

<https://doi.org/10.1038/s42005-025-02086-1>

Towards AI-driven autonomous growth of 2D materials based on a graphene case study



Leonardo Sabattini^{1,2,6}, Annalisa Coriolano^{3,4,6}, Corneel Casert^{5,6}, Stiven Forti³, Edward S. Barnard⁵, Fabio Beltram¹, Massimiliano Pontil⁴, Stephen Whitlam⁵✉, Camilla Coletti³✉ & Antonio Rossi³✉

The scalable synthesis of two-dimensional (2D) materials remains a key challenge for their integration into solid-state technology. While exfoliation techniques have driven much of the scientific progress, they are impractical for large-scale applications. Advances in artificial intelligence (AI) now offer new strategies for materials synthesis. This study explores the use of an artificial neural network (ANN) trained via evolutionary methods to optimize graphene growth. The ANN autonomously refines a time-dependent synthesis protocol without prior knowledge of effective recipes. The evaluation is based on Raman spectroscopy, where outcomes resembling monolayer graphene receive higher scores. This feedback mechanism enables iterative improvements in synthesis conditions, progressively enhancing sample quality. By integrating AI-driven optimization into material synthesis, this work contributes to the development of scalable approaches for 2D materials, demonstrating the potential of machine learning in guiding experimental processes.

The emergence of two-dimensional (2D) materials has revolutionized material science, offering promising advancements across a wide range of technological applications^{1–5}. However, achieving scalable production of high-quality single-crystal 2D materials remains a significant challenge. Although exfoliated forms of these crystals demonstrate excellent potential for addressing numerous technological challenges, non-scalable devices made from exfoliated materials often lack reproducibility and require cumbersome fabrication methods. Achieving scalability for these materials is highly desirable, not only to facilitate practical applications but also to validate the promising scientific discoveries made to date.

Graphene with high crystallinity and charge mobility comparable to that of exfoliated graphene can be produced through chemical vapor deposition (CVD)^{6–8}. This method was successfully demonstrated for large-scale growth, making it a promising approach for various industrial applications^{9–13}. In contrast, the development of other 2D materials and their heterostructures still lacks a viable and scalable method. The main challenges stem from the difficulty in producing large-area crystals and, consequently, heterostructures, with controlled thickness, minimal defects, and uniformity on a large scale. As a result, the scalable synthesis of these materials has not yet matured in terms of reproducibility and defect management^{14,15}. This

limitation restricts the practical use and replication of promising materials across various fields such as quantum computing and sensing^{16,17}.

Artificial Intelligence (AI) can provide a compelling solution to this challenge. In recent years, AI has become an indispensable part of our society, proving to be extremely effective in solving complex problems across various fields, with applications ranging from precision medicine to autonomous driving^{18–22}. In the rapidly advancing field of graphene research AI was applied to tasks such as determining the potential energy surface²³, predicting bandgaps²⁴, or forecasting crack evolution in graphene sheets²⁵. Recent advances in AI have also been used to explore density functional theory (DFT) calculations and study diffusion mechanisms based on molecular dynamics (MD)^{26–29}. One of the most striking results was shown by Merchant and coworkers who were able to predict the stability of millions of new crystal structures, many of which had never been discovered through traditional methods³⁰. In this context, the integration of AI into materials science opens a new era of innovative synthesis methodologies. Recently, there has been considerable interest in establishing autonomous laboratories that combine robotics with the exploitation of *ab initio* databases and active learning to optimize the synthesis of novel inorganic materials³¹, a process that is generally time-consuming and expensive.

¹Scuola Normale Superiore, Laboratorio NEST, Pisa, Italy. ²Dipartimento di Chimica e Chimica Industriale, Università di Pisa, Pisa, Italy. ³Center for Nano-technology Innovation, Laboratorio NEST, Istituto Italiano di Tecnologia, Pisa, Italy. ⁴Computational Statistics and Machine Learning, Istituto Italiano di Tecnologia, Genova, Italy. ⁵The Molecular Foundry, Lawrence Berkeley National Laboratory, Berkeley, CA, USA. ⁶These authors contributed equally: Leonardo Sabattini, Annalisa Coriolano, Corneel Casert. ✉e-mail: swhitlam@lbl.gov; camilla.coletti@iit.it; antonio.rossi@iit.it

Artificial Neural Networks (ANNs) are particularly suited for encoding high-dimensional objects, such as time-dependent synthesis protocols, where control parameters (e.g., temperature) vary dynamically over time. This flexibility allows ANNs to model complex relationships effectively, making them an ideal choice for autonomous optimization tasks. As demonstrated in our previous work^{32–35}, neural networks can be trained to express optimal control protocols for a wide range of physical systems.

In this study, we aim to address the following question: Can an ANN learn a time-dependent protocol in order to grow a material with desired optical and crystalline properties, without prior knowledge of growth protocols? This is a radically different approach with respect to those mentioned so far. Here, we utilize an active learning method, called adaptive Monte Carlo (aMC), in which the ANN iteratively refines its synthesis protocols without relying on historical data. This innovative approach represents a significant advancement for synthesis methodologies in general, enabling their optimization directly from experimental feedback. In this work, we illustrate its application in the specific case of graphene synthesis, demonstrating its effectiveness. As proof of principle, we focus on the relatively straightforward task of growing high-quality, homogeneous graphene through the thermal decomposition of silicon carbide (SiC). Here an ANN is tasked with proposing a protocol, a profile of temperature as a function of time, to achieve this goal. The growth of graphene from SiC is chosen as an ideal candidate for exploring the feasibility of applying ANN learning to crystal growth, owing to its simplicity and the manageable number of growth parameters that can be easily controlled (i.e., min temperature, max temperature, and ramp). This method enables the synthesis of graphene directly from SiC^{36,37}, thus eliminating the need for gaseous carbon precursors. This work showcases the feasibility and significant advantages of employing artificial intelligence to tailor and optimize the growth of 2D materials. By integrating an ANN that learns and adapts, we effectively embed a “brain” into the synthesis lab, paving the way for the autonomous synthesis of desired materials. More ambitiously, this approach holds the potential to discover methods for synthesizing high-quality materials that are currently beyond our capabilities, unlocking new frontiers in material science.

