

Semester Thesis

BEM Benchmark of the Ehrenpreis–Palamodov Gaussian Process for Maxwell Cavity Scattering

Abstract

This thesis presents a high-fidelity benchmark for the Ehrenpreis–Palamodov Gaussian Process (EPGP), a probabilistic solver whose plane-wave prior satisfies the time-harmonic Maxwell equations exactly by construction. The benchmark is an interior electromagnetic scattering problem in a perfectly electrically conducting cavity, where dipole sources on an interior sphere excite a field that is measured back on the same sphere, defining a reaction operator from dipole excitations to field responses. We reconstruct this operator with the EPGP and compare it against an independent boundary element method (BEM) reference using Bembel (Dölz et al. 2020). We investigate two cavity geometries. On the spherical cavity, where the reaction operator is available in closed form, both solvers match it to a relative Frobenius-norm error between 10^{-10} and 10^{-12} . On the ellipsoidal cavity, which admits no analytic solution, the EPGP agrees with the BEM reference to a relative Frobenius-norm error of about 10^{-8} . This establishes the EPGP as an accurate, Maxwell-consistent solver for the interior cavity problem and certifies the reaction operator as a robust reference benchmark. We further study convergence, the accuracy–runtime trade-off, and uncertainty quantification.

Luis Wirth

luwirth@ethz.ch

ethz.lwirth.com

Supervised by

Prof. Dr.-Ing. Stefan Kurz

Seminar for Applied Mathematics

1st July 2026

Table of Contents

Title	I
Table of Contents	II
1 Introduction	1
1.1 Context and Motivation	1
1.2 Cavity Scattering Benchmark	2
1.3 Contributions	2
2 Cavity Scattering	3
2.1 Time-Harmonic Maxwell's Equations	3
2.2 Scattering Problem	4
2.2.1 Transmitter Dipole and Incident Field	4
2.2.2 Scattered Field and curl-curl BVP	5
2.2.3 Receiver Dipole and Measurement	5
2.2.4 Reaction Operator	5
2.3 Geometry	7
2.3.1 Ellipsoidal Cavity	8
2.3.2 Spherical Cavity	8
3 Methods	10
3.1 Ehrenpreis–Palamodov Gaussian Process	10
3.1.1 Fundamental Principle	10
3.1.2 Gaussian Maxwell Prior	11
3.1.3 Conditioning and Posterior	12
3.1.4 Discretization	12
3.1.5 Implementation maxwellgp	14
3.1.6 EPGP for Boundary Value Problems	14
3.1.7 EPGP for Cavity Scattering	14
3.2 Boundary Element Method	14
3.2.1 Formulation	14
3.2.2 BEM for Cavity Scattering	15
3.2.3 Implementation Bembel	16
4 Benchmark Results	18
4.1 Spherical Cavity	18
4.2 Ellipsoidal Cavity	23
5 Conclusion and Outlook	30
5.1 Summary	30
5.2 Limitations	30
5.3 Future Work	31
A Implementation Source Code	32
B Document Source Code	33
C Use of Generative AI	34
Bibliography	i

Chapter 1

Introduction

1.1 Context and Motivation

Machine learning (ML) is increasingly used as a tool in the sciences and engineering, where the systems of interest obey well-established physical laws. This has given rise to the field of **scientific machine learning (SciML)**, also known as **AI4Science**. Purely data-driven models ignore this known physical structure and must rediscover the laws from data alone, which is wasteful and typically inaccurate outside the regime they were trained on. The more principled approach embeds established physical knowledge into the model from the outset. This idea appears under several names, such as **first-principles AI (FPAI)**, **physics-informed ML** or **hybrid modeling**, which all refer to essentially the same idea: letting the governing equations constrain the model rather than learning them from data.

Many physical laws take the form of **partial differential equations (PDEs)**. There are two main ways to make a model respect such a law. Weak enforcement adds the PDE residual as a penalty in the training objective, so the constraint is only satisfied approximately and is traded off against the data. **Strong enforcement** instead restricts the model to functions that satisfy the equation identically by construction. **Structure preservation** is the key advantage: Unphysical phenomena, such as spurious modes or phantom charges, cannot arise. The model can never violate the law it is meant to obey, regardless of the data. When the governing law is known exactly, strong enforcement is therefore preferable.

Among ML models, the **Gaussian process (GP)** (Rasmussen and Williams 2006, Sec. 2.2) is particularly well suited to strong enforcement. A GP is a probabilistic model over functions that returns a full posterior distribution instead of just a point estimate. This enables **uncertainty quantification**: Every prediction carries an estimate of its own reliability. If the GP prior is built so that its samples satisfy a PDE exactly, the model is strongly physics-informed and probabilistic at the same time.

The **Ehrenpreis–Palamodov (EP)** principle of Ehrenpreis (1970) provides a constructive route to such priors for **linear PDEs with constant coefficients**, inspired by the **inverse Fourier transform**. It represents every solution as a superposition of **plane waves** with wavevectors on the **characteristic variety** of the operator and amplitudes in the **kernel of the symbol**. Its GP realization, the **Ehrenpreis–Palamodov Gaussian Process (EPGP)** of Härkönen et al. (2023), uses a Gaussian prior based on this principle. The prior and the posterior both lie exactly in the solution space of the operator.

A particularly well-understood PDE system is **Maxwell’s equations**, which govern **electromagnetism (EM)**. Building a principled EPGP for these equations is part of the ongoing effort of Vittori and Kurz (2026), which constructs EPGP priors for the **time-harmonic** Maxwell system from a geometric perspective based on **differential forms**, the **de Rham complex**, and **Hertz potentials**. This yields a more geometrically grounded construction than the generic one of Härkönen et al. (2023).

The theoretical framework and a proof-of-concept exist but a convincing reference benchmark does not. This thesis provides that benchmark.

1.2 Cavity Scattering Benchmark

As a benchmark we use an **interior electromagnetic scattering problem** in a cavity with **perfectly electrically conducting (PEC)** boundaries, a setup we adapt from the cavity-scattering problem of Zeng et al. (2011). **Dipole sources** placed inside the cavity act as transmitters, exciting a field. It scatters off the PEC boundary, reflects back into the interior, and is measured at receivers. This defines a **reaction operator** that maps the dipole excitations to the corresponding field responses. This operator is the object of study.

We construct a probabilistic EPGP surrogate of this reaction operator. Unlike a purely data-driven surrogate, it satisfies the time-harmonic Maxwell equations exactly by construction and reports its own uncertainty. Its plane-wave prior spans exactly the source-free solutions of the interior Maxwell problem, matching the scattered field it must represent. The surrogate is built with the Maxwell EPGP library `maxwellgp` from the ongoing work of Vittori and Kurz (2026).

For comparison, we compute the reaction operator in a second, independent way, with a **boundary element method (BEM)** built on a **boundary integral formulation** of the same underlying boundary value problem. For this we use the BEM library `Bembe1` (Dölz et al. 2020). The two solvers share only the problem setup and nothing of their internal discretizations.

We benchmark on two cavity geometries. For a **spherical cavity**, the scattered field and the reaction operator are available in closed form. This lets us validate the EPGP and the BEM solver independently, each against the **analytic solution**, and thereby establish that both are correct. For an **ellipsoidal cavity** no closed-form solution exists, so we must rely on numerical methods alone. Since both solvers were already certified against the analytic solution on the sphere, their close agreement here is strong evidence that both are correct.

1.3 Contributions

This thesis makes the following contributions.

- Maxwell EPGP library: Improved existing JAX implementation `maxwellgp`, fixed bugs, added support for tangential traces, posterior covariance computation, and sampling.
- Cavity EPGP solver: Built a probabilistic EPGP solver `cavity-maxwellgp` for the cavity reaction operator in Python using `maxwellgp`, relying on the analytic dipole Green's function.
- Cavity BEM solver: Built a deterministic boundary element solver `cavity-bem` for the cavity reaction operator in C++ using `Bembe1`, relying on an indirect single-layer formulation.
- Analytic spherical cavity solution: Assembled the closed-form reaction operator for the spherical cavity from the standard vector spherical harmonic expansion, serving as exact ground truth.
- Benchmarking: Built a benchmark harness `cavity-benchmark`. Demonstrated convergence of the EPGP and BEM solvers against the analytic solution on the spherical cavity. Demonstrated agreement between the two solvers on the ellipsoidal cavity, which has no analytic solution. Reported a 2D convergence grid and an accuracy–runtime trade-off. Performed uncertainty quantification of the EPGP surrogate.

All solver code and the experiment scripts are openly available on GitHub, as listed in Appendix A, so every result reported here is reproducible.

Chapter 2

Cavity Scattering

We present the interior electromagnetic scattering problem: the governing field equations, the cavity geometry, and the reaction operator that both solvers compute.

2.1 Time-Harmonic Maxwell's Equations

The EPGP of Vittori and Kurz (2026) is built on the time-harmonic Maxwell's equations, which we briefly review here.

We start with the source-free Maxwell's equations in three-dimensional space $\mathbb{R}^3$

$$\begin{aligned}\mathbf{curl} \mathbf{E} &= -\partial_t \mathbf{B} & \operatorname{div} \mathbf{D} &= 0 \\ \mathbf{curl} \mathbf{H} &= +\partial_t \mathbf{D} & \operatorname{div} \mathbf{B} &= 0,\end{aligned}\tag{1}$$

together with the constitutive relations in vacuum,

$$\mathbf{D} = \varepsilon \mathbf{E} \quad \mathbf{B} = \mu \mathbf{H}.\tag{2}$$

Here the electric field $\mathbf{E}$, magnetic field $\mathbf{H}$, electric flux density $\mathbf{D}$, magnetic flux density $\mathbf{B}$, permittivity ε , and permeability μ carry their standard meanings (Jackson 1999, Ch. 6).

We choose the time-harmonic ansatz

$$\begin{aligned}\mathbf{E}(\mathbf{x}, t) &= \operatorname{Re}(\mathbf{E}(\mathbf{x})e^{-i\omega t}) \\ \mathbf{H}(\mathbf{x}, t) &= \operatorname{Re}(\mathbf{H}(\mathbf{x})e^{-i\omega t}),\end{aligned}\tag{3}$$

under which differentiation in time becomes complex multiplication,

$$\partial_t \mapsto -i\omega.\tag{4}$$

This yields the time-harmonic Maxwell's equations in the frequency domain,

$$\begin{aligned}\mathbf{curl} \mathbf{E} &= +i\omega\mu \mathbf{H} \\ \mathbf{curl} \mathbf{H} &= -i\omega\varepsilon \mathbf{E}.\end{aligned}\tag{5}$$

Eliminating the magnetic field via

$$\mathbf{H} = \frac{1}{i\omega\mu} \mathbf{curl} \mathbf{E},\tag{6}$$

gives the curl-curl equation for the electric field $\mathbf{E}$

$$\mathbf{curl} \mathbf{curl} \mathbf{E} - k^2 \mathbf{E} = 0,\tag{7}$$

with the wavenumber

$$k = \omega\sqrt{\varepsilon\mu} = \frac{\omega}{c}.\tag{8}$$

This is a Helmholtz-type equation for the electric field, with the operator $\mathcal{L} := \mathbf{curl} \mathbf{curl} - k^2$.


