

G.A.B.R.I.E.L. (*Generative Agent for Brand-Ready, Realistic Image Engineering & Lighting*)

Celso Sousa

<https://www.linkedin.com/in/celso-sousa/>

1. Strategic context of G.A.B.R.I.E.L.

Generative imaging is becoming a core infrastructure for marketing, commerce, entertainment, product development, and digital communication. However, most organizations still use it as an isolated creative tool rather than as a controlled, measurable, and scalable production capability.

G.A.B.R.I.E.L. is designed to close this gap by creating an intelligent orchestration layer between human creative intent and image-generation models.

- **Strategic transformation:** Converts generative imaging from an experimental activity into a governed corporate capability;
- **Scientific foundation:** Combines multimodal AI, agentic systems, semantic modeling, visual reasoning, and automated optimization;
- **Business scalability:** Enables companies to produce larger volumes of high-quality visual assets without proportionally expanding creative teams;
- **Kunumi positioning:** Establishes a proprietary technology platform at the intersection of scientific research, AI engineering, and commercial content production.

2. Market problem and gaps in AI-powered visual generation

Current image-generation platforms can produce impressive outputs, but they remain difficult to control in professional environments.

- **Low predictability:** Small prompt variations can produce significant changes in composition, identity, lighting, proportions, and brand alignment;
- **High experimentation cost:** Creative teams often depend on repeated trial-and-error cycles to obtain an acceptable asset;
- **Limited consistency:** Maintaining characters, products, environments, visual identity, and photographic language across multiple images remains challenging;
- **Dependence on specialists:** High-quality outputs frequently require advanced knowledge of prompting, photography, design, lighting, and model behavior;
- **Weak corporate integration:** Most tools are not designed around approval workflows, brand rules, asset management, auditability, or performance measurement;

- **Lack of objective evaluation:** Organizations have limited mechanisms for automatically assessing visual quality, semantic compliance, brand consistency, and commercial readiness.

The market does not need another image generator. It needs an intelligent system capable of planning, controlling, evaluating, and improving visual generation.

3. General objective

Develop and scientifically validate an agentic platform capable of transforming business and creative objectives into structured visual specifications, optimized prompts, and production-ready images with greater realism, consistency, technical precision, and predictability.

G.A.B.R.I.E.L. will operate as a visual intelligence layer that plans the image, selects the appropriate generation strategy, evaluates the results, identifies deviations, and performs controlled refinement cycles.

4. Specific objectives

- Develop a structured visual representation covering composition, lighting, camera, lens, color, texture, atmosphere, narrative, and brand attributes;
- Create methods for translating subjective creative intentions into technically precise visual instructions;
- Implement specialized agents for planning, prompt engineering, generation, evaluation, refinement, and quality assurance;
- Support multiple image-generation models through a vendor-independent orchestration architecture;
- Develop automated metrics for semantic alignment, realism, aesthetics, brand consistency, and asset usability;
- Create memory mechanisms capable of preserving visual context across campaigns and image collections;
- Build proprietary datasets, evaluation protocols, and benchmarks for controlled visual generation;
- Validate the platform through high-value corporate use cases and measurable business outcomes.

5. Expected strategic outcomes

- A proprietary platform for scalable visual content production;
- A portfolio of enterprise solutions, APIs, specialized agents, and industry-specific products;
- Reduced dependence on individual prompting expertise and external creative production;
- Faster campaign development, product visualization, prototyping, and content localization;
- Proprietary datasets and benchmarks that improve performance and create defensible intellectual property;

- New scientific contributions in controllable generation, multimodal evaluation, agentic optimization, and visual consistency;
- A reusable technological foundation for future products involving video, animation, 3D assets, virtual environments, and multimodal experiences;
- Stronger international recognition of Kunumi as a company capable of transforming advanced AI research into scalable products.

6. Why invest in G.A.B.R.I.E.L. now?

The current market window is defined by the convergence of rapidly improving image models, growing enterprise adoption, increasing content demand, and persistent limitations in control and production reliability.

