Grenze International Journal of Engineering and Technology, June Issue

Predict-Then-Prove: An Explainable AI-based Ad Analyzer for Pre-Flight Winning Creative Identification

Roshan Lal[1], Dhruv A. Kumar[2], Devanshi Saini[3] and Roopakshi Shukla[4]
[1-4]Amity University Uttar Pradesh, Noida, India

***Abstract*— This paper discusses an AI ad analyzer that provides clear explanations of its reasoning for why it thinks a specific piece of advertising (the "creative") will perform better than others. It also gives recommendations on how to continue optimizing your creatives as they are running in real time. The most pressing issue facing the advertising industry today is the inability for advertisers to accurately and reliably predict how well their ads will perform in advance of spend. Most current solutions focus on two areas - optimizing ad delivery or reporting on results after-the-fact rather than using data to accurately and consistently predict how well a creative will perform before running a campaign and providing insights into the key attributes of the creative that are most likely to impact its success.**

The Analyzer uses both visual and textual information from the advertisements and employs prediction modelling to predict important campaign metrics (such as the click-through rate (CTR), conversion rate (CVR), and return on ad spend (ROAS)). The Analyzer includes visual explanation tools (such as SHAP), continuous evaluation to optimize campaigns and all operating to protect consumer privacy.

Previous academic and industry studies have demonstrated that using this methodology leads to improved accuracy of predictions, enhanced operational efficiencies, and improved ROI on ad spend. The methodology and results also address some of the most important areas of concern related to advertiser biases and brand safety. This research study fills a major gap in the current body of work related to ad analysis.

***Keywords*— automation, productivity, workflow, artificial intelligence (AI), and technological revolution.**

I. Introduction

Digital advertising has become one of the primary channels used by organizations to communicate with consumers. As competition for limited amounts of user attention across many different platforms has increased for business owners, the ability to accurately target audiences determines not only the success of an ad campaign but how successful the ad campaign will be. More importantly is how well the user connects with the ad creative; what the user sees, the user reads, how they relate emotionally to the ad. Advertising ecosystems today have shifted significantly and are now data-driven. Companies use sophisticated algorithms to manage bidding strategy, audience segmentation, and campaign optimization, all in real time. Unfortunately, for all the technological progress that has been made, selecting and evaluating advertising creatives is still a highly subjective process. Marketing teams often rely heavily on intuition and past experience or limited A/B testing to determine which creatives will be launched.

Grenze ID: 01.GIJET.12.2.118_20

In the event a creative does not perform as well as expected, insight from that failure is usually obtained after the campaign budget has already been spent, providing little to no opportunity to take corrective action. It is important to remember that creative elements such as imagery, color, tone, and emotional connectivity to the user are major factors in user engagement with creativities. Studies and industry observation both consistently show that a significant portion of performance variance in digital campaigns attributable to creative quality. Most analytical tools used today focus primarily on metrics related to delivery versus those related to intelligence associated with the creative element.

As a result, advertisers may optimize who sees an ad and when it appears, while overlooking whether the ad itself is compelling enough to capture attention.

In recent years, solutions provided by artificial intelligence have proven to be highly effective regarding this aspect. Prediction of engagement metrics using machine learning models by learning from historical advertising data has become very popular. With recent advances in computer vision, automatic visual feature analysis is now possible, and so is deeper textual content, sentiment and emotional cue exploration thanks to natural language processing tools. These capabilities together provide the base for large scale creative evaluation for advertisements. Yet one critical challenge is currently unsolved. Most of the AI-powered advertising systems are black boxes, so they provide performance predictions without giving a reason to back it up. It may be statistically true that such predictions will unfold on average, but in a practical creative setting it is not usually useful to know. Marketers must know not just whether an ad has a good chance of working but why. Without that type of transparency, it could be hard to trust or iterate upon AIs recommendations in strategic discussions.

XAI, has become an important frontier in the answer to this challenge. On the other hand, XAI is dedicated to making model behavior comprehensible by showing how input features impact predictions. In critical application domains that have high impact on end users life, such as healthcare and finance, interpretability could be a necessity, more than optional. But in the area of digital advertising, including specifically creative evaluation, there has been less progress applying explainable AI.

