

ARTICLE OPEN



Advanced machine learning decision policies for diameter control of carbon nanotubes

Rahul Rao¹✉, Jennifer Carpena-Núñez^{1,2}, Pavel Nikolaev^{1,2}, Michael A. Susner¹, Kristofer G. Reyes¹✉ and Benji Maruyama¹

The diameters of single-walled carbon nanotubes (SWCNTs) are directly related to their electronic properties, making diameter control highly desirable for a number of applications. Here we utilized a machine learning planner based on the Expected Improvement decision policy that mapped regions where growth was feasible vs. not feasible and further optimized synthesis conditions to selectively grow SWCNTs within a narrow diameter range. We maximized two ranges corresponding to Raman radial breathing mode frequencies around 265 and 225 cm^{-1} (SWCNT diameters around 0.92 and 1.06 nm, respectively), and our planner found optimal synthesis conditions within a hundred experiments. Extensive post-growth characterization showed high selectivity in the optimized growth experiments compared to the unoptimized growth experiments. Remarkably, our planner revealed significantly different synthesis conditions for maximizing the two diameter ranges in spite of their relative closeness. Our study shows the promise for machine learning-driven diameter optimization and paves the way towards chirality-controlled SWCNT growth.

npj Computational Materials (2021)7:157; <https://doi.org/10.1038/s41524-021-00629-y>

INTRODUCTION

The utilization of single-walled carbon nanotubes (SWCNTs) in a wide range of applications, including “Beyond Moore’s Law” computing¹, hinges on our ability to control their physical properties². A key property is the tube diameter, which is intimately related to its electronic properties³, making diameter control of SWCNTs important for semiconducting device applications⁴. In addition, control over SWCNT diameters is desirable for applications such as molecular transport⁵, membranes⁶, field emission⁷, and sensing⁸. To date, there are many published reports on diameter-controlled growth of SWCNTs by chemical vapor deposition (CVD), the most popular synthesis method. As the SWCNT diameter is defined primarily by the size of the catalyst particle^{9,10}, most of these reports focus on synthesis from catalysts with defined sizes or compositions^{11–16}. Other reported alternatives for controlling SWCNT diameters involve careful selection of the growth substrates^{17,18} and manipulation of the type and amount of hydrocarbon feedstocks^{19–22}.

A major drawback of trying to control particle shape and size during CVD growth is the inherent difficulty in mitigating several competing effects that occur at high growth temperatures (typically 600–900 °C). These effects include diffusion of particles on and into the substrate, and particle coarsening^{23–25}, which result in the broadening of the particle size distributions and, consequently, SWCNT diameters. One way to reduce these effects is to lower the growth temperature or develop schemes to immobilize the particles, and such methods have had limited success, although often at the cost of a decrease in yield²⁶. On the other hand, manipulation of the hydrocarbon feedstock can lead to contrary results, owing to the wide range of feedstocks used for growth²⁷. For example, higher amounts of hydrocarbon precursors such as ethane, ethylene, and methane result in larger tube diameters^{19,21,22}, whereas others such as CO result in smaller diameters²². These effects are further complicated with the inclusion of additives such as NH_3 , H_2O , or CO_2 , which can

increase or decrease the average diameter of the SWCNTs through etching or selective catalyst particle deactivation^{28–30}.

The precursors and additives used in CVD can be classified as catalyst reducing (e.g., hydrocarbons and CO) or catalyst oxidizing (H_2O and CO_2). The conflicting results published in the literature suggest that the SWCNT diameter could be tuned by balancing the ratios of reductants and oxidants, in addition to altering the synthesis temperature. Unfortunately, the few systematic studies reported so far have been restricted to the influence of only one or two growth parameters^{21,31}. Here we considered three synthesis parameters, namely temperature and the amounts (partial pressures) of reductants (ethylene, acetylene, and hydrogen) and oxidants (CO_2 and system water vapor) to tune the growth of SWCNTs with diameters in a narrow range. We achieved this by employing a machine learning (ML) planner to optimize diameter-selective growth of SWCNTs from a 1 nm-thick cobalt catalyst film.

ML algorithms are increasingly being utilized in materials research and can optimize a variety of processes^{32–38}. We have also previously used ML (random forest and Bayesian optimization) to optimize the average growth rate and the yield of SWCNTs using in situ Raman spectroscopy as the feedback method^{39,40}. Here we focused on optimization of the selective growth of small-diameter carbon nanotubes (CNTs) through their diameter-dependent low-frequency Raman radial breathing modes (RBMs). An ML planner based on the Expected Improvement (EI) decision policy was used to map the growth parameter phase spaces for maximizing the intensities of two RBMs at 225 and 265 cm^{-1} , which correspond to SWCNT diameters around 1.06 and 0.91 nm, respectively. Analysis of the RBM distributions with four laser excitation wavelengths showed that roughly a third of the SWCNTs in the optimized experiments lie in the desired diameter range. Notably, the growth parameter space for maximizing the intensities of these two RBMs was significantly different in spite of their closeness in diameters. Our results show that it is indeed

¹Materials and Manufacturing Directorate, Air Force Research Laboratory, WPAFB, Dayton, OH 45433, USA. ²UES, Inc., Dayton, OH 45423, USA. ³Department of Materials Design and Innovation, University at Buffalo, Buffalo, NY 14260, USA. ✉email: rahul.rao.2@us.af.mil; kreyes3@buffalo.edu

possible to use ML to control SWCNT diameter and point the path forward towards scaled and selective synthesis of diameter-controlled CNTs.

RESULTS AND DISCUSSION

SWCNT growth

The SWCNT growth experiments were performed in ARESTM (Autonomous Research System), an automated high-throughput laser-induced cold-wall CVD system coupled to a Raman spectrometer (experimental details in the “Methods” section)³⁹. In the past, ARES has provided insights into chiral angle-dependent growth rates and defects in SWCNTs^{41,42}, optimal phase spaces for CNT growth^{43–46}, new alloy catalyst compositions for SWCNT growth⁴⁷, and into defect evolution in two-dimensional materials^{48–51}. Here we grew SWCNTs from a 1 nm Co catalyst film deposited using ion beam sputtering onto 10 nm Al₂O₃-coated Si micropillars. We used ethylene and acetylene as the hydrocarbon source. It is noteworthy that both ethylene and acetylene can reduce the catalyst film, and are therefore considered reductant species⁵². Along with the hydrocarbons, we used hydrogen as an additional reductant and CO₂ as the oxidant. We also measured the base water vapor pressure in the system, which was <1 p.p.m. The growth temperatures and pressures ranged between 500 and 900 °C, and between 2 and 50 Torr, respectively. Under these conditions, we typically grow SWCNTs in the form of bundled mats^{41,43–45}. Supplementary Fig. 1 shows a representative scanning electron microscope image of the SWCNTs.

Figure 1 shows a Raman spectrum in the low-frequency RBM range, collected from one of our typical growth experiments, which was not optimized for SWCNT diameter selectivity. This spectrum was collected post growth in ARES, with an excitation wavelength of 532 nm. At least four distinct RBMs at 162, 189, 225, and 265 cm^{−1} can be identified in the spectrum. The Raman spectrum is plotted beneath the so-called Kataura plot⁵³, which maps the electronic transition energies of each SWCNT as a function of its diameter (top horizontal axis in Fig. 1, calculated

according to the relation $\omega_{\text{RBM}} = 235.9/d_t + 5.5$ ⁵⁴, where d_t and ω_{RBM} are the diameter and RBM frequency, respectively). The Raman intensity of a SWCNT is enhanced when it is in resonance, i.e., when the excitation laser energy (denoted in the top panel in Fig. 1 by the solid horizontal line and is 2.33 eV in our ARES experiments) is close to one of the electronic transition energies of the SWCNT⁵⁵. The resonance-enhanced intensity occurs over an energy window of ± 0.1 eV, indicated by the two dashed horizontal lines in Fig. 1. Thus, all nanotubes within the window can appear in the spectrum.

