ACCELERATOR DESIGN EDUCATIONAL PRIMER – CONCEPTUALIZING AND OPTIMIZING THE HYBRID LHEC-LIKE ELECTRON-ION COLLIDER DESIGN*

A. Seryi[1†], A. Formenti[2], A. Huebl[2], R. Lehe[2], C. A. Lindstrøm[3], C. Mitchell[2], T. Pieloni[4], J. Qiang[2], L. van Riesen-Haupt[1], L. Vanhecke[4], J.-L. Vay[2], E. Zoni[2]
[1]Old Dominion University, Norfolk, VA, USA
[2]Lawrence Berkeley National Laboratory, Berkeley, CA, USA
[3]Department of Physics, University of Oslo, Oslo, Norway
[4]École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland

Abstract

The Electron-Ion Collider (EIC) Mission Need requires $\sqrt{s}$ = 20–100 GeV (upgradable to 140 GeV) and luminosity 10^{33}–10^{34} cm^{-2} s^{-1}. The current ring–ring baseline achieves the full scope, including ~10^{34} cm^{-2} s^{-1} across all energies. However, when the design is re-optimized for the lower boundary – accepting 10^{33} cm^{-2} s^{-1} and prioritizing cost – an alternative configuration emerges as more advantageous: a hybrid LHeC-like electron accelerator using multi-pass energy recovery linacs (ERL). This solution reduces electron-beam power by roughly an order of magnitude, yielding nearly a factor of two reduction in total project cost compared with the present baseline while still satisfying the minimum physics requirements. The study performs parametric cost and performance modeling, augmented by AI-driven optimization, to explore this design space. Serving primarily as an educational exercise for the next generation of accelerator experts the paper demonstrates modern design methods: rapid parametric scans, cost-driven optimization, and integration of AI tools. It examines technical feasibility, identifies critical R&D (high-current ERL operation, beam–beam effects, synchronization, etc.), and discusses how such re-optimization studies can be used to train designers in an era when AI dramatically expands exploration of complex accelerator parameter spaces.

PREAMBLE – AI AND EDUCATION

The emergence of AI-assisted coding presents both unique challenges and unprecedented opportunities for accelerator physics education. Institutions like Fermilab and CERN, which host accelerator school USPAS and CAS, need to rapidly adapt their pedagogical approaches, and the Genesis program [1] in the US further encourages AI for accelerated design of particle accelerators. A natural vehicle for developing design skills is the mini-project such as Refs. [2, 3], where students tackle a realistic challenge in a short period. This paper rehearses what such a mini-project looks like when AI tools are available – from parameter optimization and code generation to cost estimation – serving as a draft script for future AI-assisted USPAS or CAS courses.

* Developed with assistance from AI such as Claude, Gemini, etc.

† aseryi@odu.edu

EIC-LHeC MINI-PROJECT

The mini-project assignment is: *Given the EIC proton beam stored in the RHIC-derived Hadron Storage Ring, design the simplest and most cost-effective electron accelerator complex satisfying the EIC CD0 Mission Need lower boundary – $\sqrt{s} \approx$ 29–100 GeV and $\mathcal{L} \sim 10^{33}$ cm^{-2}s^{-1} – and estimate its cost relative to the ring-ring baseline.*

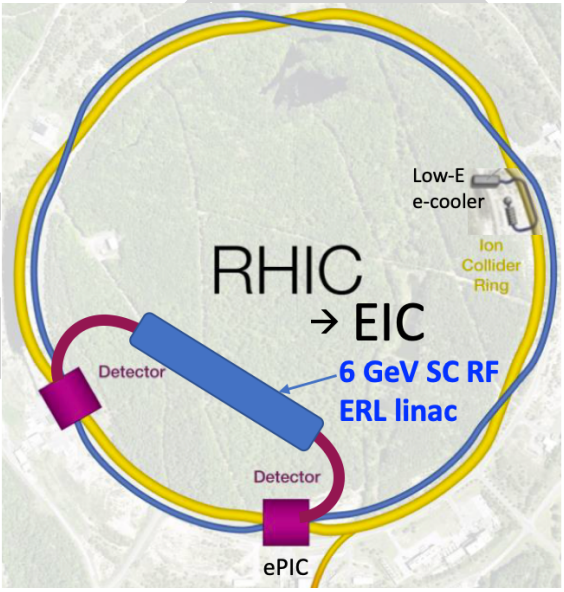

Figure 1: Schematic layout of the EIC built as an LHeC-like hybrid with energy recovery.

