
Inverse Modeling of Laser Pulse Shapes in Inertial Confinement Fusion with Auto-Regressive Models

Vineet Gundecha^{1*}, Ricardo Luna Gutierrez^{1*}, Rahman Ejaz^{2*},
Varchas Gopalaswamy^{2*}, Riccardo Betti^{2*}, Sahand Ghorbanpour^{1*}, Aarne Lees^{1*},
Soumyendu Sarkar^{1*†}

¹Hewlett Packard Enterprise ²University of Rochester
{vineet.gundecha, rluna, sahand.ghorbanpour, soumyendu.sarkar}@hpe.com
{reja, vgopalas, betti, alee}@lle.rochester.edu

Abstract

Realizing practical fusion energy remains one of society’s most significant unresolved scientific challenges, carrying profound implications for sustainable, carbon-free power. A key determinant of success in Inertial Confinement Fusion (ICF) experiments is the design of a Laser Pulse (LP) Shape capable of optimally driving implosions within strict physical limits. Conventional LP design depends on costly simulations and labor-intensive iterative tuning. To address this, we introduce the Laser Pulse Shape Design System (LPDS), a generative inverse modeling framework based on auto-regression that directly maps desired fusion outcomes and target pellet parameters to optimized LPs. We explore a multi-objective training setup to design diverse LPs that adhere to physical constraint while achieving less than 2% error in the desired implosion outcomes. In addition, we incorporate constraint-conditioning via inpainting and gradient-based editing strategies, enabling precise control over pulse characteristics during generation. This framework offers a data-driven solution for LP design in ICF, advancing the pursuit of practical, sustainable fusion energy.

1 Introduction

Inertial Confinement Fusion (ICF) offers the promise of virtually limitless clean energy. It achieves nuclear fusion with high-energy lasers that compress and heat fuel pellets to extreme temperatures and pressures, overcoming electrostatic repulsion between nuclei [1]. Achieving ignition requires a precise, symmetric implosion regulated by the laser pulse (LP) shape—the temporal energy profile controlling compression over 3 nanoseconds. The LP is structured in stages: a low-intensity foot launching timed shocks, a rising ramp managing pressure buildup, a high-intensity main pulse driving rapid implosion, and an optional tail sustaining pressure (Figure 1). Designing effective LPs is difficult due to the complex interplay between laser delivery and implosion physics, typically demanding costly simulations and weeks of manual tuning or Bayesian optimization methods [5, 6, 2, 10]. To address this bottleneck, we introduce the Laser Pulse Shape Design System (LPDS), a machine learning-based inverse modeling




Figure 1: The LP is typically characterized by 12 parameters, some of which are annotated here.

Designing effective LPs is difficult due to the complex interplay between laser delivery and implosion physics, typically demanding costly simulations and weeks of manual tuning or Bayesian optimization methods [5, 6, 2, 10]. To address this bottleneck, we introduce the Laser Pulse Shape Design System (LPDS), a machine learning-based inverse modeling

*These authors contributed equally.

†Corresponding author.

framework that generates LPs from target implosion outcomes. LPDS combines auto-regressive generative models with auxiliary objectives for outcome alignment, a physics-informed loss for feasibility, and conditioning mechanisms for fine-grained control, enabling data-driven LP design.

2 Laser Pulse Shape Design System

We present the Laser Pulse Shape Design System (LPDS) - an inverse modeling framework to generate LPs in ICF. LPDS takes desired implosion outcomes and target pellet parameters as input, aiming to generate LPs that achieve specified fusion objectives while honoring physical constraints.

2.1 Inverse Modeling

Let $\mathbf{m} = \text{LILAC}(\mathbf{l}, \mathbf{p})$ denote the vector of implosion outcomes produced by the LILAC [3, 4, 7] ICF simulator for a given LP ($\mathbf{l}$) and target pellet parameters ($\mathbf{p}$). These outcomes include quantities such as energy yield, areal density, burn width, and ion temperature. The pellet parameters are defined by attributes such as outer radius, ice thickness, ablator thickness, ice-tritium fraction, etc. To construct an inverse design model, we first generate a dataset $D_F = \{(\mathbf{l}_i, \mathbf{p}_i, \mathbf{m}_i)\}_{i=1}^I$ of 1 Million examples via systematic simulation sweeps, capturing the relationship between pulse shapes, target parameters, and resulting implosion outcomes. The LP $\mathbf{l}_i$ is characterized as a real-valued sequence of length 256. $\mathbf{p}_i, \mathbf{m}_i$ are real-valued vectors of size 5 and 12, respectively. Our goal is to learn a data-driven inverse mapping G_θ that designs a feasible LP $\mathbf{l}'$ given desired implosion outcomes $\mathbf{m}$ and pellet parameters $\mathbf{p}$:

$$\mathbf{l}' = G_\theta(\mathbf{m}, \mathbf{p}). \quad (1)$$

This inverse design problem has multiple plausible solutions, since multiple pulse shapes can lead to similar outcomes. To address this, we consider two modeling approaches - generative and predictive. Auto-regressive generative models are trained to generate from a distribution of valid LPs, thereby providing diversity and robustness, and predictive auto-regressive models to predict a single high-fidelity LP.

2.2 Auto-regressive Model

LP design can be formulated auto-regressively, generating the pulse vector $\mathbf{l} = \{l_1, l_2, \dots, l_T\}$ sequentially. For the first timestep, since no past l_t values are available, the model instead conditions on an initialization vector constructed by applying linear transformations to both the target parameters $\mathbf{p}$ and the desired outcomes $\mathbf{m}$ (see Figure 3). This embedding serves as the initial conditioning input for the auto-regressive model. We explore two paradigms of modeling for the auto-regressive models - predictive and generative. The predictive models try to reconstruct the ground-truth LP with high fidelity. The generative models trade-off between reconstruction fidelity and the ability to provide diverse LP candidates. The training objective for the predictive auto-regressive models is the minimization of the mean squared error between the predicted and true pulse values across all time-steps. The generative models output a probability distribution over the LP value at time t called $p_\theta(l_t)$. The training objective for the generative auto-regressive models is to minimize the negative log-likelihood of the true pulse under the predicted distribution $p_\theta(l_t)$: $\mathcal{L}(\theta) = \frac{1}{T} \sum_{t=1}^T \log p_\theta(l_t | l_{1..t}, \mathbf{m}, \mathbf{p})$.

We explore 3 different modeling techniques for p_θ . The simplest technique is to model this as a Gaussian distribution where the model G_θ learns to predict the mean and variance. To offer more expressiveness, we also model p_θ as a weighted mixture of K Gaussians. Here, the model learns to predict K -length vectors for mean, variance and the mixing weights. We also try a discrete approach where the pulse l is quantized to 128 discrete values and p_θ is modeled as a Categorical distribution and trained with a cross-entropy loss.




Figure 2: LPDS supports: (a) generation of diverse pulse shapes achieving desired fusion outcomes, (b) constraint-based editing of specific pulse features, and c) Inpainting specific regions of the LP.

2.3 Auxiliary Objectives and Physics-Informed Loss

While the reconstruction objective enables LP reproduction, our broader goal is scientific exploration through plausible, diverse LP generation consistent with implosion outcomes. This provides a wider array of selectable, high-performing candidates and facilitates scientific discovery by potentially identifying novel LP configurations. To achieve this, we introduce an auxiliary objective that guides the inverse model to prioritize implosion fidelity and robustness over exact replication.

We train a surrogate model³ $\mathbf{S}_\phi$ to approximate the LILAC simulator. This model learns the mapping from a LP $\mathbf{l}$ and target parameters $\mathbf{p}$ to the corresponding implosion outcomes $\mathbf{m}$. This can be represented as $\hat{\mathbf{m}} = \mathbf{S}_\phi(\mathbf{l}, \mathbf{p})$, where ϕ represents the learnable parameters of the surrogate model. The surrogate is trained on the simulation dataset D_F by minimizing the prediction error: $\text{LILAC}_{\text{surrogate}}(\phi) = \mathbb{E}_{(\mathbf{l}, \mathbf{m}, \mathbf{p}) \sim D_F} [\|\mathbf{m} - \mathbf{S}_\phi(\mathbf{l}, \mathbf{p})\|^2]$. After training, $\mathbf{S}_\phi$ is kept frozen and is used during training of the inverse model to evaluate the physical fidelity of generated LPs. Given a generated pulse $\mathbf{l}' = G_\theta(\mathbf{m}, \mathbf{p})$, the surrogate model predicts the corresponding outcomes $\hat{\mathbf{m}} = \mathbf{S}_\phi(\mathbf{l}', \mathbf{p})$, which are then compared to the target outcomes $\mathbf{m}$. Our new objective can be defined as:

$$\mathcal{G}(\theta) = \lambda_L \mathcal{L}(\theta) + \mathcal{S}(\theta) \quad (2)$$

$$\mathcal{S}(\theta) = \mathbb{E}_{\mathbf{l}' \sim G_\theta, (\mathbf{m}, \mathbf{p}) \sim D_F} [\|\mathbf{m} - \mathbf{S}_\phi(\mathbf{l}', \mathbf{p})\|^2] \quad (3)$$

Here, λ_L controls the trade-off between faithfully reconstructing the original LP and generating physically consistent alternatives. A properly tuned λ_L allows the model to generalize beyond pure reconstruction, producing diverse yet plausible LPs tailored to the specified design goals. In practice, we first train only on the reconstruction loss $\mathcal{L}(\theta)$ since at the beginning of training G_θ can produce LPs that are incoherent and out-of-distribution for $\mathbf{S}_\phi$.

Physics-informed loss: While our goal is to generate a diverse set of LPs for a given target configuration, the practical utility of these designs depends critically on their adherence to the fundamental physics of ICF and their feasibility under real-world experimental constraints. One essential constraint is energy conservation. Specifically, generated LPs should not exceed the total energy budget of corresponding reference pulses. To enforce this, we introduce a physics-based penalty term:

$\mathcal{P}(\theta) = \mathbb{E}_{\mathbf{l}' \sim G_\theta, \mathbf{l} \sim D_F} \left[\left(\int_0^T \mathbf{l}'(t) dt - \int_0^T \mathbf{l}(t) dt \right)^+ \right]$. This term penalizes excess energy, encouraging designs within experimental and physical limits. Incorporating this constraint into our overall loss function, we define the complete training objective as: $\mathcal{G}(\theta) = \lambda_L \mathcal{L}(\theta) + \lambda_S \mathcal{S}(\theta) + \lambda_P \mathcal{P}(\theta)$, where λ_P and λ_S are hyper-parameters that control the relative importance of each loss term in the overall objective.

3 Experiments

To evaluate our approach, we compare LPDS variants: a predictive auto-regressive Transformer [11] ($\text{LPDS}_{\text{Transformer}}$), and an LSTM-based [8] model ($\text{LPDS}_{\text{LSTM}}$). The generative auto-regressive models (as described in Section 2.2) are based on the LSTM architecture and are named $\text{LPDS}_{\text{GaussianAR}}$, $\text{LPDS}_{\text{MixtureOfGaussianAR}}$ and $\text{LPDS}_{\text{DiscreteAR}}$. Additionally, we evaluate ablated versions excluding the auxiliary loss $\mathcal{S}(\theta)$ (*w/o S*) and the physics-informed loss (*w/o P*).

We assess performance on four key metrics: implosion outcomes $\mathbf{m}$ error, reconstruction error, generation diversity, and energy conservation error. All models are evaluated over $R = 10$ random seeds. The implosion outcomes $\mathbf{m}$ error is computed via the surrogate loss (Equation 3) and reported as mean absolute percentage error. Reconstruction error is the L2 distance between the generated and ground-truth pulse. Energy conservation is calculated in terms of percentage as:

³The surrogate achieves an error of 1.4% across the implosion outcomes $\mathbf{m}$. This can be considered as a lower bound on the implosion error that can be achieved by LPDS.

