

Inverse modelling of laser-plasma acceleration

From particle-in-cell simulations to fast models, and from the beam back to the machine

Mykyta Lenivenko · IJCLab, Université Paris-Saclay, CNRS/IN2P3

DMLab@Paris · 22.10.2026



Reading an accelerator backwards

An operator sets the **knobs**. The accelerator gives an **electron beam**. We measure the beam on every shot.

The real conditions in the plasma are **not** exactly what the knobs say. The laser drifts. The gas pressure fluctuates.

Inverse model: from a measured beam, find the settings that made it. A virtual diagnostic of the machine, shot by shot.

Diagnose

Which knob drifted today?

Check

Is the shot record what the plasma saw?

Control

Correct the knobs to get the beam back.



Today

- 1 · Laser-plasma acceleration
- 2 · Simulation, and its price
- 3 · Surrogate models
- 4 · Forward models: PALLAS and KIT
- 5 · Why going backwards is hard
- 6 · Inverse model (MLST paper)
- 7 · Inverse models at the AISSAI hackathon

01

Laser-plasma acceleration

An accelerator a few millimetres long



1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models


4 · Forward models: PALLAS and KIT

5 · The inverse problem

6 · Inverse model

7 · Inverse models at the hackathon

Riding a plasma wave



3 · Pure helium: no new electrons. The bunch rides the back of the bubble ($E_z < 0$) and gains energy. Being faster than the wave, it slowly **catches up with the laser**: past the middle of the bubble it would start to lose energy (dephasing).

pause

restart

Two laser-plasma injectors

PALLAS · IJCLab, Orsay

cSTART · KIT, Karlsruhe

What	Laser-plasma injector test facility (EuPRAXIA context)	Laser-plasma injector (LPI) for the cSTART storage ring
Goal	A beam with the quality and stability of a conventional injector: about 200 MeV , 30 pC, energy spread below 5 %	A bunch the ring can accept: at most 90 MeV , short and stable
Laser	LaserIX: 3 J (1.5 J on target, <i>in our dreams</i>), 45 fs, now 1 Hz	In commissioning: expected 1.75 J on target, 35 fs
Target	Gas cell: higher density , strong self-focusing (J. Serhal talk)	Gas jet: lower density , weak self-focusing
Injection	Ionisation injection: helium + nitrogen, then pure helium	Ionisation injection: hydrogen with a nitrogen-doped region
My part	Simulation scans, surrogate and inverse models, closed-loop optimisation	Simulation scans and surrogate models, during my stay at KIT

Same physics, two regimes. The same machine-learning tools serve both.

02

Simulation, and its price

Particle-in-cell codes: exact, but expensive

1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models


4 · Forward models: PALLAS and KIT

5 · The inverse problem


6 · Inverse model

7 · Inverse models at the hackathon

Particle-in-cell simulation



- The laser and the plasma, solved together, particle by particle.
- Codes: Smilei, FBPIC.
- Every run needs **huge computing resources**: about an hour on hundreds of CPU cores.



The bubble. An empty cavity behind the laser. Below: the field on the axis. Yellow zone: the field that accelerates electrons.

One PALLAS-type Smilei simulation (helium + nitrogen, laser envelope, cylindrical). Window 70 μm long, moving with the laser (to the right). Blue-white: plasma electrons. Red glow: the laser. Yellow: injected electrons. Bottom: the field on the axis.

The price of one answer

300 – 12,000

core-hours for one simulation: laser envelope model (384 cores, ~45 min) or full laser wave (1,536 cores, ~8 h)

~3 million

core-hours used on Adastra by our scans in 2026

thousands

of simulations behind every question below

Tune the machine

To find good settings, you try many settings and keep the best one. Each try is one simulation.

Read a beam backwards

To explain one measured beam, you compare it with the beams of many possible settings. Again thousands of simulations, for one shot.

