

Lifting Lines and Tone: Image-space Stylization in Path-space

REX WEST, Aoyama Gakuin University, Japan
SAYAN MUKHERJEE, The University of Tokyo, Japan
YONGHAO YUE, Aoyama Gakuin University, Japan




Fig. 1. Using our method, lines and tone naturally blur with camera depth-of-field (**yellow**) and glossy reflections (**red**), tone patterns maintain their image-space structure even in curved reflections (**blue**), and rendering stylistic visuals in a light transport-based system lets hatching on the van respond to the physically-based reflection of the sitting robot (**green**). Lines and tone; lifted.

Many non-photorealistic rendering (NPR) styles, such as feature lines and hatching, are defined by image-space structure inherited from hand-drawn media. While recent path-space formulations like the stylized rendering equation (SRE) enable stylization to interact naturally with light transport, they provide no mechanism for enforcing image-space consistency.

We present a conceptual framework for *lifting* image-space stylizations into path-space in a principled, SRE-compatible manner. Our key observation is that image-space consistency can be achieved by establishing geometrically-driven mappings from image-space to path-space. We demonstrate this through two complementary stylizations: feature lines and tone.

For feature line rendering, we introduce a *conditional* lifting based on partial path-space variation, where a geometric path parametrization is combined with parallel transport to preserve image-space structure under distribution effects. This enables a curvature-aware, stochastic, geometry-driven formulation of line detection that generalizes prior ray-based methods. For tone-based styles, such as hatching and halftone, we propose a *canonical* lifting anchored to a material-independent mapping, motivated by stationary-phase arguments from wave optics. Its locally invertible approximation enables evaluation of image-space tone fields at arbitrary path vertices while preserving image-space structure under complex light transport.

Both methods are implemented as ordinary style functions under the SRE and work with existing estimation and sampling strategies; demonstrating how image-space structure can be preserved within path-space rendering, enabling a broader class of expressive, physically-grounded NPR styles.

Authors' addresses: Rex West, Aoyama Gakuin University, Japan, rexwest@gmail.com; Sayan Mukherjee, The University of Tokyo, Japan, sayan@phys.s.u-tokyo.ac.jp; Yonghao Yue, Aoyama Gakuin University, Japan, yonghao@it.aoyama.ac.jp.

© 2026 Copyright held by the owner/author(s).
This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *ACM Transactions on Graphics*, <https://doi.org/10.1145/3811359>.

CCS Concepts: • **Computing methodologies** → **Non-photorealistic rendering**; **Rendering**; **Ray tracing**.

ACM Reference Format:

Rex West, Sayan Mukherjee, and Yonghao Yue. 2026. Lifting Lines and Tone: Image-space Stylization in Path-space. *ACM Trans. Graph.* 45, 4, Article 138 (July 2026), 15 pages. <https://doi.org/10.1145/3811359>

1 INTRODUCTION

Several non-photorealistic rendering (NPR) styles—such as feature lines, hatching, and halftone—inherit their visual structure from hand-drawn media. Traditionally, line and tone stylizations are implemented independently of light transport, often as image-space post-processes. While this enforces image-space structure, it decouples stylization from lighting, materials, and global illumination.

Recent work integrates stylization directly into the light transport formulation [Doi et al. 2021; West and Mukherjee 2024]. In particular, the stylized rendering equation (SRE) [West and Mukherjee 2024] extends the rendering equation [Kajiya 1986] with parameterized (potentially non-linear) style functions applied at each bounce. This enables stylization to interact naturally with material scattering, distribution effects, and indirect illumination.

However, the SRE is fundamentally a path-space formulation and provides no mechanism for enforcing image-space structure (see Fig. 2). As a result, styles whose defining characteristics are inherently image-based remain difficult to express within this otherwise compelling framework.

Our key observation is that to achieve image-space consistency in a path-space formulation we need a way to *lift* image-space structure into path-space geometry. This geometric lifting supplies




Fig. 2. Failure modes of prior path-space stylization (e.g., West and Mukherjee [2024] (bottom; left/middle) and West [2021] (bottom; right)) versus our image-space-consistent formulation (top). Path-space hatching produces geodesic-like curves (left), halftone patterns alias with surface curvature (middle), and line widths are inconsistent in curved reflections (right). Such effects may be desirable stylistically, but should be controllable choices rather than algorithm-imposed artifacts.

the correspondences needed by image-space styles, while leaving the SRE’s style functions, and estimators, unchanged.

We achieve this by establishing mappings between image-space coordinates and *path-space geometry*. However, such mappings are not immediately well-defined under geometric distribution effects, such as glossy reflection or depth of field, where a single point on the image plane can correspond to uncountably infinite geometric paths. We address this via *conditional specularity*, a geometric path parametrization that leverages the fact that certain distribution effects can be modeled as an ensemble of specular events.

We demonstrate two such mappings for two common components of hand-drawn styles: lines and tone. In both cases, we design geometric operators that establish the mapping needed to preserve image-space structure, and then use those mappings inside ordinary SRE style functions.

For *line rendering*, we use a *conditional lifting*. The objective of line rendering is to draw lines along perceived discontinuities in the rendered image. In a path-space formulation, this boils down to detecting sufficiently large gradients of path-space quantities (e.g., depth, normals, albedo) with respect to changes in image coordinate. For purely specular transport, such gradients are not only well-defined, but there is also a notion of *image-space parallelism*, under which path-space measurements vary smoothly with image coordinate. Extending line rendering to distribution effects, conditional specularity lets us “freeze” a glossy path such that it acts specular with respect to a parametrization. Parallel transport then provides a way to smoothly move the parametrization across the image plane such that we can construct paths for other image-space coordinates that preserve image-space parallelism, account for changes in surface curvature, and establish a well-defined gradient suitable for line detection.

We then reinterpret line detection for a given path as a *stochastic search for a lower bound on (parametrized) extremal image-space gradients*. An SRE-compatible style function follows directly from the expectation of this search process, yielding a geometry-driven formulation that preserves image-space line structure under complex light transport and generalizes prior ray-based methods.

For *tone rendering*, such as hatching and halftone, we instead use a *canonical lifting*. In contrast to lines, the defining quality of tone is an *image-space pattern* whose visual identity lies in coherent spacing, orientation, and phase. Preserving this structure in a path-space formulation requires establishing a mapping that is *material-invariant*—if the mapping depends on microfacet randomness or bidirectional scattering distribution function (BSDF) parameters, the pattern would be re-parametrized by material scattering and sampling noise.

We use a style-agnostic, mirror-reflection “anchor” mapping motivated by stationary phase arguments from wave optics, in which mirror-reflections about the macro surface geometry act as canonical carriers of phase, while distribution effects correspond to neighborhood integration around this anchor. Because this mapping is generally non-invertible (due to caustics and partial domain coverage), we introduce a *locally invertible, numerically stable approximation that behaves like an analytic continuation* where sufficient geometric information is available. This enables coherent evaluation of image-space tone fields at arbitrary path vertices under indirect illumination, glossy reflection, refraction, and depth of field.

Viewed together, line and tone stylization differ mainly in how they use the image-path mappings: lines rely on *conditional lifting*, where a set of image-space parallel paths is constructed from the parametrization of a given path, while tone relies on a *canonical lifting*, where the image-space structure is anchored to a fixed reference such that distribution effects integrate naturally over it. These geometric operators do not change the underlying formulation, provide the necessary correspondences needed for line and tone rendering, and seamlessly integrate into ordinary SRE style functions.

Concretely, our contributions are:

- A stochastic, extremal-gradient formulation of feature line detection, together with an SRE-compatible line rendering style function.
- A principled treatment of distribution effects for line rendering, combining conditional specularity with Levi-Civita parallel transport to maintain image-space line placement and width through glossy transport and depth of field.
- A material-invariant mapping from image-space to path-space for tone fields.
- A robust local approximation of the inverse of this mapping that enables image-space-consistent tone patterns under complex light transport.

2 BACKGROUND AND RELATED WORK

2.1 Stylized rendering

The stylized rendering equation extends the rendering equation [Kajiya 1986] by introducing a potentially non-linear style function g_θ on the exitant radiance $L(\mathbf{x}, \mathbf{y})$ at each bounce,

$$L(\mathbf{x}, \mathbf{y}) = g_\theta \left(L_e(\mathbf{x}, \mathbf{y}) + \int_{\mathcal{V}} f_r(\mathbf{x}, \mathbf{y}, \mathbf{z}) G(\mathbf{y}, \mathbf{z}) L(\mathbf{y}, \mathbf{z}) d\mathbf{z} \right), \quad (1)$$

with an auxiliary form,

$$L(\mathbf{x}, \mathbf{y}) = L_e(\mathbf{x}, \mathbf{y}) + \int_{\mathcal{V}} f_r(\mathbf{x}, \mathbf{y}, \mathbf{z}) G(\mathbf{y}, \mathbf{z}) g_\theta(L(\mathbf{y}, \mathbf{z})) d\mathbf{z}, \quad (2)$$

where $L(\mathbf{x}, \mathbf{y})$ represents the exitant radiance from surface point \mathbf{y} towards \mathbf{x} , L_e is emitted radiance, f_r is the bidirectional scattering distribution function, and G is the geometry term. The style function $g_\theta : \mathbb{R}^k \rightarrow \mathbb{R}^k$ operates on vector-valued radiance quantities and is parameterized by θ .

Recent work focuses on improving the practicality of SRE-based rendering. Tong and Hachisuka [2025] introduce sampling algorithms based on multi-vertex path filtering [Deng et al. 2021; West et al. 2022] and neural radiance caching [Müller et al. 2021] to mitigate the exponential cost of baseline tree sampling. These advances make physically-based stylized rendering feasible for complex scenes while preserving the expressiveness of the SRE.

However, maintaining image-space consistency for stylizations whose defining characteristics are image-based—such as lines and tone—remains an open challenge.

2.2 Feature line rendering

Feature line rendering [Bénard and Hertzmann 2019; DeCarlo 2012] places lines along perceived discontinuities in surface appearance, analogous to hand-drawn line drawings [Cole et al. 2008; Hertzmann 2020, 2021]. These discontinuities typically correspond to extrema or large gradients of view- or surface-dependent quantities such as depth, surface normal, curvature, or shading. Since closed-form expressions for these gradients are rarely available, most methods rely on approximations using local samples.

Feature line methods can be broadly categorized as image-space or object-space methods. Image-space methods operate on rendered images or image-space buffers and place lines based on image-space cues such as depth, normal, or shading discontinuities [Saito and Takahashi 1990; Lee et al. 2007; Kim et al. 2008]. Their computational cost scales with image resolution rather than geometric complexity, but line placement is limited to directly visible surfaces and is sensitive to sampling artifacts. Object-space methods instead operate directly on surface geometry, using differential quantities such as curvature to place lines on silhouettes, ridges, and valleys [DeCarlo et al. 2003, 2004; Ohtake et al. 2004; Judd et al. 2007]. While effective, these approaches become increasingly expensive for complex geometry and are difficult to integrate within general light transport.

