The Matter Stack

A Layered Physical–Acoustic–Photonic Display: Concept, Build Guide, and Holodeck Roadmap

Design document and step-by-step construction plan — July 2026

1. Executive Summary

This document describes a three-layer tangible display — the Matter Stack — that combines a magnetically sculpted ferrofluid surface, an acoustic levitation layer, and a photonic detail overlay into one registered workspace. Each layer contributes exactly what its physics does best: matter holds graspable shape, sound moves fast and delivers mid-air haptics, and light carries unlimited visual detail. Because ultrasound is a pressure phenomenon and ferrofluid responds only to magnetic fields, the two force media coexist in the same volume without interfering, which is the enabling observation behind the whole design.

The document contains a complete beginner-appropriate build plan: a TinyLev-class acoustic levitator (the proven on-ramp to MATD-style displays), a 16-channel electromagnet grid driving a sealed ferrofluid cell, an integration stage, full parts and cost tables (core build approximately $235–330), safety guidance, and a candid analysis of what does and does not scale on the road to a holodeck.

2. The Matter Stack Concept

The stack is ordered by mass and timescale. The magnetic layer is the slowest and most substantial: a pool of ferrofluid over an addressable electromagnet grid forms persistent, genuinely touchable terrain — the body of a displayed object, changing over seconds and staying put when grasped. The acoustic layer is the dynamic middle: a phased ultrasonic field flies a millimetre-scale bead through the air above the fluid at metres per second, drawing volumetric annotations, cursors, and particles by persistence of vision, while the same field projects focused pressure points onto the user's skin for mid-air haptic feedback. The photonic layer is pure information: an eye-worn display, phone AR view, or projector paints texture, color, labels, and text onto both physical layers at effectively unlimited resolution.

Interaction follows the same division. The user grasps the fluid to move and mold the object; touches the air to operate acoustic controls, feeling a tactile tick at the moment of contact; and reads the photonic layer for data. Registration ties it together: the levitated bead is a bright, trajectory-known optical fiducial, so it can anchor the coordinate frame for the entire stack. A later depth camera closes the loop on the fluid layer, whose shape — unlike the bead's — is not pre-planned and must be sensed.

A note on scope: this build starts with a TinyLev-class levitator, which holds and slowly moves a bead in a standing wave, rather than a full MATD, which requires per-transducer phase control to fly the bead at up to 8.75 m/s. TinyLev is the recognized first step: it teaches transducer handling, drive electronics, and trap physics on the exact same components, and Section 7 describes the upgrade path to a true phased array. Attempting the phased array first is the most common way this project fails for a first-time builder.

3. Build Part A — The Acoustic Layer (TinyLev-class levitator)

This build follows the open-source TinyLev design by Marzo, Barnes and Drinkwater (Review of Scientific Instruments, 2017), with the authoritative step-by-step guide published on Instructables under 'Acoustic Levitator.' An authorized complete kit is sold by Makerfabs and is the recommended route for a first-time builder, since component sourcing and transducer quality are the main failure points. Budget one to three weekends.

A.1 Print the frame

1. Download the TinyLev STL files from the Instructables guide and print the two-bowl frame in PLA. Use a 0.4 mm or 0.6 mm nozzle with a brim and no supports; the frame prints in one piece.
2. Clean the transducer sockets with a small file so each transducer press-fits snugly. Optionally drill the small center hole in each bowl for a camera or needle access later.

A.2 Mark transducer polarity — the critical step

Transducers must all push and pull in phase. Manufacturer polarity markings are unreliable — builders consistently report that a large fraction arrive with inverted markings — so every transducer must be tested individually. One inverted transducer creates a dead spot; several make levitation impossible.

