



NAVIGATING THE FUTURE OF ADS: LARGE LANGUAGE MODELS AND THEIR IMPACT ON PROGRAMMATIC STRATEGIES

JAVIER CAMPOS is a visionary Chief Information Officer at Fenestra, steering the company's Product and Technology divisions with a focus on pioneering AI advancements in programmatic management. With a wealth of 28+ years in the global arena, his expertise spans finance, AI, market research, media, and technology. Notably, Javier has previously enhanced data innovation as Head of Experian DataLabs for UK&I and EMEA. His thought leadership in AI was further recognised through his permanent seat at the Bank of England & FCA Artificial Intelligence Public-Private Forum, which aims to expedite AI integration in the financial sector. An esteemed author, Javier released *How to Grow Your Business with AI* in 2023, adding depth to his role as a dynamic speaker at industry events. His career includes pivotal positions such as Global Chief Technology Officer at Kantar-WPP, Havas Media, and EMEA CIO at GroupM-WPP, marking him as a vanguard in his field.

THIS ARTICLE WILL DELVE INTO THE UNDERPINNINGS OF LARGE LANGUAGE MODELS (LLMS) THAT ARE RESHAPING MANY INDUSTRIES, INCLUDING PROGRAMMATIC ADVERTISING.

WE WILL EXPLORE THE NUANCES OF FOUNDATIONAL MODELS, DEBATE THE MERITS OF OPEN SOURCE VERSUS PROPRIETARY TECHNOLOGIES, AND DISCUSS STRATEGIC APPROACHES SUCH AS FULL TRAINING VERSUS FINE-TUNING, AS WELL AS THE USE OF RETRIEVAL-AUGMENTED GENERATION (RAG) FOR ENHANCING AD RELEVANCE AND PERSONALISATION.

PRACTICAL GUIDELINES FOR DATA SCIENTISTS IN THE ADVERTISING DOMAIN WILL BE PROVIDED, WITH A LENS ON ETHICAL IMPLICATIONS AND FUTURE DIRECTIONS.

FENESTRA 

Large language models (LLMs) are reshaping the landscape across various industries with their profound ability to mimic human language, bringing a new wave of innovation and efficiency. In the realm of advertising and programmatic media buying, this transformation is particularly pronounced. LLMs are revolutionising the field by enabling more nuanced content generation and optimisation strategies. Their capacity to accurately predict, understand, and emulate human-like text and interactions is changing the way advertising content is conceptualised, created, and delivered.

This article will explore the core technologies behind these models, understand their capabilities in the context of industry-specific applications, and discuss the strategic implications of integrating these advanced AI tools into business workflows.

As the field rapidly evolves, the choices between open source and proprietary models, as well as between full training, fine-tuning, and leveraging advanced techniques like retrieval-augmented generation (RAG), present both opportunities and challenges for data scientists. Through an exploration of these pivotal decisions, this article will provide data scientists with a roadmap to navigate the intricacies of LLMs in programmatic advertising, underscored by an understanding of ethical practices and an anticipation of future AI trends.

Programmatic advertising represents the automated buying and selling of online advertising space, where software algorithms determine the placement and price of ads in real time. This process uses data analytics and machine learning to deliver personalised advertising content to users across various digital platforms, such as websites, social media, and Connected TV (CTV).

LLMs are uniquely positioned to address several prevailing challenges in the programmatic advertising industry, thanks to their advanced capabilities in language understanding and generation. The impending obsolescence of third-party cookies, for instance, threatens the data-driven nature of targeted advertising. LLMs, with their sophisticated data processing and generation abilities, offer an alternative by enabling the creation of highly personalised ad content without relying heavily on third-party data.

LLMs can analyse user interactions and content engagement to generate targeted advertising that aligns with user interests and behaviours. This approach not only ensures ad relevance but also respects user privacy, a growing concern in the digital advertising space. Furthermore, LLMs can significantly reduce the occurrence of ad fraud. By understanding and predicting user engagement patterns, these models can identify anomalies that may indicate fraudulent activity, thereby

ensuring that advertising budgets are spent on genuine user interactions.

Another challenge in programmatic advertising is the creation and testing of a multitude of ad creatives, which requires significant resources and time. LLMs streamline this process by automating the generation of diverse and personalised ad content, from textual copy to potential image suggestions. This automation allows for rapid A/B testing and optimisation of ads, ensuring that the most effective content is delivered to the right audience at the right time.