The motivation comes from the ESPP Open Symposium [4], where it was argued that inserting a new electron ring into the LHC tunnel for LHeC would be cost-prohibitive and luminosity-limiting [5]. This prompted the question: could replacing the EIC electron storage ring with a superconducting energy-recovery linac (ERL) yield a similar cost-performance advantage? The proposed hybrid keeps the full EIC proton infrastructure intact and replaces the electron ring with a multi-pass ERL: electrons are accelerated in a superconducting radio-frequency (SRF) linac, recirculated through arcs, brought to collision, and then decelerated back through the linac for energy recovery. Key advantages over the ring-ring baseline are: electron beam power reduced by one order of magnitude [each bunch used once, such

that a large interaction-point (IP) disruption is acceptable]; no large electron storage ring, substantially reducing civil and magnet costs; and a 6 GeV linac traversed three times covering the full 18 GeV EIC energy range.

The layout is shown in Fig. 1. The ERL linac is installed in a new tunnel in the South-West part[1] of the RHIC inner area. Recirculation arcs routed through new tunnels, delivering the electron beam to the ePIC interaction region (IR).

The mini-project proceeds through the following steps, each carried out with AI assistance and described in subsequent sections: beam parameter and luminosity modeling; beam–beam evaluation; IR optics and magnet technology assessment; energy recovery analysis; cost estimation; and R&D identification. At each step, AI tools wrote Python scripts and draft text, while the "students" (here, the authors) formulated physically meaningful prompts and verified outputs against accelerator physics principles.

EIC-LHeC PARAMETERS

Approach and AI Prompting Strategy The parameter development followed an iterative dialogue with AI tools, beginning from the two parent designs and converging on a self-consistent hybrid. The initial prompts requested Python codes to compute beam parameters, beam sizes at the IP, luminosity, and beam–beam tune shifts $\xi_{x,y}$ and RMS divergences for the standard EIC and LHeC designs separately, using published parameter tables [7, 8] as inputs. Details of prompts, outputs, and generated codes are available at [9].

Baseline EIC and LHeC Reference Parameters Table 1 shows abridged parameter sets for the two parent designs; full tables are in [7,8] and review of both designs in Ref. [10].

Table 1: Abridged LHeC Reference Parameters (Phase-1, 20 GeV e on 7 TeV p) and Baseline EIC (at 105 GeV CM)

Parameter	EIC	LHeC
e / p energy, GeV	10 / 275	20 / 7000
$\sqrt{s}$, GeV	105	749
Number of bunches	1160	2808
Collision frequency, Hz	9.1×10^7	4.0×10^7
e bunch population	1.72×10^{11}	0.94×10^{10}
p bunch population	6.9×10^{10}	2.2×10^{11}
e norm. emitt. h/v, µm	391 / 26	22 / 22
p norm. emitt. h/v µm	3.3 / 0.3	2.5 / 2.5
e $\beta^*_{x/y}$ at IP, m	0.45 / 0.06	0.2 / 0.2
p $\beta^*_{x/y}$ at IP, m	0.80 / 0.07	0.35 / 0.35
e $\sigma_{x/y}$ at IP, µm	94.8 / 8.6	10.6 / 10.6
p $\sigma_{x/y}$ at IP, µm	94.9 / 8.6	10.7 / 10.7
e / p beam current, mA	2500 / 1000	60.1 / 1410
p beam-beam $\xi_{x/y}$	0.012 / 0.012	0.001 / 0.001
Peak lumi. $\mathcal{L}$, cm^{-2}s^{-1}	1.05×10^{34}	6.0×10^{33}

