

Physics-Guided and Physics-Explainable Recurrent Neural Network for Time Dynamics in Optical Resonances

Yingheng Tang

Purdue University

Jichao Fan

University of Utah

Xinwei Li

California Institute of Technology

Jianzhu Ma

Peking University <https://orcid.org/0000-0002-8236-6609>

Minghao Qi

Purdue University

Cunxi Yu

University of Utah

Weilu Gao (✉ weilu.gao@utah.edu)

The University of Utah <https://orcid.org/0000-0003-3139-034X>

Article

Keywords: recurrent neural network, machine learning, time dynamics

Posted Date: September 20th, 2021

DOI: <https://doi.org/10.21203/rs.3.rs-879817/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Version of Record: A version of this preprint was published at Nature Computational Science on March 28th, 2022. See the published version at <https://doi.org/10.1038/s43588-022-00215-2>.

Physics-Guided and Physics-Explainable Recurrent Neural Network for Time Dynamics in Optical Resonances

Yingheng Tang,^{1,2} Jichao Fan,¹ Xinwei Li,³ Jianzhu Ma,⁴ Minghao Qi,² Cunxi Yu,^{1,*} and Weilu Gao^{1,*}

¹*Department of Electrical and Computer Engineering, University of Utah, Salt Lake City, UT 84112, USA*

²*Department of Electrical and Computer Engineering, Purdue University, West Lafayette, IN 47907, USA*

³*Division of Physics, Mathematics and Astronomy, California Institute of Technology, Pasadena, CA 91125, USA*

⁴*Institute of Artificial Intelligence, Peking University, Beijing, China*

*To whom correspondence should be addressed; E-mail: cunxi.yu@utah.edu; weilu.gao@utah.edu.

Understanding the time evolution of physical systems is crucial to revealing fundamental characteristics that are hidden in frequency domain. In optical science, high-quality resonance cavities and enhanced interactions with matters are at the heart of modern quantum technologies. However, capturing their time dynamics in real-world scenarios suffers from long data acquisition and low analysis accuracy due to slow convergence and limited time window. Here, we report a physics-guided and physics-explainable recurrent neural network to precisely forecast the time-domain response of resonance features with the shortest acquired input sequence being 7% of full length, and to infer corresponding resonance frequencies. The model is trained in a two-step multi-fidelity framework for high-accuracy forecast, where the first step is based on a large amount of low-fidelity physical-model-generated synthetic data and second step involves a small set of high-fidelity application-oriented observational data. Through both simulations and experiments, we demonstrate that the model is universally applicable to a wide range of resonances, including dielectric metasurfaces, graphene plasmonics, and ultrastrongly coupled Landau polaritons, where our model accurately captures small signal features and learns essential physical quantities. The demonstrated machine learning algorithm offers a new way to accelerate the exploration of physical phenomena and the design of devices under resonance-enhanced light-matter interaction.

Forecasting time dynamics has been a key area in many contexts of scientific research and commercial decision making, such as climate modeling¹, medical science², and finance³. Classical parametric models informed by domain expertise, such as autoregressive integrated moving average⁴, enjoy merits of model simplicity and easy solutions while suffer from large prediction errors in strongly nonlinear and stochastic data. Machine learning approaches provide new perspective of learning time dynamics and underlying complex representation in a data-driven manner without fixed parameters or structures⁵. Especially, given the impressive success of recurrent neural networks (RNNs) in natural language processing and logical interpretation of time series as sequences, various RNN architectures, such as long short-term memory and gated recurrent unit (GRU), have been employed for time series forecast applications⁶⁻⁸. However, the overfitting and local minimum problems often lead to low model efficacy⁹. Moreover, the predictive power of these models is limited to forecasting short sequences, while the high accuracy long sequence forecast is more desirable for practical scenarios¹⁰. In addition, conventional machine learning approaches generally require significant amount of training data^{11,12} that is infeasible in practical context and they are also unable to extract explainable information and knowledge from the dataset. The fusion of machine learning models with empirical models is a promising effective and efficient learning philosophy to address these limitations^{13,14}, especially physics-guided machine learning methods in physical science^{15,16}.

High-quality optical cavities for enhancing light-matter interaction form the backbone of optical quantum technologies¹⁷⁻²³. Among them, one example showing the significance of capturing and understanding time dynamics is the terahertz time-domain spectroscopy (THz-TDS) of THz

resonance structures, which continue to stimulate much interest in diverse disciplines and applications such as health²⁴, sensing and imaging^{25,26}, security²⁷, computing^{28,29}, and communication³⁰, in addition to quantum applications^{31,32}. The uniqueness of THz-TDS includes the simultaneous acquisition of amplitude and phase information, broad spectral coverage, and time-resolved capability³³⁻³⁵. However, the energy decaying in resonance features is notoriously slow, leading to long data acquisition time. Moreover, theoretically infinitely long THz time signal is inevitable to be practically truncated for reasonable experimental measurement and numerical simulation time or preventing undesired phenomena. Thus, scientific and technical trade-offs between data accuracy and acquisition time have to be carefully crafted.

Here, we describe a physics-guided and physics-explainable cascaded GRU networks model to forecast long-sequence time-domain signal using short-sequence input obtained from THz-TDS numerical simulations and experiments. Instead of directly training the model using a significant amount of high-fidelity data from either computationally expensive electromagnetic solvers or time consuming experiments, we employ a two-step multi-fidelity training approach. A large number of low-fidelity physical-model-generated synthetic data is employed to first pretrain the model, which reduces the hypothesis and search space of the network for fast and efficient learning, alleviates local minimum problems for high-accuracy forecast, and generalizes model applicability. Through transfer learning using only a small set of high-fidelity application-oriented data, the pretrained model is tailored to a broad range of resonance features, including resonant dielectric metasurfaces³⁶, electro-optic graphene plasmonics³⁷, and the ultrastrong coupling between electron cyclotron resonance in a high-mobility two-dimensional electron gas (2DEG) and photons in

a high-quality-factor cavity³². The cascaded networks enable precise long sequence forecast with the shortest input time sequence being 7% of full sequence length, suggesting the best $15\times$ data acquisition speedup. Furthermore, this model accurately captures signal features that only occupy 0.01 % of total signal energy in the experimentally collected data of Laudau polaritons, and simultaneously learns resonance energies of spectral features. The polariton dispersions obtained from experimentally measured and forecast time series, as well as model-learned quantities, show excellent agreement, and the derived cooperativity from forecast spectra matched experimentally reported value well³².

Figure 1a illustrates the model of cascaded GRU networks taking short input sequence (length k) and forecasting long output sequence (length L). The initial input sequence is used in the first GRU network and then combined with forecast output sequence for the input of the next-stage network. Theoretically, the minimum required number of GRU networks is $O(\log_2(L/k))$. In a middle stage of GRU networks, a physics-explainable GRU decoder is branched out for simultaneously learning corresponding physical quantities, resonance energies, associated with time signals (golden dashed rectangle in Fig. 1a). Hidden states from the GRU encoder and the time series data from previous stage GRU network make the input for both physics-explainable GRU decoder and the forecast GRU generator to produce the time series output of this GRU stage. The physics-explainable GRU decoder suggests that hidden states obtained from time series forecast GRU networks bear actual physical meaning, and is also practically useful in experiments for accelerating not only the acquisition but also the explanation and understanding of time dynamics data.

