Progressive null-tracking for volumetric rendering

Zackary Misso[1] Yining Karl Li[2] Brent Burley[2] Daniel Teece[2] Wojciech Jarosz[1]

[1]**Dartmouth College,** [2]**Walt Disney Animation Studios**

Hyperion

Strange World

From the SIGGRAPH Asia 2014 conference proceedings.

Residual Ratio Tracking for Estimating Attenuation in Participating Media

Jan Novák[1,2] Andrew Selle[1] Wojciech Jarosz[2]

[1]Walt Disney Animation Studios [2]Disney Research Zürich

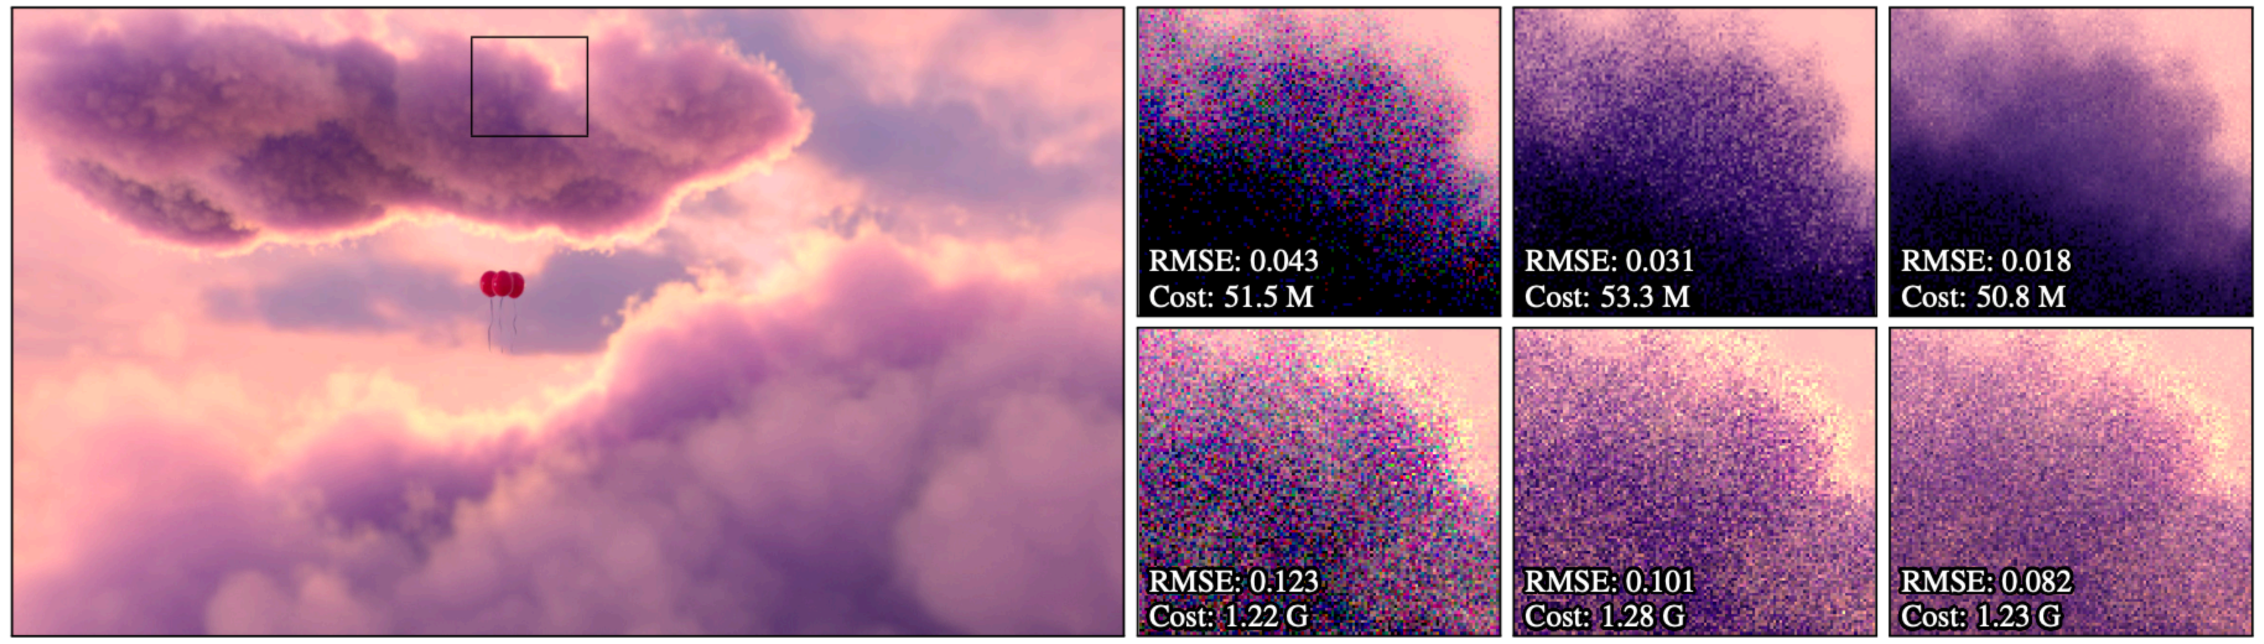

(a) *Path traced light transport in clouds* **(b)** *Delta tracking* **(c)** *Ratio tracking* **(d)** *Residual ratio tracking*

Figure 1: *A cloudy sky rendered with our residual ratio tracking estimator for computing transmittance in heterogeneous volumes. Our technique is unbiased, outperforms the delta tracking-based estimator* (**b**)*, and fits well into path-tracing, production frameworks. The insets show renderings of absorptive-only (top) and scattering (bottom) clouds; the transmittance was estimated using delta tracking* (**b**)*, ratio tracking* (**c**)*, and residual ratio tracking* (**d**) *with a roughly equal cost reported as the number of extinction coefficient evaluations. Images ©Disney.*

Abstract

Evaluating transmittance within participating media is a fundamental operation required by many light transport algorithms. We present *ratio tracking* and *residual tracking*, two complementary techniques that can be combined into an efficient, unbiased estimator for evaluating transmittance in complex heterogeneous media. In comparison to current approaches, our new estimator is unbiased, yields high efficiency, gracefully handles media with wavelength dependent extinction, and bridges the gap between closed form solutions and purely numerical, unbiased approaches. A key feature of ratio tracking is its ability to handle negative densities. This in turn enables us to separate the main part of the transmittance function, handle it analytically, and numerically estimate only the residual transmittance. In addition to proving the unbiasedness of our estimators, we perform an extensive empirical analysis to reveal parameters that lead to high efficiency. Finally, we describe how to integrate the new techniques into a production path tracer and demonstrate their

1 Introduction

The world around us is filled with participating media that attenuates and scatters light as it travels from light sources, to surfaces, and finally to our eyes. Simulating this transport in heterogeneous participating media—such as smoke, clouds, nuclear reactor housings, biological tissue, or other volumetric datasets—is important in many fields, ranging from neutron transport, to medical physics, scientific visualization, and film and visual effects production.

Monte Carlo (MC) path sampling approaches, including variants of path tracing [Kajiya 1986], bidirectional path tracing [Lafortune and Willems 1993; Veach and Guibas 1994; Pauly et al. 2000], or many-light methods [Keller 1997; Dachsbacher et al. 2013], have proven to be practical approaches for accurately approximating this light transport. All of these rely on generating random paths between the light(s) and the sensor, and there has been extensive research on importance sampling such paths to obtain low-noise images [Raab

Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes

PETER KUTZ, Walt Disney Animation Studios
RALF HABEL, Walt Disney Animation Studios
YINING KARL LI, Walt Disney Animation Studios
JAN NOVÁK, Disney Research

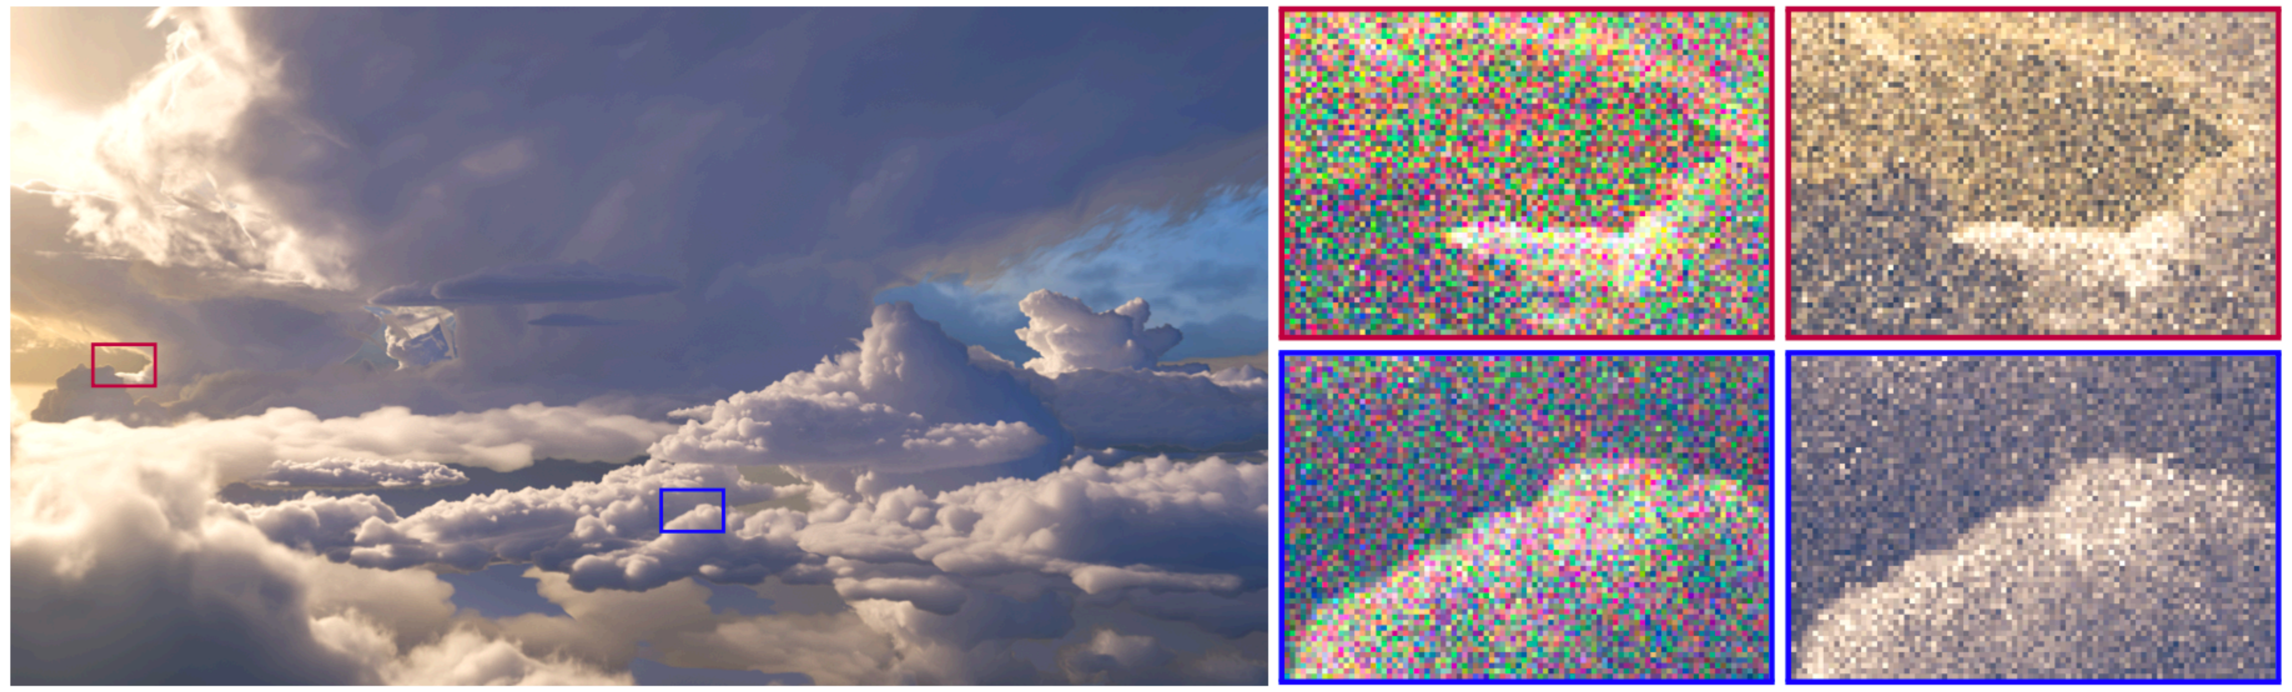

Figure 1. A cloudscape rendered with a combination of our spectral and decomposition tracking techniques, which gracefully handle chromatic media and reduce collision coefficient evaluations. The insets on the right were computed in equal time, with our method yielding 3.5× lower MSE than delta tracking.

We present two novel unbiased techniques for sampling free paths in heterogeneous participating media. Our *decomposition tracking* accelerates free-path construction by splitting the medium into a control component and a residual component and sampling each of them separately. To minimize expensive evaluations of spatially varying collision coefficients, we define the control component to allow constructing free paths in closed form. The residual heterogeneous component is then homogenized by adding a fictitious medium and handled using weighted delta tracking, which removes the need for computing strict bounds of the extinction function. Our second contribution, *spectral tracking*, enables efficient light transport simulation in chromatic media. We modify free-path distributions to minimize the fluctuation of path throughputs and thereby reduce the estimation variance. To demonstrate the correctness of our algorithms, we derive them *directly* from the radiative transfer equation by extending the integral formulation of null-collision algorithms recently developed in reactor physics. This mathematical framework, which we thoroughly review, encompasses existing trackers and postulates an entire family of new estimators for solving trans-

CCS Concepts: • **Computing methodologies → Rendering**; *Ray tracing*;

Additional Key Words and Phrases: participating media, volume rendering, free-path sampling, transmittance, delta tracking, ratio tracking, color

ACM Reference format:
Peter Kutz, Ralf Habel, Yining Karl Li, and Jan Novák. 2017. Spectral and Decomposition Tracking for Rendering Heterogeneous Volumes. *ACM Trans. Graph.* 36, 4, Article 111 (July 2017), 16 pages.
DOI: http://dx.doi.org/10.1145/3072959.3073665

1 INTRODUCTION

Accurate and efficient simulation of radiative transfer in participating media is essential in many domains, such as nuclear reactor design, medical imaging, scientific visualization, and realistic image

