

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039

RANDOM FEATURE MEAN-SHIFT

Anonymous authors

Paper under double-blind review

ABSTRACT

Locating the modes of a probability density function is a fundamental problem in many areas of machine learning. However, classical mode-seeking algorithms such as mean-shift and its variants exhibit quadratic complexity with respect to the number of data points due to exhaustive pairwise kernel computation - a well-known bottleneck that severely restricts the applicability. In this paper, we propose **Random Feature mean-shift (RFMS)**, a novel linear complexity mode-seeking algorithm. We give a sampling-based estimator using random feature kernel approximation and zeroth-order gradient method that allows us to provably achieve linear runtime per iteration, with comprehensive theoretical guarantees for mode estimation and convergence behavior. Empirical evaluations on clustering and pixel-level image segmentation tasks show RFMS is up to 12x faster when compared with other mean-shift variants, offering substantial efficiency gains while producing near-optimal results. **Overall, RFMS offers a practical and principled framework for scalable mode-seeking beyond kernel-value approximation, with explicit guarantees on the induced mode landscape and optimization dynamics.**

1 INTRODUCTION

Mode-seeking refers to the identification of maxima (“modes”) of a probability density. It provides a principled way to summarize and organize complex multimodal distributed data. This technique is routinely used in a variety of unsupervised learning tasks, notably clustering [13; 32], image processing [15; 38], and object tracking [16; 74]. For decades, Mean-shift [20] and its many variants have been the de facto algorithm for mode-seeking, the algorithm iteratively shifts each data point towards regions of higher density by computing local means within a predefined window. As a non-parametric method, mean-shift requires minimal assumptions about data distributions, making it highly versatile and effective.

The idea of mean-shift is closely related to kernel density estimation (KDE), where the gradients of the KDE function are employed to guide the points toward local maxima, representing the modes of the underlying probability density function. Figure 1 demonstrates the simple geometric intuition behind mean-shift. One natural way of interpreting the mean-shift procedure is gradient ascent, in the sense that the point will eventually converge to the modes of KDE. Despite its effectiveness, conventional mean-shift methods face significant computational limitations, primarily stemming from their quadratic complexity related to exhaustive pairwise kernel function computation required for KDE. This bottleneck restricts the applicability of mean-shift to large-scale datasets and computationally intensive applications. Existing works on efficient mean-shift algorithms have explored solutions such as discretizing the feature space via hashing [38], employing specialized data structures for distance-based queries [72], or accelerating the algorithm in hardware [35]. Despite those efforts, none have succeeded in addressing the complexity bottleneck of mean-shift asymptotically without discretizing data representation in some fashion. Noticing the close connection between kernel density

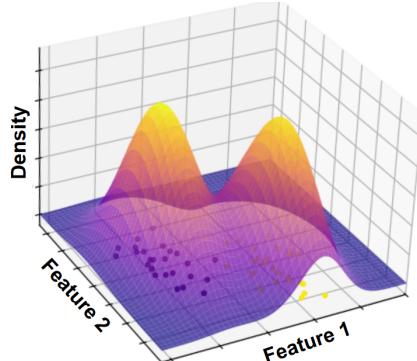


Figure 1: Kernel density estimation over synthetic 2D data (with 2 clusters) using Gaussian kernel. Points converge to the modes of empirical distribution in mean-shift.

054 estimation using shift-invariant kernels [1] and random Fourier feature kernel approximation [65; 21].
 055 In this paper, we propose a novel sampling-based estimator for the mean-shift process. Based on
 056 random feature kernel approximation and zeroth-order optimization, RFMS is useful in many areas
 057 of machine learning. Overall, our contributions are as follows:

- 058 • **Algorithm:** We propose Random Feature mean-shift (RFMS), an asymptotic complexity
 059 improvement of the standard mean-shift algorithm (from $\mathcal{O}(n^2)$ to $\mathcal{O}(n)$ with respect to the
 060 number of data points), enabling efficient density-based analysis and mode-seeking.
- 061 • **Theory:** Strong and comprehensive theoretical guarantees for RFMS are provided. We first
 062 establish high-probability concentration bounds for random feature KDE approximation
 063 (Theorem 1, 2). Furthermore, error bounds for mode estimation and tracking are also
 064 provided (Theorem 5, 6), ensuring the results obtained by RFMS are close to the actual
 065 modes of kernel density estimation.
- 066 • **Evaluation:** Experiments against other mean-shift variants on clustering and pixel-level
 067 image segmentation demonstrate RFMS’s advantage on reducing computational cost (up to
 068 12x speedup on clustering and 3x on segmentation) while providing nearly optimal results.

069 Our work saliently and naturally brings together density estimation, random feature approximation,
 070 and mode-seeking into a single end-to-end framework. To the best of our knowledge, this is the
 071 first work to provide high-probability random feature-induced perturbation guarantees for KDE, and
 072 propagate these approximation effects into principled mode-stability and convergence guarantees for
 073 an efficient mode-seeking algorithm.

074 2 RELATED WORKS

075 **Mean-shift:** The mean-shift algorithm has been extensively studied as a powerful non-parametric
 076 technique for locating the modes of a density function. The algorithmic basis of mean-shift was
 077 first introduced by Fukunaga and Hostetler [20], who proposed a method to estimate the gradient
 078 of a multivariate probability density function and perform mode-seeking iteratively. This initial
 079 formulation was later improved and popularized through the lens of clustering [13] and computer
 080 vision tasks such as edge-preserving smoothing and image segmentation [15]. The mean-shift
 081 algorithm is computationally expensive. Specifically, one iteration of mean-shift on a single point
 082 requires computing the kernel values between all other points, resulting in a computational complexity
 083 of $\mathcal{O}(n^2)$ per iteration, where n is the number of data points. Modern implementations of mean-shift
 084 algorithms typically utilize a flat kernel and efficient data structures, such as ball-tree or kd-tree, to
 085 organize distance searches over data points [62; 72]. This, however, does not improve the asymptotic
 086 complexity. More recently, mean-shift++ [39] by Jang et al. proposes to hash data points into discrete
 087 hypercubes at the beginning of every iteration, allowing mean-shift to be performed by averaging the
 088 neighboring hypercubes instead of all other points. However, as the number of hypercubes scales
 089 exponentially with data dimensionality, MS++ does not scale well into higher dimensions. **Despite**
 090 **mean-shift being the de facto and the most used mode-seeking method in practice, we note that**
 091 **other methods for finding mode of a KDE exist, notably the works by Lee et al. [54] and Luo et**
 092 **al. [54].** However, they primarily aim to find an approximately global maximizer of the KDE with
 093 dimensionality reduction plus additional heavy algorithmic machinery (e.g., solving polynomial
 094 systems). In contrast, RFMS is designed to run a mean-shift-style mode-seeking process for many
 095 modes, which is necessary in clustering and segmentation. Furthermore, RFMS is designed to
 096 use only simple and fully vectorized computation (essentially just dot products and some simple
 097 element-wise operations which emphasizes scalability and efficiency).

098 **Random Features:** Emerging from the literature on kernel methods, Random Features (RFF) [65]
 099 is a strategy to scale up kernel-based learning algorithms. The central idea behind random feature
 100 methods is to approximate a shift-invariant kernel $k(x, x') = k(x - x')$ (i.e. Gaussian, Laplacian
 101 and Cauchy kernels) using an explicit randomized feature map $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^D$ or \mathbb{C}^D such that
 102 $k(x, x') = \mathbb{E}[\langle \phi(x), \phi(x') \rangle]$, the expectation is taken with respect to the construction of ϕ . This
 103 is theoretically grounded in Bochner’s theorem, which states that any continuous, shift-invariant,
 104 positive definite kernel k on \mathbb{R}^d is the Fourier transform of a probability measure. Specifically,
 105 let $p(\omega)$ be the spectral density of the kernel, since $k(x - x') = \int_{\mathbb{R}^d} p(\omega) e^{j\omega \cdot (x - x')} d\omega$, one can
 106 approximate the integral via Monte Carlo method [57] by sampling on frequencies ω . The error of
 107 Random Feature approximation can be controlled by increasing D , the number of features. One

108 central challenge of Random Feature is to achieve a high-quality approximation with D being as
 109 small as possible. Several variants and extensions have been developed, such as orthogonal random
 110 features [78] and Quasi Monte-Carlo [36] to reduce the size of D . **Prior random feature works largely**
 111 **focus on approximating kernel values or improving the uniform-error bound behavior. In contrast, our**
 112 **work here provides an end-to-end framework for RFF-based mode-seeking, and we give theoretical**
 113 **insights on how random feature approximation induced errors can propagate into KDE approximation**
 114 **and mode stability.**

115 **Zeroth-Order Optimization:** Also known as derivative-free optimization or black-box optimization-
 116 is a class of optimization methods that do not require first-order information (i.e., without
 117 computing gradients). Such techniques are often used for problems where gradient information is
 118 unavailable, expensive, or unreliable [51]. The central challenge in Zeroth-Order Optimization is
 119 approximating the gradient using only function values. One of the most common approaches is the
 120 finite-difference gradient estimator [58; 17], which estimates the gradient $\nabla f(x)$ by querying the
 121 function at perturbed points around x via $\hat{\nabla} f(x) = \frac{f(x+\mu u) - f(x)}{\mu} u$, where u is a random direction
 122 and μ is a small scalar called smoothing parameter. Particularly, if u is sampled from multivariate
 123 Gaussian or uniformly from a sphere of radius the same as the dimension of the problem [51], then
 124 $\mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d) \text{ or } \mathcal{U}(\mathcal{S}(0, d))} [\hat{\nabla} f(x)] = \nabla f_\mu(x)$, which provides an unbiased estimation to the gradient
 125 of f_μ , the smoothed version of f , useful for Stochastic Gradient methods [26; 52].

3 NOTATION AND PRELIMINARIES

130 Consider $\mathcal{D} = \{x_1, x_2, \dots, x_n\} \subset \mathbb{R}^d$ a dataset of n points in d -dimensional euclidean space. One
 131 way to model the probability density function of the underlying data generative process is to construct
 132 a kernel density estimation (KDE) [61] over observed data points: $\hat{f}(x) = \frac{1}{nc} \sum_{i=1}^n k(x, x_i) \forall x_i \in \mathcal{D}$. Here $k : \mathbb{R}^d \rightarrow \mathbb{R}_+$ is a radially symmetric kernel function, and c is some kernel dependent
 133 constant that ensures $\int_{\mathbb{R}^d} \hat{f}(x) dx = 1$ so it is a valid probability distribution. KDE represents a
 134 smoothed version of the empirical distribution, with the kernel function distributing mass around
 135 each point x_i . Although any shift-invariant kernel can be used in RFMS, for consistency and ease of
 136 notation, in this paper we consider k to be the Gaussian kernel and denote k_h the Gaussian kernel
 137 of bandwidth h : $k_h(x_i, x_j) = \exp\left(-\frac{1}{2h^2} \|x_i - x_j\|^2\right) = k_h(\Delta)$ where $\Delta = \|x_i - x_j\|^2$, and
 138 \hat{f}_{k_h} is the KDE using k_h . As such $c = \int_{\mathbb{R}^d} k_h(\Delta) d\Delta = (2\pi h^2)^{d/2}$. One natural way of finding
 139 the modes of the KDE is to allow points to “climb” over the KDE via gradient ascent [20; 32]:
 140 $\nabla \hat{f}_{k_h}(x) = \frac{1}{h^2 cn} \sum_{i=1}^n k_h(x - x_i) \cdot (x_i - x)$. This expression shows that the gradient of the
 141 KDE points in the direction of a weighted average of the $(x_i - x) \forall x_i \in \mathcal{D}$. Notice that $\nabla \hat{f}_{k_h}$
 142 is proportional to $m(x) = \frac{\sum_{i=1}^n x_i k_h(x - x_i)}{\hat{f}_{k_h}(x)} - x$ by a multiplicative factor of $(h^2 cn \hat{f}_{k_h}(x))^{-1}$.
 143 Consider the gradient ascent $x^{(l+1)} = x^{(l)} + m(x^{(l)})$ which gives the modern fix-point iteration
 144 version of the mean-shift [15]: $x^{(l+1)} = \frac{\sum_{i=1}^n x_i k_h(x^{(l)} - x_i)}{\sum_{i=1}^n k_h(x^{(l)} - x_i)}$. In this way, since $x^{(l+1)} - x^{(l)} =$
 145 $m(x^{(l)}) = (h^2 cn \hat{f}_{k_h}(x^{(l)}))^{-1} \nabla \hat{f}_{k_h}(x^{(l)})$ we can also view the update rule as normalized gradient
 146 ascent (by the density at $x^{(l)}$) on the KDE. This allows us to define the mean-shift problem we aim to
 147 solve in this paper:

$$x_i^{(l+1)} = x_i^{(l)} + \eta \frac{\nabla \hat{f}_{k_h}(x_i^{(l)})}{\hat{f}_{k_h}(x_i^{(l)})} \quad \forall i \in \{1 \dots n\} \quad \text{with} \quad x_i^{(0)} = x_i \quad \forall x_i \in \mathcal{D} \quad (1)$$

153 Similar mean-shift formulations are also studied in previous papers [1; 32; 20]. The main overhead
 154 of this procedure comes from the density estimation via KDE (denominator) and the gradient com-
 155 putation (numerator), both of which takes $\mathcal{O}(n^2)$ since each iteration of update requires computing
 156 pairwise kernel value between $x_i^{(l)}$ and $x_i \in \mathcal{D}$, which is the main bottleneck. It is worth noting that
 157 although many variants of the mean-shift algorithm have been proposed in existing literature, such as
 158 Blurring mean-shift [15; 7], Robust mean-shift [5], they all suffer from the same quadratic complexity

162 limitation. In this work, we demonstrate that it is in fact possible to construct an efficiency estimator
 163 for $\frac{\nabla \hat{f}_{k_h}(x_i^{(l)})}{\hat{f}_{k_h}(x_i^{(l)})}$ that takes constant time per point to achieve $\mathcal{O}(n)$ complexity per iteration of update.
 164
 165

166 4 RANDOM FEATURE MEAN-SHIFT (RFMS)

168 Here, we provide algorithmic details of Random Feature mean-shift (RFMS), a novel linear complexity
 169 mode-seeking algorithm. We first establish a framework for efficient kernel density estimation
 170 using random feature method and then utilize it for zeroth-order gradient ascent for mode-seeking.
 171

172 **Random Feature Density Estimation:** As discussed in Section 3, one of the main bottleneck of
 173 mean-shift algorithm is the computation of local density as $\hat{f}_{k_h}(x)$ for any $x \in \mathbb{R}^d$ requires the
 174 pairwise kernel computation between x and every x_i in dataset \mathcal{D} . This becomes problematic if
 175 \mathcal{D} is large. However, recall that $\hat{f}_{k_h}(x) = \frac{1}{nc} \sum_{i=1}^n k_h(x, x_i)$, which is the sum of kernel values
 176 with a fixed data point. Now consider a Random Feature [65] transformation for k_h denoted as
 177 $\phi_{k_h} : \mathbb{R}^d \rightarrow \mathbb{C}^D$ with the property that $\langle \phi_{k_h}(x), \phi_{k_h}(x') \rangle \approx k_h(x, x') \forall x, x' \in \mathbb{R}^d$ for some
 178 sufficiently large D . We can then rewrite $\hat{f}_{k_h}(x)$:

$$179 \hat{f}_{k_h}(x) = \frac{1}{nc} \sum_{i=1}^n k_h(x, x_i) \approx \frac{1}{nc} \sum_{i=1}^n \langle \phi_{k_h}(x), \phi_{k_h}(x_i) \rangle = \frac{1}{nc} \left\langle \phi_{k_h}(x), \sum_{i=1}^n \phi_{k_h}(x_i) \right\rangle \quad (2)$$

182 The above derivation uses the fact that the complex inner product is linear in the second argument. We
 183 can then define $\Phi = \frac{1}{nc} \sum_{i=1}^n \phi_{k_h}(x_i)$, and density estimation can be done by $\hat{f}_{k_h}(x) \approx \langle \phi_{k_h}(x), \Phi \rangle$.
 184 As a reminder $c = (2\pi h^2)^{d/2}$ for Gaussian kernel of bandwidth h over \mathbb{R}^d and complex vector inner
 185 product $\langle a, b \rangle = (\bar{a} \cdot b) \forall a, b \in \mathbb{C}^D$. To construct ϕ_{k_h} , simply sample random features $\{\omega_1 \dots \omega_D\}$
 186 from $\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$ and we have:

$$188 \phi_{k_h}(x) = \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T x)} e^{j(\omega_2^T x)} e^{j(\omega_3^T x)} \dots e^{j(\omega_D^T x)} \right] \quad (3)$$

190 Here, j is the imaginary unit. Let $\hat{g}(x) = \frac{1}{nc} \langle \phi_{k_h}(x), \Phi \rangle$, it is an unbiased estimation of the KDE:
 191

$$192 \hat{f}_{k_h}(x) = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\hat{g}(x)] \quad \forall x \in \mathbb{R}^d \quad (\text{See Appendix E for the proof}). \quad (4)$$

193 We will later give a high-probability bound on the concentrations around this expected value. One
 194 important caveat is that although the expectation of $\hat{g}(x)$ is real, in an approximation setting where D
 195 is finite, the value of $\hat{g}(x)$ is not guaranteed to be real-valued. However, since KDE is a real-valued
 196 function, it implies that the imaginary part has zero mean therefore adds only variance, so in practice
 197 we only take the real part: $\text{Re}(\hat{g}(x)) = \text{Re}(\langle \phi_{k_h}(x), \Phi \rangle) = \frac{1}{nc} \sum_{i=1}^n \text{Re}(\langle \phi_{k_h}(x), \phi_{k_h}(x_i) \rangle)$. What
 198 this means is that when evaluating the KDE at some point in \mathbb{R}^d with respect to a dataset \mathcal{D} , instead
 199 of compute n kernel values, one can get an approximation in constant time via $\hat{g}(x)$ because the
 200 mapping $\phi_{k_h}(x)$ takes constant time per point and Φ only needs to be computed once over \mathcal{D} at the
 201 beginning and stored.