Ray-based methods bridge these categories by estimating gradients via finite differences of surface quantities evaluated at ray intersections. Choudhury and Parker [2009] trace stencils of camera rays and compare surface attributes at their first intersections, restricting line detection to directly visible geometry. Ogaki and Georgiev [2018] extend this approach to specular reflection and

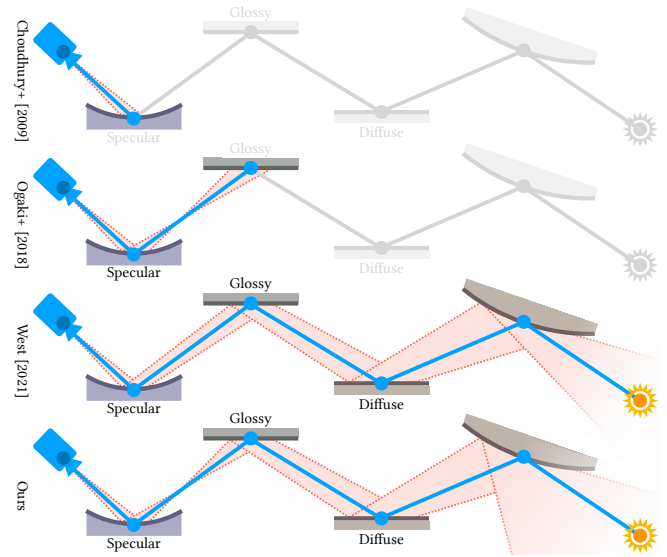


Fig. 3. For a base path sample (blue), additional paths are sampled over a projected region of image-space (red) and used to estimate line-producing gradients. Choudhury and Parker [2009] restrict line detection to directly visible geometry, while Ogaki and Georgiev [2018] extend it to perfect specular reflection/refraction, leveraging the image-space parallelism and one-to-one mapping of image-space coordinates to specular paths. However, this mapping breaks down under distribution effects. West [2021] partially overcomes this by sampling additional paths in a cone around the base path, enforcing a pseudo-parallel neighborhood and one-to-one mapping, but ignoring surface curvature, leading to inconsistent line widths (e.g. lack of an hourglass shape along the second edge). Our approach explicitly accounts for the geometric relationship of image-space to path-space, including surface curvature, resulting in image-space parallelism across distribution effects and curved surfaces, and enabling consistent line widths at arbitrary path depth (gray interactions indicate bounces not handled by earlier methods).

refraction and reduce aliasing through stochastic sampling. West [2021] further integrate ray-based line rendering into physically based rendering by treating lines as implicit light-emitting surfaces, extending line detection to distribution effects.

While the method of West [2021] successfully incorporates line rendering into physically based light transport, it lacks an explicit underlying formulation. Moreover, its cone-based sampling heuristic does not account for surface curvature and can produce inconsistent line widths on curved surfaces. These limitations highlight the need for a geometrically principled formulation for line rendering under distribution effects that preserves image-space structure. See Fig. 10 for a concrete example.

2.3 Tone rendering

Tone rendering methods in non-photorealistic rendering represent tonal variation using structured patterns. This includes styles such as hatching and halftone, which rely on consistent spatial relationships in image-space to convey form, material, and illumination.

Hatching methods [Hertzmann and Zorin 2000; Praun et al. 2001] approximate tone using sets of lines whose density, orientation, and spacing vary with shading. Some approaches [Deussen et al. 1999;

Medeiros et al. 2009] generate hatching patterns by intersecting scene geometry with families of planes or curves, producing line patterns that follow surface geometry or geodesics. While effective at conveying surface shape, such geometry-driven constructions define hatching primarily in object space, lacking controllable image-space structure.

Halftone techniques [Hall 1999; Deussen and Isenberg 2013] similarly approximate continuous tone using discrete patterns—such as dots or stipples—whose spatial distribution creates the perception of smooth shading when viewed at sufficient scale. As with hatching, the visual identity of halftone relies on consistent image-space spacing and orientation.

Traditionally, these tone-based stylizations have been implemented as a post-processes. However, this separation from light transport prevents tone style from interacting coherently with physical effects such as indirect illumination, glossy reflection, or refraction.

More recently, West and Mukherjee [2024] demonstrate examples of hatching and halftone as style functions within the SRE framework. However, because these stylizations are defined purely in path-space, the resulting patterns fail to preserve image-space structure, exhibiting geodesic artifacts, particularly in reflections and on curved surfaces (see Figs. 2 and 13). For some use-cases these artifacts can be desirable, but they should be an artistic choice, not a limitation of the underlying method.

This highlights the need for a mechanism that allows image-space-structured tone patterns to be evaluated in path-space under complex light transport.

3 LIFTING IMAGE-SPACE TO PATH-SPACE

To lift image-space \mathcal{S} stylizations into path-space \mathcal{P} , we formalize the connection between image-space structure and path-space geometry as a *parametrization of paths*. We introduce a space of conditional variables Ξ that fix a particular realization ξ of distributional effects (e.g., lens sample or microfacet selection) that lets us establish a mapping,

$$\Phi : \mathcal{S} \times \Xi \rightarrow \mathcal{P}, \quad (3)$$

such that a pair (s, ξ) of an image-space coordinate s and conditional variable ξ uniquely determines a path $\bar{x} = \Phi(s, \xi)$.

This view is closely related to path parametrizations used in primary sample space methods [Kelemen et al. 2002] and manifold exploration [Jakob and Marschner 2012]—the image-space coordinate s and conditioning variable ξ together fully specify a path sample. Fixing s and varying ξ recovers the distribution of paths contributing to a point on the image-plane, while fixing ξ and varying s defines a family of *image-space-parallel* paths.

Within this framework, lifting an image-space quantity into path-space amounts to choosing how Φ is used. We consider two complementary constructions.

Conditional lifting. A *conditional lifting* treats ξ as fixed and varies only the image-space coordinate s . This induces a family of paths $\{\Phi(s, \xi)\}$ that preserves the conditioned scattering events while moving smoothly across the image plane. Conceptually, this corresponds to a *partial variation* in path-space. Conditional lifting provides a well-defined notion of image-space variation, even in

the presence of distribution effects, since each conditioned path behaves as a specular realization. This makes it a natural tool for establishing image-space derivatives of path-space quantities.

In Section 4, we use this construction for *line rendering* where feature lines are detected by measuring gradients of geometric quantities along families of conditionally parallel paths.

Canonical lifting. A *canonical lifting* instead selects a particular conditioning variable ξ^* and defines a mapping

$$\phi(s) = \Phi(s, \xi^*), \quad (4)$$

which acts as an *anchor* relating image-space coordinates to path-space geometry. Unlike conditional lifting, which defines local variation, canonical lifting establishes a global reference parameterization. For this purpose, ξ^* should be chosen so that ϕ is coherent and independent of material-specific randomness. In our case, this corresponds to selecting a specular (mirror-like) realization of transport, yielding a material-invariant mapping.

In Section 5, we use this construction for *tone rendering* where image-space patterns are evaluated with respect to the canonical mapping, while distribution effects manifest as integration around this anchor rather than reparameterization of the pattern.

4 FEATURE LINES

Ray-based line rendering boils down to determining if a given ray passes close enough to a sufficiently large gradient of some measurable quantity, and then stylizing the light traveling down that ray with a line color—exactly the job of a style function under the SRE.

In the context of image-space consistent line rendering, the problem becomes *view-dependent* and we need to consider the geometric relationship between image-space and the quantities we want to measure a gradient for. For path-space quantities, like surface normals, we need a well-defined instance of the mapping (4) from image-space to those path-space quantities that follows from the geometry of the scene (Fig. 3).

Our goal is then to define a mapping (4) that allows us to lift line rendering from image-space into path-space, and derive a SRE-compatible style function that satisfies the following properties:

- (1) **Image-space consistency:** line width and placement are controlled in image-space and remain stable under indirect illumination, reflection, refraction, and depth of field.
- (2) **Geometric correctness:** the mapping from image-space to path-space respects local surface geometry, avoiding curvature-induced distortion.
- (3) **Distribution support:** the formulation applies uniformly to specular and non-specular (distributional) scattering.
- (4) **Estimator compatibility:** the resulting style function admits an unbiased Monte Carlo estimator with cost comparable to prior ray-based methods.

Rather than starting from previous methods, we take a step back and build our line rendering style function from the ground up over several subsections.

In Section 4.2 we decompose the geometrics of distribution effects as an aggregate over “conditionally-specular” paths. We then, in Sections 4.3 and 4.4, establish a way to move those conditionally-specular path across the image-plane that results in a well-defined




Fig. 4. A geometric path \bar{x} has an image-space coordinate s and traverses the scene up to a stylized edge yz (top). While purely specular transport provides a one-to-one mapping of image-space coordinates to paths (bottom, left), distribution effects like glossy reflections and depth-of-field result in uncountably infinite paths for point on the image-plane (bottom, right).

gradient. With a well-defined gradient, Section 4.5 shows how we reinterpret line detection as a stochastic search for extremal image-space gradients using finite differences. In Section 4.6, we turn this search process into an SRE-compatible style function and unbiased estimators. Section 4.7 then shows how the line rendering style function, and its estimators, can be recursively composed to support multiple line types. We include implementation details and a sampling algorithm in the supplemental document (Section S1.4).

4.1 Problem statement

For image-space consistent line rendering, determining if a line rendering style function $g_\theta(L(y, z))$ should return the line color or the unmodified exitant radiance is inherently *view-dependent*.

For a *geometric* base path \bar{x} (see Fig. 4 top) that traverses from the camera to the edge of stylization (y, z), and a neighborhood $S \subset \mathcal{S}$ around s , we want to determine whether there exists a point $s' \in S$ at which the image-space gradient of a measurement $f(s')$ exceeds a threshold t ,

$$\exists s' \in S \quad \text{s.t.} \quad \|\nabla_s f(s')\| \geq t. \quad (5)$$

When f is defined on path-space quantities (e.g., the terminal surface normal), evaluating $f(s')$ for $s' \neq s$ requires a mapping from image-space to paths, and, for $\nabla_s f$ to be well-defined, this mapping must vary smoothly with s .

For purely specular transport, this mapping is well defined: infinitesimal image-space perturbations induce unique, smoothly varying specular paths, yielding a notion of image-space parallelism, which prior work [Ogaki and Georgiev 2018] has successfully exploited.

However, under distribution effects, such as depth of field, glossy reflection, or rough refraction, a single image-space location maps to a distribution over paths (see Fig. 4 bottom). Differentiating expectations over these distributions produces gradients of *averaged measurements*, which is, as West [2021] put it, detecting *hard feature lines on blurry data*.

Our solution is to reinterpret distributional light transport as an aggregate of *conditionally-specular* paths.