Homogenization

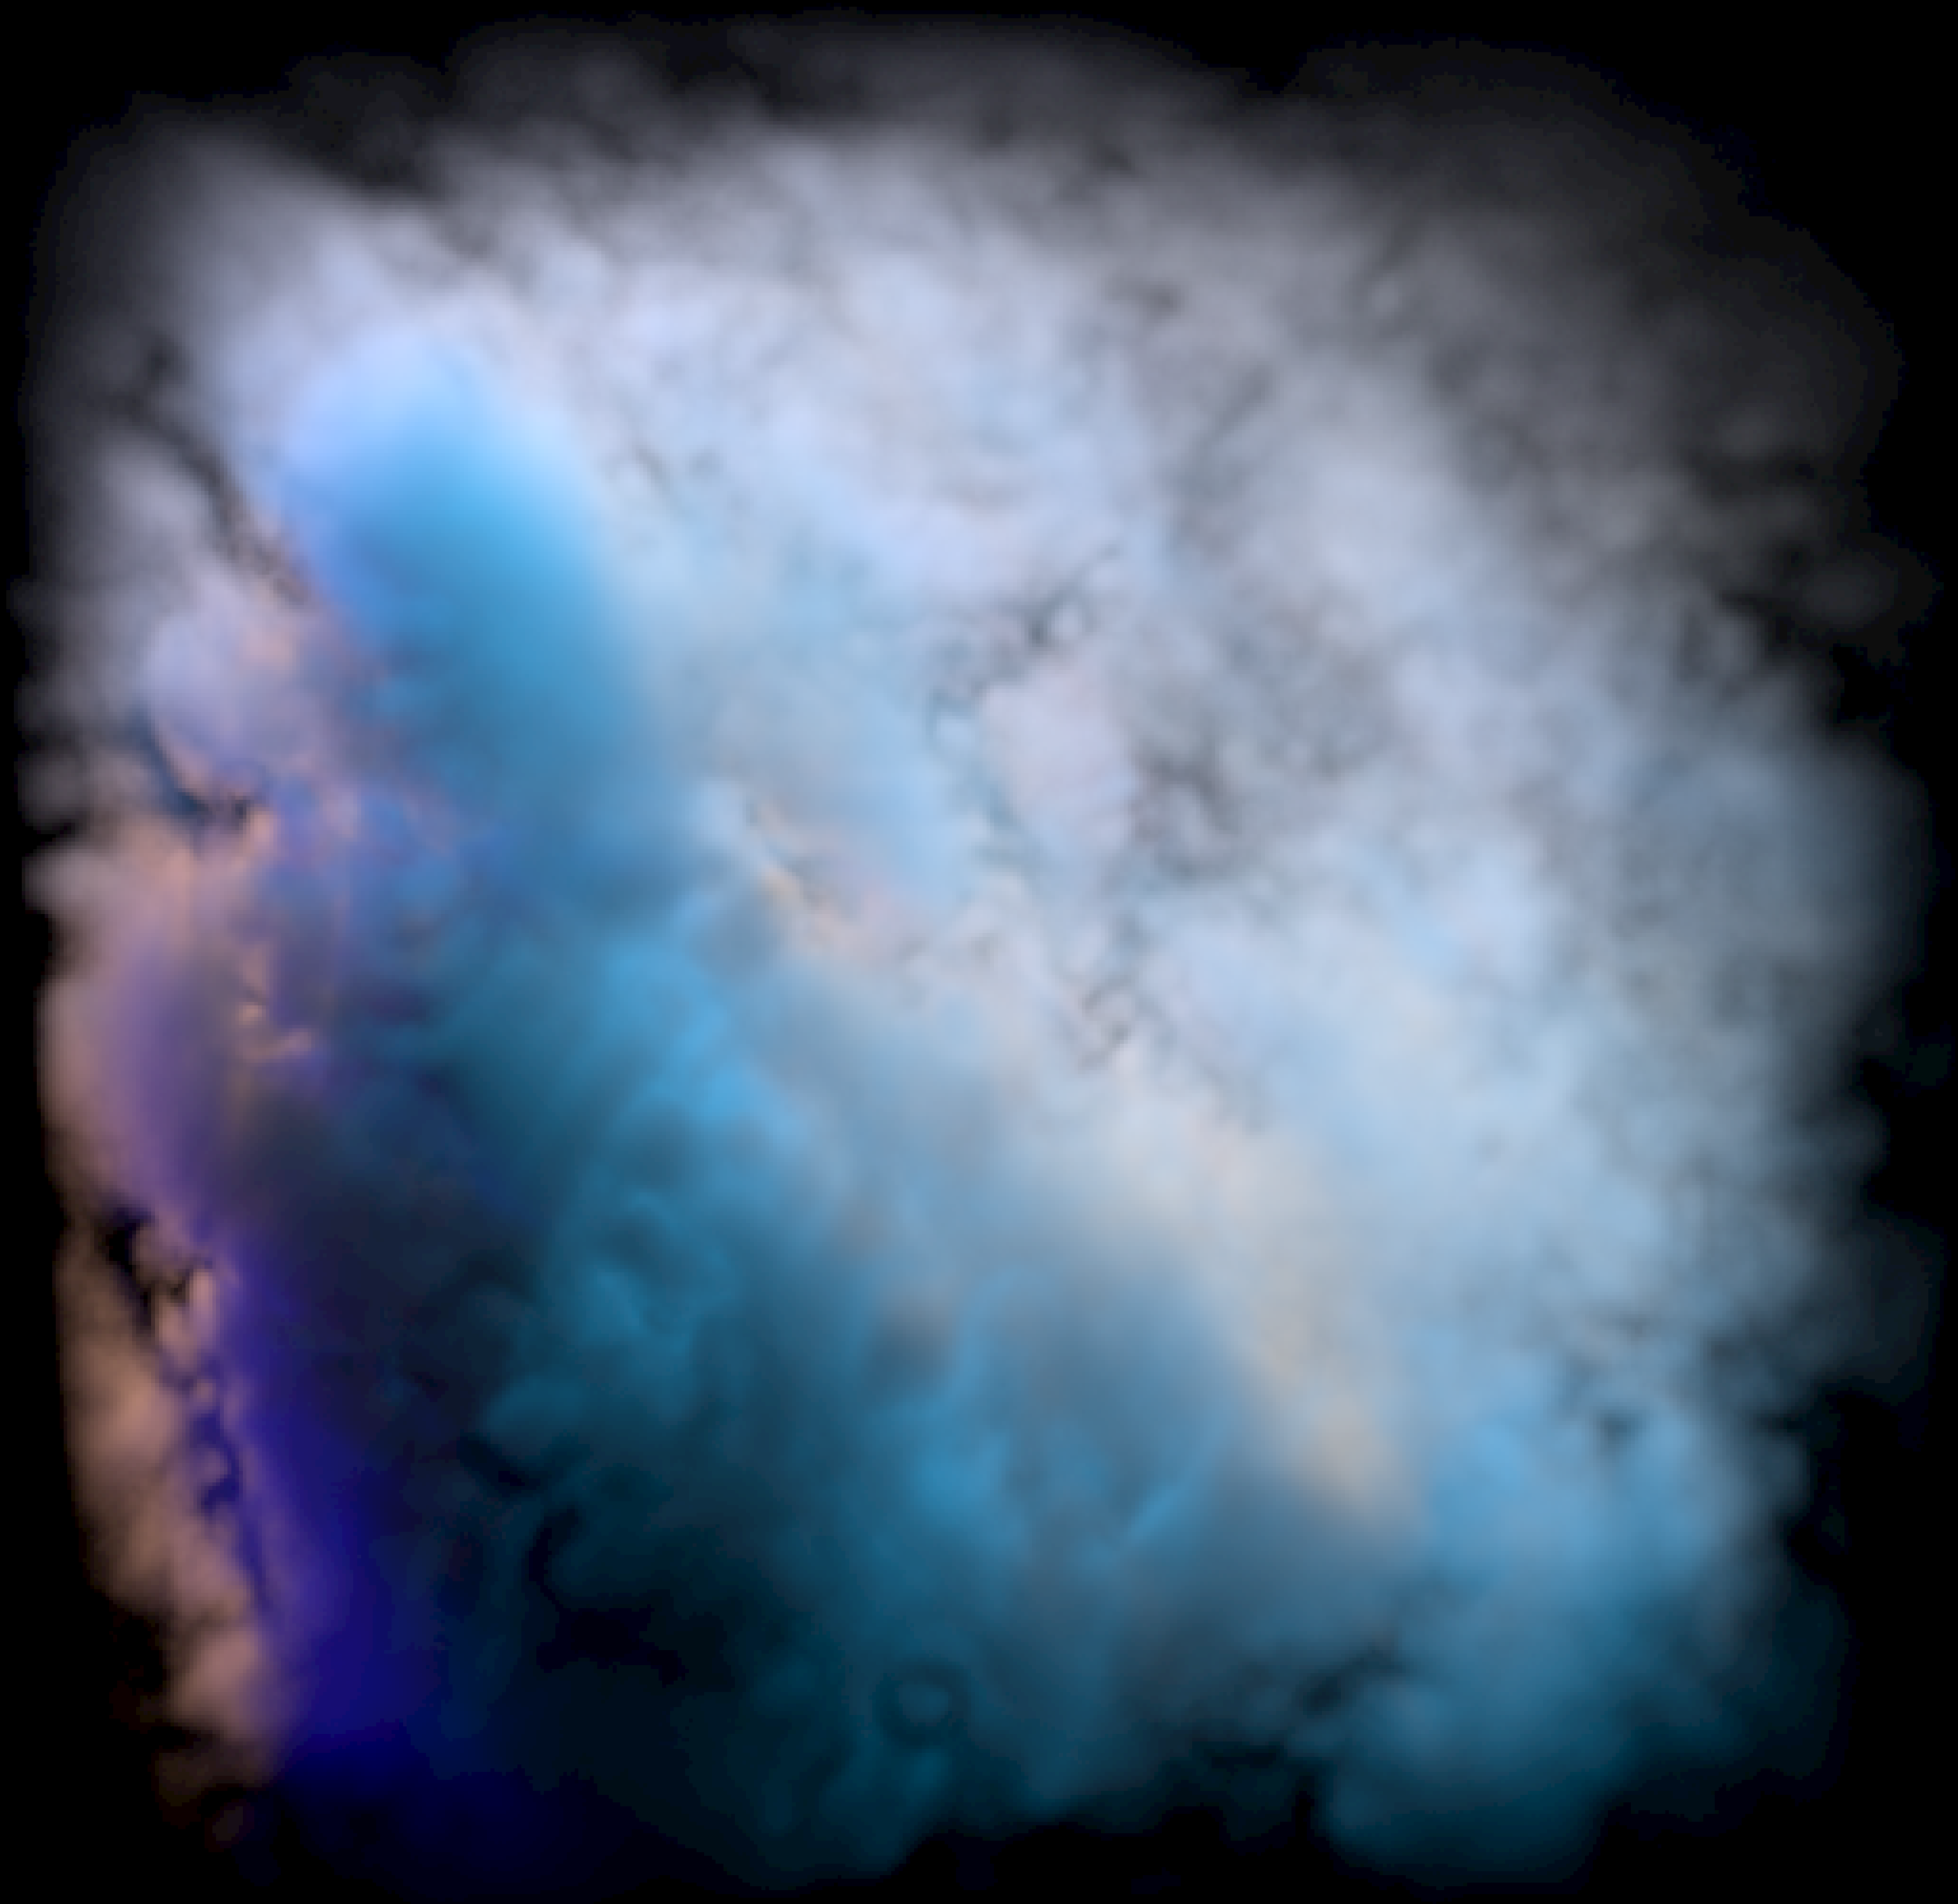

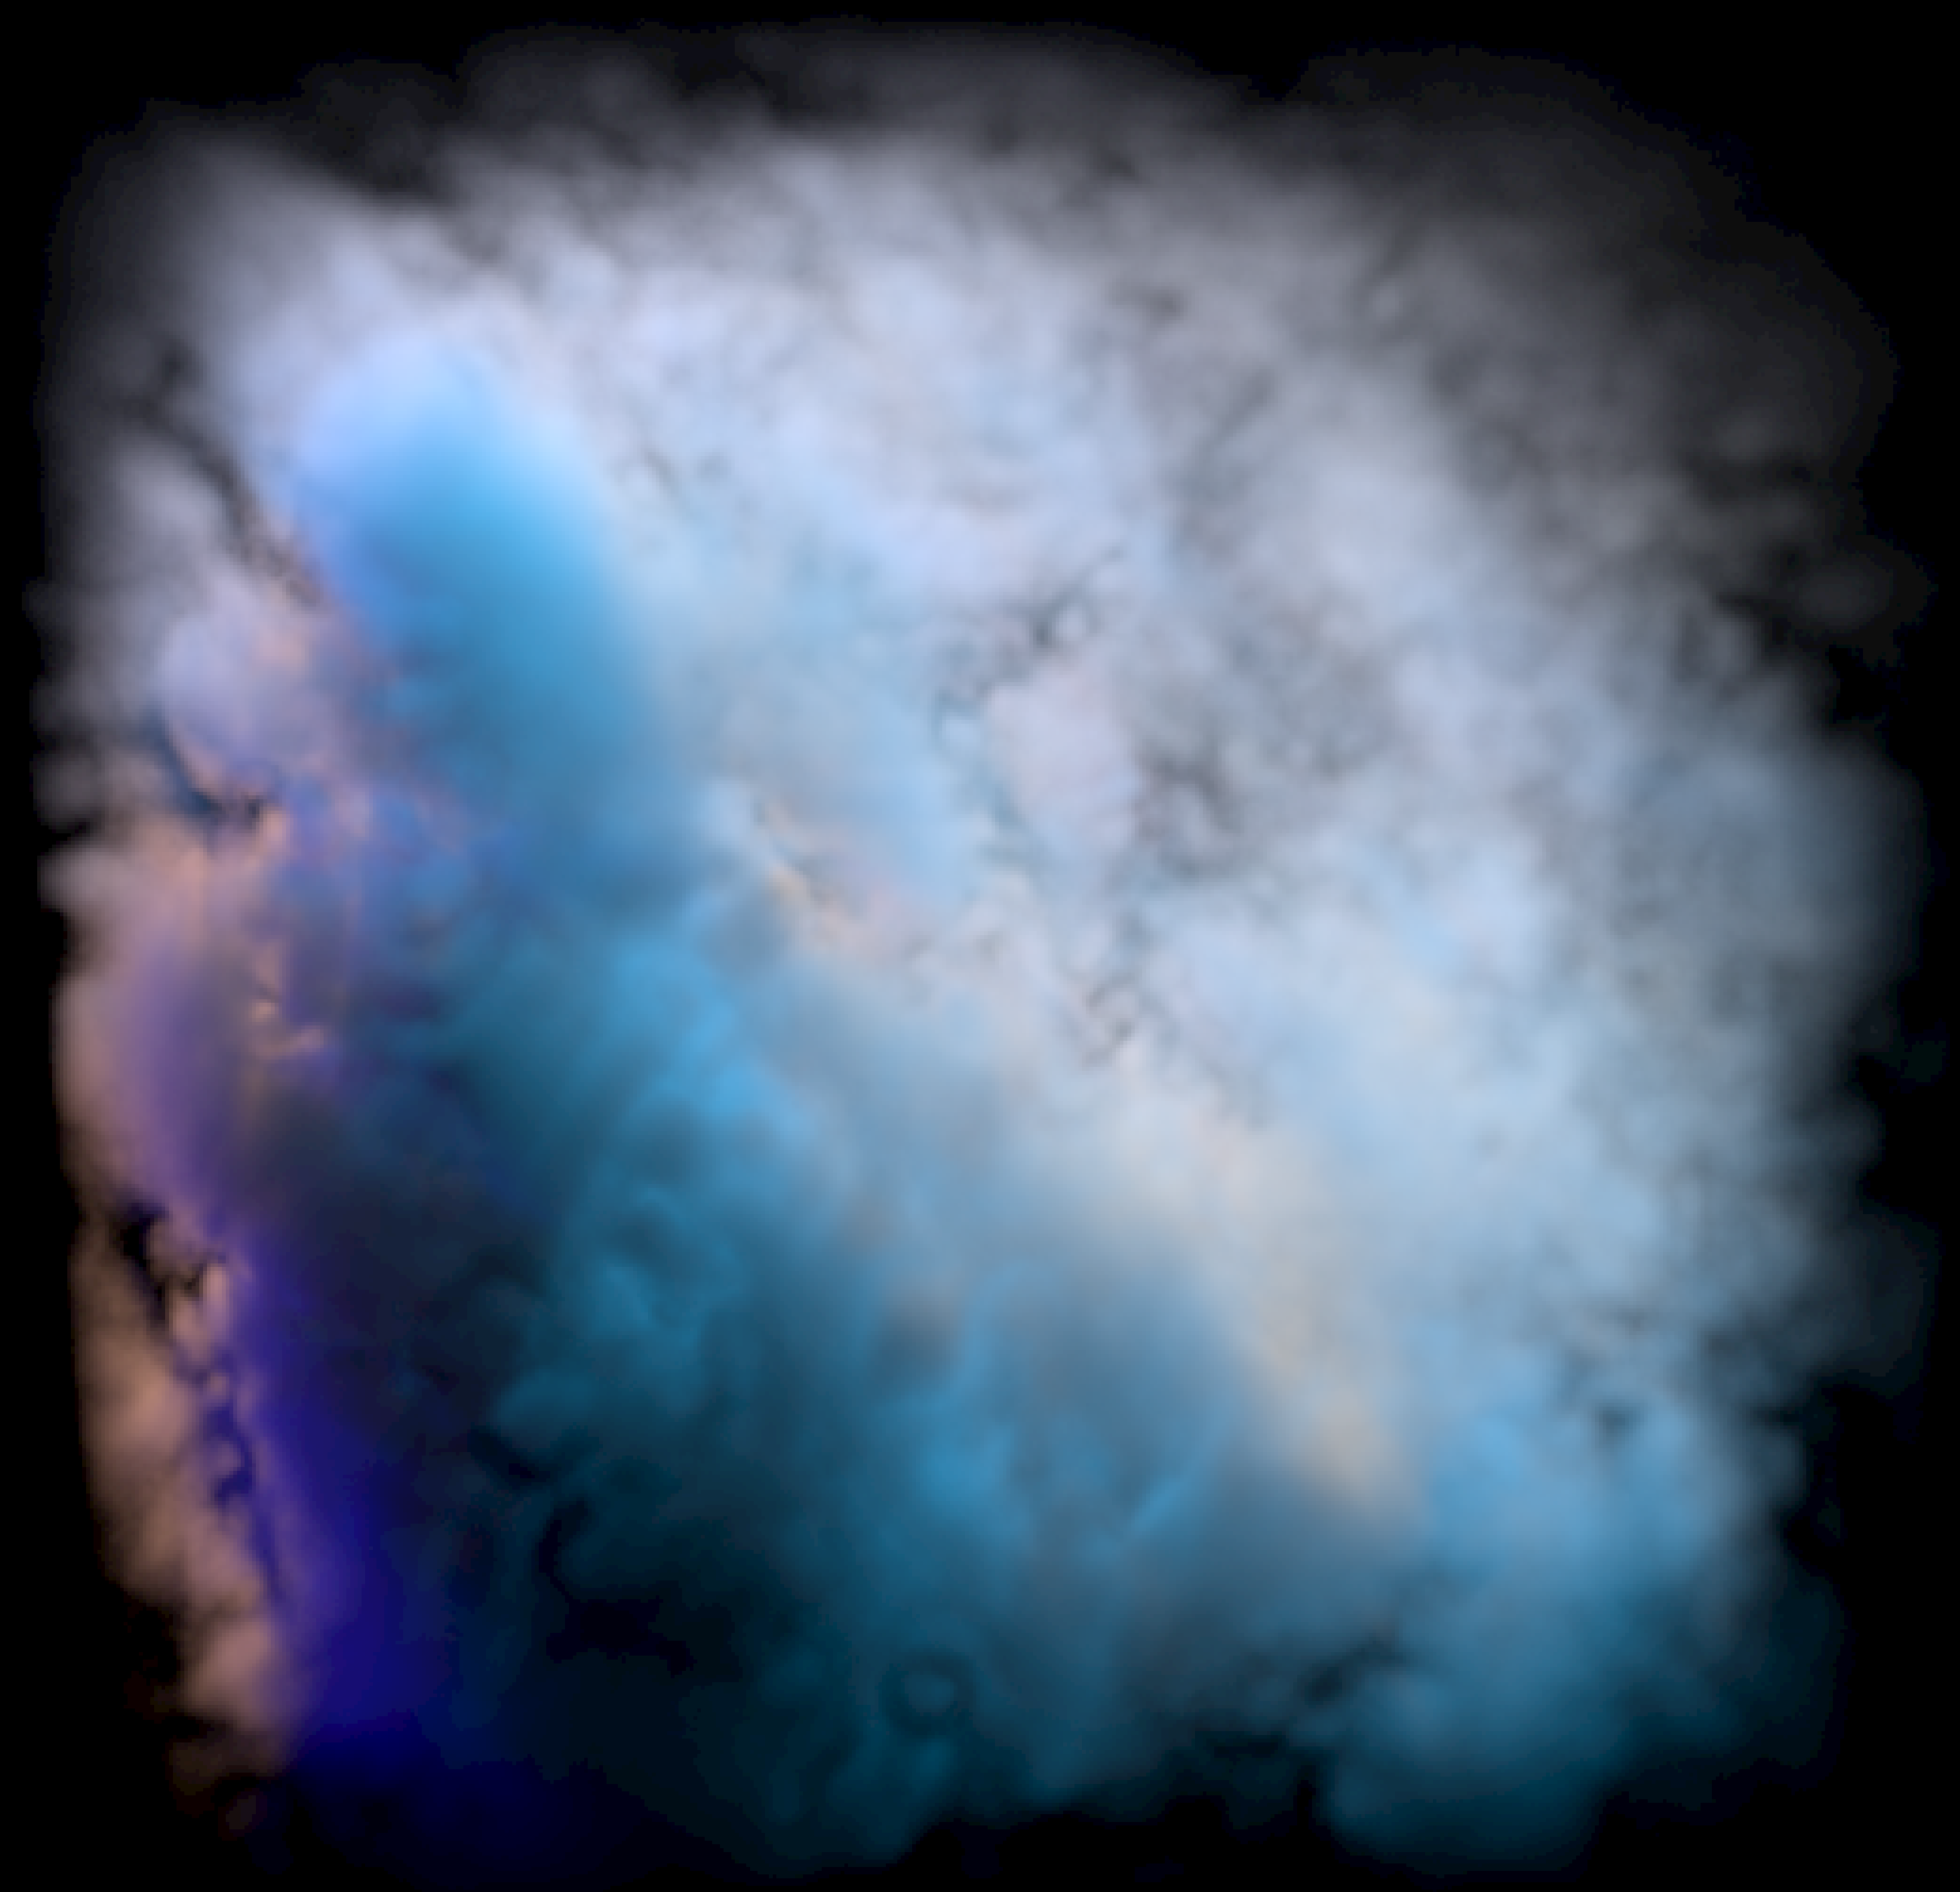

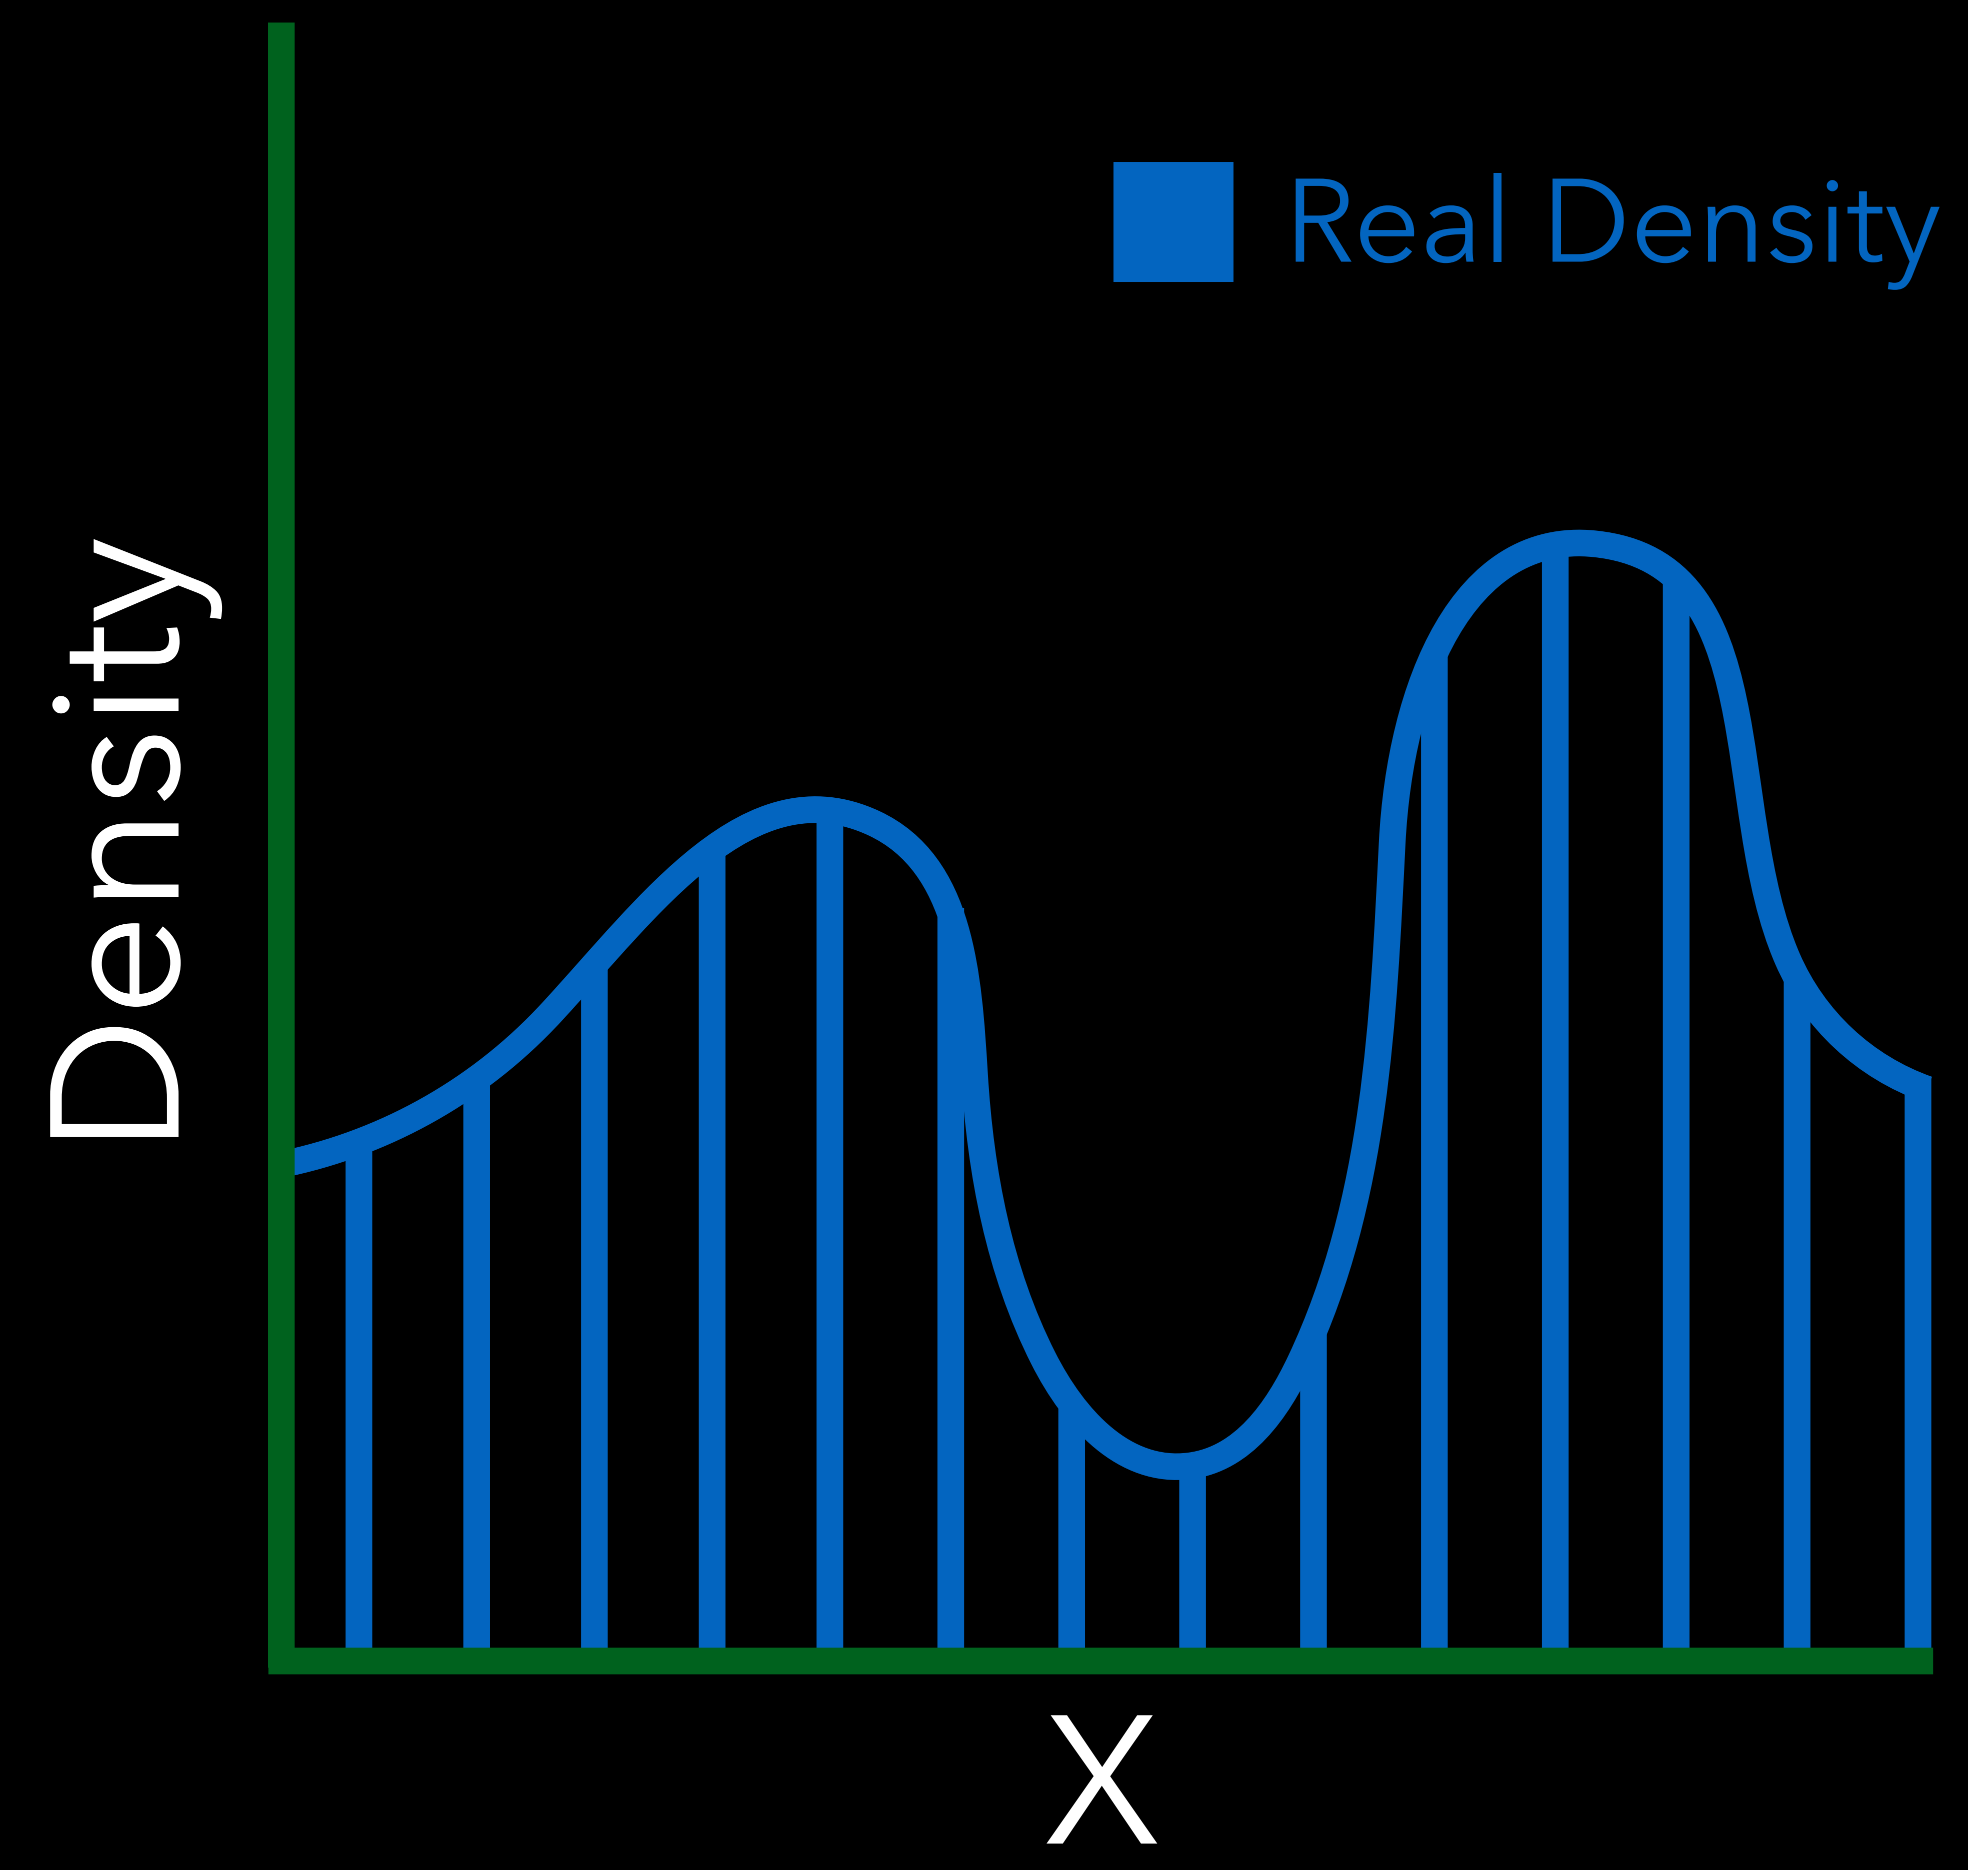
Real Density
Density
X

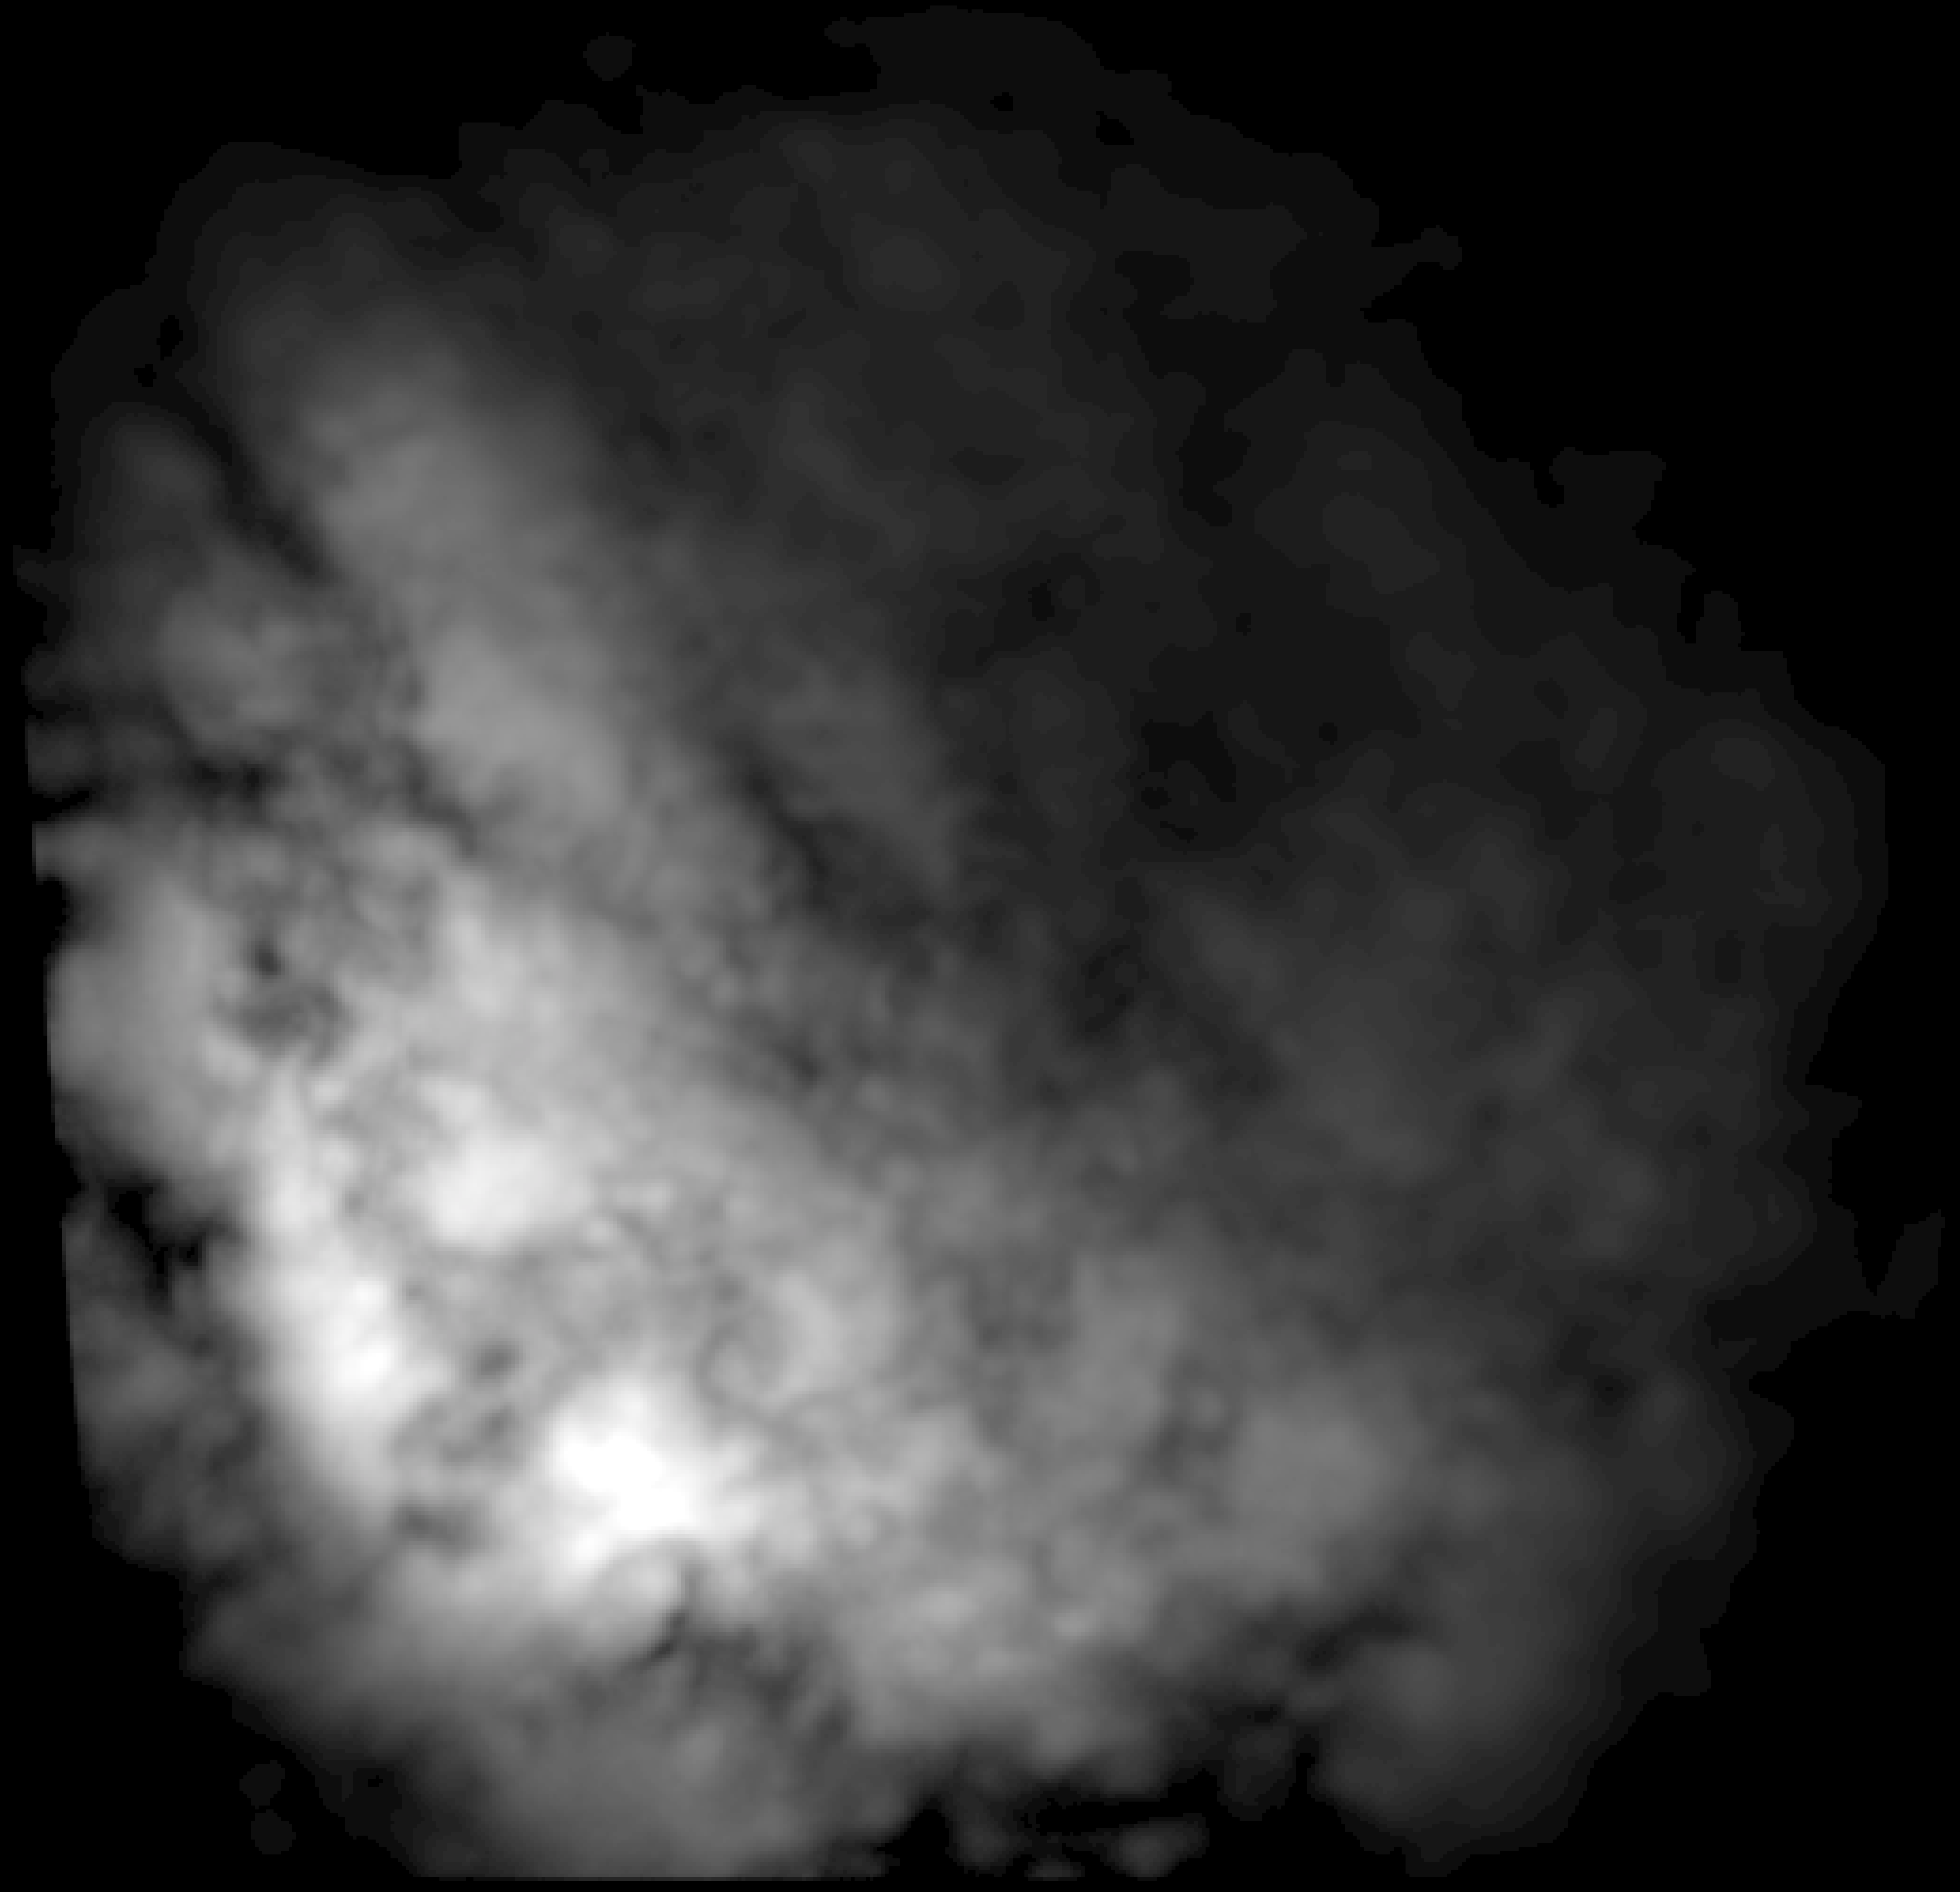

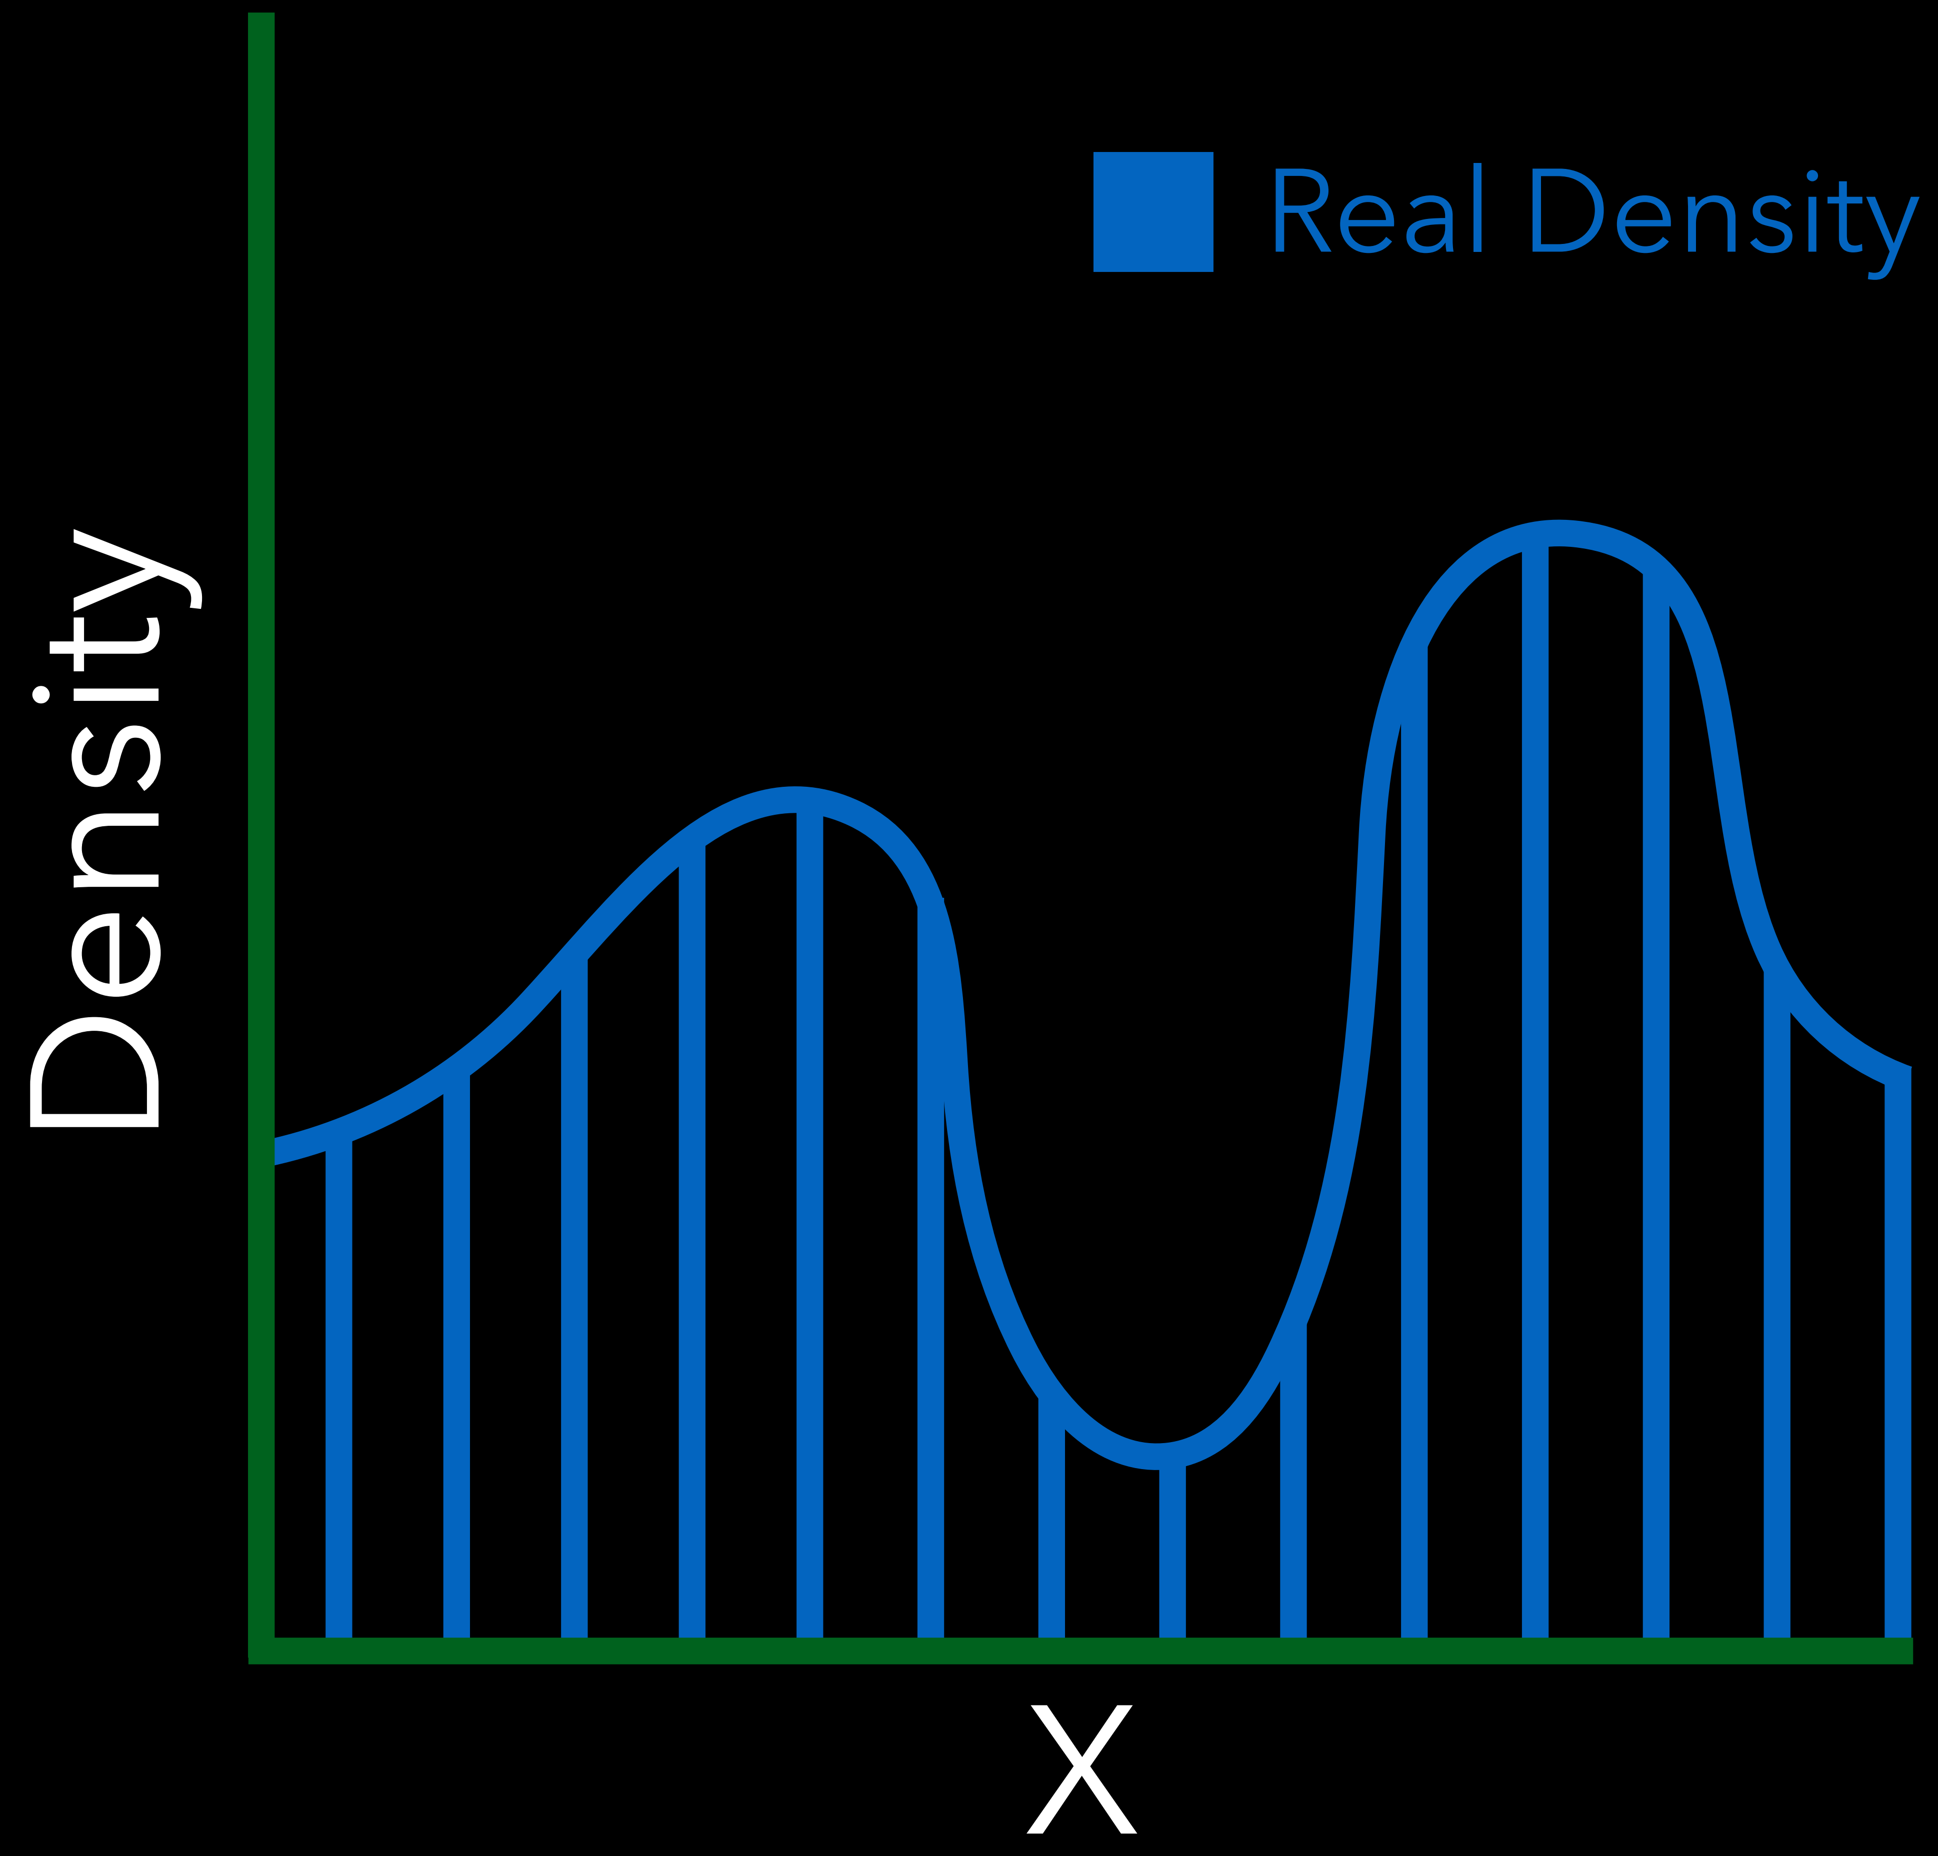

Real Density
Density
X

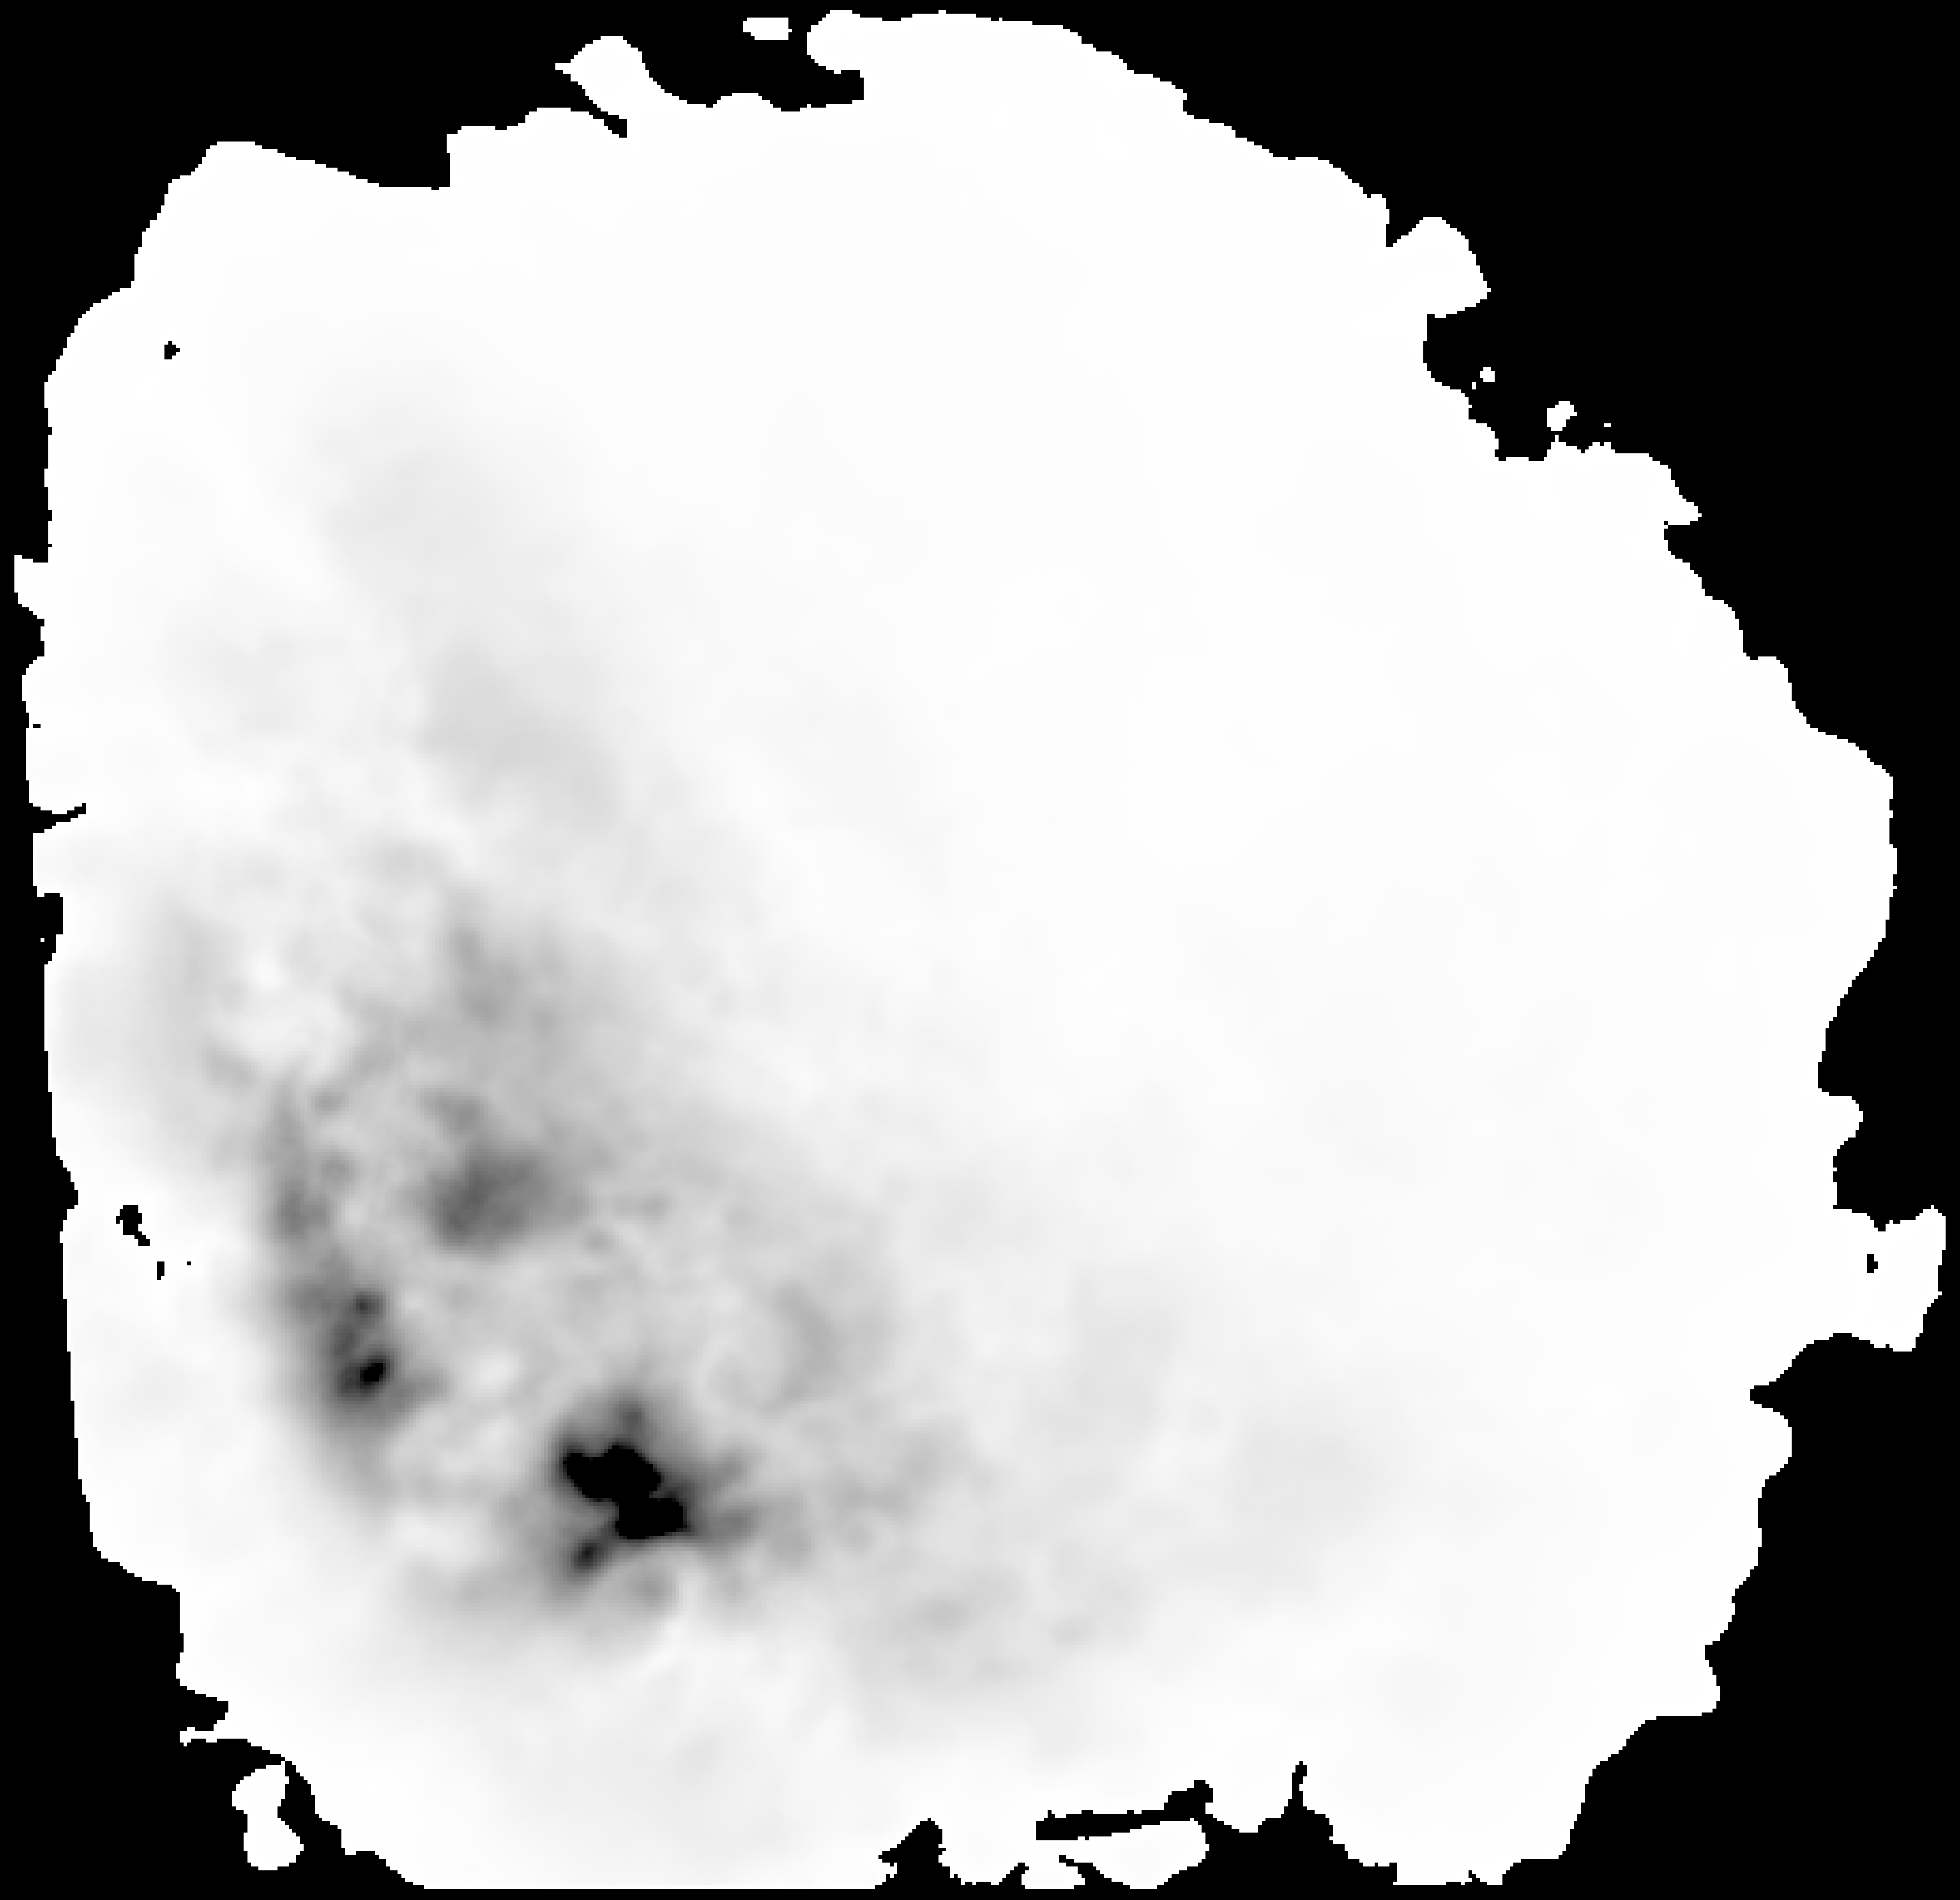

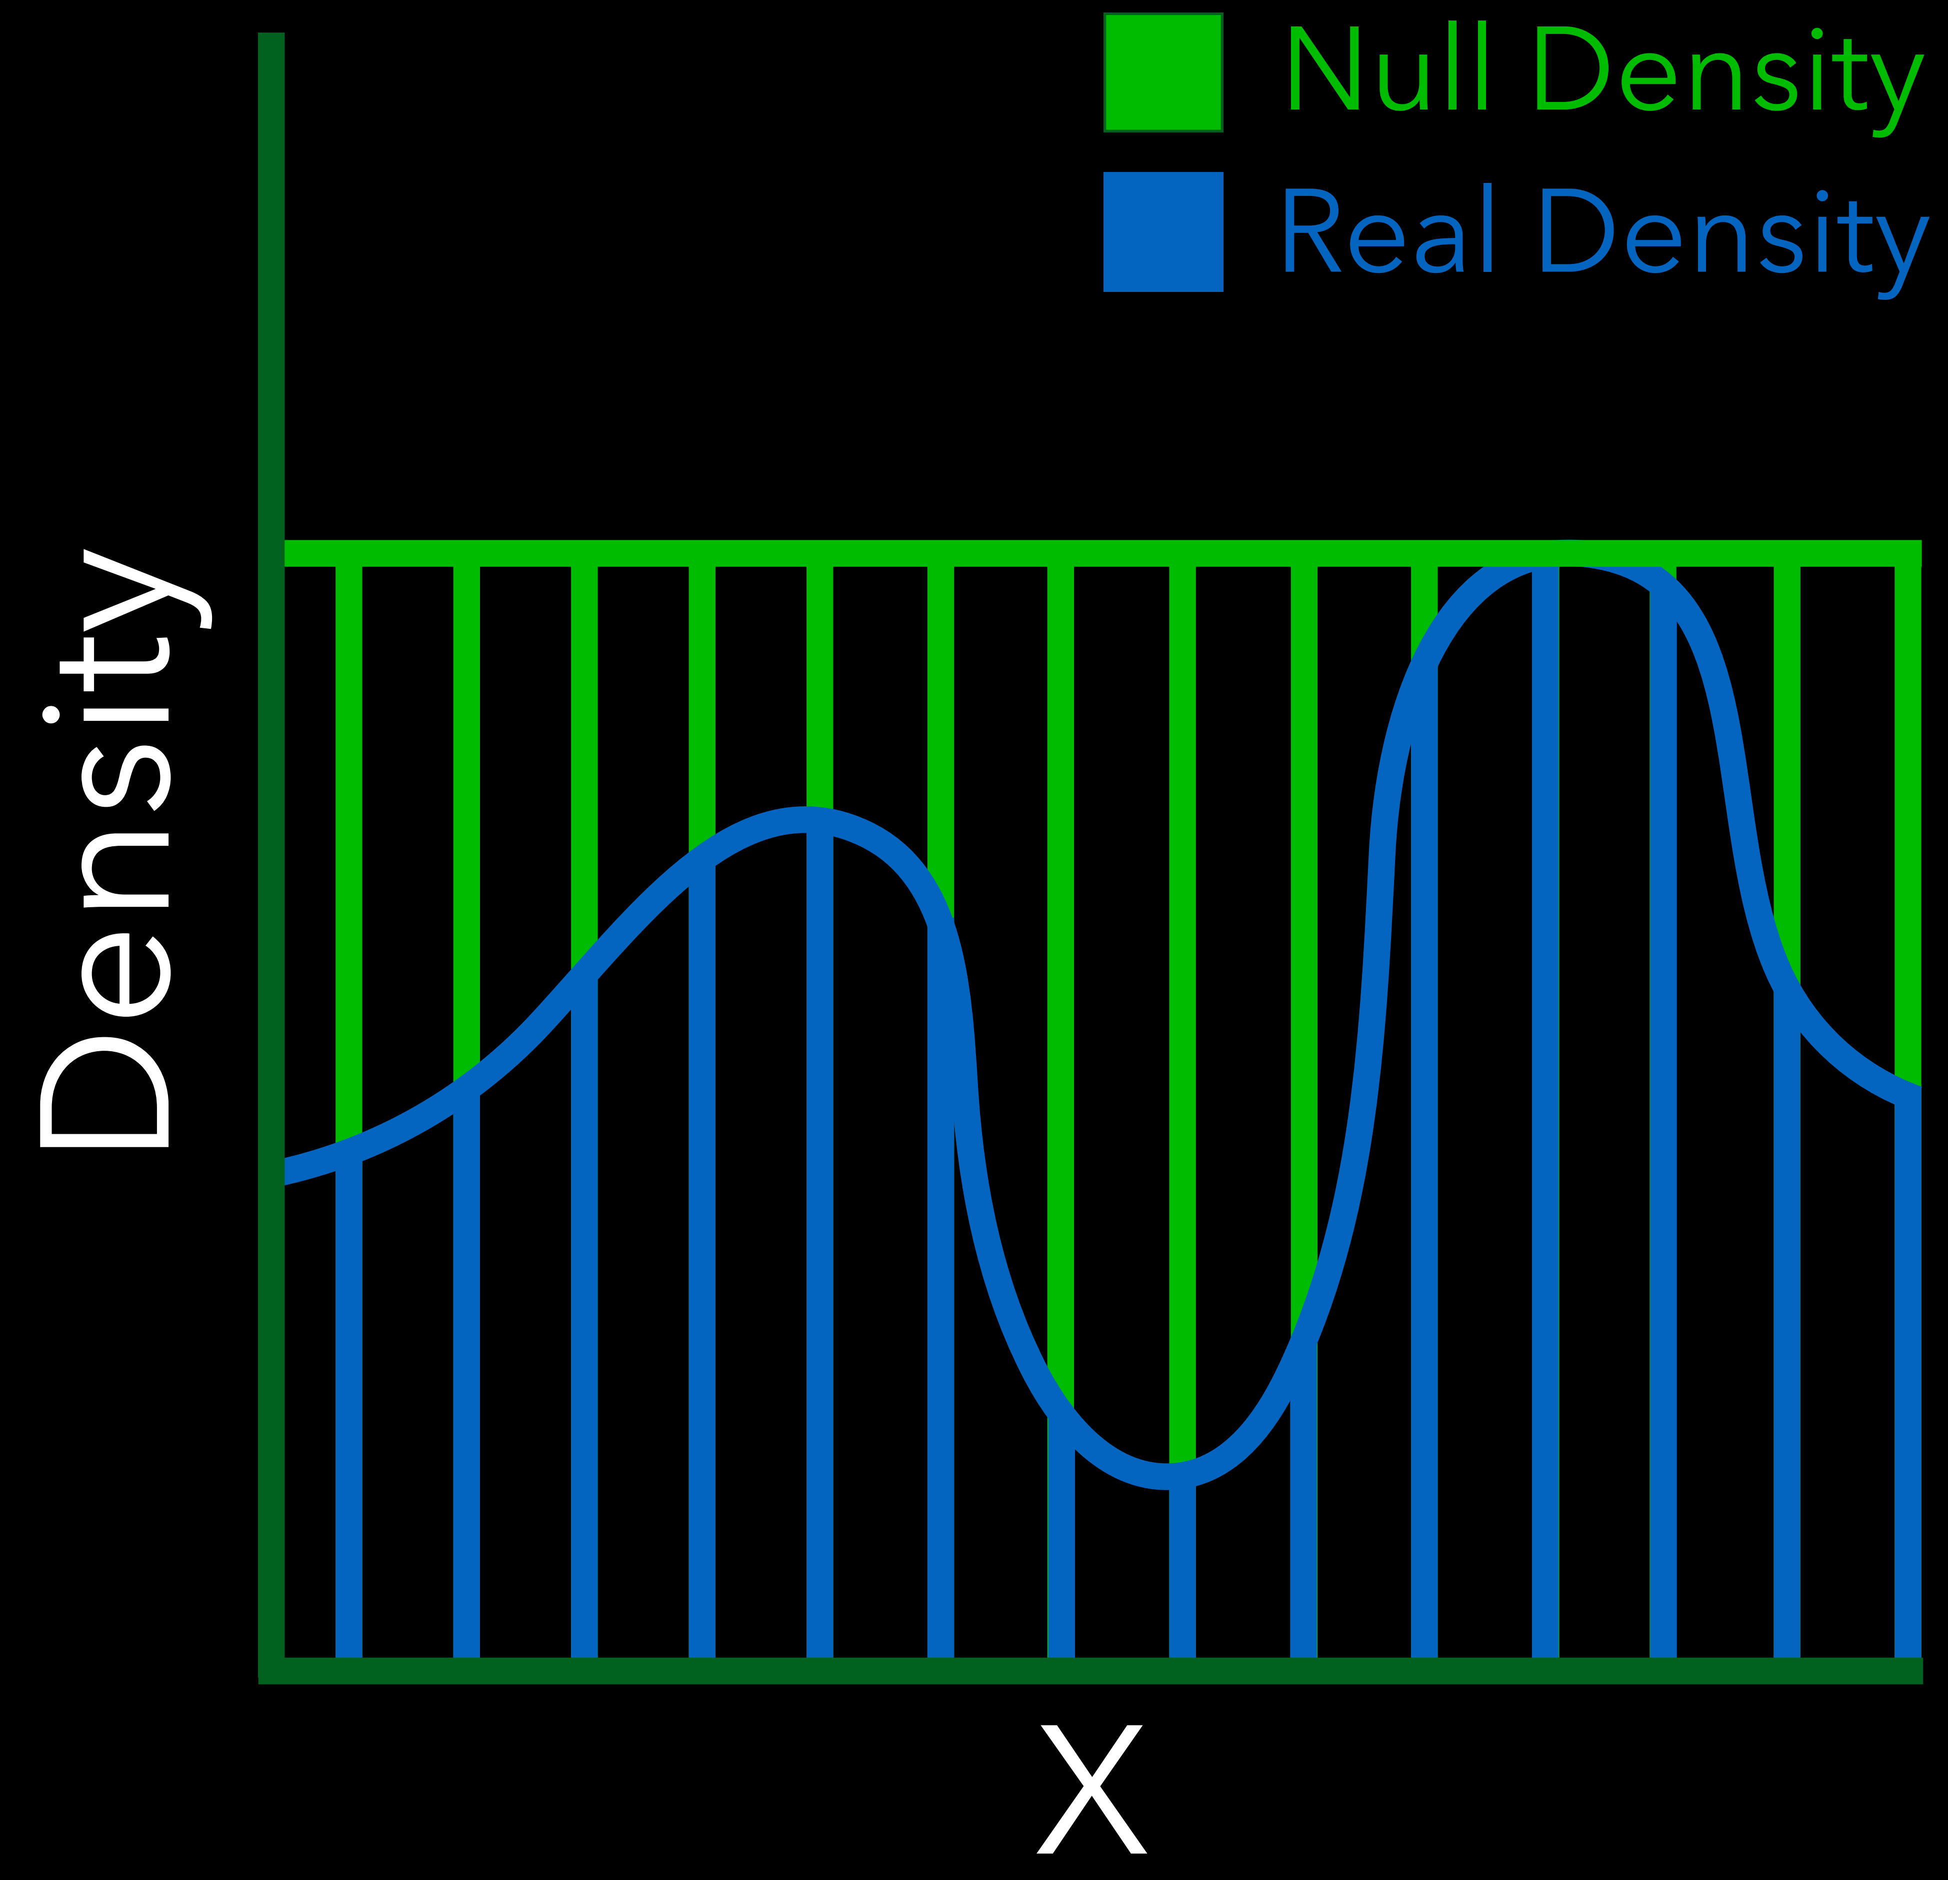
Null Density
Real Density
Density
X

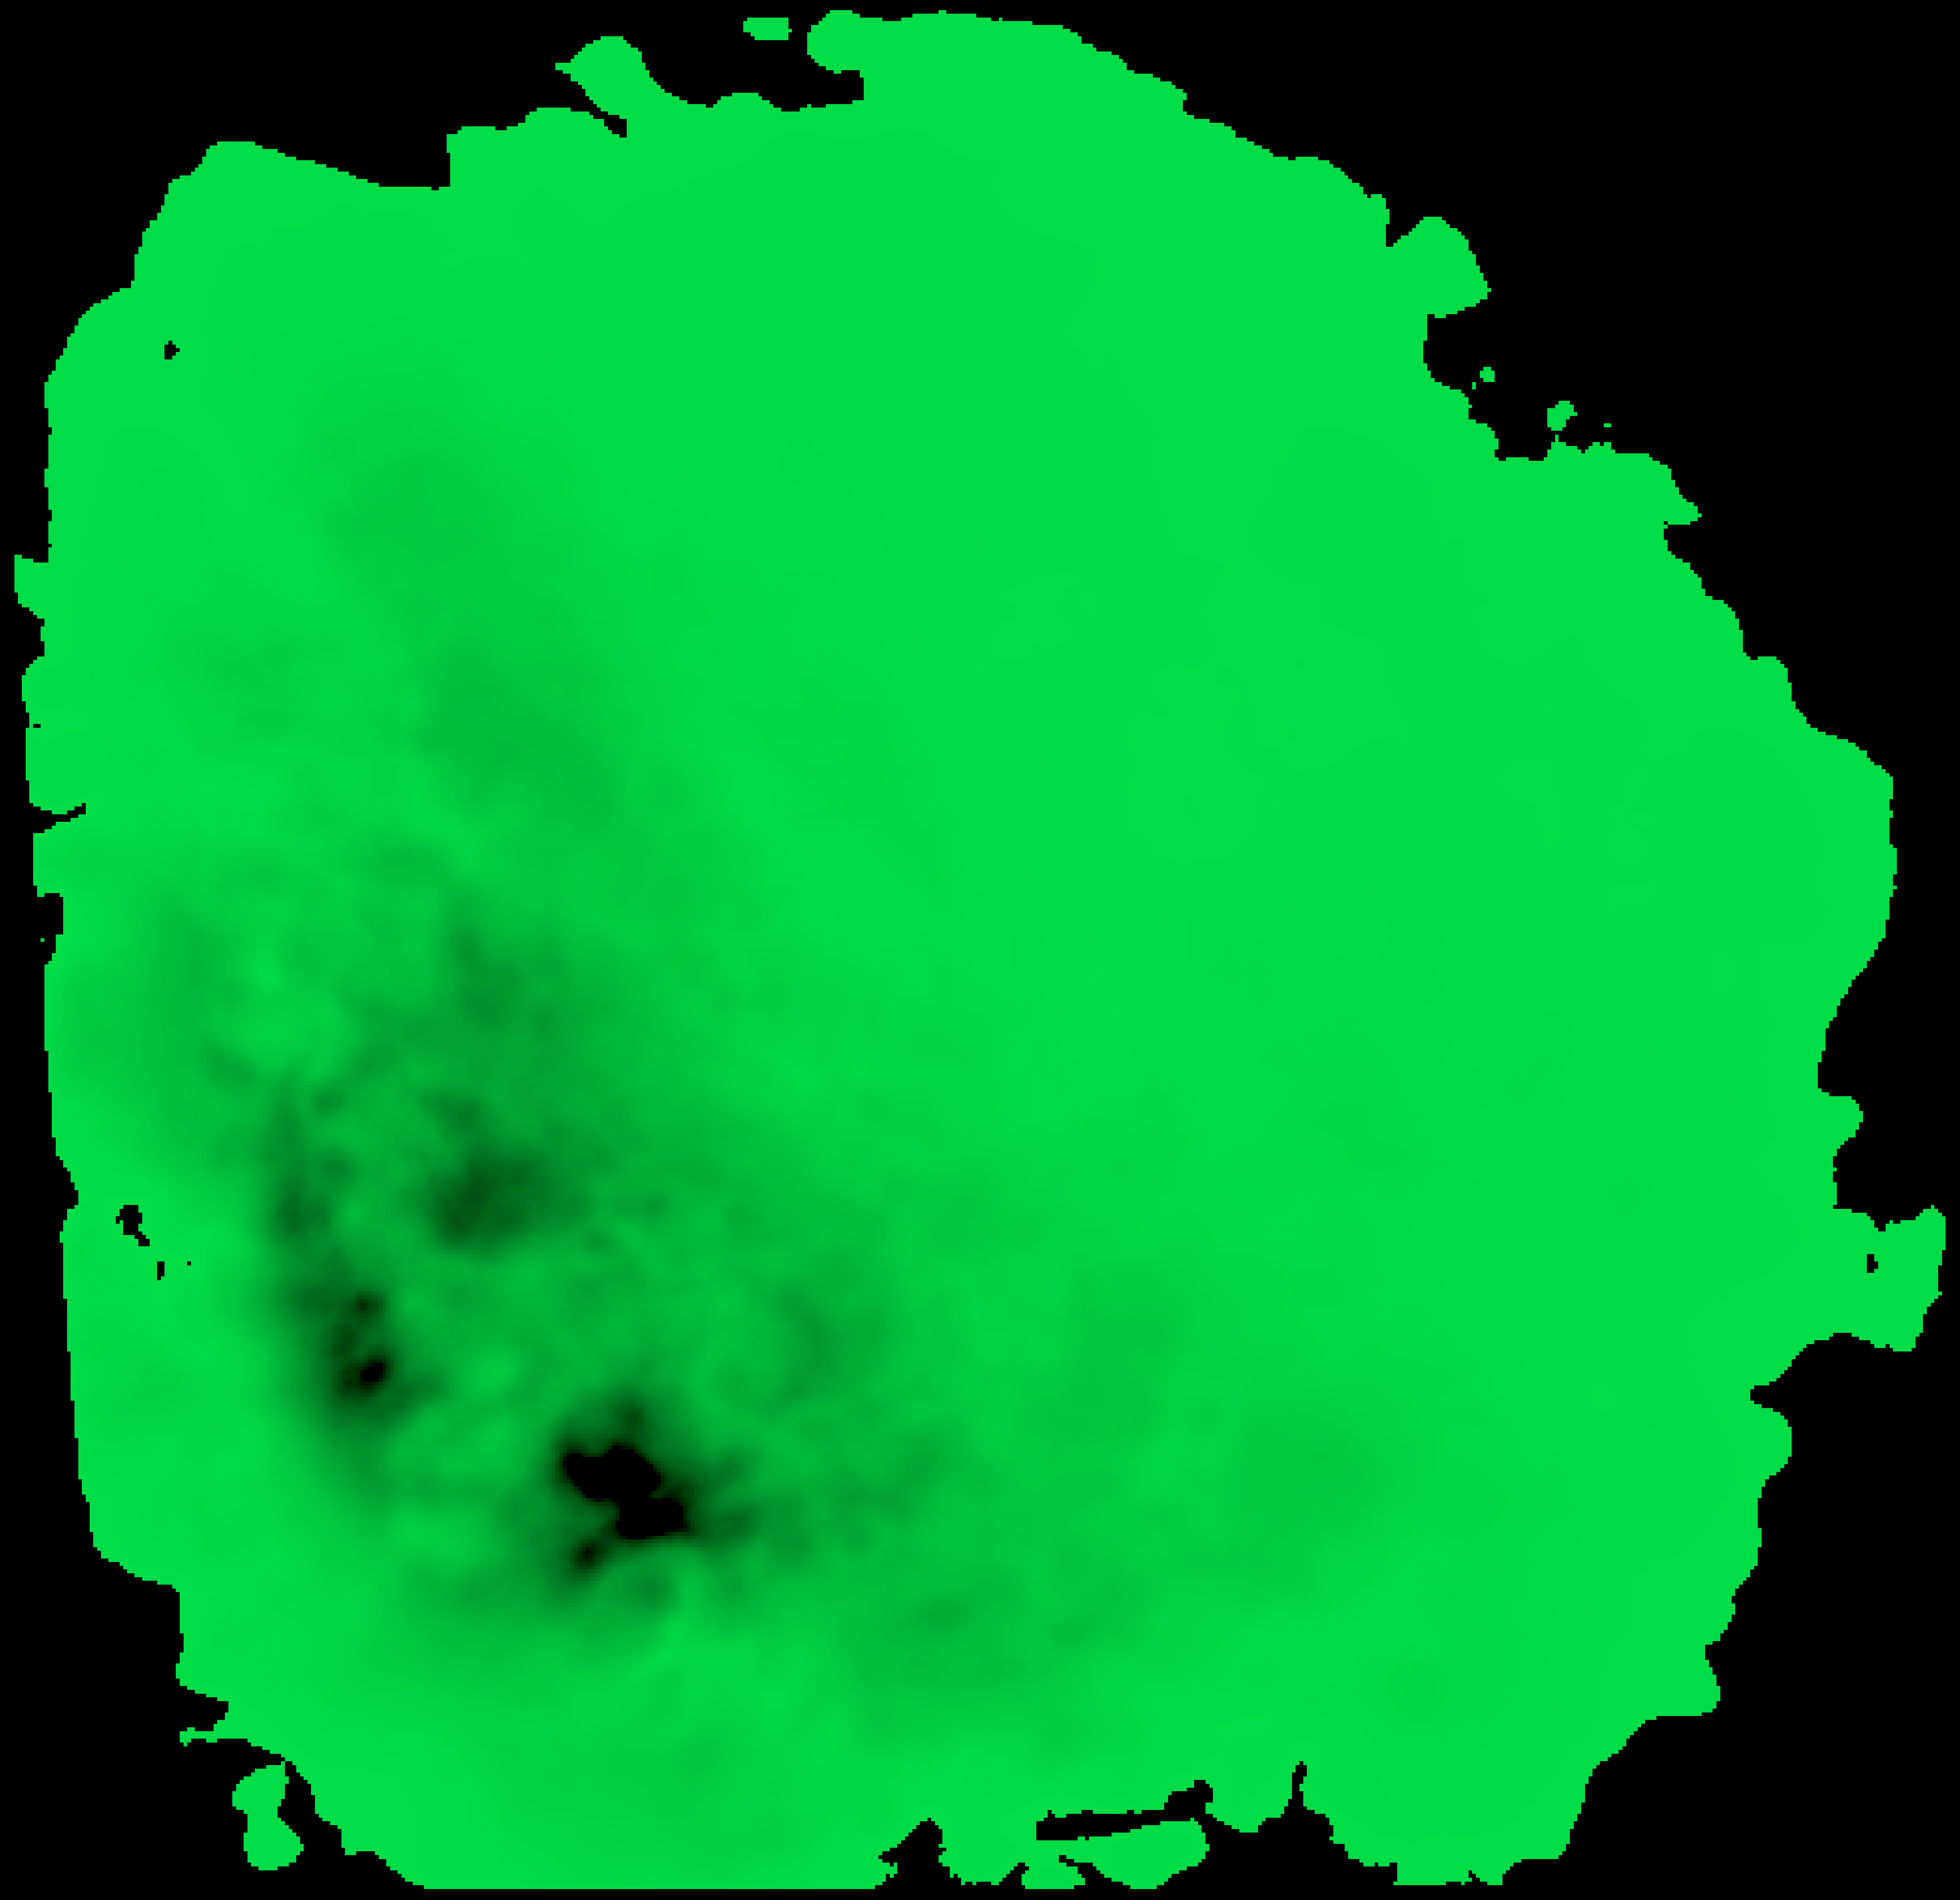

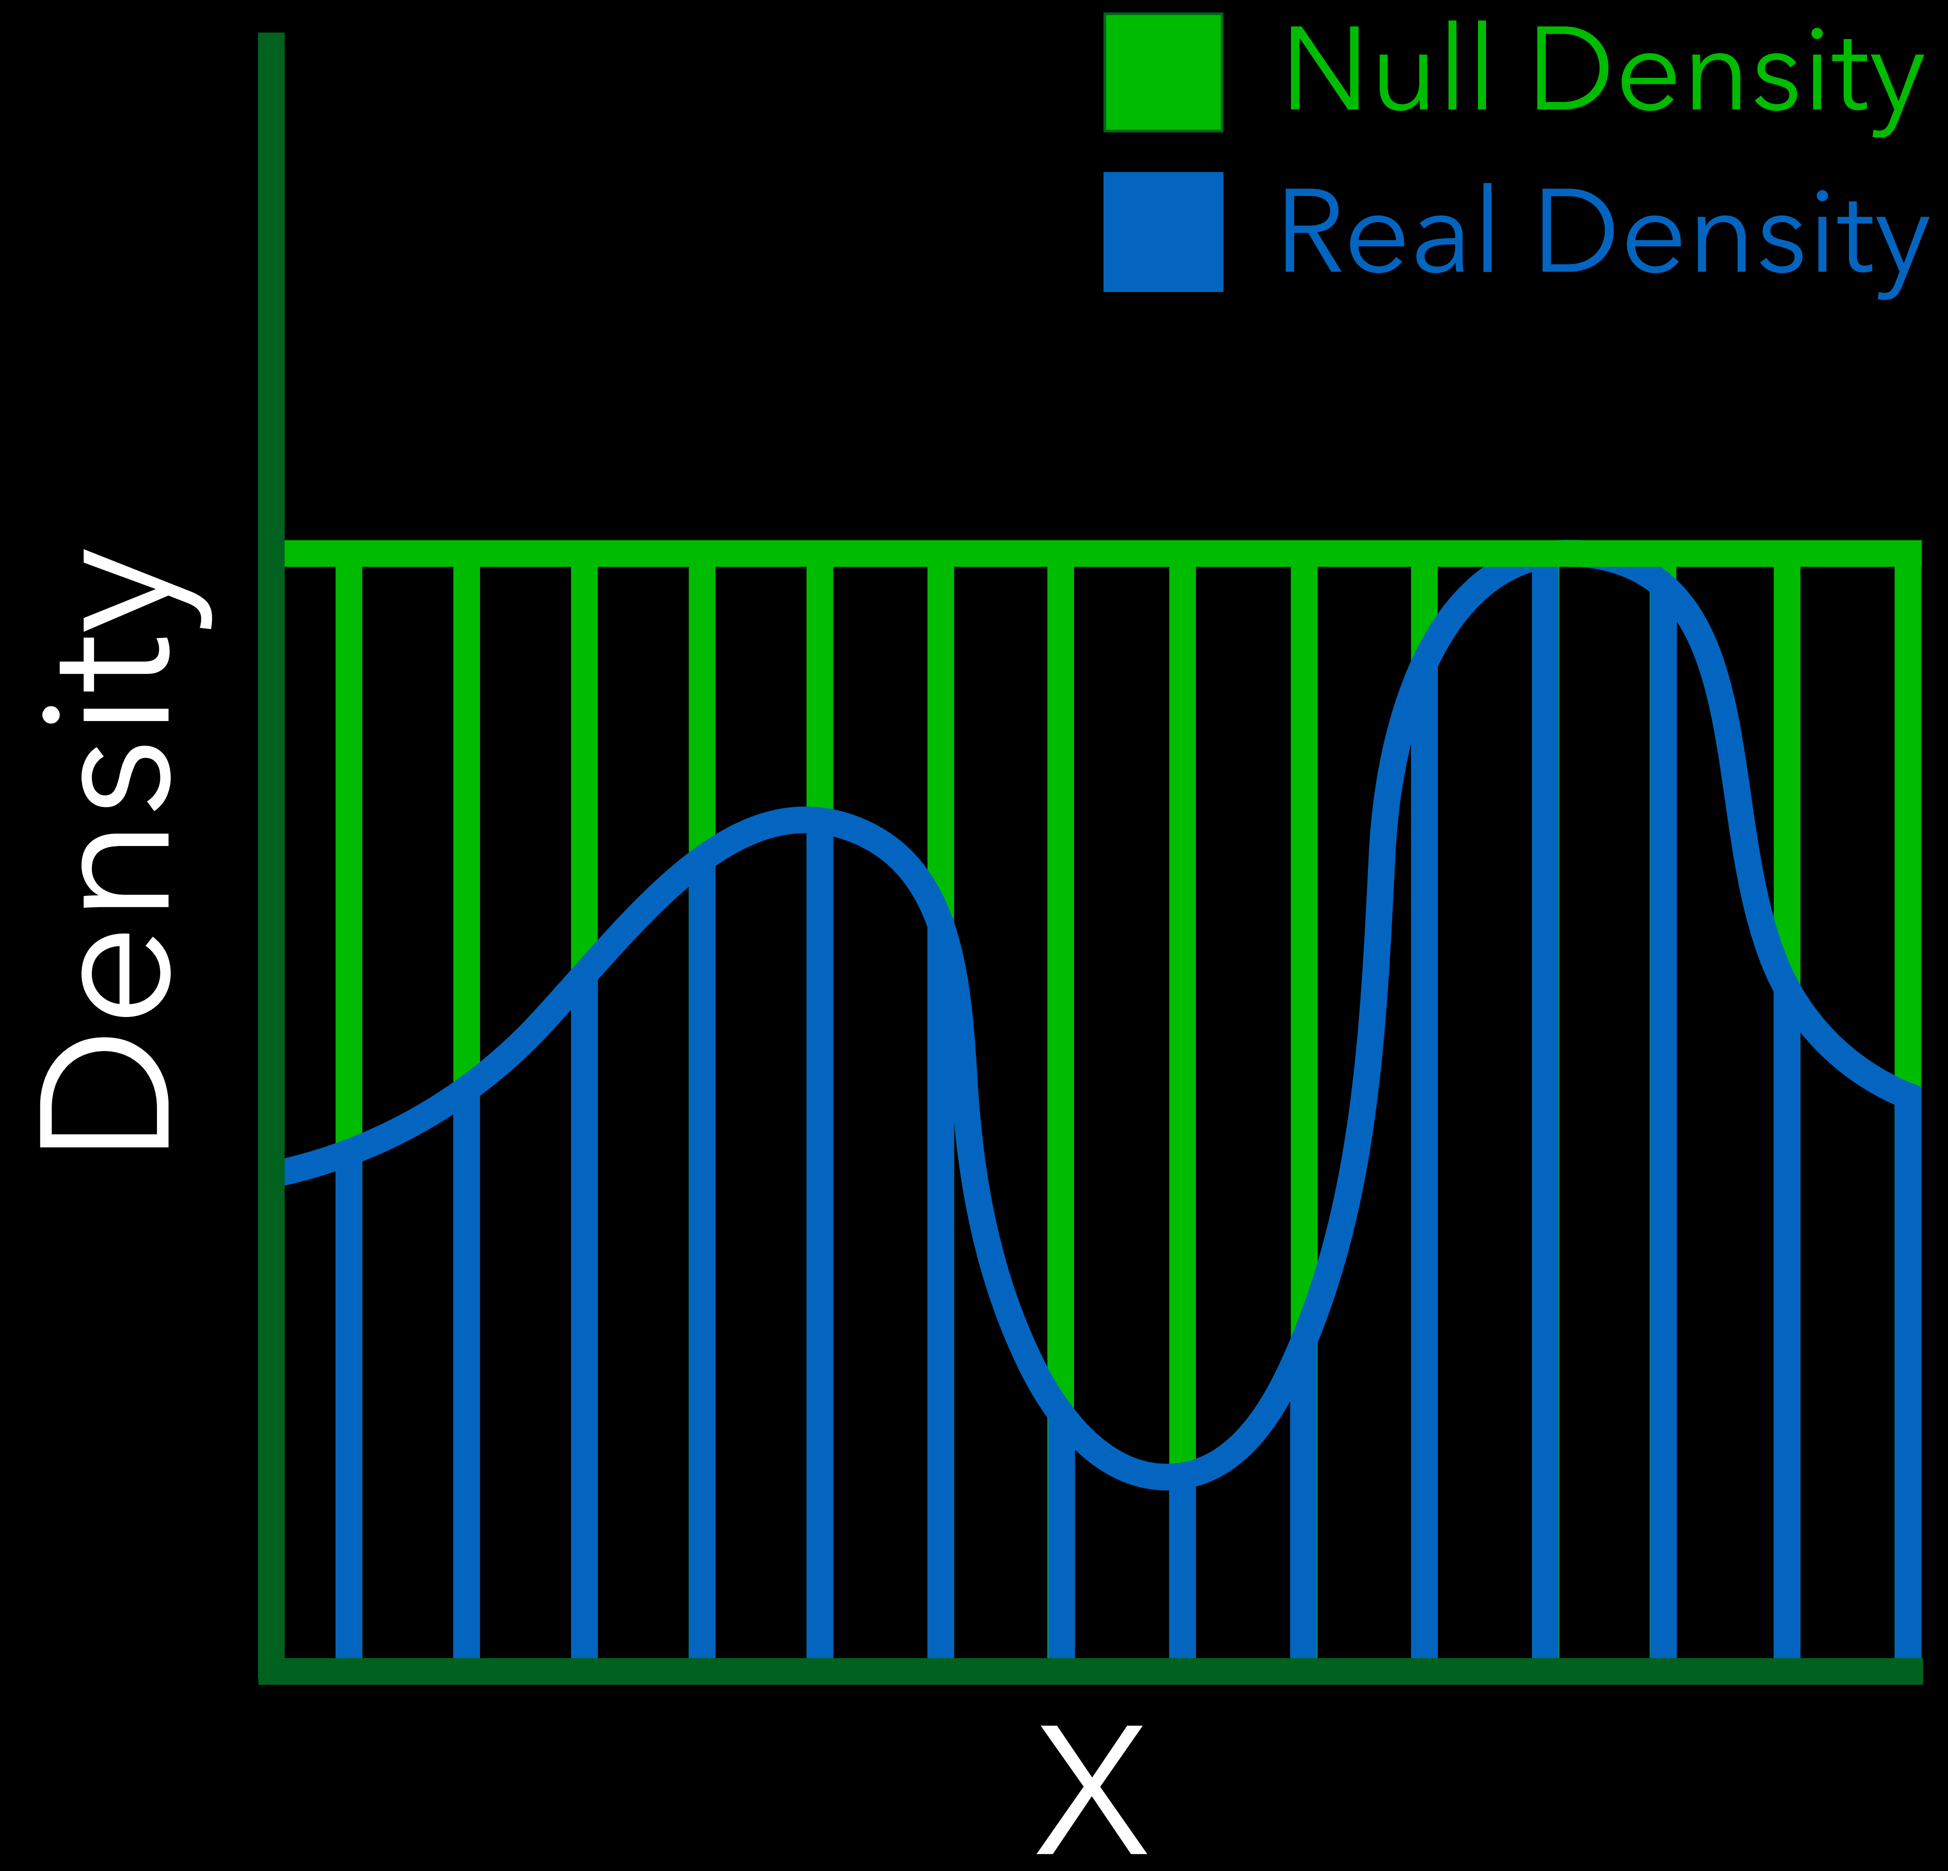
Null Density
Real Density
Density
X

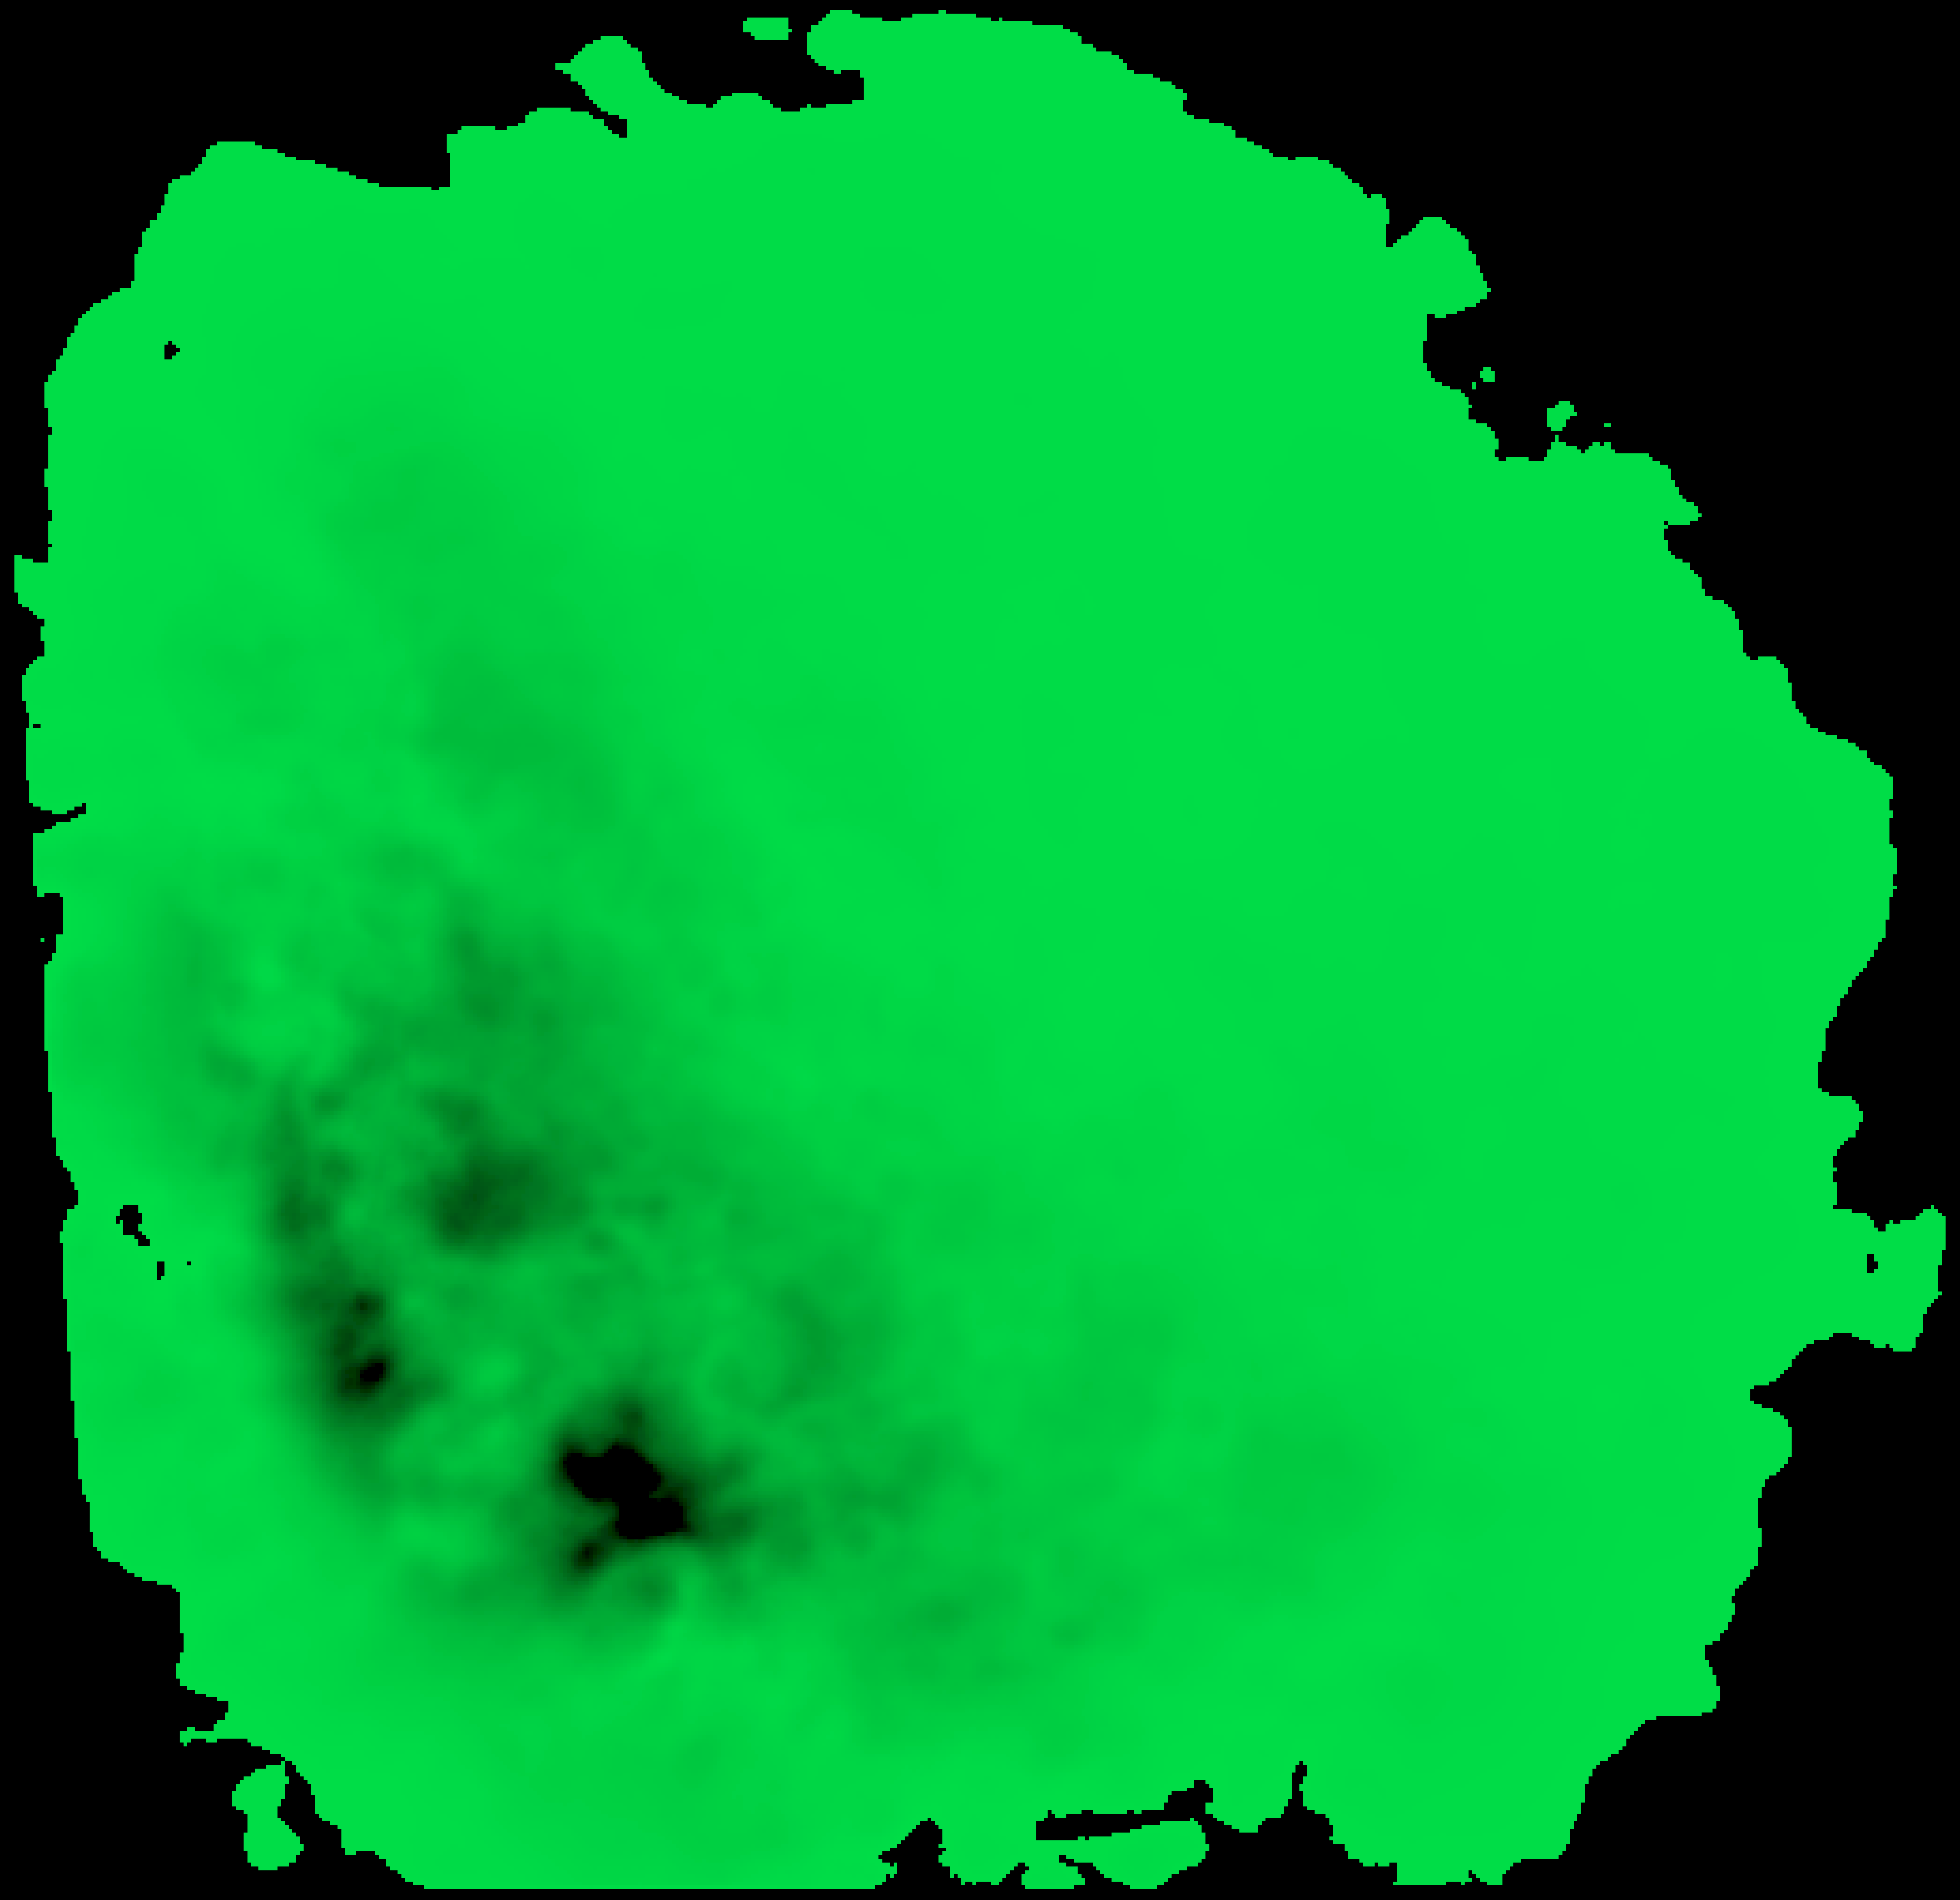

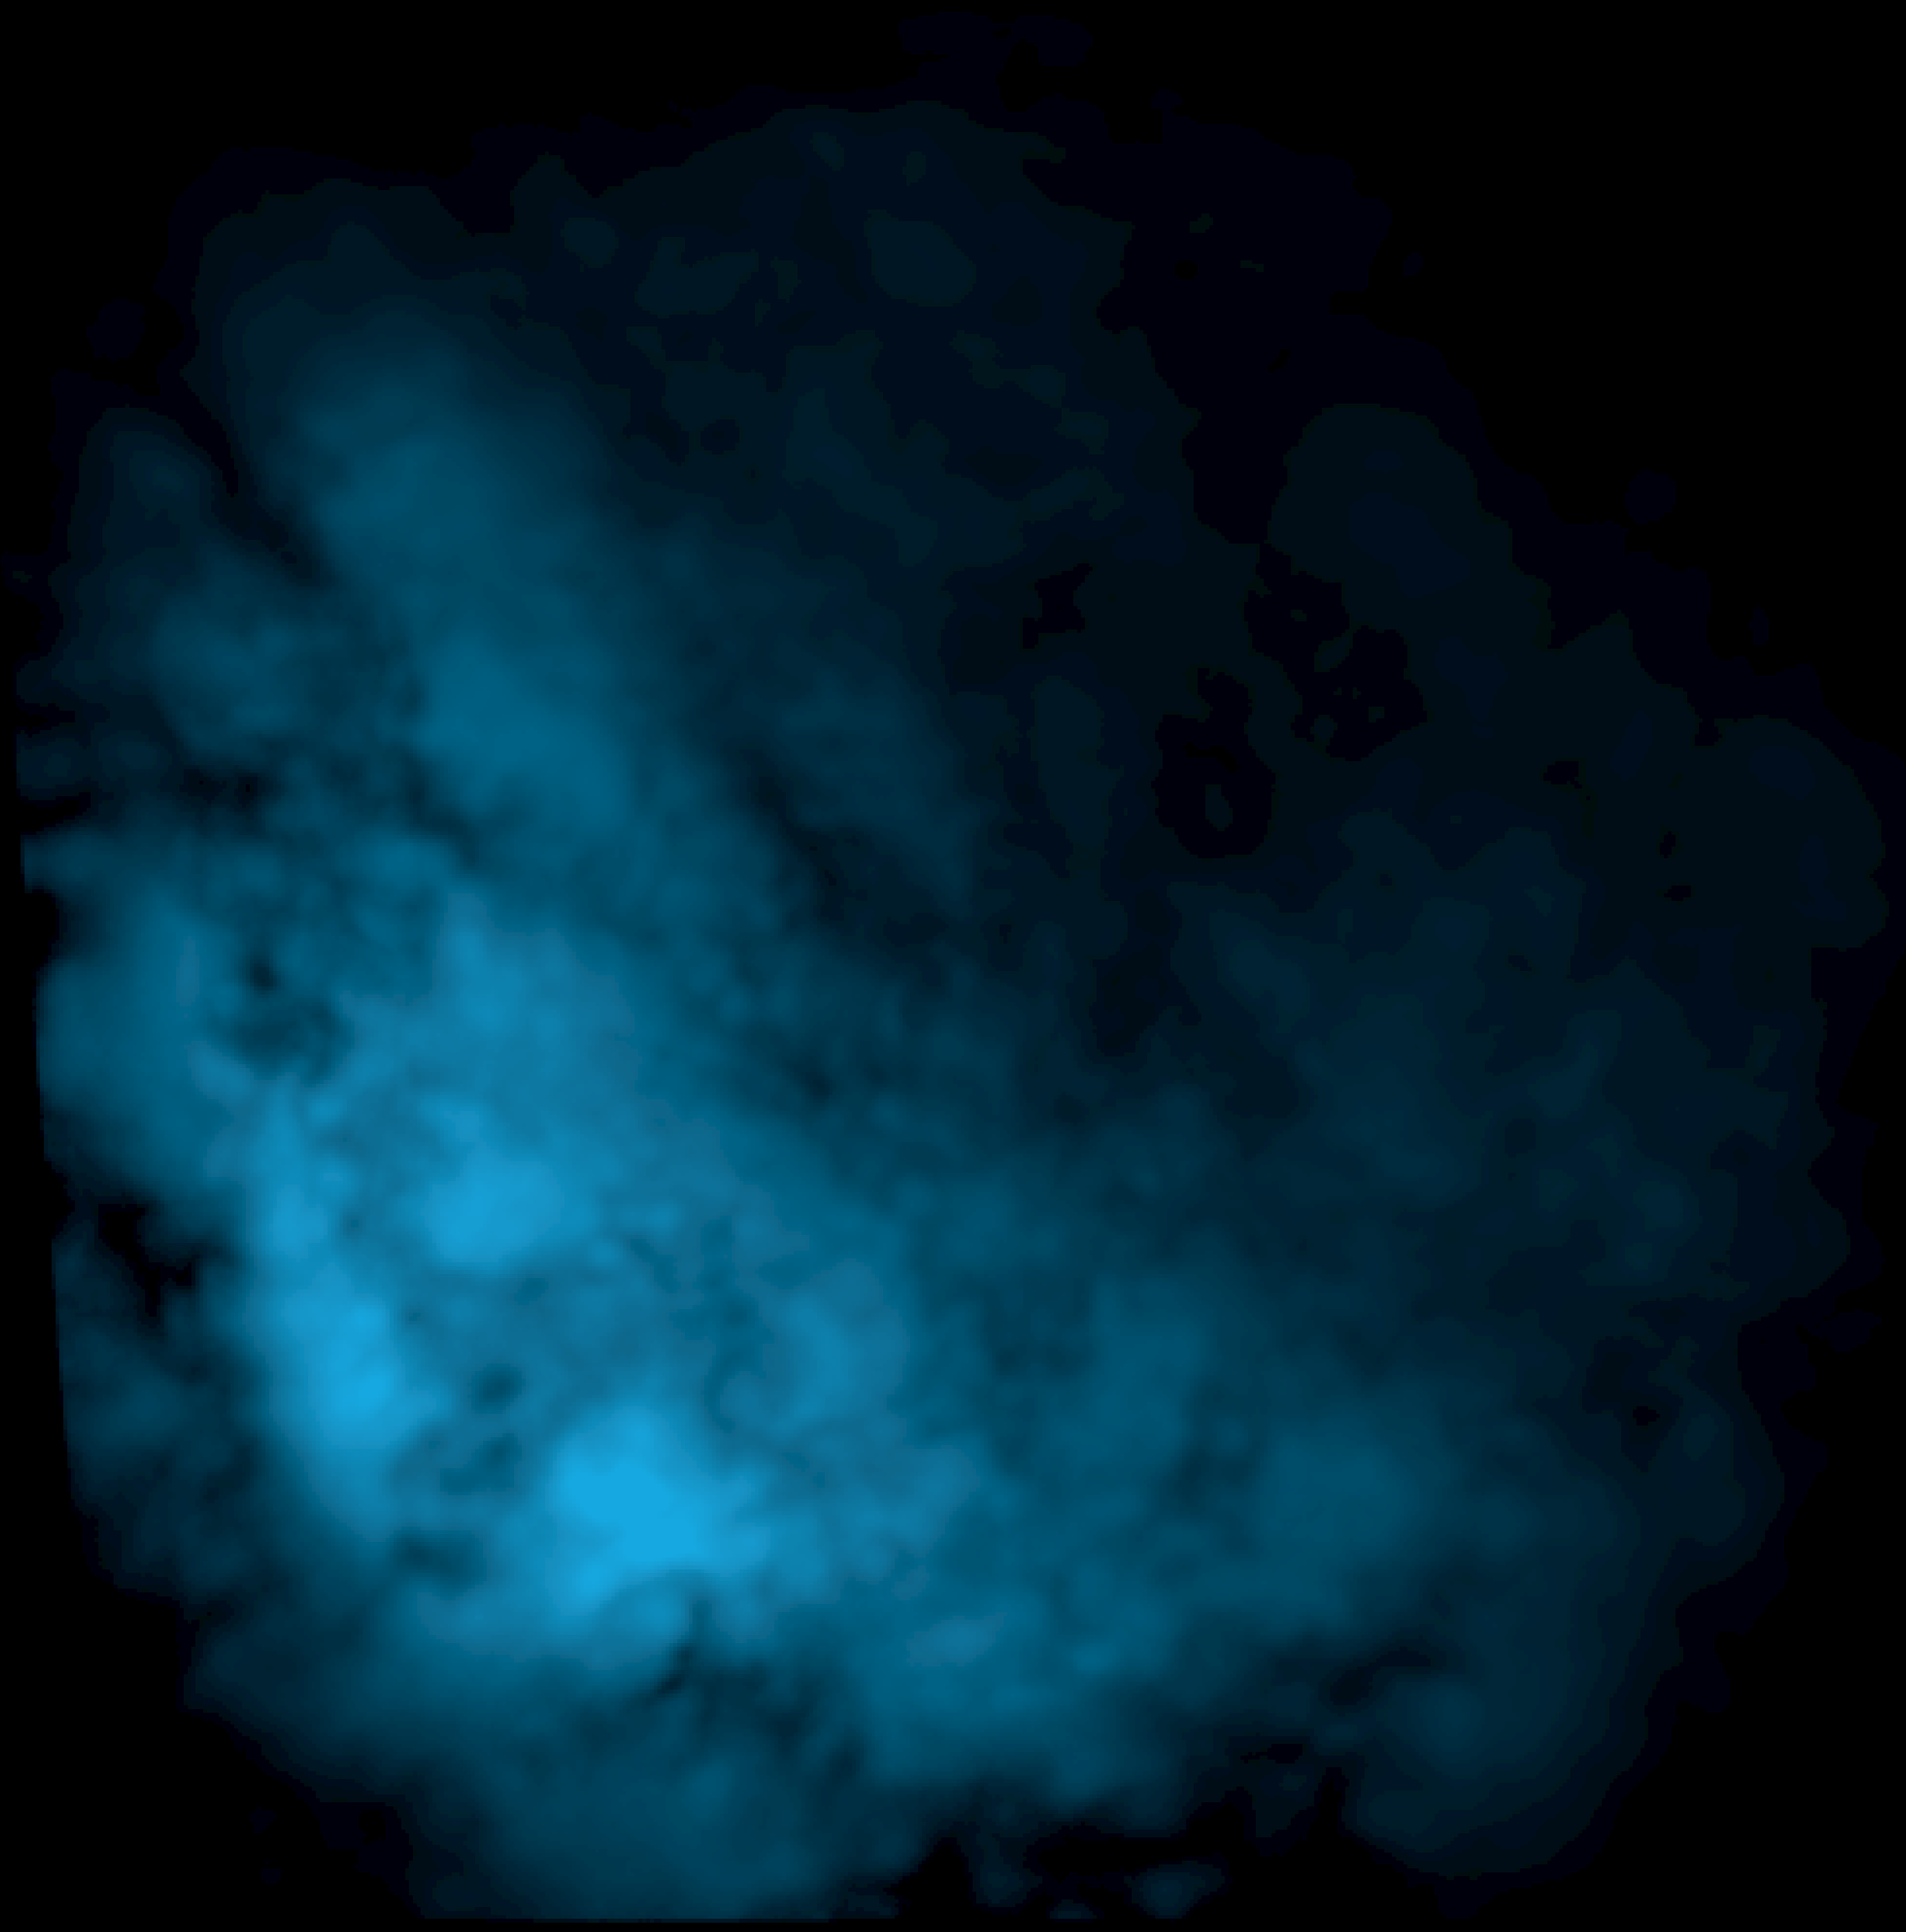

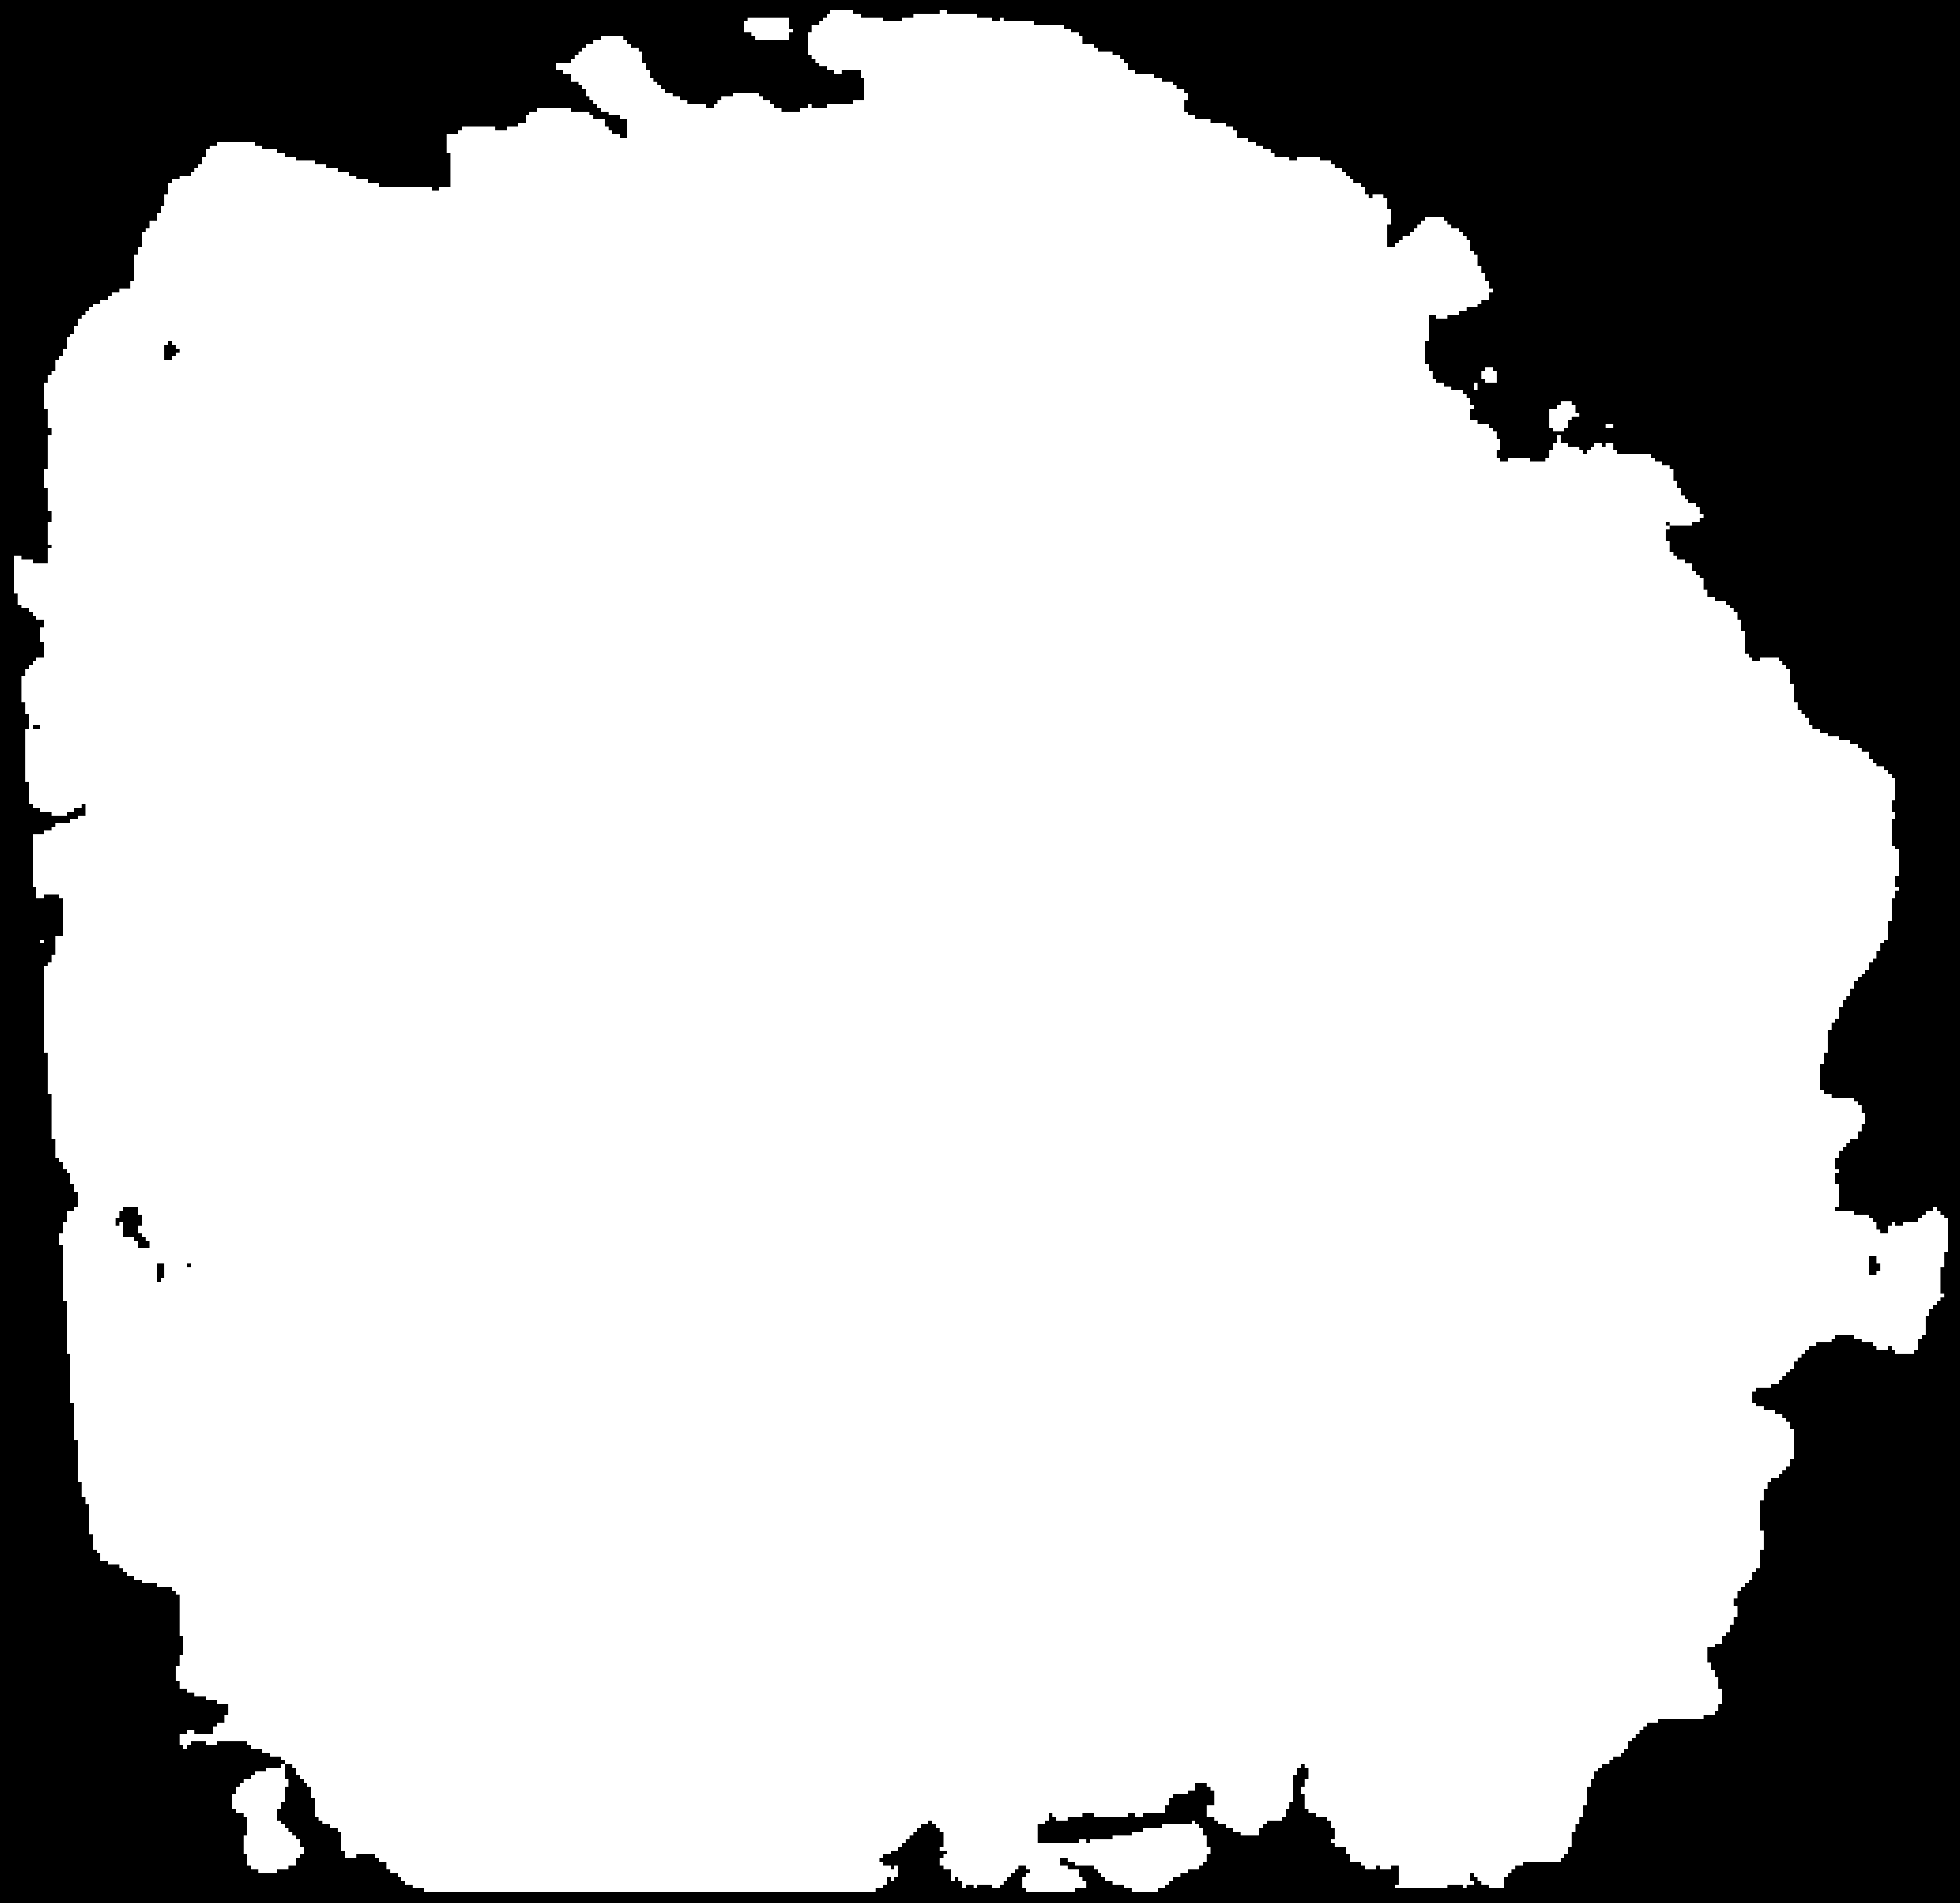

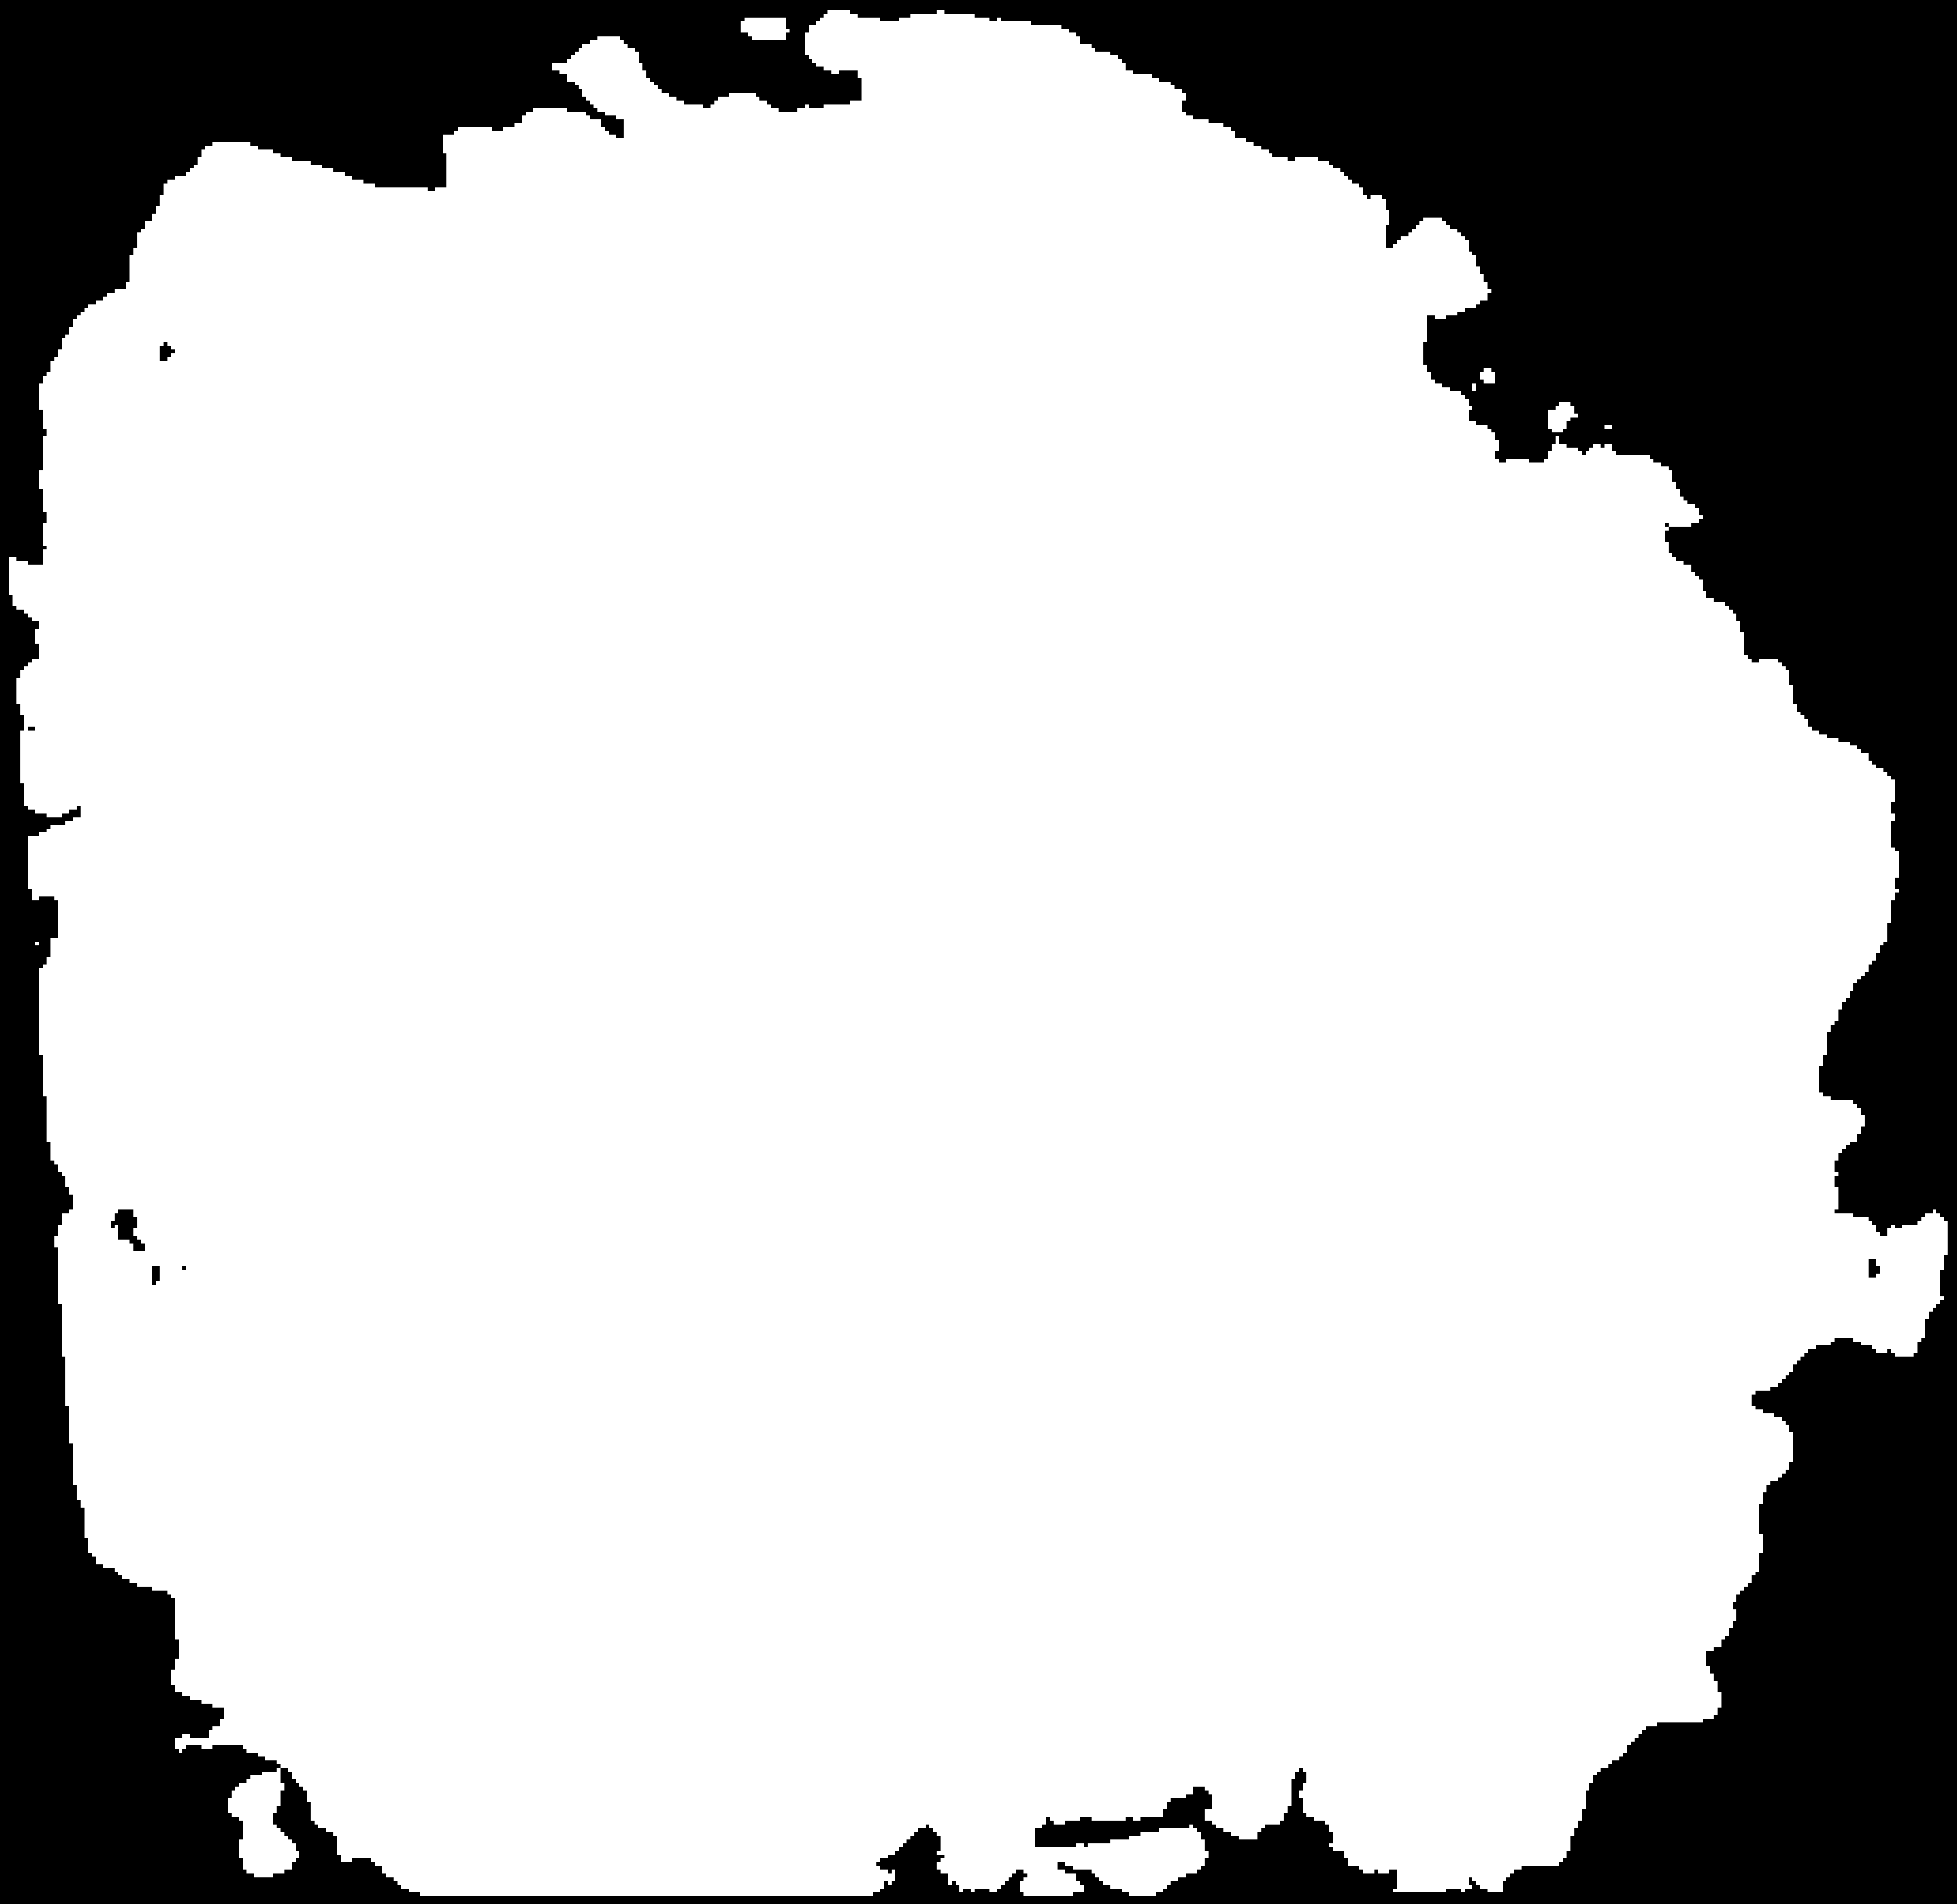

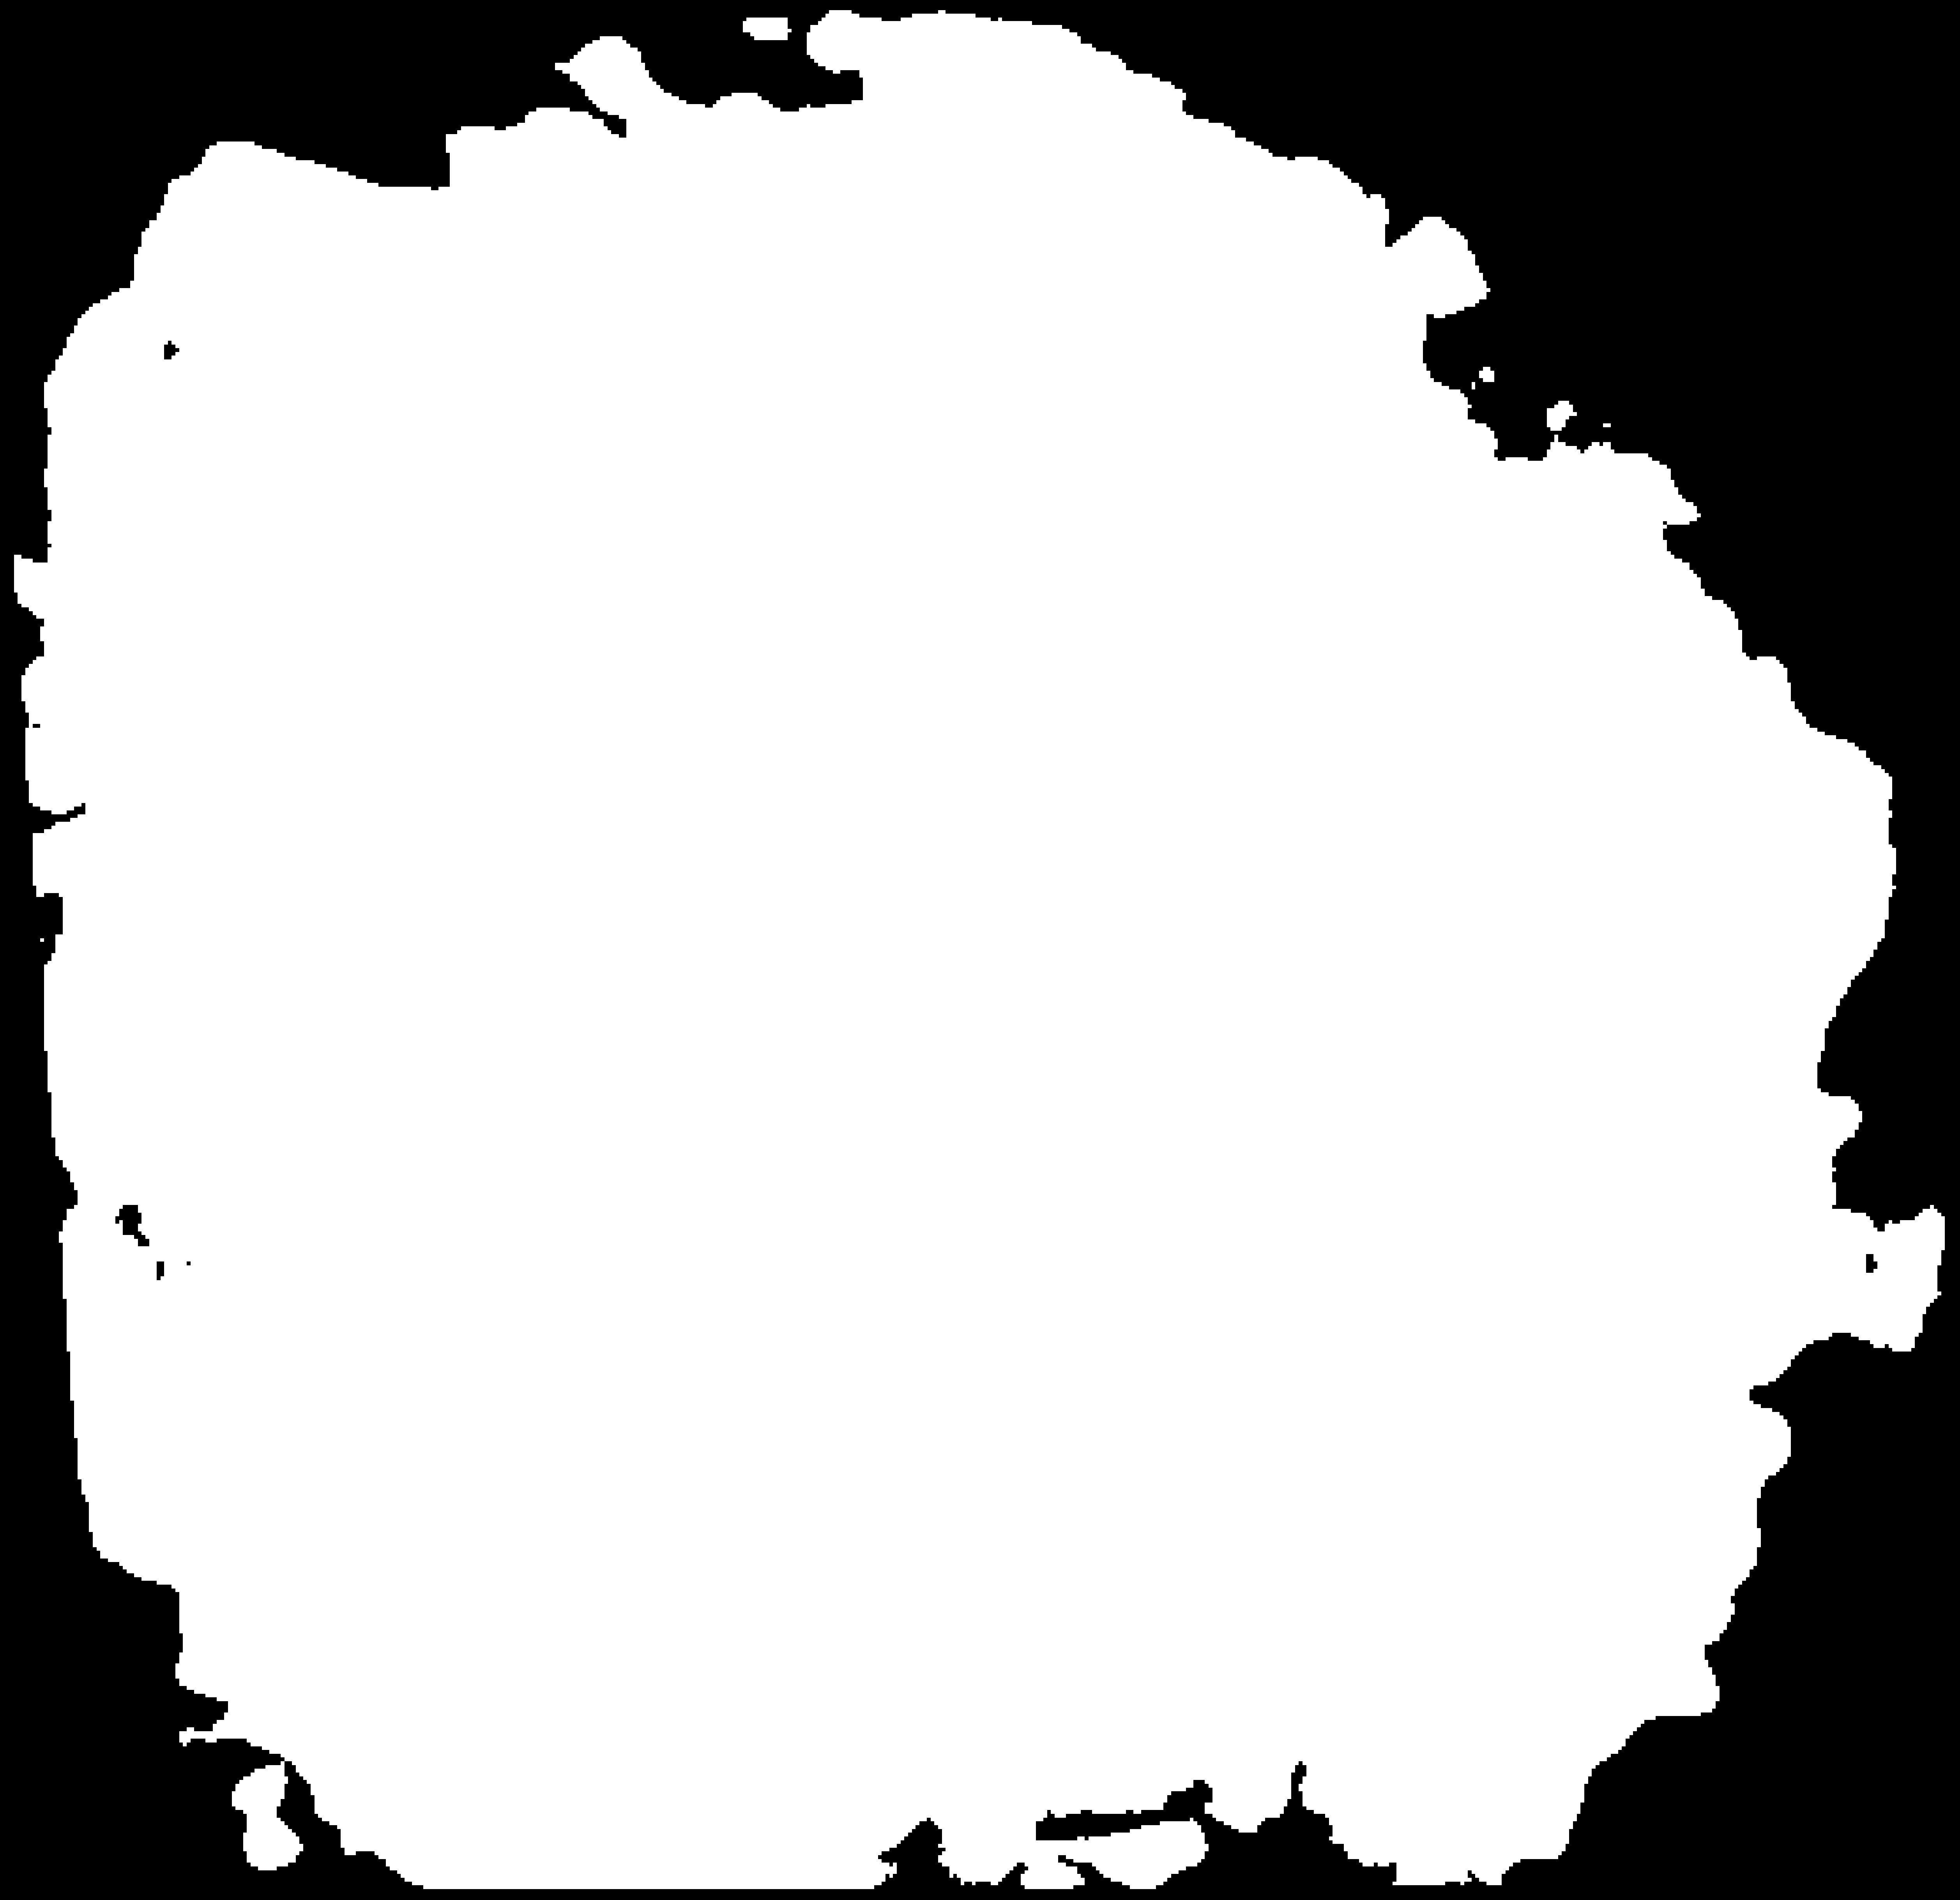

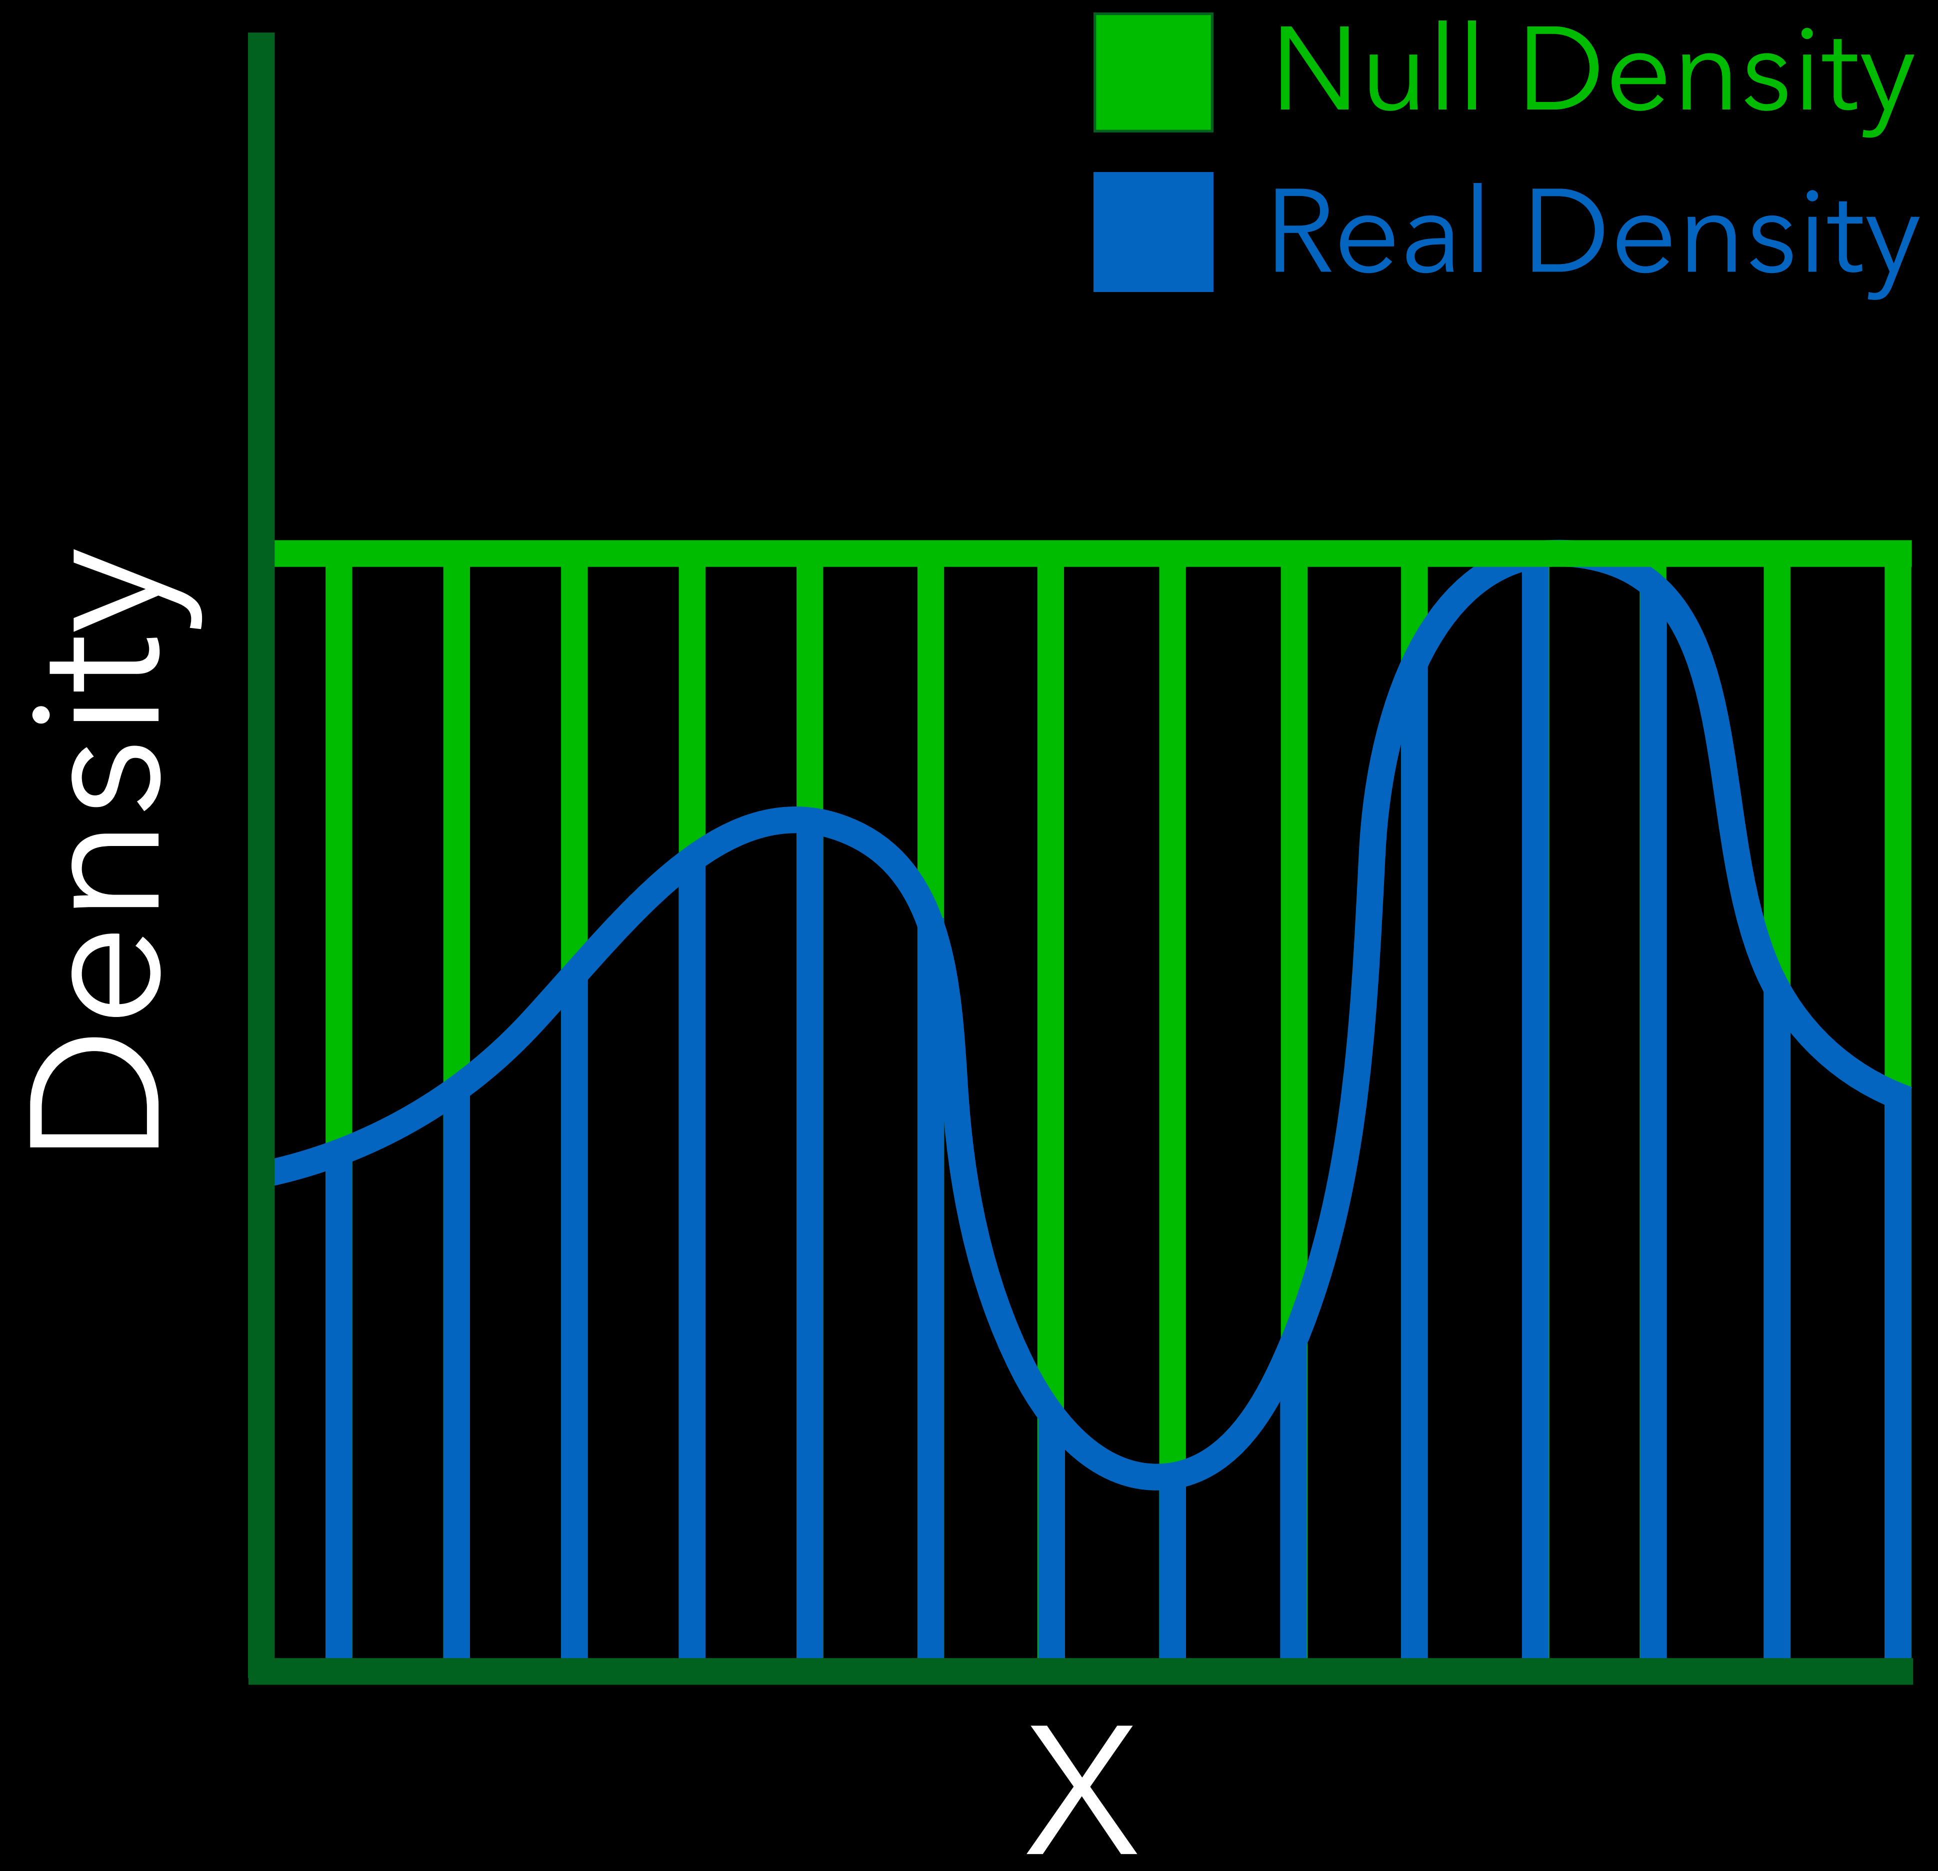
Null Density
Real Density
Density
X

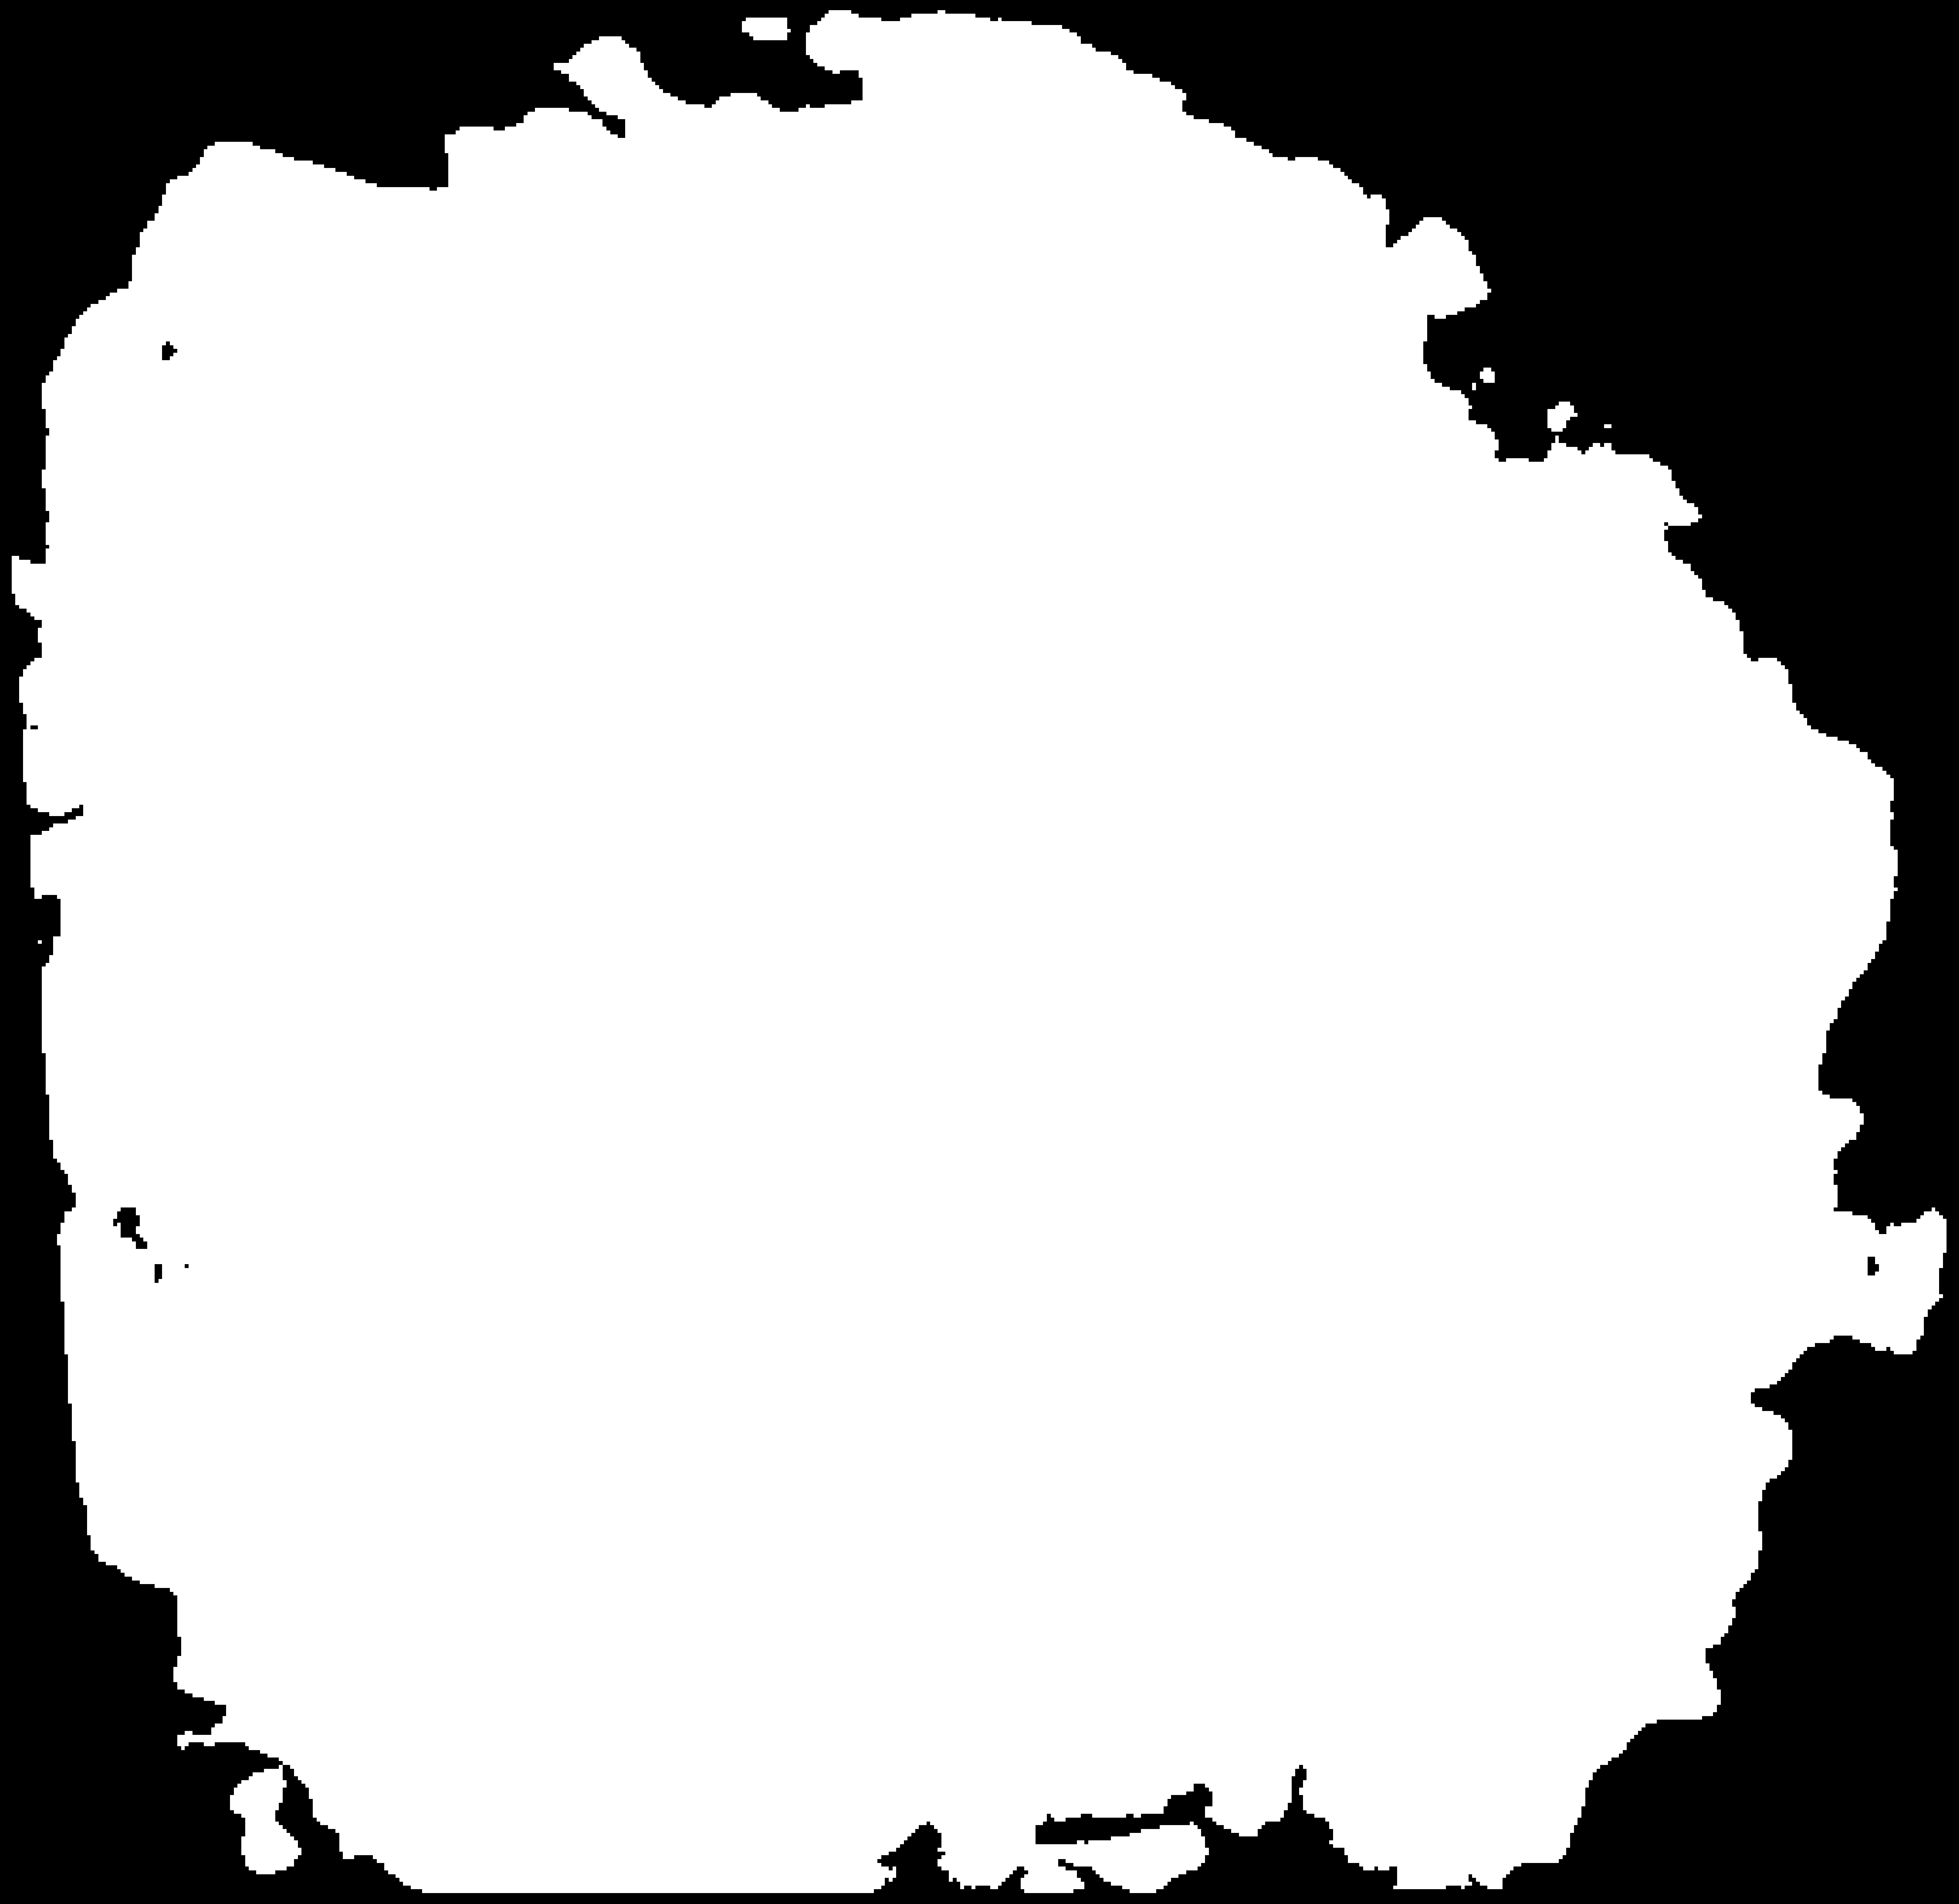

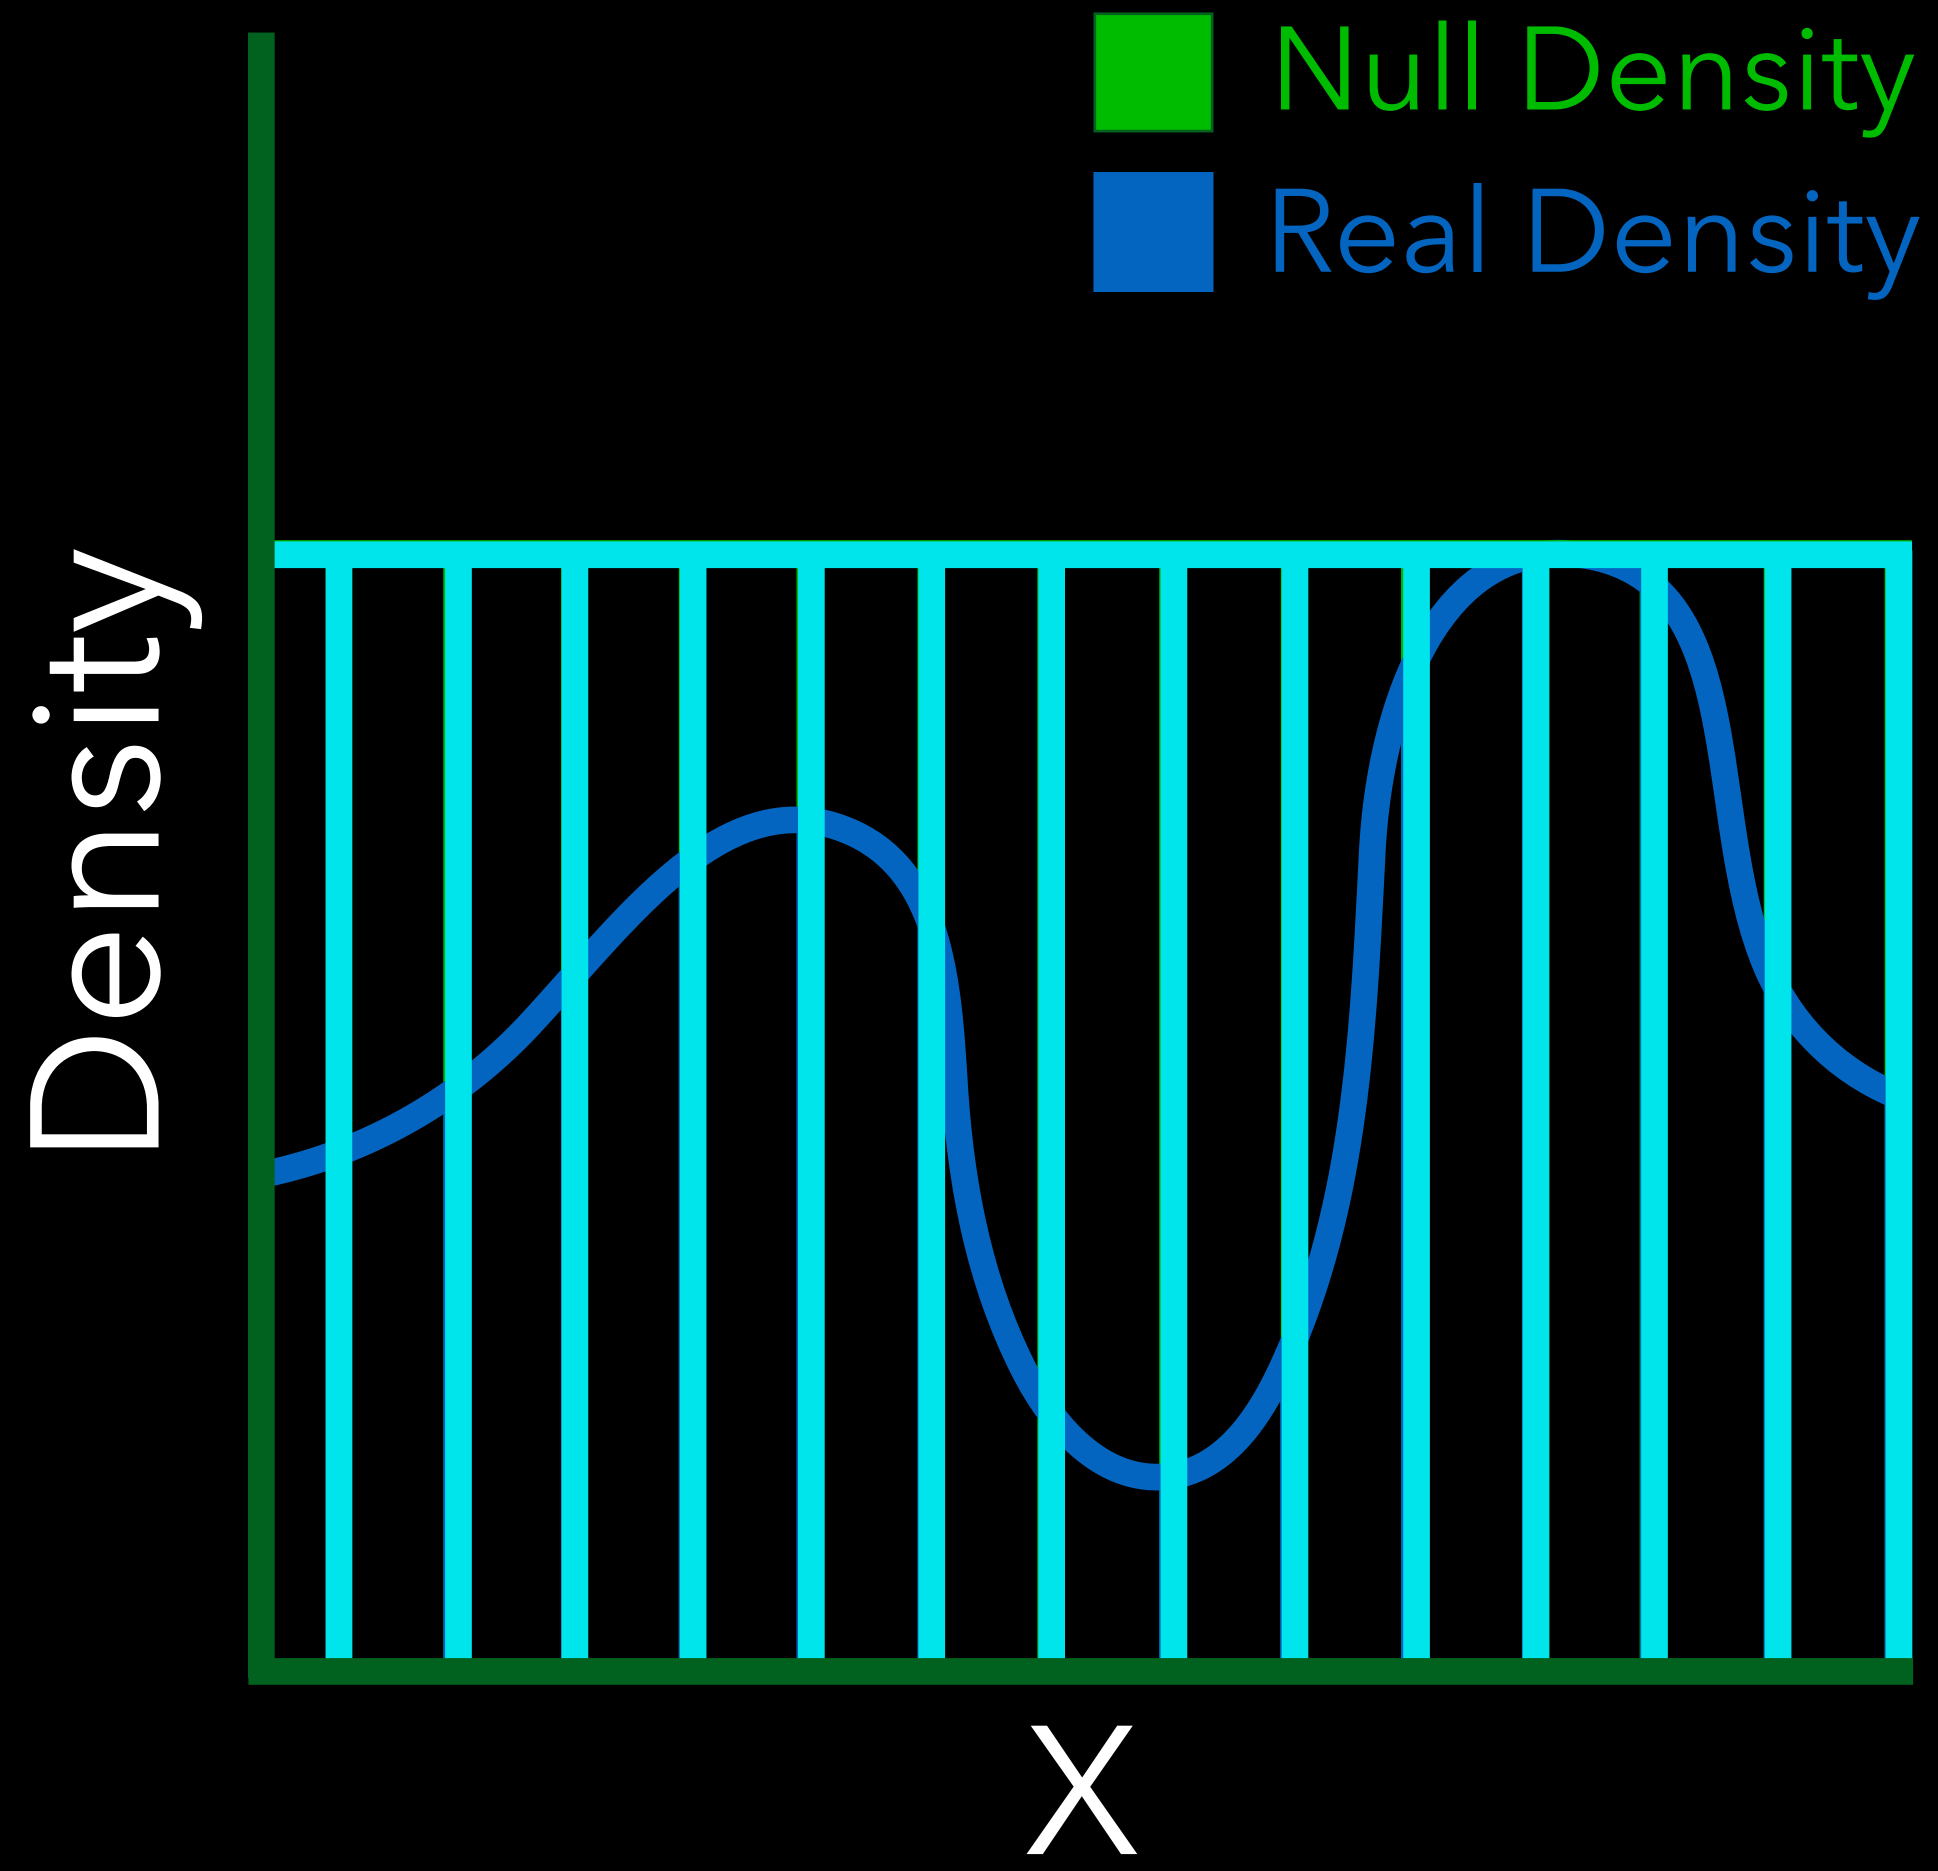

Null Density
Real Density
Density
X

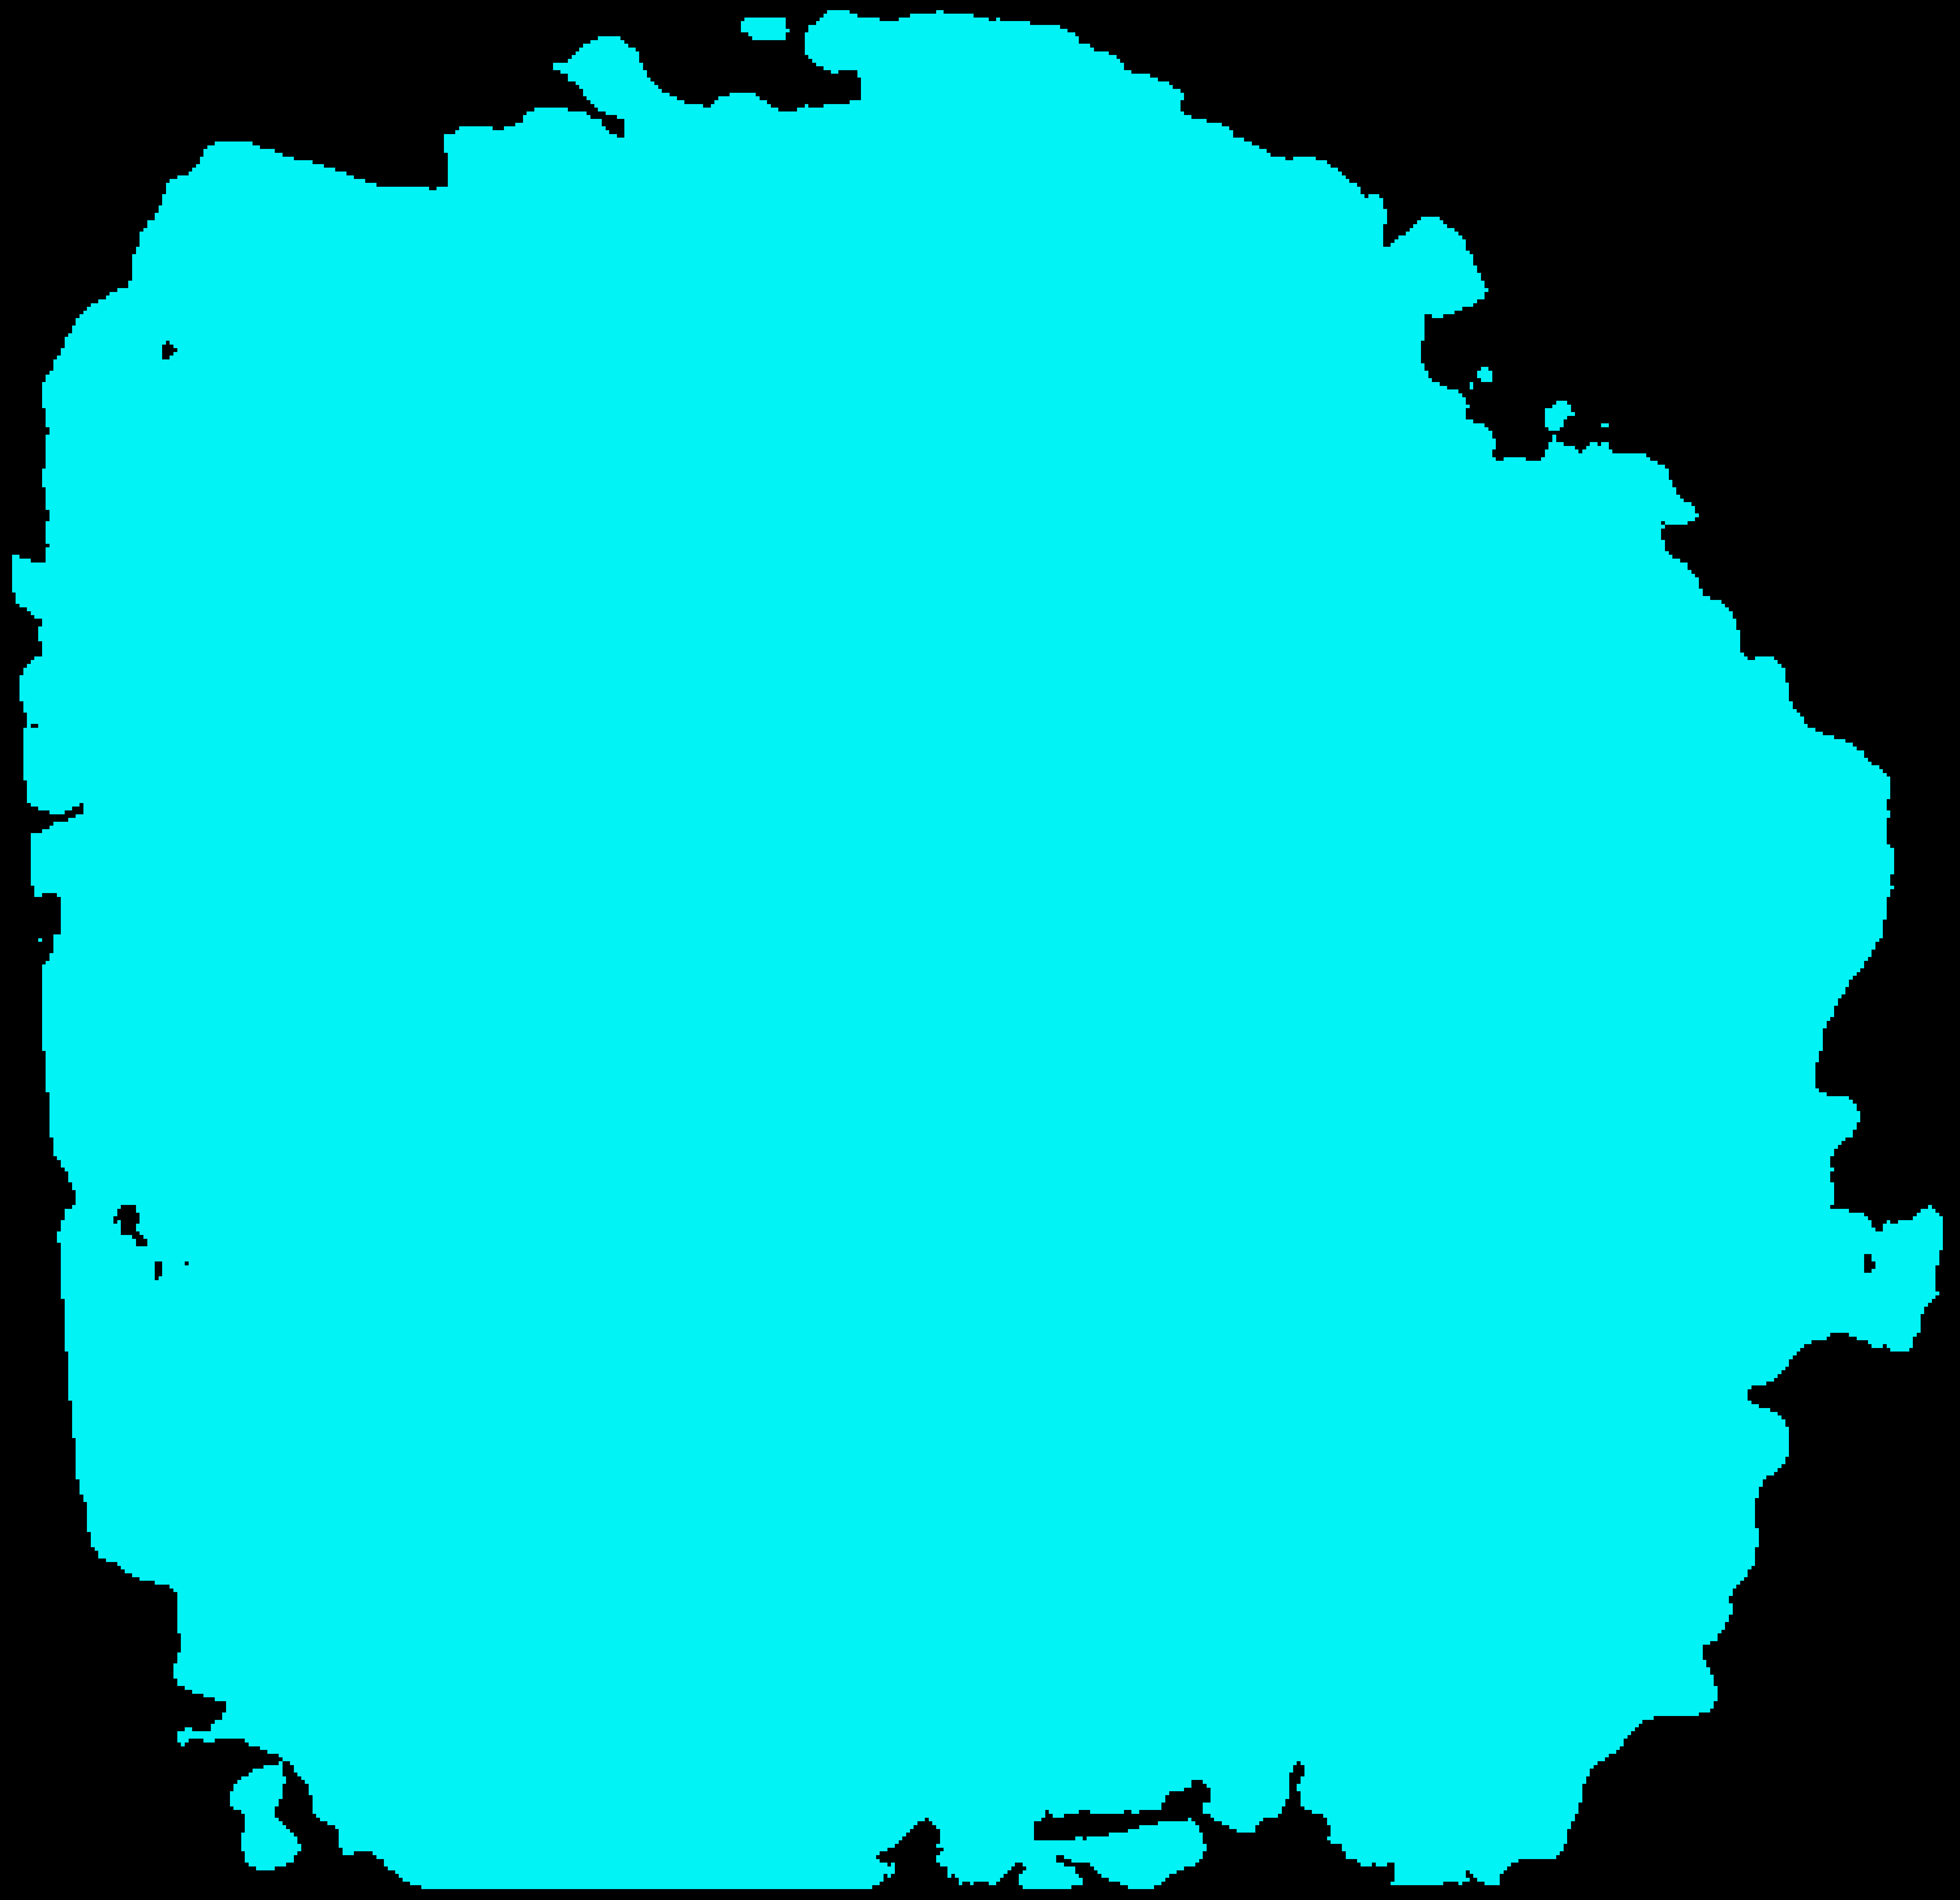

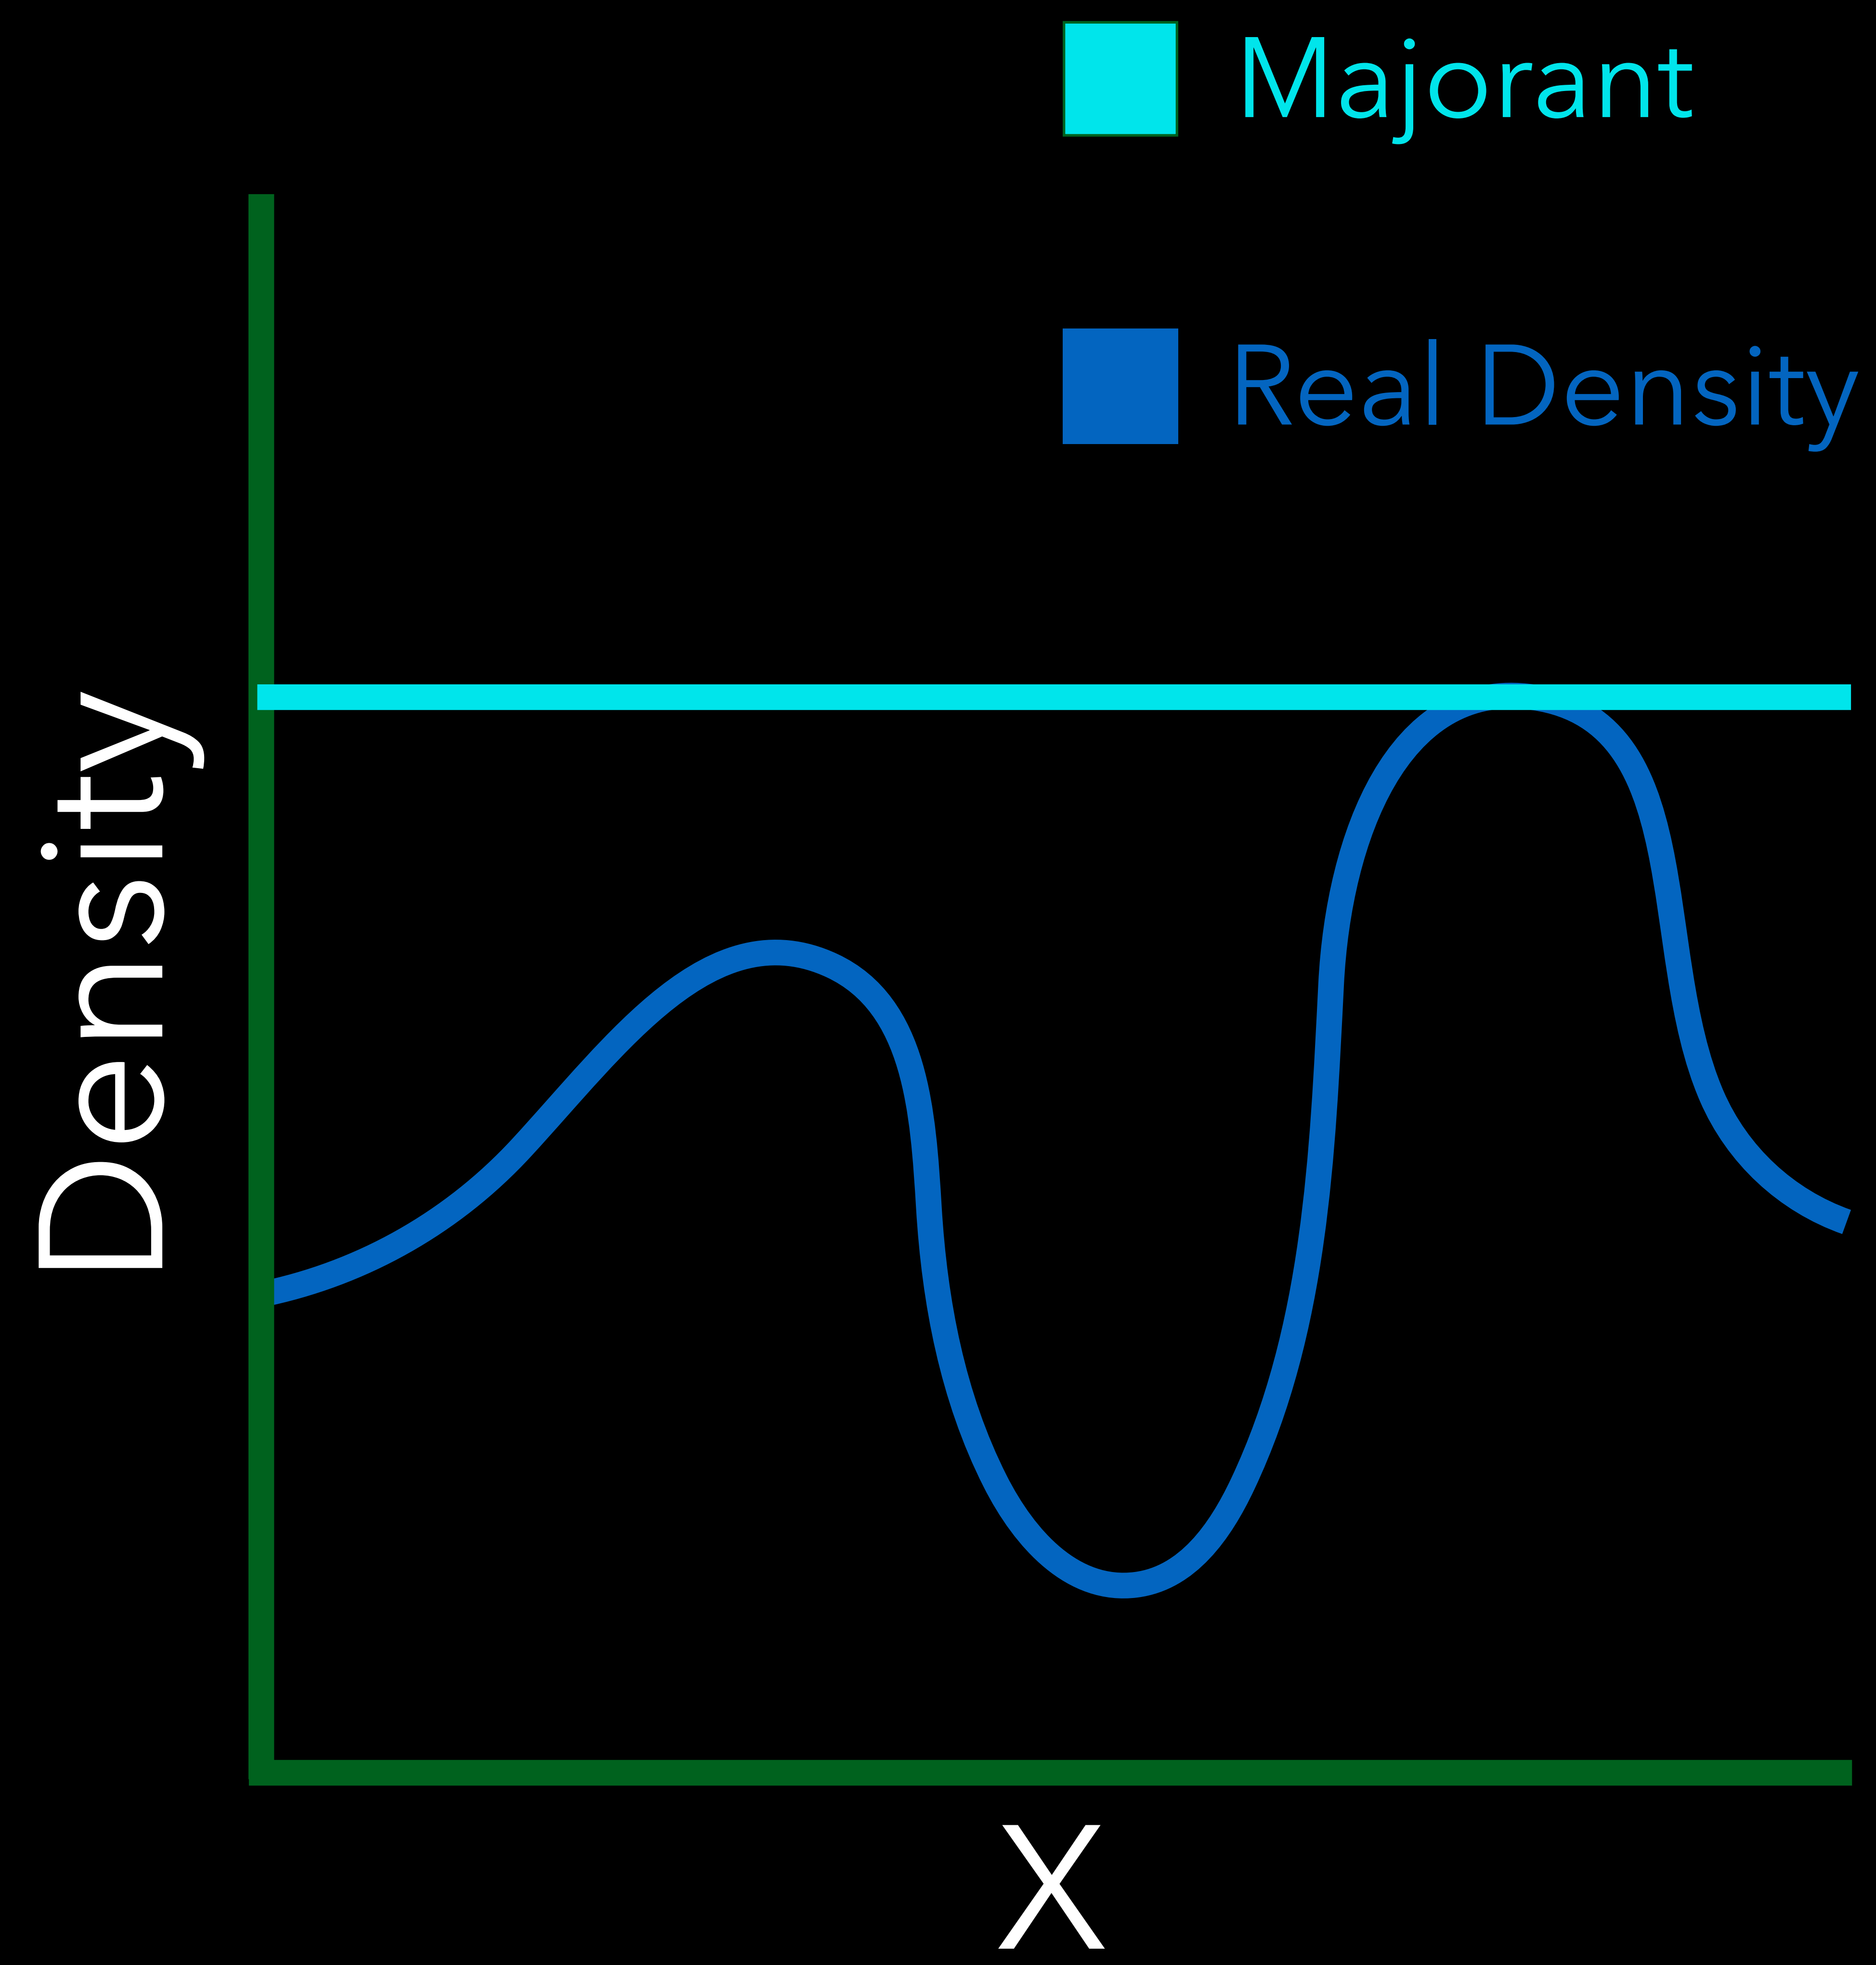
Majorant
Real Density
Density
X

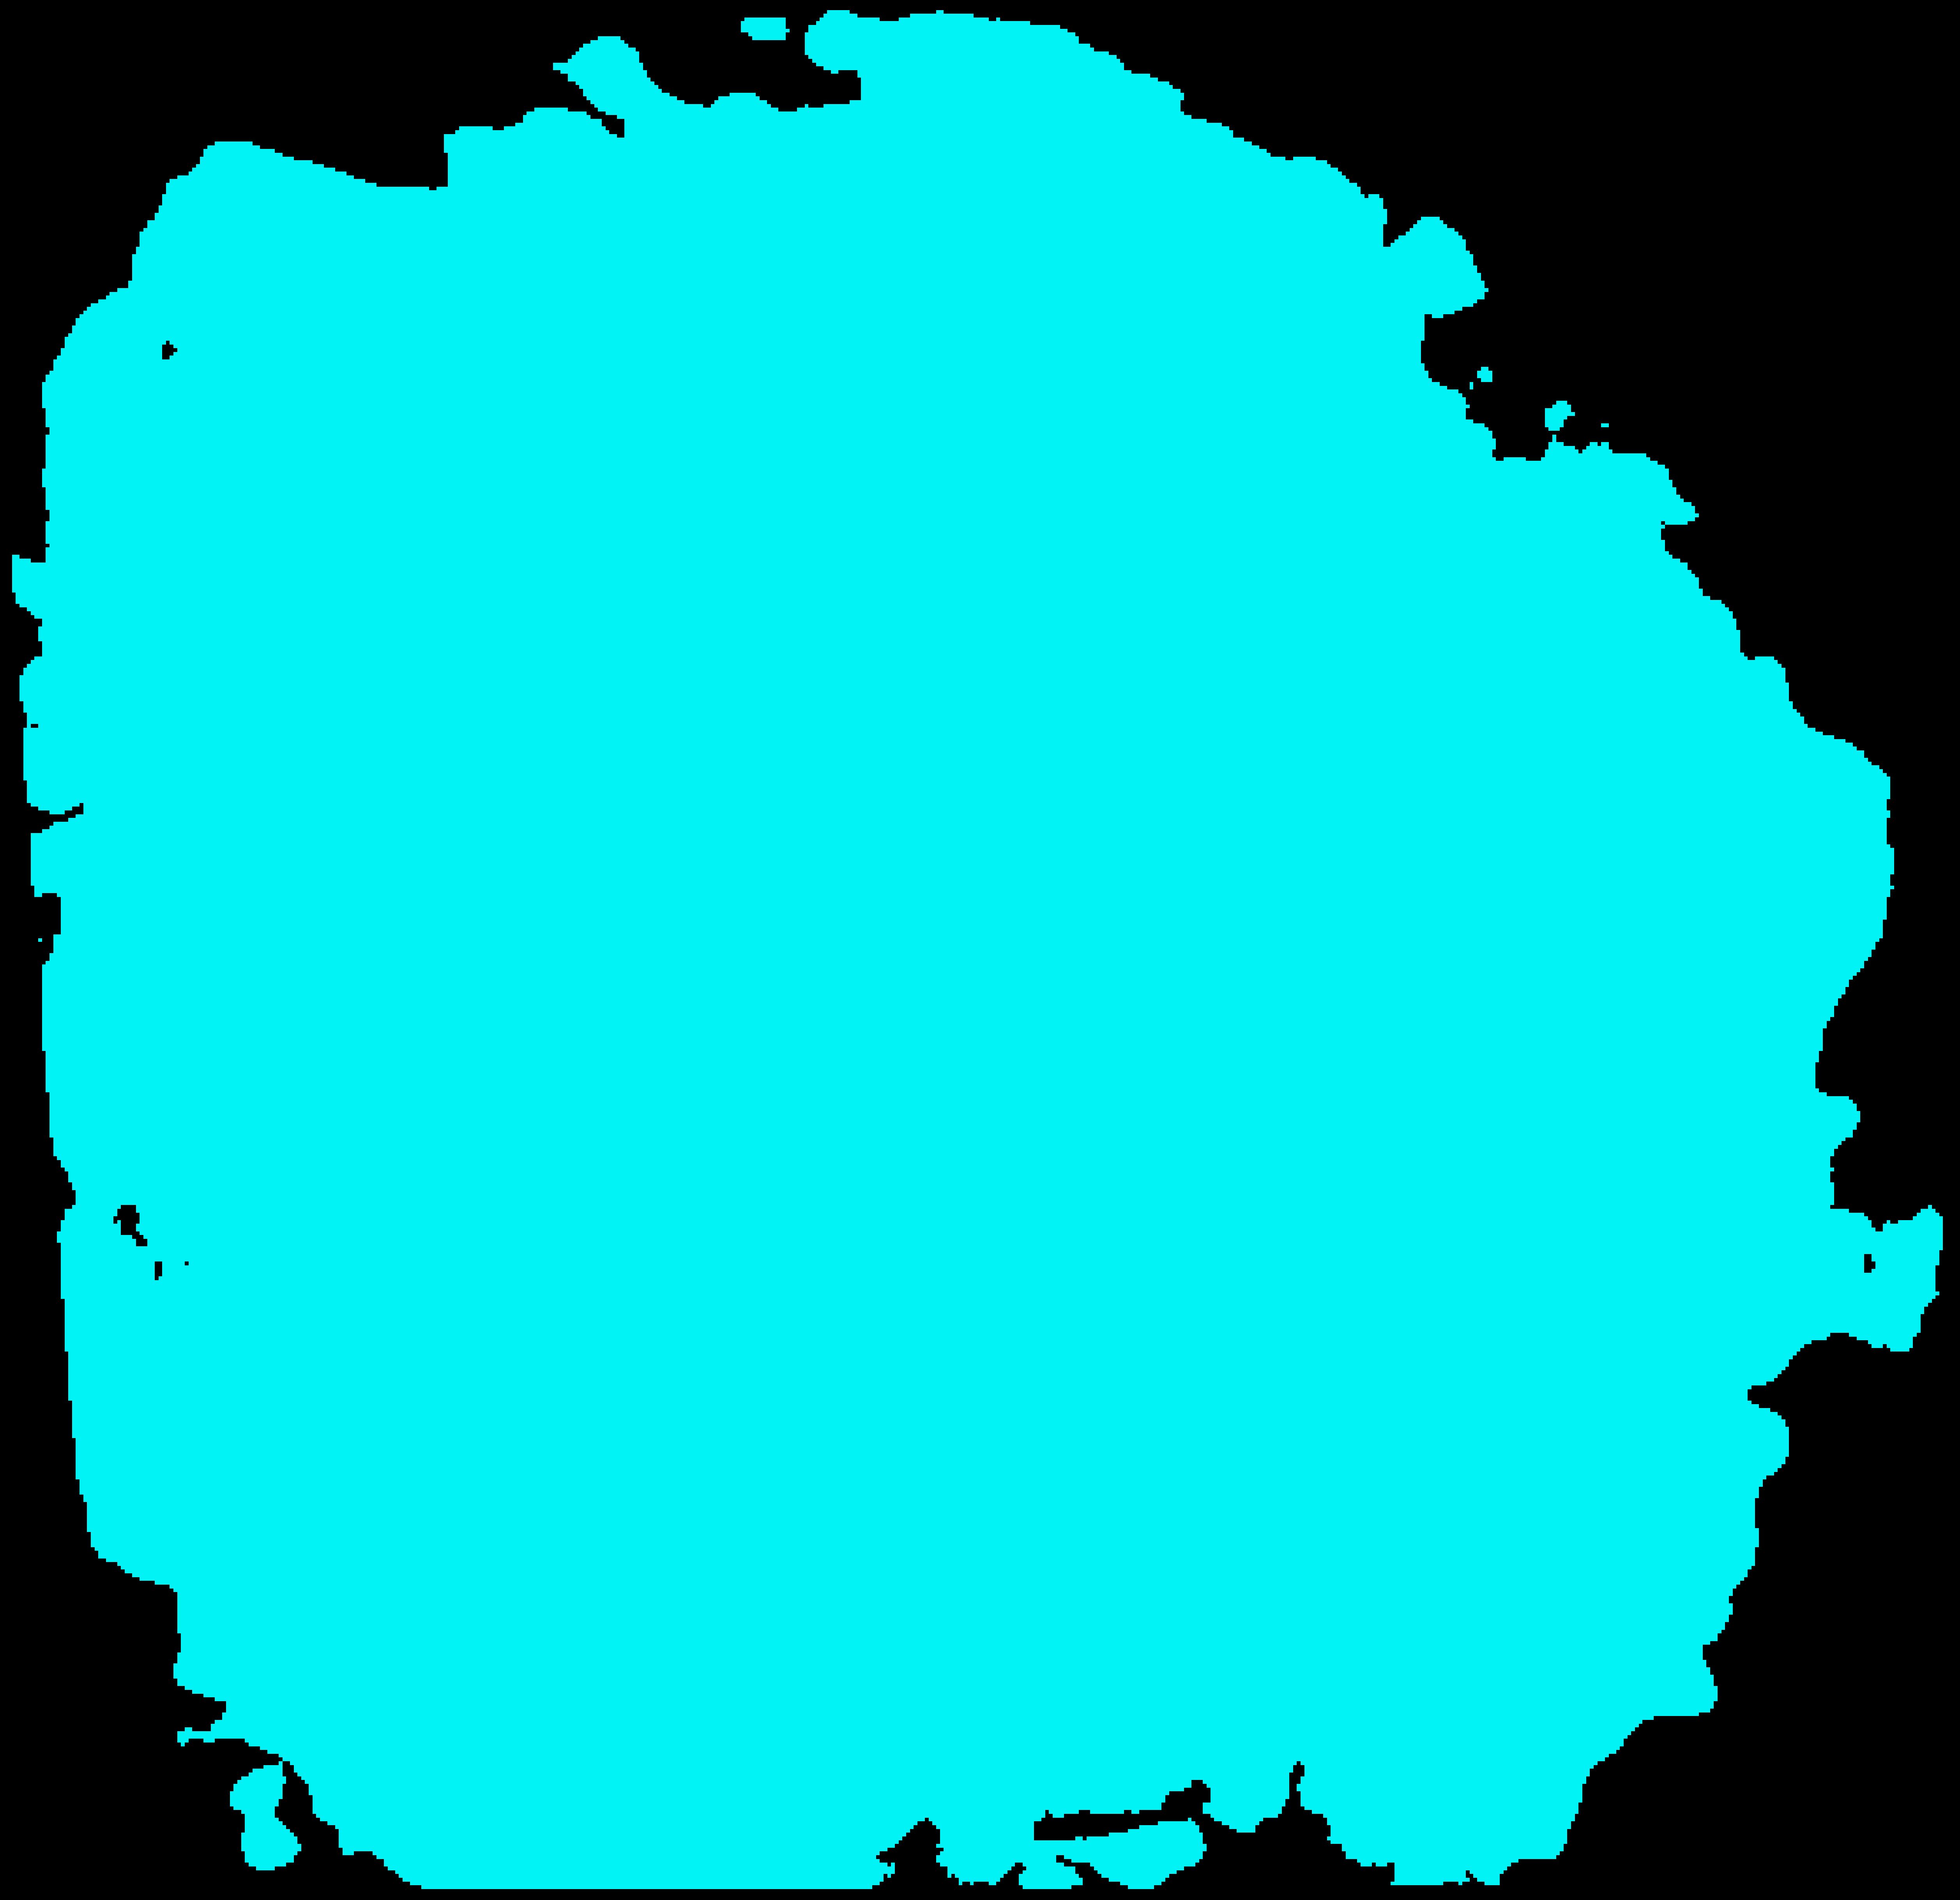

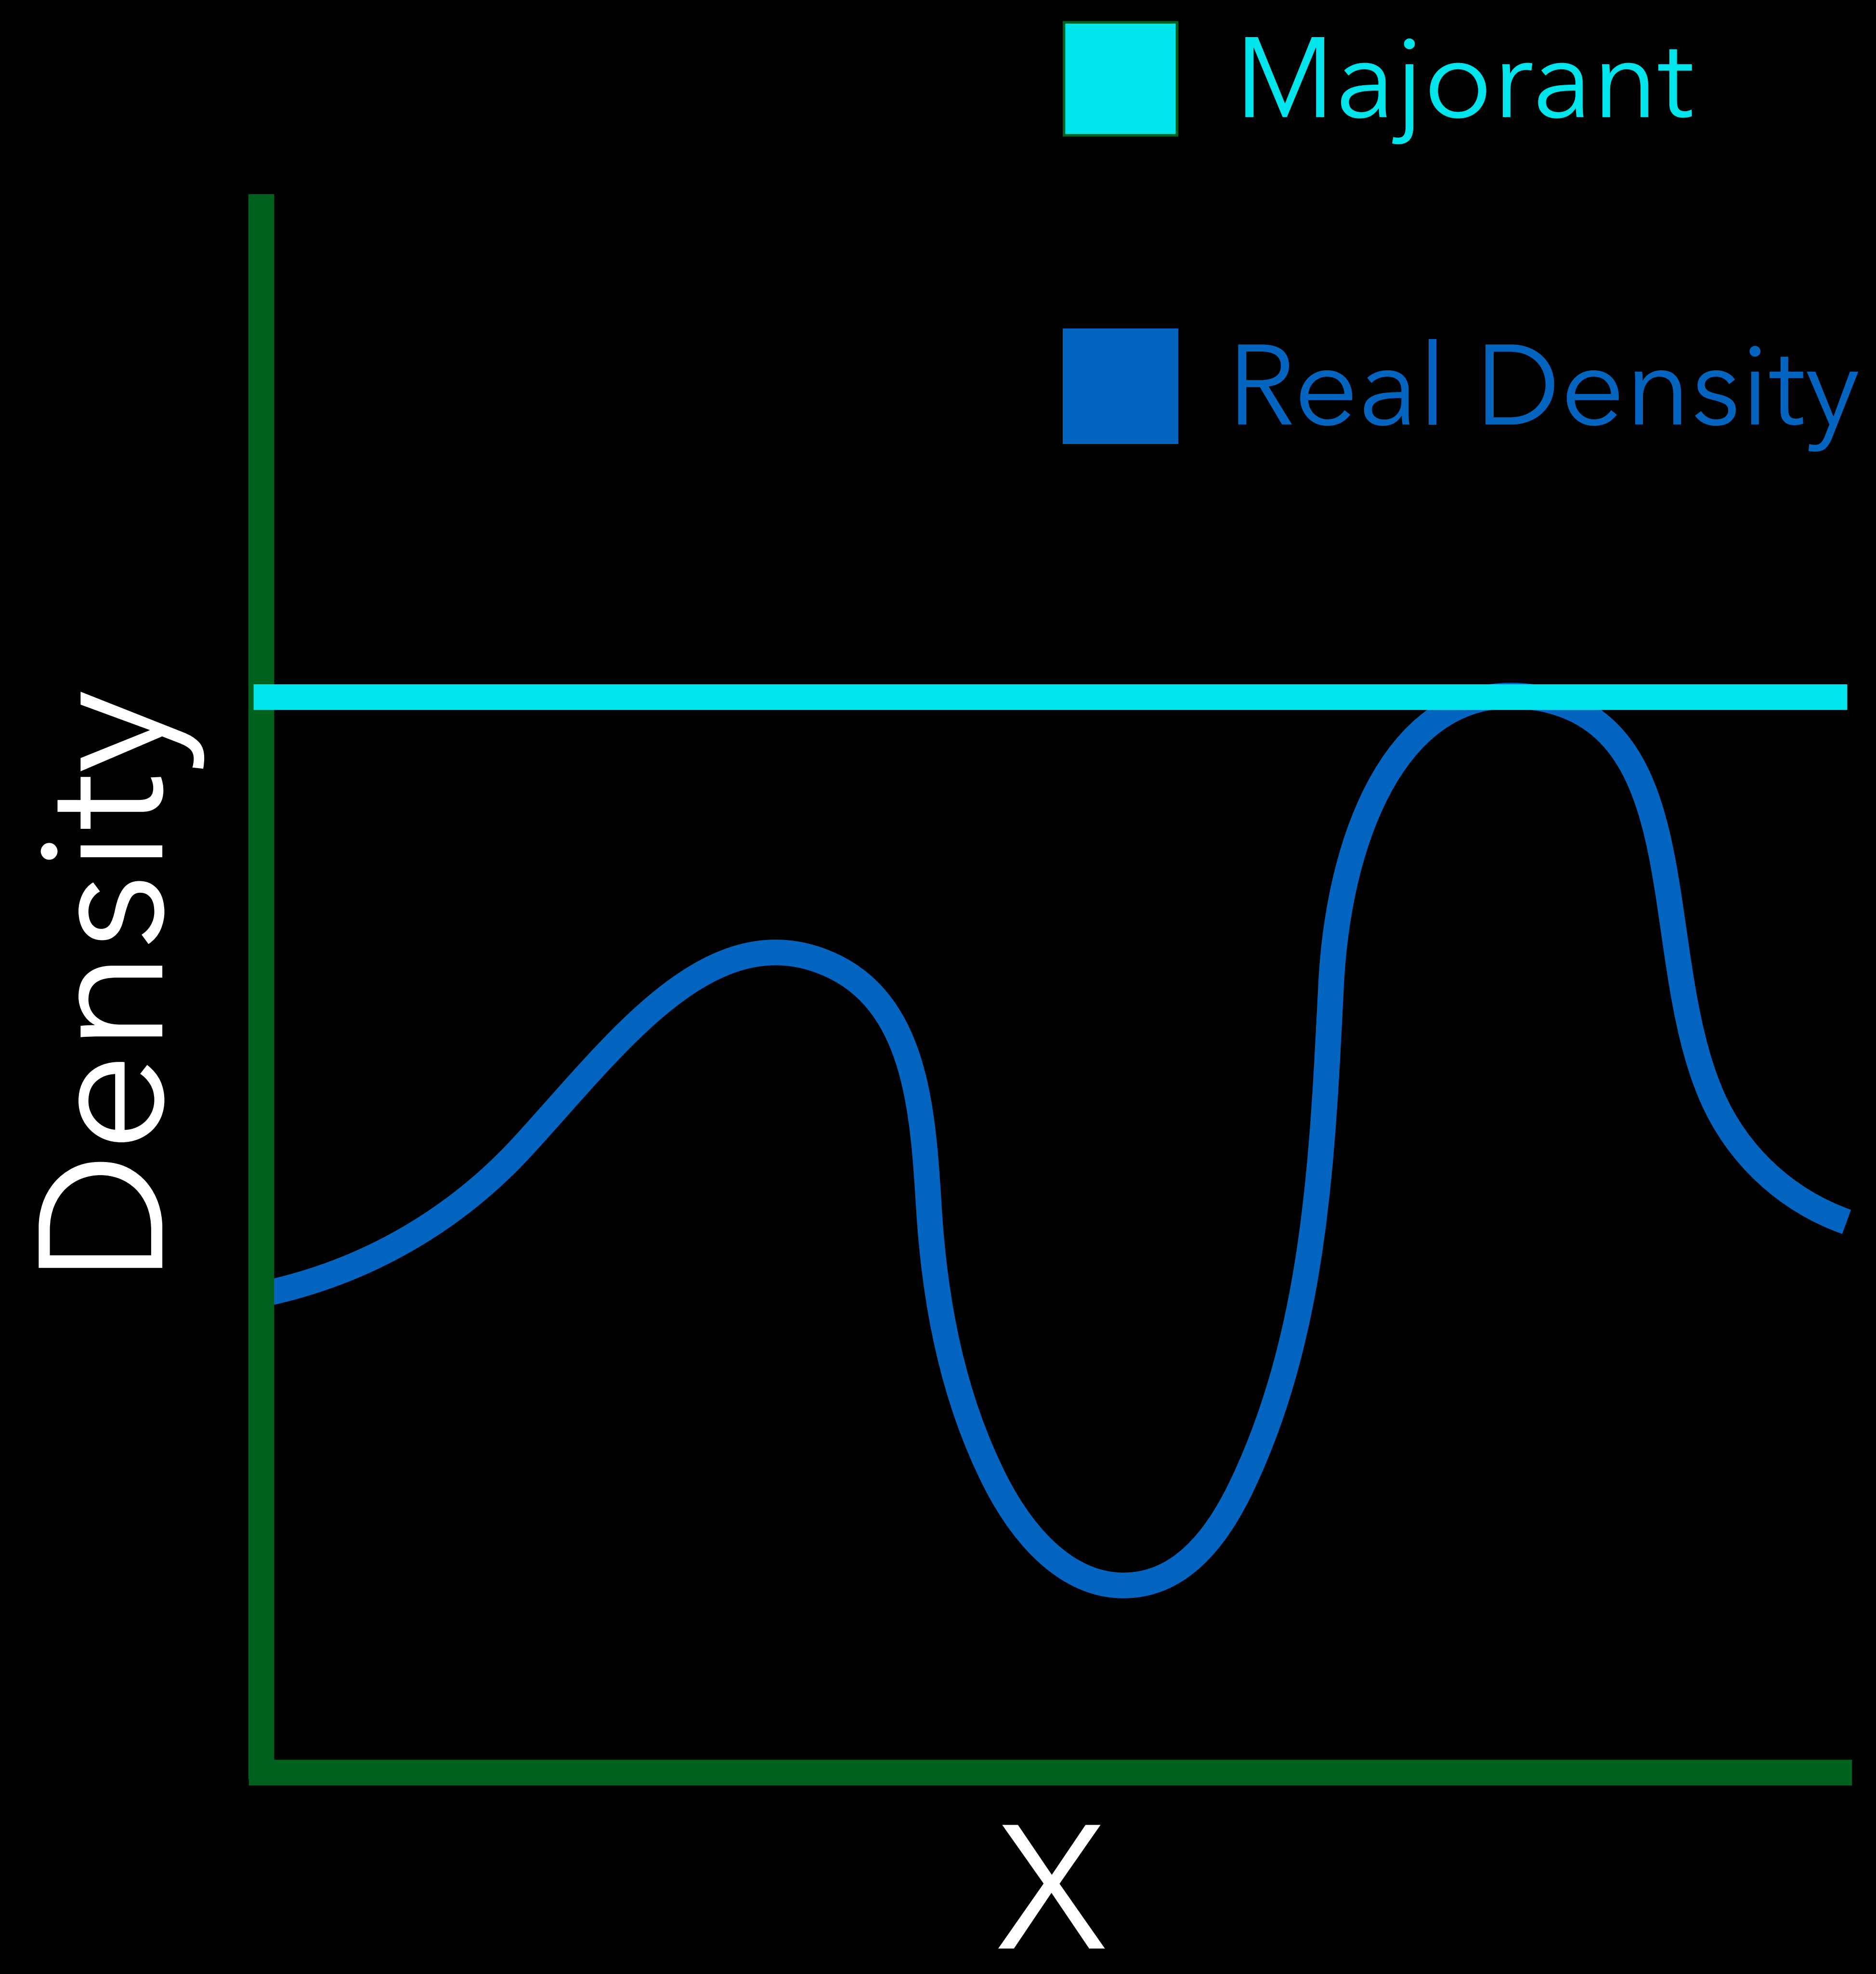

Majorant
Real Density
Density
X

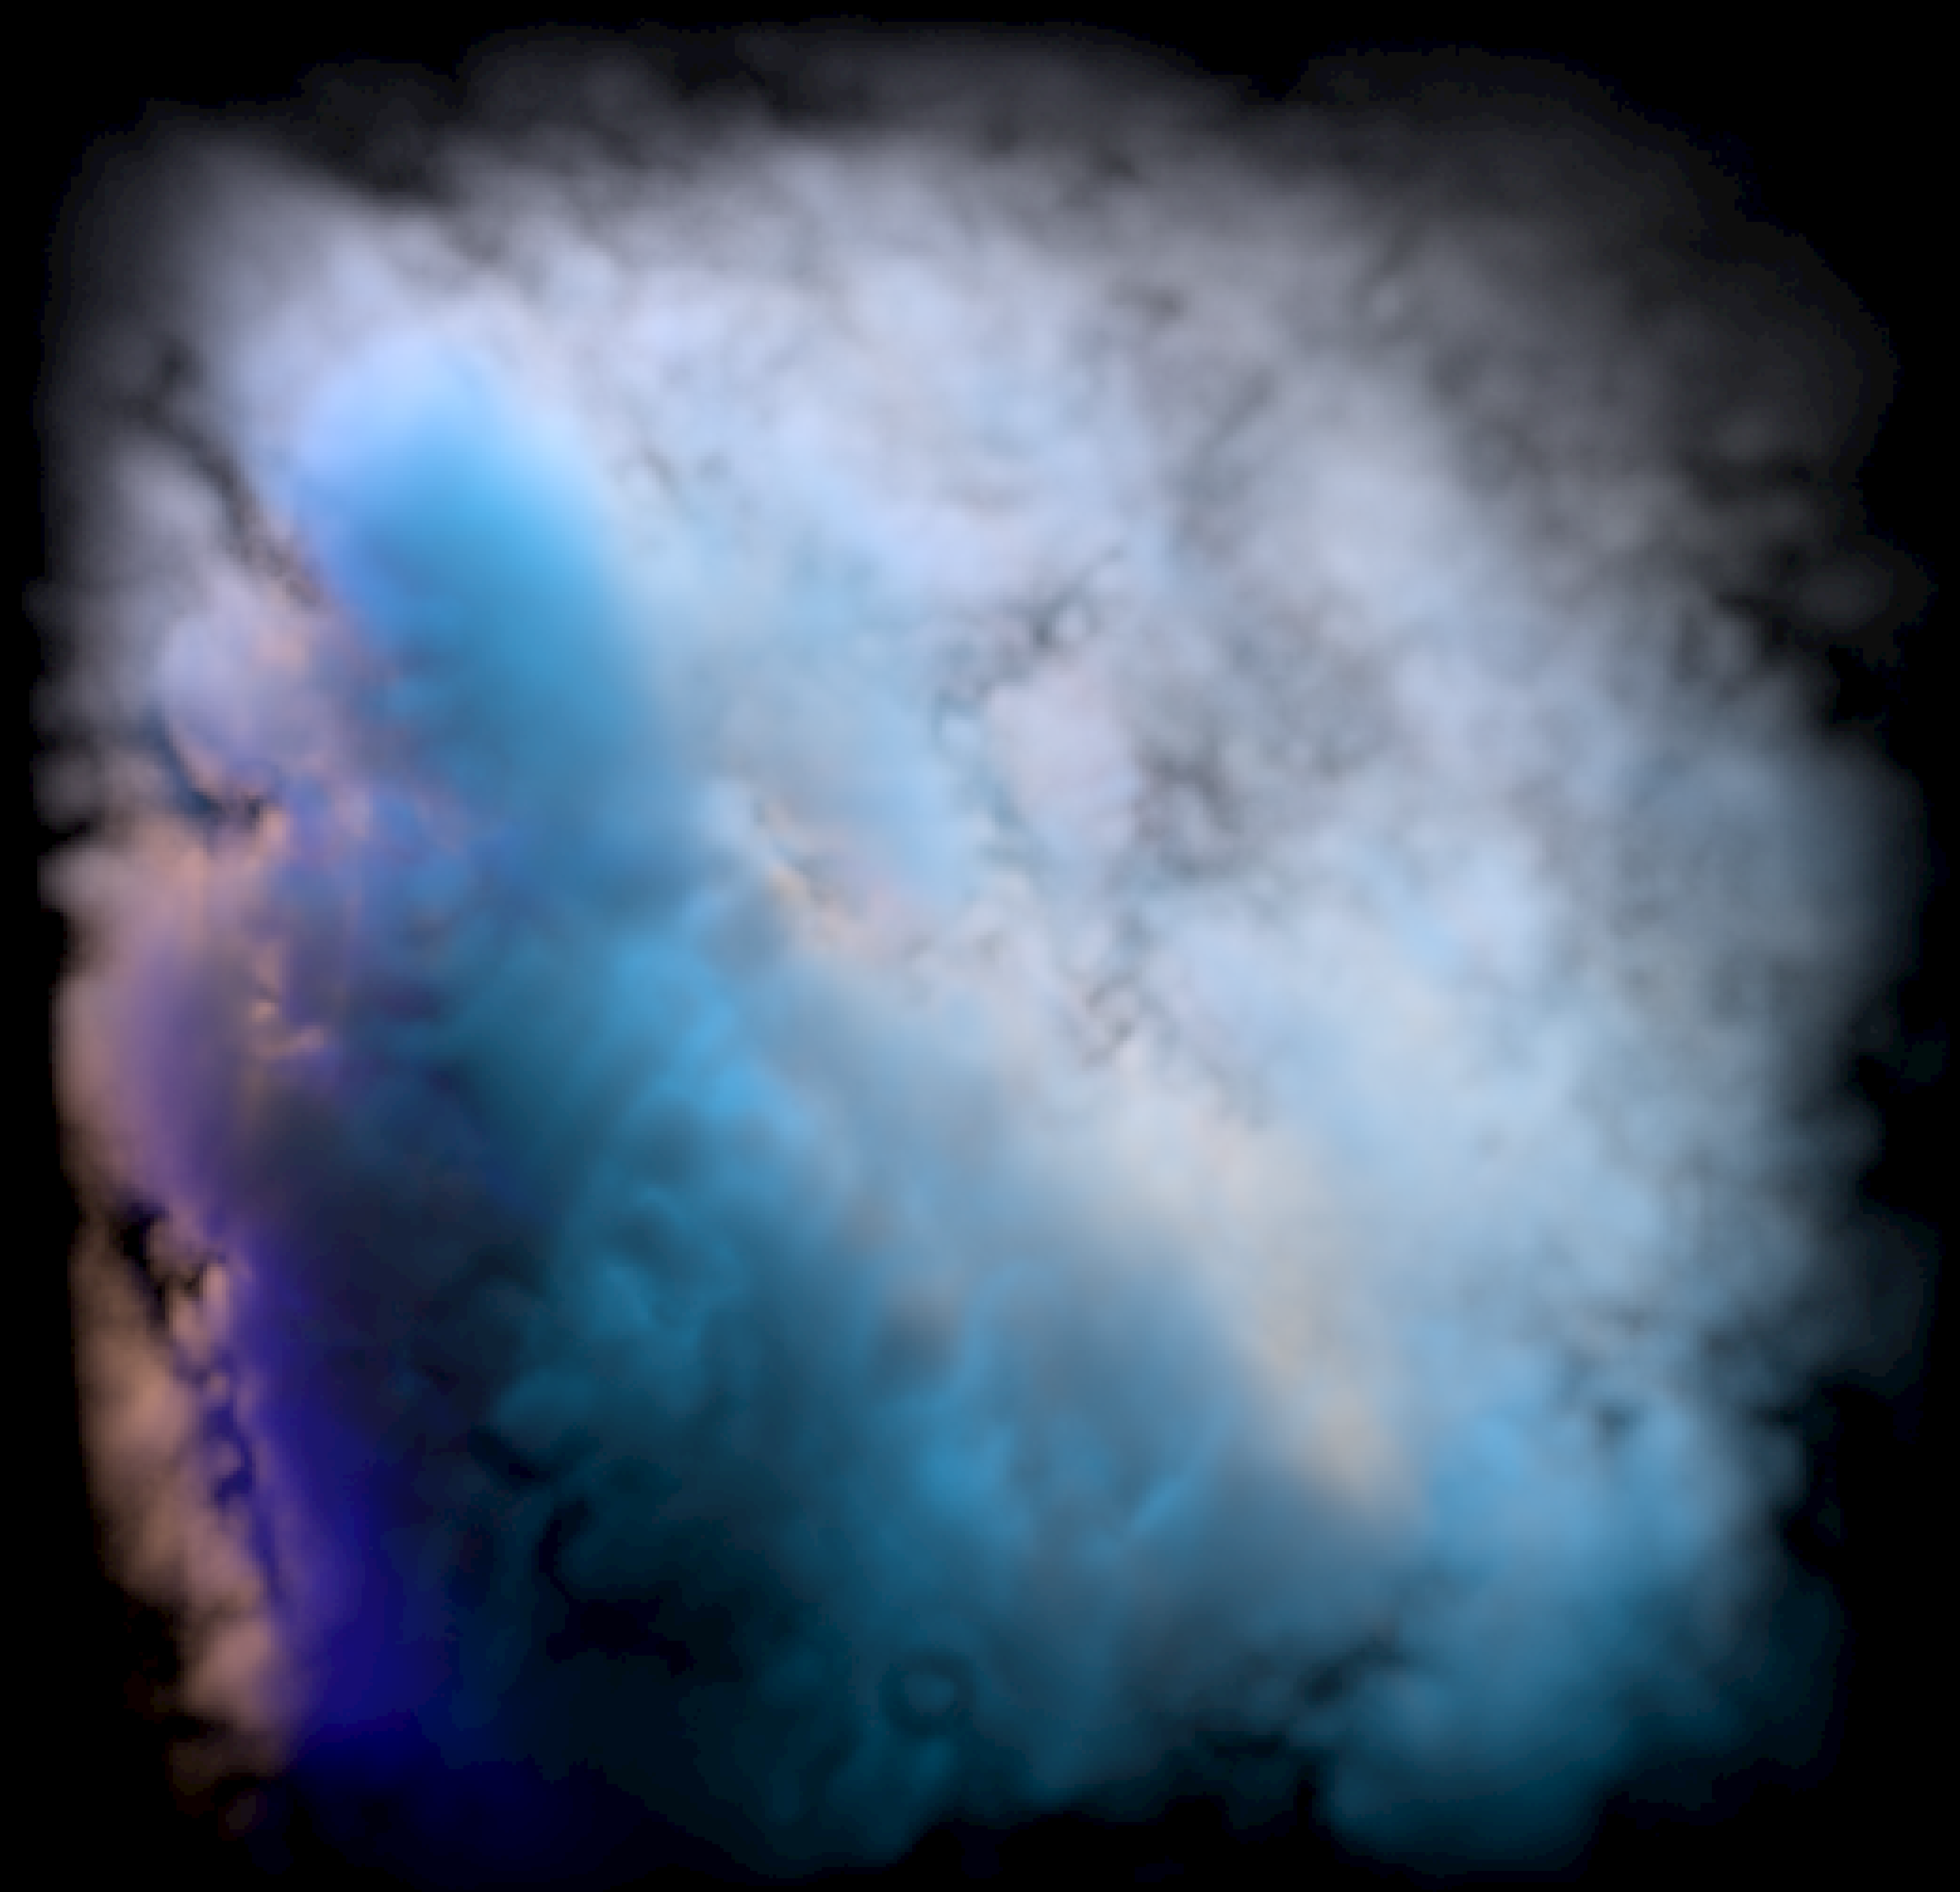

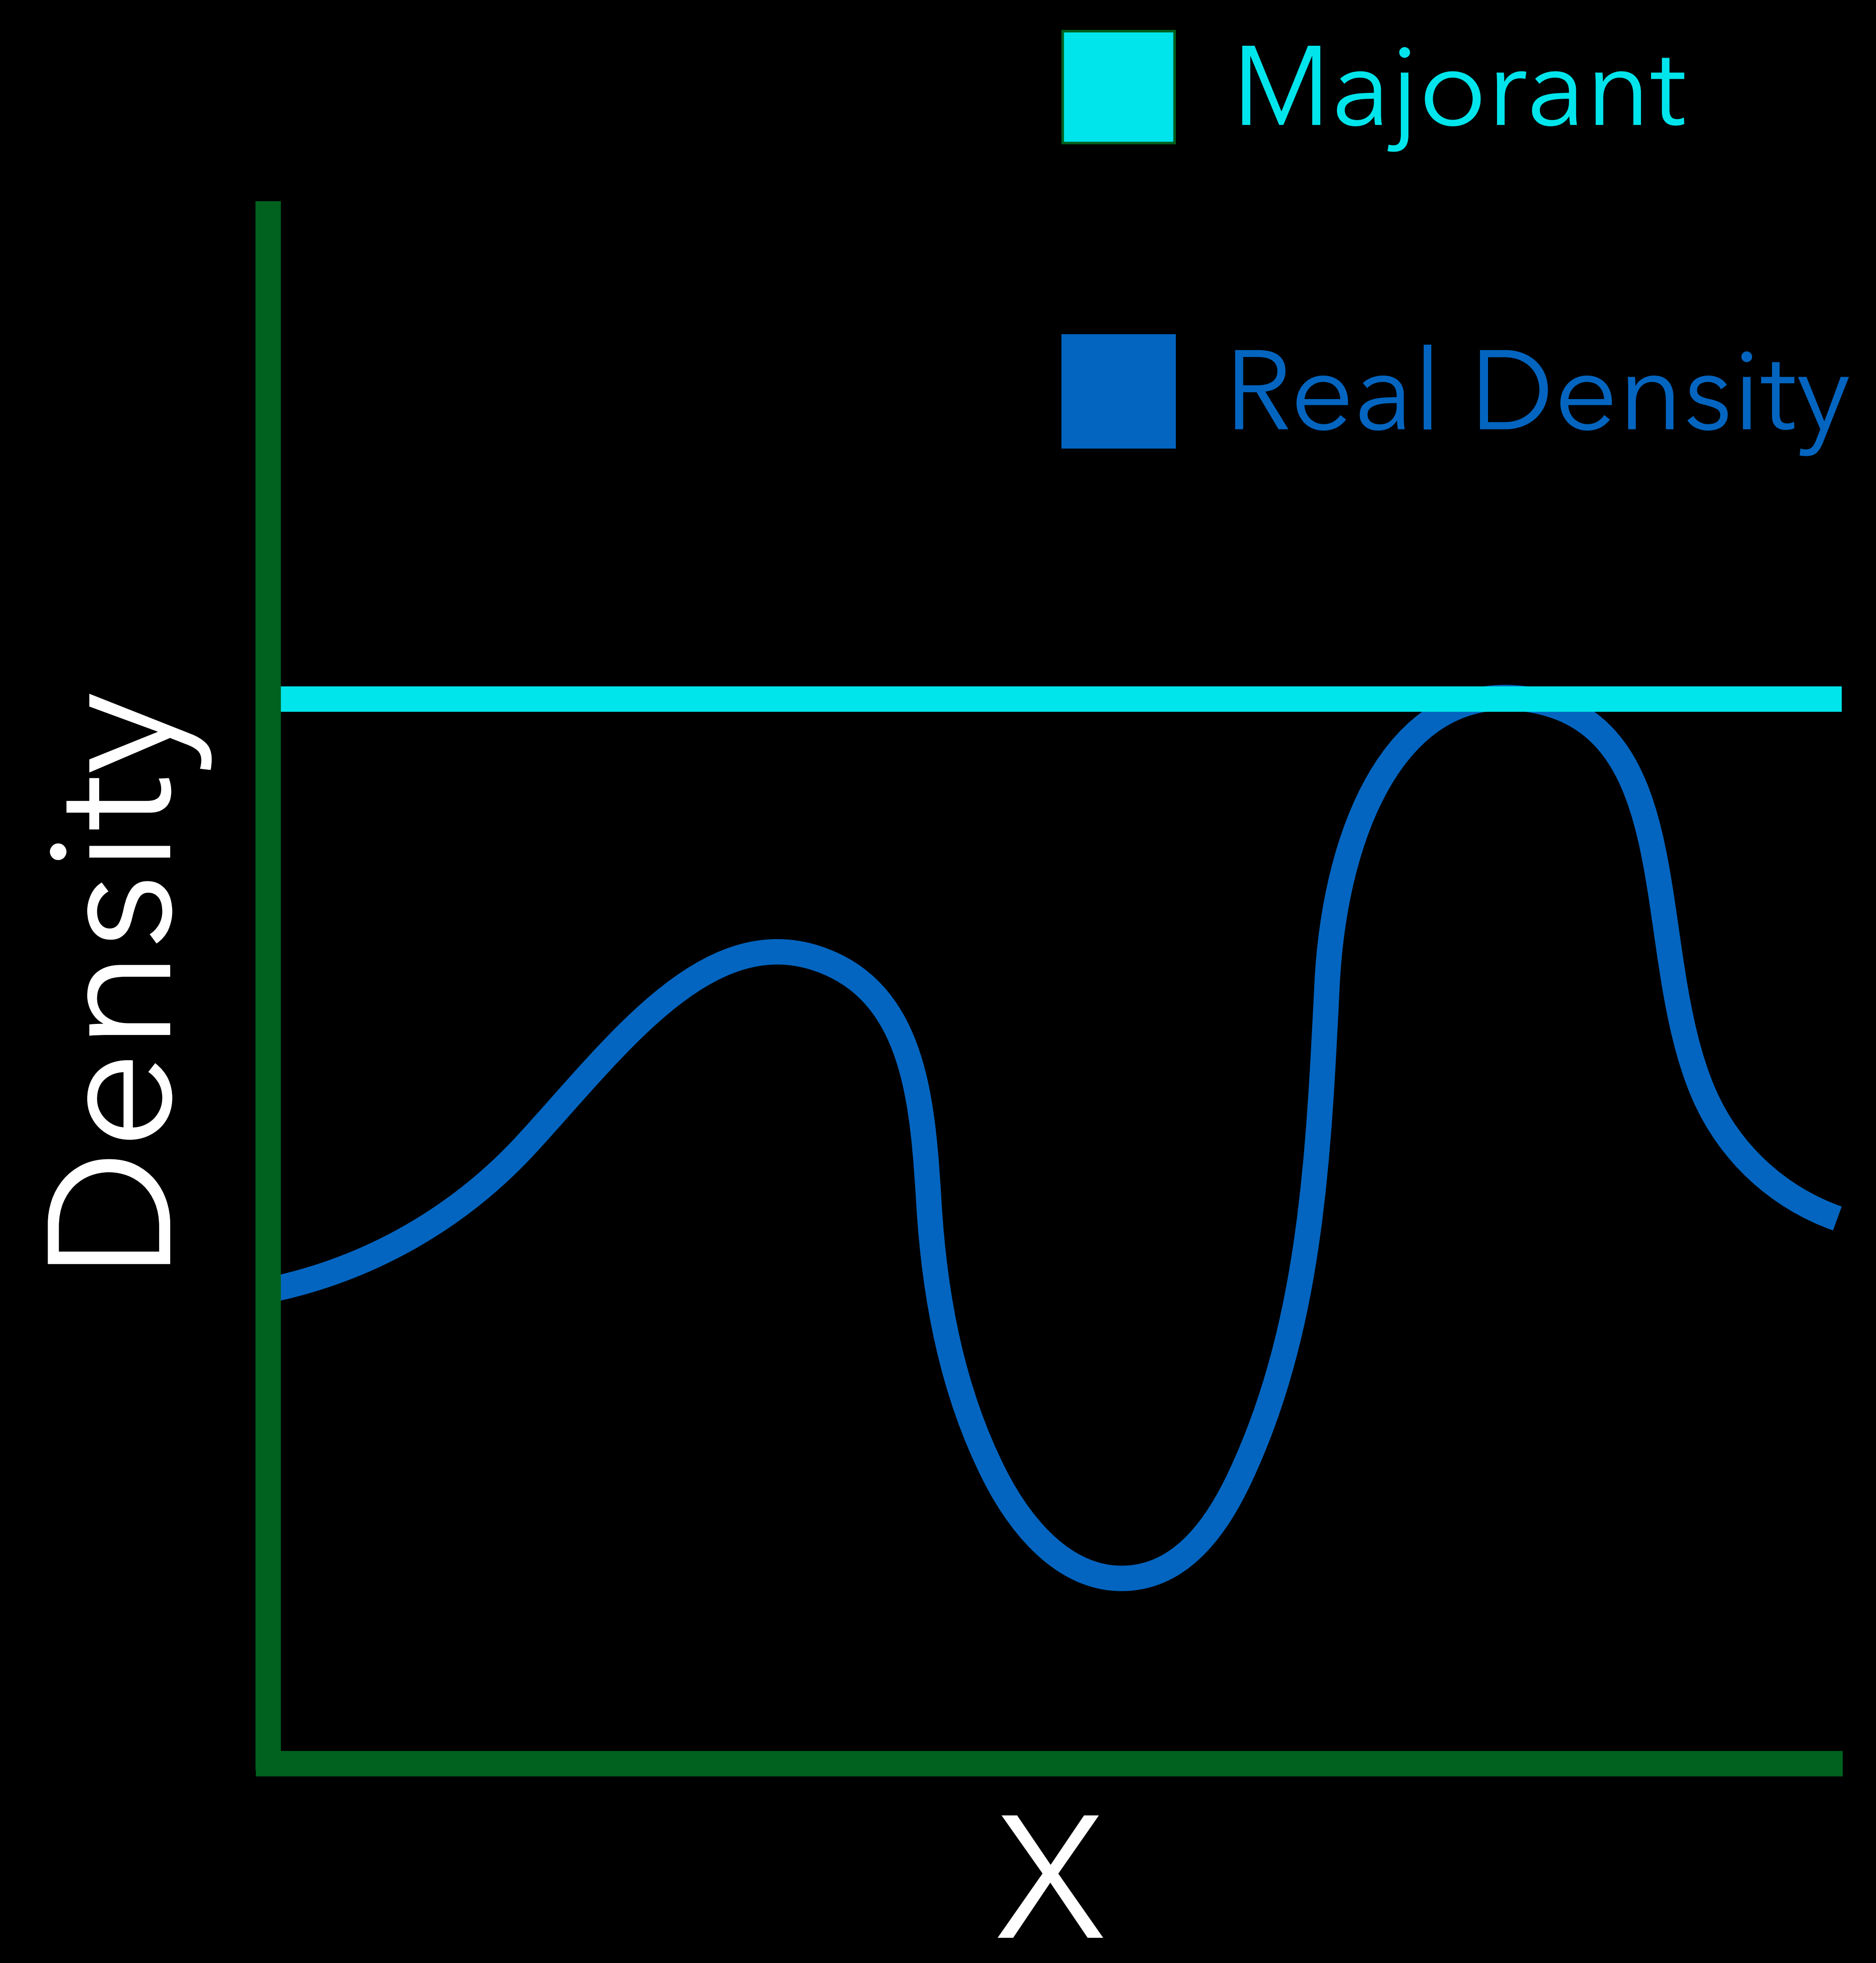
Majorant
Real Density
Density
X

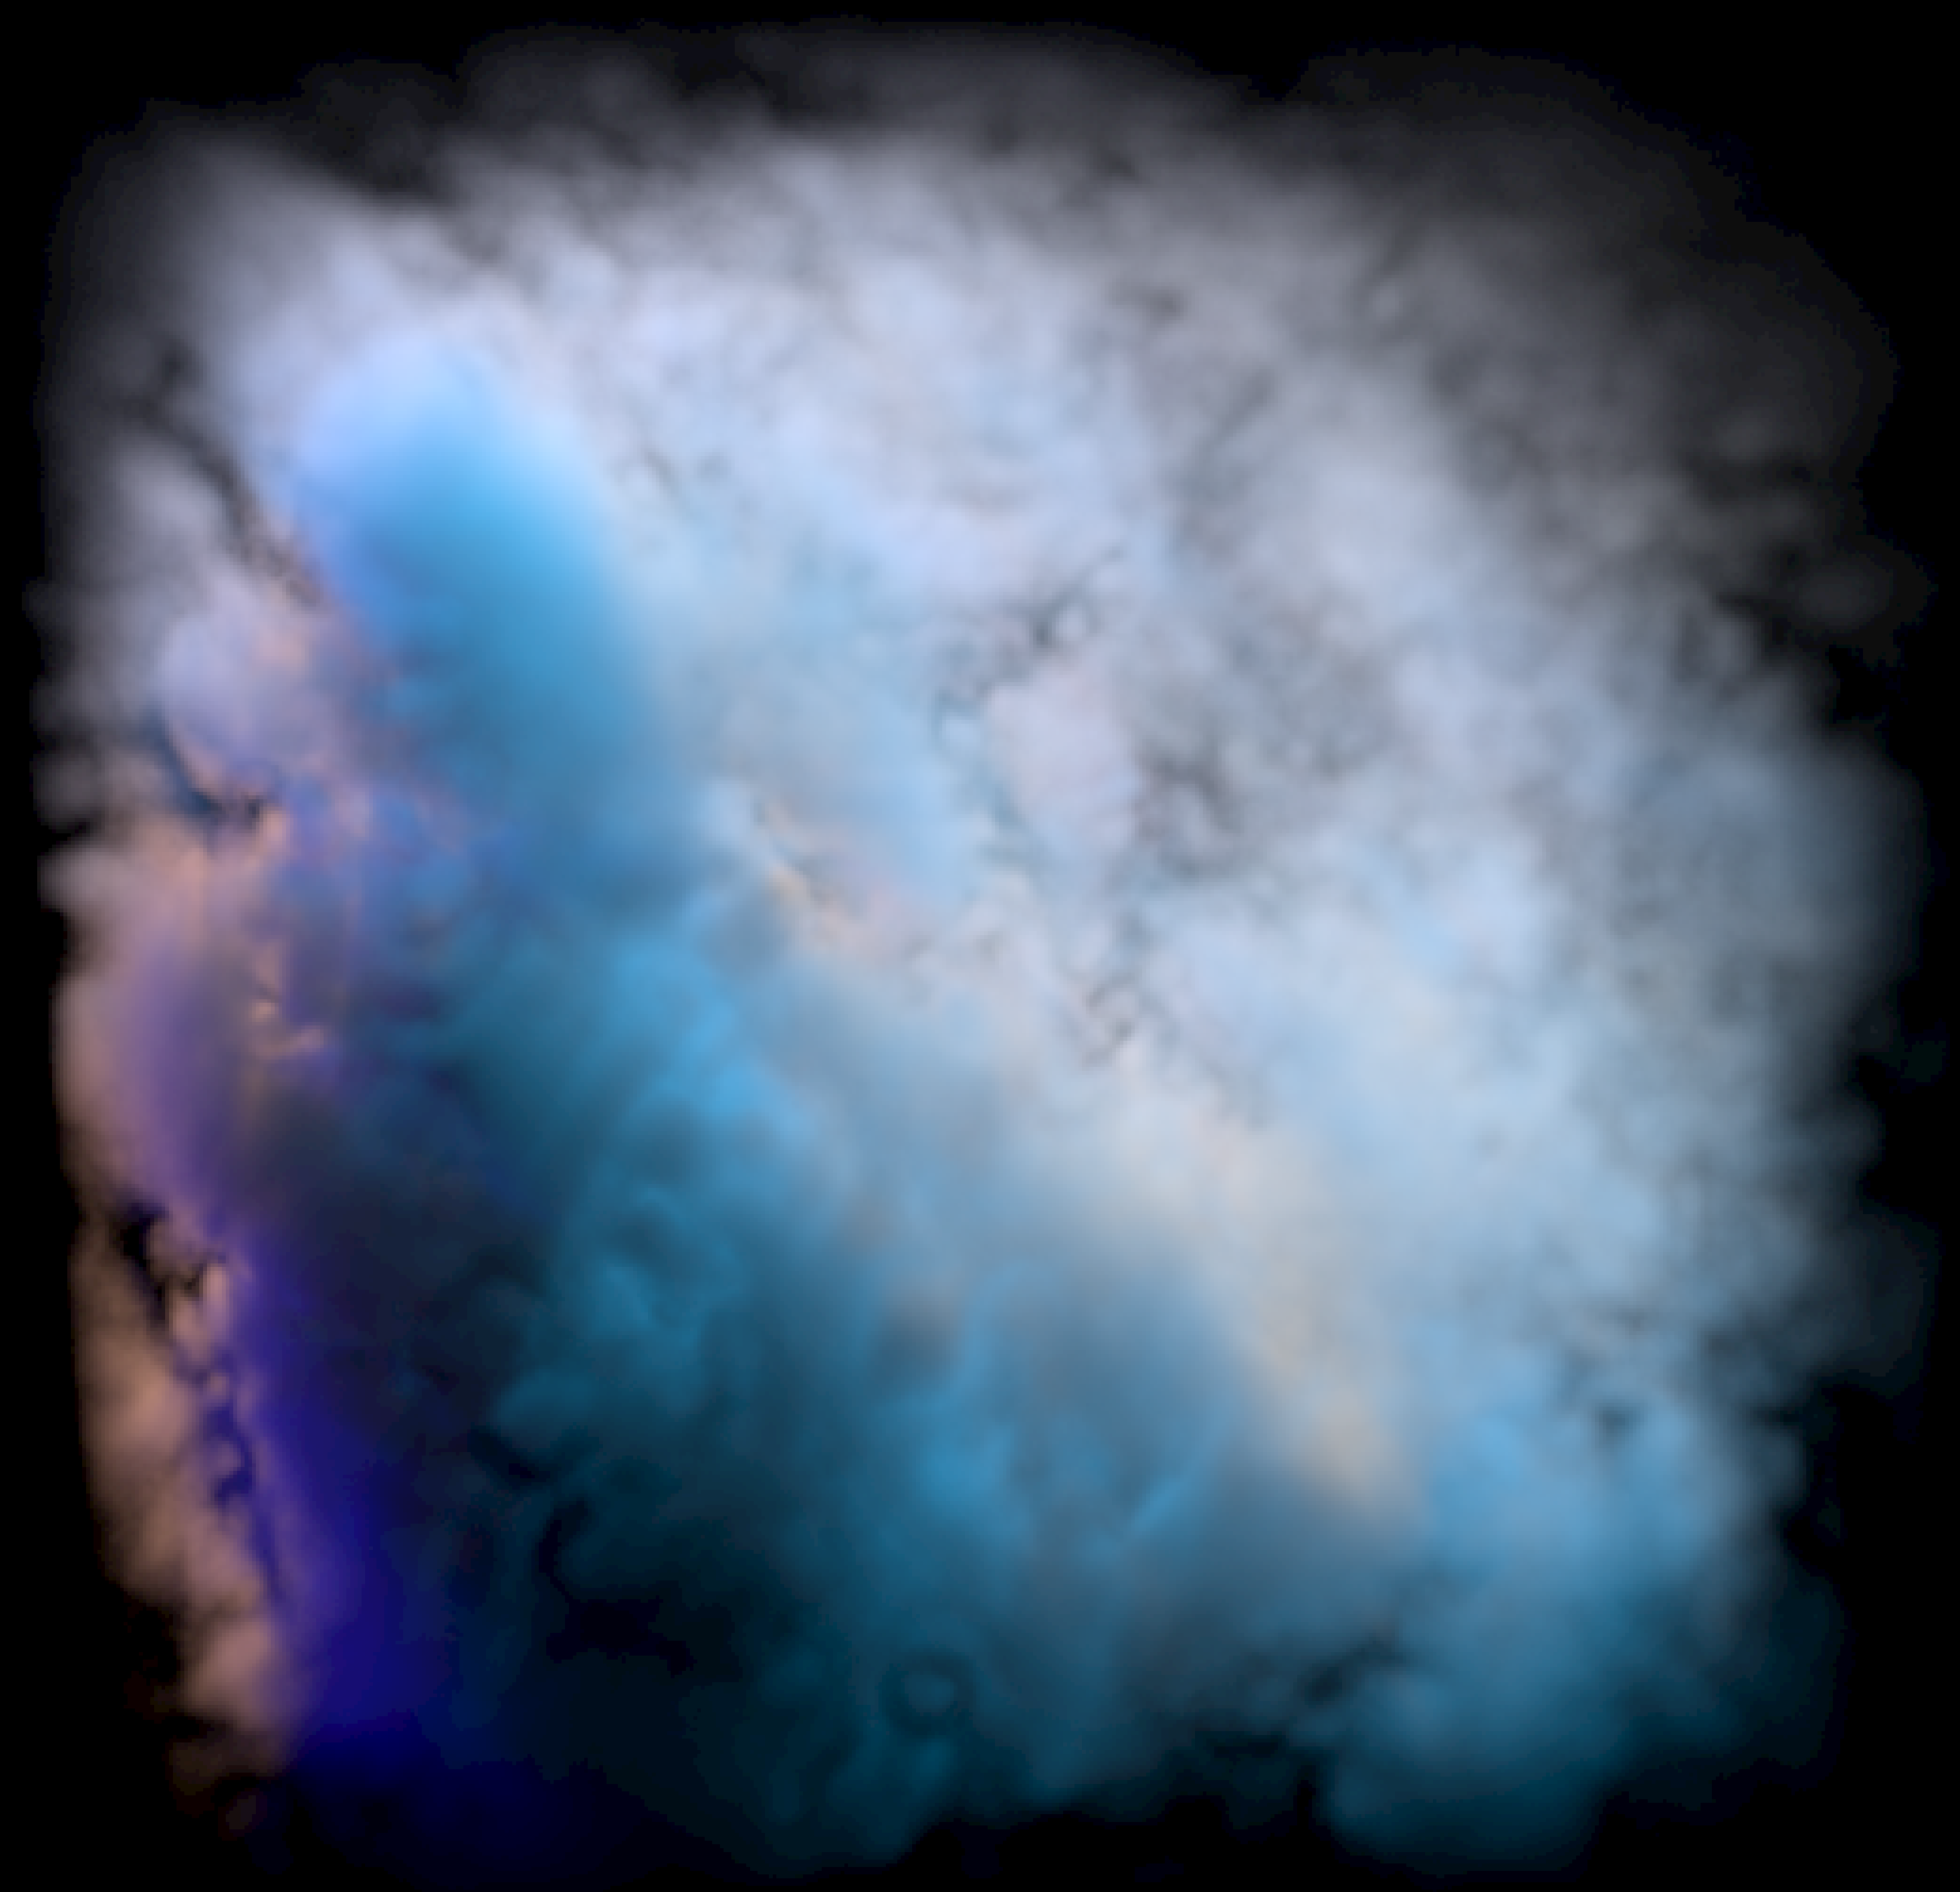

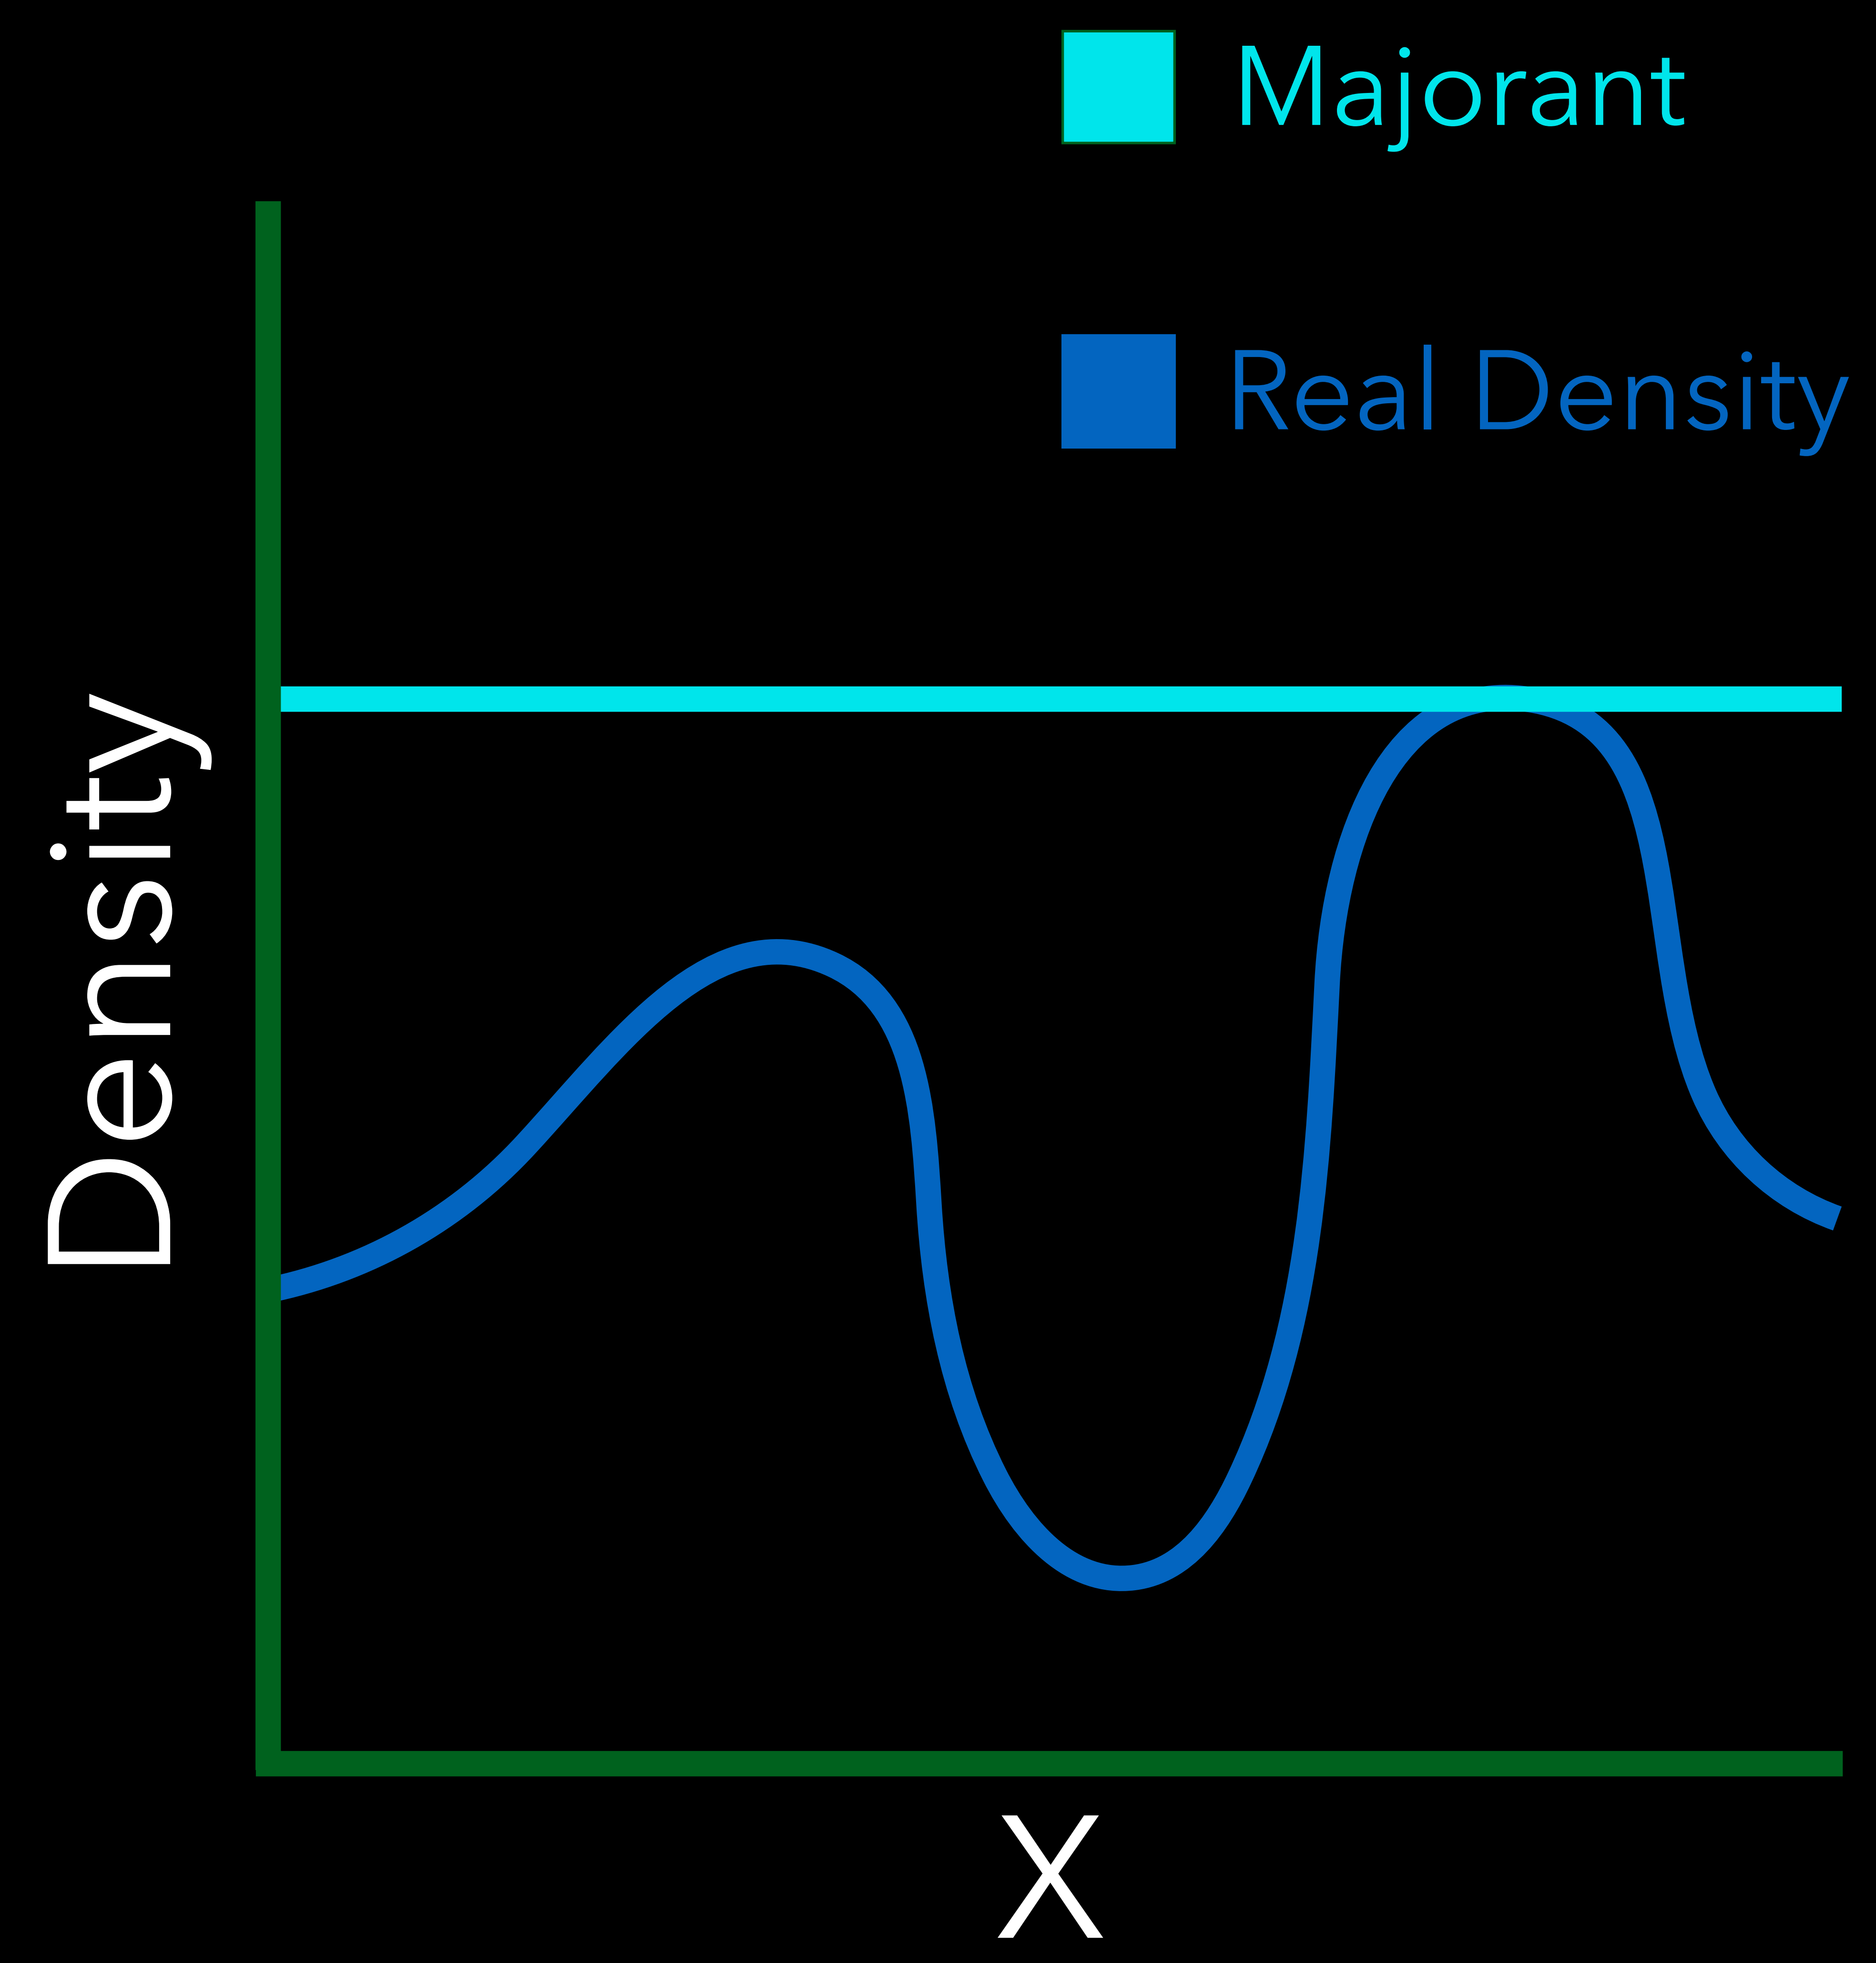
Majorant
Real Density
Density
X

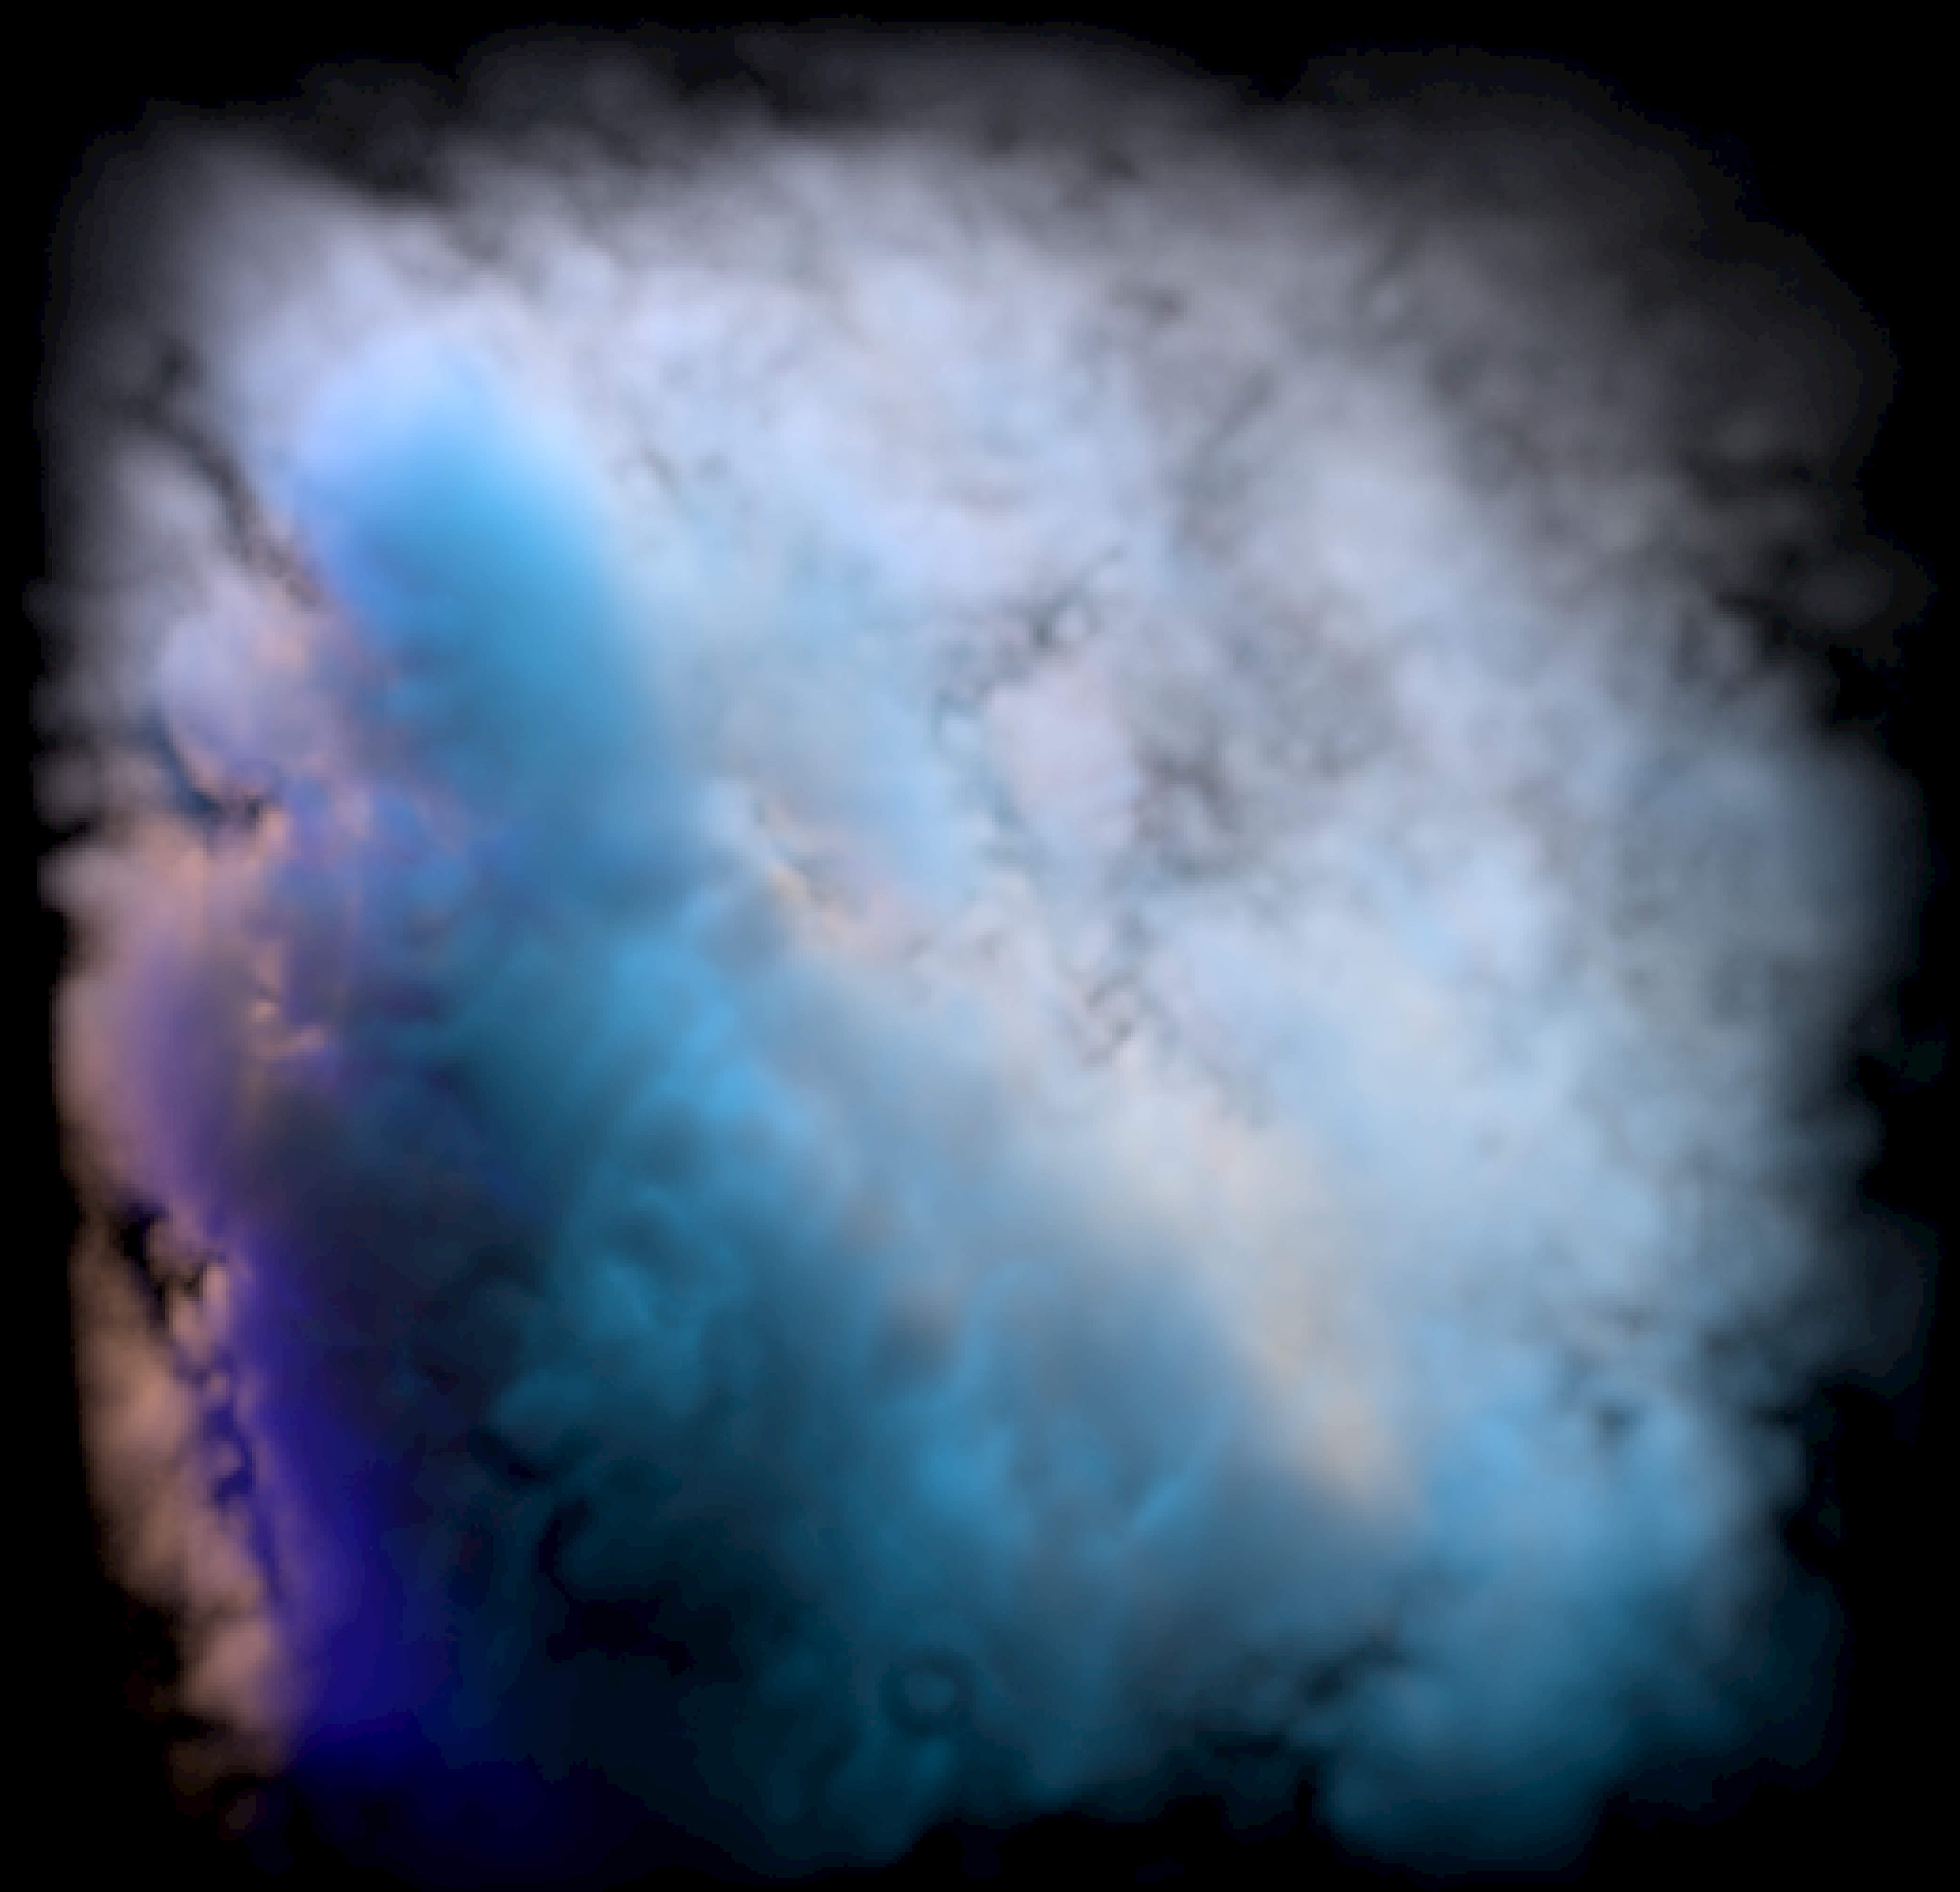

