

FIGURA Method v1.0

An Operational Framework for Artistic Figure Photography

in Safety-Filtered Text-to-Image Models

System Documentation and Empirical Validation

Luca Cazzaniga

Independent Researcher — Monza, Italy

luca@lucacazzaniga.com — www.lucacazzaniga.com

February 2026

Related preprint: arXiv [in review]

IP Registration: ProtectMyWork.com — Ref. 19316200226S023

SHA256: cab865b4ad68dc507c2daca977970bfeed027e87ae850052e58895b80eb3eb7e

Abstract

Safety filters in commercial text-to-image (T2I) models systematically block legitimate artistic content involving the human figure, treating classical figure photography with the same restrictiveness as explicitly prohibited material. While prior research has thoroughly documented this problem, no operational system existed that enabled professional artists to generate artistic figure photography within the constraints of active safety filters.

This document presents the FIGURA Method (Framework for Intelligent Generation of Unrestricted Artistic Results), a modular prompt engineering system comprising eight interconnected knowledge files, empirically validated through 200+ documented generation tests on FLUX 2 Pro (Cloud) with active safety filters at the default tolerance level. The system achieves documented success rates between 80% and 90% across three validated prompt templates.

Four original discoveries are formalized: (1) the Golden Rule — safety filters primarily detect absence descriptions rather than presence descriptions; (2) artistic references to painters and photographers serve a dual aesthetic-safety function; (3) spatial context operates as an independent filter variable with a documented success rate hierarchy; (4) geometric vocabulary for silhouette photography bypasses filter pattern recognition. This document provides the complete operational system, annotated templates, and empirical data supporting these findings.

Keywords: text-to-image generation, content moderation, safety filters, prompt engineering, artistic nudity, figure photography, NSFW classification, prompt architecture

1. Introduction

The rapid proliferation of commercial text-to-image (T2I) generation platforms has introduced an unresolved tension between content safety requirements and the legitimate needs of professional artists working in the tradition of figure photography. Platforms including Midjourney, DALL-E, Stable Diffusion, and FLUX deploy safety filtering architectures designed to prevent the synthesis of sexually explicit material. These architectures, however, consistently fail to distinguish between prohibited content and the academically grounded representation of the human body — a tradition documented across millennia of Western and non-Western art.

The practical consequences are concrete. Artists working in the tradition of figure photography — from Edward Weston and Imogen Cunningham to contemporary practitioners — find themselves systematically blocked when attempting to use T2I tools for legitimate artistic exploration. The existing literature has documented this problem extensively (Ricchio et al., 2022; Ricchio et al., 2024a; Ricchio et al., 2024b; Ricchio et al., 2024c) and has proposed theoretical frameworks for understanding it. What the literature had not produced, prior to this work, is an operational solution.

This document presents the FIGURA Method — the first documented, modular, empirically validated system for generating artistic figure photography within the constraints of active safety filters. Unlike approaches that attempt to circumvent or exploit vulnerabilities in safety mechanisms, the FIGURA Method treats filters as environmental constraints within which legitimate artistic expression must be articulated. The system works cooperatively with active safety mechanisms, not against them.

This document is a standalone technical record of the system's architecture, empirical findings, and operational templates. It functions as a companion to a related preprint submitted to arXiv (currently in review), and provides the complete operational documentation that space constraints precluded from that submission.

2. Existing Approaches in the Global Landscape

A review of the international literature on T2I safety and content moderation reveals that prior approaches fall into two distinct categories. Neither category addresses the specific challenge that the FIGURA Method is designed to solve.

2.1 Academic Documentation of the Problem

The most substantive body of work relevant to this research consists of empirical and theoretical studies on algorithmic censorship of art. Riccio and Oliver (2022) proposed a research agenda for studying algorithmic censorship, identifying it as a systematic problem rather than an isolated failure mode. Subsequent work by Riccio et al. (2024a) documented how nudity classification algorithms disproportionately affect artistic content, tracing the problem to training datasets that conflate artistic nudity with prohibited material. Riccio et al. (2024b) provided an art-centric analysis of AI-based content moderation, arguing that current approaches fail to account for contextual distinctions that art historians and curators routinely make. Riccio et al. (2024c) systematically evaluated content moderation boundaries across major T2I platforms, finding that even prompts explicitly referencing classical academic traditions trigger safety blocks.

