International Journal of Electrical and Computer Engineering (IJECE)
Vol. 15, No. 5, October 2025, pp. 4875~4888
ISSN: 2088-8708, DOI: 10.11591/ijece.v15i5.pp4875-4888

Artificial intelligence-driven integrated system for comprehensive email marketing automation

Soumaya Loukili, Lotfi Elaachak, Abderrahim Ghadi, Abdelhadi Fennan
Data and Intelligent Systems Team FSTT, Abdelmalek Essaadi University, Tetouan, Morocco

Article Info

Article history:

Received Jan 28, 2025
Revised Jun 20, 2025
Accepted Jun 30, 2025

Keywords:

Artificial intelligence
Digital marketing
Knowledge graphs
Large language models
Stable diffusion

ABSTRACT

Right in the context of digital marketing, this paper presents a comprehensive integrated system that combines the latest artificial intelligence advancement – large language models and diffusion models – to generate marketing email subjects and content that result in higher engagement. The system uses finetuned large language models for compelling email subject generation and finetuned stable diffusion model for visually appealing and convincing email content images creation. For the latter, both knowledge graphs and vector embeddings have been incorporated to improve contextual relevance. Experimental results demonstrated significant improvement in all engagement metrics that marketers rely on, including 46% growth in open rates, 56% higher click-through rates, and an 51% boost in conversion rates, all compared to human generated content. The unified approach presented by this paper outperforms standalone models and human-generated content in terms of engagement, as the comparative analysis shows. We also discuss the ethical considerations related to content bias and personalization boundaries, alongside challenges faced in this type of projects, such as computation demands and probable solutions. Finally, this paper proposes future directions to be taken, including expansion to other digital marketing channels, the use of other advanced artificial intelligence techniques, and the development of real-time content adaptation mechanisms based on user feedback.

This is an open access article under the CC BY-SA license.

Corresponding Author:

Soumaya Loukili
Data and Intelligent Systems Team FSTT, Abdelmalek Essaadi University
Tetouan, Morocco
Email: soumaya.loukili1@etu.uae.ac.ma

1. INTRODUCTION

Over the last few years, digital marketing channels have multiplied and became numerous, all with their own agenda and target customers; through this, email marketing remains one of the most important ones to engage with customers and drive conversions. However, one of the challenges of this particular domain is that crafting subjects and visual content of said emails takes time and resources to be done, with no guarantee of them being compelling enough to drive traffic, which is the objective behind the enterprise. Human-generated content, whether in text or visual form, can at times yield lower engagement because of the potential gaps in cohesion, personalization, and alignment with targeted demographics. Therefore, as more reliable artificial intelligence (AI) tools emerge, there is a growing demand for solutions placing high importance on contextual relevance, integrating text with visuals to create an experience that will better resonate with intended audiences. This research thus seeks to design and assess a unified system that integrates the large language models (LLMs) to create email subject lines with the stable diffusion models to create images. By consolidating these state-of-the-art generative AI technologies, the system is intended to create marketing content that remains contextually relevant across different segments of audiences targeted. To this end, the present research compares the performance of a unified system against traditional marketing methods and stand-alone AI models across key engagement metrics: open rates, click-through rates (CTR), and conversion rates. Beyond that, technical challenges of deploying AI-driven marketing content at scale, as well as ethical challenges facing the use of AI in marketing, are issues this research tries to address.

The structure of this paper is as follows: section 2 is a literature review on AI in marketing and content generation, including existing research around LLMs and diffusion models. Section 3 is the proposed system architecture, describing, on a higher level, the main components and their role. Section 4, Methodology, explains the technical implementation of the unified system with regards to the integration of LLMs, stable diffusion, knowledge graphs, and vector embeddings. Section 5 describes the experimental design, including data collection, fine-tuning of the models, and their evaluation. Section 6 presents results and key findings, along with performance comparison, from real-life experiment with content generated by the system. Finally, sections 7 and 8 discuss implications for future directions and a conclusion that reflect on the theoretical contributions and ethical considerations of this system, and also ideas about how it can be further extended to other digital marketing channels.

2. LITERATURE REVIEW

2.1. AI in digital marketing

The integration of AI in digital marketing has undergone through significant changes, fundamentally changing how companies communicate with consumers. Early AI implementations were limited to rule-based systems for customer segmentation and email automation tasks. These systems, though efficient, needed more adaptability and personalization. The emergence of machine learning (ML) in the late 1990s enabled more dynamic and data-driven strategies, including targeted recommendations, exemplified by early collaborative filtering algorithms used in platforms like Amazon [1].

The advent of deep learning in the 2010s marked a significant shift in AI's capabilities, introducing models such as convolutional neural networks (CNNs) and recurrent neural networks (RNNs). These technologies powered innovations in image recognition, natural language processing (NLP), and predictive analytics—to drive applications such as chatbots and ad generation [2]. Especially, chatbots showed the power of how NLP can be in enhancing customer interactions, while predictive analytics optimized campaign strategies with regard to finding high-value customers [3].

In 2017, the introduction of transformer architectures, particularly bidirectional encoder representations from transformers (BERT), revolutionized NLP. BERT's understanding of textual context improved applications such as sentiment analysis [4] and search engine optimization [5]. Building on these advancements, large language models (LLMs) like OpenAI's GPT-3.5 and Google's PaLM 2 expanded the possibilities of AI in marketing by enabling the generation of personalized and contextually relevant text [6]. These models support creating personalized email campaigns and ad copies that increase the rate of customer engagement and click-through.

Meanwhile, generative models in AI, such as generative adversarial networks (GANs) and diffusion models, including stable diffusion, have taken visual content generation to the next level. GANs dominated this space initially [7], but then came the diffusion models, which created high-resolution images by iteratively refining noisy data [8]. With stable diffusion, for instance, marketers can generate brand-aligned visuals that will be in tune with the objectives of a campaign, leading to more engagement and stronger brand identity.

All of this progress enables personalization in marketing with AI-powered text and visual content at an unprecedented level. AI-powered recommendation engines, dynamic pricing models, and predictive analytics make campaign optimization possible and the delivery of personalized customer experiences [9]. But, notwithstanding such progress, challenges remain. Advanced AI models can be prohibitively costly for smaller businesses to implement; besides, ethical considerations abound, including algorithmic bias [10] and environmental impact [11]. Moreover, how automation can be balanced with authenticity still remains a big issue for marketers. It is important to know that even though the ever-evolving face of AI has changed digital marketing, moving from rule-based systems to complex ML, deep learning, and generative AI architectures, it is also engendering many challenges which must be resolved as these rapid advances persist.

2.2. Large language models

Large language models (LLMs) have become a fundamental element in modern NLP, primarily driven by the integration of the transformer architecture first proposed by Vaswani *et al.* [12]. Unlike recurrent and convolutional neural networks, transformers use mechanisms of self-attention to process input sequences simultaneously, which significantly increases the efficiency and scalability of language models. This led to the development of these models-pretrained on large datasets and later fine-tuned for specific applications, which is a huge turnaround in NLP research.

