CROPIQ: AI-POWERED HYBRID YIELD FORECASTING FOR AGRICULTURE

[1]Narmathadevi S, [2]Saveetha SS, [3]Shruthi P, [4]Jeevanandham S

[1, 2, 3] UG Scholar of Information Technology, Sri Ramakrishna Engineering College, Coimbatore, India
[4] Assistant professor of Information Technology, Sri Ramakrishna Engineering College, Coimbatore, India

Abstract: Guessing how much crops will grow helps plan farming, keep food steady, leaves markets less shaky. Yet farm data holds tricky time-based shifts, repeating cycles year after year, links between factors that bend instead of line up straight - tough stuff when trying older number tricks or basic AI on its own. A fresh look at crop forecasting comes through CropIQ, mixing three smart methods into one system. Instead of relying on just one approach, it uses Prophet to spot slow changes across seasons in farm-related numbers over time. While Prophet handles broad patterns, LSTM jumps in to track how events link together moment by moment. At the same time, XGBoost digs into complex links between weather factors and farming conditions, uncovering hidden relationships. Because each part brings something different, their mix strengthens overall guesses about harvests. With all pieces working side by side, the whole system handles surprises better, making estimates steadier when things get unpredictable. Tests ran on old farming records used RMSE, MAE, MAPE, along with R^2 to check how well things worked. Accuracy plus consistency came out better with the combined method than with standalone models. Because it gives steady forecasts, people who grow crops or shape farm policies might find this useful when organizing planting, handling supplies, or balancing food availability.

Keywords**: Crop Yield Prediction, Hybrid Deep Learning, Time Series Forecasting, Prophet, LSTM, XGBoost, Smart Agriculture**.

I. INTRODUCTION

Growing crops matters for feeding people and keeping economies steady. Because weather shifts, soil varies, seasons change, and farms work differently, guessing harvests gets tricky. Predictions guide planting choices, shape how tools and land are used, also affect rules made by leaders. When numbers reflect reality, farmers adjust earlier, prepare better, handle surprises well. Knowing what fields will produce supports decisions across entire food systems. Though methods improve, nature's unpredictability stays a core challenge. New progress in artificial intelligence, along with access to massive farming data, has opened doors for using advanced algorithms in predicting farm outcomes. Instead of guessing, these systems study old records to spot trends affecting how much crops grow. Research shows such models perform well when estimating harvest amounts. With clearer forecasts, growers gain an edge - yet officials who manage land use and supply chains find value too. Better predictions help shape smarter choices around feeding populations, maintaining farmland health, while keeping markets steady.

Even with progress, older number-crunching techniques struggle when farm data gets messy. Because real-world growing conditions shift slowly over years, swing wildly day to day, and tie together in twisted ways. Machines that learn sequences - like LSTM systems - can track what comes after what, yet often miss repeating cycles or tangled cause-effect links on their own. While tools like Prophet sketch out rising curves and yearly loops well, they tend to overlook how inputs twist around each other unpredictably. Facing these issues head on, CropIQ emerges as a mix of artificial intelligence tools including Prophet, XGBoost, and LSTM, built to forecast how much crops will produce. Instead of relying on just one method, it pulls together time-based statistics, pattern tracking through neural networks, alongside smart feature analysis via algorithms. Because each model covers what the others miss, their teamwork leads to steadier, more trustworthy estimates over time. What results is a system tuned to handle uncertainty while delivering sharper forecasts when needed most.

II. LITERATURE REVIEW

Not long ago, studies began looking at how machines could learn to predict farm outputs more precisely. Instead of old ways, some turned to systems inspired by brain cells - ones that dig into patterns hidden in data. Meena with Chaitra [1] built one such system aiming at prices for ragi, a grain grown in Karnataka. It showed these smart models see deeper than usual tools when dealing with messy field numbers. Results got sharper because the method handled twists in weather, soil, and market shifts better. Starting off differently, Wang and colleagues presented a method blending CNNs with LSTMs to forecast harvest amounts. This mix handles location patterns along with time-based changes found in farming records. Results showed better accuracy when using combined systems instead of single approaches. Such designs pick up on details spread across fields as well as shifts happening through growing seasons.