2.2 Scattering Problem

The scattering problem of interest is inspired by Zeng et al. (2011). It is an interior problem in a bounded cavity.

The problem is set up as follows:

A transmitter dipole δ_t is placed in the interior of a cavity D and radiates an analytically known incident field E^i . Throughout, the subscripts t and r mark the transmitter and the receiver. The incident field hits the cavity boundary ∂D and is reflected by it, creating an unknown scattered field E^s . A receiver dipole δ_r then reads back the field. Physically the receiver sees the total field $E = E^i + E^s$ which is a superposition of the incident and scattered part, but the incident field is split off, leaving the scattered field as the quantity of interest, which satisfies the source-free curl-curl equation in the interior.

The object of study is the mapping from a transmitter dipole to the measured response at a receiver dipole.



2.2.1 Transmitter Dipole and Incident Field

Before introducing oscillating dipoles, we first consider the simpler oscillating monopole.

An oscillating monopole $\delta^1 = (z, q)$ is a charge source with density $\rho(x, t) = q\delta_z \exp(-i\omega t)$ at a point $z \in D$, with the unit Dirac delta distribution δ_z centered at z and a charge $q \in \mathbb{R}$ that gives its strength. The scalar potential it generates is $q\Phi$, with the free-space fundamental solution Φ

$$\Phi(x; z) = \frac{1}{4\pi} \frac{\exp(ikr)}{r}, \quad (9)$$

of the scalar Helmholtz equation

$$(-\Delta - k^2)\Phi(\cdot; z) = \delta_z. \quad (10)$$

Here the distance to the source is $r := \|\mathbf{r}\|$, with separation vector $\mathbf{r} := \mathbf{x} - \mathbf{z}$.

An oscillating Hertzian dipole $\delta^2 = (z, \mathbf{p})$ is a current source with density $\mathbf{J}(x, t) = \mathbf{p}\delta_z \exp(-i\omega t)$ at a point $z \in D$, together with a polarization vector $\mathbf{p} \in \mathbb{R}^3$ that gives the dipole's orientation and strength.

The electric field E^i radiated by a dipole is obtained by applying the curl-curl operator to the Hertz vector potential $\Phi\mathbf{p}$

$$E^i(x; \delta) = \frac{i}{k} \operatorname{curl}_x \operatorname{curl}_x (\Phi(x; z)\mathbf{p}). \quad (11)$$

We can rewrite this expression by factoring out the polarization vector $\mathbf{p}$ by linearity of the curl-curl operator,

$$E^i(x; \delta) = \mathbf{G}(x; z)\mathbf{p}. \quad (12)$$

This defines the free-space electric dyadic Green's function $\mathbf{G}$, which is the Green's function of the operator $\mathcal{L} = \operatorname{curl} \operatorname{curl} - k^2$,

$$\mathbf{G}(x; z)\mathbf{p} := \frac{i}{k} \operatorname{curl}_x \operatorname{curl}_x (\Phi(x; z)\mathbf{p}). \quad (13)$$

Its explicit form (Tai 1994, Ch. 4) is given by

$$\mathbf{G}(x; z) = ik\Phi(x; z) \left[\left(1 + \frac{i}{kr} - \frac{1}{(kr)^2}\right) \mathbf{I} - \left(1 + \frac{3i}{kr} - \frac{3}{(kr)^2}\right) \hat{\mathbf{r}}\hat{\mathbf{r}}^\top \right], \quad (14)$$

where $\hat{\mathbf{r}}\hat{\mathbf{r}}^\top$ is the outer product of the unit separation vector $\hat{\mathbf{r}} := \mathbf{r}/r$ and $\mathbf{I}$ is the 3×3 identity matrix.

2.2.2 Scattered Field and curl-curl BVP

Our cavity $D \subseteq \mathbb{R}^3$ is a smooth, bounded domain. The incident field $\mathbf{E}^i$ is reflected by the cavity boundary ∂D and creates a scattered field $\mathbf{E}^s$.

The transmitter dipole thus determines both the incident field $\mathbf{E}^i(\mathbf{x}; \delta_t)$ and the scattered field $\mathbf{E}^s(\mathbf{x}; \delta_t)$. The incident field is known analytically, but the scattered field is unknown and gives rise to a boundary value problem that must be solved.

PEC Boundary Condition

The boundary conditions of the BVP are the perfectly electrically conducting (PEC) boundary conditions. They enforce that the scattered field reflects the incident field such that their tangential components cancel on the cavity wall. Hence the total field $\mathbf{E} = \mathbf{E}^i + \mathbf{E}^s$ satisfies the PEC condition

$$\pi_t \mathbf{E} = 0 \quad \text{on } \partial D, \quad (15)$$

or equivalently, the scattered field satisfies the boundary condition

$$\mathbf{h} := \pi_t \mathbf{E}^s = -\pi_t \mathbf{E}^i \quad \text{on } \partial D. \quad (16)$$

Here the tangential projection trace π_t onto ∂D , with outward unit normal $\hat{\mathbf{n}}$, is given by

$$\pi_t \mathbf{E} := \hat{\mathbf{n}} \times (\mathbf{E} \times \hat{\mathbf{n}}) = \mathbf{E} - (\mathbf{E} \cdot \hat{\mathbf{n}}) \hat{\mathbf{n}}. \quad (17)$$

Interior BVP

The full BVP is then to find the scattered field $\mathbf{E}^s : D \rightarrow \mathbb{C}^3$ in the interior of the cavity D such that the curl-curl equation is satisfied in the interior and the PEC boundary condition is satisfied on the wall,

$$\begin{aligned} \text{curl curl } \mathbf{E}^s - k^2 \mathbf{E}^s &= 0 \quad \text{in } D \\ \pi_t \mathbf{E}^s &= \mathbf{h} \quad \text{on } \partial D. \end{aligned} \quad (18)$$

This BVP is well-defined only when k^2 is not a cavity resonance, i.e. not an eigenvalue of the curl-curl operator. Otherwise the BVP solution operator $\mathcal{S} : \mathbf{h} \mapsto \mathbf{E}^s$ becomes singular.

2.2.3 Receiver Dipole and Measurement

The receiver dipole $\delta_r = (\mathbf{z}_r, \mathbf{p}_r)$ measures the scattered field $\mathbf{E}^s$ at its location $\mathbf{z}_r$, by projecting it onto its polarization vector $\mathbf{p}_r$, yielding a single complex scalar.

Mathematically this is a linear functional $\mathcal{R} : C^{\text{oo}}(D; \mathbb{C}^3) \rightarrow \mathbb{C}$ defined as

$$\mathcal{R}(\mathbf{E}^s) := \mathbf{p}_r \cdot \mathbf{E}^s(\mathbf{z}_r). \quad (19)$$

Reciprocity

There are two dipoles, one transmitter and one receiver. Their roles are complementary: the transmitter generates the incident field, and the receiver measures the scattered field.

The measured response generated by a transmitter dipole δ_t and measured by a receiver dipole δ_r is

$$r(\delta_t, \delta_r) = \mathcal{R}(\mathbf{E}^s) = \mathbf{p}_r \cdot \mathbf{E}^s(\mathbf{z}_r; \delta_t). \quad (20)$$

Exchanging the two roles leaves the measured response unchanged. This is the Lorentz (Rayleigh–Carson) reciprocity property, a consequence of the symmetry of the curl-curl Green's operator (Colton and Kress 2019, Sec. 6.6),

$$r(\delta_t, \delta_r) = r(\delta_r, \delta_t). \quad (21)$$

2.2.4 Reaction Operator

A single transmit-receive pair yields one number, the reaction $r(\delta_t, \delta_r)$. We now extend to a whole ensemble of dipoles.

Transmitter and receiver are the same kind of object, both dipoles. By reciprocity their roles are interchangeable, so there is no reason to keep two separate families. A single set of dipoles serves both as transmitters and as receivers. We restrict all dipoles to a common surface $\Lambda \subseteq D$, with locations $\mathbf{z} \in \Lambda$ and polarization vectors $\mathbf{p} \in T_{\mathbf{z}}\Lambda$ tangent to it.

Acting together, this set of dipoles excites the field and reads it back. We are interested in the map from the full excitation to the full response, the reaction operator.

Continuous Operator

A tangential dipole density on Λ is a tangential vector field

$$\mathbf{g} : \Lambda \rightarrow T\Lambda \quad \mathbf{g}(\mathbf{z}) \in T_{\mathbf{z}}\Lambda, \quad (22)$$

assigning a polarization vector $\mathbf{g}(\mathbf{z})$ to each point of Λ . A single dipole $\delta = (\mathbf{z}_0, \mathbf{p})$ is the singular special case

$$\mathbf{g} = \mathbf{p}\delta_{\mathbf{z}_0}. \quad (23)$$

By linearity of the problem, the density excites a single incident field that superposes the fields of its constituent dipoles,

$$\mathbf{E}^i[\mathbf{g}](\mathbf{x}) = \int_{\Lambda} \mathbf{E}^i(\mathbf{x}; (\mathbf{z}, \mathbf{g}(\mathbf{z}))) d\mathbf{s}(\mathbf{z}), \quad (24)$$

which is reflected into a single scattered field,

$$\mathbf{E}^s[\mathbf{g}](\mathbf{x}) = \int_{\Lambda} \mathbf{E}^s(\mathbf{x}; (\mathbf{z}, \mathbf{g}(\mathbf{z}))) d\mathbf{s}(\mathbf{z}). \quad (25)$$

Reading the scattered field back on Λ gives the response density $\mathbf{m} : \Lambda \rightarrow T\Lambda$, the tangential projection trace of the scattered field,

$$\mathbf{m} := \pi_t^{\Lambda} \mathbf{E}^s[\mathbf{g}]. \quad (26)$$

Here the trace π_t^{Λ} is the same tangential projection as π_t , but taken on the dipole surface Λ with its own unit normal $\hat{\mathbf{n}}_{\Lambda}$, so that $\pi_t^{\Lambda} \mathbf{E} = \mathbf{E} - (\mathbf{E} \cdot \hat{\mathbf{n}}_{\Lambda}) \hat{\mathbf{n}}_{\Lambda}$ at each point of Λ .

