

Rodcast

Quantitative Morphology–Property  
Correlations in Gold Plasmonic  
Nanoparticles

Ph. D. Thesis, University of Amsterdam, 3 June 2026

*Rodcast*

*Quantitative Morphology-Property Correlations in Gold Plasmonic Nanoparticles*

Mees Dieperink

Cover design: Stray light in the backfocal plane of the objective (front) and intermediate plane (back) of the first version of our home-built quasi-darkfield microscope.

The work described in this thesis was performed between September 2021 and March 2026 at AMOLF, Science Park 104, 1098XG Amsterdam, The Netherlands, an institute of the Dutch Research Council (NWO).

ISBN: 9789465376066

A digital version of this thesis is available at <https://ir.amolf.nl>

Printed by Ridderprint

Rodcast  
Quantitative Morphology–Property  
Correlations in Gold Plasmonic  
Nanoparticles

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor  
aan de Universiteit van Amsterdam  
op gezag van de Rector Magnificus  
prof. dr. ir. P. P. C. C. Verbeek  
ten overstaan van een door het College voor Promoties ingestelde commissie,  
in het openbaar te verdedigen in de Agnietenkapel  
op woensdag 3 juni 2026, te 10.00 uur

door Mees Dieperink  
geboren te Amsterdam

*Promotiecommissie*

<i>Promotor:</i>	prof. dr. A. Polman	Universiteit van Amsterdam
<i>Copromotor:</i>	dr. W. Albrecht	AMOLF
<i>Overige leden:</i>	prof. dr. P.C.M. Planken	Universiteit van Amsterdam
	dr. J. van de Groep	Universiteit van Amsterdam
	prof. dr. E. Alarcon Llado	Universiteit van Amsterdam
	prof. dr. S. Link	University of Illinois
	prof. dr. P. Zijlstra	Technische Universiteit Eindhoven

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

# Contents

<b>1</b>	<b>Broadcast – Introduction &amp; Motivation</b>	<b>1</b>
1.1	Technology for society . . . . .	2
1.2	Plasmonic nanoparticles . . . . .	4
1.2.1	Plasmons . . . . .	4
1.2.2	Plasmons in metals . . . . .	5
1.2.3	Plasmonic nanoparticles for technology . . . . .	7
1.2.4	How to see single plasmonic nanoparticles . . . . .	10
1.3	Thesis outline . . . . .	15
<b>2</b>	<b>Cast Light – Quantitative Electromagnetic Simulations</b>	<b>17</b>
2.1	Introduction . . . . .	18
2.2	Gold as a plasmonic material . . . . .	18
2.3	Analytical modelling of the dielectric function of gold . . . . .	19
2.3.1	Intraband transitions: the Drude–Sommerfeld model . . . . .	19
2.3.2	Interband transitions: Lorentz oscillators . . . . .	20
2.3.3	Measured dielectric functions . . . . .	21
2.4	From bulk to nanoparticle . . . . .	22
2.5	From quasistatic to retarded electrodynamics: when does the Rayleigh approximation fail? . . . . .	23
2.6	Considerations for electromagnetic simulation methods . . . . .	25
2.6.1	BEM meshing for anisotropic nanoparticles . . . . .	27
2.6.2	Material and environment choices: dielectric function and substrates . . . . .	29
2.7	Conclusions and outlook . . . . .	30
2.A	Appendix . . . . .	32
2.A.1	Electromagnetic simulation methods . . . . .	33
<b>3</b>	<b>Screen-cast – From Electron Tomography to Electromagnetic Simulations</b>	<b>37</b>
3.1	Introduction . . . . .	38
3.2	Limitations of 2D imaging for morphology extraction . . . . .	40
3.3	Simulating realistic tomography data for benchmarking . . . . .	42
3.4	Comparing tomographic reconstruction algorithms . . . . .	44
3.5	Segmentation strategies and threshold selection . . . . .	46
3.6	Generating surface meshes from voxel reconstructions . . . . .	48
3.7	Mesh simplification and its effect on shape accuracy . . . . .	48

3.8	Propagation of morphological errors into optical simulations . . . . .	49
3.9	Recommended workflow and best practices . . . . .	51
3.10	Conclusion . . . . .	52
3.A	Appendix . . . . .	54
3.A.1	Experimental methods . . . . .	54
3.A.1.1	Sample preparation . . . . .	54
3.A.1.2	Electron tomography experiment . . . . .	54
3.A.2	From electron tomography to mesh . . . . .	54
3.A.2.1	Pre-processing electron tomography data . . . . .	54
3.A.2.2	Electron tomography reconstructions . . . . .	55
3.A.2.3	Segmentation of reconstructions . . . . .	55
3.A.2.4	Smoothing before and after reconstruction . . . . .	58
3.A.2.5	Fitting algorithms for shape parametrization . . . . .	60
3.A.2.6	Meshing segmented reconstructions . . . . .	62
3.A.2.7	Simplifying meshed reconstructions . . . . .	64
3.A.2.8	Initial correlation of optical measurement to 2D HAADF-STEM data . . . . .	66
<b>4</b>	<b>Cast Shadow – Quasi-Darkfield Microspectroscopy</b>	<b>73</b>
4.1	Introduction . . . . .	74
4.2	Quasi-darkfield microspectroscopy . . . . .	74
4.2.1	Introduction and motivation . . . . .	74
4.2.2	Setup overview . . . . .	75
4.2.3	Angular filtering and coupled design parameters . . . . .	76
4.2.4	Angular collection efficiency and optimal reflector size . . . . .	78
4.2.5	Nanoparticle radiation pattern and collection efficiency . . . . .	81
4.2.6	Pinhole design and stray-light suppression . . . . .	83
4.3	Considerations for spectroscopy . . . . .	84
4.3.1	Instrument response function . . . . .	85
4.3.2	Signal quality and alignment requirements . . . . .	87
4.3.3	Refractive-index matching and background suppression . . . . .	89
4.3.4	Wavelength-focusing . . . . .	91
4.4	Conclusion . . . . .	95
4.A	Appendix . . . . .	97
<b>5</b>	<b>Typecast – Quantitative Optical–Morphological Correlation: Ex- tracting Local Dielectric Environments</b>	<b>99</b>
5.1	Introduction . . . . .	100

5.2	Optical measurements and spectral analysis . . . . .	100
5.2.1	Spectral model and fit strategy . . . . .	102
5.3	Morphology extraction and environment determination . . . . .	102
5.3.1	Determining the effective environment for a single particle . . .	104
5.4	Sensitivity and reproducibility of the dielectric-environment extraction . . . . .	105
5.5	Linewidths and amplitudes: observations and model limitations . . . .	108
5.6	Conclusions and outlook . . . . .	110
<b>6</b>	<b>Melt-cast – Reshaping Dynamics of Gold Nanorods</b>	<b>111</b>
6.1	Introduction . . . . .	112
6.2	Experimental setup and imaging scheme . . . . .	113
6.3	Initial morphology and pre-illumination characterization . . . . .	114
6.4	Quantification of reshaping through projection difference mapping . .	115
6.5	Pulse-resolved dynamics and automated tracking . . . . .	117
6.6	Heat-diffusion modelling and temperature analysis . . . . .	119
6.7	Quantification of reshaping through tip-difference mapping . . . . .	121
6.8	Mass redistribution and metastable locking in dimers . . . . .	123
6.9	Cluster-mediated reshaping and long-range correlations . . . . .	123
6.10	Mechanistic interpretation: diffusion, carbon Shell, and kinetics . . . .	126
6.11	Conclusion . . . . .	128
6.A	Appendix . . . . .	130
6.A.1	Segmentation . . . . .	130
6.A.1.1	Manual . . . . .	130
6.A.1.2	Automatic . . . . .	131
6.A.2	Simulations . . . . .	135
	<b>Bibliography</b>	<b>137</b>
	<b>List of publications and author contributions</b>	<b>151</b>
	<b>Rodcast – Summary</b>	<b>153</b>
	<b>Samenvatting</b>	<b>157</b>
	<b>The Cast – Acknowledgements</b>	<b>161</b>
	<b>About the author</b>	<b>167</b>



# Chapter 1

---

## Broadcast – Introduction & Motivation

*Broadcast – to spread information to a lot of people*

**Abstract:** This chapter situates the thesis within the broader technological and societal context that increasingly shapes expectations for materials research. It highlights how long-term challenges, particularly the climate crisis, sustainable energy, and low-carbon industrial processes, depend critically on advances in nanoscale materials. Within this landscape, plasmonic nanoparticles emerge as a promising class for light management, photocatalysis, and energy-relevant surface processes.

The chapter introduces the physical origins of plasmons, explains how localized surface plasmon resonances (LSPRs) arise in metal nanoparticles, and outlines the material properties that make gold an effective visible-range plasmonic system. It then connects plasmonic dissipation pathways to technological functions in photovoltaics and sustainable catalysis, emphasizing how radiative and non-radiative damping govern scattering, hot-carrier generation, and photothermal effects. Finally, the chapter motivates the need for *single-particle* and *correlative* electron-microscopy techniques to establish quantitative structure–property relationships. Because plasmonic behaviour is highly sensitive to nanoparticle geometry, single-particle spectroscopy combined with detailed morphological characterization provides the most reliable route toward understanding and modelling plasmonic resonances. This context forms the foundation for the experimental and analytical approaches developed in the subsequent chapters.

Section 1.2.4 is inspired by M. Dieperink, F. Scalerandi, and W. Albrecht, “Correlating structure, morphology and properties of metal nanostructures by combining single-particle optical spectroscopy and electron microscopy,” *Nanoscale*, vol. 14, pp. 7460–7472, 20 2022, doi: 10.1039/D1NR08130F.

## 1.1 Technology for society

Scientific research is often driven by curiosity and the desire to understand fundamental principles. Yet, across many disciplines there is a growing expectation to highlight how even highly specialized work connects to broader societal or technological developments. For researchers in fields such as nanoscience or photonics, establishing this link is not always straightforward, because our training rarely covers adjacent areas such as energy policy, economics, or technology assessment. Nevertheless, reflecting on these connections can clarify why certain scientific questions matter beyond their disciplinary scope.

In this introduction, I therefore ask a modest guiding question: *Which long-term societal challenges are consistently identified in international assessments, and how do emerging technologies related to materials-based research contribute to addressing them?* This thesis does not attempt to solve these challenges, but this context provides a meaningful setting for understanding how materials-based research, including plasmonics, can support future technological developments.

A natural starting point is the set of global risk assessments published annually by the World Economic Forum (WEF) [1], [2], [3]. Based on surveys of experts from academia, industry, government, and civil society, these reports consistently identify long-term risks dominated by the climate crisis, including failures of climate mitigation and adaptation, extreme weather, and biodiversity loss. These long-term risks have remained remarkably stable over recent years and align with scientific assessments by the Intergovernmental Panel on Climate Change (IPCC) [4].

Complementary perspectives appear in risk or strategy reports by the Big Four global consulting firms: Deloitte [5], [6], EY [7], [8], KPMG [9], [10], and PwC [11], [12]. These reports emphasize strategic priorities around technological competitiveness, regulatory compliance, and corporate sustainability. While the WEF frames risks in broad societal terms, the Big Four focus on strategic innovation. A document that bridges these two perspectives is the 2024 Draghi report prepared for the European Commission [13]. This report argues that Europe must significantly invest in strategic technologies, including clean-energy systems such as photovoltaics, batteries, and hydrogen technologies, as well as advanced materials such as semiconductors, nanomaterials for catalysis, and materials for high-performance computing. Moreover, the report stresses that these technologies are intertwined, because energy, materials, and digital infrastructure depend on shared innovation capacities.

A clear illustration of how such technologies can simultaneously advance sustainabil-

ity and economic growth is the rapid global expansion of solar photovoltaics. In 2025, more solar capacity was installed worldwide than coal, primarily because photovoltaics have become the cheapest source of new electricity generation [14]. Similar dynamics appear in technologies where catalysis plays a central role, for example carbon capture and storage, power-to-fuels conversion, and carbon management pathways [15], [16]. These technologies are essential for long-term climate strategies and rely on continued advances in materials science and nanotechnology.

Table 1.1 summarizes the United Nations’ Sustainable Development Goals (SDGs) most relevant to the technological domains discussed above. The selected SDGs relate directly to clean energy, sustainable industrial processes, and climate action, and therefore connect most clearly to the societal context of this thesis. The example technologies are grouped into two categories, namely solar energy and sustainable catalysis, which will return later in the introduction when discussing the relevance of plasmonic nanoparticles.

<b>SDG #</b>	<b>SDG name</b>	<b>Example technology</b>
7	Affordable and clean energy	Solar energy
8	Decent work and economic growth	Solar energy
9	Industry, innovation, and infrastructure	Sustainable catalysis
12	Responsible consumption and production	Sustainable catalysis
13	Climate action	Sustainable catalysis

Table 1.1: United Nations’ Sustainable Development Goals (SDGs) most relevant for this thesis, grouped by example technologies in solar energy and sustainable catalysis.

Taken together, reports by the WEF, the Big Four, the IPCC, and international energy agencies such as the IEA and IRENA, alongside the SDG framework, emphasise that advanced materials and nanoscale technologies are central to future sustainable energy systems. They are required to improve photovoltaic efficiency, enable new catalytic pathways, enhance energy storage, and support low-carbon industrial processes.

Addressing these challenges ultimately requires progress at the level of fundamental materials research. Plasmonic nanoparticles, metallic nanostructures that can concentrate and manipulate light at the nanoscale, represent one such frontier, with direct relevance to solar-energy harvesting, photocatalysis, and sensing. Their ability to convert and redirect light with high efficiency makes them promising candidates for the energy and sustainability applications identified across the assessments above. The physical principles underlying these properties, and the specific case of gold nanoparticles, are introduced in the following section.

## 1.2 Plasmonic nanoparticles

Plasmonic nanoparticles are an example of advanced materials that hold promise for sustainable technologies, including solar-energy harvesting and sustainable catalysis. Before discussing these applications, this section introduces the physical origin of plasmons, explains how plasmonic resonances arise in nanoparticles, and outlines why materials such as gold are particularly well suited for strong light–matter interactions at the nanoscale.

### 1.2.1 Plasmons

In conductive materials, the valence electrons are not bound to individual atoms but can move collectively in response to an external electric field. When these electrons oscillate as a group, the motion is known as a *plasma oscillation* [17]. In bulk materials, these collective electron oscillations occur at a characteristic frequency set by the free-electron density. When an electromagnetic wave drives such a collective mode, the resulting optical excitation is called a *plasmonic resonance*.

The ability of a material to support strong plasmonic resonances depends on three key quantities: its free-electron density, its dielectric function, and its optical losses. The free-electron density sets the approximate plasmon frequency. The dielectric function captures both the strength of electromagnetic coupling and the dissipative losses within the material. Low optical losses and a negative real part of the dielectric function are essential requirements for bright and well-defined plasmonic resonances.

Table 1.2 summarizes several classes of materials that support plasmonic excitations, grouped according to their conductivity and free-electron density. These two quantities provide an intuitive first indication of the spectral range in which a material is expected to exhibit plasmonic behaviour. Metals such as Ag, Cu, and Au combine high conductivities with very large free-electron densities, which makes them excellent plasmonic materials in the visible and near-infrared ranges. Other classes, such as doped semiconductors or transparent conducting oxides, support plasmonic resonances at lower frequencies, typically in the mid-infrared or terahertz regimes.

The comparison in Table 1.2 shows that metals generally offer the strongest performance for visible-light plasmonics because they combine high free-electron densities with relatively low losses. This results in strong electromagnetic confinement at dimensions well below the wavelength of light when structured into nanoparticles, which is essential for applications such as field-enhanced spectroscopy, photocatalysis, and light harvesting. Later sections will discuss how size, shape, and dielectric

Material class	Conductivity (S/m)	Free-electron density (cm <sup>-3</sup> )	Typical plasmonic spectral range
Metals (Ag, Cu, Au)	10 <sup>7</sup> –10 <sup>8</sup>	10 <sup>22</sup>	Visible / Near-IR
Transition-metal nitrides	10 <sup>6</sup> –10 <sup>7</sup>	10 <sup>22</sup>	Near-IR / Mid-IR
Graphene	10 <sup>4</sup> –10 <sup>6</sup>	10 <sup>13</sup>	Terahertz / Mid-IR
MXenes	10 <sup>4</sup> –10 <sup>6</sup>	10 <sup>21</sup>	Near-IR / Mid-IR
Transparent conducting oxides	10 <sup>3</sup> –10 <sup>4</sup>	10 <sup>20</sup>	Mid-IR
Topological insulators	10 <sup>2</sup> –10 <sup>4</sup>	10 <sup>19</sup>	Mid-IR / Far-IR
Doped semiconductors	10 <sup>2</sup> –10 <sup>4</sup>	10 <sup>18</sup>	Mid-IR / THz

Table 1.2: Overview of plasmonic material classes grouped by conductivity and free-electron density, which together provide a first indication of the spectral region in which these materials support plasmonic resonances.

environment determine the optical response of gold nanoparticles.

