Journal of Management & Entrepreneurship
ISSN 2229-5348
UGC Care Group I Journal
Vol-10 Issue-01 Feb 2021

Improving Lives of Indebted Farmers Using Deep Learning: Predicting Agricultural Produce Prices Using Convolutional Neural Networks

Borusu Leela Sai Padmavathi, Ambati Teja Sai Kumar, Dr. RVVSV Prasad
SWARNANDHRA College of Engineering and Technology
Department of Information Technology

Abstract

Nations are working hard to find solutions to the growing societal issue of farmer suicides. Uncertainty and volatility in product pricing caused by changing market circumstances are the main reasons why most farmers end their lives because they can't sell their goods at the profit levels they wish. In this research, we introduce PECAD, a deep learning algorithm that accurately predicts future produce prices by analysing previous pricing and volume trends. This method aims to resolve the problem of produce price uncertainty and avoid farmer suicides. Though earlier research has offered ML algorithms for produce price prediction, these have two major flaws: (i) they use classical ML prediction models, which don't always work well with spatio-temporal datasets, and (ii) they don't take into account the spatial-temporal dependence of future prices on past data. In order to address these limitations, PECAD has three main contributions: first, we use an official Indian government website to collect real-world daily price and (produced) volume data for various crops over an 11-year period; second, we use state-of-the-art imputation techniques to pre-process this raw dataset and account for missing data; and third, we propose a novel wide and deep neural network architecture that consists of two separate convolutional neural network models trained for pricing and volume data, respectively. We find that PECAD significantly reduces the root mean squared error (RMSE) compared to state-of-the-art baselines (by 25% less) in our simulations, making it the superior technique. In the Indian state of Jharkhand, where we collaborate, a non-profit organisation is looking at the possibility of using PECAD to reduce the number of farmer suicides.

Introduction

Suicide rates among small-scale farmers have risen significantly in the past twenty years, disproportionately in developing nations like Pakistan, India, and others, as a result of agrarian distress and associated socio-economic issues like debt, loss of agricultural income, etc. Approximately 300,000 farmers in India have taken their own lives. ever since 1995. In the Indian state of Maharashtra, 60,000 farmers took their own lives in 2014, with a daily average of 10 (NCRB 2019). Crop failures, poor farm production, not being able to make a profit, inefficient cold chain management leading to wasted agricultural products, inadequate irrigation facilities, and overwhelming debt are just a few of the many issues that might cause farmers to take their own lives. One of the main causes of farmer suicides is the unpredictability of agricultural pricing and markets. For example, changes in global market circumstances may cause local prices of agricultural goods to fluctuate wildly (Barik 2018). This price volatility makes it impossible for small-scale farmers with debts to determine when and where to sell their crops, as they lack the sophisticated technical tools and understanding of global market circumstances that larger-scale farmers have. Many of these farmers end their lives because they are unable to repay their agricultural debts and obtain the profits they seek from their produce (Panagariya 2008).

(a) Farmers Protesting by Throwing their Unsold Produce (b) Huge Demand for Loan Waiver at Farmer Rally

Figure 1: Agrarian Distress in India

As a result, the problems faced by these farmers must be addressed promptly. Thanks to recent developments in ML approaches, learning algorithms may now be used effectively to many societal challenges (Tambe and Rice 2018). In order to address the issues mentioned before, this study suggests using AI and ML to find the answer to the following question: Is it possible to forecast the future price of agricultural produce at various marketplaces using data-driven methods that take into account past pricing and volume patterns? After that, these AI/ML methods can
use it to make smart decisions about when to sell their crops; for instance, based on projected prices, farmers may plan ahead and sell their crops at the peak of profit. To get to the bottom of this mystery, we need to solve a number of problems. To begin, training is made more difficult by the fact that the current datasets on price patterns1 are quite sparse, meaning they include many missing items. Secondly, it is important to create prediction models that can explicitly account for the spatio-temporal dependence between past prices and future produce prices. For example, the price of tomatoes in August 2019 could be influenced by their price in August 2018, and prices at nearby markets could be similar to those at faraway places. We validate this in our experiments, but previous work on produce price prediction algorithms (e.g., decision trees) often fails to account for the spatio-temporal dependence of future prices on past data. Additionally, these algorithms rely on classical ML prediction models, which do not take this dependence into account. Because of these flaws, these procedures aren't very practical or accurate. In order to overcome these limitations, we present PECAD, a new neural network architecture for predicting agricultural product prices in the future.

