Modeling Human Vision with Differential Geometry

Ana Dodik ANADODIK@MIT.EDU
MIT CSAIL, Cambridge, MA, USA

Editors: List of editors' names

Abstract

We describe recent efforts to tackle the problem of computationally representing impossible objects, i.e., shapes which have local geometry but cannot be globally assembled into 3D, in a manner reflective of how humans perceive them. We build off of the initial work describing a discrete representation of these objects (Dodik et al., 2025) toward a broader smooth mathematical theory independent of the parametric function class used to represent impossible objects on the computer, potentially opening doors to encoding them via, e.g., a neural network. We will also discuss implications of our work for human and machine vision research, including concrete testable hypotheses as well as some more speculative ideas.

1. Introduction

It has long been hypothesized, for example by Marr (1982), that that our visual experience proceeds as an assembly of low-level local visual cues such as orientation and connectivity into higher-level global representations involving depth and global orientation. Impossible objects, such as the Penrose triangle (1958) or many of the works of M.C. Escher (1961), thus pose a kind of edge case for the human visual system; their local geometry, as perceived by us, cannot be consistently assembled into a global construct, making them of interest to both artists and vision scientists alike. Thus, a computational model capable of representing and manipulating impossible objects would no doubt be of use to the study of the human visual system.

Thus, previous work set out to create a computational model that reflects this computational principle, one that can represent locally, but not globally, consistent geometry, that nonetheless allows us to answer geometric queries regarding global consistency and distances between points. Dodik et al. (2025) introduce the *mescher*, a mesh-based representation inspired by vision science, which serves as a computational analogue of Marr's "2.5D sketch" (1982) by encoding local orientation and connectivity without committing to a 3D embedding. This extended abstract largely consists of a summary of the model by Dodik et al. (2025), only framed in the context of vision science and using a smooth parameterization-agnostic mathematical formulation. This smooth formulation points towards other parameterizations of impossible objects, including using, e.g., constrained neural fields.

The prior work central to this extended abstract (Dodik et al., 2025) points out that differential geometry—as made computational via discrete exterior calculus (DEC) (Hirani, 2003; Desbrun et al., 2006; Crane et al., 2013)—represents a useful formalism for our purpose. Specifically, meschers are described in terms of discrete 0- and 1-forms and state the visual consistency requirements in terms of local and global integrability. Meschers can be used to compute the discrete metric of an impossible shape, unlocking practical applications that involve solving partial differential equations on the surface of impossible objects, answering distance queries, or even the inverse rendering a possible object into impossible one. Lastly, not demonstrated before is that this formalism allows us to find the closest projection in an L_2 sense of an impossible object onto the space of possible ones and thus extract out the "impossibility"—we include additional results of this projection in Figure 1.

We summarize related work on perception of impossible objects (§2) before describing meschers through a purely smooth lens which we then tie back to the original discretization (2025) (§3), concluding with a discussion of possible future directions in human and machine vision (§4).

2. Related Work

This perceptual process is fundamentally driven by the spatial locality of vision. Foveation makes it so that we perceive the world through small, localized fixations which need to be assembled into a global percept (Yarbus, 1967; Irwin, 1991). According to Marr (1982), local visual cues of shading and occlusions as extracted by the early visual processing system consist a so-called "2.5D sketch". Impossible objects providing valid local depth cues—and therefore a valid 2.5D sketch—but fail to integrate into a consistent global embedding (Freud et al., 2013, 2015; Heinke et al., 2021). Our visual system is robust enough to make partial local sense of such objects and detect the "impossibility" (Linton et al., 2023) once a contradiction is reached (Schuster, 1964). Koenderink's pictorial reliefs (1998) bare relevance as they are a computational representation of locally integrable geometry; however, he concludes that our perceptual representations must also be globally consistent, a fact contradicted by impossible objects. Consequently, a mescher (Dodik et al., 2025) is a representation that includes local geometric properties without committing to global depths.

The long-standing fascination with impossible objects (Penrose and Penrose, 1958; Escher, 1961; Hogarth, 1754; Reutersvärd, 1934) has spurred various attempts to formalize a model of human perception capable of representing them. Early theories proposed that humans "scan" a shape to assemble locally plausible regions into a globally consistent whole (Simon, 1967), or detect impossibilities through conflicting depth hypotheses (Gregory, 1970) or as a violation of spatial connectedness (Draper, 1978). Meschers (Dodik et al., 2025) unify these observations within the mathematical formalism of global integrability (or lack thereof), somewhat similar to Penrose's treatment of impossible shapes in terms of the de Rham cohomology (Penrose, 1993). Unlike Térouanne (1984) who explains impossible objects purely topologically in terms of a 2D tesselation and its connectivity, Dodik et al. (2025) include local geometry into their model. These two works ultimately represent two competing hypotheses for our perceptual representation of these objects.

