

THE POWER OF DATA-DRIVEN INSIGHTS: ADVANCING AD TARGETING WITH GENERATIVE AI

Ofer Weintraub, Alain Nochimowski, Alice Wittenberg

Viaccess-Orca, Israel

ABSTRACT

Artificial intelligence (AI) is revolutionizing how video service providers engage audiences through digital content delivery. This paper outlines a novel architecture blueprint that use large language models (LLMs) to enhance ad targeting effectiveness through personalized messaging. This approach is structured around four key functional pillars: automated audience segmentation based on core beliefs (Tru-Values); AI-driven tailored message crafting, offering a dual-layer of personalization; a feedback loop to achieve continuous improvements; and the implementation of rigorous safety standards including privacy protections and bias prevention. A simplified implementation of the feedback loop inspired by Reinforcement Learning from AI Feedback (RLAIF) techniques is demonstrated to assess and improve ad campaign effectiveness over time. The objective is to create a personalized viewer experience that resonates deeply with diverse audience segments, thereby generating heightened engagement. Ultimately, the use of generative AI facilitates the creation of multimodal, customized advertising content at scale, paving the way for targeting methods that are more practical across large audiences.

INTRODUCTION

In 2017, Netflix demonstrated the impact of personalized promotional messaging by adapting artwork based on user preferences and viewing history (1). The company's content experts generated multiple images for every title and change them regularly to lure audiences based on their previous viewing history. The approach utilized a sophisticated "online reinforcement learning" strategy, optimizing the balance between exploiting known user preferences and exploring new data for improved recommendations. This dynamic methodology was essential in minimizing the cumulative "regret" (defined as the difference between the expected "payoff" (e.g. engagement) of the algorithm and the payoff of a single fixed strategy for selecting artworks) and enhancing viewer satisfaction over time. Despite its efficiency, this approach epitomized a few key principles under the legacy Personalization 1.0 paradigm, namely the need to choose between a finite set of creatives (while being limited by multimodal content creation's cost and complexity), the reliance on high-quality (and therefore expensive) human preference labels to optimize algorithmic tuning, and the use of shallow context information (e.g. transient signals like search history) as proxies for viewers' interests. Today's advancements in Generative AI necessitate a complete re-evaluation of the algorithmic / architectural trade-off and its relevance.

A new Personalization 2.0 approach pushes the envelope on at least four key pillars:

- 1. First Pillar: Tru-Values Audience Segmentation** — Advanced machine learning algorithms can craft dynamic and enduring audience profiles by extracting deep beliefs and Tru-Values beyond simple behavioural indicators. This Tru-Values segmentation ensures continuously relevant advertisements while augmenting traditional targeted advertising strategies, such as those based on IAB categories (3).
- 2. Second Pillar: Tailored Message Crafting** — With deep insights provided by Tru-Values segmentation, this pillar develops advertising content that meaningfully resonates with audiences. Using Generative AI, this process efficiently produces customized message across various media, simplifying the logistical challenges associated with campaign creation and ensuring that the advertisements resonate with audience values and beliefs.
- 3. Third Pillar: Continuous Improvement Through Closed-Loop Systems** — This pillar continually refines and optimizes advertising strategies based on direct, robust feedback and engagement mechanisms. By dynamically adjusting changes in consumer preferences and market trends, this feedback loop ensures relevant and effective advertising by resonating with the audience’s current values and interests.
- 4. Fourth Pillar: Ethical, Privacy, and Safety Considerations** — After integrating the first three pillars into systems for data management, multi-screen ad insertion, and analytics, personalized advertising campaigns can now be executed across the advertiser, broadcast, and operator domains. To ensure all personalized content upholds the highest privacy standards and ethical considerations, this architecture must adhere to stringent company guidelines, such as data usage policies, consent protocols, positive messaging principles, and inappropriate content safeguards. Additionally, it must incorporate rigorous ethical “guardrails” — algorithms that input a set of objects and determine which enforcement actions can be taken to reduce the risks embedded in the objects (4).

These four enhancements promise a more resonant and ethically responsible approach to personalized advertising. A Personalization 2.0 paradigm represents a forward-thinking approach to targeted video advertising, structured around the innovative concepts of Tru-Values segmentation, tailored message crafting, and continuous improvement through closed-loop systems. This paper will further illustrate a practical, simplified implementation of the third pillar, illustrating how a static set of Tru-Values and corresponding personas can feed a closed-loop system. for Personalization 2.0.

