

The blind implosion-maker: Automated inertial confinement fusion experiment design

Cite as: Phys. Plasmas **26**, 062706 (2019); <https://doi.org/10.1063/1.5091985>

Submitted: 07 February 2019 . Accepted: 16 May 2019 . Published Online: 07 June 2019

P. W. Hatfield , S. J. Rose , and R. H. H. Scott 



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[On alpha-particle transport in inertial fusion](#)

Phys. Plasmas **26**, 062701 (2019); <https://doi.org/10.1063/1.5101074>

[Approaching a burning plasma on the NIF](#)

Phys. Plasmas **26**, 052704 (2019); <https://doi.org/10.1063/1.5087256>

[Making inertial confinement fusion models more predictive](#)

Phys. Plasmas **26**, 082704 (2019); <https://doi.org/10.1063/1.5108667>



NEW: TOPIC ALERTS

Explore the latest discoveries in your field of research

[SIGN UP TODAY!](#)

The blind implosion-maker: Automated inertial confinement fusion experiment design

Cite as: Phys. Plasmas **26**, 062706 (2019); doi: [10.1063/1.5091985](https://doi.org/10.1063/1.5091985)

Submitted: 7 February 2019 · Accepted: 16 May 2019 ·

Published Online: 7 June 2019




View Online



Export Citation



CrossMark

P. W. Hatfield,^{1,a)}  S. J. Rose,^{1,2}  and R. H. H. Scott³ 

AFFILIATIONS

¹Clarendon Laboratory, University of Oxford, Parks Road, Oxford OX1 3PU, United Kingdom

²Blackett Laboratory, Imperial College, London SW7 2AZ, United Kingdom

³Central Laser Facility, STFC Rutherford Appleton Laboratory, Harwell Oxford, Didcot OX11 0QX, United Kingdom

^{a)}Electronic mail: peter.hatfield@physics.ox.ac.uk

ABSTRACT

The design of inertial confinement fusion (ICF) experiments, alongside improving the development of energy density physics theory and experimental methods, is one of the key challenges in the quest for nuclear fusion as a viable energy source [O. A. Hurricane, J. Phys.: Conf. Ser. **717**, 012005 (2016)]. Recent challenges in achieving a high-yield implosion at the National Ignition Facility (NIF) have led to new interest in considering a much wider design parameter space than normally studied [J. L. Peterson *et al.*, Phys. Plasmas **24**, 032702 (2017)]. Here, we report an algorithmic approach that can produce reasonable ICF designs with minimal assumptions. In particular, we use the genetic algorithm metaheuristic, in which “populations” of implosions are simulated, the design of the capsule is described by a “genome,” natural selection removes poor designs, high quality designs are “mated” with each other based on their yield, and designs undergo “mutations” to introduce new ideas. We show that it takes $\sim 5 \times 10^4$ simulations for the algorithm to find an original NIF design. We also link this method to other parts of the design process and look toward a completely automated ICF experiment design process—changing ICF from an experiment design problem to an algorithm design problem.

Published under license by AIP Publishing. <https://doi.org/10.1063/1.5091985>

I. INTRODUCTION

Achieving controlled nuclear fusion burn in the laboratory is a key goal in the pathway to nuclear fusion as an industrial power source. One of the main potential pathways to this goal is inertial confinement fusion (ICF), in which deuterium–tritium fuel is compressed to extremely high temperatures and pressures very quickly. The world’s leading ICF facility is NIF, in California, USA, at Lawrence Livermore National Laboratory, although ICF is a worldwide endeavor with facilities around the world contributing toward this goal.

When considering different designs, normally the intention is to maximize the neutron yield. In this sense, ICF can be viewed as an optimization problem:³ we wish to maximize yield within the constraints of the experimental setup (e.g., within what experiments could feasibly be fielded on NIF). The design space of ICF experiments is very large, making exploration of the design space nontrivial. Designs have historically been found by human imagination combined with an understanding of physical principles. However, the inability of conventional designs to reach ignition, combined with the explosion of applications of algorithmic and machine learning techniques in the physics

community at large, has led to an increase in the use of computer algorithms and machine learning approaches to identify new designs.^{2,4,5}

In this work, we consider a much larger parameter space than hitherto and seek to design an ICF experiment completely “from scratch” (as opposed to within some predefined smaller parameter space). In particular, this work is motivated by the “The Surprising Creativity of Digital Evolution”⁶—which summarizes a variety of unexpected results from a range of fields, where algorithmic approaches produced completely unanticipated designs. For ICF at NIF, it would be of great interest to find if there are any “unexpected” design features that are yet to be found that could help reach ignition; conversely, if such features do not exist, we would like to increase our confidence that no such alternative designs do exist.

II. GENETIC ALGORITHMS

We consider the performance of the class of metaheuristics known as *genetic algorithms*⁷ (GAs), which are motivated by Darwinian evolution and have already been applied to some problems in fusion and laser physics.^{8,9} One particularly relevant experimental

campaign to this analysis is an effort to optimize x-ray production from an argon cluster jet at the Gemini laser facility.¹⁰ They found that they were able to double the x-ray production from their experiments with a genetic algorithm—with comparatively little knowledge of the underlying physics.