202 **Zeroth-Order Gradient Estimation:** Recall in Equation 1, in order to construct the mean-shift
 203 updates for a point $x \in \mathbb{R}^d$ using a dataset $\mathcal{D} \subset \mathbb{R}^d$, it is necessary to obtain the ascending-direction
 204 vector $\nabla \hat{f}_{k_h}(x)$, which is the average of $(x_i - x) \forall x_i \in \mathcal{D}$ weighted by $k_h(x, x_i)$. Previously, we have
 205 established that one can estimate the KDE value $f_{k_h}(x)$ by querying $\hat{g}(x) = \langle \phi_{k_h}(x), \Phi \rangle$. **Naviely**,
 206 **one could take the gradient of $\hat{g}(x)$ as a surrogate to the actual KDE gradient; however, that would**
 207 **require computing the Jacobian w.r.t every dimension of x , which will result in an complexity of**
 208 **$\mathcal{O}(Dd)$ per point per iteration. To further reduced the mean-shifit cost to $\mathcal{O}(D)$ per point per iteration,**
 209 **we apply zeroth-order gradient method instead of calculating the gradient analytically.** Therefore, the
 210 ability to query density estimation is powerful as it enables us to estimate $\nabla \hat{f}_{k_h}(x)$ via a zeroth-order
 211 gradient estimator, particularly the 2-point forward estimator taking the form:

$$212 \hat{\nabla} \hat{f}_{k_h}(x) = \frac{\text{Re}(\hat{g}(x + \mu u) - \hat{g}(x))}{\mu} u \quad (5)$$

213 Here we sample u , the random ascent direction, uniformly from the standard Gaussian $u \sim \mathcal{N}(0, \mathbf{I}_d)$.
 214 And μ is a small smoothing parameter set beforehand. This estimator works by probing the function

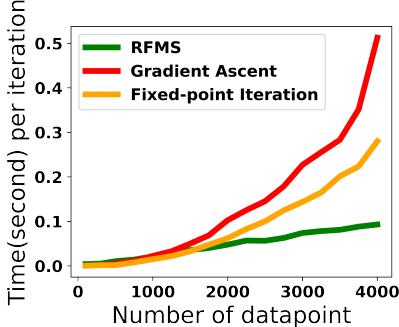


Figure 2: Time (in seconds) per iteration of mean-shift on randomly generated 2D data.

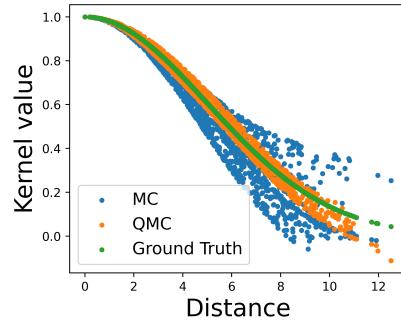


Figure 3: Kernel approximation quality using MC and QMC, $D = 30$.

in a randomly chosen direction and measuring the change in function output after taking a small step. This change is used to estimate the directional derivative. Geometrically, it captures how steeply the function rises in that sampled direction, giving a noisy ascent direction. Repeating this across iterations allows a point to follow the landscape of the function without ever computing actual gradients. Two levels of approximation are happening in Equation 5, the first one being that $\hat{g}(x)$ itself is an unbiased estimation of the KDE $\hat{f}_{k_h}(x)$. The second one is the gradient estimation of the KDE, namely the $\hat{\nabla} \hat{f}_{k_h}$, which is unbiased with respect to the gradient of a smoothed version of \hat{f}_{k_h} . With that, we can arrive at a new update rule for mean-shifting:

$$x_i^{(l+1)} = x_i^{(l)} + \eta \frac{\hat{\nabla} \hat{f}_{k_h}(x_i^{(l)})}{\text{Re}(\hat{g}(x_i^{(l)}))} \quad \forall i \in \{1 \dots n\} \quad \text{with} \quad x_i^{(0)} = x_i \quad \forall x_i \in \mathcal{D} \quad (6)$$

Here, η is the learning rate, since the core characteristic of zeroth order optimization is that it only requires evaluating function values (in the context of KDE, the \hat{f}_{k_h}), it is worth noting that other, potentially more advanced, zeroth order gradient estimator [50; 12] can be used as drop-in replacement in RFMS which can lead to further variance reduction in gradient estimation and faster convergence.

Complexity & Scalability: Essentially, RFMS estimates the mean-shift update via two density estimations; as such, by using $\text{Re}(\hat{g})$, the cost of RFMS per iteration becomes constant per data point and scales linearly in the size of \mathcal{D} . One other cost of RFMS comes from the number of random features used in ϕ_{k_h} , making the final time complexity $\mathcal{O}(nD)$ for RFMS in contrast to the $\mathcal{O}(n^2d)$ complexity of standard mean-shift algorithm (both in the gradient ascent form and fixed-point iteration form). In situations where the dataset is very small, since D is typically larger than d , standard mean-shift can potentially be faster than RFMS. However, in practical settings, RFMS is much more suited for scaling to larger tasks. In Fig. 2, we empirically demonstrate the scalability of RFMS on randomly generated 2D data, showing a minimal increase in cost (time per iteration) when compared to standard mean-shift algorithms. **RFMS inherits the standard RFF trade-off between efficiency and approximation quality as D increases, but this approximation is widely accepted because it replaces quadratic kernel computations with simple feature inner products and yields major scalability gains in large- n settings.**

5 ESTIMATION AND CONVERGENCE BOUNDS

$\text{Re}(\hat{g})$ serves as a surrogate of the actual KDE function \hat{f}_{k_h} . The quality of RFMS largely depends on how well it approximates KDE. In this section, we first establish an error bound on estimating \hat{f}_{k_h} with $\text{Re}(\hat{g})$. Then, we show conditions under which the modes of the \hat{f}_{k_h} and $\text{Re}(\hat{g})$ are close and bounded. Lastly, utilizing existing results regarding zeroth-order optimization, we show the convergence of RFMS to points near the modes of \hat{f}_{k_h} .

Error bound on estimating $\hat{f}_{k_h}(x)$ with $\hat{g}(x)$: The goal here is to give a bound on the error of estimating KDE value with \hat{g} . That is, we want to bound $|\hat{g}(x) - \hat{f}_{k_h}(x)|$. By applying Hoeffding’s inequality, we can derive the following bound:

270 **Theorem 1** For any point $x \in \mathbb{R}^d$, chose any $\delta \in (0, 1)$, with probability at least $1 - \delta$:

$$272 \quad \left| \hat{g}(x) - \hat{f}_{k_h}(x) \right| \leq \frac{4}{c} \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}} \quad (\text{See Appendix F for the proof.}) \quad (7)$$

274 Furthermore, since the $\hat{g}(x)$ concentrates around a real number, the complex part of $\hat{g}(x)$ contribute
275 only variance, so if KDE is approximated via $\text{Re}(\hat{g}(x))$, the bound can be further reduced to
276 $\frac{2}{c} \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}}$. For the remainder of this section, we will only consider the real part of $\hat{g}(x)$. Extending
277 classical results on random feature method [65], we then extend Theorem 1 to a uniform convergence
278 bound:

280 **Theorem 2** Let \mathcal{X} be a compact set over \mathbb{R}^d such that $\mathcal{D} \subset \mathcal{X}$. Denote $\text{diam}(\mathcal{X})$ the diameter of \mathcal{X} .
281 Then, for error tolerance ϵ , the following bound holds:

$$283 \quad \Pr \left(\sup_{x \in \mathcal{X}} \left| \text{Re}(\hat{g}(x)) - \hat{f}_{k_h}(x) \right| \geq \epsilon \right) \leq 2^8 \left(\frac{c\sqrt{d} \text{diam}(\mathcal{X})}{h\epsilon} \right)^2 \exp \left(- \frac{D\epsilon^2}{c^2 4(d+2)} \right) \quad (8)$$

284 *(See Appendix G for the proof.)*

287 The above yields a uniform additive bound for approximating KDE via RFF. The primary reason for
288 providing an additive error bound instead of a relative error bound is that the density can be arbitrarily
289 close to 0 in low-density regions, hence a relative error bound over the entire space is generally
290 ill-posed unless one restricts to regions with meaningful density values [10; 43]. If we pose constrain
291 to a subset $S \in \mathbb{R}^d$ such that $S = \{x : \hat{f}_{k_h}(x) \geq \tau\}$, our uniform additive error bound immediately
292 implies a relative bound: $\sup_{x \in S} \left| \left(\text{Re}(\hat{g}(x)) - \hat{f}_{k_h}(x) \right) / \hat{f}_{k_h}(x) \right| \geq \frac{\epsilon}{\tau}$.

294 **Mode Stability:** In RFMS, we use $\text{Re}(\hat{g})$ as an surrogate of the KDE \hat{f}_{k_h} . Therefore, for mode-
295 seeking purposes, we would like to show that the modes of \hat{f}_{k_h} and $\text{Re}(\hat{g})$ are close. To achieve this,
296 we first demonstrate the point-wise closeness of the gradient and the Hessian.

297 **Theorem 3** C is a universal constant, for any point $x \in \mathbb{R}^d$, chose any $\delta \in (0, 1)$, with probability
298 at least $1 - \delta$:

$$301 \quad \left\| \nabla \text{Re}(\hat{g}(x)) - \nabla \hat{f}_{k_h}(x) \right\| \leq \frac{1}{nc} \left(\frac{8Cen\sqrt{d} \ln(2/\delta)}{Dh} + \sqrt{\frac{8Cen^2 d \ln(2/\delta)}{Dh^2}} \right) \quad (9)$$

303 *(See Appendix H for the proof.)*

305 **Theorem 4** For any point $x \in \mathbb{R}^d$, chose any $\delta \in (0, 1)$, with probability at least $1 - \delta$:

$$307 \quad \left\| \nabla^2 \text{Re}(\hat{g}(x)) - \nabla^2 \hat{f}_{k_h}(x) \right\|_F \leq \frac{1}{nc} \left(\frac{8eCnd \ln(2/\delta)}{Dh^2} + \sqrt{\frac{8eCn^2 d \sqrt{d(d+2)} \ln(2/\delta)}{Dh^4}} \right) \quad (10)$$

310 *(See Appendix I for the proof.)*

311 Then, assume \hat{f}_{k_h} is a Morse function (a function with non-degenerate critical points). Denote $\text{Lip}(\cdot)$
312 the Lipschitz constant of a function, and $\lambda_{\min}(\cdot)$ the smallest eigenvalue of a square matrix. We can
313 then establish the conditions regarding the closeness between their critical points (modes):

315 **Theorem 5** Let x^* be any critical point of \hat{f}_{k_h} , define: $\alpha = \frac{\text{Lip}(\nabla^2 \text{Re}(\hat{g})) \epsilon_1}{(\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2)^2}$. If: (1). Choose
316 $\epsilon_2 \leq \lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*))$. (2). Choose ϵ_1 such that $\alpha \leq \frac{1}{2}$. (3). $\left\| \nabla \text{Re}(\hat{g}(x)) - \nabla \hat{f}_{k_h}(x) \right\| \leq \epsilon_1$
317 and $\left\| \nabla^2 \text{Re}(\hat{g}(x)) - \nabla^2 \hat{f}_{k_h}(x) \right\|_F \leq \epsilon_2$ both holds with probability at least $1 - \delta/2$. Then, with
318 probability at least $1 - \delta$, there is only one critical point \hat{x}^* of $\text{Re}(\hat{g})$ such that:

$$322 \quad \hat{x}^* \in B \left(x^*, \frac{\epsilon_1}{\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2} \right) \quad (\text{See Appendix J for the proof.}) \quad (11)$$

324 This result enables the quantification of how much each mode shifts between the KDE \hat{f}_{k_h} and
 325 the random feature approximated KDE $\text{Re}(\hat{g})$. Since we can make ϵ_1 and ϵ_2 arbitrarily small by
 326 increasing D (see theorem 3 and theorem 4), with properly chosen ϵ_1 and ϵ_2 , the mode change can be
 327 well controlled, hence making $\text{Re}(\hat{g})$ a good surrogate for mode-seeking. Our gradient/Hessian and
 328 mode-closeness guarantees (Theorems 3–5) imply that the true modes are preserved and can shift
 329 only by a small amount under the random-feature approximation. Any additional modes, if they occur,
 330 must arise from minor oscillations of the approximated density in low-density regions; such minor
 331 oscillations are typically negligible for mode-seeking and do not affect the algorithm’s behavior. To
 332 empirically substantiate these claims, Appendix B presents synthetic 1D and 2D experiments that
 333 visualize and compare the random-feature density approximation against the exact KDE surface. The
 334 results are consistent with the theory.

335 **Convergence of RFMS:** Recall the RFMS iteration in equation 6, the algorithm can be interpreted as
 336 running zeroth-order gradient ascent on $\text{Re}(\hat{g})$ with decaying step-size. This interpretation allows us
 337 to analyze the convergence of RFMS with existing results from zeroth-order optimization literature.

338 **Theorem 6** Suppose a point $x \in \mathbb{R}^d$ has a local mode \hat{x}^* of $\text{Re}(\hat{g})$ with Łojasiewicz exponent θ . Let
 339 $l^* \in \{0 \dots T\}$ be the iteration index such that $\mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d)} \|\nabla \text{Re}(\hat{g}(x^{(l)}))\|^2$ is the smallest. Then:

340
$$\mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d)} \left[\|\hat{x}^* - x^{(l^*)}\|^2 \right] \leq \mathcal{O} \left(\left(\frac{1}{\sqrt{T}} + \mu^2 d^2 \right)^{1/2\theta} \right) \text{(See Appendix K for the proof.)} \quad (12)$$

341 Combined with Theorem 5, the above result allows one to quantify how close the solution returned
 342 by RFMS is to an actual mode of the KDE. We believe this “kernel approximation + mode stability +
 343 zeroth-order optimization” synthesis is nontrivial and substantial, because it directly addresses what
 344 matters for mean-shift: not just approximating the kernel function, but preserving the mode structure
 345 that defines clusters.

350 6 IMPLEMENTATION DETAIL

351 The pseudocode for RFMS is presented in Alg. 1. It takes in a set of points in \mathbb{R}^d and outputs the
 352 shifted version of those points also in \mathbb{R}^d . We provide additional details on RFMS to enhance its
 353 efficiency and extend it to applications that require a blurring process.

356 Algorithm 1: Random Feature mean-shift (RFMS)

357 **Data:** Dataset \mathcal{D} , bandwidth h , smoothing parameter μ , Learning rate η , number of iteration T ,
 358 RFF dimension D

359 1 $\phi_{k_h} \leftarrow$ generate random feature mapping for gaussian kernel with band width h via QMC;
 360 2 $e_i \leftarrow \phi_{k_h}(x_i) \quad \forall x_i \in \mathcal{D}; \quad \text{/*Encode each data point into*/}$
 361 3 $\Phi \leftarrow \frac{1}{n_c} \sum_{i=1}^n e_i; \quad \text{/*Define } \Phi \text{ to be used in } \hat{g}(x)\text{*/}$
 362 4 $x_i^{(0)} \leftarrow x_i \quad \forall x_i \in \mathcal{D}; \quad \text{/*Initial position*/}$
 363 5 **for** $l = 0, 1, 2, 3 \dots T$ **do**
 364 6 **for** each $x_i^{(l)}$ **do**
 365 7 | $u \leftarrow \mathcal{N}(0, \mathbf{I}_d); \quad \text{/*Sample random ascent direction*/}$
 366 8 | $x_i^{(l+1)} \leftarrow x_i^{(l)} + \eta \left(\left(\hat{\nabla} \hat{f}_{k_h} \left(x_i^{(l)} \right) \right) / \text{Re} \left(\hat{g} \left(x_i^{(l)} \right) \right) \right); \quad \text{/*Update position*/}$
 367 9 **end**
 368 10 **if** Blurring **then**
 369 11 | $\Phi \leftarrow \frac{1}{n_c} \sum_{i=1}^n e_i^{(l+1)}; \quad \text{/*Update } \Phi \text{ to reflect new density function after shifting*/}$
 370 12 **end**
 371 13 **end**
 372 14 **return** $\{x_i^{(T)}\}; \quad \text{/*New points after mean-shift*/}$

373
 374
 375
 376
 377 **Construct Random Feature Mapping using Quasi Monte-Carlo:** We use the random feature
 378 method in equation 3 for kernel approximation. As discussed previously, one core challenge here is

378 to reduce the number of features used while still providing a good approximation; in other words,
 379 how to make D as small as possible while still providing a high-quality approximation. Since
 380 evaluating a shift-invariant kernel is essentially the same as evaluating an integral associated with
 381 the kernel(Bochner’s Theorem [66]), one prominent solution is to incorporate Quasi Monte-Carlo
 382 (QMC) techniques for numerical integration into the random feature framework [36; 76]. QMC uses
 383 low-discrepancy sequences (e.g., Sobol, Halton, or Faure sequences) to generate random features that
 384 cover the space more uniformly. To put it simply, instead of sampling frequencies $\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$,
 385 we can improve RFMS by sample ω from Halton sequence and apply inverse cumulative distribution
 386 function to move them into the correct distribution. The intuition behind QMC is that well-distributed
 387 deterministic sampling can outperform random sampling in integration and approximation tasks. We
 388 demonstrate the improvement of QMC In Fig. 3 where we generate random 2D points and compute
 389 pairwise kernel values with k_h, ϕ_{k_h} constructed with Monte-Carlo, and ϕ_{k_h} constructed use QMC.
 390 As shown, with the same D , QMC can produce much higher-quality approximations, especially when
 391 the actual kernel value is small. The use of QMC allows us to reduce D , hence further improving the
 392 computational efficiency of RFMS.

