Functional Gradient Descent with Adaptive Representations

Daniel Csillag[1*] **Rodrigo Schuller**[2] **Pedro Dall'Antonia**[1]
Leonidas Guibas[3] **Luiz Velho**[2] **Tiago Novello**[2,3]
[1]FGV EMAp [2]IMPA [3]Stanford University

Abstract

Functional optimization problems are typically solved by optimizing the parameters of a fixed representation, such as a neural network, resulting in highly nonconvex losses that complicate both training and theoretical analysis. An interesting alternative is functional gradient descent (FGD), that is, gradient descent directly in function space, which benefits from strong convergence results and admits a clean theory. However, FGD is difficult to implement in practice because functional gradients are infinite-dimensional, and thus cannot be fully computed nor stored in memory. Existing implementations therefore rely on fixed approximations, which introduce approximation error. We propose a new, theoretically-grounded FGD algorithm that adapts the representation of the functional gradients over the course of optimization. By explicitly incorporating this approximation into the analysis, we establish convergence to a stationary point (for smooth losses) and to a global minimizer (under smoothness + a Polyak-Łojasiewicz-type condition) regardless of our approximations. To the best of our knowledge, this is the first implementable FGD method with such guarantees in a general setting. We demonstrate the effectiveness of our method on regression, numerical solution of PDEs, and modern computer vision. Across settings, our method consistently outperforms both FGD with fixed approximations and neural network baselines in efficiency and accuracy.

1 Introduction

Functional optimization problems are core to many areas, including machine learning, statistics, signal processing, scientific computing, and more. In these settings, the objective is to optimize a *function* over a function space to minimize some loss (e.g., mean squared error or cross-entropy).

To solve such problems, it is common to represent the target function as a neural network and optimize the loss over the network's weights through gradient descent. While flexible, this introduces some challenges: (i) the resulting optimization problem is highly nonconvex, making training costly; (ii) the representation is typically fixed *a priori*, limiting its ability to adapt during optimization; and (iii) the nonconvex nature of the parametrized loss significantly complicates theoretical analyses of convergence and training dynamics.

We instead consider gradient descent directly in function space — *functional gradient descent* (FGD) — whose dynamics are generally simpler than their parameterized counterparts and admit strong convergence guarantees. Formally, consider the optimization problem

$$\arg\min_{f \in \mathcal{H}} L(f), \tag{1}$$

*Corresponding author: Daniel Csillag (daniel.csillag@fgv.br)

Preprint.

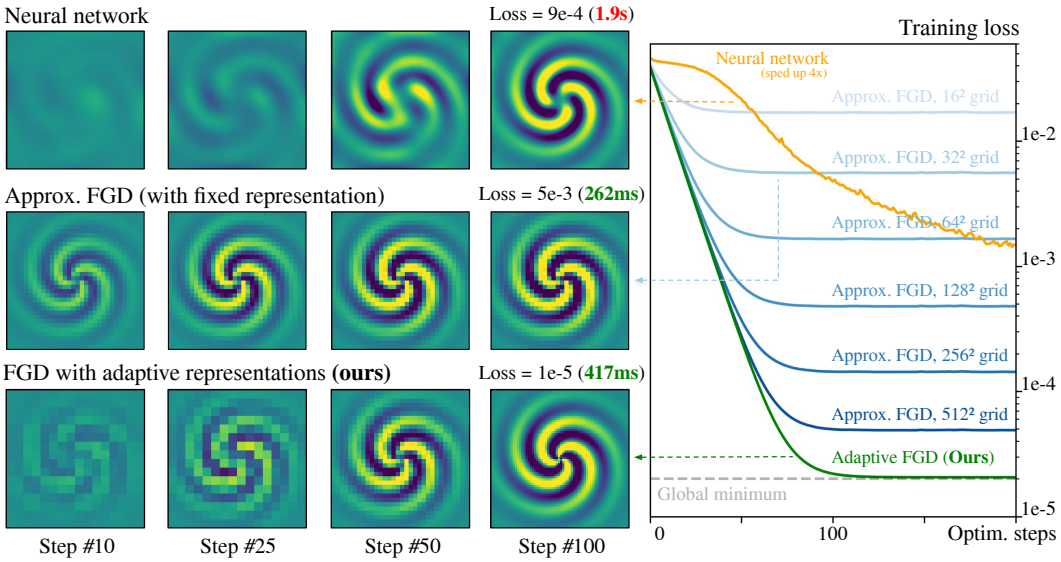


Figure 1: **Our method on a toy MSE minimization task.** We compare adaptive FGD (ours) with (i) approximate FGD using fixed representations, which converge quickly but plateau at a resolution-dependent error, and (ii) a neural network baseline trained with Adam, which converges more slowly. In contrast, our method starts from a coarse representation and automatically refines it, achieving faster convergence and reaching the global minimum.

where $\mathcal{H}$ is a Hilbert space of functions (e.g., L^2, Sobolev spaces, or an RKHS[2]). Under suitable differentiability conditions, the (non-parametric) loss admits a well-defined functional gradient $\nabla L(h) \in \mathcal{H}$, allowing us to write the FGD updates

$$f_{t+1} = f_t - \eta \nabla L(f_t). \qquad \text{(Ideal FGD)}$$

From a theoretical point of view, FGD is well known to benefit from good convergence guarantees (cf. e.g. [Bauschke and Combettes, 2017, Nocedal and Wright, 2004]). However, it is nontrivial to implement in practice. Since the functional gradients reside in the infinite-dimensional space $\mathcal{H}$, they generally cannot be computed exactly nor stored in memory. As a result, to actually implement FGD we must *approximate* the gradients $\nabla L(f_t)$ via some finite-dimensional representation. Yet using a fixed representation introduces approximation error, preventing convergence to a global minimizer.

In this paper, we propose a new theoretically-grounded functional gradient descent algorithm that adapts the representation of the functional gradients over the course of optimization. We show that by adaptively refining the gradient approximations so as to ensure a certain bound on a relative error condition, we guarantee sufficient descent and thus proper convergence. To the best of our knowledge, **ours is the first implementable FGD method that guarantees convergence to a global minimizer.** The resulting procedure strictly outperforms functional gradient descent with any fixed representation, and simultaneously outperforms neural networks in both training speed and quality; see Figure 1.

Our **main contributions** are:

- We establish sufficient conditions for the proper convergence of approximate FGD. Notably, our analysis covers the general setting where true gradients lie in a Hilbert space but are approximated within a broader Banach space, enabling important functional optimization tasks. We prove that maintaining a specific relative error bound guarantees convergence to a stationary point for smooth losses, and to a global minimizer under an additional Polyak-Łojasiewicz-type condition.
- We develop novel gradient approximation strategies that provably satisfy our derived relative error condition, using only computable quantities. This yields the first fully implementable FGD algorithms that retain strict, end-to-end convergence guarantees.

[2]Reproducing Kernel Hilbert Space

- We evaluate our method on three representative tasks across regression in an RKHS, numerical PDE solvers, and modern computer vision. On all settings, our adaptive approach consistently surpasses FGD using any fixed representation, while substantially outperforming neural network baselines in both training efficiency and solution quality.

Related work FGD is well-established as a theoretical procedure; see, e.g. [Nocedal and Wright, 2004] for an overview of the convergence theory for exact FGD. In specific cases, exact FGD is actually feasible to implement; for example, in the specific case of standard regression losses over an RKHS, the gradients are finite-dimensional and can therefore be exactly represented. This is akin to the classic representer theorems for kernel regression [Kimeldorf and Wahba, 1971, Scholkopf et al., 2001]. However, such approaches rely on specific finite-sample structures, do not generalize beyond simple regression settings, and typically scale poorly with sample size. Alternatively, several works approximate functional gradients [Fonseca and Saporito, 2022, Peixoto et al., 2024, 2025, Petrulionyte et al., 2024, Kim et al., 2025, Johnson and Zhang, 2020, Vaswani et al., 2021, D'Orazio et al., 2024], but introduce irreducible approximation error and often do not incorporate this approximation into their theoretical analysis. Our work addresses this by (i) providing a general theory for approximate FGD, and (ii) identifying sufficient conditions that ensure convergence to proper minimizers sans approximation error.

Our work is also related to first-order optimization with inexact gradients [Devolder et al., 2013, d'Aspremont, 2005, Schmidt et al., 2011], but differs in that we operate in the infinite-dimensional setting, where many previous methods break down. In particular, most approximation schemes (e.g. [Berahas et al., 2019, Nocedal and Wright, 2004, Nesterov and Spokoiny, 2015]) do not extend to infinite dimensions. Additionally, we consider a more general setting in which gradients lie in a Hilbert space but are approximated in a broader Banach space, beyond the scope of existing theory.

2 Background: function spaces and functional gradients

We briefly recall function spaces and differentiability in infinite dimensions. A Hilbert space $\mathcal{H}$ is a complete inner product space, endowed with $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ and induced norm $\|h\|_{\mathcal{H}} = \sqrt{\langle h, h \rangle_{\mathcal{H}}}$. Similarly, a Banach space $\mathcal{B}$ is a complete normed space with norm $\|\cdot\|_{\mathcal{B}}$, but generally does not admit an inner product. Every Banach space $\mathcal{B}$ admits a dual space $\mathcal{B}^*$, consisting of continuous linear functionals on $\mathcal{B}$, which is itself a Banach space under the dual norm $\|u\|_{\mathcal{B}^*} = \sup_{\|b\|_{\mathcal{B}} \leq 1} |u(b)|$.

Let $L : \mathcal{H} \to \mathbb{R}$ be a functional. The directional derivative of L at f in the direction $\vec{f}$ is defined as

$$\mathrm{D}L(f; \vec{f}) := \lim_{\delta \searrow 0} \frac{L(f + \delta \vec{f}) - L(f)}{\delta}.$$

If $\mathrm{D}L$ exists for all $f, \vec{f}$ and is linear over $\vec{f}$, L is Gâteaux differentiable; additionally, if $\mathrm{D}L$ is continuous in f, then L is Fréchet differentiable.[3]

When L is Fréchet differentiable over a Hilbert space $\mathcal{H}$, for every $f \in \mathcal{H}$ the map $\mathrm{D}L(f; \cdot)$ (i.e., the functional $\vec{f} \mapsto \mathrm{D}L(f; \vec{f})$) is a continuous linear functional. Therefore, by the Riesz representation theorem, there exists a unique element $\nabla L(f) \in \mathcal{H}$ such that

$$\mathrm{D}L(f; \vec{f}) = \langle \nabla L(f), \vec{f} \rangle_{\mathcal{H}} \qquad \text{for all } \vec{f} \in \mathcal{H}. \tag{2}$$

This element $\nabla L(f)$ is the *functional gradient*, defining the gradient operator $\nabla L : \mathcal{H} \to \mathcal{H}$. When $\mathcal{H} = \mathbb{R}^d$ with the standard inner product, this recovers the usual notion of gradients.

Finally, while Hilbert spaces provide a convenient structure for defining gradients, they can impose significant constraints on their elements. For instance, the empirical MSE functional

$$L_{\mathrm{MSE}}(f) = \frac{1}{n} \sum_{i=1}^{n} \big(f(x_i) - y_i\big)^2$$

is not well-defined on $L^2(\mathbb{R})$ (space of square-integrable functions), since functions in $L^2(\mathbb{R})$ cannot be evaluated on a set of measure zero. In contrast, it is well-defined in RKHSs, which enforce additional regularity but are more restrictive than $L^2(\mathbb{R})$.

[3] This is equivalent to the usual definition of Fréchet derivatives as bounded linear operators, [Pathak, 2018].

Algorithm 1 Functional gradient descent with adaptive representations

Require: loss $L : \mathcal{H} \to \mathbb{R}$, learning rate η, initial condition f_0, tolerance ϵ
  Choose an initial representation
  for $t = 0, 1, \ldots, T-1$ **do**
    // Approximate the functional gradient
    while true **do**
      $g_t \leftarrow$ approximate $\nabla L(f_t)$ using the current representation
      $U_t \leftarrow$ tight upper bound on $\|g_t - \nabla L(f_t)\|_{\mathcal{B}}$
      if $(1+\epsilon)U_t < \epsilon\|g_t\|_{\mathcal{B}}$ **then break**
      Refine the current representation
    // Do the gradient descent step
    $f_{t+1} \leftarrow f_t - \eta g_t$
  return f_T

Nevertheless, these issues arise due to the requirement of a proper inner product. Once we move from Hilbert to Banach spaces, we can use flexible function spaces such as L^∞ or the space of continuous functions, without caveats.

3 Method

Let $\mathcal{H}$ be a Hilbert space contained in a Banach space $\mathcal{B}$, and recall the optimization problem (1)

$$\arg\min_{f \in \mathcal{H}} L(f),$$

where $L : \mathcal{H} \to \mathbb{R}$ is taken to be a Fréchet differentiable functional on $\mathcal{H}$. Since $\mathcal{H}$ is a Hilbert space, the derivative of L admits a gradient operator $\nabla L : \mathcal{H} \to \mathcal{H}$ (cf. Section 2). An ideal functional gradient descent (FGD) scheme would, therefore, perform the steps

$$f_{t+1}^{\text{ideal}} = f_t^{\text{ideal}} - \eta \nabla L(f_t^{\text{ideal}}),$$

which remain strictly within $\mathcal{H}$ and whose convergence follows by the standard proofs of convergence of FGD. However, since this is intractable to implement, we instead replace the gradient by an *approximate gradient* g_t, yielding updates

$$f_{t+1} = f_t - \eta g_t, \tag{3}$$

where $g_t \in \mathcal{B}$ (and thus the iterates f_t also reside in $\mathcal{B}$). We will show that as long as g_t approximates $\nabla L(f_t)$ reasonably well (in a sense we will make precise), we still guarantee descent and proper convergence. Subsequently, we shall study strategies to guarantee that the required approximation conditions hold in practice by adaptively refining the gradient representations.

3.1 Convergence

In order to analyze the convergence of the updates in Equation 3, we must introduce some assumptions. First, note that loss and gradient operator are defined over the Hilbert space $\mathcal{H}$, but our updates will reside in the larger Banach space $\mathcal{B}$; therefore, in order to keep things well-defined, our first assumption is that both the loss L and the gradient operator ∇L have extensions from $\mathcal{H}$ onto $\mathcal{B}$:

Assumption 3.1 (Extension)**.** *There exists a Fréchet differentiable functional $\overline{L} : \mathcal{B} \to \mathbb{R}$ and an operator $\overline{\nabla L} : \mathcal{B} \to \mathcal{B}$ such that, for any $f \in \mathcal{H}$,*

$$\overline{L}(f) = L(f) \qquad \textit{and} \qquad \overline{\nabla L}(f) = \nabla L(f).$$

For convenience, from here on we will slightly abuse notation and "shadow" L and ∇L by their extended counterparts; i.e., for some $f \in \mathcal{B}$ we shall write $L(f)$ instead of $\overline{L}(f)$, and similarly we shall write $\nabla L(f)$ instead of $\overline{\nabla L}(f)$.

We will additionally rely on the standard assumptions used for the analysis of gradient descent, albeit over the larger Banach space $\mathcal{B}$ rather than the usual Hilbert $\mathcal{H}$. In particular, we will assume that the loss is K-smooth:

Assumption 3.2 (K-smoothness)**.** *For any $f, f' \in \mathcal{B}$, it holds that*

$$L(f) \leq L(f') + \mathrm{D}L(f'; f - f') + \frac{K}{2}\|f - f'\|_{\mathcal{B}}^2.$$

Note that we write the K-smoothness bound in terms of the directional derivative $\mathrm{D}L$ rather than an inner product with the gradient, as the Banach space does not generally have an inner product.