The reaction operator is a linear integral operator, defined as the map from the excitation density to the response density,

$$\begin{aligned} \mathcal{T} : \mathbf{g} &\mapsto \mathbf{m} \\ (\mathcal{T}\mathbf{g})(\mathbf{x}) &= \int_{\Lambda} \mathbf{T}(\mathbf{x}, \mathbf{z}) \mathbf{g}(\mathbf{z}) d\mathbf{s}(\mathbf{z}). \end{aligned} \quad (27)$$

Its integral kernel $\mathbf{T}$ is the point-to-point reaction, the response at $\mathbf{x}$ to a unit dipole $\mathbf{p}$ at $\mathbf{z}$,

$$\begin{aligned} \mathbf{T}(\mathbf{x}, \mathbf{z}) &: T_{\mathbf{z}}\Lambda \rightarrow T_{\mathbf{x}}\Lambda \\ \mathbf{T}(\mathbf{x}, \mathbf{z})\mathbf{p} &:= \pi_t^{\Lambda} \mathbf{E}^s(\mathbf{x}; (\mathbf{z}, \mathbf{p})). \end{aligned} \quad (28)$$

Reciprocity makes the kernel and the operator complex-symmetric (not Hermitian): exchanging the two points and transposing leaves the kernel unchanged,

$$\mathbf{T}(\mathbf{x}, \mathbf{z}) = \mathbf{T}(\mathbf{z}, \mathbf{x})^{\top}. \quad (29)$$

The scalar reaction is this kernel as a bilinear form in the two polarization vectors, one slot per dipole,

$$r(\delta_t, \delta_r) = \mathbf{p}_r^{\top} \mathbf{T}(\mathbf{z}_r, \mathbf{z}_t) \mathbf{p}_t. \quad (30)$$

Discretization

To compute with $\mathcal{T}$ we discretize it on a finite set of dipoles. Since each tangent space $T_z\Lambda$ is two-dimensional, we need a basis to coordinatize it. We sample N_Λ points on Λ and assign each an orthonormal tangent basis $\{\hat{e}_1(z), \hat{e}_2(z)\}$. The basis is built from a pivot vector $\mathbf{c}$, a coordinate axis kept away from parallel to the outward normal $\hat{\mathbf{n}}_\Lambda(z)$: the cross products $\hat{e}_1 \propto \mathbf{c} \times \hat{\mathbf{n}}_\Lambda$ and $\hat{e}_2 = \hat{\mathbf{n}}_\Lambda \times \hat{e}_1$ are both tangent to Λ and mutually orthogonal.

We have two polarization vectors per point, each point thus carries two dipoles, one per polarization, giving $M = 2N_\Lambda$ dipoles in total. We enumerate the dipoles by a single index $i \in \{1, \dots, M\}$, writing $\delta_i = (z_i, \hat{\mathbf{p}}_i)$ for the dipole with location z_i and unit polarization vector $\hat{\mathbf{p}}_i$. Consecutive pairs share a location.

The discretized operator is then a matrix $\mathbf{T} \in \mathbb{C}^{M \times M}$, whose entries are the kernel evaluated at the two dipoles δ_i and δ_j ,

$$\mathbf{T}_{ij} = \hat{\mathbf{p}}_i^\top \mathbf{T}(z_i, z_j) \hat{\mathbf{p}}_j = r(\delta_j, \delta_i), \quad (31)$$

inheriting the kernel's symmetry,

$$\mathbf{T} = \mathbf{T}^\top. \quad (32)$$

This matrix is the object of interest for the benchmark. Both solvers compute it, and the resulting operators are compared.

2.3 Geometry

We now fix the concrete geometry of our benchmark problem. The geometry is mostly taken from Zeng et al. (2011).

We fix the wavenumber $k = 2$. The dipole surface Λ is the unit sphere,

$$\Lambda := \{\mathbf{x} \in \mathbb{R}^3 : \|\mathbf{x}\| = 1\} \subseteq D. \quad (33)$$

We choose $N_\Lambda = 32$ dipole locations drawn from a low-discrepancy quasi-uniform Fibonacci sphere distribution. This is a deterministic point set that spaces neighboring points approximately equally while covering the surface area nearly uniformly.

A standard construction (González 2010) places the points $\mathbf{x}_j \in \mathbb{S}^2$, for $j = 0, \dots, N_\Lambda - 1$, at

$$z_j = 1 - \frac{2(j + \frac{1}{2})}{N_\Lambda}, \quad \mathbf{x}_j = \left(\sqrt{1 - z_j^2} \cos(2\pi j/\varphi), \sqrt{1 - z_j^2} \sin(2\pi j/\varphi), z_j \right), \quad (34)$$

with the golden ratio $\varphi = (1 + \sqrt{5})/2$.

Together with the 2 polarizations per point, this gives $M = 2N_\Lambda = 64$ configurations.

There are two cavity geometries: an ellipsoidal cavity and a spherical cavity.

2.3.1 Ellipsoidal Cavity

The ellipsoidal cavity is the original geometry from Zeng et al. (2011). It is a smooth, convex, and simply connected domain.

Its boundary ∂D is an ellipsoidal surface with semi-axes $\mathbf{a} = (a_1, a_2, a_3) = (4, 4, 6)$.

The interior of the ellipsoid is the domain D in which the interior BVP is solved,

$$D := \left\{ (x_1, x_2, x_3) \in \mathbb{R}^3 : \left(\frac{x_1}{a_1} \right)^2 + \left(\frac{x_2}{a_2} \right)^2 + \left(\frac{x_3}{a_3} \right)^2 < 1 \right\}. \quad (35)$$

Λ is well-separated from ∂D with a minimum distance of $d_{\min} = 3$.




Figure 1: Ellipsoidal cavity cross-section

2.3.2 Spherical Cavity

The spherical cavity is a new addition that allows for an analytic solution. It is a smooth, convex, and simply connected domain.

Its boundary ∂D is a spherical surface with radius $R = 4$. It is obtained from the ellipsoid by shrinking all semi-axes to the shortest, $a_1 = a_2 = 4$.

The interior of the sphere is the domain D in which the interior BVP is solved,

$$D := \{ \mathbf{x} \in \mathbb{R}^3 : \|\mathbf{x}\| < R \} \quad R = 4. \quad (36)$$




Figure 2: Spherical cavity cross-section

Analytic Solution

Thanks to the spherical symmetry, the interior BVP is separable and hence has a closed form solution. Therefore both scattered field $\mathbf{E}^s$ and the reaction operator $\mathbf{T}$ have analytic formulas, which serve as exact reference solutions to validate both numerical solvers.

We do not re-derive the analytic solution here. We refer to (Tai 1994, Sec. 10-4) for the full derivation and record the resulting formulas below.

The interior fields expand in the regular Hansen multipoles, the TM-type $\mathbf{N}_{lm}$ and the TE-type $\mathbf{M}_{lm}$. On a sphere of radius r their tangential traces reduce to the tangential vector spherical harmonics $\mathbf{\Psi}_{lm}$ and $\mathbf{\Phi}_{lm}$,

$$\pi_t \mathbf{N}_{lm} = \frac{\psi'_l(kr)}{kr} \mathbf{\Psi}_{lm} \quad \pi_t \mathbf{M}_{lm} = j_l(kr) \mathbf{\Phi}_{lm}, \quad (37)$$

with the spherical Bessel function j_l and the Riccati–Bessel function $\psi_l(x) = x j_l(x)$ and its derivative ψ'_l .

We expand the incident tangential trace on the cavity wall at $r = R$ in these harmonics,

$$p_{lm} = \langle \pi_t \mathbf{E}^i, \mathbf{\Psi}_{lm} \rangle_{r=R} \quad q_{lm} = \langle \pi_t \mathbf{E}^i, \mathbf{\Phi}_{lm} \rangle_{r=R}. \quad (38)$$

The scattered field expands in the same multipoles,

$$\mathbf{E}^s = \sum_{lm} (a_{lm} \mathbf{N}_{lm} + b_{lm} \mathbf{M}_{lm}). \quad (39)$$

The PEC condition $\pi_t \mathbf{E}^s = -\pi_t \mathbf{E}^i$ on the cavity wall ∂D at $r = R$ fixes the coefficients,

$$a_{lm} = -\frac{kR}{\psi'_l(kR)} p_{lm} \quad b_{lm} = -\frac{1}{j_l(kR)} q_{lm}. \quad (40)$$

Evaluating the scattered field on the dipole surface Λ at $r = r_0$ and taking its tangential trace gives the measured response,

$$\pi_t^\Lambda \mathbf{E}^s = -\sum_{lm} \left(\frac{R}{r_0} \frac{\psi'_l(kr_0)}{\psi'_l(kR)} p_{lm} \mathbf{\Psi}_{lm} + \frac{j_l(kr_0)}{j_l(kR)} q_{lm} \mathbf{\Phi}_{lm} \right). \quad (41)$$

Projecting onto the receiver polarization vector gives the reference reaction operator,

$$\mathbf{T}_{ij} = \hat{\mathbf{p}}_i \cdot \pi_t^\Lambda \mathbf{E}^s(\mathbf{z}_i; \delta_j). \quad (42)$$

Chapter 3

Methods

This chapter describes the two solvers used to compute the cavity reaction operator $\mathbf{T}$. Both represent the scattered field as a superposition of exact solutions of the interior curl-curl equation, so Maxwell's equations hold by construction and only the PEC boundary condition remains to be enforced. Beyond this shared structure, the two solvers are numerically independent: they share only the problem setup (cavity geometry, wavenumber, dipole configurations) and none of their internal discretizations.

EPGP is probabilistic, returning a posterior distribution over the solution space. It uses a volume plane-wave superposition ansatz and enforces the boundary condition by conditioning on boundary data at scattered collocation points, with no mesh required. BEM is deterministic, returning a single solution. It discretizes the cavity wall into a surface mesh, represents the field via a boundary single-layer potential, and enforces the boundary condition by solving a boundary integral equation.

3.1 Ehrenpreis–Palamodov Gaussian Process

The Ehrenpreis–Palamodov Gaussian Process (EPGP) is a probabilistic, spectral method. Its ansatz is a superposition of transverse plane waves, which satisfy Maxwell's equations exactly. It places a Gaussian prior on the plane-wave amplitude density and enforces the boundary condition by conditioning on the boundary data, yielding a posterior distribution over the solution space.

3.1.1 Fundamental Principle

A linear PDE with constant coefficients takes the form $L\mathbf{u} = 0$, where $L = L[\partial_1, \dots, \partial_n]$ is a matrix-valued polynomial in the partial derivatives. The Fourier transform turns differentiation into multiplication, $\partial_j \mapsto ik_j$, so L acts on plane waves by its symbol,

$$L(\mathbf{a} \exp(i\mathbf{k} \cdot \mathbf{x})) = P(\mathbf{k})\mathbf{a} \exp(i\mathbf{k} \cdot \mathbf{x}), \quad (43)$$

where $P(\mathbf{k}) := L[ik_1, \dots, ik_n]$ is an ordinary matrix polynomial in the wavevector $\mathbf{k}$. A plane wave is therefore a solution exactly when its amplitude lies in the kernel of the symbol, $P(\mathbf{k})\mathbf{a} = 0$.