### 1.2.2 Plasmons in metals

The interaction of light with metals depends on the characteristic length scale of the system, which leads to distinct plasmonic responses in bulk materials, at interfaces, and in nanoscale structures. This subsection introduces these regimes to clarify how localized surface plasmon resonances arise in nanoparticles, which form the central topic of this thesis.

When an electromagnetic wave impinges on a bulk metal, the free electrons collectively oscillate in response to the driving electric field. Their motion induces surface charges and currents that generate an internal electric field opposing the incident one, resulting in strong screening. The restoring force for this collective displacement is provided by the positively charged ion cores, and together with electron inertia it gives rise to a natural eigenfrequency known as the bulk plasma frequency. Because the plasma frequency of noble metals lies in or above the ultraviolet, bulk metals reflect visible and infrared light strongly. The associated skin depth causes the electric field to decay exponentially inside the metal, producing the familiar reflective appearance of macroscopic gold objects (Figure 1.1A).

Above the plasma frequency, the electrons can no longer effectively screen the external field and the real part of the metal’s permittivity becomes positive. In this regime, electromagnetic waves can propagate more easily into the metal. The corresponding collective longitudinal excitation of the electron gas is known as a *bulk plasmon*. Since freely propagating light is transverse, bulk plasmons cannot be




Figure 1.1: Illustration of plasmonic behavior across length scales in gold. (A) Bulk gold structures, where the optical response is dominated by reflection and skin-depth-limited field penetration causing a homogeneous electric field at the surface. (B) Thin gold films support propagating surface plasmon polaritons (SPPs), shown here in a 50 nm thick coating (B.a). (C) Localized surface plasmon resonances (LSPRs) in nanoscale gold structures: the characteristic red color of 50 nm colloidal gold nanoparticles used in rapid antigen tests (C.a) and an electromagnetic wave polarizing a particle at resonance (C.b). (A.b) and (B.b) are adapted from [17].

directly excited by light in a uniform metal. Instead, they are typically accessed using excitation sources with additional momentum, such as fast electrons in electron microscopes.

At a flat metal–dielectric interface, another type of collective excitation can arise: a *surface plasmon polariton* (SPP). This mode is a surface-bound electromagnetic wave coupled to a charge-density oscillation at the interface. Because SPPs carry more in-plane momentum than freely propagating photons of the same frequency, they cannot be excited directly by light incident from air. In practice, they are launched using coupling schemes that increase the in-plane photon momentum, for example through a high-index prism or a grating. Thin gold films, such as the 50 nm coating shown in Figure 1.1B.a, support these propagating modes (Figure 1.1B.b), which form the basis of widely used surface-sensitive techniques. Although SPPs illustrate how plasmonic behavior emerges at planar interfaces, they are not the focus of this thesis and are mentioned here primarily to contrast their excitation mechanism with that of localized resonances in nanoparticles.

When a metal is structured into a nanoparticle with dimensions smaller than the wavelength of light, the situation changes fundamentally. The broken translational

symmetry removes the momentum constraint that prevents free-space excitation of propagating plasmons. As a result, incident light can directly drive a non-propagating collective mode known as a *localized surface plasmon resonance* (LSPR). In this case, the electron oscillation is confined to the nanoparticle and determined by its geometry rather than by a dispersion relation. The resonance frequency depends sensitively on the particle's size, shape, and dielectric environment, as well as on the frequency-dependent dielectric function of the metal, which governs both radiative and non-radiative damping. Compared to bulk gold, the optical properties can change drastically; for example, 50 nm gold nanoparticles exhibit a strong LSPR in the visible range that gives rise to the characteristic red color exploited in rapid antigen tests (Figure 1.1C.a). LSPRs give rise to strong local field enhancement and characteristic absorption and scattering peaks, which are visible in colloidal gold and in nanoscale structures. These size- and shape-dependent resonances are the main subject of this thesis and will be analyzed quantitatively in the following chapters.

### 1.2.3 Plasmonic nanoparticles for technology

Plasmonics in general, and plasmonic nanoparticles (PNPs) in particular, are being explored in several technology areas linked to the global challenges outlined in Section 1.1. Their technological relevance follows directly from how an LSPR dissipates energy. A localized plasmon is a driven resonance of a finite metal nanoparticle and it loses energy through two main pathways: *radiative damping*, which results in far-field scattering, and *non-radiative damping*, which corresponds to absorption in the metal and produces electronic excitations (hot carriers), followed by ultrafast thermalization and local heating. Depending on materials and geometry, these dissipation channels can be harnessed for light management, carrier generation, and heat-driven processes [18], [19].

These dissipation pathways underpin distinct functionalities in photovoltaics and photocatalysis. Radiative losses favour applications that rely on far-field scattering and light redirection, such as optical path-length enhancement in thin-film photovoltaics and surface-enhanced Raman scattering (SERS). Non-radiative losses enable hot-carrier generation and photothermal effects relevant for catalysis and photoelectrochemistry. The balance between these pathways is not universal; it depends on the dielectric function of the metal and is strongly influenced by nanoparticle size, shape, and morphology. Understanding and controlling this structure–damping interplay is essential for connecting fundamental plasmonic behaviour to practical device concepts, and it motivates the structure–property analyses developed in this thesis.

## Solar energy applications (light management and niche pathways)

Plasmonic nanoparticles offer promising light-management capabilities for photovoltaic applications, though their integration into mainstream silicon, perovskite, or CdTe solar cells remains limited because parasitic losses in the metal can compete with or outweigh potential optical gains. In conventional photovoltaics (PVs), PNPs are not the primary absorber. Early investigations explored whether PNPs could improve absorption near the band edge, but in most standard PV architectures the parasitic losses in the metal reduce net benefit. Nevertheless, PNPs remain of interest in scenarios where their unique scattering or field-confinement properties can be advantageous, particularly in ultrathin absorber layers or in niche concepts.

Two light-management mechanisms are commonly discussed. First, suitably designed PNPs can enhance absorption through *far-field scattering*. By redirecting incident sunlight into larger in-material angles or guided modes, they increase the effective optical path length and improve weak absorption near the band edge. Second, PNPs can enhance absorption in the semiconducting PV material through *near-field concentration*, where the strong localized field increases the excitation rate of the nearby semiconductor, and near-field-mediated coupling (e.g. resonant energy transfer) can be beneficial close to the band edge [18], [20]. In practice, metal losses often compete with these desirable effects, which helps explain why net gains are limited in conventional devices.

Beyond light management, *hot-carrier injection* from the metal into a semiconductor has been investigated, but in conventional PVs it typically plays a niche role because ultrafast thermalization in the metal and modest injection yields limit gains in device efficiency [18], [21]. Local photothermal heating from PNPs, while generally detrimental to standard PV operation through increased recombination and reduced open-circuit voltage [22], represents one of the more promising and realistic plasmonic contributions to photovoltaics when exploited in specialized concepts such as thermophotovoltaics [23]. For comprehensive overviews of hot-carrier generation and transport in plasmonic systems, see [18], [21].

## Sustainable catalysis

The role of PNPs in catalysis is broader because plasmonic dissipation pathways, both radiative and non-radiative, can contribute more directly to reaction enhancement, without the same competing loss mechanisms that limit their impact in conventional photovoltaics.

In the first case, the *PNP itself acts as the (photo)catalyst*. Here the nanoparticle surface participates directly in the chemical reaction, and plasmonic excitation can influence surface chemistry through several mechanisms. Non-radiative decay generates energetic carriers that may transfer to adsorbates or across metal–semiconductor interfaces, enabling selective activation pathways that are not accessible thermally. In addition, photothermal heating from carrier thermalization increases the local temperature and can accelerate reaction kinetics [19]. In pulsed or strongly driven regimes, periodic modulation of surface temperatures or hot-carrier populations can transiently shift adsorption energies and reaction barriers, potentially giving rise to resonant or dynamic catalysis in which time-averaged rates exceed those achievable under steady-state conditions, an exciting prospect that may open fundamentally new pathways in plasmonic photocatalysis [19].

In the second case, the PNP functions primarily as an *optical antenna*. Here the particle is not the active catalyst but rather harvests and concentrates light for a nearby photoactive catalyst or electrode. Near-field concentration increases the local excitation rate in the photoactive component without requiring direct electronic contact. In contrast, direct hot-carrier transfer to the catalyst or to adsorbates generally requires electronic coupling (either direct contact or tunnelling-range proximity) and appropriate band or level alignment [19], [24], [25]. Photothermal heating can again modify kinetics, depending on the catalytic mechanism and thermal management of the system [19]. Representative discussions of mechanisms and design considerations include [19], [24], [25].

Although the mechanistic distinctions above are conceptually useful, practical plasmonic catalytic systems often exhibit several intertwined pathways simultaneously. Which mechanism dominates depends sensitively on the nanoparticle morphology, which controls the mode structure, the distribution of local fields, and the balance between radiative and non-radiative damping.

Crucially, this structure–damping interplay governs PNP functionality across all the technology areas discussed above. Radiative scattering, near-field localization, hot-carrier generation, and photothermal effects occur on ultrafast timescales and can coexist in a single nanoparticle. Disentangling their contributions is challenging because ensemble measurements average over particles of different shapes and sizes. This motivates the use of *single-particle* and *correlative* characterization strategies, where the optical response of an individual nanoparticle is directly linked to its precise morphology on the same particle. Such approaches are essential for quantifying how geometry influences plasmonic response, identifying dominant dissipation pathways, and establishing reliable structure–property relationships. These methods form a

central methodological pillar of this thesis and provide the basis for the experimental overview introduced in the next section.

### 1.2.4 How to see single plasmonic nanoparticles

Plasmonic nanoparticles derive their technological relevance from the way their LSPRs dissipate energy through radiative and non-radiative channels. As discussed above, these damping pathways determine whether a nanoparticle primarily scatters light, generates hot carriers, or produces local heating, and therefore which photovoltaic or catalytic functions it can support. Because all of these effects depend sensitively on particle morphology, establishing reliable structure–property relationships requires techniques that can probe the optical response of *individual* nanoparticles rather than large ensembles.

Ensemble measurements average over entire nanoparticle populations that inevitably exhibit variations in size, shape, crystallinity, and environmental conditions. These variations broaden resonances, obscure mode-specific behaviour, and often conceal the differences between radiative and non-radiative contributions. This limitation is particularly pronounced for anisotropic nanoparticles, where shape-dependent plasmon modes and orientation effects further broaden and obscure spectral features in ensemble measurements. Single-particle measurements overcome these limitations by isolating the optical signature of a *single well-defined object*, allowing direct comparison with its precise morphology [26], [27], [28].

#### Optical single-particle techniques

Optical microscopy is a natural starting point for single-particle studies because it enables broadband, minimally invasive measurements under ambient conditions. Although optical microscopy is diffraction-limited, individual nanoparticles can still be measured provided they are sufficiently separated on the substrate so that their optical responses do not spatially overlap. Under these conditions the scattered or absorbed light from a single nanoparticle can be isolated and analysed spectroscopically. Two principal optical configurations are commonly used for such measurements: bright-field (BF) extinction microscopy and dark-field (DF) scattering microscopy.

*Bright-field (BF) extinction microscopy* detects nanoparticles through attenuation of the transmitted illumination. Figure 1.2a illustrates the reflection bright-field configuration used for extinction measurements, in which the incident beam and the collected light share the same optical path. The measured signal corresponds to the

optical extinction, defined as the combined effect of absorption and scattering by the nanoparticle. In practice, however, detecting extinction from a single nanoparticle is challenging because the particle-induced change in the transmitted intensity is extremely small compared to the much larger background signal from the illumination and substrate. For small gold nanoparticles in particular, the resulting fractional intensity change often falls within the noise and background fluctuations inherent to the bright-field geometry, making reliable single-particle spectroscopy difficult.




Figure 1.2: Schematic illustration of single-particle configurations: (a) bright-field (BF) reflectance for extinction; (b) dark-field (DF) transmittance with an annular condenser and separate collection objective; (c) DF reflectance using a single objective; (d) quasi-dark-field (QDF) reflectance used in this thesis. Colours indicate illumination (blue in a,d; yellow in b,c) and collected signal (orange). In (b), forward-scattered light is collected; in (c) and (d), backward-scattered light and reflections are managed by geometry and spatial filtering. Panels (b) and (c) adapted from [29].

*Dark-field (DF) scattering microscopy* overcomes the BF limitation by suppressing the directly transmitted illumination so that only light scattered by the nanoparticle is collected. This is typically achieved using annular or highly oblique illumination, which ensures that the incident beam does not enter the objective while scattered light from the nanoparticle does. As a result, nanoparticles appear as bright, wavelength-dependent spots on a dark background. For plasmonic nanoparticles, whose resonant modes can produce large scattering cross-sections, particularly for particles comparable to or larger than several tens of nanometers, this geometry provides substantially higher signal-to-background ratios than BF measurements and forms the basis for most room-temperature single-particle plasmonic spectroscopy.

Spectra are obtained by dispersing the scattered light spectrally (e.g. a spectrometer or wavelength-resolved detection), and hyperspectral imaging approaches can additionally record spatially resolved spectra. Figures 1.2b–d illustrate several DF implementations. In conventional configurations, a dedicated DF condenser provides annular illumination while the scattered light is collected with a separate objective (Fig. 1.2b) or with the same objective (Fig. 1.2c). In the configuration used in this thesis (Fig. 1.2d), the illumination is not fully annular; instead, the reflected illumination is suppressed after the objective using spatial filtering so that primarily scattered light reaches the detector. Because this configuration does not reject the illumination before the objective, it is referred to here as a *quasi-dark-field* geometry. The implementation and performance of this approach are discussed in detail in Chapter 4.

Several practical considerations influence spectra obtained from DF measurements. High numerical apertures illuminate and collect over a wide range of angles and polarisations, so the measured response of anisotropic nanoparticles represents an average over multiple excitation conditions [30], [31]. In addition, substrate reflections, interference effects, and background scattering can affect the measured spectra, while the diffraction limit constrains spatial resolution and the measured scattering amplitude depends on the numerical aperture and collection geometry. Despite these factors, optical spectroscopy remains a powerful method for studying plasmonic nanoparticles because it directly probes the optical response of individual particles under realistic conditions and provides access to the balance between radiative and non-radiative damping [19], [20], [32], [33], [34], [35], [36], [37], [38], [39].

## Transmission electron microscopy

Electron microscopy provides structural information with spatial resolutions far beyond optical techniques because high-energy electrons have wavelengths orders of magnitude shorter than visible light. As a result, electron microscopy can resolve nanoparticle morphology, facet structure, and crystallographic features at nanometer to atomic length scales. Similar to optical microscopy, electron imaging can be performed in bright-field (BF) and dark-field (DF) configurations, where contrast arises from different scattering processes of the electron beam.

In conventional *transmission electron microscopy* (*TEM*), a nearly parallel electron beam passes through the sample and the transmitted electrons form a projection image. In the bright-field configuration, the directly transmitted electrons form the image, producing contrast from mass–thickness variations as well as diffraction effects associated with the crystal structure of the material. TEM images can therefore contain

crystallographic information, but the resulting contrast is not always straightforward to interpret quantitatively. In *scanning transmission electron microscopy (STEM)*, the electron beam is instead focused into a probe that scans across the sample. When detected with a high-angle annular dark-field (HAADF) detector, corresponding to a dark-field imaging mode, the image is formed primarily from electrons scattered to high angles. Because this scattering is largely incoherent, diffraction contrast is strongly reduced compared to conventional TEM, and the image intensity scales more directly with sample thickness and atomic number (Z-contrast). This generally makes HAADF-STEM images easier to interpret in terms of morphology and composition. Figure 1.3a,b illustrate the difference between STEM and TEM imaging geometries, while Figure 1.3c shows the annular detector typically used in STEM to collect high-angle scattered electrons, producing HAADF contrast.

