

Removal of cross-phase modulation artifacts in ultrafast pump-probe dynamics by deep learning

Original

Removal of cross-phase modulation artifacts in ultrafast pump-probe dynamics by deep learning / Bresci, A., Guizzardi, M., Valensise, C.M., Marangi, F., Scotognella, F., Cerullo, G., Polli, D.. - In: APL PHOTONICS. - ISSN 2378-0967. - 6:7(2021). [10.1063/5.0057404]

Availability:

This version is available at: 11583/2985371 since: 2024-01-25T10:07:18Z

Publisher:

American Institute of Physics

Published

DOI:10.1063/5.0057404

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Removal of cross-phase modulation artifacts in ultrafast pump–probe dynamics by deep learning

Cite as: APL Photonics 6, 076104 (2021); <https://doi.org/10.1063/5.0057404>

Submitted: 20 May 2021 . Accepted: 06 July 2021 . Published Online: 19 July 2021

 A. Bresci,  M. Guizzardi, C. M. Valensise, F. Marangi,  F. Scotognella,  G. Cerullo, and  D. Polli



View Online



Export Citation




CrossMark



yttrium iron garnet glassy carbon beamsplitters fused quartz additive manufacturing
 zeolites III-IV semiconductors gallium lump copper nanoparticles organometallics
 nano ribbons barium fluoride europium phosphors photonics infrared dyes
 epitaxial crystal growth ultra high purity materials transparent ceramics CIGS
 cerium oxide polishing powder surface functionalized nanoparticles MRE grade materials thin film
 silver nanoparticles perovskites OLED lighting solar energy
 MOCVD beta-barium borate rare earth metals quantum dots sputtering targets fiber optics
 osmium scintillation Ce:YAG h-BN deposition slugs
 refractory metals laser crystals CVD precursors photovoltaics
 anode lithium niobate InAs wafers metamaterials borosilicate glass
 dysprosium pellets MOFs AuNPs YBCO superconductors InGaAs
 chalcogenides ZnS CdTe indium tin oxide MgF2 rutile
 perovskite crystals transparent ceramics diamond micropowder optical glass

The Next Generation of Material Science Catalogs



www.americanelements.com
 © 2001-2021, American Elements is a U.S. Registered Trademark



Removal of cross-phase modulation artifacts in ultrafast pump–probe dynamics by deep learning

Cite as: APL Photon. 6, 076104 (2021); doi: 10.1063/5.0057404

Submitted: 20 May 2021 • Accepted: 6 July 2021 •

Published Online: 19 July 2021



View Online



Export Citation



CrossMark

A. Bresci,  M. Guizzardi,  C. M. Valensise, F. Marangi, F. Scotognella,  G. Cerullo,  and D. Polli^{a)} 

AFFILIATIONS

Department of Physics, Politecnico di Milano, Piazza Leonardo da Vinci 32, 20133 Milan, Italy

^{a)} Author to whom correspondence should be addressed: dario.polli@polimi.it

ABSTRACT

Pump–probe spectroscopy is a gold standard technique to investigate ultrafast electronic dynamics of material systems. Pulsed laser sources employed to pump and probe samples feature typically high peak power, which may give rise to coherent artifacts under a wide range of experimental conditions. Among those, the Cross-Phase Modulation (XPM) artifact has gathered particular attention as it produces particularly high signal distortions, in some cases hiding a relevant portion of the dynamics of interest. Here, we present a novel approach for the removal of XPM coherent artifacts in ultrafast pump–probe spectroscopy, based on deep learning. We developed XPMnet, a convolutional neural network able to reconstruct electronic relaxation dynamics otherwise embedded in artifact distortions, thus enabling the retrieval of fundamental information to characterize the material system under investigation. We validated XPMnet on Indium Tin Oxide (ITO), a heavily doped semiconductor displaying a plasmon resonance in the near-infrared, which is a key material for the development of infrared plasmonic devices. Pump–probe measurements of ITO show strong XPM artifacts that overwhelm the electronic cooling dynamics of interest due to the low optical density of the material at near-infrared photon energies. XPMnet retrieved ITO electronic dynamics in excellent agreement with expected outcomes in terms of material-specific time constants. This artificial intelligence method constitutes a powerful solution for XPM artifact removal, providing high accuracy and short execution time. We believe that this model could be integrated in real time in pump–probe setups to increase the amount of information one can derive from ultrafast spectroscopy measurements.

© 2021 Author(s). All article content, except where otherwise noted, is licensed under a Creative Commons Attribution (CC BY) license (<http://creativecommons.org/licenses/by/4.0/>). <https://doi.org/10.1063/5.0057404>

I. INTRODUCTION

Ultrafast pump–probe spectroscopy has proven to be a powerful technique to study out-of-equilibrium phenomena, being applicable over a broad range of photon energies from THz to x rays.¹ In pump–probe, a medium is first excited with a short pump pulse and the photoinduced dynamics is probed by a time-delayed broadband probe pulse. The excitation pulses commonly employed are shorter than or close to 100 fs, which leads to peak intensities higher than 1 GW/cm². Such a condition may promote the generation of several Coherent Artifacts (CAs) of considerable intensity that can completely or partially distort the first hundreds of femtoseconds of relaxation dynamics, causing loss of information about early electronic processes under investigation. Different CAs can be recognized in ultrafast pump–probe measurements: Two-Photon Absorption (TPA), Cross-Phase Modulation (XPM), Stimulated Raman

Scattering (SRS),^{2,3} and Pump Perturbed Free Induction Decay (PPFID).^{4,5}

Here, we focus on XPM since it induces stronger distortions in relaxation dynamics compared to other artifacts and it is present across the whole probe spectrum. XPM was first reported in 1986 by Alfano *et al.*⁶ it originates from the redistribution of the spectral components of the probe pulse induced by the Kerr effect, namely, a change in the medium refractive index n caused by an intense pump pulse with intensity $I_{pu}(t)$, according to $n(t) = n_0 + n_2 \cdot I_{pu}(t)$. Such a rapid refractive index change modulates the phase shift experienced by the probe pulse and causes time-dependent shifts of its spectrum, which give rise to positive/negative differential transmission ($\Delta T/T$) signals at specific probe wavelengths. XPM-related distortions are unavoidable when employing glass substrates and when samples under investigation feature a low optical density in the range of the pump pulse.