3. Meschers: Summary

While the original work (Dodik et al., 2025) understands meschers as discrete triangle meshes, we will start off with a more general formulation by defining a mescher to be a smooth oriented manifold $\mathcal{M}$. Its geometry is represented by its screen-space coordinates, which can be stated as the 0-forms (scalar fields) $x, y \in \Omega^0(\mathcal{M})$, and the *relative*, i.e., differential depths, stated as the 1-form (covector field) $\zeta \in \Omega^1(\mathcal{M})$. Intuitively, at a given point $p \in \Omega$, we can ask what change in depth is incurred if we were to move along some tangent vector $v \in T_p(\Omega)$ by consulting the 1-form $\zeta_p(v)$.

In general, not all possible assignments of values to ζ are useful. Specifically, to shade an object, we must have access to its normal vectors, necessitating each sufficiently small local neighborhood of $\mathcal{M}$ be embeddable into 3D even if they cannot all be assembled together consistently. We thus require ζ is locally integrable, meaning that after making a full loop around a small local patch, we end up at the same depth we started at; or, the sum change in depth around a local loop must be zero. In terms of exterior calculus, local integrability can be stated as the condition

$$d\zeta = 0, \tag{1}$$

where $d \in: \Omega^k(\mathcal{M}) \to \Omega^{k+1}(\mathcal{M})$ is the exterior derivative mapping k-forms to $(k+1)$-forms, an exterior calculus analogue of the curl operator. This condition ensures well-defined local geometry without enforcing global integrability, ultimately allowing meschers to represent impossible shapes.

We can characterize the space of ζ via the Hodge decomposition,

$$\zeta = d_{01}z + \omega, \tag{2}$$

where $z \in \Omega^0(\mathcal{M})$ is a primal 0-form and ω is a harmonic 1-form satisfying $d_{12}\omega = 0$ and $\star d \star \omega = 0$, with $\star$ as the Hodge star operator. If $\omega = 0$, then the mescher can be globally embedded in $\mathbb{R}^3$ and is therefore a *possible* object; if however $\omega \neq 0$, the mescher encodes an *impossible* object. Interestingly, simply discarding the harmonic component ω and keeping z as the absolute depth results in the closest L_2 projection of the impossible object onto the space of possible ones; we show this in Figure 1.

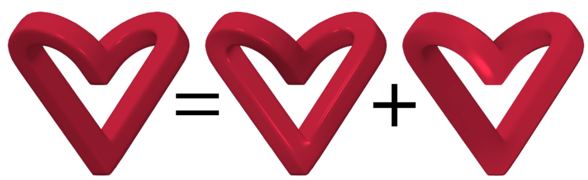

Figure 1: Impossible heart. An impossible object (left) decomposes into the closest possible object (middle) and an impossble, harmonic component (right) via the Hodge decomposition.

The work by Dodik et al. (2025) discretizes this using discrete exterior calculus (DEC) (Hirani, 2003; Desbrun et al., 2006; Crane et al., 2013) by specifying the topology of a manifold triangle mesh with vertex set $\mathbb{V}$, edge set $\mathbb{E}$, and face set $\mathbb{F}$, and orientations assigned to each face. In the discrete setting, we have per-vertex discrete 0-forms $\boldsymbol{x}, \boldsymbol{y} \in \mathbb{R}^{|\mathbb{V}|}$, and per-edge *signed* relative depth, the 1-form $\boldsymbol{\zeta} \in \mathbb{R}^{|\mathbb{E}|}$. Local integrability of $\boldsymbol{\zeta}$ now means that the sum change in depth around each triangle must be zero and can be equivalently stated as the condition $d_{12}\,\boldsymbol{\zeta} = 0$, where $d_{12} \in \{-1, 0, 1\}^{|\mathbb{F}|\times|\mathbb{E}|}$ is the discrete exterior derivative mapping 1-forms to oriented sums around faces.

Once the geometry is fully specified, we have can use the metric to build differential operators including the Laplacian and Hodge Laplacian, as well as the gradient and divergence. This allows us to solve partial differential equations on the surface of impossible objects, opening doors to answering distance queries, as well as editing operations such as surface smoothing. Meschers further allowed to solve an inverse rendering problem involving impossible objects: starting off with a possible torus, one can minimize the distance between its rendering and the image of a Penrose triangle (Figure 2).

