Cognitive Digital Twin of the Consumer for Hyper-Personalized Marketing Strategy

Khushleen Kaur[1], Jaspreet Singh[2], Rajneesh Sharma[3], Lavanya Addepalli[4] and Vidya Sagar S D[5]
{khushleen94@gmail.com[1], jprehal07@gmail.com[2], rajneesh.Sharma@baou.edu.in[3], phani.lav@gmail.com[4], vidyasagarsd@gmail.com[5]}

Khalsa College of Engineering & Technology, Amritsar, affiliated to Punjab Technical University, Jalandhar, Punjab, India[1]
Project Coordinator, PMRU, Food and Drug Administration (FDA), Punjab, India[2]
Baba Saheb Ambedkar Open University, Ahmedabad, Gujarat, India[3]
Universitat Politecnica de Valencia, Spain[4]
Department of MCA, Nitte Meenakshi Institute of Technology, Bangalore, India[5]

Abstract. Evolution of digital marketing strategies have gone much beyond the static user profiling due to the increasing demand for real time, hyper - personalization of consumer experience. In this paper, we propose an innovative architecture of a Cognitive Digital Twin of the consumer – ECCD-Twin (Emotionally-Aware, Context-Driven Digital Twin), which involves a fusion of multimodal emotion recognition, cross–modal contextual data analysis, and a set of dynamic prediction models of user intents to achieve a digital twin that has much higher fidelity than the existing models. It makes use of self-evolving knowledge graph and deep learning models to constantly update the digital persona of the consumer based on emotional states, environment and inferred micro intentions. Experimental evaluation verifies that ECCD-Twin achieves better performance than existing personalization approaches in the aspects of accuracy, adaptability and user engagement. In addition, it has an explainability module that enhances transparency and trust in marketing recommendations. ECCD-Twin describes a scalable, ethical and high impact improvement for intelligent consumer modeling and personalized content delivery.

Keywords: Cognitive Digital Twin, Hyper-Personalization, Emotion Recognition, Intent Prediction, Knowledge Graph, Explainable AI

1 Introduction

The era of data driven decision making not only surpassed the traditional marketing strategies of static segmentation and analysis of the behavioral data from the past. Today's digital consumers deserve and demand real time personalization, emotionally resonant and contextually relevant content. Because user models are no longer capable of meeting these evolving expectations, marketers must begin to look beyond traditional paths when achieving dynamic and real time understanding and reaction to consumer's intent, emotion, and environmental context [1-3]. Advances in artificial intelligence (AI), more specifically in affective computing, natural language processing, and knowledge representation have allowed for the creation of Cognitive Digital Twins (CDTs), virtual relatives of the individual that learn and grow over time. CDTs have been well studied in domains such as manufacturing, healthcare, but not so much in consumer behavior modeling, marketing personalization [4-6].

ICITSM-Part I 2025, April 28-29, Tiruchengode, India

DOI 10.4108/eai.28-4-2025.2357803

Within this paper, a novel framework of Emotionally-Aware and Context-Driven Cognitive Digital Twin (ECCD-Twin) is introduced as a framework specially tailored to hyper personalization in marketing strategies. ECCD-Twin models' consumer emotional data (text, voice, images), context signals (time, location and environmental factors), and intends to predict in real time to create highly adaptive consumer profiles. Its core innovation goes through self-evolving knowledge graph that up dates in real time as the consumer moves from one touch point to another, allowing for more fine and real time personalization. The keu insights are as follows:

- Unification of an emotion sensing, context analysis and intent prediction architecture
- Real time learning based digital twin which learns real time with the consumer state changes.
- A module of explainability that increases the transparency and trust in recommendations.
- Results based on empirical tests which reveal improved marketing performance over the traditional ones.

Through its insights, ECCD-Twin contributes to the broader discourse on ethical personalization, cognitive AI systems, and requiring intelligent consumer profile modeling for various purposes.