There is an urgent need for the development of XPM removal methods in order to retrieve the maximum amount of information from pump–probe spectroscopy. Simply measuring the XPM on the bare glass substrate and then subtracting it from pump–probe signals at every wavelength cannot be a valid solution to get rid of signal distortions. In fact, the XPM signal is affected by the pump pulse absorption and by the redistribution of the probe spectral components so that its shape would be different with or without the sample.

Recent years saw a huge rise of Artificial Intelligence (AI) applications in technology and engineering. In photonics and optics, AI-driven approaches have been mainly employed for automated image processing,^{7–9} control of adaptive optics for aberration correction,¹⁰ wavefront shaping for computational imaging,¹¹ and self-optimization of nonlinear optical systems.^{12–14} The leading principle of AI is the idea that machines can be programmed to independently learn how to efficiently execute very complex tasks.

Among the branches of AI, Machine Learning (ML)¹⁵ relies on the fact that experience and large amounts of data can be exploited by machines to learn how to model and solve problems, acting as black-box architectures without the need for any explicitly coded instruction. Learning processes are generally grouped into supervised, unsupervised, and reinforcement learning.¹⁶ Supervised algorithms train on a labeled dataset featuring pairs of inputs and ideal outputs, aiming at finding a parametric transfer function from the input to the ground truth in the context of a classification or regression problem.¹⁷ Unsupervised algorithms train on an untagged dataset, with the purpose of automatically unveiling hidden patterns and exploiting them to group data into meaningful clusters. On the other hand, reinforcement learning takes place through the interaction of the machine with the environment: according to the quality of the actions an agent takes, it will gain a related positive or negative reward, thus shaping its knowledge from experience to take fruitful actions on its own.

Here, we introduce a novel AI-driven method to remove XPM artifacts from ultrafast pump–probe dynamics. We propose XPMnet, a Neural Network (NN) model able to operate directly on raw pump–probe data and efficiently retrieve the embedded electronic dynamics of physical interest in material systems. We designed and developed XPMnet as a supervised ML model. We structured the model architecture as a Convolutional Neural Network (CNN) trained on a labeled dataset, which consisted of simulated pairs of XPM-affected pump–probe instances and related electronic time dynamics, the latter carrying the physically relevant information about the specimen under investigation. CNNs fall into a peculiar branch of ML, namely, Deep Learning (DL),¹⁸ since their structure composed of sequential layers gives rise to a fairly deep model architecture. According to the universal approximation theorem proved by Cybenko in 1989,¹⁹ NNs are able to represent any kind of transfer function within an arbitrary tolerance. Among several generalizations of the theorem, Yarotsky demonstrated its validity also for the specific case of CNNs,²⁰ providing a solid mathematical proof of the approximation capabilities of these models.

We demonstrated the capability of XPMnet to learn in a supervised fashion how to extract electronic relaxation dynamics hidden by XPM artifacts with high accuracy and short execution time. We validated the model on Indium Tin Oxide (ITO), a key material for infrared plasmonics: the electronic relaxation dynamics retrieved

via XPMnet showed excellent agreement with expected outcomes in terms of material-specific time constants of the cooling process.

II. METHODS

A. Pump–probe setup

The pump–probe experiments were performed using a regeneratively amplified Ti:sapphire laser generating 100-fs pulses at 800 nm with >1 mJ energy and 1 kHz repetition rate. The pump beam resonant with the plasmonic absorption of ITO was generated by an Optical Parametric Amplifier (OPA) pumped by the second harmonic of the Ti:sapphire laser with output tuned at 1500 nm and a pulse duration of ~90 fs. To generate the broadband probe pulse, a small portion of the fundamental beam was focused into a 2 mm sapphire crystal, producing a white-light continuum that spanned from 420 to 730 nm. The pump–probe delay was varied by a computer-controlled mechanical delay line, and the differential transmission ($\Delta T/T$) spectrum of the probe was measured by a synchronized spectrometer detecting single-shot probe spectra.²¹

B. Sample preparation

ITO specimens were employed for the experimental validation of the DL architecture for pump–probe dynamics retrieval. The samples were fabricated by means of spin-coating starting from a commercial water dispersion of ITO nanocrystals (GetNano Materials, 99.99%, 20–30 nm, In₂O₃:SnO₂, 90:10 wt. %). Fabrication of the samples begins with the dilution of the water nanodispersion from 30% to 10%, which is then placed into an ultrasonic bath for 30 min to favor mixing and separation of aggregates. Glass slides are cut into small pieces and washed following a standard procedure with different solvents (i.e., water, acetone, and isopropyl alcohol). Before the deposition of the films, the substrates undergo an oxygen plasma treatment to reduce the contact angle of the dispersion and increase adhesion of the films. A small droplet of dispersion is placed onto the substrate, which is put into rotation at high speed. This will lead to the ejection of the solvent in excess, thinning of the layer of material, and evaporation of the residual solvent. To achieve complete evaporation of the solvent and obtain a more compact film, the samples undergo thermal annealing. The annealing treatment also allows us to increase the thickness of the layers with successive depositions of the same dispersion, preventing washing away of the already deposited film.

C. ITO ultrafast electronic dynamics

ITO is a plasmonic material with a plasmon resonance in the near-infrared (NIR) that can be tuned varying the doping level. Recently, hot electron extraction from ITO to different semiconductors has been demonstrated by Sakamoto *et al.*,²² thus justifying the increasing attention in NIR plasmonic properties of this material. Femtosecond pump–probe spectroscopy allows the study of carrier dynamics in ITO, but $\Delta T/T$ signals are typically affected by XPM artifacts distorting electronic dynamics due to the low optical density of the material at NIR photon energies.

