

NEURAL SEAL

AI Safety Entrepreneurship White Paper

Kailash Balasubramaniyam
(Founder)

Mohammed Areeb Aliku
(Design and Frontend
Specialist)

Eden Simkins
(Research)

Neural Seal is an AI transparency solution that creates a standardized labeling framework—akin to “nutrition facts” or “energy efficiency ratings”—to inform users how AI is deployed in products or services.

Modern products and services frequently rely on AI at various stages of development—often without clearly communicating to end users how or where these intelligent systems are applied. When issues arise, consumers do not understand the AI’s role, leading to mistrust and uncertainty about the product or service. For instance, in **2020, Google’s Verily Health Sciences introduced an AI system in Thailand** designed to detect diabetic retinopathy. During field trials, the system encountered real-world challenges such as poor lighting conditions, low-resolution images, and cloud-based processing delays. As a result, **21%** of input images were deemed unusable, significantly reducing diagnostic accuracy and limiting the number of patients clinics could serve each day (*Heaven, 2020*).

This example underscores the need for transparent disclosures of where, when, and how AI is used, so consumers can make informed decisions and companies can proactively address potential limitations. By clearly labeling AI involvement and explaining its constraints, we believe businesses and consumers alike will have greater trust in AI-driven products and services—and a stronger incentive to use AI ethically and responsibly. By offering an **A/B/C/D rating** and a **QR code** that links to deeper metrics, Neural Seal empowers consumers and regulators to quickly identify the level of AI involvement, potential risks, and the degree of human oversight.

This transparent labeling system not only makes AI usage understandable for non-experts but also pushes companies to implement safer, more ethical AI practices—much like mandatory ingredient disclosures in the food and healthcare industries. Our approach is unique because it unifies technical explainability (e.g., how data is processed, which algorithms are used) with regulatory compliance (GDPR, EU AI Act, and other international guidelines) in a single, concise label. Commercially, Neural Seal opens doors for organizations looking to signal trustworthiness and proactively meet evolving AI regulations. By establishing an industry-wide standard, we believe our labeling framework will become the “go-to” benchmark for AI safety and accountability by 2027.

1. Problem overview

Modern AI systems influence critical decisions in finance, healthcare, social media, and beyond. Despite this, the average consumer has limited visibility into when AI is used, how it processes personal data, or which human oversight measures are in place. This opacity can lead to biases, privacy violations, and unsafe outcomes—especially when corporate incentives favor speed and profit over thorough safety checks.

Neural Seal addresses this challenge by creating a **universal labeling standard** that mirrors proven transparency strategies in other industries. For example, labeling laws revealing

carcinogenic ingredients in talcum powders helped protect consumers (*Information and Resources About Cancer: Breast, Colon, Lung, Prostate, Skin, n.d.*), and mandatory emissions disclosures compelled car manufacturers to meet stricter environmental regulations (*The European Green Deal, 2021*). These precedents demonstrate that straightforward, accessible information can shift consumer behavior and market practices toward better standards of safety and ethics.

Given the lack of a standardized mechanism for AI transparency, our approach tackles a big and meaningful problem: bridging the gap between highly technical AI operations and end-user understanding, while also enabling regulators to track compliance more efficiently. With AI adoption growing exponentially, a robust, user-friendly transparency system becomes increasingly urgent to preserve trust and mitigate risks.

2. Solution

2.1 Proposed Solution's Technical Foundation

Neural Seal incorporates a standardized **Structural Evaluation** with multi-step questionnaire that companies must complete to detail their AI usage and oversight processes across each stage of product or service development. This questionnaire evaluates factors such as data handling (future ideas), AI usage, decision-making autonomy, and human supervision, forming the basis of our universal labeling standard. By gathering consistent, verifiable information, Neural Seal ensures that consumers and regulators alike have a transparent view of AI's role and potential risks—ultimately fostering greater trust and accountability in AI-driven solutions.

The scoring model generates a **color-coded rating (A/B/C/D)** for each product or service, displayed on a concise label. A **dedicated QR code** on the label then directs users to an online dashboard that provides deeper insights into Data Collection & Privacy Safeguards (future ideas), Model Interpretability & Risk Analysis (future ideas), Human Oversight Mechanisms Across Key Workflows, Overall AI Influence in Each Operational Phase, Compliance with Relevant Regulations (e.g., EU AI Act, ISO 27001) (future ideas)

2.2 How It Addresses the Identified AI Safety Challenge

By **enforcing transparency**, Neural Seal allows consumers, partners, and regulators to see how AI decisions are made and what safeguards are in place. This discourages unethical uses of AI, and aligns development incentives with consumer protection and legal requirements.

