DiffICF: Diffusion-Driven Inverse Modeling for Laser Pulse Design in Inertial Confinement Fusion

Ricardo Luna Gutierrez[1*], Vineet Gundecha[1*], Rahman Ejaz[2*], Varchas Gopalaswamy[2*], Riccardo Betti[2*], Sahand Ghorbanpour[1*], Aarne Lees[1*], Soumyendu Sarkar[1*]

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

Abstract

Traditional design of Laser Pulse Shapes (LPs) for Inertial Confinement Fusion (ICF) is a significant bottleneck, relying on computationally expensive simulations and manual iterative refinement. We introduce Diffusion-Driven Inverse Modeling for Laser Pulse Design (DiffICF), a generative inverse model that directly maps specified implosion outcomes to tailored LPs. DiffICF incorporates a physics-informed loss function that enforces known experimental and physical constraints. Moreover, it enables fine-grained control over pulse characteristics through constraint conditioning and inpainting. The efficacy of this framework was experimentally validated for optimizing implosion outcomes, offering a scalable, data-driven design tool to accelerate progress in fusion energy.

1 Introduction

The success of Inertial Confinement Fusion (ICF) experiments hinges on a precisely designed Laser Pulse (LP). This LP is the temporal profile that drives fuel pellet compression. Current LP design relies on expensive simulations and manual exploration, often requiring weeks per high-performing LP. This bottleneck necessitates inverse design models mapping desired outcomes to promising LPs.

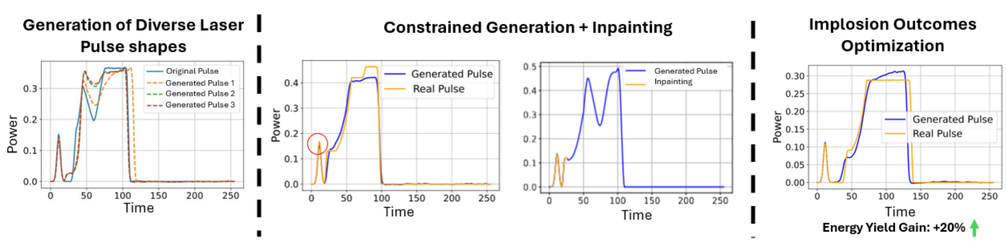

Figure 1: Overview of Diffusion-Driven Inverse Modeling for Laser Pulse Design (DiffICF). DiffICF supports diverse generation, constrained generation, inpainting and implosion outcomes optimization

We introduce Diffusion-Driven Inverse Modeling for Laser Pulse Design (**DiffICF**), a diffusion-based inverse modeling framework addressing this challenge. Our contributions are: (1) A diffusion-based inverse model for LP design with auxiliary objectives aligning generation with target outcomes, (2) Physics-informed constraints ensuring experimental feasibility, (3) Conditioning mechanisms for

*These authors contributed equally.

Machine Learning and the Physical Sciences Workshop, NeurIPS 2025.

direct control over specific LP attributes and regions, and (4) Demonstration of LP optimization improving implosion outcomes.

2 Background

2.1 ICF

Inertial Confinement Fusion (ICF) is a promising path to clean energy that uses high-power lasers to implode a fuel pellet and trigger nuclear fusion [1]. The success of this process hinges on the precise design of the Laser Pulse (LP), a complex, temporally-structured signal that governs the implosion. Due to high experimental costs, LP designs are typically evaluated using sophisticated physics simulators like LILAC [2], which map a given LP to key fusion performance metrics. Various Bayesian optimization based approaches have been used for better outcomes[3–6]. This simulation-based approach, however, remains a significant bottleneck.

2.2 Denoising Diffusion Models

Denoising Diffusion Probabilistic Models (DDPMs) [7] are generative models that use a forward diffusion process to gradually add Gaussian noise to a data sample $\mathbf{x}$ over N steps. This forward process is a Markov chain, where each step n is defined by $q(\mathbf{x}_n|\mathbf{x}_{n-1}) = \mathcal{N}(\mathbf{x}_n; \sqrt{1-\beta_n}\,\mathbf{x}_{n-1}, \beta_n\,\mathbf{I})$.