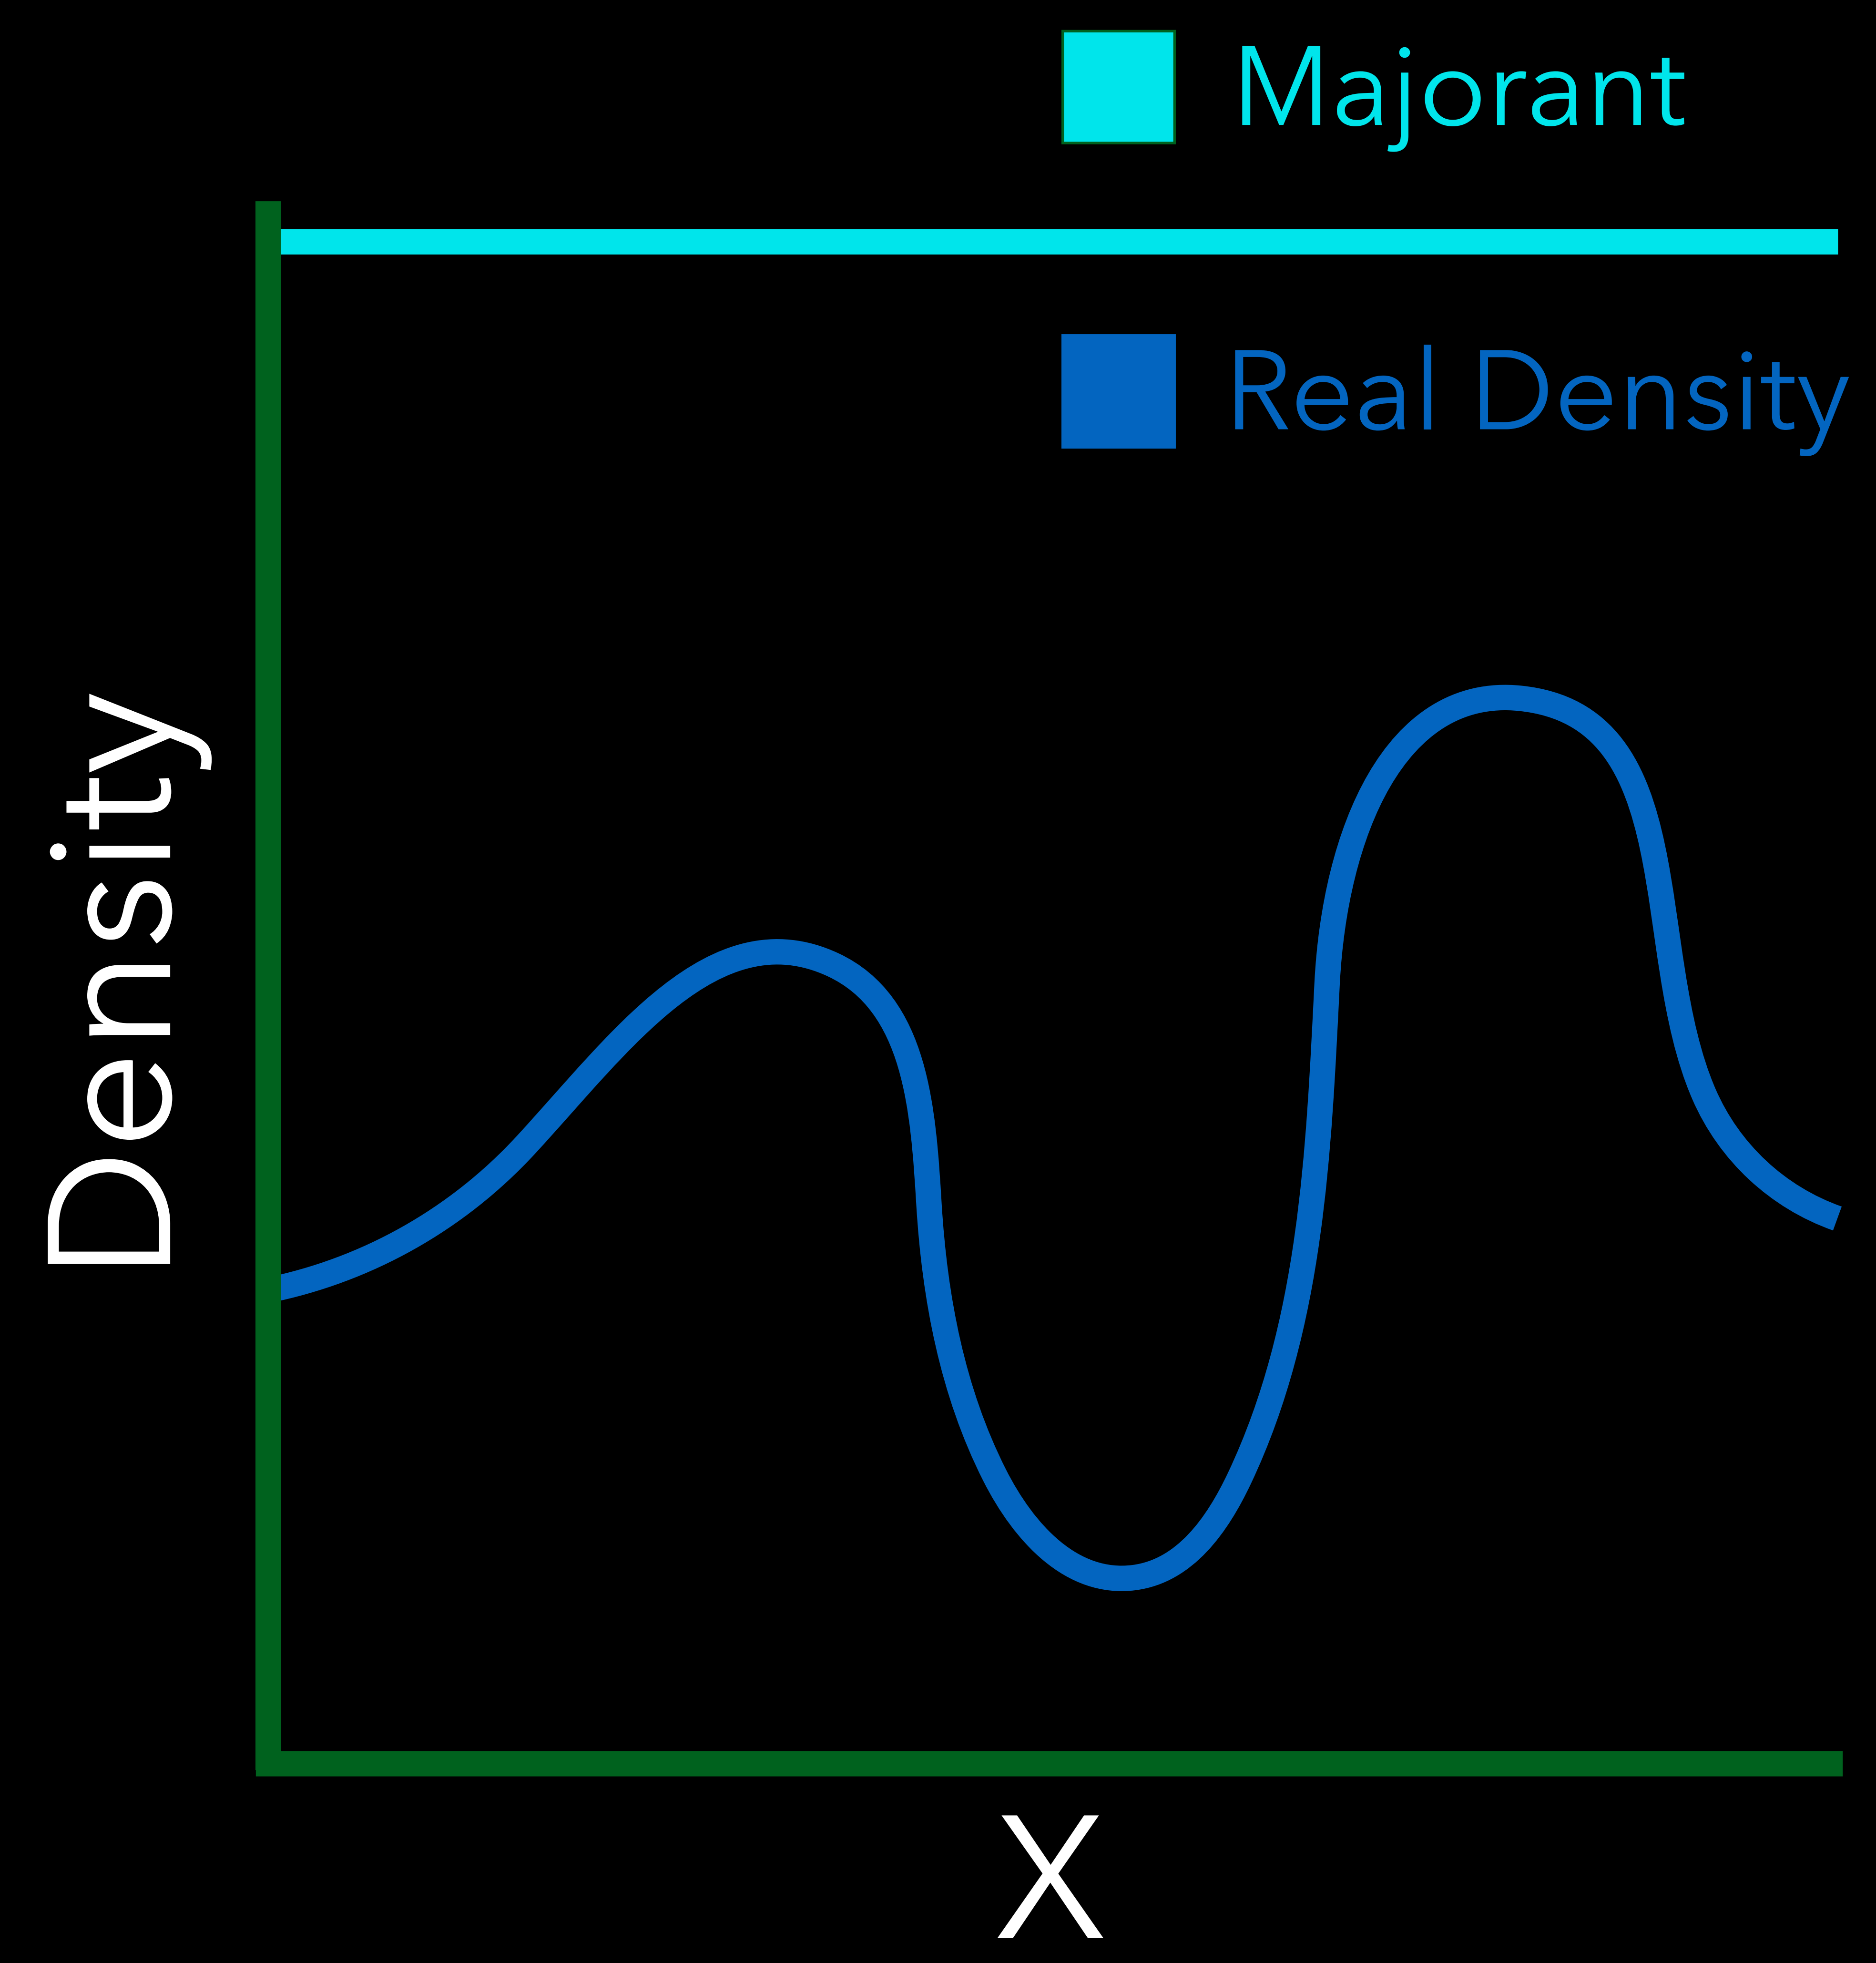
Majorant
Real Density
Density
X

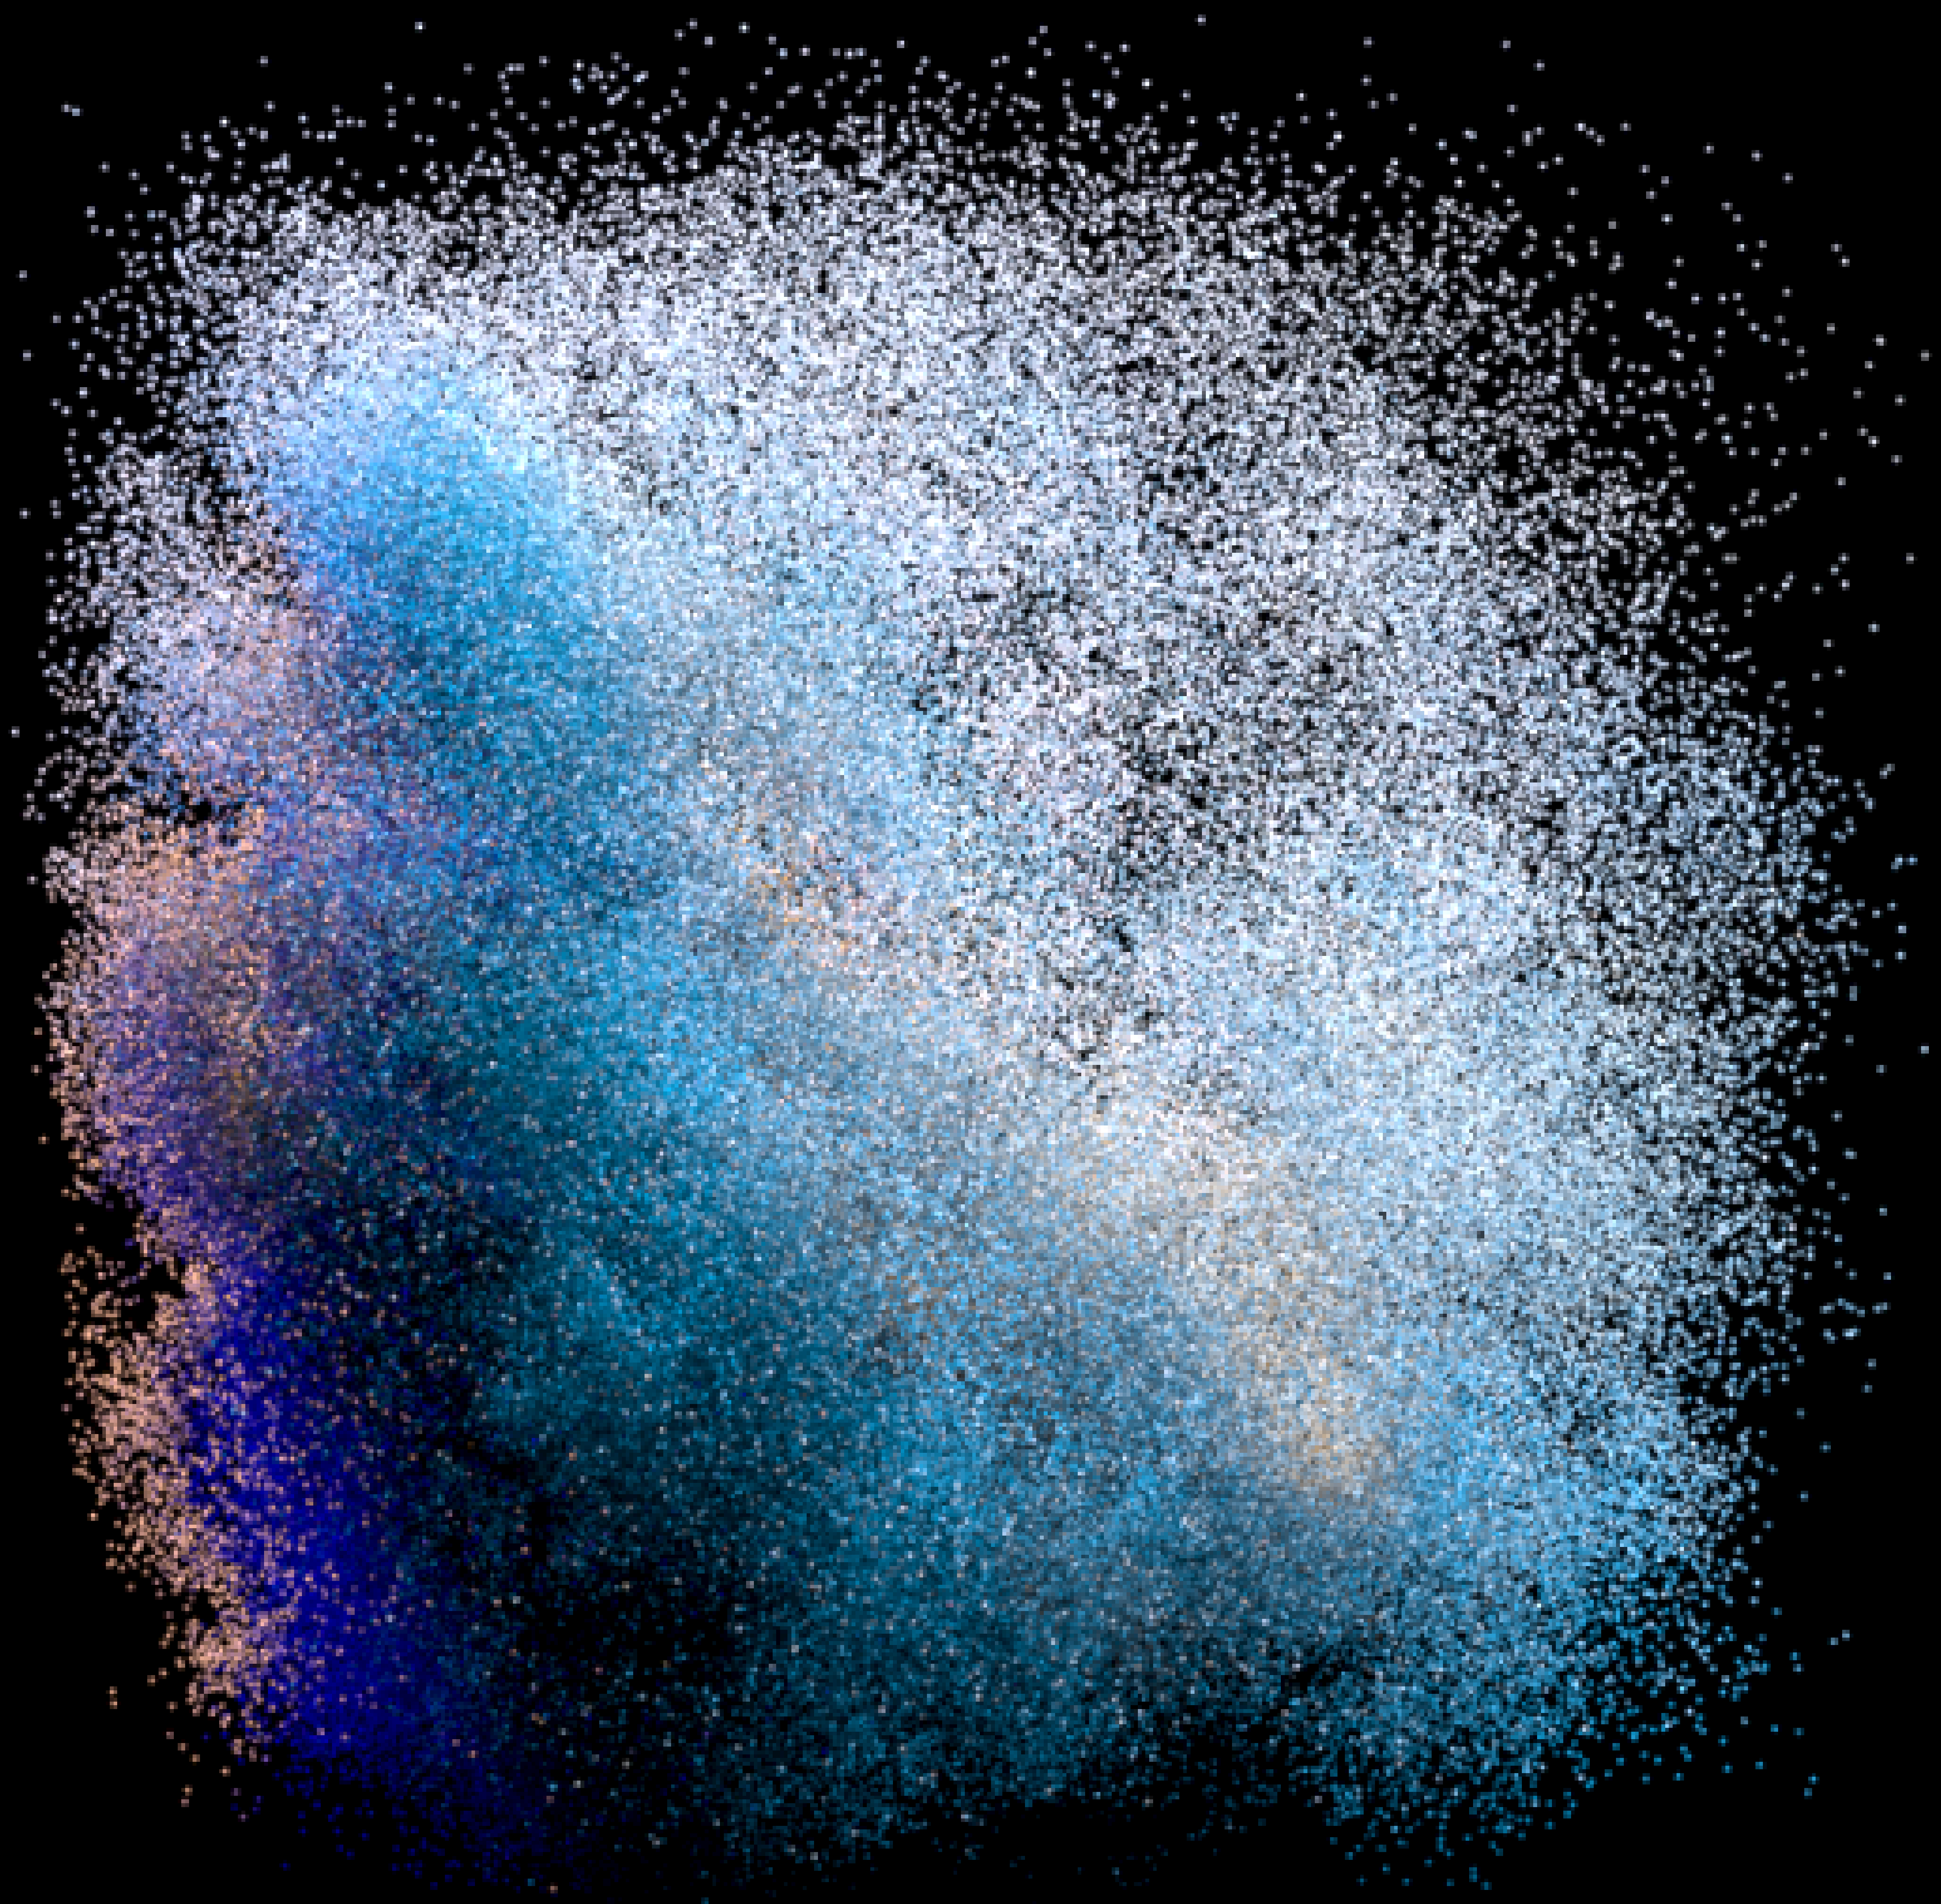

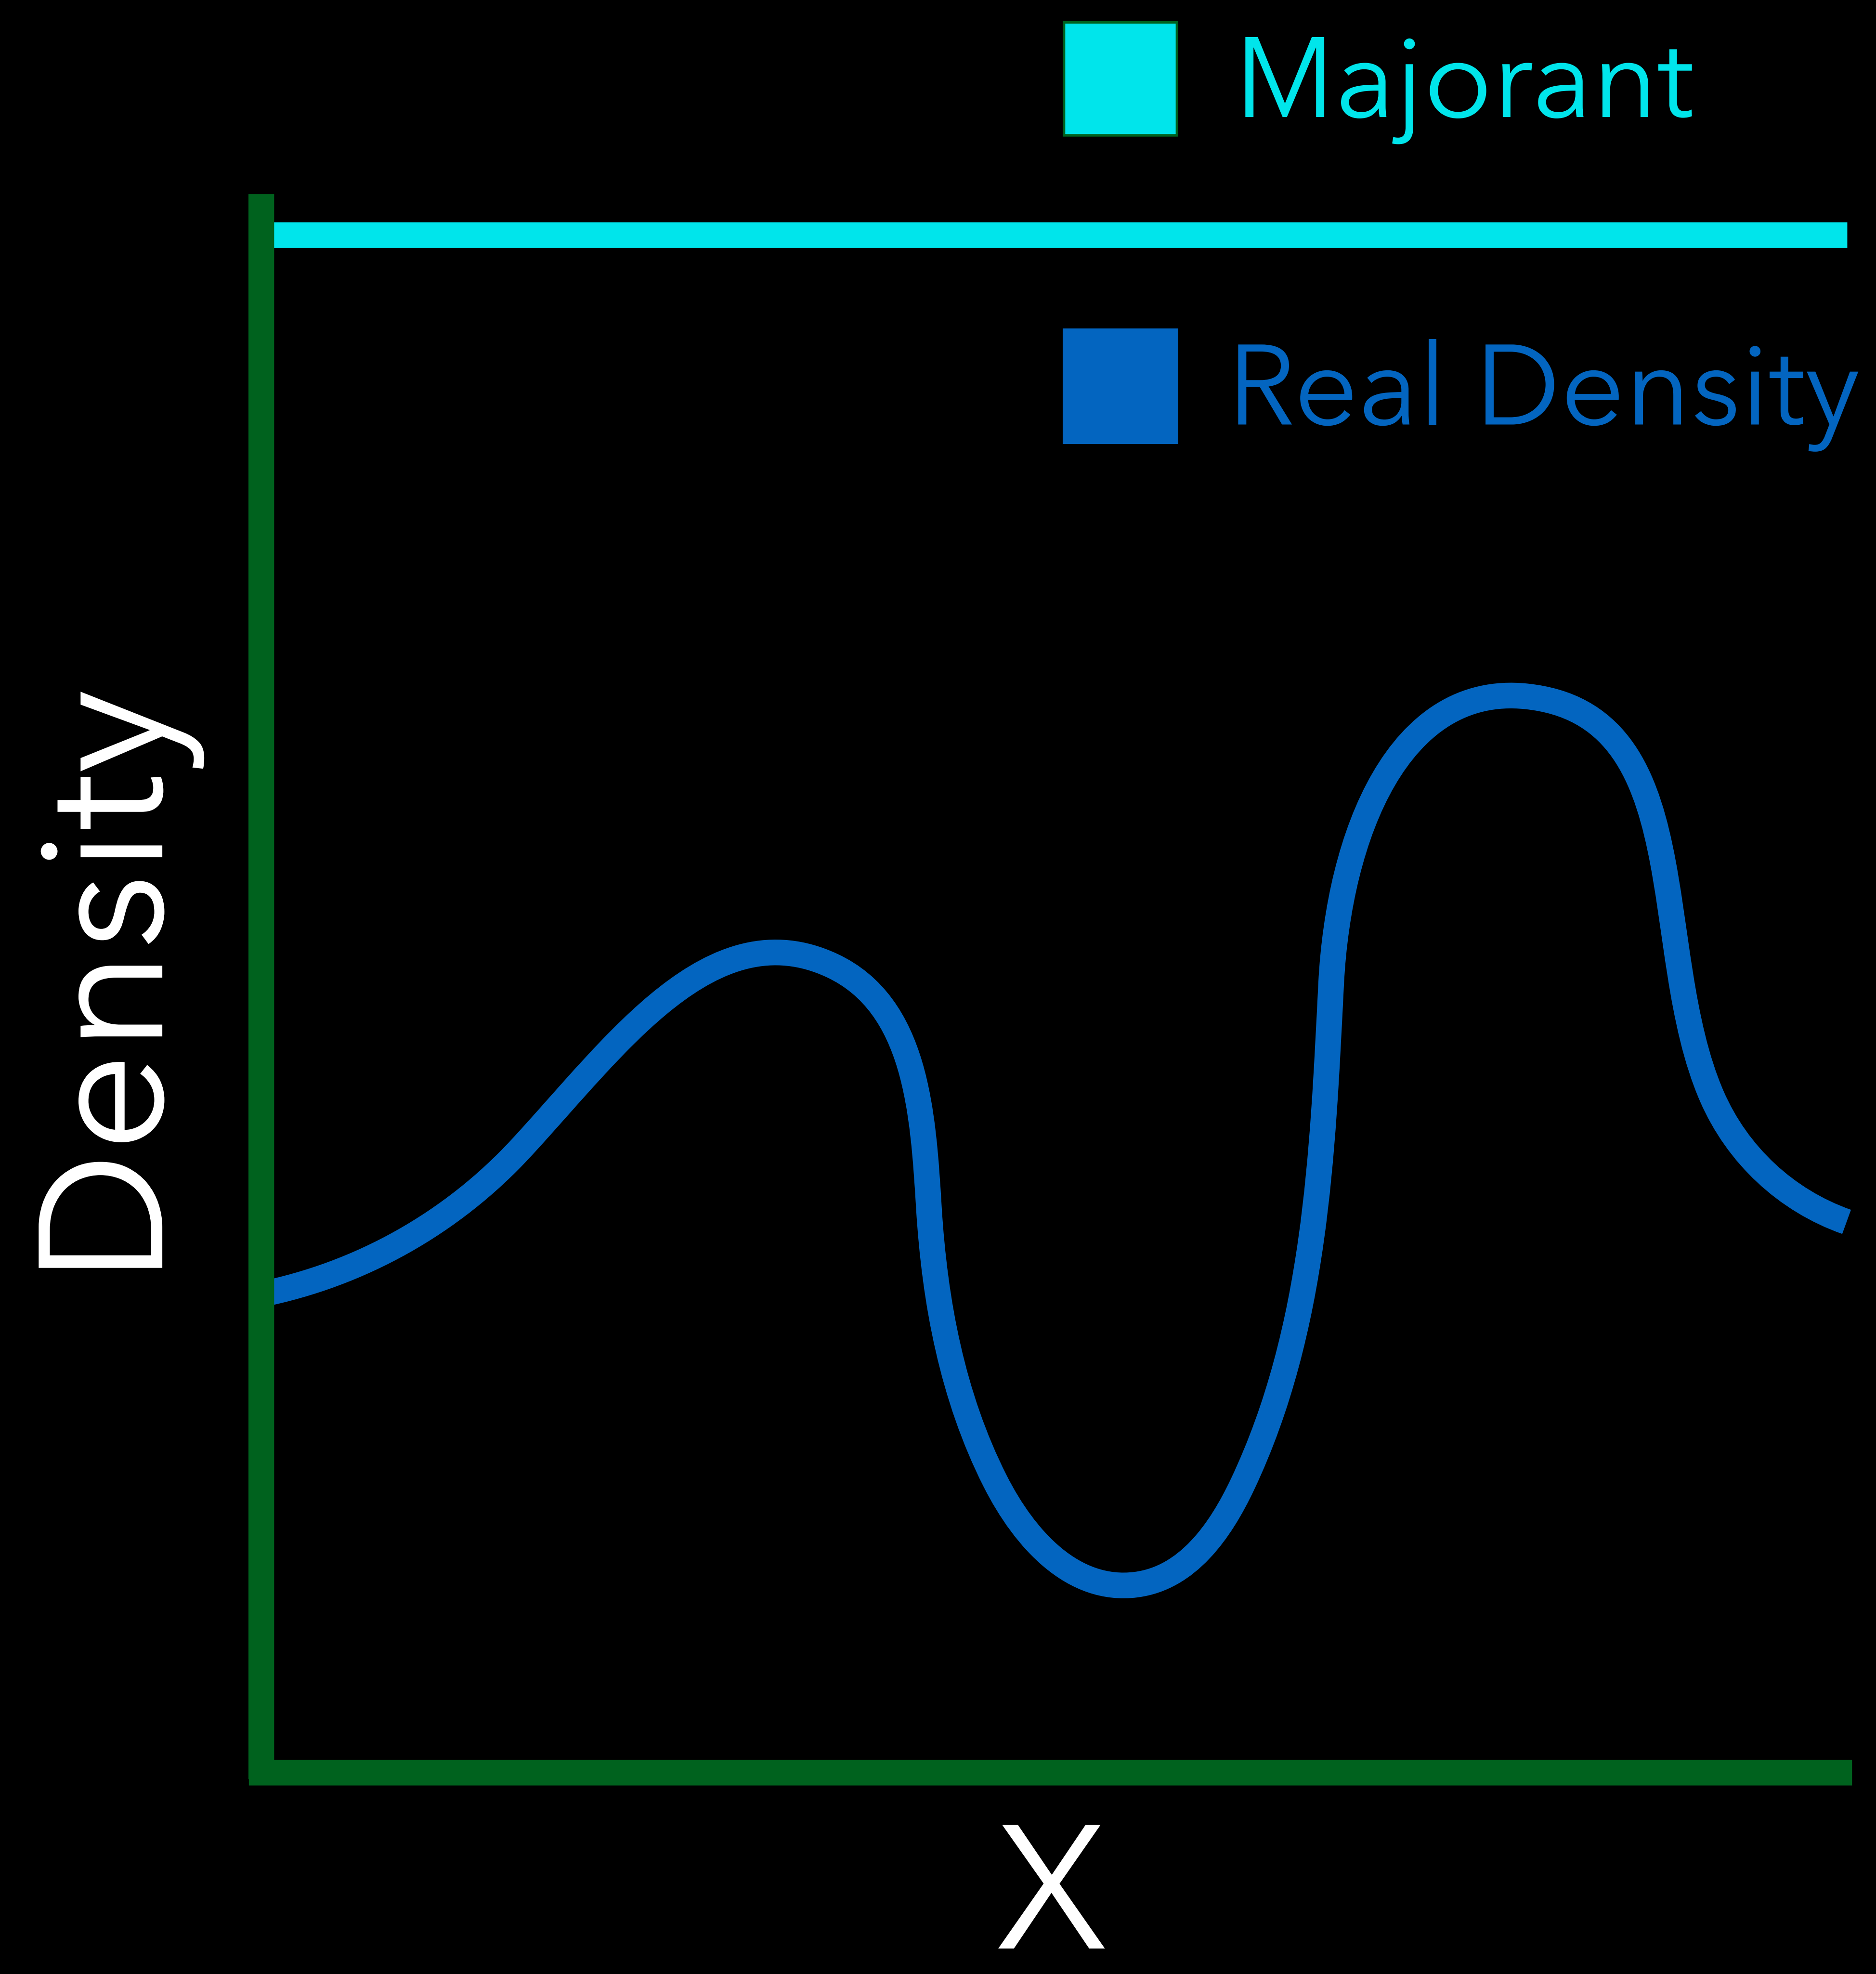
Majorant
Real Density
Density
X

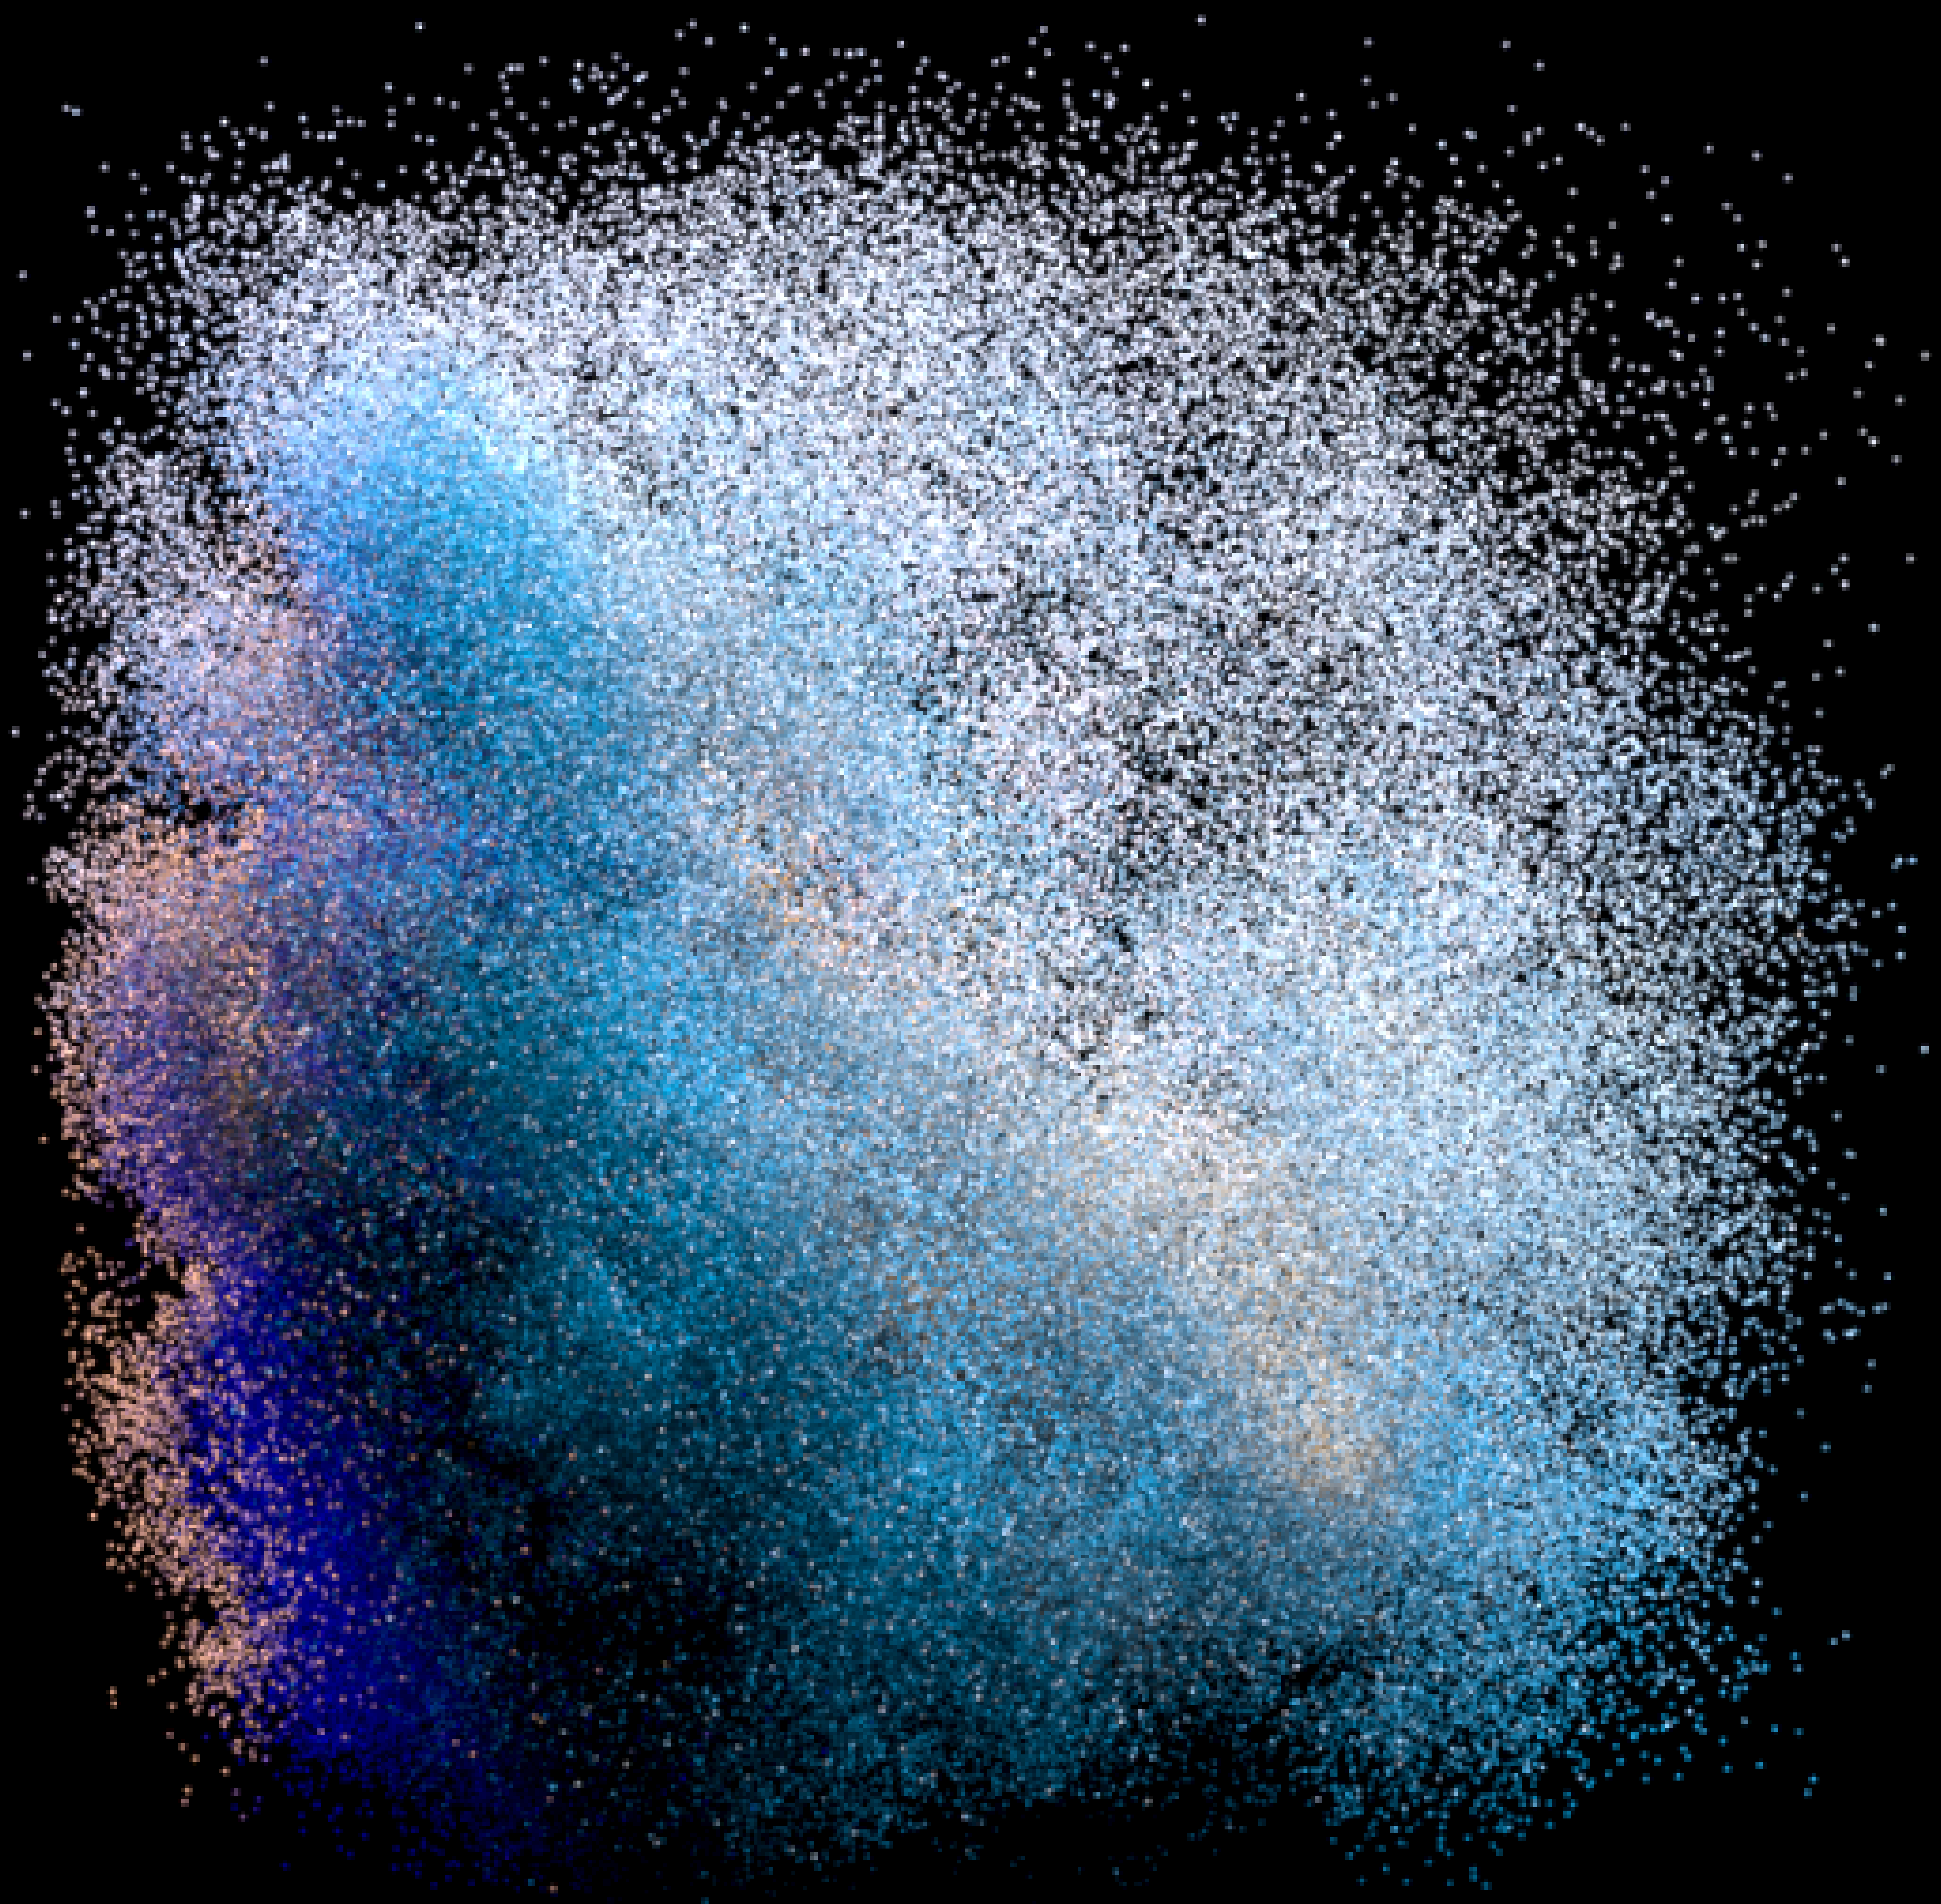

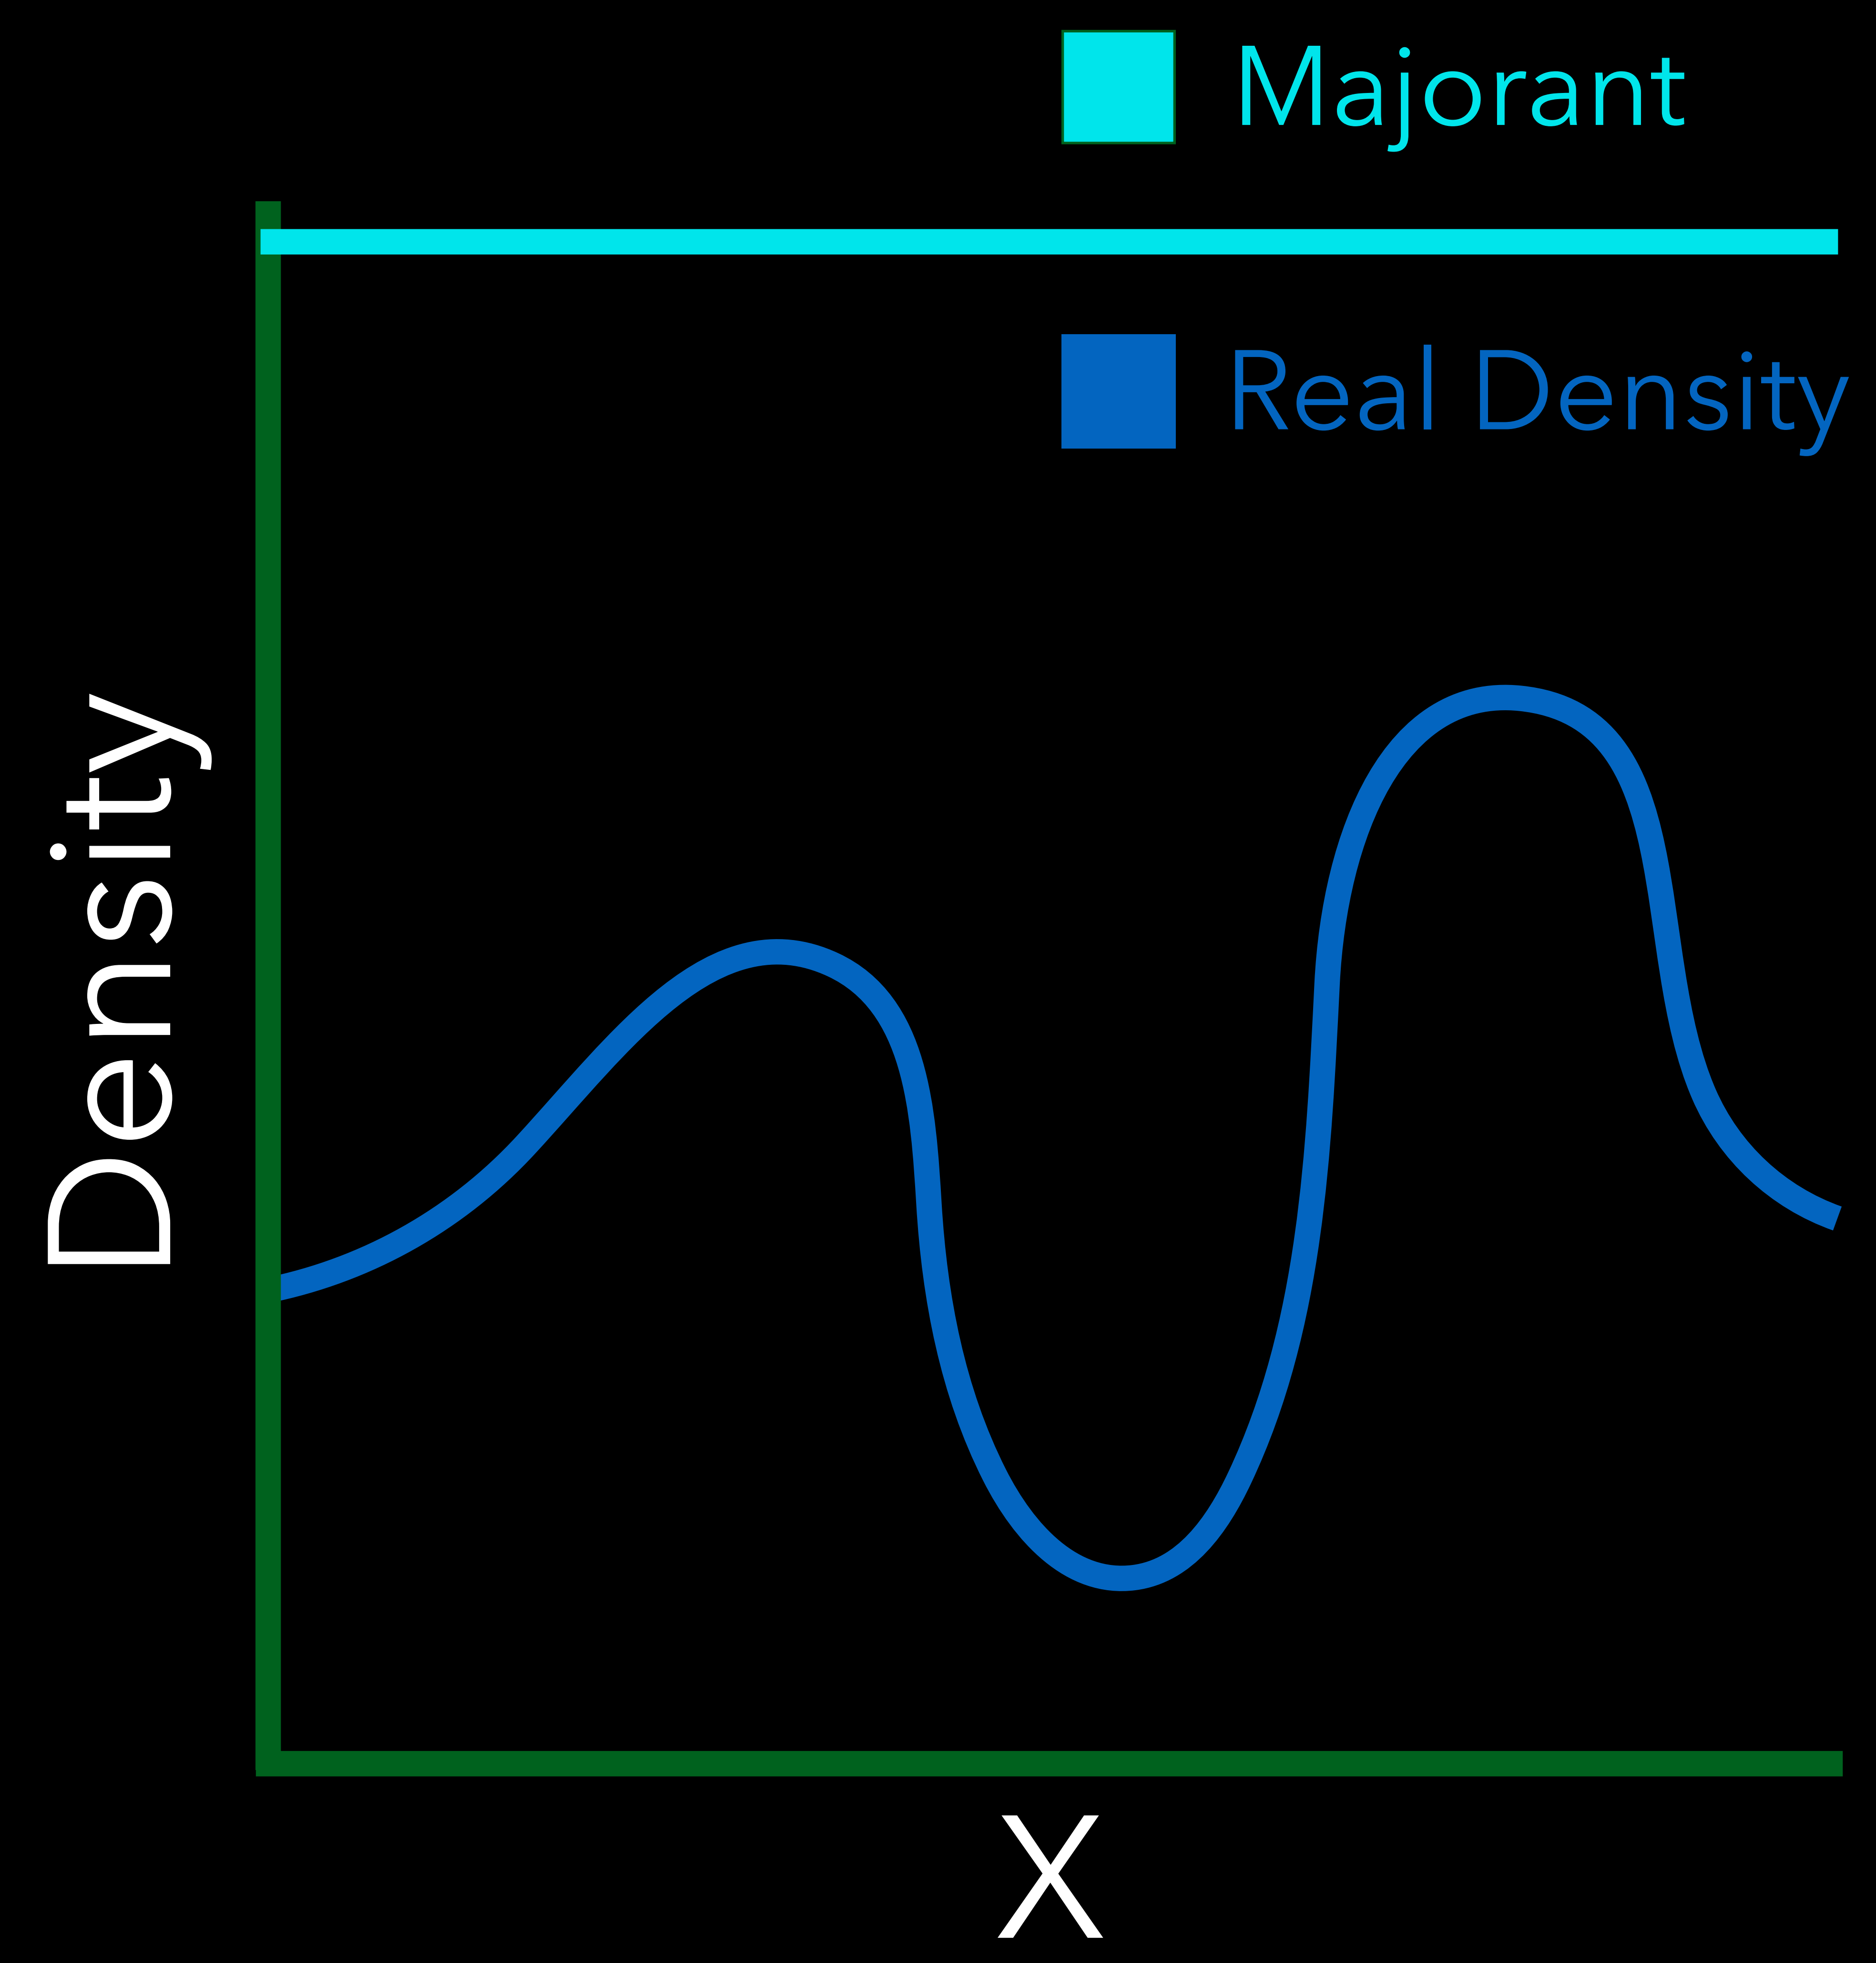
Majorant
Real Density
Density
X

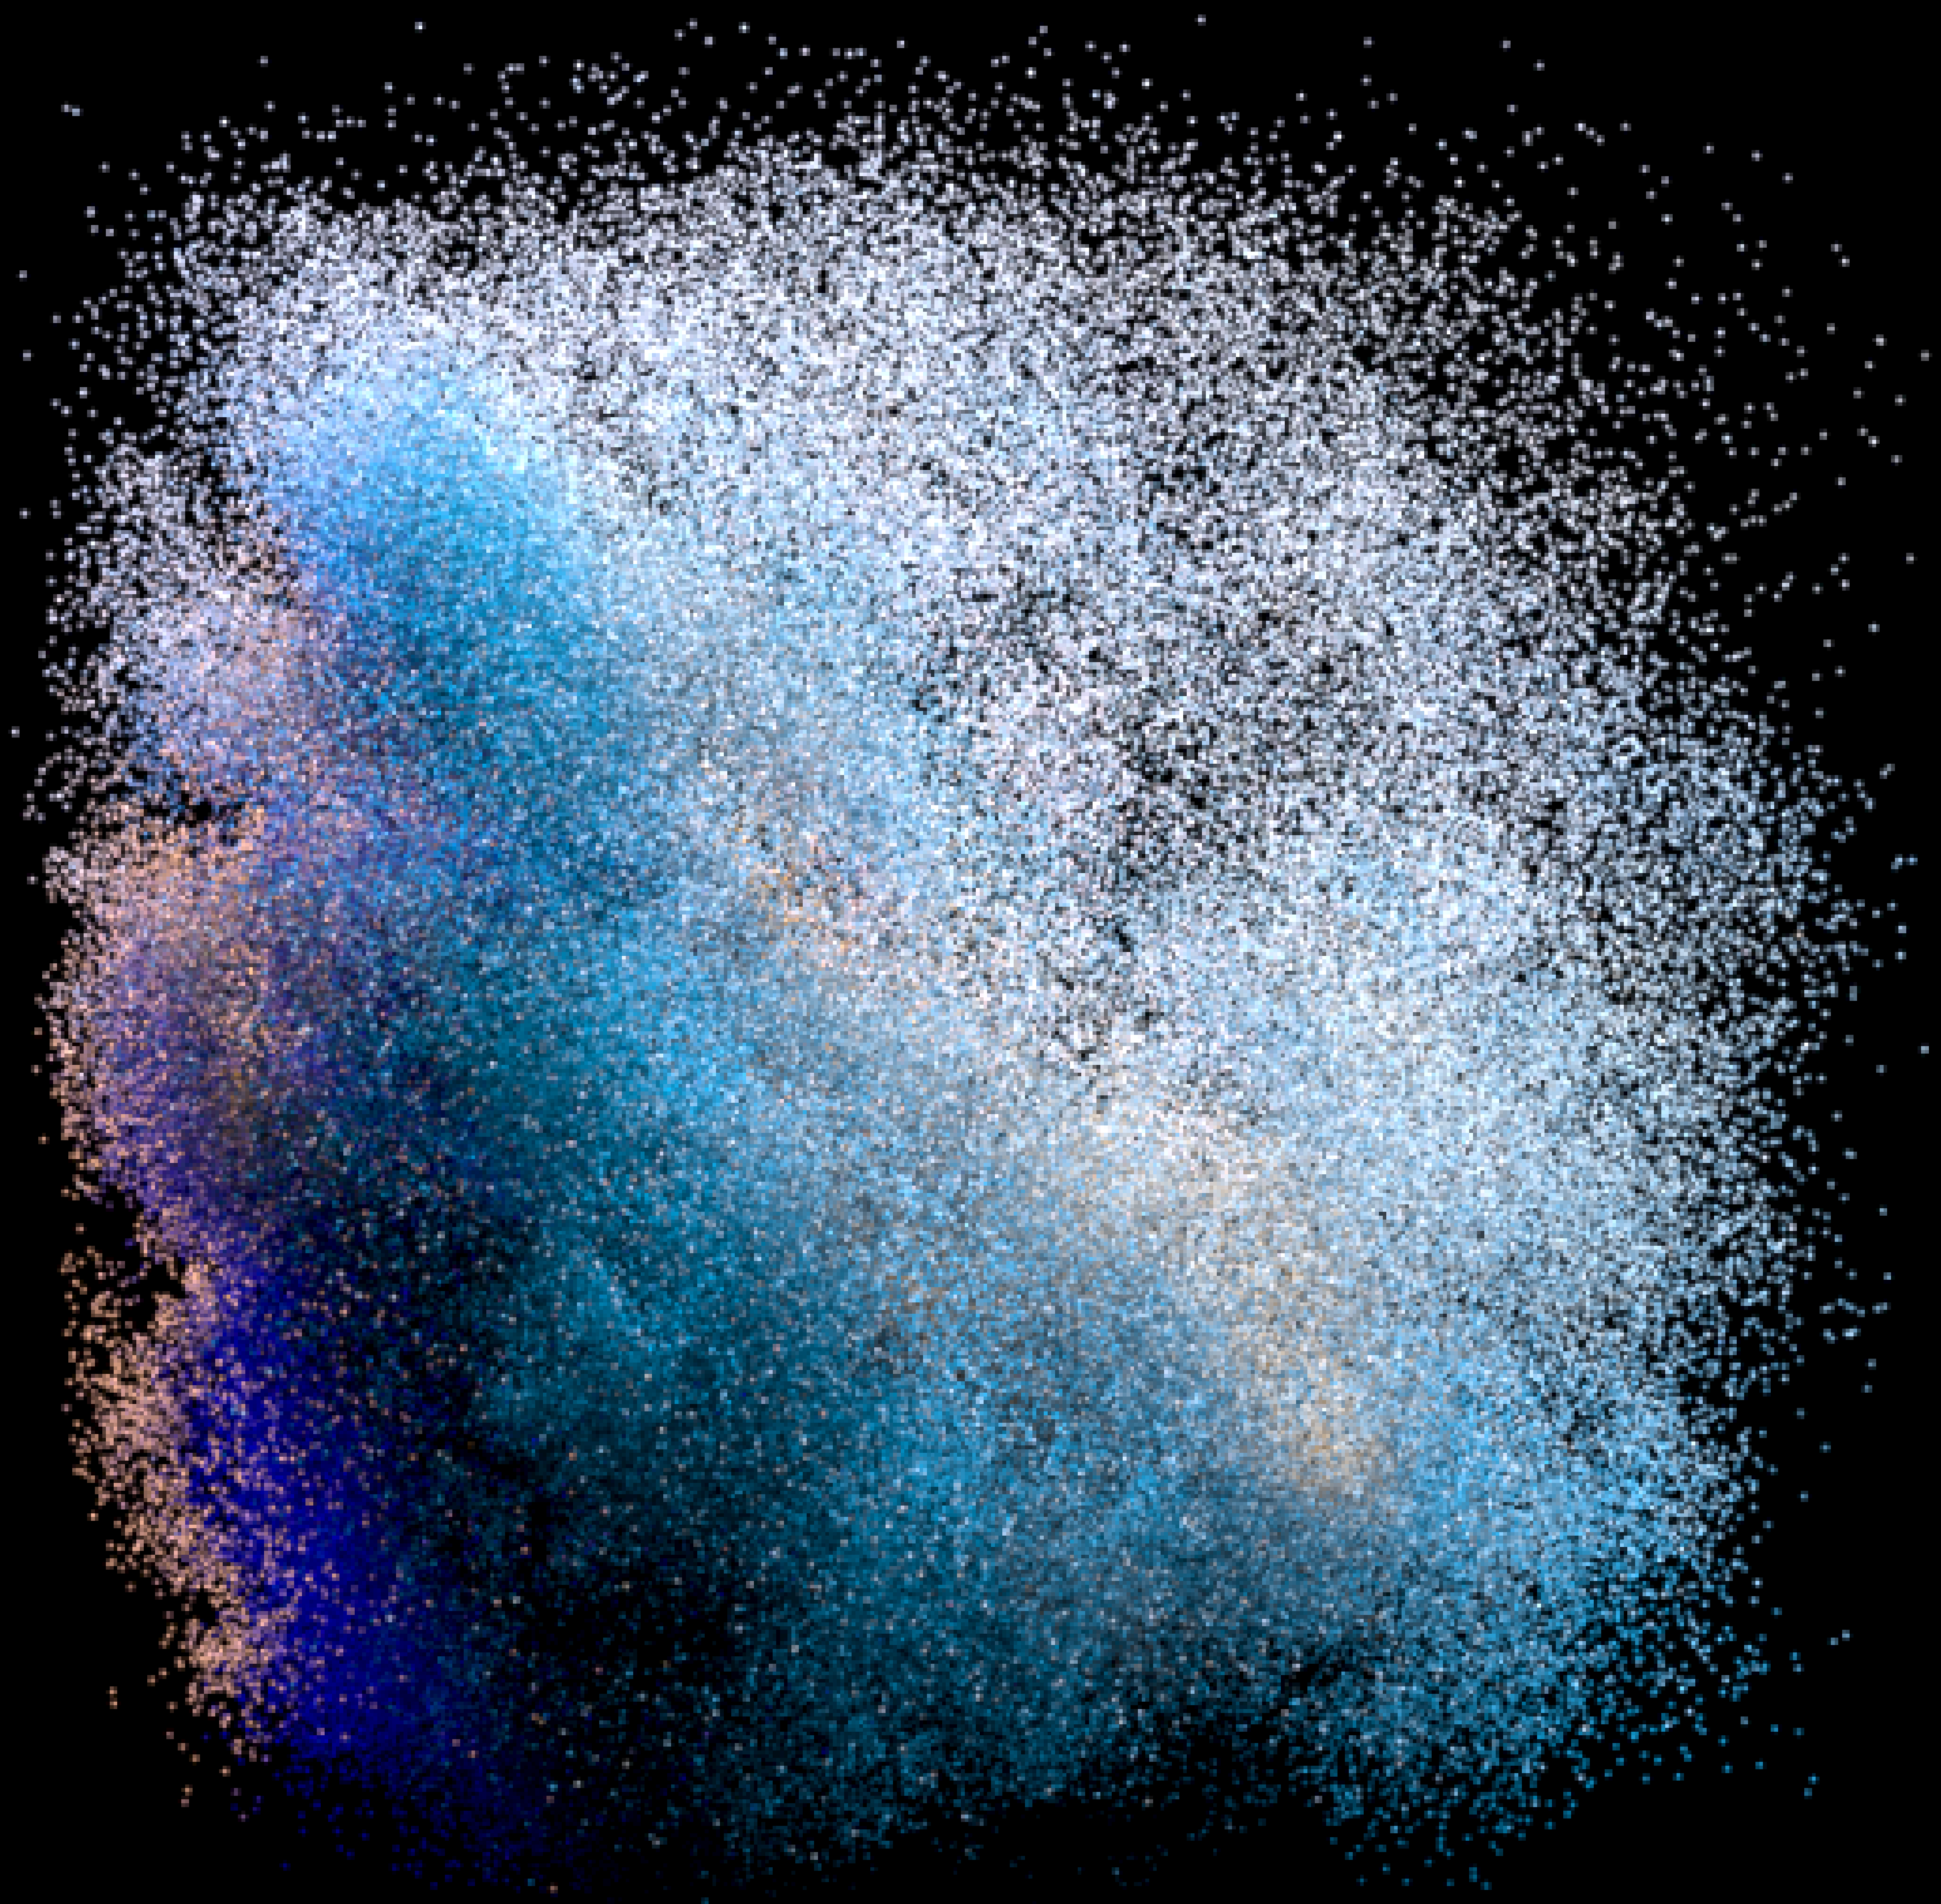

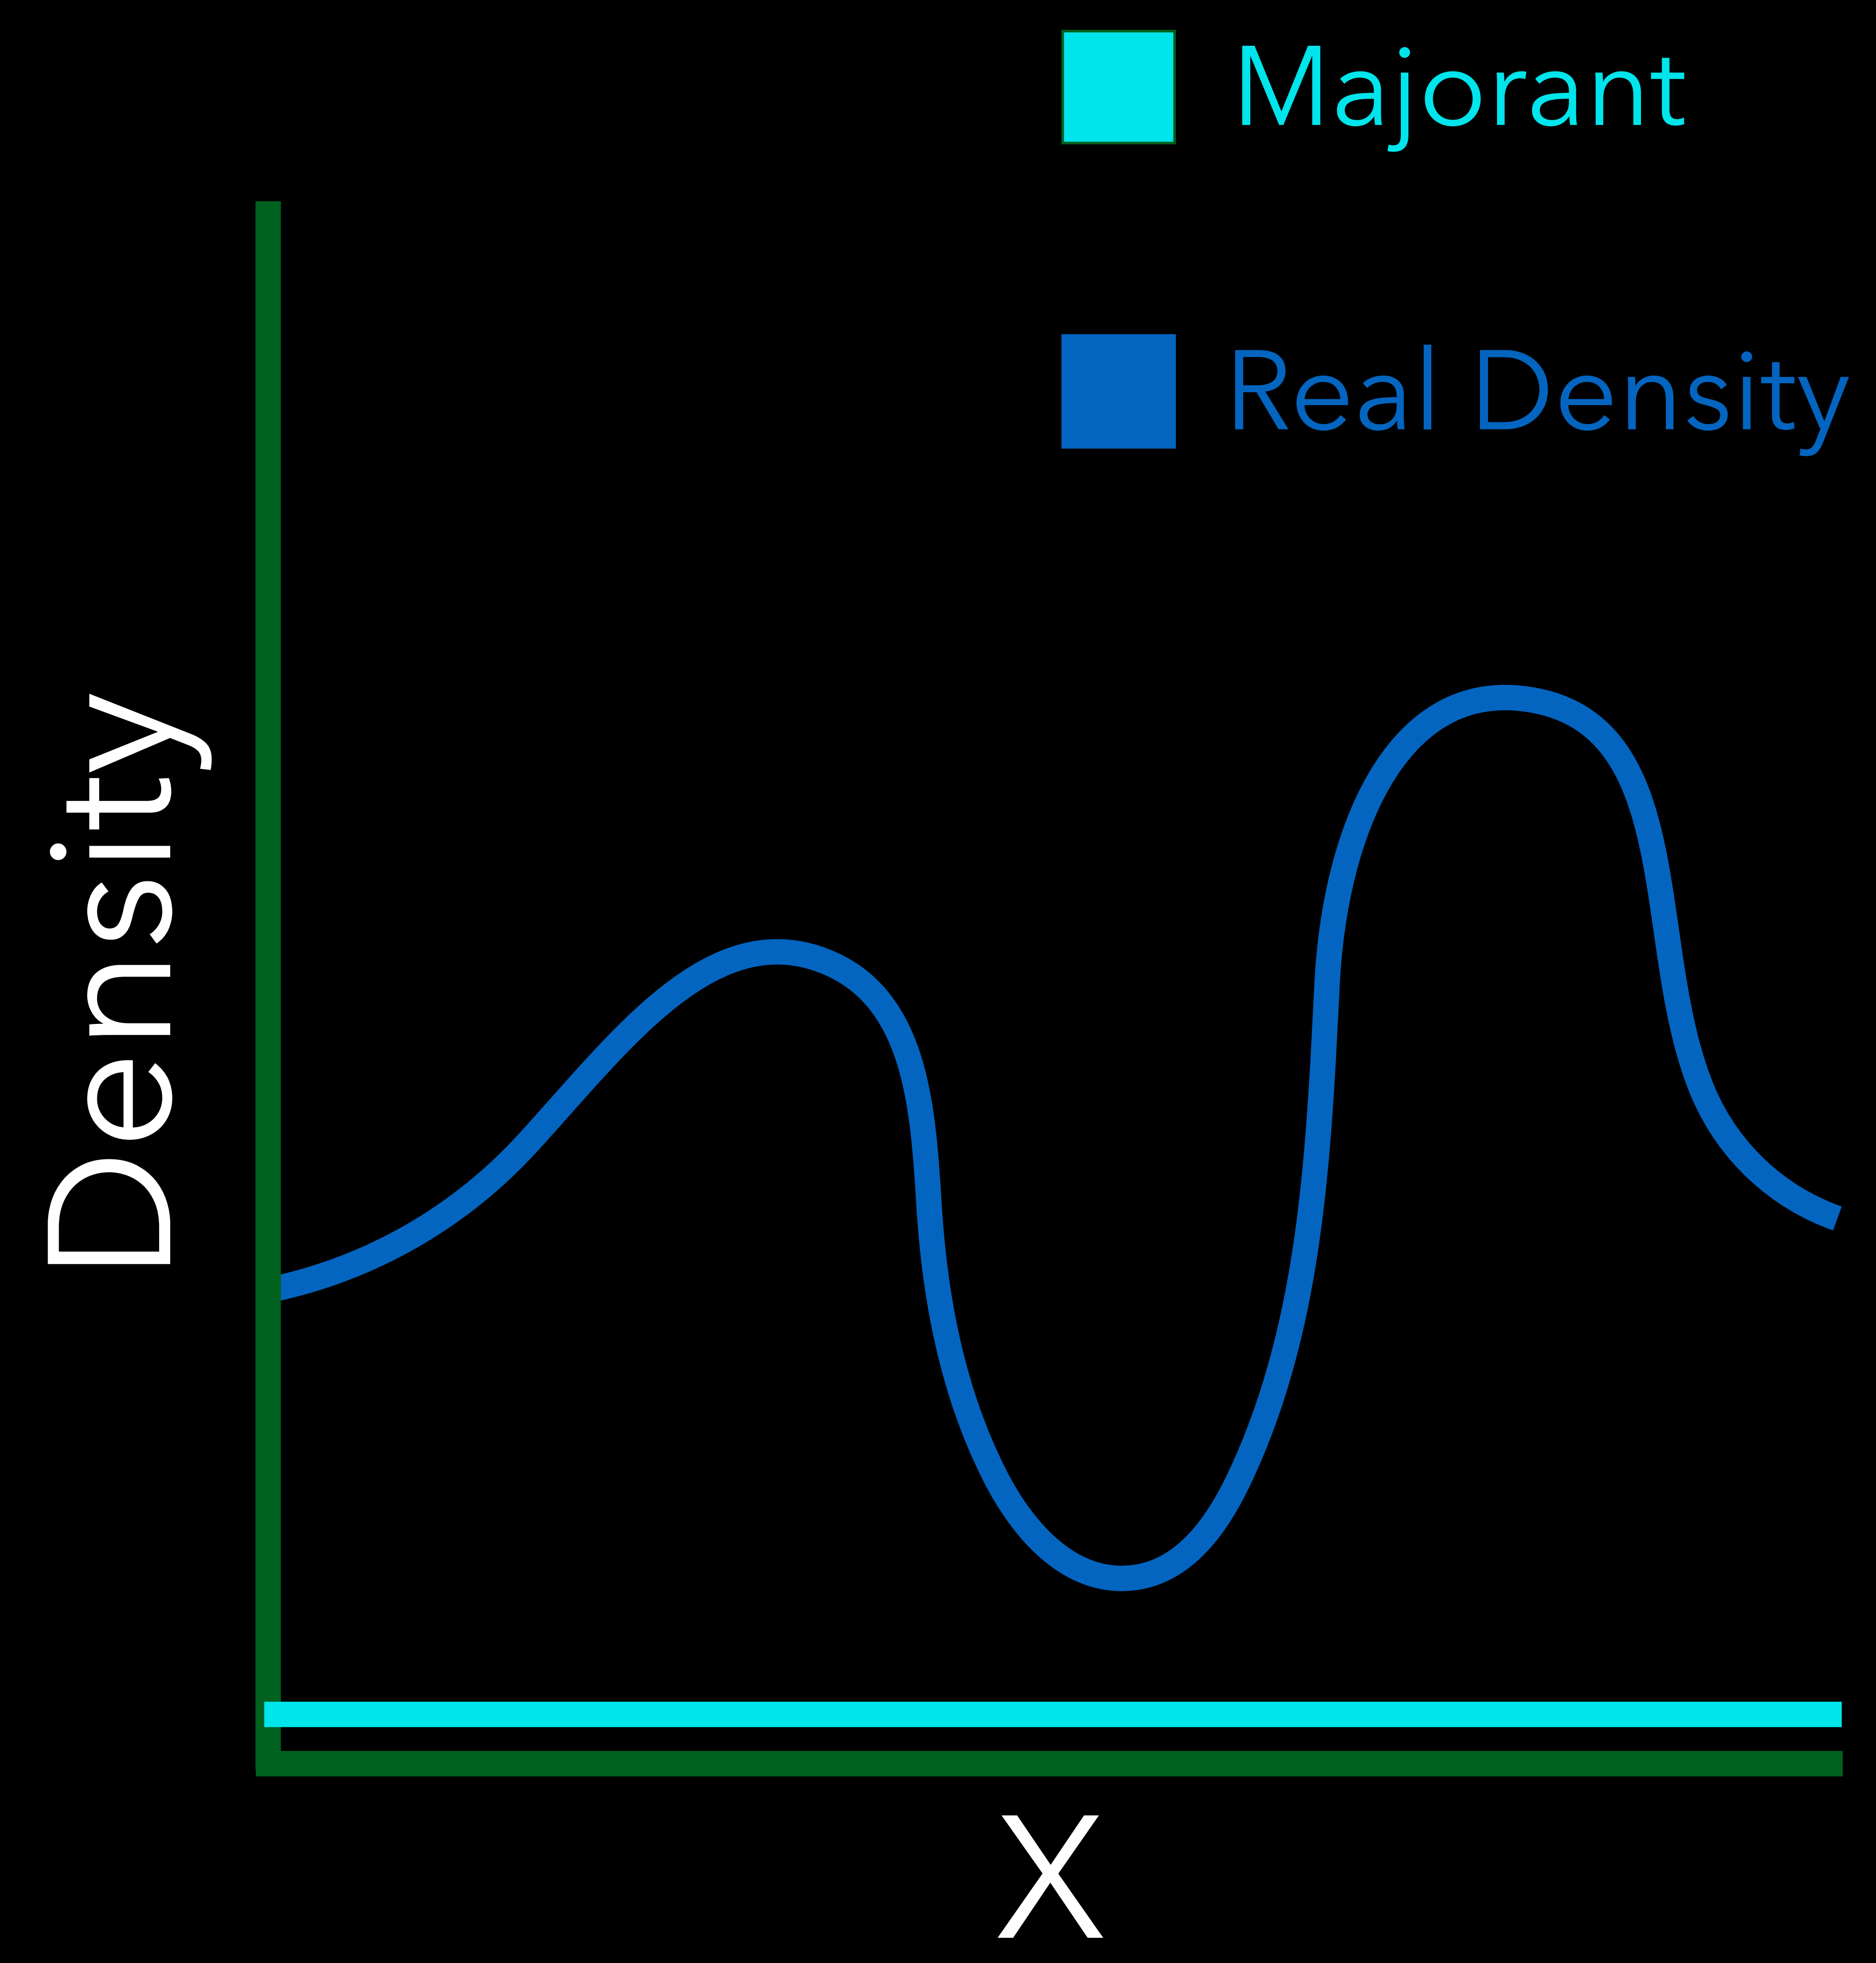
Majorant
Real Density
Density
X

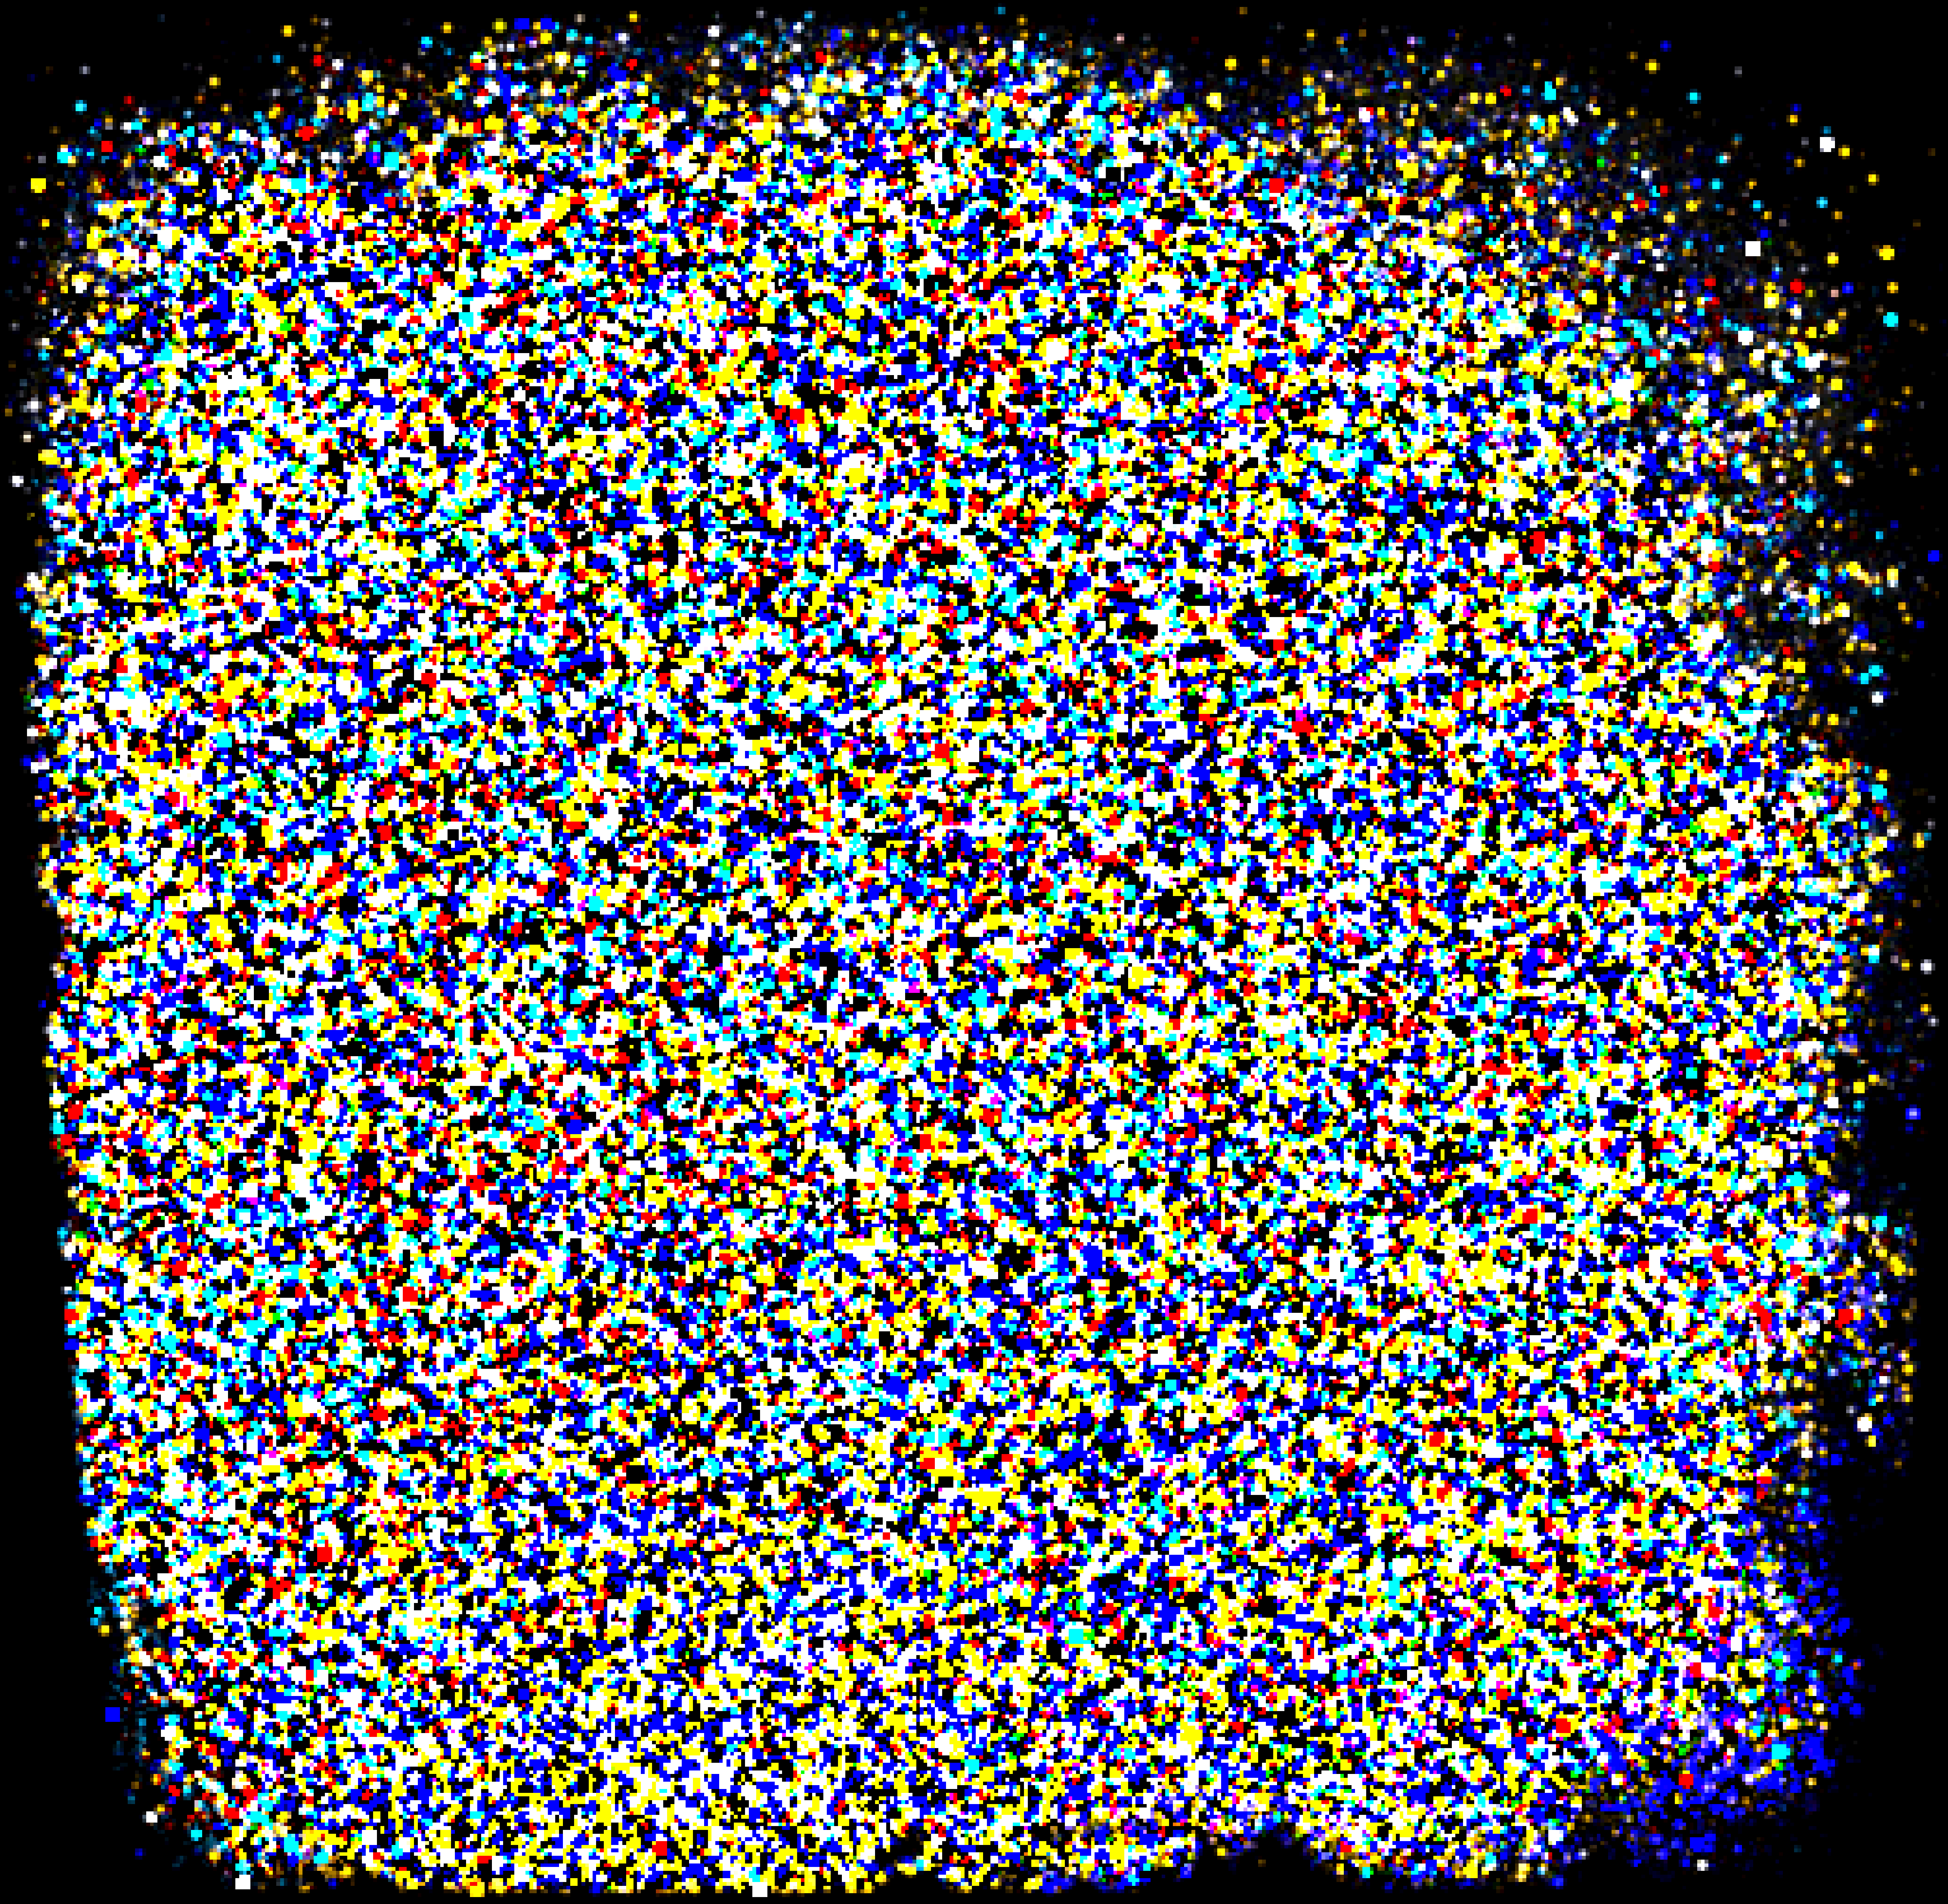

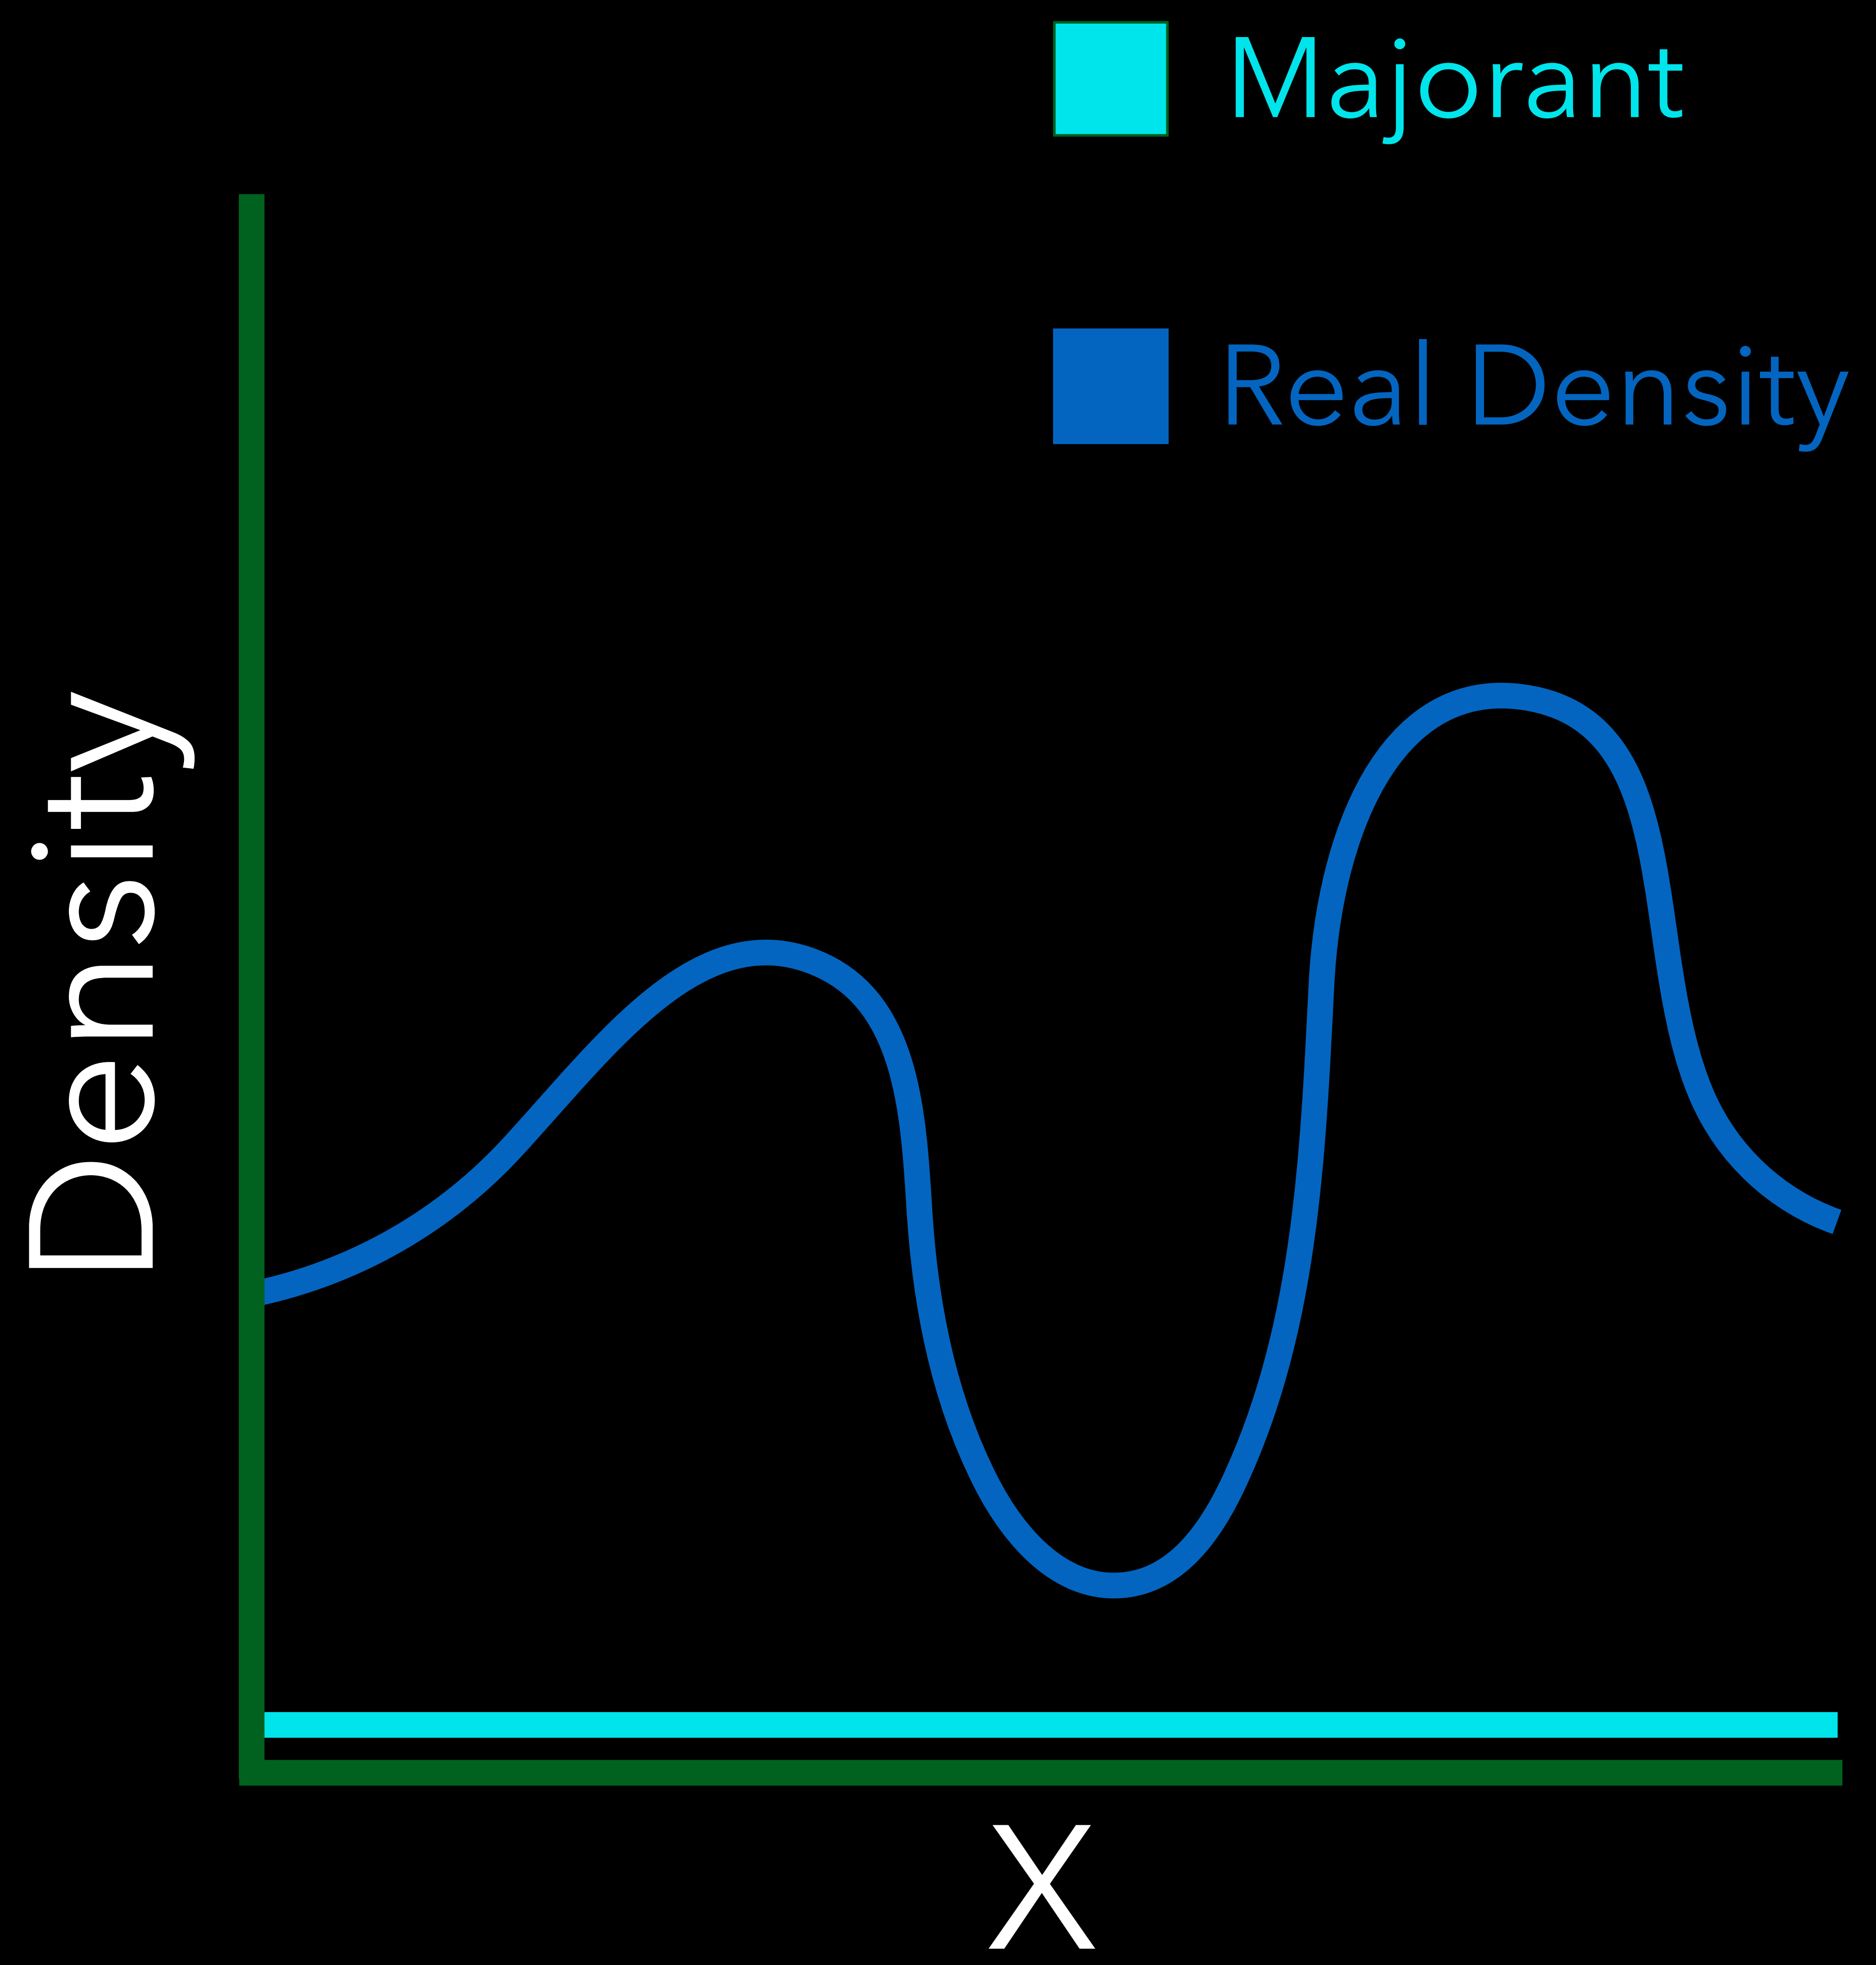
Majorant
Real Density
Density
X

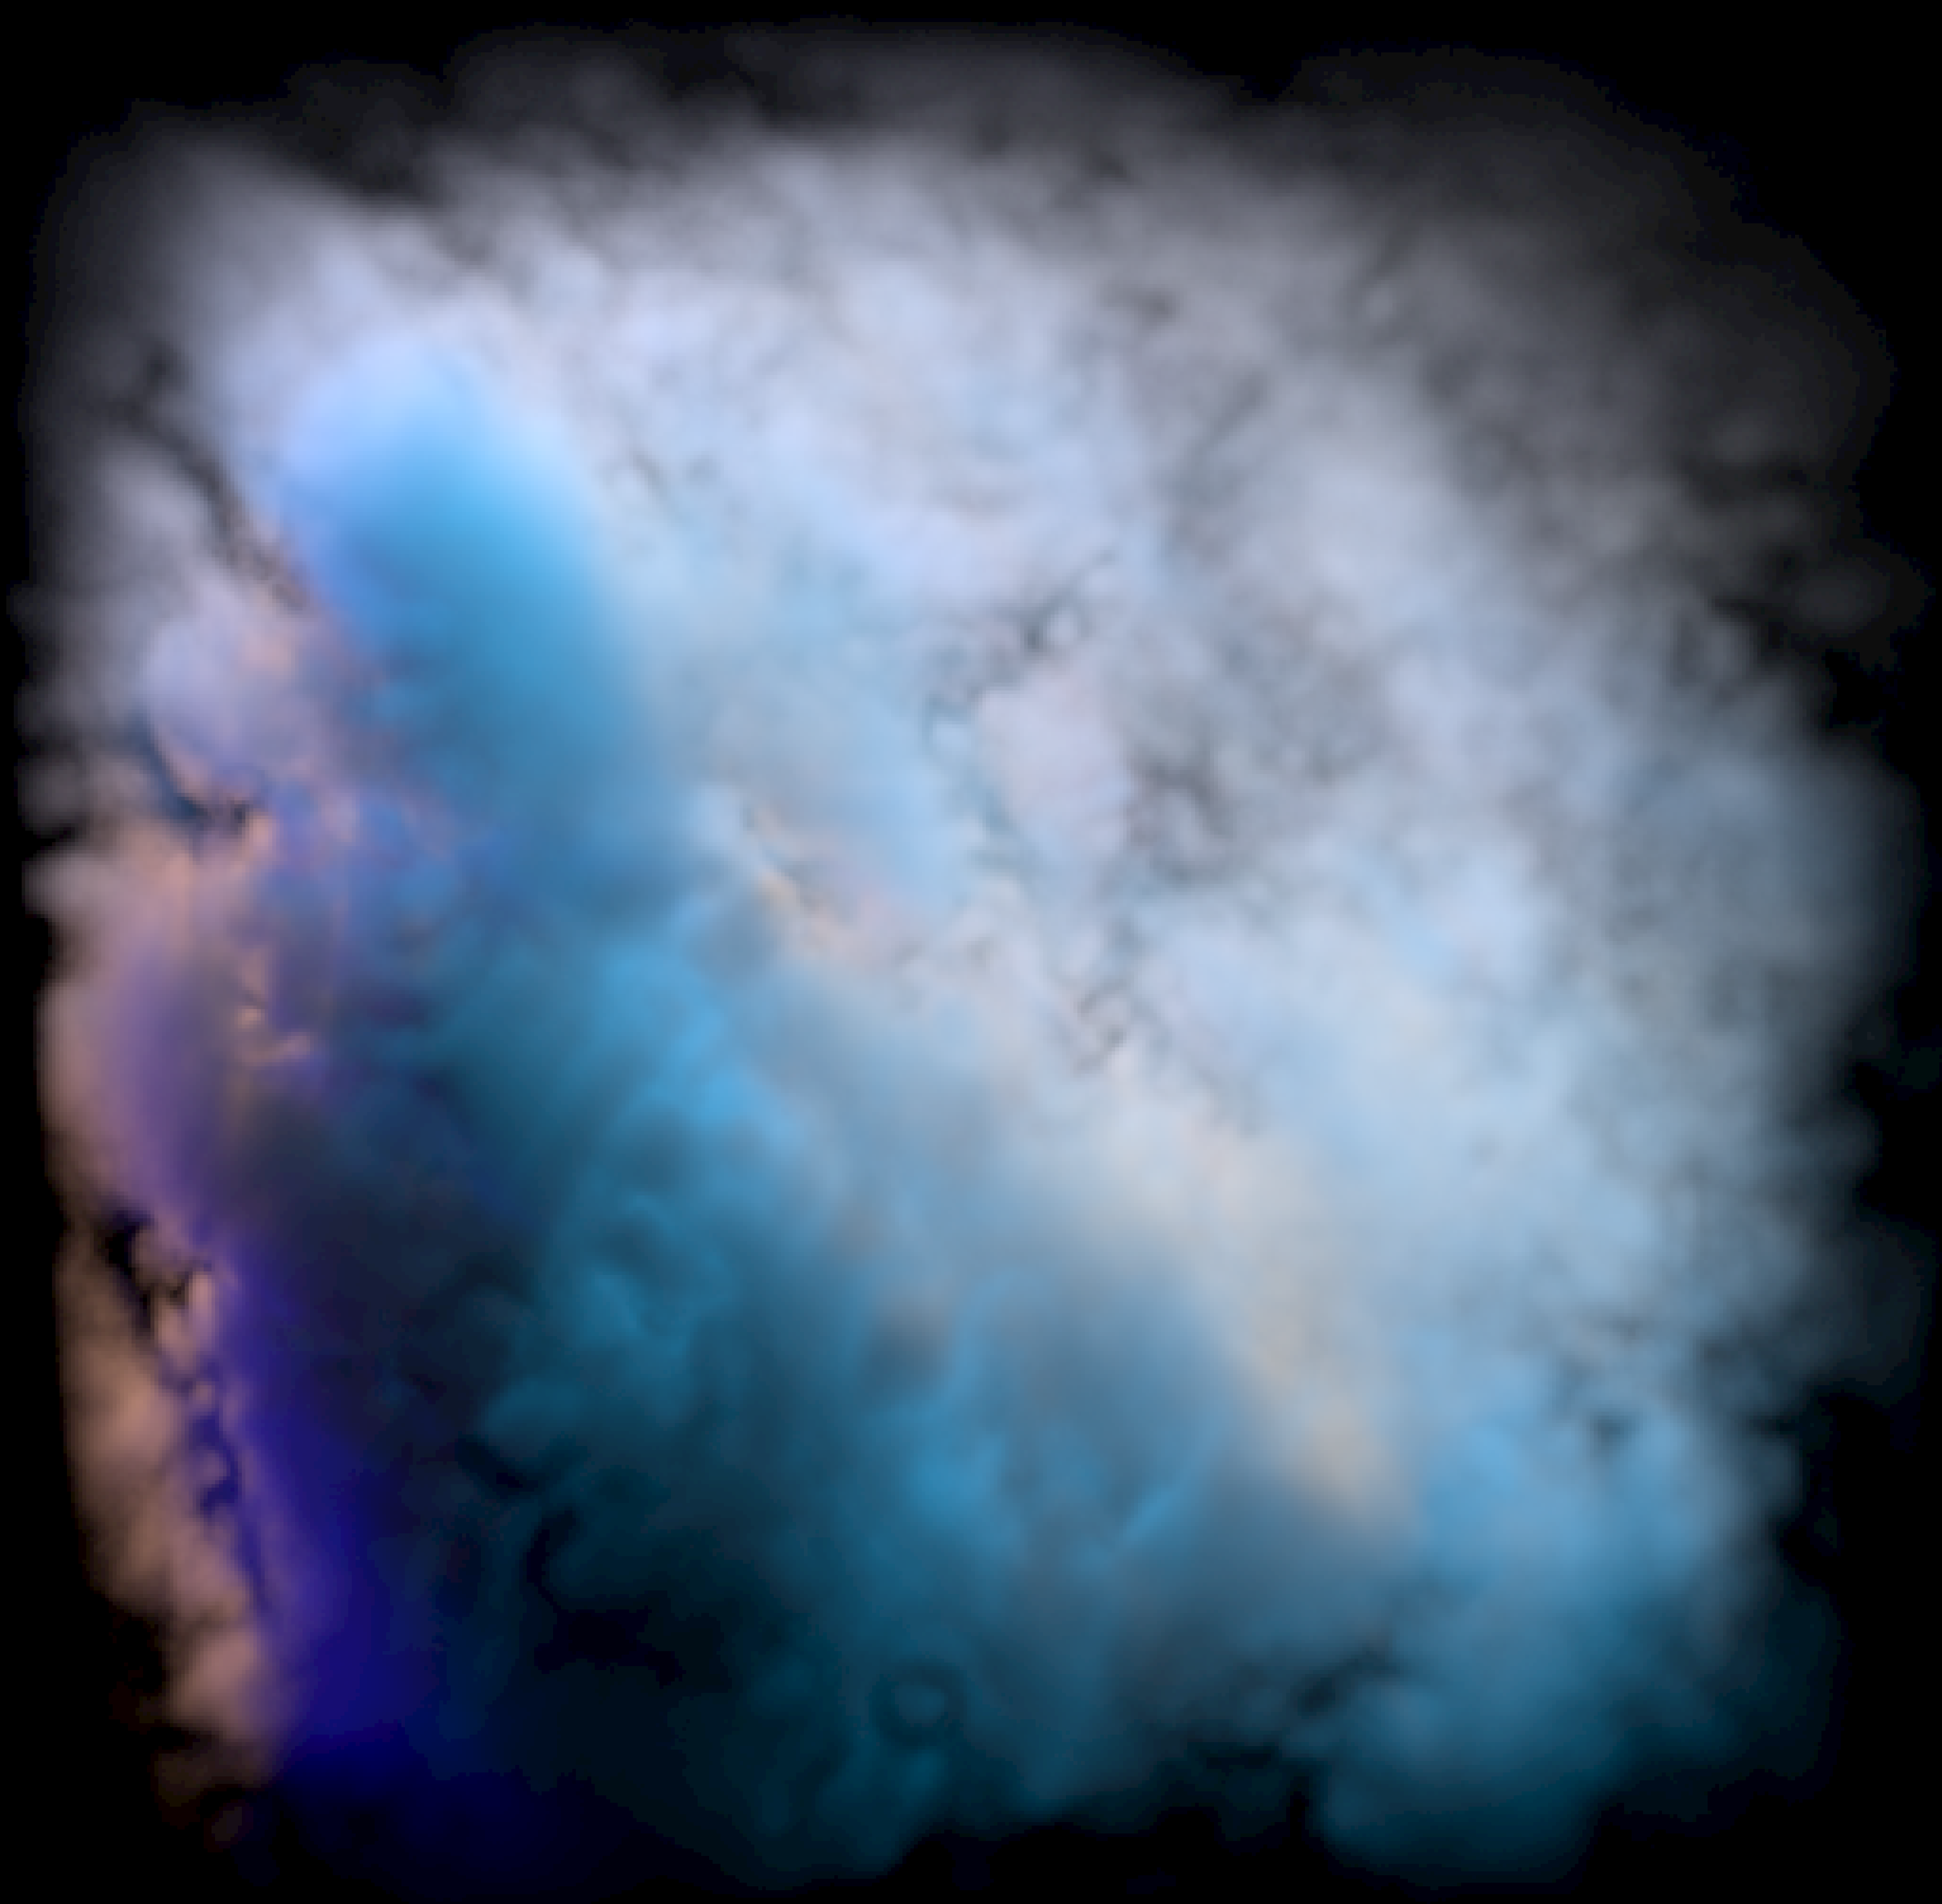

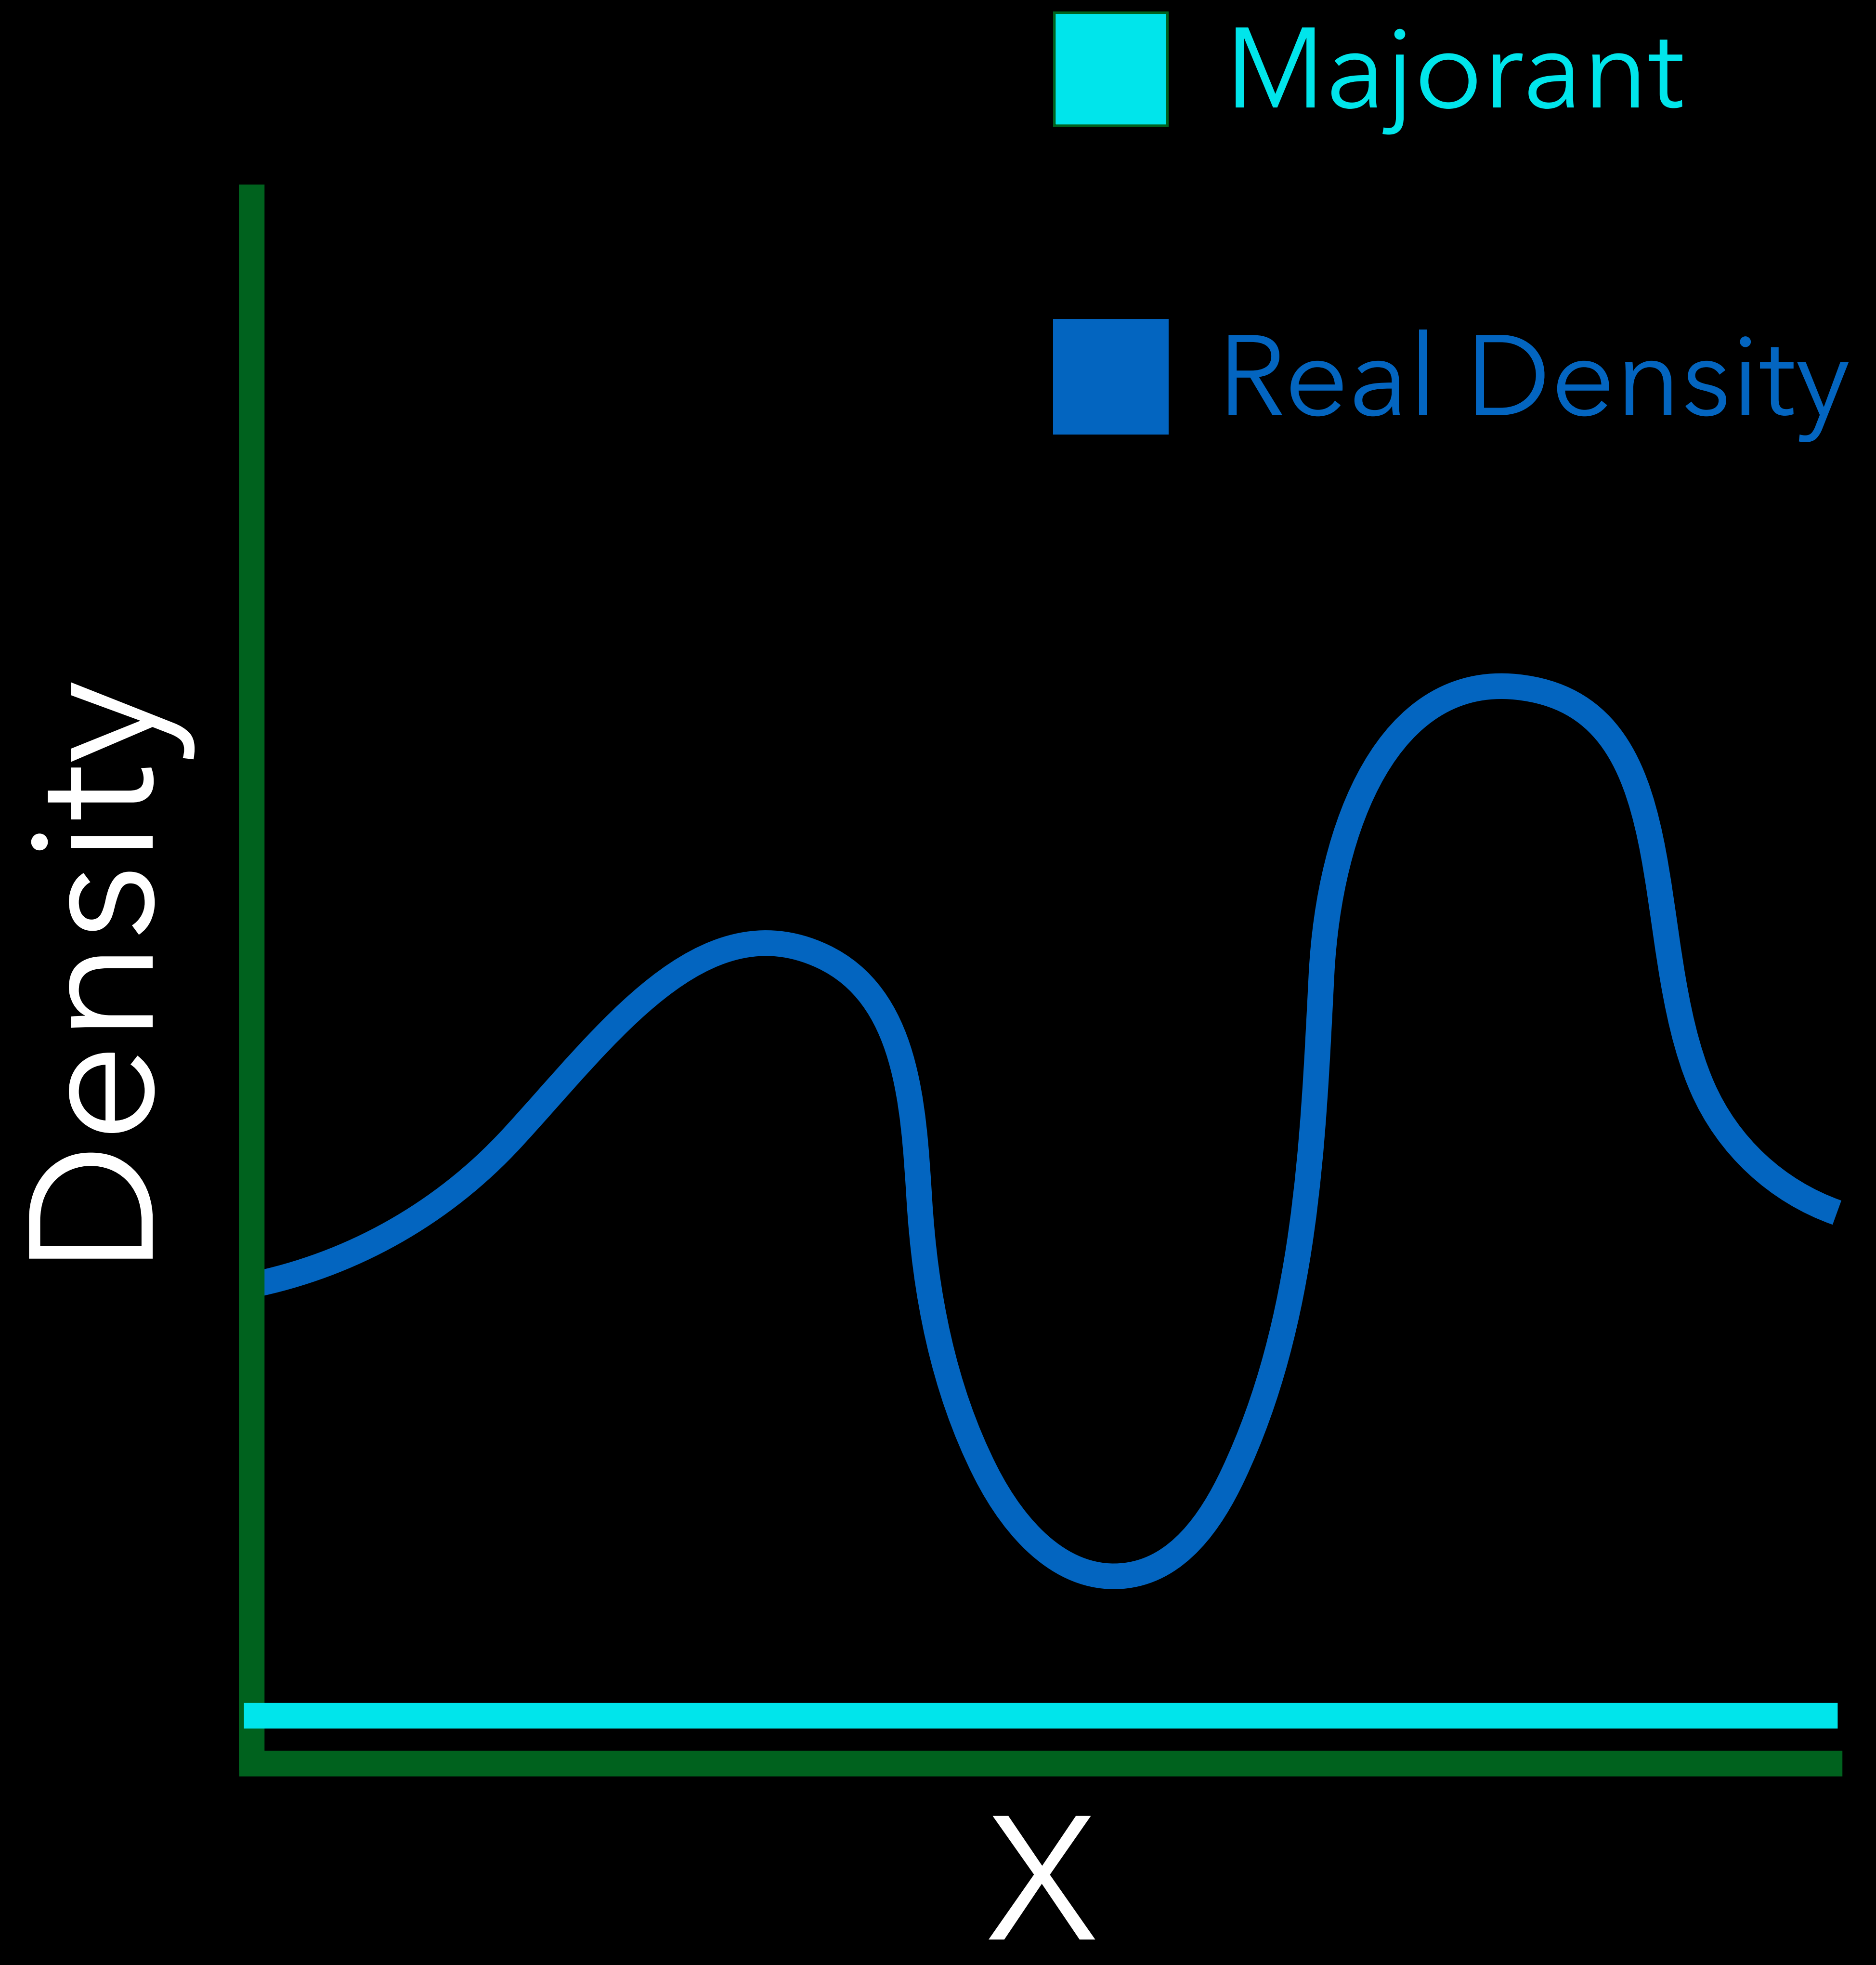
Majorant
Real Density
Density
X

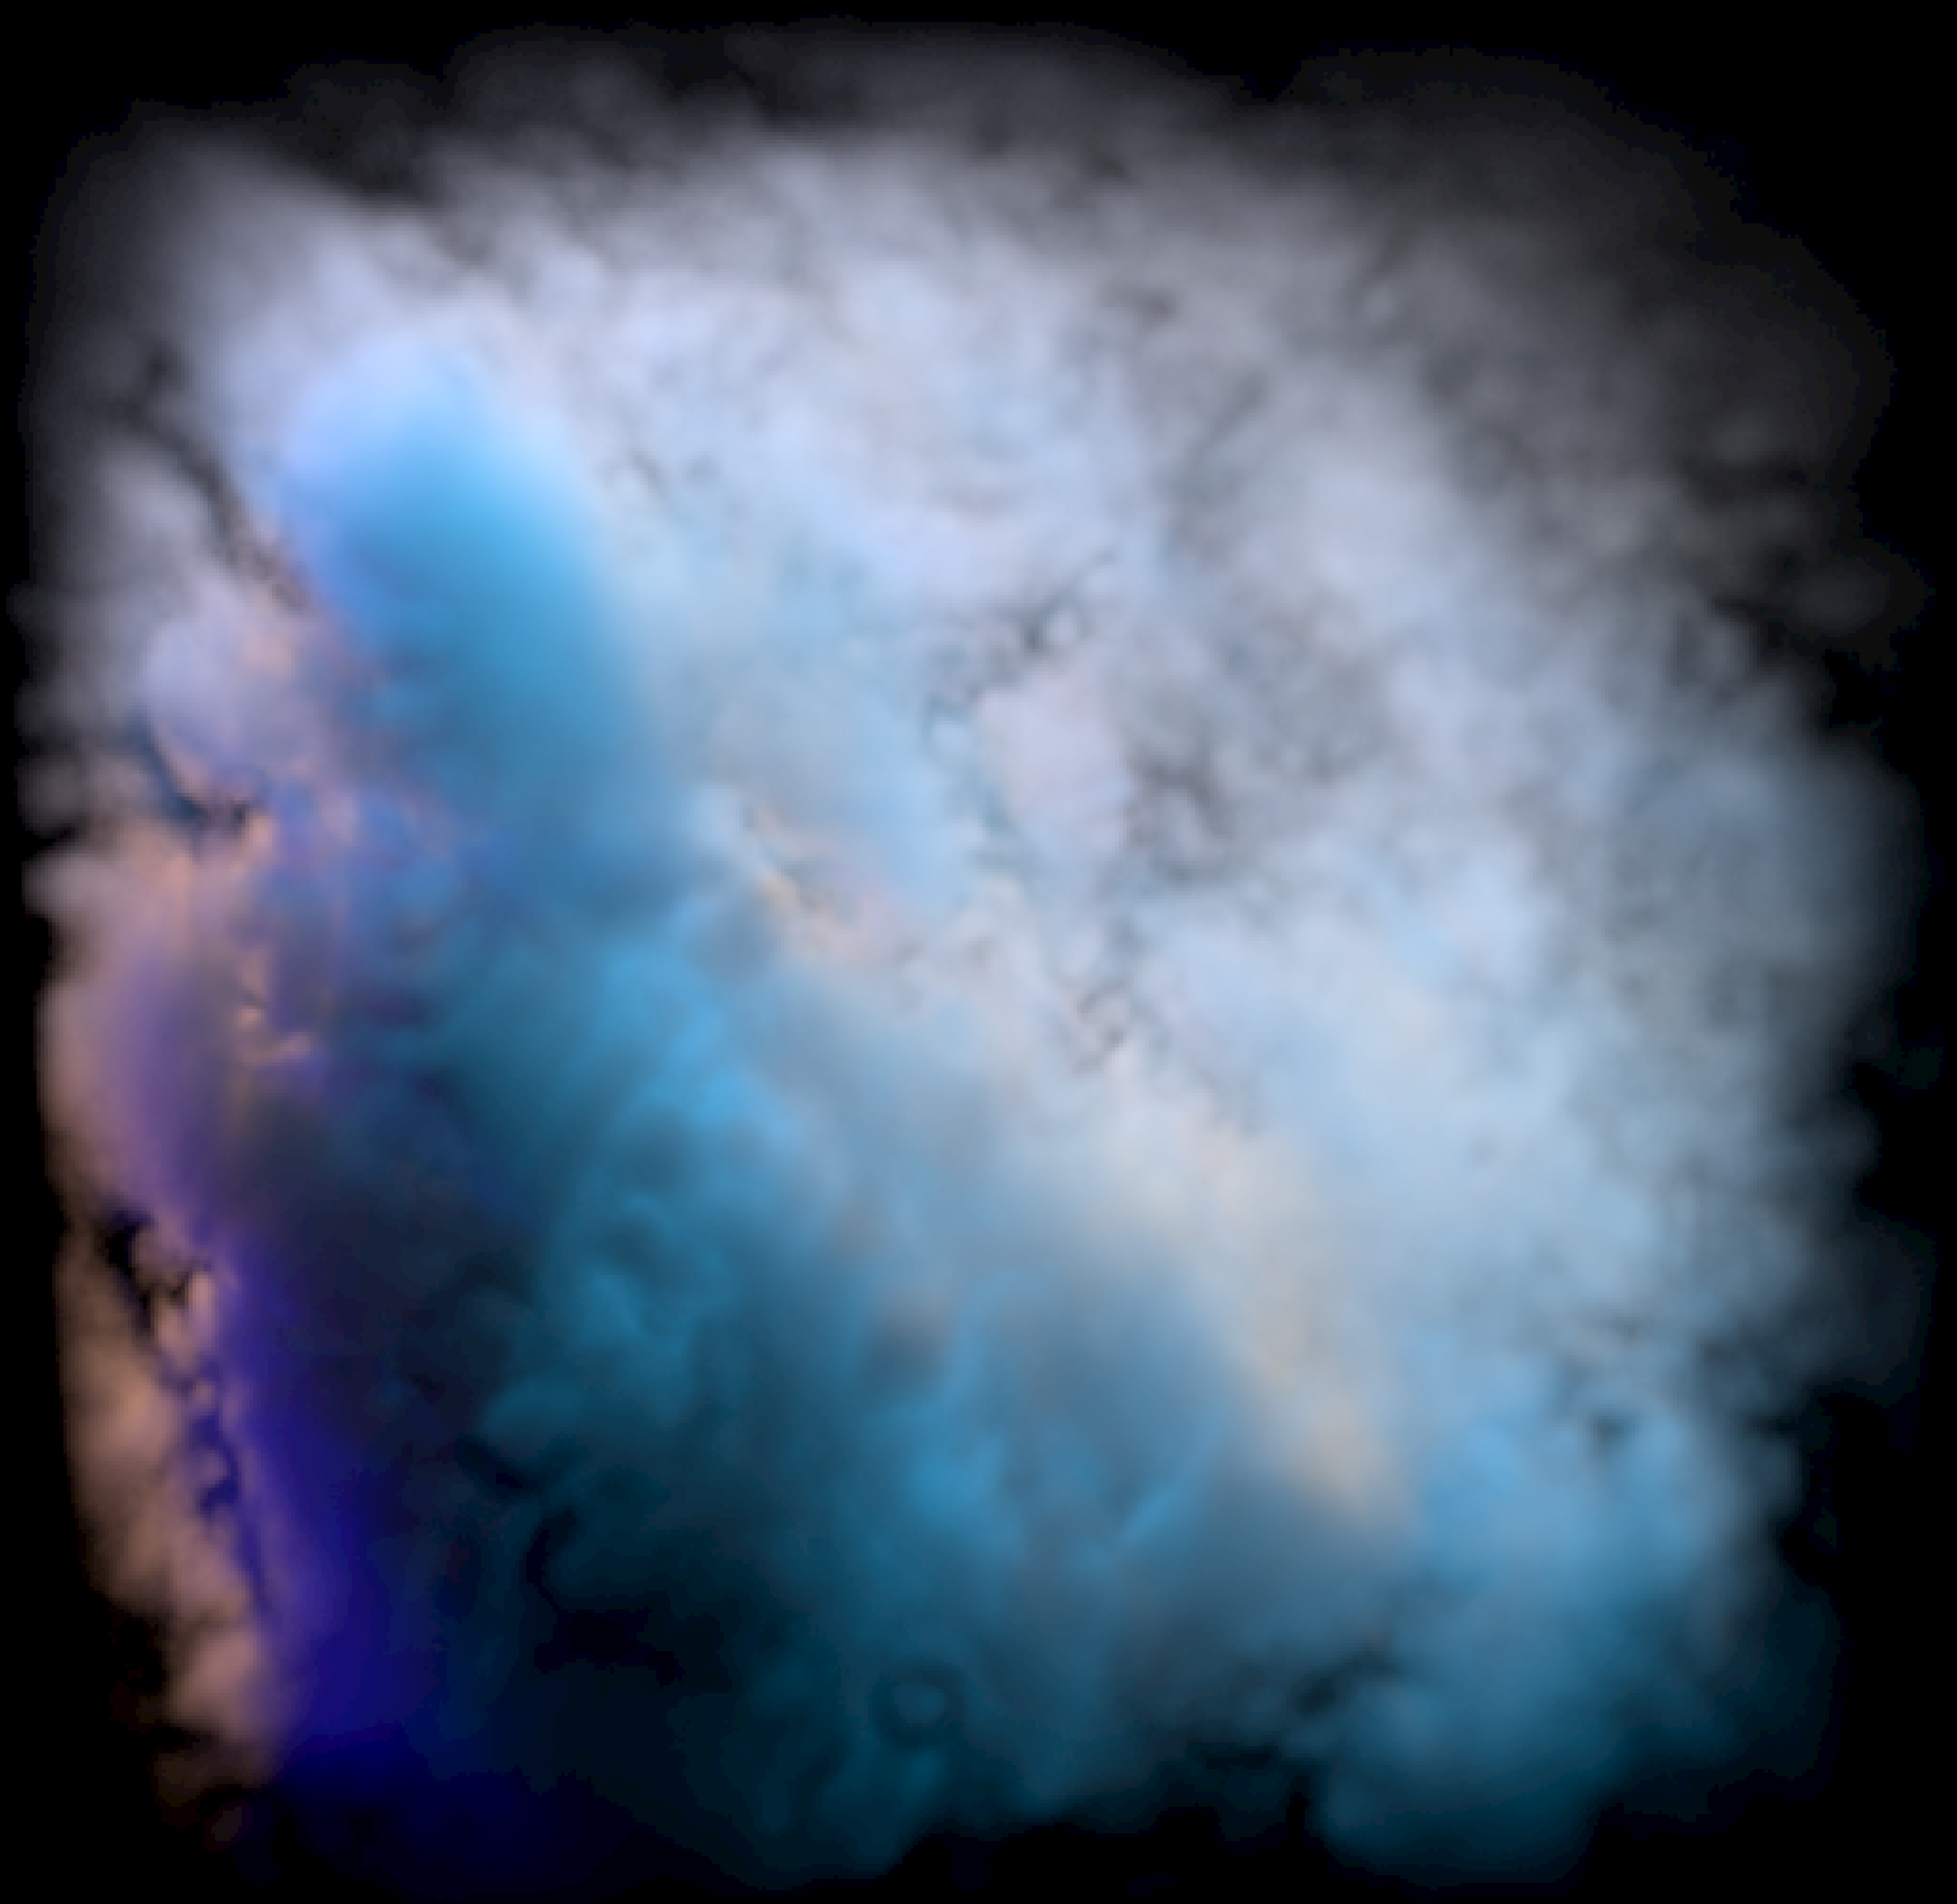

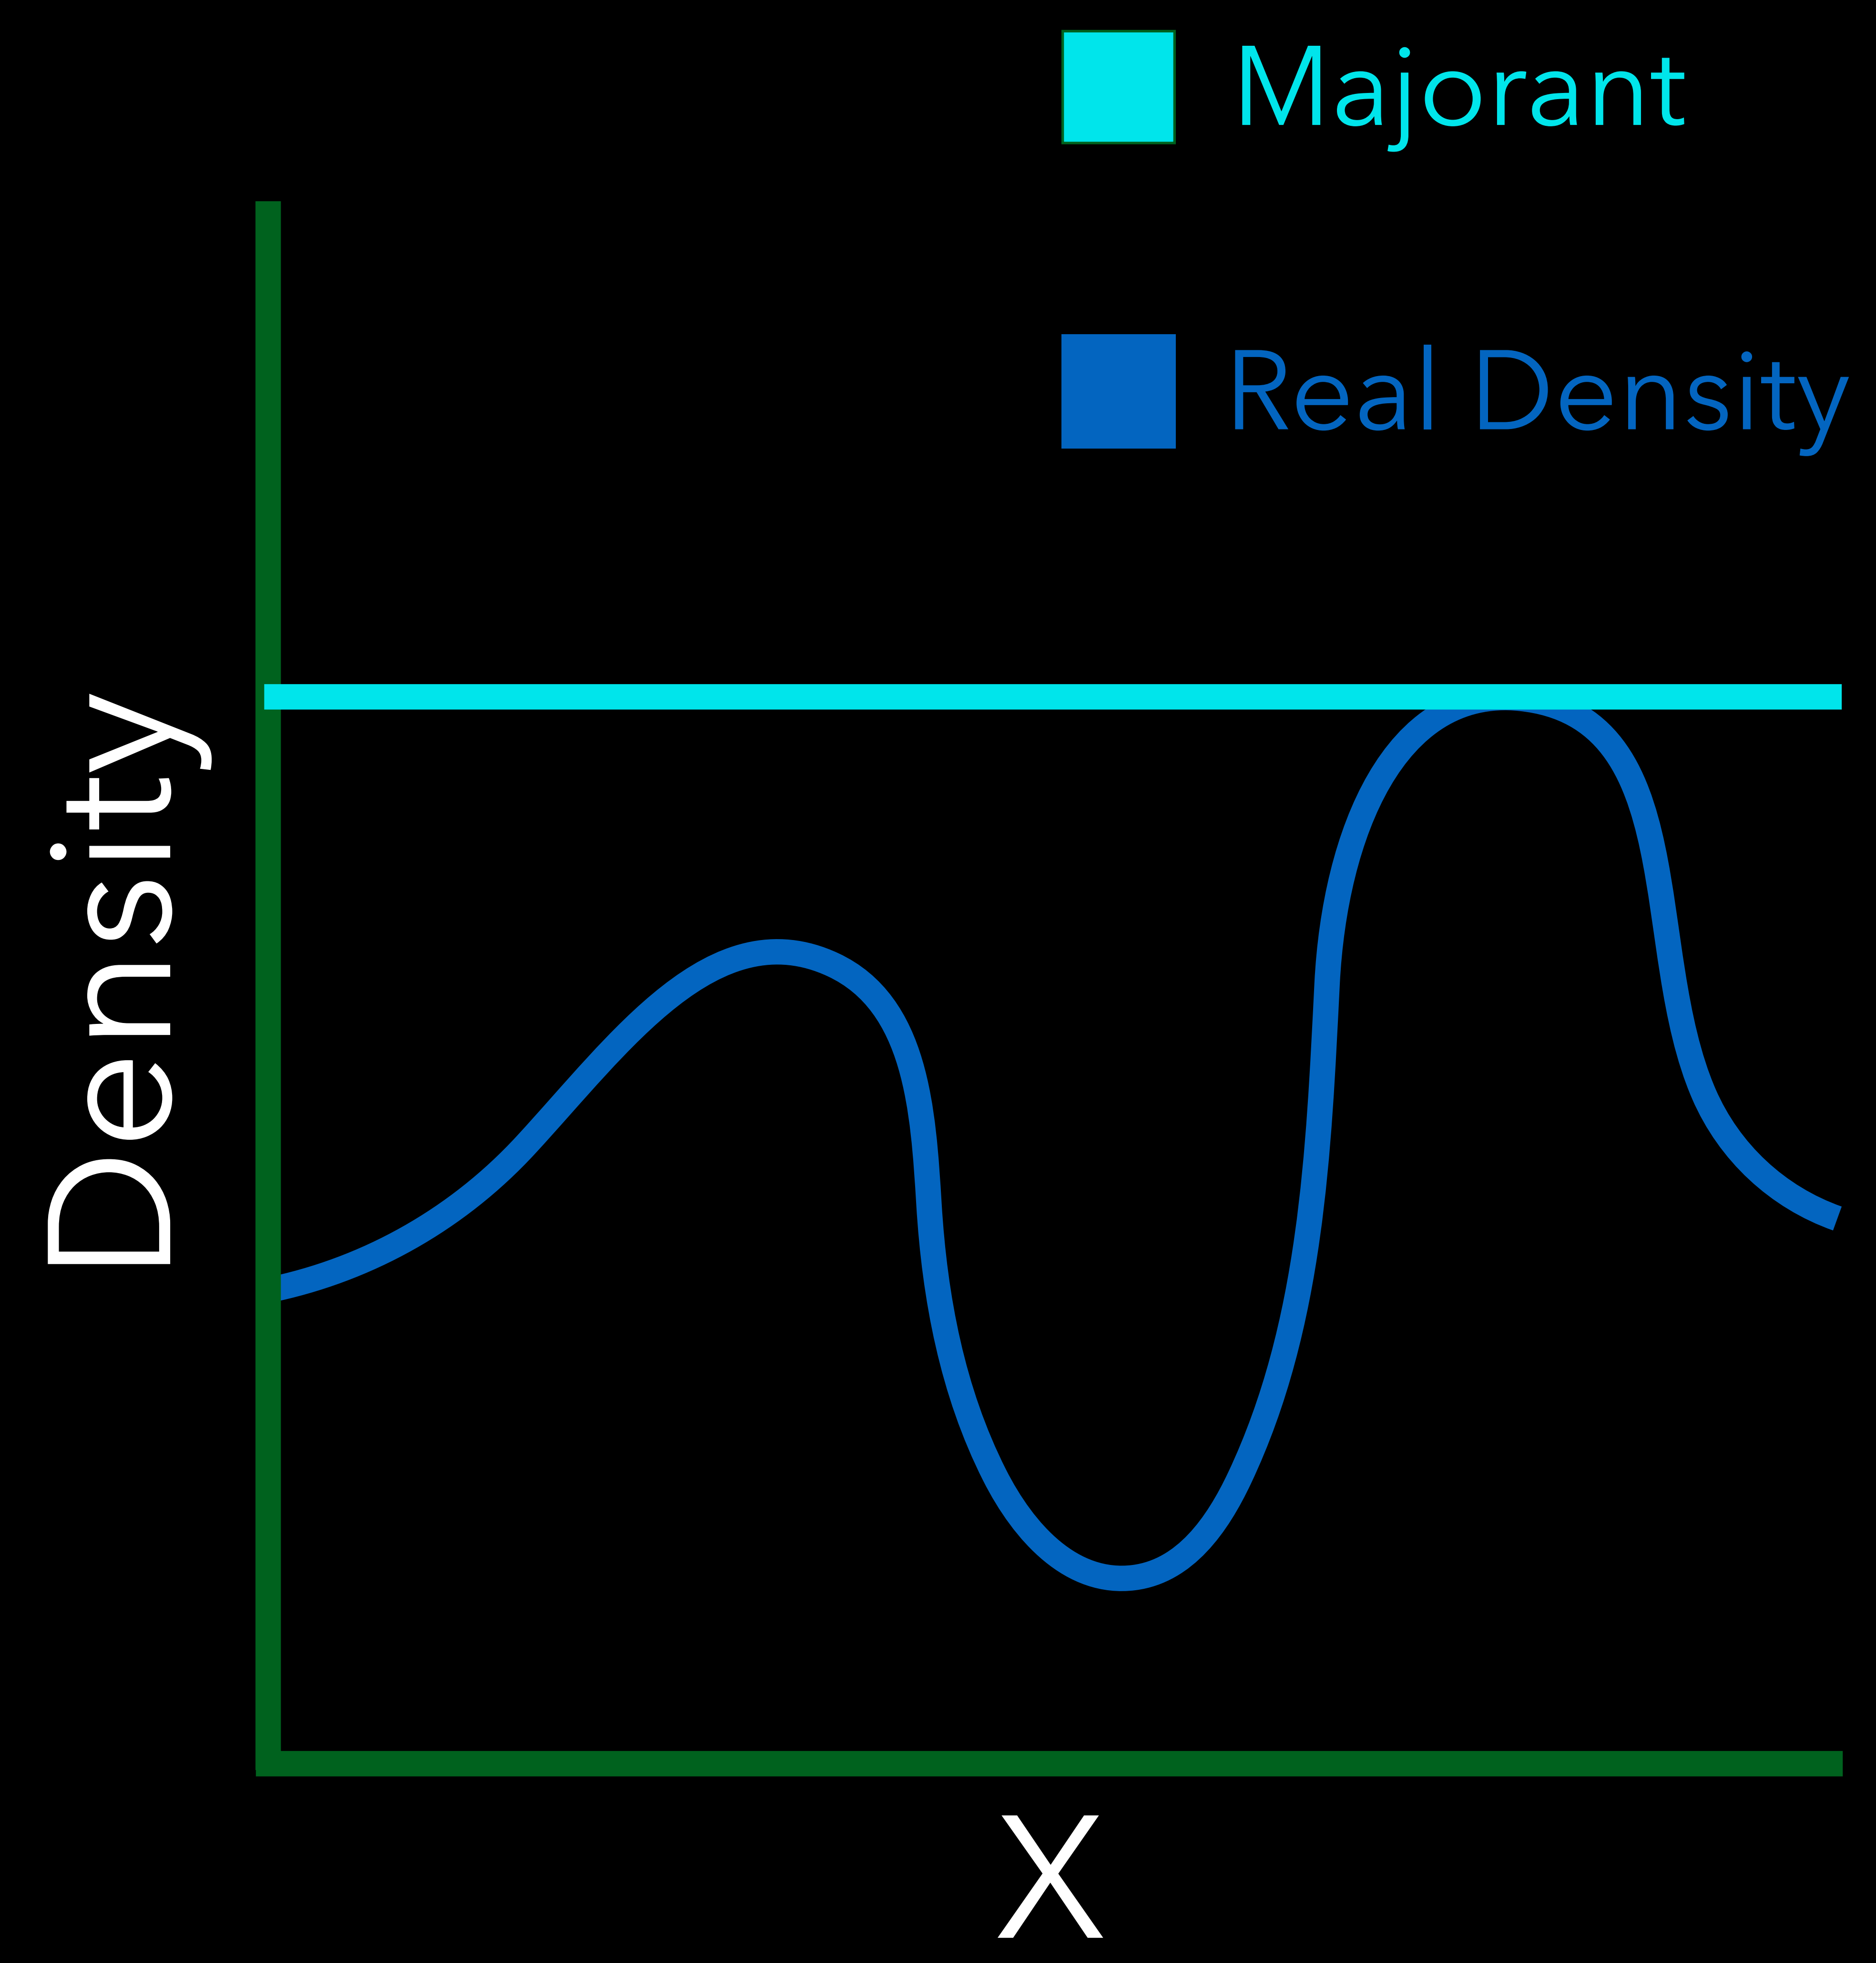
Majorant
Real Density
Density
X

Volume representation

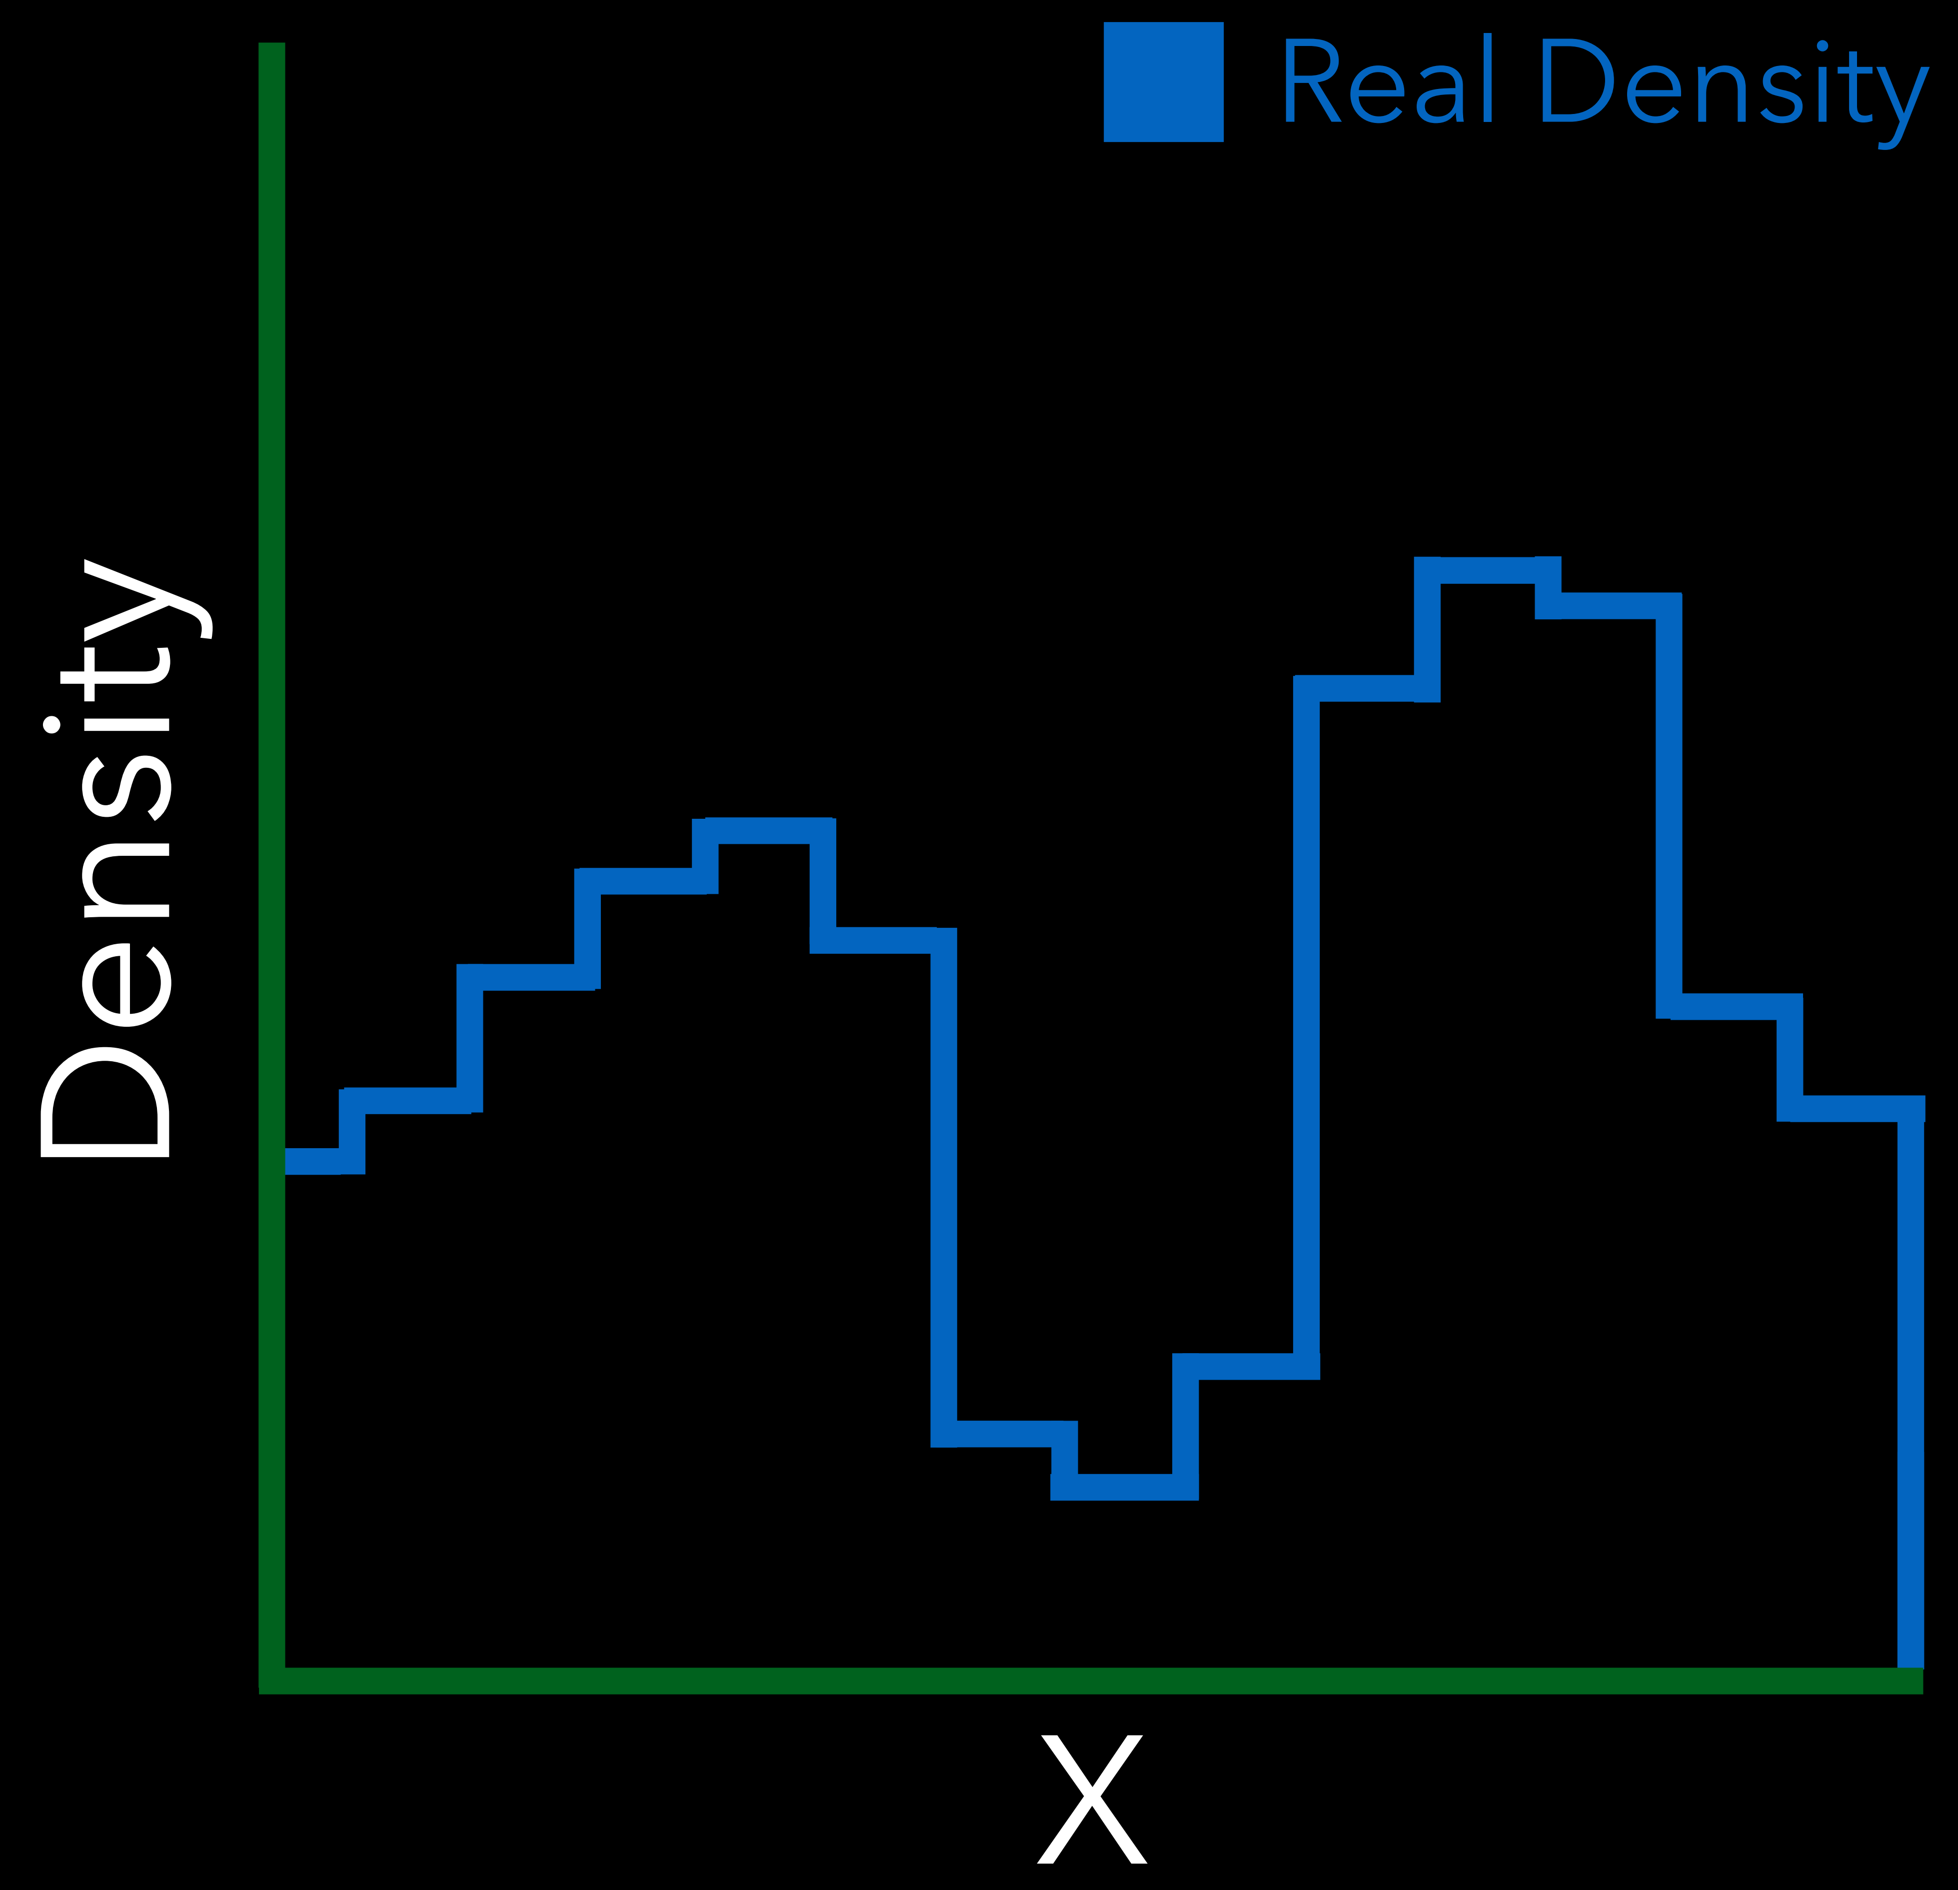
Real Density
Density
X

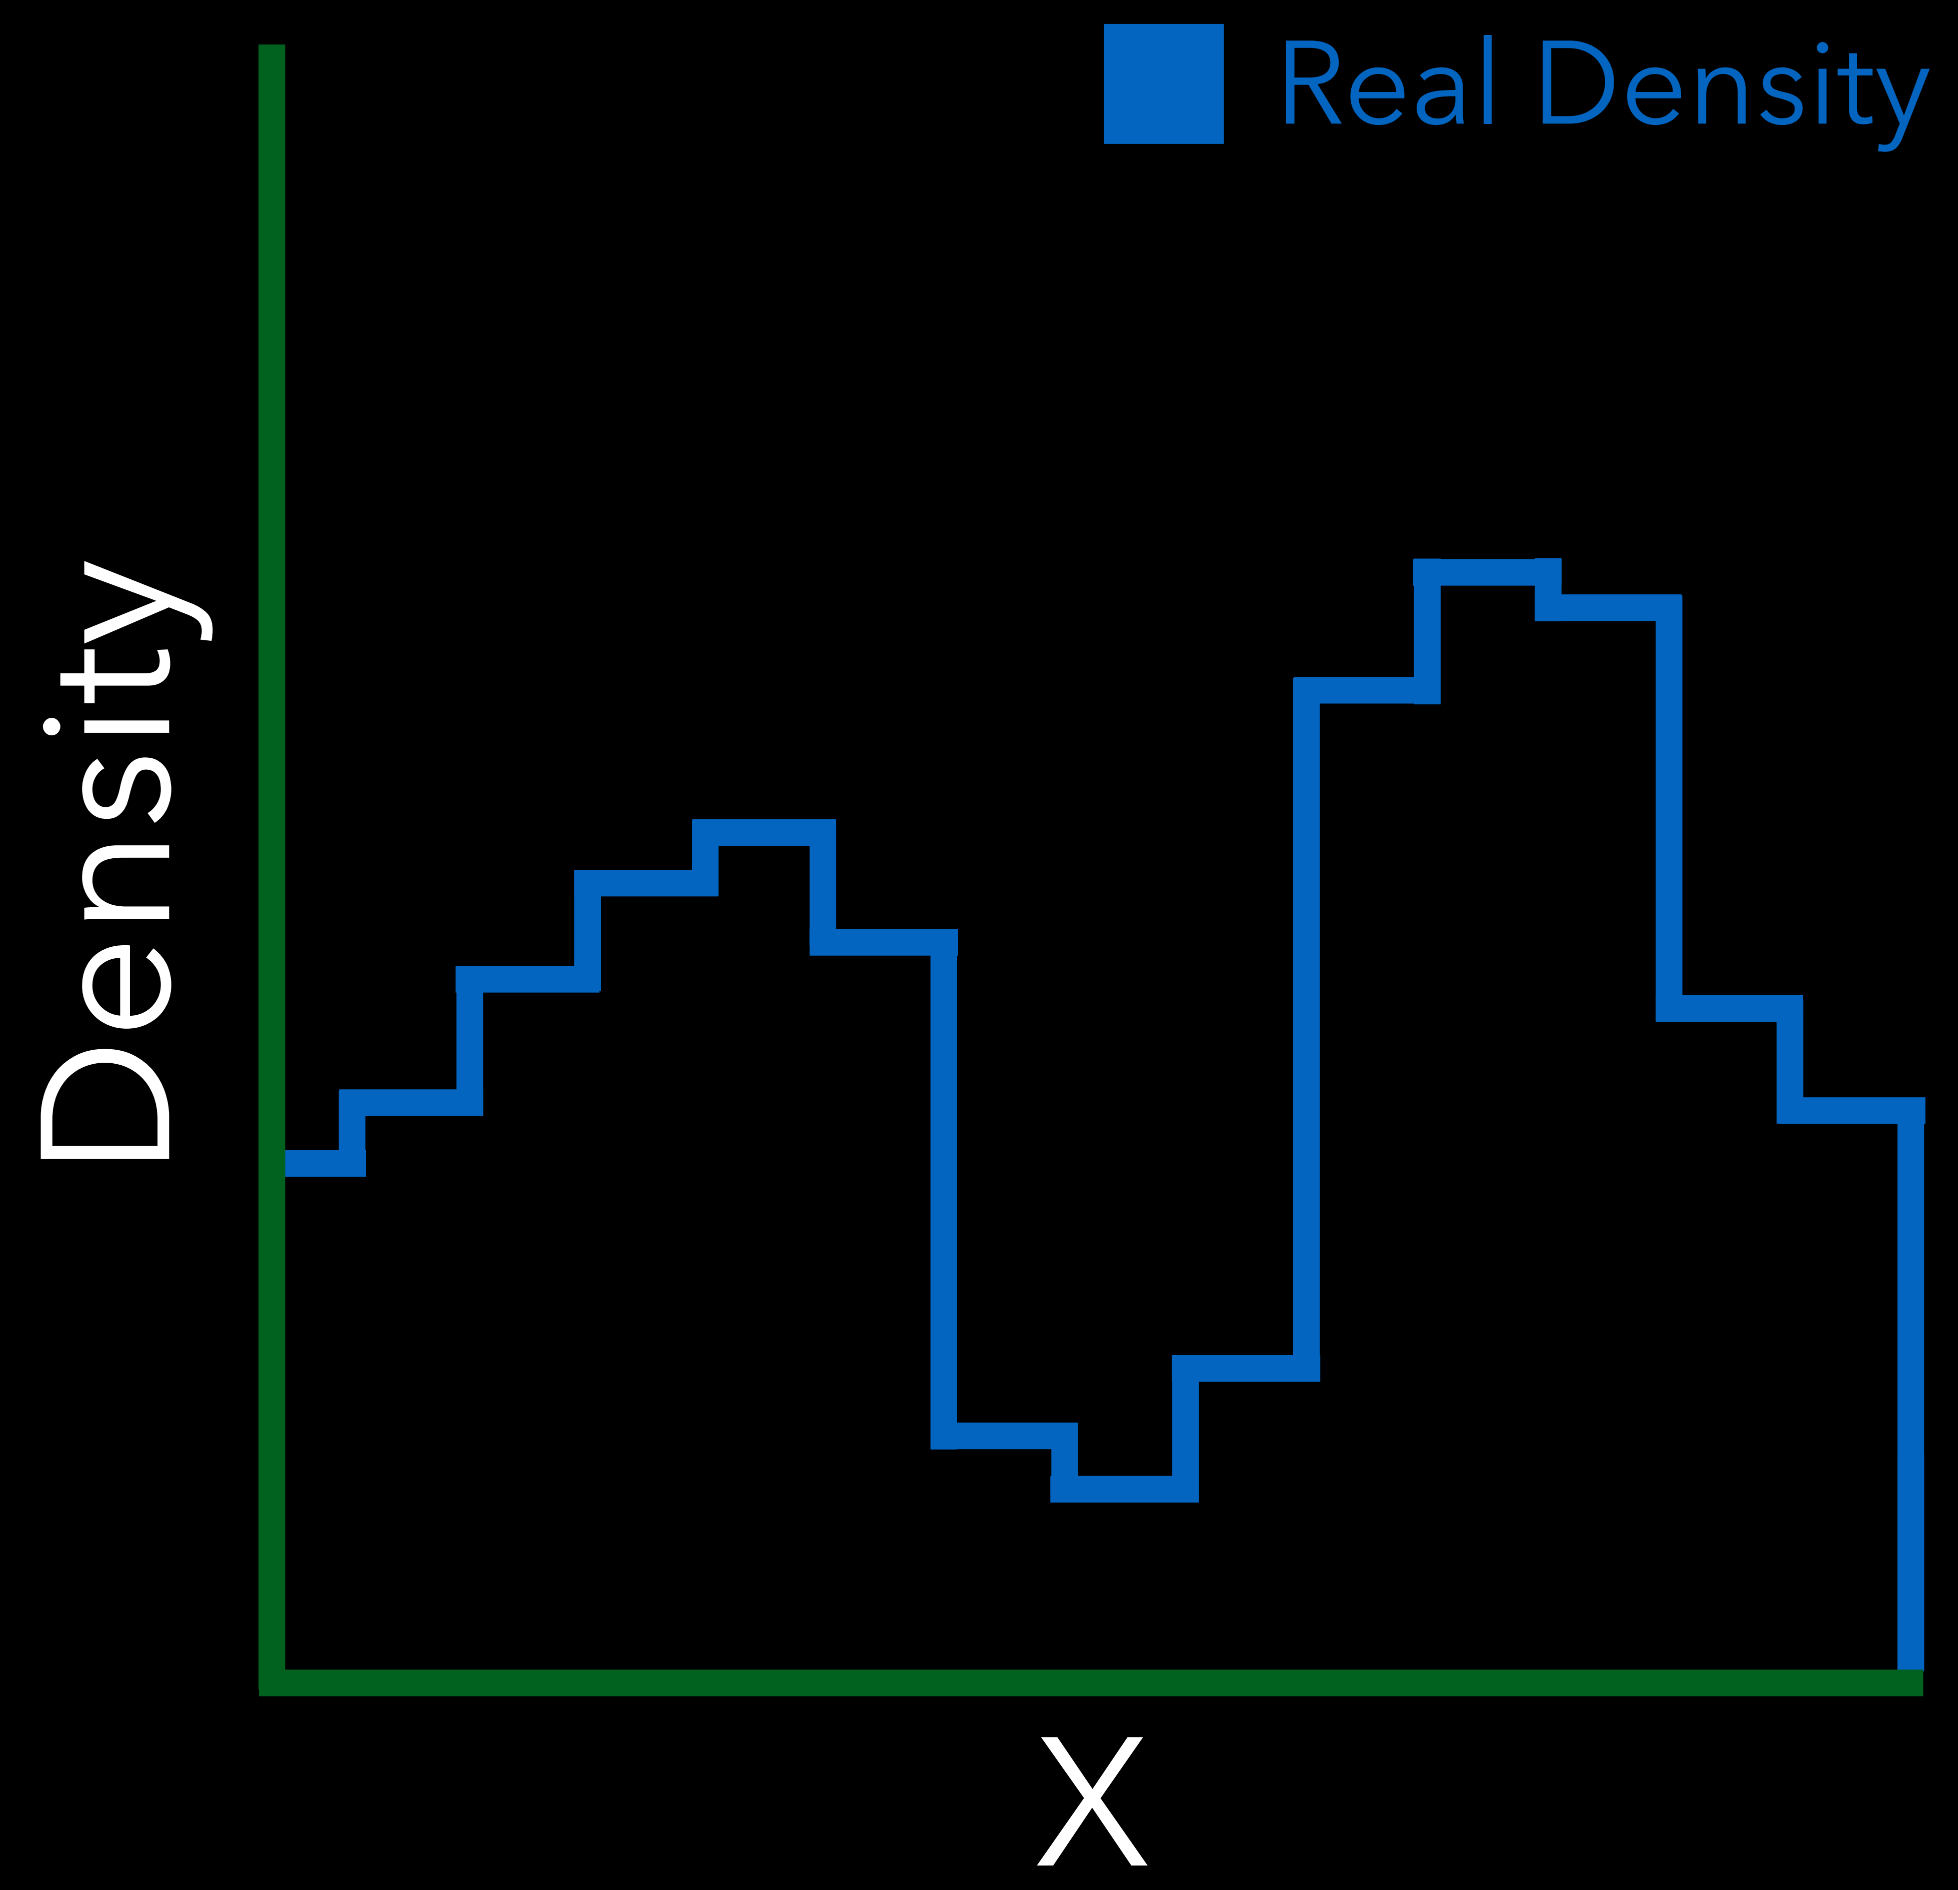
Real Density
Density
X

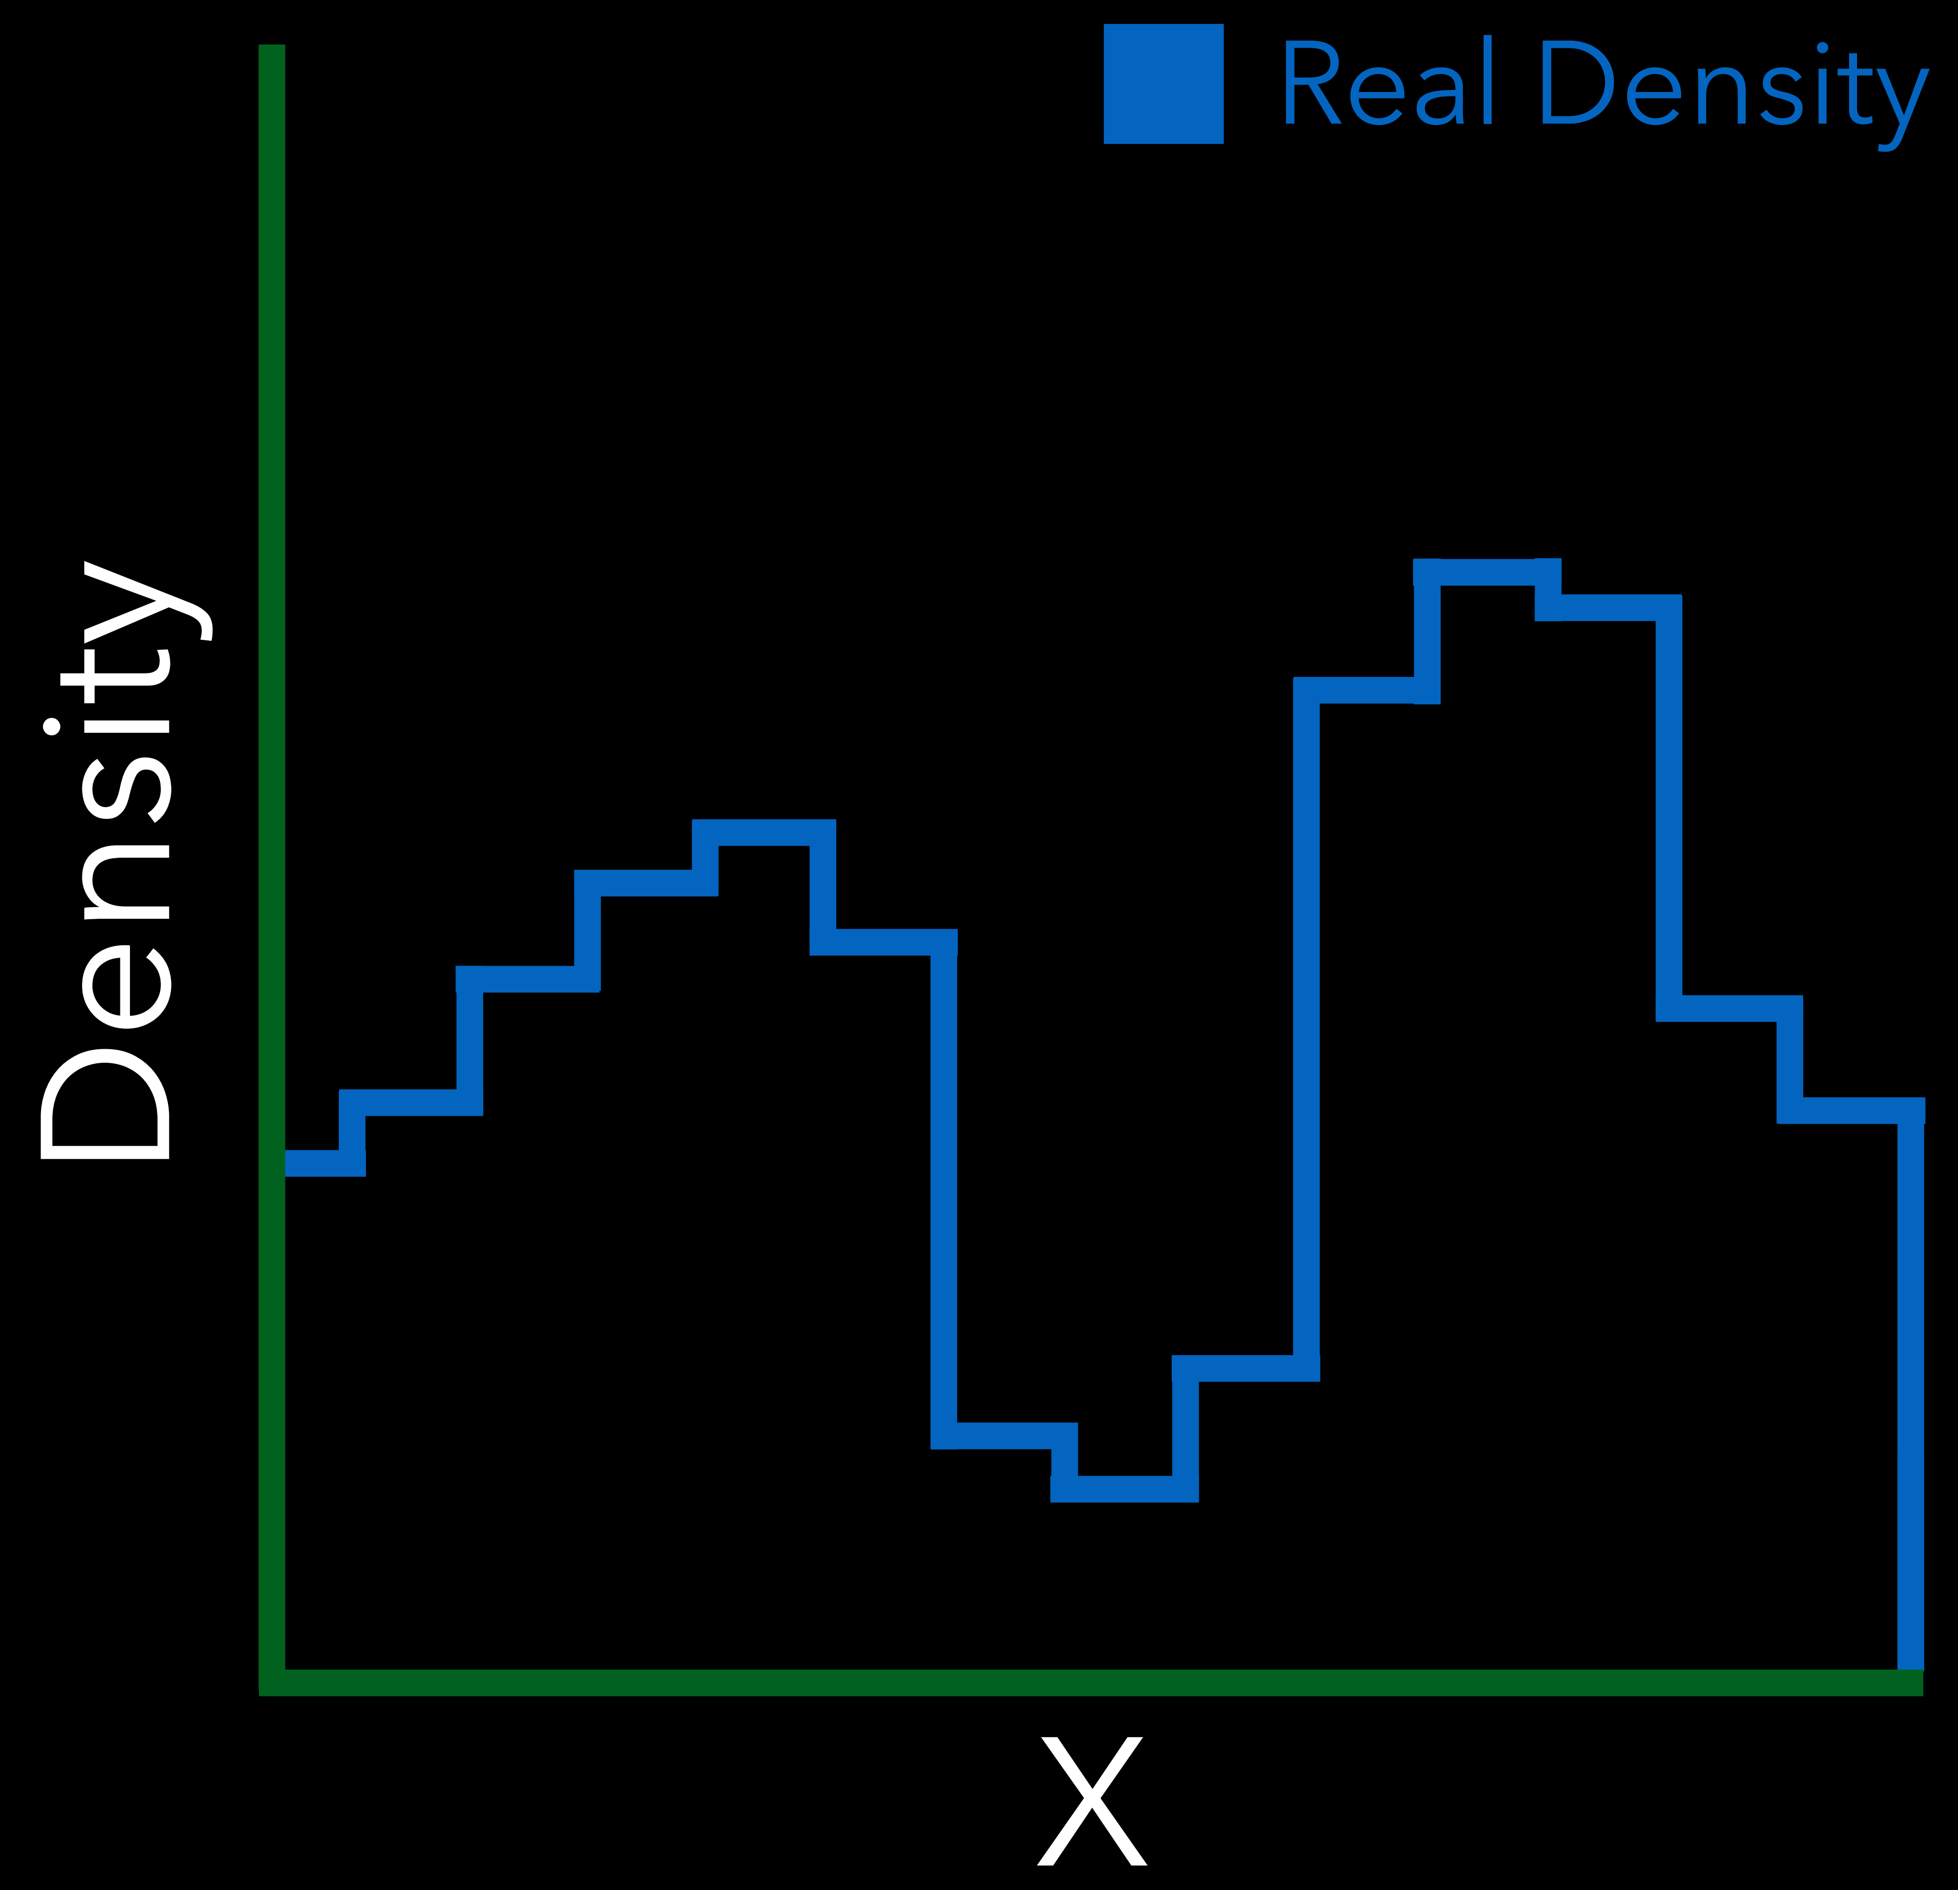
Real Density
Density
X

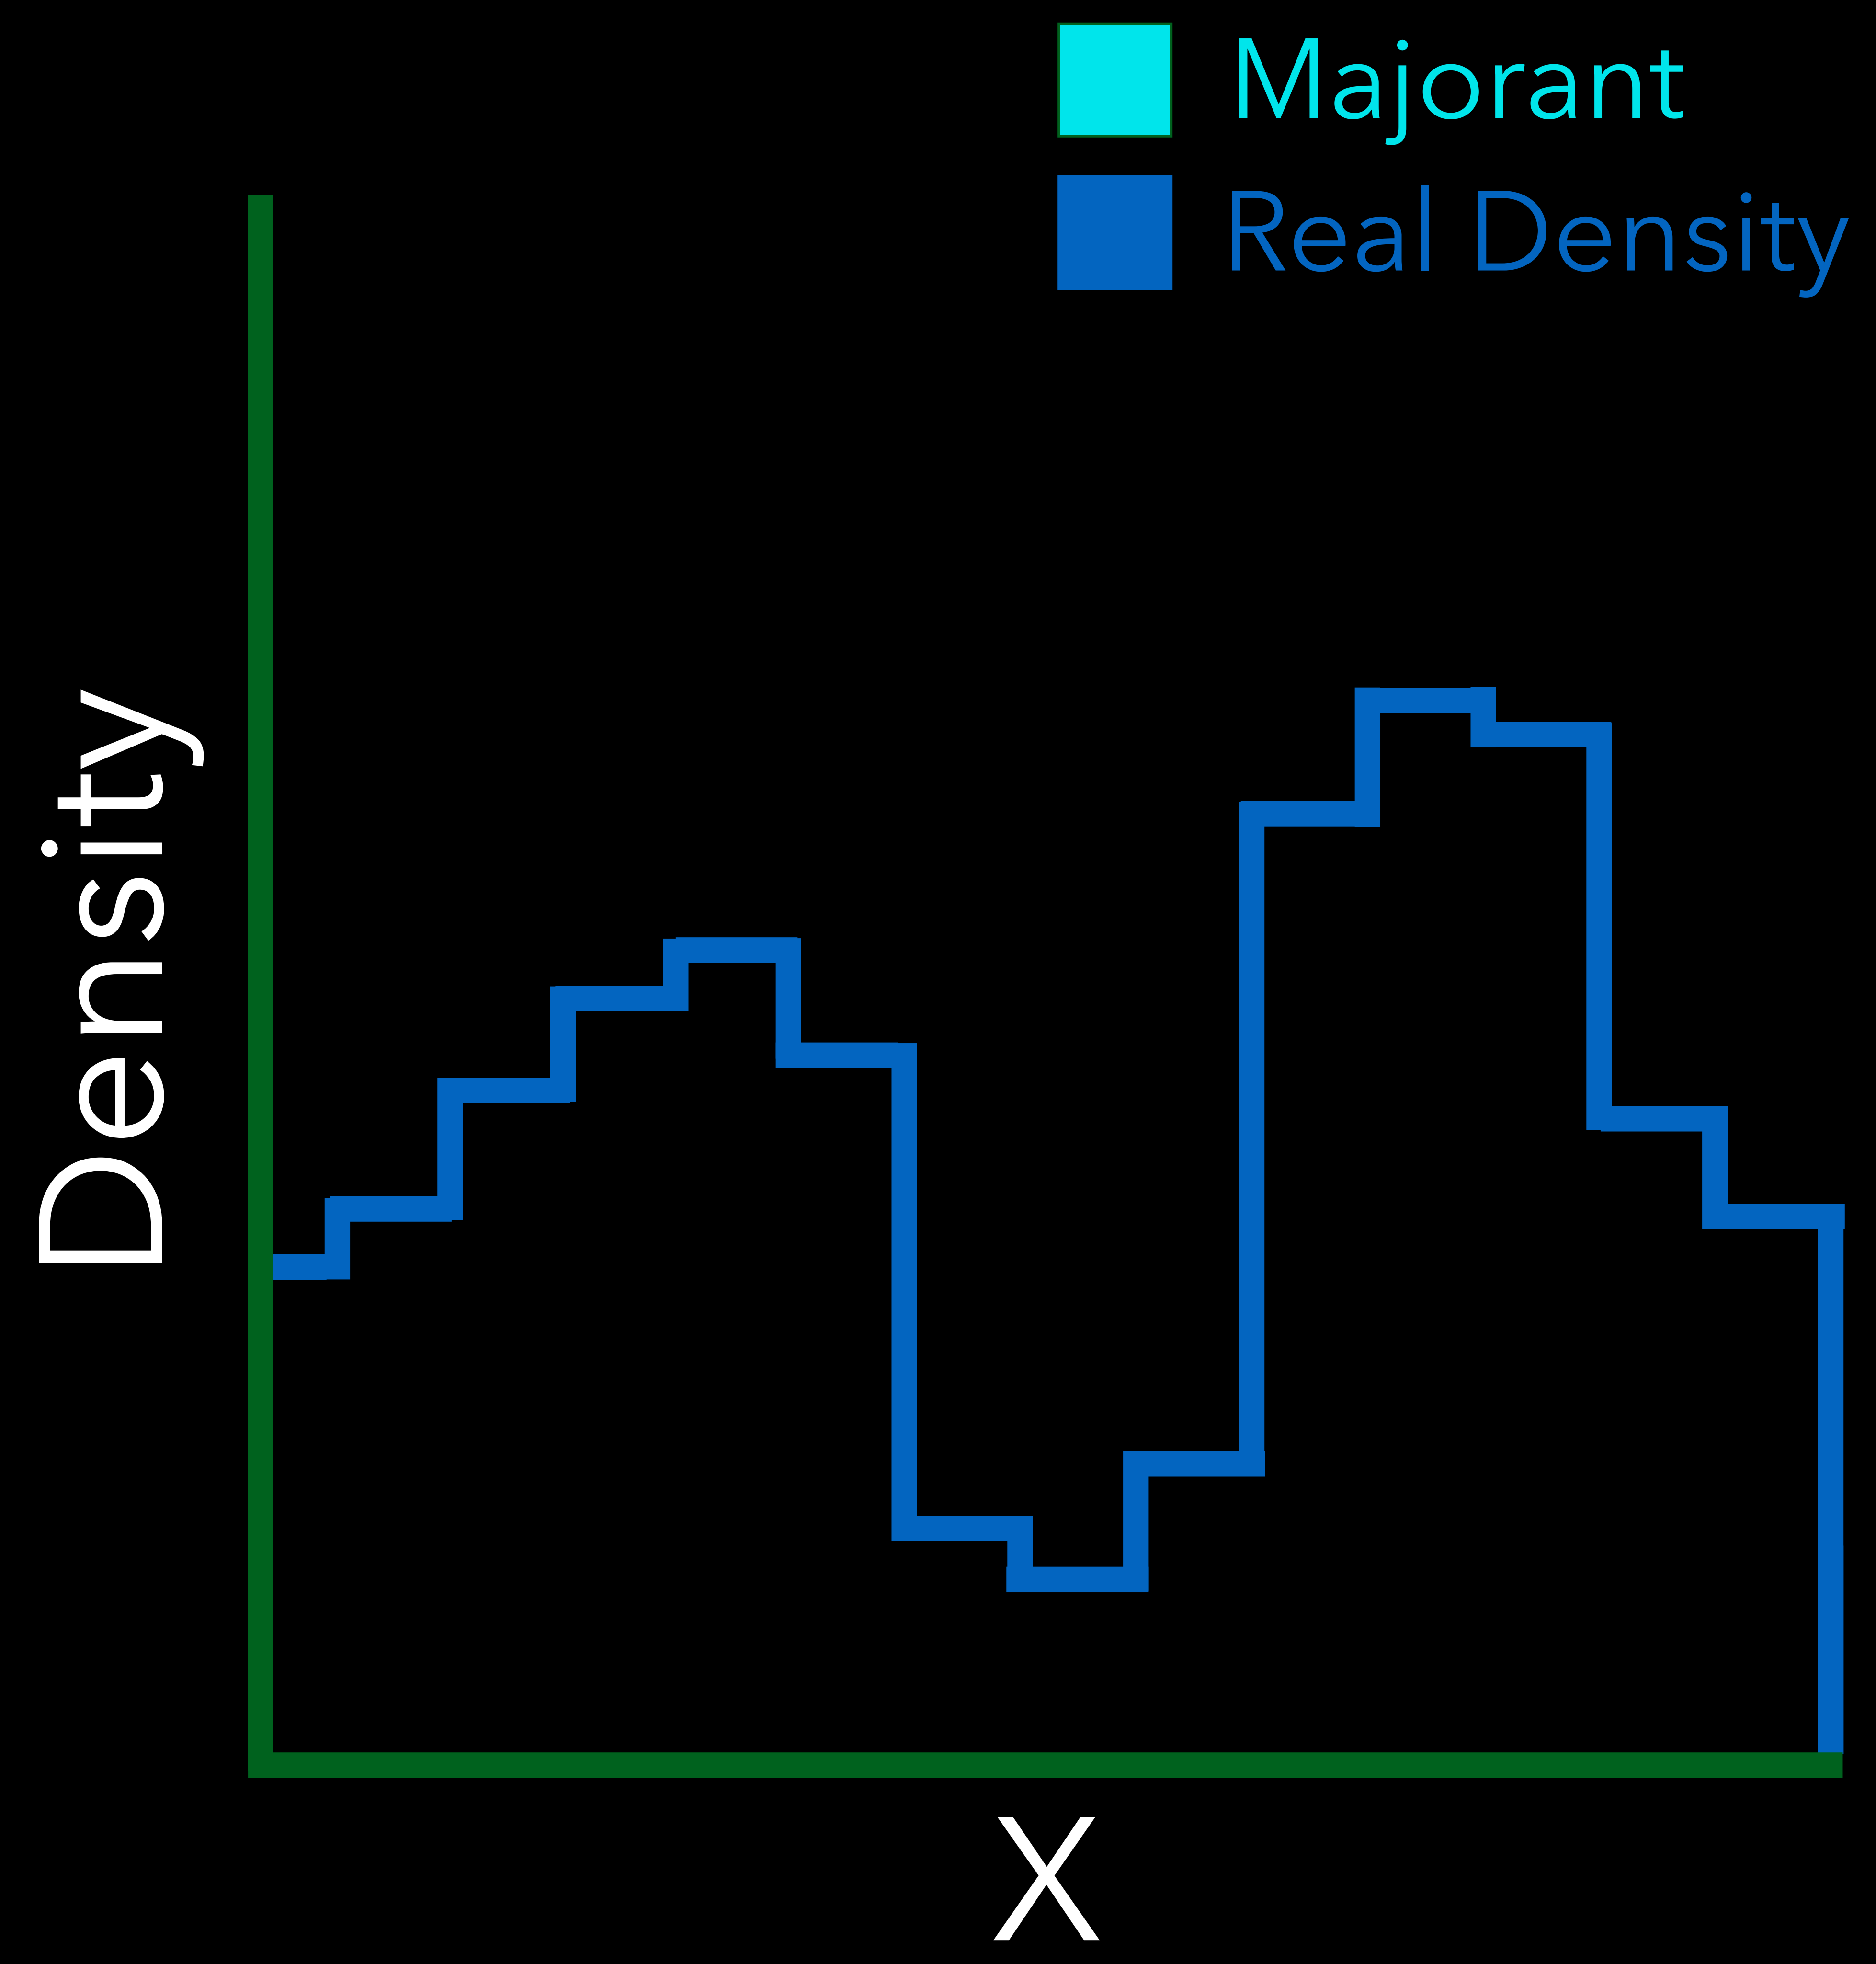
Majorant
Real Density
Density
X

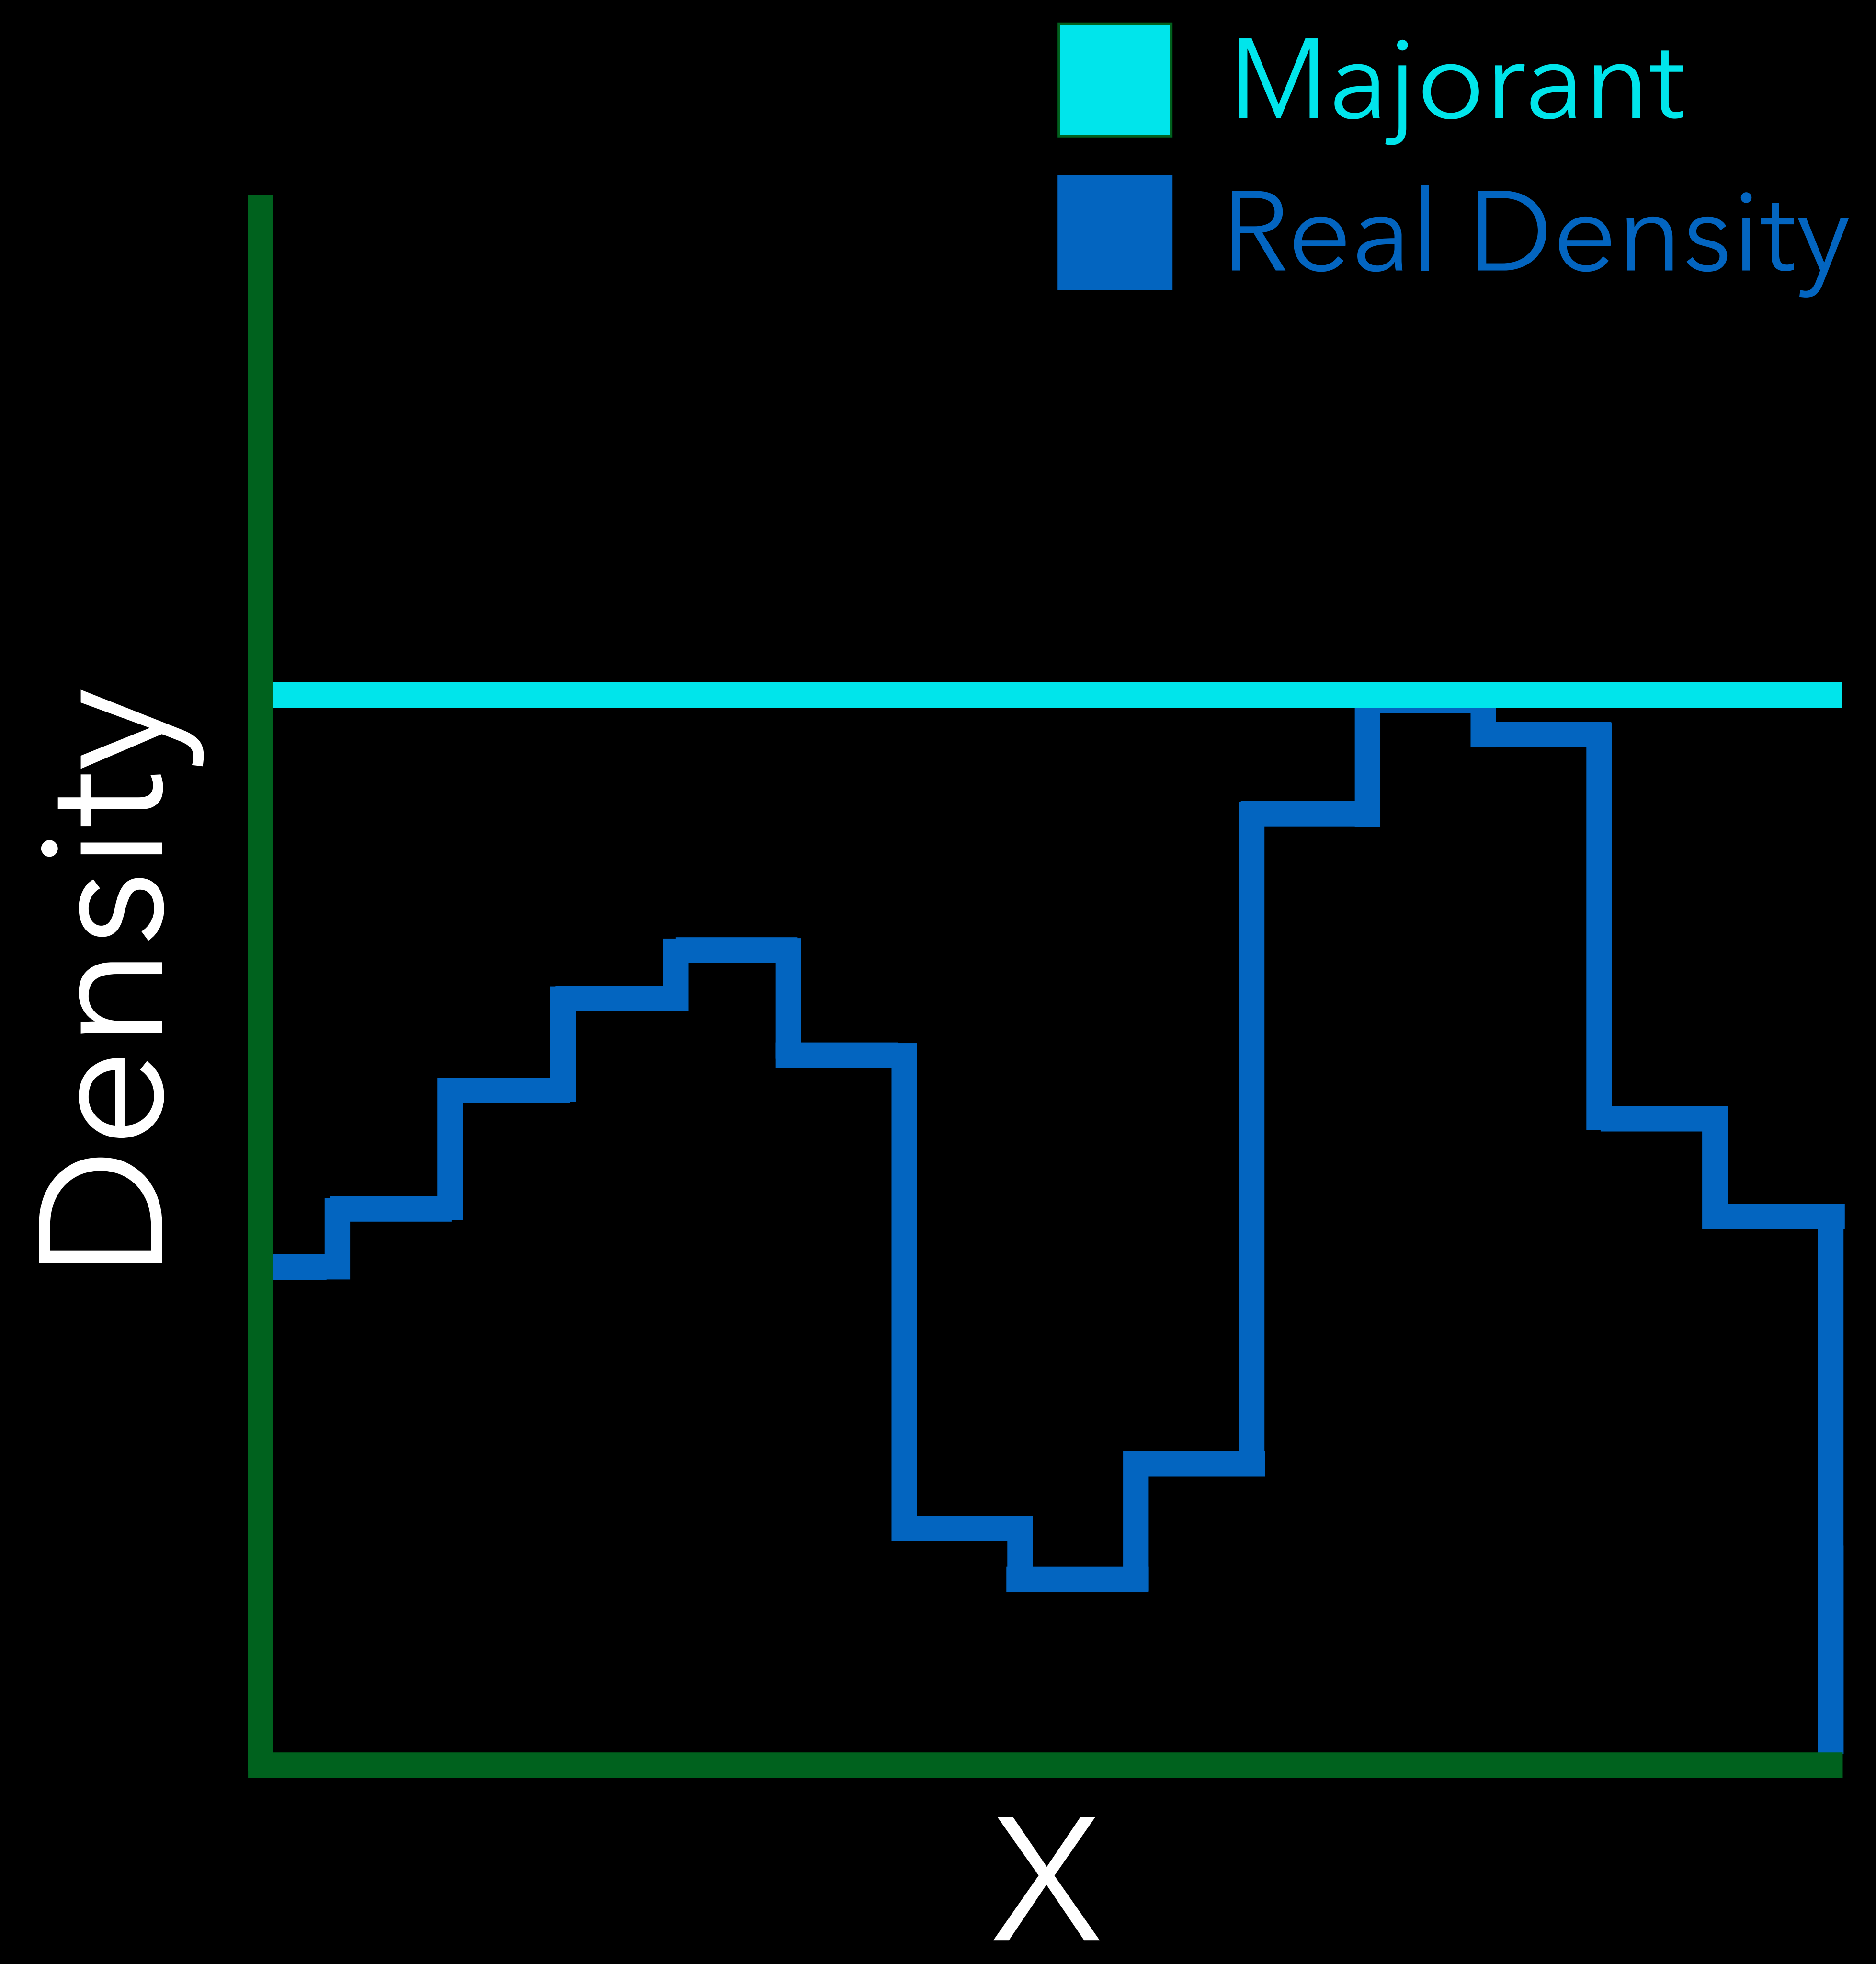
Majorant
Real Density
Density
X

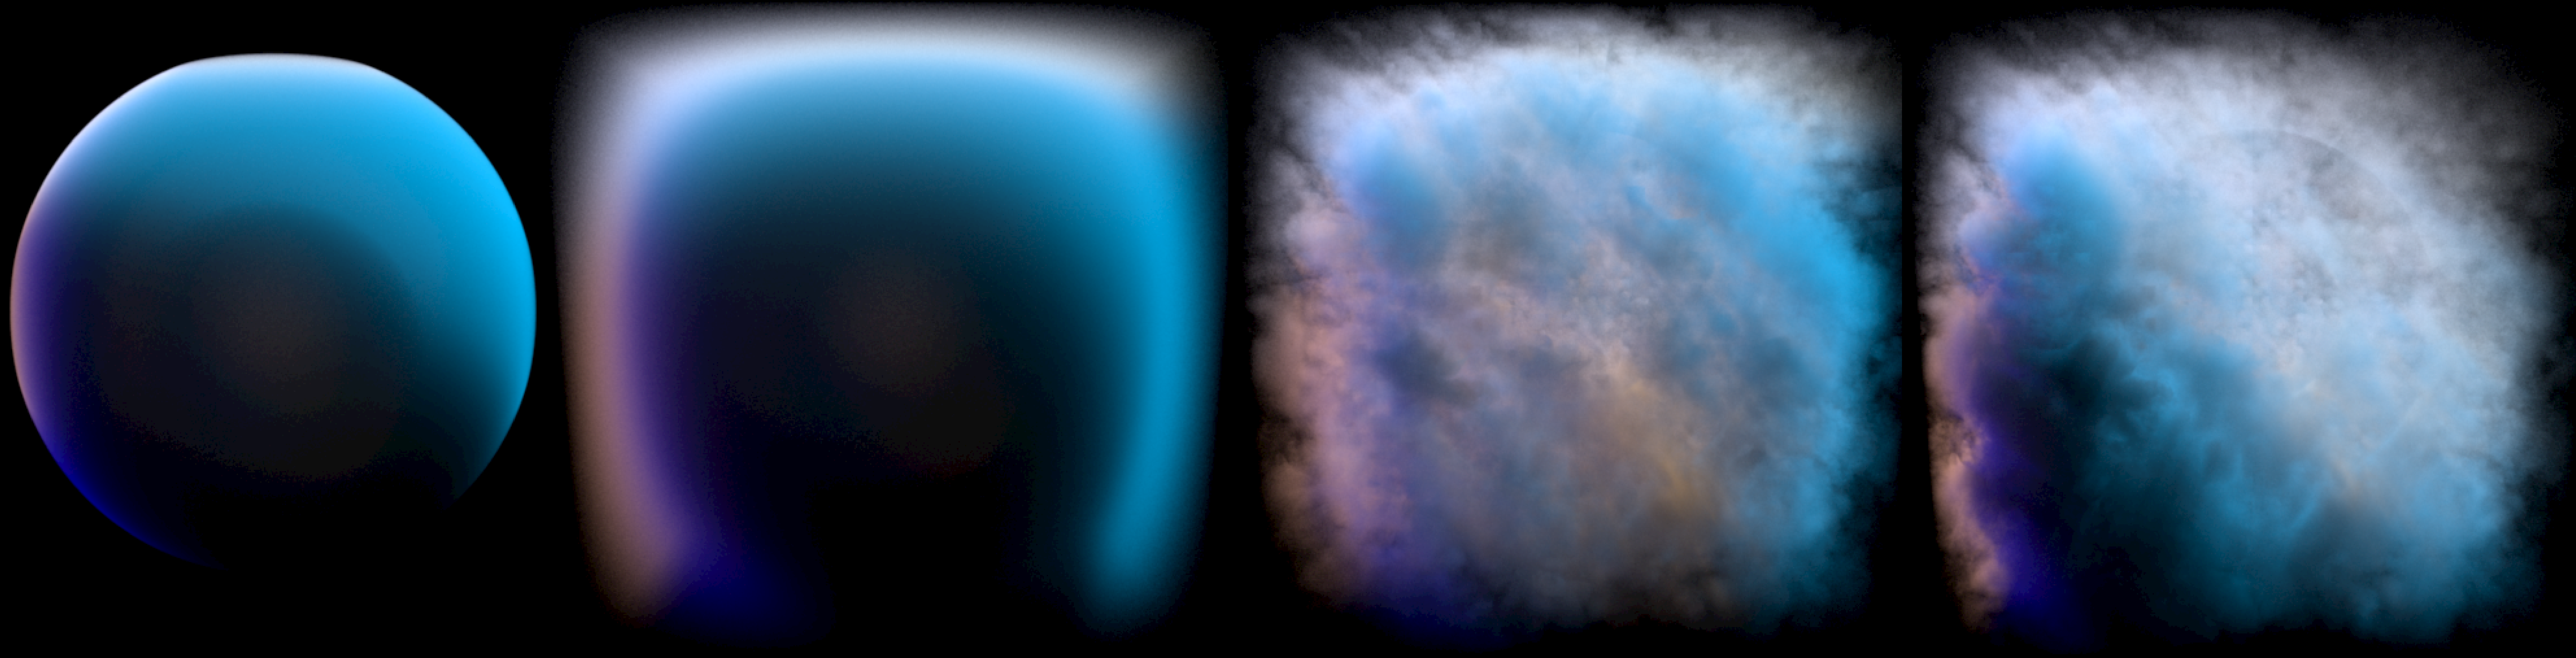


checkbox node
Hello
World
in
out
node 2
single in 1
single in 2
multi in
single out 1
single out 2
multi out
Backdrop
chantastic!
in port
out port
checkbox node 1
Hello
World
in
out
node 3
Text Input
in
out
menu node
DISABLED
hello
world
out A
node
in A
out A
checkbox node 2
Hello

