

Provenance Granularity and Executable Constraints in a Fragment-Based Computational Memorial

Tomas Laurenzo
Department of Critical Media Practices
University of Colorado Boulder
Boulder, Colorado, USA
tomas@laurenzo.net

Abstract—*Desaparecidos.uy* is an artwork in development and an ongoing research project that forms portraits of people detained-disappeared in connection with Uruguay from fragments of contemporary people, places, and street-level traversals. Implementing its three memorial propositions exposed a computing problem: diversity constraints can appear to succeed while operating at the wrong provenance unit. In a 432-run proxy benchmark, a sequential six-dimensional nearest-neighbor matcher improved structural similarity while increasing row-level source concentration. A one-tile row cap reduced that concentration by construction, yet one parent image could still provide 16% of an output; under the same greedy assignment, the cap reduced structural similarity by 9.2–18.1%. The project responds with inspectable fragment-to-source mappings, placement histories, feasibility checks, and work-specific constraints. These controls make limited technical conditions visible and contestable without claiming consent, anonymity, historical adequacy, or memorial value.

Keywords—computational art, digital memorial, provenance, source participation, artistic research, cultural heritage

I. INTRODUCTION

Enforced disappearance attacks not only a person but the possibility of establishing their fate. The International Convention for the Protection of All Persons from Enforced Disappearance defines enforced disappearance as deprivation of liberty by or with the involvement of the State, followed by refusal to acknowledge that deprivation or concealment of the person's fate or whereabouts [1]. In Uruguay, the search remains active. The 205 names read during the 2026 *Marcha del Silencio* [2] comprised 162 active and 43 clarified cases [3]. The Institución Nacional de Derechos Humanos y Defensoría del Pueblo separately reported 243 investigations: the same 162 active cases plus 81 complaints under analysis [3]. Photographs remain central to this public presence. They are carried each 20 May in the *Marcha del Silencio*, circulate through archives and digital platforms, and sustain a demand that is simultaneously familial, collective, historical, and political [4]. In 2022, *Imágenes del Silencio*, working with *Madres y Familiares*, distributed more than twenty thousand reproductions of the then 197 recognized portraits among marchers under the title *Todos somos familiares* [5].

Desaparecidos.uy begins from this unresolved condition. It uses reviewed public portraits as referents but does not infer missing facts, improve archival photographs, animate the disappeared, or claim forensic authority. It comprises three related works. The first, *Todos somos familiares*, deliberately adopts the title of the 2022 action and forms a portrait from fragments of reviewed images of contemporary people. *Están*

en todas partes forms it from reviewed images of places, surfaces, landscapes, and infrastructures. *Seguimos buscando* forms it from a visual vocabulary acquired progressively during a street-level traversal. The stills and videos are manifestations of the memorial, not demonstrations of a general-purpose tool. Collectivity, territorial implication, and continuing search must therefore be present in the image process itself rather than added later as explanatory language.



Fig. 1. Actual outputs from the current *Están en todas partes* process, reconstructing the public portrait of Lila Epelbaum Slotopolsky from fragments extracted from reviewed Montevideo street photographs. Left: regular-grid composition. Right: free-positioned composition with overlapping fragments. Epelbaum Slotopolsky was an Argentine student abducted with her brother Claudio in Punta del Este, Uruguay, on 4 November 1976. Her remains were recovered from the Cementerio de Avellaneda and identified; Claudio remains detained-disappeared. These software-generated research outputs belong to the developing artwork; they are not presented as completed exhibition works.

The project is an artwork in development and an ongoing research project. Its current software supports the three works from source review through still and video generation, but the intended corpora, public consultation, legal review, and exhibition form remain in progress. This incompleteness is not a defect to be concealed. A computational memorial concerned with unresolved disappearance should not equate a successful render with historical closure.

The research follows practice-based inquiry, in which making produces questions and knowledge rather than merely illustrating a prior argument [6]. Performance-led research in the wild likewise demonstrates how artist-driven public works can generate computing practice, studies, and theory through their development and deployment [7]. The project also continues earlier work on making technological mediation explicit as an artistic operation [8]. Implementing the three works revealed contradictions that the original concept did not resolve: a large source collection can still yield a concentrated image; neighboring fragments can partially reconstitute a contemporary source person; a nominal search can use

material not yet encountered; and a fixed grid can reduce a processual memorial to a conventional photomosaic.

