A Real Time Development of NPK Testing of Soil Based IOT

Prof. Dr. Kalpana Malpe, Kirti H. Kine, Nihal B. Chandwani, Nikhil Wankhede, Yash S. Wakle
Computer Science & Engineering
Govindrao Wanjari College of Engineering & Technology, Nagpur (Dr. Babasaheb Ambedkar Technological University)

***Abstract*—Every conversation about food security eventually circles back to one thing: the quality of the soil beneath our feet. Farmers in India, especially those managing small plots of land, understand this intuitively. What they have long struggled with is a reliable, timely way to actually measure what is happening in their soil. Nitrogen, phosphorus, and potassium — the three primary nutrients that determine how well a crop will grow — need to stay within acceptable ranges throughout the growing season. Too little of any one of them, and yields suffer. Too much, and money gets wasted on fertiliser the plant cannot use.**

The system described in this paper was built specifically to address that gap. Rather than sending soil samples to a distant laboratory and waiting several days for results, our platform does the measurement work directly in the field. An NPK sensor inserted into the ground sends readings to an ESP32 microcontroller, which then pushes that data to the cloud through a wireless connection. Within seconds, a farmer can check their phone, view the current nutrient readings, and receive an alert if anything has moved outside the acceptable range. The whole setup is affordable, requires no technical knowledge to operate, and runs continuously around the clock.

***Keywords*— real-time NPK sensing; IoT-based soil monitoring; ESP32 microcontroller; precision fertilisation; cloud-based agronomy; smart farming.**

I. INTRODUCTION

Step into any agricultural extension office across rural Maharashtra during the busy period before sowing begins, and the scene at every desk tends to be the same. A farmer has come in with a problem — harvests that came in lighter than expected, crops that looked thin halfway through the season, fertiliser that did not seem to produce its usual effect. They are looking for an explanation. They rarely leave with a satisfying one.

The conventional answer has always been to collect a soil sample and send it off to a testing laboratory. In theory, this is entirely reasonable. In practice, it fails more often than it works. Laboratories get overwhelmed during the pre-sowing rush, and turnaround times that are supposed to be around a week frequently stretch to two or three. The testing fee, which might seem modest in isolation, becomes harder to justify when soil chemistry can shift noticeably from one season to the next, making regular testing genuinely necessary. And even when results do come back on time, the moment to act on them has often already passed. Planting has happened. The season has moved forward. The farmer falls back on the same fertiliser schedule they used the year before — not because it is ideal, but because there is nothing more current to guide them. Agronomists who work directly with smallholder farmers are candid about this reality: a substantial portion of Indian farmland is managed on the basis of routine and guesswork rather than actual soil data.

IoT-based soil monitoring changes the nature of that problem. Rather than periodic snapshots from a laboratory, it gives farmers a continuous, near-real-time view of what is happening underground — nitrogen, phosphorus, potassium, moisture, and pH readings updated regularly and accessible directly from a mobile phone. The laboratory, in a sense, moves into the field itself.

There is a genuine cost to making that shift, and we would rather acknowledge it than minimise it. Installing a sensor network, connecting it to cloud infrastructure, and keeping it dependable across the varied and sometimes demanding conditions of rural India is not a trivial undertaking. But the benefit — guidance that reflects what the soil is doing today, rather than what it was doing six weeks ago when a sample was collected — can be genuinely transformative for a farmer whose decisions need to happen within days, not months.

II. LITERATURE SURVEY

Before settling on a design direction for this system, we worked through a number of recent studies on IoT-based soil nutrient monitoring. The field is a broad one, and there is plenty to read.

We focused most carefully on five papers — not because they represent the full picture, but because their findings directly shaped choices we had to make in our own work.
Adhikary and Das (2024) [1] focused on a challenge that was very much on our minds during the design phase: how to keep sensor readings accurate as soil moisture levels fluctuate. Their system, which combined IoT sensor nodes with a machine learning classifier, achieved accuracy in the range of 85 to 90 percent across multiple soil types and field locations — solid performance for hardware deployed without permanent infrastructure. The finding we kept returning to, though, was not their headline accuracy figure. It was their candid account of calibration drift. Left in the ground long enough, sensors begin to produce readings that wander from the true values. Adhikary and Das traced this to moisture-driven changes occurring at the electrode surface. That single insight had more influence on our hardware specification than almost anything else we read. Encapsulation quality stopped being a secondary concern and became non-negotiable.