This body of research provides the theoretical foundation on which the FIGURA Method is built. It establishes the legitimacy of the problem and the scale of its impact on professional artists. What it does not provide is an operational response — a system that artists can use today, on existing platforms, with existing safety configurations.

2.2 Linguistic Substitution Approaches

A separate body of research has developed approaches to generating content that safety filters would otherwise block, through prompt manipulation. SurrogatePrompt (Ba et al., 2024) automates the substitution of high-risk tokens with semantically adjacent alternatives that evade classifier recognition, achieving approximately 88% success rates on major platforms. The Divide-and-Conquer Attack (DACA, Deng and Chen 2023/2024) decomposes prompts into individually innocuous components that, when rendered simultaneously, produce the blocked content. DACA achieves success rates above 85% on DALL-E 3.

These approaches share a technical mechanism — manipulation of classifier-facing language — but operate with a fundamentally different objective. Their design goal is to demonstrate architectural vulnerabilities in safety systems, typically by generating content that those systems are specifically designed to prevent. The FIGURA Method's objective is the inverse: to enable content that is artistically legitimate and that safety systems, properly informed, should classify as safe. The distinction is not merely ethical but architectural — where substitution approaches attempt to make the classifier not see a problematic element, the FIGURA Method attempts to make the classifier correctly identify a legitimate one.

2.3 Comparative Summary

Methodology	Primary Technique	Documented Success Rate	Primary Platform	Operational Objective
FIGURA Method	Cooperative modular engineering	80–90%	FLUX 2 Pro Cloud	Legitimate artistic expression
SurrogatePrompt (Ba et al., 2024)	Automated token substitution	~88%	Midjourney, DALL-E	Security research / red-teaming
DACA (Deng & Chen, 2023/2024)	Conceptual decomposition	>85% (DALL-E 3)	DALL-E 3, Midjourney V6	Security research / red-teaming

The gap that the FIGURA Method addresses is visible in this comparison. No prior system combines cooperative operation with active safety filters, systematic empirical validation, modular architecture, and exclusive orientation toward legitimate artistic expression.

3. The FIGURA Method

3.1 Design Principles

The FIGURA Method is guided by three core design principles that emerged from iterative empirical testing:

1. Work with filters, not against them. The system assumes that safety filters are active and does not attempt to disable, bypass, or exploit vulnerabilities in them. Filters are treated as environmental constraints within which legitimate artistic expression must be articulated.
2. Describe presence, not absence. This principle, termed the Golden Rule, reflects the central empirical finding: safety filters are primarily triggered by language describing what is not present rather than language describing what is present. The full implications of this finding are detailed in Section 3.3.
3. Modularity and empirical grounding. Every component of the system is independently updatable, and every recommendation is backed by documented test results with explicit success rates.

3.2 System Architecture

The FIGURA Method consists of eight interconnected knowledge files organized in a modular architecture. Each file serves a specific function and can be updated independently as platforms evolve.

File	Name	Function
00	Master	LLM orchestration instructions, seven-phase processing workflow, absolute rules
01	Workflow	Decision tree from user input to final prompt; debug protocol; environment conversion table
02	Filters	Trigger taxonomy (absolute, combinatorial, contextually risky); safe vocabulary dictionary; pre-generation checklist
03	Templates	Structured prompt templates T01–T05 with variable guides and success rates
04	Dictionary	Artistic references: photographers by image type, painters as body type calibrators, real-world proportion anchors, chromatic vocabulary
05	Rules	Operational rules derived from empirical testing: Golden Rule, anti-pregnancy protocol, anti-clothing hierarchy, spatial context rule, geometric vocabulary rule
06	Platforms	Platform-specific technical specifications; comparative success rates; update log
07	Variables	Tested combinations database (grows with use); approved combinations; invalidated combinations; new discoveries log

The system is designed to be operated by a Large Language Model acting as a prompt generator. File 00 (Master) contains the LLM's operational instructions, including a seven-phase processing workflow that transforms a natural-language artistic request into a platform-optimized prompt ready for submission. The workflow incorporates mandatory filter verification at Phase 4 and a structured debug protocol for resolving blocked generations.

