https://doi.org/10.2352/EI.2026.38.10.HVEI-216

Improving Uniformity of Vertex Distributions in Spherical Sampling of Metamer Mismatch Body Boundaries

Alexander G. Forsythe[1,2,*], Brian Funt[1]
[1] Simon Fraser University, Vancouver, Canada
[2] Academy of Motion Picture Arts and Sciences, Hollywood, California
[*] Corresponding Author

Abstract

Metamer mismatch bodies (MMBs) quantify the extent of metamer mismatching for a given color stimulus with a change in color mechanism (i.e. change in illuminant and/or observer). Prior work has shown that the MMB boundary can be efficiently approximated by spherical sampling of unit directions in the 6D joint color-mechanism space, and for each sampled direction, maximizing the boundary point subject to the metameric cross-section constraints. Many sampled directions map to the same boundary vertex, so the number of recovered vertices is typically far smaller than the number of sampled directions. This produces a plausible approximation, but the resulting boundary vertices, expressed in sensor-response spaces (for example XYZ, LMS, or RGB), are often distributed in a highly non-uniform manner. Increasing the number of sampled directions increases the number of recovered vertices but does not improve boundary uniformity. We explored a simple post-processing workflow that builds a larger candidate pool of vertices and then selects a fixed-size subset using a spacing-driven sampling algorithm, improving vertex uniformity as measured by a nearest-neighbor metric. This approach substantially improves vertex uniformity in sensor space, but it can discard boundary-defining extreme vertices, potentially altering hull volume and other distinguishing boundary features. We therefore argue that any practical workflow for improving MMB vertex uniformity should include an explicit mechanism for retaining boundary-critical extremes prior to applying spacing-driven selection.

Introduction

A metamer mismatch body is the set of all possible color signals produced by a set of objects which yield an identical color signal under a color mechanism (i.e. the product of the illuminant spectral power distribution and the sensor spectral sensitivities), but a different set of color signals when the color mechanism is altered [5]. This set of divergent color signals produces a closed convex body in color space, which can be specified in terms of its boundary representation derived from Eqs. (1–2) [5, 1]. The points on this boundary are known as μ-optimal reflectances and are defined by object reflectances that are either zero or one at each wavelength [5].

These μ-optimal reflectances define the theoretical limits of how much color signals can disperse due to the variation in color mechanism [5, 6]. Accurately characterizing this boundary is fundamental to quantifying the potential for metamer mismatching in color science [5, 4].

Let $r(\lambda) \in [0,1]$ denote an object reflectance and let $\mathbf{s}(\lambda), \mathbf{s}'(\lambda) \in \mathbb{R}^3$ denote two color mechanisms. The corresponding color responses are:

$$\Phi(r) = \int r(\lambda)\,\mathbf{s}(\lambda)\,d\lambda, \qquad \Psi(r) = \int r(\lambda)\,\mathbf{s}'(\lambda)\,d\lambda. \quad (1)$$

Defining the joint color signal map

$$\Gamma(r) = \big(\Phi(r), \Psi(r)\big) \in \mathbb{R}^6, \quad (2)$$

the set $\Gamma(\mathscr{X})$, where $\mathscr{X} = \{r(\lambda) \in [0,1]\}$, forms a closed convex body in $\mathbb{R}^6$.

A boundary point of $\Gamma(\mathscr{X})$ in direction $\mathbf{k} \in \mathbb{R}^6$ with $\|\mathbf{k}\|_2 = 1$ can be expressed as a supporting-point maximization:

$$r_{\mathbf{k}} \in \arg\max_{r\in\mathscr{X}} \mathbf{k}^{\mathsf{T}}\Gamma(r). \quad (3)$$

For a fixed stimulus $\mathbf{z}_0 \in \mathbb{R}^3$ under the first mechanism, the corresponding metamer mismatch body is the cross-section

$$\mathscr{M}(\mathbf{z}_0) = \Big\{\Psi(r) \in \mathbb{R}^3 : \Phi(r) = \mathbf{z}_0,\ r \in \mathscr{X}\Big\}. \quad (4)$$

Equivalently, $\mathscr{M}(\mathbf{z}_0)$ is the projection onto the Ψ-coordinates of the slice of $\Gamma(\mathscr{X})$ satisfying $\Phi(r) = \mathbf{z}_0$.