ITO features a single exponential relaxation dynamics, as discussed in detail by Hartland.²³ After the excitation of the plasmon resonance with an infrared (1500 nm) pulse, four main processes take place. The first two, dephasing of the plasmon and

electron–electron scattering, are much faster than the typical time resolution (<100 fs). Electron–electron scattering generates a hot Fermi–Dirac carrier distribution, which eventually cools through electron–phonon scattering. This process is the one we are interested in measuring, and it is modeled with a mono-exponential decay dynamics. ITO has a non-parabolic conduction band,²⁴ which means that photoexcited electrons have a different effective mass. The change in the effective mass varies the plasma frequency so that the plasmon excitation causes an ultrafast change of the refractive index, as reported by different groups.^{25,26} The fourth step is the phonon–phonon scattering to cool down the lattice.

As shown in Fig. 1(a), pump–probe measurements on a bare 1-mm-thick glass slide display XPM distortions over the entire spectral range considered. Similarly, pump–probe measurements on a 100 nm thick ITO film made with two layers of nanocrystals deposited on the glass substrate show a clear XPM pattern covering a relevant portion of the underlying electronic exponential relaxation dynamics [Fig. 1(b)]. To better appreciate the artifact-related signal distortion, we display in Fig. 1(c) a single-wavelength of the $\Delta T/T$ maps: the whole dynamics, completed in less than 1 ps, is clearly affected by the artifact in the first 250 fs. We attribute the observed negative $\Delta T/T$ signal to an ultrafast increase in reflectivity due to the refractive index increase induced by hot electrons. This signal lasts as long as the electrons cool down by scattering with phonons.

D. Generation of the training dataset

One of the major issues in employing DL for the solution of experimental physics problems is the impossibility to rely only on measured data to train deep NNs due to the need of a very large amount of instances to compose the training dataset. Because of this reason, Data Augmentation (DA) became popular in DL models to enlarge the number of available instances.²⁷ The term DA refers to all the case-specific operations needed to generate an augmented but still physically meaningful dataset. In the case of the artifact-affected pump–probe dynamics here investigated, the training dataset includes noisy input instances featuring the XPM artifact embedding the exponential electronic relaxation, along with related

ideal output instances presenting the sole physically relevant electronic dynamics. The generation of the input counterparts of the dataset included the following steps: (i) the fit of experimentally measured XPM artifacts on glass substrates and their DA based on fit parameters; (ii) the simulation of pump–probe electronic temporal dynamics via the convolution between exponential decay dynamics and the Instrumental Response Function (IRF) related to the temporal duration of the pump pulse;²⁸ and (iii) the sum of XPM and pump–probe dynamics with the addition of white noise. The sole electronic dynamics simulated were employed as ground truth to pair with such inputs. The total amount of input–output pairs generated was 10^5 . The dataset was split as follows: 60% of the input–output pairs were employed for the CNN training, 20% for testing, and the remaining 20% were used to evaluate the model performance metrics on unseen data [i.e., mean squared error (MSE) and mean absolute percentage error (MAPE), R^2]. All the instances featured 200 temporal sampling points with a 5 fs sampling period so as to have a time window of 1 ps for each instance.

1. Simulation of XPM artifacts

A complete theoretical treatment of the XPM artifact has been proposed by Kovalenko *et al.*²⁹ Starting from the third-order polarization induced by the pump and probe electric fields, they described the probe as a chirped pulse with which the pump pulse interacts at a given time within a narrow-band spectral region. Assuming a Gaussian temporal profile for both pump and probe pulses, they derived a Gaussian model to fit the XPM. Different models were then proposed to best represent the artifact shape, making use of the sum of a Gaussian and its derivatives.² Among them, we employ the model developed by Baudisch,³⁰ which achieves a good XPM fit by using the sole first-order derivative of the Gaussian,

$$XPM(t, \lambda) = \cos(B(\lambda)(t - t_0(\lambda))^2 + \Phi(\lambda)) \cdot \left\{ A_0(\lambda) \exp\left[\frac{-4 \ln 2 (t - t_0(\lambda))^2}{\tau_1(\lambda)^2}\right] - A_1(\lambda) \frac{8 \ln 2}{\tau_1(\lambda)^2} \cdot (t - t_0(\lambda)) \exp\left[\frac{-4 \ln 2 (t - t_0(\lambda))^2}{\tau_1(\lambda)^2}\right] \right\}. \quad (1)$$




FIG. 1. (a) XPM time evolution measured on a 1-mm-thick soda-lime glass slide. (b) Electronic dynamics of ITO on soda-lime glass superimposed with the XPM artifact. The colorbar refers to $\Delta T/T$. (c) Single-wavelength time evolution showing the presence of the XPM CA distorting the early ITO electronic relaxation dynamics compared to the sole XPM signal on the glass substrate.

The parameters space of Eq. (1) includes the amplitudes $A_0(\lambda)$ and $A_1(\lambda)$, the full width half maximum (FWHM) duration of the Gaussian $\tau_1(\lambda)$ (fs), and $B(\lambda)$ (rad/s²) and $\Phi(\lambda)$ (rad) to reproduce artifact fringes generated when the pump pulses are shorter than 25 fs. Finally, $t_0(\lambda)$ (fs) is a wavelength-dependent delay due to the chirp of the probe pulse that causes a temporal shift of the overlap between pump and probe pulses. All the parameters exhibit wavelength dependence but are not related to physical quantities: they only provide a fitting function that does not always guarantee a realistic simulation of XPM artifacts by any arbitrary parameter variation. Hence, in order to obtain a realistic dataset, DA was performed starting from pump–probe measurements on standard 1-mm-thick soda-lime glass substrates. Different pump intensities were employed to enlarge the amount of experimental measurements at disposal. A total amount of 695 measured artifacts were fit via Particle Swarm Optimization (PSO).³¹ DA was applied on sets of fit parameters: it consisted of a $\pm 5\%$ random shift, along with the variation of $t_0(\lambda)$ from 150 to 700 fs (corresponding to the typical time zero shifting range observed experimentally employing a chirped probe pulse), thus allowing the simulation of 10^5 realistic and experimentally derived XPM instances.