In genetic algorithms, the design in question is characterized by a “genome.” A “population” of designs is simulated; the highly performing ones are given higher chances of “mating” with each other (crossover), receive “mutations,” and produce a new generation. We use a crossover method that respects the capsule (and drive) structure, see Fig. 1. This process is then repeated indefinitely, producing better and better designs. Genetic algorithms are typically most appropriate when little is known about the parameter space—different algorithms may be better if more is known about the problem. The “No Free Lunch” theorem¹¹ suggests that a perfect “black-box” optimizer likely cannot exist, so in general an appropriate algorithm for the situation must be chosen, and performance between algorithms may need to be tested empirically.

Our implosion simulations are performed using the Hyades¹² radiation-hydrodynamics simulation code, which is well benchmarked and used widely for the simulation of inertial fusion and high energy density physics applications.^{13–16} Hyades models hydrodynamics within a Lagrangian framework. Electron and ion thermal energy transport is described by a flux-limited Spitzer–Härm thermal conductivity model. Equations-of-state either use the Los Alamos SESAME tables¹⁷ or QEOS.¹⁸ Ionization levels come from a hydrogenic average-atom model or self-consistently from QEOS. Radiation transport uses

the multigroup diffusion approximation; here, we use 60 groups. A 1D spherically symmetric geometry is employed.

Neither laser-plasma interactions nor hohlraum physics are modeled; instead, we use an incoming x-ray drive imposed at the outside of the grid. The capsules are modeled within a 5 mm helium container. Simulations start in cryogenic conditions at $1.551 \times 10^{-3} \text{ eV} = 18 \text{ K}$. The Hyades runs were performed on SCARF at the Central Laser Facility at Rutherford Appleton Laboratory using 600 CPUs.

The design space is chosen to roughly correspond to what is achievable with a gold hohlraum on NIF. We consider capsule designs made of 5 layers. Each layer can be DT gas, DT ice, or plastic (CH). If the layer is DT gas, its density may be between 10 mg/cc and 200 mg/cc. Note that this lower bound is substantially higher than the DT gas density in the NIF point design, so we would not expect it to be possible to achieve the higher yields possible within other design spaces. This constraint is motivated by a desire to avoid implosions with extremely high convergence ratios; partially because 1D simulations are likely to be more realistic for low convergence ratio implosions and partially because there is some suggestion that the reliance on very high convergence ratios contributes to the difficulty in achieving ignition. DT ice and CH densities are fixed at 255 mg/cc and 1044 mg/cc, respectively. The thickness of each layer may be between 0.01 mm and 1.1 mm. The total radius of the capsule is constrained to be less than 1.1 mm. For the x-ray drive, a pure Planckian spectrum (no M-band) is used, and its time dependence is modeled by combining 15 Gaussians, each of which is parameterized by three numbers (peak temperature, time of peak, and width). The drive is constrained




FIG. 1. The Genetic Algorithm Schematic of the key components of the genetic algorithm as applied to the capsule designs. A similar process is applied to the drive designs.