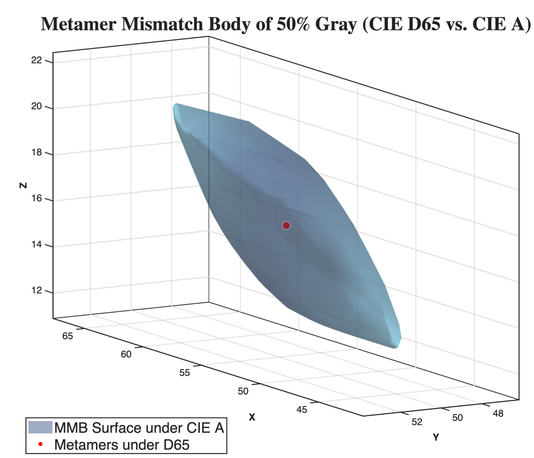

Figure 1. *Metamer mismatch body (MMB). The hull represents the set of color signals under a second mechanism for reflectances yielding $\mathbf{z}_0$ (center), forming a cross-section of the joint object-color solid defined by $\Phi(r) = \mathbf{z}_0$.*

Figure 1 illustrates this construction, showing the metamer mismatch body as a convex set in sensor-response space with the fixed color signal under the first mechanism indicated at the center.

In practice, MMBs are often constructed from the convex hull of a discrete set of vertices generated by solving optimization problems [1, 6]. One approach utilizes spherical sampling of unit directions in the 6D joint color-mechanism space where linear programming (LP) is used to identify the extreme points of the metamer set in those specific directions [8, 6, 1].

After discretizing the spectrum at q wavelengths, the reflectance becomes a vector $\mathbf{r} \in [0,1]^q$ and the maps in Eq. 1 become linear. Let $A \in \mathbb{R}^{q\times 3}$ and $B \in \mathbb{R}^{q\times 3}$ denote the discretized matrices such that $\Phi(r) = A^\top \mathbf{r}$ and $\Psi(r) = B^\top \mathbf{r}$. Then for a sampled direction $\mathbf{u} \in \mathbb{R}^3$ with $\|\mathbf{u}\|_2 = 1$, a boundary vertex of $\mathscr{M}(\mathbf{z}_0)$ can be obtained by solving

$$\max_{\mathbf{r}\in[0,1]^q} \mathbf{u}^\top B^\top \mathbf{r} \quad \text{subject to} \quad A^\top \mathbf{r} = \mathbf{z}_0. \tag{5}$$

Alternatively, a more efficient computational geometry approach first constructs a half-space representation of a 6-dimensional joint object-color solid (OCS) and calculates the MMB as a cross-section through half-space intersection [6, 7]. However, empirical analysis has shown that spherical sampling often results in a large number of duplicated boundary points and redundant vertices [8].

This duplication arises because a single vertex can optimize Eq. 5 for a nontrivial set of directions. In particular, for a polytope vertex $\mathbf{v}$ there exists a normal cone $\mathscr{N}(\mathbf{v})$ such that

$$\mathbf{u} \in \mathscr{N}(\mathbf{v}) \;\Rightarrow\; \arg\max_{\mathbf{r}} \text{Eq. } 5 = \mathbf{v}, \tag{6}$$

so many sampled directions map to the same boundary vertex.

Applications leveraging MMBs for color imaging and lighting design, such as the visualization of color-constancy limits or the calculation of mismatch indices for light source evaluation, benefit from a more even vertex coverage across the MMB boundary [5, 6, 4]. High-quality metrics for applications like camera sensor design require that the discrete sample of vertices be dense and well-distributed to avoid significantly underestimating the actual mismatch volume [6, 7]. Comparative studies between imaging systems rely on a stable, and optimally repeatable, numerical representation of MMBs [3]. MMBs comprised of vertices that are clustered unevenly or contain a high number of duplicate vertices may be suboptimal [8, 5].

In this paper we report on the observed vertex distribution behavior in spherical sampling methods proposed by Mackiewicz et al. and offer specific considerations for subset selection to enhance MMB boundary vertex uniformity. By examining the relationship between sampled normal direction vectors and the resulting vertices, we address the common practical issue that may arise from many sampling directions in k-space producing repeated or near-repeated boundary vertices [8, 6]. Accounting for this behavior during subset selection helps to avoid the clustering of boundary points. Our results suggest that any spacing-driven subset selection intended to improve vertex uniformity may benefit from the inclusion of an explicit mechanism for retaining boundary-critical extreme points before applying uniformity-driven sampling.