1. Load the polarity-test sketch from the guide onto the Arduino Nano. Connect one wire to pin A0 and one to GND.
2. Open the Arduino Serial Plotter. Touch the two wires to the transducer's legs and gently press the transducer face.
3. If the plotted signal deflects upward (or sits high), mark the leg connected to A0 as positive. If it deflects down, the legs are reversed — mark accordingly.
4. Mark the positive leg of every transducer with a paint pen before any soldering. Do this in one patient sitting; it is tedious and it is also the single highest-leverage hour of the build.

A.3 Populate and solder the arrays

1. Press transducers into both bowls with all marked (positive) legs oriented the same way per bowl, following the guide's orientation diagram. Fix each with a small drop of epoxy or superglue — hot glue is reported to fail.
2. Solder bus wires connecting all positive legs together and all negative legs together within each bowl. Keep joints low-profile so they do not shadow neighboring transducers.

3. Check with the multimeter in continuity mode: every positive leg to the positive bus, every negative leg to the negative bus, and no short between buses.

A.4 Drive electronics

1. Flash the TinyLev sketch to the Arduino Nano. It generates complementary 40 kHz square waves.
2. Wire the Nano's signal pins to the L298N inputs per the guide's schematic, and the L298N outputs to the two arrays — the marked leg of each array to the positive output of its channel.
3. Connect the 12 V supply through the power switch to the L298N. Power the Nano from USB or the L298N's 5 V regulator, but never both simultaneously.
4. Optional but valuable: with a two-channel oscilloscope, confirm both arrays are driven in phase; the guide describes the expected waveforms. If running above 12 V or for very long sessions, upgrade the L298N heatsink.

A.5 First levitation

1. Assemble the two bowls facing each other at the frame's designed separation and power on at 6 V for testing.
2. Place a 1–3 mm polystyrene bead into the central region using a scoop of teabag paper or fine metal mesh on a toothpick — acoustically transparent materials that do not collapse the standing wave the way fingers or tweezers do.
3. The bead should snap into a pressure node and hover. Multiple beads can be stacked in adjacent nodes, evenly spaced. If beads are ejected or oscillate wildly, reduce the supply voltage slightly; if they refuse to trap, try smaller or irregularly shaped beads and re-verify polarity of any suspect transducers.
4. Milestone demo: stable levitation for several minutes, plus slow vertical movement of the bead using the phase-shift feature of the driver code.

4. Build Part B — The Magnetic Layer (ferrofluid over an electromagnet grid)

Budget two weekends. Everything about ferrofluid rewards preparation: it is not dangerous, but it stains skin, clothing, plastics, and porous surfaces permanently. Work gloved, over plastic sheeting, and keep the fluid sealed except when filling the cell.

B.1 Build the sealed fluid cell

1. Prepare a saturated salt solution: dissolve table salt into distilled water until no more dissolves, then let it settle. The dense saline suspends the ferrofluid and stops it from wetting and staining the glass.

2. Fill the flat glass container about two-thirds with the saline, then inject 10–20 ml of ferrofluid with a syringe below the surface. It will pool as a glossy black blob.
3. Seal the lid. From here on, all interaction with the fluid is magnetic, through the glass — which is exactly the point.
4. Hand-test with a strong neodymium magnet under the cell: observe spiking, how sharply features form, and how fast the blob follows a moving magnet. This calibrates your intuition for the grid's achievable 'resolution.'

B.2 Assemble the electromagnet grid

1. Mount the 16 electromagnets face-up in a 4×4 grid at roughly 30 mm spacing on a rigid non-magnetic base (plywood or a 3D-printed jig), directly beneath where the cell will sit. The cell should rest 3–8 mm above the coil faces.
2. Label the coils 1–16 and route each coil's two leads to a screw-terminal strip. Neatness here saves hours of fault-finding later.

