

Aspects of Classical Density Functional Theory[†]

JEROME K. PERCUS[‡]

*Courant Institute of Mathematical Sciences and Physics Department, New York University, 251 Mercer Street,
New York, New York 10012*

Received February 14, 1994

Introduction

Chemistry is a quantitative science. “Quantitative” means that the vast territory between yes and no is decomposed into identifiable, reproducible subterritories. “Science” implies a basic theoretical framework from which all conclusions flow in a strictly logical fashion. But neither definition carries with it a recognizable picture of the way that advances in chemistry actually occur. And this is the standard frustration of the theoretical chemist: an imaginative and experienced practitioner will reliably analyze a situation by what can only charitably be called fuzzy logic and end up with a decently reliable quantitative assessment. How does one use this as a clue to develop an incisive, accurate, quantitative structure without volumes of mathematical and computational analysis which effectively bury the essentials and often improve very little on the quantitative results? Perhaps a way to earn a living, but of dubious social utility.

Some Biased History

The topic that this Account is going to settle down into is that of density functional theory, a quantitative descriptive vehicle for many aspects of many-body theory. The associated viewpoint is one that developed slowly but by the end of the 1950s was a well-understood tool among the equilibrium classical fluid aficionados.^{1,2} The problems in its application to quantum ground states—in particular the kinetic energy density functional problem—and thermal equilibrium ensembles were discussed extensively by the same group, but did not reach the outside world until the elegant review of Hohenberg and Kohn.³ Even then, it took the technical tour de force of Kohn and Sham⁴ to carry applicability past the traditional Thomas–Fermi theory and its extensions. And the fact that an in-principle constructive definition was available⁵ took even longer to reach this community.

While the quantum density functional approach was exploding, classical fluid theorists were largely basking in the brilliance of increasingly sophisticated and successful integral equation methods⁶ in the bulk fluid context. As the efforts required to deal with the nonuniform real world become more onerous, density functional methods, heavily colored by empiricism and consequently simpler to apply, became more attractive. And the idea of incorporating coarse-grained or weighted densities, which had hitherto been present

in rudimentary form, except for a widely unread publication⁷ in which it was based upon an exact model solution, put the classical fluid theorists in the lead once more.⁸ My feeling is that the bases of this development, now translated to the quantum domain as well,⁹ are not as firm as they should be, and I would like to offer some comments aimed in this direction.

Integral Equations

The thermodynamics of fluids serves as basic input to much of chemistry. It encompasses not just the traditional relationships between intensive and extensive quantities—pressure, volume, temperature—but also the relationships that recognize the spatial non-uniformity that dominates the configurations of so many real systems, that due to containment and consequent short- and long-range order, to internal interfaces, to the effect of external fields, and so on. Theoretical studies have proceeded along two major directions. One of them, the integral equation method,⁶ is primarily a development of the last few decades. In essence, one realizes that, for a system with given external and pair interaction forces alone, it should be sufficient to focus on a typical pair of particles if the environment due to all of the others can be precisely simulated. The result of focusing is to obtain the mean density of the pair, $n_2(r, r')$, a function of six variables—two triplets of coordinates—for simple fluids. If this known, then the “profile equation”, the relation between density $n(r)$ and applied force field, is obtained by just balancing thermal, external, and internal forces on a unit volume of fluid (Figure 1)—the first of the famous YBG (Yvon–Born–Green) hierarchy of relations among multiparticle densities. The “closure problem” is that of expressing the pair density in terms of the observable particle density $n(r)$, and many suggestions have been made, based on various sequential approximations, resulting in routinely reliable results for simple fluids, and similarly for complex fluids, but here at a computational cost that can be daunting even with modern parallel computers.

Density Functionals

Another way at getting at the structure of a heterogeneous fluid is by the condition that the density

[†] Based on Hildebrand Award lecture, Denver, 1993.

[‡] Supported in part by a grant from the National Science Foundation.

(1) Yvon, J. *Nuovo Cimento, Suppl.* **1958**, *9*, 144.

(2) Lebowitz, J. L.; Percus, J. K. *J. Math. Phys.* **1963**, *4*, 116.

(3) Hohenberg, P.; Kohn, W. *Phys. Rev.* **1964**, *B136*, 864.

(4) Kohn, W.; Sham, L. J. *Phys. Rev.* **1965**, *A140*, 1133.

(5) Percus, J. K. *Int. J. Quantum Chem.* **1978**, *13*, 89.

(6) Hansen, J.-P.; McDonald, I. R. *Theory of Simple Liquids*; Academic Press: New York, 1986.