Methods

In order to compute the metamer mismatch body boundary, we used the half-space intersection method described by Mackiewicz et al. [6]. The method involves generating a large set of directions on the unit hypersphere in k-space and treating those directions as candidate half-space normals [6]. We denote the number of sampled directions by N_k. After half-space intersection, multiple sampled directions may map to the same boundary vertex [8]. These normals define a 6D half-space representation of the joint OCS, which is then subjected to a set of linear inequality constraints that bound the region of feasible metameric responses [6]. The intersection of these half-spaces with the metameric slices defined by the inequality constraints produces a set of boundary vertices in the reduced three-dimensional sensor response space [6]. These vertices serve as the baseline estimate of the MMB boundary for our problem.

Diagnostics for Sampling Quality and Variability

In order to characterize the uniformity of the vertices used to compute the MMB boundary we compute several metrics for each experimental trial. First, we record the number (N_{vertices}) of intersection vertices returned by the half-space and slice intersection procedure. Measuring N_{vertices} provides a measure of how many unique boundary vertices are recovered for a given number of sampled directions (N_k). Note that many directions may yield the same vertex. A very low N_{vertices} value would indicate a failure to adequately resolve the boundary for the chosen sampling density or feasibility constraints. We hold all parameters fixed across all experimental trials including, color mechanisms, wavelength sampling, metameric slice constraints, and half-space intersection formulation. Differences between runs are caused by the random sampling of the k-space normals, which governs which boundary vertices are found. All analyses in this work are restricted to three-dimensional metamer mismatch body cross-sections defined by fixing the response under the first color mechanism. Higher-dimensional properties of the joint object–color solid are not considered here.

Second, we quantify boundary sampling density using nearest-neighbor distance statistics in the sensor response space of the second color mechanism. For each vertex, the Euclidean distance to its closest neighbor is computed. We denote the mean nearest-neighbor distance by $\bar{d}_{\text{NN}}$ and use the coefficient of variation CV_{NN} to quantify boundary non-uniformity.

Let $d_{\text{NN},i}$ denote the Euclidean distance from vertex i to its nearest neighboring vertex. We define the mean nearest-neighbor distance as

$$\bar{d}_{\text{NN}} = \frac{1}{N_{\text{vertices}}} \sum_{i=1}^{N_{\text{vertices}}} d_{\text{NN},i},$$

and the corresponding coefficient of variation as

$$\text{CV}_{\text{NN}} = \frac{\sigma(d_{\text{NN},i})}{\bar{d}_{\text{NN}}},$$

where $\sigma(d_{\text{NN},i})$ denotes the standard deviation of the nearest-neighbor distances across all boundary vertices.

Finally, in order to detect gross numerical issues or instability across runs, we compute the volume of the convex hull vertex set in the sensor space of the second color mechanism (e.g. LMS, XYZ, RGB). This hull volume is not assumed to be a ground-truth estimate of the true MMB volume, but rather serves as a consistency check for computational problems. Large deviations in hull volume across trials may indicate under-sampling, missing boundary regions in the intersection stage. All diagnostic metrics for the experiments are computed after mapping the boundary vertices into the three-dimensional sensor response space of the second color mechanism (i.e. the space in which the metamer mismatch body is defined). The MMB volume is used only as a relative diagnostic for consistency across trials and is not interpreted as an estimate of the true metamer mismatch volume.

Pool-and-subset method

To mitigate the variability inherent in a single spherical sampling run without increasing the computational cost of an extremely large set of sampled directions, we used a pool-and-subset strategy. Rather than using a single run of the spherical sampling algorithm, we repeat the sampling multiple times using independent random seeds. This ensures each run produces a partially distinct set of boundary vertices. These vertex sets are then concatenated to create a pooled candidate set. The pooled set provides a more comprehensive collection of candidate vertices than any individual run because different random sampling directions tend to populate different regions of the boundary.