The introduction of BERT by Devlin *et al.* [13] demonstrated the effectiveness of bidirectional context modeling through masked language modeling (MLM). BERT attained state-of-the-art performance on a variety of NLP benchmarks, which further highlighted the benefits of pretraining on large corpora followed by task-specific fine-tuning. Its architecture has since motivated further innovations in the field and set a standard for multitask learning within NLP.

Generative models such as GPT-3, a generative pre-trained transformer, introduced by Brown *et al.* [14], extended the capabilities of LLMs by focusing on autoregressive text generation. GPT-3's massive 175-billion-parameter architecture enabled unprecedented zero-shot and few-shot learning, outperforming smaller models in a variety of generative tasks without requiring extensive task-specific data. The later GPT-3.5 introduced improved fine-grained contextual understanding and demonstrated enhanced reasoning capabilities, proving the utility of generative LLMs for conversational AI, content generation, and beyond.

More recent advancements include pathways language model 2 (PaLM 2), which was developed by Google [15]. Using multi-modal, multilingual data, the model is said to outperform the competition in tasks related to translation, scientific reasoning, and even code generation. It is generally more efficient, balancing computing needs with how well it performs. PaLM 2 reflects continued research into LLMs designed for specific fields and generic tasks. That sets the example for training methods using resources efficiently.

Despite their success, LLMs face notable challenges. Ethical concerns surrounding these models include the perpetuation of biases embedded in training data, risks of misinformation propagation, and susceptibility to malicious use. For example, Weidinger *et al.* [16] noted risks such as discrimination, misinformation, and automation harms, calling for effective ways to mitigate these problems. Besides, technical limitations, such as hallucinations, where models confidently generate incorrect outputs, and high computational costs are also serious issues that must be addressed. The environmental impact of training and deploying LLMs, especially those with billions of parameters, raises sustainability concerns, as studied by Bender *et al.* [10] and also, Strubell *et al.* [17].

LLMs such as BERT, GPT-3.5, and PaLM 2 represent the fast-paced evolution of AI in language understanding and generation. However, the more these models are put into wide use across industries, the more urgent it will be to address their ethical, technical, and societal implications. This will expand their abilities to their fullest potential.

2.3. Stable diffusion models

Diffusion models have become a key technique in generative modeling, with applications ranging from image synthesis to video generation and beyond. The main concept of diffusion models was first proposed by Sohl-Dickstein *et al.* [18], and involves iteratively transforming data into noise through a forward diffusion process and then reversing this process to recover the original data. This reverse process is parameterized by deep neural networks and trained to produce high-quality samples from noise [18]. This technique gained prominence with the development of de-noising diffusion probabilistic models (DDPMs) in a paper by Ho *et al.* [19] that demonstrated state-of-the-art performance in image generation by employing variational inference and an architecture based on U-Net.

Stable diffusion, thus, as a significant extension of DDPMs, introduced an efficient approach to latent-space diffusion. The stable diffusion model, developed by Rombach *et al.* [20], projects data into a compressed latent space using a pre-trained autoencoder, which drastically reduces computational requirements both for training and inference. This innovation not only makes diffusion models accessible for large-scale tasks but also enables real-time applications. Stable diffusion has found wide-spread adoption because it generates photorealistic and high-resolution images whose properties can be controlled, such as style and content specificity.

The incorporation of auxiliary methods such as vector embeddings and knowledge graphs has further increased the potential of stable diffusion. Knowledge graphs model structured relations among entities, which allows more contextual relevance and interpretability with generated content [21]. Knowledge graphs integrated into stable diffusion have allowed researchers to produce from such models' outputs grounded in factual or domain-specific knowledge, hence making the end result more reliable and appropriate for specific use in some domains such as pharmaceutical drug discovery and semantic parsing in NLP [22]. Vector embeddings, on the other hand, are a fundamental representation technique in machine learning, as they encode high-dimensional data into continuous vector spaces, preserving semantic relationships. This embedding enables stable diffusion models to condition their outputs on textual prompts, which is shown in systems such as Imagen, which is a Google research text-to-image diffusion model known for its high photorealism [23], and DALL-E, a model developed by OpenAI that creates images from text prompts [24].

The integration of vector embeddings and knowledge graphs can overcome two of the most significant weaknesses in diffusion models: lack of control over the content generated and the risk of irrelevant or unstructured outputs. For instance, Feng [25] showed that embedding techniques, together with knowledge-augmented diffusion models, improved the alignment of generated visuals to their textual descriptions significantly in medical imaging. Likewise, knowledge graphs within diffusion-based frameworks have also shown their potential to align better with domain-specific requirements, as recent works on recommendation systems applied to e-commerce [26] and product design [27] have pointed out.

While stable diffusion has revolutionized generative AI by combining computational efficiency and creative flexibility, challenges remain. The latent spaces are typically high-dimensional, which presents optimization challenges; the models are sensitive to adversarial prompts, which may produce undesired or misleading outputs [28]. Moreover, similar to large-scale models in general, ethical concerns persist, especially with regard to intellectual property, misinformation, and possible misuse for generating misleading content [29]. These issues, again, require further development of model robustness, transparency, and governance frameworks. Nonetheless, stable diffusion represents an important next step toward both scaling and quality in generative modeling. Its combination with vector embeddings and knowledge graphs epitomizes a growing trend toward improving the capability of models while paying much-needed attention to the practical and ethical aspects of the deployment of generative technologies.

2.4. Existing gaps

Despite the fast development of both LLMs and diffusion models, only some are the frameworks that use both these technologies to solve particular challenges, such as those related to email marketing. Although each of these models has individually achieved great success in its own domain—text generation and visual content creation, respectively—they are rarely applied together in a single workflow. This disconnection limits the exploration of how their combined potential can enhance creative and impactful marketing strategies [20].

Large language models like GPT-3.5, PaLM 2, and BERT have considerably enhanced text generation [30], producing fluent, coherent, and contextually relevant outputs for various marketing applications. At the same time, some other powerful models, including diffusion-based models like stable diffusion, have achieved similar results in generating appealing, high-quality images that cater to wide-ranging creative needs [31]. However, these are being implemented independently, without a combined system that unifies their functionalities into one solid, high-performing mechanism for marketers. That gap must be filled by allowing marketers to apply fully the various generative AI models in email marketing campaigns before the full potential of AI in email marketing will be leveraged.

Another critical gap is how little these generative frameworks incorporate contextual and domain-specific knowledge. Knowledge graphs and vector embeddings, which have proved effective for grounding outputs into domain-specific contexts, are still vastly underutilized in systems using diffusion models. For example, structured data and relationships from knowledge graphs help improve the relevance of the outputs provided [32]. At the same time, vector embeddings can enrich diffusion models by aligning generated content with themes or branding requirements [33]. Not having this contextual integration constrains the models to be effectively adapted to the domain-specific needs, which makes them less practically applicable within marketing.

Another underdeveloped area regarding multi-model systems is the evaluation of generative models. While LLMs' evaluation metrics are related to fluency, coherence, and relevance, and diffusion models' evaluation are metrics related to visual quality and style control, there needs to be a unified methodology for assessing their combined use cases. A framework that captures the overall effectiveness of outputs for specific applications, such as marketing, remains a significant research challenge [34].