Methods

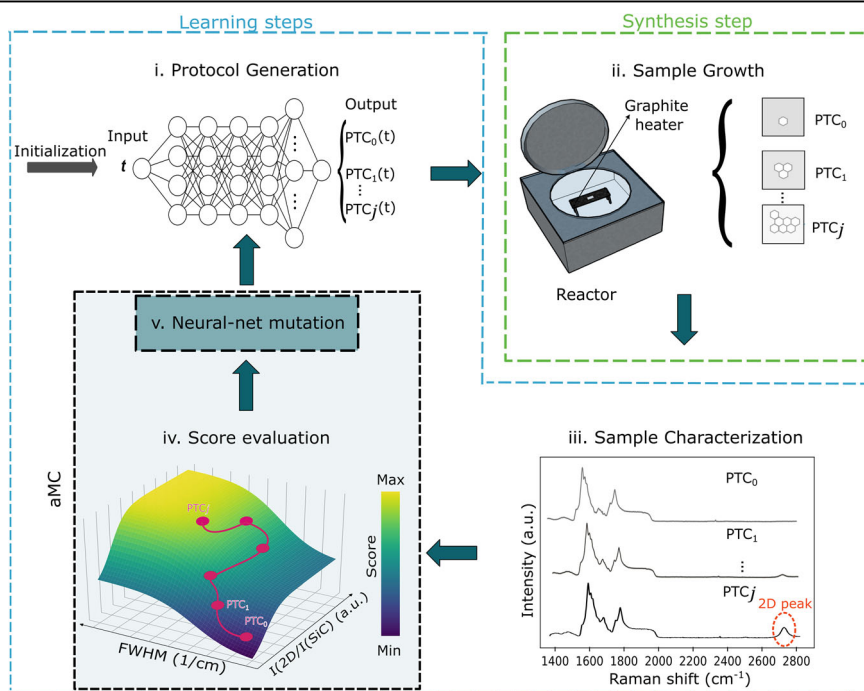
We outline here the methodology developed to train an ANN with aMC to autonomously find the most efficient growth protocol for graphene from the

thermal decomposition of SiC. Minimal input is provided by setting only the furnace working temperature range and a starting temperature. The entire ANN training procedure is schematized in Fig. 1 and consists of five iterative steps: (i) Protocol Generation: a protocol PTC , that is a temperature profile $T(t)$, is generated by an ANN and given as an input to a cold-wall reactor; (ii) Sample growth: graphene growth on SiC is performed in the reactor adopting as input the temperature profile generated in step (i); (iii) Sample Characterization: Raman spectroscopy is performed on the synthesized sample and the spectrum obtained is benchmarked with an ideal target to generate a score (iv). At the end of this process the ANN parameters are updated (v) following the aMC method described below, and a new protocol PTC_j is generated. This iterative process generates one protocol per cycle, starting with PTC_0 in the first cycle, PTC_1 in the second cycle, and so on, with the j^{th} cycle generating PTC_j .

At its core, the neural network is a nonlinear function approximator characterized by a set of parameters (weights and biases). The nodes take as input a sequence of time values $t = \{t_1, t_2, \dots\}$ and generate a temperature profile $T(t)$ that defines a protocol PTC , modeling the evolution of temperature over time based on the learned relationships between time and temperature. Each protocol PTC_j produced by the network is tested against experimental data and generates a scalar score. This score is then used in the learning step, where the weights and biases of the network are adjusted in order to minimize the error and improve the accuracy of the model's predictions.

Before describing the details of the aMC algorithm, we will discuss the main aspects of the growth of epitaxial graphene on SiC. It involves a high-temperature process in which the crystal is thermally decomposed³⁸. The most important control parameters during growth are temperature, pressure and time. Originally this technique was implemented in ultra high vacuum (UHV) chambers³⁷. It is now agreed that atmospheric-pressure growth conditions under inert gas in quartz reactors (either cold-wall or hot-wall, horizontal or vertical) provide the most favorable conditions to obtain graphene³⁹. During the thermal decomposition process, silicon atoms sublimate from the surface, and leave behind a carbon-rich layer, a graphene-like honeycomb lattice with one third of the C atoms forming covalent bonds to the SiC substrate through sp^3 -hybridized orbitals⁴⁰. This is known as buffer or zero-layer graphene (ZLG). This layer is not yet graphene and exhibits a band gap^{39,41}. The growth of graphene requires

Fig. 1 | Schematic representation of ANN training used in this work. After initialization with the parameter guess, the protocol generation is conducted (i), followed by sample growth (ii). The obtained sample is then characterized through Raman spectroscopy (iii). The extracted data are used to evaluate the score (iv). Finally, the protocol is updated with the new parameters (v) and a new protocol is generated.



reaching an optimal (higher) temperature at which an additional carbon-rich layer forms beneath the first one, effectively decoupling it from the substrate. This layer does display the typical linear dispersion of monolayer graphene (MLG)³⁹. If the process continues further, another buffer layer will form at the interface with SiC, turning the previous graphene and buffer layers into a bilayer graphene (BLG). Hence, determining a correct temperature profile is fundamental to obtain MLG with minimum ZLG and BLG inclusions. Given the fundamental importance of the temperature profile in achieving optimal MLG coverage, we sought to investigate a broad yet physically-meaningful temperature range. We set the temperature boundaries between $T_{min} = 1100$ °C and $T_{max} = 1300$ °C. The minimum temperature was chosen below the known temperatures for graphene synthesis to allow for a reasonably wide temperature range for the ANN to explore. The maximum temperature was set equal to the maximum operational capacity of our reactor, ensuring the respect of safety and equipment limitations. We set a starting temperature $T_{start} = 1200$ °C as midpoint within this range. This starting temperature is not too low to hinder the initiation of the growth process, yet not too close to the optimal synthesis temperature, thereby encouraging the ANN to explore a variety of temperature profiles. In subsequent runs, as discussed below, we shall remove this starting temperature constraint to allow the ANN to explore the entire temperature range more freely. Detailed experimental information about the growth process implemented in this work are reported in Supplementary Note 1.

aMC is an evolutionary algorithm designed to work with relatively few experiments: each learning step requires one experiment, and the algorithm learns from past successes, proposing similar moves with increased likelihood⁴².