In summary, LLMs bring scalability, efficiency, and a new level of innovation to programmatic advertising. By leveraging their ability to generate and optimise ad content, LLMs can produce a diverse range of advertisements tailored to various user segments and contexts. This not only mitigates the challenges posed

by the loss of third-party cookies, but also significantly enhances the overall user experience with ads that are highly relevant and engaging.

[LLMs'] capacity to accurately predict, understand, and emulate human-like text and interactions is changing the way advertising content is conceptualised, created, and delivered.

WHAT ARE LARGE LANGUAGE MODELS?

LLMs are advanced AI systems designed to understand, interpret, and generate human-like text. They can be conceptualised as tools that effectively compress vast amounts of internet data, distilling the essence of human communication into algorithms capable of mimicking language. The true potential of LLMs was unlocked following the landmark 2017 paper on the attention mechanism, which introduced a more efficient way for models to process and prioritise different parts of the input data.

This attention mechanism, particularly exemplified in models like the Transformer, revolutionised the field of natural language processing (NLP). It allowed for the development of more sophisticated and contextually aware models, capable of handling longer sequences of text and understanding the nuances and complexities of human language. Since then, LLMs have rapidly evolved, with notable examples like OpenAI, GPT-4 and Google's Gemini showcasing their ability to generate coherent and contextually relevant text across a variety of applications.

These models are trained on enormous datasets, often encompassing a significant portion of the publicly available text on the internet. This training enables them to learn patterns, styles, and information from a wide range of sources, effectively giving them a compressed understanding of human language as represented online. As a result, LLMs have become pivotal in numerous applications, particularly in areas like programmatic advertising, where the ability to generate personalised and relevant content at scale is crucial. Their emergence

represents a significant leap in AI's capability to interact with and understand human language, opening up new possibilities in technology and communication.

OPEN SOURCE VS PROPRIETARY LLMs

In the realm of LLMs, the distinction between open source and proprietary models presents a significant crossroad for data scientists and organisations. Open source models, like those released by Hugging Face, offer transparency and community-driven innovation. They allow researchers and practitioners to inspect the model's architecture, training data, and inner workings, fostering an environment of collaboration and trust. The agility of open source models accelerates experimentation and adoption in programmatic advertising, enabling practitioners to fine-tune models to specific domains or tasks without the constraints of licensing agreements. Hugging Face also hosts an open LLM leaderboard, where teams around the world submit new LLMs every week, which are automatically benchmarked and ranked. By regularly visiting this page, you can get the most performing open LLM - at the time of writing (December 2023) this was Microsoft's Phi-2, but it will likely be different by the time you read this article:

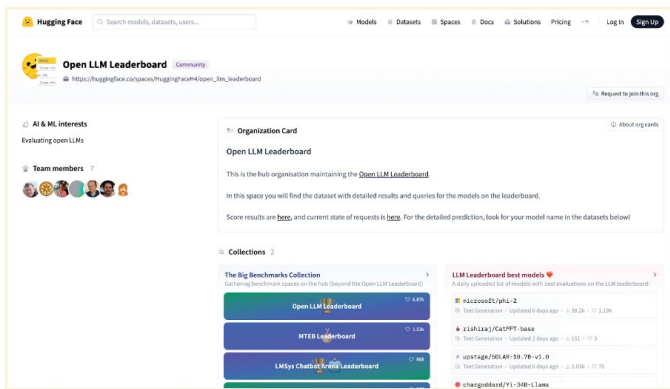


Figure 1. Open LLM leaderboard

Proprietary models, on the other hand, are developed and controlled by organisations like OpenAI. These models often come with performance benefits derived from proprietary datasets and resources unavailable to the broader community. Their closed nature can mean better integration, support, and potentially advanced features that cater to specific business needs in advertising, like enhanced personalisation and targeting capabilities.

However, the benefits of proprietary models come with trade-offs. The lack of transparency can raise concerns about the replicability of results and the ability to audit for biases or errors. The cost of access and potential usage restrictions can also be limiting factors, particularly for smaller organisations or independent developers.

Each type of LLM carries implications for scalability, innovation, and ethical considerations. Open source models enable wider accessibility and collective problem-solving, which can be crucial for tackling industry-wide

challenges such as ad fraud detection and privacy-preserving personalisation. Proprietary models may offer competitive advantages but require a commitment to vendor relationships and may necessitate navigating around black-box algorithms.