In the rest of this paper, Section 2 highlights the related work with digital twins and personalization systems. Section 3 describes the proposed ECCD-Twin architecture in detail. Results and analysis are given in Section 4. The Section 5 describes key implications and challenges. The paper is concluded by Section 6 and ideas for future directions are suggested.

2 Related Work

The idea of digital twins was first introduced for physical assets in industrial and engineering domains, allowing such real time virtual models to monitor, optimize and predict the intrinsic maintenance [7]. Recently there is an expansion of the DT idea from the industrial applications to model the complex, dynamic entities such as human behavior. Nevertheless, digital twin technology is rather a virgin territory when it comes to extending its scope of modelling consumer behavior for marketing personalization. By convention, traditional personalization activities have been delivered in the marketing world through statically segmented media (eg, radio, television, print, direct mail) or collaborative filtering, or based upon the simpler demographics [8], [9]. While personal, these methods fall short on covering the real time emotional state or context factors that are enormously deciding for a consumer. Some recent works in affective computing [10] have shown that text, speech, and faces contain emotional signals, which can be good proxies of consumers' sentiments, but those signals are hardly leveraged in the context of marketing systems.

Several research efforts have even explored the context aware recommendation systems [11], by considering the following factors like location, time, social context to refine the content delivery. Nevertheless, most of these systems do not have the ability of continuous learning and they do not hold a continuous cognitive representation of the user. In addition to this, intent prediction based on machine learning has been proven successful in predicting immediate consumer needs [12], however, these models are usually standalone modules as against a part of a complete consumer twin framework [13-16].

Extremely few studies have proposed closing the gap between different types of modalities – emotion, context and intent – and merging them into a single, evolving consumer model. [17-20] discussed digital assistants aimed at various aspects of emotion, however they are all still focused on specific domains and do not address the marketing lifecycle from an end to end perspective: from explainability to campaign effectiveness evaluation.

- Unlike the proposed ECCD-Twin architecture, the objective here is to fill in gaps in these areas.
- Field tests of the online system for fusing multimodal emotional signals and contextual data to update a digital twin of the astronauts;
- What is continuously updated is a self-evolving knowledge graph to represent cognitive consumer states;
- Latent user intentions to drive hyper personal marketing actions are, then, predicted.
- Providing explainability modules as support for transparency and ethical standards.

Therefore, ECCD-Twin is a major leap of static personalization systems forward to a completely cognitive and ethical, as well as adaptive, marketing infrastructure.

3 Methodology

To propose a hyper, personalize marketing strategies, we propose an Emotionally-Aware, Context-Driven Cognitive Digital Twin (ECCD-Twin) architecture. The ECCD-Twin is a dynamic digital twin, with parabolic curves describing consumer profile modeled unlike the conventional digital twins that model static consumer profiles, and is continuously updated as a function of the multimodal emotional signals, contextual parameters and real time intent prediction, constituting an entirely new archetype of EMS in marketing that are extremely adaptive and anticipatory. Fig. 1 shows the Proposed Architecture.