A fixed-size subset of vertices is selected from this pooled set with the goal of improving uniformity of boundary coverage. The final subset selection is performed using farthest point sampling (FPS), which iteratively adds the candidate point that maximizes its distance to the current selected set [2]. This algorithm fills gaps between existing points and increases the minimum separation between selected vertices producing a more even distribution of the boundary than a purely random selection with the same number of vertices. Farthest point sampling is repeated with multiple initial seeds, and the most uniform subset is selected.

Farthest point sampling improves boundary uniformity but may discard geometrically important boundary features when applied directly to a pooled vertex set. In particular, vertices corresponding to corners or strongly protruding regions of the metamer mismatch body may be lost, even though these points define the boundary extent.

To prevent this, we explicitly preserve extreme boundary vertices prior to subset selection. Given a pooled set of boundary vertices expressed in the sensor response space of the second color mechanism, we evaluate the boundary response along a fixed set of unit directions. For each direction, the vertex with the maximum projection is identified and retained. To ensure that the subsequent spacing stage retains sufficient degrees of freedom, we cap the number of retained extreme vertices to at most $\lfloor f_{\text{ext}} \cdot N_{\text{target}} \rfloor$ (we use $f_{\text{ext}} = 0.35$), keeping the vertices that are selected most frequently as support maximizers across the sampled directions. These vertices correspond to supporting points of the sampled boundary and capture corner-like features of the MMB.

All retained extreme vertices are force-included in the final subset and are never removed during subsequent spacing-driven selection. Any remaining subset capacity is filled using farthest point sampling applied to the non-extreme vertices, with distances initialized from the already-retained extreme vertices. This ensures that fill vertices are well-separated from extremes, which improves boundary uniformity while also preserving geometrically critical features. The support directions are generated using a Sobol sequence on the unit sphere, though the preservation mechanism itself is independent of the specific direction sampling scheme.

To address this, we include an explicit extreme-vertex retention step that ensures selected subsets preserve these boundary-critical points. While the MMB is defined by its entire boundary, these specific vertices correspond to maximal color mismatches along the analyzed three-dimensional direction vectors. In practice, these extreme candidates are identified and retained prior to subset selection. This prevents them from being discarded purely on the basis of spacing considerations. This safeguard allows improved uniformity to be achieved without distorting the geometry of the sampled metamer mismatch body. Algorithm 1 summarizes the full pool-and-subset workflow with extreme vertex preservation.

Algorithm 1 Pool-and-subset selection with extreme vertex preservation

Require: sensors $\mathbf{S}$; metameric slice constraints; number of generation runs G; normals per run N_k; target subset size N_{target}; support directions $\mathscr{U} = \{\mathbf{u}_j\}_{j=1}^{N_{\text{supp}}}$ on S^2; FPS restarts R; cap fraction f_{ext}
Ensure: selected vertex subset $\mathscr{V}_{\text{sub}}$

```
V_pool ← ∅
for g = 1, ..., G do
    Sample N_k directions and compute boundary vertices V_g
    V_pool ← V_pool ∪ V_g
end for
Map pooled vertices to mech-2 sensor space: X_pool ← map_2(V_pool)  ▷ X_pool = {x_i}, each x_i ∈ R^3
                                    ▷ Extreme vertex identification (support-function anchors)
I_ext ← ∅
Initialize counts c_i ← 0 for all x_i ∈ X_pool
for j = 1, ..., N_supp do
    i* ← arg max_i ⟨x_i, u_j⟩
    I_ext ← I_ext ∪ {i*};  c_i* ← c_i* + 1
end for
N_ext,max ← max(1, ⌊f_ext N_target⌋)
if |I_ext| > N_ext,max then
    Filter I_ext to the N_ext,max indices with largest counts c_i
end if
X_ext ← {x_i : i ∈ I_ext}                          ▷ anchors in sensor space
X_rem ← X_pool \ X_ext
if |X_ext| ≥ N_target then
    Select I_sub ⊆ I_ext as the N_target indices with largest counts c_i
else
    for r = 1, ..., R do
        X_fill^(r) ← FPS(X_rem, N_target − |X_ext|; X_ext)  ▷ FPS initialized with anchors X_ext
        X_cand^(r) ← X_ext ∪ X_fill^(r)
        Compute CV_NN(X_cand^(r))
    end for
    r* ← arg min_r CV_NN(X_cand^(r))
    Identify pooled indices I_sub of X_cand^(r*) within X_pool
end if
Retrieve V_sub ← {v_i ∈ V_pool : i ∈ I_sub}
```

Throughout this paper, we distinguish between the number of sampled support directions and the number of boundary vertices recovered after half-space intersection. We denote by N_k the number of unit directions sampled during the spherical sampling procedure, either in the joint color-mechanism space or, after slicing, in the sensor-response space. The resulting number of unique boundary vertices is denoted by N_{vertices} and is the quantity analyzed in the diagnostics and results sections.