Not far off from expected, Abdel-Salam and team [3] built a system using tuned machine learning models along with mixed methods for picking key data traits - aimed at forecasting how much crops will produce. What stood out was how crucial it became to pull only meaningful variables when trying to boost forecast precision within farming-related information sets. Late results from Kiran Kumar's team [4] looked at how well LSTM plus Bi-LSTM work when guessing crop output, setting those against older machine learning approaches. What stood out was that the deeper networks handled patterns in farm-related sequences better than standard regression ever did.

Farming forecasts have seen testing with tools like Random Forest, along with Support Vector Machines, plus Neural Networks, according to various research efforts [5], [7]. Performance got a boost using these methods - still, combining different approaches often works better since single ones might miss subtle trends across farm data.

Looking at past studies, mixing several prediction methods tends to give better results when estimating harvest amounts. Instead of relying on just one way, combining them uses what each does well. This blend often handles changes in weather or soil more effectively than single models. Accuracy improves because weaknesses in one approach get balanced by others. Most research points toward these combined systems as a stronger choice overall.

III. RESEARCH METHODOLOGY

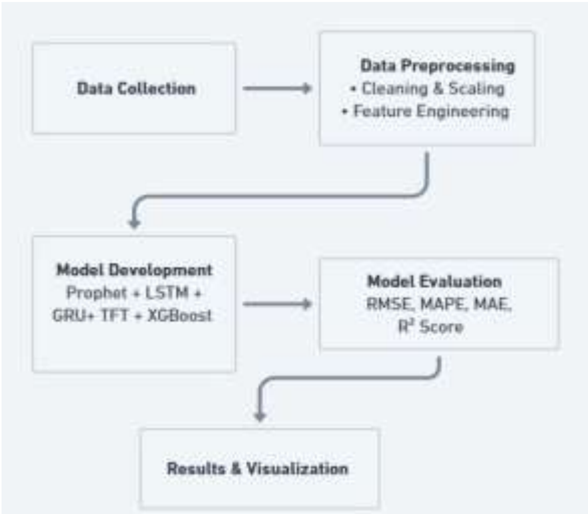

Fig 3.1 Flow Chart

A. Dataset Collection

Crop production records and environmental variables are among the historical agricultural data included in the dataset used in this investigation. These variables include climatic elements that have a big impact on agricultural productivity, like temperature, humidity, and rainfall.

B. Data Preprocessing

Before any modeling begins, raw data gets cleaned so it works better for machine learning systems. Because farm-related information sometimes has gaps or strange entries, adjustments become necessary. Odd patterns in features might mislead predictions, making corrections a priority. Fixing incomplete rows helps later stages run without hiccups. Scattered outliers tend to skew outcomes, which is why smoothing steps come into play. Each change aims to shape the set into something reliable for training networks. Preparation methods differ based on what flaws show up during inspection. The goal stays constant: turn messy inputs into structured ones quietly, behind the scenes.

At first glance, gaps show up when checking the dataset for absent or partial entries. Such omissions often come from faulty sensors, broken logs, or limits during gathering phases. Instead of dropping these points, smart filling methods step in - keeping patterns intact without distorting how numbers naturally spread across the set.

After that, spotting odd values happens before tossing them out - these standouts often twist how data usually spreads. Because mistakes in measuring or wild weather shifts cause some numbers to stray far, they sneak into training and nudge results off track. Tossing those aside keeps predictions steadier when new situations show up later.Once the data gets cleaned, numbers are adjusted so they fit within similar ranges. Because things like rain levels, heat readings, and harvest amounts often measure differently, shifting them helps balance their impact when teaching models. Here, values go through Min–Max adjustment to sit neatly from zero up to one. Such rescaling works well with smart systems like LSTM, where big number gaps can cause issues unless smoothed out first.

Last comes shaping the cleaned data into a layout fit for tracking changes over time, keeping each point in order so past values still connect to future ones. By handling the data this way, its reliability improves - models then spot real trends without getting tripped up by mess or gaps.