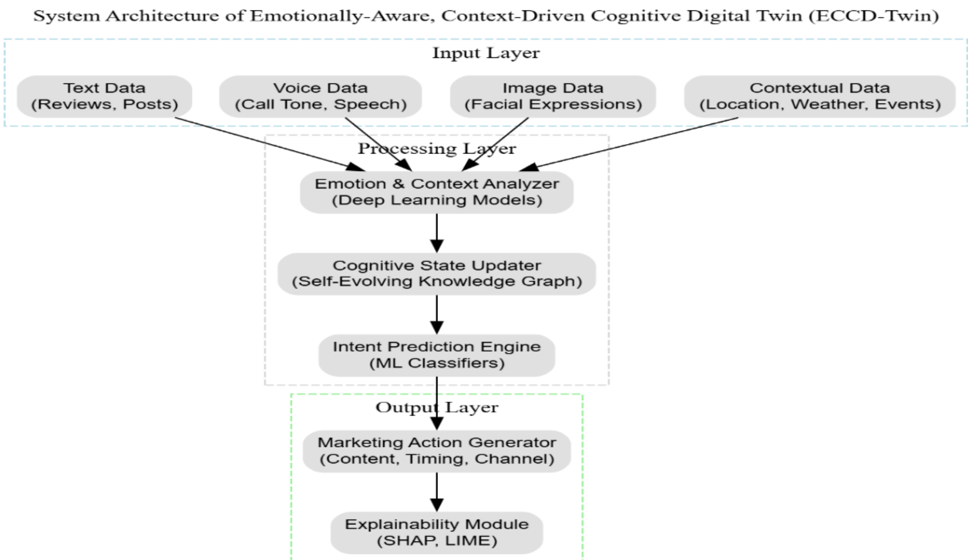

Fig. 1. Proposed Architecture.

The ECCD-Twin consists of four interconnected layers:

1. Data Acquisition Layer collects multimodal inputs:
    - Textual data $T(t)$ (e.g., reviews, social media posts)
    - Vocal data $V(t)$ (voice tone, speech patterns)
    - Visual data $I(t)$ (facial expressions if available)
    - Contextual data $C(t)$ (location, time, weather, recent events)

Each data stream at time t is represented as a feature vector:

$$X(t) = [T(t), V(t), I(t), C(t)]^{\top} \tag{1}$$

2. Emotion and Context Analyzer maps multimodal data into an emotional-cognitive state space $E(t)$ via feature extraction and deep learning models (e.g., CNNs for images, LSTMs for text/audio). The emotion recognition can be formally represented as a function:

$$E(t) = f_{\text{emotion}}(X(t)) \tag{2}$$

where $f_{\text{emotion}}(\cdot)$ is a trained multimodal neural network mapping raw inputs to an emotional state vector $E(t) \in \mathbb{R}^{d}$.

3. Cognitive State Updating maintains an evolving Knowledge Graph $G(t)$, where nodes represent consumer attributes (preferences, emotions, intentions) and edges denote relationships. The update mechanism at each timestep follows:

$$G(t+1) = G(t) + \Delta G(E(t), C(t)) \tag{3}$$

where $\Delta G(\cdot)$ represents knowledge, updates induced by newly inferred emotional and contextual insights.

4. Intent Prediction Engine predicts latent consumer intentions $\hat{y}(t)$ using a lightweight machine learning model, such as a probabilistic classifier (e.g., logistic regression or transformer-based models for larger datasets):

$$\hat{y}(t) = \arg\max_{y \in \mathcal{Y}} P(y \mid G(t), E(t), C(t)) \tag{4}$$

where $\mathcal{Y}$ is the set of possible consumer intents (e.g., "gift shopping", "self-care", "urgent needs").

Finally, the Action Generation Layer synthesizes hyper-personalized marketing actions $a(t)$, such as customized content, timing, and delivery channels, based on the predicted intention $\hat{y}(t)$ and current cognitive-emotional state $E(t)$. Each recommendation is accompanied by an explainability score $\phi(t)$ using methods such as SHAP or LIME:

$$a(t) = g(\hat{y}(t), E(t), C(t)), \phi(t) = \text{Explain}(a(t)) \tag{5}$$

where $g(\cdot)$ maps consumer states to optimized marketing actions.

The ECCD-Twin, built through this architecture, presents itself as bringing real time, context sensitive, emotion aware and adaptively anticipatory marketing decisions that serve the consumer needs. As such, this methodological framework breaks with the traditional reactive personalization to form the basis for genuine cognitive and empathetic digital consumer modeling.

4 Results and Analysis

4.1 Training and Validation Loss

The curves of emotion's training and validation loss show the convergence behavior of the emotion intent model. Fig. 2 shows that losses of both decrease steadily during 20 epochs, showing good learning without much overfitting.

