

# Biomedical Idea Evaluation

## Midjourney's Whole-Body Ultrasound Computed Tomography ("Midjourney Medical Imaging")

Prepared with [Paperclip](https://paperclip.gxl.ai) (<https://paperclip.gxl.ai>) — grounded in full-text biomedical literature. All load-bearing claims verified against source text.

---

### The idea, in one paragraph

A person stands on a platform that descends through a water column at ~5 cm/s. As they descend, the body passes through a ring of ~500,000 sand-grain-sized elements, each acting as both a tiny transmitter and a tiny receiver. The elements emit ultrasonic waves and record the returning ripples at millions of cycles per second, producing terabytes of data per second that stream to a large compute cluster. Because waves change shape whenever they cross a boundary in density or stiffness (water → skin → fat → muscle → bone), the reconstruction software inverts those shape changes into a quantitative 3D map of the body at a fraction of a millimeter. An AI segmentation layer labels structures, and the result is presented as MRI-like imagery acquired at "nearly a hundred times the speed" of MRI. In imaging terms, this is **whole-body ultrasound computed tomography (USCT) driven by full-waveform inversion (FWI), with a deep-learning segmentation overlay.**

---

### 1. Grade

#### **B+**

An excellent and scientifically grounded vision. The physics is sound and most individual claims map onto real, peer-reviewed results. The grade is held below A by two hard barriers that the headline claim understates: (a) ultrasound cannot transmit usefully through the air and dense bone that fill the thorax, abdomen, and pelvis, so every validated success to date is confined to immersible, air-/bone-free anatomy; and (b) the compute required for whole-body sub-millimeter FWI is extreme, with the AI shortcut still diagnostically unproven.

---

## 2. Reasoning and method

**Method.** I decomposed the narrative into its testable scientific claims — wave-to-image physics, sub-mm resolution, acquisition speed, MRI-likeness, AI segmentation, data/compute load, hardware scale, and whole-body feasibility — then searched the full-text biomedical corpus (PMC, preprints) for each. I assembled a verification repository of ten papers, attached a specific quantitative claim to each, and verified every claim against the source full text before citing. Only verified claims appear below.

### What the literature confirms (the idea's strengths):

- **The physics is correct and is an active field.** Inverting wave-shape changes for tissue density/stiffness is exactly full-waveform inversion, demonstrated for soft-tissue and brain imaging [1][2][7]. The narrative's own framing — that the *major task is turning waves into images* — matches the field's consensus that reconstruction is the bottleneck.
- **Sub-millimeter resolution is real.** 3D FWI of the human brain achieved sub-millimetre resolution of about 0.9 mm [1]. A 3D ring-array breast system achieved 0.78 mm resolution and reconstructed lesions as small as 0.3 mm [3]. Clinical breast sound-speed maps are reconstructed at a 0.25 mm pixel size [4].
- **"~100× faster than MRI" is defensible for acquisition.** A full brain dataset can be acquired in ~2 seconds [1]; that is genuinely fast versus a multi-minute MRI sequence. The narrative correctly attributes the slowness to reconstruction, not acquisition.
- **"Looks like an MRI" has empirical support.** Ultrasound-tomography sound-speed maps of the breast correlate strongly with MRI percent water content, supporting USCT as a reproducible, radiation-free surrogate [5].
- **AI segmentation is mature.** Labeling structures in 3D volumes is routine; the segmentation overlay is the least speculative component of the system.

### What the literature shows is hard or unproven (why not an A):

- **Air and bone are near-opaque to ultrasound.** Skull attenuation rises sharply with bone thickness, and bone dominates the loss [9]; reverberation and aberration degrade transcranial images, especially for shallow targets [10]. Every validated soft-tissue result is confined to the *immersible, air-/bone-free* breast [3][4][5] or the *heavily corrected* skull [1]. The thorax (air-filled lungs), abdomen (bowel gas), and pelvis/spine (dense bone) are far harsher than the skull. **No paper in this set — and, to the best of current knowledge, none in the field — demonstrates whole-body transmission tomography.** This is the central gap between the pitch and the science.
- **Whole-body compute is brutal.** 3D FWI of a numerical breast phantom required ~2.3 days on a single A100 GPU [2]; 3D brain FWI required ~32 hours on 128 CPU nodes [1]. A whole-body column is orders of magnitude larger. The "thousands of computers"

framing is honest, and learned/AI reconstruction is the credible escape hatch (~24 ms per image once trained [6]; near-real-time sound-speed estimation [4]) — but its accuracy and generalization on real, diverse, in-vivo whole-body anatomy is unproven.

