

LODESTONE

A low-cost instrument that finds metal objects buried in the ground and tells what they are, how large they are and how deep they lie.

Design complete · seeking grant funding · full report as PDF (74 pages)

What it is

An ordinary metal detector makes a sound when there is metal under its coil. It cannot say whether the metal is a coin, a nail or a piece of foil, how big it is, or how deep it is buried. To find out, someone has to dig.

LODESTONE is a research instrument designed to answer those questions before anyone digs. As the operator walks across the ground, it measures every buried metal object from many positions and reports on a map, for each object, the most likely type of metal, its approximate size and shape, its depth and its exact position.

Why it is needed

Most of what metal detectors find is scrap. Each signal has to be checked by hand, because the detector cannot tell a harmless object from one that matters. This makes surveys slow and expensive wherever buried metal has to be found:

- **clearance of mines and unexploded ordnance**, where most of the working time goes into digging up harmless fragments (LODESTONE is a research platform for this problem, not a clearance tool);
- **archaeology**, where every unnecessary excavation damages a site;
- **construction and utilities**, where pipes, cables and reinforcement must be located before digging;
- **environmental surveys** of contaminated or former industrial land.

Instruments that can classify buried metal exist, but they cost more than 100,000 USD and their designs and data are closed. LODESTONE aims to reach a useful part of that capability with components costing about 2,354 USD, and to publish the design, the software and a labelled dataset openly so that others can build on it.

How it works




Figure I. How LODESTONE works. The numbered markers correspond to the four steps below. Lettered objects A to D appear on the result map with their estimated type and depth.

- 1 The sensor sends a magnetic field into the ground.** Coils in the sensor head carry an alternating current at twelve frequencies at the same time, from 188 Hz to 18.75 kHz. The resulting magnetic field passes through soil and reaches buried objects.
- 2 Buried metal answers.** The field makes small electric currents flow in any metal object, and these currents produce a weak field of their own, which receive coils in the same head pick up. How strongly the object answers at each of the twelve frequencies depends on its metal, size and shape. This pattern works like a fingerprint; a stone gives no answer at all.
- 3 Depth and position are measured.** A small radar in the sensor head measures how deep objects lie. Satellite positioning, motion sensors and a camera record where the sensor head is, to about 2 cm, at every moment. The same object is therefore measured from many known positions as the operator walks past it.

- 4** A computer builds a map of identified objects. For each object it fits a physical model to all measurements, removes the effect of depth and orientation, and compares the result with a library of known objects. The output is a map showing each object with its likely type, size and depth.

about 40 cm detection range for a coin-sized object, calculated	2 cm position accuracy of every measurement	12 frequencies measured at once, 188 Hz to 18.75 kHz
3 × 3 transmit × receive coils: the full target tensor	2,354 USD equipment, including import and contingency	102,936 USD grant request for 12 months, team of four

Technical summary

The induction part uses a six-input, eight-output audio codec (CS42448) to drive three orthogonal transmit coils and read three receive coils, so that the full magnetic polarizability tensor of a target can be estimated. The radar is a stepped-frequency ground-penetrating radar from 0.5 to 3 GHz on a software-defined radio (ADALM-Pluto). RTK satellite positioning, inertial sensors and a ground-facing camera are fused in a Kalman filter, and all data share one GPS time base through the receiver's pulse-per-second signal. Every design value in the report was derived from physical models before construction; the noise budget predicts a detection range of about 40 cm for coin-sized objects, in line with commercial detectors.




Figure II. System overview: the three sensing subsystems, the codec and the processing computer. Details in Chapters 4 to 7.




Figure III. Why the frequency pattern identifies the metal: computed response of four metal spheres of 1 cm diameter (exact solution, Chapter 3). Each metal answers most strongly at a different frequency. The shaded band and the ticks show the twelve measurement frequencies.

Plan

The project runs for twelve months with a team of four in seven work packages. Each technical package ends with a measurable milestone, followed by a decision gate with a prepared fallback




Figure IV. Work packages, coloured by the role that leads them, with milestones. M1 measured spectra agree with theory; M2 sensor position known to 2 cm; M3 classifier validated on objects not seen in training; M4 radar locates a can at 35 cm; M5 fused field demonstration and public release.

Budget and funding

The request of 102,936 USD covers four full-time positions for twelve months, equipment, running costs, dissemination and contingency. The main target is a NATO Science for Peace and Security Multi-Year Project with a partner group in a NATO member country, complemented by national research funding in Armenia (Chapter 11).




Figure V. Budget request by cost category, USD. Equipment is a small share because the platform is built from inexpensive components; most of the cost is the work of the team.

About this document

Report TR-2026-01, revision 3

Date	8 October 2026
Status	Design complete; construction starts when funding is in place
Project	12 months, team of four, 7 work packages, 5 milestones
Funding sought	102,936 USD
Download	Full report as PDF (74 pages)

Contents

1	Introduction	page 2
2	State of the art	page 3
3	Physics of detection and identification	page 4
4	System architecture	page 5
5	Induction subsystem	page 6
6	Radar subsystem	page 7
7	Positioning and timing	page 8
8	Signal processing and machine learning	page 9
9	Work plan	page 10
10	Team and responsibilities	page 11
11	Budget and funding	page 12
12	Risks, safety and quality	page 13
	Appendices	page 14

Use the page numbers at the bottom of each page, or the arrow keys, to move through the report.

Introduction

Why the project exists, the questions it answers and what it will deliver.

1.1 Motivation

A metal detector tells its user that something conductive is under the coil. It does not say reliably what the object is, how large it is, or how deep it lies. Commercial detectors reduce the received signal to a single "target identity" number, and experienced users still dig up large numbers of bottle caps, foil and nails for every object of interest. In survey and clearance work the same limitation appears as a high false-alarm rate: each excavation of harmless scrap costs time.

The physics does not impose this limitation. A buried metal object responds to an alternating magnetic field in a way that depends on its size, shape, orientation, conductivity and magnetic permeability. With enough measurements, taken at several frequencies, from several directions and at known positions, most of these properties can be separated. Military programmes for unexploded ordnance have shown this with vehicle-mounted instruments costing well above one hundred thousand dollars (Chapter 2).

LODESTONE asks whether the same class of measurement can be made with inexpensive, openly available components: a multichannel audio codec, a software-defined radio, and consumer positioning sensors. The project name refers to magnetite, the naturally magnetised stone used in the first compasses. The project combines magnetic sensing with precise navigation.

1.2 Research questions

The work started from a series of practical questions. They are stated here in the form in which the report answers them.

1. **RQ1.** Can a software-defined radio be used to detect buried metal, and in which frequency range does it help? (Chapters 3 and 6.)
2. **RQ2.** Is a measurement across many frequencies more informative than a measurement at one frequency, and what exactly does it add? (Chapter 3.)
3. **RQ3.** Can a trained model predict the type, size and location of a buried object, and how many transmit and receive coils does that require? (Chapters 5 and 8.)
4. **RQ4.** How much does precise knowledge of the sensor's motion (inertial sensors, camera, RTK satellite positioning) improve the result? (Chapter 7.)

1.3 Approach

The platform has three measurement subsystems that share one time base:

- an **induction subsystem** operating between 188 Hz and 18.75 kHz, with three transmit and three receive coils driven and read by a CS42448 codec (six 24-bit inputs, eight outputs, one sample clock);
- a **radar subsystem**, a stepped-frequency continuous-wave ground-penetrating radar built on an ADALM-Pluto software-defined radio (AD9363), sweeping 0.5 to 3 GHz;
- a **positioning subsystem** combining RTK satellite positioning, an inertial measurement unit, a ground-facing camera and a height sensor.

The pulse-per-second output of the satellite receiver is recorded on a spare codec input and shown to the camera through an LED. Every sample of every sensor can therefore be placed on the same GPS time line.

Every design value in this report was calculated before any hardware is built. Where an exact solution exists (for example the response of a conducting, permeable sphere), the figures are computed from it directly in the page. The last check of the design was a noise budget. It predicts a detection range for coin-sized targets of about 40 cm, which agrees with the performance of commercial detectors.

1.4 Expected contributions

1. An open hardware and software design that recovers the full magnetic polarizability tensor of buried targets for a hardware cost of about two thousand dollars.
2. A method for recording induction, radar and positioning data on a common GPS time base with single-sample precision.
3. A labelled, openly published dataset of induction and radar measurements of known objects at known positions and depths.
4. A physics-based classification pipeline, evaluated on objects that were never seen during training.

1.5 Scope and limitations

LODESTONE is a research instrument. It is **not** a demining device, and it will not be certified or used as one. Field work is restricted to prepared test beds and to land that is known to be free of ordnance (Chapter 12). The project does not address detection of non-metallic objects by induction, and it treats the radar only as a supporting sensor for depth and shape.

1.6 Status of this document

This is revision 3 of the report. The design work is complete; construction will start when funding is in place (Chapter 11). The one piece of hardware already characterised is the CS42448 codec board, and even there the noise and full-scale values are derived from the datasheet and must be confirmed on the bench (task T1.3). Measured results will replace calculated ones as the work packages in Chapter 9 are completed.

1.7 Structure of the report

Chapter 2 reviews related work. Chapter 3 sets out the physics that the design depends on. Chapters 4 to 8 describe the system: architecture, the induction, radar and positioning subsystems, and data processing. Chapters 9 to 12 form the research plan: work packages and schedule, team and responsibilities, budget, and risks. The appendices collect key numbers, code, references and photographs.

State of the art

Published work on each part of the platform, and the gap this project addresses.

A literature and product search was carried out in September 2026 for each component of the planned platform. Each component has been demonstrated before. We found no published system that combines them, and none that is openly documented or built at low cost.

2.1 Software-defined ground-penetrating radar

Stepped-frequency continuous-wave (SFCW) radar implemented on software-defined radios is an established academic approach. A Georgia Tech thesis and the corresponding IEEE conference paper [R1, R2] implemented SFCW GPR on SDR hardware. A 2025 GNU Radio Conference paper [R3] used commercial off-the-shelf SDRs, and a hobby project imaged the subsurface with two HackRF units [R4]. These works confirm that the radar architecture in Chapter 6 is feasible with SDR hardware. None of them identifies the material of a target.

2.2 Multi-axis electromagnetic induction for unexploded ordnance

The closest related technology is the family of advanced induction sensors developed for classifying unexploded ordnance in the United States. MetalMapper [R5] uses three orthogonal transmit loops and multi-axis receivers. TEMTADS and ALLTEM follow similar principles [R6]. These systems invert a dipole model to obtain the magnetic polarizability tensor of each target and classify on it.

Reported performance at live sites is about 90% probability of detection at about 10% false positives. The mathematics of polarizability-tensor spectra as object signatures is an active research topic [R7, R8]. These instruments work in the time domain and are vehicle-towed or cued systems in the price range above 100,000 USD.

