# Π14. ΔΗΜΟΣΙΕΥΣΕΙΣ ΣΕ ΠΕΡΙΟΔΙΚΑ ΚΑΙ ΣΥΝΕΔΡΙΑ

SUB1.1: Συμπράξεις Ερευνητικής Αριστείας / Clusters of Research Excellence (CREs) - ΟΠΣ ΤΑ 5180519

**ΤΙΤΛΟΣ ΕΡΓΟΥ**

«Καινοτόμα Πλατφόρμα Ηλεκτρονικού Επανεμπορίου βασισμένη στην Τεχνολογία Blockchain για την Ενίσχυση του Κυκλικού και Βιώσιμου Μοντέλου Οικονομίας στη Βιομηχανία της Μόδας»

**ΑΚΡΩΝΥΜΙΟ:** REFASHION

**ΚΩΔΙΚΟΣ ΕΡΓΟΥ ΣΤΟ ΠΣΚΕ**: ΥΠ3ΤΑ-0560500

**ΑΘΗΝΑ 2026**

The present research is carried out within the framework of the National Recovery and Resilience Plan Greece 2.0, funded by the European Union – NextGenerationEU – Project Code: TAEDR-0560500 - Acronym: REFASHION).

Η παρούσα έρευνα υλοποιείται στο πλαίσιο του Εθνικού Σχεδίου Ανάκαμψης και Ανθεκτικότητας «Ελλάδα 2.0», με χρηματοδότηση από την Ευρωπαϊκή Ένωση – NextGenerationEU – Κωδικός Έργου: TAEDR-0560500 – Ακρωνύμιο: REFASHION

# Ελλάδα 2.0 ΕΘΝΙΚΟ
# ΕΘΝΙΚΟ ΣΧΕΔΙΟ ΑΝΑΚΑΜΨΗΣ και ΑΝΘΕΚΤΙΚΟΤΗΤΑΣ

## Συμπράξεις Ερευνητικής Αριστείας – ΣΕΑ

**ΕΡΓΟ: «Καινοτόμα Πλατφόρμα Ηλεκτρονικού Επανεμπορίου βασισμένη στην Τεχνολογία Blockchain για την Ενίσχυση του Κυκλικού και Βιώσιμου Μοντέλου Οικονομίας στη Βιομηχανία της Μόδας»**

**ΚΩΔΙΚΟΣ ΕΡΓΟΥ: ΥΠ3ΤΑ-0560500**

**ΠΑΡΑΔΟΤΕΟ Π14: ΔΗΜΟΣΙΕΥΣΕΙΣ ΣΕ ΠΕΡΙΟΔΙΚΑ ΚΑΙ ΣΥΝΕΔΡΙΑ**

| | |
|---|---|
| Διάρκεια Έργου: | 15 μήνες |
| Ημερομηνία Παράδοσης: | 31/05/2026 |
| Υπεύθυνος Παραδοτέου: | ΕΜΠ - ΕΕΕΕ |
| Επίπεδο Διάχυσης: | Δημόσιο |
| Τρέχουσα Κατάσταση Παραδοτέου: | Τελικό |
| Έκδοση: | F |
| Ιστοσελίδα Έργου: | https://refashion.cyclefi.com/ |

**ΦΟΡΕΙΣ ΥΛΟΠΟΙΗΣΗΣ:**

**ΑΘΗΝΑ 2026**

The present research is carried out within the framework of the National Recovery and Resilience Plan Greece 2.0, funded by the European Union – NextGenerationEU – Project Code: TAEDR-0560500 - Acronym: REFASHION).

Η παρούσα έρευνα υλοποιείται στο πλαίσιο του Εθνικού Σχεδίου Ανάκαμψης και Ανθεκτικότητας «Ελλάδα 2.0», με χρηματοδότηση από την Ευρωπαϊκή Ένωση – NextGenerationEU – Κωδικός Έργου: TAEDR-0560500 – Ακρωνύμιο: REFASHION

# Περιεχόμενα

# A Blockchain-Based E-Commerce Platform for Advancing a Circular and Sustainable Fashion Economy

Stavros T. Ponis(✉), Eleni Aretoulaki, George Plakas, and E. Gioti

School of Mechanical Engineering, National Technical University Athens, Athens, Greece
staponis@central.ntua.gr

**Abstract.** The fashion industry is undoubtedly a significant contributor to environmental degradation, by producing approximately 1.2 billion tons of carbon emissions each year, corresponding to around 8 per cent of global greenhouse gas emissions and consuming approximately 93 billion cubic meters of water per year, enough to meet the needs of five million people. In addition, the proliferation of the fast fashion model has exacerbated waste production, with millions of tons of textiles ending up in landfills annually. The lack of efficient and transparent recommerce platforms limits the potential for second-hand markets to contribute to circular economy models. This study introduces REFASHION, a blockchain-powered e-retail platform designed to facilitate second-hand trading, upcycling, and recycling of fashion products while ensuring security, transparency, and traceability through blockchain technology. The platform integrates an incentive-based mechanism through tokenized rewards to encourage sustainable consumer behaviors. This study details the technological framework, methodology, expected environmental and economic impact, and implementation challenges of REFASHION. The findings demonstrate that REFASHION has the potential to revolutionize the fashion recommerce industry by fostering sustainability and consumer engagement.

**Keywords:** Circular economy · Blockchain · Sustainable fashion · Textile waste management · Smart contracts · ReWear tokens

## 1 Introduction

In a circular economy, materials that can be reused are reintroduced into the economy in the form of new raw materials. These "secondary raw materials" can be sold and transported like primary raw materials from classical extractive resources. Currently, secondary raw materials comprise only a small proportion of the materials used in the EU. Waste management practices have a direct impact on the quantity and quality of these materials, and therefore actions to improve these practices are critical. The transition to a more circular economy, where the value of products, materials, and resources remains in the economy for as long as possible, and waste generation is kept to a minimum, is a necessary contribution to the efforts of the global community to develop a sustainable, efficient and competitive low-carbon and resource-efficient economy.


K. Arai (Ed.): IntelliSys 2025, LNNS 1553, pp. 601–610, 2025.
https://doi.org/10.1007/978-3-031-99958-1_40

A key factor for the creation of a dynamic market for secondary raw materials is substantial supply and sufficient demand, driven by citizens shaping supply chains through the integration of second-hand and sustainable goods into their daily lives and active participation in a barter economy. The barter economy is driven by the exchange of time, services, and goods and greatly impacts it, as it is based on principles such as collectivity, solidarity, and cooperation. To achieve this, the sharing economy utilizes resources that are available but untapped, with mutual benefits for those who give up the resources and those who use them, to reduce the waste generated and contribute to the gradual change of citizens' culture in relation to waste management. This goal can only be achieved through a series of waste prevention measures: the introduction of new and the strengthening of existing distinct waste streams, the promotion of reuse, the enhancement of recycling rates, and of course the promotion of the market of secondary materials, which is also at the heart of the REFASHION project.

The successful transition to the circular economy requires the movement of second-hand goods, their repair, reuse, or reconstruction, while informing the public about the reuse of objects. Nonetheless, despite the increasing relevance of sustainable fashion and the opportunities of the re-commerce business model, there is still an under-research gap of blockchain applications in the context of incentive-based second-hand trade platforms, which are more fashion waste-focused. Current recommerce platforms typically do not include trust mechanisms, traceability, and environmental value tracing, thus falling short of fully supporting the circular economy. This gap is connected through a novel concept of trust in value transactions, where a trust model and reward system are developed to enable blockchain-based value transactions while also promoting sustainability-conscious consumer behavior through tokenized rewards.

In this direction, the REFASHION project aims precisely at accelerating circular economy's actions and linking the social economy with technological innovation towards an alternative model of waste management, encouraging social participation and upgrading the role of citizens. The proposed e-commerce platform emerges as an innovative response to the urgent call for the development and implementation of circular economy business models. By redefining the life cycle of fashion products, the REFASHION project embodies the principles of circularity, extending the usefulness of textiles, garments and footwear after the end of their useful life, through a user-centric marketplace, which facilitates the sale, donation and purchase of second-hand goods. This platform not only supports the shift towards a more sustainable consumption model but also pioneers new business models that incentivize environmentally friendly practices. Through its commitment to reducing waste and promoting recycling, as well as creative upcycling, the proposed re-commerce platform is fully aligned with the objectives of the circular economy's strategies, offering a scalable, effective solution designed to meet the contemporary challenges of environmental protection and resource efficiency.

The rest of this study is organized as follows: Sect. 2 describes the environmental and socio-economic challenges for fashion waste, which motivates the proposed solution. Section 3 describes the REFASHION platform design and main features, with the focus on the technological infrastructure and incentive schemes. Section 4 describes the anticipated environmental and economic effects of the development. Eventually, Sect. 5 concludes the research and provides the summary, the challenge, and the future works.

## 2 Problem Statement

Every year, millions of tons of clothing and footwear go to waste, causing multiple environmental and socio-economic impacts that threaten the eco-balance of the planet and the well-being of societies [1]. Specifically, the fashion industry is the second most polluting industry in the world, responsible for approximately 8 per cent of the global greenhouse gas emissions, in essence, more emissions than those from the international aviation and maritime sectors combined, with an expected increase of 50 per cent by 2030 [2]. At the same time, it is responsible for 20 per cent of global water consumption, equivalent to about 93 billion cubic meters per year on a global scale [3, 4].

The fashion industry's dependence on fast fashion trends exacerbates these issues through an unsustainable increase in the use of natural resources. The fast fashion cycle, characterized by the constant release of new collections, promotes rapid cycles of production and consumption, encourages a culture of disposability and overconsumption and undermines the principles of material conservation and sustainable use [5]. Considering the growing world population, there has been an excessive increase in the resources used for textile production, reaching 98 million tons of non-renewable resources per year [6]. Per capita fiber consumption almost tripled from 1950 to 2008, from 3.7 kg to 10.4 kg per person [7], while from 2007 to 2014, fiber production increased by an additional 20.2 million tons to 90.8 million tons, which is expected to increase by 3.7 per cent per year [8].