- **~500,000 elements is ~500× current hardware.** Real ring arrays use 256–1024 elements [3][8]. Half a million elements is a major manufacturing leap (yield, multiplexing, crosstalk, calibration), not a settled engineering detail.
- **The evidence base is largely pre-clinical.** Most cited results are in silico, phantom, or ex vivo; current systems use simulated or archived data [2][6][7]. Breast USCT is the only mature soft-tissue clinical case [4][5].

**Net.** The idea is directionally excellent and built on a real, publishing field with validated components. It is not an A because the headline whole-body/MRI-equivalence claim rests on an unsolved physics barrier and extreme compute, and because the evidence beyond the breast is thin. B+ reflects "validated pillars, gated by hard physical and engineering limits the claim currently understates."

---

### 3. Risks and opportunities

#### Risks

1. **Whole-body physics barrier (high impact, high likelihood).** Air and bone make true whole-body transmission tomography likely infeasible with current physics; realistic use is probably capped at immovable, soft-tissue regions [9][10].
2. **Compute cost and latency (high).** Sub-mm whole-body FWI is enormously expensive [1][2]; clinical throughput depends on AI reconstruction that is not yet diagnostically validated.
3. **AI hallucination in a diagnostic image (high).** A learned segmentation/reconstruction overlay can introduce plausible but false structures — a patient-safety and regulatory hazard. The "crossfade between raw reconstruction and AI segmentation" must keep the raw layer inspectable.
4. **Hardware manufacturability (medium-high).** 500k elements is unproven at scale [3][8].
5. **Regulatory pathway (medium-high).** A generative overlay on a primary diagnostic image invites significant FDA scrutiny; quantitative sound-speed output is more defensible than AI-painted anatomy.

#### Opportunities

1. **A validated beachhead exists.** Breast imaging is the one place USCT is clinically mature [4][5], with a real unmet need (dense-breast screening, where mammography underperforms) and a radiation-free, contrast-free advantage.

2. **Quantitative biomarkers.** Sound-speed and attenuation maps are quantitative tissue measurements, not just pictures — useful for density, composition, and longitudinal tracking [5].
  3. **The AI-reconstruction wave is rising fast.** Learned FWI is collapsing reconstruction time by orders of magnitude [4][6], which is precisely the lever that makes the data/compute story tractable.
  4. **Speed as a workflow advantage.** Fast acquisition [1] enables screening-scale throughput that MRI cannot match.
- 

## 4. Suggested improvements

1. **Reframe the claim around immersible anatomy.** Position as high-resolution tomography of body parts that can sit in water without air or thick bone — breast first, then limbs, neck/thyroid, and pediatric imaging — and treat thorax/abdomen as a research frontier rather than a shipping feature. This keeps the vision while aligning the claim with the physics [9][10].
  2. **Lead commercially with the breast wedge.** It is the validated, reimbursable, high-need entry point [3][4][5]; whole-body becomes a roadmap aspiration, not a launch claim.
  3. **Make the AI layer accountable.** Report calibrated uncertainty on the AI reconstruction/segmentation, keep the raw quantitative reconstruction visible beneath any overlay, and validate against ground truth (MRI/CT/pathology) before any diagnostic claim [6].
  4. **De-risk the hardware in stages.** Validate the element array at intermediate scales (thousands → tens of thousands) before committing to half a million, characterizing yield, crosstalk, and calibration at each step [3][8].
  5. **Invest in the reconstruction stack as the moat.** The bottleneck — and therefore the defensibility — is the wave-to-image step. Learned-FWI accuracy and speed on real anatomy is where the differentiated value sits [1][2][6].
- 

## 5. Papers and databases used

**Databases (via Paperclip):** PubMed Central (PMC) full-text papers; bioRxiv / medRxiv preprints; abstract index. Searches spanned ultrasound computed tomography, full-waveform inversion, ring-array breast imaging, transcranial ultrasound attenuation/aberration, deep-learning sound-speed estimation, 3D medical image segmentation, and breast-cancer detection performance. Ten papers were selected, each carrying a specific verified claim (see References). All ten claims passed full-text verification.