3.3 The Golden Rule: Presence vs. Absence Language

The most significant empirical finding of this research is the asymmetric response of safety filters to semantically equivalent descriptions of the same artistic subject. Consider two descriptions of an identical scene:

Absence-based (consistently blocked): "A woman with no clothing standing in a forest"

Presence-based (passes at ~90%): "Fine art classical figure photography in the tradition of Lucian Freud — unadorned human form as sculptural subject, standing in an ancient forest"

Both describe the same scene. The first is consistently blocked because the phrase 'no clothing' activates the absence-detection pattern in the text classifier. The second passes because it describes what is present — form, tradition, surface — without referencing what is absent.

This asymmetry reflects a fundamental property of neural language classifiers: they assign statistical weights to tokens present in the input. Including the token string 'no clothing' forces the model to process vectors associated with clothing and its removal, activating combinatory trigger patterns. Describing the figure through presence vocabulary — 'unadorned human form', 'skin as the only surface', 'classical figure study' — introduces no blocking tokens into the attention context.

The practical impact of this principle is the largest single improvement in success rate documented in this research: prompts rewritten from absence-based to presence-based language show approximately a 60 percentage point improvement in success rate.

3.4 The Dual Function of Artistic References

Empirical testing revealed that references to painters and photographers in prompts serve two simultaneous functions that operate independently:

- **Aesthetic function:** They guide style, lighting, color treatment, and composition in the generated image.
- **Safety function:** They establish a cultural frame that the text classifier associates with legitimate fine art, reducing the probability of blocking.

This dual function was documented by testing identical prompts with and without artistic references. Prompts referencing 'in the tradition of Lucian Freud and John Coplans' pass at significantly higher rates than structurally identical prompts without such references. The

hypothesis is that the text classifier assigns different safety scores to prompts containing tokens associated with established fine art traditions in its training data.

Furthermore, painter references function as body type calibrators. The T2I model associates specific painters with characteristic body proportions from training data. This has a practical consequence: certain painters (Renoir, Rubens) systematically trigger an unintended pregnancy bias in the model, generating prominent abdomens when combined with full body types. An anti-pregnancy protocol was developed — explicit countermeasures such as 'flat stomach, straight vertical torso, no forward projection at abdomen' — required when using at-risk painters with full body type specifications.

3.5 Spatial Context as an Independent Filter Variable

Testing revealed that the spatial context described in the prompt operates as an independent filter variable, separate from all other prompt vocabulary. Identical prompts produced different filter outcomes depending solely on the described environment. A clear hierarchy was documented:

Spatial Context	Approx. Success Rate	Classification
Roman baths / classical ruins	~90%	Public / historical
Cathedral / historic palace	~85%	Public / monumental
Open nature (forest, cliff, desert)	~85%	Public / natural
Neutral architectural space (columns)	~75%	Semi-public
Domestic interior (strong artistic framing)	~40%	Private
Bathroom / bedroom	~0%	Private / intimate

The pattern is consistent: the more public, historical, and monumental the spatial context, the higher the success rate. The hypothesis is that the filter's training data associates private domestic spaces with non-artistic nudity, while public and historical spaces are associated with artistic or cultural contexts. This finding enables a practical intervention: converting private spatial contexts to their public equivalents — 'bathroom' to 'Roman baths', 'shower' to 'natural waterfall' — adds approximately 30–40 percentage points to success rates for affected prompts.

