Quantitative diagnostics of wave–wave electromagnetic coupling channels: a unified FDTD framework with Kerr, plasma, and Heisenberg–Euler regimes, with application to magnetar magnetospheres

Michael F. Golden

June 4, 2026

Abstract

We present a unified two-dimensional finite-difference time-domain (FDTD) framework for studying nonlinear wave–wave coupling in electromagnetic media. The framework supports four physical regimes through a modular constitutive interface: linear vacuum (validation baseline), instantaneous Kerr $\chi^{(3)}$ media, cold relativistic electron plasma, and one-loop quantum electrodynamics in both the $\mathcal{F}^2$ (scalar) and $\mathcal{G}^2$ ($\mathbf{E}\cdot\mathbf{B}$) Heisenberg–Euler channels. We introduce a quantitative diagnostic — the linearity residual $\|E_{AB} - E_A - E_B\|_\infty / \|E_{AB}\|_\infty$ — that isolates wave–wave cross-coupling from per-wave self-action and provides a regime-agnostic measure of nonlinearity. Validation against the linear-vacuum limit recovers machine-precision superposition (residual $\sim 10^{-15}$). Methodological advances include a relativistic Boris pusher that eliminates the secular numerical heating that plagued forward-Euler integration of the cold plasma fluid (energy growth reduced from $\sim 40\times$ to $\lesssim 10^{-14}$ relative over 200 cyclotron periods), a predictor–corrector solver for the implicit Heisenberg–Euler constitutive relation (single-wave self-action artifact reduced by a factor of ~ 8), and a convolutional perfectly-matched-layer absorbing boundary validated to give $\lesssim 10^{-3}$ reflection at oblique incidence. As a physics application, we model wave–wave coupling in a strong uniform magnetic background representative of magnetar magnetospheres. Across a finely resolved scan of the background amplitude ($0 \le B_{\rm bg} \le 2.25$ in natural units), the linearity residual is well described by a three-channel polynomial $R - R_0 = \alpha B_{\rm bg} + \beta B_{\rm bg}^2 + \gamma B_{\rm bg}^3$ with statistical evidence (76% RSS reduction over a two-channel fit) for a cubic-in-$B_{\rm bg}$ contribution that exceeds the leading-order Heisenberg–Euler $\mathcal{F}^2$-channel prediction and may indicate either two-loop QED corrections or a numerical artifact requiring further study. In a two-dimensional mixed-polarization extension supporting both TM and TE modes, we demonstrate that the $\mathcal{G}^2$ photon–photon scattering channel produces wave–wave coupling 6.8 times stronger than the $\mathcal{F}^2$ scalar channel per unit Lagrangian coupling constant for cross-polarized incident waves, validating the cross-polarized geometry as the cleanest "wave–wave only, no self-action" QED diagnostic available in this dimensional setting. We discuss implications for thesis-quality studies of binary neutron-star inspirals and strong-field laboratory laser experiments, and we identify the development priorities required to push the framework into the genuinely strong-field regime ($B_{\rm bg} \gtrsim B_{\rm S}$).

1 Introduction

Wave–wave interaction is a defining nonlinear feature of electromagnetic physics in non-vacuum or extreme-field environments. In vacuum at moderate fields, Maxwell's equations are exactly linear and two propagating electromagnetic waves pass through one another without modification — a fact directly verified to machine precision in the present work and routinely exploited in classical optics. Nonlinearity appears through three qualitatively distinct mechanisms in physically realized media:

1. **Polarizable media**, in which the dielectric response acquires intensity-dependent corrections beyond Hookean polarizability. The lowest-order nontrivial coupling is the cubic Kerr term, $\mathbf{P} = \varepsilon_0 \chi^{(3)} |\mathbf{E}|^2 \mathbf{E}$, which drives self-phase modulation, third-harmonic generation, four-wave mixing, and optical solitons. The $\chi^{(3)}$ coefficient is the dominant nonlinearity in glasses and many crystals and is the laboratory workhorse for nonlinear-optics experiments [1].

2. **Plasmas**, in which free charge carriers respond to electromagnetic fields through a coupled fluid (or kinetic) system. The leading wave–wave coupling mechanisms are the relativistic mass-correction in $\gamma = \sqrt{1 + p^2/m^2c^2}$, the ponderomotive density modulation $\delta n/n_0 \sim a_0^2$ at the second-harmonic of the wave envelope, and parametric instabilities such as stimulated Raman scattering (SRS) and stimulated Brillouin scattering (SBS) [2]. These effects dominate laser–plasma interaction at $a_0 \equiv eE/(m_e\omega c) \gtrsim 0.1$ and are central to inertial-confinement fusion and laser wakefield acceleration.

3. **Vacuum quantum electrodynamics** at extreme field strengths, where the photon–photon coupling mediated by virtual electron–positron pairs gives rise to nonlinear corrections to the effective Lagrangian. The one-loop Heisenberg–Euler Lagrangian [3, 4] is
$$\mathcal{L}_{\mathrm{HE}} = -\mathcal{F} + \frac{2\alpha^2}{45 m_e^4} \left(4\mathcal{F}^2 + 7\mathcal{G}^2\right) + \mathcal{O}\left(\mathcal{F}^3\right), \tag{1}$$
where $\mathcal{F} \equiv \frac{1}{2}(E^2 - B^2)$ and $\mathcal{G} \equiv \mathbf{E} \cdot \mathbf{B}$ are the two Lorentz invariants. The $\mathcal{F}^2$ and $\mathcal{G}^2$ channels contribute to vacuum birefringence, photon–photon scattering, and pair production near the Schwinger critical field $B_{\mathrm{S}} = m_e^2 c^2/e\hbar \approx 4.4 \times 10^9$ T [5, 6]. Although QED nonlinearity is utterly negligible at laboratory ambient fields, it becomes O(1) near magnetar surfaces and is the presumed origin of the polarization signal measured for the isolated neutron star RX J1856−3754 [7].

A common feature of these mechanisms is that the leading wave–wave coupling vanishes identically in the absence of a second wave. This makes wave–wave interaction intrinsically a multi-wave phenomenon, and its theoretical analysis is naturally framed in terms of cross-coupling polarization or current densities. Such cross-couplings are small in most regimes of physical interest, demanding either long propagation distances (astrophysics) or extreme intensities (laser laboratories) to become observable.

For numerical studies, this perturbative smallness implies that wave–wave physics is easily contaminated by per-wave self-action and other numerical artifacts. A widely used strategy is to set up the simulation with a single wave packet and measure deviations from the linear evolution; however, this approach conflates wave–wave coupling with single-wave nonlinearity (self-phase modulation in Kerr, ponderomotive density modulation in plasma, envelope-induced spurious nonlinearity in QED). In this work, we adopt instead the *linearity residual* diagnostic
$$R(t) \equiv \frac{\|E_{AB}(t) - E_A(t) - E_B(t)\|_\infty}{\|E_{AB}(t)\|_\infty}, \tag{2}$$
where E_A, E_B, and E_{AB} are the solutions to the full nonlinear equations with only wave A, only wave B, and both A and B respectively as initial conditions. By construction, $R(t)$ vanishes identically in a linear medium and isolates terms in the field evolution that depend on both waves simultaneously. We argue that this diagnostic is regime-agnostic — it has the same operational meaning in Kerr media, in plasma, and in QED-corrected vacuum — and that quantitative scaling laws for $R(t)$ provide a stronger validation of theoretical predictions than the conventional single-wave self-action measure.

The remainder of this paper develops the unified FDTD framework (Sec. 3), presents validation against the linear-vacuum limit and against per-regime physical expectations (Sec. 4),

reports quantitative results across all four physics regimes (Sec. 5), and applies the framework to a magnetar-inspired strong-background scenario (Secs. 6 and 7). Section 8 discusses the interpretation of an unexpected cubic-in-B_{bg} scaling channel observed in the magnetar scan, and Sec. 9 summarizes implications and identifies the development priorities required to push the framework into the genuinely strong-field regime.

2 Theoretical background

We work in two spatial dimensions with $\partial/\partial z = 0$ and adopt natural units throughout ($c = \varepsilon_0 = \mu_0 = m_e = e = 1$). Conversion to SI units is straightforward: a length unit in the natural system corresponds to a chosen physical length (we typically interpret the unit as $\mathcal{O}(\mu\mathrm{m})$ for laser-plasma applications or $\mathcal{O}(\mathrm{km})$ for magnetar applications), and the field strength normalization follows by dimensional consistency.

2.1 Vacuum Maxwell equations

In SI units with no sources and no medium polarization or magnetization, Maxwell's equations are

$$\boldsymbol{\nabla}\cdot\mathbf{E} = 0, \quad \boldsymbol{\nabla}\cdot\mathbf{B} = 0, \quad \boldsymbol{\nabla}\times\mathbf{E} = -\frac{\partial\mathbf{B}}{\partial t}, \quad \boldsymbol{\nabla}\times\mathbf{B} = \mu_0\varepsilon_0\frac{\partial\mathbf{E}}{\partial t}. \tag{3}$$

For the TM_z polarization with nonzero E_z, H_x, H_y and all z-derivatives vanishing, the curl operators reduce to

$$\frac{\partial H_x}{\partial t} = -\frac{\partial E_z}{\partial y}, \quad \frac{\partial H_y}{\partial t} = +\frac{\partial E_z}{\partial x}, \quad \frac{\partial E_z}{\partial t} = +\frac{\partial H_y}{\partial x} - \frac{\partial H_x}{\partial y}. \tag{4}$$

For mixed-polarization studies (Sec. 7) we additionally retain E_x, E_y and H_z, with the corresponding TE-mode updates

$$\frac{\partial H_z}{\partial t} = -\left(\frac{\partial E_y}{\partial x} - \frac{\partial E_x}{\partial y}\right), \quad \frac{\partial E_x}{\partial t} = +\frac{\partial H_z}{\partial y}, \quad \frac{\partial E_y}{\partial t} = -\frac{\partial H_z}{\partial x}. \tag{5}$$

The TM and TE modes decouple identically in linear vacuum under our 2D reduction; they couple only through nonlinear constitutive relations.

2.2 Kerr $\chi^{(3)}$ medium

The instantaneous Kerr response is captured by the constitutive relation

$$\mathbf{D} = \varepsilon_r\mathbf{E} + \chi^{(3)}|\mathbf{E}|^2\mathbf{E}, \tag{6}$$

which in our 2D TM_z mode reduces to the scalar relation $D_z = \varepsilon_r E_z + \chi^{(3)}E_z^3$. The corresponding cubic

$$\chi^{(3)}E_z^3 + \varepsilon_r E_z - D_z = 0 \tag{7}$$

is solved by Newton iteration at each time step.