Finally, ethical and operational concerns create a more complex environment for adopting these advanced models. Both large language and diffusion models pick up the biases in their training data and risk creating unintended consequences in the output. This includes, for example, generated texts from LLMs reinforcing stereotypes [35] and biased or otherwise inappropriate content with diffusion models [36]. Additionally, this high computational cost of training and deploying these models also presents challenges in sustainability, mainly for businesses aiming to deploy scalable solutions [37].

While large language and diffusion models are excellent in their respective domains, there is still a vast, untapped potential between the two as complementary technologies. Lack of integration, contextual grounding, unified evaluation frame-works, and attention to ethical and operational challenges are needed for an approach that considers everything. Enabling systems to integrate these technologies into one framework could revolutionize applications such as email marketing.

3. PROPOSED SYSTEM ARCHITECTURE

3.1. System overview

The proposed system architecture uses fine-tuned large language models, which will be responsible for generating high quality email subjects. It also employs a finetuned stable diffusion model, which will be in charge of generating compelling email visual content. This will ensure that coherent, personalized content is created and is in line with the marketing objectives and product descriptions provided by the user. The architecture follows these key stages:

3.1.1. Input data collection and preprocessing

The system collects essential input data through a form, including the product title, product description, desired email tone (chosen from multiple options), and preferred length. For image generation, a prompt is collected, detailing how the user envisions the image and the relevant product description. Text data is preprocessed—tokenized, normalized, and cleaned—ensuring consistency and accuracy before being used for generation.

3.1.2. Subject line generation via LLM module

The system uses three fine-tuned LLMs—GPT-3.5, PaLM 2, and BERT—to generate multiple subject line suggestions based on the provided product description, tone, and length preferences. Because every model offers a unique contribution, the user would have a wide range of options for subject lines. This allows flexibility in choosing that best option that would fit their marketing strategy since each model operates on the same input and molds it into its unique strengths. This multi-model approach tends to yield a more engaging subject line that best fits the targeted audience.

3.1.3. Visual content generation via stable diffusion module

Once subject lines are generated, one can proceed to create visual content - or vice versa. Here, the stable diffusion module will generate images that complement the chosen subject line in relation to the product description and user-provided prompts. For this purpose, the model also incorporates knowledge graphs and vector embeddings to make the visuals correspond to the thematic elements of the email and its tone. Quality, contextually relevant visuals are created through iterative refinements that enhance the impact of the email campaign.

3.1.4. Campaign creation and email deployment

After the generation of subject line and visual variations, the system proceeds with the creation of campaigns, where the user decides upon their preferred content. In this case, a final email package is composed, where email segments are selected based on the user's target audience. The system includes integrations of several email service providers (ESPs), such as Sendgrid, Mailgun, Mailchimp, Postmark, and Mailjet, to name a few, where deployment should take place.

It is important that the system does not comprise feedback loops in real time. However, through the analytics tools by ESPs, marketers can monitor key metrics that give them precise data about the performance. These include open rates and click-through rates which provides insight into the effectiveness of the campaign.

3.2. LLM module

The LLM module plays a vital part in this system because it generates customized and high-impact email subject lines with three fine-tuned models, namely GPT-3.5, PaLM 2, and BERT. Each has different positive features, which is a plus point the system takes advantage of. The opportunity the system offers by suggesting a variety of subject lines helps meeting the target audience’s preferences and market objectives.

3.2.1. Integration of multiple LLMs

By incorporating GPT-3.5, PaLM 2, and BERT, the system generates a range of subject line options. GPT-3.5 excels in creative, human-like outputs, PaLM 2 in conversational tones, and BERT in handling complex semantics. Together, these models ensure that the subject lines are varied and rich in context, increasing engagement potential.

3.2.2. Subject line generation process

With the input of a product title, description, tone, and subject length, each model generates a dissimilar subject. While the input data is the same, each model interprets it in its own way given their architecture. It shows the user three suggestions of a subject line, from which a user can pick the most relevant to his campaign. He could even revise the one in which he sees the most potential, taking elements from all three suggestions.

3.2.3. Model fine-tuning for marketing relevance

Each of the LLMs has been fine-tuned on email marketing-specific datasets to ensure that the generated subject lines have relevance, drive engagement, and meet best marketing practices. This will ensure the subject lines are aligned with strategies that trigger open rates, such as the use of urgency and personalization.

3.3. Stable diffusion module

The stable diffusion module generates highly engaging visual content based on product descriptions and user-provided prompts, functioning independently from subject line generation. At any point during campaign creation, users can initiate the generation of visual content. User-provided product information and the campaign's overall theme prompt the system to create images that are both aesthetically appealing and aligned with broader marketing objectives. One of the main enhancements in this module is the integration of knowledge graphs and vector embeddings. Knowledge graphs map relationships between entities within the product description, enabling the system to create images that reflect the campaign's thematic elements—whether by emphasizing unique product features or reinforcing brand identity. As for vector embeddings, they translate text-based inputs into rich numerical representations that will enable the model to create visuals capturing the subtlety of the user's prompts, including style and tone, or specific visual requirements.

Generation of images involves iteration refinement-a technique where initial noise is refined over successive steps to generate high-quality visuals. Adopting this technique ensures that the generated final images stand contextually in tune with the product description and marketing objectives. The system then generates different alternatives for the user, who can select which image best suits his campaign. It also makes sure that all visuals fit the technical requirements of the email platform in terms of file size, resolution, and aspect ratio, and the asset is ready for deployment across multiple devices and platforms.

The stable diffusion module maintains scalability and is built to handle multiple campaigns, along with large datasets, with efficiency. Knowledge graphs, vector embeddings, and the power of diffusion models assure high-quality, contextually relevant visuals. These features will enhance the overall effectiveness of the email campaign by making it more compelling to the consumer.

3.4. Integration and workflow

As shown in Figure 1, the system's architecture is built entirely within the Django framework. It is responsible for managing the flow of data between the LLM module and stable diffusion module, ensuring seamless generation and storage of content. Django handles each step of the workflow, from input collection to final campaign deployment, using Python scripts to run the generation processes and retrieve results.

