Toward operation of high-intensity compact accelerators at small facilities via an AI agent: R.E.G.I.N.A.L.D.

Daniel Winklehner and Janet Conrad
LNS GENESIS Workshop, 2026

Focus area 13.A
Enhancing Particle Accelerators for Discovery:
AI Driven Accelerator Facilities

Partnership

MIT (Lead Institution, IHE) Collaborators

Daniel Winklehner (Project Lead, accelerator surrogate modeling expert)

Janet Conrad (Supervision of Students, advised 5 AI-based theses since 2015)

Josh Villarreal (AI patent pending, Physics, Stats & Data Science GS, Course in Reinforcement Learning)

Claire Huchthausen (GS studying acc. phys, COMSOL and IBSimu expertise)

Niklas Franklin (Undergraduate summer student from U.N.M.)

David McCLain (PD with accelerator experience)

IBA RadioPharma Solutions (Industry Partner) Collaborators

Eric Forton (Technologies Team Lead, expertise on operating requirements)
Emanuele Festa (R&D Project Manager)
Laurent Maunoury (Ion Source Lead at IBA)

Notes on collaboration:

- Above collaborators are confirmed for Phase I and II.
- IBA is a multinational company with manufacturing in Virginia
- It is a strategic choice to not include a National Laboratory in Phase I. Our past proposals have reviewed well because of university-origin and independence from the labs. Our proposal will stand out as IHE-industry rather than IHE-lab.
- For Phase II, we already have two strongly interested additional partners:
  - * LANL – verbal commitment from Bruce Carlsten
  - * Draper Engineering – must run this past upper management and lawyers, but agreement is expected

Important Background Information

Our prototype cyclotron is under construction by IBA, funded by **DOE – IRP** and DOE-ARDAP.

IBA is participating in the science of the prototype tests under a no-cost collaboration agreement
A similar no-cost collaboration agreement will be put in place for this project.
which proceeds in parallel with the prototype tests.
→Existing agreement simplifies Phase I organization
and IBA does not appear in the budget.

This gives us a prototype and its subsystems for deployment of REGINALD.
The prototype timeline is:

- by 9/1/2026, Installation of the MIST-2 ion source
- by late 2026, new RFQ being installed & ready to test
- by December 2027, complete prototype accelerates to 1.5 MeV/amu

existing ion source simplifies Phase I

new RFQ arrives this month

new cyclotron installed in 2027

Phase I proposal deploys REGINALD v1 to run the ion source, controlling the quality of the extracted beam
MIT+IBA Partnership → Reduced risk, fast delivery. Project has valuable outcome to the larger community.

Phase II proposal will deploy the next generation to operate the full system (source-through-cyclotron)

Note change from blurb submitted to LNS last week: our proposal no longer includes a new ion source.
We can isolate and safely use the existing MIST-2, in-situ at the IBA test stand for Phase I as needed.

Project Overview

Vision

Long-term Vision:	**Develop and deploy REGINALD on accelerators up to 1 MW** (not just cyclotrons) **operated outside of a traditional laboratory environment** (Underground Labs, Isotope Production Facilities, Prototypical Fast n Sources, etc)
Phase I Proposal Goal:	Develop REGINALD v1 to demonstrate semi-autonomous running of the MIST-2 ion source.
Phase II Proposal Goal:	Extend REGINALD to semi-autonomously control the 30 kW cyclotron prototype

REGINALD is as-yet unfunded. Since last autumn we have been preparing proposal material.

Our plan is an exact match to focus area!

Quoting 13.A:

→ We receive Isotope Program funding (IRP).
→ Small accelerator facilities match us.
→ We will deploy the system in Phase I
→ We are experts at surrogate modelling , aka "Digital Twins"