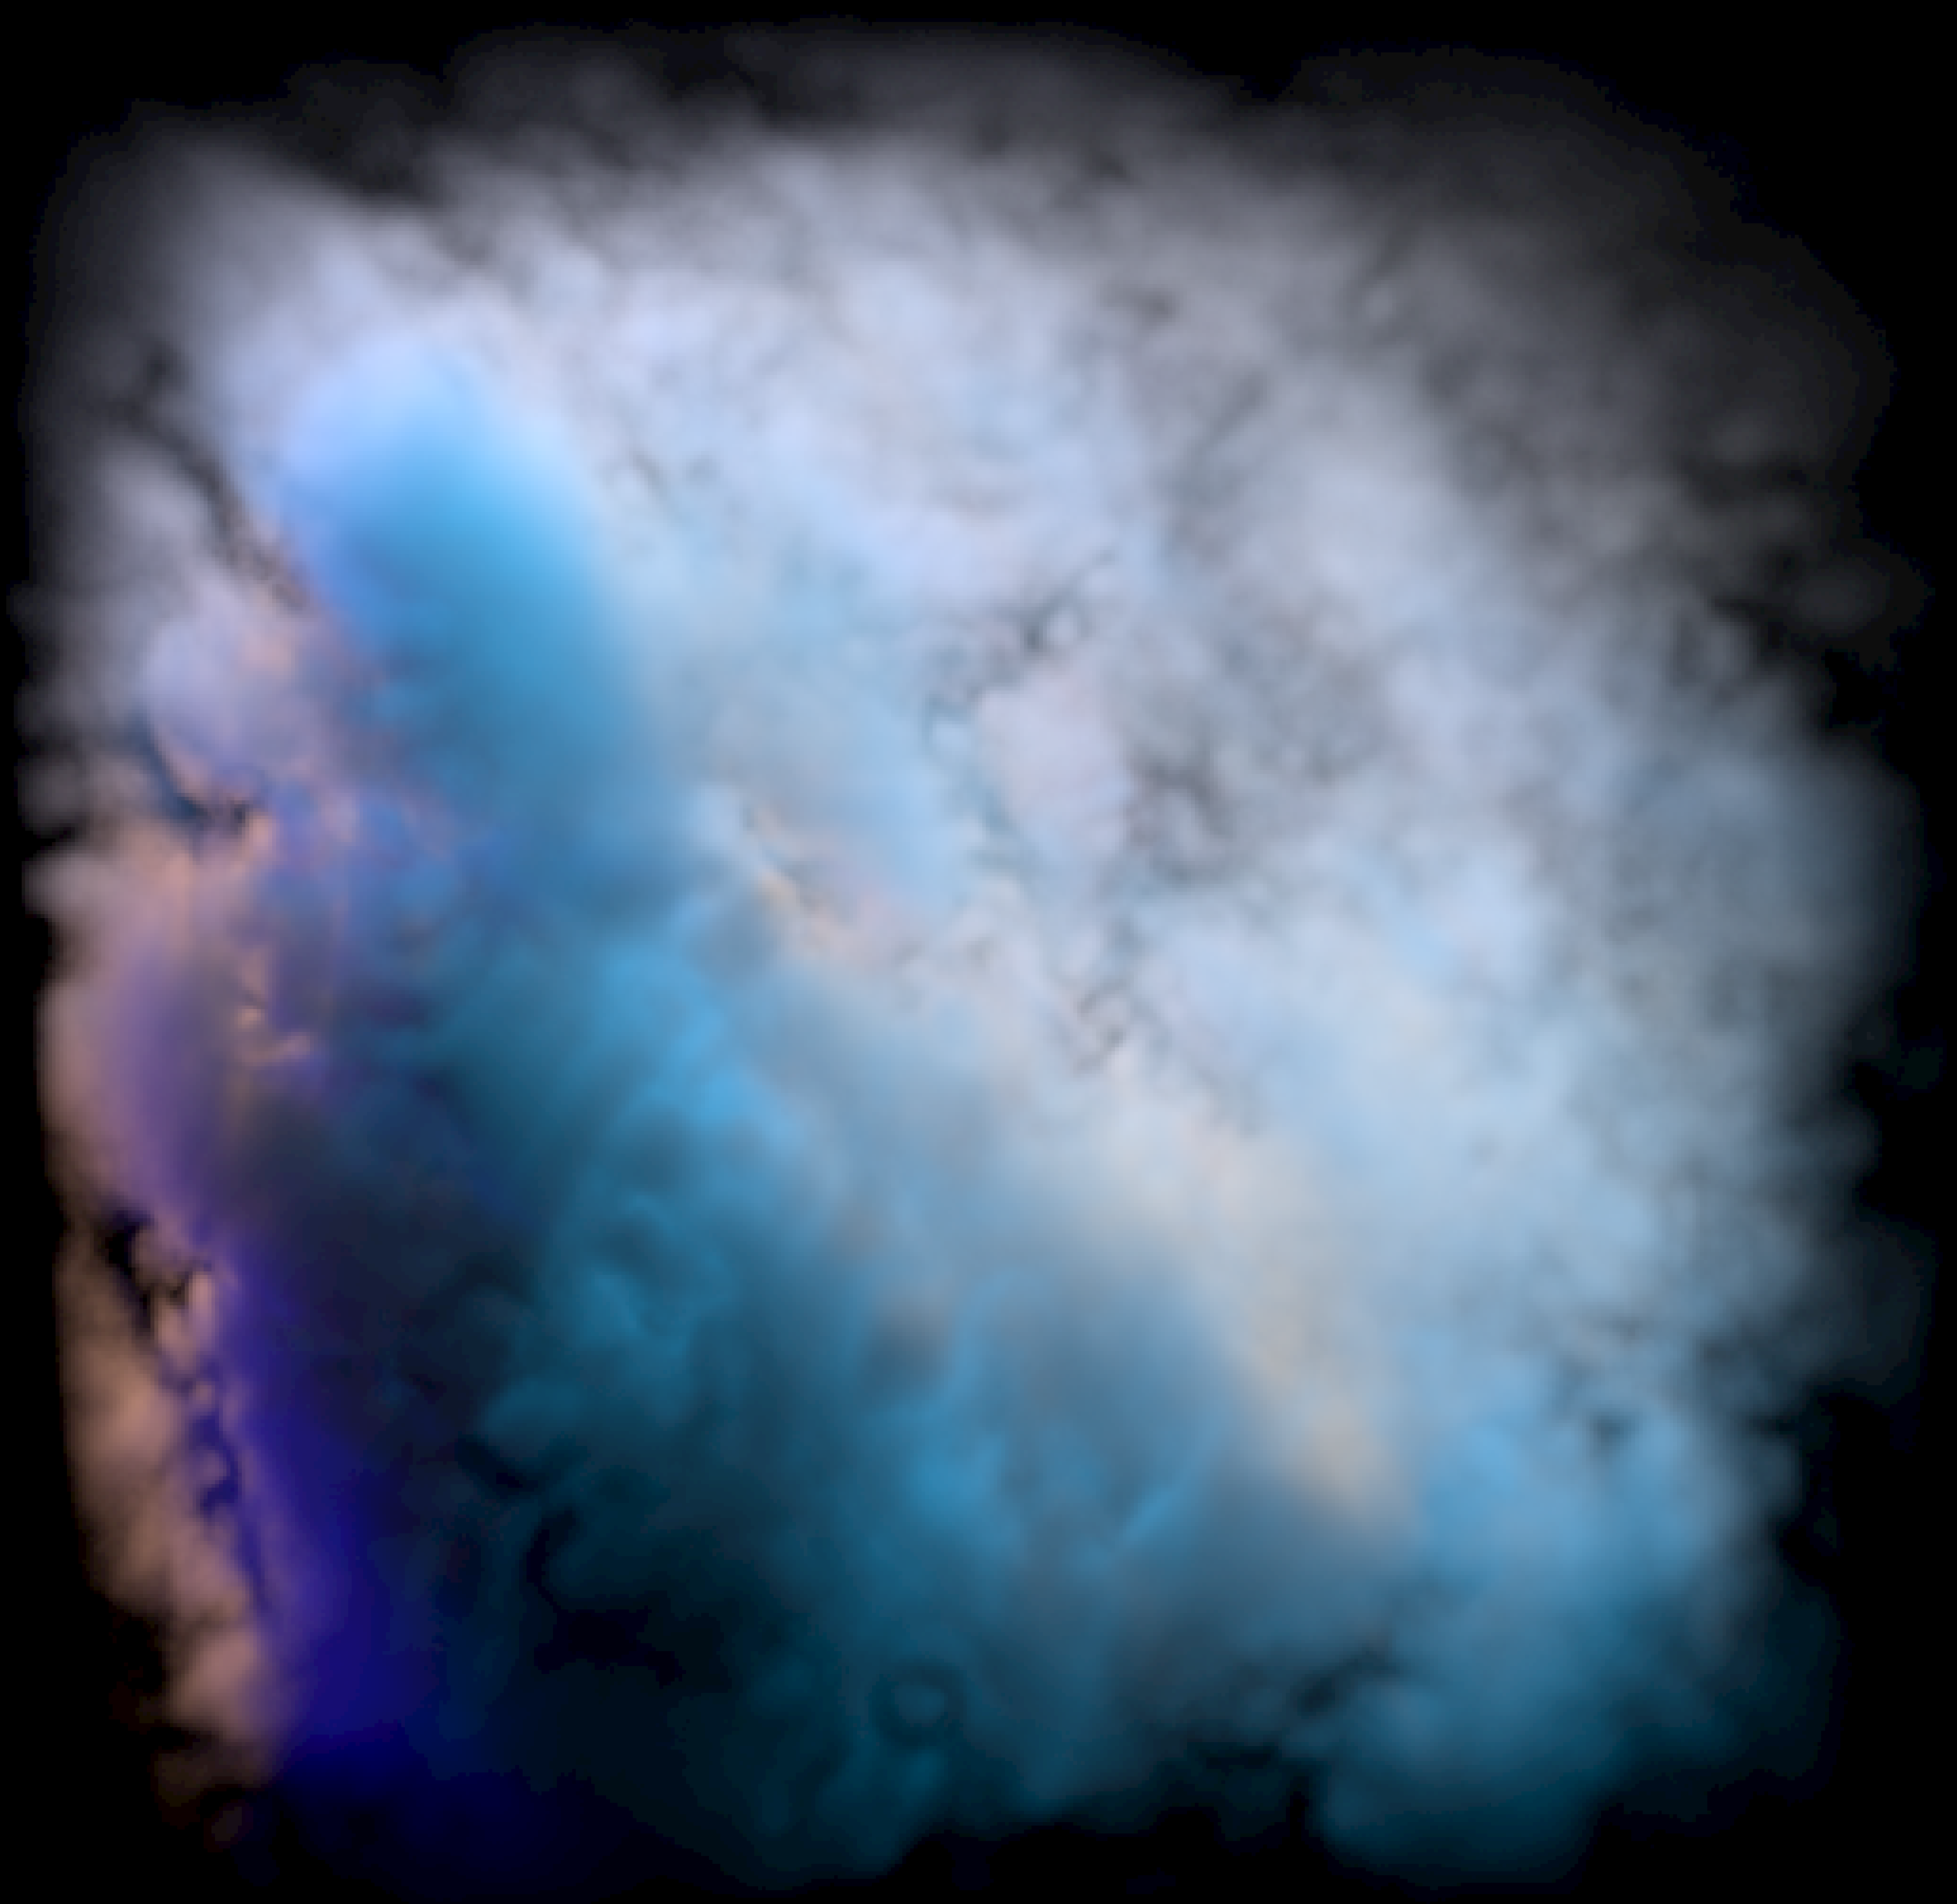

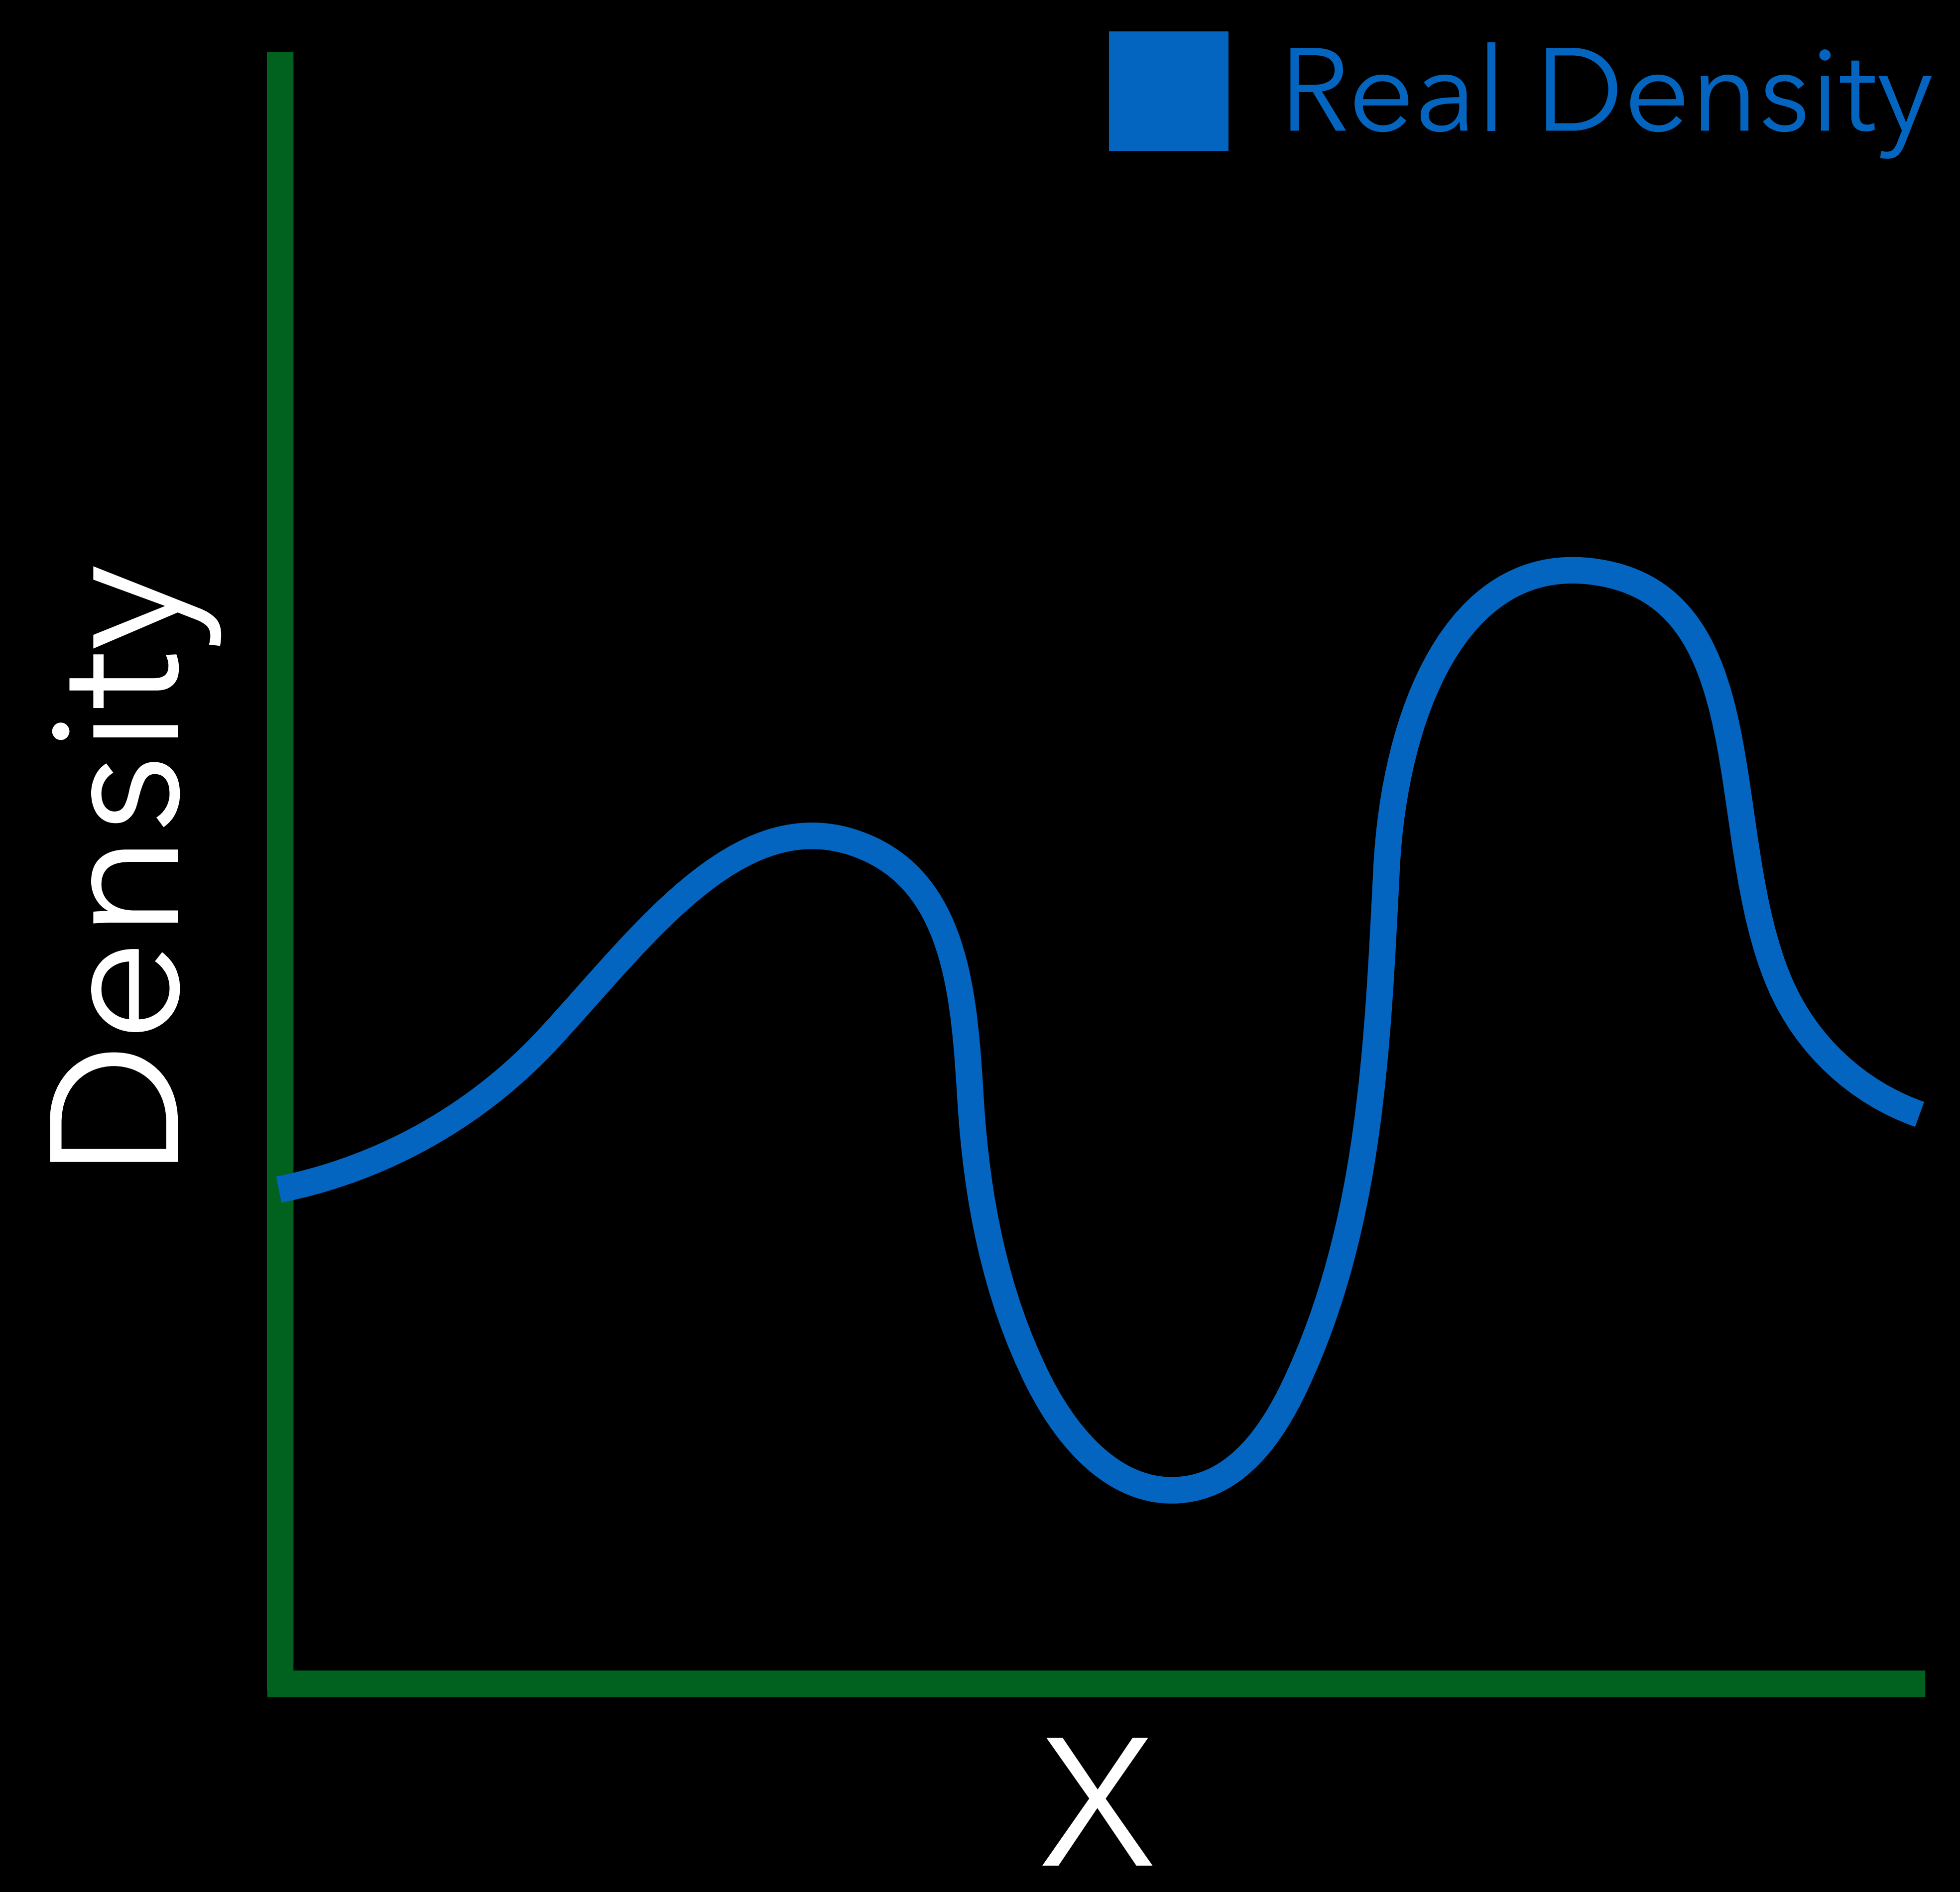
Real Density
Density
X

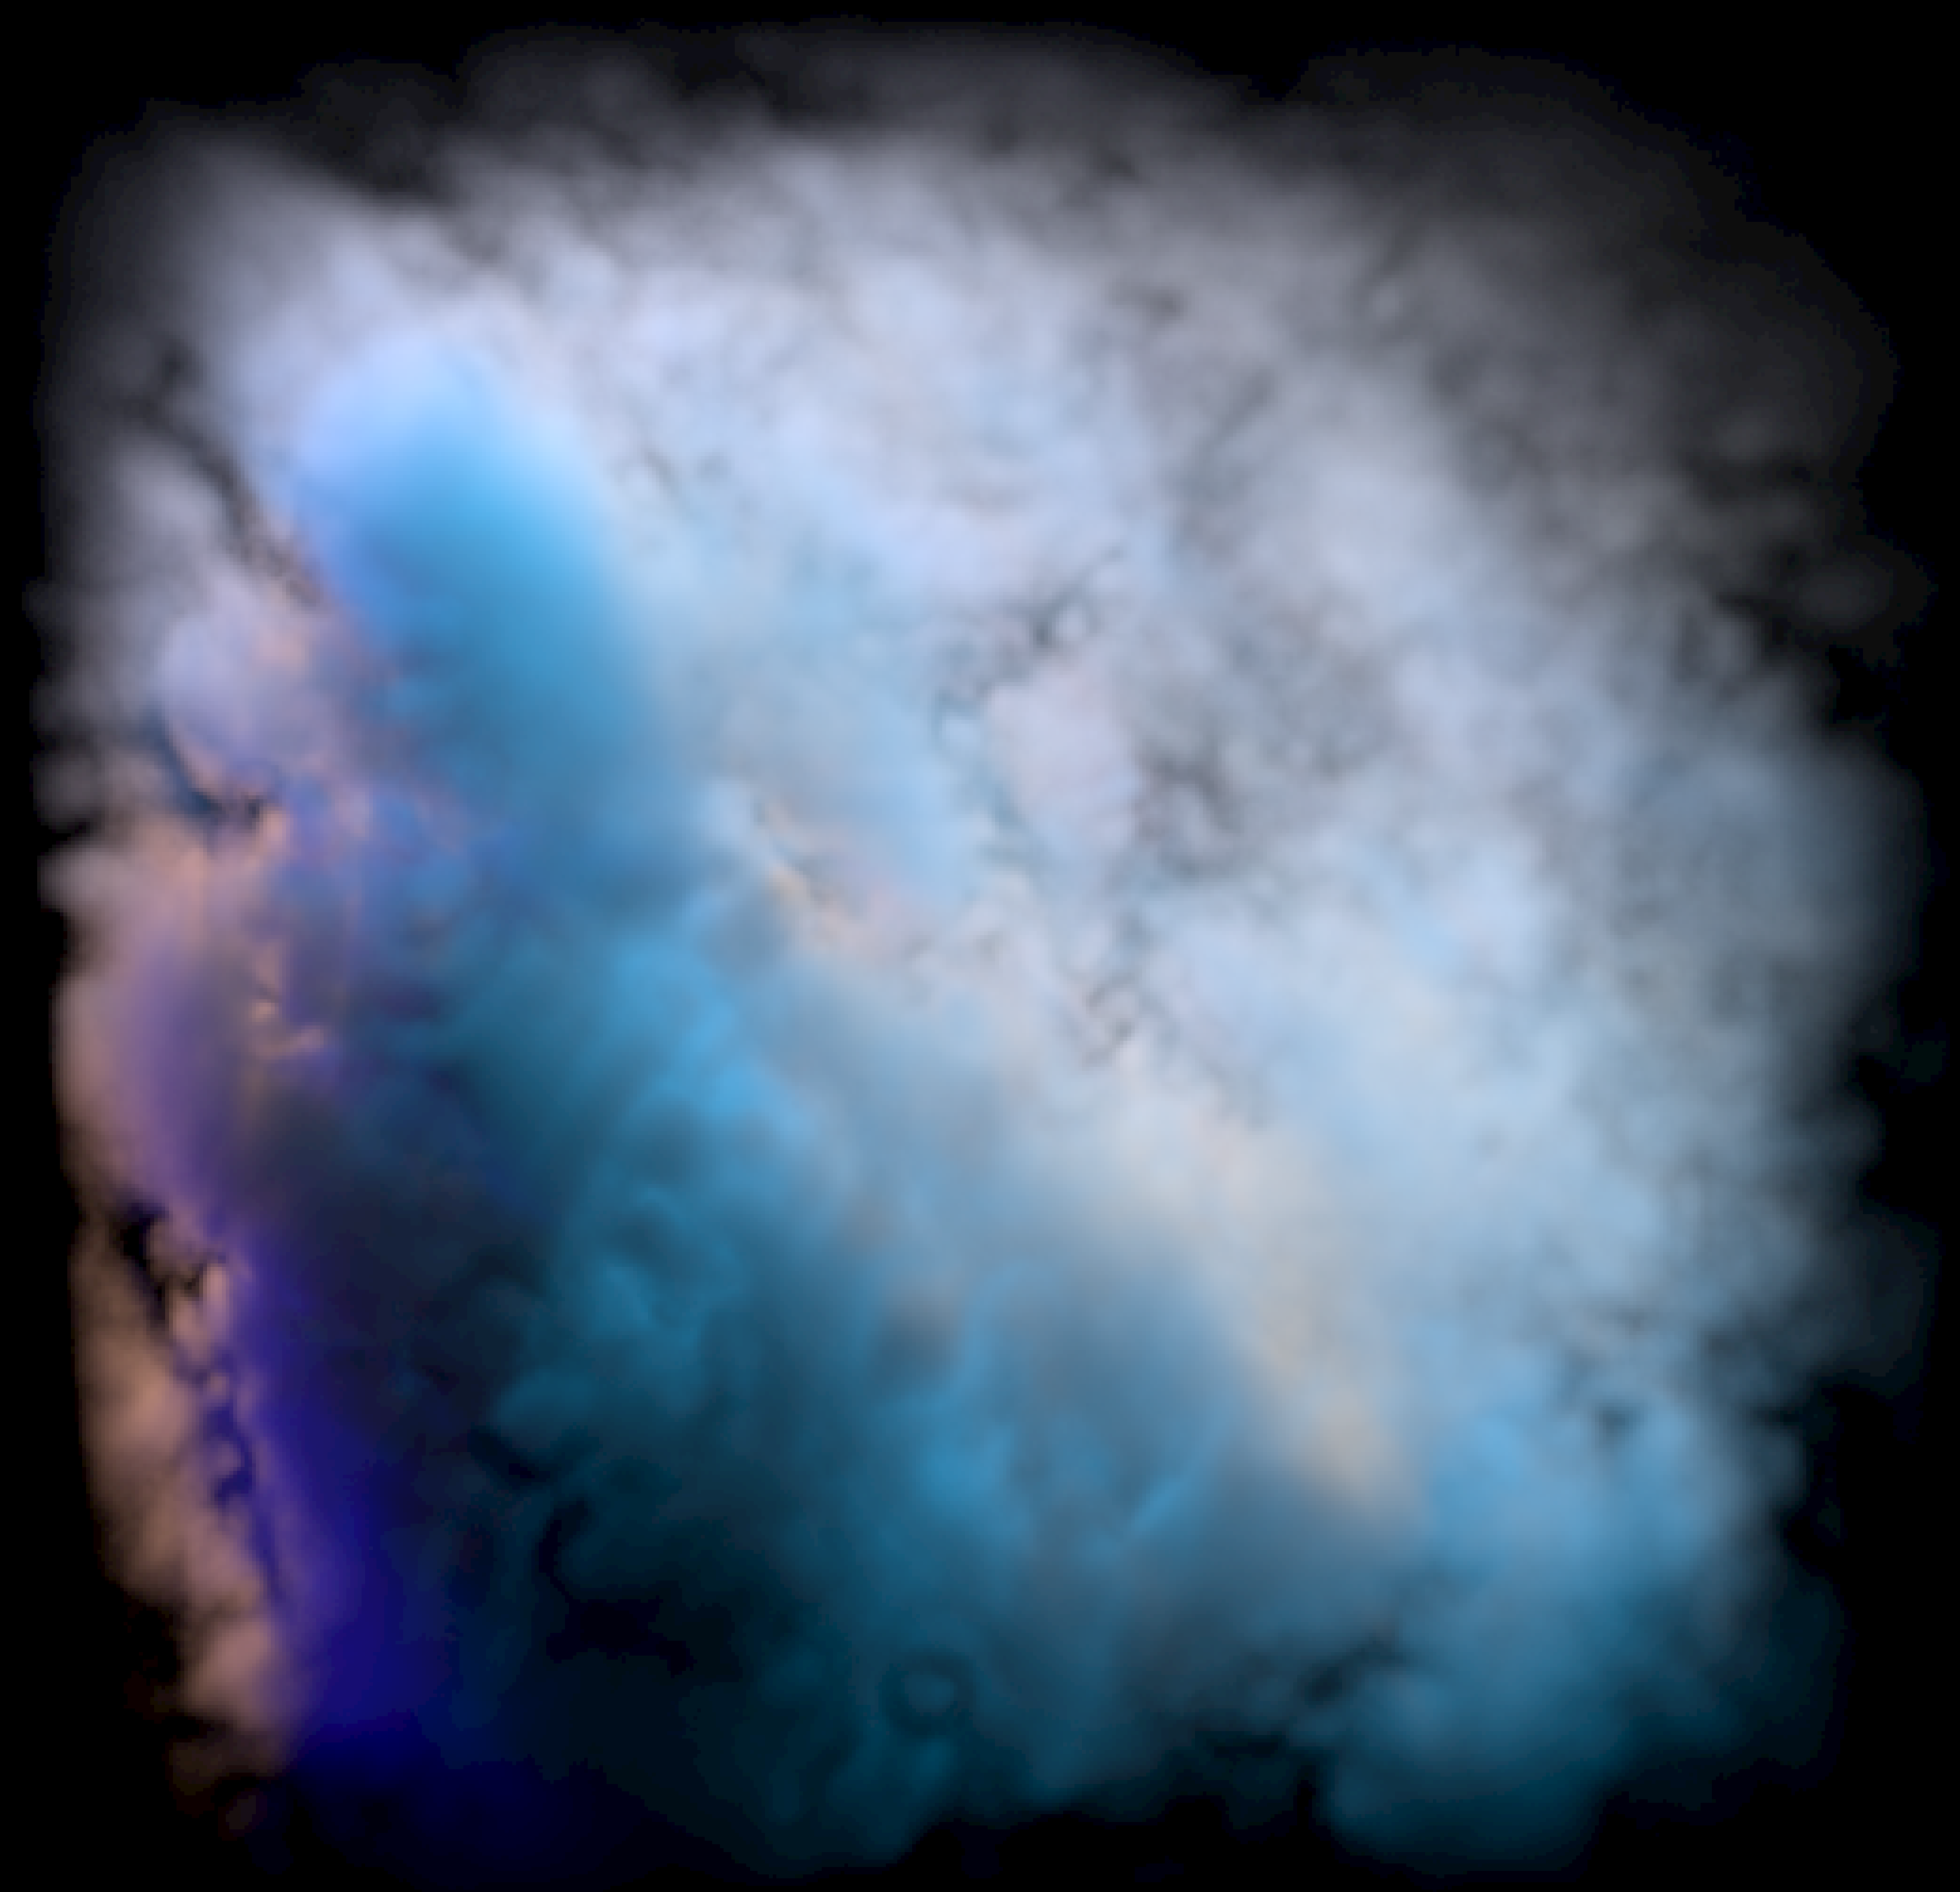

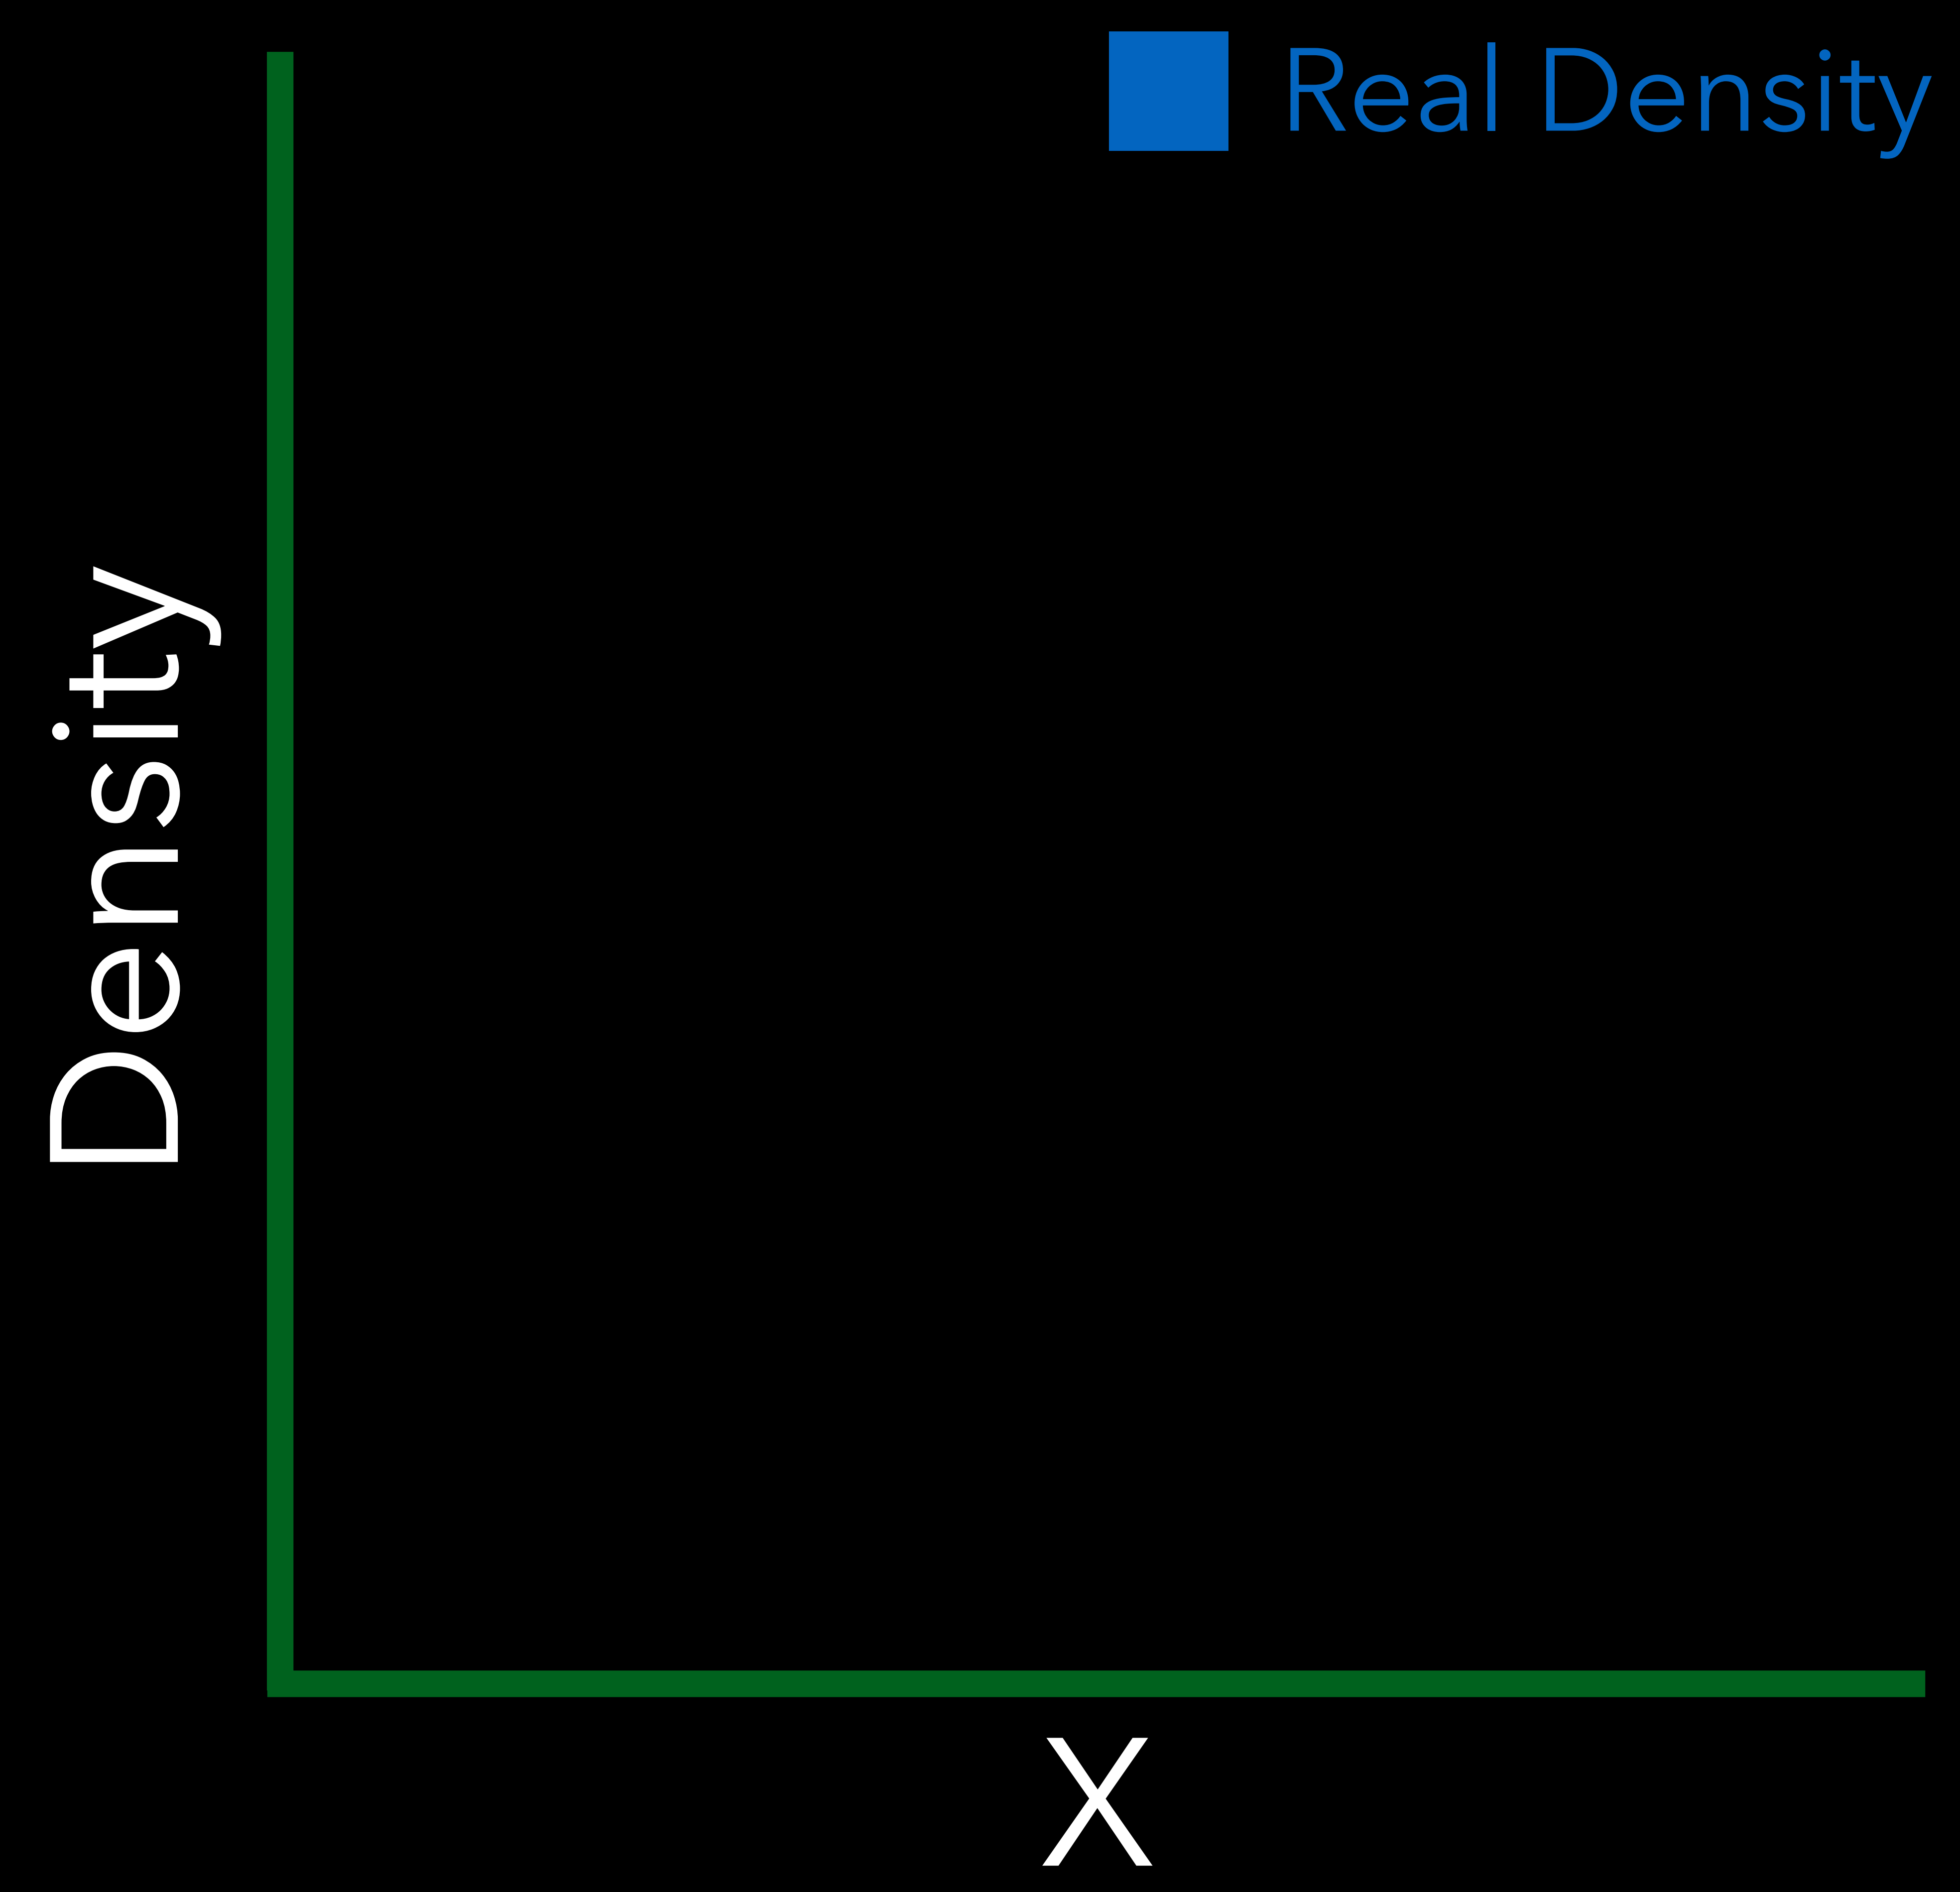
Real Density
Density
X

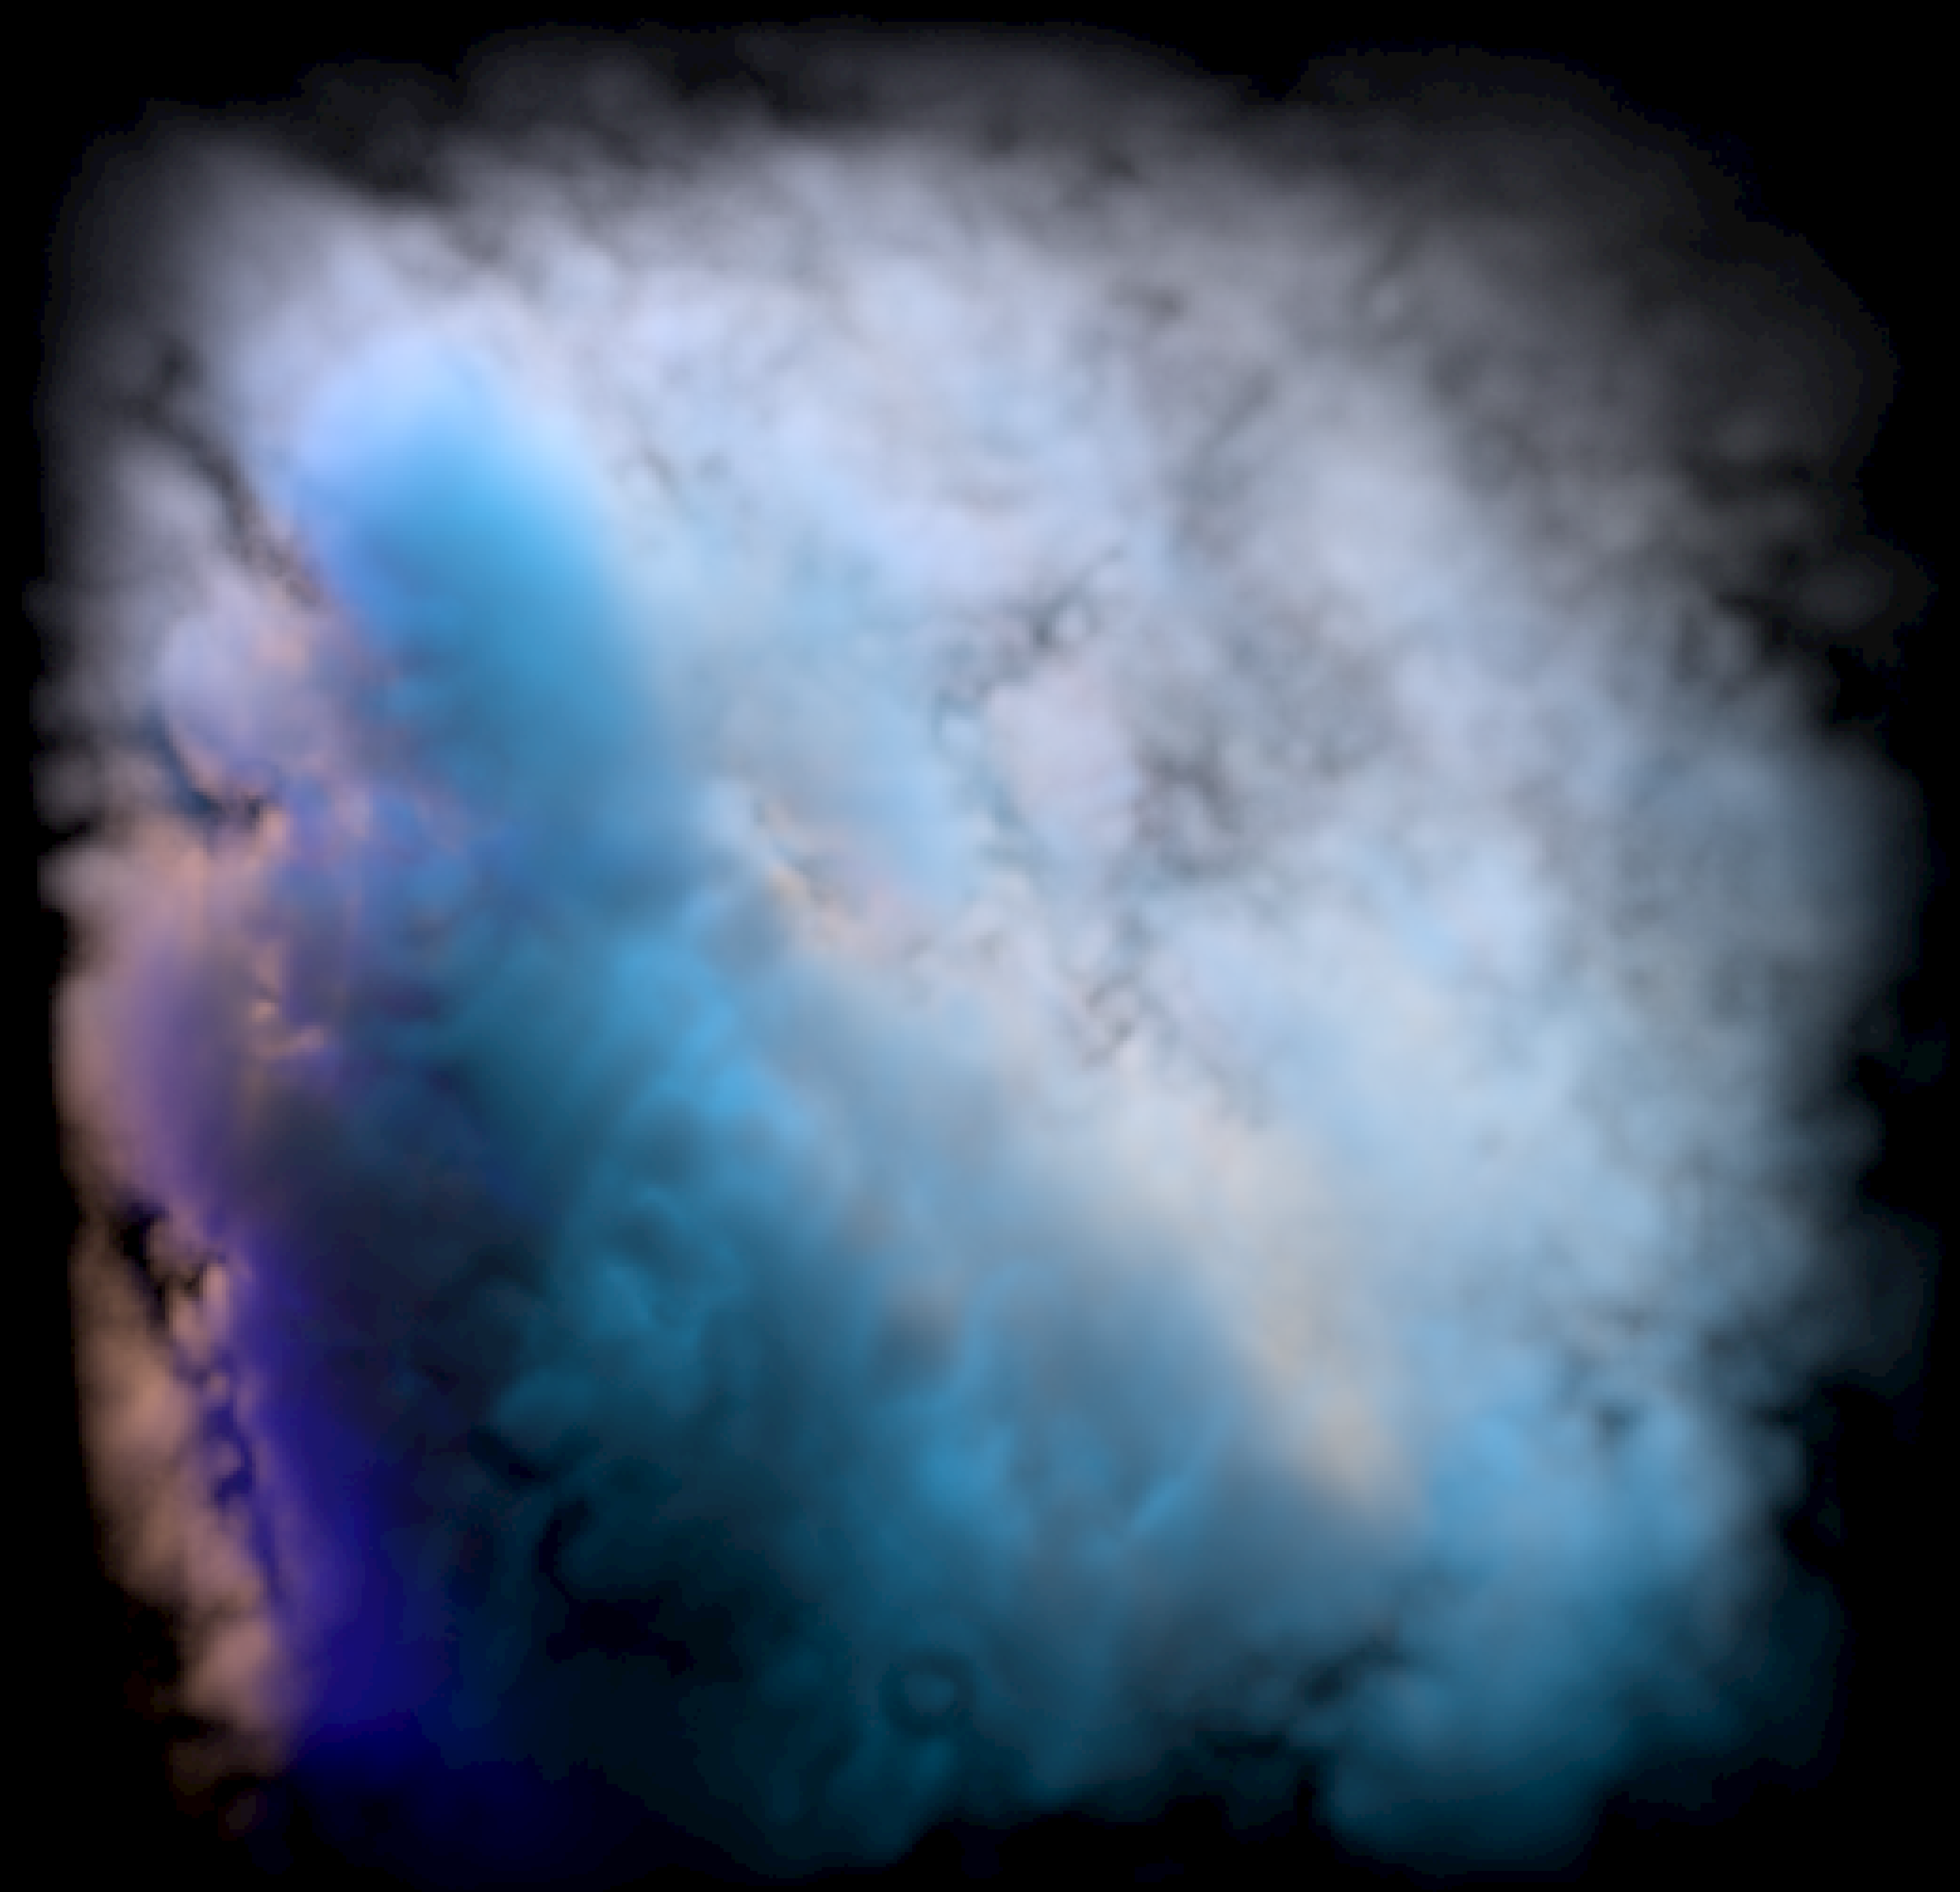

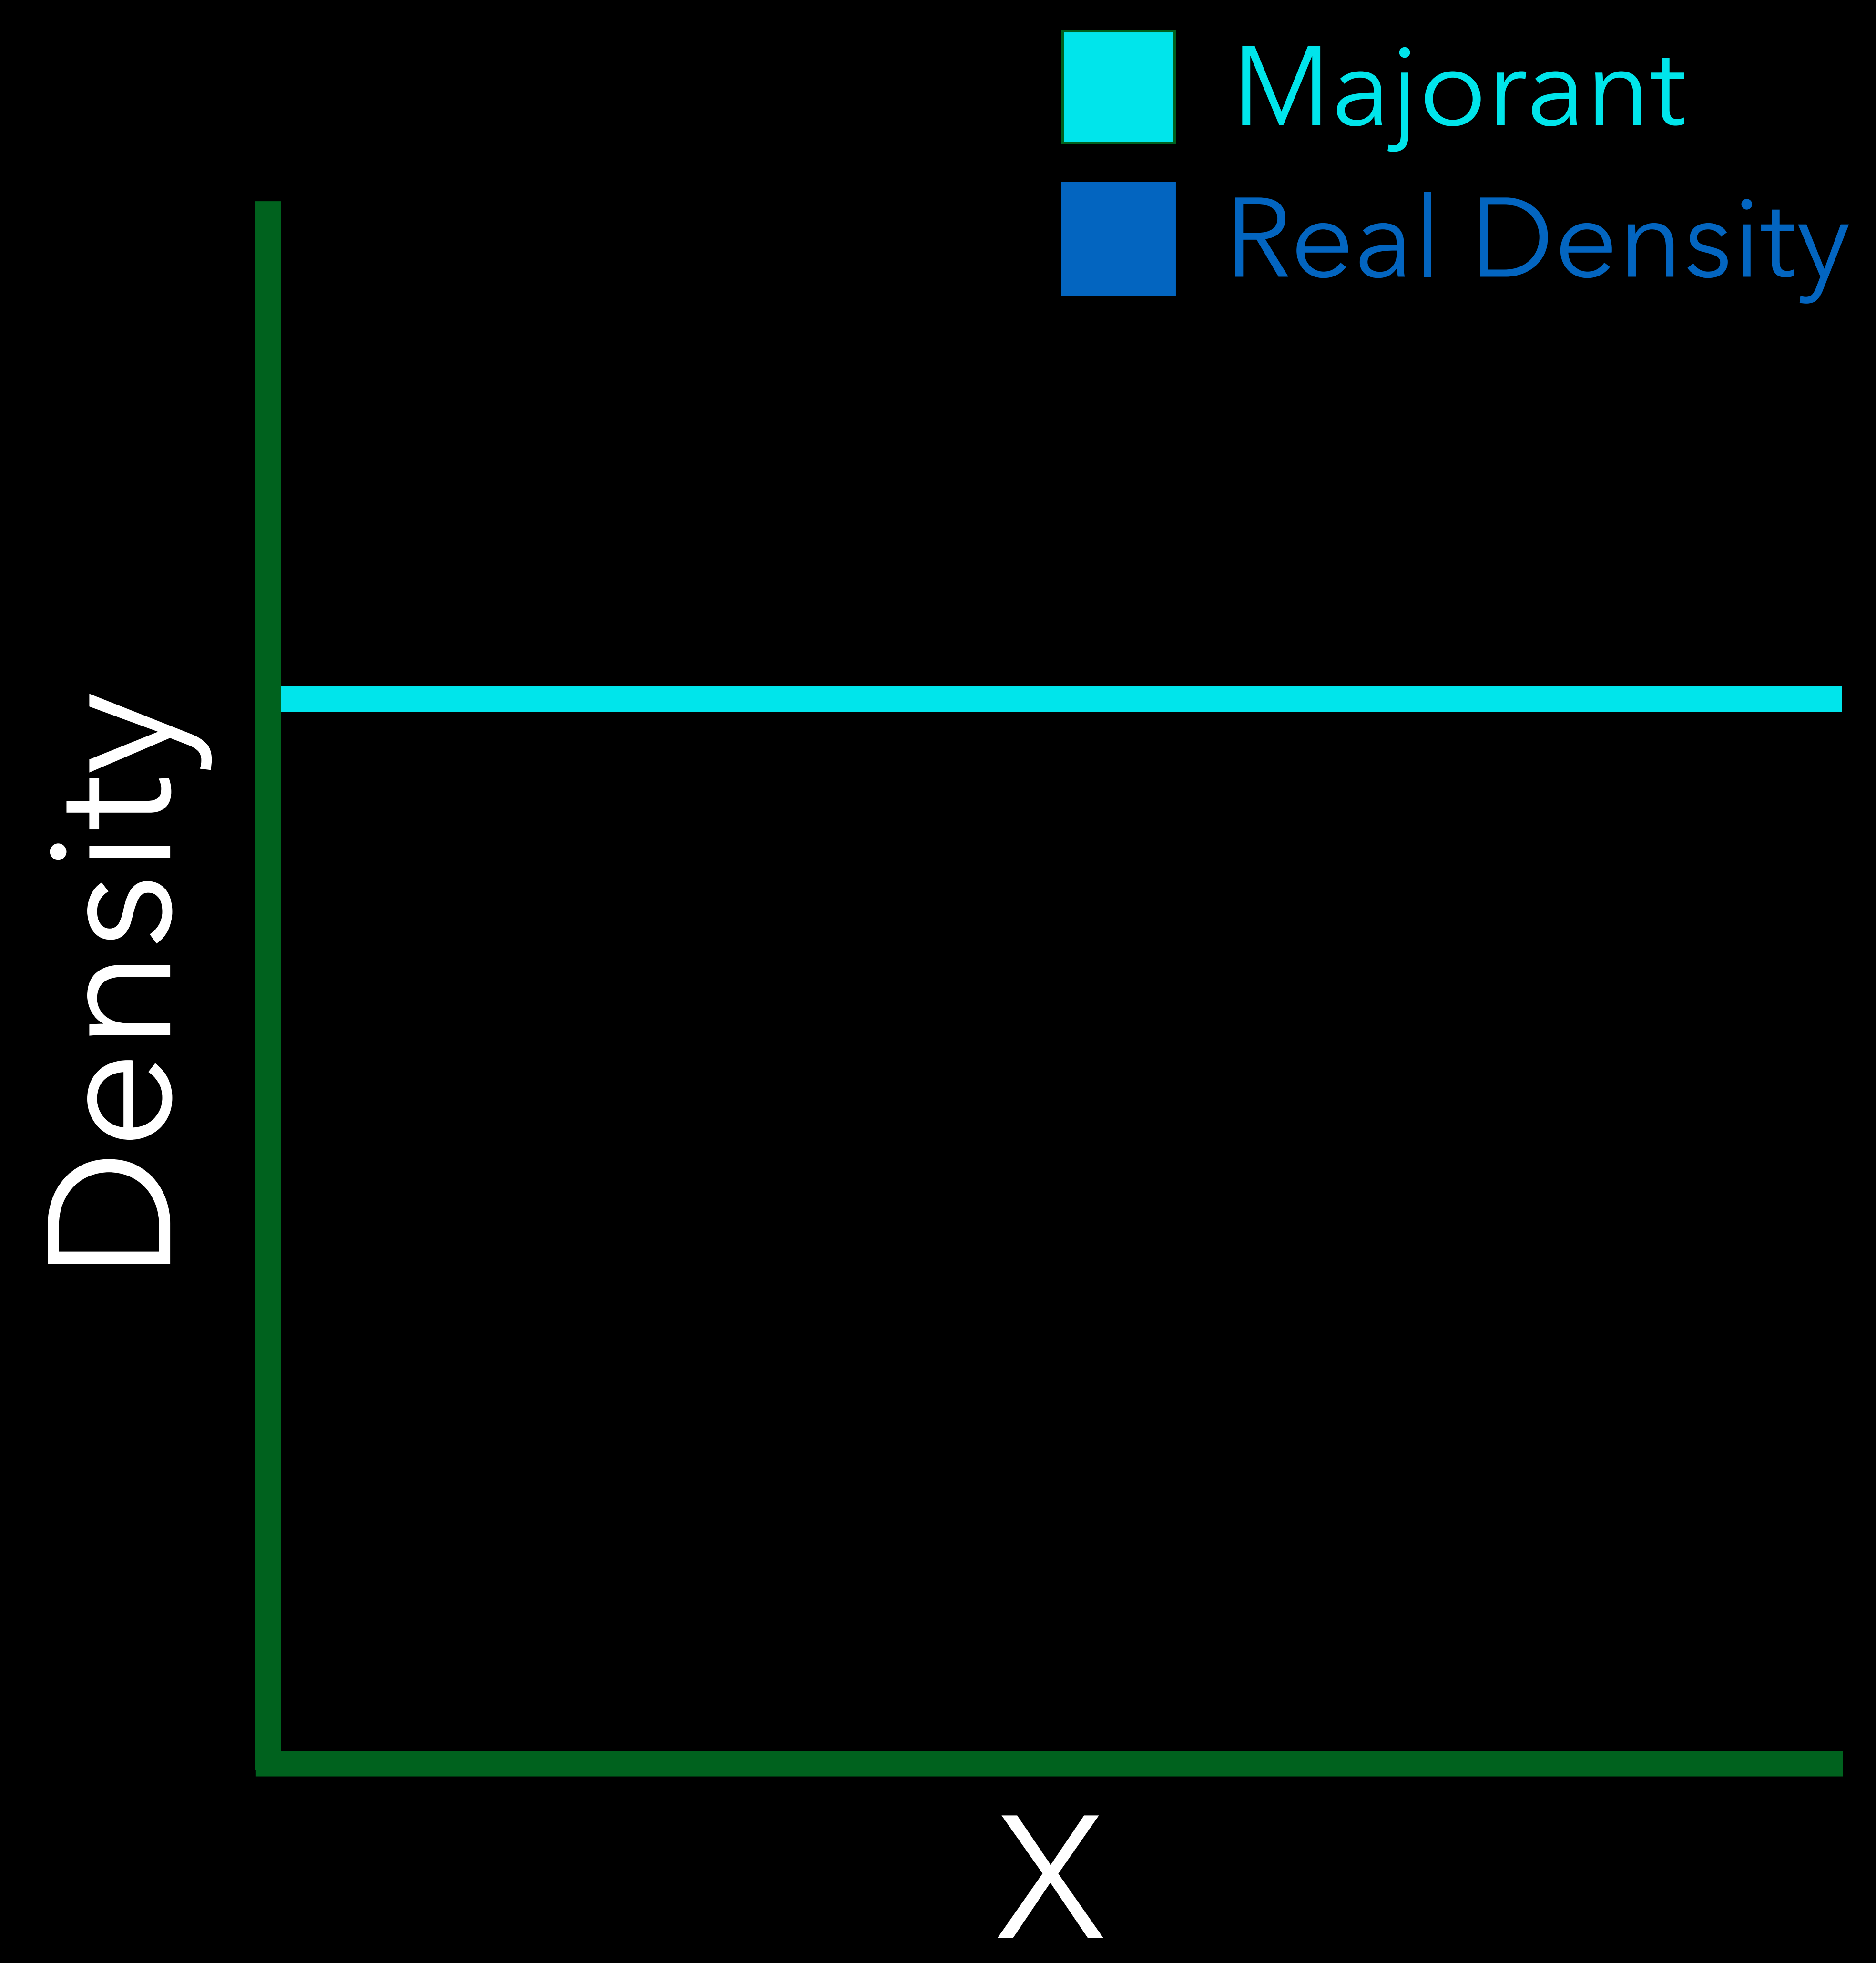
Majorant
Real Density
Density
X

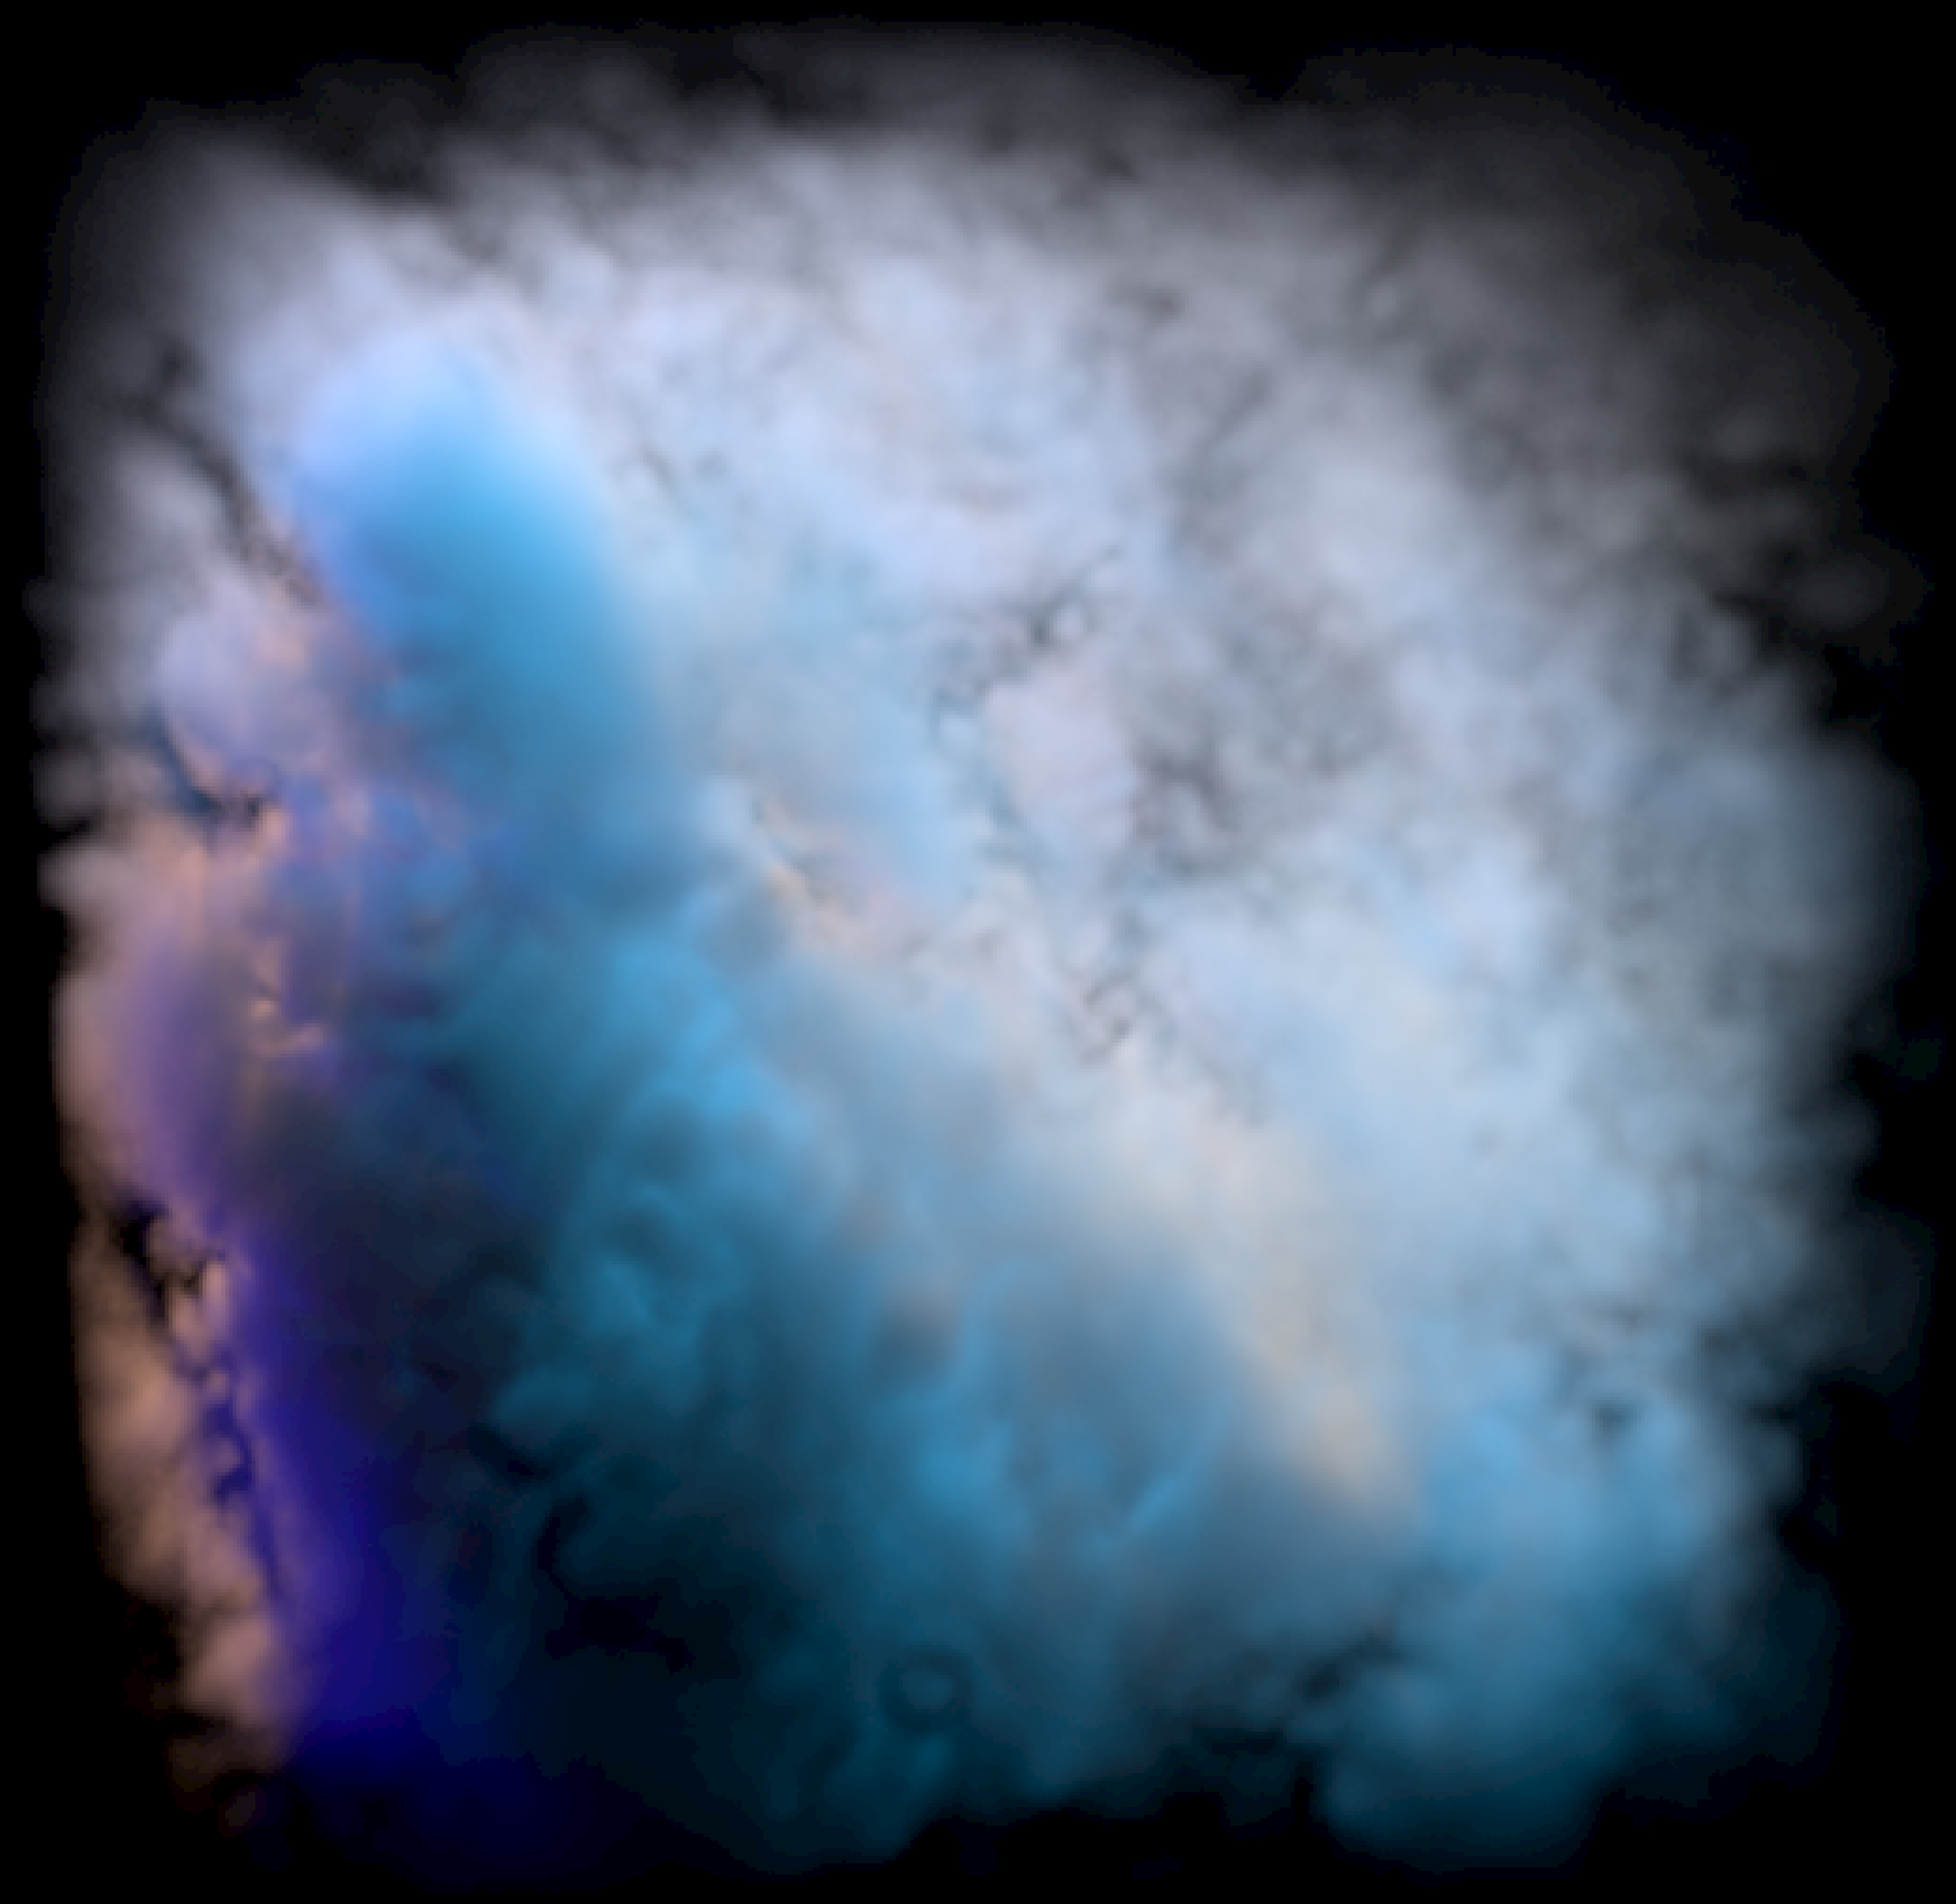

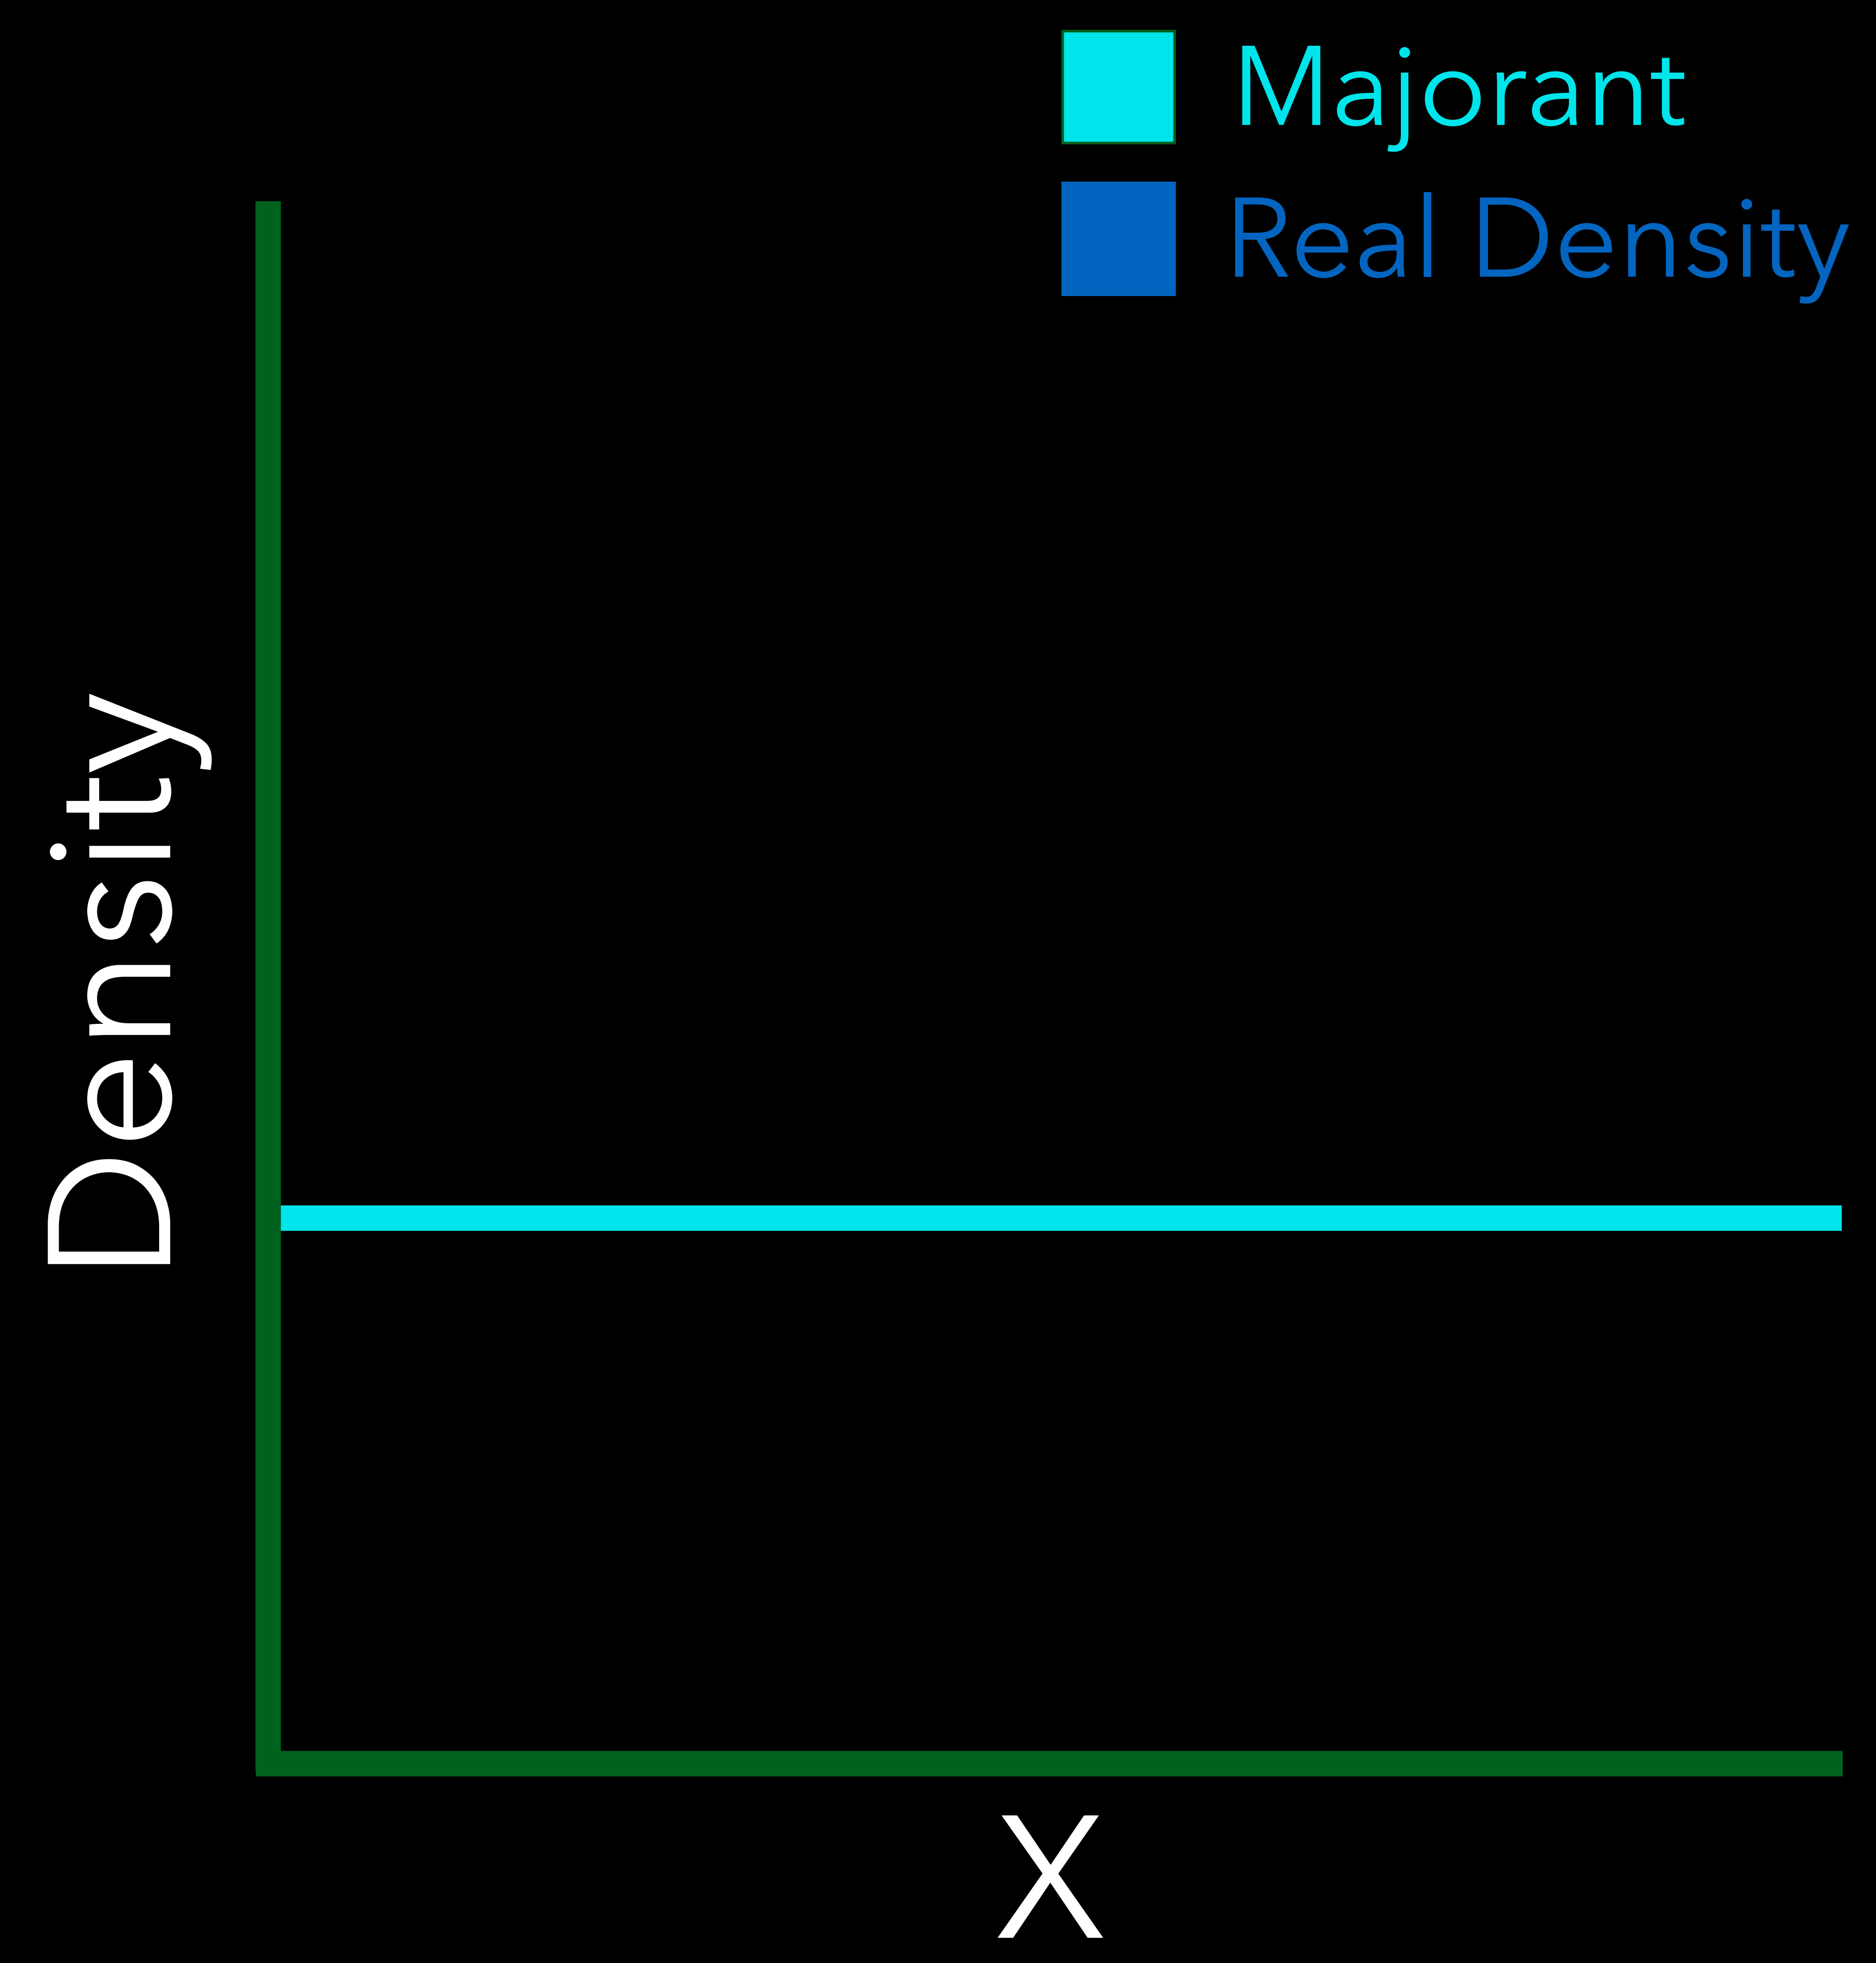
Majorant
Real Density
Density
X

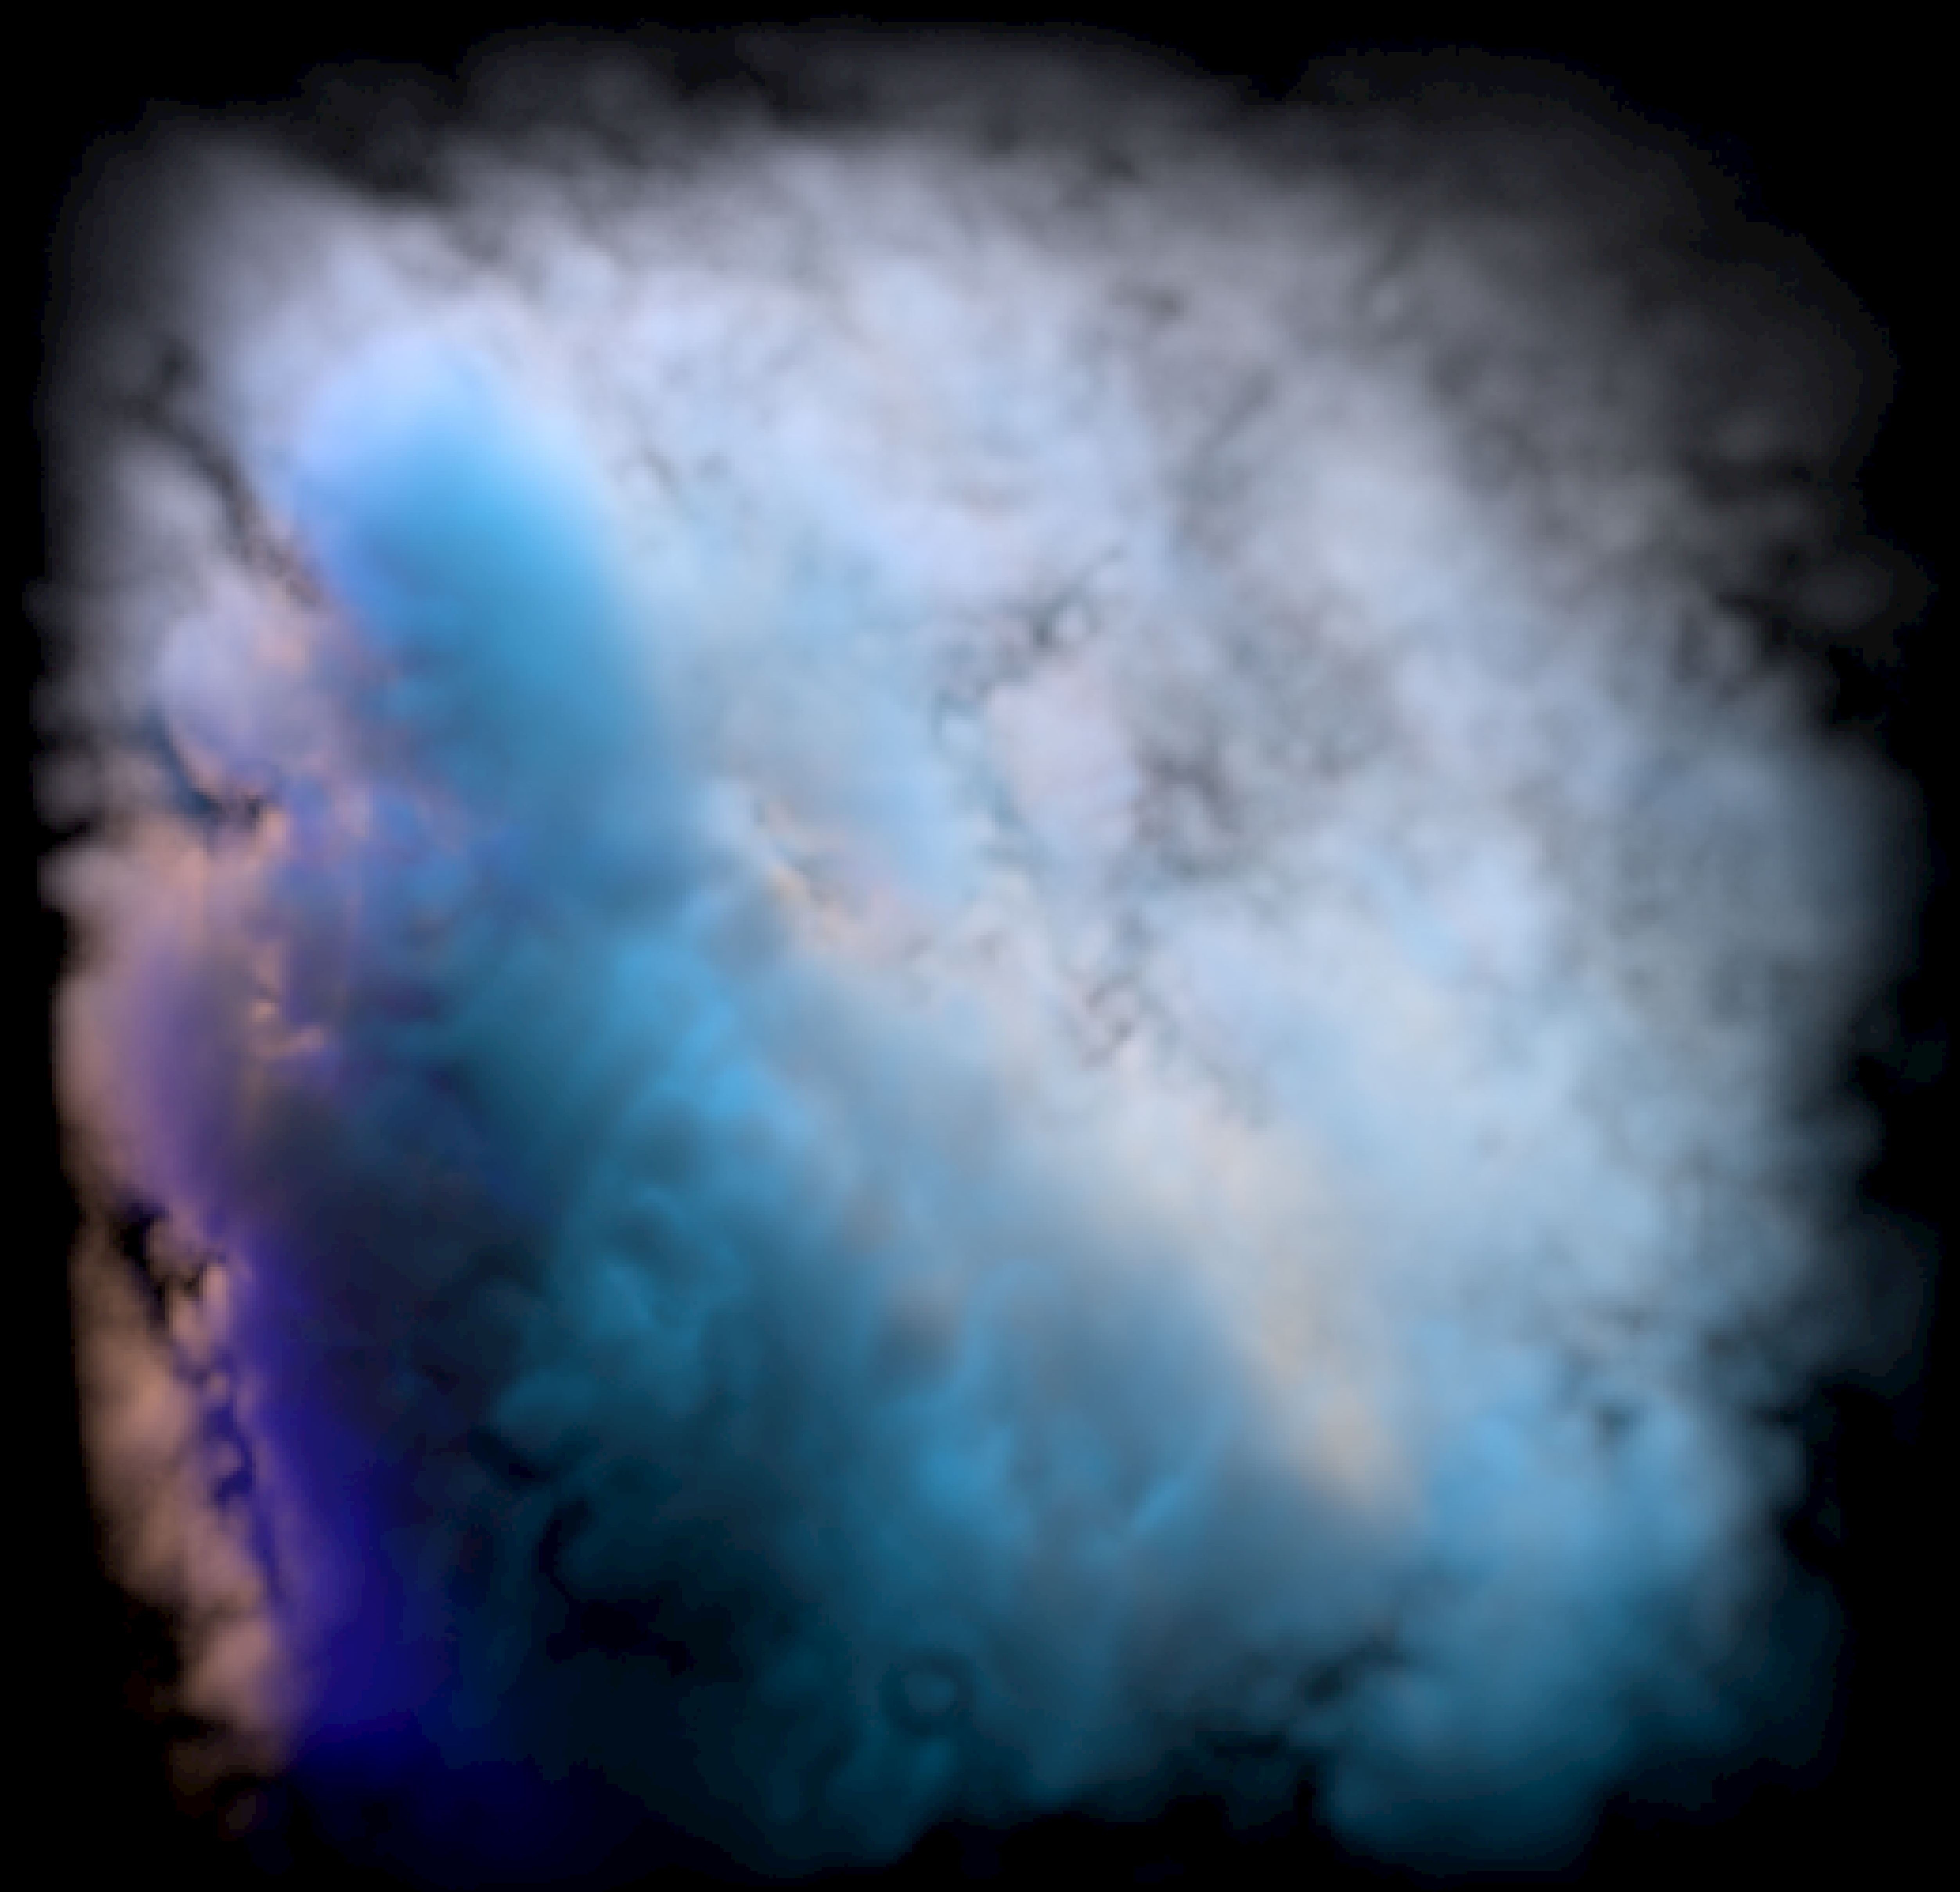

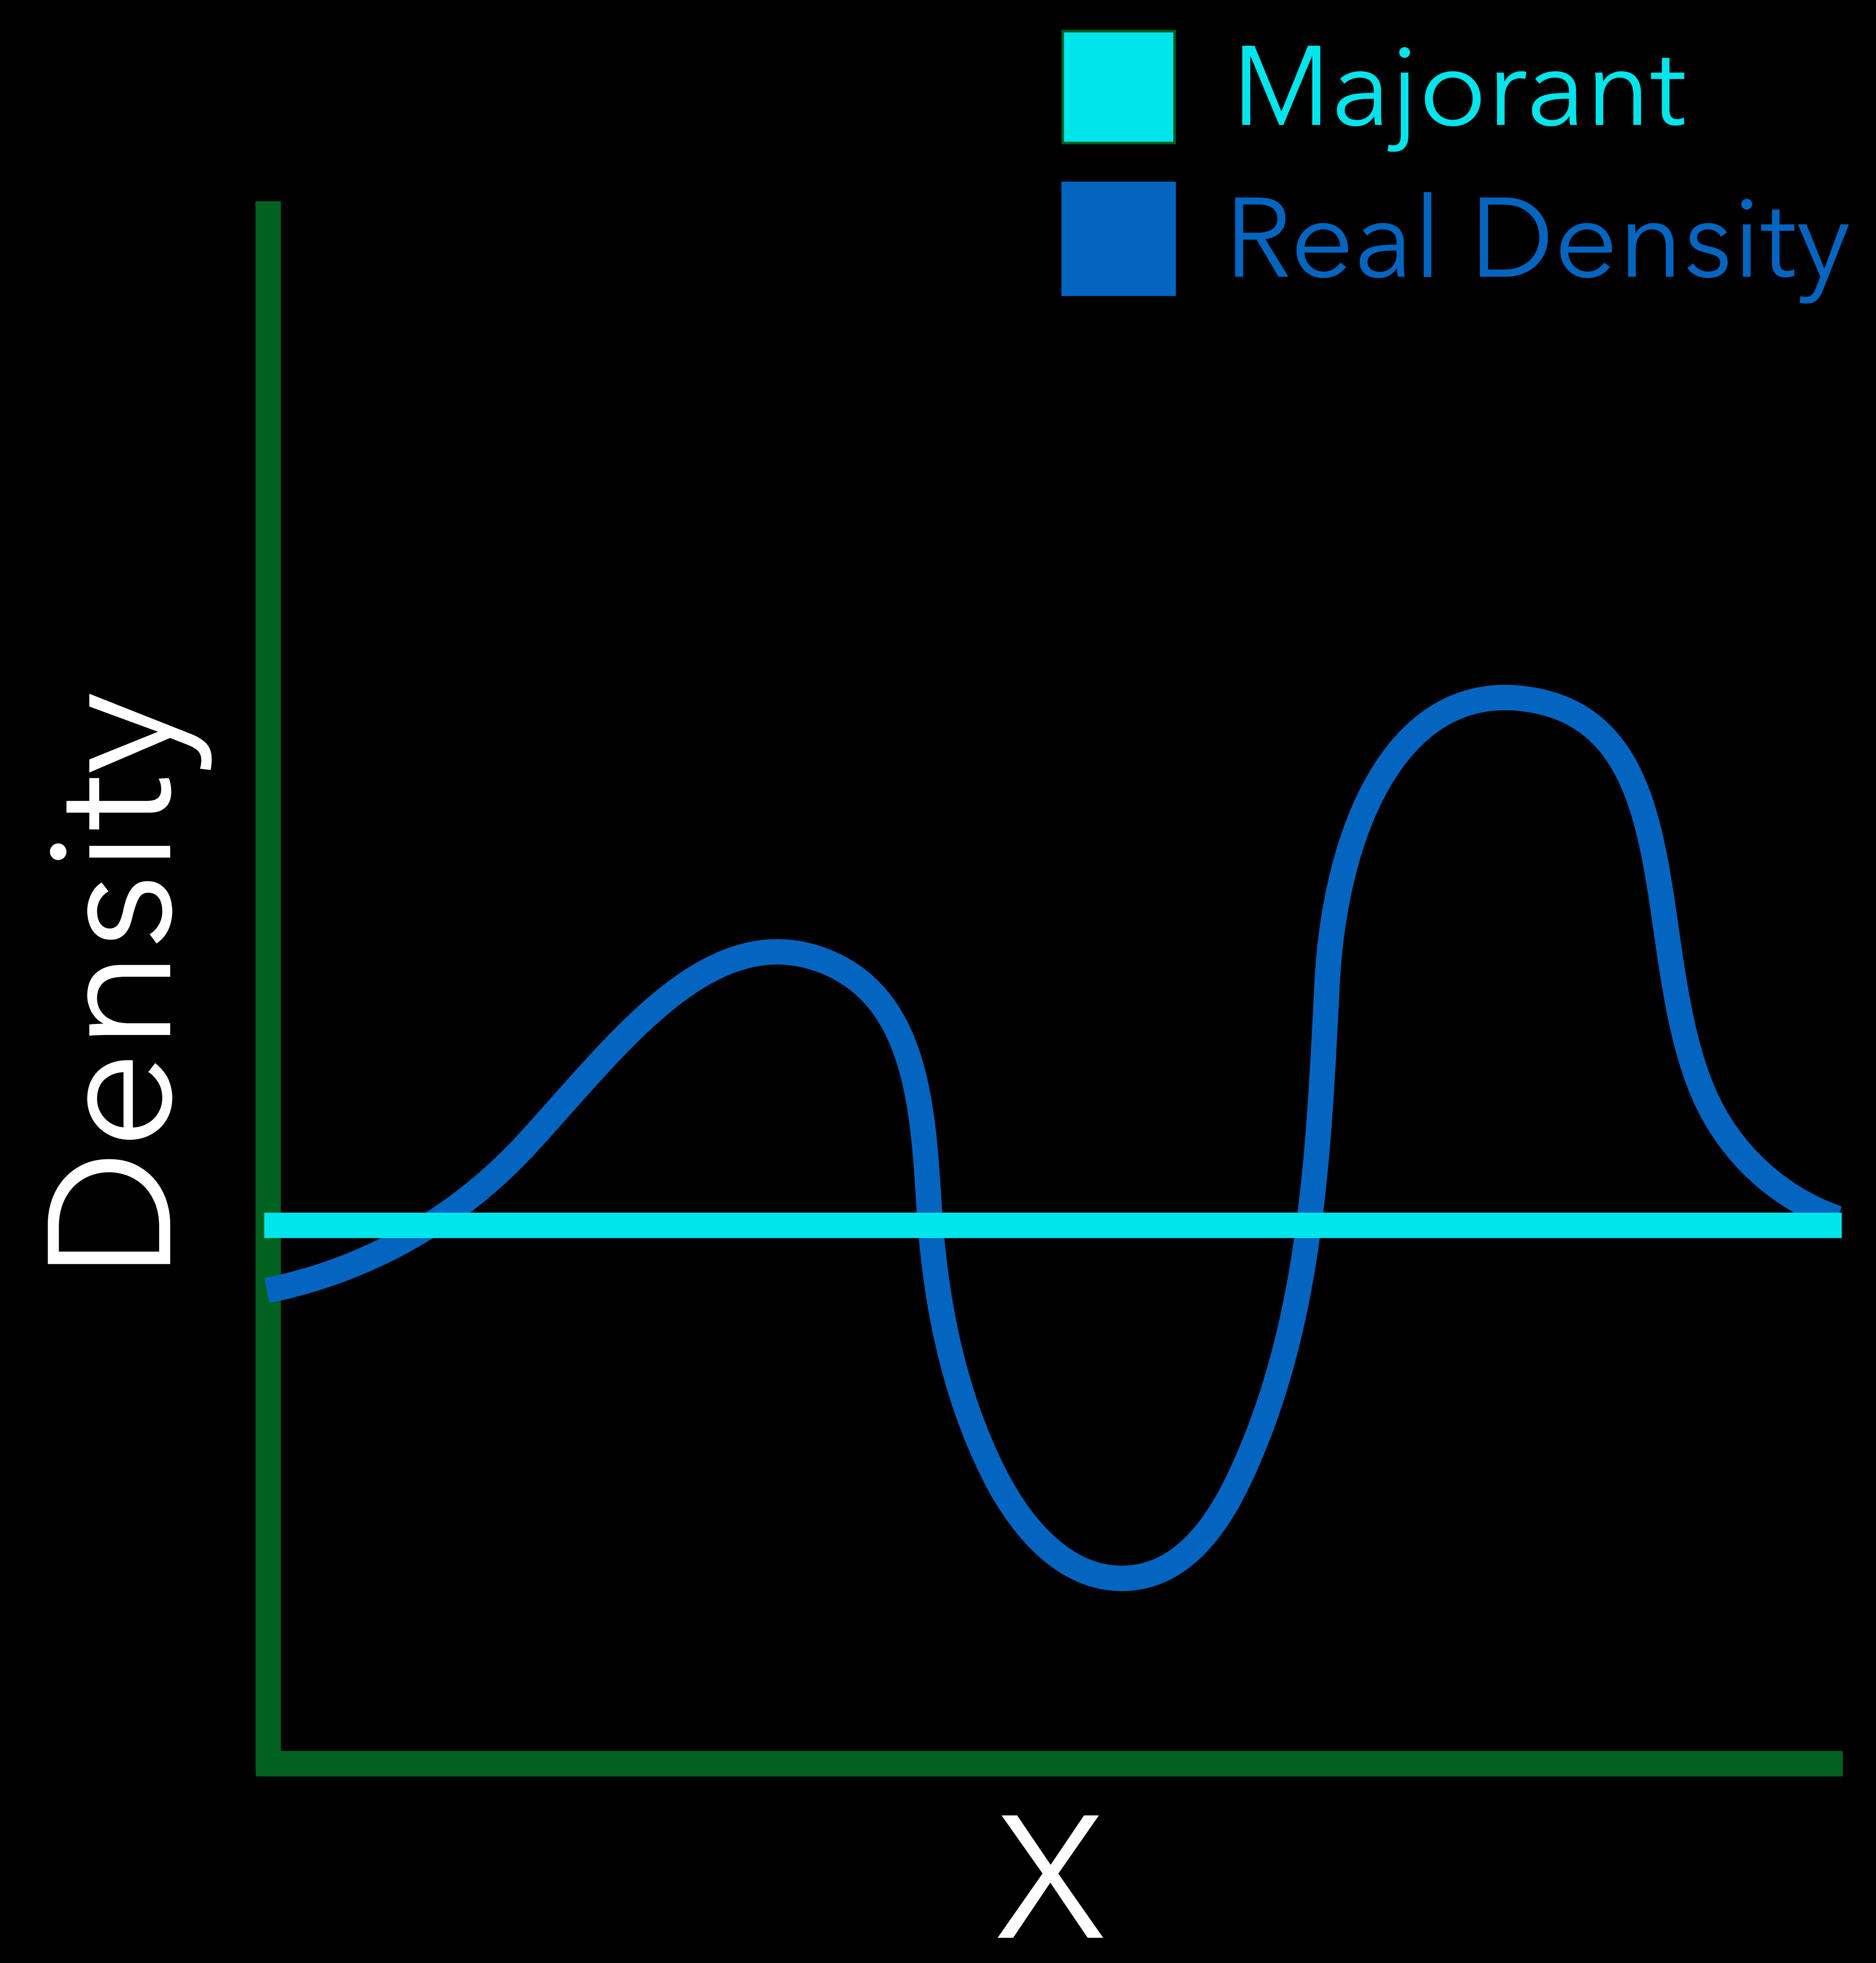
Majorant
Real Density
Density
X

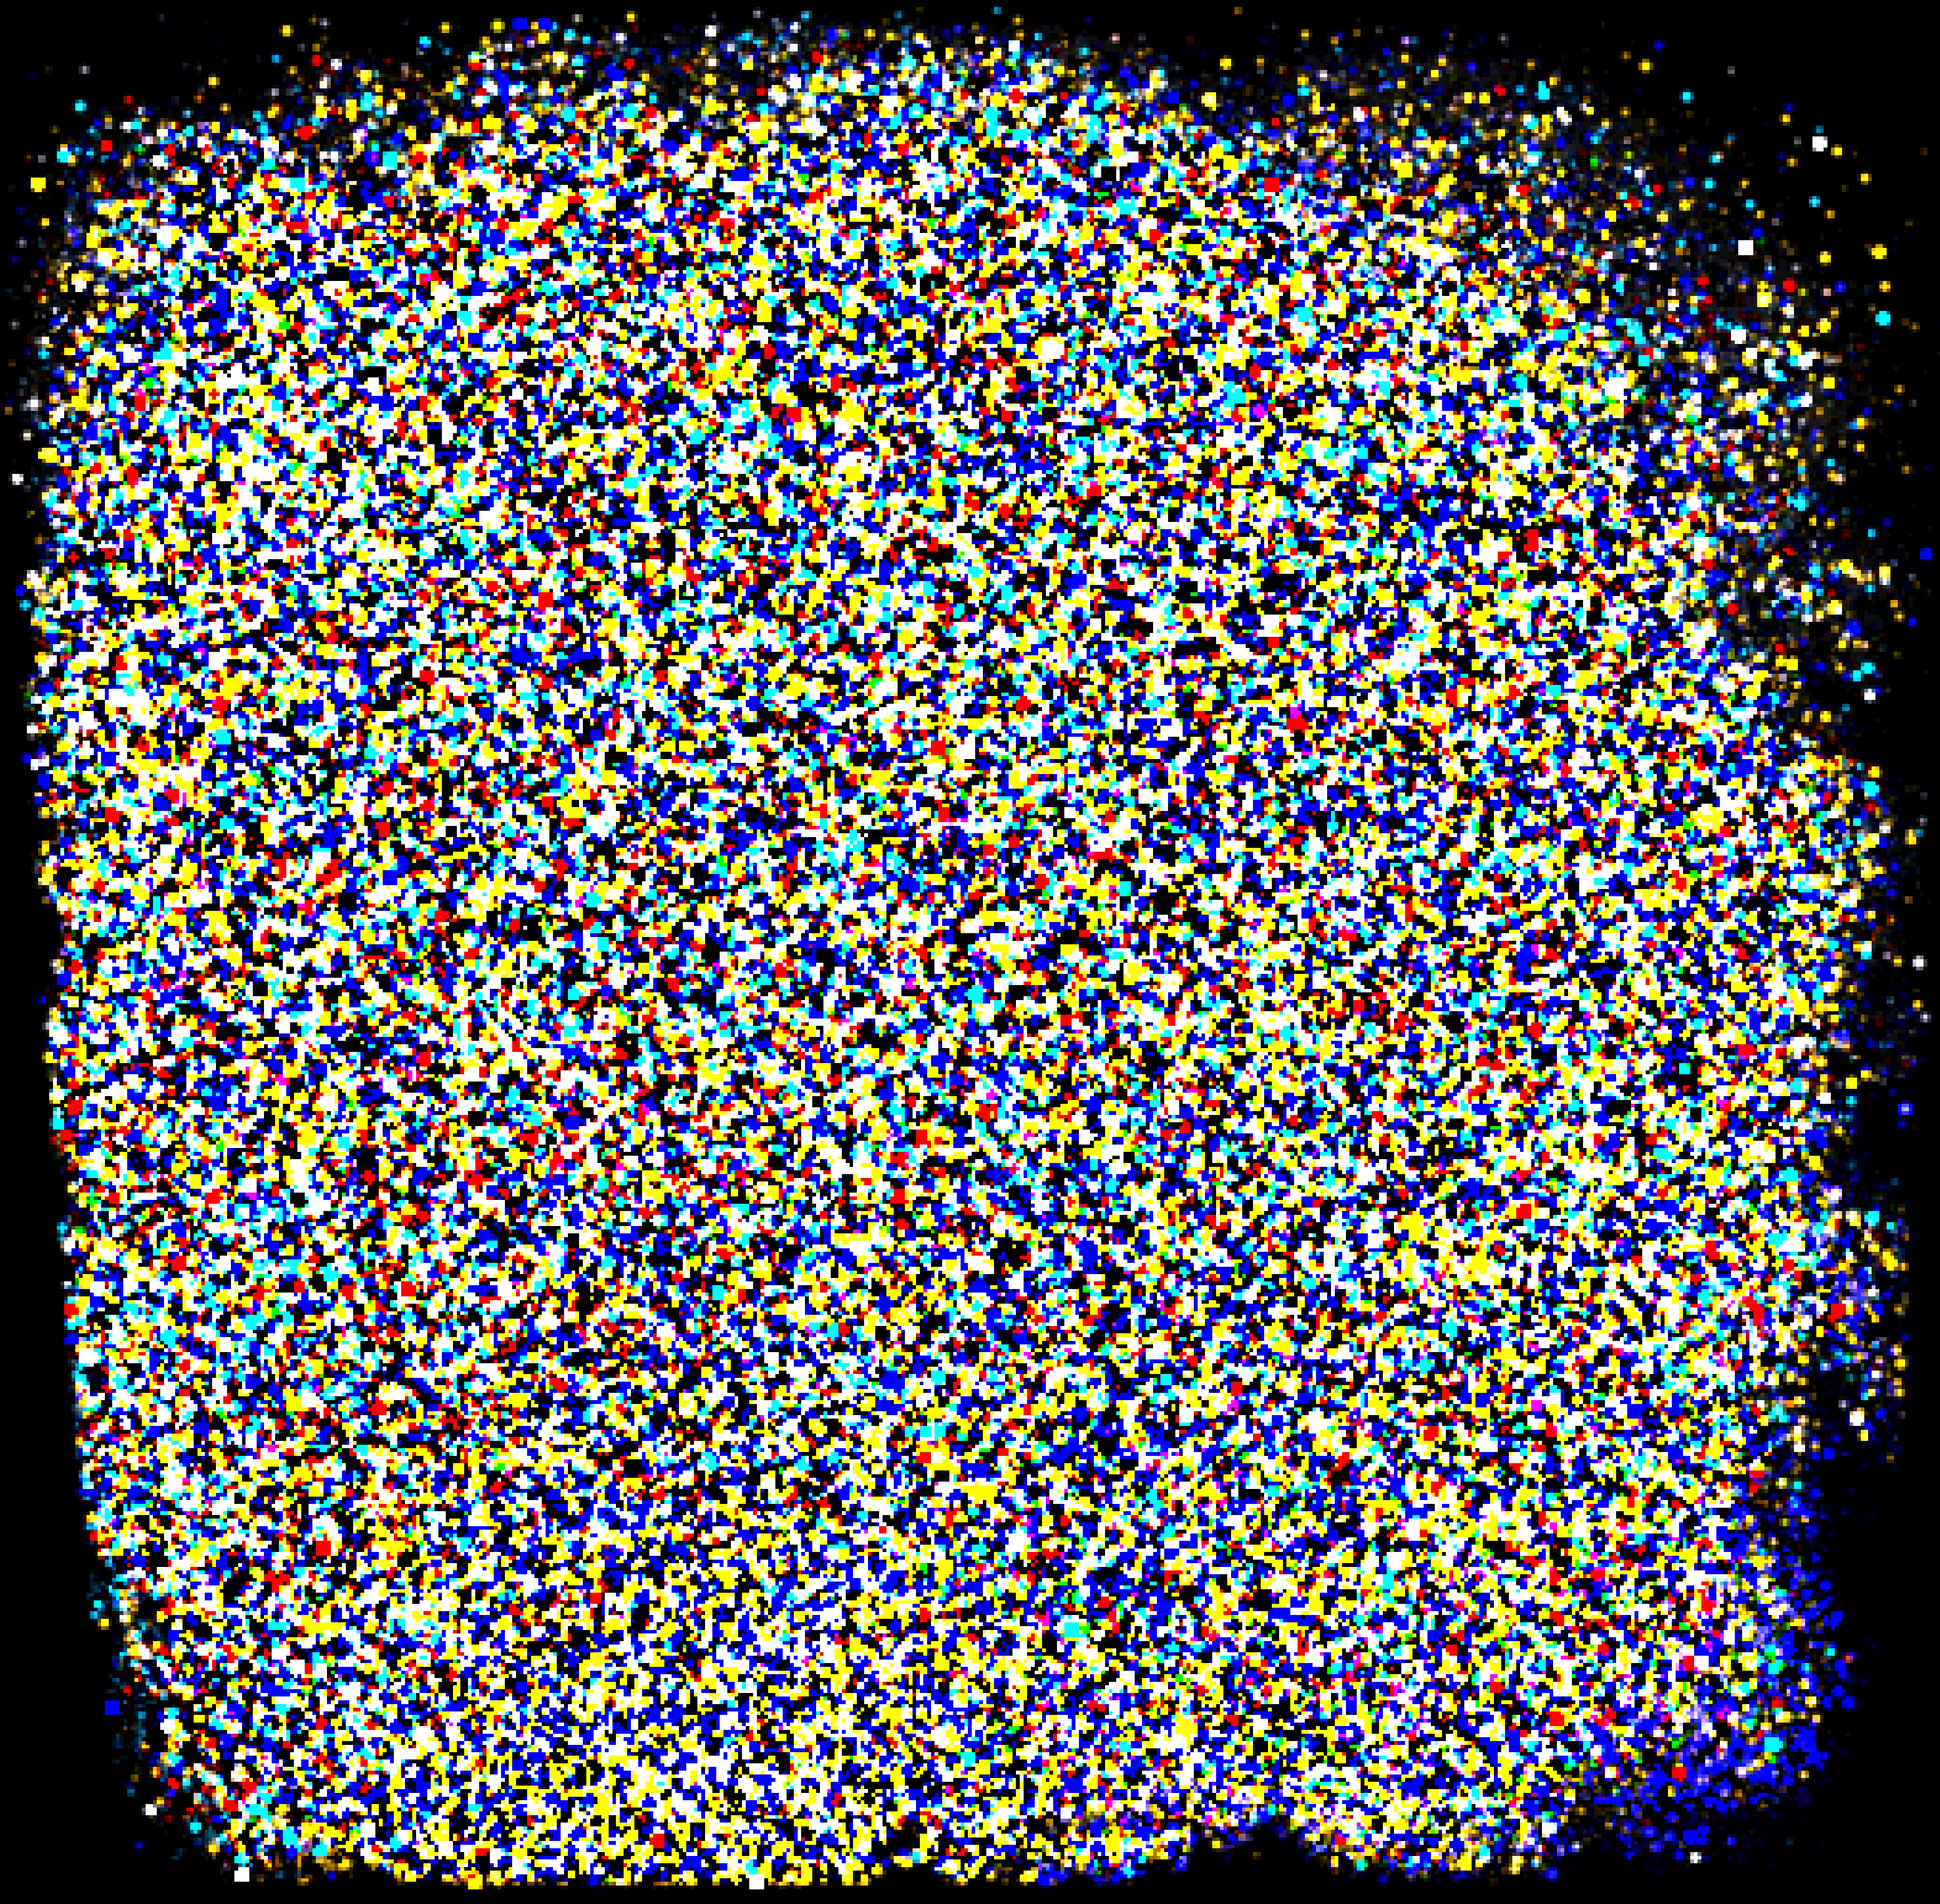

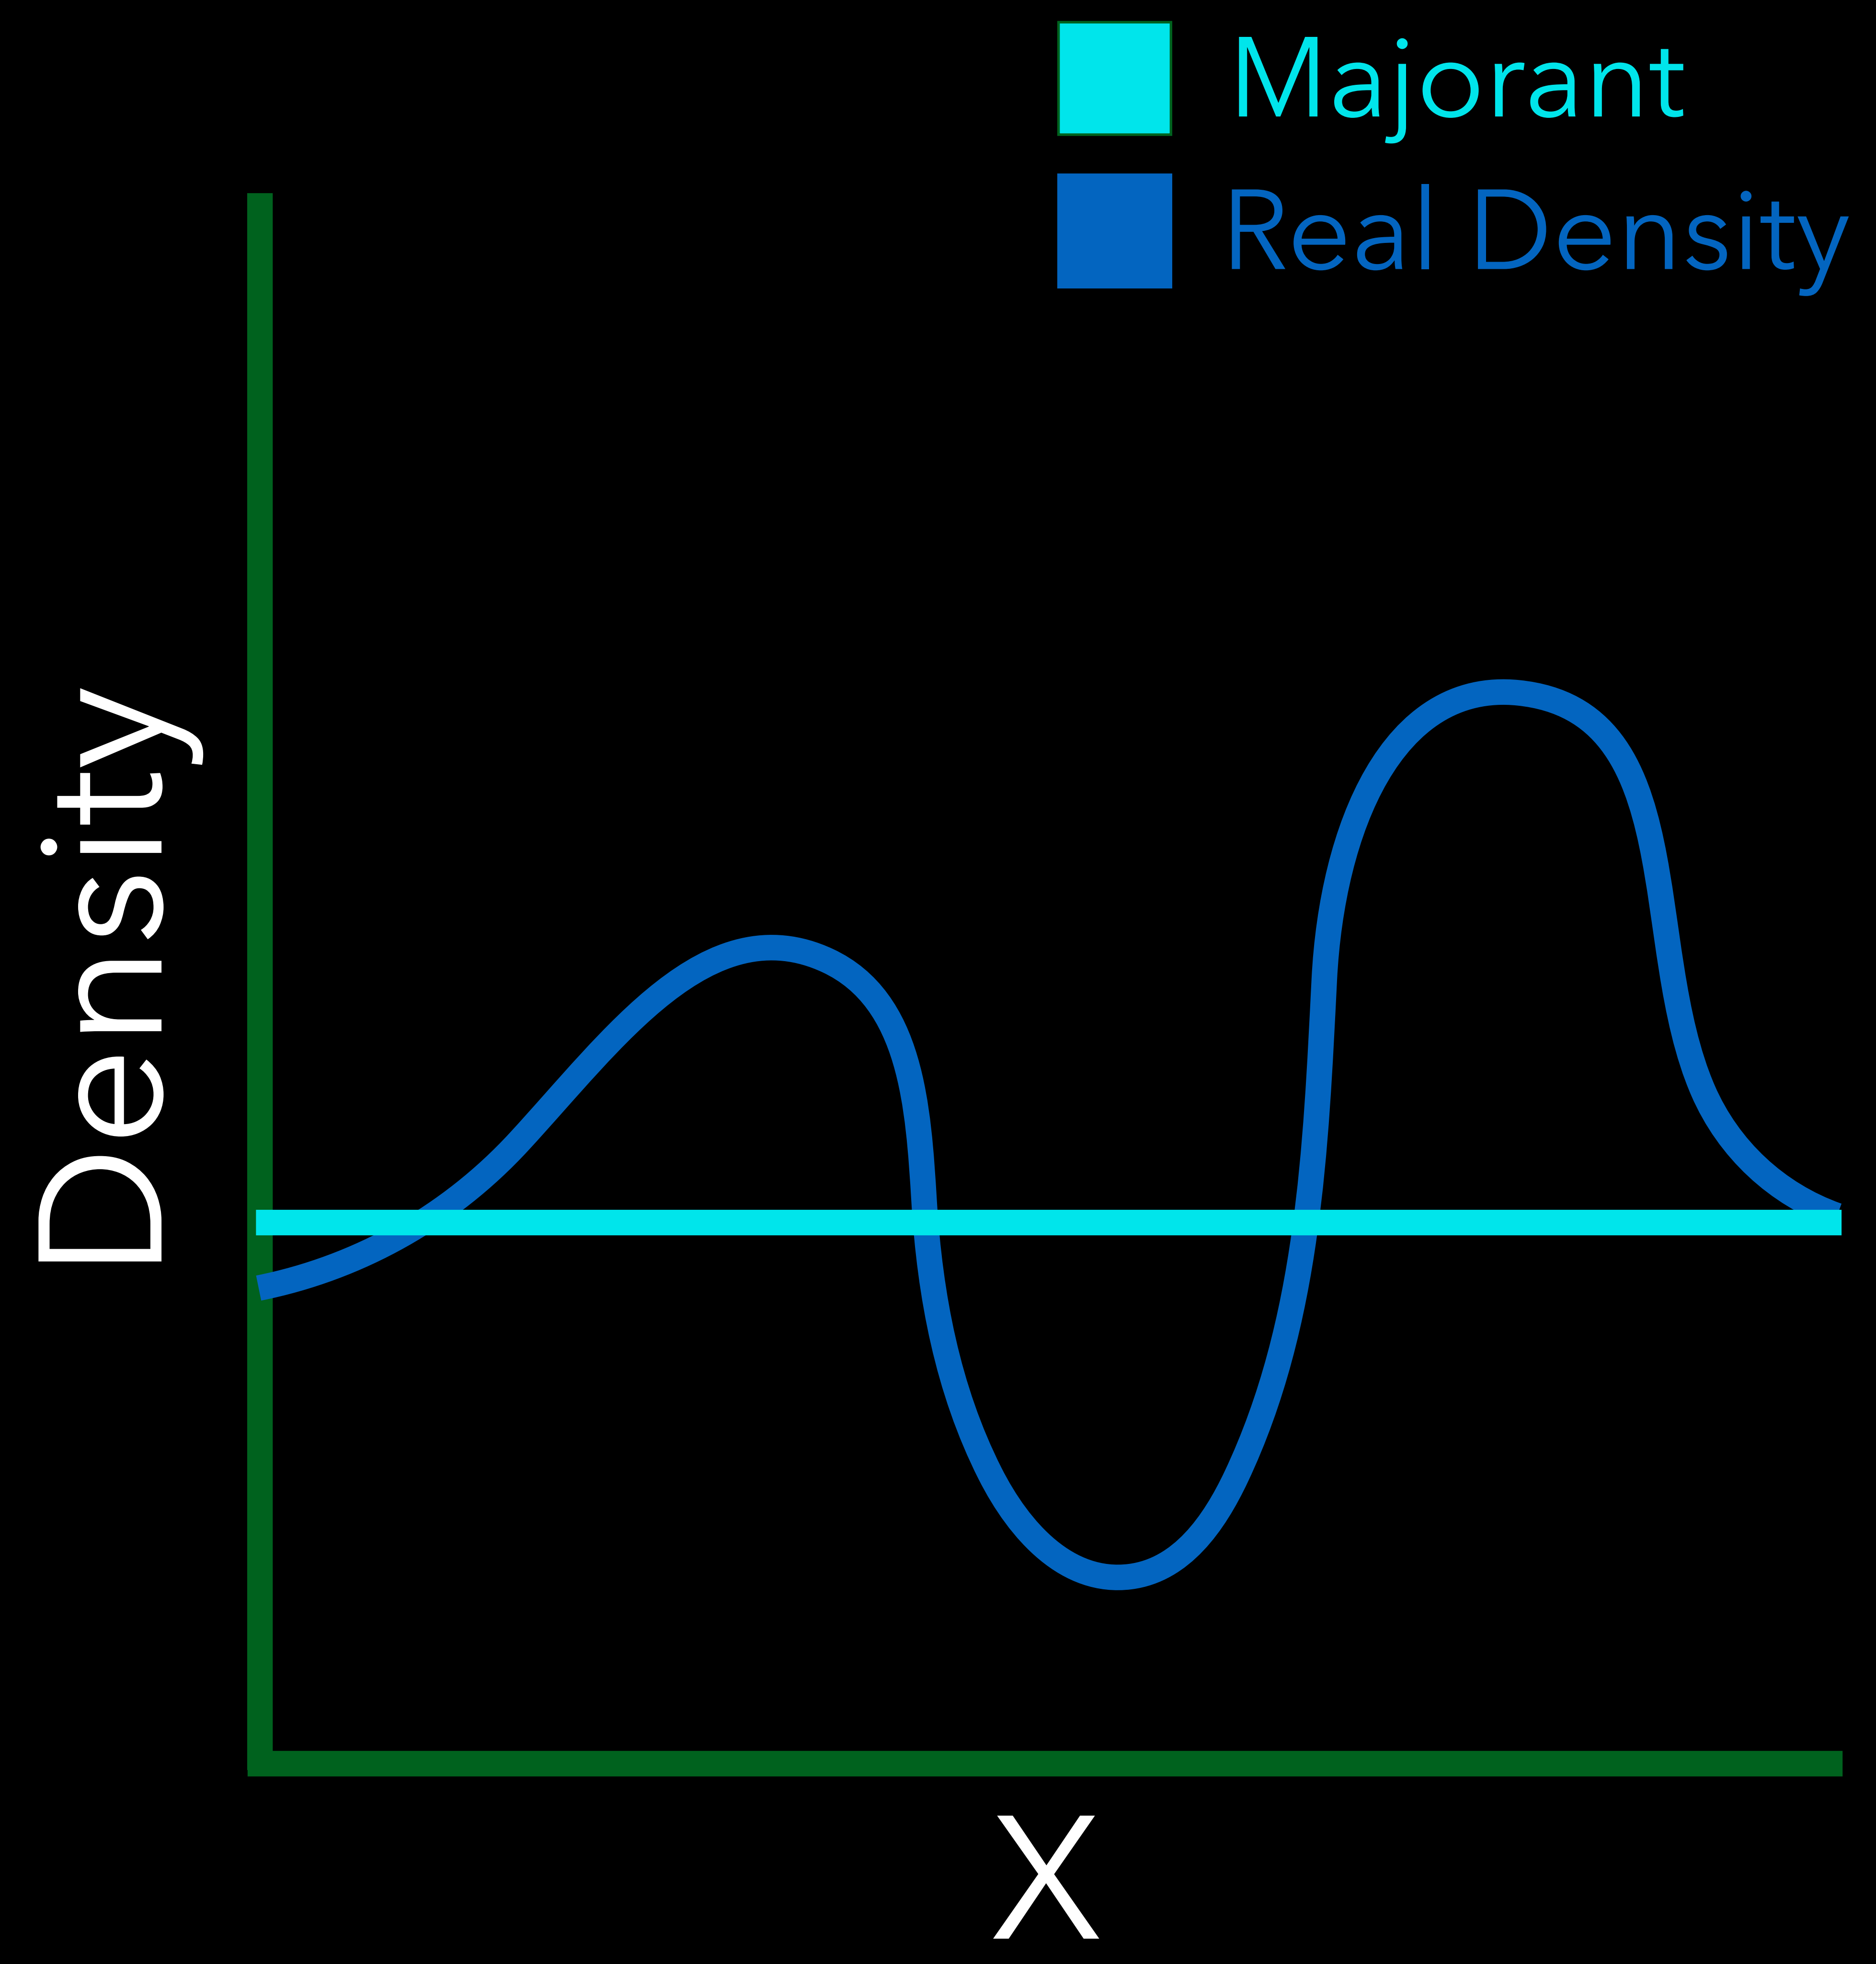
Majorant
Real Density
Density
X

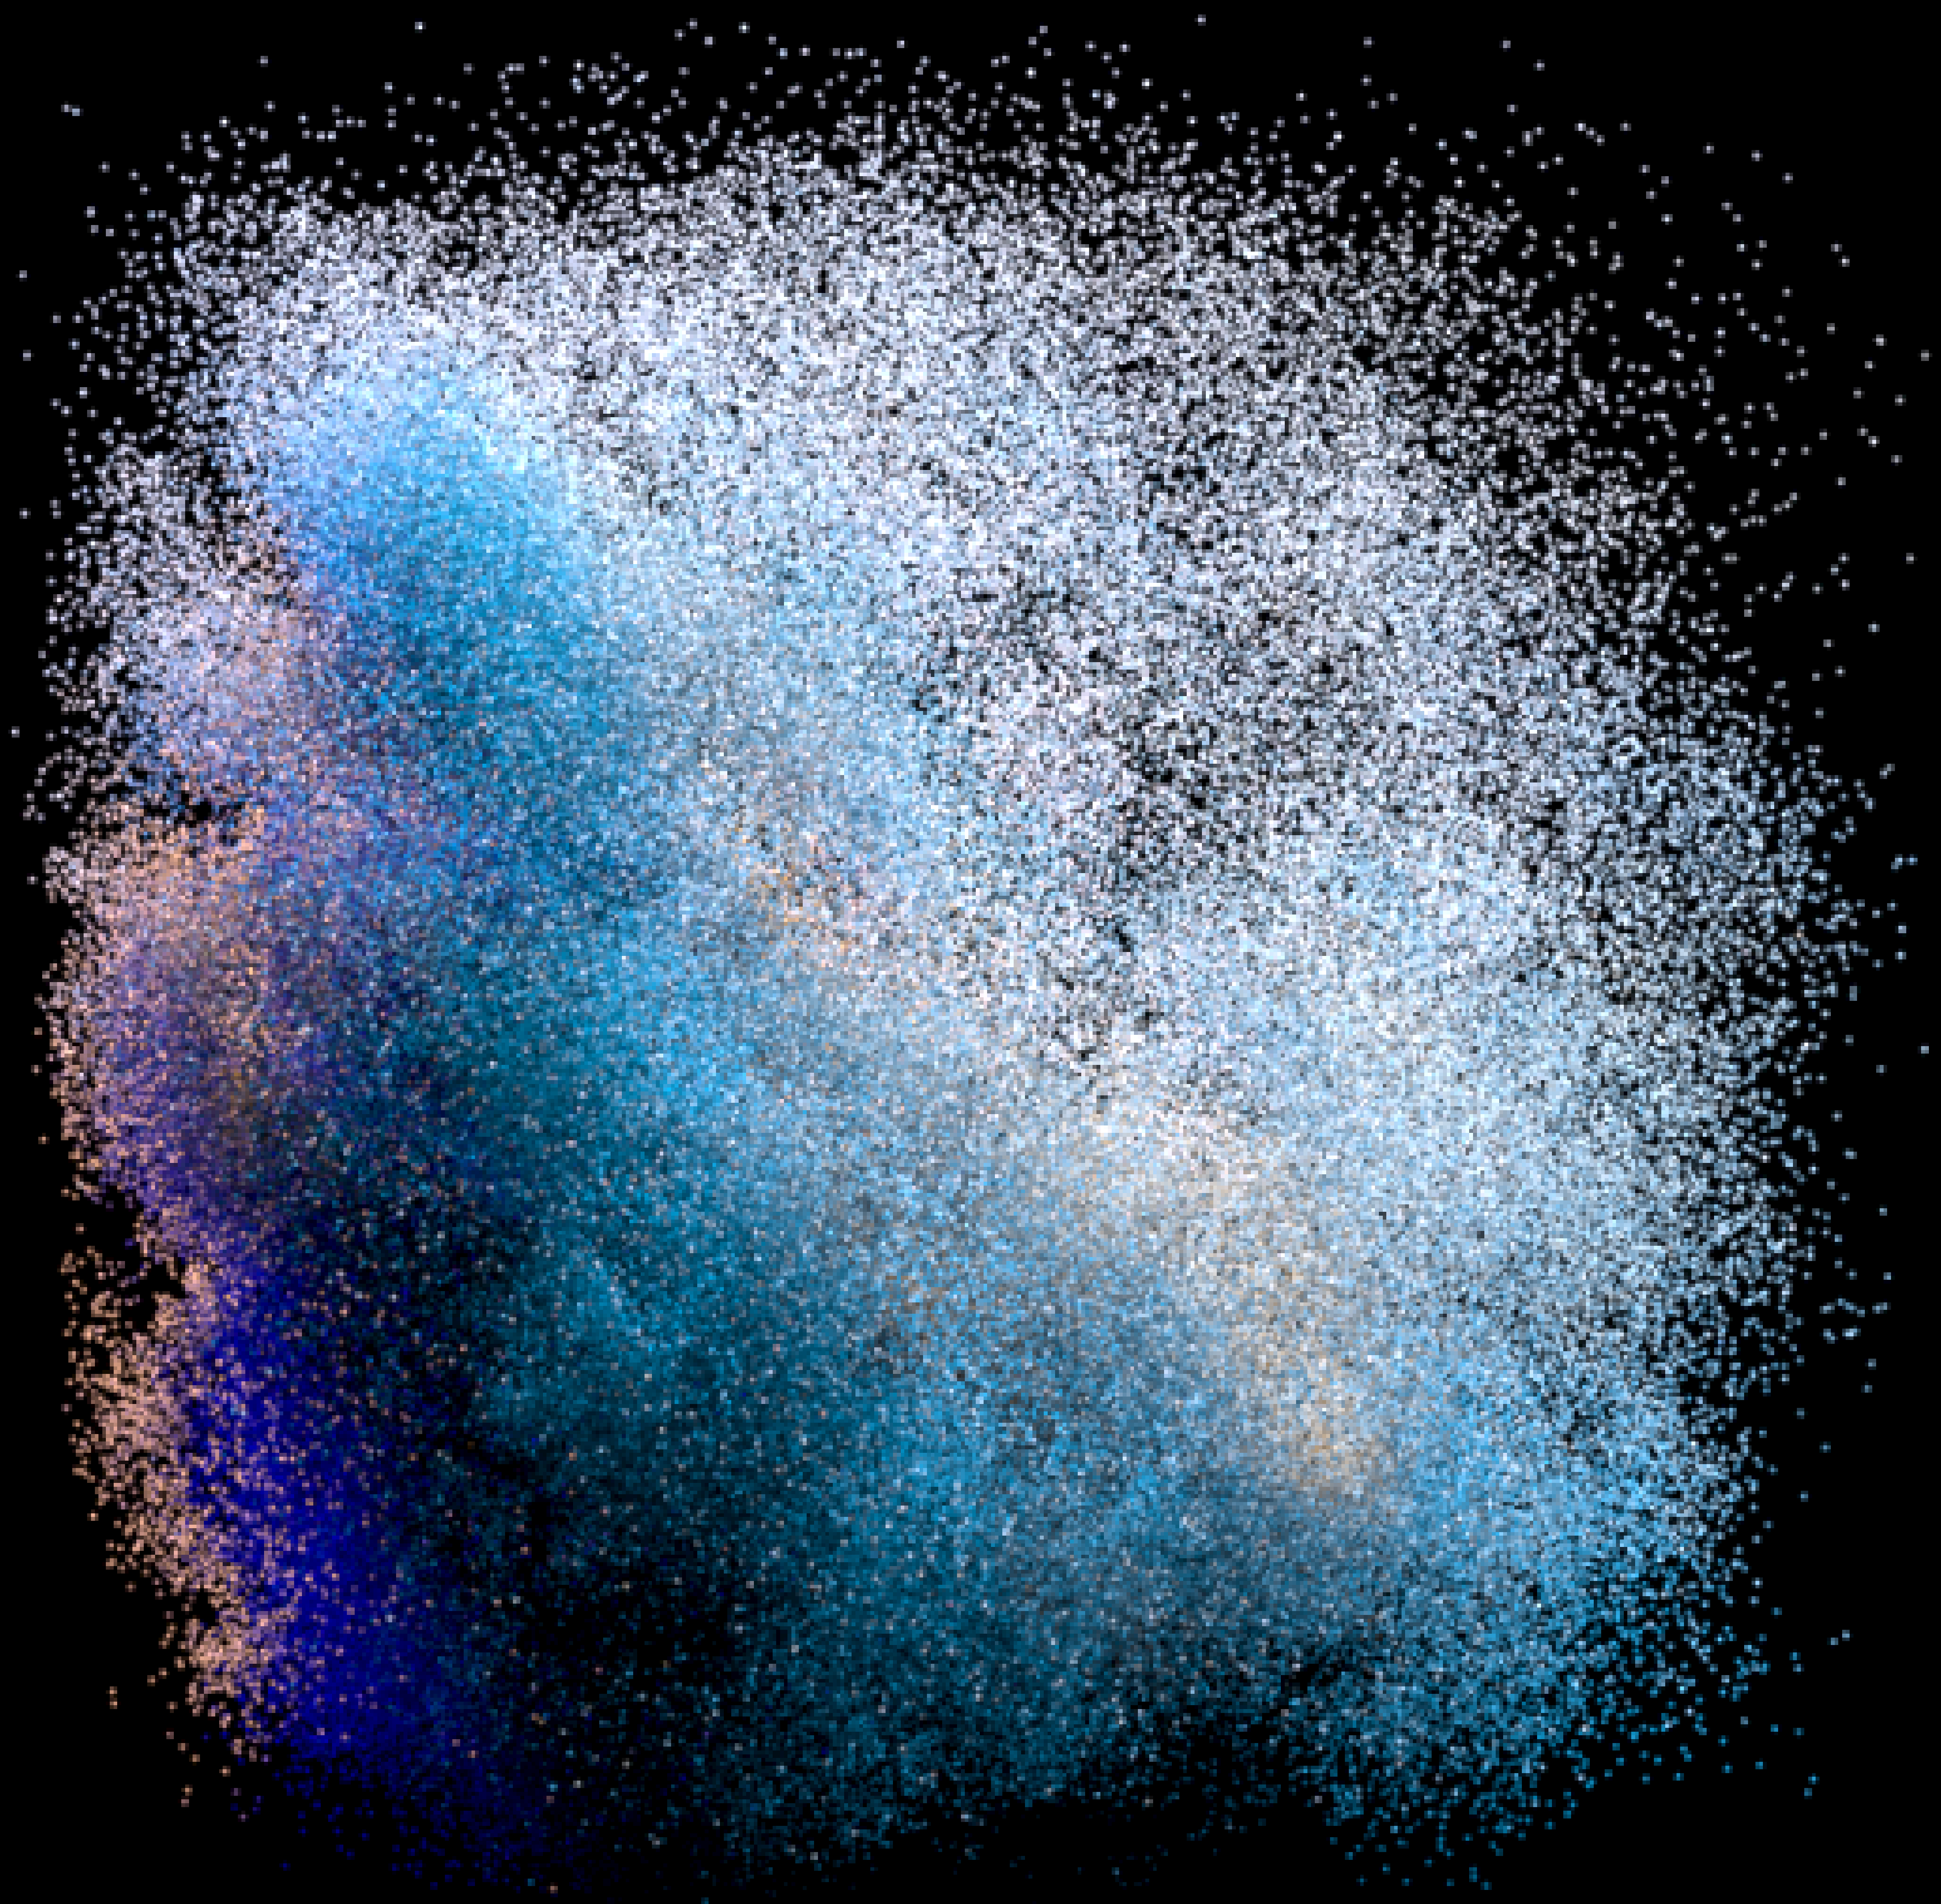

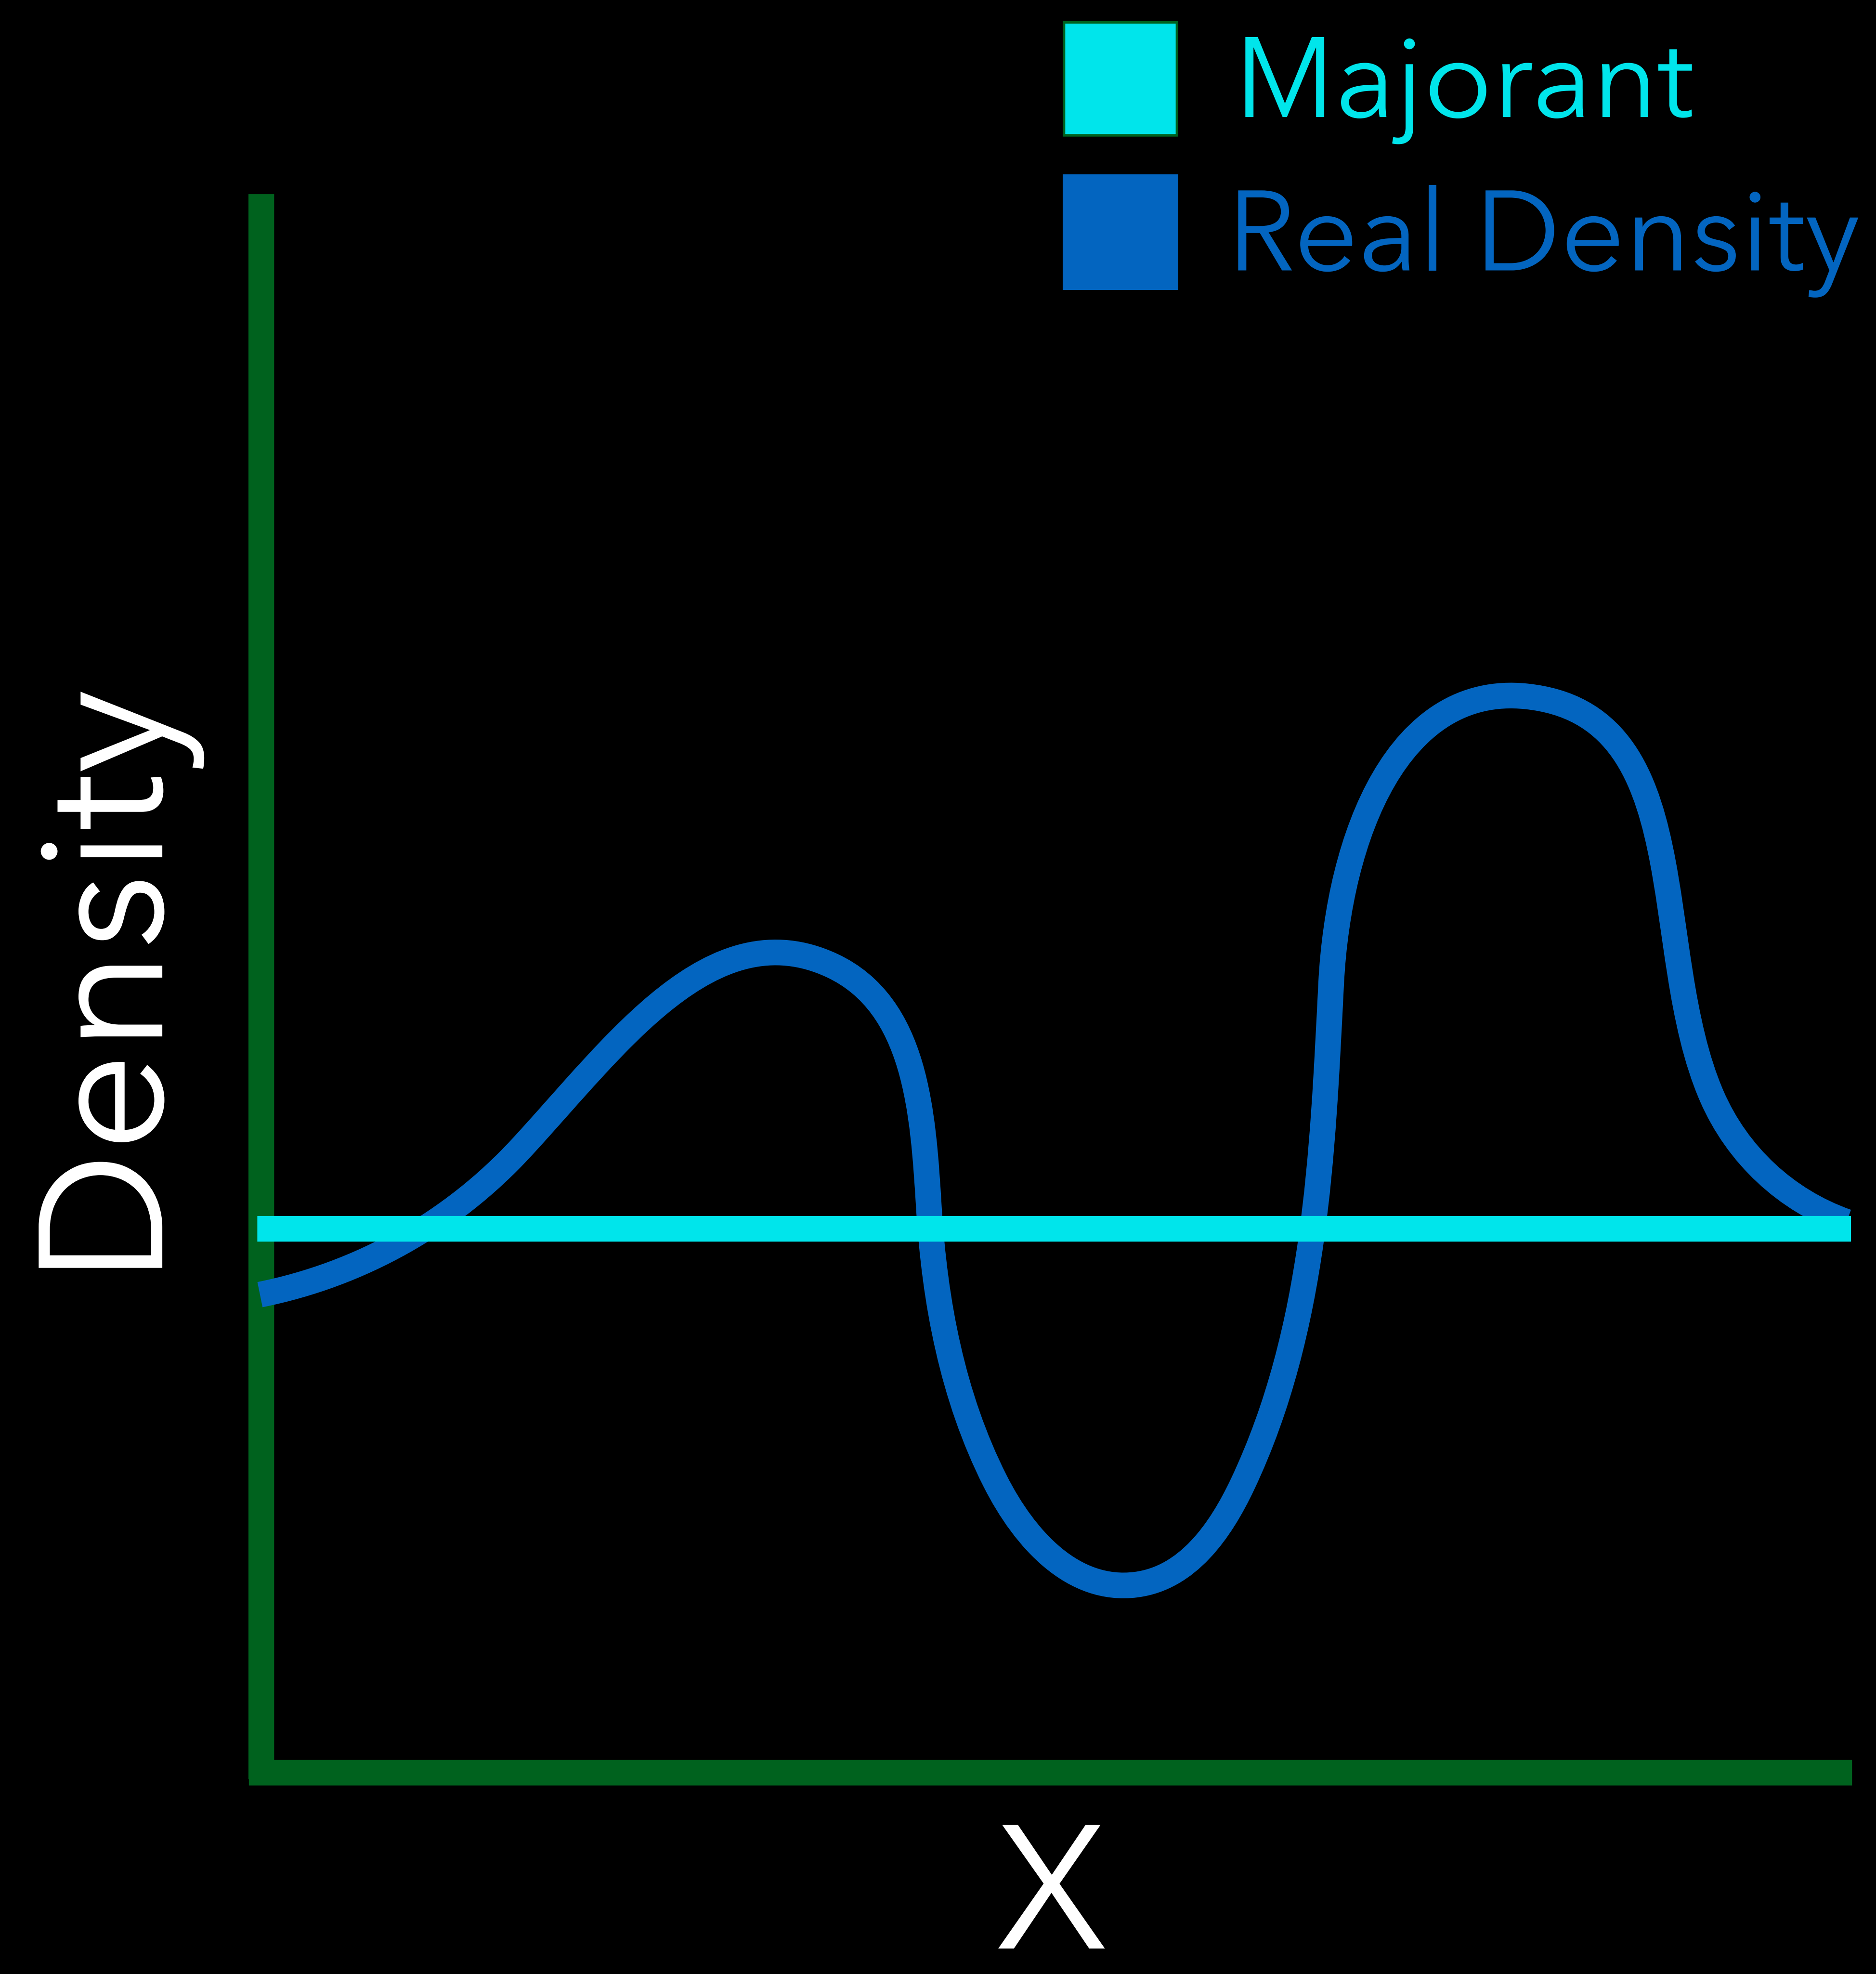
Majorant
Real Density
Density
X

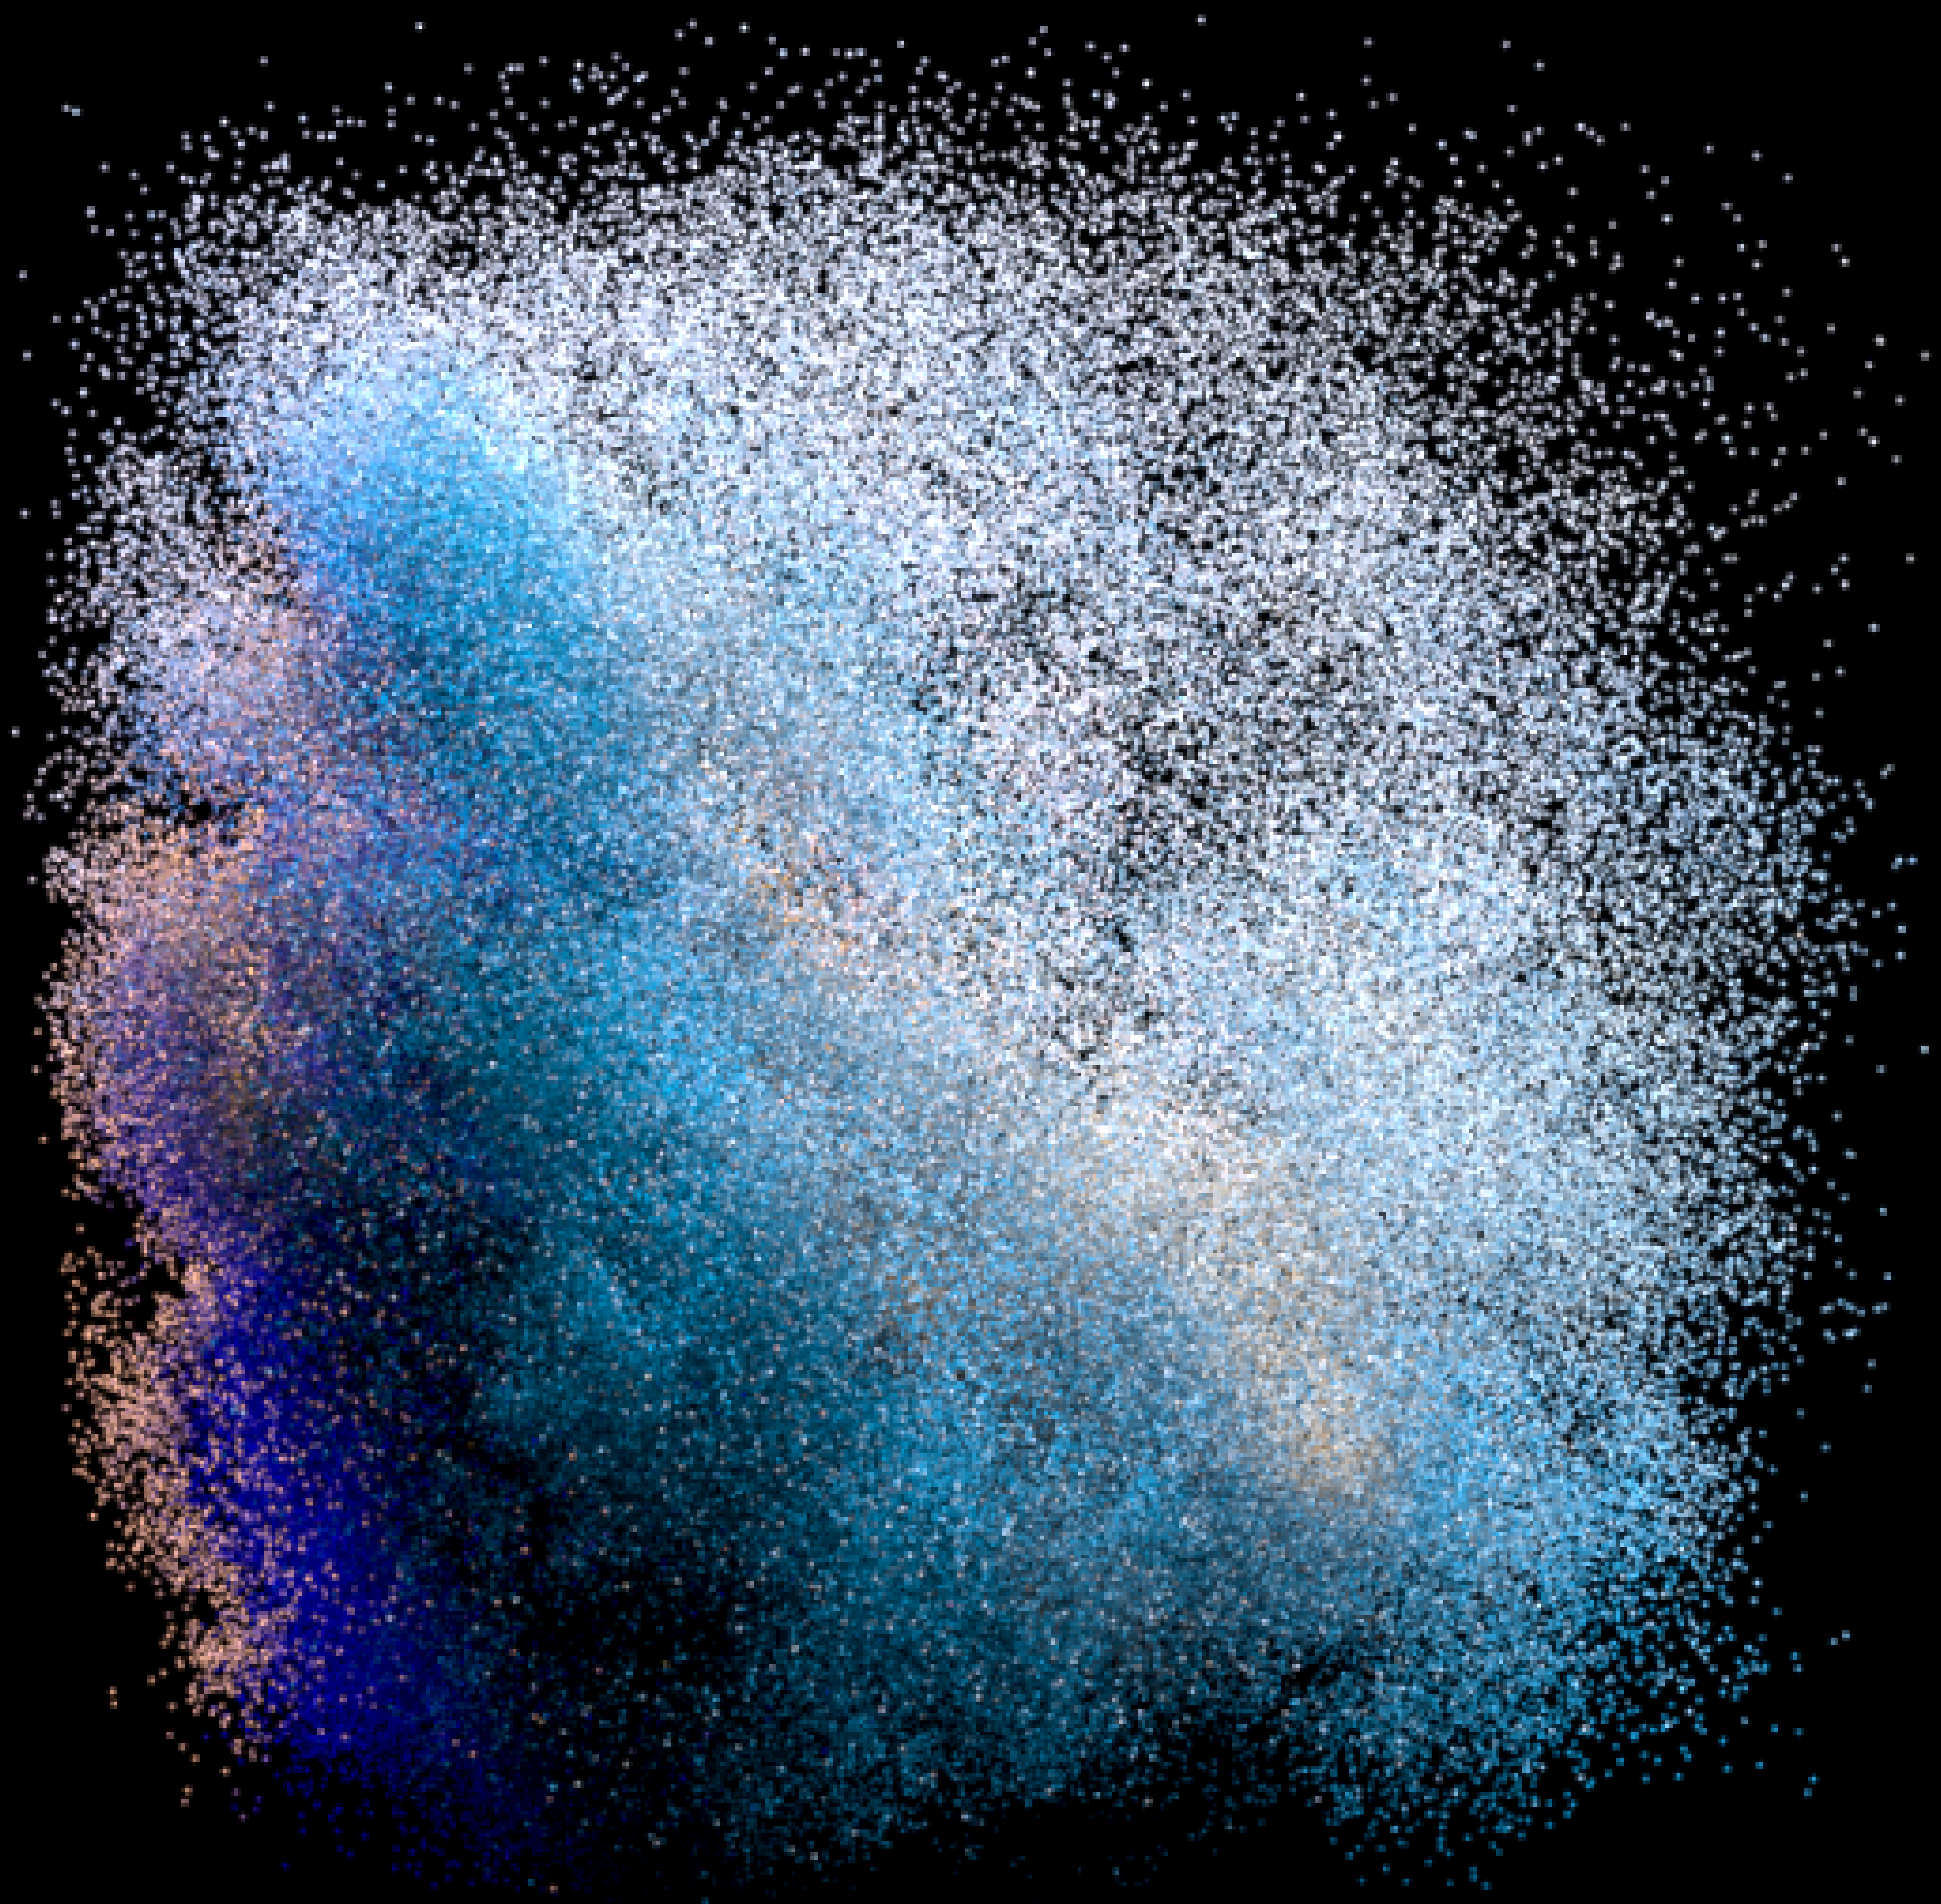

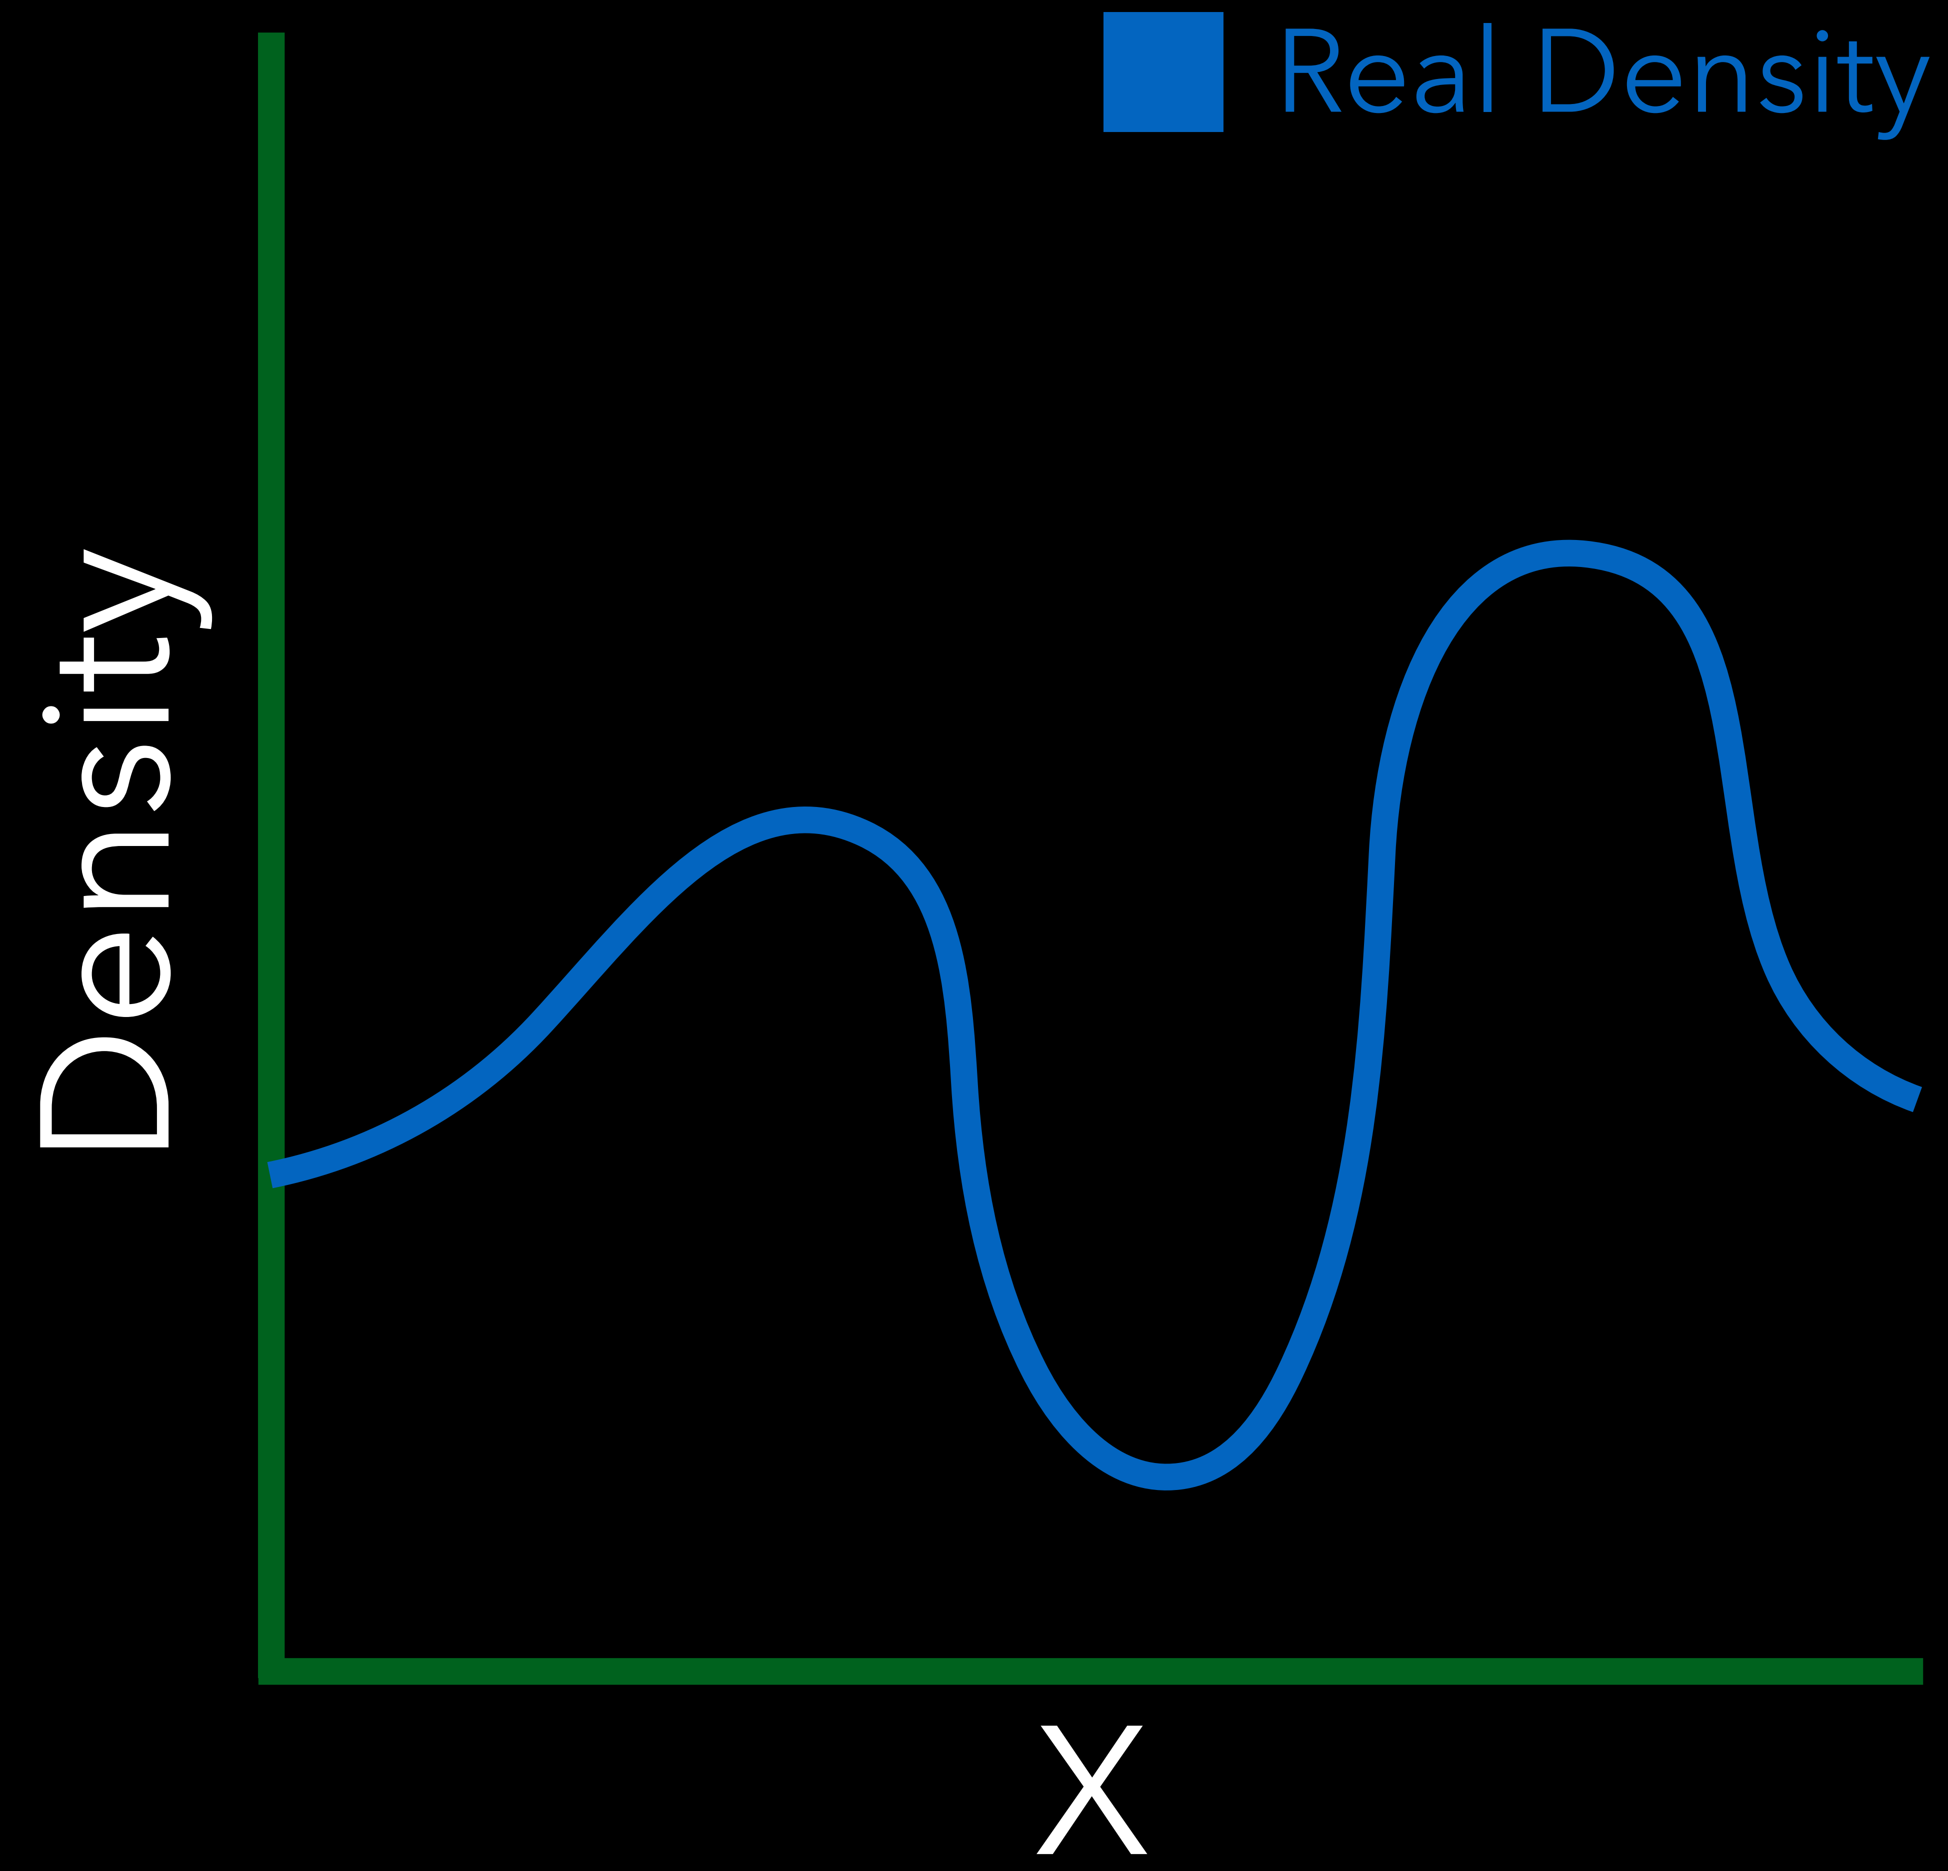

Real Density
Density
X

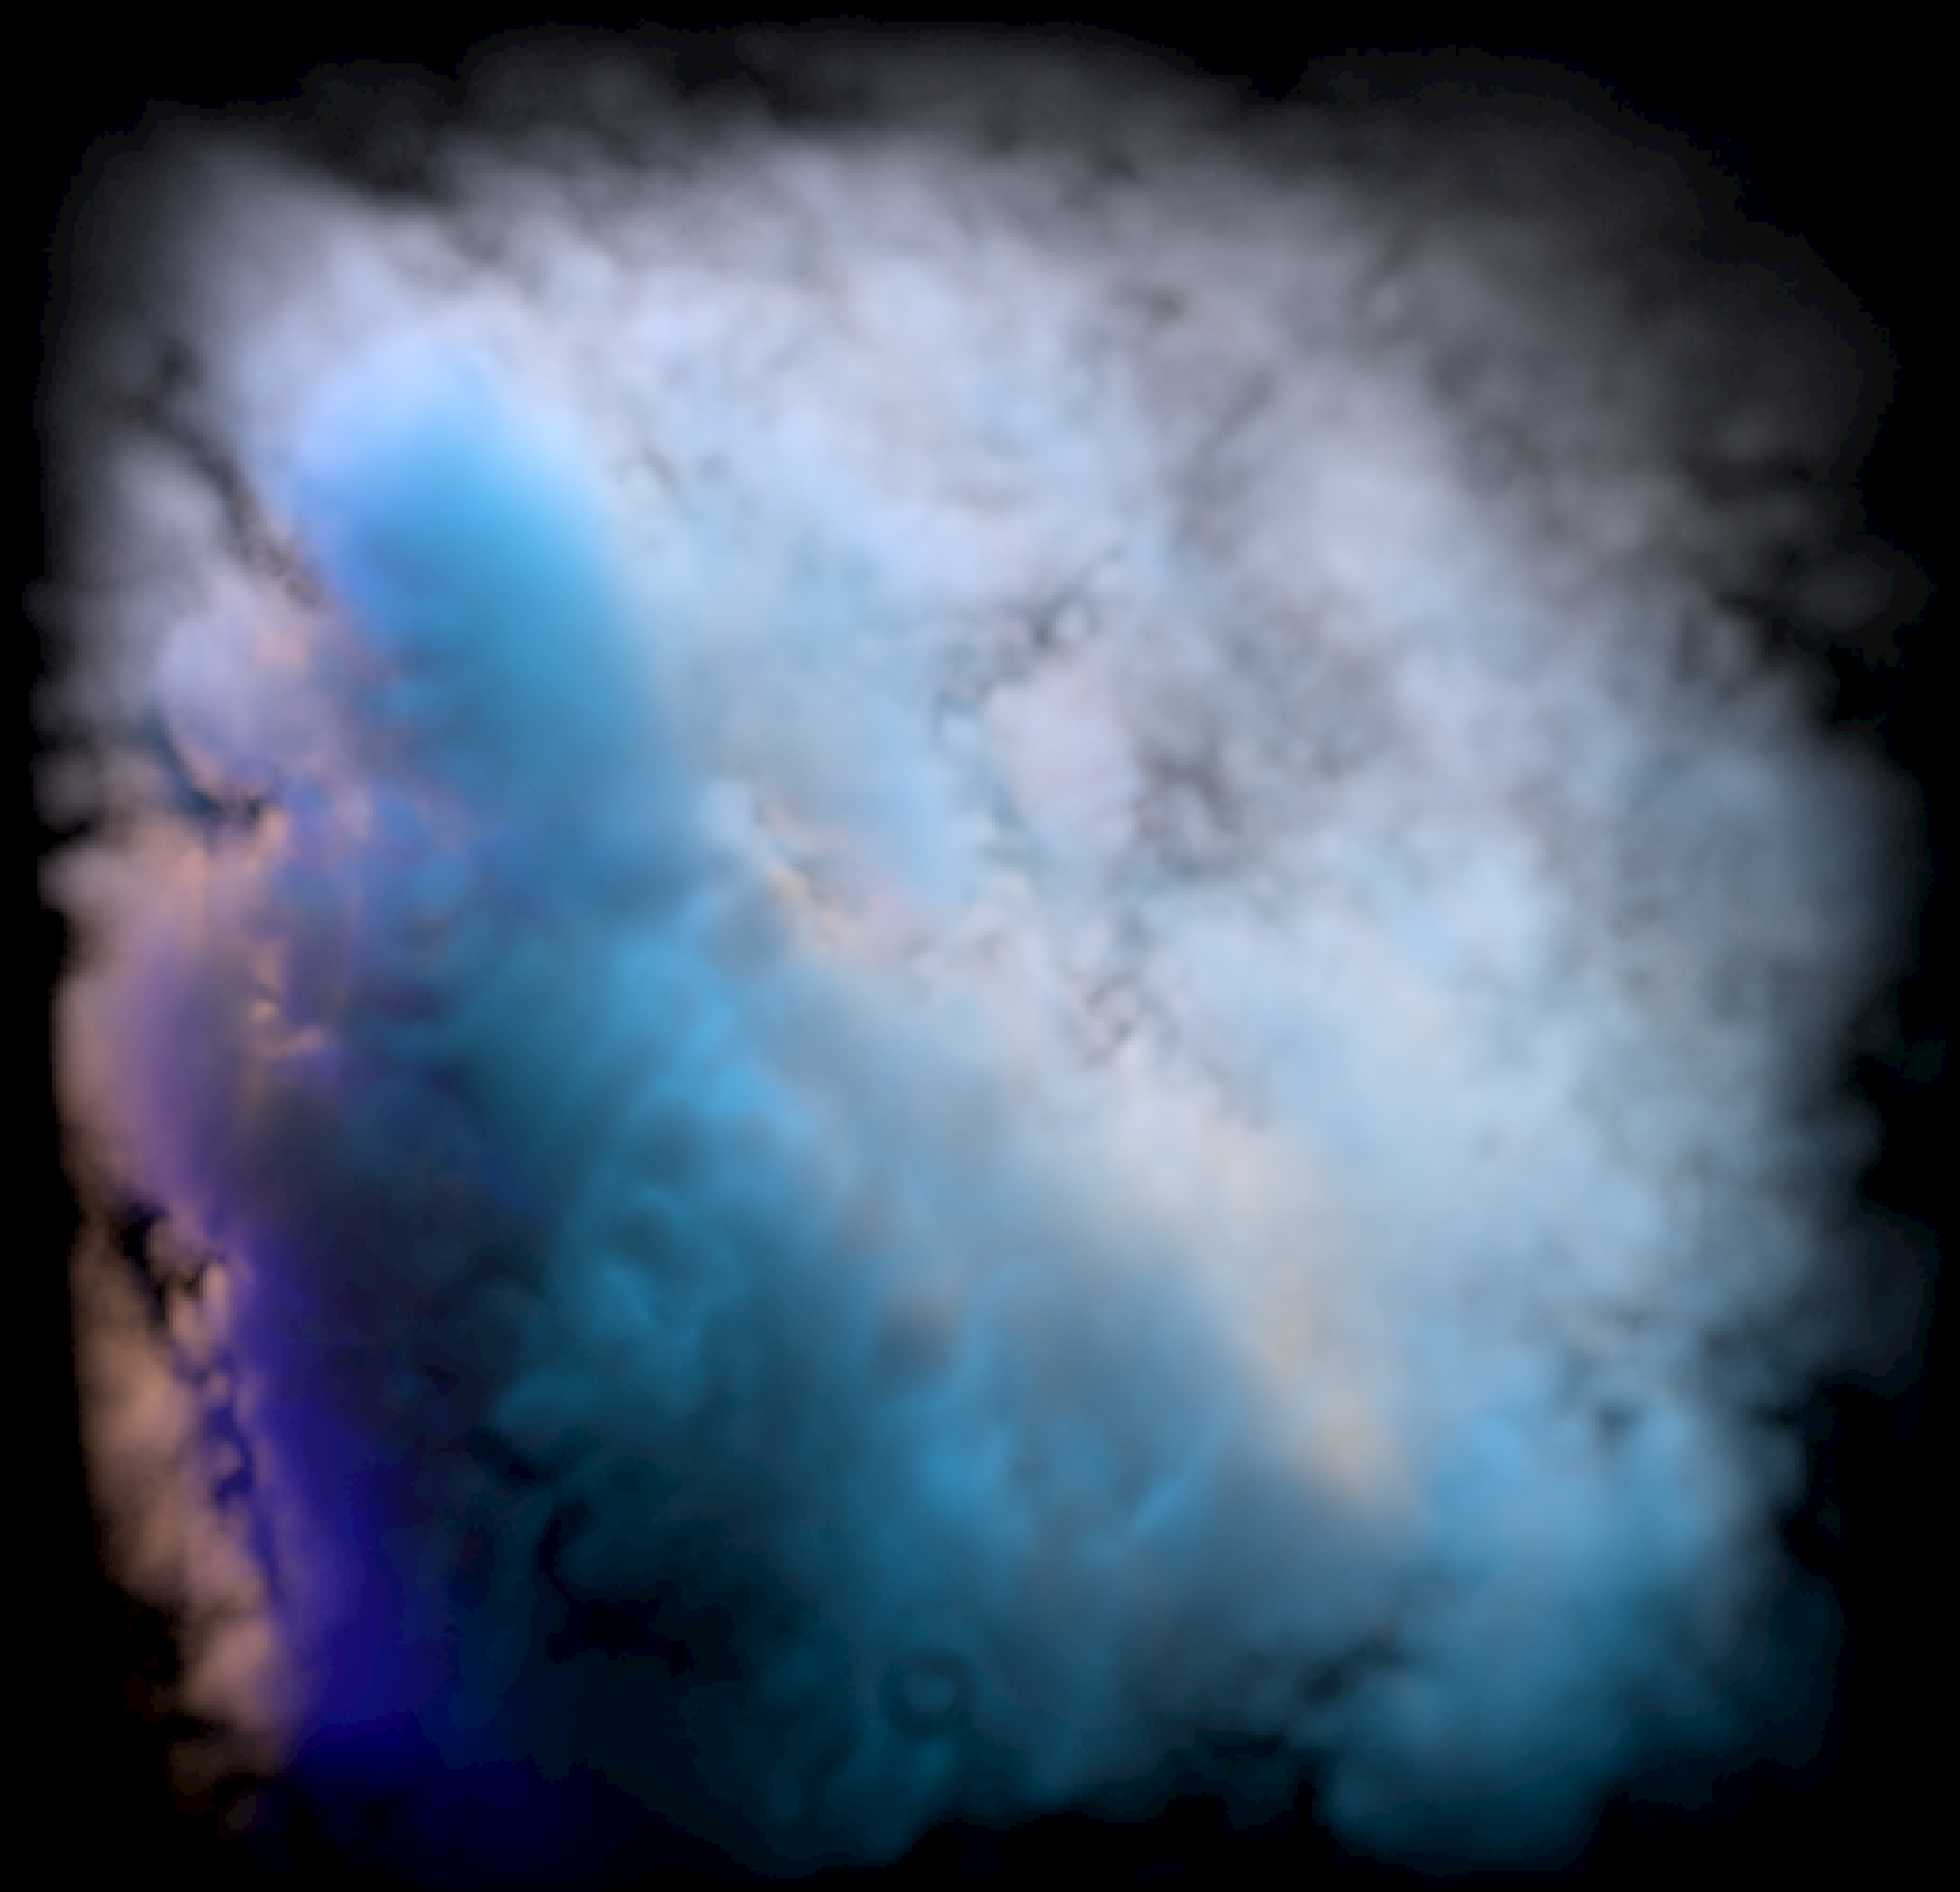

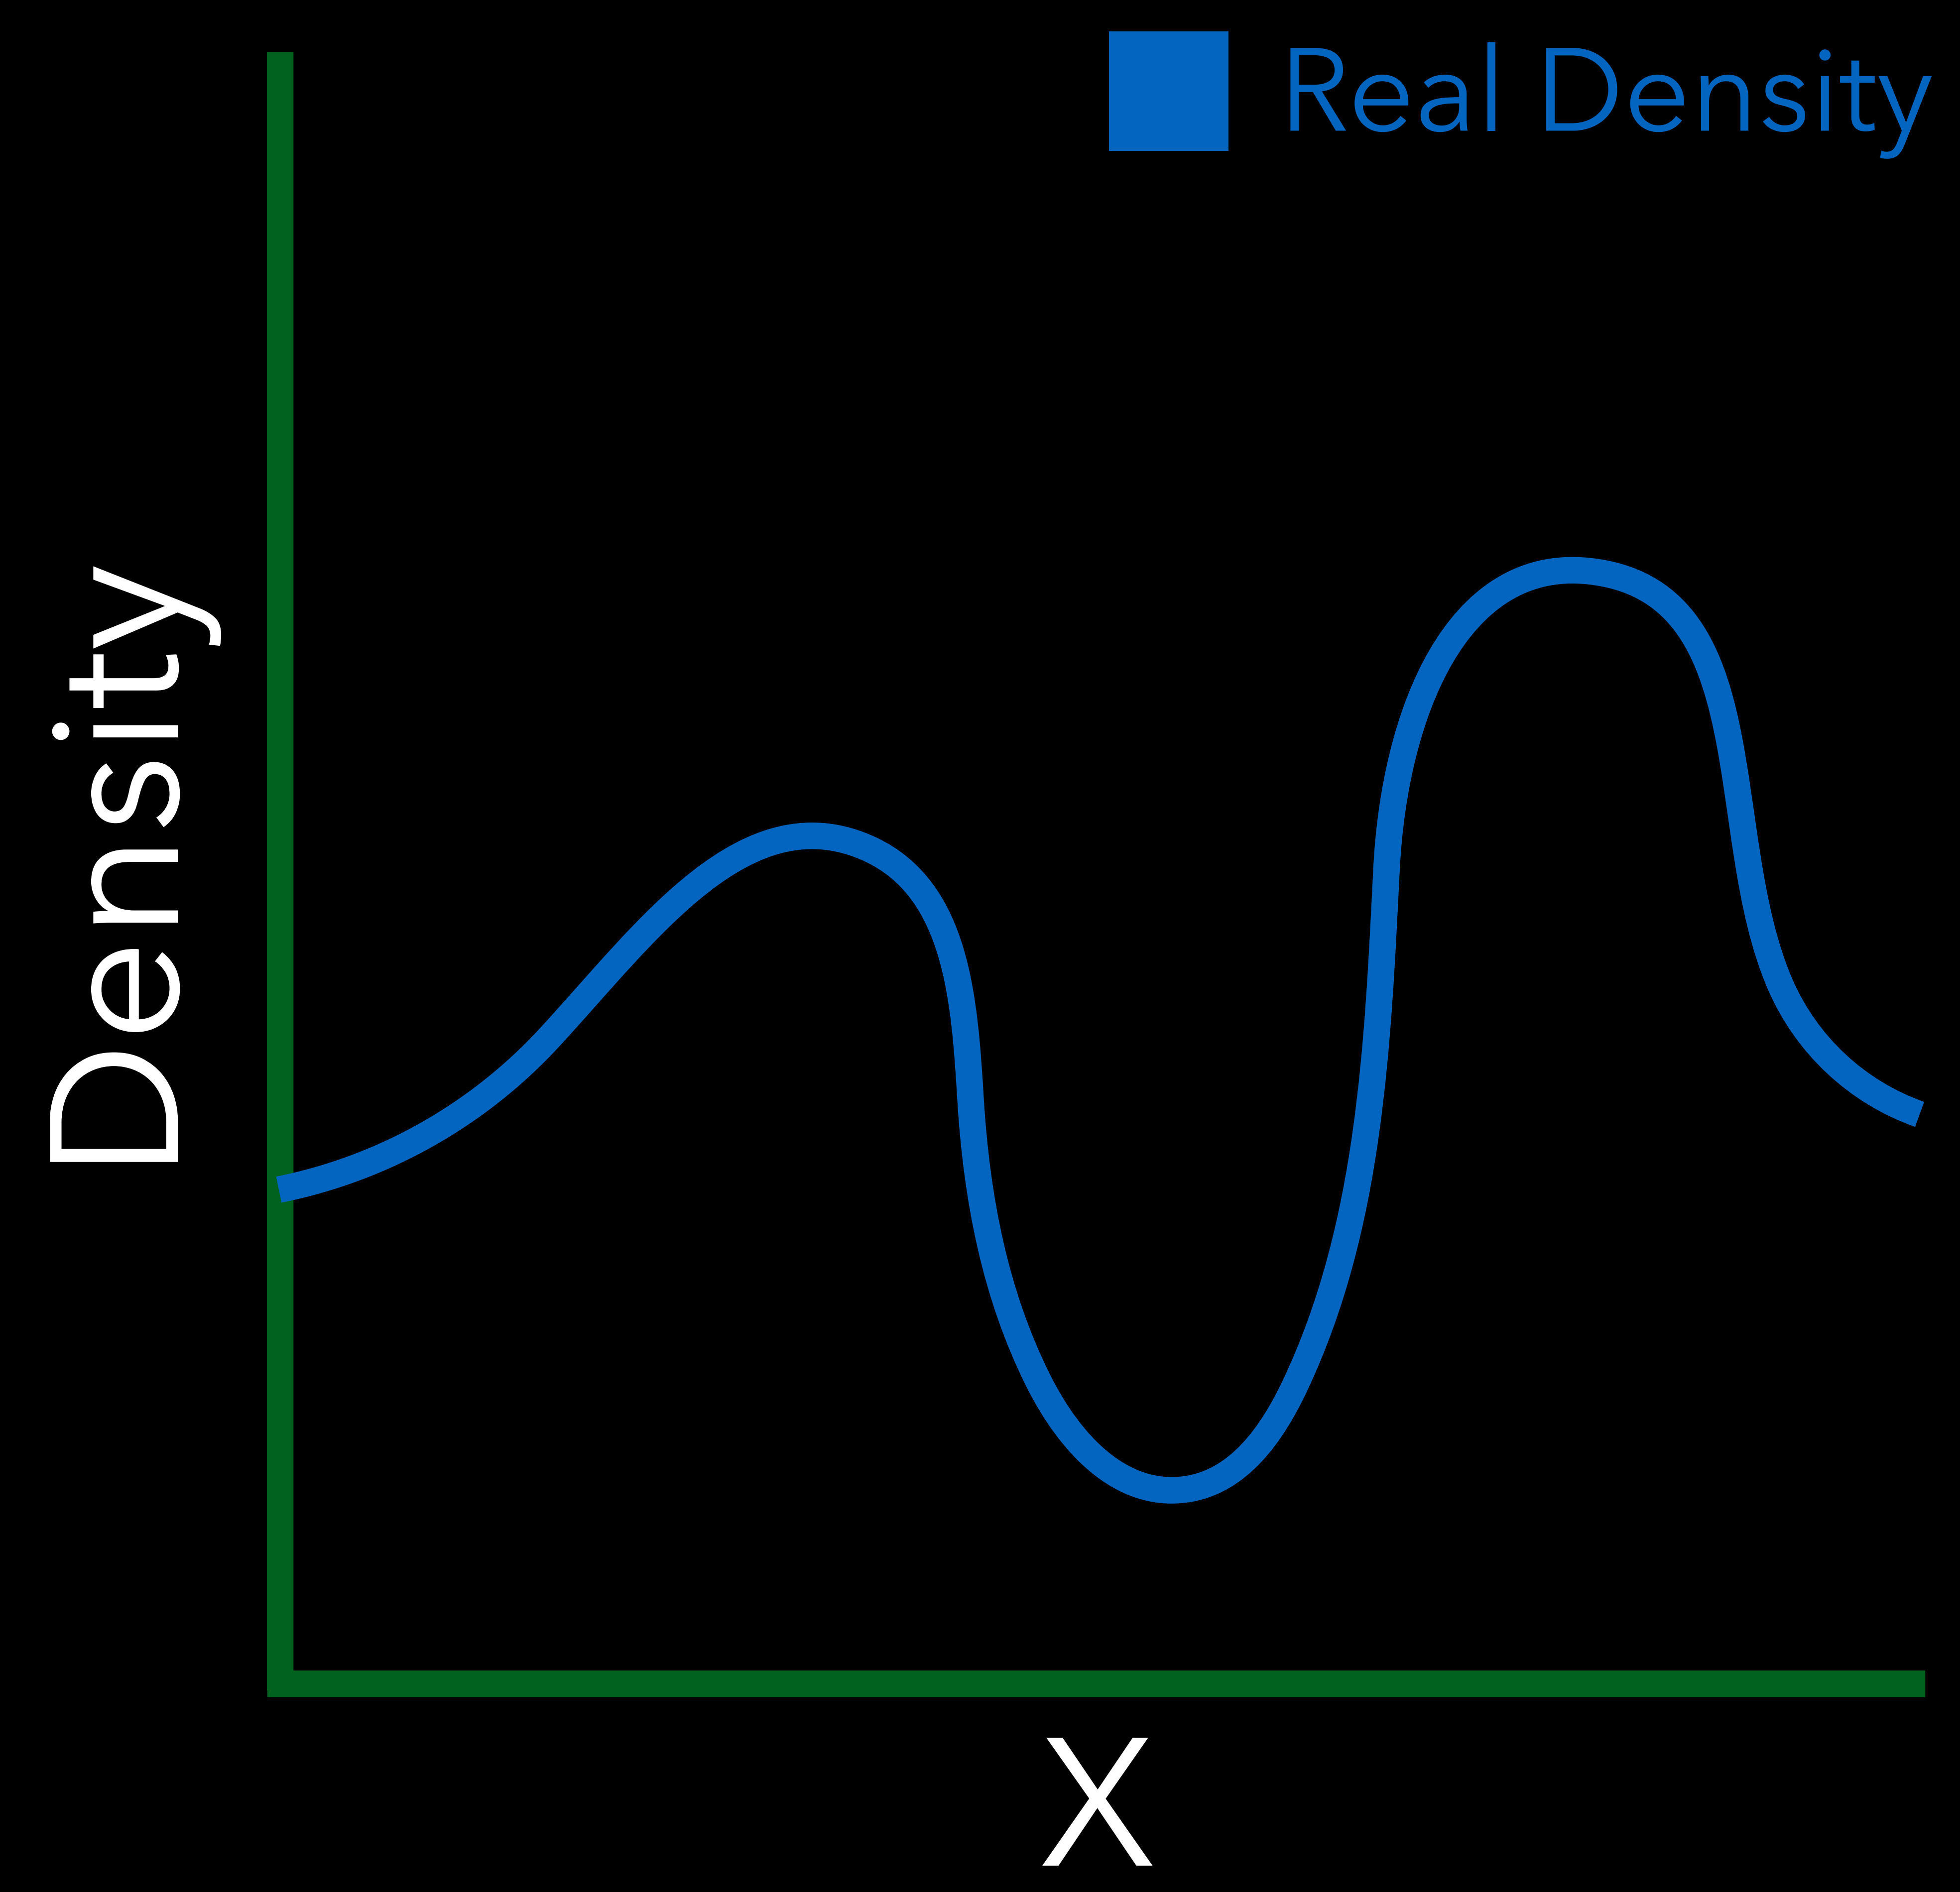
Real Density
Density
X

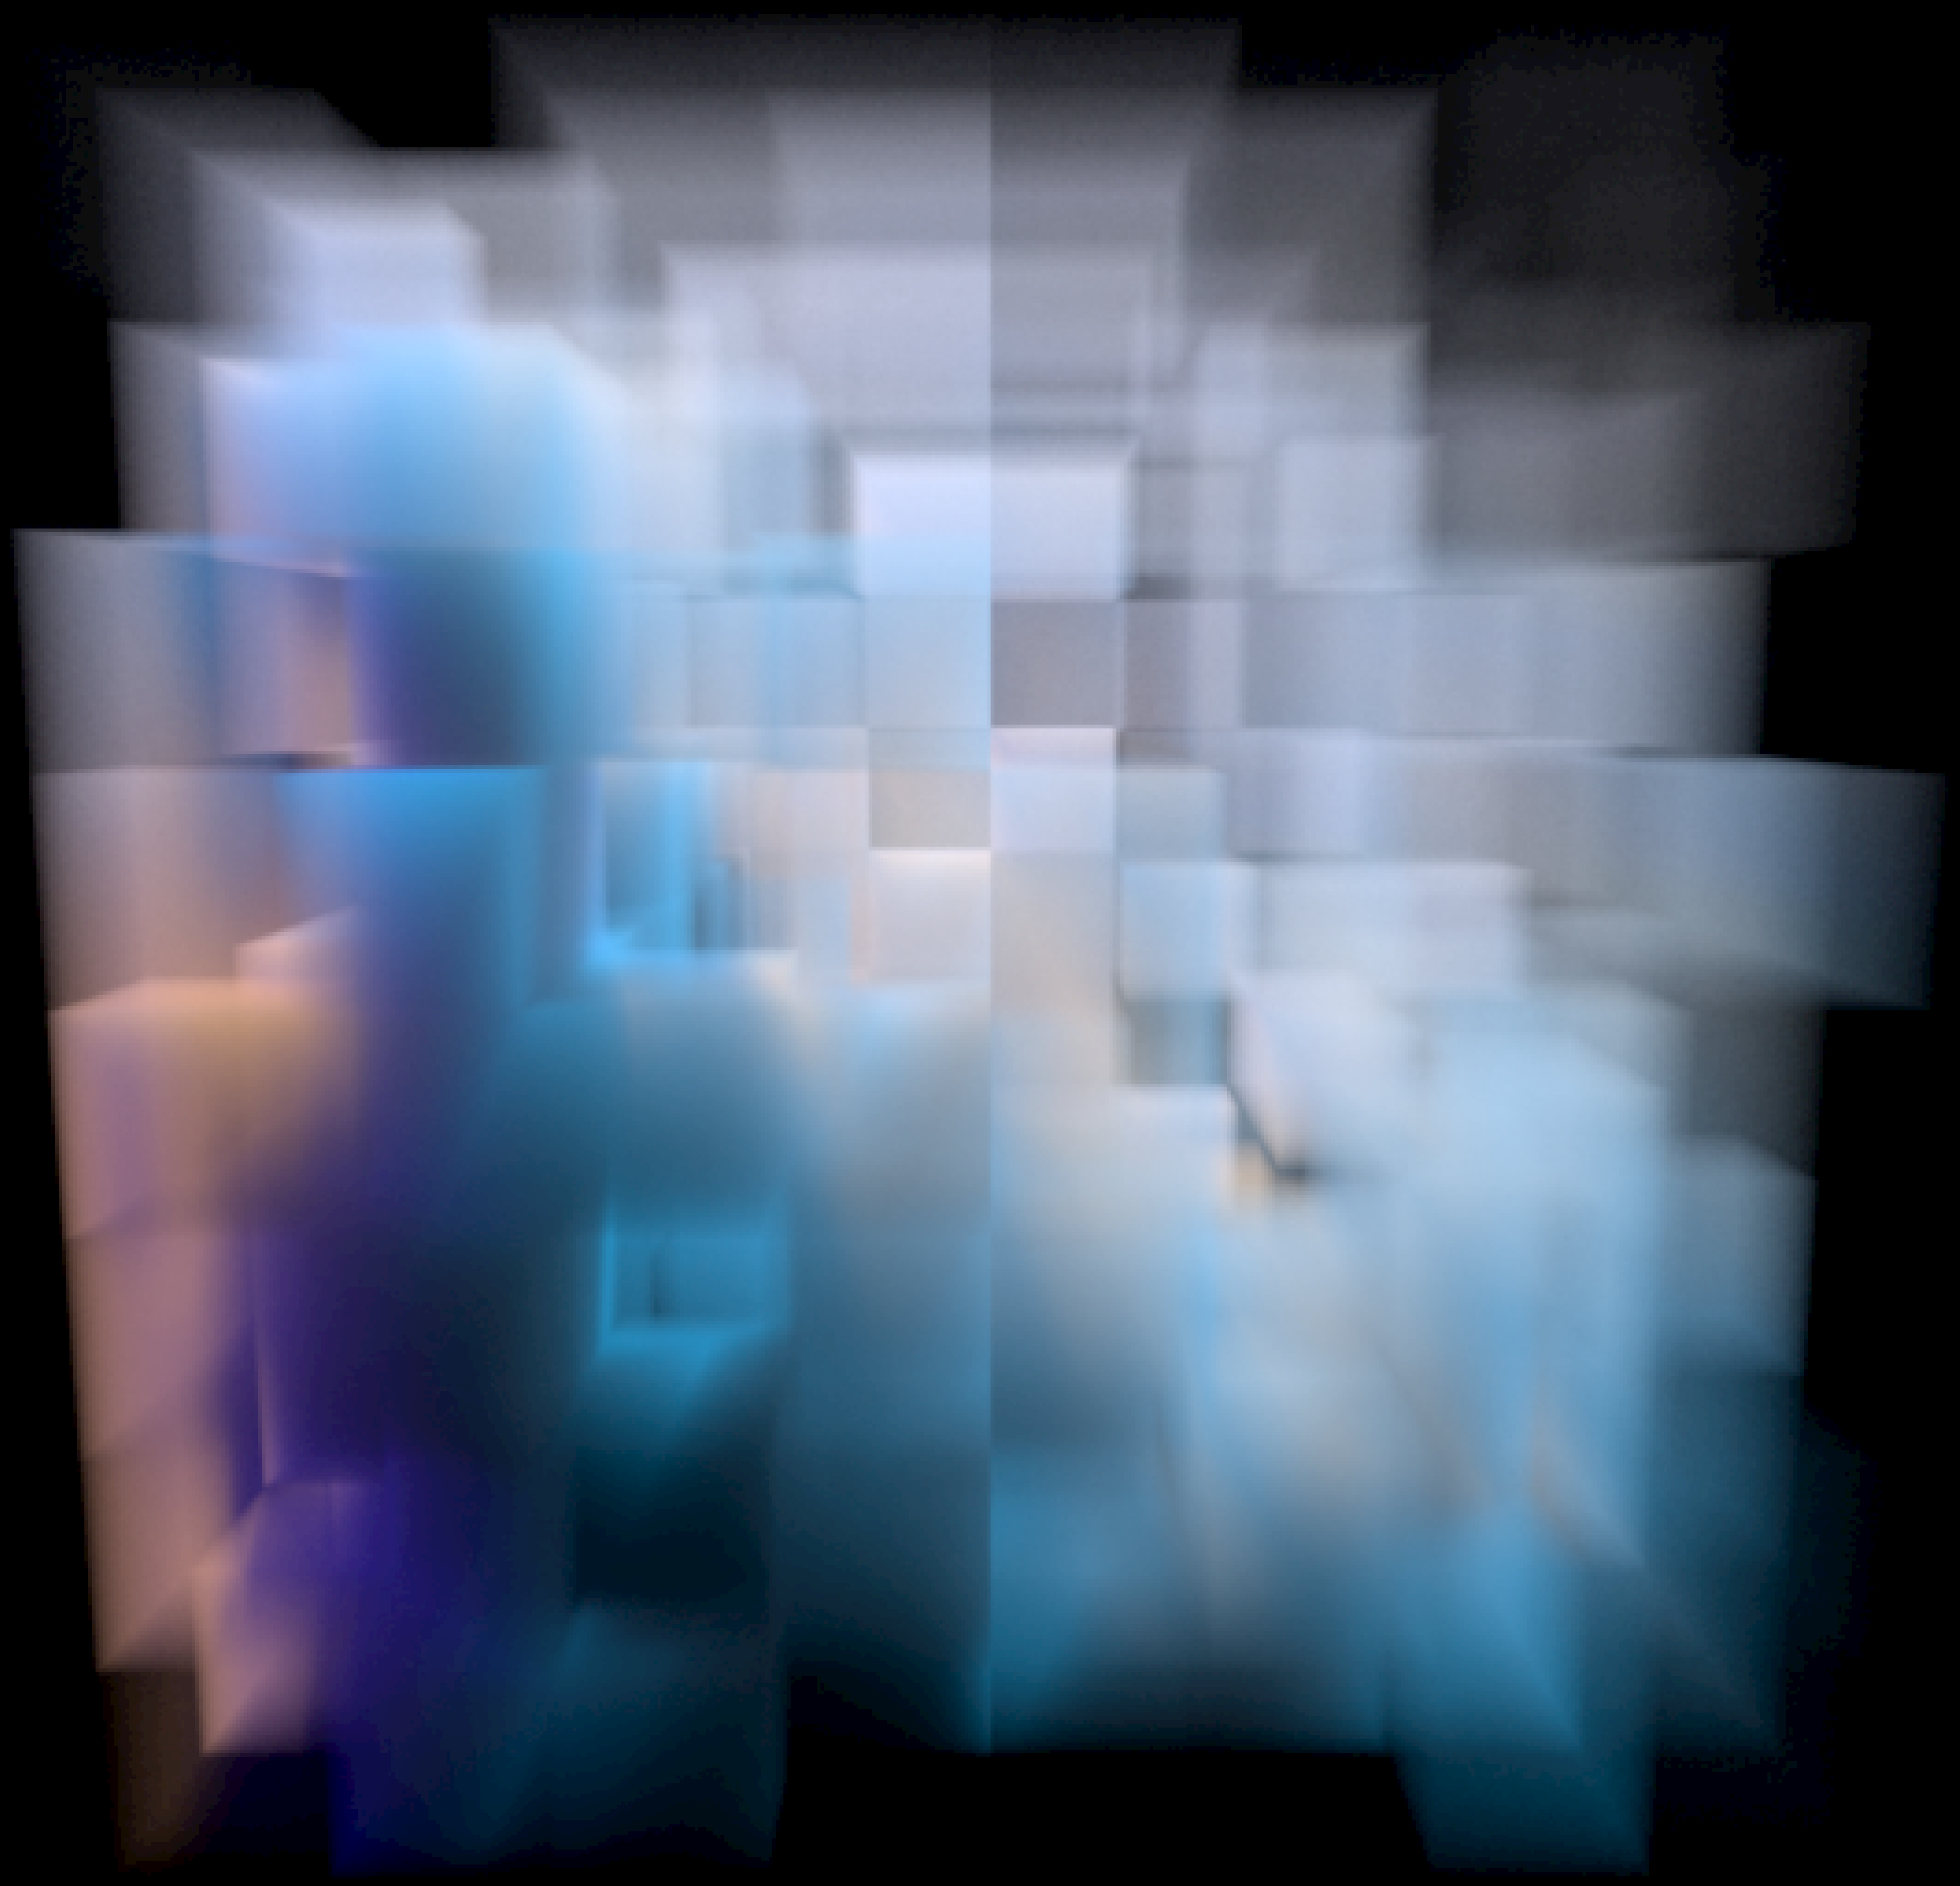

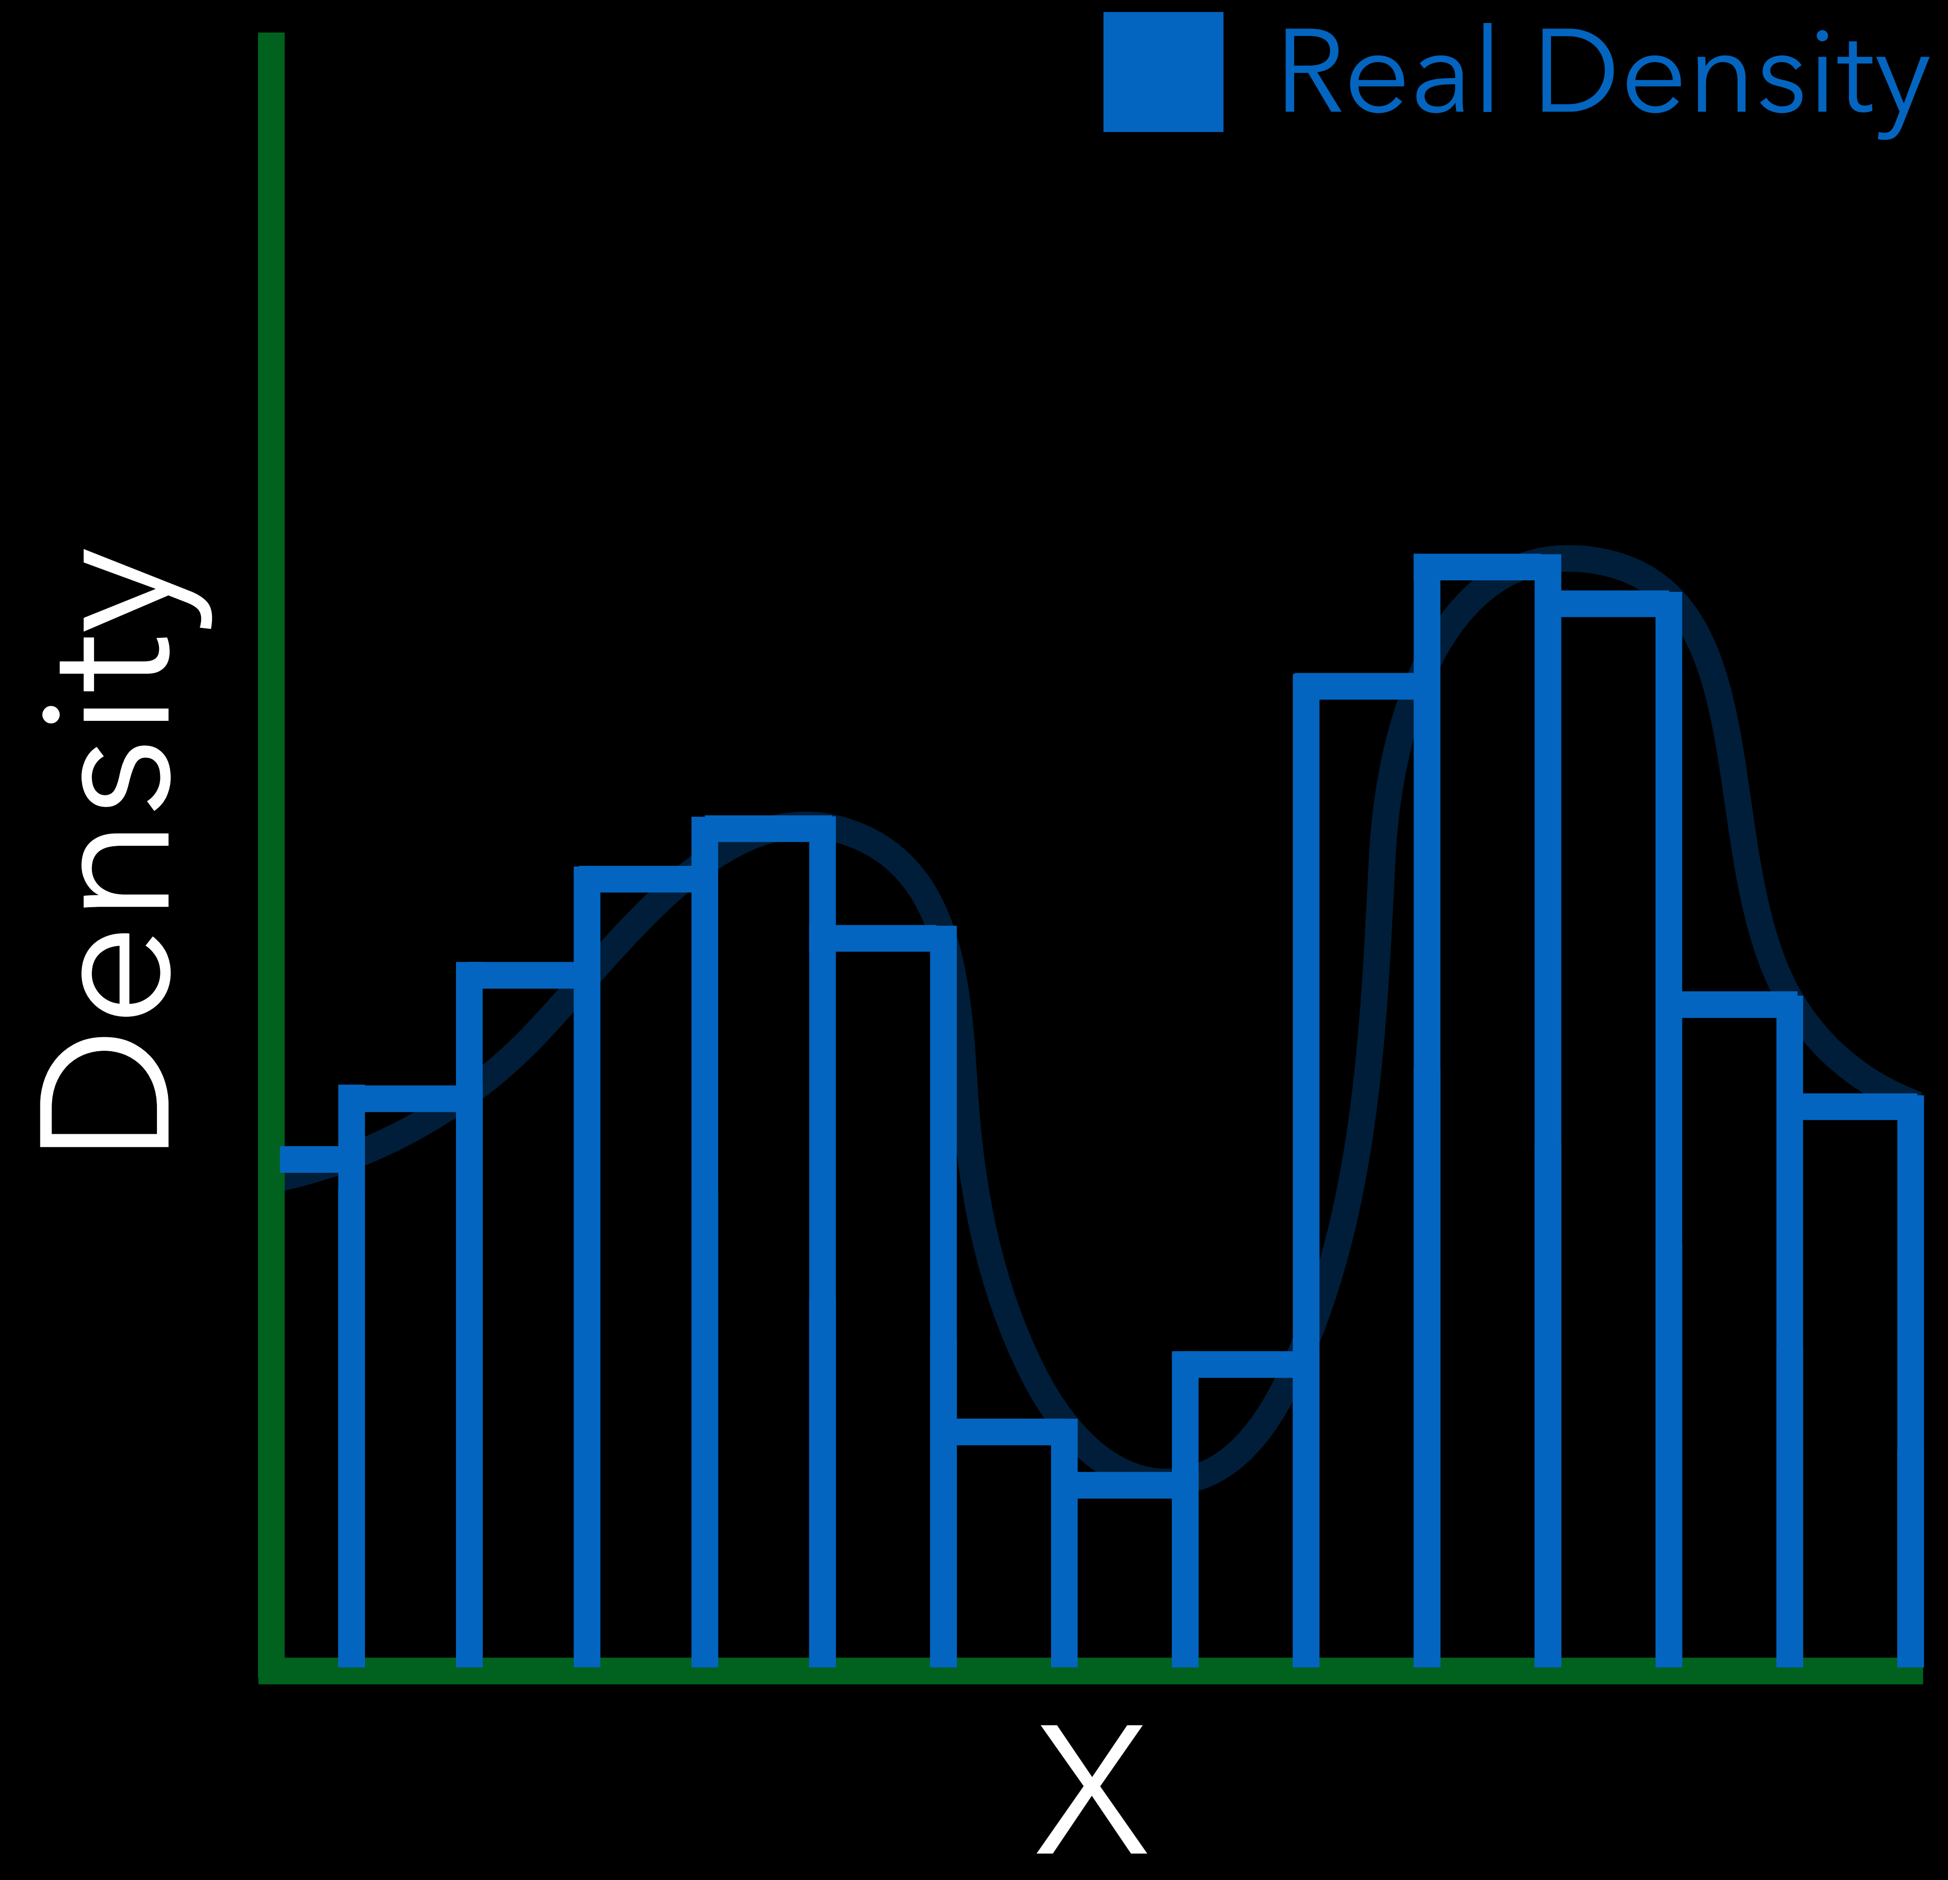
Real Density
Density
X

👍 robust

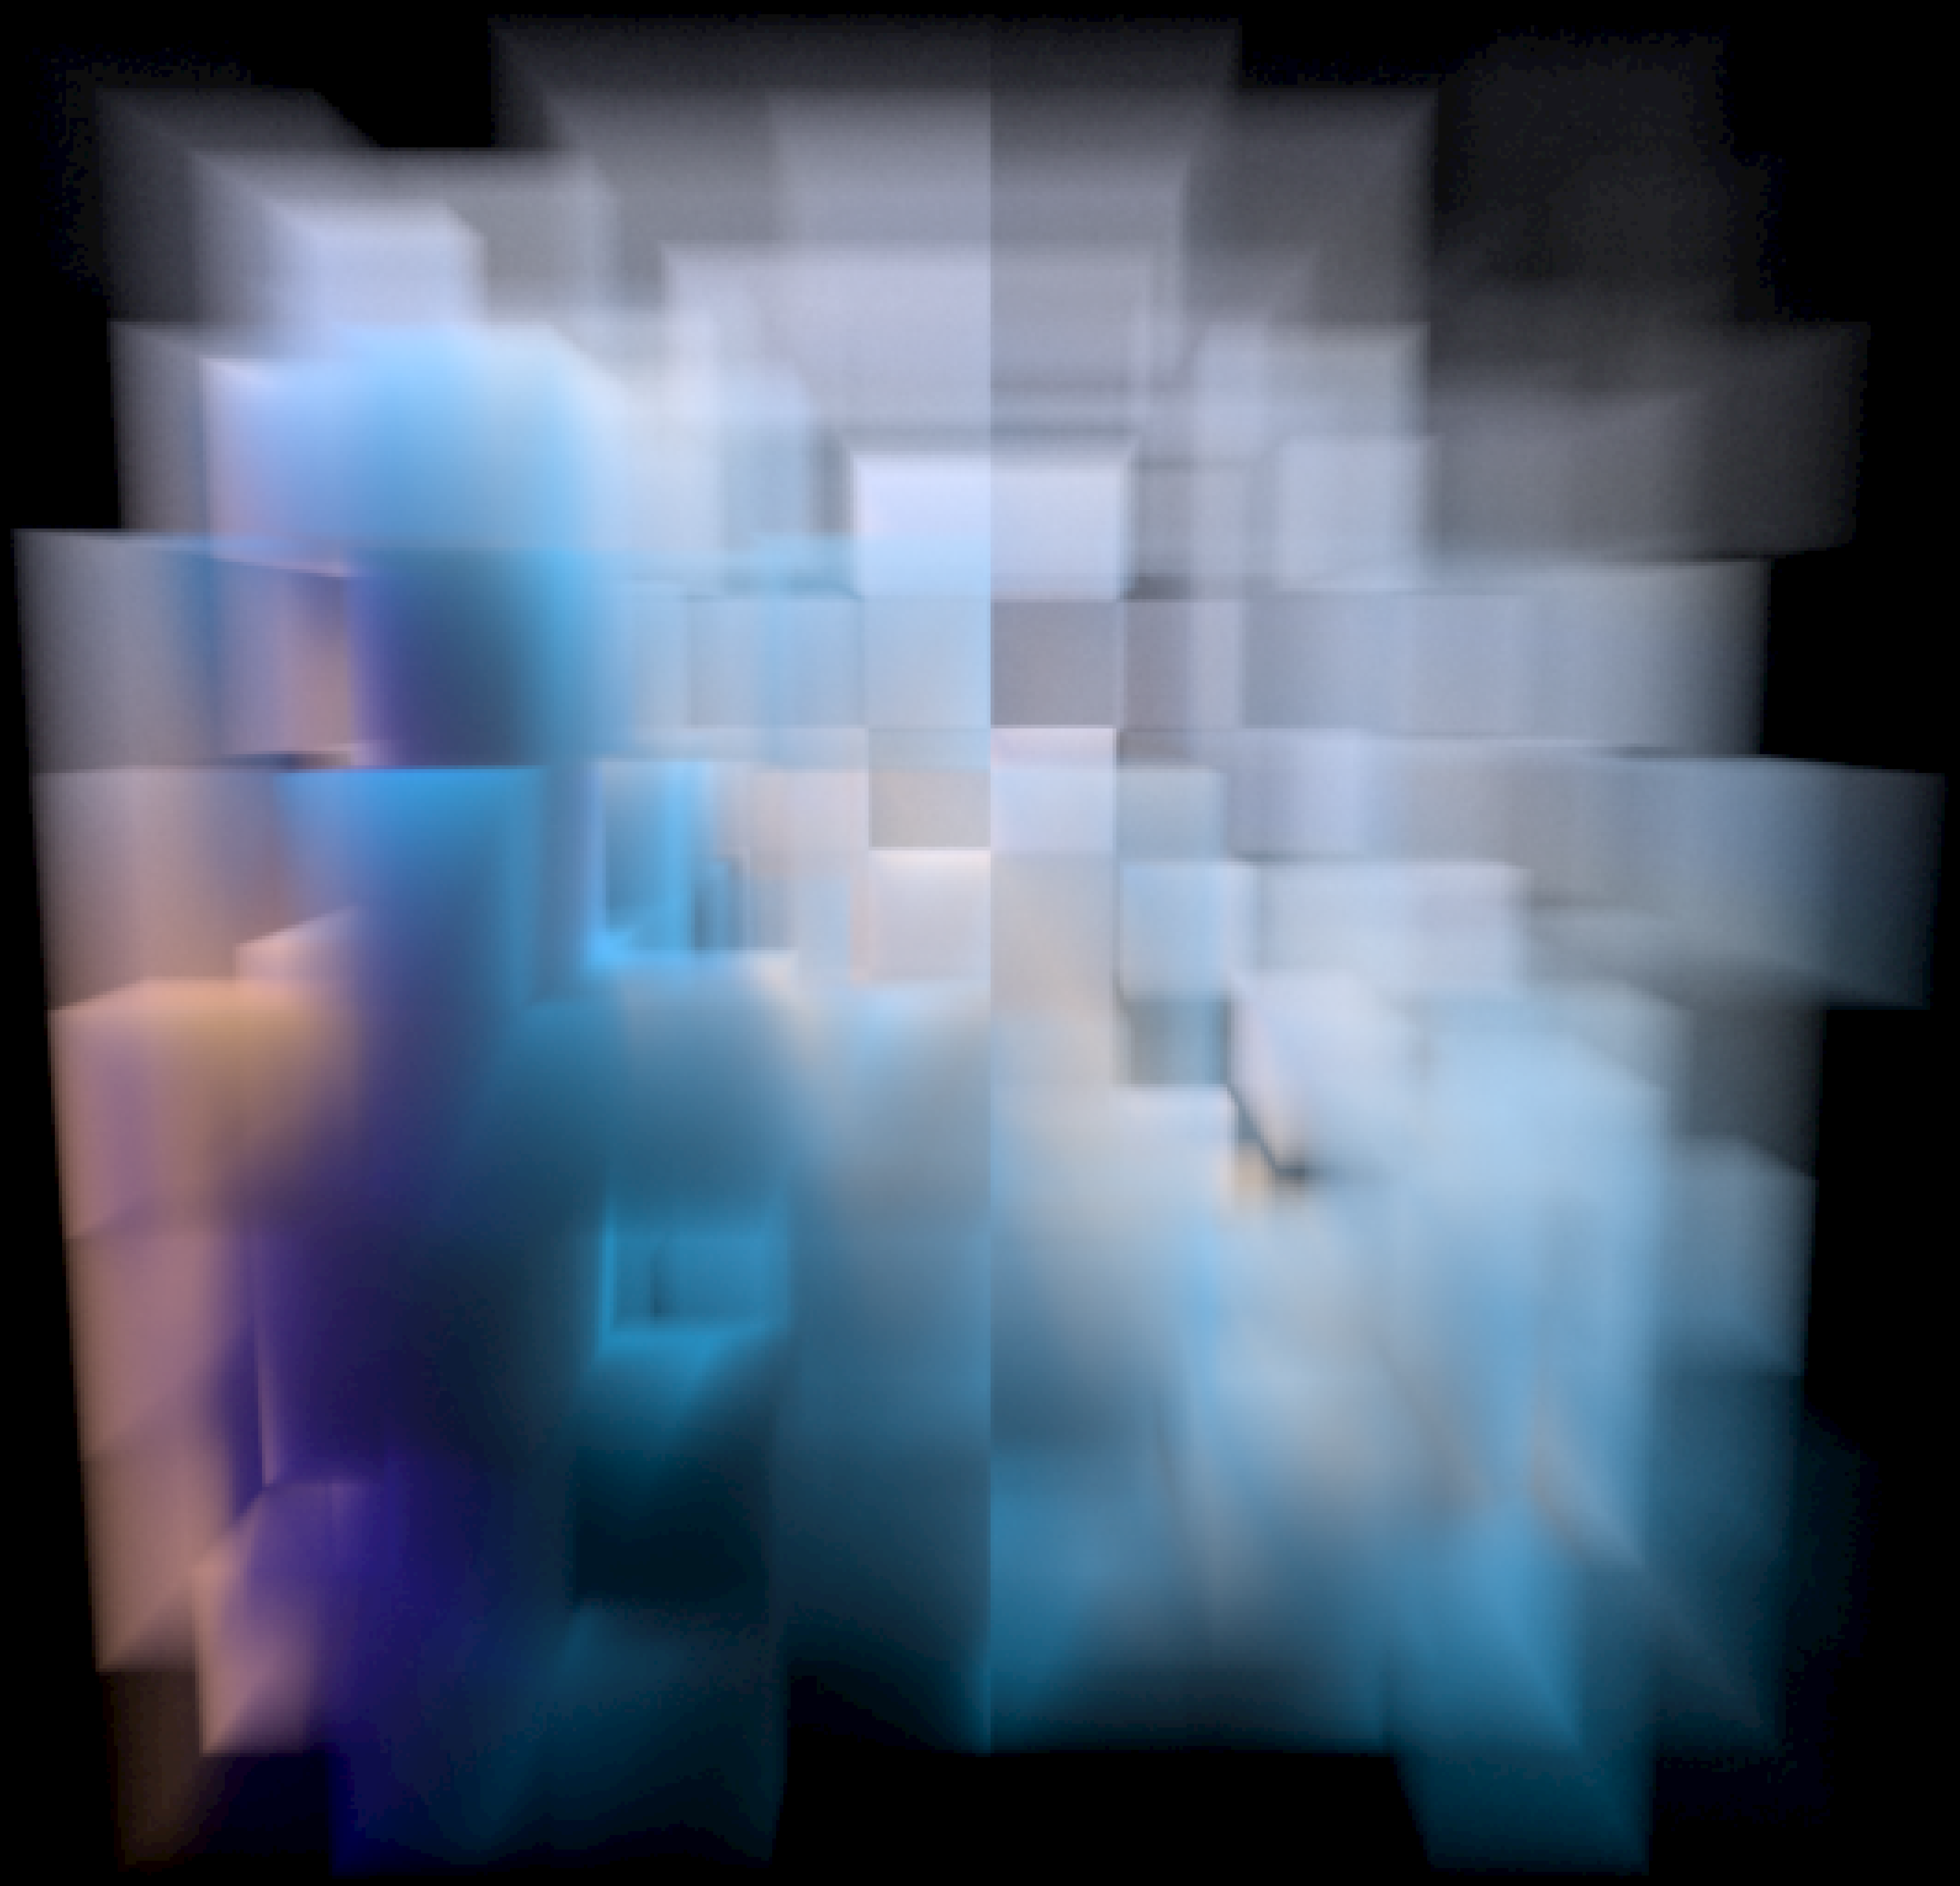

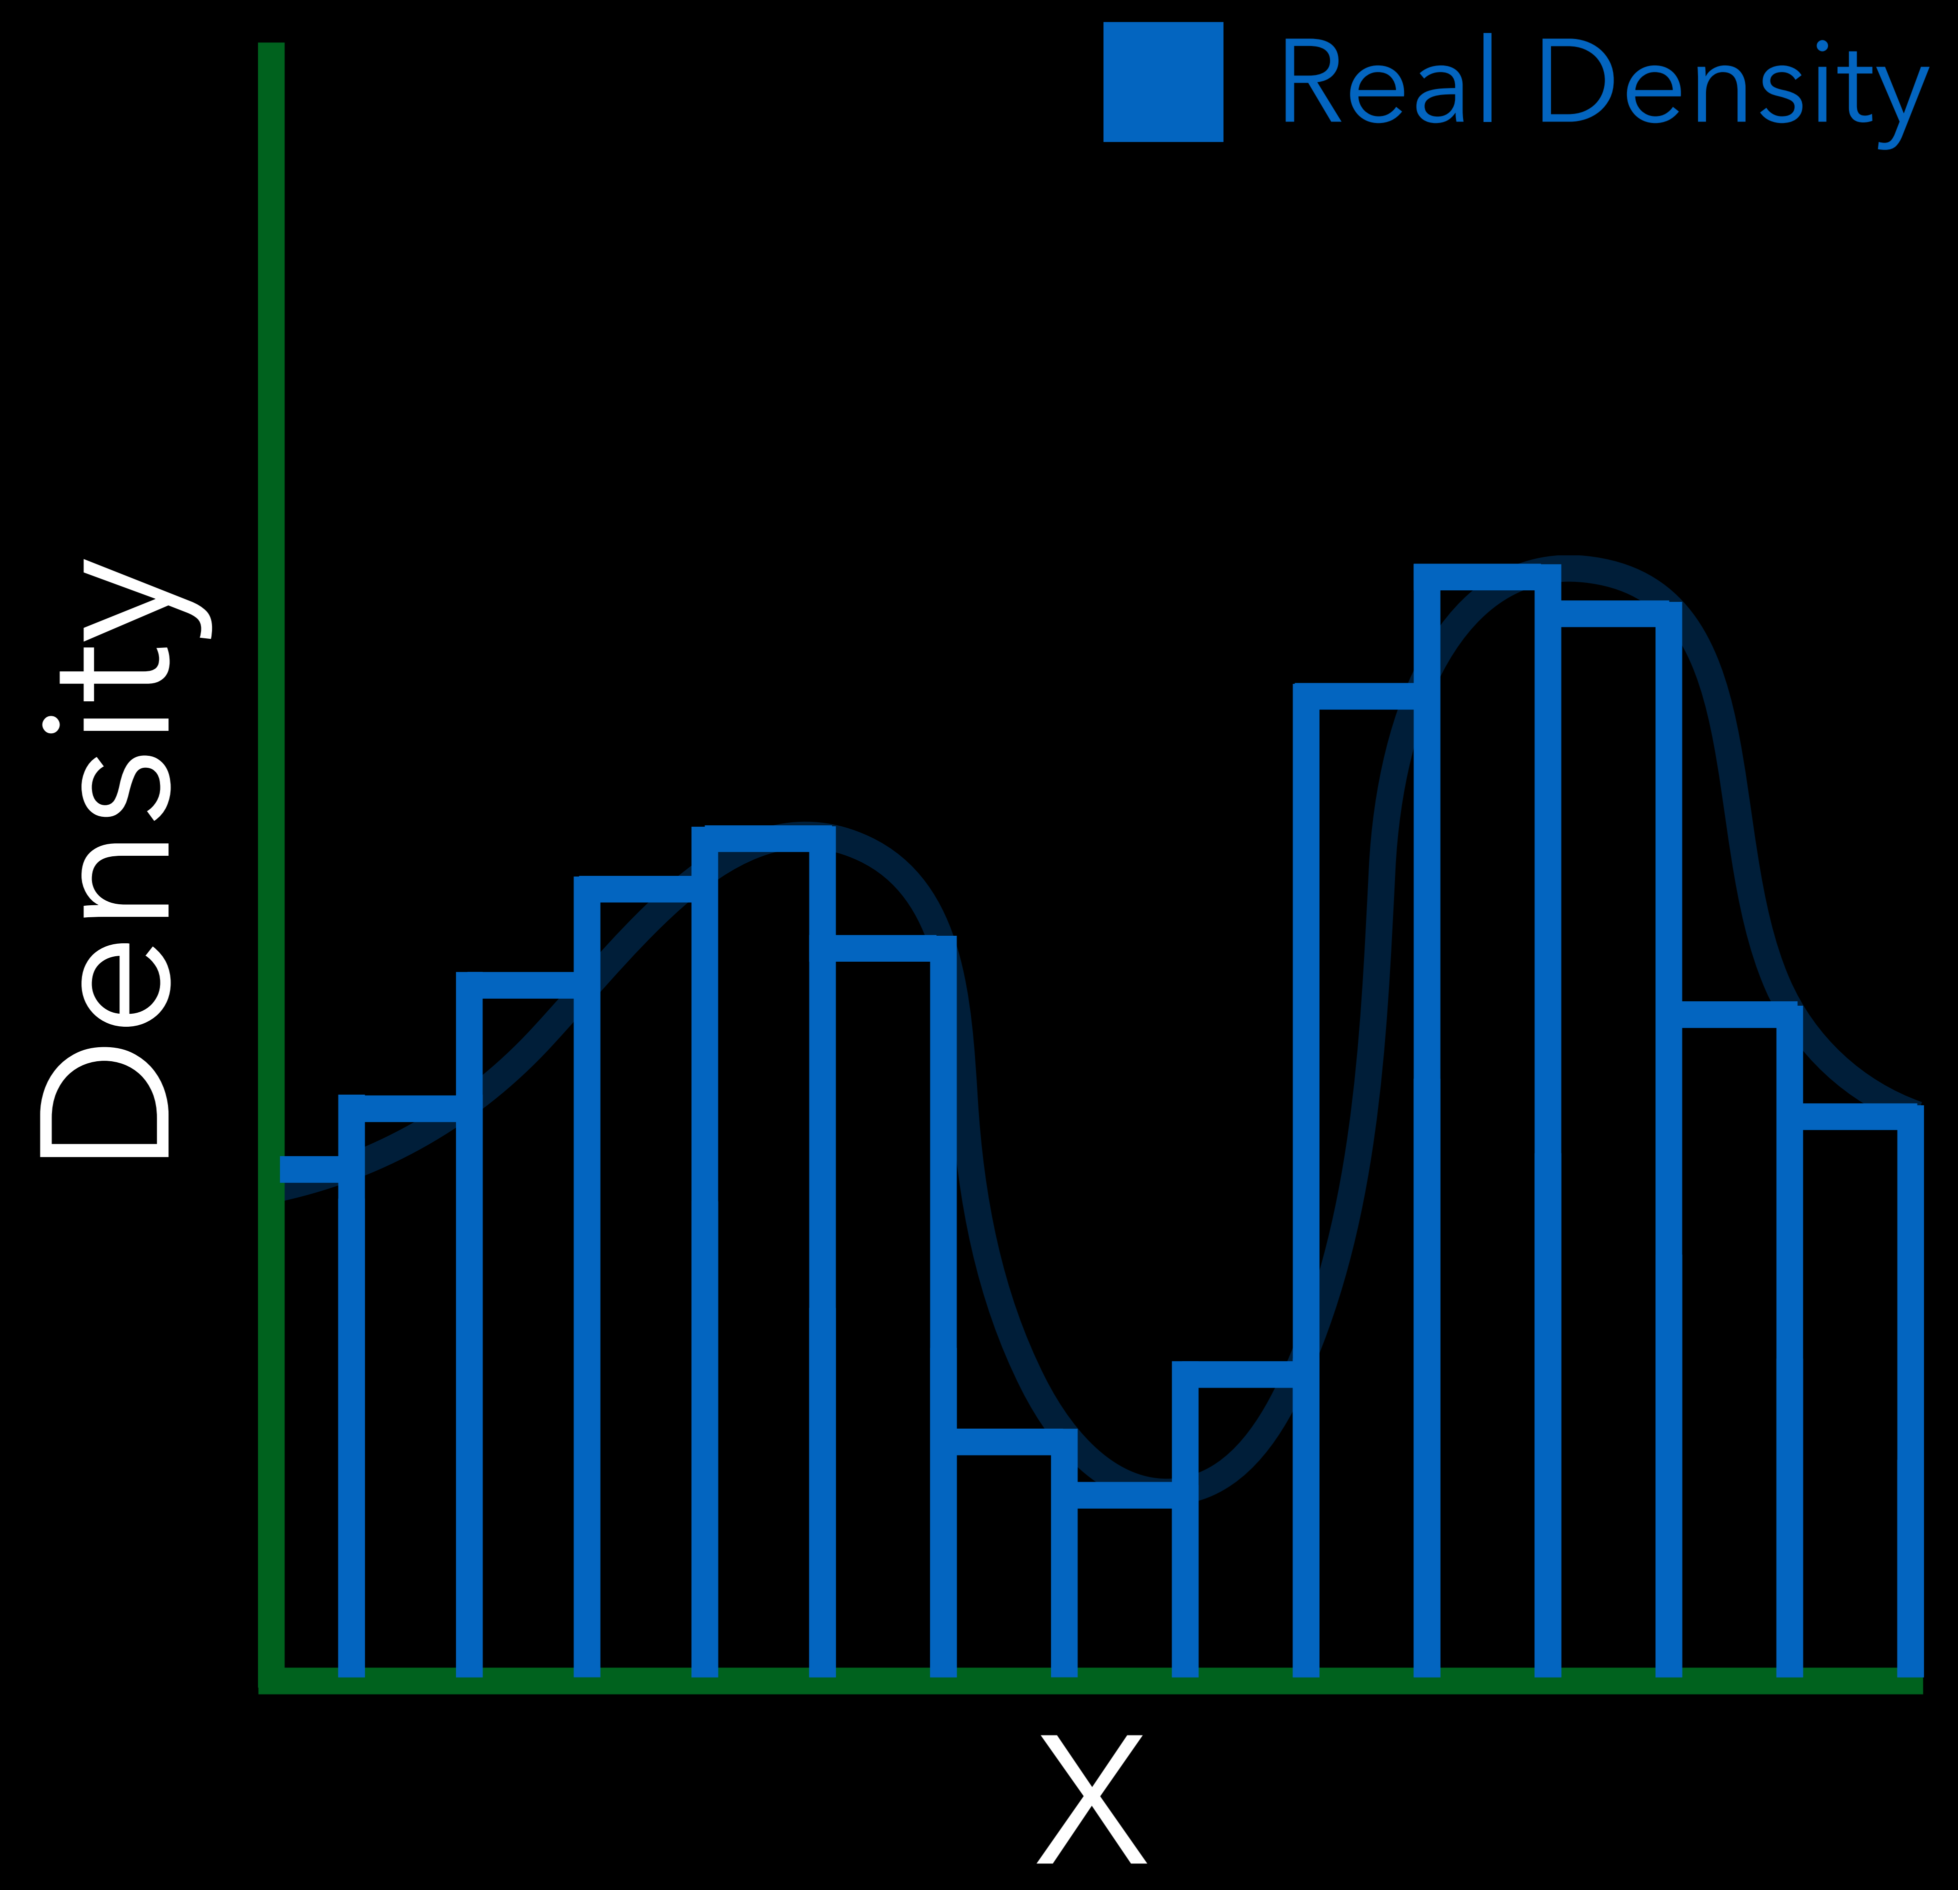

👍 tight 🐜 👍 robust

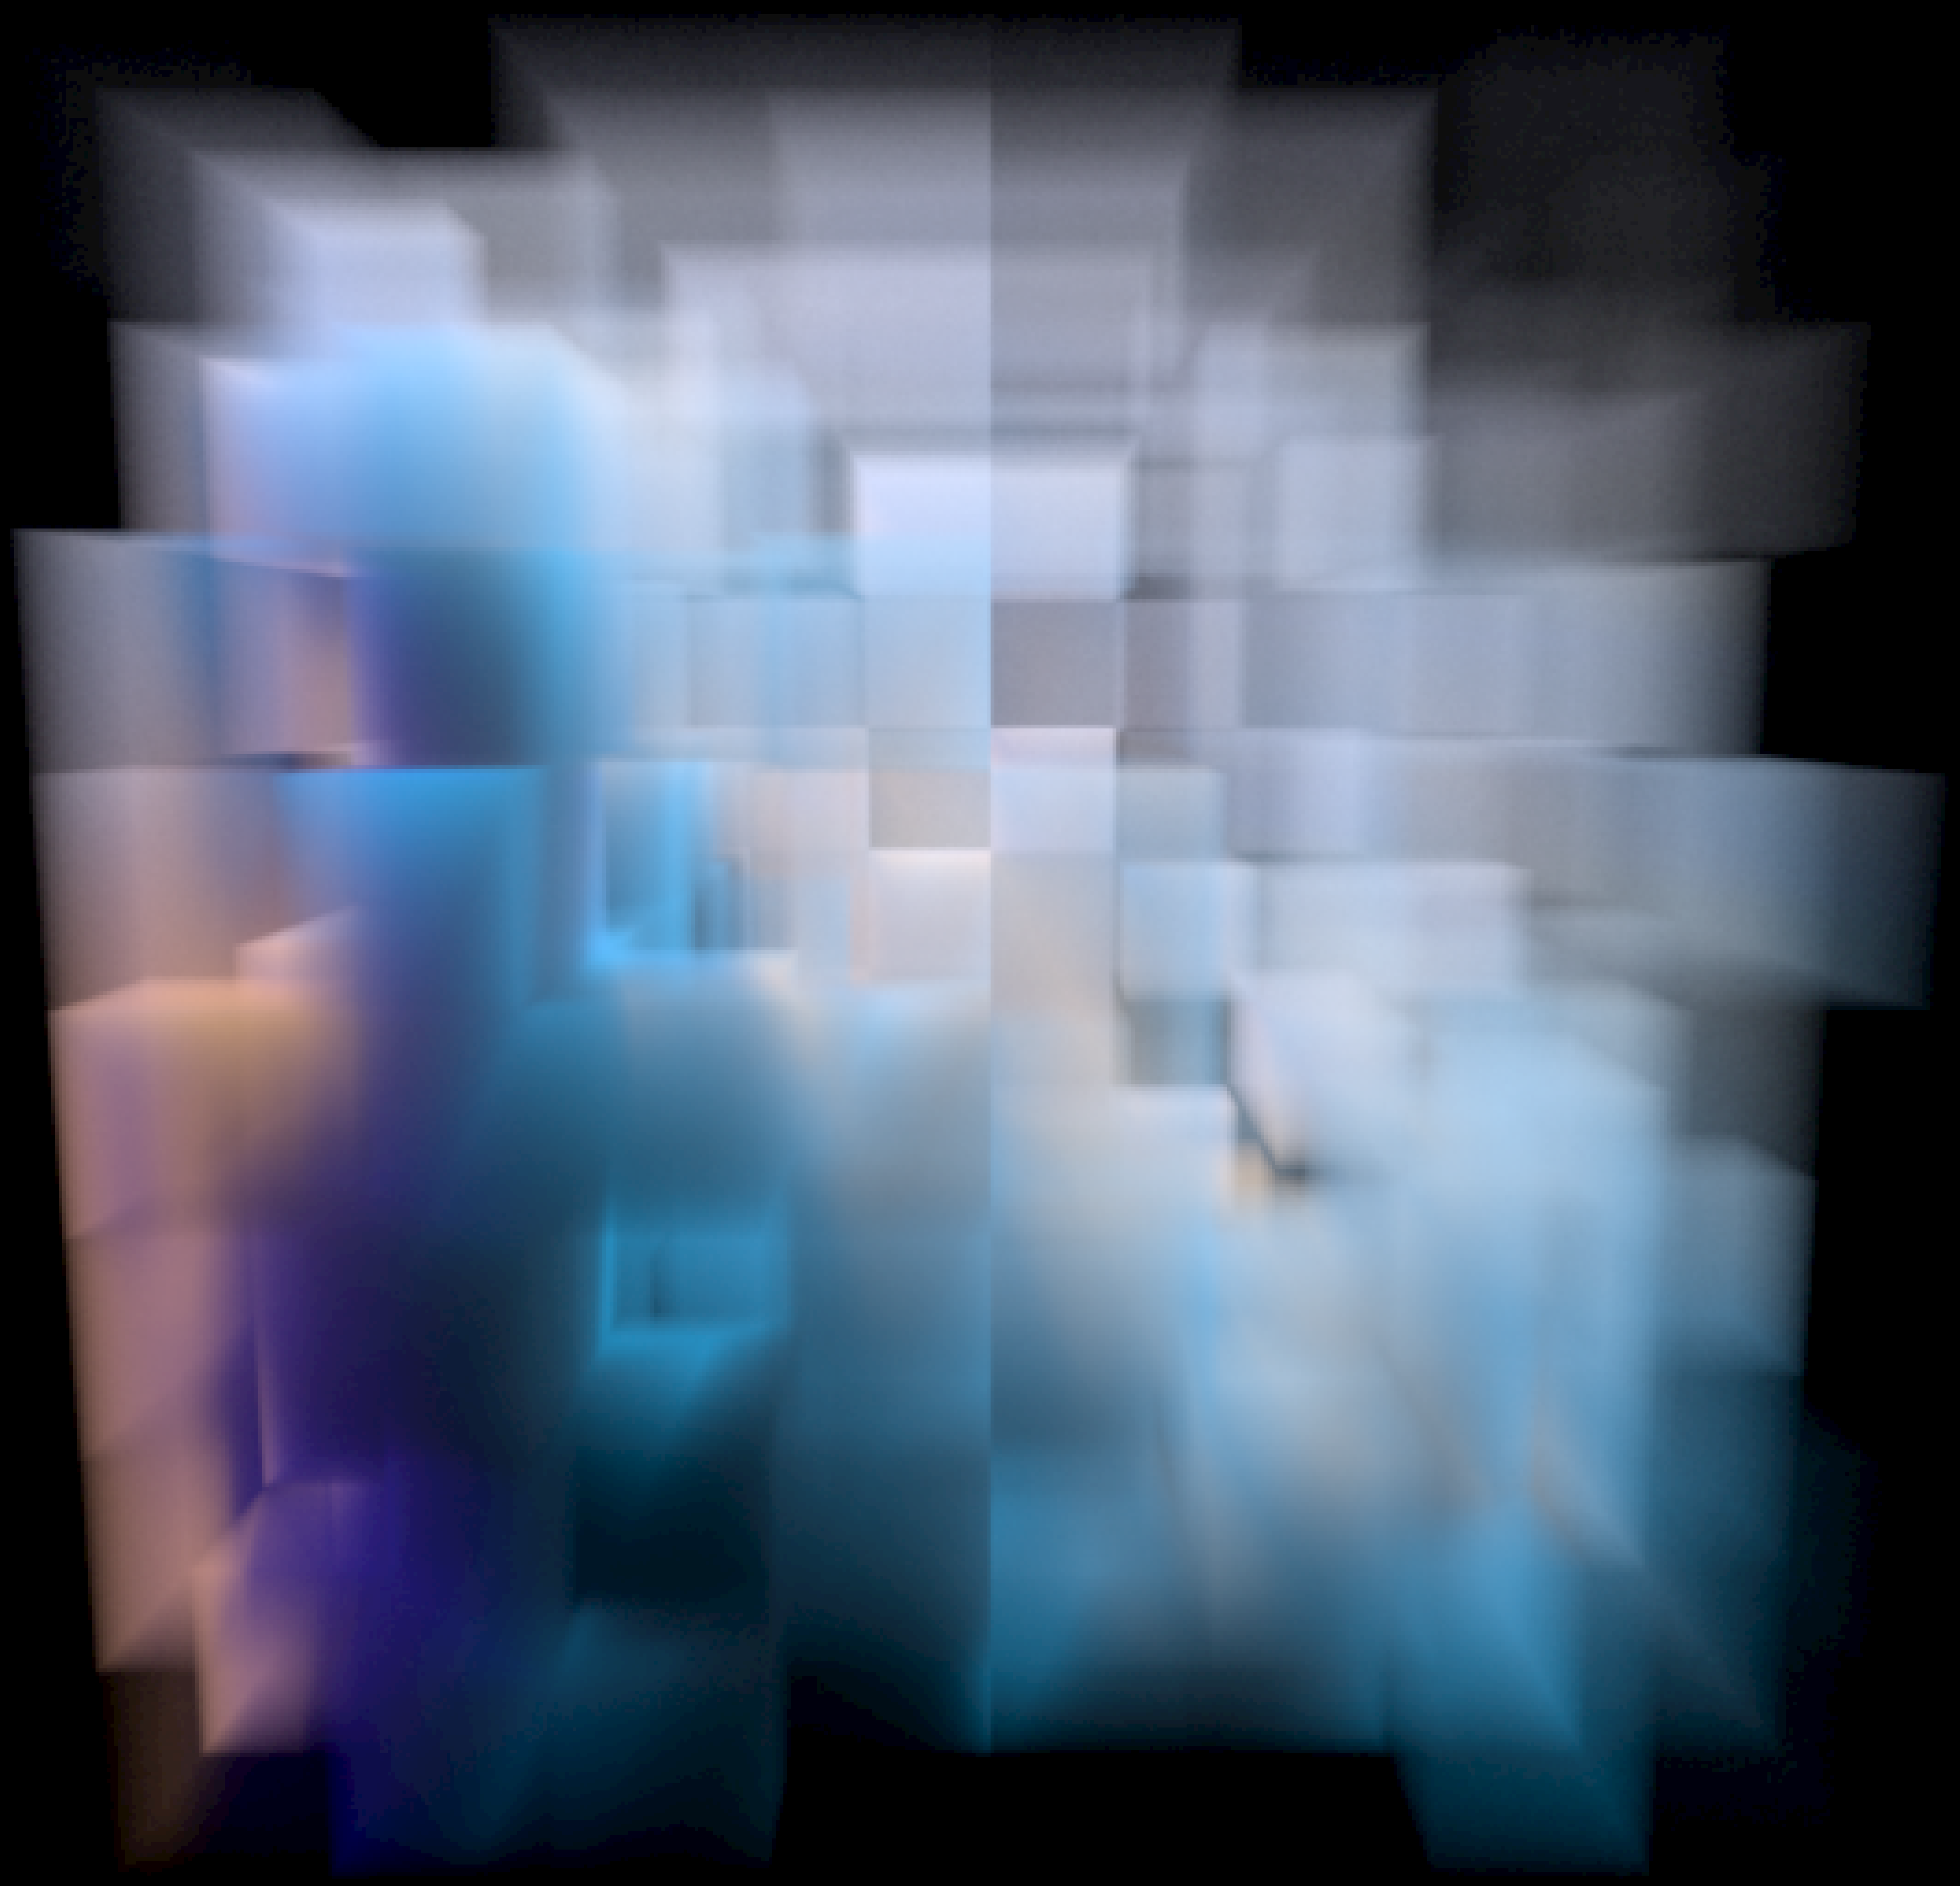

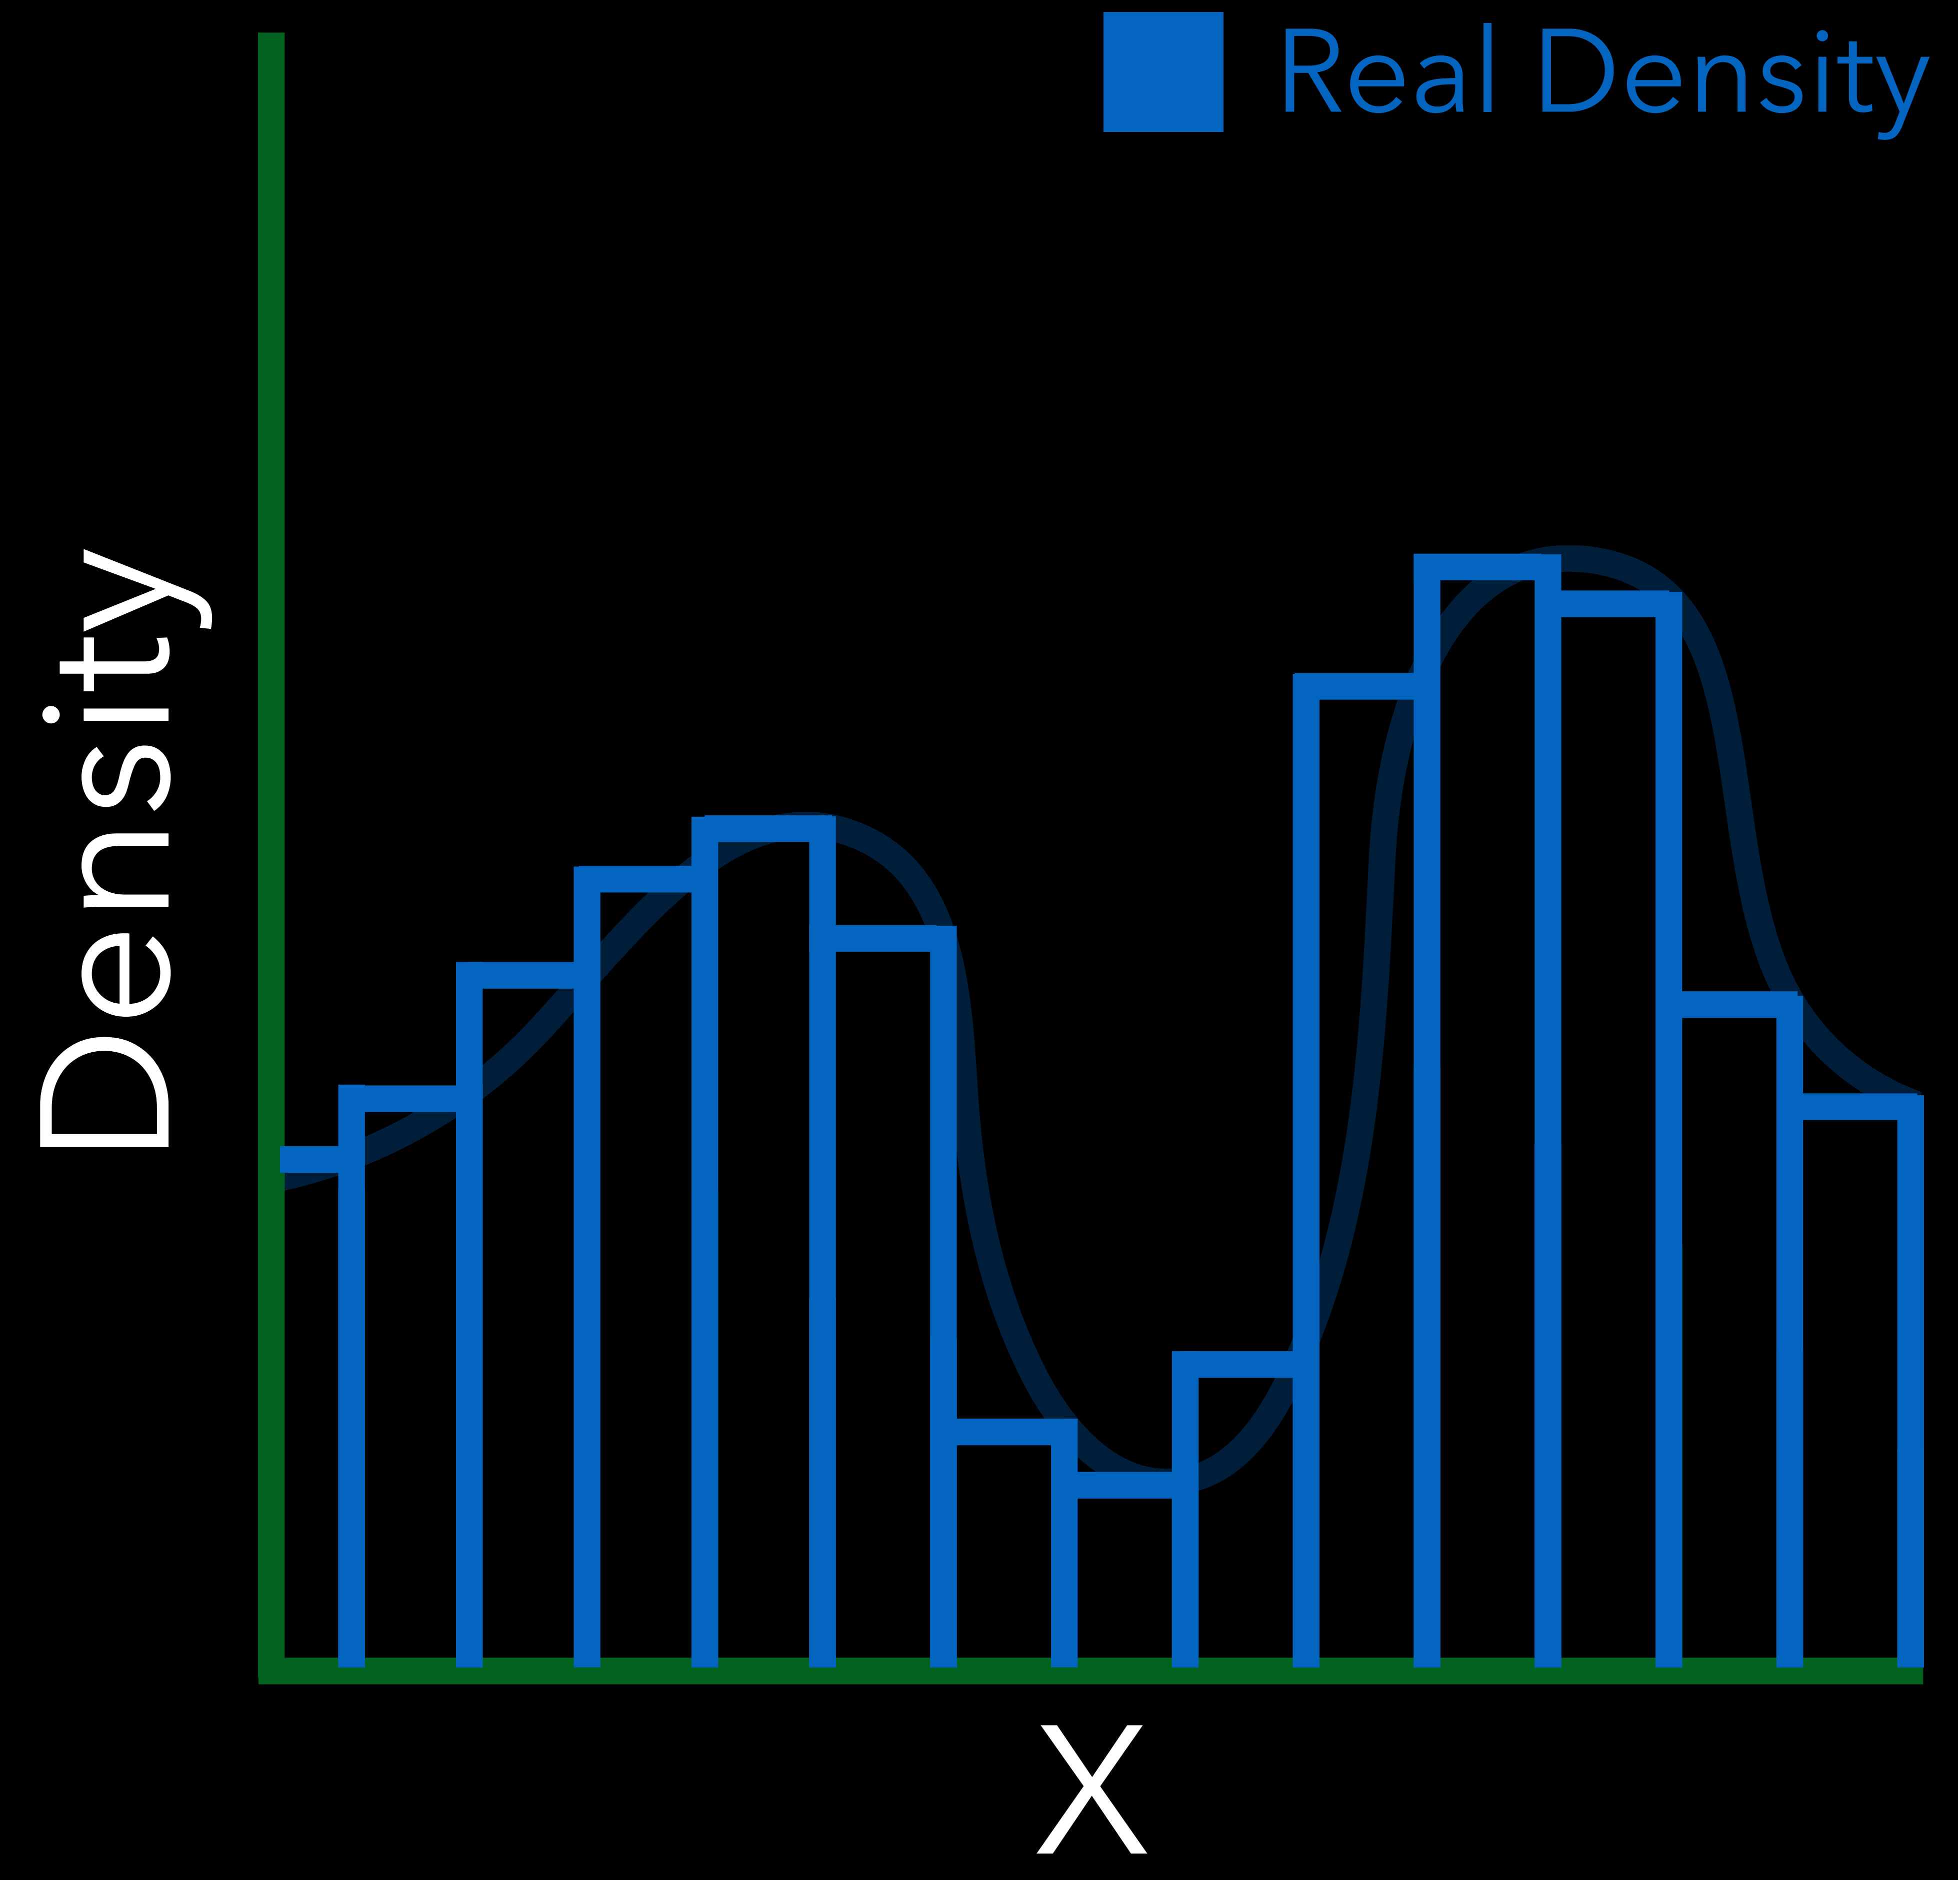

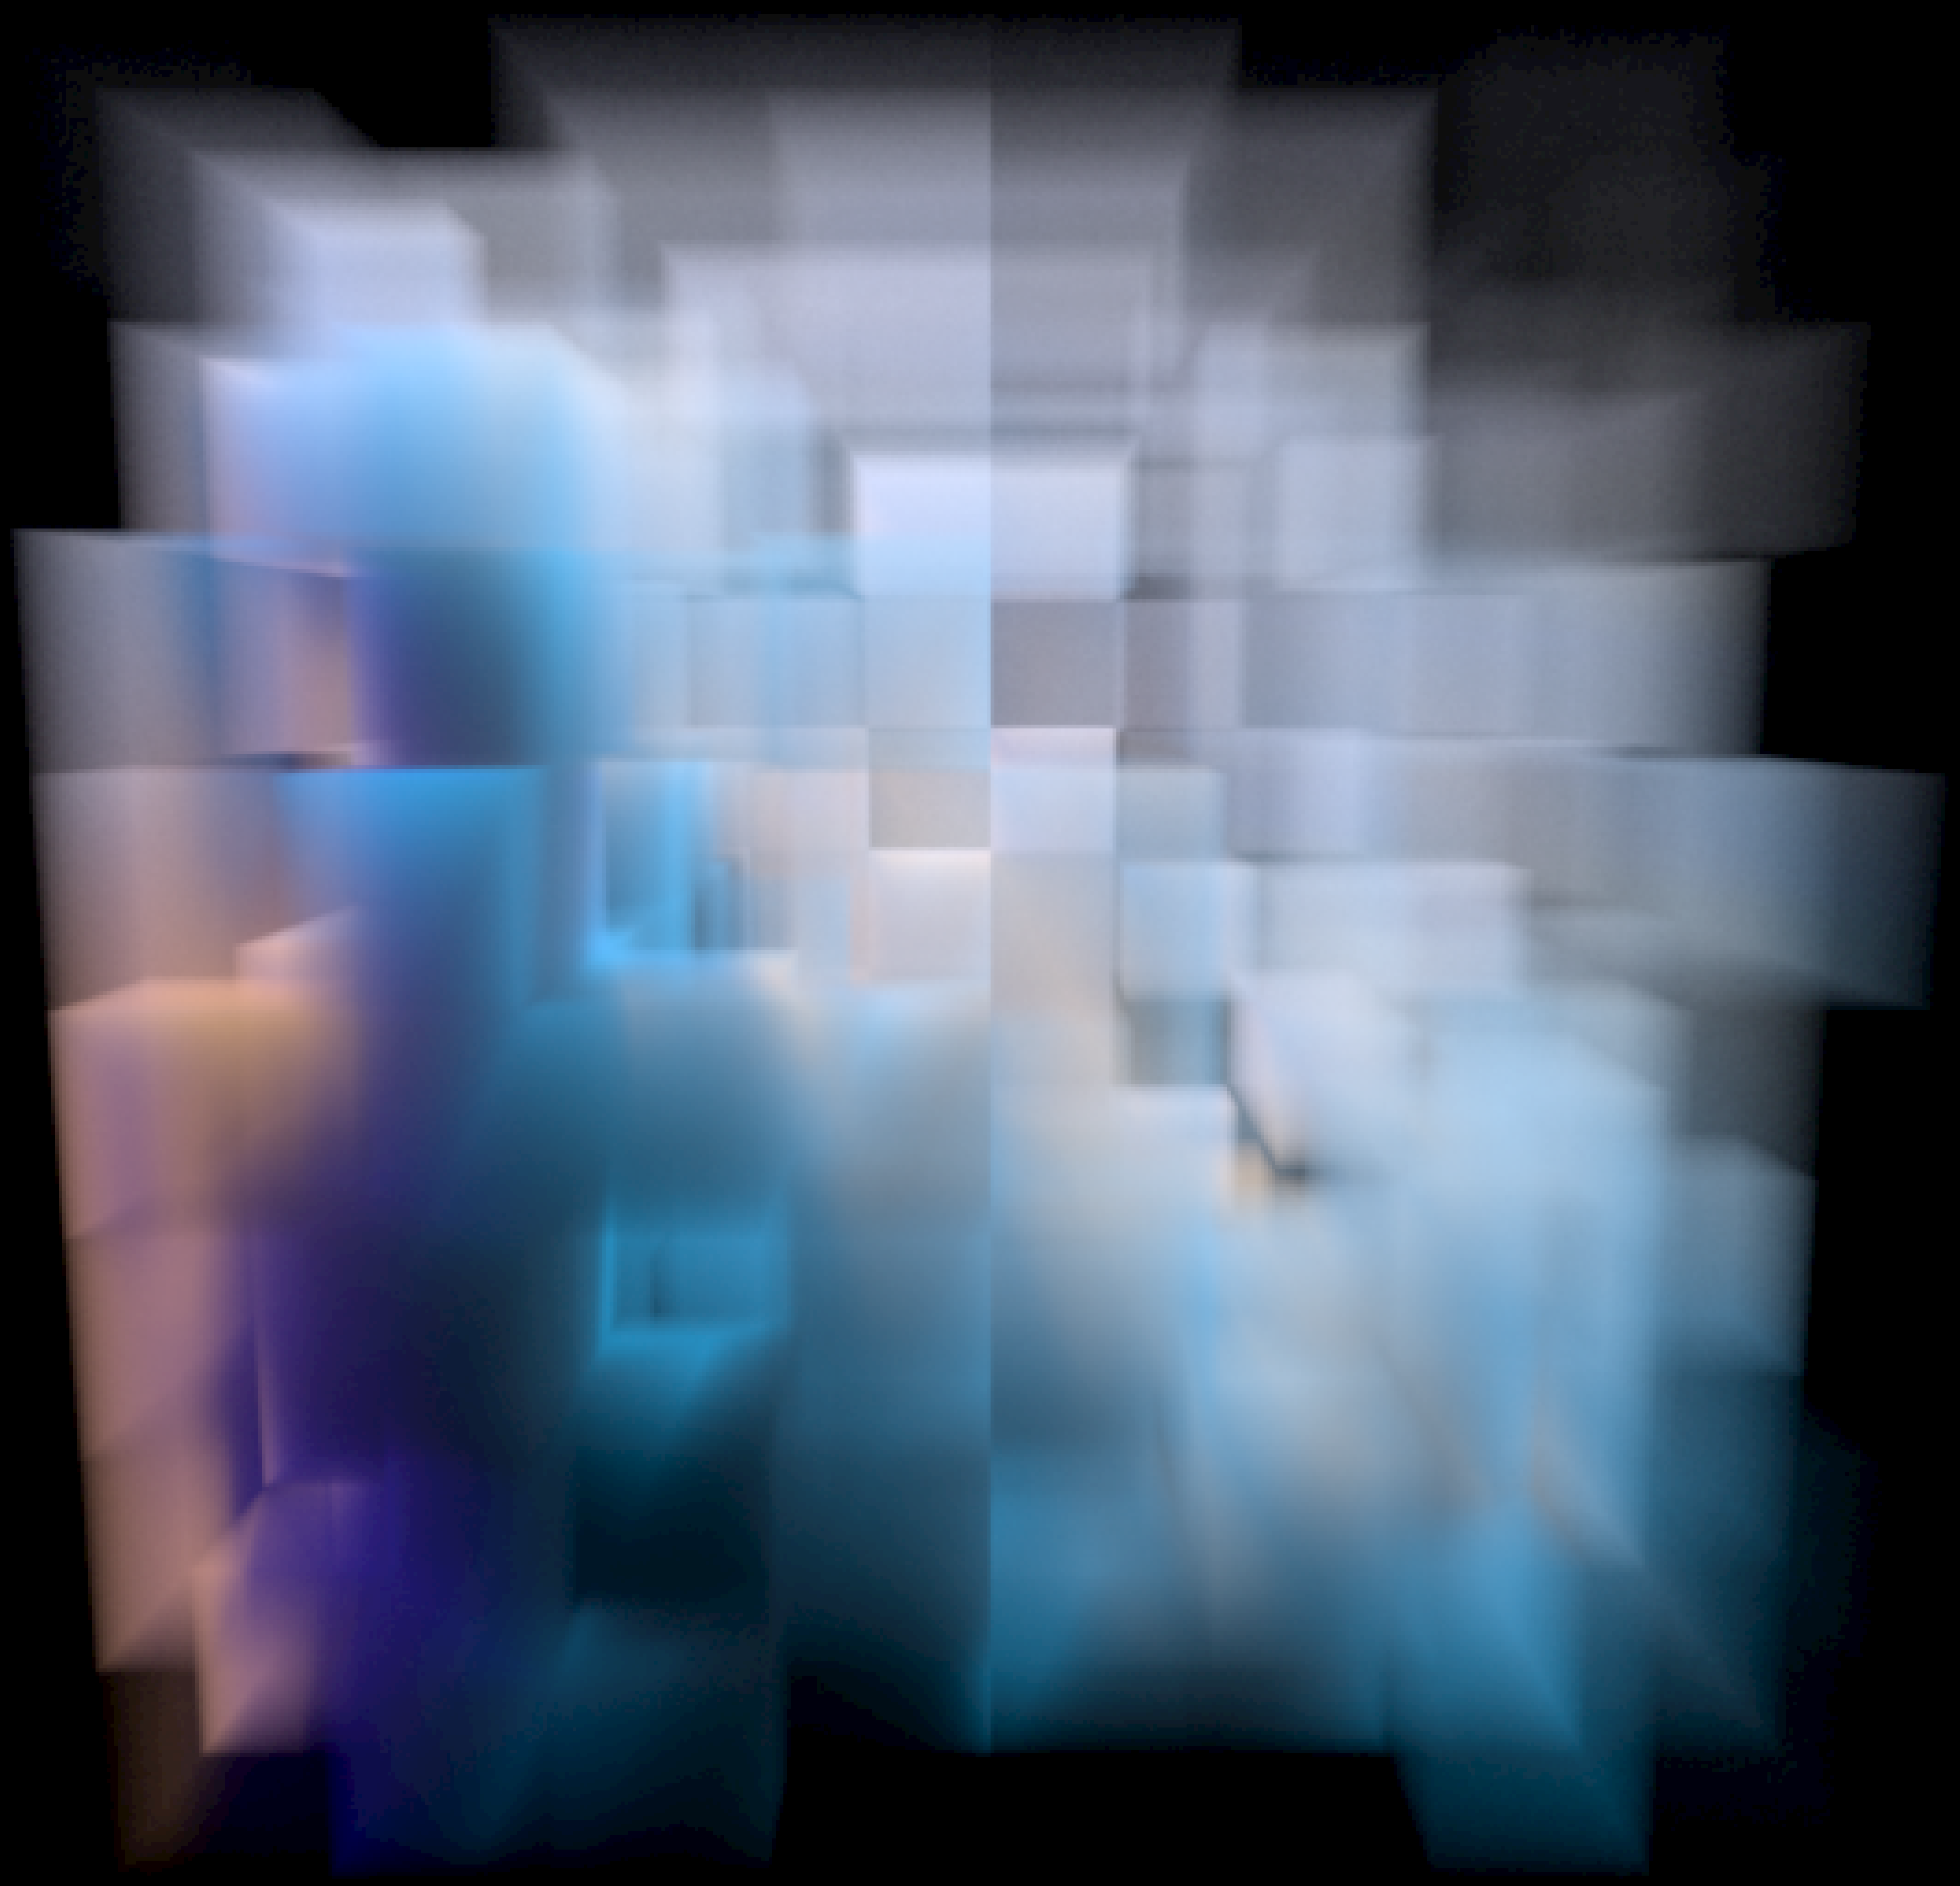

👍 tight 🐜 👍 robust

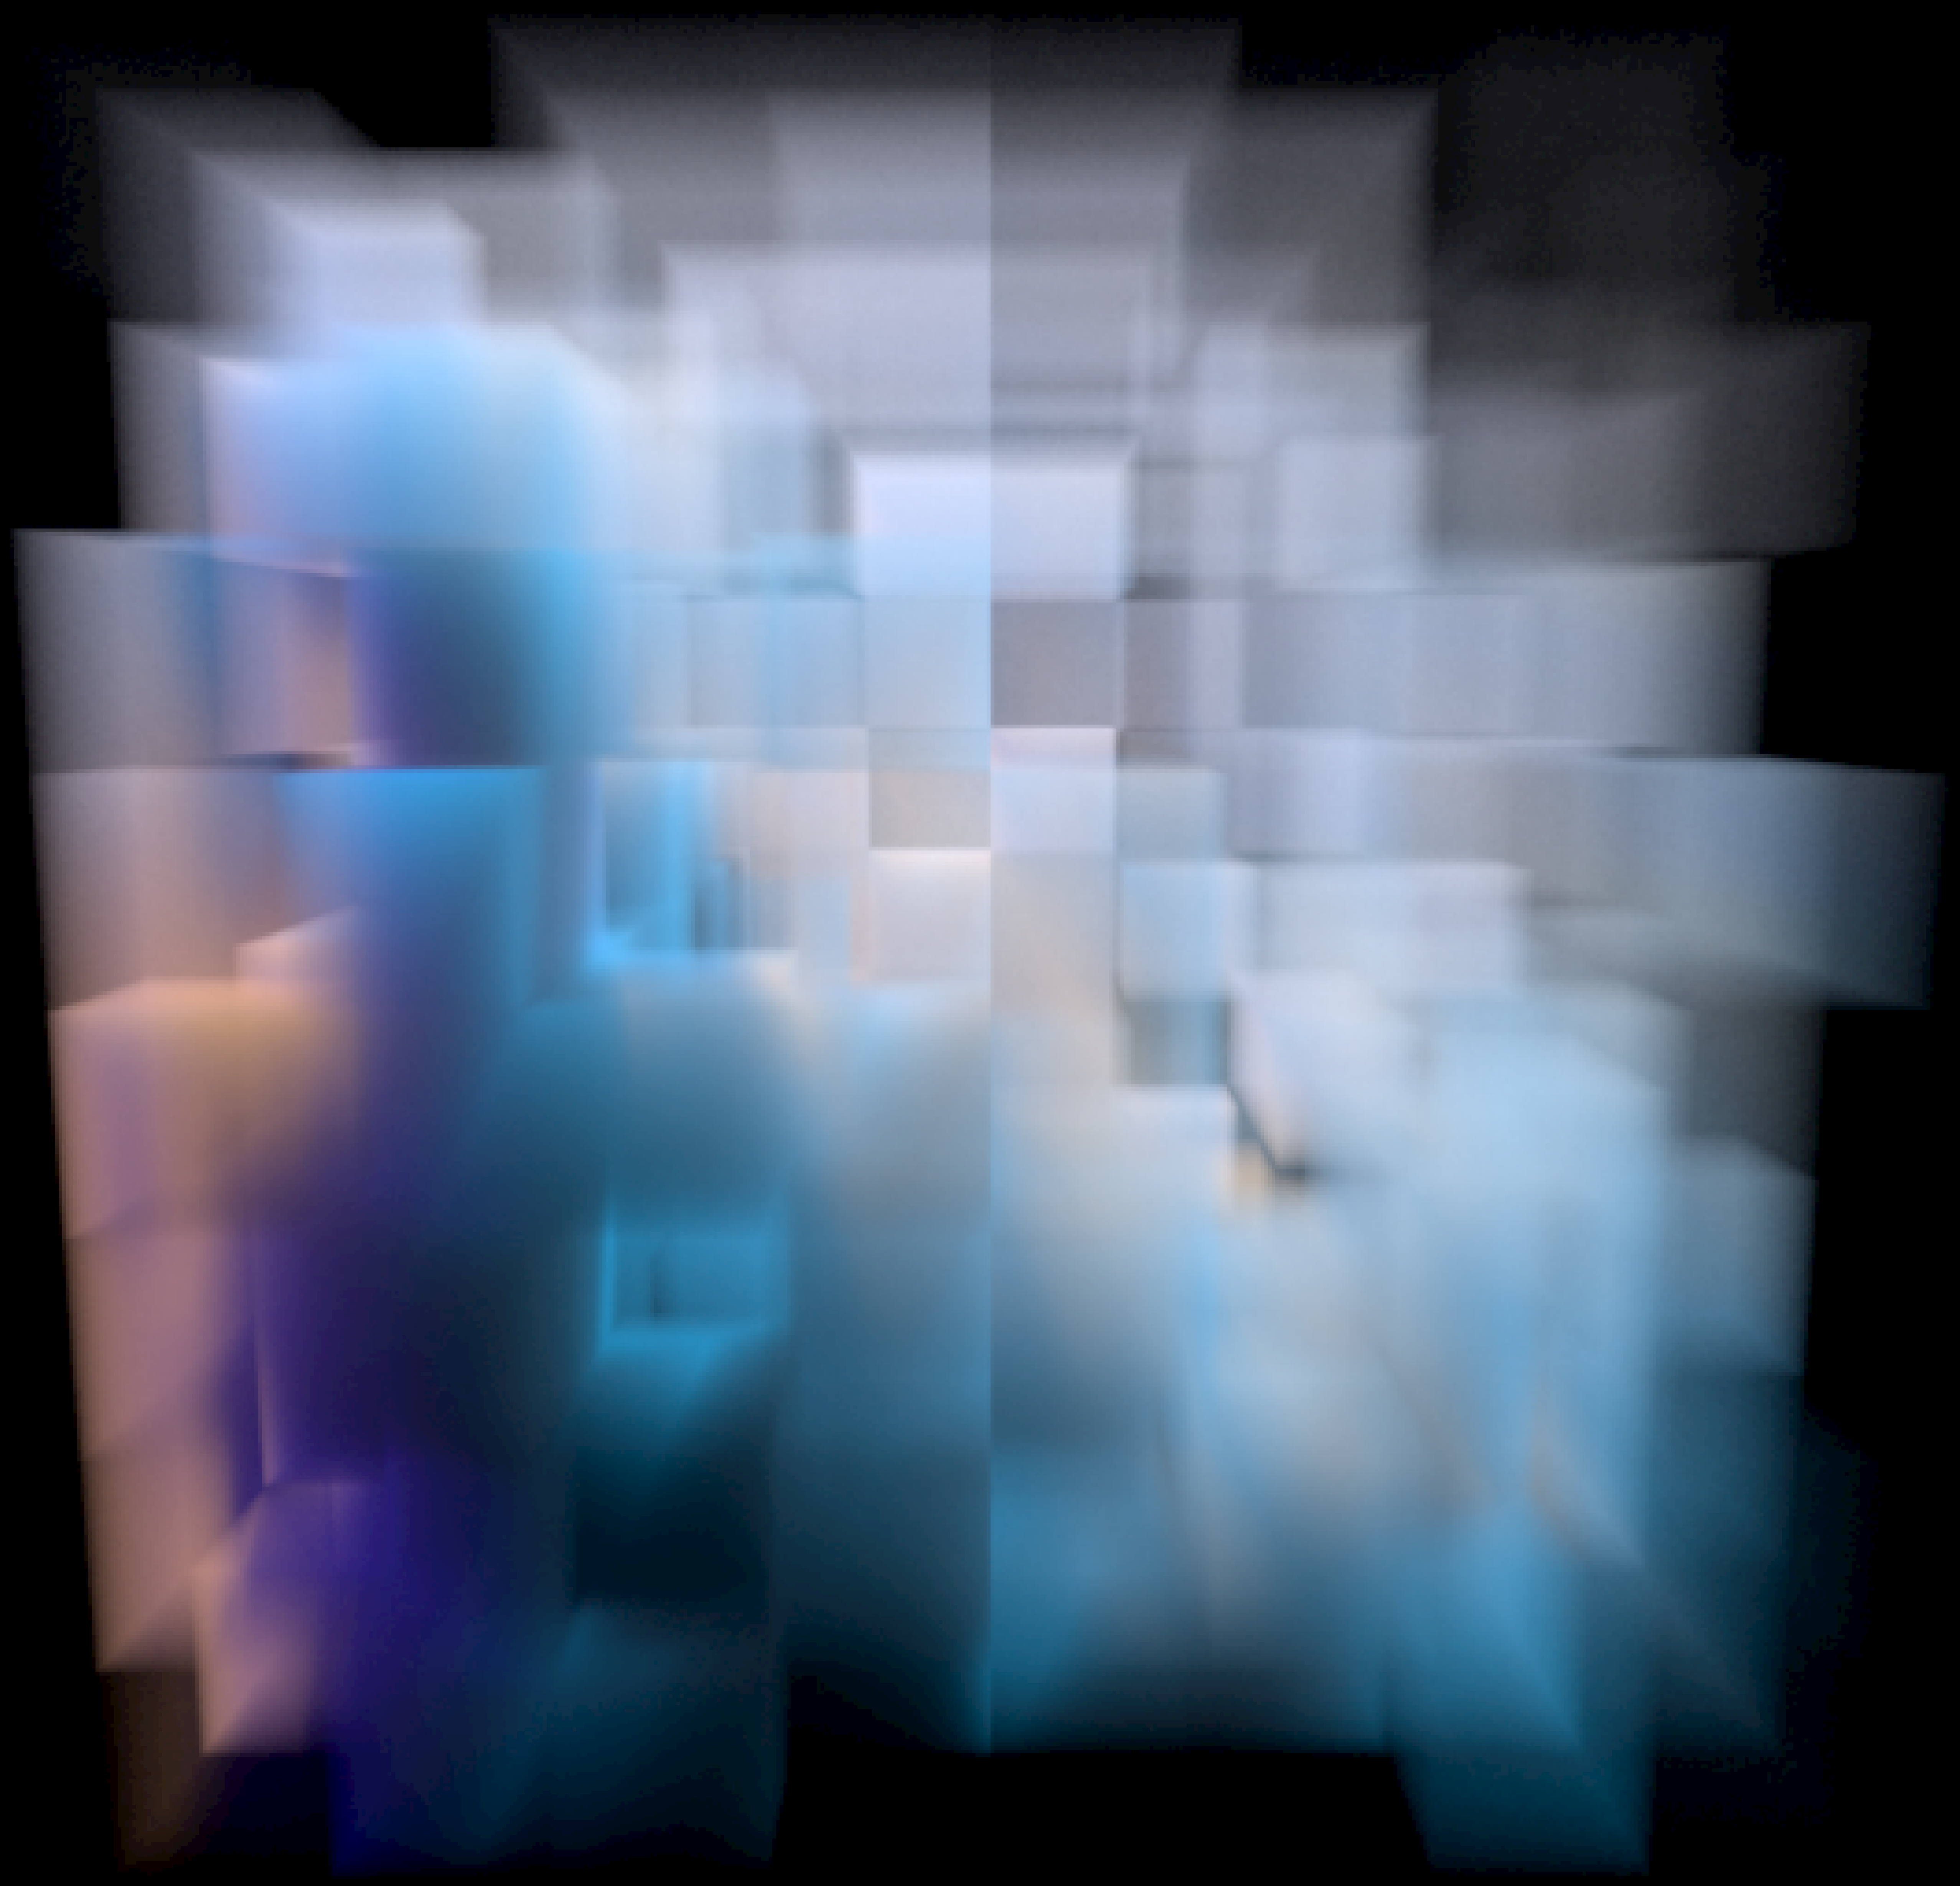

👍 tight 🐜 👍 robust

😢 visual quality

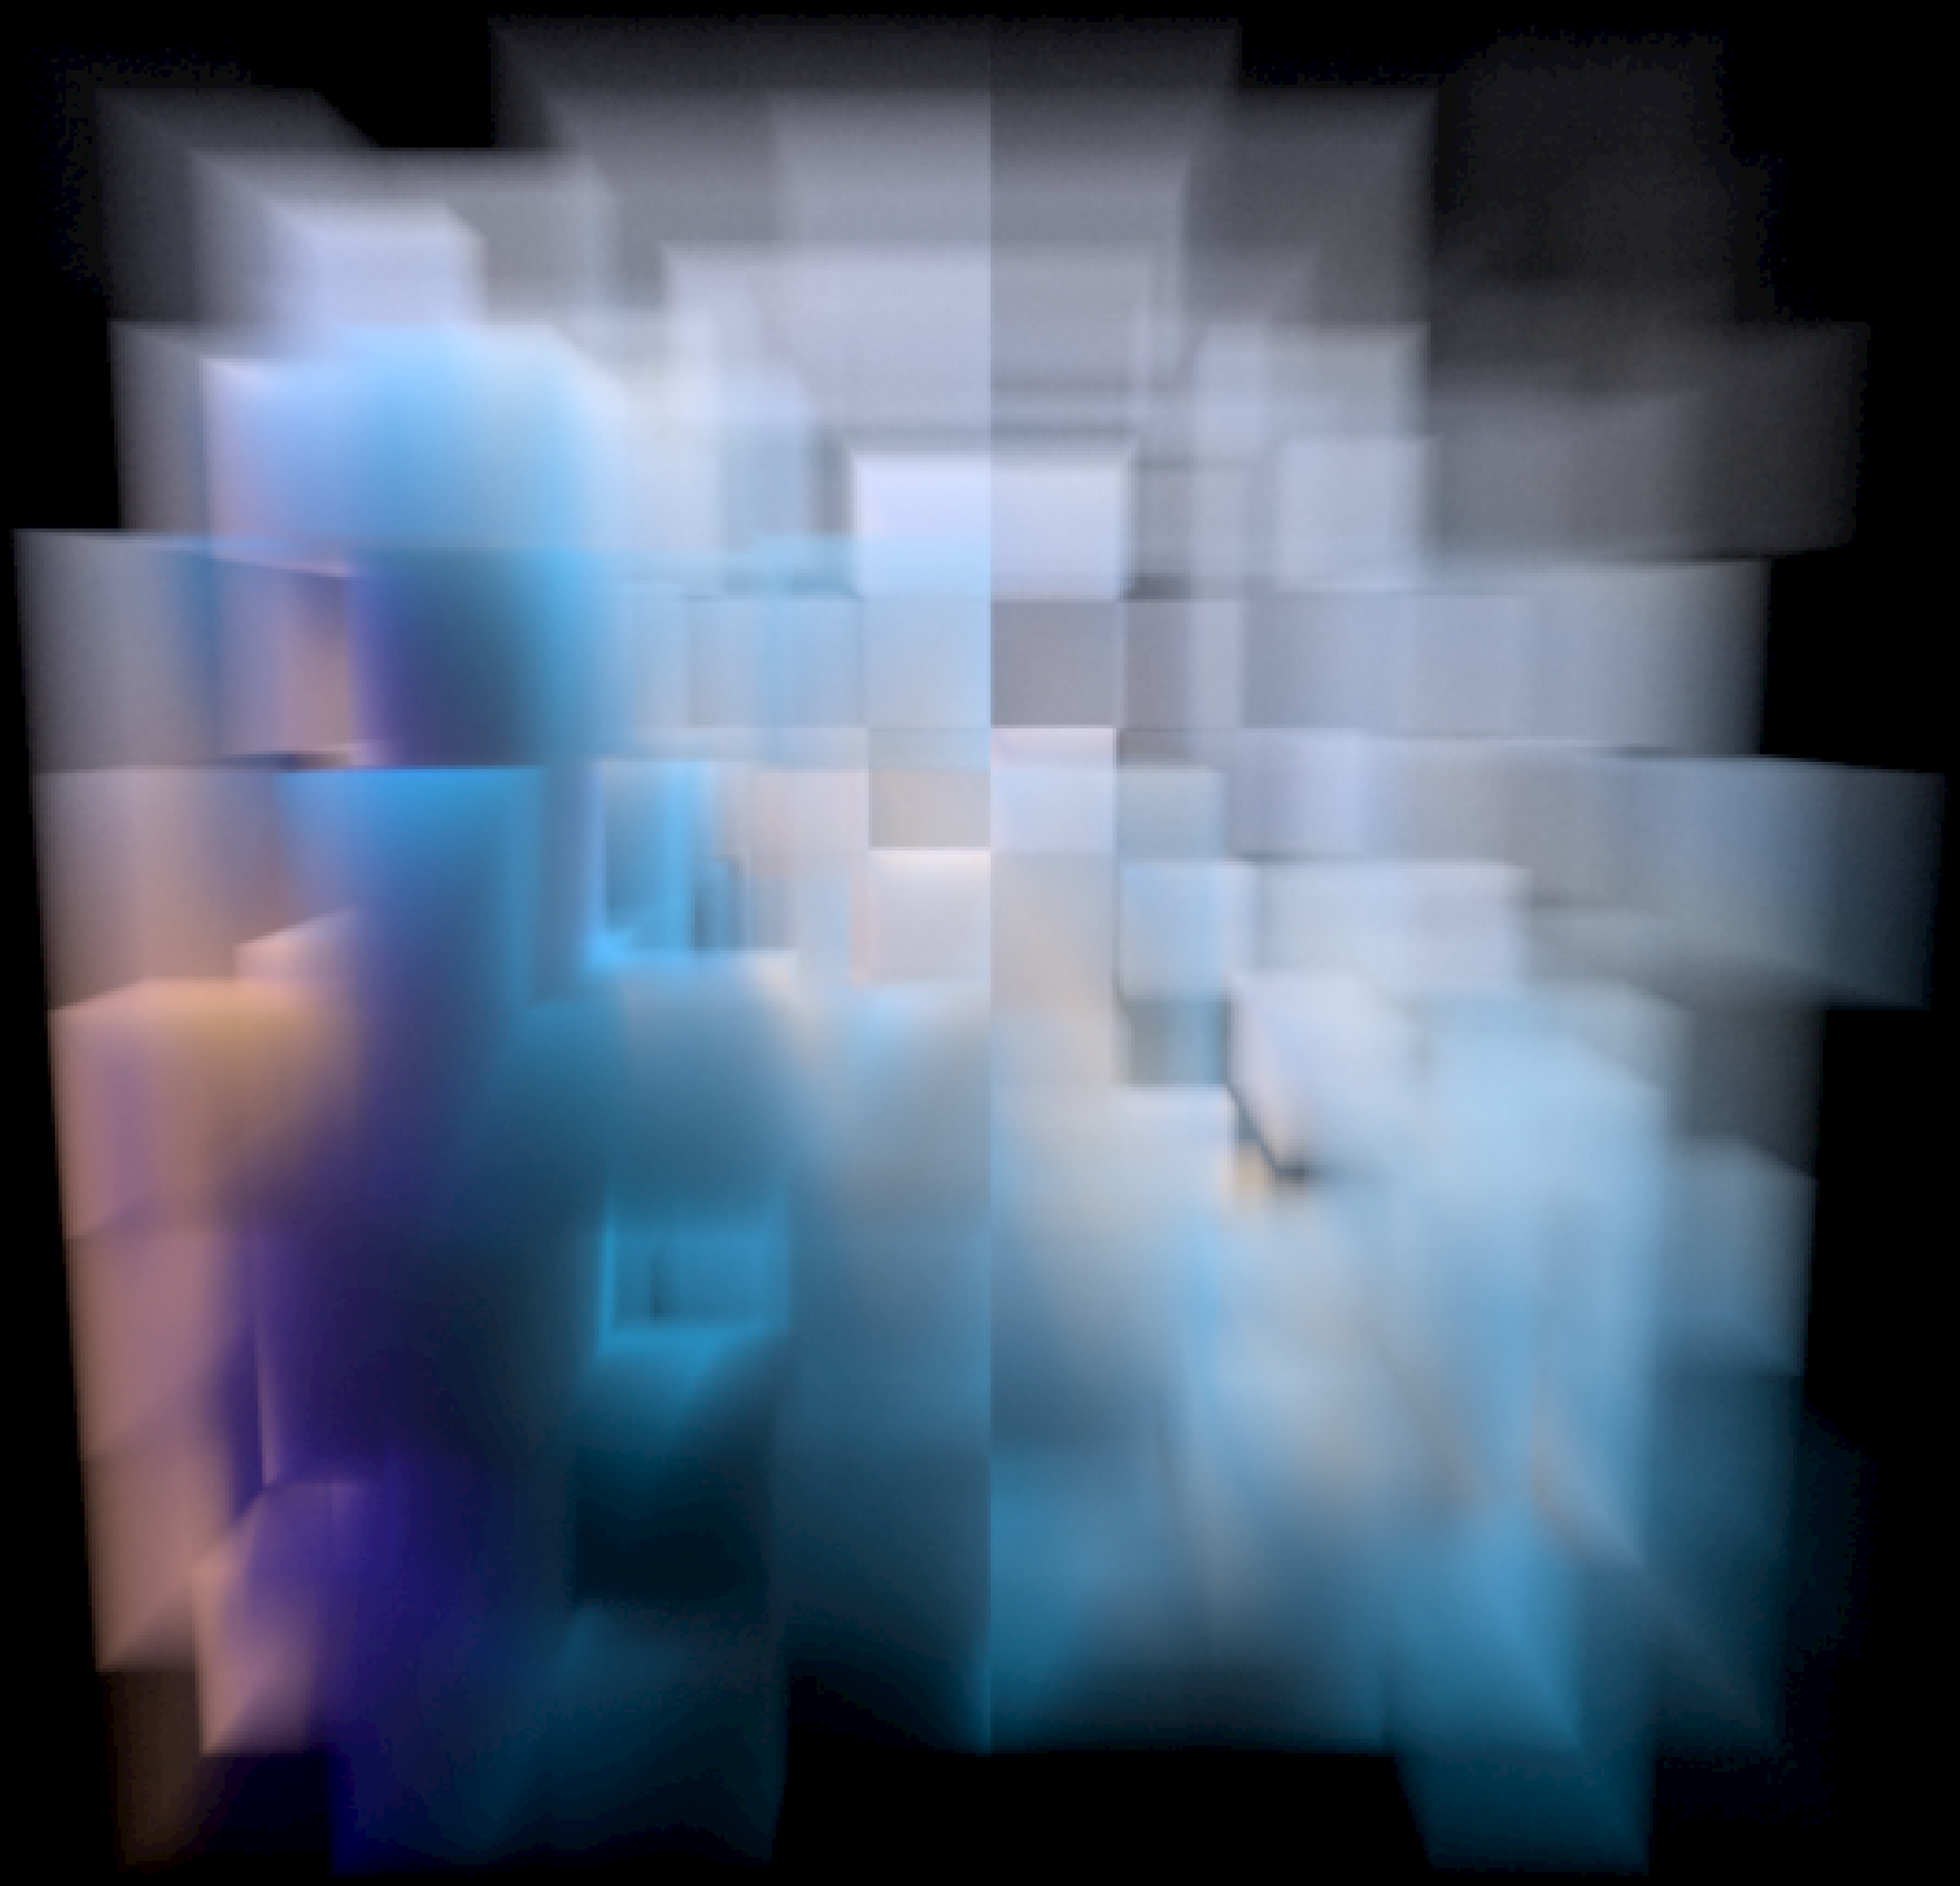

👍 tight 🐜 👍 robust

😢 visual quality

😢 preprocessing time

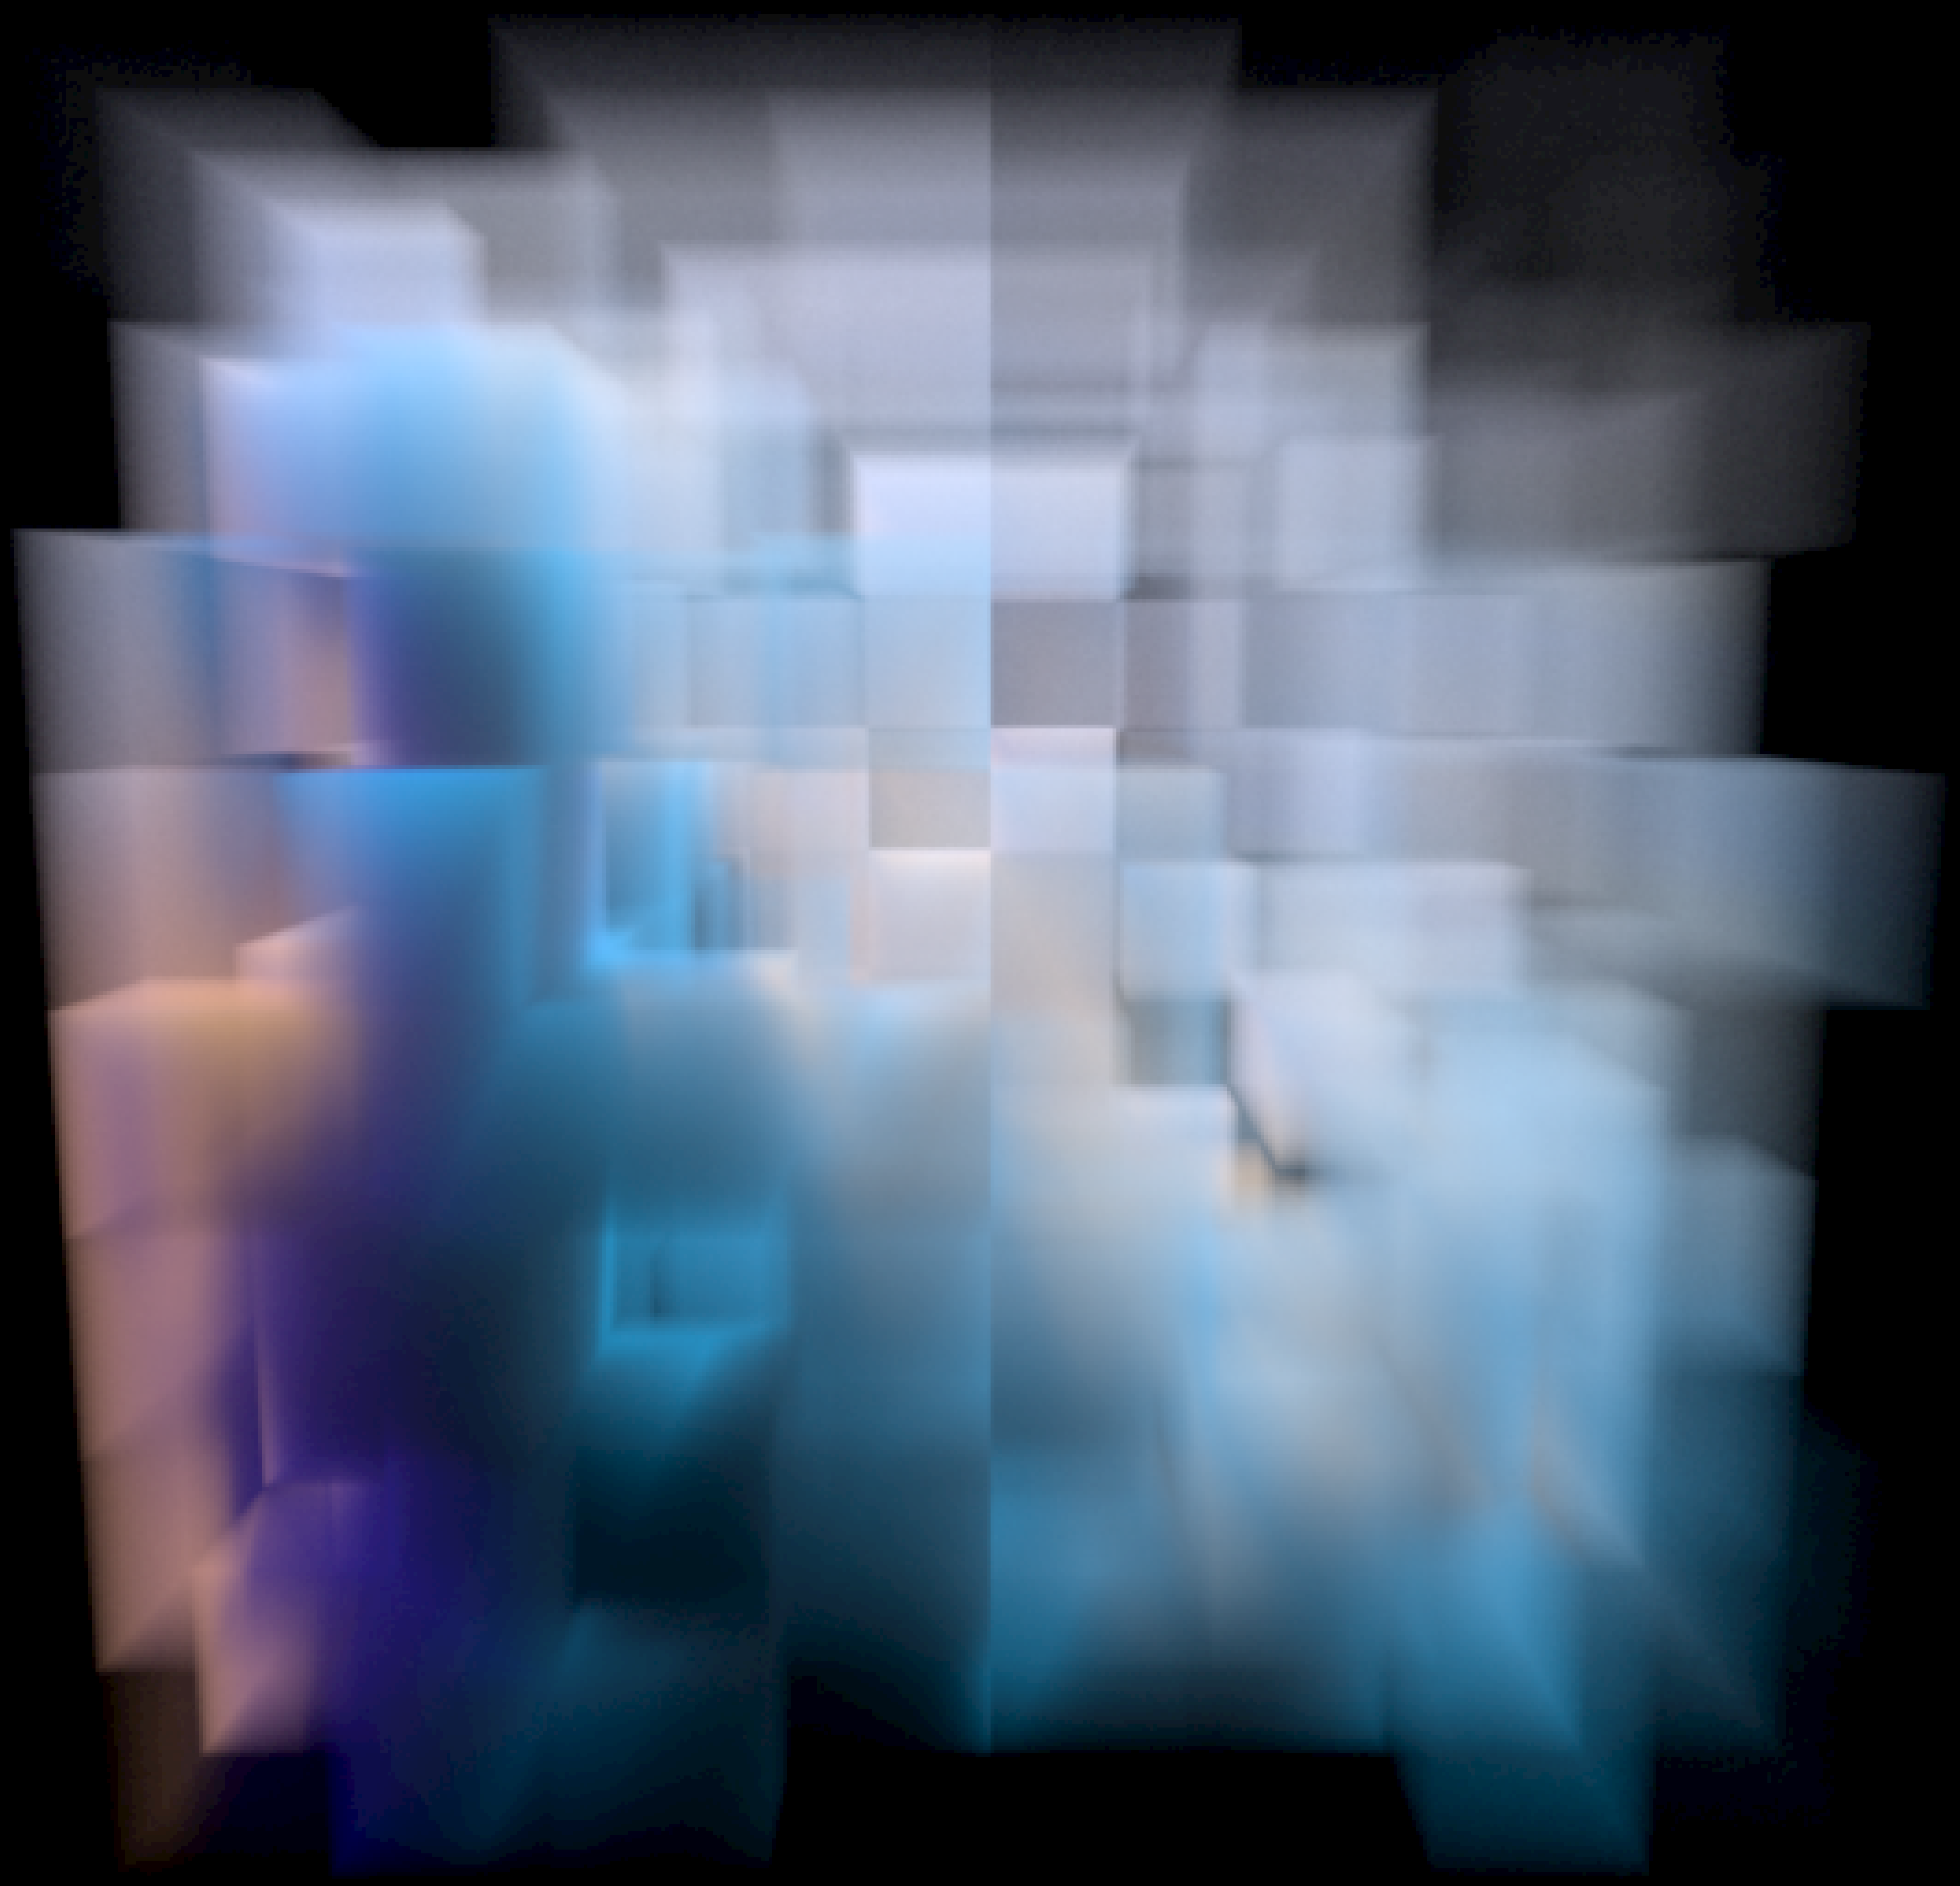

👍 tight 🐜 👍 robust

😢 visual quality

😢 preprocessing time

😢 storage / memory

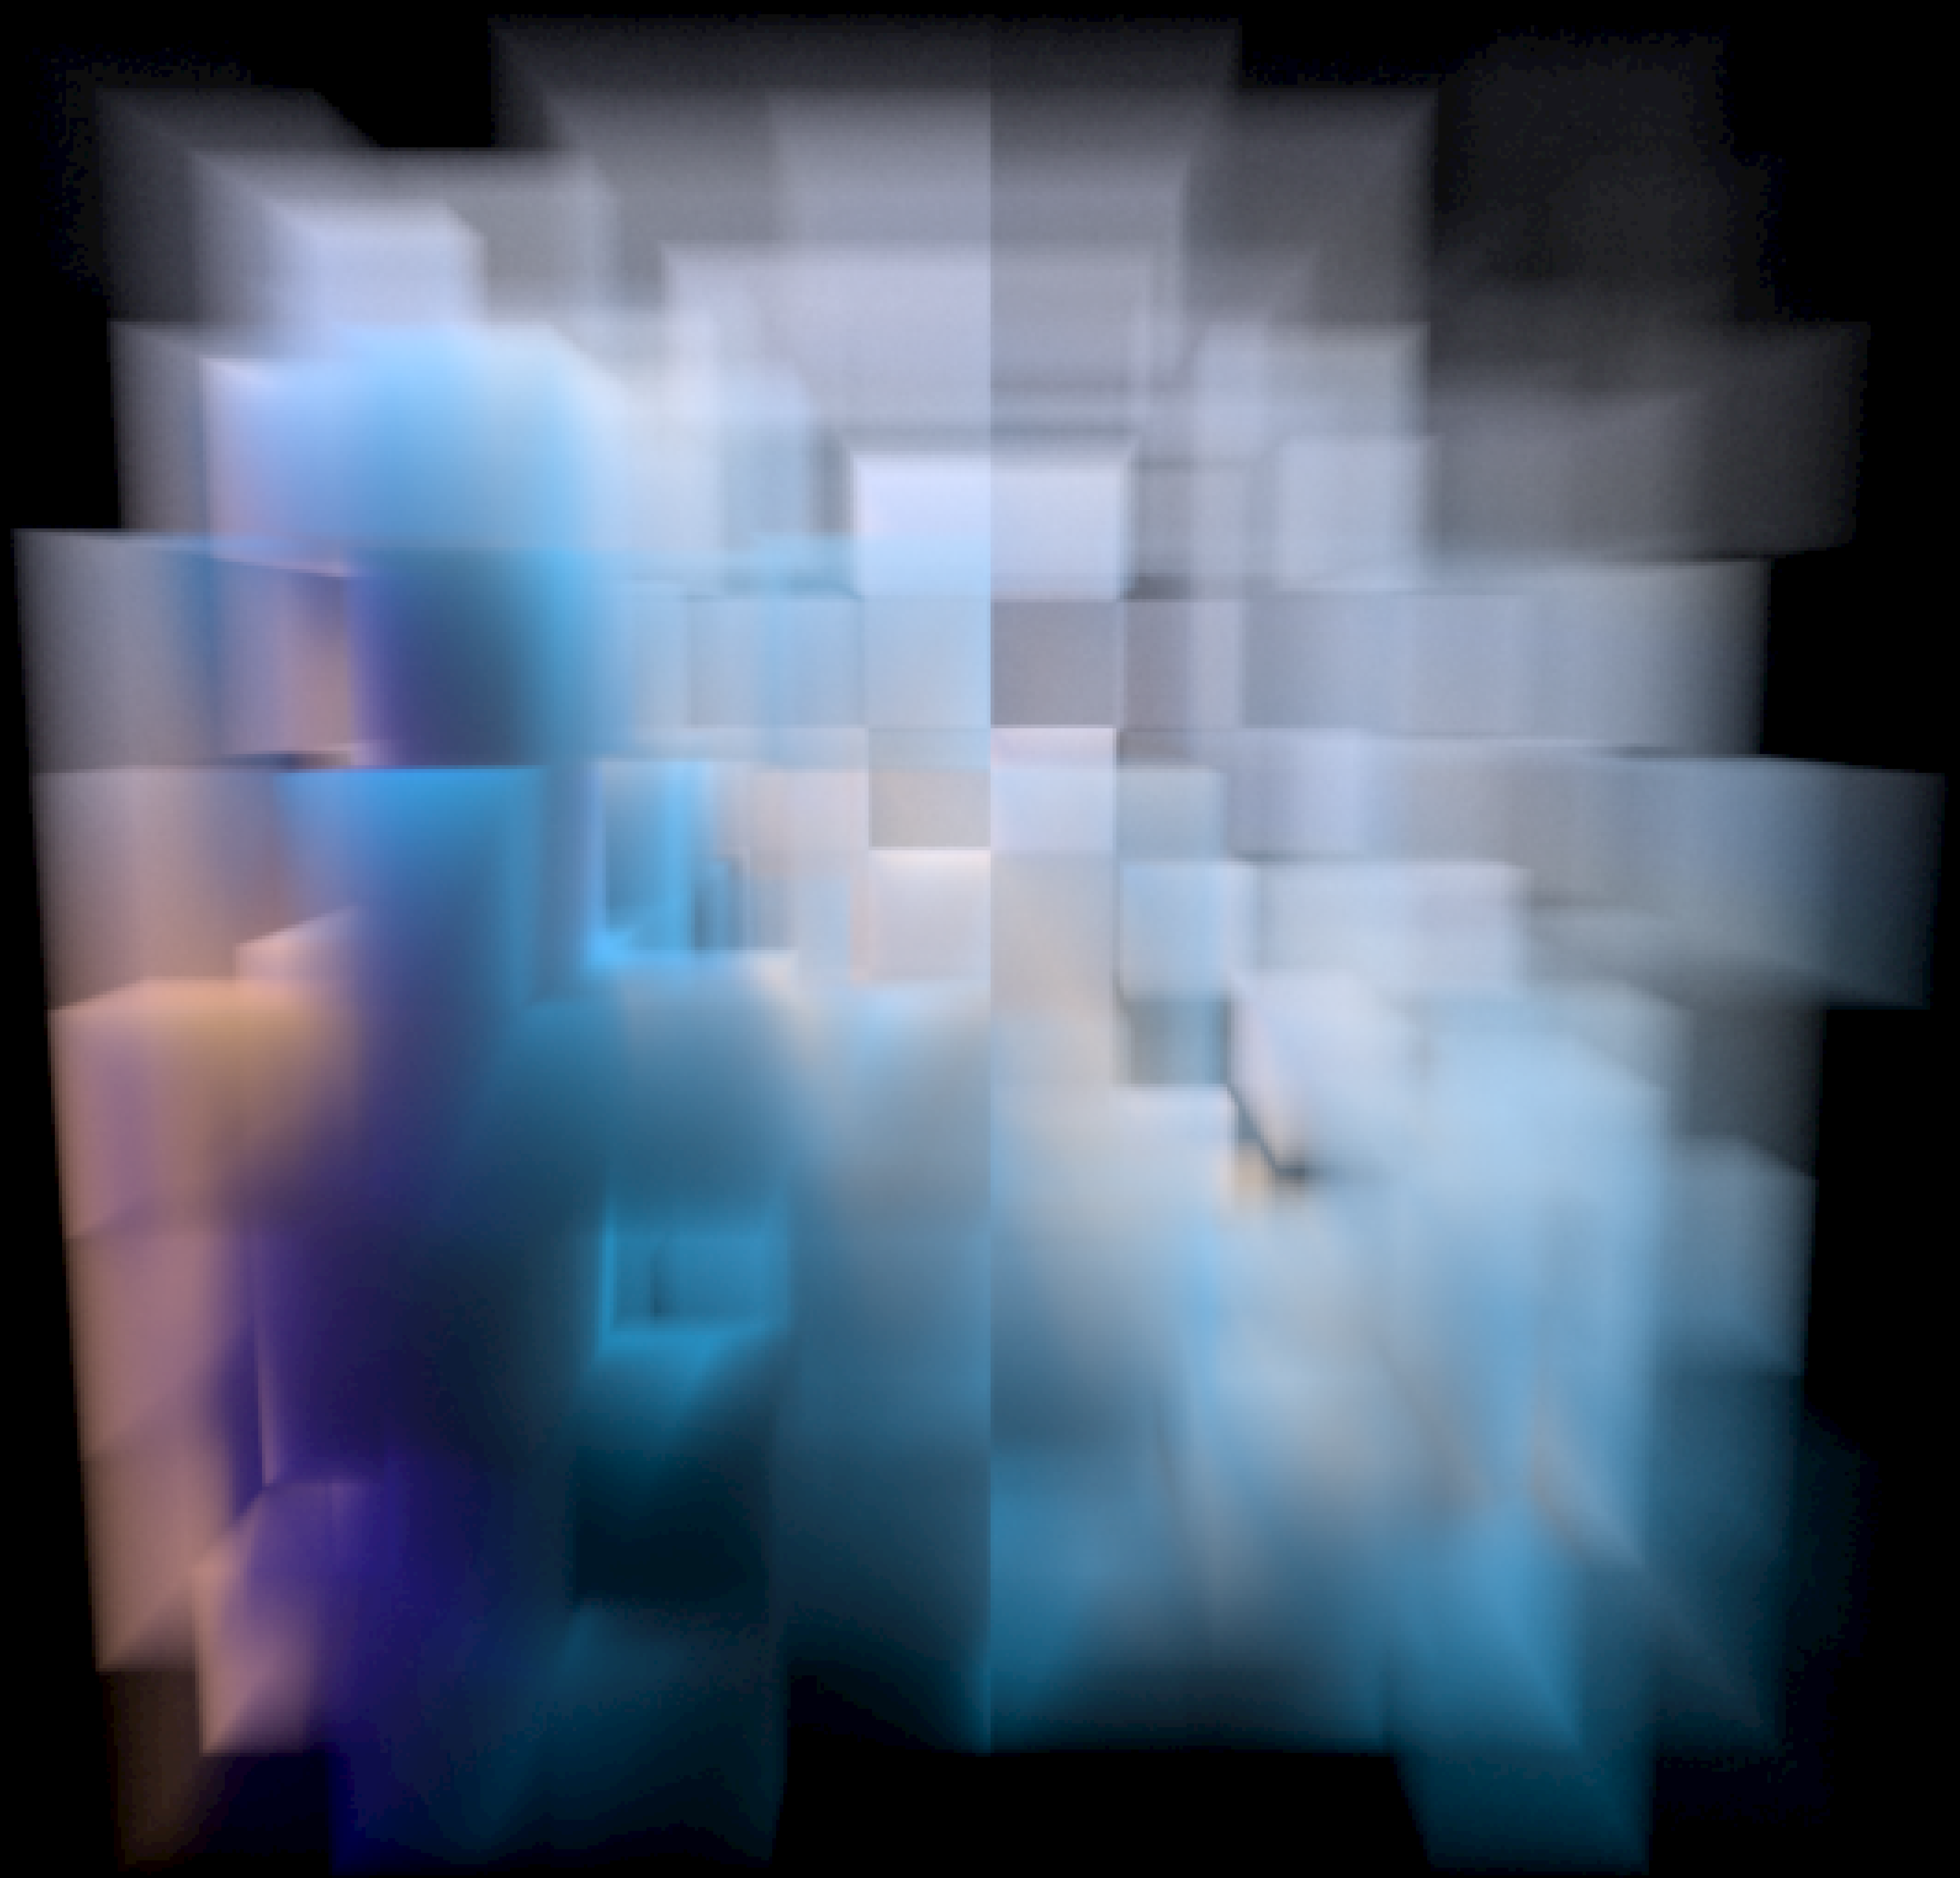

👍 tight 🐜 👍 robust

😢 visual quality

😢 preprocessing time

😢 storage / memory

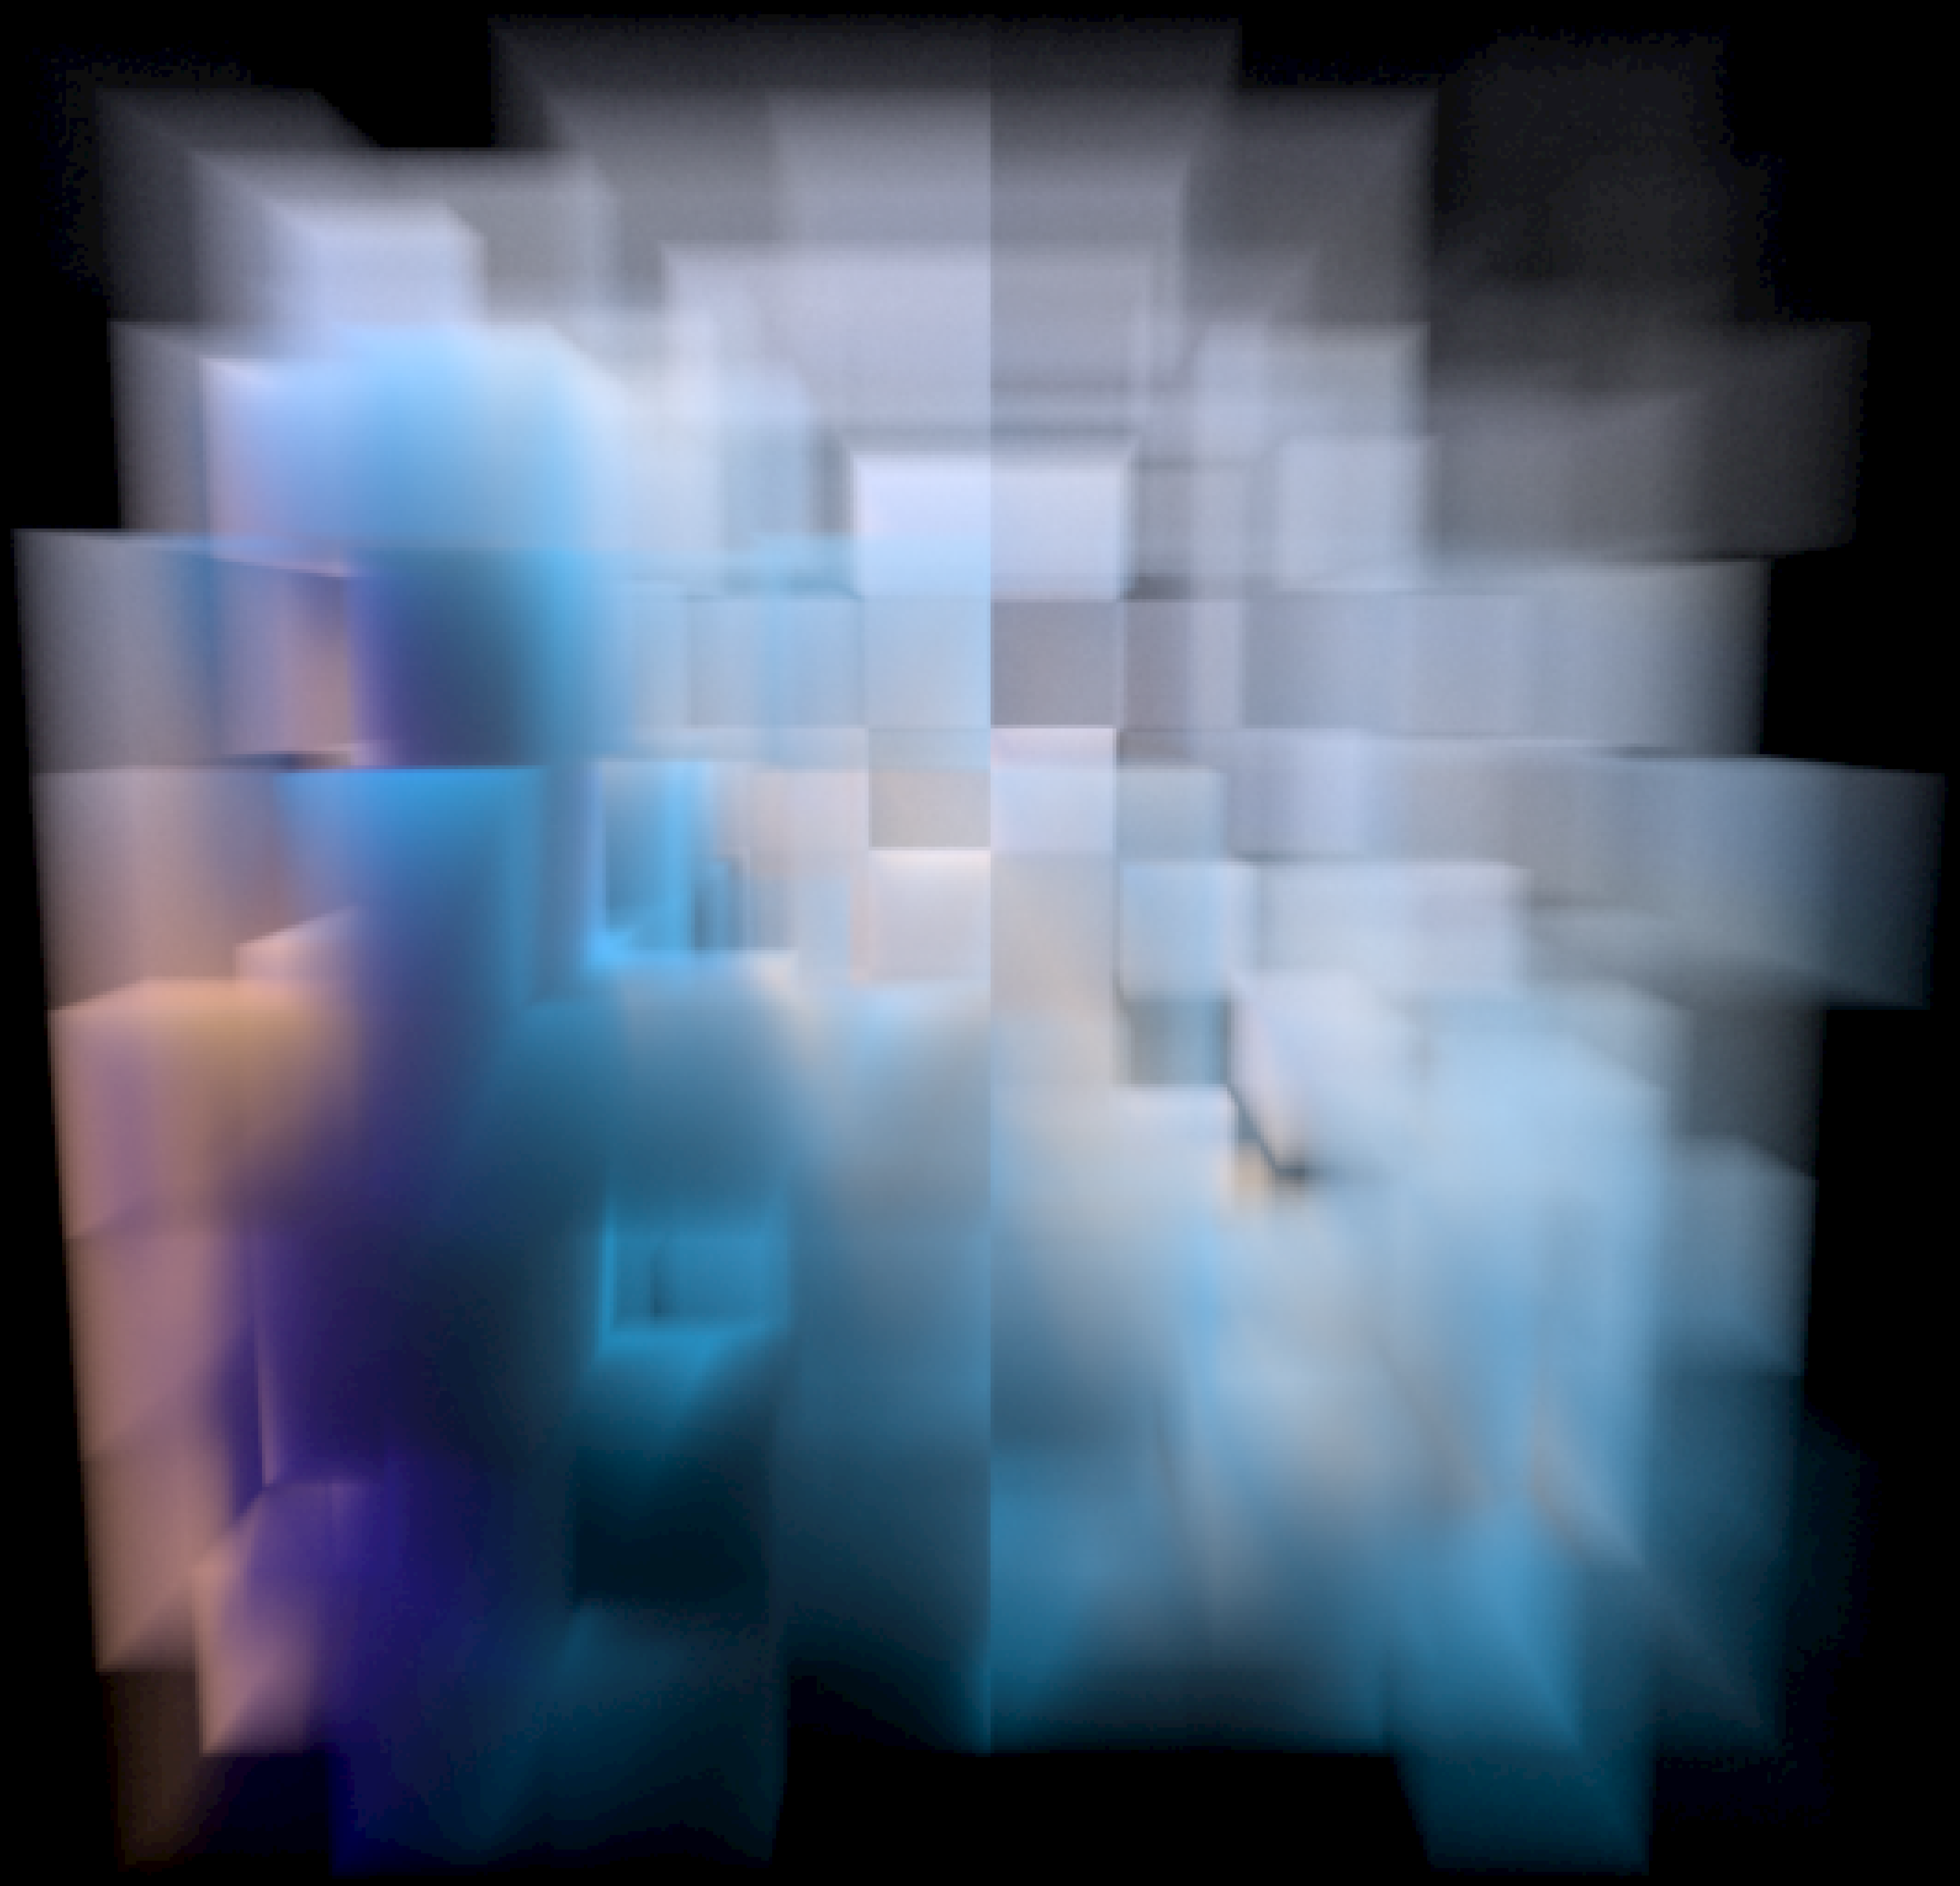

👍 tight 🐜 👍 robust

😄 works for any majorant

😢 visual quality

😢 preprocessing time

😢 storage / memory

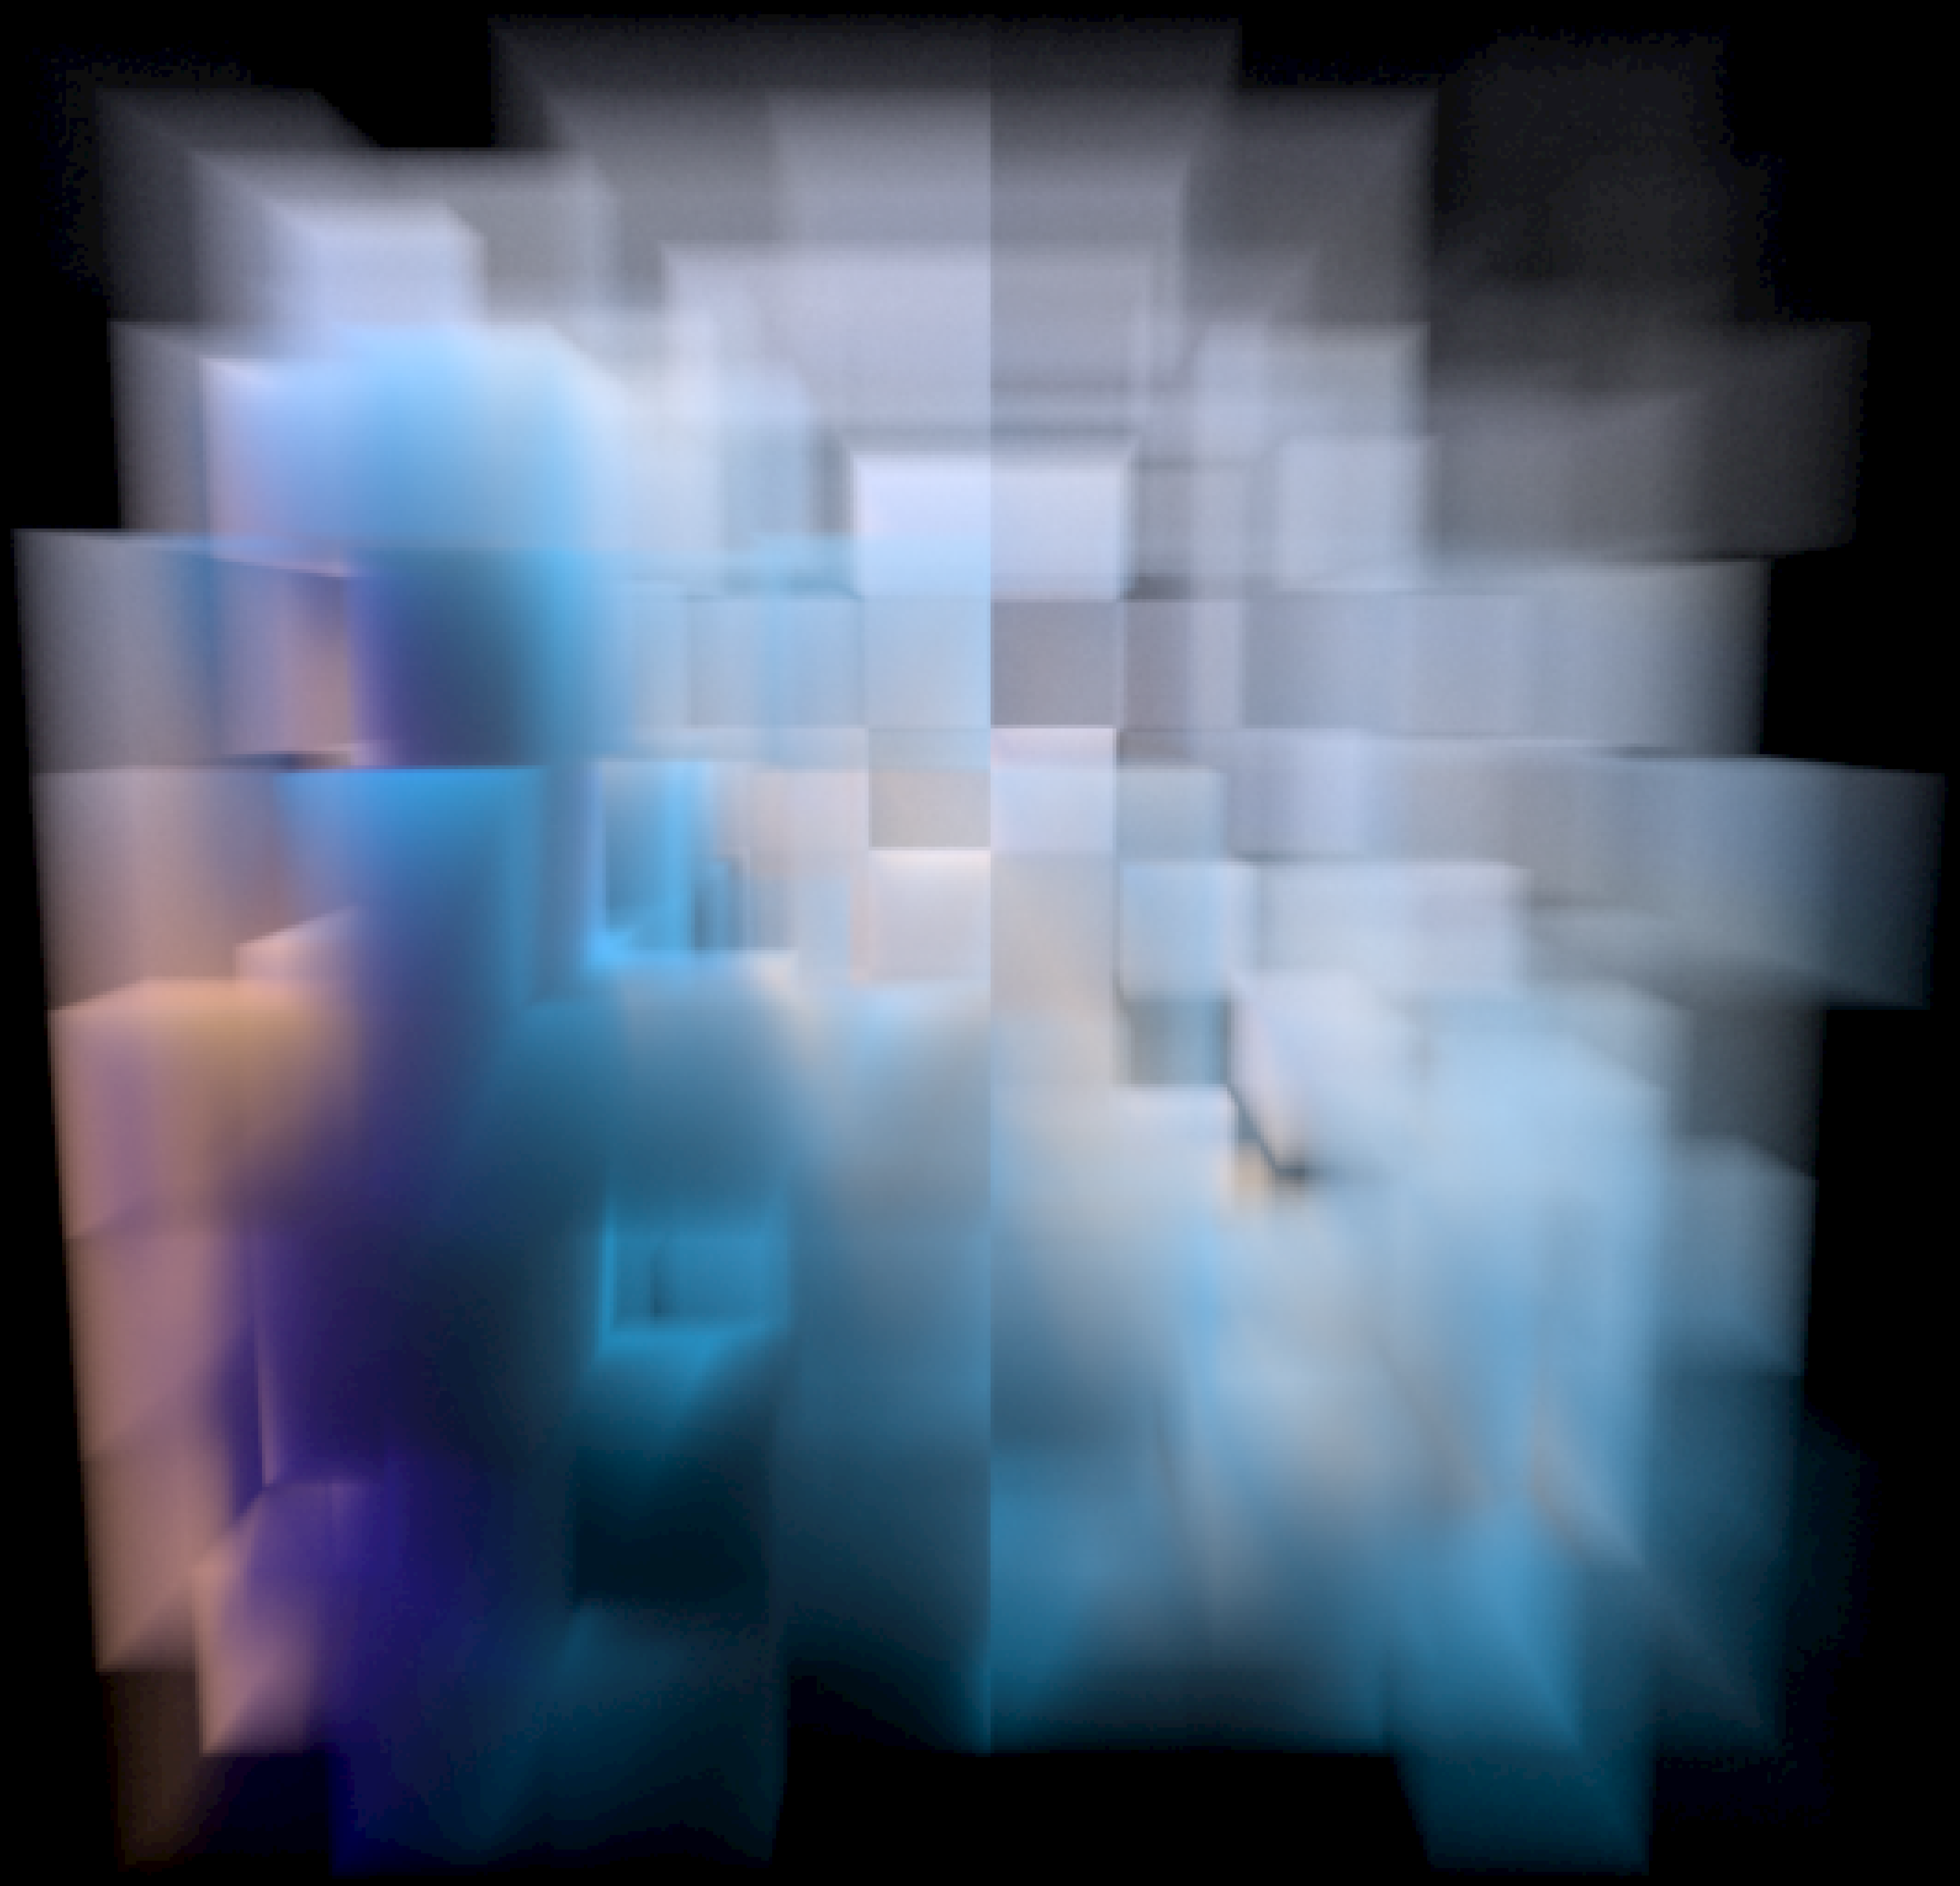

👍 tight 🐜 👍 robust

😄 works for any majorant

😢 visual quality

😢 preprocessing time

😢 storage / memory

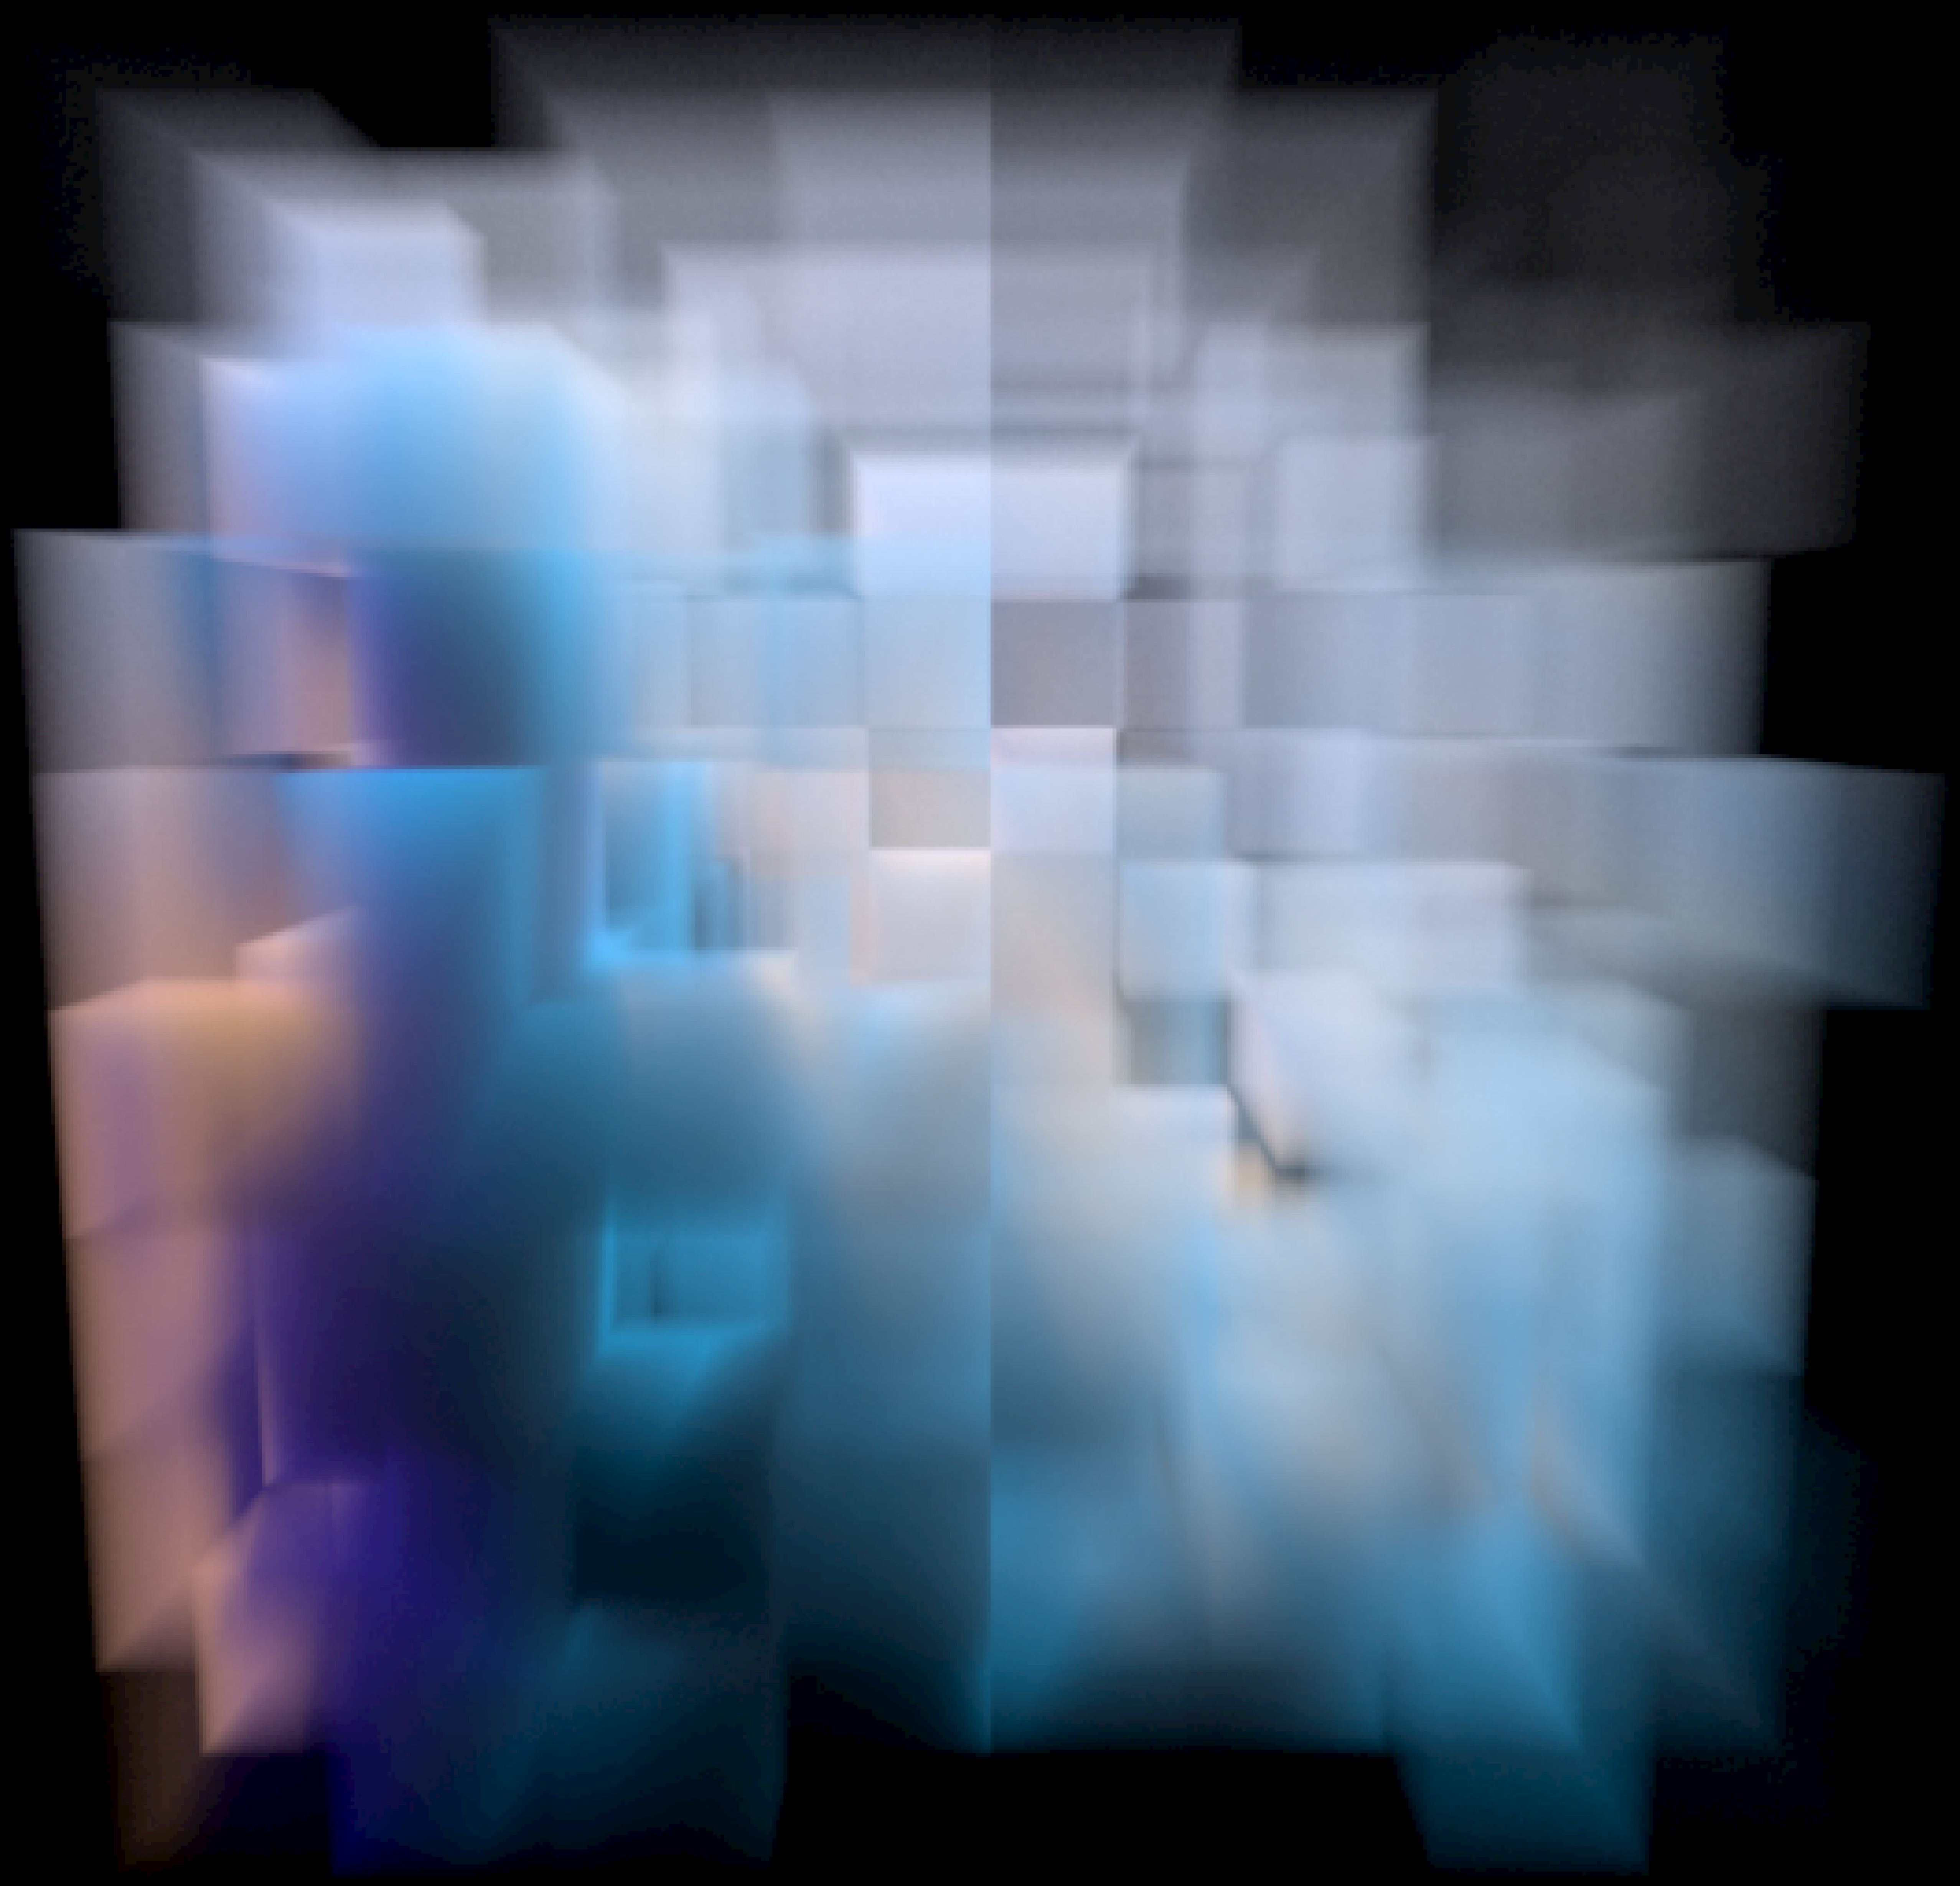

👍 tight 🐜 👍 robust

😄 works for any majorant

😢 visual quality

😢 preprocessing time

😢 storage / memory

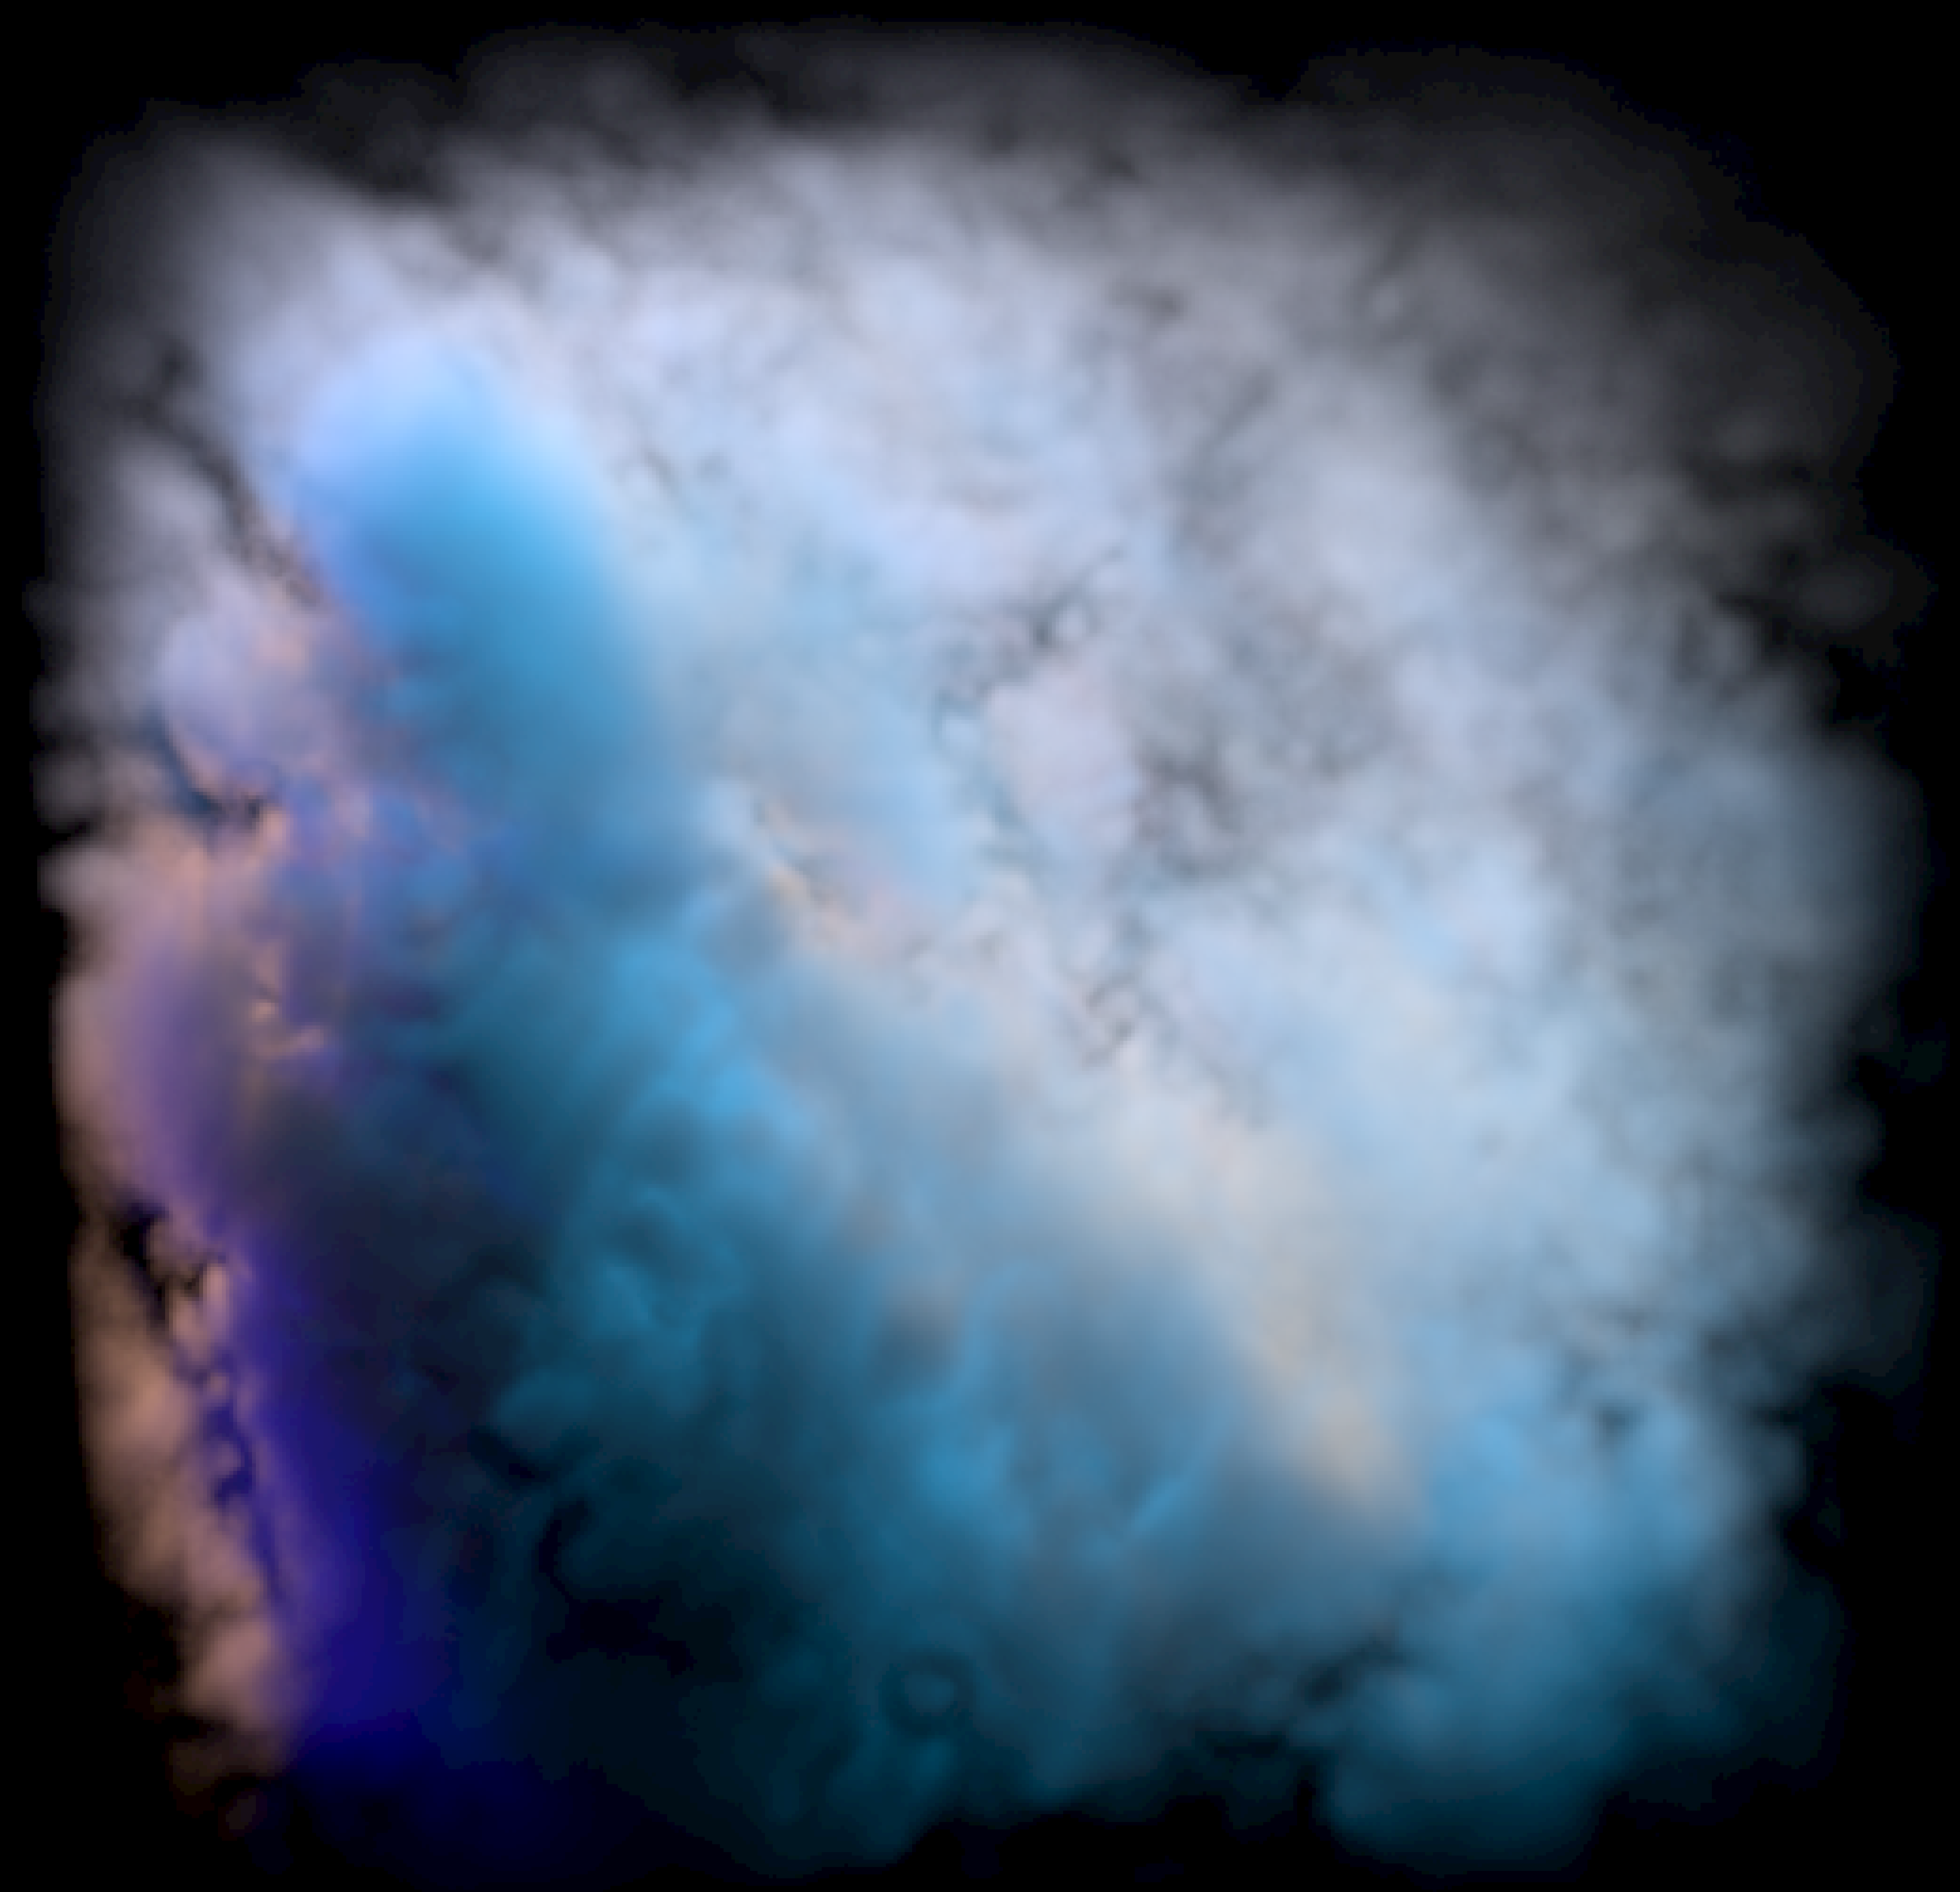

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😢 preprocessing time

😢 storage / memory

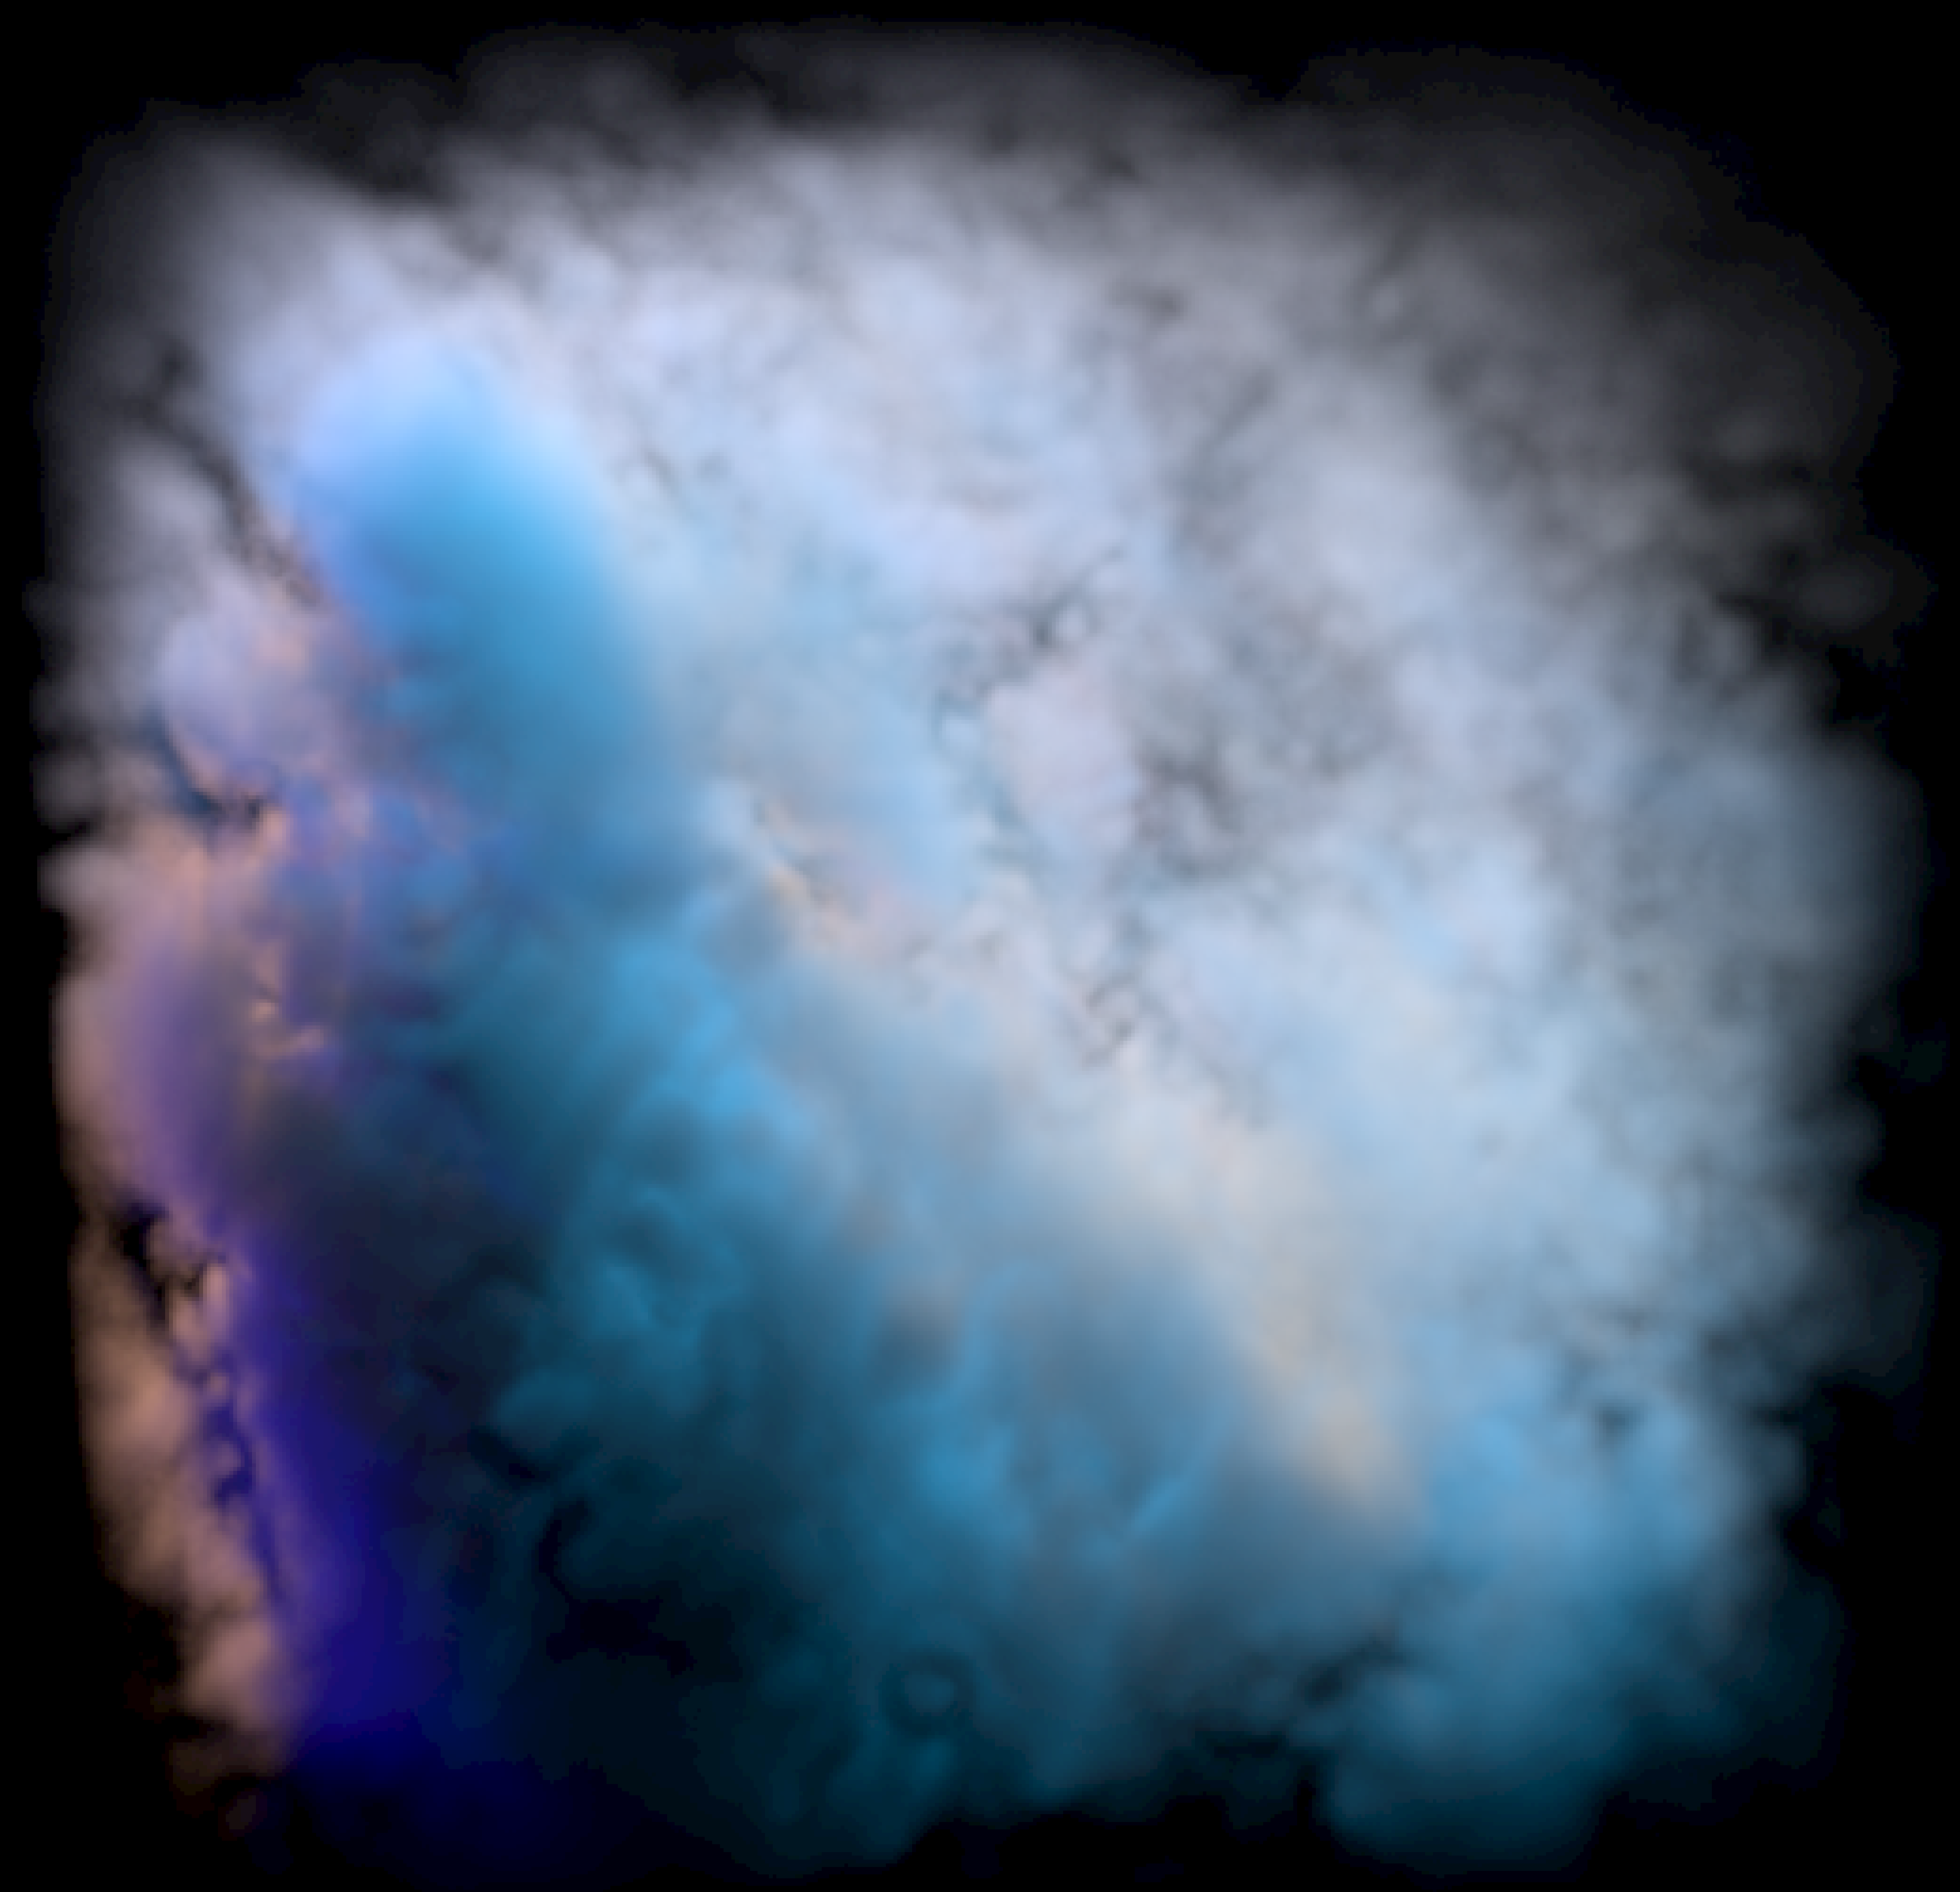

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😢 preprocessing time

😢 storage / memory

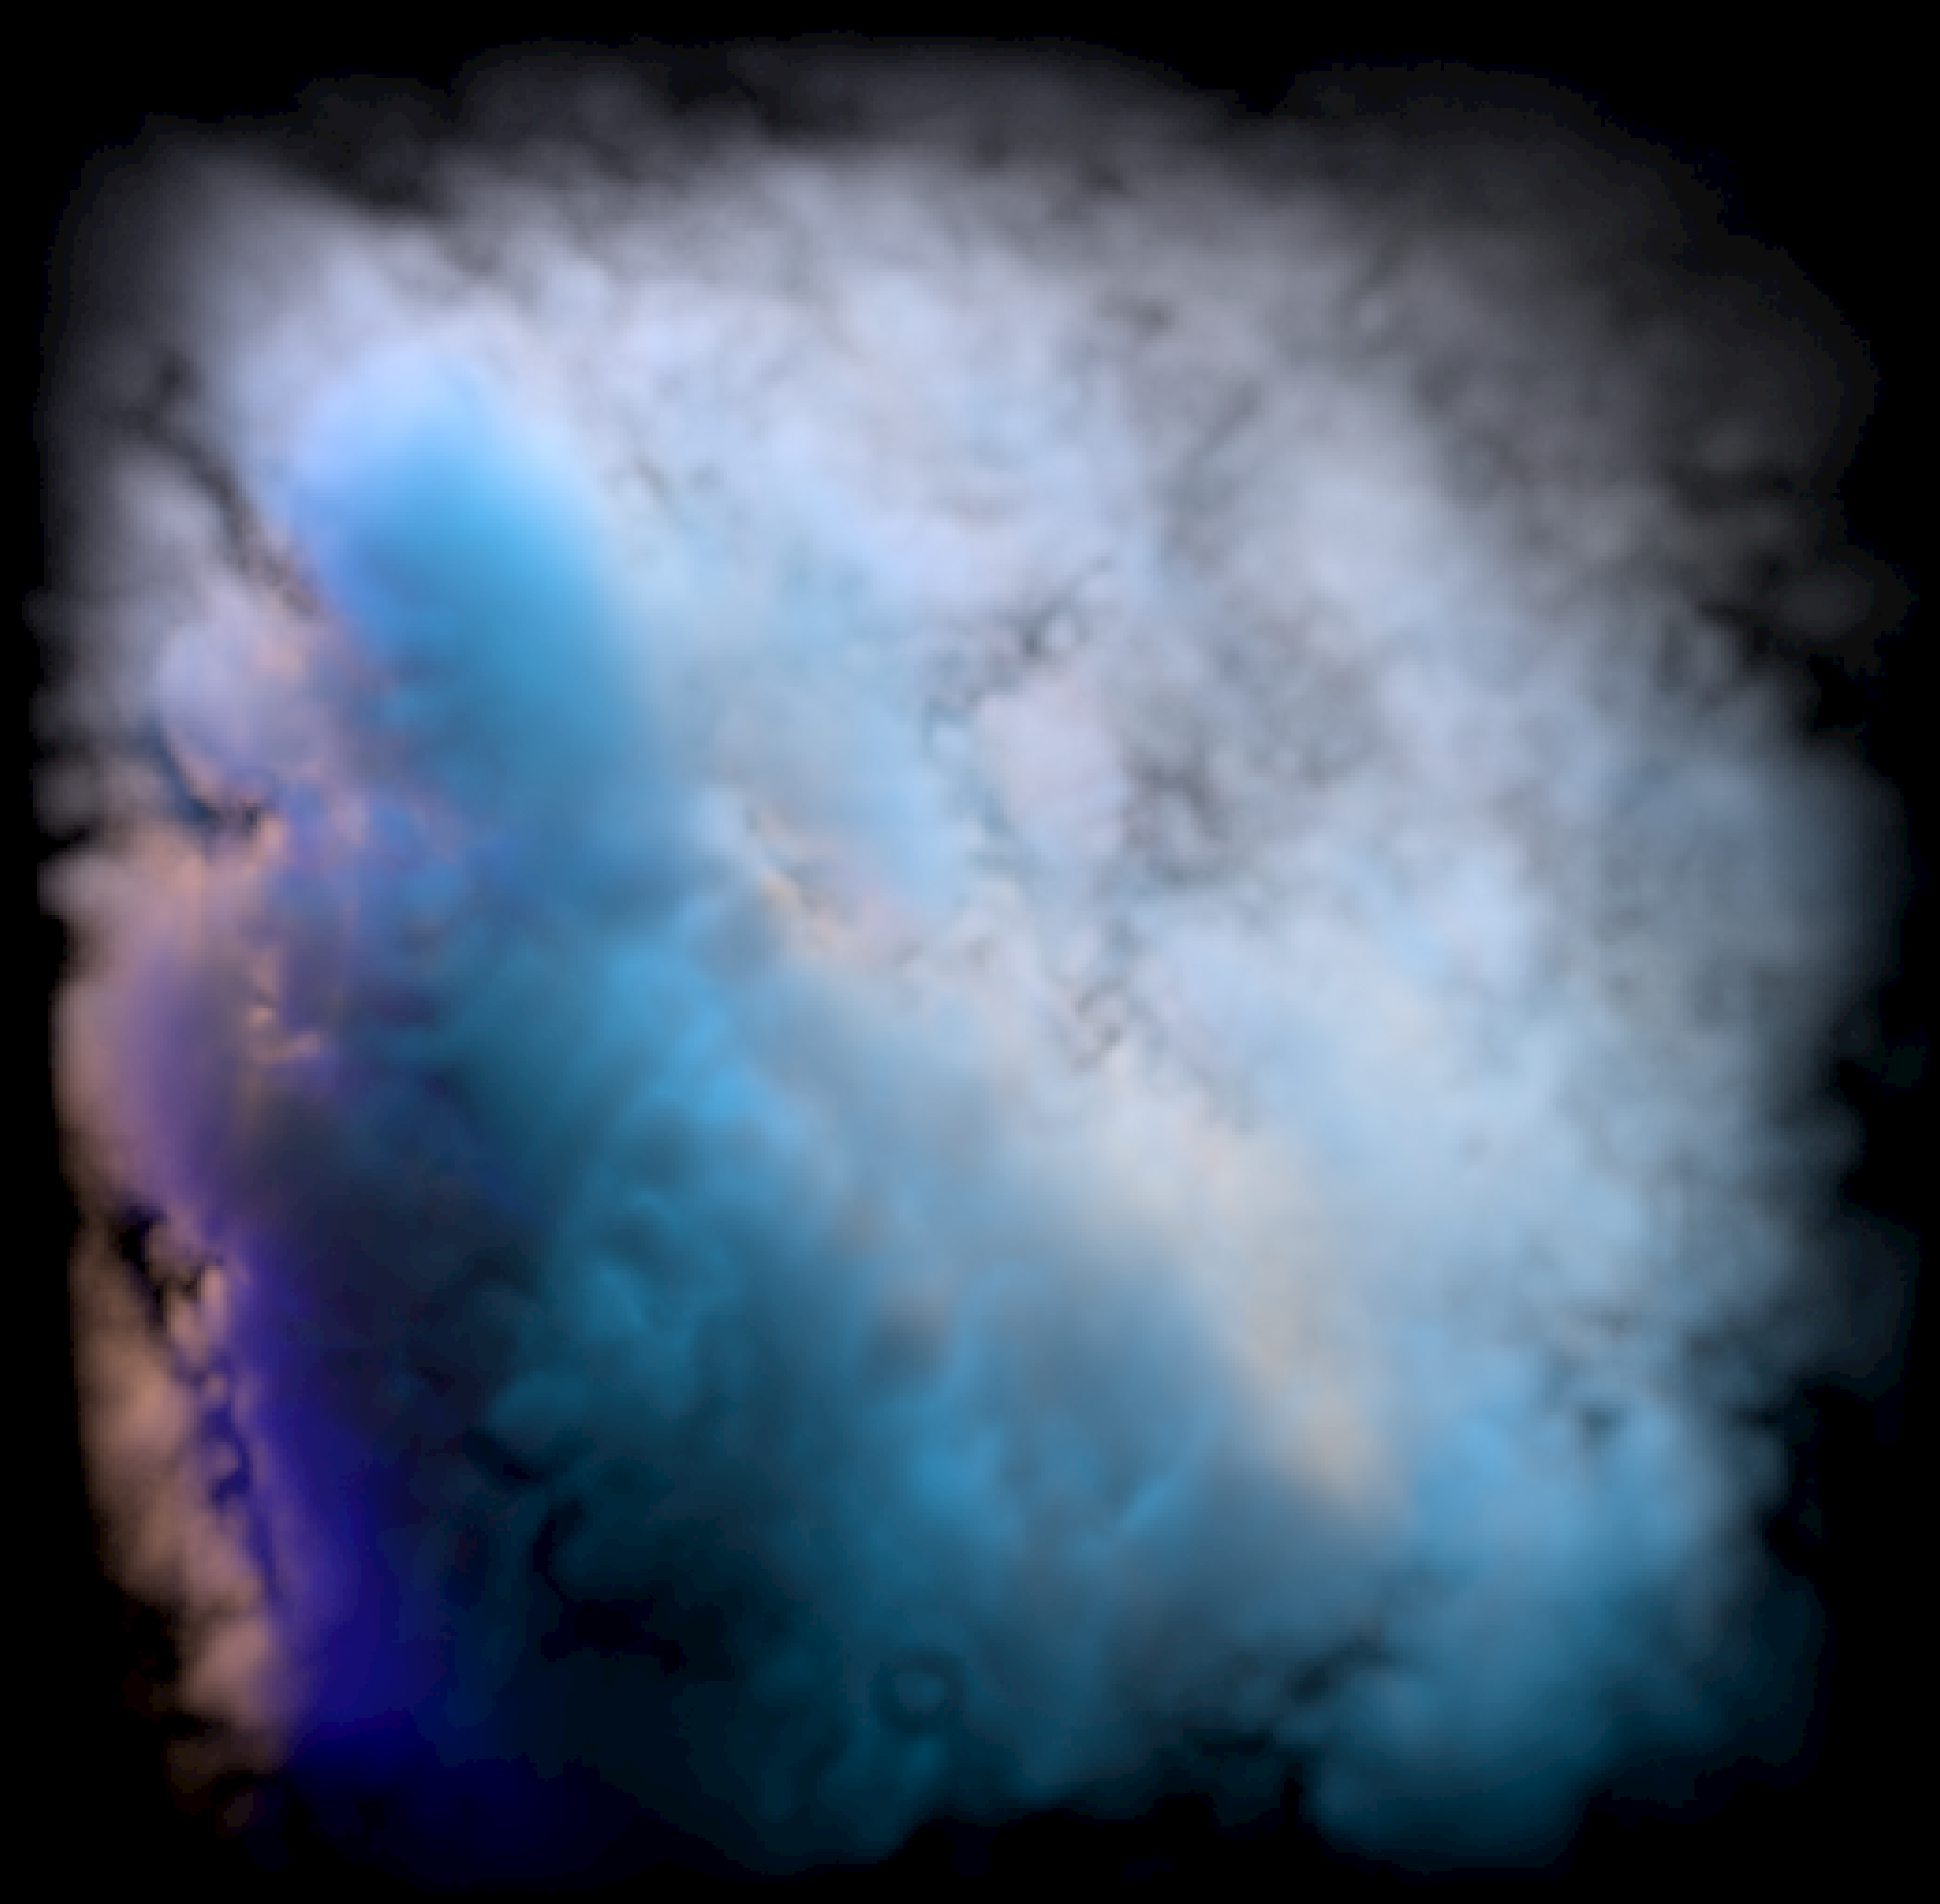

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😢 storage / memory

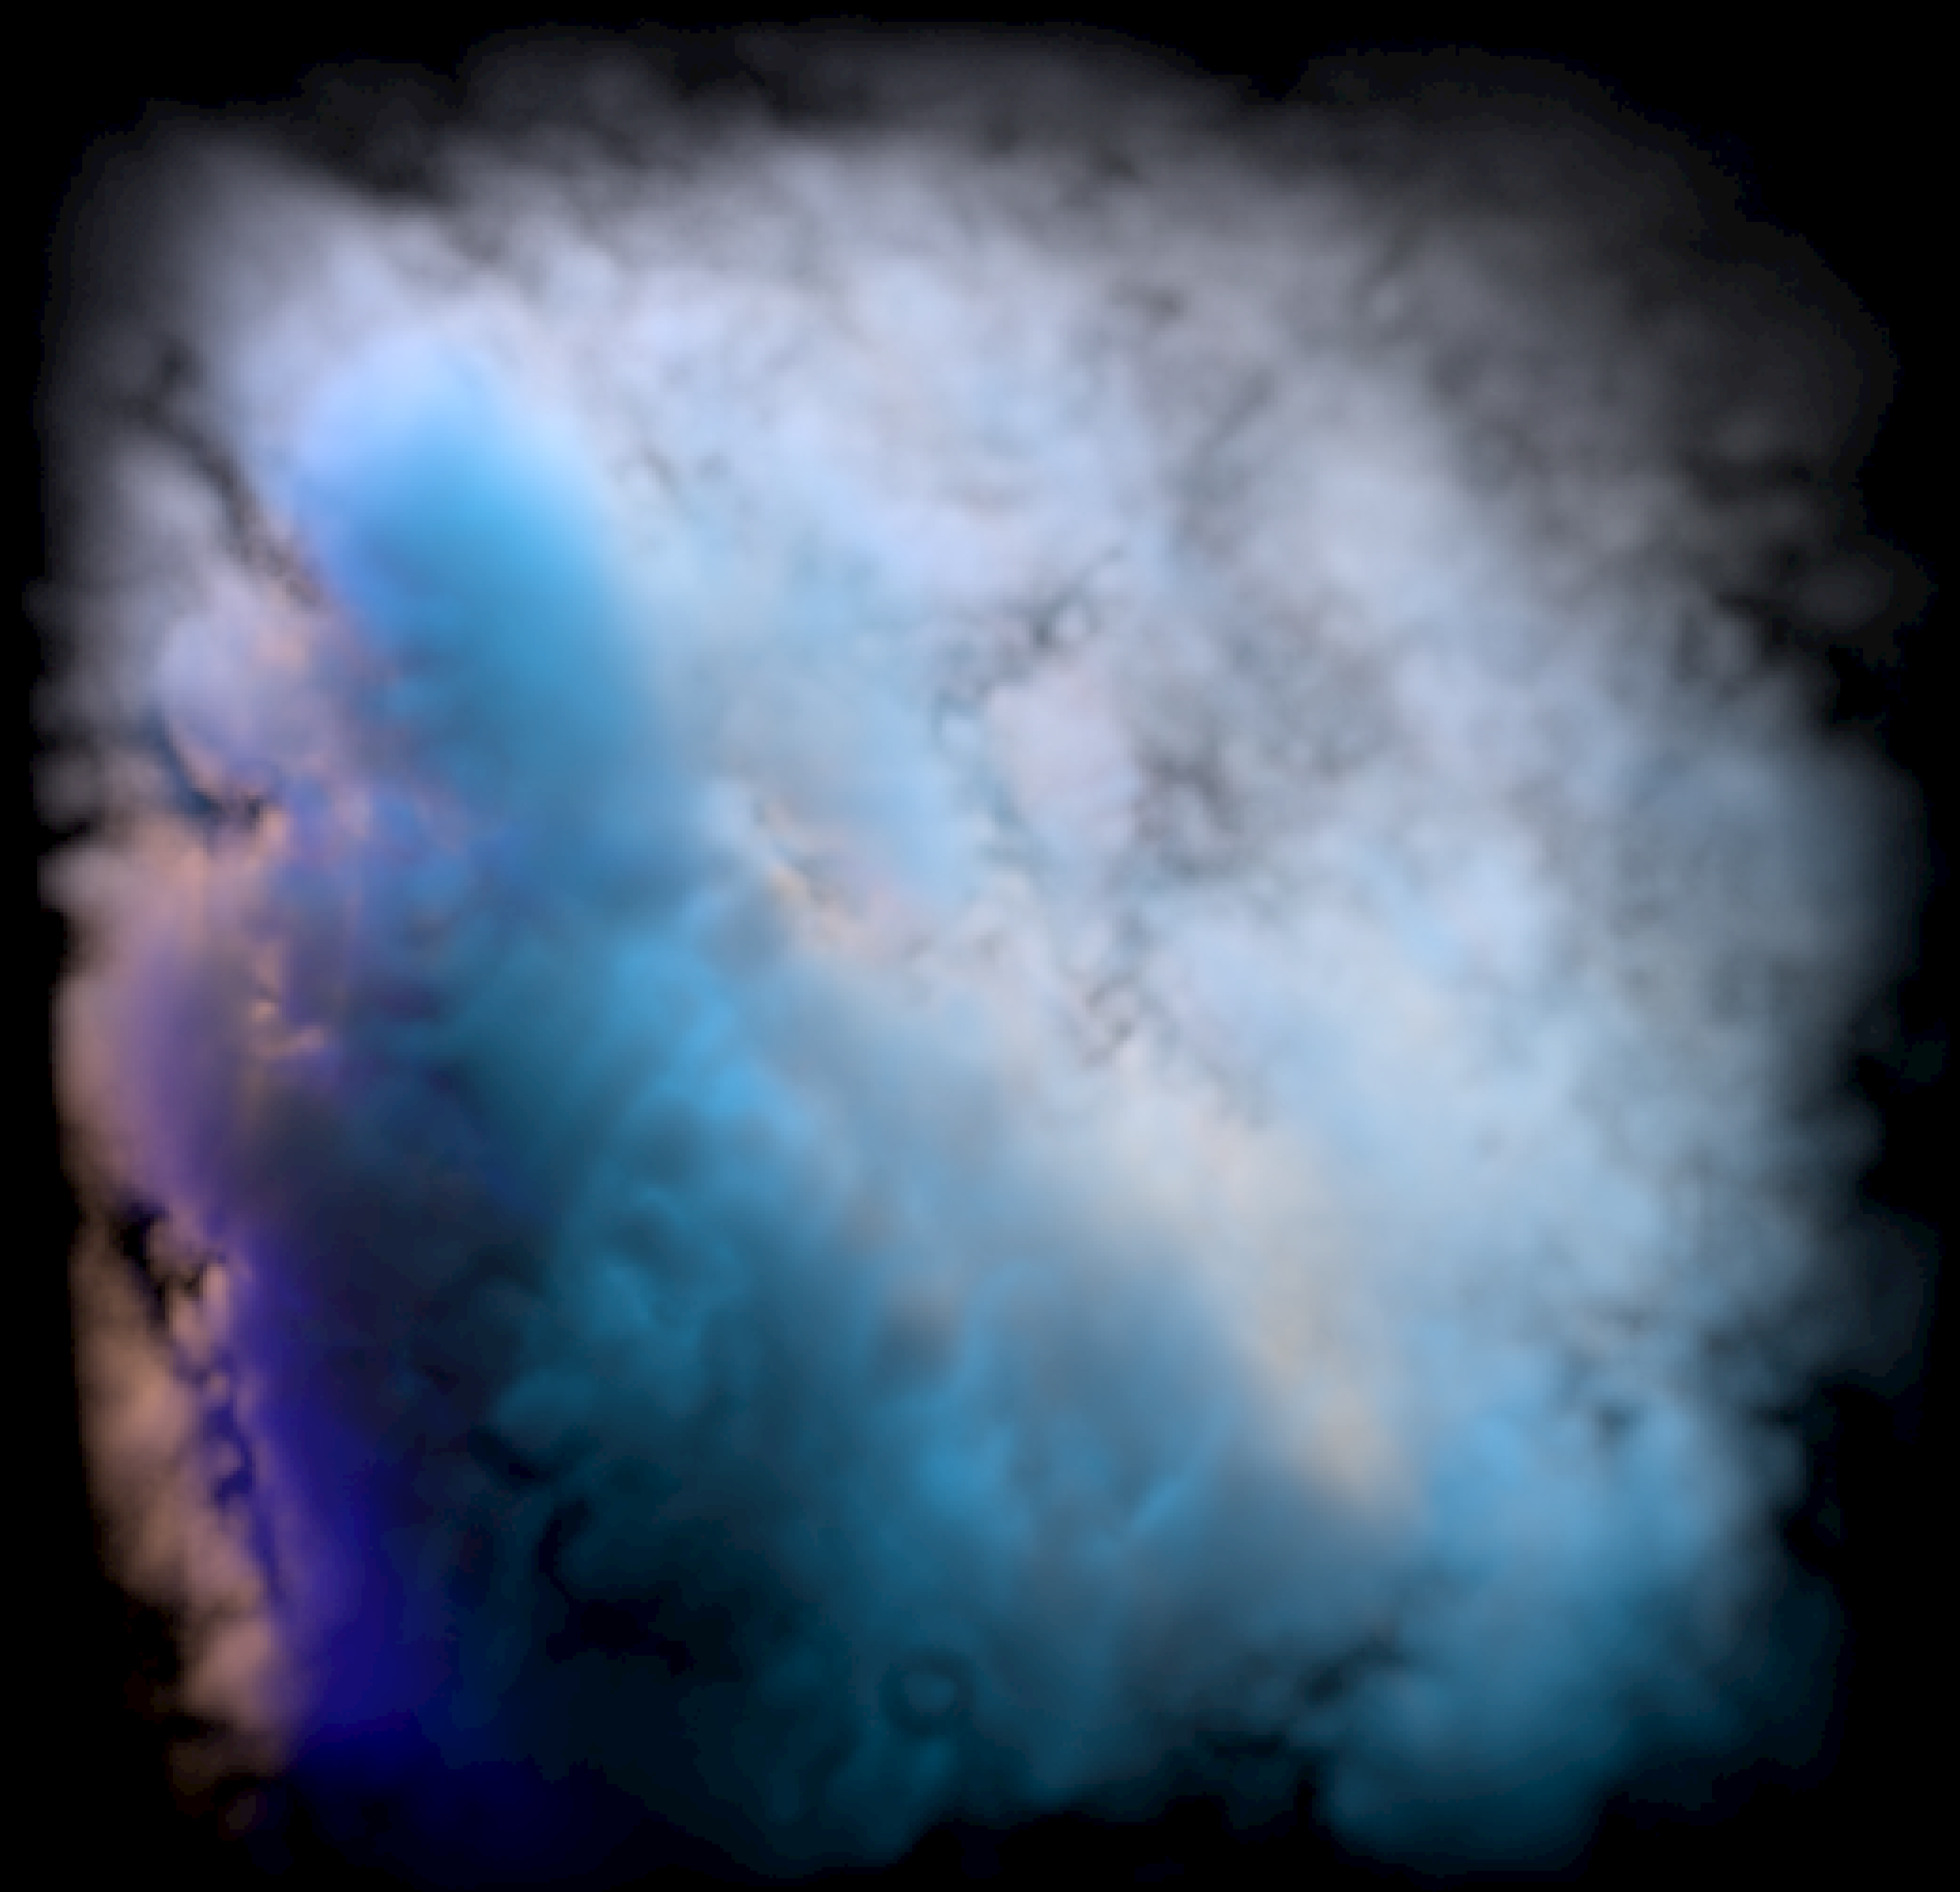

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😢 storage / memory

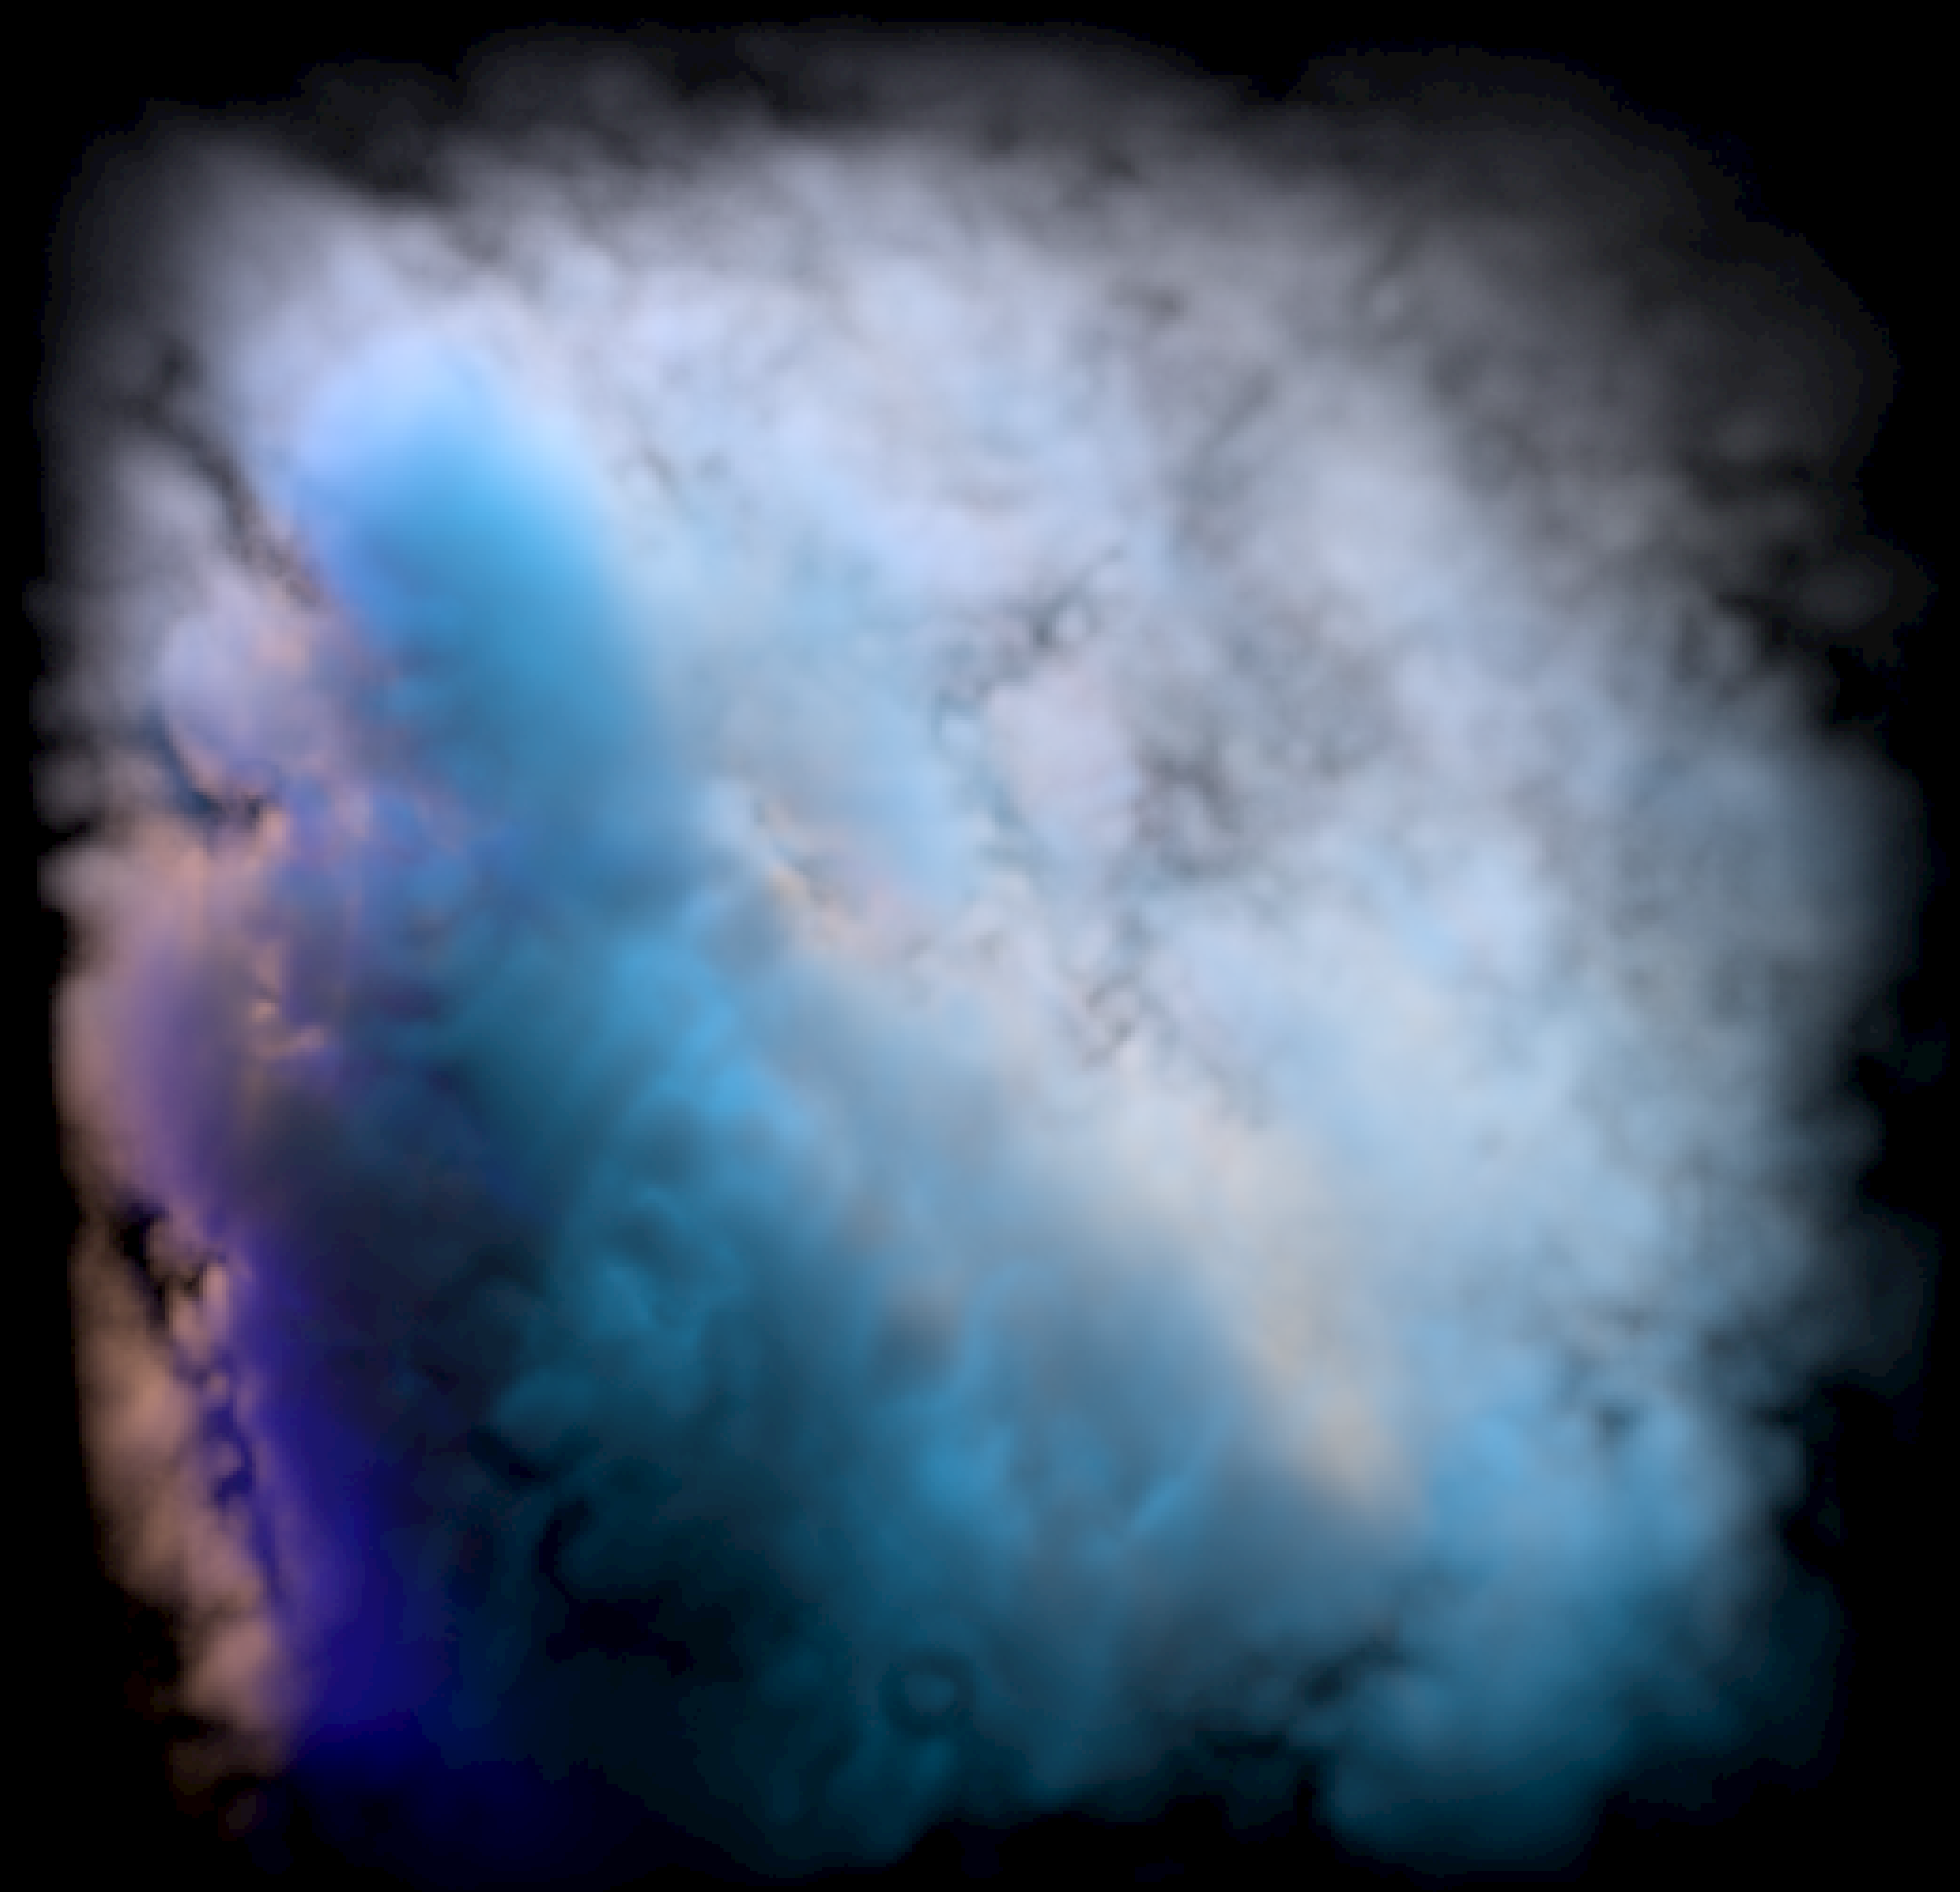

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

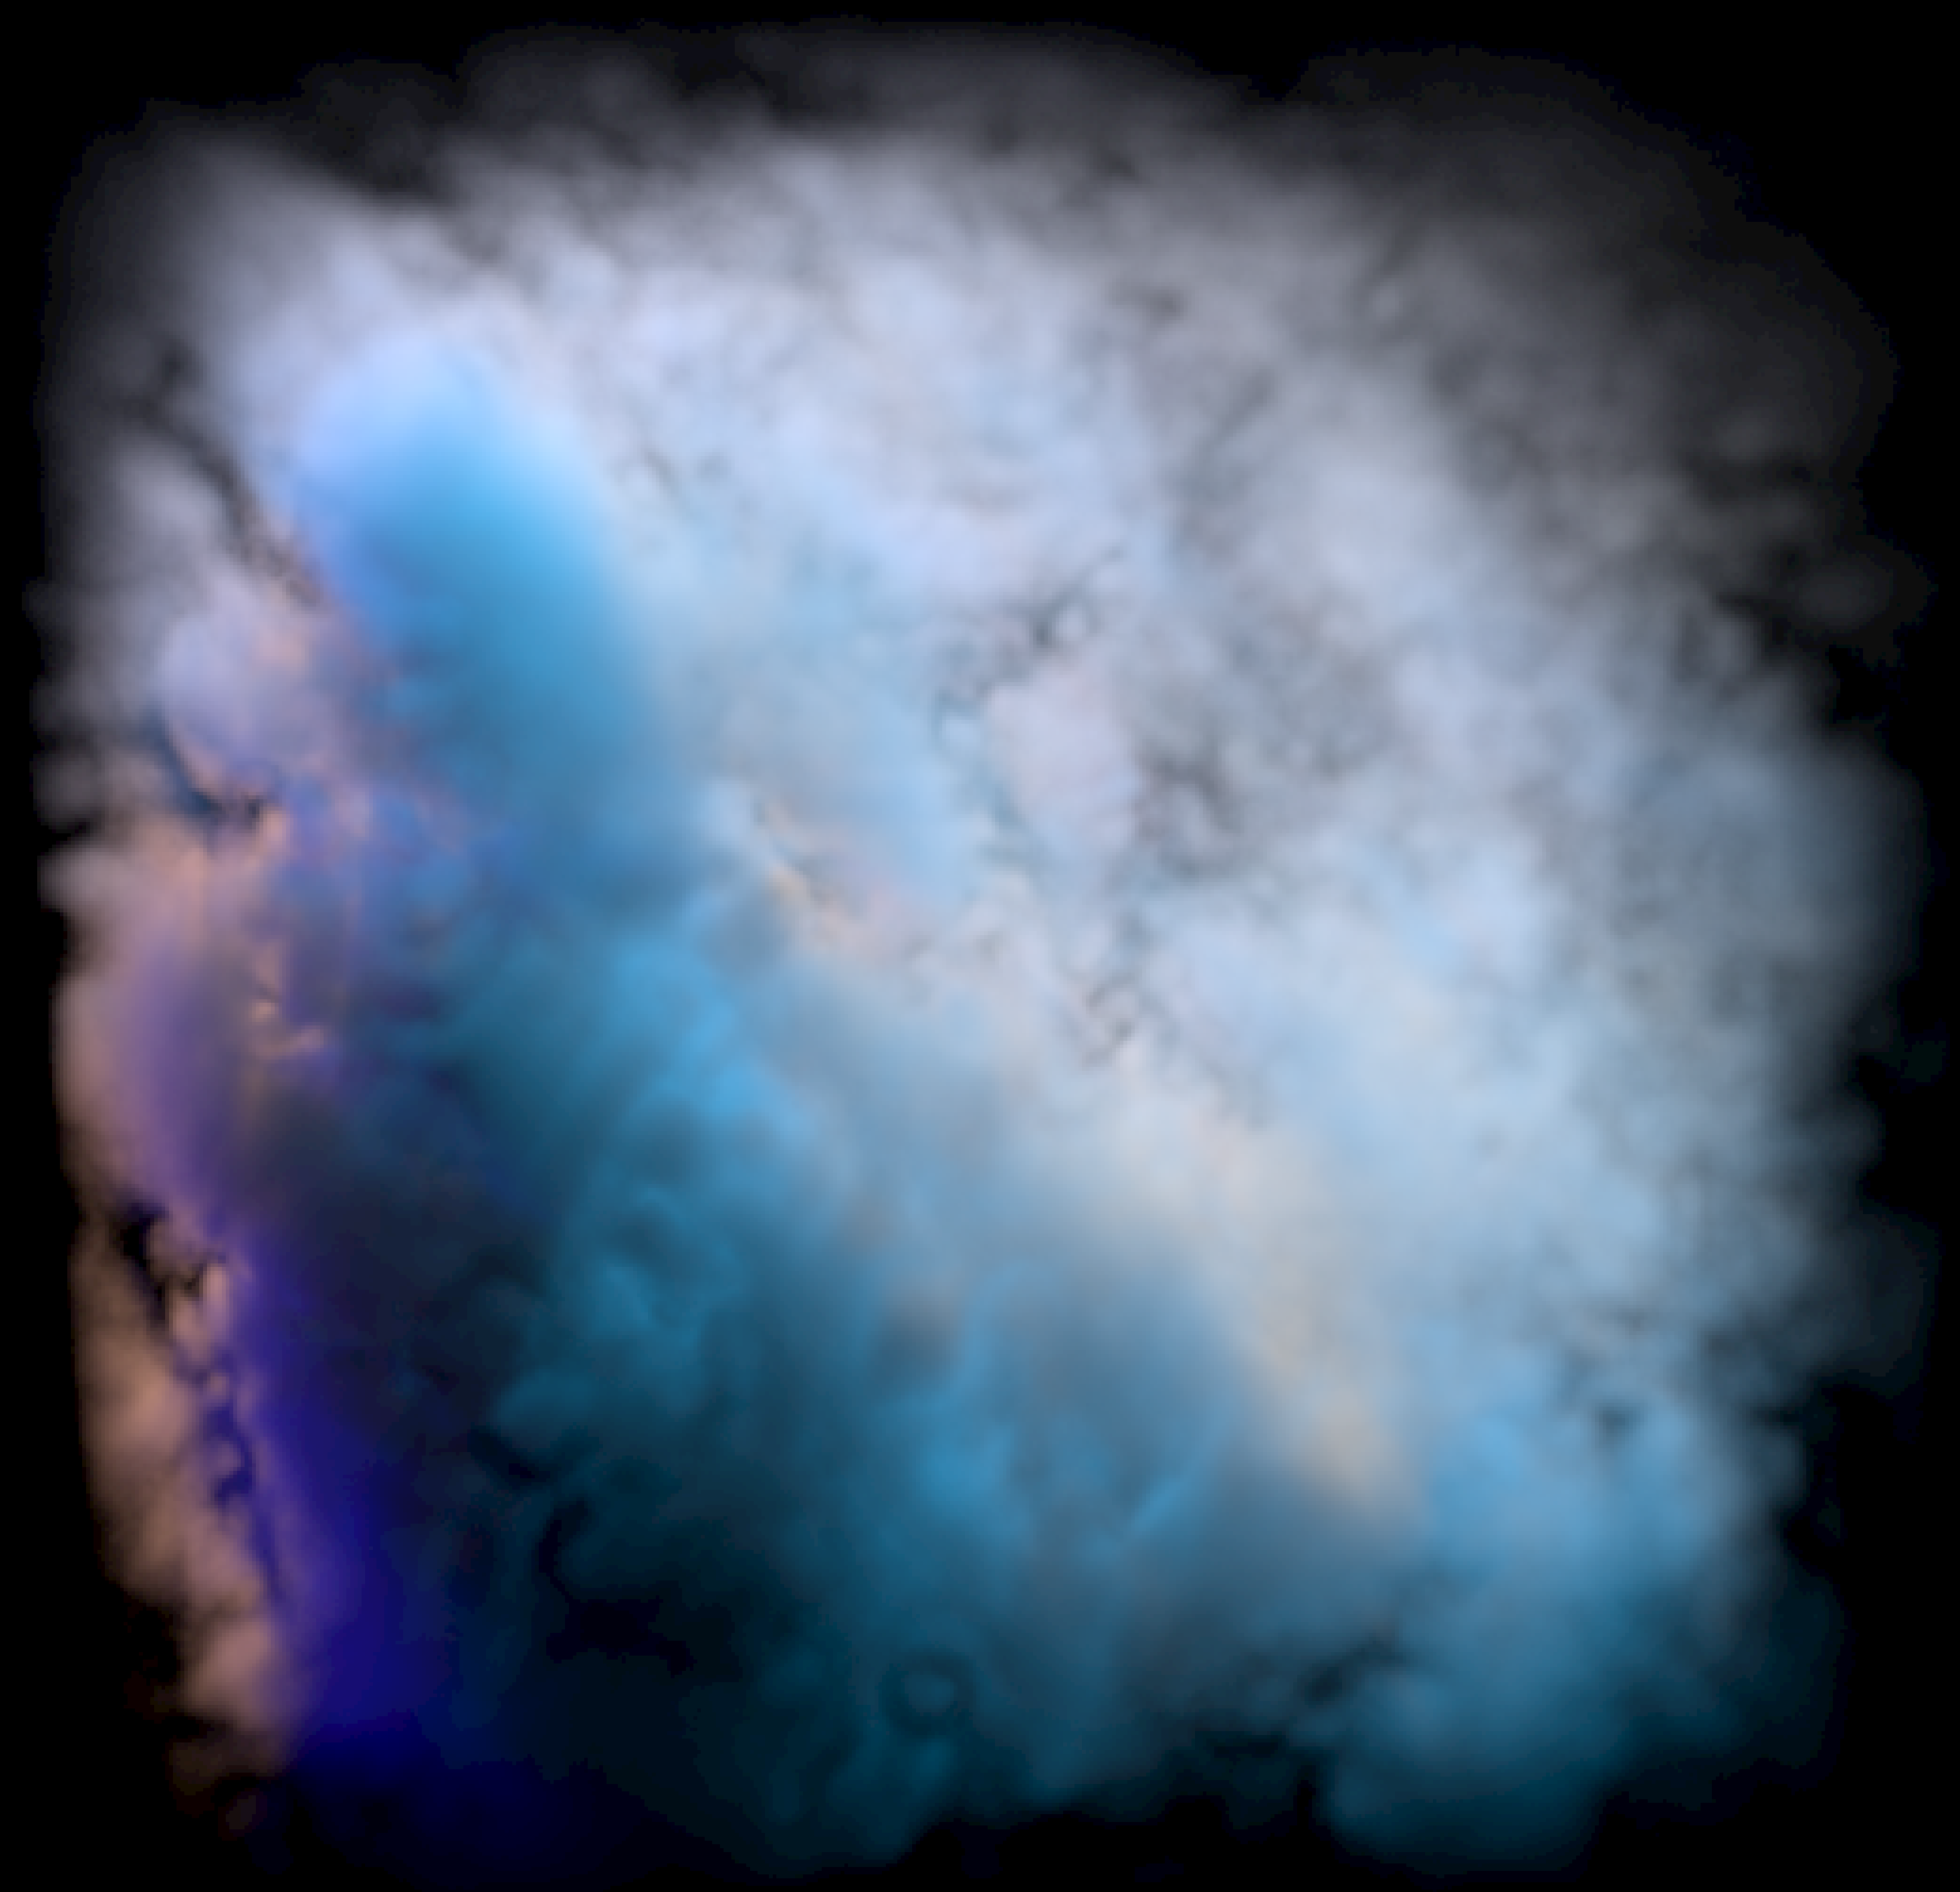

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

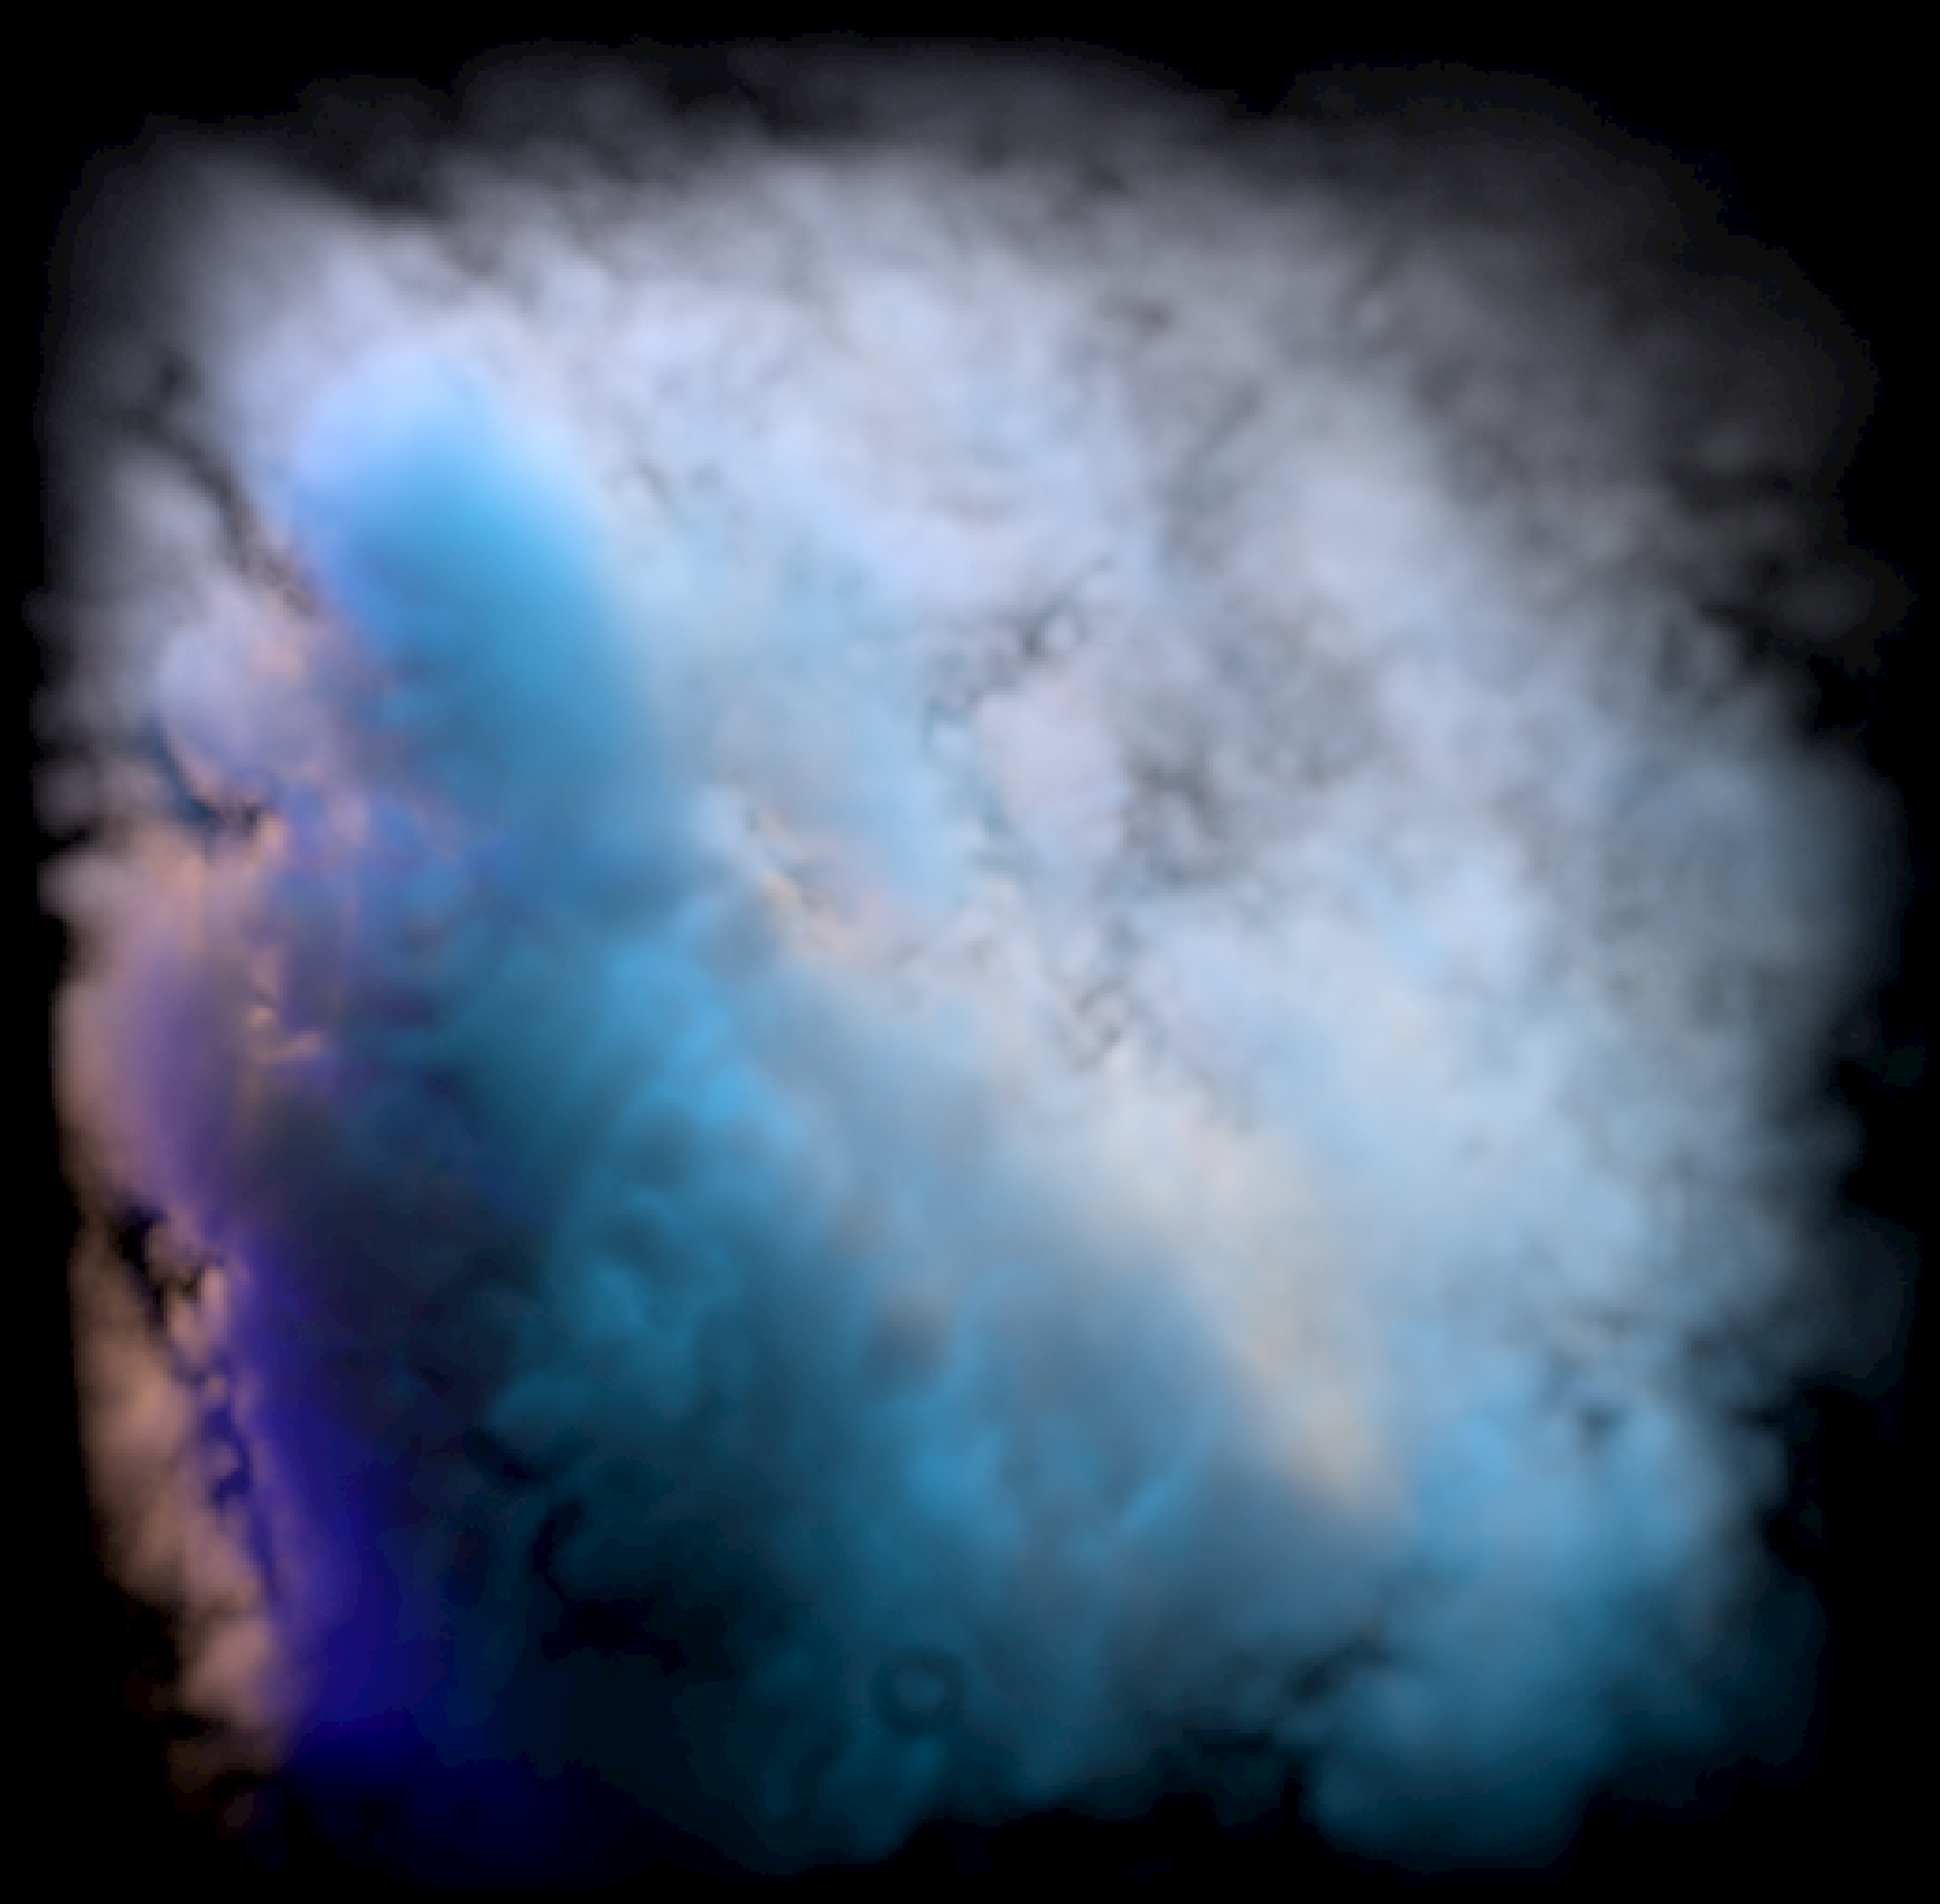

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

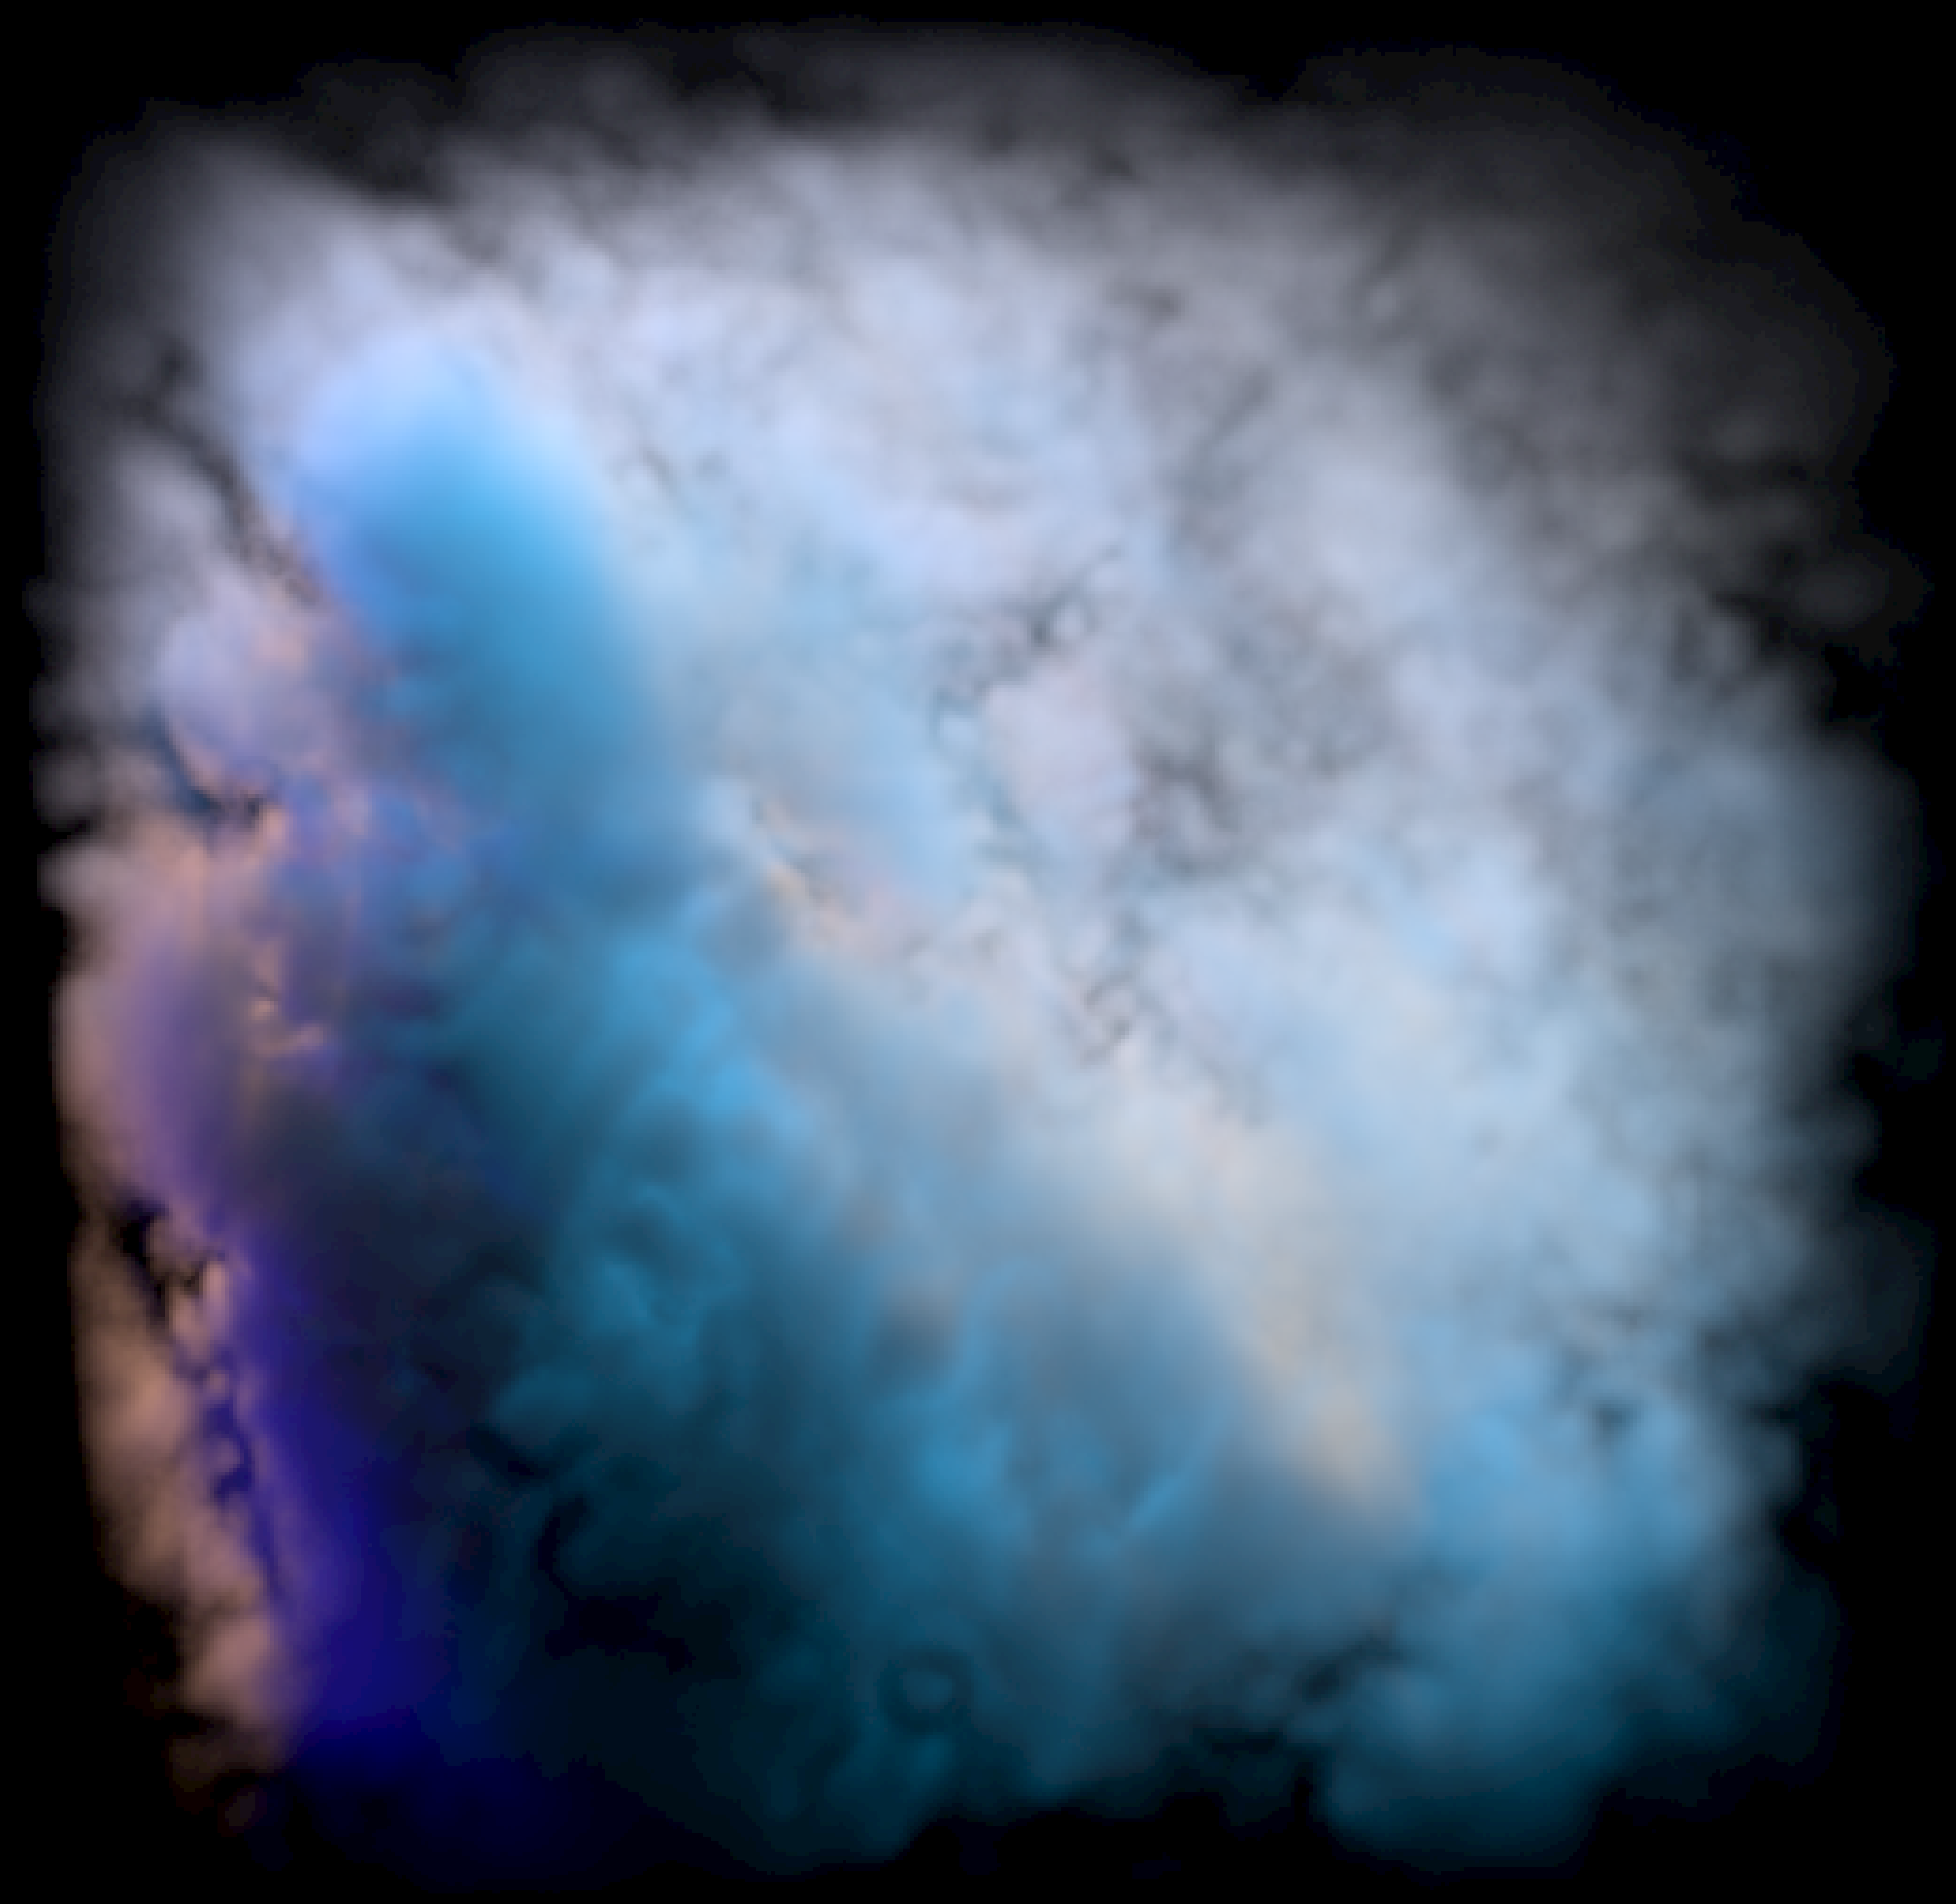

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

😄 unbiased

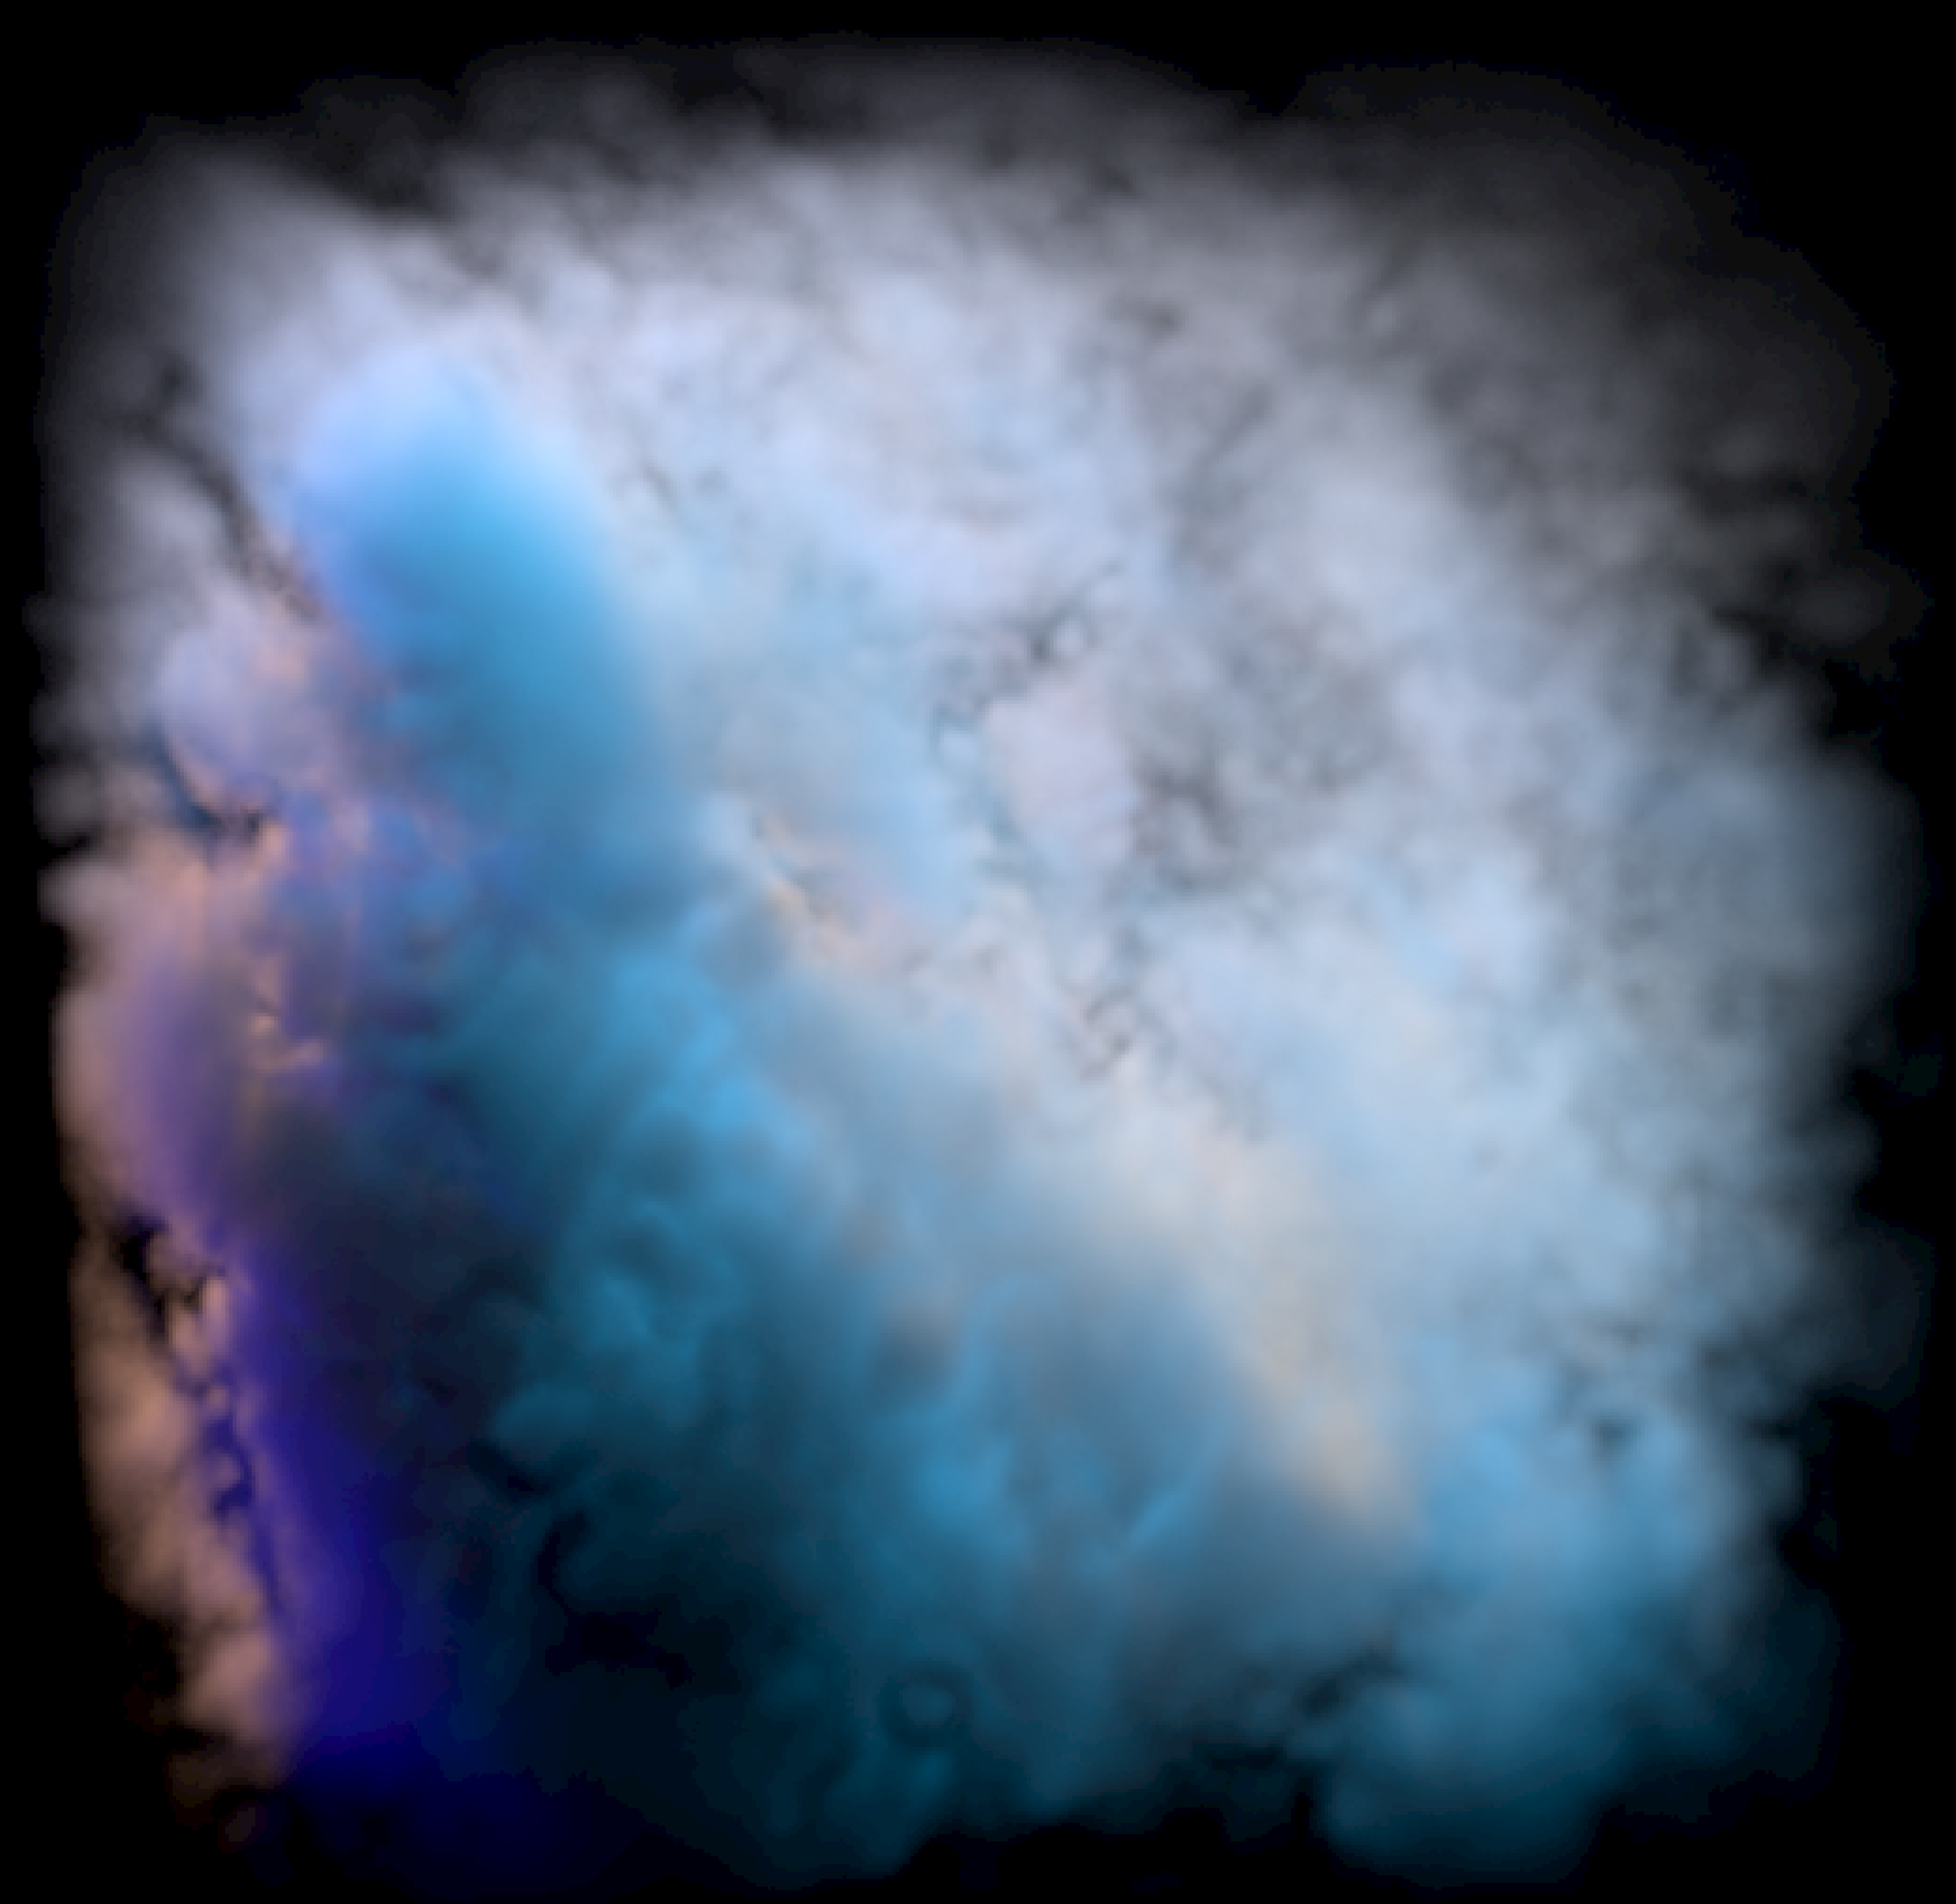

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

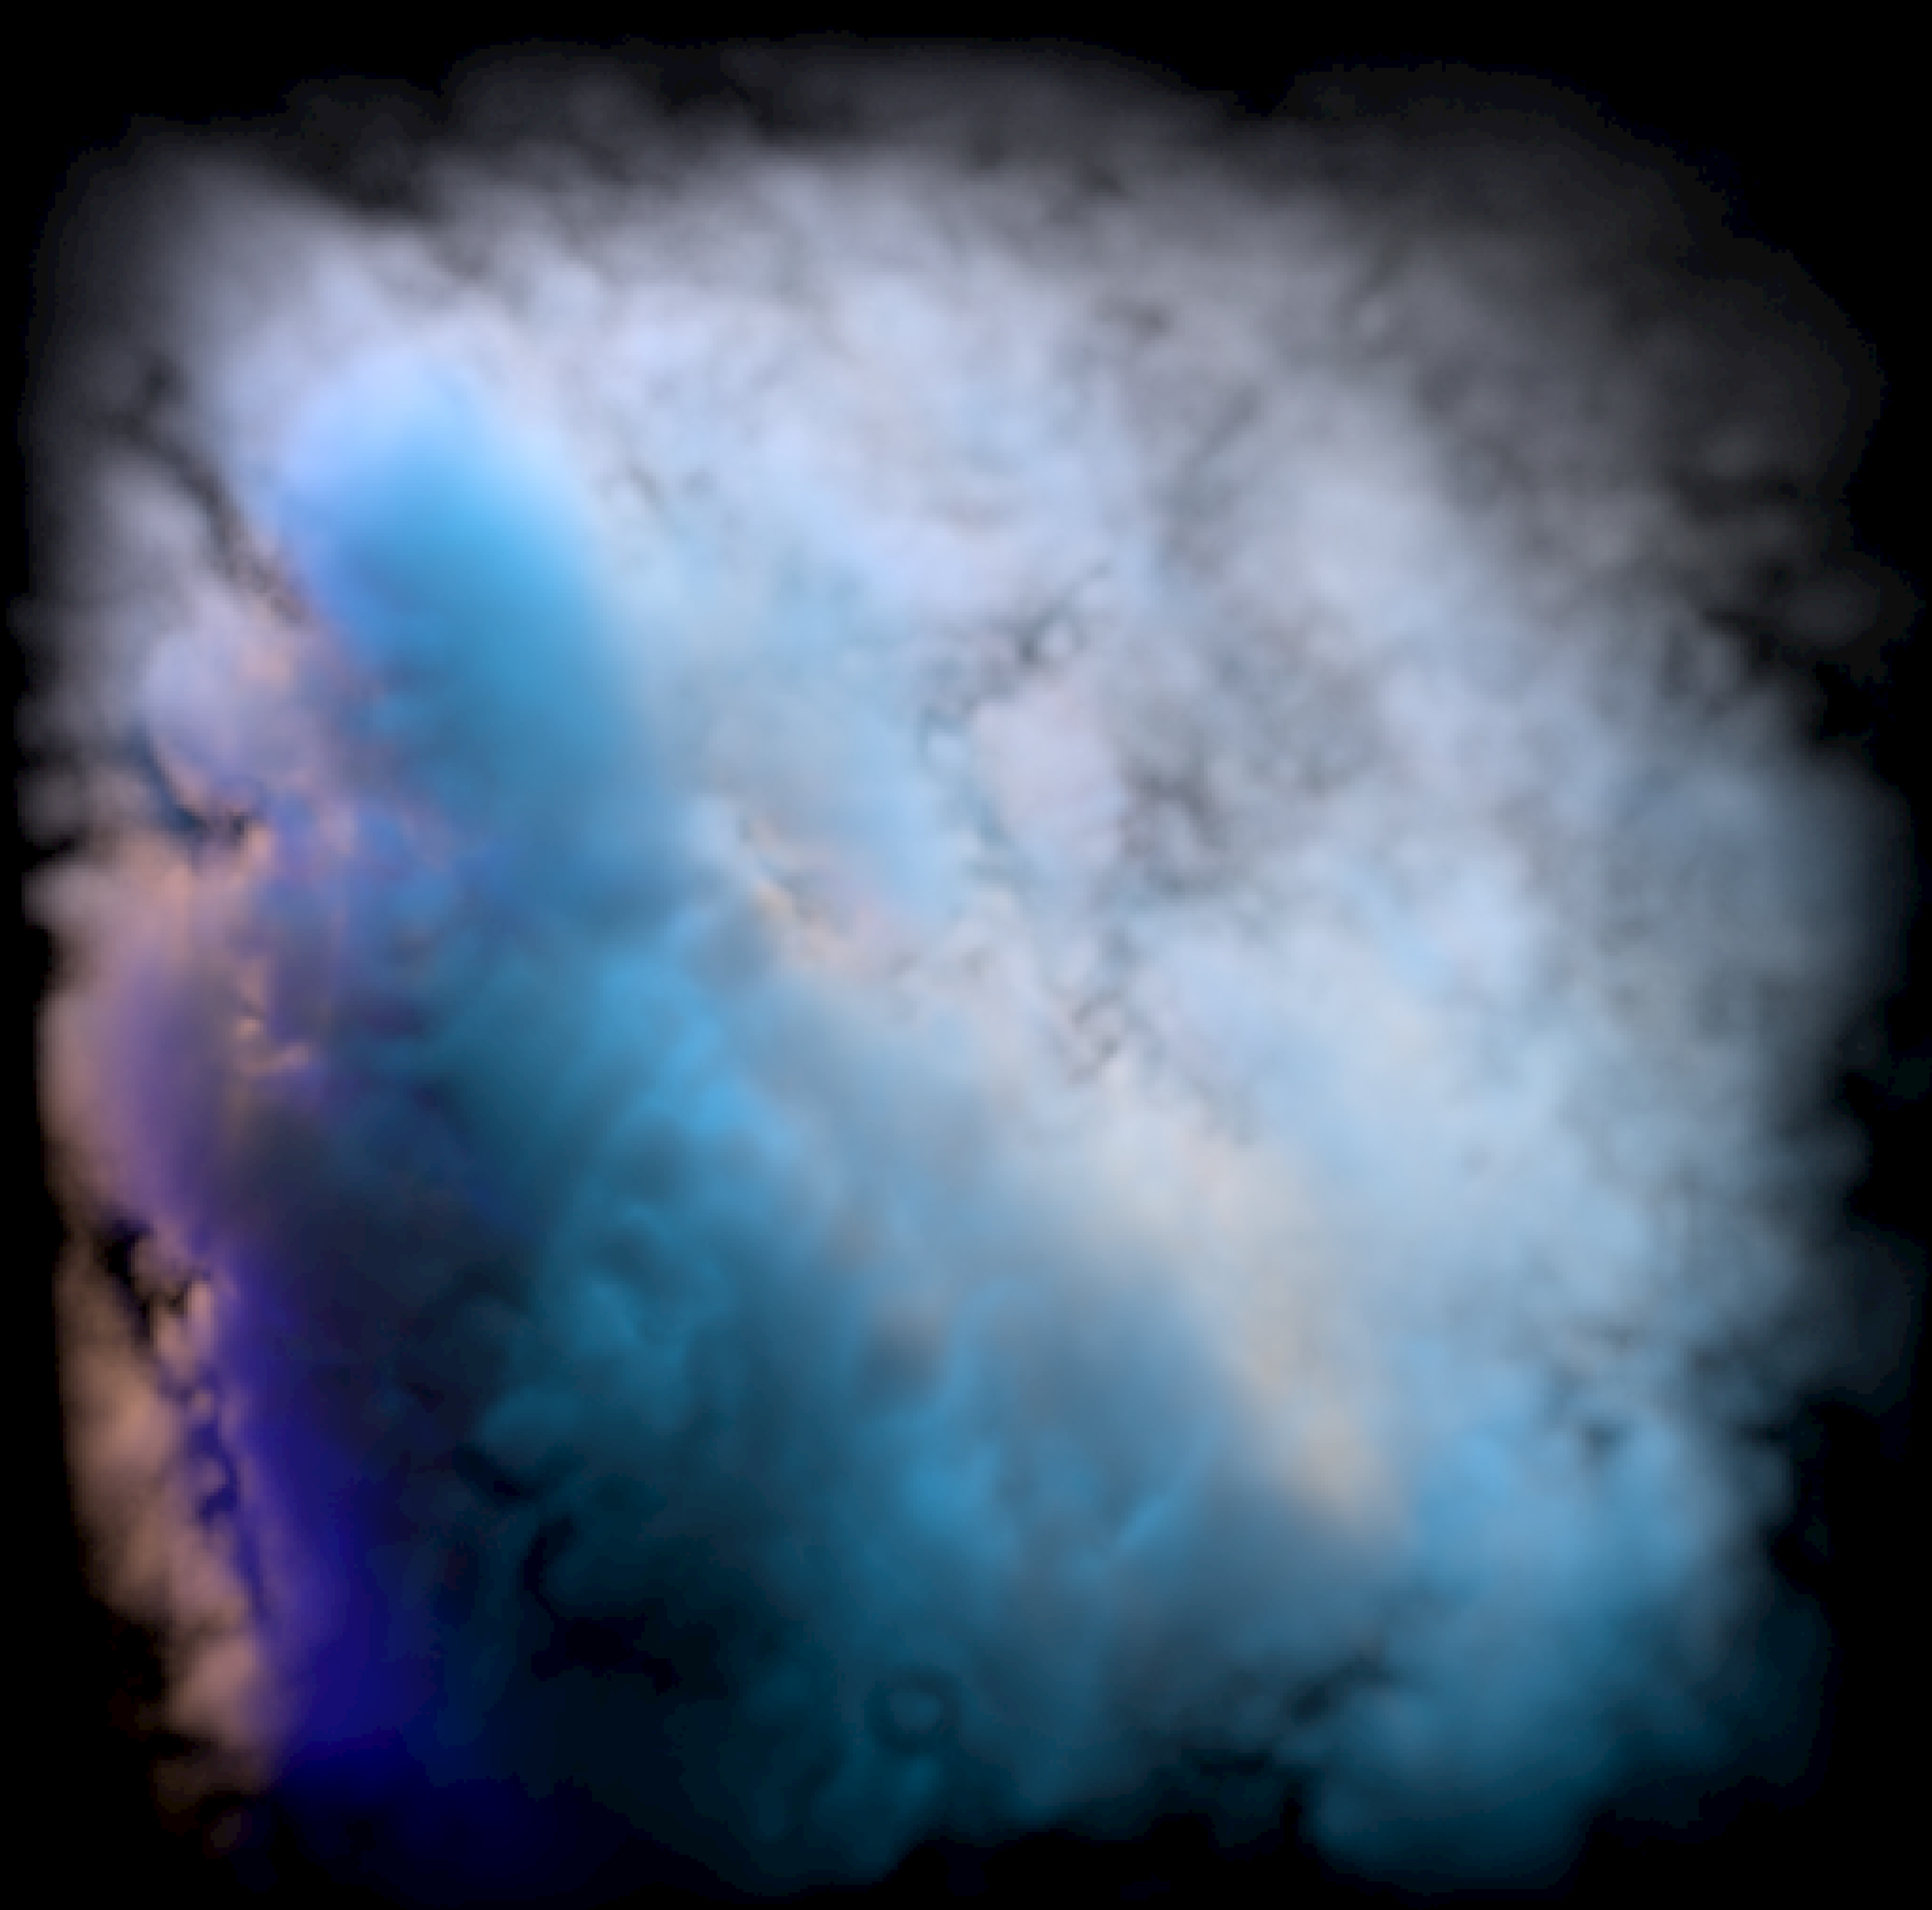

👍 tight 🐜 👍 robust

😄 works for any majorant

😄 visual quality

