ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL034

MACHINE LEARNING FOR LASER PULSE SHAPING

A.E. Pollard*, W. Okell, D.J. Dunning, E.W. Snedden
ASTeC, STFC Daresbury Laboratory, Sci-Tech Daresbury, Warrington, UK

Abstract

The temporal profile of the electron bunch is of critical importance in accelerator areas such as free-electron lasers and novel acceleration. In FELs, it strongly influences factors including efficiency and the profile of the photon pulse generated for user experiments, while in novel acceleration techniques it contributes to enhanced interaction of the witness beam with the driving electric field. Work is in progress at the CLARA facility at Daresbury Laboratory on temporal shaping of the ultraviolet photoinjector laser, using a fused-silica acousto-optic modulator (AOM). Generating a user-defined (programmable) time-domain target profile requires finding the corresponding spectral phase configuration of the shaper; this is a non-trivial problem for complex pulse shapes. Using a physically informed machine learning model, we demonstrate accurate and rapid shaping of the photo-injector laser to a wide range of arbitrary target temporal intensity profiles on the CLARA PI laser. Additionally, we discuss the utility of this expanded range of laser pulse shapes to potential applications in FELs and novel acceleration.

INTRODUCTION

In photoinjector systems, control over the longitudinal properties of the electron bunch can be achieved by temporal shaping of the laser pulse temporal profile [1]. Following the temporal shaping concept presented in [2], we developed an apparatus for temporally shaping the photoinjector laser pulses at CLARA, shown schematically in Fig. 1. The input laser pulse is spectrally dispersed by a transmission grating. A concave mirror one focal length away from the grating collimates the spectrum and focuses the laser pulse to a line focus, along which the laser wavelength varies approximately linearly. A fused-silica AOM is placed at the position of the focus, and a transducer driven with an RF waveform at 200 MHz central frequency generates an acoustic wave in the AOM, which propagates along the line focus of the laser. The laser pulses are diffracted from the induced refractive index modulation, and the spectral components are recombined using a second concave mirror and transmission grating.

To shape the laser pulse temporally, the spectral phase can be adjusted by varying the temporal phase of the acoustic wave via the temporal phase of the RF drive wave. The laser pulses can also be shaped temporally by varying the temporal amplitude of the acoustic wave; however, as this approach is lossy, all shaping must be carried out using only the phase. In order to produce a specific target pulse temporal intensity profile, we must find a suitable spectral phase mask to apply to the shaper. This is non-trivial for arbitrary

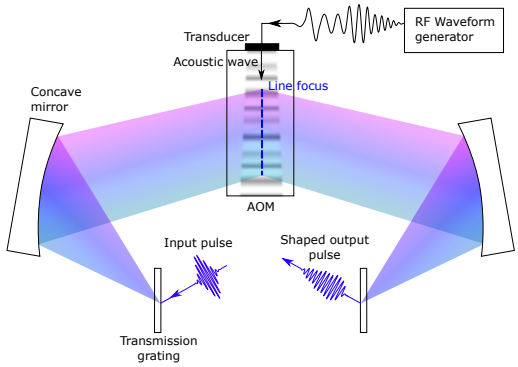

Figure 1: Schematic of the temporal pulse shaper at CLARA.

shapes, as we need both the phase and amplitude in either the spectral or temporal domain to fully define the pulse. We know only the temporal and spectral intensity, leaving the temporal and spectral phase as unknowns with many potential solutions. The complexity of real experimental systems poses additional challenges. For example, there are limitations imposed by the physical characteristics of the AOM. Modulating the spectral phase by modulating the temporal phase of the RF wave broadens the RF spectrum. The AOM has a finite acoustic bandwidth, and the RF spectrum must remain within this bandwidth for spectral phase modulations to be physically realisable.

Machine learning approaches excel for complex nonlinear problems such as this. In particular, deep neural networks are capable of approximating any function[3], and recent work has demonstrated that such networks can be used to learn and manipulate spectral, temporal, and shape properties of laser pulses[4][5]. Recent research has explored encoding physical laws into machine learning models with partial differential equations as priors[6] to reduce the data requirements of these otherwise data-intensive approaches. This approach is known as Physically Informed Neural Networks (PINNs) and can be used to constrain the output of deep neural networks to physical reality, by encoding properties such as conservation of energy.

In this paper, we present a PINN for finding the spectral phase mask required to produce a target temporal intensity profile in our photoinjector laser pulse shaper, subject to the physical limitation of the AOM bandwidth. Our approach both reduces the data requirements of our model and constrains the search space to a physically realisable range. Thus, we can be confident that the temporal intensity profiles produced by our model will be experimentally achievable.