The algorithm proceeds as follows. Let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$ be the vector of neural-network parameters, which defines a time-dependent protocol. Let $U(\mathbf{x})$ be the loss function, which quantifies the success of the synthesis resulting from that time-dependent protocol. The algorithm proposes a change of all neural-network parameters $\mathbf{x} \rightarrow \mathbf{x}'$ by small Gaussian random numbers,

$$x_i \rightarrow x'_i = x_i + \epsilon_i, \quad (1)$$

where $\epsilon_i \sim \mathcal{N}(\mu_i, \sigma^2)$. Here σ sets the scale of parameter updates, and the μ_i are momentum-like parameters that we shall specify shortly. If the synthesis outcome resulting from the new protocol is better than (or as good as) the current outcome, i.e., if $U(\mathbf{x}') \leq U(\mathbf{x})$, then we accept the update (1), and the \mathbf{x}' become the new parameters of the neural network. Otherwise we revert to the previous parameters. The update procedure is then repeated.

The momentum-like parameters μ_i are initially set to zero. Following each accepted move, they are updated as $\mu_i \rightarrow \mu_i + \eta(\epsilon_i - \mu_i)$, where η is a hyperparameter of the method. This update ensures that subsequent parameter updates are more likely to be similar to past accepted updates. Following several consecutive rejected moves, the momentum-like parameters are reset to zero, and the scale parameter σ is reduced in size⁴².

For the current synthesis, the loss function U depends on the Raman spectrum. Each measured spectrum is assigned a score that reflects how closely it matches the desired characteristics. The loss function is inversely related to this score: a lower loss function corresponds to a higher score. The highest possible score is attributed to the ideal Raman spectrum, which serves as the benchmark for our evaluations. Consequently, the objective of the ANN training is to minimize the loss function by maximizing the similarity between the measured spectrum and the ideal Raman spectrum.

One of the key advantages of Raman spectroscopy is its capability to perform detailed mapping of samples, effectively distinguishing between different forms and qualities of graphene (e.g., buffer layer, MLG, and BLG)⁴³, thereby offering valuable insights into their distribution across the sample.

The ability to identify and map these variations in graphene layers aids in optimizing growth conditions and improving the overall quality of the 2D material. The spectrum obtained from graphene typically shows two

prominent peaks: G, that is a primary in-plane vibrational mode, and 2D, the second-order peak of the in-plane vibrational mode D⁴⁴. The intensity, position, and width of these peaks are highly sensitive to the quality and structure of graphene. High-quality graphene typically exhibits sharp and intense peaks in the Raman spectrum. For our analysis we consider the 2D peak, that can be fitted with a Lorentzian function. The score (f) is then calculated with the equation (2):

$$f = I + \sigma + \frac{1}{\chi}(\sigma \cdot I) \quad (2)$$

where I is related to the 2D peak intensity, while σ is related to the Lorentzian full width at half maximum (FWHM) through hyperbolic tangent, defined as follows:

$$I = \tanh\left(\frac{I_{2D} + a}{b}\right) + 1 \quad (3)$$

and

$$\sigma = \tanh\left(-\frac{FWHM + a'}{b'}\right) + 1 \quad (4)$$

χ is the interaction coefficient parameter, while a, b, a' , and b' are empirical parameters.

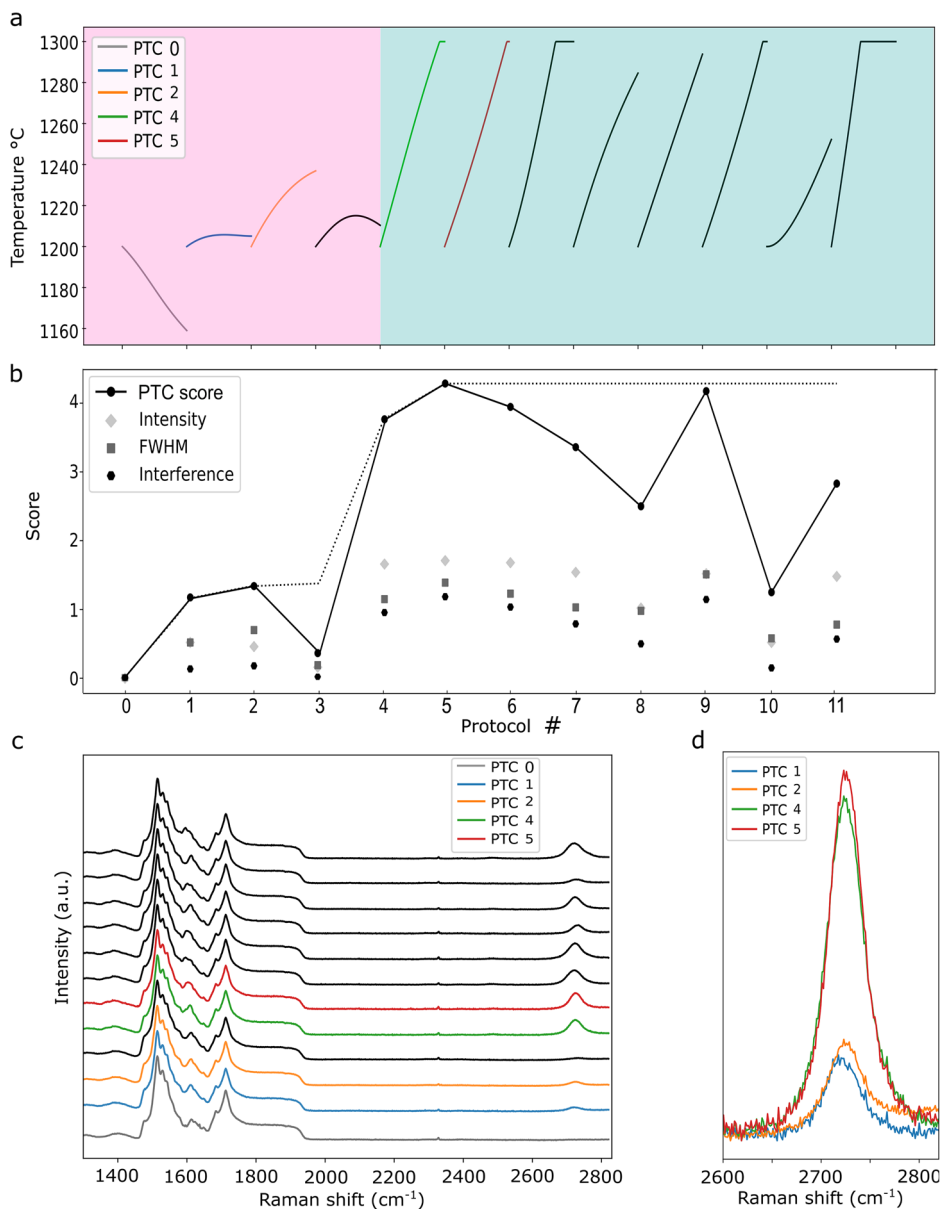
The intensity parameter is crucial because it directly correlates with the quantity of graphene on top of the substrate, while the term correlated to the FWHM aids the protocol in distinguishing between MLG and BLG, as BLG typically exhibits a larger FWHM⁴⁵. The choice of the hyperbolic function is dictated by the fact that beyond a threshold, the growth is complete; conversely, below a certain value, the growth is negligible. Moreover, the interference term $\frac{1}{\chi}(\sigma \cdot I)$ is necessary to prevent convergence in a region where only one of the two terms is optimized. Table S1 of Supplementary Note 2 reports the values of the empirical parameters used in (3) and (4). Figure S1 shows the 2D Raman peak (orange solid line) with the corresponding Lorentzian fit (blue solid line) for the Raman spectra corresponding to the protocols used in this work. A summary of the obtained values of I and σ are reported in Table S2. All the samples were analyzed using Raman spectroscopy under identical spectral sampling conditions. In particular, multiple samples were grown for each protocol to test the reproducibility of the process. We also checked the reproducibility of the single protocol, comparing the theoretical temperature profile with the measured one followed by the furnace (see Supplementary Note 3, Fig. S2).