Resolving Parameter Inconsistencies in the Hybrid When asked to combine EIC proton parameters with LHeC electron parameters, the AI correctly identified two fundamental conflicts before producing any code. (i) *Bunch rate mismatch:* the EIC ring (3834 m circumference, 1160 bunches) gives a collision frequency of ~91 MHz, while the LHeC electron linac assumes ~40 MHz; naively combining the two implies ~140 mA electron current, far beyond the 60 mA LHeC design value. (ii) *Beam size mismatch:* LHeC electron optics ($\beta^*_{x,y}$ = 0.20 m, ε_n = 22 µm) were optimized for the ~10 µm LHC proton beam, whereas EIC protons at 275 GeV give $\sigma_x \approx 48$ µm at the same β^*. The resolution was to reduce the number of proton bunches by a factor of 3 (from 1160 to 386), proportionally increasing bunch population at constant 1.0 A proton current. This lowers the collision frequency to ~30 MHz, restoring a consistent ~58 mA electron current (with 1.9 nC bunch charge, somewhat higher than 1.5 nC of LHeC), and the IP optics were re-optimized to match beam sizes (β^*_x = 1.20 m, β^*_y = 0.05 m for electrons; β^*_x = 0.12 m, β^*_y = 0.06 m for protons). The AI also noted that ξ_e is unconstrained in a linac-ring scheme since electrons are discarded after each pass; only the proton beam–beam parameter was constrained to $\xi_p < 0.01$. A further step explored $\mathcal{L} = 10^{34}$ cm^{-2}s^{-1} via: (i) proton emittance reduction through coherent electron cooling; and (ii) vertical electron emittance reduction by a factor of 10 using a magnetized gun and Derbenev's flat-beam transformer [11].

Hybrid EIC-LHeC Parameter Sets Two configurations are given in Tables 2 and 3, both for $\sqrt{s} \approx 105$ GeV (10 GeV electrons, 275 GeV protons).

Table 2: Hybrid EIC-LHeC, Initial Configuration

Parameter	e	p	Unit
E	10.0	275.0	GeV
N	1.19×10^{10}	2.1×10^{11}	
n_{bunches}	386		
f_{collis}	3.03×10^7		Hz
Norm. ε h/v	22.0 / 22.0	3.3 / 0.3	µm
$\beta^*_{x/y}$ at IP	1.20 / 0.05	0.12 / 0.06	m
σ_z	3.0	70.0	mm
$\sigma_{x/y}$ at IP	36.7 / 7.5	36.8 / 7.8	µm
$\sigma'_{x/y}$ at IP	30.6 / 149.9	306.3 / 130.6	µrad
$\xi_{x/y}$	3.474 / 0.679	0.001 / 0.002	
I_{beam}	57.6	1000	mA
Peak $\mathcal{L}$	2.1×10^{33}		cm^{-2}s^{-1}

Table 3: Hybrid EIC, optimized with e-cooling and flat-beam optics. Unchanged parameters as in Table 2.

Parameter	e	p	Unit
Norm. ε h/v	22.0 / 2.2	0.9 / 0.08	µm
$\beta^*_{x/y}$ at IP	0.25 / 0.10	0.09 / 0.05	m
$\sigma_{x/y}$ at IP	16.8 / 3.4	16.8 / 3.6	µm
$\sigma'_{x/y}$ at IP	66.8 / 32.7	182.4 / 77.8	µrad
$\xi_{x/y}$	3.474 / 6.789	0.003 / 0.007	
Peak $\mathcal{L}$	1.0×10^{34}		cm^{-2}s^{-1}

The initial hybrid satisfies the CD0 lower boundary with unmodified EIC proton emittances; $\xi_p < 0.01$ in both con-

[1] Sited to avoid the Peconic River protected watershed boundary [6].

figurations. The large electron ξ_e reflects severe single-pass disruption – acceptable since electrons are discarded after collision, but raising questions about ERL acceptance for the disrupted bunch. The optimized case reaches 10^{34} cm^{-2}s^{-1} through emittance reductions; proton divergences up to 182 μrad horizontally imply tight FD aperture requirements.

BEAM–BEAM EVALUATIONS

Beam–beam collisions, highly asymmetrical for EIC hybrid, can be evaluated using existing codes [12, 13], but can be also studied using AI-generated code. A simplified 2D beam–beam code was written by Claude AI in a few minutes from a paragraph-long prompt (the IT community describe this process as vibe coding).

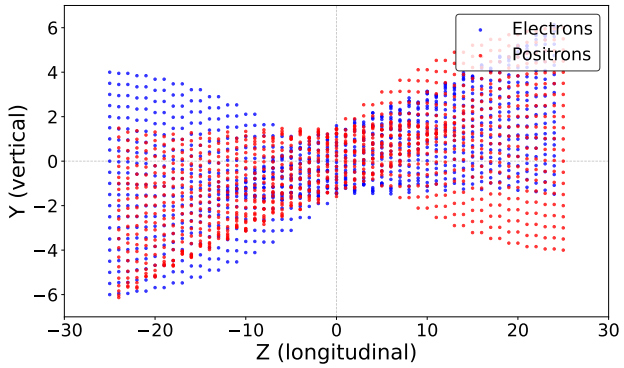

Figure 2: Beam–beam collision calculated by the AI-generated code. Normalized units. Electrons are moving right. Positrons (or protons) are moving left.