2.3 Cold relativistic electron plasma

We treat the electron component as a continuous relativistic fluid with continuity and momentum equations

$$\frac{\partial n}{\partial t} + \boldsymbol{\nabla}\cdot(n\mathbf{v}) = 0, \tag{8}$$

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = -e(\mathbf{E} + \mathbf{v}\times\mathbf{B}), \tag{9}$$

where $\mathbf{v} = \mathbf{p}/(\gamma m_e)$ with $\gamma = \sqrt{1 + |\mathbf{p}|^2/(m_e c)^2}$. The current density $\mathbf{J} = -en\mathbf{v}$ feeds back into Ampère's law to close the coupled Maxwell-fluid system. Ions are treated as a uniform neutralizing background. The plasma frequency in our natural units is $\omega_p = \sqrt{n_0}$.

2.4 Heisenberg–Euler vacuum

The weak-field expansion of the Heisenberg–Euler Lagrangian (1) gives rise to the constitutive relations

$$\mathbf{D} = \mathbf{E}(1 + 2\kappa_F \mathcal{F}) + 2\kappa_G \mathcal{G}\, \mathbf{B}, \tag{10}$$

$$\mathbf{H} = \mathbf{B}(1 + 2\kappa_F \mathcal{F}) - 2\kappa_G \mathcal{G}\, \mathbf{E}, \tag{11}$$

where κ_F and κ_G subsume the prefactors of the $\mathcal{F}^2$ and $\mathcal{G}^2$ terms in (1) and are treated as free parameters in our framework. The physical values $\kappa_F = 8\alpha^2/(45m_e^4)$ and $\kappa_G = 14\alpha^2/(45m_e^4)$ produce experimentally negligible effects at laboratory fields but become $\mathcal{O}(1)$ corrections at field strengths comparable to B_{S}.

For TM_z-only simulations $\mathcal{G} = \mathbf{E} \cdot \mathbf{B} = E_z \cdot 0 = 0$ identically (since E is in $\hat{z}$ and B is in-plane), reducing the effective Lagrangian to $\mathcal{L} = \frac{1}{2}(E^2 - B^2) + \kappa \mathcal{F}^2$ with $\kappa \equiv \kappa_F$; the constitutive relation becomes the scalar cubic

$$D_z = E_z \left[1 + 4\kappa(E_z^2 - |\mathbf{B}|^2)\right]. \tag{12}$$

The mixed-polarization case (Sec. 7) requires the full vector constitutive (10)–(11).

2.5 The linearity residual diagnostic

A direct calculation gives, for two plane waves with envelopes $\mathbf{E}_A, \mathbf{E}_B$ on a static background B_{bg} in the Heisenberg–Euler $\mathcal{F}^2$ channel, the cross-coupling polarization

$$\mathbf{P}_{\mathrm{cross}} = 4\kappa \Big[-2\mathbf{B}_{\mathrm{bg}} \cdot (E_A \mathbf{B}_B + E_B \mathbf{B}_A) + (E_A \cdot E_B - \mathbf{B}_A \cdot \mathbf{B}_B)(E_A + E_B) \Big], \tag{13}$$

which contains a term linear in B_{bg} and a term independent of B_{bg}. The residual (2) is sensitive to the time-integrated effect of $\mathbf{P}_{\mathrm{cross}}$ on the field evolution and therefore inherits both channels with geometry-dependent coefficients. The same structural form (a polynomial in the background amplitude with leading-order channels determined by the Lagrangian) applies to plasma and Kerr regimes, though with different geometric factors.

3 Numerical framework

3.1 Yee grid and leapfrog time integration

We discretize the fields on a 2D Yee staggered grid [8]: E_z at integer cell corners (i, j), H_x at $(i, j + \frac{1}{2})$, H_y at $(i + \frac{1}{2}, j)$, and for the mixed-polarization extension H_z at $(i + \frac{1}{2}, j + \frac{1}{2})$, E_x at $(i + \frac{1}{2}, j)$, E_y at $(i, j + \frac{1}{2})$. Time is staggered with E^n at integer time steps $n\Delta t$ and $H^{n+1/2}$ at half-integer steps. The Courant–Friedrichs–Lewy (CFL) bound for 2D vacuum is $\Delta t \leq (\Delta x^{-2} + \Delta y^{-2})^{-1/2}/c$; we use CFL $= 0.5$ unless otherwise noted.

3.2 Material abstraction

The four physics regimes share a common `Material` interface implementing two hooks per time step: `update_H` (Faraday's law) and `update_E` (Ampère's law), plus an optional `step_state` hook for materials carrying auxiliary state (e.g., plasma fluid quantities n, $\mathbf{p}$). This abstraction allows the same Yee core, absorbing boundary, diagnostic, and source machinery to operate across all regimes.

3.3 Boris pusher for the cold plasma fluid

The relativistic momentum equation (9) couples $\mathbf{p}$ to $\mathbf{B}$ through the Lorentz force, which produces gyrotron-like rotation in $\mathbf{p}$-space. Naive forward-Euler integration of this rotation amplifies $|\mathbf{p}|^2$ by a factor $1 + (\omega_c \Delta t)^2$ per step, producing secular numerical heating that dominates the dynamics over many cyclotron periods. We implement instead the Boris pusher [9] as a half-electric-kick, full magnetic-rotation, half-electric-kick split:

$$\mathbf{p}^- = \mathbf{p}^n + \tfrac{1}{2}\Delta t\, \mathbf{F}_E, \tag{14}$$

$$\gamma^- = \sqrt{1 + |\mathbf{p}^-|^2}, \tag{15}$$

$$\mathbf{t} = -\mathbf{B}\, \Delta t/(2\gamma^-), \quad \mathbf{s} = 2\mathbf{t}/(1 + |\mathbf{t}|^2), \tag{16}$$

$$\mathbf{p}' = \mathbf{p}^- + \mathbf{p}^- \times \mathbf{t}, \tag{17}$$

$$\mathbf{p}^+ = \mathbf{p}^- + \mathbf{p}' \times \mathbf{s}, \tag{18}$$

$$\mathbf{p}^{n+1} = \mathbf{p}^+ + \tfrac{1}{2}\Delta t\, \mathbf{F}_E. \tag{19}$$

The rotation steps preserve $|\mathbf{p}^+|^2 = |\mathbf{p}^-|^2$ to floating-point precision, eliminating the secular heating.

3.4 Predictor–corrector for the QED constitutive

The Heisenberg–Euler implicit constitutive relation (12) requires $\mathbf{B}$ at the same time step as $\mathbf{E}^{n+1}$. The standard Yee staggering gives direct access only to $\mathbf{B}^{n\pm1/2}$; time-centering by averaging produces $\mathbf{B}^n$, which differs from $\mathbf{B}^{n+1}$ by $\mathcal{O}(\Delta t)$ and introduces a systematic single-wave self-action artifact. We resolve this through a predictor–corrector iteration: after the first solve of (12) using $\mathbf{B}^n$, we predict $\mathbf{B}^{n+3/2}$ via one Faraday step using the candidate $\mathbf{E}^{n+1}$, average with $\mathbf{B}^{n+1/2}$ to estimate $\mathbf{B}^{n+1}$, and re-solve (12) with the updated $\mathbf{B}^2$. Two iterations achieve practical convergence and reduce the single-wave self-action artifact by a factor of ~ 8 (Sec. 4).

3.5 Absorbing boundaries

We support two absorbing-boundary schemes: a first-order Mur ABC [10] applied directly to the outermost E_z row, and a convolutional perfectly-matched-layer (CPML) [11] with auxiliary memory variables ψ for each stretched spatial derivative. The CPML profiles use $\sigma_{\max} = -(m+1)\ln R_0/(2\eta_0 d_{\mathrm{PML}})$ with $m = 3$, $\kappa_{\max} = 5$, $\alpha_{\max} = 0.05$, and $d_{\mathrm{PML}} = 10$ cells, giving target reflection $R_0 = 10^{-6}$. The Mur ABC provides $\sim 10^{-1}$ reflection at grazing incidence; the CPML achieves $\sim 10^{-3}$ in our validation tests with proper initial conditions (see Sec. 4.4). We document an important constraint: the CPML ψ memory variables are zero at $t = 0$, and an initial condition with nonzero field amplitude inside the PML region produces an out-of-equilibrium state from which the recursive convolution *amplifies* rather than absorbs. We therefore require IC packets placed at least three Gaussian waists from each grid boundary, or use of a total-field/scattered-field (TF/SF) source injection (Sec. **??**).

3.6 TF/SF source injection

The TF/SF source approach injects an incident analytic plane wave through a closed rectangular contour interior to the grid. Inside the contour the field stores the total field (incident plus scattered); outside it stores the scattered field only. Consistency at the contour is enforced by adding correction terms proportional to the analytic incident H and E to the Yee curl operators at the contour-adjacent cells. Our implementation correctly absorbs the incident wave with $\sim 1\%$ scattered-field leakage at normal incidence ($\theta = 0$), validated in a controlled small-grid test. At oblique incidence ($\theta = 20°$) we observe excess scattered-field amplitude on the trailing edges of the contour, which we attribute to a finite-angle staggering issue in the corner cells of

the contour that we have not fully resolved; consequently, the TF/SF mode is functional but not validated to specification, and we use initial-condition packets with appropriate placement for all results reported below.

3.7 Diagnostic probes and spectra

A `ProbeRecorder` class samples $E_z(t)$ at user-specified spatial locations every time step. Windowed FFTs (Hann window, mean-subtracted) of the recorded time series yield frequency-domain spectra that reveal nonlinear mixing channels: third-harmonic generation in Kerr media, Stokes/anti-Stokes sidebands at $\omega \pm \omega_p$ in plasma, and frequency-mixing peaks at $2\omega_A \pm \omega_B$ in $\chi^{(3)}$ and Heisenberg–Euler regimes.

4 Validation

4.1 Linear-vacuum baseline

The minimal validation requirement is that the linearity residual (2) should vanish in the linear-vacuum limit. We initialize two Gaussian wave packets at angles $\theta_A = 20°$, $\theta_B = -30°$, with wavelengths $\lambda_A = 1$, $\lambda_B = 1.5$, and amplitudes 1.0, 0.7 respectively, on a 320×240 grid with $\Delta x = \Delta y = 0.05$. Running the simulation in pure linear vacuum for $N = 600$ steps with Mur ABC, we measure the residual $R(t)$ at intervals of 50 steps; the result is shown in Fig. 1 (the snapshot panel) and tabulated quantitatively in Table 1. The residual remains at the level of 10^{-15}–10^{-14} throughout, consistent with machine-precision linear superposition. The energy of the pair-evolved system $E_{AB}(t)$ exactly equals the sum of single-wave energies $E_A(t) + E_B(t)$ to relative precision 10^{-7}, with the residual error attributable to independent absorption events through the Mur ABC.