- **Demand is growing faster than production capacity:** Companies require continuous visual content for multiple audiences, channels, formats, and markets;
- **Foundation models are becoming accessible:** Competitive differentiation is shifting from access to models toward orchestration, proprietary data, workflows, evaluation, and domain intelligence;
- **The control layer remains underdeveloped:** The company that controls how models are instructed, evaluated, and integrated can capture value independently of the underlying model provider;
- **Enterprise standards are still being defined:** Kunumi has an opportunity to influence how controlled generative imaging is designed, governed, and measured;
- **Early investment creates cumulative advantage:** Every project can generate data, visual preferences, evaluation signals, templates, and optimization knowledge that strengthen the platform.

Delaying the initiative increases the risk of entering the market after visual-generation workflows have been consolidated by international platforms.

7. High-value corporate applications

- **Advertising and marketing:** Creation of campaign concepts, key visuals, social assets, audience variations, and localized content;
- **Brand management:** Generation of assets aligned with visual identity, photographic guidelines, tone, and brand positioning;
- **E-commerce:** Product scenes, lifestyle images, catalog variations, backgrounds, seasonal campaigns, and merchandising assets;
- **Financial services:** Personalized communication, educational content, product visualization, customer engagement, and internal campaigns;
- **Retail and consumer goods:** Packaging concepts, promotional assets, store communication, and product-launch simulations;
- **Entertainment and media:** Concept art, storyboards, characters, environments, promotional materials, and pre-production;
- **Industrial design and innovation:** Early visualization of products, interfaces, services, facilities, and future scenarios;
- **Corporate communication:** Presentations, reports, training materials, recruitment campaigns, and institutional content.

8. Value proposition for enterprises

G.A.B.R.I.E.L. delivers a controlled visual production capability rather than isolated image generation.

- **Greater production speed:** Reduces the time between creative briefing and usable visual output;
- **Lower operational cost:** Decreases manual experimentation, rework, external production dependency, and repetitive execution;
- **Higher quality:** Structures technical and aesthetic attributes before generation and automatically evaluates the resulting assets;
- **Brand consistency:** Preserves visual guidelines, campaign context, product characteristics, and creative direction;
- **Operational scalability:** Enables the production of multiple formats, versions, markets, and audience variations;
- **Knowledge retention:** Transforms individual creative expertise into reusable organizational intelligence;
- **Model flexibility:** Allows organizations to use different generation providers without rebuilding their complete production workflow.

9. Market potential and monetization models

G.A.B.R.I.E.L. can support multiple revenue streams and commercialization strategies.

- **Enterprise platform licensing:** Annual contracts for corporations with large-scale visual production needs;
- **Usage-based SaaS:** Subscription plans combined with generation, evaluation, or processing volume;
- **API and agent services:** Integration of visual planning and optimization capabilities into external products and workflows;
- **White-label solutions:** Customized platforms for agencies, media companies, retailers, and technology providers;
- **Industry-specific products:** Specialized versions for banking, retail, e-commerce, advertising, entertainment, and industrial design;
- **Professional services:** Implementation, model customization, workflow redesign, governance, and creative intelligence consulting;
- **Bradesco ecosystem applications:** Internal adoption across marketing, products, customer communication, innovation, and personalized experiences;
- **Technology partnerships:** Integration with cloud providers, model developers, design platforms, content-management systems, and digital agencies.

10. Operational gains and business impact metrics

The project will be managed through measurable business and technical indicators.

- **Efficiency:** Production time per asset, number of iterations, manual effort, rework rate, and cost per approved image;

- **Productivity:** Assets produced per professional, campaign throughput, and capacity to generate variations;
- **Quality:** Approval rate, visual consistency, semantic alignment, realism, defect frequency, and brand compliance;
- **Speed:** Time from briefing to first proposal, time to final approval, and campaign time-to-market;
- **Commercial impact:** Engagement, click-through rate, conversion, campaign performance, and content utilization;
- **Adoption:** Active users, recurring workflows, number of integrated teams, and asset reuse;
- **Platform economics:** Cost per generation, inference efficiency, model-routing performance, gross margin, and revenue per customer;
- **Scientific performance:** Benchmark improvement, evaluation correlation with human judgment, robustness, and cross-model generalization.

11. Competitive advantage and barriers to entry

The project's defensibility will not depend solely on access to a specific foundation model.

