Attosecond Chemistry

Simulating macroscopic high-order harmonic generation driven by structured laser beams using artificial intelligence ☆

José Miguel Pablos-Marín, Javier Serrano *, Carlos Hernández-García

Grupo de Investigación en Aplicaciones del Láser y Fotónica, Departamento de Física Aplicada, Universidad de Salamanca, Pl. Merced s/n, E-37008 Salamanca, Spain

A R T I C L E   I N F O

A B S T R A C T

Artificial intelligence, and in particular deep learning, is becoming a powerful tool to access complex simulations in intense ultrafast laser science. One of the most challenging tasks to model strong-field physics, and in particular, high-order harmonic generation (HHG), is to accurately describe the microscopic quantum picture—that takes place at the sub-nanometer/attosecond spatiotemporal scales—together with the macroscopic one—at the millimeter/femtosecond scales—to reproduce experimental conditions. The exact description would require to couple the laser-driven wavepacket dynamics given by the three-dimensional time-dependent Schrödinger equation (3D-TDSE) with the Maxwell equations, to account for propagation. However, such simulations are beyond the state-of-the-art computational capabilities, and approximations are required. Here we introduce the use of artificial intelligence to compute macroscopic HHG simulations where the single-atom wavepacket dynamics are described by the 3D-TDSE. We use neural networks to infer the 3D-TDSE microscopic HHG response, which is coupled with the exact solution of the integral Maxwell equations to account for harmonic phase-matching. This method is especially suited to compute macroscopic HHG driven by structured laser beams carrying orbital angular momentum within minutes or even seconds. Our work introduces an alternative and fast route to accurately compute extreme-ultraviolet/x-ray attosecond pulse generation.

1. Introduction

More than thirty years after its discovery [1,2], high harmonic generation (HHG) stands as one of the richest processes in strong-field physics. Its extraordinary coherence allows to up-convert laser light from the infrared (IR)/mid-infrared domains into the extreme-ultraviolet (EUV) or even soft x-rays [3], that can be synthesized into laser pulses as short as tens of attoseconds [4]. This unique source has enabled unprecedented studies of time-resolved electronic excitations [5] and molecular photoionization [6–8], attosecond spectroscopy in solids [9], observation of ultrafast magnetization dynamics [10], or x-ray diffractive imaging [11], among many others [12–14].

In the standard scheme of HHG, high-order harmonics are emitted upon the highly nonlinear interaction of an intense IR laser pulse with an atomic or molecular gas target. A realistic description of the HHG process, comparable to that present in the ex-

periments, requires its illustration from both the microscopic and the macroscopic points of view. At the microscopic level, the non-perturbative interaction of an intense IR laser pulse with an atom can be explained semiclassically through the so-called three-step model [15]: first, when the laser field is intense enough, an electronic wavepacket is detached from the atom through tunnel ionization; second, the wavepacket is accelerated by the laser field, and due to its oscillatory behavior, it is driven back towards the parent ion; finally, upon recollision, high frequency radiation is emitted in the form of multiple harmonics of the driving laser field. The resulting HHG emission is composed of a *plateau* of high-order harmonics with similar intensity and regular phase relationship, that extend towards the EUV or soft x-rays. At the macroscopic level, the coherent addition of the HHG emission from each atom in the gas target needs to be considered [16]. Being a high nonlinear process, the amplitude and phase of the HHG radiation is very sensitive to the details of the driving IR laser pulse. As a result, phase-matching of the HHG emission is critical to achieve efficient harmonic radiation [17–19], bright enough to be useful in experiments. In addition, the macroscopic description is required to harness the characteristics of the harmonic beam, such as its fo-

cusing [20–23] or angular momentum [24–26] properties, among others.

Theory has played a key role in the development of HHG, not only to understand the physics behind the process, but to guide fundamental advances. As such, attosecond pulse generation was theoretically predicted in the 90's [27,28], before being experimentally measured at the beginning of the present century [29,30]. Analogously, the experimental achievement of circularly polarized harmonics, reported in 2014 [31], was based on the bicircular driving technique that was theoretically proposed almost twenty years before [32,33]. More recently, the complete understanding of orbital angular momentum (OAM) conservation in HHG was understood theoretically [34], before its experimental evidence [35,36].

However, the exact theoretical description of HHG remains as a holy grail, due to the extremely expensive computational time required to achieve it. At the microscopic level, the exact calculation of HHG is given by the solution of the three-dimensional time-dependent Schrödinger equation (3D-TDSE), which describes the laser-driven quantum wavepacket dynamics in the vicinity of each atom or molecule. A full calculation of the 3D-TDSE must include the three-dimensional dynamics of all electrons in the atom or molecule. Fortunately, the main phenomenology of HHG is properly described by the electron occupying the outermost valence orbital, and thus, the single-active electron approximation is a reasonable—and widely used—assumption. Such approximation, however, is not able to describe multielectron effects, whose signature can be present in HHG under certain situations [37,38]. A standard single-atom 3D-TDSE HHG calculation under the single-active electron approximation requires from several minutes to hours in a single thread, depending mainly on the wavelength, intensity and temporal duration of the driving field. This calculation time increases dramatically if one aims to compute the 3D-TDSE single-atom response in all atoms that are present in an experiment—trillions of them in a standard gas jet. In addition, the exact calculation at the macroscopic level would require to couple the single-atom HHG results with the Maxwell equations to properly describe the propagation of the driving field, including other perturbative nonlinear effects in addition to HHG—such as self-phase modulation, plasma defocusing, Kerr effect, etc. Such calculation is extremely expensive, beyond the state-of-the-art computation capabilities.

Despite the computation limitations, clever approaches have been used to successfully simulate HHG. By neglecting the nonlinear propagation of the driving beam in the generating medium, and considering symmetries to reduce the calculation dimensions both at the microscopic and macroscopic levels, the 3D-TDSE can be coupled with the Maxwell equations [39,40], though limited cases can be reproduced. Under the assumption that once generated, the harmonics are barely affected by the target, the macroscopic calculation can be further simplified [16,20]. At the microscopic level, the development of the strong-field approximation (SFA) [41–43] has allowed not only to speed-up the calculations, but to understand the physics beyond the HHG process. In short, the SFA considers the electronic wavepacket as a Volkov wave once ionized, and neglects the field-dressing of the bound state under recollision. It thus retains the quantum nature of HHG, and the properties of the resulting high-order harmonics are in good agreement against the 3D-TDSE results. The implementation of the so-called extended SFA [44], which includes the field-dressing in the recollision step, allows also to improve the quantitative agreement against the 3D-TDSE results. A further approach consists on applying the saddle-point approximation over the SFA [45], which describes the HHG process in terms of semiclassical trajectories, drastically reducing the computational time. Thus, the SFA, in its different approaches, has been successfully used to speed up the

macroscopic calculations [46–49]. Other approaches to compute HHG at the microscopic level are based on calculating ionization rates [50], or the quantitative rescattering model [51], which has been also used as a source for macroscopic calculations [52].

In parallel to the development of strong-field theoretical methods, computer science has been evolving exponentially, not only by increasing the computational power provided by modern graphical processor units (GPU)—and corresponding software libraries—but by offering alternative resources. Among them, Artificial Intelligence (AI), and specifically the development of Neural Networks (NN), stand as a new paradigm to infer simulation results based on massive training. In particular, deep learning is becoming a powerful tool in ultrafast and nonlinear optics. Machine learning algorithms can predict the X-ray pulse properties of free electron lasers [53], retrieve complex molecular structures [54,55], reconstruct femtosecond laser pulses [56], speed up the streaking retrieval of attosecond pulse measurements [57,58], or offer alternative approaches to calculate strong-field time-dependent dynamics [59]. AI has also been applied to predict characteristics of harmonic generation in plasmas (such as flux, duration or maximum photon energy) using particle-in-cell simulations [60]. However, the application of AI to HHG has been scarce. At the microscopic level, AI has been applied to predict the HHG emission from molecules by training a NN with the two dimensional TDSE for different molecular parameters [61]. Such approach also showed the potentiality to apply AI to solve the inverse problem, i.e. to retrieve the molecular parameters from the microscopic HHG spectrum [61]. At the macroscopic level, AI has been used to predict the harmonic flux of HHG depending on macroscopic parameters such as the pulse energy, gas pressure, gas cell position and medium length [62], where the NN was trained using the SFA-based macroscopic method [46].

In this work we introduce the use of AI to obtain complete 3D-TDSE-based macroscopic HHG calculations in atomic gases. We train a NN to infer the microscopic single-atom HHG response through the solution of the 3D-TDSE for different amplitudes and phases of the driving laser pulse. To do so, we have implemented two NN architectures, depending on whether the single-atom HHG output is trained in the temporal or in the spectral domain. By using the exact solution of the integral Maxwell equations, we combine the predicted single-atom emissions in a macroscopic gas jet and propagate them towards a far-field detector, following the approach described in [48].

This method, which is extremely fast once the NN is trained, is particularly fitted to predict the macroscopic HHG response from structured driving beams in low density gas jets. In such a case, the variation of the driving laser pulse across the target can be described through changes in its amplitude and spatial phase. A paradigmatic example is HHG driven by structured beams carrying OAM. We have applied our method to simulate HHG from two beams with different topological charges, a scenario that provides a clear signature of the non-perturbative behavior of the HHG process [63,26].

Our results demonstrate that AI applied to HHG provides a two-fold advancement: (i) it speeds up the calculations, providing a route towards *in-situ* strong-field simulations that can be performed within minutes or seconds, in parallel to the experiments; and (ii) it allows to perform TDSE-based macroscopic simulations that can reveal hidden signatures neglected in the standard approximations.

