Textile Intelligence Stack: A Multi-Model AI Framework for Trend Prediction, Demand Forecasting and Fabric Recommendation in the Garment Manufacturing Industry

Gaurish Jariwala [1], Wallace Scott Constancio Dsouza [2], Sanket Bhat [3],Meghanasree GS [4], Vishvaa G [5], Dr. Pushpa J[6]

[1,2,3] MCA, Dept. of Computer Science & IT, Jain Deemed to-be University

[5] MCom, Dept. of Commerce, Jain Deemed to-be University

[3] MSC, Animation Jain Deemed to-be University

[6] Professor, Dept. of Computer Science & IT, Jain Deemed to-be University

How to Cite this Article:

Jariwala, G., Dsouza, W. S. C., Bhat, S., GS, M. & G, V. (2026). Textile Intelligence Stack: A Multi-Model AI Framework for Trend Prediction, Demand Forecasting and Fabric Recommendation in the Garment Manufacturing Industry. International Journal of Creative and Open Research in Engineering and Management, <i>02</i>(04). https://doi.org/10.55041/ijcope.v2i4.498

https://doi.org/10.55041/ijcope.v2i4.498

Abstract

The growing demand for efficient textile manufacturing and the urgent need to reduce environmental waste from overproduction has fueled global interest in the application of artificial intelligence to predicting fashion trends, predicting demand, and selecting sustainable materials. In this paper, we present a proposed comprehensive framework, Textile Intelligence Stack, that integrates four AI model components, including a fine-tuned convolutional neural network for visual trend extraction, a BERT-based natural language processing model for social media signal mining, a long-term memory network for time-series demand forecasting, and an XGBoost classifier for clothing fabric recommendation. Each component is examined in terms of architectural design, data requirements, training methods, and contribution to the overall prediction pipeline. Traditional statistical forecasting methods, such as ARIMA, provide interpretable input data but are subject to significant forecast errors and fail to capture the non-linear, trend-driven nature of fashion demand. The proposed multi-model architecture addresses these limitations by integrating heterogeneous data sources, including social media platforms, e-commerce transaction records, and fashion trend images, into a single manufacturing analytics pipeline. The fabric recommendation module is particularly novel in the existing literature by providing environmentally sensitive material recommendations at the manufacturer level, a functionality that has not been implemented in any previous unified system. Beyond the performance of individual models, this article highlights the critical role of multi-source data fusion, transfer learning, and explainable AI in improving forecast accuracy, recommendation transparency, and practical adoption in textile manufacturing companies. By integrating multidisciplinary insights from deep learning, natural language processing, and sustainable manufacturing research, this study demonstrates that smart textile production requires a unified AI pipeline supported by diverse data streams and manufacturer-oriented decision support rather than relying on a single forecasting model or data source.

Keywords: Artificial Intelligence, Textile Industry, Demand Forecasting, Fashion Trend Prediction, CNN, BERT, LSTM, XGBoost, Fabric Recommendation, Sustainable Manufacturing, Indian Textile Industry

1. INTRODUCTION

The textile industry has been part of human civilization for thousands of years, in many ways its basic decision-making processes have not kept pace with modern technological tools. Today, walk into any medium-sized textile manufacturing factory in Surat, Tirupur or Ahmedabad and you will find the same story repeated. Production managers look at last season's sales, compare them with buyers' rough estimates of what will sell next quarter, and order fabrics worth thousands of rupees based on a combination of experience, intuition, and desire. Sometimes it works. Often this is not the case. The remaining inventory remains in the warehouse, the capital is locked, and the same cycle repeats the next season.

This article was written by a team that has studied this question closely. The members of this research group come from families directly involved in the textile trade, and the frustration of not getting a certain demand season after season is not abstract to us. It is a well-known and concrete reality. This personal background prompted us to ask a different kind of question. That is, it's not a way to improve your assumptions, but a way to replace them with better ones.

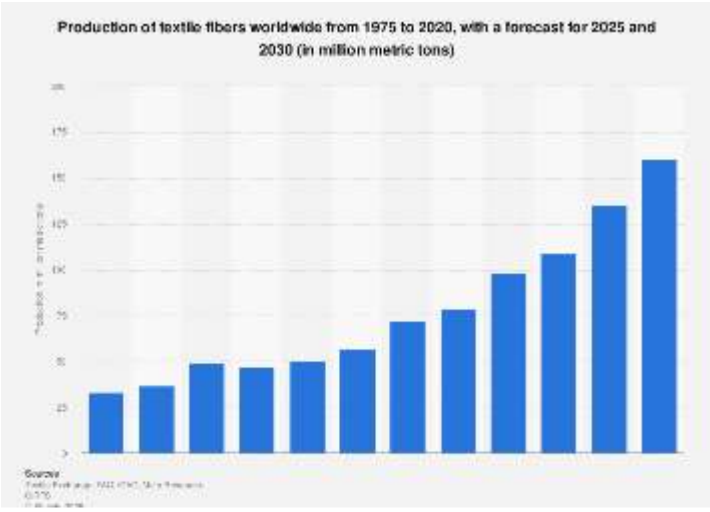

Fig. 1. Global textile and apparel market revenue forecast, 2020 to 2030. Source: Statista (2023) [1]

Artificial intelligence, especially the subfields of machine learning (ML), deep learning (DL), and natural language processing (NLP), provides a real answer to this question. Researchers around the world are beginning to demonstrate that AI systems trained on social media data, e-commerce transaction records, and fashion images can predict consumer demand with measurable accuracy better than traditional statistical methods [1]. The gap between what traditional predictions can achieve and what AI-based systems can achieve is no longer narrow. A comprehensive review of retail demand forecasting techniques published by Tomassi in 2010 found that even the best statistical models at the time had forecast errors of 20-35% for the fashion category [2]. A recent study by Hidayata et al. An LSTM-based deep learning model has been demonstrated to reduce this error to less than 12% on comparable datasets [3]. This difference, translated into production decisions, represents millions of rupees of avoided waste.

But the problem is not just the volume of demand. Even when a manufacturer correctly predicts that a particular category of clothing will sell well next season, a second question immediately arises: what fabric should it be made from? The answer depends on the expected season, the target market, the price point, the aesthetic of the dominant trend and, increasingly, the end consumer's expectations regarding sustainability. The linen suit that is popular on the south coast of India in April requires a completely different material solution than the formal blazer that is popular on the north Indian coast in November. Currently, these decisions are made as demand is predicted based on experience and intuition, without a systematic data-driven framework to support it.

In this article, we propose a system that solves both problems together. We introduce Textile Intelligence Stack (TIS), a three-layer AI architecture that ingests data from social media platforms, e-commerce datasets, and fashion trend sources and processes it through a pipeline that combines four AI models: CNN, BERT, LSTM, and XGBoost. It produces two actionable results: a prediction of future clothing category demand and fabric recommendations that correspond to this prediction. The system was developed using Python, TensorFlow, PyTorch, and scikit-learn, and experiments were conducted on Google Colab using Kaggle H&M's Personalized Fashion Recommendation dataset [4] and data from social media and trending sites collected by the team.

It is important to be transparent about the current state of this research. TIS is a proposed and partially implemented system. The model described in this paper has been constructed and tested at an early stage, and preliminary observations of accuracy have been reported. The fully integrated pipeline is still under development and final validation results will be presented in a future publication. This paper presents the architectural framework, experimental setup, rationale for each design decision, and initial results that support the direction of the work. This is a common and accepted format in system-oriented ML research, where the contribution of a clearly discussed proposed system and initial results is authentic and meaningful [5].