2. Simulation of electronic relaxation dynamics

As discussed above, the only two processes we can detect in ITO with our time resolution are electron–phonon scattering and phonon–phonon scattering. The electron–phonon scattering in ITO is fast (< 1 ps), and it can be modeled as a mono-exponential decay. The phonon–phonon scattering process is much slower (> 10 ps) so that a constant offset is added to the exponential electron–phonon relaxation to take it into account. Hence, ITO mono-exponential electronic dynamics were modeled as follows:

$$y(t, \lambda) = \left[A_2(\lambda) \exp\left[\frac{-(t - t_1(\lambda))}{\tau_2(\lambda)}\right] + C \right] \cdot H(t_1(\lambda)) \otimes \text{IRF}. \quad (2)$$

In Eq. (2), $A_2(\lambda)$ is the signal amplitude, $\tau_2(\lambda)$ (fs) is the decay time constant, C accounts for the remaining signal after relaxation, $H(t_1(\lambda))$ is the Heaviside function centered in $t_1(\lambda)$ (fs), and the wavelength-dependent time zero between pump and probe pulses. The exponential function is convolved (i.e., \otimes) with the IRF, given by a Gaussian-shaped pulse with known FWHM. To generate the simulated dataset, all the parameters were randomly varied within typical experimental ranges, as reported in Table I.

Time zeros $t_0(\lambda)$ and $t_1(\lambda)$ were chosen to be temporally superimposed with a tolerance of $\pm 5\%$ to account for a temporal offset

TABLE I. Parameters values for $y(t, \lambda)$. Ranges were chosen to accommodate a variety of typical experimental conditions.

Parameter	Range
$A_2(\lambda)$	$-1.3 \cdot 10^{-2}$ to $-1.7 \cdot 10^{-3} \Delta T/T$
$\tau_2(\lambda)$	150–300 fs
C	0 to $-10^{-5} \Delta T/T$
$t_1(\lambda)$	150–700 fs
Noise	$3 \cdot 10^{-5}$ – $10^{-4} \Delta T/T$
IRF	70–100 fs FWHM

between the XPM and the electronic relaxation dynamics. In fact, XPM is generated in the glass substrate while the signal is generated in the sample under investigation: the distance between the two explains this time difference. Training pairs were normalized in the range $[0, 1]$, which increases the model accuracy and accelerates convergence. A baseline shifting of a maximum of 30% was applied in order to train the network to keep up with a modest baseline variation in the normalized inputs.

E. CNN model architecture

The model architecture was designed and developed as a CNN operating on single-wavelength pump–probe instances: convolutional layers process the input as feature extractors, whereas fully connected layers receive convolutional feature maps to compute the final prediction (Fig. 2). The convolutional stage of the CNN includes the following layers: 128 (32,1)-shaped kernels, 96 (24,1)-shaped kernels, 64 (8,1)-shaped kernels, and eventually three layers of 8 (3,1)-shaped kernels each. The following fully connected stage features a 64-neurons layer, two 36-neurons layers, and a final 200-neurons output layer. The output neurons store numerical values representing the retrieved electronic relaxation dynamics. In the proposed model, the total amount of trainable parameters (θ) (i.e., weights and biases of kernels and neurons) was $1.3 \cdot 10^6$. The non-linear activation function (σ) for all layers was chosen as a Rectified Linear Unit (ReLU),

$$\sigma(z) = \max(0, z), \quad (3)$$

where z is the linear output of a single kernel in convolutional layers or of a single neuron in fully connected layers. Through the training process, the CNN algorithm predicts electronic dynamics y_{pred} by applying the current parametric transfer function to the input, and, thanks to the ideal output y_{true} provided in the training dataset, it computes the distance between prediction and the ground truth. Such a distance is quantified by means of a loss function $L(\theta)$, which was here chosen to be the MAPE,

$$L(\theta) = \left\langle 100 \frac{|y_{true} - y_{pred}|}{y_{true}} \right\rangle. \quad (4)$$

The model cost function $J(\theta)$ is then obtained by averaging the loss $L(\theta)$ of single training pairs over the N_{train} pairs of a mini-batch [Eq. (5)].³² Mini-batches comprised 128 input–output pairs, a batch size that accelerated and regularized convergence along with batch-normalization at the CNN input,

$$J(\theta) = \frac{1}{N_{train}} \sum_{n=0}^{N_{train}} L_n(\theta). \quad (5)$$

The goal of the algorithm is to find a set of parameters that minimizes $J(\theta)$, thus maximizing the model performance accuracy by leading to a minimum average distance between ideal and predicted outputs. According to the mini-batch gradient descent method, the opposite direction of the gradient of $J(\theta)$ is used to update the current parameters set θ_t , taking steps proportional to the learning rate α ,