Results

Baseline spherical sampling

We first examine the boundary vertex sets produced by a single application of the baseline spherical sampling method described in the Methods Section and originally proposed by Mackiewicz et al. [6]. For a fixed number of sampled directions N_k, a typical run returns on the order of several hundred to several thousand boundary vertices. The exact number of vertices recovered depends on the specific realization of the randomly sampled directions. When mapped into the sensor response space of the second color mechanism, these vertices form a visually plausible approximation of the metamer mismatch body boundary (Figure 2).

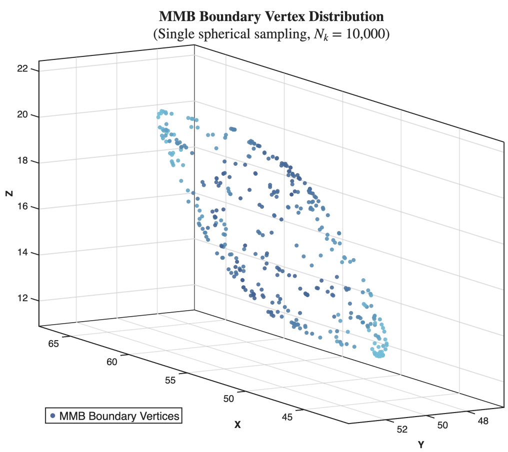

Figure 2. Boundary vertices of a metamer mismatch body obtained from a single spherical sampling run ($N_k = 10{,}000$) using the Mackiewicz half-space intersection method, shown in the sensor response space of the second color mechanism. Although the boundary is well resolved, the vertex distribution is highly non-uniform, with dense clustering and sparse coverage elsewhere (nearest-neighbor spacing coefficient of variation $\mathrm{CV} = 0.96$).

Despite this, the spatial distribution of vertices along the boundary is highly non-uniform. Certain regions of the boundary are densely populated by clusters of vertices, while other regions contain comparatively few points. This behavior is consistent with the geometry of the half-space construction of MMB boundary vertices. A single boundary vertex can optimize the supporting-point maximization for a nontrivial set of directions in k-space. As a result, many sampled directions map to the same geometric vertex.

This non-uniformity is reflected in the nearest-neighbor distance statistics. Even within a single run, the coefficient of variation (CV) of nearest-neighbor distances remains close to unity, indicating substantial variability in local point spacing. Increasing N_k typically increases N_{vertices}, but does not significantly reduce the variability in local point spacing. The observed non-uniformity is therefore not caused by insufficient sampling density, but is an inherent property of the spherical sampling procedure.

Variability across repeated baseline runs

We next assess how the baseline spherical sampling method behaves across repeated runs using independent random direction sets. Figure 3 overlays vertex sets from several representative runs and shows that, while the overall body is consistent, local boundary coverage is not, different runs populate different regions and exhibit different degrees of clustering.

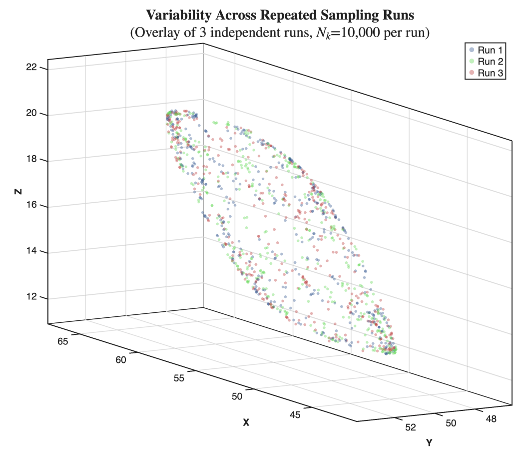

Figure 3. Variability across repeated spherical sampling runs. Overlaid boundary vertex sets in the sensor response space of the second color mechanism suggest a consistent overall hull shape but substantial run-to-run variation in local boundary coverage and vertex clustering.