Denoising Diffusion Implicit Models (DDIMs) [8] accelerate sampling by replacing DDPM's stochastic reverse process with a deterministic update. The deterministic reverse update is given by $\mathbf{x}_{n-1} = \sqrt{\bar{\alpha}_{n-1}}\,\mathbf{x}' + \sqrt{1-\bar{\alpha}_{n-1}-\sigma_n^2}\,\epsilon_\theta(\mathbf{x}_n, n) + \sigma_n\epsilon_n$, where $\epsilon_n \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and $\sigma_n \in [0,1]$ controls the stochasticity. The original input $\mathbf{x}'$ is stimated as:

$$\mathbf{x}' \approx \frac{\mathbf{x}_n - \sqrt{1-\bar{\alpha}_n}\,\epsilon_\theta(\mathbf{x}_n, n)}{\sqrt{\bar{\alpha}_n}}. \tag{1}$$

DDIM trains the noise prediction network $\epsilon_\theta(\mathbf{x}_n, n)$ following the noise prediction objective:

$$\mathcal{L}(\theta) = \mathbb{E}_{n\sim[1,N],\,\mathbf{x}\sim q(\mathbf{x}),\,\epsilon\sim\mathcal{N}(\mathbf{0},\mathbf{I})}\left[\|\epsilon - \epsilon_\theta(x_n, n)\|^2\right]. \tag{2}$$

3 Diffusion-Driven Inverse Modeling for Laser Pulse Design

In this Section, we describe (**DiffICF**). DiffICF takes desired implosion outcomes and target pellet parameters as input, aiming to generate LPs that achieve specified fusion objectives.

3.1 Inverse Modeling

Let $\mathbf{m} = \text{LILAC}(\mathbf{l}, \mathbf{p})$ denote the vector of implosion outcomes produced by the LILAC 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. 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 $\mathbf{p}$ as $\mathbf{l}' = G_\theta(\mathbf{m}, \mathbf{p})$.

We leverage DDIM [8] to build G_θ, conditioning it on both the desired implosion outcomes $\mathbf{m}$ and the target parameters $\mathbf{p}$. During the forward diffusion process, noise is gradually added to the original LP $\mathbf{l}$ over N timesteps, resulting in a noisy latent variable $\mathbf{l}_N \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. The inverse model learns a conditional noise prediction network $\epsilon_\theta(\mathbf{l}_n, n, \mathbf{m}, \mathbf{p})$ that estimates the noise added during a forward process, conditioned on the noisy LP $\mathbf{l}_n$, the timestep n, the target implosion outcomes $\mathbf{m}$, and the target parameters $\mathbf{p}$. The training objective is:

$$\mathcal{L}(\theta) = \mathbb{E}_{n\sim[1,N],(\mathbf{l},\mathbf{m},\mathbf{p})\sim D_F}\left[\|\epsilon - \epsilon_\theta(\mathbf{l}_n, n, \mathbf{m}, \mathbf{p})\|^2\right], \tag{3}$$

where $\mathbf{l}_n = \sqrt{\bar{\alpha}_n}\,\mathbf{l} + \sqrt{1-\bar{\alpha}_n}\,\epsilon$.

Auxiliary Objective While the reconstruction objective enables LP reconstruction, 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.

We train a surrogate model $\mathbf{S}_\phi$ 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 $\mathbf{m}$: $\hat{\mathbf{m}} = \mathbf{S}_\phi(\mathbf{l}, \mathbf{p})$. The surrogate is trained by minimizing:

$$\text{LILAC}_{\text{surrogate}}(\phi) = \mathbb{E}_{(\mathbf{l},\mathbf{m},\mathbf{p})\sim D_F} \left[\|\mathbf{m} - \mathbf{S}_\phi(\mathbf{l}, \mathbf{p})\|^2 \right]. \tag{4}$$

After training, $\mathbf{S}_\phi$ is kept frozen and used during training of the inverse model to evaluate the physical fidelity of generated LPs. Our new objective is:

$$\mathcal{G}(\theta) = \lambda \mathcal{L}(\theta) + \mathcal{S}(\theta) \tag{5}$$

$$\mathcal{S}(\theta) = \mathbb{E}_{\mathbf{l}'\sim G_\theta,(\mathbf{m},\mathbf{p})\sim D_F} \left[\|\mathbf{m} - \mathbf{S}_\phi(\mathbf{l}', \mathbf{p})\|^2 \right] \tag{6}$$

Here, λ controls the trade-off between faithfully reconstructing the original LP and generating physically consistent alternatives. A properly tuned λ allows the model to generalize beyond pure reconstruction, producing diverse yet plausible LPs tailored to the specified design goals.