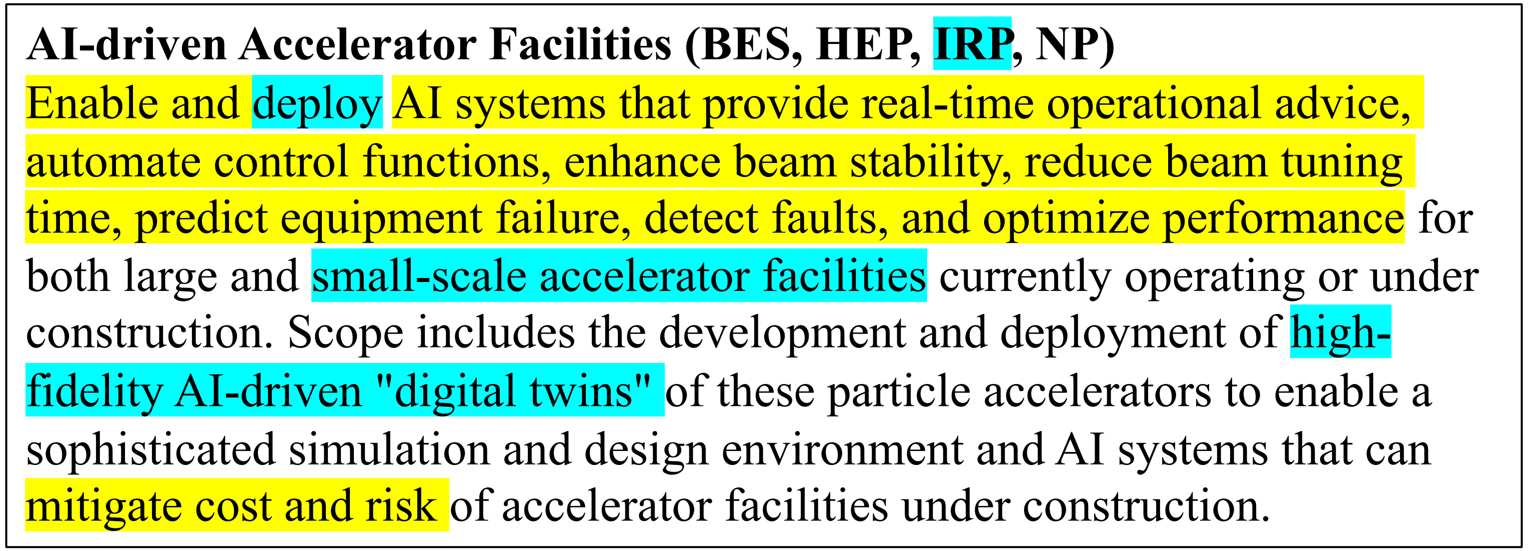

AI-driven Accelerator Facilities (BES, HEP, IRP, NP)

Enable and deploy AI systems that provide real-time operational advice, automate control functions, enhance beam stability, reduce beam tuning time, predict equipment failure, detect faults, and optimize performance for both large and small-scale accelerator facilities currently operating or under construction. Scope includes the development and deployment of high-fidelity AI-driven "digital twins" of these particle accelerators to enable a sophisticated simulation and design environment and AI systems that can mitigate cost and risk of accelerator facilities under construction.

Project Overview

“How will AI enhance our scientific and technical workflows?

Big Picture: The highest power (IsoDAR) cyclotron will deliver 60 MeV at 10 mA (600 kW) with an 80% duty factor, operatinng for 5 years in an underground environment. At this power, 200 W of beam loss drives significant activation, while larger losses risk rapid component failure; thus, machine protection and radiological safety are a priority. The system must balance sustain high availability for maintenance against beam quality to meet physics objectives.

Given the power density, complexity, and limited access, semi-autonomous operation—combining real-time diagnostics, predictive fault detection, and fast beam shutdown—is required to ensure safe, stable, long-term performance.

The “AI Edge” across the REGINALD project’s workflow

- Advantage 1: Speed. Expertise demonstrated in 2021 w/ μs-evaluation of accelerator parameters using surrogate modeling, leading to awards and intellectual leadership in Snowmass Computation for Accelerators
- Advantage 2: Accuracy. Our group’s digital twin development has been published in physics and machine learning journals alike; demonstrating <6% prediction error for key beam dynamics parameters allowing for rapid component optimization
- Advantage 3: Safety. Compared to human: Faster reaction times and better predictive power (based on small changes in several diagnostics) lead to safer operation of the system, provided safeguards against overcorrection are in place (e.g., interlocks).

It's the "Waymo edge": AI has consistently proved a safer operator than human drivers

9-months work plan (text from proposal draft in blue)

7/1-8/31 Surrogate Environment Preparation (Huchthausen, Franklin, Conrad)

Data will be collected from the **high-fidelity simulators COMSOL and IBSimu to build a surrogate model** of the plasma and beam dynamics.
The prepared digital twin will then be used to generate many operating states according to randomly sampled or prespecified heuristic policies.
Domain randomization will be injected into collected data to ensure training sample robustness.