In the context of programmatic advertising, the choice between open source and proprietary LLMs hinges on factors like budget, desired control over the model, ethical considerations, and the specific advertising goals of an organisation. As the industry continues to evolve, the interplay between these two paradigms will shape the development and deployment of AI-driven advertising strategies.

STRATEGIES FOR IMPLEMENTING LLMs

Implementing LLMs within programmatic advertising frameworks requires strategic planning to fully harness their capabilities. Full model training is a resource-intensive process, often involving significant computational power and a vast corpus of training data to produce a model that can understand and generate human-like text. The advantages of training an LLM from scratch include the ability to customise the model to highly specific advertising needs, which can lead to increased ad relevance and engagement. However, the cost, expertise, and time required to train such models can be prohibitive for many organisations.

Fine-tuning pre-trained LLMs presents a more accessible alternative. By adjusting an existing model - such as GPT-3 or BERT - on a smaller, domain-specific dataset, data scientists can imbue the model with the nuances of their target audience or specific advertising context. This method requires far less computational resources and time, allowing for rapid deployment and iteration. Fine-tuning is particularly effective when the base model is already performing well and only minor adjustments are needed to tailor the model to specific campaign objectives.

RAG is a novel approach that combines the power of LLMs with external knowledge sources to generate content that is both relevant and contextually rich. By querying a database of information during the generation process, RAG can produce ad content that is informed by the latest market data, trends, or user-specific information, making it highly adaptive and personalised. This technique is especially beneficial in scenarios where ads need to be responsive to real-time events or user interactions.

EXPLORING FOUNDATION MODELS IN PROGRAMMATIC ADVERTISING

Foundation models, particularly LLMs like GPT-3, are at the heart of the next generation of programmatic advertising. These models are transformative due to their deep understanding of language nuances and user intent. Their design goes beyond simple keyword matching, enabling them to interpret the subtleties of human

communication and generate responses or content that feels authentic and engaging.

In programmatic advertising, the role of LLMs is multifaceted. They serve as the backbone for dynamic creative optimisation, where ad content is not just personalised but also generated in real-time to match the user's current context and sentiment. LLMs can craft copy that resonates with the user's current emotional state or intent, a capability that traditional models, which lack nuanced language understanding, cannot match.

For instance, GPT-3 can analyse vast datasets of successful advertising copy and images, learning from the patterns of high-performing ads. It can then generate similar content, but with variations tailored to different user profiles and platforms. This not only increases engagement by ensuring relevance but also significantly reduces the time and resources required for ad creative development.

Moreover, the application of these models extends to improving the efficiency of ad targeting. By understanding user queries and online behaviour, LLMs can help advertisers predict the most effective touchpoints for engagement, allowing for a more strategic deployment of advertising budgets.

As we continue to unlock the capabilities of LLMs, their integration into programmatic advertising workflows is becoming more nuanced. The potential for these models to learn and adapt over time promises a continuously improving advertising ecosystem, one that becomes increasingly efficient at delivering the right message to the right user at the right time.

LLMs like GPT-3 and DALL-E are driving significant transformations in programmatic advertising, offering innovative solutions to long-standing industry challenges. These models have opened up new avenues in content creation, ad targeting, and campaign optimisation, reshaping the way programmatic advertising is conceptualised and executed.

The ongoing development of foundation models presents an exciting frontier for data scientists in

advertising. By leveraging these advanced AI tools, they can address the industry's most pressing challenges, including improving user engagement and ensuring content relevance, while navigating the ever-important issues of privacy and user consent.


ETHICAL IMPLICATIONS

The integration of LLMs into programmatic advertising necessitates careful consideration of ethical implications, particularly regarding data privacy, algorithmic bias, and transparency. Data privacy concerns revolve around the extent and manner of personal data usage by LLMs to personalise advertising content. There is a growing demand for models that not only comply with data protection regulations like GDPR and CCPA, but also align with broader ethical principles of respect for user autonomy and consent.

Algorithmic bias in LLMs is another critical ethical issue. Since these models are trained on large datasets that may contain historical biases, there is a risk of perpetuating or amplifying these biases in ad targeting and content generation. It's essential to implement measures to identify and mitigate such biases, ensuring that advertising practices are fair and do not discriminate against any group.

Transparency in the use of LLMs is about making the model's decision-making processes understandable to users and regulators. It involves explaining how personal data influences the content generated by these models and how decisions about ad targeting are made. This level of transparency is crucial for building trust among users and for the ethical use of AI in advertising.

