
Data-Driven Density Functional Theory: A case for Physics Informed Learning

Petr Yatsyshin
The Alan Turing Institute
London
pyatsyshin@turing.ac.uk

Serafim Kalliadasis
Department of Chemical Engineering
Imperial College London
s.kalliadasis@imperial.ac.uk

Andrew B. Duncan
The Alan Turing Institute
Imperial College London
London
a.duncan@imperial.ac.uk

Abstract

We propose a novel data-driven approach to solving a classical statistical mechanics problem: given data on collective motion of particles, characterise the set of free energies associated with the particle system. We demonstrate empirically that the particle data contains all the information necessary to infer the free energy of the underlying physical system. Unlike traditional physical modelling, which seeks to construct analytic approximations, our proposed approach leverages modern Bayesian computational capabilities to accomplish the same goal in a purely data-driven fashion. The Bayesian paradigm permits us to combine first principles with simulation data to obtain uncertainty-quantified predictions of the free energy, in the form of a probability distribution over the family of free energies, consistent with the observed particle data.

1 Introduction

Statistical mechanics relates macroscopic properties of matter, such as pressure, magnetisation or electric charge, to the underlying microscopic structure, given by the interactions between the constituent atomic particles. This is accomplished by casting the desired macroscopic quantities as averages over particle positions and velocities. Since in equilibrium particle velocities follow the Maxwell distribution, the central problem of equilibrium statistical mechanics is in representing the distribution over particle positions, $\rho(\mathbf{r})$. Due to correlations between the particles, this is a highly non-trivial task, even for systems with simple interparticle interactions. Statistical mechanics can be formulated as a so-called Density Functional Theory (DFT), where the free energy of the system, $F[\rho]$, is expressed as functional of $\rho(\mathbf{r})$. The equilibrium distribution $\rho(\mathbf{r})$ is the one that minimizes the free energy $F[\rho]$, subject to suitable constraints [1]. In this work we deal with the inverse problem: using observations of the instantaneous coordinates of the system's particles, find an *unknown* $F[\rho]$.

Modern DFT has a versatile plethora of computationally efficient approximate methods. In quantum DFT, these encompass effective approximations for the exchange-correlation energy of the electronic density, and in classical DFT, these constitute a library of methods to approximate the excess-over-ideal Helmholtz free energy $F[\rho]$, to account for interactions in classical many-body systems. In the present work we focus on classical DFT and develop a data-driven Bayesian approach to the derivation of $F[\rho]$ of systems with excluded volume interactions, i.e. system of hard particles which interact via elastic collisions. For such systems, exact $F[\rho]$ is known only for one-dimensional (1D)

fluid of hard rods (HR), constrained to a line [2]. Approximate DFTs are known only for a handful of relatively simple systems, such as hard spheres or parallel hard cubes [3]. In general, constructing DFT approximations from first-principles is a highly non-trivial task, and there is high need for an algorithmic hands-off method for obtaining approximate DFT functionals. Since our goal here is methodological, we work with the simplest possible system: 1D HR on a line. This system is chosen for two reasons. Firstly, access to exact $\rho(\mathbf{r})$ allows us to easily benchmark the inference scheme against the ground truth. Secondly, data generation is easy in 1D, which facilitates convergence studies and comparison to brute-force statistical inference. Application of our principled approach to more complex system is straightforward, but would require more intense computational effort.

Separating the Maxwellian part, we can express $F[\rho] = \beta^{-1} \int \rho(\mathbf{r}) (\ln \lambda^3 \rho(\mathbf{r}) - 1) d\mathbf{r} + F_{ex}[\rho]$, where β is the inverse of temperature and λ is the thermal wavelength. The excess-over-ideal part $F_{ex}[\rho]$ is what needs to be inferred from the data. The most successful existing approximations for $F_{ex}[\rho]$ of hard-particle fluids are based on a physical intuition that $F_{ex}[\rho]$ in a unit volume is a function of the geometric measures of the particles, expressed as partial densities $\{n_i\}$:

$$\beta F_{ex}[\rho] = \int \Phi(n_1(\mathbf{r}), n_2(\mathbf{r}), \dots) d\mathbf{r}, \quad n_i(\mathbf{r}) = \omega_i * \rho \equiv \int \omega_i(\mathbf{r} + \mathbf{t}) \rho(\mathbf{t}) dt \quad (1)$$