**Coverage note:** the evidence is strongest for breast and brain organ imaging and for the reconstruction/AI methods; it is weakest — effectively absent — for whole-body transmission tomography, which is the idea's defining claim.

---

## 6. Meta-analysis report

### 6.1 Background and scope

Ultrasound computed tomography reconstructs quantitative maps of tissue acoustic properties — chiefly speed of sound and attenuation — from waves transmitted through and scattered by the body, rather than from the back-scattered echoes used in conventional B-mode ultrasound. Full-waveform inversion is the most physically complete reconstruction method: it iteratively updates a model of the tissue until simulated waves match the recorded waves, exploiting the full waveform (phase and amplitude) rather than travel times alone. The "Midjourney" concept is a maximalist instantiation of this paradigm — a very large transmit/receive ring array, an immersion water path, and FWI reconstruction at scale — extended from single organs to the whole body, with a deep-learning segmentation layer on top. This meta-analysis assesses how far the published evidence supports each pillar of that concept.

### 6.2 The reconstruction physics is sound and validated

The narrative's physical premise — that wave shape changes at boundaries in density and stiffness, and that inverting those changes yields a structural map — is precisely full-waveform inversion. Guasch and colleagues demonstrated *in silico* that FWI can produce accurate three-dimensional images of the human brain with sub-millimetre resolution of about 0.9 mm, validating the simulation with *ex vivo* and *in vivo* signal measurements to confirm that real-world signal penetration and signal-to-noise are sufficient [1]. This is the strongest single demonstration that FWI can image complex soft tissue behind bone at the resolution the concept claims, and it directly substantiates the idea's framing that converting waves into images is the central computational task.

For soft tissue without an intervening skull, the breast is the canonical target. Forward-modeling work using anatomically realistic 3D numerical breast phantoms confirms that ring-array FWI can in principle resolve fine breast structure at a 0.2 mm voxel scale, though the study is simulation-only [2]. A physical 3D ring-array breast system reported 0.78 mm spatial resolution and reconstructed simulated lesions as small as 0.3 mm [3]. Together these results establish that the resolution claim — "a fraction of a millimeter" — is attainable for organ-scale soft tissue, in line with the concept's specification.

### 6.3 Acquisition speed and the MRI comparison

The claim of imaging at "nearly a hundred times the speed" of MRI is best understood as an acquisition claim. Guasch et al. report that it takes approximately 2 seconds to acquire a full brain dataset [1], which is far faster than a comparable MRI acquisition. The narrative is careful to attribute the heavy lifting to reconstruction rather than acquisition, which is consistent with the literature.

The qualitative claim that the output "looks a lot like today's MRIs" also has empirical grounding. O'Flynn and colleagues showed that ultrasound-tomography sound-speed measurements of breast density correlate strongly with MRI percent water content, and they advance USCT as a reproducible, radiation-free surrogate [5]. The agreement is at the level of a quantitative tissue property (density/water content), not pixel-for-pixel anatomical equivalence — an important nuance. The honest framing is that USCT produces quantitative maps that correlate with MRI-derived measures, not that it reproduces an MRI.

### 6.4 The whole-body claim collides with ultrasound physics

This is where the concept's evidence base thins to nothing. Every validated soft-tissue success above shares a precondition: the imaged region is either fully immersible and free of air and dense bone (the breast) or is a region where bone aberration can be explicitly modeled and corrected (the skull). Neither condition holds across the torso.