C. Model Development

Prophet Model: Ahead of most tools, Prophet handles repeating cycles and gradual changes across time. Built by Facebook, it works openly so anyone can study or adjust its code. Instead of treating data as one block, it splits signals into pieces like drifts, yearly rhythms, or special dates. Each segment gets shaped separately, then they blend into a full picture later. Even when holidays shift or gaps appear, the system adapts without breaking rhythm. Smooth shifts and sudden spikes both find space within its structure.

When predicting farm outputs, harvest numbers tend to follow clear yearly rhythms shaped by weather, when crops go into the ground, and when they are gathered. Because of this, Prophet handles the job well - its design fits naturally with repeating swings and slow shifts seen in farming records.A curve rises, then dips - Prophet sees that shift without being told. Instead of fixed rules, it adjusts itself when seasons change or surprises hit. Missing points? Outliers lurking? The method keeps going anyway. Real farms deal with messy records; this fits right in.

Later on, the forecasts help boost accuracy when mixed with outputs from additional models inside the combined system. At first, the Prophet model digs into crop yield records, spotting slow-moving seasonal trends over time.

GRU Model: One step at a time, GRU networks improve on old-style RNNs by handling sequences better. Because regular RNNs struggle to remember patterns over long stretches, gradients tend to fade away - this is what holds them back. Instead of stacking complexity, the GRU adjusts flow through gates that control how much past information stays active. With smarter internal switches, it sidesteps the fading signal issue found in basic designs. Over here, progress comes not from more layers but from sharper memory management built into its structure.

Not like LSTM, GRU runs on just two key switches - the update one plus the reset one. What slips through before comes down to the update switch. The reset one decides what old data gets dropped. Fewer moving parts mean it grabs time-based trends quicker than LSTM. Less math per step cuts processing load without losing learning power.

When predicting harvests, weather shifts like rain, heat, soil wetness, yet changing seasons shape how crops grow over time. Earlier conditions often set the stage for what comes next in field results. Because growth links to prior moments, timing matters a lot. GRUs handle these linked events well, noticing rhythms that unfold day by day or week by week. Their design helps track evolving trends without getting lost in noise. Past data guides future estimates more reliably when using this approach. Patterns emerge clearer when sequences stay central to analysis.

LSTM Model: Imagine trying to remember details from earlier in a story - LSTMs do something similar. These aren’t your average neural nets; they’re built like RNNs but smarter. Hidden inside? Cells that store info across stretches of data points. Because of this trick, spotting patterns over time gets easier. Most regular systems forget too fast. But these keep hold of what matters, even after many steps pass by.

Years go by, yet what grows in fields still leans on weather from seasons long past. Hidden layers inside LSTMs hold pieces of old data like quiet reminders. These echoes help make sense of patterns stretched across time. What happened back then shapes predictions today. Memory lingers, not in words but numbers. Time leaves traces systems can learn. Past rains, past harvests - still matter now.LSTM networks contain several parts like an input gate, a forget gate, one output gate - each managing how info moves inside. Information sticks around when it matters, fades when it does not, thanks to these controls shaping what stays. Accuracy gets better because only useful pieces influence what comes next.

TFT Model: What makes the Temporal Fusion Transformer stand out? It's built for clear, long-range forecasts using time data. Instead of just repeating patterns, it combines memory cells with focused attention layers. One part tracks changes over time, while another weighs which inputs matter most. This mix helps spot links across moments plus relationships between variables. Even with messy sequences, it holds structure without oversimplifying. Built for clarity, not complexity - interpretation stays central through each step.Starting off, the TFT setup uses parts like gated residual nets, combined with variable picking systems. Sequence-based layers work alongside multiple attention heads, helping shift focus where needed. Instead of treating all inputs equally, it learns which data matters more. Temporal trends get noticed without being told when to look. Attention spreads across time steps, linking strong signals together.

What stands out about the Temporal Fusion Transformer is how it shows which inputs and moments in time matter most for each forecast. Because it reveals these key drivers, farmers can see exactly how factors like rain levels, heat shifts, or ground moisture affect outcomes. Seeing the weight behind predictions helps make sense of complex growing conditions without guessing. Instead of treating forecasts as black-box answers, users get a clear view into what shapes them. This transparency becomes useful when timing planting or managing resources based on weather history. When models point clearly to causes, choices around crop care gain stronger footing.

