












RESEARCH ARTICLE | JUNE 04 2026

Bayesian and particle swarm approaches to inertial confinement fusion optimization

J. J. Lee ; D. Coope ; J. Redfern ; H. Martin ; R. Timmis ; E. Denis ; A. James ; R. T. Ruskov ; Z. Zhang ; P. Norreys ; R. W. Paddock 



Phys. Plasmas 33, 062702 (2026)

<https://doi.org/10.1063/5.0331142>



View
Online



Export
Citation

Articles You May Be Interested In

Hydroscaling indirect-drive implosions on the National Ignition Facility

Phys. Plasmas (June 2022)

What next: Further implosion space exploration on the path to NIF extended yield capability

Phys. Plasmas (June 2024)

Particle swarm optimization of 1D isochoric compression designs for fast ignition

Phys. Plasmas (February 2025)

HIDEN
ANALYTICAL

Trusted in Research
for over 40 years

Plasma Diagnostics for
Fundamental and Applied Research
Mass & energy analysis of ions, neutrals and radicals

Find Solutions for Your Research

Bayesian and particle swarm approaches to inertial confinement fusion optimization

Cite as: Phys. Plasmas **33**, 062702 (2026); doi: 10.1063/5.0331142

Submitted: 24 February 2026 · Accepted: 18 May 2026 ·

Published Online: 4 June 2026



View Online



Export Citation



CrossMark

J. J. Lee,^{1,a)} D. Coope,¹ J. Redfern,¹ H. Martin,¹ R. Timmis,¹ E. Denis,¹ A. James,¹
R. T. Ruskov,¹ Z. Zhang,¹ P. Norreys,¹ and R. W. Paddock²

AFFILIATIONS

¹Department of Physics, Atomic and Laser Physics Sub-Department, University of Oxford, Clarendon Laboratory, Parks Road, Oxford OX1 3PU, United Kingdom

²Central Laser Facility, UKRI-STFC, Rutherford Appleton Laboratory, Harwell Campus, Didcot OX11 0QX, United Kingdom

^{a)}Author to whom correspondence should be addressed: jordan.lee@physics.ox.ac.uk

ABSTRACT

The optimization of laser pulse shapes and target configurations is central to high-performance inertial confinement fusion (ICF) implosions, yet remains challenging due to the high dimensionality of the design space and the substantial computational and experimental cost of evaluation. This work presents, to our knowledge, the first comparison of Gaussian process-based Bayesian optimization and particle swarm optimization (PSO) frameworks augmented with physics-motivated extensions in full radiation-hydrodynamic ICF optimization under experimentally relevant constraint handling. These methods are first applied to the re-optimization of low-convergence-ratio wetted-foam implosions, providing a benchmark against traditional sequential scan approaches. Both strategies identify improved designs, with Bayesian optimization achieving the highest final performance using fewer simulations, while PSO converges more rapidly in wall-clock time. The PSO framework is then extended to a 16-dimensional fast ignition design problem, where Gaussian process-based Bayesian optimization becomes computationally impractical. In this regime, PSO efficiently identifies a compressed fuel assembly with $\rho R \approx 1.5 \text{ g/cm}^2$ under strict laser intensity and energy constraints. These results demonstrate that the presented optimization strategies outperform conventional scan-based approaches and provide a scalable platform for high-dimensional ICF optimization. Beyond numerical design studies, the same frameworks are directly applicable to experimental optimization campaigns on high-power laser facilities, where limited shot availability and high evaluation cost demand efficient search methodologies.

© 2026 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0331142>

I. INTRODUCTION

Inertial confinement fusion (ICF) is a promising approach to achieving controlled nuclear fusion and has the potential to provide a nearly limitless source of clean energy. In laser-driven ICF, a spherical target containing fusion fuel is rapidly compressed to extreme pressures and temperatures using high-intensity laser pulses, with the aim of initiating thermonuclear burn.^{1–3} Central to the success of high-performance implosions is the design and optimization of target configurations and laser pulse shaping in order to maximize energy coupling and compression while minimizing instability growth throughout the implosion.

Among the various target concepts explored in ICF, wetted-foam targets have garnered significant interest due to the unique control they provide over the hydrodynamic properties of the implosion.^{4–7} The hydrodynamic control and increased vapor fuel mass

enable igniting implosions in low-convergence-ratio (low-CR) regimes^{8–10} that are more stable against hydrodynamic instabilities,¹¹ such as Rayleigh–Taylor and Richtmyer–Meshkov instabilities, which can severely degrade implosion performance. In addition to improved stability, wetted-foam targets are cheaper, faster to fabricate, and better suited to mass production,¹² making them attractive candidates for eventual inertial fusion energy (IFE) applications.

An alternative ICF approach, known as fast ignition (FI), has also emerged as a promising route to achieving high fusion gain.^{13–15} In FI, the fuel compression and ignition phases are decoupled, with a separate, ultra-intense laser pulse used to ignite a pre-compressed fuel assembly.¹⁶ This separation allows the fuel to be compressed to high areal densities at relatively low temperatures, forming an isochoric assembly that is subsequently ignited, and enables the compression and ignition stages to be optimized largely independently. If

successfully implemented, such high fuel-mass assemblies offer the potential for significantly enhanced fusion gain.

Optimizing ICF target designs is challenging due to the high dimensionality of the design space, the nonlinear nature of implosion dynamics, and the substantial computational cost of high-fidelity radiation-hydrodynamic simulations. Performance metrics such as shock timing, fuel areal density, and gain are often noisy and can exhibit multiple physically distinct local optima associated with different implosion pathways. As a result, simple optimization strategies, such as random sampling and sequential linear scans, struggle to efficiently explore these landscapes or capture strong parameter coupling. For instance, previous optimization studies of wetted-foam implosions in the robust low-CR regime relied on sequential scans of individual parameters,¹⁷ inherently limiting sensitivity to coupled effects that can be critical in identifying globally optimal designs.

In recent years, advanced optimization methodologies have received increasing attention within the ICF community, including the application of population-based algorithms,¹⁸ Bayesian^{19,20} methods, genetic algorithms,²¹ and other derivative-free approaches such as Nelder–Mead²² method and covariance matrix adaptation evolution strategy (CMA-ES).²³ These efforts demonstrate the potential of modern optimization strategies to navigate complex, high-dimensional implosion landscapes more effectively than traditional scan-based approaches. Selecting between these approaches is non-trivial, as each exhibits distinct strengths and weaknesses in terms of sample efficiency, scalability, and ability to explore complex, multi-modal optimization landscapes. However, assessments of how such methods compare and perform across ICF regimes, and under experimentally relevant constraint handling, remains limited.

Our contribution to this emerging area is the development of dedicated Bayesian optimization and particle swarm optimization (PSO) frameworks, implemented in Python and specifically adapted for high-cost plasma simulations. These approaches enable efficient global exploration of complex, noisy, and multi-modal design spaces while accommodating the parallel execution requirements of large-scale simulation campaigns. This capability is achieved through the integration of a set of established algorithmic enhancements tailored to the demands of radiation-hydrodynamic ICF optimization. While the individual enhancements are not themselves novel, their curation and integration into unified optimization frameworks constitute a key contribution of this work. The optimization packages are available for academic research use under license from the authors' institutions, with commercial licensing arrangements available upon request.

In this work, the frameworks are first applied to the re-optimization of low-convergence-ratio (low-CR) wetted-foam implosions, providing a controlled benchmark against previously published sequential scan approaches. The particle swarm optimization framework is then extended to a substantially higher-dimensional problem of isochoric fast ignition target design, where the compression phase of the implosion is optimized. In this case, Gaussian process-based Bayesian optimization becomes computationally impractical. Fast ignition presents a different optimization landscape, characterized by distinct objective functions and constraints compared to conventional central-hot-spot implosions. Its inclusion therefore demonstrates that the optimization frameworks developed here are applicable across a broad range of design challenges. These applications utilize the radiation-hydrodynamics code HYADES²⁴ to execute the simulations.

The remainder of this paper is structured as follows: Secs. II and III describe the Bayesian and particle swarm optimization methodologies, respectively, including the key extensions added to the packages introduced for ICF applications. Section IV presents the re-optimization of wetted-foam target designs previously studied in Ref. 17, serving as a benchmark to validate performance. Section V then details the optimization of an isochoric fuel assembly for fast ignition. The main findings and implications are summarized in Sec. VI.

II. BAYESIAN OPTIMIZATION

A. Overview of Bayesian optimization

Bayesian optimization^{25,26} is an efficient strategy for optimizing expensive black-box functions under uncertainty, making it well suited to ICF optimization, where each function evaluation corresponds to a costly simulation or experiment. Rather than sampling the design space blindly or relying on local search methods, Bayesian optimization constructs a surrogate model of the objective function using the results of previous evaluations of the parameter space. This surrogate provides both a prediction of performance and an estimate of uncertainty across the parameter space, allowing new simulations to be selected in a way that balances exploration of poorly understood regions against exploitation of promising designs.

In this work, Gaussian processes (GPs)²⁷ are used to form a surrogate model, which provides a flexible, non-parametric representation of the objective function together with analytic uncertainty estimates. A GP is fully specified by a mean function, $m(\mathbf{x})$, and a kernel function, $k(\mathbf{x}_1, \mathbf{x}_2)$, which encodes correlations between points in the parameter space. These correlations reflect assumptions about the smoothness and characteristic length scales of the underlying objective function. Functions which display characteristics aligned with the kernel function are considered more plausible by the GP.

Throughout this work, a radial basis function (RBF) kernel is employed, which assumes that the objective function varies smoothly with respect to changes in the input parameters. The RBF kernel is given by

$$k_{\text{RBF}}(\mathbf{x}_1, \mathbf{x}_2) = \exp \left[-\frac{1}{2} \sum_i \frac{(x_{1,i} - x_{2,i})^2}{\ell_i^2} \right], \quad (1)$$

where ℓ_i is the characteristic length scale along parameter i that controls how rapidly correlations decay with distance in the parameter space. Smaller values of ℓ allow for more rapid variation of the objective function, while larger values enforce smoother behavior. In the context of laser-driven ICF implosions, this corresponds to the physical expectation that continuous variations in laser pulse or target parameters lead to correspondingly smooth changes in implosion performance, while still permitting strongly nonlinear global behavior.