8/31-10/31 Behavioral Grounding and Policy Pretraining (Villarreal, Winklehner)

The policy network of REGINALD will be pretrained leveraging techniques like behavioral cloning and advantage-weighted regression. Such techniques in imitation learning are important to establish safe policies and populate a manifold of reasonable and safe actions to jump-start later fine-tuning.

10/1-11/15 Offline Reinforcement Learning and Robustness Training (Villarreal, Winklehner)

REGINALD's core **reinforcement learning** infrastructure, a soft actor-critic (SAC) algorithm, will be implemented, and trained on the dataset generated from surrogate. More extreme domain randomization will be incorporated into the training dataset to enhance model robustness.
Critic learns conservative Q-function (conservative Q-learning, CQL)

11/1-11/15 Validation on High-Fidelity Beam Dynamics Simulator (Villarreal, Huchthausen, Winklehner, Forton)

REGINALD's trained CQL-SAC algorithm will be integrated with a high-fidelity beam dynamics simulator (IBSimu)
to evaluate control performance on a small number of states.

7/1-9/1,11/15-2/1 Human-in-Loop and Semi-Autonomous Operation (McClain, Winklehner, Maunoury, Forton)

REGINALD will be integrated with the EPICS control system of the ion source and run in advisory mode. A human operator will approve actions.
Resulting states are logged as a new dataset, and retraining will happen periodically. Human-only vs Semi-Autonomous Operation comparison.

12/1-3/31 Dissemination of Results and Planning for Phase II (Conrad as organizer, all involved)

Results of the REGINALD studies will be disseminated via two papers, code will be disseminated via GitHub, a workshop will be held for US manufacurers.
The Phase II proposal will be written. This will address any required diagnostic equipment upgrades and extend to operating a full 30 kW cyclotron system.

The steps in a nutshell...

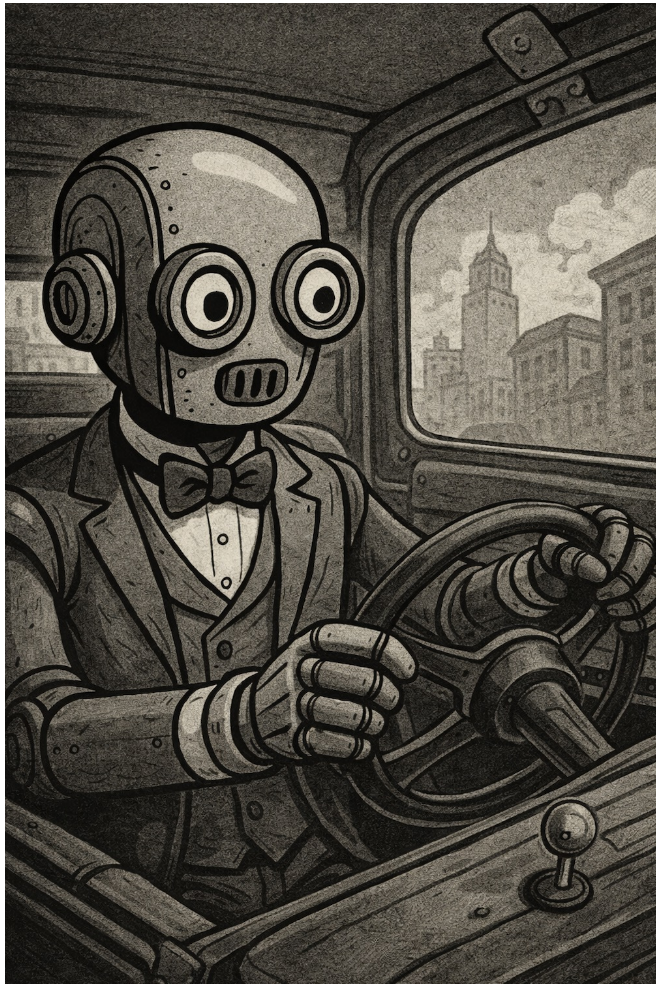

REGINALD is first trained on simulation data.

After success with that,
then REGINALD will be allowed to operate the ion source.

This will proceed in a two-step process:

Initially, we allow REGINALD to recommend actions to take with the approval of a human operator.