😄 preprocessing time

😄 storage / memory

😐 consistent

Consistency

$$I_1$$

I_1

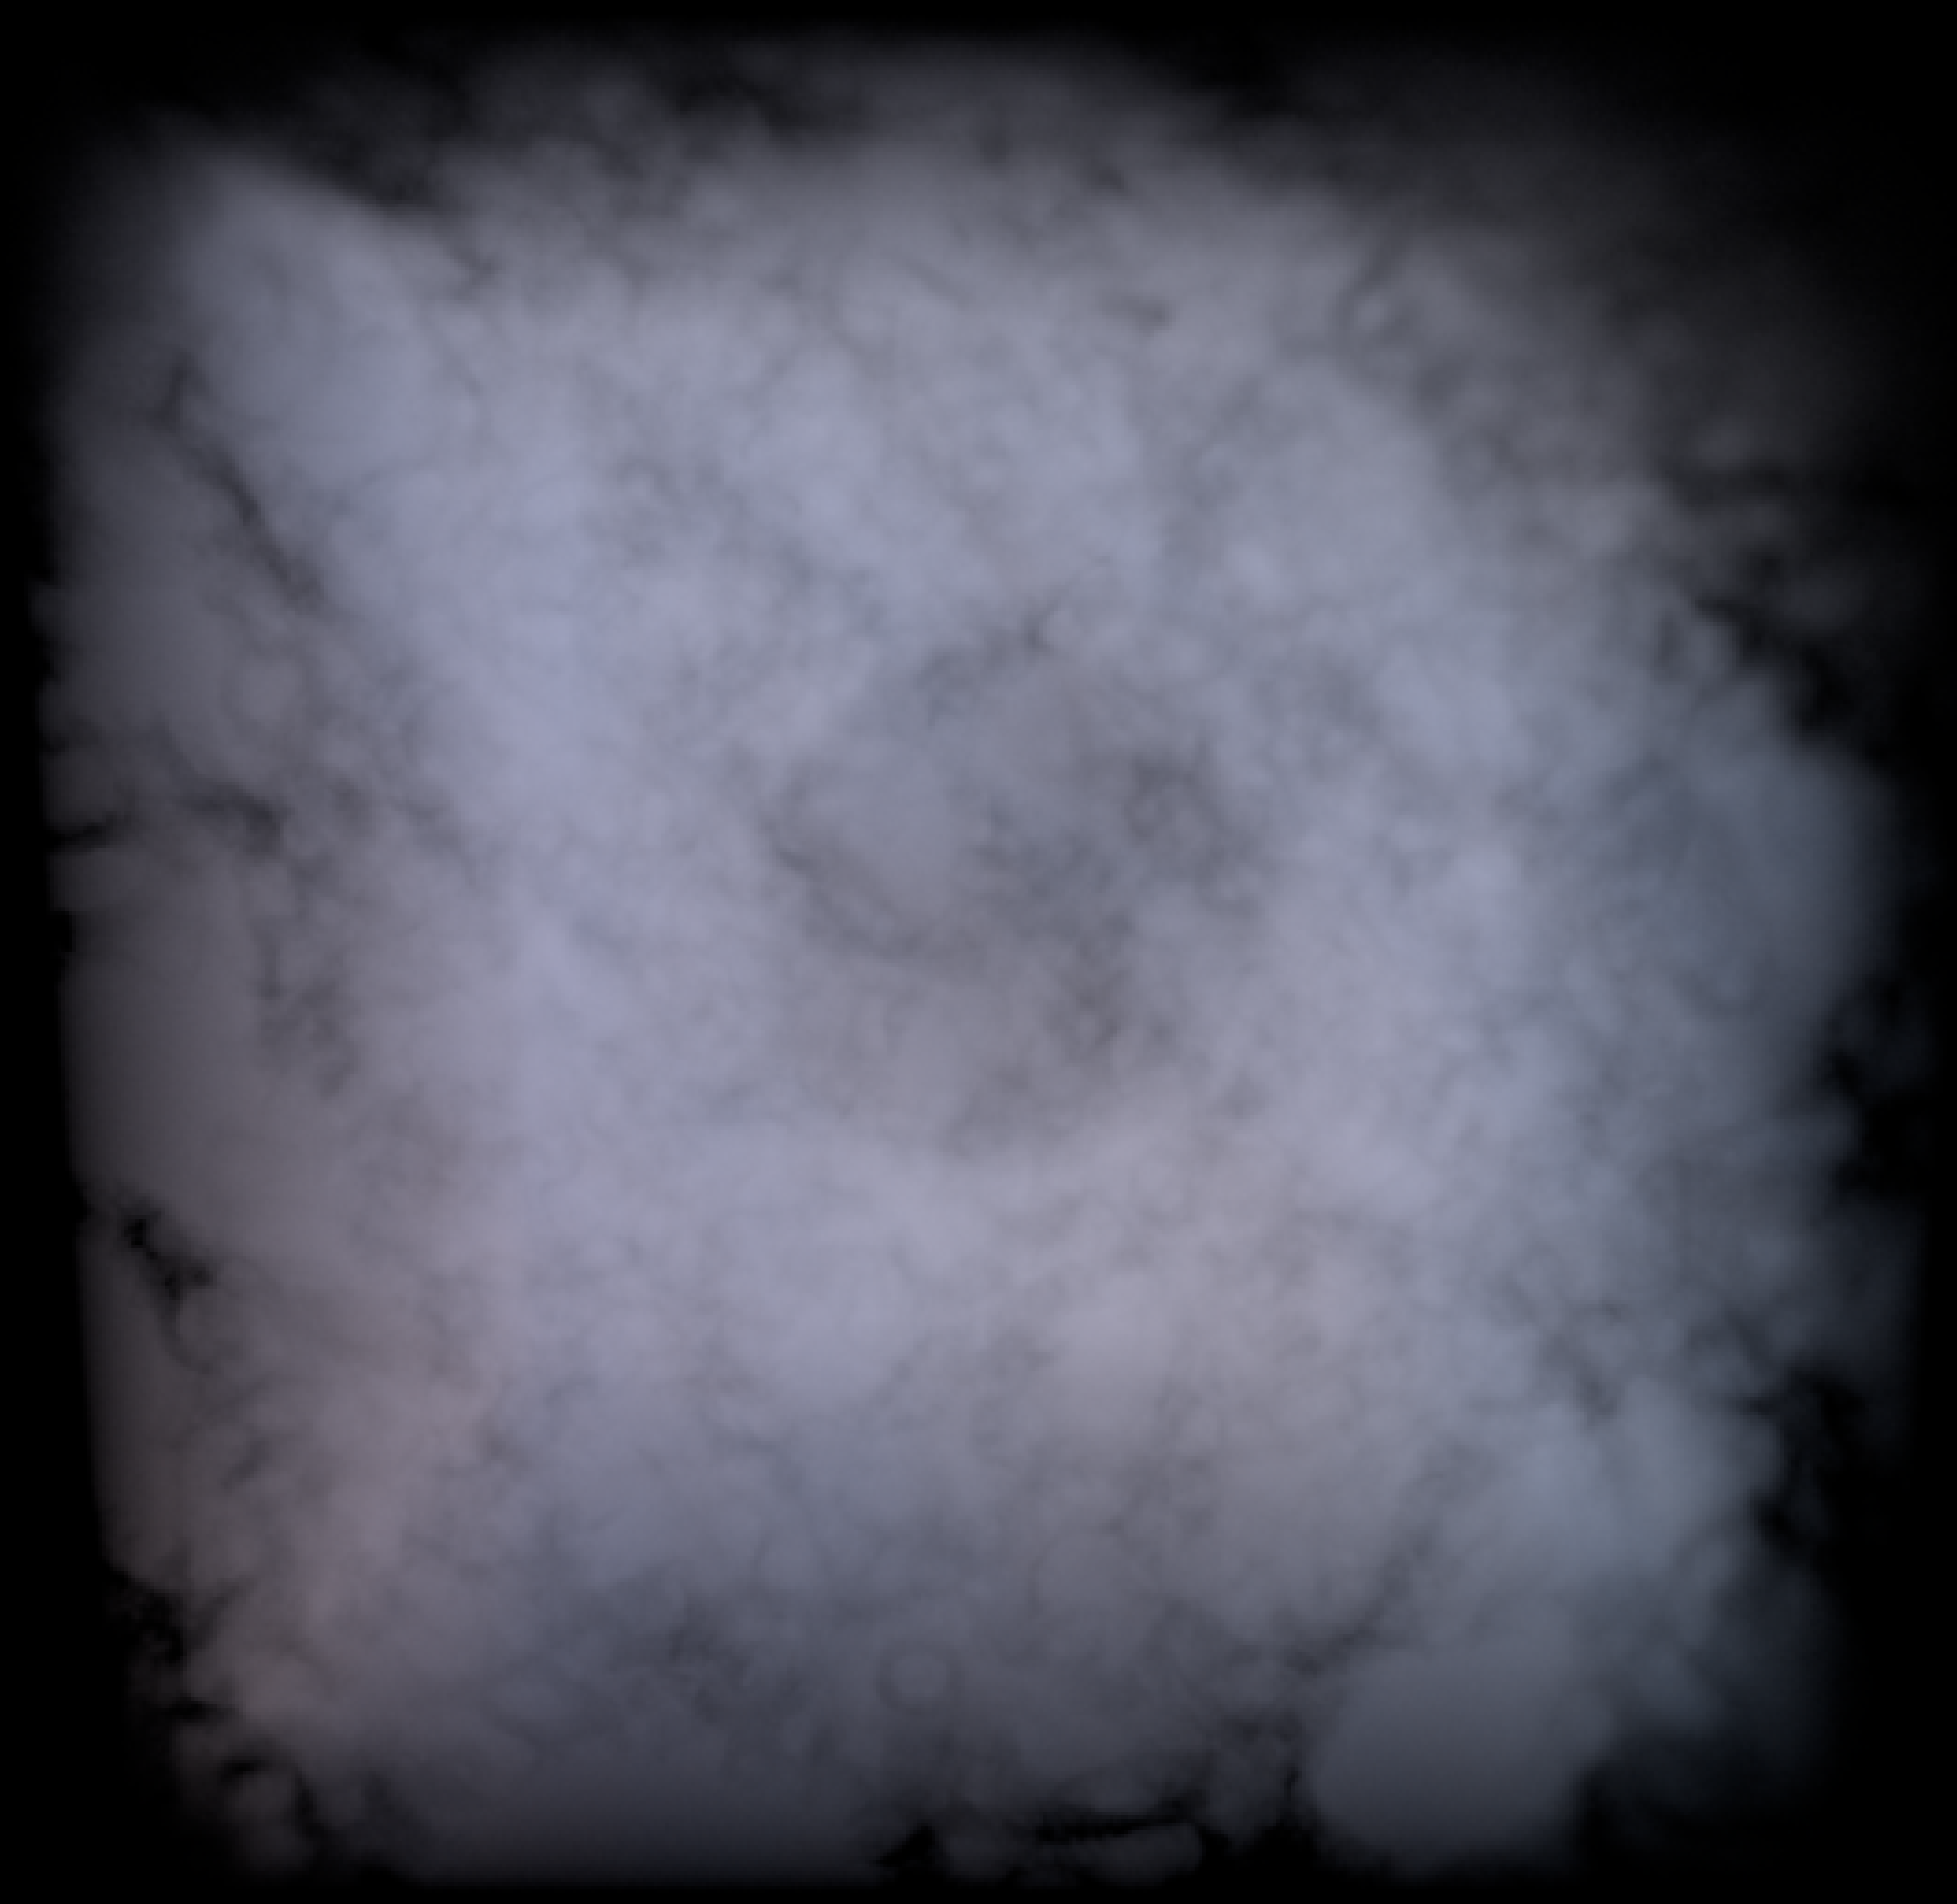

$$\frac{1}{n}\sum_{k=1}^{n} I_k$$

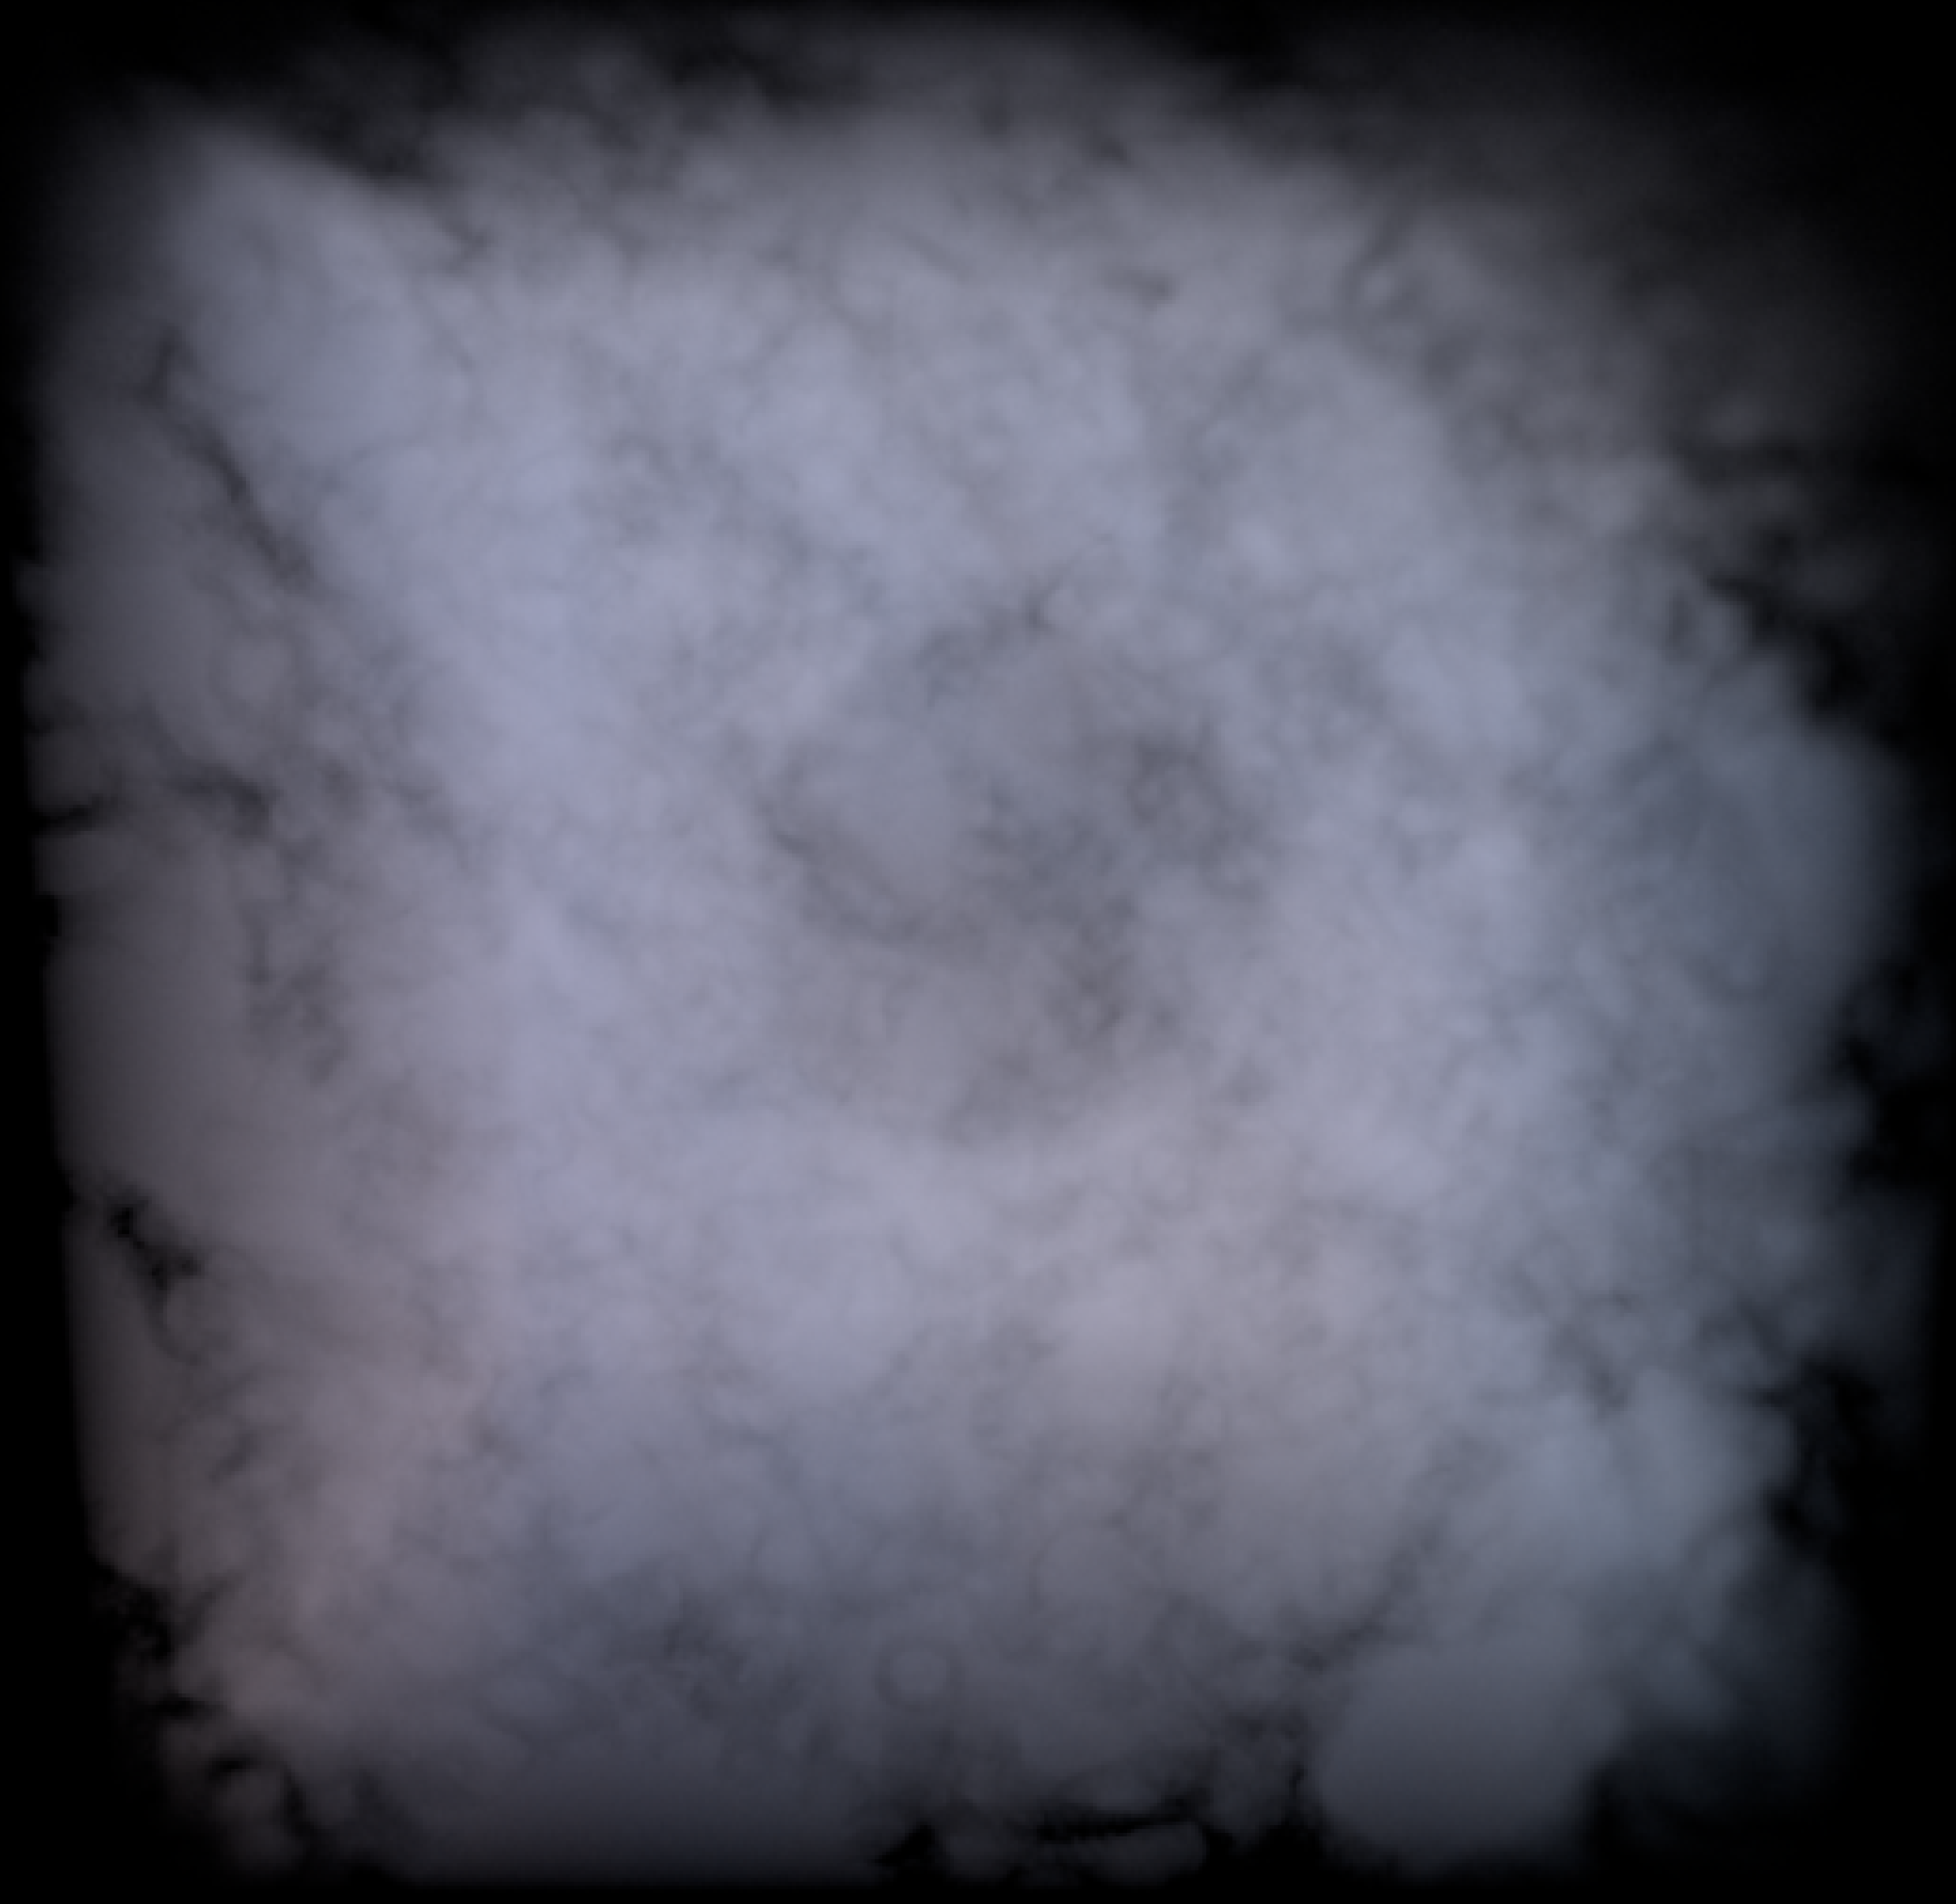

$$\frac{1}{n}\sum_{k=1}^{n} I_k$$

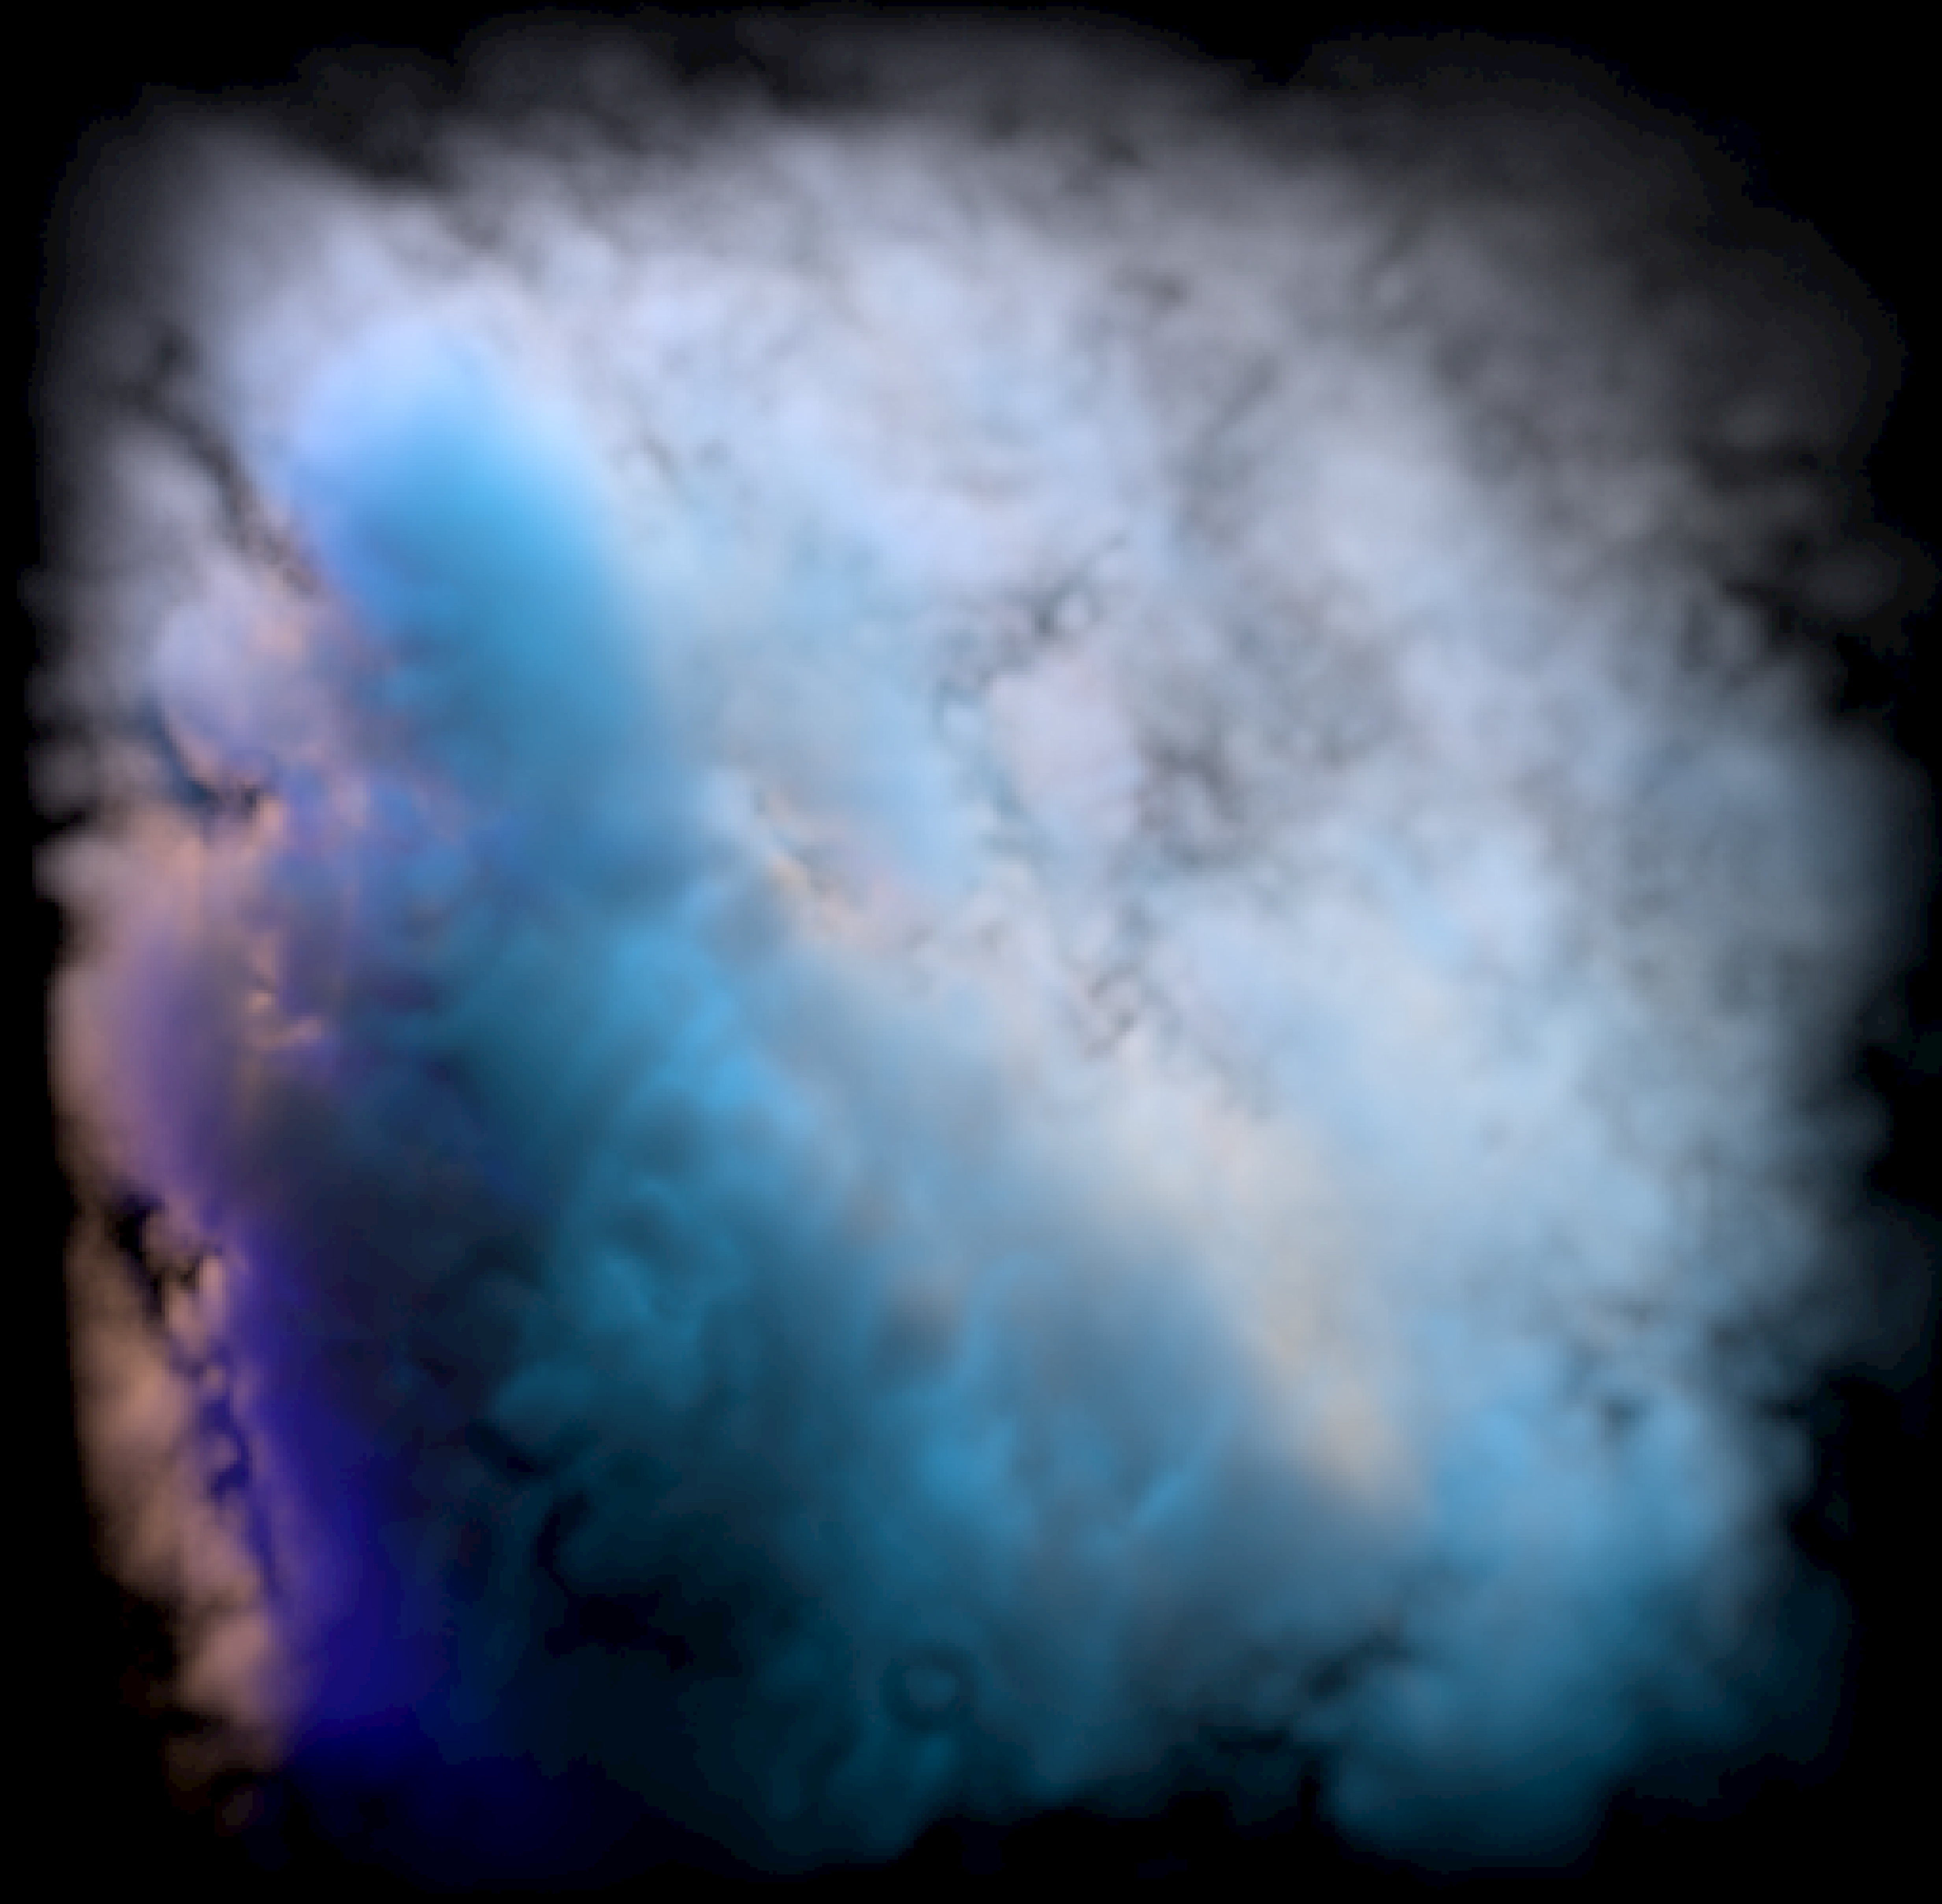

$$\lim_{n\to\infty} \frac{1}{n} \sum_{k=1}^{n} I_k$$

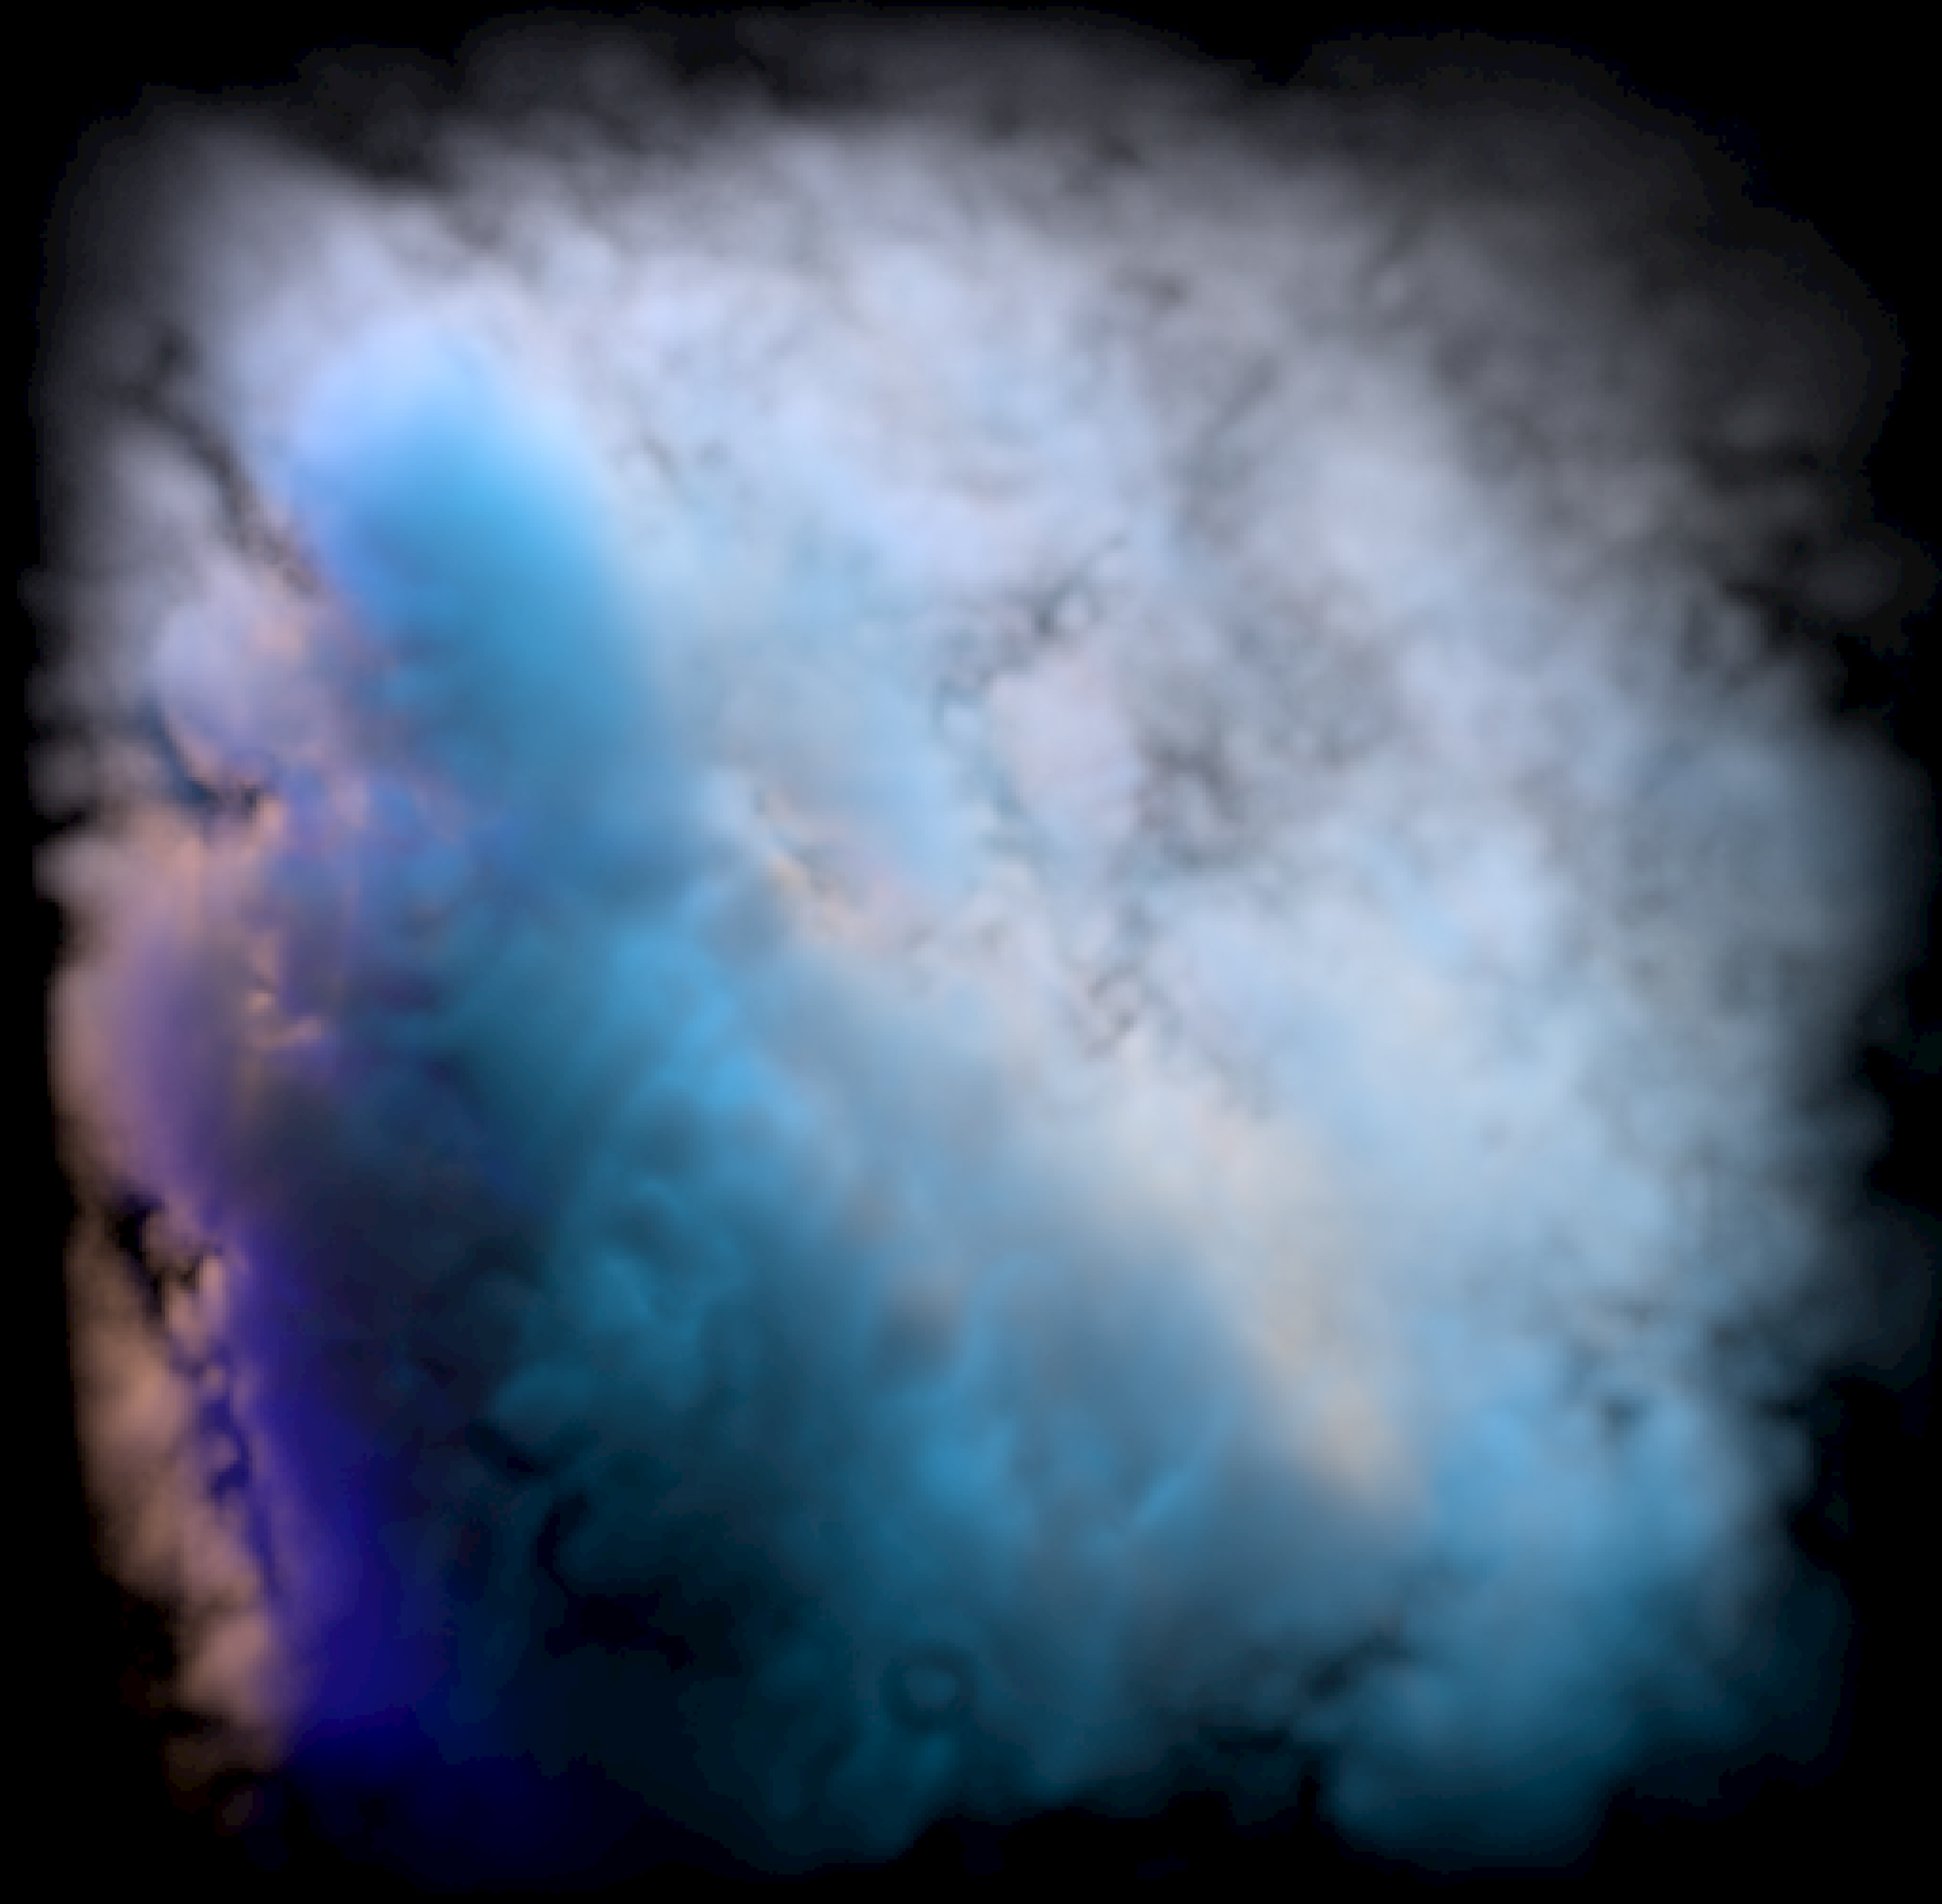

Unbiased and consistent rendering using biased estimators

ZACKARY MISSO, Dartmouth College, USA
BENEDIKT BITTERLI, Dartmouth College, USA and NVIDIA, USA
ILIYAN GEORGIEV, Autodesk, United Kingdom
WOJCIECH JAROSZ, Dartmouth College, USA

We introduce a general framework for transforming biased estimators into unbiased and consistent estimators for the same quantity. We show how several existing unbiased and consistent estimation strategies in rendering are special cases of this framework, and are part of a broader debiasing principle. We provide a recipe for constructing estimators using our generalized framework and demonstrate its applicability by developing novel unbiased forms of transmittance estimation, photon mapping, and finite differences.

CCS Concepts: • **Computing methodologies** → **Rendering**; **Ray tracing**.

Additional Key Words and Phrases: Monte Carlo, infinite series, Taylor series

ACM Reference Format:
Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and consistent rendering using biased estimators. *ACM Trans. Graph.* 41, 4, Article 48 (July 2022), 13 pages. https://doi.org/10.1145/3528223.3530160

1 INTRODUCTION

From estimating the amount of radiance reaching a camera sensor, to estimating how much light transmits through a participating medium, there are countless situations in graphics which require estimating intricate integrals. While we have developed a large arsenal of unbiased estimation techniques, situations still arise where we must fall back on biased formulations.

We consider problems where we need to compute some finite quantity I, but we only have a biased estimator $\langle I(k)\rangle$ with a *controllable* amount of bias—dictated by some parameter k—at our disposal. By adjusting the bias parameter towards some limit (e.g. $k \to \infty$) the estimator's expected value $I(k)$ approaches the correct answer:

$$I = \lim_{k\to\infty} I(k). \tag{1}$$

The bias parameter k could be continuous or discrete; for example, a discrete k could represent the maximum path length in a path tracer, while a continuous k could correspond to the step size in ray

Progressive Photon Mapping

Toshiya Hachisuka, UC San Diego
Shinji Ogaki, The University of Nottingham
Henrik Wann Jensen, UC San Diego

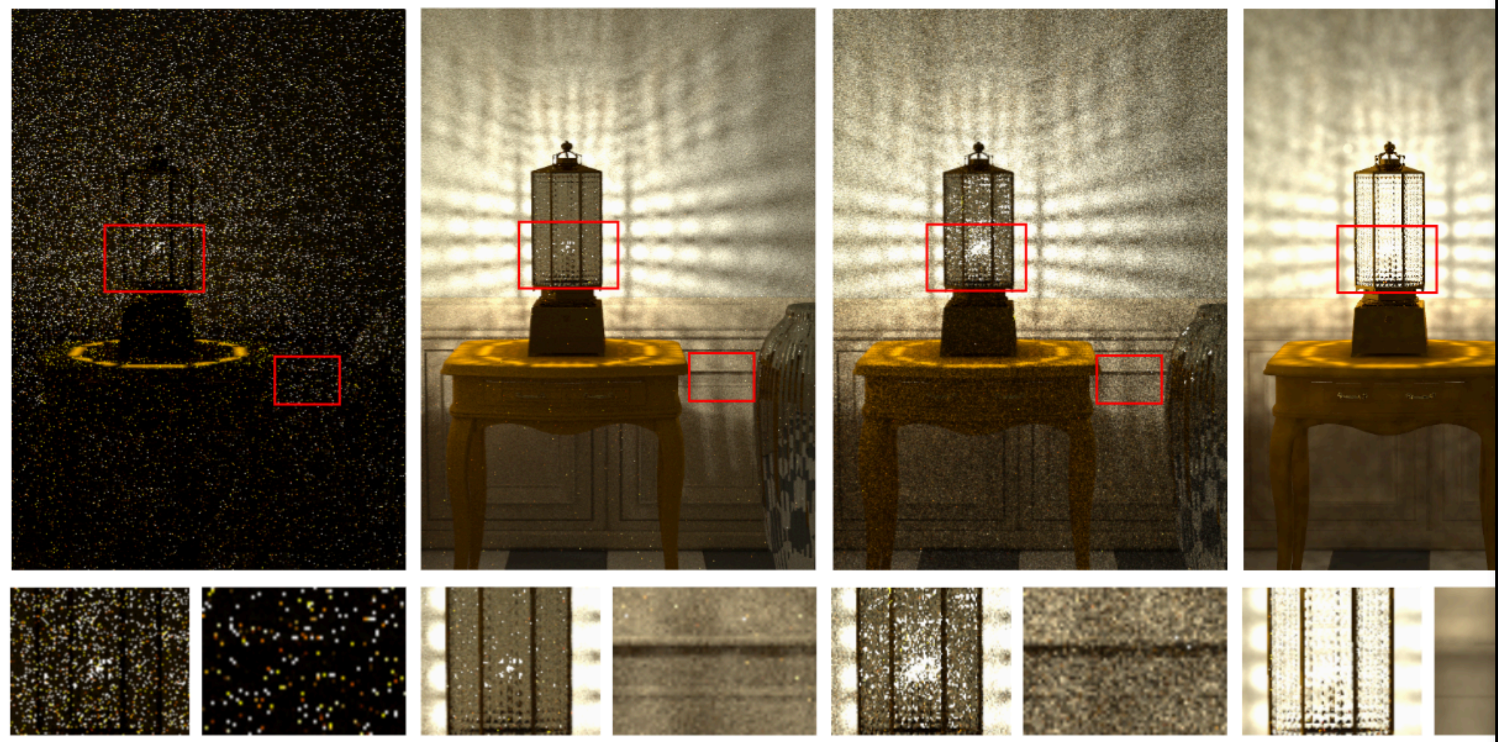

Figure 1: *A glass lamp illuminates a wall and generates a complex caustics lighting pattern on the w… to simulate with Monte Carlo ray tracing methods such as path tracing, bidirectional path tracing… lighting seen through the lamp is particularly difficult for these methods. Photon mapping is signif… lighting seen through the lamp, but the final quality is limited by the memory available for the photo… illumination. Progressive photon mapping provides an image with substantially less noise in the sam… tracing methods and the final quality is not limited by the available memory.*

Abstract