As the parameter space is sampled, the GP is conditioned on the resulting data, updating the surrogate model from a prior distribution to a posterior distribution that incorporates the evaluated samples. Conditioning modifies both the predicted mean and the associated uncertainty across the parameter space, pulling the mean toward the observed data points and reducing uncertainty in regions correlated with sampled points. This iterative updating process allows the surrogate model to progressively refine its representation of the objective function as additional data becomes available.

GP-based Bayesian optimization is able to naturally account for noise in the evaluated data. Observational noise is modeled using a Gaussian distribution, such that the observed values are assumed to be normally distributed around the measured value. Inside the model, all object evaluations are normalized between -1 and 1 , and in the optimization campaigns presented here, the magnitude of the noise was of the order 10^{-3} – 10^{-4} . This reflects the low (but present) noise in hydrodynamic simulations.

The next point to evaluate is guided by an acquisition function. The acquisition function utilizes the predicted performance of the objective function and the associated uncertainty to provide a controlled balance between exploiting regions expected to yield high performance and exploring regions that remain poorly characterized. In this work, an upper confidence bound (UCB) acquisition function is employed due to its simplicity, robustness, and transparent physical interpretation. The UCB acquisition function is defined as

$$A_{\text{UCB}}(\mathbf{x}) = m(\mathbf{x}) + \kappa \sigma(\mathbf{x}), \quad (2)$$

where $m(\mathbf{x})$ and $\sigma(\mathbf{x})$ are the posterior mean and standard deviation predicted by the GP at location \mathbf{x} , respectively. The parameter κ controls the relative weighting of exploration and exploitation, with larger values favoring sampling in regions of high uncertainty and smaller values prioritizing regions of high predicted performance.

In this work, we restrict attention to the RBF kernel and UCB acquisition function for simplicity and robustness. We note, however, that these choices are not necessarily optimal for all ICF optimization landscapes, and alternative kernels or acquisition strategies may yield improved convergence behavior depending on the structure of the objective function.^{26–28}

This describes the iterative loop of Bayesian optimization. An initial Gaussian process prior is defined to encode assumptions about the objective function, which is subsequently conditioned on sampled data to form a posterior distribution. An acquisition function is then used to identify the next most informative point in the parameter space, which is evaluated and incorporated into the dataset, and the process is repeated until a stopping criterion is met. This procedure is illustrated in Fig. 1, using an RBF kernel length scale of $\ell = 0.1$ and a UCB acquisition function $\kappa = 2$.

B. Physics-motivated extensions to the Bayesian optimization package

While standard Bayesian optimization is effective for problems with short evaluation times, several challenges arise when applying it to ICF simulations. In particular, ICF optimization involves high computational costs per evaluation, the need to exploit parallel simulation workflows, noisy objective functions, and moderately high-dimensional parameter spaces. To address these challenges, a number of physics-motivated extensions to the Bayesian optimization procedure have been implemented in the developed package, enabling reduced time-to-solution, improved efficiency, robustness, and scalability in practical ICF optimization tasks.

1. Batch sampling

In large-scale ICF optimization campaigns, simulations are typically executed on high-performance computing (HPC) clusters, where multiple simulations can be run concurrently. Standard Bayesian

optimization, which selects a single new evaluation per iteration, is therefore inefficient in this context. To better exploit available computational resources, batch sampling is implemented, allowing multiple points in the parameter space to be selected simultaneously during each optimization iteration.

Batch selection is achieved by varying the exploration-exploitation balance within the acquisition function.²⁹ In the case of the UCB acquisition function, this corresponds to sampling multiple candidate points using different values of the exploration parameter κ . Lower κ values favor exploitative sampling near the current optimum, while higher values promote exploration of uncertain regions. By combining these samples into a single batch, the optimization process efficiently spans a diverse range of candidate designs within each iteration. This approach not only fosters a richer sampling strategy but also mitigates the sensitivity of the optimization process to the choice of κ .

2. Sub-batch sampling (hallucination)

While batch sampling reduces the time-to-solution, it can lead to redundant sampling if multiple points within a batch cluster around the same region of the parameter space, as each point is selected using an identical surrogate model. To mitigate this effect, a sub-batch sampling strategy, sometimes referred to as hallucination, is employed to improve diversity within each batch.

The batch is divided into smaller subsets, which are selected sequentially. After the samples in a given subset are identified, the GP is temporarily conditioned on these locations using the surrogate model's predicted mean values.³⁰ This reduces the predicted uncertainty in regions surrounding the selected points without biasing the model toward any particular objective value. Subsequent subsets are then selected using this updated surrogate, discouraging redundant sampling and promoting exploration of distinct regions of the parameter space. Once the full batch has been assembled, the true objective function values are evaluated, and the surrogate model is updated accordingly.

3. Dynamic bounds reduction

As the dimensionality of the optimization problem increases, the efficiency of Bayesian optimization can degrade due to the exponential growth of the search space. To address this, dynamic bounds reduction is implemented, enabling the optimization to progressively focus on promising regions of the parameter space as confidence in the surrogate model increases.³¹ By adaptively contracting the bounds, the effective search volume is reduced, concentrating computational effort in regions of interest and maintaining sufficient sampling density for the GP surrogate to remain informative.

Initially, the optimization is performed over the full parameter bounds. As the optimization proceeds and improvements stagnate, the bounds are adaptively contracted around the best-performing configurations identified. This reduction is performed conservatively to avoid excluding viable optima, while substantially reducing the effective search volume. By iteratively refining the parameter bounds in this manner, the optimization process maintains global exploration early on and transitions naturally toward local refinement in later stages. This enables effective optimization in moderately

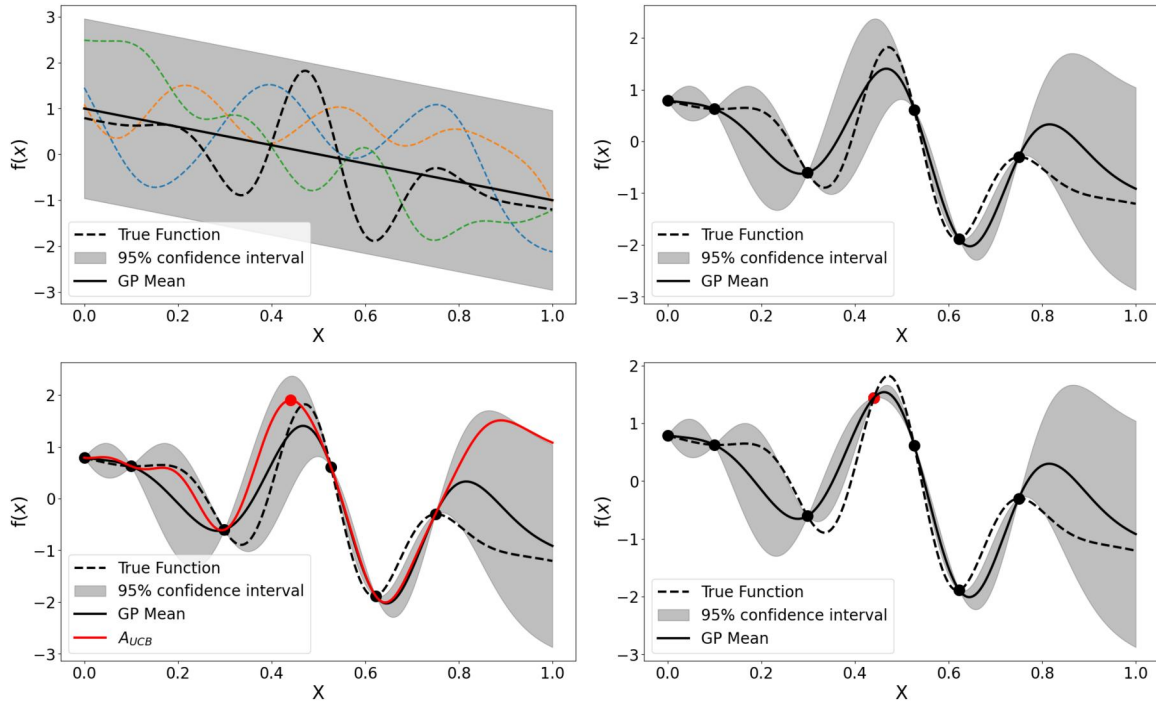


FIG. 1. Demonstration of the Bayesian optimization iterative loop. (a) A GP prior using an RBF kernel and a prior mean function of $\mathbf{m}(x) = 1 - 2x$. GP samples are indicated by the colored dashed lines and indicate the type of function the GP deems plausible. (b) Posterior GP after conditioning on six data points. The GP mean interpolates the sampled data, and the uncertainty in regions correlated with these samples is reduced. (c) Application of the UCB acquisition function to the posterior in (b). The point that maximizes the acquisition function is identified as the next sample location. (d) GP posterior after conditioning on the newly sampled data point.

higher-dimensional parameter spaces than would otherwise be possible using standard Bayesian optimization.

4. In situ hyperparameter optimization

The performance of GP surrogate models depends on kernel hyperparameters, such as characteristic length scales, which encode assumptions about the structure of the objective function. In many practical ICF optimization problems, appropriate values for these hyperparameters are not known *a priori* and may vary throughout the optimization process.

To reduce reliance on prior assumptions, kernel hyperparameters are optimized *in situ* by maximizing the log marginal likelihood of the observed data³² using a bounded quasi-Newton method. A mild regularization term is also applied to discourage extreme disparities between length-scales. This allows the surrogate model to adapt its representation of the objective function as additional data becomes available, automatically selecting length scales consistent with the observed behavior while penalizing overly complex models. This adaptive approach improves robustness to noise, reduces the need for manual tuning, and enhances the reliability of uncertainty estimates used to guide the optimization.

III. PARTICLE SWARM OPTIMIZATION

A. Overview of particle swarm optimization

Particle swarm optimization (PSO)³³ is a population-based optimization algorithm in which a swarm of candidate solutions (referred

to as particles) moves through the parameter space in search of an optimum. The particles evolve through the parameter space by combining individual experience with collective information gathered by the swarm. This population-based search enables efficient exploration of complex parameter spaces with minimal algorithmic overhead, as no explicit surrogate model is constructed and the computational cost is dominated by the objective function evaluations themselves.