That prompt is reproduced here to illustrate the level of physical detail required to obtain useful AI-generated code. The prompt specified: a 2D head-on collision (longitudinal Z and vertical Y coordinates only); each bunch represented by 50 transverse slices of 21 particles, placed equidistantly along Z with spacing $dZ = 1$ and vertical particle separation $dY = 0.5$; zero initial transverse velocities; and a simplified model for the beam–beam force in which the vertical electric field acting on a particle at position Y_1 equals the number of particles of the opposing slice with $Y > Y_1$ minus those with $Y < Y_1$ – the thin-slab approximation appropriate when the horizontal extent of the bunch is much larger than its vertical size. The kick coefficient D_Y was requested to be tuned so that a particle undergoes approximately one betatron oscillation during the full collision. Visualization of the collision dynamics was requested from the outset. The AI (Claude) produced working, animated code on the first attempt. Five short follow-up prompts then corrected the kick sign, tuned D_Y, added offset collisions, and improved the visualization display. The entire exchange – initial prompt plus follow-up adjustments – took a few minutes and required no debugging of code syntax by the author. The prompt and the resulting Python code are available at [9] and beam–beam simulation is illustrated in Fig. 2.

This capability suggests a transformative shift in how students learn. Rather than passively attending lectures or running decades-old legacy codes, students could actively create simplified simulation codes, then iteratively refine them through AI-assisted modifications to explore different physics phenomena.

For the specific example of EIC-LHeC hybrid mini-project, students could progressively enhance their understanding by first converting the code to use of physical units; then implementing realistic beam distributions with emittance and beta-functions; adjusting input parameters to match those of EIC-LHeC hybrid; adding luminosity calculations; investigating how luminosity depends on asymmetry of disruption parameters and beam offsets; measuring angular distribution of highly disrupted electron beam; evaluating the hour-glass effects or crab-crossing configurations.

IR OPTICS AND MAGNETS

Approach and AI Prompting Strategy The IR optics design focused on the Final Doublet (FD), the pair of quadrupoles that converts the diverging beam emerging from the IP into a parallel beam downstream. The prompt specified: use thick-lens transfer matrices; trace four rays with zero position at the IP and initial angles equal to $\pm\sigma'_{x,y}$ from Table 2; assume $L^* = 2$ m (IP to entrance of Q1), quadrupole length $L_q = 2$ m, and inter-quadrupole gap 0.5 m; solve for the quadrupole strengths such that the beam exits the FD as a parallel beam in both planes simultaneously (condition $M_{22} = 0$ in the full transfer matrix); try both QF–QD and QD–QF polarities and select the one that minimizes the maximum beam size inside the FD; plot X and Y beam envelopes and save the figures to PNG files. Claude generated working Python code in a single pass, requiring no manual debugging [9], and independently reproduced by Gemini.

FD Design Results For the beam with IP parameters from Table 2, for the proton beam the solver selects QF–QD polarity with $k_1 = +0.319$ m^{-2} and $k_2 = -0.167$ m^{-2}, corresponding to physical gradients $G_1 = 293$ T/m and $G_2 = 153$ T/m (for $B\rho = 917$ T·m). For the electron beam the complementary QD–QF polarity is optimal, giving $G_1 = 10.6$ T/m and $G_2 = 5.6$ T/m (at $B\rho = 33.4$ T·m). The maximum beam envelope inside the FD (shown in Fig. 3) reaches 0.82 mm and 1.43 mm (horizontal and vertical) for protons; and 0.34 mm and 0.40 mm for electrons.

Magnet Technology Assessment For a quadrupole of gradient G, the maximum achievable aperture radius for a given superconducting (SC) technology with peak bore field $B_{\rm bore}$ is $R_{\rm max} = B_{\rm bore}/G$. The aperture-to-beam-size ratios $N_X = R_{\rm max}/\sigma_x^{\rm max}$ and $N_Y = R_{\rm max}/\sigma_y^{\rm max}$ quantify the available margin; from collider IR experience $N_{X,Y} \gtrsim 15$ is required to accommodate orbit offsets, beam halo, and tolerances. Assuming that magnet technologies, NbTi at 4.2 K, NbTi at 1.9 K, and Nb_3Sn at 4.2 K, allow peak bore field of 7 T, 8.5 T and 12 T correspondingly, the EIC-LHeC hybrid proton FD (the most constrained) was analyzed with Gemini AI, varying the L^* parameter. The resulting algorithm took into account extra aperture for shielding, cold-bore, insulation and beam screen based on scaling from FCC-hh [14]