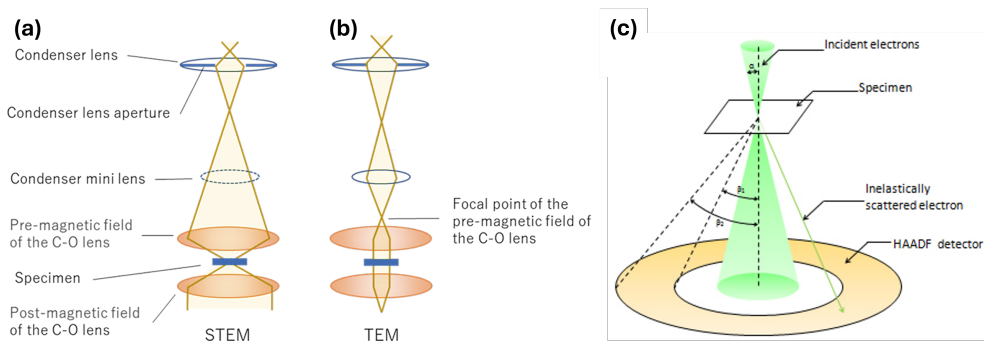


Figure 1.3: (a) Ray path of the STEM mode using the C–O lens. (b) Ray path of the TEM mode, where a parallel beam is produced on the specimen using the condenser-mini lens. (c) Relationship between the convergence semi-angle of the incident electron beam and acceptance semi-angles of the detector for HAADF-STEM. Typical inner and outer semi-angles of the detector are respectively  $\beta_1 \approx 50$  mrad and  $\beta_2 \approx 200$  mrad, detecting predominantly high-angle (largely incoherent) scattering. A typical convergence semi-angle is  $\alpha \approx 25$  mrad for a 200 kV Cs-corrected TEM. Figure adapted from jeol.com.

Despite their high spatial resolution, both TEM and STEM images represent two-dimensional projections of the three-dimensional nanoparticle structure. To recover the full morphology, *electron tomography* can be used, in which a tilt series of TEM or STEM images is reconstructed into a three-dimensional volume. In practice, STEM-based tomography is often preferred for crystalline materials because the reduced diffraction contrast leads to projection images that more closely satisfy the assumptions required for tomographic reconstruction. Tomography therefore enables detailed structural characterization of anisotropic or irregular nanoparticles, although

practical limitations such as incomplete tilt ranges (the “missing wedge”) and increased electron dose place constraints on the achievable reconstructions.

Beyond structural imaging, electron spectroscopies such as *electron energy-loss spectroscopy* (EELS) and *cathodoluminescence* (CL) can probe plasmonic excitations directly. These techniques allow plasmonic eigenmodes to be mapped with nanometer-scale spatial resolution, even in complex geometries [40], [41], [42] and sub-nanometer gaps [43], and provide insight into nanoscale energy-transfer pathways [44], [45]. Comprehensive overviews are available in Refs. [46], [47], [48], [49], [50], [51]. However, because these measurements are performed inside an electron microscope, they require vacuum conditions and rely on excitation by fast electrons rather than optical fields. As a result, beam-induced modifications can occur and the excitation conditions differ from those encountered in optical spectroscopy. For understanding plasmonic nanoparticles under realistic optical excitation, far-field optical techniques therefore remain an essential complementary approach.

### **Correlative optical–electron microscopy measurements**

Because plasmonic properties are extremely sensitive to nanoscale morphology, even subtle variations in shape, facet structure, or composition can lead to significant spectral changes. Linking optical spectra to morphology on the *same* nanoparticle enables robust, quantitative structure–property correlations. In this thesis, we combine single-particle dark-field scattering spectroscopy with TEM/STEM imaging of the same nanoparticles and compare the measured spectra with electromagnetic simulations based on the measured geometries. While many correlative optical–electron microscopy studies have relied on SEM due to its experimental accessibility, extending such approaches to TEM remains significantly more challenging. TEM-based correlation requires electron-transparent substrates, careful sample handling, and mitigation of electron-beam-induced modifications, making experiments on the exact same nanoparticle substantially more demanding. This approach allows us to connect resonance energies, linewidths, and scattering strengths to well-defined size, shape, composition, and environment, and to track light-induced morphological evolution and stability under operational conditions [52], [53], [54], [55]. As nanoparticle architectures become more complex and multimetallic [56], [57], [58], [59], [60], [61], [62], [63], such correlative EM–optical studies are essential for establishing reliable structure–property relationships and for informing nanoscale design strategies aimed at tuning catalytic activity/selectivity, enhancing light–matter coupling, or mitigating degradation [25], [40], [64], [65], [66], [67], [68], [69], [70], [71], [72], [73], [74].

Despite substantial progress in single-particle plasmonics, establishing quantitative structure–property relationships remains an open challenge. Experimentally measured spectra often agree qualitatively with electromagnetic simulations, capturing the overall mode structure and approximate resonance energies. However, persistent discrepancies in exact peak positions, linewidths, damping balances, and absolute scattering strengths reveal the limitations of modelling real nanoparticles with simplified or two-dimensional geometries. Small morphological details, facet truncation, surface roughness, rounded edges, crystallinity, or subtle distortions introduced during synthesis or measurement, can shift resonances by tens of nanometers or more, yet are difficult to incorporate faithfully into simulations. These differences obscure the ability to extract mechanistic information with confidence, including the contributions of surface scattering, chemical interface damping, or the emergence of non-classical effects at the smallest length scales. Reducing this mismatch between measured and simulated responses, and thereby enabling quantitative interpretation of plasmonic behaviour, serves as a central motivation for the work presented in this thesis.

## 1.3 Thesis outline

The societal and scientific context above converges on a single experimental challenge: to understand and control the optical response of plasmonic nanoparticles, one must be able to quantitatively link their nanoscale geometry to their measured optical properties. Achieving this requires accurate electrodynamic simulations, reliable morphology acquisition, high-quality single-particle optical measurements, and an understanding of how these structures evolve under operational conditions. This thesis builds that capability systematically, chapter by chapter.

**Chapter 2** establishes the theoretical and computational foundations needed to simulate nanoparticle optics from geometry. It analyses the dielectric response of gold (separating free-electron and interband contributions), quantifies the breakdown of quasistatic approximations for the sizes studied here, and benchmarks electrodynamic solvers against exact Mie theory, identifying the boundary element method (BEM) as the most accurate and efficient choice for our geometries. It also formalises how material datasets and local environments (e.g. substrates) propagate into simulated spectra, defining the simulation controls used throughout.

**Chapter 3** turns to morphology acquisition and its conversion into watertight simulation input. It demonstrates why two-dimensional imaging is insufficient for quantitative correlations and constructs a validated pipeline from HAADF-STEM electron tomography to simulation-grade surface meshes. By tracking both shape

error and spectrum error through to BEM, it identifies reconstruction/segmentation strategies that are most reliable and quantifies how morphology errors grow with geometric complexity, disproportionately affecting resonances governed by high-curvature features.

**Chapter 4** develops the optical platform used throughout the thesis: a quasi-dark-field reflectance geometry optimised for high signal-to-noise single-particle spectroscopy directly on TEM grids and on glass. By characterising and calibrating the instrument response (IRF) with reference scatterers and simulations, and by controlling angular collection through pupil engineering and spatial filtering, the setup yields IRF, corrected spectra that can be compared quantitatively to electrodynamic simulations.

**Chapter 5** brings these elements together to demonstrate quantitative morphology–property correlations on individual nanoparticles by combining calibrated single-particle optical spectra, TEM-based morphology, and electromagnetic simulations. It shows that a simple effective-medium description of the local dielectric environment can account for much of the particle-to-particle spectral variability, while remaining discrepancies point toward the need for full three-dimensional morphology and more realistic environment models in the most demanding cases. The result is a scalable framework: use 2D-based correlations for statistical surveys and trends, and reserve full 3D reconstructions for precision cases.

**Chapter 6** explores how plasmonic nanoparticles behave when driven far from equilibrium. Using pulsed optical excitation inside an electron microscope, it captures pulse-by-pulse structural changes of gold nanorods and reveals a rich family of reshaping behaviours governed by fast surface diffusion and kinetically biased recrystallization. Beyond its fundamental interest, this chapter connects non-equilibrium structural control to possible future routes for dynamic or resonant catalysis, where time-averaged reaction rates could exceed those achievable under steady-state conditions, an emerging prospect with direct relevance to light-driven sustainable chemistry.

Together, these contributions establish a quantitative foundation for understanding and designing plasmonic nanoparticles, with direct implications for the energy-harvesting and catalysis applications that motivated this work. By closing the loop between nanoscale geometry, optical response, and non-equilibrium dynamics, this thesis maps a route toward structure-aware control of plasmonic functionality at the nanoscale.

## Chapter 2

---

# Cast Light – Quantitative Electromagnetic Simulations

*Cast Light – to provide an explanation for a situation or problem, or information that makes it easier to understand*

**Abstract:** This chapter casts light on the theoretical and computational foundations required for quantitative simulations of plasmonic nanoparticles. It establishes gold as the primary material of study by analysing its free-electron response, interband transitions, and experimentally measured dielectric functions. Building on this, the chapter develops a hierarchy of electromagnetic models, from the Drude and Drude–Lorentz formalisms to full, frequency-dependent dielectric datasets, and examines how size-, geometry-, and interface-dependent damping mechanisms modify plasmonic behaviour as metals are reduced from bulk to nanoscale dimensions. The chapter then assesses the validity limits of the quasistatic approximation and demonstrates the necessity of retarded, full-wave electrodynamics for the nanoparticle sizes considered. A systematic benchmark of widely used electromagnetic solvers identifies the Boundary Element Method (BEM) as the most accurate and computationally efficient approach for the geometries studied. Finally, the chapter formulates practical meshing strategies and quantifies the impact of material datasets and environmental conditions, providing a robust, converged simulation framework for the single-particle structure–property analyses developed in later chapters.

Section 2.6 is published in M. Dieperink, A. Skorikov, N. Claes, S. Bals, and W. Albrecht, “Considerations for electromagnetic simulations for a quantitative correlation of optical spectroscopy and electron tomography of plasmonic nanoparticles,” *Nanophotonics*, vol. 13, pp. 4647–4665, 25 Dec. 2024, doi: 10.1515/nanoph-2024-0238.

## 2.1 Introduction

Quantitative comparison between single-particle optical measurements and electromagnetic simulations is essential for establishing reliable structure–property relationships in plasmonic systems. Such comparisons validate simulation frameworks, allow controlled inclusion or exclusion of specific damping mechanisms (surface scattering, chemical interface damping, radiation damping), and clarify whether more subtle effects such as nonlocal response or quantum confinement are required. Hence, simulation accuracy is a scientific requirement: the physical conclusions drawn from plasmonic spectra depend directly on the fidelity of the underlying model. This chapter develops the theoretical and computational foundation for such quantitative analyses.

## 2.2 Gold as a plasmonic material

Most plasmonic applications operate in the visible–near infrared (NIR) region. Suitable materials must therefore combine the correct spectral response with high chemical stability and low optical losses. Only a small set of metals satisfy these requirements. Alkali metals exhibit clean Drude-like behaviour and minimal interband losses but are too reactive for practical environments. Post-transition metals, such as aluminium or indium, offer greater stability but support resonances mainly in the ultraviolet, where interband transitions increase losses [75]. Noble metals, by contrast, combine visible–NIR plasmon resonances with comparatively low damping. Silver and copper have excellent optical performance but are prone to oxidation. Gold uniquely combines favourable optical behaviour with exceptional chemical stability, establishing it as the benchmark material in plasmonics [76], [77]. It is also the material used throughout this thesis.

The optical response of plasmonic metals is largely governed by the conduction-electron density. In the simplest approximation, these electrons form a free-electron gas described by the Drude model. The electron density is

$$n = \frac{Z\rho N_A}{M},$$

with  $Z$  the number of conduction electrons per atom,  $\rho$  the mass density,  $M$  the molar mass, and  $N_A$  Avogadro’s constant. For gold, assuming  $Z = 1$  yields  $n \approx 6 \times 10^{28} \text{ m}^{-3}$ , about 60 electrons per  $\text{nm}^3$ . A 50 nm gold sphere therefore contains  $\sim 4 \times 10^6$  conduction electrons.

The free-electron density sets the bulk plasma frequency

$$\omega_p = \sqrt{\frac{n e^2}{\varepsilon_0 m_e}},$$

where  $e$  is the elementary charge,  $\varepsilon_0$  is the vacuum permittivity, and  $m_e$  is the electron mass. For gold, this yields  $\omega_p \approx 9$  eV, corresponding to a wavelength of approximately 140 nm. For photon energies below  $\omega_p$ , the real part of the dielectric function becomes negative: the field is screened over a short skin depth, bulk gold strongly reflects visible and NIR light, and surface plasmon modes may form at interfaces. Confined geometries further support localized surface plasmon resonances (LSPRs).

However, the Drude response alone cannot describe gold's optical behaviour in the visible range, where interband transitions contribute substantially to both real and imaginary parts of the dielectric function. The next section therefore discusses how intraband and interband processes are modelled analytically.

## 2.3 Analytical modelling of the dielectric function of gold

Analytical dielectric models are valuable because they can be evaluated at arbitrary frequencies, allow physical contributions to be selectively varied, and provide smooth, compact representations suitable for parameter studies. We discuss the Drude–Sommerfeld model for intraband transitions, its extension with Lorentz oscillators for interband contributions, and the use of experimental dielectric datasets.

### 2.3.1 Intraband transitions: the Drude–Sommerfeld model

The classical Drude model [17], [78] treats conduction electrons as free carriers with relaxation time  $\tau$ , yielding

$$\varepsilon(\omega) = \varepsilon_\infty - \frac{\omega_p^2}{\omega^2 + i\gamma\omega}, \quad (2.1)$$

with

$$\varepsilon'(\omega) = \varepsilon_\infty - \frac{\omega_p^2}{\omega^2 + \gamma^2}, \quad \varepsilon''(\omega) = \frac{\gamma\omega_p^2}{\omega(\omega^2 + \gamma^2)}. \quad (2.2)$$

The damping parameter  $\gamma = 1/\tau$  may be estimated from the DC resistivity via

$$\tau = \frac{m_e}{n e^2 \rho_{\text{DC}}}. \quad (2.3)$$

For gold, this gives  $\tau \approx 25$  fs.

Sommerfeld refined the model by imposing Fermi–Dirac statistics [79]. Since only electrons within  $\sim k_B T$  of  $E_F$  contribute to momentum relaxation, this yields a slightly longer relaxation time ( $\tau \approx 30$  fs). However, Drude–Sommerfeld still underestimates losses in the visible–NIR because interband transitions dominate the damping.

Figure 2.1 compares Drude and Drude–Sommerfeld models to experimental thin-film data from McPeak *et al.* [80]. Both underestimate  $\epsilon''$ , especially below  $\sim 700$  nm. A single best-fit Drude model fitted only to an extra-NIR window ( $\lambda > 1200$  nm) reproduces the intraband tail but absorbs non-Drude losses into an artificially large  $\gamma$ . No single-parameter Drude model can reproduce the full visible–NIR behaviour, necessitating explicit treatment of interband transitions.




Figure 2.1: Real (a) and imaginary (b) parts of the dielectric function of gold. Blue: experimental data [80]. Dashed: Drude. Dash-dotted: Drude–Sommerfeld. Dotted: best-fit Drude model for  $\lambda > 1200$  nm.

### 2.3.2 Interband transitions: Lorentz oscillators

Interband transitions are modelled by augmenting the Drude term with Lorentz oscillators,

$$\epsilon_{DL}(\omega) = \epsilon_{\infty} - \frac{\omega_{D,p}^2}{\omega^2 + i\gamma_D\omega} + \sum_{j=1}^N \frac{\omega_{L,p,j}^2}{\omega_{L,j}^2 - \omega^2 - i\gamma_{L,j}\omega}. \quad (2.4)$$

where the Drude parameters ( $\omega_{D,p}, \gamma_D$ ) describe intraband motion and each Lorentz term ( $\omega_{L,j}, \omega_{L,p,j}, \gamma_{L,j}$ ) encodes a band-to-band transition with resonance energy, strength, and linewidth, respectively. For gold, the dominant interband absorption stems from  $5d \rightarrow 6s/6p$  transitions starting near  $\sim 1.8$  eV and peaking around  $\sim 2.4$  eV.

To determine the number of Lorentz oscillators required, we monitor convergence

using the spectral deviation metric

$$\Delta\sigma(N_i) = \sqrt{\frac{\int_{\lambda_{\min}}^{\lambda_{\max}} [\sigma(N_i, \lambda) - \sigma(N_{i-1}, \lambda)]^2 d\lambda}{\int_{\lambda_{\min}}^{\lambda_{\max}} \sigma^2(N_i, \lambda) d\lambda}}. \quad (2.5)$$

where  $N_i$  is the number of Lorentz oscillators at iteration  $i$ ,  $\sigma(N_i, \lambda)$  denotes the wavelength-dependent fitted dielectric response (either component), and the integration limits match the fit window. Figure 2.2 shows that  $\Delta\sigma$  saturates at  $N \approx 11$  for a 300–1700 nm window. The resulting Drude–Lorentz model captures magnitude and dispersion well, though over-parameterization may introduce narrow artefacts. In narrower windows, substantially fewer oscillators suffice. For quantitative work, we use either directly measured datasets or window-optimized Drude–Lorentz fits when analytic continuity is required.