4. Discussion

A discrete mesh representation is only one approach to making impossible objects a reality on a computer. We are particularly excited by other possible parameterizations unlocked by our

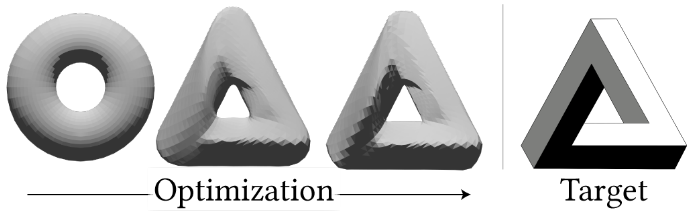

Figure 2: Penrose triangle. Meschers allow us to recover an impossible object by starting from a possible torus (left) and optimizing to match a target image (right). Image from Dodik et al. (2025).

smooth formulation from Section 3, including neural fields and Gaussian splats, both more amenable representations to inverse rendering. Moreover, a representation akin to an occupancy field may also be possible as long as global depth remains unspecified.

From the perspective of vision science, meschers point towards a computational model of perception (Marr, 1982). Indeed, any model of human vision needs to account for the fact that the human visual system can successfully reason about impossible objects, reflecting tolerance towards the global contradiction (Linton et al., 2023). Because meschers parameterize the impossible component via the harmonic part of the depth 1-form, they enable smooth interpolation between possible and impossible geometry, suggesting a new class of stimuli for psychophysical and neuroimaging experiments in which the degree of impossibility is systematically varied to test where and how neural representations diverge from those of possible objects. They also open the door to asking whether geometric reasoning actually enters into perception of impossible objects, or if a purely topological model such as that proposed by Térouanne (1984) suffices.

A further line of investigation concerns the dynamics of integration within our visual system. Applying the Hodge decomposition and discarding the harmonic component is only one approach to assembling local depth cues into a global structure, but the presence of saccades in the visual system certainly points towards alternative incremental strategies more akin to a propagating front and similar to Simon's hypothesis (1967). Investigating how and to which degree our visual system performs global integration via various mescher stimuli certainly seems of interest. Furthermore, meschers can allow us to investigate how complex, be it topologically or geometrically, an impossible object can become before we cease to detect the impossibility, and whether that changes if the impossibility is contained only in a small local region of the shape.

Meschers may be as exciting for computer vision as they are for human vision research. Knowing that human vision can operate only using locally consistent geometry raises the question of how much 3D computer vision can be accomplished using local information alone. It is already known that the task of monocular surface normal prediction is can be done in a more parameter-efficient way than that of monocular depth prediction. This makes sense: the depth network has to assemble local shading information into 3D by performing a conceptual equivalent of a Hodge decomposition within its weights. Essentially, this means that predicting a mescher-like representation is much easier than predicting global depth. Are there computer vision tasks which currently rely on depth prediction which could effectively be accomplished with a mescher-like structure? We foresee certain applications in robotics as being particularly impacted by such differential representations. Specifically, estimating the local geometry and distances of nearby surfaces may come in handy for navigation and contact planning even if we forgo absolute depth.

In summary, geometry representations which exist only in the context of our perception help expand the frontier of knowledge on human perception and offer paths forward for computer vision. Meschers—be it discrete or smooth—are only one such representation; other components of the human visual system are yet to be explored through the formal lens of geometry.

References

Keenan Crane, Fernando de Goes, Mathieu Desbrun, and Peter Schröder. Digital geometry processing with discrete exterior calculus. In *ACM SIGGRAPH 2013 Courses*, SIGGRAPH '13, New York, NY, USA, 2013. Association for Computing Machinery. ISBN 9781450323390. doi: 10.1145/2504435.2504442. URL https://doi.org/10.1145/2504435.2504442.

Mathieu Desbrun, Eva Kanso, and Yiying Tong. Discrete differential forms for computational modeling. In *ACM SIGGRAPH 2006 Courses*, SIGGRAPH '06, page 39–54, New York, NY, USA, 2006. Association for Computing Machinery. ISBN 1595933646. doi: 10.1145/1185657.1185665. URL https://doi.org/10.1145/1185657.1185665.