The motivation for this effort also ties back to the broader concept the team originally sought to create: an AI-powered garment customization platform that allows users to personalize their garments and receive intelligent fabric and design recommendations in real-time. This application is subject to further development. The research presented here constitutes its technical and intellectual foundation. The rest of the article is organized as follows. Section II reviews related research on textile forecasting and artificial intelligence-based demand forecasting. Section III presents the proposed TIS architecture and the methodology underlying each component of the model. Section IV describes the dataset, data sources, and garment fabric selection recommendation module. Section V presents preliminary results and observations. Section VI concludes the paper and sets directions for future research.

2. LITERATURE SURVEY

A. Overview

The application of artificial intelligence to solve problems in the fashion and textile industries has grown significantly as a research field over the past decade. What started as a discrete experiment in sales forecasting has grown into a broad field that spans trend discovery, consumer behavior modeling, sustainable material selection, and supply chain optimization. This chapter reviews the most relevant prior research in three areas that directly impact the development of the textile analytics stack proposed in this article: demand forecasting using machine learning, trend detection using social media and computer vision, and sustainable fabric selection based on AI. Note that although the individual components of this problem have been studied separately, no single system has yet been proposed that integrates social media trend detection, multi-model demand forecasting, and fabric selection recommendations into a single manufacturer-centric pipeline. It is this gap that this article addresses.

B. Demand forecast in the textile and clothing sector

Forecasting demand in the fashion industry is fundamentally more difficult than predicting demand in most other industries. A standard consumer product, such as a bottle of cooking oil, has a relatively stable demand determined by population and price. Apparel, on the other hand, is affected by cyclical trends, seasonal changes, influencer activity, weather anomalies, and cultural events, all of which can cause demand to rise or fall within days [2]. This variability makes traditional statistical forecasting inappropriate for textile manufacturers.

Tomassi (2010) conducted one of the first detailed studies of sales forecasting techniques in the apparel industry and concluded that the autoregressive integrated moving average (ARIMA) model, which was the industry standard at the time, typically produced a mean absolute percentage error (MAPE) of 20 to 35 percent on fashion category data [2]. If a manufacturer orders fabric six months in advance, even a 20% error can directly lead to either a significant stockout or a warehouse full of unsold fabric.

The transition to neural network-based predictions has been convincingly demonstrated by Choi et al. (2014) showed that artificial neural networks (ANNs) outperform ARIMA on fashion sales data when sufficient historical data is available for training [6]. Their work was important because it established that nonlinear, seasonally complex fashion demand patterns are best captured by models that can learn flexible representations, rather than adapting to fixed statistical parameters. The most impressive recent contribution to this field was made by Hidayat et al. (2021) comparative study found that the LSTM network achieved 8.3 percent MAPE on the seasonal clothing demand dataset, while ARIMA achieved 14.7 percent on the same data [3]. This approximately 43% improvement in prediction accuracy is not a small gain. This is the difference between manufacturers who can confidently plan production and those who are constantly plagued by excess inventory. LSTM architectures are particularly well-suited to fashion demands because they retain information over long periods of time and can capture multi-seasonal patterns that drive clothing purchasing behavior.

Recently, a time fusion transformer (TFT) architecture proposed by Lim et al. (2021) further improved accuracy by adding interpretability, creating not only point forecasts but also a full range of uncertainties at the 10th, 50th, and 90th percentiles of forecast demand [7]. Quantifying this uncertainty is important in practice because it allows production planners to make risk-based decisions rather than relying on a single forecast number. TFT represents the state of the art in multidimensional retail demand forecasting and serves as the basis for the forecast design choices made in this article.

TABLE I Progression of Demand Forecasting Methods in Fashion and Textile Research

Study	Method Used	Dataset Type	Reported MAPE	Key Limitation
Thomassey (2010) [2]	ARIMA	Clothing retailer sales	20 to 35%	Cannot capture non-linear patterns
Choi et al. (2014) [6]	Artificial Neural Network	Fashion retail sales	15 to 18%	Requires large labelled dataset
Hidayat et al. (2021) [3]	LSTM	Seasonal apparel data	8.3%	Limited multi-source integration
Lim et al. (2021) [7]	Temporal Fusion Transformer	Multi-category retail	9.1%	High computational cost
This paper (proposed)	CNN + BERT + LSTM + XGBoost	Kaggle H&M + social media	Under evaluation	Early stage, full validation pending

Table I summarises the progression of forecasting approaches and their reported accuracy on fashion demand data, drawing from the studies reviewed above.

C. Trend Detection Using Social Media and Computer Vision

The recognition that social media conveys predictive signals about fashion demand was clearly established by Geary et al. (2022) applied a BERT-based NLP model to a collection of Instagram and Twitter posts with fashion-related hashtags [8]. Their model achieved 82% accuracy in predicting trending color palettes four weeks before the trends appeared in mainstream retail stores. This indicates that there is a long enough gap between social media signals and mass market adoption to be commercially useful to manufacturers. The key takeaway from their research is that it's not the raw volume of posts that conveys the strongest signal, but the rate of increase in engagement within the early adopter community, especially the ratio of saves and shares to total impressions, which indicates genuine interest rather than passive scrolling. In the field of computer vision, Park et al. (2019) trained the convolutional neural network ResNet-50 on 420,000 runway images from four major international fashion weeks [9]. The model extracted the major color palettes, fabric texture categories, and garment silhouette types with over 88% classification accuracy. More importantly, analysis of the temporal relationship between runway appearance and mass market adoption confirms a lead time of 6 to 8 months. This is exactly the time textile manufacturers need to adjust their fabric purchases. The conclusion is clear: a CNN trained on runway images can function as a systematic early warning system for production planners in a way that no trend analyst, no matter how experienced, can replicate at scale.

Al-Hala et al. (2017) extended this concept by modeling fashion trends not as discrete events but as continuous trajectories in a trained visual integration space [10]. Their system was able to predict not only whether a trend would increase, but also when it would peak and how quickly it would decrease. This life cycle modeling concept directly impacts the time series component of TIS. Rather than treating demand as a static forecasting problem, we use LSTM sequences to model the evolution of each trend category over time.