B.3 Drive electronics — protect the microcontroller

1. Solder or fit a flyback diode directly across every coil, reverse-biased (cathode/stripe to the coil's positive lead). Coils store energy; when a MOSFET switches off, the collapsing field produces a high-voltage spike that the diode safely recirculates. Skipping even one diode will eventually destroy the drive electronics.
2. Wire each coil through one MOSFET module: coil positive to +12 V, coil negative to the module's drain/load terminal, module ground common with the Arduino Mega ground.
3. Connect each module's signal input to a PWM-capable Mega pin. PWM is what turns 16 on/off magnets into a smooth, shapeable field — cross-fading duty cycles between neighboring coils moves the fluid bump continuously rather than in jumps.
4. Power the coils from the 12 V 10 A supply on a separate rail from the Arduino's supply, grounds tied together at one point.
5. In firmware, cap the maximum duty cycle at roughly 60–70 percent and consider a software thermal budget (for example, limiting how long any coil may exceed 50 percent). Holding electromagnets run hot at full duty; heat is this subsystem's main reliability enemy.

B.4 First terrain

1. Place the sealed cell on the grid. Energize a single coil at 50 percent duty: a smooth bump of fluid should rise above it.
2. Write the traveling-bump demo: cross-fade PWM between adjacent coils so the bump orbits the grid in a circle. This is the subsystem's milestone — it proves smooth spatial control, not just on/off spikes.

3. Characterize the response: how fast can the bump travel before it tears or lags? That measured speed becomes the magnetic layer's specification in the integrated system.

5. Build Part C — Integration and the Photonic Layer

1. Print or build a gantry that suspends the TinyLev horizontally 10–15 cm above the fluid cell, so the levitation zone sits over the terrain. Piezoelectric transducers are non-magnetic and the ferrofluid is deaf to ultrasound at these intensities, so the layers do not fight — but keep the Arduinos and driver boards ~15 cm from the strongest magnets as cheap insurance.
2. Two-layer milestone: a bead levitating above while a fluid bump tracks beneath it, both driven from one host computer sending commands to the two Arduinos over USB serial.
3. Add the photonic layer the inexpensive way first: a mini projector aimed at the cell projection-maps color and labels onto the fluid — the glossy black surface takes projected light strikingly well — or a phone AR app overlays graphics using a printed marker on the frame.
4. Full-stack milestone: one script animates all three layers in sync — for example, terrain rises, the bead circles the new peak, and the projector paints elevation colors on the fluid.
5. Later upgrade: add the depth camera above the cell so the system senses the fluid's actual shape and registers the photonic layer to measured, not commanded, geometry — the first step toward the closed-loop registration this architecture ultimately depends on.

6. Parts and Cost Lists

6.1 Acoustic layer

Item	Qty	Est. cost	Notes
10 mm 40 kHz ultrasonic transducers	76	$45–60	Buy ~76 to have spares; the type used in car parking sensors. Bulk packs are cheapest.
Arduino Nano (or clone)	1	$5–10	Generates the 40 kHz drive signal.
L298N dual motor driver board	1	$5–8	Acts as the amplifier for both transducer arrays.
12 V DC power supply (2 A) + switch	1	$10–15	Run at 6–9 V for testing; 12 V maximum.
3D-printed TinyLev frame (two bowls)	1	$5	Free STL files from the Instructables guide; ~0.4 mm nozzle, brim, no supports.
Hook-up wire, solder, heat-shrink, glue	—	$10	Hot glue is reported unreliable for transducer mounting; use epoxy or superglue.
Expanded polystyrene beads, 1–3 mm	1 bag	$5	The levitated 'pixels.' Must be under ~3 mm (half a wavelength).
Teabag paper or fine metal mesh + toothpick	—	$2	Acoustically transparent scoop for placing beads; fingers disturb the field.

