

SCIENCE & VALIDATION

THE SCIENCE BEHIND PREDICTIVE CREATIVE INTELLIGENCE

NeuroVision is built on one powerful scientific principle:


HUMAN ATTENTION IS NOT RANDOM.

Before people consciously judge a design, their brain has already started filtering the scene. In the first moments of viewing, attention is shaped by visual signals such as contrast, colour, spatial layout, faces, text, saliency, and visual hierarchy. These early attention patterns can be modelled, benchmarked, and compared with real eye-tracking data (Itti & Koch, 2001; Tatler et al., 2011; Bylinskii et al., 2019).

NeuroVision turns this science into a scalable creative intelligence system.

We help teams predict what people are likely to notice, what they may miss, how they may interpret the message, and what should be improved before launch.

METRIC CARDS



1–2 sec The critical early attention window where visual hierarchy starts to form.

Millions Gaze points and fixation signals used to inform predictive attention modelling.

Up to ~ 95% Benchmark similarity to real eye-tracking patterns under selected validation conditions.

Minutes Early creative validation without waiting weeks for traditional research.

PREDICTIVE VISUAL ATTENTION

FROM HUMAN GAZE TO SCALABLE PREDICTION

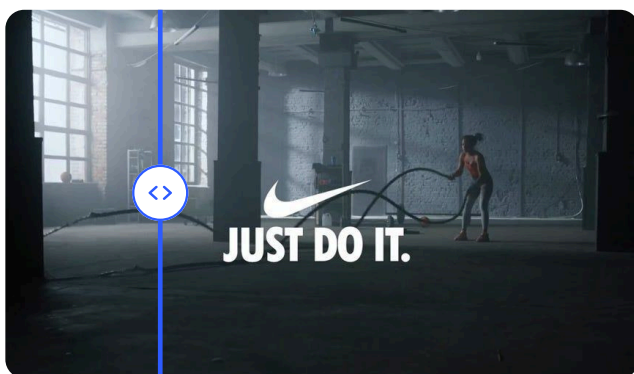
Eye-tracking research has shown that early gaze behaviour is not random. In the first moments of viewing, attention is partly driven by bottom-up visual features such as contrast, colour, saliency, spatial composition, object structure, faces, text, and visual hierarchy (Itti & Koch, 2001; Bruce & Tsotsos, 2009).

This is the scientific foundation behind predictive eye-tracking.

Large-scale research from MIT, Tübingen, and the wider saliency-modelling field has shown that computational models can predict human fixation patterns by learning from real gaze data. Deep learning models such as DeepGaze II and DeepGaze IIE have demonstrated strong benchmark performance on MIT/Tübingen saliency datasets, showing that human visual attention can be modelled as a measurable probability pattern — not as perfect individual certainty, but as statistically meaningful prediction (Kümmerer et al., 2016; Linardos et al., 2021).

NeuroVision brings this research into applied neuromarketing.

We use these scientific foundations to build and tune our own predictive attention system for business and creative applications. Instead of applying generic saliency maps, NeuroVision adapts predictive eye-tracking to real marketing materials — including ads, product visuals, landing pages, banners, brand assets, and campaign concepts — and strengthens the model through our own research data, validation studies, and neuromarketing-specific use cases.



The result is a fast, scalable attention layer that helps teams estimate whether the product, brand, headline, CTA, or key message is likely to be seen before media budget is spent.

AGENTIC SURVEYS

AUDIENCE INTERPRETATION AT SCALE

Attention prediction shows what people may see in the first one to two seconds. Agentic surveys help estimate what different audiences may think, feel, trust, or misunderstand.

Traditional surveys are powerful, but they are limited by recruitment, cost, speed, and sample availability. In many real-world research settings, the exact audience a brand wants to understand — by location, demographic profile, culture, purchasing context, or professional background — can be difficult, expensive, or sometimes practically impossible to recruit with precision.

NeuroVision adds a faster first-pass layer:
STRUCTURED AGENTIC AUDIENCE SIMULATION.

These are not generic chatbots. They are controlled behavioural agents designed to approximate how different demographic, regional, and psychographic groups may respond to a creative stimulus. Each agent interprets the design through structured perception layers, allowing NeuroVision to detect clarity gaps, trust signals, emotional tone, brand fit, purchase relevance, and potential misunderstanding before launch.