Figure 1b displays the training process of GRU networks. Conventional training approach (orange path in Fig. 1b) requires a large training dataset, which generally takes long time to be generated through simulations or experiments and frequently leads to suboptimal trained networks because of local minimum issues. In contrast, we utilize a physics-guided multi-fidelity two-step training approach (blue path in Fig. 1b), where random networks are first pretrained with a large number of low-fidelity synthetic data instantaneously generated from analytical physical models based on domain expertise. The pretrained network is then fine trained through transfer learning with a small high-fidelity specific dataset from various applications, which is obtained from either numerical simulations or experiments. The fully trained network is optimal and has superior performance over the one trained through conventional approach.

We first demonstrate our model and training methodology for time series forecast in an example of resonant dielectric metasurfaces, which consists of a periodic array of unit cells with four dielectric cylindrical pillars of varying heights (h) and diameters (d)³⁶; see Fig. 2a. We generated training data by numerically simulating electrical field time response of different structures with randomly generated combinations of h and d ; detailed data generation procedure is described in *Methods*. In addition to bright photonic modes, the coupling between neighboring pillars can generate sharp Fano resonances. Figure 2b displays the training loss as a function of training epochs using conventional and two-step approaches. Despite extensive hyper-parameter tuning in conventional training approaches (see more details on hyper-parameters in *Methods*), we clearly observe local minimum issue that is prevailing in non-convex optimization problems and leads to poor forecast (blue dashed line with cross markers).

In order to solve this problem and improve forecast accuracy, we employ a multi-fidelity training framework¹⁵. Most physical resonance features and structures render themselves as a sum of damped oscillations in time signals, which follow a general mathematical form $\sum_i A_i e^{-\alpha_i t} \sin(\omega_i t + \phi_i)$, where A_i is the amplitude factor, $e^{-\alpha_i t}$ describes harmonics envelope decay, $\sin(\omega_i t + \phi_i)$ is an oscillating carrier with a resonance frequency ω_i and initial phase ϕ_i , and i enumerates all resonance features. We employed this analytical model to generate low-fidelity synthetic data, by randomly choosing ω_i within the frequency of interest, and arbitrarily selecting α_i , ϕ and the number of resonances; detailed generation procedures, data pre-processing, and exemplary synthetic data are in *Supplementary Information Sections 1, 2, and Fig. S1*. All GRU networks were pretrained with these synthetic data and then fine trained via transfer learning with a small set of high-fidelity time-domain data of the dielectric metasurface obtained through full-wave finite domain time difference (FDTD) simulations (see *Methods* for training data generation). The ratio of the length of final sequence over that of input sequence, L/k , is ~ 15 . As clearly indicated in Figs. 2b and c, Mean-Squared-Error (MSE) and the accuracy of forecast time series utilizing the multi-fidelity training approach outperforms those obtained using conventional approach. Furthermore, in the frequency domain, the spectra that is calculated through the Fourier transformation of forecast time signals obtained using the model trained by the multi-fidelity approach also show excellent agreement with the target spectra calculated by running full-time simulations.

The transfer learning technique greatly generalizes the applicability of pretrained networks. The similar damped oscillation signature of the time-domain signals in most resonance features physically guarantees the feasibility of such generalization. We demonstrated the approach gen-

eralization to two physically distinct resonance features, active graphene plasmonics³⁷ and ultra-strongly coupled Landau polaritons^{31,32}, in addition to dielectric metasurfaces. Specifically, periodically patterned monolayer graphene ribbons, which are expected to support localized plasmonic THz resonance from bounded carriers by ribbon boundaries, are simulated to obtain time-domain signals; see Fig. 3a. The ribbon width and graphene Fermi level are arbitrarily selected; details can be found in *Methods*. With similar pretrained networks and a small set of short input sequences ($L/k \sim 4$), both time and frequency domain forecast show excellent agreement with the target time signal (Fig. 3b) and spectra (Fig. 3c) obtained from FDTD simulations.

Moreover, an ultrahigh-mobility 2DEG inside a high-quality-factor one-dimensional THz photonic cavity under magnetic fields displayed ultra-narrow Landau polaritons³²; see Fig. 3d. Their spectra were experimentally measured with THz-TDS under high magnetic fields. There are in total 71 spectra under various magnetic fields from 0 T to 4.5 T, and we used 24 spectra as training data to fine-tune synthetic-data-trained model and the rest 47 spectra as test data. These 24 spectra were chosen as ones taken under magnetic fields either uniformly or randomly distributed between 0 T and 4.5 T. We then used the model trained by selected 24 measurements to forecast the rest 47 time signals under other magnetic fields with corresponding short input sequence to reveal full dispersion of Landau polaritons. Compared with the other two simulation examples, experimental data are much more noisy and the number of dataset is very limited. Moreover, our Landau polariton features locate inside the defect mode of the stopband of photonic crystals, and the signal energy of interest only occupies 0.01 % of total signal energy; see Fig. S2 in *Supplementary Information*. All these factors make the forecast using conventional direct training approach

nearly impossible in practice. The multi-fidelity two-step approach instead significantly improves the predictive power and accuracy, and thus small but important resonance features are captured. The loss function is also the MSE loss calculated from predicted and target time signals; details on the loss function selection are in *Supplementary Information Section 3 and Fig. S3*. Figures 3e and f display one representative time signal and corresponding spectra from Fourier transformation both aligning well with target experimental results.

For Laudau polariton experiments, in addition to cascaded GRU networks for time-domain signal forecast, we added a physics-explainable GRU network as shown in Fig. 1a in the middle (2nd) GRU network to infer resonance energies associated with time signals. The cascaded GRU networks are first trained for time series forecast, and then the physics-explainable GRU network is trained with the hidden states and intermediate sequences from time series forecast GRU networks as input; see *Methods* for detailed training process. After training, the short time signals from test data are input into the full model, and both forecast time signals and corresponding resonance energies are simultaneously generated.

Figures 4a and b display all 47 experimental and predicted spectra obtained from the Fourier transformation of corresponding time signals, where L/k is ~ 6 for predicted time sequences. All essential physical features expected in linearly polarized transmission spectra are well reproduced, including cyclotron-resonance-active lower polaritons (CRA-LPs), CRA upper polaritons (CRA-UPs), and CR-inactive (CRI) modes³²; see *Supplementary Information Section 4 and Figs. S4 and S5* for more data in both time and frequency domains. In a stark contrast, the spectra obtained

directly from short input sequences with appropriate zero padding, as shown in Fig. 4c, display completely random patterns with all features missing. The distinct difference between Figs. 4a, 4b, and 4c highlights the necessity of long data acquisition to capture essential THz features if no prediction is employed, as well as the high predictive power and accuracy of our model of cascaded GRU networks and training approach. Note that the results do not depend on how 24 training spectra are selected and the average validation loss displays small variance when training spectra are randomly shuffled for different magnetic fields; see *Fig. S6 and Supplementary Information Section 4*.