Figure 2.2: Drude–Lorentz modelling of interband transitions in gold. (a) Convergence of the spectral deviation  $\Delta\sigma(N)$  [Eq. (2.5)] as a function of the number of Lorentz oscillators, saturating around  $N \approx 11$  for a 300–1700 nm window. (b) Real part  $\epsilon'(\lambda)$ : measured data (McPeak *et al.* [80]) vs. the converged Drude–Lorentz fit. (c) Imaginary part  $\epsilon''(\lambda)$ : same comparison. Detailed one-oscillator comparisons and intermediate multi-oscillator fits are provided in the Appendix.

### 2.3.3 Measured dielectric functions

For fully quantitative simulations, experimentally measured dielectric functions are generally preferred, as they incorporate band structure, interband transitions, and film microstructure by construction. Figure 2.3 compares widely used datasets [80], [81], [82], [83]. Systematic variations in  $n(\lambda)$  and  $k(\lambda)$  arise from differences in film thickness, grain size, roughness, annealing, deposition methods, and the ellipsometry model used to extract optical constants [80], [83]. Because  $k$  directly controls loss, such variations translate into differences in simulated resonance energies and linewidths (see Fig. 2.6).

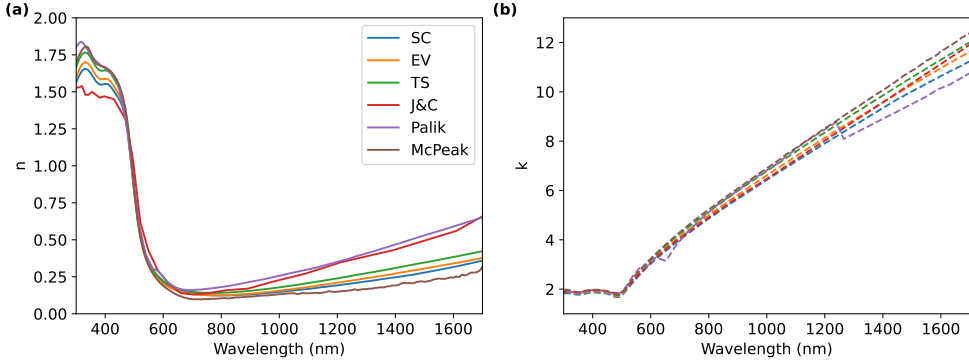


Figure 2.3: Measured dielectric functions of gold from Refs. [80], [81], [82], [83]. (a)  $n(\lambda)$ ; (b)  $k(\lambda)$ .

Dataset selection in this thesis is guided by the fabrication method and structural quality of the measured nanoparticles. Low-loss, large-grain datasets (e.g. template-stripped films [80]) best represent high-quality crystals, while others may better capture evaporated or rougher films. The quantitative impact of dataset choice is analysed in Sec. 2.6. Bulk dielectric functions, however, do not include nanoparticle-specific damping processes such as surface scattering, chemical interface damping, or radiation damping. These additional contributions are introduced in the next section.

## 2.4 From bulk to nanoparticle

Reducing a metal from a bulk film to a nanoscale particle introduces additional decay pathways that modify the plasmon linewidth beyond the intrinsic bulk damping  $\gamma_b$  [84], [85], [86]. In bulk metals, the effective damping rate is dominated by

- **Landau damping** ( $\tau_{La} \sim 1\text{--}10$  fs): plasmon decay into electron–hole pairs [18];
- **electron–electron scattering** ( $\tau_{e-e} \sim 10\text{--}100$  fs);
- **electron–phonon scattering** ( $\tau_{e-p} \sim 0.1\text{--}1$  ps);
- **impurity/defect scattering** ( $\tau_{im} \sim 0.1\text{--}10$  ps).

These contributions add approximately as

$$\gamma_b = \gamma_{La} + \gamma_{e-e} + \gamma_{e-p} + \gamma_{im}.$$

In nanoparticles, additional mechanisms appear. **Radiation damping** ( $\gamma_{rad}$ )

## 2.5. From quasistatic to retarded electrodynamics: when does the Rayleigh approximation fail?

---

describes re-emission of electromagnetic energy by the dipole moment, scaling roughly with particle volume and becoming significant above  $\sim 30\text{--}40$  nm. **Surface-assisted Landau damping** ( $\gamma_{\text{surf}}$ ) arises when dimensions approach the electron mean free path ( $\sim 40$  nm for gold), enabling boundary-assisted transitions into electron–hole pairs. **Chemical interface damping** ( $\gamma_{\text{CID}}$ ) occurs when adsorbates or nearby molecules accept charge or energy from the plasmon [85].

The total damping is thus

$$\gamma_{\text{tot}} = \gamma_b + \gamma_{\text{rad}} + \gamma_{\text{surf}} + \gamma_{\text{CID}}.$$

For gold nanoparticles of 30–100 nm, the size range examined in this thesis, the dominant contributions are intrinsic bulk damping, radiation damping (especially for larger particles), and modest surface scattering; chemical interface damping depends on surface chemistry. A goal of this work is to isolate and quantify these channels experimentally via correlative tomography, full-wave simulations, and single-particle spectroscopy. The next section examines when quasistatic approximations break down and when full-wave electromagnetic simulations are required.

## 2.5 From quasistatic to retarded electrodynamics: when does the Rayleigh approximation fail?

Before turning to full-wave methods, we assess when the quasistatic (Rayleigh) approximation accurately describes nanoparticle optics. In the limit  $r \ll \lambda$ , the incident field is uniform across the particle. For a spherical particle, the dipole moment is

$$p = \varepsilon_0 \varepsilon_m E_0 \alpha,$$

where  $\varepsilon_0$  is the vacuum permittivity,  $\varepsilon_m$  is the dielectric constant of the surrounding medium,  $E_0$  is the incident field, and  $\alpha$  is the polarizability. The quasistatic polarizability is given by

$$\alpha = 3V \frac{\varepsilon(\omega) - \varepsilon_m}{\varepsilon(\omega) + 2\varepsilon_m},$$

and the LSPR condition

$$\text{Re}[\varepsilon(\omega)] = -2\varepsilon_m.$$

In this regime, the resonance position is size-independent, while the linewidth follows the intrinsic damping. Absorption and scattering cross sections are [87]

$$C_{\text{ext}} = C_{\text{sca}} + C_{\text{abs}} = \frac{k^4}{6\pi} |\alpha|^2 + k \text{Im}(\alpha), \quad k = \frac{2\pi}{\lambda} \sqrt{\varepsilon_m}.$$

For spheres, Mie theory [88] provides the exact electrodynamic solution and serves as a benchmark. For non-spherical particles, such as the nanorods studied in this thesis, no analytical solution exists, and numerical methods are required even at small sizes. To quantify breakdown of the quasistatic approximation, Fig. 2.4 compares quasistatic and fully retarded BEM simulations for a 50 nm gold sphere in media of different refractive indices. Retardation causes red-shifting, linewidth broadening, and redistribution of energy into higher multipoles. The spectral deviation metric (Eq. (2.5)) shows that deviations increase with size but may exhibit non-monotonic behaviour due to partial cancellation between errors in position and width (notably in oil).

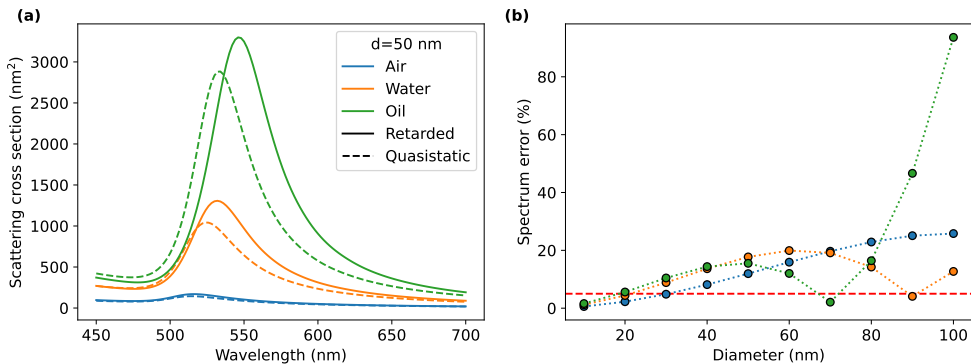


Figure 2.4: (a) Quasistatic (dashed) and fully retarded (solid) BEM scattering spectra for a 50 nm gold sphere in air, water, and oil. (b) Spectrum deviation of the quasistatic approximation relative to the retarded solution, as a function of sphere diameter (using the spectral deviation metric defined in Eq. (2.5)). The red dashed line indicates a 5% deviation threshold.

Figure 2.4(b) shows that for spheres, a 5% deviation threshold corresponds to a diameter of  $\sim 20$  nm under typical conditions. Above this, full electrostatics is required. To clarify the contributions to this deviation, Fig. 2.5 compares Lorentzian fit parameters. Retardation red-shifts the resonance (Fig. 2.5b), broadens it via radiation damping and multipole excitation (Fig. 2.5c), and reduces the relative dipole peak height (Fig. 2.5a) due to energy redistribution.

For anisotropic particles, retardation becomes important when any characteristic

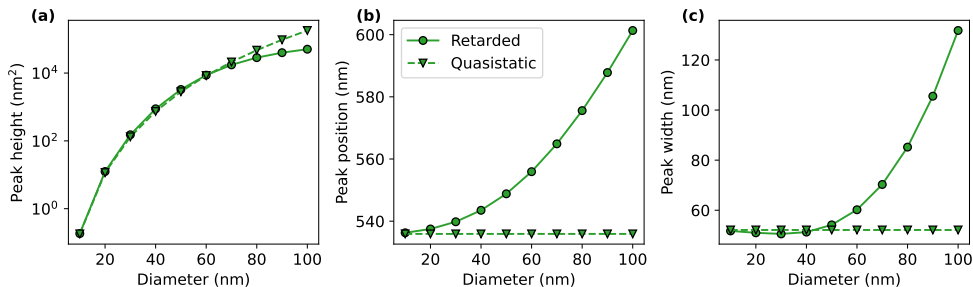


Figure 2.5: Lorentzian parameters for quasistatic (dashed) and retarded (solid) simulations in oil: (a) peak height, (b) resonance wavelength, (c) linewidth.

dimension approaches a significant fraction of the wavelength. For nanorods, the long axis dictates this threshold, and retardation becomes relevant even when the short axis lies well below 20 nm. The particles studied in this thesis therefore require retarded full-wave simulations. The next section benchmarks different electrodynamic solvers against the exact Mie solution to identify the most accurate and efficient method for the geometries relevant here.

## 2.6 Considerations for electromagnetic simulation methods

Having established that retarded full-wave simulations are required for the particle sizes and geometries studied in this thesis, where the quasistatic approximation breaks down for spheres and, critically, no analytical solution exists for non-spherical shapes, the practical question becomes: *which numerical method delivers accurate spectra at reasonable computational cost, and how can convergence be certified rigorously and reproducibly?* To address this, we compare four widely used electromagnetic solvers: the Discrete Dipole Approximation (DDA), the Finite-Difference Time-Domain (FDTD) method, the Discontinuous Galerkin Time-Domain (DGTD) method (a finite-element family), and the Boundary Element Method (BEM) [34], [54], [55], [89].

For a fair, quantitative comparison, a spherical Au nanoparticle is adopted as the test case because the analytical Mie solution provides an exact reference [88]. Each method is optimized by sweeping its principal discretization parameter (e.g., finite-element edge length for DGTD, number of surface triangles for BEM, number of dipoles for DDA, and grid spacing together with propagation settings for FDTD)

until convergence is achieved according to a *discretization* error metric,

$$\Delta\sigma(d_i) = \sqrt{\frac{\int [\sigma(d_i, \lambda) - \sigma(d_{i-1}, \lambda)]^2 d\lambda}{\int \sigma^2(d_i, \lambda) d\lambda}}, \quad (2.6)$$

and a *reference* error with respect to Mie theory,

$$\Delta\sigma_{\text{ref}}(d_i) = \sqrt{\frac{\int [\sigma(d_i, \lambda) - \sigma_{\text{ref}}(\lambda)]^2 d\lambda}{\int \sigma_{\text{ref}}^2(\lambda) d\lambda}}. \quad (2.7)$$

Here,  $d_i$  is a monotone discretization parameter,  $\sigma$  denotes the simulated scattering spectrum, and  $\sigma_{\text{ref}}$  is the Mie spectrum. Additional parameters and implementation details for all solvers are provided in Appendix 2.A.

In practice, different methods expose multiple interdependent settings (e.g., FDTD simulation time and auto-shutoff criteria) whose naive choices can yield apparently smooth spectra that are nonetheless unconverged. The convergence framework above separates (i) *internal* discretization error,  $\Delta\sigma$ , certifying stability with respect to refinement, from (ii) *external* reference error,  $\Delta\sigma_{\text{ref}}$ , certifying quantitative agreement with the analytical solution. Using both is essential for reproducible accuracy claims.

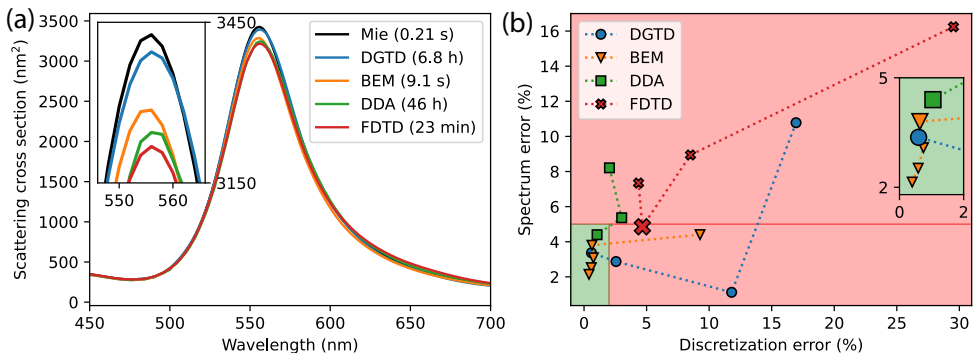


Figure 2.6: Comparison to Mie theory of simulated scattering spectra of a Au sphere (50 nm diameter) immersed in oil ( $n = 1.51$ ) using DGTD, BEM, DDA, and FDTD. (a) Qualitative comparison of spectra. Legend times correspond to the runtimes required to produce the spectra highlighted in (b) that were the first to reach  $\Delta\sigma < 2\%$  and  $\Delta\sigma_{\text{ref}} < 5\%$  (green rectangle), except FDTD which did not reach  $\Delta\sigma < 2\%$  for the swept parameters. (b) Quantitative convergence map: discretization error [Eq. 2.6] versus reference error [Eq. 2.7], with enlarged markers corresponding to the spectra emphasized in (a).

Figure 2.6 summarizes the comparison. Panel 2.6a shows scattering spectra for a 50 nm Au sphere in oil. For the quantitative comparison in Panel 2.6a, each method

is refined until it first attains  $\Delta\sigma < 2\%$ ; the corresponding  $\Delta\sigma_{\text{ref}}$  is then evaluated against Mie theory. This provides a like-for-like basis: spectra of comparable *internal* convergence are contrasted by their *external* reference error and the runtime needed to reach that state. As expected, all solvers can approach the reference solution when sufficiently refined, but their convergence behaviour differs:

1. **BEM** shows a highly regular decrease in both  $\Delta\sigma$  and  $\Delta\sigma_{\text{ref}}$  with mesh refinement, and reaches the target accuracy orders of magnitude faster than the volume/time-domain methods. The speed advantage arises because BEM enforces the field solution only on the particle boundary rather than throughout the enclosing volume or a time-propagation domain [89].
2. **DGTD** can achieve very small  $\Delta\sigma_{\text{ref}}$  in specific settings, but at substantially higher computational cost; moreover, its  $\Delta\sigma_{\text{ref}}$  trend can be non-monotonic (finer meshes not strictly reducing the reference error), making the refinement path less predictable.
3. **DDA** exhibits irregularities in the discretization error  $\Delta\sigma$  across dipole-density refinements, necessitating careful parameter sweeps to identify stable regimes.
4. **FDTD** is further penalized by the need to propagate fields in a relatively large domain (absorbing boundaries, padding), and its convergence depends sensitively on coupled parameters (grid spacing, simulation time, auto-shutoff). In our sweeps it did not reach  $\Delta\sigma < 2\%$  within practical runtimes for this geometry.

These observations are consistent with prior qualitative reports on method performance [54], but here are supported by a *quantitative* convergence framework. For the purposes of this thesis, where absolute cross sections and reproducible convergence are required, BEM offers the most favourable accuracy–runtime trade-off. We therefore adopt BEM for all subsequent nanoparticle simulations.