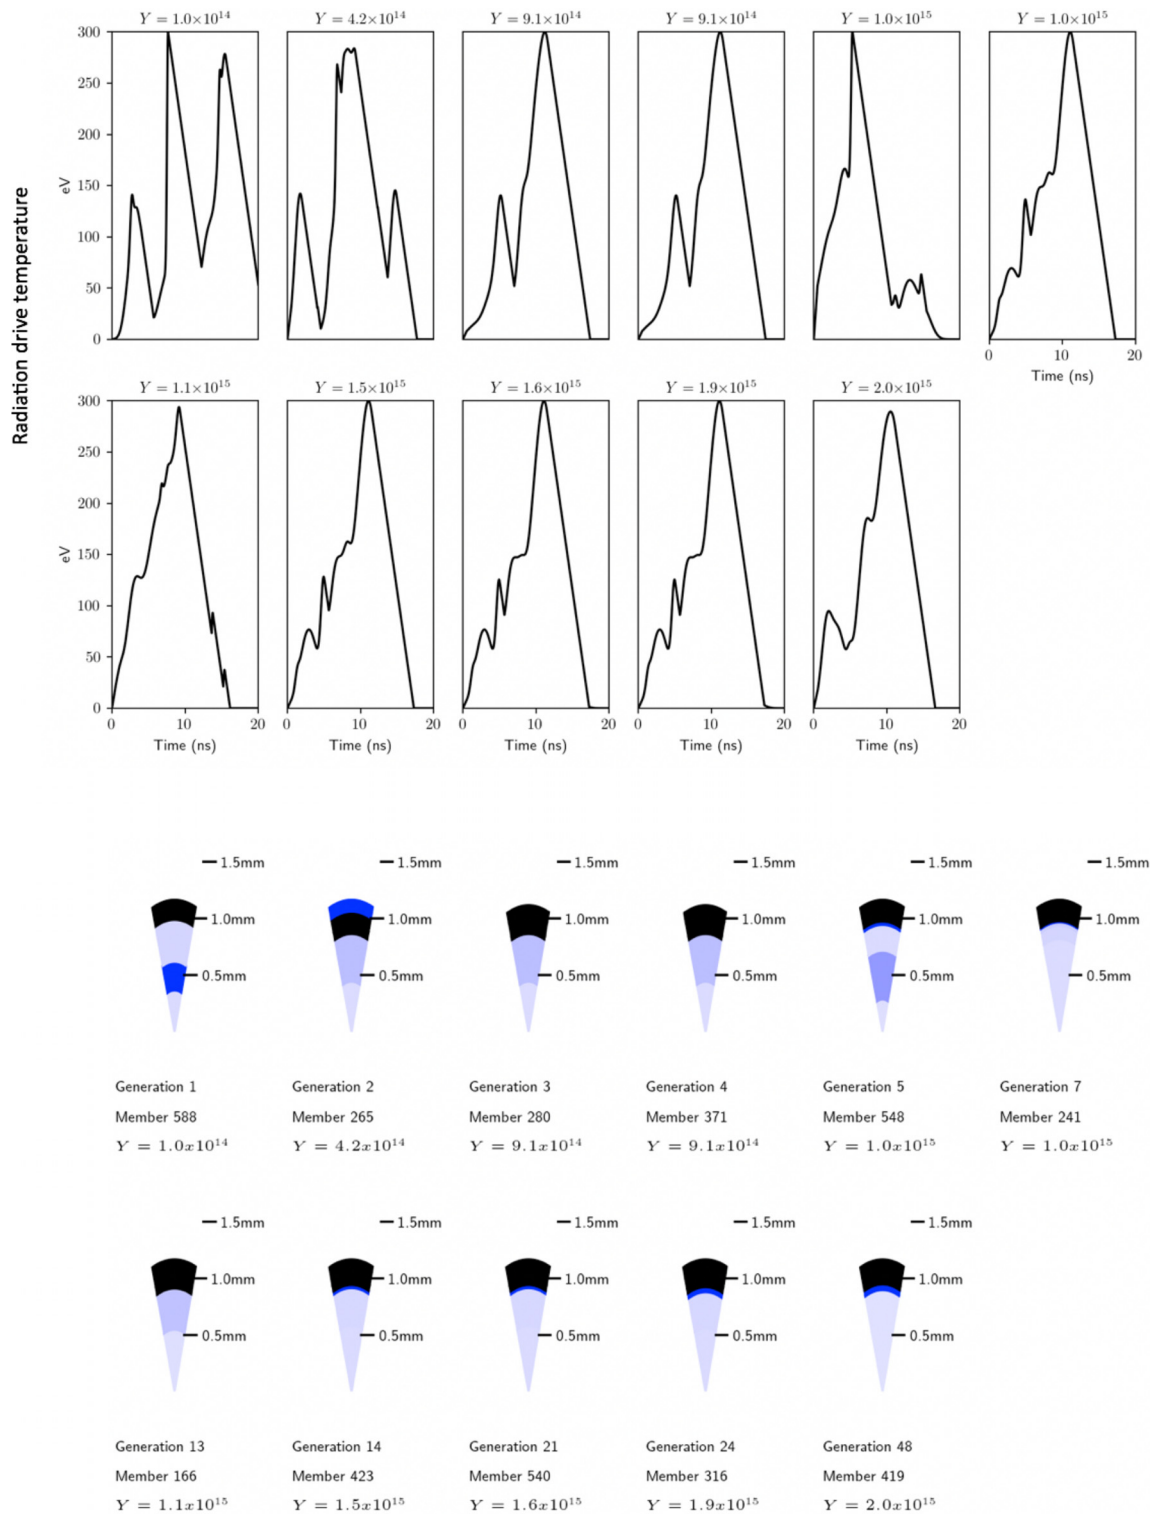


FIG. 2. Survival of the fittest—designs over time. The top panel shows capsule designs and the bottom panel shows the corresponding drive designs. Designs are only shown when a new generation improves over the previous generation. Pale blue represents DT gas (with darker corresponding to higher densities), royal blue to DT ice, and black to CH.




FIG. 3. Yield improvement with a number of generations. The highest yield achieved as a function of population generation.

to never go above 300 eV and to not deliver more than 300 kJ to a 1 mm radius sphere, and its cooling rate is constrained to be below 50 eV/ns, all of which are considered realistic constraints for capsule implosion in a hohlraum at NIF. Rigorous modeling of hohlraum physics to determine precisely what drives are possible is beyond the scope of this text, but these rough constraints are consistent with other experiments fielded at the facility.^{1,19,20} This is a very large parameter space: of order 60 discrete or continuous parameters.