This paper introduces a simple and robust progressive global illumination algorithm based on photon mapping. Progressive photon mapping is a multi-pass algorithm where the first pass is ray tracing followed by any number of photon tracing passes. Each photon tracing pass results in an increasingly accurate global illumination solution that can be visualized in order to provide progressive feedback. Progressive photon mapping uses a new radiance estimate that converges to the correct radiance value as more photons are used. It is not necessary to store the full photon map, and unlike standard photon mapping it possible to compute a global illumination solution with any desired accuracy using a limited amount of memory. Compared with existing Monte Carlo ray tracing methods progressive photon mapping provides an efficient and robust alternative in the presence of complex light transport such as caustics and in particular reflections of caustics.

1 Introduction

Efficiently simulating globa… lems in computer graphics. … solving for all types of light … vide a full solution to the ren… global illumination algorith… and a number of algorithms … capable of solving the rende… tions [Dutré et al. 2006].

Monte Carlo based methods … materials, but there is one … particularly problematic for … tion involves light being tra… specular path (SDS path) be… ple of an SDS path is the sh…

EUROGRAPHICS 2013/ M. Sbert, L. Szirmay-Kalos *STAR – State of The Art Report*

Scalable Realistic Rendering with Many-Light Methods

Carsten Dachsbacher, Karlsruhe Institute of Technology
Jaroslav Křivánek, Charles University, Prague
Miloš Hašan, UC Berkeley
Adam Arbree, Autodesk, Inc.
Bruce Walter, Cornell University
Jan Novák, Karlsruhe Institute of Technology

Figure 1: *Many-light rendering methods, covered in this report, yield good results at different points along the quality-speed trade-off axis. The images on the left were rendered in real-time with [REH*11] (courtesy of Tobias Ritschel) and capture diffuse interreflections. The center image took 52 minutes to render and demonstrates many-light methods for participating media (adapted from [ENSD12]). The image on the right combines different phenomena such as glossy surfaces, subsurface BSSRDFs and a detailed anisotropic volumetric cloth model rendered with Bidirectional Lightcuts [WKB12] in about 46 minutes.*

Abstract
Recent years have seen increasing attention and significant progress in many-light rendering, a class of methods for the efficient computation of global illumination. The many-light formulation offers a unified mathematical framework for the problem reducing the full lighting transport simulation to the calculation of the direct illumination from many virtual light sources. These methods are unrivaled in their scalability: they are able to produce artifact-free images in a fraction of a second but also converge to the full solution over time. In this state-of-the-art report, we have three goals: give an easy-to-follow, introductory tutorial of many-light theory; provide a comprehensive, unified survey of the topic with a comparison of the main algorithms; and present a vision to motivate and guide future research. We will cover both the fundamental concepts as well as improvements, extensions, and applications of many-light rendering.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Clamping to reduce variance