393 **Non-blurring vs. Blurring mean-shift:** In non-blurring setting, each point climbing a hill (mode)
 394 based on the fixed landscape: $x^{(l+1)} = \frac{\sum_{i=1}^n x_i k_h(x^{(l)} - x_i)}{\sum_{i=1}^n k_h(x^{(l)} - x_i)}$. The landscape is fixed in the sense that
 395 kernels are computed with unshifted points. However, in many application scenarios such as image
 396 smoothing [15], data consolidation [7], or structure-preserving denoising [34], blurring mean-shift:
 397 $x^{(l+1)} = \frac{\sum_{i=1}^n x_i k_h(x^{(l)} - x_i^{(l)})}{\sum_{i=1}^n k_h(x^{(l)} - x_i^{(l)})}$ is preferred for its faster convergence due to data contraction. This
 398 blurring process can be easily integrated into RFMS. Simply view the blurring as gradient ascent over
 399 a new KDE based on shifted points at every iteration. We can update Φ at the end of every iteration,
 400 so the $\text{Re}(\hat{g}(x))$ would produce an estimated KDE value over shifted points.

401 **Representing shift via element-wise multiplication:** Alg. 1 requires going back and forth between
 402 \mathbb{R}^d and \mathbb{C}^d . However, with the help of complex number properties, it is possible to run RFMS entirely
 403 on the encoded version of the data points. Consider ϕ_{k_h} the encoding function, we can then represent
 404 translation(shift) in \mathbb{R}^d via element-wise multiplication in \mathbb{C}^d based on the property that:

$$405 \phi_{k_h}(x + x') = \frac{D}{\sqrt{D}} \phi_{k_h}(x) \otimes \phi_{k_h}(x') \quad \forall x, x' \in \mathbb{R}^d \quad (\text{See Appendix L for the proof}). \quad (13)$$

406 Where \otimes denotes element-wise multiplication. This is the primary reason we chose to use the complex
 407 version of the random feature instead of the real-valued version, as the real-valued version is unable
 408 to achieve the same results due to the periodic nature of the cosine function. In this way, the original
 409 data can be discarded after the encoding, and subsequent operations can be performed exclusively on
 410 the encoded version of the data:

$$411 e_i^{(l+1)} = \frac{D}{\sqrt{D}} e_i^{(l)} \otimes \phi_{k_h} \left(\eta \frac{\hat{\nabla} \hat{f}_{k_h}(x^{(l)})}{\text{Re}(\hat{g}(x^{(l)}))} \right) \quad (14)$$

412 The capability of updating encodings in place is appealing as it further simplifies the algorithm.
 413 Although instead of shifted points, the algorithm will give $\{e_i^{(T)}\}$ with pairwise inner products
 414 approximating k_h over shifted points. Since there exists a one-to-one correspondence between kernel
 415 value and distance, this is sufficient for any subsequent kernel or distance-based algorithms.

420 7 EXPERIMENTS

421 We verify the effectiveness and applicability of RFMS, we first directly inspect the mode-seeking
 422 behavior of RFMS using randomly synthesized clusters of different variance and cluster shape. The
 423 data points at different iterations are visualized and shown in Figure 4. As demonstrated, RFMS
 424 indeed achieves the intended mode-seeking functionality. It is also worth noting that the observations
 425 here closely match the theoretical insights we provided in Section 5. With a relatively small $D = 200$,
 426 we observe the points converge to a point very close to the actual modes (Theorem 5). Since directly
 427 evaluating mode-seeking algorithms is difficult, we instead apply RFMS in two applications where
 428 mean-shift algorithms are often applied - (1). In section 7.1, the RFMS algorithm is evaluated
 429 against other mean-shift algorithms in the context of clustering. We report clustering quality and
 430 time consumption of different methods. (2). In section 7.2, we apply RFMS to pixel-level image
 431 segmentation, a practical area of interest in computer vision. We use QMC for both experiments.

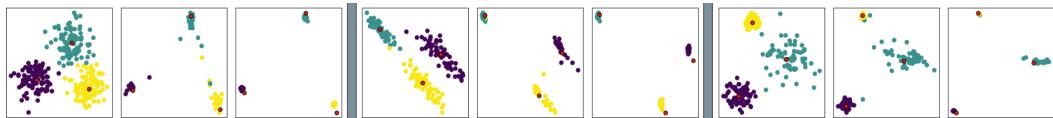


Figure 4: Trajectory Visualization of RFMS at $T = \{0, 30, 100\}$ and $D = 200$ on three separate examples with varying variance and cluster shape. Red points mark actual modes of clusters.

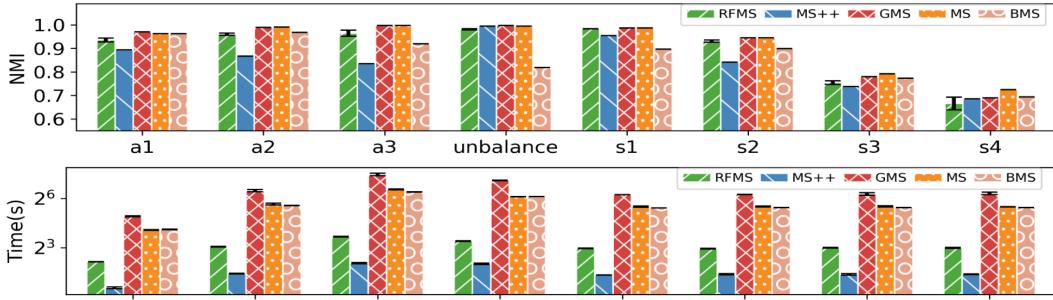


Figure 5: Comparison of different mean-shift algorithms in terms of time (measured in seconds in log scale) and normalized mutual information(NMI).

7.1 MEAN-SHIFT CLUSTERING

Mean-Shift algorithms are routinely used as clustering algorithms. It is particularly useful as it does not require the number of clusters to be predefined and can discover arbitrarily shaped clusters, given cluster forms a density peak. RMFS, as an efficient approximator of the classical mean-shift formulation, can also be used in this way. Here, we compare RFMS against other types of mean-shift algorithms: **(MS)**.Fixed point iteration [9; 47; 3; 5]. **(GMS)**. Gradient ascent over KDE [32; 20]. **(BMS)**. Blurring mean-shift [9; 28; 71]. **(MS++)**.Grid-based hashing [38]. Points that converge to the same mode are considered a cluster. After applying mean-shift, the connected component algorithm [8] is used to assign data points into different clusters. We conduct experiments on eight clustering benchmarking datasets from Fänti et al. [18]. We use $T = 100$ across all methods and $D = 300$ for RFMS. Each experiment was run 5 times, and we report the mean and standard deviation. The main clustering results are presented in Figure 5. Another huge advantage RFMS has is that its computational complexity scales well with d , the dimension of the data. The same is not true for hashing-based methods like MS++. To verify this, we conduct additional experiments on higher-dimensional real datasets [19; 40; 4]. The NMI and time results are shown in Table 1.

Results: We observe that RFMS can produce nearly optimal clustering results (on average 0.03 NMI drop-off compared with best NMI on each dataset) while being significantly more efficient than conventional mean-shift algorithms (MS, BMS, and GMS), with up to a 12x speed-up. We also observe that the efficiency benefit becomes more significant as the dataset grows larger, which is due to the asymptotic complexity improvement of RFMS, thereby verifying its scalability. Despite RFMS being slightly slower than MS++ (On average, 5.6s slower), RFMS produced better clustering quality in 6 out of 8 datasets tested in terms of normalized mutual information. In Table 1, on additional datasets, with increasing ambient dimension d , the efficiency of MS++ drops significantly. Particularly, when $d = 7$ (WirelessLocalization dataset), RFMS is over 70x faster than MS++. On the WallRobots dataset, where both d and n are large, RFMS show the overall best efficiency performance. The results above demonstrate the good scalability and mode-seeking quality of RFMS in comparison with previous mean-shift approaches.

7.2 PIXEL-LEVEL IMAGE SEGMENTATION

Mean-shift is also a popular vision algorithm commonly used for pixel-level segmentation [15; 38]. It is useful in generating initial region proposals or superpixels for deep semantic segmentation networks [55; 6; 45; 59; 77; 80]. Adapting a similar evaluation setup as MS++ [38], we conduct experiments on the Berkeley Segmentation Dataset Benchmark (BSDS500) [56], which contains 500 images with human-labelled segments.

Method	WallRobot		UserKnowledge		WirelessLocalization	
	$d = 4, n = 5456, D = 2000$	NMI Time(s)	$d = 5, n = 403, D = 500$	NMI Time(s)	$d = 7, n = 2000, D = 750$	NMI Time(s)
MS	33.8 \pm 0.0	32.9 \pm 2.8	34.3 \pm 0.0	0.2 \pm 0.0	74.5 \pm 0.0	4.5 \pm 0.6
MS++	33.5 \pm 0.0	27.7 \pm 1.2	32.6 \pm 0.0	8.6 \pm 0.5	70.8 \pm 0.0	499.8 \pm 129.5
RFMS	31.4 \pm 0.7	9.5 \pm 0.3	33.6 \pm 1.7	2.8 \pm 0.1	71.1 \pm 2.9	6.9 \pm 1.5

Table 1: Comparison of different mean-shift algorithms on additional higher-dimensional real datasets.

Each image is processed into a dataset containing 154401 three-dimensional points representing pixels in LAB color space. In addition to MS, BMS, and MS++ baselines, we also include QuickShift [69]; another popular segmentation algorithm based on mean-shift that jointly considers spatial and color features. We use the blurring version of RFMS. Due to the inefficiency of the conventional mean-shift algorithm, MS and BMS were run on images 1/36 of the original size, all other methods were run on full resolution. For RFMS, we set $D = 10$. All methods were run until convergence or a maximum of 100 iterations.

Results: We observe that RFMS, MS++, MS, and BMS all perform equally well on segmentation tasks.

Despite MS and BMS being run on lower-resolution sampled images, RFMS still achieves 3x speedup when compared with MS and BMS, and 2x speedup when compared to QuickShift. In contrast to the clustering experiments, we also observed that RFMS is slightly faster than MS++. This is due to the fact that the MS++ algorithm does not scale well to higher-dimensional input because the number of neighboring hypercubes increases exponentially with dimensionality. RFMS, however, is not affected by the dimension of the data.

Additional information: Full experimental details regarding baseline algorithms, important hyperparameters, and additional results can be found in Appendix N and O. We also provide a comprehensive ablation study regarding the sensitivity of RFMS hyperparameters (D, T, h, η, μ) and the effects of using MC and QMC sampling for RFMS in Appendix M. Furthermore, we also provided useful discussions on the significance of mean-shift algorithms, limitations & future works of RFMS in Appendix C and D.

8 CONCLUSION

Mean-shift is the de facto algorithm for mode-seeking - a fundamental procedure in many areas. In this paper, we propose Random Feature mean-shift (RFMS) for mode-seeking over kernel density estimation. Built on top of Random Feature method and zeroth-order optimization, RFMS is an asymptotic complexity improvement over the classical mean-shift algorithm. Theoretically, we show that the modes RFMS produces are close to the actual modes of the kernel density estimation, making RFMS an effective and efficient mode-seeking algorithm. **Rather than presenting standard RFF concentration bounds, we develop a complete pipeline tailored to mode seeking.** The key significance and novelty here is connecting random-feature approximation to the preservation of modes and mode-seeking dynamics, which, to our knowledge, is not addressed by prior RFF analyses that focus on kernel/value approximation. Empirically, RFMS matches the best clustering NMI within 0.03 while delivering up to 12x speedups. Similarly, on BSDS500, it attains 2–3x speedup compared to the baselines. This advancement broadens the practical applicability of mean-shift algorithms to domains previously limited by high computational demands.

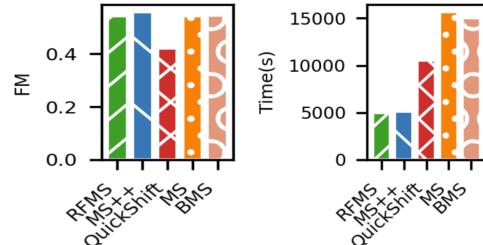


Figure 6: Comparison of different mean-shift image segmentation algorithms in terms of time and Fowlkes-Mallows Score(FM). Example segmented images can be found in Appendix O.

540 REFERENCES
541

542 [1] Ery Arias-Castro, David Mason, and Bruno Pelletier. On the estimation of the gradient lines of
543 a density and the consistency of the mean-shift algorithm. *The Journal of Machine Learning
544 Research*, 17(1):1487–1514, 2016.

545 [2] Siavash Arjomand Bigdeli, Matthias Zwicker, Paolo Favaro, and Meiguang Jin. Deep mean-shift
546 priors for image restoration. *Advances in neural information processing systems*, 30, 2017.

548 [3] Gaël Beck, Tarn Duong, Mustapha Lebbah, Hanane Azzag, and Christophe Cérin. A distributed
549 approximate nearest neighbors algorithm for efficient large scale mean shift clustering. *Journal
550 of Parallel and Distributed Computing*, 134:128–139, 2019.

551 [4] Rajen Bhatt. Wireless Indoor Localization. UCI Machine Learning Repository, 2017. DOI:
552 <https://doi.org/10.24432/C51880>.

554 [5] Claude Cariou, Steven Le Moan, and Kacem Chehdi. A novel mean-shift algorithm for data
555 clustering. *IEEE Access*, 10:14575–14585, 2022.

556 [6] Joao Carreira and Cristian Sminchisescu. Cpmc: Automatic object segmentation using con-
557 strained parametric min-cuts. *IEEE transactions on pattern analysis and machine intelligence*,
558 34(7):1312–1328, 2011.

560 [7] Miguel A Carreira-Perpiñán. Fast nonparametric clustering with gaussian blurring mean-shift.
561 In *Proceedings of the 23rd international conference on Machine learning*, pages 153–160, 2006.

562 [8] Miguel A Carreira-Perpiñán. A review of mean-shift algorithms for clustering. *arXiv preprint
563 arXiv:1503.00687*, 2015.

564 [9] Saptarshi Chakraborty, Debolina Paul, and Swagatam Das. Automated clustering of high-
565 dimensional data with a feature weighted mean shift algorithm. In *Proceedings of the AAAI
566 conference on artificial intelligence*, volume 35, pages 6930–6938, 2021.

568 [10] Moses Charikar and Paris Siminelakis. Hashing-based-estimators for kernel density in high
569 dimensions. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science
(FOCS)*, pages 1032–1043. IEEE, 2017.

571 [11] Frédéric Chazal, Leonidas J Guibas, Steve Y Oudot, and Primoz Skraba. Persistence-based
572 clustering in riemannian manifolds. *Journal of the ACM (JACM)*, 60(6):1–38, 2013.

574 [12] Xiangyi Chen, Sijia Liu, Kaidi Xu, Xingguo Li, Xue Lin, Mingyi Hong, and David Cox. Zo-
575 adamm: Zeroth-order adaptive momentum method for black-box optimization. *Advances in
576 neural information processing systems*, 32, 2019.

578 [13] Yizong Cheng. Mean shift, mode seeking, and clustering. *IEEE transactions on pattern analysis
579 and machine intelligence*, 17(8):790–799, 1995.

580 [14] Sua Choi, Dahyun Kang, and Minsu Cho. Contrastive mean-shift learning for generalized
581 category discovery. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
582 Pattern Recognition*, pages 23094–23104, 2024.

584 [15] Dorin Comaniciu and Peter Meer. Mean shift: A robust approach toward feature space analysis.
585 *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619, 2002.

586 [16] Dorin Comaniciu, Visvanathan Ramesh, and Peter Meer. Real-time tracking of non-rigid objects
587 using mean shift. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition.
588 CVPR 2000 (Cat. No. PR00662)*, volume 2, pages 142–149. IEEE, 2000.

590 [17] John C Duchi, Michael I Jordan, Martin J Wainwright, and Andre Wibisono. Optimal rates for
591 zero-order convex optimization: The power of two function evaluations. *IEEE Transactions on
592 Information Theory*, 61(5):2788–2806, 2015.

593 [18] Pasi Fränti and Sami Sieranoja. K-means properties on six clustering benchmark datasets.
594 *Applied intelligence*, 48:4743–4759, 2018.