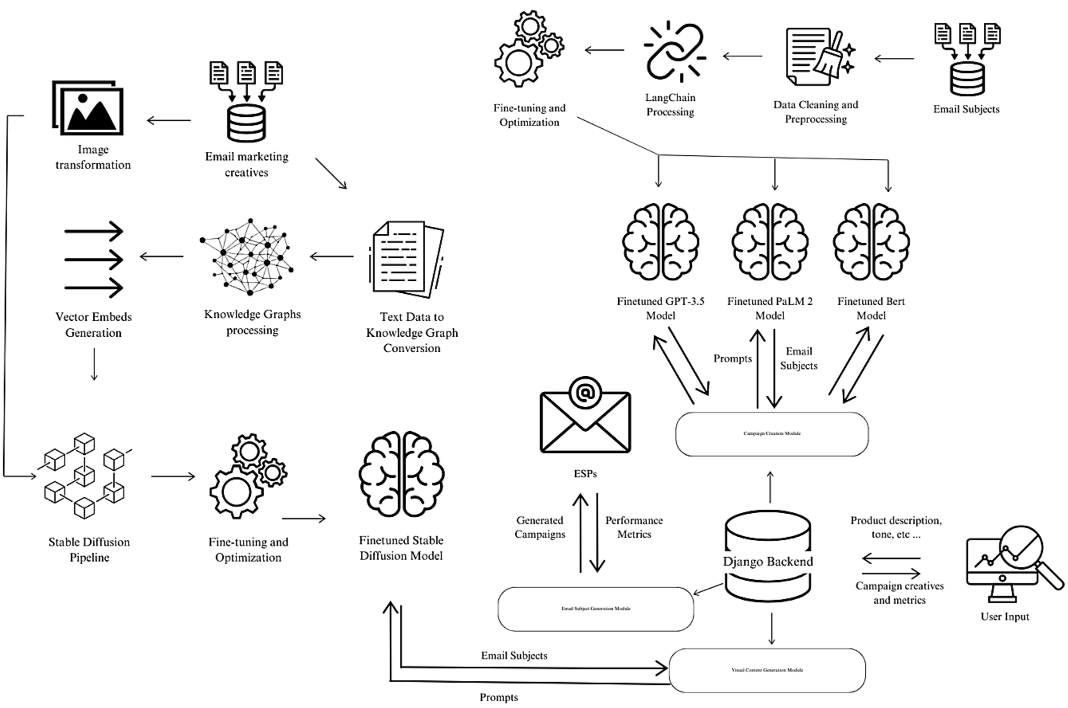

Figure 1. AI-driven system architecture

User inputs - including product descriptions, visual prompts, and tone preferences - are collected via Django forms and processed by Django's backend. Both the email subject lines and visual content are generated by Python scripts triggered through Django. These scripts interact with the LLM Module for subject line generation and the stable diffusion module for visual content creation. Results produced by these Python processes are subsequently collected and saved in a central database through Django's object-relational mapping (ORM) system.

Django's ORM efficiently handles the data flow: every subject line and image generated is correctly linked to its campaign and its product. The system gives users flexibility to create content in any order, with Django maintaining up-to-date storage and synchronization of each generated item. Such a setup will enable users to handle and access the subject lines and visuals seamlessly as they go along creating a campaign.

Django also takes care of the longer-running process scheduling and execution, like image generation, via a parallel execution tool called Celery. It is important that content generation is efficient. It is also crucial that, in the background, we maintain overall responsiveness within the system for user experience while handling many generation operations.

Once the content is generated and saved in the database, Django packages up the selected subject lines and visuals for sending. The system is integrated with several ESPs, including Sendgrid, Mailgun, and Mailchimp, that support direct sending of campaigns from Django to the target audience. Each one of the integrations with the different ESPs are already pre-set, hence seamless deployment of email campaigns is possible without needing to deal with any application programming interface (API) management from the outside.

4. METHODOLOGY

4.1. Data collection and preprocessing

This study adopts an applied experimental methodology, combining fine-tuning of LLMs and stable diffusion models with comparative evaluation based on real-world campaign data. Fine-tuning the LLMs and the stable diffusion models required highly curated datasets, that went through strict preprocessing workflows, each designed for the particular needs of each architecture. The dataset for the LLM consisted of 644,600 email subject lines sourced from a partner marketing company; this dataset included relevant metadata such as product identifiers, delivery statistics, and open rates, thus covering a broad view of the effectiveness of the campaigns. The linguistic diversity of the subject lines represented several tones and styles that constitute a strong foundation for fine-tuning models like GPT-3.5, PaLM 2, and BERT for generating high-quality and context-sensitive subject lines.

Preprocessing of the LLM dataset began with the most thorough cleaning process. In an effort to prevent redundancy, we made sure to eliminate duplicates and missing values that were not essential for further analysis. The next step was text normalization, which ensures that punctuation and capitalization are consistent unless they have se-mantic meaning, like when using words like "FREE" to convey urgency. The most important preprocessing steps in LLMs was tokenization, since long text streams had to be reduced to smaller units -words or subwords- to ensure the models could deal with the input. That way, it allowed the models to take in the syntactic structure and semantic nuances necessary in marketing language. Finally, in order to convert words to continuous vector representations, we used pre-trained embeddings -Word2Vec-. This enabled the models to understand the fine-grained relationships between words like "discount" and "promotion".

The stable diffusion model was trained on a dataset of 46,326 marketing creatives. These are images of products with metadata: for every product, there is a product description and performance indicators, such as click-through rates. This dataset aligned the visual and the textual data in a way that it allowed the model to generate contextually relevant images given textual prompts. To standardize the input dimensions, all images were resized to a 9:16 aspect ratio. Not only is this size the better fitting for email marketing, but it also retained critical visual details from different sizes. We also normalized pixel values to a mean of 0.5 and a standard deviation of 0.5 across all color channels. This step minimized biases caused by variations in lighting, color intensity, and contrast, creating a uniform dataset for training. Data augmentation techniques, such as random cropping, flipping, rotation, and color adjustments, were applied to introduce variability and improve the model's ability to generalize. These augmentations enhanced the robustness of the model by simulating diverse real-world scenarios and visual styles commonly found in marketing materials.

Preprocessing pipelines for datasets for the LLM and stable diffusion training were carefully designed so that they correspond with the respective data inputs. On one hand, LLMs rely on tokenization and embeddings in order to extract semantic depth in textual data, and on the other, stable diffusion requires normalization and augmentation techniques to maintain visual integrity. This gives all models the ability to better grasp the psychology of the user, and understand which words or writing style is most compelling to them. These extensive preprocessing steps enhanced the models' performance and also prepared a robust foundation for further training and inference operations.

4.2. Model fine-tuning

In fine-tuning for textual generation, the goal was to adapt large language models such as GPT-3.5, PaLM 2, and BERT to generate email subject lines for better-performing marketing campaigns. Such fine-tuning involved the identification of the important hyperparameters to be considered for model performance optimization. The first selected hyperparameter is the learning rate, which was set to 0.001, minimizing abrupt adjustments in the models' learning trajectory and ensuring a smooth gradual convergence. Up next, we have the batch size, selected to be 16. This modest size allowed the models to process data in manageable segments, taking into consideration computational restrictions.

In the interest of avoiding overfitting, the models were fine-tuned for three epochs—a number determined based on the dataset's complexity—thus allowing the models to learn domain-specific patterns while avoiding memorization of the training data. All three models were optimized using the Adam optimizer, which is extremely efficient in handling sparse gradients and noisy data. This optimizer improves the convergence rates of the models by adaptively changing the learning rates during training. Gradient clipping further stabilized the learning process for GPT-3.5, specifically to avoid the exploding gradients that may be bothersome in training models on longer sequences, as in the case of email subject generation.