PECAD stands for Price Estimation for Crops using the Application of Deep Learning. The following innovative additions are made by PECAD to address the inadequacies of earlier work. An official Indian government website, Agmarknet.gov.in1, was used to compile the actual pricing and (produced) volume of various commodities at 1,350 agricultural markets in India from 2008 to 2018. Second, in order to prepare this raw information for analysis, PECAD employs cutting-edge imputation (and other) methods to fill in any gaps in the data. Finally, PECAD takes this data and suggests a new way of building neural networks that combines deep learning with broad linear models (Cheng et al., 2016). Nevertheless, PECAD employs an innovative approach by combining two distinct convolutional neural network (CNN) models—one for pricing data and the other for volume data—for the crop in question. These CNN models are then fed into the wide linear model, rather than cross-product feature transformations. While baseline approaches obtain a coefficient of variance that is 25% lower than PECAD, our simulation findings demonstrate that PECAD significantly outperforms them. This highlights the importance of
while processing data collected by remote sensing. But their method depends on collecting field photos from satellites, which may be costly in underdeveloped nations with limited resources. To forecast future rates, our team uses publicly accessible pricing and volume data. Then, a software and hardware solution was suggested by (Chen, Nowocin, and Marathe 2017) to lessen agricultural produce spoiling. Our work is most closely connected to that of (Ma et al. 2019), as they, too, use the same data source to construct a model for predicting crop prices1. We demonstrate in our studies that their low performance accuracy is due to their failure to take use of the spatio-temporal features of price and (produced) volume data for different crops. We find spatio-temporal relationships in price and volume data by using certain kinds of convolutional neural networks.

Dataset Construction

Data Collection There are two main databases that we use. A website maintained by the Indian government's Ministry of Agriculture and Farmers Welfare, Agmarknet.gov.in1, provides us with all of the raw data we need on agricultural products (produce). This website has twelve years' worth of daily pricing and volume data from thirteen hundred and fifty-two agricultural marketplaces throughout India. Using data gathered from all marketplaces over an 11-year period (2008–2018), we analysed the price and volume of three crops: brinjal, tomatoes, and chilies. We successfully scraped the market data from this website using a multi-process crawler script that was installed on two cloud servers. It took one week to finish this tyre scraping operation. Furthermore, we supplement this information by gathering geographic features, such as the exact location of each agricultural market. Since crop prices in nearby marketplaces tend to be comparable, we gather this information to record the spatial correlations between these markets. We use the Google Maps API to extract the geometric coordinates (latitude and longitude) for every market in the Agmarknet.gov.in database as the website does not provide any geographical information about the 1352 marketplaces that are there. In addition, we use sparse one-hot encoding vectors to capture this characteristic, and we provide a unique ID to each market and crop.

Data Preprocessing Let $\boldsymbol{M}$ denote the set of all 1352 markets in our dataset, $\boldsymbol{C}$ denote the set of produce types(we collect data for three crops, so $\boldsymbol{C}$ = 3), and $\boldsymbol{T}$ de- note the set of all dates (timesteps) for which we have priceand volume entries. For each crop $c \in \boldsymbol{C}$, we define P^c andV^c as $\boldsymbol{M} \times \boldsymbol{T}$ price and volume matrices (respectively). Forportance of explicitly modeling the spatio-temporal depen-each $m \in \boldsymbol{M}$ and $t \in \boldsymbol{T}$, P^c indicates the price of crop c dence of future prices on past data inside our ML algorithm. in market m on day t, whereas V^c indicates the volume of Our work is done in collaboration with a non-profit agencythat works on preventing farmer suicides in the Indian stateof Jharkhand (name withheld for anonymity), and PECAD iscurrently being reviewed by them for potential deployment.**Related Work** We discuss prior AI/ML research that as- sists in alleviating agrarian distress. (You et al. 2017) pro-posed deep Gaussian processes to predictcrop yields us-crop c (in metric tonnes) that arrived in market m on day t.