In addition to the constraints of the physical system, we can also encode the inherent symmetries of the underlying physical system. By employing the principles of geometric deep learning[7], we can exploit those symmetries to reduce

* amelia.pollard@stfc.ac.uk

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL034

the complexity of the underlying parameter space. In particular, we can exploit the translation-invariant nature of the temporal pulse shape, since the structure of the pulse signal is our primary concern and pulse timing can be adjusted without consequence. This allows us to consider alternative loss functions, such as the Pearson correlation coefficient, rather than the mean squared error functions of other works, and significantly improve our results by expanding the potential search space.

APPLICATIONS TO NOVEL ACCELERATION

Arbitrary laser pulse shaping enables several technologies, which we consider here. Future work will likely see these methods applied to CLARA.

The FEBE arc system[8] contains a mask array for longitudinal shaping of the beam. This enables the generation of a variety of longitudinal bunch distributions such as drive/main bunch-pairs with variable delay, a single ultrashort (fs) low charge bunch, and a drive bunch with a train of witness bunches. Having the PI laser pulse 'precondition' the electron beam will make the FEBE mask more efficient by allowing more charge transport through the mask.

Echo-Enabled Harmonic Generation (EEHG) is used in FELs for the production of high-intensity, narrow bandwidth X-ray photons. A low-energy laser pulse is used to induce a sinusoidal energy modulation in the bunch which is then sent through a chicane to convert the energy modulation into a density modulation. This is then repeated with a second laser pulse to produce microbunching suitable for high coherence lasing in an undulator. Fluctuations in laser power can significantly affect performance, and present a significant obstacle in the EEHG regime[9]. However we may be able to use this property with the temporal laser pulse shaping for the production of multi-colour EEHG by carefully crafting the pulse to modulate different regions of the bunch by different amounts so that each region ends up microbunched at a different harmonic of the seed laser.

For laser-driven wakefield acceleration, the driving laser pulse shape has significant impact[10][11]. Optimisation of laser-plasma interactions for laser-driven particle acceleration could therefore be an interesting future direction for research. However, as AOM shaping is not possible at extremely high energies due to the energetic limits of the optics, the pulse would likely need to be shaped while in a low energy regime.

METHODOLOGY

Using simulated data, we developed and tested a machine learning model to find the phase mask required to achieve a particular target pulse temporal profile. The simulated laser pulses used for training and testing the model have a spectral intensity with Gaussian shape in wavelength, central wavelength of 266 nm, and FWHM bandwidth of 1.5 nm. As pulses in the CLARA photoinjector laser system are temporally stretched in a grating stretcher before entering the shaper, our simulated unshaped pulses have 8×10^4 fs^2 of spectral phase applied. We generated 10^5 pairs of spectral phase profiles and corresponding temporal intensity profiles for the training set, and 10^3 pairs for the test set. Each pair consists of a spectral phase profile of 2283 samples over 5.33 nm and a temporal intensity profile of 271 samples over 12 ps. In order to improve our results on human-crafted targets, 25% of the dataset is generated by compositing a selection of shapes with randomised parameters, an example of which can be seen in Fig. 2.

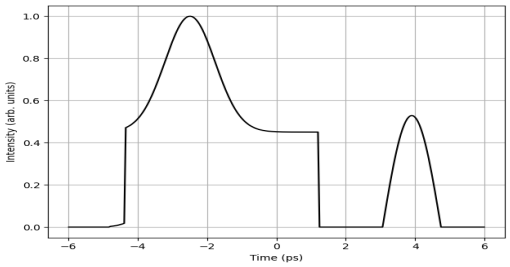

Figure 2: Example of training sample generated by compositing shapes.

To constrain our model to the physical limits of the AOM bandwidth, we consider the effect of AOM bandwidth on the spectral phase mask. Modulating the temporal phase of the RF drive wave broadens its spectrum; the instantaneous RF frequency at a particular point in time is given by the gradient of its temporal phase at that point. The AOM bandwidth limits the gradient of the acoustic wave temporal phase modulation, and consequently the limits on the available phase modulation per unit length along the acoustic wave propagation direction are

$$\frac{\mathrm{d}\varphi}{\mathrm{d}x} = \pm\frac{\pi\Delta f_{\mathrm{ac}}}{v_{\mathrm{ac}}}, \tag{1}$$

where φ is the phase modulation, x is the spatial coordinate across the AOM window, $v_{\mathrm{ac}} = 5968\,\mathrm{m\,s^{-1}}$ is the acoustic velocity in fused silica, and $\Delta f_{\mathrm{ac}} \approx 100\,\mathrm{MHz}$ is the AOM acoustic bandwidth. The change in laser wavelength per unit length across the AOM window is

$$\frac{\mathrm{d}\lambda}{\mathrm{d}x} \approx \frac{\Delta\lambda}{W}, \tag{2}$$

where $\Delta\lambda \approx 5\,\mathrm{nm}$ is the optical bandwidth covered by the AOM window, and $W = 20\,\mathrm{mm}$ is the width of the AOM window. From Eq. 1 and Eq. 2, the limits on the laser spectral phase gradient are therefore

$$g_\varphi = \frac{\mathrm{d}\varphi}{\mathrm{d}\lambda} \approx \pm\frac{\pi\Delta f_{\mathrm{ac}} W}{v_{\mathrm{ac}}\Delta\lambda}. \tag{3}$$

For our experimental parameters, $g_\varphi = \pm\pi/0.015$ rad/nm.