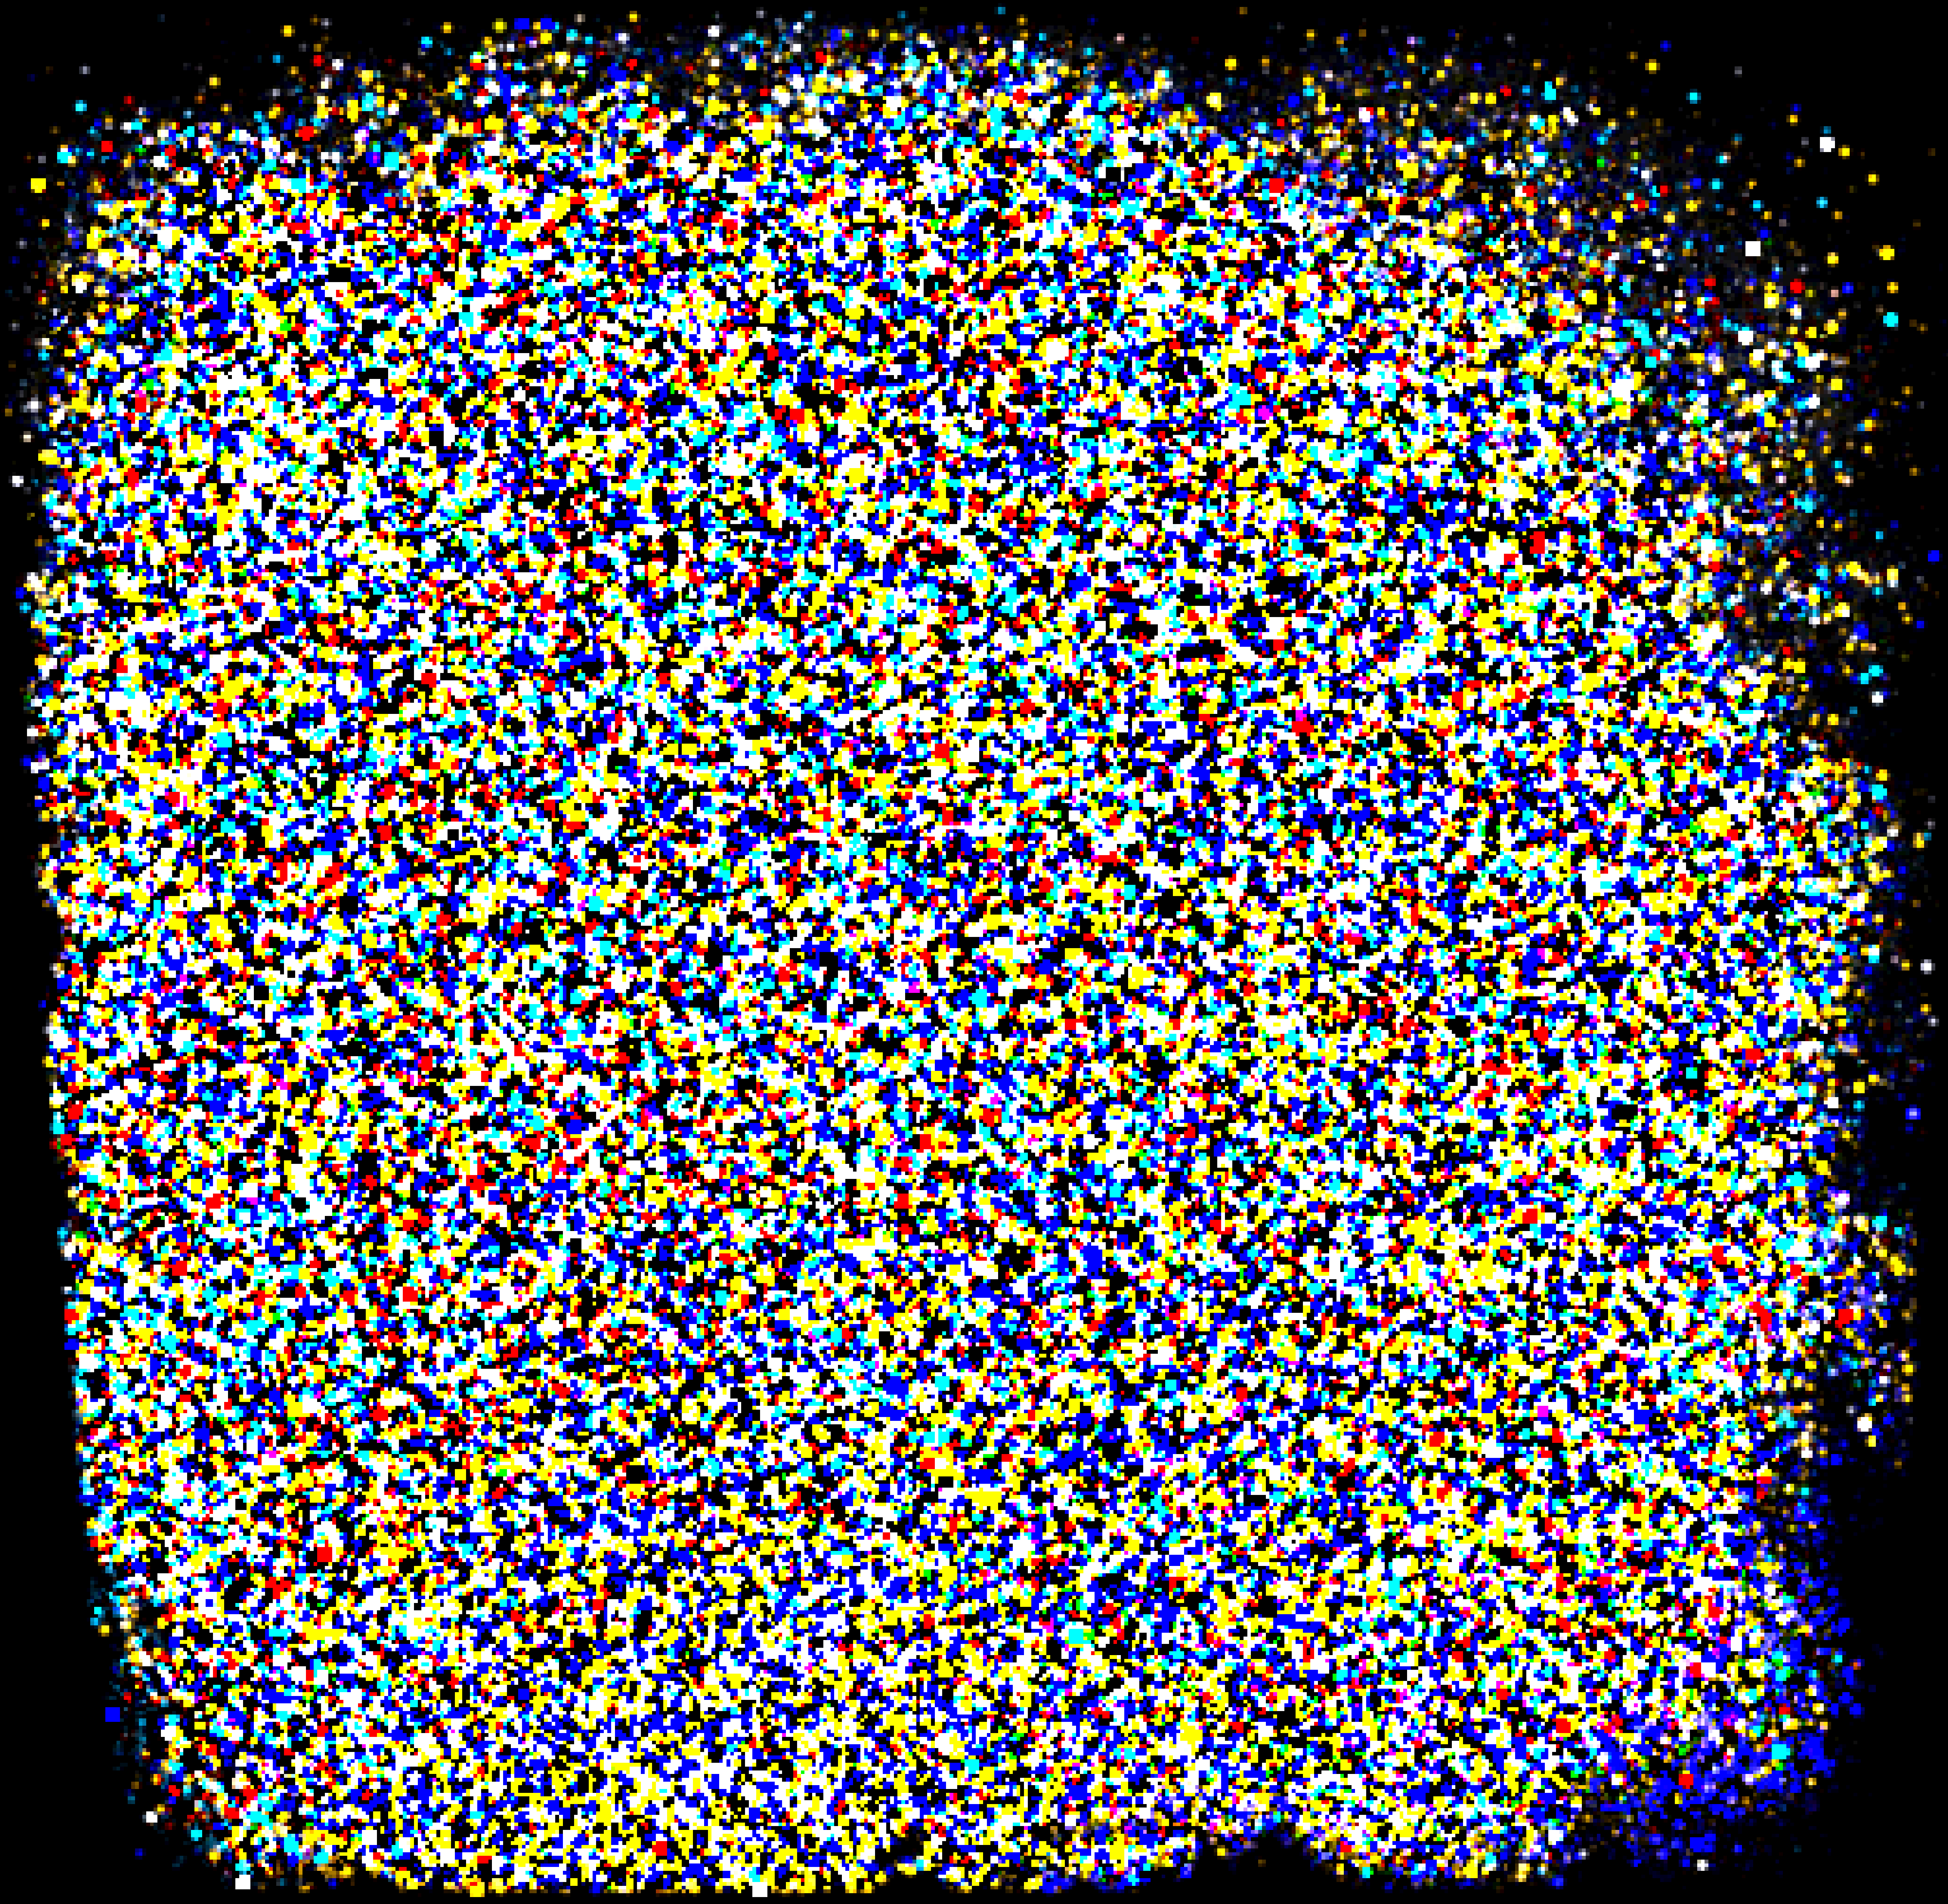

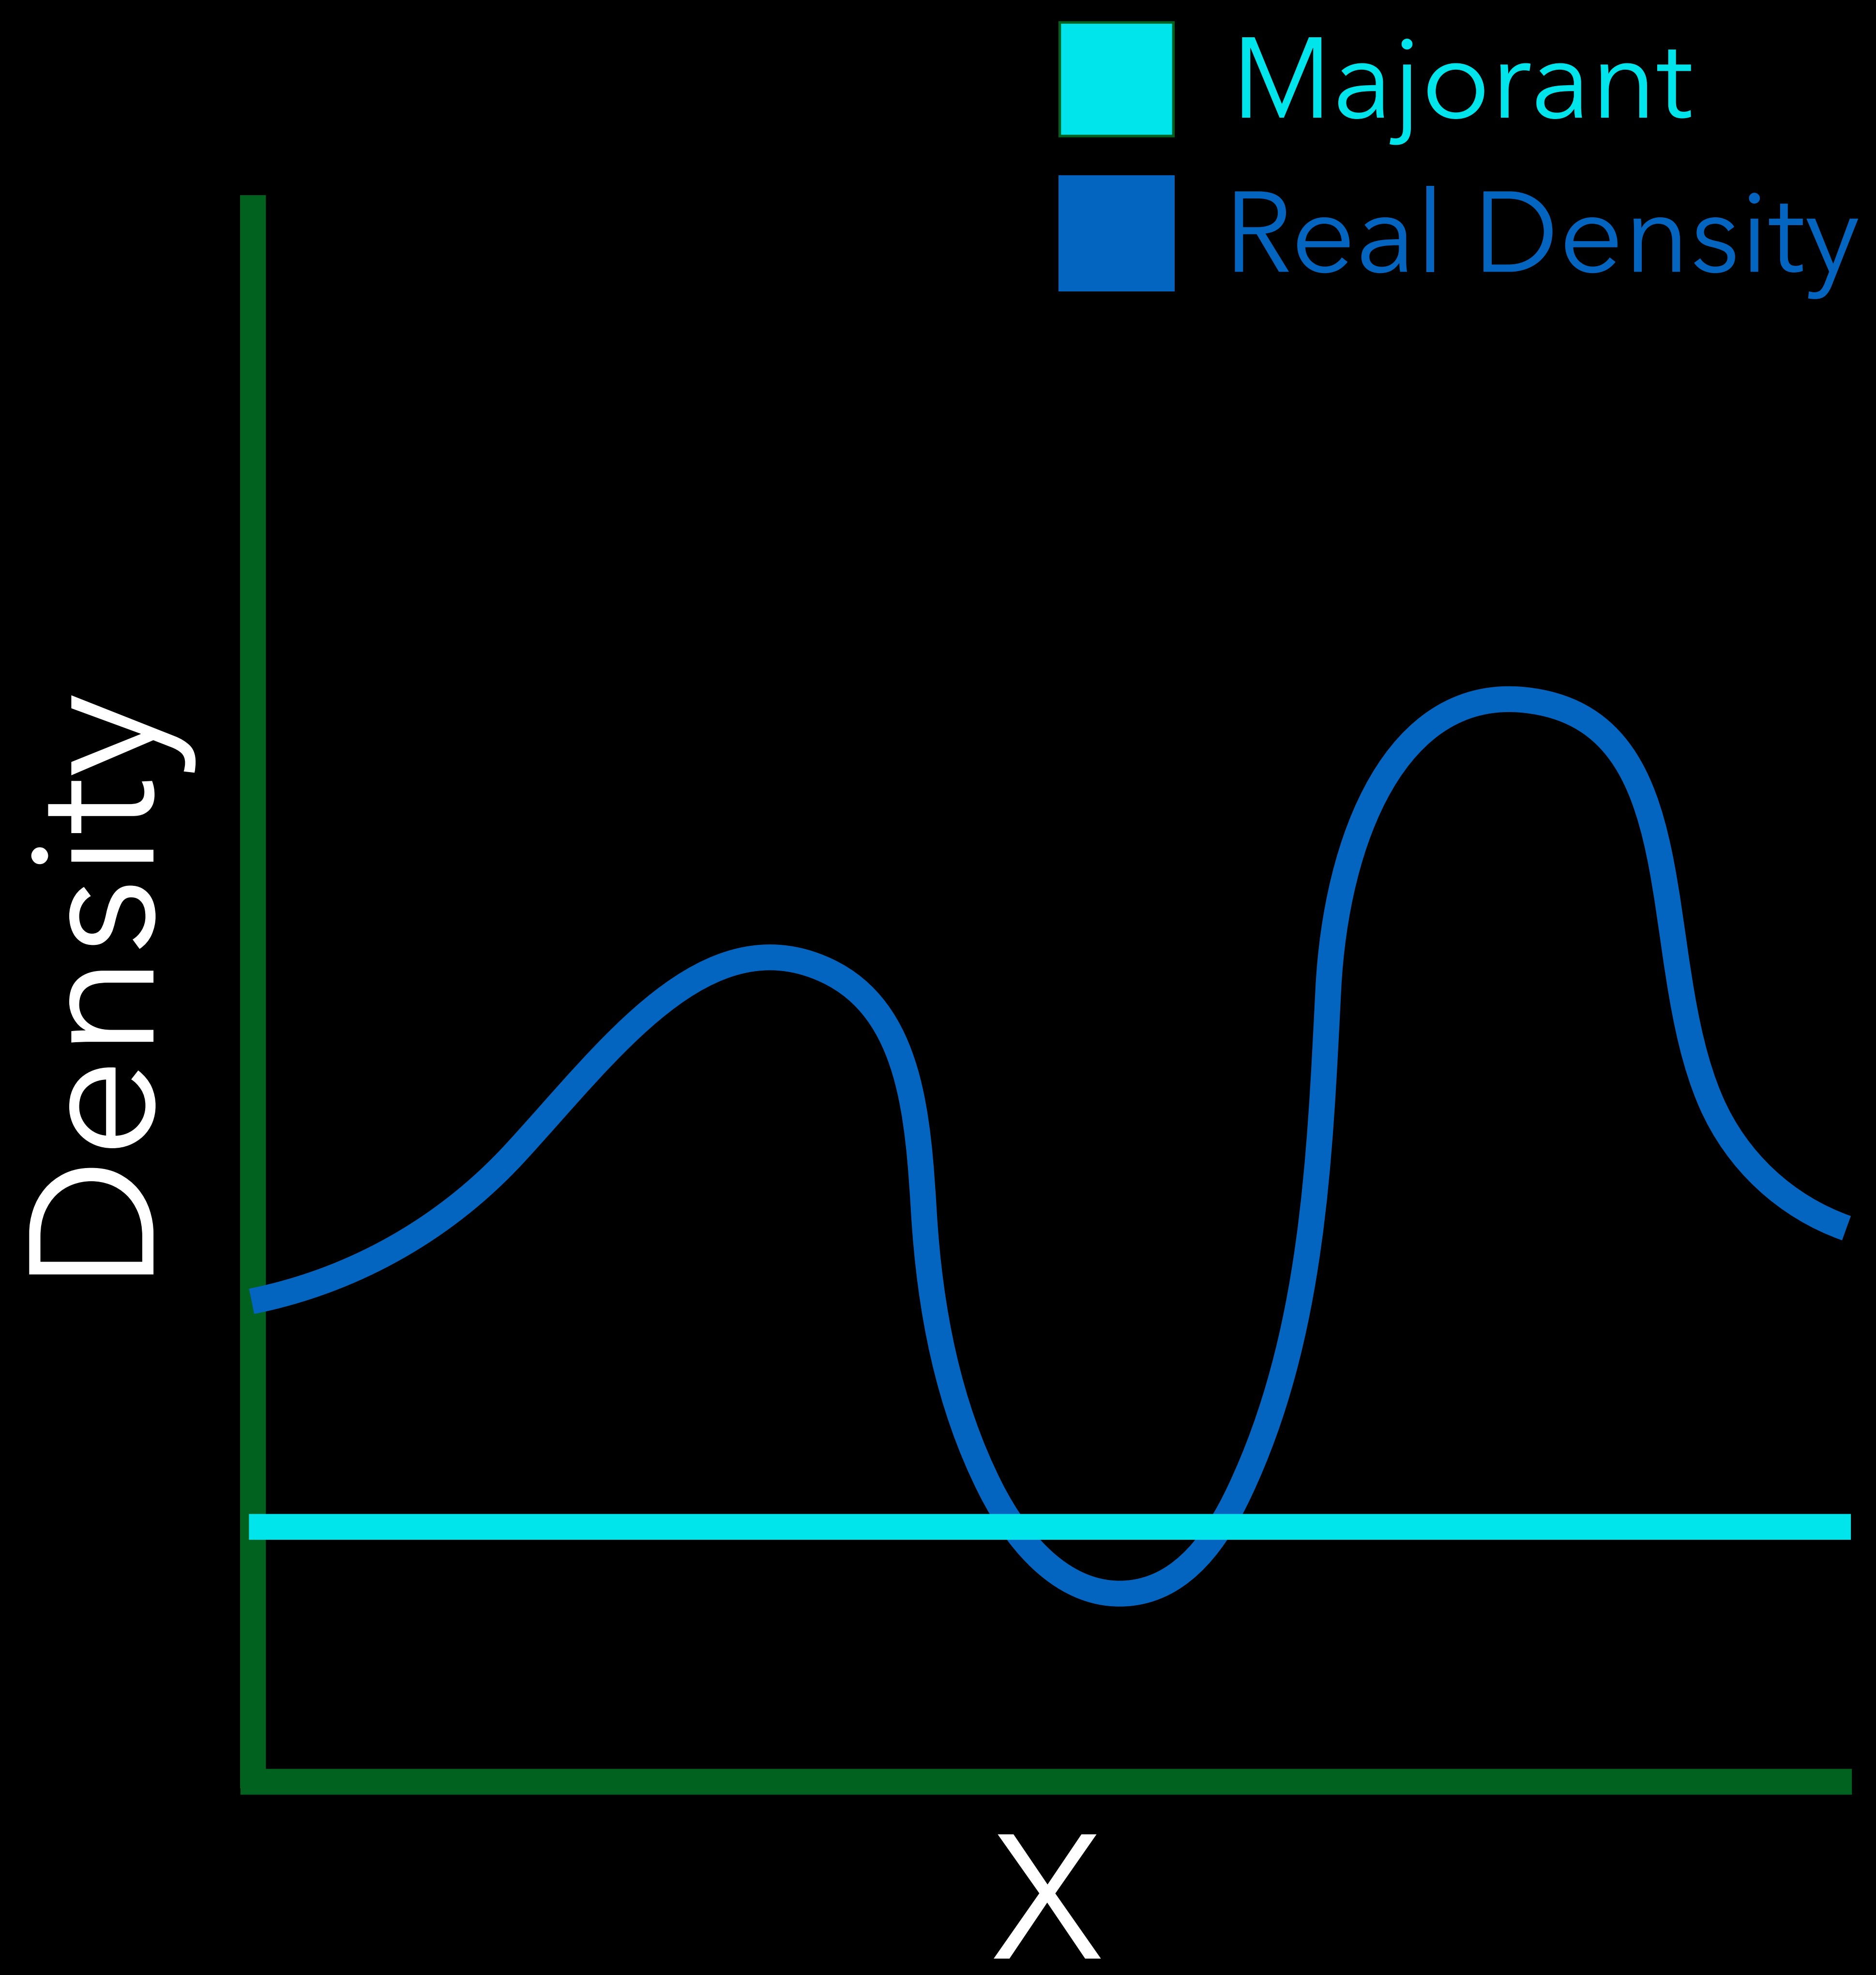
Majorant
Real Density
Density
X

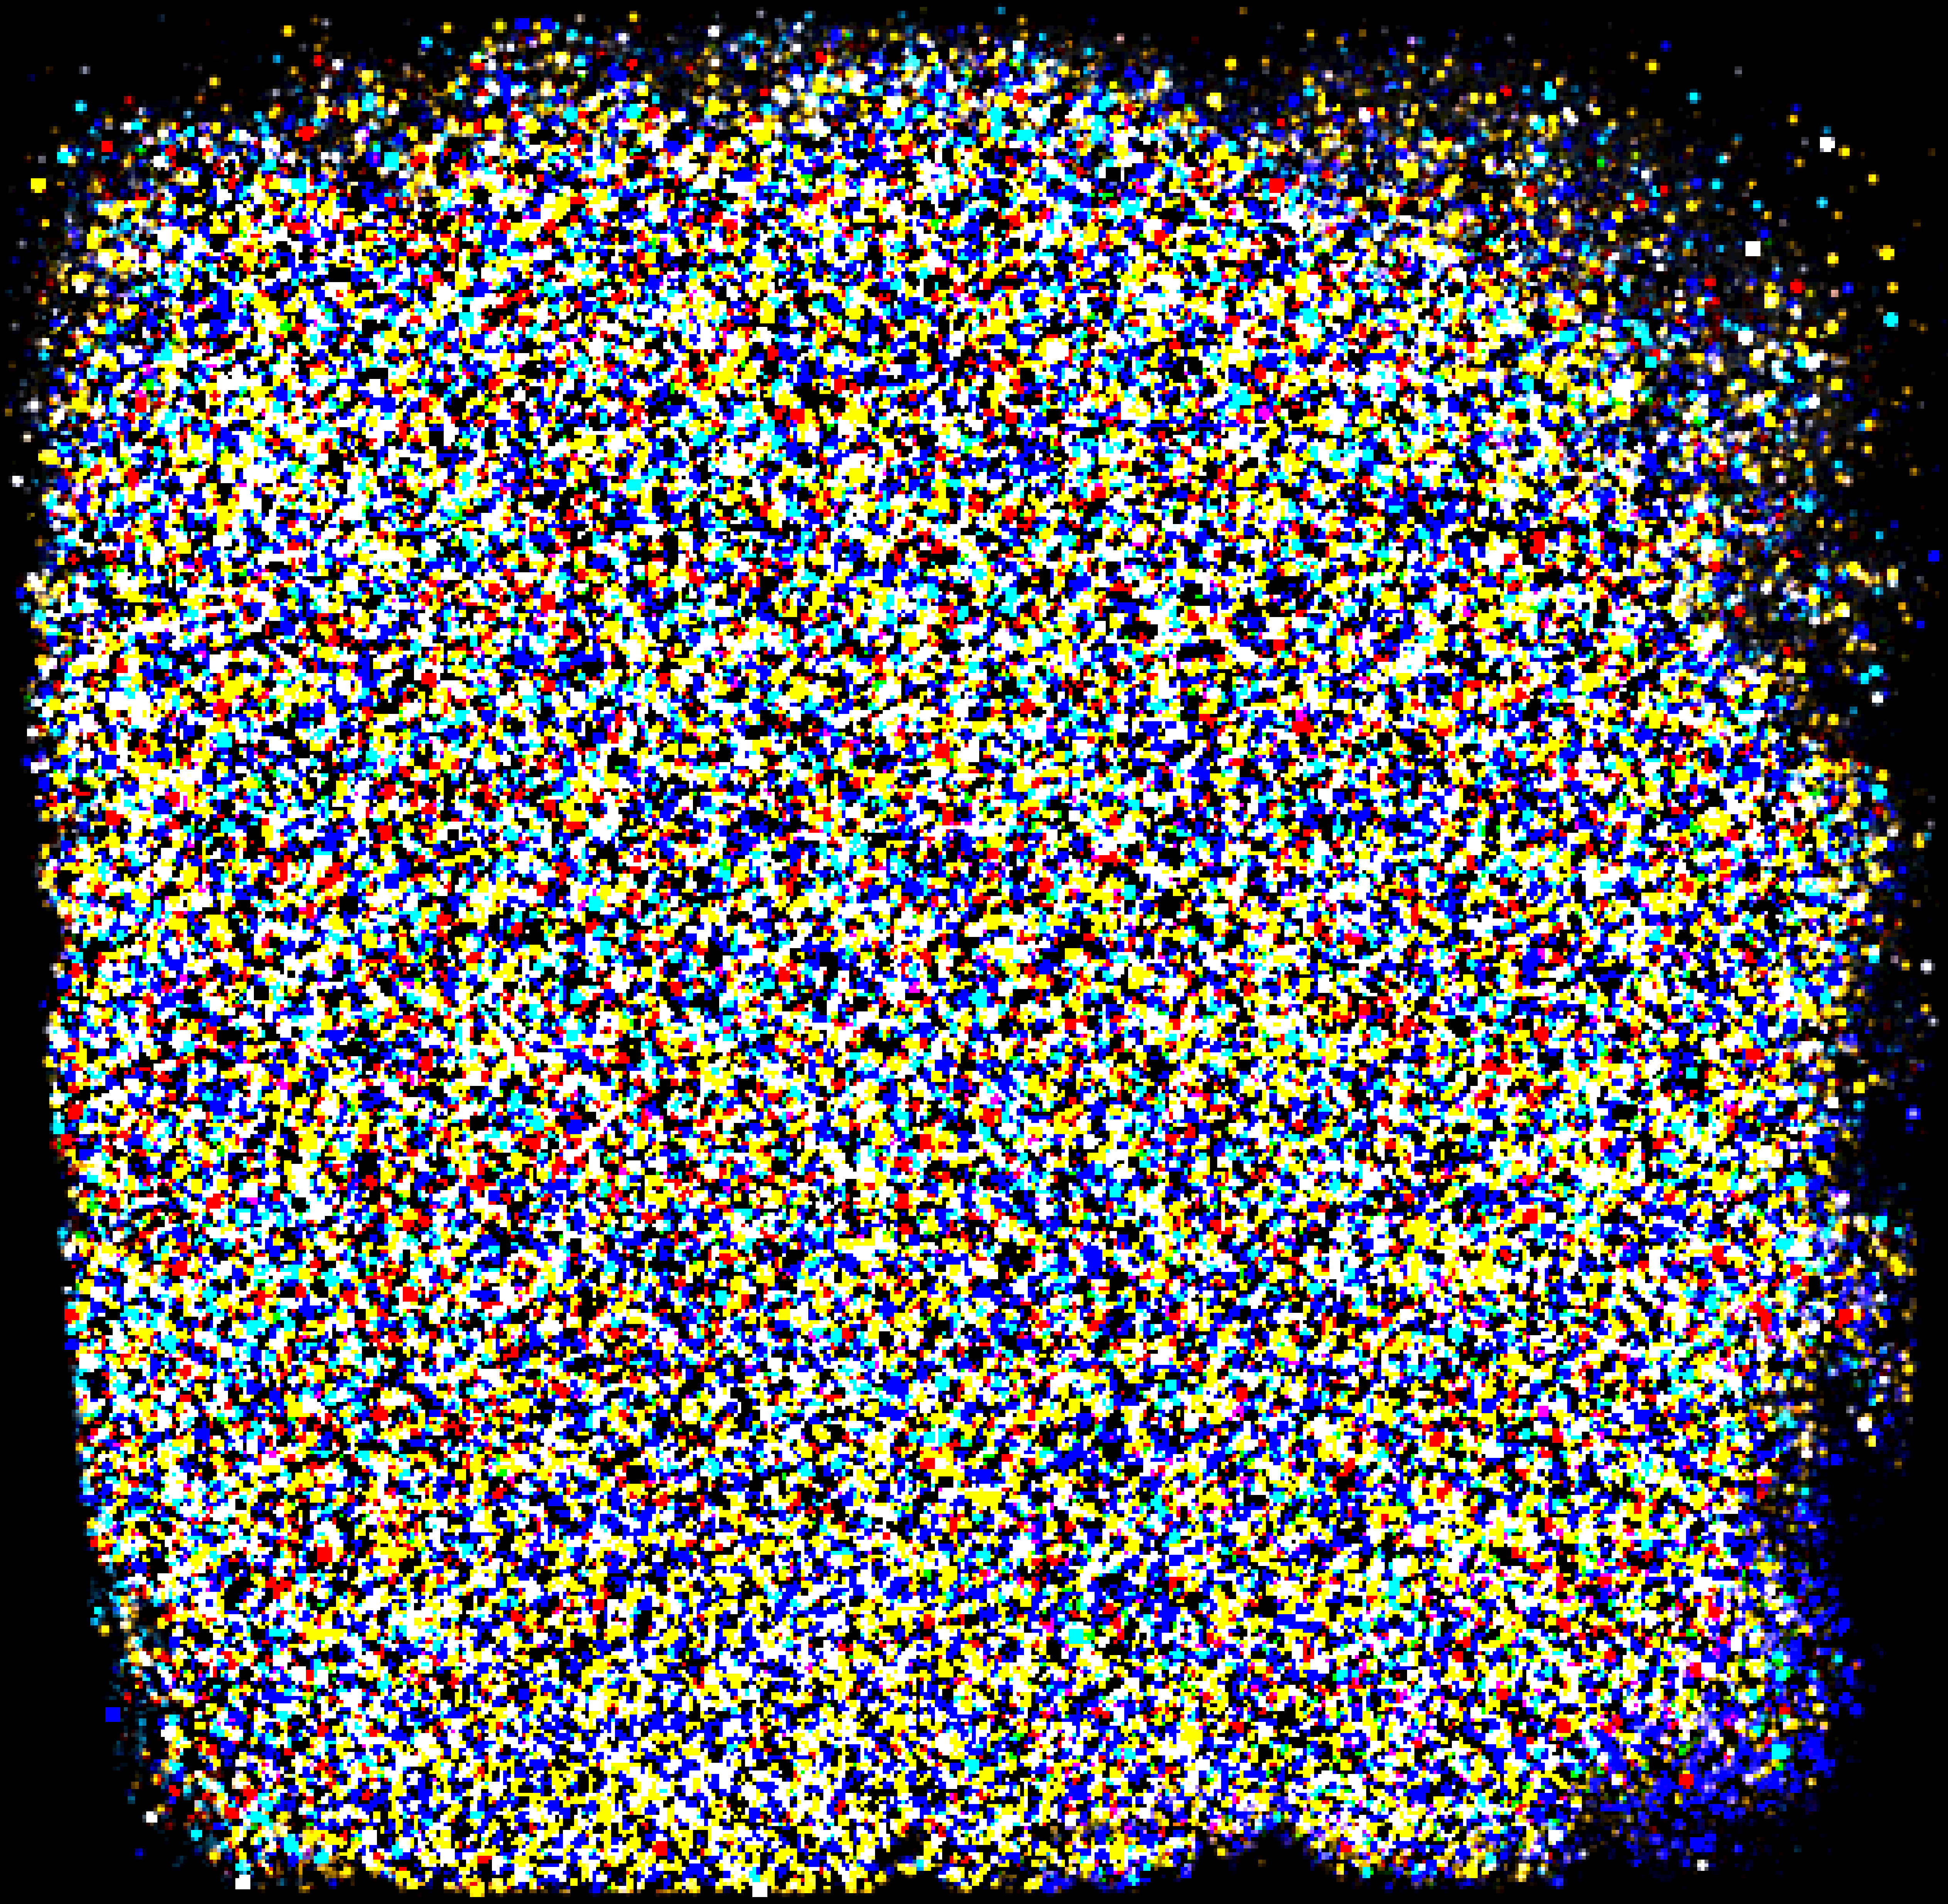

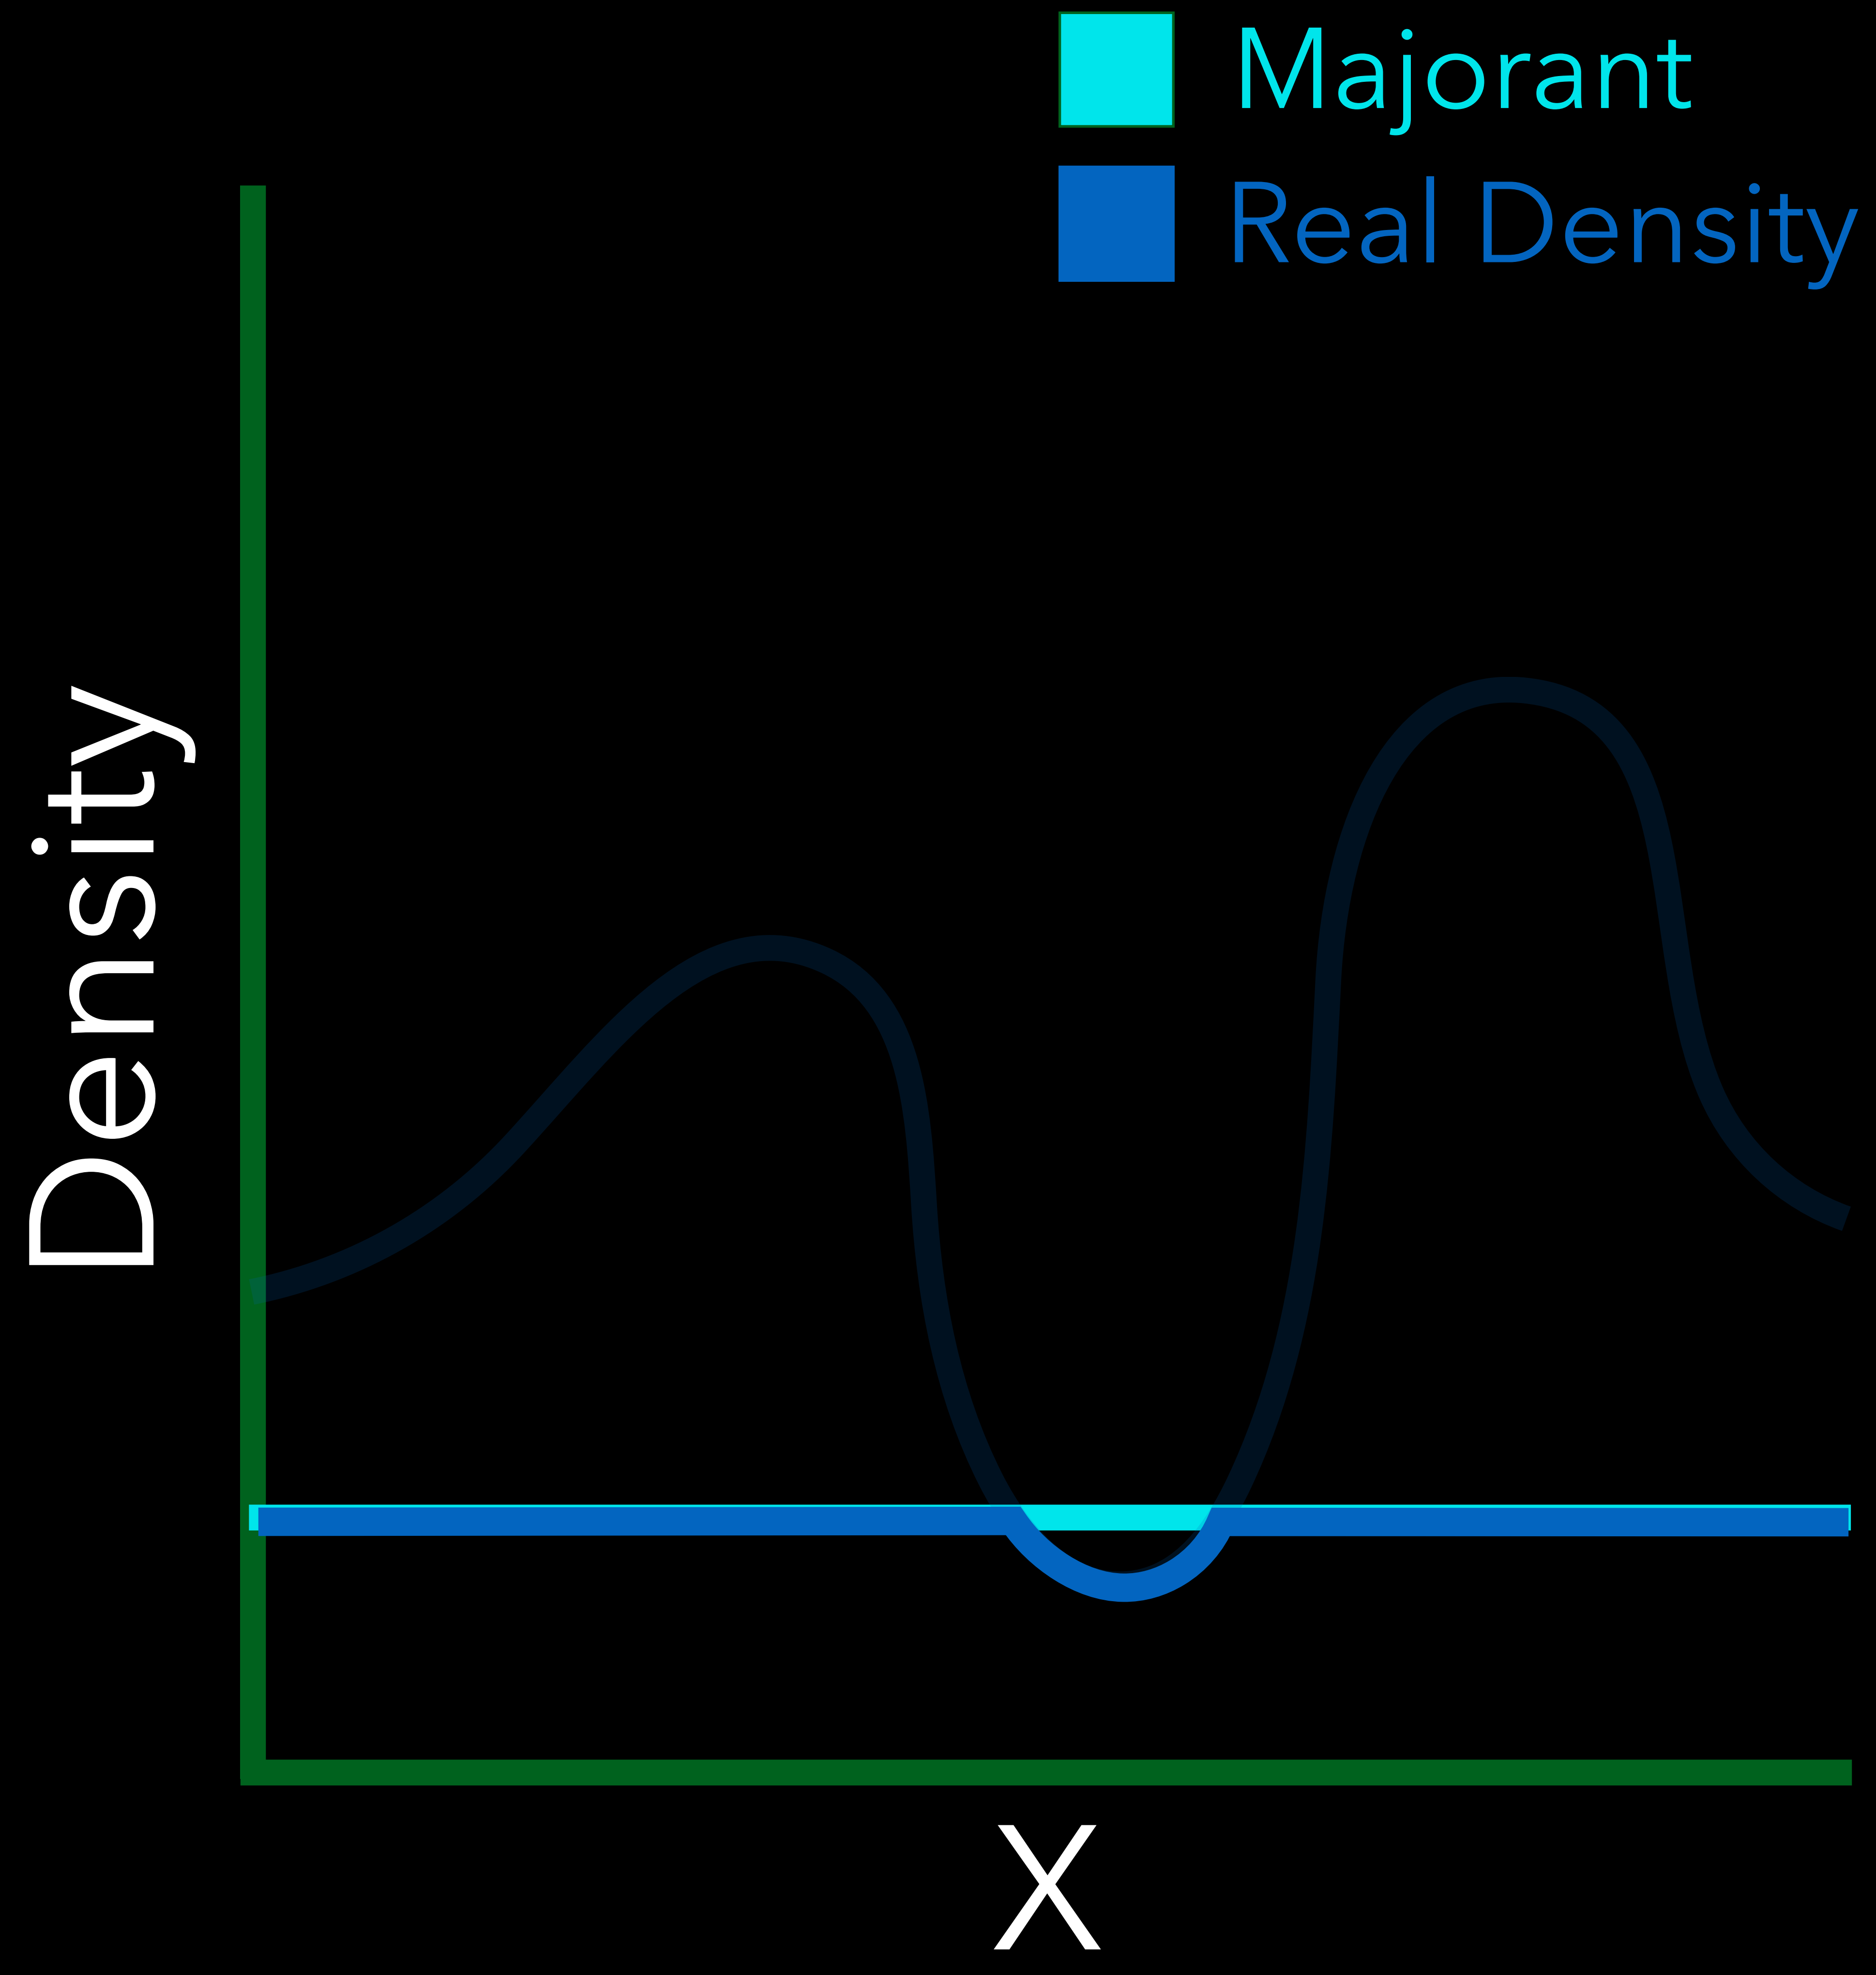
Majorant
Real Density
Density
X

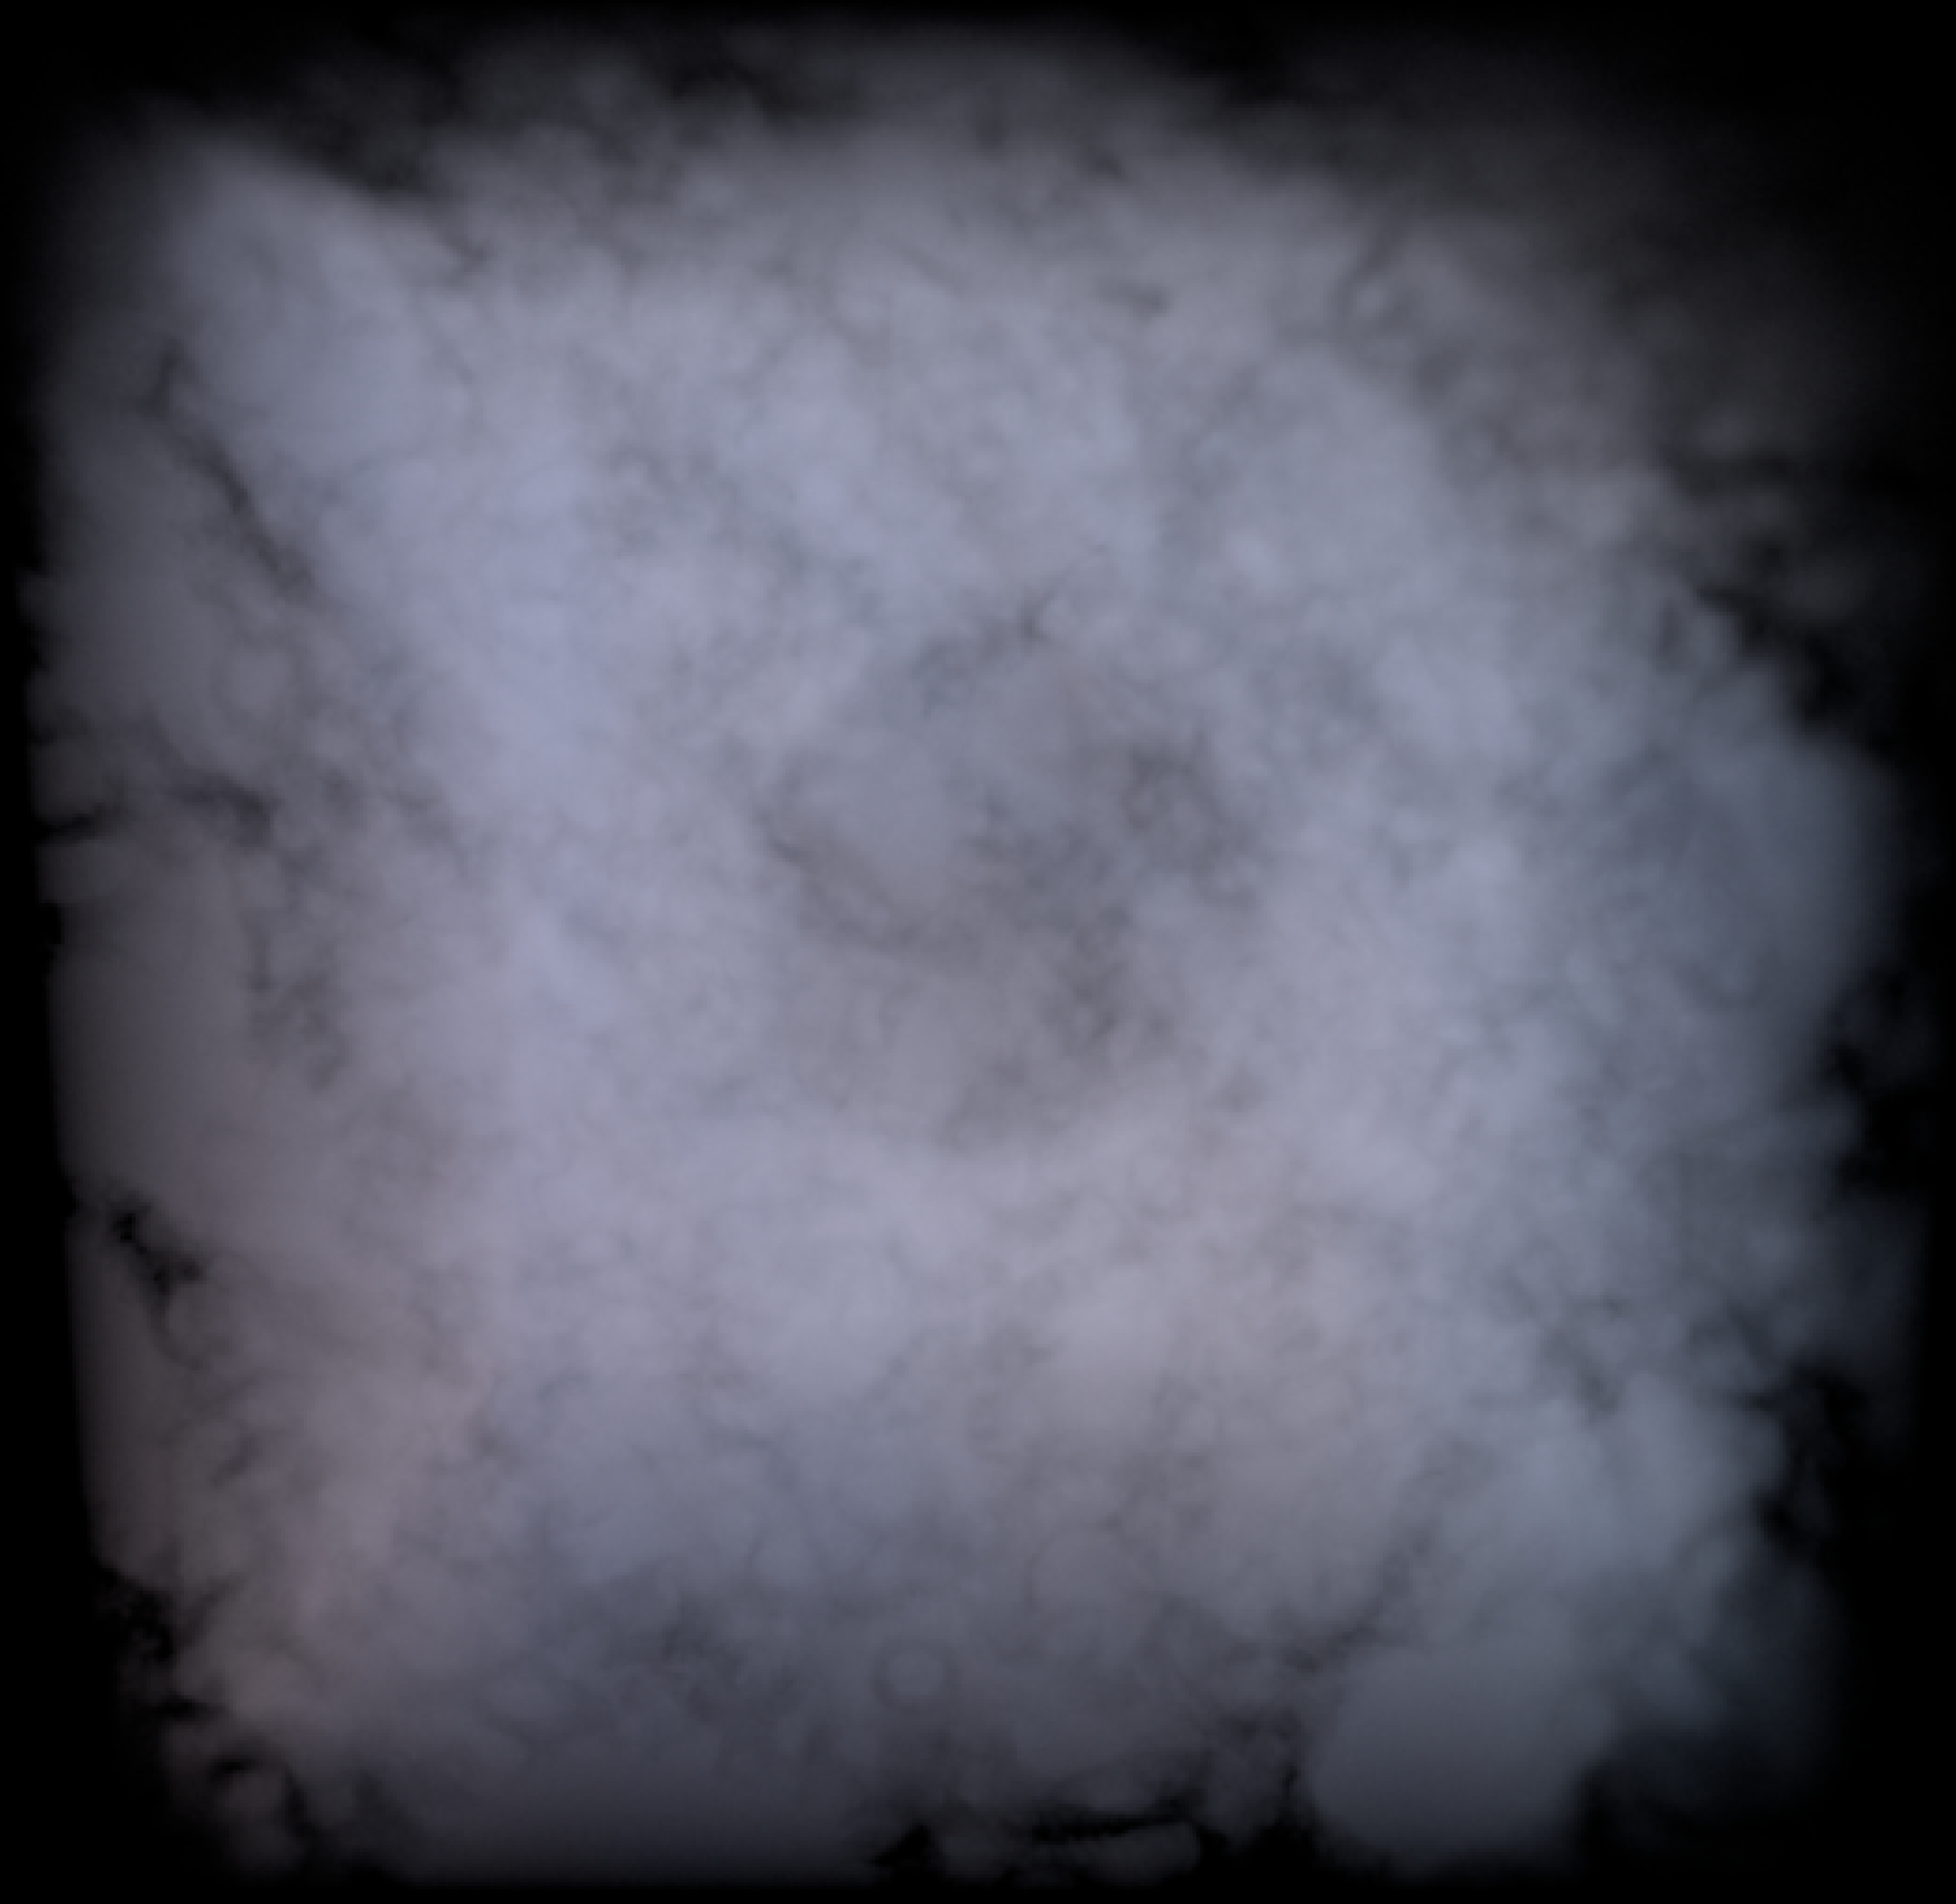

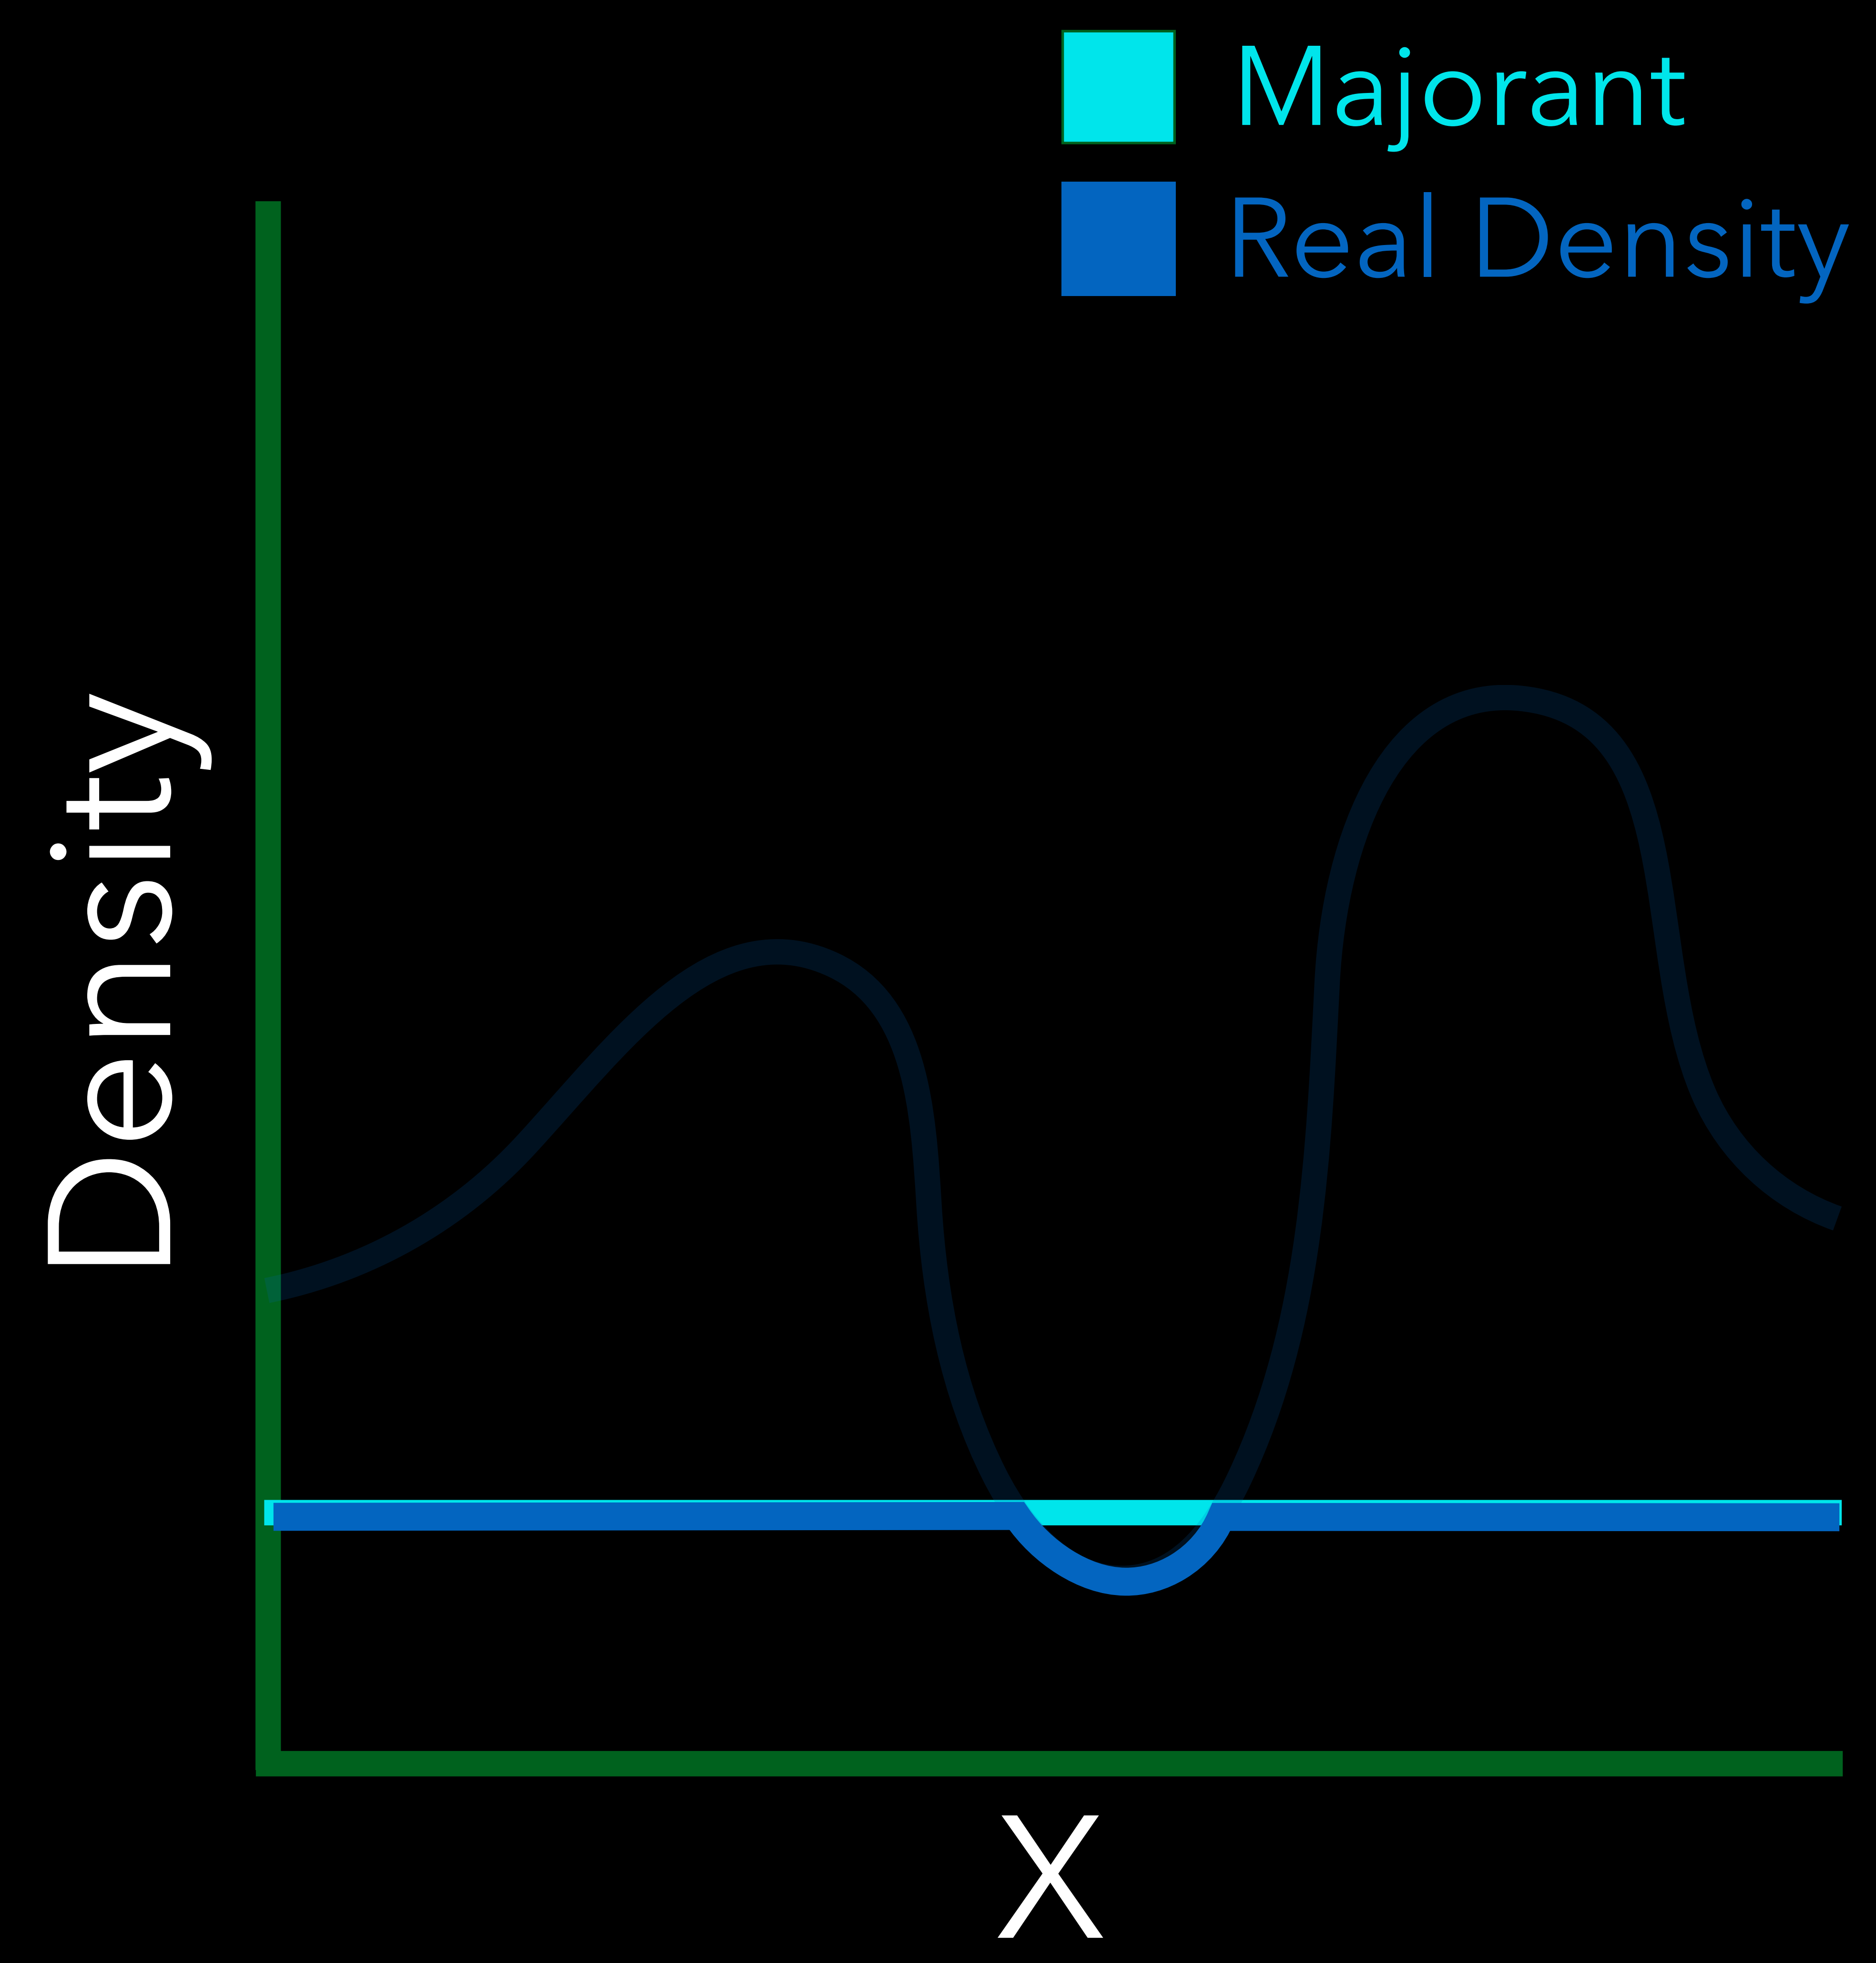
Majorant
Real Density
Density
X

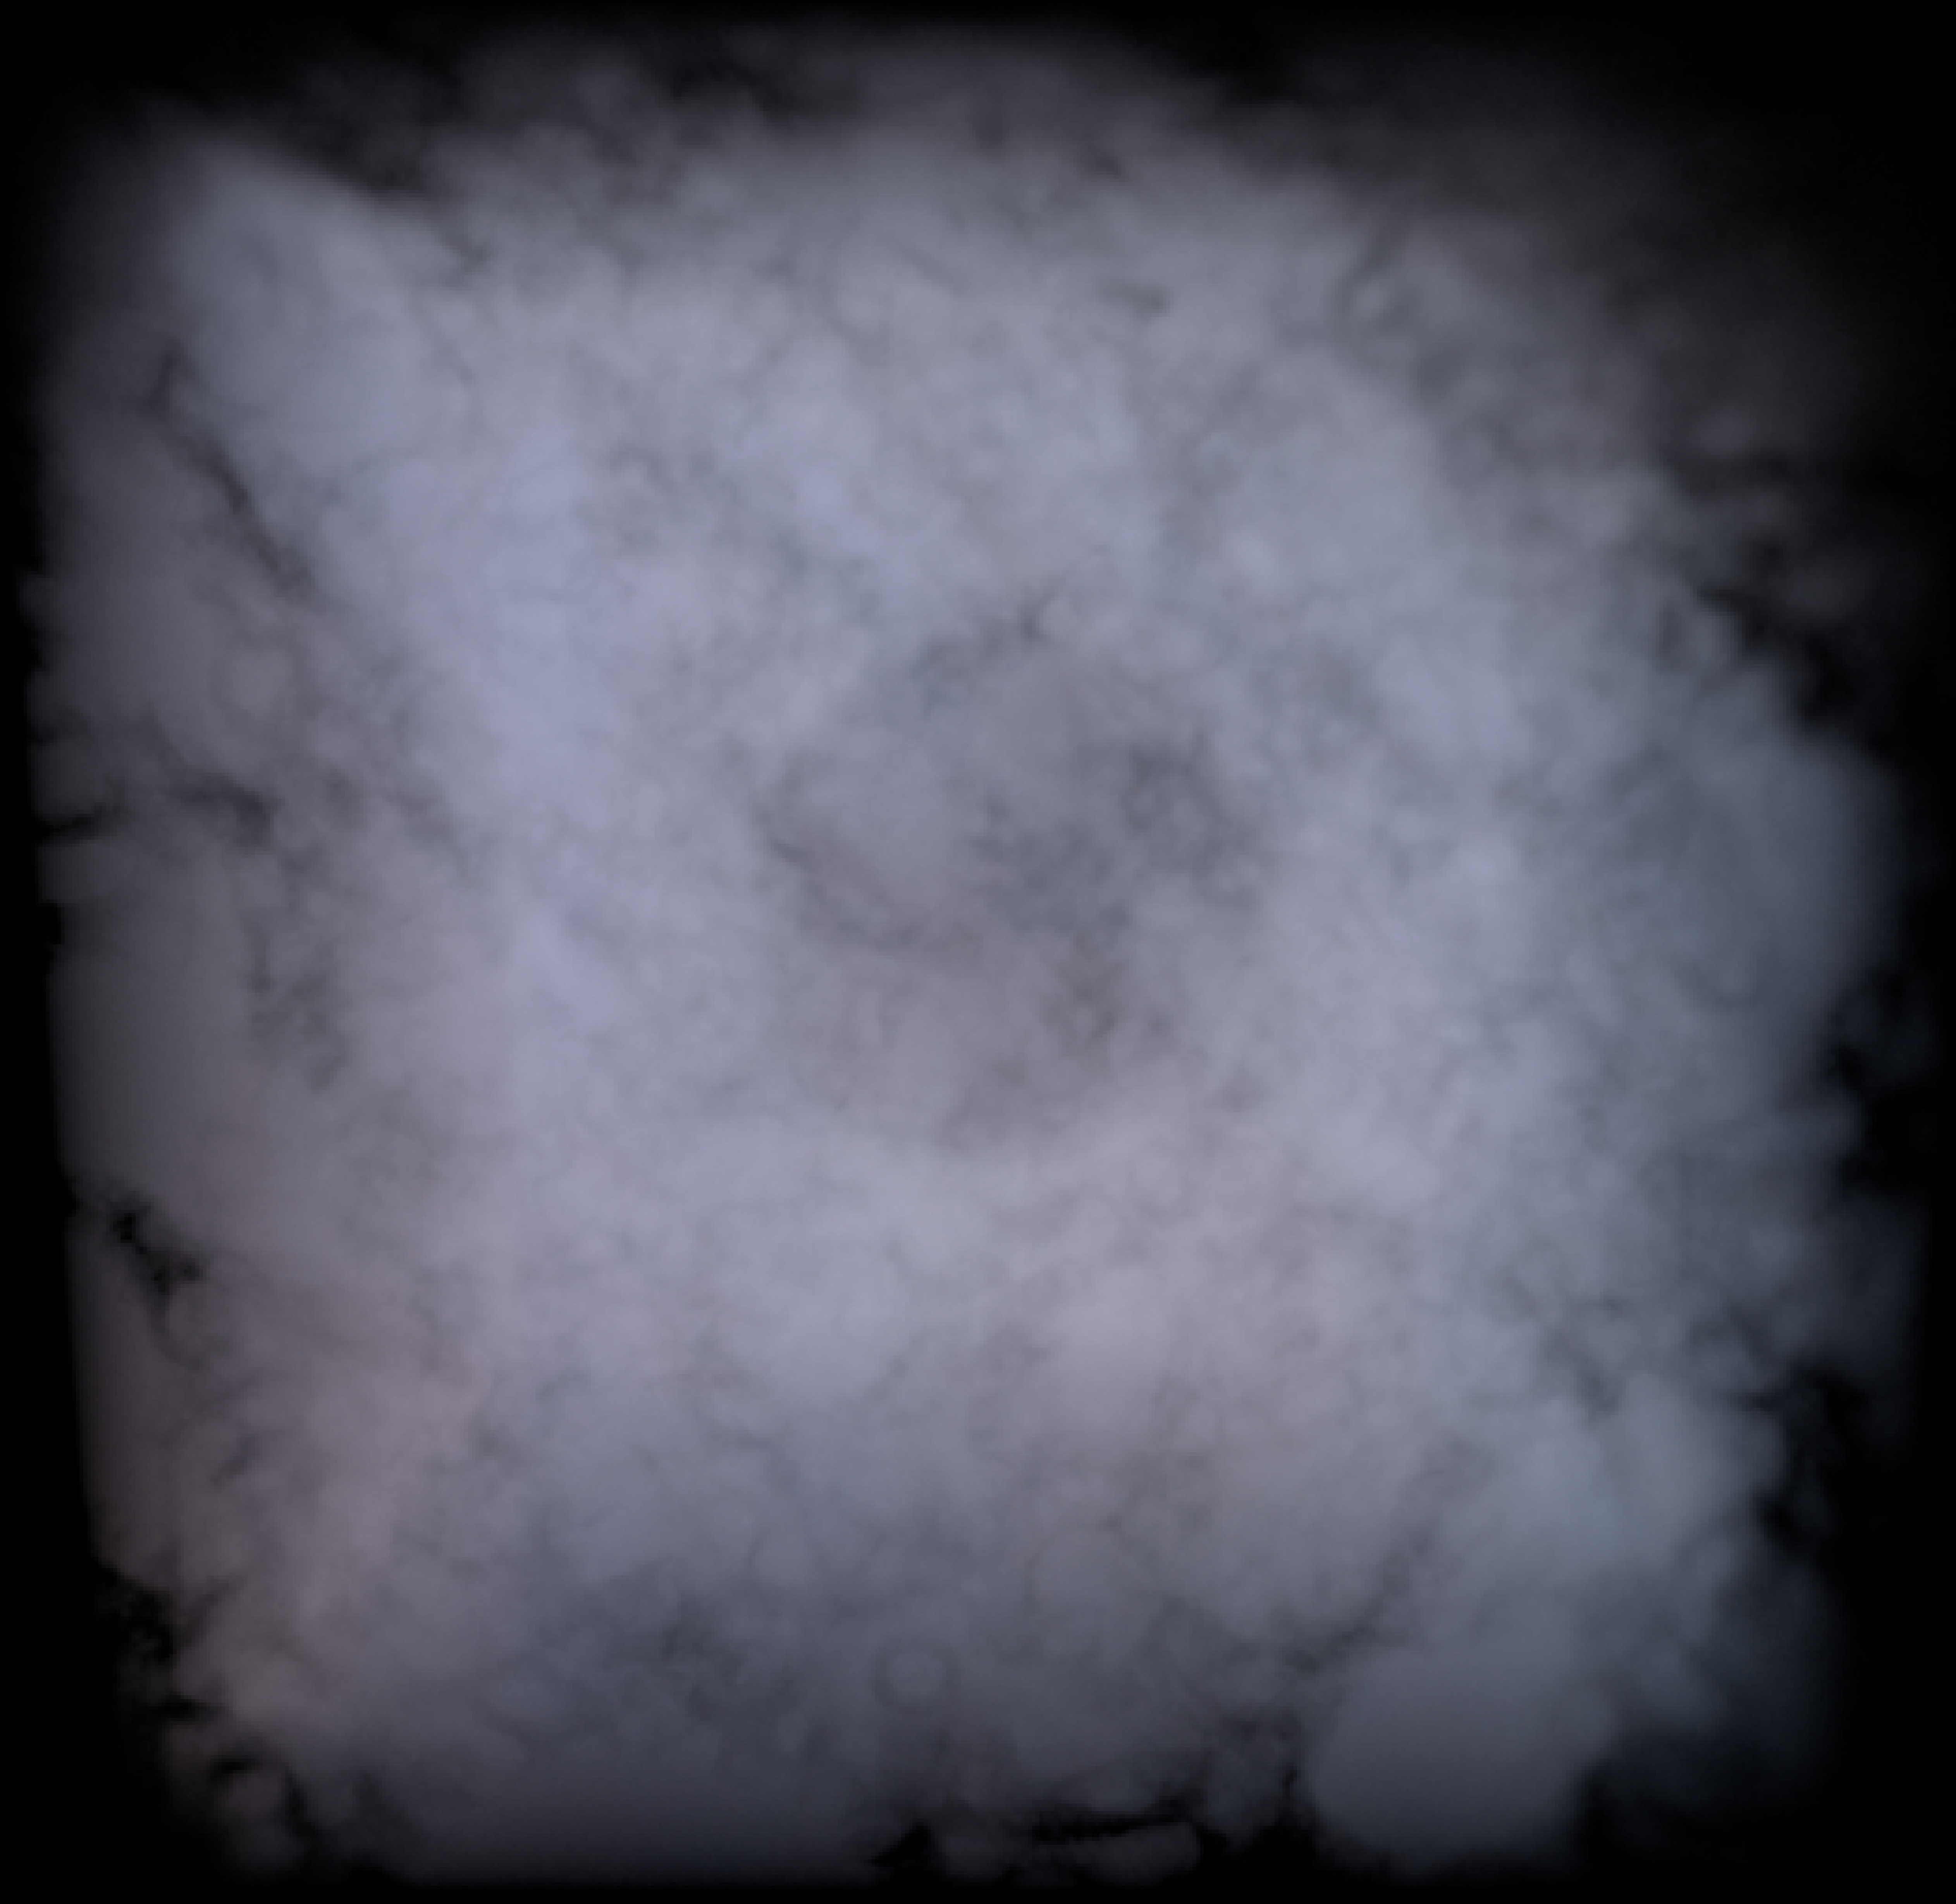

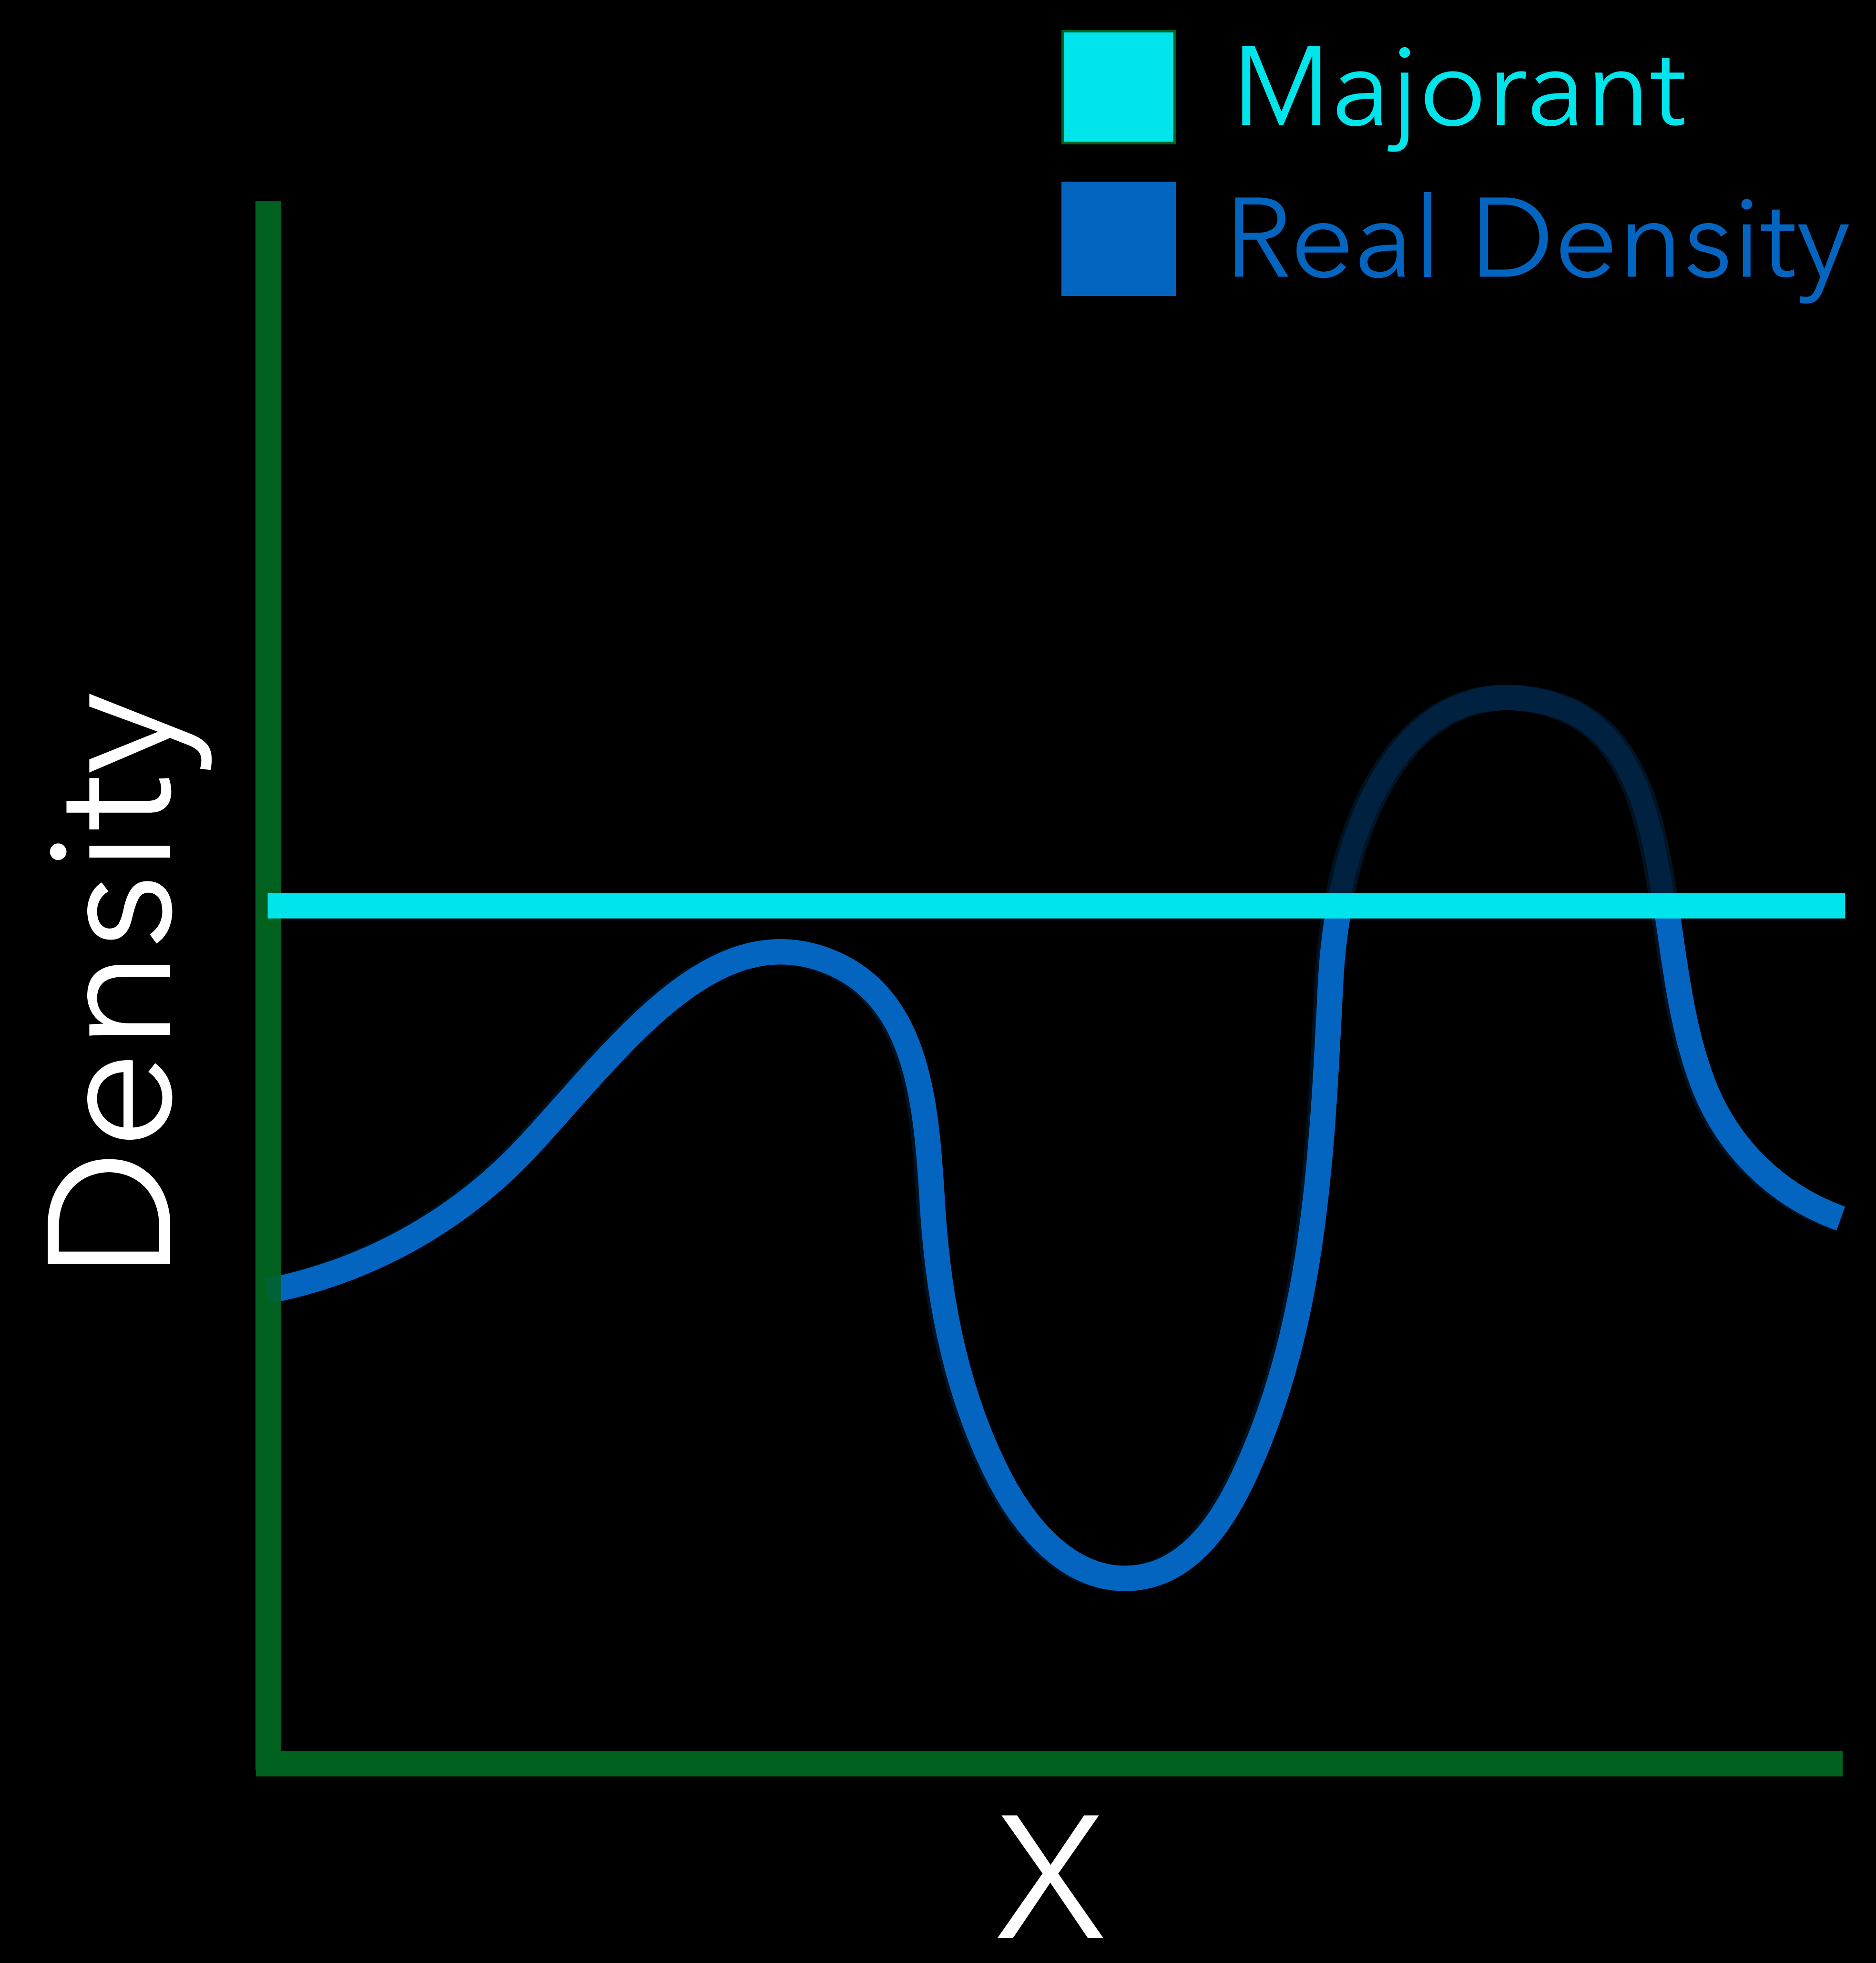
Majorant
Real Density
Density
X

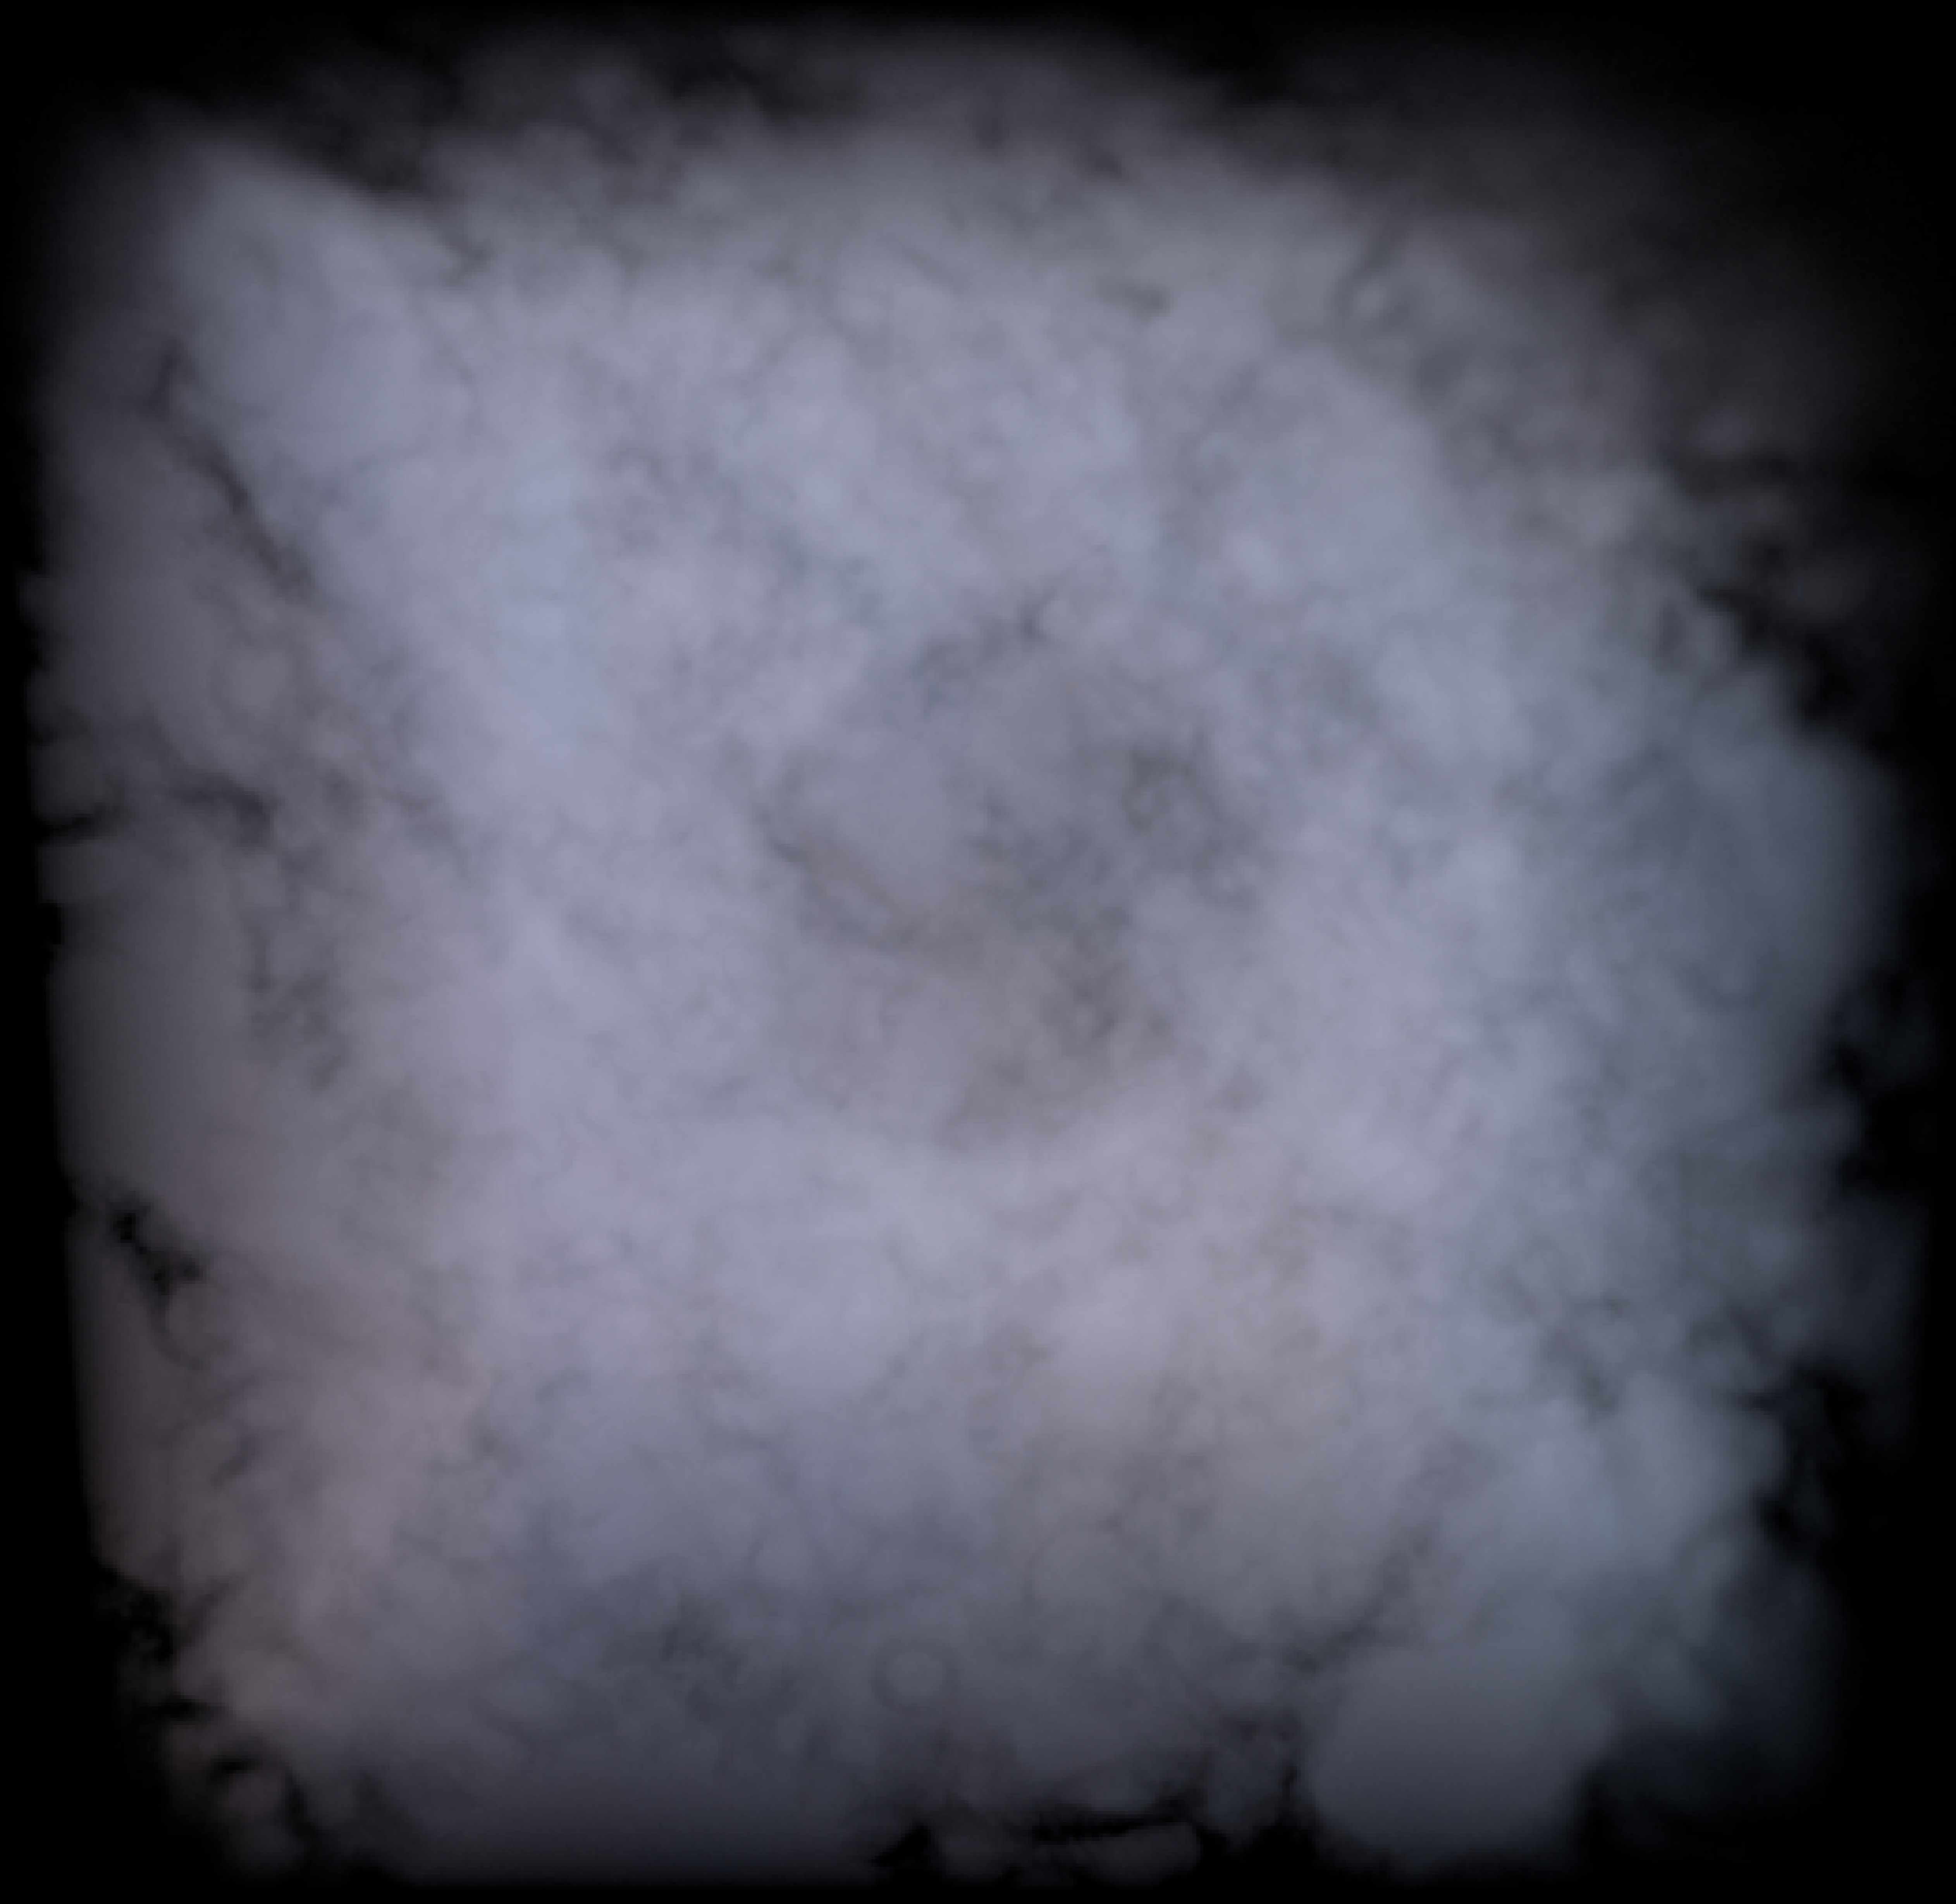

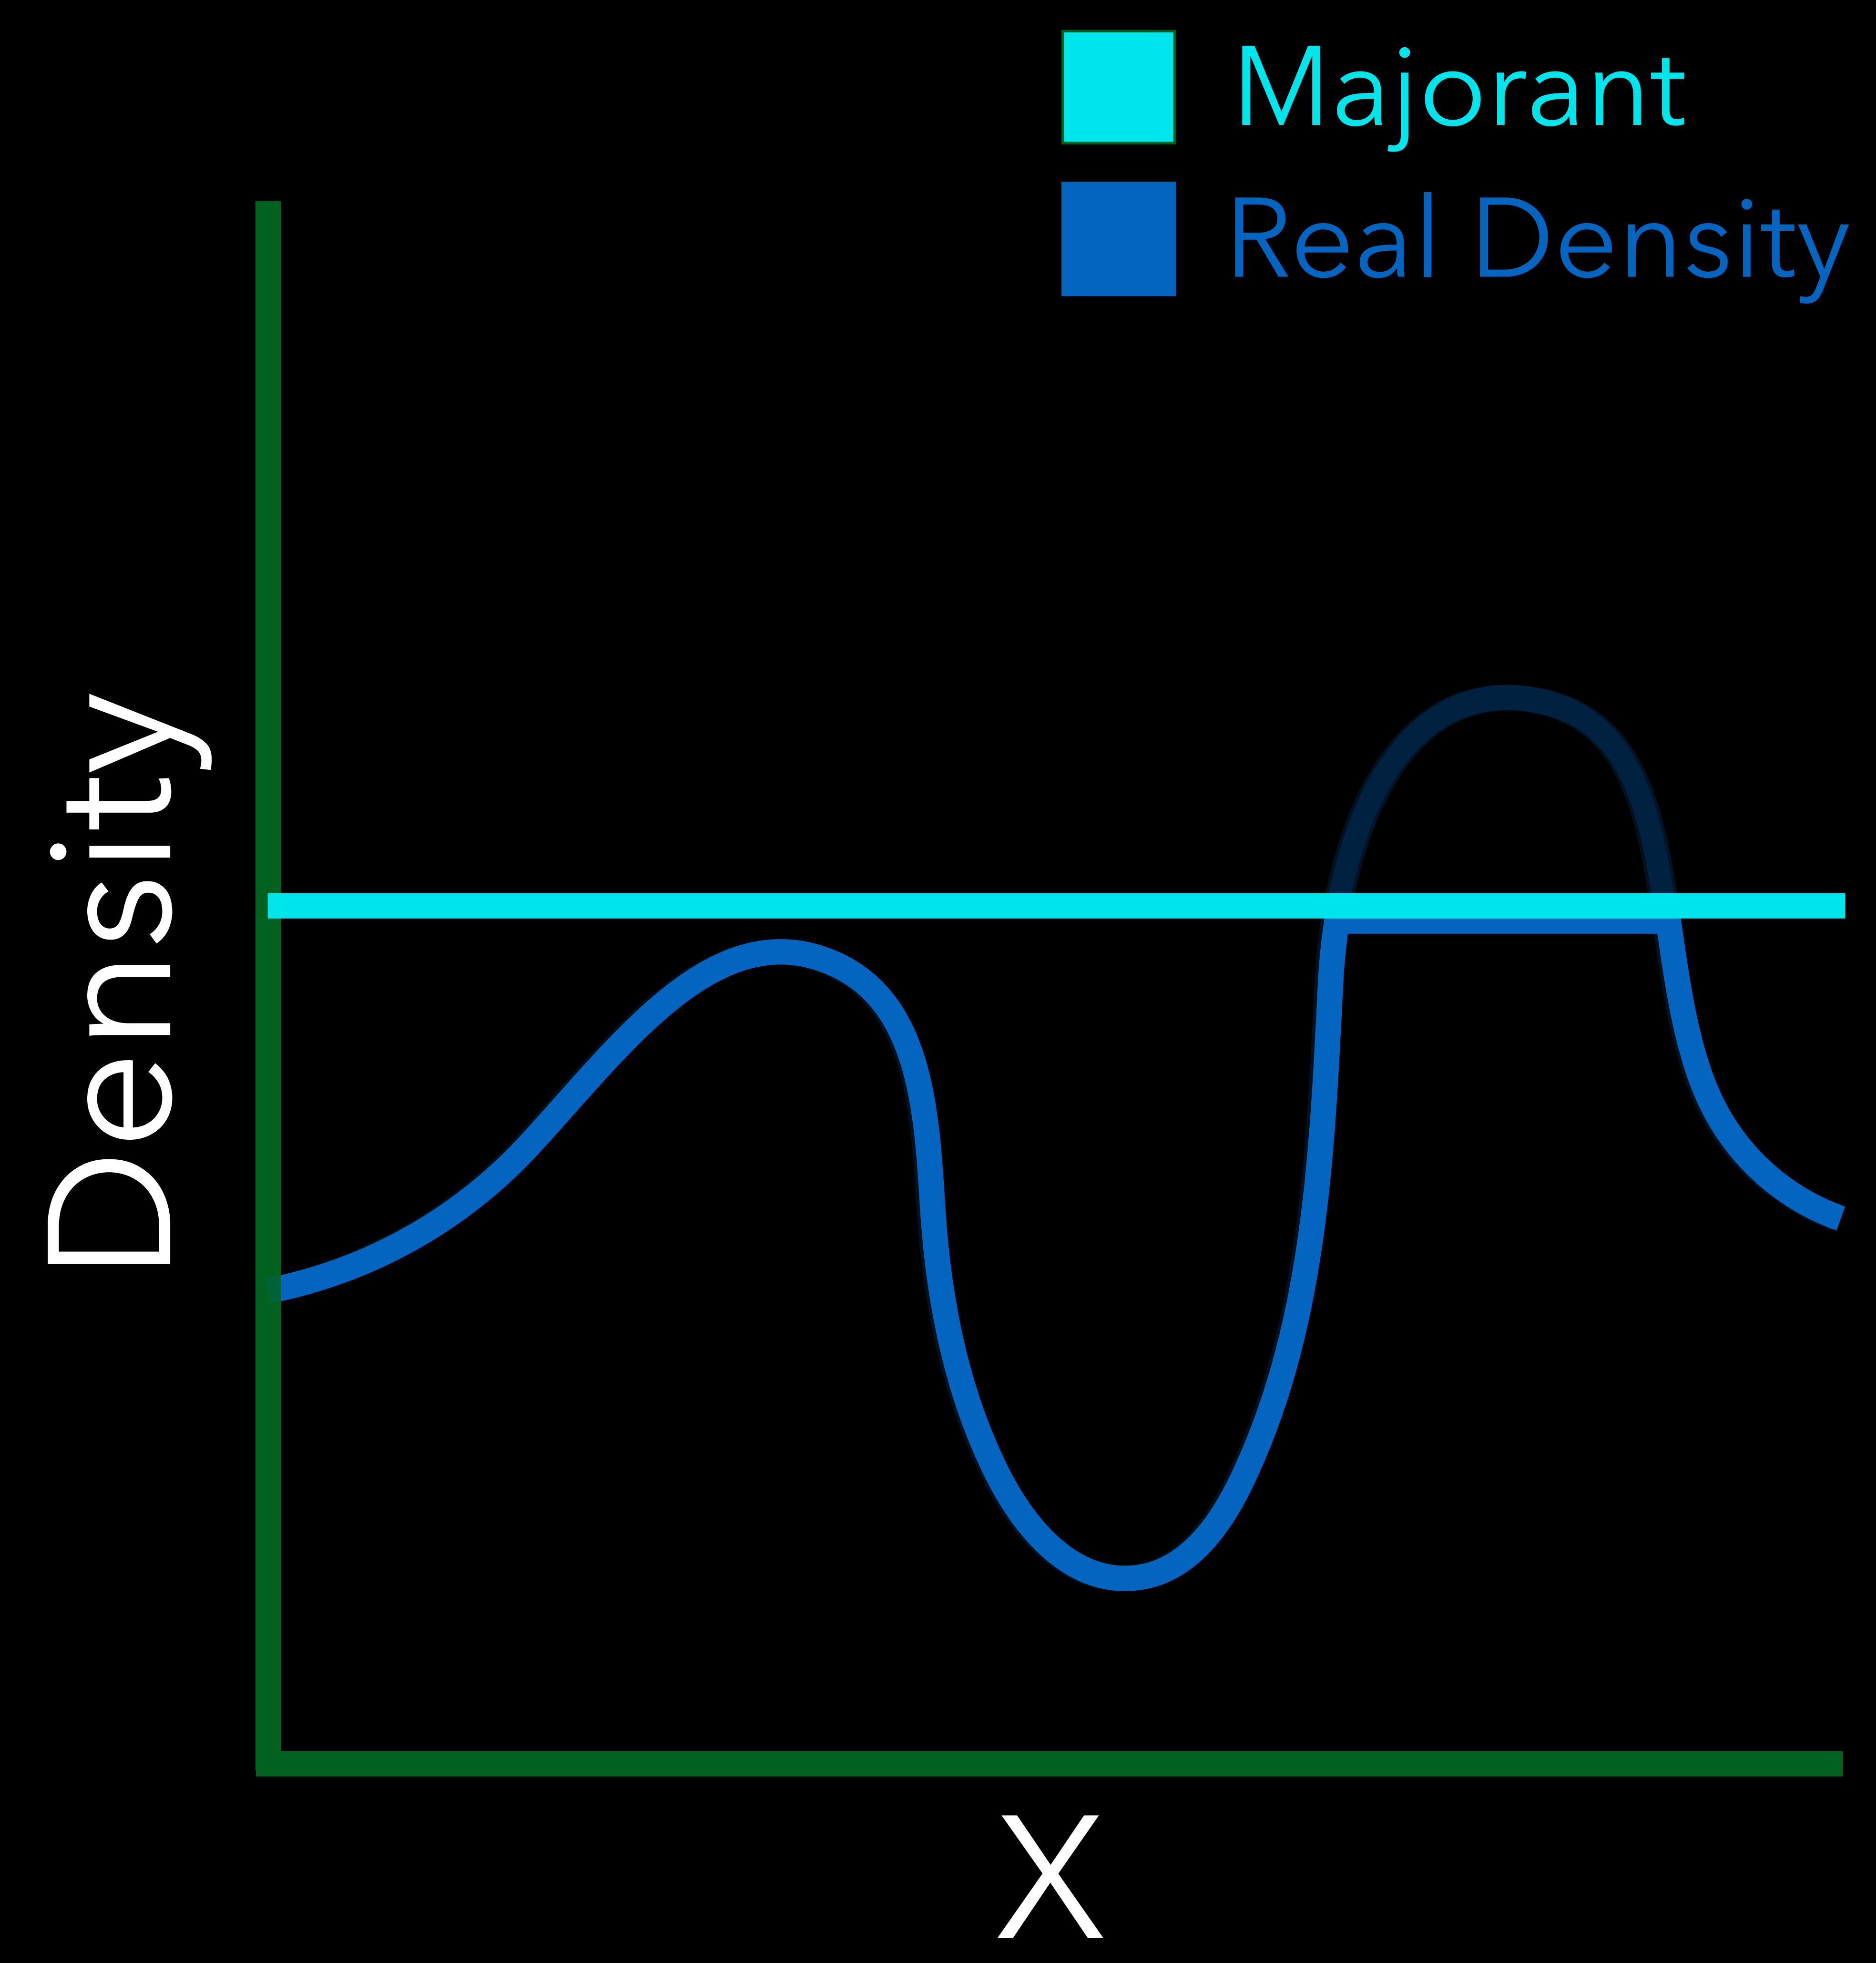
Majorant
Real Density
Density
X

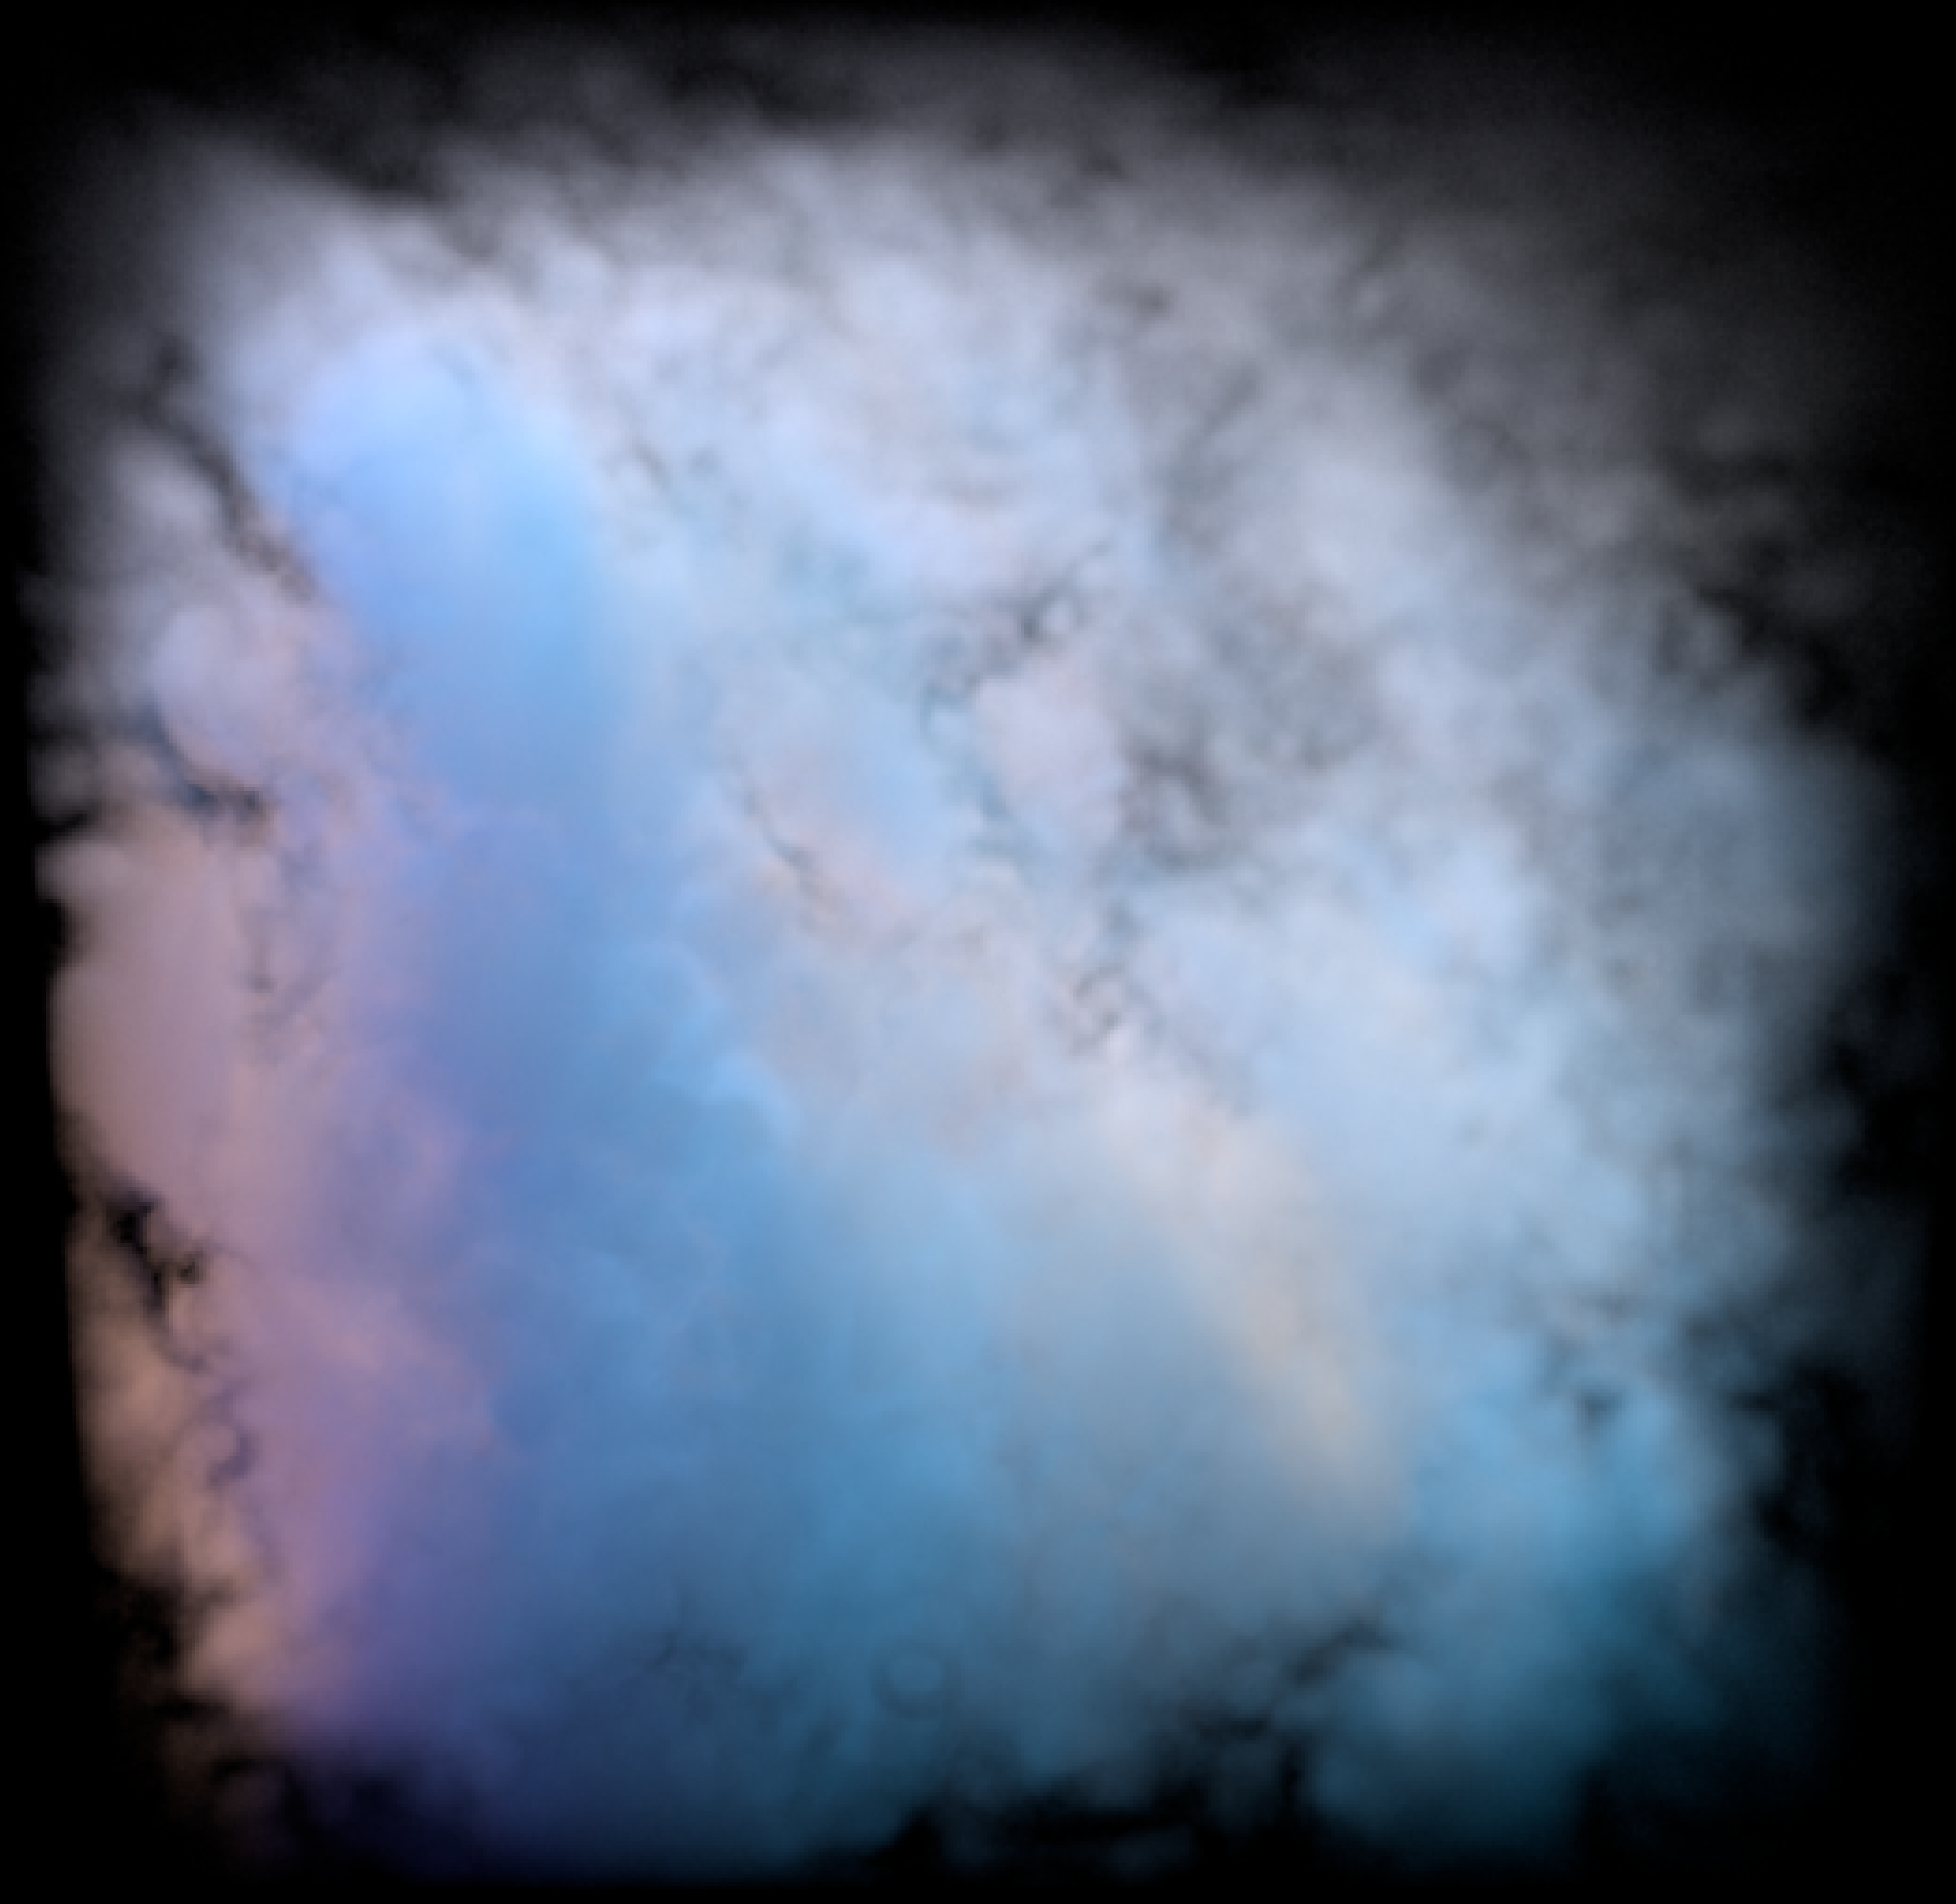

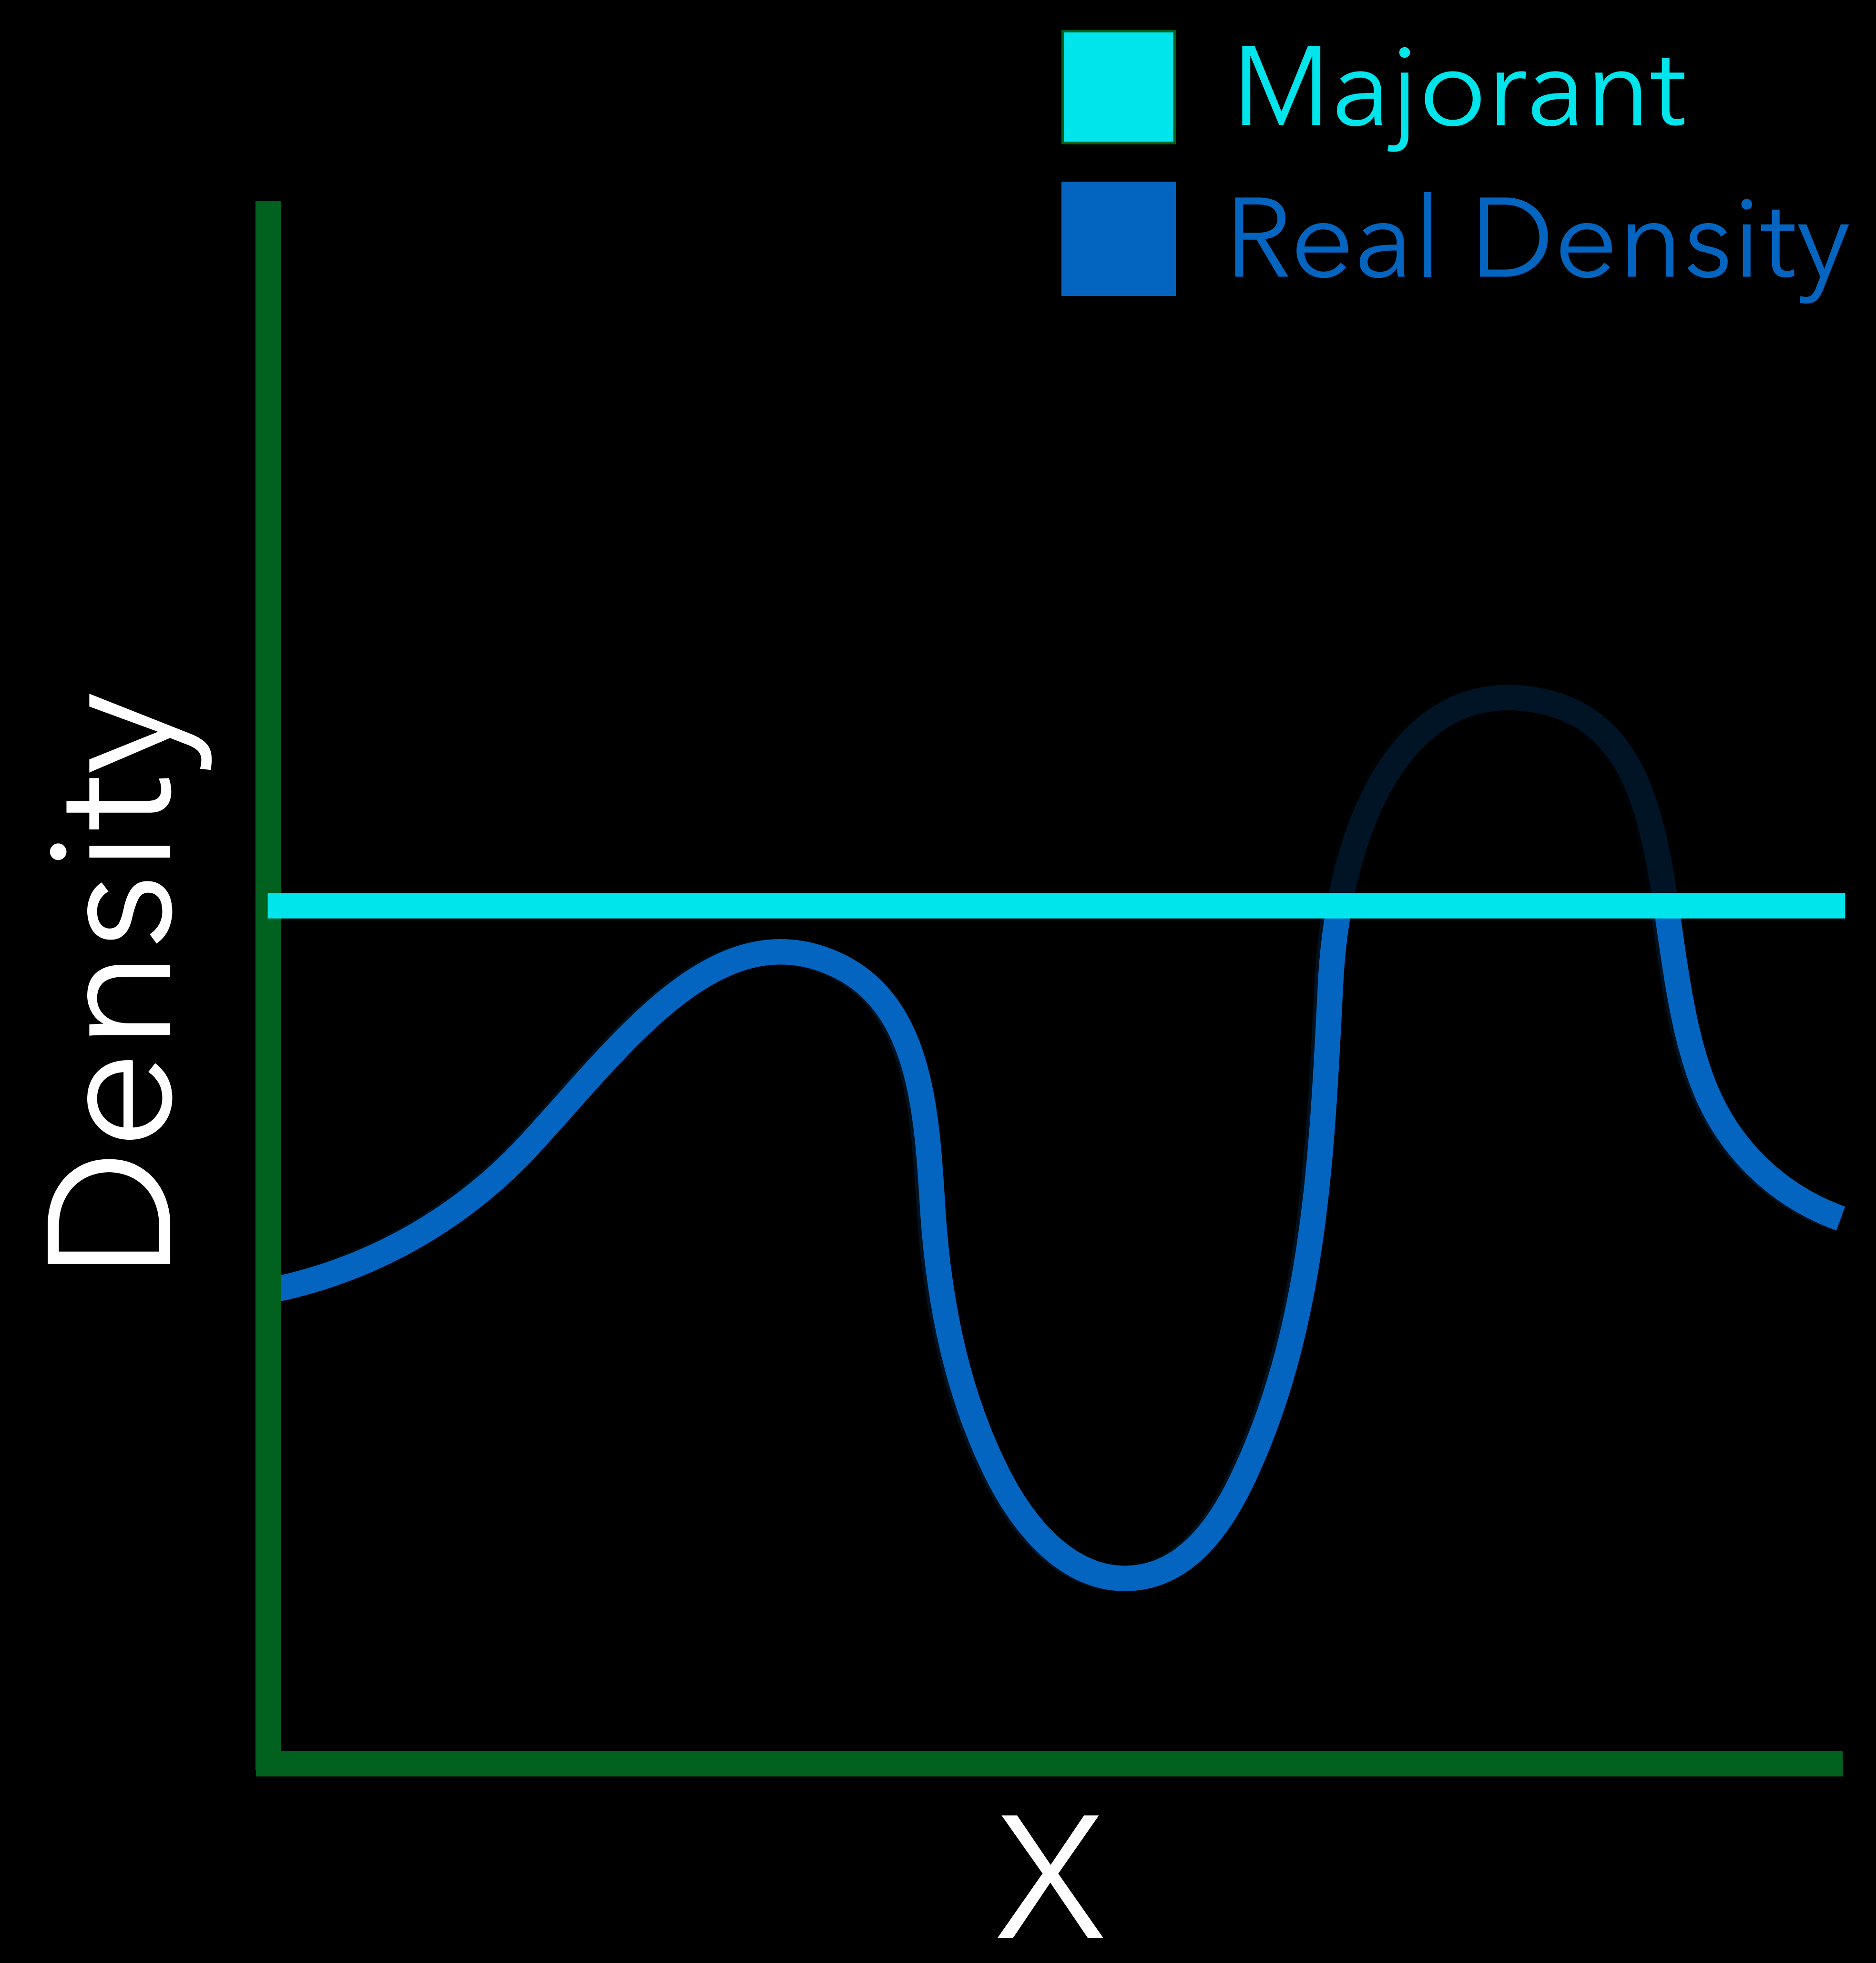
Majorant
Real Density
Density
X

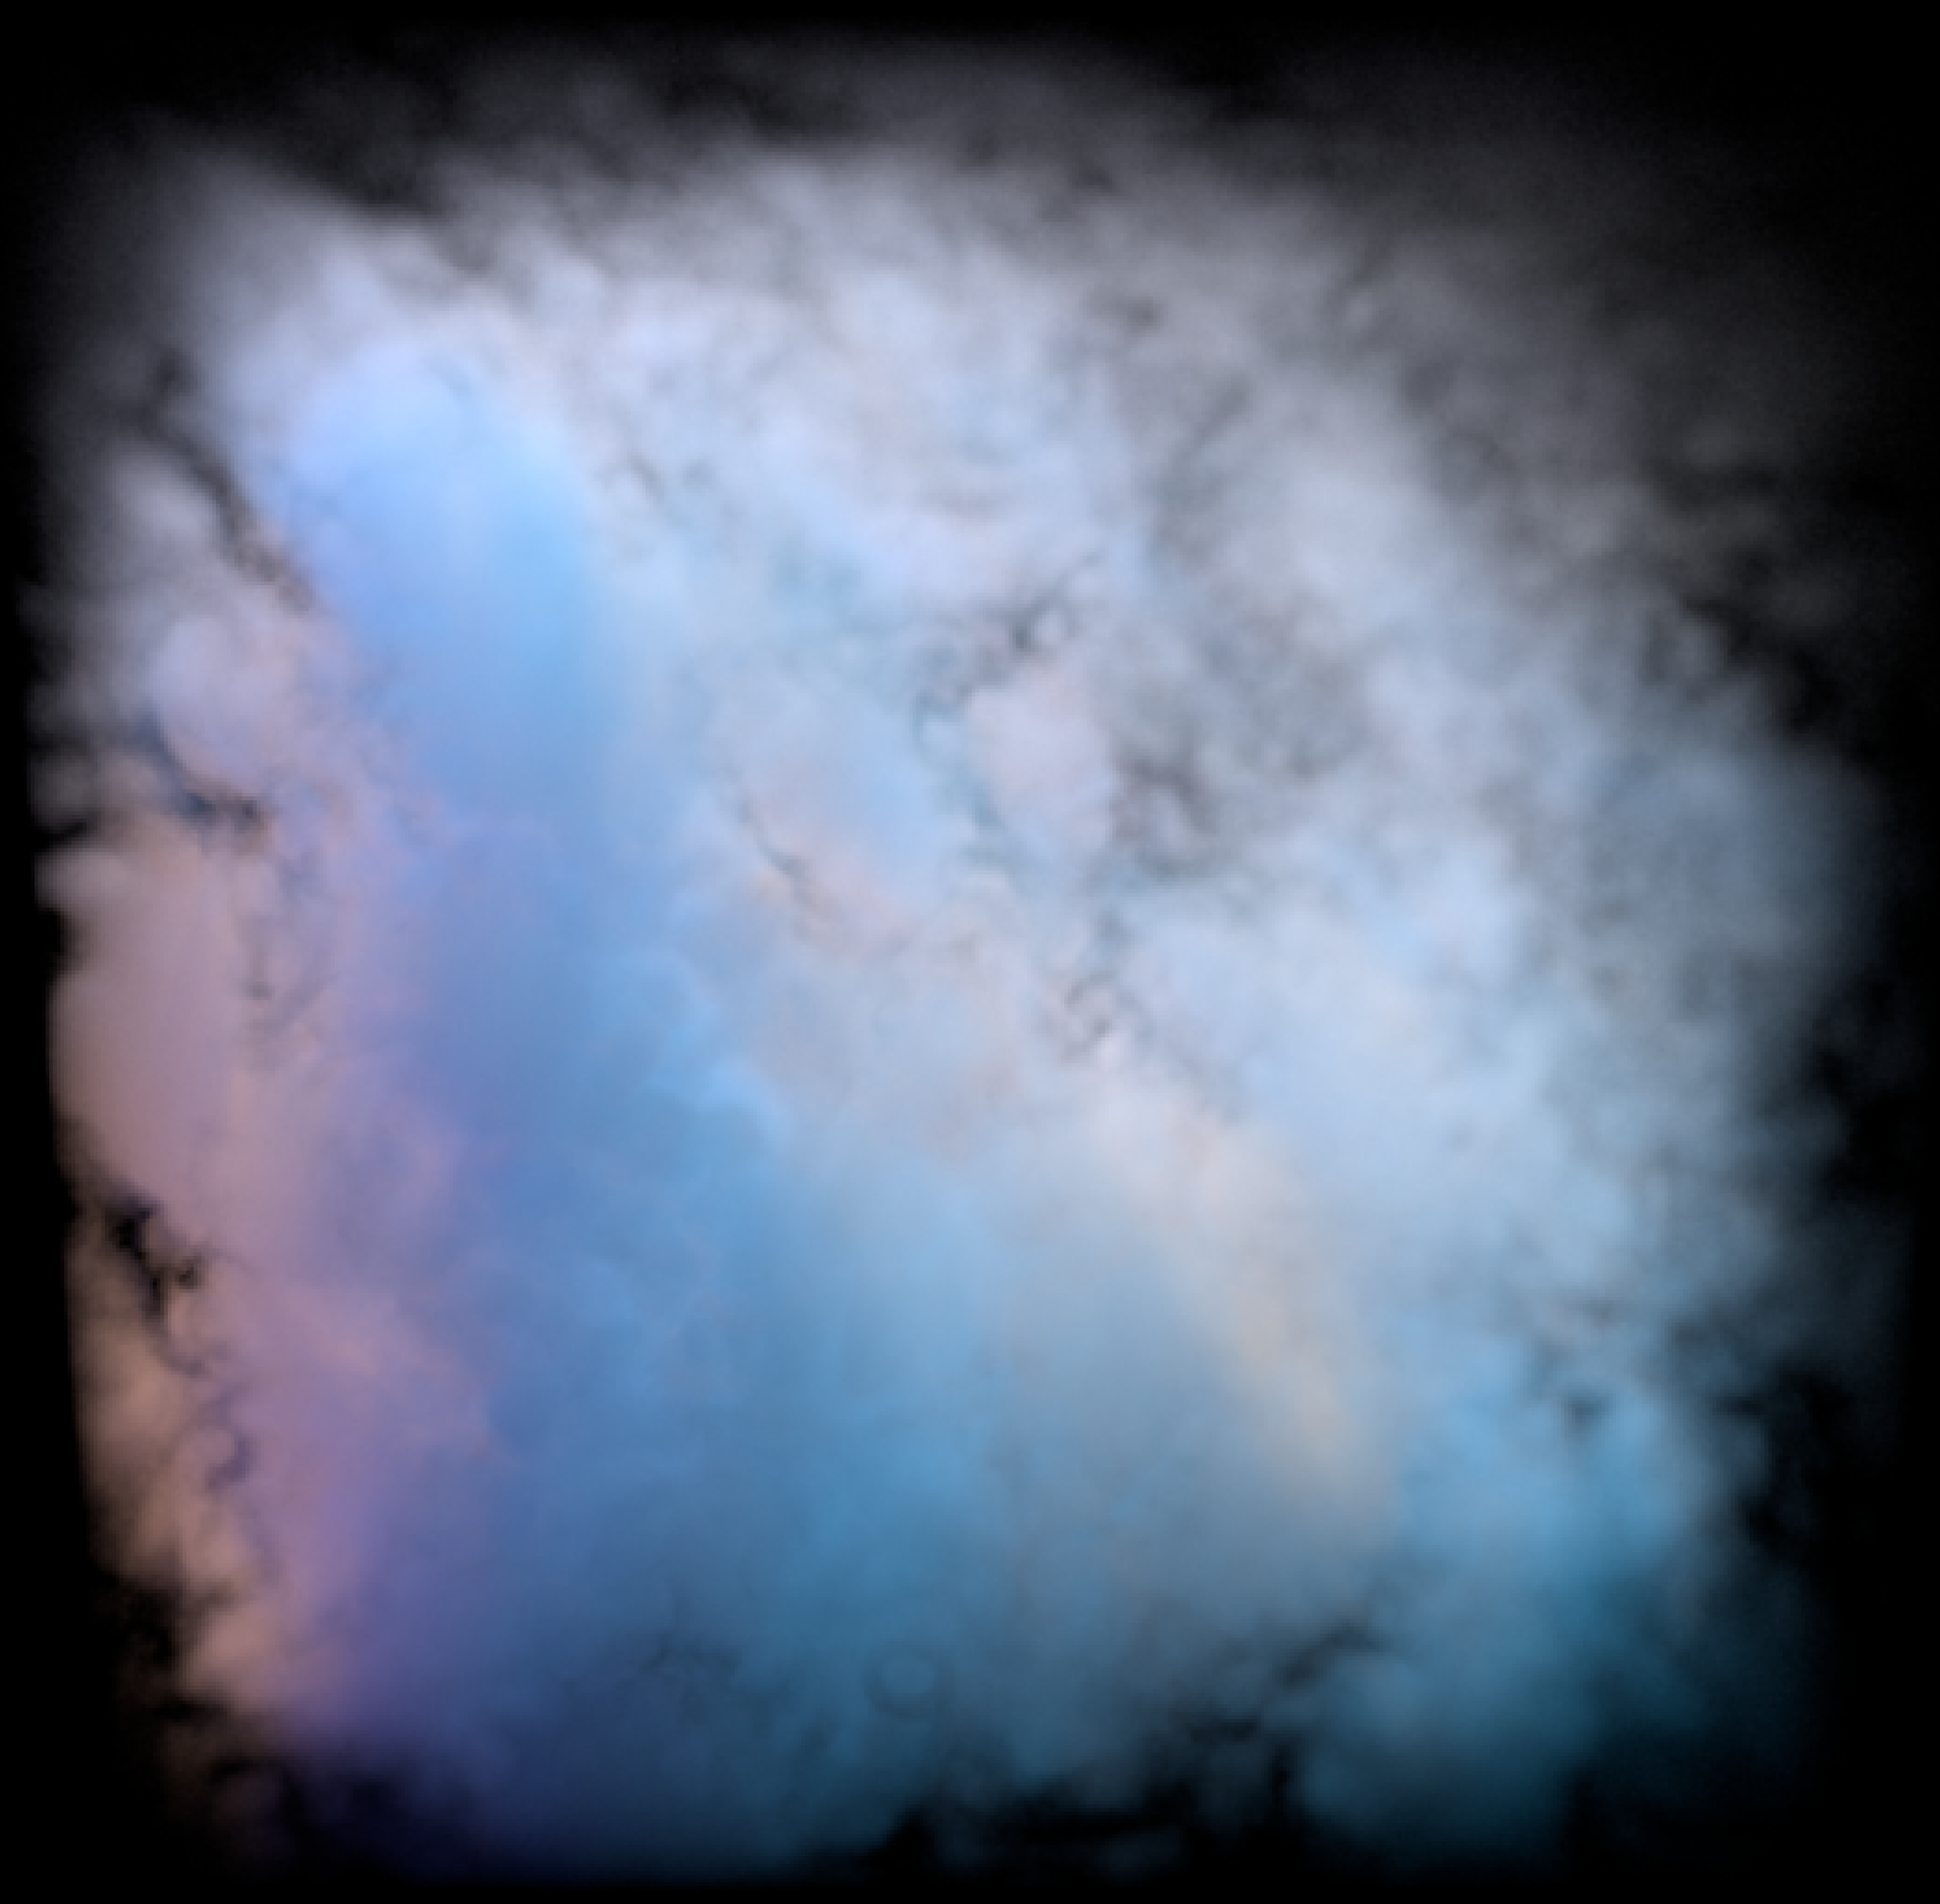

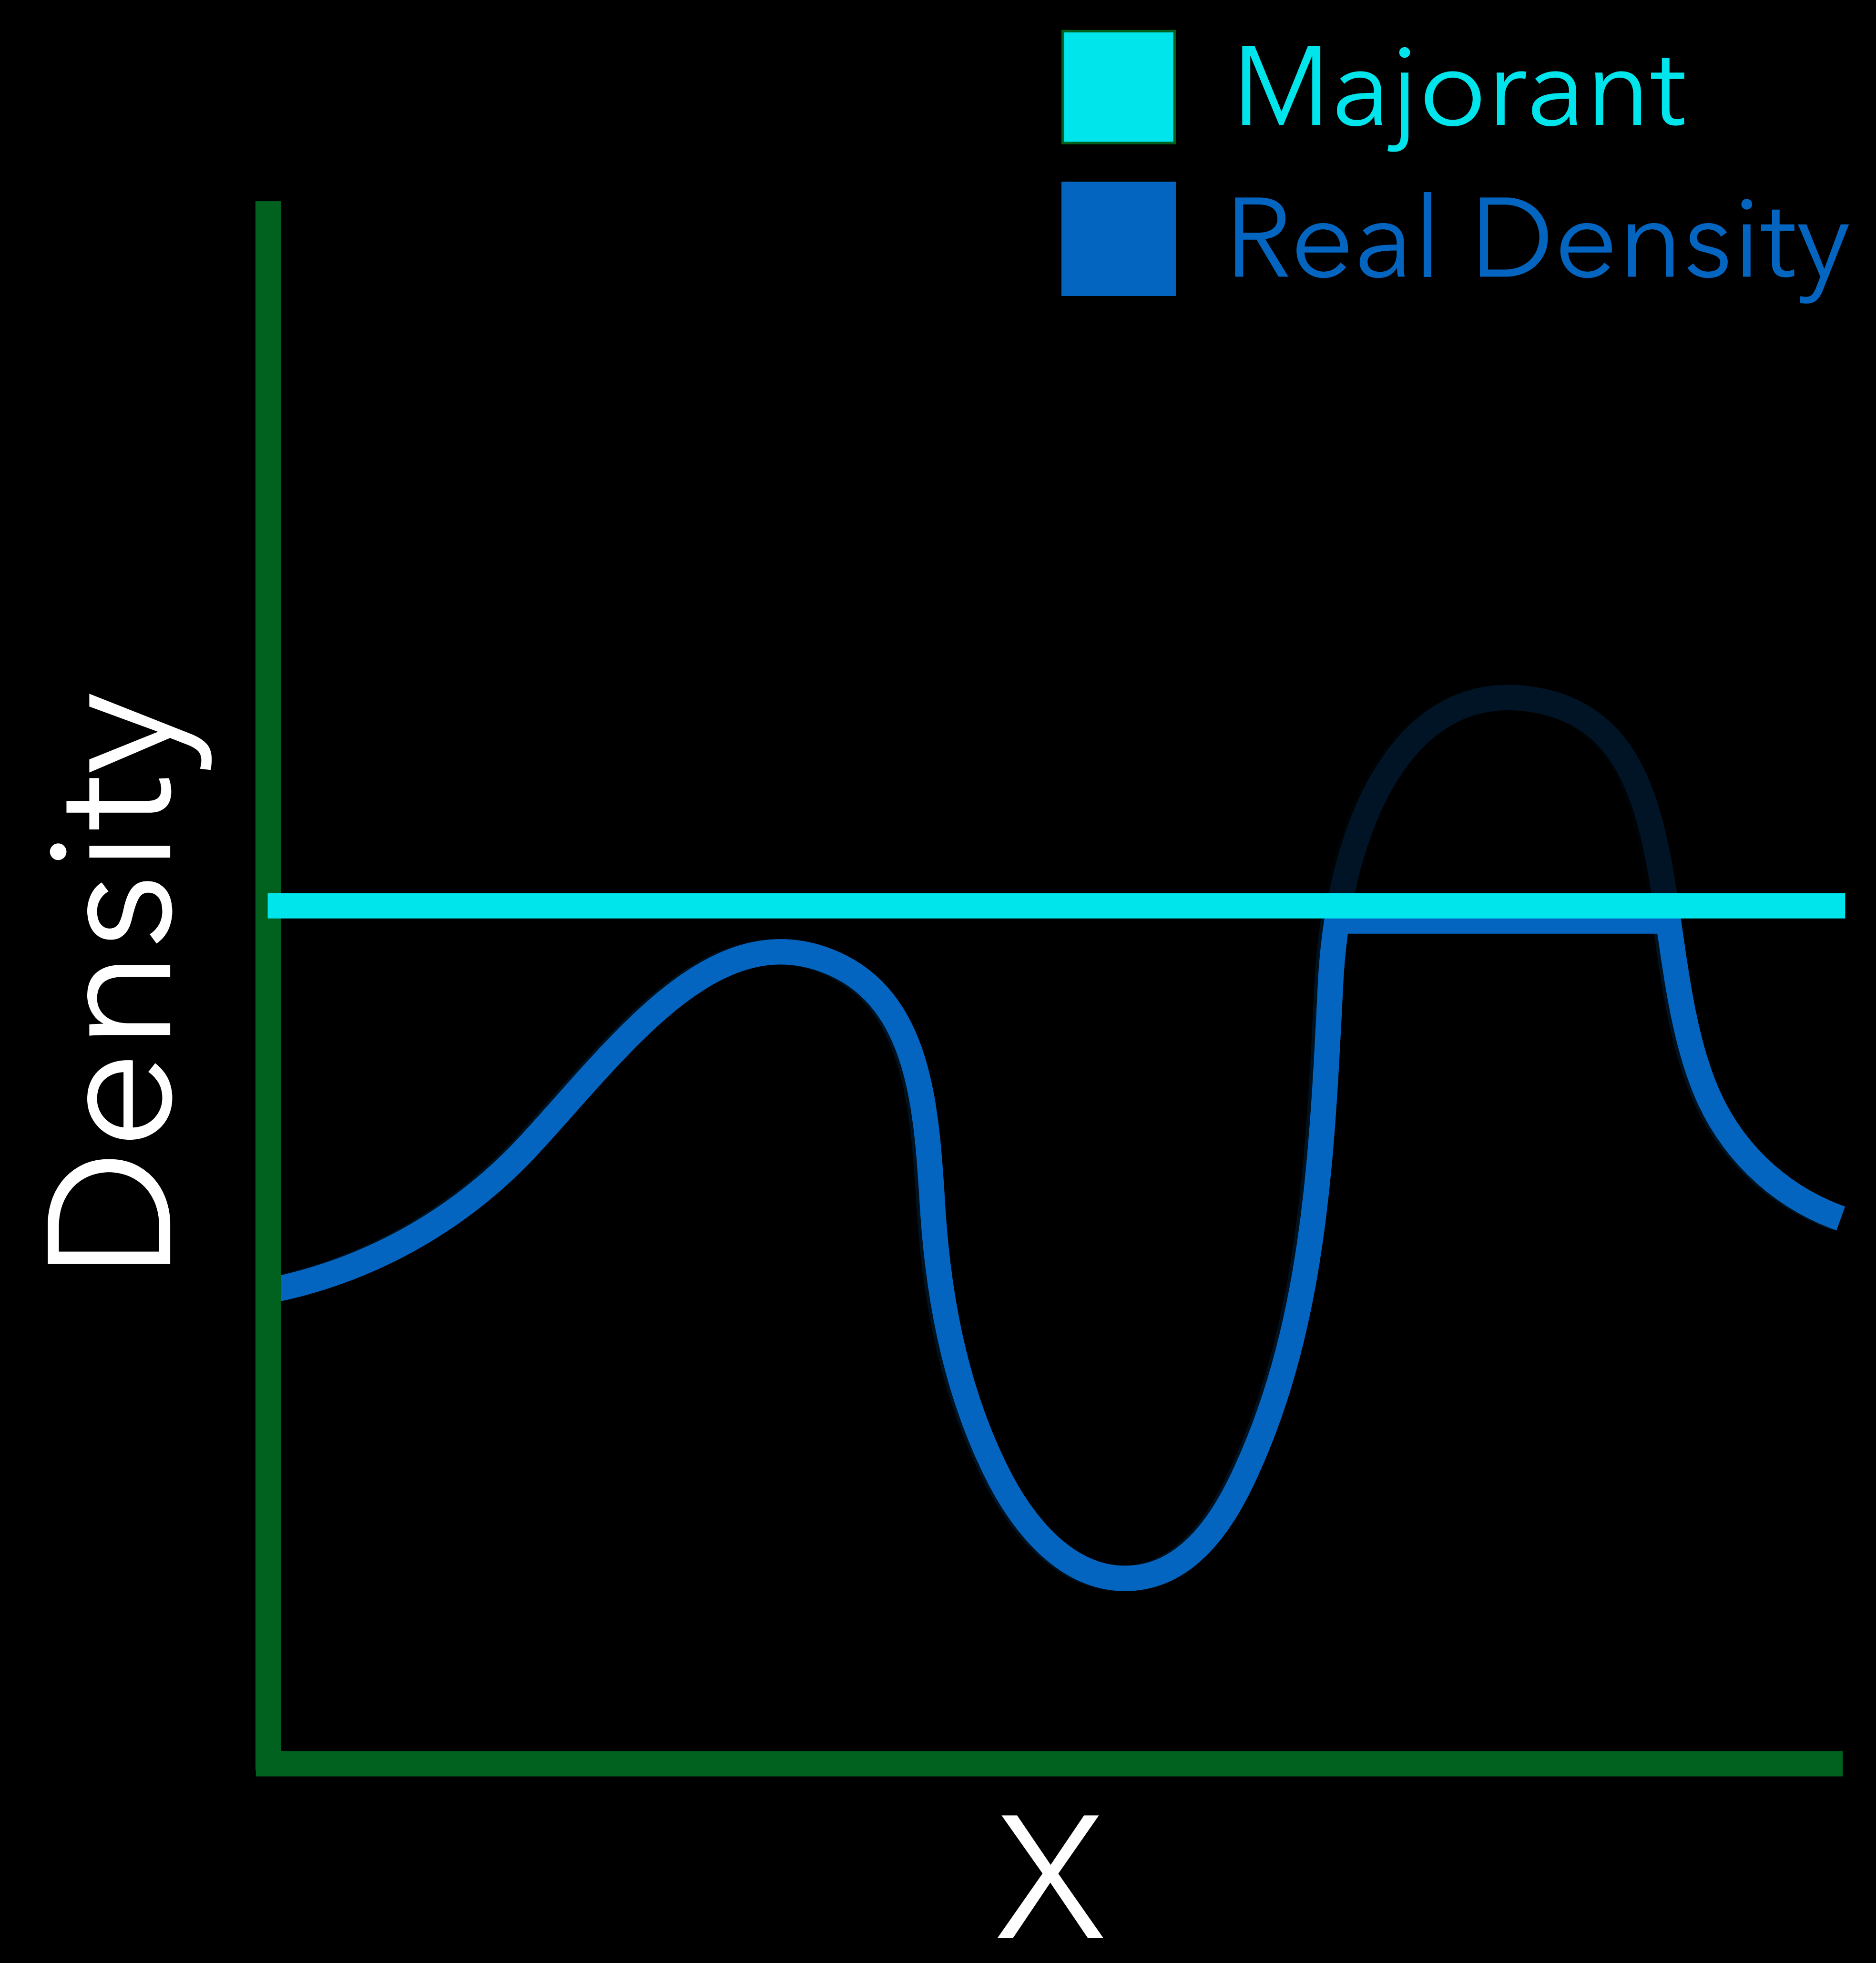
Majorant
Real Density
Density
X

Majorant increases from left to right

Majorant increases from left to right

Majorant increases from left to right

$$\sum_{k=1}^{n} \frac{I_k}{n}$$

$$\sum_{k=1}^{j} \frac{I_k}{n} + \sum_{k=j+1}^{n} \frac{I_k}{n}$$

$$\boxed{\sum_{k=1}^{j} \frac{I_k}{n}} + \sum_{k=j+1}^{n} \frac{I_k}{n}$$

Biased

$$\sum_{k=1}^{j} \frac{I_k}{n} + \boxed{\sum_{k=j+1}^{n} \frac{I_k}{n}}$$

Biased Unbiased

$$\lim_{n\to\infty}\left[\sum_{k=1}^{j}\frac{I_k}{n}+\sum_{k=j+1}^{n}\frac{I_k}{n}\right]$$