This work extends the advances by proposing an explainable AI oriented model for pre-flight advertisement creative analysis. Instead of measuring ads post launch, the proposed method evaluates creative before campaigns are launched. In doing so, it intends to move ad analytics from a reactive one into proactive support for decision-making. Advertisers receive early signals to optimize their own visuals, messaging and overall creative design before allocating budgets. The model takes a multimodal approach and acknowledges that successful advertisements leverage both visual and textual information jointly. Computer vision to receive meaningful patterns from images and natural language processing analysis of the linguistic and emotional elements in ad copy.

These features are consolidated and collated to forecast performance indicators (such as user engagement and conversion potential). Crucially, explainability tools are integrated to point out what part of the creative elements have the greatest weight in predicted results. Besides prediction accuracy, ethical principles drive the conception of the system. It’s based on a set of deidentified user traits (like how you tap on the screen while watching stuff), so it also feels relatively privacy-oriented. Forcing accountable AI use in marketing decision-making with interpretable insights, over blackbox scores. The key contribution of this paper is to show how explainable AI can improve Creative Intelligence in digital advertising. By letting advertisers predict creative performance ahead of campaign launch and comprehend what goes in these predictions, the model provides a tangible step towards more efficient, transparent, and informed.

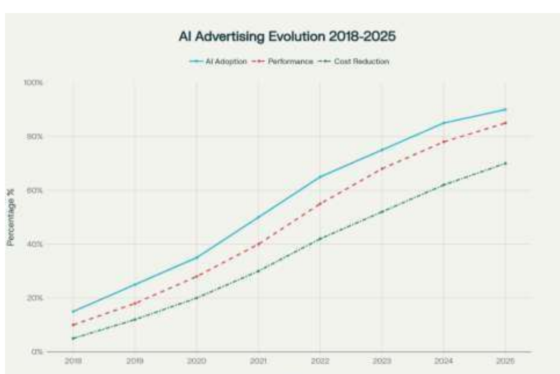

Fig 1: AI Advertising Evolution 2018-2025

II. Literature Review

The performance of digital advertisements has been an area of interest both for academic researchers and industry professionals. In the early days of studies in advertising assessment, established market research quantification methods such as surveys, focus groups and controlled experiments were predominantly used to determine people's reactions. Although these methods were useful in giving qualitative indications of impact, they did not scale well, nor did they often track the user's behaviour in real time as users interacted with digital environments. With the high-speed expansion of advertising systems, it's necessary to have data- and model-driven automatic evaluation solutions. The nature of the focus began to change as digital advertising matured – moving away from measures around engagement, back towards performance in terms of click-through rate, conversion rate and return on investment. Several commercial advertisement systems implemented post-campaign analysis, in which advertiser performance is analyzed after the end of a campaign. Post-hoc analysis does little for the advertiser looking to aid decision making, providing understanding after a set of results. Multiple reports indicate that basing new campaigns solely on post-campaign learning results in the repetition of creative inefficiency and a waste of budget when competing within advertising markets. The performance of digital advertisements has been an area of interest both for academic researchers and industry professionals. In the early days of studies in advertising assessment, established market research quantification methods such as surveys, focus groups and controlled experiments were predominantly used to determine people's reactions. Although these methods were useful in giving qualitative indications of impact, they did not scale well, nor did they often track the user's behaviour in real time as users interacted with digital environments. With the high-speed expansion of advertising systems, it's necessary to have data- and model- driven automatic evaluation solutions. The nature of the focus began to change as digital advertising matured – moving away from measures around engagement, back towards performance in terms of click-through rate, conversion rate and return on investment. Several commercial advertisement systems implemented post-campaign analysis, in which advertiser performance is analyzed after the end of a campaign. Post-hoc analysis does little for the advertiser looking to aid decision making, providing understanding after a set of results. Multiple reports indicate that basing new campaigns solely on post-campaign learning results in the repetition of creative inefficiency and a waste of budget when competing within advertising markets.