Furthermore, yellow stars in Fig. 4d, orange circles in Fig. 4e, and gold triangles in Fig. 4f show extracted peak positions from experimental measured spectra that are the Fourier transformation of time signals, model-learned resonance frequencies, and extracted peaks from predicted spectra, respectively. We utilized transfer matrix formalism to calculate transmission spectra for samples under various magnetic fields, through which the coupling ratio g can be extracted; details can be found in *Supplementary Information Section 5* and Ref.³². Figs 4d, e, and f also display calculated transmittance colour contour map to match extracted and learned peaks. Despite noisy CRA-UP branch in predicted spectra, the fitting can be done uniquely. The extracted coupling ratios, g_{exp} , g_{explain} , and g_{pred} from experimental spectra, physics-explainable GRU decoder, and predicted spectra, are 150.1 GHz, 145.6 GHz and 147.9 GHz. This excellent agreement confirms not only the precise forecast of cascaded GRU networks but also generated hidden states bear actual and explainable physical meanings, which are resonance energies.

With obtained predicted dispersion, the magnetic field corresponding to the conditions at zero detuning and far away from zero detuning can be determined. As shown in Figs. 4g, h, and i, we predicted spectra of CRI mode at 3 T and the CRA-UP and CRA-LP peaks at zero detuning using short input sequence with $L/k \sim 2$ using 24 training spectra and determined their spectral line widths to be 4.6 GHz, 5.2 GHz and 5.2 GHz, respectively. Thus, the cooperativity from predicted sequence is 3663, which is very close to the experimental value 3513³². In addition to the acceleration of experimental data acquisition for mapping out Landau polariton dispersion and determining corresponding physical quantities, our model is useful to assist and guide experiments by dynamically forecasting time signals through active learning. We utilized measured spectra from 4.5 T to a magnetic field close to zero detuning (e.g., $B_m = 1.4$ T) as training spectra to forecast the spectrum under next lower magnetic field (B_{m-1}). We then augmented the training spectra by adding the measured spectra under B_{m-1} for forecasting the spectra under B_{m-2} , and this process keeps moving to the lowest magnetic field. The anomalous increase of training loss during this active learning process indicates new time signal and spectral features, such as additional transmission peaks in the stopband around zero detuning, and helps experimentalists to make their decisions; see *Fig. S7 and Supplementary Information Section 4* for more details.

Our physics-guided and physics-explainable cascaded GRU networks combine complementary advantages of both physics-based models and machine learning approaches for fast, efficient, and accurate forecast of long time signals in optical resonances with a minimum of 7% input sequence, which corresponds to a best $15\times$ speedup of data acquisition. The multi-fidelity two-step training framework enables the model to be generalized to varieties of contexts of resonance

features and structures because of the similarity of their physical appearance. This model is especially efficient in practical experimental scenarios, where long data acquisition is inevitable and the collection of a large number of data is almost infeasible. The incorporation of synthetic data in training process together with transfer learning technique significantly reduces the required number of training data to 24 spectra for forecasting experimental spectra of Landau polaritons, and improves the predictive power and accuracy for capturing signal features that only occupy 0.01 % of total signal energy. This approach is promising to accelerate the discovery of new phenomena in complex systems through the analogy with more accessible optical system³⁸ and the exploration of device functionalities under resonant light-matter interactions.

Methods

Training data generation for dielectric metasurface and graphene plasmonics. Finite domain time difference (FDTD) time-domain simulations implemented in commercial Ansys Lumerical software are used for the training data generation in the examples of the dielectric metasurface and graphene plasmonics. For the dielectric metasurface example, we used four silicon cylindrical rods as the base structure with periodic boundary condition. Different data were generated by randomly selecting the radii of four cylindrical rods. The radius is chosen in the range from $39.5 \mu\text{m}$ to $44.5 \mu\text{m}$ with a step resolution of $0.25 \mu\text{m}$. We generated a total number of 5500 samples and using 5000 of them as the training set and the rest 500 as the validation set. For the graphene plasmonics example, graphene monolayer is modeled as a 2D rectangle conducting sheet in Lumerical material library, including both interband and intraband contributions. Fermi level and scattering rate are two parameters used to calculate dielectric constants used for the software. The dataset is generated by randomly sweeping graphene ribbon width and chemical potential. The width range is between $3.8 \mu\text{m}$ to $13.8 \mu\text{m}$ with a step of $0.2 \mu\text{m}$. The chemical potential ranges from 0.18 eV to 0.41 eV with a step of 0.01 eV . We generated a total number of 1000 samples and using 800 of them as the training set and the rest 200 as the validation set.

Training of recurrent neural networks with GRU. The multi-fidelity two-step training contains two steps: pretrain and fine-train (transfer learning). During the pretrain stage, the synthetic data are divided into pieces with different sequence length and fed into the model. During the fine-train process, the actual data (simulations and experimental measurements) are used to further update all parameters of the model. The GRU network we used here is a four-layer structure with a size

of 200 hidden states. In the pretrain stage, batch size is set as 128 and the total epoch number is 80. The Adam optimizer is used with the initial learning rate set at 5×10^{-4} . The learning rate decays every 30 epochs with the decay rate 0.01. The detailed hyper-parameters are summarized in the *Supplementary Information Section 6 Table 1*.

In the example of Landau polaritons, once the model for time dynamics forecast is fully trained, a physics-explainable GRU decoder is connected to the trained encoder of the 2nd GRU to receive the encoded hidden states and processed time signal outputs. We trained this physics-explainable GRU network with resonance frequency labels associated with time-domain signals, while we kept weights of the rest cascaded GRU networks unchanged and only updated this branch GRU decoder. The decoder is a four-hidden layer GRU network similar to the decoders in time signal forecasting networks. In limited 71 experimental spectral data, we utilized 24 data for training and 47 data for testing. This physics-explainable network was also trained using the Adam optimizer with a learning rate of 5×10^{-5} . The total epoch for the training is 300 with the batch size of 1. The learning rate decays every 100 epochs with the decay rate 0.1.

Acknowledgements J. F. and W. G. thank the support from the University of Utah start-up fund. C. Y. thanks the support from grants NSF-2019336 and NSF-2008144.

Author contributions C. Y. and W. G. conceived the idea and designed the project. Y. T. performed the modeling and calculations with the help of J. F., J. M., M. Q., C. Y. and W. G. X. L. helped with the analysis of Landau polariton data. Y. T. and W. G. wrote the manuscript. All authors discussed the manuscripts and provided the feedback.

Competing Interests The authors declare that they have no competing financial interests.

Correspondence Correspondence and requests for materials should be addressed to Cunxi Yu (email: cunxi.yu@utah.edu) and Weilu Gao (email: weilu.gao@utah.edu).

1. Mudelsee, M. Trend analysis of climate time series: A review of methods. *Earth-Sci. Rev.* **190**, 310–322 (2019).
2. Topol, E. J. High-performance medicine: the convergence of human and artificial intelligence. *Nat. Med* **25**, 44–56 (2019).
3. Lim, B. & Zohren, S. Time series forecasting with deep learning: A survey. *arXiv preprint arXiv:2004.13408* (2020).
4. Box, G. E., Jenkins, G. M., Reinsel, G. C. & Ljung, G. M. *Time series analysis: forecasting and control* (John Wiley & Sons, 2015).
5. Goodfellow, I., Bengio, Y., Courville, A. & Bengio, Y. *Deep learning*, vol. 1 (MIT press Cambridge, 2016).
6. Fu, R., Zhang, Z. & Li, L. Using lstm and gru neural network methods for traffic flow prediction. In *2016 31st Youth Academic Annual Conference of Chinese Association of Automation (YAC)*, 324–328 (IEEE, 2016).
7. Kong, W. *et al.* Short-term residential load forecasting based on lstm recurrent neural network. *IEEE Transactions on Smart Grid* **10**, 841–851 (2017).
8. Nelson, D. M., Pereira, A. C. & de Oliveira, R. A. Stock market's price movement prediction with lstm neural networks. In *2017 International joint conference on neural networks (IJCNN)*, 1419–1426 (IEEE, 2017).
9. Hyndman, R. J. A brief history of forecasting competitions. *Int. J. Forecast.* **36**, 7–14 (2020).