(7) Percus, J. K. *Chem. Phys.* **1981**, *75*, 1316.

(8) Curtin, W. A.; Ashcroft, N. W. *Phys. Rev.* **1985**, *A32*, 2909.

(9) See, e.g.: Moroni, S.; Senatore, G. *Phys. Rev.* **1991**, *B44*, 9864.

Jerome K. Percus was born in New York City on June 21, 1926. He received his B.S. in electrical engineering, M.A. in mathematics, and Ph.D. in physics from Columbia University. After teaching at Columbia University and research at Stevens Institute of Technology, he joined the Courant Institute of Mathematical Sciences of New York University, where he is now Professor, jointly with the Physics Department. His research interests have been centered around many-body theory, but have included econometrics, computer architecture, and a broad swath of mathematical biology.




Figure 1. Force balance on element of volume.

profile at a given external force field, temperature, and chemical potential adjusts itself to minimize the (Helmholtz) free energy of the system. In fact, since the variational procedure required to check this results precisely in an expression in which the external potential is given in terms of the density, this (inverse) profile equation can be "integrated back" to the free energy by turning up the density pattern, say from 0, in which everything is trivial, to the actual system profile. The result of course cannot depend on the intermediate details of turning up the density, in which case the profile equation is termed "integrable"; it had better be! It is the development of expressions for the free energy in terms of—as a *functional* of—the density that has dominated the current course of theoretical studies of nonuniform fluids. This functional can involve all sorts of nonlinear and nonlocal (e.g., integrals depending on the density at several points at once) combinations. *But* the result is a set of recipes that yield to relatively painless computations, perhaps without the typical high accuracy of the integral equation approaches. The recipes almost universally make use of reliable data on uniform fluids to fill in unknown parameters or even functions in model expressions for free energy obtained by experience, by mathematical artifice, and often by rather rapid hand-waving. Of course, intelligent interpolation has a habit of working out much better quantitatively than one could justify a priori, but it in some sense says that one knows the answer from the outset and, thus, exerts a bias against qualitative advances. I would like to describe here a growing current activity in which the class of recipes¹⁰ is severely restricted both by quite general requirements and by the necessity of reducing properly to exact expressions when these are available.

van der Waals Model

A discussion of free energy density functionals has no choice other than that of starting with van der Waals's elegant work of a century ago¹¹ (the great James Clerk Maxwell is said to have taken off a year to study the Dutch language in order to fully appreciate van der Waals's doctoral thesis!). From our current sophisticated viewpoint, it is simplicity itself: the interaction between particles in a simple fluid is divided into two parts (Figure 2): a short-range repulsive core ϕ_c so short in range that a fluid of cores alone can be regarded as locally uniform, with its free energy just the sum of all uniform local pieces. Then crucially, there is an attractive tail ϕ_T , long in range compared to the correlation range of fluctuations in the fluid, so that the pair density is just the product




Figure 2. van der Waals decomposition of interaction.




Figure 3. Liquid-gas interface profile.



Figure 4. Density profile near hard wall.

of the component densities, and the tail energy is computed in this fashion. Finally, the energy due to the external fields is added, and that is the end of the story. This may be simple, but it is remarkably effective. The density profiles resulting from it do a good job of representing two-phase interfaces, wetting layers, electrode-bounded electrolytes, and many other situations.

Defects

The van der Waals theory is hardly perfect, but is there anything fundamentally wrong with it? Well, yes: at long range, at short range, and in the middle. To start, imagine a planar liquid-gas interface. The van der Waals profile (Figure 3) is a not unreasonable rapid transition—a couple of interparticle spacings—from gas density to liquid density. But something is missing. Any child knows that the interface will be riddled with thermally excited surface waves (i.e., capillary waves), and many children would guess that the larger the area of interface, the "softer" it is, and the larger the amplitude of the waves, leading to a broadening of the profile in the long-exposure snapshot that constitutes the definition of an equilibrium profile. The interface softening due to these correlated long-range fluctuations¹² is absent in the van der Waals theory.

What is wrong at short range? Suppose that the fluid in question is bounded by an infinite hard wall, itself only a model of reality, of course. The infinite force felt by a particle at the wall effectively decouples it from the other particles, so that the fluid density must jump from 0 inside the wall to its ideal gas value P/RT , where P is the system pressure (Figure 4). This extremely rapid local change is simply not allowed in the van der Waals picture, in which everything tends to be averaged over the range of the tail potential. In fact, in the opposite case of a highly attractive wall, one would expect the particles to really pile up at the wall, creating an arbitrarily high density. But the van der Waals approximation never allows the density to be higher than its maximum for a uniform fluid, and,

(10) Evans, R. In *Inhomogeneous Fluids*; Henderson, D., Ed.; Dekker: New York, 1991; Chapter 5.

