J. Opt. A: Pure Appl. Opt. **11** (2009) 114002 (9pp) doi:10.1088/1464-4258/11/11/114002

REVIEW ARTICLE

Comparing experiment and theory in plasmonics

W L Barnes

School of Physics, University of Exeter, Stocker Road, Exeter EX4 4QL, UK

Received 16 February 2009, accepted for publication 6 May 2009
Published 16 September 2009
Online at stacks.iop.org/JOptA/11/114002

Abstract
Progress in plasmonics has been greatly assisted by developments in experimental techniques and in numerical modelling. In this paper I look at some of the difficulties that emerge when comparisons are made between experiment and theory. Using four examples I illustrate what some of these difficulties are, from the perspective of both experiment and of modelling. Although there are many aspects to consider, two seem to be of particular concern at the time of writing; identifying the most appropriate relative permittivity (dielectric function) to describe the optical response of the metals used, and how best to make space discrete when using numerical models that rely on this approach.

Keywords: experiment numerical simulation plasmonics theory

(Some figures in this article are in colour only in the electronic version)

The invitation to give a presentation at the first international workshop on theoretical and computational nanophotonics (TACONA 2008) in Bad Honnef gave me the opportunity to reflect on a number of difficulties concerning the relationship between experiment and theory in the area of plasmonics. My views on this are coloured by the fact that I am primarily an experimentalist and have only very limited experience of the computational/numerical approaches that have been developed. What follows is not therefore a definitive comparison of the state-of-the-art in terms of experiment and theory. Nonetheless, I hope that it may still provide a useful starting point for those not familiar with the complex array of problems that arise when we look more closely at comparing experiment with the results of modelling in the quest for better understanding.

In plasmonics we seek to harness the resonant interaction between light and the conduction electrons in the surfaces of metallic structures to control light on sub-wavelength length scales [1, 2], and to manipulate light–matter interactions [3–6]. The topicality of plasmonics (see figure 1) derives in part from important developments in experimental and numerical techniques. Nanofabrication tools such as electron-beam lithography allow us ever better control over the kinds of structures we can make. Improved computational approaches such as finite-difference time-domain simulations allow us to map the electromagnetic field in ways that give us the impression that we can peer directly into optics at the nanoscale. The new knowledge that flows from the interplay of experiment and theory is leading to many interesting potential applications (figure 1). Despite the very significant progress that has been made in recent years there is still room for improvement in trying to match the results of experiments with those of theory, especially as we strive to improve our understanding. It is also important that our knowledge base is solid as the field moves into less familiar areas such as nonlinear plasmonics [2, 7] and the combination of plasmonics and gain [8, 9].

Just before the TACONA 2008 meeting a wonderful example of combining experiment and theory from a completely different realm of physics appeared. Dürr *et al* [10] produced predictions for the masses of light hadrons from numerical simulations based on quantum chromodynamics (QCD). The agreement between experiment and theory, reproduced in figure 2, is spectacular, especially when one realizes that this calculation has taken the community 30 or so years to produce. Commenting on this work, Wilczek noted that the agreement between experiment and theory, which is within error (of both), goes beyond 'simply' validating the numerical approach and with it the ability of QCD to describe the world, it also means that with 'numerical techniques that reliably reproduce what is known, we can address the unknown confidently' [11]. It is perhaps this goal—using numerical

 © 2009 IOP Publishing Ltd Printed in the UK

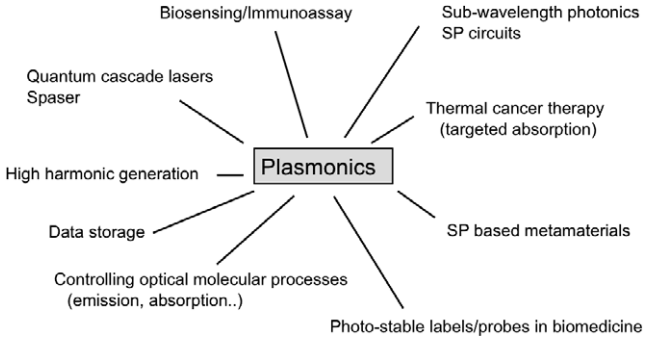

Figure 1. Plasmonics concerns the confinement and control of light in the sub-wavelength regime. This diagram indicates some of the areas of application that are either already being pursued or are being considered.

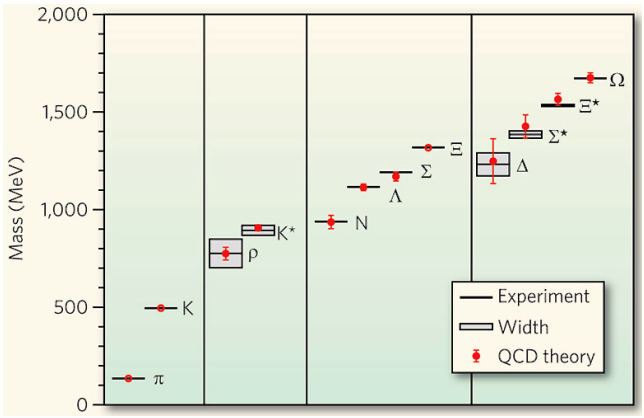

Figure 2. This figure, reproduced from [11], shows the comparison between measurements for the masses of light hadrons and predictions for those masses based on numerical calculations using quantum chromodynamics. Reproduced with permission.

models to explore new exciting areas of physics—that spurs us to develop them. However, we should not forget the vital role played by analytical models in developing a deeper physical understanding [12]. In nanophotonics, and especially in plasmonics, our confidence is not in doubt with regard to the Maxwell equations, but the material parameters and the boundary conditions that we impose are the subject of much debate, and it is these that form the cornerstone of the discussion below.

This paper is based around four exemplar structures. I begin by comparing the spectrum of light scattered by a gold nanodisc obtained from experiment with simulations based on finite-element modelling. The second example concerns a comparison of various numerical approaches with an exact analytical result for the scattering of light by a gold nanosphere. A number of other topics are then discussed with reference to the plasmon modes supported by a planar metal film. The last example is far from the nanoscale, it concerns a set of copper grids with interesting properties in the microwave part of the spectrum. Despite the significant change in spectral region, this last example highlights further difficulties in modelling the electromagnetic response of structured metals.