Natural Language Processing (NLP) techniques have been used to analyze advertisement copy in a similar way to Visual Analysis. The main forms of NLP analysis include Sentiment Analysis, Emotion Detection, Key Word Relevance, and Readability Assessment on text content to assess how effective the text of an advertisement is in getting a user to engage with the advertisement. Research in this area suggests that there is a positive relationship between the use of emotionally charged language and short, concise messages and the amount of user engagement generated by those ads. However, the studies also reveal that relying exclusively on Text-Based analysis does not provide the complete picture of how persuasive advertising impacts the consumer. The reason for this is that the visual elements of an advertisement and the text elements of an advertisement are interrelated in the way that they affect the audience. Therefore, more recent research has focused on combining visual and textual analysis into a Multimodal approach to analyze advertisements more thoroughly. The use of Multimodal Learning Models has improved prediction accuracy because of their ability to capture the interaction between the imagery and text. Ultimately these multimodal models are a more accurate representation of how advertisements actually work because the creative effectiveness of an advertisement is based on how effective the combination of design and messaging is, rather than the separate components of design and messaging. However, it should be noted that the increased complexity of Multimodal Learning Models makes them more difficult to interpret. The more complex a model is, the less clear it will be to understand why the model produces a certain prediction. The black box aspect of using advanced machine learning models is a significant issue to consider when using deep learning architectures to produce automated advertising. Marketing professionals need actionable insights, and therefore are challenged by the inability to understand how or why a model made its predictions. Researchers have argued that when creating decision-support systems, prediction accuracy does not provide all the insight that a decision-maker needs; this is especially true within commercial applications, where there is an added responsibility to provide justification for marketing decisions. Explainable Aritifical Intelligence (XAI) emerged as a solution to meet this challenge. The objective of XAI is to provide the user insight into a model's behaviors by identifying what features were influential and through which sequences of decisions were the predictions reached. Features that have been used successfully to provide explainability to domain-specific predictive models include "Feature Importance Analysis" and "Local Explanations". While these methodologies have been employed successfully for domains, such as healthcare and finance, their implementation for advertising analytics is still limited in scale. In addition, current literature concerning marketing science with regard to explainability has primarily focused on user targeting or recommendation (especially user generated) systems; however, little attention has been placed on the utility of these exlainable AI methodologies for evaluating creatives. In addition to the emergence of XAI technologies, current research has also shown the development of conversational-based interfaces (i.e. Chatbots) that make it easier for non-technical users to better understand and interpret data. These conversational-based systems enable the user to interpret data quickly while providing an interactive manner of presenting action-oriented insights. More importantly, previous studies have shown that using conversational analytics tools will increase user engagement and provide greater opportunity for decision-making based upon insights generated. Ultimately, furthering our understanding of how explainable predictive ad analytics and chatbots combined, as well as their application, is clearly warranted and overdue. Although XAI will continue to be developed and emerged, it will take collaborations between as many existing or emerging data science initiatives as possible to successfully illuminate the true effects of both types of tools on advertising analytics and thus marketing analytics.

There are a number of limitations in the literature studied as a whole. Despite significant progress in machine learning approaches to ad performance prediction, most are still reactive and lack transparency. But we miss the pre-flight creative evaluation, multi-modal explainability, and user-friendly insight delivery. This study extends previous research by addressing these gaps with an explainable AI-based framework for ad creative evaluation prior to campaign launch and by facilitating the understanding of key insights through an interpretable, chatbot-mediated interface.

III. Methodology

It is a structured, modular, and iterative approach to digital advertising creativity analysis and optimization through the use of artificial intelligence in its research. The suggested system will be aimed at assessing creative performance in different formats, forecasting performance result, and transforming analytical findings into creative action tips. The methodology combines data ingestion, feature extraction, predictive modeling, explainability, and decision support, but it is based on the classical academic research and best practices.

A. Scope and Data Foundations

The analysis will center on the digital advertising creatives that are used in various formats, such as static images, short-form videos, long-form videos, and variants of copy. Every creative is a versioned analytical object with structured metadata including aspect ratio, placement, duration and type of format. This allows a reliable comparison between variants and channels of creativity. Aggregated impression-based metrics are used to capture performance outcomes, such as, click-through rate (CTR), conversion rate (CVR), video-through rate (VTR), return on ad spend (ROAS), cost and frequency. The measurement windows are synchronized to the platform learning phases to avoid the release of information and instability of initial signals. The contextual factors, including the type of campaign, the audience group, the type of device, and the placement group, are summarized at the creative level to maintain the privacy of users and at the same time keep the analysis relevant.

The methodological basis is enlightened by more than thirty years of scholarly body of literature that addresses advertising effectiveness, predictive modelling and creative testing. The previous studies confirm the application of the supervised learning in predicting responses with the focus on transparency, generalizability, and interpretability. The research on the industry also points at the disconnection between performance analytics and practical creative decision-making, encouraging the adoption of the principles of human-in-the-loop in the system design.