Results and discussion

Unlike methodologies used in other works³¹, which rely heavily on pre-trained models and vast amounts of historical synthesis data to guide experiments, our approach distinguishes itself by requiring minimal input and by learning dynamically through direct experimental feedback. Additionally, previous methods focus on optimizing a limited state space of initial conditions, and do not treat time-dependent protocols. By contrast, our ANN expresses the protocol as a function of time, learning from experimental outcomes and progressively improving the quality of the graphene produced.

To train ANNs, gradient-based optimization methods like backpropagation are commonly used; however, these require explicit gradient information, which is often unavailable in experimental setups where the loss function depends on physical observables, such as Raman spectroscopy results. This limitation renders standard backpropagation unsuitable for our application. Therefore, training the neural network must rely on evolutionary methods^{46–48}. Standard genetic algorithms⁴⁸ are too costly for this purpose, because they require many runs of the experiment to train. Instead, we employ evolutionary methods, aMC, that uses only a single run of the experiment for each learning step. aMC effectively perform noisy gradient

Fig. 2 | Evolution of the Raman score-based protocol. **a** Temperature profiles in function of the protocol. Protocols that enhance the score are highlighted by colored lines. The first region (pale magenta) reflects ANN's trial phase, while the second (pale cyan) shows it learning a successful trend, favoring mostly increasing temperature profiles. **b** Score evolution with the generation: The solid black line represents the score calculated using Equation 2, while the dashed black line indicates the accepted score threshold. Additionally, the intensity (diamonds), σ (squares), and Interference parameters (circles) are depicted relative to the generation. **c** Raman spectra collected from samples used to train the ANN. Protocols that contribute to score improvement are identified by colored lines. The spectra are vertically spaced for clarity. **d** Zoomed-in Energy Range of the 2D-band Raman Peak.



descent without requiring explicit gradient calculation^{34,42,49}. This approach enables robust ANN training directly from experimental feedback, a critical capability for refining dynamic protocols. Alternative machine learning approaches, such as Bayesian optimization using Gaussian process regression, have shown success in static parameter optimization, leveraging pre-computed data to reduce the number of required simulations. However, their reliance on pre-existing training data and computationally expensive Gaussian processes limits their adaptability in high-dimensional, dynamic systems. As prior studies demonstrate, Bayesian optimization tends to explore regions of parameter space with promising outcomes but converges more slowly in the absence of training samples³¹. In contrast, our use of aMC facilitates efficient optimization of time-dependent protocols without the need for extensive historical data or pre-computed models, utilizing an evolutionary algorithm that resembles the zero-temperature Metropolis Monte Carlo algorithm with an adaptive, momentum-like component⁴². We start with a neural network whose parameters specify a time-dependent protocol. We add random numbers from a normal distribution to all weights, evaluate the new protocol resulting from this change, and accept (or reject) the new protocol if the outcome is better (or worse) than the previous protocol, as determined by a score function. This approach proceeds iteratively, with the algorithm remembering weight changes that led to

accepted protocols and making such changes with higher probability. In this way, we can achieve meaningful learning within a few tens of experiments, making the approach readily accessible to laboratory studies.

In our approach, score function evaluation relies on Raman spectroscopy measurements, a highly effective method extensively used for the characterization of 2D materials owing to its versatility and precision^{43,50-52}. Additionally, it is a rapid and non-invasive characterization technique that can cover an area of tens of micrometers in just a few minutes. For these reasons, it provides crucial feedback that aids in iteratively optimizing the synthesis protocols of artificial neural networks, thereby steadily improving the quality of graphene.

To further assess and validate our Raman-based scoring system, the structural, chemical, and electronic quality of graphene are characterized using supplementary techniques such as Atomic Force Microscopy (AFM), X-ray Photoelectron Spectroscopy (XPS), and Angle-Resolved Photoemission Spectroscopy (ARPES). These analyses offer comprehensive insight into the material quality, retrospectively validating the effectiveness of the AI-driven synthesis approach.

The results obtained are summarized in Fig. 2. The temperature profile for each generation is reported in Fig. 2a. Each curve is a tentative temperature profile or protocol, proposed by the ANN and implemented into