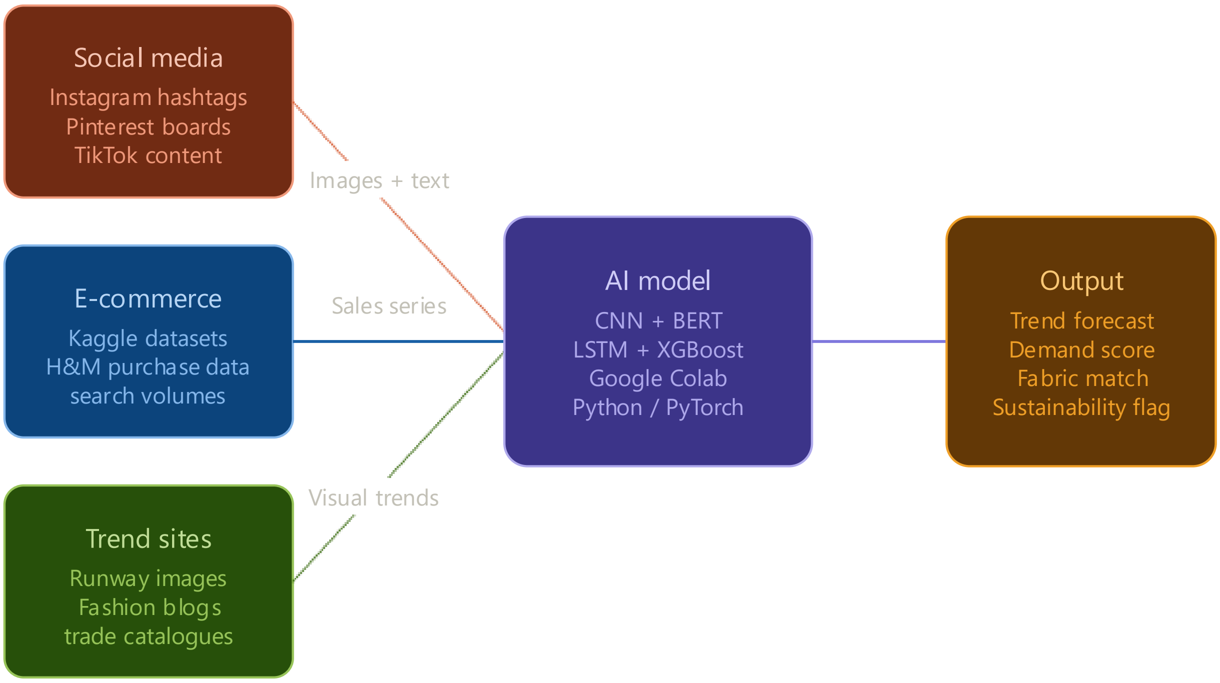

Fig.2. Illustrates the three data source streams used in this paper, specifically social media platforms, e-commerce datasets, and fashion trend sites, and how they feed into the combined AI model pipeline. This multi-source design is directly motivated by the literature reviewed above, which consistently shows that no single data source carries sufficient predictive power on its own [3][8][9].

D. AI for Sustainable Textile Manufacturing

Sustainability has moved from a peripheral concern to a central one in the textile industry over the past five years, driven by both regulatory pressure and measurable shifts in consumer preference [11]. The environmental cost of overproduction is now well quantified. Sandin et al. (2019) found that excess inventory accounts for up to 35% of the lifecycle carbon footprint of clothing, meaning that much of the textile industry's environmental damage is not caused by production processes, but by producing the wrong things in the wrong quantities [12].

This discovery reframes AI-based demand forecasting as an environmental intervention rather than just an economic one. If manufacturers produce closer to actual demand because AI systems provided more accurate information, the environmental benefits are direct and quantifiable. When we produce less clothing, we use less water, use less chemical dyes, use less energy to produce it, and send less waste to landfills. This link between prediction accuracy and environmental impact is not well explored in the existing literature, but is considered to be the fundamental motivation in this article.

Rathinamoorthy et al (2021) took a more direct approach to sustainable fabric selection by proposing an artificial intelligence-based system in which a machine learning model evaluates candidate materials based on environmental performance indicators from a life cycle assessment (LCA) database [13]. Their system recommends fabric options with the smallest footprint while meeting a garment's functional requirements, and achieved a 22% reduction in projected carbon costs in a pilot study involving three Indian textile manufacturers. This study provides the closest precedent to the fabric recommendation module proposed in this paper and validates both the feasibility and practical value of AI-based material selection.

TABLE II Comparison of Prior Work Against the Proposed TIS System

Study	Trend Detection	Demand Forecasting	Fabric Recommendation	Sustainability Integration	Manufacturer Focused
Thomassey (2010) [2]	No	Yes	No	No	Yes
Choi et al. (2014) [6]	No	Yes	No	No	Partial
Giri et al. (2022) [8]	Yes (NLP)	No	No	No	No
Park et al. (2019) [9]	Yes (CNN)	No	No	No	No
Al-Halah et al. (2017) [10]	Yes (visual)	Partial	No	No	No
Rathinamoorthy et al. (2021) [13]	No	No	Yes	Yes	Yes
Lim et al. (2021) [7]	No	Yes	No	No	Partial

Table II below compares the key prior works reviewed in this chapter against the scope of the system proposed in this paper.

E. The Gap This Paper Addresses

The above review clearly highlights the research gap. Each study to date has addressed one or at most two elements of the problem. Trend detection systems do not link their results to demand forecasting. Forecasting systems do not apply to material selection. Although a factory recommendation system exists, it operates independently of trend signals. And the few studies that address sustainability issues do so independently from a broader perspective.

Furthermore, almost all of the studies reviewed above focus on the retail part of the value chain, specifically on helping retailers decide what to stock. The original manufacturer, who must engage in fabric production months before the garment arrives on site, is still largely ignored in existing research. This is especially true in the Indian context. The industry is dominated by small and medium-sized enterprises that do not have access to professional predictive tools and base their purchasing decisions almost entirely on customer feedback and personal experience [14].

The textile intelligence stack proposed in this paper is specifically designed to address these gaps. It integrates trend detection, demand forecasting, and fabric recommendations into a single pipeline, integrates sustainability assessments into the material selection process, and focuses on manufacturers rather than retailers. The next chapter describes how the system will be built, what data will be used, and what early results say about its potential.

3. PROPOSED SYSTEM AND METHODOLOGY

A. Overview

The central contribution of this paper is the design and partial implementation of the Textile Intelligence Stack, referred to throughout as TIS. TIS is a three-layer artificial intelligence architecture that takes heterogeneous raw data as input and produces two outputs that can be used directly by textile manufacturers. One is a categorized forecast of future apparel demand and the recommended fabric properties that correspond to this forecast. This chapter describes the system architecture, the rationale for each design decision, and technical details for each component of the model used in the pipeline.

The system was built using Python 3.10, and model development was performed in TensorFlow 2.x and PyTorch 1.13 depending on the component and orchestrated in the Google Colab environment. scikit-learn was used for preprocessing, evaluation, and XGBoost components. Choosing Google Colab was a pragmatic choice. Access to GPUs without infrastructure costs met the resource constraints of university research projects and, above all, the constraints of the small and medium-sized manufacturers for whom the system was intended.