2.3 Combined induction and radar detectors

Hand-held detectors that combine induction with radar are standard equipment in humanitarian demining. The US Army's AN/PSS-14 (HSTAMIDS) has been standard issue since 2006 [R9]. The Vallon-Cobham Minehound and Minelab's MDS-10 and MDS-20 [R10, R11] pair simultaneous multi-frequency induction with ultra-wideband radar. Research on combining electromagnetic induction spectroscopy with GPR for anti-personnel mines continues [R12]. The main result of this line of work is that the two sensors fail on different kinds of clutter, so fusion lowers the false-alarm rate. These products are closed: raw data and tensor information are not available to the user. Minelab's

consumer detectors (Multi-IQ) also transmit several frequencies simultaneously, which is the excitation principle chosen for LODESTONE.

2.4 Machine learning on detector data

Recent work applies deep learning to induction data for classifying mines and small metal content in mineralised soil [R13]. It applies one-dimensional convolutional networks to time-domain polarizability features [R14] and machine learning to pulse-induction data [R15]. Methods that start from physically meaningful features, such as the polarizability tensor, are common, and the pipeline in Chapter 8 follows that direction. A limitation mentioned in almost all of these works is the small size and closed nature of the available labelled datasets.

2.5 Comparison

Table 2.1. Capabilities of related systems compared with the planned platform.

Capability	MetalMapper / TEMTADS	AN/PSS-14, MDS-20	SDR GPR (academic)	LODESTONE
Full-tensor induction (3 TX × 3 RX)	Yes	No	No	Yes
Multi- frequency induction spectra	Yes, time domain	Yes, closed	No	Yes, 12 tones
Radar	No	Yes	Yes	Yes
Centimetre- level positioning	GPS on vehicle	No	No	RTK, IMU, camera
Open hardware and data	No	No	Publications only	Yes
Approximate cost class	above 100,000 USD	20,000 to 50,000 USD	1,000 to 5,000 USD	about 2,300 USD

2.6 Position of this project

We did not find a platform that combines (1) multi-frequency, full-tensor induction on an audio codec, (2) SDR radar on the same time line, (3) centimetre positioning with hardware time synchronisation, (4) classification on orientation-independent tensor spectra, and (5) an open labelled dataset. The individual components are each supported by the literature above. The contribution of LODESTONE is their integration at low cost with full documentation.

From the related work we adopt the dipole-model inversion and tensor-based classification of the ordnance programmes, simultaneous multi-frequency excitation, the principle of fusing induction with radar to reduce false alarms, and the spectral-tensor mathematics of [R7, R8]. The dataset is designed into the project from the first work package in response to the data shortage noted above.

The search covered English-language publications and product literature. Patents and literature in other languages have not yet been reviewed; this is listed as task T6.5.

Physics of detection and identification

How metal targets respond to alternating magnetic fields, and what a frequency sweep adds.

This chapter summarises the physical principles that the design relies on. All curves are computed in the browser from the equations given in the text.

3.1 Electromagnetic induction

A transmit coil carrying alternating current produces a magnetic field. The field induces eddy currents in any conductor nearby, and the eddy currents produce a secondary field that a receive coil detects. At the frequencies used here the wavelength is kilometres long, so all coupling takes place in the magnetic near field; no radio waves are involved.

Table 3.1. Principal types of metal detector.

Type	Frequency	Principle	Typical strength
VLF induction balance	3 to 30 kHz	Continuous tone; amplitude and phase of the secondary field are measured	Discrimination between metals
Pulse induction	pulsed, μs decay	Transmitter switched off; decay of eddy currents is timed	Depth, tolerance of mineralised ground
Beat frequency oscillator	about 100 kHz	Metal changes coil inductance and shifts an oscillator	Simplicity

Low frequencies are used because a field entering a conductor decays over the skin depth

$$\delta = \sqrt{\frac{2\rho}{\omega\mu}},$$

which in copper is about 0.65 mm at 10 kHz and about 2 μm at 1 GHz. Soil is a weak conductor and follows the same law, so induction at kilohertz frequencies penetrates the ground while radio frequencies are attenuated near the surface. A VLF detector already has the structure of a software-defined receiver. It transmits a tone, multiplies the received signal by the sine and cosine of the reference (I/Q demodulation), and reads amplitude and phase. The design in Chapter 5 implements this with an audio codec.

3.2 Response of a conducting target

A simple target, such as a ring, can be modelled as a loop with inductance L_t and resistance R_t . Its response to a field at angular frequency ω is

$$H(\omega) = \frac{j\omega\tau}{1 + j\omega\tau}, \quad \tau = \frac{L_t}{R_t} \sim \mu\sigma\ell^2,$$

where ℓ is a characteristic size. This is a relaxation, not a resonance: the magnitude increases monotonically with frequency. The quadrature (90°) component has a maximum at $\omega = 1/\tau$, and this maximum is how the time constant is measured in practice.




Figure 3.1. Response of a single-time-constant target. The magnitude (dashed) has no maximum; the quadrature component peaks at $\omega\tau = 1$.

Because $\tau \sim \mu\sigma\ell^2$ contains both material and size, a large aluminium object and a smaller copper object can have the same time constant. This is the basic ambiguity of metal identification by induction.

3.3 The conducting, permeable sphere

For a solid sphere of radius a , conductivity σ and relative permeability μ_r in a uniform field, the induced dipole moment is known exactly [R16]:

$$\frac{M}{2\pi a^3} = \frac{2\mu_r (x \coth x - 1) - (1 + x^2 - x \coth x)}{\mu_r (x \coth x - 1) + (1 + x^2 - x \coth x)}, \quad x = a\sqrt{j\omega\mu_0\mu_r\sigma}.$$

The expression was checked against three limits: at zero frequency it gives the magnetostatic value $2(\mu_r - 1)/(\mu_r + 2)$; at high frequency it tends to -1 (a perfect diamagnet); and at low frequency

for a non-magnetic sphere it reduces to the eddy-current result $-x^2/10$. All three agree to numerical precision.




Figure 3.2. In-phase and quadrature response of five spheres, normalised by volume, from 10 Hz to 1 MHz. The shaded band is the range covered by the codec at 48 kHz sampling. The dashed aluminium curve (radius 6.3 mm) coincides with the copper curve (radius 5 mm).

For non-magnetic spheres the quadrature peak lies at $f_{\text{peak}} \approx 1.85/(\mu_0 \sigma a^2)$. This is about 1.0 kHz for a copper sphere of 5 mm radius, 1.6 kHz for aluminium and 41 kHz for stainless steel. An aluminium sphere of 6.3 mm radius has the same σa^2 as the 5 mm copper sphere, so their normalised spectra are identical; only the absolute amplitude differs, by a factor of $(6.3/5)^3 \approx 2$. A steel sphere ($\mu_r \approx 100$) remains on the ferrous side of the response over the whole audio band; its in-phase component changes sign only near 280 kHz.

3.4 One frequency compared with many

A detector working at one frequency measures one complex number per position: an amplitude, which combines size with depth, and a phase. In the complex plane every target traces a curve as frequency increases, and a single-frequency measurement is one point on that curve.




Figure 3.3. Response loci in the complex plane. All non-magnetic solid spheres lie on one common curve; conductivity and size only determine where on it a given frequency falls. A ring follows a semicircle. Points mark a 10 kHz measurement.

A frequency sweep records the whole curve. Its shape carries information that a single point cannot. The clearest example is geometry. A ring has one eddy-current path and one time constant. A solid object has many superposed modes, which widens its response.




Figure 3.4. Quadrature response of a ring and of a solid sphere, each shifted so that its peak is at the centre. The sphere's response is lower and about 40% wider (1.56 against 1.13 decades at half height).

Table 3.2. What one frequency and a sweep can determine.

Question	One frequency	Sweep
Ferrous or non-ferrous	Yes (phase)	Yes
Product σa^2	Only for simple shapes	Yes, by fitting the full curve
Ring, solid or flat object	No	Yes, from curve width and shape
Mineralised ground	Suppressed at one phase, together with targets at that phase	Separable by spectral shape
Two objects under the coil	No	Partly
Copper or aluminium, size unknown	No	No; requires size from geometry (Chapter 7)

Magnetically viscous soil, which contains very fine magnetic grains, shows a quadrature response that is nearly constant with frequency. Its in-phase response falls roughly linearly with the logarithm of frequency. No metallic target behaves like this, so with a sweep the ground response can be fitted and removed instead of being suppressed at one phase.

3.5 Dipole model and distance dependence

A small target in the field of a transmit coil becomes a magnetic dipole. Its strength and direction are set by a symmetric 3×3 polarizability tensor $\mathbf{M}(\omega)$. The voltage induced in a receive coil is

$$V(\omega) \propto \mathbf{h}_{\text{rx}}^T(\mathbf{r}) \mathbf{M}(\omega) \mathbf{h}_{\text{tx}}(\mathbf{r}),$$

where $\mathbf{h}_{\text{tx}}$ is the transmitter's field at the target and $\mathbf{h}_{\text{rx}}$ the receiver's sensitivity there. Both fall off as $1/r^3$, so for a small target

$$|V| \propto \frac{|\mathbf{M}|}{r^6}.$$

Doubling the depth reduces the signal by a factor of 64. In the target's own axes the tensor is diagonal:

$$\mathbf{M}(\omega) = \mathbf{R}^T \text{diag}(\beta_1(\omega), \beta_2(\omega), \beta_3(\omega)) \mathbf{R}.$$

The three curves $\beta_i(\omega)$ describe the object independently of its orientation. A sphere gives three equal curves; elongated and flat objects give unequal ones. Recovering the β_i is the central measurement task of the project. Chapter 5 shows which coil arrangement makes it possible.

System architecture

The three sensing subsystems, their interfaces and the common time base.

This chapter gives an overview of the platform before the subsystems are described in detail in Chapters 5 to 7.

4.1 Overview




Figure 4.1. System overview. Three sensing subsystems feed one processing computer. The GNSS pulse-per-second signal is recorded on a codec input and shown to the camera, which places induction samples, video frames and position fixes on one time line.

4.2 Design principles

- One clock per subsystem, one time base overall.** All induction channels are sampled by one codec with one master clock, so their relative phases are exact. The radar derives its time from the host computer, and the GNSS pulse ties both to GPS time.
- Calculate before building.** Every component value is derived from a model, and each work package starts by checking one prediction against a measurement (Chapter 9).

3. **Physics before learning.** Known relations such as the $1/r^6$ distance law and tensor rotation are handled by an explicit model. Machine learning is applied only to the orientation-independent quantities that remain (Chapter 8).
4. **Open and reproducible.** Designs, code, data and this report are published under open licences at the end of the project.

4.3 Subsystems

Table 4.1. Subsystems, main hardware and procurement status.