Once we've established confidence in REGINALD's ability, only then do we allow it to operate the ion source
with human over-ride (but not human guidance).
i.e. "Semi-autonomous operation"

It's the Waymo model of how to teach driving

9-months deliverables & Metrics (text from proposal draft in blue)

7/1-8/31 Surrogate Environment Preparation (Huchthausen, Franklin, Conrad)

- Deliverable: Neural network-based digital twin of high-fidelity beam dynamics simulators.

  Measure of success: high prediction accuracy (>95% R^2) on out-of-sample collected trajectories from IBSimu.
- Deliverable: Collected offline trajectory dataset covering stable operating region, respecting constraint boundaries, and few frequent failure modes

8/31-10/31 Behavioral Grounding and Policy Pretraining (Villarreal, Winklehner)

- Deliverable: Pre-trained policy network which suggests actions statistically consistent with data set collected in previous phase.

  Measures of success: Low action prediction error (NRMSE < 10%) on validation sample, qualitative action distribution consistency, qualitative conditional policy distribution consistency

10/1-11/15 Offline Reinforcement Learning and Robustness Training (Villarreal, Winklehner)

- Deliverable: Trained QCL-SAC algorithm and integration into online environment simulation.

  Measures of success: Reasonable training stability and convergence (>10% improvement in cumulative reward over heuristic controllers)

11/1-11/15 Validation on High-Fidelity Beam Dynamics Simulator (Villarreal, Huchthausen, Winklehner, Forton)

- Deliverable: Demonstration of QCL-SAC success at interacting with high-fidelity beam dynamics simulation.

  Measures of success: Failure-free operation, >10% improvement in cumulative reward over human controller to ID "AI edge"
- Deliverable: RL methods paper, targeting *Engineering Applications of Artificial Intelligence (IF: 8.0)*.

7/1-9/1,11/15-2/1 Human-in-Loop and Semi-Autonomous Operation (McClain, Winklehner, Maunoury, Forton)

- Deliverable: (for 7/1 period) MIST-2 installed with isolation equipment to allow operation; x-ray radiation rates measured at safe levels.
- Deliverable: (for 11/15 period) Integration of REGINALD with EPICS ion source control system.

  Measure of success: REGINALD operates the ion source for 12 continuous hours at 10 mA w/o human intervention.

12/1-3/31 Dissemination of Results and Planning for Phase II (Conrad as organizer, all involved)

- Deliverable: List of improvements to diagnostic equipment and to REGINALD assembled for Phase II.
- Deliverable: Results dissemination via paper *(e.g. in Phys. Rev. Research)* and workshop; code tool kit world-available.

Where will we be at the end of 6 months (end of December) for the go/no-go decision?

We will be 6 weeks into the Human-in-Loop and Semi-autonomous Operation

We will have first results from REGINALD running on the ion source.
We will also have a paper on our Reinforcement Learning methodology on the arXiv.

At least 7 deliverables will be checked off, and 4 metrics will be met.

We will also be able to report on progress on the full prototype construction.

→ This should be sufficient for the Phase II go/no-go decision.

The remaining 3 months will go into polishing and disseminating the final Phase I results, and developing the Phase II proposal.

Workforce Development Plan

What specific workforce development does the Phase I plan involve?

State-of-the-art computing experience for students:

- Development of state-of-the-art AI application for GS Villarreal in the Phys/Stats/Data Science program. This will be his thesis chapter related to instrumentation.
- Apprenticeship in AI techniques for GS Huchthausen (now a 1st year student) for lead role in Phase II
- Mastery of COMSOL and IBSimu for GS Huchhausen and UG Franklin

Training a student who, from day 1, has intended a career in industry:

- GS Huchthausen intends a career developing and using novel accelerators for medical isotope production.

Beyond-MIT Geographic representation and training beyond the field of physics:

- UG Franklin is from New Mexico and attends UNM. He is an engineering student

Unique contact with industry for the students and postdoc:

- Extended collaboration with IBA, in-person for PD McClain, by zoom for others
- Interactions/networking during final in-person workshop.

Publications with students as first authors:

- Engineering Appl. of A.I. and Phys Rev Res. papers planned, one thesis chapter

Our workforce development is great,
but…
Our biggest impact on US industry
is through technology transfer.
Because this is an open science project,
our impact is across the industry.