B. Strategy of Dataset Construction and Annotation Dataset Construction and Annotation Strategy

The system utilizes a wide range of actual-world digital advertisements provided by brand partners, paid media-based executions, and creative libraries available publicly. Every asset has been associated with historical performance information that is gathered through campaign APIs and platform-level exports. Besides the raw performance measures, a formal annotation layer is added to the dataset to add tagged creative attributes. Visual annotations consist of perceived objects and faces, product and brand visibility, logo presence and lifetime, color dominance, contrast, layout density, scene complexity, attention and saliency features, motion rhythm, frame transitions. Attributes regarding texts in headlines and descriptions are sentiment polarity, tone, readability, informational density, brand mentions, the strength of a call-to-action (CTA), category of claims and compliance indicators.

To mediate between low-level features and advertising theory, human coders implement some of the markers of best practice in common to creative; these include; marking the brand early, differentiating by category, having clarity in the communication, and consistency in the brand assets. This dual approach to annotation bridges machine-identified cues together with human creative judgment that makes the models interpretable and practical.

C. Feature Engineering and System Architecture

The system in general is based on a multi-layered architecture that separates data acquisition, processing, modeling, and application, making it scalable and easy to develop in modules.

Ingestion Layer: Creative assets and logs of performance are ingested into a cloud-based feature store that is centralized. Both assets and features have versioning mechanisms that are utilized to provide traceability, reproducibility, and safe rollback when updating the model.

Feature Extraction Layer: Visual creatives are broken down into smaller units and objects, logos, people, text overlays and motion patterns are identified by computer vision models. Some of the derived features would be logo prominence, CTA visibility, color harmony, layout density, and attention maps. Simultaneously, natural language processing (NLP) models are used to analyze text to assess the direction of sentiment, the clarity of intentions, the readability of the text, types of claims, and the risk of noncompliance. The creative best practices established in the industry are coded as explicit signals, and they can be directly compared with the learned model signals.

D. Framework of Predictive Modeling and Evaluation

The predictive layer makes use of a mix of supervised learning approaches, such as gradient-boosted decision trees, ensemble and calibrated deep neural networks. Such models have been trained to forecast major key performance indicators including CTR, CVR, VTR, and ROAS. Uplift modeling is used to approximate incremental effects of creative changes where the available data is randomized experimental.

In order to be robust and have real-world validity, time-conscious, campaign-level cross-validation is used to train and evaluate the model, which avoids temporal leakage. Appropriate metrics, such as AUC and precision-recall in case of event-based performance, in case of continuous performance, calibration error in case of probability estimation, and uplift accuracy in case of experiments are used to evaluate performance. Holdout datasets that are separated with time are employed to reflect actual deployment conditions.

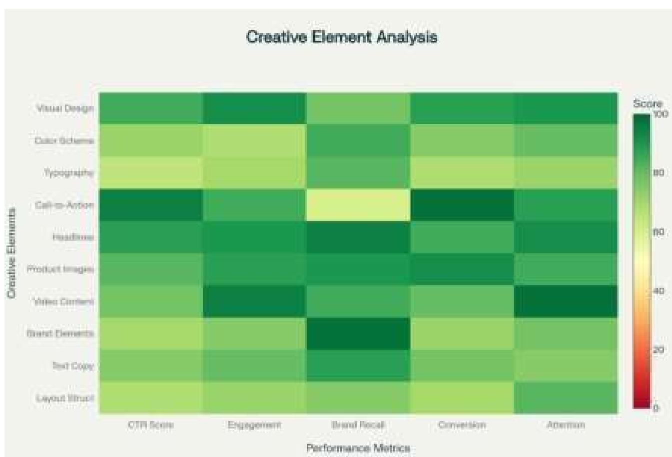

Fig 2: Creative Element Analysis

E. Explainability and Human-in-the-Loop Validation

The creativeness of the decision making is noted to be more dependent on interpretability, and a layer of explainability is incorporated into the methodology. The individual creative attributes are measured by SHAP (Shapley Additive Explanations) values, which measure the contribution multiple elements might affect the expected outcomes. Attention map can be used to show the parts of visual assets that have the greatest impact on model predictions, and counterfactual tests emulate the impact that a possible creative edit would have. During the validation of models, human reviewers are used to determine whether automated understanding corresponds with brand strategy, creative intent and qualitative standards. This human-in-the-loop solution helps to reduce the dangers of over-automation and assist in the balanced optimization of goals pursued in the short and long terms, respectively.