$$\theta_{t+1} = \theta_t - \alpha \nabla J(\theta_t). \quad (6)$$

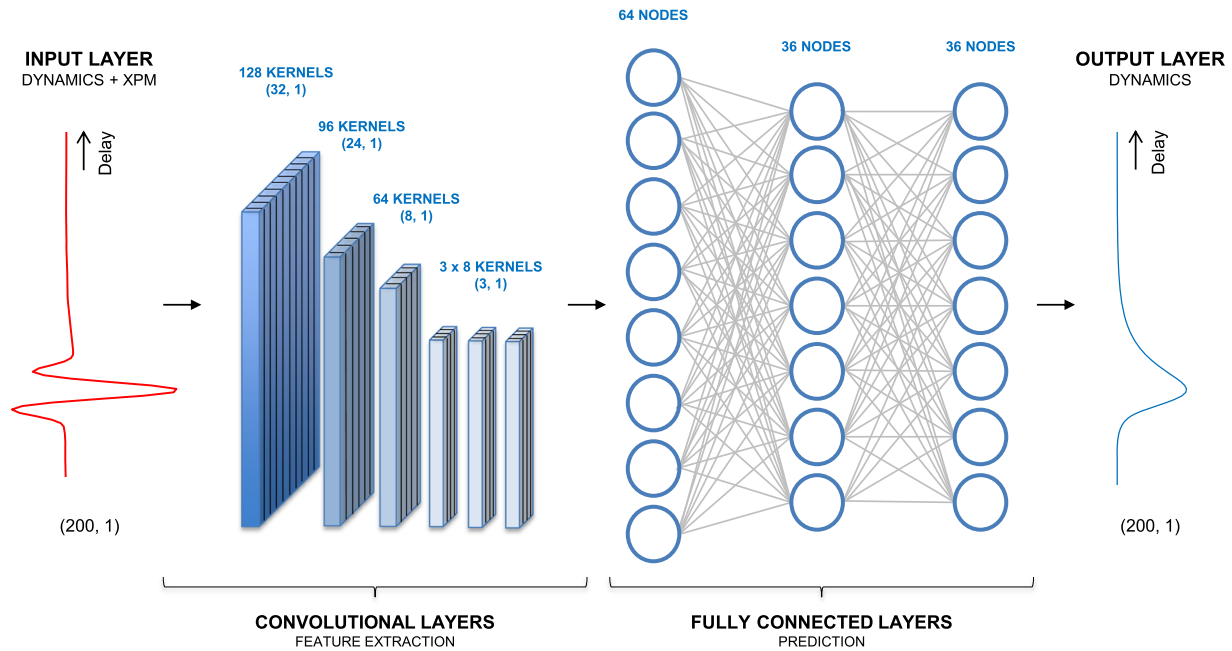


FIG. 2. XPMnet model architecture. The noisy XPM-distorted pump–probe dynamics is fed into the network as a column vector; the output layer presents the retrieved electronic relaxation dynamics. The feed-forward architecture consists in pyramid-shaped convolutional layers for feature extraction, followed by fully connected layers to predict the output from convolutional feature maps.

The XPMnet learning rate was tuned according to the adaptive moment estimation optimization technique (ADAM).³³ In order to avoid overfitting on training data and increase the CNN generalization capability, a 35% drop-out³⁴ was applied ahead of the output layer and L_2 weight regularization of fully connected layers³⁵ was introduced with 10^{-1} as the regularization factor. The model was trained for 100 epochs. The entire algorithm was developed in Python employing the TensorFlow³⁶ platform and Keras³⁷ library. For further details about the algorithm implementation, the reader can find the code available online.³⁸

III. RESULTS

A. XPMnet performance

Upon training, the XPMnet featured excellent figures of merit: a MSE of $5 \cdot 10^{-5}$, a MAPE of 1%, and an R^2 of 0.99 on unseen data. The training process took 8 s per epoch, whereas the time required for XPM artifact removal, averaged over 10^5 instances, was $3 \cdot 10^{-2}$ s. A Tesla K80 graphics processing unit (GPU) was employed.

The results of applying the XPMnet algorithm on simulated and experimental instances of noisy XPM-affected pump–probe dynamics are reported in Figs. 3(a)–3(c) for simulated and Figs. 3(d)–3(f) for experimental instances. As expected, the XPM artifact highly distorts and covers the relaxation dynamics in pump–probe measurements. The CNN model is able to efficiently and accurately retrieve the CA and subtract it in order to extract the relevant dynamics: the difference between the simulated ideal output and the network prediction is negligible ($MSE \sim 10^{-5}$), as it can be appreciated in

Figs. 3(a)–3(c). Thanks to the flexibility of the DA process to handle several scenarios, the CNN performs efficiently regardless of the wavelength of the probe beam. This is crucial when processing sequential measurements with chirped laser pulses: $t_0(\lambda)$ and $t_1(\lambda)$ must not affect the network performance. The model reconstructs exponential relaxation dynamics from which decay time constants τ can be easily derived by fitting, and it is able to provide a correct baseline for the output so that the prediction can be perfectly superimposed over the input.

The results of applying XPMnet on experimental measurements are shown in Figs. 3(d)–3(f). We employed different ITO samples to test the model: a single-layer film of ITO nanocrystals [Fig. 3(d)], a four-layer film of ITO nanocrystals [Fig. 3(e)], and bulk sputtered ITO [Fig. 3(f)]. The cooling dynamics were reconstructed via XPMnet by isolating them from artifact and noise-affected experimental data. The overlap of predicted dynamics with input instances is perfectly managed by the algorithm in terms of baseline and $t_0(\lambda)$ and $t_1(\lambda)$ variations. When employing thicker ITO samples, a higher absorption causes a decrease of the artifact amplitude with respect to the electron dynamics. When varying the ITO thickness from one layer ~ 40 nm [Fig. 3(d)] to four layers ~ 160 nm [Fig. 3(e)], the CNN predicts relaxation dynamics with the same material-specific τ of 190 fs, despite the signals show different amplitudes along with different starting times. Interestingly, the time constant retrieved by the XPMnet for bulk sputtered ITO [~ 200 nm, Fig. 3(f)] is slightly shorter ($\tau = 150$ fs), in agreement with the fact that the size and shape of ITO nanoparticles affect the electronic cooling and modify the decay time.³⁹ The values of τ obtained by fitting the XPMnet predictions with mono-exponential