For the evaluation of the fine-tuned models, advanced metrics such as bilingual evaluation understudy (BLEU) and recall-oriented understudy for Gisting evaluation (ROUGE) were used. BLEU was selected because of its appropriateness for measuring the precision of generated email subject lines compared with human-written reference lines. BLEU did a very good job in capturing the accuracy of the subject lines in maintaining grammatical structure and coherence. In addition, the ROUGE-LSum metric was used to judge the unigram overlap between generated text and the reference subject lines, with key phrase recall in focus. These metrics allowed for a strong output comparison of the models against human-generated content, therefore quantifying improvements in subject line fluency and relevance.

The fine-tuning for personalized marketing images with the stable diffusion model capitalized on knowledge graphs and vector embeddings to enhance the contextual relevance of generated images. It uses product metadata information, including but not limited to product names, product descriptions, and campaign performance metrics like click-through rates, to create these knowledge graphs. Knowledge graphs allowed the model to learn the relationships between marketing elements and hence enabled the generation of images that could meet the requirements of the marketing objectives. Also, vector embeddings allowed representations of textual data in a continuous vector space. These embeddings enabled the semantic richness of product descriptions to be captured by the model for better and coherent visuals.

Perceptual loss was used during fine-tuning to enhance the visual fidelity of the generated images. As opposed to traditional pixel-wise loss functions, perceptual loss compares high-level features between the generated and reference images. With the use of this approach, the model produced images that were not only realistic-looking but also semantically very accurate, making them more relevant to the marketing campaign objectives.

Besides the described methods, advanced neural network architecture was introduced to further improve the model's capability. A graph convolutional network (GCN) was used for knowledge graph processing, able to capture complicated relation-ships among entities in marketing data. This allowed the model to generate contextually enriched visuals that exactly match the associated product information. Integration of the CLIPTextModel -from OpenAI's contrastive language–image pretraining (CLIP)- was applied -a dual encoder architecture-, that allowed for alignment between textual and visual aspects of marketing creatives. It aligned latent representations of text and images, which helped CLIP enable the model to generate marketing visuals appealing and contextually correct, closely reflecting product descriptions given in the dataset

This fine-tuning process for both LLMs and the stable diffusion model demonstrated significant improvement in their generation capabilities. This is especially true for the production of personal and contextually relevant content related to email marketing campaigns. By incorporating advanced evaluation metrics, knowledge graphs, vector embeddings, and perceptual loss, high-quality outputs are generated by the models that aligned with textual and visual demands of modern marketing strategies.

4.3. System integration

The architecture relies on Django as the core framework, integrating the LLM and stable diffusion modules. Django automatically fires up background Python scripts, which generate subject lines and images. Python scripts also interface with LLMs and stable diffusion models that are responsible for the content creation, pulling it in through Django's ORM into a central database. This means good synchronization between generated content and the marketing campaign behind it.

It provides this through Celery, which integrates with resource-intensive tasks, such as generating images that can run asynchronously, keeping the system ready to take on many requests being serviced.

Celery facilitates asynchronous execution of tasks-queuing them in a message broker called RabbitMQ for background tasks like image creation to be distributed for performance and horizontal scaling. Further reduction in latency is accomplished by Redis-based caching that minimizes redundant computation by storing previously generated content. We use Django's ORM for efficient handling of database interactions and handling storing all the content generated, relating them back to their respective campaigns. The system architecture is modular by nature, hence scalable, because each component can be scaled independently to manage the increased load. At deployment, Django is integrated with ESPs like Sendgrid and Mailchimp for efficient and direct distribution of e-mail campaigns.

5. EXPERIMENTAL DESIGN AND EVALUATION

5.1. Experimental setup

The performance of the proposed system for the generation of email subject lines and visual content was evaluated using a controlled experimental framework that was designed to be fair, reliable, and relevant for realistic conditions of marketing. Such an experiment was designed to comprehensively test the capabilities of the system in creating contextually appropriate content across a wide variety of product categories and audience profiles, with strict control over variables affecting performance metrics. The aim was to compare it to content produced by humans, more specifically email marketing experts.

For this experiment, we picked 5 different products to promote, each from a distinct category. The goal was to test out the subjects and visual creatives generated by the system across several different categories, so the system’s adaptability to different domains is evaluated. Such diverse selection would make sure that the flexibility and the contextual relevance of the system could be thoroughly tested. These categories included consumer electronics, fashion, skincare, household appliances and insurance. Each product had different requirements related to marketing, including specific tone preferences and visual aesthetic considerations. For instance, consumer electronics required a professional and technical tone, while fashion products needed more creative and visually dynamic content.

For each product, the system suggested three different subject lines —one by each LLM—and many visual designs. A professional email marketer was responsible of describing the tone and impression of the campaign, and then reviewed the generated content and selected a subject line and visual creative to use. This step ensured that the content complies with marketing standards while showcasing the system's capabilities in generating high-quality, compelling outputs.

When it comes to the recipients of the promotional emails, audience segments were formed for each product based on demographic and behavioral data, as is common practice in marketing. Demographic variables included age, gender, and location, while behavioral information consisted of past engagement rates, purchasing history, and browsing activities. Such segmentation criteria allowed for the generation of very specific audience profiles to simulate real-world conditions for the most accurate marketing efforts. System-generated content was sent to these segments, mirroring the segments targeted by human-generated content, thus allowing a fair comparison.

For even more fairness, other campaign parameters were set to be the same between the system-generated and human-generated content. All emails were sent at the same time of day and over a consistent period to minimize the underlying influence brought about by timing on open and click-through rates. The email templates, delivery methods, and supplementary design elements remained identical for both test groups. Furthermore, external factors such as promotions were avoided during the testing phase to minimize their impact on results, ensuring a fair evaluation.

This experimental setup ensured a rigorous and unbiased evaluation of the system’s capabilities. By testing five different categories of products and holding other conditions constant, this experiment provided good data to determine the impact that the system has on user engagement and marketing outcomes. Furthermore, the comparison with human-generated content provided useful insights about its ability to match, if not outperform, current industry standards and set it as a potentially revolutionary tool for email marketing campaigns.

5.2. Evaluation metrics

The system’s performance effectiveness was evaluated using key marketing engagement metrics to measure the impact of AI-generated content on user interaction and campaign success.

a. Engagement metrics
   - Open rates: This metric measures the percentage of recipients who opened emails, indicating how well AI-generated subject lines capture user attention, as it is the first thing the user sees and determines whether or not he’ll continue exploring the email.
   - Click-through rates: CTR assesses the percentage of users who clicked links in the email, reflecting the engagement level with the visual content of the email.

- Conversion rates: Conversion rates track the percentage of users who complete actions like making a purchase or subscribing to a newsletter, providing an overall measure of a marketing campaign success.
- These metrics are real-life, precise numbers that allowed direct comparisons between AI and human-generated content in different marketing scenarios.

b. A/B testing

A/B testing compared AI-generated content to human-crafted alternatives, with random groups receiving either AI or human content. A t-test was conducted to determine if differences in open, click-through, and conversion rates were statistically significant. This combination of metrics and statistical validation comprehensively evaluated the system’s performance.