Two contributions follow. First, the paper demonstrates empirically that source diversity depends on the provenance unit to which a constraint is attached. A quota applied to manifest rows or crops can manufacture apparent plurality while concentration persists at the parent-image level. Second, it presents an artwork-derived architecture in which fragment provenance, final placement, and process time remain inspectable, allowing the three memorial propositions to impose different technical conditions without collapsing them into a single optimization objective.

II. ARTISTIC AND TECHNICAL CONTEXT

The most direct artistic precedent is Uruguayan. Juan Ángel Urruzola's *Álbum de memoria* assembles small identification-style photographs of detained-disappeared people into the enlarged face of another disappeared person [9]. *Desaparecidos.uy* reverses the source field: the disappeared emerge not from one another but from fragments of the contemporary country. The project is also distinct from the author's earlier conversational-AI research on collective memories of the disappeared [10], which concerned testimony and language rather than image synthesis.

Photomosaic research has addressed efficient retrieval, adaptive tiling, and repetition constraints [11]–[13]. Those techniques matter here, but resemblance is not the sole objective. Source constitution, the visibility of process, the risk of reconstituting a living contributor, and the temporal order of search are part of the artwork's meaning. Thanatosensitive design established death-facing and memorial systems as legitimate objects of computing research [14]. The present case extends that concern to state violence, archival politics, and a memorial whose historical demand remains open.

Public availability does not erase contextual integrity [15]. Under Uruguayan law, biometric data are personal data obtained through specific technical processing of physical, physiological, or behavioral characteristics that permits or confirms unique identification; processing of this kind requires a prior data-protection impact assessment [16]. The project does not perform image recognition or other processing intended to identify source persons. Facial images are nevertheless treated as personal source material requiring contextual, legal, and rights review. The system separates technical eligibility from consent, rights clearance, community authorization, and permission to publish.

III. FROM MEMORIAL PROPOSITIONS TO EXECUTABLE CONSTRAINTS

A. Shared Representation

Fig. 2 summarizes the shared computational representation, and Table I the constraints retained by each work. A target record identifies a disappeared person and a reviewed public portrait used as referent. A source record identifies one reviewed image or region, its origin, review state, and available crop coordinates. Approved source regions are divided into square fragments. The reproducible benchmark and canonical artwork runtime describe each fragment and target section through normalized mean color, luminance contrast, and horizontal and vertical luminance variation, producing a six-dimensional descriptor. The current interactive renderer used for Fig. 1 additionally provides a

coarse spatial matcher: 4×4 mean-color blocks, an 8×8 luminance map, and horizontal and vertical luminance gradients form a 240-dimensional hand-designed descriptor. Neither mode uses learned embeddings, face recognition, identity inference, restoration, or a post-compositional smoothing pass.




Fig. 2. The three works share a common representation of fragments, placements, and history, while each imposes constraints derived from its own memorial proposition.

The matcher selects fragments at minimum Euclidean descriptor distance while enforcing configured fragment-reuse and source-participation capacities. A one-use-region option sets the fragment-reuse limit to one. Before matching, requested capacities are checked against target positions, available fragments, and available sources; an infeasible request fails rather than silently relaxing a cap. In the interactive *Están en todas partes* path, exact region reuse and parent-image participation are deliberately separate controls: several distinct regions may come from one reviewed photograph. This distinction is central to the provenance argument developed below.

Matching is separated from final geometry. A placement binds a source fragment to its source rectangle and matched target rectangle and can record its rendered position, size, rotation, opacity, layer, and process time. The current interactive interface exposes regular-grid and free-positioned composition; the latter uses the overlap placement grammar. The canonical artwork runtime also retains grid, irregular, and overlap grammars for controlled comparison. Its output sidecars contain aggregate source use and a versioned placement history supporting deterministic replay, dependency queries, excluded-source regeneration, geometric comparison, and later video rendering. Traceability records technical dependence; it does not legitimate the underlying use.