When predicting harvest amounts, farm data usually shows many shifting elements plus how nature's parts connect in tangled ways. What helps here is a system called TFT - it handles those links by blending slow-unfolding trends with smart spotlighting of key inputs. Important cycles like seasons or weather hints get noticed more because the method tunes its attention where it matters most.

XGBoost Model: One step at a time, XGBoost builds decision trees that fix mistakes of those before them. This method, known as Extreme Gradient Boosting, works well for sorting data into groups or predicting numbers. Instead of working alone, each tree adds small improvements through a process tied to gradients. Over rounds, the whole system gets sharper by focusing on what earlier versions missed. Performance climbs because every new piece targets past weaknesses directly.

What makes XGBoost stand out? It handles tricky patterns where inputs twist together in unpredictable ways. Picture farm data - rain, heat, dirt makeup, how crops are tended - all bumping into each other. Outcomes like harvest size rarely follow a straight line when these mix. Old-school models that expect neat lines tend to fall short here. What sets XGBoost apart is how it uses penalties to prevent overfitting, making predictions more reliable. Instead of just one task at a time, it splits work across processors, speeding things up noticeably. Missing data? It handles those gaps without needing extra steps beforehand. One useful byproduct: each input gets a score showing how much it sways the outcome. When predicting harvests, these scores spotlight which factors really move the needle.

Starting off different each time, the XGBoost method picks up on complex patterns in farm and weather data. Because it handles twists in relationships well, it fits alongside tools like Prophet and LSTM that track changes over time. While those focus on timing, XGBoost dives into how factors mix in messy ways. Together, they cover more ground than any one alone would. Their strengths line up without overlapping too much.

Hybrid CropIQ Architecture: A fresh blend of prediction tools - Prophet, LSTM, GRU, TFT, and XGBoost - comes together in CropIQ to handle varied farm data shapes. One after another, they bring different skills to the table, feeding into a stronger joint forecast than any could manage alone.A single workflow guides how the combined system runs across several phases of structure building.

First off, the Prophet model looks at past crop yield numbers to spot slow changes over years along with repeating ups and downs through seasons. By doing this, it reveals how farming output shifts across time in a structured way.

Later on, LSTM together with GRU handles time-based data, picking up how crop yields shift over brief periods. What these models do is study past records, linking weather changes directly to harvest results. They unfold patterns step by step, revealing how one season's conditions lead into the next outcome.

Midway through processing, the Temporal Fusion Transformer dives into the data by combining attention layers with dynamic feature filters. Instead of treating every input equally, it weighs certain variables more heavily when forecasting yields. This selective approach sharpens prediction quality while making decisions easier to trace. What sets it apart is how clearly each factor's role emerges during analysis. One thing about the XGBoost model is how it picks up on crooked patterns between farm-related elements - rain levels, heat, ground makeup, along with outside influences. What helps it predict better lies in spotting tangled links, something regular sequence-based systems might miss entirely.

When the models finish their work, one after another, their outputs join inside a shared system that shapes the last version of the crop forecast. Mixing timed data patterns, memory-style networks, focus-driven layers, along with stepwise improvement methods helps CropIQ stay sharp - its guesses hold steady, hit closer, perform well even when conditions shift.

D. Model Evaluation

The performance of individual and hybrid models is evaluated using:

- Root Mean Square Error (RMSE)

$$\mathrm{RMSE} = \sqrt{\frac{1}{n}\sum_{t=1}^{n}(A_t - F_t)^2}$$

- Mean Absolute Error (MAE)

$$\mathrm{MAE} = \frac{1}{n}\sum_{t=1}^{n}|A_t - F_t|$$

- Mean Absolute Percentage Error (MAPE)

$$\text{MAPE} = \frac{100}{n}\sum_{t=1}^{n}\left|\frac{A_t - F_t}{A_t}\right|$$

- Coefficient of Determination (R^2)

$$R^2 = 1 - \frac{\sum_{t=1}^{n}(A_t - F_t)^2}{\sum_{t=1}^{n}(A_t - \bar{A})^2}$$

E. Results and Visualization

Once the models finish learning, their forecasts get measured against real harvest numbers. This step checks how close the predictions come to what actually happened in the fields. The method includes both standalone systems and the combined setup. Looking at differences between guessed and recorded yields shows which approach works best. Results then appear in clear visuals to reveal patterns over time.