By comparing the experimental Raman spectrum to the Kataura plot, we see that the RBM at 225 cm^{−1} can be matched to one of three SWCNTs of (*n,m*) chiralities: (11,5), (12,3), and (9,6), which are within a frequency range of 225 ± 20 cm^{−1}, whereas the number of matches for the RBM at 265 cm^{−1} is lower—(8,5) or (9,3). The number of matches for the RBM at 189 cm^{−1} is higher—(14,4), (13,6), (15,2), or (16,0)—and there are many possible matches for the RBM at 162 cm^{−1}. Owing to the lower number of possible matches, in this study we chose to maximize the growth of SWCNTs with RBMs in the range of 265 ± 10 cm^{−1} (diameter 0.91 ± 0.4 nm) and 225 ± 20 cm^{−1} (diameter 1.07 ± 0.11 nm), hereafter referred to as RBMs at 265 and 225 cm^{−1}, respectively.

Support vector machine-constrained EI policy

The campaign’s objective was to identify optimal CNT synthesis conditions that achieve two scientific goals, namely maximizing the number of SWCNTs of a particular diameter and selectively minimizing the synthesis of different-diameter SWCNTs. To achieve this, we framed the experimental parameter search as an optimization of a scalar objective function (OF) that considers the two goals described above and used a constrained version of the EI policy to select experiments. The policy was additionally constrained by a feasible set of synthesis conditions dynamically learned using a support vector machine (SVM) classifier in an attempt to explicitly avoid low OF values, which is an augmentation of the natural balance that EI achieves.

The desired outcome of our experiments was diameter selectivity and optimal yield. To measure how well this was achieved for SWCNTs synthesized with growth conditions \mathbf{x} , we considered a scalar utility OF

$$f(\mathbf{x}) = \frac{1}{1 + \exp[-\beta(\frac{A}{\sigma} - \alpha)]} \cdot \frac{A}{B - A} \quad (1)$$

where A is the area under the Raman curve for a particular RBM peak of interest (i.e., 265 ± 10 or 225 ± 20 cm^{−1}) and B is the area under the broader RBM spectrum region (100–350 cm^{−1}). The non-dimensionalized quantity $A/(B - A)$ measures selectivity of peaks of interest with respect to the rest of the broader region. We treated both A and B as functions of the synthesis parameters \mathbf{x} . A prefactor was incorporated with the OF to diminish the value for cases where the signal-to-background for A was low. The selectivity was balanced by the sigmoidal prefactor, taking values between 0 and 1. This depended on the absolute size of the peak A and its value was near 1 if the normalized quantity $\frac{A}{\sigma}$ surpassed a predefined threshold α , which was set to 100 counts. In Eq. (1), σ was set to the background level of the Raman spectrum, whereas the inverse length-scale parameter $\beta = 1/10$ (counts^{−1}) determined how fast the prefactor scaled from 0 to 1 for values $\frac{A}{\sigma}$ near the threshold α (further discussion on these model hyperparameters is presented in Supplementary Section 1). Although we have selected to fix the values based on visual inspection of the rankings the OF produced, these and other hyperparameters could be alternatively optimized through maximum likelihood optimization.

Five experimental inputs \mathbf{x} were used, including the partial pressures of the reductants (ethylene, acetylene, and hydrogen), oxidants (water vapor and carbon dioxide), and the growth

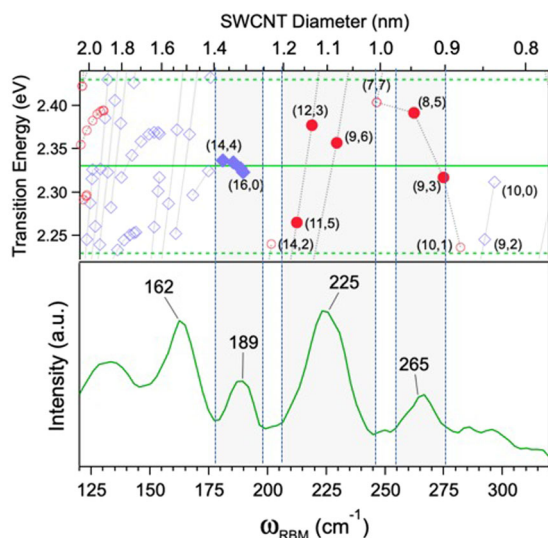


Fig. 1 Radial breathing modes matched to SWCNT chiral indices. Bottom: representative Raman spectrum in the low-frequency RBM range from a typical non-optimized growth experiment in ARES (laser excitation at 532 nm). Top: Kataura plot showing the possible matches for SWCNTs (solid red data points) within the chosen frequency ranges (shaded areas). There are three matches for the RBM at 225 cm^{−1} and two for the 265 cm^{−1} peak. The dashed horizontal lines in the Kataura plot represent the resonance window (± 0.1 eV) for the SWCNTs.

temperature. As mentioned above, water vapor was found in trace amounts in our system (<1 p.p.m.) and was also accounted for in our partial pressure calculations. We have previously shown that the growth of SWCNTs is heavily mediated by a balance of two reduced order parameters forming the input vector \mathbf{z} , which is derived from such experimental inputs \mathbf{x} : (1) temperature itself and (2) the ratio between the partial pressures of reducing and oxidizing species (pressure ratio)⁵⁶. By combining the partial pressures of all the active precursors into a ratio of reductants over oxidants (hereafter called the pressure ratio), we were able to parse the growth results in terms of a reduced two-dimensional space (temperature vs. pressure ratio). These two features could therefore be used to determine whether SWCNTs grew on a catalyst nanoparticle or whether oxidation effects prevented such growth outright. Moreover, the critical region of growth could be delineated in this reduced-dimensional space such that the transition from the feasible “growth” phase to the infeasible “no-growth” phase was sharp. Throughout the campaign, we assumed that the location of this phase boundary in the reduced-dimensional space was unknown a priori, and that planned experiments had to be selected and run, to learn the true feasible region in addition to the primary goal of discovering optimal synthesis. Typically, phase-change phenomena leading to abrupt differences in experimental responses due to varying synthesis conditions across phase boundaries may not be properly modeled using Gaussian processes⁵⁷, because they often presume smoothness in the functions they model. Thus, avoiding such “no-growth” experiments and data lying on the phase boundary itself can help train models in which such smoothness constraints are fully satisfied by the data properly inside the “growth” region, far from the “no-growth” region. Hence, when deciding what experiments to run, our planner avoided suggesting synthesis conditions \mathbf{x} that mapped $\mathbf{z} = h(\mathbf{x})$ to the low-dimensional features \mathbf{z} lying in the as-predicted “no-growth” region of the reduced-dimensional space.

To model the phase boundary in the reduced-dimensional space, we used a kernel SVM classifier and automatically labeled the experimental data into “growth” or “no-growth” classes using a thresholding parameter applied to the observed OF values. That is, given the experimental data $(\mathbf{x}, f(\mathbf{x}))$, we trained the kernel SVM classifier on the data $(h(\mathbf{x}), \chi_{f(\mathbf{x}) > f_{\min}})$, where $\chi_{f(\mathbf{x}) > f_{\min}}$ is an indicator function

$$\chi_{f(\mathbf{x}) > f_{\min}} = \begin{cases} 1 & \text{if } f(\mathbf{x}) \geq f_{\min}, \\ -1 & \text{if } f(\mathbf{x}) < f_{\min}. \end{cases} \quad (2)$$

We used the squared-exponential SVM kernel and a threshold value of $f_{\min} = 0.01$ (arbitrary units).

