

WORLD ROBOT OLYMPIAD 2026

FUTURE INNOVATORS · SENIOR

WRO

STRATASCAN ON

Reading the past, layer by layer

Reading the past, **layer by layer.**

A non-invasive robotic system that estimates the age of buried structures from soil-moisture gradients — so archaeologists know where to dig before the first spade goes in.

SEASON THEME

Recording & Protecting Cultural Places

Project report · 20 counted pages · English

Team Sensor ON

Cîmpian Darius · Șandru Anisia · Paici Ștefan

Inventika Place · Bucharest, Romania

STRATASCAN ON

Reading the past, layer by layer

Powered by

stratascanon.inventikaplace.ro

The live dashboard — open it while you read this report.

CONTENTS

Table of contents

Team presentation	Who we are, our roles, eight seasons together	01
Summary project idea	The problem, the solution and why it matters	02

THE ROBOTIC SOLUTION

How the idea was born	From a shared interest to a defined problem	03
Research: what already exists	LiDAR, GPR, ERT — and the gap between them	04
Five iterations	What each version taught us	05
Iteration 5 — leaving LEGO behind	Why we rebuilt on the SparkFun XRP	06
Field research at Jidava	Two days on a Roman fort with archaeologists	07
From moisture to age	Topp's equation, the drying curve and the time constant	08
The sensor suite	Every sensor, what it measures, why it is there	09
Mechanics of the probe	How the probe gets into the ground, and how that changed	10
3D design and printed parts	Designing our own attachments for the XRP	11
Software architecture	MicroPython, the scan loop and autonomous decisions	12
The StrataScan ON application	Live telemetry and geolocated readings, in any browser	13
Challenges, failures and AI use	What broke, what we fixed, what AI did	14

SOCIAL IMPACT & INNOVATION

Impact on society	Who benefits — and the risks we take seriously	15
Three sites, three time constants	Our calibration set, and what it does not prove	16
A tried, practical use case	The rural settlement, where nothing is mapped	17
Cost structure	What one unit costs to build and to bring to market	18
Revenue streams and market	Who buys this, and at what price	19
Resources, partners and next steps	What the project rests on, and what follows	20

Page count. This report contains exactly 20 counted pages, numbered 01–20 in the bottom-right corner. The cover, this table of contents and the list of sources are not counted, in line with the WRO 2026 Future Innovators General Rules (section 6.5). Section limits: team 1 · summary 1 · robotic solution 12 · social impact & innovation 6.

TEAM PRESENTATION

Three students, eight seasons, one team name

We are **Sensor ON**, a team of three high-school students from **Inventika Place**, a robotics club in Bucharest, Romania. The team was founded in 2018 and named by Ștefan, who was nine at the time and wanted a name that said what the robots actually did: they sensed things. Eight of us have passed through the team since; two of us have been here the whole way.



Cîmpian Darius

18 years old

MECHANICS DESIGN

Chassis, drivetrain and the probe-insertion tower. Designed the rack-and-pinion mechanism and led the move to a 3D-printed structure.



Șandru Anisia

16 years old

RESEARCH AND MARKETING

Soil science, the calibration protocol and the measurement dataset. Owns contact with the archaeologists, the brand and the business model.



Paici Ștefan

15 years old

PROGRAMMING AND ELECTRONICS

MicroPython firmware, sensor wiring and the I²C bus, the vision pipeline, the Bluetooth protocol and the dashboard.

Eight years, in nine photographs

Left to right: Ștefan and Darius in 2018, the year the team was named, through to the national stage of WRO Romania in June 2026.



2018



2018



2019



2020



2021



2022



2024



2025



2026

Where the project comes from. StrataScan ON did not start this season. The first prototype was built for a different robotics programme — **FIRST® LEGO® League**, whose 2025–26 season was archaeology-themed — where it won the Innovation Project award at the regional stage in February 2026. We declare this openly because the WRO rules ask teams to say when a project builds on earlier work, and because what we bring to Ontario is not that robot. The controller, the language, the structure, the sensor suite, the insertion mechanism and the data pipeline have all changed. Pages 05 and 06 set out exactly what carried over and what did not.

How we divide the work. We work in weekly build sessions at the club plus asynchronous work between them. Each of us owns a domain and is accountable for it, but no decision that affects the whole robot is taken alone — mechanics, firmware and measurement protocol are too tightly coupled for that.

Field days are the exception: at Jidava all three of us did the same work — pegging out the grid, drilling pilot holes, sieving soil and recording readings — because that is how the archaeologists work, and we wanted the data collected the way they would collect it.

SUMMARY PROJECT IDEA

Scan first. Dig only where it matters.

The problem. Excavation is slow, expensive and irreversible. An archaeological team gets one chance at a layer: once soil is removed, the stratigraphy that dates it is gone. Yet a large share of every season is spent opening trenches that turn out to be empty, because until you dig you rarely know whether what lies below is a Roman wall, a modern pipe, or nothing at all. In Romania alone there are hundreds of recorded sites with almost no funding, where a wasted week is a season lost.

Our solution. StrataScan ON is an autonomous ground robot that surveys a marked grid and measures how the soil holds and releases water at each cell. Compacted soil above a buried structure dries more slowly than the ground around it, so the *rate* of drying carries a readable signature of what is underneath — and of how long it has been there. The robot drives the grid, pushes a probe into the ground at each node, reads moisture, soil temperature, air conditions and GPS position, and streams the result to a dashboard. Nothing is excavated.

0 cm	173	3	~$1,050
excavation required to produce a reading	field measurements logged in a single session	sites measured, including one known to be empty	bill of materials for one complete unit



Working the grid at the Roman fort of Jidava, August 2026.

The value, if this were real. A survey that today costs a team days of digging and a permanent loss of stratigraphy becomes an afternoon of scanning and a map. The decision "is this worth opening?" moves from intuition to evidence.

The people who gain most are not the well-funded excavations — they already have geophysics budgets. It is the rural site, the rescue dig before a road, the county museum with one archaeologist and no equipment. Our target price puts a unit inside a small institutional budget rather than a research grant.

Why it matters to us. We are not proposing to replace archaeologists. We are proposing to give them one more thing to look at before they commit a season to a hole in the ground.

OUR SLOGAN

Reading the past, layer by layer.

SEE IT RUNNING

stratascanon.inventikaplace.ro

WHY THIS PROJECT EXISTS

A LiDAR survey tells an archaeologist where something is buried. Nothing on the market tells them, without digging, roughly how long it has been there. That single missing number is what StrataScan ON measures — and it is the number that decides whether a trench is worth opening.

HOW THE IDEA WAS BORN

We wanted to build for archaeology. We had to find the right problem.

The starting point was interest, not a problem statement: all three of us liked archaeology and wanted to point robotics at it. That is a bad way to begin, and it showed. Our first ideas were solutions looking for a use.

Ideas we investigated and dropped. A robot that cleans and catalogues finds after excavation — already done, and done better, by conservation labs. A drone that photographs trenches for photogrammetry — commercial products exist and are cheap. A sensor that detects metal objects — a metal detector, reinvented, worse.