FIG. 3. (a)–(c) Results of applying XPMnet on unseen simulated data. Simulations were performed in the parameter ranges reported in Table I. XPMnet approximates the ground truth of simulated dynamics not employed in the training process with high accuracy, as it can be clearly seen from the superposition of green (ground truth) and orange (prediction) curves. (d)–(f) Real experimental data of ITO samples on the soda-lime substrate at 500 nm (d) and 600 nm (e) and (f) probe wavelength. XPMnet applied on real experimental data retrieves the embedded relaxation dynamics of electrons, unveiling characteristic time constants [(d) and (e) $\tau = 190$ fs, (f) $\tau = 150$ fs] in agreement with the physics of the material system under investigation.

curves are thus in very good agreement with expected outcomes. Figure 3(d) displays an XPM that is split due to group velocity mismatch between pump and probe: when they travel with a significant velocity difference due to the dispersion of the material, their relative position changes along the medium. Such a condition causes a broadening of the interaction time interval of the pulses in the glass substrate, thus inducing the splitting of the artifact. In our training dataset, such an artifact was not present, but nonetheless our model was able to manage this complex input shape and retrieve the correct decay dynamics.

The AI-driven model here presented is able to deal with a variety of experimental conditions in terms of absorption coefficients, artifact and pump–probe dynamics amplitude ratio, probe pulse chirp causing $t_0(\lambda)$ and $t_1(\lambda)$ variations, as well as baseline shifting due to a case-specific input normalization. In addition, it is able to provide a quick and highly accurate prediction, thus increasing in real time the amount of information available from pump–probe measurements.

B. Convolutional layers: Un-boxing and interpretation

It can be argued that CNN models, despite being extremely powerful for the solution of complex non-linear physical problems,

act in a black-box manner: the user’s knowledge about how exactly CNNs process the input to achieve the final result is very limited. It is well known from the literature that higher complexity features can be extracted by deeper convolutional layers: progressively more elaborated patterns are pointed out as the input traverses sequentially deeper kernels.⁴⁰ Similarly, feature maps generated by deep convolutional layers are very unintuitive and can hardly be interpreted. In the present work, we investigated how kernels in convolutional layers process the input and unveil its most relevant patterns and their location. For this purpose, once the model is trained, its architecture is cut so that the feature maps of selected convolutional layers are given as output. Such un-boxed feature maps are thus ranked by their activation level, considering more active the ones featuring higher absolute values.

Figure 4 displays the most active feature maps of the first and second convolutional layers. The XPMnet feature extraction logic in the first convolutional layer can be associated both with the artifact and the relaxation dynamics. The maps highlight with spatial correspondence the regions of the input featuring fairly positive [Fig. 4(b)] or negative [Fig. 4(c)] values of the second derivative of the signal time evolution. On the other hand, the map in Fig. 4(d) shows active regions in correspondence with the spatial location of the electronic dynamics, interestingly unveiling the rising and decaying shapes in this early convolutional stage. In the second

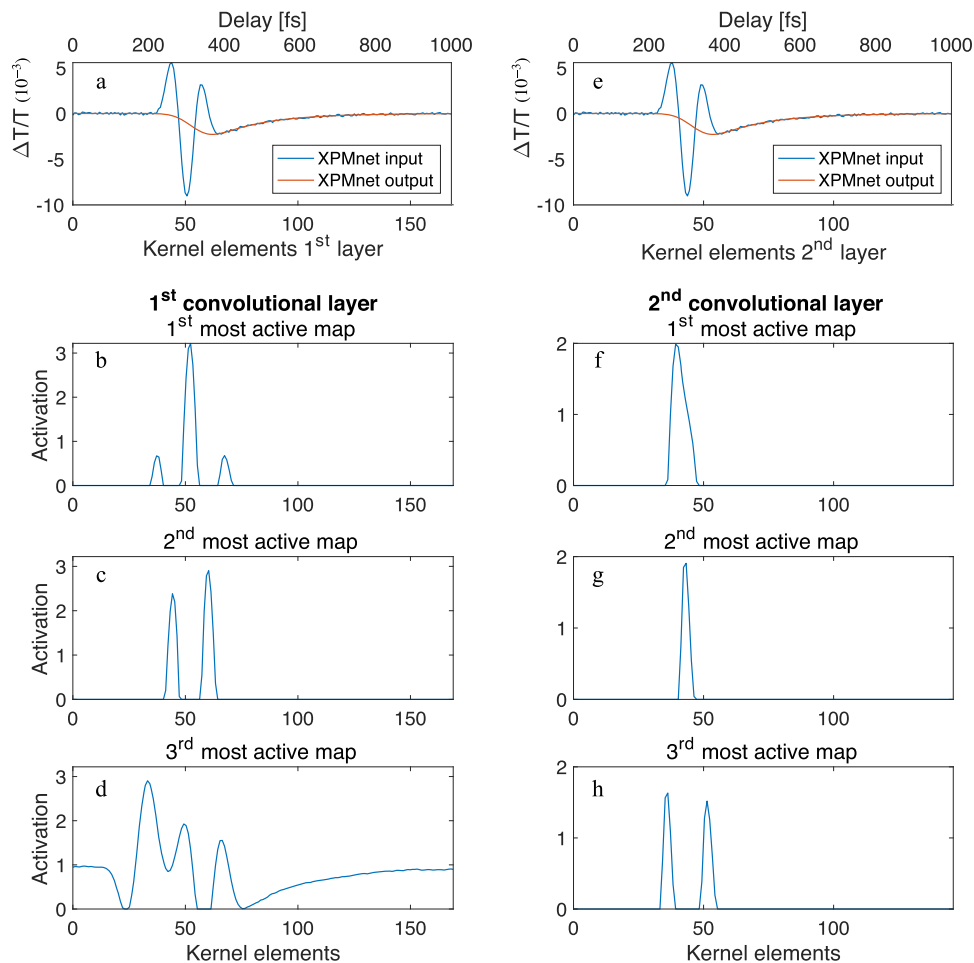


FIG. 4. Convolutional layers un-boxing. (a) and (e) XPMnet input and output signals considered for the extraction of feature maps. Kernel elements of the first (a) and second (e) convolutional layers are indicated to ease their relation with delay times. Early convolutional layers [i.e., first layer [(b)–(d)] and second layer [(f)–(h)]] were investigated in terms of feature extraction logic. (b) and (c) First layer maps pointing out regions of the $\Delta T/T$ signal with a fairly positive and negative second derivative, respectively. (d) First layer map showing the rising and decay shapes of the electronic dynamics. (f) and (g) Second layer maps pointing out the region with lower $\Delta T/T$ signal. (h) Second layer map pointing out the regions with higher $\Delta T/T$ signal.