Each particle represents a point in the parameter space and is associated with both a position and a velocity. During the optimization process, particles are influenced by their own best-known position (\mathbf{L}_p^n), as well as the best-known position identified by the swarm as a whole (\mathbf{G}^n). The velocity, \mathbf{v}_p , and position, \mathbf{x}_p , of each particle are then be updated as

$$\mathbf{v}_p^{n+1} = h_i \cdot \mathbf{v}_p^n + r_l^n h_l \cdot (\mathbf{L}_p^n - \mathbf{x}_p^n) + r_g^n h_g \cdot (\mathbf{G}^n - \mathbf{x}_p^n), \quad (3)$$

$$\mathbf{x}_p^{n+1} = \mathbf{x}_p^n + \mathbf{v}_p^{n+1}, \quad (4)$$

where subscript p denotes particle-specific values, and superscript denotes the iteration.

Equation (3) is comprised of three terms, each introducing a hyperparameter, h_i , h_l , and h_g . In order, these terms can be described as an inertial term, which determines how much the particle remains on its original trajectory, a local attraction term, which determines how strongly the particle is attracted to its own best-known position, and a global attraction term, which determines how strongly the particle is pulled toward the swarm's best found location. In this

description, the inertia weight controls the extent to which particles explore new regions of the parameter space, while the local and global hyperparameters control the strength of the attraction particles feel toward promising areas already explored by themselves or by the swarm. It is through this hyperparameter selection that the explore/exploit trade-off can be adjusted.

The local and global attraction terms also include random numbers, r_l^n and r_g^n , uniformly distributed in $[0, 1]$. The use of these random numbers adds a level of stochasticity to the particle motion, which helps to maintain diversity in the swarm. As a result, particles approach their local and global best positions with differing magnitudes and directions, which helps to preserve swarm diversity and mitigates premature convergence to suboptimal solutions. Importantly, while individual particle updates are stochastic, the expected value of the attraction terms remains biased toward the best-known positions, ensuring the particles are encouraged toward promising regions of the parameter space.

B. Physics-motivated extensions to particle swarm optimization

The absence of an explicit surrogate model makes PSO computationally lightweight and naturally parallelizable, as the objective function can be evaluated independently for each particle. However, in complex optimization landscapes, the swarm can converge prematurely or stagnate around sub-optimal solutions, particularly in the presence of noise or multiple competing optima. Additionally, striking a suitable balance between exploitation and exploration can be difficult. To address these challenges and better adapt PSO to the demands of ICF optimization, several physics-motivated extensions to the standard PSO framework are implemented in the developed package.

1. Particle species

In standard particle swarm optimization, all particles share a common set of hyperparameters, resulting in a uniform search behavior across the swarm. However, this uniform behavior is highly sensitive to the exact choice of hyperparameters. To address this, a multi-species formulation of PSO is introduced, in which the swarm is composed of multiple particle species, each defined by a distinct set of hyperparameters governing inertia, local attraction, and global attraction.^{34,35} Particles are assigned to species according to specified probabilities, allowing multiple search behaviors to coexist simultaneously within a single optimization run. This enables the swarm to maintain global exploration through more exploratory species, while concurrently performing local refinement through more exploitative species.

By allowing different species to operate concurrently, the optimizer can more effectively refine promising configurations while reducing the risk of premature convergence to sub-optimal configurations. The use of multiple species also reduces sensitivity to any single hyperparameter choice, improving robustness across a range of optimization problems.

2. Exploration term

While the introduction of multiple particle species enables a range of search behaviors to coexist within the swarm, standard PSO

dynamics still rely exclusively on attraction toward previously identified promising locations. As a result, even exploratory species can become biased toward regions already sampled by the swarm, particularly once a strong global best has emerged. This behavior can limit the ability of the swarm to systematically probe poorly explored regions of the parameter space.

To address this, an explicit exploration term³⁶ is introduced into the velocity update equation [Eq. (3)], providing a mechanism by which particles can be actively encouraged to sample regions poorly explored. This term introduces an additional attraction toward locations in the parameter space that are distant from previously evaluated points, thereby promoting exploration independently of the current best-known solutions. The velocity update equation now reads

$$\mathbf{v}_p^{n+1} = h_i \cdot \mathbf{v}_p^n + r_l^n h_l \cdot (\mathbf{L}_p^n - \mathbf{x}_p^n) + r_g^n h_g \cdot (\mathbf{G}^n - \mathbf{x}_p^n) + h_E (\mathbf{E}_p^n - \mathbf{x}_p^n), \quad (5)$$

where h_E is the weight of the exploration term and \mathbf{E}_p^n is the region of low sample density identified for particle p at iteration n .

In practice, regions of low sample density are identified by randomly sampling across the parameter space and determining the sample furthest from any previous samples. By explicitly encoding exploration in this manner, the optimizer is able to maintain coverage of the broader parameter space even in the presence of strong local or global attractors.

In the context of ICF optimization, such a mechanism may be valuable for identifying alternative implosion pathways that may not be accessible through incremental refinement of existing designs. When combined with the multi-species framework, the explicit exploration term enables a subset of particles to prioritize global search while others focus on local optimization, improving robustness in complex, multi-modal landscapes. Although this explicit exploration term is now implemented within the PSO framework, it was introduced after completion of the optimization studies presented in this paper. Consequently, it was not activated in the optimization campaigns reported in this work. The results presented, therefore, correspond to the multi-species PSO formulation without explicit exploration forcing.

3. Velocity boosting

Despite the introduction of multi-species behavior and an explicit exploration term, particle swarm optimization can still experience stagnation when particles converge into narrow regions of the parameter space and velocities decay toward zero. In such cases, the swarm may remain trapped near sub-optimal configurations, particularly in noisy or weakly informative optimization landscapes where small gradients provide little guidance.

To mitigate this behavior, a velocity-boosting mechanism is introduced, designed to actively dislodge particles that have stagnated for an extended period.³⁷ When a particle's velocity falls below a specified threshold for a certain number of iterations, its velocity is boosted, effectively redistributing the particle within the local region of the parameter space. This controlled perturbation allows particles to escape shallow basins of attraction or noise-induced traps while preserving the overall swarm dynamics.

In the context of ICF optimization, velocity boosting is particularly useful for navigating rugged objective landscapes where multiple

implosion configurations exhibit comparable performance. By intermittently re-energizing particles that have stalled, the optimizer is better able to identify alternative pathways and avoid premature convergence. When combined with multi-species behavior and explicit exploration, velocity boosting provides a robust mechanism for maintaining long-term swarm activity without compromising convergence toward high-performing solutions.

IV. LOW-CONVERGENCE-RATIO OPTIMIZATION

A. Previous optimization of the low-CR regime

Previous efforts to optimize wetted-foam implosions have focused on maximizing performance in robust regimes that minimize the influence of hydrodynamic and parametric instabilities.¹⁷ This has led to the development of the low-convergence-ratio (low-CR) regime, which imposes restrictions on key implosion parameters. Experimental observations^{6,9,10} indicate that, in satisfying these constraints, the implosion remains stable, and that one-dimensional (1D) radiation-hydrodynamic simulations accurately reproduce experimental results. Consequently, the radiation-hydrodynamic simulations employed here are anticipated to remain predictive for experimental performance within this regime. All simulations in the previous approach and this work are performed using the HYADES²⁴ radiation-hydrodynamics code. The key restrictions that define the low-CR regime in Ref. 17 are as follows:

- Convergence ratio ≤ 16 : The convergence ratio is defined as

$$CR(t) = \frac{R_{\text{initial}}}{R_{\text{HS}}(t)}, \quad (6)$$

where R_{initial} is the initial radius of the foam/ablator boundary, and $R_{\text{HS}}(t)$ is the inner radius of the imploding shell (the hotspot radius). The value used is the maximum CR, providing a measure of the peak compression of the hotspot.

- In-flight aspect ratio (IFAR) ≤ 30 : The IFAR is defined as the ratio of the shell's outer radius to its thickness during implosion. This is evaluated when the outer shell radius is 2/3 of the initial vapor/foam boundary radius.
- Implosion velocity ≤ 400 km/s: This is the mass-averaged velocity of the imploding shell.
- $I\lambda^2 \leq 10^{14} \text{ W}\mu\text{m}^2/\text{cm}^2$: This condition restricts the product of laser intensity (I) and the square of the wavelength (λ).




FIG. 2. Four-pulse laser profile and capsule design construction as described in Ref. 17. Parameters shown in turquoise are free parameters to be optimized.

The hydrodynamic constraints depend on the definition of the shell, which follows that used in Ref. 17. The shell is defined using the density profile of the implosion. The outer shell radius is taken as the radius at which the density falls to $1/e^2$ of the peak density. The inner shell radius (also the hotspot radius) is defined as the radius at which the difference between the local density and the central density first falls to $1/e^2$ of its maximum value. The shell thickness is then defined as the difference between the outer and inner shell radii.

In Ref. 17, five different capsule sizes, each corresponding to a different energy scale, were explored. Each capsule size was optimized for both a three-pulse and four-pulse laser sequence, resulting in a total of ten optimized target configurations. The construction of capsule designs and laser pulse profiles is shown in Fig. 2, with the free parameters to be optimized (laser pulse timings, shell radii, vapor density) indicated in turquoise. Consequently, the three-pulse optimizations spanned a six-dimensional (6D) parameter space, and the four-pulse optimizations were 7D.

The optimization approach employed began from a “reasonable” configuration, followed by a sequential linear scan of the parameter space. The sequential nature of this approach prevents exploration of the interaction between parameters, potentially leading to suboptimal results. Moreover, the laser power of each pulse was fixed and not treated as an optimization parameter. This severely limits the flexibility of the design space, biasing the optimization toward a narrow class of pulse shapes and excluding potentially high-performing configurations, such as those associated with shock-augmented ignition.³⁸

B. Changes to the optimization approach

The re-optimization of these configurations introduces several key differences from the previous work. First, we retain the optimized capsule designs from the previous study, but focus on optimizing the laser pulse profiles for these fixed capsule configurations. In contrast to the earlier study, laser powers are now treated as free parameters, allowing the optimization process to determine their optimal values. This new optimization parameter space is shown in Fig. 3(a) for a three-pulse laser profile, with the parameters to be optimized again indicated in turquoise. While capsule parameters are no longer being optimized, the dimensionality of the parameter space remains similar, with six dimensions in the three-pulse case and eight dimensions for in the four-pulse case.