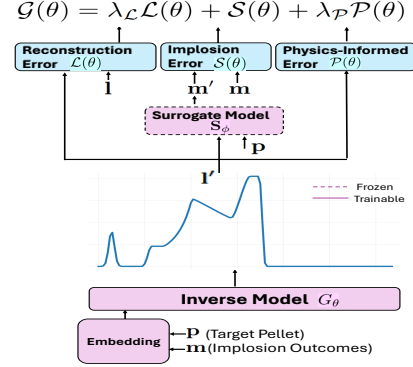


Figure 3: Training setup for LPDS. The loss term is comprised of multiple terms described in Section 2.3

$\frac{1}{E} \sum_{e=1}^E \frac{1}{R} \sum_{r=1}^R \frac{(\int_0^J l'_{e,r}(j) dj - \int_0^J l_e(j) dj)}{\int_0^J l_e(j) dj} \times 100$. Diversity is measured by average pairwise L2 distance across generated samples. The estimated upper bound for diversity is 1.9, which is the diversity value obtained when comparing randomly selected LPs from $D_{F_{\text{test}}}$.

The evaluation results in Table 1 highlight the effectiveness of LPDS. The auto-regressive model **LPDS_{LSTM}**, achieves the highest reconstruction fidelity, accurately replicating original LPs. Among the

Approach	Diversity $\uparrow$	m Error $\downarrow$	Reconstruction Error $\downarrow$	Energy Conservation $\downarrow$
LPDS_{LSTM}	—	1.65%	0.0001	0.66%
LPDS_{Transformer}	—	1.94%	0.0008	0.95%
LPDS_{GaussianAR}	0.42	$1.89\% \pm 0.01$	$0.0005 \pm 2e^{-5}$	$0.58\% \pm 0.004$
LPDS_{MixtureOfGaussianAR}	0.56	$1.95\% \pm 0.09$	$0.0006 \pm 8e^{-5}$	$1.58\% \pm 0.006$
LPDS_{CategoricalAR}	0.39	$2.01\% \pm 0.04$	$0.0009 \pm 5e^{-5}$	$1.18\% \pm 0.08$
LPDS_{LSTM, w/o S}	—	3.9%	0.0004	0.69%
LPDS_{Transformer, w/o S}	—	4.4%	0.001	1.23%
LPDS_{LSTM, w/o P}	—	1.85%	0.0001	0.95%
LPDS_{Transformer, w/o P}	—	2.1%	0.0009	1.3%

Table 1: LPDS model performance. $\pm$ denotes standard deviation over seeds. For the predictive auto-regressive models (**LPDS_{LSTM}**, **LPDS_{Transformer}**), diversity is not defined since they are deterministic.

generative variants, **LPDS_{MixtureOfGaussians}** offers the most diversity across samples due to better expressiveness of the multi-Gaussian distribution. Both techniques achieve less than 2% error in the desired implosion outcomes (**m-Error**). Refer Appendix Section B for example pulse shapes generated by the models.

4 Pulse Shape Constrained Design

In many ICF experimental design scenarios, scientists require control over specific regions or attributes of the LP. A differentiable mapping function $\mathcal{M}$ (details in Appendix) is used to project an LP $\mathbf{l}$ onto $M = 12$ parameters $C = \mathcal{M}(\mathbf{l}) = \{c^1, \dots, c^{12}\}$. Each c^m corresponds to a physically interpretable property of the LP (Figure 1). When constructing new designs, scientists fix one or two of these values within C , while allowing the remaining parameters, and thus the shape of the pulse, to vary. However, our current inverse model G_θ lacks a mechanism to enforce such partial constraints during inference. We use two different methods to add controllability -

Gradient-Based Adaptation: The LP can be adapted at post-design without additional finetuning. To honor constraints in $\mathcal{C}$ while maintaining desired implosion outcomes, a loss combining the pulse constraining loss $\mathcal{J}(\theta) = \mathbb{E}_{l', c^m} [\|c_{l'}^m - c^m\|^2]$ and the LILAC surrogate loss $\mathcal{S}(\theta)$ (Equation 3) is formulated. Since both the losses are differentiable w.r.t the pulse l' (now referred to as $\mathcal{J}(l')$, $\mathcal{S}(l')$), we minimize the following loss by adapting the LP with gradient descent - $\mathcal{T}(l') = \mathcal{S}(l') + \lambda_{\mathcal{J}} \mathcal{J}(l')$. We focus our evaluation on two physically meaningful parameters: picket power and foot power (see Figure 1). To quantify adaptation accuracy, we get an mean absolute percentage error (MAPE) of 2.5% between the specified constraint value and the corresponding value extracted from the generated pulse l' (**m - Error** remains unaffected).