10. Zhou, H. *et al.* Informer: Beyond efficient transformer for long sequence time-series forecasting. *arXiv preprint arXiv:2012.07436* (2020).
11. Jiang, J., Chen, M. & Fan, J. A. Deep neural networks for the evaluation and design of photonic devices. *Nat. Rev. Mater.* 1–22 (2020).
12. Ma, W. *et al.* Deep learning for the design of photonic structures. *Nat. Photonics* 1–14 (2020).
13. Rangapuram, S. S. *et al.* Deep state space models for time series forecasting. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems (NIPS 2018)*, 7796–7805 (2018).
14. Salinas, D., Flunkert, V., Gasthaus, J. & Januschowski, T. Deepar: Probabilistic forecasting with autoregressive recurrent networks. *Int. J. Forecast.* **36**, 1181–1191 (2020).
15. Willard, J., Jia, X., Xu, S., Steinbach, M. & Kumar, V. Integrating physics-based modeling with machine learning: A survey. *arXiv preprint arXiv:2003.04919* (2020).
16. Karniadakis, G. E. *et al.* Physics-informed machine learning. *Nat. Rev. Phys.* **3**, 422–440 (2021).
17. Pellizzari, T., Gardiner, S. A., Cirac, J. I. & Zoller, P. Decoherence, continuous observation, and quantum computing: A cavity qed model. *Phys. Rev. Lett.* **75**, 3788 (1995).
18. Cirac, J. I., Zoller, P., Kimble, H. J. & Mabuchi, H. Quantum state transfer and entanglement distribution among distant nodes in a quantum network. *Phys. Rev. Lett.* **78**, 3221 (1997).

19. Duan, L.-M., Lukin, M. D., Cirac, J. I. & Zoller, P. Long-distance quantum communication with atomic ensembles and linear optics. *Nature* **414**, 413–418 (2001).
20. Mabuchi, H. & Doherty, A. Cavity quantum electrodynamics: coherence in context. *Science* **298**, 1372–1377 (2002).
21. Englund, D. *et al.* Controlling cavity reflectivity with a single quantum dot. *Nature* **450**, 857–861 (2007).
22. O’Brien, J. L., Furusawa, A. & Vučković, J. Photonic quantum technologies. *Nat. Photonics* **3**, 687–695 (2009).
23. Forn-Díaz, P., Lamata, L., Rico, E., Kono, J. & Solano, E. Ultrastrong coupling regimes of light-matter interaction. *Rev. Mod. Phys.* **91**, 025005 (2019).
24. Siegel, P. H. Terahertz technology in biology and medicine. *IEEE Trans. Microw. Theory Techn* **52**, 2438–2447 (2004).
25. Tonouchi, M. Cutting-edge terahertz technology. *Nat. Photonics* **1**, 97–105 (2007).
26. Mittleman, D. M. Twenty years of terahertz imaging. *Opt. Express* **26**, 9417–9431 (2018).
27. Liu, H.-B., Zhong, H., Karpowicz, N., Chen, Y. & Zhang, X.-C. Terahertz spectroscopy and imaging for defense and security applications. *Proc. IEEE* **95**, 1514–1527 (2007).
28. Lin, X. *et al.* All-optical machine learning using diffractive deep neural networks. *Science* **361**, 1004–1008 (2018).

29. Li, Y., Chen, R., Sensale-Rodriguez, B., Gao, W. & Yu, C. Real-time multi-task diffractive deep neural networks via hardware-software co-design. *Sci. Rep.* **11** (2021).
30. Kleine-Ostmann, T. & Nagatsuma, T. A review on terahertz communications research. *J. Infrared Millim. Terahertz Waves* **32**, 143–171 (2011).
31. Zhang, Q. *et al.* Collective non-perturbative coupling of 2D electrons with high-quality-factor terahertz cavity photons. *Nat. Phys.* **12**, 1005–1011 (2016).
32. Li, X. *et al.* Vacuum bloch–siebert shift in landau polaritons with ultra-high cooperativity. *Nat. Photonics* **12**, 324–329 (2018).
33. Ulbricht, R., Hendry, E., Shan, J., Heinz, T. F. & Bonn, M. Carrier dynamics in semiconductors studied with time-resolved terahertz spectroscopy. *Rev. Mod. Phys.* **83**, 543 (2011).
34. Jepsen, P. U., Cooke, D. G. & Koch, M. Terahertz spectroscopy and imaging—modern techniques and applications. *Laser & Photonics Reviews* **5**, 124–166 (2011).
35. Neu, J. & Schmuttenmaer, C. A. Tutorial: An introduction to terahertz time domain spectroscopy (thz-tds). *J. Appl. Phys.* **124**, 231101 (2018).
36. Nadell, C. C., Huang, B., Malof, J. M. & Padilla, W. J. Deep learning for accelerated all-dielectric metasurface design. *Opt. Express* **27**, 27523–27535 (2019).
37. Ju, L. *et al.* Graphene plasmonics for tunable terahertz metamaterials. *Nat. Nanotechnol.* **6**, 630–634 (2011).

38. Li, X. *et al.* Observation of dicke cooperativity in magnetic interactions. *Science* **361**, 794–797 (2018).

Figure 1 Cascaded GRU networks and two-step training approach. (a) Forecast of a L -long time signal with a k -long time signal using $O(\log_2(L/k))$ cascaded gated recurrent units (GRUs). The output sequence of the GRU at each stage is combined with the input sequence to serve as the input for the next-stage GRU. In experiments, a physics-explainable GRU network is connected to a GRU encoder in the middle stage of cascaded GRU networks that are for time series forecast. (b) Slow and suboptimal conventional training approach (orange path) and our fast, broadly applicable, and optimal two-step training approach (blue path). Conventional training approach needs a large number of training datasets. The generation of such large datasets requires long time of numerical simulations and experiments and sometimes is not even feasible in experimental settings. In contrast, in our multi-fidelity training approach, the model is first trained with fast generated low-fidelity synthetic data from physical models and then tailored to specific applications using only a small number of high-fidelity datasets from either numerical simulations or experiments.

Figure 2 Dielectric metasurfaces. (a) A generic dielectric metasurface with the unit cell consisting of four cylindrical pillars of different dimensions. (b) The escape from local minimum in the two-step training approach. (c) Forecast time signal using conventional training approach and two-step approach. The clear better forecast performance is observed in the two-step approach. (d) THz spectra obtained through the Fourier transformation of the full time signal and the forecast time signal using short input time sequence.