Given a data set of n experimental inputs \mathbf{x} and the corresponding observed (and presumed noisy) OF responses $f(\mathbf{x})$, we formed a Gaussian process Bayesian belief $\mathcal{GP}(\mu^n, \Sigma^n)$, where $\mu^n(\mathbf{x})$ is the time- n prediction of the true OF and $\Sigma^n(\mathbf{x}, \mathbf{x}')$ is the corresponding covariance function that provided the predicted statistical relationship between the OF function values at pairs of inputs \mathbf{x} and \mathbf{x}' . This Gaussian process belief was used for the EI decision-making policy (discussed in greater detail in section Supplementary Section 2)⁵⁸, which attempted to select the next experiment to run with the goal of balancing between resolving uncertainties in the prediction of the OF (called exploration) vs. focusing on promising regions of the response, based on current beliefs (exploitation). Such a balance between exploration and exploitation must be maintained in iterative, closed-loop experimental campaigns due to the lack of data available for most campaigns. Initial beliefs, which are based on a limited set of seed data, may have inaccurate predictions. Therefore, it would be premature to trust such predictions, in particular optimal regions delineated by such predictions. Conversely, the campaign had a specific optimization objective and so

a limited experimental budget should not be wasted in generic learning of the response in order to build a globally accurate surrogate model. This balance was captured by the EI policy acquisition function $\text{Acq}(\mathbf{x})$, which evaluated potential experiments to run next according to a predefined measure of the explorative and exploitative value of running the experiment. The policy selected the next experiment as the one that maximized $\text{Acq}(\mathbf{x})$, subject to the constraint that the experimental inputs mapped to the feasible region in \mathbf{z} -space predicted by the SVM classifier. That is, the policy selected an experiment by solving the constrained optimization

$$\mathbf{x}^* = \text{Acq}(\mathbf{x}), \quad (3)$$

where \mathcal{X}^n was the set of synthesis conditions \mathbf{x} such that $\mathbf{z} = h(\mathbf{x})$ was predicted to be in the “growth” region by an SVM classifier trained on the n available data points. If the training data does not contain examples of both growth and no-growth experiments, which is typical early in the campaign, the feasible set is simply defined to be the entire domain and the optimization of the acquisition function is unconstrained. Supplementary Fig. 2 shows the evolution of the SVM decision boundary over the first 25 experiments and demonstrates the fact that the SVM is only trained and used to constrain the optimization after 10 experiments.

Maximization of the 265 and 225 cm^{-1} RBMs

Our first target was to maximize the growth of SWCNTs with RBMs around 265 cm^{-1} . Three initial seed growth experiments were conducted in ARES, with growth conditions selected such that they corresponded to our typical experiments as published previously³⁹. Including the three seed experiments, we performed a total of 74 growth experiments. Most of the experiments selected by the planner used the EI policy (learning mode). Exploitation experiments were selected every four to five experiments (exploitation mode) to track the evolution of the estimated best growth conditions. Following each growth experiment, a Raman spectrum was collected at room temperature with a minimal laser power of 0.2 W and 30 s acquisition time. The low-frequency RBM regions between 100 and 350 cm^{-1} in these post-growth spectra were fitted with Lorentzian peaks (an example of a fitted spectrum in the RBM range is shown in Supplementary Fig. 3), enabling the calculation of the quantities A and B , and the OF as described above.

The results of all the growth experiments maximizing the area of the 265 cm^{-1} RBM peak are plotted in Fig. 2a, which shows a heat map of the OF against the growth temperature and the pressure ratio (partial pressures of reductants/oxidants). It is clear from Fig. 2a that the experimental conditions that maximized the 265 cm^{-1} RBM peak were confined to a relatively narrow range (indicated by the red dashed ellipse in Fig. 2a). These conditions corresponded to an average growth temperature around 700°C and oxidizing conditions (pressure ratio ~ 0.015). Figure 2b shows a Raman spectrum collected post growth from a growth experiment corresponding to the highest OF (0.33). It is interesting to note that even for the highest OF, there are two RBMs observed in the spectrum—at 225 and 265 cm^{-1} .

Next, we performed 41 additional experiments to maximize the area of the RBM at 225 cm^{-1} and the result of these growth experiments is shown in Fig. 2c. The post-growth Raman spectrum corresponding to the highest OF is shown in Fig. 2d. Comparing the data in Figs. 2a, c, it is immediately apparent that the optimal temperature and the pressure ratios were different for the two RBMs. Unlike the case for the 265 cm^{-1} peak, the growth conditions that maximized the 225 cm^{-1} peak were more reducing (ratio of reductant/oxidant pressures around 1) and at higher temperatures (average growth temperature $\sim 800^\circ\text{C}$). The differences between the data presented in Fig. 2a, c are

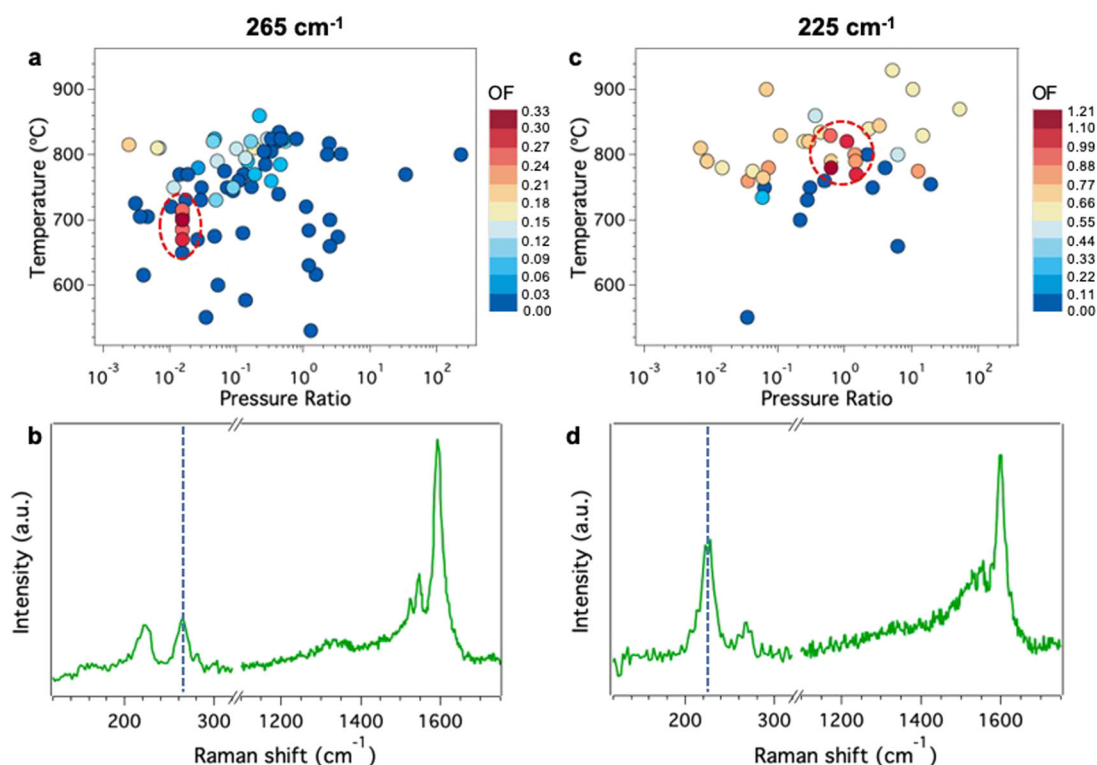


Fig. 2 Map of the reduced dimensionality synthesis space. Temperature vs. pressure ratio for maximizing the areas of the RBM peak at **a** 265 and **c** 225 cm^{-1} . The circled areas in **a** and **c** correspond to the synthesis parameters that produced the highest objective functions (OFs). Post-growth Raman spectra (collected in ARES with 532 nm excitation) showing the low-frequency RBMs and the G band regions for growth experiments that exhibited the highest OF values for RBMs at **b** 265 and **d** 225 cm^{-1} , respectively.

remarkable considering that the diameters of the SWCNTs corresponding to these two RBMs, namely 225 and 265 cm^{-1} , are very close, ~ 1.06 and 0.91 nm, respectively. The observation of totally different synthesis parameters for maximizing SWCNTs with similar diameters highlights the power of our EI decision policy planner, which produced these results in under a hundred experiments. Such a result would normally have taken over several months to produce with traditional synthesis methods, whereas the ARES growth experiments were performed over a couple of weeks.

To map the chiral distribution of SWCNTs in the optimized growth experiments, we performed additional characterization of the RBMs using multi-excitation Raman spectroscopy in a Renishaw inVia Raman system. For this, we used three other laser excitations—514.5, 633, and 785 nm; these three laser excitations cover the majority of SWCNTs with respect to their electronic transition energies and are commonly used for characterizing SWCNT populations using resonance Raman spectra^{59,60}. Two-dimensional Raman spectral maps were collected on a number of ARES micropillars ($11 \times 11 \mu\text{m}$ areas with $1 \mu\text{m}$ spacing between points) and the average spectrum from each map was obtained. To emphasize the effectiveness of our planner, in Fig. 3a, b we show average multi-excitation Raman spectra corresponding to optimization of the 265 cm^{-1} RBM peak with a high (OF = 0.33) and low (OF = 0.03) OF, respectively. Additional multi-excitation Raman spectra from other experiments are shown in Supplementary Fig. 4. All spectra in the RBM region in Fig. 3a, b were normalized with respect to the highest intensity peak within the RBM frequency range ($100\text{--}320 \text{ cm}^{-1}$). The vertical dashed lines in Fig. 3a, b correspond to the RBM range targeted in our experiments: $255\text{--}275 \text{ cm}^{-1}$.