15 June 2026 11:29:54




FIG. 3. (a) Three-pulse laser profile with the parameters to be optimized indicated in turquoise. (b) Penalty imposed on the gain of the implosion due to exceeding the low-CR IFAR constraint of 30.

In addition to the changes in the optimization parameter space, several changes have been made in the analysis of the implosions. The implosion velocity has been redefined to be the more physically meaningful mass-averaged velocity of the imploding shell, rather than the volume-averaged velocity used previously. Furthermore, the definition of gain was updated from the neutron fusion energy released (approximately 80% of the total fusion energy for DT reactions) normalized by the laser energy, to the total fusion energy released divided by the laser energy. This is the more common definition of gain found in the literature and remains consistent when accounting for fusion reactions beyond DT (such as DD), although these are significantly less frequent, or when using alternate fusion fuel.

It should also be noted that, since the optimization scanned a large parameter space, for a reasonable number of implosions, the CR was infinite, corresponding to cases where the inner shell boundary reached the center of the implosion. Instead of returning invalid values for these implosions, which can confuse the optimizer, the CR was instead defined to use the outer shell boundary, which is always well-defined, with a threshold value of 13 which was selected empirically. This alternative definition preserves the physical intent of the low-CR constraint while remaining well defined across the entire parameter space.

Finally, rather than enforcing hard cutoff values for the low-CR regime parameters, a penalty-based system has been introduced. This allows for smoother exploration of the parameter space, penalizing solutions more heavily as they deviate from the low-CR constraints. This penalty based method is particularly useful due to the ignition instability associated with ICF. Within the low-CR regime, the implosion may be on the cusp of ignition yet still report little gain. This approach enables testing the boundaries of the low-CR regime, potentially surpassing the threshold of ignition slightly and resulting in capsules with significant gain. Figure 3(b) illustrates the penalty imposed for exceeding the low-CR thresholds, using the IFAR as an example. Notably, the penalty applied to implosion velocity uses a more stringent exponential length scale (by a factor of 2), reflecting its stronger influence on hydrodynamic instabilities. The functional form of this multiplier is provided in the supplementary material. Taking these penalties into account, the function to be optimized is given by

$$\text{Result} = \text{gain} \cdot \text{CR}_M \cdot \text{IFAR}_M \cdot v_M \cdot (I\lambda^2)_M, \quad (7)$$

where subscript M refers to the penalty multiplier.

C. Results

Using the optimization frameworks described in Secs. II and III, all five capsule sizes were re-optimized using both three- and four-pulse laser profiles under the revised parameterization and constraint handling outlined above. Owing to the modifications outlined above, the re-optimization results are not strictly directly comparable to those of the original sequential scan study. The benchmark should therefore be interpreted as a qualitative reference point rather than a controlled one-to-one comparison. Nevertheless, it provides a useful indication that the optimization frameworks identify plausible optimal implosion configurations within the revised design space.

The simulation process follows the same structure as the previous optimization. The 1D Lagrangian radiation-hydrodynamics code HYADES²⁴ is used. As such, multidimensional effects such as hydrodynamic instabilities are not captured, and, due to the hydrodynamic nature of the simulations, parametric instabilities are also not captured. The simulations include multi-group radiation diffusion for radiation transport and use the standard HYADES thermonuclear burn model, which couples fusion reactions to local plasma conditions. Electron and ion heat conduction are treated using flux limiters of 0.05 and 0.4, respectively.

The capsule is discretized using 150 equally spaced cells in the inner vapor region, with subsequent cells defined such that the relative mass difference between adjacent cells is less than 2% and the final cell thickness is less than 1 μm . The low-density vapor and wetted-foam regions are modeled as a 50–50 DT mixture using a quotidian equation of state. As in the previous study, the wetted foam is approximated as homogeneous DT and therefore does not explicitly include carbon content. The outer shell consists of a 50–50 CD mixture described using the SESAME equation of state tables.

The resulting optimal configurations are shown in Fig. 4, where they are displayed alongside the sequential linear scan results for reference. In all cases, both the Bayesian and particle swarm optimization identify optima that exceed the results obtained using the sequential scan approach, with comparable final performance between the two advanced methods.

The largest performance improvements are observed for the ~ 2 and ~ 4 MJ implosions, where the achieved gain is improved by a factor of two. These capsules lie close to the ignition threshold, where fusion burn is highly sensitive to small changes in implosion configuration. Ignition is a thermal instability process, which, once started, drastically improves the performance of an implosion. There is

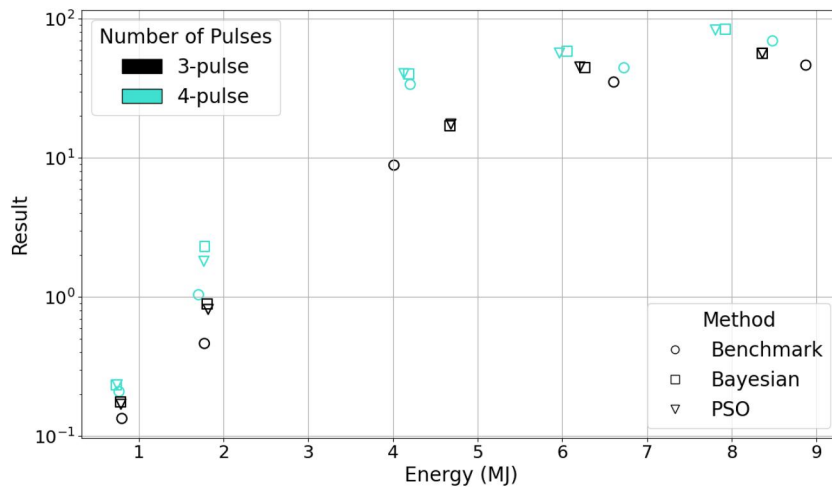


FIG. 4. Comparison of final results from the sequential linear scan, particle swarm optimization, and Bayesian optimization approaches for all five capsule sizes using both three-pulse and four-pulse laser profiles. Both advanced optimization approaches identify improved optimum compared to the sequential linear scan, with comparable performance between the two advanced optimization techniques.

therefore a stark transition from non-igniting, low-gain implosions to high-gain, igniting configurations. This narrow transition region is referred to as the cliff-edge of ignition, and the ~ 2 and ~ 4 MJ implosions sit within the range of this cliff. It is therefore expected that minor configuration adjustments can lead to drastic variations in implosion performance, and capsules operating within this regime are therefore highly sensitive to optimization improvements.

By contrast, the lowest energy capsule (~ 0.7 MJ) remains too far from ignition for optimization to yield substantial gains. Ignition is not possible in the parameter space, and deviations in this space result in minimal change in performance. Similarly, the highest-energy capsules (~ 6.5 and ~ 8.5 MJ) are already robustly igniting with large burn-up fractions, rendering their performance relatively insensitive to further refinement of the design space.

A representative example of the optimization trajectory for the ~ 2 MJ capsule using a three-pulse laser profile is shown in Fig. 5, which displays both the sampled objective values and their locations in the parameter space. Distinct behaviors are evident for the two optimization approaches. In particular, the Bayesian optimizer typically exhibits a slower, more incremental improvement pathway. This behavior arises from the use of a surrogate model, which typically promotes exploitative sampling through small, local adjustments around the currently identified optimum. The application of dynamic bounds reduction further reinforces this behavior, as the effective search region contracts smoothly around the best-performing configuration.

The PSO approach contrasts sharply with this behavior, with the swarm sampling broadly across the parameter space during the early iterations before converging rapidly once a high-performing region is identified. When a distinct region yields improved performance, the swarm quickly reconfigures and coalesces around the newly discovered optimum, as illustrated in Fig. 5 after ~ 3000 simulations.

Two key metrics for assessing the performance of these optimization approaches are the time-to-solution and the computational cost. These two metrics are related, but have important distinctions. Both approaches utilize parallel processes, where a number of simulations are run concurrently to form a batch of simulations per iteration. The time-to-solution, therefore, depends on the number of

iterations required to locate the optimum, whereas the computational cost is determined by the total number of simulations performed.

The two approaches differ considerably in this regard. Bayesian optimization typically prefers smaller batch sizes, with the most efficient approach being single-sample batches. This maximizes the amount of information available before sampling, leading to a more accurate surrogate model from which samples are determined. PSO sits at the other end of the spectrum, preferring larger numbers of particles in the swarm. This allows the swarm to explore the parameter space thoroughly before settling into an optimum. As the swarm size decreases, the particles become more likely to settle into local optima due to insufficient exploration. As a result, the two metrics, time-to-solution and computational cost, are biased toward a specific approach. PSO often converges to an optimum in fewer iterations; however, it requires more simulations to reach this optimum.

In this work, the optimizer is deemed to have identified the optimum once the objective value first reaches 90% of the final optimum achieved during the run. As can be seen in Fig. 5, beyond this threshold, the sampled points occupy only a very small region of the parameter space, indicating that the optimization has effectively converged. Differences between configurations beyond this threshold correspond to minor parameter variations and would be indistinguishable within typical shot-to-shot experimental variability. We have verified that adopting alternative thresholds (e.g., 80% or 95%) produces negligible changes in the inferred time-to-solution, confirming that the conclusions are not sensitive to the precise numerical choice.