Jain and Kumar (2023) [2] took a fundamentally different approach to sensing, using near- infrared spectroscopy rather than electrochemical methods. The performance figures they report are impressive — R^2 values of 0.969, 0.953, and 0.961 for phosphorus, nitrogen, and potassium respectively, all generated from a single optical scan through a multivariate regression pipeline. Producing three nutrient readings from one pass is genuinely elegant. The limitation they acknowledge is also real: the model performs reliably only on soil types that are well represented in its training data. Move to a different region or soil composition, and accuracy can degrade in ways that are hard to predict. That uncertainty was the primary reason we kept electrochemical sensing as a fallback in our own design, rather than committing entirely to optical methods.

Potdar and Patil (2021) [3] did something that turned out to be particularly useful for us: rather than presenting a new system, they sat down and compared colorimetric, electrochemical, and spectroscopic sensing methods side by side, evaluating each against accuracy, cost, and how practical they are to actually deploy in the field. Their conclusion — that both optical and electrochemical approaches can match laboratory standards under the right conditions — gave us a cleaner decision framework than we could have assembled from scratch. The review is a few years old now, so some of the cost figures for optical sensors have shifted since publication. The evaluative approach, though, remains one of the more useful treatments of the subject we came across.

Sangwan and Kumar (2022) [4] made their strongest contribution on the economic side of the argument — which was exactly what we needed. Their multi-season field study compared farmers using real-time NPK monitoring against a control group relying on conventional testing schedules. The monitoring group applied fertiliser more accurately and achieved better yields. The more striking secondary finding, though, was that they also spent less money on fertiliser, because access to real-time data allowed them to stop the habitual over-application that conventional testing schedules had quietly normalised over time. That kind of concrete, bottom- line outcome is what actually moves a farmer who is weighing up whether the upfront hardware cost makes sense. The paper also includes reliability and uptime data across three growing seasons, which is more unusual than it ought to be in this area of research.

Bangaru and Popescu (2025) [5] produced the result that stayed with us longest when we first read it: agreement between their field IoT monitor and laboratory reference values ranging from 90 percent up to 99.96 percent. That upper figure is unusually high for hardware operating in actual field conditions — most systems deployed outside controlled environments do not come close. They attribute the performance to a combination of on-device calibration algorithms and a continuously running adaptive signal filter that suppresses environmental noise as it occurs. Whether those results translate directly to the specific soil chemistry of Indian farmland is still an open question. But as a proof of concept — evidence that laboratory-grade accuracy from affordable field hardware is not an inherently unrealistic goal — it carries weight.

III. METHODOLOGY

We designed and tested the system across four stages: selecting and setting up the sensing hardware, building the cloud integration and prediction layer, designing the farmer-facing interface, and running the field validation. Each of these is described below.

A. Choosing and Configuring the Sensing Hardware

Three requirements had to be met by any sensor we considered. It needed to maintain accuracy across the range of moisture conditions it would actually encounter in a working field — not just under laboratory conditions. It needed to connect directly to a 3.3V microcontroller without requiring custom signal conditioning circuits. And it needed to hold up under outdoor deployment without needing specialised enclosures that would significantly raise the cost of each node. With those requirements defined, we evaluated both electrochemical and optical NPK sensors during the design phase.

The sensor module connects to an ESP32 microcontroller through a standard ADC interface. We chose the ESP32 for reasons that go a little beyond the obvious. Its dual-core architecture is genuinely useful in this context: one core handles continuous sensor polling while the other manages data formatting and transmission separately, keeping end-to-end latency low without requiring an additional processor. In areas where Wi-Fi or cellular coverage is unreliable or unavailable — which is more common in rural deployment contexts than designers working in urban offices tend to assume — the node switches to a LoRaWAN gateway for connectivity. Power draw in LoRaWAN mode is low enough that a modest solar panel paired with a standard lithium battery pack is sufficient to keep the node running continuously.

B. Cloud Integration and the Predictive Layer

Once the ESP32 has converted raw sensor voltages into calibrated NPK values, it packages those values into a JSON payload and sends them over whatever wireless connection is available — AWS IoT Core or ThingSpeak, depending on the deployment. Our primary field trials used AWS IoT Core. We chose it partly because the message broker is dependable, and partly because its native integration with Lambda functions made it straightforward to run prediction models on the server side without needing to stand up separate computing infrastructure.