4.2 A conditionally-specular view of distributions

Distribution effects can be expressed as integrals over families of specular interactions. Each realization of such a distribution corresponds to a specific choice of conditional variables ξ that fully determine the geometric behavior of a path. For example, ξ may encode a lens sample in a thin-lens camera or a sequence of microfacet half-vectors in a reflectance model.

An image-space location s and conditioning variable ξ fully specifies a unique path $\bar{x}(s, \xi)$. Paths constructed in this way are *conditionally specular*: conditioned on ξ , the path behaves as a classical specular path, even though the unconditioned transport is distributional. This lets us rewrite the path-space measurement as a function of image-space conditioned on ξ , where $f(\bar{x}(s, \xi)) = f(s, \xi)$.

With this conditionally-specular interpretation, we see that the unconditioned measurement $f(s)$ in Eq. (5) corresponds to an expectation over ξ , such that, $f(s) = \int_{\Xi} f(s, \xi) p(\xi) d\xi$, and differentiating this expectation yields,

$$\nabla_s f(s) = \int_{\Xi} \nabla_s f(s, \xi) p(\xi) d\xi, \quad (6)$$

under standard regularity assumptions (dominated convergence / Leibniz rule).

This expectation comes as a consequence of the distribution effects of light transport—exactly the effects we want to affect the rendered lines. Thus *we care about the integrand, not the integral*, and need to establish the partial gradient $\nabla_s f(s, \xi)$ for fixed ξ .

4.3 Parallel transport for conditioning variables

To make the gradient $\nabla_s f(s, \xi)$ well-defined for a given ξ , we need a way to construct a family of *auxiliary paths* $\{\bar{x}(s', \xi)\}$ for $s' \in S$ that both smoothly moves the conditioned specular events across the image-plane and preserves image-space *parallelism*.

Inspired by differential geometry, we establish this by introducing a *path transport operator* [Pinkall and Polthier 1993; Polthier and Schmie 1998; Crane et al. 2010; Knöppel et al. 2013; Sharp et al. 2019]. The operator enforces a partial-variation semantics: the image-space coordinate s varies while the conditioning variable ξ is held constant. Under this interpretation, two paths are *parallel* if they correspond to the same conditioned sequence of specular events realized at different neighboring image-space locations.

Formally, we define an operator \mathcal{T} that maps an image-space location s and conditioning variable ξ to a geometric path \bar{x} ,

$$\mathcal{T}(s, \xi) := \bar{x}(s, \xi), \quad (7)$$

and a corresponding differential operator $\mathcal{T}_{\delta s}$ that produces auxiliary paths under an infinitesimal image-space perturbation,

$$\mathcal{T}_{\delta s}(\bar{x}(s, \xi)) := \bar{x}(s + \delta s, \xi). \quad (8)$$

By construction, $\mathcal{T}_{\delta s}$ preserves the conditioned degrees of freedom encoded in ξ and modifies only those aspects of the path that are forced to change by moving the image-space coordinate.

With this operator in hand, path-space quantities evaluated along $\bar{x}(s, \xi)$ become smooth functions of image-space almost everywhere (barring those s for which $\bar{x}(s, \xi)$ intersects object boundaries, which




Fig. 5. The differential operator $\mathcal{T}_{\delta s}$ provides a way to move the tangent vector \mathbf{u} along surface path between two points \mathbf{x} and \mathbf{x}' . Using the surface normal \mathbf{n}' and the transport tangent vector \mathbf{u}' we can reconstruct the half-vector \mathbf{h}' at \mathbf{x}' .

have measure 0 in \mathcal{S}), and their partial gradients $\nabla_s f(s, \xi)$ are well-defined. The remainder of this subsection specifies how $\mathcal{T}_{\delta s}$ is realized for common distribution effects, beginning with depth of field and microfacet scattering.

Thin lens. Depth of field introduced by a thin-lens model can be interpreted as a distribution over camera ray origins on the lens aperture. Selecting a particular lens sample t fixes the camera frame and for any point s in image-space \mathcal{S} yields a purely specular ray from the camera onward.

In other words, for a given ξ , we simply need to copy the lens sample t from the base path $\bar{\mathbf{x}}(s, \xi)$ to any auxiliary path $\bar{\mathbf{x}}(s', \xi)$.

A microfacet view of surface scattering. Microfacet models [Torrance and Sparrow 1967] represent a wide range of materials—from diffuse-like surfaces [Oren and Nayar 1994] to glossy metals [Cook and Torrance 1982; Ashikmin et al. 2000] and rough dielectrics [Walter et al. 2007]—as distributions of infinitesimally small specular facets. Surface scattering is described by a normal distribution function (NDF) that specifies the statistical distribution of specular facet orientations [Beckmann and Spizzichino 1963; Trowbridge and Reitz 1975; Walter et al. 2007]. Although the aggregate surface interaction is stochastic and non-specular, each individual microfacet interaction is perfectly specular. A single surface interaction then corresponds to selecting a microfacet parametrized by its half-vector (i.e. facet normal) $\mathbf{h} \in \mathbb{S}^2$.

Parallel transport of parametrized paths. Similar in spirit to the primary sample space [Kelemen et al. 2002] or half-vector [Jakob and Marschner 2012] parametrizations for paths, we can specify a unique path from the image-space coordinate s and the conditioning variable ξ .

A geometric path under depth-of-field and glossy or rough scattering can be viewed as a conditionally-specular sequence containing the half-vectors \mathbf{h}_i and the reflection/refraction choices β_i ,

$$\bar{\mathbf{x}} = [(\mathbf{x}_1, t), (\mathbf{x}_2, \mathbf{h}_2, \beta_2) \dots (\mathbf{x}_k, \mathbf{h}_k, \beta_k)], \quad (9)$$

from which we can express the facet-normals \mathbf{h}_i in a local tangent space, denoted as \mathbf{h}_i^* ,

$$\xi = [t, \mathbf{h}_2^*, \beta_2, \dots, \mathbf{h}_k^*, \beta_k], \text{ such that,} \quad (10)$$

$$\bar{\mathbf{x}}(s, \xi) = [(\mathbf{x}_1, t), (\mathbf{x}_2, \mathbf{h}_2, \beta_2) \dots (\mathbf{x}_k, \mathbf{h}_k, \beta_2)]. \quad (11)$$

Given a parametrization ξ and a small perturbation δs , we can then construct an auxiliary path,

$$\bar{\mathbf{x}}(s + \delta s, \xi) = \mathcal{T}_{\delta s}(\bar{\mathbf{x}}(s, \xi)) = [(\mathbf{x}_1, t), (\mathbf{x}'_2, \mathbf{h}'_2, \beta_2) \dots (\mathbf{x}'_k, \mathbf{h}'_k, \beta_2)], \quad (12)$$

such that each \mathbf{h}'_i and β_i represents the *same conditioned specular event*, now expressed at the displaced vertex \mathbf{x}'_i .

To construct \mathbf{h}'_i in a geometrically consistent manner, we build parallel transport directly into the operator $\mathcal{T}_{\delta s}$. As a path vertex moves along a surface in response to an image-space perturbation, we transport the microfacet degrees of freedom so that the conditioned event remains the “same” across tangent spaces.

At the i -th vertex, conditioning fixes the cosine tilt,

$$c_i = \mathbf{h}_i \cdot \mathbf{n}_i, \quad s_i = \sqrt{1 - c_i^2}, \quad (13)$$

and we decompose the half-vector into normal and tangential parts,

$$\mathbf{h}_i = c_i \mathbf{n}_i + s_i \mathbf{u}_i, \quad \mathbf{u}_i \in T_{\mathbf{x}_i} \mathcal{S}, \quad \|\mathbf{u}_i\| = 1. \quad (14)$$

We then define the transported tangential direction by Levi-Civita parallel transport along the surface from \mathbf{x}_i to \mathbf{x}'_i ,

$$\mathbf{u}'_i = \text{PT}_{\mathbf{x}_i \rightarrow \mathbf{x}'_i}(\mathbf{u}_i), \quad (15)$$

where $\text{PT}_{\mathbf{x}_i \rightarrow \mathbf{x}'_i}$ is the differential operator $\mathcal{T}_{\delta s}$ applied along a path on the surface geometry from \mathbf{x}_i to \mathbf{x}'_i (see Fig. 5), and reconstruct the transported half-vector as,

$$\mathbf{h}'_i = c_i \mathbf{n}'_i + s_i \mathbf{u}'_i. \quad (16)$$

The choice of Levi-Civita transport is natural for our partial-variation interpretation—it is the only surface metric-compatible connection that is torsion-free (i.e. it preserves lengths and angles in the tangent plane) and *minimal-twist* (i.e. no rotation beyond what is forced by the change in local chart). The continuous differential derivation, including the uniqueness of $\mathcal{T}_{\delta s}$, is given in Section S1.1 of the supplemental document.

4.4 Choosing a local chart

In implementation we realize (16) with a convenient local chart choice. At the base vertex \mathbf{x}_i , we define a view-oriented local frame, $\mathbf{t}_{1,i} = \mathbf{n}_i \times \mathbf{v}_i$, and $\mathbf{t}_{2,i} = \mathbf{t}_{1,i} \times \mathbf{n}_i$, assuming the view direction \mathbf{v}_i is facing away from the surface, and represent the tangential component \mathbf{u}_i as,

$$\mathbf{u}_i = \cos \alpha_i \mathbf{t}_{1,i} + \sin \alpha_i \mathbf{t}_{2,i}, \quad \alpha_i = \text{atan2}(\mathbf{u}_i \cdot \mathbf{t}_{2,i}, \mathbf{u}_i \cdot \mathbf{t}_{1,i}). \quad (17)$$

At the auxiliary vertex \mathbf{x}'_i , we construct the corresponding view-oriented frame, $\mathbf{t}'_{1,i} = \mathbf{n}'_i \times \mathbf{v}'_i$, and $\mathbf{t}'_{2,i} = \mathbf{t}'_{1,i} \times \mathbf{n}'_i$, and reconstruct

$$\mathbf{u}'_i = \cos \alpha_i \mathbf{t}'_{1,i} + \sin \alpha_i \mathbf{t}'_{2,i}, \quad (18)$$

from which we obtain \mathbf{h}'_i via (16).

The outgoing direction \mathbf{d}'_i at the auxiliary vertex follows from either specular microfacet reflection,

$$\mathbf{d}'_i = \mathbf{v}'_i - 2(\mathbf{v}'_i \cdot \mathbf{h}'_i) \mathbf{h}'_i, \quad (19)$$

or refraction,

$$\mathbf{d}'_i = \eta \mathbf{v}'_i - \left(\eta (\mathbf{v}'_i \cdot \mathbf{h}'_i) + \sqrt{1 - \eta^2 (1 - (\mathbf{v}'_i \cdot \mathbf{h}'_i)^2)} \right) \mathbf{h}'_i, \quad (20)$$

depending on β_i , where $\eta = \eta_i / \eta_t$ for the incident η_i and transmitted η_t indices of refraction at the base path vertex \mathbf{x}_i . In the context of paths, this relationship is recursive: the outgoing direction \mathbf{d}'_i becomes the view direction \mathbf{v}'_{i+1} for the next auxiliary vertex. Fig. 6 visualizes this recursive path construction process.