Unfortunately, the P^c and V^c matrices for each crop$c$ $\boldsymbol{C}$ (which we construct after data collection) are ex- tremely sparse, i.e., they have several missing entries. On Agmarknet.gov.in, these missing entries are created due to avariety of reasons, e.g., a particular market might have beenclosed on a given day t $\boldsymbol{T}$, no produce was sold in a mar- ket on a given day, or simply the data for that market was never recorded due to human errors. In particular, we ob-

Feature	Explanation	Notation
Market	Unique identifier for each market	$m \in \boldsymbol{M}$
Crop	Unique identifier for each crop	$c \in \boldsymbol{C}$
Price	Denotes the price of crop c in market m on day t	$P^c_{m,t}$
Volume	Denotes the volume of crop c in market m on day t	$V^c_{m,t}$
Geo location	Denotes the geographical latitude/longitude of market m	$[lat_m, lon_m]$

Table 1: Features and notations in our paper.

Consider that there is a severe lack of data for some markets; that is, there are relatively few legitimate (non-empty) data points for certain markets, and the data from these markets does not contribute much to the learning process as a whole. Hence, we exclude all markets from the data set if their validation data is missing for more than 10% of the days in T. The next step is to fill in the remaining blanks using efficient data imputation techniques. Naive ̃imputation techniques, such as hot-deck and mean substitution, cannot be used since the crop price and volume data is sparse. Nevertheless, in order to fill in the gaps in our dataset, we use SoftImpute, a cutting-edge collaborative filtering approach (Hastie et al., 2015), which takes into account the spatial correlations between crop prices (and volumes) at nearby marketplaces.Each crop C has fully filled Pc and V c matrices after data imputation. The inability of most sequential neural networks to learn long-term temporal dependencies is due to vanishing or exploding gradients (Sutskever, Vinyals, and Le 2014). To circumvent this, we compress the Pc and V c matrices by treating a time window of w days as a single time step. Specifically, we average crop prices and quantities for each non-overlapping consecutive block of w days in order to get compressed 1 minus (t+1)w Deep and Wide Networks Figure 2 shows the wide and deep network model, which is made up of deep neural networks and jointly trained wide linear models. This model is great for large-scale regression problems with sparse inputs, like features with a lot of possible values, like categorical features (Cheng et al. 2016). Since our price prediction dataset includes feature vectors for identifying markets and crops that are very sparse and encoded using one-hot encoding, the broad and deep learning paradigm is a perfect fit for PECAD. As seen in Figure 2 (right), the deep component is a feed-forward neural network. In the first layer of the deep component, sparse one-hot encoding vectors with large dimensions are transformed into dense real-valued vectors with low dimensions; these are called embedding vectors. The neural network's hidden layers are then fed these dense embedding vectors (see to the right side of Figure 2 for details).

Figure 2 (left) shows the broad component as a generalised linear model (GLM) using the equation y = wT x + b. A vector of d characteristics is represented by x = [x1, x2,..., xd], the model parameters are w = [w1, w2,..., wd], the bias is b, and y is the prediction. Crucially, the broad linear model's feature vector x incorporates non-linearity derived from cross-product transformation features that record interactions between the input binary characteristics.

Data Characteristics Our final dataset has 40000 data-points, each consisting of a feature vector and a continuous label. A single feature vector for the t^{th} time-step at market m consists of historical price and volume pairs for the last n time-steps from our compressed $\hat{P}^c$ and $\hat{V}^c$ matrices, along with market latitude/longitude coordinates, and market and crop identifiers. The ground-truth label (which we want to predict) is the price of crop c at market m on the $(t+1)^{th}$ time-step (i.e., the crop price in the next time-step). Table 1 describes a list of all features in our dataset.

Deep Learning Algorithm

Our innovative deep learning architecture, PECAD, is based on broad and deep networks, as described by Cheng et al. (2016). Our PECAD design is based on broad and deep networks, as well as temporal convolutional networks (TCN) (Bai, Kolter, and Koltun 2018). To ensure completeness, we first provide a brief description of these networks.