Using the definitions, the number of iterations (time-to-solution) and the total number of simulations (computational cost) for each optimization approach are summarized in Table I. Each simulation required approximately 3 h of wall-clock time on a single processor, as the code is not parallelized. With an average of 2000 simulations performed before convergence per optimization, this corresponds to a total computational cost of ~ 6000 processor-hours per capsule configuration, or $\sim 60\,000$ processor-hours for all ten low-CR re-optimizations. As batches were run concurrently, 35 to 50 iterations correspond to roughly one working week of wall-clock time per optimization. Given the substantial duration of each optimization campaign, the full optimization process was performed once per

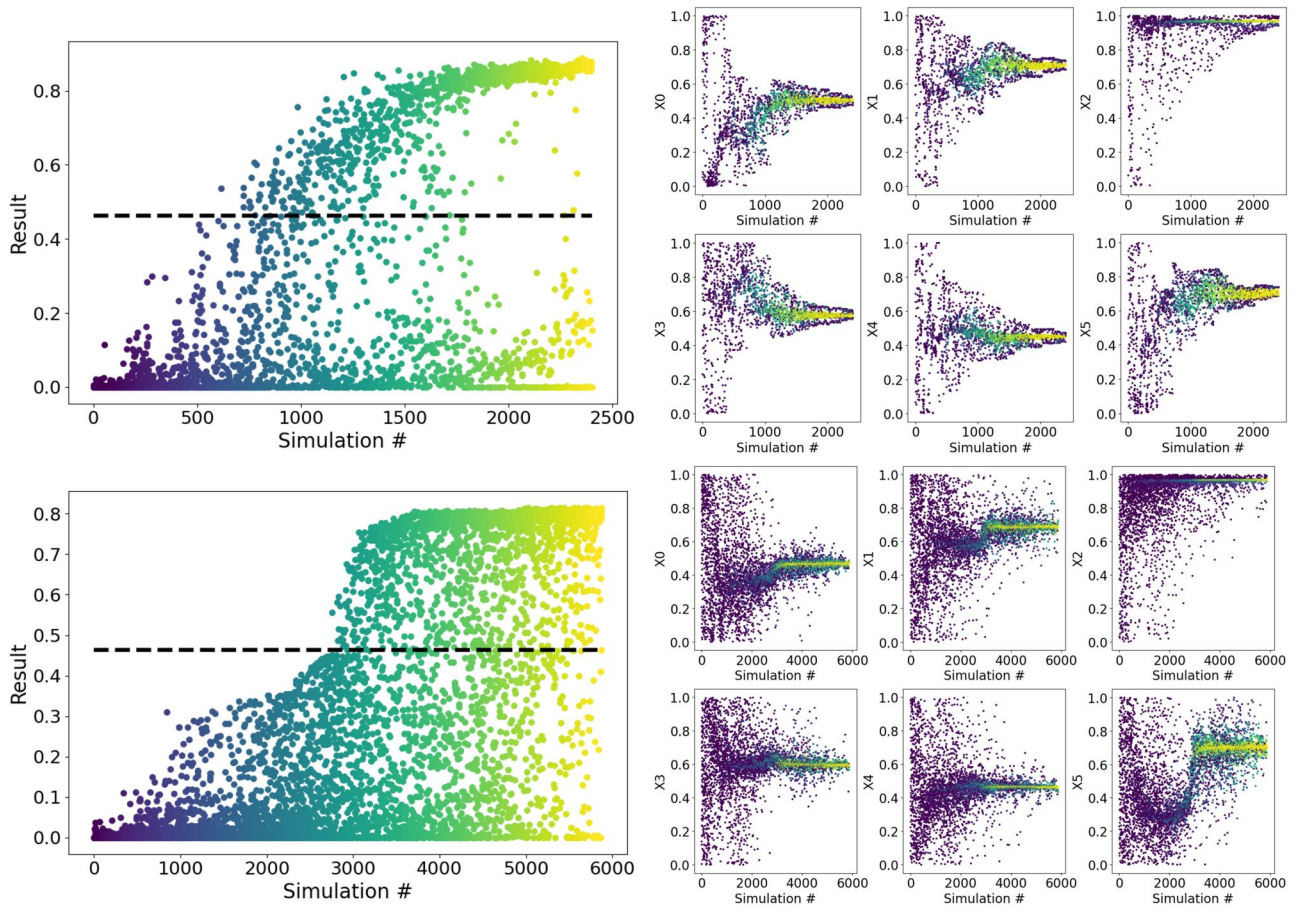


FIG. 5. Evolution of the optimization approaches through the parameter space (right) and the associated evaluation values (left). Bayesian (top) exhibits more gradual improvement, while PSO (bottom) converges rapidly on identified optima. The horizontal dashed line indicates the benchmark optimized value.

TABLE I. Comparison of Bayesian optimization and PSO convergence performance. The Bayesian approach converges more efficiently, on average in less than half the number of function evaluations, while the PSO approach converges in fewer iterations.

Pulse	Capsule size	BO		PSO	
		Iterations	Simulations	Iterations	Simulations
Three-pulse	0.5	45	1078	17	897
	0.65	52	1205	56	3067
	0.85	50	1177	13	736
	1.0	35	952	20	821
	1.1	26	713	20	1795
Four-pulse	0.5	53	1328	61	5687
	0.65	66	6162
	0.85	47	1123	29	2482
	1.0	89	2028	41	3441
	1.1	47	1116	35	2904
	Average	49.3	1190	35.8	2800

15 June 2026 11:29:54

capsule configuration. As a result, repeated stochastic restarts and extensive hyperparameter sensitivity scans were not performed for each capsule configuration. However, both optimization strategies independently converged to comparable optima across all energy scales, and the optimization trajectories exhibited stable late-stage convergence behavior, with sampling concentrated in narrow regions of parameter space once high-performing configurations were identified. Together, these observations suggest that the reported optima are robust within the explored design space.

For comparison, in the benchmark sequential linear scan, approximately 20 000 simulations were required to optimize all ten implosion designs, corresponding to an average of ~ 2000 simulations per design. A batch size of 100 simulations was used per iteration, resulting in a total of ~ 20 iterations. It should be noted that the sequential scan was initialized from a reasonable starting configuration and excluded large regions of poorly performing parameter space as a consequence of this choice. In comparison, the optimization processes presented here started with no prior bias or knowledge, making direct comparison difficult.

Comparisons between the two optimization strategies indicate that Bayesian optimization generally requires significantly fewer simulations, while particle swarm optimization typically converges in fewer iterations, making it the faster method in terms of wall-clock time. Consequently, the preferred approach depends on the optimization context: Bayesian optimization is better suited to simulation-limited studies where sample efficiency is critical, whereas PSO is advantageous in scenarios where rapid convergence is prioritized.

It should also be noted that standard GP-based Bayesian optimization incurs additional computational overhead associated with surrogate model training and kernel hyperparameter optimization. While this overhead is relatively low for the six- and eight-dimensional parameter spaces considered here, it increases rapidly with dimensionality and becomes a limiting factor beyond $\gtrsim 12$ dimensions. In such higher-dimensional regimes, Gaussian process-based Bayesian optimization becomes increasingly inefficient, and particle swarm optimization is therefore the more practical choice. The GP training time during these optimizations scales $\mathcal{O}(n^3)$ with the number of samples, n , and ranged from less than a second in early iterations to approximately 3 min in 6D cases and up to ~ 30 min in 8D cases after thousands of samples, representing a small fraction of the overall wall-clock time dominated by simulation execution.

The training of a surrogate model does, however, have advantages. By providing an estimate of the objective function across the full parameter space, the surrogate can be used to quantify the stability of the identified optimum, identify alternative optima, flag unphysical solutions, and assess parameter dominance. As an illustrative example, a stability analysis was performed for the ~ 0.7 MJ, three-pulse optimum. This analysis consists of sampling the objective function along each parameter axis through the optimal point and comparing these evaluations with the predictions of the Bayesian surrogate model, thereby assessing its local accuracy. The results of this analysis are shown in Fig. 6.

The surrogate reproduces the objective function well in the vicinity of the optimum, reflecting the high density of samples and the associated high confidence, accumulated in this region during optimization. Given the large volume of the parameter space, the surrogate is not expected to provide accurate predictions across the entire

domain; rather, its primary utility lies in resolving the local structure around the optimum, where sample density is greatest, and optimization decisions are made. Crucially, the analysis demonstrates that the optimum exhibits significant width along all parameter directions, indicating that the implosion performance is tolerant to realistic experimental variation. The dominant parameter, defined as the parameter with the smallest characteristic length scale, is x_2 , corresponding to the power of the final laser pulse. This behavior is physically intuitive, as this parameter controls the strength of the final and strongest shock driven through the target, making the implosion performance particularly sensitive to its value.

Minor discrepancies between the surrogate and sampled objective values are present, most notably along parameters x_4 and x_5 . Along x_4 , the GP prediction falls below zero. As the objective function is strictly non-negative by construction, this is a result of the unconstrained nature of the GP. Interpreting negative predictions as zero largely removes this discrepancy. Along x_5 , the deviation between the GP prediction and sampled values is more pronounced; however, the location and width of the maximum are captured accurately. Since the purpose of this analysis is to assess local stability and parameter sensitivity rather than exact agreement across the entire parameter space, the surrogate model is considered sufficiently accurate in the vicinity of the optimum.

A comprehensive summary of the results for each optimization strategy is provided in the [supplementary material](#), together with the corresponding optimization hyperparameters. This includes the final laser pulse profiles, low-CR constraint metrics, and gain. These tables enable direct, quantitative comparison between the different optimization approaches and clearly illustrate the improvements achieved using Bayesian and particle swarm optimization.

A key observation that can be made from these results is that the highest-energy implosions (~ 6.5 and ~ 8.5 MJ) deviate from expectations based on simple hydrodynamic scaling (hydroscaling)³⁹ of the smaller capsules. Hydroscaling provides a set of scaling relations that allow a capsule and laser drive configuration to be translated between energy scales and is commonly used to estimate performance on larger laser facilities. However, the optimized pulse shapes for the largest capsules require substantially lower peak laser power (and, to a lesser extent, total laser energy) than would be predicted by directly scaling the smaller configurations.

To investigate this behavior, test simulations were performed in which the four-pulse, ~ 1.8 MJ configuration identified by the Bayesian optimizer (the highest-performing implosion at this energy scale) was hydroscaled to match the outer radius of the ~ 8 MJ capsule. While the resulting implosion remains viable, achieving an objective function value of 57 compared to 84 for the directly optimized configuration, it is clearly suboptimal. The procedure was then reversed by hydroscaling the optimized ~ 8 MJ configuration down to the radius of the ~ 1.8 MJ capsule. In this case, the degradation is more pronounced, with the objective function decreasing from 2.13 to 0.14. These results indicate that, although hydroscaling can generate reasonable configurations, an optimized solution at one scale cannot be hydroscaled to form an optimized solution at another scale.