(11) van der Waals, J. D. Z. *Phys. Chem.* **1894**, *13*, 657.

(12) Buff, F. P.; Lovett, R. A.; Stillinger, F. H. *Phys. Rev. Lett.* **1965**, *15*, 621.




Figure 5. One-dimensional lattice gas.

for example, for really hard core interactions, this can never exceed the reciprocal volume of the core.

Perhaps the trouble is that the fluid that the van der Waals model is modeling is not quite the one we thought, but is a more complicated object, with 3, 4—any many—body forces. Then at least we would know precisely what approximation is being made. But no, we may in fact not be using any microscopic model whatsoever! This requires a more detailed analysis, simplest to carry out when the repulsive core is omitted from the discussion. It is then possible to work backwards and find the many-particle probability distribution that the model has to come from, only to discover that it typically contains regions of negative probability,¹³ a clear impossibility.

Lattice Gas

We had better build up our approximate models with greater care. A reliable way of doing so is to work with toy models or caricatures that can be solved exactly—so we know where we stand—and then extrapolate to reality. What is the simplest nontrivial fluid model that exists? It depends. If one really demands little balls and breakable springs, that is one thing, and capable of getting quite complex. But we will start instead with a much more rudimentary set of models, that of “lattice gases”. The prototypical lattice gas is simply a spatial grid, with each grid point or site being either occupied by a particle or not occupied, and a rule telling you how much you gain or lose in energy by occupying two adjacent sites. Because only one particle at a time can occupy a given site, there is the equivalent of a hard core built in, and only the strength of interaction remains as a parameter. Of course, the state of the system is also controlled by thermodynamic variables of chemical potential and temperature. Even this system can be solved exactly only in special cases (e.g., not on a full three-dimensional space), but the big complication is that, as in the van der Waals approach, we want to be able to accommodate an arbitrary external force field in order to be able to produce an arbitrary density pattern $\{n_x\}$ of mean number of particles at the sites $\{x\}$.

Let us crawl a bit before we start walking. Imagine a one-dimensional integer lattice in an external potential field $\{h_x\}$ (including chemical potential), and with given next neighbor interaction potential (Figure 5). Can we find its free energy? This is not trivial, but can be done,¹⁴ and it ends up with an *intrinsic* free energy (subtracting out the obvious $h_x n_x$ per site) that is simply a sum of pair contributions $f(n_x, n_{x+1})$ just like the microscopic energy, and very much in the van der Waals mold. The expression for $f(n_x, n_{x+1})$ is a little complicated, but may be found by solving the trivial two-site problem. But how do we tackle lattices that are not just one-dimensional? One intermediate step is to consider a Bethe lattice of coordination number q . Here each site is joined to q other sites ($q = 3$ shown in Figure 6), and the lattice is built up in




Figure 6. Unit of Bethe lattice.




Figure 7. Two-row triangular ladder.

this fashion. It looks locally like a regular lattice—hexagonal in this case—but only locally because it has no closed loops at all. Nonetheless, it is a start, and it *can* be solved.¹⁵ Doing so, one finds exactly the same kind of sum of pair contributions as in the one-dimensional case: better and better.

Entropy Functional

A more systematic way to build up, say, a two-dimensional triangular lattice might be to just add row after row. But matters come to a grinding halt almost at once, at the two-row or ladder level. Here, try as we will, there seems to be no neat way of writing the intrinsic free energy in terms of the site densities. What has gone wrong is that the ladder has feedback loops: site x can communicate with z not only directly but also via y ; technically, the ladder is no longer simply connected. There are at least two ways of proceeding. First, suppose the loops are small, as in the ladder (Figure 7). The interactions around the loop can be regarded as a special case of a multisite interaction, i.e., ϕ_{xyz} in the ladder, so that the topology of the loop is irrelevant. The only effective way of including multisite interactions is to place them at the same level as one-site interactions, and any two-site interactions must come along for the ride. This introduces the concept of the entropy functional.