Temporal Convolutional Networks According to Bai, Kolter, and Koltun (2018), the TCN model handles sequential data using convolutional layers. Like regular RNN models, TCNs may accept input sequences of any length and produce output sequences of the same length. The TCN model achieves this by using a 1D fully-convolutional network (FCN) design. In this design, the input layer and each hidden layer have the same length, and to ensure that following layers have the same length, a zero padding of length (filter size 1) is added. To further guarantee that no information leaks from the future into the past, TCN employs causal convolutions, where an output at time t is convolved exclusively with components from time t and earlier in the prior layer. Lastly, TCN employs dilated convolutions (which causes the receptive field of convolutional filters to grow exponentially) to extract correlations from long-term sequences. TCN has shown to be more effective than cutting-edge recurrent architectures like LSTM on a variety of tests and benchmarks. Because TCN has so many useful features, we include it as a foundational element in PECAD.

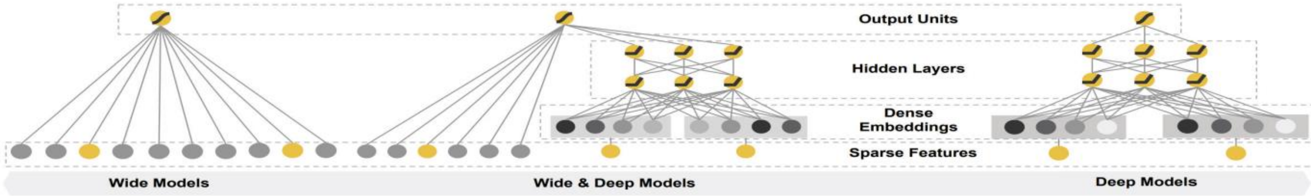

Figure 2: Wide and Deep Network Architectures (Cheng et al. 2016)

PECAD: Deep Learning Architecture

After this, we will go into PECAD's deep learning architecture and its wide-ranging features. While PECAD's deep neural networks generalise to previously undiscovered feature interactions via low-dimensional embeddings, the broad linear models memorise long-term sequential pricing/volume information. Keep in mind that in conventional deep and wide networks, the wide component's feature vector x contains cross-product transformation features (Cheng et al. 2016). Standard wide and deep networks struggle to learn from the produce price prediction issue because the number of characteristics used grows exponentially with the duration of the price/volume history being considered. Instead of adding an exponential number of cross-product transformation features, PECAD trains two independent TCN models for price and volume, which adds non-linearity to the broad component (GLM). This is a significant innovation within PECAD. Through our tests, we have shown that the broad component model (GLM) benefits from having the TCN models as input. We begin with a brief summary of PECAD's overall architecture. We then go on to explain PECAD's embedding layer and the inner workings of its deep and broad networks in depth.

Architecture Overview You can see the whole PECAD architecture in Figure 3. Figure 3, on the left, shows the results from the deep network, and on the right, shows the results from the broad network (GLM). In order to train separate TCN models that use raw historical price and volume patterns as input (respectively) and construct complicated non-linear features that compose the feature vectors for the wide network, we avoid using the raw input features in the wide network GLM's feature set. In contrast, the deep network's feature vector contains market and crop embedding vectors, geographical characteristics of agricultural marketplaces (such as geo-graphical latitude and longitude coordinates), and The sparse high-dimensional datapoints, which include market and crop one-hot encoding vectors, are transformed into low-dimensional embedding vectors ve,i = W · vi by applying an embedding layer. Here, ve,i ∈ Rde represents the ith datapoint's embedding vector and vi Rdx represents the ith datapoint with one-hot encoding (de < dx). To reduce loss as much as possible, the embedding parameter matrix W RdeÏdx is first set up at random and then updated while the model is being trained.

PECAD Deep Network A feed-forward deep neural network (DNN) is what the deep network uses to process the low-dimensional embedding vectors. The DNN has three fully connected hidden layers that are rectified linear units (ReLUs). One of these layers is the (l + 1)-th hidden layer, which is represented as hl+1 = ReLU(Wlhl + bl). The weights and bias for the l-th fully-connected layer are bl Rdl+1×dl and Wl Rdl+1×dl, respectively.