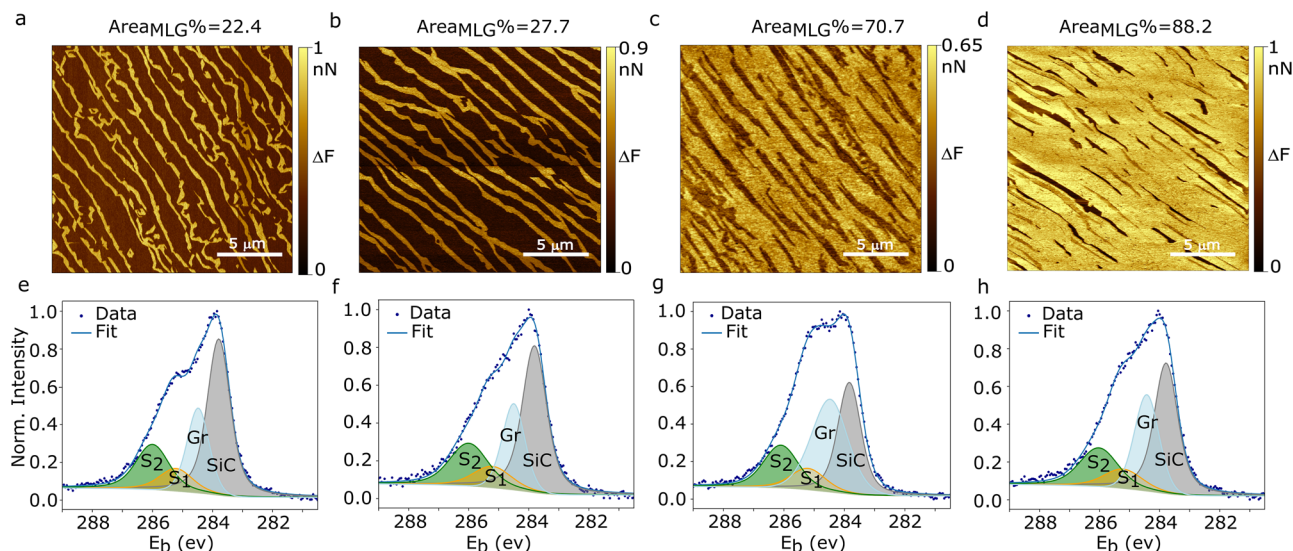


Fig. 3 | Adhesion force mapping and XPS analysis. Adhesion force map on samples obtained from different protocols (PTC): PTC1 (a) PTC2 (b) PTC4 (c) and PTC5 (d). The percentage of the monolayer graphene area ($Area_{MLG}$), calculated using this

method, is displayed at the top of each map. C1s XPS spectra recorded on samples obtained from PTC1 (e) PTC2 (f) PTC4 (g) and PTC5 (h).

the growth set up. In Fig. 2b, the solid black line represents the score calculated using Eq. 2, while the dashed black line indicates the accepted score threshold. Additionally, the intensity (diamonds), σ (squares), and interference (circles) partial scores are shown in relation to the generation. The plots in Fig. 2a, b are divided into two main regions. The first region (pale magenta) shows various attempts by the ANN to identify a protocol with a favorable score function. In contrast, the second region (pale cyan) demonstrates that the ANN has learned a successful trend, and suggests temperature profiles that are mostly monotonically increasing functions from 1200 °C to the temperature upper limit (1300 °C). This indicates that when the ANN encounters a protocol with a high score, it tends to keep moving in that direction. This score evolution can be visualized by analyzing the sequence of the Raman spectra reported in Fig. 2c, d. At the beginning, the ANN proposes a completely random trend: for protocol (PTC) 0 (solid gray line), the temperature falls below 1200 °C, which is the lower experimental limit necessary for graphene growth. This is evident from the absence of the 2D peak in the Raman spectra depicted in Fig. 2c (solid gray line). In contrast, for protocols 1 and 2 (represented by the solid blue and orange lines, respectively, in Fig. 2), there is an improvement: the 2D peak emerges, albeit not as intensely or sharply as in high-quality graphene growth. Subsequently, through additional generations, the ANN further enhances the graphene quality, as evidenced by protocols 4 and 5 (indicated by the solid green and red lines, respectively, in Fig. 2). In these instances, a distinct and sharp 2D peak is observed. Figure 2d displays a zoomed-in energy range of the 2D Raman Peak for the protocols that yield improved scores, normalized with respect to the SiC peak. It's noteworthy that while the intensity value of the 2D peak for protocol 4 closely resembles that of protocol 5, the score for protocol 4 consistently falls below that of protocol 5 (Solid black line in Fig. 2b). This discrepancy can likely be attributed to the broader FWHM of the 2D peak in protocol 4, indicating the presence of a BLG. Thus, this reaffirms the significance of the interference term within the score evaluation formula. ANN training is also conducted by removing the lower limit temperature constraint (see Supplementary Note 4, Fig. S3). It is evident that, even lifting this constraint, the ANN successfully learns the desired patterns in a similar manner with a slightly higher number of protocols. Fluctuations in the score during the later stages of training are an intrinsic aspect of the aMC algorithm's stochastic nature. Rather than indicating instability, these variations are a key feature that helps prevent premature convergence on suboptimal solutions. Initially, the temperature profiles generated by the ANN start at T lower than 1200 °C. However, over time, the ANN adapts and begins to suggest temperature profiles that

correspond to effective graphene synthesis, generating more consistent protocols characterized by a steep, monotonic temperature increase with expected stochastic variations, leading to a lower score. The convergence of the ANN to an optimal growth condition is reflected in its ability to consistently pick up such a trend with occasional stochastic variations. This adaptation is noticeable, as in the previous run, in the pale red region of the plot, where the suggested temperature profiles align with known effective growth conditions.

To visualize the learning process and validate our scoring mechanism, we conducted a series of cross-checked experiments using AFM, XPS and ARPES, providing a comprehensive set of surface characterization tools. Starting with AFM, adhesion force maps are shown in Fig. 3a–d. Adhesion AFM depends on the interactions between the probe tip and the sample surface, which can be influenced by factors such as surface roughness, chemical composition, and the presence of contaminants or adsorbates⁵³. Hence, it can serve as a valuable method to quantify the amount and quality of graphene on the surface.

The adhesion force is strictly correlated to the material Young's modulus⁵⁴. The darkest areas (i.e., low adhesion force) represent the SiC surface, while the lightest area (i.e., high adhesion force) represents the graphene surface. It is possible to note that with the score improvement there is an increase of the graphene area from 22.4% (protocol 1, Fig. 3a) to 88.2% (protocol 5, Fig. 3d). For protocol 4, the adhesion map shows a high percentage coverage of graphene. However, the relative adhesion force values are less distinct. Generally, the Young's modulus of MLG is higher than that of BLG⁵⁵, thus providing an explanation for the different adhesion contrast and confirming the presence of BLG on the sample obtained from protocol 4. The percentage of the graphene surface as a function of the protocol number is reported in Fig. S4 of the Supplementary Note 5.