594 [19] Veloso Marcus Freire, Ananda and Guilherme Barreto. Wall-Following Robot Navigation Data.
 595 UCI Machine Learning Repository, 2009. DOI: <https://doi.org/10.24432/C57C8W>.

596

597 [20] Keinosuke Fukunaga and Larry Hostetler. The estimation of the gradient of a density function,
 598 with applications in pattern recognition. *IEEE Transactions on information theory*, 21(1):32–40,
 599 1975.

600 [21] Joseph A Gallego, Juan F Osorio, and Fabio A Gonzalez. Fast kernel density estimation with
 601 density matrices and random fourier features. In *Ibero-American Conference on Artificial
 602 Intelligence*, pages 160–172. Springer, 2022.

603

604 [22] Ross W Gayler. Vector symbolic architectures answer jackendoff’s challenges for cognitive
 605 neuroscience. *arXiv preprint cs/0412059*, 2004.

606

607 [23] Lulu Ge and Keshab K Parhi. Classification using hyperdimensional computing: A review.
 608 *IEEE Circuits and Systems Magazine*, 20(2):30–47, 2020.

609

610 [24] Lulu Ge and Keshab K Parhi. Robust clustering using hyperdimensional computing. *IEEE
 611 Open Journal of Circuits and Systems*, 5:102–116, 2024.

612

613 [25] Georgescu, Shimshoni, and Meer. Mean shift based clustering in high dimensions: A texture
 614 classification example. In *Proceedings Ninth IEEE International Conference on Computer
 615 Vision*, pages 456–463. IEEE, 2003.

616

617 [26] Saeed Ghadimi and Guanghui Lan. Stochastic first-and zeroth-order methods for nonconvex
 618 stochastic programming. *SIAM journal on optimization*, 23(4):2341–2368, 2013.

619

620 [27] Leslie Greengard and John Strain. The fast gauss transform. *SIAM Journal on Scientific and
 621 Statistical Computing*, 12(1):79–94, 1991.

622

623 [28] Carlo Grillenzoni. Design of blurring mean-shift algorithms for data classification. *Journal of
 624 Classification*, 33(2):262–281, 2016.

625

626 [29] Victor Guillemin and Alan Pollack. *Differential topology*, volume 370. American Mathematical
 627 Soc., 2010.

628

629 [30] Charles R Harris, K Jarrod Millman, Stéfan J Van Der Walt, Ralf Gommers, Pauli Virtanen,
 630 David Cournapeau, Eric Wieser, Julian Taylor, Sebastian Berg, Nathaniel J Smith, et al. Array
 631 programming with numpy. *Nature*, 585(7825):357–362, 2020.

632

633 [31] Alejandro Hernández-Cano, Cheng Zhuo, Xunzhao Yin, and Mohsen Imani. Reghd: Robust
 634 and efficient regression in hyper-dimensional learning system. In *2021 58th ACM/IEEE Design
 635 Automation Conference (DAC)*, pages 7–12. IEEE, 2021.

636

637 [32] Alexander Hinneburg and Hans-Henning Gabriel. Denclue 2.0: Fast clustering based on kernel
 638 density estimation. In *International symposium on intelligent data analysis*, pages 70–80.
 639 Springer, 2007.

640

641 [33] Michael Hintermüller. Semismooth newton methods and applications. *Department of Mathe-
 642 matics, Humboldt-University of Berlin*, 2010.

643

644 [34] Guofei Hu, Qunsheng Peng, and A Robin Forrest. Mean shift denoising of point-sampled
 645 surfaces. *The Visual Computer*, 22:147–157, 2006.

646

647 [35] Miaoqing Huang, Liang Men, and Chenggang Lai. Accelerating mean shift segmentation
 648 algorithm on hybrid cpu/gpu platforms. *Modern Accelerator Technologies for Geographic
 649 Information Science*, pages 157–166, 2013.

650

651 [36] Zhen Huang, Jiajin Sun, and Yian Huang. Quasi-monte carlo features for kernel approximation.
 652 In *Forty-first International Conference on Machine Learning*, 2024.

653

654 [37] Zhenyang Hui, Shuanggen Jin, Yuanping Xia, Yunju Nie, Xiaowei Xie, and Na Li. A mean
 655 shift segmentation morphological filter for airborne lidar dtm extraction under forest canopy.
 656 *Optics & Laser Technology*, 136:106728, 2021.

648 [38] J Jang and H MeanShift+ Jiang. Extremely fast mode-seeking with applications to segmentation
649 and object tracking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and*
650 *Pattern Recognition (CVPR), Nashville, TN, USA*, pages 20–25, 2021.

651 [39] Jennifer Jang and Heinrich Jiang. Meanshift++: Extremely fast mode-seeking with applications
652 to segmentation and object tracking. In *Proceedings of the IEEE/CVF Conference on Computer*
653 *Vision and Pattern Recognition*, pages 4102–4113, 2021.

654 [40] Colak Ilhami Kahraman, Hamdi and Seref Sagiroglu. User Knowledge Modeling. UCI Machine
655 Learning Repository, 2009. DOI: <https://doi.org/10.24432/C5231X>.

656 [41] Pentti Kanerva. Hyperdimensional computing: An introduction to computing in distributed
657 representation with high-dimensional random vectors. *Cognitive computation*, 1:139–159, 2009.

658 [42] Leonid Vitalevich Kantorovich and Gleb Pavlovich Akilov. *Functional analysis*. Elsevier, 2014.

659 [43] Matti Karppa, Martin Aumüller, and Rasmus Pagh. Deann: Speeding up kernel-density
660 estimation using approximate nearest neighbor search. In *International Conference on Artificial*
661 *Intelligence and Statistics*, pages 3108–3137. PMLR, 2022.

662 [44] Soroush Abbasi Koohpayegani, Ajinkya Tejankar, and Hamed Pirsiavash. Mean shift for self-
663 supervised learning. In *Proceedings of the IEEE/CVF International Conference on Computer*
664 *Vision*, pages 10326–10335, 2021.

665 [45] Fengkai Lang, Jie Yang, Shiyong Yan, and Fachao Qin. Superpixel segmentation of polarimetric
666 synthetic aperture radar (sar) images based on generalized mean shift. *Remote Sensing*,
667 10(10):1592, 2018.

668 [46] Samuel Lanthaler and Nicholas H Nelsen. Error bounds for learning with vector-valued random
669 features. *Advances in Neural Information Processing Systems*, 36:71834–71861, 2023.

670 [47] Itshak Lapidot. Stochastic mean-shift clustering. *arXiv preprint arXiv:2312.15684*, 2023.

671 [48] Dongryeol Lee, Alexander G Gray, and Andrew W Moore. Dual-tree fast gauss transforms.
672 *arXiv preprint arXiv:1102.2878*, 2011.

673 [49] Jasper CH Lee, Jerry Li, Christopher Musco, Jeff M Phillips, and Wai Ming Tai. Finding the
674 mode of a kernel density estimate. *arXiv preprint arXiv:1912.07673*, 2019.

675 [50] Sijia Liu, Pin-Yu Chen, Xiangyi Chen, and Mingyi Hong. signsdg via zeroth-order oracle. In
676 *International conference on learning representations*, 2019.

677 [51] Sijia Liu, Pin-Yu Chen, Bhavya Kailkhura, Gaoyuan Zhang, Alfred O Hero III, and Pramod K
678 Varshney. A primer on zeroth-order optimization in signal processing and machine learning:
679 Principals, recent advances, and applications. *IEEE Signal Processing Magazine*, 37(5):43–54,
680 2020.

681 [52] Sijia Liu, Bhavya Kailkhura, Pin-Yu Chen, Paishun Ting, Shiyu Chang, and Lisa Amimi. Zeroth-
682 order stochastic variance reduction for nonconvex optimization. *Advances in neural information*
683 *processing systems*, 31, 2018.

684 [53] Yangxiao Lu, Yuqiao Chen, Nicholas Ruozzi, and Yu Xiang. Mean shift mask transformer for
685 unseen object instance segmentation. In *2024 IEEE International Conference on Robotics and*
686 *Automation (ICRA)*, pages 2760–2766. IEEE, 2024.

687 [54] Xinyu Luo, Christopher Musco, and Cas Widdershoven. Dimensionality reduction for general
688 kde mode finding. In *International Conference on Machine Learning*, pages 23067–23082.
689 PMLR, 2023.

690 [55] Tomasz Malisiewicz and Alexei A Efros. Improving spatial support for objects via multiple
691 segmentations. 2007.

692 [56] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images
693 and its application to evaluating segmentation algorithms and measuring ecological statistics.
694 In *Proc. 8th Int'l Conf. Computer Vision*, volume 2, pages 416–423, July 2001.

702 [57] Nicholas Metropolis. The monte carlo method. *Journal of the American statistical association*,
 703 44(247):335–341, 1949.

704

705 [58] Yurii Nesterov and Vladimir Spokoiny. Random gradient-free minimization of convex functions.
 706 *Foundations of Computational Mathematics*, 17(2):527–566, 2017.

707

708 [59] Tsz Ching Ng, Siu Kai Choy, Shu Yan Lam, and Kwok Wai Yu. Fuzzy superpixel-based image
 709 segmentation. *Pattern Recognition*, 134:109045, 2023.

710

711 [60] James M Ortega. The newton-kantorovich theorem. *The American Mathematical Monthly*,
 712 75(6):658–660, 1968.

713

714 [61] Emanuel Parzen. On estimation of a probability density function and mode. *The annals of
 715 mathematical statistics*, 33(3):1065–1076, 1962.

716

717 [62] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion,
 718 Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. Scikit-
 719 learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830,
 720 2011.

721

722 [63] Jeff M Phillips and Wai Ming Tai. Near-optimal coresets of kernel density estimates. *Discrete
 723 & Computational Geometry*, 63(4):867–887, 2020.

724

725 [64] Tony A Plate. *Holographic Reduced Representation: Distributed representation for cognitive
 726 structures*, volume 150. CSLI Publications Stanford, 2003.

727

728 [65] Ali Rahimi and Benjamin Recht. Random features for large-scale kernel machines. *Advances
 729 in neural information processing systems*, 20, 2007.

730

731 [66] Walter Rudin. *Fourier analysis on groups*. Courier Dover Publications, 2017.

732

733 [67] Sahand Salamat, Mohsen Imani, Behnam Khaleghi, and Tajana Rosing. F5-hd: Fast flexible
 734 fpga-based framework for refreshing hyperdimensional computing. In *Proceedings of the 2019
 735 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 53–62,
 736 2019.

737

738 [68] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne,
 739 Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image
 740 processing in python. *PeerJ*, 2:e453, 2014.

741

742 [69] Andrea Vedaldi and Stefano Soatto. Quick shift and kernel methods for mode seeking. In
 743 *Computer Vision–ECCV 2008: 10th European Conference on Computer Vision, Marseille,
 744 France, October 12–18, 2008, Proceedings, Part IV 10*, pages 705–718. Springer, 2008.

745

746 [70] Roman Vershynin. *High-dimensional probability: An introduction with applications in data
 747 science*, volume 47. Cambridge university press, 2018.

748

749 [71] Weiran Wang and Miguel A Carreira-Perpinán. Manifold blurring mean shift algorithms for
 750 manifold denoising. In *2010 IEEE Computer Society Conference on Computer Vision and
 751 Pattern Recognition*, pages 1759–1766. IEEE, 2010.

752

753 [72] Chunxia Xiao and Meng Liu. Efficient mean-shift clustering using gaussian kd-tree. In
 754 *Computer Graphics Forum*, volume 29, pages 2065–2073. Wiley Online Library, 2010.

755

756 [73] Yang, Duraiswami, and Gumerov. Improved fast gauss transform and efficient kernel density
 757 estimation. In *Proceedings ninth IEEE international conference on computer vision*, pages
 758 664–671. IEEE, 2003.

759

760 [74] Changjiang Yang, Ramani Duraiswami, and Larry Davis. Efficient mean-shift tracking via a
 761 new similarity measure. In *2005 IEEE Computer Society Conference on Computer Vision and
 762 Pattern Recognition (CVPR’05)*, volume 1, pages 176–183. IEEE, 2005.

763

764 [75] Jiawei Yang, Susanto Rahardja, and Pasi Fränti. Mean-shift outlier detection and filtering.
 765 *Pattern Recognition*, 115:107874, 2021.

756 [76] Jiyan Yang, Vikas Sindhwani, Haim Avron, and Michael Mahoney. Quasi-monte carlo feature
757 maps for shift-invariant kernels. In *International Conference on Machine Learning*, pages
758 485–493. PMLR, 2014.

759 [77] Xujiong Ye, Gareth Beddoe, and Greg Slabaugh. Automatic graph cut segmentation of lesions in
760 ct using mean shift superpixels. *International journal of biomedical imaging*, 2010(1):983963,
761 2010.

762 [78] Felix Xinnan X Yu, Ananda Theertha Suresh, Krzysztof M Choromanski, Daniel N Holtmann-
763 Rice, and Sanjiv Kumar. Orthogonal random features. *Advances in neural information processing
764 systems*, 29, 2016.

765 [79] Tianyang Yu, Bi Wu, Ke Chen, Gong Zhang, and Weiqiang Liu. Fully learnable hyperdi-
766 mensional computing framework with ultratiny accelerator for edge-side applications. *IEEE
767 Transactions on Computers*, 73(2):574–585, 2023.

768 [80] Yong-mei Zhou, Sheng-yi Jiang, and Mei-lin Yin. A region-based image segmentation method
769 with mean-shift clustering algorithm. In *2008 Fifth International Conference on Fuzzy Systems
770 and Knowledge Discovery*, volume 2, pages 366–370, 2008.

771

772

773

774

775

776

777

778

779

780

781

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

810
811 The appendix here provides additional details for the ICLR 2026 submission, titled “Random Feature
812 Mean-Sift”. The appendix is organized as follows:
813
814

- 815 • **A - List of Notation**
816
817
- 818 • **B - RFMS Density Estimation Visualization**
819
820
- 821 • **C - Discussion**
822
823
- 824 • **D - Limitation & Future Work**
825
826
- 827 • **E. Proof of Equation 4**
828
829
- 830 • **F - Proof of Theorem 1**
831
832
- 833 • **G - Proof of Theorem 2**
834
835
- 836 • **H - Proof of Theorem 3**
837
838
- 839 • **I - Proof of Theorem 4**
840
841
- 842 • **J - Proof of Theorem 5**
843
844
- 845 • **K - Proof of Theorem 6**
846
847
- 848 • **L - Proof of Equation 13**
849
850
- 851 • **M - Ablation Study**
852
853
- 854 • **N - Additional Details on Clustering Experiments**
855
856
- 857 • **O - Additional Details on Image Segmentation Experiments**
858
859
- 860 • **P - Reproducibility / Code Availability**
861
862
- 863 • **Q - LLM Usage**

864
865
866
867
868
869
870 A LIST OF NOTATION871
872
873
874
875
876
877
878
879
880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899 We hereby provide a list of notations used in this paper and accompanying proofs:

Symbol	Meaning
d, D	Data dimension, number of random feature
C	Universal constant
x	data point in \mathbb{R}^d
\mathcal{D}	Dataset
n	Number of data points in \mathcal{D}
k	Positive symmetric kernel function.
h	Gaussian kernel bandwidth
k_h	Gaussian kernel of bandwidth h
f	Data generating density function for \mathcal{D}
\hat{f}_{k_h}	Kernel density estimation of f using k_h
c	Normalizing constant making sure the integral of \hat{f}_{k_h} is 1
$x_i^{(l)}$	Point x_i after l iterations of mean-shift
ϕ_{k_h}	Random Feature transformation for kernel k_h
$\nabla \hat{f}_{k_h}, \hat{\nabla} \hat{f}_{k_h}$	Gradient and Estimated gradient of KED
Φ	Summation of random feature transformed points
\hat{g}	Random Feature estimation of KDE
ω	Frequencies for constructing Random Feature Mapping
T	Total number of iteration
η	Learning rate for gradient ascent
μ	Smoothing parameter in zeroth-order optimization
j	Imaginary unit
\mathcal{X}	Compact set over \mathbb{R}^d
$\text{Lip}(\cdot)$	Lipschitz constant of a function
$\lambda_{\min}(\cdot)$	Smallest eigenvalue of a square matrix
$B(x, r)$	Closed ball centered at x with radius r
θ	Łojasiewicz exponent
$\ \cdot\ $	L^2 norm
$\ \cdot\ _{\psi_1}$	Sub-exponential Orlicz norm
$\ \cdot\ _F$	Frobenius norm
$\ \cdot\ _{op}$	Operator norm