Subsystem	Function	Main hardware	Status
Induction	Multi-frequency, full-tensor measurement of the target response	CS42448 codec board, 3 TX and 3 RX coils, class-D drivers, ×33 preamplifiers	Codec owned; rest to build
Radar	Depth and shape by stepped-frequency GPR, 0.5 to 3 GHz	ADALM-Pluto (AD9363), RF amplifiers, two Vivaldi antennas	Pluto and amplifiers owned; antennas to build
Positioning	Sensor pose at every sample, centimetre level	2 × RTK GNSS receivers, IMU, magnetometer, global-shutter camera, time-of-flight ranger	To buy
Host	Streams codec data to the PC	Teensy 4.1 (baseline) or Raspberry Pi	To decide in T1.2
Processing	Demodulation, pose estimation, inversion, classification, fusion	Laptop PC, Python	Available

4.4 Interfaces and data rates

Table 4.2. Data interfaces.

Link	Interface	Rate	Notes
Codec to host	TDM serial audio	48 kHz × 6 ch	Master clock shared by all channels
Host to PC	USB	9.2 Mbit/s	6 channels × 32 bit × 48 kHz
Radar to PC	USB 2.0 (IIO)	up to about 4 MS/s IQ	One capture per frequency step
GNSS to PC	UART / USB (UBX)	10 Hz fixes	PPS wired to codec input 5
IMU to PC	I ² C or SPI through host	200 Hz	Accelerometer, gyroscope, magnetometer
Camera to PC	USB	60 frames/s	PPS-driven LED in field of view

4.5 Host decision

Two hosts are possible for the codec board. A Teensy 4.1 drives the CS42448 over TDM with existing library support; because its standard USB audio class is limited to two channels, the six input channels are streamed as a raw serial stream at 9.2 Mbit/s, which is well within its USB 2.0 High Speed port. A Raspberry Pi supports the CS42448 in its mainline audio drivers and presents it as an ordinary multichannel sound card, which simplifies the software. The choice will be made in task T1.2, after both have been tested with the actual board. The rest of the design does not depend on it.

Induction subsystem

Coils, drivers, receive chain, digitiser and excitation for full-tensor measurement.

The induction subsystem measures the complex response of buried targets at twelve frequencies, from up to three transmit directions and at three receivers. It is the part of the platform that identifies what a target is made of and what shape it has.

5.1 Requirements

- Frequency coverage that contains the quadrature peaks of coin-sized non-ferrous targets (about 1 to 6 kHz for 2 to 5 mm copper and aluminium).
- Phase-coherent sampling of all transmit current references and all receivers.
- Input-referred noise low enough that the ground, not the electronics, limits detection range.
- A coil arrangement that allows the full polarizability tensor to be recovered.

5.2 Why no radio-frequency front end is needed

Radio-frequency design rules (transmission lines, impedance matching, mixers) apply when conductors are comparable to the wavelength, usually taken as longer than $\lambda/10$. A 2 m coil cable reaches that length at 15 MHz. In the induction band, below 19 kHz, $\lambda/10$ is more than 1.5 km. All connections are therefore ordinary lumped circuits. The analog-to-digital converter samples the band directly, so no local oscillator or intermediate frequency is needed. The first amplifier is specified by its input noise density, as in audio design.




Figure 5.1. Wire length at which transmission-line effects begin ($\lambda/10$) as a function of frequency. The induction band lies four decades below the limit for a 2 m cable; the radar band lies above it.

5.3 Transmit chain

Each transmit coil is driven by a class-D audio power amplifier through a series resistor, which keeps the load impedance above the amplifier's minimum rating at every tone. The coil is not tuned with a capacitor: a resonant coil would emphasise one frequency and suppress the others. A small sense resistor in series with each coil is read by a codec input. Dividing every receiver signal by this current reference removes gain and phase drift in the digital-to-analog converter and the amplifier. With 20 turns, 12 cm radius and 1 A, the coil's magnetic moment is $0.9 \text{ A}\cdot\text{m}^2$.

5.4 Receive coils

Table 5.1. Reference receive coil.

Parameter	Value	Basis
Turns, diameter	100, 12 cm	Trade-off between signal and self-resonance
Inductance	4.58 mH	$L = \mu_0 N^2 r [\ln(8r/a) - 2]$, wire diameter 0.3 mm
Winding resistance	12.8 Ω	37.7 m of wire; thermal noise 0.46 nV/ $\sqrt{\text{Hz}}$
Self-resonance	429 kHz	Assumed winding capacitance 30 pF; 23 times the highest tone
Tuning	none	Flat response across all tones
Electrostatic shield	slotted foil	Blocks capacitive pickup; the slot prevents a shorted turn

The coil can be read in voltage mode, where the output rises in proportion to frequency, or by a transimpedance amplifier, which gives a flat response above 446 Hz. Voltage mode needs no extra components and its slope is removed by calibration. The transimpedance stage is kept as an option in case the lowest tones turn out to be limited by dynamic range.




Figure 5.2. Response of the reference coil in voltage mode and transimpedance mode, normalised at 10 kHz. Across the tone band the voltage-mode response changes by a factor of about 100.

5.5 Digitiser: CS42448

The CS42448 provides six 24-bit analog-to-digital converters and eight digital-to-analog converters behind one TDM interface. All channels run from one master clock, so inter-channel phase is fixed by construction. The board available to the project ("6IN8OUT") has line-level, AC-coupled inputs without microphone preamplifiers.

Table 5.2. CS42448 properties relevant to the design. Noise and full-scale values are derived from the datasheet and will be measured in T1.3.

Property	Value	Consequence
Converters	6 in, 8 out, 24-bit	Sufficient for 3 TX, 3 current references and 3 RX
Clocking	one master clock, TDM	All channels sample-synchronous
ADC dynamic range	about 105 dB (A-weighted)	Input noise about 80 nV/√Hz at 2 V _{rms} full scale
Input type	line level, AC-coupled	An external low-noise stage is required per receiver
Sample rate	up to 192 kHz; 48 kHz planned	96 kHz to be tested (T1.3)




Figure 5.3. Channel assignment in the two build phases. Phase A uses one transmitter and leaves an input for the GNSS pulse. Phase B uses all six inputs for the full-tensor arrangement; the pulse is then summed onto the first current reference. A spare output provides a permanent calibration tone.

5.6 Noise budget

The input noise of the codec, about 80 nV/√Hz, is roughly 36 times higher than that of a good microphone preamplifier. The echo of a 1 cm copper sphere at 20 cm depth is about 20.5 μV at the coil and falls as $1/d^6$.




Figure 5.4. Echo of a 1 cm copper sphere against the noise floors. With 4.3 s averaging the codec alone is adequate. For single 85 ms blocks during movement its floor coincides with the estimated ground response. A $\times 33$ preamplifier restores the margin.

Each receiver therefore gets a non-inverting $\times 33$ stage with a low-noise operational amplifier. With an OPA1612 the input-referred noise of the chain is 2.6 nV/ $\sqrt{\text{Hz}}$; with an NE5532 it is 5.5 nV/ $\sqrt{\text{Hz}}$. In both cases the ground response, estimated at about 300 times the electronic floor, sets the practical detection range of about 40 cm for coin-sized targets.

5.7 Excitation

All tones are transmitted at once as a multi-tone waveform. At 48 kHz and a block length of 4096 samples the frequency resolution is 11.72 Hz and one block lasts 85.3 ms. Twelve tones are placed on exact FFT bins with logarithmic spacing:

Table 5.3. Tone plan for 48 kHz sampling.

Quantity	Value
FFT bins (N = 4096)	16, 24, 37, 56, 85, 130, 197, 300, 456, 693, 1053, 1600
Frequencies	188, 281, 434, 656, 996 Hz; 1.52, 2.31, 3.52, 5.34, 8.12, 12.3, 18.75 kHz
Crest factor	4.90 with zero phases; 3.06 with optimised phases (4.1 dB more power per tone)
Phases (rad)	2.589, 5.287, 1.692, 2.610, 3.376, 2.940, 1.805, 1.120, 0.966, 3.592, 5.042, 0.208




Figure 5.5. (a) Tone comb. (b, c) One block of the transmit waveform with zero phases and with optimised phases, scaled to equal power. The peak value, which limits the amplifier, falls from 4.90 to 3.06 times the RMS value.

Simultaneous transmission is not free. At equal amplifier peak, twelve simultaneous tones give about 6.7 dB less signal-to-noise ratio per tone than stepping a single tone through the same frequencies in the same total time. We therefore use stepped tones for static measurements and simultaneous tones when the sensor moves. For moving acquisition the block length is reduced to 1024 samples (21.3 ms, about 1 cm of travel at walking speed).




Figure 5.6. Quadrature response of the reference targets compared with the band available at 48 kHz. The peaks of coin-sized copper and aluminium targets lie inside the band. Stainless steel and the ferrous peak lie above it; ferrous discrimination does not depend on the peak.

5.8 Coil configurations

The tensor $\mathbf{M}$ is symmetric and has six independent elements per frequency. One transmitter excites the target along one direction $\mathbf{t}$; three receivers then measure the vector $\mathbf{M}\mathbf{t}$, which is one column of the tensor. Three transmitters along independent directions give all columns and over-determine the six unknowns.




Figure 5.7. Information obtained with one and with three transmitters. Colours in the right-hand matrix mark pairs of equal elements.

Table 5.4. Coil configurations.

Configuration	Out / in	Location	Tensor from one position	Use
1 TX + 1 RX	1 / 2	Only by scanning	1 element	Conventional detector
1 TX + 3 RX	1 / 4	Yes	One column	Phase A; full tensor by moving the sensor (Chapter 7)
3 TX + 3 RX	3 / 6	Yes	Full tensor	Phase B

The three transmit coils are oriented along x, y and z. Coils lying side by side in one plane produce nearly parallel fields at a deep target and add little information. Orthogonal coils also have almost no mutual coupling. All three transmit at the same time on interleaved tone sets (for example bins 16, 24, ... for coil x; 17, 25, ... for y; 18, 26, ... for z). Tones on distinct FFT bins are orthogonal over one block, so a single FFT separates the three transmitters.

Two of the receive coils can also be operated as a gradiometer. Their outputs are subtracted, so that distant fields such as mains hum and the bulk ground response cancel, while the field of a nearby target is retained.




Figure 5.8. Gradiometer. Distant fields reach both receive coils equally and cancel in the difference. The dipole field of a nearby target falls as $1/r^3$ and is much stronger in the lower coil.

5.9 Coupling and nulling

The field of the transmit coil at a receive coil is 10^5 to 10^6 times larger than a target echo. This direct coupling has to be reduced geometrically, by overlapping the coils so that the net flux through the receiver is close to zero, before any electronics can help. The design target is a residual of 10^{-3} or less at all twelve tones simultaneously (task T2.2). The remaining coupling is measured without a target and subtracted as a background spectrum. This nulling is the most demanding analog task of the project.

Radar subsystem

Stepped-frequency ground-penetrating radar on the ADALM-Pluto.

The radar subsystem measures depth and shape. It does not identify the metal, but it provides the independent depth estimate that the induction measurement needs, and its false alarms have different causes from those of induction.