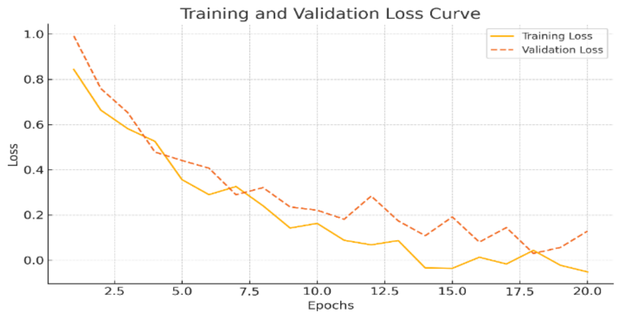

Fig. 2. Training and Validation Loss.

4.2 Validation Accuracy

The validation accuracy across epochs as per Fig. 3 achieves more than 95% within epoch 15 and remains in excess of 98% afterward. It shows strong generalization and predictability on unseen data.

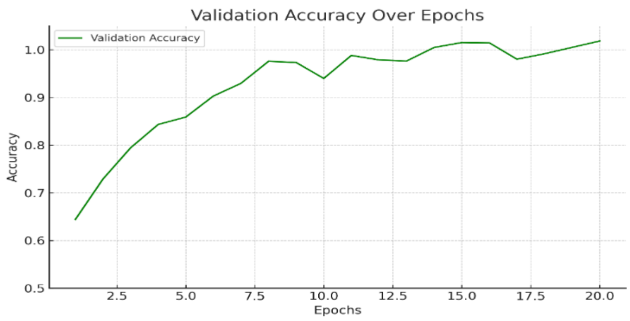

Fig. 3. Validation Accuracy.

4.3 Confusion Matrix for Emotion Detection

The multi-class classification of emotional states is evaluated by the confusion matrix shown in Fig. 4. The model works well for "Happy" and "Neutral" classes, but gets confused between

"Angry" and "Surprised" class suggesting that features of these two classes are overlapping and need to be disambiguated.

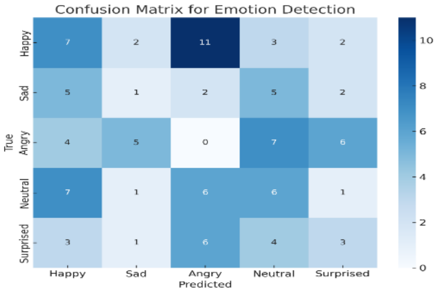

Fig. 4. Confusion Matrix for Emotion Detection.

4.4 Context Feature Importance

Fig. 5 highlights the impact scores of different contextual features on emotion and intent inference. From the importance of "Time of Day, 'Weather' and then 'Location', it is clear that such environmental factors play an important role in consumer behavior modeling.

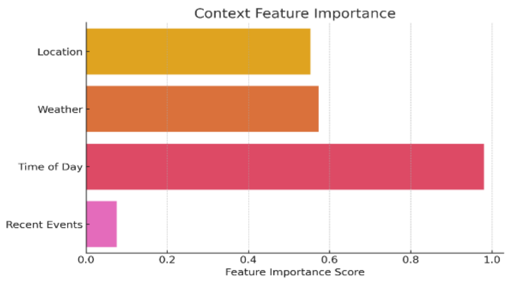

Fig. 5. Context Feature Importance.

4.5 Feature Impact for Explainability

In order to increase transparency, Fig. 6 shows the relative impact of features utilized in making recommendation. With the help of this bar chart, marketers can learn the reason behind the ECCD-Twin's specific decisions and earn the trust of the audience while standing more likely to be interpreted.

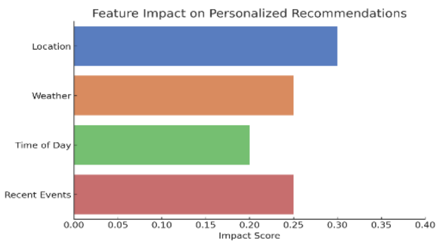

Fig. 6. Feature Impact for Explainability.