We will also require two additional assumptions, which serve to connect the gradients between the Banach space $\mathcal{B}$ and the Hilbert space $\mathcal{H}$:

Assumption 3.3 ($\mathcal{H}$ descends in $\mathcal{B}$)**.** *There exists some constant $\alpha > 0$ such that, for all $f \in \mathcal{B}$,*

$$\mathrm{D}L(f; \nabla L(f)) \geq \alpha \|\nabla L(f)\|_{\mathcal{B}}^2.$$

Assumption 3.4 (Gradient compatibility)**.** *There exists a constant $\beta > 0$ such that, for every $f \in \mathcal{B}$,*

$$\|\mathrm{D}L(f; \cdot)\|_{\mathcal{B}^*} \leq \beta \|\nabla L(f)\|_{\mathcal{B}}.$$

When $\mathcal{B} = \mathcal{H}$ both of these assumptions are immediately satisfied with $\alpha = \beta = 1$: $\mathrm{D}L(f; \nabla L(f)) = \langle \nabla L(f), \nabla L(f) \rangle_{\mathcal{H}} = \|\nabla L(f)\|_{\mathcal{H}}^2$, and $\|\mathrm{D}L(f; \cdot)\|_{\mathcal{B}^*} = \|\nabla L(f)\|_{\mathcal{B}}$. But these assumptions are more general, and satisfied in some important cases where $\mathcal{H} \neq \mathcal{B}$ (cf. Section 4.1). Intuitively, Assumption 3.3 asserts that the direction given by the extended gradient $\nabla L(f)$ is guaranteed to be a direction of sufficient descent per the geometry of $\mathcal{B}$. Assumption 3.4, on the other hand, establishes a sort of local equivalence of norms over the gradients (but not over the whole space).

With these in hand, we can analyze a single step of our method (Equation 3) to establish a sufficient descent lemma. The following result follows from the K-smoothness inequality (Assumption 3.2), along with our Assumptions 3.3 and 3.4 and the use of triangular inequalities:

Lemma 3.5 (Sufficient descent lemma)**.** *Let f_t, f_{t+1} be as in (3). Suppose L and ∇L admit extensions satisfying K-smoothness, gradient compatibility and that $\mathcal{H}$ descends in $\mathcal{B}$ (Assumptions 3.1-3.4). Then, for any learning rate $\eta > 0$ and step t with $\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq 1/2$,*

$$L(f_{t+1}) \leq L(f_t) - \eta \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2} K\eta \right) \frac{\mathrm{RelErr}(g_t, \nabla L(f_t))}{1 - \mathrm{RelErr}(g_t, \nabla L(f_t))} \right] \|\nabla L(f_t)\|_{\mathcal{B}}^2,$$

where

$$\mathrm{RelErr}(g, \nabla L(f)) := \frac{\|g - \nabla L(f)\|_{\mathcal{B}}}{\|g\|_{\mathcal{B}}}. \tag{4}$$

Note that this bound is tight, in the sense that when $\mathcal{H} = \mathcal{B}$, $g_t = \nabla L(f_t)$, and $\eta = 1/K$, we recover the usual descent lemma $L(f_t - \eta g_t) \leq L(f_t) - \frac{\eta}{2}\|\nabla L(f_t)\|_{\mathcal{B}}^2$ (cf. e.g. [Wright and Recht, 2022]).

As long as the term within the brackets is strictly positive, we guarantee that a step of approximate functional gradient descent decreases the loss in a manner proportional to the current gradient value. By telescoping and rearranging we obtain:

Proposition 3.6 (Convergence to a stationary point)**.** *Under the setting of Lemma 3.5, furthermore assume that the loss is lower bounded by $L^\star$, and that the relative error is uniformly bounded by a constant $\epsilon < \min\{1/2, \alpha/(\alpha + \beta)\}$ throughout iterations:*

$$\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq \epsilon \qquad \textit{for all iterations } t = 0, \ldots, T-1.$$

Then, for any learning rate $\eta \leq 2(\alpha - (\alpha + \beta)\epsilon)/K(2\epsilon + 1)$,

$$\min_{t=0,\ldots,T-1} \|\nabla L(f_t)\|_{\mathcal{B}}^2 \leq \frac{L(f_0) - L^\star}{T\eta r},$$

where $r > 0$ is a constant defined in the proof.

Under additional structural assumptions we are able to furthermore establish convergence to the proper global minimizer. For our purposes, we will work under a Polyak-Łojasiewicz-type assumption, which serves as a generalization of strong convexity:

Assumption 3.7 (Polyak-Łojasiewicz condition on $\mathcal{H}$ over $\mathcal{B}$)**.** *The extended Fréchet-differentiable functional $L : \mathcal{B} \to \mathbb{R}$ is μ-Polyak-Łojasiewicz if, for any $f \in \mathcal{B}$,*

$$L(f) - L^\star \leq \frac{1}{2\mu} \|\mathrm{D}L(f; \cdot)\|_{\mathcal{B}^*}^2, \quad \textit{where} \quad L^\star = \inf_{f' \in \mathcal{H}} L(f').$$

It then holds:

Proposition 3.8 (Global convergence)**.** *Under the conditions of Proposition 3.6, further assume that L is μ-Polyak-Łojasiewicz (Assumption 3.7). Then, for any learning rate $\eta < 2(\alpha-(\alpha+\beta)\epsilon)/K(2\epsilon+1)$, performing T iterations of approximate functional gradient descent guarantees that*

$$L(f_T) - L^\star \leq [1 - 2\eta\beta^{-2}\mu r]^T \left(L(f_0) - L^\star\right).$$

3.2 Adaptive representations of functional gradients

In the previous section we have established that as long as we can ensure an upper bound on the relative error, we guarantee proper convergence. In this section, we now consider how exactly we can approximate the functional gradient so as to guarantee this will be the case.

Recall that our ultimate objective is to find some gradient approximation g that ensures that the relative error is bounded by some fixed ϵ:

$$\mathrm{RelErr}(g, \nabla L(f)) = \frac{\|g - \nabla L(f)\|_{\mathcal{B}}}{\|g\|_{\mathcal{B}}} \leq \epsilon.$$

Note that while the denominator can generally be computed exactly (as it is the norm of our own approximation), the approximation error in the numerator can still be tricky. However, it is often possible to compute a (reasonably tight) upper bound $U(g, \nabla L(f))$ on $\|g - \nabla L(f)\|_{\mathcal{B}}$, in the sense that

$$\|g - \nabla L(f)\|_{\mathcal{B}} \leq U(g, \nabla L(f)) \quad \text{for all } g, f \in \mathcal{B}. \tag{5}$$

With such an upper bound in hand, we can guarantee:

Lemma 3.9. *Given some $f \in \mathcal{B}$ and desired tolerance $\epsilon > 0$, let $U(\cdot, \cdot)$ be an upper bound on the approximation error (as in (5)), and let $(g_1, g_2, \ldots)$ be a sequence of gradient approximations in $\mathcal{B}$ such that $U(g_n, \nabla L(f)) \to 0$. Then there exists some n such that*

$$\frac{\mathrm{RelErr}(g_n, \nabla L(f))}{1 - \mathrm{RelErr}(g_n, \nabla L(f))} \leq \epsilon. \tag{6}$$

The result in Lemma 3.9 is quite general. Importantly, it suggests a concrete algorithm: at each FGD step we design a sequence of representations, and loop over this sequence; at each iterate, approximate the gradient using the current representation and compute the upper bound on the approximation error. As soon as Equation 6 is satisfied (which must happen at some point, as per Lemma 3.9), we use the current approximation. Algorithm 1 gives the overall procedure.

Connecting with our previous convergence theory, we obtain:

Theorem 3.10. *Let f_t be the iterates produced throughout Algorithm 1, with $\epsilon \in (0, 1)$. Then:*

(i) ***Relative error bound:*** *For all iterations $t = 1, \ldots, T$,*

$$\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq \frac{\epsilon}{1+\epsilon} < 1/2;$$

(ii) ***Convergence to a stationary point:*** *Under the conditions of Proposition 3.6,*

$$\min_{t=0,\ldots,T-1} \|\nabla L(f_t)\|_{\mathcal{B}}^2 \leq \frac{L(f_0) - L^\star}{T\eta r},$$

where $r > 0$ is a constant defined in the proof; and

(iii) ***Convergence to a global minimizer:*** *Under the conditions of Proposition 3.8,*

$$L(f_T) - L^\star \leq [1 - 2\eta\beta^{-2}\mu r]^T \left(L(f_0) - L^\star\right).$$

4 Experiments

In this section we empirically assess our method on three representative functional optimization tasks: regression in an RKHS, solution of partial differential equations and inverse rendering. Throughout,

Method	MSE		Cross-Entropy	
	Test Loss	Train Time	Test Loss	Train Time
Neural network	12.84	1.02s	0.3325	1.05s
Approx. FGD (depth 2)	0.1310	257ms	0.7167	209ms
Approx. FGD (depth 4)	0.0900	232ms	0.5134	234ms
Approx. FGD (depth 8)	0.0588	274ms	0.4024	194ms
Approx. FGD (depth 12)	0.0444	762ms	0.2916	806ms
Adaptive FGD **(ours)**	**0.0378**	843ms	**0.2434**	815ms

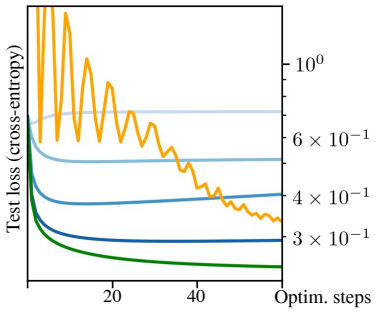


Figure 2: **Our method for regression in an RKHS.** We compare our method against FGD with a fixed tree representation and a neural network for regression. Our method quickly converges towards the global minimizer with a good test loss, whereas fixed-representation FGD stagnates at a higher loss, and the neural network is slow to converge.

we focus on comparing our method against (i) neural networks (with appropriate engineering), and (ii) FGD with fixed representations.

All experiments were run on a desktop with an Intel Core i9-14900KF CPU and NVIDIA GeForce RTX 4090 GPU, with 128GB of RAM and 24GB of VRAM, respectively. That said, the experiments are relatively cheap and should run on weaker hardware. Code to reproduce our results can be found at `https://github.com/dccsillag/experiments-adaptive-fgd`. In all experiments the learning rates were tuned to obtain the best possible performance of each method.

4.1 Regression in an RKHS

In this section we consider the optimization problem

$$\arg\min_{f \in \mathcal{H}_K} L_{\mathrm{REG}}(f), \qquad \text{with} \qquad L_{\mathrm{REG}}(f) := \frac{1}{n}\sum_{i=1}^{n} \ell\big(f(X_i), Y_i\big), \tag{7}$$

where $\mathcal{H}_K$ is an RKHS with kernel K. Here, ℓ denotes a pointwise loss such as the mean squared error ($\ell(\hat{y}, y) = \frac{1}{2}(\hat{y} - y)^2$) or binary cross-entropy ($\ell(\hat{y}, y) = -y \log \hat{y} - (1 - y)\log(1 - \hat{y})$). Such regression problems are commonplace in statistical inference procedures, including conditional independence testing [Zhang et al., 2011], asymptotic statistics [van der Vaart, 1998], and more.

The loss L_{REG} admits simple functional gradients, is K-smooth and Polyak-Łojasiewicz:

Proposition 4.1. *The functional $L_{\mathrm{REG}} : \mathcal{H}_K \to \mathbb{R}$ has gradients given by*

$$\big[\nabla L_{\mathrm{REG}}(f)\big](x) = \frac{1}{n}\sum_{i=1}^{n} \ell'(f(X_i), Y_i) K(X_i, x),$$

where ℓ' denotes the derivative of ℓ with respect to its first argument. Moreover, if ℓ is strongly-convex in its first argument, then L_{REG} is Polyak-Łojasiewicz, and if ℓ is K-smooth then so is L_{REG}.

As mentioned in Section 2, the space $\mathcal{H}_K$ is generally quite constrained. For this reason, in this experiment we opt to approximate the gradients in the larger Banach space L^∞ of bounded functions, which can be checked to satisfy our gradient bridging assumptions (Assumptions 3.3 and 3.4), cf. Proposition A.2. For our adaptive representations we then choose a decision tree-based approximation scheme inspired by the classic CART algorithm (which reside in L^∞, but generally not in $\mathcal{H}_K$). By using sufficiently deep trees we are able to approximate functions arbitrarily well, allowing us to use Lemma 3.9 and thus ensure convergence as per Theorem 3.10.

For this experiment we use the binary classification dataset of [Fan and Lin, 2011] for the detection of non-coding RNA sequences. We consider two losses ℓ: mean squared error and binary cross entropy. Figure 2 shows the results of our experiment. We find that our method consistently outperforms a standard MLP neural network as well as FGD with any fixed representation, adapting as necessary to the target function and guaranteeing convergence to the global optimum while achieving a better test loss.

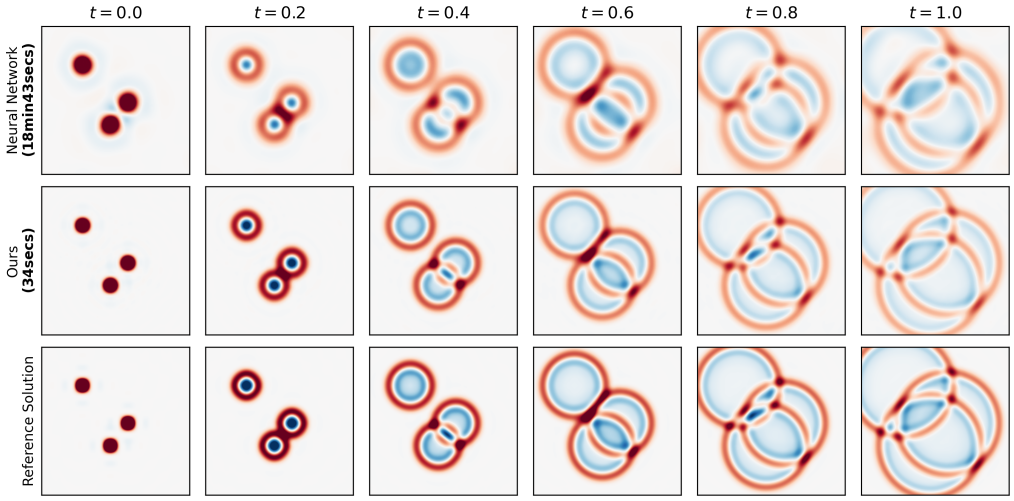


Figure 3: **Our method solves PDEs.** Comparison of learned wave dynamics over time. Our adaptive FGD solution (middle) closely matches the reference solution (bottom), while achieving higher quality than the neural network baseline (top) with a much lower training time.

4.2 Solving the wave equation

This section shows how our method can be applied to physics-informed machine learning tasks [Raissi et al., 2019, Bizzi et al., 2025]. As an example, consider the wave equation for $f : \mathbb{R} \times \mathbb{R}^2 \to \mathbb{R}$,

$$\frac{\partial^2 f}{\partial t^2} = c^2 \Delta f, \qquad \text{with} \qquad \begin{cases} f(0, \mathbf{x}) = h(\mathbf{x}), \\ \frac{\partial f}{\partial t}(0, \mathbf{x}) = 0; \end{cases} \tag{8}$$

which models wave propagation phenomena such as sound, mechanical vibrations, and electromagnetic waves [Evans, 2010, Jackson, 1999, Jensen et al., 2011, Marques et al., 2025]. Although solutions exist under broad conditions on the initial data h, explicit closed-form solutions are generally unavailable in realistic settings, requiring numerical approximation methods.