Equilibrium statistical mechanics is formally described by the grand partition function Ξ , or by the grand potential $\Omega = -kT \ln \Xi$, which is $-PV$ for a *uniform* fluid in volume V . For a nonuniform fluid, say a lattice gas, Ω also serves as a universal generating function in the sense that $n_x = \partial \Omega / \partial h_x$, as well as $n_{xy} = \partial \Omega / \partial \phi_{xy}$ for two-site interactions, $n_{xyz} = \partial \Omega / \partial \phi_{xyz}$ for three-site interactions, etc.; here of course one has many independent variables, $\Omega(\{n_x\}, \{n_{xy}\}, \{n_{xyz}\})$. The intrinsic Helmholtz free energy is obtained by subtracting out the average one-site energy: $\bar{F} = \Omega - \sum h_x n_x$, so that now $h_x = -\partial \bar{F} / \partial n_x$, and various intermediate quantities arise by subtracting out other energy components. In particular, the (negative) entropy $-TS = \Omega - \sum h_x n_x - \sum \phi_{xy} n_{xy} - \sum \phi_{xyz} n_{xyz} \dots$ has

(13) Percus, J. K. *J. Stat. Phys.* **1986**, *42*, 921.

(14) Percus, J. K. *J. Stat. Phys.* **1977**, *16*, 299.

(15) Samaj, L. *Physica* **1988**, *A153*, 517.




Figure 8. Ring lattice.

all components removed and so is a purely geometric object, with the interactions appearing as the derived quantities $h_x = \partial TS/\partial n_x$, $\varphi_{xy} = \partial TS/\partial n_{xy}$, etc.

Is the entropy functional much more complicated? Not necessarily. For a simply connected network, such as a Bethe lattice with nearest neighbor interaction, $-S$ is just the sum of two-site contributions $\sigma(n_x, n_y, n_{xy})$. And then, more importantly, if the lattice is built up of triangles, simply connected in the sense that exactly one independent path of triangles connects two triangles—as in the triangular ladder (Figure 7)—then $-S$ is simply the sum of three-site contributions $\sigma(n_x, n_y, n_z, n_{xy}, n_{yz}, n_{xyz})$, a bit more complicated but conceptually just as simple.¹⁶ In particular, if there is actually no three-site interaction, $\partial S/\partial n_{xyz} = 0$ allows us to get rid of the triplet density.

Multiconnected Lattices

The entropy approach just described is impractical if the interaction loops are too large. The prototype for such a nonsimply connected lattice is a ring of sites (Figure 8), so in the second approach, we try really hard to find the intrinsic free energy \bar{F} for this model. By the way, the fact that the full free energy—including the external $h_x n_x$ terms—is a minimum, shows that $h_x = -\partial \bar{F}/\partial n_x$, so that the profile equations are automatically obtained. At any rate, solving the ring turns out to be *very* nontrivial, until one notices that there is a collective variable C , a combination of local potentials and densities, that has the same value at each site. Then, if we are willing to write $\bar{F}(\{n_x\}, C)$ as a function of the n_x and C , i.e., more variables than the number of degrees of freedom of the ring, things get much clearer.¹⁷ \bar{F} can be chosen so that $h_x = -\partial \bar{F}/\partial n_x$ still holds, but now \bar{F} is stationary with respect to C : $\partial \bar{F}/\partial C = 0$, and when this is done, \bar{F} is still a sum of pair terms $f(n_x, n_{x+1}, C)$ with just an additional collective amplitude term $\Delta(C)$. In other words, by expanding our viewpoint to include descriptions both in the small and in the large, the system takes on a simplified structural form, and the relation between the small and the large is encoded in the stationary properties of the free energy.

Once you are sensitized to them, multilevel descriptions appear repeatedly in studies of lattice gases. For this purpose, of course, more complex networks must be investigated. One technique for doing so¹⁸ starts by associating a collective amplitude with each string of ($q = 2$) sites and then combines parallel channels, reducing the network to a skeleton of bifurcating ($q > 2$) sites to which the same technique can be applied. Another one introduces subsidiary variables to close loops of an initially simply connected network or to identify sites of a spatially overlapping Bethe lattice. But the take-home message is always the same: one ends up by expanding the space of local densities $\{n_x\}$




Figure 9. Effective van der Waals interaction.

to include a set $\{C_\alpha\}$ of nonlocal amplitudes which are related by the stationarity of a free energy on the expanded space: $\partial \bar{F}(\{n_x\}, \{C_\alpha\})/\partial C_\alpha = 0$, with profile still given by $h_x = -\partial \bar{F}/\partial n_x$. And entropy generalizes in a very similar fashion.

van der Waals Revisited

Our excursion into lattice gases has provided us with invaluable background for the continuum fluids that we are really interested in. But the relation between discrete and continuous versions has to be spelled out. On the one hand, we can imagine the lattice as a coarse grid sampling of the continuum, with the advantage that interaction potentials need only be defined on a few grid points; models of the solidification transition, in which there is a natural grid attached to the system, have been constructed along these lines. On the other hand, we may use the lattice simply as an accounting device, with the limit of vanishing grid spacing always in mind. In the latter case, a clearly suggested strategy would be to take the continuum limit of lattice models that approximate a full grid decomposition of physical space.