### 2.6.1 BEM meshing for anisotropic nanoparticles

Given BEM’s favourable performance for metals, it is essential to understand how surface meshing controls both convergence and cost for anisotropic shapes. We use a gold nanorod as a standard test geometry because it supports polarization-dependent modes and concentrates fields at high-curvature tips. In the MNPBEM toolbox a spherocylindrical rod is parameterized by three discrete meshing integers ( $n_\phi, n_\theta, n_z$ ) for the azimuthal, cap, and axial directions, respectively [90]. For practical comparison,

we relate them to approximate physical edge lengths of the surface triangles,

$$dis_\phi \approx \frac{\pi(d+1)}{n_\phi}, \quad dis_\theta \approx \frac{d+1}{n_\theta}, \quad dis_z \approx \frac{l-d+1}{n_z},$$

with rod diameter  $d$  and length  $l$ ; smaller  $dis_\phi, dis_\theta, dis_z$  imply finer meshes.

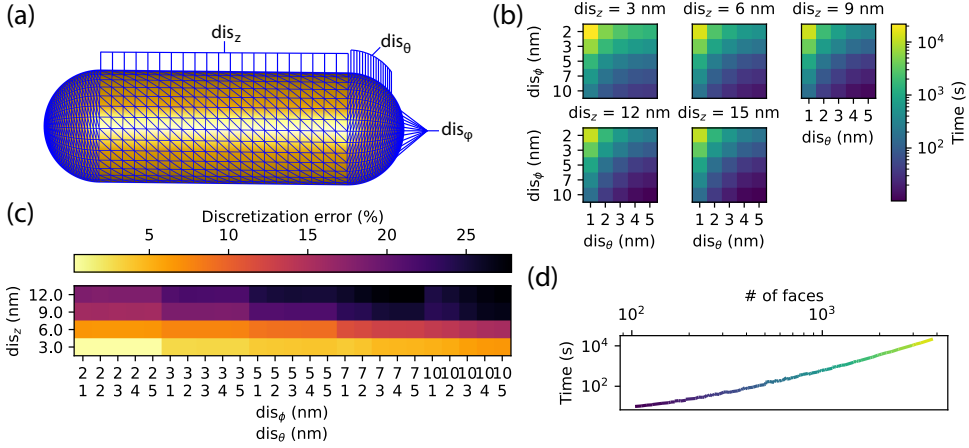


Figure 2.7: Effect of BEM surface discretization for a Au nanorod (diameter 30.0 nm, length 96.5 nm; AR=3.2). (a) Example surface mesh with  $dis_z \approx 3$  nm,  $dis_\theta \approx 1$  nm, and  $dis_\phi \approx 3$  nm. (b) Runtimes and (c) corresponding relative differences in discretization error [Eq. 2.6] as a function of individual meshing parameters, and (d) versus the total number of faces.

Figure 2.7 reveals three practical guidelines: (i) reducing  $dis_z$  (axial meshing) often yields the largest initial drop in  $\Delta\sigma$ , establishing a low-error baseline at modest cost; (ii) further reducing  $dis_\phi$  and  $dis_\theta$  (circumferential and cap meshing) is required to reach sub-percent  $\Delta\sigma$ , but this substantially increases face count and runtime; (iii) strongly non-equilateral triangles (e.g., large disparity between  $dis_z$  and  $dis_\phi$ ) degrade accuracy, consistent with general finite-element meshing best practices [91], [92]. Empirically, the face count follows

$$n_{\text{faces}} \approx n_\phi (n_z + 2n_\theta), \quad (2.8)$$

which helps anticipate runtimes for a target  $\Delta\sigma$ . For the rod studied here, a practical accuracy–speed balance is obtained near  $dis_z \approx 3$  nm,  $dis_\phi \approx 3$  nm,  $dis_\theta \approx 1$  nm (few minutes;  $\Delta\sigma < 4\%$ ), whereas sub-percent  $\Delta\sigma$  requires fine circumferential/cap meshing at higher computational cost. Throughout this chapter, we therefore use

$$dis_\phi \approx 3 \text{ nm}, \quad dis_\theta \approx 1 \text{ nm}, \quad dis_z \approx 3 \text{ nm}$$

unless stated otherwise, to keep comparisons general across rods while maintaining predictable runtimes.

## 2.6.2 Material and environment choices: dielectric function and substrates

With numerical convergence under control, the next dominant uncertainties are the *material* response and the *environment*. As discussed in Sec. 2.3.3, literature dielectric functions for gold differ due to measurement protocol, microstructure (grain size), surface roughness, annealing, and the inversion model used in ellipsometry [80], [81], [82], [83]. Even modest changes in  $k(\lambda)$  lead to substantial differences in simulated cross sections, linewidths, and LSPR energies.

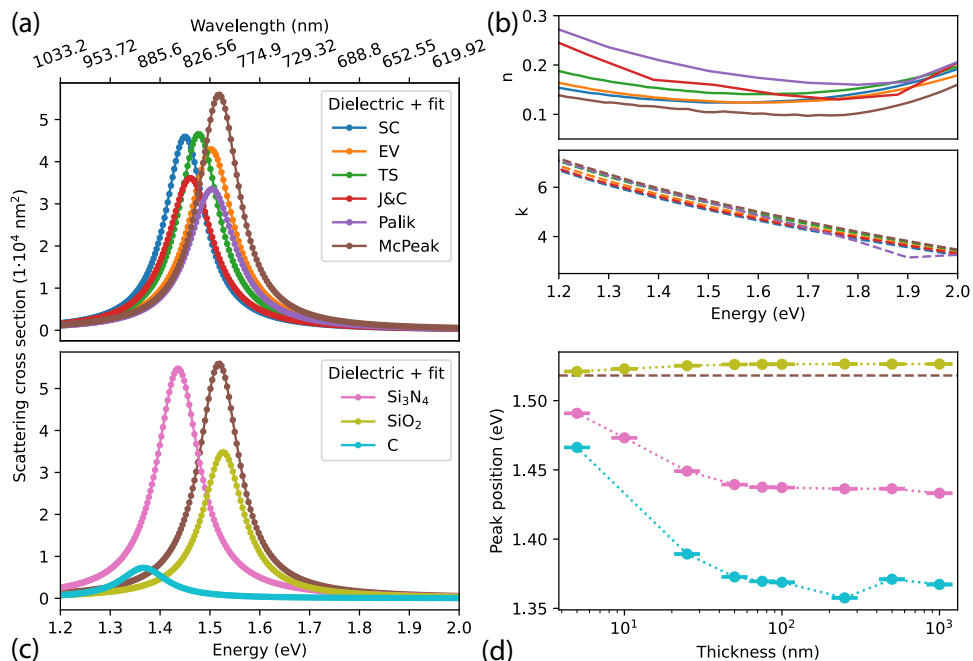


Figure 2.8: Material and environment effects for a Au nanorod (diameter 30.0 nm, length 96.5 nm; AR=3.2). (a) BEM scattering spectra using dielectric functions from Olmon *et al.* (single-crystalline, evaporated, template-stripped variants), Johnson & Christy, Palik, and McPeak *et al.*; mesh:  $dis_z \approx 3 \text{ nm}$ ,  $dis_\theta \approx 1 \text{ nm}$ ,  $dis_\phi \approx 3 \text{ nm}$ ; medium: oil,  $n = 1.51$ . (b) Corresponding  $n(E)$  and  $k(E)$ ; plotting on an *energy* axis facilitates direct comparison with (a). (c) Effect of different substrates ( $\text{SiO}_2$ ,  $\text{Si}_3\text{N}_4$ , c) on the spectra for an oil environment and an optically thick substrate. (d) LSPR peak position versus substrate thickness for the systems in (d); the brown dashed line indicates the no-substrate result.

Figure 2.8a compares BEM spectra for a single nanorod using several widely cited Au datasets; Panel 2.8b plots the corresponding  $n(E)$  and  $k(E)$  on an energy axis for direct visual comparison with (a). Small differences in  $n, k$  translate into noticeable shifts of the longitudinal LSPR and significant changes in peak magnitude. Trends correlate with microstructural indicators from the original works [80], [83]: larger grains and smoother films generally reduce loss, yielding higher-energy and higher-magnitude resonances for rods of the same morphology. In this thesis we adopt McPeak *et al.* as the *primary* dataset because it best reflects the high-quality, low-loss films representative of our samples; we note that single-particle extractions are ideal but limited in spectral range for now [93].

The *environment* must also be modelled explicitly. Most single-particle measurements are performed on substrates that modify local fields and radiative coupling. Figures 2.8c,d show the effects of substrate material and thickness: carbon (high index and lossy in the visible–NIR) induces the strongest red-shift and damping; SiO<sub>2</sub> is better index-matched to immersion oil and produces minimal thickness-dependent LSPR shifts [31], [94], [95], [96]. Consequently, quantitative comparisons should (i) include realistic substrates and thicknesses and (ii) choose a dielectric dataset consistent with sample quality and preparation.

## 2.7 Conclusions and outlook

Based on the analyses in Secs. 2.6–2.7: (i) retarded full-wave simulations are required for the particle sizes studied here, as the Rayleigh approximation fails above  $\sim 20$  nm for spheres and at even smaller volumes for anisotropic particles where the long axis drives retardation; (ii) BEM is adopted as the primary solver given its favourable accuracy–runtime trade-off and predictable convergence; (iii) the McPeak dielectric dataset is used as the primary material input, reflecting low-loss films most representative of our nanoparticles; (iv) BEM surface meshing for nanorods should prioritise axial refinement ( $dis_z$ ) to establish a low-error baseline, followed by circumferential/cap refinement ( $dis_\phi, dis_\theta$ ) to reach target accuracy, while avoiding strongly non-equilateral triangles; (v) substrate effects must be included explicitly, with SiO<sub>2</sub> preferred for oil-immersion measurements due to index matching and minimal thickness-dependent shifts.

Numerical convergence is a controllable error source that can be reduced systematically; the convergence framework used here provides rigorous criteria for certifying accuracy (internal  $\Delta\sigma$  and external  $\Delta\sigma_{\text{ref}}$ ). The choice of dielectric dataset introduces a systematic uncertainty that cannot be eliminated but can be minimised by selecting

datasets representative of the sample quality. Substrate effects are non-negligible and must be modelled explicitly. After addressing these, the dominant remaining uncertainty in single-particle comparisons is *morphology*: small deviations in shape, tip geometry, or aspect ratio can shift LSPR position and linewidth by amounts comparable to or larger than the spread induced by dataset choice.

The simulation framework established here is necessary but not sufficient for quantitative structure–property correlations: given an accurate morphological description, the optical response can be computed reliably. What remains is to obtain that morphology with sufficient fidelity and in a watertight mesh form suitable for simulation input. The next chapter develops a correlative framework that derives surface meshes directly from electron tomography and quantifies how morphological uncertainties propagate into simulated spectra, completing the foundation for the single-particle correlations that are the scientific goal of this thesis.

## 2.A Appendix




Figure 2.A1: Single-Lorentz-oscillator extensions to the Drude model. (a) Onset-matched model reproducing the  $5d \rightarrow 6s/6p$  interband threshold near  $\sim 1.8$  eV. (b) Full-range optimized fit. Measured data (McPeak *et al.* [80]) shown in blue. The onset-matched model captures the sharp rise in  $\epsilon''$  but fails at higher energies, while the full-range fit captures overall amplitude but misses the onset, showing that a single oscillator cannot reproduce both features.

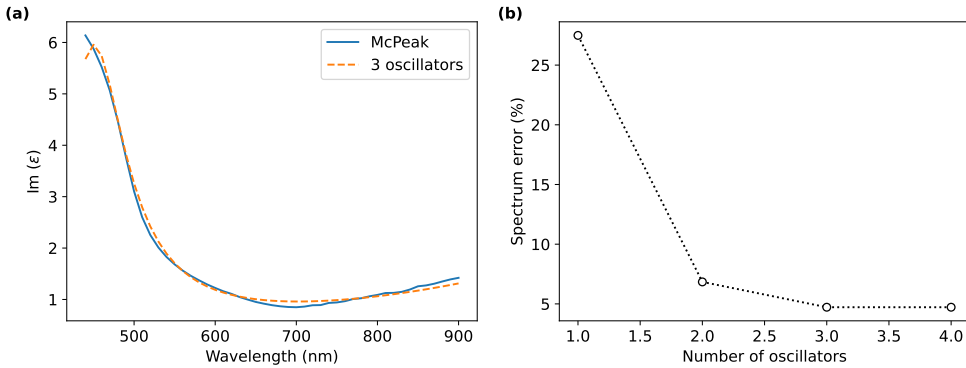


Figure 2.A2: Three-oscillator Drude-Lorentz fit in the visible-NIR range. (a)  $\epsilon''(\lambda)$  for gold (blue; McPeak *et al.* [80]) and corresponding fit (dashed), showing good agreement over the restricted range. (b) Spectral deviation  $\Delta\sigma(N)$  vs. number of oscillators, with strong improvement up to  $N = 3$  and minimal gain beyond. Fewer oscillators suffice for windowed fits compared to full-range modelling.




Figure 2.A3: Simulations and Lorentzian fits for Figure 2.5.

## 2.A.1 Electromagnetic simulation methods

For electromagnetic simulations of plasmonic nanoparticles, four main methods are commonly used that solve Maxwell's equations numerically for a given geometry. Below a short description for each of the methods is given:

- **Discrete Dipole Approximation (DDA)**. Nanostructures' volumes are divided into a set of interacting radiating dipoles. To obtain the absorption and scattering cross sections the interaction matrix needs to be solved. The accuracy of the method depends on the number of dipoles and the spacing between them. Often, more than a million closely spaced dipoles are needed to reach a desired accuracy. This makes the DDA method computationally intensive.
- **Finite-Difference Time-Domain (FDTD)**. The volume is meshed into Yee cells. Maxwell's equations are simplified into difference equations and solved via leapfrog time-stepping. Two constraints apply:
  - the dielectric function must be converted into the time domain,
  - rectilinear meshing introduces staircasing errors for curved shapes.
- **Discontinuous Galerkin Time-Domain (DGTD)**. Uses tetrahedral meshes to avoid staircasing. Local Helmholtz problems are solved in each element, coupled by numerical flux. Arbitrary shapes can be simulated with reduced computational volume.
- **Boundary Element Method (BEM)**. Only the surfaces of nanostructures

are discretized. Surface charges are solved via integral equations. This enables efficient simulations for plasmonic metals.

In the following, details on all simulation methods are described.

## DDA

DDA simulations were performed with DDSCAT 7.3.3 using Ubuntu Linux under Oracle VirtualBox [97]. The BiConjugate Gradient with Stabilization (PBCGS2) solver was used. The GPFAFT FFT routine was required (MKL unavailable). Polarizability was described using the Lattice Dispersion Relation (GKLDR).

The memory region was set to  $130 \times 130 \times 130$  nm for a 50 nm Au sphere. Interdipole spacings of 0.8, 0.5, and 0.4 nm were used, corresponding to 127,832; 523,599; and 1,022,654 dipoles, respectively. Other parameters followed Table 2.A1.

## FDTD

FDTD simulations used Lumerical 2022 R1.3. The dielectric function of Olmon et al. was fitted between 400–750 nm with tolerance 0.01. Auto-shutoff ( $10^{-8}$ ) controlled simulation termination. The simulation region spanned 200 nm on each axis with auto-nonuniform meshing. Conformal mesh variant 0 was used. Default 8-layer PML was retained.

The mesh size was swept from 5 nm down to 0.25 nm for Figure 2.6. Remaining parameters followed Table 2.A1.

## DGTD

DGTD simulations used Lumerical 2023 R1.2 due to a material-handling bug in 2022 R1.3. The simulation boundary was a 250 nm radius sphere with absorbing boundaries. The source region radius was 70 nm. Maximum tetrahedral element sizes of 15, 10, 5, 2, and 1 nm were swept. Other parameters followed Table 2.A1.

## BEM

BEM simulations used the MNPBEM17 MATLAB toolbox [90]. Retarded solvers with curved surfaces were used. Surface meshes were refined from 32 to 1444 faces for convergence testing in Figure 2.6. Remaining parameters followed Table 2.A1.

## Dielectric constants of the surrounding media

Figure 2.A4 shows refractive indices  $n(\lambda)$  and  $k(\lambda)$  for  $\text{SiO}_2$ ,  $\text{Si}_3\text{N}_4$ , C, and immersion oil. The peak shifts in Figure 2.8 correlate with  $n$ , peak height with  $k$ , and linewidth with combined effects. No additional damping mechanisms were included due to large particle sizes.

Table 2.A1: Parameters kept constant for all numerical methods used in Figure 2.6 of the main text.

Parameter	Value
Dielectric function (Au)	Olmon et al. [83]
Environment	Immersion oil ( $n = 1.51$ )
Spectral range	400–750 nm (2 nm steps)
Particle size	Diameter = 50 nm
Excitation	Plane wave, $x$ -polarized, incident from $+z$




Figure 2.A4: Refractive indices  $n$  and  $k$  of  $\text{SiO}_2$ ,  $\text{Si}_3\text{N}_4$ , carbon, and immersion oil.