convolutional layer, most active kernels act as intensity filters: feature maps point out with spatial correspondence the regions of the input instance featuring relatively low [Figs. 4(f) and 4(g)] or high [Fig. 4(h)] intensity values.

The interpretation of the filtering principles employed by kernels upon training is not meant to be rigorous. Nevertheless, such an approach can be used to gain some understanding of the ratio underlying the black-box architecture of DL models. It can be exploited to explicitly visualize, test, and assess the model feature extraction capability.

IV. CONCLUSION

We proposed a novel AI-driven method—XPMnet—for the removal of XPM CAs in ultrafast pump–probe spectroscopy signals, providing a powerful tool to reveal electronic relaxation

dynamics otherwise highly distorted in a wide variety of common experimental conditions. Indeed, ordinary glass substrates for ultrafast spectroscopy experiments, as well as samples with low optical density in the spectral range of the pump, typically show a relevant XPM-related distortion in the pump–probe dynamics measured. XPMnet is able to process noisy artifact-affected pump–probe data and retrieve embedded exponential electronic dynamics of physical interest for the characterization of the material system.

The DL model here presented performs with high accuracy ($R^2 = 0.99$, $MSE = 5 \cdot 10^{-5}$) in a very short execution time ($3 \cdot 10^{-2}$ s): it could be easily integrated into pump–probe spectroscopy setups to process data in real time for artifact removal. The method can be adapted with ease to a variety of experimental conditions by tuning the time resolution of the data, the XPM fit, and the electronic dynamics simulations. In particular, future research could investigate the handling of non-exponential and/or

multi-exponential electronic dynamics to generalize the applicability of the model on a broad diversity of material systems. We believe that XPMnet could serve to significantly boost the power of ultrafast pump-probe spectroscopy by increasing the amount and quality of information one can derive from the measured signals.

AUTHORS' CONTRIBUTIONS

A.B. and M.G. contributed equally to this work.

ACKNOWLEDGMENTS

This project has received funding from the European Union project CRIMSON under Grant Agreement No. 101016923 and from the Regione Lombardia project NEWMED under Grant Agreement No. POR FESR 2014–2020.