PECAD Wide Network A GLM receives its output from two TCNs in the broad network, which were trained independently on sequential price and volume data, respectively. The standard practice is to combine several time series variables into one TCN network; for example, a price and volume TCN model might be trained using the same data. Nonetheless, PECAD teaches two separate TCN models to memorise consecutive price and volume data over the long term. In our studies, we compare the predictive performance of PECAD—which uses two independent TCN models—against that of a version of PECAD—PECAD-Single TCN—in Table 2. This allows us to experimentally support the decision to use such a configuration.

Training Procedure We create a training set and a test set by dividing the data using time. The training data consists of price and volume records from 2008 to 2016, in addition to the other characteristics listed in Table 1. Utilising this training data, we educate PECAD. The test set consists of data collected between 2017 and 2018. Last but not least, we use an L2 to process the complex non-linear feature vectors that were input into the broad network (refer to Figure 3). As a last point, the loss function, i.e., $L_2 = \Sigma_n (\hat{y} - y_{predicted})^2$.

wide and deep networks is combined and fed into a single fully connected layer, which outputs a prediction of the produce price on the next day.

Embedding Layer We assign unique identifiers for each of the M = 1352 markets and C = 3 crops, and rep- resent this feature using extremely sparse one-hot encoding feature vectors, which leads to poor learning performance.

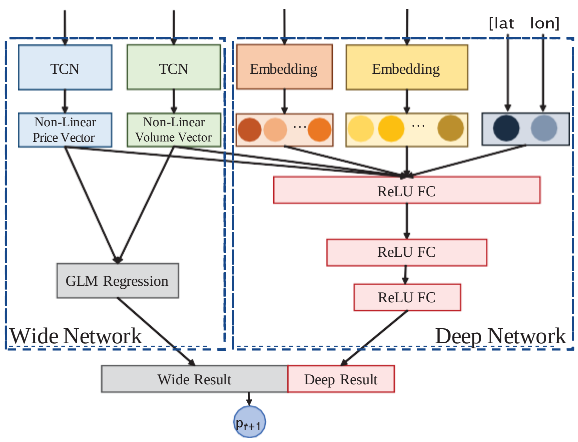

Figure 3: Architecture of our deep learning model

Experimental Evaluation

An Amazon Machine Image (AMI) server running Ubuntu version 24.0 was used for all tests using deep learning. All of our experiments make use of the feature space in our dataset, which includes price and volume data over the previous n = 360 days. The performance of all deep learning models, including PECAD and others, is evaluated after 150 epochs of training.

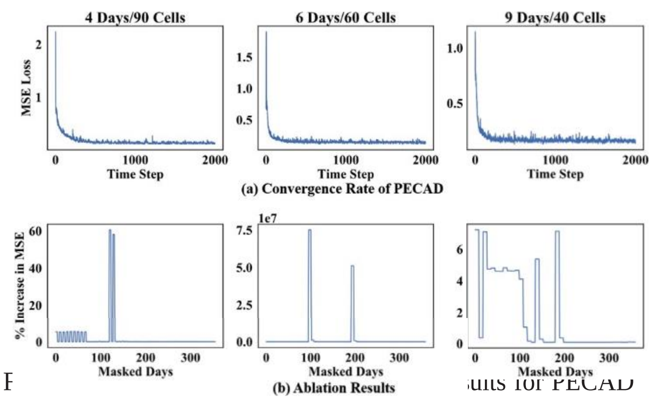

F ults for PECAD

hit 20 balls each inning. We use the "coefficient of variation" (i.e., the root mean squared error (RMSE) divided by the mean produce price) to compare the predictive performance of various algorithms, rather than RMSE values, which cannot be meaningfully compared across different models to ascertain which model offers better outcome predictions (Sørensen 2002).

Baselines Two traditional ML models, random forests (RF) and gradient tree boosting (GTB), are used for comparison. Since these two baselines are the most effective algorithms for predicting the prices of product, we utilise them (Ma et al. 2019). There are four deep learning models that we also compare against; these models make use of spatio-temporal features: (i) PECAD using a single TCN for both price and volume sequences (PECAD-Single TCN); (ii) attention-LSTM networks (Sutskever, Vinyals, and Le 2014); (iii) standard wide and deep networks (Standard Wide & Deep) with cross-product transformation features (Cheng et al. 2016); and (iv) standard (TCN) model.