## Chapter 3

---

# Screen-cast – From Electron Tomography to Electromagnetic Simulations

*Screen-cast – the act of making an image, a piece of film, etc. appear on a screen by sending it electronically from another device such as a computer or phone*

**Abstract:** This chapter develops the methodological bridge between three-dimensional electron tomography and quantitative electromagnetic simulations. While Chapter 2 established that plasmonic resonances are highly sensitive to nanoscale morphology, here we examine how accurately those morphologies can be retrieved in practice, and how errors introduced during tomography, segmentation, and meshing propagate into simulated optical responses. Using both experimental and realistically simulated HAADF–STEM tilt series, we systematically compare reconstruction algorithms, thresholding strategies, and surface-meshing approaches. We quantify the resulting shape deviations using voxel-level error metrics and evaluate their optical impact through full-wave Boundary Element Method (BEM) simulations. By analysing nanorods alongside more complex geometries, we identify the processing steps that most critically affect electromagnetic accuracy and formulate a robust workflow, from raw projections to watertight surface meshes, that minimizes morphology-induced uncertainties. This framework provides the structural fidelity required for reliable single-particle structure–property correlations in subsequent chapters.

The content of this chapter is published in M. Dieperink, A. Skorikov, N. Claes, S. Bals, and W. Albrecht, “Considerations for electromagnetic simulations for a quantitative correlation of optical spectroscopy and electron tomography of plasmonic nanoparticles,” *Nanophotonics*, vol. 13, pp. 4647–4665, 25 Dec. 2024, doi: 10.1515/nanoph-2024-0238.

### 3.1 Introduction

In Chapter 2 we established the numerical foundations required for quantitative single-particle simulations: a validated and convergent Boundary Element Method (BEM) framework, together with a careful treatment of dielectric functions and substrate geometry. We concluded that, once these numerical uncertainties are controlled, morphology remains the dominant source of discrepancy in structure–property correlations. Even nanometer-scale variations in tip curvature, facet truncation, or aspect ratio can shift localized surface plasmon (LSPR) energies and linewidths by amounts comparable to, or exceeding, those caused by different dielectric datasets. This makes the accurate retrieval of three-dimensional nanoparticle geometries essential for reliable simulations. The present chapter therefore develops and validates a pipeline that extracts such morphologies from electron tomography data and quantifies how uncertainties introduced at each processing step propagate into the resulting electromagnetic response.

Plasmonic nanoparticles (PNPs) can be exploited in a large range of optical applications [98], such as optical data storage [65], [66], sensing [67], [68], or photocatalysis [25], [40], [69], [70], [71], [72] owing to their highly tuneable LSPRs ranging from the UV to the IR region. Next to the plasmonic and surrounding material, the morphology of the PNP is the key ingredient in defining the optical response [73], as detailed in Chapter 2. For that reason, a lot of effort has gone into developing new protocols for the colloidal synthesis of PNPs with varying shapes and compositions [74], [99]. In particular, for gold NPs an impressive control over morphology has been achieved, and highly anisotropic nanoparticles can now be routinely synthesized [56], [57]. Examples include platelets [58], platonic solids [59], stars [60], [61], and even twisted [62] or wrinkled nanoparticles [26], [27], [63]. The more complex the morphology, the less straightforward it becomes to correlate structural and optical properties, as several LSPR modes emerge and sample polydispersity increases. This broadens ensemble spectra and can mask correlations [100].

Consequently, analysis at the level of single particles has become increasingly important for understanding structure–property relationships [40], [101], [102], [103], [104]. Single-particle optical properties are most commonly measured via scattering techniques such as dark-field scattering spectroscopy [29], [105], [106]. For a full picture, such optical data are ideally correlated with structural information of the same nanoparticle [32], [33], [34], [41], [42], [43], [44], [107] as outlined in Section 1.2.4 of the Introduction. Due to the small dimensions of typical PNPs, electron microscopy is indispensable for retrieving the necessary morphological information. Because of simplified sample preparation and measurement workflows, SEM is often preferred over

TEM, but both provide only two-dimensional projections. For increasingly complex morphologies, 2D information becomes insufficient, and electron tomography (ET) has been established as a powerful technique to visualize and quantify three-dimensional structural properties [35], [45], [46], [47], [49], [50], [51], [108]. For crystalline materials, high-angle annular dark-field scanning TEM (HAADF-STEM) is typically used, as its signal approximately satisfies the projection requirement for tomographic reconstruction [36]. During ET, the grid is tilted over the maximal achievable range, often around  $\pm 75^\circ$  due to sample-holder shadowing [37], and each tilt angle yields a projection image. These projections are subsequently combined into a three-dimensional reconstruction.

However, incorporating ET into a single-particle correlation workflow introduces additional complexity. The TEM grids most suited for ET (thin carbon-based supports) perform poorly for optical scattering experiments, while grids more suitable for optical measurements (e.g. ultrathin  $\text{SiO}_2$ ) can complicate tomography due to charging and increased shadowing at high tilt angles [38], [44]. Furthermore, reconstructions need to be segmented and often meshed before they can be used for quantitative analysis. A variety of reconstruction [109], segmentation [110], and meshing algorithms [111] exist, each introducing different artifacts or uncertainties. So far, no comprehensive quantitative comparison of these factors has been available.

In this chapter, we therefore systematically examine the workflow from raw ET projections to segmented 3D morphology and triangulated surface meshes. Using simulated tilt-series data based on a known ground-truth particle, we quantify how different reconstruction algorithms, segmentation strategies, and meshing approaches affect the resulting morphology. In addition, we perform electromagnetic simulations on these meshes to assess how morphology deviations, introduced at each processing step, translate into errors in the optical response. While Chapter 2 focused on the numerical foundations and convergence behaviour of the simulation methods themselves, the emphasis here is on how reconstruction and processing choices influence the simulated spectra through their impact on the underlying morphology. The analysis is carried out for nanorods and extended to more complex shapes such as a triangle and an octopod to assess robustness across geometries. Ultimately, this pipeline serves the broader goal of the thesis: establishing a correlative framework in which discrepancies between measured and simulated spectra can be confidently attributed to genuine physical effects, such as nonlocal response, anomalous damping, or chemical interface interactions, rather than to inaccuracies in the reconstructed geometry.

## 3.2 Limitations of 2D imaging for morphology extraction

When correlating optical and structural properties of single plasmonic nanoparticles, the most straightforward approach is to use 2D SEM or TEM images to retrieve the structure of the nanoparticle. As discussed in the introduction, this approach is not applicable to complex shapes, and one needs to resort to electron tomography. But even for seemingly symmetric particles extracting a 3D structure from 2D images might lead to incorrect estimation of morphological parameters, complicating the structure-property interpretation. An example is shown in Figure 3.1 for a Au nanorod imaged by HAADF-STEM (Figure 3.1a). In order to extract the length and width of the Au nanorod, the image needs to be segmented to differentiate the particle from the background signal. The corresponding pixel intensity histogram of the image in Figure 3.1a is displayed in green in Figure 3.1c, revealing two peaks in the intensity distribution. The peak at lower intensities corresponds to the background, while the peak at higher intensities corresponds to the foreground.

One problem for 2D data like the one in Figure 3.1a is the rather smooth transition between the two distributions due to the thickness-dependent HAADF-STEM intensity and the strong influence of Poisson noise, which makes it difficult to segment the particle [112]. Several algorithms exist for segmentation, with the most common one being Otsu’s method [113], [114], but none of them are designed for the case of a smooth transition between two intensity distributions. Therefore, choosing a threshold value in this is not straightforward as demonstrated in Figure 3.1b,c. The vertical lines in Figure 3.1c correspond to threshold values using Otsu’s method (green line) or fractions thereof (0.5: blue, 0.75: orange, 1.25: red, 1.5: purple). All these values could seem like a reasonable choice for separating the two distributions. However, the effect of this choice can be seen directly in Figure 3.1b, where the colors correspond to the threshold choice. In combination with the finite pixel size, the choice of threshold for the 2D HAADF-STEM image resulted in highly changing extracted dimensions of the nanorod. The changing nanorod sizes directly influenced the simulated LSPR of the PNP tremendously, which is shown in Figure 3.1d by performing BEM simulations with models of fitted sizes using the 2D fitting algorithm described in Section 3.A.2.5. This illustrates the uncertainty of relying on 2D data when aiming for accurate simulations of electromagnetic properties of plasmonic nanoparticles.

We proceeded by comparing the 2D pixel intensity histograms to the voxel intensity histogram of the 3D data set. The latter was obtained by acquiring a set of 2D

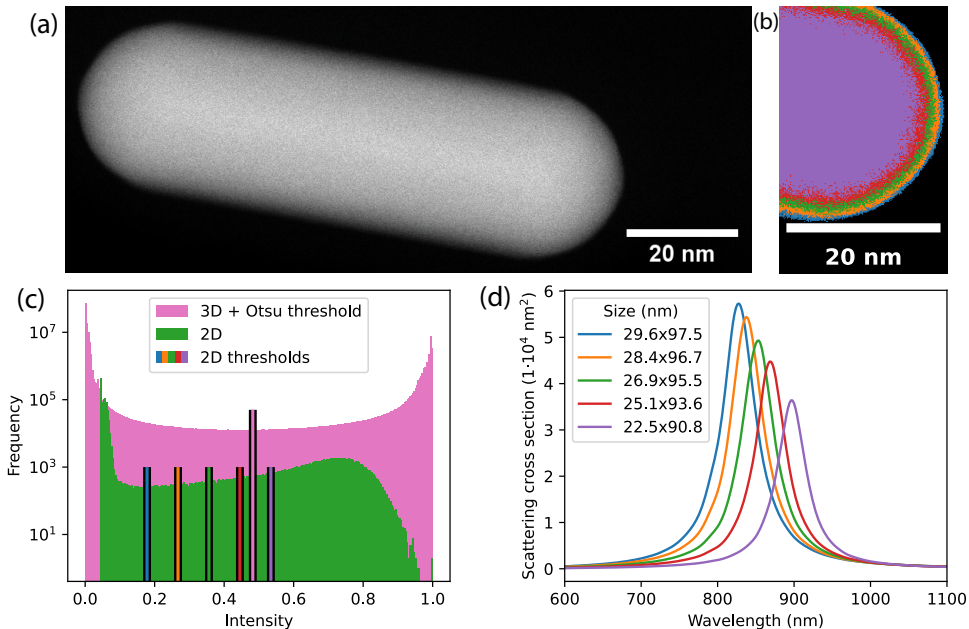


Figure 3.1: The effect of segmentation of high-angle annular dark-field scanning TEM (HAADF-STEM) projections and a reconstruction thereof. (a) HAADF-STEM image of a Au nanorod on a homemade holey-Cu TEM grid at  $0^\circ$  tilt angle. (b) Different segmentations of the tip for which the colors correspond to (c) and (d). (c) Pixel intensity distributions with corresponding Otsu thresholds or fractions thereof (0.5: blue, 0.75: orange, 1: green, 1.25: red, 1.5: purple) of reconstructed data from the complete 3D data set (pink) and the 2D data shown in (a) (green). (d) Simulated BEM scattering spectra of Au nanorods with sizes corresponding to the legend, resulting from the different segmentation thresholds in (b) and our 2D fitting algorithm (Section 3.A.2.5), discretized with  $dis_z = 3$  nm,  $dis_\theta = 1$  nm, and  $dis_\phi = 3$  nm and immersed in oil ( $n = 1.51$ ) using the gold dielectric function measured by McPeak et al. [80].

projection images in the tilt range of  $-77^\circ$  to  $+72^\circ$  (details in Section 3.A.2), which were subsequently reconstructed using the total variation minimization (TVM) algorithm. The influence of the choice of reconstruction algorithm and segmentation method will be detailed later. For now, the pink histogram in Figure 3.1c demonstrates another advantage of using ET in addition to providing the realistic 3D morphology: The separation of the background and foreground became significantly clearer after reconstructing an experimental tilt series of the Au nanorod, reducing the uncertainty in the segmentation process. Segmenting and fitting the 3D reconstruction of our experimental example resulted in a diameter of 30.0 nm and a length of 96.5 nm (AR=3.2). It should be noted that fitting the nanorod shape to the 3D data was done

for the sake of comparing the sizes to the 2D results. However, in the following we use the 3D output of the tomographic reconstruction directly. This approach becomes particularly important for simulations of PNPs with complex geometries, where it is not possible to guess a 3D shape from 2D images. To make the output of the tomographic reconstruction suitable for BEM simulations, the voxelized reconstruction needs to be transformed into a triangular surface mesh requiring segmentation of the particle as an intermediate step. In the following, we discuss considerations for an optimized workflow to achieve this.

### 3.3 Simulating realistic tomography data for benchmarking

To evaluate the importance of the possible errors that are introduced during the different steps along the way, we used a well-defined ground truth. For that, we simulated electron tomography data for a nanorod using the sizes from the fit to the experimental 3D data from Figure 3.1. Since electron microscopy data contains noise and image artefacts, we need to account for this when simulating the 2D projection images. The most prominent contributions to this are Gaussian blurring caused by defocus and astigmatism, and Poisson noise arising from the discrete nature of the recorded signal [112]. Additionally, the background signal from the sample support needs to be taken into account for a realistic representation of STEM images. The STEM images were simulated by forward projecting a voxelized model of a nanorod with the fitted sizes using the ASTRA toolbox 2.1.0 for the experimental tilt angles (details in Section 3.A.1) [115]. Then, for each 2D projection a Gaussian filter was applied to model blurring, followed by simulating the background signal from the sample support. To stay as close to experimental parameters as possible, we modeled the relative background signal level by calculating the mean of the background values for every experimental projection image from the Au nanorod from Figure 3.1. We did that by first removing the particle from the 2D image through segmentation using a threshold that made sure that the whole particle was removed and calculating the mean of the remaining image. Figure 3.2a shows that the background level increased with increasing tilt angle because of carbon contamination deposition throughout the experiment, which can be clearly observed when comparing a 2D HAADF-STEM image taken before and after the tilt series (see insets). The sharp increase at the first and last tilt angle can be attributed to detector shadowing. The estimated background level was added to each simulated projection image independently, and Poisson noise was applied on a pixel-by-pixel basis after manually tuning the scaling of the signal

for the particle and the background to match the experimental noise levels. To assess the result of noise addition to the simulated data, line profiles were compared between simulated and experimental data. The insets in Figure 3.2b show representative projections from both experimental and simulated electron tomography data and the extracted line profiles show a good qualitative match.




Figure 3.2: Simulating electron tomography data with realistic experimental input. (a) shows the mean values of the background of the data from Figure 3.1 at different tilt angles. The inset shows a projection image at 0° before and after tilting with a logarithmic intensity scale to show the background more clearly. (b) compares representative line profiles of experimental (blue) and simulated (orange) 2D projection data. The inset shows the corresponding projection images and lines along which the profiles were extracted.

Using the tomography data simulated for the ground truth shape, it becomes possible to compare different methods for the different steps in our processing pipeline. For a quantitative comparison we used the ground truth shape on a voxel grid with the same voxel size as used in the experiments as a reference for calculating the shape error  $E_S$  induced by the different choices in the data processing steps:

$$E_S = \frac{\sum |Vox_{sim} - Vox_{ref}|}{\sum |Vox_{ref}|} \cdot 100\% \quad (3.1)$$

The shape error takes into account misclassified voxels and therefore reveals shape deviations and possible misalignments of two structures even if the total volume is the same. In the case of our simulated data, there is no effect of misalignment but it needs to be kept in mind when using experimental data, for which careful object registration needs to be performed first for an accurate shape error calculation [49]. To calculate the shape errors, the reconstructions needed to be segmented. For the comparison here, we used the Minimum method [110], which outperformed other methods for the

nanorod shape as detailed below. A more in-depth discussion of the segmentation method is provided in Appendix 3.A.2.3 with Figures 3.A2 and 3.A3.

### 3.4 Comparing tomographic reconstruction algorithms

As a first step, we compared different pre-processing methods applied to the input projection images before performing tomographic reconstruction. It was observed that pre-processing had negligible influence on the final result due to the high signal-to-noise ratios for our data (details in Section 3.A.2.4). Interestingly, smoothing the input projection images resulted in the same marginally improved shape error as smoothing the 3D reconstruction as a whole in the case of a low noise reconstruction (see Figure 3.A4). Therefore, in the following discussion we used the unprocessed projection data. Next, three common iterative tomographic reconstruction algorithms were compared: expectation-maximization (EM) [116], simultaneous iterative reconstruction technique (SIRT) and total variation minimization (TVM) [109]. These algorithms utilize different assumptions about the reconstructed object: EM and SIRT algorithms produce maximum likelihood reconstructions in case the input data are coming from Poisson or normal distributions, respectively. Both of these algorithms minimize the discrepancy between the input data and the projection of the reconstructed object, and TVM incorporates an additional objective of minimizing intensity variations in the solution, thereby promoting smooth, piecewise-constant reconstructions.