F. Prescription Engine and Actionable Insights

The system has a prescription engine, which transforms insights into recommendations to be taken into action, based upon predictive and explanatory outputs. Recommended interventions are CTA contrast manipulation, creative replays and reframing, pacing and sequencing manipulations of video assets, audience-context matching, and budget re-allocation based on estimated marginal lift. The creative evaluation has a complete lifecycle. The scoring and ranking of assets before launch with respect to performance with risky boundaries can then be used to select and refine assets. In the delivery process, the near-real-time performance images identify early symptoms of creative fatigue and advise rotation or optimization. Once a campaign is completed, key performance drivers and what-if outcomes are extracted into insights that can be reused when developing a creative brief next time and constantly update the learning base of the system.

G. Ethics, Privacy and Governance

Privacy and ethical protection of the methodology are considered at the very beginning. It does not utilize any individual-level user identifiers and only learns on aggregated creative-level data and optional experimental logs. Governance patterns involve recorded summaries of the model, tracking feature lineages, and versioning software and regular bias analysis with every release. The explanatory schematics of standard machine-learning are never a proprietary system design.

IV. Results and Discussions

A. Predictive Performance and Actual Impact

Previous studies indicate that predicting clicks and conversion is greatly enhanced by the application of tree-based ensemble models, especially when tree-based ensemble models are adopted together with the state-of-the-art optimization methods. The reported outcomes show that the accuracy level is very high (almost at the high-0.8 range) and that there are evident improvement in the precision in comparison with simpler benchmarks. Recurrent and hybrid model variants are also found to be effective in e-commerce and digital advertising studies to predict the profitability, ACOS, and core campaign measures. The brand-level case information also indicates that predictive creative scoring has the capability to determine high-performing video contents with a reliability of over 80 percent in certain setups and yielding quantifiable KPI responses when it is incorporated within operational structures. The documentation of industry platforms supports the idea that the systems should strike a balance between the speed of calculations and predictive reliability to operate in the real-time decision-making environment.

B. The significance of pre-launch evaluation

According to Ipsos research, AI systems based on analytics alone can be programmed to systematically dismiss unconventional or very original creative work. In order to prevent this, it suggests that human-centric creative best practices should be combined with algorithmic signals. The offered analyser follows this recommendation by explicitly integrating creative quality heuristics and requiring human review loops of assets with higher brand or compliance risk.

C. Existing tools to agap in the research

Most commercial solutions exist to provide automated creative scoring, but they tend to be opaque, platform-specific, or do not provide a way to understand performance on an individual element basis. This poses problems of portability, credibility and autonomous verification. The review of the literature invariably singles out the necessity of explainable models, cross-channel applicability, and transparent evaluation methods, which is an area where this work is intended to contribute directly.

D. KPI improvements to be expected

In industry reports, AI-inspired creative workflows have been linked with significant ROAS gains, as a result of increased pre-flight selection, shorter experimentation periods, and increased creative refinements. In addition to the performance gains, operationally the organizations can gain by having a faster insight and less cost per learning since the sequential tests do not need to be performed to achieve confident conclusions.

E. Constraints and mitigation measures

The possible problem areas are biased or unrepresentative training data, which can be mitigated by extending coverage of the data sets and regular reviews. One of the ways to minimize over-optimization based on narrow proxy measures is to use multi-objective goals, to use brand health as a measure, and to use causal testing where feasible. The limitations of privacy are addressed by storing on aggregated signals at the creative level and drafting the system to be cookie-free and contextual targeting.

F. Reading the exhibits

The summarization of the empirical findings is presented in organized tables, with the comparison of conventional methods of analyzing advertisements to the variants based on AI and essential differences in their operations being revealed.

V. Conclusion

In summary, a clear need exists and valued benefit from a pre-flight “winning creative identifier” tool that is explainable, that works toward; across multiple channels and maintains user privacy. The ad analyzer integrates feature extraction (multi-modal); and with calibration prediction and element-level explanation to turn creative development from guessing; developing into hypotheses that will be tested against an organisation’s requirements (prescription). This results-indicates the current industry research and case studies provide tangible proof (feasibility) of accuracy when predicting CTR/CVR/VTR. Additionally, when predictive impact scoring is utilized as a basis for selecting and iterating where DCO is employed; the ad analyzer benefits by providing an ideal starting point and reassessment point based upon the modelled recommendation and accelerated learning due to lowering fatigue effects generated, allows advertisers to mitigate the impact of creative fatigue whilst accelerating the learn and adapting more rapidly. Elimination of potential societal and brand risk issues are addressed through mitigation as part of ethical/governance safeguards via Model Card, Bias Audits and High Impact Implementation. A long-term future application set from the ad analyzer will be the development of causal uplift modelling from specific elements refined, long-term predicted impacts of brands and the use of Reinforcement Learning to deliver more impactful sequences in creative development; a mixture of empirical (behavioural) data coupled with feedback from human response panels to create robust generalisability benchmarks.