Physics-Informed Loss Although our goal is to generate diverse LPs, their practical utility depends on adherence to fundamental physics and real-world experimental constraints. One essential constraint is energy conservation. Generated LPs should not exceed the total energy budget of corresponding reference pulses. For any $\mathbf{l}_i \in D_F$, we impose the constraint on a generated LP $\mathbf{l}'_i$ as $\int_0^T \mathbf{l}'_i(t)\,dt \leq \int_0^T \mathbf{l}_i(t)\,dt, \quad t \in [0, T]$, where T is the total pulse duration. To enforce this, we introduce a physics-based penalty term:

$$\mathcal{P}(\theta) = \mathbb{E}_{\mathbf{l}'\sim G_\theta,\mathbf{l}\sim D_F} \left[\left(\int_0^T \mathbf{l}'(t)\,dt - \int_0^T \mathbf{l}(t)\,dt \right)^+ \right] \tag{7}$$

The complete training objective becomes $\mathcal{G}(\theta) = \lambda_\mathcal{L}\mathcal{L}(\theta) + \lambda_\mathcal{S}\mathcal{S}(\theta) + \lambda_\mathcal{P}\mathcal{P}(\theta)$

Pulse Shape Constrained Design In many ICF experimental scenarios, scientists require control over specific regions or attributes of the LP. A mapping function $\mathcal{M}$ projects an LP $\mathbf{l}$ onto $M = 12$ parameters $C = \mathcal{M}(\mathbf{l}) = \{c^1, \ldots, c^{12}\}$. Each c^m corresponds to a physically interpretable property of the LP. When constructing new designs, scientists fix one or two of these values within C, while allowing the remaining parameters to vary.

To enable constraint conditioning, we introduce an embedding network $\mathcal{B}_\delta$ that encodes constraint parameters $c \in C$. We enforce equal embedding dimensionality for timestep n' and constraint $c' = \mathcal{B}_\delta(c)$, enabling their combination into $u' = n' + c'$. This allows conditioning as $\epsilon_\theta(\mathbf{l}_t, u', \mathbf{m}', \mathbf{p}')$ without architectural changes.

For rapid adaptation to new constraints, we fine-tune only $\mathcal{B}_\delta$ parameters according to the constraining loss:

$$\mathcal{J}(\theta) = \mathbb{E}_{\mathbf{l}',c^m} \left[\|c^m_{l'} - c^m\|^2 \right] \tag{8}$$

where $c^m_{l'} \in C_{l'}, C_{l'} = \mathcal{M}(\mathbf{l}')$. This strategy facilitates efficient, stable adaptation while preserving model integrity.

Inpainting For additional controllability, we support inpainting-based generation, enabling scientists to specify desired LP regions directly. Users provide partial LP segments, and the model conditions its denoising process on specified segments, producing coherent, physically valid completions (details in Appendix).

4 Experiments

Performance Evaluation We compare **DiffICF** against state-of-the-art generative models: Tabsyn [9] for tabular data and Variational Latent Diffusion (VLD) [10] for high-energy physics. Addition-

Approach	**Diversity** $\uparrow$	$\mathbf{m}$ **Error** $\downarrow$	**Recon. Error** $\downarrow$	**Energy Cons.** $\downarrow$
$\mathbf{DiffICF}$	0.67	$1.95\% \pm 0.009$	0.007 ± 0.0001	$1.67\% \pm 0.005$
Tabsyn	0.69	$15.4\% \pm 0.2$	0.021 ± 0.0003	$5.67\% \pm 0.014$
VLD	0.39	$17.12\% \pm 0.3$	$0.016 \pm 9.7e^{-5}$	$1.88\% \pm 0.027$
$\mathbf{DiffICF}_{\text{w/o } \mathcal{S}}$	0.62	$4.5\% \pm 0.014$	$0.0057 \pm 8.8e^{-5}$	$1.71\% \pm 0.005$
$\mathbf{DiffICF}_{\text{w/o } \mathcal{P}}$	0.69	$1.99\% \pm 0.01$	0.0084 ± 0.0001	3.21 ± 0.007

Table 1: ICF model performance comparison. $\pm$ denotes standard deviation over seeds.

ally, we evaluate ablated versions excluding the auxiliary loss $\mathcal{S}(\theta)$ ($w/o\, \mathcal{S}$) and the physics-informed loss ($w/o\, \mathcal{P}$).

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 on a test set disjoint from training data. The implosion outcomes error is computed via the surrogate loss and reported as mean average percentage error (MAPE). Reconstruction error is: $\frac{1}{E}\sum_{e=1}^{E}\frac{1}{R}\sum_{r=1}^{R}\|\mathbf{l}_e - \mathbf{l'}_{e,r}\|$. Energy conservation is calculated as percentage deviation from original energy budget. Diversity is measured by average pairwise L2 distance across generated samples. The evaluation results highlight the effectiveness of DiffICF. $\mathbf{DiffICF}$ achieves superior performance with 1.95% implosion error while maintaining strong diversity (0.67). Ablation studies confirm the critical role of both the auxiliary loss $\mathcal{S}(\theta)$ and the physics-informed loss $\mathcal{P}(\theta)$ in enhancing performance. Experimental details are presented in the Appendix.