The Ehrenpreis–Palamodov fundamental principle reverses this reasoning. It states that *all* solutions $\mathbf{u}$ of linear PDEs with constant coefficients can be recovered by an inverse Fourier transform, as a continuous superposition of such plane-wave solutions

$$\mathbf{u}(\mathbf{x}) = \int_V \mathbf{a}(\mathbf{k}) \exp(i\mathbf{k} \cdot \mathbf{x}) d\mathbf{k}, \quad (44)$$

with wavevectors $\mathbf{k}$ on the characteristic variety V of the symbol P

$$V = \{\mathbf{k} \in \mathbb{R}^d : \det P(\mathbf{k}) = 0\}, \quad (45)$$

and amplitude densities $\mathbf{a}$ in its kernel,

$$\mathbf{a}(\mathbf{k}) \in \ker P(\mathbf{k}). \quad (46)$$

The variety collects the wavevectors at which the symbol is singular, and the kernel condition restricts each amplitude to the directions the symbol annihilates.

We can remove the kernel constraint by projection. Writing the orthogonal projector onto $\ker P(\mathbf{k})$ as $\Pi_{\mathbf{k}}$, any unconstrained weight $w(\mathbf{k})$ yields an admissible amplitude $\mathbf{a}(\mathbf{k}) = \Pi_{\mathbf{k}} w(\mathbf{k})$. Folding the projector into a matrix-valued feature map,

$$\Phi_{\mathbf{k}}(x) = \Pi_{\mathbf{k}} \exp(i\mathbf{k} \cdot x), \quad (47)$$

every solution becomes a superposition of features with a free weight,

$$u(x) = \int_V \Phi_{\mathbf{k}}(x) w(\mathbf{k}) d\mathbf{k}. \quad (48)$$

Time-Harmonic Maxwell's Equations

We specialize to the time-harmonic Maxwell's equations, equivalently the curl-curl Helmholtz equation. Its symbol, with the wavenumber k , is

$$P(\mathbf{k}) = (\|\mathbf{k}\|^2 - k^2)\mathbf{I} - \mathbf{k}\mathbf{k}^\top, \quad (49)$$

whose determinant vanishes on a k -scaled 2-sphere, the characteristic variety

$$V_k = \{\mathbf{k} \in \mathbb{R}^3 : \|\mathbf{k}\| = k\} = k\mathbb{S}^2. \quad (50)$$

On the variety the symbol is singular and reduces to a scaled rank-one projector onto the wavevector,

$$P(\mathbf{k}) = -\mathbf{k}\mathbf{k}^\top \quad \mathbf{k} \in V_k, \quad (51)$$

whose kernel is the plane perpendicular to the wavevector,

$$\ker P(\mathbf{k}) = \{\mathbf{a} \in \mathbb{C}^3 : \mathbf{k} \cdot \mathbf{a} = 0\}. \quad (52)$$

Enforcing Maxwell's equations therefore amounts to a transversality constraint on the amplitude,

$$\mathbf{a}(\mathbf{k}) \perp \mathbf{k}. \quad (53)$$

The orthogonal projector onto the kernel is the transverse complement of that rank-one projector,

$$\Pi_{\mathbf{k}} = \mathbf{I} - \frac{\mathbf{k}\mathbf{k}^\top}{\|\mathbf{k}\|^2}. \quad (54)$$

The general feature map and EP representation thus specialize to a superposition of transverse plane waves with a free weight,

$$\Phi_{\mathbf{k}}(x) = \Pi_{\mathbf{k}} \exp(i\mathbf{k} \cdot x) \quad E(x) = \int_{V_k} \Phi_{\mathbf{k}}(x) w(\mathbf{k}) d\mathbf{k}. \quad (55)$$

3.1.2 Gaussian Maxwell Prior

Since the field is linear in the weight, a zero-mean Gaussian measure on the weight makes the field a zero-mean Gaussian process,

$$E \sim \mathcal{GP}(0, \mathbf{K}). \quad (56)$$

The weight has covariance $\mathbf{W}$, a positive-semidefinite operator on the weight space that sets the prior strength of each spectral direction. The matrix covariance kernel is then

$$\mathbf{K}(x, y) = \int_{V_k} \Phi_{\mathbf{k}}(x) \mathbf{W} \Phi_{\mathbf{k}}^H(y) d\mathbf{k}. \quad (57)$$

The isotropic choice $\mathbf{W} = \mathbf{I}$ weights all spectral directions equally and recovers the plain feature kernel. We adopt it in the Hyperparameters section but keep the weight covariance general here. Because the projector sits inside every feature, the prior is supported entirely on the solution space: every sample satisfies Maxwell exactly.

3.1.3 Conditioning and Posterior

We turn the prior into a solver by conditioning on observations of the field. An observation is a linear functional $\mathcal{R}$ of the field, for instance its value or its tangential trace at a point. We condition on N_b observations at points X_b , collected with their measured values in a data vector $\mathbf{h}$.

After conditioning the posterior is again a Gaussian process

$$\mathbf{E} \mid \mathbf{h} \sim \mathcal{GP}(\mathbf{E}_*, \mathbf{K}_*), \quad (58)$$

with posterior mean $\mathbf{E}_*$ and posterior covariance $\mathbf{K}_*$.

We introduce a Tikhonov regularization parameter $\sigma_n^2 > 0$. It trades off fitting the data against trusting the prior: as $\sigma_n^2 \rightarrow 0$ the posterior mean interpolates the observations exactly, while $\sigma_n^2 > 0$ relaxes this to a regression that smooths the fit.

The posterior mean field is the regularized best fit to the data, tempered by the prior and the regularization,

$$\mathbf{E}_*(\mathbf{x}) = \mathbf{K}(\mathbf{x}, X_b)(\mathbf{K}_{bb} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{h}. \quad (59)$$

Here the Gram matrix $\mathbf{K}_{bb} := \mathbf{K}(X_b, X_b)$ is the kernel evaluated at all pairs of the N_b observation points.

The posterior covariance is the Schur complement of the conditioning block, measuring how underdetermined the field remains after conditioning,

$$\mathbf{K}_*(\mathbf{x}, \mathbf{y}) = \mathbf{K}(\mathbf{x}, \mathbf{y}) - \mathbf{K}(\mathbf{x}, X_b)(\mathbf{K}_{bb} + \sigma_n^2 \mathbf{I})^{-1} \mathbf{K}(X_b, \mathbf{y}). \quad (60)$$

This is exact, infinite-dimensional GP regression: prior and posterior both live on the solution space.

3.1.4 Discretization

The posterior of the previous section is exact but not yet computable: the kernel entries are integrals over V_k with no closed form. We discretize the EP integral by quadrature, which both makes the kernel computable and expresses the field through explicit finite features.

Finite Spectral Features

We approximate the integral over V_k by a finite sum over N_s spectral directions $\mathbf{k}_j \in V_k$. The characteristic variety is itself a sphere, $V_k = k\mathbb{S}^2$, so the Fibonacci construction that placed the dipoles on Λ serves here unchanged, rescaled from the unit sphere to radius k . The two uses are the same problem on two different spheres, one in space and one in the Fourier domain.¹

For each direction we pick an orthonormal basis $\mathbf{a}_{j1}, \mathbf{a}_{j2}$ of the transverse plane, by the same pivot-and-cross-product construction used for the tangent basis on Λ , with the wavevector now in the role of the normal.² Taking as pivot the coordinate axis $\mathbf{c}_j$ least aligned with $\mathbf{k}_j$, we set $\mathbf{a}_{j1} \propto \mathbf{k}_j \times \mathbf{c}_j$ and $\mathbf{a}_{j2} \propto \mathbf{k}_j \times \mathbf{a}_{j1}$. This recovers the projector as an outer-product sum,

$$\Pi_{\mathbf{k}_j} = \sum_{a=1}^2 \mathbf{a}_{ja} \mathbf{a}_{ja}^H, \quad (61)$$

so each spectral direction contributes two scalar plane-wave features $\varphi_{ja}(\mathbf{x}) = \mathbf{a}_{ja} \exp(i\mathbf{k}_j \cdot \mathbf{x})$, giving $F = 2N_s$ features in total. The field becomes a finite superposition with one scalar coefficient per feature,

$$\mathbf{E}(\mathbf{x}) = \sum_{j=1}^{N_s} \sum_{a=1}^2 w_{ja} \varphi_{ja}(\mathbf{x}). \quad (62)$$

These coefficients carry the prior. Gathering them in $\mathbf{w} \in \mathbb{C}^F$, the Gaussian measure becomes a finite complex Gaussian,

¹The two are implemented separately, in cavity-benchmark for Λ and in maxwellgp for V_k . Their golden-angle forms differ by a full turn per index, which on the half-integer indices offsets the azimuth by π , so the two point sets agree up to a rotation about the z -axis.

²The two implementations pick the pivot differently: on Λ a fixed coordinate axis is swapped for another once it comes close to parallel with the normal, while here the least-aligned axis is taken outright. Either rule yields an admissible pivot.

$$\mathbf{w} \sim \mathcal{CN}(0, \mathbf{W}). \quad (63)$$

Every feature is a transverse plane wave, so this finite prior is still supported entirely on the solution space.

The choice of transverse basis does not matter. Any other orthonormal pair spanning the same plane gives the same outer-product sum, and with isotropic weights $\mathbf{W} = \mathbf{I}$ the kernel sees the features only through it, so prior and posterior are unchanged.

Sampling Λ and sampling V_k are thus one construction carried out twice: place points on a sphere, then frame the plane perpendicular to the radius at each of them. That plane carries the dipole polarizations in space and the wave amplitudes in the Fourier domain, and its dimension is the factor of two in both $M = 2N_\Lambda$ and $F = 2N_s$.

Weight-Space Posterior

The explicit features admit a second, equivalent form of the same posterior. Rather than conditioning the process through its kernel, we infer the weight vector $\mathbf{w}$ in the expansion

$$\mathbf{E}(\mathbf{x}) = \Phi(\mathbf{x})^H \mathbf{w}, \quad \mathbf{w} \sim \mathcal{CN}(0, \mathbf{W}). \quad (64)$$

Evaluating the features at the observation points through $\mathcal{R}$ gives the design matrix

$$\Phi = \Phi(X_b) \in \mathbb{C}^{F \times N_b}, \quad (65)$$

so that the observation model becomes

$$\mathbf{h} = \Phi^H \mathbf{w} + \boldsymbol{\eta}, \quad \boldsymbol{\eta} \sim \mathcal{CN}(0, \sigma_n^2 \mathbf{I}), \quad (66)$$

in which the regularization σ_n^2 of the previous section reappears as the observation-noise variance. This is a standard Bayesian linear regression for $\mathbf{w}$. Its posterior is Gaussian,

$$\mathbf{w} \mid \mathbf{h} \sim \mathcal{CN}(\mathbf{w}_*, \mathbf{A}^{-1}), \quad (67)$$

with posterior precision and mean

$$\mathbf{A} = \mathbf{W}^{-1} + \frac{1}{\sigma_n^2} \Phi \Phi^H, \quad \mathbf{w}_* = \frac{1}{\sigma_n^2} \mathbf{A}^{-1} \Phi \mathbf{h}. \quad (68)$$

Mapping the inferred weights back through the features gives the posterior mean field $\mathbf{E}_*(\mathbf{x}) = \Phi(\mathbf{x})^H \mathbf{w}_*$.