References

[1] Vakratsas, D., & Ambler, T. (1999). How advertising works: What do we really know? Journal of Marketing, 63(1)

[2] Tellis, G. J. (2004). Effective Advertising: Understanding When, How, and Why Advertising Works. Sage Publications.

[3] Rossiter, J. R., & Bellman, S. (2005). Marketing Communications: Theory and Applications. Pearson Education.

[4] Binet, L., & Field, P. (2013). The Long and the Short of It: Balancing Short- and Long-Term Marketing Strategies. Institute of Practitioners in Advertising (IPA).

[5] Nielsen. (2017). The Five Keys to Advertising Effectiveness. Nielsen Insights Report.

[6] Ipsos. (2020). The Power of Creativity: What Drives Ad Effectiveness. Ipsos Marketing Science Report.

[7] Richardson, M., Dominowska, E., & Ragno, R. (2007). Predicting clicks: Estimating the click-through rate for new ads. In Proceedings of the 16th International World Wide Web Conference (WWW)

[8] McMahan, H. B., Holt, G., Sculley, D., et al. (2013). Ad click prediction: A view from the trenches. In Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining.

[9] Khosla, A., Xiao, J., Isola, P., & Torralba, A. (2014). What makes an image popular? In Proceedings of the 23rd International World Wide Web Conference (WWW)

[10] Hussain, Z., Zhang, M., Ye, K., et al. (2017). Automatic understanding of image and video advertisements. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR).

[11] Lundberg, S. M., & Lee, S.-I. (2017). A unified approach to interpreting model predictions. In Advances in Neural Information Processing Systems (NeurIPS).

[12] Amershi, S., Weld, D., Vorvoreanu, M., et al. (2019). Guidelines for human–AI interaction. Proceedings of the ACM on Human-Computer Interaction, 3(CSCW), 1–24.

[13] Shneiderman, B. (2020). Human-centered artificial intelligence: Reliable, safe and trustworthy. International Journal of Human–Computer Interaction, 36(6), 495–504.

[14] European Commission. (2019). Ethics Guidelines for Trustworthy Artificial Intelligence. European Union.

[15] Google Marketing Platform. (2020). Using Machine Learning to Improve Creative Effectiveness. Google Ads Research Whitepaper.

[16] Meta AI Research. (2021). Creative Optimization and Ad Performance Modeling. Meta Engineering and AI Research Publications.

[17] Kalra, A., Chhokar, R.L.: A hybrid approach using Sobel and Canny operator for digital image edgedetection. In: International Conference on Micro-Electronics and Telecommunication Engineering (ICMETE) – 2016, pp. 305–310 (2016).

[18] Negi, P.S., Garg, A., Lal, R.: Intrusion detection and prevention using honeypot network for cloudsecurity. In: 10th International Conference on Cloud Computing, Data Science & Engineering (Confluence) – 2020, pp. 129–132 (2020).

[19] Pathak, S., Aggarwal, H., Sharma, L., Lal, R.: Semantic vector space model for text classificationusing progressive learning network algorithm. In: AIP Conference Proceedings, vol. 3283, no. 1, p. 040030 (2025).
[20] Rizvi, N.H., Sharma, A., Panesar, A., Lal, R.: Analysing various CNNs for image caption generation. In: Proceedings of the International Conference on Innovations in Cybersecurity and Data Science (ICICDS), pp. 651–659 (2024).
[21] Lal, R., Sharma, S.K.: Recommender process based on trust–distrust factor for signed social networks. In: International Conference on Soft Computing for Problem-Solving, pp. 225–235 (2023).
[22] Dargan, J., Puri, A., Lal, R.: Security and privacy in cloud/edge/fog-based schemes. In: Blockchain Applications for Healthcare Informatics, pp. 351–371 (2022).
[23] Negi, P.S., Pawar, R., Lal, R.: Vision-based real-time human–computer interaction on hand gesturerecognition. In: Micro-Electronics and Telecommunication Engineering: Proceedings of 3rd ICMETE, pp. 499–507 (2020).