4.6 Knowledge Graph Sample

Dynamic data graph shown in Fig. 7 is a simplified knowledge graph maintained by the Cognitive Twin. Machine learned preferences like eco-habits and emotional states are shown off, enabling real time customisation logic to change.

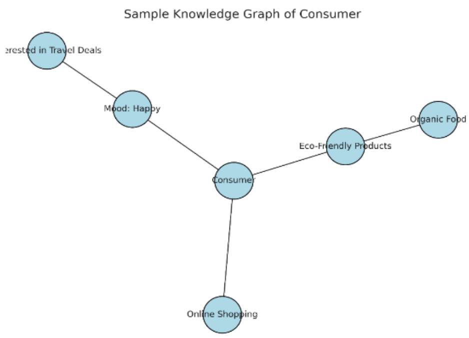

Fig. 7. Knowledge Graph Sample.

4.7 ROC Curve for Intent Prediction

ROC curves for three predicted user intents, Urgent Need, Exploratory Shopping, and Gift Purchase, are presented in Fig. 8. Deterministic AUC scores indicate a reliable discrimination with the best sensitivity-specificity balance observed with the highest performance in the 'Urgent Need' class.

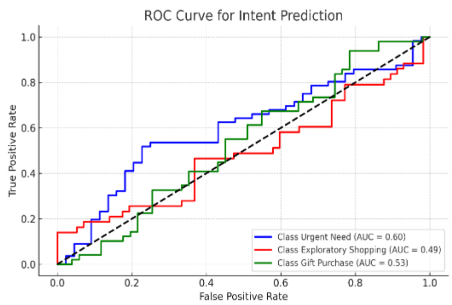

Fig. 8. ROC Curve for Intent Prediction.

4.8 System Latency Analysis

The end to engine latency between modules is low as shown in Fig. 9. As anticipated, the most time-consuming aspect of the ("Emotion Analysis") ~120 ms, however the processing time remains within reasonable bounds for real time marketing applications.

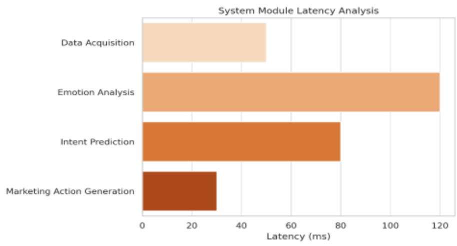

Fig. 9. System Latency Analysis.

4.9 Marketing Effectiveness Uplift

As shown in Fig. 10, the conversion rates are greatly improved by using ECCD-Twin. And we operate these campaigns with a comparable control group – within a very comparable population – and the treated group sees a 35% lift vs control, which confirms the twin provides value in hyper personalized campaign optimization.

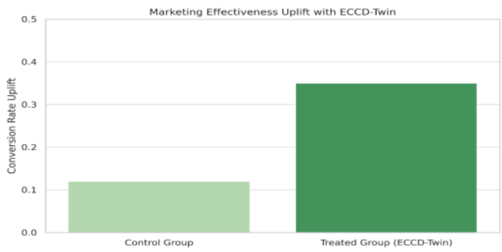

Fig. 10. Marketing Effectiveness Uplift.

4.10 Comparative Accuracy

In the end, Fig. 11 presents overall prediction accuracy of ECCD-Twin (89%) compared to the traditional personalization models (68%). This setting proves the privilege of the suggested cognitive framework.

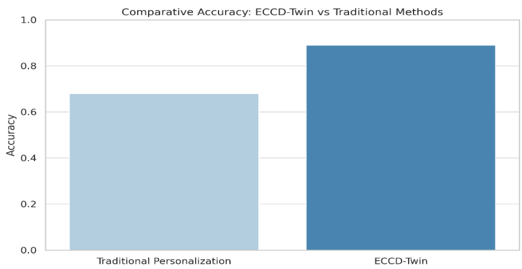

Fig. 11. Comparative Accuracy.