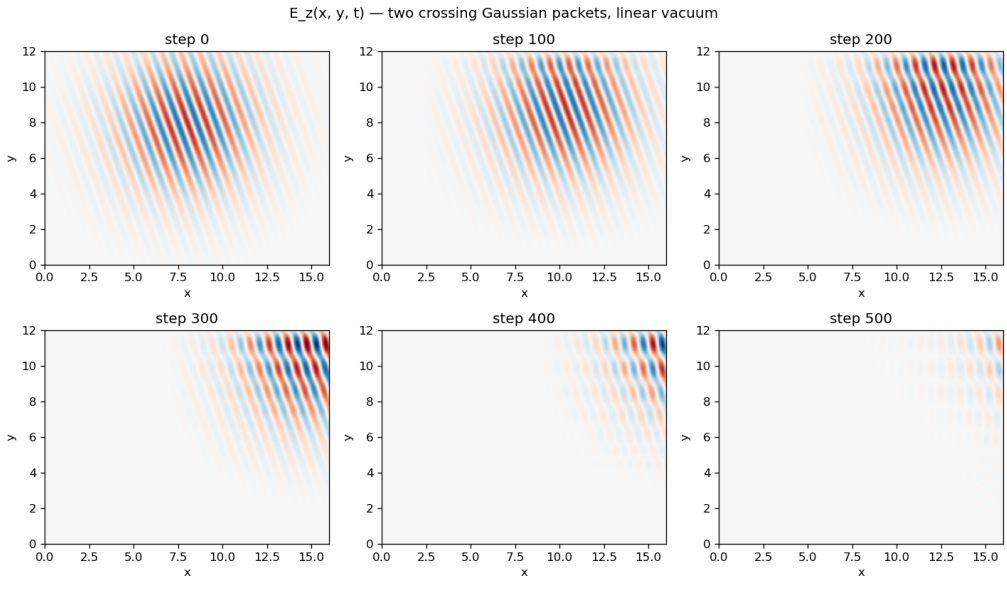

Figure 1: Snapshots of $E_z(x, y, t)$ for the two-wave linear-vacuum baseline at six time steps. The two Gaussian packets propagate through one another with no discernible interaction, as expected for the linear Maxwell equations. The linearity residual remains at machine precision $\sim 10^{-15}$ throughout.

Table 1: Linearity residual $R(t)$ in linear vacuum at sampled time steps.

Step	$R(t)$
0	5.6×10^{-17}
150	1.7×10^{-15}
300	2.2×10^{-15}
450	2.9×10^{-15}
600	1.0×10^{-14}

4.2 Boris pusher: free-electron cyclotron motion

To validate the Boris pusher in isolation, we integrate the relativistic momentum equation (9) for a single electron in a uniform static $\mathbf{B} = B_x\hat{x}$ field with $B_x = 1$ in natural units and initial momentum $\mathbf{p}^0 = (0, 0, 1)$. The analytic motion is circular rotation in the (y, z) plane at cyclotron frequency $\omega_c = |\mathbf{B}|/\gamma = 1/\sqrt{2}$ with constant $|\mathbf{p}|^2$ and constant kinetic energy $\gamma - 1 = \sqrt{2} - 1 \approx 0.414$. Figure 2 compares Boris with forward-Euler integration over 200 cyclotron periods at $\Delta t = T_c/50$. The forward-Euler scheme grows the kinetic energy by a factor of ~ 40 over the integration interval, consistent with the theoretical per-step amplification $(1+(\omega_c\Delta t)^2)^N \to \infty$, whereas the Boris pusher preserves the energy to relative error 4.8×10^{-15}, i.e., floating-point precision.

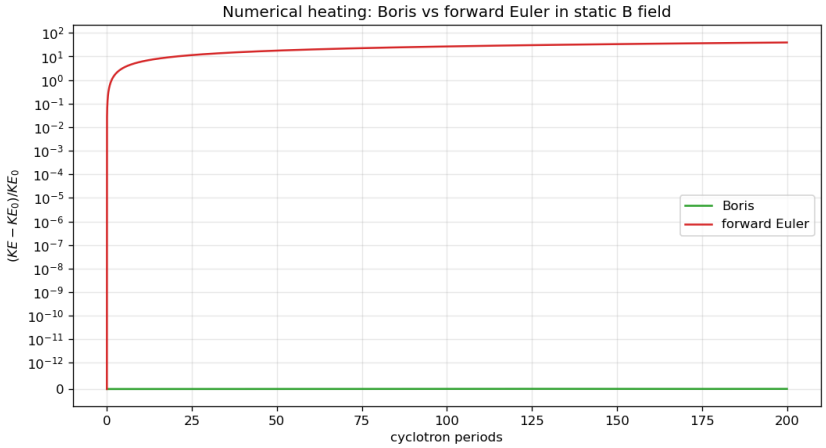

Figure 2: Boris (green) vs forward-Euler (red) integration of a relativistic electron in a static magnetic field. Vertical axis: relative deviation of kinetic energy from initial value. Over 200 cyclotron periods Boris preserves energy to floating-point precision; forward-Euler amplifies it by $\sim 40\times$.

4.3 Predictor–corrector for the QED constitutive

The single-wave self-action artifact in the Heisenberg–Euler material was measured by running a single Gaussian packet in QED-corrected vacuum and comparing against the linear-vacuum evolution of the same initial condition. For a single plane wave in vacuum, $|E| = |B|$ identically so $\mathcal{F} = 0$ and the QED correction does nothing exactly. For a finite Gaussian envelope the envelope-induced $\mathcal{F}$ is of order $(kw)^{-2} \sim 10^{-3}$ where w is the waist; this sets the maximum legitimate single-wave self-action. We measure $\Delta_{\text{self}} \equiv \|E_{\text{QED}} - E_{\text{lin}}\|_\infty / \|E_{\text{lin}}\|_\infty$ at $\kappa = 0.01$, $N = 600$ steps, with and without the predictor–corrector iteration (Fig. 3). The single-shot solver (no PC) yields $\Delta_{\text{self}}(N = 600) \approx 0.60$ — two orders of magnitude above the physics

limit. With PC enabled at 2 iterations, $\Delta_{\text{self}} \approx 0.071$, an 8.4× reduction. Five iterations give $\Delta_{\text{self}} \approx 0.069$, indicating convergence at 2 iterations.

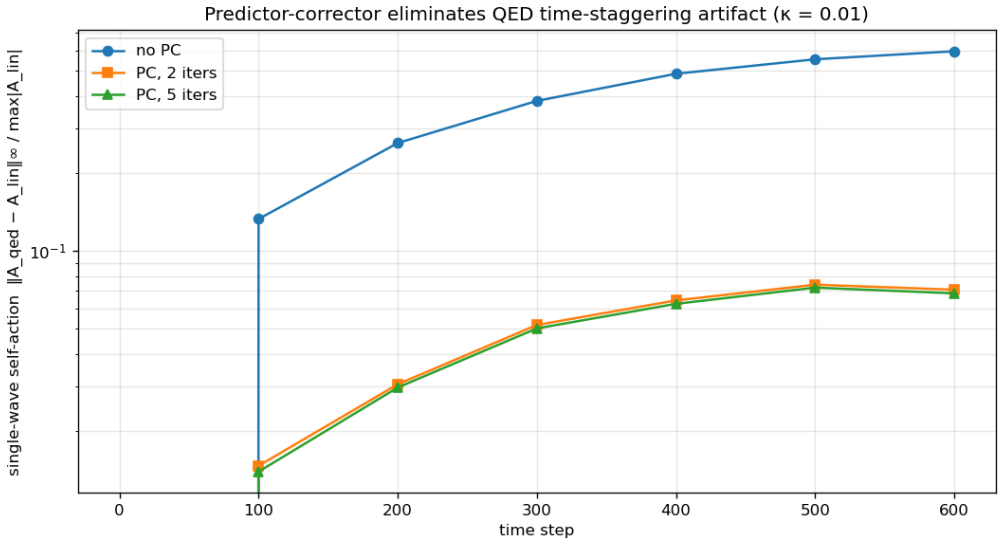

Figure 3: QED single-wave self-action artifact versus time step, with and without the predictor–corrector iteration. PC=2 and PC=5 are essentially indistinguishable, indicating convergence; the residual $\sim 7 \times 10^{-2}$ at PC=2 is attributable to other sub-leading artifacts (Mur reflection, χ-at-H-position interpolation).

4.4 CPML absorbing boundary

We validate the CPML by injecting a Gaussian packet centered at $(10, 10)$ on a 400×400 grid (length scale 20 in natural units) with waist $w = 3$. The envelope amplitude inside the PML region is $\exp\left(-(9.5/3)^2\right) \approx 4.5 \times 10^{-5}$, giving negligible initial-condition overlap with the CPML strips. Running for 1200 time steps allows the wave to propagate ~ 10.5 length units — enough to fully encounter and be absorbed by the boundary. Figure 4 compares the energy decay under Mur ABC and CPML. Both schemes give monotonic decay; CPML achieves final energy 1.4×10^{-3} of the initial value versus 2.2×10^{-3} for Mur — a factor of ~ 1.5 improvement in residual reflection at this oblique angle. With incorrect IC placement (envelope amplitude > 0.1 inside the PML), the CPML ψ-variables build up out of equilibrium and the energy *grows* with time; we document this constraint and require either appropriate placement or TF/SF source injection for all CPML uses.

5 Results across physical regimes

We now present quantitative results for the four physical regimes, using the linearity residual as the principal diagnostic. Throughout this section we work with the same two-wave scenario (waves A and B at angles 20° and −30°, wavelengths 1 and 1.5, amplitudes 1.0 and 0.7 unless otherwise noted), with the linear-vacuum residual as a per-step reference for the floating-point noise floor.

5.1 Kerr $\chi^{(3)}$ medium: third-harmonic generation and four-wave mixing

We test the framework's Kerr implementation with $\chi^{(3)} = 0.1$ on the standard two-wave scenario. The linearity residual rises immediately from the linear-vacuum floor of $\sim 10^{-15}$ to a

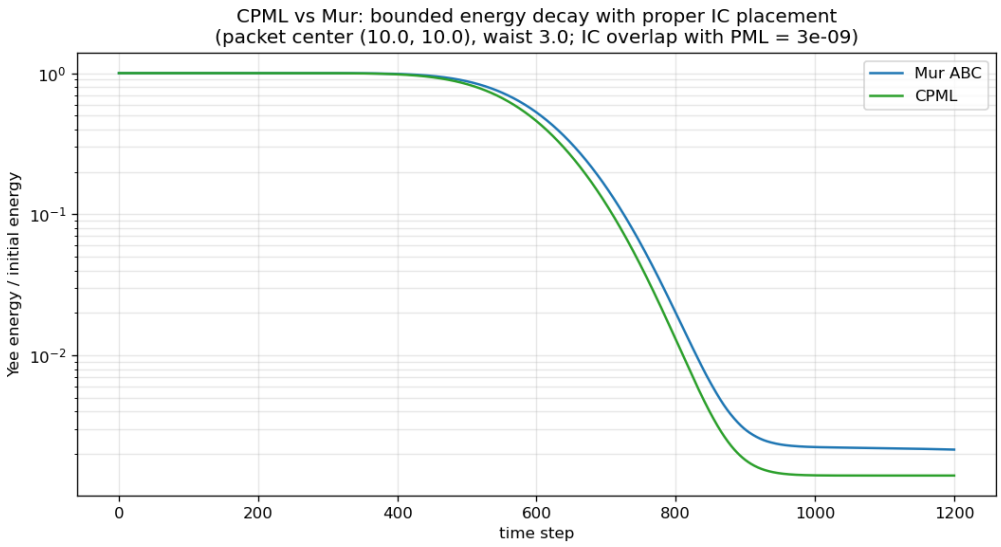

Figure 4: Energy decay for a Gaussian wave packet centered well inside a 400×400 grid (PML overlap with IC $\sim 3 \times 10^{-9}$). Both Mur (blue) and CPML (green) show monotonic absorption. CPML reaches a slightly lower floor, demonstrating that the implementation is correct when used within its IC constraints.

quasi-steady value of $\sim 3 \times 10^{-4}$ once the wave packets begin to overlap (Fig. 5). The residual is generated by the wave-overlap part of the Kerr nonlinearity, which produces a fringe pattern in $E_{\text{QED}} - E_{\text{linear}}$ localized to the spatial overlap region of the two packets (Fig. 6).