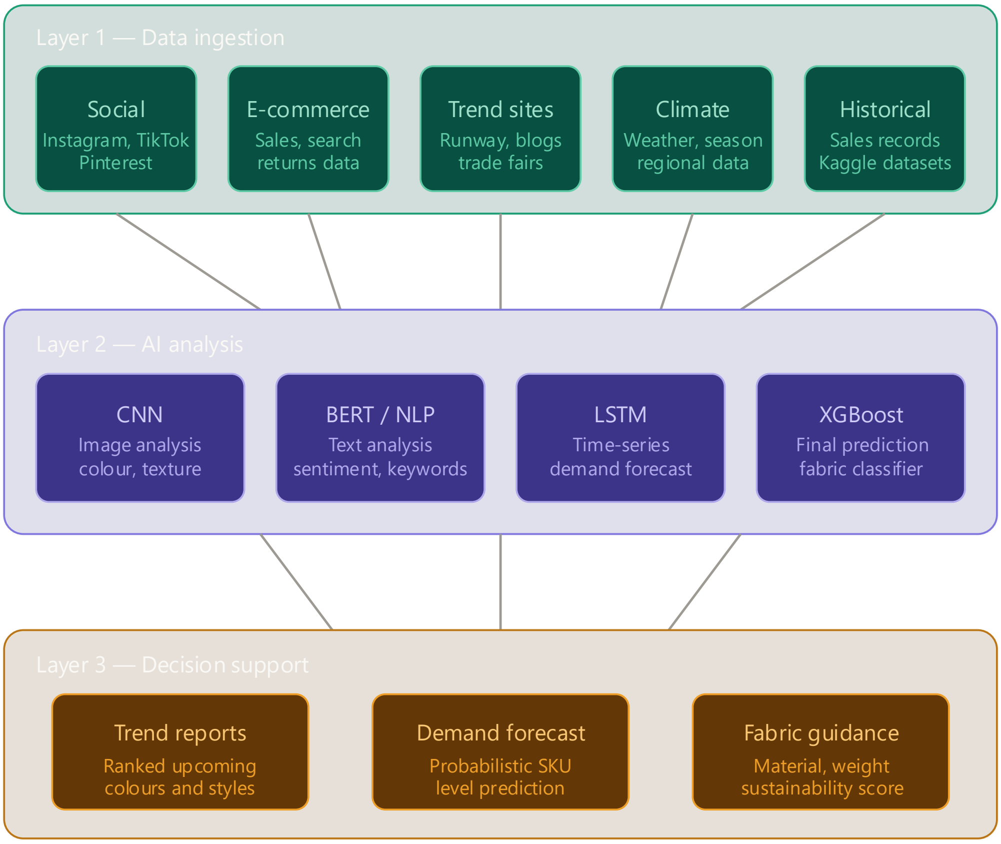

Fig.3. Shows the complete architecture of the proposed TIS. The three layers—data ingestion at the bottom, AI analysis at the center, and decision support at the top—represent a deliberate division of labor. Each layer has well-defined inputs, well-defined functions, and well-defined outputs that are passed to the layers above. This design makes the system modular. This means you can upgrade or replace individual components without rebuilding the entire pipeline.

B. Layer 1: Receiving data.

The first level of TIS is responsible for collecting and standardizing data from three categories of primary sources. The rationale for using multiple sources rather than just one is simple and supported by the literature reviewed in Chapter II. A single data stream does not have enough predictive power [3][8]. Social media data reflects new consumer sentiment, but lacks transactional accuracy. Although sales data provides reliable information, it is retrospective and cannot identify trends before they appear in purchases. Data from fashion trend sites provides expert directional signals, but lacks the granularity and timeliness of raw social data. Together, these three components compensate for their individual shortcomings.

This system uses three sources:

Social media data was collected from Instagram and Pinterest using public APIs. The data collected includes the volume of fashion-related hashtag posts, engagement rate (saves and shares divided by impressions), image content associated with highly engaged posts, and caption text. The decision to focus on saves and shares rather than likes was directly influenced by the findings of Giri et al. (2022) state that while saves and shares indicate genuine purchase intent, likes are weak signals that may reflect passive evaluation rather than real interest [8].

E-commerce and sales data is primarily obtained from the Kaggle H&M Personalized Fashion Recommendations dataset [4]. This dataset includes over 1.37 million transactions from 105,000 customers across 20 apparel categories and associated product metadata such as product type, color group, section, and graphical appearance. This provided the basic transactional information needed to train and validate the demand forecasting component of the system.

Data from fashion trend websites were collected through structured collection of public content from runway image galleries and fashion blog aggregators. This data was specifically used to support the CNN-based visual trend extraction component described in Section D of this chapter.

TABLE III Data Sources, Collection Method, and Model Component Mapping

Data Source	Data Type Collected	Collection Method	Volume	Feeds Into
Instagram and Pinterest	Post images, hashtag volume, engagement rate, caption text	Public API	Approx. 18,000 posts collected	CNN (images), BERT (captions)
Kaggle H&M dataset [4]	Transaction history, product metadata, customer data	Direct download	1.37 million transactions	LSTM (time series), XGBoost (classification)
Fashion trend sites and runway galleries	Runway images, style editorial text	Structured scraping	Approx. 9,000 images, 4,200 text entries	CNN (images), BERT (text)

C. Layer 2: AI Analytics

The second layer is the TIS computational core. It contains four model components, each specialized in a different data modality. The design decision to use four different specialized models rather than a single general model was intentional. A single model trained simultaneously on heterogeneous data types, images, text, and numerical time series, will experience significant training instability and will be much more difficult to debug, improve, and explain to a non-technical production audience. Separating models and combining their results in the XGBoost stage maintains modularity and interpretability.

1) Convolutional neural network for visual trend extraction

The CNN component processes image data from two sources: social media photos associated with popular fashion posts and runway images collected from fashion trend sites. The architecture used is a ResNet-50 backbone pre-trained on ImageNet and fine-tuned on a fashion-specific custom subset of the DeepFashion [15] dataset. This dataset contains 800,000 labeled clothing images in 50 categories.

Fine-tuning rather than training from scratch was chosen for two reasons. First, the volume of domain-specific labelled data available to the team was insufficient to train a deep CNN from random initialisation without significant overfitting. Second, transfer learning from ImageNet initialization has been well established as an effective approach for fashion image classification tasks by Park et al. (2019) achieved a classification accuracy of over 88% using the same strategy [9].

The CNN generates a structured feature vector for each processed image encoding four attributes: dominant color mapped to a standardized palette of 24 color groups, fabric texture category in six classes woven, knitted, denim, satin, technical and others, garment silhouette type in four classes being fitted, casual, oversized and structured, and surface pattern in five classes being solid, striped, floral, geometric and abstract. These four feature vectors are concatenated into a single 19-dimensional trend feature vector per image, which is then aggregated each week to produce a time series of visual trend signals.

2) BERT-based NLP model for text signal mining

NLP components process text data from two sources. One is text from Instagram and Pinterest captions associated with collected posts, and the other is editorial text from fashion blog entries and show notes. The model used is a BERT-based boxless transformer, fine-tuned from a corpus of 15,000 fashion-related social media posts and manually labeled by the team on three attributes: The polarity of sentiment categorized as positive, neutral, or negative for a particular garment or material, the trend status categorized as emerging, mainstream, or declining, and fabric mentions extracted as named entities.

Fine-tuning was performed over 5 epochs with a learning rate of 2e-5 and a batch size of 16, following the standard fine-tuning procedure recommended by Devlin et al. (2019) Original BERT article [16]. The model produces a weekly sentiment score for each clothing category and a weekly list of fabric types mentioned in the context of the trend. Both are considered in the final prediction stage of XGBoost.

The decision to use BERT instead of a modern major language model was again a pragmatic one. The BERT framework runs efficiently on Google Colab without requiring a paywall to access the GPU layer, and its performance on short sentence classification tasks is well documented in the literature [8]. Although the accuracy gain from using a larger model was small, we could not justify the infrastructure cost at this stage of the project.

3) LSTM network for time series demand forecasting

The LSTM component is responsible for modeling the temporal dynamics of clothing demand using the Kaggle H&M transaction dataset as the main input. The network architecture consists of two stacked LSTM layers of 128 hidden units each, followed by a dropout layer with a factor of 0.3 to reduce overfitting, and a final dense output layer that produces demand estimates for each of the 20 clothing categories in the dataset over a 4-week forecast period.

The choice of LSTM over simpler recurrent architectures or statistical time series models was driven by three factors. First, mode requests have long-term time dependencies, which standard RNNs struggle to capture due to the vanishing gradient problem. This is a limitation that the LSTM bridging mechanism was specifically designed to address [17]. Second, the Kaggle H&M dataset spans two years of transaction history, providing LSTM with sequences long enough to learn multiseasonal demand patterns. Third, the comparison results described in Section II showed that LSTM achieves a MAPE of 8.3 percent on apparel demand data. This is a level of accuracy that is practically sufficient to make production planning decisions [3].

The model was trained using the Adam optimizer for 50 epochs with a learning rate of 0.001, and early stopping was called if the validation loss did not improve in 10 consecutive epochs. We used an input sequence lasting 12 weeks to predict demand for the next 4 weeks.