6.1 Why not induction on the SDR

The ADALM-Pluto uses the AD9363 transceiver, specified from 325 MHz to 3.8 GHz and commonly operated from 70 MHz to 6 GHz with the AD9364 firmware profile. Induction needs kilohertz frequencies, four orders of magnitude lower. The SDR is therefore used as a radar, and induction is handled by the codec (Chapter 5).

6.2 Stepped-frequency continuous-wave radar

An SFCW radar transmits one frequency at a time and records the complex echo at each step. An inverse FFT of the sweep gives the echo as a function of range. The method needs neither a fast pulse generator nor a fast analog-to-digital converter, which suits an SDR.




Figure 6.1. Radar subsystem. The power amplifier and the low-noise amplifier are used in their intended band. Both antennas are shielded and backed with absorber so that they look into the ground.




Figure 6.2. Simulated SFCW measurement: 251 steps from 0.5 to 3 GHz, three reflectors in soil with $\epsilon_r = 9$. After a Hann window and inverse FFT the ground surface, a metal object at 35 cm and a stone at 80 cm are separated.

Table 6.1. Radar parameters.

Quantity	Relation	Value
Sweep	0.5 to 3.0 GHz, 10 MHz steps	251 steps
Range resolution in soil	$\Delta R = v/2B, v = c/\sqrt{\epsilon_r}, \epsilon_r = 9$	2.0 cm
Range resolution in air	$c/2B$	6.0 cm
Unambiguous range in soil	$v/2\Delta f$	5 m
Sweep time	limited by synthesiser retuning	about 0.3 s
Averaging gain	$\sqrt{N}$	+20 dB for 100 sweeps

Low frequencies penetrate further and high frequencies resolve more detail. Because the whole sweep is recorded, the band can be windowed after the measurement. Typical penetration is more than a metre in dry sand and a few tens of centimetres in wet clay.

6.3 Phase coherence

The AD9363 has separate transmit and receive synthesisers locked to one reference. After each retune their relative phase can change, and SFCW processing depends on phase. The second receive channel of the AD9363, available on Pluto revision C and later after enabling the two-channel firmware mode, is therefore used as a reference. A directional coupler feeds a small part of the transmit signal into it, and each echo is divided by the reference:

$$S(f_k) = \frac{\text{RX1}(f_k)}{\text{RX2}(f_k)}.$$

Both receive channels share one synthesiser, so the division cancels the retuning phase.

6.4 Structural resonances

A metal object that is comparable in size to half a wavelength resonates. For an elongated object of length L the first resonance is approximately

$$f_1 \approx \frac{c}{2L\sqrt{\epsilon_r}}.$$




Figure 6.3. First structural resonance as a function of object length in air and in soil ($\epsilon_r \approx 9$), with the AD9363 bands shaded. Objects from a few centimetres to about one metre resonate within the band when buried.

The late part of a radar echo is a sum of damped oscillations, one per resonance. Their poles depend on the size and shape of the object but not on the viewing angle, which makes them usable as a signature. They are estimated with the matrix-pencil method. Metal reflects almost completely with inverted phase ($\Gamma \approx -1$), while stones and voids reflect partly without inversion. The echo polarity is therefore a second test for metal. Resonances of coin-sized objects lie near 2.5 GHz in soil, where attenuation is high; the radar is aimed at objects of about 5 cm and larger.

6.5 Antennas and amplifiers

Two Vivaldi antennas on FR4 cover the sweep. The lowest usable frequency requires an aperture of about half a wavelength: roughly 30 cm for 0.5 GHz. The Pluto delivers up to about +7 dBm, so a power amplifier is used on transmit and a low-noise amplifier on receive. The antennas are placed close to the ground and shielded on top, which improves the signal and also keeps the emission directed into the soil (Chapter 12).




Figure 6.4. Fusion of radar and induction. Each sensor answers questions the other cannot. The combination follows established dual-sensor detectors (Chapter 2).

Positioning and timing

Centimetre-level pose of the sensor head on a common GPS time base.

The positioning subsystem records where the sensor head is, and how it is oriented, at every measurement block. This turns a moving sensor into a set of measurements from many known positions. It also gives an estimate of target depth that does not depend on the signal amplitude.

7.1 Why position matters

When the sensor passes over a target at constant height, the received signal as a function of *position* forms a peak. Its width is set by geometry. For a coincident coil over a small sphere, the full width at half maximum equals 0.90 times the depth.




Figure 7.1. Computed signal along a straight pass over a small sphere at 10, 20 and 30 cm depth, each normalised to its own maximum. The width is proportional to depth.

Depth obtained from the width does not depend on the size of the target. This matters because signal amplitude alone cannot separate a small shallow object from a larger deep one (Chapter 3). With depth known, amplitude gives size, and size together with the product σa^2 from the spectrum gives an estimate of conductivity.

The measurement also depends strongly on coil height. From $|V| \propto r^{-6}$, a change in distance dr changes the signal by $dV/V = -6 dr/r$. At 20 cm, raising the coil by 1 cm reduces the signal by about 30%. Coil height is therefore measured directly with a time-of-flight ranger.

With 21.3 ms blocks and a walking speed of 0.5 m/s the sensor moves 1.1 cm per block, which samples the peak adequately. With pose known for every block, a head with one transmitter and three receivers views the target from many directions during a pass. This provides the tensor columns that it cannot measure from a single position (Chapter 5).

7.2 Sensors

7.2.1 Inertial measurement unit

Integrating acceleration twice converts an uncorrected accelerometer bias b into a position error of $\frac{1}{2}bt^2$. For a bias of 1 mg this reaches 2 cm after 2 s, 12 cm after 5 s and 49 cm after 10 s. The IMU is therefore used only to interpolate between absolute position fixes and to measure tilt.




Figure 7.2. Position error against time since the last fix, for inertial dead reckoning with 1 mg and 5 mg bias, for standard GNSS and for RTK GNSS. Between RTK fixes at 10 Hz the IMU error stays below 1 cm.

7.2.2 Magnetometer

The magnetometer measures heading from the Earth's static field. Close to the transmit coil the coil's own alternating field dominates: with a moment of $0.9 \text{ A}\cdot\text{m}^2$ it is $181 \mu\text{T}$ at 10 cm, compared with the Earth's $50 \mu\text{T}$. The coil field falls as $1/r^3$. At 50 cm it is $1.45 \mu\text{T}$, which limits the heading error to 3.3° , and at 1 m it is $0.18 \mu\text{T}$ (0.4°). The magnetometer is therefore mounted at least 50 cm above the coil. Heading is also disturbed by steel near the survey area and is treated as a slowly corrected estimate.




Figure 7.3. Field of the transmit coil along the pole compared with the Earth's field. Within about 15 cm the coil field exceeds the Earth's field.

7.2.3 Satellite positioning

Table 7.1. Satellite positioning options.

Receiver	Horizontal error	Suitability
Standard GNSS (phone class)	2 to 5 m	Unsuitable; larger than the signal peak by two orders of magnitude
RTK GNSS, dual-frequency receiver with base station	1 to 2 cm	Required
PPK (same hardware, corrections applied afterwards)	1 to 2 cm	Alternative without a radio link; no live map

The RTK antenna is mounted at the top of the pole for a clear view of the sky, about 1.5 m above the coil. An unmeasured tilt θ moves the coil by $L \sin \theta$ relative to the antenna, which is 2.6 cm per degree and 13 cm at 5°. The IMU tilt estimate is used to project the antenna position down to the coil.

7.2.4 Camera

A camera facing the ground tracks surface texture from frame to frame (visual odometry). Its drift grows with distance travelled, typically 0.5 to 1%, and not with time. Combined with the IMU it continues to work under tree cover, where RTK degrades. A single camera cannot measure absolute scale, but the time-of-flight ranger measures the camera height and provides it. A global-shutter camera is used because a rolling shutter distorts images under vibration. Visual odometry fails on featureless or moving surfaces and in poor light; the IMU bridges such periods.

7.3 Mechanical arrangement




Figure 7.4. Sensor pole. The RTK antenna is at the top, the IMU and magnetometer at least 50 cm above the coil, and the coil head with the height sensor at the bottom. The faint outlines show the same pole at earlier times during a pass.

The pole and all mountings near the coils are non-metallic. The lever arm from antenna to coil centre is measured to the millimetre during assembly and verified by the rotation test in task T3.1.

7.4 Timing and fusion

Position fixes come from the GNSS clock, induction samples from the codec clock and video frames from the camera clock. The GNSS pulse-per-second signal connects them. It is recorded on a codec input, which places GPS time in the audio stream with one-sample resolution ($20.8 \mu\text{s}$ at 48 kHz). It also drives an LED in the camera's field of view, so each frame can be assigned a GPS time. The AC coupling of the codec input converts the pulse into a pair of spikes; edges are detected rather than levels.




Figure 7.5. Fusion and synchronisation. An extended Kalman filter combines IMU, RTK, magnetometer, camera and height data into the pose of the coil. Pose is attached to every induction block. Repeated passes over one target provide a consistency check.

Table 7.2. Positioning performance targets (milestone M2).

Quantity	Target	Verification
Horizontal position of coil centre	≤ 2 cm (1σ)	Ten passes over a marked target from different headings
Coil height above ground	≤ 3 mm	Comparison with a measured reference surface
Timing between induction and pose	≤ 1 sample	PPS edge in audio stream against GNSS time tag
Depth from peak width at 20 cm	± 1.5 cm	Buried calibration sphere at known depth

Signal processing and machine learning

From raw samples to target position, class and size.

This chapter describes how raw samples become target estimates: demodulation, background removal, inversion of the dipole model, classification, fusion with radar, and evaluation.

8.1 Data flow

1. **Demodulation.** Each block of codec samples is transformed by one FFT. The complex values at the twelve tone bins of each receiver are divided by those of the corresponding transmit current reference. The result is a response matrix per block: transmitters $\times$ receivers $\times$ tones.
2. **Time stamping.** The PPS edges in the audio stream assign GPS time to each block; the pose estimate (Chapter 7) is interpolated to that time.
3. **Background removal.** A background spectrum measured without a target is subtracted. Ground response is estimated along the track with a moving median and fitted with the viscous-soil model (Chapter 3).
4. **Detection.** Candidate targets are local maxima of the residual signal on the position map.
5. **Inversion.** For each candidate, all blocks within about one metre are used to fit the dipole model.
6. **Classification and fusion.** Orientation-independent features are classified and combined with radar results.

8.2 Dipole inversion

The unknowns for one target are its position (3), its orientation (3) and the three principal polarizabilities $\beta_1, \beta_2, \beta_3$ at each of the twelve tones (36 complex values). The forward model computes the coil fields $\mathbf{h}_{tx}$ and $\mathbf{h}_{rx}$ from the known coil geometry and pose, and predicts each measurement from the tensor model in Chapter 3. Position and orientation are found by non-linear least squares (Levenberg–Marquardt, several starting points). For fixed geometry the polarizabilities follow by linear least squares. The quality of the fit and the covariance of the estimates are stored with each result.