6. RESULTS AND DISCUSSION

6.1. Results overview

Table 1 compares the open rates obtained through human-written subject lines against those that are generated by our framework across different product categories. The data shows that AI-generated content significantly improves open rates across all product categories, with an average increase of 46.34%. Consumer electronics experienced the highest improvement (+71.35%), while Fashion had the lowest (+30.58%). These results demonstrate the effectiveness of AI in crafting engaging subject lines that capture user interest. Table 2 presents the CTR for the same campaigns, allowing a comparative analysis between human-written and AI-generated content.

AI-driven content increased click-through rates across all product categories, with an overall average improvement of 56.66%. Fashion recorded the highest gain (+64.05%), followed closely by Skincare (+60.14%). These results indicate that AI-generated visuals and subject lines create more compelling email content. We summarize the conversion rates associated with each approach in Table 3.

Conversion rates improved notably with the use of AI, with an average increase of 51.5%. Consumer electronics achieved the highest improvement (+60.6%), while Insurance saw the lowest gain (+43.7%). The data highlights AI's impact in motivating users to complete desired actions, such as making purchases.

Table 1. Open rates by product

Product category	Open rate (Manual)	Open rate (AI)	Improvement (%)
Consumer electronics	19.2%	32.9%	+71.3%
Fashion	27.8%	36.3%	+30.5%
Skincare	26.1%	37.2%	+42.5%
Household appliances	20.4%	32.4%	+58.8%
Insurance	16.7%	22.9%	+37.1%
Overall average	22.0%	32.2%	+46.3%

Table 2. Click-through rates (CTR) by product

Product category	Click-through rate (Manual)	Click-through rate (AI)	Improvement (%)
Consumer electronics	12.5%	19.7%	+57%
Fashion	15.3%	25.1%	+64.0%
Skincare	13.8%	22.1%	+60.1%
Household Appliances	14.1%	20.2%	+43.2%
Insurance	9.6%	15.2%	+58.3%
Overall Average	13.0%	20.4%	+56.6%

Table 3. Conversion rates by product

Product category	Conversion rate (Manual)	Conversion rate (AI)	Improvement (%)
Consumer electronics	6.1%	9.8%	+60.6%
Fashion	7.6%	11.5%	+51.3%
Skincare	6.2%	9.6%	+54.8%
Household appliances	5.9%	8.7%	+47.4%
Insurance	4.8%	6.9%	+43.7%
Overall average	6.1%	9.3%	+51.4%

6.2. Comparative analysis

The experimental results demonstrated significant improvements in engagement, click-through rates, and conversion rates when utilizing the integrated system, which combines LLM-generated subject lines with stable diffusion-generated visuals. Compared to traditional, manually created marketing content, the AI-driven system consistently outperformed across all key performance indicators, therefore optimizing marketing outcomes.

Open rates for email campaigns sent via the AI system saw a growth of 46.3%. This shows that LLM-generated subject lines capture user interest more, by resonating with different audience segments. More notably, consumer electronics emails improved with the highest open rates, reaching +71.3%, demonstrating AI's effectiveness in improving technical and professional content creation.

The CTR increased by 56.6% on average, showing that images designed by AI were more attractive and compelling to customers. Products from the fashion category showed the highest CTR increase at +64.06%, underlining the important role that engaging visuals play in convincing recipients to further explore the contents of an email. The results prove that AI can deliver consistent and dynamic marketing material that encourages greater user interaction.

Conversion rates, which measure the percentage of users completing a desired action, saw an average increase of 51.4%. This indicates that the personalized subject lines and engaging visuals generated by the AI system directly influenced purchasing decisions and other targeted actions. Consumer electronics campaigns saw the biggest conversion improvement at +60.6%, while Insurance campaigns improved with more modest gains of +43.7%, indicating that further refinements may be needed for more conservative domains – or perhaps those asking for personal information from users.

Apart from engagement improvements, the system showed remarkable time efficiency: content creation was 70% faster, and A/B testing processes finished 75% faster than traditional methods. This efficiency enables marketers to deploy campaigns more rapidly without compromising quality, offering a competitive edge in fast-paced markets. Finally, a breakdown of engagement metrics by content element revealed that AI-generated main visuals performed exceptionally well, with a 46% increase in click rates compared to manual visuals. Subject lines and call-to-actions (CTAs) also showed improvements of 56% and 51%, respectively, further affirming the AI's holistic impact on user engagement.

All in all, the proposed system with fine-tuned LLMs and stable diffusion models outperformed human effort in creating compelling subject lines and visual content for email marketing. It showed significant improvement in all main metrics: open, click-through, and conversion rates, while increasing efficiency, at least time-wise. These results confirm the role of generative AI as a game-changer in marketing, given that it can produce high-quality, contextually correct material.

6.3. Theoretical contributions

Knowledge graph and vector embedding integration with generative AI represents a serious theoretical leap to diffusion models. Diffusion models, as understood thus far, would create visuals based only on the text inputted into them. It is from the knowledge graphs that relationships of ideas are realized by the system, while the vector embedding translates this to forms that the models can understand. This will allow the creation of much more contextually relevant and personally targeted marketing visuals. In return, the LLMs benefit by producing linguistically correct subject lines, deeply aligned with the purposes of the campaign. Each field has its own specific details one should curate so it is impactful, and inculcating email marketing to already highly trained LLMs gives them the edge to generate marketing textual content.

While these developments enable more personalization, they do raise ethical concerns, mainly with respect to user privacy and autonomy. Too much personalization makes users feel victimized and may thus lead to trust issues. Finding the balance between offering personalization and respecting the user's privacy is crucial when it comes to training these models.

Another important risk is bias in training data. If the training data from both the LLMs and diffusion models are not diverse, then these models run the risk of reflecting unintended biases. Again, in the case of marketing, biased content alienates or misrepresents certain groups of demographics. Ensuring that our training data is diverse and introducing mechanisms for the detection of bias is crucial for ensuring the fairness of our AI-generated content. To prevent this from happening, it is important to carefully curate the training datasets so they include a broad range of demographics, cultural contexts, and behavioral patterns. To minimize such discriminatory outputs, techniques such as fairness constraints and bias detection algorithms should be integrated in the development pipeline. It is also key to set up periodic audits of generated content and get feedback from diverse user groups, so emerging biases are identified the moment they appear. It is a proactive approach that will not only mitigate risks but also enhance user trust and brand reputation in the long term. Also, beyond addressing ethical risks, this study contributes to theory by integrating LLMs and diffusion models into a unified multimodal marketing framework, bridging AI research and practical deployment.

6.4. Challenges and limitations

Some challenges and limitations indeed arose during the implementation of the integrated system. Most outstanding among these is model bias, especially within LLM content generation and in the Stable

Diffusion model. This would mean that when biases were encoded into the training data, sometimes less diverse or stereotype-reinforcing content would be generated, especially in the visuals. This involves increasing the diversity of training datasets and integrating bias detection algorithms that minimize such effects. Advanced techniques along the lines of adversarial training or fairness-aware models would potentially help decrease the bias in the generated output of future work.

Another challenge was the computational resources required to have multiple large models running in parallel. This was particularly true for image generation in stable diffusion, which being very computationally intensive, resulted in slower performance and computing at great costs. Possible remedies involve model compression through pruning or quantization, and maybe more efficient architecture, such as latent diffusion models. Future work should also consider various distributed computing and cloud scaling strategies that might enable broader access with much greater efficiency. While the integrated models indeed bring several advantages, challenges in terms of model bias and high computational costs further offer scope for improvement. Only then it will be possible to refine the system and expand its practical applications into the domains of marketing and beyond.