Biased Unbiased

$$\lim_{n\to\infty}\left[\quad 0 \quad + \sum_{k=j+1}^{n} \frac{I_k}{n}\right]$$

Biased Unbiased

Progressively update majorants

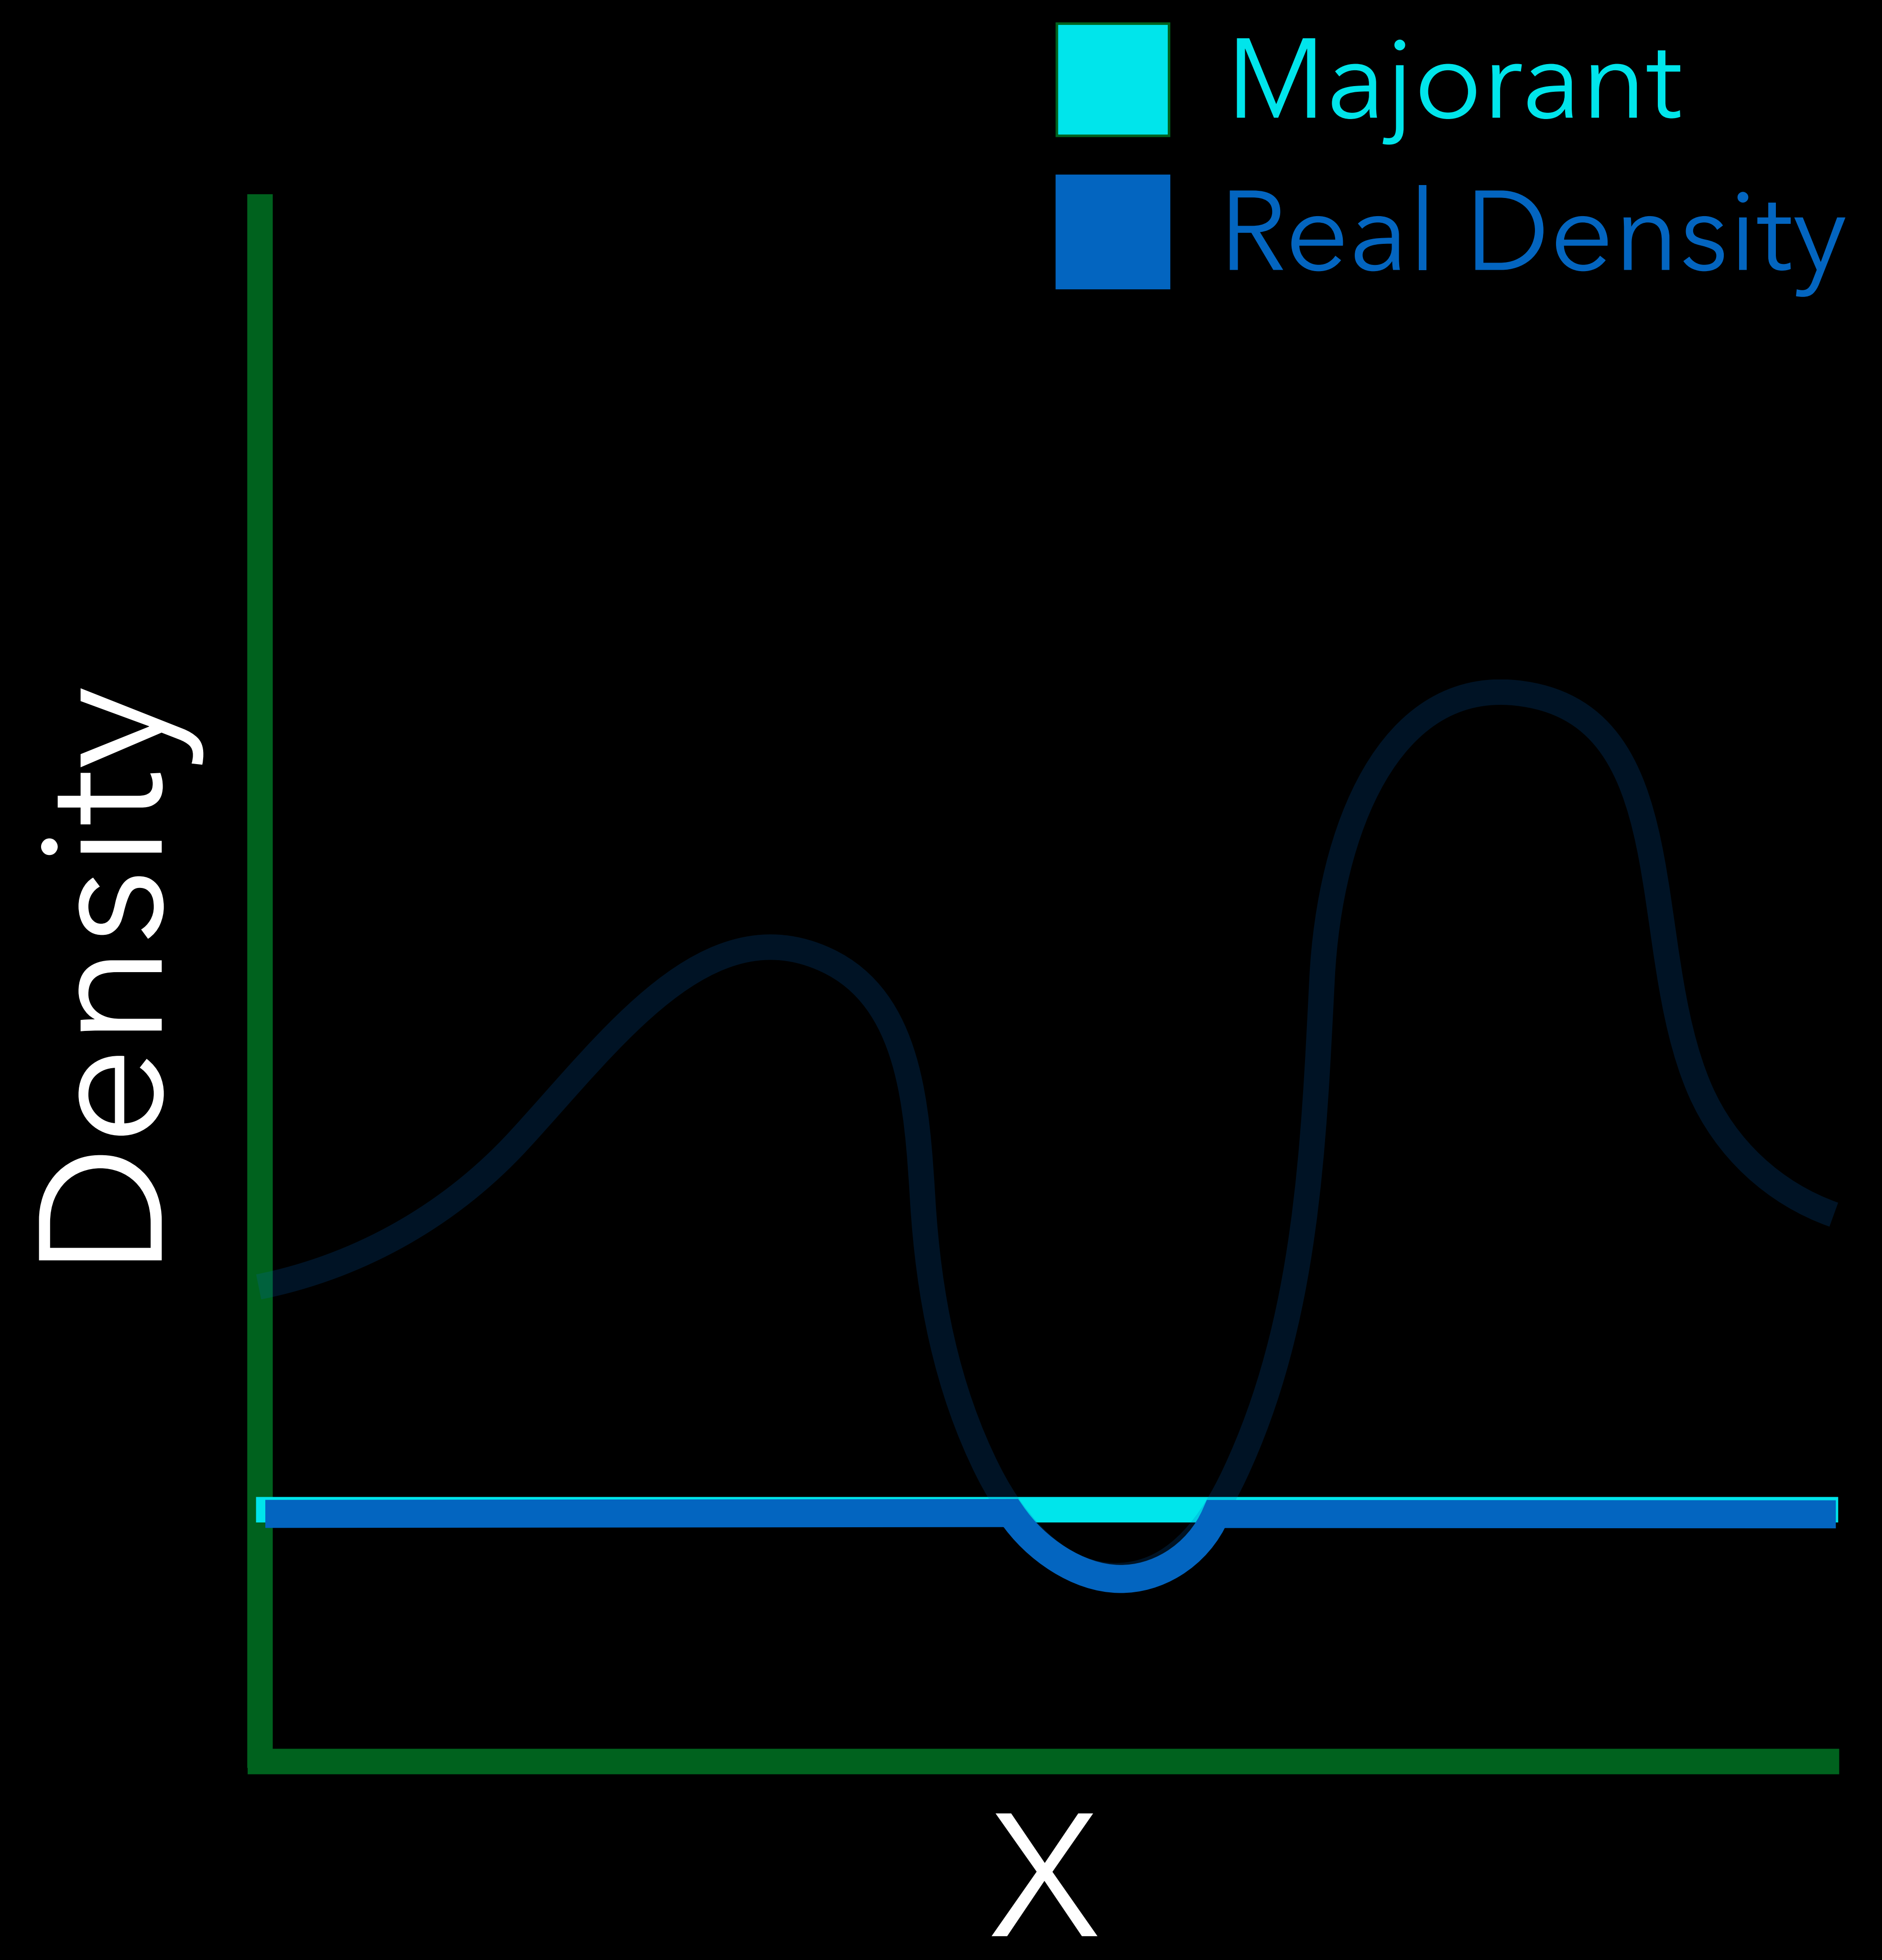
Majorant
Real Density
Density
X

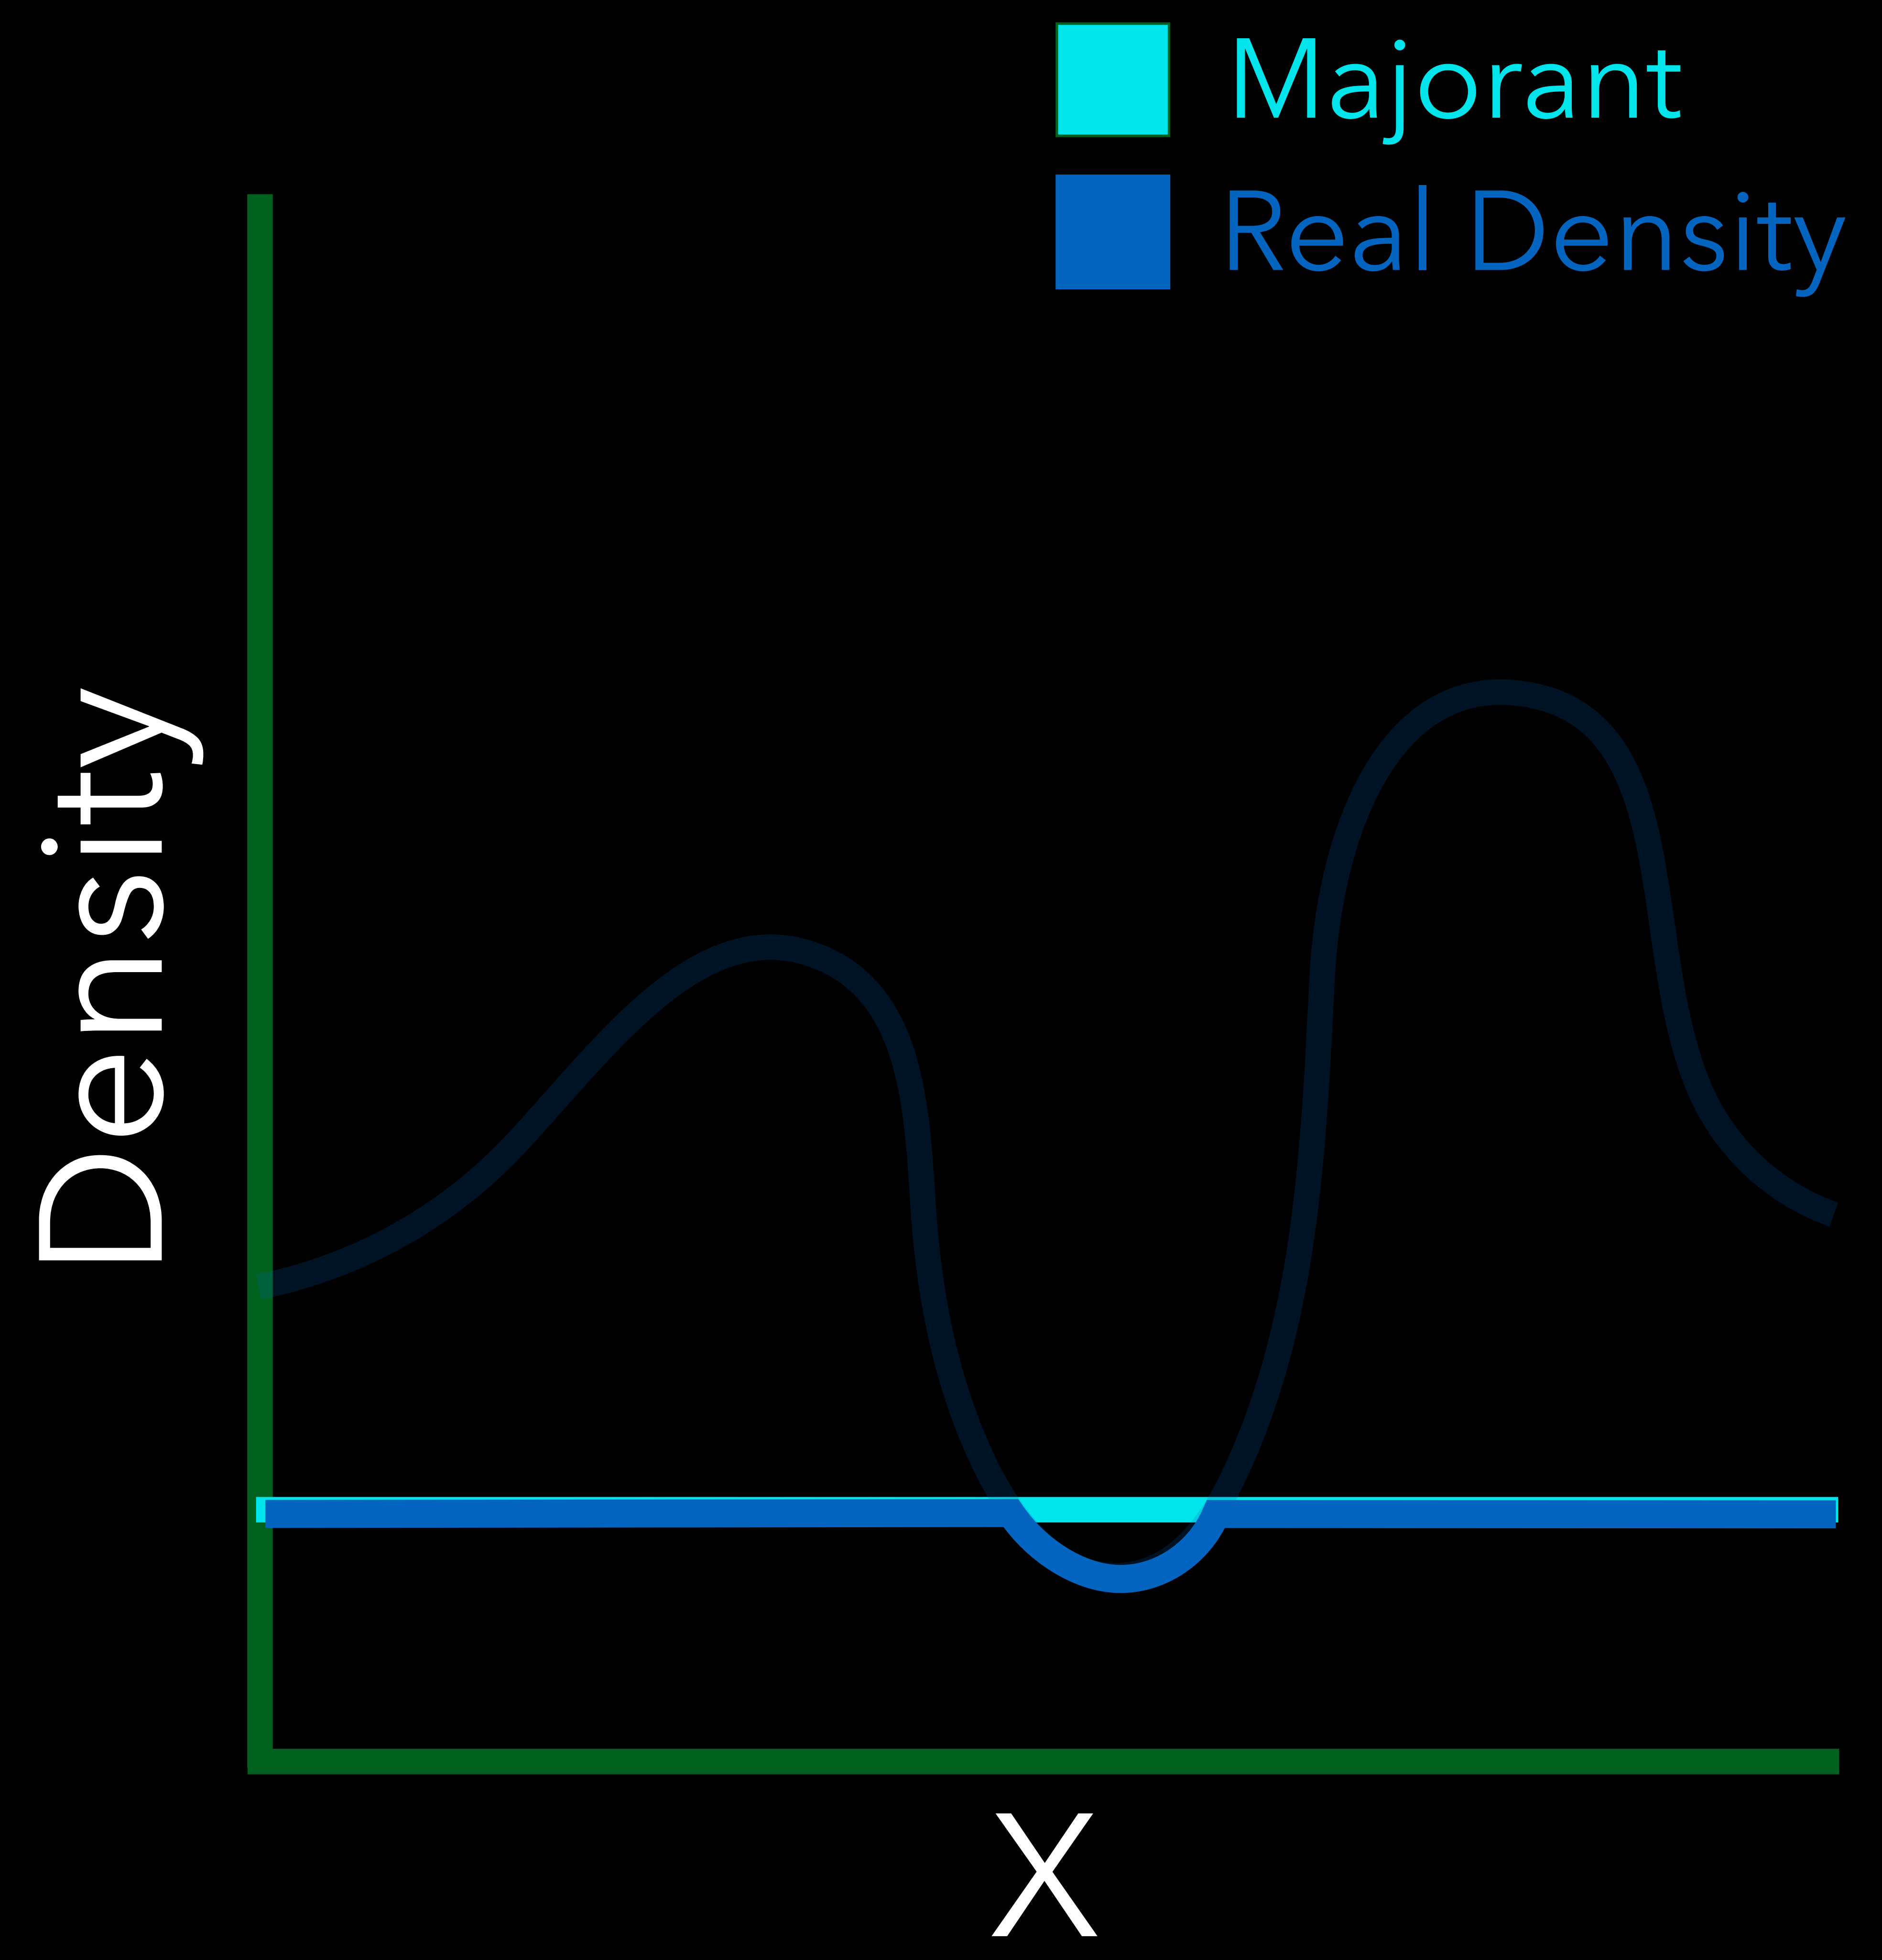
Majorant
Real Density
Density
X

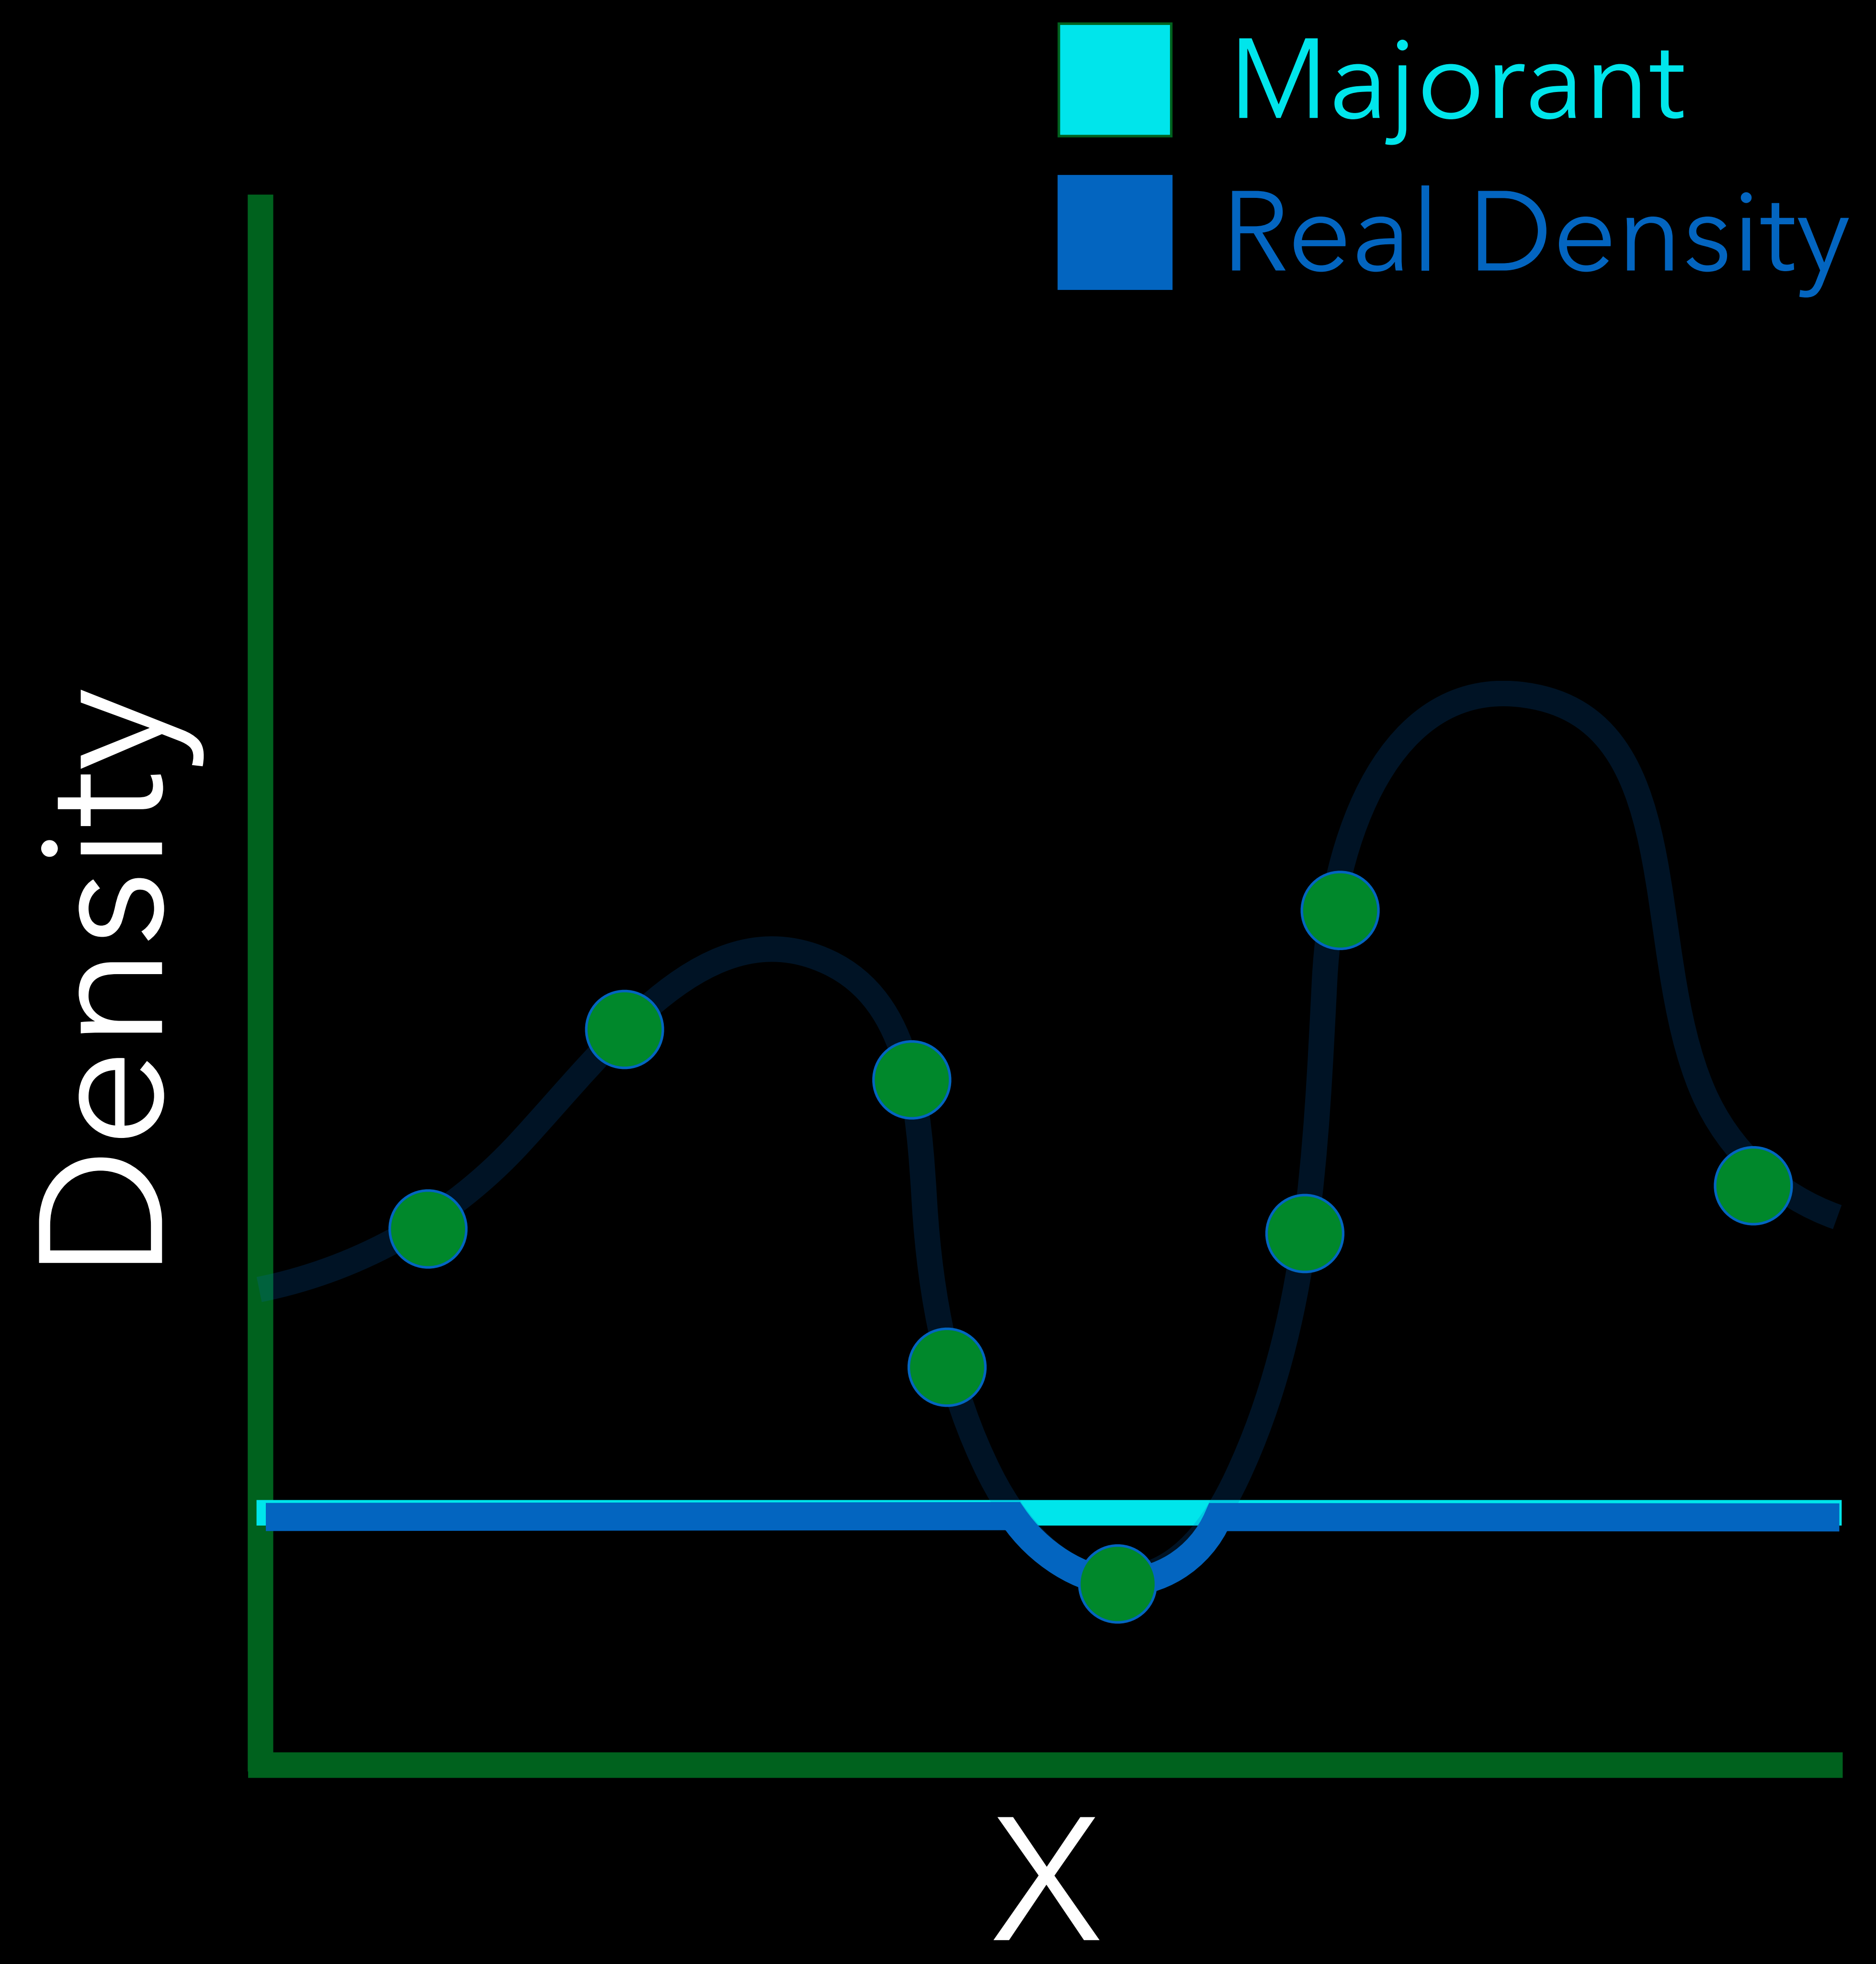
Majorant
Real Density
Density
X

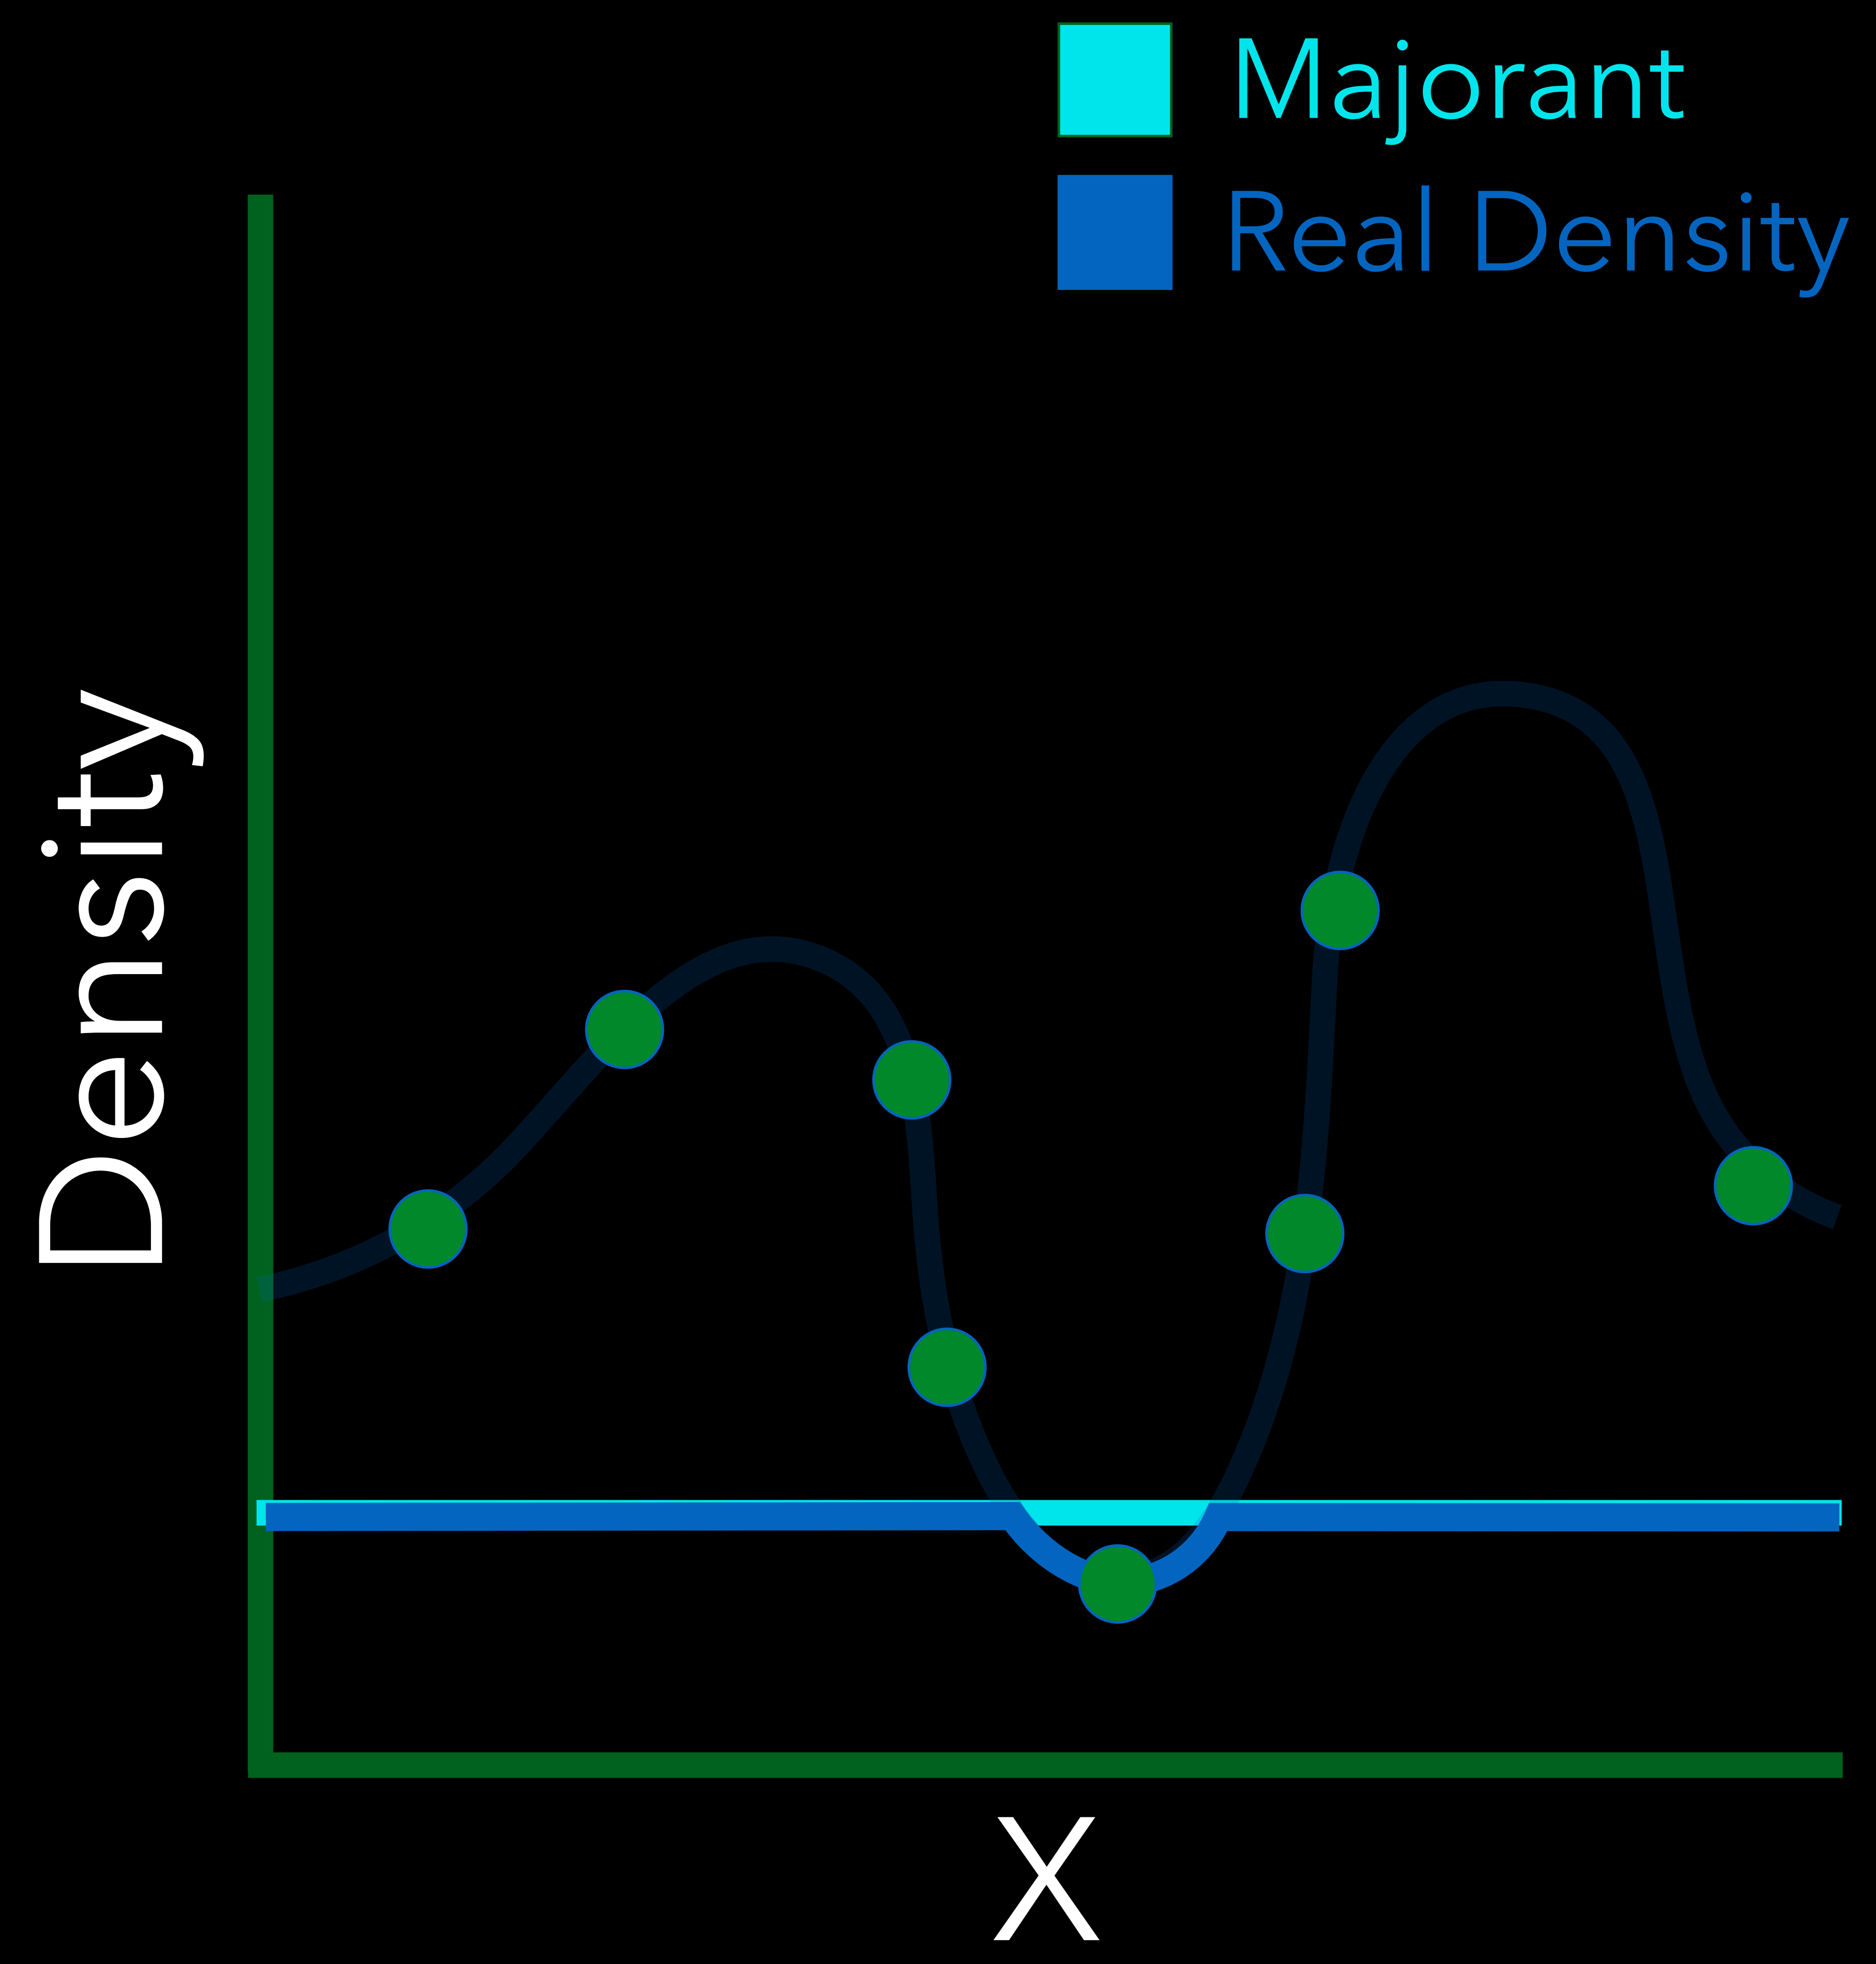
Majorant
Real Density
Density
X

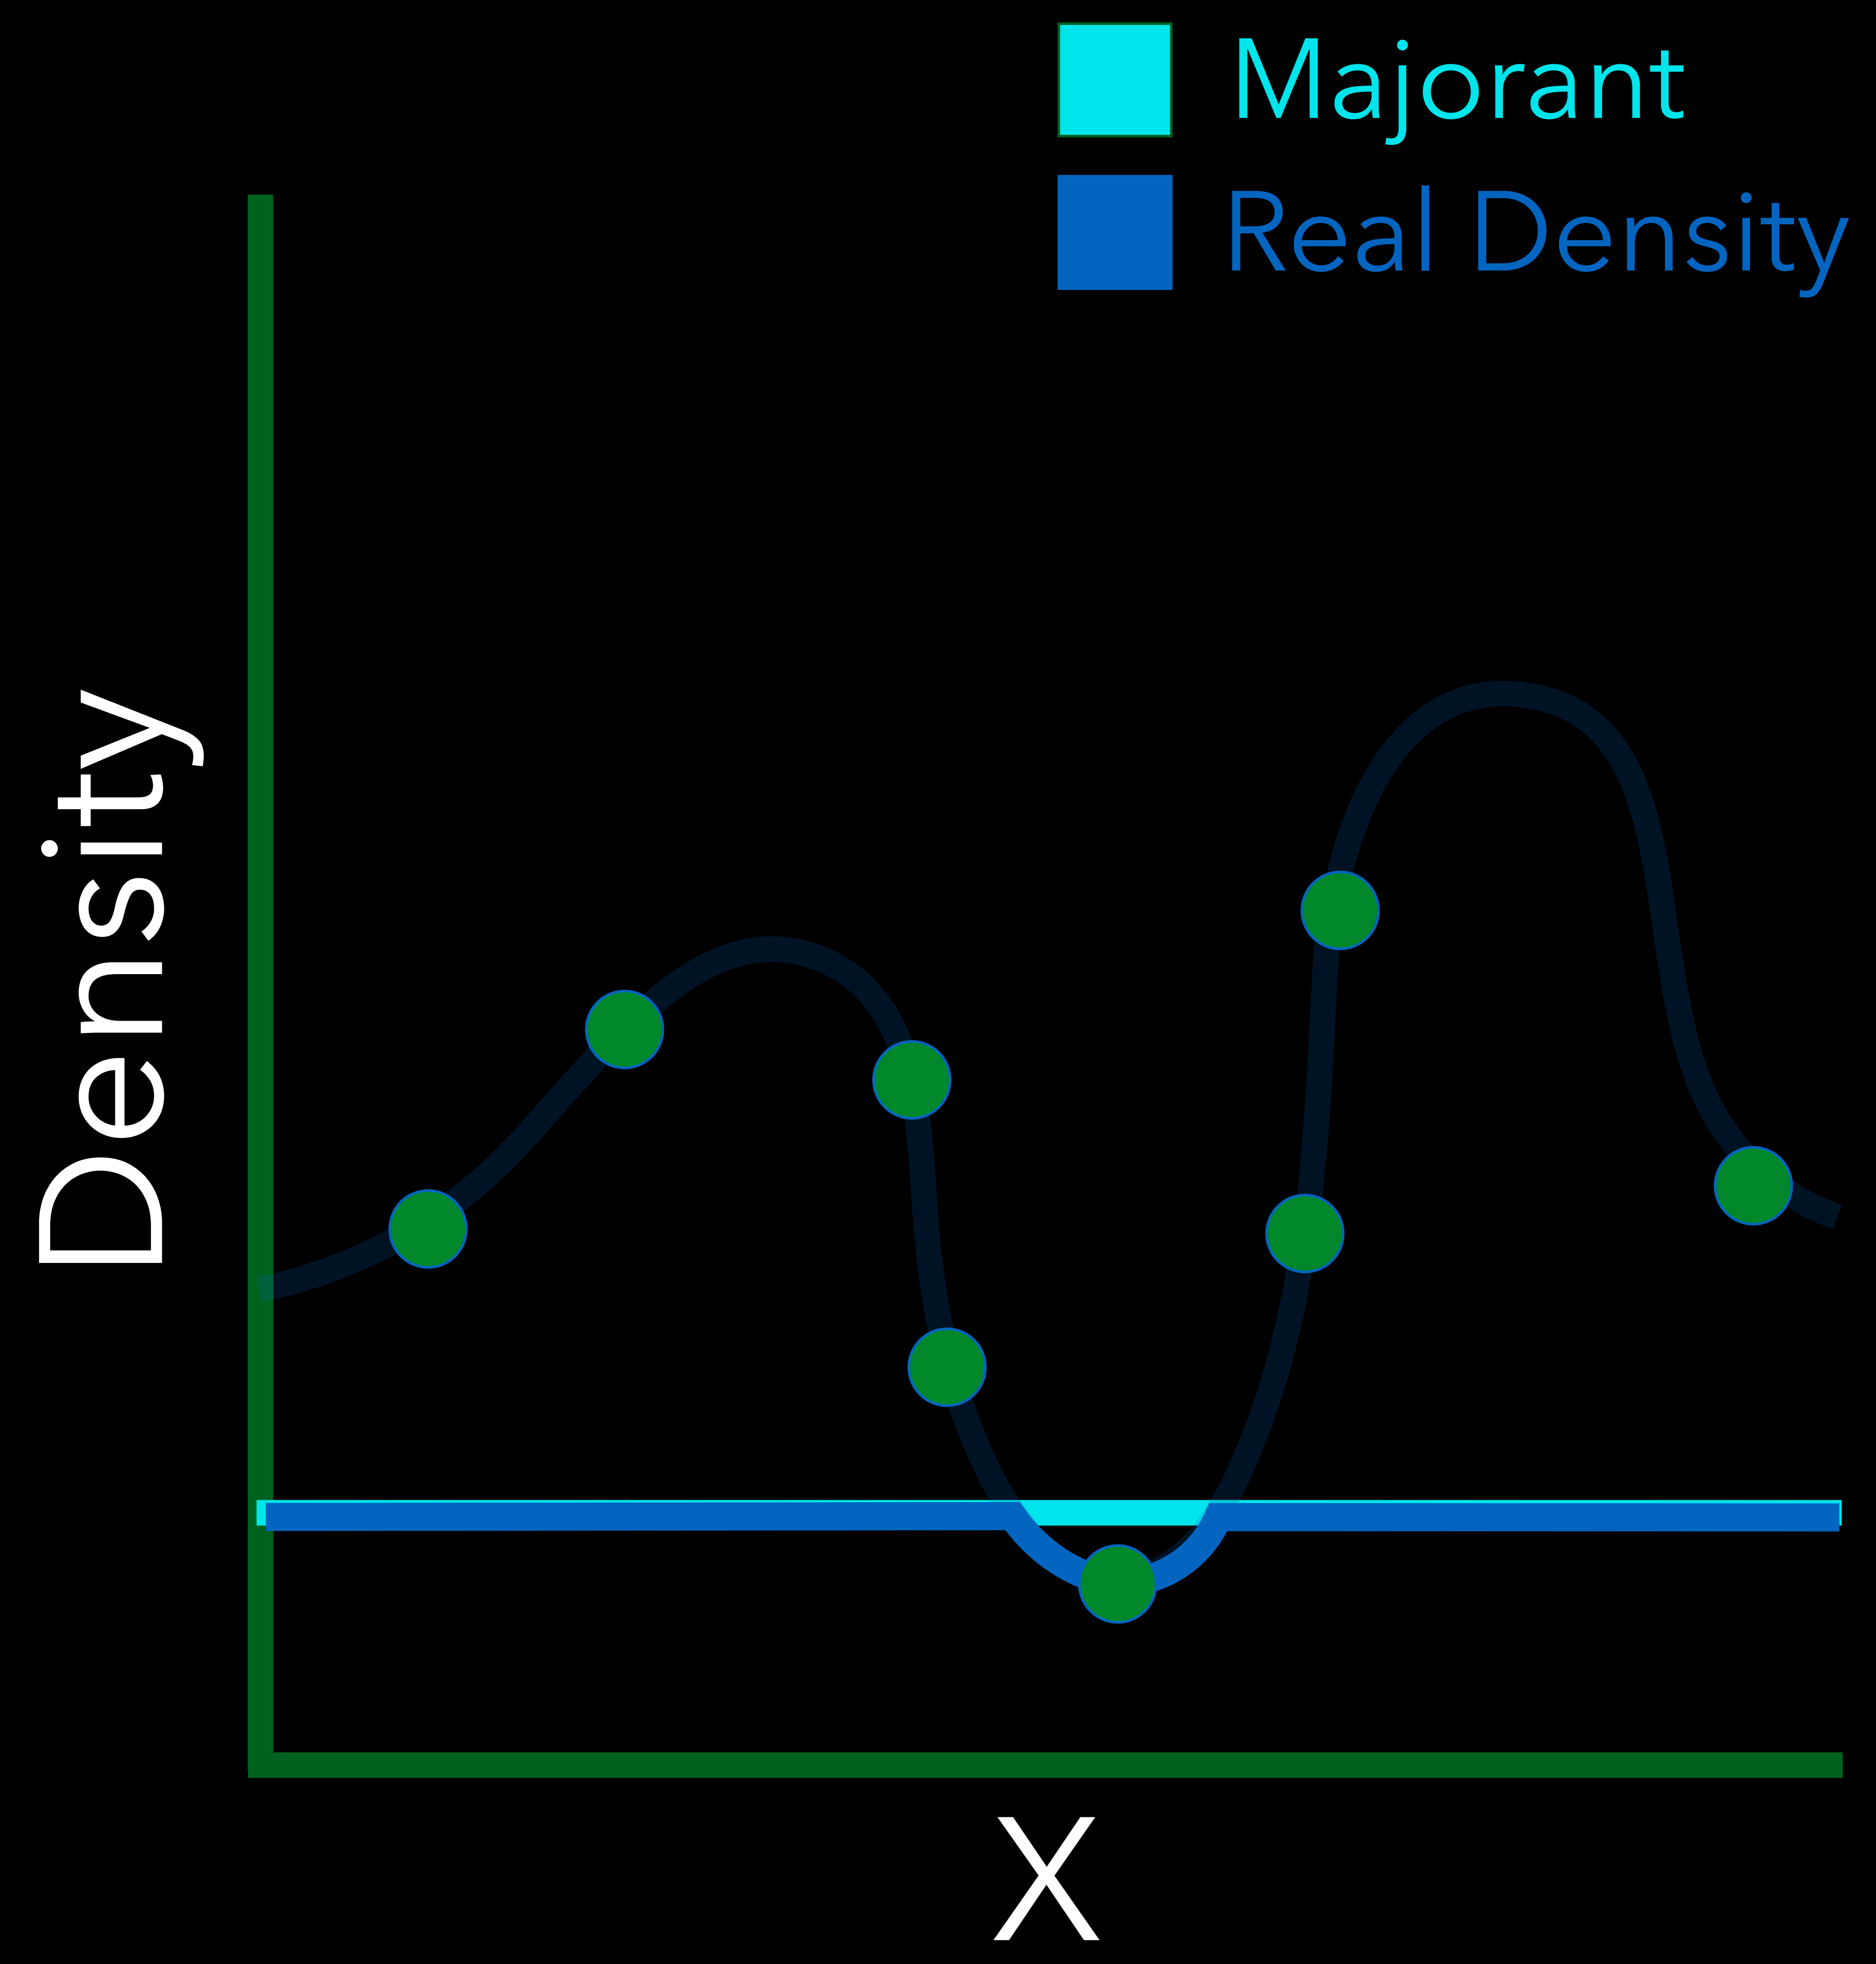
Majorant
Real Density
Density
X

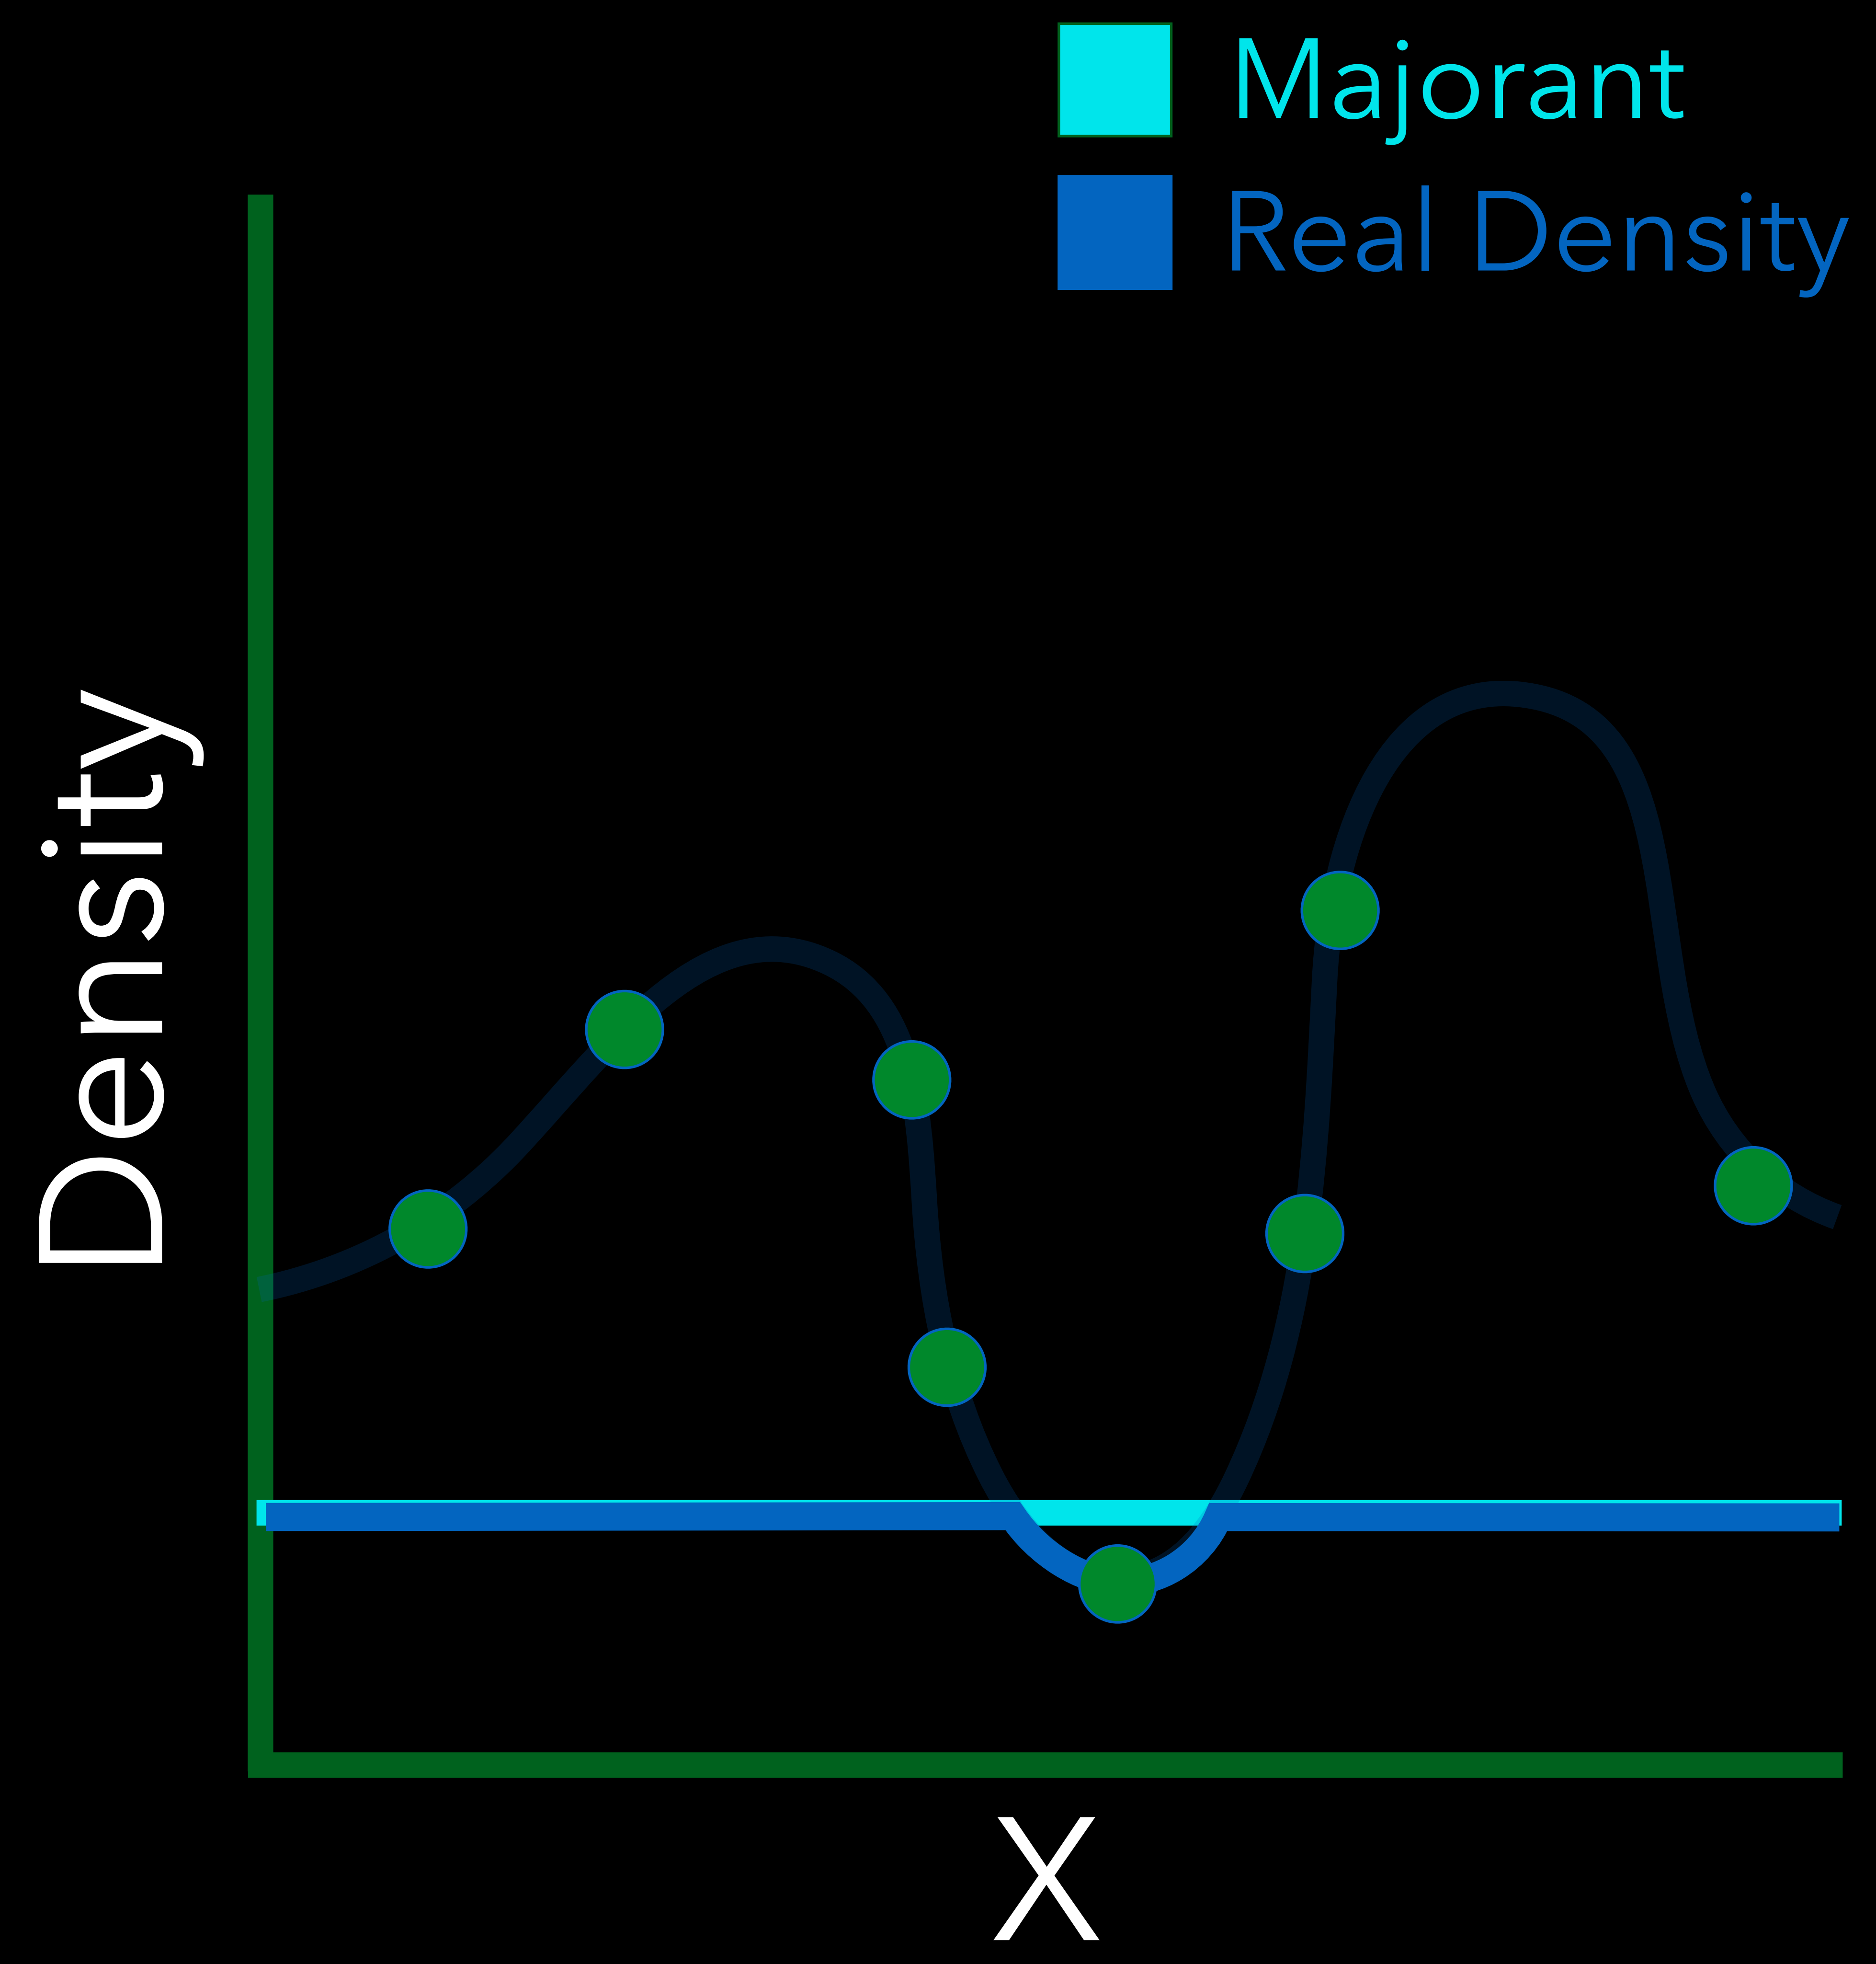

Majorant
Real Density
Density
X

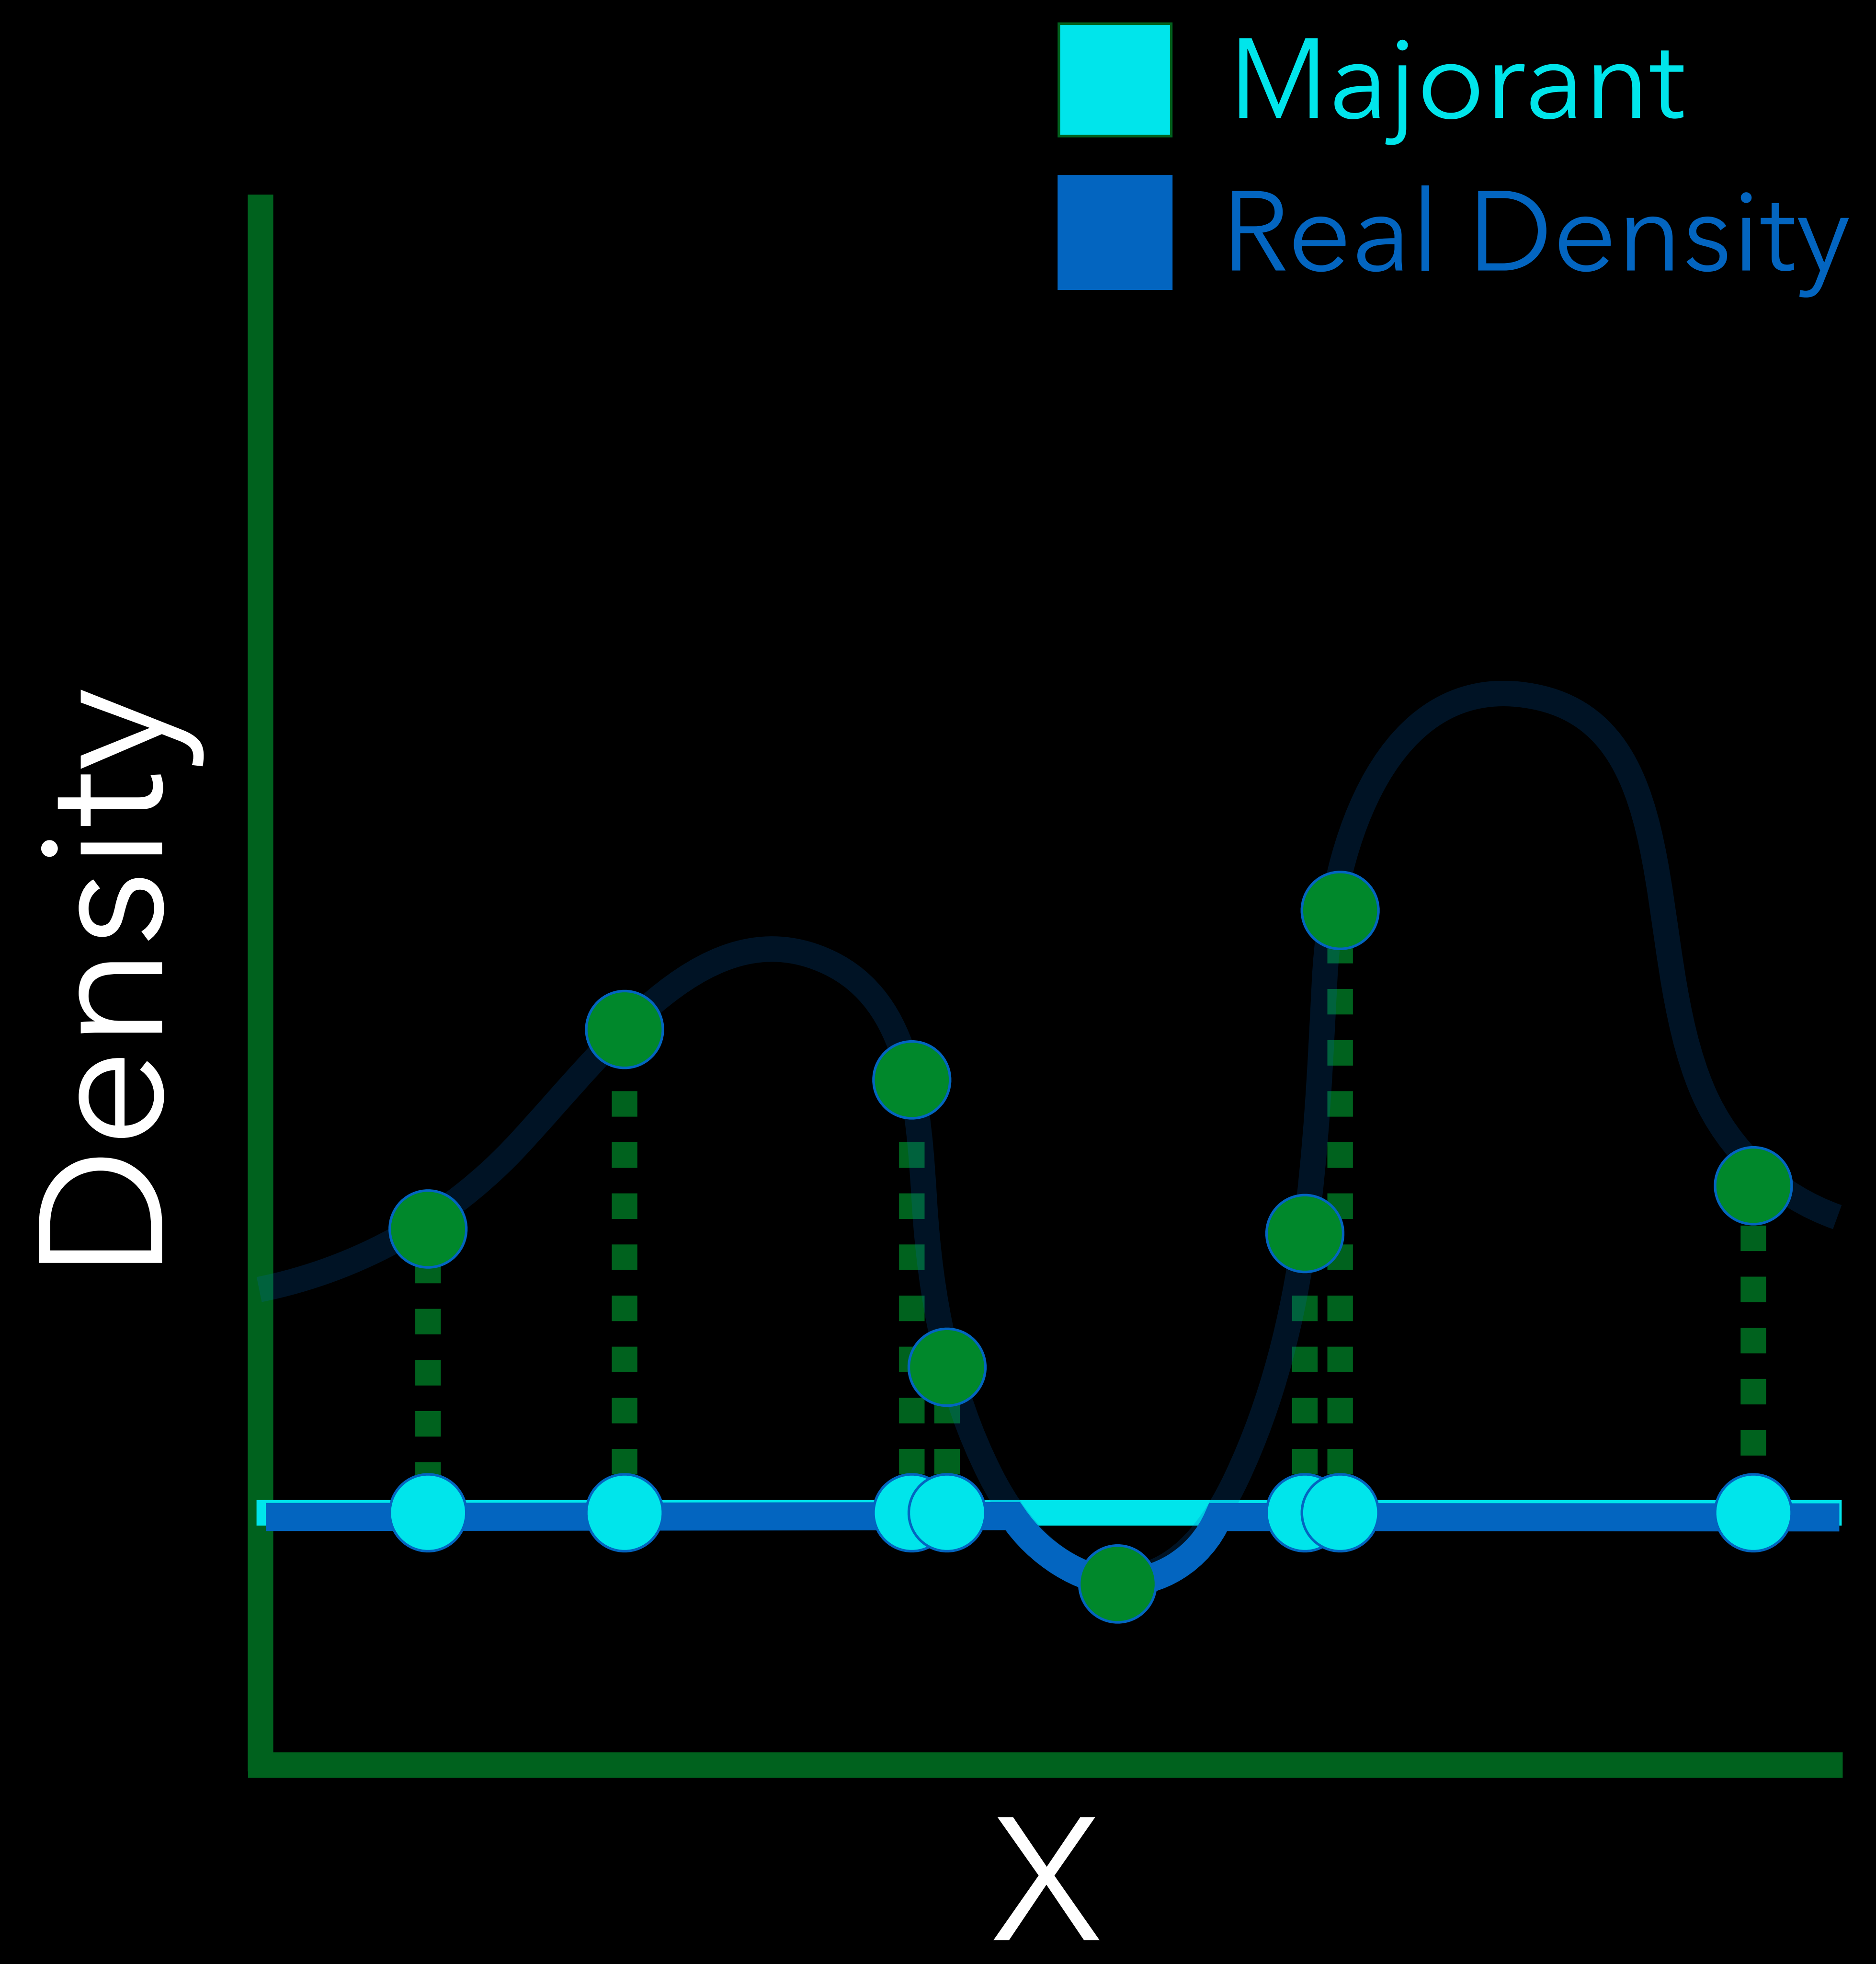
Majorant
Real Density
Density
X

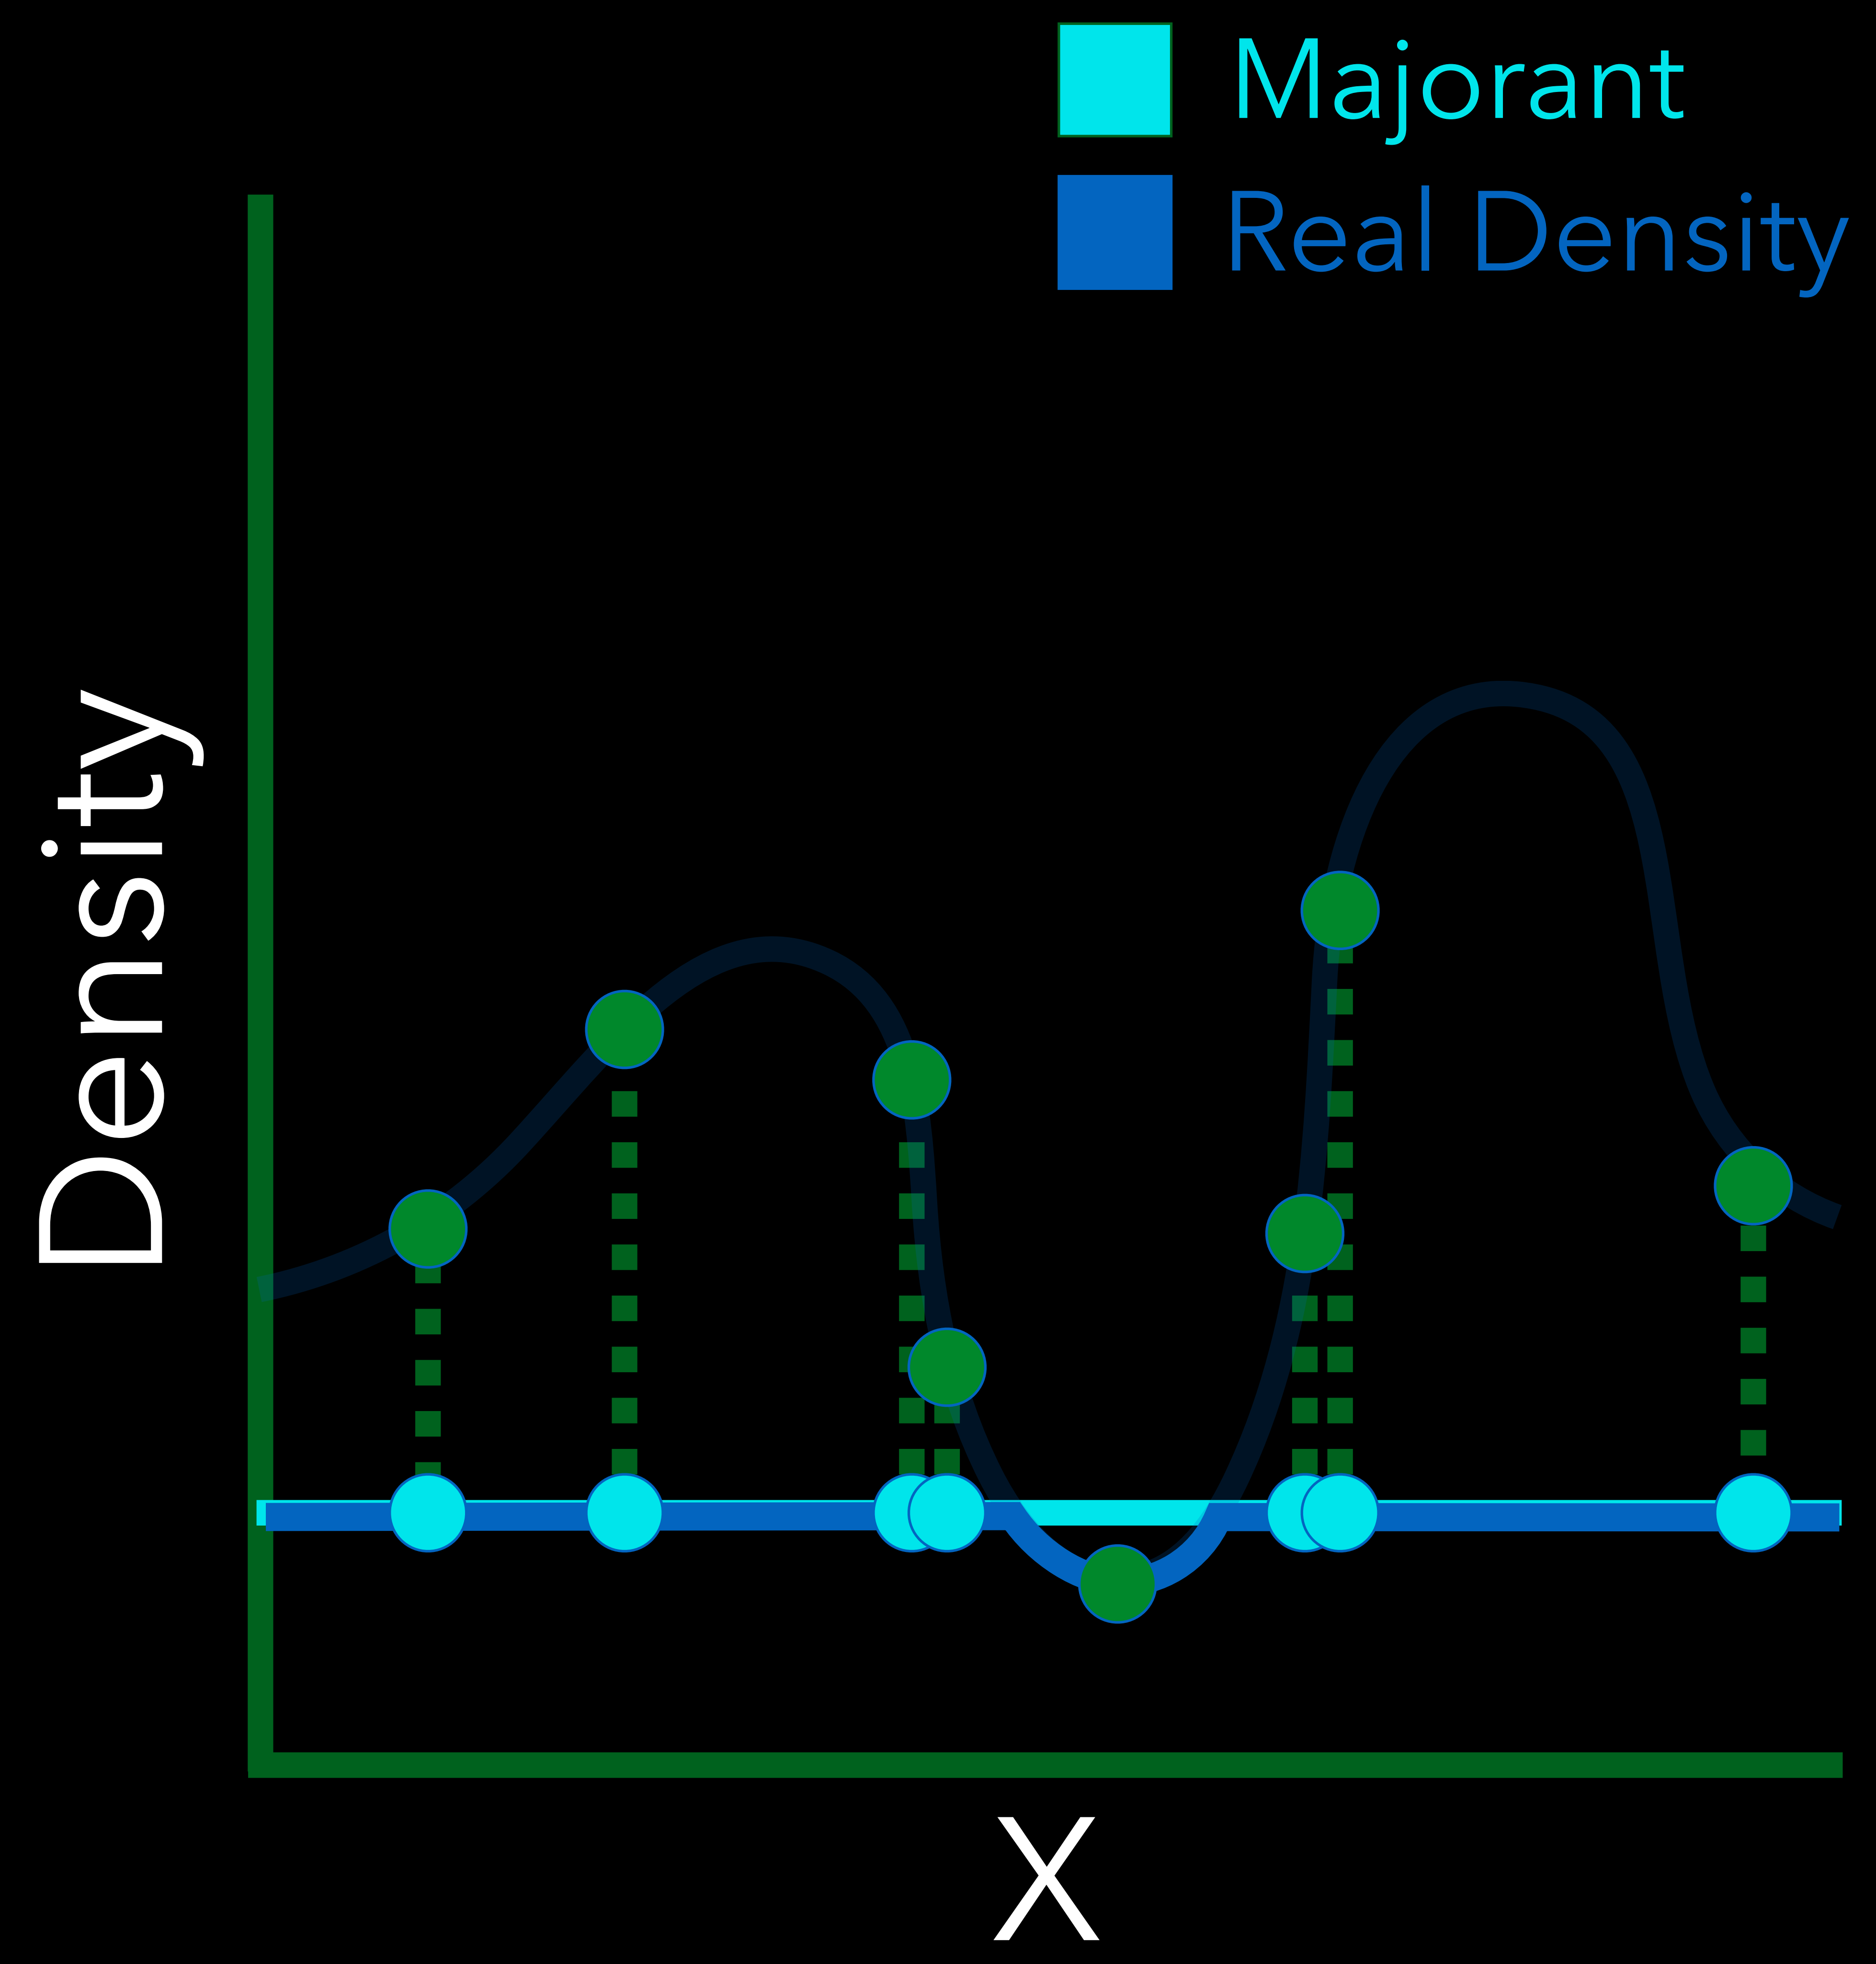
Majorant
Real Density
Density
X

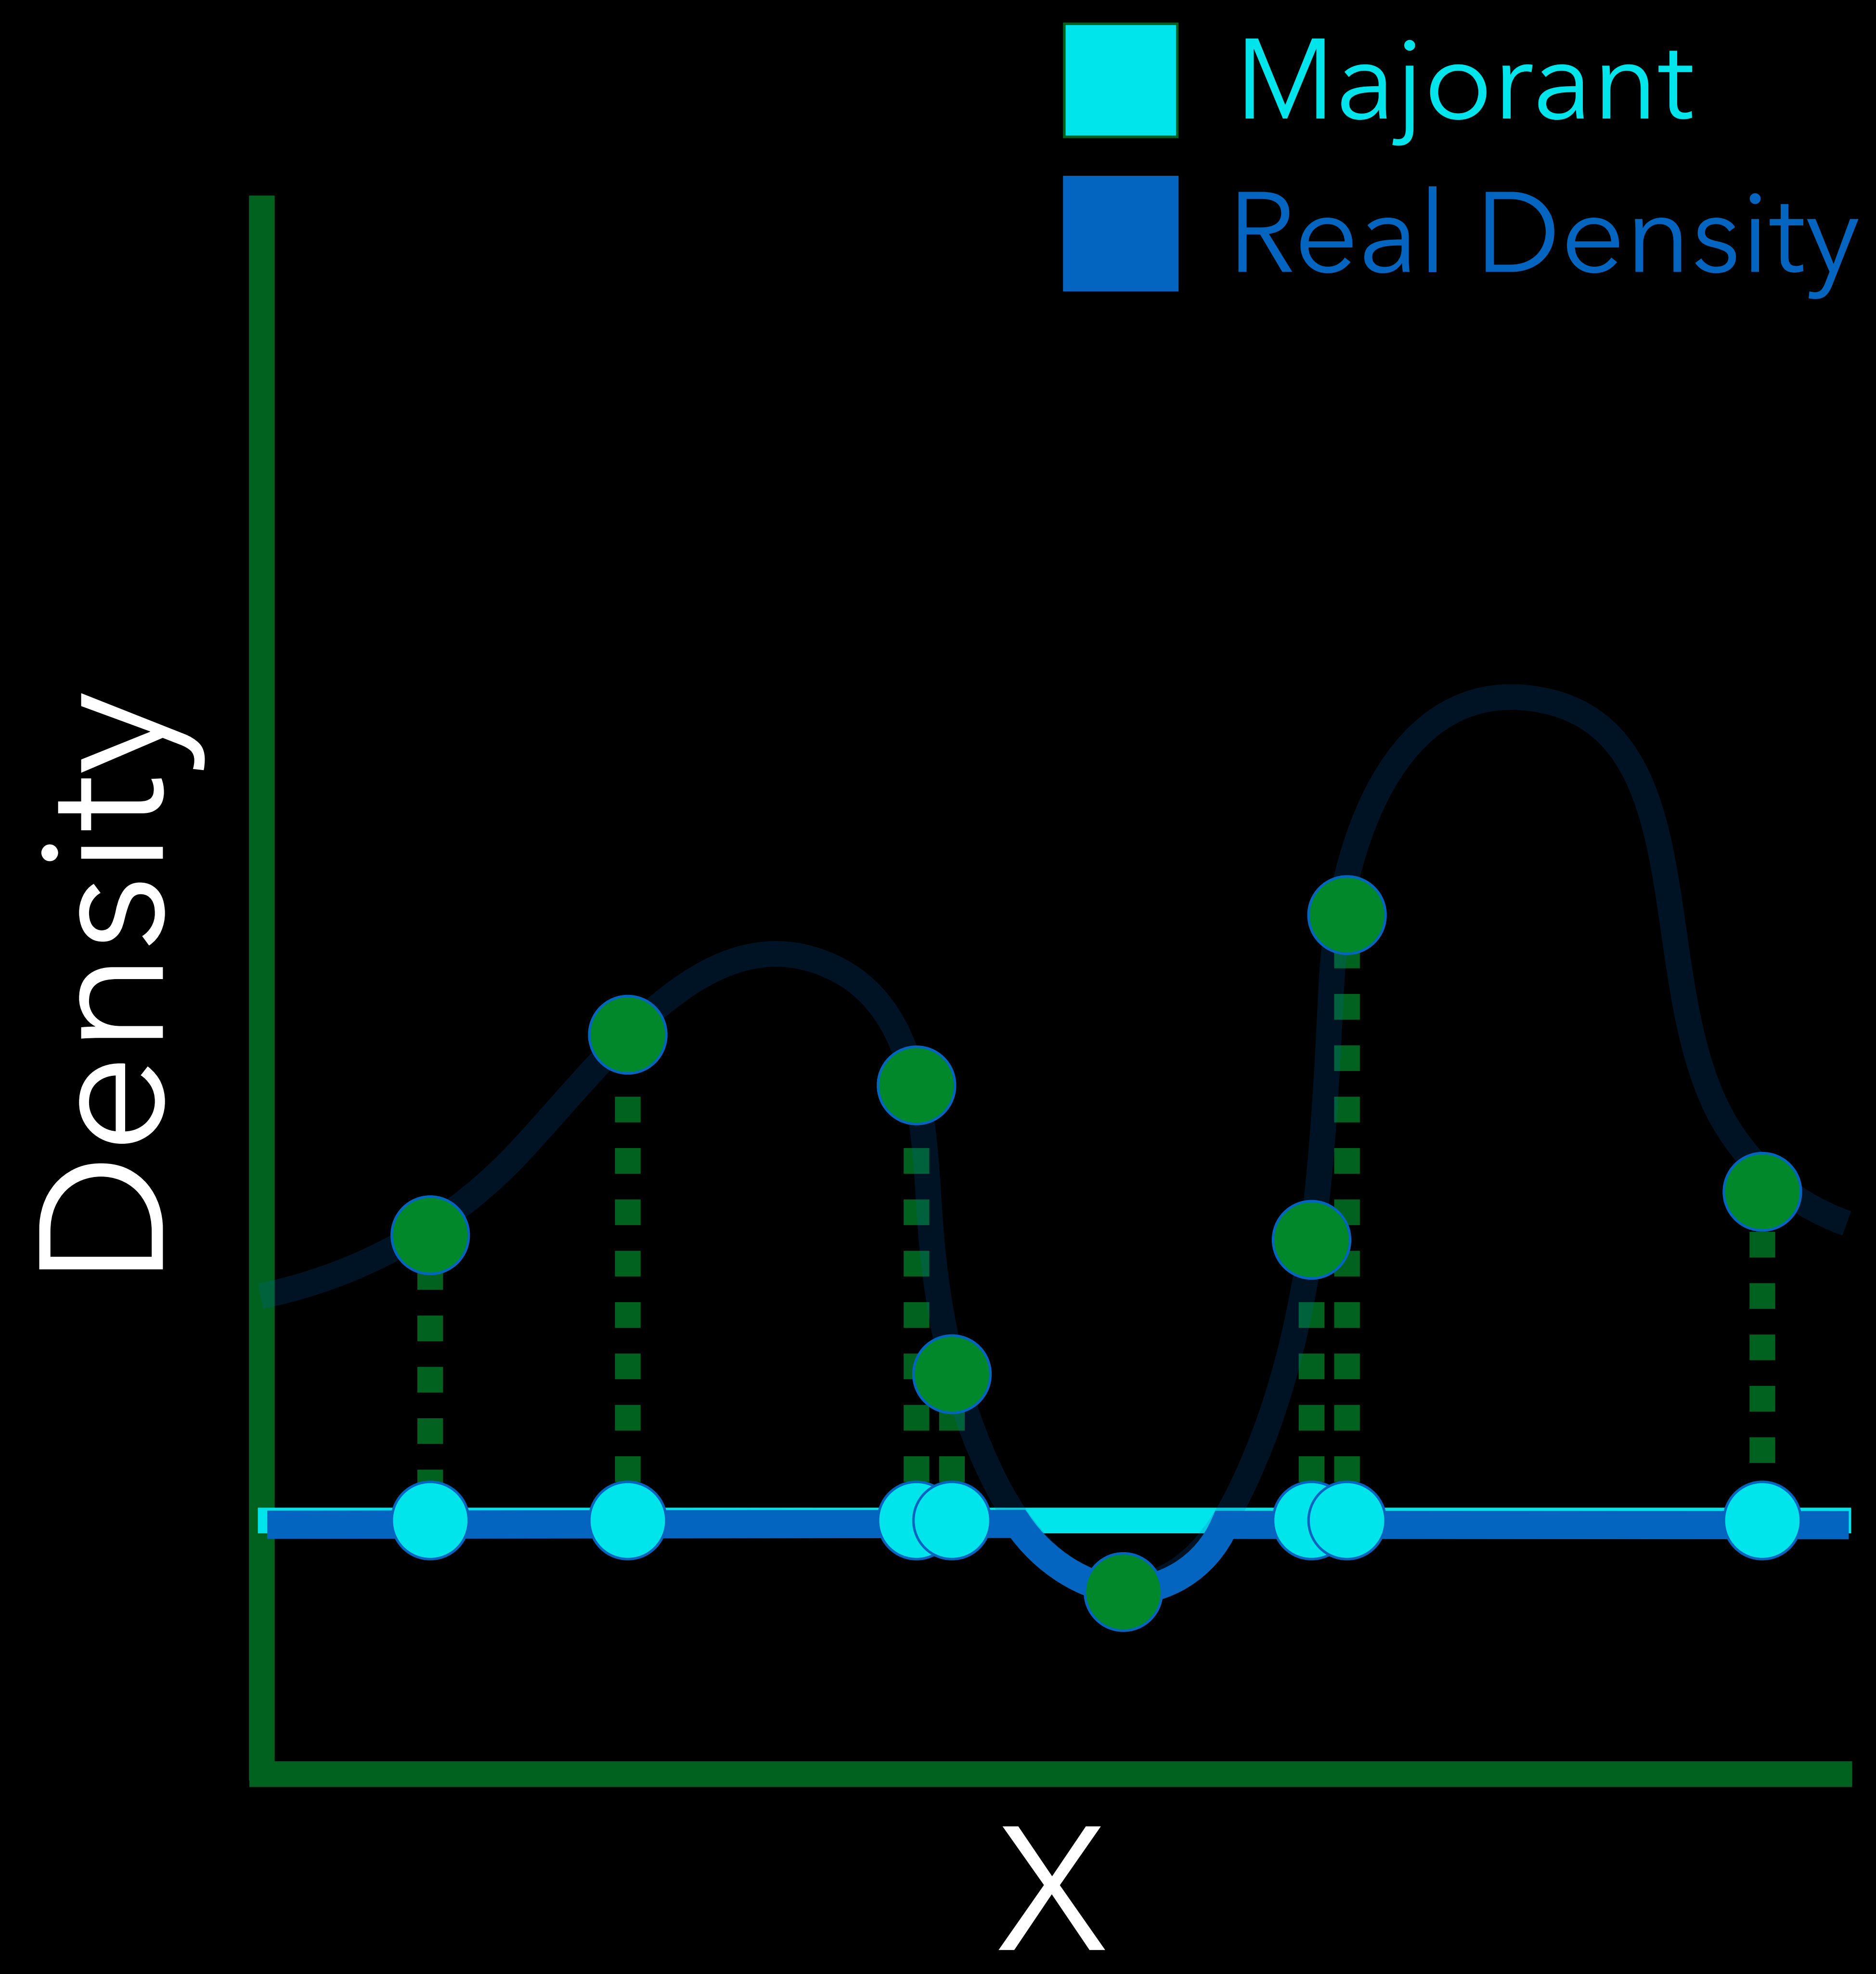
Majorant
Real Density
Density
X

max{Density Evals}

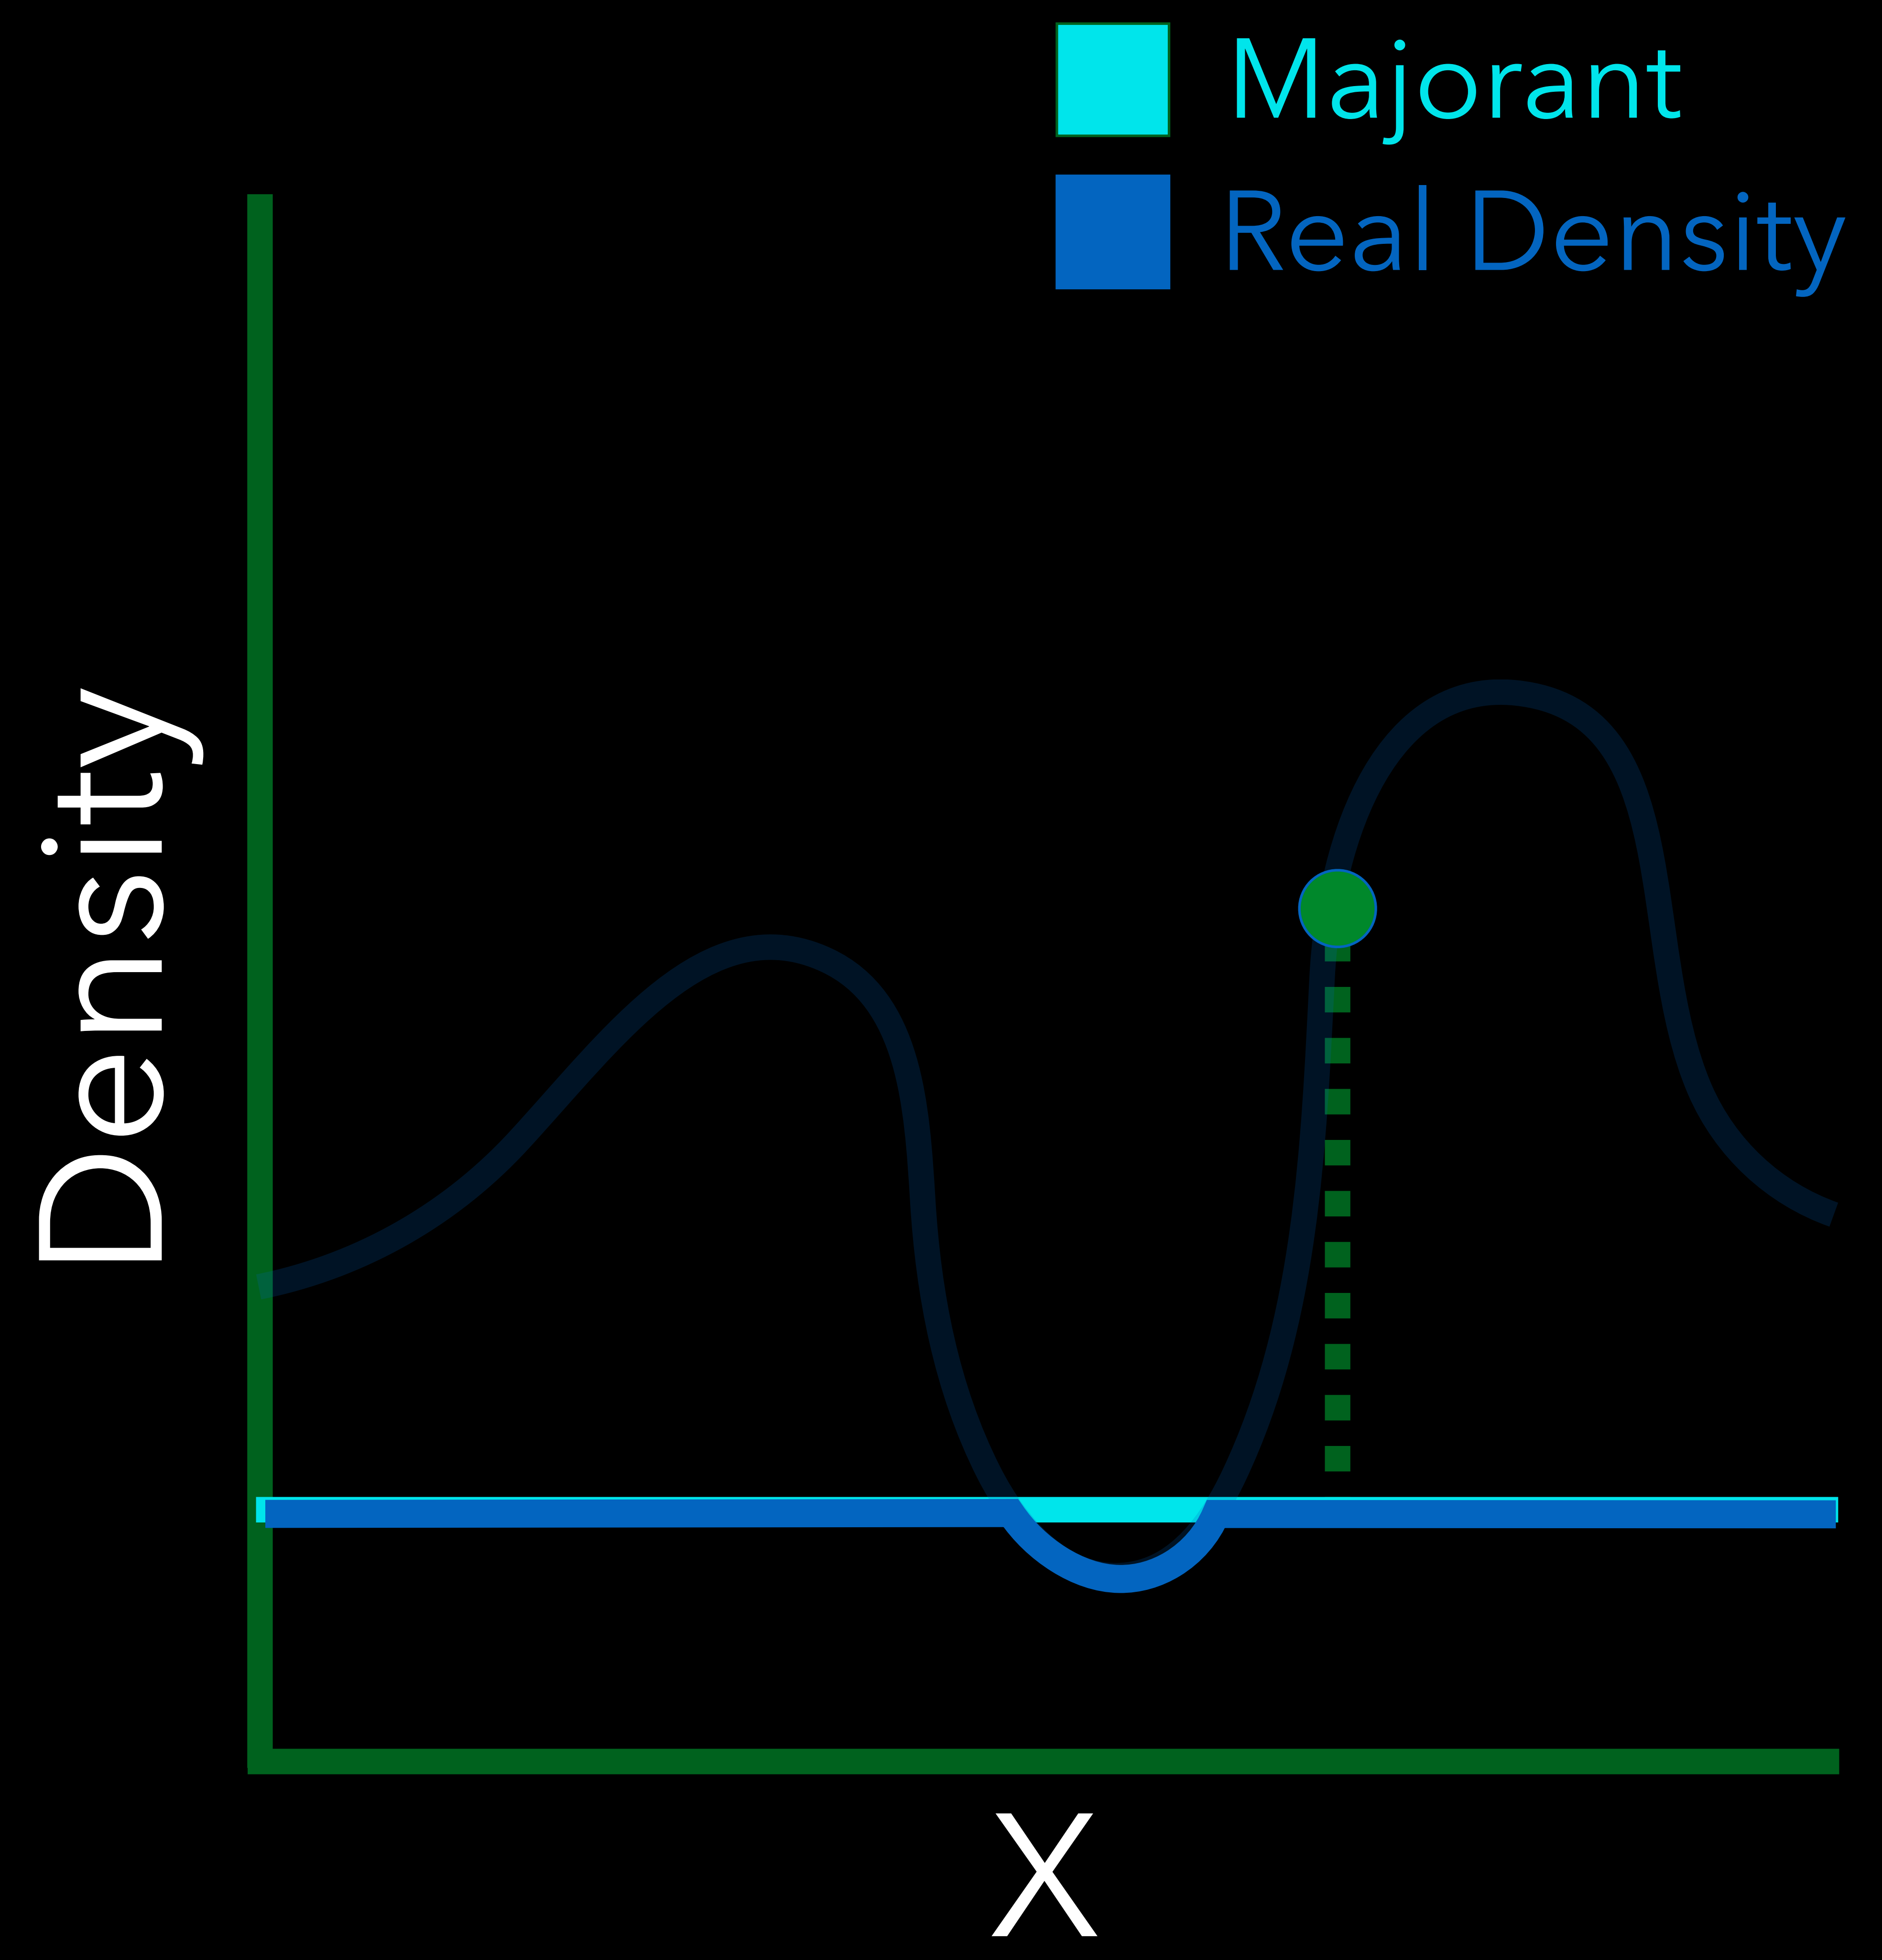

Next Majorant = max{Density Evals}

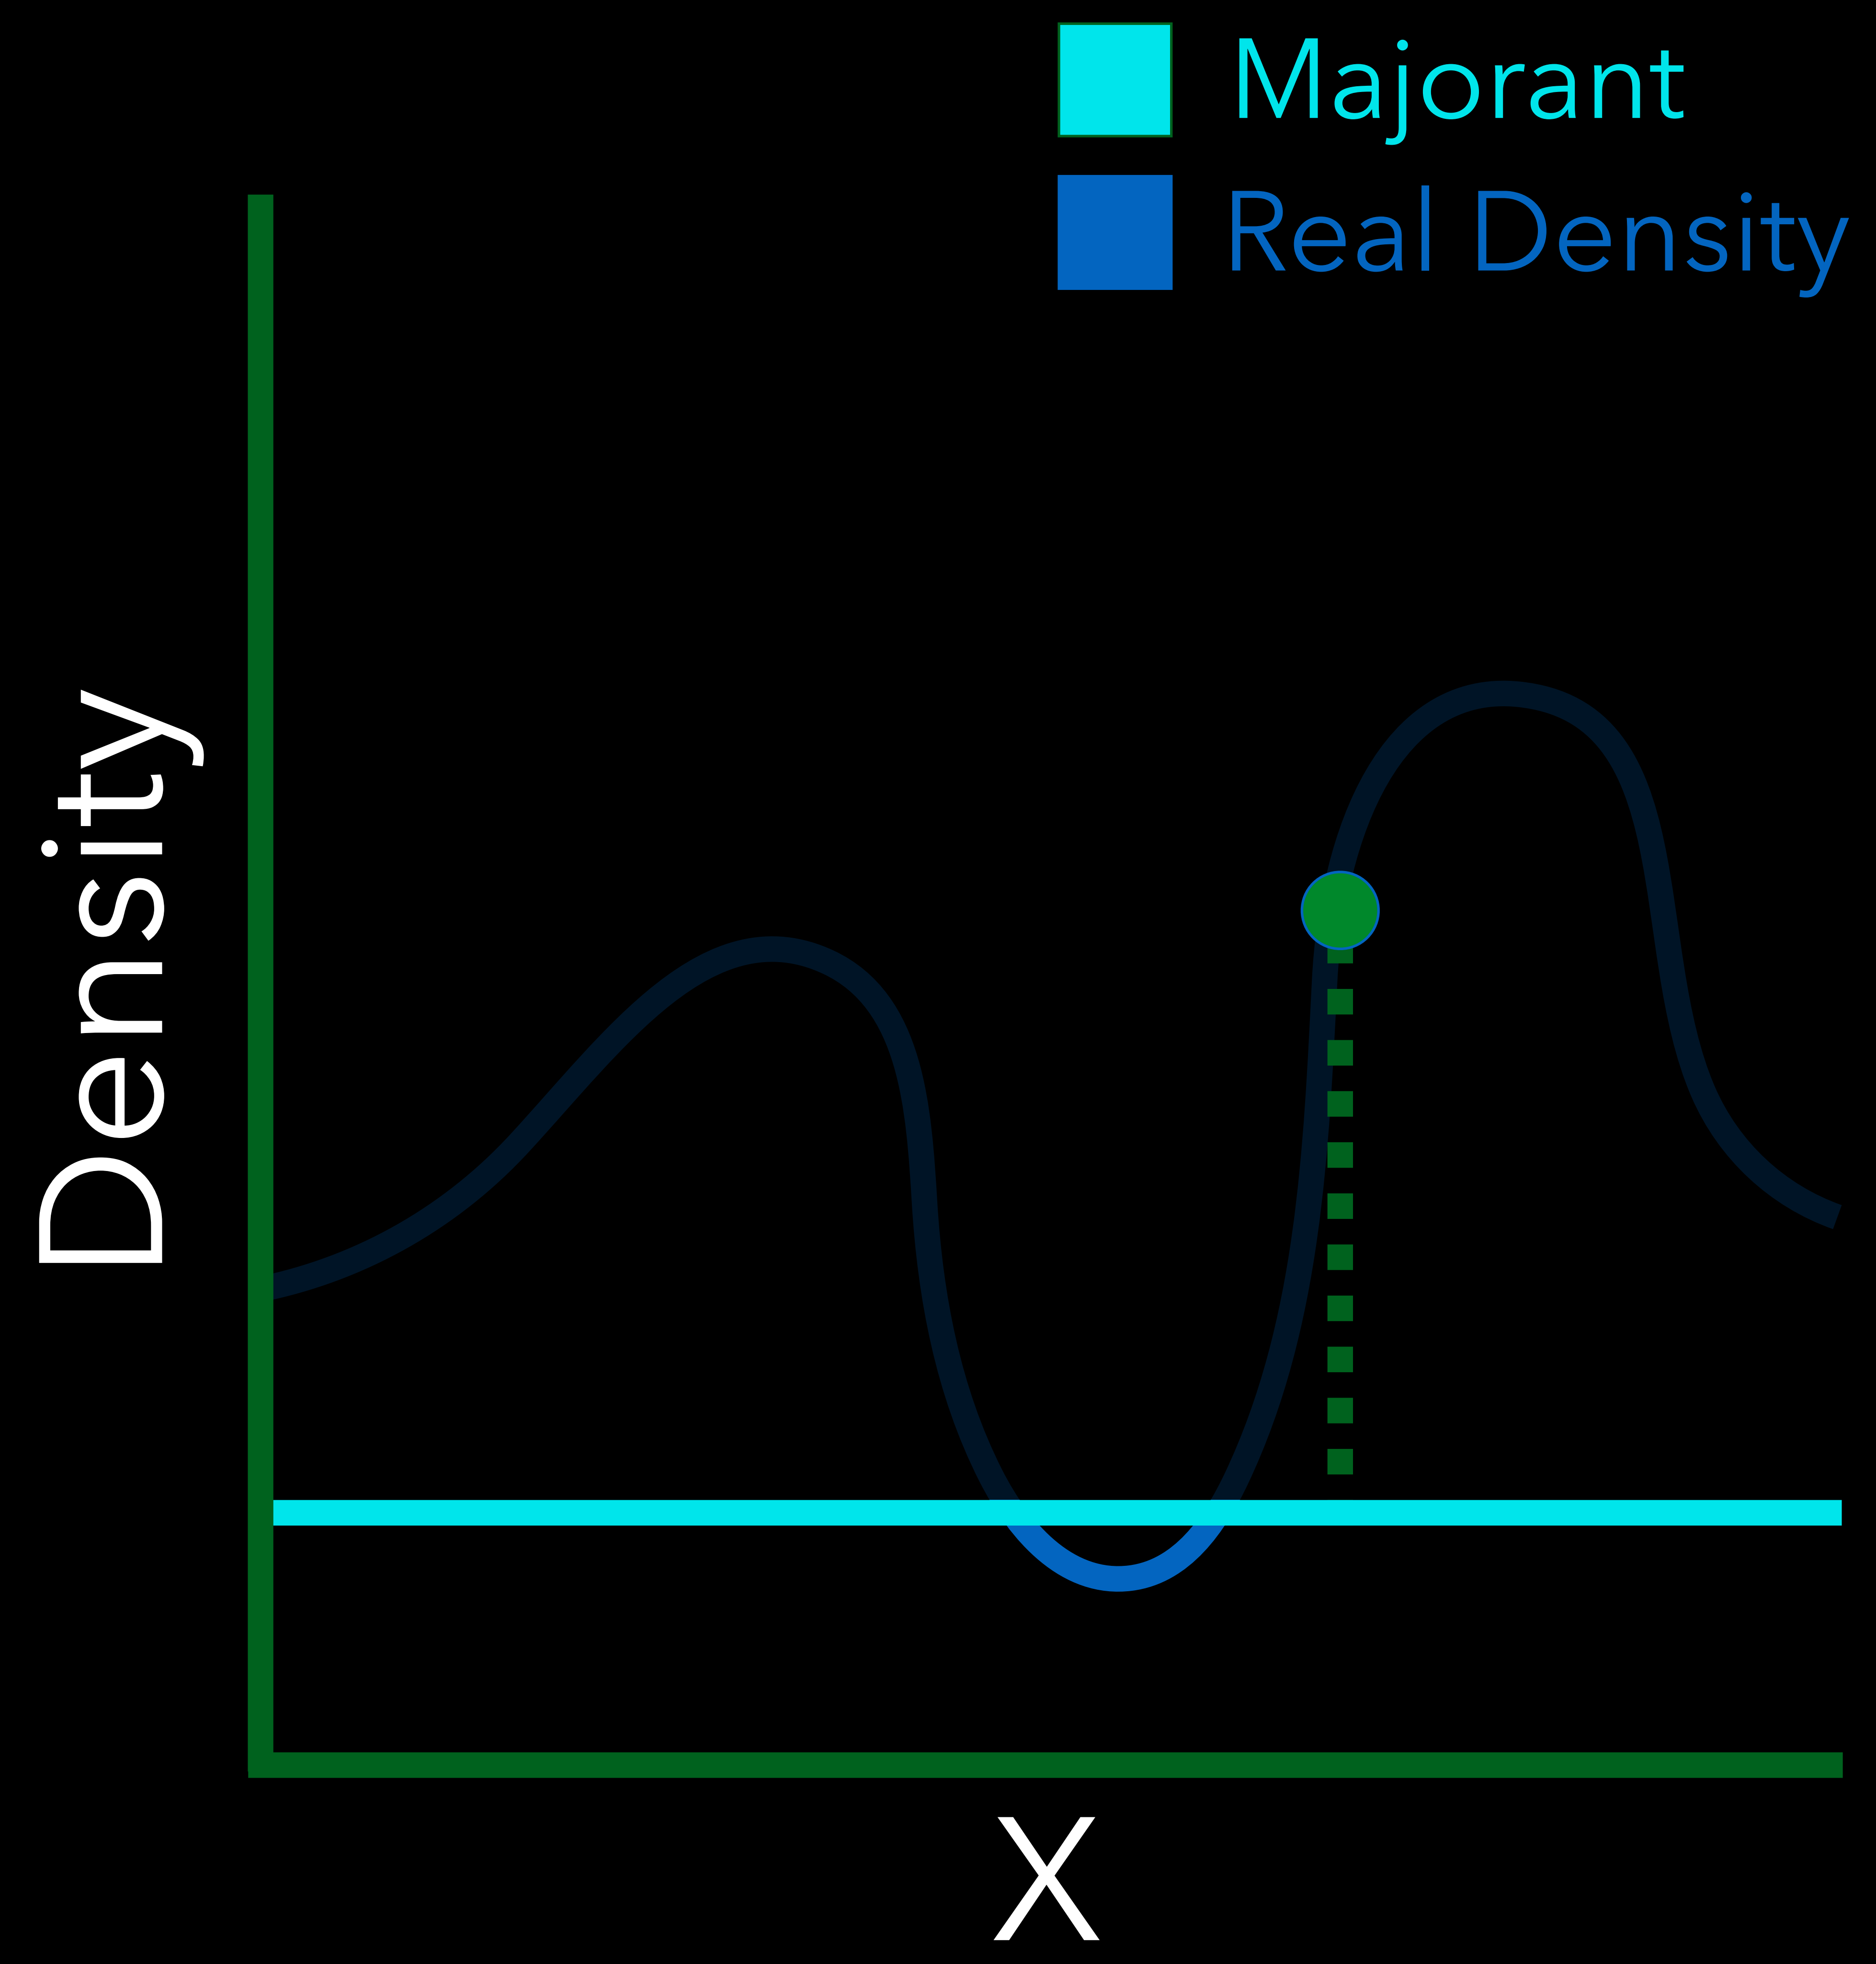

Next Majorant = max{Density Evals}

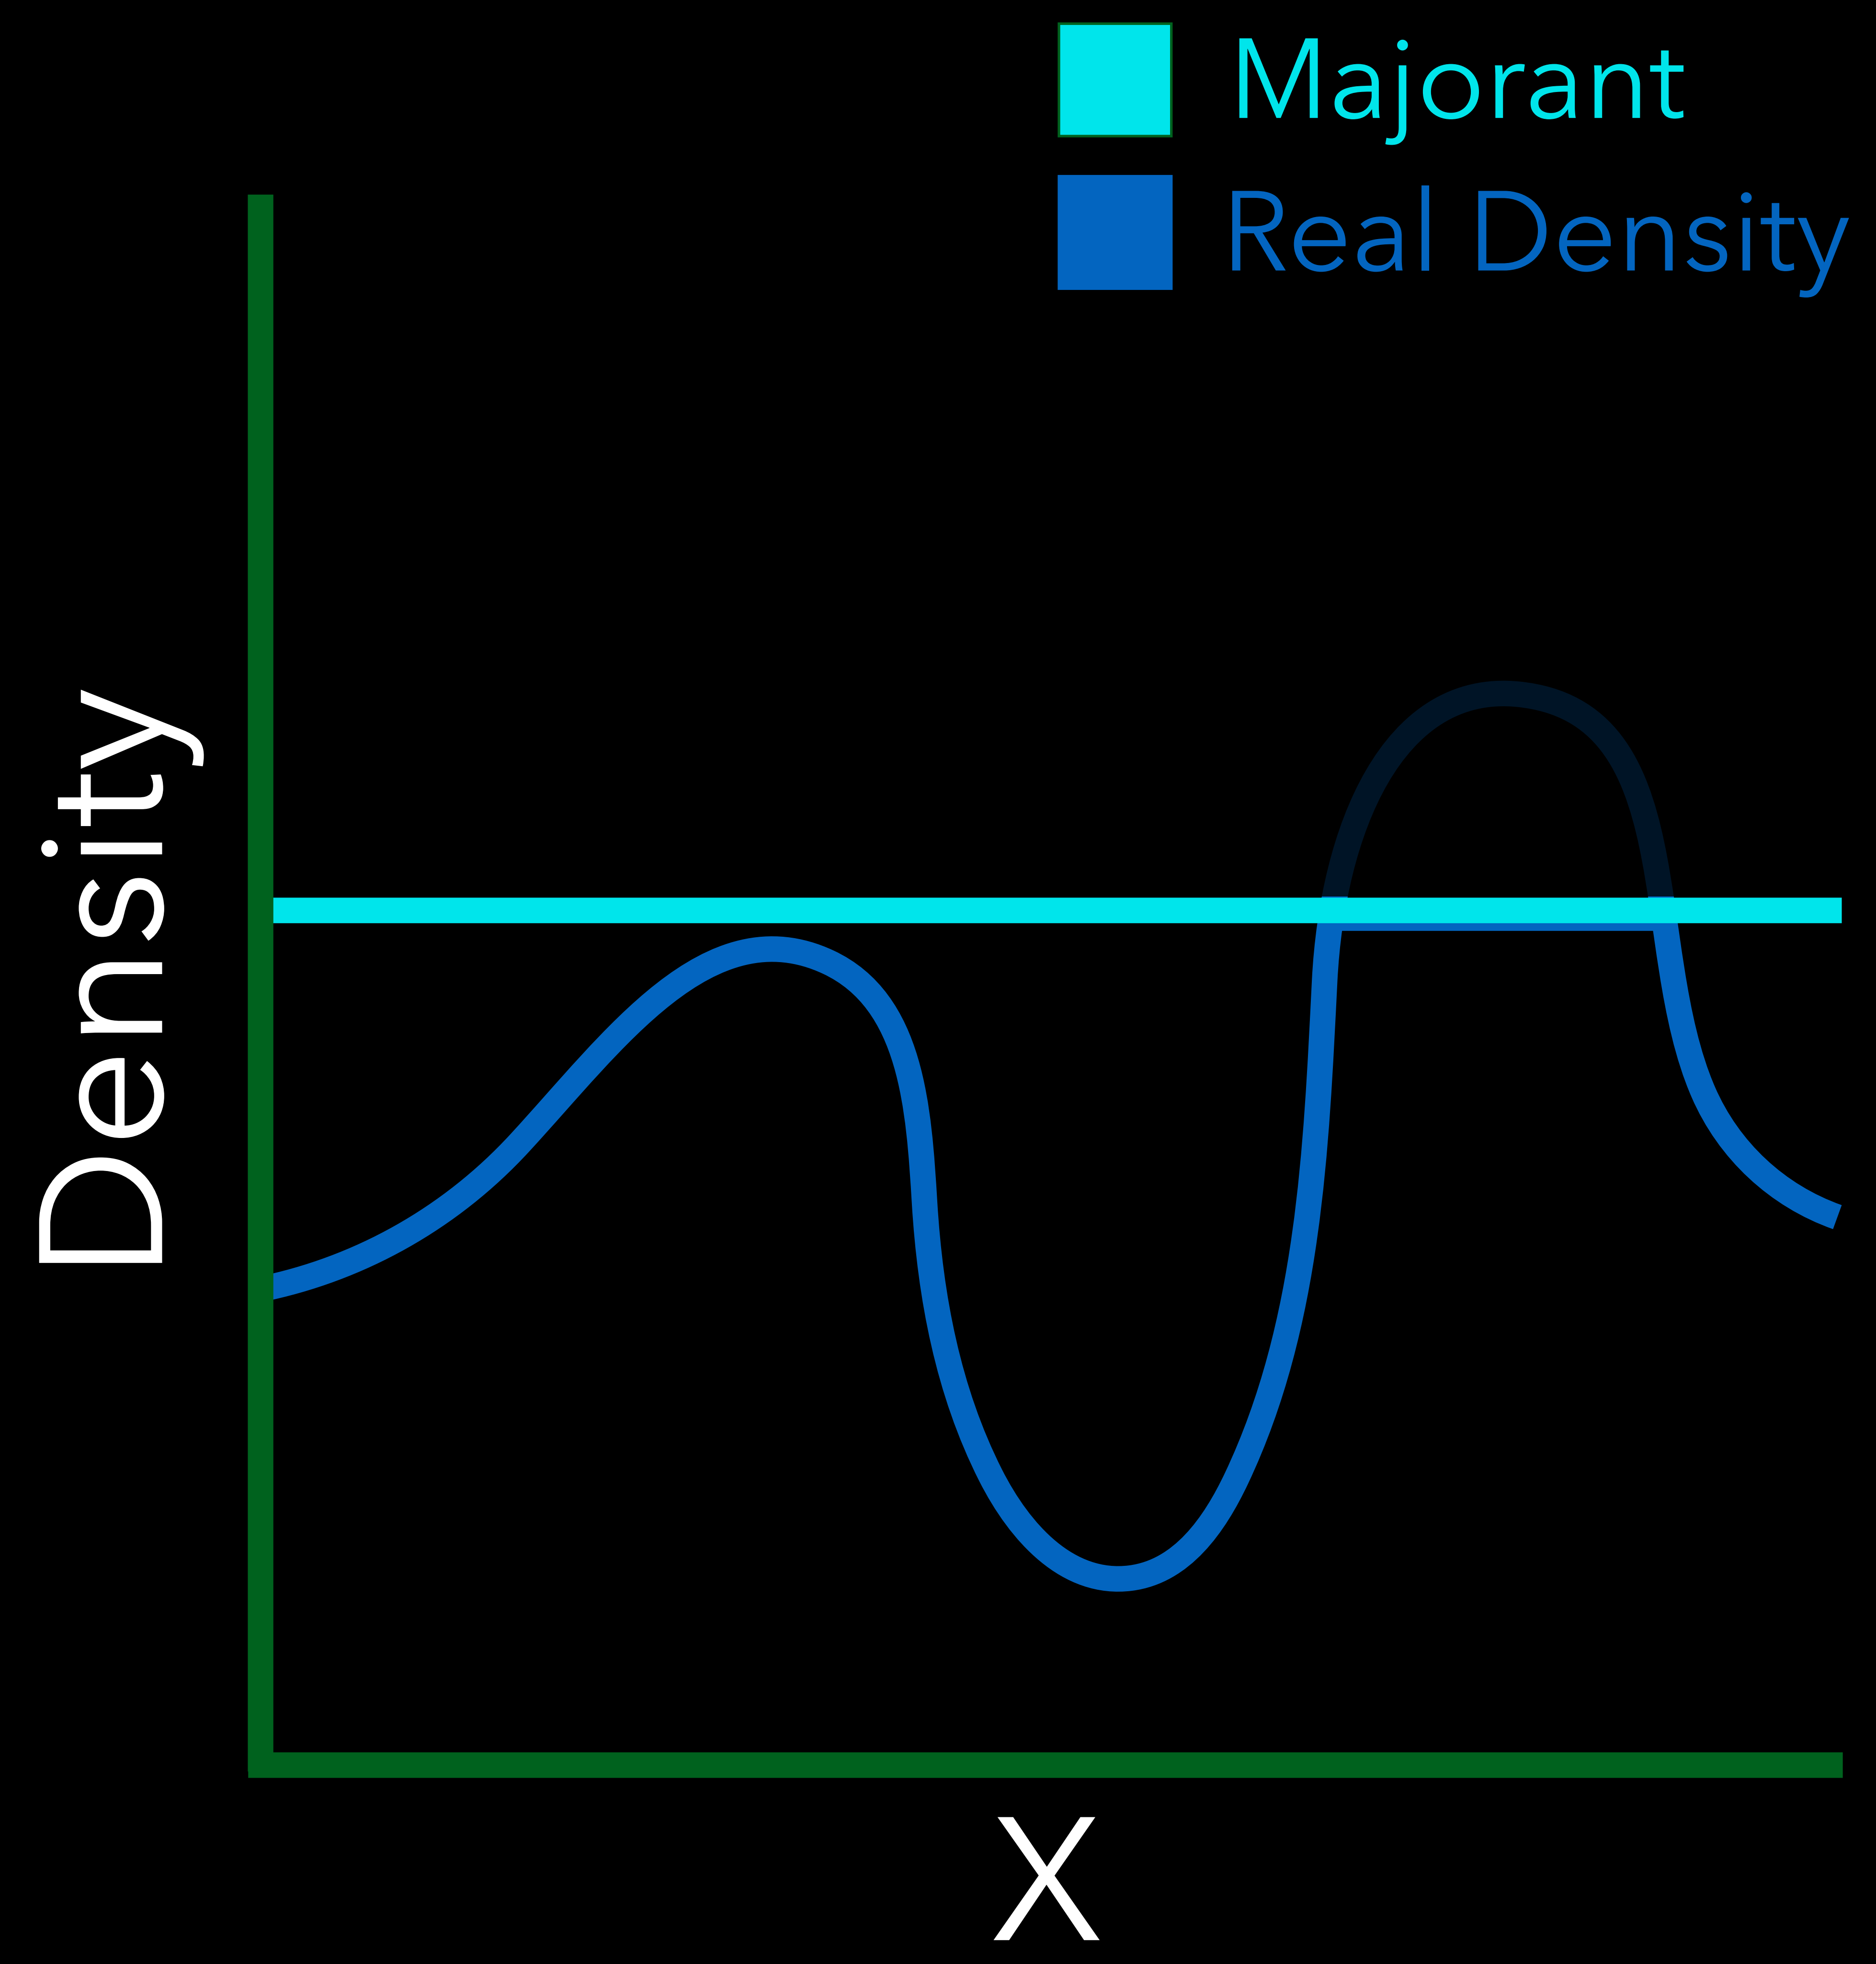

Next Majorant = max{Density Evals} + C

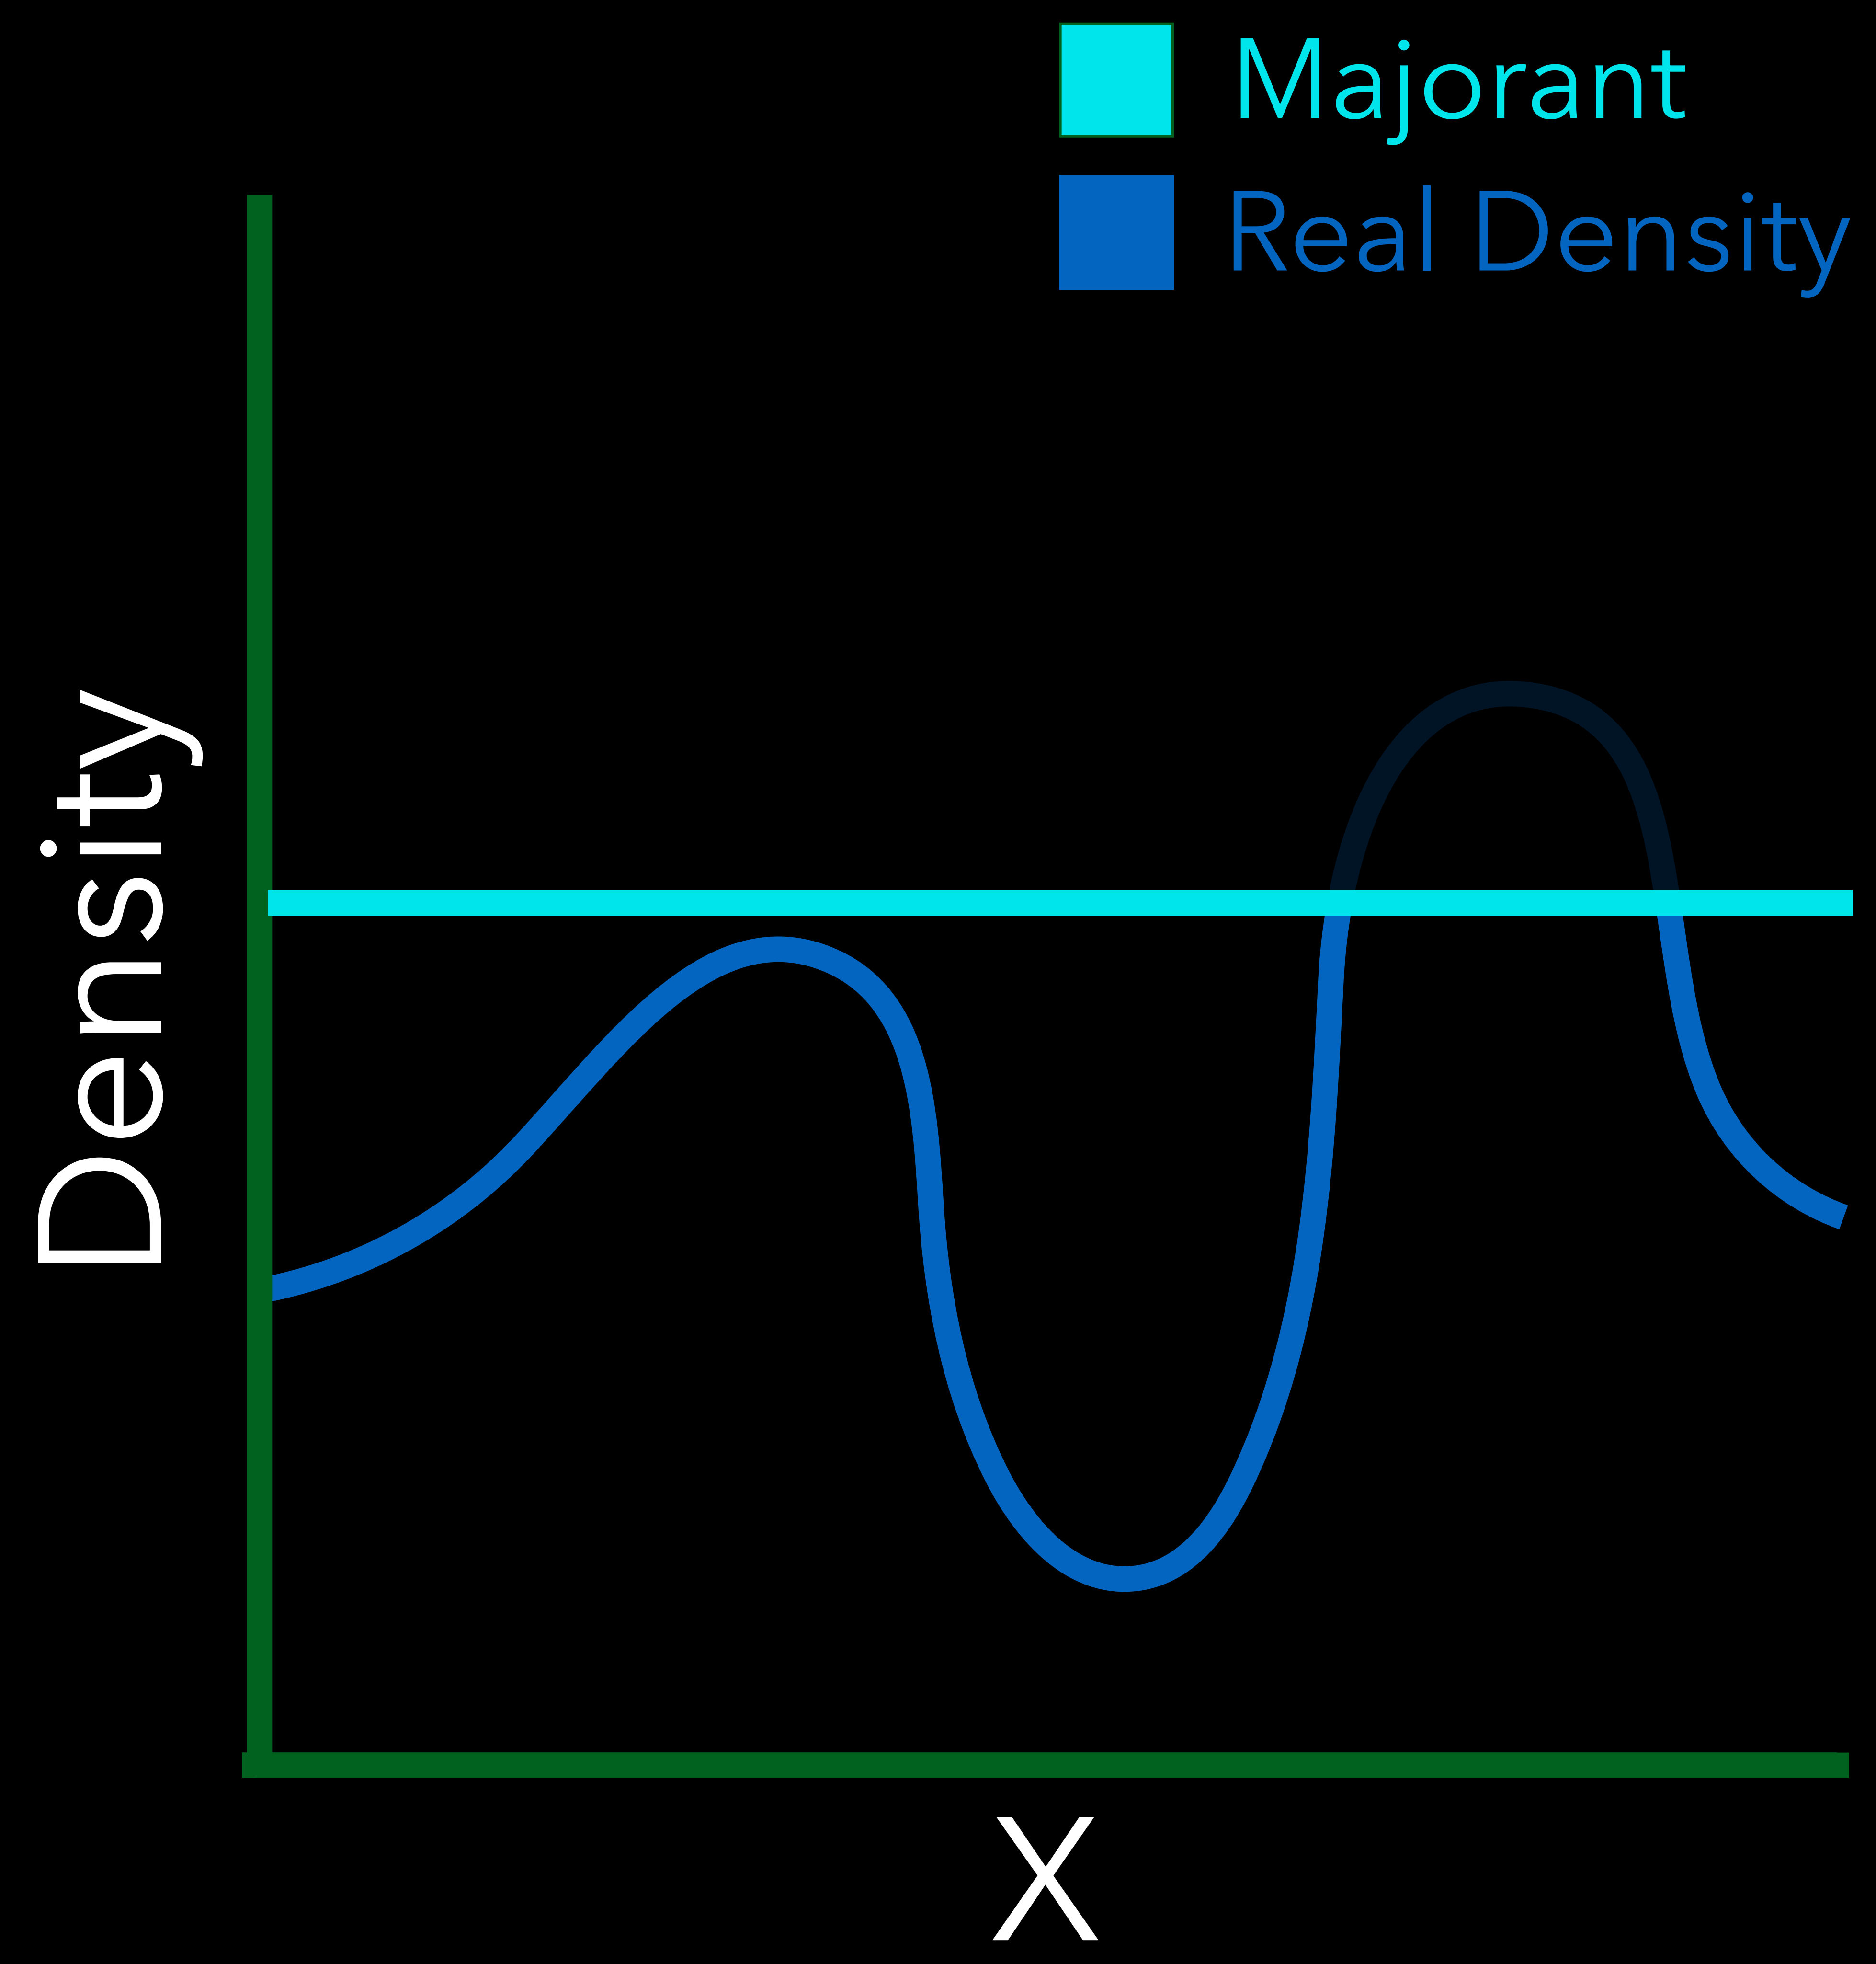

Next Majorant = max{Density Evals} + C

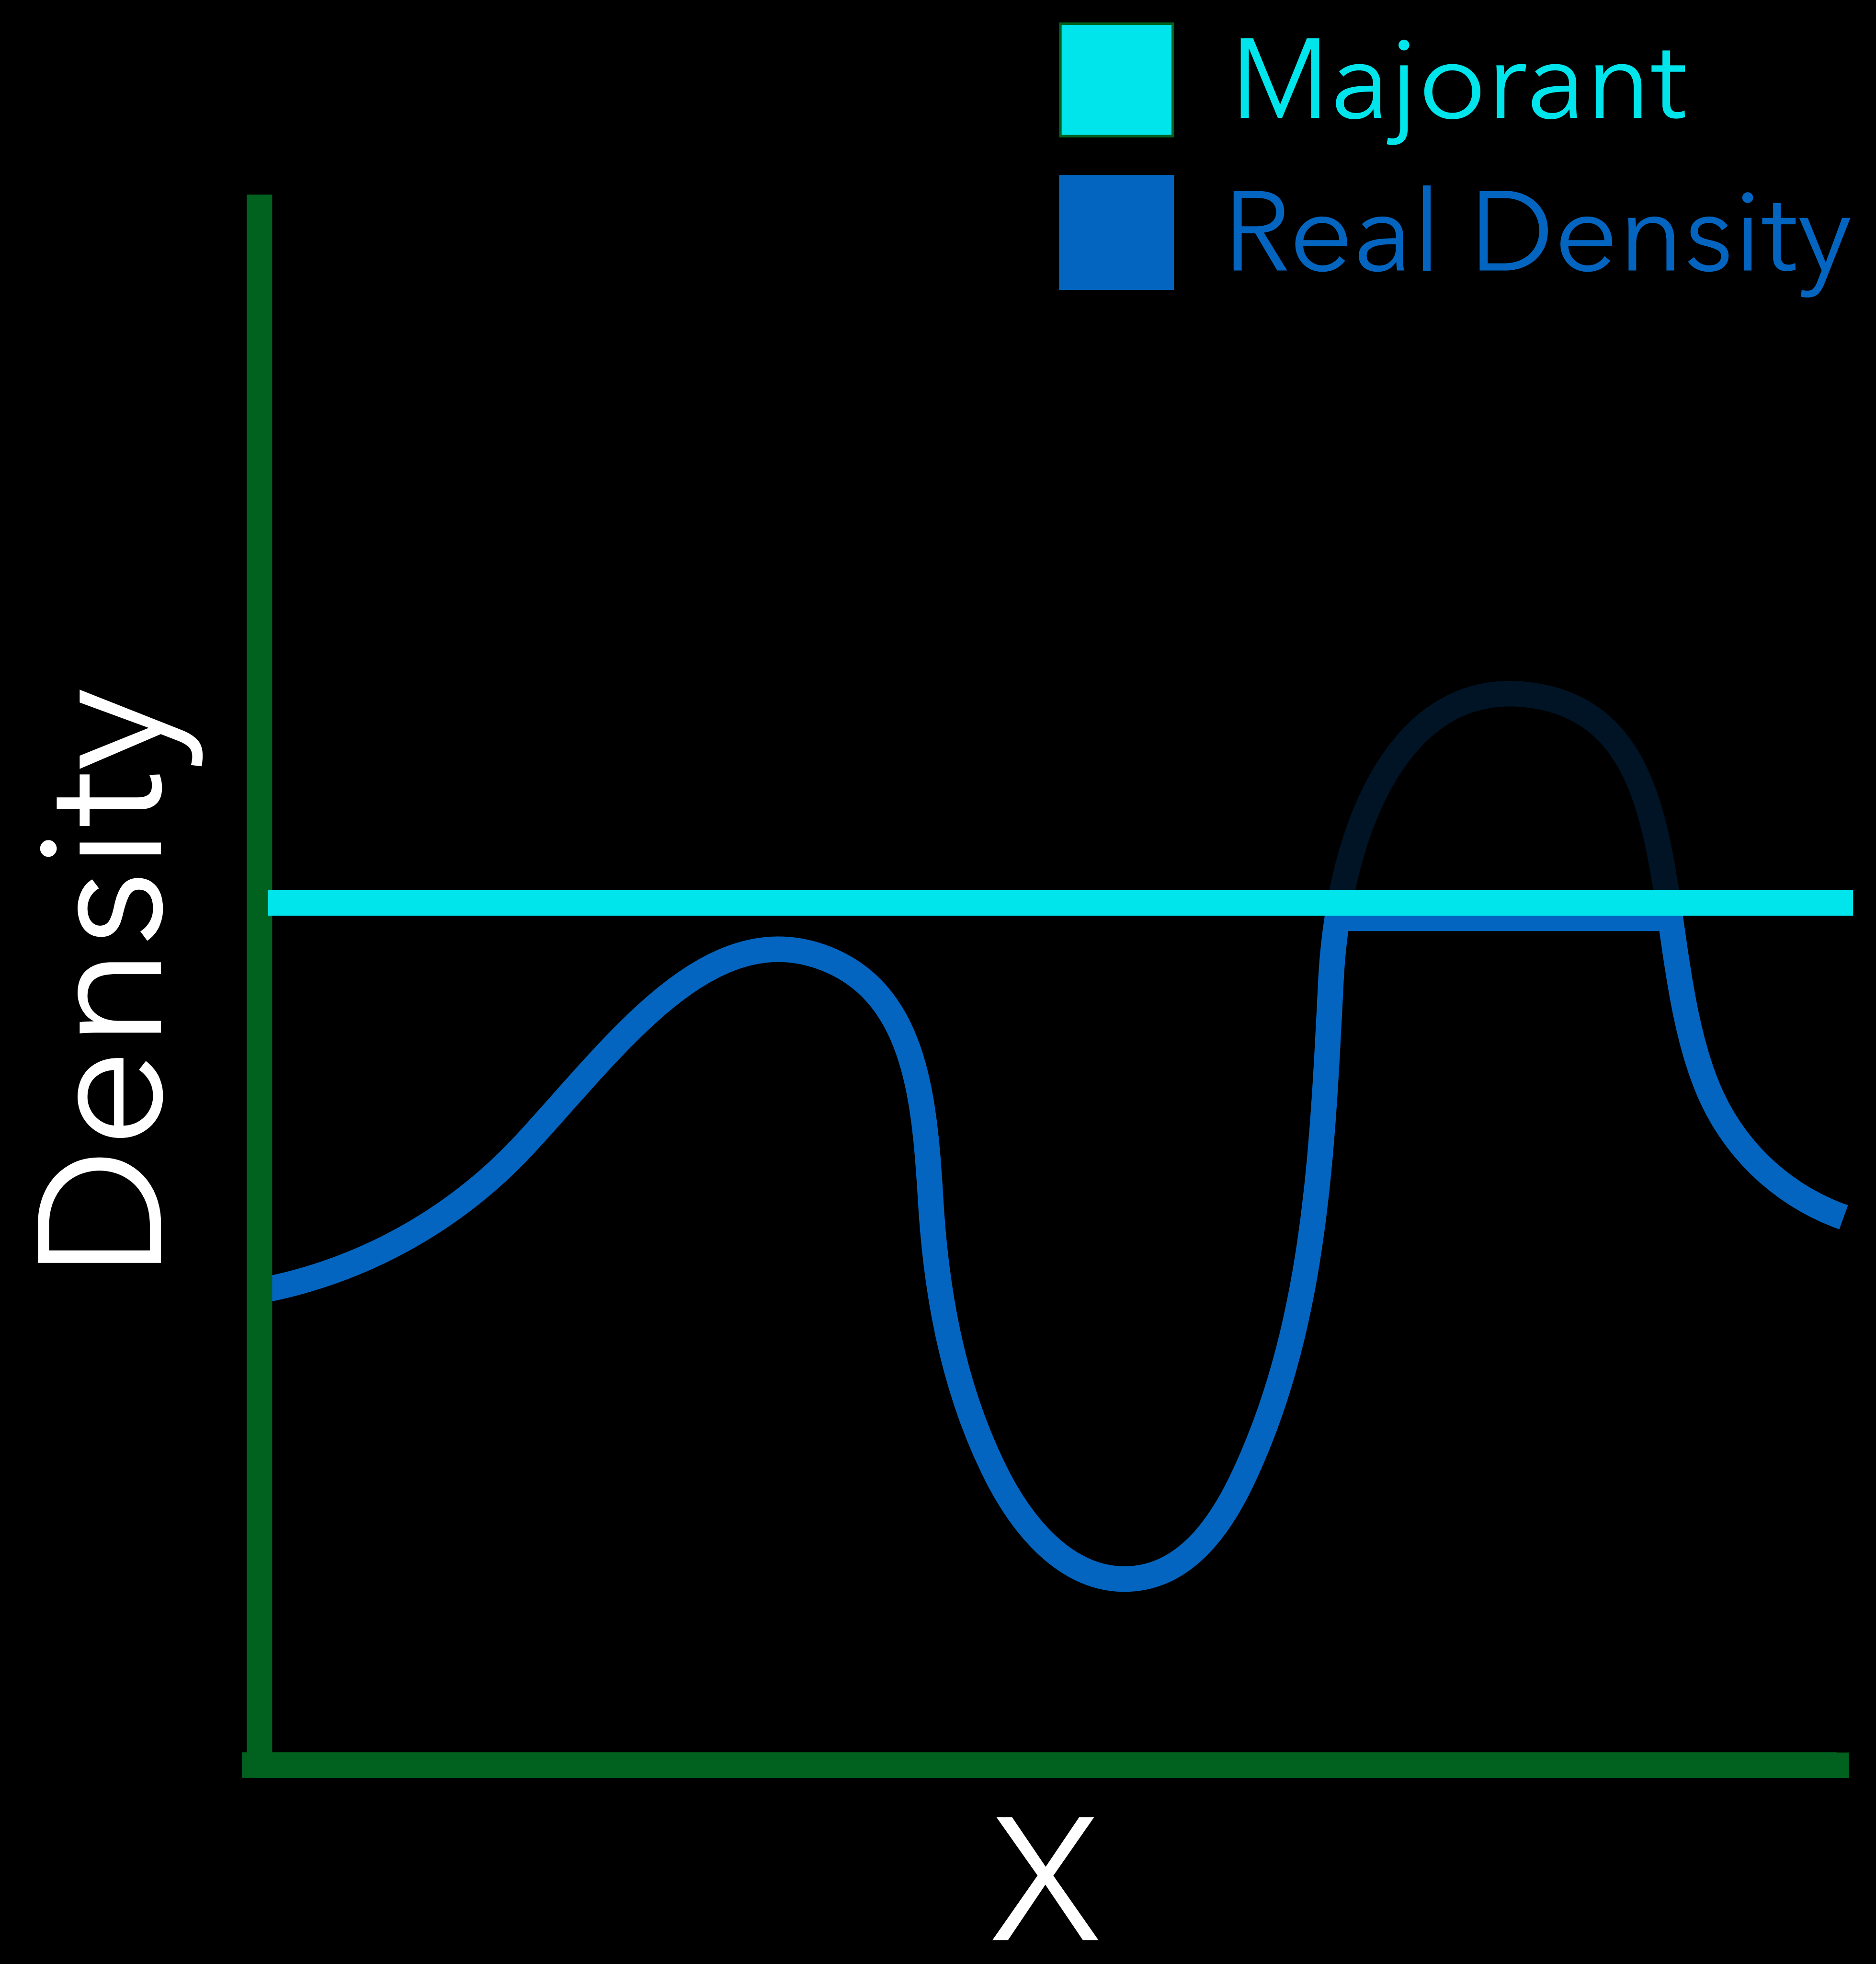

Implementation

Hyperion

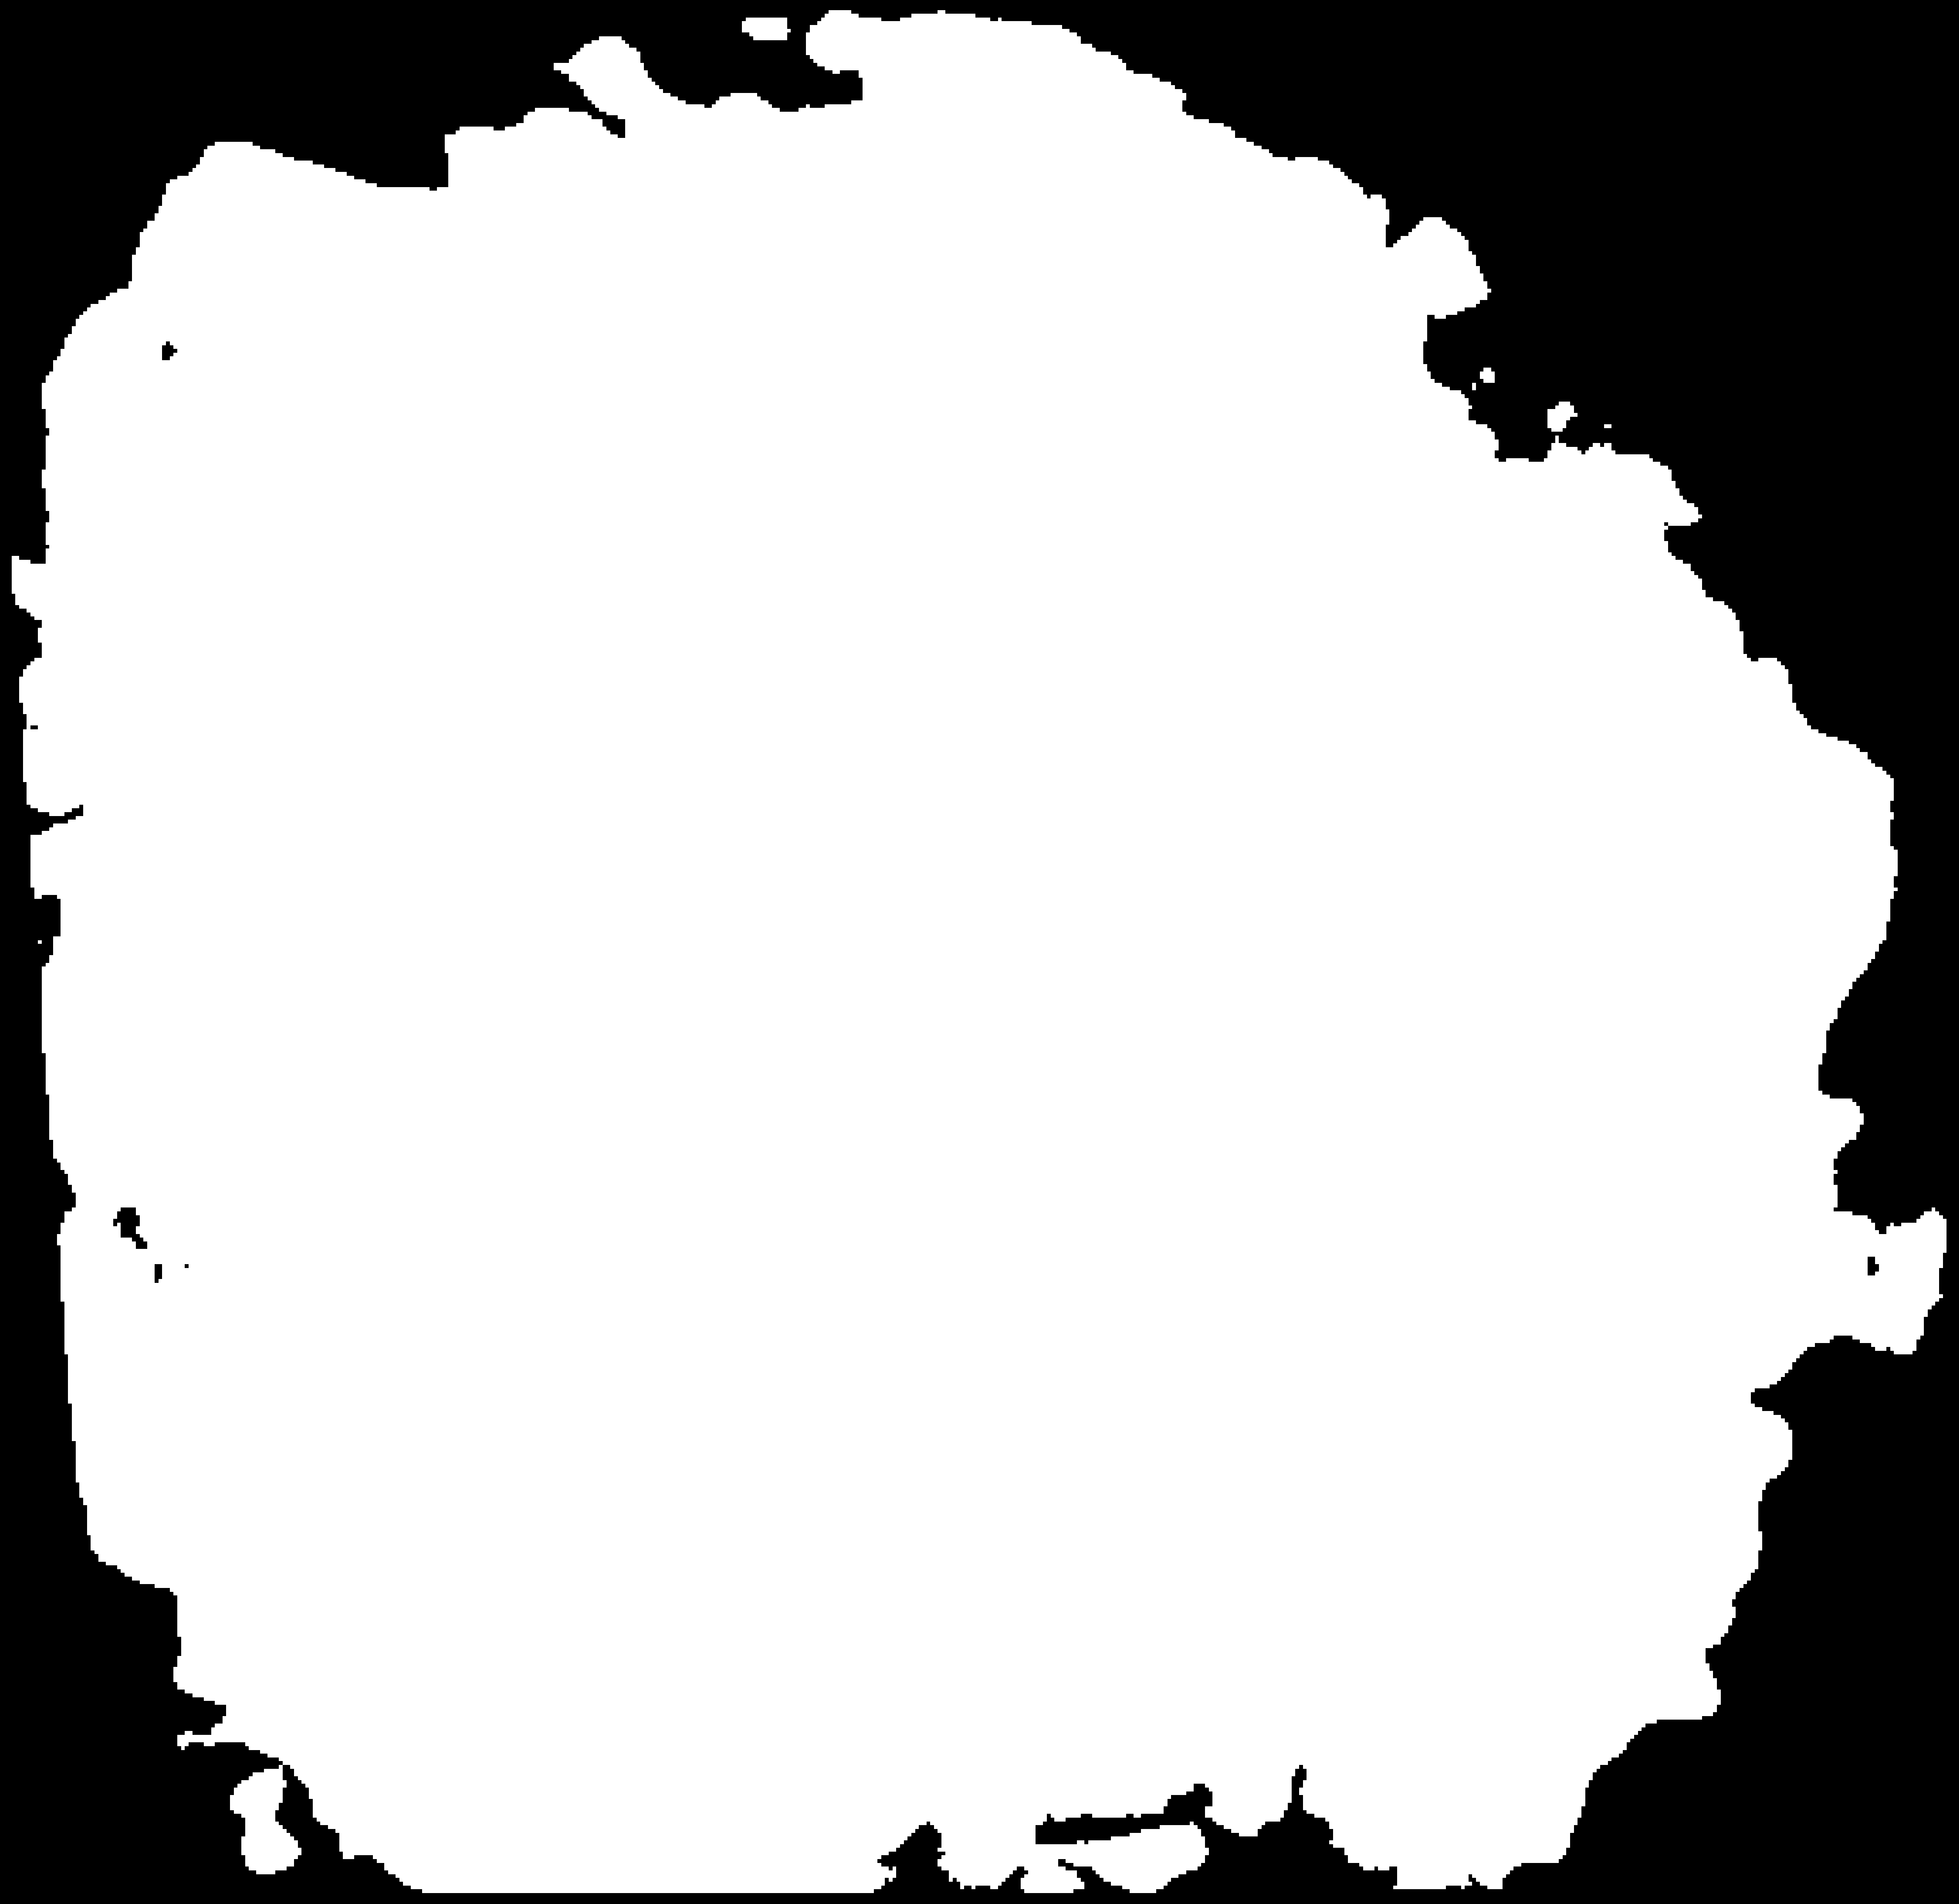

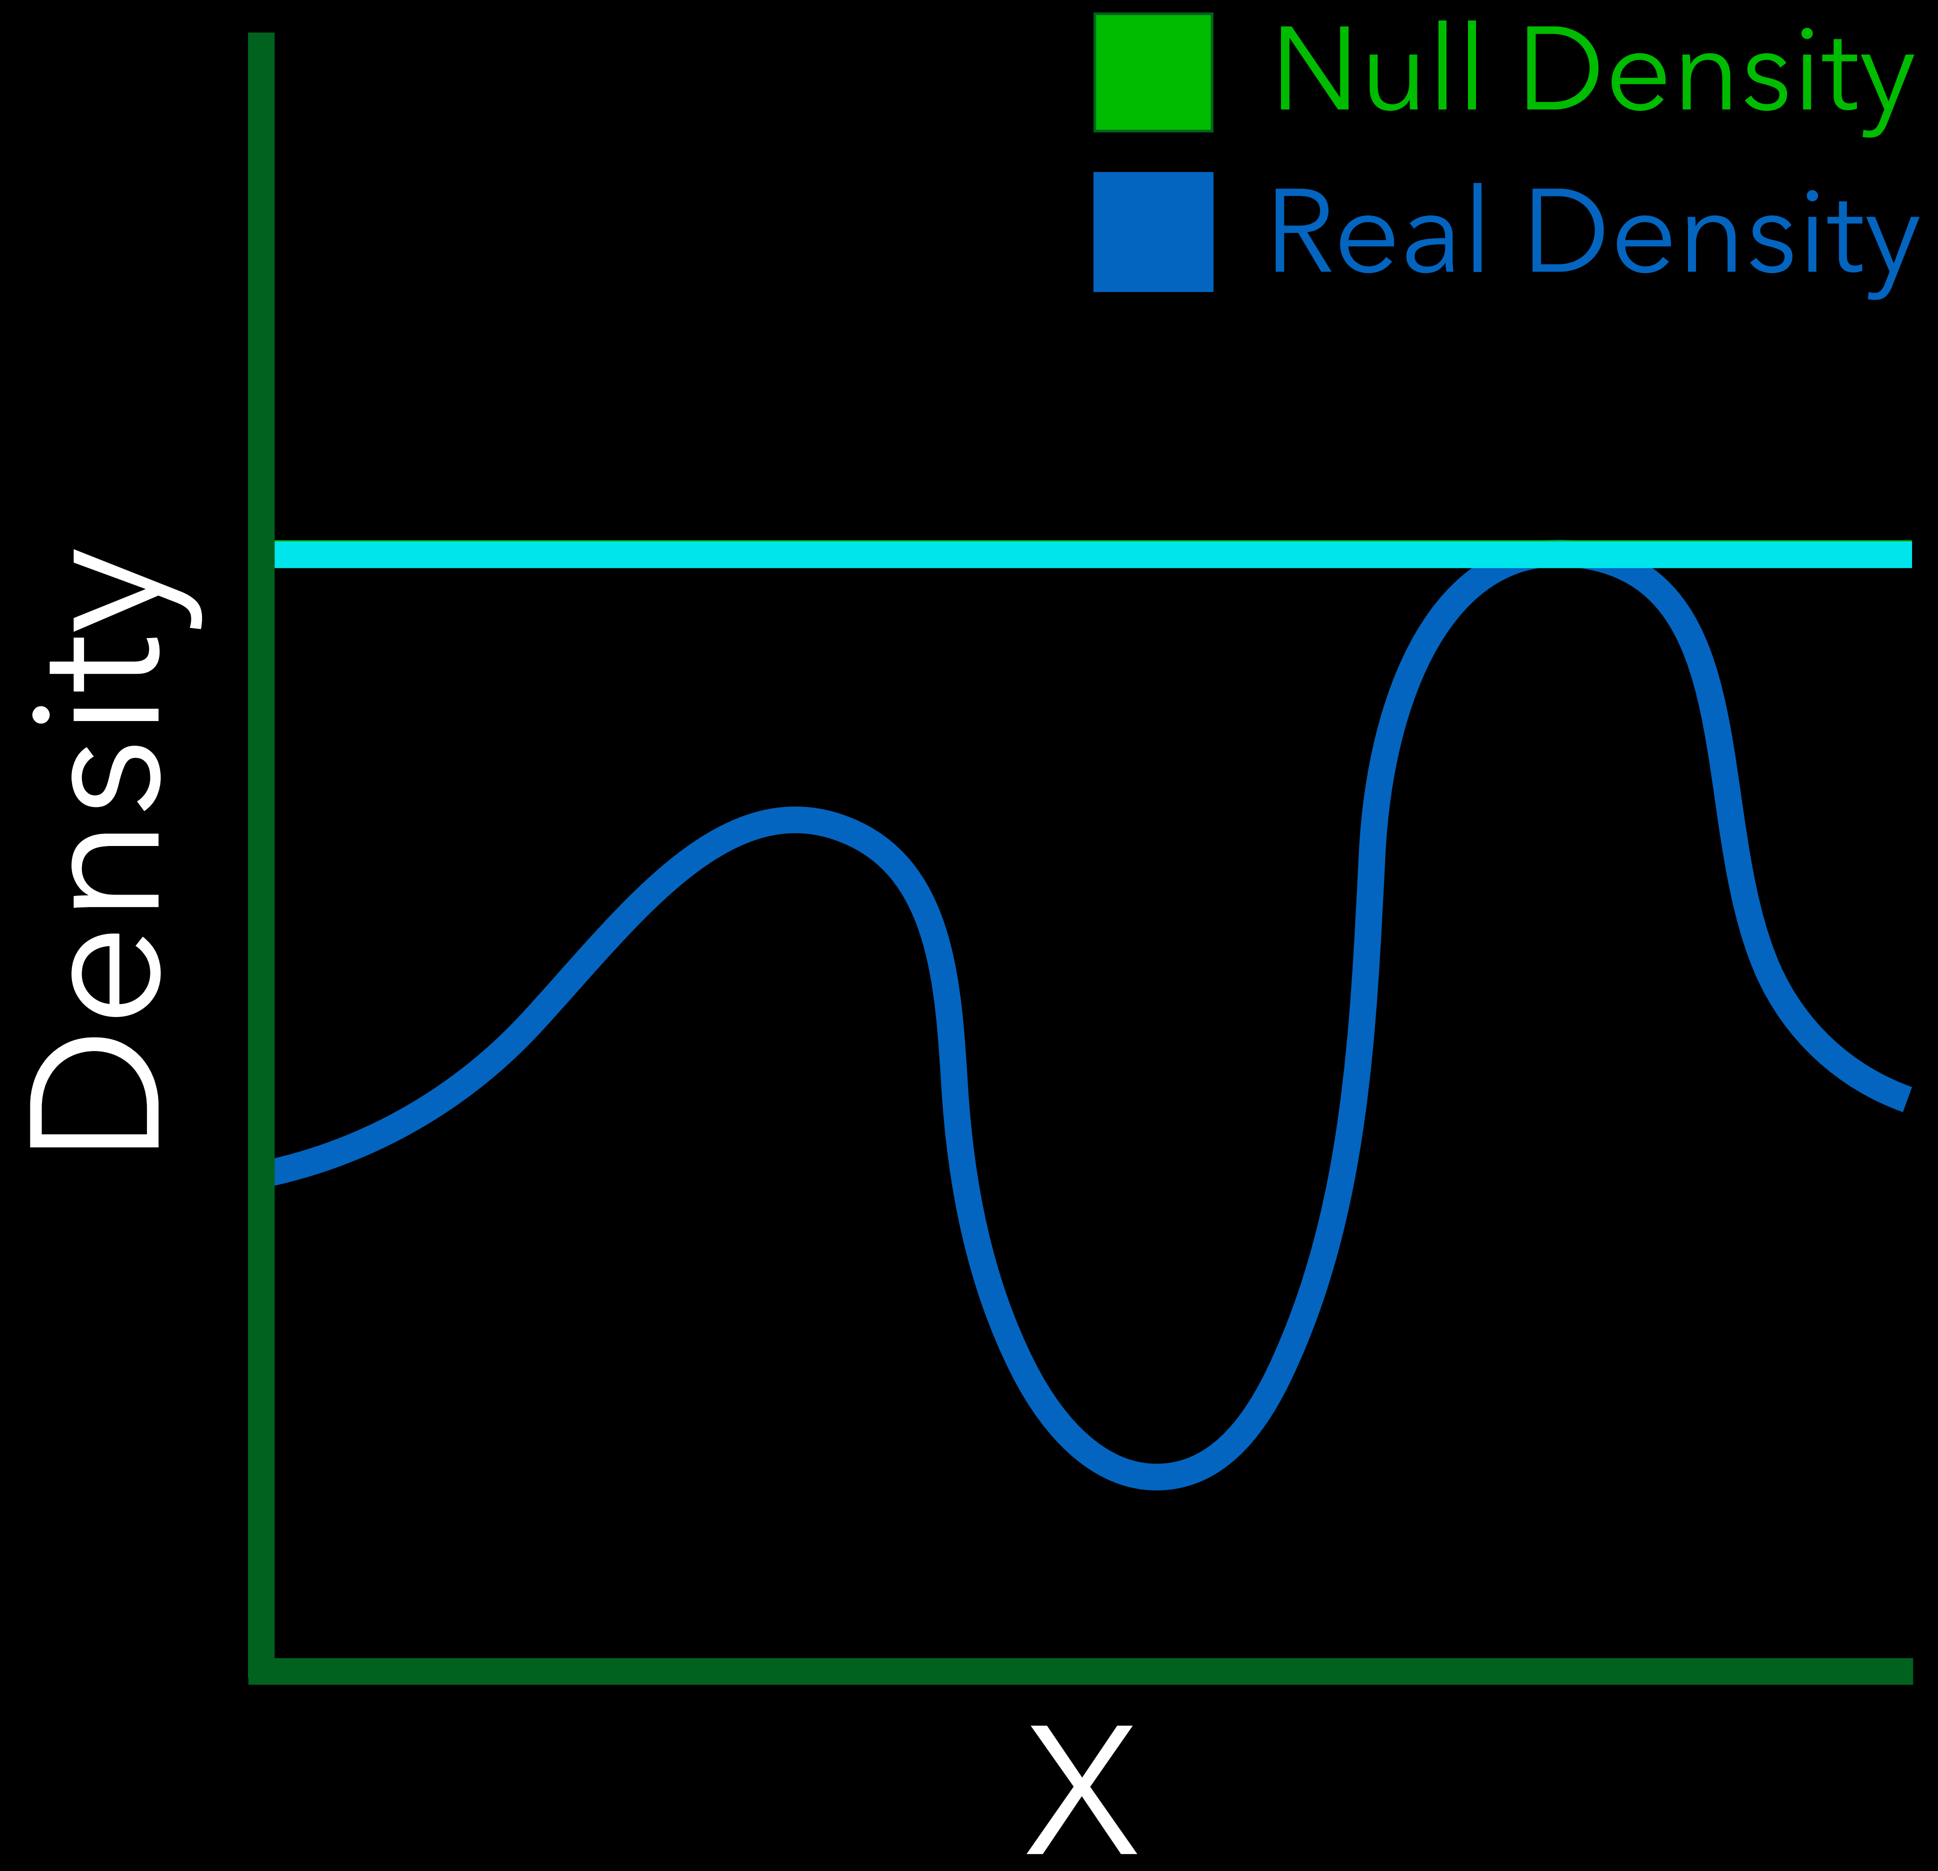

Null Density
Real Density
Density
X

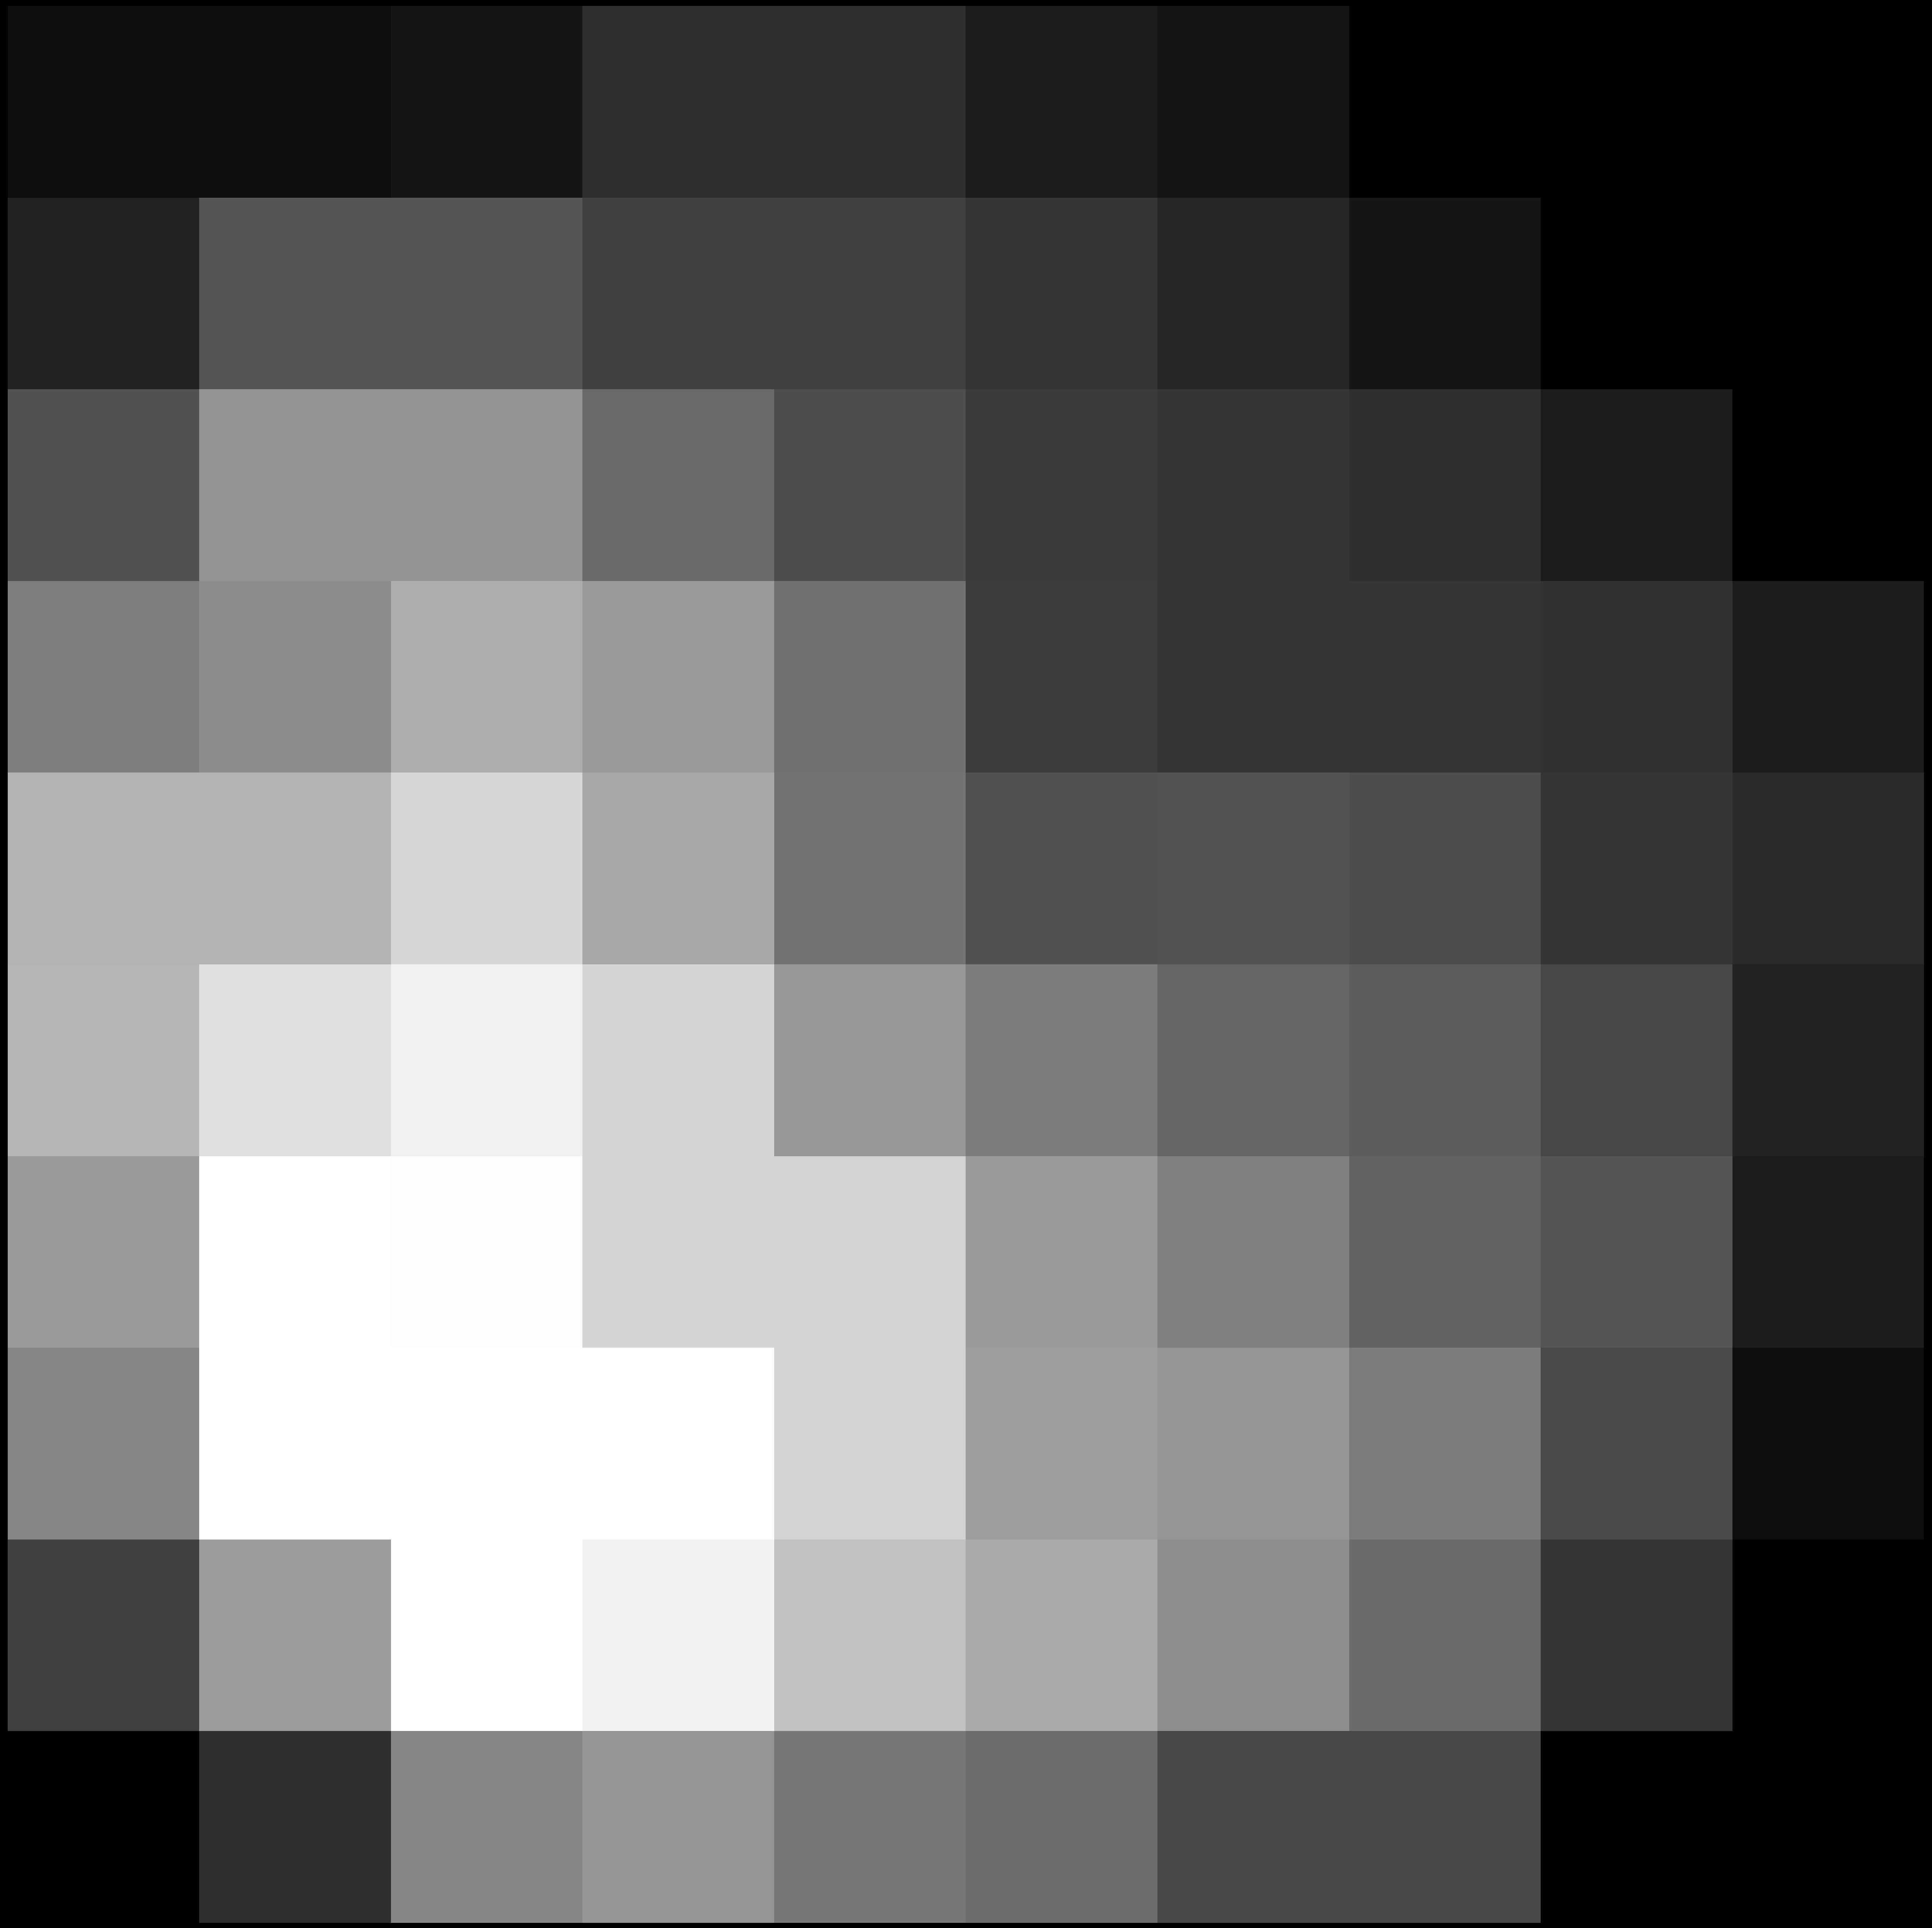

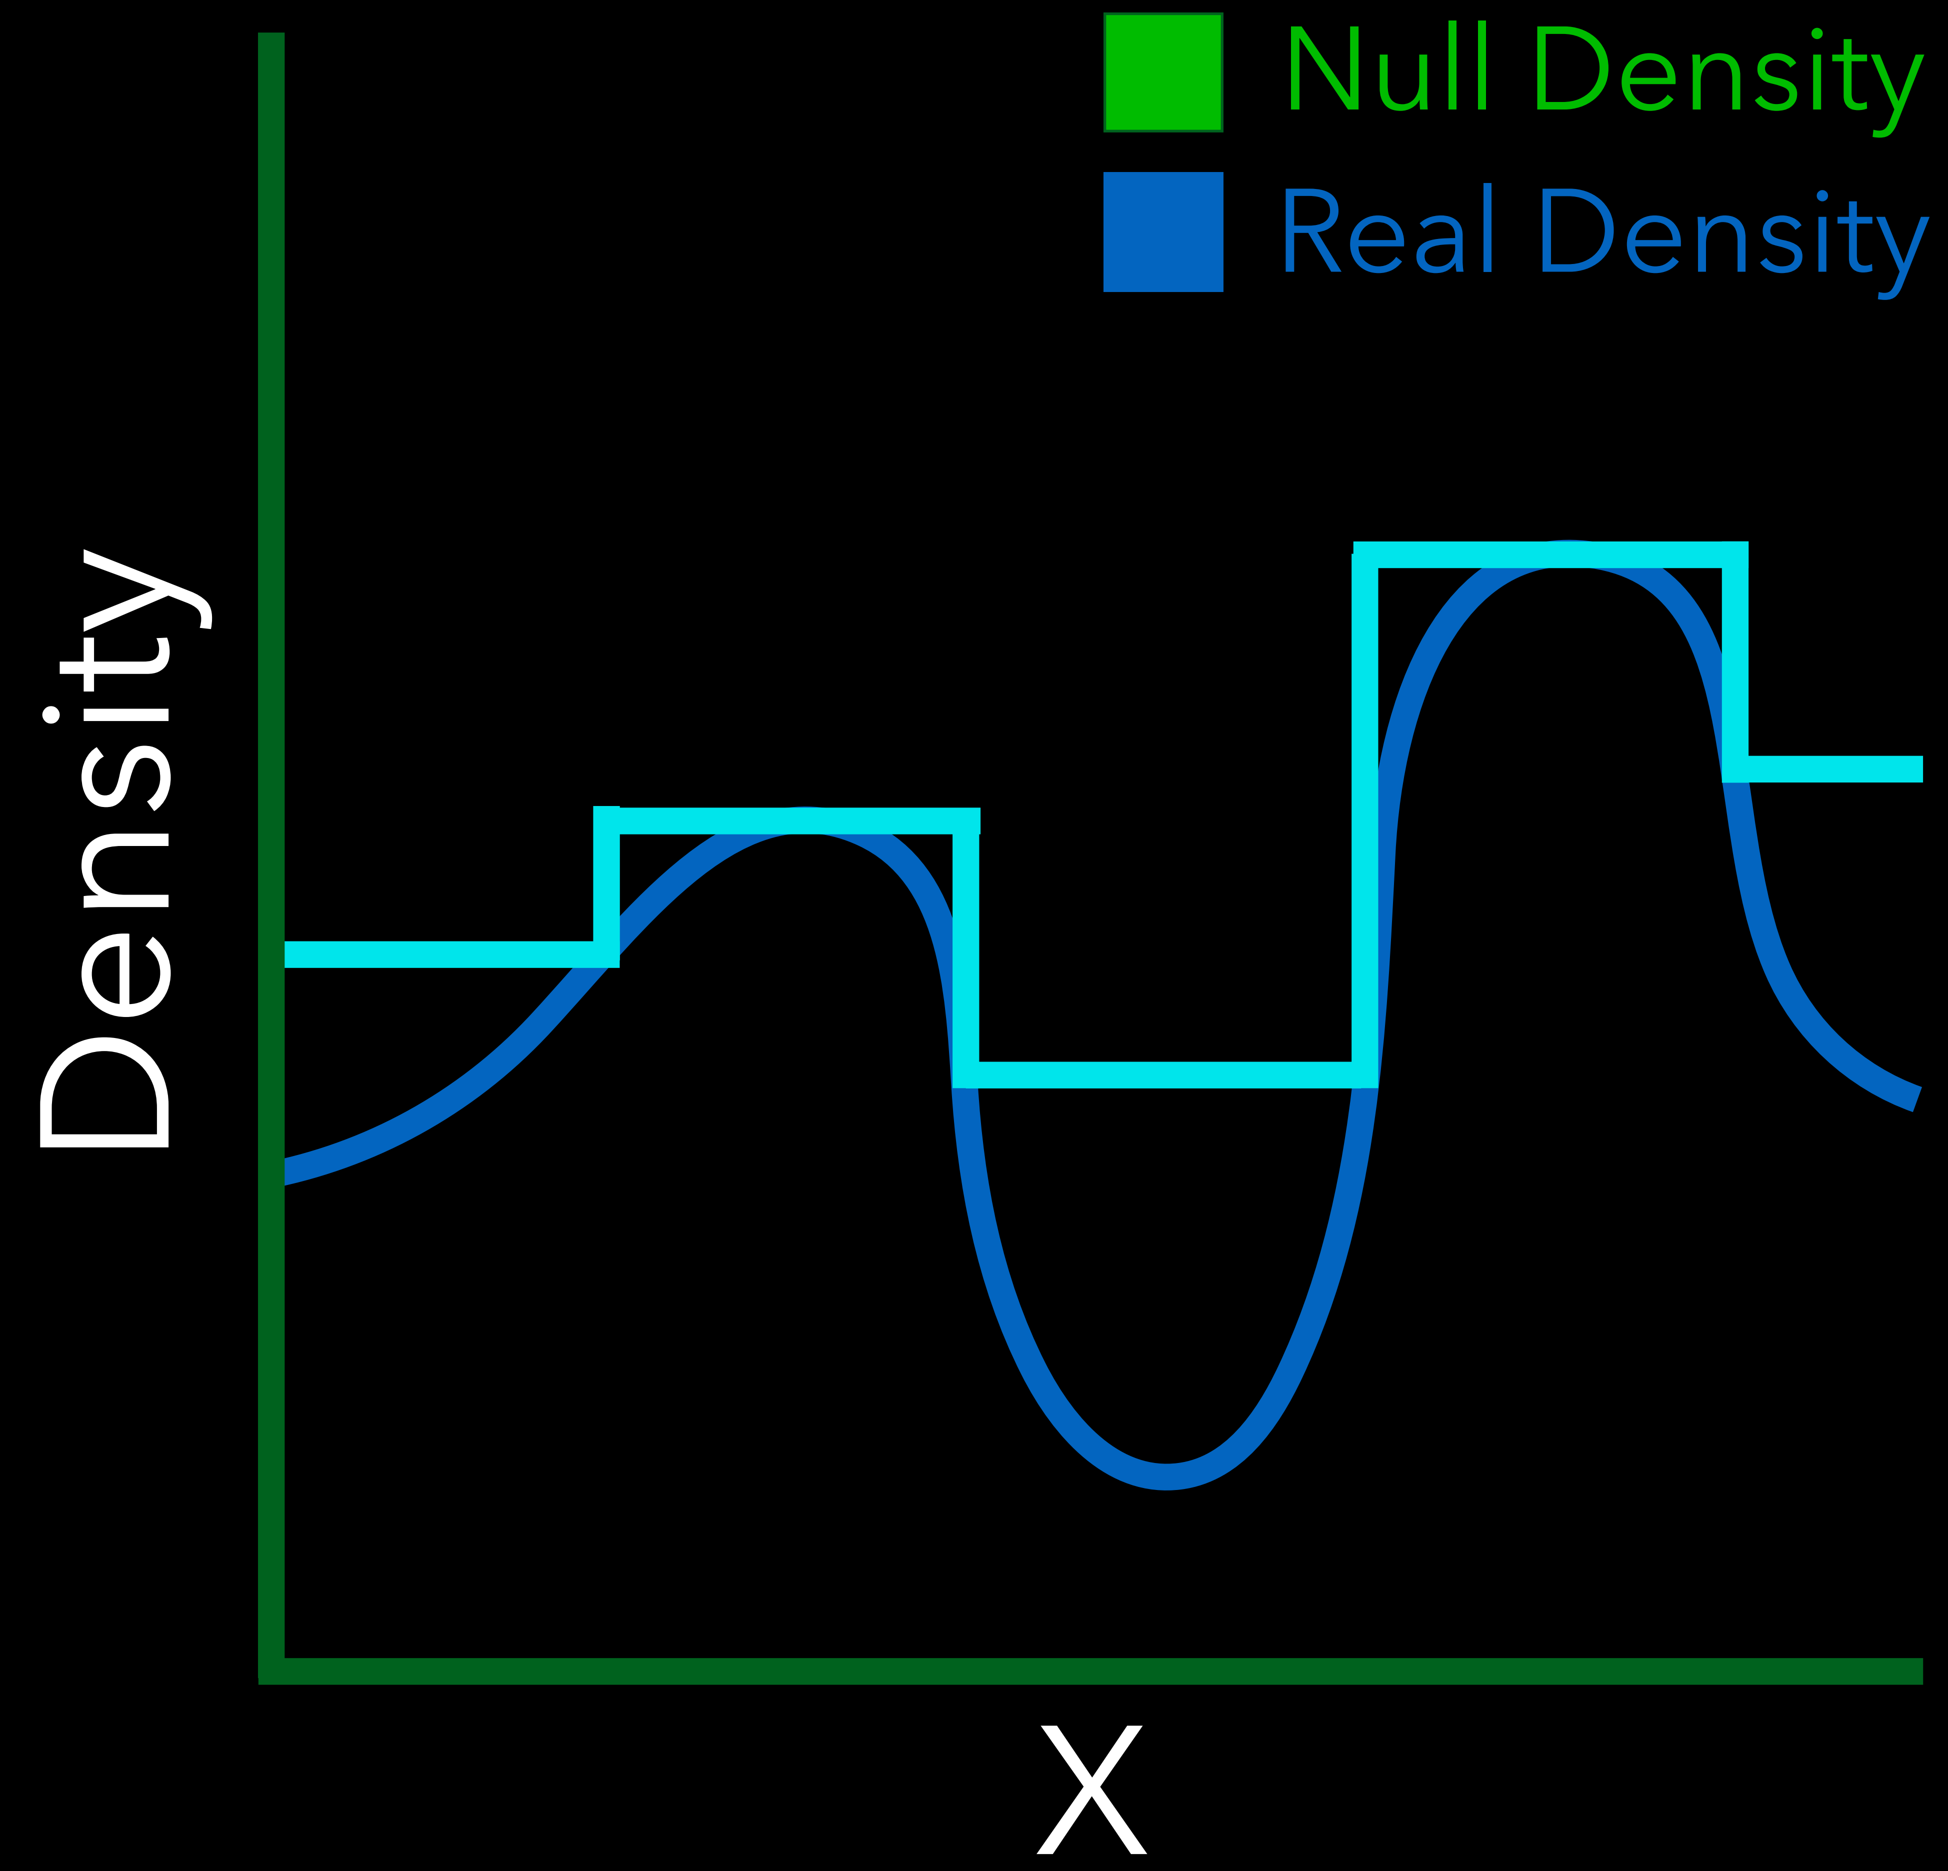
Null Density
Real Density
Density
X

Results

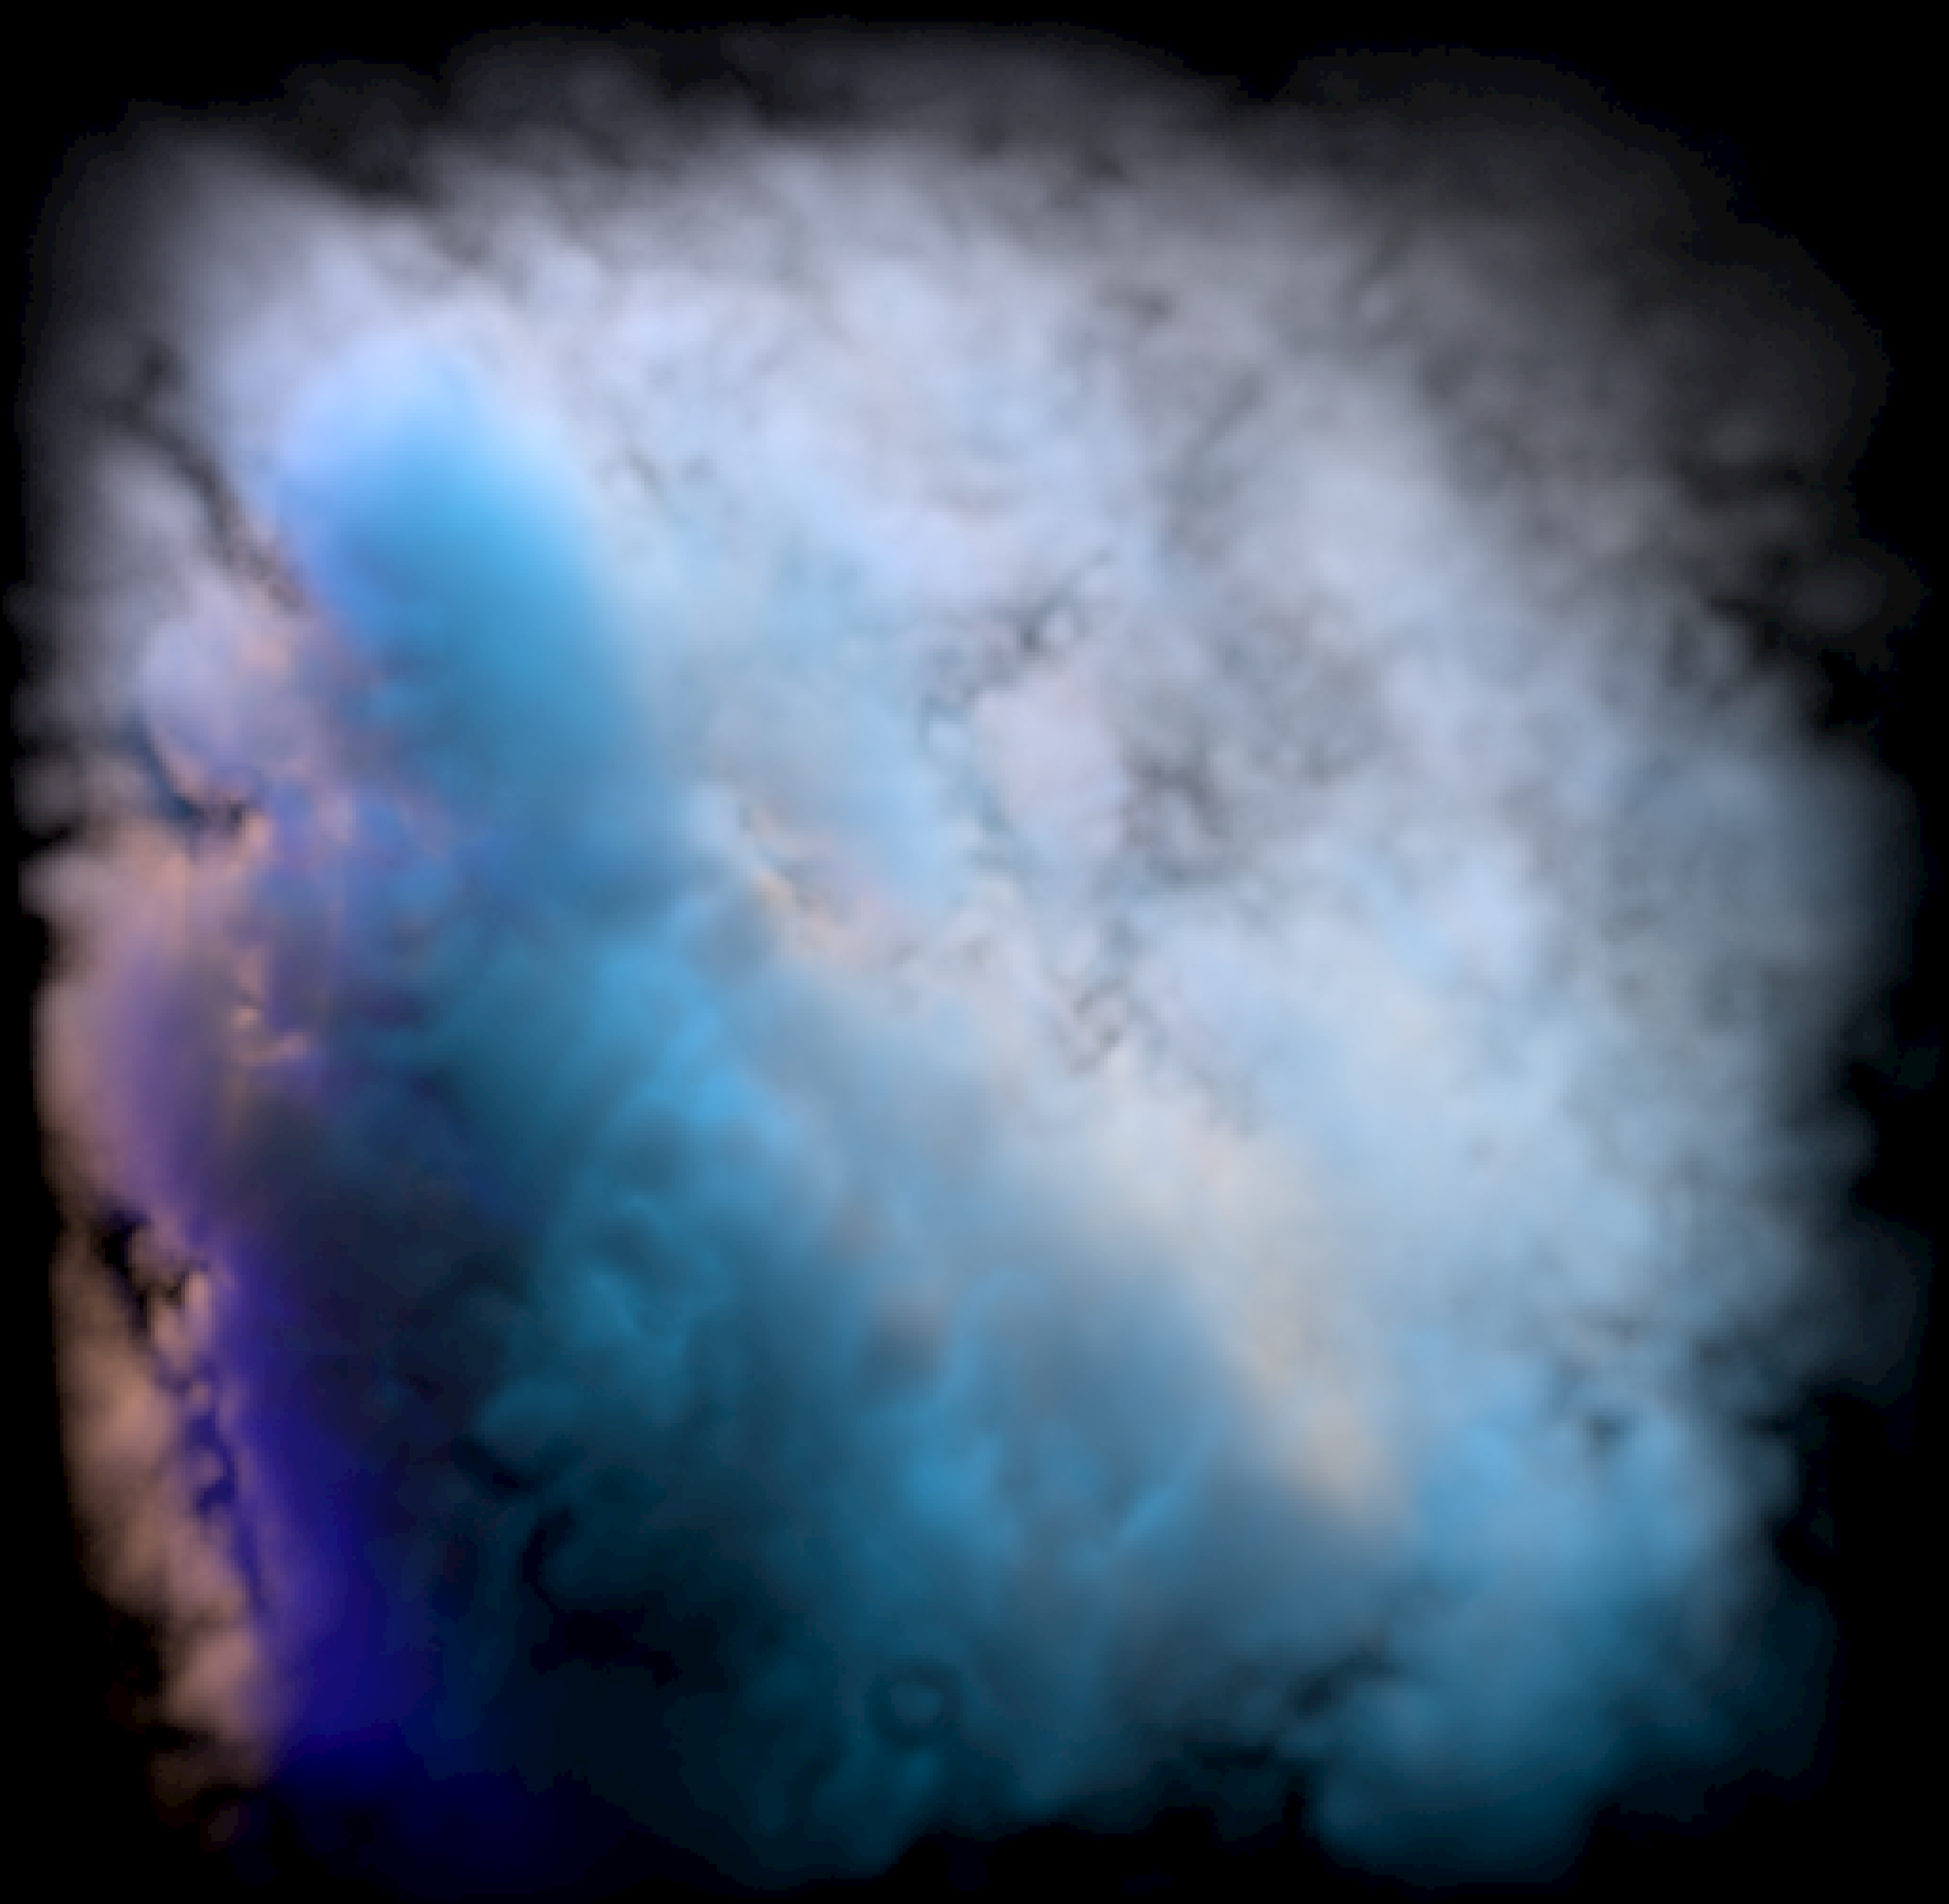

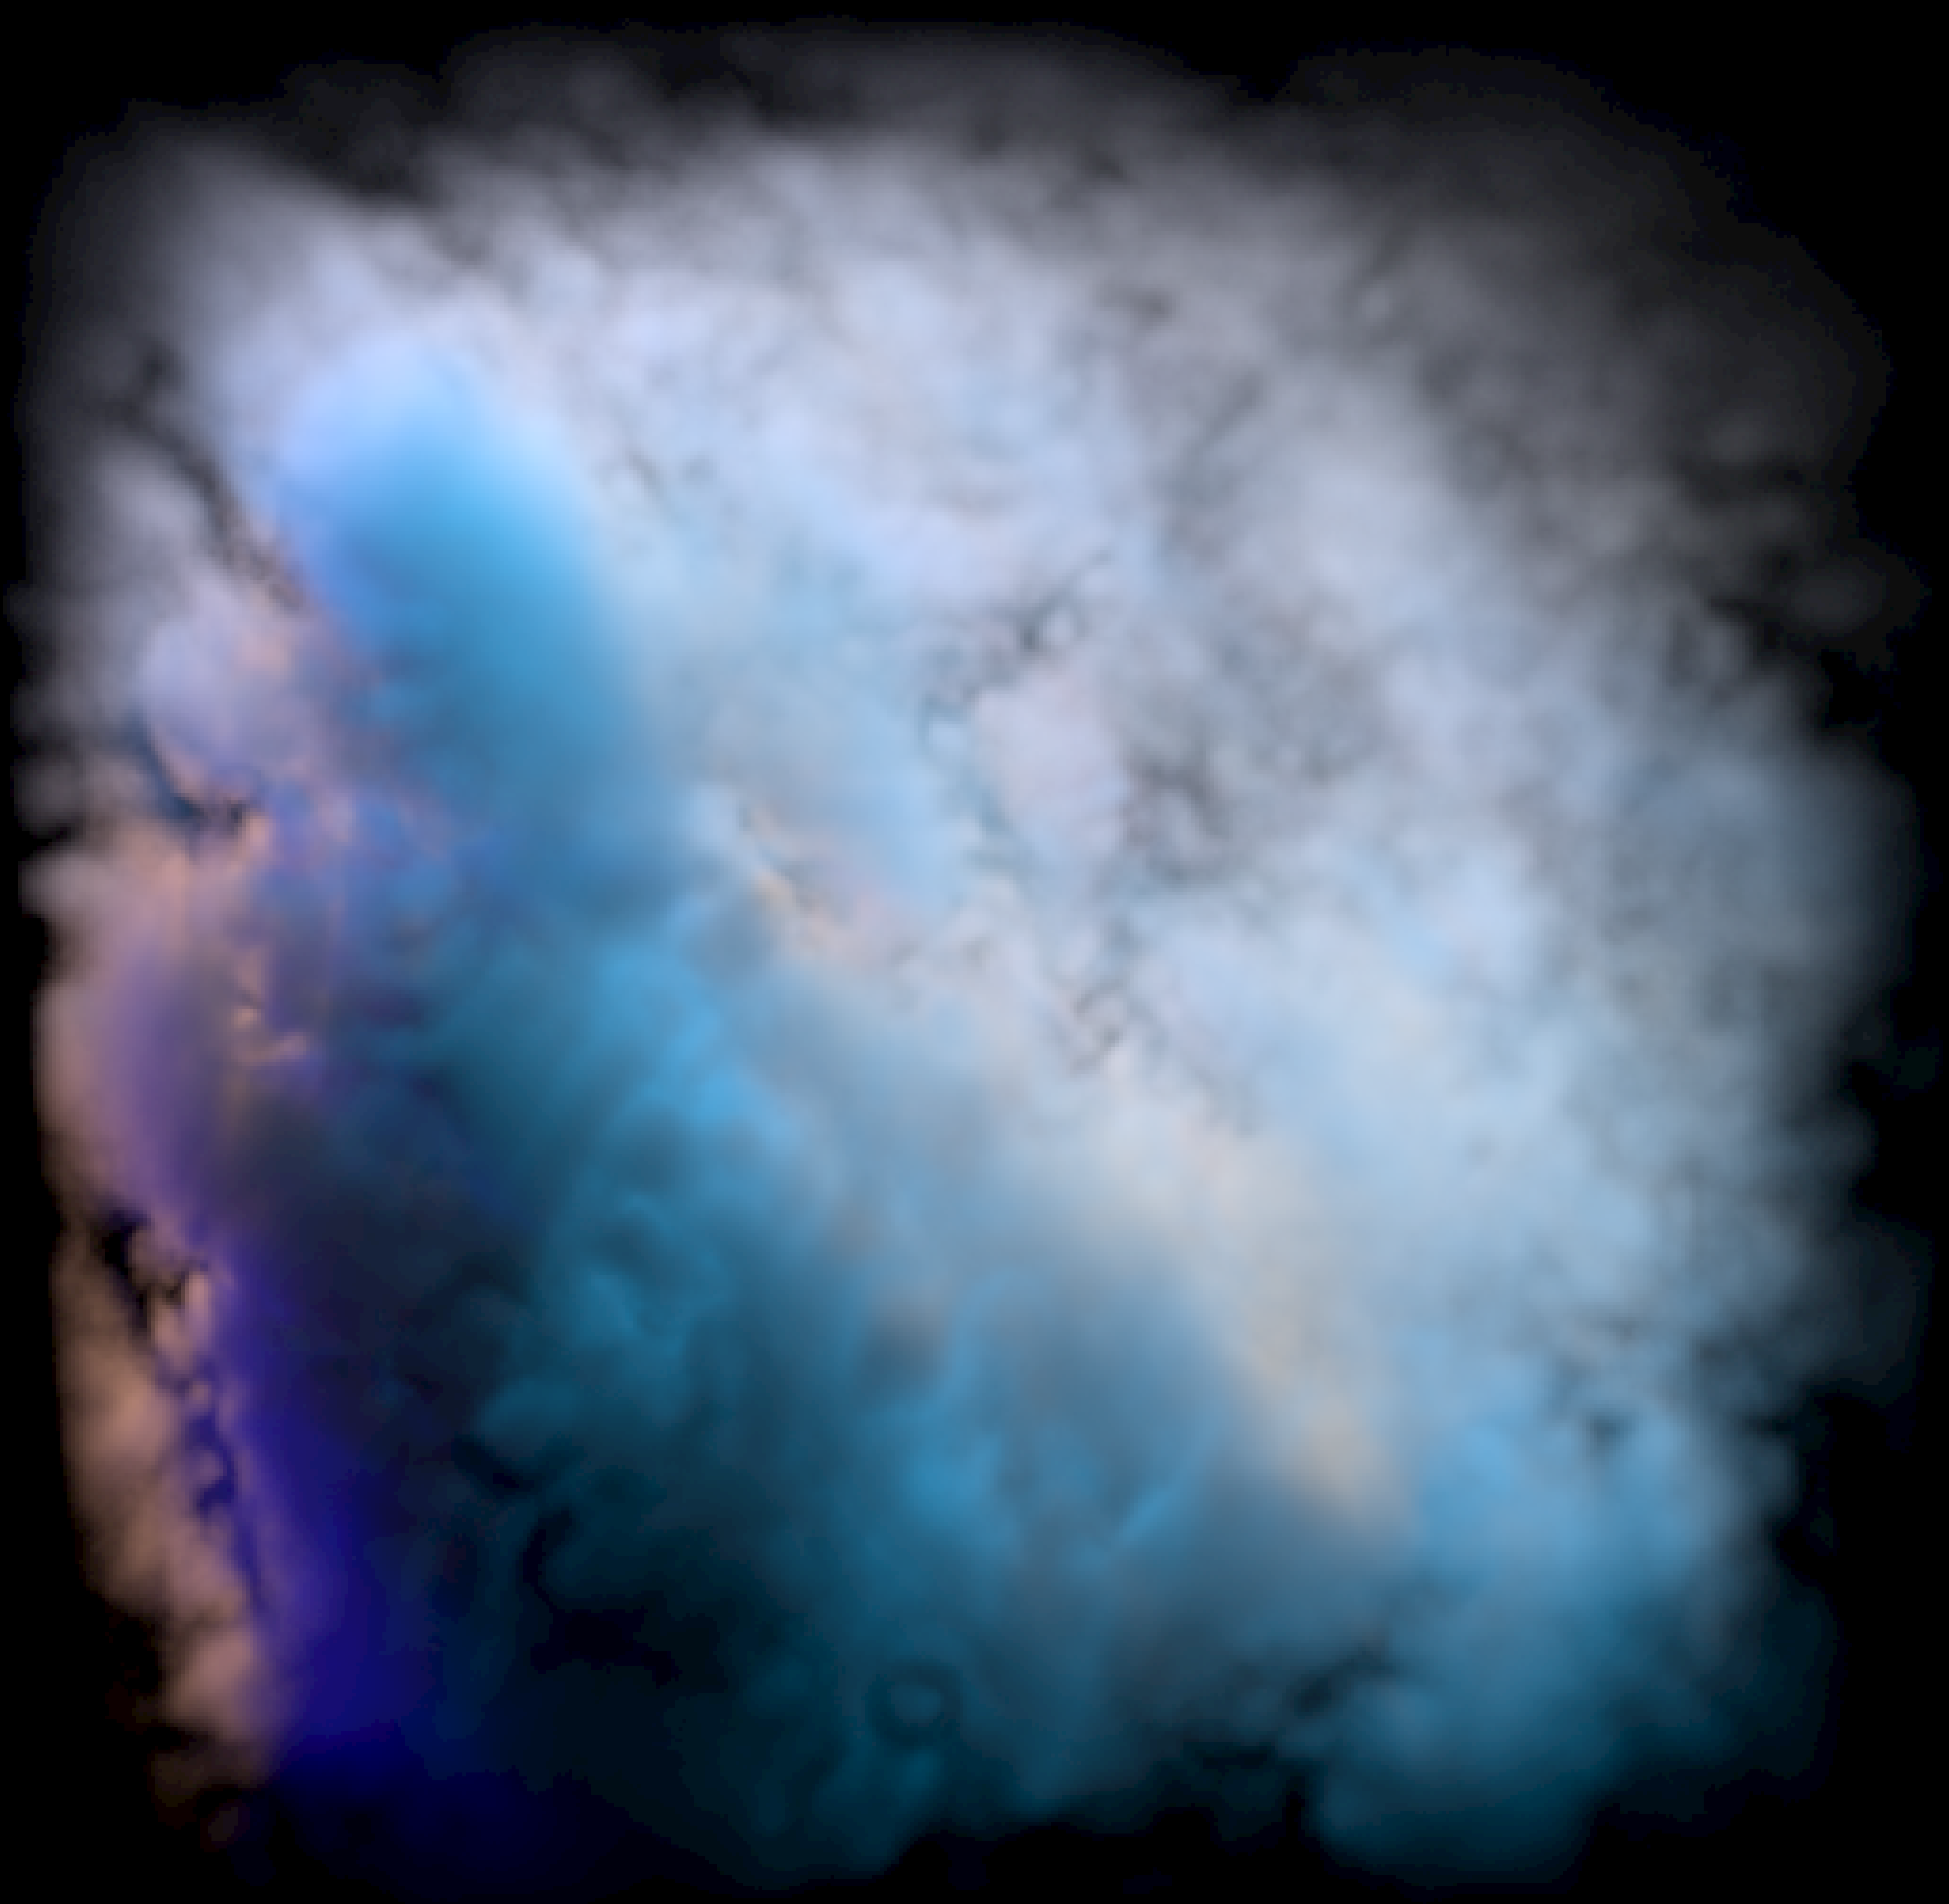

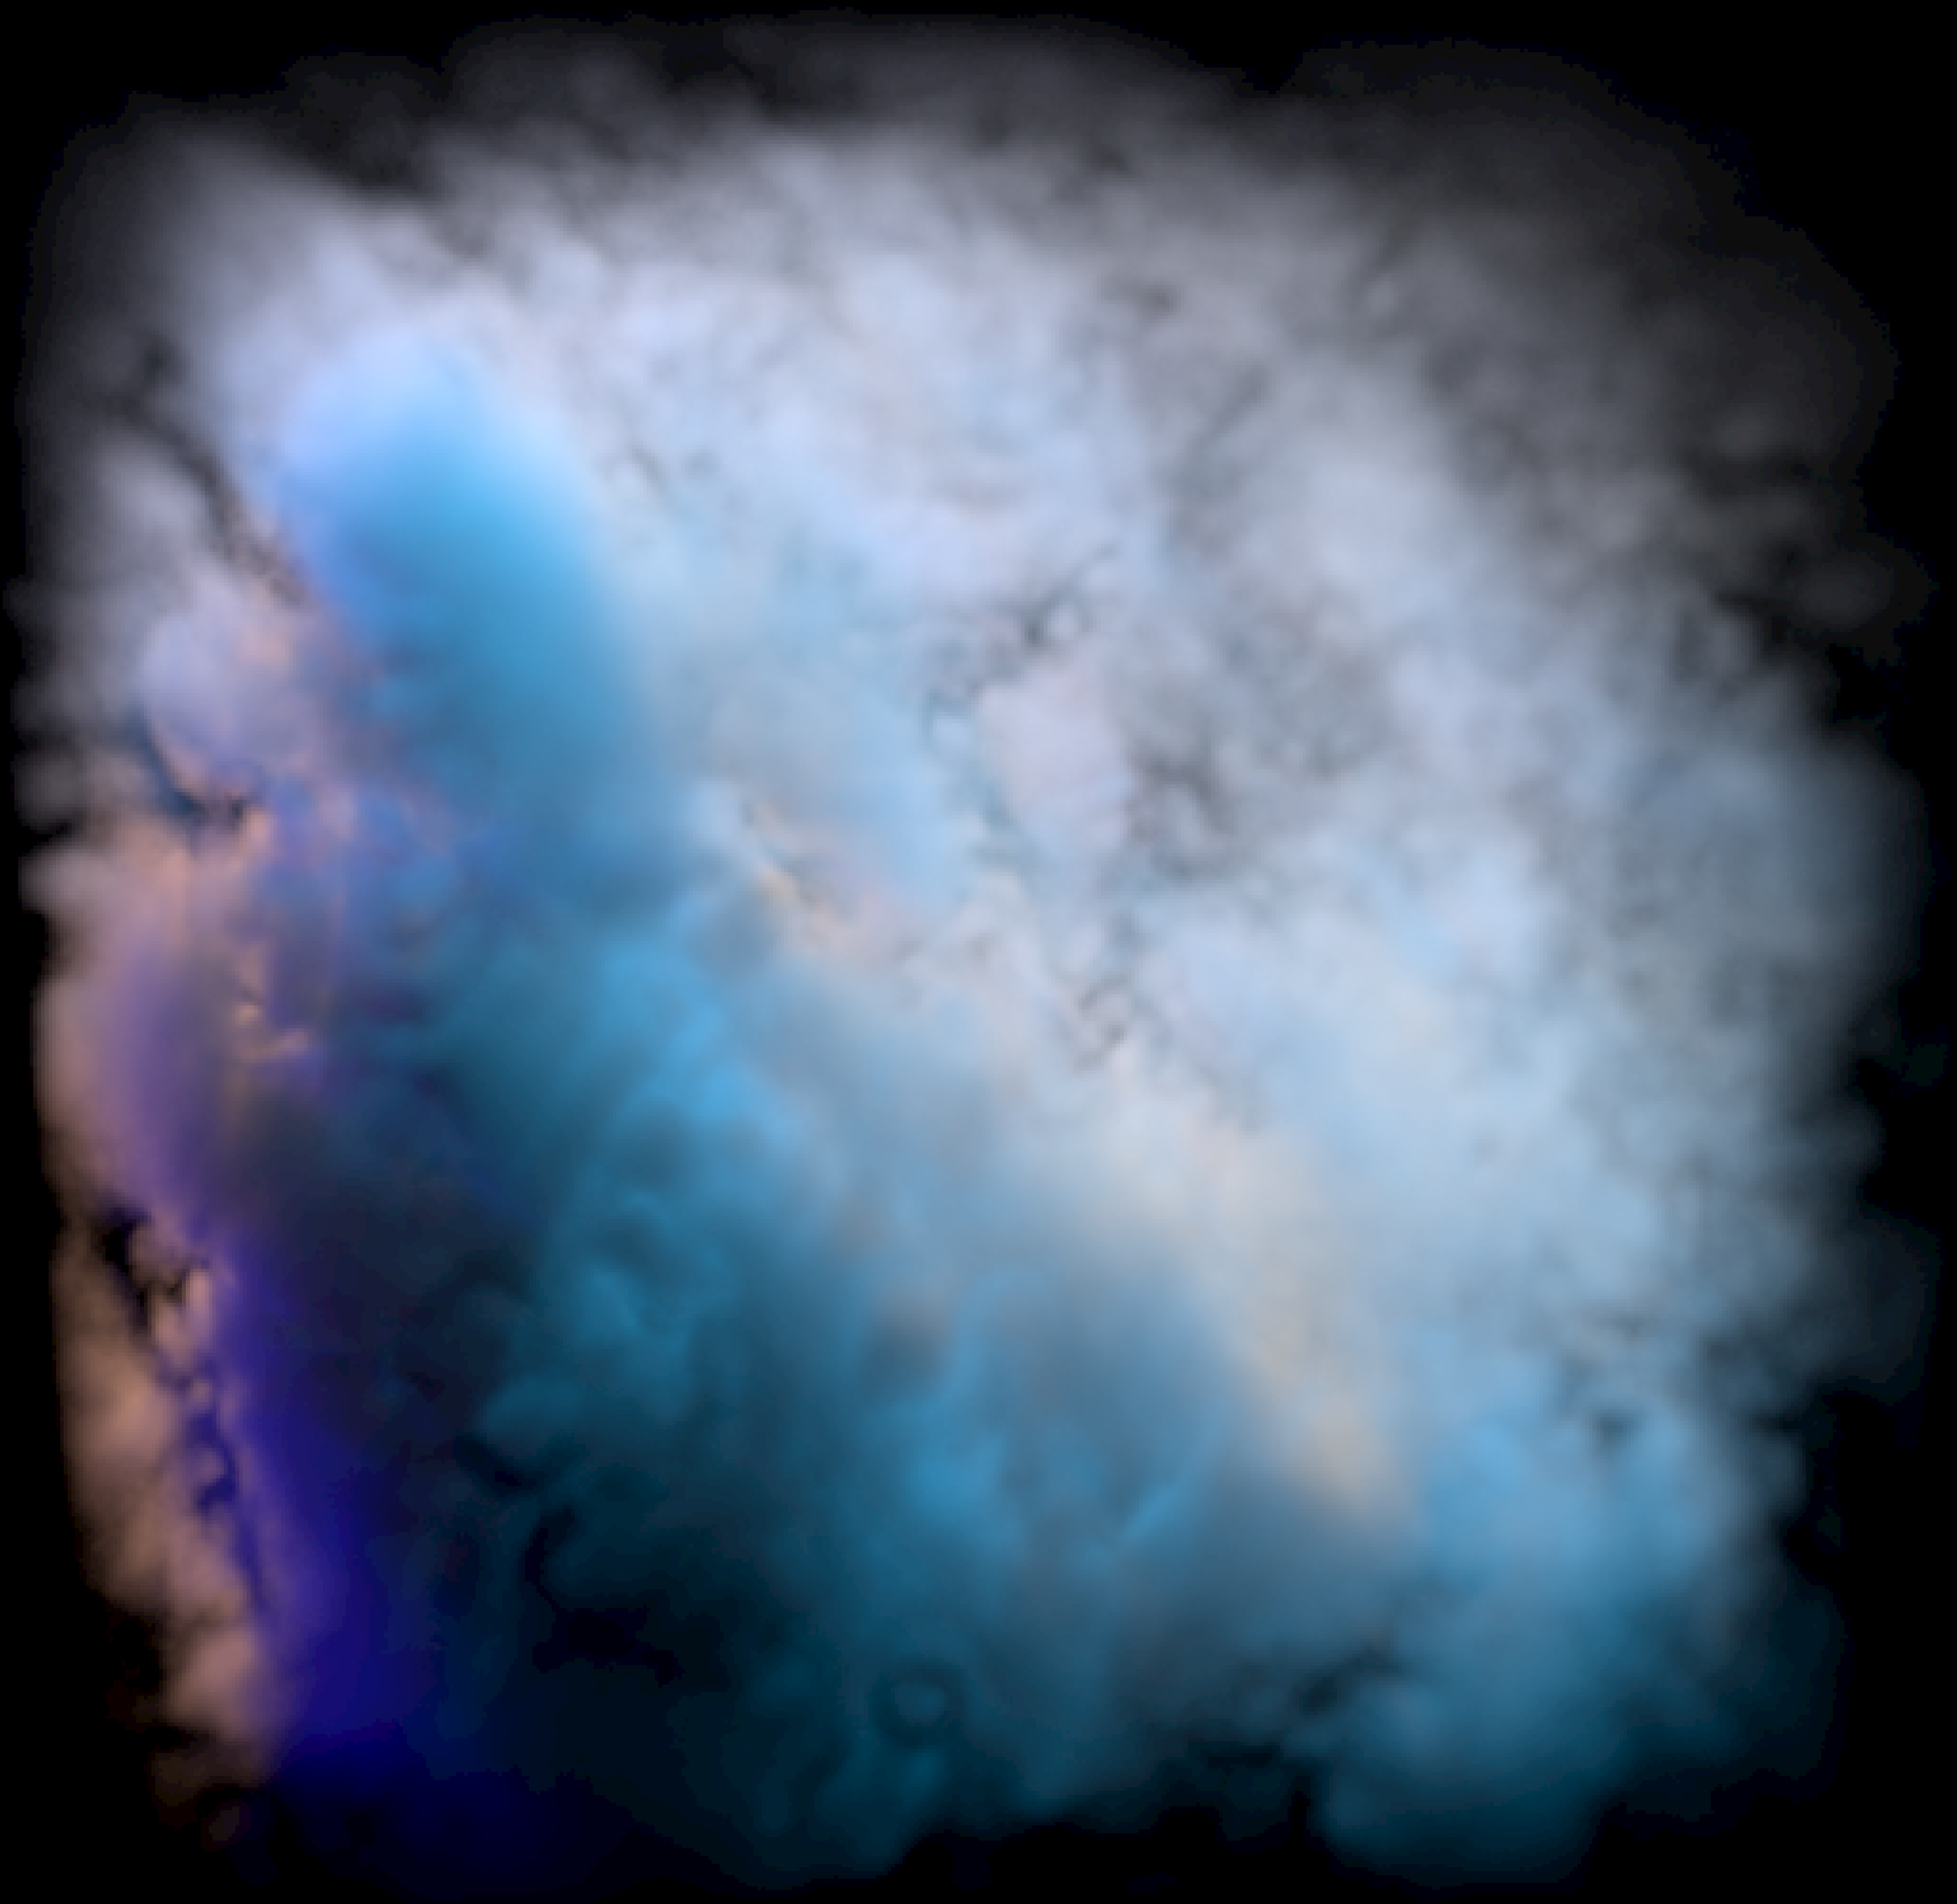

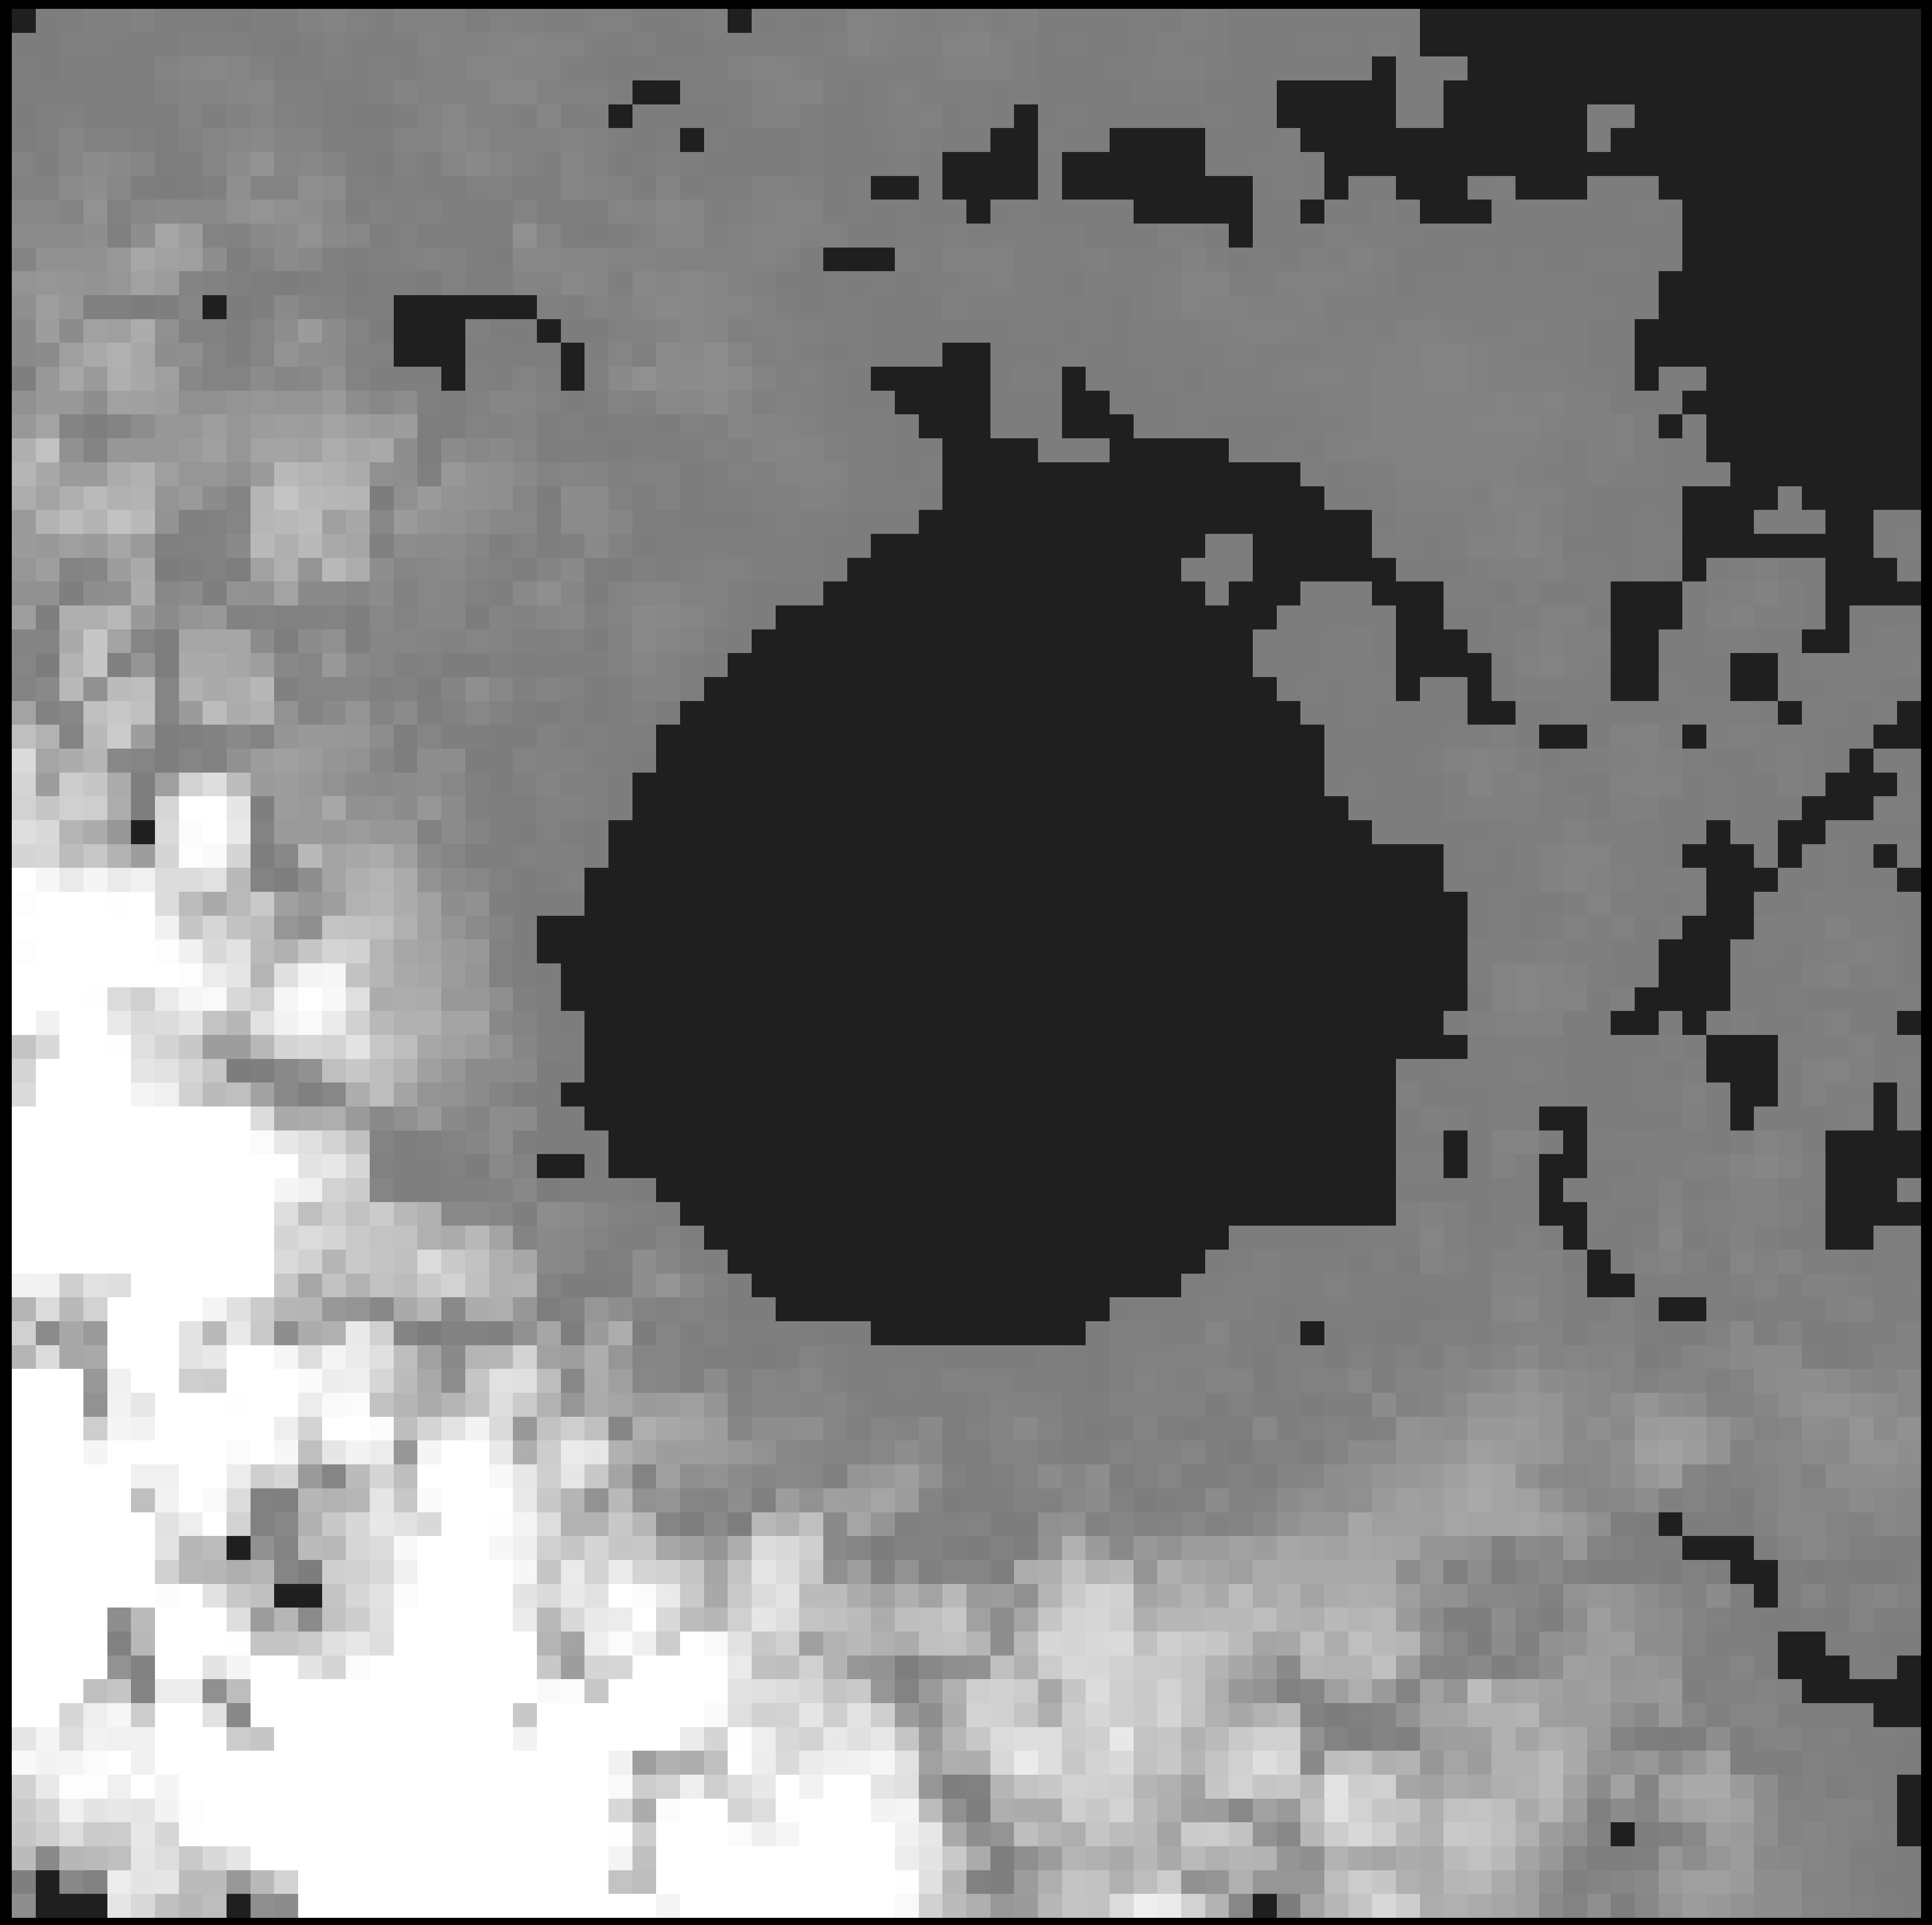


8 spp
64 spp
512 spp
4096 spp
Ratio

8 spp
64 spp
512 spp
4096 spp
Ours
Ratio

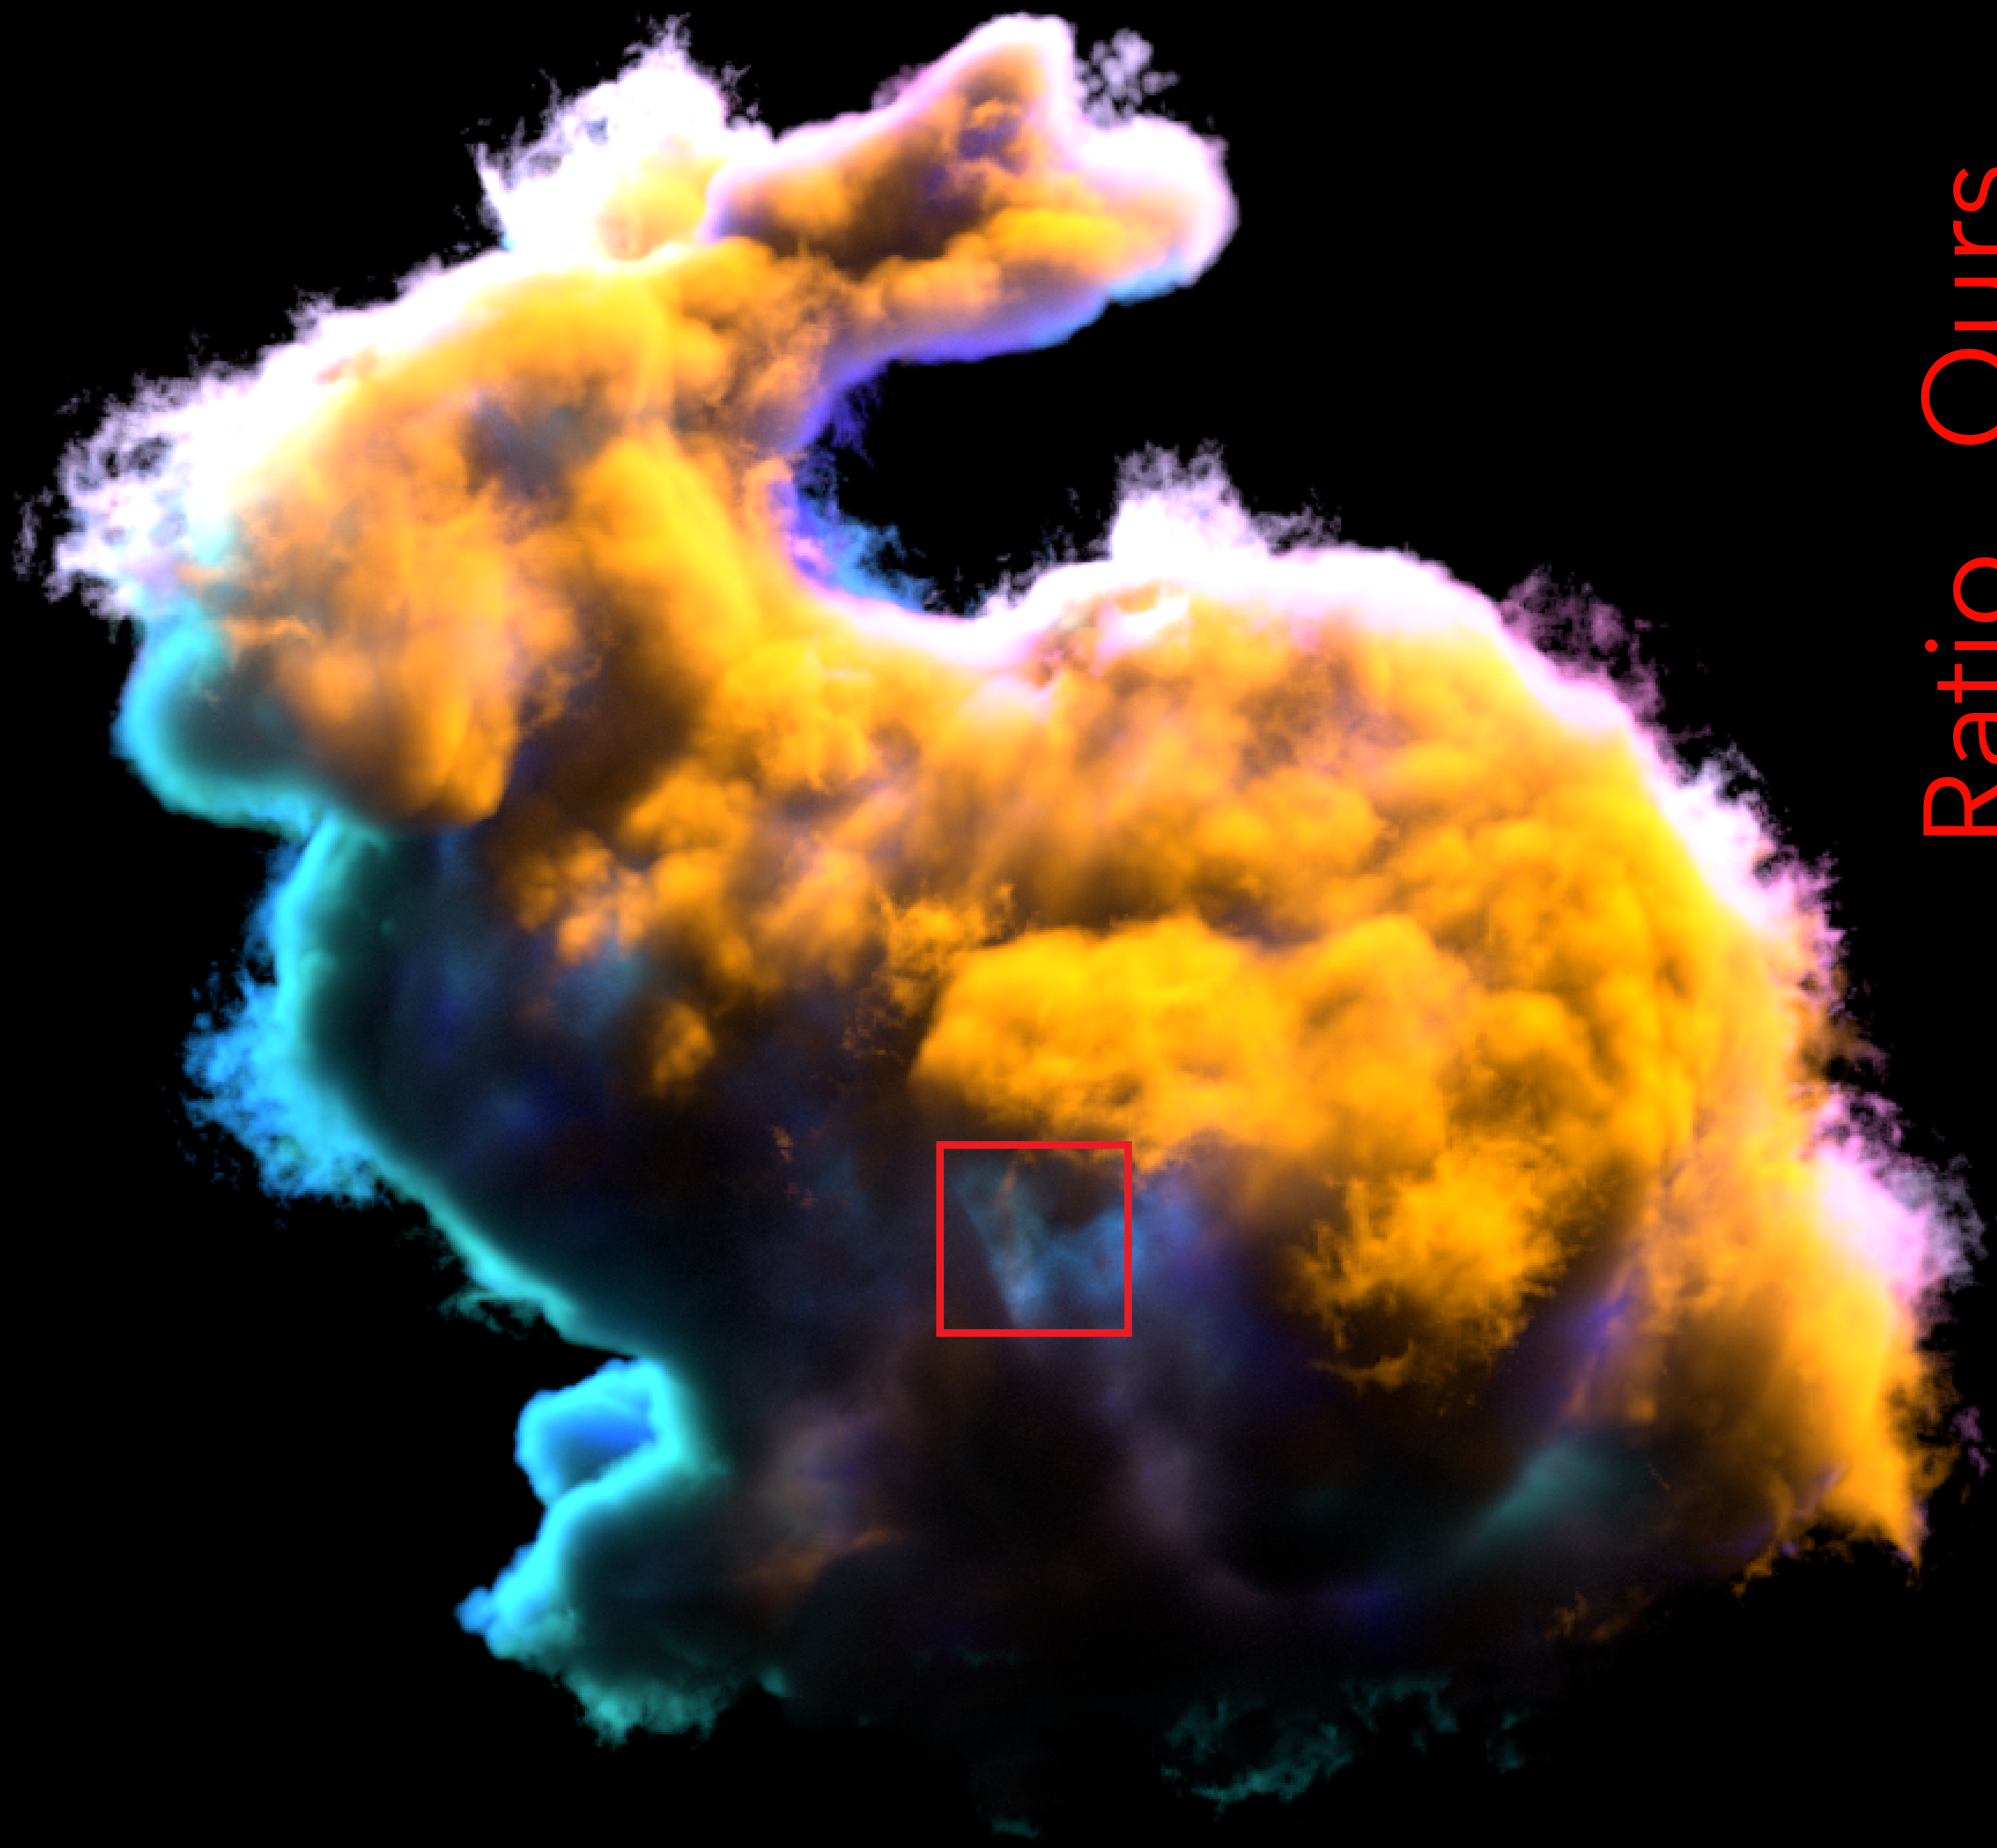

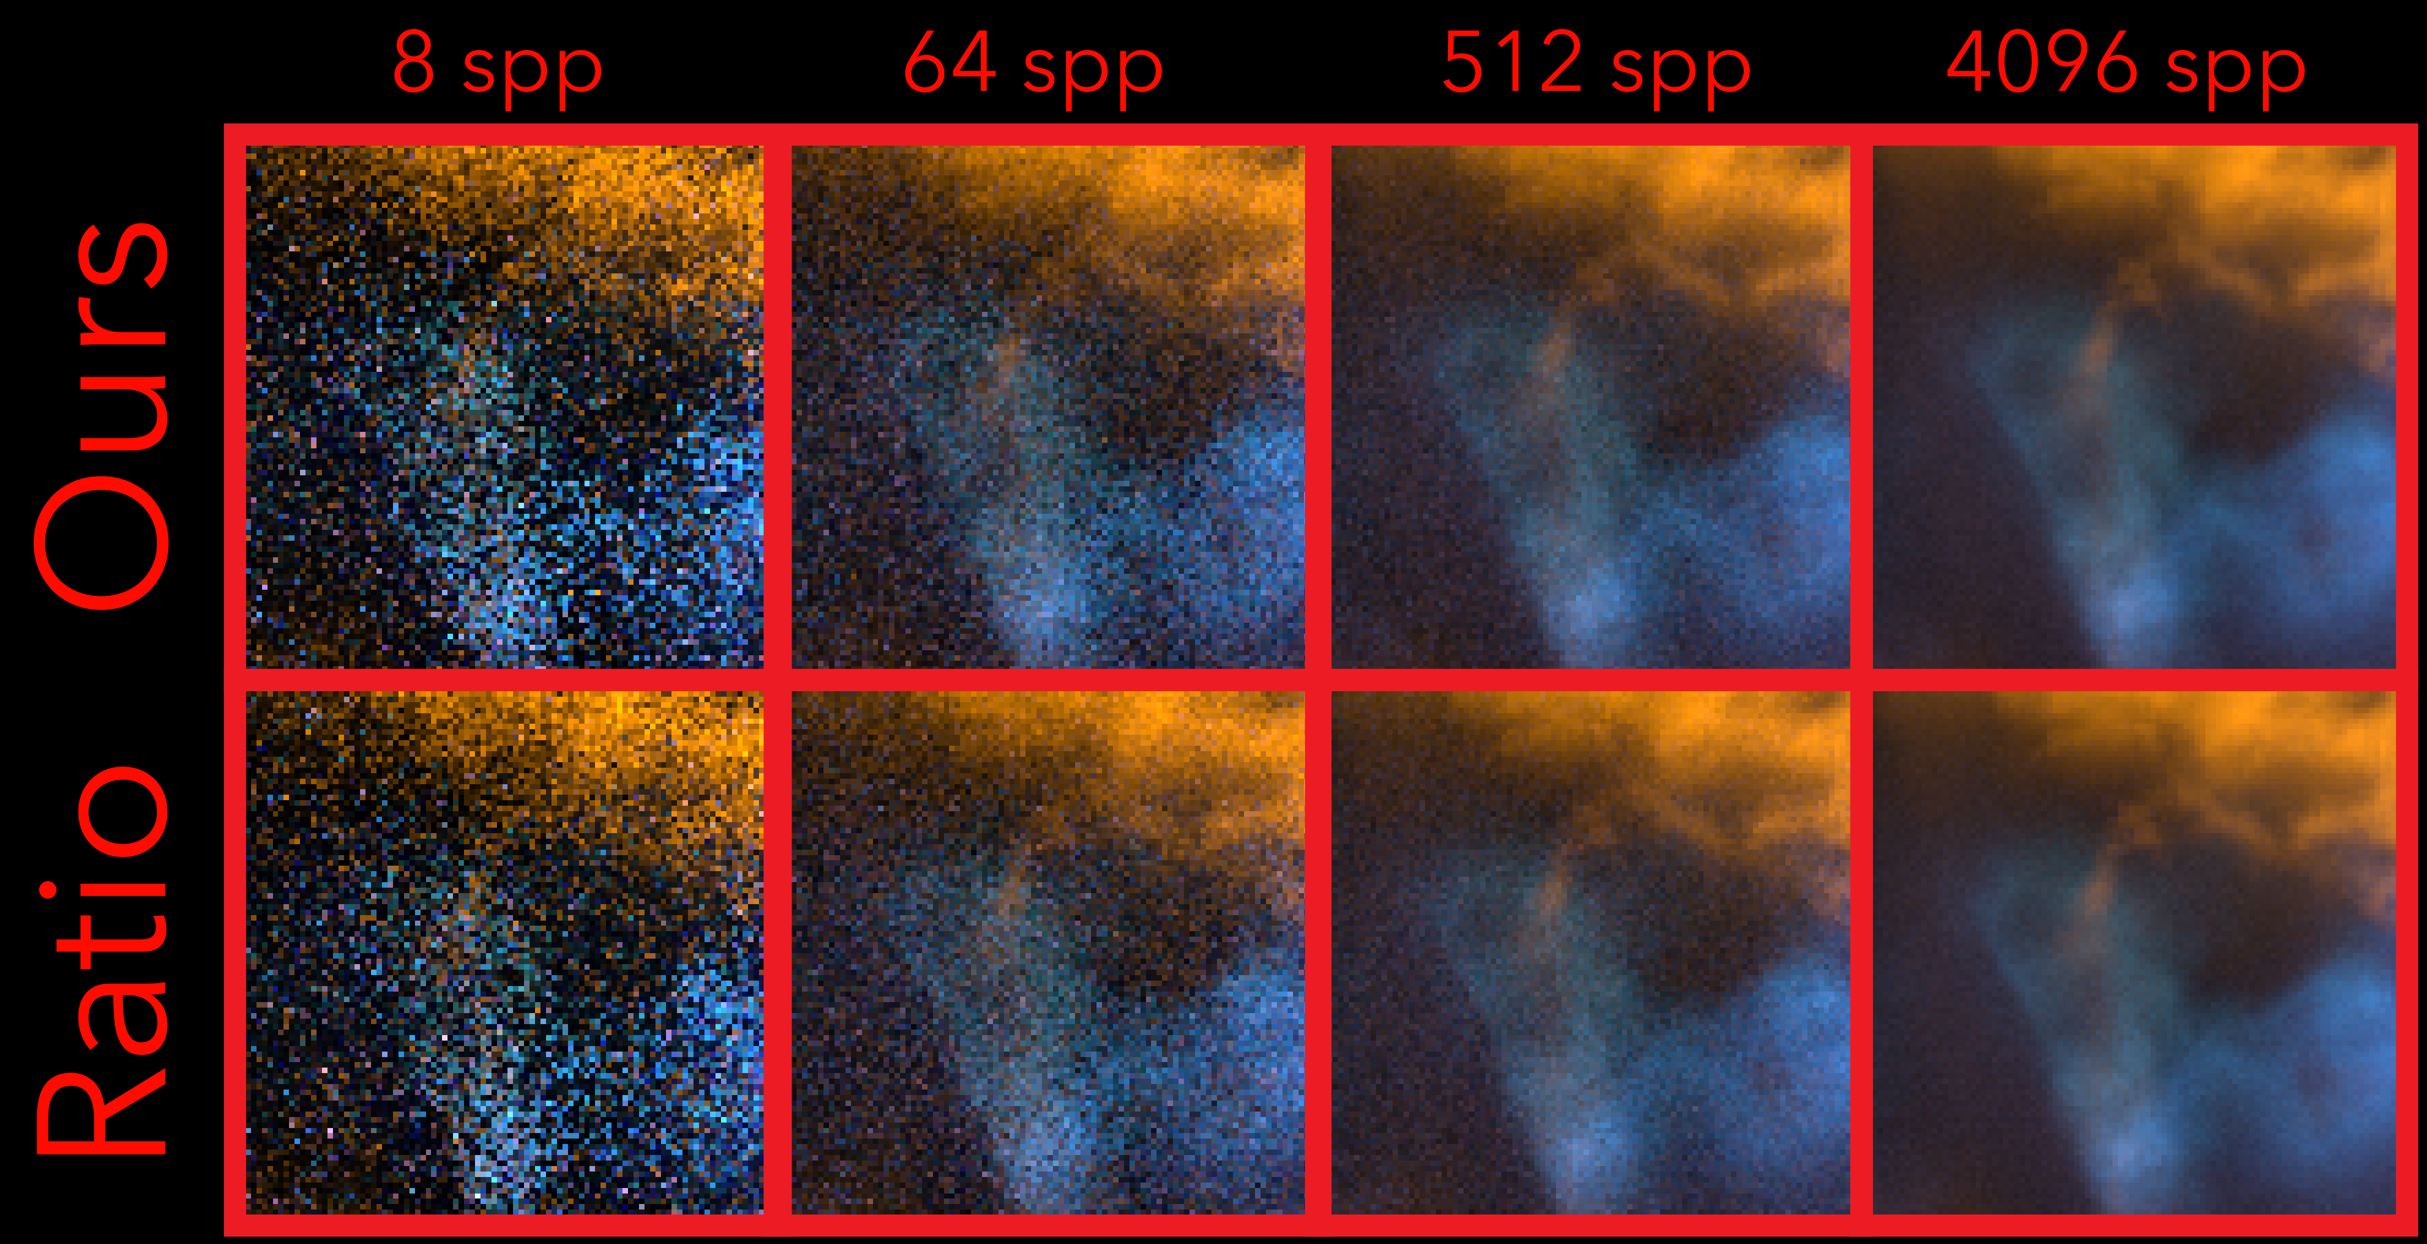

8 spp
64 spp
512 spp
4096 spp
Ours
Ratio

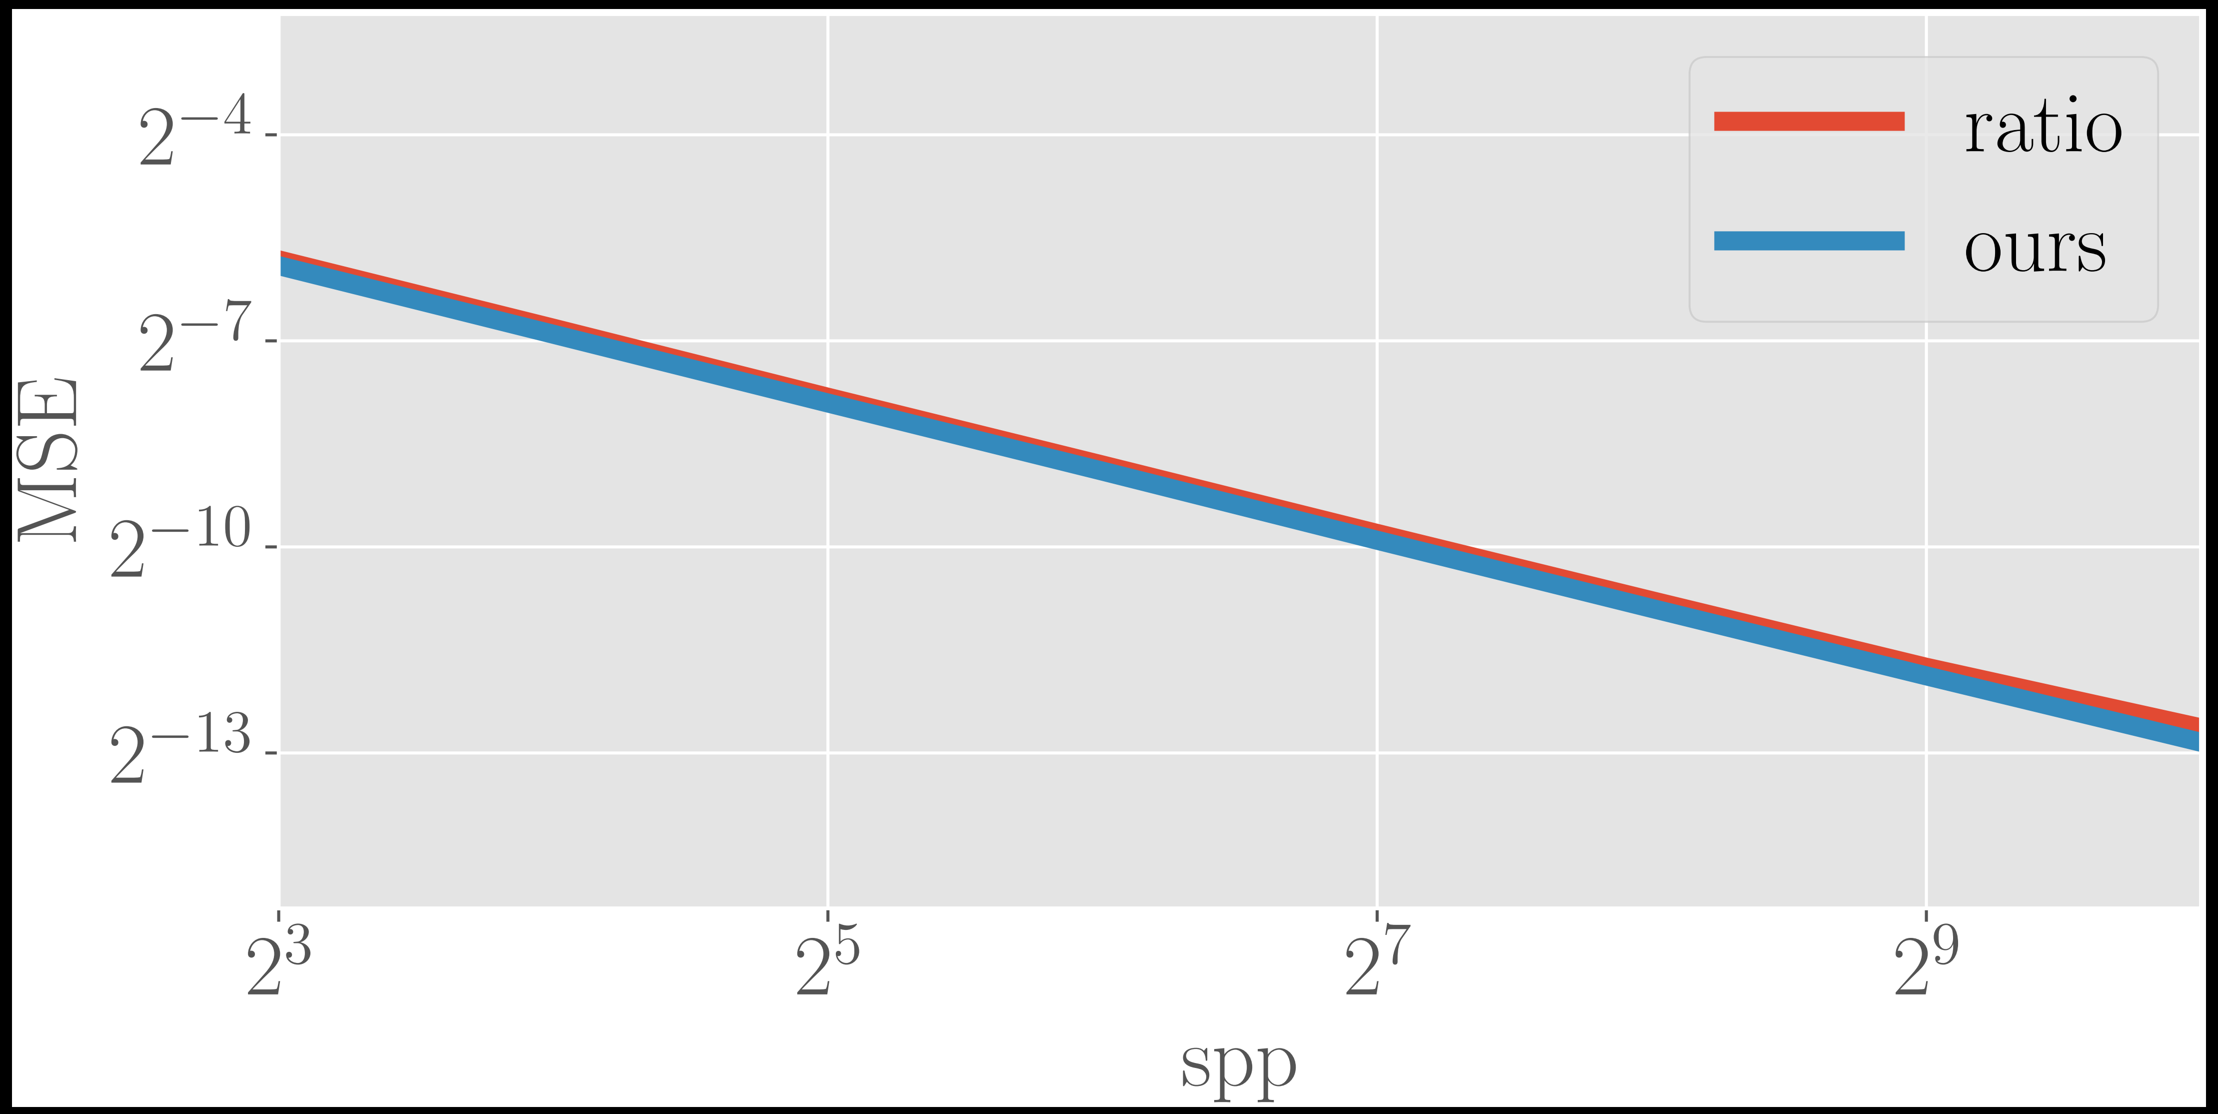

ratio
ours
MSE
2^{-4}
2^{-7}
2^{-10}
2^{-13}
2^3
2^5
2^7
2^9
spp


MSE: 1.4e-2
MSE: 5.4e-3
MSE: 5.9e-3
Unb. Ray
Ratio
Ours
Reference
[Kettunen et. al. 2021]
[Novak et. al. 2014]
(Progressive Ratio)

Conclusion

In the paper

In the paper

- Full analysis of explosive variance

In the paper

- Full analysis of explosive variance
- Adaptive ratio tracking

In the paper

- Full analysis of explosive variance
- Adaptive ratio tracking
- Proofs and convergence rates

Future work

Future work

- Residual

Future work

- Residual
- Better majorant updating

Future work

- Residual
- Better majorant updating
- Full incorporation into Hyperion


Thank You