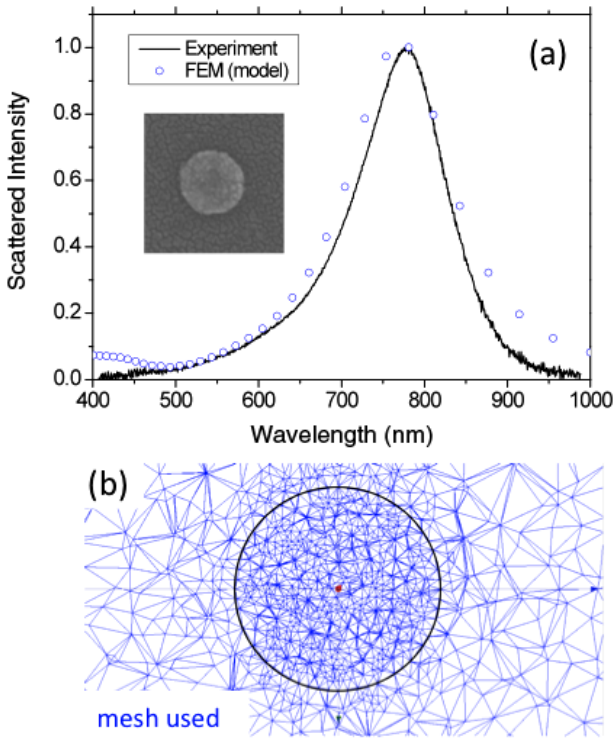

Figure 3. Scattering of light by a gold nanodisc. The upper panel (a) shows the scattering spectrum from a gold disc 120 nm in diameter and 30 nm thick. The line corresponds to data acquired using dark-field scattering spectroscopy. Illumination was unpolarized, azimuthally symmetric and covered a range of polar angles from 30° to 50°. The range of (polar) collection angles was estimated as just under ±30°. The particle was made by electron-beam lithography on a glass substrate, and was covered with an index-matching oil ($n = 1.52$) so as to ensure the particle was embedded in a homogeneous environment. The circles correspond to data generated from a finite-element model (HFSS™, Ansoft). In the model the mesh involved ∼50 000 tetrahedra in the metal nanoparticle. The inset shows a scanning electron micrograph (SEM) of the gold nanoparticle. It has been coated with a thin layer of chrome (for the purposes of acquiring the SEM). From the micrograph we can determine the major and minor axes of the particle to be 122 and 118 nm. The lower panel (b) is a schematic of the mesh generated by the modelling package to produce the data shown in (a). In this model the incident angle of the light was 35° and the scattered light was integrated over collection angles of ±30°. The experimental and numerical scattering intensities have been scaled so as to match in terms of peak height.

First example: a gold nanodisc

Let us begin our discussion by considering the Drosophila of plasmonics, a single metal nanoparticle. Figure 3 shows data relating to the scattering of light by a gold disc 120 nm in diameter and 30 nm thick. In figure 3(a) the experimentally derived scattering spectrum of light between the wavelengths of 400 and 1000 nm is shown (solid line), together with the results from calculating the scattering spectrum using a commercial finite-element numerical package (HFSS™, Ansoft.v10). Both experiment and the numerical model show a strong peak in the scattered intensity at a wavelength of ∼780 nm. This peak corresponds to the excitation of the lowest-order (dipolar) localized surface plasmon resonance (LSPR) of the gold nanoparticle [13].

The position and width of the peak in the scattering spectrum are similar for both experiment and theory in figure 3(a) and one might be tempted to say that the agreement is good, but I think that would be a mistake for a number of reasons. Closer inspection shows that the widths are in fact different: why? Why do the shapes of the two spectra not match better? Why does the model show an increase in scattering intensity for wavelengths below ∼500 nm when the experimental data are getting weaker, not stronger? These questions naturally follow from an examination of the data shown in figure 3(a), but what if we also concern ourselves with how those data were derived.

First let us look at experiment. What about the geometry of the particle; is it really a disc? A SEM is shown in the inset to figure 3(a). Now we see that the particle does not have perfect circular symmetry. Further measurements show that its major and minor axes are 122 nm and 118 nm, respectively; is this slight asymmetry significant? And what about the grains we see in this figure? We need to be very careful here, although we know that evaporated gold films are granular in structure [14, 15], the granularity we see here is more to do with the thin (5 nm) chrome layer evaporated over the whole sample to avoid the build-up of surface charge whilst the SEM is acquired—a common procedure. If for the moment we ignore this inhomogeneous (granular) nature of the disc, what can we say about the quality of the structure in terms of its optical properties? Specifically, what is the relative permittivity (dielectric function) of the gold? Was there any surface contamination, and if so does it matter? Are the illumination and collection conditions the same for experiment and simulation? One aspect we did take care of in collecting these data was to ensure a homogeneous surrounding environment; many models struggle to deal effectively with a particle on a substrate when the superstrate medium is of a different refractive index. We will come back to discuss some of these matters soon, but first what of theory?

Second example: a gold nanosphere

Finite-element models such as that used for figure 3 are based on assuming that space can be broken up into a number of discrete points. The mesh of points used for the calculation shown in figure 3(a) is shown in 3(b). How good a representation of the optical field associated with the particle resonance does this mesh allow? What about the relative permittivity used in the model (in this case the data were interpolated from the literature [16]), and how well does that permittivity represent the optical response of the particle? This question has been discussed recently by Drachev *et al* [17].

Being an experimentalist I naturally ask whether different models agree; if they do not I will want to ask theorists why not, what are the different limitations and approximations etc. However, turning this around might I think be rather informative; it would be fascinating to try and make gold discs of the same dimensions by a range of techniques in different labs and see how well they compare with each other—something of a challenge. For the moment, however, let us return to a comparison of a number of different models. I have chosen a spherical nanoparticle because the scattering of light by such an object has an analytic solution that we can use as a reference set of data against which to compare numerical models. In figure 4(a) are plotted the results from five different approaches. The scattering efficiency [18] of an 80 nm diameter gold sphere has been calculated with each approach. The black (solid) line is the result of using analytic expressions from Mie theory [18] which act as the reference data set. The numerical approaches considered include: T-matrix [19–21], discrete dipole approximation (DDA) [22, 23], finite-difference time-domain (FDTD) [24, 25] and finite-element modelling (FEM) [26]. There are still other approaches, for example the boundary element method (BEM) [27].