Let us start then with the simplest model representation, that of a simply connected tree-like structure, e.g., a nearest neighbor Bethe lattice, as in Figure 6. But we want the branches to go an arbitrary distance in an arbitrary direction; the mesh spacing here is the length of the smallest interaction link. As the result of this construction, continued indefinitely, any two points in physical space are arbitrarily close to two sites that are “nearest neighbors”. The density of paths connecting two points in this way is underestimated, but the interaction amplitude can be suitably increased. The result is a weird but not nonsensical approximate topology. Rather than try to describe the structure of the associated interaction, let us simply go to the entropy viewpoint in which, after the resulting approximate geometry is accounted for, the interactions came along automatically.

As we have seen, the entropy for this model is just the sum of (one-site and) two-site entropies for those sites that interact—here, every pair, in the limit. There is no difficulty in taking the limit, with the obvious correspondence $\tau \sum_r \rightarrow \int dn_r$, $n_r/\tau \rightarrow n(r)$, $n_{rr}/\tau^2 \rightarrow n_2(r, r')$ for mesh volume τ . We find that, to within an irrelevant additive constant, $-S$ consists of the anticipated one-body part $\int n(r) \ln n(r) dr$ and two-body part $(1/2) \int \int n_2(rr') \ln(n_2(rr')/n(r)n(r')) dr dr'$, together with an extra correlation contribution $(-1/2) \int \int (n_2(r, r') - n(r)n(r')) dr dr'$. Consequently, we can write down the $\varphi(r, r'), n_2(r, r')$ “profile equation” and use it to eliminate $n_2(r, r')$, obtaining an intrinsic free energy $\bar{F}[n]$. The result is very close to the coreless mean field van der Waals free energy, but not quite: the mean interaction uses the combination $\bar{\varphi} = kT(1 - \exp(-\varphi/kT))$ instead of φ (where k is the Boltzmann constant) (Figure 9). Thus it can deal directly with infinite hard particle potentials—van der Waals

(16) Percus, J. K. *J. Stat. Phys.* **1990**, *60*, 221.

(17) Percus, J. K.; Zhang, M. Q. *Phys. Rev.* **1988**, *B38*, 11737.

(18) Samaj, L.; Percus, J. K. *J. Stat. Phys.* **1993**, *73*, 235.

cannot—due to being based on a valid, if curious, physical model.

The Next Step

The inclusion of small feedback loops is certainly the next step in a program of representing a full spatial grid by an approximating network. We have found that this process is easy to start, by the simple expedient of “thickening” the Bethe lattice so that it becomes a tree composed of triangles. The entropy is then a sum of local triplet contributions, and we can take the continuum limit just as we did before, resulting in the addition of an obvious triplet contribution of the form $(1/6) \int \int \int n_3 \ln(n_3 nnn/n_2 n_2 n_2)$ and the subtraction of an integrated triplet correlation. If there is no explicit triplet potential, the condition $0 = \delta S/\delta n_3$ allows us to solve for and eliminate n_3 . Interestingly, the solution obtained in this way¹⁹ is precisely the famous Kirkwood superposition approximation that, with obvious arguments, $n_3 = n_2 n_2 n_2 / nnn$. Following this, the continuum equivalent of $h = \partial TS/\partial n$, $\varphi_2 = \partial TS/\partial n_2$ gives us a pair of integral equations to solve.

The end result of this process is a slightly modified “CHNC” equation, one of the more reliable integral equation approximations (the result of a diagrammatic expansion technique in which only diagrams—i.e., networks—constructed recursively by series and parallel combinations are retained). Is this good? Not really, because if our improved lattice-continuum approximation is only able to reproduce a reasonable integral equation—with all of the attendant computational difficulties—our enthusiasm for the approach tends to wane. Perhaps the systematic increase in local complexity is not the best use of the insights we have gained from lattices.