Each of these sat at the *end* of the archaeological chain, where the problems are well funded and largely solved. What changed our direction was asking a professional where the money actually goes.



Consultation with Dr. Florian Matei-Popescu, "Vasile Pârvan" Institute of Archaeology



THE CONVERSATION THAT SET THE PROBLEM

In February 2026 we met **Dr. Florian Matei-Popescu**, a specialist in the Roman period at the "Vasile Pârvan" Institute of Archaeology in Bucharest. We asked what wastes the most time on a dig. The answer was not conservation or recording. It was opening ground that turns out to be empty — and the fact that once you have opened it, the evidence that would have dated it is destroyed whether you found something or not. He also told us the honest limit of the idea: soil moisture will never give you a calendar date. It can give you a *relative* signal — this ground has been undisturbed and compacted for a long time, that ground has not. For deciding where to dig, relative is enough.

We also checked the physics with someone outside archaeology. Adriana Sfetcu, a horticulturist, confirmed that compaction changes both water retention and the rate of evaporation in a way that is measurable with ordinary agricultural probes — the same principle used to map soil compaction in farming. That was the point at which the project stopped being an idea and became a measurement problem with a known instrument.

The problem statement we settled on

Given a piece of ground that a team is considering excavating, produce — without digging — a map showing where the subsurface behaves anomalously, and an indication of whether that anomaly is old or recent, at a cost a county museum can afford.

WHAT CHANGED OUR DIRECTION

We did not start from a robot and look for a use. We started from a professional telling us what wastes his season, and built backwards from that sentence. Every design decision in the rest of this report traces to it.

RESEARCH INTO SIMILAR IDEAS

What already exists, and the gap we found

Archaeological prospection is a mature field. We did not expect to find empty space, and we did not. What we found is that every established method answers **where** and **how deep**, and none of them answers **how long has it been there** — which is the question that decides whether a trench is worth opening.

Method	What it measures	Answers "where?"	Answers "how old?"	Cost to a small team
LiDAR (airborne)	Surface micro-relief under vegetation	Yes — plan of earthworks	No	High; usually contracted
GPR — ground-penetrating radar	Dielectric contrasts vs. depth	Yes — depth and shape	No	€15,000–€40,000 per unit
ERT — electrical resistivity	Bulk resistivity of the soil column	Yes — sections	Indirectly, poorly	€8,000–€25,000
Magnetometry	Local magnetic anomalies	Yes — fired features, ditches	No	€6,000–€20,000
Soil coring / test pits	The soil itself	Yes	Yes — but destructive	Cheap, slow, irreversible
StrataScan ON	**Moisture retention and drying rate**	**Coarse — grid anomalies**	**Relative, non-destructively**	**Target < €1,500**

The closest prior art. Soil-moisture mapping is used in precision agriculture, and moisture contrast is a known confounder in archaeological GPR — practitioners wait for the right soil conditions precisely because water content changes what they see. What we could not find was anyone treating the moisture field itself as the primary archaeological signal, on a low-cost autonomous platform, aimed at the pre-excavation decision.

We take that carefully: absence of a published prototype is not proof of novelty, and we say so to judges. What we can defend is that the combination — grid autonomy, penetrating probe, drying-curve capture and a sub-€1,500 target — did not appear in anything we found.

COMPLEMENTARY, NOT COMPETING

We do not think StrataScan ON replaces geophysics. The workflow we propose runs the other way round:

1 · LiDAR or GPR maps *where* anomalies are and how deep.
2 · StrataScan ON visits those anomalies and asks whether the ground above them has been undisturbed long enough to be worth the trench.
3 · The team digs the cells that pass both filters.

On a site with no geophysics budget at all — which is most of them — step 2 becomes step 1, and a coarse moisture map is still infinitely more than nothing.

THE GAP, IN ONE LINE

Every established prospection method answers **where** and **how deep**. None of them answers **how long**. We are not trying to beat GPR at its own question — we are answering the one nobody else is answering, at a price that puts it in reach of the teams who have no geophysics at all.

EVOLUTION OF THE PROJECT

Five iterations, and what each one taught us

i1

Proof of concept

LEGO EV3 · a moisture sensor taped to a moving arm

Could a moving robot read soil at all? Yes — but readings drifted with contact pressure. Won the Innovation Project award at a regional competition in February 2026.



i2

Repeatability

LEGO EV3 · fixed-depth probe holder, manual grid

Fixing insertion depth removed most of the drift. We learned that the measurement is only as good as the mechanical contact — a lesson that drove every later design.

i3

Automation

EV3 + Arduino Uno · Cartesian X–Y scanning frame

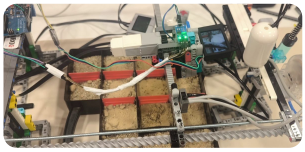
The first version that surveyed a grid unattended. Two lead-screws, an Arduino reading the analogue probe, data stored in an X–Y matrix and rendered as heatmaps.

i4

Perception and rigidity

EV3 · HuskyLens AI camera, 4 mm steel rods, MicroPython

Vision to avoid stones and respect no-go markers; steel rods with lock nuts replaced flexing plastic axles. This is the robot that placed first at the national stage in June 2026.

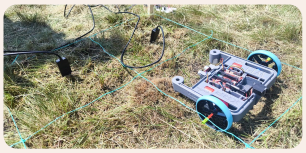


i5

A real robot

SparkFun XRP (RP2350) · full sensor suite, GPS, own 3D-printed structure

The current build. We stopped adapting an education kit and started designing a machine: proper I²C bus, geolocated readings, rack-and-pinion insertion and a wireless data pipeline.



What actually carried over from i1 to i5. Three things: the hypothesis, the grid protocol, and the calibration routine. Nothing physical. The chassis, controller, language, probe mount, insertion mechanism and data path have all been replaced at least once. We think that is the honest answer to "is this the same project?" — the question is the same, the machine is not.

What we would do differently. We spent iterations 1 to 4 inside a LEGO ecosystem because it was what the club had and what we knew. In hindsight the constraint that hurt most was not stiffness — it was the sensor interface. Reading one analogue probe through a second microcontroller and a serial link cost us a season of noise we did not need to have.

WHY FIVE ITERATIONS

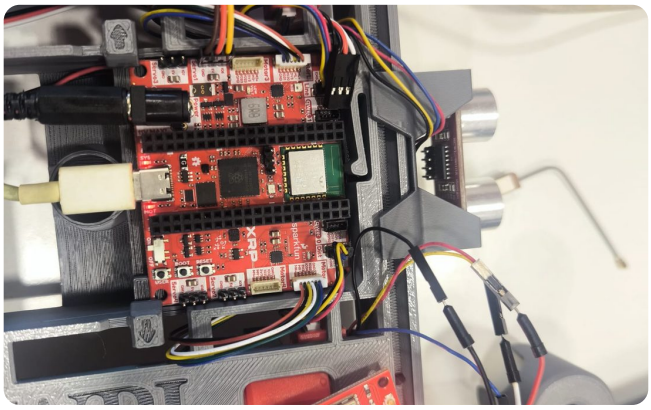
Five rebuilds in one year is not a sign that we could not decide. Each one closed a specific measurement error we could name: contact pressure, insertion depth, grid drift, terrain blindness, and finally the sensor interface itself.