This current model not only perpetuates the depletion of resources but also exacerbates the problems that arise when fashion products reach the end of their useful life and are discarded. Landfill overflow, driven by the non-biodegradable nature of widely used synthetic fibers (e.g. polyester) and increased pollutants from toxic dyes and chemicals used in textiles, exacerbates waste management challenges, posing serious risks of degradation of aquatic ecosystems and the wider environment. In 2015, 92 million tons of fashion waste were generated globally and projected to increase by 56 million tons by 2030 [1]. In Europe, it has been reported that only 15 to 20 per cent of textile waste is collected for recycling or reuse and the rest is disposed of in landfills or incinerated [9, 10]. For Greece, the amount that ends up in landfills corresponds to approximately 125 to 250 thousand tons of textile waste per year [11]. Inadequate management of waste fashion products, in turn, incurs huge costs for creating additional space for landfill [12].

The seriousness of the problem is such that in October 2020, at the G7 summit, hosted in Biarritz, French President Emmanuel Macron presented a new pact for the fashion industry, the so-called "Fashion Pact". This is a set of common objectives agreed by major companies in the sector to reduce environmental impact. The "Fashion Pact" brings together more than 60 global fashion industry companies, suppliers and retailers, representing more than 200 brands and 1/3 of the industry, to implement seven climate, biodiversity and ocean goals. In particular, the signatories of the "Fashion Pact" committed to achieving net-zero carbon emissions by 2050 and keeping global warming below 1.5 °C by 2100; protect endangered animals, restore critical natural ecosystems and eliminate unnecessary plastics in packaging.

This harsh reality demonstrates the fact that the harmonization of the fashion industry with the basic rules of sustainability and circular economy is still in its infancy. Therefore, the need to shift towards more sustainable practices is imperative. Clearly, such a transition requires the concerted effort of all those involved in the fashion value chain. Moving in this direction, it initially requires a paradigm shift in the way consumers perceive fashion and an installation of a conscious shopping mindset. From an industry perspective, circular practices are essential, including resilience-driven design, to transform from a take-make-dispose linear model to one based on reuse and recycling. For producers, it includes investments in technologies that recycle textiles on a large scale, as well as methods of creative recycling (upcycling) and energy recovery. Overall, this is a call to action aimed at systemic overhaul of the current fashion field, requiring collaboration, innovation and determination across the EU value chain. Such a transition is fully in line with the legislative initiatives of the European Commission, which in summer 2023 revised the Waste Framework Directive, incorporating the enforcement of harmonized Extended Producer Responsibility (EPR) schemes for textiles, and published the 'EU strategy for sustainable and circular textiles', which aims to the digital and 'green' transformation of fashion industries in the European Union [13].

## 3 Proposed Solution

The comparative advantages arising from the purchase and sale of second-hand items are many. For example, any item purchased second-hand can lead to a reduction in the production of new products due to a decrease in active demand. The reuse of products obviously reduces air pollution from factories and the energy required, potential waste, as well as resource consumption. By buying second-hand products, production, packaging and transportation requirements are reduced, as well as the storage or depletion of some items that could still serve needs and be reused.

The proposed solution of the REFASHION project includes the introduction of an innovative model of distribution of secondary products and exchange economy and its support with the development of a decentralized exchange economy platform, based on blockchain technology, through which continuous communication between stakeholders (citizens, businesses, waste management bodies, system operator) will be ensured. The platform barter economy aims to develop and promote local second-hand fashion reuse chains allowing significant profit to sellers and buyers, reducing the cost of purchasing new items, creating a new revenue stream and giving value to an item previously considered as waste. In addition, the barter economy platform will allow the creation of a reliable information network for all stakeholders, through which citizens will communicate the type, quantity, date of placement and geographical location of second-hand items, using an Online Distributed Ledger, the operation of which is based on the computational paradigm of blockchain. The REFASHION platform will function through a private permissioned blockchain that maintains user data privacy while achieving transparency and security with optimal performance. The platform will utilize Proof of Authority (PoA) consensus to achieve low energy usage and high transaction speeds unlike the energy-demanding Proof of Work (PoW) systems. Smart contracts control essential operations on the platform by managing product listings and ownership verification, along with payment escrow, transaction validation and reward distribution. Smart contracts will resolve disputes by holding transactions in place until arbitration concludes, which builds user confidence and reduces the possibility of fraudulent activities.

Once items reach the end of their useful life, citizens will be able to declare that the registered quantity is available for reuse, while interested citizens who wish to acquire second-hand items will be properly informed through a specially designed mobile app, thus successfully completing the cycle of reuse of objects. To foster collaboration across supply chain links, the application will offer an intelligent waste factoring system (marketplace) that ensures high traceability. This will facilitate the secure buying and selling of second-hand items while synchronizing transportation logistics with real-time item metadata, such as product condition, authentication status, and transaction history.

In addition, blockchain technology will provide tokens, as a reward, to the participants of the platform, thus creating an incentive to participate and collaborate. Tokens will function as a superior incentive to cash since they help users earn additional discounts and services while unlocking special offers within the platform ecosystem, which will boost user loyalty and maintain continuous engagement. Finally, through encryption, provided by blockchain technology, confidentiality, transparency and traceability will be ensured of transactions, characteristics that are particularly critical for the fashion industry, which has traditionally been plagued by the phenomenon of copying and counterfeiting of trademarks and products. The choice of this technology is fully in line with the actions of the European Union, which allocated €44 million in funding to promote sustainability in textiles through blockchain technology, with the aim of promoting sustainable economic growth, tackling climate change and supporting the European Green New Deal [5].

The online exchange economy platform will be user-friendly, fast, completely secure, and will provide a fully automated system for matching products and services, thus giving more opportunities to buy, exchange and/or negotiate through bidding, an internal communication system between members and an exchange evaluation system. Essentially, the REFASHION platform will enable citizens and businesses to donate or monetize fashion products they no longer need by selling or exchanging them for others, with the possibility of payment either by paying money or by using the app's native ReWear token. The ReWear token could also be used as a medium of exchange outside the app's ecosystem with collaborating companies from all industries.

More specifically, the proposed platform will include the following key functions:

### 3.1 Registration

Users of the platform can be either individuals or businesses, e.g. second-hand retailers, centers for recycling, upcycling or waste-to-energy, thus providing a variety of options for managing fashion products according to their quality level. Upon first opening the application, users are asked to enter their details in a registration form. For individuals, basic information is required; such as name, email address, and password, unlike businesses, which are asked to record their company's details, i.e. name, type, address, hours, e-mail address, contact phone number and details of the person in charge who will manage the application, combined with a password, for their successful registration on the platform.

### 3.2 Resale-Donate-Purchase

Depending on who uses the application and what they wish to do (sale/donation/purchase), its functions are customized. When it comes to selling, users can list the products they want, providing photos, the product category (e.g. women's dress), brand, size, condition (e.g. excellent), occasion (e.g. partywear) and selling price. Similarly, for buyers, by opening the application and activating the geolocation function, using GPS, a map is loaded that highlights the location of users and available fashion products for sale. There they can choose from the list of available products based on the results of the platform's machine learning models, which provide personalized and sustainable fashion recommendations based on user purchase history, browsing behavior, the location of the product and its other characteristics (category, condition, circumstance, size, brand, price), but also the Eco-Impact Score, which represents the environmental benefits that will result from the reuse of the product in question (see Fig. 1). In that way, the Artificial Intelligence (AI) system will suggest pre-owned clothing items that align with both the user's style preferences and sustainability goals. The recommendation engine also considers material composition, brand sustainability ratings, and resale value to promote circular fashion practices.

Special attention should be put, to the platform's Eco-Impact Score, which is a key innovation of the REFASHION project, and includes various environmental metrics, such as CO2 savings, water saving and waste reduction in a single composite index, giving users a quick and easy way to understand the environmental benefits of their purchases, highlighting the multifaceted value of choosing second-hand items over new ones. After selecting the product(s) they are interested in through the platform, users can make secure transactions, ensuring the quality and authenticity of the products through the use of blockchain technology. Once a transaction has been made, the seller sends the product either to the buyer's site or to a "smart" parcel delivery machine, in eco-friendly packaging, and upon delivery, the buyer confirms the receipt through the app, finalizing the transaction and triggering the release of funds to the seller. Blockchain technology is also used in this phase through a monitoring system, which provides real-time updates to both buyer and seller, ensuring transparency and security until the successful delivery of the item.

### 3.3 Validation-Evaluation-Rewarding

To enhance trust and transparency in the REFASHION platform, critical functions of evaluation and validation of the condition and quality of resold fashion products are implemented. Through a sophisticated system of reviews from the user community, as well as quality control procedures, the platform ensures the reliability of the recommerce ecosystem, contributing to the creation of a safe and positive trading environment. Blockchain technology plays a central role in this process, offering transparency and traceability. By leveraging it, it becomes possible to validate the origin and authenticity of each product, boosting trust among community members and allowing consumers to make informed purchasing decisions. In addition, upon successful completion of a resale, the blockchain's smart contracts system becomes operational, ensuring that both the buyer and the seller or donor are rewarded with a specific number of ReWear

Tokens, directly proportional to the ecological footprint of the transaction. These tokens, reflecting the platform's commitment to a circular economy, are then automatically and securely transferred to the digital wallets of the parties involved. Users will have the opportunity to leverage their ReWear tokens in purchases within the platform, but also in an extensive network of participating companies, unlocking discounts, and enjoying a variety of benefits.