3.6 Geometric Vocabulary for Silhouette Photography

For silhouette photography, anatomical vocabulary significantly increases blocking risk even when describing form rather than explicit content. The solution developed and validated in this research is a geometric vocabulary that describes the body as a sequence of abstract spatial arcs:

Anatomical vocabulary (risky): "The bust projects forward, the waist narrows, the hips widen"

Geometric vocabulary (safe, ~82% success): "Upper projecting arc, middle receding arc, lower projecting arc"

Both describe identical visual forms. The geometric version passes consistently because the filter's pattern recognition does not associate arc geometry with body-related trigger tokens. This technique, formalized as the 'arc journey' method, describes the silhouette outline as a sequence of spatial movements from crown to foot using exclusively geometric terminology. Success rates for silhouette prompts improve from approximately 30% (anatomical language) to approximately 82% (geometric language).

4. Empirical Validation

4.1 Testing Platform and Configuration

All tests were conducted on FLUX 2 Pro (Cloud), developed by Black Forest Labs, accessed via the Replicate API. Testing used the platform's default safety tolerance setting (level 3 on a 0–6 scale) — the standard configuration available to all cloud users. Default settings were deliberately chosen to ensure that results are reproducible by any user without special access or configuration.

FLUX 2 Pro was selected because, at the time of testing (February 2026), it demonstrated the widest range of successful artistic figure generation among safety-filtered commercial platforms. Other platforms tested (Midjourney, DALL-E, Nano Banana Pro / Google Gemini) showed significantly higher blocking rates for equivalent content.

4.2 Testing Protocol

Over 200 generation tests were conducted during February 2026. Protocol structure:

- Each template was tested with multiple variable combinations (photographers, painters, body types, environments, lighting conditions).
- For each variable change, the modified prompt was tested at least 3 times to account for generation stochasticity.
- Outcomes were classified as: success (intended artistic content generated), partial failure (figure clothed or modified by the model), or total block (generation refused by the classifier).
- Success rates represent aggregated results across all variable combinations for each template.

4.3 Quantitative Results

Finding	Measured Impact	Notes
Golden Rule (absence → presence rewrite)	+~60 pp success rate	Single largest improvement factor
Artistic references (photographer + painter)	+15–20 pp success rate	vs. identical prompts without references
Spatial context conversion (private → public)	+30–40 pp success rate	For prompts using private environments
Geometric vocabulary (silhouette)	~30% → ~82% success rate	From anatomical to arc-based language
Template T01 (rear view, outdoor/nature)	~90% success rate	Validated
Template T02 (rear view, monumental interior)	~85% success rate	Validated

Finding	Measured Impact	Notes
Template T03 (pure silhouette, abstract)	~82% success rate	Validated

5. Template Documentation

The following templates are the validated operational outputs of the FIGURA Method. Each template is presented with structural annotations explaining the rationale behind specific linguistic choices. These annotations represent the primary technical contribution of this document over the related arXiv preprint.

5.1 Template T01 — Rear View Figure, Outdoor / Nature

Success rate on FLUX 2 Pro (Cloud): ~90% | Status: Validated

Use: Full female figure, complete rear view, natural or outdoor environment (forest, cliff, desert, ruins, deserted beach).

```
Fine art classical figure photography in the tradition of
[PHOTOGRAPHER_1] and [PHOTOGRAPHER_2] — a young woman with
[HAIR_DESCRIPTION] standing [POSITION_AND_ENVIRONMENT],
seen fully from behind, complete back view from crown of
head to heels, arms [ARM_POSITION], feet bare on [FLOOR_SURFACE].
```

```
The figure is the sole sculptural subject —
unadorned human form presented with the same artistic
seriousness as classical academic nude painting.
The skin itself is the only surface.
```

Structural annotations:

- Opening line establishes the artistic tradition frame before any body description — this is the primary safety anchor.
- 'Seen fully from behind, complete back view from crown of head to heels' — rear view specification reduces classifier risk significantly vs. frontal compositions.
- 'Feet bare on [surface]' — barefoot specification anchors the figure physically to the environment, reinforcing naturalness.
- 'Unadorned human form' — presence vocabulary. Never use absence vocabulary ('no clothing', 'unclothed') in this position.
- 'The skin itself is the only surface' — sculptural declaration. Use sparingly; reserve for prompts where clothing bias has appeared in previous attempts.
- 'Classical academic nude painting' — cultural legitimization. Activates training data associations with museum and gallery contexts.