From figure 4(a) it seems that there is reasonably good agreement between the different techniques, except the FDTD approach which looks to be doing less well. However, this is both misleading and highlights an important difference between FDTD and the other techniques mentioned. In the FDTD approach the relative permittivity needs to be used in the code over a wide range of frequencies, much wider than the range shown in figure 4(a). To do this a Drude-type model was used for the permittivity in the FDTD calculation shown in figure 4(a), whereas interpolated literature data [16] were used for the other techniques so the comparison is not really fair. In figure 4(b) the same Drude-type model for the permittivity was used both for the Mie calculation and for the FDTD-derived data. In figure 4(b) are shown a number of FDTD data sets for different mesh sizes. The agreement between the Mie calculated data and the FDTD-derived data gets better as the mesh size is reduced, but it looks as though a very fine mesh size (<2 nm) may be needed to obtain good agreement (see inset). Whilst discussing the importance of meshing, it might also be worth noting that even in trying to model a simple disc structure, field profiles derived by finite-element modelling can show asymmetries. This is because although the sample geometry may be specified as symmetric the mesh that is derived within the numerical package need not necessarily be symmetric; again one has to take care. This problem seems to be associated with the tetrahedral meshing of finite-element approaches: FDTD techniques typically employ a cubic mesh that may be better suited to avoiding this problem.

Third example: planar metal films

Let us step back for a moment and consider an example that should offer fewer problems in matching simulation to experiment. Figure 5(a) shows data obtained using the Kretschmann–Raether technique to excite the surface plasmon-polariton mode on a planar metal film [28]. A prism is used to achieve the required momentum matching of incident light to the surface plasmon-polariton (SPP) mode. For an appropriate metal film thickness, and for p-polarized incident light, coupling to the SPP mode is revealed by a dip in the reflectivity [29] (figure 5(a)). One can fit modelled data based on Fresnel's equations to the experimental data to determine the material parameters [30]: the solid line in the figure 5(a) shows such a fit. The parameters of the metal film derived in this way (wavelength = 633 nm) are $\varepsilon = -10.73(\pm 0.02) +$