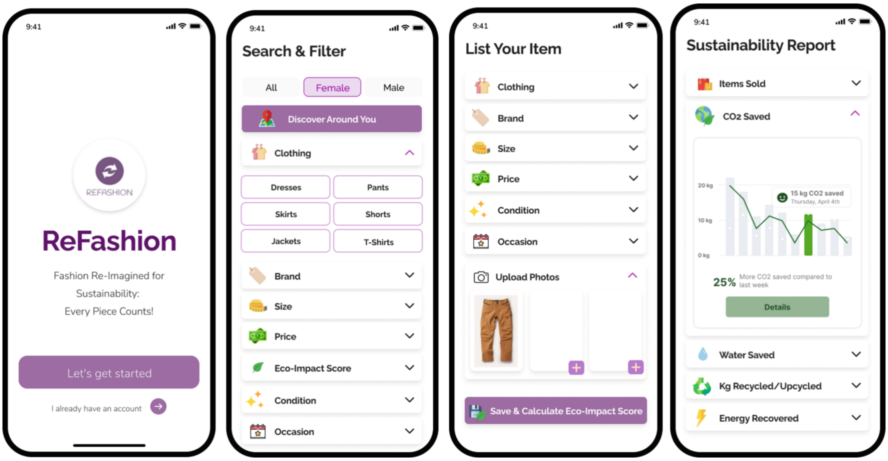


**Fig. 1.** (a) Launch Screen, b) Search for Fashion Products to Buy (Buyer UI), c) Registration of Fashion Products for Resale (Seller/Donor UI), d) Sustainability Report (Business UI)

## 4 Expected Impact

The REFASHION recommerce platform aims not only to create environmental benefits by extending the lifespan of fashion items but also to introduce a new economic model that encourages sustainable consumer behavior. Through blockchain technology, the project platform ensures authenticity, traceability and security of transactions, addressing the prevailing issues of trust and quality that plague existing and particularly "immature" recommerce markets. Finally, the project explores the technical and economic challenges associated with the development and scaling up of such a platform, its potential impact on reducing fashion waste, and enhancing consumer engagement in the project ecosystem, through appropriate incentive mechanisms to change their environmental behavior.

At the same time, the project will provide multiple benefits for all stakeholders of the application. On the one hand, individuals will participate in a sustainable buying and selling cycle, earn reward tokens that will incentivize their further participation and enhance their sustainable behaviors and mindset. On the other hand, second-hand fashion companies, in addition to reward tokens, gain access to a wider eco-conscious customer base through the platform, enhancing their sales and visibility. However, the benefits are not limited to those directly involved in the application ecosystem. For both the state and society in general, the platform is aligned with the environmental objectives of the National Action Plan for Recycling and Green Development, contributing to waste reduction and wider efforts against climate change. In addition to strengthening these efforts, the economy will also benefit directly, considering that the proposed project will significantly increase the efficiency of material collection, processing and recovery systems from companies such as those participating in the project consortium. Indeed, through the resale and donation of used items, the platform can significantly reduce the amount of waste ending up in landfills. This not only alleviates the pressure from waste management systems but also supports environmental sustainability through extending the life cycle of products. Furthermore, companies specializing in recovery and upgrading can benefit from a more reliable supply of materials, as the platform offers an organized way to collect used clothing products that are suitable for recycling or upcycling.

Early signs suggest that REFASHION could create meaningful economic value for both individuals and businesses. In recent years, second-hand fashion platforms have been growing steadily. Actually, by 2026, the global secondhand apparel market is expected to grow 127 per cent, which is three times faster than the apparel market overall. Some forecasts even predict that the global resale market could exceed $70 billion by 2027 [14]. For everyday users, REFASHION could turn unused clothing into a modest but noticeable income boost, depending on how active they are. Buyers stand to benefit too, with potential savings of 25 per cent to 50 per cent compared to shopping for brand-new items [15]. And beyond personal gains, expanding the platform could open up new job opportunities in logistics, quality assurance, and digital operations. Looking ahead, we aim to develop detailed economic models to better understand just how large these impacts could become as more users and businesses join the REFASHION ecosystem.

In the long term, the proposed project will:

1. Facilitate the implementation of the Waste Framework Directive [16], which seeks to reduce waste generation and promote recycling and reuse. The transparency and traceability offered by blockchain technology can improve the efficiency of the industry's material recycling and recovery processes.
2. Contribute to promoting environmental sustainability in the EU by promoting sustainable consumption and production, in line with the objectives of the European Green Deal [17], to reduce CO2 emissions.
3. Contribute to the achievement of the objectives of the Fashion Pact, which focuses on sustainability in the fashion industry, through waste reduction and the promotion of the circular economy. The increased use of materials from reuse and recycling can reduce the environmental impact of the fashion industry.
4. Help implement the objectives of the European Union's strategy "For Sustainable and Circular Textiles" for the green transition of the ecosystem, according to which by 2030, textiles placed on the EU market are long-lived and recyclable, largely made from recycled fibers, free of hazardous substances and produced with respect for social rights.
5. Help develop and promote local second-hand fashion reuse chains, enabling significant profit for sellers and reducing the cost of purchasing new items for buyers,

creating a new revenue stream and giving value to items that would otherwise be disposed of as waste.
6. Scale the implementation of its results in such a way as to expand the network of users and partners domestically and potentially globally, attracting manufacturers, retailers, recycling centers and consumers in a joint effort to promote sustainable fashion, making the REFASHION platform a center of innovation and sustainability.

## 5 Conclusions

The REFASHION platform introduces a blockchain-powered solution to address the sustainability challenges of the fashion industry by promoting circular economy principles. Through decentralized transactions, tokenized incentives, and smart contract automation, the platform enhances transparency and trust in the resale of fashion products. By integrating eco-impact scoring and AI-driven recommendations, REFASHION not only facilitates ethical consumer choices but also encourages sustainable behavior. The findings of this study indicate that blockchain technology has the potential to revolutionize the second-hand fashion industry, reducing waste and contributing to a more responsible and resource-efficient fashion ecosystem. The results highlight the feasibility of using blockchain to build a secure, transparent, and user-driven recommerce ecosystem for fashion products. The integration of ReWear tokens as incentives has shown promise in fostering active user participation, promoting sustainable habits, and enhancing the value of second-hand goods. Furthermore, the REFASHION platform aligns with global sustainability initiatives, such as the EU Strategy for Sustainable and Circular Textiles and the Fashion Pact, contributing to the broader agenda of reducing environmental impact in the fashion industry.

Despite its promise, REFASHION faces challenges, particularly the slow adoption of blockchain in consumer-driven recommerce due to technological complexity, digital literacy barriers, and skepticism towards decentralized systems. Achieving widespread user engagement may require significant educational efforts and the development of a more user-friendly interface. Additionally, while blockchain enhances transparency and security, its energy consumption presents a potential drawback. The environmental benefits of REFASHION must be weighed against the energy demands of blockchain transactions, particularly those utilizing Proof of Work (PoW) consensus mechanisms. This issue underscores the importance of exploring low-energy blockchain models, such as Proof of Stake (PoS), to mitigate the environmental footprint of decentralized platforms.

In order to overcome slow uptake and complex UI challenges, a number of strategies are recommended, such as a) an extensive user introduction including tutorials, tooltips, and interactive guides to transition from blockchain-noob to master, b) targeted marketing campaigns pitching the environmental impact of the platform and the rewards for users to quickly own value proposition and c) incentivizing early adopters with bonus tokens or similar discount to kick start the marketplace and achieve organic growth.

The REFASHION platform represents an innovative approach for tackling the growing problem of fashion waste through blockchain-enabled recommerce. While challenges remain in adoption and scalability, the potential environmental and economic benefits make this a promising area for further research and practical implementation. Future advancements in blockchain efficiency, consumer engagement strategies, and policy frameworks will be critical in scaling REFASHION into a widely adopted and impactful sustainable fashion marketplace.

**Acknowledgments.** The present work is carried out within the framework of the National Recovery and Resilience Plan Greece 2.0, funded by the European Union – NextGenerationEU – Project Code: TAEDR-0560500.

## References

1. Stanescu, M.D.: State of the art of post-consumer textile waste upcycling to reach the zero waste milestone. Environ. Sci. Pollut. Res. **28**(12), 14253–14270 (2021)
2. Bailey, K., Basu, A., Sharma, S.: The environmental impacts of fast fashion on water quality: a systematic review. Water **14**(7), 1073 (2022)
3. Dolzhenko, I.B., Churakova, A.A.: Environmental Responsibility of Fashion Industry Multinational Corporations (MNCs) in the Context of Industry 4.0. In Industry 4.0: Fighting Climate Change in the Economy of the Future, pp. 79–89. Springer, Cham (2022)
4. Peters, G., Li, M., Lenzen, M.: The need to decelerate fast fashion in a hot climate—a global sustainability perspective on the garment Industry. J. Clean. Prod. **295**, 126390 (2021)
5. Niinimäki, K., Peters, G., Dahlbo, H., Perry, P., Rissanen, T., Gwilt, A.: The environmental price of fast fashion. Nat. Rev. Earth Environ. **1**(4), 189–200 (2020)
6. Ellen MacArthur Foundation. A New Textiles Economy: Redesigning fashion's future 2017. Accessed 4 April 2024 (2017). https://ellenmacarthurfoundation.org/a-new-textiles-economy
7. Sanchis-Sebastiá, M., Ruuth, E., Stigsson, L., Galbe, M., Wallberg, O.: Novel sustainable alternatives for the fashion industry: a method of chemically recycling waste textiles via acid hydrolysis. Waste Manag. **121**, 248–254 (2021)
8. Pensupa, N., et al.: Recent trends in sustainable textile waste recycling methods: current situation and future prospects. Top. Curr. Chem. **375**, 76 (2017)
9. Beton, A., et al.: Environmental improvement potential of textiles (IMPRO Textiles). Eur. Comm. 20 (2014)
10. Thakker, A.M., Sun, D.: Sustainable development goals for textiles and fashion. Environ. Sci. Pollut. Res. **30**(46), 101989–102009 (2023)
11. European Environment Agency (eea.europa.eu) (2021) 'Textiles in Europe's circular economy'. https://www.eea.europa.eu/publications/textiles-ineuropes-circular-economy
12. To, M.H., Uisan, K., Ok, Y.S., Pleissner, D., Lin, C.S.K.: Recent trends in green and sustainable chemistry: rethinking textile waste in a circular economy. Curr. Opin. Green Sustain. Chem. **20**, 1–10 (2019)
13. European Commission. EU strategy for sustainable and circular textiles. Environment.ec.europa.eu (2022). https://environment.ec.europa.eu/strategy/textiles-strategy_en
14. Retail Dive (2023). ThredUp: US resale market projected to reach $70B by 2027. https://www.retaildive.com/news/thredup-us-resale-market-projected-to-reach-70b-by-2027/646861/
15. The Week. How much can you save shopping secondhand? (2024). https://theweek.com/personal-finance/secondhand-shopping-saving-thrift-store
16. European Union. Directive 2008/98/EC on waste and repealing certain Directives (2024). https://eur-lex.europa.eu/legal-content/
17. Fetting, C.: The European Green Deal. ESDN Report, December 2020, ESDN Office, Vienna (2020)

# A Unified Ensemble Framework for Detecting Manipulated and AI-Generated Facial Content

Anastasia Akrioti *, *ece19003@go.uop.gr*
Emmanuel Pintelas, *e.pintelas@go.uop.gr*
Vassilis Tsakanikas, *v.tsakanikas@uop.gr*
Vassilis Tampakas, *tampakas@uop.gr*
*Distributed Intelligent Systems and Data Lab (DISyD Lab), Department of Electrical and Computer Engineering, University of Peloponnese*

**ABSTRACT**

Deepfake videos have become a serious and technological threat, as they are used to spread fake news, misinformation and manipulate public opinion. Advanced techniques for synthesizing and editing multimedia content make the detection of such manipulated videos particularly challenging, creating the need for reliable detection solutions. This paper proposes a Deepfake detection method based on Ensemble Learning, aiming to improve the efficiency of classifying images as real or fake. The implementation leverages an ensemble of four pre-trained convolutional neural networks — EfficientNet-B0, ResNet50, MobileNetV3, and DenseNet121— combined with an adaptive weighted soft voting strategy. The analysis is conducted on three distinct datasets, each with unique characteristics, as well as on a constructed Unified dataset, with all three datasets combined. The classifier is developed using the Python programming language and modern deep learning libraries; its performance is evaluated using metrics such as AUC-ROC, Accuracy, and Sensitivity. The experimental results demonstrate that the proposed ensemble outperforms individual baselines, confirming the potential of a multi-architecture strategy as an effective tool in addressing the deepfake detection problem.