A population size of 600 was run for 80 generations: 48 000 Hyades simulations in total, taking ~ 27 days in total to run, and finding a recognizable design. The evolution in drive and capsule structure leading up to the final design is shown in Fig. 2 and the corresponding increase in yield in Fig. 3. The time-radius profile of the final implosion is shown in Fig. 4. The algorithm was run three times (each time with different seeds), and the best performing final design had a total neutron yield of $\sim 2.1 \times 10^{15}$ neutrons (~ 5 kJ). This yield is achieved with a modest convergence ratio. We note that in future implementations, the design need not be optimized over yield. For example, instead of being treated as a constraint, the total energy in the drive could be minimized, with the constraint of $Y > 10^{18}$ (say) imposed.

In terms of algorithmic convergence, each run did not converge to exactly the same design, see Fig. 5. All three achieve a factor of ten improvement very rapidly in the first ~ 5 – 10 generations and then get a further 50%–100% improvement in the subsequent 70 generations much more slowly. Smaller mutations likely would make the initial rapid increase in yield with generation slower, and would increase the risk of being stuck in a local maximum, but would let the fine tuning of the design proceed more quickly at later stages. This shows that the




FIG. 4. Structure of the final implosion. The top panel of the left half of the figure shows the radiation drive for the final design as a function of time, the middle panel shows thermonuclear burn energy rate as a function of time, and the bottom panel shows a radius-time plot of the final implosion. Pale blue represents DT gas (with darker corresponding to higher densities), royal blue to DT ice, black to CH, and red to He. The right half of the figure shows the final capsule design.

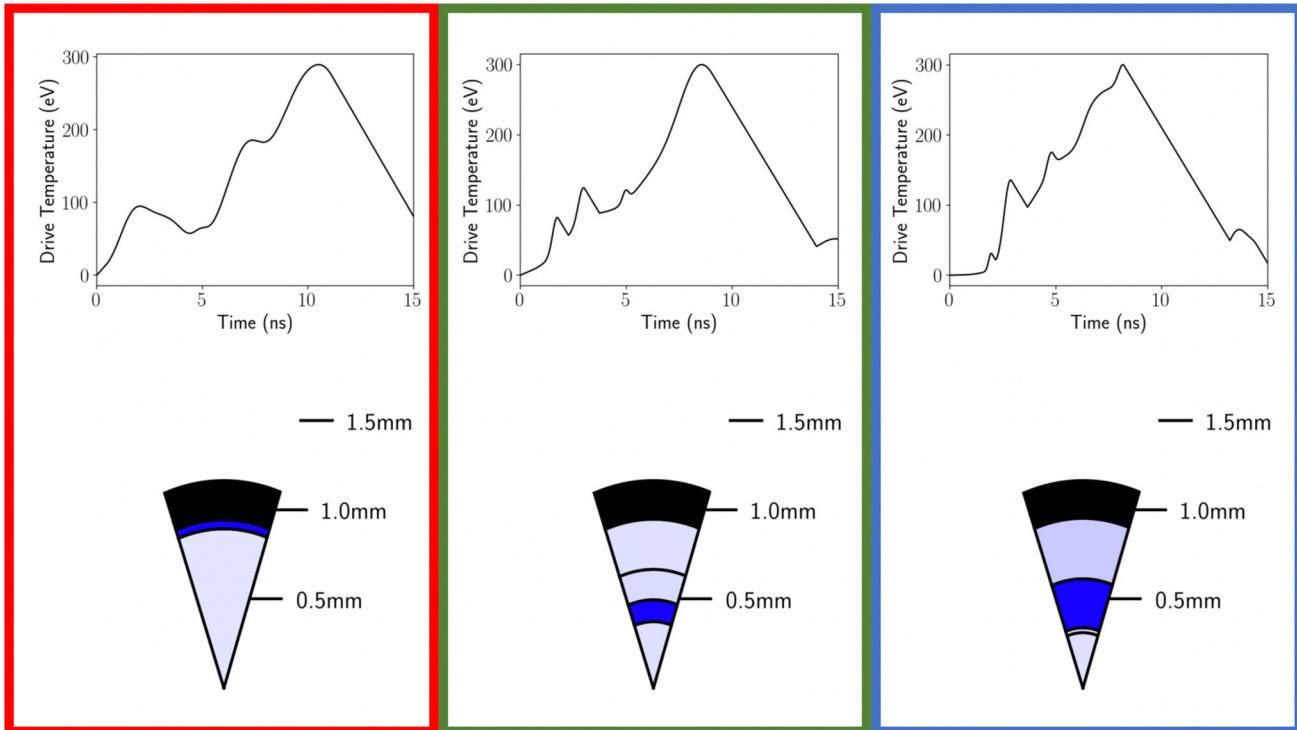


FIG. 5. Different Designs. The final capsule and drive designs of each of the runs of the algorithms, box colors match to Fig. 3.