Inpainting/Prompting: For additional controllability, we support inpainting-based generation (Figure 2c), enabling scientists to specify desired LP design directly in the LP space. Rather than conditioning on the $\mathcal{C}$ constraining parameters, users can provide a partial LP such as prefixes or fixed regions. For the auto-regressive models, providing a LP prefix is akin to prompting it to complete the rest of the pulse. A prefix comprising 10% of the pulse (this corresponds to the first peak) is provided as the prompt. We get a reconstruction error of 0.0003 and **m - Error** of 1.55%.

5 Conclusion

We introduced LPDS, a data-driven inverse modeling framework for generating diverse ICF laser pulse shapes satisfying scientist-defined objectives and physical constraints. LPDS has the potential to significantly advance ICF research, bringing us closer to realizing virtually limitless clean energy, a breakthrough with profound environmental and societal implications.

References

- [1] R. Betti and O. A. Hurricane. Inertial-confinement fusion with lasers. *Nature Physics*, 12(5): 435–448, 2016. doi: 10.1038/nphys3736. URL <https://doi.org/10.1038/nphys3736>.
- [2] S. Ghorbanpour, R. L. Gutierrez, V. Gundecha, D. Rengarajan, A. R. Babu, and S. Sarkar. Llm enhanced bayesian optimization for scientific applications like fusion. In *NeurIPS 2024 Workshop on Physical Sciences*, 2024.
- [3] V. Gopalaswamy, R. Betti, J. P. Knauer, N. Luciani, D. Patel, K. M. Woo, A. Bose, I. V. Igumenshchev, E. M. Campbell, K. S. Anderson, K. A. Bauer, M. J. Bonino, D. Cao, A. R. Christopherson, G. W. Collins, T. J. B. Collins, J. R. Davies, J. A. Delettrez, D. H. Edgell, R. Epstein, C. J. Forrest, D. H. Froula, V. Y. Glebov, V. N. Goncharov, D. R. Harding, S. X. Hu, D. W. Jacobs-Perkins, R. T. Janezic, J. H. Kelly, O. M. Mannion, A. Maximov, F. J. Marshall, D. T. Michel, S. Miller, S. F. B. Morse, J. Palastro, J. Peebles, P. B. Radha, S. P. Regan, S. Sampat, T. C. Sangster, A. B. Sefkow, W. Seka, R. C. Shah, W. T. Shmyada, A. Shvydky, C. Stoeckl, A. A. Solodov, W. Theobald, J. D. Zuegel, M. G. Johnson, R. D. Petrasso, C. K. Li, and J. A. Frenje. Tripled yield in direct-drive laser fusion through statistical modelling. *Nature*, 565(7741):581–586, Jan. 2019.
- [4] V. Gopalaswamy, C. A. Williams, R. Betti, D. Patel, J. P. Knauer, A. Lees, D. Cao, E. M. Campbell, P. Farmakis, R. Ejaz, K. S. Anderson, R. Epstein, J. Carroll-Nellenbeck, I. V. Igumenshchev, J. A. Marozas, P. B. Radha, A. A. Solodov, C. A. Thomas, K. M. Woo, T. J. B. Collins, S. X. Hu, W. Scullin, D. Turnbull, V. N. Goncharov, K. Churnetski, C. J. Forrest, V. Y. Glebov, P. V. Heuer, H. McClow, R. C. Shah, C. Stoeckl, W. Theobald, D. H. Edgell, S. Ivancic, M. J. Rosenberg, S. P. Regan, D. Bredesen, C. Fella, M. Koch, R. T. Janezic, M. J. Bonino, D. R. Harding, K. A. Bauer, S. Sampat, L. J. Waxer, M. Labuzeta, S. F. B. Morse, M. Gatu-Johnson, R. D. Petrasso, J. A. Frenje, J. Murray, B. Serrato, D. Guzman, C. Shulderberg, M. Farrell, and C. Deeney. Demonstration of a hydrodynamically equivalent burning plasma in direct-drive inertial confinement fusion. *Nature Physics*, 20(5):751–757, May 2024.
- [5] V. Gundecha, R. L. Gutierrez, S. Ghorbanpour, R. Ejaz, V. Gopalaswamy, R. Betti, D. Rengarajan, and S. Sarkar. Meta-learned bayesian optimization for energy yield in inertial confinement fusion. In *NeurIPS 2024 Workshop on Physical Sciences*, 2024. URL https://ml4physicalsciences.github.io/2024/files/NeurIPS_ML4PS_2024_4.pdf.
- [6] R. L. Gutierrez, S. Ghorbanpour, V. Gundecha, R. Ejaz, V. Gopalaswamy, R. Betti, A. Naug, D. Rengarajan, A. R. Babu, P. Faraboschi, et al. Explainable meta bayesian optimization with human feedback for scientific applications like fusion energy. In *NeurIPS 2024 Workshop on Tackling Climate Change with Machine Learning*, 2024.
- [7] E. C. Hansen, D. H. Barnak, R. Betti, E. M. Campbell, P.-Y. Chang, J. R. Davies, V. Y. Glebov, J. P. Knauer, J. Peebles, S. P. Regan, and A. B. Sefkow. Measuring implosion velocities in experiments and simulations of laser-driven cylindrical implosions on the omega laser. *Plasma Physics and Controlled Fusion*, 60(5):054014, apr 2018. doi: 10.1088/1361-6587/aab73f. URL <https://dx.doi.org/10.1088/1361-6587/aab73f>.
- [8] S. Hochreiter and J. Schmidhuber. Long short-term memory. *Neural computation*, 9(8): 1735–1780, 1997.
- [9] D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [10] A. Shmakov, A. Naug, V. Gundecha, S. Ghorbanpour, R. L. Gutierrez, A. R. Babu, A. Guillen, and S. Sarkar. Rtdk-bo: High dimensional bayesian optimization with reinforced transformer deep kernels. In *2023 IEEE 19th International Conference on Automation Science and Engineering (CASE)*, pages 1–8, 2023. doi: 10.1109/CASE56687.2023.10260520.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, volume 30, 2017.