where ω_i are window-functions, which allow us to compute particle volume, surface area, curvatures and so on by taking convolutions as above [3]. In our case, rods of diameter $2R$ are moving on a line, interacting via elastic collisions [$\mathbf{r} \equiv x$, see Fig. 1(a)]. These particles are described by volume (i.e., length) and surface area (i.e., the positions of HR ends). The respective window-functions are $\omega_v(x) = \Theta(R - |x|)$ and $\omega_s(x) = \delta(R - |x|)/2$, giving rise to two partial densities, $\eta(x) = \int_{x-R}^{x+R} \rho(t) dt$ and $n_0(x) = (\rho(x - R) + \rho(x + R)) / 2$.

A one-dimensional system of hard rods inside a pore of width L is sketched at the top of Fig. 1(a) for $2R = 1$. In a so-called grand-canonical ensemble of statistical mechanics, the HR system is considered open and connected to a reservoir or particles, held at chemical potential μ . This means that there is a μ -dependent probability for insertion or deletion of a particle in the pore. In the present paper we work in a grand-canonical ensemble, because the exact free energy functional of HRs is known for this case and can be used for benchmarking. On the other hand, the collective behaviour of HRs can be simulated in the grand-canonical ensemble using standard Monte-Carlo techniques [4]. This produces the training dataset \mathcal{D}_μ , which consists of the expected number of particles $\langle N_\mu \rangle$ in the pore and a set of M instantaneous particle positions $\{y_i\}_{i=1}^M$:

$$\mathcal{D}_\mu = (\mu, \{y_i\}_{i=1}^M, \langle N_\mu \rangle). \quad (2)$$

When M is sufficiently large, the histogram of $\{y_i\}_{i=1}^M$ should approximate the DFT density profile $\rho(x)$, which minimizes the DFT free energy, and thus satisfies the following Euler-Lagrange equation:

$$\ln \rho(x) + \left(\omega_v * \frac{\partial \Phi}{\partial \eta(x)} + \omega_s * \frac{\partial \Phi}{\partial n_0(x)} \right) - \beta \mu = 0, \quad \text{subject to } \int_{-L/2}^{L/2} \rho(x) dx = \langle N_\mu \rangle. \quad (3)$$

We attempt the inference of the term $\Phi(\eta, n_0)$ above, in terms of a polynomial parametrisation:

$$\Phi(n_0, \eta) \equiv \Phi(n_0, \eta | Q) = (a_{N_1} n_0(x)^{N_1} + a_{N_1-1} n_0(x)^{N_1-1} + a_0) (b_{N_2} \eta(x)^{N_2} + \dots b_0), \quad (4)$$

where the model parameter set $Q = (a_{N_1}, \dots, a_0, b_{N_2}, \dots, b_0)^T$ has $N_Q = N_1 + N_2 + 2$ elements. The exact grand-canonical HR functional is given by $\Phi_X(n_0, \eta) = -n_0 \log(1 - \eta)$ in Eq. (1) [2]. Thus, the ground-truth $\rho(x) \equiv \rho_{X|\mu}(x)$ solves (3) with $\Phi \equiv \Phi_X(n_0, \eta)$. Notice that since $\Phi(n_0, \eta)$ is not unique [any function which integrates to zero can be added to it in Eq. (1)], we do not expect to infer $\Phi_X(n_0, \eta)$ precisely. We will characterise the posterior spread by bursts of samples of $\rho(x | Q)$, and the point-predictions – by maximum a-posteriori estimator (MAP), $\rho_{\text{MAP}}(x) = \rho(x | \text{argmax}_Q P(Q))$. During training, we solve Eq. (3) numerically for a given Q , which is done using a Newton scheme. The most numerically stable way to discretize Eq. (3) for our purposes is with a simple linear interpolant, using trapezium rule for the quadratures.