Predictive Performance In Table 2, we can see the performance of several ML models for brinjal, tomatoes, and chillis over three separate time frames (w=4,6, and 9 days). Consideration of the spatio-temporal dependency of future product prices on previous data yields beneficial results, as shown in Table 2. When compared to traditional ML techniques, our deep learning models that do this task considerably improve the accuracy of produce price predictions. We see that deep learning models always have the lowest coefficient of variation (bold) for all crops and w values. Deep learning models outperform the two classic ML models by 12.5% when it comes to the coefficient of variance.

Further, with a 25% decrease in coefficient of variation relative to their average case performance, PECAD substantially surpasses the other four deep learning models in Table 2. By achieving a 13% lower coefficient of variation, the work model (Cheng et al., 2016) demonstrates the limits of using traditional cross-product transformation characteristics in price prediction when compared to the industry-standard deep and wide net. Differences in performance of 13.5% between PECAD and PECAD-Single TCN demonstrate the benefits of training two separate TCN models in the PECAD architecture. As can be seen from the image, PECAD is the go-to programme for projecting product prices into the future.

Convergence Results The rate of convergence of PECAD is examined in Figure 4a for various time window widths, averaging all three crops. Mean squared error (MSE) training loss is shown on the Y-axis, while rising time epochs are shown on the X-axis. It is evident from this image that PECAD reaches locally optimum solutions quite fast.

Ablation Studies We investigate how different components of our feature space affect PECAD's prediction accuracy. As a result, we try out several ablations of our PECAD model that we get by gradually removing important features from the feature space. For each of the past n = 360 days, we mask the price and volume inputs to construct ablated models. Ablating various regions of the feature space results in the effects seen in Figure 4b. To get an ablated model, the X-axis displays the day for which the price/volume data are masked. Across all three crops, the average % increase in MSE loss due to ablation is shown on the Y-axis. For instance, in the PECAD model trained with a 4-day time frame, the MSE loss rises by 5% when the most recent day's price and volume inputs are masked (X-axis label = 1). Figure 4b reveals an unexpected finding: the MSE loss increases dramatically when price/volume entries of days in the past around days 100 to 150 are masked, suggesting that these days have a major role in forecasting future product prices, and this holds true across all three time window widths. The fact that our three crop varieties—brindaal, tomato, and chilli—have set planting dates each year and an average growing duration of three to four months (DARD 2019) suggests that new supplies of food are introduced to the market every three to four months (90 to 120 days)—which might account for the observed outcome. Therefore, entries for fresh product deliveries recorded three to four months ago, together with their accompanying volumes and prices, may be useful indicators of crop prices for the day after.

	4 Days/90 Cells			6 Days/60 Cells			9 Days/40 Cells		
	Brinjal	**Tomato**	**Chilli**	**Brinjal**	**Tomato**	**Chilli**	**Brinjal**	**Tomato**	**Chilli**
RF	21.12	22.88	19.45	23.47	38.48	21.60	24.50	44.30	23.54
GTB	21.38	20.85	18.99	22.88	26.18	18.58	23.64	31.55	20.36
Attention-LSTM	19.88	20.52	17.49	21.98	24.36	18.44	**21.00**	31.94	21.04
TCN	20.59	**19.87**	17.36	54.42	33.25	27.69	27.59	98.02	81.83
Standard Wide & Deep	23.63	24.47	19.07	24.34	28.22	18.67	27.36	34.29	21.10
PECAD - Single TCN	21.90	23.50	17.77	29.43	30.86	20.21	26.26	33.65	20.46
PECAD	**19.64**	21.62	**17.07**	**21.14**	**24.20**	**17.65**	21.75	**28.46**	**19.31**

Table 2: Coefficient of Variation of different ML models with varying time window sizes