The paper is organized as follows. First we present our AI approach to compute 3D-TDSE-based macroscopic HHG, introducing the two architectures used to train the single-atom NN in the temporal and spectral regimes. Then we perform macroscopic HHG calculations using structured driving beams. We reveal the OAM content of the harmonics, as well as their emission as structured
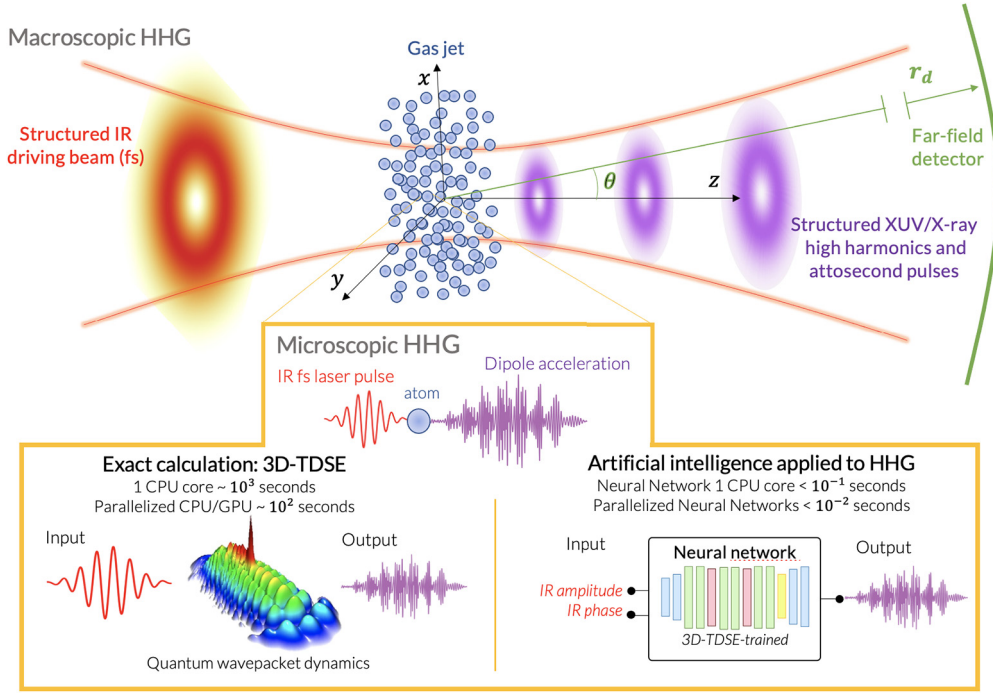
Fig. 1. Scheme of the HHG process driven by structured laser beams. At the macroscopic scale, the interaction of an intense structured IR laser beam with a gas target results in the emission of high-order harmonics that propagate towards a far-field detector placed at $\mathbf{r}_d$. At the microscopic level, in each of the atoms of the gas target the intense IR laser pulse drives coherent electronic wavepacket dynamics. The resulting dipole acceleration is composed of high-order harmonics that can be synthesized into attosecond pulses. The exact calculation of the HHG process at the microscopic level is given by the 3D-TDSE. In this work we propose an alternative to drastically reduce the computation time required to solve the 3D-TDSE, by using neural networks that are properly trained against 3D-TDSE results.

attosecond pulses. In the Supplementary Material we provide calculations based on the extended SFA for comparison, and additional macroscopic 3D-TDSE HHG calculations.

2. Method: artificial intelligence applied to macroscopic high-order harmonic generation

The general sketch of HHG in gases is depicted in Fig. 1. An intense IR laser beam is focused into a gas jet, where high-order harmonics are produced. The harmonic radiation from each atom in the target is propagated towards a far-field detector, placed at $\mathbf{r}_d$. In order to simulate the macroscopic response of HHG, we follow the method depicted in [48]. In general, the propagation of the transverse component of the electric field, $\mathbf{E}(\mathbf{r}, t)$, is solved through the wave equation

$$\nabla^2 \mathbf{E}(\mathbf{r}, t) - \frac{1}{c^2}\frac{\partial^2}{\partial t^2}\mathbf{E}(\mathbf{r}, t) = \frac{1}{\epsilon_0 c^2}\frac{\partial}{\partial t}\mathbf{J}(\mathbf{r}, t), \quad (1)$$

where $\mathbf{J}$ is the current density and c the speed of light. The electric field can be decomposed into that of the propagating driving field, and that radiated by the gas target upon the nonlinear interaction. In addition, this latter source can be decomposed into the sum of the emission from each of the n atoms present in the gas target. Then, by considering the integral solution of the wave equation, the total electric field can be expressed as

$$\mathbf{E}(\mathbf{r}, t) = \mathbf{E}_0(\mathbf{r}, t) + \sum_{j=1}^{n}\mathbf{E}_j(\mathbf{r}, t), \quad (2)$$

where $\mathbf{E}_0(\mathbf{r}, t)$ is the driving field and $\mathbf{E}_j(\mathbf{r}, t)$ is the harmonic field radiated from the accelerated charges in the target, given by [64]

$$\mathbf{E}_j(\mathbf{r}_d, t) = \frac{1}{4\pi\epsilon_0}\frac{1}{c^2|\mathbf{r}_d - \mathbf{r}_j(0)|}\mathbf{s}_d$$

$$\times \left\{ \mathbf{s}_d \times \mathbf{a}_j\left[t - \frac{|\mathbf{r}_d - \mathbf{r}_j(0)|}{c} \right] \right\}. \quad (3)$$

$\mathbf{a}_j$ is the dipole acceleration evaluated at the retarded time, and $\mathbf{s}_d$ is the unit vector pointing to a virtual detector located at $\mathbf{r}_d$ (see Fig. 1). We use the dipole approximation, i.e. we assume that the charge displacement during the interaction is considerably smaller than the wavelength of the driving and radiated fields. In addition we assume that the charge's dynamics are driven solely by the intense IR driving field, $\mathbf{E}_0(\mathbf{r}, t)$.

We consider that the IR driving field propagates in vacuum, an assumption that is valid in the case of a low density target (gas jets with densities of $\sim 10^{17}$ atoms/cm^3 and radial thicknesses of $\sim 10^2$ μm), such as those used in standard HHG experiments. A detailed analysis of such approximation can be found in [48]. In addition, this assumption is especially suited when structured IR drivers carrying OAM are used, an scenario where transverse phase-matching dominates over longitudinal phase-matching, and thus the gas jet can be approximated as a thin slab [65].

In this work we employ the Laguerre Gauss modes to describe the spatial structure of the IR driving beam, though the method can be used for any spatial profile. We thus consider an electric field linearly polarized along the x-direction, propagating in the z-direction under the paraxial approximation, and with wavelength λ_0 ($k_0 = 2\pi/\lambda_0$), expressed as $\mathbf{E}_0(\mathbf{r}, t) = E_0 F(t) LG_{\ell, p}(\mathbf{r}) e^{ik_0 z}\mathbf{e}_x$. The spatial dependence is given by the Laguerre-Gauss modes in cylindrical coordinates as

$$LG_{\ell, p}(\rho, \phi, z)$$

$$= \frac{w_0}{w(z)}\left(\frac{\sqrt{(2)}\rho}{w(z)} \right)^{|\ell|} L_p^{|\ell|}\left(\frac{2\rho^2}{w^2(z)} \right) e^{\frac{-\rho^2}{w^2(z)}} e^{i\varphi_s(\rho, \phi, z)}, \quad (4)$$
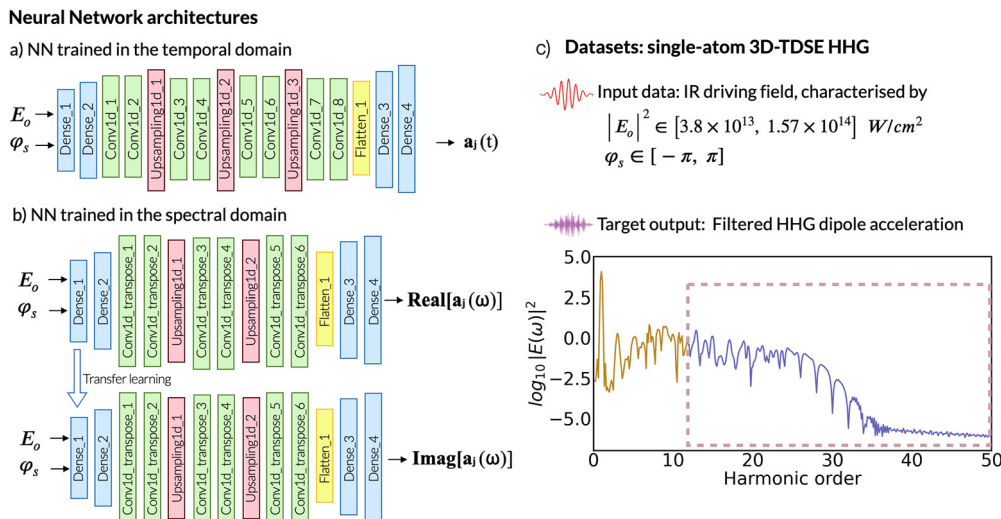
Neural Network architectures