Several factors likely contribute to this. First, the optimization objective function incorporates stability constraints and performance metrics that do not scale trivially with the configuration scale. Second, ignition-relevant burn physics introduces nonlinear sensitivity to




FIG. 6. Stability analysis of the Bayesian-optimized ~ 0.7 MJ, three-pulse wetted-foam implosion. The objective function is sampled (red points) along each optimization parameter axis through the identified optimum, while the corresponding Bayesian surrogate model predictions are shown by the solid blue curves. The width of the optimum along each axis indicates a stable optimum, with the most sensitive parameter (smallest width) given by parameter x_2 , corresponding to the power of the final pulse.

pulse structure and timing. As a result, these observations suggest that while hydroscaling remains a useful guide for estimating performance trends, independent optimization at each energy scale is necessary to identify truly optimal configurations.

V. FAST IGNITION OPTIMIZATION

A. Fast ignition outline and previous optimization attempts

To demonstrate the versatility and broader applicability of advanced optimization techniques to ICF design, the particle swarm

optimization (PSO) framework was applied to the optimization of fast ignition (FI) implosions. In contrast to the low-CR regime, FI is characterized by different objective functions and constraints, providing a distinct optimization landscape. The dimensionality of the design space, discussed below, however, renders GP-based Bayesian optimization computationally implausible for this problem, inhibiting its use for this task.

FI is an alternative ICF approach in which the fuel compression and ignition phases are decoupled, offering the potential for higher gain and reduced driver energy requirements. In this scheme, a

primary driver compresses the fuel to high areal densities (ρR) after which a secondary, ultra-intense laser pulse generates a beam of energetic particles to rapidly heat the compressed fuel to sufficiently high temperatures to initiate fusion burn.

Achieving high areal density during the compression phase of a fast ignition implosion is critical for both ignition and subsequent burn. Large ρR values ensure that alpha particles produced in fusion reactions redeposit their energy within the fuel, enabling the self-heating required to sustain burn. This behavior underpins the burn-up fraction, defined as the fraction of fuel undergoing fusion reactions, which can be approximated as⁴⁰

$$\Phi = \frac{\rho R}{\rho R + 7}, \tag{8}$$

where ρR is in units of g/cm^2 . During the ignition phase, high ρR also enhances the collisionality of the fuel, allowing energetic particle beams to deposit energy more efficiently to ignite the fuel. High ρR is therefore essential for both initiating fusion and maximizing energy yield in fast ignition.

While high ρR is a necessary condition for effective FI performance, it is not sufficient on its own. Additional constraints must be placed on the density of the compressed fuel. If the density is too high, energetic particle beams are unable to penetrate the fuel, depositing energy predominantly at the periphery of the compressed region. Conversely, if the density is too low, the energy deposition volume becomes too large, preventing the formation of temperatures sufficient for ignition. Consequently, the compressed fuel density is typically constrained to the range $300\text{--}450\text{ g}/\text{cm}^3$.^{18,41–44}

The optimal density within this window depends on the available short-pulse energy. For fixed ρR , increasing the compressed fuel density reduces the characteristic size of the assembly and therefore decreases the total compressed fuel mass, which scales as $M \propto \rho^{-2}$. Since the energy required for ignition is expected to scale approximately with the mass of fuel that must be heated to ignition conditions, the ignition energy requirement similarly decreases with increasing density ($E \propto \rho^{-1.85}$).⁴⁵ The optimal compressed density is, therefore, facility-dependent and reflects a trade-off between achievable compression, ignition energy requirements, and beam penetration considerations.

A number of promising fast ignition implosion designs have been explored previously, including the self-similar configurations proposed by Clark and Tabak.⁴⁶ While these results are compelling, the optimized laser pulse shapes reported exceed the parametric limit of $I\lambda^2 \leq 10^{14}\text{ W}\mu\text{m}^2/\text{cm}^2$. As discussed in the low-CR context, exceeding this limit increases susceptibility to parametric instabilities and reduces the reliability of one-dimensional radiation-hydrodynamic simulations. Similar issues apply to the optimization study presented in Ref. 18 where a particle swarm approach was used to refine self-similar designs within a highly restricted parameter space.

B. Our optimization approach

To address these limitations, the PSO framework presented here is applied to the fast ignition design landscape. The optimization presented here considers only the compression phase of the fast ignition scheme. As fast ignition inherently separates the compression and ignition stages, these phases can be investigated independently. The

simulations are therefore restricted to the one-dimensional compression of the fuel assembly and do not include modeling of the short-pulse ignition beam, fast electron transport, or coupling to the fuel.

The ignition phase of fast ignition involves strongly kinetic and multidimensional processes, including non-local electron transport and beam-plasma interactions, which are not captured within a one-dimensional radiation-hydrodynamic framework. Rather than introducing simplified or inconsistent models of these processes, the present study focuses on the compression phase, where the employed modeling approach remains appropriate. While thermonuclear burn is included, the configurations do not reach ignition conditions, and the resulting fusion yield is negligible.

The capsule configuration is fixed to that proposed by Clark and Tabak,⁴⁶ as shown in Fig. 7, while the laser pulse profile is parameterized using eight time-power pairs between which the pulse is linearly interpolated. This parameterization permits the optimizer to explore a vast volume in configuration space, and does not bias the optimization process to any preferred pulse shape. An upper bound on laser intensity is set by the constraint $I\lambda^2 \leq 10^{14}\text{ W}\mu\text{m}^2/\text{cm}^2$, while the lower bound is chosen to be three orders of magnitude smaller. In addition to the laser intensity constraint, a constraint is also applied to the total laser energy. If the laser energy of a proposed pulse is above 485 kJ, then the whole pulse power profile is reduced so that the pulse energy is 485 kJ. This energy limit fits anticipated FI laser system sizes of around 1 MJ energy⁴⁷ being split between both long pulse compression and short pulse ignition laser pulses.

The objective function is comprised of two components: the compressed fuel areal density, ρR , and a density suitability metric designed to penalize density values outside of the $300\text{--}450\text{ g}/\text{cm}^3$ range. To identify the compressed fuel region, a shock-finding metric defined in Ref. 18 is employed to locate the outgoing shock, which coincides with the outer boundary of the compressed fuel. A shock is characterized by a discontinuous change in material properties such as pressure and density, and its location can be robustly identified by combining the local pressure and density gradients. Evaluating these gradients on a logarithmic scale further enhances the sensitivity to the shock boundary. The shock metric used to identify the index of the




FIG. 7. Schematic of the fast ignition capsule design based on Ref. 46. The ablator, DT ice, and DT vapor regions are shown with their corresponding densities and initial radii.

15 June 2026 11:29:54




FIG. 8. Optimized FI density (solid black) and temperature (light gray) (a) and laser pulse (b) profiles obtained from the PSO process. The density distribution exhibits a $\sim 20 \mu\text{m}$ hotspot followed by a density profile within the suitable density range of $300\text{--}450 \text{ g/cm}^3$, achieving $\rho R = 1.53 \text{ g/cm}^2$.

outgoing shock, and hence the outer radius of the compressed fuel region, is defined as

$$M = \max \left(\left| \frac{\ln(P_{i+1}) - \ln(P_i)}{r_{i+1} - r_i} \right|, \left| \frac{\ln(\rho_{i+1}) - \ln(\rho_i)}{r_{i+1} - r_i} \right| \cdot m_i \right), \quad (9)$$

where P is the pressure, ρ is the density, m is the mass, and subscript i indicates the cell index. The inclusion of the cell mass biases the metric toward larger radii, ensuring the outgoing shock is identified rather than internal shocks. Using this shock index, the compressed fuel ρR can be calculated at each time step.

A density suitability multiplier is also evaluated to assess whether the compressed fuel lies within the desired density window. For each cell within the compressed region, a target density $\rho_{*,i}$ is assigned. If the cell density lies within the ideal range of $300\text{--}450 \text{ g/cm}^3$, the target density is set equal to the cell density and no penalty is incurred. For densities below 300 or above 450 g/cm^3 , the target density is set to the nearest bound. The suitability penalty, S , is then calculated using a target-density- and cell-width-normalized mean square error

$$\text{MSE}(M) = \sum_{i=0}^M \left(\frac{\rho_i - \rho_{*,i}}{\rho_{*,i}} \right)^2 \frac{(r_{i+1} - r_i)}{r_M}, \quad S = 1 - \sqrt{\text{MSE}(M)}, \quad (10)$$

where r_M denotes the outer radius of the compressed region. The suitability metric is bounded between 0 and 1 (assuming $\rho \leq 900 \text{ g/cm}^3$), with $S = 1$ corresponding to all compressed fuel lying within the target density range. The identification of the compressed fuel radius, r_M (vertical dotted line) and the target density window (horizontal dotted lines) is illustrated in Fig. 8(a).

The final objective function combines the compressed fuel areal density with the density suitability metric. To prioritize physically viable compressed density profiles over artificially high ρR values, the suitability metric is squared, biasing the optimization toward configurations with favorable density distributions

$$\text{Result} = \rho R \cdot S^2. \quad (11)$$

Without this weighting, highly unsuitable density profiles can appear deceptively attractive due to excessively large ρR values alone.

C. Results

A swarm consisting of 999 particles was used to explore the 16D parameter space defined by the eight time-power pairs describing the

laser pulse. This is the largest swarm size permitted on the HPC cluster, and was chosen to ensure broad coverage of the high-dimensional design space. The full swarm configuration and hyperparameter choices are provided in the [supplementary material](#).

The optimized laser pulse profile and the resulting compressed fuel density distribution are shown in Fig. 8. In the density profile, the outer shock radius, $r_M = 63.7 \mu\text{m}$, identified using the outgoing shock metric [Eq. (9)], is indicated by the vertical dotted line. The target density window for fast ignition ($300\text{--}450 \text{ g/cm}^3$) is indicated by the horizontal dotted lines.

The innermost $\sim 20 \mu\text{m}$ of the compressed fuel lies below the target density threshold of 300 g/cm^3 . This behavior is difficult to avoid under the imposed constraint on the maximum allowable laser intensity. As the capsule is compressed, the central pressure resisting the implosion increases, and once this pressure exceeds the applied drive pressure, the implosion stagnates and subsequently rebounds. The imposed limit on laser intensity, therefore, constrains the maximum achievable implosion pressure, making it challenging to further increase the central fuel density.