A number of observations can be made from Fig. 3a. First, all the spectra contain more than one RBM. Second, the spectrum collected with only one of the excitations—514.5 nm—contains a

predominant peak at 265 cm^{-1} . The multi-excitation Raman spectra show that most of the RBMs lie within a relatively narrow range of frequencies ($\sim 220\text{--}280 \text{ cm}^{-1}$), corresponding to a diameter range between 0.86 and 1.1 nm. This range is slightly wider than the range requested in the ARES optimization experiments— 0.91 ± 0.04 nm. In contrast, there is no prominent peak at 265 cm^{-1} for the growth experiment that exhibited a low OF (0.03, Fig. 3b); in fact, all the RBMs appear at lower frequencies, indicating that poor optimization (low OF values) resulted in the growth of larger diameter SWCNTs.

To get a clearer idea of the chiral distribution of these SWCNTs, we recast the data from Fig. 3a, b onto the Kataura plots in Fig. 3c, d, where the relative intensities of the peaks in Fig. 3a, b are reflected in the size of the circular data points. We note that the data in these Kataura plots do not show the actual abundance of the chiral distribution, owing to the difficulty in estimating accurate SWCNT densities from the bundled mats on the surface of the Si micropillar. Rather, we show these plots for the purpose of visualizing the average chiral distributions as obtained from the multi-excitation Raman spectra. It is immediately apparent from Fig. 3c, d that we grew both semiconducting and metallic SWCNTs, indicating that there is no preference for conducting type in both populations. Moreover, a tally of all the RBMs by conducting type results in a semiconducting to metallic SWCNT ratio of $65:35$, similar to what is commonly observed in ensembles of CVD-grown SWCNTs⁵⁵. The highlighted rectangular boxes in Fig. 3c, d denote regions with the highest RBM peak intensities and the sharp contrast between the high and low OF can be seen clearly in the overlap (or lack thereof) between the highlighted box and the vertical dashed lines indicating the targeted RBM range.

The results presented in Fig. 3a–d can be compared to the Raman spectra collected from the campaign of experiments to maximize the 225 cm^{-1} RBM peak. The multi-excitation spectra for

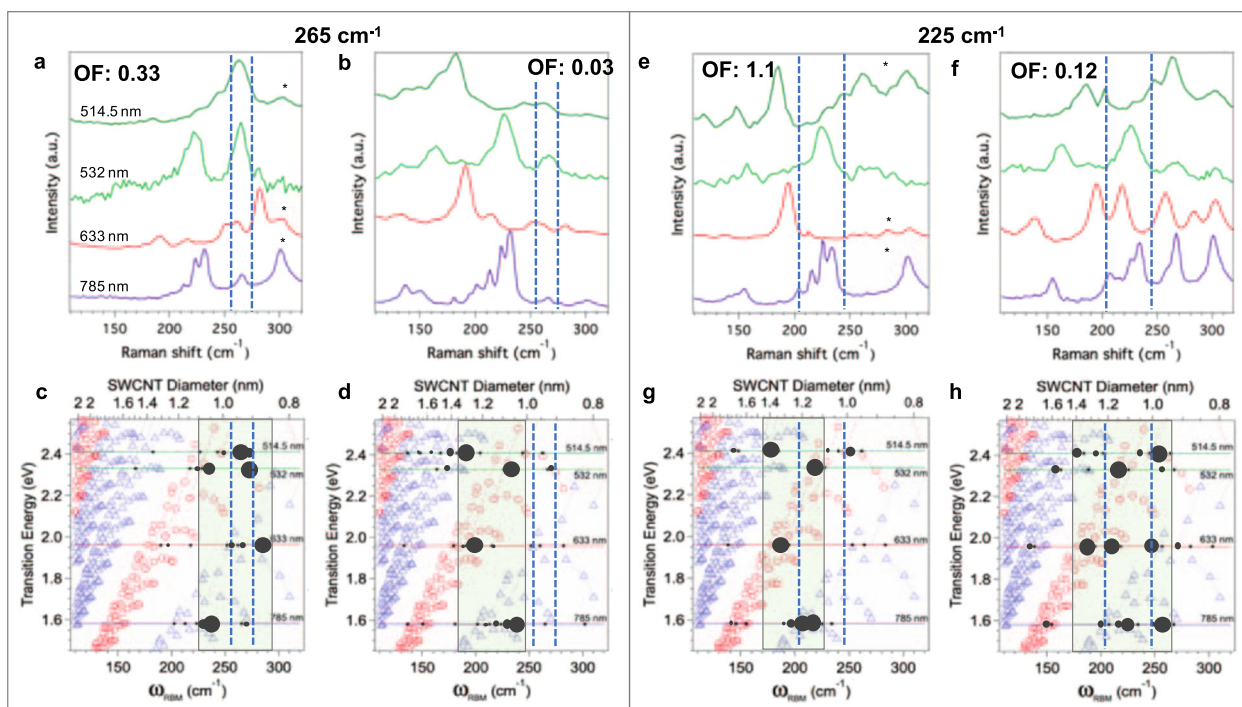


Fig. 3 Multi-excitation Raman spectra. Average Raman spectra in the RBM region from the growth experiments exhibiting high and low OFs for optimization of the **a, b** 265 cm^{-1} and **e, f** 225 cm^{-1} RBM peaks, respectively. The asterisks indicate the second-order peak from the silicon substrate. **c, d, g, h** The peak intensities of the spectra in **a, b, e, f** respectively, plotted on a Kataura plot and showing the average SWCNT chiral distribution. The red circles and blue triangles in the Kataura plot indicate metallic and semiconducting SWCNTs, respectively. The dashed vertical lines in the graphs indicate the RBM frequency ranges targeted by the planner (i.e., $255\text{--}275$ and $205\text{--}245\text{ cm}^{-1}$) and the highlighted boxes in **c, d, g, h** show regions with the most intense RBMs. The good overlaps between the highlighted boxes and the dashed vertical lines in **c** and **g** show a higher degree of diameter selectivity in the experiments that exhibited a high OF.

experiments that exhibited high and low OFs are shown in Fig. 3e, f, and the corresponding Kataura plots in Fig. 3g, h. Similar to the case of the 265 cm^{-1} RBM, the data in Fig. 3e reveals high-intensity RBMs close to 225 cm^{-1} , indicating a narrow diameter distribution. The corresponding Kataura plot in Fig. 3f shows that the high-intensity RBMs correspond to SWCNTs with diameters that range between 1.05 and 1.35 nm . However, the big difference between the high and low OF growths in the case of the 225 cm^{-1} RBM campaign is the much wider SWCNT diameter distribution in the growths with low OF. This can be seen clearly in the larger size of the shaded rectangle in Fig. 3h.

By analyzing the intensities of the RBMs in the multi-excitation Raman spectra for the two growths that exhibited high OFs (Fig. 3a, e for RBMs at 265 and 225 cm^{-1} , respectively), we found that $\sim 31\%$ of the total Raman intensity was in the range $265 \pm 10\text{ cm}^{-1}$. If we expanded the range to $220\text{--}280\text{ cm}^{-1}$ (i.e., SWCNT diameters ranging between 0.86 and 1.1 nm), the RBM intensities accounted for $\sim 71\%$ of the total intensity between 100 and 350 cm^{-1} . This trend is similar in the case of the 225 cm^{-1} RBM: $\sim 35\%$ of the RBM intensity lay between our selected range ($225 \pm 20\text{ cm}^{-1}$) and $\sim 63\%$ if we expanded the range to $180\text{--}230\text{ cm}^{-1}$.