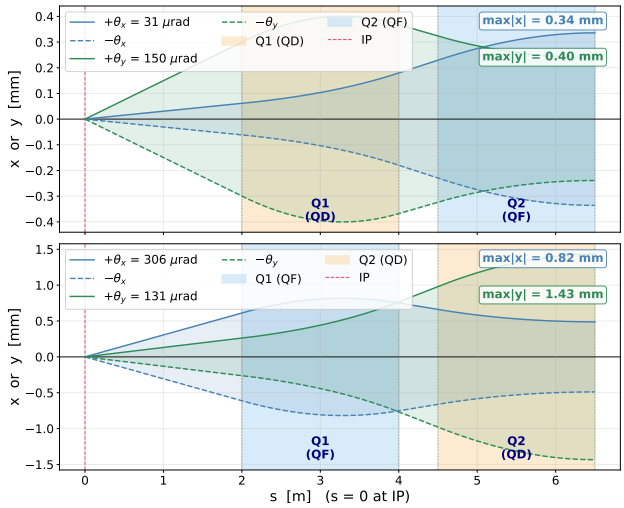

Figure 3: Hybrid EIC-LHeC FD produced via vibe-coding. Electrons (10 GeV, top plot) and protons (275 GeV, bottom).

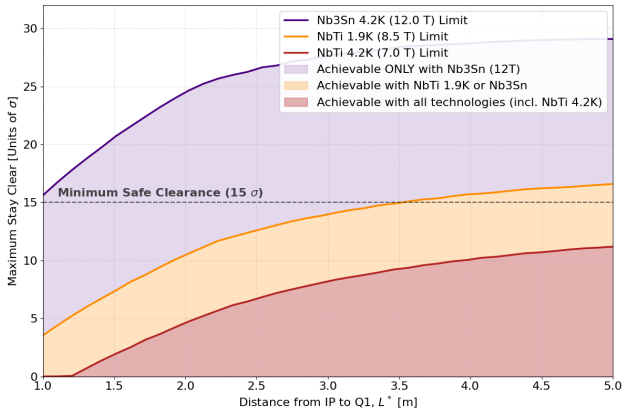

Figure 4: Hybrid EIC-LHeC FD stay-clear vs L^* and SC magnet technology analyzed via vibe-coding (with 19 mm shielding and internals subtracted).

and redistributed the 4 m magnetic length to maximize beam stay clear, the result is shown in Fig. 4, indicating the preference for NbTi at 1.9 K. For the electron FD, the required gradients are so low that SC magnets are in principle unnecessary; however SC approach may be required for practical purposes of integration within tightly constrained FD.

ENERGY RECOVERY EVALUATIONS

Approach and AI Prompting Strategy A detailed prompt was given to Claude specifying all relevant beam parameters for the EIC-LHeC hybrid: 1.5 nC bunch charge, 1.5 mm bunch length, 7 MeV injection energy, 10 GeV SRF linac gain, 801.58 MHz RF frequency, and 40.08 MHz bunch repetition rate, yielding 60 mA in the arcs and 120 mA in the linac (accelerating plus decelerating bunches). The prompt requested a comprehensive simulation of a single-pass ERL and a two-pass ERL (5 GeV per pass), including: longitudinal phase-space tracking through acceleration, an electron-proton beam–beam collision, 180-degree phase reversal in the return arc, and deceleration; energy recovery efficiency as a function of deceleration-phase error; a detailed wall-plug power budget; and comparison with a conventional storage ring.

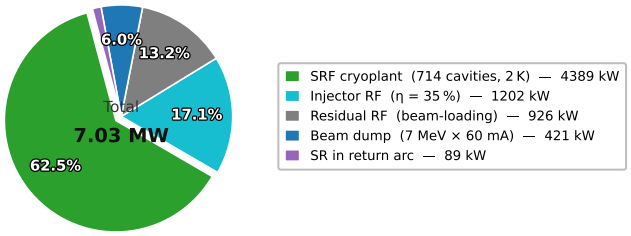

Figure 5: Single-pass ERL wall-plug power budget (10 GeV, 60 mA, f_{RF} = 801.58 MHz) produced by vibe-coding.