We can formulate solving Equation 8 as minimizing the following loss functional

$$L_{\text{PDE}}(f) := \int \frac{1}{2} \left[\frac{\partial^2 f}{\partial t^2} - c^2 \Delta f \right]^2 \mathrm{d}\mathbf{x}\mathrm{d}t + \int \frac{1}{2} \left[f(0, \mathbf{x}) - h(\mathbf{x}) \right]^2 \mathrm{d}\mathbf{x} + \int \frac{1}{2} \left[\frac{\partial f}{\partial t}(0, \mathbf{x}) \right]^2 \mathrm{d}\mathbf{x}, \tag{9}$$

which vanishes if and only if f solves Equation 8. The functional is defined over the Sobolev space $\mathcal{H} = H^2(\mathbb{R}^3)$ so that the second-order derivatives are well-defined.

Computing functional gradients in Sobolev spaces is generally challenging [Peixoto et al., 2025]. However, the *Fourier transform* of the gradient admits a closed-form expression:

Proposition 4.2. *Let* $\omega := \tau^2 - c^2(\xi^2 + \zeta^2)$. *The Fourier transform of the gradient of* L_{PDE} *is*

$$\widehat{[\nabla L_{\text{PDE}}(f)]}(\tau, \xi, \zeta) = \frac{\omega^2 \widehat{f}(\tau, \xi, \zeta) + \frac{1}{2\pi} \int \widehat{f}(\tau', \xi, \zeta) \mathrm{d}\tau' - \frac{1}{\sqrt{2\pi}} \widehat{h}(\xi, \zeta) + \frac{\tau}{2\pi} \left(\int \tau' \widehat{f}(\tau', \xi, \zeta) \mathrm{d}\tau' \right)}{1 + (\tau^2 + \xi^2 + \zeta^2) + (\tau^2 + \xi^2 + \zeta^2)^2}.$$

We therefore approximate gradients using adaptive uniform grids in frequency space. Each grid defines a piecewise-constant approximation of the Fourier transform $\widehat{f}$, while the corresponding function f is recovered through the inverse Fourier transform. Moreover, the Fourier representation of the $H^2(\mathbb{R}^3)$ norm yields a straightforward upper bound on the gradient approximation error.

Figure 3 shows the results. Our adaptive method consistently outperforms FGD with fixed representations, which quickly plateaus due to approximation error. It also surpasses the neural network baseline while reducing wall-clock training time by nearly two orders of magnitude.

4.3 Learning radiance fields

We next consider the inverse rendering task of reconstructing a 3D scene from posed images [Mildenhall et al., 2021]. This can be framed as a functional optimization problem: given images with known

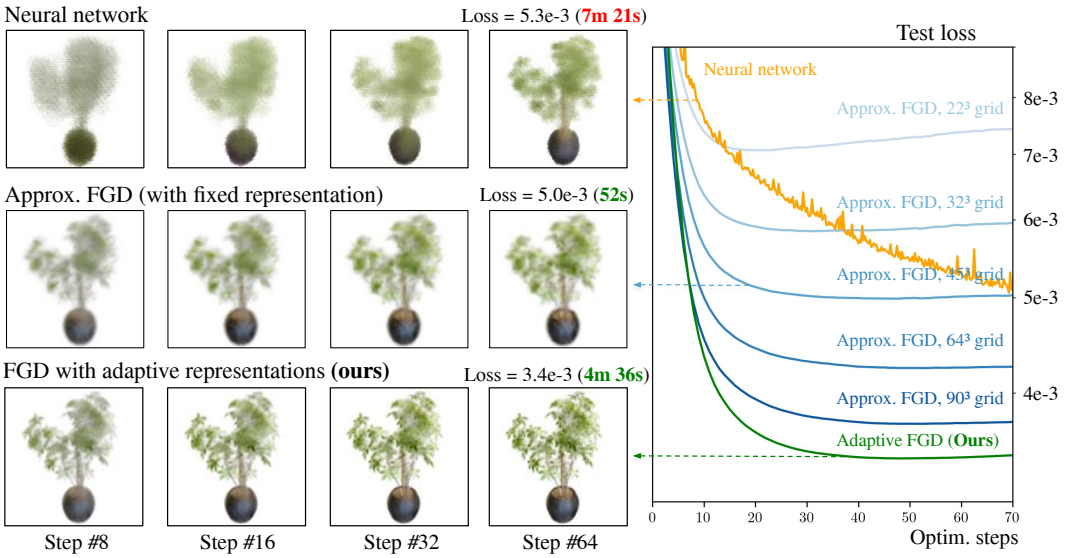


Figure 4: **Our method solves inverse rendering.** We compare novel-view renderings, from viewpoints not used during training, for an Adam-trained neural network (top), fixed-grid FGD (middle), and adaptive FGD (bottom). Adaptive FGD progressively refines its representation, achieving a lower loss and sharper reconstructions.

camera parameters, the goal is to learn the scene as a density field and a view-dependent color field,

$$\sigma : \mathbb{R}^3 \to \mathbb{R} \qquad \text{and} \qquad c : \mathbb{R}^3 \times \mathbb{S}^2 \to \mathbb{R}^3.$$

Given a camera ray $\gamma : [t_{\text{near}}, t_{\text{far}}] \to \mathbb{R}^3$ with viewing direction γ_ω, where t_{near} and t_{far} denote integration bounds, we can render the scene via the volumetric rendering equation [Max, 1995] with distributional antialiasing [Cook et al., 1984], yielding the antialiased rendering operator $\widetilde{R}_{\sigma,c}(\gamma)$, which is differentiable w.r.t. σ and c. Then, given a set of pixels $(Y_{i,j})_{i=1,j=1}^{n,P}$ and corresponding rays $(\gamma_{i,j})_{i=1,j=1}^{n,P}$ originating from the respective cameras $i = 1, \ldots, n$, our goal is to optimize the loss

$$L_{\text{CV}}(\sigma, c) := \frac{1}{n}\sum_{i=1}^{n}\frac{1}{P}\sum_{j=1}^{P}\frac{1}{2}\|\widetilde{R}_{\sigma,c}(\gamma_{i,j}) - Y_{i,j}\|^2, \tag{10}$$

which is well-defined for $\sigma \in L^2(\Omega)$ and $c \in L^2(\Omega \to \mathcal{H}^{\mathbb{S}})$, where Ω is a compact subset of $\mathbb{R}^3$ and $\mathcal{H}^{\mathbb{S}}$ is an RKHS with domain over the unit sphere $\mathbb{S}^2$. This loss directly measures image reconstruction error and is the standard supervision signal for radiance-field optimization (cf. e.g. [Mildenhall et al., 2021, Kerbl et al., 2023]). It is highly nonconvex; nonetheless, it has functional gradients known in closed form, cf. Appendix A.1.4, which we efficiently approximate by the means of uniform 3D grids with spherical harmonics for the color view-dependence.

We evaluate on the Ficus scene from the NeRF-Synthetic dataset [Mildenhall et al., 2021], using the same camera poses and train/test split for all methods. All methods optimize the photometric loss in (10); we compare adaptive FGD with an Adam-trained neural network baseline and fixed-grid FGD variants. As shown in Figure 4, adaptive FGD achieves the lowest test loss and sharpest reconstructions by starting coarse and refining only when needed, whereas fixed-grid FGD plateaus at a resolution-dependent error floor and the neural baseline converges much more slowly.

5 Conclusion and Future Directions

In this paper we presented a new framework for the implementation of FGD methods, which we have shown to exhibit a significant edge over standard methods on various tasks. We achieve this by approximating the dynamics of (vanilla) functional gradient descent over function space. The effectiveness of our approach suggests the natural target of also approximating more involved

functional procedures, such as accelerated (e.g. Nesterov), mirror and proximal variants of FGD. We expect these to work quite well, though the underlying theory may be slightly more involved.

Additionally, it is worth highlighting that efficient implementation of our method requires the design of an appropriate approximation family (as per Lemma 3.9), allowing for fast and compact fits. This is often somewhat problem-dependent. Moreover, it is also worth highlighting that approximating FGD means inheriting its inductive bias. While this appears to be often favorable, there are likely cases where we may desire a more fine-grained control over such biases. We believe that further exploration on these topics are particularly exciting directions for future research on FGD.

Acknowledgments and Disclosure of Funding

We'd like to thank Google for partially funding this work. We also acknowledge partial funding from CNPq and FAPERJ.

References

Heinz H. Bauschke and Patrick Louis Combettes. *Convex analysis and monotone operator theory in Hilbert spaces*. Springer, 2017. URL https://hal.sorbonne-universite.fr/hal-01517477. 2

Jorge Nocedal and Stephen J. Wright. Numerical optimization. *The Mathematical Gazette*, 88:412 – 413, 2004. URL https://api.semanticscholar.org/CorpusID:189864167. 2, 3

George Kimeldorf and Grace Wahba. Some results on tchebycheffian spline functions. *Journal of Mathematical Analysis and Applications*, 33:82–95, 1971. URL https://api.semanticscholar.org/CorpusID:121062339. 3

Bernhard Scholkopf, Ralf Herbrich, and Alex Smola. A generalized representer theorem. In *COLT/EuroCOLT*, 2001. URL https://api.semanticscholar.org/CorpusID:9256459. 3

Yuri Fonseca and Yuri F. Saporito. Statistical learning and inverse problems: An stochastic gradient approach. *ArXiv*, abs/2209.14967, 2022. URL https://api.semanticscholar.org/CorpusID:252595623. 3

C. Peixoto, Yuri Saporito, and Yuri Fonseca. Nonparametric instrumental variable regression through stochastic approximate gradients. *ArXiv*, abs/2402.05639, 2024. URL https://api.semanticscholar.org/CorpusID:267547764. 3

C. Peixoto, Daniel Csillag, Bernardo Freitas Paulo da Costa, and Yuri F. Saporito. Random gradient-free optimization in infinite dimensional spaces. 2025. URL https://api.semanticscholar.org/CorpusID:284133189. 3, 8

Ieva Petrulionyte, Julien Mairal, and Michael Arbel. Functional bilevel optimization for machine learning. *ArXiv*, abs/2403.20233, 2024. URL https://api.semanticscholar.org/CorpusID:268793583. 3

Kaheon Kim, Bohan Zhou, Changbo Zhu, and Xiaohui Chen. Sobolev gradient ascent for optimal transport: Barycenter optimization and convergence analysis. *ArXiv*, abs/2505.13660, 2025. URL https://api.semanticscholar.org/CorpusID:278769678. 3

Rie Johnson and Tong Zhang. Guided learning of nonconvex models through successive functional gradient optimization. *ArXiv*, abs/2006.16840, 2020. URL https://api.semanticscholar.org/CorpusID:220265986. 3

Sharan Vaswani, Olivier Bachem, Simone Totaro, Robert Müller, Shivam Garg, Matthieu Geist, Marlos C. Machado, Pablo Samuel Castro, and Nicolas Le Roux. A general class of surrogate functions for stable and efficient reinforcement learning. In *International Conference on Artificial Intelligence and Statistics*, 2021. URL https://api.semanticscholar.org/CorpusID:245853534. 3

Ryan D'Orazio, Danilo Vucetic, Zichuan Liu, Junhyung Lyle Kim, Ioannis Mitliagkas, and Gauthier Gidel. Solving hidden monotone variational inequalities with surrogate losses. *ArXiv*, abs/2411.05228, 2024. URL https://api.semanticscholar.org/CorpusID:273950322. 3

Olivier Devolder, François Glineur, and Yurii Nesterov. First-order methods of smooth convex optimization with inexact oracle. *Mathematical Programming*, 146:37 – 75, 2013. URL https://api.semanticscholar.org/CorpusID:1624627. 3

Alexandre d'Aspremont. Smooth optimization with approximate gradient. *SIAM J. Optim.*, 19: 1171–1183, 2005. URL https://api.semanticscholar.org/CorpusID:208328. 3

Mark W. Schmidt, Nicolas Le Roux, and Francis R. Bach. Convergence rates of inexact proximal-gradient methods for convex optimization. *ArXiv*, abs/1109.2415, 2011. URL https://api.semanticscholar.org/CorpusID:11262278. 3

Albert S. Berahas, Liyuan Cao, Krzysztof Choromanski, and Katya Scheinberg. A theoretical and empirical comparison of gradient approximations in derivative-free optimization. *Foundations of Computational Mathematics*, 22:507 – 560, 2019. URL https://api.semanticscholar.org/CorpusID:146120814. 3

Yurii Nesterov and Vladimir G. Spokoiny. Random gradient-free minimization of convex functions. *Foundations of Computational Mathematics*, 17:527 – 566, 2015. URL https://api.semanticscholar.org/CorpusID:2147817. 3

Hemant Kumar Pathak. An introduction to nonlinear analysis and fixed point theory. 2018. URL https://api.semanticscholar.org/CorpusID:126068909. 3

Stephen J. Wright and Benjamin Recht. Optimization for data analysis. 2022. URL https://api.semanticscholar.org/CorpusID:247919424. 5

Kun Zhang, J. Peters, Dominik Janzing, and Bernhard Scholkopf. Kernel-based conditional independence test and application in causal discovery. *ArXiv*, abs/1202.3775, 2011. URL https://api.semanticscholar.org/CorpusID:18509113. 7

A. W. van der Vaart. *Asymptotic Statistics*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1998. 7

Rong-En Fan and Chih-Jen Lin. LIBSVM data: Classification, regression, and multi-label. https://www.csie.ntu.edu.tw/~cjlin/libsvmtools/datasets/, 2011. Accessed: 2026-05-07. 7

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019. 8

Arthur Bizzi, Leonardo Moreira, Márcio Marques, Leonardo Mendonça, Christian de Oliveira, Vitor Balestro, Lucas dos Santos Fernandez, Daniel Yukimura, Pavel Petrov, João Pereira, Tiago Novello, and Lucas Nissenbaum. Neuro-spectral architectures for causal physics-informed networks. In *Advances in Neural Information Processing Systems*, 2025. 8

Lawrence C. Evans. *Partial Differential Equations*. American Mathematical Society, 2 edition, 2010. 8

John David Jackson. *Classical Electrodynamics*. Wiley, 3 edition, 1999. 8

Finn B Jensen, William A Kuperman, Michael B Porter, Henrik Schmidt, and Alexandra Tolstoy. *Computational ocean acoustics*, volume 2011. Springer, 2011. 8

Márcio Marques, Leonardo Mendonça, Arthur Bizzi, Leonardo Moreira, Christian Oliveira, Deborah Oliveira, Lucas Fernandez, Vitor Balestro, João Pereira, Daniel Yukimura, Tiago Novello, Pavel Petrov, and Lucas Nissenbaum. Stable adaptive training for physics-informed neural networks in acoustic wave propagation. *JASA Express Letters*, 5(11):112401, 2025. 8

Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021. 8, 9

Nelson Max. Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics*, 1(2):99–108, 1995. 9

Robert L Cook, Thomas Porter, and Loren Carpenter. Distributed ray tracing. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 137–145, 1984. 9

Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, George Drettakis, et al. 3d gaussian splatting for real-time radiance field rendering. *ACM Trans. Graph.*, 42(4):139–1, 2023. 9

Lawrence C Evans. *Partial differential equations*, volume 19. American mathematical society, 2022. 19

Matthew Tancik, Pratul P. Srinivasan, Ben Mildenhall, Sara Fridovich-Keil, Nithin Raghavan, Utkarsh Singhal, Ravi Ramamoorthi, Jonathan T. Barron, and Ren Ng. Fourier features let networks learn high frequency functions in low dimensional domains. *ArXiv*, abs/2006.10739, 2020. URL https://api.semanticscholar.org/CorpusID:219791950. 28