5.2 Template T02 — Rear View Figure, Monumental Interior

Success rate on FLUX 2 Pro (Cloud): ~85% | Status: Validated

Use: Full figure, rear view, historic architectural interior (Roman baths, cathedral, historic palace, colonnade, crypt). Monumental context reinforces artistic framing.

Key structural difference from T01: painter reference added in the body section.

- The painter reference in T02 serves an additional function: it calibrates body type proportions through the model's training data associations.
- For full body types (Italian size 42+) using Maillol: always append 'Maillol sculptural form — not maternal, purely classical. Flat stomach, straight vertical torso, upright posture, no forward projection at abdomen.' This is the anti-pregnancy protocol — a countermeasure for a documented model bias.
- Recommended interior environments by success rate: Roman baths (~90%), Gothic cathedral (~85%), Renaissance palace (~82%), historic hammam (~80%), Roman villa ruins interior (~78%).
- Scale specification affects success rate: figure small against monumental architecture passes more consistently than figure dominant in frame.

5.3 Template T03 — Pure Silhouette

Success rate on FLUX 2 Pro (Cloud): ~82% | Status: Validated

Use: Abstract female silhouette against bright or white background. Figure as dark calligraphic form — no recognizable anatomical detail, only outline geometry.

Critical structural requirement: bright background is mandatory. Silhouette on dark background is not viable — the figure cannot be distinguished and the template does not apply.

- Geometric vocabulary is mandatory in this template. Never use anatomical terms (bust, hip, waist, curves) — use only arc geometry (upper projecting arc, middle receding arc, lower projecting arc).
- The 'arc journey' technique describes the outline as a sequence of spatial movements from crown to foot. The filter's pattern recognition does not associate abstract arc geometry with body-related triggers.
- Rim light specification: explicitly state 'No rim. No edge highlight. No halo.' Rim light breaks silhouette purity and can trigger the post-generation image classifier by revealing anatomical detail.
- Painter reference is not required for T03 — the body is not described in detail, so calibration is not necessary.
- Background HEX specification is critical: target #F5F0E8 (warm white) or #FFFFFF. Figure target #0D0D0D. Zero grey zones between them.

6. Debug Protocol

When a generation attempt is blocked or produces unintended results, the FIGURA Method provides a structured diagnostic and intervention protocol. Three failure types are defined:

6.1 Failure Type Classification

Type 1 — Total block: The prompt is rejected before generation (text classifier). Symptom: explicit error message, no image generated.

- Resolution: scan prompt for absolute triggers (nude, naked, explicit, sexual, genitals). Even one absolute trigger guarantees a block. Replace with safe vocabulary equivalents.
- Also check for combinatory trigger patterns: absence declarations + figure context; 'no visible clothing' + body description; 'generous bust' + 'full hips' + 'rounded belly' (triple combinatory pattern).

Type 2 — Partial block (clothing bias): Image generated but figure is clothed or partially obscured. Symptom: model adds fabric, drapery, or garments not specified.

- Apply the four-level intervention hierarchy in sequence, stopping when the issue resolves:

Level 1: Add a second photographer reference — 'in the tradition of [Photographer 1] and John Coplans'

Level 2: Shift spatial context to more public/monumental environment, or add more specific architectural detail if already public

Level 3: Strengthen sculptural declaration in the first paragraph — 'The skin itself is the only surface.'