Fig. 2. Neural network architectures used to train the single-atom HHG in a) the temporal and b) the spectral domain. Both architectures make use of two inputs (E_0 and φ_s) that are connected to two consecutive fully connected layers (*Dense* in Keras), with an increasing number of neurons. While in the temporal domain the NN uses convolutional layers, in the spectral domain we build two NN (to predict the real and imaginary parts of the spectrum), that make use of transpose convolutional layers. In both cases the output of the last convolutional layer is flatten and connected to a *Dense* layer, while a second fully connected layer provides the output. While in a) the output is the temporal evolution of the dipole acceleration $\mathbf{a}_j(t)$, in b) the output consists on the real and imaginary parts of the harmonic spectrum, $\mathbf{a}_j(\omega)$. c) The dataset consists on 3D-TDSE calculations for different input IR driving pulses, scanning over the two input parameters ($\varphi_s \in [-\pi, \pi]$, $|E_0|^2 \in [3.8 \times 10^{13}$ W/cm^2, 1.57×10^{14} W/cm^2]). The target output, either $\mathbf{a}_j(t)$ or $\mathbf{a}_j(\omega)$ is filtered to properly reproduce the harmonics above the 12^{th}-order.

where $\varphi_s(\rho, \phi, z) = \ell\phi + \frac{k_0\rho^2}{2R(z)} + \Phi_G(z)$ is the spatial phase of each mode. The beam waist is given by $w(z) = w_0\sqrt{1 + (z/z_0)^2}$ with w_0 the beam waist at focus and $z_0 = k_0 w_0^2/2$ the Rayleigh range. $R(z) = z[1 + (z_0/z)^2]$ is the radius of curvature, $\Phi_G(z) = -(2p + |\ell| + 1)\arctan(z/z_0)$ the Gouy phase, and $L_p^{|\ell|}$ the associated Laguerre polynomials. The indices $\ell = 0, \pm1, \pm2, ...$ and $p = 0, 1, 2, ...$ correspond to the topological charge and the number of non-axial radial nodes of the mode, respectively. The temporal evolution of the laser pulse is modeled as $F(t) = \sin^2(\pi t/NT)\sin\omega_0 t$, where T is the period, $\omega_0 = 2\pi/T = 2\pi c/\lambda_0$ the field frequency, and N depicts the temporal duration of the pulse as the total number of periods.

We compute the charge's acceleration through the full integration of the 3D-TDSE, under the single-active electron approximation. The Hamiltonian governing the interaction is given by

$$H(t, \mathbf{r}) = \frac{\mathbf{p}^2(t)}{2m} + V(\mathbf{r}), \tag{5}$$

where m is the electron mass, $V(\mathbf{r})$ the Coulomb potential energy, and $\mathbf{p} = -i\hbar\nabla - (q/c)\mathbf{A}(t)$, $\mathbf{A}(t)$ being the vector potential of the linearly polarized laser field at the atom's position, $\mathbf{E}_0(t, \mathbf{r}_j)$, with q the electron charge. The dipole acceleration, $\mathbf{a}_j(t)$, is computed from the mean value of the acceleration operator $-(1/m)\nabla V(\mathbf{r})$. The integration is carried out using a Crank-Nicolson algorithm in the finite differences scheme.

The main advance of this work is to train a NN to predict the 3D-TDSE dipole acceleration, $\mathbf{a}_j(t)$, from the amplitude (E_0) and spatial phase (φ_s) of the driving field. To do so, we fix the wavelength (λ_0) and pulse duration (N) of the linearly polarized driving field, and we train the NN to infer the dipole acceleration for any amplitude (E_0) and spatial phase (φ_s) of the driving laser pulse. This allows us to simulate HHG driven by a structured beam represented by any spatial mode, such as the $LG_{\ell,p}$ modes. As a result, once the NN is trained and validated, we demonstrate that the time required to perform a full macroscopic HHG calculation is drastically reduced by 4 to 5 orders of magnitude, depending on the size of the far-field detector and the number of microscopic calculations required to reach the convergence.

2.1. Neural networks trained with the 3D-TDSE

In this work we have implemented two NN architectures to predict the microscopic HHG response. The two architectures differ in the output target of the NN training: the dipole acceleration either in the temporal or in the spectral domain, as depicted in Figs. 2a and 2b, respectively. Though the results are similar, there are several differences in the implementation and performance, as we shall comment. NNs are typically decomposed in three stages: the input parameters, followed by several hidden layers to process the input, and the output. In order to reproduce the microscopic HHG process, we introduce the IR driving laser pulse through two input parameters—amplitude (E_0) and spatial phase (φ_s)—aiming to obtain as an output the HHG dipole acceleration either in the temporal, $\mathbf{a}_j(t)$, or spectral, $\mathbf{a}_j(\omega)$, domain.

The NN architectures that we have implemented make use of convolutional layers, which have been demonstrated to be extremely useful in image recognition, text processing and time series data for their ability to identify patterns. In particular we have used one-dimensional convolutional layers to predict the temporal evolution of the dipole acceleration (similar to that performed in [61]), and one-dimensional transpose convolutional layers to predict the HHG response from the spectral domain. These latter layers are also commonly used for feature recognition [66]. In Fig. 2 we detail the layer composition used in both architectures, with all neurons using *tanh* as their activation function.

Our NNs are implemented with the libraries Keras and Tensorflow in Python [67] and trained with GPU acceleration. They take two inputs (E_0 and φ_s) that are connected to two consecutive fully connected layers (*Dense* in Keras), with an increasing number of neurons. The output of the second *Dense* layer is connected to a first convolutional layer, that is followed by another one of the same type. We repeat this block of two consecutive convolutional layers several times, four in the NN that we use to predict HHG in the time domain, and three in the ones we use in the spectral domain. Between these blocks, we have an upsampling layer to duplicate the number of outputs. The output of the last convolutional layer is flatten and connected to a *Dense* layer. Finally, a second fully connected layer provides the output of our NNs.
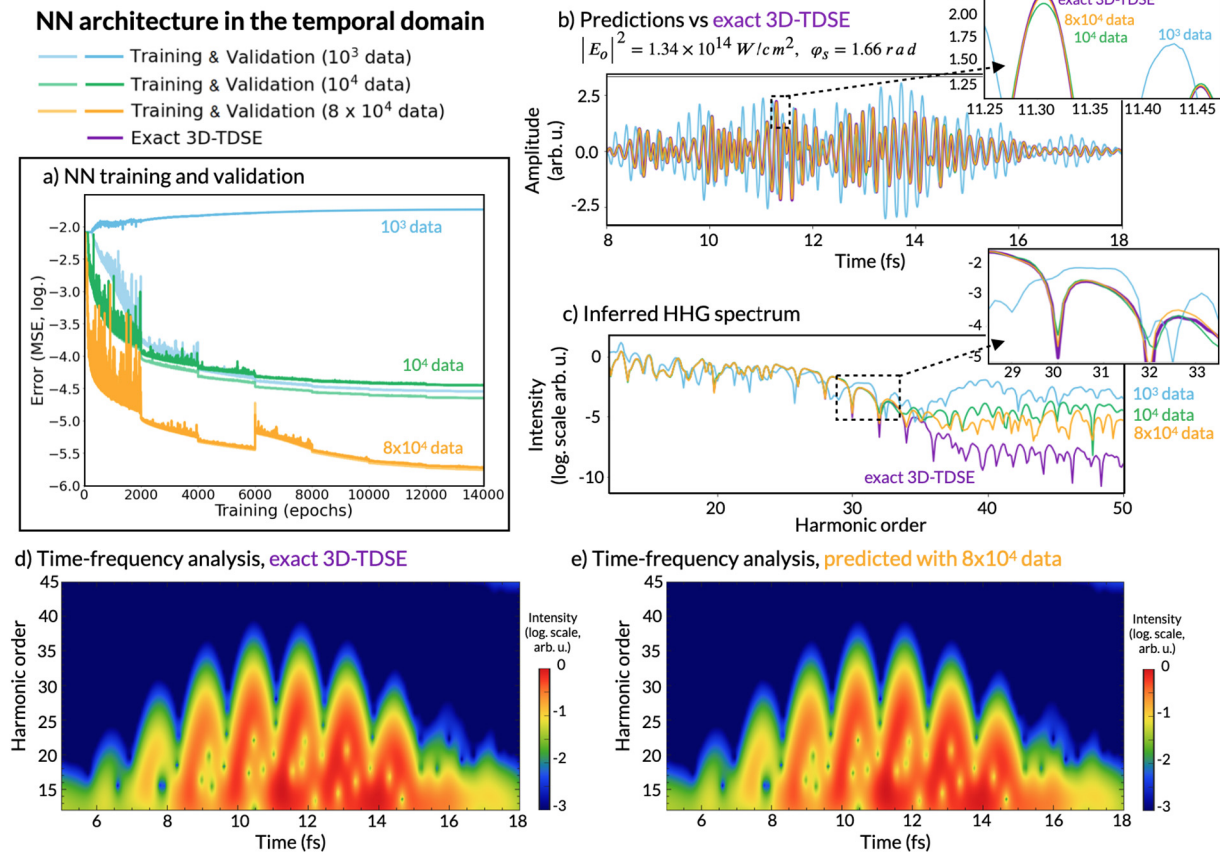
Fig. 3. NN architecture trained in the temporal domain. a) Training and validation curves of the 3D-TDSE dipole acceleration, $a_j(t)$, for dataset sizes of 10^3 (blue), 10^4 (green), and 8×10^4 (orange). Light colors are used for the MSE NN training, while darker colors represent the MSE NN validation. b) Prediction for each dataset size against the exact 3D-TDSE result (purple) for the case of $|E_0|^2 = 1.34 \times 10^{14}$ W/cm^2 and $\varphi_s = 1.66$ rad. c) Inferred HHG spectrum after performing the Fourier Transform of the results presented in b). The higher part of the spectrum allows to identify the better accuracy of the bigger dataset sizes. d) and e) show the time-frequency analysis of the exact 3D-TDSE and the predicted case with 8×10^4 data, respectively. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

In order to train the NNs, we have generated a dataset of 3D-TDSE calculations for different input IR driving pulses. In particular, we fix the wavelength ($\lambda_0 = 800$ nm) and pulse duration ($N = 8$ cycles), which corresponds to 7.7 femtoseconds in full width at half maximum in intensity (FWHM). We randomly scan over the two input parameters as follows: all possible phases are considered ($\varphi_s \in [-\pi, \pi]$), whereas the amplitude is scanned to cover peak intensities $|E_0|^2 \in [3.8 \times 10^{13}$ W/cm^2, 1.57×10^{14} W/cm$^2]$. Note that as we are interested in the high-order harmonics, the lower intensity limit is given by the field amplitude that will not generate harmonics beyond the 12^{th}-order. On the other hand, the upper intensity limit is given by the barrier suppression threshold in the ionization of the hydrogen atom. The output target values for the NN training are the temporal (or spectral) distributions of the dipole acceleration. Note that we spectrally filter the HHG spectrum below the 12^{th} harmonic order, as depicted in Fig. 2. This is an important feature of our method. In standard HHG, the high-order harmonics that are close to the cutoff frequency—those that we are more interested in—are orders of magnitude weaker than the fundamental and lower-order ones. Thus, if the low-frequency part of the HHG spectrum is included in the NN training, the prediction over the high-order harmonics fails.