900
901
902
903
904
905
906
907
908
Table 2: List of notations.909
910
911
912
913
914
915
916
917 B RFMS DENSITY ESTIMATION VISUALIZATION918
919
920
921
922
923
924
925
926
927
928
929
930
931
932
933
934
935
936
937
938
939
940
941
942
943
944
945
946
947
948
949
950
951
952
953
954
955
956
957
958
959
960
961
962
963
964
965
966
967
968
969
970
971
972
973
974
975
976
977
978
979
980
981
982
983
984
985
986
987
988
989
990
991
992
993
994
995
996
997
998
999
This section provides a qualitative sanity check of the random-feature density approximation used by RFMS. While Theorems 3–5 establish that the approximation preserves mode locations up to a small perturbation (and that any spurious modes must be confined to low-density regions), visualizing the estimated density surfaces offers an intuitive confirmation of these claims.1000
1001
1002
1003
1004
1005
1006
1007
1008
1009
1010
1011
1012
1013
1014
1015
1016
1017
1018
1019
1020
1021
1022
1023
1024
1025
1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040
1041
1042
1043
1044
1045
1046
1047
1048
1049
1050
1051
1052
1053
1054
1055
1056
1057
1058
1059
1060
1061
1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073
1074
1075
1076
1077
1078
1079
1080
1081
1082
1083
1084
1085
1086
1087
1088
1089
1090
1091
1092
1093
1094
1095
1096
1097
1098
1099
1100
1101
1102
1103
1104
1105
1106
1107
1108
1109
1110
1111
1112
1113
1114
1115
1116
1117
1118
1119
1120
1121
1122
1123
1124
1125
1126
1127
1128
1129
1130
1131
1132
1133
1134
1135
1136
1137
1138
1139
1140
1141
1142
1143
1144
1145
1146
1147
1148
1149
1150
1151
1152
1153
1154
1155
1156
1157
1158
1159
1160
1161
1162
1163
1164
1165
1166
1167
1168
1169
1170
1171
1172
1173
1174
1175
1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187
1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
1242
1243
1244
1245
1246
1247
1248
1249
1250
1251
1252
1253
1254
1255
1256
1257
1258
1259
1260
1261
1262
1263
1264
1265
1266
1267
1268
1269
1270
1271
1272
1273
1274
1275
1276
1277
1278
1279
1280
1281
1282
1283
1284
1285
1286
1287
1288
1289
1290
1291
1292
1293
1294
1295
1296
1297
1298
1299
1300
1301
1302
1303
1304
1305
1306
1307
1308
1309
1310
1311
1312
1313
1314
1315
1316
1317
1318
1319
1320
1321
1322
1323
1324
1325
1326
1327
1328
1329
1330
1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398
1399
1400
1401
1402
1403
1404
1405
1406
1407
1408
1409
1410
1411
1412
1413
1414
1415
1416
1417
1418
1419
1420
1421
1422
1423
1424
1425
1426
1427
1428
1429
1430
1431
1432
1433
1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468
1469
1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488
1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506
1507
1508
1509
1510
1511
1512
1513
1514
1515
1516
1517
1518
1519
1520
1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532
1533
1534
1535
1536
1537
1538
1539
1540
1541
1542
1543
1544
1545
1546
1547
1548
1549
1550
1551
1552
1553
1554
1555
1556
1557
1558
1559
1560
1561
1562
1563
1564
1565
1566
1567
1568
1569
1570
1571
1572
1573
1574
1575
1576
1577
1578
1579
1580
1581
1582
1583
1584
1585
1586
1587
1588
1589
1590
1591
1592
1593
1594
1595
1596
1597
1598
1599
1600
1601
1602
1603
1604
1605
1606
1607
1608
1609
1610
1611
1612
1613
1614
1615
1616
1617
1618
1619
1620
1621
1622
1623
1624
1625
1626
1627
1628
1629
1630
1631
1632
1633
1634
1635
1636
1637
1638
1639
1640
1641
1642
1643
1644
1645
1646
1647
1648
1649
1650
1651
1652
1653
1654
1655
1656
1657
1658
1659
1660
1661
1662
1663
1664
1665
1666
1667
1668
1669
1670
1671
1672
1673
1674
1675
1676
1677
1678
1679
1680
1681
1682
1683
1684
1685
1686
1687
1688
1689
1690
1691
1692
1693
1694
1695
1696
1697
1698
1699
1700
1701
1702
1703
1704
1705
1706
1707
1708
1709
1710
1711
1712
1713
1714
1715
1716
1717
1718
1719
1720
1721
1722
1723
1724
1725
1726
1727
1728
1729
1730
1731
1732
1733
1734
1735
1736
1737
1738
1739
1740
1741
1742
1743
1744
1745
1746
1747
1748
1749
1750
1751
1752
1753
1754
1755
1756
1757
1758
1759
1760
1761
1762
1763
1764
1765
1766
1767
1768
1769
1770
1771
1772
1773
1774
1775
1776
1777
1778
1779
1780
1781
1782
1783
1784
1785
1786
1787
1788
1789
1790
1791
1792
1793
1794
1795
1796
1797
1798
1799
1800
1801
1802
1803
1804
1805
1806
1807
1808
1809
1810
1811
1812
1813
1814
1815
1816
1817
1818
1819
1820
1821
1822
1823
1824
1825
1826
1827
1828
1829
1830
1831
1832
1833
1834
1835
1836
1837
1838
1839
1840
1841
1842
1843
1844
1845
1846
1847
1848
1849
1850
1851
1852
1853
1854
1855
1856
1857
1858
1859
1860
1861
1862
1863
1864
1865
1866
1867
1868
1869
1870
1871
1872
1873
1874
1875
1876
1877
1878
1879
1880
1881
1882
1883
1884
1885
1886
1887
1888
1889
1890
1891
1892
1893
1894
1895
1896
1897
1898
1899
1900
1901
1902
1903
1904
1905
1906
1907
1908
1909
1910
1911
1912
1913
1914
1915
1916
1917
1918
1919
1920
1921
1922
1923
1924
1925
1926
1927
1928
1929
1930
1931
1932
1933
1934
1935
1936
1937
1938
1939
1940
1941
1942
1943
1944
1945
1946
1947
1948
1949
1950
1951
1952
1953
1954
1955
1956
1957
1958
1959
1960
1961
1962
1963
1964
1965
1966
1967
1968
1969
1970
1971
1972
1973
1974
1975
1976
1977
1978
1979
1980
1981
1982
1983
1984
1985
1986
1987
1988
1989
1990
1991
1992
1993
1994
1995
1996
1997
1998
1999
2000
2001
2002
2003
2004
2005
2006
2007
2008
2009
2010
2011
2012
2013
2014
2015
2016
2017
2018
2019
2020
2021
2022
2023
2024
2025
2026
2027
2028
2029
2030
2031
2032
2033
2034
2035
2036
2037
2038
2039
2040
2041
2042
2043
2044
2045
2046
2047
2048
2049
2050
2051
2052
2053
2054
2055
2056
2057
2058
2059
2060
2061
2062
2063
2064
2065
2066
2067
2068
2069
2070
2071
2072
2073
2074
2075
2076
2077
2078
2079
2080
2081
2082
2083
2084
2085
2086
2087
2088
2089
2090
2091
2092
2093
2094
2095
2096
2097
2098
2099
2100
2101
2102
2103
2104
2105
2106
2107
2108
2109
2110
2111
2112
2113
2114
2115
2116
2117
2118
2119
2120
2121
2122
2123
2124
2125
2126
2127
2128
2129
2130
2131
2132
2133
2134
2135
2136
2137
2138
2139
2140
2141
2142
2143
2144
2145
2146
2147
2148
2149
2150
2151
2152
2153
2154
2155
2156
2157
2158
2159
2160
2161
2162
2163
2164
2165
2166
2167
2168
2169
2170
2171
2172
2173
2174
2175
2176
2177
2178
2179
2180
2181
2182
2183
2184
2185
2186
2187
2188
2189
2190
2191
2192
2193
2194
2195
2196
2197
2198
2199
2200
2201
2202
2203
2204
2205
2206
2207
2208
2209
2210
2211
2212
2213
2214
2215
2216
2217
2218
2219
2220
2221
2222
2223
2224
2225
2226
2227
2228
2229
2230
2231
2232
2233
2234
2235
2236
2237
2238
2239
2240
2241
2242
2243
2244
2245
2246
2247
2248
2249
2250
2251
2252
2253
2254
2255
2256
2257
2258
2259
2260
2261
2262
2263
2264
2265
2266
2267
2268
2269
2270
2271
2272
2273
2274
2275
2276
2277
2278
2279
2280
2281
2282
2283
2284
2285
2286
2287
2288
2289
2290
2291
2292
2293
2294
2295
2296
2297
2298
2299
2300
2301
2302
2303
2304
2305
2306
2307
2308
2309
2310
2311
2312
2313
2314
2315
2316
2317
2318
2319
2320
2321
2322
2323
2324
2325
2326
2327
2328
2329
2330
2331
2332
2333
2334
2335
2336
2337
2338
2339
2340
2341
2342
2343
2344
2345
2346
2347
2348
2349
2350
2351
2352
2353
2354
2355
2356
2357
2358
2359
2360
2361
2362
2363
2364
2365
2366
2367
2368
2369
2370
2371
2372
2373
2374
2375
2376
2377
2378
2379
2380
2381
2382
2383
2384
2385
2386
2387
2388
2389
2390
2391
2392
2393
2394
2395
2396
2397
2398
2399
2400
2401
2402
2403
2404
2405
2406
2407
2408
2409
2410
2411
2412
2413
2414
2415
2416
2417
2418
2419
2420
2421
2422
2423
2424
2425
2426
2427
2428
2429
2430
2431
2432
2433
2434
2435
2436
2437
2438
2439
2440
2441
2442
2443
2444
2445
2446
2447
2448
2449
2450
2451
2452
2453
2454
2455
2456
2457
2458
2459
2460
2461
2462
2463
2464
2465
2466
2467
2468
2469
2470
2471
2472
2473
2474
2475
2476
2477
2478
2479
2480
2481
2482
2483
2484
2485
2486
2487
2488
2489
2490
2491
2492
2493
2494
2495
2496
2497
2498
2499
2500
2501
2502
2503
2504
2505
2506
2507
2508
2509
2510
2511
2512
2513
2514
2515
2516
2517
2518
2519
2520
2521
2522
2523
2524
2525
2526
2527
2528
2529
2530
2531
2532
2533
2534
2535
2536
2537
2538

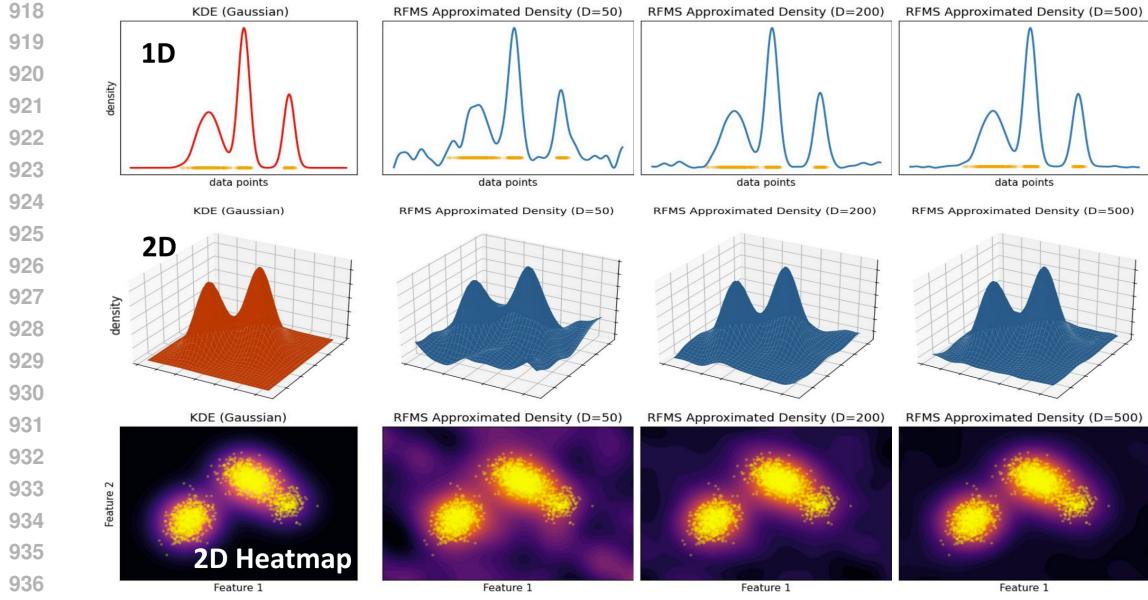


Figure 7: Density visualization on synthetic 1D and 2D points using Gaussian KDE and random feature density estimation. (Top) Actual and RFF approximated density function using 800 1D points with 3 clusters and different variances. (Middle) Actual and RFF approximated density function using 2000 2D points with 2 clusters and different variances. (Bottom) Heatmap of 2D examples.

Observations: Across both 1D and 2D examples, the RFMS density approximation accurately reproduces the dominant basins of attraction and preserves the number and locations of the high-density modes. Small local discrepancies mostly visible in the tails where the true density is near zero. In these low-density regions, the RFMS surface may exhibit mild ripples; consistent with our theory, such oscillations can introduce visually small, isolated extrema that do not correspond to meaningful modes and do not affect mode-seeking trajectories initialized in moderate-to-high density regions. As D increases, the RFMS surface becomes progressively smoother and converges visually to the KDE surface: peak locations stabilize, and tail oscillations diminish. This qualitative trend aligns with the approximation guarantees, where the gradient/Hessian error decreases with larger D , implying improved stability of critical points and their local geometry. Overall, these visualizations support the theoretical guarantees that RFMS preserves the relevant mode structure of the KDE and that any approximation-induced artifacts are limited to low-density regions where they have minimal impact on the practical mode-seeking behavior.

C DISCUSSION

C.1 COMPARISON WITH OTHER FAST KDE APPROACHES

A rich body of literature accelerates KDE through techniques such as coresets, locality sensitive hashing (LSH), and specialized data structures (e.g., space partitioning) [63; 10; 43; 27; 73; 39; 48]. Our focus here is different in that we aim for an end-to-end mode-seeking algorithm whose per-iteration cost is linear in n , without discretizing the domain or relying on search structures that can weaken in higher dimensions (e.g., tree, grid-style space partitions, or fast Gauss transform-based approaches [27; 73; 39; 48] whose cost grows rapidly with ambient data dimension). More importantly, those prior methods primarily target fast evaluation of KDE at a query point, whereas our analysis is designed to support mode-seeking correctness: we bound how the modes of the approximate KDE move relative to true KDE modes (Theorems 3–5) and prove convergence of a stochastic/ascent-style iteration to a neighborhood of those modes (Theorem 6). Propagating RFF-based KDE approximation error through to mode stability and tracking guarantees is a central

972 contribution of our work. We also note a compatibility issue with blurring mean shift: when
 973 points move each iteration, data-structure or coreset-based accelerations may require rebuilding or
 974 substantial updates, potentially eroding their efficiency. In contrast, RFMS only updates and sums the
 975 evolving feature encodings, which remains inexpensive and highly parallelizable. Overall, we view
 976 RFMS as a complementary point in the broader design space of fast KDE approximations, alongside
 977 coresets, LSH, and partitioning methods. Distinguished by RFMS’s simplicity in computation and a
 978 well-established RFF theoretical foundation.

980 C.2 COMPATIBLE KERNEL FUNCTIONS

982 We note that the class of shift-invariant kernels is extensive — including Gaussian, Laplacian, Cauchy,
 983 Matérn, and other widely used non-negative similarity measurements. Furthermore, RFMS is not
 984 restricted to Fourier-based kernels. It can accommodate other random feature constructions, such as
 985 Polynomial kernels via random Maclaurin expansions or dot-product kernels. In this sense, RFMS is
 986 a general mode-seeking framework over random feature approximated kernel densities, and is not
 987 inherently limited to any particular kernels.

990 C.3 BOUNDS TIGHTNESS

992 Theorem 1 attains the canonical Monte Carlo RFF rate where the KDE approximation error decays as
 993 $\mathcal{O}(1/\sqrt{D})$ [65] in the number of features D . This dependence is standard (and essentially optimal)
 994 for vanilla i.i.d. feature sampling. Theorems 3-4 extend the same $\mathcal{O}(1/\sqrt{D})$ dependence to the
 995 gradient and Hessian. The extra factors are the usual cost of controlling the derivatives uniformly, not
 996 a deterioration in the RFF sampling rate. We note that the constants in the bounds can be conservative
 997 due to the use of Hoeffding-type inequality arguments, but the rates in D are the key notion of
 998 tightness here and are sufficient for our mode-tracking guarantees. Practically, this yields a clean
 999 accuracy-efficiency knob; the RFF-KDE is unbiased w.r.t. feature randomness, so increasing D
 1000 reduces variance and improves approximation, while runtime scales roughly linearly in D and T .
 Thus, the approximation can be made arbitrarily tight by choosing D large enough.

1002 C.4 RELATION TO HDC/VSA

1004 Hyperdimensional computing (HDC)—also known as Vector Symbolic Architecture (VSA)— is
 1005 a class of computational models that represent and manipulate structured information using high-
 1006 dimensional vectors. The characteristic of VSA/HDC is that it first encodes data as high-dimensional
 1007 vectors and operates on encoded data using a set of simple algebraic operations that are efficient
 1008 and highly parallelizable [41; 64]. HDC is connected to theoretical neuroscience as its mathematical
 1009 framework closely resembles models of neural coding in the brain [22]. As a result, various machine
 1010 learning algorithms based on HDC/VSA have been proposed, such as classification [23], clustering
 1011 [24], and regression [31]. However, there is a lack of density-based analysis methods in existing
 1012 HDC/VSA literature. The notion of high-dimensional, distributed, and compositional representation
 1013 of HDC/VSA aligns closely with RFMS. In that sense, RFMS can be viewed functionally as an
 1014 HDC/VSA algorithm, therefore filling in the gap between HDC/VSA and density-based analysis.
 Furthermore, HDC/VSA has also been extensively studied, especially within the hardware community.
 Various types of accelerators [79; 67] have been proposed for HDC/VSA workloads. This explicit
 connection between RFMS and HDC/VSA, and their computational similarity, can potentially lead to
 the use of existing HDC/VSA accelerators for RFMS, providing practical benefits.