Figure 3.3 shows the obtained shape errors with reference to the voxelized ground truth for a variety of different reconstructions for which the reconstruction method, the number of algorithm iterations, the object shape, and the angular sampling range were varied. Figure 3.3a displays the effect of the number of iterations, illustrated using the EM algorithm. Increasing the number of iterations from 15 to 25 decreased the shape error but more iterations led to its increase. This effect is common to the iterative algorithms that minimize the discrepancy between the reconstruction and the noisy input data. At lower iterations the algorithm converges to the true solution, but eventually the reconstruction is overfitted to experimental noise, and the error compared to the ground truth increases. Thereby, there is an optimal number of iterations depending on the noise level in the input data, which in our case was around 25 iterations [116]. The effect of the limited angular sampling range was evaluated by comparing the shape errors for the reconstructions obtained for different tilt ranges with 25 iterations of EM (Figure 3.3b). Non-surprisingly, the shape error increased when the number of tilt angles decreased. The ‘residual’ error of 1.18% for a full tilt range of  $\pm 90^\circ$  is a combination of the discrete tilt step and segmentation process [37].

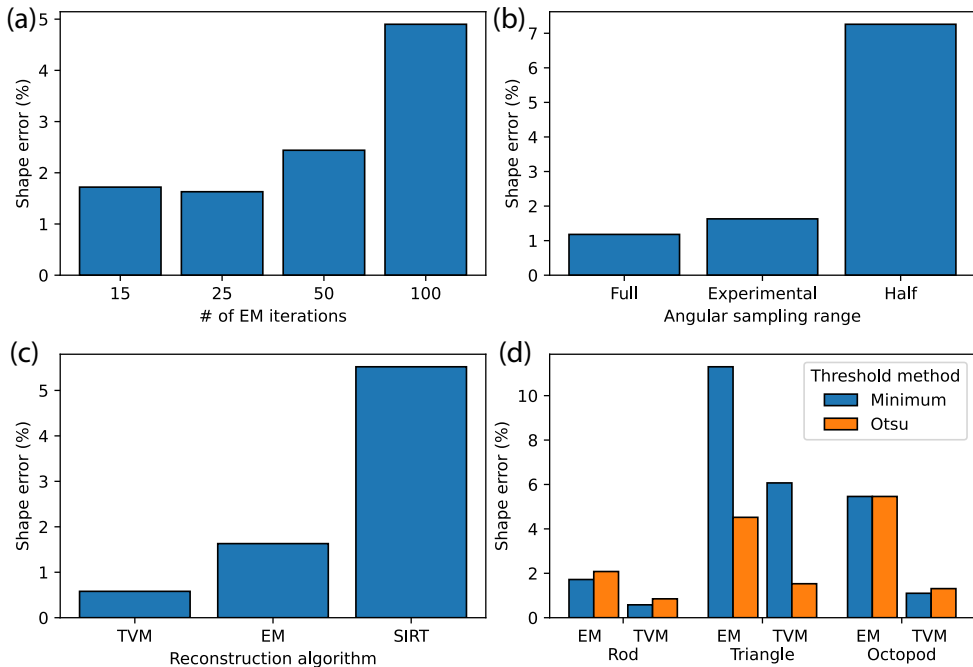


Figure 3.3: The effect of the number of iterations (a), the angular sampling range (b), and the reconstruction algorithm on the shape error (c) when reconstructing our simulated data from Figure 3.2b. (d) Comparison of EM and TVM for three different shapes using the Minimum and Otsu segmentation methods.

For the experimental tilt range of  $-77^\circ$  to  $+72^\circ$  the shape error increased marginally to 1.63%. However, when the tilt range was significantly decreased to  $\pm 45^\circ$ , the shape error severely increased to 7.26%.

When comparing different reconstruction algorithms, EM significantly outperformed the more commonly used SIRT algorithm (Figure 3.3c). This is not surprising since it is suitable for Poisson distributed data typical in STEM imaging, whereas SIRT is based on normally distributed data [109]. EM was even more outperformed by TVM with a remarkably low resulting shape error of 0.58% for a tilt range of  $-77^\circ$  to  $+72^\circ$ . This is expected, since TVM incorporates additional prior knowledge about the smoothness of the reconstructed object, which allows for compensating the noise and limited angular sampling range artefacts. The same conclusion was drawn for more complex reconstructed object shapes, such as a triangle and an octopod (Figure 3.3d). The triangle served as an example of a shape that is more susceptible to the limited angular sampling range artefacts, which stems from the alignment of the particle with respect to the tilt axis. Whereas for elongated shapes like nanorods,

the missing information can be reduced by positioning the nanorod perpendicular to the tilt axis as done here, a triangle cannot be rotated in a similar optimal manner. An octopod, on the other hand, is an example of a shape with smaller and sharper geometrical features. For both of these more challenging shapes, utilizing TVM led to the reconstructions with the smallest shape error similar to the nanorod case.

### 3.5 Segmentation strategies and threshold selection

Figure 3.3d also demonstrates that the choice of segmentation method becomes crucial when the limited angular sample range produces larger artefacts as is the case for the triangle. In the case of the nanorod or the octopod, using the Minimum or Otsu segmentation method resulted in similar shape errors, although the Minimum method slightly outperformed the Otsu one for both the TVM and EM reconstructions. However, for the nanotriangle the segmentation method had a significant influence. Using the Minimum threshold almost tripled the shape error compared to the Otsu method for the TVM and EM reconstructions. The reason behind this is detailed in Section 3.A.2.3. In short, the Minimum method used here calculates the minimum in the smoothed intensity histogram, which is much more sensitive to noise in the reconstruction and therefore produces less predictable segmentation results. The Otsu method, on the other hand, minimizes the inter-class variance, which is significantly more robust in the case of noisier and lower quality data. Hence, the Minimum method can be assumed to work less well for noisier data, which includes shapes that suffer from a larger influence due to a limited angular range, and should be applied to high signal-to-noise data only. It is advisable to look at the actual histograms to help with the judgement (see Figure 3.A3).

The 3D visualizations of the final segmented TVM reconstructions of the three simulated particle shapes are displayed in Figure 3.4a–c. The high quality of the reconstruction and segmentation as evidenced by the low shape errors is clearly visually reproduced. Figure 3.4d–i demonstrates why TVM led to a smaller shape error, in particular for the triangle and the octopod. Representative slices of reconstructions using either EM (Figure 3.4d–f) or TVM (Figure 3.4g–i) for the three different shapes are shown. Strikingly, the TVM reconstructions had a significantly higher subjective quality than the EM reconstructions, as they were less noisy and displayed a less significant effect of the limited angular sampling range. This is also reflected in the voxel intensity histograms, which displayed clearer separation of the foreground and background compared to EM (Figure 3.A4). Consequently, segmentation (segmented boundaries are displayed in light blue) was easier and more robust on the TVM data.

As a result, the quantitative shape errors obtained for the TVM reconstructions were surprisingly low even for the more challenging shapes.




Figure 3.4: 3D visualizations of the (a) nanorod reconstructed with TVM and segmented with the Minimum threshold, (b) triangle reconstructed with TVM and segmented with Otsu threshold and (c) the octopod reconstructed with TVM and segmented with the Minimum threshold. Slices through the (d-f) EM and (g-i) TVM reconstructions before segmentation highlight the differences in the reconstruction methods. The slices were taken along the blue planes in (a-c). The light blue outlines in (d-i) correspond to the segmentations mentioned above and the double-sloped colormap is used to aid in visualizing both background and foreground noise in these segmentations.

## 3.6 Generating surface meshes from voxel reconstructions

To proceed with electromagnetic simulations based on the BEM method, the obtained reconstructions need to be converted to surface meshes. One possible approach is to fit a 3D model of a particle to the reconstruction data as we did for the experimental data to obtain the length and diameter for our ground truth simulations (Figure 3.A5). However, this introduces an additional shape error because it is just an approximation of the shape. In fact, for the experimentally measured nanorod in Figure 3.1, the discrepancy between shape fitting and directly meshing of the particle resulted in a shape error of almost 5%. It is therefore beneficial to create the surface mesh from voxel data directly. The most popular algorithm for achieving this task is the marching cubes method [117], [118]. In this algorithm, segmented 3D data on a voxel grid are converted into a mesh by placing triangles at the boundary of the object with their orientations determined from the local arrangement of voxels in the segmentation. Surface meshing of the reconstructions did not result in significantly larger shape errors compared to Figure 3.3, see Section 3.A.2.7 and Figure 3.A7 for details. We observed small but noticeable differences in obtained shape error for different implementations of the marching cubes algorithm (see Section 3.A.2.7 for a full discussion). It should be noted that we had to slightly smooth the reconstructions with Gaussian of pixel size 1 to be able to create surface meshes for all reconstructions presented in Figure 3.3. Without smoothing, some of the created meshes contained otherwise holes, which could not always be fixed. Whereas we did not see an effect of smoothing of the reconstructed 3D data set for the less noisy reconstructions (Figure 3.A4), smoothing led to a significantly decreased shape error for the reconstructions performed by SIRT and 100 iterations of EM, which were noisier compared to the rest. In that case of a more limited angular tilt range of  $\pm 50^\circ$ , the actual missing information could non-surprisingly not be retrieved through smoothing (Figure 3.A8).

## 3.7 Mesh simplification and its effect on shape accuracy

The marching cubes algorithm usually produces a mesh with the same resolution as the input voxel data, which leads to a number of triangles on the order of  $10^6$  in our case. In fact, the 3D visualizations in Figure 3.4a–c are these surface meshes. Such a large mesh size makes it computationally intractable to perform electromagnetic simulations

[44]. For this reason, we used a mesh simplification algorithm that reduced the number of triangles to a user specified value [119]. After comparing several algorithms in terms of the shape error introduced by mesh simplification (see Figure 3.A7), we chose to use the so-called fast simplification algorithm, a quadric error metric-based algorithm, which iteratively removes mesh edges that contribute the least to the final simplification error. For this fast simplification algorithm, an aggression parameter needs to be chosen, which determines how aggressively faces are removed from the mesh. We found that 7 was a suitable aggression parameter (Figure 3.A6).

### 3.8 Propagation of morphological errors into optical simulations

The final test is to identify how the different processing steps influence the simulated scattering cross sections, which is displayed in Figure 3.5. Figure 3.5a displays the scattering cross sections and Figure 3.5b plots the corresponding spectrum errors (Equation 2.7 in Chapter 2) as a function of the shape errors with respect to the voxelized ground truth for the nanorod (same as reported in Figure 3.3). The ground truth for the spectrum error was based on a spherocylinder mesh with the dimensions of 30 nm x 96.5 nm and optimal discretization (see Figure 2.7 in Chapter 2), corresponding to 4960 triangles. For a direct comparison, all other meshes were simplified to the same number of triangles. Note that because of mesh simplification, even the ground truth model for spectrum error has a shape error of about 1%.

We first compared the simplified surface meshes based on the segmented reconstructions obtained with different tilt ranges, reconstruction algorithms and segmentation methods from Figure 3.3 for the nanorod shape. Both spectrum and shape errors for the majority of cases were very low, below 1% and 2%, respectively. This is because a nanorod is a simple, symmetric shape and different investigated data processing steps, such as reconstruction smoothing and mesh simplification, are effective in removing artifacts originating from noise and suboptimal reconstruction parameters. In turn, the remaining small shape deviations do not significantly influence the spectral response, and there is no clear correlation between the shape and spectral error in this regime. In contrast to the data processing parameters, limited input data, as in the case of strongly restricted angular range reconstruction (pentagon symbol in Figure 3.5b), led to significant shape and spectrum errors.

The fact that the spectrum errors for the different reconstruction and data processing parameters were mainly below 1% with our workflow demonstrates that our

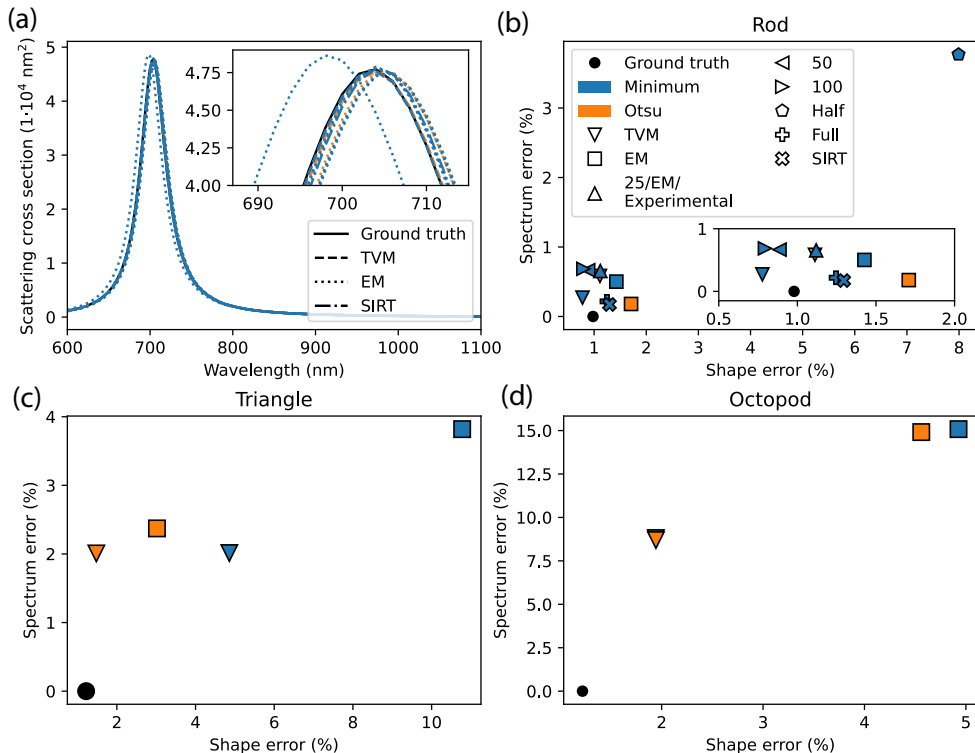


Figure 3.5: The effect of the shape error on the spectrum error. (a) Simulated BEM scattering spectra of meshes created by reconstructing, segmenting, smoothing, meshing, and simplifying simulated HAADF-STEM data of the rod. (b) Their corresponding shape and spectrum errors. (c) The shape and spectrum errors for the triangle meshes and (d) for the octopod meshes. It should be noted that all results from Figure 3.3 are included, which means that there are two different thresholded results for TVM (dashed line and triangle markers) and EM (dotted line and square markers), i.e. using the Minimum (blue) and Otsu (orange) methods. The numbered labels correspond to results that were reconstructed using a different number of EM iterations. 'Half' and 'Full' correspond to the used angular sampling range for the reconstruction. The labels '50' and '100' also correspond to the number of iterations for the EM reconstruction using the experimental angular range, which was the same as for 25 iterations, labelled here '25/EM/Experimental'.

meshing pipeline is rather robust and can create low spectrum errors even in the case of sub-optimal reconstruction choices. However, from Figure 3.3d we know that the nanorod is actually the most forgiving shape in terms of reconstruction and segmentation workflow. The situation is indeed different for the more challenging shapes of the triangle and octopod shown in Figure 3.5c and d (with their corresponding scattering spectra in Figure 3.A9a,b). In both cases we meshed the voxelized ground truth from

Figure 3.4b,c to use as the reference for the spectrum error. Same as for the nanorod shape, mesh simplification led to a small shape error of below 1%. For the ground truth mesh two general observations can be made for these more complex shapes. First, the higher shape errors compared to the nanorod shape led to significantly higher spectrum errors. Second, even with comparable shape errors, the spectrum errors were significantly higher for shapes with higher complexity. Whereas a 5% shape error for the nanotriangle still resulted in a spectrum error around 2%, for the octopod the spectrum error increased to 15% for a similar shape error. The reason is that more important morphological features are affected by the missing shape information. For triangles, the reconstructed shape inaccuracy mainly resulted in thickness variations (Figure 3.4e). For octopods, a higher shape error was connected to blunting of the tips, which blue-shifted and decreased the scattering cross section (Figure 3.A9b). Thus, the more complex the shape, the better the reconstruction needs to be for a successful electromagnetic simulation. In our comparison, TVM performed significantly better than other algorithms because of incorporating additional prior knowledge about the reconstructed object. A promising future direction could be employing reconstruction methods based on mesh representation [120], which would allow for minimizing shape errors stemming from mesh simplification.

### 3.9 Recommended workflow and best practices

The overall proposed workflow and main findings are summarized in Figure 3.6. For optimal results, and in particular for complex shapes, we recommend using TVM as a reconstruction algorithm together with Otsu segmentation, which proved to be more generally robust compared to other segmentation methods. To transform the segmented reconstruction into a surface mesh, the marching cubes algorithm performed well for all shapes analyzed here. We recommend smoothing the segmented reconstruction with a 1 px Gaussian kernel before meshing to create a watertight surface. We further recommend using the fast simplification algorithm with an aggression parameter of 7 to reduce the number of surface elements, which is needed to ensure that the electromagnetic simulations can be performed within feasible computation times.