Appendix

A Experimental Details

Num of layers	4
Number of hidden units	512
Learning rate	$1e^{-5}$
Batch size	128

Table 2: Network Hyperparameters used in the LSTM auto-regressive model.

Num of layers	4
D_{model}	256
$D_{feedforward}$	512
Num. attention heads	16
Learning rate	$1e^{-5}$
Batch size	128

Table 3: Network Hyperparameters used in the Transformer model.

$\lambda_{\mathcal{L}}$	0.5
$\lambda_{\mathcal{P}}$	0.25
$\lambda_{\mathcal{E}}$	0.25

Table 4: Hyperparameters used for our inverse model loss function $\mathcal{G}(\theta)$.

A.1 LILAC surrogate

The LILAC surrogate ($\mathbf{S}_{\phi}$) is a MLP with the following hyper-params. It is trained with the Adam optimizer [9], with a batch size of 128, learning rate of $1e^{-5}$ for 100 epochs.

Num. of layers	4
Num. of hidden units	256
Activation function	ReLU

Table 5: Hyperparameters used for the LILAC surrogate $\mathbf{S}_{\phi}$

The surrogate has an average error of 1.4% across all the outputs.

A.2 Mapping Function

The mapping function ($\mathcal{M}$) is a 1 layer bidirectional LSTM with 64 hidden units. We concatenate the hidden state from both directions, before passing it through an output linear layer. It is trained with the Adam optimizer, with a batch size of 64, learning rate of $1e^{-5}$ for 250 epochs. The model achieves an average error of 1.1% across all 12 parameters in $\mathcal{C}$

B Example Pulse Shape Generation

We present some of the pulses generated by the different approaches.




Figure 4: LPs generated by LPDS_{LSTM} model




Figure 5: LPs generated by the $\text{LPDS}_{\text{Transformer}}$ model




Figure 6: LPs generated by $\text{LPDS}_{\text{GaussianAR}}$ model




Figure 7: LPs generated by the $\text{LPDS}_{\text{Mixture of Gaussian AR}}$ model




Figure 8: LPs generated by the model when doing inpainting using $\text{LPDS}_{\text{LSTM}}$.