Follow the machine

The laser and the gas drift during a day. The answer is needed in seconds. One simulation takes an hour on hundreds of cores.

simulations needed



1,000

envelope

full wave

with PIC (one envelope run $\approx$ 300 core-hours)

300,000 core-hours

= 33 days on 384 cores, one run after the other

with a surrogate (one answer $\approx$ 1 ms)

1 second

on one laptop core

We cannot run a simulation every time we have a question. We need a fast stand-in.

03

Surrogate models

Learn the simulator once, ask it a million times

1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models

4 · Forward models: PALLAS and KIT

5 · The inverse problem

6 · Inverse model

7 · Inverse models at the hackathon

Learn the simulator once

Design a scan

choose thousands of settings



Run PIC

once, on a supercomputer



Extract the beam

numbers, curves, spectra



Train a network

settings → beam



Ask it anything

about 1 ms per answer

hours → ms

per answer, on one CPU core

forward

settings → beam: scan, optimise, explain

inverse

beam → settings: diagnose, control

PALLAS · old random scan

Laser and gas knobs: pressure, laser amplitude, nitrogen, focus. About **13,000** simulations, the beam at the end.

PALLAS · analytical-profile scan

Laser fixed, the shape of the gas profile varies. About **6,000** simulations, the beam at 113 positions.

KIT · cSTART scan

Six plasma-profile knobs. About **1,500** simulations, the beam at every time step.


A surrogate is only as good as its data. Before training, look at what the data can tell you.

What is in the data

PALLAS · old random scan (laser + gas)

PALLAS · analytical-profile scan

KIT · plasma-profile scan



What does laser amplitude change in the beam?

energy
energy spread
charge
peak current
emittance
bunch length
beam size
divergence



laser amplitude moves no beam number much: it will be **hard to read back** from the beam.

$\rho = +1$: rise together · $\rho = -1$: one rises, the other falls · $\rho \approx 0$: no visible link. PALLAS · old random scan (laser + gas), 10,751 simulations.

Matrix: rank correlation ρ between every pair of numbers. **Amber: knobs**, **cyan: beam**. Pink: rise together. Blue: one rises, the other falls. Click a cell: the right plot shows what that knob changes in the beam.

0 4

Forward models: PALLAS and KIT

Settings in, the beam along the whole accelerator out



1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models

4 · **Forward models: PALLAS and KIT**

5 · The inverse problem

6 · Inverse model

7 · Inverse models at the hackathon

The beam along the accelerator

Trajectory: the 6 beam numbers at 113 positions. Inside the plasma: acceleration. After it: free drift.



Model: FiLM decoder. 5 networks; their spread gives the error bar σ .

Test runs: hackathon hidden test, 357 simulations never used in training.

Divergence: post-processed incorrectly in the data this model was trained on (fixed since). A likely cause of its poor error bar.

a typical beam

- a low-charge beam
- a high-charge beam
- a high-energy beam



Five architectures, one answer

MLP




Plain network: one row per (settings, z).

Bayesian NN



Bayesian by dropout: the spread of 100 random passes is the error bar.

DLinear




Time-series model: writes the whole curve from the settings.

Residual MDN



A network, plus a mixture that learns how wrong it is.

FiLM decoder



The settings steer a network that reads the position: any z, smooth curves.

What we learned

- Very different designs: all within **0.01** in R^2 .
- Inside the training range, a plain network is enough.
- What is left: the hard outputs, and **outside the range**.


Beyond the training density

DeepONet

FiLM ensemble

Bayesian NN

● inside the training range ◆ above the density cut (never seen)



Train only on simulations **below** a gas density. Test on simulations **above** it, which the model never saw.