Fig. 6. Auxiliary paths are constructed recursively. Starting from the fixed lens sample and new image-space coordinate s' (a), a ray is cast to determine the first hit point x'_1 (b), the half-vector h'_1 is transported to x'_1 (c), and the outgoing direction d'_1 is computed (d). The outgoing direction d'_1 then becomes the incoming direction at x'_2 , and the process repeats.

Given our fixed ξ , base path \bar{x} , and transport operator $\mathcal{T}_{\delta s}$, we can now map points s in some region S of image-plane to geometric paths that give us access to path-space quantities in an image-space consistent manner. For example, we can define a measurement function $f(s', \xi) = \mathbf{n}(\bar{x}(s', \xi))$ that returns the surface normal at the last vertex of the path $\bar{x}(s', \xi)$.

The specular case. Purely specular paths require no special handling: specular reflection has no azimuthal degree of freedom, and the operator $\mathcal{T}_{\delta s}$ reduces to classical image-space parallelism in the specular limit. Our formulation therefore recovers prior ray-based line rendering behavior as a special case.

4.5 The search for extremal gradients

Now that we have an image-space consistent measurement function $f(s, \xi)$, we need to establish a gradient that we can use to determine if we draw the line color or not.

As the transport operator $\mathcal{T}_{\delta s}$ considers an infinitesimal perturbation δs , it gives rise to a well-defined partial gradient $\nabla_s f(s', \xi)$ *almost everywhere*. Interestingly, there is a measure-zero subset of image-space where this gradient is undefined, corresponding to object silhouettes, grazing-angle microfacet degeneracies, etc.—exactly where we want to draw lines. This poses no issue in finding line-producing gradients, as regions of undefined gradient correspond to discontinuities in the transport which we treat as a limiting case on arbitrarily large gradients (see Supplemental Section S1.2).

In practice, there is often no closed-form for the partial gradient $\nabla_s f(s', \xi)$, making it impossible to formulate closed-form expressions for the threshold problem (e.g. computing the intersection of the threshold plane with the gradient field). We can, however, compute the primal value $f(s', \xi)$ at arbitrary points $s' \in S$, allowing us to use point sampling methods.

One obvious choice would be to estimate the gradient magnitude at point samples in the region S , but this can be slow to converge and produce false positives (see Supplemental Section S1.3).

Finite difference as a lower bound. A more stable, and faster converging option, is to use the finite difference $\frac{|f(s', \xi) - f(s, \xi)|}{\|s' - s\|}$ of point samples s and s' to establish a *lower bound* on the gradient $\|\nabla_s f(s^*, \xi)\|$. By the median value theorem, there exists a point s^* on the line segment between s and s' such that,

$$\|\nabla_s f(s^*, \xi)\| \geq |D_{s \rightarrow s'}^{(s)} f(s^*, \xi)| = \frac{|f(s', \xi) - f(s, \xi)|}{\|s' - s\|}. \quad (21)$$

Note that the finite difference here is not an approximation, it is exactly equivalent to the directional derivative $|D_{s \rightarrow s'}^{(s)} f(s^*, \xi)|$ and strict lower bound on the gradient at s^* . It follows that if any finite difference lower bound in S exceeds the threshold t , then necessarily,

$$\max_{y \in S} \|\nabla_s f(y, \xi)\| \geq t, \quad (22)$$

satisfying Eq. (5) (see Supplemental Section S1.2).

The search continues. While a single instance of finite difference provides a lower bound, that lower bound may not be tight (i.e. it is not the infimum), and it can fail to prove the existence of an extremal gradient. Let us quantify this probability of failure and explore how it diminishes as more sample pairs are drawn.

We can represent whether a single instance of finite-difference fails to exceed the target threshold t as a metric function m_ξ for a fixed ξ that returns 1 if the threshold is not crossed, else 0,

$$m_\xi(u, v) = \begin{cases} 1, & \frac{|f(v, \xi) - f(u, \xi)|}{\|v - u\|} < t \\ 0, & \text{otherwise.} \end{cases} \quad (23)$$

In practice, the metric m_ξ can be almost any function that returns 0 or 1, providing significant flexibility in specifying where lines are detected.

If we stochastically select a pair of points u and v from an image-space region S around the point s , we can then formulate the expectation of the metric m_ξ as,

$$\begin{aligned} \mathbb{E}[m_\xi(u, v)] &= \int_{S^2} m_\xi(u, v) p(u, v) \, dudv \\ &= P \left(\frac{|f(v, \xi) - f(u, \xi)|}{\|v - u\|} < t \right), \end{aligned} \quad (24)$$

which is exactly the likelihood of the finite difference for a single pair (u, v) *not* exceeding the threshold t , i.e. not finding a line-producing gradient. If we then draw n independent pairs of (u, v) , the expectation of failure is simply the power of expectation,

$$\mathbb{E}_n[m_\xi(u, v)] = (\mathbb{E}[m_\xi(u, v)])^n = \left(\int_{S^2} m_\xi(u, v) p(u, v) \, dudv \right)^n, \quad (25)$$

which falls off exponentially fast in n (see Fig. 7). Conversely, the probability of finding at least one pair (u, v) with a line-producing gradient in n tries converges quickly towards 1, and is high when $n \geq \frac{1}{1 - \mathbb{E}[m_\xi(u, v)]}$.

At the limit (i.e. exhaustive search), the lower bound established by finite differences is guaranteed to converge to the infimum which is equivalent to the maximal gradient. We provide a thorough discussion and proof in the supplemental document.

4.6 A line rendering style function

By considering a simple process where we output the line color $c(y, z)$ if we successfully find a line-producing gradient in n sample pairs, else the exitant radiance $L(y, z)$, we come to an expected stylized exitant radiance,

$$g_\theta(L(y, z)) = (1 - \mathbb{E}_n[m_\xi(u, v)]) c(y, z) + \mathbb{E}_n[m_\xi(u, v)] L(y, z), \quad (26)$$

arriving at what is effectively an SRE-compatible style function that follows from the search process.

General form. We can think of the expectation (25) of not finding a line-producing gradient in n tries as an equivalent “blend weight” w_θ . In other words, w_θ specifies where lines are drawn and how opaque they are.

This allows us to decouple the sample count from line opacity, making n an artist-controllable parameter n_θ ,

$$w_\theta = \left(\int_{S_\theta^2} m_{\theta,\xi}(u,v) p_\theta(u,v) dudv \right)^{n_\theta}; \quad (27)$$

we note here the term $p_\theta(u,v)$ can feasibly be any artist-controllable function that satisfies $\int_{S^2} p_\theta(u,v) dudv = 1$. However, the same effect can be achieved when $p_\theta(u,v)$ is the pdf of sample pairs by controlling the sample distribution.

Replacing the expectations in Eq. (26) with blend weights (27), we arrive at the general form of the line rendering style function,

$$g_\theta(L(y,z)) = (1 - w_\theta) c_\theta(y,z) + w_\theta L(y,z), \quad (28)$$

a simple linear interpolation of two colors.

Let us note that we are not limited to this specific form. The blend weight w_θ expresses where lines are drawn, allowing arbitrary non-linear stylizations at the location of lines. For example, a saturation stylization that makes the silhouette of objects more colorful.

Estimation. An unbiased estimator of the style function (28) can be constructed for a specified power n by drawing n samples $(u,v) \sim p_\theta(u,v)$ of the metric function $m_{\theta,\xi}(u,v)$, one per instance of the inner-integral, and taking their product,

$$\langle w_\theta \rangle = \prod_{i=1}^n m_{\theta,\xi}(u_i, v_i), \quad (29)$$

such that,

$$\langle g_\theta(L(y,z)) \rangle = \left(1 - \prod_{i=1}^n m_{\theta,\xi}(u_i, v_i) \right) c_\theta(y,z) + \prod_{i=1}^n m_{\theta,\xi}(u_i, v_i) L(y,z), \quad (30)$$

which is an unbiased estimator of Eq. (28) for almost any $p(u,v)$.

Coming full circle, this estimator corresponds exactly to the search process of drawing n sample pairs, and returning the line color $c_\theta(y,z)$ if any of their metrics are zero, $m_{\theta,\xi}(u,v) = 0$.

Sample count n . In the case that $N > n$ samples are drawn for a fixed n , a significantly lower variance estimate of w_θ can be computed using the variance-minimal recurrence relation of Kettunen et al. [2021]. Additionally, while one can feasibly consider fractional values for n , the added utility is minimal, and requires much more heavy-handed estimation strategies (e.g. stochastic truncation of a Taylor-series expansion [Georgiev et al. 2019; Misso et al. 2022]).

Bias. Although the estimators introduced above are unbiased for any fixed power n , the formulation itself exhibits a bias relative to the ideal of perfectly hard lines. As line detection relies on stochastic point sampling over the region S , truly hard lines can only be recovered in finite time when the n -sample finite-difference expectation (25) evaluates to zero for finite n . For a fixed sampling

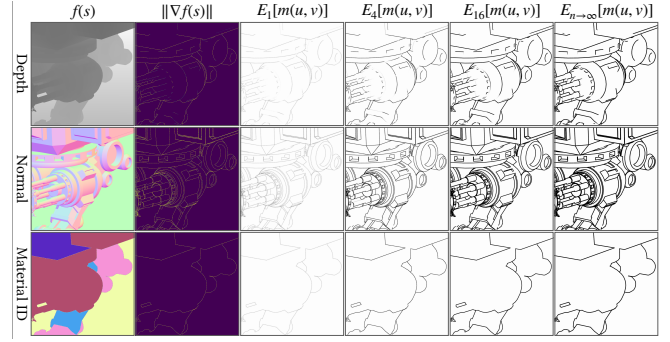


Fig. 7. Rows use depth, surface normals, and material IDs as the measurement field $f(s)$, and we visualize the corresponding gradient magnitude $\|\nabla f(s)\|$ estimated via a brute force finite-difference approximation. The remaining columns visualize the expectation of the n -sample formulation in Eq. (25), $E_n[m_\xi(u,v)]$, for increasing n (and a converged large- n reference). In this example, $n = 16$ already produces a search expectation that is visually close to the large- n result, while smaller n results in softer line edges. The appropriate choice of n is density-dependent (i.e., it depends on the area of the line stencil). The same extremal-gradient construction works for various line producing features like geometric discontinuities (depth), orientation changes (normals), or semantic boundaries (material IDs).

distribution $p(u,v)$, the probability of failing to detect a line decays super-linearly with n , mirroring the line-softening observed in ray-based methods such as West [2021]. Practical implementations will set n sufficiently high, e.g. $n \geq (1 - E[m_\xi(u,v)])^{-1}$, to meet artistic goals. In practice, we found 12 – 16 samples works well, with larger line widths tending to require more samples.