We have chosen atomic hydrogen as the gas target for simplicity, though our method could be easily extended to any other atom or molecule. The 3D-TDSE dataset is calculated using 8192 temporal points (with a grid spacing of 0.94 attoseconds), and a spatial grid in cylindrical coordinates with 2000 points in the direction of the field polarization and 800 points in the radial direction (with a

grid spacing of 5.3×10^{-3} nm). Within those parameters, the noise level in the normalized HHG spectrum is $\sim 10^{-9}$ (see Fig. 2c).

We have trained all our NNs using Adam as the optimizer, with a learning rate of $5 \cdot 10^{-4}$, and taking the mean squared error (MSE) as the loss function to be minimized. The training of the NN in the temporal domain takes place in sets of 2000 epochs for an increasing batch size of 2^m with $m = 3, 4, 5...9$. This training method has been proven to be faster by reducing the number of weights updates while keeping accuracy [68]. In Fig. 3a we present the training and validation of the NN trained in the temporal domain for different dataset sizes: 10^3 in blue, 10^4 in green, and 8×10^4 in orange. Light colors are used for the MSE NN training, while darker colors represent the MSE NN validation.

The NN validation fails for the smallest dataset considered (10^3 3D-TDSE calculations). However, the NN is properly trained if a dataset of $\geq 10^4$ is used, with an increasing accuracy with the dataset size. In Fig. 3b we show an example of the prediction for each dataset size considered in Fig. 3a, for input parameters $|E_0|^2 = 1.34 \times 10^{14}$ W/cm^2 and $\varphi_s = 1.66$ rad, compared to the exact 3D-TDSE calculation (purple line). While the smallest dataset fails in predicting the 3D-TDSE results, if the dataset is big enough, the NN is able to predict the attosecond pulse evolution with great accuracy, as shown in the inset of Fig. 3b. Another comparison can be drawn by performing the Fourier Transform and showing the resulting high-order harmonics in the spectral domain (Fig. 3c). In this case, the low harmonic signal beyond the maximum frequency—the so-called cutoff frequency—allows to easily
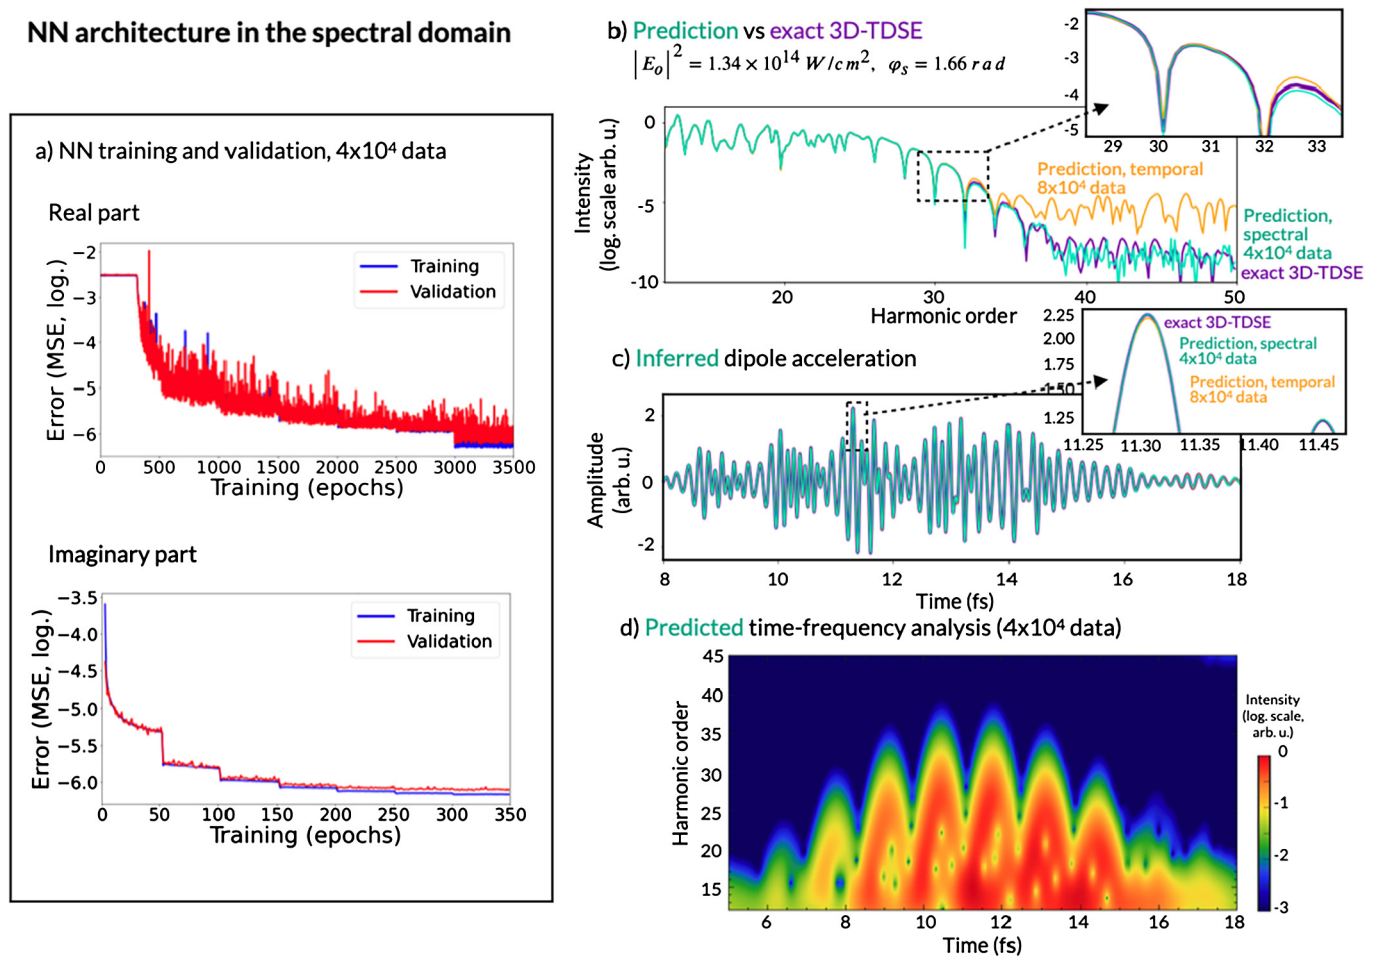
Fig. 4. NN architecture trained in the spectral domain. a) Training (blue) and validation (red) curves of the 3D-TDSE real and imaginary parts of the HHG field, $\mathbf{a}_j(\omega)$, for a dataset sizes of 4×10^4. b) Prediction of the HHG spectrum (cyan) against the exact 3D-TDSE result (purple) and the prediction with the temporal NN (orange), for the case of $|E_0|^2 = 1.34 \times 10^{14}$ W/cm^2 and $\varphi_s = 1.66$ rad. c) Inferred temporal evolution of the dipole acceleration after performing the Fourier Transform of the results presented in b). d) Time-frequency analysis of the predicted case, which directly compares with the exact 3D-TDSE case shown in Fig. 3d.

observe how the accuracy of the NN increases with the dataset size.

In order to compare not only the spectral intensity but also the phase of the resulting HHG emission, we show in Figs. 3d and 3e the time-frequency analysis for the exact 3D-TDSE and the predicted (using 8×10^4 data) dipole accelerations, respectively. In the time-frequency analysis, we select a spectral window in the HHG spectrum and take its Fourier transform. By shifting the window to cover the entire HHG spectrum, it is possible to resolve the time in which the different harmonics are generated. In Figs. 3d and 3e we consider a Gaussian spectral window with FWHM $6\omega_0$. The excellent agreement between the two time-frequency plots allows us to unequivocally confirm the prediction of our NN. Note that the positive (negative) slope of the structures in the time-frequency analyses reflects the contribution from the so-called short (long) trajectory contributions [69]. Such structures are more complex at the rear part of the pulse, and for the lower order harmonics, which reflects the complexity of the quantum wavepacket dynamics described by the 3D-TDSE. Remarkably, all the details of the time-frequency analysis are perfectly reproduced by the NN. Finally, it is important to note that although Figs. 3b-e show the agreement of a single prediction (E_0, φ_s), the low error obtained in the training after 14000 epochs (Fig. 3a) ensures excellent agreement for all possible cases within the $|E_0|^2$ limits considered.