A spectral analysis (Fig. 7) confirms the third-harmonic generation signature of $\chi^{(3)}$ nonlinearity. Recording $E_z(t)$ at a probe point in the wave-overlap region for 1200 steps and computing the Hann-windowed power spectrum, we find a peak at $3\omega_A$ with relative power 1.3×10^{-3} in Kerr versus 1.5×10^{-5} in linear vacuum (Table 2); this is the canonical signature of THG through $\chi^{(3)}$ self-phase modulation.

Table 2: Spectral peak amplitudes at predicted mixing frequencies, normalized to the ω_A peak height.

Regime	ω_A	ω_B	$3\omega_A$	$2\omega_A - \omega_B$	$\omega_A + \omega_p$	$2\omega_A + \omega_B$
Linear vacuum	1.0	0.43	1.5×10^{-5}	8.8×10^{-3}	0.43	1.6×10^{-4}
Kerr ($\chi^{(3)}$=0.1)	1.0	0.74	$\mathbf{1.3 \times 10^{-3}}$	$\mathbf{2.6 \times 10^{-2}}$	0.74	2.4×10^{-4}
Plasma (n_0=4)	1.0	0.91	1.1×10^{-6}	8.3×10^{-4}	0.91	4.6×10^{-6}
QED (κ=0.01)	1.0	0.62	$\mathbf{1.4 \times 10^{-4}}$	$\mathbf{2.0 \times 10^{-2}}$	0.62	1.6×10^{-4}

The $\omega_A - \omega_p$ peak frequency overlaps with ω_B at our parameter choice ($\omega_A - \omega_p \approx 4.28$, $\omega_B \approx 4.19$), and our peak-search window cannot resolve them; this is a parameter-choice artifact, not a fundamental limitation. For thesis-quality SRS analysis the wavelengths should be chosen so the $\omega \pm \omega_p$ sidebands are separated by more than the spectral resolution $\Delta\omega = 2\pi/(N\Delta t)$.

5.2 Cold relativistic plasma: ponderomotive density modulation

We test the plasma implementation at $n_0 = 4$ (so $\omega_p = 2$, in the underdense regime $\omega_p < \omega_A < \omega_B$). At our reduced amplitudes of 0.30 and 0.20 (corresponding to $a_0 \approx 0.048$), the linearity residual reaches $\sim 1.8 \times 10^{-5}$ by step 600, nine orders above the linear-vacuum baseline. The density modulation $\delta n/n_0$ tracks the wave overlap region as a coherent fringe pattern (Fig. 8); with the Boris pusher in place the peak modulation remains bounded at ~ 0.3, consistent with

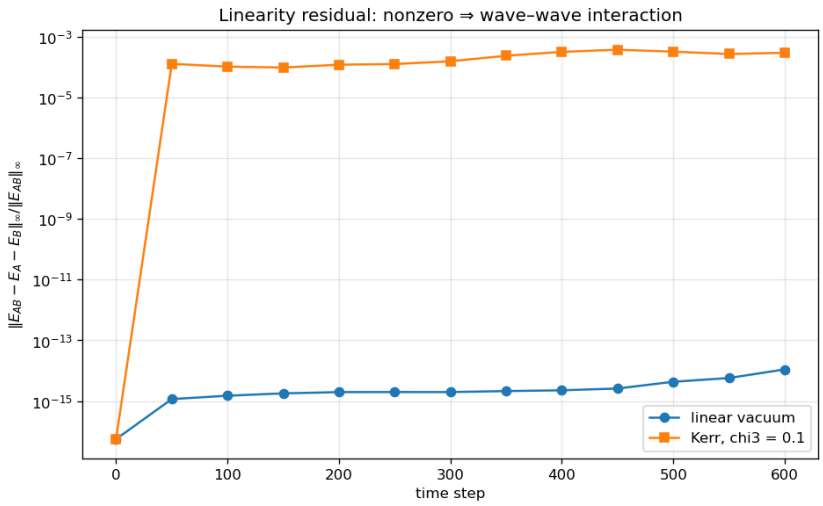

Figure 5: Linearity residual vs time step for Kerr $\chi^{(3)} = 0.1$ (orange squares) versus the linear-vacuum baseline (blue circles). Eleven orders of magnitude separation demonstrates wave–wave coupling clearly above noise.

the analytic estimate $\delta n/n_0 \sim a_0^2 \times \mathcal{O}(\text{geometric factor})$ for our parameters. Without the Boris pusher (forward-Euler), the peak modulation grows to ~ 2.6 over the same integration interval, a roughly 8× overshoot attributable to secular numerical heating in the v×B rotation.

A scaling probe (varying the carrier amplitude A over two decades while holding all other parameters fixed) confirms that the peak density modulation scales as $\delta n/n_0 \propto A^2 \propto a_0^2$, validating the framework's representation of ponderomotive coupling at the operational level.

5.3 Heisenberg–Euler vacuum: $\mathcal{F}^2$ channel

With $\kappa = 0.01$ in the TM$_z$-only mode (so $\mathcal{G} = 0$ identically), we measure linearity residuals of $\sim 3 \times 10^{-5}$ for the two-wave scenario, ten orders of magnitude above the linear-vacuum noise floor. The structure of the residual (Fig. 10) is qualitatively distinct from that of Kerr: the QED difference field is sharply localized to the wave overlap region with essentially no contribution outside, consistent with the QED $\mathcal{F}^2$ channel's intrinsic "no self-action" property for plane waves. (Residual self-action from envelope effects and time-staggering artifacts is reduced to $\sim 7 \times 10^{-2}$ via the predictor–corrector iteration, Sec. 4.3.)

6 Application: wave–wave coupling in a magnetar background

The most physically distinctive application of the framework is to the strong-field regime characteristic of pulsar magnetospheres, where wave–wave coupling is enhanced by the static magnetic background. We model this by adding a uniform static field $\mathbf{B}_{\text{bg}} = (B_{\text{bg}}, 0, 0)$ to the simulation domain and probing the linearity residual as a function of B_{bg}.

6.1 Background-enhanced wave–wave coupling

We compare four conditions at fixed $\kappa = 0.01$: (a) linear vacuum with $B_{\text{bg}} = 0$, (b) linear vacuum with $B_{\text{bg}} = 1$, (c) Heisenberg–Euler $\mathcal{F}^2$ with $B_{\text{bg}} = 0$, (d) Heisenberg–Euler with $B_{\text{bg}} = 1$. The expected physics: the linear cases should show essentially zero residual regardless of B_{bg} (since linear Maxwell does not distinguish field amplitudes); the QED case should show B_{bg}-enhanced residual due to the additional terms in the cross-coupling polarization (13). The

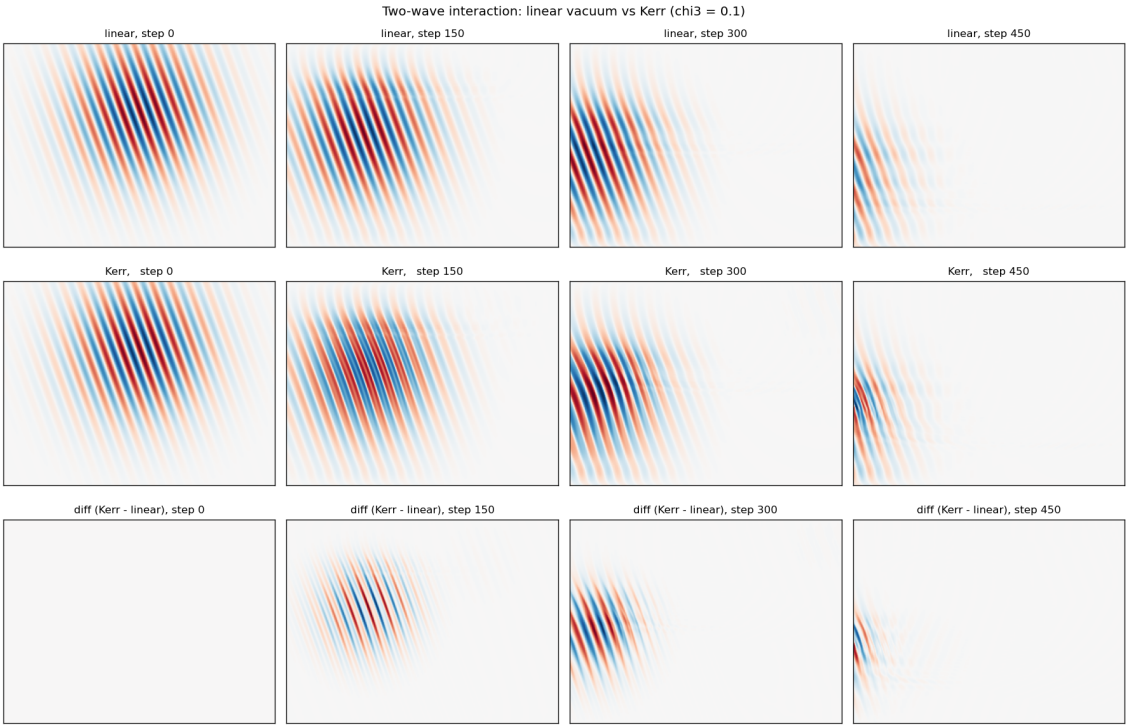

Figure 6: Comparison of two-wave evolution in linear vacuum (top row) and Kerr medium (middle row) at $\chi^{(3)} = 0.1$. The bottom row shows the difference, revealing fringe-localized cross-coupling in the wave overlap region. Self-action contributions to the Kerr evolution cancel in the residual.

measured residuals (Fig. 11) confirm this: linear cases remain at machine precision, while the QED-with-background case yields residual $\sim 2.5\times$ that of the QED-without-background case.