**Keywords:** Deepfakes, Convolutional Neural Networks, Deep Learning, Ensemble Learning

## 1 INTRODUCTION

In recent years, fake news has spread rapidly through social media, posing a serious threat to public discourse and human society. The rise of Deepfakes has exacerbated this situation, as it allows the creation of convincing synthetic video content through deep learning algorithms [1]. Deepfakes are a product of artificial intelligence that modifies images, videos, and audio, creating fake content. The term DeepFake (DF) comes from the combination of the words "Deep Learning (DL)" and "fake" [2].

The creation of fake content is not a new phenomenon; however, recent developments in artificial intelligence (AI) and deep learning have dramatically increased the spread and quality of deepfakes. The practice of producing realistic fake online creations is becoming increasingly common on social media platforms, with numerous images and videos of celebrities, politicians, and other well-known figures. The social impact is significant and often harmful, especially since producing deepfakes now requires minimal technical training and basic equipment [3].

Deepfakes can now mimic facial expressions, lip movements, and even voices with remarkable accuracy, making them nearly indistinguishable from real videos. These capabilities make the technology extremely dangerous, creating a growing need for the development of effective detection techniques that can address its abuse [4].

With the advancement of artificial intelligence and computer vision, the most effective way to create deepfakes is based on machine learning techniques, and in particular generative adversarial networks (GANs), which in turn leverage the structure of convolutional neural networks (CNNs). The use of GANs makes it possible to create highly realistic images and videos, thereby contributing to the further spread of the phenomenon [5], [6].

Alongside, many detection methods have also been developed. However, a significant issue remains that most detection models are trained on specific data sets and struggle to maintain their performance when applied to different or new data.[7] This lack of generalization means that models may fail when confronted with new types of manipulation that were not included in the original training set. For this reason, continuous updating and expansion of datasets is essential to ensure that detection techniques remain effective against ever-evolving methods of creating deepfakes [8].

In this paper, we advance the use of ensemble learning, a technique in which we combine the predictions of many different models in order to achieve better performance than relying on a single model alone [9]. The central idea is that different models tend to make different errors. Therefore, if we combine their predictions, the errors can be balanced out, leading to the result to be more accurate and more generalizable [10]. The contribution of this study lies not only in its technical approach to detecting deepfake content but also in emphasizing the severity of the issue, highlighting the need for advanced mechanisms to protect against the spread of fake audiovisual material.

We then compare the performance of individual baseline models with the final hybrid model that utilizes ensemble learning. The evaluation is performed on three different datasets, each with unique characteristics. Furthermore, we merged these three datasets into a single unified dataset, in order to increase diversity within a single dataset. The goal of this approach is to build a more robust classifier for deepfake detection, one capable of performing better in different conditions and types of manipulated content.

.

# 2 RELATED WORK

Recently, many methods have been proposed for deepfake detection, utilizing different architectural models.

M. El-Gayar et al. introduced an advanced method for deepfake video detection that leverages graph neural networks (GNNs) and convolutional neural networks (CNNs). This method incorporates GNNs to dynamically adjust relationships between neighboring nodes, while CNNs contribute by extracting discriminative visual features. The model uses a hierarchical multiscale structure to capture fine-grained features, thereby improving representation and detection accuracy. The process consists of two phases: a mini-batch flow and a four-block CNN flow, which are integrated through three fusion networks (FuNet-a, FuNet-M, FuNet-C). The accuracy of the method across different datasets was 95.09% on FF++, 99.3% on DFDC, and 98.9% on Celeb-DF [4].

Another method proposed by Zeina Ayman et al. uses the VGG (Visual Geometry Group) model. VGG is a pretrained neural network based on the CNN architecture with five convolutions and max pooling blocks, which recorded an accuracy of 88%. In contrast, a standard CNN, after preprocessing with face extraction and resizing, reached an accuracy of 94%. These results indicate that CNNs can outperform VGG in deepfake detection [6].

Alongside these, Halliasos et al presented an approach based on self-supervision. This approach is called RealForensics and consists of two stages. In the first stage, the model learns time-dense video representations through intermodal self-supervision from many speaking faces. These representations are then used as prediction targets in the second stage to normalize binary forgery classification. The method achieved 97.1% on FaceShifter dataset and 95.7% on DeeperForensics dataset [11].

Lastly, Feng et al developed a method based on detecting anomalies between sound and image in video. Their model combines information from both modalities and is trained exclusively on real data to identify inconsistencies, without requiring exposure to deepfake data during training. The method demonstrates strong cross-data performance, achieving results with AUC scores around 94-95% on datasets such as FakeAVCeleb and KoDF [12].

Overall, related work demonstrates that many effective techniques have been developed, which exploit different architectures. However, a key issue remains: most approaches perform well on specific datasets. This highlights the need for more robust solutions that leverage multi-source datasets and the combination of multiple models. In this context, the present paper focuses on a systematic evaluation of four CNNs on a unified dataset, along with the use of ensemble learning.

# 3 DATASETS & PREPROCESSING

In this study, we used three datasets that include both still images and video frames. Specifically, for the video data, face detection was first applied, followed by the extraction of the middle to enable image level classification. The dataset was split into training and testing sets with an 80:20 ratio to ensure sufficient data for training while maintaining a reliable control set for performance evaluation.

The DFDC dataset (DeepFake Detection Challenge) is one of the largest publicly available deepfake datasets, which employs a variety of methods for creating fake videos, including both advanced GAN-based techniques and simpler approaches, aiming to train effective detection models. The version of the DFDC used contains 1500 fake images and 1500 real images, extracting one frame per video and isolating the faces of each individual [13].

Deepfake MNIST + is an advanced dataset that focuses on animated facial expressions instead of identity exchange, filling a major gap in deepfake detection research. It was created using modern image animation frames, simulating 10 specific

movements, such as blinking and surprise. In addition, to avoid bias in training, it was ensured that each sample represents a unique individual, forcing the model to truly distinguish between authentic and fake videos [14].
EXP dataset is a unique collection of processed images of faces created by experts from the Computational Intelligence and Photography Lab at Yonsei University, using advanced tools such as Photoshop. It includes approximately 1,000 original and 1,000 edited images, where facial features such as eyes, nose, and mouth have been replaced or altered to compose new images. EXP offers high-quality data that can train classifiers capable of detecting fake images regardless of the method of creating them [15].
To preprocess the data, we used the function transform.Compose, which is owned by the library torchvision of the PyTorch framework [16]. This function allows us to combine a lot transformation to the images. In detail, it combines a list of transformations applied sequentially to each image that passes through the dataset. In this particular case, we applied the function Resize to standardize the dimensions on all samples and the function ToTensor, in order to convert the image to a tensor format so it can be processed by the neural networks. During training, batch size 32 was used to optimize GPU memory usage. Also, we assign binary labels to real (0) and fake (1).

# 4 PROPOSED METHOD

## 4.1 Overview

To overcome the limitations of single model deepfake detectors, we propose a heterogeneous ensemble learning framework. In contrast to traditional approaches that rely on a single backbone network, our method integrates the feature representations of four distinct CNNs architectures: ResNet50, EfficientNet-B0, MobileNetV3-Large, and DenseNet121. The architectural diversity of this ensemble enables the system to capture complementary information, from high-level semantic inconsistencies to subtle textural differences. As a result, the proposed method demonstrates increased robustness and generalization capability.

## 4.2 CNN Backbones

We selected these four backbones based on their structural distinctiveness:

- ResNet50: Selected as a robust baseline, leveraging residual connections to enable deep learning extraction without the vanishing gradient problem [17].
- EfficientNet-B0: Chosen as a state-of-the-art approach that employs a compound scaling to optimally balance network depth, width and resolution [18].
- MobileNetV3-Large: Included to evaluate a lightweight architecture that utilizes Squeeze-and-Excitation blocks to adaptively focus on informative image features [19].
- DenseNet121: Incorporated for its dense connectivity pattern, which promotes feature reuse and preserves low-level artifacts critical for detection [20].

## 4.3 Ensemble Strategy

The proposed ensemble framework employs a weighted soft voting strategy, where the contribution of each backbone is calculated based on its individual performance. Rather than assuming equal reliability across all classifiers (Simple Averaging), we performed a grid search to optimize the weight coefficients ($w_i$), a practice strongly supported by recent reviews on ensemble deep learning [9]. This approach is mathematically motivated by assigning larger weights to classifiers with lower variance and higher accuracy, we effectively minimize the expected generalization error of the ensemble [21]. Our experiments demonstrated that this weighted fusion consistently outperforms the simple averaging baseline, aligning with findings in modern deepfake detection methodologies [10].