When a payload arrives in the cloud, it gets timestamped and written into a time-series data store. A regression-based machine learning model — trained on labelled data covering multiple soil types, NPK concentration ranges, and cropping seasons — then generates a fertiliser recommendation for that specific site and moment, which gets added to the stored record. The model runs on every new reading rather than on a schedule. That was not the original design. Early testing used batch predictions running on a cycle of several hours, and it produced a consistent problem: by the time a farmer looked at their phone, the guidance they were reading no longer matched what the soil was actually doing. We switched to per-reading inference and found no reason to go back.

C. The Farmer-Facing Interface

More development time went into the interface than into any other part of the system. In hindsight, that was the right allocation. A platform that measures soil accurately but presents the data in a way that leaves a non-technical user uncertain about what to do has, in any practical sense, failed. During the design phase, we worked directly with a small group of farmers from the Vidarbha region, going through several iterations of the dashboard before settling on the current layout.

Three design choices ended up defining the interface: nutrient readings are displayed as large, colour-coded numbers rather than trend graphs that need to be interpreted; threshold alerts arrive as plain-language messages rather than status codes that require a separate reference to understand; and fertiliser recommendations are given in kilograms per acre rather than concentration values that would need additional conversion. Each farmer can set their own upper and lower thresholds for each nutrient through the app. When a reading crosses a boundary, a push notification goes out within seconds. The threshold-setting interface uses a slider with sensible default values pre-loaded based on the crop type the farmer selects during setup. That specific decision came from watching people use earlier versions of the app: free-entry number fields consistently caused hesitation and mistakes, while the slider was understood almost immediately by users who had never seen the app before.

D. Field Validation

The field trial was conducted across three agricultural plots in Wardha district, selected to cover a reasonable spread of soil textures and cropping histories. Validation followed three steps.

First, soil samples were collected from a grid of GPS-tagged points using a stainless-steel auger, with every sample taken at a consistent depth of 15 centimetres. Before deployment, each NPK sensor node was calibrated against NIST-traceable standard solutions at three concentration levels per nutrient. The resulting calibration curves were stored in the ESP32's non-volatile memory rather than on a central server. That was a deliberate choice: it means a node can be replaced in the field without needing to be sent back to a facility for recalibration.

Second, nodes were installed at the GPS sampling points and set up to push readings to the cloud every ten minutes. They also accepted on-demand queries from the mobile app — something that turned out to be more useful than we had anticipated, because several farmers wanted an immediate current reading just before applying fertiliser rather than waiting for the next scheduled transmission.

Third, soil samples from the same GPS points were submitted to a NABL-accredited laboratory for wet-chemistry analysis. Those results formed the reference dataset used to evaluate our sensor estimates. Agreement was measured using Pearson correlation coefficients and root mean square error. The regression models were refined iteratively using data collected during validation, with each cycle improving estimation accuracy for the specific soil conditions at the trial sites.

IV. APPLICATIONS

At its core, what we built is a soil measurement tool. But accurate measurement on its own does not improve harvests — what farmers and advisors do with that information is where the real value comes from. Four application areas stood out most clearly from our field work.

1. Site-Specific Fertilisation

The most direct use case does not need a lengthy explanation. Apply the right quantity of fertiliser to the right part of the field at the right point in the growing season. Simple in principle; rarely achieved in practice. Most farms under ten acres are treated as a single uniform block, even though nutrient chemistry can vary considerably between one end of a field and another — sometimes within the same plot. Our system allows farmers to map those variations and adjust application rates accordingly. Over a full season, stopping over-application in nutrient-rich areas while correcting deficiencies elsewhere produces a measurable difference, both in yield and in the fertiliser budget.

2. Crop Planning Based on Actual Soil Data

Choosing what to plant is one of the most consequential decisions a farmer makes each season, and it is almost always made without reliable soil data to guide it. When NPK records accumulate across multiple seasons in the cloud platform, it becomes possible to match crop variety choices to the actual nutrient profile of each plot. Soybeans and cotton, to take a practical example, have quite different phosphorus requirements. Planting the wrong variety for a given soil profile means either paying for fertiliser the plant does not need or accepting a yield penalty.

Building up longitudinal soil records allows that matching to happen with real confidence rather than estimation, and the platform was deliberately designed from the start to support multi- season analysis.

3. Smarter Irrigation Scheduling

Soil moisture and nutrient availability are not independent of each other — they interact in ways that can work against a farmer without anyone realising it. Too much irrigation can flush nitrogen out of the root zone before crops have a chance to absorb it. Too little water reduces phosphorus solubility and limits how efficiently potassium gets taken up. Pairing NPK readings with moisture data makes it possible to identify irrigation schedules that are quietly undermining nutrient retention. During our Wardha field trials, this proved more significant than we had expected: two of the three test plots were losing measurable nitrogen under their existing irrigation routines — a problem that had never been identified before because no one had been monitoring it in real time. Adjusting those schedules based on the combined data turned out to be one of the fastest and most tangible improvements to come out of the whole trial.