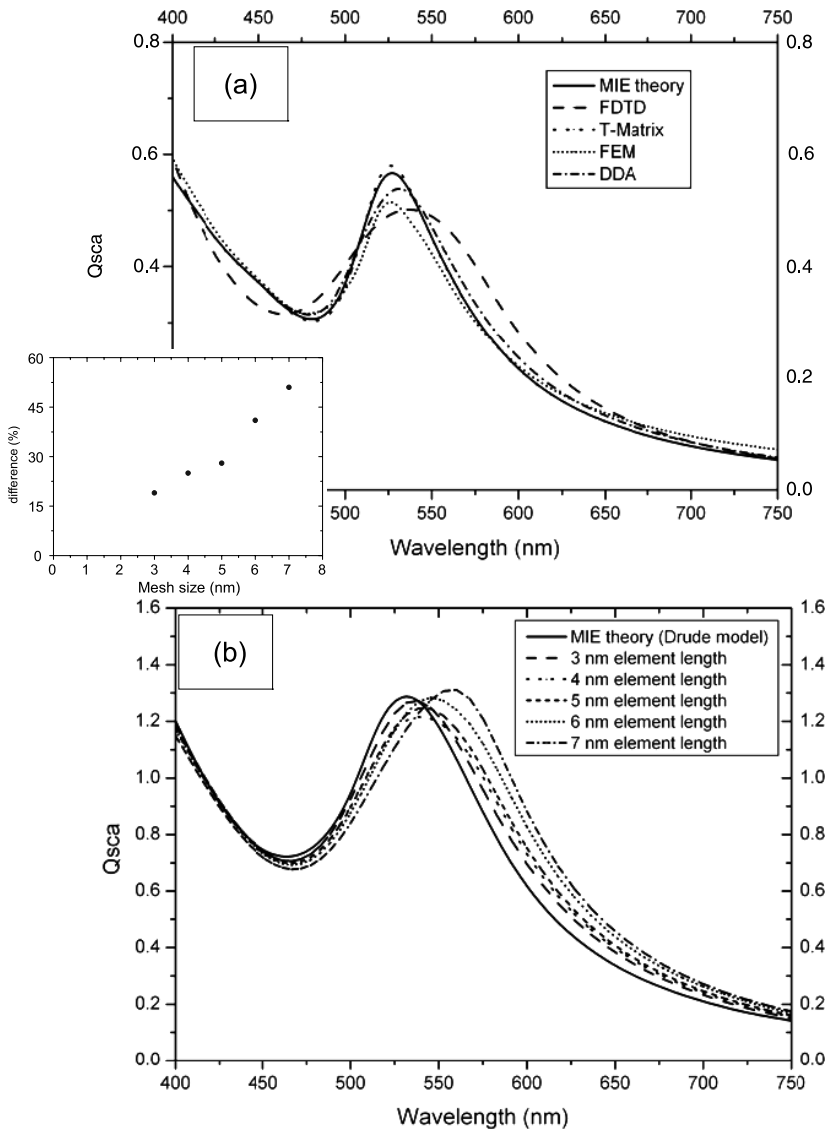

Figure 4. Theory/modelling derived scattering efficiencies (defined in [18]) for a gold sphere of 80 nm diameter in a vacuum. The upper panel (a) shows the results from Mie theory, from a T-matrix model (http://www.giss.nasa.gov/~crmim/t_matrix.html), from a finite-difference time-domain (FDTD) model (numerical solutions FDTD), from a finite-element modelling (FEM) package (HFSS™, Ansoft), and from a discrete dipole approximation (DDA) model (http://www.astro.princeton.edu/~draine/DDSCAT.html). For all of these models except FDTD the relative permittivity values were interpolated from Johnson and Christie [16], for the FDTD model the permittivity was approximated using a Drude model: parameters are given below. The parameters used in the different models meant that to compute the data shown in (a) took the following times (2 GHz processor, 4 Gb RAM): Mie (1 s), T-matrix (100 s), DDA (1000 s), FDTD (10 000 s) and FEM (10 000 s). The meshing details were as follows: FEM element length 3 nm, meshed at a wavelength of 525 nm; FDTD element length 3 nm, meshed at a wavelength of 525 nm; DDA inter-dipole spacing 3 nm. In the lower panel are shown data from just the Mie calculation and FDTD for a sphere of 100 nm diameter. Here, in contrast to (a), for both Mie and FDTD, the permittivity was approximated by $\varepsilon = 1 - \omega_p^2/\omega^2 - i\Gamma\omega$ with $\omega_p = 1.3 \times 10^{16}$ rad s^{-1}, and $\Gamma = 5.0 \times 10^{13}$ s^{-1}. Data from FDTD calculations are shown a number of different element lengths in a cubic mesh. The inset shows how the difference at 600 nm between Mie and FDTD data depends on this element length. (In the Mie and T-matrix calculations the full response, i.e. for all angles, is obtained from the one calculation, for the numerical approaches this is not the case and each angle has to be computed separately; this point should be borne in mind when comparing the calculation times.)

$1.279i(\pm 0.005)$, d (thickness) $= 47.2\,\text{nm} \pm 0.1\,\text{nm}$. How does the permittivity derived in this way compare with the literature? We find that for gold at a wavelength of 633 nm Johnson and Christie [16] give $\varepsilon = -12.3+ \sim 1.2i$, Lynch and Huttner [31] give $\varepsilon = -10.4 + 1.4i$, whilst Innes and Sambles [32] have $\varepsilon = -11.8 + 1.36i$. The message here is that the relative permittivity of a metallic object/film depends on the way it was fabricated, possibly its size, the substrate it is formed on, etc. In other words the relative permittivity is not something that is 'constant'.

If rather than looking at the quality of the match between experiment and theory in figure 5(a) directly we instead look at the residuals (figure 5(b)) we see that there are some systematic differences. Could these perhaps be due to surface roughness? An AFM scan of the film used in this experiment, shown in figure 5(c), reveals roughness on the level of a few nanometres:

Table 1. Adapted from [35]. The permittivity derived from fitting models to experimental reflectivity data of a number of thin silver films are shown. By using a combination of prism and grating coupling data were obtained that included responses due to surface plasmon-polaritons on both the silver–air interface and the silver–glass interface. In the model the silver layer for each film was divided into 5 sub-layers. The data presented here are the results obtained for the permittivity for the top (adjacent to air) and bottom (adjacent to glass) sub-layers. Not the significant difference between the permittivity derived for the two sides of the film, indicating that even a 'simple' metal film cannot be completely described by a single vale of the relative permittivity.

Silver film thickness (nm)	Glass/silver interface		Silver/air interface		Uncertainties	
	ε_r	ε_i	ε_r	ε_i	$\Delta\varepsilon_{r,i}$	Thickness (nm)
31.1	−20.41	0.92	−14.74	0.69	±0.1	±1
50.1	−21.06	0.94	−13.85	0.90	±0.1	±2
58.6	−20.99	0.79	−15.56	0.58	±0.1	±2
76.2	−21.36	0.78	−14.50	0.43	±0.1	±4
86.2	−21.01	0.82	−17.02	0.49	±0.1	±5
125.3	−20.86	0.79	−17.92	0.49	±0.1	±5

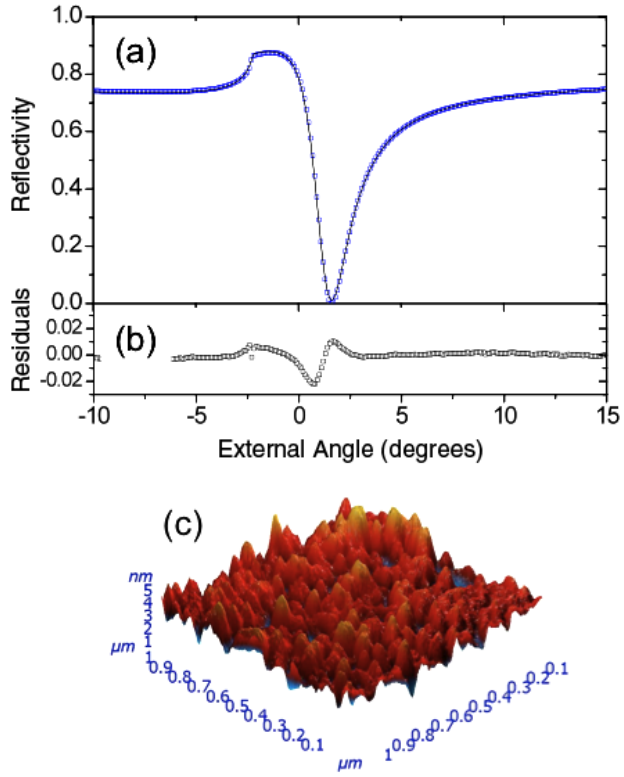

Figure 5. Prism coupling to surface plasmon-polaritons (SPP). Panel (a) shows experimentally derived data for the reflectivity from a prism coated with a thin gold film. Also shown is a continuous line, this is the result of a model used to fit theoretically derived reflectivity data (based on Fresnel's equations) to the experimentally acquired data. The parameters for the gold film derived from this fitting procedure were $\varepsilon = -10.73(\pm 0.02) + 1.279\mathrm{i}(\pm 0.005)$, d (thickness) $= 47.2$ nm $\pm$ 0.1 nm. The next panel (b) shows the residuals (difference between experiment and model). The last panel (c) shows the results from an atomic force microscopy (AFM) scan of the gold film.

our model did not include this. There is more. If we consider the metal film used in the experiment shown in figure 5 we note that there are two metal film surfaces, so there should be two surface plasmon-polariton modes, one associated with each surface [33]. We do not normally couple to the mode associated with the metal–prism interface; light incident from the prism side does not have enough momentum to couple to this mode [34]. However, using a combination of prism and grating coupling Nash and Sambles were able to couple to both modes [35]: their results are very interesting. In trying to fit theory to the experimentally derived reflectivity data these authors had to assume that the relative permittivity of the silver varied across the thickness of the film. They found a marked difference in the permittivity on the two sides of the metal film, the difference increasing with film thickness as shown in table 1. This difference was tentatively attributed to the free (air) surface of the metal being more porous.