This variability is reflected in the diagnostics summarized across 500 runs in Figure 4. Even with the sampling density and all other parameters held fixed (N_k=10,000 directions per run), N_{vertices} varies substantially from run to run (Figure 4a). This variability is induced by the random draw of k-space directions, which governs which supporting half-spaces are encountered and therefore which boundary vertices are recovered.

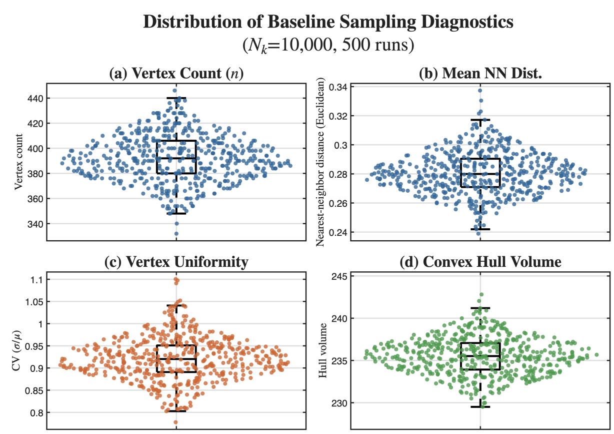

Figure 4. Diagnostics across 500 spherical sampling runs ($N_k = 10{,}000$): (a) N_{vertices}, (b) mean nearest-neighbor distance, (c) CV, and (d) hull volume. In these swarm plots points are jittered horizontally to show local density (x-axis is arbitrary). Black rectangles (box plots) denote the median and interquartile range.

Notably, higher vertex counts do not correspond to more uniform boundary coverage. The mean nearest-neighbor distance and its coefficient of variation remain similar across runs (Figure 4b,c), indicating that runs with larger N_{vertices} primarily contain additional redundant vertices rather than a more even distribution over the boundary.

Finally, the convex hull volume computed from each recovered vertex set is comparatively stable (Figure 4d). This quantity is used only as a relative consistency check, but its limited variation supports the interpretation suggested by Figure 3, repeated baseline runs recover broadly similar boundary extents even though the specific sampled vertices and local coverage differ.

Figures 3 and 4 show that repeating the baseline method, or increasing the number of sampled directions, does not address the non-uniformity of the recovered vertex distributions. This suggests a post-processing strategy that explicitly targets spacing while preserving the geometric extent of the sampled metamer mismatch body would be useful.

Effect of sampling density on boundary uniformity

Table 1 summarizes the effect of increasing the number of sampled directions N_k on the baseline spherical sampling results. As N_k increases, the number of boundary vertices N_{vertices} grows much more slowly, with large increases in sampled directions yielding only modest increases in unique vertices. The mean nearest-neighbor distance decreases monotonically with N_k, indicating denser sampling of the boundary. However, the coefficient of variation of nearest-neighbor distances remains high and shows no systematic reduction with increasing N_k. This indicates that additional sampled directions primarily generate redundant vertices in already well-sampled regions rather than improving uniformity of boundary coverage. These results demonstrate that increasing N_k alone does not resolve the intrinsic non-uniformity of vertex distributions produced by baseline spherical sampling.

Table 1. Baseline spherical sampling behavior as a function of the number of sampled directions N_k. Reported values are mean ± standard deviation across 50 independent runs.

N_k	N_{vertices}	$\bar{d}_{\text{NN}}$	CV
2000	218.76 ± 14.35	0.41 ± 0.03	0.88 ± 0.06
5000	300.84 ± 19.12	0.33 ± 0.02	0.90 ± 0.06
10000	394.60 ± 17.19	0.28 ± 0.02	0.91 ± 0.05
50000	696.92 ± 26.65	0.20 ± 0.01	0.98 ± 0.04

Table 2. Pool + subset results (Algorithm 1), reported using the same diagnostics as Table 1. Unlike the baseline method, the pool+subset procedure targets a fixed subset size N_{target} for each N_k, so N_{vertices} is fixed by construction (std=0).