2 Bayesian inference procedure

Clearly, not every Q yields an admissible free energy functional. Hence, we choose a Gaussian prior $\mathcal{N}(Q | \bar{Q}, \Sigma_Q)$ with mean $\bar{Q} = 0$ and a diagonal covariance matrix Σ_Q . With this prior on Q ,

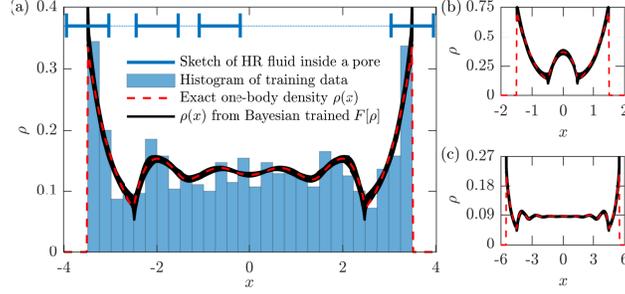


Figure 1: Illustration of the trained HR functional. Simulated HRs have width $2R = 1$ and interact via elastic collisions. They are confined to a pore of width $L = 8$ and are held at chemical potential $\mu = 2$, so that $\langle N_\mu \rangle = 4.6$. The histogram in (a) shows the training dataset from Eq. (2) with $M = 1000$ simulated HR coordinates, used to train $F[\rho]$ with $N_1 = N_2 = 5$ in Eq. (4). The black “curve” shows 200 profiles $\rho(x)$, obtained from samples of the trained DFT functional. The spread of these profiles characterises the uncertainty of the Bayesian scheme about $F[\rho]$. (b) and (c) show 200 samples in pores with $L = 4$ and $L = 12$, obtained from the same functional as (a). In (a)-(c), the dashed red curve shows the ground truth in terms of exact distribution $\rho_{X|\mu}(x)$.

$\Phi(\eta, n_0)$ is a Gaussian random field with polynomial features in η and n_0 . The prior variances on Q are chosen to constrain the components of Q to be sufficiently close to zero. The expression for Bayesian likelihood follows from the physical interpretation of $\rho(x | Q)$ as the probability density function: $P(\mathcal{D}_\mu | Q) = \prod_{i=1}^M \rho(y_i | Q)$. Posterior distribution over Q follows from the Bayes rule, $P(Q | \mathcal{D}_\mu) \propto \mathcal{N}(Q | \bar{Q}, \Sigma_Q) P(\mathcal{D}_\mu | Q)$, and yields predictive posterior distributions over $F[\rho | Q]$ and $\rho(x | Q)$. Since $P(Q | \mathcal{D}_\mu)$ is not analytically tractable and known up to a normalising constant, approximate methods of inference must be considered.

We implement a Hamiltonian Monte-Carlo (HMC) algorithm to generate samples from the target distribution [5]. This algorithm requires computing the gradient of the log-posterior $\nabla_Q \log P(Q)$ at every iteration of the chain, and therefore the gradient of log-likelihood $P(\mathcal{D}_\mu | Q)$. The latter necessitates computing $\nabla_Q \rho$, the Jacobian of the solution of Eq. (3) with respect to the parameter Q . Computing this gradient directly would require solving Eq. (3) $N_Q + 1$ times to evaluate $\nabla_Q \log P(Q)$ rendering it computationally infeasible even for moderate N_Q . Significant improvement can be achieved by using adjoint differentiation methods, which allow us to relate the numerical solution of Eq. (3) with its Jacobian with respect to Q via a simple linear system.

After tuning the HMC algorithm step-size and burn-in parameters to ensure that the output is stationary and sufficiently fast mixing, we generate samples from the HMC using a sufficiently long run of 4 independent chains. The empirical predictive distribution for $F_{ex}[\rho]$ is then obtained from these samples analytically via Eqs. (4) and (1) and gives the distribution over free energies, consistent with the observed simulation data. We illustrate this in Fig. 1. Fig. 1 (a) shows the histogram of the training data, the minimising densities $\rho(x)$ of the trained $F[\rho]$ (black), and the truth, given by the exactly known distribution $\rho_{X|\mu}(x)$ (red). Noticeably, we are able to approximate $\rho_{X|\mu}(x)$ with relatively few particle trajectories, attesting to the method’s ability to combine essential physical assumptions with simulation data. Figs. 1 (b) and (c) illustrate the generalisation of the trained functional with the pore width L . The functional trained at $L = 8$ is used for computations at $L = 2$ and $L = 12$ (at the same $\mu = 2$), with superimposed $\rho_{X|\mu}(x)$. The spread of the prediction curves $\rho(x)$ indicates the standard deviation and captures local uncertainty. This seems largest around turning points of the profiles. The uncertainty can be reduced by increasing the size of the training dataset. For the wide pore in Fig. 1(c) the effects of the side walls are lost in the pore center, and the fluid near the pore center should behave like bulk fluid. The fact that the correct plateau of $\rho_{X|\mu}(x)$ is reproduced by the trained $F[\rho]$ means that the trained functional correctly captures the physics of the bulk fluid and its thermodynamic equation of state. This result is physically remarkable considering the fact that we were training on the data of a highly confined fluid, represented by the histogram in Fig. 1(a).