As mentioned above and seen in Fig. 2a, c, the growth conditions for the experiments that resulted in low and high OFs were different. For the 265 cm^{-1} RBM, temperatures for the high and low OF experiments were 700 and 750°C , respectively, whereas they were 820 and 830°C for the 225 cm^{-1} RBM. Although the differences in temperatures could account for shifts in SWCNT diameter distributions, it is likely that the results presented in Fig. 3 were caused by the significantly different pressure ratios between the high and low OF growth experiments. For the growth experiments maximizing the 265 cm^{-1} RBM, the pressure ratios for the high and low OF experiments were 0.015 and 0.33 , respectively, whereas these values were 1.08 and 0.009 ,

respectively, for the 225 cm^{-1} RBM. The orders of magnitude difference between the pressure ratios highlights the effectiveness of tuning SWCNT diameter distributions through control over the feed rates of the gaseous precursors.

SWCNT yield

Efforts to increase selectivity during SWCNT growth typically suffer from a drop in the yield, owing to a variety of factors such as low catalyst activities and difficulty in controlling the size and morphology of a large number of particles during high-temperature CVD growth²⁶. As mentioned above, in our experiments we grew sparse SWCNT bundles on the ARES micropillars, making it difficult to accurately estimate the growth yield through counting or weight gain measurements. However, the overall intensity of the SWCNT G band (G_{max}), collected at the end of each experiment in the post-growth Raman spectrum could be used as a proxy for the yield. In order to see the effect of the planner on the diameter selectivity and yield, the data in Fig. 2 were replotted by multiplying the OFs with the respective G_{max} values (normalized by the intensity of the Si substrate peak at 520 cm^{-1}). Figures 4a, b show the heat maps of the $\text{OF} \times G_{\text{max}}$ against the growth temperatures and pressure ratios for growth experiments maximizing RBMs at 265 and 225 cm^{-1} , respectively. In the case of the 265 cm^{-1} RBM, the high-yield area corresponds to the same temperature and pressure ratios as the area that exhibited high OFs or diameter-selective growth (Fig. 2a). However, the highest yields for the 225 cm^{-1} RBM correspond to different synthesis conditions compared to the regions of high diameter selectivity. As discussed above, the 225 cm^{-1} RBM was maximized at a temperature of $\sim 800^\circ\text{C}$ and a pressure ratio around 1 (Fig. 2c). When we take the SWCNT yields into account, we see that the highest yields occur for more oxidizing conditions,

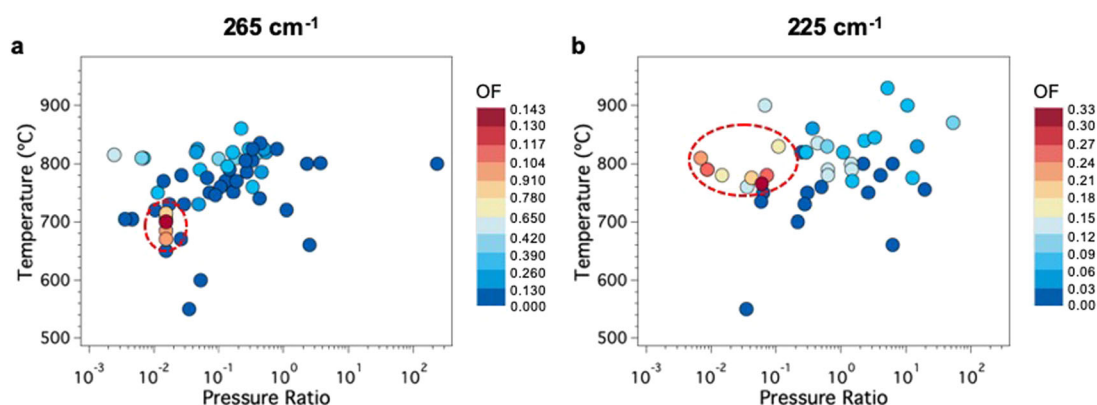


Fig. 4 Map of the reduced dimensionality space accounting for SWCNT yield. Temperature vs. pressure ratio for maximizing the areas of the RBM peak at **a** 265 cm^{-1} and **b** 225 cm^{-1} . The color scale corresponds to the calculated OFs multiplied by the maximum G band intensities. The circled areas in **a** and **b** correspond to the synthesis parameters that produced the highest yield and diameter selectivity.

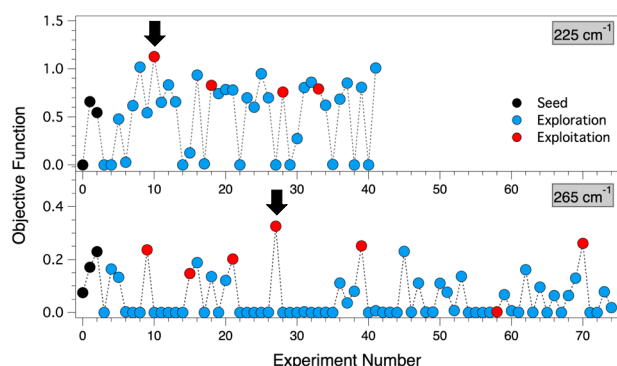


Fig. 5 Progression of the objective function during the campaigns. Calculated OF values over the experimental campaign to maximize RBMs at 265 cm^{-1} (bottom) and 225 cm^{-1} (top).

for pressure ratios between 0.1 and 0.01. In fact, these partial pressures are close to those for the 265 cm^{-1} RBM (~ 0.015 , Fig. 2a). Multi-excitation Raman analysis of the RBMs showed that the intensity of the 225 cm^{-1} RBM ($\pm 20 \text{ cm}^{-1}$) dropped down to $\sim 26\%$. This result, and the differences in the pressure ratios shown in Figs. 4b and 2c corroborate the inverse relationship between yield and selectivity. However, importantly, our studies, enabled by the exploration of a wide range of precursor pressures, show that it is possible to find conditions in the synthesis phase space leading to high diameter selectivity and high yields.

Exploration vs. exploitation analysis

To analyze the efficiency of the planner, in Fig. 5 we plot the OFs over the course of the experimental campaigns to maximize the 265 and 225 cm^{-1} RBMs. As mentioned previously, the first three experiments of both campaigns were manually chosen seed experiments and are indicated by the black data in Fig. 5. Experiments conducted in exploration and exploitation modes are indicated in Fig. 5 by blue and red data points, respectively. Looking at the progression of the OF during the campaign to maximize the RBM at 265 cm^{-1} (bottom panel in Fig. 5), it is clear that the highest OF was achieved in the 27th experiment in the exploitation mode. The experiments that exhibited the highest OFs are indicated by the arrows in Fig. 5. Other than one of the experiments (#58), the OFs in the exploitation mode were also consistently high, averaging 0.24 ± 0.06 through the 74-experiment campaign. The high OFs obtained in the exploitation mode are not surprising considering that the planner chose the temperature and pressure ratios such that they produced high OF values based on the three seed experiments (average 0.16 ± 0.08).

However, it is interesting that the high OF values in both experimental campaigns were greater than the OFs in the seed experiments. On the other hand, not surprisingly, the experimental conditions were chosen over a much wider range of conditions in the learning mode, leading to a large variety in OF values. Similar trends can be seen in the OF values for the other experimental campaign to maximize the RBM at 225 cm^{-1} (top panel in Fig. 5). Four experiments were conducted in the exploitation mode, with an average OF around 0.88 ± 0.17 ; the highest OF (1.12) was achieved in the tenth experiment.

There is a similar balance between exploration and exploitation that occurs during the learning and utilization of the correct growth vs. no-growth classifier. We note here that the decision-making policy is attempting to solely optimize the OF. In particular, we make no attempt in selecting experiments that balance between this optimization and the learning of an optimal classifier in the two-dimensional feature space. There are various heuristics and approaches in attempting to augment the decision-making procedure to accomplish this auxiliary learning, including performing Bayesian optimization with unknown constraints⁶¹ and a more generic framework called the Mean Objective Cost of Uncertainty⁶². In the present setting, however, the balance is achieved in an emergent, albeit not necessarily optimal, manner. Specifically, the choice of experiments in the growth region leads to potential class imbalance favoring the “growth” classes. This imbalance in turn leads to an overestimation of the growth region, which may lead to the selection of no-growth experiments, mitigating such an imbalance. The optimization of this balance could potentially increase the effectiveness of this technique. More broadly, the use of the SVM classifier in this data-limited setting to specify feasible regions over which we make decisions under uncertainty represents a mixture of Bayesian and frequentist methods. This mixture merits further discussion, which we have presented in Supplementary Section 3.