ITERATION 5

Why we rebuilt the robot on the SparkFun XRP

Iteration 4 placed first nationally, and it was still the wrong machine. Every sensor we wanted to add had to be bolted onto a controller that was never designed to host it: one analogue input, borrowed through an Arduino Uno over serial, no I²C bus we could trust, no way to know *where on Earth* a reading had been taken.

The **SparkFun XRP Controller** — an RP2350-based board built for exactly this kind of student robot — solved all four at once. It runs MicroPython natively, exposes a Qwiic I²C bus, has dedicated motor and servo drivers, and leaves eight GPIO free for everything else.



NEO-6M GPS with its external antenna — the sensor that made readings geolocated

	Iteration 4 — LEGO EV3	Iteration 5 — SparkFun XRP
Controller	EV3 brick + Arduino Uno over serial	XRP Controller, RP2350 dual-core
Language	MicroPython on EV3, C++ on the Arduino	MicroPython throughout, one codebase
Analogue inputs	1, via the Arduino	4 usable ADC channels on-board
Sensor bus	None — bit-banged, unreliable	Qwiic I²C with hardware pull-ups
Position awareness	Dead reckoning from wheel counts only	GPS fix + IMU + wheel encoders
Structure	LEGO beams, later 4 mm steel rods	XRP chassis + our own printed PLA parts
Data out	USB cable, offline CSV	Bluetooth LE to a live web dashboard
Mass / footprint	≈ 2.4 kg, 34 × 30 cm	≈ 1.3 kg, 22 × 19 cm

THE HONEST TRADE-OFF

We lost things too. The EV3 build had far more torque and could push its probe into hard ground that stops the XRP — which is why one LEGO part deliberately stayed: an EV3 large motor now drives the insertion tower, as a geared actuator rather than a controller (page 10). The EV3 also had a wider wheelbase and was more stable on a slope, and it was, frankly, more finished — iteration 5 is the machine we are still building, with the 3D-printed structural parts being revised as we write. We chose the platform we can keep developing.

THE DECISION

Iteration 4 won the national stage. We rebuilt it anyway, because a robot that cannot tell you where a reading was taken is producing numbers, not data.

FIELD RESEARCH

Two days on a Roman fort, with the people who would actually use this

On **12 and 13 August 2026** we took the robot to **Jidava**, a Roman auxiliary fort near Câmpulung Muscel in Argeș County, and worked alongside **arch. Ion Dumitrescu** and **arch. dr. Florian Matei-Popescu** of the "Vasile Pârvan" Institute of Archaeology.

We chose Jidava deliberately. It is a site where the subsurface is already documented — the walls, their depth and the fort's occupation dates are published. That makes it the one kind of place where you can check whether a robot is telling the truth, because you already know the answer.

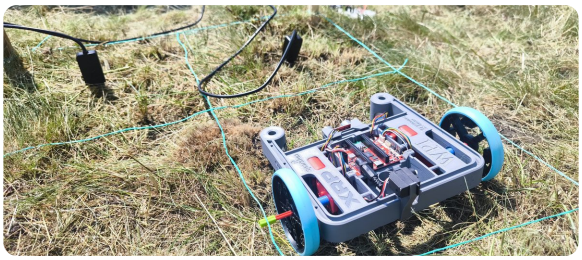


The excavated walls of the fort, looking across the site.

Zone 1 — ≈ 15 cm

Administrative buildings of the nearby mine

Roughly one hundred years old. A control at the shallow end: recent, known, and still a buried structure. If the method cannot separate a century-old foundation from undisturbed ground, it cannot do anything harder.



Zone 2 — ≈ 30 cm

Defensive walls of the Roman fort

Occupied between AD 300 and 358. Seventeen centuries of compaction above stone, in an open section where the archaeologists could show us exactly what was underneath our grid.



WHAT THE ARCHAEOLOGISTS TOLD US

"It would be genuinely useful for us."

Their qualification mattered more than the compliment. They said the method would have the greatest value not on known fortifications like Jidava — where everyone already knows where to dig — but on **rural settlements**, where nothing is mapped, the layout follows no standard plan, and a team can spend a season finding the edge of a site. They also warned us that soil type dominates the signal, so a calibration taken on one site cannot be carried to another. We have built that limitation into how we describe the product.

WHAT THE FIELD DAY WAS FOR

Two days at Jidava taught us more than four months in the workshop. Not because the robot performed well — it performed adequately — but because the people who would use it told us which half of our assumptions were wrong.

Three steps: a percentage, a curve, a constant

The chain from a raw voltage to an estimate of age has three distinct stages, and it is worth separating them, because only the first is textbook physics. The second is a standard curve fit. The third is the part we are still building, and the part we are most careful about claiming.

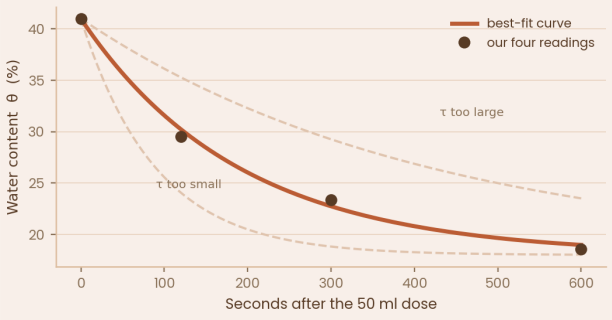
1 Permittivity → water content, with Topp's equation

The probe measures the **dielectric permittivity** of the soil around its blade, not water. Water has a permittivity near 80, dry mineral soil about 4, air 1 — so permittivity is dominated by water content. Topp, Davis and Annan (1980) give the empirical polynomial that converts one into the other, and it is what turns a meaningless ADC value into the concrete percentage the sensor shows us.

$$\theta = -5.3\times10^{-2} + 2.92\times10^{-2}\,\varepsilon_r - 5.5\times10^{-4}\,\varepsilon_r^2 + 4.3\times10^{-6}\,\varepsilon_r^3$$

θ = volumetric water content · ε_r = relative dielectric permittivity · valid for mineral soils

2 Water content → a drying curve, and a time constant τ



Four readings, and the exponential that fits them best.

One reading tells us how wet a cell is today, which depends mostly on last week's weather. What we actually want is how *fast* it dries. So the robot applies a controlled 50 ml dose and reads the same cell at fixed intervals — currently at 0, 2, 5 and 10 minutes.

Those points are fitted to an exponential decay. The parameter we are solving for is the **time constant** τ: the robot searches candidate values of τ, computes the sum of squared residuals against our four points, and keeps the τ with the smallest error. A binary search over the plausible range converges in around a dozen evaluations, which the RP2350 does between one cell and the next.

```
theta(t)  =  theta_inf  +  (theta_0 - theta_inf) * exp(-t / tau)
```

fit_decay() returns the tau minimising the residual against the four sampled points

3 τ → age, which needs a calibration set we do not yet have

Compacted ground releases water more slowly, so τ grows with the length of time the soil has been undisturbed. Turning τ into a number of years requires a function fitted to ground of *known* age — and that function is only as good as the number of known points behind it. We currently have three, described on page 16. Three points give a direction, not a curve, and we say so to judges rather than presenting an age estimate we cannot defend.