8.3 Features

- The three curves $\beta_i(\omega)$, normalised by their high-frequency values: they describe the shape and material class independently of orientation and depth.

- The quadrature peak frequency and width of each curve (Chapter 3): relaxation time and geometry.
- The sign of the low-frequency in-phase response: ferrous or non-ferrous.
- The ratios $\beta_1 : \beta_2 : \beta_3$: sphere, elongated or flat object.
- The absolute magnitude together with the fitted depth: size.
- Where available, radar features: echo polarity, estimated resonance poles and depth.

8.4 Classification




Figure 8.1. Processing pipeline. The physics model removes position and orientation before classification. An end-to-end network on raw measurements is kept as a comparison but needs considerably more data.

The baseline classifier is gradient boosting on the features above. It works with a few hundred examples. A one-dimensional convolutional network on the $\beta_i(\omega)$ curves is trained for comparison. Class labels are ferrous / non-ferrous and a shape class (coin or disc, ring, rod or nail, can or hollow body, foil, other). Size is estimated by regression. Every prediction is reported with a calibrated probability.

8.4.1 Simulation

The forward model is cheap to evaluate, so large training sets can be generated: exact solutions for spheres, analytical approximations for spheroids and finite-element models for selected real shapes. Depth, orientation, position, noise and soil response are drawn at random. Models are trained on simulated data and calibrated and validated on measured data. Performance is only reported on measured data.

8.5 Fusion with radar

The division of work between the two sensors is shown in Figure 6.4. Fusion is carried out at the level of candidate targets. A candidate is kept if both sensors detect it, or if one sensor detects it with high confidence. Radar depth is used as a prior in the induction inversion. Because the two sensors are on the same GPS time line and pose is known, their detections can be associated by position.

8.6 Evaluation protocol

- **Split by object.** Training and test sets never share a physical object. Otherwise the model learns individual objects and the test score is meaningless.
- **Blind test.** Test objects are buried by a team member who is not involved in the analysis; positions and identities are revealed only after the predictions have been submitted.
- **Metrics.** Classification accuracy with 95% confidence intervals, confusion matrices, calibration of probabilities, size error, position error, and detection rate against false alarms per square metre.
- **Reference.** The same test set is also assessed with the target-identity reading of a commercial detector, to measure the improvement.

8.7 Data management

Every capture is stored as raw codec samples (FLAC, lossless) with a JSON metadata file. The metadata records time, operator, site, soil condition, coil configuration, tone plan, pose track, ground truth (object identifier, depth, orientation, position) and the software version. Files are named `site_date_run_capture`, checksummed, and copied to the project server nightly. Object records include photographs, dimensions, mass and material. The released dataset carries a data sheet describing collection, known biases and licence. Code is kept in the project's Git repository on the team server; every result in the report refers to a tagged commit.

Work plan

Work packages, tasks, schedule, milestones and decision gates.

The project runs for 52 weeks with a team of four (Chapter 10). The work is divided into seven work packages: a continuous management package (WP0) and six technical packages. Each technical package ends with a measurable milestone. Four decision gates define what happens if a milestone is missed.

9.1 Structure and logic

The order of the work follows the dependencies between measurements. The receive chain must work (WP1) before the static rig can be calibrated against theory (WP2). The full-tensor dataset (WP4) needs a calibrated rig and, for tracked captures, the positioning platform (WP3). The radar (WP5) is developed in parallel because it has no dependency on the induction hardware. Fusion, field trials and release (WP6) close the project.

Table 9.1. Work packages.

WP	Title	Lead	Weeks	Objective
WP0	Project management	PI	1–52	Planning, reviews, procurement, documentation, reporting
WP1	Bring-up and receive chain	HW	1–8	Working codec, coils and preamplifiers with measured characteristics
WP2	Static induction rig and calibration	PI	6–18	Measured sphere spectra agreeing with the exact model
WP3	Positioning and timing platform	SW	10–26	Pose of the coil at every block to 2 cm
WP4	Full tensor, dataset and classification	DS	18–36	Labelled dataset and a classifier validated on unseen objects
WP5	Radar subsystem	HW	22–40	SFCW radar imaging metal targets at 35 cm
WP6	Fusion, field trials and release	PI	36–52	Fused system tested in the field; public release

9.2 Schedule




Figure 9.1. Schedule at task level. Bars are coloured by the role that leads the task; diamonds mark milestones; weeks 1 to 52.

9.3 Tasks

Each task has one responsible owner (see the role codes in Chapter 10), a time window and a defined result. The acceptance criterion states how completion is checked.

Table 9.2. Tasks with owner, supporting role, time window and acceptance criterion.

ID	Task	Lead	Support	Weeks	Acceptance criterion / result
WP0 · Project management (lead PI, weeks 1–52)					
T0.1	Kick-off; project handbook, repository and data conventions	PI	all	1–3	Handbook D0.1 agreed by all four team members
T0.2	Procurement and budget control	PI	SW	1–52	Orders placed per procurement schedule; monthly spending report
T0.3	Weekly meetings, monthly reviews, risk and decision logs	PI	all	1–52	Minutes and logs up to date at every monthly review
T0.4	Maintenance of this report; quarterly reports	PI	DS	1–52	Report revised monthly; quarterly reports on time
WP1 · Bring-up and receive chain (lead HW, weeks 1–8)					
T1.1	Bench set-up: instruments, supplies, single-point grounding	HW	–	1–2	Bench ready; ground scheme documented
T1.2	Codec bring-up on Teensy 4.1 and Raspberry Pi; host decision	SW	HW	1–4	All 14 channels verified in loopback; no lost samples in 1 h; host chosen
T1.3	Codec characterisation: full scale, input noise, coupling corner, 96 kHz test	HW	SW	3–5	Measured values replace datasheet values in Chapter 5 (D1.1)
T1.4	Wind and measure transmit and receive coils	HW	–	2–5	Self-resonance at least 400 kHz; inductance within 10% of 4.58 mH
T1.5	Preamplifier and calibration-injector PCB: design, fabrication, test	HW	SW	3–8	Input noise at most 3 nV/√Hz; channel gains matched within 1% (D1.2)
T1.6	Acquisition software: tone generation, block demodulation, logging	SW	DS	2–7	12-tone loopback stable to 0.1% amplitude and 0.1° phase over 10 min
WP2 · Static induction rig and calibration (lead PI, weeks 6–18)					
T2.1	Transmit driver: class-D amplifier, series resistor, current sense, soft ramps	HW	–	6–9	1 A per coil without clipping at crest factor 3.1
T2.2	Coil geometry and nulling fixture	HW	PI	8–13	Direct coupling at most 10^{-3} at all 12 tones (D2.2)
T2.3	Forward model: exact sphere, dipole model, coil geometry	PI	–	6–10	Model reproduces the three analytic limits and published

ID	Task	Lead	Support	Weeks	Acceptance criterion / result examples
T2.4	Calibration measurements with four sphere types at 10, 20 and 30 cm	PI	DS	13–16	Complete data set including repeat measurements
T2.5	Comparison with the model; systematic error budget	PI	DS	15–18	M1 criterion evaluated; calibration report D2.1
T2.6	Ground survey at two sites; ground-response factor	DS	HW	14–18	Background spectra and measured ground factor for both sites
WP3 · Positioning and timing platform (lead SW, weeks 10–26)					
T3.1	Sensor pole: non-metallic structure, mounts, lever arm	HW	–	10–13	Rotation test: computed coil position stays within a 2 cm circle
T3.2	RTK base and rover, correction link, availability survey	SW	–	10–14	Fixed solution for at least 95% of the time at the test sites
T3.3	PPS timing into codec and camera	SW	–	12–16	Timing error at most one sample against GNSS time tags
T3.4	IMU and magnetometer integration; interference test with transmitter on	SW	HW	13–16	Heading change at most 1° when the transmitter switches on
T3.5	Camera visual-inertial odometry with height scaling	DS	SW	14–20	Drift at most 1% of distance on a 20 m test track
T3.6	Extended Kalman filter: RTK, IMU, camera, magnetometer, height	SW	DS	16–24	Continuous pose at block rate (D3.1)
T3.7	Tracked-pass validation: peak width against depth, repeatability	PI	DS	22–26	M2 criterion evaluated; validation report D3.2
WP4 · Full tensor, dataset and classification (lead DS, weeks 18–36)					
T4.1	Three-axis transmit head with interleaved tone sets	HW	SW	18–24	Cross-talk between transmitters at most -60 dB after FFT separation
T4.2	Tensor inversion software	PI	DS	20–28	Simulated tensors recovered within 5% at 30 dB SNR (D4.2)

ID	Task	Lead	Support	Weeks	Acceptance criterion / result
T4.3	Simulator: spheres, spheroids, finite-element shapes; noise and soil	DS	PI	20–30	10,000 simulated captures; agreement with calibration data
T4.4	Dataset collection: at least 30 objects, 500 captures, 2 soils	DS	all	24–34	Dataset v1 with metadata and data sheet (D4.1)
T4.5	Classifiers: gradient boosting baseline, 1D-CNN comparison	DS	PI	28–35	Cross-validated accuracy with confidence intervals
T4.6	Blind test on hidden objects	DS	HW	34–36	M3 criterion evaluated; blind-test report D4.3
WP5 · Radar subsystem (lead HW, weeks 22–40)					
T5.1	Pluto set-up: firmware profile, two-channel mode, coupler loopback	HW	SW	22–25	Reference-corrected phase stable within 2° over 100 retunes
T5.2	Vivaldi antennas: simulation, fabrication, measurement	HW	–	22–30	S11 at most -10 dB from 0.5 to 3 GHz
T5.3	Sweep software, background subtraction, range profiles	SW	PI	24–32	Metal plate in air at the correct range within 2 cm
T5.4	Amplifiers, shielding and emission check	HW	–	30–33	Emission within applicable limits; radar hardware D5.1
T5.5	Sandbox radar tests along tracked passes	PI	DS	32–40	M4 criterion evaluated; radar report D5.2
WP6 · Fusion, field trials and release (lead PI, weeks 36–52)					
T6.1	Fusion of induction and radar detections	DS	PI	36–44	Fused false-alarm rate lower than either sensor alone
T6.2	Live map and operator interface	SW	–	38–46	Map updated within 1 s of acquisition during a pass
T6.3	Field trials at two sites with seeded targets	PI	all	42–48	Trial report with ground truth by excavation (D6.1)
T6.4	Documentation of hardware, software and dataset	SW	all	40–52	An outside reader rebuilds the receive chain from the documents
T6.5	Review of patents and non-English literature	PI	–	40–44	Review added to Chapter 2
T6.6	Paper preparation and submission	PI	all	44–52	Paper submitted (D6.2)

ID	Task	Lead	Support	Weeks	Acceptance criterion / result
T6.7	Public release of designs, code, dataset and report	SW	all	50–52	Release published (D6.3); M5

9.4 Milestones

Table 9.3. Milestones and acceptance criteria.