Current ethical guidelines, such as those proposed by AI ethics boards and industry groups, emphasise principles like fairness, accountability, and transparency. Applying these guidelines to the use of LLMs in advertising means ensuring that models are audited for biases, data usage is transparent and consensual, and there are mechanisms for accountability in cases of misuse or harm.



FUTURE DIRECTIONS AND RESEARCH

The future of LLMs in advertising is poised for significant advancements, with potential developments including more advanced personalisation techniques, improved user privacy protections, and enhanced model interpretability. Personalisation is likely to become more nuanced, with LLMs being able to understand and adapt to complex user preferences and behaviours while respecting privacy boundaries.

Research in model explainability will be crucial, as there is a growing need to make AI decision-making processes transparent, especially when they impact consumer experiences and choices. This involves developing methods to interpret how LLMs generate specific content and how they decide on particular ad placements. Such research will not only aid in compliance with evolving privacy regulations, but also help in building user trust.

Another vital area of research is finding the balance between personalisation and privacy. This includes developing techniques for data anonymisation and synthetic data generation, which can help in training LLMs without relying on sensitive personal information.

The call to action for the data science community is to actively engage in these research areas. Collaboration between industry practitioners, academics, and regulatory bodies will be key in advancing these technologies in an ethical and sustainable manner. Additionally, ongoing monitoring of the societal impacts of LLMs in advertising is necessary to ensure that the benefits of AI are realised broadly and equitably.

The exploration of LLMs in the context of programmatic advertising uncovers a landscape rich with potential and challenges.

CONCLUSION

The exploration of LLMs in the context of programmatic advertising uncovers a landscape rich with potential and challenges. LLMs, as foundational models, are redefining the norms of content creation and ad personalisation,

heralding a new era in digital marketing. The shift from traditional data-driven approaches to AI-centric methods, particularly in a post-cookie world, emphasises the need for innovative solutions to maintain ad relevance and user engagement.

The comparison between open source and proprietary LLMs reveals a trade-off between transparency, cost, innovation, and control. While open source models offer accessibility and collaborative advancement opportunities, proprietary models bring tailored solutions with competitive advantages. Implementing these models, whether through full training or fine-tuning, requires a strategic approach, balancing the campaign's objectives with available resources. The integration of techniques like RAG promises enhanced responsiveness and contextual relevance in ad content.

Ethical considerations, particularly concerning data privacy, bias, and transparency, remain at the forefront. The application of LLMs in advertising must adhere to ethical guidelines that prioritise fairness, accountability, and user consent. This ethical framework is not static; it evolves with the technology and societal values, necessitating continuous vigilance and adaptation from data scientists.

As we look to the future, the potential advancements in LLMs beckon a wave of innovation in advertising. Research in model explainability, balancing personalisation with privacy, and the mitigation of algorithmic bias will be pivotal. For data scientists in this domain, the journey is one of continuous learning and ethical introspection, ensuring that the advancements in AI are leveraged responsibly and beneficially.

On a closing note, our AI team at Fenestra.io stands at the forefront of innovation in programmatic advertising, leveraging advanced technologies to automate and optimise processes for traders. Specialising in transforming the complex landscape of digital ad trading, Fenestra.io harnesses the power of LLMs to revolutionise the industry. By automating intricate processes, from campaign optimisation to campaign reporting, the company provides traders with more time to focus on strategic decision-making and creative aspects of advertising campaigns.

REFERENCES

1. Campos Zabala, Francisco Javier (2023) *"How to grow your Business with AI"* Apress/Nature Springer.
2. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Nets. *"Advances in Neural Information Processing Systems"*.
3. Kingma, D. P., & Welling, M. (2013). Auto-Encoding Variational Bayes. *"arXiv preprint arXiv:1312.6114"*.
4. Vaswani, A., et al. (2017). Attention Is All You Need. *"31st Conference on Neural Information Processing Systems"*.
5. Radford, A., et al. (2019). Language Models are Unsupervised Multitask Learners. *"OpenAI Blog"*.
6. Brown, T. B., et al. (2020). Language Models are Few-Shot Learners. *"arXiv preprint arXiv:2005.14165"*.
7. Lewis, P., et al. (2020). Retrieval-Augmented Generation for Knowledge-Intensive NLP Tasks. *"arXiv preprint arXiv:2005.11401"*.
8. GDPR. (2018). General Data Protection Regulation.
9. CCPA. (2018). California Consumer Privacy Act.
10. Bender, E. M., et al. (2021). On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? *"Proceedings of FAccT"*.