Under the same quadrature the prior kernel collapses to the finite feature product

$$\mathbf{K}(\mathbf{x}, \mathbf{y}) = \Phi(\mathbf{x})^H \mathbf{W} \Phi(\mathbf{y}), \quad (69)$$

which is the discretization used in the function-space formulation. Substituting it into the function-space posterior and applying the Woodbury matrix identity reproduces the expressions above, so the two formulations give identical posterior means and covariances.

The distinction is computational. The function-space solve inverts the $N_b \times N_b$ Gram matrix $\mathbf{K}_{bb} + \sigma_n^2 \mathbf{I}$, while the weight-space solve inverts the $F \times F$ precision $\mathbf{A}$. Since the explicit-polarization construction gives $F = 2N_s$ features, the weight-space solve is the cheaper one whenever $F < N_b$, trading a cost set by the number of observations for one set by the number of features.

Hyperparameters

The model has three hyperparameters: the spectral directions $\mathbf{k}_j$, the prior weights $\mathbf{W}$, and the regularization parameter σ_n . In a Gaussian process these can be tuned by maximizing the marginal likelihood of the data, usually by gradient descent on its negative logarithm. We instead fix them on principled grounds. The directions come from the Fibonacci sphere, whose even coverage we prefer to keep. The prior weights are set to $\mathbf{W} = \mathbf{I}$, treating all spectral directions equally. The regularization parameter is held fixed, since maximizing the marginal likelihood would drive it toward the floor set by the ill-conditioned feature system rather than toward a value that reflects the data, as discussed in the conclusion. In all benchmark runs it is fixed to $\sigma_n^2 = e^{-12} \approx 6.1 \times 10^{-6}$, that is $\sigma_n = e^{-6} \approx 2.5 \times 10^{-3}$.

3.1.5 Implementation maxwellgp

The prior and posterior above are implemented in the maxwellgp library of Vittori and Kurz (2026), written in Python and JAX, general and problem-independent. The observation functional $\mathcal{R}$ is supplied to the library and realized through the feature map. Only $\mathcal{R}$ changes between applications, while the plane-wave features, the prior weights, the directions, and the posterior solve are identical.

3.1.6 EPGP for Boundary Value Problems

The prior already solves the PDE in the interior, so conditioning turns it into a boundary value problem solver: observing the prescribed boundary trace yields the posterior field consistent with it. The relevant functional is the tangential trace $\mathcal{R} = \pi_t$, evaluated at N_b points on the boundary ∂D , each carrying its outward unit normal $\hat{n}$.

Although the prescribed boundary data is exact, we keep $\sigma_n^2 > 0$. The finite feature space cannot represent the trace exactly, and the plane-wave Gram matrix is ill-conditioned, so a small noise acts as a Tikhonov regularizer that stabilizes the solve.

Convergence is governed by two parameters, the number of spectral directions N_s and the number of boundary conditioning points N_b .

3.1.7 EPGP for Cavity Scattering

The cavity-maxwellgp layer specializes the BVP solver to the cavity scattering problem. The cavity enters only through the boundary data, and the prior is unchanged. Each transmitter dipole sets a scattered-field boundary trace $\mathbf{h} = -\pi_t \mathbf{E}^i$. Conditioning on it at the N_b boundary points enforces the PEC condition and yields the posterior scattered field $\mathbf{E}_*^s$, whose tangential trace at the receivers fills one column of the reaction operator.

The $M = 2N_\Lambda$ transmitters share the same N_b conditioning points and differ only in their boundary values $\mathbf{h}$. The conditioning matrix $\mathbf{A}$ is therefore identical across transmitters, and only the right-hand side changes. We factor $\mathbf{A}$ once and reuse the factorization for all M excitations. The posterior covariance depends only on the conditioning points, not on the boundary values, so it is shared by all transmitters and computed once.

3.2 Boundary Element Method

The boundary element method (BEM) is a deterministic, boundary-integral method, described in Colton and Kress (2019) and Buffa and Hiptmair (2003). Its ansatz is a single-layer potential, which satisfies Maxwell's equations exactly. It leaves the surface density unknown and enforces the boundary condition by solving the resulting electric field integral equation, yielding a single solution. BEM is well established and serves here only as the framework for the reference solution, so we present the main ideas rather than develop it in full.

3.2.1 Formulation

Maxwell Single-Layer Potential

The Maxwell single-layer potential Ψ_{SL} is defined as

$$\begin{aligned} \Psi_{\text{SL}} : H^{-\frac{1}{2}}(\text{div}_\Gamma, \partial D) &\rightarrow H(\text{curl}, D) \\ (\Psi_{\text{SL}} \mathbf{j})(\mathbf{x}) &= ik \int_{\partial D} \mathbf{G}(\mathbf{x}; \mathbf{y}) \mathbf{j}(\mathbf{y}) \text{d} s(\mathbf{y}), \end{aligned} \quad (70)$$

where $\mathbf{G}$ is the free-space electric dyadic Green's function of the Helmholtz curl-curl equation and $\mathbf{j} : \partial D \rightarrow T(\partial D)$ is a tangential boundary density.

The single-layer potential is a solution of the interior curl-curl equation for any density,

$$\text{curl curl}(\Psi_{\text{SL}} \mathbf{j}) - k^2(\Psi_{\text{SL}} \mathbf{j}) = 0 \quad \text{in } D. \quad (71)$$

Rotated Tangential Trace

There is a more natural trace for BEM, the rotated tangential trace $\gamma_{\times}$,

$$\gamma_{\times} \mathbf{u} := \hat{\mathbf{n}} \times \mathbf{u}. \quad (72)$$

It differs from the projection tangential trace π_t by a quarter turn in the tangent plane,

$$\gamma_{\times} = \hat{\mathbf{n}} \times \pi_t. \quad (73)$$

Maxwell Single-Layer Operator

The Maxwell single-layer operator $\mathcal{V}$ is obtained by applying the rotated tangential trace $\gamma_{\times}$ to the single-layer potential,

$$\begin{aligned} \mathcal{V} : H^{-\frac{1}{2}}(\text{div}_{\Gamma}, \partial D) &\rightarrow H^{-\frac{1}{2}}(\text{curl}_{\Gamma}, \partial D) \\ \mathcal{V} &:= \gamma_{\times} \Psi_{\text{SL}}. \end{aligned} \quad (74)$$

Electric Field Integral Equation

In the indirect formulation, the electric field is represented by a Maxwell single-layer potential Ψ_{SL} with an unknown density $\mathbf{j}$,

$$\mathbf{E} = \Psi_{\text{SL}} \mathbf{j}. \quad (75)$$

By taking the rotated tangential trace of both sides,

$$\gamma_{\times} \mathbf{E} = \gamma_{\times} \Psi_{\text{SL}} \mathbf{j}, \quad (76)$$

we obtain the boundary integral equation for the density,

$$\mathcal{V} \mathbf{j} = \mathbf{h}_{\times} \quad \text{on } \partial D, \quad (77)$$

where $\mathbf{h}_{\times} = \gamma_{\times} \mathbf{E}$ is the boundary forcing.

This is an electric field integral equation (EFIE) in indirect single-layer form. As a first-kind Fredholm integral equation, it is inherently ill-conditioned.

The electric field solution is obtained by applying the single-layer potential to the density, and the BVP solution operator is then formally

$$\mathcal{S} = \Psi_{\text{SL}} \mathcal{V}^{-1}. \quad (78)$$

3.2.2 BEM for Cavity Scattering

PEC Boundary Condition

In the rotated trace the PEC condition reads

$$\begin{aligned} \gamma_{\times} \mathbf{E} &= \gamma_{\times} (\mathbf{E}^i + \mathbf{E}^s) = 0 \\ \gamma_{\times} \mathbf{E}^s &= -\gamma_{\times} \mathbf{E}^i. \end{aligned} \quad (79)$$

Scattered Field

The boundary data is

$$\mathbf{h}_{\times} = -\gamma_{\times} \mathbf{E}^i. \quad (80)$$

This gives us the boundary integral equation

$$\mathcal{V} \mathbf{j} = -\gamma_{\times} \mathbf{E}^i \quad \text{on } \partial D. \quad (81)$$

The scattered field is then recovered by applying the single-layer potential to the density,

$$\mathbf{E}^s = \Psi_{\text{SL}} \mathbf{j} = \mathcal{S}(-\gamma_{\times} \mathbf{E}^i). \quad (82)$$

Measurement

The receiver measurement is obtained by projecting the scattered field onto the polarization of the receiver dipole $\delta = (z, \mathbf{p})$

$$r = \mathbf{p} \cdot (\pi_t^\Lambda \mathbf{E}^s)(z) = \mathbf{p} \cdot (\pi_t^\Lambda \mathcal{S}(-\gamma_\times \mathbf{E}^i))(z). \quad (83)$$

Notice the trace asymmetry: the density is solved against the rotated trace $\gamma_\times$ on the cavity boundary ∂D , but the measurement reads the tangential projection trace π_t^Λ on the dipole surface Λ .

Reaction Operator

Each of the $M = 2N_\Lambda$ transmitter dipoles, two polarizations per point on Λ , generates a different incident field and hence a different right-hand side $\mathbf{h}_\times = -\gamma_\times \mathbf{E}^i$ of the EFIE.

By using an LU decomposition of the single-layer matrix, we can reuse the expensive factorization for each right-hand side.

Evaluating the resulting scattered field at the receivers and projecting onto the polarization frame fills one column of the reaction operator.

3.2.3 Implementation Bembe1

The formulation is discretized with the Bembe1 library of Dölz et al. (2020), described by its authors as:

[Bembe1 is] the C++ library featuring higher order isogeometric Galerkin boundary element methods for Laplace, Helmholtz, and Maxwell problems. Bembe1 is compatible with geometries from the Octave NURBS package, and provides an interface to the Eigen template library for linear algebra operations. For computational efficiency, it applies an embedded fast multipole method tailored to the isogeometric analysis framework and a parallel matrix assembly based on OpenMP.

— (Dölz et al. 2020)

Geometry

Bembe1 uses a NURBS boundary representation, where our cavity geometries can be represented exactly. We use a NURBS representation of a sphere and scale it to the cavity semi-axes for both our geometries.