ID	Week	Milestone	Criterion	Verification	Owner
M1	18	Calibration agrees with theory	Spectra of the four calibration spheres agree with the exact model within $\pm 10\%$ in amplitude and $\pm 5^\circ$ in phase at all 12 tones, at 10 and 20 cm	Calibration report D2.1, reproduced independently from raw data by DS	PI
M2	26	Positioning validated	Ten passes over a marked target from different headings give positions within 2 cm (1σ); depth from peak width within ± 1.5 cm at 20 cm	Validation report D3.2	SW
M3	36	Classifier validated on unseen objects	Blind test: ferrous / non-ferrous accuracy at least 95%, shape class at least 85%, size within $\pm 25\%$ for at least 80% of objects	Blind-test report D4.3; objects hidden by HW	DS
M4	40	Radar imaging	A metal can of about 10 cm buried at 35 cm in dry sand is detected and located; depth error at most 3 cm; echo polarity identified as metal	Radar report D5.2	HW
M5	52	Fused demonstration and release	Live fused survey of a seeded 5 m $\times$ 5 m plot; designs, software, dataset and report published	Public repository; field-trial report D6.1	PI

9.5 Decision gates

At each gate the team reviews the milestone evidence and chooses one of the prepared options. Decisions and their reasons are recorded in the decision log.

Table 9.4. Decision gates.

Gate	Week	Condition	If met	If not met
G1	18	M1 met	Full-tensor work (WP4) starts as planned	Up to 4 weeks of diagnosis of model or hardware; WP4 start moves accordingly, WP3 continues. If still not met in week 22, the dataset uses empirically calibrated spectra and the discrepancy is reported.
G2	26	M2 met	Tracked dataset collection with the hand-held pole	A wheeled cart with odometry replaces the hand-held pole; RTK used for start and end points of each track.
G3	36	M3 met	WP6 proceeds as planned	Data collection extended by 4 weeks; the operator interface (T6.2) is reduced to a minimum.
G4	40	M4 met	Radar included in fusion and field trials	Induction-only results are published; the radar is reported as a partial result and future work.

9.6 Critical path and dependencies

The critical path runs T1.2 → T1.5 → T2.2 → T2.4 → M1 → T4.1 → T4.4 → T4.5 → M3 → T6.1 → T6.3 → M5. Its slack is small: a delay in nulling the coils (T2.2) or in the calibration measurements (T2.4) moves the dataset collection and all later work. The positioning platform (WP3) and the radar (WP5) have about four weeks of slack each. The main external dependency is procurement of the RTK receivers and printed circuit boards, which is started in week 1 (T0.2).

Table 9.5. Procurement schedule.

Item	Order in week	Needed in week	Remark
Teensy 4.1, op-amps, passives, connectors	1	2	Local or fast-shipping distributors
Magnet wire, foil, coil formers	1	2	
Preamplifier PCB, first run	4	6	Design completed in T1.5
Class-D amplifier modules, sense resistors	3	6	
RTK receivers, antennas, radio link	1	10	Longest lead time, import
IMU, magnetometer, camera, ranger	6	12	
Calibration spheres, test objects, sandbox	6	12	
Vivaldi antenna PCBs, RF cables, coupler	20	24	After simulation in T5.2
Second PCB run (revised preamp, tensor head)	16	20	

9.7 Deliverables

Table 9.6. Deliverables.

ID	Deliverable	Type	Due week	Owner
D0.1	Project handbook: lab-log template, data conventions, safety rules	Document	3	PI
D1.1	Codec bring-up and characterisation report	Report	5	SW
D1.2	Preamplifier and calibration-injector PCB, revision A	Hardware	8	HW
D2.2	Static induction rig (1 TX + 3 RX), documented	Hardware	13	HW
D2.1	Calibration report: measured sphere spectra against the model	Report	18	PI
D3.1	Positioning and timing platform with fusion software	Hardware, software	24	SW
D3.2	Tracked-pass validation report	Report	26	PI
D4.2	Tensor inversion software	Software	28	PI
D5.1	Radar hardware: antennas, amplifiers, coupler	Hardware	33	HW
D4.1	Labelled dataset, version 1	Data	34	DS
D4.3	Classifier and blind-test report	Report	36	DS
D5.2	Radar software and test report	Software, report	40	PI
D6.1	Field-trial report of the fused system	Report	48	DS
D6.2	Paper submitted to a journal or conference	Publication	52	PI
D6.3	Public release: designs, code, dataset, report	Release	52	SW

9.8 Reporting and reviews

- **Weekly:** one-hour team meeting (progress against tasks, blockers, next week's measurements). Short written notes are added to the lab log.
- **Every two weeks:** demonstration of a working result on real hardware or real data.
- **Monthly:** review of schedule, budget and risk register; update of this report.
- **At each milestone:** milestone review with a written report, followed by the decision gate where one applies.
- **Quarterly:** progress report to the funder with spending against budget.

Team and responsibilities

Four roles, who does what, effort and working practices.

The project needs four people with complementary skills. This chapter defines the four roles, what each person is responsible for, how effort is distributed and how the team works together. Names are entered when the positions are filled; throughout the report the roles are referred to by their codes.

10.1 Roles

Table 10.1. The four roles.

Code	Role and profile	Responsibilities	Ownership
PI	Principal investigator (physicist) Electromagnetic modelling, measurement design and data analysis; scientific Python; laboratory practice.	Scientific direction and project management (WP0). Forward models and calibration against exact solutions (WP2). Tensor inversion and radar analysis (WP4, WP5). Fusion method, field-trial design and publication (WP6). Chairs reviews and decision gates.	Leads WP0, WP2, WP6 13 tasks Owns D0.1, D2.1, D3.2, D4.2, D5.2, D6.2
HW	Hardware engineer (analog and RF) Analog circuit design, PCB layout, coil winding, RF measurement with a vector network analyser, mechanical fabrication.	Coils, transmit drivers, preamplifiers and nulling (WP1, WP2). Sensor pole mechanics (WP3). Three-axis coil head (WP4). Antennas, amplifiers and radar front end (WP5). Safety officer; places objects for blind tests.	Leads WP1, WP5 11 tasks Owns D1.2, D2.2, D5.1
SW	Embedded and software engineer Embedded C/C++, real-time acquisition, Linux, Kalman filtering, software engineering practice.	Codec bring-up and host streaming (WP1). Acquisition software. GNSS, IMU and timing integration and the fusion filter (WP3). Radar sweep software (WP5). Live map, data pipeline, repository and public release (WP6).	Leads WP3 10 tasks Owns D1.1, D3.1, D6.3
DS	Data scientist and field engineer Machine learning, statistics, simulation, experimental design; field work.	Test objects and acquisition support (WP1). Ground survey (WP2). Camera odometry (WP3). Simulator, dataset collection, classifiers and blind evaluation (WP4). Sensor fusion and dataset release (WP6).	Leads WP4 7 tasks Owns D4.1, D4.3, D6.1

10.2 Responsibilities by work package

Each work package has one accountable lead and at least one second person who reviews and can take over. The matrix uses the usual RACI convention: **R** responsible (does the work), **A** accountable (owns the result and signs off), **C** consulted, **I** informed.

Table 10.2. Responsibility matrix by work package.

WP	Title	PI	HW	SW	DS
WP0	Project management	A/R	C	R	C
WP1	Bring-up and receive chain	C	A/R	R	R
WP2	Static induction rig and calibration	A/R	R	C	R
WP3	Positioning and timing platform	R	R	A/R	R
WP4	Full tensor, dataset and classification	R	R	C	A/R
WP5	Radar subsystem	R	A/R	R	C
WP6	Fusion, field trials and release	A/R	R	R	R

10.3 Effort

Each person works full time for twelve months, giving 48 person-months in total. The table shows the planned distribution. It is used to cost the work packages in Chapter 11.

Table 10.3. Effort in person-months.

Role	WP0	WP1	WP2	WP3	WP4	WP5	WP6	Total
PI	1.5	0.5	3.0	1.0	2.0	1.0	3.0	12.0
HW	0.3	3.0	2.5	1.2	1.5	3.0	0.5	12.0
SW	0.3	2.5	0.5	4.5	0.5	1.7	2.0	12.0
DS	0.2	0.8	1.5	1.5	5.0	0.5	2.5	12.0
Total	2.3	6.8	7.5	8.2	9.0	6.2	8.0	48.0

10.4 Focus by quarter

Table 10.4. Main activities of each role per quarter.

Role	Q1 (weeks 1–13)	Q2 (14–26)	Q3 (27–39)	Q4 (40–52)
PI	Forward model, calibration plan, procurement	Calibration against theory, M1 review	Tensor inversion, radar analysis	Fusion, field trials, paper
HW	Coils, preamplifiers, TX driver	Nulling, pole mechanics, tensor head	Radar antennas and amplifiers	Field hardware, documentation
SW	Codec bring-up, acquisition software	RTK, timing, sensor fusion	Radar software, data pipeline	Live map, release
DS	Acquisition software, test objects	Ground survey, camera odometry	Simulator, dataset, classifier	Blind test, fusion, dataset release

10.5 Working practices

- **Two people per result.** Every measurement or analysis that feeds a milestone is checked by a second team member before it is reported.
- **Lab log.** Each measurement session is recorded with date, set-up, settings, data file names and observations, including failed attempts.
- **Code review.** All code is merged through reviewed pull requests in the project repository; analysis scripts are versioned with the data they produced.
- **Decision log.** Design choices and gate decisions are recorded with the alternatives considered and the reason for the choice.
- **Blind testing.** The person who buries test objects for an evaluation (normally HW) does not take part in the analysis of that test.
- **Safety officer.** HW is responsible for the electrical and field safety rules in Chapter 12 and briefs each new participant.

10.6 Substitution

If one person becomes unavailable, the second owner of each of that person's tasks (the RACI "R" entries of another role) continues the work. For PI, the DS role takes over analysis and SW takes over coordination. For HW, PI covers measurements and an external contractor may be engaged for PCB work (contingency budget). The written lab log and decision log make such transfers possible without loss of information.

Budget and funding

Cost of the twelve-month programme and the grant programmes it is aimed at.

The project is planned for grant funding: four people working full time for twelve months, with the equipment, running costs and overheads needed to complete all six technical work packages. This chapter gives the cost basis of the request and the funding programmes it is aimed at. Prices are indicative street prices in US dollars as of October 2026, excluding VAT, with an allowance for shipping and import to Armenia.

11.1 Equipment and materials

Table 11.1. Equipment and materials. Prices in USD.