Conclusion: Assessment Q&A

1. Alignment of proposed work with focus area priorities? The alignment is exact. See talk.

2. Likelihood that the project produces a demonstrated advantage in workflows through AI incorporation within the initial 9 month timeframe (Phase I proposals)? Very High. This is low risk/high reward.

   We understand the steps well and have experience with surrogate modeling for accelerators.
   Group has delivered AI tool kits of similar scale this on a similar schedule several times.
   The ion source has been run, some data published, only installation at the teststand is required.
   At the end of the Phase I plan, we plan deployment for real world measurements.

3. An achievable Phase II transition plan? In place.

   The broad outline of the plan, involving application of REGINALD to the cyclotron, is already clear.
   Time is allocated at end of Phase I for collecton of detailed new information for the proposal.
   Two new collaborating institutions (one lab, one industry), with clear roles, have volunteered for Phase II.

4. Strength of the collaboration partnerships? Very strong.

   Built on a decade of collaboration, demonstrated with publications
   and a directly relevant, immediately on-going collaboration on the prototype that simplifies organization.

5. Workforce development? (recognizing the unique role academic institutions play in this space) Detailed in talk.

   More discussion of our credibility on this topic appears on a backup slide.
   But in terms of value to industry, we point also to our tech transfer, even at the end of Phase I

Back up slides

Comment on the new scoring sheet.

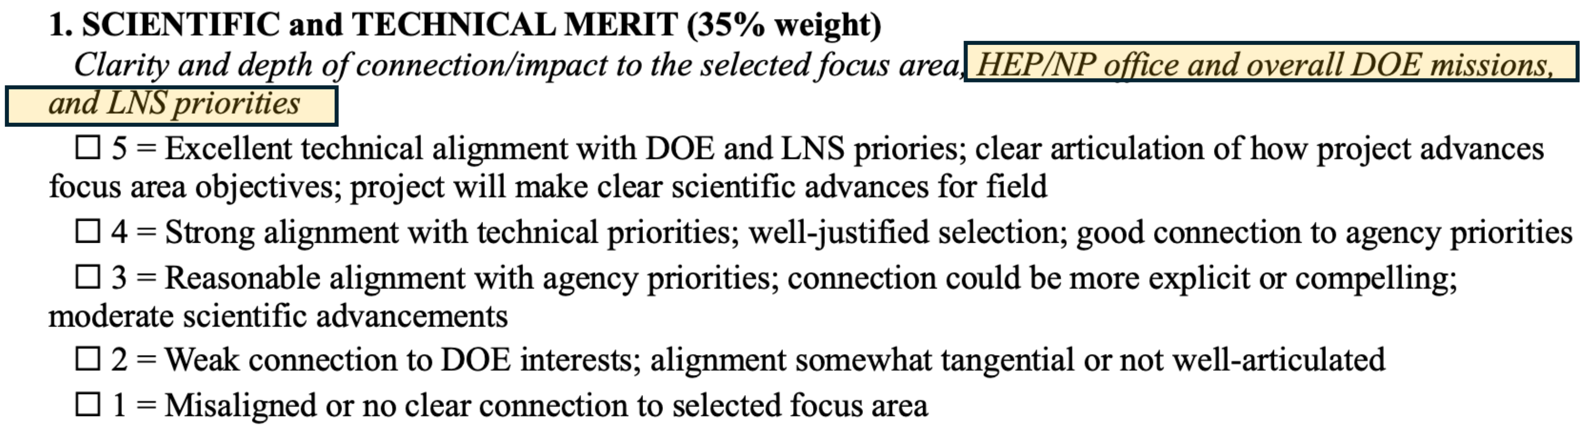

1. SCIENTIFIC and TECHNICAL MERIT (35% weight)

Clarity and depth of connection/impact to the selected focus area, HEP/NP office and overall DOE missions, and LNS priorities

☐ 5 = Excellent technical alignment with DOE and LNS priories; clear articulation of how project advances focus area objectives; project will make clear scientific advances for field

☐ 4 = Strong alignment with technical priorities; well-justified selection; good connection to agency priorities

☐ 3 = Reasonable alignment with agency priorities; connection could be more explicit or compelling; moderate scientific advancements

☐ 2 = Weak connection to DOE interests; alignment somewhat tangential or not well-articulated

☐ 1 = Misaligned or no clear connection to selected focus area

Highlighted text narrows the FOA for area 13 by adding the requirements HEP/NP and to "LNS priorities"

w/r/t HEP/NP ... this should really be **Participating Offices (BES, HEP, IRP, NP)**

This Phase I proposal aligns with the IRP mission/priority of moving to higher power very well

(As a secondary office, this proposal also aligns with ARDAP program that is now under HEP.)

w/r/t "LNS priorities" ... this competition is bigger than LNS

This is an **MIT-wide competition** for areas 13 and 14 that LNS is hosting,

It is important to cast a wide net for ideas that will succeed in getting funded.