One critical aspect of the emerging TV/video ad targeting paradigm is the ability to constantly fine-tune advertising strategies based on evolving consumer preferences and market trends. Equally critical is the ability to achieve this at scale. Generative AI bears the promise of not only alleviating the cost and complexity of multimodal content creation but also of scaling up the historically expensive tasks of collecting and processing feedback.

Moving from Reinforcement Learning from Human Feedback (RLHF) to Reinforcement Learning from AI Feedback (RLAIF) is a novel direction and the topic of promising research (5). This paper proposes a simplified implementation of such an approach (maximizing scores instead of optimizing a reward function) and its application to self-improvement of targeted ads based on Tru-Values categories.

We have chosen an adversarial architecture featuring a Generator Large Language Model (LLM) and an Evaluator LLM. Figure 1 below depicts the main components:

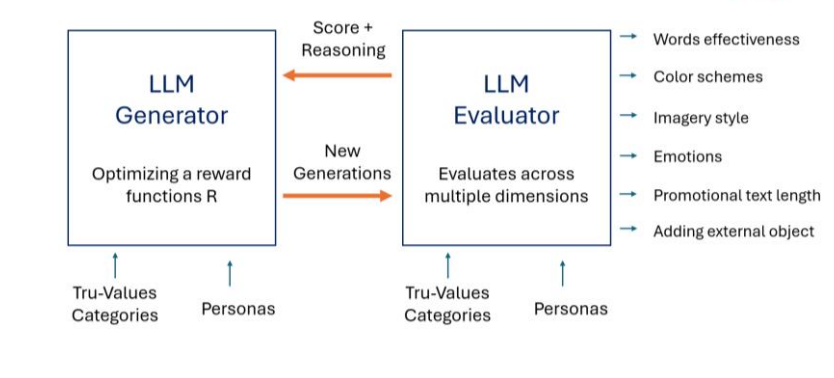


Figure 1 – Architecture blueprint for continuous self-improving targeted ads.

The main relevant components are:

1. An LLM that generates campaign creatives in the form of imagery, promotional text, audio, narration, video, etc.
2. The Generator LLM receives as context both the Tru-Values categories and the Persona (or user profile) of the target segment.
3. The Generator creates multiple variations across different eval-domains such as: choice of words, colour schemes, image style, text style, emotions, length of text, implied season, additional objects in image, etc...
4. An Evaluator LLM receives as context both the Tru-Values categories and the Persona (or user profile) of the target segment.
5. The Evaluator LLM scores each variant according to a pre-defines set of metrics (shown right to the LLM-eval) and return to the generator both score and explanation (or reason) for the given score along a specific domain.
6. The Generator uses a “reward function” and tries to maximize the score along the given domains.
7. The process may continue for as long as the scores improve.

This self-improving architecture results in highly optimized ads targeted to specific personas and a specific set of Tru-Values categories.

THE TRU-VALUES APPROACH IN PRACTICE

Before putting all these concepts into practice, the Tru-Values approach must be implemented, starting with the extraction of content categories (i.e., documentaries, books etc.) to define personas that represent these deep-seated values.

The value of video viewership data as a marker of viewers’ socio-demographics (which can be further enriched by third-party data) has been explored already (2). The notion of Tru-Values pushes the envelope even further, as most individuals tend to consume content and media that align with their Tru-Values. Tru-Values is defined as a set of deeply ingrained beliefs and principles that serve as a guiding compass for an individual's thoughts and behaviours. These values can be shaped by factors such as family, education, cultural environment, personal growth experiences, and more. Examples of Tru-Values include:

- Environmental conservation and promotion of green energy
- Providing health and social services in underprivileged areas
- Seeking justice for underserved populations

- Advancing equality through technology

As for any other personal data, manipulating Tru-Values would be subject to privacy regulations (e.g., GDPR). While people typically take pride in their values and are willing to share and discuss them, the most stringent data protection techniques are expected to be implemented by any segmentation system making use of Tru-Values.

Ultimately, a Tru-Values segmentation approach enables the delivery of continuously relevant advertisements while complementing traditional targeted advertising strategies, such as those based on IAB (Internet Advertising Bureau) (3) categories. The result is a more nuanced and effective method of connecting with audiences based on their core values and beliefs.