Ana Dodik, Isabella Yu, Kartik Chandra, Jonathan Ragan-Kelley, Joshua Tenenbaum, Vincent Sitzmann, and Justin Solomon. Meschers: Geometry processing of impossible objects. *ACM Transactions on Graphics*, 2025. doi: https://doi.org/10.1145/3731422.

Stephen W Draper. The penrose triangle and a family of related figures. *Perception*, 7(3):283–296, 1978. doi: 10.1068/p070283. URL https://doi.org/10.1068/p070283. Pmid: 693228.

M.C. Escher. *The Graphic Work of M.C. Escher*. Oldbourne, 1961. URL https://books.google.com/books?id=1T5gwgEACAAJ.

Erez Freud, Tzvi Ganel, and Galia Avidan. Representation of possible and impossible objects in the human visual cortex: Evidence from fmri adaptation. *NeuroImage*, 64:685–692, 2013. ISSN 1053-8119. doi: https://doi.org/10.1016/j.neuroimage.2012.08.070. URL https://www.sciencedirect.com/science/article/pii/S1053811912008798.

Erez Freud, Bat-Sheva Hadad, Galia Avidan, and Tzvi Ganel. Evidence for similar early but not late representation of possible and impossible objects. *Frontiers in Psychology*, 6, 2015. ISSN 1664-1078. doi: 10.3389/fpsyg.2015.00094. URL https://www.frontiersin.org/journals/psychology/articles/10.3389/fpsyg.2015.00094.

R. L. (Richard Langton) Gregory. *The intelligent eye, by R. L. Gregory.* Weidenfeld & Nicolson, London, 1970. ISBN 0297000217.

Dietmar Heinke, Peter Wachman, Wieske van Zoest, and E. Charles Leek. A failure to learn object shape geometry: Implications for convolutional neural networks as plausible models of biological vision. *Vision Research*, 189:81–92, 2021. ISSN 0042-6989. doi: https://doi.org/10.1016/j.visres.2021.09.004. URL https://www.sciencedirect.com/science/article/pii/S0042698921002017.

Anil N. Hirani. *Discrete Exterior Calculus*. PhD thesis, California Institute of Technology, May 2003. URL http://resolver.caltech.edu/CaltechETD:etd-05202003-095403.

William Hogarth. Satire on false perspective, 1754. URL https://publicdomainreview.org/collection/william-hogarth-satire-on-false-perspective/. Frontispiece to Joshua Kirby's "Dr. Brook Taylor's Method of Perspective Made Easy".

David E Irwin. Information integration across saccadic eye movements. *Cognitive Psychology*, 23(3): 420–456, 1991. ISSN 0010-0285. doi: https://doi.org/10.1016/0010-0285(91)90015-G. URL https://www.sciencedirect.com/science/article/pii/001002859190015G.

Jan J Koenderink. Pictorial relief. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 356(1740):1071–1086, 1998.

Paul Linton, Michael J Morgan, Jenny CA Read, Dhanraj Vishwanath, Sarah H Creem-Regehr, and Fulvio Domini. New approaches to 3d vision, 2023.

David Marr. *Vision: A computational investigation into the human representation and processing of visual information*. MIT press, 1982.

Lionel S Penrose and Roger Penrose. Impossible objects: a special type of visual illusion. *British Journal of Psychology*, 1958.

Roger Penrose. On the cohomology of impossible figures. *Leonardo*, 25:245–247, 1993. URL https://api.semanticscholar.org/CorpusID:125905129.

Oscar Reutersvärd. First impossible triangle, 1934. URL https://sis.modernamuseet.se/objects/28602/. Earliest known impossible figure, later reproduced as lithographs.

Donald H Schuster. A new ambiguous figure: A threestick clevis. *The American Journal of Psychology*, 1964.

Herbert A. Simon. An information-processing explanation of some perceptual phenomena. *British Journal of Psychology*, 58(1-2):1–12, 1967. doi: https://doi.org/10.1111/j.2044-8295.1967.tb01051.x. URL https://bpspsychub.onlinelibrary.wiley.com/doi/abs/10.1111/j.2044-8295.1967.tb01051.x.

E. Térouanne. Impossible objects and inconsistent interpretations. In E. Degreef and J. Van Buggenhaut, editors, *Trends in Mathematical Psychology*, volume 20 of *Advances in Psychology*, pages 105–120. North-Holland, 1984. doi: https://doi.org/10.1016/S0166-4115(08)62082-8. URL https://www.sciencedirect.com/science/article/pii/S0166411508620828.

A. L. Yarbus. *Eye Movements and Vision*. Plenum. New York., 1967.