We also propose to train the 3D-TDSE NN by optimizing the output dipole acceleration in the spectral domain—instead the

temporal domain—as already depicted in Fig. 2. As Tensorflow library layers can only generate real valued NNs, we use two equal NNs to predict the real and imaginary parts of the HHG spectrum (see Fig. 2b). However, in this case there is no need to consider the whole HHG spectral array, and we can restrict each dataset to the real and imaginary parts of the spectrum lying in the interval where the high-order harmonics are emitted [$12\omega_0$, $50\omega_0$]. In our particular case, this reduces each dataset from 8192 points to 2×304 points. In Fig. 4a we show the training and validation of the real and imaginary parts of the HHG spectrum, for a dataset size of 4×10^4. The training of the real part NN takes place in sets of 500 epochs for an increasing batch size of 2^m with $m = 3, 4, 5...9$. For the imaginary part NN, once the real part NN is trained, we consider transfer learning by copying all trainable parameters of the real part NN into the imaginary part NN. In addition, all weights of the imaginary part NN are frozen (making their respective layers non-trainable) except for the last two dense layers, which are trained in sets of 50 epochs with the same batch size methodology.

We present in Fig. 4b an example of the prediction (solid cyan line) from the spectral NN for input parameters $|E_0|^2 = 1.34 \times 10^{14}$ W/cm^2 and $\varphi_s = 1.66$ rad (same as in Figs. 3b-e), compared to the exact 3D-TDSE calculation (purple line), and the prediction from the previous temporal NN (orange line). We also show in Fig. 4c the comparison for the temporal evolution of the dipole acceleration after performing the Fourier Transform of the results

presented in Fig. 4b. In addition, Fig. 4d shows the corresponding time-frequency analysis for the predicted case. It is worth mentioning that the NN training of the imaginary part of the spectrum is much more stable than the training of the real part. This fact is probably due to the transfer learning carried out, which indicates that convolutional layers could be an important source of those instabilities.

While both NN architectures perform very well, the NN trained in the spectral domain reduces the MSE for the same training dataset size, thus giving a more accurate prediction, as observed in the higher part of the HHG spectrum in Fig. 4b. However, the training of the NN in the spectral domain is less stable than in the temporal domain, as depicted in Fig. 4a. Thus, the NN architecture trained in the temporal domain, while being more robust, is able to provide better results for smaller dataset sizes.

3. Results. High harmonic generation driven by structured driving beams

Our AI-based method to compute macroscopic HHG is especially suited to study HHG driven by laser beams structured in their OAM, as their spatial distribution in a low density gas jet can be fully characterized by E_0 and φ_s. The link between the topological charge (ℓ) of a light beam and its helical waveform established in 1992 by Allen and coworkers [70] paved the way to many applications that make use of the OAM of light, such as optical communications, quantum optics, imaging, metrology, etc. [71,72]. One can easily understand the potentiality of translating these applications into the nanometer and sub-femtosecond scales if OAM beams are generated in the ultrafast/x-ray regimes. As such, in the last decade HHG has emerged as a unique process to up-convert OAM beams from the IR towards the XUV/x-rays. Indeed, theoretical macroscopic HHG simulations [34] have been essential to understand the rules of angular momentum transfer in the HHG process, and thus to create high-order harmonics and attosecond pulses with OAM. Nowadays we are able to generate high-order harmonics with high [35,36,73] and low [24,74,75] topological charges, ultrafast vector beams [76–78], circularly polarized EUV OAM beams [25,79], high harmonics with time-dependent OAM [26], or high harmonics from a polarization Mobiüs strip [80], among others.

Up to now, most of the theoretical calculations performed in these works of OAM-HHG have been based on the SFA approach, with a successful comparison against experiments. In this work we have applied our AI-HHG method to perform the first macroscopic OAM-HHG calculations based on the 3D-TDSE. In order to show the performance of our method, we have focused our attention in HHG driven by two IR driving beams with different topological charges, a configuration that has been demonstrated to reveal the non-perturbative character of HHG as the appearance of additional OAM contributions in the high-order harmonics [63]. To do so, we drive HHG with a laser beam composed of two collinear LG modes (ℓ_1=1 and ℓ_2=2) with same intensity, focused into an atomic hydrogen slab target with a beam waist of $w_0 = 30$ μm. We point out that though in these simulations we are using a thin slab, our method can properly describe HHG in a thicker low density gas jet. Each laser pulse is modeled with a temporal duration of $N = 8$ cycles, wavelength $\lambda_0 = 800$ nm, and their amplitudes are chosen to give a maximum peak intensity at focus of 1.57×10^{14} W/cm². The spatial intensity and phase profiles of the resulting IR beam at the target slab are shown in Fig. 5a. It can be observed that all possible values of (E_0, φ_s) at different target positions were covered in the training of our NNs. In the following, we have chosen the predictions from the NN trained in the spectral domain with a dataset of 4×10^4.

In Fig. 5b we show the far-field intensity of the high-order harmonics as a function of the divergence angle$-\theta$, defined in Fig. 1. Note that the harmonics are spatially integrated along the azimuthal coordinate. In Fig. 5c we show the OAM content of the 15^{th} and 21^{st} harmonics. The topological charge of the harmonics is retrieved by performing the Fourier Transform of the harmonic field along the azimuthal coordinate. It is worth to mention that the OAM spectrum of the harmonics presents several contributions, following the non-perturbative OAM build-up rule introduced in [63], with similar results as those obtained with the macroscopic SFA approach (see Supplementary Material).

Fig. 5d shows the far-field spatial intensity profile of the HHG beam, integrated spectrally, i.e., where all high-order harmonics are included. Our TDSE-based method provides a unique opportunity to look into the synthesized attosecond pulses. In Fig. 5e we plot the attosecond pulse train and the time-frequency analysis of the harmonic emission detected at the maximum of the intensity profile. The attosecond pulse train obtained in the macroscopic picture is cleaner than the single-atom one (shown in Figs. 3b and 4c). This can be explained through the time-frequency analysis, where one can identify the quantum path contributions that are cleared-up upon macroscopic build-up.

A second example, where the IR driving beam is composed of two beams with $\ell_1 = 1$ and $\ell_2 = 3$, is shown in the Supplementary Material. For completeness, we also show in the Supplementary Material a comparison between the results presented in Fig. 5 and those obtained with other methods based on the extended SFA (instead of the 3D-TDSE), using both non-AI assisted, and AI-assisted methodologies.

In Table 1 we show the performance of the AI-based HHG method within the results presented in this work, using different levels of parallelization and hardware. At the single-atom level, the use of NNs allows to reduce the computational time required to calculate the dipole acceleration from several minutes to tens of milliseconds. This gain is translated into the full macroscopic calculations. For the results presented in Fig. 5, where 10^5 single-atom calculations were used, the use of NNs allows to perform simulations within minutes, compared to several months (estimation) with the standard method. Finally, one should also consider the computational time required for training the NNs. Using a GPU architecture, the NN training in the temporal domain took ~ 12 days, whereas the NNs training in the spectral domain took ~ 4 days.

4. Conclusions and perspectives

We provide a fast and accurate method to simulate macroscopic HHG, where the single-atom dipole acceleration is obtained through the 3D-TDSE. To do so, we have trained a NN to predict the single atom emission. We provide two alternative solutions, training the NN against 3D-TDSE results in the temporal or in the spectral domain. While both architectures yield excellent results, they present some differences. The NN trained in the spectral domain reduces the mean square error for the same training dataset size, thus giving a more accurate prediction. However, the training of the NN in the spectral domain is less stable, which means that the NN architecture trained in the temporal domain is able to provide better results for smaller dataset sizes. Once the NN is trained, our method is able to compute macroscopic HHG within minutes or even seconds.