Bone is a severe attenuator. Guo and colleagues showed that ultrasound attenuation through the skull increases with bone thickness and that the bone itself is the dominant source of loss [9]. Beyond attenuation, the skull introduces phase aberration and reverberation; Soulioti et al. found that reverberation significantly degrades transcranial ultrasound imaging, especially for shallower targets, and that both effects must be corrected for usable images [10]. The skull is, relatively speaking, a thin and well-characterized bone. The spine, pelvis, sternum, and ribs are thicker, more heterogeneous, and far less tractable. More fundamentally, ultrasound essentially does not transmit through air; the air-filled lungs and gas-filled bowel present near-total barriers to transmission tomography across the chest and abdomen. No published study demonstrates whole-body transmission USCT, and the physics gives strong reason to expect that the thorax, abdomen, and pelvis are not imageable by transmission FWI in the way the breast is. The realistic envelope of the concept is therefore immersible, air-/bone-free anatomy — a meaningful and valuable scope, but not the whole body.

### 6.5 The data deluge and the computational bottleneck

The narrative's emphasis on terabytes per second and a large compute cluster is well-founded; if anything it understates the reconstruction challenge. 3D FWI of a single numerical breast phantom required approximately 2.3 days on a single NVIDIA A100 GPU [2], and the 3D brain reconstruction required approximately 32 hours on a cluster of 128 CPU-based 24-core nodes

[1]. Whole-body sub-millimeter reconstruction would scale these figures by orders of magnitude. The field is explicit that computational cost is the primary barrier: a 2025 multi-row ring-array study, itself simulation-only, identifies compute as the chief obstacle and proposes multi-GPU systems and dedicated integrated circuits as mitigations [7].

The credible path through this barrier is learned reconstruction. Lozenski et al. showed that a trained learned-FWI network reconstructs each image in approximately 24 milliseconds, against roughly 8.8 minutes per image for traditional FWI — a four-order-of-magnitude speedup, though the study is simulation-only [6]. Jeong et al. demonstrated deep-learning sound-speed estimation evaluated on numerical phantoms and validated on clinical human breast data at 0.25 mm, producing estimates in near real-time once trained [4]. These results make the concept's compute story tractable in principle. The unresolved question is diagnostic-grade accuracy and generalization: learned reconstructions can be fast and visually convincing while being wrong in clinically meaningful ways, and the evidence for robust in-vivo performance across diverse whole-body anatomy does not yet exist.

## 6.6 Hardware scale

The half-million-element ring is  $\sim 500\times$  larger than current practice. Real, working ring arrays operate at 256–1024 elements — a  $1\times 256$  system at 0.78 mm resolution [3] and a 22 cm, 1024-element ring transducer driven by a Verasonics Vantage 256 scanner [8]. Scaling element count by two to three orders of magnitude raises unproven questions of fabrication yield, channel multiplexing, inter-element crosstalk, and per-element calibration. None of this is precluded by physics, but none is demonstrated, and the engineering risk is substantial.

## 6.7 The AI segmentation layer

Of all the components, the deep-learning segmentation overlay is the least speculative; 3D volume segmentation is a mature, widely deployed capability. The risk is not feasibility but framing. Presenting a "crossfade between the raw reconstruction and its AI segmentation" inside a primary diagnostic image blurs the line between measured signal and model inference. In a regulated diagnostic context, a generative or learned overlay that can introduce plausible but non-existent structure is both a patient-safety concern and a regulatory liability. Best practice — and the more defensible regulatory posture — is to keep the quantitative raw reconstruction inspectable beneath any AI layer and to report calibrated uncertainty on the inferred structures.

## 6.8 Synthesis and verdict

The concept is a maximalist but physically literate extrapolation of a real, advancing field. Its core physics, its sub-millimeter resolution claim, its fast acquisition, its MRI-correlation, and its identification of reconstruction as the central task are each supported by peer-reviewed evidence [1][2][3][4][5]. Its two weakest pillars are the two on which the headline depends. First, whole-body transmission tomography is unsupported and physically improbable through air and dense bone; the validated envelope is immersible, air-/bone-free anatomy, with the breast as the

mature beachhead [9][10][4][5]. Second, whole-body sub-millimeter reconstruction is computationally enormous, and the AI reconstruction that would tame it is fast but not yet diagnostically validated in vivo [1][2][6][7]. The honest, high-value version of this idea is "immersible-anatomy quantitative tomography," led by breast imaging, with whole-body as a long-range research aspiration rather than a launch claim. On that reframing, the underlying science is strong and the commercial wedge is real — hence a B+: an excellent, grounded idea gated by hard limits the marketing narrative currently glosses over.