Simulations

To further explore this system, we performed statistical simulations of subsequent experimental campaigns. Taking existing data obtained from the growth experiments on the 1 nm Co catalyst, we fitted a Gaussian process (GP) posterior belief $f(\mathbf{x}) \sim \text{GP}(\mu(\mathbf{x}), \Sigma(\mathbf{x}, \mathbf{x}'))$ on the observed OF values from the data. From this belief, we sampled ground truths and by doing so we obtained a family of sampled OFs $f_i(\mathbf{x})$, which were consistent with existing experimental data. For each such sample of the ground truth OF, we simulated an experimental campaign in a manner identical to the physical campaigns outlined above. However, in place of running an actual, physical experiment on ARES for a selected set of experimental conditions \mathbf{x} , we instead sampled a

noisy observation $f_i(\mathbf{x}) + W$, where W is an additive noise term drawn from a normal distribution with mean 0 and variance σ_W^2 . Apart from using a statistically simulated experimental observation, such simulated campaigns proceeded identically to their physical analog on the ARES experiments, including the belief modeling, decision-making policies, and the use of classification models to define nonlinear feasible sets over which to optimize acquisition functions of the decision-making policies in a constrained manner. In particular, after n steps of the simulated campaign, we could predict the optimal synthesis conditions $\mathbf{x}^{*,n}$ based on the n simulated experimental data points available at the time. From this, we calculated the sampled ground truth OF value at this predicted optimum, $y^{*,n} = f_i(\mathbf{x}^{*,n})$. As a measure of improvement, we then compared this best value to the best observed value from the existing experimental data set, $y^{*,\text{data}}$. For every simulation we ran, we calculated this difference to report a percent increase statistic:

$$D(n) = 100 \times (y^{*,n} - y^{*,\text{data}}) / y^{*,\text{data}} \quad (4)$$

By running several simulations, we obtained distributions of this statistics, allowing us to make inferences about average- and extreme-case behavior to be expected from further simulations.

Such simulations allow us to therefore study the impact of various modeling or policy choices on the effectiveness of a closed-loop campaign. As a first study, we ran simulated campaigns to consider the effectiveness of the EI policy compared to a Pure Exploration policy that simply selects a feasible synthesis condition uniformly at random, and the Maximum Variance policy, which selects the experiment whose predicted OF has highest uncertainty, which is often used for Active Learning campaigns. We also considered the Upper Confidence Bound (UCB) policy, which selects experiments that maximize an optimistic estimate of the OF response given by an upper confidence bound⁶³. For all policies, we ran 100 simulated campaigns each, sampling a different ground truth OF value for each campaign, followed by calculation of the $D(n)$ statistic. Figure 6a shows the median values of $D(n)$ vs. n for both the EI and Exploration policies, calculated over the 100 simulated campaigns. In addition, shaded regions indicating the 25th and 75th percentiles are shown. The figure indicates that the EI policy will, on average, result in a 13% increase of the best observed value over 20 experiments, with a spread of between 2% and 30% increase observed over the simulations. In contrast, pure exploration policies do not improve over an initial 3–4% increase, with a spread between 12% increase and a 6% decrease of the predicted best value observed over the simulations. We observe a similar performance from the maximum variance policy. Here, such a decrease would suggest campaigns in which, due to noisy observations of the sampled ground truth, an inaccurate belief of the ground truth results in a predicted best synthesis condition $\mathbf{x}^{*,n}$ whose corresponding ground truth value $y^{*,n} = f_i(\mathbf{x}^{*,n})$ actually yields poor performance. From these simulations, we see that the EI policy does outperform pure exploration, i.e., random sampling. The UCB policy performs the best overall. Interestingly, the UCB policy is a policy typically used for the multi-armed bandit problem, within which we wish to optimize the cumulative rewards (here, the sum of observed OF values) obtained throughout a campaign. Due to this, the policy, although attempting to balance between exploration and exploitation, more heavily favors the former, suggesting enough experiments are present for the campaign to transition more heavily to exploitation, even more than the natural balance present in the EI algorithm.

For another simulation study, we considered the impact of constraining experimental suggestions to a predicted feasible set on the campaign performance. Specifically, recall that an SVM classifier was trained on experimental data that was projected to the two-dimensional space of temperature and the ratio between oxidizers and reducers. These two quantities were used as features

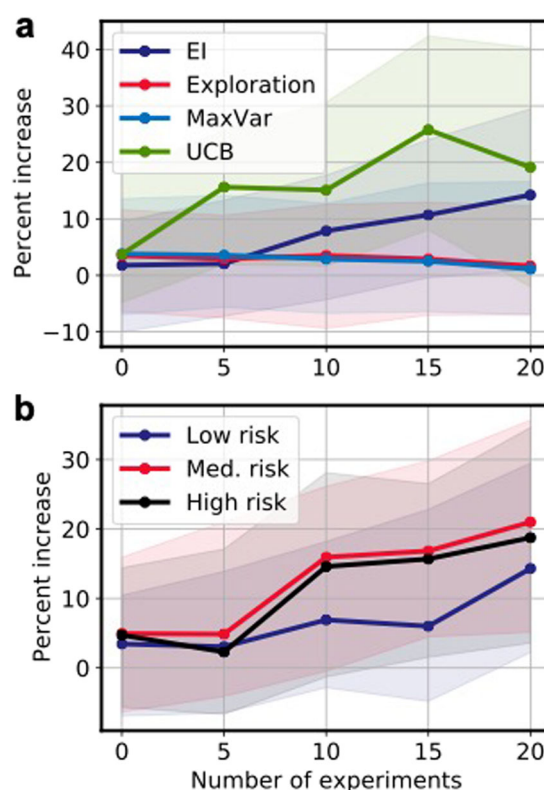


Fig. 6 Results from simulations of experimental campaigns. **a** Median and the interquartile range for the Percent Increase statistic for an EI-driven and Exploration-driven campaign. The EI campaign resulted in a median increase of around 12%, whereas the Exploration campaign resulted in no median increase. **b** Percent Increase for different risk levels, including risk-averse (threshold = 0.1) and more risk-tolerant (threshold = −0.1 and unconstrained) campaigns. In this context, risk-tolerant campaigns performed better than risk-averse ones.

of the classifier. Class labels were obtained by thresholding the observed OF value here using Z-normalized data. Such normalized data points where this value $y > y_{\text{thresh}}$ surpassed a predefined threshold were labeled with a “growth” class, whereas those that did not surpass this threshold were labeled with a “no-growth” class. That is, y_{thresh} is analogous to f_{\min} defined above, but applied to normalized data. Larger values of y_{thresh} implied more constraints in selecting experimental actions, meaning selected experimental actions within a more constrained feasible set were predicted to result in more substantial growth. Conversely, smaller values implied less constraint and indicated a willingness to select experimental conditions that may result in low growth. Thus, a selection of y_{thresh} can be considered as a specification of a risk tolerance level: larger values imply a risk-averse campaign, whereas smaller values imply a risk-tolerant one.

We ran simulation studies to consider the impact of such a threshold, considering threshold values of 0.1, −0.1, and $-\infty$, which corresponds to risk-averse, risk-tolerant, and unconstrained selection of synthesis conditions. For each threshold value, we ran 100 simulations and calculated the performance metric $D(n)$, which is shown in Fig. 6b. We observe that for the risk-averse ($y_{\text{thresh}} = 0.1$) setting, campaigns obtained a 14% improvement over the existing data set over 20 experiments, on average. However, we observe that both the risk-tolerant and unconstrained settings perform better, obtaining an increase of around 19% and 20% improvement. For future experimental runs, the simulations suggest an increased risk tolerance would therefore be beneficial. These simulations offer a manner to calibrate the risk tolerance of a campaign when

determining experimental feasibility, along with other hyperparameters. Moreover, such a parameter could be dynamically tuned along with the changing beliefs of the response function itself. An implementation of such a dynamically tuned (via simulations) risk tolerance factor could improve the performance of closed-loop campaigns in which such experimental feasibility plays a significant role.