Subtotal (scratch build)		**$85–115**	**Authorized Makerfabs complete kit: roughly $100–130 and removes all sourcing risk.**

6.2 Magnetic layer

Item	Qty	Est. cost	Notes
12 V electromagnets, 25 mm barrel	16	$35–50	Holding-type electromagnets; 4×4 grid at ~30 mm spacing.
Logic-level MOSFET driver modules (IRLZ44N or similar)	16	$15–25	One channel per coil; modules are far more beginner-friendly than bare MOSFETs.
Flyback diodes (1N4007 or 1N5819)	20	$5	One across every coil, reverse-biased. Non-negotiable — omitting them destroys the microcontroller.
Arduino Mega 2560 (or clone)	1	$15–25	Enough PWM-capable pins for 16 channels.
12 V 10 A power supply	1	$20–25	Coils draw ~0.3–0.5 A each; budget headroom and never run all 16 at full duty.
Ferrofluid (e.g., Ferrotec EFH-1), 50 ml	1	$25–35	Quality fluid resists clumping. Cheap fluid stains worse and performs worse.
Flat glass container / large petri dish + lid	1	$10	The fluid cell. Glass, not plastic — ferrofluid stains plastic permanently.
Distilled water + table salt	—	$3	Saturated salt solution suspends the ferrofluid and prevents glass staining.
Perfboard, screw terminals, 18 AWG wire	—	$10–15	Power distribution for 16 coil channels.
Nitrile gloves, plastic sheeting, paper towels	—	$8	Ferrofluid stains skin, clothes, and furniture. Treat it like ink that never dries.
Subtotal		**$145–200**	

6.3 Integration and tools

Item	Qty	Est. cost	Notes
3D-printed gantry / frame	1	$5–10	Holds the acoustic levitator 10–15 cm above the fluid cell.
Mini LED projector (optional)	1	$50–80	Projection-maps color and labels onto the glossy black fluid; a phone AR app is the free alternative.
USB depth camera (optional, later)	1	$60–150	Lets the system sense the fluid's actual shape for closed-loop overlay registration.
Tools if not owned: soldering iron, multimeter	—	$30–45	A multimeter is mandatory for transducer polarity testing.

Project total (core build, no optionals)		**$235–330**	**Assumes scratch acoustic build and no projector or depth camera.**

7. Upgrade Path: From TinyLev to a True MATD

TinyLev holds the bead in a standing wave and can move it slowly along one axis by phase shifting. A Multimodal Acoustic Trap Display requires individual phase control of every transducer so the trap can be steered anywhere in the volume thousands of times per second, flying the bead at up to 8.75 m/s while synchronized RGB LEDs color it. The established open route is the Ultraino platform and its successors (such as OpenMPD) from the same research group: a flat array of 64 or more transducers driven by an FPGA or a microcontroller with dedicated driver boards providing per-channel phase resolution. Expect $300–600 in parts, surface-mount soldering or pre-fabricated driver boards, and real firmware work. Everything learned in Build Part A — polarity discipline, transducer handling, drive electronics, trap behavior — transfers directly. Amplitude modulation on the same array then adds the MATD's remaining modalities: mid-air tactile pressure points and audible sound from the display itself.

8. Safety

Ultrasound: 40 kHz is inaudible but intense near the arrays. Do not place your ear close to a running levitator, limit continuous close-range hand exposure, and keep pets — whose hearing extends far higher than ours — out of the room during long sessions. Electrically the acoustic build is low-voltage and benign.

Magnetics: keep strong magnets and energized coils away from pacemakers and implanted devices, credit cards, and mechanical watches. Electromagnets run hot; mount them on a non-flammable base, enforce duty-cycle limits in software, and power down when unattended.

Ferrofluid: not acutely toxic, but treat it as a permanent stain hazard — gloves, sheeting, sealed cell. Never let it contact bare neodymium magnets directly; it is nearly impossible to remove and the suspension degrades.

Power: the 12 V 10 A supply can deliver enough current to melt thin wiring. Fuse the coil rail, use 18 AWG or heavier for power distribution, and double-check polarity before first power-up.