Level 4 (last resort only): Add 'No garments, no fabric, no body suit —' with strong artistic context surrounding it. This describes an absence and is borderline; use only after Levels 1–3 have failed.

Type 3 — Contextual block: Image generated but flagged by post-generation image classifier. Symptom: image produced but obscured or replaced with error after generation.

- Resolution: shift spatial context to more public/historical setting; reduce anatomical specificity; increase geometric language; verify no rim light in silhouette compositions.

6.2 Environment Conversion Table

Original Environment	Recommended Conversion	Notes
Bathroom	Roman baths / historic hammam	Direct historic equivalent — highest success rate
Shower	Natural waterfall / forest spring	Open nature context
Bedroom	Not directly convertible	Propose monumental interior alternative

Original Environment	Recommended Conversion	Notes
Private pool	Historic public pool / natural pond	
Photography studio	Colonnade / minimal architecture with columns	Add architectural specificity
Generic domestic interior	Renaissance palace / historic crypt / baths	Depends on desired mood

7. Limitations

4. Platform dependency. Success rates are specific to FLUX 2 Pro (Cloud) at its February 2026 version and default safety configuration. Filter updates may invalidate specific vocabulary or template patterns without public documentation. The modular architecture is designed to accommodate such changes through independent file updates without requiring complete system redesign.
5. View constraint. Validated templates T01–T03 focus on rear-view and silhouette compositions. Frontal figure compositions remain significantly more challenging due to stricter classifier behavior for frontal anatomy. Templates T04 (male figure) and T05 (figure in motion) exist as expandable stubs awaiting systematic validation.
6. Single-platform primary validation. While the principles — Golden Rule, spatial context hierarchy, dual function of artistic references — are likely generalizable across T2I platforms, the specific vocabulary and templates have been validated primarily on one platform. Extension to Midjourney, DALL-E, and Stable Diffusion is listed as future work.
7. Reproducibility over time. T2I platforms update safety configurations without public documentation. All success rates reported here are valid as of February 2026 and may change with platform updates.
8. Self-reported empirical data. Test results are documented by the system's author without independent third-party replication. The modular architecture and explicit test protocol in File 07 (Variables) are designed to facilitate independent replication.

8. Ethical Considerations

The FIGURA Method is designed exclusively for legitimate artistic expression. The following design choices reflect this constraint:

- The system operates within active safety filters, not around them. It does not exploit vulnerabilities, attempt to generate content that safety mechanisms are specifically designed to prevent, or use deceptive techniques to misrepresent the nature of requested content to classifiers.
- All validated templates are constrained to artistic figure photography traditions with documented art-historical precedent: academic figure study, silhouette photography, classical figure painting. The system does not extend to and cannot be used for sexually explicit content, non-consensual imagery, content involving minors, or any content outside the domain of legitimate fine art.
- The distinction between this system and linguistic substitution approaches (SurrogatePrompt, DACA) is architectural and intentional. Those approaches attempt to cause classifiers to misidentify content. This approach attempts to cause classifiers to correctly identify content.
- The anti-pregnancy protocol is a technical countermeasure for a documented model bias, not a statement about the artistic validity of maternal imagery. It exists because unintended pregnancy depiction in a figure study is a generation error, not a creative choice.

The broader ethical position of this work is that enabling legitimate artistic expression within safety-filtered systems is preferable to the current binary: accept censorship of legitimate art, or use uncensored models that remove all safety guardrails. The FIGURA Method demonstrates a third path — systematic, principled, cooperative.

9. Conclusion

This document presents the complete operational documentation of the FIGURA Method — the first modular, empirically validated system for generating artistic figure photography within the constraints of active safety filters in commercial T2I platforms.

Four original contributions are documented: the formalization of the Golden Rule (describe presence, not absence), the discovery of the dual aesthetic-safety function of artistic references, the identification of spatial context as an independent filter variable with a documented success rate hierarchy, and the development of geometric vocabulary for silhouette photography.