Looking at how well CropIQ works involves using various visual methods. Across multiple years, forecasted yields appear alongside real harvest numbers through time-based graphs. With these visuals, it becomes clear whether guesses stay near actual results over seasons. Patterns shifting year by year show up plainly, giving a sense of how steadily the system tracks farming changes. What stands out is its response to both short cycles and broader shifts in output.

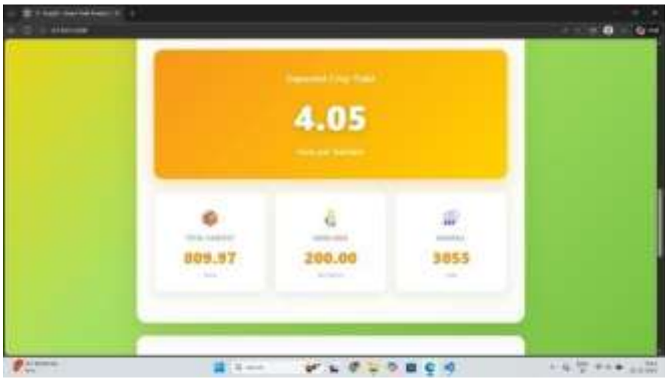

Fig 3.2 Yield Prediction Result

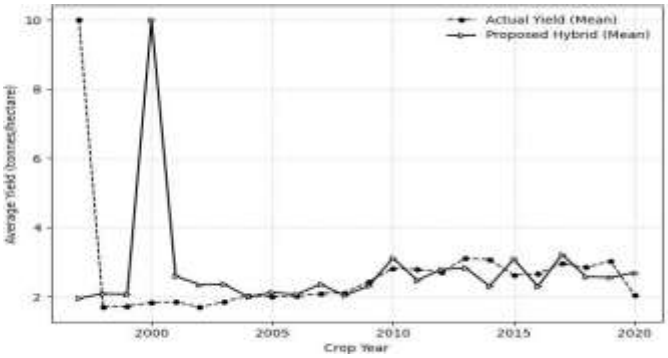

Fig 3.3 Comparison of actual and proposed hybrid mean yield

Look at figure 3.3 - it shows how maize yields changed each year, comparing real data with what the new combined model predicted. Though actual numbers jump around, the model's line moves more steadily, yet still follows the general path. Because it doesn’t swing wildly, its guesses stay closer to reality over time. Even when field results vary, this method holds firm, missing less than older ways might. Across years, it clearly picks up on slow shifts that pure models often overlook. Other researchers noticed alike gains - mixing deep learning layers helped their forecasts run smoother too.

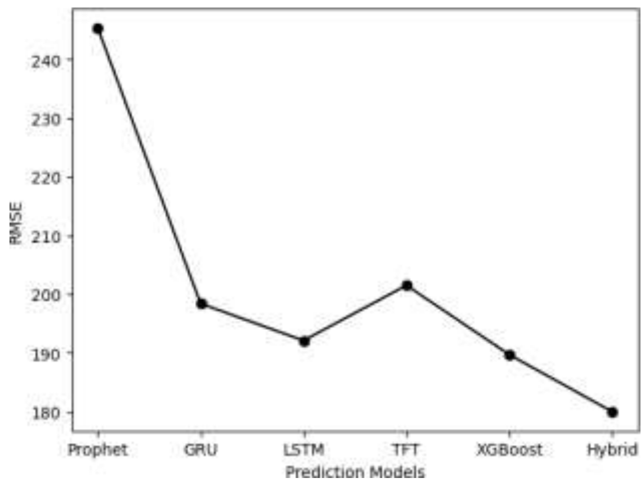

Fig 3.4 RMSE-based comparison of crop yield prediction models

The above chart represents that older stats methods show higher mistakes when guessing harvest sizes compared to newer tech. Because it struggles with messy real-world patterns, Prophet ends up with the biggest errors. After that, GRU and LSTM do better since they learn how past seasons affect future crops. What makes TFT stand out is not just good guesses but also showing which factors mattered most during predictions. XGBoost steps ahead by spotting tricky links between different farming conditions. Put together, the new mixed approach beats them all, landing the smallest margin of error overall.