TABLE I. MEMORIAL PROPOSITIONS AND TECHNICAL CONDITIONS

Work	Source material	Technical condition
<i>Todos somos familiares</i>	Reviewed images of contemporary people	Positive per-source cap; optional exclusion of same-source neighbors; no raw source reveal
<i>Están en todas partes</i>	Reviewed places, surfaces, landscapes, infrastructures	Territorial grouping and interleaving; no fabricated coverage claim
<i>Seguimos buscando</i>	Ordered, reviewed traversal frames	No placement may depend on a frame not yet encountered

B. Todos Somos Familiares

The work requires an asymmetry between target and source. The documented portrait must remain perceptible enough to sustain naming, while a contemporary person contributing fragments must not become the represented subject. People-derived generation therefore requires a positive per-source cap and can exclude adjacent target sections from using the same source. The canonical artwork video reveals the emergence of fragments without displaying the complete source photograph. These measures reduce structural opportunities for reconstitution but do not establish anonymity. People-derived outputs remain internal until contextual, legal, and visual review.

C. Están en Todas Partes

Here the contemporary territory is material rather than background. Source records can be grouped and interleaved through reviewed territorial labels, while missing or unlocated material remains explicit. The system records which groups were available and which were realized in an output; it does not infer national coverage from a bounded collection.

Fig. 1 compares two current output geometries. The grid preserves the target lattice. Free-positioned composition enlarges, rotates, displaces, and overlaps the same class of matched fragments, producing a less stable field without changing the historical referent. Its target, Lila Epelbaum Slotopolsky, was Argentine and was abducted in Uruguay. Her remains were recovered from the Cementerio de Avellaneda and identified, while her brother Claudio remains detained-disappeared [17]. Her inclusion makes explicit that the project's relation to Uruguay is territorial and historical rather than a restriction based on nationality, and that the referent corpus includes both active and clarified cases. The proposition *están en todas partes* is political and memorial, not a statistical claim that a particular source corpus represents every territory equally.

D. Seguimos Buscando

Traversal frames form an ordered sequence. Let $F(I_j)$ denote the fragments obtained from frame I_j . At frame k , the available vocabulary is $F_k = \bigcup_{j=1}^k F(I_j)$. A placement created at frame k may draw only from this accumulated vocabulary, never from a future frame. The placement history records the encounter index associated with each fragment. Route and encounter order therefore affect the constitution of the portrait rather than functioning as interchangeable background footage. This condition is simple but decisive: a process that claims to continue searching cannot begin with everything already found.

IV. EVALUATION

Evaluation follows the artwork's requirements rather than a single image-quality score. For N placements with per-source counts n_s , participation is described by source count, maximum share, $HHI = \sum_s (\frac{n_s}{N})^2$ and effective source count. Target-side concentration is described by same-source adjacency and the largest connected same-source region. Geometry is described by overlap, rotation, and opacity. Temporal validity is evaluated by comparing each placement's encounter index with the point at which its source entered the traversal. These measures identify conditions requiring interpretation; they do not determine consent, recognizability, historical adequacy, or memorial value.

A. Provenance Granularity and the Cost of Caps

A 432-run proxy benchmark examined the reach and visual cost of source limits. It used six native 25×25 target portraits from *scikit-image*'s *lfw_subset*, upsampled to 200×200 [18], [19]; three deterministic run seeds; two 128-row source corpora; three matchers; four row caps; a fragment-reuse limit of eight; and 100 target sections per output. The mixed-LFW corpus used indices 20–147: 80 face images and 48 non-face background samples. The heterogeneous-example corpus used deterministic crops from eight *scikit-image* example images, including astronomical and synthetic material. The first condition is matched-domain and the second cross-domain, so their absolute SSIM values are not a controlled corpus comparison. No face-recognition model was trained or evaluated, and no identity inference was performed. The baselines were deterministic modular selection (labeled random), mean-RGB nearest neighbor, and the six-dimensional descriptor. SSIM [20] and CIEDE2000 [21] measured structural similarity and color difference. Repeated seeds are algorithmic runs rather than independent samples; the deterministic six-dimensional matcher produced identical outputs across seeds.

Uncapped, the six-dimensional descriptor improved SSIM over mean-RGB and deterministic modular selection: 0.409 against 0.278 and 0.147 for mixed LFW, and 0.283 against 0.172 and 0.147 for heterogeneous examples. Row-level HHI rose from 0.0174 to 0.0196 to 0.0266 across those matchers for mixed LFW, and from 0.0174 to 0.0435 to 0.0619 for heterogeneous examples. The identical modular-selection HHI reflects the two equal 3,200-fragment pools. Unconstrained six-dimensional matching on the heterogeneous corpus used an average of 31.7 of 128 rows; the largest row supplied 12.3% of an output on average and 15% in the most concentrated target.