The protocol in the field: peg the grid, string the cells, read — and compare against what the excavation already shows.

THE SENSOR SUITE

Every sensor, what it measures, why it is there

Sensor	Measures	Interface	Pins on the XRP	Why it is on the robot
SEN0308 capacitive probe ×2	Soil dielectric permittivity	Analogue	Line port — GPIO44 / GPIO45	The primary instrument. Two probes let us take a differential reading and reject a stuck sensor.
DS18B20	Soil temperature at probe depth	1-Wire	GPIO12 + 4.7 kΩ pull-up	Permittivity is temperature-dependent. Without this the moisture figure drifts by several percent across a hot afternoon.
BME280	Air temperature, humidity, pressure	I²C (0x76)	Qwiic — GPIO4 / GPIO5	Evaporation rate depends on the air above the soil. Needed to compare drying curves between morning and afternoon.
NEO-6M + external antenna	Latitude, longitude, time	UART, 9600 baud	GPIO16 / GPIO17	Makes a reading a point on Earth rather than a cell in an anonymous grid, and lets a site be re-surveyed years later.
HuskyLens	Terrain classification, markers	UART / I²C	GPIO13 / GPIO14	Detects stones, surface cracking and red no-go markers before the probe descends.
On-board IMU (LSM6DSO)	Attitude, angular rate	I²C (0x6B)	Internal	Keeps the probe vertical on slopes and flags a reading taken while the chassis was tilted.
Wheel encoders ×2	Odometry	Quadrature	Motor ports L / R	Grid navigation between GPS fixes, which are too coarse on their own at 30 cm cell spacing.
HC-SR04 ultrasonic	Obstacle range	Digital	Distance port — GPIO0 / GPIO1	Safety stop. A trench edge is a two-metre drop and the robot has to see it.

CONTROLLER

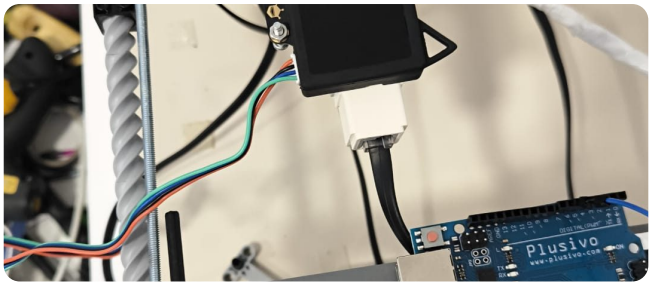
SparkFun XRP Controller — RP2350 dual-core at 150 MHz, 520 kB SRAM, MicroPython. Two DC motor drivers with encoder inputs, four servo headers, two Qwiic I²C connectors, on-board IMU and 2.4 GHz radio for the Bluetooth link.

We use 8 of the 12 free GPIO. The remaining four are deliberately left open — every previous iteration ran out of pins before it ran out of ideas.

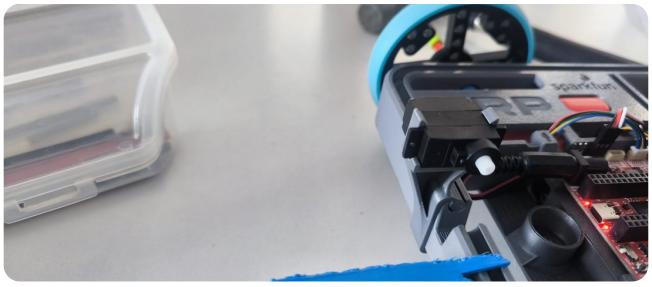
POWER — SEPARATE RAILS

The insertion drive draws far more current than the controller's servo rail can supply, so it runs from its own pack; only the **signal** wire goes to the board and the grounds are bonded.

We learned this the hard way — for two weeks a servo buzzed and the board reset under load, and the cause was a floating ground, not a faulty servo.



Sensor wiring and the probe head during assembly.


The printed MG996 mount on the XRP chassis rails.

MECHANICS OF THE PROBE

Getting a sensor into the ground, 200 times, to exactly the same depth

This is the part of the project that has been rebuilt most often, because it decides whether the data is worth anything. A capacitive probe reads the soil in a small volume around its blade. If the blade is 2 cm deep in one cell and 4 cm in the next, the two readings are not comparable — and the difference will look exactly like an archaeological anomaly.

ITERATION 4

Pneumatic insertion

A small pneumatic cylinder, charged by an on-board pump, delivered a short high-force stroke that drove the blade through the crust. It worked, and it proved that force — not travel — was the limiting factor.

What it cost us: an air reservoir, a pump, tubing and a valve, all of which take space on a small chassis; and a charging step that in practice someone had to supervise. A survey that needs a human standing next to it is not an autonomous survey.

ITERATION 5 — CURRENT

Motor-driven rack and pinion

The pneumatics are gone. A pinion drives a printed rack on a vertical rail, and the sled carrying the probe descends under motor control, with the depth commanded in millimetres rather than hoped for.

It is **more compact** — no reservoir, no pump, no tubing — and it runs **without any human intervention** between cells. For the torque the pneumatic stroke used to supply, we are **fitting a LEGO EV3 large motor back onto the rack**. We have not abandoned the EV3: it returns in the one role where it is still the best part we own — a geared actuator that survives being stalled against hard ground while it drives the probe in.

Status as we write: the printed housing that carries the

1 Depth is commanded, not dropped

Encoder counts on the pinion give closed-loop position over 60 mm of travel. A current limit detects the moment the blade meets resistance, so insertion always starts from a known surface contact rather than a fixed height above uneven terrain.

2 Archimedes screws ×2

Two printed augers clear loose material from the entry point before insertion, so the probe meets undisturbed soil instead of the spoil it just displaced. Carried over from iteration 4, where they were the single change that most improved repeatability.

3 Every reading is a coordinate

The NEO-6M fixes each cell to a latitude and longitude, and the reading is pushed with its coordinates to the dashboard. A survey stops being a grid on a piece of paper and becomes a dataset that another team can return to in five years and repeat.

4 Levelling

The IMU holds the tower vertical to within 3° and flags any reading taken outside that window, because a probe entering at an angle samples a different volume of soil.

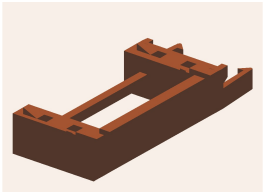
ONE CELL, START TO FINISH

The cell cycle, end to end: arrive · level · vision check · clear · insert to commanded depth · dose 50 ml · read at 0, 2, 5 and 10 minutes · fit τ · retract · log with GPS · move. About eleven minutes per cell in drying-curve mode, or forty seconds in single-reading mode when a team wants a fast coarse map first.