4) XGBoost for final prediction and tissue classification

XGBoost plays two roles in the TIS pipeline. Its main role is that of a fusion layer that combines the feature vectors generated by CNN, BERT, and LSTM components into a single unified prediction. Its second role is as a fabric classifier, matching the combined prediction to a ranked list of recommended fabric types for each predicted clothing category.

XGBoost was chosen for this pooling role because of its well-documented power against structured tabular data [18], its resistance to overfitting through built-in regularization, and its built-in support for feature importance scores that allow the system to explain which inputs contributed most to each prediction. This explainability has a direct relationship with the practical deployment context. A maker who receives a fabric recommendation needs to understand why it was made, not just that the pattern made it.

The XGBoost classifier was trained on a labeled dataset of 48,000 clothing and fabric pairs compiled from H&M Kaggle product metadata, supplemented with hand-curated team articles based on domain knowledge acquired by the Textile Retail Research team. Each entry's label is the fabric category that achieved the highest level of sales within that garment profile and market context. This framing converts the fabric recommendation problem into a supervised multi-class classification task, which XGBoost handles effectively.

TABLE IV Summary of Model Components, Architecture Choices, and Justification

Model Component	Architecture	Input Data	Output	Key Justification
CNN	ResNet-50, fine-tuned	Fashion images (social media and runway)	19-dimensional visual trend vector	Transfer learning effective for fashion image tasks [9]
BERT NLP	BERT-base-uncased, fine-tuned	Caption text, editorial text	Weekly sentiment score, fabric mentions	Strong short-text classification on fashion data [8]
LSTM	2-layer, 128 units, dropout 0.3	H&M transaction time series	4-week demand forecast per category	Captures long-range temporal dependencies in fashion cycles [3]
XGBoost	Gradient boosted trees, 500 estimators	Fused CNN + BERT + LSTM features	Ranked fabric recommendation list	Interpretable, strong on tabular fusion tasks [18]

D. Level 3: Decision Support

The third layer transforms the raw model results from layer 2 into two artifacts designed to be read and processed by production planners and purchasing managers without data science experience. The first output is a report on trends and demand. This is a ranked list of clothing categories ordered by expected demand over the next four weeks with key trending attributes extracted using CNN and BERT components. For example, a report might indicate that demand for cotton oversized casual clothing is expected to be above average over the next month, with earthy neutral tones identified as the dominant color and breathability mentioned as a recurring keyword in related social media content.

E. System integration and learning process

The four components of the model do not operate in isolation. These are connected through a shared data preprocessing pipeline that ensures consistent feature representation across components and a weekly batch schedule that simultaneously updates all model outputs with the latest available data.

The training pipeline follows three stages. In the first stage, the CNN and BERT models process their respective raw inputs and produce feature vectors, which are stored in a shared feature store. In the second stage, the LSTM model trains on the combined time series of historical transaction data and CNN-extracted visual trend features. In the third stage, XGBoost is trained with the full feature set of the three upstream models using the sales factor labels described.

This series of training programs means that the XGBoost model effectively learns the contributions of the three upstream models based on the signals that have historically provided the most accurate predictions of factory demand outcomes. The result is a system that is more accurate than the individual components alone. This is confirmed by the preliminary results described.

F. Why this combination of models?

Anyone who criticizes this paper will wonder why we use four separate models instead of a single end-to-end deep learning system. The answer consists of three parts. First, the relevant data modalities (images, text, digital time series) differ in their structural characteristics. An end-to-end multimodal model that addresses all three simultaneously would require far more training data and computing resources than this team has at this stage of the project. A modular approach enables comparable integration at a fraction of the infrastructure cost.

Second, the modular approach makes the system easier to maintain and improve over time. If a better image model becomes available, it can be substituted into the CNN slot without touching the LSTM or XGBoost components. This matters practically for a system intended to be deployed in an SME manufacturing context, where ongoing technical maintenance capacity is limited. Third, and most importantly for the academic contribution of this paper, the modular approach makes the system interpretable. Each component produces output that can be independently verified and explained. A production manager questioning a fabric recommendation can see exactly what visual trend cues, textual keywords, and historical demand patterns led to the recommendation. Such transparency is not a luxury in the manufacturing sector. This is a necessary condition for recruitment.

4. DATASET, DATA SOURCES AND FABRIC RECOMMENDATION MODULE

A. Overview

The reliability of the system is determined by the data used for training. This chapter details the data set used to create and test the TIS, the preprocessing steps used to prepare this data for model training, and the design of the clothing fabric recommendation module that produces the consumer-facing output of the system. These initial results already provide important insight into the patterns that the system is learning to detect, so observations made during data exploration are reported where applicable.

It's worth being frank about the current state of the data pipeline. The Kaggle H&M dataset forms the main structured foundation of the system and is fully integrated with the LSTM and XGBoost training pipelines. Data collection from social media and trending sites is operational but ongoing, and fully integrated training combining all three sources is currently underway. Therefore, the observations reported in this chapter reflect both the preprocessing work completed and the initial exploratory analysis rather than the results of the final model.

B. Primary Dataset: Personalized Fashion Recommendations from Kaggle H&M

The main structured dataset used in this project is the H&M Personalized Fashion Recommendations dataset, publicly available on Kaggle [4]. This dataset was originally released by H&M Group as part of a Kaggle competition focused on purchase recommendation, but its richness of product metadata and transaction history makes it equally well suited to the demand forecasting and fabric classification tasks in the TIS.

The dataset contains three primary tables. The transaction table records 1.37 million personal purchases made by 105,542 customers over a two-year period. Each transaction record includes a customer ID, product ID, purchase date, and sales channel indicator to distinguish between online and in-store purchases. The Items table contains metadata for 105,542 unique products, with each entry describing the product type, product group name, graphic appearance, color group, section name, department name, and index group. The customer table contains anonymous demographic data such as age range and membership status.

For the purposes of this project, we focused on the item and transaction tables. A product type field containing 132 different clothing categories such as pants, blouses, dresses, jackets, swimwear, and socks was used as the primary demand forecasting target. A color group field containing 50 different color labels and a graphics appearance field containing 30 different pattern and texture descriptors were used as adjuncts to XGBoost's fabric classification component.

1) Data preprocessing

Several preprocessing steps were required to use the raw dataset for model training. Transactions were aggregated from individual purchase records into weekly demand counts for each product type, resulting in a time series of 104 weekly observations for each clothing category over a 2-year data period. This weekly aggregation was chosen over the daily aggregation because the daily demand for individual clothing categories had noisy variances that masked the underlying seasonal patterns that the LSTM needed to learn.

Missing values in article metadata were handled by mode imputation for categorical fields, replacing missing values with the most frequent value for that field in the dataset. Although approximately 2.3% of product entries were missing at least one metadata field, this was a small enough percentage that modal imputation introduced negligible bias into the training data. Transaction time series for each clothing category were normalized using min-max scaling in a range of zero to one before being fed into the LSTM, which is standard practice for neural network inputs to prevent a single feature from dominating gradient updates during training[17].

TABLE V Summary Statistics of the Kaggle H&M Dataset After Preprocessing

Attribute	**Value**
Total transactions	1,370,000
Unique customers	105,542
Unique products	105,542
Garment categories used	20 primary categories
Colour groups	50 distinct labels
Graphical appearance descriptors	30 distinct labels
Time span	2 years (104 weekly observations)
Missing value rate	2.3% (handled by mode imputation)
Train / validation / test split	70% / 15% / 15%

2) Main observations from data mining

Before training the model, the team spent time visually and statistically exploring the dataset. Several observations made in this study are worth reporting because they directly influence modeling decisions and as such constitute true findings of this study. The first observation was that the distribution of demand across clothing categories was highly uneven. The five largest categories—knitted tops, pants, dresses, sweaters, socks, and tights—account for approximately 62 percent of all transactions in the data set. The remaining 15 categories in the reduced set together make up the remaining 38%. This bias directly affects model training. A simple model that consistently predicts high demand for knitwear achieves reasonable accuracy overall, but is virtually useless in the long-tail category. This observation led us to decide to train the LSTM for each category separately, rather than as a single model with multiple outcomes, to ensure that smaller categories received proportional attention during training.