We cannot expect the model used of Fig. 1 to generalise with μ . After all, the training dataset currently includes only a single μ -point. To achieve generalisation with μ we must extend the learning in two ways: (i) provide training data for multiple values of μ and (ii) extend the model to accommodate

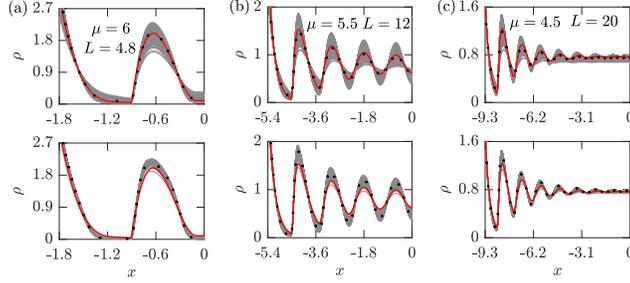


Figure 2: Generalisation of the trained functionals with chemical potential μ , as expressed by Eqs. (5)–(6). Two functionals, both with $N_1 = 3$ and $N_2 = 8$ (but different M) are trained at $K = 8$ integer values of $\mu = -2, \dots, 5$, using $M_n = 10^4$ particle coordinates per μ -point. Top and bottom panels in (a)–(c) correspond to $M = 0$ (no μ dependence) and $M = 1$ (linear μ dependence). Depicted density profiles minimise DFT at the specified μ and L , and due to symmetry are shown for $-L/2 \leq x \leq 0$. The posterior spread is illustrated by 400 samples (grey). The MAP estimators (black dots) lie close to the ground truth, $\rho_{X|\mu}(x)$ (red). Notice that the linear in μ model has less uncertainty, but the μ -independent model has better MAP estimators, particularly in (a) and (b), where μ extrapolates from the training set. This exemplifies over-fitting.

μ -dependence. Thus, we generalise $\Phi(n_0, \eta | Q)$ to $\Phi(n_0, \eta | Q(\mu))$ by representing each element of Q as a polynomial of degree M . The new parameter set is represented by the $N_Q \times (M + 1)$ matrix A of polynomial coefficients:

$$\Phi(n_0, \eta | \mu, \alpha) \equiv \Phi(n_0, \eta | Q(\mu | \alpha)), \quad Q(\mu | \alpha) = A (\mu^M, \mu^{M-1} \dots 1)^T, \quad (5)$$

where $\alpha = (\alpha_1, \dots, \alpha_{N_\alpha})^T$, $N_\alpha = N_Q(M + 1)$, is the (row-wise) flattened matrix A . The training data set \mathcal{D} and the likelihood function for this extended model become:

$$\mathcal{D} = \{\mathcal{D}_{\mu_n}\}_{n=1}^K \equiv \left\{ \left(\mu_n, \{y_{i|n}\}_{i=1}^{M_n}, \langle N_n \rangle \right) \right\}_{n=1}^K, \quad P(\mathcal{D} | \alpha) = \prod_{n=1}^K \prod_{i=1}^{M_n} \rho(y_{i|n} | \mu_n, \alpha), \quad (6)$$

where $M_n \equiv M_{\mu_n}$, $y_{i|n}$ is the i -th simulated particle coordinate at μ_n , $\langle N_n \rangle \equiv \langle N_{\mu_n} \rangle$, and $\rho(x | \mu_n, \alpha)$ is the solution of Eq. (3) at $\mu = \mu_n$ and Φ given by Eqs. (4) and (5). Now a joint Gaussian prior on the coefficients α induces a Gaussian random field prior on the space of functions of n_0, η and μ .