3D DESIGN AND PRINTED PARTS

We stopped buying brackets and started designing them

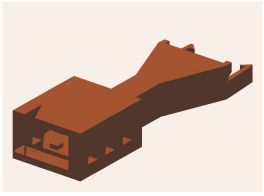
The XRP chassis is a flat plate with a standard slot pattern. Nothing we needed to mount — an MG996 servo, a BME280 on a stalk clear of the chassis heat, a wheel that accepts a LEGO axle — existed as a part. So we modelled them ourselves and printed them in PLA on a Prusa MK2S with a 0.6 mm nozzle. **This work is still in progress:** the parts below are printed and fitted, while the **housing that carries the probe tower** — the part that also has to take the EV3 drive for the rack — is on the printer as we write.



Servo mount, v3

`71 × 27.6 × 12 mm`

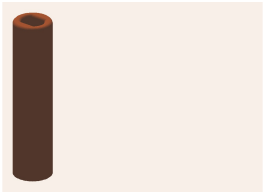
Holds the MG996 above the chassis with the output shaft on the tower axis. Three revisions: v1 clamped from the side and fouled the cable; v2 extended the chassis prongs and no longer fitted; v3 leaves the original prongs untouched and cradles the servo from below.



BME280 arm, v2

`91.5 × 27.6 × 17 mm`

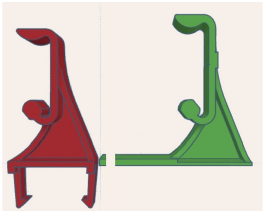
Lifts the air sensor 90 mm clear of the electronics so it reads ambient air, not the board's own waste heat. v1 was unprintable — the slicer flagged floating bridges and poor bed adhesion; v2 sits on a continuous base plate with 1,548 mm² of bed contact and vent slots along the arm.



Wheel with LEGO axle

`Ø10 × 42.3 mm boss`

A cross-axle socket added to the standard XRP wheel so the Archimedes screws and probe linkage from our earlier LEGO builds still fit. One stud of insertion depth, which is enough to transmit the torque involved and lets us keep parts we already own.



Probe-arm bracket

`two variants, in trial`

The newest part, and the reason page 10 calls the build open. Both carry the same hooked arm and triangular web; they differ only in how they meet the chassis — one snaps onto the rails, the other bolts flat. We print both and keep whichever survives a field day.

HOW WE MODELLED THEM

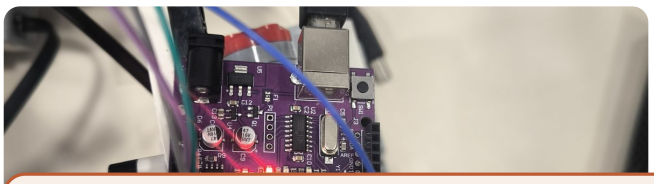
We did not use a CAD package. Every part is generated by a short Python script that emits an STL directly — boxes, cylinders and a triangle-clipping routine for the cut-outs. It is an odd choice, but it means each part is a parametric, version-controlled file: changing a slot width is one number, and the change is visible in the diff.

Scripts: `mg996c.py`, `bme_arm2.py`, `add_lego.py`.

THE MISTAKE WORTH RECORDING

Version 3 of the servo mount looked correct on screen and had a small cube floating in mid-air in the slicer preview. The cause: our screw slots cut across the full wall width, splitting the wall into three disconnected strips.

We now run a connectivity check on every STL before printing — a winding-number voxel test that confirms the mesh is a single solid. Our first version of that check used an even-odd rule and silently passed overlapping solids.





WHY WE PRINT OUR OWN

Designing our own parts was not about saving money — the brackets cost pennies either way. It was about being able to change the machine on a Saturday instead of waiting a week for something that almost fits.

CODING THE SOLUTION

One language, one codebase, four autonomous decisions

Iteration 5 runs **MicroPython** on the RP2350, end to end. Iterations 1–4 split logic between an EV3 in one language and an Arduino in another, connected by a serial link that was the source of most of our bugs. Collapsing that into a single codebase removed a whole class of failure and made the drying-curve timing reliable for the first time.

READING THE PROBE

```
from machine import ADC, Pin, I2C, UART
import time, math

probe_a = ADC(Pin(44))      # Line L
probe_b = ADC(Pin(45))      # Line R
i2c     = I2C(0, sda=Pin(4), scl=Pin(5))
gps     = UART(1, 9600, tx=Pin(16), rx=Pin(17))

AIR, WATER = 47200, 19800   # per-probe calibration

def permittivity(raw):
    # linear map of raw ADC to eps_r, then
    # Topp (1980) for volumetric water content
    f = (AIR - raw) / (AIR - WATER)
    return 1.0 + f * 79.0

def theta(eps):
    return (-5.3e-2 + 2.92e-2*eps
            - 5.5e-4*eps**2 + 4.3e-6*eps**3)
```

THE CELL CYCLE

```
def scan_cell(node):
    level()                         # IMU to < 3 deg
    if vision.blocked(): return SKIP
    auger.clear(); tower.descend(DEPTH_MM)
    dose(50)                        # ml, controlled
    curve = []
    for t in (0, 120, 300, 600):  # seconds
        wait_until(t)
        curve.append(theta(permittivity(
                    probe_a.read_u16())))
    tower.retract()
    tau = fit_decay(curve)          # time constant
    node.store(curve, tau, gps.fix(), soil_temp())
    ble.push(node)                  # live to dashboard
    return OK
```

The four decisions the robot makes without us

Can I scan here?

HuskyLens classifies the cell. Stones become obstacles, red markers become forbidden zones, and both are written to the map rather than silently skipped.

How hard do I push?

Servo current is monitored during descent. Soft ground gets the servo alone; when current crosses threshold before target depth, the pneumatic stroke is triggered.

Is this reading trustworthy?

A reading is rejected if the IMU shows tilt beyond 3°, if the two probes disagree by more than 8% θ, or if soil temperature moved more than 1.5 °C during the curve.

Where do I go next?

The grid is traversed in boustrophedon order, but a cell whose decay constant is more than 2σ from the running mean is re-queued for a second pass before the robot leaves the site.



The MicroPython project in the editor — the scan loop and navigation.

Code quality, honestly. About 1,400 lines across nine modules, with the sensor drivers separated from the scan logic so a probe can be swapped without touching the state machine. What we do not have yet is a test suite — we test on the robot, which is slow and means some bugs only appear in a field session. Building a simulator that replays a recorded CSV through the scan logic is the first item on our software backlog.

WHAT AUTONOMY MEANS HERE

The robot is autonomous in the sense that matters on a site: it decides where it may scan, how hard to push, whether a reading can be trusted, and which cells deserve a second look — without anyone standing over it.

THE STRATASCAN ON APPLICATION

The data has to leave the robot to be useful

Through iteration 4 the survey lived on the robot until someone plugged in a USB cable. On a real site that is unworkable: the archaeologist wants to see the map forming while the robot is still on the grid, so they can stop it, move it, or ask for a second pass.

Ștefan built **StrataScan ON**, a web application that pairs with the robot over **Bluetooth Low Energy** directly from the browser — no driver, no installation, no cable. It streams live telemetry, records the session, and exports the result as CSV for whatever the institution already uses.