To encode this physical limit associated with the AOM bandwidth into the network, we developed a regulariser which acts to limit the gradient of the spectral phase profile to a physical limit of $\pi/0.015$ rad/nm, corresponding to a maximum phase change per wavelength step of $\delta\varphi \approx 0.153\pi$

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL034

rad/step. For the purposes of limiting the gradient over a discrete sampling, we define the discrete gradient as $\Delta_+(f) = f_i - f_{i+1}$. To account for the cyclic nature of angular frequency in a differentiable manner, we calculate the gradient by projecting into the complex plane. We then take the absolute value of the discrete gradient of the network's output vector. We multiply the resulting function by a high-gradient sigmoid function with an offset of $\delta\varphi$ to provide a differentiable approximation of a step function. Since $Re(e^{i\theta})$ bounds between -1 and 1, and the input is between $-\pi$ and π, we divide $\delta\varphi$ by π to arrive at the final definition for the regulariser, shown in Eq. 4. Note that the first and last elements are masked out from the difference calculation, since we are not concerned with forcing the spectral phase profile to begin and end at 0 rads.

The loss function encodes the translation invariant nature of temporal pulse shaping by calculating the Pearson correlation coefficient of the target temporal pulse profile against the temporal pulse profile simulated from the spectral phase profile output by the network. The simulation code is made differentiable by the Keras[12] framework, allowing the network training to be guided by the gradient of the underlying function space. This further encodes physical laws into the network.

$$\frac{1}{N}\sum |\Delta_+(e^{i\varphi(\omega)})| * \eta\sigma(|\Delta_+(e^{i\varphi(\omega)})| - \delta\varphi/\pi); \eta = 100 \quad (4)$$

The network architecture is a simple three layer ReLu[13] DNN, with batch normalisation between the layers. The final output layer uses a linear activation function. We use the Adam optimiser[14] with a learning rate schedule decaying from 0.001 at a factor of 0.8 per 50 epochs when the validation score plateaus.

RESULTS

By applying the principles of physically informed networks, we enable the model to learn to extrapolate appropriate spectral phase profiles, without non-physical phase transitions, for arbitrary temporal pulse shapes in linear time. This allows users to specify arbitrary temporal pulse profiles and receive an input for the AOM that provides that profile within milliseconds, a significant advantage over algorithmic and iterative methods.

As in other works, we calculated the mean squared error (MSE) of the output temporal intensity profile against those in the test set, and find strong agreement ($6.4e{-3} \pm 3.7e{-5}$ MSE over 10^3 samples). Indeed, our results match simulation extremely well, as shown in Fig. 3. We improved upon previous results for arbitrary shapes by introducing human-like data in the training set, as seen in Fig. 2.

Without the limitations of the SLM it is possible to achieve very high quality matches to the target patterns, however these are not physically achievable since they require spectral phase transitions well beyond what is physically possible. However, by imposing physical limitations through the PINN, we achieved high quality matches to arbitrary physically realisable temporal intensity profiles.

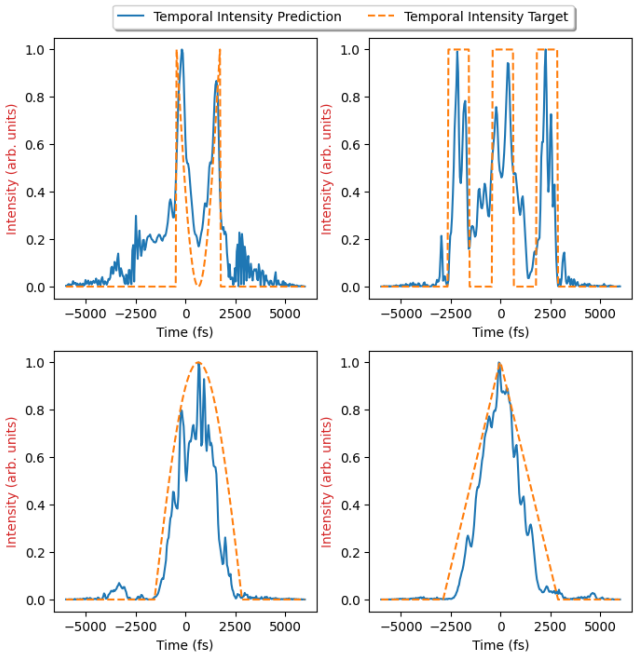

Figure 3: Demonstration of solutions found for arbitrary pulse shapes. Note that these are physically realisable due to the gradient constraint.