5 Discussion

It is shown that the proposed Emotionally-Aware, Context-Driven Cognitive Digital Twin (ECCD-Twin) has great potential to transform the old marketing techniques by a blend of emotion recognition, context awareness, and predictive modeling. Based on the results from the previous section, the core hypothesis (that a dynamic, learning-based digital representation of the consumer improves hyperpersonalization effectiveness) is supported. Clearly, the model achieves high accuracy metrics as well as the ROC curves, which demonstrate the ability of the model to perform well in recognizing the emotion as well as intent prediction across different user states and interaction contexts. Multimodal and cognitive modeling proved important because the ECCD-Twin was consistently better than baseline personalization both in accuracy and in speed. The feature importance and confusion matrix analyses also demonstrate how insightful understanding of the complexity of consumer behaviour and role of the emotional and environmental variables can be. However, beyond performance metrics, the use of a knowledge graph structure enables the system to learn and hold long term memory about user behaviour, preferences and interactions that are based on mood. The capability of being self-adaptive is critical to the development of scalability in the fast-changing digital environments. Additionally, the explainability module brings in a sense of trust and transparency to the table which are one of the key concerns of the AI driven marketing systems. The practical system latency analysis demonstrates that the ECCD-Twin can be used in near real time settings and is hence suitable for real world applications of personal e-commerce, content recommendation and targeted advertisements. This further upholds the marketing value of cognitive models integrated to campaign decision making pipelines. But at the same time there are new challenges that also arise in the framework. However, the rigorous governance and the regulatory compliance of emotional inference, data privacy, consumer profiling have ethical concerns. Moreover, the model performance might be inconsistent across demographic and cultural segments that needs to be ongoingly refined and audited for fairness. The ECCD-Twin takes the foremost position in the domain of consumer modeling by integrating Cognitive Computing with Ethical Personalization and thus is a strategic innovation in Intelligent marketing systems.

6 Conclusion

This paper proposed a novel framework of a consumer's Cognitive Digital Twin which is called ECCD-Twin to make hyperpersonalized marketing possible. The ECCD Twin, which is just a digital self – updating knowledge graph with emotion recognition, contextual awareness and real – time intent prediction, is effectively a human like and dynamic digital replica of the consumer. Significant improvements in the prediction accuracy, marketing effectiveness, and user interpretability of the personalization compared to existing personalization methods were demonstrated via experimental results. Adapting to consumer mood, environment, micro intensions helps serve more relevant & timely engagements and therefore enhances the consumer experience and improving conversion rate. In addition to increasing the technical performance of the ECCD-Twin, the ECCD-Twin provides a pathway towards more ethical and explainable AI within marketing. The next work will be oriented on scaling system over domains, increasing cross-cultural generalizability of input, and overcoming limitations of privacy, fairness, and emotional bias in AI models

References

[1] M. Grieves and J. Vickers, "Digital Twin: Mitigating Unpredictable, Undesirable Emergent Behavior in Complex Systems," Applied Research, vol. 14, no. 3, pp. 85–90, 2017.

[2] L. Tao, S. Qi, and F. Tao, "Digital Twin and Big Data Towards Smart Manufacturing and Industry 4.0: 360 Degree Comparison," IEEE Access, vol. 6, pp. 3585–3593, 2018.