LIVE AT

stratascanon .inventikaplace.ro

Judges are welcome to open it during the judging session — on a phone, from the booth. The public pages describe the system; pairing a robot requires an account.

The data path, end to end

Probe analogue › **RP2350** , , GPS › **BLE GATT** 20 B/notify › **Browser** Web Bluetooth › **Dashboard** live heatmap › **CSV** export

Wireless pairing

Web Bluetooth GATT. The robot advertises a custom service; the browser subscribes to notifications. Nothing to install on a museum laptop.

Live telemetry

Moisture, soil and air temperature, humidity, pressure and attitude, plotted as they arrive. A stalled sensor is visible in seconds rather than at the end of the day.

Geolocated readings

Each cell arrives with the GPS fix taken at the moment of insertion, so the map on screen is anchored to real coordinates and can be re-surveyed years later.

Stored history

Every session is recorded server-side, so a survey can be replayed, compared against an earlier visit to the same coordinates, and exported.

Account-based access

Site data is sensitive — an open map of where undisturbed archaeology lies is a looter's shopping list. Access is authenticated per account.

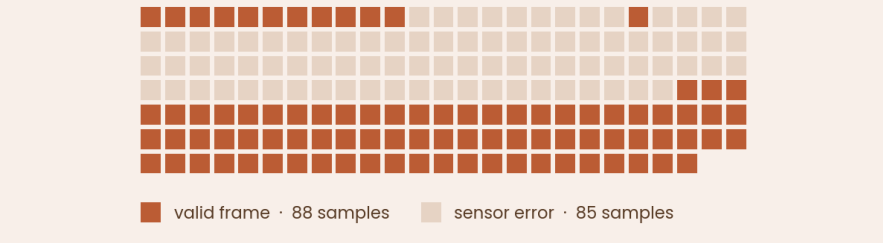
DECLARATION — HOW THIS APPLICATION WAS BUILT

The application was generated with **Base44**, an AI application builder. Ștefan wrote the specification, the Bluetooth service definition and the data model, and directed and corrected the generated code; the AI produced the React front-end and the persistence layer from that brief. We are stating this plainly because the 2026 rules require it and because it would be dishonest to present the interface as hand-written. The firmware on the robot is not AI-generated — it is ours, line by line.

CHALLENGES DURING DEVELOPMENT

Four failures we can prove

A 173-sample session recorded on 7 August 2026, one reading per minute for 1 h 39 min. The BME280 returned a valid frame in only **88 of 173 samples** — a 49.1% failure rate. Each square below is one sample; the pale ones are lost readings.



I²C dropouts on the air sensor

The failure above. Cause: the BME280 was on 20 cm of unshielded jumper wire routed alongside the motor leads, with no ground plane. The failure mode was not random — it correlated with motor current. Fix: the sensor moved onto the XRP's Qwiic bus with a proper twisted cable and a ground return, and onto the printed arm that keeps it away from the motors. This single problem is the strongest argument for leaving the iteration-4 wiring behind.

The GPS that never got a fix

Days of debugging the UART, the baud rate and the parser. The module was fine. The antenna ships separate from the board and we had never attached it to the u.FL connector. The lesson we actually took from it: check the physical layer first, and read the whole datasheet including the parts that look like packaging notes.

A servo that buzzed and reset the board

The MG996 stalled, vibrated and periodically dropped the controller's USB connection. We assumed a faulty servo and ordered a second one, which behaved identically. The cause was that its separate battery pack shared no ground with the controller, so the PWM signal had no reference. Two wires and it worked.

A printed part that was three parts

Servo mount v3 sliced with a cube floating in mid-air. Our screw slots ran the full width of the wall and had quietly cut it into three disconnected strips. We now voxel-test every mesh for single connectivity before it goes to the printer.

DECLARATION OF AI USE · REQUIRED BY RULE 6.5

System	Used for	Extent
HuskyLens (on-robot)	Terrain classification — stones, cracking, no-go markers	Runs on the robot; models trained by us on our own images
Cursor	Code editor with AI completion, used while writing the MicroPython firmware	Assisted; all logic designed, reviewed and tested by Ștefan
Base44	Generated the StrataScan ON web application front-end and data layer	Substantial — see page 13. Specification and corrections by the team
Claude	Structuring this report, checking our physics against the literature, English wording	Editorial and fact-checking; all claims, data and decisions are ours

IMPACT ON SOCIETY

Who this helps — and who it could hurt

Locally. Romania's National Archaeological Record lists on the order of **21,000 sites**, against an estimated hundred thousand that exist. Most are in the care of county museums with one or two archaeologists, no geophysics equipment and a budget that does not stretch to contracting a survey. For them the choice today is to dig on judgement or not to dig at all.

A €1,850 instrument that a single person can carry to a field and operate in an afternoon does not compete with a €30,000 GPR. It competes with *nothing*, which is what most of these teams currently have. That is the honest measure of our impact: not better data than the best labs, but usable data where there is none.



With the excavation team at Jidava, August 2026.

Globally. The same shape of problem exists wherever heritage is dense and funding is thin — much of the Mediterranean, the Balkans, North Africa, Latin America and South-East Asia. Rescue archaeology ahead of roads and pipelines is a particular case: the team has weeks, not seasons, and every trench that comes up empty is time taken from one that would not have. Reducing unnecessary excavation is also, straightforwardly, conservation — an undug layer is a layer that is still there for a better-equipped generation.

Possible negative effects, which we take seriously

Looting

A map of undisturbed archaeology is exactly what an illegal detectorist wants. This is the risk we consider most serious. Mitigation: survey data is account-bound and never public by default; we would not sell to private individuals; coordinates are stored at reduced precision unless the account is institutional.

False confidence

A moisture map is a probability, not a fact. A team that treats a cold cell as "nothing here" may skip something real. Every export carries the calibration conditions and an explicit statement that the method does not detect voids, metal or organic remains.

Deskilling

Instruments can displace the judgement they were meant to support. We designed the output as an additional layer on the archaeologist's own plan, not as a verdict — and the dashboard shows the raw curves, not only the interpreted heatmap.

Environmental cost

The robot adds water to the ground and PLA to the world. The dose is 50 ml per cell — under a litre for a full survey — and the printed parts are recyclable PLA. Small, but not zero, and worth stating.

THREE SITES, THREE TIME CONSTANTS

Our calibration set — and what it does not prove

To turn a time constant into an age you need ground whose age you already know. We have measured three such places, chosen so that they span nearly three orders of magnitude in time and so that one of them contains, with certainty, nothing at all.

Bucharest

≈ 3 years

THE EMPTY CONTROL

The ground beside our own club. It was fully excavated three years ago when the building that houses the club was put up, so we know two things with certainty: there is nothing archaeological under it, and the soil has been disturbed within living memory.