Beyond the central $\sim 20 \mu\text{m}$, the remaining compressed fuel lies predominantly within the desired density range, with densities clustered around 300 g/cm^3 . This region constitutes the majority of the compressed fuel mass, accounting for 98.6% of the total compressed fuel mass. The resulting compressed areal density is $\rho R = 1.53 \text{ g/cm}^2$, approaching the 2 g/cm^2 threshold commonly cited^{48,49} as a requirement for high-gain fast ignition designs. If successfully ignited, Eq. (8) implies that this compressed areal density, together with a total compressed mass of $3.2 \times 10^{-4} \text{ g}$, would correspond to a burn-up fraction of $\sim 18\%$ and a fusion energy release of $\sim 20 \text{ MJ}$. Assuming the compressed assembly can be ignited using a total laser energy budget of 1 MJ (includes both long and short pulse lasers), this would imply a gain of ~ 20 . It should be emphasized that the optimization presented here concerns the compression phase only. The generation, transport, divergence, and coupling efficiency of the short-pulse ignition beam are not modeled. This gain estimate, therefore, represents that an implied performance assuming successful ignition is possible with 500 kJ of short-pulse energy. Promising routes to further improve this result would be to include the capsule design into the optimization parameter space, increase the available energy budget, or shift to a higher frequency laser pulse, such as ArF⁵⁰ or KrF.⁵¹

The optimized laser power profile used to drive the implosion is generally an increasing function, reaching a maximum laser intensity of $1.6 \times 10^{14} \text{ W/cm}^2$, corresponding to $I\lambda^2 = 8 \times 10^{13} \text{ W } \mu\text{m}^2/\text{cm}^2$, and utilizes the full available energy budget of 485 kJ. While the pulse




FIG. 9. Modified optimized laser pulse, with (a) the relaxation feature removed, and (b) a smoothed laser pulse. Both modifications lead to reduced performance.

shape is necessarily coarse due to the parameterization using only eight time-power pairs, several physically meaningful features can be identified. For much of the drive, the laser power increases gradually, corresponding to ramped compression in which the applied pressure is steadily increased. At ~ 13 ns, the pulse deviates from this trend with a negative power gradient, allowing the plasma to partially relax before being driven again. Since the drive pressure is closely related to the power gradient, this relaxation enables a larger pressure to be applied than would otherwise be possible. This behavior is qualitatively similar to shock-augmented ignition schemes, in which intentional reductions in laser power are used to facilitate the formation of stronger shocks at later times.

Two additional simulations were performed to analyze the physical importance of the precise structure of the optimized laser pulse. The laser pulses used in these simulations are shown in Fig. 9. In the first case, the reduction in laser power associated with releasing the plasma was removed, resulting in a monotonically increasing drive. In the second case, the pulse was smoothed using spline interpolation through the original time-power pairs, which preserves the structure of the original pulse.

In both cases, implosion performance was degraded relative to the optimized pulse. This sensitivity demonstrates that the relaxation feature identified by the optimizer plays a crucial role in the implosion dynamics. In removing or smoothing this feature, the compression of the fuel is hindered. The fact that this feature is not biased or imposed *a priori*, but instead emerges from the optimization, highlights the ability of the optimization process to identify non-intuitive yet physically meaningful drive strategies.

To assess the convergence behavior of the optimization, the location of the optimum in parameter space is again defined as the iteration at which the objective function first reaches 90% of its final value. Using this definition, the swarm identified the optimum within 14 iterations, corresponding to approximately 13 000 simulations. Each fast ignition simulation required approximately 5 h of wall-clock time, resulting in a total computational cost of ~ 65 000 processor-hours for the full optimization. Given the dimensionality of the parameter space and the complexity of the fast ignition design landscape, this demonstrates that the PSO framework is able to efficiently locate physically meaningful optima within a computationally feasible number of evaluations.

VI. CONCLUSION

This work has presented the development and application of Bayesian optimization and particle swarm optimization frameworks

specifically adapted for high-cost, physics-driven optimization problems in inertial confinement fusion. A suite of physics-motivated extensions was introduced to improve robustness, efficiency, and scalability, enabling effective optimization in noisy, multi-modal design spaces subject to stringent physical constraints.

Re-optimization of low-convergence-ratio (low-CR) wetted-foam implosions demonstrated that both approaches significantly outperform traditional sequential scan methods, achieving improved gain while reducing the total number of simulations required. Both advanced optimization approaches achieved comparable final optima, with Bayesian optimization requiring fewer evaluations, and particle swarm optimization converging more rapidly in terms of iteration count and wall-clock time. These results highlight a clear trade-off between sample efficiency and convergence speed, indicating that the optimal choice of optimizer depends on the available computational resources and the specific optimization context.

The particle swarm optimization framework was further applied to a substantially higher-dimensional fast ignition design landscape. In this regime, PSO proved capable of efficiently identifying a physically viable design that achieves a high compressed areal density of 1.5 g/cm². This high areal density was achieved using a laser pulse profile under strict intensity and energy constraints.

While Gaussian process-based Bayesian optimization becomes increasingly computationally expensive in highly dimensional parameter spaces, particle swarm optimization remains well suited to such high-dimensional problems. The complementary strengths of these two approaches suggest that no single optimization strategy is universally optimal for ICF design, but rather that different regimes benefit from different algorithmic choices.

Looking forward, these optimization frameworks provide a foundation for more sophisticated design studies, including capsule design, hybrid optimization strategies, and experimental optimization on high-power laser facilities. As ICF designs continue to grow in complexity and computational cost, physics-aware optimization methods such as those presented here will play an increasingly important role in navigating complex design landscapes and accelerating progress toward high-performance fusion implosions.

SUPPLEMENTARY MATERIAL

See the [supplementary material](#) for additional details supporting the optimization studies presented in this work. This includes the functional form of the constraint multipliers, optimization hyperparameter choices, and summary tables of the final optimized low-CR configurations for all capsule sizes and pulse shapes. It also includes

details of the fast ignition optimization, including swarm hyperparameters and the evolution of the optimization process.

ACKNOWLEDGMENTS

This work has been funded by the EPSRC HEC grant under Grant No. EP/X035336/1, the AWAKE2 Grant No. ST/X005518/1, the JAI Grant No. ST/V001655/1, the Oxford-Living Optics and the Oxford-IBM Computational Discovery grants, and the Clarendon Fund Scholarship at the University of Oxford. The authors thank all of the staff of the Central Laser Facility and Scientific Computing Department of UKRI Rutherford Appleton Laboratory for the use of the SCARF supercomputing facility in support of this research.

AUTHOR DECLARATIONS

Conflict of Interest

The authors have no conflicts to disclose.

Author Contributions

J. J. Lee: Conceptualization (lead); Data curation (lead); Formal analysis (lead); Investigation (lead); Methodology (lead); Supervision (equal); Writing – original draft (lead); Writing – review & editing (lead). **D. Coope:** Data curation (equal); Methodology (equal); Writing – review & editing (equal). **J. Redfern:** Data curation (equal); Methodology (equal); Writing – review & editing (equal). **H. Martin:** Writing – review & editing (equal). **R. Timmis:** Writing – review & editing (equal). **E. Denis:** Writing – review & editing (equal). **A. James:** Writing – review & editing (equal). **R. T. Ruskov:** Writing – review & editing (equal). **Z. Zhang:** Writing – review & editing (equal). **P. Norreys:** Conceptualization (supporting); Supervision (equal); Writing – original draft (supporting); Writing – review & editing (equal). **R. W. Paddock:** Conceptualization (supporting); Writing – review & editing (equal).

DATA AVAILABILITY

The data that support the findings of this study are available from the corresponding author upon reasonable request.