0.987

inside the range

0.969

above the cut

0.018

lost outside (the gap)


- Energy and charge survive the jump.
- **Bunch length** suffers most, for every model.
- Models with structure (ensembles, FiLM, Bayesian) lose less than a plain MLP.

cSTART at KIT

- Laser-plasma injector for the **cSTART** storage ring: at most **90 MeV**.
- Same idea as PALLAS, but a gas jet at lower density: **much less self-focusing**.
- About **1,350** Smilei runs that inject a bunch. The beam is stored at **every** time step: the whole trajectory, not only the end.
- Trajectory model and dashboard delivered to the cSTART team.

Work done during my stay at KIT. **Thank you, DMLab, for supporting this visit.**

How accurate along the run (held-out simulations)



The first millimetre after injection is hard: a few particles set every number there. Further on, the model is as good as at the end.

The KIT trajectory model, live

laser focus

H plateau length

H down-ramp

H density

N start

N density

1.25 mm

0.55 mm

3.50 mm

×3.2

0.93 mm

×3.2

Load a Smilei run (◆ = its simulated beam at the end):

run 1

run 2


run 3

run 4

run 5

run 6

The real model runs in this page: 5 networks, one answer per position, in milliseconds. Grey zone: the first 1.5 mm, less reliable.



Asking for a beam

The forward model is fast enough to **search**.

1. Name the beam you want: energy, spread, charge, bunch length.
2. Scan about **235,000** settings in about **2 s**.
3. Keep every setting whose beam matches within the model's own error.

What the search tells you

- The best working points, and the beam each one gives.
- Which targets **cannot** be reached.
- Which knobs are **pinned** by the target, and which stay **free**.

Ask for a 75 MeV, low-spread bunch: the hydrogen density is pinned to a few % of its range. The plateau length can roam over a third of it. **The beam does not see that knob.** That is the inverse problem, already.

05

The inverse problem

From the beam back to the machine

1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models


4 · Forward models: PALLAS and KIT

5 · The inverse problem

6 · Inverse model

7 · Inverse models at the hackathon

Why going backwards is hard



Not unique

Different settings can give the **same beam**. A stronger laser in a thinner gas makes the same beam as a weaker laser in a denser gas. From the beam alone we cannot tell A from B.

Not stable



Every diagnostic has an error bar. Inside it, the beam is compatible with settings that are **far apart**. A small change in the measured beam can move the answer a lot.

Not complete

Some settings change the beam **less than the measurement error**. The measured beam carries almost no information about them.

The inverse problem is ill-posed. One number is not an honest answer: we need the distribution of the possible settings.

Same beam, different machine



beams equal within  **x2**
× the paper's measurement error (3 % on energy, 10 % on the others), same focus.

7 other simulations give the same beam. Laser amplitude **1.23–1.85**, pressure **16–37 mbar**, nitrogen **10.6–12.0 %**.

Every point is one PALLAS simulation. Click a ringed one: the amber points are the other simulations whose 8 beam numbers all match it. Right: how far their beams are from it (dots), against the error band.

06

Inverse model

Mykyta Lenivenko *et al.*, Machine Learning: Science and Technology (in review)

1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models


4 · Forward models: PALLAS and KIT

5 · The inverse problem


6 · Inverse model

7 · Inverse models at the hackathon

Predict a distribution, not a number




- **In:** 8 beam numbers + the focus set point.
- Trained **with the measurement noise** of each diagnostic.
- **Out:** a mixture of 5 Gaussians over pressure, laser amplitude and nitrogen.
- An ensemble of 5 networks.



One test shot. Shades: the predicted distribution. The star is the truth.

Two candidates, with weights 0.68 and 0.32. The truth sits near the first one, but the second one is a real possibility.

Does it work?




- **Pressure** and **nitrogen**: recovered well.
- **Laser amplitude**: the hardest. It hides in the valley.
- **The error bars are honest**: the predicted intervals contain the truth as often as they claim (panel d, close to the diagonal).

Good on average. But the average of two valid answers can be no answer at all.