1019 C.5 SIGNIFICANCE OF MEAN-SHIFT ALGORITHMS

1021 Mean-Shift, as a geometry-respecting procedure, is broadly useful across machine learning and data
 1022 analysis, some notable examples including vision, anomaly detection, self-supervised learning and
 1023 more [38; 44; 75]. Despite known limitations such as bandwidth sensitivity and the applicability
 1024 in high-dimensional data, mean-shift remains meaningful because it is often used as an algorithmic
 1025 primitive and is still being actively used by recent research [2; 53; 44]. Furthermore, the regimes
 where mean-shift performs well are well understood [25; 15], and our approach provides a faithful

1026 approximation with asymptotic complexity improvements, making it desirable for many existing and
 1027 emerging applications.
 1028

1029 D LIMITATION & FUTURE WORK

1030 D.1 MORSE FUNCTION ASSUMPTION

1031 In Theorem 5, our results regarding the stability of critical points have the assumption that the KDE
 1032 is a Morse function, meaning a smooth function with non-degenerate critical points. This is, in fact, a
 1033 standard assumption in density-based mode analysis [1; 11]. Moreover, the general body of literature
 1034 regarding the mean-shift algorithm assumes that gradient ascent on a KDE surface is well-behaved.
 1035 This implicitly assumes the density function has isolated, non-degenerate modes, aligning with the
 1036 Morse function.
 1037

1038 D.2 SEMANTIC SEGMENTATION

1039 In the context of image processing, mean-shift is a non-parametric unsupervised algorithm based on
 1040 low-level features such as color and/or spatial proximity. It operates on the pixel level and groups
 1041 pixels based on local density in a feature space, not on high-level semantic categories. Consequently,
 1042 mean-shift can segment coherent regions but cannot segment regions based on semantic information.
 1043 Despite its limitations, mean-shift can still be helpful in roles like generating superpixels or region
 1044 proposals, which can be a key step in semantic segmentation with deep neural networks [55; 6; 45;
 1045 59; 77].
 1046

1047 D.3 MODE-SEEKING FOR HIGH-DIMENSIONAL DATA

1048 In the high-dimensional regime, both RFMS and classical mean shift (or any KDE-based method)
 1049 are fundamentally limited by the curse of dimensionality. Our goal here is not to fix this statistical
 1050 issue, but to provide a computationally scalable approximation to classical mean-shift in regimes
 1051 where it is still used, which leads to substantial runtime gains in large n settings for the low-to-
 1052 moderate dimensional data. Despite this limitation, it remains useful and a powerful primitive in many
 1053 applications [38; 44; 75; 2; 53; 44; 14]. We believe RFMS can be incorporated with other methods
 1054 (e.g., dimensionality reduction methods such as in [54; 49]) and be explored in higher-dimensional
 1055 regimes. Our present goal, however, is to lay the theoretical foundations and analyze frameworks for
 1056 this sampling-based mean-shift estimator to support any future extensions.
 1057

1058 D.4 FUTURE WORK

1059 The primary aim of this paper is to establish the algorithmic and theoretical foundation of RFMS.
 1060 We view this work as a principled first step toward scalable, kernel-based mode-seeking for domains
 1061 such as tracking and point cloud, as these applications have been addressed by inefficient forms of
 1062 mean-shift [37]. RFMS is also designed to be modular and extensible: different random-feature maps
 1063 and zeroth-order gradient strategies can be used depending on the application. We believe RFMS can
 1064 be incorporated into other learning pipelines (e.g., deep neural networks) [2; 53; 44]. It is also well
 1065 known that the mean-shift algorithms work best in a low-to-moderate dimension regime, as kernel
 1066 density estimation suffers from the curse of dimensionality [25]. As a potential future direction, we
 1067 would also like to extend RFMS into high-dimensional regimes [10; 9]. In any case, this paper can
 1068 serve as a theoretical foundation for this sampling-based mean-shift estimator in support of any future
 1069 extensions.
 1070

1071 E PROOF OF EQUATION 4

1072 *Proof.* Want to show $\hat{g}(x)$ is an unbiased estimation of the kernel density estimation $\hat{f}(x)$ with
 1073 respect to the randomness in $\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$. Recall the definition of ϕ_{k_h} :

$$1080 \\ 1081 \quad \phi_{k_h} = \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T x)} e^{j(\omega_2^T x)} e^{j(\omega_3^T x)} \dots e^{j(\omega_D^T x)} \right] \quad (15) \\ 1082 \\ 1083$$

1084 We start by showing that the complex inner products between encodings produced by ϕ_{k_h} are an
1085 unbiased estimation of the kernel function k_h . For any $x, x' \in \mathbb{R}^d$, the results follow directly after
1086 linearity of expectation and Bochner's theorem:

$$1087 \\ 1088 \quad \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\langle \phi_{k_h}(x), \phi_{k_h}(x') \rangle] \\ 1089 \\ 1090 \quad = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\frac{1}{D} \sum_{k=1}^D \overline{e^{j(\omega_k^T x)}} e^{j(\omega_k^T x')} \right] \\ 1091 \\ 1092 \quad = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\overline{e^{j(\omega^T x)}} e^{j(\omega^T x')} \right] \quad (16) \\ 1093 \\ 1094 \quad = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[e^{j\omega^T (x' - x)} \right] \\ 1095 \\ 1096 \quad = \int_{\mathbb{R}^d} p(\omega) e^{j\omega \cdot (x' - x)} d\omega \\ 1097 \\ 1098 \quad = k_h(x', x) = k_h(x, x') \\ 1099$$

1100 Recall that $\Phi = \frac{1}{nc} \sum_{i=1}^n \phi_{k_h}(x_i) \forall x_i \in \mathcal{D}$ and $\hat{g}(x) = \langle \phi_{k_h}(x), \Phi \rangle$, so:

$$1101 \\ 1102 \\ 1103 \quad \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\langle \phi_{k_h}(x), \Phi \rangle] \\ 1104 \\ 1105 \quad = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\overline{\phi_{k_h}(x)} \cdot \left(\frac{1}{nc} \sum_{i=1}^n \phi_{k_h}(x_i) \right) \right] \\ 1106 \\ 1107 \quad = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\frac{1}{nc} \sum_{i=1}^n \left(\overline{\phi_{k_h}(x)} \cdot \phi_{k_h}(x_i) \right) \right] \quad (17) \\ 1108 \\ 1109 \quad = \frac{1}{nc} \sum_{i=1}^n k_{k_h}(x, x_i) = \hat{f}_{k_h}(x) \\ 1110 \\ 1111 \\ 1112 \\ 1113 \\ 1114 \\ 1115$$

F PROOF OF THEOREM 1

1116 *Proof.* First expand $|\hat{g}(x) - \hat{f}_{k_h}(x)|$:

$$1117 \\ 1118 \\ 1119 \quad \left| \hat{g}(x) - \hat{f}_{k_h}(x) \right| = \left| \langle \phi_{k_h}(x), \Phi \rangle - \frac{1}{nc} \sum_{i=1}^n k_h(x, x_i) \right| \\ 1120 \\ 1121 \quad = \left| \left\langle \phi_{k_h}(x), \frac{1}{nc} \sum_{i=1}^n \phi_{k_h}(x_i) \right\rangle - \frac{1}{nc} \sum_{i=1}^n k_h(x, x_i) \right| \\ 1122 \\ 1123 \quad = \left| \frac{1}{nc} \sum_{i=1}^n \langle \phi_{k_h}(x), \phi_{k_h}(x_i) \rangle - \frac{1}{nc} \sum_{i=1}^n k_h(x, x_i) \right| \quad (18) \\ 1124 \\ 1125 \\ 1126 \\ 1127 \\ 1128 \\ 1129 \\ 1130 \\ 1131 \\ 1132 \\ 1133$$

1134 We define an estimator:
 1135

1136
 1137
$$Z_k = \sum_{i=1}^n e^{j\omega_k^T (x_i - x)} \quad (19)$$

 1138

1139 Since Random Feature is unbiased for approximating individual kernel values, we know that:
 1140

1141
 1142
$$\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] = \sum_{i=1}^n k_h(x, x_i) \quad (20)$$

 1143

1144 Since the sampling of random features ω in the construction of ϕ_{k_h} is uniform i.i.d., the approximation
 1145 is an average of i.i.d. complex-valued random variables Z_k , with expectation equal to the target
 1146 kernel sum. Split Z_k into real and imaginary parts using sine and cosine:
 1147

1148
 1149
$$Z_k = A_k + jB_k \quad \text{where: } A_k = \sum_{i=1}^n \cos(\omega_k^T (x_i - x)), B_k = \sum_{i=1}^n \sin(\omega_k^T (x_i - x)) \quad (21)$$

 1150

1151 Since sine and cosine functions are bounded between $[-1, 1]$, it is clear that $A_k, B_k \in [-n, n]$, which
 1152 is bounded. Use Hoeffding's inequality for both A_k and B_k . Chase any $\delta \in (0, 1)$ with probability if
 1153 at least $1 - \delta/2$:

1154
 1155
$$\begin{aligned} \left| \frac{1}{D} \sum_{k=1}^D A_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [A_k] \right| &\leq 2n \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}} \\ \left| \frac{1}{D} \sum_{k=1}^D B_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [B_k] \right| &\leq 2n \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}} \end{aligned} \quad (22)$$

 1156

1157 Since k_h is real valued, meaning $\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [B_k] = 0$, so $\mathbb{E}[A_k] = \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] =$
 1158 $\sum_{i=1}^n k_h(x, x_i)$. Recall that for complex number Z_k , $|Z_k| = \sqrt{A_k^2 + B_k^2} \leq |A_k| + |B_k|$, and the
 1159 probability of both the imaginary and real inequality holds is $(1 - \delta/2)^2 > 1 - \delta$, so:
 1160

1161
 1162
$$\begin{aligned} \left| \frac{1}{D} \sum_{k=1}^D Z_k - \sum_{i=1}^n k_h(x, x') \right| &\leq \left| \frac{1}{D} \sum_{k=1}^D A_k - \sum_{i=1}^n k_h(x, x_i) \right| + \left| \frac{1}{D} \sum_{k=1}^D B_k \right| \\ &\leq 4n \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}} \end{aligned} \quad (23)$$

 1163

1164 Finally:
 1165

1166
 1167
$$\begin{aligned} \left| \hat{g}(x) - \hat{f}_{k_h}(x) \right| &= \frac{1}{nc} \left| \frac{1}{D} \sum_{k=1}^D Z_k - \sum_{i=1}^n k_h(x, x') \right| \\ &\leq \frac{4}{c} \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}} \quad \text{With probability at least } 1 - \delta \end{aligned} \quad (24)$$

 1168

1169 So, for any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the above error bound holds. Additionally,
 1170 suppose KED are estimated via $\text{Re}(\hat{g}(x))$. In that case, there will be no additional variance from the
 1171 imaginary part, further reducing the bound to $\frac{2}{c} \sqrt{\frac{1}{2D} \ln \frac{4}{\delta}}$.
 1172

1188 G PROOF OF THEOREM 2
11891190 *Proof.* Let \mathcal{X} be a compact set over \mathbb{R}^d such that $\mathcal{D} \subset \mathcal{X}$. The goal is to extend Theorem 1 to a
1191 uniform convergence statement over the entire \mathcal{X} .
1192

1193
1194
$$\sup_{x \in \mathcal{X}} \left| \left(\operatorname{Re}(\hat{g}(x)) - \hat{f}_{k_h}(x) \right) \right| = \sup_{x \in \mathcal{X}} \left| \frac{1}{nc} \left| \sum_{i=1}^n \operatorname{Re}(\langle \phi_{k_h}(x), \phi_{k_h}(x_i) \rangle) - \sum_{i=1}^n k_h(x, x_i) \right| \right| \quad (25)$$

1195
1196
1197
$$\leq \frac{1}{c} \sup_{x, y \in \mathcal{X}} \left| \left| \operatorname{Re}(\langle \phi_{k_h}(x), \phi_{k_h}(y) \rangle) - k_h(x, y) \right| \right|$$

1198
1199

1200 Ignoring the constant $\frac{1}{c}$ for now, the remaining is the uniform convergence of random Fourier features,
1201 which has been studied by Rahimi and Recht [65], for an error tolerance epsilon ϵ :
1202

1203
1204
$$\Pr \left(\sup_{x, y \in \mathcal{X}} \left| \left| \operatorname{Re}(\langle \phi_{k_h}(x), \phi_{k_h}(y) \rangle) - k_h(x, y) \right| \right| \geq \epsilon \right)$$

1205
1206
$$\leq 2^8 \left(\frac{\sqrt{\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\langle \omega, \omega \rangle]} \operatorname{diam}(\mathcal{X})}{\epsilon} \right)^2 \exp \left(-\frac{D\epsilon^2}{4(d+2)} \right) \quad (26)$$

1207
1208
1209

1210 Next, solve $\sqrt{\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\langle \omega, \omega \rangle]}$:
1211
1212

1213
1214
$$\sqrt{\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\langle \omega, \omega \rangle]} = \sqrt{\frac{d}{h^2}} = \frac{\sqrt{d}}{h} \quad (27)$$

1215

1216 Putting everything back together:
1217
1218

1219
1220
$$\Pr \left(\sup_{x \in \mathcal{X}} \left| \operatorname{Re}(\hat{g}(x)) - \hat{f}_{k_h}(x) \right| \geq \epsilon \right) \leq 2^8 \left(\frac{c\sqrt{d} \operatorname{diam}(\mathcal{X})}{h\epsilon} \right)^2 \exp \left(-\frac{D\epsilon^2}{c^2 4(d+2)} \right) \quad (28)$$

1221
1222
1223
1224

H PROOF OF THEOREM 3

1225 *Proof.* Here, we would like to show a bound between the gradient of the Random Feature approxi-
1226 mated KDE and the gradient of the actual KDE. Specifically, since we are using real-valued version
1227 of $\hat{g}(x)$, we what to show a bound on $\left\| \nabla \operatorname{Re}(\hat{g}(x)) - \nabla \hat{f}_{k_h}(x) \right\|$. We first expand:
1228
1229

1230
1231
$$\left\| \nabla \operatorname{Re}(\hat{g}(x)) - \nabla \hat{f}_{k_h}(x) \right\| = \frac{1}{nc} \left\| \sum_{i=1}^n \frac{1}{D} \sum_{k=1}^D \nabla \operatorname{Re} \left(e^{j\omega_k^T(x_i - x)} \right) - \sum_{i=1}^n \nabla k_h(x, x_i) \right\| \quad (29)$$

1232
1233
1234
1235
1236
1237

1238 We define an estimator:
1239

1240
1241
$$Z_k = \sum_{i=1}^n \nabla \cos(\omega_k^T(x_i - x)) \quad (30)$$

1242

Since Differentiation commutes with expectation:

1243

1244

1245

1246

$$\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] = \nabla \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\sum_{i=1}^n \cos(\omega_k^T (x_i - x)) \right] = \sum_{i=1}^n \nabla k_h(x, x_i) \quad (31)$$

1247

To bound the deviation of Z_k from its expectation, we apply Bernstein's inequality:

1248

1249

1250

1251

1252

1253

However, since $\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$ is unbounded, random variable Z_k is also unbounded, so the standard Bernstein inequality does not apply. However, notice that the ω is sub-exponential, we apply the Bernstein inequality for a tail-heavy random variable.