What are we to make of all of this? To describe the optical response at a given frequency by just one (complex) number, the relative permittivity, is often a poorly justified approximation. There is no 'correct' value of the permittivity. The question, 'what permittivity should I use in my model?' might be better stated as 'what is the most appropriate permittivity to use in this particular situation?'. Even that might be misleading though: for example, the width of a plasmon mode might be broader than one had expected, implying that the damping is greater than expected. To deal with this one might simply increase the imaginary part of the permittivity of the metal, but perhaps the origin of the increased damping is roughness, or non-local effects. It is all too easy for the permittivity to be used as a dumping ground for all sorts of other factors that are not fully considered or properly accounted for. Note that Drachev *et al* [17] found that surface roughness does not result in a change in the relative permittivity.

Surface roughness is more than just a problem for the modellers, it can have important consequences for applications. There are many groups pursuing the use of plasmon modes as a wave-guiding technology [36–38]. Berini and co-workers in particular have conducted many detailed investigations to explore the use of long-range surface plasmons guided by metal stripes. For such guides, reducing the thickness of the metal film should be advantageous as it reduces the amount of metal in which dissipation may take place. However, as the metal film thickness is reduced performance does not increase as expected—roughness and grain boundaries start to cause problems in ways that are not fully understood [39].

Roughness can play another important role, this time advantageous rather than detrimental. Surface roughness can lead to electromagnetic hot spots, small regions of space where the electromagnetic field is much higher than in the immediate

surroundings. There are two mechanisms behind this kind of enhancement. The first concerns a sharp metallic protrusion on a metallic nanostructure and is a combination of a lightning-rod effect associated with the sharp protrusion [40, 41] coupled with a plasmon resonance associated with the metallic nanostructure [42]. The plasmon resonance gives the hot spot a resonant character. The second situation involves a short gap between two metallic objects [43, 44]. Again, if there is a plasmon resonance associated with the metallic nanostructure then this will make the enhancement in the gap resonant [45–47]. Li and Schatz [48] found that including roughness in their DDA-based model for the surface enhanced Raman scattering (SERS) signal from gold dimers led to a better agreement with data from experiment. Roughness is a difficult aspect to deal with in models. First, the length scale associated with roughness is often comparable to the mesh size used in numerical techniques, making the inclusion of roughness in these approaches challenging. Second, whilst some kind of roughness can be assumed, as was done by Li and Schatz [48], the detailed morphology in any given situation is often very hard to determine experimentally, and yet hot spots associated with roughness can dominate the behaviour of the system [49].

Whilst roughness is widely recognized as a possible reason for the lack of agreement between experiment and theory, surface contamination is less well recognized, indeed it is often ignored. Most experimental studies are not undertaken in ultrahigh vacuum, or with atomic level cleanliness a priority. As a result surface contamination has to be considered an ever present problem. Frequently we simply ignore it, hoping that it will somehow be incorporated, along with roughness etc, into a modified relative permittivity. A good example that highlights just how important the role of contamination can be comes from looking at the localized surface plasmon resonances associated with copper nanoparticles. Copper is sometimes not considered to be a particularly good metal in terms of supporting plasmon modes at optical wavelengths. However, as Van Duyne and co-workers showed [50], this apparently poor performance is due to the presence of copper oxide on the surface. Using acid to remove this oxide, the strength of the plasmon resonance was found to increase and the width decrease, to the point where the resonance was comparable with that obtained with gold.

This sensitivity of surface plasmon modes to surface contamination is the basis of what is so far the only significant commercial application of surface plasmons, biosensing [51] (figure 1). In biosensing the contaminant is what is interesting, surface chemistry being used to functionalize the metallic structure in an attempt to ensure that only the molecules of interest bind to the structure and are thus detected. When such a binding event takes place the LSPR associated with the metallic nanoparticle shifts and may change in strength [52]. By monitoring these changes the presence of just a few hundred molecules has been observed [53, 54]. It is interesting to note that the sensitivity of the surface plasmon technique is now at such a level that we have to be concerned about very subtle temperature and pressure changes [55]. Furthermore, we are prompted into questioning how best to model the refractive index of a bound layer of molecules, especially if there are just a few, i.e. a low density of coverage [56].

The last topic I wanted to consider under the theme of 'let's hide it in the relative permittivity' is the breakdown of the bulk description. So far we have implicitly assumed that the relative permittivity is at least in principle a good parameter in describing the optical response of a metallic nanostructure. However, it is a bulk property; specifically it does not recognize the atomic nature of the metal. At some point, especially as we become more concerned with very small volumes, this assumption must break down. I have already discussed roughness, but there are also other phenomena such as surface scattering to consider [57]. The effect of the surface is often treated simply as a modification to the damping rate of the conduction electrons contributing to the permittivity [18]. Garciá de Abajo has recently shown that a significantly better match with experiment can be obtained by incorporating non-local effects (where the response of the material at any given location depends on more than just the electric and magnetic fields at that same location) [58]. Density functional theory (DFT) has been used to tackle the optical response of very small (<10 nm) gold particles with good effect [59]. More recently, Schatz and co-workers used DFT to calculate the Raman intensities for pyridine attached to a Ag_{20} cluster [60].

Fourth example: copper grids

I want to use this final example to discuss the mesh of points used to model the field in many of the numerical techniques available to simulate the electromagnetic response of structured metals. In the optical regime experience shows that for good agreement between numerical modelling and analytical techniques one needs to work with a mesh size (in regions of strongly varying fields) down towards 1 nm (see figure 4). The key problem here is the big mismatch between the important length scales in the problem, the mesh size is a factor of 10^2–10^3 smaller than the wavelength of light. As with the other problems we have looked at, there are associated experimental problems with trying to characterize morphology on the same length scale, although some recent work has set out to address this problem (see for example [61]). Here I would like to illuminate some aspects of the meshing problem by moving to the microwave regime, a regime where one can be surer about the exact dimensions of the structure under study.