Under the sequential six-dimensional matcher, a one-tile row cap made row-level HHI approximately 0.0100 by construction but reduced SSIM by 9.2% for mixed LFW and 18.1% for heterogeneous examples; mean CIEDE2000 increased by 4.5% and 33.4%. This is not an intrinsic cost of caps. In reduced cap-one fixtures, the greedy objective was 15.7% and 28.1% above a min-cost-flow optimum, respectively; moreover, cap one improved heterogeneous-corpus SSIM under the weaker RGB and modular baselines. The reported resemblance loss is therefore specific to a matcher and assignment regime able to exploit concentration.

The heterogeneous-example corpus exposes the more important limitation. Its 128 rows descend from eight parent images. Under a one-tile row cap, an output may truthfully report 100 contributing rows while one parent still supplies sixteen crops, or 16% of the image. Parent-level HHI has the theoretical lower bound $1/8 = 0.125$ for this corpus; with 100 discrete placements, the attainable minimum is 0.1252. A uniform parent cap below $\lceil 100/8 \rceil = 13$ is infeasible. The measured parent-level HHI at the one-tile row cap was 0.1459 ± 0.0054 , mean $\pm$ population SD across six targets, compared with 0.2648 without a row cap. The constraint reduced concentration, but the corpus imposed a floor that row-level reporting alone concealed.

Table II summarizes the principal results.

The result is not specific to memorial art. Any quota, fairness constraint, or participation audit governs only the entity to which it is attached. A claim of plurality must

therefore name the relevant provenance unit and disclose the hierarchy beneath it. In *desaparecidos.uy*, this means distinguishing fragments, manifest rows, parent images, contributors, territorial groups, and traversal frames rather than treating them as interchangeable sources.

TABLE II. KEY RESULTS FOR THE SIX-DIMENSIONAL MATCHER

Corpus	SSIM	Row HHI: Mod. ^a ; RGB; 6-D	Cap-one ΔSSIM; ΔE00	Parent HHI; greedy gap
Mixed LFW	0.409	0.0174; 0.0196; 0.0266	-9.2%; +4.5%	— ^b ; 15.7%
Heterogeneous	0.283	0.0174; 0.0435; 0.0619	-18.1%; +33.4%	0.1459 ± 0.0054; floor 0.125; 28.1%

^a SSIM and row-HHI columns report uncapped runs. “Mod.” denotes the deterministic modular baseline. Parent-level HHI applies only to the heterogeneous corpus, whose 128 source rows derive from eight parent images; greedy-gap values come from reduced 32-source cap-one fixtures, whose LFW corpus contains faces only; ^b “—” indicates not applicable.

B. Placement-history and Temporal Checks

The archived evaluator operates on versioned placement histories rather than final pixels alone. Its pinned test fixture contains four placements on a 48×48 canvas and verifies exact participation metrics, same-source adjacency, nonzero geometric overlap, minimum opacity 0.9, and the explicit result that these structural measures do not establish anonymity.

The same four-placement history initially produces no causal violations. Changing one placement's encounter index so that it precedes the source's position in the recorded sequence produces exactly one future-source violation. The test establishes that the condition is executable and that its breach is detectable. It does not establish that a particular route is historically significant or that a technically causal traversal constitutes an adequate memorial.

The 432-run script, reduced min-cost-flow comparison, evaluator, and tests are archived at revision 3c1a0a4 under the benchmark's working title, *Bounding and Auditing Source Participation in Fragment-Based Image Synthesis* [22]. The script's per-run records contain the values cited here; the committed *reported-results.json* stores a principal subset rather than every prose statistic. Fig. 1 was produced with the later interactive runtime audited at revision 700bc1f [23], so the benchmark supports the provenance-granularity claim but is not a quantitative evaluation of Fig. 1. The reproducible artifact contains only proxy images; the memorial portrait corpus, reviewed people-derived corpus, and crawler records are not distributed with it.

V. DISCUSSION AND LIMITATIONS

Computing is constitutive of this memorial. A source schema determines what counts as a participant. Capacity and adjacency constraints determine how participation may appear. Placement determines whether an image reads as a regular mosaic, an unstable accumulation, or a layered field. Encounter history determines whether search is genuinely causal. These are engineering decisions, but they also determine whether the work remains faithful to its propositions.