7. FUTURE WORK

The groundwork for the finetuned LLMs and stable diffusion model can go further than email marketing, and be extended to other digital channels, such as social media and other personalized web content. These models could adapt social media postings, targeted ads, and website banners for personal attention in all the many ways a user can interact with the system. It will open up the possibility of wider spread within the digital marketing ecosystem through better user engagement across many formats.

Embedding sophisticated AI techniques in this regard might provide better results with the system. Reinforcement learning provides opportunities for constant optimization of marketing strategies and dynamically adjusts the content depending on user feedback and engagement metrics. Further integration of generative adversarial networks will enable the generation of more diverse and creative visuals to extend the variety in marketing content and ensure capturing the user’s attention.

Of course, this requires great ethical consideration as the system expands. Federated learning can be applied at this level to ensure that artificial intelligence techniques preserve the users' data security and privacy while still fostering personalized content generation. Data protection law compliance-for example, general data protection regulation (GDPR) - will be of prime importance, and this must be set well in stone as the system further goes onto various platforms. Setting guidelines on how to detect and eliminate bias when generating content will help the system keep itself fair and unbiased.

Finally, the integration of real-time feedback loops will enable dynamic adaptation to constantly changing user interactions and market trends. This would make the system more responsive to changes in user behavior and preference by conducting real-time analysis of engagement. It would also mean running flexible campaigns where marketers can adjust content based on direct feedback and live performance insights that come in handy in fine-tuning the system for each market situation.

8. CONCLUSION

This research further proved the efficacy of integrated AI in personalized email marketing content, coming both from LLMs and diffusion models. The system overcame existing methods based on improved open rates, click-through rates, and conversion rates, confirming once more the value of integrating AI in this field, for a better generation of text and visuals. In such a way, knowledge graphs and vector embeddings can allow contextually relevant and cohesive content; this is a huge leap in generative AI for marketing. This system enables marketers and retailers to automate content creation while improving personalization, campaign efficiency, and engagement - saving both time and cost.

However, some of the limitations that were noted included model biases and computational demands. The biases in the training data carry severe risks in terms of diversity and representation for the generated content, whereas the compute-intense nature of generating images has potential scaling limitations. Future work will thus try to mitigate these issues by using more inclusive datasets and leveraging optimization techniques on models, such as pruning or lightweight architectures.

Personalization and privacy remain some of the key issues to be handled with due care in the interest of protection of customer data and over-personalization, which can destroy users' trust. In the future, the system will be extended into other digital marketing channels, including social media; more advanced AI techniques will be employed, such as reinforcement learning and GANs, which would further enhance its capabilities. Moreover, real-time feedback loops will enable dynamic content adaptation to users' interactions for the first time and make it even more responsive to market trends.

In the future, the integrated system will surely enable state-of-the-art services within the area of digital marketing and provide a powerful means for personalized content and engagement. In fact, continued research will quite precisely be very necessary in fine-tuning the system for its challenges ahead and in finding new applications within the various digital marketing channels.

REFERENCES

[1] G. Linden, B. Smith, and J. York, "Amazon.com recommendations: item-to-item collaborative filtering," *IEEE Internet Computing*, vol. 7, no. 1, pp. 76–80, Jan. 2003, doi: 10.1109/MIC.2003.1167344.

[2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436–444, May 2015, doi: 10.1038/nature14539.

[3] P. A. Okeleke, D. Ajiga, S. O. Folorunsho, and C. Ezeigweneme, "Predictive analytics for market trends using AI: A study in consumer behavior," *International Journal of Engineering Research Updates*, vol. 7, no. 1, pp. 036–049, Aug. 2024, doi: 10.53430/ijeru.2024.7.1.0032.

[4] B. Wan, P. Wu, C. K. Yeo, and G. Li, "Emotion-cognitive reasoning integrated BERT for sentiment analysis of online public opinions on emergencies," *Information Processing & Management*, vol. 61, no. 2, p. 103609, Mar. 2024, doi: 10.1016/j.ipm.2023.103609.

[5] C. Anastasiu, H. Behnke, S. Lück, V. Malesevic, A. Najmi, and J. Poveda-Panter, "Leveraging BERT to generate search engine optimized headlines," *arXiv:2107.10935*, 2021.

[6] I. Lodovico Molina, V. Švábenský, T. Minematsu, L. Chen, F. Okubo, and A. Shimada, "Comparison of large language models for generating contextually relevant questions," *arXiv:2407.2057*8, 2024.

[7] C. Abou Akar, R. Abdel Massih, A. Yaghi, J. Khalil, M. Kamradt, and A. Makhoul, "Generative adversarial network applications in industry 4.0: A review," *International Journal of Computer Vision*, vol. 132, no. 6, pp. 2195–2254, Jun. 2024, doi: 10.1007/s11263-023-01966-9.

[8] F.-A. Croitoru, V. Hondru, R. T. Ionescu, and M. Shah, "Diffusion models in vision: A survey," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 45, no. 9, pp. 10850–10869, Sep. 2023, doi: 10.1109/TPAMI.2023.3261988.

[9] M. A. AL Qudah and L. Muradkhanli, "The use of generative artificial intelligence for customer services," *Problems of Information Technology*, vol. 15, no. 1, pp. 10–17, Jan. 2024, doi: 10.25045/jpit.v15.i1.02.

[10] E. M. Bender, T. Gebru, A. McMillan-Major, and S. Shmitchell, "On the dangers of stochastic parrots," in *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, Mar. 2021, pp. 610–623, doi: 10.1145/3442188.3445922.

[11] D. Patterson *et al.*, "Carbon emissions and large neural network training," *arXiv:2104.10350*, 2021.

[12] A. Vaswani *et al.*, "Attention is all you need," *Advances in neural information processing systems*, 2017.

[13] J. Devlin, M.-W. Chang, K. Lee, K. T. Google, and A. I. Language, "BERT: pre-training of deep bidirectional transformers for language understanding," *Naacl-Hlt 2019*, no. Mlm, pp. 4171–4186, 2018.

[14] T. B. Brown *et al.*, "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.

[15] A. Chowdhery, E. Reif, D. Ippolito, D. Luan, H. Lim, and B. Zoph, "PaLM: scaling language modeling with pathways," *Journal of Machine Learning Research*, vol. 24, pp. 1–113, 2023.

[16] L. Weidinger *et al.*, "Taxonomy of risks posed by language models," in *2022 ACM Conference on Fairness Accountability and Transparency*, Jun. 2022, pp. 214–229, doi: 10.1145/3531146.3533088.

[17] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for modern deep learning research," *AAAI 2020 - 34th AAAI Conference on Artificial Intelligence*, pp. 1393–13696, 2020, doi: 10.1609/aaai.v34i09.7123.

