with super-resolution." Scientific Reports, 2025.
9. Bland et al. "Geometric distortion of area in medical ultrasound images." arXiv:1502.04611, 2015.