The provenance experiment provides the paper's strongest general result. A system can satisfy a numerical diversity rule

and still concentrate material at a higher level that the metric ignores. The remedy is not simply to add stricter caps. Constraints must follow the provenance hierarchy, and feasibility must be checked before a system claims that a requested distribution is possible. The current interactive *Están en todas partes* path illustrates the distinction: its one-use option prevents exact fragment repetition, while permitting multiple distinct regions from the same parent photograph. Unique regions are not unique sources.

The implementation does not yet fully encode the hierarchy that the benchmark shows to be necessary. Production manifests record a source row's URL, page, review status, local path, hashes, territorial label, and available crop or face region, but they do not represent every parent asset, contributor, or institutional collection as a dedicated first-class entity. The proxy corpus's parent relations were known by construction. Parent- and contributor-level controls in the artwork therefore require further schema work before they can be audited with the same precision as fragment- and row-level participation.

The controls are also deliberately incomplete. A low source share does not prove that a contemporary person is unrecognizable. Territorial interleaving does not establish national representativeness. A causal traversal does not validate the route. A target may remain insufficiently perceptible while every constraint passes, or become falsely restorative while scoring well on low-level similarity. Successful generation is consequently separated from release: an output can be technically valid and still remain unsuitable for publication or exhibition.

The evidence is bounded in several ways. The benchmark uses public proxy data rather than the memorial corpus so that it can be reproduced without distributing sensitive images. Its mixed-LFW and heterogeneous-example conditions are not domain-matched, and it does not measure audience reception or memorial adequacy. Fig. 1 uses a bounded, reviewed Montevideo source corpus and demonstrates two current composition modes; it does not claim territorial breadth. The spatial matcher and portrait-priority ordering are explicit heuristics rather than validated models of visual salience. Final exhibition loops have not yet been produced from the complete intended corpora. People-derived public presentation remains conditional on rights review, consultation, and frame-level inspection. None of the reported metrics establishes anonymity.

These limits describe the state of the work rather than a separate future project. *Desaparecidos.uy* is already an artwork because its images and temporal processes constitute the memorial proposition. It is also research in progress because making continues to expose technical and ethical questions that alter both the system and the work. The paper does not claim closure over either domain.

VI. CONCLUSION

Desaparecidos.uy forms portraits of people detained-disappeared in connection with Uruguay through three related processes: collective contribution, territorial implication, and continuing search. Its principal computing contribution is not a new image descriptor but the demonstration that diversity depends on provenance granularity, together with an architecture that keeps fragment-to-source participation, final

geometry, and encounter time available for inspection while making the remaining provenance gaps explicit.

The project also establishes a methodological point for computational art. Situated artistic propositions can generate precise technical questions without being reduced to engineering requirements. Here, collectivity produced the problem of concentration; the nonrepresentation of contemporary contributors produced capacity, adjacency, and reveal controls; continuing search produced temporal causality; and historical incompleteness produced a system in which a finished render is never treated as proof that the memorial task is finished. Technical completion is not historical closure, traceability is not consent, structural risk controls are not anonymity, and image metrics are not memorial adequacy.