Galerkin Discretization

We discretize the EFIE by a Galerkin method, which discretizes the operator equation, not merely the unknown. The density is expanded in a finite basis $\varphi_1, \dots, \varphi_N$,

$$\mathbf{j} \approx \sum_{i=1}^N j_i \varphi_i, \quad (84)$$

and the EFIE is tested against the same basis. The continuous equation $\mathcal{V}\mathbf{j} = \mathbf{h}_\times$ becomes a linear system for the coefficients,

$$\mathbf{V}\mathbf{j} = \mathbf{b} \quad b_i = \int_{\partial D} \varphi_i \cdot \mathbf{h}_\times ds, \quad (85)$$

in which the single-layer operator $\mathcal{V}$ is replaced by its Galerkin matrix $\mathbf{V}$. Its entries are the single-layer bilinear form on pairs of basis functions. Integrating the gradients of the dyadic Green's function by parts onto the basis functions reduces it to a double surface integral over the scalar fundamental solution Φ (Buffa and Hiptmair 2003, Sec. 5),

$$V_{ab} = ik \int_{\partial D} \int_{\partial D} \Phi(\mathbf{x}, \mathbf{y}) \left(\varphi_a(\mathbf{x}) \cdot \varphi_b(\mathbf{y}) - \frac{1}{k^2} \operatorname{div}_\Gamma \varphi_a(\mathbf{x}) \operatorname{div}_\Gamma \varphi_b(\mathbf{y}) \right) ds(\mathbf{y}) ds(\mathbf{x}). \quad (86)$$

The surface divergence $\operatorname{div}_\Gamma \varphi$ appears in the bilinear form, so the basis must keep it square-integrable. This requires div-conforming elements, whose normal component is continuous across element edges. Bembe1 uses div-conforming

isogeometric B-splines, the higher-order spline analogue of Raviart–Thomas elements, matching the $H^{-\frac{1}{2}}(\operatorname{div}_\Gamma, \partial D)$ trace space of $\mathcal{V}$.

Linear Solver

Bembel provides no preconditioner, such as a Calderón projector, so an iterative solver is unattractive. We rely on direct solves instead.

Furthermore we do not use any matrix compression techniques like fast-multipole or hierarchical matrices, because we need high-fidelity reference solutions to compare with the EPGP solutions. Hence we assemble the dense matrix and solve it by a direct LU factorization.

Time-Harmonic Convention

Bembel uses the opposite time-harmonic convention to ours, e^{-ikr} instead of e^{+ikr} . To reconcile this, we conjugate the Bembel output to restore the e^{+ikr} convention.

Refinement and Convergence

Convergence is governed by the polynomial degree p and the refinement level m , equivalently the mesh width h . For a smooth boundary and smooth solution, h -refinement converges algebraically and p -refinement geometrically. The ill-conditioning of the first-kind operator worsens under mesh refinement.

Chapter 4

Benchmark Results

We present results for two cavity geometries, which differ in what we can validate against. The spherical cavity admits a closed-form operator and is used to check each solver against ground truth. The ellipsoidal cavity admits none and is treated by cross-validation. Each geometry is treated in turn, covering the EPGP solver, the BEM solver, and, for the ellipsoid, a direct comparison of the two.

Two relative error measures are used throughout. The reciprocity error ρ is reference-free: reciprocity forces the true operator to be symmetric, and ρ measures the violation of that symmetry. It is a necessary but not sufficient check. The reference error ε is the distance to a trusted reference operator $\mathbf{T}_{\text{ref}}$, which is the analytic operator $\mathbf{T}_{\text{anal}}$ on the sphere and the high-fidelity BEM operator $\mathbf{T}_{\text{BEM}}$ on the ellipsoid,

$$\rho := \frac{\|\mathbf{T} - \mathbf{T}^\top\|}{\|\mathbf{T}\|} \quad \varepsilon := \frac{\|\mathbf{T} - \mathbf{T}_{\text{ref}}\|}{\|\mathbf{T}_{\text{ref}}\|}. \quad (87)$$

We use the Frobenius norm to quantify distance between operators. The operator norm is bounded by the Frobenius norm,

$$\|\mathbf{T}\|_{\text{op}} \leq \|\mathbf{T}\|_{\text{F}}, \quad (88)$$

so convergence in ε implies convergence in the operator norm (Horn and Johnson 2013, Sec. 5.6).

4.1 Spherical Cavity

The spherical cavity is a PEC sphere of radius $R = 4$. It admits a closed-form reaction operator of unlimited accuracy, which serves as ground truth for both solvers.

4.1.1 EPGP Field

We begin with a qualitative look at the EPGP solution for a single transmitter dipole at $\mathbf{z} = (0, 0, 1)$ on Λ , polarized along x . Being a probabilistic solver, the EPGP returns a full posterior over the scattered field, a Gaussian summarized by its mean and its covariance: the mean is the point estimate of the field, and the covariance is its uncertainty. We visualize both on the xz -plane at $y = 0$ through the cavity, the dipole location marked in green. The field was computed at the highest grid resolution $N_s = 1024$, $N_b = 8192$.

Mean

Figure 3 shows the posterior mean field. The top row is the real part of the x -component as a heatmap, the middle row a line-integral-convolution (LIC) texture whose streaks follow the local field direction, and the bottom row the time-averaged Poynting vector $\mathbf{S} = \frac{1}{2} \text{Re}(\mathbf{E} \times \overline{\mathbf{H}})$, which measures the flow of energy. The three columns are the incident, scattered, and total field. The incident field is the dipole near field, sharply localized at the source. Scattering off the wall produces the scattered field, and together they form the total field.

The spherical symmetry of the cavity is clearly visible in the first two rows: the scattered wavefronts are concentric and the LIC texture forms smooth rings.

The last row is a physical check on the solution. The incident field carries energy radially outward from the dipole, and the scattered field forms a recirculating pattern of its own, but in the total field the two combine to give $\mathbf{S} = 0$ throughout the cavity. This is the expected steady state of a lossless PEC cavity at a non-resonant wavenumber:

energy can neither escape through the wall nor be absorbed inside, so no net power flows and the field is a pure standing wave. The cancellation is not termwise, since S is quadratic in the field and the two parts interfere. Nothing in the solver enforces it, so recovering it is a consistency check on the posterior mean.




Figure 3: EPGP field on the spherical cavity slice.

Uncertainty

The mean is paired with a posterior uncertainty, the standard deviation of the field about that mean. Only the scattered field is inferred and thus carries uncertainty. The incident field is exact, so the total field's uncertainty is identical to the scattered field's. Figure 4 maps this standard deviation over the same slice as Figure 3. It is on the order of 10^{-3} in the chosen normalization, in which the dipole sources carry unit polarization, so it bears the same physical dimension as the field it quantifies. It forms concentric rings that grow from the center toward the wall. A possible explanation is that we condition only on the tangential part of the field, so at the wall its normal component is left free and stays uncertain, while in the interior the field equations and the plane-wave prior tie the components together.

This pattern is fixed by the conditioning geometry alone, independent of the measured data. The map therefore shows how well the boundary data determine the field, not the size of the true reconstruction error.




Figure 4: EPGP scattered-field uncertainty on the spherical cavity slice.

4.1.2 EPGP Operator

Having inspected the field, we turn to the reaction operator itself. Figure 5 shows its magnitude alongside the posterior uncertainty. The left panel is a heatmap of the magnitude $|\mathbf{T}|$, a 64×64 matrix over receiver i and transmitter j . The 64 rows and columns correspond to the 32 surface points on Λ , each carrying two tangential polarizations, so consecutive index pairs share a point on Λ . The largest entries lie along the diagonal, consistent with each dipole coupling most strongly to itself and its near neighbors. The right panel is a heatmap of the per-receiver posterior standard deviation $\sigma_i = \sqrt{\Sigma_{ii}}$, which depends on the receiver but not the transmitter.

This is because the posterior covariance is fixed only by where we condition and where we evaluate, not by the measured values. We always condition at the same boundary points for every transmitter, and only the boundary values change, so all transmitters share the same covariance. The receiver, by contrast, is the evaluation point, so the uncertainty can in general vary from receiver to receiver. On the sphere, symmetry makes all receivers equivalent, so here σ is uniform across the panel.




Figure 5: Reaction operator $|\mathbf{T}|$ and posterior uncertainty σ , spherical cavity.

Convergence

We now quantify the operator's accuracy as the resolution grows. Figure 6 plots the reference error ε against the number of spectral features N_s , one curve per boundary-point count N_b . The figure exhibits a crossover between two error regimes. For small N_s all curves coincide: the error is dominated by truncation of the spectral representation and is essentially independent of the number of boundary observations. Once the spectral approximation is sufficiently

rich, past $N_s \approx 200$, the error becomes limited by the boundary discretization, producing an N_b -dependent floor. Both axes are logarithmic, so the plot does not resolve whether the decay before that floor is algebraic or exponential in N_s . The finest curve reaches $\varepsilon \approx 1.3 \times 10^{-10}$ at $N_s = 1024$ and $N_b = 8192$.




Figure 6: EPGP convergence on the spherical cavity versus analytic reference.

Noise Influence

A free parameter of the solver is the assumed boundary noise σ_n , which regularizes the conditioning. Figure 7 sweeps it. We show the sweep for the sphere, where the true reconstruction error is available against the analytic operator. The ellipsoid behaves the same. More noise means stronger regularization, so both the reconstruction error and the predicted uncertainty grow. On the log-log axes the growth is nearly a straight line, indicating a power-law dependence on σ_n .

Both rise because σ_n drives them: it inflates the predicted variance directly and, through stronger regularization, also worsens the reconstruction. Their covariation does not mean that uncertainty tracks the reconstruction error. The predicted uncertainty is thus governed largely by the assumed noise, which is what makes its calibration delicate.




Figure 7: Reconstruction error and predicted uncertainty versus assumed noise, spherical cavity.

4.1.3 BEM Operator

The BEM is deterministic, so it returns the operator without an uncertainty estimate. We assess its accuracy by convergence against the analytic reference.

Convergence

Figure 8 splits this convergence into two panels, each a family of curves over one fixed parameter. The left panel performs h -refinement, varying the mesh level m at each fixed polynomial degree p , plotted against degrees of freedom on log-log axes. Algebraic convergence appears as a straight line there, and each curve is one, with a slope that steepens with p . A $\propto N^{-3/2}$ reference line is included for comparison. The right panel performs p -refinement, varying p at each fixed mesh level, plotted against p on a semilog axis. Geometric convergence appears as a straight line there, and each curve is one. The two refinements thus behave as the theory predicts, and the finest resolutions reach $\varepsilon \approx 3 \times 10^{-12}$. The only exception is the $p = 5, m = 4$ point, which ticks back up as it reaches the floor of attainable accuracy.




Figure 8: BEM convergence on the spherical cavity versus analytic reference.

4.1.4 Wavenumber Sweep

The benchmark fixes $k = 2$. A cavity resonance, where the interior boundary value problem is not uniquely solvable, would make the reconstruction ill-posed, so we check that $k = 2$ avoids one. We detect resonances with the subspace-angle diagnostic (Betcke and Trefethen 2005). At wavenumber k we evaluate the F plane-wave features both on the boundary and in the interior: the matrix $\Phi_b \in \mathbb{C}^{3N_b \times F}$ collects their tangential traces at N_b boundary points, and $\Phi_i \in \mathbb{C}^{6N_i \times F}$ their full six-component fields at N_i interior points. We orthonormalize the columns of the stacked matrix by a thin QR factorization and partition the orthonormal factor conformally,

$$\begin{bmatrix} \Phi_b \\ \Phi_i \end{bmatrix} = \mathbf{Q}\mathbf{R}, \quad \mathbf{Q} = \begin{bmatrix} \mathbf{Q}_b \\ \mathbf{Q}_i \end{bmatrix}, \quad (89)$$

then record the smallest singular value of the boundary block,

$$\sigma_{\min} := \sigma_{\min}(\mathbf{Q}_b). \quad (90)$$

Including the interior block normalizes the boundary trace against the field's interior magnitude, so $\sigma_{\min}$ measures how nearly some non-trivial interior field attains a zero boundary trace. It therefore collapses toward zero at a cavity resonance and stays bounded away from it otherwise.