4.7 Multiple line types

Feature line renderings typically consist of multiple line types, each defined by, for example, a distinct measurement function f , threshold t , measurement region (i.e. line stencil) S , and color c_θ . Existing ray-based methods handle this by evaluating several candidate line types per ray and resolving visibility through an ordering rule, commonly based on depth along the ray and optionally augmented with user-defined priorities [Ogaki and Georgiev 2018; West 2021].

This behavior arises naturally in the SRE by treating each line type as its own style function and composing these functions in an order consistent with the desired occlusion rule.

Given n (potentially) different line types, each with an instance of the line style function g_{θ_i} in Eq. (28) and their line-specific parametrization θ_i , we can define a composite style function,

$$\bar{g}_{n,\Theta} = (g_{\theta_1} \circ g_{\theta_2} \circ \dots \circ g_{\theta_n}), \quad \Theta = \{\theta_1, \dots, \theta_n\}, \quad (31)$$

where the composition order reflects the chosen occlusion ordering (e.g. increasing distance along the ray, or artist-chosen line priority).

Expanding the composition yields,

$$\bar{g}_{n,\Theta}(L) = \sum_{i=1}^n \left(\prod_{j=1}^{i-1} w_{\theta_j} \right) (1 - w_{\theta_i}) c_{\theta_i} + \left(\prod_{m=1}^n w_{\theta_m} \right) L, \quad (32)$$

where, for readability, the dependence on the edge (y,z) is omitted.

Stochastic composition. In practice, there is no way to know a priori the exact set of line types that a given edge (y, z) will intersect. Prior works dealt with this by implicitly using point samples s'_1, \dots, s'_n (often the same as those used in large gradient detection) to search for line-producing surfaces. Although this process is often described algorithmically, it defines a mapping between point samples and line types, and produces a stochastic selection of ordered line types.

We can employ a similar approach under the proposed formulation, giving us a stochastic composition of style functions. The expectation of the stochastic composition quantifies the behavior observed in prior works where lines of different types blend at intersections and overlapping regions following from an expectation,

$$g_\theta(L(y, z)) = E_n[\bar{g}_{n,\Theta}(L(y, z))] \\ = \int_{S_\theta^n} \bar{g}_{n,\Theta}(L(y, z)) p_\theta(s'_1 \dots s'_n) ds'_1 \dots ds'_n, \quad (33)$$

where each image-plane point s'_i is mapped to a style function $g_{\theta_i} \in \bar{g}_{n,\Theta}$, and $p_\theta(s'_1 \dots s'_n)$ can be any weighting function satisfying $\int_{S_\theta^n} p_\theta(s'_1 \dots s'_n) ds'_1 \dots ds'_n = 1$, though having it match the probability density of samples is convenient when estimating.

Estimation. An estimator for the composite style function can be constructed by estimating the blend weight w_{θ_i} of each of the i style functions g_{θ_i} ,

$$\langle g_\theta(L) \rangle = \sum_{i=1}^n \left(\prod_{j=1}^{i-1} \langle w_{\theta_j} \rangle \right) (1 - \langle w_{\theta_i} \rangle) c_{\theta_i} + \left(\prod_{m=1}^n \langle w_{\theta_m} \rangle \right) L, \quad (34)$$

where we assume the joint probability density of samples $p_\theta(s'_1 \dots s'_n)$ follows from the sampling process and the term gracefully vanishes.

Note that, while correlation resulting from sample reuse across weight estimates is partially masked by early termination of the product, strict unbiasedness requires the samples of the blend weight estimates $\langle w_{\theta_i} \rangle$ in each of the $n + 1$ product terms, $\prod_{j=1}^{i-1} \langle w_{\theta_j} \rangle (1 - \langle w_{\theta_i} \rangle)$ and $\prod_{m=1}^n \langle w_{\theta_m} \rangle$ to be independent.

When using a binary metric $m_{\theta,\xi}(u, v)$, each product will evaluate to 0 or 1, and the composite estimator reduces to selecting exactly the color c_{θ_i} of the first line type in the ordering whose search detects a sufficiently large finite difference—all subsequent terms vanish, and no additional sampling is required. The unstylized radiance $L(y, z)$ is the *only* recursive term in (34), and is only ever computed if every weight estimate evaluates to 1, potentially early terminating recursive sampling. See supplemental Section S1.4 for further details.

This recovers the behavior of prior methods: the rendered line color corresponds to the shallowest (or highest-priority) intersected line type along the ray. Importantly, this behavior emerges from style function composition and does not require special-case logic.

Overall, the proposed feature line framework lets us fuse the consistency of image-space line rendering with the expressiveness of a path-space formulation. By interpreting surface scattering and camera effects as conditionally specular events, we preserve a principled notion of image-space parallelism, letting us draw lines of arbitrary color, width, and falloff across complex geometry, depth of




Fig. 8. Conceptually, a conformal one-to-one map ϕ_k from the image space S to a manifold \mathcal{M} (left) admits a *unique* conformal one-to-one extension from an enlarged domain to an enlarged connected subset of \mathcal{M} (right). We leverage this concept to extend tone rendering beyond the region of path-space covered by the anchor mapping.

field, and distributional effects; all while remaining compatible with ordinary style functions under the stylized rendering equation.

5 TONE RENDERING

Tone-based styles such as hatching and halftone convey form, material, and illumination through *image-space-consistent structure*. Their visual identity arises from coherent spatial patterns—e.g., line spacing, orientation, or dot density—that are defined and perceived in the image-plane.

When tone stylization is expressed directly as a path-space style function under the stylized rendering equation, this inherent image-space structure is lost. As style functions are evaluated independently along light paths, the resulting patterns are parameterized by surface geometry and distribution effects like surface scattering rather than image-space coordinates. In practice, this leads to characteristic artifacts such as geodesic distortion and depth-dependent variation in pattern scale.

Our goal is therefore to lift image-space tone fields into a path-space formulation while preserving their defining image-space structure. In contrast to line rendering, which relies on conditional correspondence between nearby image-space locations, tone rendering requires a *canonical lifting* that anchors patterns to a material-independent image-space parameterization. This lifting must associate path vertices \mathbf{x} on the scene manifold \mathcal{M} with a meaningful image-space coordinate at which the tone field can be queried.

Moreover, this association must remain valid under complex light transport, including indirect illumination, glossy reflection, refraction, and depth of field. In particular, tones observed through distribution effects should preserve their image-space organization while being naturally affected (e.g. blurred) by the transport.

We frame this as the construction of a mapping $\phi : \mathcal{S} \rightarrow \mathcal{P}$, $\phi = \Phi(\cdot, \xi^*)$ for a canonical conditioning variable $\Xi = \{\xi^*\}$ of image-space coordinates s to paths, whose vertices \mathbf{x} lie at mapping destinations along the scene manifold \mathcal{M} , together with evaluating a practical approximation of its inverse ϕ^{-1} at arbitrary path vertices.

While, in principle, any locally coherent and invertible projection could serve this purpose, we argue that ϕ must satisfy two additional criteria: it should be *material-invariant*, so that tone structure is not tied to specific surface properties, and it should be *aligned with light transport*, so that distribution effects are handled consistently with the underlying path sampling process.

5.1 The mirror reflection projection ϕ

A well-motivated and style-agnostic choice for ϕ is the mapping induced by mirror reflection about *macro* surface geometry. That is, for a given image-space coordinate, ϕ traces a path obtained by mirror reflection at every surface interaction, independent of the actual material. This construction provides a canonical, material-invariant correspondence between image-space and path-space.

From a physical standpoint, this choice aligns with the geometry of light transport. In a wave-optics interpretation, the mirror reflection path corresponds to a stationary-phase (and typically dominant) contribution of light, and distribution effects effectively integrate around this path. Under the Huygens–Fresnel formulation, the complex field at a point $\mathbf{x} \in \mathcal{M}$ after interaction with a surface $M \subset \mathcal{M}$ can be written, up to slowly varying amplitude and visibility terms, as

$$U(\mathbf{x}) \propto \int_M U(\mathbf{y}) A(\mathbf{x}, \mathbf{y}) e^{i\nu d(\mathbf{x}, \mathbf{y})} dA(\mathbf{y}), \quad \nu = \frac{2\pi}{\lambda},$$

where $d(\mathbf{x}, \mathbf{y})$ denotes optical path length. In the high-frequency regime relevant to tone patterns, stationary-phase arguments imply that the dominant contribution arises from points $\mathbf{y}^* \in M$ satisfying

$$\nabla_M d(\mathbf{x}, \mathbf{y}^*) = 0,$$

i.e., the gradient of d projected onto the tangent space of M vanishes, precisely the specular (Fermat) condition [Born and Wolf 2002]. For multi-bounce transport, the accumulated optical path length,

$$d(\mathbf{x}_0, \dots, \mathbf{x}_k) = \sum_{i=0}^{k-1} d(\mathbf{x}_i, \mathbf{x}_{i+1}),$$

(and thereby phase) must be stationary with respect to each intermediate vertex $\mathbf{x}_i \in \mathcal{M}$, yielding perfectly specular reflection along the entire path.

This motivates the mirror-reflection projection as a natural *anchor* for lifting image-space tone structure into the scene. It is locally coherent, independent of material parameters, and aligned with the characteristic flow of light propagation. Once tones are anchored to this specular structure via ϕ , distribution effects such as glossy blur or depth of field can be interpreted as neighborhood integration around this stationary path, rather than as deformations of the underlying image-space pattern.

Despite these desirable properties, the mirror-reflection projection has three fundamental limitations:

- (1) **Partial coverage:** ϕ is not defined for all path vertices; many points in \mathcal{M} admit no valid pre-image in image-space.
- (2) **Non-bijection:** The mapping is neither one-to-one nor onto: multiple image-space points may map to the same scene point, while large regions of \mathcal{M} may not be reachable.
- (3) **Implicit form:** ϕ is only accessible through forward path construction; it generally has no closed-form expression and no analytic inverse.