REFERENCES

- [1] United Nations, *International Convention for the Protection of All Persons from Enforced Disappearance*. New York, NY, USA, 2006.
- [2] Presidencia de la República Oriental del Uruguay, “Marcha del Silencio: Una multitud marchó en reclamo de memoria, verdad y justicia,” May 20, 2026. [Online]. Available: <https://www.gub.uy/presidencia/comunicacion/noticias/marcha-silencio-2026>
- [3] Institución Nacional de Derechos Humanos y Defensoría del Pueblo, “Equipo de Investigación de la Búsqueda de INDDHH trabaja en 243 casos por denuncias de desaparición forzada,” May 14, 2026. [Online]. Available: <https://www.gub.uy/institucion-nacional-derechos-humanos-uruguay/comunicacion/noticias/equipo-investigacion-busqueda-inddhh-trabaja-243-casos-denuncias-desaparicion>
- [4] S. Bustamante-Brauning, “[Re]-appearances online: Photography, mourning and new media ecologies for representing the Southern Cone’s disappeared on two digital memory platforms,” *Photographies*, vol. 15, no. 3, pp. 381–403, 2022, doi: 10.1080/17540763.2022.2096679.
- [5] Imágenes del Silencio, “Todos somos familiares,” 2022. [Online]. Available: <https://imagenesdelsilencio.uy/projects/todos-somos-familiares-2022/>
- [6] L. Candy and E. Edmonds, “Practice-based research in the creative arts: Foundations and futures from the front line,” *Leonardo*, vol. 51, no. 1, pp. 63–69, 2018, doi: 10.1162/LEON_a_01471.
- [7] S. Benford et al., “Performance-led research in the wild,” *ACM Trans. Comput.-Hum. Interact.*, vol. 20, no. 3, Art. no. 14, pp. 14:1–14:22, 2013, doi: 10.1145/2491500.2491502.
- [8] T. Laurenzo, “Media appropriation and explication,” *Journal of Science and Technology of the Arts*, vol. 8, no. 2, pp. 27–36, 2016, doi: 10.7559/citarj.v8i2.181.
- [9] Dodecá, “Expone Juan Ángel Urruzola,” Oct. 9, 2009, quoting T. Oroño on *Álbum de memoria*. [Online]. Available: <https://dodeca.org/2009/10/expone-ja-urruzola/>
- [10] T. Laurenzo, “Capturing Collective Memories of the Disappeared with Artificial Intelligence,” in *Advances in Artificial Intelligence – IBERAMIA 2024*, L. Correia, A. Rosá, and F. Garijo, Eds., LNCS 15277. Cham, Switzerland: Springer, 2025.
- [11] G. Di Blasi and M. Petralia, “Fast photomosaic,” in *Poster Proc. 13th Int. Conf. Central Europe Computer Graphics, Visualization and Computer Vision*, 2005, pp. 15–16.
- [12] D. Pavić, U. von Ceumern, and L. Kobbelt, “GIzMOs: Genuine image mosaics with adaptive tiling,” *Computer Graphics Forum*, vol. 28, no. 8, pp. 2244–2254, 2009, doi: 10.1111/j.1467-8659.2009.01437.x.
- [13] H. Narasimhan and S. Sathesh, “A randomized iterative improvement algorithm for photomosaic generation,” in *Proc. NABIC*, 2009, pp. 777–781, doi: 10.1109/NABIC.2009.5393882.
- [14] M. Massimi and A. Charise, “Dying, death, and mortality: Towards thanatosensitivity in HCI,” in *CHI ’09 Extended Abstracts*, New York, NY, USA: ACM, 2009, pp. 2459–2468, doi: 10.1145/1520340.1520349.
- [15] H. Nissenbaum, “Privacy as contextual integrity,” *Washington Law Review*, vol. 79, no. 1, pp. 119–157, 2004.
- [16] Uruguay, *Ley No. 18.331, Protección de Datos Personales y Acción de Habeas Data*, arts. 4(Ñ) and 18 bis; Unidad Reguladora y de Control de Datos Personales, Resolución No. 30/020, May 12, 2020.
- [17] Sitios de Memoria Uruguay, “Epelbaum Slotopolsky, Lila.” [Online]. Available: <https://sitiosdememoria.uy/epelbaum-slotopolsky-lila>
- [18] S. van der Walt et al., “scikit-image: Image processing in Python,” *PeerJ*, vol. 2, e453, 2014, doi: 10.7717/peerj.453.
- [19] G. B. Huang, M. Ramesh, T. Berg, and E. Learned-Miller, “Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments,” Univ. of Massachusetts Amherst, Amherst, MA, USA, Tech. Rep. 07-49, 2007.
- [20] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, “Image quality assessment: From error visibility to structural similarity,” *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, 2004, doi: 10.1109/TIP.2003.819861.
- [21] G. Sharma, W. Wu, and E. N. Dalal, “The CIEDE2000 color-difference formula: Implementation notes, supplementary test data, and mathematical observations,” *Color Research & Application*, vol. 30, no. 1, pp. 21–30, 2005, doi: 10.1002/col.20070.
- [22] T. Laurenzo, “desaparecidos.uy: Artwork runtime and URUCON 2026 research artifact,” GitHub repository, rev. 3c1a0a4d42a98c3ecc630e555845a326e83a2bc9, 2026. [Online]. Available: <https://github.com/krahd/desaparecidos.uy/tree/3c1a0a4d42a98c3ecc630e555845a326e83a2bc9>
- [23] T. Laurenzo, “desaparecidos.uy: Current interactive artwork runtime,” GitHub repository, rev. 700bc1f820ee08074bf341bf7714f9aca43c4783, 2026. [Online]. Available: <https://github.com/krahd/desaparecidos.uy/tree/700bc1f820ee08074bf341bf7714f9aca43c4783>