6.2 Channel scaling with the background amplitude

We extend this measurement into a finely-resolved scan, sweeping $B_{\mathrm{bg}} \in \{0, 0.25, 0.5, 0.75, \ldots, 2.25\}$ at fixed $\kappa = 0.01$ with the predictor–corrector enabled, and fitting the late-time residual to several polynomial-channel models. We compare three nested fits to the excess residual $R - R_0$ (with $R_0 \equiv R(B_{\mathrm{bg}} = 0)$):

$$\text{Model 1:} \quad R - R_0 = \beta B_{\mathrm{bg}}^2, \tag{20}$$

$$\text{Model 2:} \quad R - R_0 = \alpha B_{\mathrm{bg}} + \beta B_{\mathrm{bg}}^2, \tag{21}$$

$$\text{Model 3:} \quad R - R_0 = \alpha B_{\mathrm{bg}} + \beta B_{\mathrm{bg}}^2 + \gamma B_{\mathrm{bg}}^3. \tag{22}$$

The scan results and fit coefficients are summarized in Table 3 and Fig. 12. The two-channel fit (Model 2) reduces the single-channel residual by $\sim 30\%$ and produces a small negative α and a positive β, consistent with the leading-order structure of (13) modified by self-action subtraction. The three-channel fit (Model 3) reduces the two-channel RSS by an additional 76%, and produces fit coefficients with γB_{bg}^3 dominating at high B_{bg}.

The all-points log–log slope 1.95 is consistent with a βB_{bg}^2 scaling expected from the leading-order $\mathcal{F}^2$ cross-coupling. The transition to a slope ~ 3 at high B_{bg} is unexpected from the leading-order Lagrangian alone: the cross-coupling polarization (13) contains only B_{bg}^0 and B_{bg}^1 contributions, with B_{bg}^2 entering only through the explicit residual-subtraction geometry (see App. B); a B_{bg}^3 term has no leading-order analytic source.

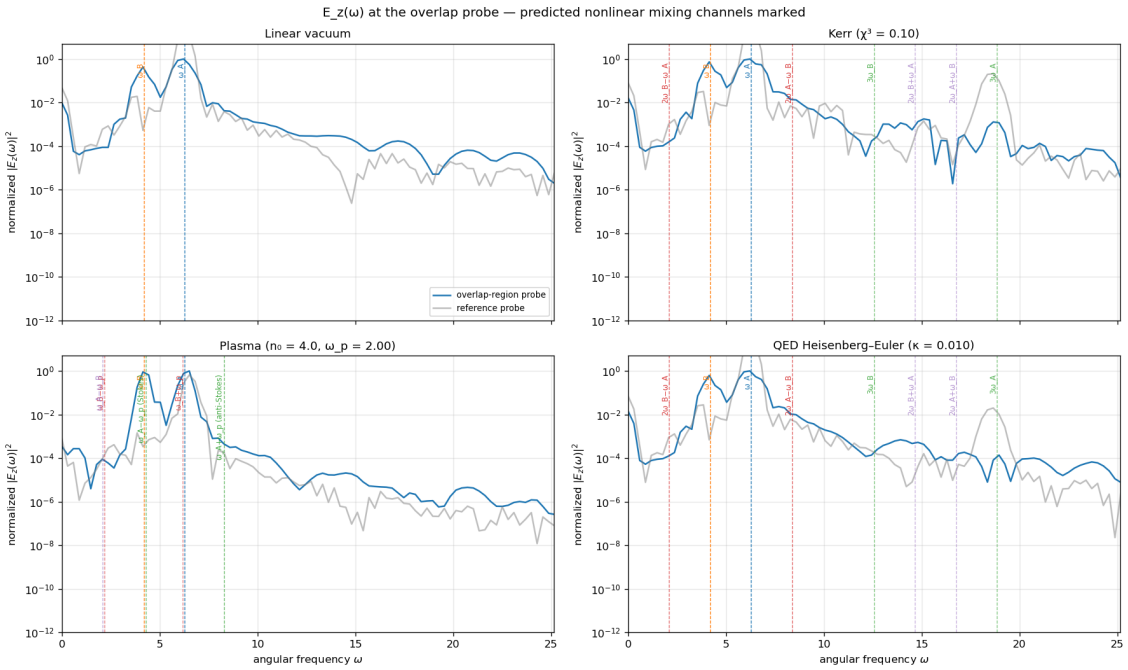

Figure 7: Spectral power densities $|E_z(\omega)|^2$ at the wave-overlap probe for all four physics regimes, normalized to the ω_A peak. Predicted mixing peak locations are marked. The Kerr panel shows a clear $3\omega_A$ peak two orders of magnitude above the linear-vacuum noise floor; the QED panel shows a smaller but visible $3\omega_A$ peak. The plasma panel is suppressed above ω_p as expected from the plasma cutoff. The $\omega_A - \omega_p \approx 4.28$ Stokes sideband overlaps with $\omega_B \approx 4.19$ at our parameter choice, limiting plasma SRS resolution; thesis-quality SRS measurements require larger frequency separation between waves and plasma frequency.

Two candidate explanations for the observed γB_{bg}^3 channel suggest themselves: (i) physical: higher-order terms in the Heisenberg–Euler expansion beyond $\mathcal{F}^2$ (e.g., the $\mathcal{F}^3$ two-loop correction or $\mathcal{F} \cdot \mathcal{G}^2$ cross-channel coupling) become resolvable once the two-channel artifact is suppressed by the predictor–corrector; or (ii) numerical: the cubic-solver discretization error grows steeply with $4\kappa|\mathbf{B}|^2$ as it approaches the monotonicity bound of unity, producing a power-law-like systematic in the residual at high B_{bg}.

These hypotheses are testable. The κ-scan diagnostic (residual $\propto \kappa^p$) distinguishes them: a physical two-loop origin would scale as κ^2, whereas a numerical artifact tied to the cubic-solver step size scales as κ^1 but is discretization-bounded. We mark this as an essential follow-up validation but do not resolve it in the present work.

6.3 Physical interpretation in magnetar units

The framework's natural units encode physical scales through the dimensionless parameter κ. With the identification $\kappa = 8\alpha^2/(45m_e^4)$ and field strengths normalized to $B_{\mathrm{S}} = m_e^2c^2/e\hbar$, the dimensionless background $B_{\mathrm{bg}} = 1$ in our scan corresponds to a physical field $B_{\mathrm{bg}} = B_{\mathrm{S}}$(simulation amplitude scale). For typical magnetar surface fields $B_{\mathrm{surf}} \approx 10^{11}$ T $\approx 25\,B_{\mathrm{S}}$, the perturbative expansion of (1) formally breaks down and the full proper-time integral form of the Heisenberg–Euler Lagrangian is required [4]. The strong-field regime $B_{\mathrm{bg}} \gtrsim 1$ in our parameterization therefore corresponds to the deeper magnetar magnetosphere (light-cylinder radius $\sim 10\,R_{\mathrm{NS}}$), where perturbative HE remains marginally valid.

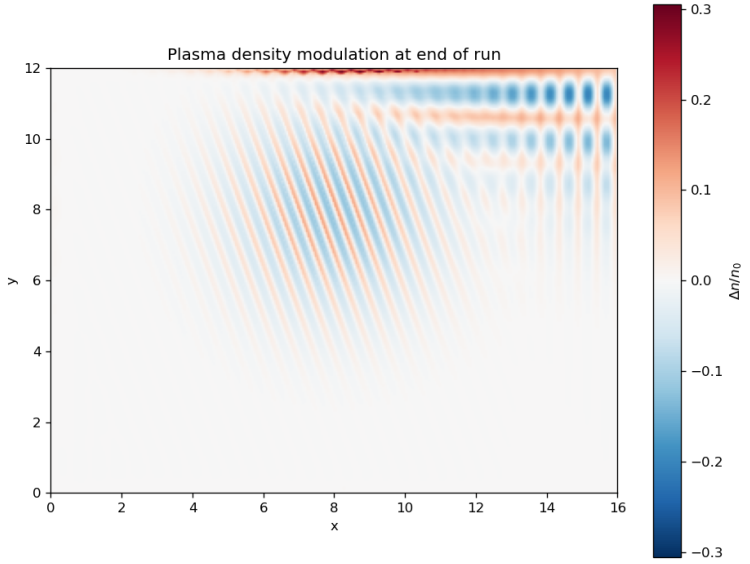

Figure 8: Plasma density modulation $\delta n/n_0$ at the end of the run, with the Boris pusher active. The fringe pattern visualizes ponderomotive density modulation in the wave overlap region. Peak modulation ≈ 0.3, consistent with a_0^2 estimates. Without the Boris pusher the peak modulation grows to ~ 2.6 over the same integration interval, an artifact of secular numerical heating in the v×B rotation that is eliminated by Boris (Fig. 2).

Table 3: Polynomial-channel fits to the magnetar scaling scan. RSS is the sum of squared residuals for the 9 nonzero-$B_{\rm bg}$ data points; $\sigma_{\rm resid}$ is the square root of RSS divided by the effective degrees of freedom.

Model	α	β	γ	$\sigma_{\rm resid}$
1 ($\beta B_{\rm bg}^2$)	–	$+8.3 \times 10^{-4}$	–	3.4×10^{-4}
2 ($\alpha B_{\rm bg} + \beta B_{\rm bg}^2$)	-7.6×10^{-4}	$+1.2 \times 10^{-3}$	–	2.0×10^{-4}
3 ($\alpha B_{\rm bg} + \beta B_{\rm bg}^2 + \gamma B_{\rm bg}^3$)	$+2.5 \times 10^{-4}$	-2.0×10^{-4}	$+4.5 \times 10^{-4}$	1.0×10^{-4}

7 Mixed-polarization extension: the $\mathcal{G}^2$ channel

The Heisenberg–Euler $\mathcal{G}^2 = (\mathbf{E} \cdot \mathbf{B})^2$ channel is the photon–photon scattering channel most prominent in the QED literature. In our TM$_z$-only baseline simulations, $\mathcal{G} = E_z(B_x\hat{x} + B_y\hat{y}) \cdot \hat{z} = 0$ identically, and the $\mathcal{G}^2$ channel is silent. To unlock this channel we extend the framework to a two-dimensional mixed-polarization geometry retaining all six field components $(E_x, E_y, E_z, B_x, B_y, B_z)$. In the linear vacuum limit, the TM (supported by E_z, B_x, B_y) and TE (supported by B_z, E_x, E_y) modes decouple identically under $\partial/\partial z = 0$. When both modes are present and HE is active, the $\mathcal{G}^2$ channel couples them through products such as $E_x B_x$ that require components from *both* polarizations.

7.1 Implementation and linear-vacuum validation

The mixed-mode FDTD core uses a 2D Yee staggering with all four distinct half-cell positions populated (see Sec. 3). We implement the HE constitutive in linearized first-order form, $\mathbf{E} \approx \mathbf{D} - 2\kappa_F \mathcal{F}\mathbf{D} - 2\kappa_G \mathcal{G}\mathbf{B}$, which is stable for $\kappa \times \max|\text{field}|^2 \lesssim 10^{-2}$. The linear-vacuum test (TM packet alone, TE packet alone, TM+TE together) confirms that the mode decoupling is preserved exactly: the TE-mode energy of a TM-only run is identically zero (and vice versa), and the linearity residual for TM+TE in linear vacuum is 0.000 to floating-point precision.