Item	Qty	Unit	Total	Status	WP
Owned equipment					
ADALM-Pluto software-defined radio (AD9363)	1	–	0	owned	WP5
RF power amplifier and low-noise amplifier modules	1	–	0	owned	WP5
CS42448 codec board (6IN8OUT)	1	–	0	owned	WP1
Induction					
Teensy 4.1 (one spare)	2	32	64	buy	WP1
Raspberry Pi 5, 4 GB (host alternative)	1	60	60	buy	WP1
OPA1612 dual operational amplifiers	6	6	36	buy	WP1
Sense resistors, protection diodes, passive components	1	40	40	buy	WP1
PCB fabrication and assembly parts, two runs	2	60	120	buy	WP1
Magnet wire 0.3 mm, 500 g	1	30	30	buy	WP1
Coil formers (3D-printing filament)	2	25	50	buy	WP1
Copper foil tape and potting epoxy	1	35	35	buy	WP1
Class-D amplifier boards, two channels	2	15	30	buy	WP2
Radar					
Vivaldi antenna PCBs, large-format FR4	2	30	60	buy	WP5
Directional coupler, SMA cables and adapters	1	70	70	buy	WP5
NanoVNA for antenna measurement	1	70	70	buy	WP5
Positioning					
RTK GNSS receiver, dual-band	2	200	400	buy	WP3
Multi-band GNSS antenna	2	50	100	buy	WP3
Correction radio link (pair)	1	60	60	buy	WP3
IMU, 9-axis	1	25	25	buy	WP3
High-resolution magnetometer	1	40	40	buy	WP3
Global-shutter USB camera	1	60	60	buy	WP3
Time-of-flight distance sensor (one spare)	2	15	30	buy	WP3
Mechanics and test					
Non-metallic pole and mounts	1	80	80	buy	WP3

Item	Qty	Unit	Total	Status	WP
Battery pack and DC-DC converters	1	70	70	buy	WP3
Calibration spheres: steel, copper, aluminium, stainless	1	60	60	buy	WP2
Sandbox: container and dry sand	1	60	60	buy	WP2
Test objects: coins, rings, nails, cans, foil	1	30	30	buy	WP4
Cables, connectors, enclosures	1	100	100	buy	WP1
Purchased equipment			1,780		
Shipping and import allowance (15%)			267		
Contingency on equipment (15%)			307		
Equipment total			2,354		
Optional instruments (if not available to the team)					
LCR meter	1	120	120	optional	WP1
Oscilloscope, used	1	350	350	optional	WP1
Running costs					
Cloud computing for simulation and training	1	150	150		WP4
Data storage (external drives)	1	60	60		WP4
Field-trip transport	1	250	250		WP6

The largest single equipment item is the pair of RTK receivers. Items marked "owned" are already available to the project and are listed at zero cost; their replacement value is about 335 USD and they form part of the team's own contribution.

11.2 Total request

Personnel costs are based on gross monthly salaries in Yerevan for the respective qualifications and include employer contributions. They will be replaced by the actual employment terms of the host organisation when the proposal is prepared.

Table 11.2. Budget request, 12 months. Prices in USD.

Item	Per month	Months	Total
Personnel			
PI · Principal investigator (physicist)	1,800	12	21,600
HW · Hardware engineer (analog and RF)	1,500	12	18,000
SW · Embedded and software engineer	1,600	12	19,200
DS · Data scientist and field engineer	1,500	12	18,000
Personnel			76,800
Overhead: premises, administration, accounting (15% of personnel)			11,520
Equipment including shipping and contingency (Table 11.1)			2,354
Instruments and running costs (Table 11.1)			930
Dissemination: one conference, publication charges			2,500
Contingency (10% of personnel and overhead)			8,832
Total request			102,936

Personnel accounts for about 75% of the request. This is typical for an applied research project: the hardware is inexpensive by design, and the cost lies in the work of building, calibrating, measuring and analysing.

11.3 Cost by work package

Personnel costs are allocated by the effort matrix (Chapter 10), and materials by the work package in which they are first needed.

Table 11.3. Cost by work package, USD.

WP	Title	Personnel	Overhead	Materials	Other	Total
WP0	Project management	3,930	590	0	–	4,520
WP1	Bring-up and receive chain	10,600	1,590	1,178	–	13,368
WP2	Static induction rig and calibration	12,200	1,830	198	–	14,228
WP3	Positioning and timing platform	13,050	1,958	1,144	–	16,151
WP4	Full tensor, dataset and classification	14,150	2,122	250	–	16,522
WP5	Radar subsystem	9,770	1,466	264	–	11,500
WP6	Fusion, field trials and release	13,100	1,965	250	2,500	17,815
	Contingency (not allocated)					8,832
	Total	76,800	11,520	3,284	2,500	102,936

11.4 Cash flow

Table 11.4. Quarterly cash flow, USD. "Other" is contingency and dissemination.

Quarter	Personnel and overhead	Materials	Other	Total	Cumulative
Q1 (weeks 1–13)	22,080	2,520	2,208	26,808	26,808
Q2 (weeks 14–26)	22,080	364	2,208	24,652	51,460
Q3 (weeks 27–39)	22,080	150	2,208	24,438	75,898
Q4 (weeks 40–52)	22,080	250	4,708	27,038	102,936
Total	88,320	3,284	11,332	102,936	

Equipment spending is concentrated in the first two quarters; about 26% of the total is needed in the first quarter.

11.5 Funding programmes

Three kinds of programme fit the project. They differ in scale and in the partners they require.

Table 11.5. Candidate funding programmes. Conditions as published in October 2026; current calls must be checked before application.

Programme	Scope	Size	Conditions
NATO Science for Peace and Security, Multi-Year Project [R19]	Applied research and development with civil application in the security field	€250,000 to €400,000 per project	Joint project with a director in a NATO member country and a co-director in a partner country; Armenia is an eligible partner. A 2027 call has been announced.
Higher Education and Science Committee of Armenia research grants [R20]	National competitive and thematic research programmes	per call	Armenian research organisation as host; formal calls with fixed deadlines
FAST Advance Research Grants [R21]	Research by teams in Armenia	per call	Conditions of the current call
Horizon Europe , e.g. EIC Pathfinder Open [R22]	High-risk research towards new technologies	consortium scale	At least three entities in three countries, including one EU member state; Armenia has been associated since 2021

11.5.1 Strategy

- 1. Primary: NATO SPS Multi-Year Project.** LODESTONE forms the Armenian part of a joint project with a group in a NATO member country that works on electromagnetic induction or ground-penetrating radar. The request in Table 11.2 covers the Armenian work packages and lies well inside the programme's range, leaving room for the partner's share.
- 2. In parallel: national funding.** An application to a suitable HESC or FAST call can fund an earlier start of WP1 and WP2 and strengthens the international proposal with first measured results.
- 3. Later: Horizon Europe.** If the platform performs as planned, a consortium proposal can extend it to field-scale systems and larger datasets.

11.5.2 Steps before submission

1. Identify and contact two or three research groups in NATO member countries working on induction sensing, GPR or humanitarian demining technology.
2. Agree the division of work: the partner typically contributes field-trial sites, reference instruments or simulation expertise.

3. Confirm a host organisation in Armenia that can employ the team and administer the grant.
4. Adapt this report into the programme's proposal template; Chapters 9 to 12 map directly onto the work plan, budget and risk sections.

11.6 Comparison and value

The closest instruments with comparable measurement capability (full-tensor induction classifiers of the MetalMapper type, Chapter 2) cost more than 100,000 USD each. The requested budget pays for the complete research programme, including the team, for about the price of one such instrument, and the results, including the dataset, are published openly.

Risks, safety and quality

Risk register, safety rules and quality assurance.

This chapter lists the main technical and organisational risks, the safety rules for laboratory and field work, and the quality measures that support the results.

12.1 Risk register

Probability (P) and impact (I) are rated from 1 (low) to 3 (high); the score is their product. The register is reviewed monthly (WP0).

Table 12.1. Risk register, sorted by score.

ID	Risk	P	I	Score	Mitigation	Owner
R2	Coil nulling does not reach 10^{-3} at all tones	2	3	6	Adjustable overlap fixture; background subtraction per tone; gradiometer receiver pair	HW
R4	Measured spectra do not match the model at M1	2	3	6	Gate G1 with diagnosis time; independent reproduction; empirical calibration as fallback	PI
R7	Dataset too small for the classification targets	2	3	6	Simulation; object-level augmentation; extension at gate G3	DS
R15	Grant not awarded, or awarded later or smaller than requested	2	3	6	Applications to several programmes in parallel (Chapter 11); a national grant funds WP1 and WP2 first; work packages can start in sequence as funds arrive	PI
R1	Codec input noise or full-scale voltage worse than derived	2	2	4	The $\times 33$ preamplifier is part of the design from the start; lower-noise op-amp option; measured in T1.3	HW
R3	Ground response larger than estimated, reducing range	2	2	4	Measured at two sites early (T2.6); spectral ground fitting; shallower test beds if needed	DS
R5	RTK fix unavailable under trees or near buildings	2	2	4	Open test sites; camera and IMU bridging; wheeled cart with odometry (G2)	SW
R9	Second AD9363 receive channel unavailable or phase unstable	2	2	4	Reference by switched loopback; radar scope reduced at gate G4	HW
R10	Procurement or customs delays	2	2	4	Long-lead items ordered in week 1; local alternatives; slack in WP3 and WP5	PI
R8	Copper and aluminium cannot be separated	3	1	3	Not required by any milestone; achieved separability is reported	PI
R11	Loss of a team member	1	3	3	Two owners per work package; lab and decision logs; contingency for a contractor	PI
R14	Safety incident during field work	1	3	3	Field rules in Chapter 12; only cleared land; safety briefing for every participant	HW
R6	Magnetometer disturbed by the	2	1	2	Mounted at least 50 cm above the coil; heading from RTK track and camera	SW

ID	Risk	P	I	Score	Mitigation	Owner
	transmitter or by steel on site				odometry	
R12	Host cannot sustain six-channel streaming	1	2	2	Both hosts tested in T1.2; reduced sample rate as fallback	SW
R13	Radar emission outside permitted limits	1	2	2	Low power, shielded antennas facing the ground; regulations checked in T5.4	HW

12.2 Safety

12.2.1 Field work and ordnance

LODESTONE is a research instrument and is not a mine or ordnance detector. Several regions have known or suspected contamination with landmines and unexploded ordnance. The following rules apply without exception:

- Field work takes place only on prepared test beds or on land that is known to be free of ordnance, with the written permission of the owner.
- No work is carried out in areas marked or suspected as hazardous, or in their surroundings.
- If an object is found that might be ordnance, work stops, the location is marked from a distance and the relevant authorities are informed. The object is not touched.
- Excavation near archaeological or protected sites requires the permits required by national law; such sites are avoided.

12.2.2 Electrical

- Transmit currents of up to 1 A per coil: connectors and wiring are rated accordingly and the drivers have current limiting.
- Lithium batteries are charged only with suitable chargers and transported in protective bags.
- No phantom power is connected to any coil input.

12.2.3 Radio emissions

The radar sweeps 0.5 to 3 GHz and crosses bands used by mobile networks and satellite navigation. Ground-penetrating radar is a regulated device class in most countries. Transmit power is kept at the minimum needed, antennas are shielded and directed into the ground, and sweeps run only during measurements. The applicable national regulations are checked before field use (T5.4).