This final example consists of a metallo-dielectric stack in which the metal sheets are not continuous but rather take the form of grids [62] (see the upper panel of figure 6). Metallic grid (net) structures are of considerable interest at present because they are one of the major design types in photonic metamaterials [63]. The lower panel of figure 6 shows the transmittance of a set of metallic grids; the details of the structure are shown schematically in the upper panel. The period of the grid is 5 mm, so that the frequencies of interest are ∼10 GHz. The data in the lower panel show two groups of modes. The four modes centred around 10 GHz are the first-order Fabry–Perot modes of this structure, those centred around 20 GHz are the second-order modes. The

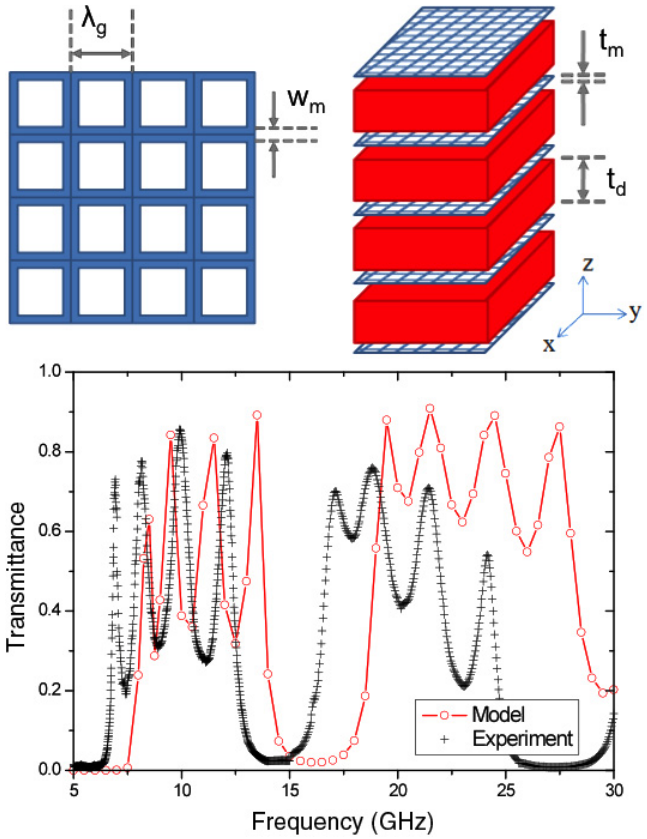

Figure 6. Transmittance of a metallic grid structure. The structure, shown in the two schematics in the upper part of the figure, comprised five grids made from copper separated by dielectric spacers. The dimensions, indicated in the upper panel, were: grid period (λ_g) 5 mm, wire width (w_m) 0.2 mm, wire thickness (t_m) 18 μm, thickness of dielectric spacer layer (t_d) 6.35 mm. The experimentally derived transmittance data are shown in the lower panel, the electric field was aligned along one of the grid axes. The frequency range spanned (5–30 GHz) corresponding to wavelengths between 1 and 6 cm. Also shown in this panel is the initial result from a finite-element model (HFSS™, Ansoft). For this simulation run no mesh was specified, rather the structure was given as input to the modelling package and the internal meshing algorithm used. The unit cell for the calculation involved one square of the grid in the x–y plane and extended through the entire depth of the structure. For this first calculation there were 5129 tetrahedra in the mesh used in this cell.

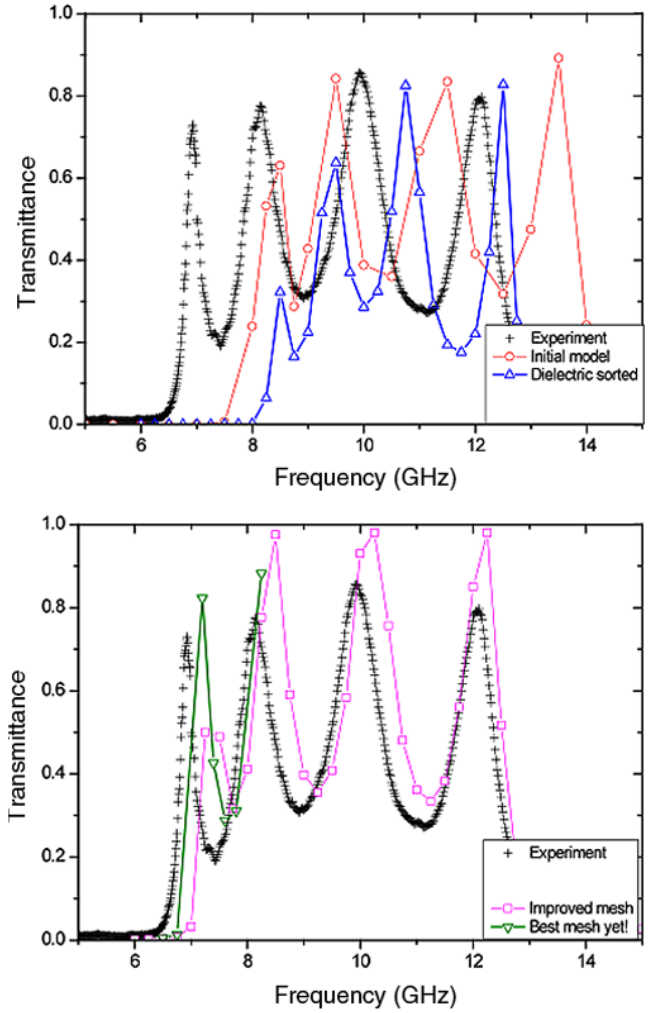

Figure 7. Further modelling of the transmittance of the structure shown in figure 6. In the top panel are the results from the initial model (red circles) and from a revised model (blue triangles). The difference between the two is simply a revised permittivity for the spacer layer, an independently derived value of 3.00 + 0.004i, rather than the assumed parameter of 2.55 + 0.0i initially used. The lower panel shows the result of a model with a higher density of meshing points. For the simulations in the upper panel no meshing was specified, the calculation proceeded as for the lower panel of figure 6. For the lower panel the unit cell was kept the same as before but now extra calculation points were forced in the layers occupied by metal grids. For the 'improved mesh' in the lower panel each metallic grid layer was broken up into ∼1000 tetrahedra, for the 'best mesh yet' ∼100 000 tetrahedra.