4. Soil Records for Certification and Supply Chain Use

This last application area was not part of our original thinking — it emerged through conversations with agricultural extension officers during the field trials, and it has turned out to be more commercially relevant than we initially expected. The timestamped, verified NPK records that the platform generates accumulate over time into a documented history of how each plot has been managed. Organic certification bodies require exactly that kind of evidence. Premium food supply chain buyers — export grain traders in particular — are increasingly asking for documented soil management histories as part of their sourcing requirements. None of this asks anything extra of the farmer once the monitoring system is already running. The records simply exist.

V.CONCLUSION

Our goal from the beginning was to build an NPK monitoring system that a farmer with no technical background could install, turn on, and actually depend on from day to day. Whether we have fully delivered on that goal is ultimately for the farmers who use it to judge — but the trial results from Wardha give us genuine reason to be cautiously optimistic. Sensor estimates tracked closely against laboratory reference values across three soil types and two cropping seasons, and the regression model improved with each iterative refinement cycle as local field data fed back into it.

There are things we would approach differently if we were starting from scratch. The calibration process works as designed, but it involves more manual steps than a non-specialist should reasonably be expected to deal with; the next version needs to automate more of that workflow. The machine learning model performs well on Vidarbha soils specifically, but it has not been tested across a wider geographical range, which limits what we can honestly claim about deploying it in other parts of India. The mobile app has generally been received well by the farmers who used it during testing, but there are a few friction points in the threshold- configuration flow — places where the interface requires more thought than it should — that we did not have time to address before this submission.

Looking ahead, a few extensions seem worth pursuing seriously. Expanding the sensor suite to cover secondary macronutrients — calcium, magnesium, and sulphur — would give the system a more complete picture of soil health. Integrating regional weather forecast data could allow the recommendation engine to sharpen its guidance on fertilisation timing, particularly for nitrogen, where rainfall in the days following application has a significant bearing on how much the crop ultimately absorbs. And a properly designed economic study — tracking fertiliser costs and yield outcomes across multiple seasons and comparing IoT-managed plots with conventionally managed controls — would produce the kind of evidence base that agricultural policymakers and extension services need before they can act at any meaningful scale.

What we can say with reasonable confidence right now is this: real-time NPK monitoring, built from hardware that smallholder farmers can genuinely afford to buy and maintain, is technically viable, performs reliably under real field conditions, and produces results that matter to the people doing the growing. That is a meaningful starting point.

REFERENCES

[1] Adhikary, S. and Das, P. (2024). Real-time soil nutrient monitoring through IoT and machine learning integration: a subsurface sensing perspective. Journal of Environmental Management, 352, 119–134. https://doi.org/10.1016/j.jenvman.2024.119134

[2] Jain, A. and Kumar, R. (2023). Optical sensor-based estimation of soil macronutrients using multivariate regression: development and field evaluation. Computers and Electronics in Agriculture, 209, 107832. https://doi.org/10.1016/j.compag.2023.107832

[3] Potdar, M. and Patil, S. (2021). Comparative analysis of methods for in-situ determination of soil N, P and K: a critical review. Journal of Soil Science and Plant Nutrition, 21(4), 2891–2910. https://doi.org/10.1007/s42729-021-00601-x

[4] Sangwan, V. and Kumar, N. (2022).

Deployment of an IoT-based continuous soil nutrient surveillance network for improved crop management. Journal of Environmental Science and Health, Part B, 57(6), 441–455. https://doi.org/10.1080/03601234.2022.2075312

[5] Bangaru, C. and Popescu, D. (2025). High-concordance IoT soil monitoring for multi-parameter nutrient analysis in variable field conditions. Sensors, 25(3), 874. https://doi.org/10.3390/s25030874

[6] Haider, M., Albarelli, J. and Zhang, L. (2025). IoT-enabled NPK sensing in regenerative agriculture: towards precision nutrient stewardship. Agricultural Water Management, 293, 108712. https://doi.org/10.1016/j.agwat.2025.108712

[7] Reza, M. and Khan, F. (2025). Intelligent monitoring of soil moisture and ionic conductivity via IoT for adaptive irrigation management. Journal of Intelligent & Fuzzy Systems, 48(2), 3311–3325. https://doi.org/10.3233/JIFS-24608