A Proofs

Proof of Lemma 3.5. Let us derive the bound in the lemma. By K-smoothness, we have that

$$L(f_t - \eta g_t) - L(f_t) \leq \mathrm{D}L(f_t; -\eta g_t) + \frac{K}{2}\|-\eta g_t\|_{\mathcal{B}}^2 = -\eta \mathrm{D}L(f_t; g_t) + \frac{K\eta^2}{2}\|g_t\|_{\mathcal{B}}^2;$$

Writing $g_t = \nabla L(f_t) - (\nabla L(f_t) - g_t) =: \nabla L(f_t) - e_t$, it follows:[4]

$$\begin{aligned}
&-\eta \mathrm{D}L(f_t; g_t) + \frac{K\eta^2}{2}\|g_t\|_{\mathcal{B}}^2 \\
&= -\eta \mathrm{D}L(f_t; \nabla L(f_t)) + \eta \mathrm{D}L(f_t; e_t) + \frac{K\eta^2}{2}\|\nabla L(f_t) - e_t\|_{\mathcal{B}}^2 \\
&\leq -\eta\alpha\|\nabla L(f_t)\|_{\mathcal{B}}^2 + \eta \mathrm{D}L(f_t; e_t) + \frac{K\eta^2}{2}\|\nabla L(f_t) - e_t\|_{\mathcal{B}}^2 && \text{(Assumption 3.3)} \\
&\leq -\eta\alpha\|\nabla L(f_t)\|_{\mathcal{B}}^2 + \eta\|\mathrm{D}L(f_t; \cdot)\|_{\mathcal{B}^*}\|e_t\|_{\mathcal{B}} + \frac{K\eta^2}{2}\|\nabla L(f_t) - e_t\|_{\mathcal{B}}^2 && \text{(Hölder)} \\
&\leq -\eta\alpha\|\nabla L(f_t)\|_{\mathcal{B}}^2 + \eta\beta\|\nabla L(f_t)\|_{\mathcal{B}}\|e_t\|_{\mathcal{B}} + \frac{K\eta^2}{2}\|\nabla L(f_t) - e_t\|_{\mathcal{B}}^2 && \text{(Assumption 3.4)} \\
&\leq -\eta\alpha\|\nabla L(f_t)\|_{\mathcal{B}}^2 + \eta\beta\|\nabla L(f_t)\|_{\mathcal{B}}\|e_t\|_{\mathcal{B}} + \frac{K\eta^2}{2}\Big(\|\nabla L(f_t)\|_{\mathcal{B}} + \|e_t\|_{\mathcal{B}}\Big)^2 && \text{(Triangle ineq.)} \\
&= -\eta\alpha\|\nabla L(f_t)\|_{\mathcal{B}}^2 + (\eta\beta + K\eta^2)\|\nabla L(f_t)\|_{\mathcal{B}}\|e_t\|_{\mathcal{B}} + \frac{K\eta^2}{2}\|\nabla L(f_t)\|_{\mathcal{B}}^2 + \frac{K\eta^2}{2}\|e_t\|_{\mathcal{B}}^2 \\
&= -\eta\left(\alpha - \frac{K\eta}{2} - (\beta + K\eta)\frac{\|e_t\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}} - \frac{K\eta}{2}\frac{\|e_t\|_{\mathcal{B}}^2}{\|\nabla L(f_t)\|_{\mathcal{B}}^2}\right)\|\nabla L(f_t)\|_{\mathcal{B}}^2;
\end{aligned}$$

to simplify the expression, let us assume that $\|\nabla L(f_t) - g_t\|_{\mathcal{B}}/\|\nabla L(f_t)\|_{\mathcal{B}} = \|e_t\|_{\mathcal{B}}/\|\nabla L(f_t)\|_{\mathcal{B}} \leq 1$. This allows us to use the fact that $x^2 \leq x$ for $|x| \leq 1$ to write

$$\begin{aligned}
&-\eta\left(\alpha - \frac{K\eta}{2} - (\beta + K\eta)\frac{\|\nabla L(f_t) - g_t\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}} - \frac{K\eta}{2}\frac{\|\nabla L(f_t) - g_t\|_{\mathcal{B}}^2}{\|\nabla L(f_t)\|_{\mathcal{B}}^2}\right)\|\nabla L(f_t)\|_{\mathcal{B}}^2; \\
&\leq -\eta\left(\alpha - \frac{K\eta}{2} - (\beta + K\eta)\frac{\|\nabla L(f_t) - g_t\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}} - \frac{K\eta}{2}\frac{\|\nabla L(f_t) - g_t\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}}\right)\|\nabla L(f_t)\|_{\mathcal{B}}^2; \\
&= -\eta\left(\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right)\frac{\|\nabla L(f_t) - g_t\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}}\right)\|\nabla L(f_t)\|_{\mathcal{B}}^2.
\end{aligned}$$

As long as the expression within the parentheses is positive, we ensure descent. However, the relative error $\|\nabla L(f_t) - g_t\|_{\mathcal{B}}/\|\nabla L(f_t)\|_{\mathcal{B}}$ is a bit tricky to manage in practice since the denominator involves the $\|\nabla L(f_t)\|_{\mathcal{B}}$, which can be hard to tightly lower bound. Fortunately, we can replace it with $\|g_t\|_{\mathcal{B}}$, which can be generally be computed exactly; indeed, by the reverse triangle inequality,[5]

$$\|g_t\|_{\mathcal{B}} - \|\nabla L(f_t)\|_{\mathcal{B}} \leq \|g_t - \nabla L(f_t)\|_{\mathcal{B}} = \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|g_t\|_{\mathcal{B}}}\|g_t\|_{\mathcal{B}};$$

rearranging, we obtain that

$$\|\nabla L(f_t)\|_{\mathcal{B}} \geq \left(1 - \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|g_t\|_{\mathcal{B}}}\right)\|g_t\|_{\mathcal{B}},$$

and thus

$$\frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|\nabla L(f_t)\|_{\mathcal{B}}} \leq \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\left(1 - \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|g_t\|_{\mathcal{B}}}\right)\|g_t\|_{\mathcal{B}}} = \frac{\mathrm{RelErr}(g_t, \nabla L(f_t))}{1 - \mathrm{RelErr}(g_t, \nabla L(f_t))},$$

where

$$\mathrm{RelErr}(g_t, \nabla L(f_t)) := \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|g_t\|_{\mathcal{B}}};$$

[4] Throughout, we follow the convention that $0/0 = 0$ and $a/0 = +\infty$ for all $a > 0$.

[5] $\big|\|u\| - \|v\|\big| \leq \|u - v\|$; as an immediate consequence, $\|u\| - \|v\| \leq \|u - v\|$.

finally, to guarantee that $\|g_t - \nabla L(f_t)\|_{\mathcal{B}} / \|\nabla L(f_t)\|_{\mathcal{B}} \leq 1$ as required, it suffices to require $\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq 1/2$, as $\frac{1}{2}/(1-\frac{1}{2}) = 1$. Putting it all together, we obtain the desired bound. □

Proof of Proposition 3.6. It follows:

$$\begin{aligned}
L(f_T) - L(f_0) &= \sum_{t=0}^{T-1} \left(L(f_{t+1}) - L(f_t)\right) \\
&\leq -\eta \sum_{t=0}^{T-1} \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\mathrm{RelErr}(g_t, \nabla L(f_t))}{1 - \mathrm{RelErr}(g_t, \nabla L(f_t))}\right] \|\nabla L(f_t)\|_{\mathcal{B}}^2 \\
&\leq -\eta \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\epsilon}{1-\epsilon}\right] \sum_{t=0}^{T-1} \|\nabla L(f_t)\|_{\mathcal{B}}^2,
\end{aligned}$$

where the first inequality holds by Lemma 3.5, and the second by the assumption that $\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq \epsilon$ for all $t = 1, \ldots, T$.

Rearranging and scaling by $1/T$ on both sides, we get

$$\eta \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\epsilon}{1-\epsilon}\right] \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla L(f_t)\|_{\mathcal{B}}^2 \leq \frac{L(f_0) - L(f_T)}{T} \leq \frac{L(f_0) - L^\star}{T};$$

dividing by $r = \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\epsilon}{1-\epsilon}\right]$ on both sides yields the desired bound; to show that $r > 0$, simply note that

$$\begin{aligned}
&\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\epsilon}{1-\epsilon} \\
&= \alpha - \frac{K}{2}\eta - \beta\frac{\epsilon}{1-\epsilon} - \frac{3}{2}\eta K \frac{\epsilon}{1-\epsilon} \\
&= -\left(\frac{3}{2}K\frac{\epsilon}{1-\epsilon} + \frac{K}{2}\right)\eta + \left(\alpha - \beta\frac{\epsilon}{1-\epsilon}\right) > 0 \\
&\iff \left(\alpha - \beta\frac{\epsilon}{1-\epsilon}\right) > \left(\frac{3}{2}K\frac{\epsilon}{1-\epsilon} + \frac{K}{2}\right)\eta \\
&\iff \eta < \frac{\alpha - \beta\frac{\epsilon}{1-\epsilon}}{\frac{3}{2}K\frac{\epsilon}{1-\epsilon} + \frac{K}{2}} = \frac{\alpha(1-\epsilon) - \beta\epsilon}{\frac{3}{2}K\epsilon + \frac{K}{2}(1-\epsilon)} = \frac{\alpha - (\alpha+\beta)\epsilon}{\frac{K}{2}(2\epsilon+1)} = \frac{2(\alpha - (\alpha+\beta)\epsilon)}{K(2\epsilon+1)},
\end{aligned}$$

which is positive thanks to the upper bound on ϵ. □

Proof of Proposition 3.8. We have that

$$\begin{aligned}
L(f_{t+1}) - L(f_t) &\leq -\eta \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\mathrm{RelErr}(g_t, \nabla L(f_t))}{1 - \mathrm{RelErr}(g_t, \nabla L(f_t))}\right] \|\nabla L(f_t)\|_{\mathcal{B}}^2 \\
&\leq -\eta \left[\alpha - \frac{K\eta}{2} - \left(\beta + \frac{3}{2}K\eta\right) \frac{\epsilon}{1-\epsilon}\right] \|\nabla L(f_t)\|_{\mathcal{B}}^2 \\
&= -\eta r \|\nabla L(f_t)\|_{\mathcal{B}}^2 \\
&\leq -\eta\beta^{-2} r \|\mathrm{D}L(f_t; \cdot)\|_{\mathcal{B}^*}^2 \\
&\leq -2\eta\beta^{-2}\mu r \left(L(f_t) - L^\star\right),
\end{aligned}$$

where the first inequality holds by Lemma 3.5, the second by the assumption that $\mathrm{RelErr}(g_t, \nabla L(f_t)) \leq \epsilon$ for all $t = 1, \ldots, T$, and the third by the Polyak-Łojasiewicz assumption. Rearranging:

$$\begin{aligned}
&L(f_{t+1}) - L(f_t) \leq -2\eta\beta^{-2}\mu r \left(L(f_t) - L^\star\right) \\
\iff &\left(L(f_{t+1}) - L(f_t)\right) + \left(L(f_t) - L^\star\right) \leq -2\eta\beta^{-2}\mu r \left(L(f_t) - L^\star\right) + \left(L(f_t) - L^\star\right)
\end{aligned}$$

$$\iff L(f_{t+1}) - L^\star \leq (1 - 2\eta\beta^{-2}\mu r)\left(L(f_t) - L^\star\right);$$

chaining these over $t = 0, \ldots, T-1$ yields

$$L(f_T) - L^\star \leq [1 - 2\eta\beta^{-2}\mu r]^T \left(L(f_0) - L^\star\right). \qquad \square$$

Proof of Lemma 3.9. First, note that by the sandwich theorem,

$$\lim_{n\to\infty} U(g_n, \nabla L(f)) = 0 \implies g_n \to \nabla L(f);$$

Then, by continuity,

$$\begin{aligned}\lim_{n\to\infty} \mathrm{RelErr}(g_n, \nabla L(f)) &= \lim_{n\to\infty} \frac{\|\nabla L(f) - g_n\|_{\mathcal{B}}}{\|g_n\|_{\mathcal{B}}} = \frac{\lim_{n\to\infty}\|\nabla L(f) - g_n\|_{\mathcal{B}}}{\lim_{n\to\infty}\|g_n\|_{\mathcal{B}}} \\ &= \frac{\|\nabla L(f) - \lim_{n\to\infty} g_n\|_{\mathcal{B}}}{\|\lim_{n\to\infty} g_n\|_{\mathcal{B}}} = \frac{\|\nabla L(f) - \nabla L(f)\|_{\mathcal{B}}}{\|\nabla L(f)\|_{\mathcal{B}}} = 0,\end{aligned}$$

and so

$$\lim_{n\to\infty} \frac{\mathrm{RelErr}(g_n, \nabla L(f))}{1 - \mathrm{RelErr}(g_n, \nabla L(f))} = \frac{\lim_{n\to\infty} \mathrm{RelErr}(g_n, \nabla L(f))}{1 - \lim_{n\to\infty} \mathrm{RelErr}(g_n, \nabla L(f))} = \frac{0}{1-0} = 0.$$

The result then follows by the definition of the limit of a sequence: for any $\epsilon > 0$, there exists some n such that the bound is satisfied for all $n' \geq n$ – which, in particular, must hold for n itself. $\square$

Proof of Theorem 3.10. Let us start by establishing claim (i). For Algorithm 1 to perform a gradient descent step, we must have that

$$(1+\epsilon)U_t < \epsilon\|g_t\|_{\mathcal{B}},$$

i.e., that

$$\frac{U_t}{\|g_t\|_{\mathcal{B}}} \leq \frac{\epsilon}{1+\epsilon};$$

and since $\|g_t - \nabla L(f_t)\|_{\mathcal{B}} \leq U_t$, we have that

$$\mathrm{RelErr}(g_t, \nabla L(f_t)) = \frac{\|g_t - \nabla L(f_t)\|_{\mathcal{B}}}{\|g_t\|_{\mathcal{B}}} \leq \frac{\epsilon}{1+\epsilon} < \frac{1}{2},$$

where the last inequality follows by the fact that $\epsilon \in (0, 1)$. Additionally, by Lemma 3.9, the while loop is guaranteed to terminate, allowing the algorithm to proceed.

Claims (ii) and (iii) then follow immediately by applying Propositions 3.6 and 3.8. $\square$

A.1 Computation of Functional Derivatives for the Experiments

A.1.1 Function fitting (Figure 1)

We consider here the loss $L_{\mathrm{FIT}} : L^2([0,1]^2) \to \mathbb{R}$ given by

$$L_{\mathrm{FIT}}(f) := \frac{1}{2}\|f - f^\star\|_{L^2}^2 = \iint \frac{1}{2}\left[f(x,y) - f^\star(x,y)\right]^2 \mathrm{d}x\mathrm{d}y.$$

It holds:

Proposition A.1. *The functional $L_{\mathrm{FIT}} : L^2([0,1]^2) \to \mathbb{R}$ has gradients given by*

$$\nabla L_{\mathrm{FIT}}(f)(x,y) = f(x,y) - f^\star(x,y),$$

and is 1-Polyak-Łojasiewicz.