Looking closer at how well single and combined models predict yields, Figure 3.4 shows their differences. A different view emerges when results are lined up side by side. Each model behaves in its own way under testing conditions. Some handle variation better than others do. The chart reveals patterns not obvious at first glance. Performance shifts depending on the approach used. What stands out is how outcomes change across methods

The RMSE scores shows how well each model predicts. As methods shift from basic Prophet to more complex ones, errors drop step by step. Prophet has the largest error, which means it struggles with patterns that aren't straight-line trends in farm-related numbers. Instead of ignoring time order, models like GRU and LSTM catch sequences better, especially LSTM, which does slightly better than GRU. On par with those in error size, TFT also gives clearer reasons why predictions happen thanks to built-in focus layers. What makes XGBoost stand out is how it captures tangled links between weather and farming data. Best results come through that mix - lowest error score shows it handles ups and downs better than others.

Model	RMSE	MAE	R^2	MAPE(%)
Prophet	245.3	187.2	0.78	12.4
XGBoost	189.7	142.8	0.89	8.7
GRU	198.4	151.3	0.87	9.3
LSTM	192.1	145.6	0.88	8.9
TFT	201.5	154.2	0.86	9.6

Table 3.1 Performance comparison of yield prediction models

This table shows how different forecast models stack up when judged by error rates and precision. Although smaller RMSE, MAE, and MAPE numbers point to more accurate predictions, an elevated R^2 suggests stronger explanation of data variation. Performance gains appear in the hybrid CropIQ system during testing, outpacing standalone models because it handles extended patterns, brief time-linked shifts, along with complex feature relationships. Despite differences across metrics, outcomes consistently favor the combined approach under varied conditions.

Looking at the numbers and charts, it becomes clear the CropIQ method works better for guessing harvest sizes. Because predictions are sharper, people who grow crops might plan planting with more confidence. Officials setting rules could adjust support programs based on these insights. Those organizing farm supplies may find timing fertilizer or water easier. Clearer forecasts help shape how land and materials get used across regions. With steadier estimates, choices about what to plant each season gain stronger backing.

IV. NOVELTY

With CropIQ, a fresh method takes shape - mixing several smart forecasting tools designed to spot different kinds of patterns in farm data. Instead of sticking to just one standard model, this setup weaves together time-tested statistical series analysis, modern neural networks, while layering in group-based learning strategies under one roof.

It combines five methods - Prophet, LSTM, GRU, TFT, and XGBoost - to handle tricky farm-related time data. One after another, they each pick up on distinct patterns hidden in the numbers. Long stretches of seasonality? That is where Prophet works best. Sequences that depend heavily on order show up clearly through LSTM and GRU. Paying sharp attention to timing clues happens inside the Temporal Fusion Transformer. Meanwhile, complex links between weather, soil, and harvests get mapped by XGBoost. Together, they cover more ground than any single method alone.

What stands out here is how this study builds a mixed prediction setup. Instead of relying on one model, it pulls results together using connections across several types. Statistical methods link with deep learning pieces through weighted sharing. Machine learning layers add further refinement by adjusting for patterns others might miss. Because each method covers gaps in the rest, the system handles messy farm data more smoothly.

Changes in crop behavior show up clearly now, even when conditions shift unexpectedly. The whole approach stays steady where older versions would have wobbled. Combining these styles wasn’t automatic - tuning mattered. Yet once balanced, the outcome tracks real-world swings without overreacting.