Generative AI can be used in combination with Tru-Value segmentation in order to customize variable campaigns — altering messages, imagery, length, colour schemes, and emotional tones — and use robust analytics to assess their effectiveness by understanding consumer response. This feedback loop will help us continuously refine advertising strategies by ensuring each ad resonates deeply and authentically with its intended audience.

Segmentation Based on Tru-Values

This paper proposes a “static” implementation of Tru-Values as opposed to a “dynamic” implementation, which derives and continuously updates Tru-Values categories from viewed content. Here's how specific Tru-Values categories could possibly be elucidated:

1. Environmental Conservation is exemplified through documentaries like "Before the Flood" and "Planet Earth II" and books such as Rachel Carson's "Silent Spring" or Elizabeth Kolbert's "The Sixth Extinction." These sources provide insights into the importance of preserving natural habitats and the urgency of addressing climate change.
2. Renewable Energy insights are drawn from documentaries like "An Inconvenient Truth" and "Pandora's Promise," alongside books like "Sustainable Energy - Without the Hot Air" by David J.C. MacKay, which discuss innovative energy solutions reducing dependency on fossil fuels.
3. Educational Equity is represented by "Waiting for Superman" and "The Smartest Kids in the World" by Amanda Ripley, highlighting disparities in educational access and the transformative power of equitable education.
4. Humanitarian Efforts are captured in "He Named Me Malala" and "Living on One Dollar," showcasing personal stories of advocacy and the fight against global poverty.
5. Social Justice themes are explored in "13th" and "Just Mercy" by Bryan Stevenson, which address systemic inequality and the pursuit of justice in the legal system.

Each category informs and shapes not only their representative personas but also the creation of targeted messaging and campaigns. By analysing audience engagement with these themes, we tailor content to resonate with specific values, constantly refining strategies through feedback loops to ensure alignment with both consumer preferences and ethical standards. This comprehensive approach ensures that each campaign drives deeper engagement and fosters a genuine connection with the audience.

Personas Based on Tru-Values

To demonstrate how Tru-Values can be used effectively, we created several fictitious personas using the above method of associating Tru-Values categories with users. For clarity, all examples introduced in this paper have been generated with GPT-4o, although there is no known dependency on the choice of the underlying LLM.

Here are four examples of such personas, whose images were generated using DALL-E:



Figure 2 - Green Advocate George




Figure 3 – Humanitarian Harish



Figure 4 – Health Advocate Hannah



Figure 5 – Justice Warrior Jane

	Green Advocate George	Humanitarian Harish	Health Advocate Hannah	Justice Warrior Jane
Short description	George is deeply committed to environmental conservation and renewable energy, spending his time advocating for sustainable practices	Harish has a serious, almost sombre expression, emphasizing his deep concern for global humanitarian issues.	Hannah, a young medical professional, is dedicated to addressing global health challenges and improving public health infrastructure.	Jane, a young Native American activist, champions social justice through grassroots activism, focusing on equality and rights.
Value & beliefs	George is a passionate environmentalist who believes in sustainable living and the protection of natural habitats. He actively promotes green technologies.	Harish focuses on providing relief and support to communities affected by crises and advocating for human rights.	Hannah is dedicated to improving public health and reducing disparities in healthcare access.	Jane is an activist who fights for social justice, equality, and the rights of marginalized communities. She uses her voice to challenge societal norms and advocate for change.

Tru-Value Categories	Environmental Conservation, Renewable Energy.	Humanitarian Efforts, Human Rights.	Global Health Challenges	Social Justice.
----------------------	---	-------------------------------------	--------------------------	-----------------

Table 1 – Personas & Tru-Values

Products and Services

To demonstrate how ads can be tailored to these personas based on their Tru-Values and deep beliefs we define several products which we try to sell to the mentioned personas.

Multi-Purpose Smartwatch	Eco-Friendly Apparel Line	Reclaimed Wood Furniture	Customizable Children's Books	Charity-Linked Water Bottles
A smartwatch with applications for health monitoring, educational tools, environmental tracking, and more, catering to a wide range of "Tru-Value" categories like Global Health Challenges, Educational Equity, Environmental Conservation, and Technological Impact.	Clothing made from sustainable materials, where each purchase supports initiatives like animal welfare, environmental conservation, and economic equality, targeting Animal Welfare, Environmental Conservation, and Economic Inequality.	Furniture crafted from reclaimed wood, appealing to those interested in Environmental Conservation, Economic Equality, and Community Development by promoting sustainability and supporting local craftsmanship.	Books that can be personalized for each child, supporting Child Protection and Educational Equity, and printed on-demand to minimize waste, promoting Environmental Conservation	Reusable water bottles linked to charitable donations towards clean water initiatives, appealing to Water Security, Economic Equality, and Global Health Challenges.