Energy Recovery Results The AI model produced a code at the first iteration, rich with physics, including Yokoya-Chen beam–beam approximation, SR losses, comparison of single or two pass ERL with rings of different size – offering vast opportunities for discussion with students about validity and applicability of approaches [15]. The simulation confirms an energy recovery efficiency of 99.93%, reducing the wall-plug power from 602 MW (no recovery) to 7.0 MW. Figure 5 shows the single-pass power budget breakdown.

COST ESTIMATIONS

Cost estimates for the hybrid electron complex, IR, and detector were developed using AI-assisted parametric scaling from the LHeC study [16], excluding upgrades to proton infrastructure. Two linac configurations were compared: a single-pass 10 GeV SRF linac and a recirculating 6 GeV linac traversed twice (or three times for 18 GeV), analogous to CEBAF and the LHeC design. The recirculating option approximately halves the dominant SRF, civil, and cryo costs while adding modest arc costs (70 MCHF, 100 PY), saving ~62 MCHF and ~250 person-years (PY) overall – making it the preferred option. Estimated totals are: electron complex 280–450 MCHF (450–720 PY), IR 50 MCHF (50 PY), detector 250 MCHF (100 PY), giving a grand total of 580–750 MCHF (600–870 PY). These ranges were obtained from estimates by two AI models [17]; IR and detector line items were manually increased above the raw AI output.

Converting to USA DOE accounting and adding 60% contingency reflecting the preliminary nature of the estimate, yields the summary in Table 4. The hybrid EIC (with 6 GeV recirculating linac) total of 1.4–1.9 B USD is comfortably below the upper bound of the CD1-approved EIC cost range of $1.7–2.8 B – confirming the earlier stated intuitive motivation that led to evaluation of this hybrid EIC concept.

A method for more advanced cost optimization for this hybrid machine that could be investigated by students is that used by the plasma–RF hybrid collider concept HALHF [18]. This method, described in Refs. [19, 20], first builds a simplified physics model of all the major subsystems of the collider, then uses a small number of key input parameters (energies, collision frequency, accelerating gradients, etc.) to calculate lengths, power and other requirement for subsystems, and finally converts this to an overall cost via cost estimates from similar machines. The global cost optimum is then found

Table 4: Hybrid EIC cost range, estimated by two independent AI-assisted scalings of the LHeC CDR cost. The two estimates agree to within ∼30 % at the pre-contingency level.

Item	Value (MUSD)
Material costs (580–750 MCHF × 1.25)	725 – 935
Labor costs (600–870 PY × 300k USD)	180 – 260
Pre-contingency total	905 – 1195
Contingency (60%)	545 – 720
Total with contingency	**1450 – 1915**

(approximately) using Bayesian optimization; a computationally economical machine-learning optimization method. While for HALHF this was performed using the start-to-end simulation framework ABEL [21, 22], a more simplified code for the same purpose is available at [23] – the latter being more suitable for AI training and code generation.

REQUIRED R&D AND NEXT STEPS

The hybrid EIC-LHeC concept rests on several assumptions requiring dedicated R&D before the design can be considered credible at project-proposal level. The principal challenges, identified in dialogue with AI, are the following. *High-current ERL operation:* fractional energy loss $\lesssim 10^{-4}$ is required at 58 mA and 10 GeV; a staged CEBAF test (high energy, low current) and PERLE (low energy, high current) would bound the two critical parameters. *Polarized source current:* 58 mA far exceeds the state of the art (∼1 mA CW); multiplexing N sources via RF mergers is a plausible path once 10 mA is demonstrated at Jefferson Lab. *IR final doublet:* small proton β^* (0.09–0.12 m) drives IP divergences of ∼300 µrad, requiring large-aperture SC quadrupoles, tight collimation, and MDI compatibility with ePIC. *Intra-beam scattering (IBS):* the factor-of-three increase in proton bunch charge raises the IBS rate; electron cooling must maintain the emittances assumed in the 10^{34} configuration. *Disrupted-beam energy recovery:* with $\xi_y^e \approx 6.8$, the electron bunch is severely disrupted; whether the resulting phase space fits within ERL return-arc acceptance requires dedicated simulation. Beam halo, instability, and their mitigation in the ERL, are to be studied as well.