Figure 9 shows the sweep on the sphere over $k \in [1.5, 2.5]$. Every dip of $\sigma_{\min}$ coincides with an analytic resonance of the sphere, shown dashed, so the diagnostic locates resonances correctly. The benchmark wavenumber $k = 2$ sits at a local maximum, far from any resonance, confirming that it is non-resonant.




Figure 9: Wavenumber sweep, spherical cavity.

4.2 Ellipsoidal Cavity

The ellipsoidal cavity has semi-axes $(4, 4, 6)$, keeping the same interior surface Λ and wavenumber $k = 2$. Unlike the sphere, it admits no analytic operator, so the high-fidelity BEM solution serves as the reference.

4.2.1 EPGP Field

We repeat the qualitative inspection of the previous section, visualizing the EPGP posterior for the same transmitter dipole at $z = (0, 0, 1)$ on Λ , polarized along x , on the same xz -plane through the cavity. The elongated geometry breaks the spherical symmetry, and the contrast with the sphere is the point of interest.

Mean

Figure 10 shows the posterior mean field, the real part of the x -component as a heatmap on top and the LIC texture below, for the incident, scattered, and total field. The broken symmetry is clearly visible. The incident field is the same as before, since it is the same dipole, but the scattered field no longer shows the concentric wavefronts of the sphere. Instead it has a more intricate pattern, stretched along the long axis of the cavity, and the same is seen in the LIC texture. The total field is again the sum of the two.




Figure 10: EPGP field on the ellipsoidal cavity slice.

Uncertainty

Figure 11 maps the posterior uncertainty over the same slice. The broken symmetry makes it richer than on the sphere: instead of concentric rings the standard deviation forms separate off-axis lobes, and its scale is an order of magnitude larger, of order 10^{-2} . As on the sphere, the pattern is fixed by the conditioning geometry alone and reports how well the boundary data determine the field rather than the true reconstruction error.




Figure 11: EPGP scattered-field uncertainty on the ellipsoidal cavity slice.

4.2.2 EPGP Operator

As on the sphere, Figure 12 shows the operator magnitude beside the posterior uncertainty. The magnitude $|\mathbf{T}|$ keeps the same diagonal structure of strong self- and near-coupling. The uncertainty, however, is no longer uniform: the elongated geometry breaks the equivalence of receivers, so σ now varies from receiver to receiver, visible as the horizontal banding. It grows for receivers deeper inside the cavity, toward the elongated z -axis and away from the wall, suggesting that the boundary data constrains receivers close to the wall more strongly than those deeper inside the cavity. As in the field maps, this uncertainty is set by where we condition and evaluate, not by the measured data.




Figure 12: Reaction operator $|\mathbf{T}|$ and posterior uncertainty σ , ellipsoidal cavity.

Convergence

With no analytic operator available, we track accuracy through the reciprocity error ρ , which needs no reference. Figure 13 plots ρ against the number of spectral features N_s . It decreases steadily and floors at $\rho \approx 3 \times 10^{-11}$, confirming that the EPGP operator is symmetric to that level. The cross-validation reference error ε against the BEM reference is reported in the table below. Table 2 lists the corresponding runs at the finest N_b , together with their runtime, memory, and conditioning.




Figure 13: EPGP reciprocity error on the ellipsoidal cavity versus N_s .

4.2.3 BEM Operator

The BEM operator is the reference for this geometry. As for the EPGP, the absence of an analytic operator leaves only the reciprocity error ρ as a self-contained measure.

Convergence

Figure 14 plots the BEM reciprocity error ρ over the $p \times m$ grid, and Table 1 lists every run with its degrees of freedom, runtime, memory, and conditioning. Here only ρ is available, which sees the antisymmetric part of the error and so cannot certify a convergence order, but its decay is consistent with the algebraic h -refinement and geometric p -refinement established on the sphere. The most refined run, $p = 5$ and $m = 4$ with 4800 degrees of freedom, reaches $\rho \approx 1.4 \times 10^{-10}$ and serves as the reference operator $\mathbf{T}_{\text{BEM}}$ for this geometry.




Figure 14: BEM reciprocity error on the ellipsoidal cavity.

4.2.4 Wavenumber Sweep

The k -sweep on the ellipsoid is shown in Figure 15. Unlike the sphere the benchmark wavenumber sits close to one: $k = 2$ lies only $\Delta k \approx 0.0018$ above an interior resonance near $k \approx 1.998$, where $\sigma_{\min}$ dips to $\approx 1.3 \times 10^{-3}$. The dip is small but nonzero, so the boundary value problem remains solvable and both solvers converge cleanly. The benchmark is thus operationally non-resonant, though by a narrow margin.




Figure 15: Wavenumber sweep, ellipsoidal cavity.

4.2.5 Comparison

Having validated each solver on its own, we now compare them directly, first in accuracy and then in cost.

Cross-Validation

The two solvers share no numerical machinery, so their agreement is strong evidence for both. We quantify it by the reference error ε of the EPGP operator against the BEM reference $\mathbf{T}_{\text{BEM}}$. It decreases monotonically with N_s , reaching $\varepsilon \approx 9 \times 10^{-9}$ at $N_s = 1024$, and is still decreasing there, so 9×10^{-9} is an upper bound on the relative difference: the two operators are at least this close.

Together with the independent certification of each solver on the analytic sphere, this is the central result of the benchmark. Two methodologically unrelated solvers, each validated on its own and then shown to agree on a geometry with no ground truth, give strong combined evidence that both compute the correct reaction operator.

Accuracy–Runtime Trade-off

All runs were carried out on the Euler cluster³, each on a single exclusive AMD EPYC 7742 node with 128 cores, and wall time was recorded per run. For the EPGP the recorded time includes the JAX just-in-time compilation and the Python startup, so it is a conservative measure.

Figure 16 plots reciprocity error against wall time for both solvers. The faint points are all grid runs and the solid line is each solver’s Pareto front. The two fronts have very different shapes. The EPGP front is nearly vertical, between ≈ 6 and 18 s: its reciprocity error improves by orders of magnitude at almost fixed wall time, because the cost is dominated by the one-off factorization. The BEM front is a staircase that descends only with large increases in wall time, reaching $\rho \approx 1.4 \times 10^{-10}$ at ≈ 1600 s at the finest grid run, with $p = 5$ and $m = 4$.

At the same reciprocity error the EPGP is two to three orders of magnitude faster, reaching $\rho \approx 10^{-9}$ to 10^{-10} in about 10 s. The BEM is cheaper only at loose tolerance. For any demanding accuracy the EPGP is the cheaper solver.

These wall times reflect two particular implementations, a compiled C++ solver and a Python/JAX one, so the absolute factor mixes implementation quality with algorithmic cost. The implementation-independent comparison is asymptotic. Both solvers factor their system once and reuse it across the $M = 64$ dipole right-hand sides. The dense

³Euler is the central high-performance computing cluster of ETH Zürich, see scicomp.ethz.ch/wiki/Euler.

BEM factorization costs $\mathcal{O}(N^3)$ in the $N \sim p^2 4^m$ surface degrees of freedom, while the EPGP system assembly costs $\mathcal{O}(N_s^2 N_b)$ and its factorization $\mathcal{O}(N_s^3)$. Since $N_b > N_s$ in practice, assembly dominates.




Figure 16: Reciprocity error versus wall time for BEM and EPGP on the ellipsoidal cavity.

p	1	1	1	1	2	2	2	2	3	3	3	3	4	4	4	4	5	5	5	5
m	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4	1	2	3	4
DOFs	48	192	768	3072	108	300	972	3468	192	432	1200	3888	300	588	1452	4332	432	768	1728	4800
mem [GiB]	0	0	0.3	3.2	0	0.1	0.9	13.5	0	0.2	2.7	39.7	0	0.4	6.2	94	0.1	0.9	12.5	191.7
t [s]	0	0	4	11	0	0	3	36	0	1	10	159	1	3	32	527	1	8	104	1593
cond $\mathbf{V}$	7.38×10^1	4.15×10^3	8.54×10^4	1.91×10^4	1.58×10^3	1.56×10^4	2.37×10^4	2.91×10^4	1.27×10^4	1.44×10^5	5.78×10^4	1.59×10^5	7.46×10^4	1.08×10^5	3.44×10^5	1.07×10^6	3.07×10^6	8.42×10^5	2.07×10^6	6.33×10^6
ρ	1.65×10^{-1}	6.54×10^{-2}	5.41×10^{-3}	4.78×10^{-4}	9.93×10^{-2}	5.13×10^{-2}	2.66×10^{-4}	4.65×10^{-6}	5.25×10^{-2}	6.27×10^{-3}	3.06×10^{-5}	7.54×10^{-8}	5.35×10^{-2}	8.07×10^{-3}	7.45×10^{-6}	2.01×10^{-9}	5.76×10^{-3}	4.86×10^{-4}	1.83×10^{-6}	1.37×10^{-10}

Table 1: BEM convergence on the ellipsoidal cavity.

N_s	16	32	64	128	192	256	384	512	768	1024
features	32	64	128	256	384	512	768	1024	1536	2048
t [s]	7	6	8	8	8	8	8	9	9	18
mem [GiB]	0.7	0.7	0.9	0.9	1	1	1.3	1.6	2.1	2.7
cond $\mathbf{A}$	3.08×10^0	6.58×10^0	3.98×10^1	4.43×10^2	4.12×10^4	1.17×10^7	2.7×10^{10}	5.51×10^{10}	8.27×10^{10}	1.1×10^{11}
ρ	4.88×10^{-1}	3.11×10^{-1}	1.57×10^{-1}	9.23×10^{-4}	1.08×10^{-4}	8.29×10^{-7}	1.15×10^{-10}	5.29×10^{-11}	3.87×10^{-11}	3.24×10^{-11}
ε	9.55×10^{-1}	8.73×10^{-1}	6.87×10^{-1}	6.63×10^{-3}	3.66×10^{-3}	5.54×10^{-6}	2.24×10^{-8}	1.7×10^{-8}	1.16×10^{-8}	8.95×10^{-9}

Table 2: EPGP convergence on the ellipsoidal cavity ($N_b = 8192$).

Chapter 5

Conclusion and Outlook

5.1 Summary

We established a high-fidelity benchmark for the Maxwell EPGP method. The cavity reaction operator was computed by two methodologically unrelated solvers, a spectral plane-wave EPGP and a boundary-element BEM, and the two were compared.

On the spherical cavity, where the analytic operator is available, BEM matches the ground truth to $\approx 3 \times 10^{-12}$ and EPGP to $\approx 1.3 \times 10^{-10}$ in relative Frobenius-norm error ε , certifying each independently. On the ellipsoidal cavity, where no analytic solution exists, the two solvers agree to about 10^{-8} in the same relative Frobenius norm. Since both were already certified on the sphere, this agreement carries over to a geometry with no ground truth. The agreement of a boundary-integral and a spectral plane-wave method, which share no numerical machinery, validates the EPGP on this geometry and certifies the reaction operator as a robust reference benchmark.