Implementation Challenges & Conclusion

When non-profit organisations assisting farmers with debts install PECAD, there are a few issues with implementation that must be addressed. To start, past weather patterns may influence future crop availability (and, by extension, crop pricing), thus including them into PECAD's prediction performance might be a good idea. As a result of the superior accuracy of physical models, deep learning approaches are hardly used for weather modelling in the actual world. As a result, physical weather prediction models should be connected with PECAD (in future study). Moreover, farmers with low levels of education may be suspicious of advanced deep learning methods (like PECAD) that attempt to forecast future commodity prices. To allay these concerns and encourage participation, the agencies involved in this programme should launch public awareness efforts. On top of that, purchasing sophisticated computer gear (to teach and operate PECAD) is usually not a top priority for non-profit groups due to their low finances. As a result, we advise the agencies to utilise PECAD independently by implementing it as a web service. Last but not least, PECAD is only one piece of the jigsaw that requires fixing in order to stop farmer suicides. If long-term crop price and volume trends are not available, for instance, PECAD's implementation will fail miserably. Although this data is accessible for Indian markets via Agmarknet.gov.in, no analogue data archives exist for other emerging nations.Using historical pricing and volume trends, this article introduces PECAD, a deep learning system that accurately predicts future crop prices. Prior ML systems for crop price prediction suffered from serious flaws due to their failure to take into account the spatial-temporal dependency of future prices on historical data. To address these problems, PECAD suggests a new deep learning architecture that uses two convolutional neural network models—one for price data and the other for volume data (pertaining to the crop in question). By reducing the coefficient of variation by 25% compared to state-of-the-art baseline approaches, PECAD beats them in our simulation results. One of our partners is a non-profit in the Indian state of Jharkhand whose mission is to reduce the number of farmer suicides; they are now considering using PECAD in their work.

References

This is a 2018 publication by Bai, Kolter, and Koltun. Modelling sequences with generic convolutional and recurrent networks: an empirical evaluation. published as a preprint on arXiv:1803.01271. Narik, Barik. 2018. Research on ways to alleviate farmer stress in Rajasthan. This document may be accessed at this URL: https://www.copenhagenconsensus.com/sites/default/files/raj farmer distress sm.pdf. Marathe, N.; Nowocin, K.; and Chen, G. H. 2017. An integrated software and hardware approach to lowering crop spoiling and raising small-scale farmers' incomes in India. Accepted for publication in the proceedings of the 2009 International Conference on ICT&D. The authors of the article are Cheng, H.-T., Koc, L., Harmsen, J., Shaked, T., Chandra, T., Arad-hye, H., Anderson, G., Corrado, G., Chai, W., Ispir, M., and others. 2016. Wide & deep learning for recommender systems. Volume 1, Issue 1, Pages 7–10, 2019 (ACM DARD/Deep Learning Workshop on Recommender Systems). Crop Duration in KwaZulu-Natal: A Look at Vegetable Production. the duration of the growing season.pdf, found at https://www.kzndard.gov.za/images/Documents/Horticulture/Veg prod. In 2015, Hastie, Mazumder, Lee, and Zadeh published a paper along these lines. Quick alternating least squares for matrix completion and low-rank svd techniques. Volume 16, Issue 1, Pages 3367–3402, Journal of Machine Learning Research. These authors are Ma, W., Nowocin, K., Marathe, N., and Chen, G. H. The year 2019. Using collaborative filtering and adaptive closest neighbours, an interpretable system is developed to anticipate crop prices for small and marginal farmers in India. Chapter 6 of the ACM NCRB 2019 Proceedings: Tenth International Conference on Information and Communication Technologies and Development. The National Crime Records Bureau operates. Visit http://ncrb.gov.in/site. Arapana, A. (2008). Indian Subcontinent: A Rising Power. Press of Oxford University.According to Sørensen (2002). The coefficient of variation, its use, and its abuse in the field of organisational demography. Results from the field of sociological study, volume 30, issue 4, pages 475–491. In 2014, Sutskever, Vinyals, and Le published a paper. Neural networks for sequence-to-sequence learning. On pages 3104–3112 of the book Advances in Neural Information Processing Systems. M. Tambe and E. Rice published in 2018. Machine Learning and Human Services. Applying AI for the Greater Good. Press of Cambridge University. In 2017, You, J., Li, X., Low, M., Lobell, D., and Ermon, S. Utilising remote sensing data, a deep gaussian method is used to forecast agricultural yields. Pages 4559–4565 in the proceedings of the 31st annual AAAI conference on artificial intelligence. Published by AAAI.