DISCUSSION

We used an ML planner based on the EI policy to map the growth parameter phase spaces for maximizing the intensities of two RBMs at 225 and 265 cm^{-1} , which correspond to SWCNT diameters around 1.06 and 0.91 nm, respectively. As optimization parameters, we considered the growth temperature and the partial pressures of reductants (C_2H_4 , C_2H_2 , and H_2) and oxidants (water vapor and CO_2). By considering the ratio of pressures of reductants to oxidants, the reduced dimensionality allowed us to construct growth condition maps from our results, which revealed significantly different optimal conditions for maximizing the intensity of RBMs around 265 vs. 225 cm^{-1} . These differences in growth conditions highlight the effectiveness of our planner, which achieved results in less than a hundred experiments. Our EI decision policy also enabled us to assess the effect of experimental risk on the rate of convergence and found that it had a significant impact. In the future, this approach could be extended to include thermodynamic models of catalyst reduction and rate equations of catalytic feedstock dissociation. Although the best diameter selectivity was around 35%, our methodology paves the way forward for achieving the ultimate goal, namely diameter- and chirality-controlled SWCNT growth.

METHODS

CNT growth in ARES

CNT growth was performed in our custom-built in situ system named ARES. In ARES, a 6 W 532 nm laser (Verdi) serves as both the heat source and Raman excitation source, and is focused on a silicon substrate consisting of patterned micropillars on an SiO_2 underlayer (10 μm in diameter and height, fabricated by reactive ion etching). The substrates are loaded into a miniature high-vacuum chamber with an optical window whose environment can be controlled through automated pressure and gas mass flow controllers. Heating of the thermally isolated micropillars is achieved by varying the laser power, allowing rapid increases (within microseconds) in temperatures up to 1200 $^\circ\text{C}$. The scattered light from the micropillars is coupled to a spectrometer through focusing optics and a notch filter, enabling in situ Raman measurements. The micropillar temperature is estimated from the redshifted Raman peak frequency of the silicon micropillar. For this study, we first deposited a 10 nm alumina barrier layer onto the silicon micropillars by atomic layer deposition. Subsequently, 1 nm Co films were sputtered on the substrates using magnetron sputtering. The micropillar substrate was loaded into the growth chamber and evacuated to a base pressure of 10^{-6} Torr, followed by backfilling with the growth gases. As described in the main text, we used the ratio of partial pressures of reductants (ethylene, acetylene, and hydrogen) and oxidants (water vapor and carbon dioxide), and labeled the pressure ratio, as chosen by the planner. Before and after each growth experiment, Raman spectra (pre- and post scans) were collected at room temperature using a low laser power (0.2 mW) and a 30 s acquisition time. During the course of each experiment, spectra were also collected in real time (every 3 s), and the RBM peak fitting was performed on the post scans following each experiment. These RBM frequencies were immediately fed into the planner to compute the parameters (i.e., the temperature and pressure ratios) for the next growth experiment.

DATA AVAILABILITY

The data generated or analyzed during the current study are available from the corresponding authors on reasonable request.

Received: 1 May 2021; Accepted: 8 September 2021;

Published online: 01 October 2021

REFERENCES

- Cavin, R. K., Lugli, P. & Zhirnov, V. V. Science and engineering beyond Moore's law. *Proc. IEEE* **100**, 1720–1749 (2012).
- Rao, R. et al. Carbon nanotubes and related nanomaterials: critical advances and challenges for synthesis toward mainstream commercial applications. *ACS Nano* **12**, 11756–11784 (2018).
- Saito, R., Fujita, M., Dresselhaus, G. & Dresselhaus, M. S. Electronic structure of chiral graphene tubules. *Appl. Phys. Lett.* **60**, 2204–2206 (1992).
- Li, J. et al. Importance of diameter control on selective synthesis of semi-conducting single-walled carbon nanotubes. *ACS Nano* **8**, 8564–8572 (2014).
- Tunuguntla, R. H. et al. Enhanced water permeability and tunable ion selectivity in subnanometer carbon nanotube porins. *Science* **357**, 792–796 (2017).
- Guo, S., Meshot, E. R., Kuykendall, T., Cabrini, S. & Fornasiero, F. Nanofluidic transport through isolated carbon nanotube channels: advances, controversies, and challenges. *Adv. Mater.* **27**, 5726–5737 (2015).
- Jung, H. Y., Jung, S. M., Kim, L. & Suh, J. S. A simple method to control the diameter of carbon nanotubes and the effect of the diameter in field emission. *Carbon* **46**, 969–973 (2008).
- Cao, Q. & Rogers, J. A. Ultrathin films of single-walled carbon nanotubes for electronics and sensors: a review of fundamental and applied aspects. *Adv. Mater.* **21**, 29–53 (2009).
- Nasibulin, A. G., Pikhitsa, P. V., Jiang, H. & Kauppinen, E. I. Correlation between catalyst particle and single-walled carbon nanotube diameters. *Carbon* **43**, 2251–2257 (2005).
- Fiawoo, M. F. C. et al. Evidence of correlation between catalyst particles and the single-wall carbon nanotube diameter: a first step towards chirality control. *Phys. Rev. Lett.* **108**, 195503 (2012).
- Han, S. et al. Diameter-controlled synthesis of discrete and uniform-sized single-walled carbon nanotubes using monodisperse iron oxide nanoparticles embedded in zirconia nanoparticle arrays as catalysts. *J. Phys. Chem. B* **108**, 8091–8095 (2004).
- Ishida, M., Hongo, H., Nihey, F. & Ochiai, Y. Diameter-controlled carbon nanotubes grown from lithographically defined nanoparticles. *Jpn. J. Appl. Phys.* **43**, L1356 (2004).
- Chiang, W.-H. & Sankaran, R. M. Linking catalyst composition to chirality distributions of as-grown single-walled carbon nanotubes by tuning NiFe1-x nanoparticles. *Nat. Mater.* **8**, 1–5 (2009).
- Song, W. et al. Synthesis of bandgap-controlled semiconducting single-walled carbon nanotubes. *ACS Nano* **4**, 1012–1018 (2010).
- Yamada, K., Kato, H. & Homma, Y. Narrow diameter distribution of horizontally aligned single-walled carbon nanotubes grown using size-controlled gold nanoparticles. *Jpn. J. Appl. Phys.* **52**, 5105 (2013).
- Bouanis, F. Z. et al. Diameter controlled growth of SWCNTs using Ru as catalyst precursors coupled with atomic hydrogen treatment. *Chem. Eng. J.* **332**, 92–101 (2018).
- Loebick, C. Z. et al. Selective synthesis of subnanometer diameter semiconducting single-walled carbon nanotubes. *J. Am. Chem. Soc.* **132**, 11125–11131 (2010).
- Carpena-Núñez, J. et al. Zeolite nanosheets stabilize catalyst particles to promote the growth of thermodynamically unfavorable, small-diameter carbon nanotubes. *Small* **16**, 2002120 (2020).
- Lu, C. & Liu, J. Controlling the diameter of carbon nanotubes in chemical vapor deposition method by carbon feeding. *J. Phys. Chem. B* **110**, 20254–20257 (2006).
- Saito, T. et al. Selective diameter control of single-walled carbon nanotubes in the gas-phase synthesis. *J. Nanosci. Nanotechnol.* **8**, 6153–6157 (2008).
- Picher, M., Anglaret, E., Arenal, R. & Jourdain, V. Processes controlling the diameter distribution of single-walled carbon nanotubes during catalytic chemical vapor deposition. *ACS Nano* **5**, 2118–2125 (2011).
- He, M., Jiang, H., Kauppinen, E. I. & Lehtonen, J. Diameter and chiral angle distribution dependencies on the carbon precursors in surface-growth single-walled carbon nanotubes. *Nanoscale* **4**, 7394–7398 (2012).
- Sakurai, S. et al. Role of subsurface diffusion and Ostwald ripening in catalyst formation for single-walled carbon nanotube forest growth. *J. Am. Chem. Soc.* **134**, 2148–2153 (2012).
- Kim, S. M. et al. Catalyst and catalyst support morphology evolution in single-walled carbon nanotube supergrowth: growth deceleration and termination. *J. Mater. Res.* **25**, 1875–1885 (2011).
- Kim, S. M. et al. Evolution in catalyst morphology leads to carbon nanotube growth termination. *J. Phys. Chem. Lett.* **1**, 918–922 (2010).