Beyond matching the reference, the EPGP satisfies the time-harmonic Maxwell equations exactly by construction, needs no surface mesh, and provides a posterior uncertainty estimate as a byproduct. It is also the cheaper solver at any demanding tolerance, reaching a given reciprocity error two to three orders of magnitude faster than the BEM in our experiments. This speed comparison should be read with the caveats in the Limitations: the BEM was run unoptimized, without fast-multipole acceleration or matrix compression, so the absolute factor mixes implementation quality with algorithmic cost. The benchmark thus establishes the EPGP as a Maxwell-consistent, uncertainty-aware surrogate for the interior cavity reaction operator.

5.2 Limitations

5.2.1 Validation Scope

The validation rests most firmly on the spherical cavity, where each solver is checked directly against the analytic operator. On the ellipsoidal cavity there is no ground truth, so the evidence is the agreement between the two unrelated solvers, which is mutual consistency rather than a direct measure of accuracy. The reciprocity error cannot close this gap, since it constrains only the antisymmetric part of the error. Splitting the error operator $\mathbf{T} - \mathbf{T}_{\text{ref}}$ into a symmetric and an antisymmetric part, the reciprocity error $\rho = \|\mathbf{T} - \mathbf{T}^T\|/\|\mathbf{T}\|$ depends only on the antisymmetric part: the true operator is symmetric, so a symmetric component of the error leaves the operator symmetric and the reciprocity error does not see it. A solver could therefore carry a large symmetric error while its reciprocity error stays small, so a small reciprocity error is necessary but not sufficient for accuracy.

Both test geometries, and the dipole surface Λ , are fairly symmetric. A defect that appears only on strongly asymmetric domains would therefore not be caught here, though we have no specific reason to expect one.

The benchmark is also confined to a single wavenumber and geometry pair. A wavenumber sweep confirms that $k = 2$ avoids the cavity resonances, but the accuracy itself is validated only at $k = 2$.

5.2.2 Cost Comparison

The cost comparison should be read with care. Both solvers ran on identical hardware, but the absolute wall times reflect two particular implementations, a compiled C++ solver and a Python/JAX one, so the runtime factor mixes implementation quality with algorithmic cost. The implementation-independent statement is the asymptotic scaling.

Several factors could shift the comparison in either direction. The EPGP wall time includes the JAX compilation and Python startup, which likely make up its near-constant offset, so timing only the numerical solve would isolate the algorithmic cost. The EPGP also ran on CPU, while JAX can target GPUs, which we did not test and which could lower its cost further. On the other side, the BEM is unoptimized: it uses a dense direct solve with no iterative solver, preconditioner, or matrix compression for fidelity, so a tuned BEM could be considerably faster. The two solvers also parallelize differently, and we did not study how the comparison scales with core count. A fairer comparison, controlling for these factors and counting operations or memory rather than wall time, is left to future work.

5.2.3 Uncertainty Quantification

The uncertainty quantification is descriptive and correctly ranks where the field is well or poorly determined by the boundary data, growing in the regions the data constrains least. However what it does not represent is the true reconstruction error. This is because the dominant error comes from truncating the plane-wave expansion at a finite number of spectral features, which is a systematic bias, and a Gaussian variance cannot represent a bias.

The scale of the uncertainty is set by the assumed noise σ_n , so one might try to calibrate it by tuning σ_n to the data. This does not help. Optimizing σ_n by maximizing the marginal likelihood drives it toward a floor rather than toward the true error level. The plane-wave features are overcomplete and strongly correlated, so as σ_n decreases the precision matrix $\mathbf{A} = \mathbf{W}^{-1} + \Phi\Phi^H/\sigma_n^2$ becomes increasingly dominated by the ill-conditioned feature term $\Phi\Phi^H/\sigma_n^2$. Its conditioning then approaches that of the feature Gram matrix $\Phi\Phi^H$, whose spectrum sets the limiting condition number, rather than following a clean $1/\sigma_n^2$ law. Below a certain noise level the system becomes too ill-conditioned to solve stably, and the floating-point floor, rather than the data, determines the smallest usable σ_n . A more principled quadrature might relieve this, though better integration accuracy does not by itself guarantee a better-conditioned feature matrix.

5.3 Future Work

We see two natural directions to extend this work, one that sharpens the EPGP’s spectral discretization and one that changes how the reaction operator is assembled.

5.3.1 Spherical Designs and Lebedev Quadrature

The EPGP draws its wavevector directions from a Fibonacci sphere, a quasi-Monte Carlo rule. A more principled choice is the Lebedev quadrature (Lebedev and Laikov 1999), which integrates spherical harmonics exactly up to a chosen degree. We expect it to reach the same kernel accuracy with fewer directions, giving a smaller system. Whether it also yields a better-conditioned system is an open question: good integration accuracy does not by itself imply a well-conditioned feature matrix, since the conditioning is governed by the correlations among the plane-wave features rather than by the accuracy of the quadrature. Should the conditioning improve, the lower condition number would in turn let the assumed noise drop further, which is the main obstacle to calibrated uncertainty.

5.3.2 Operator-Learning Perspective

The reaction operator is currently assembled one column at a time. Each transmitter dipole is a separate conditioning of the same shared prior, solved independently. An operator-learning view would instead treat the reaction operator itself as the object to learn directly, producing the whole operator at once instead of rebuilding each column in isolation.

Appendix A

Implementation Source Code

All software for this thesis is open-source on GitHub. Each repository carries a README documenting its purpose, setup, and usage, so the code remains the single source of truth and this thesis does not restate implementation details. The numerical results reported here are reproducible from these repositories through the scripts described in their READMEs.

To pin the thesis to a fixed state, the links below point to the semester-thesis git tag of each repository. Later development continues on the main branch without affecting these references.

A.1 maxwellgp

The general Maxwell-constrained EPGP library: plane-wave feature maps, the GP regression core, and tangential-trace conditioning.

[github:luiswirth/maxwellgp](https://github.com/luiswirth/maxwellgp)

A.2 cavity-maxwellgp

A thin cavity-specific layer over maxwellgp, owning the analytic dipole physics and the reaction-operator assembly.

[github:luiswirth/cavity-maxwellgp](https://github.com/luiswirth/cavity-maxwellgp)

A.3 cavity-bem

The deterministic boundary-element reference solver, built on Bembel.

[github:luiswirth/cavity-bem](https://github.com/luiswirth/cavity-bem)

A.4 cavity-benchmark

Performs post-processing on the BEM and EPGP operators. Compares to the analytic sphere reference and cross-validates the two solvers on the ellipsoid. Generates all visualizations, such as plots and animations.

[github:luiswirth/cavity-benchmark](https://github.com/luiswirth/cavity-benchmark)

Appendix B

Document Source Code

The thesis document has been written using the new type-setting language Typst, see typst.app.

The source code for this Typst document is open-source and available on GitHub. As above, the link points to the `semester-thesis` tag, the submitted version. Any later corrections appear on the main branch.

[github:luiswirth/cavity-thesis](https://github.com/luiswirth/cavity-thesis)

To compile the document into a PDF, the Typst compiler needs to be installed and then the `./build.sh` script can be run from the root project directory.

Appendix C

Use of Generative AI

Generative AI assistants were used as a productivity aid in two places.

- Software development: drafting, refactoring, and debugging parts of the accompanying code (`maxwellgp`, `cavity-maxwellgp`, `cavity-bem`, and `cavity-benchmark`), including the plotting and benchmark harness.
- Writing: drafting, restructuring, and copy-editing the prose of this document, and the Typst typesetting.

All numerical results, figures, and data were produced by the solvers and scripts described in the preceding appendices, not by an AI tool. Every AI-assisted output was reviewed and, where necessary, corrected by the author, who takes full responsibility for the content of this thesis. No factual or physical claim rests on an AI tool without verification against the cited literature or the computed results.

Bibliography

- Betcke, Timo, and Lloyd N. Trefethen. 2005. "Reviving the Method of Particular Solutions." *SIAM Review* 47 (3): 469–91. <https://doi.org/10.1137/S0036144503437336>.
- Buffa, Annalisa, and Ralf Hiptmair. 2003. "Galerkin Boundary Element Methods for Electromagnetic Scattering." In *Topics in Computational Wave Propagation: Direct and Inverse Problems*, vol. 31 of *Topics in Computational Wave Propagation: Direct and Inverse Problems*. Lecture Notes in Computational Science and Engineering. Springer. https://doi.org/10.1007/978-3-642-55483-4_3.
- Colton, David, and Rainer Kress. 2019. *Inverse Acoustic and Electromagnetic Scattering Theory*. 4th ed. Vol. 93, pt. 93. Applied Mathematical Sciences. Springer.
- Dölz, Jürgen, Helmut Harbrecht, Stefan Kurz, Michael D. Multerer, Sebastian Schöps, and Felix Wolf. 2020. "Bembel: The Fast Isogeometric Boundary Element C++ Library for Laplace, Helmholtz, And Electric Wave Equation." *Software* 11 : 100476.
- Ehrenpreis, Leon. 1970. *Fourier Analysis in Several Complex Variables*. Wiley-Interscience.
- González, Álvaro. 2010. "Measurement of Areas on a Sphere Using Fibonacci and Latitude–Longitude Lattices." *Mathematical Geosciences* 42 (1): 49–64. <https://doi.org/10.1007/s11004-009-9257-x>.
- Härkönen, Marc, Markus Lange-Hegermann, and Bogdan Raiță. 2023. "Gaussian Process Priors for Systems of Linear Partial Differential Equations with Constant Coefficients." In "International Conference on Machine Learning." Special issue, *International Conference on Machine Learning*, 12587–615.
- Horn, Roger A., and Charles R. Johnson. 2013. *Matrix Analysis*. 2nd ed. Cambridge University Press.
- Jackson, John David. 1999. *Classical Electrodynamics*. 3rd ed. Wiley.
- Lebedev, V. I., and D. N. Laikov. 1999. "A Quadrature Formula for the Sphere of the 131st Algebraic Order of Accuracy." *Doklady Mathematics* 59 (3): 477–81.
- Rasmussen, Carl Edward, and Christopher K. I. Williams. 2006. *Gaussian Processes for Machine Learning*. MIT Press.
- Tai, Chen-To. 1994. *Dyadic Green Functions in Electromagnetic Theory*. 2nd ed. IEEE Press.
- Vittori, Félix, and Stefan Kurz. 2026. "Learning Electromagnetic Fields: Gaussian Process Priors for Maxwell's Equations." Unpublished manuscript.
- Zeng, Fang, Fioralba Cakoni, and Jiguang Sun. 2011. "An Inverse Electromagnetic Scattering Problem for a Cavity." *Inverse Problems* 27 (12): 125002. <https://doi.org/10.1088/0266-5611/27/12/125002>.