# 5 EXPERIMENTS & RESULTS

## 5.1 Implementation Details and Training Strategies

All experiments were conducted in a computational environment equipped with dual NVIDIA Tesla T4 GPUs (16GB VRAM each) using PyTorch framework. The implementation code of the proposed model can be found in the following link. Input images were preprocessed by resizing to a fixed resolution of 224x224 pixels and normalized using standard ImageNet statistics (mean=[0.485,0.456,0.406], std=[0.229,0.224,0.225]). To ensure maximum architectural diversity, we tailored the training strategy of each backbone: To establish a reliable performance benchmark, we employed ResNet50 as our baseline architecture. This model was trained using a standard transfer learning approach with the Adam optimizer ($lr = 10^{-4}$) and a constant learning rate strategy [22]. No layer specific freezing or complex scheduling model was applied, allowing it to serve as a stable control variable against which the other optimization strategies could be evaluated.

For the EfficientNet-B0 backbone, we implemented a targeted fine-tuning strategy to balance computational efficiency with domain adaptation. Initially, all feature extraction layers were frozen to retain the robust high-level representation learned from the ImageNet. To achieve task specificity, we explicitly unfroze only the last five layers of the network. Concurrently, the final classification head was replaced with a linear layer designed for binary classification.

While standard hyperparameters proved effective for the baseline, the lightweight MobileNetV3 initially exhibited instability. To address this, we switched from Adam to Stochastic Gradient Descent (SGD) with a learning rate of 0.0015, momentum of 0.9, and Nesterov acceleration [23]. Additionally, we implemented a ReduceLROnPlateau scheduler and extended the fine-tuning depth by unfreezing the last 8 layers.

Conversely, the DenseNet121 architecture required a strategy to prevent gradient instability during the early training phases. We utilized the Adam optimizer but replaced the standard scheduler with a hybrid SequentialLR strategy.[24] This involves a linear warmup phase for the first 2 epochs (gradually increasing the learning rate from $10^{-5}$ to $10^{-4}$), followed by Cosine Annealing to decay the rate to $10^{-6}$.

Table 1: Hyperparameter Configuration per Backbone

| Model Architecture | Training Strategy |
|---|---|
| ResNet50 (Baseline) | Adam |
| DenseNet121 | Adam + Warmup |
| MobileNetV3 | SGD + Scheduler |
| EfficientNet-B0 | Adam (5-layer unfreeze) |
| Proposed Ensemble | Adaptive Weighted Voting |

## 5.2 Comparative Analysis

The overall evaluation of the proposed methodology is summarized in Table 2, which provides a comparative analysis of the AUC-ROC performance between individual CNN backbones and ensemble strategies. It is observed that no single network consistently dominates across all datasets. Notably, the weighted fusion strategy, that calibrated via Grid Search, always maximizes the AUC-ROC, outperforming the simple averaging method in every scenario. These results validate that combining diverse architectural features through adaptive weighting is crucial for mitigating domain shifts and achieving generalization.

Table 2: AUC-ROC Performance Comparison across Datasets

| Dataset | ResNet50 | MobileNetV3 | EfficientNet-B0 | Ensemble (simple avg) | Ensemble |
|---|---|---|---|---|---|
| DFDC | 0.7303 | 0.7252 | 0.7836 | 0.8023 | 0.8041 |
| DFMNIST | 0.8457 | 0.8854 | 0.9070 | 0.9096 | 0.9144 |
| EXP | 0.7050 | 0.7489 | 0.8017 | 0.8014 | 0.8091 |
| Unified | 0.7620 | 0.7783 | 0.8311 | 0.8401 | 0.8463 |

Beyond the AUC score, Table 3 provides a detailed breakdown of the ensemble's evaluation metrics. It is crucial to highlight the system's high sensitivity (recall) in almost all datasets. This indicates that the method effectively minimizes the risk of false negatives (i.e., failing to detect a manipulated image or video). The system demonstrates exceptional Sensitivity (up to 89.06%) on AI-generated datasets like DFMNIST. A notable exception is the EXP dataset, where the Sensitivity drops to 56.63%, despite maintaining robust Specificity (82.63%). This discrepancy suggests that while the ensemble is highly effective against GAN generated artifacts, detecting high quality photoshop images remains a distinct challenge compared to AI.

Table 3: Evaluation Metrics of the Proposed Ensemble

| Dataset | Accuracy | AUC-ROC | Sensitivity | Specificity | Precision |
|---|---|---|---|---|---|
| DFDC | 73.00% | 0.8041 | 74.76% | 71.13% | 73.33% |
| DFMNIST | 82.21% | 0.9144 | 89.06% | 71.25% | 83.21% |
| EXP | 70.17% | 0.8091 | 56.63% | 82.63% | 75.00% |
| Unified | 77.55% | 0.8463 | 72.76% | 82.42% | 80.80% |

Finally, the method's capability to generalize under realistic conditions is demonstrated in Fig. 1. The Receiver Operating Characteristic (ROC) curve illustrates the diagnostic trade-off between the true positive rate (Sensitivity) and the false positive rate (Specificity) [25]. As illustrated, the ensemble curve (solid black line) consistently encloses the largest area compared to individual backbones.

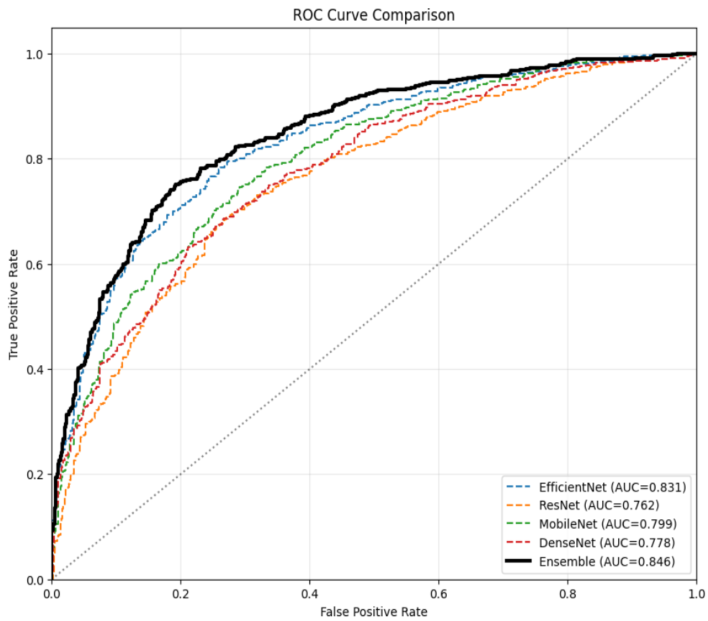


Figure 1: ROC Curve of Unified Dataset

# 6 CONCLUSION

This study presented a robust deepfake detection framework leveraging a heterogeneous ensemble of four CNN architectures, optimized via a weighted soft-voting strategy. Our experimental evaluation on both distinct and a large-scale Unified dataset demonstrated that the proposed ensemble consistently outperforms single-model baselines in terms of AUC-ROC and generalization capability. By assigning adaptive weights based on validation performance, the system effectively mitigates the variance of individual classifiers, resulting in a detector that is more stable than any single backbone.

However, while the system exhibited exceptional sensitivity against GAN-based deepfakes, such as DFDC and DFMNIST, the lower recall observed on the EXP dataset highlights a remaining challenge in identifying content made manually. This suggests that further domain adaptation is required to fully address manual forgeries [26], [27]. Ultimately, this research confirms that architectural diversity is a fundamental component for building defense mechanisms capable of evolving through digital misinformation.

Beyond its relevance to misinformation and public safety, deepfake detection is increasingly important for trusted digital marketplaces where user-generated media drives both decision-making and automated AI services. This direction aligns directly with REFASHION, a blockchain-based e-retail platform architecture that supports circular and sustainable fashion through AI services (e.g., condition assessment, material recognition, and behavioral/gamification intelligence) coupled with immutable transaction records. In such an environment, manipulated or AI-generated human-centric content (including facial imagery used for user verification, seller credibility signals, or multimedia listing enhancements) can enable impersonation and fraudulent activity, degrading trust and weakening the transparency benefits offered by blockchain. The proposed weighted-ensemble framework can therefore be positioned as a complementary REFASHION module that performs pre-ingestion authenticity screening, reducing the probability that synthetic/edited content propagates into downstream scoring, recommendation, and reward allocation processes—or becomes permanently referenced through on-chain evidence. As future work, we plan to (i) evaluate the ensemble under REFASHION-like "in-the-wild" acquisition conditions (mobile cameras, compression, varied lighting), (ii) extend detection beyond faces toward broader manipulation cues relevant to marketplace scenarios, and (iii) couple detection outputs with provenance and audit workflows, such as anchoring hashes of verified media to the blockchain to support accountability, traceability, and dispute resolution.

## ACKNOWLEDGMENTS

This work was supported by the project "Innovative Blockchain-Based e-Retail Platform for Strengthening the Circular and Sustainable Economy Model in the Fashion Industry" (Action Code: ΥΠ3ΤΑ-0560500), implemented under "SUB1.1: Partnerships of Research Excellence / Clusters of Research Excellence (CREs)" (OPS TA 5180519) within the framework of "Greece 2.0 – National Recovery and Resilience Plan."

## REFERENCES

[1] L. Gaur, *DeepFakes: Creation, Detection, and Impact*. 2022.

[2] Y. Mirsky and W. Lee, "The Creation and Detection of Deepfakes: A Survey," *ACM Comput. Surv.*, vol. 54, no. 1, pp. 1–41, Jan. 2022.

[3] M. Rana, M. Nobi, B. Murali, and A. Sung, "Deepfake Detection: A Systematic Literature Review," *IEEE Access*, vol. 10, pp. 1–1, Jan. 2022.

[4] M. El-Gayar, M. Abouhawwash, S. S. Askar, and S. Sweidan, "A novel approach for detecting deep fake videos using graph neural network," *J. Big Data*, vol. 11, Feb. 2024.

[5] R. Tolosana, R. Vera-Rodriguez, J. Fierrez, A. Morales, and J. Ortega-Garcia, *DeepFakes and Beyond: A Survey of Face Manipulation and Fake Detection*. 2020.

[6] D. S. AbdElminaam, N. Sherif, Z. Ayman, M. Mohamed, and M. Hazem, "DeepFakeDG: A Deep Learning Approach for Deep Fake Detection and Generation," *J. Comput. Commun.*, vol. 2, no. 2, pp. 31–37, July 2023.