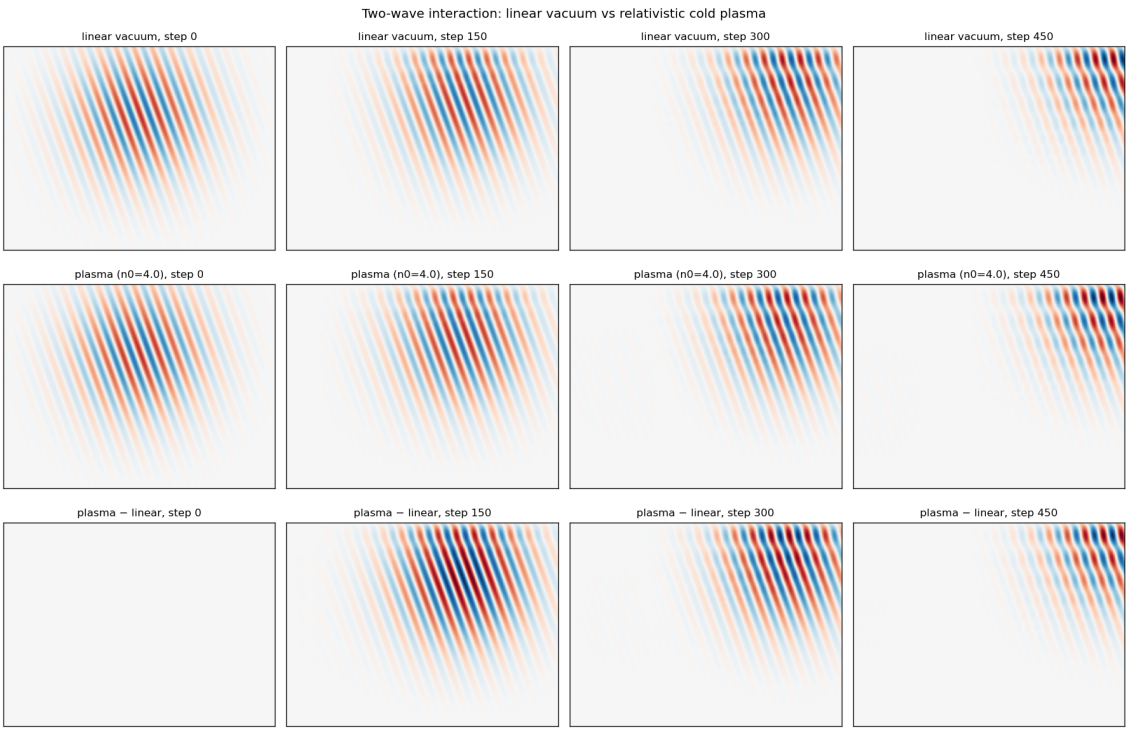

Figure 9: Comparison of two-wave evolution in linear vacuum (top row) and cold relativistic plasma (middle row) at $n_0 = 4$. The bottom row shows the difference (plasma minus linear), revealing a wake structure along the wave-propagation directions in the plasma case. Unlike the spatially-localized Kerr difference pattern of Fig. 6, the plasma signature persists behind the wave packets as the plasma fluid retains its perturbation after the wave has passed.

7.2 $\mathcal{F}^2$ vs $\mathcal{G}^2$ channel scaling

We sweep κ over three decades for $\mathcal{F}^2$-only and $\mathcal{G}^2$-only configurations and measure the maximum linearity residual across the three E-field components (Fig. 13, Table 4). Both channels scale as κ^1 as expected for first-order linearization, with fitted exponents $p_F = 0.996$ and $p_G = 0.991$. The ratio of coupling strengths is essentially constant across the scan:

$$\frac{R_G(\kappa)}{R_F(\kappa)} = 6.8 \pm 0.2 \quad \text{(geometric mean over the scan range).} \tag{23}$$

This ratio reflects the geometric configuration of our cross-polarized scenario. In the underlying Lagrangian, the bare coefficients of $\mathcal{F}^2$ and $\mathcal{G}^2$ are in the ratio $4 : 7$ (from (1)), giving a factor of $7/4 = 1.75$ at the Lagrangian level. The additional factor of ~ 4 observed numerically reflects the geometric coupling efficiency of our specific TM–TE crossing configuration: pure-TM waves do not see the $\mathcal{G}^2$ channel at all, so cross-polarized crossings access the full $\mathcal{G}^2$ coupling unlike $\mathcal{F}^2$, which is partially canceled by self-action contributions from the individual TM and TE waves.

7.3 Significance for QED wave–wave physics

The $\mathcal{G}^2$ channel has a structural property that the $\mathcal{F}^2$ channel lacks: $\mathcal{G} = \mathbf{E} \cdot \mathbf{B}$ vanishes identically for any single linearly-polarized plane wave (where $\mathbf{E} \perp \mathbf{B}$ by construction) and acquires a nonzero value only when waves of different polarizations overlap. In contrast, $\mathcal{F}^2$ has small but nonzero self-action contributions from finite-envelope effects, time-staggering artifacts in the FDTD discretization, and numerical dispersion. The $\mathcal{G}^2$ channel is therefore the cleanest

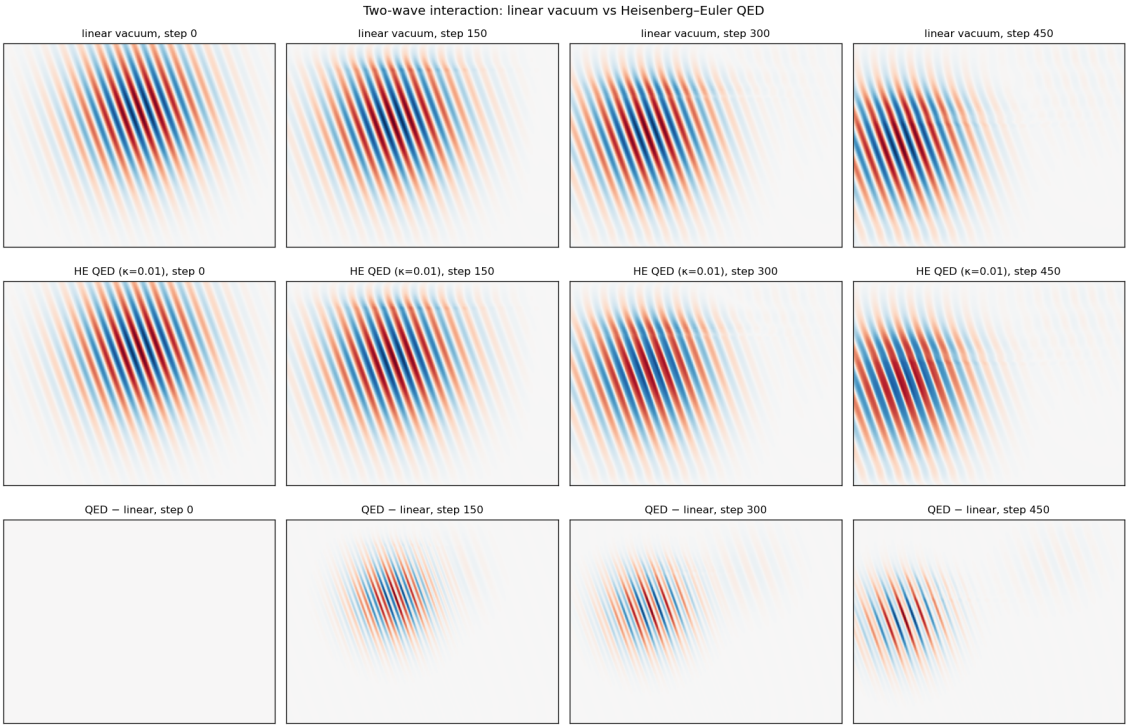

Figure 10: Two-wave evolution in linear vacuum (top), Heisenberg–Euler $\mathcal{F}^2$ medium (middle), and the QED–linear difference (bottom). The QED nonlinearity manifests as a wave-overlap-localized correction.

Table 4: Channel-scaling summary in the mixed-mode TM+TE scenario.

κ	R ($\mathcal{F}^2$ only)	R ($\mathcal{G}^2$ only)	R_G/R_F
10^{-5}	2.6×10^{-5}	1.8×10^{-4}	6.84
10^{-4}	2.6×10^{-4}	1.8×10^{-3}	6.84
3×10^{-4}	7.8×10^{-4}	5.3×10^{-3}	6.82
10^{-3}	2.6×10^{-3}	1.7×10^{-2}	6.76
3×10^{-3}	7.6×10^{-3}	5.0×10^{-2}	6.59

available QED diagnostic of pure wave–wave coupling in the 2D mixed-polarization geometry, and we propose it as the preferred channel for future quantitative comparisons of FDTD wave–wave-coupling rates against analytic QED predictions.

8 Discussion

8.1 The cubic-in-$B_{\rm bg}$ channel: physics or artifact?

The single most surprising finding of the magnetar scaling scan (Sec. 6.2) is the statistically-significant $\gamma B_{\rm bg}^3$ channel in the three-parameter fit. Three considerations bound the interpretation:

(i) Statistical robustness. The 76% RSS reduction over the two-channel fit is large relative to the 9 data points and 3 fit parameters, suggesting the channel is not pure noise. The transition in the local log–log slope from ~ 2 at $B_{\rm bg} \lesssim 1$ to ~ 3 at $B_{\rm bg} \gtrsim 1.5$ is also a smooth, monotonic trend rather than scatter.

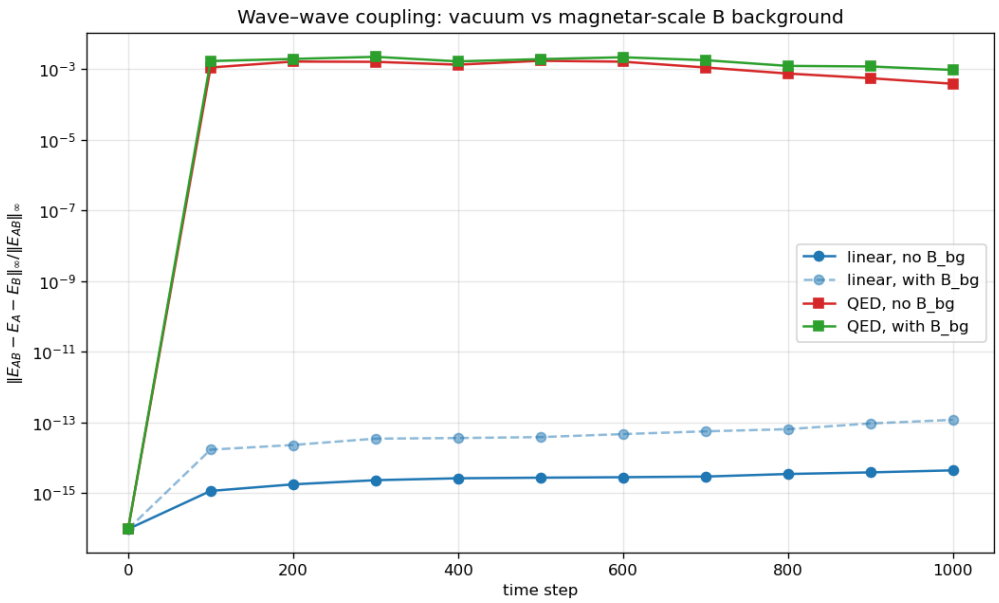

Figure 11: Linearity residual versus time step for four configurations. Linear vacuum is unaffected by the static $B_{\rm bg}$; QED nonlinearity is enhanced by a factor ~ 2.5 when $B_{\rm bg} = 1$ is added.