Eqs. (5)–(6) allow us to use the same posterior to describe the fluid on a set of μ , including dilute (small μ) and highly structured (large μ) fluid configurations. As an example, consider a training dataset $\{\mathcal{D}_{\mu_n}\}_{n=1}^K$ with $K = 8$ integer μ -points, $\mu = -2 \dots 5$, and $M_n = 10^4$ particle coordinates per μ -point in a pore of width $L = 8$. We train two functionals: a μ -independent one with $M = 0$ in Eq. (5), and a linear one in μ with $M = 1$. Both functionals have the same form of Eq. (4) with $N_1 = 3$ and $N_2 = 8$. The trained functionals are represented in Fig. 2 in terms of samples of minimising $\rho(x)$, for a variety of pore widths and chemical potentials outside of the training dataset. The top and bottom plots in (a)–(c) correspond to the μ -independent model and the linear model, respectively. Dotted curves show the MAP estimators. The uncertainty of the inferred $F[\rho]$ is illustrated by plotting 400 samples from the posterior (grey). The exact $\rho_{X|\mu}(x)$ is superimposed in red and demonstrates a good agreement of the trained functionals with the ground truth.

As before, the model generalises well with the pore width L . Notice that μ -dependent model shows significantly more confidence in its prediction than the μ -independent one, with predictive posterior distribution forming a narrower band around the MAP estimator on the top panels than on the bottom ones. A more subtle difference concerns the two possibilities for the test μ -points: either the chosen μ extrapolates from the training set [Figs. 2 (a) and (b)] or interpolates it [Fig. 2(c)]. We see that during extrapolation, μ -independent model gives slightly better MAP estimators than the linear model, in spite of the fact that the former has higher uncertainty. Moreover, when μ interpolates the training set, this difference between the MAP estimators of the two functionals nearly vanishes. We can attribute the higher certainty of the linear model to its higher flexibility in fitting the dependencies. At the same time, the slightly worse accuracy of the MAP estimator from the more complex linear

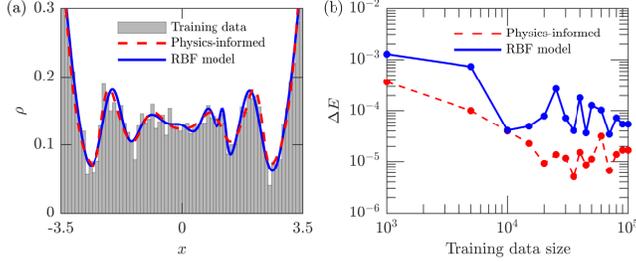


Figure 3: Comparison of physics-informed and black-box inference at fixed μ . (a) MAP of the physics-informed model from Eq. (4) with $N_1 = N_2 = 6$ (dashed red), and the RBF model from Eq. (7) with $N_f = 21$ (blue). Both models are trained at $\mu = 3$ and $L = 8$, on a fixed- μ dataset $\{\mathcal{D}_\mu\}$ of size $M = 5 \times 10^3$, represented by the histogram. (b) Energy distance to ground-truth [6].

model, observed during extrapolation suggests over-fitting. The actual $\Phi(n_0, \eta)$ is independent of μ , so artificially relaxing the μ -dependence may fit the training set with more certainty, but sacrifices generalisation.

3 Data efficiency

We assess the data efficiency of physics-informed inference approach by comparing it to a baseline black-box mixture model of Gaussian radial basis functions (RBF):

$$\rho(x | \mu) = \sum_{i=1}^{N_f} \alpha_i(\mu) \exp\left(-\frac{(x - p_i)^2}{w_i^2(\mu)}\right), \quad (7)$$

where $\alpha_i(\mu) \geq 0$ for all i and $\sum_j^{N_f} \alpha_j(\mu) = 1$. The Gaussian means p_i are fixed to be equispaced inside the computational domain to speed-up the training. For the fixed- μ model, we place a Gaussian prior on w_i and a rectified prior distribution with mean 1 and variance 0.1 on α_i , forcing α_i to be non-negative. Fig. 3 illustrates our comparison.