- **Proprietary visual representation:** A structured framework for converting creative objectives into machine-executable visual specifications;
- **Specialized agent architecture:** Coordinated agents for planning, generation, evaluation, correction, and governance;
- **Accumulated visual intelligence:** Data from prompts, outputs, human evaluations, corrections, and business results;
- **Corporate workflow integration:** Connection with brand guidelines, asset repositories, approval systems, and campaign processes;
- **Model-independent orchestration:** Ability to select and combine models according to quality, cost, latency, privacy, and use case;
- **Scientific benchmarks:** Proprietary methods for measuring realism, alignment, consistency, and production readiness;
- **Domain-specific expertise:** Visual knowledge adapted to sectors, products, audiences, cultural contexts, and regulatory requirements.

These assets create a cumulative moat based on data, science, engineering, processes, and customer integration.

12. Exploratory innovation with disruptive potential

G.A.B.R.I.E.L. creates a foundation for visual systems that move beyond prompt generation.

- Agents capable of analyzing a business objective and autonomously proposing a complete visual strategy;
- Automatic decomposition of campaigns into scenes, formats, variations, and audience-specific assets;
- Multimodal agents capable of inspecting generated images, detecting failures, and initiating targeted correction cycles;

- Persistent visual memory for maintaining characters, products, environments, and artistic direction across projects;
- Optimization based on campaign performance, allowing the platform to learn which visual characteristics generate better business outcomes;
- Coordinated generation across text, image, video, audio, 3D, and interactive environments;
- Simulation of virtual focus groups to assess visual perception before real-world publication;
- Generative systems capable of balancing creativity, brand identity, cost, compliance, and conversion goals.

13. Long-term strategic advantage for Kunumi

G.A.B.R.I.E.L. can become more than a single product. It can become the visual intelligence foundation of a broader multimodal portfolio.

- **Product expansion:** The architecture can evolve toward video, animation, virtual presenters, 3D visualization, and immersive experiences;
- **International scalability:** Visual-production challenges exist across industries, languages, and geographic markets;
- **Recurring revenue:** Enterprise integration, accumulated context, and proprietary workflows support long-term contracts;
- **Research-to-market engine:** The project creates a repeatable process for converting scientific advances into commercial capabilities;
- **Strategic independence:** Kunumi can orchestrate multiple foundation models while retaining ownership of customer experience, data, evaluation, and business logic;
- **Brand leadership:** A successful platform strengthens Kunumi's ambition to become the leading AI company in Latin America;
- **Bradesco synergy:** The initiative can generate immediate internal use cases while creating products for external commercialization.

14. Scientific foundation of generative visual engineering

The project will investigate visual generation as a structured optimization problem rather than as an informal prompting activity.

- Semantic representation of creative intent and visual requirements;
- Hierarchical modeling of scenes, objects, relationships, styles, and narrative context;
- Computational representation of photography, lighting, composition, color, texture, and depth;
- Alignment between natural-language instructions and visual outputs;
- Multimodal reasoning for identifying inconsistencies and proposing corrections;
- Human preference modeling for aesthetics, realism, and brand suitability;
- Multi-objective optimization balancing quality, fidelity, creativity, cost, latency, and safety;
- Contextual memory for preserving visual decisions across iterative workflows.

This scientific foundation will support explainable, testable, and continuously improvable visual generation.

15. State of the art in image generation and multimodal agents

Current approaches combine foundation models for image generation, vision-language models, prompt enhancement, reference-image conditioning, editing mechanisms, and agentic workflows.

Despite rapid progress, important gaps remain:

- Prompt optimization is commonly treated as a generic language task rather than a specialized visual-engineering discipline;
- Most platforms provide limited reasoning about photographic and compositional decisions;
- Automated evaluation remains insufficiently correlated with human and business judgment;
- Cross-image consistency remains difficult in complex campaigns and long production sequences;
- Systems are usually tied to individual models and offer limited model-routing intelligence;
- Brand guidelines, production workflows, governance, and performance data are rarely integrated into a unified architecture;
- Few solutions connect scientific experimentation directly to product, market, and operational metrics.

G.A.B.R.I.E.L. addresses these gaps through an integrated agentic platform.

16. Scientific and technological differentiators

- A structured visual language capable of representing creative intent with technical precision;
- Hierarchical prompt construction based on scene, subject, camera, lighting, composition, material, atmosphere, and narrative;
- Specialized visual agents rather than a single general-purpose assistant;
- Multimodal evaluation loops that compare requirements with generated results;
- Automated diagnosis of visual failures and targeted prompt correction;
- Persistent memory for campaign, character, product, and brand consistency;
- Dynamic model selection based on task requirements and operational constraints;
- Human-in-the-loop mechanisms for capturing expert feedback and improving the system;
- Proprietary benchmarks connecting perceptual quality with production and business outcomes;
- Continuous experimentation as a central component of product evolution.