Two answers on most shots

where


how different



Fraction of shots with a strong second answer, over the scan.

one target beam

which combination



Mismatch to one target beam over (pressure, laser). Dots: simulations that give the same beam.

The two answers differ in every knob, but keep the **self-focusing strength $a_0^2 \cdot p_1$** almost fixed. The beam measures this combination, not the knobs one by one.



Breaking the degeneracy

do both answers fit?

picking one

the idea

tested on PIC



Ask a forward model which answer fits. For most shots **both** answers reproduce the measured beam. The beam alone cannot decide.

Measure again at another gas density. Each shot has its own valley. The valleys cross at one point: the true setting.

The limit is the diagnostics, not the network. A second, deliberately changed shot breaks the degeneracy.

07

Inverse models at the hackathon

AISSAI Accelerator Hackathon, GANIL, October 2026 · team Plasma Pirates

1 · Laser-plasma acceleration

2 · Simulation, and its price

3 · Surrogate models

4 · Forward models: PALLAS and KIT


5 · The inverse problem

6 · Inverse model

7 · Inverse models at the hackathon


One week, seven tasks

AISSAI Accelerator Hackathon, GANIL, 12–16 October 2026. About 30 people, 3 challenges. I was **challenge holder** of *Surrogate models of PALLAS*, and a member of team **Plasma Pirates**.



Inverse B

8 beam numbers + focus → peak pressure, nitrogen, injector length, density dip




The linear model was only the bar to beat. Our internal mixture model (MDN) was already better on pressure, nitrogen and injector length, but it could not get the density dip. The team model recovers most of it (next slide).

Clean beams: what is left without noise


Old random scan: beam + focus → pressure, laser, nitrogen

◆ clean beam (hackathon MLP) ● with measurement noise (paper) ■ linear model



Analytical-profile scan: beam + focus → gas-profile knobs

■ linear model ■ MLP ■ ensemble + mixture



- Exact beam numbers: a plain network **almost solves it**, laser amplitude included.
- With realistic noise the laser amplitude falls: the noise pushes the answer along the valley.
- The **density dip** is almost invisible to a linear model.
- An ensemble with a mixture output finds most of it.

Without noise, one answer is often enough. With real diagnostics, we need the distribution, and a strategy to break ties.

Conclusion

Forward: an easy target

A direct model learns the simulator well, even the beam along the whole accelerator. It answers in milliseconds instead of hours.

Data first

With good data, even a plain MLP does most of the job. It fails outside the training data. Models with more structure still predict there.

Inverse: ill-posed

Different settings give the same beam: a degeneracy. Extra diagnostics or a forward model can break it. For the experiment it is also a plus: several working points give the same beam, so we can pick the most stable one.

Predict a distribution

One number hides the second answer. A calibrated distribution shows every setting compatible with the beam, and how sure we are.

Next: an ML-accelerated pipeline for PIC codes. Learned models replace the most expensive parts of the simulation, from the laser to the bunch.

Thank you

DMLab, for supporting my visit to KIT.

KIT / cSTART colleagues, for the collaboration.

Plasma Pirates (AISSAI hackathon): H. Guler, A. Hill, S. Joly, K. H. Kim, G. Rodríguez Llorente.

Co-authors of the MLST paper, and K. Cassou.

GENCI: computing time on Adastra (CINES) and Irene (TGCC).

mykyta.lenivenko@ijclab.in2p3.fr

- M. Lenivenko *et al.*, *Inverse modelling of a laser-plasma accelerator*, MLST (in review)
- G. Kané *et al.*, *Mach. Learn.: Sci. Technol.* **7**, 030502 (2026)
- J. Derouillat *et al.*, *Comput. Phys. Commun.* **222**, 351 (2018)
- E. Esarey, C. Schroeder, W. Leemans, *Rev. Mod. Phys.* **81**, 1229 (2009)
- A. Döpp *et al.*, *High Power Laser Sci. Eng.* **11**, e55 (2023)

Questions?