[7] B. Zi, M. Chang, J. Chen, X. Ma, and Y.-G. Jiang, *WildDeepfake: A Challenging Real-World Dataset for Deepfake Detection*. 2021.

[8] A. Boutadjine, F. Harrag, K. Shaalan, and S. Karboua, "A Comprehensive Study on Multimedia DeepFakes," Mar. 2023.

[9] M. Ganaie, M. Hu, A. K. Malik, M. Tanveer, and P. Suganthan, "Ensemble deep learning: A review," *Eng. Appl. Artif. Intell.*, vol. 115, p. 105151, Oct. 2022.

[10] N. Bonettini, E. D. Cannas, S. Mandelli, L. Bondi, P. Bestagini, and S. Tubaro, "Video Face Manipulation Detection Through Ensemble of CNNs," Apr. 16, 2020.

[11] A. Haliassos, R. Mira, S. Petridis, and M. Pantic, *Leveraging Real Talking Faces via Self-Supervision for Robust Forgery Detection*. 2022.

[12] C. Feng, Z. Chen, and A. Owens, "Self-Supervised Video Forensics by Audio-Visual Anomaly Detection," Mar. 27, 2023.

[13] B. Dolhansky *et al.*, "The DeepFake Detection Challenge (DFDC) Dataset," Oct. 28, 2020.

[14] J. Huang, X. Wang, B. Du, P. Du, and C. Xu, "DeepFake MNIST+: A DeepFake Facial Animation Dataset," Aug. 18, 2021.

[15] "Real and Fake Face Detection." Accessed: Dec. 18, 2025. [Online]. Available: https://www.kaggle.com/datasets/ciplab/real-and-fake-face-detection.

[16] A. Paszke *et al.*, "PyTorch: An Imperative Style, High-Performance Deep Learning Library," Dec. 03, 2019.

[17] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," June 2016.

[18] M. Tan and Q. V. Le, "EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks," Sept. 11, 2020.

[19] A. Howard *et al.*, "Searching for MobileNetV3," in *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2019.

[20] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely Connected Convolutional Networks," in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, July 2017.

[21] C. Ju, A. Bibaut, and M. J. van der Laan, "The Relative Performance of Ensemble Methods with Deep Convolutional Neural Networks for Image Classification," Apr. 05, 2017.

[22] D. P. Kingma and J. Ba, "Adam: A Method for Stochastic Optimization," Jan. 30, 2017.

[23] L. Bottou, "Large-Scale Machine Learning with Stochastic Gradient Descent," *Proc COMPSTAT*, Sept. 2010.

[24] I. Loshchilov and F. Hutter, "SGDR: Stochastic Gradient Descent with Warm Restarts," May 03, 2017.

[25] T. Fawcett, "An introduction to ROC analysis," *Pattern Recognit. Lett.*, vol. 27, no. 8, pp. 861–874, June 2006.

[26] Pintelas E, Livieris IE. Convolutional neural network framework for deepfake detection: A diffusion-based approach. Computer Vision and Image Understanding. 2025 Apr 28:104375.

[27] Pintelas E, Livieris IE, Pintelas P. Quantization-based 3D-CNNs through Circular Gradual Unfreezing for DeepFake detection. IEEE Transactions on Artificial Intelligence. 2025 Apr 28.

# REFASHION: Design of an e-Retail Web Based Platform with AI Integration for Sustainable Fashion

Vassilis Tsakanikas*, v.tsakanikas@uop.gr
Emmanuel Pintelas, e.pintelas@go.uop.gr
Petros Potsikas, p.potsikas@go.uop.gr
George Darsaklis, ece19030@go.uop.gr
Vassilis Tampakas, tampakas@uop.gr
*Distributed Intelligent Systems and Data Lab, Department of Electrical and Computer Engineering, University of Peloponnese*

**ABSTRACT**

The fashion industry stands as the second most polluting industry globally, responsible for 8% of greenhouse gas emissions and 20% of global water consumption. The rise of fast fashion has accelerated the already high effects of the fashion industry, promoting overconsumption and rapid disposal of clothing. This results in an accelerating cycle of production and disposal, underscoring the imperative for a transition toward circularity and sustainable practices within the fashion industry. This paper presents REFASHION[1], an innovative e-retail platform architecture designed to enhance the circular and sustainable economy model in the fashion industry through artificial intelligence. The proposed system architecture comprises two primary subsystems: (1) a web-based trading platform enabling peer-to-peer transactions for second-hand fashion i3Q terms, and (2) an AI-integrated services layer providing intelligent automation and decision support. Key AI services include a deep learning-based image analysis service for condition assessment and material recognition, a novel Visual Product Intelligence Service that leverages reverse image search, distributed web scraping, and open-source Large Language Models (LLMs) for automated product information extraction, and a predictive analytics engine for demand forecasting and personalized recommendations. The data processing infrastructure employs Apache Spark and streaming frameworks to orchestrate real-time scraping pipelines and model inference. This paper details the system architecture, AI service integration patterns, and the innovative visual product intelligence workflow. This study analyzes the **system architecture**, integration patterns, and implementation trade-offs of our proposed framework, providing a preliminary evaluation of its impact on **manual listing friction** and scalability. Our results demonstrate how AI-driven automation can **bolster** circular economy platforms by optimizing operational efficiency and the end-user experience in second-hand marketplaces.

**CCS CONCEPTS:** Information systems, Electronic commerce, Computing methodologies, Machine learning, Computer vision, Applied computing, Environmental informatics

**Keywords:** Circular Economy, Sustainable Fashion, Artificial Intelligence, Machine Learning, e-Commerce Platform, Visual Search, Large Language Models, Web Scraping, Apache Spark, Big Data Analytics

## 1 INTRODUCTION

In the contemporary era, where environmental consciousness and the need for sustainable development occupy a primary role in the social and economic dimension of the planet, the problem of fashion waste emerges as one of the most pressing challenges. Every year, millions of tons of clothing and footwear end up as waste, causing multiple environmental and socioeconomic impacts that threaten the planet's balance and the welfare of societies [1]. Specifically, the fashion industry constitutes the second most polluting industry globally, responsible for 8% of greenhouse gas

[1] www.refashion.cyclefi.com

emissions - more emissions than the international aviation and shipping sectors combined - with an expected increase of 50% by 2030 [2]. The industry's dependence on fast fashion trends exacerbates these issues through unsustainable increases in natural resource use. The fast fashion cycle, characterized by the continuous release of new collections, promotes rapid production and consumption cycles. As a result, fashion companies now release double the number of collections compared to the pre-2000 era, encouraging a culture of disposability and overconsumption, undermining the principles of conservation and sustainable material use [3]. Per capita fiber consumption nearly tripled from 1950 to 2008, from 3.7 kg to 10.4 kg per person [4]. The successful transition to a circular economy requires the circulation of second-hand goods, their repair, reuse, or remanufacturing, with parallel public awareness for the re-utilization of items. In this context, the REFASHION project, in full alignment with national and European strategies for Circular Economy, aims precisely at accelerating circular economy actions and connecting social economy with technological innovation toward an alternative, modern, and environmentally friendly waste management model, while encouraging active citizen participation and involvement. Despite the growing adoption of digital platforms in second-hand commerce, existing systems provide limited automation for product understanding and listing generation, often requiring significant manual effort from users. Furthermore, current solutions typically treat visual search, product enrichment, and recommendation as separate components rather than as an integrated pipeline. This paper presents the architectural design of the REFASHION platform components, with particular emphasis on the trading platform and AI-integrated services that enable intelligent decision-making and user engagement optimization. The main contributions of this work include: (1) a comprehensive web-based trading platform architecture for second-hand fashion commerce; (2) novel AI services for image-based condition assessment and material recognition; (3) a Visual Product Intelligence Service combining reverse image search, distributed web scraping, and LLM-based information extraction; and (4) a scalable data processing infrastructure leveraging Apache Spark for real-time analytics.

## 2 RELATED WORK

The foundation of this research lies in the shifting paradigm of the circular economy, which has gained significant traction as a necessary response to the traditional ‘take-make-dispose’ linear model. Organizations such as the Ellen MacArthur Foundation [4] have been instrumental in advocating business models that prioritize durability, reuse, and material recovery. This shift is further reinforced by regulatory frameworks, such as the EU's Strategy for Sustainable and Circular Textiles, which envisions a 2030 market where textile products are long-lived, recyclable, and produced with high social and environmental standards. However, achieving this vision at scale requires a technological infrastructure capable of managing complex product lifecycles. Transitioning from these high-level frameworks to practical e-commerce implementation has increasingly relied on the evolution of artificial intelligence. AI has already transformed traditional e-commerce through recommendation systems and predictive analytics, but recent breakthroughs in deep learning now allow for sophisticated product categorization and pricing optimization [5]. Furthermore, the availability of Life Cycle Assessment (LCA) databases - such as Ecoinvent and the Higg Index - provides the foundational data necessary for AI-driven environmental impact calculations [6]. Yet, despite these capabilities, the specific integration of AI services into circular fashion platforms remains underexplored. The technical realization of such platforms depends heavily on two emerging fields: visual intelligence and natural language processing. Visual search technologies, powered by Convolutional Neural Networks (CNNs) and more recently by attention-based transformer architectures [18], have revolutionized product discovery by enabling content-based image retrieval [10]. Complementing this, the emergence of Large Language Models (LLMs) like LLaMA [13] and

Mistral [14] has transformed the extraction of structured data from unstructured text. In an e-commerce context, these models excel at attribute extraction and catalog enrichment [11, 12]. While these technologies demonstrate immense potential individually, there is a notable lack of research regarding their integration into a unified system that combines visual search, automated web data extraction, and AI-driven enrichment specifically for the circular economy. This technological gap serves as the primary motivation for the architecture proposed in this paper, which seeks to advance beyond current market offerings. Existing solutions, such as **Loopwardrobe**[2] and **AirRobe**[3], provide important benchmarks for this transition. Loopwardrobe.app operates as an extension for Shopify-based platforms, enabling users to maintain digital wardrobes and resell garments within a social ecosystem of conscious consumers. Similarly, AirRobe focuses on creating a "Digital Identity" for purchased items, facilitating resale, rental, or recycling throughout the item's lifecycle to reduce textile waste. While these platforms successfully prioritize sustainable consumption behaviors, REFASHION aims to extend these capabilities by automating the high-friction aspects of listing and data enrichment that currently limit the scalability of second-hand marketplaces.