The AFM analysis is cross-checked looking at the chemical properties of the samples measuring the core-level spectrum of each sample via XPS. The results are presented in Fig. 3e–h. The spectral intensities are normalized to facilitate comparison among the different samples. We found that the sp^2 position peak occurs at around 284.4 ± 0.1 eV. The SiC component is identified at 283.8 ± 0.1 eV, while for the components associated with the buffer layer, S1 and S2, the peaks are centered at 285.2 ± 0.1 eV and 286.0 ± 0.1 eV, respectively, in good agreement with the literature^{36,56}. The core-level fitting procedure used in this work is detailed in the Supplementary Note 6, "XPS fitting procedure" section (see Fig. S5 and Table S3). The area subtended by the graphene curve is 18.94% for protocol 1 (Fig. 3e), 19.60% for protocol 2 (Fig. 3f), 31.30% for protocol 4 (Fig. 3g), and 23.10%

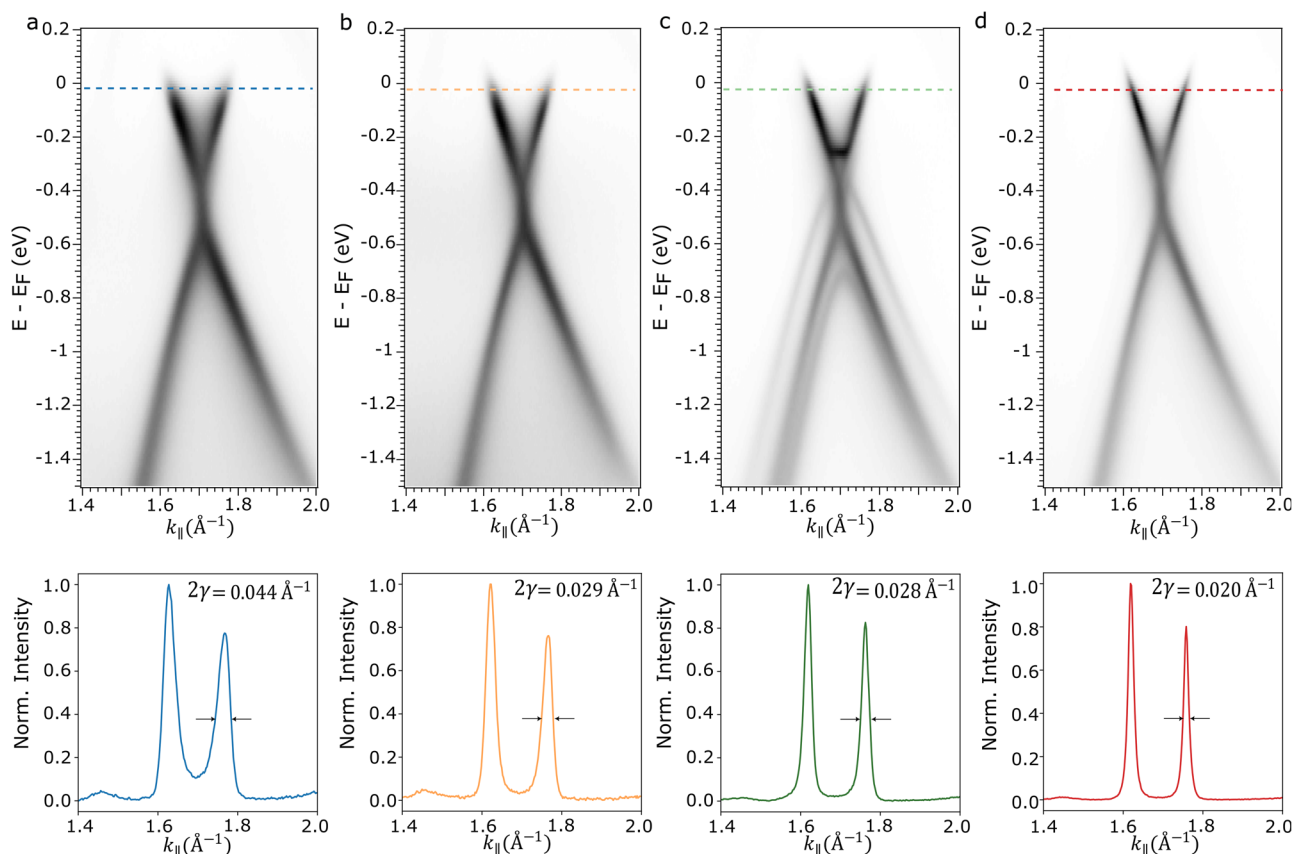


Fig. 4 | ARPES characterization. (Top Panels) ARPES intensity maps collected on samples obtained from different protocols (PTC): PTC1 (a) PTC2 (b) PTC4 (c) and PTC5 (d). (Bottom Panels) The corresponding normalized momentum distribution

curves (MDC) spectra, obtained by integrating the signal at 50 meV below the Fermi level selected to maximize signal to noise level (indicated with colored dashed lines in the ARPES spectra).

for protocol 5 (Fig. 3h), also observing in this case an improvement in the quality of the graphene with the increase in the score. It is important to note that the FWHM of the graphene curve for protocol 4 is slightly higher compared to the other protocols, due to the Van der Waals interactions between two graphene layers, further confirming the presence of BLG in the sample obtained through protocol 4 in great agreement with the AFM analysis. In Fig. S6 we report the trend of the percentage of the graphene area obtained from the XPS fit as a function of the protocol number. The dashed black line indicates the threshold of the area for a fully covered sample. This means that for a percentage of area higher than $\sim 26\%$, there is the presence of BLG on the substrate, as shown in the case of the sample obtained with PTC4.

Finally, the band structure of the as-grown samples is investigated using ARPES (see Fig. 4). This technique is ideally suited to probe the electronic band structure of a material and can provide a definitive assessment of graphene quality. The spectra are collected at the \bar{K} -point of the graphene Brillouin zone, along the $\Gamma\bar{K}$ direction in reciprocal space. We use the sharpness of the graphene bands as a metric to quantify the quality of the relative protocol. We fit the momentum distribution curves (MDC), extracted 50 meV below the Fermi level (colored dashed lines in Fig. 4a–d) to achieve a better fit of the curves, reported in the bottom panels of Fig. 4: the FWHM (2γ) of the Voigt function used to fit the peaks gradually decreases from 0.044 \AA^{-1} (protocol 1, Fig. 4a) to 0.020 \AA^{-1} (protocol 5, Fig. 4d). As expected, the ARPES spectrum collected on the sample obtained from protocol 4, displayed in Fig. 4c, shows the classic band dispersion of a BLG on SiC⁵⁷, overlapping to the MLG bands, with a splitting of the π band originated from the interlayer interaction between the two graphene layers.