Logical next steps are: beam–beam parameter scans with strong-strong codes; conceptual IR lattice and MDI design; ERL lattice and disrupted-beam acceptance study; IBS and cooling-power modeling; bottom-up cost re-estimation; and programmatic compatibility assessment with the EIC CD schedule. The hybrid concept illustrates how AI-assisted exploration can rapidly map an alternative design space and identify the critical questions – accelerating the early conceptual phase without replacing rigorous engineering analysis.

VERIFICATION OF AI RESULTS

Verification of AI-generated results follows a three-pronged approach, all of which reinforce our belief that students will learn accelerator physics from this educational style rather than be displaced by it.

The first prong is human verification by the student. Returning to the beam–beam example introduced earlier, the student progresses from a toy 2D simulation to increasingly realistic versions, where at each step the student is encouraged to design diagnostic tests that interrogate the AI-generated code: does momentum balance? Does the code reproduce the analytic Yokoya–Chen pinch enhancement? Does the disruption scale as expected with bunch charge? Such tests turn the AI from an oracle into a collaborator whose output the student must judge against first principles.

The second prong is cross-verification by a second, independent AI model. The parameter tables of this paper were initially generated with an AI model in summer 2025 and subsequently cross-checked with Claude, which flagged a small number of internal inconsistencies that have since been corrected [24]. Independent re-derivation by a different model is cheap, fast, and surprisingly effective at catching unit slips, mixed-convention errors, and copy-paste residues that single-model workflows tend to propagate [25, 26].

The third prong is benchmarking against established first-principles simulations. AI-generated results are compared against well-validated codes whose outputs have been extensively benchmarked against analytic theory and experimental data from previous colliders, providing an independent, quantitative standard against which the AI-assisted workflow must stand. This closes the verification loop: student judgment, cross-model consistency, and fidelity to the existing simulation heritage all reinforce one another.

CONNECTION TO GENESIS PROGRAM

This educational primer directly supports the Genesis Mission by demonstrating AI-accelerated **design** and **optimization** – two of the three pillars (design, optimization, and operations) emphasized in Challenge Area 13 ("Enhancing Particle Accelerators for Discovery") of the DOE RFA [27].

By employing AI for parameter definition, simulations, and cost estimations, the mini-project compresses the conceptual design cycle, identifies cost-effective configurations (such as the multi-pass ERL EIC concept), and trains the next generation of accelerator scientists – precisely the transformative capabilities sought by the Genesis Mission to advance U.S. leadership in particle physics and energy science.

SELF-ASSESSMENT AND NEXT STEPS

The approach sketched here is a first step, not a finished recipe – pedagogy, mini-project selection, and verification protocols all need refinement through actual classroom use. With those caveats, we believe it points in the right direction for education in the era of human-AI partnership.

The authors acknowledge support by the Old Dominion University and CHART Program [28]. Supported in part by the U.S. DOE's Genesis Mission and the Office of Science, Office of Advanced Scientific Computing Research's ModCon under Contract No. AC02-05CH11231.

REFERENCES

[1] Genesis – a national mission to accelerate science through Artificial Intelligence, 2025, https://genesis.energy.gov/

[2] M. Turner *et al.*, "Compact ring-based x-ray source with on-orbit and on-energy laser-plasma injection", in *Proc. NA-PAC'16*, Chicago, IL, USA, pp. 435–438, Oct. 2016. doi:10.18429/JACoW-NAPAC2016-TUA3CO03

[3] G. Tiwari *et al.*, "EUV FEL light source based on energy recovery linac with on-orbit laser plasma injection", in *Proc. IPAC'24*, Nashville, TN, USA, pp. 408–411, Jul. 2024. doi:10.18429/JACoW-IPAC2024-MOPG66

[4] The European Strategy for Particle Physics (ESPP) open symposium, Jul. 2025, https://agenda.infn.it/event/44943/

[5] J. D'Hondt, Presentation of LHeC project at ESPP 2025, questions and answers, Jul. 2025, https://youtu.be/NX1eYfol5JI?list=PLbsqUzxZlcP5p4izedCWb7EOP-ZkSH495%5C&t=2703

[6] Peconic river protected watershed boundary, https://www.peconicestuary.org/news-and-events/maps-gis/maps-watershed-boundaries/

[7] S. Peggs and T. Satogata, Electron Ion Collider parameter list summary, Apr. 2026, https://eic.jlab.org/Documents/EIC-General/EIC%5C_ParameterList.pdf