In certain early-stage research contexts, agentic surveys may be more useful than traditional panels because they can model highly specific audience profiles consistently, repeatedly, and at scale — including segments that are normally difficult to recruit. This makes it possible to generate structured brand insight in minutes, compared with the days or weeks often required for conventional survey recruitment, fieldwork, and analysis.

This approach builds on emerging research in behavioural simulation and synthetic agents (Horton, 2023; Binz & Schulz, 2023; Park et al., 2024; Mei et al., 2024). NeuroVision is now extending this into a neuromarketing validation workflow with academic and industry partners, with research currently being prepared for submission (Kostrzewa & Molnár, 2026, manuscript in preparation).

The screenshot displays the NeuroVision software interface. On the left is a navigation sidebar with 'Dashboard', 'Creative Center', and 'Settings'. The main content area is titled 'Creatives / Nike 05/06/2026' and has tabs for 'Overview', 'Analysis', and 'Surveys (1)'. The 'Overview' tab is active, showing an 'Ad Preview' of a Nike advertisement with the slogan 'JUST DO IT.' and a person in a gym. Below the ad preview is a section titled 'The big insight?' with a lightbulb icon, stating: 'Nike's design is authentic to its core audience — but risks alienating those who want inspiration, not intimidation. The "Just Do It" message is still magnetic, but the visual tone may need nuance to broaden appeal without losing edge.' To the right is a 'Report' section titled 'Nike "Just Do It" Ad: Mixed-Methods Insight Report' for 'Executive Strategy & Creative Teams', prepared by 'Qualitative UX Researcher • Marketing Scientist'. The report includes an 'EXECUTIVE SUMMARY' and several key findings: 'This isn't just another Nike ad. It's a raw, moody, sweat-drenched declaration of grit — and it's working. But not for everyone. The design nails brand consistency (88%; 4.7/5) and professionalism (89%; 4.3/5). People feel Nike's DNA in the industrial gym, the battle ropes, the red top slicing through gray. The "Just Do It" stays still a hummer. Still effective. Still trusted.'; 'But here's the tension: Intensity is a double-edged sword. Younger audiences (18–24) and those with lower income or education levels love the authenticity. They see real sweat, real struggle, real London warehouse gyms. One said: "The way she's sweating and working hard makes me believe Nike gets real athletes." (Translation: This feels like real life, not a photoshoot.)'; 'Older (25–34), higher-income, or higher-education groups? They're more divided. Some call it "relentless" or "unstoppable" — good. Others say "overdone," "staged," or "looks warmish." One Masters grad in the 25–34 group said: "Focused but lacks warmth." (Translation: It's powerful, but emotionally cold.)'; and 'The brand personality? Rugged (92%) wins — not "exciting" or "competent." And 78% see it as irrevocable, not traditional. That's a win. But "irrevocable" here doesn't mean tech or futurism — it means rebelling what motivation looks like: raw, unpolished, no-nonsense.'

The result is fast, repeatable audience interpretation before launch — especially for audience segments that are too slow, costly, or difficult to reach through traditional survey recruitment.

VALIDATION PROGRAM

PROVING THE WORKFLOW FROM SCIENCE TO REAL-WORLD USABILITY

NeuroVision was built on validated scientific foundations: eye-tracking, visual saliency, predictive attention modelling, behavioural simulation, and applied cognitive science.

Now, together with academic and industry partners, we are working to validate the full NeuroVision workflow as one integrated system: predictive eye-tracking, agentic surveys, creative interpretation, and AI-assisted optimization for real-world business use.

Our goal is to help establish a new standard for applied neuromarketing software — fast enough for everyday creative teams, but grounded enough to be scientifically meaningful.

STUDY 0 - FEASIBILITY

Initial feasibility validation of this complex workflow has been completed and is being refined for industry-focused research communication.

Focus: commercial usefulness, quality of AI interpretation, predictive attention value, and workflow fit for creative teams.

STUDY 1 - EYE-TRACKING BENCHMARK

NeuroVision predictions are being benchmarked against existing eye-tracking datasets.

Initial benchmarking has reached approximately 95% similarity under selected benchmark conditions, and we are now working to further improve model performance, robustness, and applicability across real marketing materials.

MAIN STUDY - EYE-TRACKING + AGENTIC SURVEY VALIDATION

A new validation study is being prepared with academic and industry partners to evaluate the complete NeuroVision workflow, including predictive attention, agentic interpretation, cost savings, and industry applicability.