GA, as implemented here, does not achieve the global maximum every time (which is likely very hard in such a large parameter space) but does consistently find plausible designs in the region of what is likely to be the global optimal. All three designs have similar ablaters, a similar peak to the drive with some preshocks, and a mix of DT gas and ice in the interior of the capsule (in fact, it is interesting that it looks like it may be possible to get comparable yields when the ice is not on the inner surface of the ablator). The final yields are within 25% of each other, reasonably similar considering the ~ 20 -fold improvement that is found from the first generation. It is also in general quite hard to judge when the algorithm has converged; for the “blue” run, it might have been tempting to stop the algorithm at around 40 generations. More fundamentally, however, finding the absolute global maximum is not the intention; GAs are known to be very slow in the final stages of convergence, but critically we would not expect the global maximum in this simplified 1D setup to be the maximum for real experiments. The intention is to find plausible outlines of designs in an acceptable amount of time that can then be further refined; see Sec. III. Large amounts of computational resources could be wasted trying to find the absolute maximum in 1D when it is likely that such a design would not be the experimental optimal.

III. AUTOMATING DESIGN

Although the design found here is worthy of greater analysis, the key result is that finding plausible ICF designs can be automated. Studies similar to the analysis presented here will likely (for the foreseeable future) only be able to do the large number of simulations presented here in 1D and with simplified physics. For this reason, we view

metaheuristics as one stage in the “Data Scientists’ Algorithm” approach to the ICF design. Figure 6 shows a schematic for this alternative approach where metaheuristics are used to generate basic designs.^{21,22} Human inspection removes designs that are clearly unphysical or impractical and reduces the dimensionality of the designs to make the subsequent steps tractable. These designs then form a “pool” of plausible designs alongside those that are produced by humans. Each of these designs is then studied in greater depth using machine learning regression based surrogates⁷ and designs that start to fail when more realistic physics is included can be removed (e.g., perhaps a design had high yield in 1D but completely fails in 2D). These surrogates can then be calibrated with more computationally expensive state-of-the-art 3D simulations (where it would be computationally unfeasible to do more than a few full calculations; in addition, the increase in fidelity need not be 1D to 3D, it could also be single to multigroup diffusion, etc.), and then calibrated again with real NIF shots.^{5,23,24} These experimentally calibrated surrogates give data-based improved predictions over theory without learning why theory and experiment disagree in a physical way. However, if new physical insight is learnt from the experiments, then it can be fed back into the simulations used in the process. The experiments used to calibrate the surrogate should be chosen using Bayesian experimental design/Bayesian optimization.^{25,26}

We give a brief example of how this proposed workflow might work. Our genetic algorithm metaheuristic earlier gave three designs. Human inspection would note that the “green” and “blue” designs are relatively similar, and the “red” design is slightly different. A human designer might in this case judge that the red design is more worthy of further testing, both as it had a higher yield, and that it had fewer

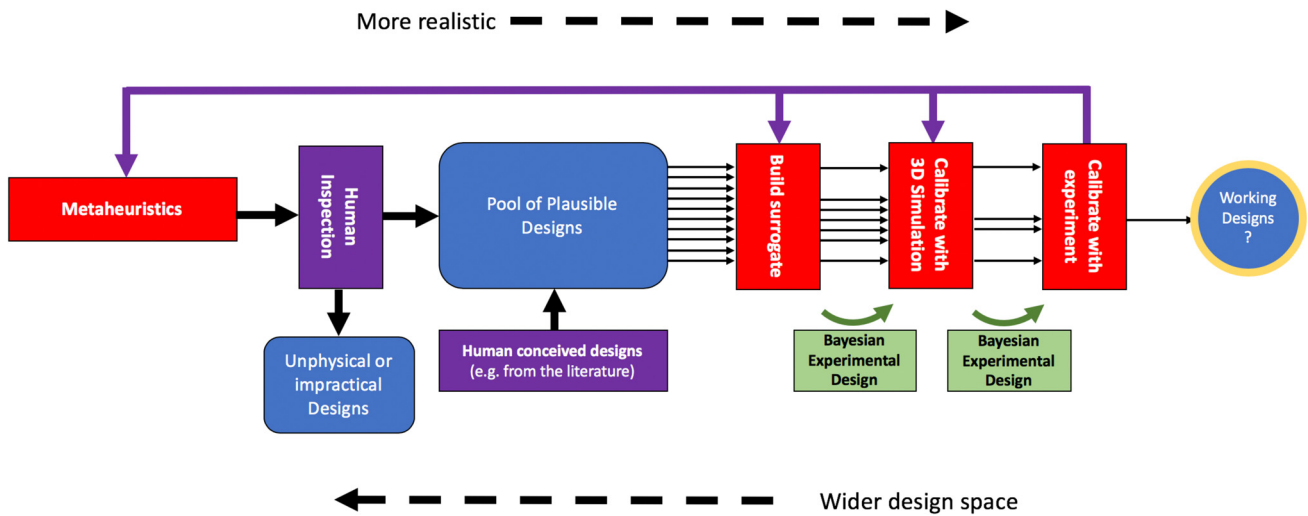


FIG. 6. Schematic of “The Data Scientists’ Algorithm.” Blue boxes show pools of designs, purple sections and lines show processes requiring human input, red boxes are algorithmic processes, green lines and boxes show decision making processes, and black lines show the flow of designs.