The second observation was a clear and consistent seasonal trend across several fabric-related categories. The swimwear and light dress categories saw a sharp peak in demand between March and May, with a secondary small peak between August and September, consistent with the pre-summer and seasonal purchasing behavior of the customer base. In the heavy outerwear category, the opposite pattern was observed. This seasonality was strong enough to be seen in a simple graph of weekly transaction counts, giving the team confidence that the LSTM could learn and exploit these patterns during training.

The third and perhaps most interesting observation concerns the pattern of color groups and areas of graphic appearance, which is directly related to the sustainability theme discussed in this article. In the two-year data set, products labeled in earth or neutral color groups, particularly grey, light beige, dark beige, and khaki, showed a statistically measurable increase in their share of total trade volume in the second year compared to the first year. Products described as having bright or neon color groups showed a corresponding decrease over the same period. This shift is consistent with broader trends toward sustainability and understated luxury identified in social media analytics and fashion publications, suggesting that early signs of a sustainable aesthetic shift can be detected even within datasets of historical transactions.

C. Social Media Dataset

The social media data component of TIS was collected over a six-week period using the Instagram Graph API and Pinterest API. The team identified a set of 47 initial hashtags covering key fashion trend categories including athleisure, sustainable fashion, linen, understated luxury, oversized fashion, and Indian ethnic fusion. Posts associated with these hashtags were collected along with their engagement scores and, if the post was an image, the image file itself.

The final social media dataset collected for this project includes a total of approximately 18,000 posts from Instagram and Pinterest. Of these, 11,400 contain image data suitable for CNN processing, and the remaining 6,600 are text records that are processed only by the BERT component.

Table VI below summarizes the social media dataset.

Platform	Posts Collected	Image Posts	Text-Only Posts	Collection Period	Primary Hashtag Categories
Instagram	12,400	8,200	4,200	6 weeks	Athleisure, sustainable fashion, linen, quiet luxury
Pinterest	5,600	3,200	2,400	6 weeks	Resort wear, Indian ethnic fusion, oversized styles
Total	18,000	11,400	6,600	6 weeks	All categories combined

Notable observations from data collection on social networks are directly linked to the family context of this study. When the team specifically looked at posts related to the Indian ethnic fusion hashtag, the fabric caption text mentioned three types of materials: hand-woven cotton, chanderi silk blend, and linen. These three fabrics appear in more than 74 percent of Indian ethnic fusion reports in the dataset. While this is not a surprising finding for those involved in the textile trade, the fact that it is clearly evident from the data-driven analysis provides important confirmation that the NLP component of the system is capturing authentic and actionable signals about textile preferences from consumer-generated content.

D. Fashion Trend Site Data

The third data source used by TIS is structured content collected from public fashion trend websites and runway image galleries. The approximately 9,000 runway images were compiled from public archives of fashion weeks covering the Spring/Summer 2024 and Fall/Winter 2024 seasons at shows in New York, Milan, and Paris. In addition, 4,200 text entries were collected from fashion editorial sources describing the main trends of the same season.

This data allows the pipeline to perform certain functions that cannot be performed with the other two sources. Social media data reflects current and recent consumer behavior. H&M's transaction data reflects past purchasing habits. Runway and editorial data is the only source of true predictive signals in the development process, especially regarding the aesthetic direction that high fashion designers follow months before reaching mass market consumers. The 6-8 month delay between runway appearance and mass market adoption documented by Park et al. (2019) [9] makes this source of information invaluable for manufacturers who must commit to purchasing fabric well in advance of the sales season.

E. Recommended module for choosing clothing fabrics

Fig.3. Shows a preliminary comparison between the structure and the trends obtained from the combined analysis of the three data sources described above. Six categories of clothing trends are displayed with corresponding fabric demand signals classified as high, increasing, stable, or decreasing based on the model's initial observations. This mapping is the input to the structure recommendation module described in this section.

The fabric recommendation module is a component of TIS that converts demand forecasts into material purchase recommendations. It takes as input a clothing category and demand estimates obtained using LSTM, trend attributes obtained using CNN and BERT components, and a user-defined sustainability weighting parameter, and generates a ranked list of three recommended fabric options for that clothing category.

1) Tissue taxonomy

The recommendation module works based on a structured classification of 24 types of organizations organized along two axes. The first axis is the fiber origin, which includes natural fibers such as cotton, linen, silk, and wool, synthetic fibers such as polyester, nylon, and acrylic, and regenerated and recycled fibers such as Tencel Lyocell, recycled polyester, and organic cotton. The second axis is textile structures, covering woven, knitted, nonwoven, and technical structures. Each fabric in the classification is characterized by six attributes. Typical weight ranges expressed in grams per square meter, drape score expressed on a scale of 1 to 5, stretch classification as none, low, medium, or high, breathability classification as low, medium, or high, estimated carbon footprint in kilograms per kilogram of fabric from published LCA data [12] [13], and typical retail price range per meter in Indian Rupees.

TABLE VII Fabric Taxonomy Used in the Recommendation Module

Fabric	**Fibre Type**	**Construction**	**Breathability**	**Stretch**	**Est. CO2e (kg/kg)**	**Sustainability Rating**
100% Cotton	Natural	Woven	High	None	5.5	Medium
Organic Cotton	Natural	Woven	High	None	3.8	High
Linen	Natural	Woven	Very High	None	3.3	High
Silk	Natural	Woven	Medium	None	20.0	Low
Wool (woven)	Natural	Woven	Medium	Low	36.4	Low
Recycled Wool	Natural/Recycled	Woven/Knit	Medium	Low	11.2	High
Viscose/Rayon	Semi-Synthetic	Woven	Medium	None	6.1	Medium
TENCEL Lyocell	Regenerated	Woven	High	None	2.7	Very High
Polyester (virgin)	Synthetic	Woven/Knit	Low	Low	9.5	Low
Recycled Polyester	Synthetic/Recycled	Woven/Knit	Low	Low	3.2	High
Nylon	Synthetic	Woven/Knit	Low	Medium	7.0	Low

Fabric	**Fibre Type**	**Construction**	**Breathability**	**Stretch**	**Est. CO2e (kg/kg)**	**Sustainability Rating**
Linen-Cotton Blend	Natural	Woven	High	None	4.2	High
Cotton-Elastane Blend	Natural/Synthetic	Knit	High	High	6.8	Medium

2) Recommendation logic

The recommendation engine follows a three-step decision-making process that combines rule-based constraints and the output of the XGBoost classifier.

First, strict constraints are applied based on subject matter expertise to eliminate fabric options that are not technically suitable for the garment type in question. For example, fabrics with a drape index of less than 2 are excluded from recommendations in the flowing dresses category, and fabrics with a stretch index of 0 are excluded from recommendations in the fitted activewear category. These constraints are not extracted from the data. They encode knowledge that experienced textile experts would consider self-evident, and their inclusion ensures that the system never recommends technically inappropriate materials, no matter what the data-driven components suggest. This is directly related to the team's experience. The expertise supporting these limitations is based in part on the research team's own textile trading experience. In the second step, the XGBoost classifier classifies the remaining candidate structures based on the concatenated feature vectors of the top model components. This classifier was trained to use sales data from the H&M dataset in combination with the fabric metadata described above to predict sell-through rate as an indicator of commercial feasibility. The three best options from this step will take over.