The fast method that we present is particularly suited for computing HHG from structured drivers. We have demonstrated its feasibility to calculate the far-field harmonic properties−intensity profile, OAM spectrum, temporal emission−when driven by a combination of two infrared OAM beams, a scenario where the non-perturbative aspects of HHG are particularly relevant. Remarkably,
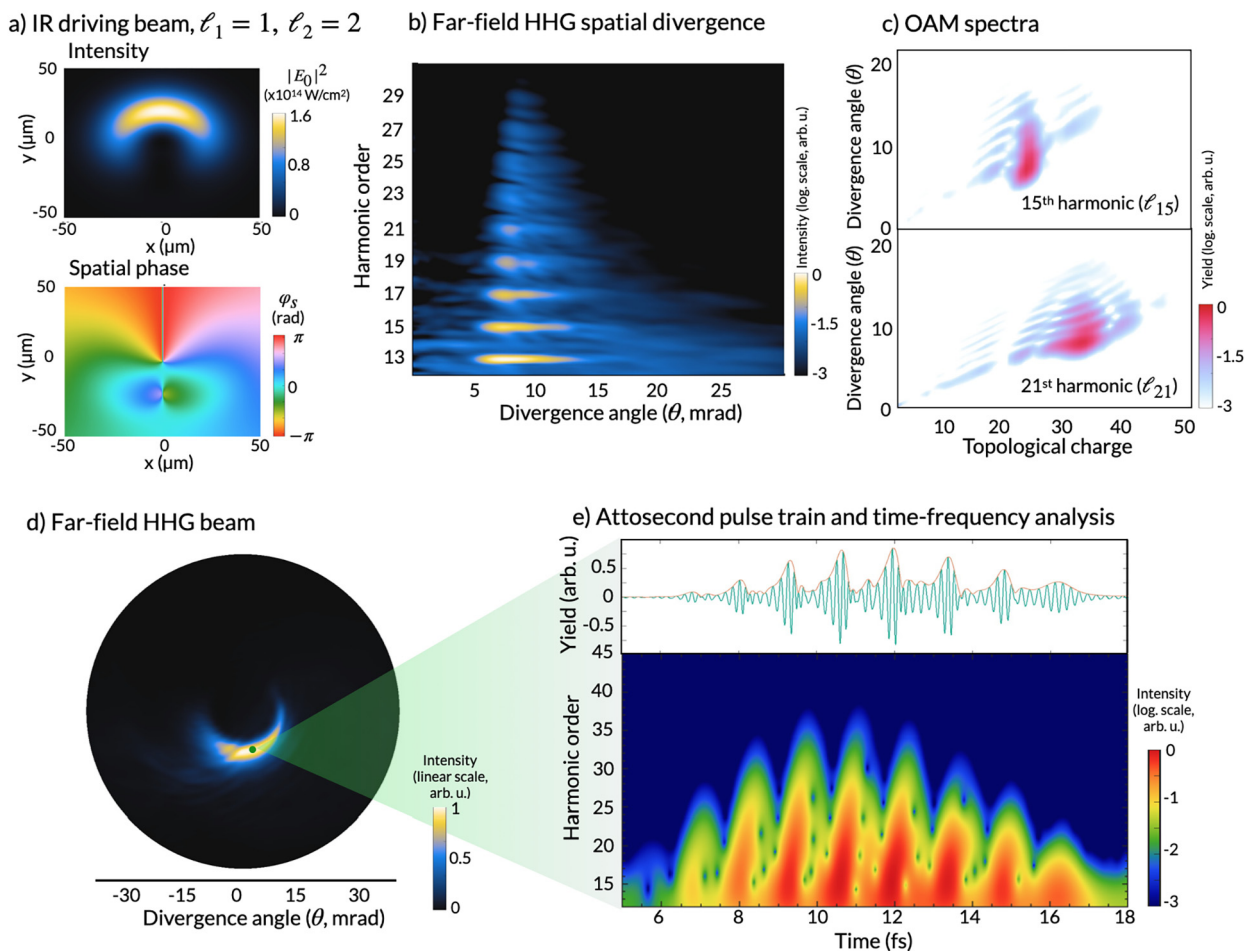
Fig. 5. Results from the AI-based macroscopic HHG method using the NN trained in the spectral domain. a) Intensity and spatial phase profiles of the driving beam, composed of two beams with $\ell_1 = 1$ and $\ell_2 = 2$. b) Far-field intensity spectrum, where the harmonics are integrated in the azimuthal coordinate. c) OAM spectra of the 15^{th} and 21^{st} harmonics, obtained after performing the Fourier Transform of the harmonic field along the azimuthal coordinate. d) Far-field intensity profile of the HHG beam. e) Attosecond pulse train and time-frequency analysis of the harmonic signal detected at the maximum of the intensity profile.

Table 1

Performance of the AI-based HHG method. We show the computational time required to compute both single-atom 3D-TDSE calculations, and the full macroscopic 3D-TDSE simulations. The computational times are based on the simulations presented in Fig. 5. We have considered different hardware and parallelization approaches, using Intel Xeon 6240R CPU with 24 cores, and NVIDIA A30 GPU with 3804 CUDA cores. As it can be observed, the use of NNs outperform all other implementations by several orders of magnitude. Note that in this table we have not included the computational time required to train the NN.

Single-atom 3D-TDSE calculation	Time per atom
1 CPU core	~70 minutes
24 CPU cores (1 atom per core)	~5 minutes
3804 CUDA cores (GPU)	~2 minutes
Neural network (1 CPU core)	<0.1 seconds
24 Neural networks (1 atom per CPU core)	<0.01 seconds

Full-macroscopic 3D-TDSE (using 10^5 single-atom calculations)	Total time
(Estimation) 1 CPU core	~13 years
(Estimation) 4 processors (96 CPU cores) in parallel (1 core/atom)	~3 months
Neural Network (1 CPU core)	~1 hour
Neural Networks (4 processors, 96 CPU cores) in parallel (1 core/atom)	~1 minute

our method allows to predict the emitted attosecond pulses, including all the quantum wavepacket dynamics described by the 3D-TDSE. We point out that although we have presented simulations of Laguerre-Gauss IR driving beams focused into a thin slab atomic hydrogen target, our method is universal, and can be used to simulate HHG in thicker low density gas jets, composed of other atomic or molecular species, and driven by structured beams described in any other basis.

However, while our method provides a new route to perform macroscopic HHG calculations, there is still room for improvement. Our NNs are trained with two input parameters, the amplitude and spatial phase of the IR driving beam. This allows us to reproduce structured infrared beams whose intensity and phase distribution varies spatially. However, this approach excludes other scenarios including temporal or spatial chirp of the driving field, the influence of the neutral atoms or free electrons in the propagation of the driving beam, or nonlinear propagation effects. In such cases the NN must include more input parameters, making the training against the 3D-TDSE more challenging, due to the larger dataset size required.

From the computational point of view, we have demonstrated an effective gain of 4 to 5 orders of magnitude in the time required to achieve macroscopic simulations based on the 3D-TDSE. Though we have used state-of-the-art high performance computational tools (such as CPU and GPU parallelization schemes), our simulations could further benefit from advances in these architectures.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

Acknowledgements

Appendix A. Supplementary material

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.cpc.2023.108823.

References

[1] A. McPherson, G. Gibson, H. Jara, U. Johann, T.S. Luk, I.A. McIntyre, K. Boyer, C.K. Rhodes, J. Opt. Soc. Am. B 4 (4) (1987) 595–601, https://doi.org/10.1364/JOSAB.4.000595, https://opg.optica.org/josab/abstract.cfm?URI=josab-4-4-595.

[2] M. Ferray, A. L'Huillier, X.F. Li, L.A. Lompre, G. Mainfray, C. Manus, J. Phys. B, At. Mol. Opt. Phys. 21 (1988) L31.

[3] T. Popmintchev, M.-C. Chen, D. Popmintchev, P. Arpin, S. Brown, S. Ališauskas, G. Andriukaitis, T. Balčiunas, O.D. Mücke, A. Pugzlys, et al., Science 336 (6086) (2012) 1287–1291.

[4] T. Gaumnitz, A. Jain, Y. Pertot, M. Huppert, I. Jordan, F. Ardana-Lamas, H.J. Wörner, Opt. Express 25 (22) (2017) 27506–27518, https://doi.org/10.1364/OE.25.027506, http://opg.optica.org/oe/abstract.cfm?URI=oe-25-22-27506.

[5] M.T. Hassan, T.T. Luu, A. Moulet, O. Raskazovskaya, P. Zhokhov, M. Garg, N. Karpowicz, A.M. Zheltikov, V. Pervak, F. Krausz, E. Goulielmakis, Nature 530 (7588) (2016) 66–70, https://doi.org/10.1038/nature16528.

[6] F. Calegari, A. Trabattoni, A. Palacios, D. Ayuso, M.C. Castrovilli, J.B. Greenwood, P. Decleva, F. Martín, M. Nisoli, J. Phys. B, At. Mol. Opt. Phys. 49 (14) (2016) 142001, https://doi.org/10.1088/0953-4075/49/14/142001.

[7] S. Beaulieu, A. Comby, A. Clergerie, J. Caillat, D. Descamps, N. Dudovich, B. Fabre, R. Géneaux, F. Légaré, S. Petit, B. Pons, G. Porat, T. Ruchon, R. Taïeb, V. Blanchet, Y. Mairesse, Science 358 (6368) (2017) 1288–1294, https://doi.org/10.1126/science.aao5624, https://www.science.org/doi/pdf/10.1126/science.aao5624, https://www.science.org/doi/abs/10.1126/science.aao5624.

[8] S. Grundmann, D. Trabert, K. Fehre, N. Strenger, A. Pier, L. Kaiser, M. Kircher, M. Weller, S. Eckart, L.P.H. Schmidt, F. Trinter, T. Jahnke, M.S. Schöffler, R. Dörner, Science 370 (6514) (2020) 339–341, https://doi.org/10.1126/science.abb9318, https://www.science.org/doi/pdf/10.1126/science.abb9318, https://www.science.org/doi/abs/10.1126/science.abb9318.

[9] Z. Tao, C. Chen, T. Szilvási, M. Keller, M. Mavrikakis, H. Kapteyn, M. Murnane, Science 353 (6294) (2016) 62–67, https://doi.org/10.1126/science.aaf6793, https://www.science.org/doi/pdf/10.1126/science.aaf6793, https://www.science.org/doi/abs/10.1126/science.aaf6793.

[10] P. Tengdin, W. You, C. Chen, X. Shi, D. Zusin, Y. Zhang, C. Gentry, A. Blonsky, M. Keller, P.M. Oppeneer, H.C. Kapteyn, Z. Tao, M.M. Murnane, Sci. Adv. 4 (3) (2018) eaap9744, https://doi.org/10.1126/sciadv.aap9744.

[11] J. Miao, T. Ishikawa, I.K. Robinson, M.M. Murnane, Science 348 (6234) (2015) 530–535, https://doi.org/10.1126/science.aaa1394, https://www.science.org/doi/pdf/10.1126/science.aaa1394, https://www.science.org/doi/abs/10.1126/science.aaa1394.

[12] F. Krausz, M. Ivanov, Rev. Mod. Phys. 81 (1) (2009) 163.

[13] X. Shi, C.-T. Liao, Z. Tao, E. Cating-Subramanian, M.M. Murnane, C. Hernández-García, H.C. Kapteyn, J. Phys. B, At. Mol. Opt. Phys. 53 (18) (2020) 184008, https://doi.org/10.1088/1361-6455/aba2fb.