5.2 Analytic continuation

To build intuition about the behavior of the mirror-reflection projection ϕ , we consider a simplified setting rather than a fully rigorous one. Suppose we look at a small region S of the image plane, and

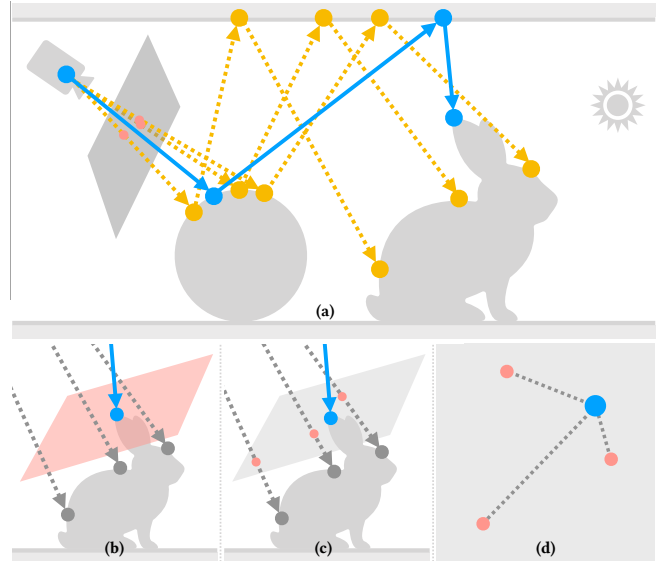


Fig. 9. Even though the mirror anchor only assigns image-space coordinates to a subset of vertices (a), these coordinates can be propagated across a connected region of the scene via overlapping local neighborhoods (yellow paths), yielding a consistent coordinate for a target vertex (blue path). In practice, each continuation step reduces to a local interpolation/extrapolation problem (b–d): given nearby vertices with known image-space coordinates, we estimate the query vertex’s coordinate by working in a local tangent-plane chart and fitting a smooth, locally invertible map (e.g. via linear moving least squares).

follow its mapping onto the scene manifold \mathcal{M} after exactly k specular bounces, which we denote by ϕ_k . *Locally*, this mapping behaves smoothly and preserves angles (i.e. conformal), meaning that small image-space neighborhoods are transferred to the scene without shearing or folding.

If this local mapping is also one-to-one, classical results from complex analysis tell us that such a map can be extended beyond its original domain in a consistent way: knowing how ϕ_k behaves on S effectively determines how it should behave in a larger surrounding region. Intuitively, the mapping can be “continued” outward while preserving its local structure, yielding a consistent correspondence between nearby parts of the scene and image-space (see Fig. 8).

As a result, every point $\mathbf{x} \in \mathcal{M}$ connected to $\phi_k(s)$ would, in principle, admit a well-defined image-space coordinate $\tilde{\phi}_k^{-1}(\mathbf{x})$. A formal statement and proof are given in Theorem S2.1 of the Supplemental.

While this idealized scenario does not arise in practical rendering, it motivates an intuitive observation: for tone stylization, ϕ_k need only be *locally* coherent and approximately invertible. In practice, given a path vertex $\mathbf{x} \in \mathcal{M}$ and a set of nearby points $\mathbf{x}_1, \dots, \mathbf{x}_\ell$ in the image of ϕ_k , it suffices to assign a consistent image-space coordinate to \mathbf{x} as long as all points lie within the same connected region of \mathcal{M} . On sufficiently small patches of image-space, the mirror-reflection mapping is smooth, locally one-to-one, and conformal, and thus locally invertible by the inverse function theorem.




Fig. 10. Comparison of Ogaki and Georgiev [2018] (left), West [2021] (middle), and our method (right) on a scene containing depth of field and both specular and glossy reflections. Under directly visible surfaces (no DoF) all methods agree up to minor implementation differences (**red**). For *flat* glossy mirror reflections, our method preserves the characteristic softened appearance of West [2021] (**yellow**). Under depth-of-field the method of Ogaki and Georgiev [2018] breaks down while West [2021] and our method remain coherent (**green**). For *curved* specular reflections, our method preserves image-space line width similarly to Ogaki and Georgiev [2018], while West [2021] exhibits width drift (**blue**).

5.3 Local approximations of analytic continuity

Our objective is therefore not to recover a true inverse of ϕ_k , but to construct a local, numerically stable approximation that behaves like an analytic continuation where sufficient information is available, and degrades gracefully elsewhere. Concretely, given a path vertex $\mathbf{x} \in M$ for which $\phi_k^{-1}(\mathbf{x})$ is unknown, and a set of nearby vertices $\mathbf{x}_1, \dots, \mathbf{x}_\ell \in M$ with known image-space coordinates $\phi_k^{-1}(\mathbf{x}_i)$, our goal is to assign a consistent image-space location $\tilde{\phi}_k^{-1}(\mathbf{x}) \in \mathbb{R}^2$. This problem can be viewed as a local interpolation or extrapolation task on the scene manifold, guided by the mirror-reflection geometry.

A variety of strategies have appeared in graphics for similar value estimation problems, including barycentric extrapolation, Möbius transformations, and mean-value coordinates. While these methods are inexpensive and easy to implement, they either lack conformal behavior, are unstable under extrapolation, or exhibit visible artifacts near boundaries. We summarize these alternatives in the supplemental material (Section S2.2 and Section S2.4), and focus here on the method we found most robust in practice.

Moving Least Squares. We adopt a Moving Least Squares (MLS) [Levin 1998] formulation to locally approximate ϕ_k^{-1} . While technically not conformal, MLS provides a smooth, locally invertible mapping that is stable under moderate extrapolation and degrades gracefully as samples become sparse.

For the projection $\mathbf{p}_i = \pi(\mathbf{x}_i) \in \mathbb{R}^2$ of nearby vertices onto a local tangent plane at \mathbf{x} , their known image-space coordinates $\mathbf{s}_i = \phi^{-1}(\mathbf{x}_i) \in \mathbb{R}^2$, and the projected query point $\mathbf{p} = \pi(\mathbf{x})$, MLS fits a local affine map that best explains the correspondence $\mathbf{p}_i \mapsto \mathbf{s}_i$ in a weighted least-squares sense.

Specifically, we solve,

$$\min_{\mathbf{A}, \mathbf{b}} \sum_{i=1}^{\ell} w_i(\mathbf{x}) \|\mathbf{A} \mathbf{p}_i + \mathbf{b} - \mathbf{s}_i\|^2 \rightarrow \mathbf{s} = \mathbf{A} \mathbf{p} + \mathbf{b}, \quad (35)$$

where the weights $w_i(\mathbf{x})$ decay with the distance $\|\mathbf{p} - \mathbf{p}_i\|$. The computed image-space coordinate \mathbf{s} can then be used to query any

image-space quantities (e.g. tone patterns) or compute image-space shaders at path vertices.

In practice, we observed linear MLS strikes a favorable balance between stability and visual quality, producing significantly fewer fold-overs and boundary artifacts than alternatives. Higher-order variants, such as quadratic MLS [Joldes et al. 2015], improve local accuracy but tend to amplify noise and introduce geodesic artifacts when extrapolating beyond the convex hull of samples.

In Section S2.2 of the supplemental document we provide comparisons of several local inversion methods, and specifics on the choice of tangent plane and methods for selecting nearby vertices.

Overall, MLS provides a practical, locally invertible approximation to analytic continuation. It preserves image-space coherence where samples are available, degrades smoothly in under-sampled regions, and integrates naturally into Monte Carlo path sampling without requiring global solves or explicit parameterization.

6 RESULTS

We demonstrate our proposed image-space-consistent line and tone stylizations on a variety of scenes with indirect illumination, specular and glossy reflection, and depth-of-field blur. Figs. 10 to 12 (Section 6.1) validate the feature-line contribution, Figs. 13 and 14 (Section 6.2) then validate the tone contribution, while Fig. 1 (Section 6.3) illustrates that these line and tone correspondences also compose cleanly within a single stylized transport pipeline.

All results were rendered with our CPU-based RGB renderer using the tree-based sampling algorithm of West and Mukherjee [2024] with our sampling modifications used to construct a cache of paths for the canonical mapping used in tone rendering, and the auxiliary path sampling along specific path traversals of the tree-sample for line rendering (see Supplemental Section S1.4 and Section S2.3).

All timings are based on rendering performance of a Macbook Pro with an M3 Max processor and 128GB RAM. Unless otherwise noted, comparisons use the same sampling budgets to isolate how each method behaves under the same transport phenomena.




Fig. 11. Our line formulation supports a range of artist-controllable line looks within a single rendering. From left to right: **(blue)** layers two distinct line styles, **(purple)** combines a very thick silhouette with a thin inner line and uses a square line stencil to sharpen corners, **(magenta)** uses displaced lines for a sketchy/offset appearance, **(pink)** uses a different line color for glossy reflections, and **(red)** shows a hatched line style.

6.1 Feature line rendering

Fig. 10 compares our feature-line rendering against the methods of Ogaki and Georgiev [2018] and West [2021] under a combination of direct visibility, glossy and specular reflection, and depth of field. For directly visible surfaces (red), all methods agree up to minor implementation differences and sampling noise. Method differences become more pronounced when lines are seen *through* transport phenomena: Ogaki and Georgiev [2018] rely on specular image-space parallelism that breaks down under glossy reflection (yellow) and depth of field (green), while West [2021] remains coherent under distribution effects but exhibits noticeable line-width drift in curved reflections (blue). Our method preserves consistent image-space line width across curved geometry and through both glossy and specular reflection, while still allowing lines to respond naturally to physical blur from glossy transport and depth of field.

Fig. 11 highlights that our line formulation does not constrain artistic expressiveness: the same geometric correspondence supports a range of artist-controllable line looks within a single rendering. We show layered line styles (blue), silhouette/inner-line decomposition with stencil-sharpened corners (purple), displaced lines for a sketchy offset appearance (magenta), reflection-specific line coloring under glossy transport (pink), and hatched line fills (red). These variations are realized through ordinary SRE style functions on top of the same auxiliary-path construction, so style design remains flexible while image-space line structure remains stable.

Fig. 12 provides a more controlled diagnosis of the line-width behavior of each method by isolating two reflector geometries (flat vs. curved) and two scattering regimes (perfectly specular vs. glossy). In the specular cases, our construction matches the expected reference behavior and preserves line width on both flat and curved reflectors. In the glossy cases, we retain the characteristic softened appearance produced by distributional transport, while avoiding the additional curvature-induced thinning/thickening that arises when auxiliary paths do not account for surface curvature. Importantly, effects

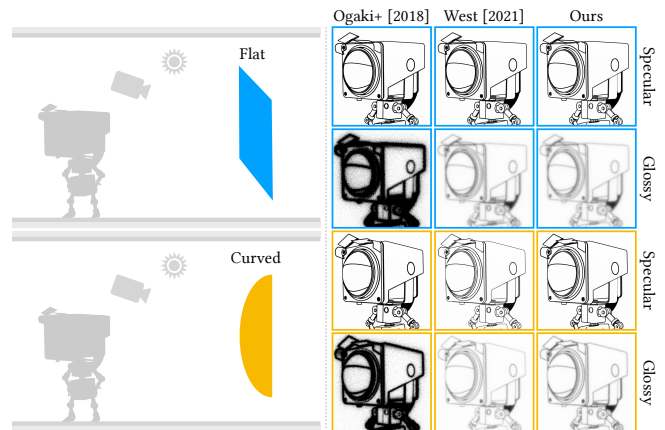


Fig. 12. Detailed comparison of *image-space line width* under reflection. Left shows the scene layouts with a flat and curved mirror respectively. For each, we test a perfectly specular mirror (top) and a glossy mirror (bottom). Columns compare Ogaki and Georgiev [2018], West [2021], and ours. Our method matches the specular reference behavior on both flat and curved reflectors, and in the glossy case preserves West’s characteristic softening while avoiding curvature-induced line-width drift.

such as curvature-following distortion or changes in line width can be aesthetically desirable in some cases—our method makes those effects controllable stylistic choices rather than unavoidable consequences of the underlying formulation.