1254

1255

1256

1257

1258

Consider the version of Bernstein inequality presented by Lanthaler et al. [46] which states: Let Z be a sub-exponential random variable in a separable Hilbert space, choose any $\delta \in (0, 1)$, with probability at least $1 - \delta$, the following holds:

1259

1260

1261

1262

$$\left\| \frac{1}{D} \sum_{k=1}^D Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] \right\| \leq \frac{2b \ln(2/\delta)}{D} + \sqrt{\frac{2\sigma^2 \ln(2/\delta)}{D}} \quad (33)$$

1263

Where:

1264

1265

1266

1267

1268

$$\sigma^2 = 4e \sqrt{\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k]\|^2]} \|Z_k\|_{\psi_1} \quad \text{and} \quad b = 4e \|Z_k\|_{\psi_1} \quad (34)$$

1269

1270

$\|\cdot\|_{\psi_1}$ denotes sub-exponential Orlicz norm. Since \mathbb{R}^d is a separable Hilbert space, it is directly applicable here. First bounding:

1271

1272

1273

1274

1275

1276

1277

1278

$$\begin{aligned} & \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\left\| Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] \right\|^2 \right] \\ &= \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|^2] - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|^2] \\ &\leq \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|^2] \end{aligned} \quad (35)$$

1279

Since the sine function is bounded:

1280

1281

1282

1283

$$\|Z_k\| \leq n \|\omega_k\| \quad (36)$$

1284

Recall that ω_k are drawn from $\mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$, so:

1285

1286

1287

1288

$$\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|\omega_k\|^2] \leq n^2 \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|\omega_k\|^2] \leq \frac{n^2 d}{h^2} \quad (37)$$

1289

Now bounding $\|Z_k\|_{\psi_1}$:

1290

1291

1292

$$\|Z_k\|_{\psi_1} = \|\|Z_k\|\|_{\psi_1} \leq n \|\omega_k\|_{\psi_1} = n \|\omega_k\|_{\psi_1} = n \|\omega_k\|_{\psi_1} = \frac{n}{h} \|z\|_{\psi_1} \quad (38)$$

1293

1294

Where $z \sim \mathcal{N}(0, I)$. We know that [70]:

1295

$$\|z\|_{\psi_1} \leq C \sqrt{d} \quad (39)$$

1296 Where C is a universal constant, so:
 1297

$$1298 \quad 1299 \quad 1300 \quad \|Z_k\|_{\psi_1} \leq \frac{Cn\sqrt{d}}{h} \quad (40)$$

1301 Which means:
 1302

$$1303 \quad 1304 \quad 1305 \quad \sigma^2 \leq \frac{4eCn^2d}{h^2} \quad \text{and} \quad b \leq \frac{4eCn\sqrt{d}}{h} \quad (41)$$

1306 We can bound the deviation of Z_k from its expectation:
 1307

$$1308 \quad 1309 \quad 1310 \quad \left\| \frac{1}{D} \sum_{k=1}^D Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] \right\| \leq \frac{8Cen\sqrt{d} \ln(2/\delta)}{Dh} + \sqrt{\frac{8Cen^2d \ln(2/\delta)}{Dh^2}} \quad (42)$$

1311 The above holds with probability at least $1 - \delta$. And finally, also with probability at least $1 - \delta$:
 1312

$$1313 \quad 1314 \quad 1315 \quad \left\| \nabla \text{Re}(\hat{g}(x)) - \nabla \hat{f}_{k_h}(x) \right\| \leq \frac{1}{nc} \left(\frac{8Cen\sqrt{d} \ln(2/\delta)}{Dh} + \sqrt{\frac{8Cen^2d \ln(2/\delta)}{Dh^2}} \right) \quad (43)$$

■

1320 I PROOF OF THEOREM 4

1321 *Proof.* We also want to bound the Frobenius norm of the difference between the Hessian:
 1322

$$1323 \quad 1324 \quad 1325 \quad \left\| \nabla^2 \text{Re}(\hat{g}(x)) - \nabla^2 \hat{f}_{k_h}(x) \right\|_F = \frac{1}{nc} \left\| \frac{1}{D} \sum_{k=1}^D \sum_{i=1}^n \nabla^2 \cos(\omega_k^T (x_i - x)) - \sum_{i=1}^n \nabla^2 k_h(x, x_i) \right\|_F \quad (44)$$

1326 Define an estimator:
 1327

$$1328 \quad 1329 \quad 1330 \quad Z_k = \sum_{i=1}^n \nabla^2 \cos(\omega_k^T (x_i - x)) \quad (45)$$

1331 Again, because differentiation commutes with expectation, know that:
 1332

$$1333 \quad 1334 \quad 1335 \quad \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] = \nabla^2 \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\sum_{i=1}^n \cos(\omega_k^T (x_i - x)) \right] = \sum_{i=1}^n \nabla^2 k_h(x, x_i)$$

1336 (46)

1337 Use the proving technique as seen in the proof of Theorem 3. We start by bounding σ^2 and b :
 1338

$$1339 \quad 1340 \quad 1341 \quad Z_k = \sum_{i=1}^n -\cos(\omega_k^T (x_i - x)) \omega_k \omega_k^T \quad (47)$$

1342 Then:
 1343

$$\begin{aligned}
& \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} \left[\left\| Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] \right\|_F^2 \right] \\
&= \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|_F^2] - \|\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k]\|_F^2 \\
&\leq \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|_F^2]
\end{aligned} \tag{48}$$

Since the cosine function is bounded:

$$\|Z_k\|_F \leq n \|\omega_k \omega_k^T\|_F \tag{49}$$

Recall that ω_k are drawn from $\mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)$, so:

$$\mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|Z_k\|_F^2] \leq n^2 \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [\|\omega_k \omega_k^T\|_F^2] \leq \frac{n^2 d(d+2)}{h^4} \tag{50}$$

Now bounding $\|Z_k\|_{\psi_1}$:

$$\|Z_k\|_{\psi_1} = \|\|Z_k\|_F\|_{\psi_1} \leq \|n \|\omega_k \omega_k^T\|_F\|_{\psi_1} = n \|\|\omega_k \omega_k^T\|_F\|_{\psi_1} = n \|\omega_k \omega_k^T\|_{\psi_1} \leq \frac{Cnd}{h^2} \tag{51}$$

Which means:

$$\sigma^2 \leq \frac{4eCn^2d\sqrt{d(d+2)}}{h^4} \quad \text{and} \quad b \leq \frac{4eCnd}{h^2} \tag{52}$$

Bound the deviation of Z_k from its expectation:

$$\left\| \frac{1}{D} \sum_{k=1}^D Z_k - \mathbb{E}_{\omega \sim \mathcal{N}(0, \frac{1}{h^2} \mathbf{I}_d)} [Z_k] \right\|_F \leq \frac{8eCnd \ln(2/\delta)}{Dh^2} + \sqrt{\frac{8eCn^2d\sqrt{d(d+2)} \ln(2/\delta)}{Dh^4}} \tag{53}$$

The above holds with probability at least $1 - \delta$. And finally, also with probability at least $1 - \delta$:

$$\left\| \nabla^2 \text{Re}(\hat{g}(x)) - \nabla^2 \hat{f}_{k_h}(x) \right\|_F \leq \frac{1}{nc} \left(\frac{8eCnd \ln(2/\delta)}{Dh^2} + \sqrt{\frac{8eCn^2d\sqrt{d(d+2)} \ln(2/\delta)}{Dh^4}} \right) \tag{54}$$

■

J PROOF OF THEOREM 5

Proof. The aim here is to show the critical points between $\text{Re}(\hat{g})$ and \hat{f}_{k_h} are close under a mild assumption. Assume \hat{f}_{k_h} is a Morse function, i.e., a smooth function with non-degenerate critical points [29], let x^* be any critical point of \hat{f}_{k_h} , we have:

$$\nabla f_{k_h}(x^*) = 0 \quad \text{and} \quad \nabla^2 f_{k_h}(x^*) \text{ is invertible with all positive eigenvalues} \tag{55}$$

In Theorem 3 and 4, we have shown that $\text{Re}(\hat{g})$ and \hat{f}_{k_h} , are point-wise close in terms of their gradient and hessian for all $x \in \mathbb{R}^d$, suppose:

1404

1405

1406

$$\begin{aligned} \|\nabla \operatorname{Re}(\hat{g}(x)) - \nabla f_{k_h}(x)\| &\leq \epsilon_1 \quad \text{With probability at least } 1 - \delta/2 \\ \|\nabla^2 \operatorname{Re}(\hat{g}(x)) - \nabla^2 f_{k_h}(x)\|_F &\leq \epsilon_2 \quad \text{With probability at least } 1 - \delta/2 \end{aligned} \quad (56)$$

1408

1409

Where ϵ_1, ϵ_2 can be made arbitrarily small by increasing D (see Theorem 3 and Theorem 4).

1410

1411

We're looking for \hat{x}^* near x^* such that:

1412

1413

$$\nabla \operatorname{Re}(\hat{g}(\hat{x}^*)) = 0 \quad (57)$$

1414

1415

\hat{x}^* can be found via Newton iteration:

1416

1417

$$T(x) = x - [\nabla^2 \operatorname{Re}(\hat{g}(x))]^{-1} \nabla \operatorname{Re}(\hat{g}(x)) \quad (58)$$

1418

1419

Start at x^* , we can show the distance between x^* and \hat{x}^* using Newton–Kantorovich theorem [42], which gives the optimality and convergence result of Newton's Method.

1420

1421

We first bound the gradient residual at x^* :

1422

1423

$$\|\nabla \operatorname{Re}(\hat{g}(x^*))\| = \left\| \nabla \operatorname{Re}(\hat{g}(x^*)) - \nabla \hat{f}_{k_h}(x^*) \right\| \leq \epsilon_1 \quad (59)$$

1424

1425

Then, show the invertibility of $\nabla^2 \operatorname{Re}(\hat{g}(x^*))$ by showing all its eigenvalue are positive. Let $\lambda_{\min}(\cdot)$ denote the smallest eigenvalue of a matrix, by Weyl's inequality:

1426

1427

1428

$$\begin{aligned} &\left| \lambda_{\min}(\nabla^2 \operatorname{Re}(\hat{g}(x^*))) - \lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) \right| \\ &\leq \|\nabla^2 \operatorname{Re}(\hat{g}(x^*)) - \nabla^2 f_{k_h}(x^*)\|_{op} \\ &\leq \|\nabla^2 \operatorname{Re}(\hat{g}(x^*)) - \nabla^2 f_{k_h}(x^*)\|_F \\ &\leq \epsilon_2 \end{aligned} \quad (60)$$

1429

1430

Which implies:

1431

1432

1433

1434

1435

$$\lambda_{\min}(\nabla^2 \operatorname{Re}(\hat{g}(x^*))) \geq \lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2 \quad (61)$$

1436

1437

This means, in order for $\nabla^2 \operatorname{Re}(\hat{g}(x^*))$ to be invertible, simply chose $\epsilon_2 \leq \lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*))$.

1438

1439

Further more we can bound the operator norm of $[\nabla^2 \operatorname{Re}(\hat{g}(x^*))]^{-1}$:

1440

1441

1442

1443

1444

1445

1446

1447

1448

Since we successfully bounded the operator norm, we also know that:

1449

1450

1451

1452

$$\left\| [\nabla^2 \operatorname{Re}(\hat{g}(x^*))]^{-1} \nabla \operatorname{Re}(\hat{g}(x^*)) \right\| \leq \frac{\epsilon_1}{\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2} \quad (63)$$

1453

1454

1455

1456

1457

Define:

$$\alpha = \frac{\operatorname{Lip}(\nabla^2 \operatorname{Re}(\hat{g})) \epsilon_1}{\left(\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2 \right)^2} \quad (64)$$

1458 If $\alpha < \frac{1}{2}$, Newton–Kantorovich theorem [60; 33] states that:
 1459

$$1460 \quad 1461 \quad 1462 \quad 1463 \quad \hat{x}^* \in B \left(x^*, \frac{\epsilon_1}{\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2} \right) \quad (65)$$

1464 $B(x, r)$ means a closed ball centered at x with radius r , and \hat{x}^* is the only critical point in the region.
 1465

1466 Overall, what this means is that for any critical point x^* of \hat{f}_{k_h} , if we choose:
 1467

- 1468 1. $\epsilon_2 \leq \lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*))$
- 1469 2. ϵ_1 such that $\alpha \leq \frac{1}{2}$
- 1470 3. $\|\nabla \operatorname{Re}(\hat{g}(x)) - \nabla f_{k_h}(x)\| \leq \epsilon_1$ and $\|\nabla^2 \operatorname{Re}(\hat{g}(x)) - \nabla^2 f_{k_h}(x)\|_F \leq \epsilon_2$ both holds with
 1471 probability at least $1 - \delta/2$.

1472 Then, with probability least $1 - \delta$, there is only one critical point \hat{x}^* of $\operatorname{Re}(\hat{g})$ such that: $\hat{x}^* \in$
 1473 $B \left(x^*, \frac{\epsilon_1}{\lambda_{\min}(\nabla^2 \hat{f}_{k_h}(x^*)) - \epsilon_2} \right)$.
 1474

■

1480 K PROOF OF THEOREM 6

1481 The convergence of zeroth-order gradient methods using two-point gradient estimation over nonconvex but L -smooth function is established by Nesterov and Spokoiny [58], who showed the method
 1482 converges to approximate stationary points at a sublinear rate. Specifically, in our case, with decaying
 1483 step-size, the average squared gradient norm satisfies:

$$1484 \quad 1485 \quad \min_{0 \leq l \leq T} \mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d)} \left\| \nabla \operatorname{Re} \left(\hat{g} \left(x^{(l)} \right) \right) \right\|^2 = \mathcal{O} \left(\frac{1}{\sqrt{T}} + \mu^2 d^2 \right) \quad (66)$$

1486 The second term comes from the fact that gradient estimation is biased (due to the smoothing
 1487 parameter μ) but close to the real gradient. We can further derive a bound to quantify the result
 1488 of zero-th order gradient ascent over $\operatorname{Re}(\hat{g})$: Let $l^* \in \{0 \dots T\}$ be the iteration index such that
 1489 $\mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d)} \left\| \nabla \operatorname{Re} \left(\hat{g} \left(x^{(l)} \right) \right) \right\|^2$ is the smallest. Since $\operatorname{Re}(\hat{g})$ is real analytical, it satisfies Łojasiewicz
 1490 inequality around a local mode \hat{x}^* :

$$1491 \quad 1492 \quad 1493 \quad \operatorname{Re}(\hat{g}(\hat{x}^*)) - \operatorname{Re} \left(\hat{g} \left(x^{(l^*)} \right) \right) \leq \mathcal{O} \left(\left\| \nabla \operatorname{Re} \left(\hat{g} \left(x^{(l^*)} \right) \right) \right\|^{2/2\theta} \right) \quad (67)$$

1494 Where θ is the Łojasiewicz exponent. Combine with the descent lemma:
 1495

$$1496 \quad 1497 \quad 1498 \quad \frac{1}{2 \operatorname{Lip}(\nabla \operatorname{Re}(\hat{g}))} \left\| \hat{x}^* - x^{(l^*)} \right\|^2 \leq \operatorname{Re}(\hat{g}(\hat{x}^*)) - \operatorname{Re} \left(\hat{g} \left(x^{(l^*)} \right) \right) \leq \mathcal{O} \left(\left\| \nabla \operatorname{Re} \left(\hat{g} \left(x^{(l^*)} \right) \right) \right\|^{2/2\theta} \right) \quad (68)$$

1500 Taking the expectation on both sides yields:
 1501

$$1502 \quad 1503 \quad 1504 \quad \mathbb{E}_{u \sim \mathcal{N}(0, \mathbf{I}_d)} \left[\left\| \hat{x}^* - x^{(l^*)} \right\|^2 \right] \leq \mathcal{O} \left(\left(\frac{1}{\sqrt{T}} + \mu^2 d^2 \right)^{1/2\theta} \right) \quad (69)$$

■

1512 **L PROOF OF EQUATION 13**
15131514 *Proof.* Want to show that $\phi_{k_h}(x + x') = \frac{D}{\sqrt{D}} \phi_{k_h}(x) \otimes \phi_{k_h}(x') \quad \forall x, x' \in \mathbb{C}^d$. Recall that $\phi(\cdot)$ is:
1515

1516
$$\phi_{k_h}(x) = \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T x)} e^{j(\omega_2^T x)} e^{j(\omega_3^T x)} \dots e^{j(\omega_D^T x)} \right] \quad (70)$$

1517

1518 So:
1519

1520
$$\begin{aligned} \frac{D}{\sqrt{D}} \phi_{k_h}(x) \otimes \phi_{k_h}(x') &= \frac{D}{\sqrt{D}} \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T x)} e^{j(\omega_2^T x)} e^{j(\omega_3^T x)} \dots e^{j(\omega_D^T x)} \right] \\ &\otimes \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T x')} e^{j(\omega_2^T x')} e^{j(\omega_3^T x')} \dots e^{j(\omega_D^T x')} \right] \\ &= \frac{D}{\sqrt{D} D} \left[e^{j(\omega_1^T x)} e^{j(\omega_1^T x')} \dots e^{j(\omega_D^T x)} e^{j(\omega_D^T x')} \right] \\ &= \frac{D}{\sqrt{D} D} \left[e^{j(\omega_1^T (x+x'))} \dots e^{j(\omega_D^T (x+x'))} \right] \\ &= \frac{1}{\sqrt{D}} \left[e^{j(\omega_1^T (x+x'))} \dots e^{j(\omega_D^T (x+x'))} \right] \\ &= \phi_{k_h}(x + x') \end{aligned} \quad (71)$$

1521
1522
1523
1524
1525
1526
1527
1528
1529
1530
1531
1532

■

1533 **M ABLATION STUDY**
15341535 **M.1 HYPERPARAMETER SENSITIVITY**
15361537 Here, we provide a sensitivity analysis on important RFMS hyperparameters, including the mapped
1538 dimension(d), number of iterations (T), kernel bandwidth(h), learning rate(η), and smoothing
1539 parameter(μ). Experimental results here are done on S4 dataset. Unless specified in the table below,
1540 we use the default $D = 500$, $T = 100$, $h = 0.2$, $\eta = 0.003$, $\mu = 5e - 4$.
15411542
1543
1544

S4 dataset	$D = 50$	$D = 100$	$D = 200$	$D = 300$	$D = 500$
NMI	64	67	68	69	69
Time(s)	1.7	2.7	6.2	10.3	17.8

15451546
1547
1548
1549
1550

S4 dataset	$T = 10$	$T = 50$	$T = 100$	$T = 150$	$T = 200$
NMI	0	35	70	70	71
Time(s)	2.1	8.6	17.4	25.8	33.7

15511552
1553
1554
1555

S4 dataset	$h = 0.05$	$h = 0.1$	$h = 0.2$	$h = 0.4$	$h = 0.8$
NMI	50	62	67	10	0

15561557
1558
1559
1560

S4 dataset	$\eta = 0.00075$	$\eta = 0.0015$	$\eta = 0.003$	$\eta = 0.006$	$\eta = 0.1$
NMI	49	54	69	71	28

15611562
1563
1564
1565

S4 dataset	$\mu = 5e - 6$	$\mu = 5e - 5$	$\mu = 5e - 4$	$\mu = 5e - 3$	$\mu = 1$
NMI	69	70	70	64	5

1566 We introduce η as a step-size parameter so the mean-shift update matches the familiar gradient ascent
 1567 form. This also lets us control how far we move each iteration. In practice, the kernel affects the
 1568 scale of the update through bandwidth, and η provides a convenient way to absorb/adjust that step
 1569 magnitude and to improve stability under approximation/zeroth-order noise.