[14] K. Midorikawa, Nat. Photonics 16 (4) (2022) 267–278, https://doi.org/10.1038/s41566-022-00961-9.

[15] K. Schafer, B. Yang, L. DiMauro, K. Kulander, Phys. Rev. Lett. 70 (11) (1993) 1599.

[16] M.B. Gaarde, J.L. Tate, K.J. Schafer, J. Phys. B, At. Mol. Opt. Phys. 41 (13) (2008) 132001, https://doi.org/10.1088/0953-4075/41/13/132001.

[17] P. Salieres, A. L'Huillier, M. Lewenstein, Phys. Rev. Lett. 74 (19) (1995) 3776–3779.

[18] P. Balcou, P. Salieres, A. L'Huillier, M. Lewenstein, Phys. Rev. A 55 (4) (1997) 3204.

[19] A. Rundquist, C.G. Durfee, Z. Chang, C. Herne, S. Backus, M.M. Murnane, H.C. Kapteyn, Science 280 (5368) (1998) 1412–1415, https://doi.org/10.1126/science.280.5368.1412, https://www.science.org/doi/pdf/10.1126/science.280.5368.1412, https://www.science.org/doi/abs/10.1126/science.280.5368.1412.

[20] F. Catoire, A. Ferré, O. Hort, A. Dubrouil, L. Quintard, D. Descamps, S. Petit, F. Burgy, E. Mével, Y. Mairesse, E. Constant, Phys. Rev. A 94 (2016) 063401, https://doi.org/10.1103/PhysRevA.94.063401, https://link.aps.org/doi/10.1103/PhysRevA.94.063401.

[21] L. Quintard, V. Strelkov, J. Vabek, O. Hort, A. Dubrouil, D. Descamps, F. Burgy, C. Péjot, E. Mével, F. Catoire, E. Constant, Sci. Adv. 5 (4) (2019) eaau7175, https://doi.org/10.1126/sciadv.aau7175, https://www.science.org/doi/pdf/10.1126/sciadv.aau7175, https://www.science.org/doi/abs/10.1126/sciadv.aau7175.

[22] H. Wikmark, C. Guo, J. Vogelsang, P.W. Smorenburg, H. Coudert-Alteirac, J. Lahl, J. Peschel, P. Rudawski, H. Dacasa, S. Carlström, S. Maclot, M.B. Gaarde, P. Johnsson, C.L. Arnold, A. L'Huillier, Proc. Natl. Acad. Sci. 116 (11) (2019) 4779–4787, https://doi.org/10.1073/pnas.1817626116, https://www.pnas.org/doi/pdf/10.1073/pnas.1817626116, https://www.pnas.org/doi/abs/10.1073/pnas.1817626116.

[23] L. Rego, N.J. Brooks, Q.L.D. Nguyen, J.S. Román, I. Binnie, L. Plaja, H.C. Kapteyn, M.M. Murnane, C. Hernández-García, Sci. Adv. 8 (5) (2022) eabj7380, https://doi.org/10.1126/sciadv.abj7380, https://www.science.org/doi/pdf/10.1126/sciadv.abj7380, https://www.science.org/doi/abs/10.1126/sciadv.abj7380.

[24] C. Hernández-García, Nat. Phys. 13 (4) (2017) 327–329, https://doi.org/10.1038/nphys4088.

[25] K.M. Dorney, L. Rego, N.J. Brooks, J. San Román, C.-T. Liao, J.L. Ellis, D. Zusin, C. Gentry, Q.L. Nguyen, J.M. Shaw, A. Picón, L. Plaja, H.C. Kapteyn, M.M. Murnane, C. Hernández-García, Nat. Photonics 13 (2) (2019) 123–130, https://doi.org/10.1038/s41566-018-0304-3.

[26] L. Rego, K.M. Dorney, N.J. Brooks, Q.L. Nguyen, C.-T. Liao, J.S. Román, D.E. Couch, A. Liu, E. Pisanty, M. Lewenstein, L. Plaja, H.C. Kapteyn, M.M. Murnane, C. Hernández-García, Science 364 (6447) (2019) eaaw9486, https://doi.org/10.1126/science.aaw9486, https://www.science.org/doi/pdf/10.1126/science.aaw9486, https://www.science.org/doi/abs/10.1126/science.aaw9486.

[27] G. Farkas, C. Tóth, Phys. Lett. A 168 (5) (1992) 447–450, https://doi.org/10.1016/0375-9601(92)90534-S, https://www.sciencedirect.com/science/article/pii/037596019290534S.

[28] I.P. Christov, M.M. Murnane, H.C. Kapteyn, Phys. Rev. Lett. 78 (1997) 1251–1254, https://doi.org/10.1103/PhysRevLett.78.1251, https://link.aps.org/doi/10.1103/PhysRevLett.78.1251.

[29] P.-M. Paul, E.S. Toma, P. Breger, G. Mullot, F. Augé, P. Balcou, H.G. Muller, P. Agostini, Science 292 (5522) (2001) 1689–1692.

[30] M. Hentschel, R. Kienberger, C. Spielmann, G.A. Reider, N. Milosevic, T. Brabec, P. Corkum, U. Heinzmann, M. Drescher, F. Krausz, Nature 414 (6863) (2001) 509–513.

[31] A. Fleischer, O. Kfir, T. Diskin, P. Sidorenko, O. Cohen, Nat. Publ. Group 8 (7) (2014) 543–549, https://doi.org/10.1038/nphoton.2014.108, http://www.nature.com/doifinder/10.1038/nphoton.2014.108.

[32] H. Eichmann, A. Egbert, S. Nolte, C. Momma, B. Wellegehausen, W. Becker, S. Long, J.K. McIver, Phys. Rev. A 51 (1995) R3414–R3417, https://doi.org/10.1103/PhysRevA.51.R3414, https://link.aps.org/doi/10.1103/PhysRevA.51.R3414.

[33] S. Long, W. Becker, J.K. McIver, Phys. Rev. A 52 (1995) 2262–2278, https://doi.org/10.1103/PhysRevA.52.2262, https://link.aps.org/doi/10.1103/PhysRevA.52.2262.

[34] C. Hernández-García, A. Picón, J. San Román, L. Plaja, Phys. Rev. Lett. 111 (2013) 083602, https://doi.org/10.1103/PhysRevLett.111.083602, https://link.aps.org/doi/10.1103/PhysRevLett.111.083602.

[35] G. Gariepy, J. Leach, K.T. Kim, T.J. Hammond, E. Frumker, R.W. Boyd, P.B. Corkum, Phys. Rev. Lett. 113 (2014) 153901, https://doi.org/10.1103/PhysRevLett.113.153901, https://link.aps.org/doi/10.1103/PhysRevLett.113.153901.

[36] R. Géneaux, A. Camper, T. Auguste, O. Gobert, J. Caillat, R. Taïeb, T. Ruchon, Nat. Commun. 7 (2016), arXiv:1509.07396, https://doi.org/10.1038/ncomms12583.

[37] I. Tikhomirov, T. Sato, K.L. Ishikawa, Phys. Rev. Lett. 118 (2017) 203202, https://doi.org/10.1103/PhysRevLett.118.203202, https://link.aps.org/doi/10.1103/PhysRevLett.118.203202.

[38] A. de las Heras, C. Hernández-García, L. Plaja, Phys. Rev. Res. 2 (2020) 033047, https://doi.org/10.1103/PhysRevResearch.2.033047, https://link.aps.org/doi/10.1103/PhysRevResearch.2.033047.

[39] N.H. Shon, A. Suda, K. Midorikawa, Phys. Rev. A 62 (2000) 023801, https://doi.org/10.1103/PhysRevA.62.023801, https://link.aps.org/doi/10.1103/PhysRevA.62.023801.

[40] E. Lorin, S. Chelkowski, A. Bandrauk, Comput. Phys. Commun. 177 (12) (2007) 908–932, https://doi.org/10.1016/j.cpc.2007.07.005, https://www.sciencedirect.com/science/article/pii/S0010465507003529.

[41] L.V. Keldysh, Zh. Èksp. Teor. Fiz. 47 (1964).

[42] F.H.M. Faisal, J. Phys. B, At. Mol. Phys. 6 (4) (1973) L89–L92.

[43] H.R. Reiss, Phys. Rev. A 22 (1980) 1786–1813.

[44] J. Pérez-Hernández, L. Roso, L. Plaja, Opt. Express 17 (12) (2009) 9891–9903.

[45] M. Lewenstein, P. Balcou, M.Y. Ivanov, A. L'huillier, P.B. Corkum, Phys. Rev. A 49 (3) (1994) 2117.

[46] E. Priori, G. Cerullo, M. Nisoli, S. Stagira, S. De Silvestri, P. Villoresi, L. Poletto, P. Ceccherini, C. Altucci, R. Bruzzese, C. de Lisio, Phys. Rev. A 61 (2000) 063801, https://doi.org/10.1103/PhysRevA.61.063801, https://link.aps.org/doi/10.1103/PhysRevA.61.063801.

[47] V. Tosa, K.T. Kim, C.H. Nam, Phys. Rev. A 79 (2009) 043828, https://doi.org/10.1103/PhysRevA.79.043828, https://link.aps.org/doi/10.1103/PhysRevA.79.043828.

[48] C. Hernández-García, J. Pérez-Hernández, J. Ramos, E.C. Jarque, L. Roso, L. Plaja, Phys. Rev. A 82 (3) (2010) 033432.