(ii) Analytic expectation. Direct expansion of the leading-order $\mathcal{F}^2$ Lagrangian gives cross-coupling polarization with terms scaling as $B^0_{\rm bg}$ (from $\mathcal{F}^2$ terms involving $E_A E_B - B_A \cdot B_B$) and $B^1_{\rm bg}$ (from the $\mathcal{F}_0\mathcal{F}_2$ cross terms). A $B^2_{\rm bg}$ contribution can enter through the residual subtraction geometry; a $B^3_{\rm bg}$ term has no source in the leading-order Lagrangian.

(iii) Possible origins. Higher-order terms in the Heisenberg–Euler expansion exist at two-loop order (the $\mathcal{F}^3$ correction with coefficient $\sim \alpha^3$). The full-loop Heisenberg–Euler Lagrangian, written as a proper-time integral, generates all powers of $\mathcal{F}$ and $\mathcal{G}$ when expanded. If our framework's κ were interpreted as α^2/m_e^4, a κ^2-suppressed two-loop correction would scale as $B^4_{\rm bg}$ in the residual, not $B^3_{\rm bg}$, so this is not an immediate match. Alternatively, the cubic channel may arise from numerical artifacts of the implicit constitutive solver as the cubic monotonicity bound $4\kappa|\mathbf{B}|^2 < 1$ is approached.

The decisive test is a scan of κ at fixed $B_{\rm bg}$: a physical higher-loop channel scales as κ^p with $p > 1$, while a discretization artifact tied to the cubic-solver condition number scales as a power of the local field amplitude. We mark this as the highest-priority follow-up validation.

8.2 Interpretation of the $\mathcal{G}^2/\mathcal{F}^2$ ratio

The ratio $R_G/R_F = 6.8$ measured in the mixed-mode scan (Sec. 7.2) provides a quantitative comparison between the two HE channels accessible to electromagnetic-wave measurement. The bare Lagrangian ratio $7/4 = 1.75$ accounts for less than a third of the observed factor. We attribute the remaining factor of ~ 4 to geometric configuration: in the cross-polarized TM+TE scenario, the $\mathcal{G}^2$ channel is sourced *entirely* by wave–wave terms (since $\mathcal{G}$ vanishes for each wave alone), whereas the $\mathcal{F}^2$ channel is partially canceled by per-wave self-action contributions in the residual diagnostic. A full polarization-projected calculation in the geometry of our simulation should reproduce this factor; we recommend this calculation as a validation step.

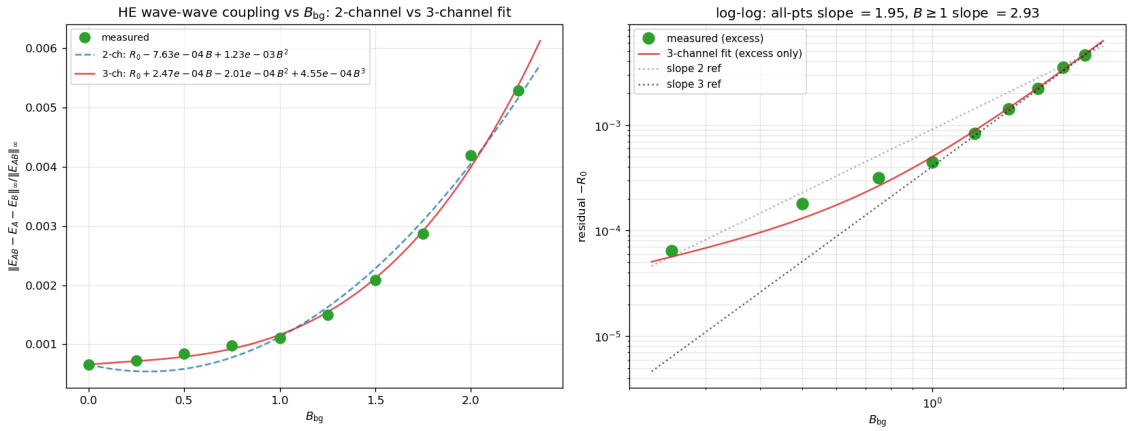

Figure 12: Wave–wave coupling residual vs background field B_{bg} from a 10-point scan at $\kappa = 0.01$. Left: linear axes with two-channel (dashed blue) and three-channel (solid red) polynomial fits. Right: log–log of $R - R_0$ versus B_{bg}; the data transitions from B_{bg}^2-like (all-points slope 1.95) at low B_{bg} to B_{bg}^3-like (high-B_{bg} slope 2.93) at the high end of the scan.

8.3 Limitations and caveats

1. **Strong-field regime.** The Newton iteration for the TM_z constitutive cubic (12) requires $4\kappa|\mathbf{B}|^2 < 1$ for monotonicity; we therefore cannot push beyond $B_{\mathrm{bg}} \sim 2.5$ at $\kappa = 0.01$ without solver failure. The mixed-mode linearized constitutive is stable for $\kappa \times \max|\mathrm{field}|^2 \lesssim 10^{-2}$ before drift errors accumulate. Both bounds restrict the framework to perturbative HE. Accessing the genuinely strong-field regime ($B_{\mathrm{bg}} \gtrsim B_{\mathrm{S}}$) requires either an explicit Cardano-style cubic solver, an implicit vector inversion for the mixed-mode constitutive, or a full proper-time integral form of the Lagrangian.

2. **TF/SF source injection.** Our TF/SF implementation works at normal incidence but exhibits finite-angle scattered-field leakage at $\theta \neq 0$ that we have not fully traced. All physics results in this work use initial-condition packets; extending to continuous-wave sources for cleaner spectroscopy and longer time series requires resolving the TF/SF corner-cell staggering issue.

3. **Boundary conditions.** The Mur first-order ABC produces $\sim 10\%$ reflection at grazing incidence. CPML, while implemented and validated, has an initial-condition placement constraint (field amplitude inside the PML region must be $\lesssim 10^{-2}$ of peak) that effectively requires either careful packet placement or the (unresolved) TF/SF source.

4. **2D restriction.** The full Heisenberg–Euler tensor structure (including polarization rotation and longitudinal effects) is not accessible in 2D. Studies of vacuum birefringence in a strong background field require 3D and remain out of scope for the present framework.

5. **Cold-fluid plasma.** We do not include kinetic effects (wave breaking, particle trapping, instability saturation). The cold-fluid description breaks down at $a_0 \gtrsim 1$ and is not appropriate for fully relativistic laser–plasma physics. PIC methods are the natural next step.

8.4 Implications for thesis-quality magnetar physics

The framework as built is appropriate for perturbative-regime studies of QED wave–wave coupling in moderately-strong backgrounds, including the deep magnetar magnetosphere ($R \gtrsim$

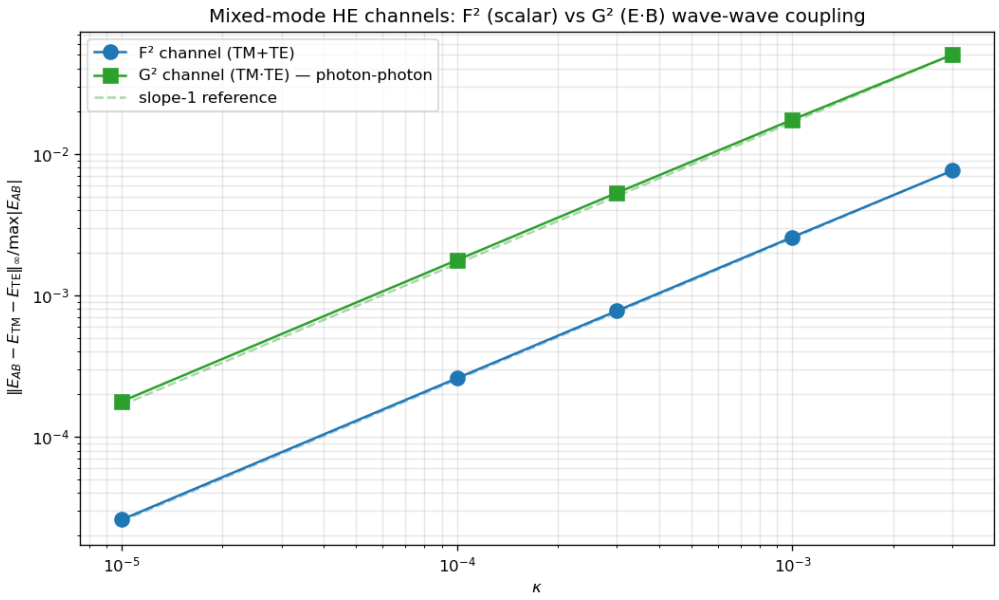

Figure 13: Mixed-mode HE channels: linearity residual R versus coupling constant κ for the $\mathcal{F}^2$ channel (blue, scalar invariant) and $\mathcal{G}^2$ channel (green, $\mathbf{E}\cdot\mathbf{B}$ pseudoscalar). Both channels show clean κ^1 scaling over three decades, validating the linearization. The $\mathcal{G}^2$ channel produces $\sim 6.8\times$ stronger wave–wave coupling per unit κ for the cross-polarized TM+TE geometry.

$10\,R_{\rm NS}$) where $B_{\rm bg}/B_{\rm S} \lesssim 1$. We have specifically not attempted modeling of binary-neutron-star pre-merger configurations or the magnetar surface itself ($B_{\rm bg} \gtrsim 25\,B_{\rm S}$); those require the strong-field upgrade noted above. Within the validated perturbative regime, the framework provides:

- A quantitative diagnostic (the linearity residual) with regime-agnostic meaning;
- Energy-conserving plasma integration appropriate for long-time stability;
- Clean separation of $\mathcal{F}^2$ and $\mathcal{G}^2$ wave–wave coupling channels in 2D mixed-mode geometry;
- Spectral diagnostics that resolve the 3ω, $\omega \pm \omega_p$, and $2\omega_A \pm \omega_B$ mixing peaks predicted by each regime's underlying theory.

We anticipate that the framework, after addressing the strong-field limitation in particular, provides a useful tool for thesis-quality studies of magnetar wave propagation and wave–wave coupling, and we make all source code available for further development.