2.3 Competitive Advantages & IP Potential

Neural Seal ensures universal applicability across diverse sectors by offering industry-specific add-ons (healthcare, finance, retail). Its proprietary scoring algorithm—weighted to measure data privacy, fairness, and oversight—can be patented, positioning it as a first-mover standard that establishes high brand equity, trust, and an industry-wide benchmark for AI transparency.

2.4 Technical Diagrams / Proof-of-Concept Results





Figure 1 – Preliminary label design with color-coded rating and QR code.



Detailed Breakdown				
Area	AI Usage	Key AI Tech	Human Oversight	Impact Score
1. R&D	Heavy	Predictive modeling & generative design	Yes (final decisions)	9/10
2. Supply Chain	Moderate	Demand forecasting, route optimization	Some (exception cases)	7/10
3. Marketing	High	Recommendation engines, personalized campaigns	Yes (brand alignment)	8/10
4. Operations	Moderate	Robotic process automation, scheduling	Yes	6/10
5. Quality Control	Heavy	Computer vision for defect detection	Minimal (auto alerts)	9/10
6. Customer Service	Low	Basic chatbots, FAQ automation	Yes (escalations)	5/10
7. Analytics	High	Real-time sales dashboards, advanced ML	Yes (executive signoff)	8/10
8. Maintenance	Moderate	Predictive maintenance, sensor data analysis	Yes (final confirmation)	7/10

⚠ Disclaimer: All data displayed here is purely illustrative and does not represent any real automotive company's operations. It's a glimpse into how you might structure your own AI usage dashboard in Notion, complete with charts, tables, and mind maps.

Figure 2 – Dashboard mockup that presents in-depth metrics when the user scans the QR code.

2.5 Innovation and Scalability: Neural Seal merges:

Neural Seal leverages explainable AI techniques like LIME and SHAP to boost interpretability (*LIME Vs SHAP, n.d.; Nieto Juscafresa, 2022*), supports compliance with key standards such as the EU AI Act, GDPR, and ISO 27001 (*EU AI Act, 2023; ISO/IEC 27001:2022*), and features a simple A–D grading design for consumers. Its modular architecture ensures it can adapt to regulatory changes, scale across industries, and be localized for different markets.

3. Pilot Experiment

3.1 Outline of Hackathon Build : Over the weekend, we aimed to develop an Angular-based web app to guide businesses through the Dynamic Structural Evaluation. Due to limited frontend experience, we adapted existing form templates rather than coding everything from scratch. Despite this pivot, we achieved a functional prototype embedded in our main site.

3.2 Screenshots & Prototype Results




Figure 1 – multi-step questionnaire design

Business Type Selection

Q1: Which of the following best describes your business?

☐ Product-Based Industry
 ☐ Service-Based Industry

☐ Hybrid (Product + Service)

Next →

Industry-Specific Questions

Q2: Select the product category that best represents your business?

☐ Food & Beverage
 ☐ Consumer Electronics
 ☐ Pharmaceuticals

☐ Automotive
 ☐ Apparel / Fashion
 ☐ Furniture

☐ Cosmetics / Personal Care
 ☐ Toys & Games
 ☐ Household Cleaning Supplies

☐ Construction Materials (e.g., Cement, Steel)

Next →

Lifecycle stages

Q3: How much is AI involved in these different lifecycle stages?
(Select all that apply)*

	None	Minimal	Moderate	Heavy
Client Outreach	☐	☐	☐	☐
Needs Assessment	☐	☐	☐	☐
Solution Design	☐	☐	☐	☐
Implementation	☐	☐	☐	☐
Performance Tracking	☐	☐	☐	☐
Reporting	☐	☐	☐	☐
Billing	☐	☐	☐	☐
Ongoing Support	☐	☐	☒	☐

[Clear](#)

Next →

Lifecycle stages

Q3: How much is AI involved in these different lifecycle stages?
(Select all that apply)

	Is there human oversight?	Does AI make the final decision?
Client Outreach	☐	☐
Needs Assessment	☐	☐
Solution Design	☐	☒
Implementation	☐	☐
Performance Tracking	☒	☐
Reporting	☐	☐
Billing	☐	☐
Ongoing Support	☐	☐

[Clear](#)

Submit

Figure 2 – Demo Form: Demonstrates how a company can fill out AI usage information.

3.3 Key Technical Challenges & Solutions

Frontend Complexity: Angular configuration was initially time-consuming. We mitigated this by using pre-built components.

Neural Seal

[Home](#) | [Result](#)

home works!