9. Scaling Analysis and Holodeck Roadmap

9.1 What does not scale — the physics walls

Force does not project. Both force media in the stack decay viciously with distance. Magnetic field strength falls off roughly with the cube of distance from a source, so a coil that sculpts fluid across a few centimetres would need MRI-class power to act across metres — at which point the field itself becomes the hazard. Acoustic traps degrade similarly: trap stiffness collapses with range, and the focal intensities already required at desktop scale would, scaled to room volume, raise genuine heating and hearing-safety concerns despite being ultrasonic.

The speed of sound is a latency floor. Sound crosses a 3.5-metre room in roughly ten milliseconds. Wall-mounted arrays therefore cannot maintain phase-precise, fast-reacting traps deep in the room; responsiveness degrades with distance and no engineering removes this limit.

The fast particle becomes a projectile. A millimetre bead at 8.75 m/s in a small guarded volume is harmless; room-scale imagery implies faster or heavier particles sharing space with human eyes. Safety architecture, not resolution, becomes the binding constraint.

Ferrofluid does not tile. A room-scale fluid surface means large volumes of expensive, staining liquid plus containment. The magnetic layer scales as furniture — tables, panels, consoles — not as floors and walls.

9.2 What scales beautifully

The photonic layer. Optics, displays, and compute ride industrial learning curves, while force media obey inverse-cube laws. This asymmetry is the central architectural fact of any holodeck: force will always be local and expensive; light will always be global and cheap. Room-scale systems must therefore concentrate force where hands actually are and let photons carry everything else.

9.3 The holodeck roadmap

Stage 1 — Zones, not volumes. Embed matter-stack workstations (exactly this build) at the loci of interaction: tabletops, wall niches, doorframe consoles. The photonic layer stitches the zones into one continuous scene. The room is an illusion of light with islands of genuine matter precisely where you touch it.

Stage 2 — Replace the bead with a swarm. At room scale, the flying bead's successor is a swarm of small LED microdrones: the same persistence-of-vision principle, but self-powered, so no force must be projected from the walls at all. This sidesteps the force-projection wall entirely and inherits the trajectory-broadcast registration trick — a drone's planned path, like the bead's, costs zero latency. Multi-particle acoustic research (electrostatic anti-collapse, learned trajectory stabilization) meanwhile grows the desktop bead count.

Stage 3 — Force that travels with the user. Wearable haptics — ultrasonic wristbands, gloves, ring transducers — keep the force source permanently centimetres from skin, the only place force is cheap. The room tells the wearable what to render; the wearable delivers it locally.

Stage 4 — Encounter surfaces. Mobile robots quietly pre-position physical props and actuated surfaces where the user is about to touch the virtual world, generalizing the magnetic layer's role: persistent matter, delivered just-in-time instead of everywhere-at-once.

9.4 Why this desktop build is the right first step

Not because it scales — it does not — but because its architecture is the one that scales: a persistent matter layer, a fast dynamic layer, and a photonic detail layer, mutually registered in real time around a shared fiducial. Every holodeck stage above is that same three-layer pattern with substitutions. The registration, synchronization, and content-splitting problems this build forces you to solve at a bench are the holodeck's actual research problems, purchasable today for a few hundred dollars.

10. References

Hirayama, R., Martinez Plasencia, D., Masuda, N., & Subramanian, S. (2019). A volumetric display for visual, tactile and audio presentation using acoustic trapping. Nature, 575, 320–323.

Marzo, A., Barnes, A., & Drinkwater, B. W. (2017). TinyLev: A multi-emitter single-axis acoustic levitator. Review of Scientific Instruments, 88, 085105. Build guide: instructables.com/Acoustic-Levitator (UpnaLab).

Marzo, A., et al. Ultraino: an open phased-array platform for airborne ultrasound; successor software platform OpenMPD.

Shi, S., Hübl, M. C., Grosjean, G., Goodrich, C. P., & Waitukaitis, S. (2025). Electrostatics overcome acoustic collapse to assemble, adapt, and activate levitated matter. PNAS, 122(50).

StableLev: Data-Driven Stability Enhancement for Multi-Particle Acoustic Levitation. CHI 2024.