[8] K. Piotrzkowski *et al.*, Phase-one LHeC submission to ESPP 2026 strategy update, 2026, https://indico.cern.ch/event/1439855/contributions/6461558/

[9] EIC-LHeC vibe-coding prompts and codes, 2025, https://github.com/aseryi/vibe-coding/

[10] O. Brüning, A. Seryi, and S. Verdú-Andrés, "Electron-Hadron Colliders: EIC, LHeC and FCC-eh", *Front. Phys.*, vol. 10, p. 886473, 2022. doi:10.3389/fphy.2022.886473

[11] Y. Derbenev, "Adapting Optics for High Energy Electron Cooling", Univ. of Michigan, Ann Arbor, MI, USA, Rep. UM-HE-98-04, Feb. 1998.

[12] J. Qiang, M. Furman, and R. Ryne, "A parallel particle-in-cell model for beam-beam interactions in high energy ring colliders", *J. Comp. Phys.*, vol. 198, pp. 278–294, 2004. doi:10.1016/j.jcp.2004.01.008

[13] X. Buffat *et al.*, COMBI, (COherent Multibunch Beam-beam Interactions), 2018, https://twiki.cern.ch/twiki/bin/view/ABPComputing/COMBI

[14] L. van Riesen-Haupt, J. L. Abelleira, E. C. Alaniz, and A. Seryi, "An optimised triplet for the final focus of the FCC-HH with a 40 m final drift", in *Proc. IPAC'18*, Vancouver, Canada, Apr.-May 2018, pp. 364–367. doi:10.18429/JACoW-IPAC2018-MOPMK007

[15] ERL simulation – model assumptions & discussion guide, AI generated summary, Apr. 2026, https://github.com/aseryi/vibe-coding/blob/main/EIC-LHeC/ERL_simulation_assumptions.md

[16] O. Bruning, "LHeC cost ostimate", CERN, Geneva, Switzerland, Rep. CERN-ACC-2018-0061, 2018. https://cds.cern.ch/record/2652349/

[17] EIC-LHeC costs verification – see sub-folder verify-EIC-LHeC-cost, Apr. 2026, https://github.com/aseryi/vibe-coding/

[18] B. Foster, R. D'Arcy, and C. A. Lindstrøm, "A hybrid, asymmetric, linear Higgs factory based on plasma-wakefield and radio-frequency acceleration", *New J. Phys.*, vol. 25, p. 093037, 2023. doi:10.1088/1367-2630/acf395

[19] E. Adli *et al.*, "HALHF: a hybrid, asymmetric, linear Higgs factory using plasma- and RF-based acceleration. Backup document", *arXiv*, Mar. 2025. doi:10.48550/arXiv.2503.23489

[20] C. A. Lindstrøm *et al.*, "Updated baseline design for HALHF: the hybrid, asymmetric, linear Higgs factory", in *Proc. IPAC'25*, Taipei, Taiwan, pp. 53–56, Jun. 2025. doi:10.18429/JACoW-IPAC2025-MOCD1

[21] J. Chen *et al.*, "ABEL: the adaptable beginning-to-end linac simulation framework", in *Proc. IPAC'25*, Taipei, Taiwan, pp. 1412–1415, Jun. 2025. doi:10.18429/JACoW-IPAC2025-TUPS012

[22] ABEL cost modelling implementation, 2025, https://github.com/abel-framework/ABEL/blob/main/abel/classes/cost%5C_modeled.py

[23] A simple collider cost model, 2025, https://github.com/carlandreaslindstrom/ColliderCostModel

[24] EIC-LHeC parameters verification – see sub-folder verify-EIC-LHeC-pars, Apr. 2026, https://github.com/aseryi/vibe-coding/

[25] M. L. Arbuzov, A. A. Shvets, and S. Beir, "Beyond Exponential Decay: Rethinking Error Accumulation in Large Language Models", *arXiv preprint arXiv:2505.24187v2*, 2026. doi:10.48550/arXiv.2505.24187

[26] Z. Gan, Y. Liao, and Y. Liu, "Rethinking external slow-thinking: from snowball errors to probability of correct reasoning", in *Proc. ICML'25*, vol. 267, Vancouver, Canada, pp. 18170–18188, Jul. 2025. doi:10.48550/arXiv.2501.15602

[27] "The Genesis mission: transforming science and energy with AI", Announcement No. DE-FOA-0003612, Mar. 2026,

[28] Swiss accelerator research and technology (CHART) program, https://www.chart.ch