12.3 Quality assurance

- **Model checks.** Every analytical model is verified against known limits before use, as done for the sphere solution in Chapter 3.
- **Prediction before measurement.** For each calibration measurement the expected result is written down before the measurement is taken.
- **Calibration objects.** Steel, copper, aluminium and stainless-steel spheres with known dimensions are measured at the start of each measurement day to detect drift.
- **Independent checking.** Results that feed a milestone are reproduced by a second team member from the raw data.
- **Blind evaluation.** Classification is only reported on objects that were not used in training, hidden by a person not involved in the analysis.
- **Traceability.** Every figure of measured data refers to raw files and a tagged software version.

12.4 Open technical questions

1. Actual input noise and full-scale voltage of the CS42448 board (T1.3).
2. Whether 96 kHz sampling is reliable on the chosen host (T1.3).
3. The ground-response factor at the test sites, which sets the detection range (T2.6).
4. Achievable depth accuracy from peak width, which determines whether copper and aluminium can be distinguished (T3.7).
5. Phase behaviour of the AD9363 synthesisers on the available Pluto unit and availability of the second channel (T5.1).
6. RTK availability at the test sites (T3.2).

Appendices

Key numbers, code listings, glossary, references, photographs and document history.

A Key numbers

Values calculated in this report, with the chapter where each is derived. They are replaced by measured values as the work proceeds.

Table A.1. Principal calculated values.

Quantity	Value	Chapter
Target physics		
Quadrature peak, copper sphere, radius 5 mm	1.0 kHz	3
Quadrature peak, aluminium sphere, radius 5 mm	1.6 kHz	3
Quadrature peak, stainless steel sphere, radius 5 mm	41 kHz	3
In-phase sign change, steel sphere ($\mu_r = 100$), radius 5 mm	about 280 kHz	3
Quadrature peak width, ring / solid sphere	1.13 / 1.56 decades	3
Signal change per doubling of depth	$\div 64$	3
Induction subsystem		
Tone band at 48 kHz	188 Hz – 18.75 kHz	5
Crest factor, zero / optimised phases	4.90 / 3.06	5
SNR penalty, simultaneous vs stepped tones (peak-limited)	6.7 dB	5
Receive coil: inductance, resistance, self-resonance	4.58 mH, 12.8 Ω , 429 kHz	5
Codec input noise (derived)	about 80 nV/ $\sqrt{\text{Hz}}$	5
Chain noise with $\times 33$ OPA1612 stage	2.6 nV/ $\sqrt{\text{Hz}}$	5
Echo, 1 cm copper sphere at 20 cm	20.5 μV	5
Detection range, coin-sized target (ground-limited)	about 40 cm	5
Host data rate, 6 channels	9.2 Mbit/s	4
Radar subsystem		
Range resolution in soil ($\epsilon_r = 9$, 2.5 GHz bandwidth)	2.0 cm	6
Unambiguous range in soil, 10 MHz steps	5 m	6
Positioning		
Peak width over a small target	$0.90 \times \text{depth}$	7
Signal change per 1 cm height change at 20 cm	about 30%	7
IMU position error, 1 mg bias, after 2 / 5 / 10 s	2 / 12 / 49 cm	7
Transmit-coil field at 50 cm / 1 m	1.45 / 0.18 μT	7
Coil offset per degree of tilt, 1.5 m pole	2.6 cm	7
Timing resolution of PPS in audio stream	20.8 μs	7

B Code listings

Both listings are starting points for tasks T1.6 and T5.3 and have not yet been run on the hardware.

B.1 Induction acquisition (48 kHz, codec as sound card)

```
import numpy as np, sounddevice as sd

FS, N = 48_000, 4096                                # 11.72 Hz bins, 85.3 ms blocks
BINS = np.array([16, 24, 37, 56, 85, 130, 197, 300, 456, 693, 1053, 1600])
PH = np.array([2.589, 5.287, 1.692, 2.610, 3.376, 2.940,
               1.805, 1.120, 0.966, 3.592, 5.042, 0.208])

n = np.arange(N)
block = np.cos(2 * np.pi * np.outer(n, BINS) / N + PH).sum(axis=1)
block *= 0.5 / np.abs(block).max()
tx = np.tile(block, 60)                             # about 5.1 s

# OUT 1 -> amplifier -> TX coil; IN 1 = TX current, IN 2-4 = RX coils, IN 5 = PPS
# the device name depends on the host driver
rec = sd.playrec(tx, FS, device="CS42448",
                 input_mapping=[1, 2, 3, 4, 5], output_mapping=[1])
sd.wait()

blocks = rec[N * 4:].reshape(-1, N, 5)               # drop start-up
spec = np.fft.rfft(blocks[:, :, :4], axis=1)[:, BINS, :]
H = (spec[:, :, 1:] / spec[:, :, :1]).mean(axis=0)   # 12 tones x 3 receivers
pps = np.flatnonzero(np.diff(rec[:, 4]) > 0.2)       # PPS edges, sample index
```

B.2 Radar sweep with phase reference

```
import adi, numpy as np

sdr = adi.ad9361("ip:192.168.2.1")          # Pluto in two-channel mode
sdr.rx_enabled_channels = [0, 1]           # RX1 echo, RX2 coupled reference
sdr.sample_rate = int(2e6)
sdr.rx_buffer_size = 4096
sdr.tx_cyclic_buffer = True

fs, f_if, N = sdr.sample_rate, 100e3, sdr.rx_buffer_size
tone = np.exp(2j * np.pi * f_if * np.arange(N) / fs)
sdr.tx(0.5 * tone * 2**14)

def phasor(x):
    return np.vdot(tone, x) / N

S = []
for f in np.arange(500e6, 3000e6 + 1, 10e6): # 251 steps
    sdr.tx_lo = sdr.rx_lo = int(f)
    for _ in range(2):
        sdr.rx()                                # let the synthesisers settle
    echo, ref = sdr.rx()
    S.append(phasor(echo) / phasor(ref))

S = np.array(S)                                # subtract a background sweep he
profile = np.abs(np.fft.ifft(S * np.hanning(len(S)), n=4096))
```

C Glossary

Term	Meaning
Crest factor	Ratio of peak to RMS value of a waveform
EKF	Extended Kalman filter
EMI	Electromagnetic induction (as a sensing method)
FFT bin	One frequency point of a discrete Fourier transform; spacing = sample rate / block length
GPR	Ground-penetrating radar
IMU	Inertial measurement unit (accelerometer and gyroscope)
Polarizability tensor	3×3 matrix relating the applied field to the induced magnetic dipole of a target
PPS	Pulse per second, a timing output of GNSS receivers aligned to GPS time
PPK / RTK	Post-processed / real-time kinematic satellite positioning, centimetre level
RACI	Responsible, accountable, consulted, informed
SDR	Software-defined radio
SFCW	Stepped-frequency continuous wave
SRF	Self-resonant frequency of a coil
TDM	Time-division multiplexed serial audio interface
ToF	Time-of-flight distance sensor
UXO	Unexploded ordnance
VIO	Visual-inertial odometry
VLF	Very low frequency; here, continuous-wave induction detectors at 3 to 30 kHz

D References

1. [R1] Software defined radio for stepped-frequency, ground-penetrating radar. Thesis, Georgia Institute of Technology. repository.gatech.edu
2. [R2] Software defined radio for stepped-frequency, ground-penetrating radar. IEEE conference paper. IEEE Xplore 8128082
3. [R3] Software-defined ground penetrating radar using COTS hardware. GNU Radio Conference 2025. events.gnuradio.org

4. [R4] Ground-penetrating radar with HackRF software-defined radios. rtl-sdr.com
5. [R5] MetalMapper: a multi-sensor TEM system for UXO detection and classification. ResearchGate 235186078
6. [R6] Dipole models for UXO discrimination at live sites. SERDP/ESTCP project. serdp-estcp.mil
7. [R7] The spectral properties of the magnetic polarizability tensor for metallic object characterisation. arXiv:1906.00382
8. [R8] Characterisation of multiple conducting permeable objects in metal detection by polarizability tensors. arXiv:1809.08916
9. [R9] Detecting landmines for a safer world. Ingenia. ingenia.org.uk
10. [R10] MDS-10 dual sensor detector. Minelab. minelab.com
11. [R11] MDS-20 dual-sensor detector. EPE. epequip.com
12. [R12] Combining electromagnetic spectroscopy and ground-penetrating radar for the detection of anti-personnel landmines. Sensors 19(15):3390, 2019. mdpi.com
13. [R13] Deep learning-based classification of anti-personnel mines and sub-gram metal content in mineralized soil (DL-MMD). Scientific Reports, 2024. nature.com
14. [R14] Detection and classification of landmines using machine learning applied to metal detector data. ResearchGate 339640290
15. [R15] Landmine identification from pulse induction metal detector data using machine learning. ResearchGate 373286697
16. [R16] J. R. Wait, A conducting sphere in a time varying magnetic field. Geophysics 16(4), 666–672, 1951.
17. [R17] Cirrus Logic, CS42448 108 dB, 192 kHz 6-in, 8-out TDM CODEC, datasheet.
18. [R18] Analog Devices, ADALM-Pluto SDR active learning module and AD9363 documentation.
19. [R19] NATO Science for Peace and Security Programme: overview and Multi-Year Projects. nato.int; programme size and partner rules: Science|Business
20. [R20] Higher Education and Science Committee of Armenia, research grants. hesc.am
21. [R21] Foundation for Armenian Science and Technology, Advance Research Grants. fast.foundation
22. [R22] European Commission, list of participating countries in Horizon Europe. ec.europa.eu

E Photographs

Photographs of the hardware are added as it is built. To add one, copy a file with the name shown into the `images/` folder of the web server; it appears here without any other change.

CS42448 codec board

`images/6in8out.jpg`

Codec board (6IN8OUT).

Radar bench set-up

`images/pluto-setup.jpg`

ADALM-Pluto with amplifiers.

Loopback test

`images/codec-loopback.jpg`

T1.2 loopback.

Receive coil

`images/rx-coil.jpg`

Receive coil with shield.

Coil head

`images/coil-head.jpg`

Transmit and receive coils.

Sensor pole

`images/pole-rig.jpg`

Positioning rig.

Test bed


`images/sandbox.jpg`

Sandbox with calibration targets.

First measured spectrum

`images/first-spectrum.png`

Compared with Chapter 3.



Compared with Chapter 6.

F Document history

Revision	Date	Changes
1	September 2026	Design notes published as seven separate web pages.
2	26 September 2026	Consolidated into this report. Detailed work plan, team structure and budget added. Tone plan and noise budget revised for the CS42448 codec at 48 kHz. Radar sweep set to 0.5–3 GHz to match the antenna size.
3	7 October 2026	Project planned for grant funding only; self-funded option removed. Funding programmes and strategy added (Chapter 11), funding risk added. Summary figures added to the cover page.