These supplementary chemical and structural analyses show a trend that align closely with the calculated score function, validating the use of Raman characterization for the score function in ANN training.

Conclusions

We have demonstrated the potential for autonomous synthesis by using an adaptive learning algorithm to train an artificial neural network that encodes a time-dependent synthesis protocol. The neural network has iteratively and autonomously learned to synthesize high-quality graphene, with minimal initial input, thus demonstrating its ability to learn and adapt. This capability resulted in progressively improved graphene quality, verified through comprehensive surface characterization techniques. The success of this approach highlights the ANN's robust learning mechanisms and adaptability, underscoring its capacity to handle complex material synthesis tasks. By embedding intelligent decision-making into the synthesis process, our work paves the way for future applications of AI in the growth of various 2D materials. This advancement not only promises significant efficiency and quality improvements in material production but also opens new frontiers in material science, where AI-driven methods can explore and optimize previously unattainable synthesis pathways.

Data availability

All data supporting the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The code used in this study is available on GitHub at https://github.com/LeonardoSabattini/amc_oven.git.

Received: 30 October 2024; Accepted: 2 April 2025;

Published online: 25 April 2025

References

- Kumbhakar, P. et al. Prospective applications of two-dimensional materials beyond laboratory frontiers: A review. *iScience* **26**, 106671 (2023).
- Ahn, E. C. 2D materials for spintronic devices. *npj 2D Mater. Appl.* **4**, 17 (2020).
- Jeong, G. H. et al. Nanoscale assembly of 2d materials for energy and environmental applications. *Adv. Mater.* **32**, 1907006 (2020).
- Zhang, Z. et al. Growth and applications of two-dimensional single crystals. *2D Mater.* **10**, 032001 (2023).
- Fan, F. R., Wang, R., Zhang, H. & Wu, W. Emerging beyond-graphene elemental 2d materials for energy and catalysis applications. *Chem. Soc. Rev.* **50**, 10983–11031 (2021).
- Banszerus, L. et al. Ultrahigh-mobility graphene devices from chemical vapor deposition on reusable copper. *Sci. Adv.* **1**, e1500222 (2015).
- Pezzini, S. et al. High-quality electrical transport using scalable CVD graphene. *2D Mater.* **7**, 041003 (2020).
- Gebeyehu, Z. M. et al. Decoupled high-mobility graphene on Cu(111)/sapphire via chemical vapor deposition. *Adv. Mater.* **36**, e2404590 (2024).
- Bhaviripudi, S., Jia, X., Dresselhaus, M. S. & Kong, J. Role of kinetic factors in chemical vapor deposition synthesis of uniform large area graphene using copper catalyst. *Nano Lett.* **10**, 4128–4133 (2010).
- Misekic, V. et al. Deterministic patterned growth of high-mobility large-crystal graphene: a path towards wafer scale integration. *2D Mater.* **4**, 021004 (2017).
- Giambra, M. A. et al. Wafer-scale integration of graphene-based photonic devices. *ACS nano* **15**, 3171–3187 (2021).
- Sun, B. et al. Synthesis of wafer-scale graphene with chemical vapor deposition for electronic device applications. *Adv. Mater. Technol.* **6**, 2000744 (2021).
- Jiang, B., Wang, S., Sun, J. & Liu, Z. Controllable synthesis of wafer-scale graphene films: challenges, status, and perspectives. *Small* **17**, 2008017 (2021).
- Alam, S., Chowdhury, M. A., Shahid, A., Alam, R. & Rahim, A. Synthesis of emerging two-dimensional (2d) materials—advances, challenges and prospects. *FlatChem* **30**, 100305 (2021).
- Xu, X. et al. Growth of 2d materials at the wafer scale. *Adv. Mater.* **34**, 2108258 (2022).
- Liu, X. & Hersam, M. C. 2d materials for quantum information science. *Nat. Rev. Mater.* **4**, 669–684 (2019).
- Turunen, M. et al. Quantum photonics with layered 2d materials. *Nat. Rev. Phys.* **4**, 219–236 (2022).
- Kumar, Y., Gupta, S., Singla, R. & Hu, Y.-C. A systematic review of artificial intelligence techniques in cancer prediction and diagnosis. *Arch. Comput. Methods Eng.* **29**, 2043–2070 (2022).
- Khan, J. et al. Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks. *Nat. Med.* **7**, 673–679 (2001).
- Bhinder, B., Gilvary, C., Madhukar, N. S. & Elemento, O. Artificial intelligence in cancer research and precision medicine. *Cancer Discov.* **11**, 900–915 (2021).
- Fujiyoshi, H., Hirakawa, T. & Yamashita, T. Deep learning-based image recognition for autonomous driving. *IATSS Res.* **43**, 244–252 (2019).
- Atakishiyev, S., Salameh, M., Yao, H. & Goebel, R. Explainable artificial intelligence for autonomous driving: A comprehensive overview and field guide for future research directions. *IEEE Access*, **12**, 101603–101625 (2024).
- Singh, A. & Li, Y. Reliable machine learning potentials based on artificial neural network for graphene. *Comput. Mater. Sci.* **227**, 112272 (2023).
- Dong, Y. et al. Bandgap prediction by deep learning in configurationally hybridized graphene and boron nitride. *npj Comput. Mater.* **5**, 26 (2019).
- Elapolu, M. S., Shishir, M. I. R. & Tabarraei, A. A novel approach for studying crack propagation in polycrystalline graphene using machine learning algorithms. *Comput. Mater. Sci.* **201**, 110878 (2022).
- Elbaz, Y., Furman, D. & Caspary Toroker, M. Modeling diffusion in functional materials: from density functional theory to artificial intelligence. *Adv. Funct. Mater.* **30**, 1900778 (2020).
- Behler, J. & Parrinello, M. Generalized neural-network representation of high-dimensional potential-energy surfaces. *Phys. Rev. Lett.* **98**, 146401 (2007).
- Li, H. et al. Improving the accuracy of density-functional theory calculation: The genetic algorithm and neural network approach. *J. Chem. Phys.* **126**, 144101 (2007).
- Schleder, G. R., Padilha, A. C., Acosta, C. M., Costa, M. & Fazzio, A. From DFT to machine learning: recent approaches to materials science—a review. *J. Phys. Mater.* **2**, 032001 (2019).
- Merchant, A. et al. Scaling deep learning for materials discovery. *Nature* **624**, 80–85 (2023).
- Szymanski, N. J. et al. An autonomous laboratory for the accelerated synthesis of novel materials. *Nature* **624**, 86–91 (2023).
- Barros, N., Whitelam, S., Ciliberto, S. & Bellon, L. Learning efficient erasure protocols for an underdamped memory. *Phys. Rev. E* **111**, 044114 (2025).
- Whitelam, S. How to train your demon to do fast information erasure without heat production. *Phys. Rev. E* **108**, 044138 (2023).
- Whitelam, S. Demon in the machine: learning to extract work and absorb entropy from fluctuating nanosystems. *Phys. Rev. X* **13**, 021005 (2023).
- Casert, C. & Whitelam, S. Learning protocols for the fast and efficient control of active matter. *Nat. Commun.* **15**, 9128 (2024).
- Emtsev, K. V. et al. Towards wafer-size graphene layers by atmospheric pressure graphitization of silicon carbide. *Nat. Mater.* **8**, 203–207 (2009).
- Berger, C. et al. Ultrathin epitaxial graphite: 2d electron gas properties and a route toward graphene-based nanoelectronics. *J. Phys. Chem. B* **108**, 19912–19916 (2004).
- Van Bommel, A., Crombeen, J. & Van Tooren, A. Leed and auger electron observations of the sic(0001) surface. *Surf. Sci.* **48**, 463–472 (1975).
- Emtsev, K., Speck, F., Seyller, T., Ley, L. & Riley, J. D. Interaction, growth, and ordering of epitaxial graphene on sic{0001} surfaces: a comparative photoelectron spectroscopy study. *Phys. Rev. B* **77**, 155303 (2008).
- Varchon, F. et al. Electronic structure of epitaxial graphene layers on sic: effect of the substrate. *Phys. Rev. Lett.* **99**, 126805 (2007).
- Goler, S. et al. Revealing the atomic structure of the buffer layer between sic(0001) and epitaxial graphene. *Carbon* **51**, 249–254 (2013).
- Whitelam, S., Selin, V., Benlolo, I., Casert, C. & Tamblyn, I. Training neural networks using metropolis monte carlo and an adaptive variant. *Mach. Learn. Sci. Technol.* **3**, 045026 (2022).
- Ferrari, A. C. & Basko, D. M. Raman spectroscopy as a versatile tool for studying the properties of graphene. *Nat. Nanotechnol.* **8**, 235–246 (2013).
- Saito, R., Hofmann, M., Dresselhaus, G., Jorio, A. & Dresselhaus, M. Raman spectroscopy of graphene and carbon nanotubes. *Adv. Phys.* **60**, 413–550 (2011).
- Röhrl, J. et al. Raman spectra of epitaxial graphene on SiC(0001). *Appl. Phys. Lett.* **92**, 201918 (2008).
- Mitchell, M. *An Introduction to Genetic Algorithms* (MIT Press, 1998).
- Holland, J. H. Genetic algorithms. *Sci. Am.* **267**, 66–73 (1992).
- Whitelam, S. & Tamblyn, I. Learning to grow: control of material self-assembly using evolutionary reinforcement learning. *Phys. Rev. E* **101**, 052604 (2020).