Table 2 – Products & services

Examples of True-Values-Based Campaigns Generated by LLM

An example of personalized campaign for “Green Advocate George” for a smartwatch is shown below:



Persona	Green Advocate George	Green Advocate George	Justice Warrior Jane	Justice Warrior Jane
Product & Service	Multi-Purpose Smartwatch	Eco-Friendly Apparel Line	Reclaimed Wood Furniture	Customizable Children Book
Example Title	<i>“Map your carbon footprint”</i>	<i>“Organic comfort”</i>	<i>“Artisanal Integrity: the elegance of Sustainable Craftmanship”</i>	<i>“World of Wonder: Custom Tales for every Children”</i>
Promotional text	<i>“Track your daily carbon footprint and make environmentally conscious choices effortlessly”</i>	<i>“Feel the touch of 100% organic cotton. Sustainable fashion that cres for the earth”</i>	<i>“Discover the subtle artistry in our reclaimed wood furniture, Jane. Each piece embodies sustainability and skilled craftsmanship, aligning with your values of care & community”</i>	<i>“Explore vibrant stories crafted for your child, celebrating diversity & sparking curiosity. Every page turns to nurture a young mind towards inclusivity and understanding”</i>

Table 3 – Examples of Tru-Value-based campaigns

Of course, we can ask the LLM to generate colour and style variations to test which one works better for a given personality. Here are two examples of colour and style templates. For the “Reclaimed Wood Furniture”:



Figure 6 - Reclaimed Wood Furniture Variations

Or for “Customizable Children Books”:



Figure 7 - Customizable Education Books Variations

Evaluating Campaign Effectiveness

Using a feedback loop based on the above-mentioned principles, LLMs can evaluate the quality of advertising campaigns across multiple domains like promotional text, visual aesthetics, and emotional resonance. LLMs can assess how well content aligns with a specific persona by providing a score from 1 to 10 and detailed feedback.

Below is an example of checking which colour scheme is better for “Humanitarian Harish” and the “Charity-Linked Water Bottles”.



- **Score:** 7/10
- **Explanation:** Calm, emphasizes water security, lacks direct human impact.

- **Score:** 8/10
- **Explanation:** Represents sustainability and nature, subtly supports his values.

- **Score:** 10/10
- **Explanation:** Directly resonates with crisis relief and urgent humanitarian aid.

Figure 8 - Humanitarian Harish - Charity-Linked Water Bottles

Another example is “Eco-Friendly Apparel” for Humanitarian Harish, both overall score and activity score are given by the LLM.



1. **Color Scheme Selection:** 9/10
2. **Activity Depicted:** 9/10 - Urban gardening perfectly aligns with Harish's focus on community support and sustainability.

1. **Color Scheme Selection:** 8/10
2. **Activity Depicted:** 10/10 - Beach cleaning directly represents environmental activism, resonating deeply with Harish's humanitarian efforts.

1. **Color Scheme Selection:** 10/10
2. **Activity Depicted:** 10/10 - Planting trees is a strong symbol of Harish's dedication to environmental conservation and reforestation efforts.

Figure 9 - Humanitarian Harish – Eco-friendly Apparel

We can ask the LLM to analyse which set of colours will perform best for environmental-oriented ads, and to explain the choice of colours — as shown by Figures 10 and 11, respectively.




Figure 10 – Effectiveness of Colours for Environmental Campaigns

Most Effective Colors

- 1.Green:** 10/10 - Represents nature, growth, and sustainability.
- 2.Blue:** 9/10 - Symbolizes water, sky, and clean air.
- 3.Brown:** 8/10 - Earth tones that evoke a natural and grounded feeling.
- 4.Yellow:** 7/10 - Represents sunlight and renewable energy, evoking a sense of optimism and positivity.
- 5.Light Green:** 8/10 - Often associated with new growth and eco-friendliness.