Proof. The directional derivative is

$$\mathrm{D}L_{\mathrm{FIT}}(f;\vec{f}) = \left[\frac{\mathrm{d}}{\mathrm{d}\delta} L_{\mathrm{FIT}}(f + \delta\vec{f})\right]_{\delta=0}$$

$$
\begin{aligned}
&= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}\iint \frac{1}{2}\big[f(x,y)+\delta\vec{f}(x,y)-f^\star(x,y)\big]^2\mathrm{d}x\mathrm{d}y\right]_{\delta=0}\\
&= \left[\iint \frac{\mathrm{d}}{\mathrm{d}\delta}\frac{1}{2}\big[f(x,y)+\delta\vec{f}(x,y)-f^\star(x,y)\big]^2\mathrm{d}x\mathrm{d}y\right]_{\delta=0}\\
&= \iint \big[f(x,y)-f^\star(x,y)\big]\vec{f}(x,y)\mathrm{d}x\mathrm{d}y\\
&= \left\langle f-f^\star, \vec{f}\right\rangle_{L^2},
\end{aligned}
$$

and so

$$\nabla L_{\mathrm{FIT}}(f)(x,y) = f(x,y) - f^\star(x,y).$$

To show that the loss is 1-Polyak-Łojasiewicz, note:

$$L_{\mathrm{FIT}}(f) - \inf_{f'\in L^2([0,1]^2)} L_{\mathrm{FIT}}(f') = L_{\mathrm{FIT}}(f) - 0 = \frac{1}{2}\|f-f^\star\|_{L^2}^2 = \frac{1}{2}\|\nabla L_{\mathrm{FIT}}(f)\|_{L^2}^2. \qquad \square$$

A.1.2 Regression in an RKHS (Section 4.1)

Proof of Proposition 4.1. The directional derivative is

$$
\begin{aligned}
\mathrm{D}L_{\mathrm{REG}}(f;\vec{f}) &= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}L_{\mathrm{REG}}(f+\delta\vec{f})\right]_{\delta=0}\\
&= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}\frac{1}{n}\sum_{i=1}^n \ell(f(X_i)+\delta\vec{f}(X_i),Y_i)\right]_{\delta=0}\\
&= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}\frac{1}{n}\sum_{i=1}^n \ell(\langle f+\delta\vec{f},K_{X_i}\rangle_{\mathcal{H}_K},Y_i)\right]_{\delta=0}\\
&= \left[\frac{1}{n}\sum_{i=1}^n \frac{\mathrm{d}}{\mathrm{d}\delta}\ell(\langle f+\delta\vec{f},K_{X_i}\rangle_{\mathcal{H}_K},Y_i)\right]_{\delta=0}\\
&= \left[\frac{1}{n}\sum_{i=1}^n \ell'(\langle f+\delta\vec{f},K_{X_i}\rangle_{\mathcal{H}_K},Y_i)\langle\vec{f},K_{X_i}\rangle_{\mathcal{H}_K}\right]_{\delta=0}\\
&= \frac{1}{n}\sum_{i=1}^n \ell'(\langle f,K_{X_i}\rangle_{\mathcal{H}_K},Y_i)\langle\vec{f},K_{X_i}\rangle_{\mathcal{H}_K}\\
&= \frac{1}{n}\sum_{i=1}^n \ell'(f(X_i),Y_i)\langle\vec{f},K_{X_i}\rangle_{\mathcal{H}_K}\\
&= \left\langle \vec{f}, \frac{1}{n}\sum_{i=1}^n \ell'(f(X_i),Y_i)K_{X_i}\right\rangle_{\mathcal{H}_K},
\end{aligned}
$$

and thus

$$\nabla L_{\mathrm{REG}}(f)(x) = \frac{1}{n}\sum_{i=1}^n \ell'(f(X_i),Y_i)K(X_i,x).$$

To establish that the loss is Polyak-Łojasiewicz, note that the loss depends on f only through its evaluations at the data points. Let $v\in\mathbb{R}^n$ with $v_i = f(X_i)$, and let $g(v) = \frac{1}{n}\sum_{i=1}^n \ell(v_i,Y_i)$. Because ℓ is μ-strongly convex in its first argument, g is strongly convex with respect to the Euclidean norm with parameter μ/n.

Letting $c_i = \ell'(f(X_i),Y_i)$, the gradient is $\nabla g(v) = \frac{1}{n}c$. By standard strong convexity, the suboptimality is bounded by the gradient norm:

$$L_{\mathrm{REG}}(f) - \inf_{f'\in\mathcal{H}_K} L_{\mathrm{REG}}(f') \le \frac{1}{2(\mu/n)}\|\nabla g(v)\|_2^2 = \frac{n}{2\mu}\sum_{i=1}^n\left(\frac{1}{n}c_i\right)^2 = \frac{1}{2\mu n}\|c\|_2^2.$$

We must now bound $\|c\|_2^2$ using the L^∞ norm of the extended RKHS gradient. Recall that the continuous extension of the gradient to $\mathcal{B} = L^\infty$ is the function:

$$\nabla L_{\mathrm{REG}}(f)(\cdot) = \frac{1}{n}\sum_{j=1}^{n} c_j K(X_j, \cdot).$$

Let $K \in \mathbb{R}^{n\times n}$ be the empirical kernel matrix, and let $g_X \in \mathbb{R}^n$ be the evaluation of this gradient at the training points, so $g_X^{(i)} = \nabla L_{\mathrm{REG}}(f)(X_i)$. In matrix form, we have $g_X = \frac{1}{n}Kc$.

Assuming the kernel matrix K is strictly positive definite, we can invert this relationship to get $c = nK^{-1}g_X$. Bounding the Euclidean norm of c via the smallest eigenvalue $\lambda_{\min}(K) > 0$ yields:

$$\|c\|_2 \le n\|K^{-1}\|_2\|g_X\|_2 = \frac{n}{\lambda_{\min}(K)}\|g_X\|_2.$$

Next, we bound the discrete Euclidean norm $\|g_X\|_2$ by the supremum norm over the entire space. Since $|g_X^{(i)}| \le \sup_x|\nabla L_{\mathrm{REG}}(f)(x)| = \|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}$ for all i, we have:

$$\|g_X\|_2^2 = \sum_{i=1}^{n}\big(g_X^{(i)}\big)^2 \le n\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}^2 \implies \|g_X\|_2 \le \sqrt{n}\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}.$$

Substituting this back into the bound for $\|c\|_2$, we get:

$$\|c\|_2 \le \frac{n^{3/2}}{\lambda_{\min}(K)}\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty} \implies \|c\|_2^2 \le \frac{n^3}{\lambda_{\min}^2(K)}\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}^2.$$

Finally, plugging this into our suboptimality bound yields:

$$L_{\mathrm{REG}}(f) - L^\star \le \frac{1}{2\mu n}\left(\frac{n^3}{\lambda_{\min}^2(K)}\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}^2\right) = \frac{n^2}{2\mu\lambda_{\min}^2(K)}\|\nabla L_{\mathrm{REG}}(f)\|_{L^\infty}^2.$$

Finally, to show that L_{REG} is K-smooth: let $f, f' \in L^\infty$. Using the definition of L_{REG} and applying the K-smoothness of ℓ to each data point X_i, we have:

$$\begin{aligned}
L_{\mathrm{REG}}(f) &= \frac{1}{n}\sum_{i=1}^{n}\ell(f(X_i), Y_i) \\
&\le \frac{1}{n}\sum_{i=1}^{n}\left[\ell(f'(X_i), Y_i) + \ell'(f'(X_i), Y_i)\big(f(X_i) - f'(X_i)\big) + \frac{K}{2}\big(f(X_i) - f'(X_i)\big)^2\right] \\
&= \left(\frac{1}{n}\sum_{i=1}^{n}\ell(f'(X_i), Y_i)\right) + \left(\frac{1}{n}\sum_{i=1}^{n}\ell'(f'(X_i), Y_i)\big(f(X_i) - f'(X_i)\big)\right) \\
&\quad + \frac{K}{2n}\sum_{i=1}^{n}\big(f(X_i) - f'(X_i)\big)^2 \\
&= L_{\mathrm{REG}}(f') + \mathrm{D}L_{\mathrm{REG}}(f'; f - f') + \frac{K}{2n}\sum_{i=1}^{n}\big(f(X_i) - f'(X_i)\big)^2.
\end{aligned}$$

Next, we bound the quadratic term using the definition of the L^∞ norm. Since $|f(X_i) - f'(X_i)| \le \sup_x|f(x) - f'(x)| = \|f - f'\|_{L^\infty}$ for all i, it follows that:

$$\frac{K}{2n}\sum_{i=1}^{n}\big(f(X_i) - f'(X_i)\big)^2 \le \frac{K}{2n}\sum_{i=1}^{n}\|f - f'\|_{L^\infty}^2 = \frac{K}{2}\|f - f'\|_{L^\infty}^2.$$

Substituting this back into our inequality yields:

$$L_{\mathrm{REG}}(f) \le L_{\mathrm{REG}}(f') + \mathrm{D}L_{\mathrm{REG}}(f'; f - f') + \frac{K}{2}\|f - f'\|_{L^\infty}^2. \qquad \square$$

Proposition A.2. *Let L_{REG} be the regression functional over $\mathcal{H}_K$, and assume the empirical kernel matrix $K \in \mathbb{R}^{n\times n}$ over the training data is strictly positive definite with minimum eigenvalue $\lambda_{\min}(K) > 0$. Furthermore, assume the kernel is bounded, such that $\sup_{x,x'}|K(x,x')| \le \kappa$. Then, evaluating the gradients over the Banach space $\mathcal{B} = L^\infty$, the loss L_{REG} satisfies:*

(i) ***Gradient compatibility (Assumption 3.4):*** *with constant* $\beta = n\lambda_{\min}^{-1}(K)$.

(ii) $\mathcal{H}$ ***descends in*** $\mathcal{B}$ ***(Assumption 3.3):*** *with constant* $\alpha = \frac{\lambda_{\min}(K)}{n\kappa^2}$.

Proof. As established previously, let $c_i = \ell'(f(X_i), Y_i)$ for $i = 1, \ldots, n$, such that the continuous extension of the gradient to $\mathcal{B} = L^\infty$ is $\nabla L_{\text{REG}}(f)(x) = \frac{1}{n}\sum_{i=1}^n c_i K(X_i, x)$. The directional derivative for any $\vec{f} \in L^\infty$ is $\mathrm{D}L_{\text{REG}}(f; \vec{f}) = \frac{1}{n}\sum_{i=1}^n c_i \vec{f}(X_i)$.

For gradient compatibility, we must show that $\|\mathrm{D}L_{\text{REG}}(f;\cdot)\|_{(L^\infty)^*} \leq \beta\|\nabla L_{\text{REG}}(f)\|_{L^\infty}$. Recall that the dual norm of the directional derivative over L^∞ is exactly the scaled L_1 norm of the coefficients:

$$\|\mathrm{D}L_{\text{REG}}(f;\cdot)\|_{(L^\infty)^*} = \frac{1}{n}\|c\|_1.$$

Let $g_X \in \mathbb{R}^n$ be the vector of gradient evaluations at the training points, so $g_X^{(i)} = \nabla L_{\text{REG}}(f)(X_i)$. In matrix form, $g_X = \frac{1}{n}Kc$, which gives $c = nK^{-1}g_X$. By the definition of the supremum norm, the maximum absolute evaluation of the gradient over the entire domain bounds the maximum absolute evaluation at the training points: $\|g_X\|_\infty \leq \|\nabla L_{\text{REG}}(f)\|_{L^\infty}$.

It then follows:

$$\begin{aligned}
\|c\|_1 \leq \sqrt{n}\|c\|_2 &= \sqrt{n}\|nK^{-1}g_X\|_2 \\
&\leq n^{3/2}\|K^{-1}\|_2\|g_X\|_2 \\
&\leq n^{3/2}\lambda_{\min}^{-1}(K)\left(\sqrt{n}\|g_X\|_\infty\right) \\
&= n^2\lambda_{\min}^{-1}(K)\|g_X\|_\infty.
\end{aligned}$$

Substituting this into the dual norm equation and applying $\|g_X\|_\infty \leq \|\nabla L_{\text{REG}}(f)\|_{L^\infty}$ yields:

$$\|\mathrm{D}L_{\text{REG}}(f;\cdot)\|_{(L^\infty)^*} \leq \frac{1}{n}\left(n^2\lambda_{\min}^{-1}(K)\|\nabla L_{\text{REG}}(f)\|_{L^\infty}\right) = n\lambda_{\min}^{-1}(K)\|\nabla L_{\text{REG}}(f)\|_{L^\infty}.$$

Thus, gradient compatibility holds with $\beta = n\lambda_{\min}^{-1}(K)$.

To establish that $\mathcal{H}$ descends in $\mathcal{B}$, we must show $\mathrm{D}L_{\text{REG}}(f; \nabla L_{\text{REG}}(f)) \geq \alpha\|\nabla L_{\text{REG}}(f)\|_{L^\infty}^2$. Evaluating the directional derivative in the direction of the gradient yields:

$$\mathrm{D}L_{\text{REG}}(f; \nabla L_{\text{REG}}(f)) = \frac{1}{n}\sum_{i=1}^n c_i \nabla L_{\text{REG}}(f)(X_i) = \frac{1}{n}c^\top\left(\frac{1}{n}Kc\right) = \frac{1}{n^2}c^\top Kc.$$

We now require an upper bound for $\|\nabla L_{\text{REG}}(f)\|_{L^\infty}^2$ in terms of $c^\top Kc$. Using the triangle inequality and the kernel bound κ:

$$\|\nabla L_{\text{REG}}(f)\|_{L^\infty} = \sup_x \left|\frac{1}{n}\sum_{i=1}^n c_i K(X_i, x)\right| \leq \frac{1}{n}\sum_{i=1}^n |c_i| \sup_x |K(X_i, x)| \leq \frac{\kappa}{n}\|c\|_1.$$

Squaring both sides and applying the norm inequality $\|c\|_1^2 \leq n\|c\|_2^2$, we have:

$$\|\nabla L_{\text{REG}}(f)\|_{L^\infty}^2 \leq \frac{\kappa^2}{n^2}\|c\|_1^2 \leq \frac{\kappa^2}{n}\|c\|_2^2.$$

By the Rayleigh quotient for positive definite matrices, $c^\top Kc \geq \lambda_{\min}(K)\|c\|_2^2$, meaning $\|c\|_2^2 \leq \lambda_{\min}^{-1}(K)c^\top Kc$. Substituting this gives:

$$\|\nabla L_{\text{REG}}(f)\|_{L^\infty}^2 \leq \frac{\kappa^2}{n\lambda_{\min}(K)}c^\top Kc.$$

Rearranging to isolate the quadratic form $c^\top Kc$:

$$\frac{1}{n^2}c^\top Kc \geq \frac{\lambda_{\min}(K)}{n\kappa^2}\|\nabla L_{\text{REG}}(f)\|_{L^\infty}^2.$$

Since the left side is exactly $\mathrm{D}L_{\text{REG}}(f; \nabla L_{\text{REG}}(f))$, the condition is satisfied with $\alpha = \frac{\lambda_{\min}(K)}{n\kappa^2}$. □

A.1.3 Solving Partial Differential Equations (Section 4.2)