(Similarly, for example, we would not want RLE asserting their priorities on areas 7,8 and 9.)

In summary: The FOA should be setting the definitions for item 1 on the score-sheet.

What are our strategies to make this proposal highly fundable,
beyond those listed on the 3/30 score-sheet?

1. We are aiming at IRP for our technical review office --- we have a strong relationship there.
2. IRP will have Phase I slots, but the competition will not be as fierce as in the larger offices
3. Within IRP, very few other proposals will focus on high power (a direction that they want to grow)
3. The IHE-industry partnership will be very well regarded in IRP.
4. We will deploy by the end of Phase I – we think most proposals will not deploy until Phase II.
5. REGINALD v1 running an ion source is a "product" that a wide community can actually use.
6. We are emphasizing technology transfer, which is a goal of GENESIS.
7. In Phase II there is well-reasoned growth –each partner brings something specific.

Workforce Development Plan

Will reviewers see us as already dedicated to accelerator community workforce development integrating AI?

Yes. We have a clear record of workforce development that aligns with GENISIS 13.A that makes our case credible.

We are well-regarded for training accelerator physicists and, equally importantly, particle physicists that understand the issues issues of accelerators that deliver their beams in depth.

We are one of the few groups at LNS that admit students whose goal from day 1 is a career in industry.

Example: Loyd Waites, Draper Fellow whose thesis included our surrogate modeling techniques.

We are well known at National Laboratories for bringing novel ideas – we are not a subordinate partner.

Our students and postdocs collaborate with national lab physicists from LANL, BNL, FNAL (and soon ANL).

Two past postdocs now working in accelerator divisions of US national labs (BNL and Cornell)

We are unique for our strong collaborations across the cyclotron industry (IBA, BCS, AIMA, etc).

Crucial to this: Our open science policy. We publish and we do not patent.

We have a record of (award-winning) leadership in integrating AI into accelerator design.

This is connects our decade-long commitment to AI integration into experimental particle physics.

We have multiple publications on AI-specific, but not experiment-specific techniques (see backup slides)

Why an IHE-Industry only Partnership in Phase I

Why not with a laboratory?

We think the University-Industry collaboration will make this proposal stand out in Phase I

- Today's accelerator program managers want true university-based accelerator physics development.
- They want the university PIs to be bringing new ideas to the labs.
- We have seen this in DOE-IRP (on 13.A list), DOE-FES, and DOE-ARDAP proposals.
- There will be ~10 Large facility and ~3 small facility National Labs partnering on proposals
  But very few accelerator manufacturers (IBA is only partnering with us).

There is a clear path to adding labs in for Phase II

- eg LANL proposes a fast beam chopper that will be integral to the safety mechanisms for autonomus running.
- Why does LANL love us? Because about 3 years ago realized that there is DOE money for small accelerators and they want to move into cyclotrons (where they will be relatively unique among labs).
- They need us and IBA to reach their goal.

If LANL were in Phase I, it would reduce the effectiveness of the "small facility" argument.

- The "small" category is going to have many fewer proposals than the "large" increasing our chances.

Our partnership with industry is well established,

- Our collaborations with IBA cover more than a decade. This is a very solid collaboration.
- We actually collaborate with ~5 different accelerator industries that we could potentially add in Phase II

Is it an issue that the industry partner is not in the budget in Phase I?

No. The FOA says:
"partners must provide intellectual contributions to the proposed project but do not need to be funded by DOE. "

For Phase I, the simplest thing to do is to build on our existing no-cost collaboration.
We need a well understood process that is implementable on the short timescale of Phase I.
IBA is happy with this plan.

For Phase II we are likely to have industry partners in the budget (both IBA and Draper).
We are also likely to have a lab partner in the budget (LANL)
But the longer timescale for Phase II allows time to work out a clear plan and budget.