REFERENCES

- J. Nuckolls, L. Wood, A. Thiessen, and G. Zimmerman, *Nature* **239**, 139–142 (1972).
- O. A. Hurricane, P. K. Patel, R. Betti, D. H. Froula, S. P. Regan, S. A. Slutz, M. R. Gomez, and M. A. Sweeney, *Rev. Mod. Phys.* **95**, 025005 (2023).
- R. Betti and O. Hurricane, *Nat. Phys.* **12**, 435 (2016).
- J. Biener, C. Dawedait, S. Kim, T. Braun, M. Worsley, A. Chernov, C. Walton, T. Willey, S. Kucheyev, S. Shin, Y. Wang, M. Biener, J. Lee, B. Koziolowski, T. van Buuren, K. Wu, J. Satcher, and A. Hamza, *Nucl. Fusion* **52**, 062001 (2012).
- R. E. Olson, R. J. Leeper, S. A. Yi, J. L. Kline, A. B. Zylstra, R. R. Peterson, R. Shah, T. Braun, J. Biener, B. J. Koziolowski, J. D. Sater, M. M. Biener, A. V. Hamza, A. Nikroo, L. B. Hopkins, D. Ho, S. LePape, and N. B. Meezan, *J. Phys.: Conf. Ser.* **717**, 012042 (2016).
- R. E. Olson, R. J. Leeper, J. L. Kline, A. B. Zylstra, S. A. Yi, J. Biener, T. Braun, B. J. Koziolowski, J. D. Sater, P. A. Bradley, R. R. Peterson, B. M. Haines, L. Yin, L. F. Berzak Hopkins, N. B. Meezan, C. Walters, M. M. Biener, C. Kong, J. W. Crippen, G. A. Kyrala, R. C. Shah, H. W. Herrmann, D. C. Wilson, A. V. Hamza, A. Nikroo, and S. H. Batha, *Phys. Rev. Lett.* **117**, 245001 (2016).
- R. Olson, C. Thomas, B. Haines, and M. Schmitt, *Nucl. Fusion* **66**, 026002 (2026).
- R. E. Olson, M. J. Schmitt, B. M. Haines, G. E. Kemp, C. B. Yeaman, B. E. Blue, D. W. Schmidt, A. Haid, M. Farrell, P. A. Bradley, H. F. Robey, and R. J. Leeper, *Phys. Plasmas* **28**, 122704 (2021).
- B. M. Haines, S. A. Yi, R. E. Olson, S. F. Khan, G. A. Kyrala, A. B. Zylstra, P. A. Bradley, R. R. Peterson, J. L. Kline, R. J. Leeper, and R. C. Shah, *Phys. Plasmas* **24**, 072709 (2017).
- A. B. Zylstra, S. A. Yi, B. M. Haines, R. E. Olson, R. J. Leeper, T. Braun, J. Biener, J. L. Kline, S. H. Batha, L. Berzak Hopkins, S. Bhandarkar, P. A. Bradley, J. Crippen, M. Farrell, D. Fittinghoff, H. W. Herrmann, H. Huang, S. Khan, C. Kong, B. J. Koziolowski, G. A. Kyrala, T. Ma, N. B. Meezan, F. Merrill, A. Nikroo, R. R. Peterson, N. Rice, J. D. Sater, R. C. Shah, M. Stadermann, P. Volegov, C. Walters, and D. C. Wilson, *Phys. Plasmas* **25**, 056304 (2018).
- S. Atzeni and J. Meyer-Ter-Vehn, *The Physics of Inertial Fusion: Beam Plasma Interaction, Hydrodynamics, Hot Dense Matter* (Oxford University Press, 2004). https://academic.oup.com/book/0/chapter/198103983/chapter-ag-pdf/44602688/book_27812_section_198103983.ag.pdf.
- C. Walters, E. Alger, S. Bhandarkar, K. Boehm, T. Braun, F. Espinosaloza, B. Haid, R. Heredia, J. Kline, B. Koziolowski, J. Kroll, D. Malone, A. Nikroo, P. Opsahl, J. Sater, and A. Zylstra, *Fusion Sci. Technol.* **73**, 305 (2018).
- M. Tabak, J. Hammer, M. E. Glinsky, W. L. Kruer, S. C. Wilks, J. Woodworth, E. M. Campbell, M. D. Perry, and R. J. Mason, *Phys. Plasmas* **1**, 1626 (1994).
- A. Caruso and V. Pais, *Nucl. Fusion* **36**, 745 (1996).
- M. Roth, T. E. Cowan, M. H. Key, S. P. Hatchett, C. Brown, W. Fountain, J. Johnson, D. M. Pennington, R. A. Snavely, S. C. Wilks, K. Yasuike, H. Ruhl, F. Pegoraro, S. V. Bulanov, E. M. Campbell, M. D. Perry, and H. Powell, *Phys. Rev. Lett.* **86**, 436 (2001).
- J. M. ter Vehn, *Plasma Phys. Controlled Fusion* **43**, A113 (2001).
- R. W. Paddock, H. Martin, R. T. Ruskov, R. H. H. Scott, W. Garbett, B. M. Haines, A. B. Zylstra, R. Aboushelbaya, M. W. Mayr, B. T. Spiers, R. H. W. Wang, and P. A. Norreys, *Philos. Trans. R. Soc. A* **379**, 20200224 (2021).
- M. Brönnner, S. Atzeni, D. Callahan, J. Gaffney, P. Gibbon, L. C. Jarrott, A. Mateo, L. Savino, N. Schott, W. Theobald, and M. Roth, *Phys. Plasmas* **32**, 022710 (2025).
- V. Gopalaswamy, A. Lees, R. Ejaz, C. A. Thomas, T. J. B. Collins, K. S. Anderson, W. Ebmeyer, and R. Betti, *Phys. Rev. Res.* **7**, 013009 (2025).
- A. J. Crilly, P. W. Moloney, D. Shi, and E. A. Ferdinandi, *Phys. Plasmas* **33**, 012702 (2026).
- P. W. Hatfield, S. J. Rose, and R. H. H. Scott, *Phys. Plasmas* **26**, 062706 (2019).
- J. A. Delettrez, T. J. B. Collins, and C. Ye, *Phys. Plasmas* **26**, 062705 (2019).
- N. Hansen, “The CMA evolution strategy: A comparing review,” in *Towards a New Evolutionary Computation: Advances in the Estimation of Distribution Algorithms*, edited by J. A. Lozano, P. Larrañaga, I. Inza, and E. Bengoetxea (Springer, Berlin/Heidelberg, 2006), pp. 75–102.
- T. Larsen and S. M. Lane, *J. Quant. Spectrosc. Radiat. Transfer* **51**, 179 (1994).
- B. Shahriari, K. Swersky, Z. Wang, R. P. Adams, and N. de Freitas, *Proc. IEEE* **104**, 148 (2016).
- R. Garnett, *Bayesian Optimization* (Cambridge University Press, 2023).
- C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (The MIT Press, 2005).
- C. E. Rasmussen and H. Nickisch, *J. Mach. Learn. Res.* **11**, 3011 (2010), see <https://www.semanticscholar.org/paper/A-Diversity-Guided-Particle-Swarm-Optimizer-the-Riget-Vesterstrom/e91e20f33c37bc370fd266172440cd508688ad63>.
- A. D. Palma, C. Mandler-Dünner, T. Parnell, A. Anghel, and H. Pozidis, “Sampling acquisition functions for batch Bayesian optimization,” *arXiv:1903.09434* (2019).
- J. Snoek, H. Larochelle, and R. P. Adams, *Advances in Neural Information Processing Systems*, edited by F. Pereira, C. Burges, L. Bottou, and K. Weinberger (Curran Associates, Inc., 2012), Vol. 25.
- D. Eriksson, M. Pearce, J. R. Gardner, R. Turner, and M. Poloczek, “Scalable global optimization via local Bayesian optimization,” *arXiv:1910.01739* (2020).
- C. E. Rasmussen, “Gaussian processes in machine learning,” in *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia, February 2–14, 2003, Tübingen, Germany, August 4–16, 2003, Revised Lectures*, edited by O. Bousquet, U. von Luxburg, and G. Rätsch (Springer, Berlin/Heidelberg, 2004), pp. 63–71.

- ³³J. Kennedy and R. Eberhart, in *Proceedings of ICNN'95—International Conference on Neural Networks* (IEEE, 1995), Vol. 4, pp. 1942–1948.
- ³⁴A. Ratnaweera, S. Halgamuge, and H. Watson, *IEEE Trans. Evol. Comput.* **8**, 240 (2004).
- ³⁵W. Du and B. Li, *Inf. Sci.* **178**, 3096 (2008).
- ³⁶J. Riget and J. Vesterstrom, A diversity-guided particle swarm optimizer - the ARPSO (2002), available at <https://www.semanticscholar.org/paper/A-Diversity-Guided-Particle-Swarm-Optimizer-the-Riget-Vesterstrom/e91e20f33c37bc370fd266172440cd508688ad63#citing-papers>.
- ³⁷X.-F. Xie, W.-J. Zhang, and Z.-L. Yang, “A dissipative particle swarm optimization,” [arXiv:cs/0505065](https://arxiv.org/abs/cs/0505065) (2005).
- ³⁸R. H. H. Scott, D. Barlow, W. Trickey, A. Ruocco, K. Glize, L. Antonelli, M. Khan, and N. C. Woolsey, *Phys. Rev. Lett.* **129**, 195001 (2022).
- ³⁹K. L. Baker, O. Jones, C. Weber, D. Clark, P. K. Patel, C. A. Thomas, O. L. Landen, R. Nora, G. J. Anderson, J. Gaffney, S. MacLaren, D. T. Casey, T. Döppner, E. L. Dewald, R. Tommasini, B. K. Spears, J. Salmonson, M. Hohenberger, S. Khan, A. Zylstra, A. Kritcher, P. Amendt, V. Smalyuk, J. Lindl, C. Young, J. S. Ross, D. Ho, O. A. Hurricane, D. A. Callahan, T. Woods, J. L. Milovich, D. J. Strozzi, B. Bachmann, R. Bionta, P. M. Celliers, D. Fittinghoff, R. Hatarik, M. Gatun Johnson, K. Meaney, M. Millot, P. L. Volegov, and C. Wilde, *Phys. Plasmas* **29**, 062705 (2022).
- ⁴⁰S. Atzeni and M. Tabak, *Plasma Phys. Controlled Fusion* **47**, B769 (2005).
- ⁴¹M. H. Key, K. Akli, F. Beg, R. Betti, D. S. Clark, S. N. Chen, R. R. Freeman, S. Hansen, S. P. Hatchett, D. Hey, J. A. King, A. J. Kemp, B. F. Lasinski, B. Langdon, T. Ma, A. J. MacKinnon, D. Meyerhofer, P. K. Patel, J. Pasley, T. Phillips, R. B. Stephens, C. Stoeckl, M. Foord, M. Tabak, W. Theobald, M. Storm, R. P. J. Town, S. C. Wilks, L. VanWoerkom, M. S. Wei, R. Weber, and B. Zhang, *J. Phys.: Conf. Ser.* **112**, 022056 (2008).
- ⁴²S. Atzeni, M. Temporal, and J. Honrubia, *Nucl. Fusion* **42**, L1 (2002).
- ⁴³S. Atzeni, A. Schiavi, J. J. Honrubia, X. Ribeyre, G. Schurtz, P. Nicolai, M. Olazabal-Loumé, C. Bellei, R. G. Evans, and J. R. Davies, *Phys. Plasmas* **15**, 056311 (2008).
- ⁴⁴R. Betti, A. A. Solodov, J. A. Delettrez, and C. Zhou, *Phys. Plasmas* **13**, 100703 (2006).
- ⁴⁵S. Atzeni, *Phys. Plasmas* **6**, 3316 (1999).
- ⁴⁶D. Clark and M. Tabak, *Nucl. Fusion* **47**, 1147 (2007).
- ⁴⁷J. Nuckolls, *J. Phys.: Conf. Ser.* **244**, 012007 (2010).
- ⁴⁸T. Johzaki, K. Mima, and Y. Nakao, *Plasma Fusion Res.* **2**, 041 (2007).
- ⁴⁹H. Nagatomo, T. Johzaki, R. Takizawa, and S. Fujioka, *Nucl. Fusion* **64**, 106019 (2024).
- ⁵⁰S. P. Obenschain, A. J. Schmitt, J. W. Bates, M. F. Wolford, M. C. Myers, M. W. McGeoch, M. Karasik, and J. L. Weaver, *Philos. Trans. R. Soc. A* **378**, 20200031 (2020).
- ⁵¹S. Obenschain, R. Lehberg, D. Kehne, F. Hegeler, M. Wolford, J. Sethian, J. Weaver, and M. Karasik, *Appl. Opt.* **54**, F103 (2015).