potential engineering difficulties than the green/blue design. Once the red design was selected, the human designer might then judge that varying the gas fill density and the ablator thickness were the most important parameters and start to apply a surrogate building methodology.² We demonstrate this by doing 1000 simulations over this 2D parameter space and letting gas fill density vary between 10 mg/cc and 30 mg/cc and ablator thickness vary between 0.1 mm and 0.3 mm (keeping total

radius and DT ice thickness fixed but also letting the DT gas radius vary in tandem). We then applied the machine learning code GPz^{27–29} to these data (100 data points in the training sample, 100 in the validation sample, and 800 in the testing sample); the results in Fig. 7 show that this modest training set is sufficient for this small parameter space. With the surrogate, one can (a) see that a slightly higher yield is indeed possible (an improvement from $\sim 2.1 \times 10^{15}$ to $\sim 2.5 \times 10^{15}$ neutrons) and (b) now with the ability to quickly predict yield in this parameter, subspace stability, etc., can easily be tested.² If the designer had then satisfied themselves that this was a sufficiently interesting design, they could then go forward with a transfer learning⁵ or alternative statistical approach²⁴ to incorporate higher fidelity simulations and finally experiments.

Finally, we note that although the algorithm discussed here is as “blind” as the process that produced the natural world, we speculate that one day it might be possible for the metaheuristic stage in this process to be replaced by a genuine artificial intelligence algorithm that really can imagine and conceive of new designs in the same way that the human brain can.

ACKNOWLEDGMENTS

We thank the referee for helpful comments which have improved the quality of this paper. P.W.H. and S.J.R. acknowledge funding from the Engineering and Physical Sciences Research Council. Many thanks to Warren Garbett at the Atomic Weapons Establishment, Aldermaston, Kris McGlinchey, and Jeremy Chittenden at Imperial College, London, and Jim Gaffney, Luc Peterson, and Kelli Humbird at Lawrence Livermore National Laboratory for useful discussions. The original idea for this article came from <http://boxcar2d.com>.

APPENDIX: GENETIC ALGORITHM DETAILS

The capsule designs in the starting population are seeded by, for each layer, sampling a material (DT gas, DT ice, or CH, each with equal probability), sampling a density to be used if the material




FIG. 7. Surrogate predictions. The graph shows the performance of our machine learning surrogate, the x-axis shows the Hyades yield, and the y-axis shows what the machine learning surrogate predicted (with uncertainties) for that point in design space. Light blue diagonal line shows a one-one relationship.

is DT gas (uniformly in log-space; the least informative prior for a scale parameter), and sampling a thickness (again uniformly in log-space). If the total radius exceeds the allowed amount, the design is scaled down to satisfy the size constraint. The drive designs (as a function of time) in the starting population are expressed as a sum of 15 Gaussians. Each Gaussian has its center in consecutive 1 ns periods of time and is characterized by three parameters, one for where in the nanosecond the peak occurs, one for the peak temperature, and one for the standard deviation for the Gaussian. The location of the peak is sampled uniformly and the component peak temperature and standard deviation are sampled uniformly in log-space. Finally, the drive temperature is then uniformly scaled down so that the peak temperature and total energy in the drive are within the prescribed limits and modified so the drive temperature does not decrease faster than permitted.

The selection method was fitness proportionate selection, with a fitness (f) function of $f = Y^\alpha$, with $\alpha = \frac{1}{2}$ chosen. Typically, a higher α corresponds to a faster rate of convergence at the cost of being more likely to get caught in local minima. Crossover is performed separately for the capsule and the drive, so that the crossover respects the structure of the problem and is neutral toward drive-capsule interactions. For the capsule, an integer i is randomly sampled between the zero and the number of sections (5 here) inclusive, and 1-point crossover is used at that point in the capsule design. Crossover for the drive is done in the same way with the Gaussian components. Mutations for the capsule are carried out by randomly selecting one of the sections, and with equal probabilities resampling either the material, thickness, or DT gas density in the same way as was used to create the starting population. A similar process is applied to the drive. Finally, we use “elitism;” the best in a population automatically survives unmodified to the next generation.

The heuristic described here is a simple implementation of a genetic algorithm to illustrate the concept that metaheuristics can produce plausible designs from extremely large parameter spaces. It is possible that performance can be improved by including more algorithmic features, e.g., reducing the mutation rate over time, or by using more complex evolutionary algorithms.^{30,31}