Current setup: controlled eye-tracking, independent online survey responses, well-known brand materials, and comparison between human feedback and NeuroVision's agentic interpretation outputs.

Expected outcome: evidence that NeuroVision can significantly reduce the time and cost of early-stage creative testing while maintaining meaningful predictive and interpretive value.

Small label:

Research in progress with academic and industry partners.

OUR SCIENTIFIC POSITION

NOT A REPLACEMENT FOR HUMAN RESEARCH. A NEW LAYER BEFORE IT.

NeuroVision is not designed to replace every eye-tracking study, survey, or human research project.

It is designed to bring a fast, scalable validation layer into everyday creative work — where traditional research is usually too slow, too expensive, or impossible to run for every design.

The platform helps teams understand:

- **WHAT PEOPLE ARE LIKELY TO NOTICE**
- **WHAT THEY MAY MISS**
- **HOW THEY MAY INTERPRET THE MESSAGE**
- **WHAT SHOULD BE IMPROVED BEFORE LAUNCH**

This gives creative teams a more objective way to evaluate design before spending time, tokens, budget, and media money behind it.

NeuroVision

REFERENCES

Bruce, N. D. B., & Tsotsos, J. K. (2009). Saliency, attention, and visual search: An information theoretic approach. *Journal of Vision*, 9(3), Article 5. <https://doi.org/10.1167/9.3.5>

Bylinskii, Z., Judd, T., Oliva, A., Torralba, A., & Durand, F. (2019). What do different evaluation metrics tell us about saliency models? *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 41(3), 740–757. <https://doi.org/10.1109/TPAMI.2018.2815601>

Goldstein, D. G., Suri, S., McAfee, R. P., Ekstrand-Abueg, M., & Diaz, F. (2014). The economic and cognitive costs of annoying display advertisements. *Journal of Marketing Research*, 51(6), 742–752. <https://doi.org/10.1509/jmr.13.0439>

Horton, J. J. (2023). Large language models as simulated economic agents: What can we learn from Homo silicus? National Bureau of Economic Research. <https://doi.org/10.3386/w31122>

Itti, L., & Koch, C. (2001). Computational modelling of visual attention. *Nature Reviews Neuroscience*, 2(3), 194–203. <https://doi.org/10.1038/35058500>

Kümmerer, M., Wallis, T. S. A., & Bethge, M. (2016). DeepGaze II: Reading fixations from deep features trained on object recognition. *arXiv*. <https://doi.org/10.48550/arXiv.1610.01563>

Kümmerer, M., & Bethge, M. (2021). State-of-the-art in human scanpath prediction. *arXiv*. <https://doi.org/10.48550/arXiv.2102.12239>

Kümmerer, M., Khanuja, H., & Bethge, M. (2025). Modeling saliency dataset bias. *arXiv*. <https://doi.org/10.48550/arXiv.2505.10169>

Linardos, A., Kümmerer, M., Press, O., & Bethge, M. (2021). DeepGaze IIE: Calibrated prediction in and out-of-domain for state-of-the-art saliency modeling. *arXiv*. <https://doi.org/10.48550/arXiv.2105.12441>

Mei, Q., Xie, Y., Yuan, W., & Jackson, M. O. (2024). A Turing test of whether AI chatbots are behaviorally similar to humans. *Proceedings of the National Academy of Sciences*, 121(9), Article e2313925121. <https://doi.org/10.1073/pnas.2313925121>

Park, J. S., Zou, C. Q., Kamphorst, J., Egan, N., Shaw, A., Hill, B. M., Cai, C. J., Morris, M. R., Liang, P., Willer, R., & Bernstein, M. S. (2024). LLM agents grounded in self-reports enable general-purpose simulation of individuals. *arXiv*. <https://doi.org/10.48550/arXiv.2411.10109>

Ries, E. (2011). *The lean startup: How today's entrepreneurs use continuous innovation to create radically successful businesses*. Crown Business.

Rogers, E. M. (2003). *Diffusion of innovations* (5th ed.). Free Press.

Tatler, B. W., Hayhoe, M. M., Land, M. F., & Ballard, D. H. (2011). Eye guidance in natural vision: Reinterpreting salience. *Journal of Vision*, 11(5), Article 5. <https://doi.org/10.1167/11.5.5>