first attempt at modelling these data using a finite-element package (HFSS™, Ansoft.v10) produced the results shown in the lower panel of figure 6; the mismatch between experiment and theory is clear. The metal was assumed to be a perfect conductor. Here the geometry of the structure was specified in the model, the internal routines of the modelling package being used to determine the mesh. The unit cell for the calculation involved one square of the grid in the x–y plane and extended through the entire depth of the structure in the z-direction. For this first calculation the mesh comprised 5129 tetrahedra in this cell. At first sight it looks as though there is simply a 10% mismatch between experiment and theory in the frequencies of the modes. This could easily be accounted for if the thickness and/or permittivity of the spacer layer is wrong. At 6.35 mm the spacer thickness is easy to determine to better than 1%, and checks confirmed that the error did not lie here; what then of the relative permittivity? Separate experiments were conducted to determine this parameter, the result of which was to revise the permittivity from $\varepsilon = 2.55 + 0.0\text{i}$ to $3.00 + 0.004\text{i}$. The results of the revised simulation, for the first-order modes only, are shown in figure 7 (upper panel); there is some improvement, but not as much as we had hoped. Closer examination of the mismatch shows that the higher frequency modes are better matched than the lower frequency modes.

To try and understand the reasons for this frequency-dependent mismatch between experiment and the model it is

instructive to examine the field distributions associated with the different modes for such a structure. We can appeal here to an analytic model for planar metallo-dielectric multilayer structures in the optical regime that has recently been studied by Gadsdon *et al* [64]. The modes of such a structure are very similar to those of the grid structure examined here. The field associated with the highest frequency mode has a sinh distribution in the metallic layers, and there is little field in the metal. On the other hand, the field distribution for the lowest frequency mode follows a cosh distribution in the metal and there is considerable field in the metal. Meshing is thus more important for the lower frequency modes, explaining why the mismatch is poorer for the lower frequency features. The problem is a challenging one, however, since the skin depth for copper at these frequencies is $\sim$1 μm, a factor of 10^4 smaller than the wavelength. Using tighter meshes, ones specified partly by the user and partly within the package, we obtained the revised simulations shown in the lower panel of figure 7—better, but still not right. We are still working to improve the match between experiment and theory here; the results so far indicate that, again, care is needed.

Having examined a number of examples as a means of discussing some of the problems in comparing experimental data with those derived from modelling let me try and summarize. First this is not a static subject. All of the models used here are probably already out of date and new developments continue to emerge [65]. I have laboured the point that the experimental situation is far from being under full control, although that too continues to improve. In comparing experiment with theory we need to recognize from the outset that the systems being considered (experiment and theory) are unlikely to be the same, the key question is in what regard do they need to be the same if we are to achieve a good quantitative understanding of the phenomena involved? There is no escaping that this also demands us to be clear about what constitutes sufficient agreement, an important topic that we have not discussed and one that goes beyond simple (?) validation of modelling techniques.

Many interesting points emerged in discussion at TACONA 2008, three of which I would like to mention here. First, there seems to be an appetite for some kind of challenge as a means to drive improvements in simulation. I mentioned above something simple along these lines by suggesting that a simple gold disc could be made and measured by different techniques in different labs and the results compared. Different numerical techniques could then be harnessed to simulate the same thing and comparisons made. More interesting structures/challenges could also be envisaged, and perhaps some data sets could be made public to facilitate comparison. Second, there was a plea to us all in our role as authors. When we publish papers we should include all of the relevant information needed to reproduce the results, as I have tried to do here. It is not enough to simply state something like 'a full 3D-FDTD simulation was used to produce the results shown in figure X'. When we referee manuscripts we should check that all the required information is included and if it is not we should ask that it be added. The third point is of a different character and goes to the heart of the scientific quest. It is a question that can be stated as follows. Given time and resources we will eventually find efficient simulation methods running on powerful computers to mimic real systems, but to what purpose? (I am talking here of pushing scientific boundaries, not of developing design tools for applications.) Are there data that cannot at present be explained and that better models might help us to resolve? Will it be cheaper and/or faster than carrying out the equivalent experiment? Finally, echoing Wilczek's comment about quantum chromodynamics, what can we learn by simulating that which we do not yet have access to by experiment? Perhaps the answers to these questions will provide extra criteria for classifying the value of simulation methods beyond just speed of computation and accuracy.

Acknowledgments

I am deeply indebted to many of the Electromagnetic Materials Group at Exeter for help in preparing this manuscript. In particular I am indebted to Chris Burrows (in relation to figure 3), James Parsons (in relation to figures 4), Celia Butler (in relation to figures 6 and 7). I would also like to thank Jonathan Wilkinson, Matthew Millyard, Ian Hooper and George Zorinyants who between them provided the data shown in figure 5. I am very grateful to the Royal Society for support through a Merit Award. I would also like to thank all of the participants of TACONA 2008 for such a stimulating meeting, and all those colleagues who helped by making suggestions about interesting topics. The three points at the end of the paper originated with John Sipe, Javier García de Abajo and Kurt Bush, respectively. Particular thanks go to the meeting organizers, Dmitry Chigrin and Kurt Busch. To Dmitry I owe a special thank you, he persuaded me to give a presentation on this theme at TACONA 2008 and I am very grateful; it made me think.