Weighted Density Functionals

The inception of global nonsimple connectivity is certainly a more dramatic alteration of the structure of a lattice gas network. But we know that the simplest way of representing this more complex structure is by expanding the roster of variables on the network to include “collective” amplitudes associated with large groups of sites. Translated to continuum fluids, we would have augmented the local density by densities at poorer levels of resolution. In fact, one of the very few nonuniform continuum fluids for which the free energy density functional is exactly known—the gas of one-dimensional hard rods—has a free energy that is most simply represented by regarding the average density over the volume of a core, $\bar{n}(x) = \int n(y) \omega(x-y) dy$ [where $\omega(x-y) = 1$ for $x-y$ inside a core, else 0], as a coarse-grained density field to be used in a joint description. One possibility that this suggests, for non-hard-core interactions in other than one-dimensional space, is that we devise a model built upon the same kind of *weighted density*, but with a specifically appropriate weight function⁷ $w(r-r')$; indeed, there is no reason why the weight cannot involve the density as well:⁸ $w(r-r', \bar{n}(r))$. So now everything depends upon having enough (but not too much!) information about known properties of the system, available to find $w(r-r', \bar{n}(r))$ once the structure of the model, i.e., the way that $n(r)$ and $\bar{n}(r)$ enter into it, is

decided upon. The major source of information routinely used is the pair distribution (technically, the direct correlation function) of the *uniform* fluid as a function of pair distance and density. This general approach has indeed been very effective in studies of both simple and complex fluids, in a wide variety of nonuniform states.

But let us survey the scene from Olympian heights. The trouble with guessing a reasonable form for the coarse-grained $\bar{n}(r)$ is that while it seems to allow full play for one’s intuitive feelings as to the essentials of the structure, it amounts in the end to being guided solely by mathematical and computational convenience rather than by any sophisticated structural concepts. This is both ad hoc and potentially dangerous; for example, the usual recipes fail in the very important case of Coulomb forces and require further modification.²⁰ However, we know from our lattice gas studies that (a) the additional “collective variables” can be very complex, but (b) the implicit equations that they satisfy are most readily obtained in the form of stationary free energy or entropy conditions. This format, e.g., $\bar{F}[n, \bar{n}]$, is not novel in the continuum framework either,²¹ but again employs a general form mathematically extrapolated from weak nonuniformity expansions. So it is easy to criticize, but that is not necessarily constructive. What more rational suggestions are there? Now one must equivocate. On the one hand, judging from the few exactly solved systems, evidence for the necessity of a hierarchical structure, such as $\bar{F}[n, \bar{n}, \bar{\bar{n}}, \dots]$ is compelling. But on the other hand, even though the nature of the required \bar{n} will come automatically from stationarity in the form $\delta \bar{F}/\delta \bar{n} = 0$, $\delta \bar{F}/\delta \bar{\bar{n}} = 0$, ... this is not helpful until one has established the nature of the functional form of \bar{F} , presumably from solvable models and limiting cases. And no preconceptions allowed: the coarse-grained \bar{n} for a two-phase interface enters only²² via derivatives!

The Future

The above certainly suggests a strategy, but it is scarcely the only one. What is holy about the free energy density functional, the entropy density and pair density functional, and similar entities? Even if one focuses on observable conjugate pairs such as density and external field—not an unreasonable restriction—conjugate pairs (whose cross-derivatives satisfy integrability conditions) can be constructed in other ways, e.g., $\{n(r) \exp(-\beta\mu(r))\}$ and $\{\exp(\beta\mu(r))\}$, which might make for a simpler description. And of course, if one wants to introduce other indirectly observable quantities such as $n_2(r, r')$, why not consider such as nearest neighbor pair distributions, which are much closer to any topology one might like to graft onto the system? It will be interesting to see how the next decades relegate to embarrassing obviousness and/or oblivion our proud edifices of the moment.

Science is a social activity, and the roster of scientists who have motivated, guided, or otherwise strongly biased the direction of the research reported here is long and illustrious. The reference list is not intended to supply more than a small subset, unduly focuses on the present writer, and is woefully incomplete.

(20) Tarazona, P.; Evans, R. *Mol. Phys.* **1984**, *52*, 847. Jones, R.; Gunnarson, O. *Rev. Mod. Phys.* **1989**, *61*, 689.

(21) Groot, R. D.; Van der Eerden, J. P. *Phys. Rev.* **1987**, *A36*, 4356.

(22) Percus, J. K. *J. Chem. Phys.* **1984**, *81*, 452.

(19) Laird, B. B.; Wang, J.; Haymet, A. D. *J. Phys. Rev.* **1993**, *E47*, 2491.