Câmpulung

≈ 100 years

MINE ADMINISTRATIVE BUILDINGS

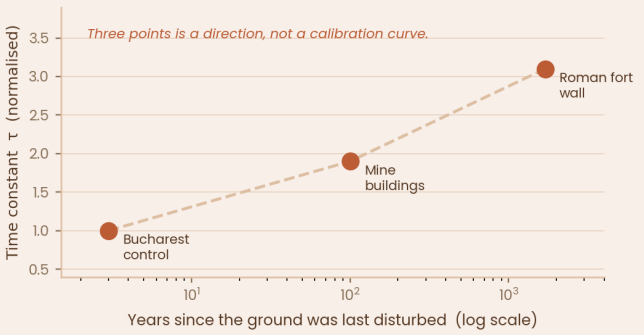
Zone 1 at Jidava, at about 15 cm. A real buried structure, recent enough that the compaction above it is measured in decades rather than centuries.

Jidava

AD 300–358

ROMAN FORT DEFENSIVE WALL

Zone 2, at about 30 cm. Stone, documented, and undisturbed for roughly seventeen centuries — the far end of the range we can currently reach.



Three known points. The direction is right; the curve is not yet a curve.

What we observed. The readings separate in the order the hypothesis predicts: the Bucharest control sits highest, the century-old mine foundations in the middle, and the Roman wall lowest — drier ground, drying more slowly, above the oldest structure.

What that is not. The three sites have *different soils*. Topp's equation is calibrated for mineral soils, but the texture at Câmpulung is not the texture in Bucharest, and texture affects retention independently of compaction. So the ordering across sites is suggestive, not proof. The evidence we actually trust is the **contrast within a single site** — cells above the wall against cells two metres away, in the same soil, on the same afternoon, under the same sky.

WHY THIS MATTERS MOST FOR RURAL SITES

A Roman fort is the easy case. Its plan is standard, its walls are straight, and an archaeologist can predict where the next one runs. A **rural settlement** has none of that: no standard layout, no visible edge, occupation layers spread thinly over hectares. It is precisely where a coarse map of "this ground has been undisturbed for a long time" changes the decision — and precisely where the archaeologists at Jidava told us, unprompted, that they would want it. Our next measurements should therefore not be on another fort.

THE STATE OF OUR EVIDENCE

One known-empty control, one century-old structure and one seventeen-century-old wall. Three points, honestly reported, with the caveat that they sit in three different soils. Building this into a usable calibration is the single most valuable thing we can do next.

ONE TRIED, PRACTICAL USE CASE

The rural settlement, where nothing is mapped

When we asked the archaeologists at Jidava where a tool like ours would actually change their work, they did not say fortifications. They said **rural settlements** — the villages and farmsteads that surrounded every Roman fort and that are, by an enormous margin, the least mapped category of site in the country.

The reason is structural. A fort has walls, and walls show up in aerial photography, in cropmarks, in nineteenth-century plans. A rural settlement is post-holes, hearths and thin occupation layers spread over hectares with no visible edge and no standard plan. Teams find them by opening test trenches on a hunch and moving twenty metres when the hunch is wrong.



Pegging out and opening a grid cell — the unit of work in a survey.

How a survey would run, concretely

DAY 0	The team pegs a 10 × 10 m grid over the suspected area — the same string grid they already use, at 30 cm cell pitch. The robot is calibrated on two control cells of known undisturbed soil from the same field, so the calibration never crosses a soil boundary.
DAY 1, MORNING	Fast pass: a single reading per cell, roughly 40 seconds each, at 3 × 3 sub-sampling. A 100 m² block takes a working day.
DAY 1, AFTERNOON	The dashboard shows clusters of slow-drying cells. The archaeologist marks the twenty most anomalous for the slow pass.
DAY 2	Drying-curve pass on the marked cells, eleven minutes each, fitting τ. Cells that are slow because they are compacted separate from cells that are simply dry.
OUTCOME	The team opens two trenches instead of eight, on evidence rather than on a hunch. Six trenches not opened is roughly a fortnight of labour and six pieces of stratigraphy still in the ground.

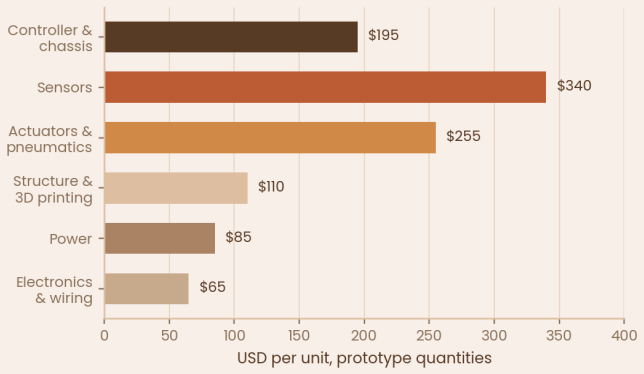
WHAT WE ACTUALLY PROVED, AND WHAT WE DID NOT

Proved. The robot operates on a real excavation, in dust and uneven ground, for two full days. It produced consistent readings at two depths, and a third dataset on known-empty ground in Bucharest. The archaeologists judged the output legible and the concept useful.

Not proved. We have not yet demonstrated that a slow-drying cell reliably corresponds to a buried structure on a site where the answer is unknown to us. Jidava told us the instrument works. It did not tell us the instrument is right — that needs a blind test, on a rural site.

COST STRUCTURE

What one unit costs, and what a company costs



The robot is cheap. The company is not. A complete unit costs us about **$1,050** to build one at a time — the figure on the left, from the parts we actually bought. At fifty units, with the printed structure injection-moulded and the electronics consolidated onto one board instead of six modules, we estimate **€620**.

That is the easy part. Everything that turns a working prototype into something a public institution is allowed to buy costs an order of magnitude more, and it is where a student project usually stops. We would rather show the number than pretend it is not there.

Bill of materials for one complete robot at prototype quantities.

Cost area	One-off	Recurring / year	Note
Bill of materials — prototype qty	$1,050 per unit	—	Verified against our own purchase records
Bill of materials — at 50 units	≈ €620 per unit	—	Estimate: moulded structure, consolidated PCB
R&D — hardware and soil-science validation	€300 spent to date	€1,200	Sensors, spares, filament, field consumables
Software development	€0 cash (our time)	€400	Hosting, domain, moving off the AI builder platform
Field testing with archaeologists	€500–€1,000	€1,500	Travel, permits, consumables per campaign
Marketing and go-to-market	€1,500	€2,500	Conference stand, printed materials, EN documentation
CE marking and EU safety compliance	€7,000–€15,000	—	Mandatory before any sale to an EU institution
Patent — Romania only	€5,000–€7,000	€500	Renewal fees after grant
Patent — international (PCT, key markets)	€20,000–€50,000	€2,000	Deferred; only worth it once sales exist

READING THE NUMBERS

Building the robot is roughly 3% of the cost of being allowed to sell it. Our conclusion is that the sensible first commercial step is **not** selling hardware at all — it is selling *surveys* as a service, which needs no CE mark, generates the validation dataset page 16 says we still lack, and funds certification from revenue rather than from investment.