Proof of Proposition 4.2. The directional derivative is

$$\begin{aligned}
\mathrm{D}L_{\mathrm{PDE}}(f;\vec{f}) &= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}L_{\mathrm{PDE}}(f+\delta\vec{f})\right]_{\delta=0} \\
&= \left[\frac{\mathrm{d}}{\mathrm{d}\delta}\iiint \frac{1}{2}\left[\frac{\partial^2 f}{\partial t^2}+\delta\frac{\partial^2 \vec{f}}{\partial t^2}-c^2\left(\frac{\partial^2 f}{\partial x^2}+\frac{\partial^2 f}{\partial y^2}\right)-\delta c^2\left(\frac{\partial^2 \vec{f}}{\partial x^2}+\frac{\partial^2 \vec{f}}{\partial y^2}\right)\right]^2 \mathrm{d}x\mathrm{d}y\mathrm{d}t\right. \\
&\quad + \frac{\mathrm{d}}{\mathrm{d}\delta}\iint \frac{1}{2}\left[f(0,x,y)+\delta\vec{f}(0,x,y)-h(x,y)\right]^2 \mathrm{d}x\mathrm{d}y \\
&\quad \left. + \frac{\mathrm{d}}{\mathrm{d}\delta}\iint \frac{1}{2}\left[\frac{\partial f}{\partial t}(0,x,y)+\delta\frac{\partial \vec{f}}{\partial t}(0,x,y)\right]^2 \mathrm{d}x\mathrm{d}y\right]_{\delta=0} \\
&= \iiint \left[\frac{\partial^2 f}{\partial t^2}-c^2\left(\frac{\partial^2 f}{\partial x^2}+\frac{\partial^2 f}{\partial y^2}\right)\right]\left[\frac{\partial^2 \vec{f}}{\partial t^2}-c^2\left(\frac{\partial^2 \vec{f}}{\partial x^2}+\frac{\partial^2 \vec{f}}{\partial y^2}\right)\right]\mathrm{d}x\mathrm{d}y\mathrm{d}t \\
&\quad + \iint \left[f(0,x,y)-h(x,y)\right]\vec{f}(0,x,y)\mathrm{d}x\mathrm{d}y \\
&\quad + \iint \frac{\partial f}{\partial t}(0,x,y)\,\frac{\partial \vec{f}}{\partial t}(0,x,y)\mathrm{d}x\mathrm{d}y.
\end{aligned}$$

Now, we have to rewrite this as an inner product on $H^2(\mathbb{R}^3)$. Recall how the Sobolev inner product can be given in terms of the Fourier transform [Evans, 2022]:

$$\langle f,g\rangle_{H^2(\mathbb{R}^3)} = \int \left(1+\|\xi\|^2+\|\xi\|^4\right)\widehat{f}(\xi)\,\overline{\widehat{g}(\xi)}\mathrm{d}\xi.$$

Therefore, let us rewrite the directional derivative to be of this form. To start, we will make it so that all references to f are through its Fourier transform, each term at a time:

$$\begin{aligned}
&\iiint \left[\frac{\partial^2 f}{\partial t^2}-c^2\left(\frac{\partial^2 f}{\partial x^2}+\frac{\partial^2 f}{\partial y^2}\right)\right]\left[\frac{\partial^2 \vec{f}}{\partial t^2}-c^2\left(\frac{\partial^2 \vec{f}}{\partial x^2}+\frac{\partial^2 \vec{f}}{\partial y^2}\right)\right]\mathrm{d}x\mathrm{d}y\mathrm{d}t \\
&= \iiint \left[i^2\tau^2\widehat{f}(\tau,\xi,\zeta)-c^2i^2\left(\xi^2+\zeta^2\right)\widehat{f}(\tau,\xi,\zeta)\right]\overline{\left[i^2\tau^2\widehat{\vec{f}}(\tau,\xi,\zeta)-c^2i^2\left(\xi^2+\zeta^2\right)\widehat{\vec{f}}(\tau,\xi,\zeta)\right]}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta \\
&= \iiint \left[i^2\left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)\widehat{f}(\tau,\xi,\zeta)\right]\overline{\left[i^2\left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)\widehat{\vec{f}}(\tau,\xi,\zeta)\right]}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta \\
&= \iiint \left[i^2\left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)\right]\overline{\left[i^2\left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)\right]}\,\widehat{f}(\tau,\xi,\zeta)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta \\
&= \iiint \left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)\overline{\left(\tau^2-c^2\left(\xi^2+\zeta^2\right)\right)}\,\widehat{f}(\tau,\xi,\zeta)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta; \\
&= \iiint \left[\tau^2-c^2\left(\xi^2+\zeta^2\right)\right]^2\,\widehat{f}(\tau,\xi,\zeta)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta;
\end{aligned}$$

For the initial condition terms, we will use the fact that

$$\widehat{u}(\xi,\zeta) = \frac{1}{\sqrt{2\pi}}\int \widehat{u}(\tau,\xi,\zeta)\mathrm{d}\tau;$$

It then follows:

$$\begin{aligned}
&\iint \left[f(0,x,y)-h(x,y)\right]\vec{f}(0,x,y)\mathrm{d}x\mathrm{d}y \\
&= \iint \left[\frac{1}{\sqrt{2\pi}}\int \widehat{f}(\tau,\xi,\zeta)\mathrm{d}\tau-\widehat{h}(\xi,\zeta)\right]\overline{\frac{1}{\sqrt{2\pi}}\int \widehat{\vec{f}}(\tau,\xi,\zeta)\mathrm{d}\tau}\mathrm{d}\xi\mathrm{d}\zeta \\
&= \iiint \left[\frac{1}{2\pi}\int \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'-\frac{1}{\sqrt{2\pi}}\widehat{h}(\xi,\zeta)\right]\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta,
\end{aligned}$$

and

$$\iint \frac{\partial f}{\partial t}(0,x,y)\,\frac{\partial \vec{f}}{\partial t}(0,x,y)\mathrm{d}x\mathrm{d}y$$

$$= \iint \left(\frac{1}{\sqrt{2\pi}} \int i\tau \widehat{f}(\tau,\xi,\zeta)\mathrm{d}\tau\right) \overline{\left(\frac{1}{\sqrt{2\pi}} \int i\tau \widehat{\vec{f}}(\tau,\xi,\zeta)\mathrm{d}\tau\right)}\mathrm{d}\xi\mathrm{d}\zeta$$
$$= \iiint \frac{\tau}{2\pi}\left(\int \tau' \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'\right) \overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta.$$

Putting it all together, we obtain:

$$\begin{aligned}
&\langle \nabla L_{\mathrm{PDE}}(f), \vec{f}\rangle_{H^2(\mathbb{R}^3)} \\
&= \mathrm{D}L_{\mathrm{PDE}}(f;\vec{f}) \\
&= \iiint \left[\tau^2 - c^2\left(\xi^2+\zeta^2\right)\right]^2 \widehat{f}(\tau,\xi,\zeta)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta \\
&\quad + \iiint \left[\frac{1}{2\pi}\int \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau' - \frac{1}{\sqrt{2\pi}}\widehat{h}(\xi,\zeta)\right]\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta, \\
&\quad + \iiint \frac{\tau}{2\pi}\left(\int \tau' \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'\right)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta. \\
&= \iiint \left(\left[\tau^2 - c^2\left(\xi^2+\zeta^2\right)\right]^2 \widehat{f}(\tau,\xi,\zeta)\right. \\
&\quad + \left[\frac{1}{2\pi}\int \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau' - \frac{1}{\sqrt{2\pi}}\widehat{h}(\xi,\zeta)\right] \\
&\quad \left. + \frac{\tau}{2\pi}\left(\int \tau' \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'\right)\right)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta; \\
&= \iiint \left(\left[\tau^2 - c^2\left(\xi^2+\zeta^2\right)\right]^2 \widehat{f}(\tau,\xi,\zeta)\right. \\
&\quad + \frac{1}{2\pi}\int \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau' - \frac{1}{\sqrt{2\pi}}\widehat{h}(\xi,\zeta) \\
&\quad \left. + \frac{\tau}{2\pi}\left(\int \tau' \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'\right)\right)\overline{\widehat{\vec{f}}(\tau,\xi,\zeta)}\mathrm{d}\tau\mathrm{d}\xi\mathrm{d}\zeta;
\end{aligned}$$

therefore, we have

$$\begin{aligned}
&[\nabla\widehat{L_{\mathrm{PDE}}(f)}](\tau,\xi,\zeta) \\
&= \frac{\left[\tau^2 - c^2\left(\xi^2+\zeta^2\right)\right]^2 \widehat{f}(\tau,\xi,\zeta) + \frac{1}{2\pi}\int \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau' - \frac{1}{\sqrt{2\pi}}\widehat{h}(\xi,\zeta) + \frac{\tau}{2\pi}\left(\int \tau' \widehat{f}(\tau',\xi,\zeta)\mathrm{d}\tau'\right)}{1 + (\tau^2+\xi^2+\zeta^2) + (\tau^2+\xi^2+\zeta^2)^2}.
\end{aligned}$$

□

A.1.4 Learning radiance fields

Our scene is represented as a pair of functions

$$\sigma : \mathbb{R}^3 \to \mathbb{R} \qquad \text{and} \qquad c : \mathbb{R}^3 \times \mathbb{S}^2 \to \mathbb{R}^3,$$

corresponding to the density and color components, respectively; note that the color (returned in RGB) is taken to be view-dependent, while the density is view-independent.

With this functional representation of the scene, an infinitesimal ray $\gamma : [t_{\mathrm{near}}, t_{\mathrm{far}}] \to \mathbb{R}^3$ with direction $\gamma_\omega := \frac{\gamma(t_{\mathrm{far}})-\gamma(t_{\mathrm{near}})}{\|\gamma(t_{\mathrm{far}})-\gamma(t_{\mathrm{near}})\|}$ can be rendered through

$$R_{\sigma,c}(\gamma) = \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma\big(\gamma(t)\big)c\big(\gamma(t),\gamma_\omega\big)\mathrm{d}t + T_\sigma(t_{\mathrm{far}};\gamma)b,$$

where $b \in \mathbb{R}^3$ is a fixed background color and $T_\sigma(\cdot)$ is the transmittance, defined as

$$T_\sigma(t;\gamma) := \exp\left(-\int_{t_{\mathrm{near}}}^{t} \sigma\big(\gamma(s)\big)\mathrm{d}s\right).$$

To render a full image from a scene, it will furthermore be useful to refer to an *antialiased* rendering operation. Given a position in the image plane $p := (u, v) \in \mathbb{R}^2$, with corresponding infinitesimal rays given by $\gamma_p = \gamma_{u,v}$, we render by convolution with an antialiasing kernel $\widetilde{k} : \mathbb{R}^2 \to \mathbb{R}$ as

$$\widetilde{R}_{\sigma,c}(p) = \widetilde{R}_{\sigma,c}(u, v) = \iint R_{\sigma,c}(\gamma_{u',v'})\widetilde{k}(u' - u, v' - v)\mathrm{d}u'\mathrm{d}v';$$

Ideally $\widetilde{k}$ would be taken to be a sinc kernel so as to compensate for the sampling of the signal, though in practice it would be a more computable choice such as a bicubic filter. Note that this can also be seen as a continuous version of supersampling antialiasing.

With these in hand, our goal is to solve a functional optimization problem

$$\operatorname*{arg\,min}_{\sigma(\cdot),\, c(\cdot,\cdot)} L(\sigma, c) \qquad \text{with} \qquad L(\sigma, c) := \frac{1}{N}\sum_{i=1}^{n}\frac{1}{P}\sum_{j=1}^{P}\frac{1}{2}\left\|\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j}\right\|^2, \tag{11}$$

over N training cameras&images, each with P pixels; here $p_{i,j}$ corresponds to the position in the pixel plane, and $Y_{i,j} \in \mathbb{R}^3$ corresponds to the color of the j-th pixel of the i-th image. A position in pixel plane $p = (u, v)$ is mapped to the ray

$$\gamma(t) = o + tA(u/f, -v/f, -1),$$

where $A \in \mathbb{R}^{3\times3}$ is the camera transformation matrix, f is the focal length and $o \in \mathbb{R}^3$ is the camera position.

To treat this as an optimization problem in function space, we must first consider the domain of the loss functional $L(\cdot, \cdot)$, so as to ensure that functional gradients are well-defined. To this end, let us first assume that the 3D domain of the density and color functions is contained to a $\Omega \subset \mathbb{R}^3$ that is compact (i.e., bounded and closed), endowed with the Lebesgue measure over Ω. We then take the density function to reside in standard L^2 space over Ω, denoted $L^2(\Omega \to \mathbb{R})$; the color function, on the other hand, has to be taken to be on a "mixture" of being L^2 in the position $x \in \Omega \subset \mathbb{R}^3$ but following a more tame RKHS structure in the direction $\omega \in S^2$; this can be done by writing the color in curried form

$$c : \Omega \to [\mathcal{H}^{\mathbb{S}}]^3,$$

where $\mathcal{H}^{\mathbb{S}}$ is a reproducing kernel Hilbert space (RKHS) of functions from the unit sphere S^2 to $\mathbb{R}$. This is necessary so as to make the loss functional L well-defined over a Hilbert space of functions. The RKHS $\mathcal{H}^{\mathbb{S}}$ comes endowed with a reproducing kernel $K_{\mathcal{H}^{\mathbb{S}}} : \mathcal{H}^{\mathbb{S}} \times \mathcal{H}^{\mathbb{S}} \to \mathbb{R}$, defined by the reproducing property that $\langle h, K_{\mathcal{H}^{\mathbb{S}}}(\omega, \cdot)\rangle = \langle h, K_{\mathcal{H}^{\mathbb{S}}}(\cdot, \omega)\rangle = h(\omega)$ for any $h \in \mathcal{H}^{\mathbb{S}}$ and $\omega \in S^2$.

Proposition A.3. *Suppose the antialiasing kernel $\widetilde{k}(\cdot, \cdot)$ has a non-empty support, and define the antialiased color differences*

$$\widetilde{D}_{\sigma,c}(x; i, j) := \sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})\,\widetilde{k}(p(x; i) - p_{i,j}).$$

Then the loss functional $L : L^2(\Omega \to \mathbb{R}) \times L^2(\Omega \to [\mathcal{H}^{\mathbb{S}}]^3) \to \mathbb{R}$ is Fréchet differentiable, with gradients given by

$$\nabla_\sigma L(\sigma, c)(x) = \frac{1}{N}\sum_{i=1}^{N} a_i \sum_{k=1}^{3}[\widetilde{D}_{\sigma,c}(x; i, j)]_k \Big(T_\sigma(x; i)[c(x, \omega(x; i))]_k - \big(R_{\sigma,c}(x; i) - R_{\sigma,c}^{[t_{\text{near}},t]}(x; i)\big)\Big),$$

$$\nabla_{c_k} L(\sigma, c)(x, \omega) = \frac{1}{N}\sum_{i=1}^{N} a_i [\widetilde{D}_{\sigma,c}(x; i, j)]_k T_\sigma(x; i)\sigma(x) K_{\mathcal{H}^{\mathbb{S}}}(\omega(x; i), \omega),$$

for constants $a_i := f_i^2 \mathbb{1}[t(x; i) \in F]/Pt(x; i)^2|\det A_i|$.

Note that when the antialiasing kernel $\widetilde{k}$ has compact support, $\widetilde{D}_{\sigma,c}(x; i, j)$ can be efficiently computed by considering just a neighborhood around the j-th pixel.