N_k	N_{vertices}	$\bar{d}_{\text{NN}}$	CV
2000	219 ± 0	0.89 ± 0.01	0.44 ± 0.03
5000	301 ± 0	0.73 ± 0.01	0.45 ± 0.03
10000	395 ± 0	0.62 ± 0.01	0.46 ± 0.02
50000	697 ± 0	0.45 ± 0.00	0.47 ± 0.01

By contrast, as shown in Table 2, the pool+subset approach yields larger mean nearest-neighbor distances and a lower coefficient of variation at comparable N_{vertices}, indicating improved uniformity of boundary vertex spacing.

Improved Boundary Coverage via Pool-and-Subset Selection

For the comparison in Figure 5, the number of sampled directions N_k was adjusted independently for the baseline and pooled workflows in order to match the final number of recovered boundary vertices N_{vertices}.

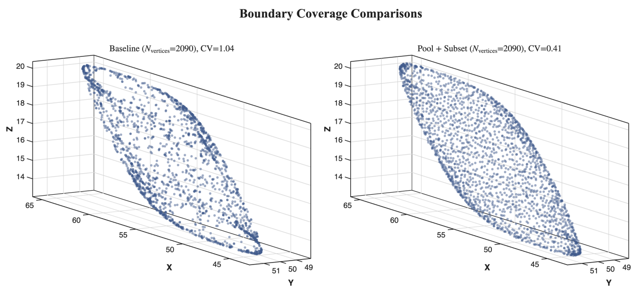

Figure 5. Comparison of boundary vertex distributions with matched final vertex counts (N_{vertices}). Left: boundary vertices from a single baseline spherical sampling run. Right: boundary vertices obtained by pooling multiple spherical sampling runs and selecting a subset using support-direction extreme retention followed by farthest point sampling. For this comparison, the number of sampled directions N_k was chosen independently for the baseline and pooled workflows so that both panels contain the same final number of unique boundary vertices ($N_{\text{vertices}} = 2090$). Both panels are shown in the sensor response space of the second color mechanism for the same fixed metameric slice.

Discussion

This work examined the behavior of baseline spherical sampling methods for constructing metamer mismatch body (MMB) boundaries and evaluated a simple post-processing strategy for improving boundary vertex uniformity. Our results show that, while spherical sampling efficiently recovers a geometrically plausible approximation of the MMB boundary, the resulting vertex distributions are highly non-uniform. This non-uniformity persists across repeated runs and is not resolved by increasing the number of sampled directions in k-space.

Variability and reproducibility considerations

Repeated baseline runs using identical problem parameters but independent random direction samples produce substantially different vertex sets, even though the overall hull shape remains stable. The observed variability in N_{vertices} and local boundary coverage underscores that a single spherical sampling run cannot be assumed to provide a representative or repeatable discretization of the MMB boundary.

This variability is not accompanied by large changes in convex hull volume, suggesting that global geometric extent is relatively robust to sampling variability. However, applications that rely on local boundary properties, such as identifying regions of maximal metamer mismatch or computing distance-based metrics, may be sensitive to the uneven and run-dependent vertex distributions produced by baseline sampling. These findings highlight the need for explicit strategies to improve vertex placement consistency when MMBs are used for comparative or quantitative analysis.

Implications for increasing sampling density

The results presented here demonstrate that increasing N_k alone is not an effective strategy for improving boundary uniformity. Although larger N_k values increase computational cost and modestly increase N_{vertices}, the additional vertices do not reduce clustering or improve spacing regularity. From a practical standpoint, this implies diminishing returns for brute-force sampling approaches and suggests that computational effort is better spent on post-processing strategies rather than further increasing directional sampling density.

This observation also provides a useful framework for prior empirical findings reporting large numbers of duplicate or near-duplicate vertices in spherical sampling–based MMB constructions. The redundancy observed in such methods is not an implementation artifact, but rather an expected outcome of the underlying geometry where sampled directions repeatedly map to the same vertex. This is defined by the associated normal cone.

Effectiveness of pool-and-subset selection

The pool-and-subset strategy explored in this work addresses the limitations of baseline sampling by decoupling boundary discovery from boundary discretization. Pooling vertices from multiple independent spherical sampling runs increases the diversity of candidate boundary points, while spacing-driven subset selection enforces more uniform coverage of the boundary surface.