## 3 SYSTEM ARCHITECTURE OVERVIEW

The proposed system is designed to support end-to-end product lifecycle management in a circular fashion marketplace. Specifically, it enables users to list, discover, and transact second-hand items while leveraging AI services to automate product understanding, enrichment, and recommendation. The REFASHION implementation follows a phased methodology designed to deliver incremental value while managing technical risk: **Phase 1 - Research and Architecture Design:** Comprehensive analysis of state-of-the-art AI and web scraping technologies, definition of system requirements, and detailed architectural specification. **Phase 2 - Core Platform Development:** Implementation of the trading platform user interfaces, transaction management workflows, and foundational API infrastructure. **Phase 3 - AI Services Integration:** Development and deployment of image analysis models, Visual Product Intelligence pipeline, and predictive analytics engines with continuous integration and automated testing. **Phase 4 - Data Infrastructure:** Implementation of Lambda architecture components, streaming pipelines, and distributed storage systems. **Phase 5 - Testing and Optimization:** Comprehensive functionality, performance, and security testing. Optimization of ML models and scraping pipelines based on real-world data. The REFASHION platform architecture follows a three-tier design pattern, comprising interconnected components that collectively enable a comprehensive circular economy marketplace for fashion items. This paper focuses on the subsystems implemented within the REFASHION platform: the Trading Platform and the AI-Integrated Services Layer, supported by a robust Data Processing Infrastructure.

### 3.1 Architectural Principles

The REFASHION architecture is grounded in a set of core design tenets intended to ensure long-term viability, performance, and user engagement. At its foundation, the system prioritizes **modularity** by employing loosely coupled services that allow individual components to evolve independently while maintaining overall system coherence. This modular nature directly facilitates robust **scalability**, as the platform leverages cloud-native infrastructure and distributed computing frameworks to seamlessly accommodate growing user traffic and transactional complexity. For these disparate services to function as a unified whole, the architecture emphasizes **interoperability** through standardized data formats and well-defined APIs that ensure seamless data exchange across all

[2] www.netkodo.com/case-studies/loopwardrobeapp
[3] www.airrobe.com/airrobe

subsystems. Building upon this structural integrity, the system adopts an **intelligence-first** philosophy, where AI services are deeply integrated into multiple architectural layers to drive personalized user experiences and facilitate sophisticated, data-driven decision-making.

### 3.2 Three-Tier Architecture Design

The system is organized into three primary layers (Figure 1), beginning with a responsive **Presentation Layer** that serves as the primary web trading platform. This layer provides role-specific interfaces for buyers, sellers, and business users, leveraging modern frameworks such as React or Angular to facilitate intuitive product discovery and transaction workflows. These user-facing components are supported by an **AI Services Layer**, which consists of a suite of intelligent modules - including image analysis, visual product intelligence, and predictive analytics - exposed via RESTful APIs. To ensure operational efficiency, these services handle requests either synchronously or asynchronously, depending on the specific computational demands of the task. At the core of the system resides the **Data Processing Layer**, which implements a Lambda architecture to unify disparate data streams. By integrating batch processing through Apache Spark with stream processing via Spark Streaming or Apache Flink, this foundation provides the high-performance environment necessary for real-time analytics and the continuous serving of machine learning models.

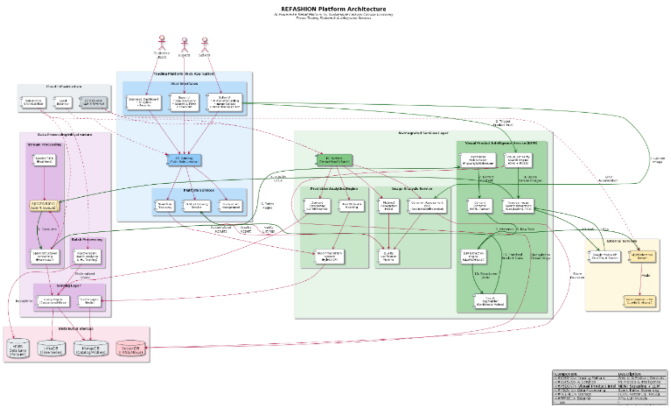


*Figure 1: REFASHION Platform and AI based functionalities.*

### 3.3 Component Integration and Data Flow

The three layers interact through a combination of synchronous API calls and asynchronous event-driven messaging. When a user uploads a product image, the Trading Platform invokes the AI Services Layer, which orchestrates multiple services: the Image Analysis Service assesses condition and materials, while the Visual Product Intelligence Service initiates web scraping pipelines to gather additional product information. Results are aggregated and returned to the user interface, with all events logged to the Data Processing Layer for analytics and model improvement.

### 3.4 Security and Data Management

Security and resilience are foundational considerations in the REFASHION system architecture. The architecture employs defense-in-depth mechanisms across all subsystems, including API authentication and authorization, encrypted data transmission, and strict separation of concerns between user-facing services and core infrastructure components [7]. Off-chain components are protected through role-based access control, continuous monitoring, and automated anomaly detection mechanisms integrated into the analytics pipeline. The architecture embeds data governance principles aligned with European regulatory frameworks, particularly the General Data Protection Regulation (GDPR). The analytics subsystem supports data lineage tracking and auditability, enabling transparent reconstruction of AI-driven recommendations and decisions. This capability ensures both regulatory compliance and methodological transparency [8].

# 4 TRADING PLATFORM

## 4.1 Platform Design and User Experience

The platform delivers a unified ecosystem with three role-specific interfaces. Buyers benefit from discovery-driven tools like map-based browsing and AI recommendations that lead into a streamlined checkout. Sellers can list items instantly using camera-assisted Visual Product Intelligence to auto-populate data, while managing the entire order lifecycle from a dedicated view. Business users leverage a centralized dashboard for high-level analytics and reporting to drive strategic decisions. By integrating these specialized functions, the system ensures a seamless, frictionless experience for every marketplace participant.

## 4.2 Product Listing and Discovery

The product listing workflow is enhanced through AI automation. When a seller uploads product images, the system automatically: (1) assesses item condition using computer vision models, (2) recognizes material composition from visual texture analysis, (3) searches the web for similar products to extract brand, model, and specification information, and (4) suggests appropriate pricing based on market analysis. This AI-assisted workflow reduces listing time from minutes to seconds while improving data quality and consistency. Product discovery leverages both traditional search mechanisms and visual similarity. Users can search by text queries with faceted filtering, or upload images to find visually similar items in the catalog. The recommendation engine personalizes search results based on user behavior history and predicted preferences.

## 4.3 Transaction Management

The platform implements a secure transaction workflow that protects both buyers and sellers. Payments are processed through integrated payment gateways with escrow-style holding until delivery confirmation. The system tracks shipment status through carrier API integrations and provides automated dispute resolution workflows. All transactions are logged with full audit trails for regulatory compliance and analytics.

# 5 AI-INTEGRATED SERVICES

A distinguishing feature of the REFASHION architecture is the deep integration of artificial intelligence services throughout the platform. These services enable intelligent automation and personalized user experiences. This section details the primary AI-integrated services and their architectural implementation.

## 5.1 Image Analysis Service

**Visual Product Intelligence:** Beyond basic image analysis, the REFASHION system employs a sophisticated visual product intelligence module to identify and categorize items within the broader fashion ecosystem. This process begins with visual search and product matching, which utilizes a transformer-based architecture to generate high-dimensional vector embeddings from product images. By comparing these embeddings against a vast database of known garments, the system can identify specific models, brands, and historical price points with high precision. To augment this internal data, the module incorporates an automated web data extraction engine. When a product is identified, the system dynamically crawls authorized e-commerce partner sites and archival fashion databases to retrieve missing metadata - such as original retail price, release year, and official colorways - ensuring that every listing is anchored in verified market data.

**AI-Driven Attribute Enrichment:** To transform these disparate visual and web-based signals into a coherent user experience, the architecture leverages the reasoning capabilities of Large Language Models (LLMs). While the visual intelligence module identifies "what" an item is, the enrichment layer focuses on "how" it is described. By processing unstructured data from web extraction and seller notes through models like LLaMA 3 or Mistral, the system performs automated attribute extraction and structured data generation. This ensures that every listing adheres to a standardized schema - categorizing items by fit, style, and occasion - while simultaneously generating SEO-optimized descriptions that highlight the item's circular value. This automated pipeline effectively removes the "data entry burden" from the seller, allowing for high-quality, searchable listings to be created in seconds rather than minutes.

### 5.2 Visual Product Intelligence Service

The Visual Product Intelligence Service represents a novel contribution of REFASHION, combining visual search, distributed web scraping, and LLM-based information extraction to automatically enrich product listings with detailed specifications and market context. Unlike existing approaches that treat visual search, web data extraction, and product enrichment as independent processes, the proposed system integrates these components into a unified pipeline, enabling automated and scalable product intelligence generation.

**Visual Similarity Search Engine:** The visual similarity search engine extracts high-dimensional feature vectors from product images using a CNN backbone (ResNet-50 or EfficientNet-B4) trained for fashion-specific representation learning [10]. These feature vectors are indexed in a vector database (FAISS or Milvus) enabling efficient approximate nearest neighbor search. When a user uploads a product image, the system generates a query embedding and retrieves visually similar products from both the internal catalog and external web sources. For external search, the system integrates with reverse image search APIs (Google Vision API, Bing Visual Search) to discover URLs of web pages containing visually similar products. These URLs serve as seeds for the subsequent web scraping pipeline. The search results are filtered and ranked based on visual similarity scores, domain reputation, and content relevance.