6.2 Tone rendering

Fig. 13 compares tone (e.g. hatching) rendering using our method, the Deussen et al. [1999]-inspired world-space algorithm of West and Mukherjee [2024], and the equivalent image-space stylization performed directly on the output buffers. All methods produce nearly identical visuals for directly viewed surfaces close to the camera




Fig. 13. We render the same scene twice: once *purely specular* (left half of each main view) and with *depth of field + glossy reflections* (right half). We compare an image-space post-process baseline (left), the world-space tone formulation of West and Mukherjee [2024] (middle), and our image-space-consistent tone rendering for the same shader as the image-space post-process (right). Insets correspond to the colored boxes: red is a reflection in a flat perfect mirror, yellow is an in-focus region (no DoF), green lies where DoF gradually develops, and blue is a reflection in a curved glossy mirror. For this scene, all three methods match closely in the in-focus region (**yellow**). We see that West and Mukherjee [2024]’s world-space construction (plane intersections with the scene) is tuned to appear camera-aligned from the main view, but when the same geometry is seen at a grazing angle, the intersection directions are no longer camera-aligned and the pattern no longer reads as hatching (**red**). Our method preserves the intended image-space orientation and scale across both renders, while the DoF/glossy case exhibits the expected transport-induced blur without re-parameterizing the underlying pattern (**green/blue**).

(yellow). However, world-space-based tone stylization can result in geodesic artifacts in off-angles and reflections (red). In contrast, image-space stylizations run on top of the rendered output and, while fast, produce hard lines on naturally blurry regions (blue). Our method captures the structure of the image-space tone, and inherits the visual queues of complex light transport, but at the cost of additional overhead. Artists can weigh the trade-off between support for distribution effects and rendering overhead.

Additionally, Fig. 13 demonstrates that our line and tone constructions compose cleanly in a single rendered image. Feature lines remain stable in width across curved geometry and through distribution effects, while tone patterns preserve their image-space structure even when seen through reflective transport and under global illumination. As our method is based on ordinary (possibly nonlinear) SRE style functions, the same pipeline supports combining multiple line types with multiple tone fields (e.g., hatching and half-tone) without sacrificing transport-aware behavior or coherent image-space structure.

Fig. 14 demonstrates image-space-consistent tone in a single scene with five objects assigned distinct tone fields (pattern, tint, and isotropy/anisotropy). Across all objects, tone appearance varies with shading: highlights and shadows affect tone intensity (e.g., pattern brightness and/or coverage) without re-parameterizing the underlying image-space structure. This makes per-object tone choices readable and stable while still reflecting illumination differences.

The glossy floor further illustrates transport-consistent tone behavior in reflection. As the reflected contribution becomes more distributional, the tone pattern exhibits a smooth transition from crisp to blurred appearance, matching the expected glossy spread while preserving coherent image-space scale and orientation. In other words, tone blurs because transport blurs—not because the pattern is stretched or distorted by surface parameterization. In addition, just like any ordinary style function, our tone rendering style

functions can change depending on their parametrization. For example, in Fig. 14 the reflection of the right-most robot is color-shifted and the hatching orientation is intentionally flipped to look like a reflection—simulating the look-and-feel of a world-space stylization.

6.3 Combined stylizations

The teaser in Fig. 1 provides a glimpse at some of the stylization that can be achieved with the proposed line and tone rendering methods. Beyond line and tone rendering it employs a few other subtle style functions, leveraging that our methods can be composed and rendered cohesively with any other style function. Here we used several brightness-level controlling style functions on emission and reflected lighting to surgically adjust the look-and-feel of the render, bringing more light to the foreground while maintaining high contrast, and lowering the intensity of background regions—the tone stylization naturally responds to these changes in lighting. The background is also entirely real geometry, not a billboard, demonstrating a unified look between near-field and far-field line rendering.

7 LIMITATIONS AND FUTURE WORK

Feature line rendering. Like prior ray-based line rendering methods, our approach operates through *local* stochastic tests along sampled paths (finite-difference detection of extremal gradients), rather than extracting explicit line primitives. Consequently, we can preserve line *appearance* (placement and width in image-space) under complex transport, but we do not directly recover global stroke structure. In other words, we have no explicit representation of line connectivity, endpoints, or topology. Enriching the method to construct editable line primitives (e.g., with support for operations such as trimming, merging, or higher-level stroke textures) is a very interesting direction of future work.

Our estimator also inherits the usual accuracy–cost tradeoffs of Monte Carlo methods. With limited samples, line detection can miss




Fig. 14. Five objects rendered with different tone fields (pattern, tint, and isotropy/anisotropy), illustrating per-object tone control within a single scene. Across all objects, the tone pattern (brightness and/or pattern thickness) follows the underlying shading, so highlights and shadows modulate the same image-space pattern rather than re-parameterizing it. The glossy floor additionally demonstrates a smooth, transport-driven transition in tone blur in reflection, preserving coherent image-space structure while exhibiting the expected glossy spread.

very thin features, and the resulting lines exhibit a controllable softening. Future work could explore importance sampling or guiding strategies for finite-difference pairs (sampling where gradients are likely to be large) to reduce variance and improve robustness.

Finally, our auxiliary-path construction is principled when distribution effects admit a *conditionally specular* interpretation (e.g., microfacet BSDFs and thin-lens). For broader scattering models—such as purely diffuse reflection, layered materials, subsurface scattering, or participating media—the notion of “fixing the realized event” is less direct. Developing analogous event-preserving correspondences for these cases (for instance, by designing connections in an appropriate scattering-parameter space, or deriving transport rules for more general distributions) would greatly broaden applicability.

Tone rendering. Our tone formulation relies on a material-invariant anchor mapping and on evaluating an approximation of its inverse at arbitrary path vertices using local fitting. This is necessarily approximate: the mirror reflection-based anchor mapping is partial and non-bijective, and the local inverse can become ill-conditioned near occlusion boundaries, in sparsely sampled regions, or when extrapolating far from any valid mirror pre-image. While our Moving Least Squares construction and fallback hierarchy are robust in practice, they do not guarantee global injectivity or strict conformality; in extreme cases this can produce fold-overs or distortions.

In addition, our anchor mapping is motivated by geometric optics and stationary-phase behavior. It does not model wave phenomena such as diffraction or interference, which may become relevant for extremely fine or near-Nyquist tone patterns. Follow-up work could investigate alternative anchor mapping constructions (that incorporate additional transport information) or adaptive sample selection for local inversion (including visibility-aware or depth-aware neighborhoods) to improve robustness.

Finally, while we focus on feature lines and tone as representative image-space styles, extending the same correspondence-driven

design principles to additional image-based stylizations (e.g., blue-noise stippling, brush-stroke fields, simulated watercolor painting, etc.) remains a promising direction.

8 CONCLUSION

We introduced path-space formulations for feature line and tone rendering that preserve coherent *image-space* structure under physically-based light transport. For feature lines, we expressed line detection as a stochastic extremal-gradient process and showed how conditional specularity together with Levi-Civita parallel transport yields a principled notion of image-space parallelism and a well-defined gradient, preserving line placement and width under distribution effects. For tone, we anchored image-space tone fields to a material-invariant mapping motivated by stationary-phase arguments, and evaluated them at arbitrary path vertices via a robust local inverse approximation based on Moving Least Squares. Together, these methods integrate naturally into the stylized rendering equation as unbiased Monte Carlo style functions, enabling line and tone stylization to interact meaningfully with global illumination while maintaining their defining image-space structure.

ACKNOWLEDGMENTS

We express our gratitude to the reviewers and the following Sketch-Fab and Blend Swap users: *dark_igorek* for the iron howl (truck) in Fig. 1. *Repka788rus* for the robot rk11 in Fig. 1. *komov* for the walking robot in Fig. 1. *YadroGames* for the house in Fig. 1. *TuppsM* for the hover bike in Fig. 1. *rodrigascar* for the sitting robot in Fig. 1. *dark_igorek* for the kansai spirit (van) in Fig. 1. *Aliksandr.melas* for the badlands landscape in Fig. 1. *kagamine.anjellalice99* for the five robots in Figs. 1, 11 and 14. *Stanford Computer Graphics Laboratory* for the Stanford bunny in Fig. 2. *ruwo* for the camboto in Figs. 2 and 12. *wundersound* for the battle sprinter in Figs. 7 and 13. *Black Snow* for the Nissan Skyline GTR in Fig. 10.

This work has been partially funded by JST FOREST Program (JPMJFR206R, Japan), the Center of Innovations for Sustainable Quantum AI (JST Grant Number JPMJPF2221, Japan), and JSPS KAKENHI Grant Number 24K22830 (Japan).