26. Sakurai, S. et al. A phenomenological model for selective growth of semi-conducting single-walled carbon nanotubes based on catalyst deactivation. *Nanoscale* **8**, 1015–1023 (2015).
27. Khalilov, U., Vets, C. & Neyts, E. C. Molecular evidence for feedstock-dependent nucleation mechanisms of CNTs. *Nanoscale Horiz.* **4**, 674–682 (2019).
28. Zhu, Z., Jiang, H., Susi, T., Nasibulin, A. G. & Kauppinen, E. I. The use of NH₃ to promote the production of large-diameter single-walled carbon nanotubes with a narrow (n,m) distribution. *J. Am. Chem. Soc.* **133**, 1224–1227 (2011).
29. Thurakitseree, T. et al. Diameter controlled chemical vapor deposition synthesis of single-walled carbon nanotubes. *J. Nanosci. Nanotechnol.* **12**, 370–376 (2012).
30. Liao, Y. et al. Tuning geometry of SWCNTs by CO₂ in floating catalyst CVD for high-performance transparent conductive films. *Adv. Mater. Interfaces* **5**, 1801209 (2018).
31. Farhat, S. et al. Diameter control of single-walled carbon nanotubes using argon–helium mixture gases. *J. Chem. Phys.* **115**, 6752–6759 (2001).
32. Mueller, T., Kusne, A. G. & Ramprasad, R. Machine learning in materials science: Recent progress and emerging applications. *Rev. Comput. Chem.* **29**, 186–273 (2016).
33. Reyes, K. G. & Maruyama, B. The machine learning revolution in materials? *MRS Bull.* **44**, 530–537 (2019).
34. Sun, W. et al. Machine learning–assisted molecular design and efficiency prediction for high-performance organic photovoltaic materials. *Sci. Adv.* **5**, eaay4275 (2019).
35. Gongora, A. E. et al. A Bayesian experimental autonomous researcher for mechanical design. *Sci. Adv.* **6**, eaaz1708 (2020).
36. Huo, H. et al. Semi-supervised machine-learning classification of materials synthesis procedures. *npj Comput. Mater.* **5**, 1–7 (2019).
37. Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O. & Walsh, A. Machine learning for molecular and materials science. *Nature* **559**, 547–555 (2018).
38. Epps, R. W. et al. Artificial chemist: an autonomous quantum dot synthesis bot. *Adv. Mater.* **32**, 2001626 (2020).
39. Nikolaev, P. N. et al. Autonomy in materials research: a case study in carbon nanotube growth. *NPJ Comput. Mater.* **2**, 16031 (2016).
40. Chang, J. et al. Efficient closed-loop maximization of carbon nanotube growth rate using Bayesian optimization. *Sci. Rep.* **10**, 9040 (2020).
41. Rao, R., Liptak, D., Cherukuri, T., Yakobson, B. I. & Maruyama, B. In situ evidence for chirality-dependent growth rates of individual carbon nanotubes. *Nat. Mater.* **11**, 1–4 (2012).
42. Rao, R., Islam, A. E., Pierce, N., Nikolaev, P. & Maruyama, B. Chiral angle-dependent defect evolution in CVD-grown single-walled carbon nanotubes. *Carbon* **95**, 287–291 (2015).
43. Rao, R. et al. Revealing the impact of catalyst phase transition on carbon nanotube growth by in situ Raman spectroscopy. *ACS Nano* **7**, 1100–1107 (2013).
44. Nikolaev, P., Hooper, D., Perea-López, N., Terrones, M. & Maruyama, B. Discovery of wall-selective carbon nanotube growth conditions via automated experimentation. *ACS Nano* **8**, 10214–10222 (2014).
45. Rao, R. et al. Maximization of carbon nanotube yield by solid carbon-assisted dewetting of iron catalyst films. *Carbon* **165**, 251–258 (2020).
46. Everhart, B. M. et al. Efficient growth of carbon nanotube carpets enabled by in situ generation of water. *Ind. Eng. Chem. Res.* **59**, 9095–9104 (2020).
47. Kluender, E. J. et al. Catalyst discovery through megalibraries of nanomaterials. *Proc. Nat. Acad. Sci. USA* **116**, 40–45 (2019).
48. Islam, A. E. et al. Photothermal oxidation of single layer graphene. *RSC Adv.* **6**, 42545–42553 (2016).
49. Rao, R., Islam, A. E., Campbell, P. M., Vogel, E. M. & Maruyama, B. In situ thermal oxidation kinetics in few layer MoS₂. *2D Mater.* **4**, 025058 (2017).
50. Share, K. et al. Nanoscale silicon as a catalyst for graphene growth: mechanistic insight from in situ Raman spectroscopy. *J. Phys. Chem. C* **120**, 14180–14186 (2016).
51. Vilá, R. A. et al. In situ crystallization kinetics of two-dimensional MoS₂. *2D Mater.* **5**, 011009 (2017).
52. Carpena-Núñez, J. et al. Isolating the roles of hydrogen exposure and trace carbon contamination on the formation of active catalyst populations for carbon nanotube growth. *ACS Nano*, <https://doi.org/10.1021/acsnano.9b01382> (2019).
53. Kataura, H. et al. Optical properties of single-wall carbon nanotubes. *Synth. Met.* **103**, 2555 (1999).
54. Zhang, D. et al. (n, m) Assignments and quantification for single-walled carbon nanotubes on SiO₂/Si substrates by resonant Raman spectroscopy. *Nanoscale* **7**, 10719–10727 (2015).
55. Saito, R., Dresselhaus, G. & Dresselhaus, M. S. *Physical Properties of Carbon Nanotubes*. Vol. 4 (World Scientific, 1998).
56. Park, C. et al. Sequential adaptive design for jump regression estimation. Preprint at <https://arxiv.org/abs/1904.01648> (2020).
57. Rasmussen, C. E. & Williams, C. K. I. *Gaussian Processes for Machine Learning* (Univ. Press Group Limited, 2006).
58. Jones, D. R., Schonlau, M. & Welch, W. J. Efficient global optimization of expensive black-box functions. *J. Glob. Optim.* **13**, 455–492 (1998).
59. Tian, Y., Jiang, H., Laiho, P. & Kauppinen, E. I. Validity of measuring metallic and semiconducting single-walled carbon nanotube fractions by quantitative Raman spectroscopy. *Anal. Chem.* **90**, 2517–2525 (2018).
60. Zhang, X. et al. High-precision solid catalysts for investigation of carbon nanotube synthesis and structure. *Sci. Adv.* **6**, eabb6010 (2020).
61. Gelbart, M. A., Snoek, J. & Adams, R. P. Bayesian optimization with unknown constraints. Preprint at <https://arxiv.org/abs/1403.5607> (2014).
62. Yoon, B.-J., Qian, X. & Dougherty, E. R. Quantifying the objective cost of uncertainty in complex dynamical systems. *IEEE Trans. Signal Process.* **61**, 2256–2266 (2013).
63. Lai, T. L. & Robbins, H. Asymptotically efficient adaptive allocation rules. *Adv. Appl. Math.* **6**, 4–22 (1985).

ACKNOWLEDGEMENTS

We acknowledge funding from the Air Force Office of Scientific Research (LRIR 16XCOR322).

AUTHOR CONTRIBUTIONS

R.R., P.N., J.C.N., K.G.R., and B.M. conceived the study and devised the objective function. R.R., J.C.N., M.A.S., and K.G.R. performed the experiments. R.R. and K.G.R. wrote the manuscript with edits and revisions from all authors.

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/s41524-021-00629-y>.

Correspondence and requests for materials should be addressed to Rahul Rao or Kristofer G. Reyes.

Reprints and permission 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 4.0 International License, which permits use, sharing, adaptation, 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 license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license 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 license, visit <http://creativecommons.org/licenses/by/4.0/>.

This is a U.S. government work and not under copyright protection in the U.S.; foreign copyright protection may apply 2021