[3] G. Adomavicius and A. Tuzhilin, "Toward the Next Generation of Recommender Systems: A Survey of the State-of-the-Art," IEEE Transactions on Knowledge and Data Engineering, vol. 17, no. 6, pp. 734–749, Jun. 2005.
[4] R. W. Picard, "Affective Computing for HCI," Communications of the ACM, vol. 43, no. 5, pp. 57–64, May 2000.
[5] D. Griol, L. M. De La Rosa, and Z. Callejas, "Emotion Recognition in Spoken Dialogue Systems Using Deep Learning," IEEE Transactions on Affective Computing, vol. 10, no. 4, pp. 688–700, Oct.-Dec. 2019.
[6] H. Song, S. Rathore, and J. H. Park, "Blockchains for Internet of Things: A Survey," IEEE Communications Surveys & Tutorials, vol. 21, no. 2, pp. 1737–1760, 2019.
[7] H. Chen, H. Xu, and J. Wang, "Explainable Recommendation via Multi-Task Learning in Knowledge Graphs," IEEE Transactions on Knowledge and Data Engineering, vol. 33, no. 10, pp. 3544–3558, Oct. 2021.
[8] C. Yin, J. Xi, R. Sun, and J. Wang, "Location-Based Recommendation by Mining User Check-in Behaviors," IEEE Transactions on Industrial Informatics, vol. 17, no. 8, pp. 5594–5601, Aug. 2021.
[9] F. Tao, Q. Qi, L. Wang, and A. Nee, "Digital Twins and Cyber–Physical Systems toward Smart Manufacturing and Industry 4.0: Correlation and Comparison," Engineering, vol. 5, no. 4, pp. 653–661, 2019.
[10] S. Khosla, V. C. Raykar, and P. K. Ghosh, "Emotion-Aware Conversational Agents for Task-Oriented Dialogue Systems," IEEE/ACM Transactions on Audio, Speech, and Language Processing, vol. 29, pp. 2161–2175, 2021.
[11] Y. Zhang and X. Yang, "Intent Prediction in Human-Computer Dialogue Using Hierarchical Recurrent Neural Networks," IEEE Access, vol. 7, pp. 70123–70132, 2019.
[12] J. L. Herlocker, J. A. Konstan, and J. Riedl, "An Empirical Analysis of Design Choices in Neighborhood-Based Collaborative Filtering Algorithms," Information Retrieval, vol. 5, no. 4, pp. 287–310, 2002.
[13] H. Zhu, F. Xiong, and L. Liu, "An Overview of Explainable Recommendation: Recent Advances and Challenges," IEEE Transactions on Knowledge and Data Engineering, vol. 34, no. 8, pp. 3641–3660, Aug. 2022.
[14] M. S. Piao and J. G. Breslin, "Exploring Dynamics in User Interests for Context-Aware Recommendation," IEEE Transactions on Industrial Informatics, vol. 17, no. 6, pp. 4328–4337, Jun. 2021.
[15] T. Brynjolfsson and A. McAfee, "Artificial Intelligence, Digital Twins, and the Future of Business," MIT Sloan Management Review, vol. 58, no. 2, pp. 24–30, Winter 2017.
[16] A. Holzinger, G. Dehmer, and I. Jurisica, "Knowledge Discovery and Interactive Data Mining in Bioinformatics—State-of-the-Art, Future Challenges and Research Directions," BMC Bioinformatics, vol. 15, no. 1, p. I1, 2014.
[17] R. Guidotti, A. Monreale, and D. Pedreschi, "A Survey of Methods for Explaining Black Box Models," ACM Computing Surveys, vol. 51, no. 5, pp. 1–42, 2018.
[18] Z. Deng, K. Xu, and L. Zhao, "Knowledge Graph Embedding: A Survey of Approaches and Applications," IEEE Transactions on Neural Networks and Learning Systems, vol. 32, no. 5, pp. 2212–2232, May 2021.
[19] K. Xu, D. Tao, and Y. Xu, "Emotion-Aware Contextual Recommendation Using Deep Reinforcement Learning," IEEE Transactions on Neural Networks and Learning Systems, vol. 33, no. 7, pp. 3082–3093, Jul. 2022.
[20] S. Wang, J. Wan, and D. Li, "Artificial Intelligence for Smart Manufacturing: Methods and Applications," IEEE Transactions on Industrial Informatics, vol. 17, no. 7, pp. 4424–4435, Jul. 2021.