REFERENCES

- Michael Ashikmin, Simon Premoze, and Peter Shirley. 2000. A Microfacet-Based BRDF Generator. In *Proc. SIGGRAPH '00*. 65–74. <https://doi.org/10.1145/344779.344814>
- Petr Beckmann and André Spizzichino. 1963. *The Scattering of Electromagnetic Waves from Rough Surfaces*. Pergamon Press, NY.
- Max Born and Emil Wolf. 2002. *Principles of Optics: Electromagnetic Theory of Propagation, Interference and Diffraction of Light* (7th ed.). Cambridge University Press.
- Pierre Bénard and Aaron Hertzmann. 2019. Line Drawings from 3D Models: A Tutorial. *Foundations and Trends® in Computer Graphics and Vision* 11, 1-2 (2019), 1–159. <https://doi.org/10.1561/06000000075>
- A. N. M. Imroz Choudhury and Steven G. Parker. 2009. Ray Tracing NPR-Style Feature Lines. In *Proc. NPAR '09* (New Orleans, Louisiana). 5–14. <https://doi.org/10.1145/1572614.1572616>
- Forrester Cole, Aleksey Golovinskiy, Alex Limpaecher, Heather Stoddart Barros, Adam Finkelstein, Thomas Funkhouser, and Szymon Rusinkiewicz. 2008. Where Do People Draw Lines? *ACM Trans. Graph.* 27, 3 (Proc. SIGGRAPH 2008), Article 88 (Aug. 2008), 11 pages. <https://doi.org/10.1145/1360612.1360687>
- Robert L. Cook and Kenneth E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24. <https://doi.org/10.1145/357290.357293>
- Keenan Crane, Mathieu Desbrun, and Peter Schröder. 2010. Trivial Connections on Discrete Surfaces. *Computer Graphics Forum* 29, 5 (2010), 1525–1533. <https://doi.org/10.1111/j.1467-8659.2010.01761.x>
- Doug DeCarlo. 2012. Depicting 3D Shape Using Lines. In *Proc. SPIE 8291, Human Vision and Electronic Imaging XVII*. Article 829116, 16 pages. <https://doi.org/10.1117/12.916463>
- Doug DeCarlo, Adam Finkelstein, and Szymon Rusinkiewicz. 2004. Interactive Rendering of Suggestive Contours with Temporal Coherence. In *Proc. NPAR '04* (Annecy, France). 15–24, 145. <https://doi.org/10.1145/987657.987661>
- Doug DeCarlo, Adam Finkelstein, Szymon Rusinkiewicz, and Anthony Santella. 2003. Suggestive Contours for Conveying Shape. *ACM Trans. Graph.* 22, 3 (Proc. SIGGRAPH 2003) (July 2003), 848–855. <https://doi.org/10.1145/882262.882354>
- Xi Deng, Miloš Hašan, Nathan Carr, Zexiang Xu, and Steve Marschner. 2021. Path Graphs: Iterative Path Space Filtering. *ACM Trans. Graph.* 40, 6 (Proc. SIGGRAPH Asia 2021), Article 276 (Dec. 2021), 15 pages. <https://doi.org/10.1145/3478513.3480547>
- Oliver Deussen, Jörg Hamel, Andreas Raab, Stefan Schlechtweg, and Thomas Strothotte. 1999. An Illustration Technique Using Hardware-Based Intersections and Skeletons. In *Proc. Graphics Interface '99* (Kingston, Ontario, Canada). 175–182. <https://doi.org/10.20380/GI1999.23>
- Oliver Deussen and Tobias Isenber. 2013. Halftoning and Stippling. In *Image and Video-Based Artistic Stylization*, Paul Rosin and John Collomosse (Eds.). Springer London, London, 45–61. https://doi.org/10.1007/978-1-4471-4519-6_3
- Kohji Doi, Yuki Morimoto, and Reiji Tsuruno. 2021. Global Illumination-Aware Stylized Shading. *Computer Graphics Forum* 40, 7 (Proc. Pacific Graphics 2021) (Oct. 2021), 11–20. <https://doi.org/10.1111/cgf.14397>
- Iliyan Georgiev, Zackary Misso, Toshiya Hachisuka, Derek Nowrouzezahrai, Jaroslav Krivánek, and Wojciech Jarosz. 2019. Integral Formulations of Volumetric Transmittance. *ACM Trans. Graph.* 38, 6 (Proc. SIGGRAPH Asia 2019), Article 154 (Nov. 2019), 17 pages. <https://doi.org/10.1145/3355089.3356559>
- Peter Hall. 1999. Nonphotorealistic Rendering by Q-mapping. *Computer Graphics Forum* 18, 1 (March 1999), 27–39. <https://doi.org/10.1111/1467-8659.00300>
- Aaron Hertzmann. 2020. Why Do Line Drawings Work? A Realism Hypothesis. *Perception* 49, 4 (2020), 439–451. <https://doi.org/10.1177/0301006620908207>
- Aaron Hertzmann. 2021. The Role of Edges in Line Drawing Perception. *Perception* 50, 3 (2021), 266–275. <https://doi.org/10.1177/0301006621994407>
- Aaron Hertzmann and Denis Zorin. 2000. Illustrating Smooth Surfaces. In *Proc. SIGGRAPH '00*. 517–526. <https://doi.org/10.1145/344779.345074>
- Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. *ACM Trans. Graph.* 31, 4 (Proc. SIGGRAPH 2012), Article 58 (July 2012), 13 pages. <https://doi.org/10.1145/2185520.2185554>
- Grand Roman Joldes, Habibullah Amin Chowdhury, Adam Wittek, Barry Doyle, and Karol Miller. 2015. Modified Moving Least Squares with Polynomial Bases for Scattered Data Approximation. *Appl. Math. Comput.* 266 (2015), 893–902. <https://doi.org/10.1016/j.amc.2015.05.150>
- Tilke Judd, Frédéric Durand, and Edward H. Adelson. 2007. Apparent Ridges for Line Drawing. *ACM Trans. Graph.* 26, 3 (Proc. SIGGRAPH 2007), Article 19 (2007), 7 pages. <https://doi.org/10.1145/1276377.1276401>
- James T. Kajiya. 1986. The Rendering Equation. *Comput. Graph.* 20, 4 (Proc. SIGGRAPH '86) (Aug. 1986), 143–150. <https://doi.org/10.1145/15886.15902>
- Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A Simple and Robust Mutation Strategy for the Metropolis Light Transport Algorithm. *Computer Graphics Forum* 21, 3 (Proc. EUROGRAPHICS 2002) (Sept. 2002), 531–540. <https://doi.org/10.1111/1467-8659.t01-1-00703>
- Markus Kettunen, Eugene D'Eon, Jacopo Pantaleoni, and Jan Novák. 2021. An Unbiased Ray-Marching Transmittance Estimator. *ACM Trans. Graph.* 40, 4 (Proc. SIGGRAPH 2021), Article 137 (July 2021), 20 pages. <https://doi.org/10.1145/3450626.3459937>
- Yongjin Kim, Jingyi Yu, Xuan Yu, and Seungyong Lee. 2008. Line-Art Illustration of Dynamic and Specular Surfaces. *ACM Trans. Graph.* 27, 5 (Proc. SIGGRAPH Asia 2008), Article 156 (Dec. 2008), 10 pages. <https://doi.org/10.1145/1409060.1409109>
- Felix Knöppel, Keenan Crane, Ulrich Pinkall, and Peter Schröder. 2013. Globally Optimal Direction Fields. *ACM Trans. Graph.* 32, 4 (Proc. SIGGRAPH 2013), Article 59 (July 2013), 10 pages. <https://doi.org/10.1145/2461912.2462005>
- Yunjin Lee, Lee Markosian, Seungyong Lee, and John F. Hughes. 2007. Line Drawings via Abstracted Shading. *ACM Trans. Graph.* 26, 3 (Proc. SIGGRAPH 2007), Article 18 (July 2007), 5 pages. <https://doi.org/10.1145/1275808.1276400>
- David Levin. 1998. The Approximation Power of Moving Least-Squares. *Math. Comput.* 67, 224 (Oct. 1998), 1517–1531. <https://doi.org/10.1090/S0025-5718-98-00974-0>
- Jonatas Medeiros, Mario Sousa, Luiz Velho, and Carla Freitas. 2009. Perspective Contouring in Illustrative Visualization. In *Proc. SIBGRAPI 2009*. 48–55. <https://doi.org/10.1109/SIBGRAPI.2009.49>
- Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and Consistent Rendering Using Biased Estimators. *ACM Trans. Graph.* 41, 4 (Proc. SIGGRAPH 2022), Article 48 (July 2022), 13 pages. <https://doi.org/10.1145/3528223.3530160>
- Thomas Müller, Fabrice Rousselle, Jan Novák, and Alexander Keller. 2021. Real-Time Neural Radiance Caching for Path Tracing. *ACM Trans. Graph.* 40, 4 (Proc. SIGGRAPH 2021), Article 36 (July 2021), 16 pages. <https://doi.org/10.1145/3450626.3459812>
- Shinji Ogaki and Iliyan Georgiev. 2018. Production Ray Tracing of Feature Lines. In *SIGGRAPH Asia 2018 Technical Briefs* (Tokyo, Japan) (SA '18). Article 15, 4 pages. <https://doi.org/10.1145/3283254.3283273>
- Yutaka Ohtake, Alexander Belyaev, and Hans-Peter Seidel. 2004. Ridge-Valley Lines on Meshes via Implicit Surface Fitting. *ACM Trans. Graph.* 23, 3 (Proc. SIGGRAPH 2004) (Aug. 2004), 609–612. <https://doi.org/10.1145/1015706.1015768>
- Michael Oren and Shree K. Nayar. 1994. Generalization of Lambert's Reflectance Model. In *Proc. SIGGRAPH '94*. 239–246. <https://doi.org/10.1145/192161.192213>
- Ulrich Pinkall and Konrad Polthier. 1993. Computing Discrete Minimal Surfaces and Their Conjugates. 2, 1 (1993), 15–36. <https://doi.org/10.1080/10586458.1993.10504266>
- Konrad Polthier and Markus Schmies. 1998. Straightest Geodesics on Polyhedral Surfaces. In *Mathematical Visualization: Algorithms, Applications and Numerics*, Hans-Christian Hege and Konrad Polthier (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 135–150. https://doi.org/10.1007/978-3-662-03567-2_11
- Emil Praun, Hugues Hoppe, Matthew Webb, and Adam Finkelstein. 2001. Real-Time Hatching. In *Proc. SIGGRAPH '01*. 581–586. <https://doi.org/10.1145/383259.383328>
- Takafumi Saito and Tokiichiro Takahashi. 1990. Comprehensive Rendering of 3-D Shapes. *Comput. Graph.* 24, 4 (Proc. SIGGRAPH '90) (Sept. 1990), 197–206. <https://doi.org/10.1145/97879.97901>
- Nicholas Sharp, Yousuf Soliman, and Keenan Crane. 2019. The Vector Heat Method. *ACM Trans. Graph.* 38, 3, Article 24 (June 2019), 19 pages. <https://doi.org/10.1145/3243651>
- Xiaochun Tong and Toshiya Hachisuka. 2025. Practical Stylized Nonlinear Monte Carlo Rendering. In *SIGGRAPH 2025 Conference Papers*. Article 84, 11 pages. <https://doi.org/10.1145/3721238.3730686>
- Kenneth E. Torrance and Ephraim M. Sparrow. 1967. Theory for Off-Specular Reflection from Roughened Surfaces. *J. Opt. Soc. Am.* 57, 9 (Sept. 1967), 1105–1114. <https://doi.org/10.1364/JOSA.57.001105>
- T. S. Trowbridge and K. P. Reitz. 1975. Average Irregularity Representation of a Rough Surface for Ray Reflection. *J. Opt. Soc. Am.* 65, 5 (May 1975), 531–536. <https://doi.org/10.1364/JOSA.65.000531>
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction through Rough Surfaces. In *Rendering Techniques (Proc. EGSR '07)* (Grenoble, France). 195–206. <https://doi.org/10.2312/EGWR/EGSR07/195-206>
- Rex West. 2021. Physically-based Feature Line Rendering. *ACM Trans. Graph.* 40, 6 (Proc. SIGGRAPH Asia 2021), Article 246 (Dec. 2021), 11 pages. <https://doi.org/10.1145/3478513.3480550>
- Rex West, Iliyan Georgiev, and Toshiya Hachisuka. 2022. Marginal Multiple Importance Sampling. In *SIGGRAPH Asia 2022 Conference Papers*. Article 42, 8 pages. <https://doi.org/10.1145/3550469.3555388>
- Rex West and Sayan Mukherjee. 2024. Stylized Rendering as a Function of Expectation. *ACM Trans. Graph.* 43, 4 (Proc. SIGGRAPH 2024), Article 96 (July 2024), 19 pages. <https://doi.org/10.1145/3658161>