17. Research hypotheses

- Structured visual planning will produce higher semantic alignment than unstructured natural-language prompting;
- Specialized agents will reduce the number of iterations required to achieve an approved image;
- Automated multimodal evaluation will identify relevant visual defects with meaningful agreement with expert human evaluation;
- Persistent context and memory will improve consistency across image collections and campaign assets;
- Domain-specific visual representations will outperform generic prompt-enhancement approaches in enterprise applications;
- Multi-model orchestration will improve the balance between quality, cost, latency, and reliability;
- Human feedback combined with experimental optimization will create cumulative performance gains over time;
- Brand-aware generation will increase approval rates and reduce downstream editing and rework.

18. Conceptual architecture of the solution

- **Intent and briefing layer:** Captures objectives, audience, channel, brand, message, constraints, and expected outcome;
- **Visual planning layer:** Decomposes the request into scenes, subjects, relationships, composition, camera, lighting, color, and narrative;
- **Agentic orchestration layer:** Coordinates specialized agents, tools, models, memory, and decision policies;
- **Prompt-engineering layer:** Produces model-specific instructions, negative constraints, parameters, and reference strategies;
- **Generation layer:** Connects multiple image models, editing tools, and enterprise infrastructure;
- **Evaluation layer:** Measures semantic alignment, realism, aesthetics, consistency, safety, and brand compliance;
- **Refinement layer:** Diagnoses deviations and performs controlled improvement cycles;
- **Governance layer:** Manages access, policies, copyright controls, provenance, logs, approvals, and auditability;
- **Integration layer:** Connects the platform with creative tools, asset repositories, APIs, campaign systems, and corporate workflows;
- **Analytics layer:** Tracks technical performance, operational efficiency, user behavior, and business outcomes.

19. Research methodology and experimental validation

The project will follow an evidence-driven methodology that connects scientific hypotheses with commercial deliverables.

- Establish baselines using current prompting practices, expert workflows, and leading generation platforms;

- Build representative datasets covering products, people, environments, lighting conditions, industries, and visual styles;
- Develop controlled experiments comparing generic prompts, expert prompts, and G.A.B.R.I.E.L.-optimized workflows;
- Combine automated metrics with blind human evaluation performed by designers, photographers, marketers, and business users;
- Measure semantic alignment, realism, consistency, approval rate, iteration count, production time, and cost;
- Conduct corporate pilots in high-value scenarios with clear operational and commercial indicators;
- Use explicit criteria for model promotion, rollback, risk mitigation, and transition to production;
- Publish scientific findings when strategically appropriate while protecting proprietary methods, datasets, and intellectual property;
- Convert validated research components into reusable platform modules and commercial products.

20. Governance, risks, scale roadmap, and investment rationale

- **Governance by design:** The platform will include access controls, audit logs, approval workflows, model documentation, content provenance, and responsible-use policies;
- **Intellectual property protection:** Contracts, datasets, model configurations, prompts, outputs, and training processes will be managed according to ownership and licensing requirements;
- **Risk management:** Technical, scientific, legal, reputational, operational, and market risks will be monitored through staged validation and controlled deployment;
- **Human oversight:** High-impact or brand-sensitive content will remain subject to defined review and approval mechanisms;
- **Initial roadmap:** Establish the scientific core, baseline architecture, proprietary visual representation, and first specialized agents;
- **Validation roadmap:** Execute controlled benchmarks and pilots in selected Bradesco and enterprise use cases;
- **Productization roadmap:** Strengthen security, scalability, integrations, observability, service levels, and user experience;
- **Commercial scale roadmap:** Launch vertical solutions, APIs, strategic partnerships, and international commercialization;
- **Investment rationale:** A program with potential annual funding of up to R\$100 million is justified by the opportunity to build a defensible multimodal platform, generate multiple products, secure recurring enterprise revenue, and establish long-term technological leadership;
- **Strategic urgency:** Investment must occur while the orchestration and control layer of generative imaging remains open for leadership. The objective is not to follow the visual AI market, but to define its enterprise standard.