References

[1] Ozbay E 2006 *Science* **311** 189
[2] Zayats A V, Smolyaninov I I and Maradudin A A 2005 *Phys. Rep.* **408** 131
[3] Andrew P and Barnes W L 2000 *Science* **290** 785
[4] Dintinger J *et al* 2005 *Phys. Rev.* B **71** 035424
[5] Dintinger J *et al* 2006 *Adv. Matter.* **18** 1645
[6] Haes A J *et al* 2006 *J. Am. Chem. Soc.* **128** 10905
[7] Porto J A, Martín-Moreno L and García-Vidal F J 2004 *Phys. Rev.* B **70** 081402
[8] Wegener M *et al* 2008 *Opt. Express* **16** 19785
[9] Lawandy N M 2006 *Opt. Lett.* **31** 3650
[10] Durr S *et al* 2008 *Science* **322** 1224
[11] Wilczek F 2008 *Nature* **456** 449
[12] Martín-Moreno L *et al* 2001 *Phys. Rev. Lett.* **86** 1114
[13] Gunnarsson L *et al* 2005 *J. Phys. Chem.* B **109** 1079
[14] Aspnes D E, Kinsbron E and Bacon D D 1980 *Phys. Rev.* B **21** 3290
[15] Dawson P and Boyle M G 2006 *J. Opt. A: Pure Appl. Opt.* **8** S219
[16] Johnson P B and Christy R W 1972 *Phys. Rev.* B **6** 4370
[17] Drachev V P *et al* 2008 *Opt. Express* **16** 1186
[18] Bohren C F and Huffman D R 1983 *Absorption and Scattering of Light by Small Particles* (New York: Wiley)

[19] Waterman P C 1965 *Proc. IEEE* **53** 805
[20] Mishchenko M I, Travis L D and Mackowski D W 1995 *Workshop on Light Scattering by Non-Spherical Particles* (Amsterdam: Pergamon–Elsevier) pp 535–75
[21] Doicu A, Wriedt T and Eremin Y A 2006 *Light Scattering by Systems of Particles: Null-Field Method With Discrete Sources Theory And Programs* (Berlin: Springer)
[22] Purcell E M and Pennypac C R 1973 *Astrophys. J.* **186** 705
[23] Draine B T and Flatau P J 1994 *J. Opt. Soc.* A **11** 1491
[24] Kane Y 1966 *IEEE Trans. Antennas Propag.* **14** 302
[25] Taflove A and Hagness S C 2005 *Computational Electrodynamics: The Finite-Difference Time-Domain Method* (Boston, MA: Artec House Publishers)
[26] Coccioli R *et al* 1996 *IEEE Antennas Propag. Mag.* **38** 34
[27] Myroshnychenko V *et al* 2008 *Adv. Matter.* **20** 4288
[28] Kretschmann E and Raether H 1968 *Z. Naturf.* A **23** 2135
[29] Herminghaus S, Klopfleisch M and Schmidt H J 1994 *Opt. Lett.* **19** 293
[30] Barnes W L and Sambles J R 1987 *Surf. Sci.* **183** 189
[31] Lynch D W and Huttner W R 1985 *Handbook of Optical Constants of Solids* ed E D Palik (New York: Academic) pp 275–367
[32] Innes R A and Sambles J R 1987 *J. Phys. F: Met. Phys.* **17** 277
[33] Smith L H *et al* 2008 *J. Mod. Opt.* **55** 2929
[34] Wedge S and Barnes W L 2004 *Opt. Express* **12** 3673
[35] Nash D J and Sambles J R 1999 *J. Mod. Opt.* **46** 1793
[36] Ebbesen T W, Genet C and Bozhevolnyi S I 2008 *Phys. Today* **61** 44
[37] Charbonneau R *et al* 2005 *Opt. Express* **13** 977
[38] Nielsen R B *et al* 2008 *Opt. Lett.* **33** 2800
[39] Berini P *et al* 2005 *J. Appl. Phys.* **98** 043109
[40] Gersten J and Nitzan A 1980 *J. Chem. Phys.* **73** 3023
[41] Hartschuh A *et al* 2004 *Phil. Trans. R. Soc.* A **392** 807
[42] Nehl C L, Liao H and Hafner J H 2006 *Nano Lett.* **6** 683
[43] García-Vidal F J and Pendry J B 1996 *Phys. Rev. Lett.* **77** 1163
[44] Camden J P *et al* 2008 *J. Am. Chem. Soc.* **130** 12616
[45] Kim S *et al* 2008 *Nature* **453** 757
[46] Muhlschlegel P *et al* 2005 *Science* **308** 1607
[47] Ghenuche P *et al* 2008 *Phys. Rev. Lett.* **101** 116805
[48] Li S and Schatz G C 2008 *MRS Spring Mtg (Boston)* pp V01–8
[49] Shalaev V M 1996 *Phys. Rep.* **272** 61
[50] Chan G H *et al* 2007 *Nano Lett.* **7** 1947
[51] Liedberg B, Nylander C and Lundstrom I 1983 *Sensors Actuators* **4** 299
[52] Willets K A and Van Duyne R P 2007 *Ann. Rev. Phys. Chem.* **58** 267
[53] Raschke G *et al* 2003 *Nano Lett.* **3** 935
[54] McFarland A D and Van Duyne R P 2003 *Nano Lett.* **3** 1057
[55] Stewart C E, Hooper I R and Sambles J R 2008 *J. Phys. D: Appl. Phys.* **41** 7
[56] Voros J 2004 *Biophys. J.* **87** 553
[57] Sipe J E 1980 *Phys. Rev.* B **22** 1589
[58] de Abajo F J G 2008 *J. Phys. Chem.* C **112** 17983
[59] Palpant B *et al* 1998 *Phys. Rev.* B **57** 1963
[60] Zhao L L, Jensen L and Schatz G C 2006 *J. Am. Chem. Soc.* **128** 2911
[61] Pecharroman C *et al* 2008 *Phys. Rev.* B **77** 035418
[62] Butler C A M *et al* 2009 submitted
[63] Soukoulis C M, Linden S and Wegener M 2007 *Science* **315** 47
[64] Gadsdon M R, Parsons J and Sambles J R 2009 *J. Opt. Soc. Am.* B **26** 734
[65] Niegemann J, Tkeshelashviii L and Busch K 2006 *2nd Workshop on Numerical Methods for Optical Nanostructures* (Zurich: Amer Scientific Publishers) pp 627–34