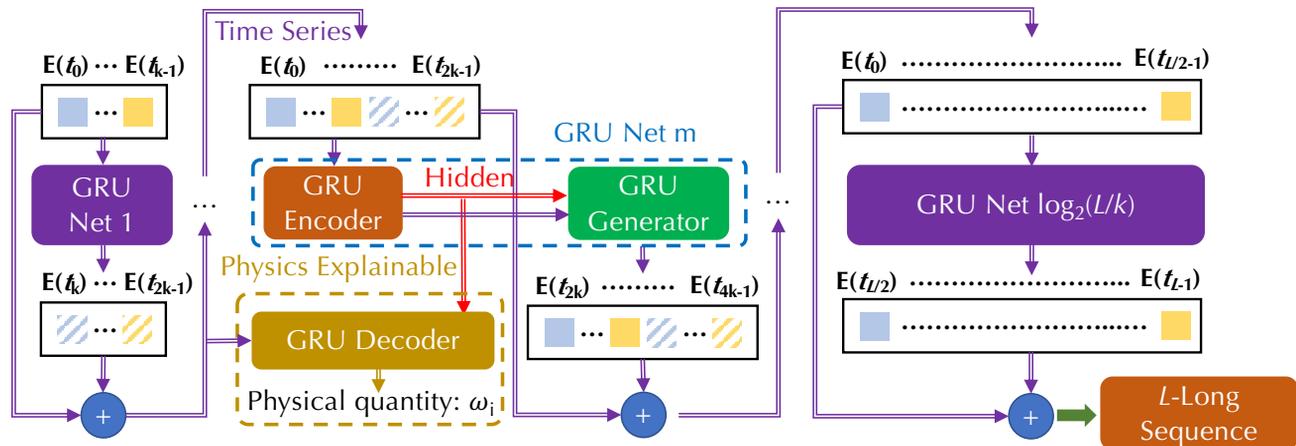
Excellent agreement between forecast spectra and target spectra confirms the superior performance of our two-step approach.

Figure 3 Model transfer to graphene plasmonics and Landau-polaritons. (a) – (c) Graphene plasmonics. (a) Schematics of graphene ribbons supporting localized plasmonic resonance. (b) Time domain and (c) corresponding frequency domain response for predicted time signal and full target time signal, respectively. (d) – (f) Landau-polaritons in strongly coupled photons in an one-dimensional photonic crystal cavity with the electron cyclotron resonance in high-mobility two-dimensional electron gases (2DEG). (e) Time domain and (f) corresponding frequency domain response under specific magnetic field for predicted time signal and full target time signal, respectively.

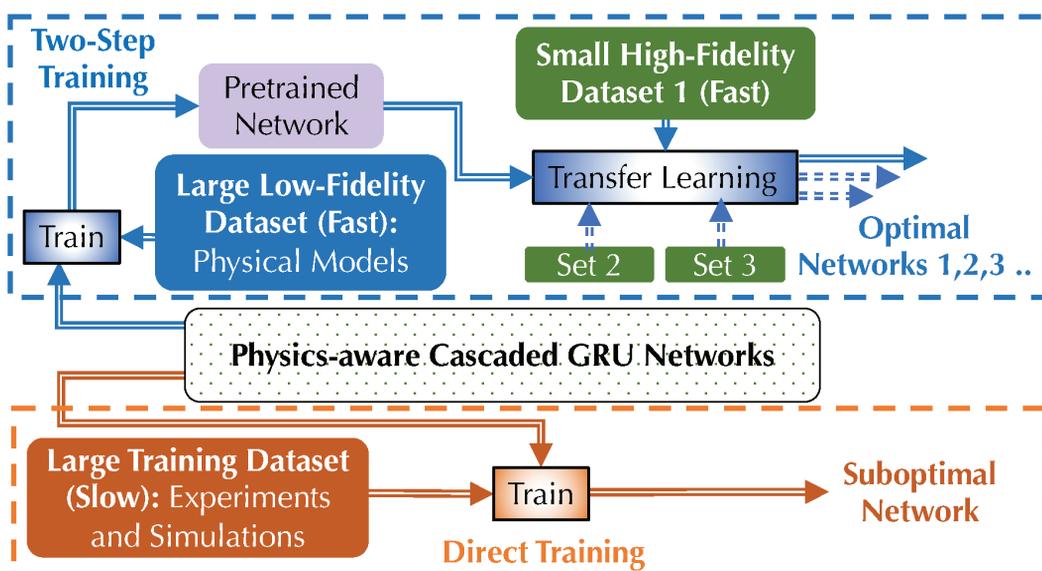
Figure 4 Experimental verification of models in Landau-polaritons. Linearly polarized THz transmittance spectra under various magnetic fields obtained through the Fourier transformation of (a) experimentally acquired time signals, (b) predicted time signals with a much shorter input experimental time sequence, and (c) short input experimental time signals with zero padding. CRA-LP and CRI branches clearly observed in experimental spectra are reproduced in predicted spectra, where the CRA-UP branch is more noisy. The spectra obtained from zero-padded input time signal is completely random. Simulated transmittance colour contour map using transfer matrix method to fit (d) experimental spectra shown in (a), (e) model-learned resonance frequencies, and