[18] J. Sohl-Dickstein, E. A. Weiss, N. Maheswaranathan, and S. Ganguli, "Deep unsupervised learning using nonequilibrium thermodynamics," in *32nd International Conference on Machine Learning, ICML 2015*, vol. 3, pp. 2246–2255, 2015.

[19] J. Ho, A. Jain, and P. Abbeel, "Denoising diffusion probabilistic models," *Advances in Neural Information Processing Systems*, vol. 2020-December, 2020.

[20] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, "High-Resolution Image Synthesis with Latent Diffusion Models," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2022, vol. 2022-June, pp. 10674–10685, doi: 10.1109/CVPR52688.2022.01042.

[21] A. Hogan *et al.*, "Knowledge graphs," *ACM Computing Surveys*, vol. 54, no. 4, pp. 1–37, May 2022, doi: 10.1145/3447772.

[22] C. Liu *et al.*, "Generative diffusion models on graphs: methods and applications," in *IJCAI International Joint Conference on Artificial Intelligence*, pp. 6702–6711, 2023, doi: 10.24963/ijcai.2023/751.

[23] S. Jamal, H. Wimmer, and C. M. Rebman Jr, "Perception and evaluation of text-to-image generative AI models: a comparative study of DALL-E, Google Imagen, GROK, and stable diffusion," *Issues In Information Systems*, 2024, doi: 10.48009/2_iis_2024_123.

[24] A. Ramesh *et al.*, "Zero-shot text-to-image generation," in *Proceedings of Machine Learning Research*, vol. 139, pp. 8821–8831, 2021.

[25] C. M. Feng, "Enhancing label-efficient medical image segmentation with text-guided diffusion models," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 15008 LNCS, pp. 253–262, 2024, doi: 10.1007/978-3-031-72111-3_24.

[26] Y. Jiang, Y. Yang, L. Xia, and C. Huang, "DiffKG: Knowledge graph diffusion model for recommendation," in *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, Mar. 2024, pp. 313–321, doi: 10.1145/3616855.3635850.

[27] Z. Wang, C. Fu, S. Niu, and S. Hu, "Research on Bio-inspired product design based on knowledge graph and semantic fusion diffusion model," *Advanced Engineering Informatics*, vol. 62, p. 102797, Oct. 2024, doi: 10.1016/j.aei.2024.102797.

[28] M. Chen, S. Mei, J. Fan, and M. Wang, "Opportunities and challenges of diffusion models for generative AI," *National Science Review*, vol. 11, no. 12, Nov. 2024, doi: 10.1093/nsr/nwae348.

[29] S. Peng, Y. Chen, C. Wang, and X. Jia, "Intellectual property protection of diffusion models via the watermark diffusion process," *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, vol. 15437, pp. 290–305, 2025, doi: 10.1007/978-981-96-0567-5_21.

[30] F. Sahinuç, I. Kuznetsov, Y. Hou, and I. Gurevych, "Systematic task exploration with LLMs: A study in citation text generation," in *Proceedings of the Annual Meeting of the Association for Computational Linguistics*, 2024, vol. 1, pp. 4832–4855, doi: 10.18653/v1/2024.acl-long.265.

[31] C. Zhang *et al.*, "Taming stable diffusion for text to 360° panorama image generation," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2024, pp. 6347–6357, doi: 10.1109/CVPR52733.2024.00607.

[32] S. Schramm, C. Wehner, and U. Schmid, "Comprehensible artificial intelligence on knowledge graphs: A survey," *Journal of Web Semantics*, vol. 79, p. 100806, Dec. 2023, doi: 10.1016/j.websem.2023.100806.
[33] S. Gu *et al.*, "Vector quantized diffusion model for text-to-image synthesis," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2022, vol. 2022-June, pp. 10686–10696, doi: 10.1109/CVPR52688.2022.01043.
[34] J. Burden, "Evaluating AI evaluation: perils and prospects," *arXiv:2407.09221*, 2024.
[35] S. Dai, C. Xu, S. Xu, L. Pang, Z. Dong, and J. Xu, "Bias and unfairness in information retrieval systems: new challenges in the LLM era," in *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, Aug. 2024, pp. 6437–6447, doi: 10.1145/3637528.3671458.
[36] Y. Qu, X. Shen, X. He, M. Backes, S. Zannettou, and Y. Zhang, "Unsafe diffusion: on the generation of unsafe images and hateful memes from text-to-image models," in *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, Nov. 2023, pp. 3403–3417, doi: 10.1145/3576915.3616679.
[37] M. Xu *et al.*, "A survey of resource-efficient LLM and multimodal foundation models," *arXiv:2401.08092*, 2024.

BIOGRAPHIES OF AUTHORS

Soumaya Loukili is a dedicated Ph.D. student, passionate about the synergy between marketing, artificial neural networks, and artificial intelligence. Her research explores innovative approaches in these fields, showcasing a commitment to staying at the forefront of cutting-edge research. Soumaya's work reflects a fresh perspective on marketing, leveraging the power of artificial neural networks and AI for insightful strategies. She also has a keen interest in sentiment analysis and did multiple studies in this field, intersecting with artificial intelligence as well. She can be contacted via email at: soumaya.loukili1@etu.uae.ac.ma.

Lotfi Elaachak is an associate professor at Faculté des Sciences et Techniques de Tanger. With a strong focus on artificial intelligence, he has a profound understanding of algorithms and artificial neural networks. His expertise extends to higher education, where he actively engages in innovative teaching methods. Proficient in machine learning, Lotfi has made significant contributions to the field, particularly in the realms of neural networks and artificial intelligence. Beyond academia, he is involved in the realm of serious games, video games, game design, and game development, showcasing a diverse skill set. Lotfi is passionate about incorporating these skills into teaching methods, emphasizing the integration of digital games, game-based learning, and new technologies like virtual reality and augmented reality. His proficiency also extends to Arabic natural language processing (NLP), reflecting his commitment to advancing research and education in these dynamic domains. He can be contacted via email at: lelaachak@uae.ac.ma.

Abderrahim Ghadi is a professor at Université Abdelmalek Essaadi in Tetouan, Morocco. His research focuses on computer science, with particular emphasis on artificial intelligence, graph theory, information systems security, and complex networks. Ghadi's expertise extends to areas such as smart cities, social media analysis, and big data. He has contributed to various academic publications, including work on data-driven marketing and machine learning applications. With a Maîtrise-Master-Doctorat degree, Ghadi plays a significant role in advancing research and education in his field. His collaborative efforts with colleagues like Lotfi Elaachak have resulted in impactful studies on topics such as neural network algorithms for cloud computing CPU scheduling. He can be contacted at email: aghadi@uae.ac.ma.

Abdelhadi Fennan is a professor in the Department of Computer Science at Abdelmalek Essaadi University, Tangier, Morocco. He is a distinguished academic with a focus on artificial intelligence, software engineering, and information systems (business informatics). His expertise extends to disciplines such as higher education, educational technology, and curriculum theory. Abdelhadi's skills encompass e-learning, online learning, semantic web, sentiment analysis, and technology-enhanced learning. With a strong commitment to research and academia, he plays a vital role in advancing knowledge and innovation in these domains. He can be contacted via email at: afennan@uae.ac.ma.