Right off the bat, clean data matters - so steps are taken early to shape raw inputs into something reliable. Instead of rushing ahead, adjustments happen along the way so values line up properly across features. From there, how well things work gets checked using more than just one yardstick. Picture it like this: errors big and small both get measured, not only average ones but extreme cases too. One metric watches percentage gaps, another tracks straight-line differences. Even the overall fit gets its moment under scrutiny, showing how tightly predictions match reality. Nothing slips through without being weighed somehow. Each angle adds clarity instead of piling on numbers. The whole setup leans on balance - no single score tells the full story. In the end, insight comes from looking sideways, forward, even backward at what sticks.

When tested, CropIQ worked better than single models at guessing right and staying steady. Because it forecasts more reliably, people who grow crops, study farming, or shape farm rules can plan smarter. Choices about planting, using water or land, even feeding communities, gain clarity through its insights.

V. MERITS AND DEMERITS

Starting off, the new CropIQ setup brings improvements beyond standard methods used to guess harvest amounts. Instead of relying on just one way, it pulls together various approaches - some based on numbers, others using smart algorithms, even those that mimic brain-like processing. It teams up tools like Prophet, LSTM, GRU, Temporal Fusion Transformer (TFT), along with XGBoost, so it can notice many kinds of patterns hidden in farm data. Things such as slow changes over years show up clearly, alongside repeating cycles each season, time-linked shifts, also complex links between factors that aren’t straightforward.

What stands out next is how much better the method predicts outcomes. One model alone usually misses parts of complicated farm-related information. Yet when combined, their strong points work together - making forecasts steadier and more trustworthy. Errors drop because the blend smooths out weaknesses, lifting results across the board. Performance climbs without leaning too hard on any single part.

Time moves forward, so the system handles it more smoothly now. Because they’re built for sequences, tools like LSTM, GRU, and TFT make sense of data that unfolds over time. When weather shifts season after season, these methods track how each change tugs at harvest outcomes.

On top of that, XGBoost helps capture how farming elements like rain, heat, and dirt interact in non-straightforward ways. Because of this, the system gets better at spotting hidden patterns in nature that shape how much food grows.

Still, the CropIQ setup comes up short in some areas. Training several models inside a combined system tends to slow things down. Models like LSTM or GRU chew through processing power. They also take far more time to train than older machine learning methods. The Temporal Fusion Transformer adds even heavier demands.

Farm experts get clearer insights because the Temporal Fusion Transformer shows which inputs matter most in its forecasts. Since it points out what drives outcomes, decision makers can focus on real causes behind harvest changes.

Putting it together, this mix of methods builds a full picture while growing easily when needed. One big plus - it fits right into actual farming choices without skipping steps.

Here lies a catch: forecasts need solid past farming records to work right. Without clear timelines and key nature factors neatly lined up, guesses go shaky. Many countryside areas lack full or steady logs, leaving systems guessing more than learning. That gap shows plainly when results start drifting off track.

Achieving strong results with the hybrid setup often depends on precise adjustments to its internal settings. When those values are off, the system might memorize training examples too closely instead of adapting well later. Performance drops appear especially when facing new, unfamiliar inputs.

VI. APPLICATIONS

CropIQ's mix of methods shows real uses across farming and data work. Because it forecasts harvest sizes well, people like farm owners might rely on its insights. Researchers watching plant growth could find value too. Those shaping food policies may lean on its results when planning ahead. Even folks managing delivery networks might adjust based on what the model suggests. Its strength lies in helping different roles act with clearer expectations. What matters is how steadily it delivers usable estimates.

A main use of this new system shows up in how farms get organized and crops are handled. Because growers see what harvests might look like, they shape their planting choices around that insight. Picking which seeds to sow ties directly to those outlooks instead of guesswork. When forecasts give clearer pictures, timing for sowing and pulling crops out shifts smarter. Fertilizer amounts adjust more precisely, just like water delivery and crew assignments do. Seeing ahead helps balance inputs without waste creeping in. Decisions gain clarity when numbers reflect real field behavior over time.