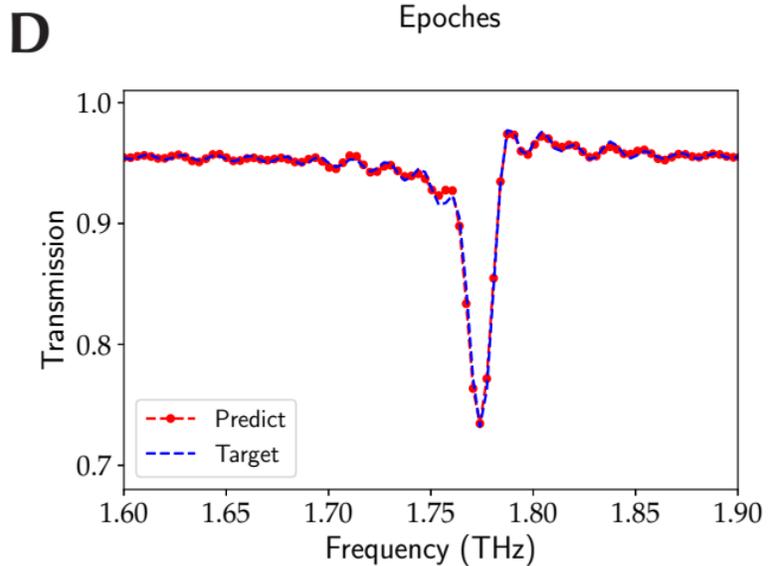
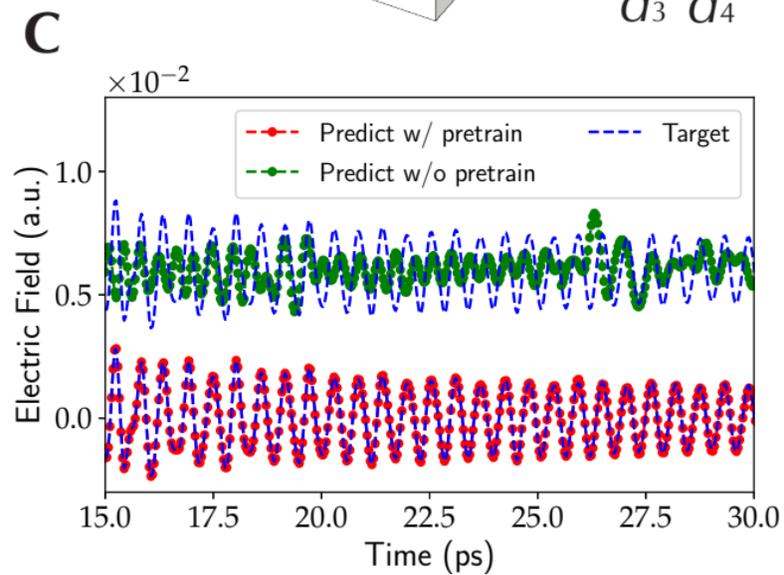
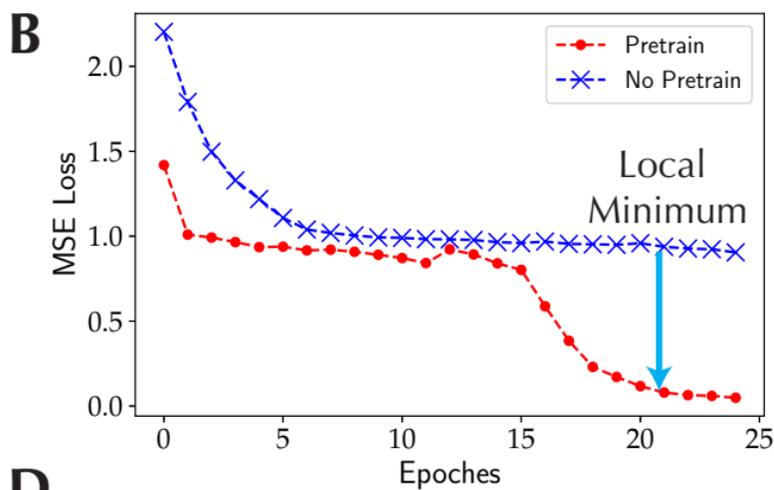
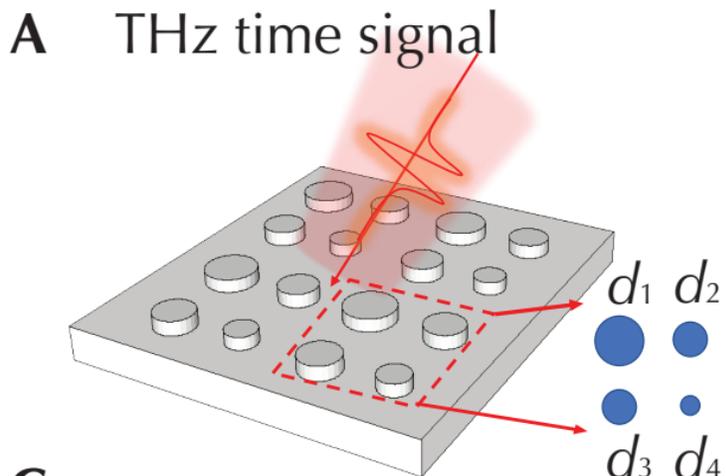
(e) predicted spectra shown in (b). Yellow stars in (d), orange circles in (e), and gold triangles in (f) mark the peak positions extracted from experimental spectra, learned from physics-explainable GRU network, and extracted from predicted spectra. The extracted coupling rates g are 150.1 GHz, 145.6 GHz, and 147.9 