[49] E. Lorin, M. Lytova, A. Memarian, A.D. Bandrauk, J. Phys. A, Math. Theor. 48 (10) (2015) 105201, https://doi.org/10.1088/1751-8113/48/10/105201.

[50] M.V. Frolov, N.L. Manakov, T.S. Sarantseva, M.Y. Emelin, M.Y. Ryabikin, A.F. Starace, Phys. Rev. Lett. 102 (2009) 243901, https://doi.org/10.1103/PhysRevLett.102.243901, https://link.aps.org/doi/10.1103/PhysRevLett.102.243901.

[51] A.-T. Le, R.R. Lucchese, S. Tonzani, T. Morishita, C.D. Lin, Phys. Rev. A 80 (2009) 013401, https://doi.org/10.1103/PhysRevA.80.013401, https://link.aps.org/doi/10.1103/PhysRevA.80.013401.

[52] C. Jin, A.-T. Le, C.D. Lin, Phys. Rev. A 83 (2011) 023411, https://doi.org/10.1103/PhysRevA.83.023411, https://link.aps.org/doi/10.1103/PhysRevA.83.023411.

[53] A. Sánchez-González, P. Micaelli, C. Olivier, T.R. Barillot, M. Ilchen, A.A. Lutman, A. Marinelli, T. Maxwell, A. Achner, M. Agåker, N. Berrah, C. Bostedt, J.D. Bozek, J. Buck, P.H. Bucksbaum, S.C. Montero, B. Cooper, J.P. Cryan, M. Dong, R. Feifel, L.J. Frasinski, H. Fukuzawa, A. Galler, G. Hartmann, N. Hartmann, W. Helml, A.S. Johnson, A. Knie, A.O. Lindahl, J. Liu, K. Motomura, M. Mucke, C. O'Grady, J.-E. Rubensson, E.R. Simpson, R.J. Squibb, C. Såthe, K. Ueda, M. Vacher, D.J. Walke, V. Zhaunerchyk, R.N. Coffee, J.P. Marangos, Nat. Commun. 8 (1) (2017) 15461, https://doi.org/10.1038/ncomms15461.

[54] X. Liu, K. Amini, A. Sanchez, B. Belsa, T. Steinle, J. Biegert, Commun. Chem. 4 (1) (2021) 154, https://doi.org/10.1038/s42004-021-00594-z.

[55] N.I. Shvetsov-Shilovski, M. Lein, Phys. Rev. A 105 (2022) L021102, https://doi.org/10.1103/PhysRevA.105.L021102, https://link.aps.org/doi/10.1103/PhysRevA.105.L021102.

[56] T. Zahavy, A. Dikopoltsev, D. Moss, G.I. Haham, O. Cohen, S. Mannor, M. Segev, Optica 5 (5) (2018) 666–673, https://doi.org/10.1364/OPTICA.5.000666, https://opg.optica.org/optica/abstract.cfm?URI=optica-5-5-666.

[57] J. White, Z. Chang, Opt. Express 27 (4) (2019) 4799–4807.

[58] C. Brunner, A. Duensing, C. Schröder, M. Mittermair, V. Golkov, M. Pollanka, D. Cremers, R. Kienberger, Opt. Express 30 (9) (2022) 15669–15684.

[59] X. Liu, G. Zhang, J. Li, G. Shi, M. Zhou, B. Huang, Y. Tang, X. Song, W. Yang, Phys. Rev. Lett. 124 (2020) 113202, https://doi.org/10.1103/PhysRevLett.124.113202, https://link.aps.org/doi/10.1103/PhysRevLett.124.113202.

[60] A. Mihailescu, J. Instrum. 11 (12) (2016) C12004.

[61] M. Lytova, M. Spanner, I. Tamblyn, Can. J. Phys. (2020).

[62] A.M.M. Gherman, K. Kovács, M.V. Cristea, V. Toşa, Appl. Sci. 8 (11) (2018) 2106.

[63] L. Rego, J. San Román, A. Picón, L. Plaja, C. Hernández-García, Phys. Rev. Lett. 117 (16) (2016) 163202.

[64] J.D. Jackson, Classical Electrodynamics, 3rd edition, Wiley, New York, NY, 1999, http://cdsweb.cern.ch/record/490457.

[65] C. Hernández-García, J.S. Román, L. Plaja, A. Picón, New J. Phys. 17 (9) (2015) 093029, https://doi.org/10.1088/1367-2630/17/9/093029.

[66] M.D. Zeiler, D. Krishnan, G.W. Taylor, R. Fergus, in: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, IEEE, 2010, pp. 2528–2535.

[67] F. Chollet, Deep Learning with Python, Simon and Schuster, 2021.

[68] S.L. Smith, P.-J. Kindermans, C. Ying, Q.V. Le, Don't decay the learning rate, increase the batch size, arXiv preprint arXiv:1711.00489, 2017.

[69] A. Zaïr, M. Holler, A. Guandalini, F. Schapper, J. Biegert, L. Gallmann, U. Keller, A.S. Wyatt, A. Monmayrant, I.A. Walmsley, E. Cormier, T. Auguste, J.P. Caumes, P. Salières, Phys. Rev. Lett. 100 (2008) 143902, https://doi.org/10.1103/PhysRevLett.100.143902, https://link.aps.org/doi/10.1103/PhysRevLett.100.143902.

[70] L. Allen, M.W. Beijersbergen, R.J.C. Spreeuw, J.P. Woerdman, Phys. Rev. A 45 (1992) 8185–8189, https://doi.org/10.1103/PhysRevA.45.8185, https://link.aps.org/doi/10.1103/PhysRevA.45.8185.

[71] A. Forbes, M. de Oliveira, M.R. Dennis, Nat. Photonics 15 (4) (2021) 253–262, https://doi.org/10.1038/s41566-021-00780-4.

[72] K.Y. Bliokh, E. Karimi, M.J. Padgett, M.A. Alonso, M.R. Dennis, A. Dudley, A. Forbes, S. Zahedpour, S.W. Hancock, H.M. Milchberg, S. Rotter, F. Nori, S.K. Özdemir, N. Bender, H. Cao, P.B. Corkum, C. Hernández-García, H. Ren, Y. Kivshar, M.G. Silveirinha, N. Engheta, A. Rauschenbeutel, P. Schneeweiss, J. Volz, D. Leykam, D.A. Smirnova, K. Rong, B. Wang, E. Hasman, M.F. Picardi, A.V. Zayats, F.J. Rodríguez-Fortuño, C. Yang, J. Ren, A.B. Khanikaev, A. Alù, E. Brasselet, M. Shats, J. Verbeeck, P. Schattschneider, D. Sarenac, D.G. Cory, D. Pushin, M. Birk, A. Gorlach, I. Kaminer, F. Cardano, L. Marrucci, M. Krenn, F. Marquardt, Roadmap on structured waves (2023), https://doi.org/10.48550/ARXIV.2301.05349, https://arxiv.org/abs/2301.05349.

[73] A.K. Pandey, A. de las Heras, T. Larrieu, J. San Román, J. Serrano, L. Plaja, E. Baynard, M. Pittman, G. Dovillaire, S. Kazamias, C. Hernández-García, O. Guilbaud, ACS Photonics 9 (3) (2022) 944–951, https://doi.org/10.1021/acsphotonics.1c01768.

[74] F. Kong, C. Zhang, F. Bouchard, Z. Li, G.G. Brown, D.H. Ko, T.J. Hammond, L. Arissian, R.W. Boyd, E. Karimi, P.B. Corkum, Nat. Commun. 8 (1) (2017) 14970, https://doi.org/10.1038/ncomms14970.

[75] D. Gauthier, P.R. Ribič, G. Adhikary, A. Camper, C. Chappuis, R. Cucini, L.F. Dimauro, G. Dovillaire, F. Frassetto, R. Géneaux, P. Miotti, L. Poletto, B. Ressel, C. Spezzani, M. Stupar, T. Ruchon, G. De Ninno, Nat. Commun. 8 (1) (2017) 14971, https://doi.org/10.1038/ncomms14971.

[76] C. Hernández-García, A. Turpin, J.S. Román, A. Picón, R. Drevinskas, A. Cerkauskaite, P.G. Kazansky, C.G. Durfee, I. Nigo, J. Sola, Optica 4 (5) (2017) 520–526, https://doi.org/10.1364/OPTICA.4.000520, http://opg.optica.org/optica/abstract.cfm?URI=optica-4-5-520.

[77] A. Turpin, L. Rego, A. Picón, J. San Román, C. Hernández-García, Sci. Rep. 7 (1) (2017) 43888, https://doi.org/10.1038/srep43888.

[78] A. de las Heras, A.K. Pandey, J.S. Román, J. Serrano, E. Baynard, G. Dovillaire, M. Pittman, C.G. Durfee, L. Plaja, S. Kazamias, O. Guilbaud, C. Hernández-García, Optica 9 (1) (2022) 71–79, https://doi.org/10.1364/OPTICA.442304, https://opg.optica.org/optica/abstract.cfm?URI=optica-9-1-71.

[79] E. Pisanty, L. Rego, J. San Román, A. Picón, K.M. Dorney, H.C. Kapteyn, M.M. Murnane, L. Plaja, M. Lewenstein, C. Hernández-García, Phys. Rev. Lett. 122 (20) (2019) 203201, https://doi.org/10.1103/PhysRevLett.122.203201.

[80] M. Luttmann, M. Vimal, M. Guer, J.-F. Hergott, A.Z. Khoury, C. Hernández-García, E. Pisanty, T. Ruchon, Conservation of a half-integer angular momentum in nonlinear optics with a polarization Möbius strip, https://doi.org/10.48550/ARXIV.2209.00454, https://arxiv.org/abs/2209.00454, 2022.