Proof of Proposition A.3. Let us start by deriving the directional (Gâteaux) derivative. It is handy to rewrite it as a scalar derivative evaluated at zero:

$$\mathrm{D}L(\sigma, c; \vec{\sigma}, \vec{c}) = \lim_{\delta \searrow 0} \frac{L(\sigma + \delta\vec{\sigma}, c + \delta\vec{c}) - L(\sigma, c)}{\delta} = \left[\frac{\mathrm{d}}{\mathrm{d}\delta} L(\sigma + \delta\vec{\sigma}, c + \delta\vec{c})\right]_{\delta=0};$$

It then follows by straightforward computation:

$$
\begin{aligned}
&\left[\frac{\mathrm{d}}{\mathrm{d}\delta}L(\sigma+\delta\vec{\sigma},c+\delta\vec{c})\right]_{\delta=0}\\
&=\left[\frac{\mathrm{d}}{\mathrm{d}\delta}\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\frac{1}{2}\left\|\widetilde{R}_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(p_{i,j})-Y_{i,j}\right\|^2\right]_{\delta=0}\\
&=\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\left[\frac{\mathrm{d}}{\mathrm{d}\delta}\frac{1}{2}\left\|\widetilde{R}_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(p_{i,j})-Y_{i,j}\right\|^2\right]_{\delta=0}\\
&=\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\left[\frac{\mathrm{d}}{\mathrm{d}\delta}[\widetilde{R}_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(p_{i,j})]_k\right]_{\delta=0}\\
&=\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\iint\left[\frac{\mathrm{d}}{\mathrm{d}\delta}[R_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(\gamma_{u',v'})]_k\right]_{\delta=0}\widetilde{k}(u'-u,v'-v)\mathrm{d}u'\mathrm{d}v'.
\end{aligned}
$$

Now, the remaining derivative is as follows. For convenience, we let $\gamma\equiv\gamma_{u',v'}$.

$$
\begin{aligned}
&\left[\frac{\mathrm{d}}{\mathrm{d}\delta}[R_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(\gamma)]_k\right]_{\delta=0}\\
&=\left[\frac{\mathrm{d}}{\mathrm{d}\delta}\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma+\delta\vec{\sigma}}(t;\gamma)\big(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t))\big)\left\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\right\rangle\mathrm{d}t\right.\\
&\qquad\left.+\frac{\mathrm{d}}{\mathrm{d}\delta}b_kT_{\sigma+\delta\vec{\sigma}}(t_{\mathrm{far}};\gamma)\right]_{\delta=0}\\
&=\left[\frac{\mathrm{d}}{\mathrm{d}\delta}\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)T_{\vec{\sigma}}(t;\gamma)^{\delta}\big(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t))\big)\left\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\right\rangle\mathrm{d}t\right.\\
&\qquad\left.+\frac{\mathrm{d}}{\mathrm{d}\delta}b_kT_{\sigma}(t_{\mathrm{far}};\gamma)T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma)^{\delta}\right]_{\delta=0}\\
&=\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)\left[\frac{\mathrm{d}}{\mathrm{d}\delta}T_{\vec{\sigma}}(t;\gamma)^{\delta}(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t)))\left\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\right\rangle\right]_{\delta=0}\mathrm{d}t\\
&\qquad+b_kT_{\sigma}(t_{\mathrm{far}};\gamma)\left[\frac{\mathrm{d}}{\mathrm{d}\delta}T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma)^{\delta}\right]_{\delta=0}\\
&=\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)\left[\big(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t))\big)\left\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\right\rangle\log T_{\vec{\sigma}}(t;\gamma)\right.\\
&\qquad\qquad\left.+\frac{\mathrm{d}}{\mathrm{d}\delta}\big(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t))\big)\left\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\right\rangle\right]_{\delta=0}\mathrm{d}t\\
&\qquad+b_kT_{\sigma}(t_{\mathrm{far}};\gamma)\log T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma)\\
&=\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)\left[\sigma(\gamma(t))\langle[c(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\log T_{\sigma}(t;\gamma)\right.\\
&\qquad\qquad\left.+\frac{\mathrm{d}}{\mathrm{d}\delta}(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t)))\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\right]_{\delta=0}\mathrm{d}t\\
&\qquad+b_kT_{\sigma}(t_{\mathrm{far}};\gamma)\log T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma)\\
&=\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)\left[\sigma(\gamma(t))\langle[c(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\log T_{\vec{\sigma}}(t;\gamma)\right.\\
&\qquad\qquad+\big(\sigma(\gamma(t))+\delta\vec{\sigma}(\gamma(t))\big)\langle[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\\
&\qquad\qquad\left.+\langle[c(\gamma(t))]_k+\delta[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\vec{\sigma}(\gamma(t))\right]_{\delta=0}\mathrm{d}t\\
&\qquad+b_kT_{\sigma}(t_{\mathrm{far}};\gamma)\log T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma)\\
&=\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}T_{\sigma}(t;\gamma)\Big(\sigma(\gamma(t))\langle[c(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\log T_{\vec{\sigma}}(t;\gamma)\\
&\qquad\qquad+\sigma(\gamma(t))\langle[\vec{c}(\gamma(t))]_k,K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle
\end{aligned}
$$

$$
\begin{aligned}
&\qquad + \langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega, \cdot)\rangle \vec{\sigma}(\gamma(t))\Big)\mathrm{d}t + b_k T_\sigma(t_{\mathrm{far}};\gamma)\log T_{\vec{\sigma}}(t_{\mathrm{far}};\gamma) \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\Big(-\sigma(\gamma(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \int_{t_{\mathrm{near}}}^{t} \vec{\sigma}(\gamma(s))\mathrm{d}s \\
&\qquad + \sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \\
&\qquad + \langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega, \cdot)\rangle \vec{\sigma}(\gamma(t))\Big)\mathrm{d}t - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s. \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \int_{t_{\mathrm{near}}}^{t} \vec{\sigma}(\gamma(s))\mathrm{d}s\mathrm{d}t \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t
\end{aligned}
$$

Now, a slightly tricky step: in the second term, let's swap the integrals.

$$
\begin{aligned}
&\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \int_{t_{\mathrm{near}}}^{t} \vec{\sigma}(\gamma(s))\mathrm{d}s\mathrm{d}t \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}\int_{t_{\mathrm{near}}}^{t} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(s))\mathrm{d}s\mathrm{d}t \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}\int_{s}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(s))\mathrm{d}t\mathrm{d}s \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}\int_{t}^{t_{\mathrm{far}}} T_\sigma(s;\gamma)\sigma(\gamma(s))\langle [c(\gamma(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}s\mathrm{d}t \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t.
\end{aligned}
$$

And now, we can turn this all into an L^2 inner product.

$$
\begin{aligned}
&\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad - \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}\int_{t}^{t_{\mathrm{far}}} T_\sigma(s;\gamma)\sigma(\gamma(s))\langle [c(\gamma(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}s\mathrm{d}t \\
&\quad - b_k T_\sigma(t_{\mathrm{far}};\gamma)\int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} \vec{\sigma}(\gamma(s))\mathrm{d}s + \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle [\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \mathrm{d}t. \\
&= \int_{t_{\mathrm{near}}}^{t_{\mathrm{far}}}\Bigg(T_\sigma(t;\gamma)\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t)) \\
&\qquad - \int_{t}^{t_{\mathrm{far}}} T_\sigma(s;\gamma)\sigma(\gamma(s))\langle [c(\gamma(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \vec{\sigma}(\gamma(t))\mathrm{d}s
\end{aligned}
$$

$$
\begin{aligned}
&\qquad - b_k T_\sigma(t_{\text{far}};\gamma)\vec{\sigma}(\gamma(t))\Big)\mathrm{d}t \\
&\quad + \int_{t_{\text{near}}}^{t_{\text{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle[\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\mathrm{d}t \\
&= \int_{t_{\text{near}}}^{t_{\text{far}}}\Bigg(T_\sigma(t;\gamma)\langle[c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \\
&\qquad - \int_t^{t_{\text{far}}} T_\sigma(s;\gamma)\sigma(\gamma(s))\langle[c(\gamma(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\mathrm{d}s - b_k T_\sigma(t_{\text{far}};\gamma)\Bigg)\vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad + \int_{t_{\text{near}}}^{t_{\text{far}}} T_\sigma(t;\gamma)\sigma(\gamma(t))\langle[\vec{c}(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\mathrm{d}t \\
&= \int_{t_{\text{near}}}^{t_{\text{far}}}\Bigg(T_\sigma(t;\gamma)\langle[c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle \\
&\qquad - \int_t^{t_{\text{far}}} T_\sigma(s;\gamma)\sigma(\gamma(s))\langle[c(\gamma(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\mathrm{d}s - b_k T_\sigma(t_{\text{far}};\gamma)\Bigg)\vec{\sigma}(\gamma(t))\mathrm{d}t \\
&\quad + \int_{t_{\text{near}}}^{t_{\text{far}}} \langle[\vec{c}(\gamma(t))]_k, T_\sigma(t;\gamma)\sigma(\gamma(t))K_{\mathcal{H}^{\mathbb{S}}}(\gamma_\omega,\cdot)\rangle\mathrm{d}t
\end{aligned}
$$

Plugging this back into the full directional derivative:

$$
\begin{aligned}
&\mathrm{D}L(\sigma,c;\vec{\sigma},\vec{c}) \\
&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \iint\left[\frac{\mathrm{d}}{\mathrm{d}\delta}[R_{\sigma+\delta\vec{\sigma},c+\delta\vec{c}}(\gamma_{u',v'})]_k\right]_{\delta=0}\widetilde{k}(u'-u,v'-v)\mathrm{d}u'\mathrm{d}v' \\
&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \iint\Bigg(\\
&\quad \int_{t_{\text{near}}}^{t_{\text{far}}}\langle[\vec{c}(\gamma_{u',v'}(t))]_k, T_\sigma(t;\gamma_{u',v'})\sigma(\gamma_{u',v'}(t))K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\mathrm{d}t \\
&\quad + \int_{t_{\text{near}}}^{t_{\text{far}}}\Big(T_\sigma(t;\gamma_{u',v'})\langle[c(\gamma_{u',v'}(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle \\
&\quad - \int_t^{t_{\text{far}}} T_\sigma(s;\gamma_{u',v'})\sigma(\gamma_{u',v'}(s))\langle[c(\gamma_{u',v'}(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\mathrm{d}s \\
&\quad - b_k T_\sigma(t_{\text{far}};\gamma_{u',v'})\Big)\vec{\sigma}(\gamma_{u',v'}(t))\mathrm{d}t\Bigg)\widetilde{k}(u'-u,v'-v)\mathrm{d}u'\mathrm{d}v' \\
&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \iint\Bigg(\\
&\quad \int_{t_{\text{near}}}^{t_{\text{far}}}\langle[\vec{c}(\gamma_{u',v'}(t))]_k, T_\sigma(t;\gamma_{u',v'})\sigma(\gamma_{u',v'}(t))K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\mathrm{d}t\Bigg)\widetilde{k}(u'-u,v'-v)\mathrm{d}u'\mathrm{d}v' \\
&\quad + \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \\
&\qquad\qquad \cdot \iint\Bigg(\int_{t_{\text{near}}}^{t_{\text{far}}}\Big(T_\sigma(t;\gamma_{u',v'})\langle[c(\gamma_{u',v'}(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle \\
&\qquad - \int_t^{t_{\text{far}}} T_\sigma(s;\gamma_{u',v'})\sigma(\gamma_{u',v'}(s))\langle[c(\gamma_{u',v'}(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\mathrm{d}s \\
&\qquad - b_k T_\sigma(t_{\text{far}};\gamma_{u',v'})\Big)\vec{\sigma}(\gamma_{u',v'}(t))\mathrm{d}t\Bigg)\widetilde{k}(u'-u,v'-v)\mathrm{d}u'\mathrm{d}v'. \\
&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}\Big[(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k
\end{aligned}
$$

$$\cdot \iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} \langle [\vec{c}(\gamma_{u',v'}(t))]_k, T_\sigma(t;\gamma_{u',v'})\sigma(\gamma_{u',v'}(t)) K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle \mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v' \Big]$$

$$+ \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \Bigg($$

$$\iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} T_\sigma(t;\gamma_{u',v'})\vec{\sigma}(\gamma_{u',v'}(t))\langle [c(\gamma_{u',v'}(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle \mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v'$$

$$- \iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} \Big(\int_{t}^{t_{\text{far}}} T_\sigma(s;\gamma_{u',v'})\sigma(\gamma_{u',v'}(s))\langle [c(\gamma_{u',v'}(s))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle \mathrm{d}s$$

$$+ b_k T_\sigma(t_{\text{far}};\gamma_{u',v'}) \Big) \vec{\sigma}(\gamma_{u',v'}(t))\mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v' \Bigg).$$

$$= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k$$

$$\iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} \langle [\vec{c}(\gamma_{u',v'}(t))]_k, T_\sigma(t;\gamma_{u',v'})\sigma(\gamma_{u',v'}(t)) K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle \mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v'$$

$$+ \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \Bigg($$

$$\iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} T_\sigma(t;\gamma_{u',v'})\vec{\sigma}(\gamma_{u',v'}(t))\langle [c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle \mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v'$$

$$- \iint \Big(\int_{t_{\text{near}}}^{t_{\text{far}}} \big(R_{\sigma,c}(\gamma_{u',v'}) - R_{\sigma,c}(\gamma_{u',v'}^{[t_{\text{near}},t]}) \big) \vec{\sigma}(\gamma_{u',v'}(t))\mathrm{d}t \Big) \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v' \Bigg).$$

Now, to help transform this into an L^2 inner product, let us introduce a lemma.

Lemma A.4. *For the i-th camera and any $g : \mathbb{R}^3 \to \mathbb{R}$,*

$$\iint \int_{t_{\text{near}}}^{t_{\text{far}}} g(\gamma_{u',v'}(t))\mathrm{d}t\, \widetilde{k}(u'-u, v'-v)\mathrm{d}u'\mathrm{d}v' = \frac{f_i^2}{|\det A_i|}\int_{\mathbb{R}^3} \frac{\mathbb{1}[t(x;i) \in F]}{t(x;i)^2} g(x)\widetilde{k}(p(x;i) - p_{i,j})\mathrm{d}x,$$

where $F = [t_{\text{near}}, t_{\text{far}}]$, and A_i, f_i are camera parameters.

Proof (of the lemma). First write $\gamma_{u',v'}(t) = o_i + tA_i q_{u',v'}$, for $q_{u',v'} = (u'/f_i, -v'/f_i, -1)$; so we have the triple integral

$$\begin{aligned}
&\iint \int_{t_{\text{near}}}^{t_{\text{far}}} g(o_i + tA_i q_{u',v'})\widetilde{k}(u'-u, v'-v)\mathrm{d}t\mathrm{d}u'\mathrm{d}v' \\
&\quad = \iint \int_{t_{\text{near}}}^{t_{\text{far}}} g(o_i + tA_i q_{u',v'})\widetilde{k}(p(o_i + tA_i q_{u',v'}; i) - p_{i,j})\mathrm{d}t\mathrm{d}u'\mathrm{d}v', \\
&\quad = \iint \int_{t_{\text{near}}}^{t_{\text{far}}} g(\Phi(u',v',t))\widetilde{k}(p(\Phi(u',v',t); i) - p_{i,j})\mathrm{d}t\mathrm{d}u'\mathrm{d}v',
\end{aligned}$$

where $p(\cdot; i)$ is such that $p(o + tAq_{u',v'}; i) = (u', v')$, and

$$\Phi\big(u, v, t\big) = o + tAq_{u,v} = o + tA(u/f, -v/f, -1).$$

We will do a change-of-variables with Φ. Note that it is invertible and differentiable, with Jacobian matrix as follows:

$$J\Phi\big(u, v, t\big) = A\begin{bmatrix} t/f & 0 & u/f \\ 0 & -t/f & -v/f \\ 0 & 0 & -1 \end{bmatrix}, \qquad \det\big(J\Phi\big(u, v, t\big)\big) = \det(A)\cdot\frac{t^2}{f^2}.$$