GHz, respectively. Lorentzian fits (red dashed line) of predicted spectra for (g) the UP peak at zero detuning (1T) (h) the LP peak at zero detuning, and (i) the CRI mode at 3T. The obtained full-width at half-maximum values are indicated by green arrows.

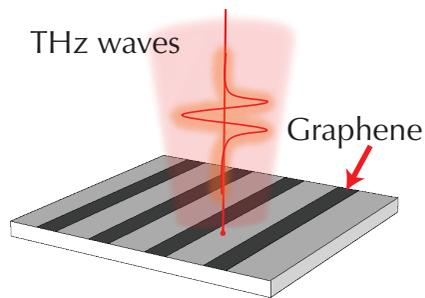
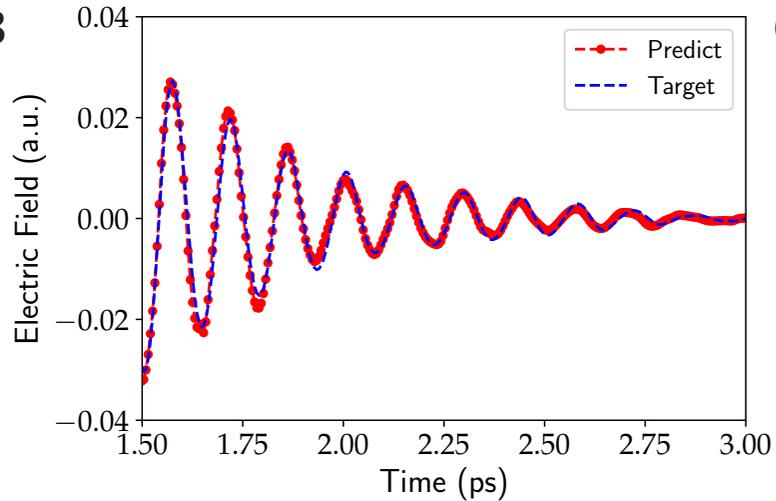
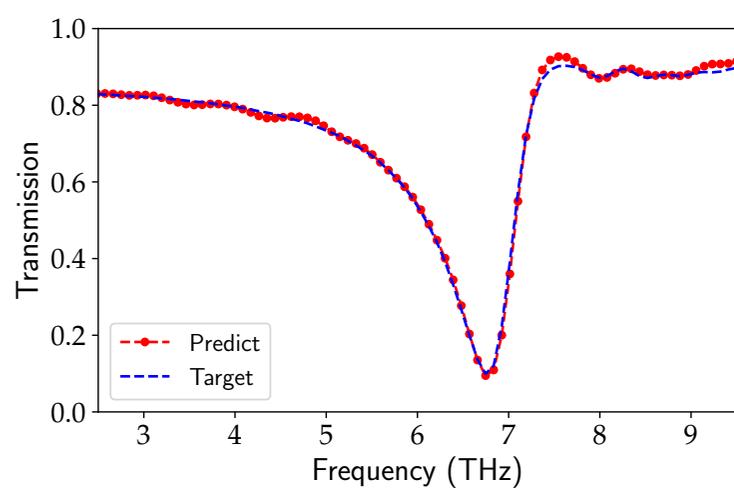
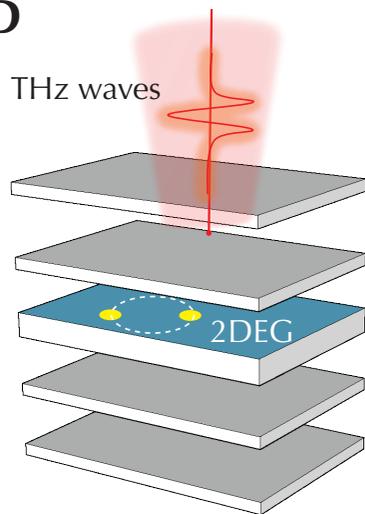
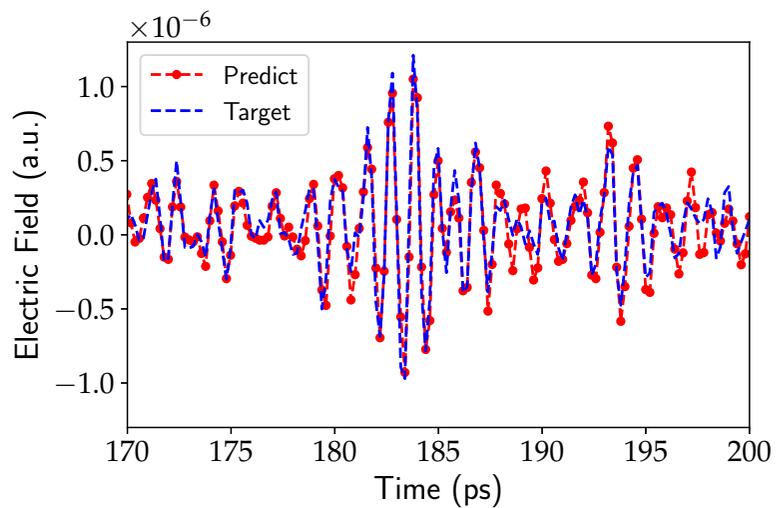
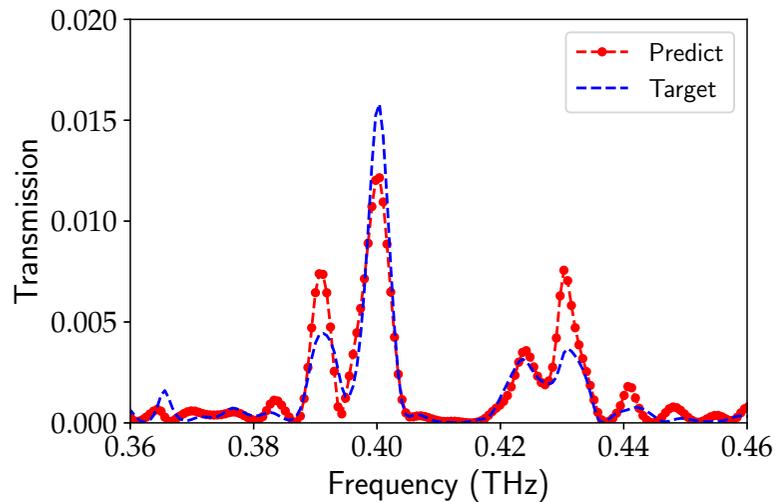
A