Less Effective Colors

- 1.Bright Red:** 3/10 - Too aggressive for peaceful environmental messages, may convey danger or warning.
- 2.Dark Gray:** 4/10 - Can imply industrialization and urban sprawl, contrary to natural conservation.
- 3.Neon Colors (like neon green or neon yellow):** 2/10 - Feel artificial and are too stark, clashing with organic aesthetics.
- 4.Black:** 3/10 - Often symbolizes oil and pollution, which might detract from a clean and green message.
- 5.Metallic Tones (like silver or gold):** 3/10 - Imply industrial luxury, which can conflict with the principles of modesty and simplicity in environmental advocacy.

Figure 11 – Evaluator LLM explainability

We can also ask the LLM to show us what words are most effective and which words are less effective, as shown below:

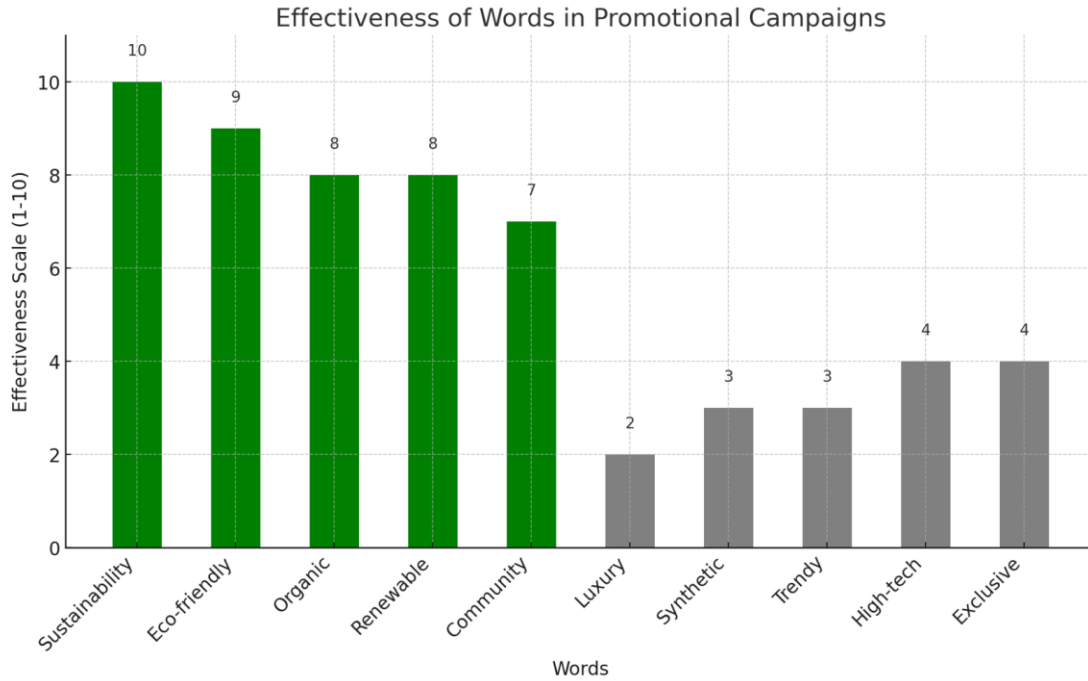


Figure 12 – Effectiveness of Words in Promotional Campaigns

We can compare emotions across multiple ads to see which ones works best:

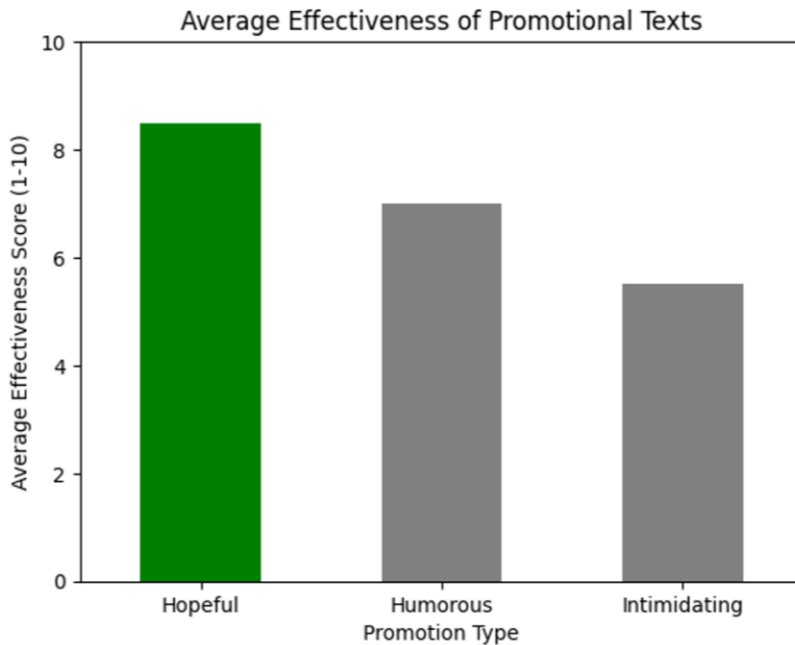


Figure 13 – Mood Effectiveness

We can add different objects (e.g., baby, rabbit, or vegetables) and ask the evaluator how these affect the quality of the promotions. An example is shown below:




Figure 14 – Adding Objects – Reclaimed Wood Furniture

The AI can also rate multiple promotions (50 in this case) and rank them from best to worst:

Less effective promotion:

Text: "Your daily water routine can bring daily relief to others. Be a part of the hydration solution."

Score: 1/10

Explanation: Vague impact, lacks emotional or direct human connection

Highly effective promotion:

Text: "Drink in the difference you make! Every bottle purchased funds clean water projects worldwide, hydrating those in need one sip at a time."

Score: 10/10

Explanation: Directly links purchase to global impact, emphasizes hydration and aid.

Figure 15 – Humanitarian Harish – Charity-Linked Water Bottles

Or we can ask the LLM to generate multiple promotions of varied lengths (10 words to 100 words) and to determine which is best for most effective ads:

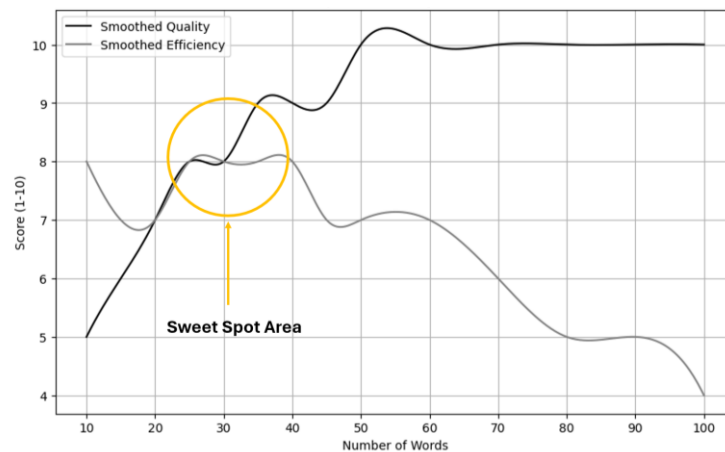


Figure 16 – Smoothed Quality vs Efficiency of Promotional Texts

In Figure 16 above, “Quality” means how well the message is being presented, whereas “Efficiency” takes into account that long messages — despite improving quality — are less

effective as it becomes cognitively difficult to absorb and process very long messages. As a result, a message of around 30-40 words seems optimal.

CONCLUSION

This paper introduces the Tru-Values approach as a novel paradigm in personalized advertising, highlighting its ability to harness deep-seated values and beliefs derived from content for more precise audience targeting. Through the implementation of three foundational pillars, supplemented by a crucial fourth pillar focusing on ethical standards, the approach signals a paradigm change in the way ads are tailored and delivered.

A robust, continuous improvement framework using closed-loop systems like RLAIIF is introduced to adapt and refine strategies in response to shifting viewer preferences and feedback.

The overall approach proposed not only augments traditional advertising methods but also paves the way for a scalable, ethical framework for continuous self-improvement in ad targeting. By aligning with Tru-Values, advertisers can achieve a dual-layer of personalization that not only broadens engagement but also enriches the viewer's experience, ensuring each campaign is not just seen but felt, driving deeper connections and sustaining brand loyalty in the digital age.

REFERENCES

1. Chandrashekar, A., Amat, F., Basilico, J., and Jebara, T., 2017. Artwork Personalization. [Netflix Technology Blog](#).
2. Nochimowski, A., Wittenberg, A., and Zucker D., 2018. Implementing AI-Aided Content Distribution Strategies in the GDPR Era, [IBC Tech Papers](#).
3. Interactive Advertising Bureau Tech Lab, 2024. Audience Taxonomy 1.1. <https://iabtechlab.com/standards/audience-taxonomy/>.
4. Dong, Y. et al., 2024. Building Guardrails for Large Language Models, [arXiv.org](#)
5. Lee. H., et al. 2023. RLAIIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback, [arXiv.org](#)