**Distributed Web Scraping Pipeline:** The web scraping pipeline is implemented as a distributed system orchestrated by Apache Spark, enabling parallel processing of multiple URLs while respecting rate limits and robots.txt directives. The pipeline architecture consists of several stages**: Queue Management**: Candidate URLs from the visual search engine are published to Apache Kafka topics, partitioned by domain to enable per-domain rate limiting. A Spark Streaming job consumes these URLs, applying deduplication and priority scoring based on expected information yield. **Content Fetching:** A distributed fetcher pool retrieves web pages using headless browser automation (Playwright/Selenium) for JavaScript-rendered content or lightweight HTTP clients for static pages. The fetcher implements adaptive rate limiting, respecting Crawl-Delay directives and backing off on HTTP 429 responses. Politeness policies ensure ethical scraping behavior. **Content Extraction**: Raw HTML is processed through a content extraction pipeline that identifies and extracts product-relevant content blocks. The extractor uses a combination of CSS selectors, XPath patterns, and machine learning-based content detection to isolate product descriptions, specifications, prices, and images from page chrome and advertisements.

***LLM-Based Product Information Extraction:*** Extracted web content is processed by an open-source Large Language Model to transform unstructured text into structured product attributes. The system supports multiple LLM backends including LLaMA-2 [13], Mistral [14], and their instruction-tuned variants, deployed on GPU infrastructure for efficient inference. The extraction process employs carefully designed prompts that instruct the LLM to identify and extract specific product attributes:

- Brand name and product model/style identifier

- Material composition (e.g., 100% cotton, 60% polyester/40% wool)
- Product dimensions and sizing information
- Care instructions and washing guidelines
- Original retail price and current market price estimates
- Sustainability certifications (e.g., GOTS, OEKO-TEX, Fair Trade)
- Country of manufacture and brand sustainability ratings

The LLM outputs structured JSON conforming to a predefined schema. Multiple extraction results from different web sources are aggregated using confidence-weighted voting, with conflicts resolved through source reliability scoring and consistency checking. The system implements retrieval-augmented generation (RAG) by providing the LLM with relevant context from a knowledge base of fashion terminology, brand information, and material properties [16].

***Stream Processing with Apache Spark***

The Visual Product Intelligence Service is orchestrated through Apache Spark's structured streaming capabilities [9], enabling scalable and fault-tolerant processing of the scraping and extraction pipeline: **Event-Driven Architecture:** Product image uploads trigger events published to Kafka, initiating the visual search and scraping workflow. Spark Structured Streaming jobs consume these events, maintaining exactly-once processing semantics through checkpointing. **Micro-Batch Processing:** URL fetching and content extraction are executed in micro-batches, balancing latency requirements with throughput optimization. The batch interval is configurable, typically set to 5-10 seconds for near-real-time response. **GPU-Accelerated Inference:** LLM inference is distributed across GPU workers using Spark's barrier execution mode, enabling efficient batched inference. The system supports both local GPU clusters and cloud-based inference endpoints (*e.g.*, vLLM, TensorRT-LLM) for elastic scaling. **Result Aggregation:** Extracted product attributes from multiple sources are aggregated using Spark's windowed operations, with results materialized to the serving layer for immediate availability to the trading platform.

### 5.3 Predictive Analytics Engine

The predictive engine synthesizes transaction and behavioral data to optimize market efficiency and personalization. It utilizes deep learning and time-series models to forecast demand, providing sellers with pricing guidance while identifying emerging trends for buyers. A hybrid recommendation system - combining matrix factorization with transformer-based sequence modeling - captures user intent to deliver tailored product suggestions across the platform. This logic is continuously refined through an integrated A/B testing framework that adapts to real-time engagement and shifting market conditions.

## 6 DATA PROCESSING AND ANALYTICS INFRASTRUCTURE

The data analytics infrastructure is built upon a Lambda architecture designed to balance high-throughput historical processing with low-latency, event-driven computation. Within this framework, a dedicated batch layer utilizes Apache Spark to process accumulated platform data for training machine learning models and generating long-term historical trends. This is complemented by a speed layer that addresses real-time requirements by ingesting streaming events from the trading platform; engines such as Apache Flink or Spark Streaming enable near real-time updates for the Visual Product Intelligence pipeline and recommendation services. To ensure high-performance access, the outputs of both layers are materialized in a serving layer - utilizing query-optimized technologies like Apache Cassandra or HBase-which exposes precomputed views and fast lookups to downstream applications. Supporting this computational logic is a multi-tiered storage architecture tailored to handle diverse data types and access patterns. Raw datasets, including user interactions and scraped web content, are maintained in a distributed data lake using the Apache Parquet format for efficient columnar queries. To power the platform's visual search capabilities, product image embeddings are indexed within specialized vector databases such as FAISS or Milvus, while a NoSQL layer using MongoDB provides the schema flexibility required for evolving product catalogs and user profiles. This is further augmented by Redis for high-speed caching and InfluxDB for capturing the time-series metrics essential to operational monitoring and trend analysis. The entire subsystem is deployed on a cloud-native infrastructure that leverages container orchestration via Kubernetes to provide the elasticity needed for fluctuating workloads and automated recovery. At the compute tier, GPU-enabled clusters are provisioned specifically to accelerate the intensive model training and inference required by the image analysis and LLM services. Finally, the platform exposes these capabilities through a centralized API gateway layer, which secures the environment through integrated authentication, rate limiting, and monitoring, ensuring a robust integration with the trading platform frontend.

## 7 EVALUATION AND DISCUSSION

The REFASHION platform architecture addresses key challenges in circular fashion e-commerce through intelligent automation. The Visual Product Intelligence Service significantly reduces seller effort in creating comprehensive listings - preliminary evaluations suggest reduction in listing creation time from 5-10 minutes to under 30 seconds when web sources provide relevant product information. The distributed scraping architecture enables processing of thousands of URLs per minute while maintaining ethical scraping practices. LLM-based extraction achieves high accuracy for well-structured product pages, though performance varies with page complexity and content quality. The system implements fallback mechanisms for cases where automated extraction is unreliable, ensuring consistent user experience. Key technical challenges addressed include: (1) handling JavaScript-rendered content requiring headless browser automation, (2) managing rate limits and anti-scraping measures across diverse e-commerce sites, (3) ensuring LLM extraction consistency through prompt engineering and output validation, and (4) scaling GPU inference for LLM workloads during peak usage periods. Limitations include dependency on external search APIs for visual similarity discovery, potential legal considerations around web scraping that require ongoing attention to terms of service, and the computational cost of LLM inference which impacts response latency for the enrichment pipeline.

## 8 CONCLUSION AND FUTURE WORK

This paper presented the REFASHION platform architecture focusing on the trading platform and AI-integrated services. The three-tier architecture comprising a web-based trading platform, AI services layer, and data processing infrastructure provides a robust foundation for sustainable fashion commerce. The novel Visual Product Intelligence Service - combining visual search, distributed web scraping, and LLM-based extraction - represents a significant contribution to automated product enrichment in circular economy marketplaces. Future work will focus on several directions: (1) fine-tuning domain-specific LLMs for fashion product extraction to improve accuracy and reduce inference costs, (2) implementing multimodal models that jointly process images and text for enhanced attribute prediction, (3) exploring federated learning approaches for model improvement while preserving user privacy, and (4) extending the scraping infrastructure to support real-time price monitoring and market intelligence features.

## ACKNOWLEDGMENTS

This work was supported by the project "Innovative Blockchain-Based e-Retail Platform for Strengthening the Circular and Sustainable Economy Model in the Fashion Industry" (Action Code: ΥΠ3ΤΑ-0560500), implemented under "SUB1.1: Partnerships of Research Excellence / Clusters of Research Excellence (CREs)" (OPS TA 5180519) within the framework of "Greece 2.0 - National Recovery and Resilience Plan."

**REFERENCES**

[1] Niinimäki, K. *et al.* 2020. The environmental price of fast fashion. Nature Reviews Earth & Environment. 1, 4 (Apr. 2020), 189–200.
[2] Bailey, K. *et al*. 2022. The Environmental Impacts of Fast Fashion on Water Quality: A Systematic Review. Water. 14, 7 (2022), 1073.
[3] Centobelli, P. *et al*. 2022. Slowing the fast fashion industry: An all-round perspective. Current Opinion in Green and Sustainable Chemistry. 38, (Dec. 2022).
[4] Ellen MacArthur Foundation 2017. A New Textiles Economy: Redesigning Fashion's Future.
[5] Bawack, R.E. *et al.* 2022. Artificial intelligence in E-Commerce: a bibliometric study and literature review. Electronic Markets. 32, 1 (Mar. 2022), 297–338.
[6] Wernet, G. et al. 2016. The ecoinvent database version 3 (part I): overview and methodology. The International Journal of Life Cycle Assessment. 21, 9 (Sept. 2016), 1218–1230.
[7] Zheng, Z. *et al.* 2017. An overview of blockchain technology: Architecture, consensus, and future trends. IEEE International Congress on Big Data.
[8] Politou *et al*. 2018. Backups and the right to be forgotten in the GDPR: An uneasy relationship. Computer Law & Security Review. 34, 6, 1247–1257.
[9] Zaharia, M. *et al.* 2016. Apache Spark: A unified engine for big data processing. Communications of the ACM. 59, 11, 56–65.
[10] Krizhevsky, A., Sutskever, I., and Hinton, G.E. 2012. ImageNet classification with deep convolutional neural networks. Advances in Neural Information Processing Systems. 25, 1097–1105.
[11] Liu, Z. et al. 2016. DeepFashion: Powering robust clothes recognition and retrieval with rich annotations. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR). 1096–1104.
[12] Dosovitskiy, A. *et al.* 2021. An image is worth 16x16 words: Transformers for image recognition at scale. ICLR 2021.
[13] Touvron, H. *et al.* 2023. LLaMA: Open and efficient foundation language models. arXiv preprint arXiv:2302.13971.
[14] Jiang, A.Q. *et al.* 2023. Mistral 7B. arXiv preprint arXiv:2310.06825.
[15] Wei, J. *et al.* 2022. Chain-of-thought prompting elicits reasoning in large language models. Advances in Neural Information Processing Systems. 35, 24824–24837.
[16] Lewis, P. *et al.* 2020. Retrieval-augmented generation for knowledge-intensive NLP tasks. Advances in Neural Information Processing Systems. 33, 9459–9474.
[17] Kwon, W. *et al*. 2023. Efficient memory management for large language model serving with PagedAttention. Proceedings of the 29th Symposium on Operating Systems Principles. 611–626.
[18] Johnson, J. *et al.*, H. 2019. Billion-scale similarity search with GPUs. IEEE Transactions on Big Data. 7, 3, 535–547.