1570 Among the hyperparameters, only D and T will affect the runtime of the algorithm. As claimed in
 1571 the paper, runtime scales linearly in both T and D . Larger T and D will almost certainly produce
 1572 better results, which aligns with Theorems 5 and 6 in the paper that quantify the closeness between
 1573 the result returned by RFMS and actual modes of KDE. RFMS can be sensitive to h, η, μ . Similar to
 1574 many machine learning approaches, those hyperparameters need to be chosen empirically based on
 1575 the on-hand data. The advantage of RFMS is that it allows for much faster hyperparameter-tuning
 1576 due to the algorithm's efficiency.

1578 M.2 MONTE-CARLO(MC) vs. QUASI MONTE-CARLO(QMC)

1580 Here, we add additional experiments to disentangle the effect of Monte-Carlo sampling (MC) and
 1581 Quasi Monte-Carlo sampling (QMC) in RFMS at varying D , better NMI are highlighted:

NMI ($D = 50$)	a1	a2	a3	unbalance	s1	s2	s3	s4
MC	0.83	0.86	0.85	0.98	0.92	0.86	0.67	0.65
QMC	0.85	0.86	0.86	0.98	0.94	0.87	0.71	0.66

NMI ($D = 100$)	a1	a2	a3	unbalance	s1	s2	s3	s4
MC	0.87	0.87	0.89	0.98	0.96	0.89	0.7	0.64
QMC	0.89	0.89	0.90	0.98	0.96	0.91	0.74	0.68

NMI ($D = 200$)	a1	a2	a3	unbalance	s1	s2	s3	s4
MC	0.89	0.91	0.92	0.98	0.97	0.91	0.74	0.63
QMC	0.92	0.93	0.94	0.98	0.98	0.93	0.74	0.68

NMI ($D = 300$)	a1	a2	a3	unbalance	s1	s2	s3	s4
MC	0.92	0.94	0.95	0.98	0.98	0.92	0.75	0.68
QMC	0.95	0.97	0.96	0.98	0.98	0.93	0.76	0.69

1603 At the same D , QMC shows a consistent improvement over MC sampling as expected.

1606 N ADDITIONAL DETAILS ON CLUSTERING EXPERIMENTS

1608 All clustering experiments were run on an Intel i5-11400 CPU. All algorithms are implemented in
 1609 Python and use the NumPy [30] library. We provide further details on the clustering experiments
 1610 below. For the cluster assignment after mean-shift, we construct a graph over shifted data points where
 1611 an edge exists between x and x' if $k_h(x, x') > 0.9$ and run the connected component algorithm.

1613 N.1 DIFFERENT MEAN-SHIFT ALGORITHMS

1615 Gradient mean-shift(GMS):

$$1617 \quad x^{(l+1)} = x^{(l)} + \eta \frac{\nabla f_{k_h}(x^{(l)})}{f_{k_h}(x^{(l)})} \quad (72)$$

1619 Fixed-point iteration mean-shift(MS):

$$x^{(l+1)} = \frac{\sum_{i=1}^n k_h(x_i - x^{(l)}) x_i}{\sum_{i=1}^n k_h(x_i - x^{(l)})} \quad (73)$$

Blurring mean-shift(BMS):

$$x^{(l+1)} = \frac{\sum_{i=1}^n k_h(x_i^{(l)} - x^{(l)}) x_i^{(l)}}{\sum_{i=1}^n k_h(x_i^{(l)} - x^{(l)})} \quad (74)$$

For MS++, we refer the readers to the original paper by Jiang et al. [38] for a more detailed discussion.

N.2 DATASET DETAILS

	a1	a2	a3	unbalance	s1	s2	s3	s4
<i>n</i>	3000	5250	7500	6500	5000	5000	5000	5000
# Cluster	20	35	50	8	15	15	15	15
<i>d</i>	2	2	2	2	2	2	2	2

	WallRobot	UserKnowledge	WirelessLocalization
<i>n</i>	5456	403	2000
# Cluster	4	4	4
<i>d</i>	4	5	7

Table 3: Dataset information for clustering experiments.

N.3 HYPERPARAMETER INFORMATION

We fix the number of iterations across all methods. Parameters such as h , μ , η are chosen empirically. For RFMS, D is chosen to be as small as possible while not degrading the result significantly.

	a1	a2	a3	unbalance	s1	s2	s3	s4
Random Feature dimension D	300	300	300	300	300	300	300	300
Bandwidth h	0.1	0.1	0.1	0.5	0.2	0.2	0.2	0.2
Iteration T	100	100	100	100	100	100	100	100
Smoothing parameter μ	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4	5e-4
Learning rate η	1e-3	1e-3	1e-3	1e-3	3e-3	3e-3	3e-3	3e-3

	WallRobot	UserKnowledge	WirelessLocalization
Random Feature dimension D	500	2000	750
Bandwidth h	0.2	0.5	0.7
Iteration T	100	100	100
Smoothing parameter μ	5e-4	5e-4	5e-4
Learning rate η	1e-2	3e-2	5e-2

Table 4: RFMS hyperparameters.

	a1	a2	a3	unbalance	s1	s2	s3	s4
Bandwidth	0.15	0.15	0.15	0.3	0.15	0.15	0.15	0.15
<hr/>								
	WallRobot	UserKnowledge	WirelessLocalization					
Bandwidth	0.2	0.5	0.5					

Table 5: MS++ hyperparameter.

For MS, GMS, and BMS, we use the same hyperparameter whenever applicable.

N.4 ADDITIONAL EXPERIMENTAL RESULTS

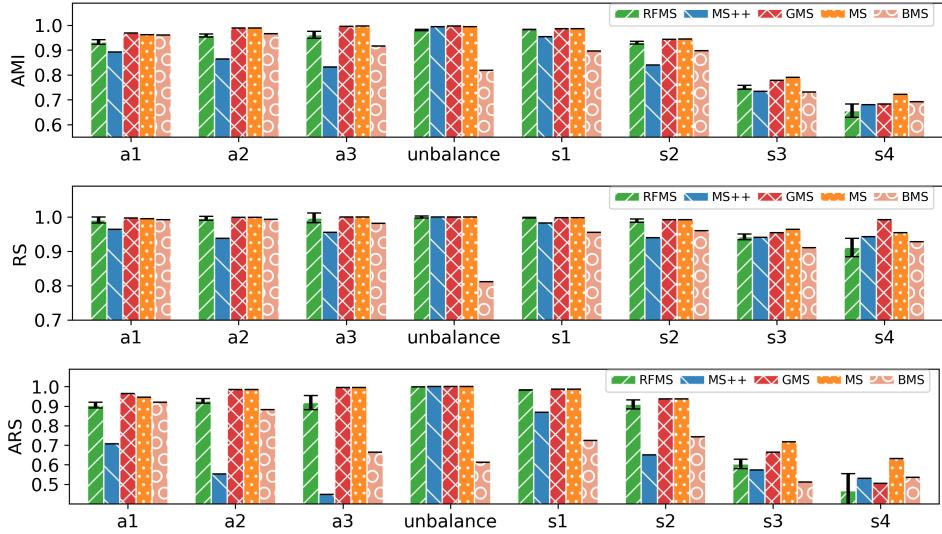


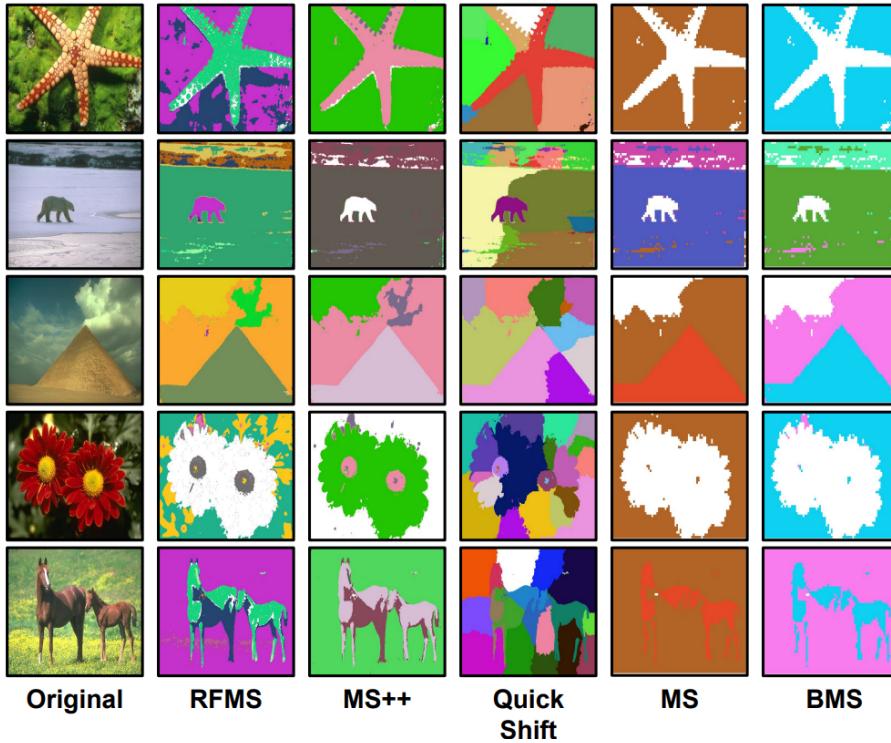
Figure 8: Adjusted mutual information(AMI), Rand Score(RS) and Adjusted Rand Score(ARS)

Metric	Method	a1	a2	a3	unbalance	s1	s2	s3	s4
AMI	RFMS	0.93 ± 0.01	0.96 ± 0.01	0.96 ± 0.01	0.98 ± 0.00	0.98 ± 0.00	0.93 ± 0.01	0.75 ± 0.01	0.66 ± 0.03
	GMS	0.97 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.94 ± 0.00	0.78 ± 0.00	0.68 ± 0.00
	MS	0.96 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.79 ± 0.00	0.72 ± 0.00
	BMS	0.96 ± 0.00	0.97 ± 0.00	0.92 ± 0.00	0.82 ± 0.00	0.90 ± 0.00	0.90 ± 0.00	0.73 ± 0.00	0.69 ± 0.00
	MS++	0.89 ± 0.00	0.86 ± 0.00	0.83 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.84 ± 0.00	0.73 ± 0.00	0.68 ± 0.00
NMI	RFMS	0.94 ± 0.01	0.96 ± 0.01	0.96 ± 0.01	0.98 ± 0.00	0.98 ± 0.00	0.93 ± 0.01	0.75 ± 0.01	0.67 ± 0.03
	GMS	0.97 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.78 ± 0.00	0.69 ± 0.00
	MS	0.96 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.79 ± 0.00	0.72 ± 0.00
	BMS	0.96 ± 0.00	0.97 ± 0.00	0.92 ± 0.00	0.82 ± 0.00	0.90 ± 0.00	0.90 ± 0.00	0.73 ± 0.00	0.69 ± 0.00
	MS++	0.89 ± 0.00	0.87 ± 0.00	0.84 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.84 ± 0.00	0.74 ± 0.00	0.69 ± 0.00
ARS	RFMS	0.91 ± 0.02	0.93 ± 0.01	0.92 ± 0.04	1.00 ± 0.00	0.98 ± 0.01	0.91 ± 0.02	0.60 ± 0.02	0.47 ± 0.09
	GMS	0.96 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.94 ± 0.00	0.67 ± 0.00	0.51 ± 0.00
	MS	0.95 ± 0.00	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.94 ± 0.00	0.71 ± 0.00	0.63 ± 0.00
	BMS	0.92 ± 0.00	0.88 ± 0.00	0.66 ± 0.00	0.61 ± 0.00	0.72 ± 0.00	0.74 ± 0.00	0.51 ± 0.00	0.54 ± 0.00
	MS++	0.71 ± 0.00	0.55 ± 0.00	0.45 ± 0.00	1.00 ± 0.00	0.87 ± 0.00	0.65 ± 0.00	0.57 ± 0.00	0.53 ± 0.00
RS	RFMS	0.99 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.94 ± 0.01	0.91 ± 0.03
	GMS	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.95 ± 0.00	0.92 ± 0.00
	MS	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	1.00 ± 0.00	0.99 ± 0.00	0.96 ± 0.00	0.95 ± 0.00
	BMS	0.99 ± 0.00	0.99 ± 0.00	0.98 ± 0.00	0.81 ± 0.00	0.96 ± 0.00	0.96 ± 0.00	0.91 ± 0.00	0.93 ± 0.00
	MS++	0.96 ± 0.00	0.96 ± 0.00	0.96 ± 0.00	1.00 ± 0.00	0.98 ± 0.00	0.94 ± 0.00	0.94 ± 0.00	0.94 ± 0.00
Time(s)	RFMS	4.44 ± 0.05	8.36 ± 0.07	12.71 ± 0.17	10.63 ± 0.12	7.80 ± 0.14	7.71 ± 0.12	8.03 ± 0.10	7.93 ± 0.15
	GMS	30.24 ± 0.54	88.36 ± 4.21	175.98 ± 8.26	136.82 ± 2.19	75.12 ± 0.29	75.42 ± 0.47	76.82 ± 4.22	79.19 ± 3.96
	MS	16.98 ± 0.24	50.31 ± 2.04	93.03 ± 1.65	68.54 ± 0.36	45.21 ± 0.93	45.63 ± 0.74	45.80 ± 1.46	44.96 ± 0.34
	BMS	17.28 ± 0.24	47.16 ± 0.29	84.12 ± 0.51	68.34 ± 0.09	42.9 ± 0.24	43.21 ± 0.05	43.56 ± 0.04	43.37 ± 0.12
	MS++	1.48 ± 0.05	2.70 ± 0.05	4.17 ± 0.12	4.10 ± 0.09	2.54 ± 0.02	2.63 ± 0.07	2.61 ± 0.07	2.61 ± 0.04

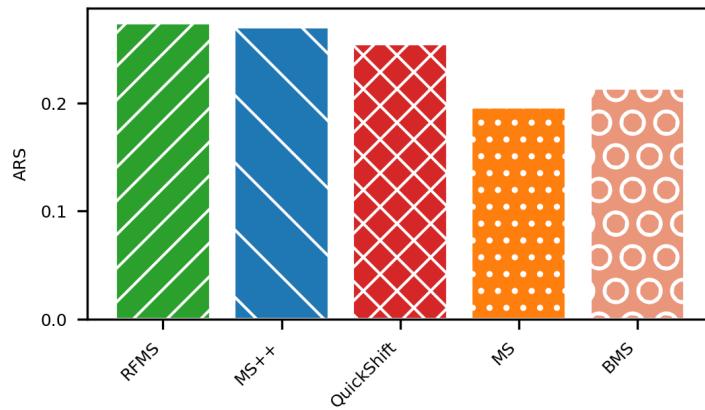
Table 6: Full clustering experimental results in table format under different metrics.

1728 **O ADDITIONAL DETAILS ON IMAGE SEGMENTATION EXPERIMENTS**
1729

1730 Segmentation experiments were run on an Intel i5-11400 CPU. RFMS, MS, BMS, and MS++ are
 1731 implemented in Python using the Numpy [30] library. QuickShift implementation comes from the
 1732 Scikit-Image library [68]. Since RFMS, MS, and BMS might produce noisy clustering, we divide
 1733 LAB color space into hypercubes of side length 100 after shifting for clustering assignment. For
 1734 RFMS, we use $D = 10$, $\mu = 0.1$ and $\eta = 30$. RFMS, MS, and BMS use $h = 25$ whereas bandwidth
 1735 for MS++ and Quickshift is 30 and 20, respectively.

1736 **O.1 EXAMPLE SEGMENTATION RESULTS**
1737

1762 Figure 9: Visualization of segmentation result using different mean-shift based algorithms on
 1763 BSDS500 dataset. Pixels belonging to the same cluster are marked with the same color.
 1764

1765 **O.2 ADDITIONAL EXPERIMENTAL RESULT**
1766

1779 Figure 10: Adjusted Rand Score(ARS)
180

FM (Fowlkes-Mallows score)				
RFMS	MS++	Quickshift	MS	BMS
0.54	0.56	0.42	0.54	0.54
ARS (Adjusted Rand Index)				
RFMS	MS++	Quickshift	MS	BMS
0.27	0.27	0.25	0.19	0.21
Time(s)				
RFMS	MS++	Quickshift	MS	BMS
5000.11	5111.94	10557.48	15703.66	15019.48

Table 7: Full image segmentation experimental results in table format under different metrics.

P REPRODUCIBILITY / CODE AVAILABILITY

We value the availability and reproducibility of our work. The code and all the hyperparameters used in the experiment section are supplied as part of the supplemental material. We will also make our code publicly available upon acceptance of the paper.

Q LLM USAGE

Large Language Models (LLMs) were used during the preparation of this paper for assistance. Usage includes grammar, phrasing correction, polishing writing, and searching for or discovering related papers. All ideation, algorithms, technical novelties, and details are done by the authors. All LLM outputs were carefully reviewed and validated before inclusion in the manuscript.