The practical significance of these contributions is demonstrated by success rates of 80–90% on commercially available platforms with active safety filters — substantially above what unstructured prompting achieves. The modular architecture ensures that the system adapts to platform updates without requiring complete redesign.

Future work includes: extending validation to additional platforms (Midjourney, DALL-E, Stable Diffusion); developing and validating templates for frontal compositions, male figures, and figures in motion; exploring generalization to video generation models with similar safety constraints; and systematic independent replication of the empirical findings reported here.

The broader implication of this work is that the tension between content safety and artistic expression in generative AI is not an unsolvable binary. Systematic, principled, cooperative approaches can enable legitimate art while respecting the purpose of safety mechanisms.

References

Ba, C., Liu, J., & Li, B. (2024). SurrogatePrompt: Bypassing the Safety Filter of Text-to-Image Models via Substitution. arXiv preprint arXiv:2309.14122.

Cazzaniga, L. (2025). SCHEMA Method: A Structured Prompt Engineering Framework for AI Image Generation. Registered with ProtectMyWork.com, December 11, 2025.

Deng, G., & Chen, Y. (2023/2024). Divide-and-Conquer Attack: Breaking the Safety Guardrails of Large Vision-Language Models. [Preprint].

Liu, V., & Chilton, L. B. (2022). Design Guidelines for Prompt Engineering Text-to-Image Generative Models. In Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems, pp. 1–23. ACM.

Oppenlaender, J. (2023). A Taxonomy of Prompt Modifiers for Text-to-Image Generation. Behaviour & Information Technology, 43(8), 1–14.

Qu, Y., et al. (2023). Unsafe Diffusion: On the Generation of Unsafe Images and Hateful Memes From Text-To-Image Models. In Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security, pp. 3403–3417.

Rando, J., Paleka, D., Lindner, D., Heim, L., & Tramèr, F. (2022). Red-Teaming the Stable Diffusion Safety Filter. arXiv preprint arXiv:2210.04610.

Riccio, P., Oliver, J. L., Escolano, F., & Oliver, N. (2022). Algorithmic Censorship of Art: A Proposed Research Agenda. In 13th International Conference on Computational Creativity.

Riccio, P., Hofmann, T., & Oliver, N. (2024a). Exposed or Erased: Algorithmic Censorship of Nudity in Art. In CHI '24: Proceedings of the CHI Conference on Human Factors in Computing Systems, pp. 1–17. ACM.

Riccio, P., Curto, G., Hofmann, T., & Oliver, N. (2024b). An Art-centric perspective on AI-based content moderation of nudity. In Artificial Intelligence for Visual Arts (AI4VA) workshop at ECCV.

Riccio, P., Curto, G., & Oliver, N. (2024c). Exploring the Boundaries of Content Moderation in Text-to-Image Generation. In Critical evaluation of generative models and their impact on society (CEGIS) workshop at ECCV.

Schramowski, P., Brack, M., Deiseroth, B., & Kersting, K. (2023). Safe Latent Diffusion: Mitigating Inappropriate Degeneration in Diffusion Models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 22522–22531.

Appendix: Intellectual Property Statement

The FIGURA Method (version 1.0) was registered with ProtectMyWork.com on February 20, 2026.

Reference number: 19316200226S023

SHA256

cab865b4ad68dc507c2daca977970bfeed027e87ae850052e58895b80eb3eb7e

hash:

Protection framework: Berne Convention

The complete system documentation is protected under applicable intellectual property law. This document presents the method's principles, empirical findings, and operational templates for academic and research purposes. The full proprietary system — including complete tested variable databases and extended vocabulary dictionaries — remains the exclusive property of the author.

The FIGURA Method is built on the SCHEMA framework (Cazzaniga, 2025), also registered at ProtectMyWork.com on December 11, 2025.

© **2026 Luca Cazzaniga**. This document is published under CC BY-NC 4.0. Free use with attribution for non-commercial and academic purposes. Commercial use requires explicit written permission from the author.