We developed a control system for the AOM that enables us to set spectral phase profiles and retrieve the resulting intensity profile at 10Hz. This allowed us to generate arbitrary spectral phase profiles that can be applied to the AOM directly from the ML model. After applying the spectral phase profiles to the AOM, we measured the resulting temporal intensity profile using a high-speed photodetector and compared the measured profile to the profile generated by the ML model. We observed that there were small discrepancies between the two profiles. To address this, we used the measured profile to fine-tune the ML model by performing a second round of training using the difference between the simulated and observed intensity profiles as the training signal. This resulted in a modest improvement to the overall matching between the ML model's output and the AOM's.

CONCLUSION

By using physically informed networks we can build better machine learning models which more accurately model the reality of the target system. In doing so, we develop a model for predicting spectral phase profiles for the photocathode laser at CLARA, enabling arbitrary specification of temporal intensity profiles for fine control over the bunch profile. We discussed potential applications of this capability to longitudinal masking, EEHG, and laser driven plasma wakefield accelerators and proposed these as future research directions.

ISBN: 978-3-95450-231-8 ISSN: 2673-5490 doi: 10.18429/JACoW-IPAC2023-THPL034

REFERENCES

[1] G. Penco *et al.*, "Experimental demonstration of electron longitudinal-phase-space linearization by shaping the photoinjector laser pulse," *Physical Review Letters*, vol. 112, p. 044 801, 2014. DOI: 10 . 1103 / PhysRevLett . 112 . 044801

[2] C. W. Hillegas, J. X. Tull, D. Goswami, D. Strickland, and W. S. Warren, "Femtosecond laser pulse shaping by use of microsecond radio-frequency pulses," *Optics Letters*, vol. 19, no. 10, pp. 737–739, 1994. DOI: 10.1364/ol.19.000737

[3] K. Hornik, M. Stinchcombe, and H. White, "Multilayer feedforward networks are universal approximators," *Neural networks*, vol. 2, no. 5, pp. 359–366, 1989. DOI: 10.1016/0893-6080(89)90020-8

[4] S. Boscolo and C. Finot, "Artificial neural networks for nonlinear pulse shaping in optical fibers," *Optics & Laser Technology*, vol. 131, p. 106 439, 2020. DOI: 10.1016/j. optlastec.2020.106439

[5] C. Xu *et al.*, "Machine learning based spatial light modulator control for the photoinjector laser at flute," in *12th Int. Particle Accelerator Conf.(IPAC'21), Campinas, Brazil*, 2021. DOI: 10.18429/JACoW-IPAC2021-WEPAB289

[6] M. Raissi, P. Perdikaris, and G. E. Karniadakis, "Physics informed deep learning (part i): Data-driven solutions of nonlinear partial differential equations. arxiv 2017," *arXiv preprint arXiv:1711.10561* DOI: 10.48550/arXiv.1711.10561

[7] M. M. Bronstein, J. Bruna, Y. LeCun, A. Szlam, and P. Vandergheynst, "Geometric deep learning: Going beyond euclidean data," *IEEE Signal Processing Magazine*, vol. 34, no. 4, pp. 18–42, 2017. DOI: 10.1109/MSP.2017.2693418

[8] A. Bainbridge, D. Angal-Kalinin, J. Jones, T. Pacey, Y. Saveliev, E. Snedden, *et al.*, "The design of the full energy beam exploitation (febe) beamline on clara," in *Proceedings of the International Particle Accelerator Conference*, 2022.

[9] D. Xiang and G. Stupakov, "Echo-enabled harmonic generation free electron laser," *Physical Review Special Topics-Accelerators and Beams*, vol. 12, no. 3, p. 030 702, 2009.

[10] C. Schroeder *et al.*, "Frequency chirp and pulse shape effects in self-modulated laser wakefield accelerators," *Physics of Plasmas*, vol. 10, no. 5, pp. 2039–2046, 2003.

[11] W. Leemans *et al.*, "Electron-yield enhancement in a laser-wakefield accelerator driven by asymmetric laser pulses," *Physical review letters*, vol. 89, no. 17, p. 174 802, 2002.

[12] F. Chollet *et al.* "Keras." (2015), https://github.com/fchollet/keras

[13] A. F. Agarap, "Deep learning using rectified linear units (relu)," *arXiv preprint arXiv:1803.08375*, 2018. DOI: 10 . 48550/arXiv.1803.08375

[14] D. Kingma and J. Ba, "Adam: A method for stochastic optimization," *International Conference on Learning Representations*, Dec. 2014. DOI: 10.48550/arXiv.1412.6980