Our results show that, for matched final vertex counts, the pool-and-subset approach yields substantially more uniform boundary coverage than a single baseline run. However, naive application of spacing-based selection methods such as farthest point sampling can discard boundary-defining extreme vertices, potentially altering the geometric extent of the recovered MMB. Incorporating an explicit extreme-vertex retention step mitigates this risk and preserves critical supporting points while still improving overall uniformity.

These findings suggest that effective MMB discretization requires balancing two competing objectives: uniform boundary coverage and faithful preservation of extreme boundary geometry. The proposed workflow provides a simple and flexible mechanism for achieving this balance.

Limitations and future directions

Several limitations of this study should be noted. First, all analyses were restricted to three-dimensional MMB cross-sections defined by fixing the response under the first color mechanism. While this reflects common practical use cases, higher-dimensional properties of the joint OCS were not examined. Second, nearest-neighbor statistics provide a convenient and interpretable measure of boundary uniformity, but other metrics, such as geodesic distances or curvature-aware measures, may offer additional insight.

Future work could explore adaptive or geometry-aware sampling strategies that directly target uniform boundary coverage during vertex generation, rather than relying on post hoc subset selection. Additionally, extending the analysis to other color mechanisms, illuminant pairs, or observer models would help assess the generality of the observed behaviors.

Conclusions

We evaluated spherical sampling methods for constructing metamer mismatch body (MMB) boundaries and showed that the resulting vertex distributions are inherently non-uniform. Increasing the number of sampled directions in k-space increases the number of recovered boundary vertices only modestly and does not improve boundary uniformity, as measured by nearest-neighbor statistics.

We showed that pooling vertices from multiple independent sampling runs and applying spacing-driven subset selection yields more uniform boundary representations for a fixed vertex budget. When combined with explicit retention of boundary-critical extreme vertices, this approach improves coverage while preserving the boundry structure of the MMB.

These results clarify the distinction between k-space sampling density and boundary resolution in sensor space and provide practical guidance for constructing stable MMB discretizations. Where boundary uniformity and repeatability are important, increasing the number of sampled directions alone is insufficient, and post-processing strategies such as pooling and subset selection are required. The source code for the proposed method is available at github.com/sfu-cs-vision-lab/uniform-spherical-sampling

References

[1] P. Centore. An open-source algorithm for metamer mismatch bodies. *Technical report*, 2017.

[2] Y. Eldar, M. Lindenbaum, M. Porat, and Y.Y. Zeevi. The farthest point strategy for progressive image sampling. *IEEE Transactions on Image Processing*, 6(9):1305–1315, 1997.

[3] Tony Esposito, Michael P. Royer, and Kevin W. Houser. Measures of illuminant-induced metameric mismatch: theory, comparative analysis, and implications for application. *Opt. Express*, 30(9):14686–14708, Apr 2022.

[4] Alexander Forsythe and Brian Funt. Metamer Mismatch Bodies: Foundations, Methods, and Applications in Color Science. *Journal of Physics: Conference Series*, 3128(1):012011, October 2025.

[5] Alexander D. Logvinenko, Brian Funt, and Christoph Godau. Metamer mismatching. *IEEE Transactions on Image Processing*, 23(1):34–43, 2014.

[6] Michal Mackiewicz, Hans Jakob Rivertz, and Graham D. Finlayson. Spherical sampling methods for the calculation of metamer mismatch volumes. *Journal of The Optical Society of America A-optics Image Science and Vision*, 36(1):96–104, January 2019.

[7] Tarek Stiebel and Dorit Merhof. A robust algorithm for computing boundary points of the metamer mismatch body. In *Proceedings of the IS&T 26th Color and Imaging Conference*, pages 134–139, 2018.

[8] Kuan Xu, Long Ma, and Peng Li. Metamer mismatch volume calculation method based on high-dimensional spherical sampling. In Xiang Zhang, editor, *Intelligent Photonics (Meta) Technology Symposium (IPTS2023)*, page 7, Wuhan, China, August 2023. SPIE.

Author Biography

Alex Forsythe is pursuing a Ph.D. at Simon Fraser University and is Senior Director of Science and Technology at the Academy of Motion Picture Arts and Sciences. He holds B.S. and M.S. degrees from Rochester Institute of Technology.

Brian Funt is Professor Emeritus, School of Computing Science, Simon Fraser University. His research focuses on computational models of colour perception.