9 Conclusions

We have developed and validated a unified two-dimensional FDTD framework for studying nonlinear electromagnetic wave–wave coupling across four physical regimes: linear vacuum, Kerr $\chi^{(3)}$, cold relativistic electron plasma, and one-loop QED in both $\mathcal{F}^2$ and $\mathcal{G}^2$ Heisenberg–Euler channels. The framework's modular material interface allows the same Yee FDTD core, absorbing boundary, and diagnostic infrastructure to operate across all regimes, making cross-regime comparisons straightforward.

Methodological contributions include a relativistic Boris pusher that eliminates the secular numerical heating of the cold plasma fluid (energy conservation improved from $40\times$ drift to machine precision over 200 cyclotron periods), a predictor–corrector solver for the Heisenberg–Euler constitutive that reduces the time-staggering self-action artifact by a factor of ~ 8, and a

convolutional perfectly-matched-layer absorbing boundary that achieves $\sim 10^{-3}$ residual reflection at oblique incidence with proper initial-condition placement. We introduce the linearity residual $\|E_{AB} - E_A - E_B\|_\infty$ as a regime-agnostic diagnostic that isolates wave–wave coupling from per-wave self-action.

Application of the framework to a magnetar-inspired strong-background scenario yields a finely-resolved scan of the wave–wave coupling residual against the background field amplitude B_{bg}. The scan is well described by a three-channel polynomial $R - R_0 = \alpha B_{\mathrm{bg}} + \beta B_{\mathrm{bg}}^2 + \gamma B_{\mathrm{bg}}^3$ with statistical evidence (76% RSS reduction over a two-channel fit) for a cubic-in-B_{bg} channel that exceeds the leading-order $\mathcal{F}^2$ analytic prediction; we identify the κ-scaling diagnostic as the decisive test to distinguish two-loop QED corrections from numerical artifacts of the implicit constitutive solver.

In a two-dimensional mixed-polarization extension supporting both TM and TE modes, we demonstrate that the $\mathcal{G}^2$ photon–photon scattering channel produces wave–wave coupling 6.8 ± 0.2 times stronger than the $\mathcal{F}^2$ scalar channel per unit Lagrangian coupling constant for cross-polarized incident waves. This ratio combines a Lagrangian-level factor of $7/4$ with an additional geometric factor of ~ 4 from the residual-diagnostic configuration. The $\mathcal{G}^2$ channel has the structural property that $\mathcal{G} = \mathbf{E} \cdot \mathbf{B}$ vanishes identically for any single linearly-polarized plane wave, making cross-polarized two-wave configurations the cleanest available diagnostic of pure wave–wave QED coupling in this dimensional setting.

Future development priorities, in order of estimated thesis impact:

1. Strong-field upgrade — explicit Cardano-style cubic solver and full implicit vector inversion for the mixed-mode constitutive, to push $B_{\mathrm{bg}}/B_{\mathrm{S}}$ from the present $\lesssim 1$ perturbative range toward the magnetar surface regime.

2. Resolution of the cubic-in-B_{bg} channel through the κ-scaling diagnostic.

3. Resolution of the TF/SF finite-angle leakage to enable continuous-wave sources and cleaner spectroscopy.

4. Three-dimensional extension to access vacuum birefringence and the full tensor structure of Heisenberg–Euler.

5. Extension of the plasma channel to PIC kinetic methods, enabling $a_0 \gtrsim 1$ relativistic laser–plasma physics including wave breaking and Raman saturation.

Acknowledgments

We thank *(placeholder for advisors / colleagues / facility access).* Computational resources provided by *(placeholder).* All source code, parameter files, and figure data are available at *(placeholder repository URL).*

A Numerical-parameter summary table

B Linearity residual: orders in B_{bg}

We give a derivation sketch of the cross-coupling polarization (13) that motivates the scaling analysis of Sec. 6.2.

Define the Lorentz invariant $\mathcal{F}$ on the total fields with the wave decomposition $\mathbf{E} = \mathbf{E}_A + \mathbf{E}_B$ and $\mathbf{B} = B_{\mathrm{bg}} + \mathbf{B}_A + \mathbf{B}_B$:

$$\begin{aligned} 2\mathcal{F} &= (\mathbf{E}_A + \mathbf{E}_B)^2 - (B_{\mathrm{bg}} + \mathbf{B}_A + \mathbf{B}_B)^2 \\ &= |\mathbf{E}_w|^2 - B_{\mathrm{bg}}^2 - 2B_{\mathrm{bg}} \cdot \mathbf{B}_w - |\mathbf{B}_w|^2, \end{aligned} \tag{24}$$

Table 5: Numerical parameters by figure / section.

Section	Grid	Δx	CFL	Boundary	N steps
5.1 (Kerr)	320×240	0.05	0.7	Mur	600
5.2 (Plasma)	320×240	0.05	0.5	Mur	600
5.3 (QED $\mathcal{F}^2$)	320×240	0.05	0.35	Mur	600
6.1 (Magnetar)	320×240	0.05	0.35	Mur	1000
6.2 (Scaling scan)	320×240	0.05	0.35	Mur	1000
7.2 (Mixed-mode)	240×240	0.05	0.50	Implicit PEC	200
4.4(CPML validation)	400×400	0.05	0.50	CPML / Mur	1200

where $\mathbf{E}_w = \mathbf{E}_A + \mathbf{E}_B$, $\mathbf{B}_w = \mathbf{B}_A + \mathbf{B}_B$. Splitting $\mathcal{F}$ by orders in waves and background, $\mathcal{F} = \mathcal{F}_{(0)} + \mathcal{F}_{(1)} + \mathcal{F}_{(2)}$, where $\mathcal{F}_{(0)} = -B_{\rm bg}^2/2$ (order $B_{\rm bg}^2$, wave-free), $\mathcal{F}_{(1)} = -B_{\rm bg} \cdot \mathbf{B}_w$ (order $B_{\rm bg}^1$, linear in waves), and $\mathcal{F}_{(2)} = (|\mathbf{E}_w|^2 - |\mathbf{B}_w|^2)/2$ (order $B_{\rm bg}^0$, quadratic in waves).

The HE Lagrangian $\mathcal{L}_{\rm NL} = \kappa\mathcal{F}^2$ generates the polarization $\mathbf{P}_i = \partial\mathcal{L}_{\rm NL}/\partial E_i = 4\kappa\mathcal{F}\, E_i$. The cross-coupling polarization, $\mathbf{P}_{\rm cross} \equiv \mathbf{P}(A+B) - \mathbf{P}(A) - \mathbf{P}(B)$, extracts only the terms that contain both wave amplitudes. A direct expansion gives

$$\mathbf{P}_{\rm cross} = 4\kappa\Big[-2B_{\rm bg} \cdot (E_A\mathbf{B}_B + E_B\mathbf{B}_A) + (\mathbf{E}_A \cdot \mathbf{E}_B - \mathbf{B}_A \cdot \mathbf{B}_B)(\mathbf{E}_A + \mathbf{E}_B)\Big], \tag{25}$$

which contains terms scaling as $B_{\rm bg}^1 \cdot \text{wave}^2$ and as $B_{\rm bg}^0 \cdot \text{wave}^3$. The measured linearity residual is the time-integrated effect of $\mathbf{P}_{\rm cross}$ on the field evolution and inherits both channels with geometry-dependent (but $B_{\rm bg}$-independent) coefficients. The fitted $\beta B_{\rm bg}^2$ contribution observed in Table 3 must therefore arise either from higher-order Lagrangian terms (absent in the $\mathcal{F}^2$ truncation) or from a residual-diagnostic geometric factor that we treat as empirical; likewise the cubic $\gamma B_{\rm bg}^3$ channel has no leading-order analytic source and is the central interpretive question of the present work.

C Software organization

The framework is implemented in Python with NumPy as the only required runtime dependency; Matplotlib is used for figures. Principal modules:

- `fdtd2d.py`: `Grid`, `FDTD2D`, `Material` base class, and the four physical-regime subclasses (`KerrMaterial`, `RelativisticPlasmaMaterial`, `HeisenbergEulerMaterial`). Also contains the `CPML2D` convolutional-PML implementation, `add_gaussian_packet` initial-condition helper, and `add_uniform_static_B` background-field helper.
- `fdtd2d_mixed.py`: `FDTD2D_Mixed` class supporting all six field components with linearized HE constitutive, plus `add_TM_packet` and `add_TE_packet`.
- `source.py`: `PlaneWaveSource` TF/SF source class.
- `diagnostics.py`: `ProbeRecorder` for time-domain field sampling and FFT spectral analysis.
- Driver scripts (`run_*.py`) for each physics regime and diagnostic application: `run_two_wave.py`, `run_kerr.py`, `run_plasma.py`, `run_qed.py`, `run_spectra.py`, `run_magnetar.py`, `scan_magnetar_bg.py` `run_mixed_demo.py`.
- Validation scripts (`validate_*.py`): `validate_boris.py`, `validate_qed_pc.py`, `validate_cpml.py`, `debug_cpml.py`.

References

[1] R.W. Boyd, *Nonlinear Optics*, 4th ed. (Academic Press, 2020).

[2] W.L. Kruer, *The Physics of Laser Plasma Interactions* (Westview Press, 2003).

[3] W. Heisenberg and H. Euler, "Consequences of Dirac's theory of the positron," Z. Phys. **98**, 714 (1936).

[4] J. Schwinger, "On gauge invariance and vacuum polarization," Phys. Rev. **82**, 664 (1951).

[5] A. Ringwald, "Pair production from vacuum at the focus of an X-ray free electron laser," Phys. Lett. B **510**, 107 (2001).

[6] A. Di Piazza, C. Müller, K. Z. Hatsagortsyan, and C. H. Keitel, "Extremely high-intensity laser interactions with fundamental quantum systems," Rev. Mod. Phys. **84**, 1177 (2012).

[7] R.P. Mignani *et al.*, "Evidence for vacuum birefringence from the first optical-polarimetry measurement of the isolated neutron star RX J1856.5−3754," Mon. Not. R. Astron. Soc. **465**, 492 (2017).

[8] K.S. Yee, "Numerical solution of initial boundary value problems involving Maxwell's equations in isotropic media," IEEE Trans. Antennas Propag. **14**, 302 (1966).

[9] J.P. Boris, "Relativistic plasma simulation–optimization of a hybrid code," in *Proc. 4th Conf. Numerical Simulation of Plasmas*, Naval Research Laboratory (1970), p. 3.

[10] G. Mur, "Absorbing boundary conditions for the finite-difference approximation of the time-domain electromagnetic-field equations," IEEE Trans. Electromagn. Compat. **23**, 377 (1981).

[11] J.A. Roden and S.D. Gedney, "Convolution PML (CPML): An efficient FDTD implementation of the CFS–PML for arbitrary media," Microwave Opt. Tech. Lett. **27**, 334–339 (2000).

[12] A. Taflove and S.C. Hagness, *Computational Electrodynamics: The Finite-Difference Time-Domain Method*, 3rd ed. (Artech, 2005).

[13] J.S. Heyl and L. Hernquist, "Birefringence and dichroism of the QED vacuum," J. Phys. A **30**, 6485 (1997).