REFERENCES

- O. A. Hurricane, “Overview of progress and future prospects in indirect drive implosions on the national ignition facility,” *J. Phys.: Conf. Ser.* **717**, 012005 (2016).
- J. L. Peterson, K. D. Humbird, J. E. Field, S. T. Brandon, S. H. Langer, R. C. Nora, B. K. Spears, and P. T. Springer, “Zonal flow generation in inertial confinement fusion implosions,” *Phys. Plasmas* **24**, 032702 (2017).
- A. Giunta, M. Eldred, T. Trucano, and S. Wojtkiewicz, “Optimization under uncertainty methods for computational shock physics applications,” in *Proceedings of the 43rd AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference* (American Institute of Aeronautics and Astronautics, Reston, Virginia, 2002).
- E. A. Baltz, E. Trask, M. Binderbauer, M. Dikovsky, H. Gota, R. Mendoza, J. C. Platt, and P. F. Riley, “Achievement of sustained net plasma heating in a fusion experiment with the optimetrist algorithm,” *Sci. Rep.* **7**, 6425 (2017).
- K. D. Humbird, J. L. Peterson, and R. G. McClarren, “Transfer learning to model inertial confinement fusion experiments,” e-print [arXiv:1812.06055](https://arxiv.org/abs/1812.06055).
- J. Lehman, J. Clune, D. Misevic, C. Adami, J. Beaulieu, P. J. Bentley, S. Bernard, G. Beslon, D. M. Bryson, P. Chrabaszcz, N. Cheney, A. Cully, S. Doncieux, F. C. Dyer, K. O. Ellefsen, R. Feldt, S. Fischer, S. Forrest, A. Frénoy, C. Gagné, L. L. Goff, L. M. Grabowski, B. Hodjat, F. Hutter, L. Keller, C. Knibbe, P. Krcak, R. E. Lenski, H. Lipson, R. MacCurdy, C. Maestre, R. Miikkulainen, S. Mitri, D. E. Moriarty, J.-B. Mouret, A. Nguyen, C. Ofria, M. Parizeau, D. Parsons, R. T. Pennock, W. F. Punch, T. S. Ray, M. Schoenauer, E. Shulte, K. Sims, K. O. Stanley, F. Taddei, D. Tarapore, S. Thibault, W. Weimer, R. Watson, and J. Yosinski, “The surprising creativity of digital evolution: A collection of anecdotes from the evolutionary computation and artificial life research communities,” e-print [arXiv:1803.03453](https://arxiv.org/abs/1803.03453).
- J. H. Holland, *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence* (MIT Press, 1975), p. 211.
- W. Miner, P. Valanju, S. Hirshman, A. Brooks, and N. Pomphrey, “Use of a genetic algorithm for compact stellarator coil design,” *Nucl. Fusion* **41**, 1185–1195 (2001).
- Z.-H. He, B. Hou, V. Lebailly, J. Nees, K. Krushelnick, and A. Thomas, “Coherent control of plasma dynamics,” *Nat. Commun.* **6**, 7156 (2015).
- M. J. V. Streeter, S. J. D. Dann, J. D. E. Scott, C. D. Baird, C. D. Murphy, S. Eardley, R. A. Smith, S. Rozario, J.-N. Gruse, S. P. D. Mangles, Z. Najmudin, S. Tata, M. Krishnamurthy, S. V. Rahul, D. Hazra, P. Pourmoussavi, J. Osterhoff, J. Hah, N. Bourgeois, C. Thornton, C. D. Gregory, C. J. Hooker, O. Chekhlov, S. J. Hawkes, B. Parry, V. A. Marshall, Y. Tang, E. Springate, P. P. Rajeev, A. G. R. Thomas, and D. R. Symes, “Temporal feedback control of high-intensity laser pulses to optimize ultrafast heating of atomic clusters,” *Appl. Phys. Lett.* **112**, 244101 (2018).
- D. Wolpert and W. Macready, “No free lunch theorems for optimization,” *IEEE Trans. Evol. Comput.* **1**, 67–82 (1997).
- J. T. Larsen and S. M. Lane, “HYADES—A plasma hydrodynamics code for dense plasma studies,” *J. Quant. Spectrosc. Radiat. Transfer* **51**, 179–186 (1994).
- K. A. Keilty, E. P. Liang, T. Ditmire, B. A. Remington, K. Shigemori, and A. M. Rubenchik, “Modeling of laser-generated radiative blast waves,” *Astrophys. J.* **538**, 645–652 (2000).
- D. Leibbrandt, R. Drake, and J. Stone, “Zeus-2D simulations of laser-driven radiative shock experiments,” in *High Energy Density Laboratory Astrophysics* (Springer-Verlag, Berlin/Heidelberg, 2005), pp. 273–276.
- K. L. Sequoia, M. S. Tillack, and H. A. Scott, “A comparison of Hyades and Cretin for modeling laser absorption in underdense plasmas,” UCSD Center for Energy Research internal report December 2006.
- R. H. H. Scott, F. Perez, J. J. Santos, C. P. Ridgers, J. R. Davies, K. L. Lancaster, S. D. Baton, P. Nicolai, R. M. G. M. Trines, A. R. Bell, S. Hulin, M. Tzoufras, S. J. Rose, and P. A. Norreys, “A study of fast electron energy transport in relativistically intense laser-plasma interactions with large density scalelengths,” *Phys. Plasmas* **19**, 053104 (2012).
- S. X. Hu, B. Militzer, V. N. Goncharov, and S. Skupsky, “FPEOS: A first-principles equation of state table of deuterium for inertial confinement fusion applications,” *Phys. Rev. B* **84**(22), 224109 (2011).
- R. M. More, K. H. Warren, D. A. Young, and G. B. Zimmerman, “A new quotidian equation of state (QEOS) for hot dense matter,” *Phys. Fluids* **31**, 3059 (1988).
- O. A. Hurricane, D. A. Callahan, D. T. Casey, E. L. Dewald, T. R. Dittrich, T. Döppner, M. A. Barrios García, D. E. Hinkel, L. F. Berzak Hopkins, P. Kervin, J. L. Kline, S. L. Pape, T. Ma, A. G. MacPhee, J. L. Milovich, J. Moody, A. E. Pak, P. K. Patel, H.-S. Park, B. A. Remington, H. F. Robey, J. D. Salmonson, P. T. Springer, R. Tommasini, L. R. Benedetti, J. A. Caggiano, P. Celliers, C. Cerjan, R. Dylla-Spears, D. Edgell, M. J. Edwards, D. Fittinghoff, G. P. Grim, N. Guler, N. Izumi, J. A. Frenje, M. Gatu Johnson, S. Haan, R. Hatarik, H. Herrmann, S. Khan, J. Knauer, B. J. Koziowski, A. L. Kritcher, G. Kyrala, S. A. Maclaren, F. E. Merrill, P. Michel, J. Ralph, J. S. Ross, J. R. Rygg, M. B. Schneider, B. K. Spears, K. Widmann, and C. B. Yeaman, “The high-foot implosion campaign on the national ignition facility,” *Phys. Plasmas* **21**, 056314 (2014).
- Y. Ping, V. A. Smalyuk, P. Amendt, R. Tommasini, J. E. Field, S. Khan, D. Bennett, E. Dewald, F. Graziani, S. Johnson, O. L. Landen, A. G. MacPhee, A. Nikroo, J. Pino, S. Prisbrey, J. Ralph, R. Seugling, D. Strozzi, R. E. Tipton, Y. M. Wang, E. Loomis, E. Merritt, and D. Montgomery, “Enhanced energy coupling for indirectly driven inertial confinement fusion,” *Nat. Phys.* **15**, 138–141 (2019).
- P. Kittara, A. Jiralucksanawong, G. Yassin, S. Wangsuya, and J. Leech, “The design of potter horns for THz applications using a genetic algorithm,” *Int. J. Infrared Millimeter Waves* **28**, 1103–1114 (2007).