REVENUE STREAMS AND MARKET

Three ways this makes money, in the order we would actually try them

1 Survey as a service
€400 per field day

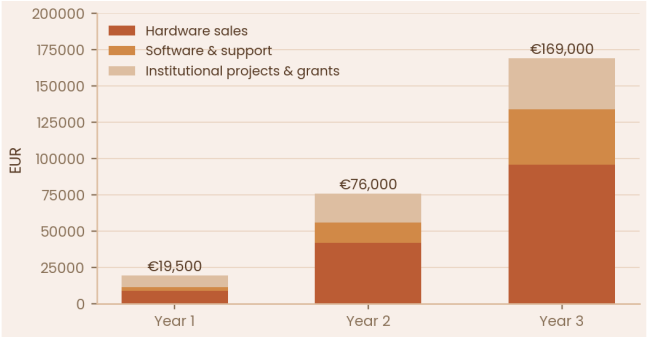
We bring the robot and run the survey. No certification needed, no capital cost to the customer, and every job adds to the validation dataset. This is how we would start, and it is the only stream a student team could realistically operate.

2 Hardware sales
€1,850 per unit · 55–66% gross

Once CE-marked. Sold to county museums, university archaeology departments and heritage consultancies. Priced deliberately below a departmental procurement threshold so it can be bought without a tender.

3 Software subscription
€25 per month per institution

The dashboard, storage and re-survey comparison. Small, but it is the only recurring revenue and it is what keeps us in contact with users after the sale.



A conservative three-year projection, service-led.

The buyer we are actually aiming at. Not the ministry and not the university — the **county museum**. There are roughly forty in Romania with an archaeology section, most with a capital budget that will not absorb a €30,000 instrument but will absorb €1,850.

Our route in is the institution we already work with: the "Vasile Pârvan" Institute of Archaeology, whose specialists have seen the robot work on their own site.

Market, sized honestly

Segment	Approximate size	Reachable in 3 years	Why they would buy
Romanian county museums with archaeology	≈ 40	8–12	No geophysics at all today; price fits a capital line
Romanian university archaeology departments	≈ 12	3–5	Teaching instrument as much as a survey tool
Rescue-archaeology contractors (RO)	≈ 25 firms	4–6	Time pressure is their whole business model
EU heritage institutions (same profile)	hundreds	pilot only	Same problem, larger budgets, harder to reach

RESOURCES, PARTNERS AND NEXT STEPS

What the project rests on, and what follows

Physical

The robot: XRP controller, sensor suite, insertion tower and our own printed structure. Deliberately built from parts anyone can buy.

Intellectual

The measurement protocol — dose, sample, fit τ — and the calibration routine. This, not the hardware, is what would be worth protecting.

Human

Three students across mechanics, research and software, and a coach. Our real constraint: we are also in school, and the project moves at the speed of weekends.

Relational

Access to a working excavation and to archaeologists who will tell us when we are wrong. Almost impossible to buy, and the thing we would protect first.

Partners

Partner	Type	What they contribute
"Vasile Pârvan" Institute of Archaeology	Scientific	Domain validation, site access, and the specialists who defined the problem — Dr. Florian Matei-Popescu and arch. Ion Dumitrescu
Inventika Place / Association ARTESA	Institutional	Workshop, equipment, coaching and the legal entity through which the project is funded
Waldevar Energy	Technical & financial	Sponsor; discussions under way about mounting our sensor payload on a walking robot platform for rough terrain
Schneider Electric, Exim Banca Românească, WeltPixel, Sara Octomed, Conterra, Grand View, RoDesign	Financial	Fund the build, the field campaigns and the team's participation in international competition

WALDEVAR
WE BUILD THE ENERGY OF THE FUTURE

Waldevar Energy — sponsor and technical partner for the walking-platform work.

Next steps

BEFORE ONTARIO

Close out the build

Finish the printed structural frame that is still in revision and complete the motor-driven insertion tower, so the robot the judges see is the robot this report describes.

AUTUMN 2026

The blind test

Survey a rural site whose subsurface is documented but unknown to us, hand our map to the archaeologists, and let them score it. Until that is done we claim plausibility, not accuracy.

AFTER THAT

Widen the calibration set

Every additional point of known age turns the three points on page 16 closer to a usable curve. Paid surveys are the cheapest way to collect them.

WHERE WE ARE, IN ONE PARAGRAPH

StrataScan ON is a working prototype that has been operated on a real excavation, judged useful by professional archaeologists, and rebuilt five times to get there. It is not a product and we are not going to describe it as one. What it is, is an instrument that measures something real, cheaply, where nothing is measured today — and a clear, costed, unglamorous plan for finding out whether it is right.

Reading the past, layer by layer.

LIST OF SOURCES

Sources and people we spoke with

People

Name	Role	What we discussed
Arch. Dr. Florian Matei-Popescu	"Vasile Pârvan" Institute of Archaeology, Roman period specialist	February 2026 consultation that defined the problem; field work at Jidava, 12–13 August 2026
Arch. Ion Dumitrescu	"Vasile Pârvan" Institute of Archaeology	Field work at Jidava; the assessment that the method suits rural settlements
Adriana Sfetcu	Horticulturist	Soil compaction, water retention and evaporation rate — confirmation that the physics is sound

Literature and technical documentation

- Topp, G. C., Davis, J. L. and Annan, A. P. (1980). "Electromagnetic determination of soil water content: measurements in coaxial transmission lines." *Water Resources Research*, 16(3), 574–582.
- Clark, A. (1996). *Seeing Beneath the Soil: Prospecting Methods in Archaeology*. Routledge.
- Gaffney, C. and Gater, J. (2003). *Revealing the Buried Past: Geophysics for Archaeologists.* Tempus.
- Hillel, D. (2004). *Introduction to Environmental Soil Physics.* Elsevier.
- World Robot Olympiad Association (2026). *WRO 2026 Future Innovators — General Rules*, version 15 January 2026. wro-association.org
- Asociația de Robotică Educațională și Știință WRO — World Robot Olympiad Romania, national organizer. *Season 2026 regulations, national-stage results and category guidance.* wroromania.ro
- SparkFun Electronics. *SparkFun XRP Controller Hookup Guide.* docs.sparkfun.com/SparkFun_XRP_Controller
- Open-STEM. *XRP mechanical CAD repository.* github.com/Open-STEM/XRP_mechanical
- DFRobot. *SEN0308 Capacitive Soil Moisture Sensor — datasheet and calibration notes.*
- Bosch Sensortec. *BME280 combined humidity and pressure sensor — datasheet, rev. 1.23.*
- DFRobot. *HuskyLens AI Machine Vision Sensor — user manual.*
- National Institute of Heritage, Romania. *Repertoriul Arheologic Național (RAN).* ran.cimec.ro
- Institute of Archaeology "Vasile Pârvan". *Jidava (Câmpulung Muscel) — Roman auxiliary fort, site documentation.*
- Team Sensor ON (2018–2026). *Eight seasons of robotics and innovation projects, and the awards that came with almost every one of them.* Unpublished, but we were there.

THE LIVE DASHBOARD

stratascanon.inventikaplace.ro

A note on our own data. Every measurement quoted in this report comes from our own field sessions and is available as raw CSV on request: the 173-sample session of 7 August 2026 quoted on page 14, the Bucharest control readings, and the Jidava survey of 12–13 August 2026. We have not smoothed, trimmed or reprocessed any of it for presentation — the 49% sensor failure rate on page 14 is in the file exactly as it came off the robot.