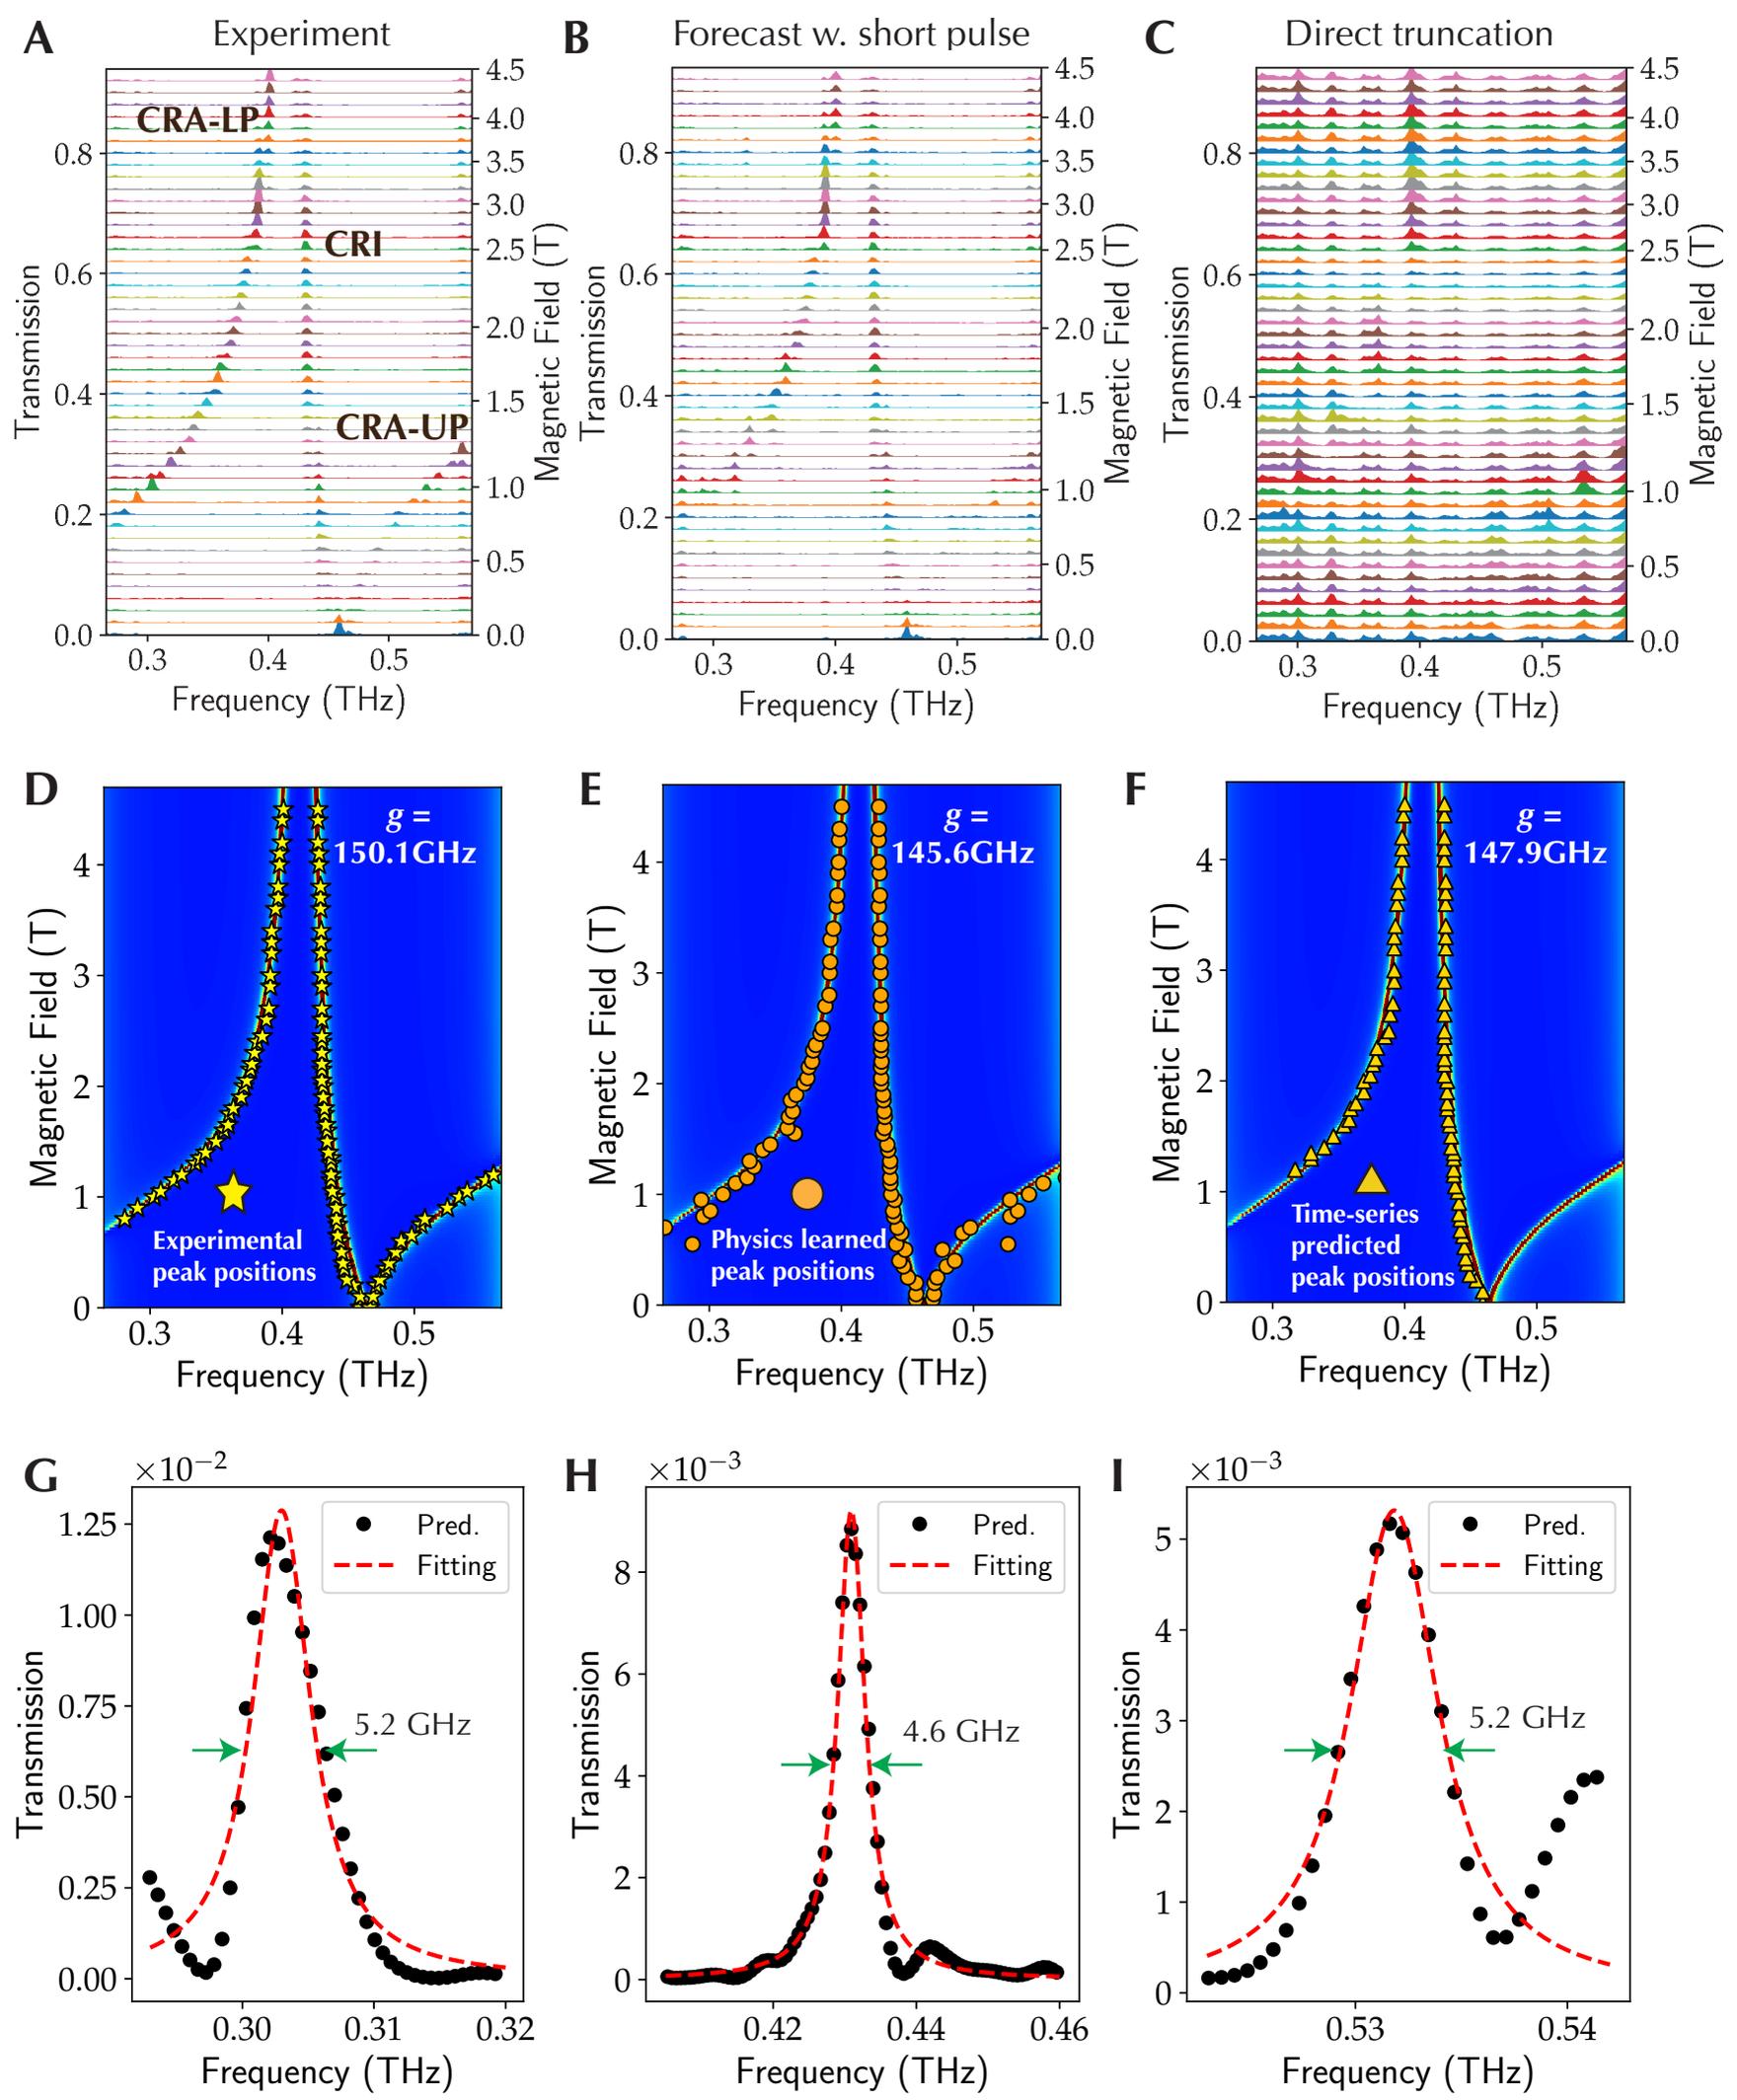


B





A**B****C****D****E****F**



Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SI.pdf](#)