- ²²A. Hector, A. V. Trifonov, E. Tong, P. Grimes, and G. Yassin, "A smooth walled four pixel feed horn array operating at 1.4 THz," in *28th International Symposium on Space Terahertz Technology*, Cologne, Germany, 13–15 March 2017.
- ²³J. Gaffney *et al.*, "Making ICF models more predictive: Combining simulations, experiments and expert knowledge using machine learning and Bayesian statistics," in *60th Annual Meeting of the APS Division of Plasma Physics* (APS, Portland, OR, 2018).
- ²⁴V. 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. Deletrez, 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. Mannon, A. Maximov, F. J. Marshall, D. T. Michel, S. Miller, S. F. B. Morse, J. Palastro, J. Peebles, P. B. Radha, S. P. Regan, S. Sampat, T. C. Sangster, A. B. Sefkow, W. Seka, R. C. Shah, W. T. Shmyada, A. Shvydky, C. Stoeckl, A. A. Solodov, W. Theobald, J. D. Zuegel, M. G. Johnson, R. D. Petrasso, C. K. Li, and J. A. Frenje, "Tripled yield in direct-drive laser fusion through statistical modelling," *Nature* **565**, 581–586 (2019).
- ²⁵A. Marco, F. Berkenkamp, P. Hennig, A. P. Schoellig, A. Krause, S. Schaal, and S. Trimpe, "Virtual vs. real: Trading off simulations and physical experiments in reinforcement learning with Bayesian optimization," in *Proceedings of the 2017 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2017), pp. 1557–1563.
- ²⁶T. M. Dieb and K. Tsuda, "Machine learning-based experimental design in materials science," in *Nanoinformatics* (Springer, Singapore, 2018), pp. 65–74.
- ²⁷I. A. Almosallam, S. N. Lindsay, M. J. Jarvis, and S. J. Roberts, "A sparse Gaussian process framework for photometric redshift estimation," *Mon. Not. R. Astron. Soc.* **455**, 2387–2401 (2016).
- ²⁸I. A. Almosallam, M. J. Jarvis, and S. J. Roberts, "GPz: Non-stationary sparse Gaussian processes for heteroscedastic uncertainty estimation in photometric redshifts," *Mon. Not. R. Astron. Soc.* **462**, 726–739 (2016).
- ²⁹P. Hatfield, S. Rose, R. Scott, I. Almosallam, S. Roberts, and M. Jarvis, "Using Gaussian Processes for predicting robust Inertial Confinement Fusion implosion yields," (2019).
- ³⁰Y. Borenstein and A. Moraglio, *Theory Principled Methods for the Design of Metaheuristics* (Springer, 2014); available at <https://link.springer.com/book/10.1007/978-3-642-33206-7>.
- ³¹N. Hansen and A. Ostermeier, "Completely derandomized self-adaptation in evolution strategies," *Evol. Comput.* **9**, 159–195 (2001).