In Fig. 3(a) we superimpose two MAP estimators in the fixed- μ setting: the RBF distribution (blue) and the DFT minimiser of the MAP functional (red dashed). Both models were trained on the same small dataset, represented by the histogram. As we saw earlier with similar examples, the physics-informed model performs very well in low-data regimes. In fact, the physics-informed $\rho(x)$ visually coincides with the ground truth $\rho_{X|\mu}(x)$ everywhere, allowing us to omit the plot of $\rho_{X|\mu}(x)$. The quality of the RBF model is comparatively worse, as seen, e.g., by its lack of symmetry. There is simply not enough training data to produce an equally good representation of $\rho_{X|\mu}(x)$. Increasing the data size will improve the quality of the black-box model. We quantify this in Fig. 3(b) in terms of the energy distance [6], ΔE , between $\rho_{X|\mu}(x)$ and $\rho(x)$, as a function of the training data size. Observe that the physics-informed model remains at least an order of magnitude closer to the ground truth than the black-box model. Additionally, since important physical relationships are built into the physics-informed model, even when the inferred model becomes inaccurate, it still yields symmetric distributions which satisfy statistical-mechanical sum-rules [1]. For the same reason, the physics-informed model by construction generalises with L . We obviously cannot expect this from the black-box model. The dependence on L must be explicitly built into Eq. (7), and then even more data, spanning different L , will be needed for training. In the end, the cost of training a black-box model may be several orders of magnitude higher than training a physics-informed model.

4 Conclusion

We presented an inference methodology for finding the distributional model of the free energy of a many-body system using simulated particle trajectories. Considering the prototypical case of a one-dimensional fluid was a natural first step for physics-informed inference in classical many-body systems. The generalisation to higher dimensions is conceptually straightforward by considering functionals of the form $\Phi(\{n_i\})$. However, to describe phase transitions the parametric form of $\Phi(\{n_i\})$ should allow for singular behaviour.

Acknowledgments and Disclosure of Funding

PY was supported by Wave 1 of The UKRI Strategic Priorities Fund under the EPSRC Grant EP/T001569/1, particularly the “Digital Twins for Complex Engineering Systems” theme within that grant, and The Alan Turing Institute. ABD was supported by the Lloyds Register Foundation Programme on Data Centric Engineering and by The Alan Turing Institute under the EPSRC grant [EP/N510129/1].

References

- [1] J. F. Lutsko. Recent Developments in Classical Density Functional Theory. In *Adv. Chem. Phys.*, page 1. John Wiley & Sons, 2010.
- [2] J. K. Percus. Equilibrium State of a Classical Fluid of Hard Rods in an External Field. *J. Stat. Phys.*, 15:505, 1976.
- [3] P. Tarazona, J. A. Cuesta, and Y. Martinez-Raton. Density Functional Theories of Hard Particle Systems. In A. Mulero, editor, *Theory and Simulations of Hard-Sphere Fluids and Related Systems. Lecture Notes in Physics 753*, page 251. Springer, Berlin Heidelberg, 2008.
- [4] M. P. Allen and D. J. Tildesley. *Computer Simulations of Liquids*. Oxford Science Publications, 1989.
- [5] Steve Brooks. *Handbook of Markov chain Monte Carlo*. Taylor & Francis, Boca Raton, 2011. OCLC: 753969788.
- [6] Gábor J. Székely and Maria L. Rizzo. Energy statistics: A class of statistics based on distances. *Journal of Statistical Planning and Inference*, 143(8):1249–1272, August 2013.

Broader Impact

Classical density functional theory offers an incredibly versatile and user-friendly framework for modelling many-body systems. Yet, its potential for widespread adoption across a variety of soft-matter fields, including biology, nanofluidics and chemical engineering – to name but a few – is held back by the simple fact that accurate free energy functionals are known only for a handful of rather special systems. Present work is, to our knowledge, a first attempt to develop an algorithmic data-driven inference method for classical DFT functionals, equipped with full uncertainty quantification. Our approach yields a humanly interpretable analytic description of a many-body system with complex interactions in terms of a distribution over possible Helmholtz free energy functionals. The only required input is small amounts of particle simulation data. Using the Born-Oppenheimer principle, intermolecular effects happening on different spatiotemporal scales can be decoupled and accounted for separately by, e.g. coarse-graining complicated interactions and using perturbation expansions. This fact makes our work readily applicable to the broad range of problems dealing with atomistic and continuum mechanical descriptions of soft condensed matter. In a broader scope, our results can be extended to the development of an inferential modelling framework for many-body systems, where small-scale simulations are used to algorithmically capture essential patterns of collective behaviour. Thus, yielding an analytic description that can be scaled to system sizes beyond simulation capabilities. Broader societal consequences are associated with potential positive impact on nano-technology industries, such as chemical and bio-engineering, design and operation of miniature lab-on-a-chip chemical reactors and nanofluidic devices, DNA separation, vapor-liquid-solid growth of semiconductor nanowires, design of nano-patterned superhydrophobic surfaces and targeted drug delivery.