More on Phase II Partners

DRAPER

Point of Contact: Loyd Waites

- Thesis included surrogate modeling for RFQ design for the prototype – he brings specific expertise
- Collaboration would build the Draper AI/ML portfolio being developed by his team.
  - → Surrogate Modelling and autonomous control are transferable to many other areas at Draper
- Draper wants to build up their Education portfolio
  - → Naturally leads to strong interest in partnership with a university on AI/ML
- Eager to sit down to discuss options and opportunities

Los Alamos National Laboratory

Points of Contact: Bruce Carlston, En-Chuan Huang, Alex Scheinker

- LANL intends to grow capability with cyclotrons. This has motivated a multi-year collaboration w/ us already.
- A specific long term goal is a cyclotron producing neutrons when the linac is not running.
- They can provide a fast chopper for Phase II – an instrument with both physics and safety uses.
- Surrogate Modelling and autonomous control are transferable to many other areas at LANL
- We have verbal agreement from Carlston that they will participate in Phase II – he was quite eager.

We will add partners if they address specific Phase II goals.

Our Publications on AI for Accelerators

★ **Neural networks as effective surrogate models of radio-frequency quadrupole particle accelerator simulations** J. Villarreal (MIT), D. Winklehner (MIT), D. Koser, J. M. Conrad (MIT)
e-Print: 2210.11451 [physics.comp-ph] Published in: Mach.Learn.Sci.Tech. 5 (2024) 2, 025009

★ **New directions for surrogate models and differentiable programming for High Energy Physics detector simulation** Andreas Adelmann (PSI, Villigen), et al. e-Print: 2203.08806 [hep-ph]

★ **Input Beam Matching and Beam Dynamics Design Optimizations of the IsoDAR RFQ Using Statistical and Machine Learning Techniques**
D. Koser, L. Waites, D. Winklehner, M. Frey, A. Adelmann et al. e-Print: 2112.02579 [physics.acc-ph]
Published in: Front.in Phys. 10 (2022), 875889

Order-of-magnitude beam current improvement in compact cyclotrons
D. Winklehner (MIT, Cambridge, CTP), et al., e-Print: 2103.09352 [physics.acc-ph]
Published in: New J.Phys. 24 (2022) 2, 023038

Modeling of advanced accelerator concepts
J.-L. Vay et al.
Published in: *JINST* **16** T10003 (2021)

Simulations of future particle accelerators: issues and mitigations
D. Sagan et al.
Published in: *JINST* **16** T10002 (2021)

Snowmass 2021 Computational Frontier CompF03 Topical Group Report: Machine Learning
Phiala Shanahan (MIT), et al., e-Print: 2209.07559 [physics.comp-ph]

Snowmass21 Accelerator Modeling Community White Paper
S. Biedron (New Mexico U.), et al., e-Print: 2203.08335 [physics.acc-ph]

★ = our papers explicitly on surrogate modeling for accelerators

These publications are directly relevant to the proposal and demonstrate our expertise.

Other Recent Publications of AI Tools from Our Larger Group

(we also have many publications from AI/ML-based analyses not listed here)

Machine Learning-Informed 3+1 Sterile Neutrino Global Fits using Posterior Density Estimation of Electron Disappearance Data J. Villarreal (MIT), J. Woodward (MIT), J. Hardin (MIT), J. Conrad (MIT)
e-Print: 2512.05784 [hep-ph]

A frequentist simulation-based inference treatment of sterile neutrino global fits
J. Villarreal (MIT), J. Woodward (MIT), J. M. Hardin (MIT), J. M. Conrad (MIT)
e-Print: 2507.01153 [hep-ph] Published in: Mach.Learn.Sci.Tech. 6 (2025) 3, 035053

NuBench: An Open Benchmark for Deep Learning-Based Event Reconstruction in Neutrino Telescopes
Rasmus Ørsøe (Munich, Tech. U.), et al. e-Print: 2511.13111 [hep-ex]

Demonstrates range of experience with producing AI toolkits.

Reference for MIST-2 Ion Source

The MIST-1 and MIST-2 multicusp ion sources for high-current H2+ beams D. Winklehner, et al.,
ePrint: 2507.03155 [physics.acc-ph] Published in: J.Phys.G 52 (2025) 11, 115001

Establishes that we have a working source