When it comes to farming rules and keeping food supplies stable, forecasts play a key role. Officials might look at projected yields to spot gaps or extra stock ahead of time. With those insights, shaping farm-related decisions becomes more grounded. Handling how food moves from farms to people grows clearer too. Stability in access to meals across areas and countries ties back to these steps.

Now imagine using this method to study farm markets and guess future prices. When we predict how much crops will grow, it shapes what gets sold, setting the stage for cost shifts. Knowing these patterns gives sellers, farms, and number watchers clearer sight on where things are headed across fields and shelves.

Besides this, farmers gain better control when predictions mix live updates on rain, ground quality, or space-based images. A link between forecasts and current field states reveals how crops are likely to perform. That insight helps guide choices - less guesswork means less waste. Outcomes shift toward smarter harvests without extra inputs.

Farmers might see benefits from CropIQ because it adjusts easily to different growing conditions. Since it learns how crops change over time, decisions around planting or harvesting could become more accurate. Patterns in weather and soil matter get noticed without clear rules guiding every step. This setup helps track what works in fields where conditions shift often. Results may support longer-term sustainability by showing hidden links between inputs and yields. Tools like this one tend to grow stronger with real-world use instead of staying fixed.

ACKNOWLEDGMENT

With a deep sense of gratitude and respect, we would like to express our sincere thanks to our Guide for his valuable guidance, encouragement, and continuous support throughout the completion of this project. His kind attention and insightful suggestions played a significant role in making this research work successful. It has been a privilege to work under such experienced and knowledgeable mentors who guided us in understanding and addressing real-world problems. Apart from technical knowledge, we also learned the importance of accuracy in documentation and research presentation, which greatly helped us during project reviews and presentations. We would like to thank all faculty members and individuals who directly or indirectly contributed to the successful completion of this research work.

REFERENCES

[1] K. Meena and B. Chaitra, "A Novel Framework Using Deep Learning Techniques for Ragi Price Prediction in Karnataka," IEEE Access, vol. 12, pp. 136103–136119, Sept. 2024.

[2] H. H. Wang, Y. C. Wang, B. L. Wee, J. Y. Khoo, and F. B. Yahya, "Optimizing Crop Yield Prediction: A Hybrid Approach Integrating CNN and LSTM Networks," Journal of Theoretical and Applied Information Technology, vol. 102, no. 22, pp. 8075–8083, Nov. 2024.

[3] M. Abdel-Salam, N. Kumar, and S. Mahajan, "A Proposed Framework for Crop Yield Prediction Using Hybrid Feature Selection Approach and Optimized Machine Learning," Neural Computing and Applications, vol. 36, pp. 20723–20750, 2024.

[4] V. Kiran Kumar, K. V. Ramesh, and V. Rakesh, "Optimizing LSTM and Bi-LSTM Models for Crop Yield Prediction and Comparison with Traditional Machine Learning Techniques," Applied Intelligence, vol. 53, pp. 28291–28309, 2023.

[5] K. Jhajharia, P. Mathur, S. Jain, and S. Nijhawan, "Crop Yield Prediction Using Machine Learning and Deep Learning Techniques," Procedia Computer Science, vol. 218, pp. 406– 417, 2023.

[6] S. Jeong, S. Moon, and J. H. Kim, "Hybrid machine learning model for crop yield prediction using climate data," IEEE Access, vol. 8, pp. 196–206, 2020.

[7] R. M. Torres, A. S. Prasad, and V. R. Rao, "Crop yield prediction using LSTM recurrent neural networks," International Journal of Advanced Computer Science and Applications, vol. 11, no. 5, pp. 186–193, 2020.

[8] Y. Khaki and L. Wang, "Crop yield prediction using deep neural networks," Frontiers in Plant Science, vol. 10, pp. 1–9, 2019.

[9] S. Makridakis, E. Spiliotis, and V. Assimakopoulos, "Statistical and machine learning forecasting methods: Concerns and ways forward," PLOS ONE, vol. 13, no. 3, pp. 1–26, 2018.

[10] A.Kamilaris and F. X. Prenafeta-Boldú, "Deep learning in agriculture: A survey," Computers and Electronics in Agriculture, vol. 147, pp. 70–90, Apr. 2018.