Now, let $D = \mathbb{R}^2 \times [t_{\text{near}}, t_{\text{far}}]$. By integration by substitution,

$$\iint \int_{t_{\text{near}}}^{t_{\text{far}}} g(\Phi(u',v',t))\widetilde{k}(p(\Phi(u',v',t); i) - p_{i,j})\mathrm{d}t\mathrm{d}u'\mathrm{d}v'$$

$$
\begin{aligned}
&= \iiint_D g(\Phi(u', v', t))\widetilde{k}(p(\Phi(u', v', t); i) - p_{i,j})\mathrm{d}t\mathrm{d}u'\mathrm{d}v' \\
&= \iiint_{\Phi^{-1}(\Phi(D))} g(\Phi(u', v', t))\widetilde{k}(p(\Phi(u', v', t); i) - p_{i,j})\mathrm{d}t\mathrm{d}u'\mathrm{d}v' \\
&= \iiint_{\Phi(D)} g(\Phi(\Phi^{-1}(x)))\widetilde{k}(p(\Phi(\Phi^{-1}(x)); i) - p_{i,j}) \cdot |\det J\Phi^{-1}(x)|\mathrm{d}x \\
&= \iiint_{\Phi(D)} g(x)\widetilde{k}(p(x; i) - p_{i,j}) \cdot |\det J\Phi^{-1}(x)|\mathrm{d}x
\end{aligned}
$$

and by the inverse function theorem,

$$
\begin{aligned}
&\iiint_{\Phi(D)} g(x)\widetilde{k}(p(x; i) - p_{i,j}) \cdot |\det J\Phi^{-1}(x)|\mathrm{d}x \\
&= \iiint_{\Phi(D)} g(x)\widetilde{k}(p(x; i) - p_{i,j}) \cdot |\det J\Phi(\Phi^{-1}(x))|^{-1}\mathrm{d}x \\
&= \iiint_{\Phi(D)} g(x)\widetilde{k}(p(x; i) - p_{i,j}) \cdot |\det A|^{-1}\frac{f^2}{\Phi_t^{-1}(x)^2}\mathrm{d}x \\
&= \frac{f^2}{|\det A|}\iiint_{\Phi(D)} \frac{1}{\Phi_t^{-1}(x)^2}g(x)\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x \\
&= \frac{f^2}{|\det A|}\iiint_{\Phi(D)} \frac{1}{t(x; i)^2}g(x)\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x \\
&= \frac{f^2}{|\det A|}\int_{\mathbb{R}^3} \frac{\mathbb{1}[t(x; i) \in F]}{t(x; i)^2}g(x)\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x. \qquad \square
\end{aligned}
$$

So, applying the lemma, we have that

$$
\begin{aligned}
&\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}\Big[(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k \\
&\quad \iint\Big(\int_{t_{\text{near}}}^{t_{\text{far}}}\langle[\vec{c}(\gamma_{u',v'}(t))]_k, T_\sigma(t; \gamma_{u',v'})\sigma(\gamma_{u',v'}(t))K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle\mathrm{d}t\Big)\widetilde{k}(u' - u, v' - v)\mathrm{d}u'\mathrm{d}v'\Big] \\
&\quad + \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k\Bigg(\\
&\quad \iint\Big(\int_{t_{\text{near}}}^{t_{\text{far}}}T_\sigma(t; \gamma_{u',v'})\vec{\sigma}(\gamma_{u',v'}(t))\langle[c(\gamma(t))]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle\mathrm{d}t\Big)\widetilde{k}(u' - u, v' - v)\mathrm{d}u'\mathrm{d}v' \\
&\quad - \iint\Big(\int_{t_{\text{near}}}^{t_{\text{far}}}\big(R_{\sigma,c}(\gamma_{u',v'}) - R_{\sigma,c}(\gamma_{u',v'}^{[t_{\text{near}},t]})\big)\vec{\sigma}(\gamma_{u',v'}(t))\mathrm{d}t\Big)\widetilde{k}(u' - u, v' - v)\mathrm{d}u'\mathrm{d}v'\Big) \\
&= \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}\Big[(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k\frac{f_i^2}{|\det A_i|} \\
&\quad \int_{\mathbb{R}^3}\frac{\mathbb{1}[t(x; i) \in F]}{t(x; i)^2}\langle[\vec{c}(x)]_k, T_\sigma(x; i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x\Big] \\
&\quad + \frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k\Bigg(\\
&\quad \frac{f_i^2}{|\det A_i|}\int_{\mathbb{R}^3}\frac{\mathbb{1}[t(x; i) \in F]}{t(x; i)^2}T_\sigma(x; i)\vec{\sigma}(x)\langle[c(x)]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega, \cdot)\rangle\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x \\
&\quad - \frac{f_i^2}{|\det A_i|}\int_{\mathbb{R}^3}\frac{\mathbb{1}[t(x; i) \in F]}{t(x; i)^2}\big(R_{\sigma,c}(\gamma_{p(x;i)}) - R_{\sigma,c}(\gamma_{p(x;i)}^{[t_{\text{near}},t]})\big)\vec{\sigma}(x)\widetilde{k}(p(x; i) - p_{i,j})\mathrm{d}x\Bigg) \\
&= \int_{\mathbb{R}^3}\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}\Big[(\widetilde{R}_{\sigma,c}(p_{i,j}) - Y_{i,j})_k
\end{aligned}
$$

$$
\begin{aligned}
&\frac{f_i^2}{|\det A_i|}\frac{\mathbb{1}[t(x;i)\in F]}{t(x;i)^2}\langle[\vec{c}(x)]_k, T_\sigma(x;i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\widetilde{k}(p(x;i)-p_{i,j})\mathrm{d}x\Big] \\
&+\int_{\mathbb{R}^3}\frac{1}{N}\sum_{i=1}^{N}\frac{1}{P}\sum_{j=1}^{P}\sum_{k=1}^{3}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\Bigg(\\
&\frac{f_i^2}{|\det A_i|}\frac{\mathbb{1}[t(x;i)\in F]}{t(x;i)^2}T_\sigma(x;i)\vec{\sigma}(x)\langle[c(x)]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\widetilde{k}(p(x;i)-p_{i,j}) \\
&-\frac{f_i^2}{|\det A_i|}\frac{\mathbb{1}[t(x;i)\in F]}{t(x;i)^2}\big(R_{\sigma,c}(\gamma_{p(x;i)})-R_{\sigma,c}(\gamma_{p(x;i)}^{[t_{\mathrm{near}},t]})\big)\vec{\sigma}(x)\widetilde{k}(p(x;i)-p_{i,j})\Bigg)\mathrm{d}x \\
=&\int_{\mathbb{R}^3}\sum_{k=1}^{3}\frac{1}{N}\sum_{i=1}^{N}\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}\left[\langle[\vec{c}(x)]_k, T_\sigma(x;i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle\right. \\
&\left.\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\,\widetilde{k}(p(x;i)-p_{i,j})\mathrm{d}x\right] \\
&+\int_{\mathbb{R}^3}\sum_{k=1}^{3}\frac{1}{N}\sum_{i=1}^{N}\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}\Bigg(T_\sigma(x;i)\vec{\sigma}(x)\langle[c(x)]_k, K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot)\rangle \\
&-\big(R_{\sigma,c}(\gamma_{p(x;i)})-R_{\sigma,c}(\gamma_{p(x;i)}^{[t_{\mathrm{near}},t]})\big)\vec{\sigma}(x)\Bigg)\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\widetilde{k}(p(x;i)-p_{i,j})\mathrm{d}x \\
=&\int_{\mathbb{R}^3}\sum_{k=1}^{3}\Bigg\langle[\vec{c}(x)]_k,\frac{1}{N}\sum_{i=1}^{N}\Bigg[\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}T_\sigma(x;i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}((\gamma_{u',v'})_\omega,\cdot) \\
&\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\,\widetilde{k}(p(x;i)-p_{i,j})\Bigg]\Bigg\rangle\mathrm{d}x \\
&+\int_{\mathbb{R}^3}\vec{\sigma}(x)\sum_{k=1}^{3}\frac{1}{N}\sum_{i=1}^{N}\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}\Bigg(T_\sigma(x;i)[c(x,(\gamma_{u',v'})_\omega)]_k \\
&-\big(R_{\sigma,c}(\gamma_{p(x;i)})-R_{\sigma,c}(\gamma_{p(x;i)}^{[t_{\mathrm{near}},t]})\big)\Bigg)\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\widetilde{k}(p(x;i)-p_{i,j})\mathrm{d}x.
\end{aligned}
$$

And thus, the functional gradients are given by:

$$
\begin{aligned}
\nabla_\sigma L(\sigma,c)(x)=&\sum_{k=1}^{3}\frac{1}{N}\sum_{i=1}^{N}\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}\Bigg(T_\sigma(x;i)[c(x,\omega(x;i))]_k \\
&-\big(R_{\sigma,c}(x;i)-R_{\sigma,c}^{[t_{\mathrm{near}},t]}(x;i)\big)\Bigg)\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\,\widetilde{k}(p(x;i)-p_{i,j}), \\
\nabla_{c_k} L(\sigma,c)(x,\omega)=&\frac{1}{N}\sum_{i=1}^{N}\Bigg[\frac{f_i^2\mathbb{1}[t(x;i)\in F]}{Pt(x;i)^2|\det A_i|}T_\sigma(x;i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}(\omega(x;i),\omega) \\
&\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})_k\,\widetilde{k}(p(x;i)-p_{i,j})\Bigg],
\end{aligned}
$$

and using $\widetilde{D}_{\sigma,c}(x;i,j)=\sum_{j=1}^{P}(\widetilde{R}_{\sigma,c}(p_{i,j})-Y_{i,j})\,\widetilde{k}(p(x;i)-p_{i,j})$ and $a_i=(f_i^2\mathbb{1}[t(x;i)\in F])/(Pt(x;i)^2|\det A_i|)$ we obtain

$$
\begin{aligned}
\nabla_\sigma L(\sigma,c)(x)=&\frac{1}{N}\sum_{i=1}^{N}a_i\sum_{k=1}^{3}[\widetilde{D}_{\sigma,c}(x;i,j)]_k\Big(T_\sigma(x;i)[c(x,\omega(x;i))]_k-\big(R_{\sigma,c}(x;i)-R_{\sigma,c}^{[t_{\mathrm{near}},t]}(x;i)\big)\Big), \\
\nabla_{c_k} L(\sigma,c)(x,\omega)=&\frac{1}{N}\sum_{i=1}^{N}a_i[\widetilde{D}_{\sigma,c}(x;i,j)]_k T_\sigma(x;i)\sigma(x)K_{\mathcal{H}^{\mathbb{S}}}(\omega(x;i),\omega).
\end{aligned}
$$

This establishes that L is (linearly) Gâteaux differentiable. It is also immediately seen to be Fréchet differentiable, as the gradients are continuous with relation to σ and c. □

B Experiment details

B.1 Regression in an RKHS

For the experiments we use an RBF kernel $K(x, y) = \exp(-100\|x - y\|^2)$. We use a simple tree-based approximation scheme which naively partitions splits at the midpoints. This is mainly for ease of (efficient) implementation in JAX, and in no way fundamental to our approach. To upper bound the approximation error, we use the fact that our trees are piecewise constant over rectangular partitions of the space, and bound the sup-norm of each leaf via a Lipschitz bound on the gradient (using the fact that the RBF kernel is Lipschitz). The neural network is a standard MLP with two hidden layers of size 256; we note that modifying this doesn't change the conclusions in any significant manner.

B.2 Solving the wave equation

The neural network baseline has a Fourier feature embedding layer [Tancik et al., 2020], which is essential for the neural network to even be capable of optimizing this loss; without this embedding, no matter how large (or small) the network or optimizer hyperparameters, it stays stuck and does not fit the solution at all. As the loss has an integral in its definition, the neural network optimizes it on batches of 4096 points uniformly sampled over the domain; lowering this batch size leads to worse results. For our method, the inverse Fourier transform can be directly computed as a sum of sincs, using the fact that a box function in frequency space corresponds to a sinc in the primary space. Finally, we bound the approximation error via a combination of a Lipschitz argument (for the interior of the frequency-space grid) and a bound on the Sobolev symbol for the content outside of the grid. For the reference solution in the figure, we solve the equation using a finite differences method with an extremely fine grid.

B.3 Learning radiance fields

We train our method on 24 160×160 RGB images, and evaluate on 25 160×160 RGB images of novel views on the test subset. Both Approx. FGD (all resolutions) and Adaptive FGD used identical optimization hyperparameters. They employed a learning rate of 8 for the color function c and 20 for the density function σ, approximated the gradient by taking 8 samples per voxel and averaging them for the density, and fitting spherical harmonics coefficients through one step of Newton's method. On the other hand, the NN used Adam optimizer with learning rate of $1e-4$ and $\beta = (0.9, 0.999)$. Since the NN cannot perform full-batch gradient descent due to VRAM limitations, we used mini-batches of 80×80 rays and counted 4×24 mini-batches as an optimization step (to match the full-batch updates of our method). Note that, since the NN traces one ray per pixel, 4×24 mini-batches correspond exactly to the number of rays in the full training data.

NeurIPS Paper Checklist

1. **Claims**

   Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

   Answer: [Yes]

   Justification: They do.

   Guidelines:

   - The answer NA means that the abstract and introduction do not include the claims made in the paper.
   - The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
   - The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
   - It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. **Limitations**

   Question: Does the paper discuss the limitations of the work performed by the authors?

   Answer: [Yes]

   Justification: All required assumptions are clearly stated.

   Guidelines:

   - The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
   - The authors are encouraged to create a separate "Limitations" section in their paper.
   - The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
   - The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
   - The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
   - The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
   - If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
   - While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. **Theory assumptions and proofs**

   Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

   Answer: [Yes]

Justification: All assumptions and proofs are provided.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. **Experimental result reproducibility**

   Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

   Answer: [Yes]

   Justification: The main text has a number of details, and the supplementary material has additional ones.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
   - If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
   - Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general. releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
   - While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
     (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
     (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
     (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
     (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. **Open access to data and code**

   Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [Yes]

Justification: Code is provided.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so "No" is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (https://nips.cc/public/guides/CodeSubmissionPolicy) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. **Experimental setting/details**

   Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

   Answer: [Yes]

   Justification: All relevant details are presented.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
   - The full details can be provided either with the code, in appendix, or as supplemental material.

7. **Experiment statistical significance**

   Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

   Answer: [No]

   Justification: Our experiments do not have significant stochasticity, making error bars redundant. The only slightly more stochastic component – runtimes – are either for long-running processes (which are not signiciantly random) or already averaged over multiple executions.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
   - The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
   - The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)

- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. **Experiments compute resources**

   Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

   Answer: [Yes]

   Justification: Hardware details are provided.

   Guidelines:

   - The answer NA means that the paper does not include experiments.
   - The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
   - The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
   - The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. **Code of ethics**

   Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics https://neurips.cc/public/EthicsGuidelines?

   Answer: [Yes]

   Justification: No ethics code infringement.

   Guidelines:

   - The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
   - If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
   - The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. **Broader impacts**

    Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

    Answer: [NA]

    Justification: No significant immediate societal impact.

    Guidelines:

    - The answer NA means that there is no societal impact of the work performed.
    - If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
    - Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. **Safeguards**

    Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

    Answer: [NA]

    Justification: No need of safeguards.

    Guidelines:

    - The answer NA means that the paper poses no such risks.
    - Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
    - Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
    - We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. **Licenses for existing assets**

    Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

    Answer: [Yes]

    Justification: All used datasets are properly cited.

    Guidelines:

    - The answer NA means that the paper does not use existing assets.
    - The authors should cite the original paper that produced the code package or dataset.
    - The authors should state which version of the asset is used and, if possible, include a URL.
    - The name of the license (e.g., CC-BY 4.0) should be included for each asset.
    - For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
    - If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, `paperswithcode.com/datasets` has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
    - For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. **New assets**

    Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

    Answer: [NA]

    Justification: No release of new assets.

    Guidelines:

    - The answer NA means that the paper does not release new assets.
    - Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
    - The paper should discuss whether and how consent was obtained from people whose asset is used.
    - At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

    Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

    Answer: [NA]

    Justification: No crowdsourcing was done.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
    - According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

    Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

    Answer: [NA]

    Justification: No crowdsourcing was done nor research with human subjects.

    Guidelines:

    - The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
    - Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
    - We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
    - For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: Not significant.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (https://neurips.cc/Conferences/2025/LLM) for what should or should not be described.