49. Whitelam, S., Selin, V., Park, S.-W. & Tamblyn, I. Correspondence between neuroevolution and gradient descent. *Nat. Commun.* **12**, 6317 (2021).
 50. Zhang, X., Tan, Q.-H., Wu, J.-B., Shi, W. & Tan, P.-H. Review on the Raman spectroscopy of different types of layered materials. *Nanoscale* **8**, 6435–6450 (2016).
 51. Paillet, M., Parret, R., Sauvajol, J.-L. & Colombari, P. Graphene and related 2d materials: an overview of the Raman studies. *J. Raman Spectrosc.* **49**, 8–12 (2018).
 52. Cong, X., Liu, X.-L., Lin, M.-L. & Tan, P.-H. Application of Raman spectroscopy to probe fundamental properties of two-dimensional materials. *npj 2D Mater. Appl.* **4**, 13 (2020).
 53. Lee, F., Tripathi, M., Lynch, P. & Dalton, A. B. Configurational effects on strain and doping at graphene-silver nanowire interfaces. *Appl. Sci.* **10**, 5157 (2020).
 54. Johnson, K. L., Kendall, K. & Roberts, A. Surface energy and the contact of elastic solids. *Proc. R. Soc. Lond. A Math. Phys. Sci.* **324**, 301–313 (1971).
 55. Cao, Q. et al. A review of current development of graphene mechanics. *Crystals* **8**, 357 (2018).
 56. Forti, S. et al. Electronic properties of single-layer tungsten disulfide on epitaxial graphene on silicon carbide. *Nanoscale* **9**, 16412–16419 (2017).
 57. Forti, S. & Starke, U. Epitaxial graphene on sic: from carrier density engineering to quasi-free standing graphene by atomic intercalation. *J. Phys. D: Appl. Phys.* **47**, 094013 (2014).
- F.B. All authors contributed to the scientific discussion and provided feedback on the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s42005-025-02086-1>.

Correspondence and requests for materials should be addressed to Stephen Whitelam, Camilla Coletti or Antonio Rossi.

Peer review information *Communications Physics* thanks the anonymous reviewers for their contribution to the peer review of this work. A peer review file is available.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License, which permits any non-commercial use, sharing, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if you modified the licensed material. You do not have permission under this licence to share adapted material derived from this article or parts of it. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by-nc-nd/4.0/>.

© The Author(s) 2025

Acknowledgements

We acknowledge the project PNRR MUR Project PE000013 CUP J53C22003010006 Future Artificial Intelligence Research (FAIR) and PNRR MUR Project PE000023 CUP J53C22003200005 National Institute of Quantum Science and Technology (NQSTI) funded by the European Union – Next Generation EU.

Author contributions

A.R., S.W., E.S.B., and C.Co. conceived the study and designed the experiments. A.R., L.S., and A.C. carried out sample growth, Raman spectroscopy, data analysis, and ANN training. C.Co. and S.W. developed the code that implements the learning algorithm. XPS and ARPES measurements were performed by A.R. and S.F. Data analysis was conducted by A.C., S.F., and A.R. The manuscript was written by A.R., L.S., S.W., and A.C. Financial support was provided by C.Co., M.P., S.W., and