**Innodata.**

Generative AI Test + Evaluation Beta Program

Innodata's beta program offers a comprehensive and rigorous evaluation process to assess the safety of AI models. With advanced testing methodologies, we identify potential vulnerabilities and risks, helping you proactively strengthen your AI's performance and reliability.

How It Works.

We connect your LLM's API endpoint to our evaluation platform, where we continuously test against your safety criteria. Our dashboard allows you to monitor how frequently the model encounters issues. We're offering this beta program to you free of charge to provide insights and support your AI governance efforts.

What We Need From You.

To get started, simply provide:

- **Your API details** and guidelines.

- **Conversation intents** reflecting typical use cases.

Program Phases

1 — Integration

We connect to your LLM to conduct simulated adversarial attacks and assess performance. Our streamlined API integration enables seamless testing.

2 — Testing Criteria Selection

Together, we define the scope, testing volume, and safety criteria. We recommend distributing conversation volume for balanced coverage and cost-efficiency.

Select intents and focus areas, such as toxicity, bias, and data leakage, ensuring prioritization of the key safety vectors that matter most to your application.

3 — Evaluation + Reporting

Receive a scorecard highlighting vulnerabilities, model deficiencies, and downloadable conversation logs for full transparency.

Get Started Today.

Begin with our AI governance framework for free in our beta program. Proactively test and evaluate your model's safety, setting a strong foundation for responsible AI deployment.

innodata.com