Constrained Design Performance We evaluate DiffICF's ability to rapidly adapt to specific design constraints using the diffusion-based adaptation approach. The model is fine-tuned for only 10 epochs on new constraint parameters, requiring minimal computational overhead.

Focusing on two physically meaningful parameters (picket power and foot power), we compute the mean absolute percentage error (MAPE) between specified constraint values and corresponding values extracted from generated pulses. The diffusion-based method yields a MAPE of $\mathbf{3.4\% \pm 0.14}$. Importantly, the adaptation method preserves DiffICF's performance on target implosion outcomes.

Inpainting Results When evaluating inpainting with random LP segments specified and remainder masked, DiffICF achieves $2.0\% \pm 0.012$ implosion error and $0.002 \pm 2.32 \times 10^{-6}$ reconstruction error, demonstrating effective coherent completion capabilities.

LP Optimization We evaluate DiffICF for energy-efficient optimization of initial laser pulse designs. Given initial parameters $\mathbf{p}_e$ and implosion outcomes $\mathbf{m}_e$, we define a new target by increasing energy yield by 10%: $\tilde{\mathbf{m}}_e^{\text{yield}} = 1.10 \times \mathbf{m}_e^{\text{yield}}$. Using this target, we generate optimized pulse $\tilde{\mathbf{l}}'_e = G_\theta(\mathbf{p}_e, \tilde{\mathbf{m}}_e)$. The objective is to produce a new LP that: (i) achieves predicted energy yield at least as high as $\tilde{\mathbf{m}}_e^{\text{yield}}$ and (ii) does not exceed the original total energy budget. Because $\tilde{\mathbf{m}}_e$ may lie outside the training distribution, we include a frequency-domain similarity constraint to ensure the generated pulse remains close to the original design (details in Appendix).

For each initial design, we generate 20 candidate pulses and select the one that maximizes energy gain while satisfying all constraints. We validate this approach using the LILAC simulator on 60 unique initial designs. $\mathbf{DiffICF}$ achieves a mean energy gain of **17.74%**, without increasing total energy input. Strikingly, $\mathbf{DiffICF}$ surpasses the 10% target, underscoring its efficiency in exploring and identifying superior LP designs for ICF optimization.

5 Conclusion

We introduced DiffICF, a data-driven inverse modeling framework for generating diverse ICF laser pulse shapes satisfying scientist-defined objectives and physical constraints. DiffICF supports constrained pulse design and implosion outcome optimization. DiffICF demonstrates significant optimization potential (17.74% energy gain) while maintaining experimental feasibility. This work provides a generative tool for accelerating fusion research through machine learning.

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] 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 Gatu 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, January 2019.

[3] Vineet Gundecha, Ricardo Luna Gutierrez, Sahand Ghorbanpour, Rahman Ejaz, Varchas Gopalaswamy, Riccardo Betti, Desik Rengarajan, and Soumyendu 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`.

[4] Ricardo Luna Gutierrez, Sahand Ghorbanpour, Vineet Gundecha, Rahman Ejaz, Varchas Gopalaswamy, Riccardo Betti, Avisek Naug, Desik Rengarajan, Ashwin Ramesh Babu, Paolo 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.

[5] Sahand Ghorbanpour, Ricardo Luna Gutierrez, Vineet Gundecha, Desik Rengarajan, Ashwin Ramesh Babu, and Soumyendu Sarkar. Llm enhanced bayesian optimization for scientific applications like fusion. In *NeurIPS 2024 Workshop on Physical Sciences*, 2024.

[6] Alexander Shmakov, Avisek Naug, Vineet Gundecha, Sahand Ghorbanpour, Ricardo Luna Gutierrez, Ashwin Ramesh Babu, Antonio Guillen, and Soumyendu 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.

[7] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.

[8] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.

[9] Hengrui Zhang, Jiani Zhang, Zhengyuan Shen, Balasubramaniam Srinivasan, Xiao Qin, Christos Faloutsos, Huzefa Rangwala, and George Karypis. Mixed-type tabular data synthesis with score-based diffusion in latent space. In *The Twelfth International Conference on Learning Representations*, 2024. URL `https://openreview.net/forum?id=4Ay23yeuz0`.

[10] Alexander Shmakov, Kevin Greif, Michael James Fenton, Aishik Ghosh, Pierre Baldi, and Daniel Whiteson. End-to-end latent variational diffusion models for inverse problems in high energy physics. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2023.