DATA AVAILABILITY

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

REFERENCES

- 1 M. Maiuri, M. Garavelli, and G. Cerullo, "Ultrafast spectroscopy: State of the art and open challenges," *J. Am. Chem. Soc.* **142**, 3–15 (2019).
- 2 M. Lorenc, M. Ziolk, R. Naskrecki, J. Karolczak, J. Kubicki, and A. Maciejewski, "Artifacts in femtosecond transient absorption spectroscopy," *Appl. Phys. B* **74**, 19–27 (2002).
- 3 K. Ekvall, P. Van Der Meulen, C. Dhollande, L.-E. Berg, S. Pommeret, R. Naskrecki, and J.-C. Mialocq, "Cross phase modulation artifact in liquid phase transient absorption spectroscopy," *J. Appl. Phys.* **87**, 2340–2352 (2000).
- 4 P. Hamm, "Coherent effects in femtosecond infrared spectroscopy," *Chem. Phys.* **200**, 415–429 (1995).
- 5 C. H. Brito Cruz, J. P. Gordon, P. C. Becker, R. L. Fork, and C. V. Shank, "Dynamics of spectral hole burning," *IEEE J. Quantum Electron.* **24**, 261–269 (1988).
- 6 R. R. Alfano, Q. X. Li, T. Jimbo, J. T. Manassah, and P. P. Ho, "Induced spectral broadening of a weak picosecond pulse in glass produced by an intense picosecond pulse," *Opt. Lett.* **11**, 626–628 (1986).
- 7 J. Yang, J. Xu, X. Zhang, C. Wu, T. Lin, and Y. Ying, "Deep learning for vibrational spectral analysis: Recent progress and a practical guide," *Anal. Chim. Acta* **1081**, 6–17 (2019).
- 8 K. Ghosh, A. Stuke, M. Todorović, P. B. Jørgensen, M. N. Schmidt, A. Vektari, and P. Rinke, "Deep learning spectroscopy: Neural networks for molecular excitation spectra," *Adv. Sci.* **6**, 1801367 (2019).
- 9 S. Li, W. Song, L. Fang, Y. Chen, P. Ghamisi, and J. A. Benediktsson, "Deep learning for hyperspectral image classification: An overview," *IEEE Trans. Geosci. Remote Sens.* **57**, 6690–6709 (2019).
- 10 S. L. S. Gómez, C. González-Gutiérrez, E. D. Alonso, J. D. Santos, M. L. S. Rodríguez, T. Morris, J. Osborn, A. Basden, L. Bonavera, J. G.-N. González, and F. J. de Cos Juez, "Experience with artificial neural networks applied in multi-object adaptive optics," *Publ. Astron. Soc. Pac.* **131**, 108012 (2019).
- 11 G. Barbastathis, A. Ozcan, and G. Situ, "On the use of deep learning for computational imaging," *Optica* **6**, 921 (2019).
- 12 G. Genty, L. Salmela, J. M. Dudley, D. Brunner, A. Kokhanovskiy, S. Kobtsev, and S. K. Turitsyn, "Machine learning and applications in ultrafast photonics," *Nat. Photonics* **15**, 91–101 (2020).
- 13 C. M. Valensise, A. Giuseppi, F. Vernuccio, A. De la Cadena, G. Cerullo, and D. Polli, "Removing non-resonant background from CARS spectra via deep learning," *APL Photonics* **5**, 061305 (2020).
- 14 C. M. Valensise, A. Giuseppi, G. Cerullo, and D. Polli, "Deep reinforcement learning control of white-light continuum generation," *Optica* **8**, 239 (2021).
- 15 P. Mehta, M. Bukov, C.-H. Wang, A. G. R. Day, C. Richardson, C. K. Fisher, and D. J. Schwab, "A high-bias, low-variance introduction to machine learning for physicists," *Phys. Rep.* **810**, 1–124 (2019).
- 16 T. O. Ayodele, "Types of machine learning algorithms," *New Adv. Mach. Learn.* **3**, 19–48 (2010).
- 17 R. Caruana and A. Niculescu-Mizil, "An empirical comparison of supervised learning algorithms," in *Proceedings of the 23rd International Conference on Machine Learning, ICML 06* (ACM Press, 2006).
- 18 F. Emmert-Streib, Z. Yang, H. Feng, S. Tripathi, and M. Dehmer, "An introductory review of deep learning for prediction models with big data," *Front. Artif. Intell.* **3**, 4 (2020).
- 19 G. Cybenko, "Approximation by superpositions of a sigmoidal function," *Math. Control Signals Syst.* **2**, 303–314 (1989).
- 20 D. Yarotsky, "Universal approximations of invariant maps by neural networks," *Constr. Approx.*, 1–68 (2021).
- 21 D. Polli, L. Lüer, and G. Cerullo, "High-time-resolution pump-probe system with broadband detection for the study of time-domain vibrational dynamics," *Rev. Sci. Instrum.* **78**, 103108 (2007).
- 22 M. Sakamoto, T. Kawawaki, M. Kimura, T. Yoshinaga, J. J. M. Vequizo, H. Matsunaga, C. S. K. Ranasinghe, A. Yamakata, H. Matsuzaki, A. Furube *et al.*, "Clear and transparent nanocrystals for infrared-responsive carrier transfer," *Nat. Commun.* **10**, 1879 (2019).
- 23 G. V. Hartland, "Optical studies of dynamics in noble metal nanostructures," *Chem. Rev.* **111**, 3858–3887 (2011).
- 24 P. Guo, R. D. Schaller, J. B. Ketterson, and R. P. H. Chang, "Ultrafast switching of tunable infrared plasmons in indium tin oxide nanorod arrays with large absolute amplitude," *Nat. Photonics* **10**, 267–273 (2016).
- 25 M. Guizzardi, S. Bonfadini, L. Moscardi, I. Kriegel, F. Scotognella, and L. Criante, "Large scale indium tin oxide (ITO) one dimensional gratings for ultrafast signal modulation in the visible spectral region," *Phys. Chem. Chem. Phys.* **22**, 6881–6887 (2020).
- 26 M. A. Blemker, S. L. Gibbs, E. K. Raulerson, D. J. Milliron, and S. T. Roberts, "Modulation of the visible absorption and reflection profiles of ITO nanocrystal thin films by plasmon excitation," *ACS Photonics* **7**, 1188–1196 (2020).
- 27 L. Taylor and G. Nitschke, "Improving deep learning with generic data augmentation," *2018 IEEE Symposium Series on Computational Intelligence (SSCI)* (IEEE, 2018), pp. 1542–1547.
- 28 D. Polli, D. Brida, S. Mukamel, G. Lanzani, and G. Cerullo, "Effective temporal resolution in pump-probe spectroscopy with strongly chirped pulses," *Phys. Rev. A* **82**, 053809 (2010).
- 29 S. A. Kovalenko, A. L. Dobryakov, J. Ruthmann, and N. P. Ernstring, "Femtosecond spectroscopy of condensed phases with chirped supercontinuum probing," *Phys. Rev. A* **59**, 2369 (1999).
- 30 B. Baudisch, "Time resolved broadband spectroscopy from UV to NIR," Ph.D. thesis, LMU, 2018.
- 31 J. Kennedy and R. Eberhart, "Particle swarm optimization," in *Proceedings of ICNN'95-International Conference on Neural Networks* (IEEE, 1995), Vol. 4, pp. 1942–1948.
- 32 M. Li, T. Zhang, Y. Chen, and A. J. Smola, "Efficient mini-batch training for stochastic optimization," in *Proceedings of the 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (ACM, 2014), pp. 661–670.
- 33 D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv:1412.6980* (2014).
- 34 N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: A simple way to prevent neural networks from overfitting," *J. Mach. Learn. Res.* **15**, 1929–1958 (2014).
- 35 A. Y. Ng, "Feature selection, L_1 vs. L_2 regularization, and rotational invariance," in *Twenty-First International Conference on Machine Learning, ICML 04* (ACM Press, 2004).
- 36 M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner,

I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, "TensorFlow: Large-scale machine learning on heterogeneous systems," software available from <https://www.tensorflow.org/>, 2015.

³⁷F. Chollet, "keras," <https://github.com/fchollet/keras>, 2015.

³⁸A. Bresci, "XPMnet," <https://github.com/ariannabresci/XPMnet>, 2021.

³⁹Y. Nam, L. Li, J. Y. Lee, and O. V. Prezhdo, "Size and shape effects on charge recombination dynamics of TiO₂ nanoclusters," *J. Phys. Chem. C* **122**, 5201–5208 (2018).

⁴⁰J. Gu, Z. Wang, J. Kuen, L. Ma, A. Shahroudy, B. Shuai, T. Liu, X. Wang, G. Wang, J. Cai, and T. Chen, "Recent advances in convolutional neural networks," *Pattern Recognit.* **77**, 354–377 (2018).