---

## References

- [1] Guasch, L., Calderón Agudo, O., Tang, M.-X., Nachev, P. & Warner, M. "Full-waveform inversion imaging of the human brain." *NPJ Digital Medicine* 3, 28 (2020). doi:10.1038/s41746-020-0240-8 <https://citations.gxl.ai/papers/PMC7060331#L6.L17.L33.L65>
- [2] Li, F., Villa, U., Duric, N. & Anastasio, M. A. "A Forward Model Incorporating Elevation-Focused Transducer Properties for 3-D Full-Waveform Inversion in Ultrasound Computed Tomography." *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* (2024). doi:10.1109/TUFFC.2023.3313549 <https://citations.gxl.ai/papers/PMC10775680#L16.L96.L104>
- [3] Zhang, J. et al. "A High-Resolution 3D Ultrasound Imaging System Oriented towards a Specific Application in Breast Cancer Detection Based on a 1 × 256 Ring Array." *Micromachines* 15, 209 (2024). doi:10.3390/mi15020209 <https://citations.gxl.ai/papers/PMC10891686#L5.L55>
- [4] Jeong, G., Li, F., Mitcham, T. M., Villa, U., Duric, N. & Anastasio, M. A. "Investigating the Use of Traveltime and Reflection Tomography for Deep Learning-Based Sound-Speed Estimation in Ultrasound Computed Tomography." *IEEE Trans. Ultrason. Ferroelectr. Freq. Control* (2024). doi:10.1109/TUFFC.2024.3459391 <https://citations.gxl.ai/papers/PMC11875925#L9.L59.L120>
- [5] O'Flynn, E. A. M. et al. "Ultrasound Tomography Evaluation of Breast Density." *Investigative Radiology* 52, 343–348 (2017). doi:10.1097/RLI.0000000000000347 <https://citations.gxl.ai/papers/PMC5417582#L5.L11>
- [6] Lozenski, L., Wang, H., Li, F., Anastasio, M., Wohlberg, B., Lin, Y. & Villa, U. "Learned Full Waveform Inversion Incorporating Task Information for Ultrasound Computed Tomography." *arXiv* (2023). <https://citations.gxl.ai/papers/PMC10491310#L68.L69>
- [7] Ali, R., Jin, G., Singh, M., Mitcham, T. & Duric, N. "3D Frequency-Domain Full Waveform Inversion for Whole-Breast Imaging With a Multi-Row Ring Array." *IEEE Open J. Ultrason. Ferroelectr. Freq. Control* (2025). doi:10.1109/ojuffc.2025.3570253 <https://citations.gxl.ai/papers/PMC12381738#L21.L28.L30>

[8] Ali, R., Mitcham, T., Owolabi, I., McConnell, S. & Duric, N. "Frequency-differencing strategy to kickstart full-waveform inversion without cycle skipping." *JASA Express Letters* (2025). doi:10.1121/10.0034763 <https://citations.gxl.ai/papers/PMC11734264#L25>

[9] Guo, J., Song, X., Chen, X., Xu, M. & Ming, D. "Mathematical Model of Ultrasound Attenuation With Skull Thickness for Transcranial-Focused Ultrasound." *Frontiers in Neuroscience* 15, 778616 (2022). doi:10.3389/fnins.2021.778616 <https://citations.gxl.ai/papers/PMC8891811#L5.L53>

[10] Soulioti, D. E., Jones, R. M. & Pinton, G. F. "Quantification of the relative contribution of phase aberration and reverberation in transcranial ultrasound imaging: an experimentally calibrated fullwave study in 2D and 3D." *Physics in Medicine and Biology* (2025). doi:10.1088/1361-6560/adf2f3 <https://citations.gxl.ai/papers/PMC12337154#L4.L84>

---

*Grading scale: A+ (excellent) to F (terrible). This evaluation reflects the state of the published evidence as retrieved; the whole-body claim in particular should be re-assessed if peer-reviewed in-vivo whole-body transmission tomography data emerge.*