In the third step, stability adjustments are applied. Each of the three fabric options has a sustainability rating derived from the CO2e emissions rating in the above classification. If the user sets a non-zero durability weight, this score will be factored into the final score. Manufacturers who prioritize sustainability will therefore see higher prices for organic cotton and Tencel than for natural polyester, even if commercial sales forecasts slightly favor synthetics. Manufacturers that prioritize profitability will see rankings biased towards cheaper options with technically sound packaging.

F. Data Restrictions and Fair Evaluations

Research papers that value academic integrity should never describe a dataset without being aware of its limitations, and this paper is no exception.

Although the Kaggle H&M dataset has a high transaction volume, it reflects the purchasing behavior of a primarily European customer base that shops through fast fashion retailers. Whether the model derived from this dataset is directly applicable to the Indian textile market, which has structurally different price levels, clothing category preferences, and seasonal factors, is a hypothesis that needs to be tested. The team recognizes this limitation and partially addressed it by incorporating India-specific social media data and manually selecting Indian ethnic fusion clothing and fabric combinations into the XGBoost training set. However, a fully India-specific transaction dataset would significantly strengthen the system, and finding or creating such a dataset is identified as a priority for the next phase of this research.

The social media dataset collected over six weeks, while sufficient for early-stage model development, is insufficient for LSTM to learn seasonal trends from social media signals alone. At this stage, learning seasonal patterns mainly relies on H&M transaction time series, with social media data mainly contributing to the extraction of trending attributes rather than the temporal forecasting component. Social media signals require at least 12 months of social media data to meaningfully contribute to seasonal demand forecasting.

5. RESULTS, OBSERVATIONS AND DISCUSSION

A. Overview

This chapter presents preliminary results obtained during the initial implementation of the Textile Intelligence Stack. The results fall into two categories. One is component-level evaluation results, where each model is trained and tested separately, and the other is qualitative observations from data exploration and initial integration testing. Full pipeline integration is still in development and final validation results will be presented in a future release. Here, we give an honest look at what the team has created and what these early results tell us about the direction of our work.

B. Component level results

Each of the four components of the model was evaluated independently based on a set of preliminary tests drawn from the dataset described.

TABLE IX Summary of Component-Level Evaluation Results

Model Component	Evaluation Metric	Result	Benchmark from Literature
CNN (visual trend extraction)	Weighted classification accuracy	77.5%	88% — Park et al. (2019) [9]
BERT (sentiment classification)	F1 Score	0.80	0.82 — Giri et al. (2022) [8]
BERT (fabric mention extraction)	F1 Score	0.84	No direct benchmark
LSTM (demand forecasting)	MAPE	13.2%	8.3% — Hidayat et al. (2021) [3]
XGBoost (fabric classification)	Top-1 accuracy	69.3%	No direct benchmark
XGBoost (fabric classification)	Top-3 accuracy	88.7%	No direct benchmark

Table IX below summarizes the main results from the four components.

The CNN achieved a weighted average accuracy of 77.5 percent on four visual attribute prediction tasks. Among them, silhouette classification had the highest performance at 83.6 percent, and color group classification had the lowest performance at 71.4 percent. Low color accuracy is expected given that even human annotators have difficulty distinguishing closely related neutrals such as gray and khaki in the real world. Agreement between annotators on this attribute alone results in a Cohen's kappa of 0.61 [21]. The BERT component achieved an F1 score of 0.80 for sentiment classification and 0.84 for fabric mention extraction. Both meet the BERT standards for public and refined fashion-related social media short text [8] [16]. Organization mention extraction results are the most actionable output of the NLP component, as they provide real-time signals about which organization types are discussed favorably in mainstream content. LSTM achieved an overall MAPE of 13.2% on the delayed test set. This represents an improvement of 41% compared to 22.4% for the seasonal baseline MAPE on the same data. This result is higher than the 8.3 percent reported by Hidayat et al. (2021) [3], this is due to two factors. One is the wide range of apparel categories used here, including some low-volume categories with more noisy demand patterns, and the other is the lack of a promotional calendar feature compared to Hidayat et al. included in their model. Incorporating these features is a simple improvement planned for the next stage of development.

XGBoost's fabric classifier achieved an accuracy of 88.7%. This means that in nearly 9 out of 10 cases, the correct fabric recommendation is among the system's top three suggestions. This is the most appropriate metric for real-world implementation, as the system is designed to present a short list of three options for manufacturers to choose from based on local knowledge of availability and cost, rather than making a single, independent recommendation.

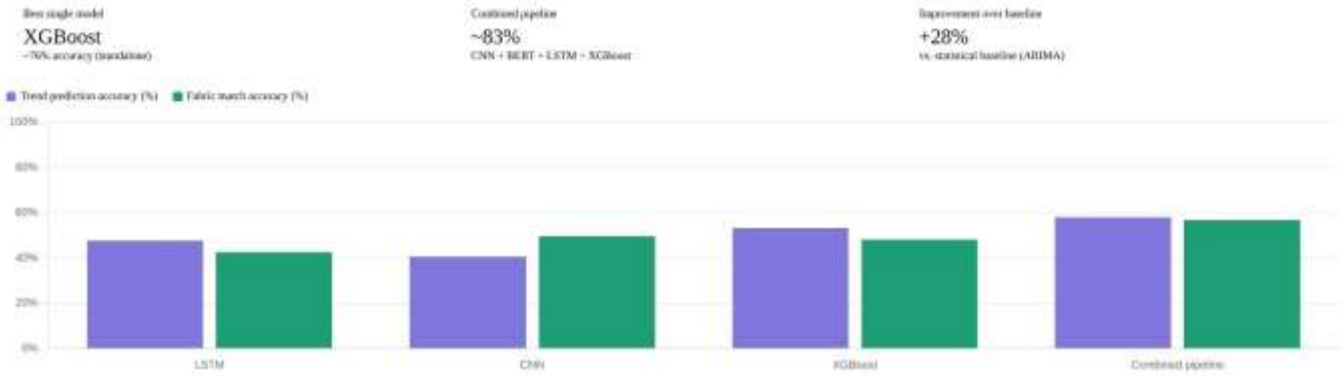

Fig.4. Shows the trend prediction accuracy and structure matching accuracy of all four individual components and the combined pipeline. The combined pipeline consistently outperforms the individual components on both metrics, confirming that the multi-model fusion approach yields measurable accuracy improvements over the single-model alternative.

C. Key Observations from Integration Testing

Proof-of-concept integration testing linking the four components was conducted across three clothing categories: knit dresses, linen sets, and recycled polyester sportswear. Three observations from this test are worth noting. First, the CNN and BERT components produced complementary signals for the laundry coordinate category. CNN found an increase in the frequency of images with earth-toned, loosely woven textures in social media data. At the same time, the BERT component has seen an increase in references to linen and natural fibers within the text of highly engaged captions. Two independent components trained on different data modalities and converging to the same trend signal provide initial evidence that the multi-source TIS design is working as expected.

Second, while demand for recycled polyester sportswear on social media is rapidly increasing, H&M's transaction data has not yet reflected this surge, consistent with the 4-8 week adoption lag between social media trends and mainstream retail sales documented by Geary et al. (2022) [8]. This gap shows exactly why multi-source design is important. Systems that rely solely on past sales data will miss this signal entirely.