Neural Seal Home

Select Type:

Select Industry:

Figure showing our initial angular frontend page

Metric Standardization: Creating a universal scoring approach required cross-industry research; we divided the framework into “core” vs. “industry-specific” modules.

3.4 Website Link : <https://uplifted-function-846918.framer.app/>

3.5 Limitations & Future Development Needs : Companies often face challenges with complex AI supply chains—like subcontractors or outsourced processes—making comprehensive reporting difficult. There’s also a risk of companies gaming the system by underreporting or portraying AI usage more favorably, which necessitates credible audits. Additionally, consumer interpretations of “high AI usage” can vary, so labels must present neutral, factual information. Relying on manual data entry today highlights the need for future automated audits, while improvements in UI/UX and robust cloud infrastructure are essential for scalability.

4. Process

Timeframe	What will you do?
Next 3 months	Refine the MVP’s UI/UX and improve Angular codebase- Begin pilot testing with a small set of partner companies
2025	Integrate automated compliance checks (GDPR, EU AI Act) and explainable AI techniques (e.g., LIME, SHAP)- Expand into targeted sectors like finance & healthcare
2026	Seek standardization endorsements from global bodies (ISO, IEEE)- Scale infrastructure for international adoption
2027	Achieve widespread industry recognition as the go-to AI labeling standard- Embed AI usage labeling into large consumer-facing products/service globally

5. Impact on AI safety & key risks

5.1 Direct and Indirect Safety Benefits: A transparency-based label directly discourages negligent or harmful AI deployments by requiring companies to disclose their decision-making processes. Indirectly, increased public awareness and competitive pressure drive safer AI design, much like how vehicle emission disclosures spurred the development of greener cars (Bradley, 2023).

5.2 Potential Risks & Mitigation: Inaccurate self-reporting is mitigated by independent third-party audits and spot checks, while regulatory changes are managed through a modular scoring system that adapts to new laws (e.g., EU AI Act updates, national legislation). Additionally, label saturation is addressed with a clear, concise design supplemented by optional in-depth details via a QR code.

5.3 Alignment with Industry Safety Standards : Neural Seal will incorporate known frameworks (*ISO/IEC 27001:2022*, n.d.) and is designed to adapt as new AI-specific standards emerge.

5.4 Monitoring & Evaluation : Each company is required to update its disclosures annually, while users can report any discrepancies through a central portal (future ideas).

5.5 Success Metrics & KPIs : Neural Seal's impact is measured by the adoption rate among companies, consumer awareness through surveys on AI usage in daily products, and a reduction in harmful incidents by tracking reported misuses or bias claims in labeled products.

6. Appendix

6.1 Market Research Details : Surveys indicate strong consumer desire for clarity on AI-driven decisions, paralleling shifts in data privacy awareness (Massachusetts Institute of Technology, 2024).

6.2 References and Sources

1. Heaven, W. D. (2020, December 10). Google's medical AI was super accurate in a lab. Real life was a different story. *MIT Technology Review*.
<https://www.technologyreview.com/2020/04/27/1000658/google-medical-ai-accurate-lab-real-life-clinic-covid-diabetes-retina-disease/>
2. *Information and Resources about Cancer: Breast, Colon, Lung, Prostate, Skin*. (n.d.). American Cancer Society. <https://www.cancer.org/>
3. *The European Green Deal*. (2021, July 14). European Commission.
https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/european-green-deal_en
4. *LIME vs SHAP: A Comparative Analysis of Interpretability Tools*. (n.d.).
<https://www.markovml.com/blog/lime-vs-shap?t>
5. Nieto Juscafresa, A. (2022). An introduction to explainable artificial intelligence with LIME and SHAP. In *GRAU DE MATEMÀTIQUES*.
6. *EU AI Act: first regulation on artificial intelligence | Topics | European Parliament*. (2023, August 6). Topics | European Parliament.
<https://www.europarl.europa.eu/topics/en/article/20230601STO93804/eu-ai-act-first-regulation-on-artificial-intelligence>
7. *ISO/IEC 27001:2022*. (n.d.). ISO. <https://www.iso.org/standard/27001>
8. Bradley, A. (2023, February 10). How climate change policy has impacted the car industry. *Nationwide Vehicle Contracts*.
<https://www.nationwidevehiclecontracts.co.uk/blog/how-climate-change-policy-has-impacted-the-car-industry?t>
9. Massachusetts Institute of Technology. (2024, September 24). *Artificial intelligence disclosures are key to customer Trust | MIT Sloan Management Review*. MIT Sloan Management Review.
<https://sloanreview.mit.edu/article/artificial-intelligence-disclosures-are-key-to-customer-trust/?t>