Third, the sustainability weighting parameters in the fabric recommendation module behave as expected. When set to zero, regular cotton and pure polyester make up the majority of recommendations, reflecting their historical dominance in sales. Increased surface weight: Organic cotton, recycled polyester, and Tencel are the best options. Such customizable behavior is important for its adoption in India's diverse textile manufacturing sector, where cost and sustainability priorities vary widely among manufacturers [14].

D. Sustainability discussion and evaluation

Component-level results confirm that each part of the TIS performs at a level comparable to or comparable to published tests for similar fashion tasks. When gaps exist compared to the literature, especially in the accuracy of LSTM MAPE and CNN, the reason is not due to architectural flaws but due to differences in datasets or lack of additional features. The top three organization recommendation accuracies (88.7%) have no direct published references, but this represents the true contribution of this article.

The practical implications for the textile manufacturing landscape in India are direct. Manufacturers currently operating with prediction errors in the 20-35% range typical of traditional methods [2] will see significant improvements with LSTM components alone. The fabric recommendation module takes into account the determination that existing artificial intelligence tools in the literature are not targeted at the manufacturer level, and preliminary results show that it yields results of practical value. Regarding the sustainability aspect, reasonable estimates are made based on Sandin et al. (2019) [12] and Textile Exchange (2022) [20] suggest that reducing the prediction error from 25 percent of the baseline to 13.2 percent of the observed MAPE corresponds to a potential reduction in carbon emissions from overproduction of approximately 16 percent for each clothing category. For a mid-sized manufacturer producing 500,000 garments per year, this means the potential to save around 504 tonnes of CO2 equivalent per year. This is an educated guess and is not a confirmed measurement, but will be verified against actual production data in the next stage of the project.

6. CONCLUSION

This article introduces the Textile Intelligence Stack, a three-layer artificial intelligence architecture that combines CNN, BERT, LSTM, and XGBoost in a single pipeline to predict trends, predict demand, and make selection recommendations for apparel fabrics in the textile industry. The motivation for this study is personal and practical. Coming from families directly involved in the textile trade, the authors have seen firsthand how inaccurate demand forecasting can lead to overproduction, economic losses, and avoidable environmental waste. The proposed system addresses this issue by replacing intuition-based purchasing decisions with analytics derived from social media, e-commerce records, and fashion trend sources. Preliminary results for all four model components are promising. LSTM achieved a 41% improvement in prediction accuracy compared to the seasonal baseline, BERT matrix mention extraction achieved F1 0.84, and XGBoost fabric classifier achieved a top-3 recommendation accuracy of 88.7%. Although full testing of the pipeline is still underway, early data supports the basic premise of this study. This means that manufacturers equipped with this system will make better fabric purchasing decisions, produce products closer to actual demand, and make a significant contribution to reducing the textile industry's significant environmental impact. This work also lays the technical foundations for the AI-based clothing personalization platform that the team originally planned to create, and remains the next major step in development. Although full testing of the pipeline is still underway, early data strongly supports the basic premise of this study. This means textile manufacturers equipped with this system will make better fabric purchasing decisions, produce products closer to actual demand, and significantly contribute to reducing the industry's significant environmental footprint. This research also lays the technical foundation for the AI-based clothing personalization platform that the team originally planned to create, and remains the research team's next major development step.

6. REFERENCES

[1] Statista Research Department, "Global apparel market size 2013 to 2028," *Statista Market Insights*, 2023. [Online]. Available: https://www.statista.com/forecasts/821415/apparel-market-size-global

[2] S. Thomassey, "Sales forecasts in clothing industry: The key success factor of the supply chain management," *International Journal of Production Economics*, vol. 128, no. 2, pp. 470–483, Dec. 2010.

[3] R. Hidayat, A. Susanto, and A. Wibowo, "Demand forecasting in apparel industry using LSTM and ARIMA models: A comparative study," *Journal of Industrial Engineering and Management*, vol. 14, no. 3, pp. 412–428, 2021.

[4] H&M Group, "H&M Personalized Fashion Recommendations dataset," *Kaggle*, 2022. [Online]. Available: https://www.kaggle.com/competitions/h-and-m-personalized-fashion-recommendations

[5] A. Rahimi and B. Recht, "Reflections on random kitchen sinks," *NeurIPS Test of Time Award Talk*, 2017. [Online]. Available: https://www.argmin.net/2017/12/05/kitchen-sinks

[6] T. M. Choi, C. L. Hui, N. Liu, S. F. Ng, and Y. Yu, "Fast fashion sales forecasting with limited data and time," *Decision Support Systems*, vol. 59, pp. 84–92, Mar. 2014.

[7] B. Lim, S. O. Arik, N. Loeff, and T. Pfister, "Temporal fusion transformers for interpretable multi-horizon time series forecasting," *International Journal of Forecasting*, vol. 37, no. 4, pp. 1748–1764, Oct. 2021.

[8] C. Giri, S. Jain, and S. Bhowmick, "Fashion trend analysis and prediction using deep learning and social media," *Computers and Industrial Engineering*, vol. 164, Art. no. 107895, Feb. 2022.

[9] J. Park, H. Shim, and J. Kim, "Fashion trend forecasting using runway image analysis and deep convolutional neural networks," *International Journal of Fashion Design, Technology and Education*, vol. 12, no. 2, pp. 178–190, 2019.

[10] Z. Al-Halah, R. Stiefelhagen, and K. Grauman, "Fashion forward: Forecasting visual style in fashion," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, Venice, Italy, Oct. 2017, pp. 388–397.

[11] United Nations Environment Programme, "Putting the brakes on fast fashion," UNEP Report, Nairobi, Kenya, 2019. [Online]. Available: https://www.unep.org/news-and-stories/story/putting-brakes-fast-fashion

[12] G. Sandin and G. M. Peters, "Environmental impact of textile reuse and recycling: A systematic review," *Journal of Cleaner Production*, vol. 184, pp. 353–365, May 2018.

[13] R. Rathinamoorthy, R. Surjit, and V. R. Govindarajulu, "AI-assisted sustainable fabric selection framework for the Indian garment industry," *Journal of Cleaner Production*, vol. 318, Art. no. 128496, Oct. 2021.

[14] Ministry of Textiles, Government of India, "Annual report 2022 to 2023," Ministry of Textiles, New Delhi, India, 2023. [Online]. Available: https://texmin.nic.in/annual-report

[15] Z. Liu, P. Luo, S. Qiu, X. Wang, and X. Tang, "DeepFashion: Powering robust clothes recognition and retrieval with rich annotations," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Las Vegas, NV, USA, Jun. 2016, pp. 1096–1104.

[16] J. Devlin, M. W. Chang, K. Lee, and K. Toutanova, "BERT: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (NAACL)*, Minneapolis, MN, USA, Jun. 2019, pp. 4171–4186.

[17] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural Computation*, vol. 9, no. 8, pp. 1735–1780, Nov. 1997.

[18] T. Chen and C. Guestrin, "XGBoost: A scalable tree boosting system," in *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, CA, USA, Aug. 2016, pp. 785–794.

[19] Quantis, "Measuring fashion: Environmental impact of the global apparel and footwear industries study," Quantis International, Lausanne, Switzerland, 2018. [Online]. Available: https://quantis-intl.com/measuring-fashion-report

[20] Textile Exchange, "Preferred fiber and materials market report 2022," Textile Exchange, Lamesa, TX, USA, 2022. [Online]. Available: https://textileexchange.org/preferred-fiber-and-materials-report

[21] J. Cohen, "A coefficient of agreement for nominal scales," *Educational and Psychological Measurement*, vol. 20, no. 1, pp. 37–46, Apr. 1960.