For the simulations themselves, BEM provided the best combination of accuracy, speed, and predictable convergence behaviour for our purposes. Finally, special attention should be given to the dielectric function and accurate description of the local dielectric environment. This is underlined in Section 3.A.2.8, where we show preliminary data on a dark-field scattering measurement correlated to 2D HAADF-STEM data of an individual Au nanorod. By removing the uncertainty in morphology, the

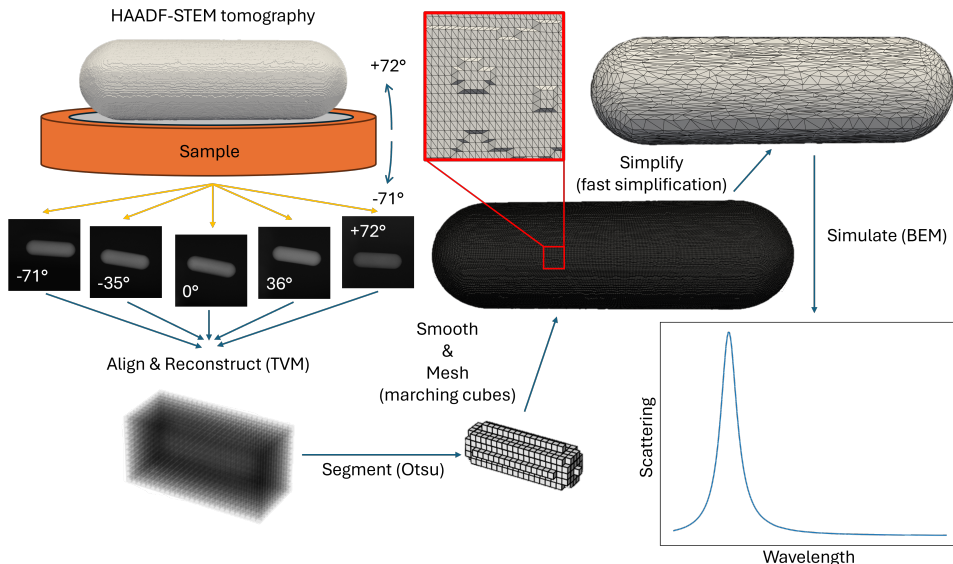


Figure 3.6: The overall proposed workflow in which HAADF-STEM tomography is performed on a nanorod. The resulting projections are aligned and reconstructed with TVM. These voxelized data are segmented with the Otsu method and smoothed before meshing with marching cubes. The resulting mesh is simplified using the fast simplification algorithm and the result is used as input for a simulation with BEM.

dielectric environment can be determined directly, enabling exact single-particle optical correlation.

### 3.10 Conclusion

In conclusion, performing electromagnetic simulations of plasmonic nanoparticles is an intricate interplay between different factors that play a role. In this chapter, we quantified possible error sources in a simulation workflow taking gold nanoparticles as an example system. We demonstrated that even parameters often considered secondary, such as meshing accuracy, segmentation thresholds, and reconstruction settings, can strongly influence the final morphology used as simulation input, and therefore the predicted optical response. In addition, the choice of metal dielectric function ideally reflects the experimental system, as it has a significant influence on the simulated optical cross sections, and for accurate results, the substrate needs to be included as well. A central challenge is that electromagnetic simulations require a well-defined 3D morphology, while the direct output of electron tomography is voxelized and contains noise, missing-wedge artefacts, and segmentation ambiguities. We showed

that morphologies obtained from electron tomography can effectively overcome this problem when processed carefully, and we developed an optimal workflow to turn a voxel-based reconstruction into a surface mesh by quantifying the introduced shape errors at each step. Although volume-based simulation methods can in principle use voxelized tomography output directly, the less predictable convergence behaviour of such approaches may not always be favourable. In contrast, converting the reconstruction into a surface mesh enables consistent and controllable simulation input once the morphology pipeline is properly tuned.

In terms of reconstruction algorithm, for all nanoparticle shapes, TVM significantly outperformed EM and SIRT. The optimal segmentation method depended on the nanoparticle shape: in general, the Otsu method was more robust and is likely the best method for most single nanoparticle shapes studied here, while for high-quality and low-noise data, the Minimum method performed slightly better, although it is more difficult to evaluate its performance without knowing the ground truth. We demonstrated that these different processing steps can alter the final input morphology, which in turn results in errors when simulating the optical response. Although slight smoothing of the reconstruction and the necessary surface mesh simplification could additionally lower the shape error of the object, we observed that the best approach is to enforce object smoothness during the reconstruction process rather than afterwards.

We also observed that the same shape error did not translate into a similar spectrum error for different nanoparticle shapes, in particular when high-curvature features are affected by the shape inaccuracies. The discussed topics in our work can therefore help to achieve more accurate simulations by minimizing artificial discrepancies stemming from sub-optimal morphology retrieval, thereby potentially allowing for a more accurate retrieval of nanoparticle morphology from optical data alone. Similar considerations are valid for correlating electron-based spectroscopies with electron tomography data, and our workflow can be applied in that case as well.

Taken together, Chapters 2 and 3 establish the complete simulation foundation required for quantitative structure–property correlations. With the electromagnetic solver validated, its numerical behaviour charted, and the morphology pipeline benchmarked and optimized, both sides of the simulation framework are now under control. Equally important, we have a quantitative understanding of the uncertainties that remain at each stage of the workflow. This positions us to apply the correlative framework to experimental single-particle data with confidence that remaining discrepancies between measured and simulated spectra arise from genuine physical effects rather than artefacts of numerical settings or morphology retrieval. Obtaining such reliable single-particle scattering data is the task of the next chapters.

## 3.A Appendix

### 3.A.1 Experimental methods

#### 3.A.1.1 Sample preparation

The Au NRs were purchased from Nanoseedz (NR-20-780, OD=5). According to the specifications the rods had an average diameter of 20 nm and an average length of 75 nm (AR=3.8) with a shape monodispersity of > 95%. After sonicating this solution for 30 minutes, 5  $\mu$ L was drop casted on a homemade holey-C Cu TEM grid. After 1 minute the remaining droplet was removed via absorption using a Kimtech paper towel.

#### 3.A.1.2 Electron tomography experiment

The electron tomography experiment was performed with a Thermo Fischer Scientific TitanCubed 300 kV D3213 SuperTwin. The X-Feg gun was operated at 300 kV to project the sample in STEM mode with a D-Corr probe aberration corrector onto a Fischione Model 3000 HAADF detector. A camera length of 115 mm and a magnification of 640000x were used to acquire images. First, images were acquired with alpha tilts from  $-77^\circ$  to  $-2^\circ$  with  $3^\circ$  steps. Subsequently, images were acquired with alpha tilts from  $0^\circ$  to  $+72^\circ$  with  $3^\circ$  steps. Also, an image at  $+1^\circ$  was acquired, which resulted in a total of 52 angles.

### 3.A.2 From electron tomography to mesh

#### 3.A.2.1 Pre-processing electron tomography data

Before performing reconstructions, it is possible to denoise and/or sharpen the input 2D projection images. A plethora of algorithms is available to do this, but here we limited ourselves to the ones available in the “restoration” module from the Python-based scikit-image library. The specific ones that were used are: bilateral denoising, non-local means denoising, total variation denoising using split-Bregman optimization or Chambolle’s algorithm, wavelet denoising, Richardson–Lucy deconvolution, the normal Wiener–Hunt and unsupervised Wiener–Hunt deconvolution.

### 3.A.2.2 Electron tomography reconstructions

Before performing reconstructions on the experimental data, the stack of 2D projection images was aligned using an algorithm based on cross-correlation, cropped, and normalized globally. Then, the following corrections were performed: background removal based on subtracting the median of the intensity histogram, detector shadowing compensation based on ensuring the preservation of total intensity between different projections, and a cupping correction algorithm of Van Der Broek et al. with  $I_0$  parameter manually adjusted to result in uniform intensities in the reconstruction [37]. As a double-check, the slices were realigned again using cross correlation, before performing TVM reconstructions using the Chambolle–Pock algorithm implemented using the ASTRA Toolbox and PyTorch library on a computer with a NVIDIA Titan RTX graphical card with 24 GB of GPU memory [115]. The value of the regularization parameter was optimized manually to result in subjectively smooth reconstructions. For the simulated data no cropping was needed and neither detector shadowing compensation nor cupping corrections were applied as these effects were not included when creating the simulated projection data. EM and SIRT reconstructions were performed on a computer with a NVIDIA GeForce RTX 3050 graphical card with 8 GB of GPU memory using ASTRA Toolbox 2.1.0.

### 3.A.2.3 Segmentation of reconstructions

After reconstructing the electron tomography data, the 3D data set needs to be segmented to distinguish the object of interest, i.e. the nanoparticle, from the background. A plethora of algorithms is available to do this, but here we limited ourselves to the ones available in the “filters” module from the Python-based scikit-image library. These are all global segmentation methods, in which a single intensity threshold value is applied to the whole reconstruction. For the exact differences and an extensive comparison, we refer the reader to the work of Glasbey et al. [110].

Only after segmentation is it possible to calculate the shape error (see Equation 3.1 in the main text). Since we knew the ground truth shape of the simulated data, we could calculate the shape error for the 3D data sets obtained for the different pre-processing (Section 3.A.2.1) and segmentation methods. The effect of the reconstruction algorithm is discussed in the main text. To compare the shape errors with respect to the ground truth, all restored data in Section 3.A.2.1 were reconstructed with 15 iterations of EM and segmented by the Minimum, Otsu, Li, and Mean threshold method. We found that these four segmentation methods resulted in significantly lower shape errors compared to the other methods available in the filters module. Although the TVM

reconstruction algorithm resulted in generally lower shape errors, the trends with respect to nanoparticle shape and angular sampling range were the same as for EM. Since EM is considerably faster, we used it for the comparisons here. Figure 3.A1 displays the obtained shape errors with respect to the ground truth for the different pre-processing and segmentation methods. Note that the shape error for the Minimum threshold in Figure 3.A1a is the same as the 15 EM iteration shape error in Figure 3.3a in the main text. From Figure 3.A1 it can be seen that, segmenting the reconstruction with the Minimum threshold most often resulted in the lowest shape error for our simulated nanorod for different pre-processing methods although the improvement with respect to the Otsu threshold was not significant for those cases. However, it was less robust than Otsu as it sometimes resulted in a significantly higher shape error, such as for the data deconvoluted with the Richard-Lucy algorithm, for which the segmentation failed, resulting in a 100% shape error.

Interestingly, most pre-processing methods did not lead to a reduced shape error compared to the non-processed data (Figure 3.A1a). The lowest shape error of 1.17% was obtained for deconvoluting the projections with the normal Wiener–Hunt algorithm, which is slightly lower than without pre-processing the data (1.72%). The main reason is likely that our data had a high signal to noise ratio but some of the methods might also need additional optimization, which we did not do here. To figure out what influenced the shape error for the different pre-processing steps it is helpful to look at the intensity distributions of both the initial projections and the reconstructed data (Figure 3.A2). For the pre-processing method that led to significantly worse shape errors, i.e. Richardson–Lucy deconvolution, the intensity distribution was severely influenced by the pre-processing. The foreground peak was heavily shifted to the left and almost coincided with the background peak (left in Figure 3.A2g). However, the non-local means denoised data also had a higher shape error (Figure 3.A1c), although not as dramatic. The intensity distribution of the reconstruction exhibited a shoulder next to the foreground peak (right in Figure 3.A2c). A similar observation can be made for the bilaterally denoised data, where a splitting of the background peak can be observed (Figure 3.A2b). This underlines once more that the intensity distribution heavily affects the quality of segmentation.

For these outliers with clearly different intensity histograms the shape error increased for the data segmented with the Minimum method compared to the Otsu method. A similar observation can be made for the data in the main text (Figure 3.4): the EM distribution of the triangle was flatter than those for the rod and octopod (Figure 3.A3), probably arising from the artefacts introduced by the limited angular sampling range. Moreover, the TVM distribution contained more frequency fluctua-

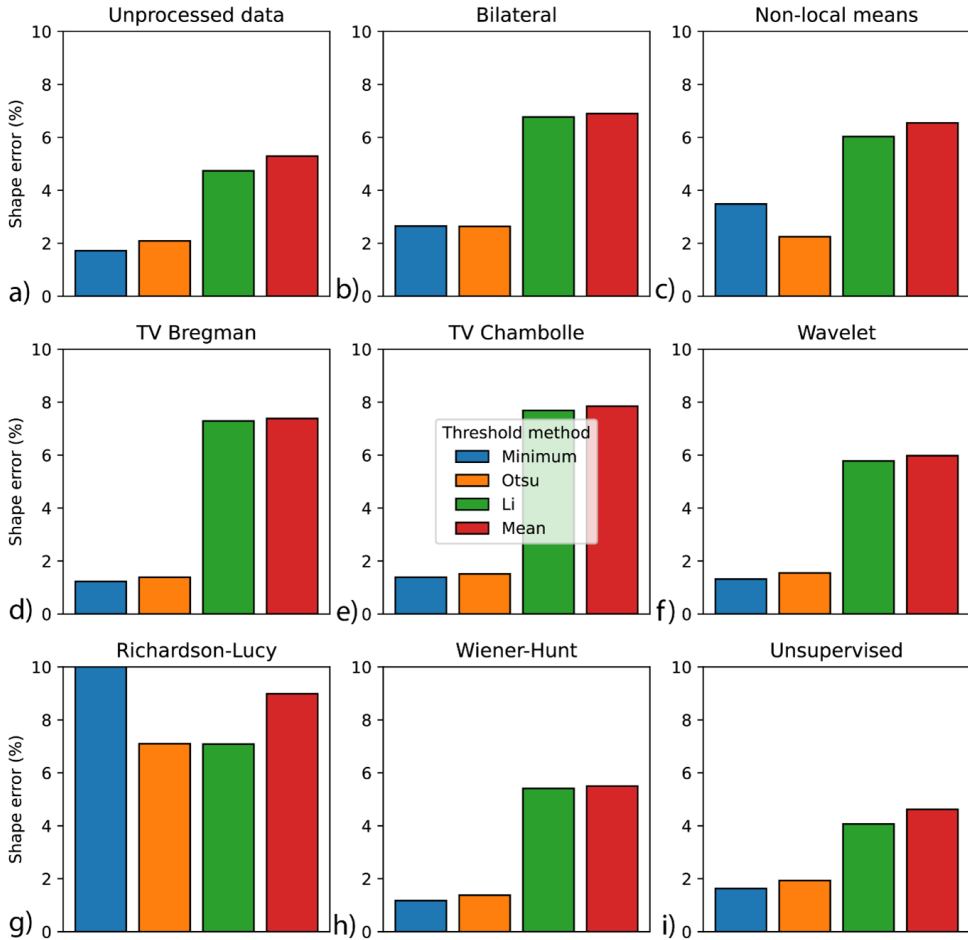


Figure 3.A1: Shape errors using different segmentation thresholds (see legend) for data pre-processed with the filters mentioned in Section 3.A.2.1. The shape error for the Minimum threshold in (g) exceeded the range because it was 100%, which indicates a failed segmentation.

tions. Since the Minimum method is susceptible to these spurious minima [110], this might lead to lower segmentation quality. It seems that the fewer frequency fluctuations are present in the intensity distribution, the closer the Minimum and Otsu threshold values are. Although this is a preliminary indication of which threshold method to use, i.e. Minimum for low-noise and Otsu for high-noise data, more experiments need to be performed to confirm this and one should always pay careful attention to the quality of preprocessing, reconstructing, and segmenting data. It is advisable to look at the actual histograms to help with the judgement.




Figure 3.A2: Intensity distributions of the projections from Figure 3.A1 (green) and their reconstructions (pink). (a) Unprocessed data, (b) bilateral denoising, (c) non-local means denoising, total variation denoising using split-Bregman optimization (d) or Chambolle’s algorithm, (f) wavelet denoising, (g) Richardson–Lucy deconvolution, the normal Wiener–Hunt (h) and unsupervised Wiener–Hunt deconvolution (i).

### 3.A.2.4 Smoothing before and after reconstruction

Next to the pre-processing methods described in Section 3.A.2.1, the data can also be smoothed before or after reconstruction. As an example, we first smoothed the projection input images with a Gaussian with a standard deviation of 3. After reconstructing the smoothed projection images with 15 iterations of EM and segmenting the reconstruction with either the Minimum or Otsu threshold, we obtained in both cases a shape error of 1.32%, which was lower than the unprocessed data (1.72%).

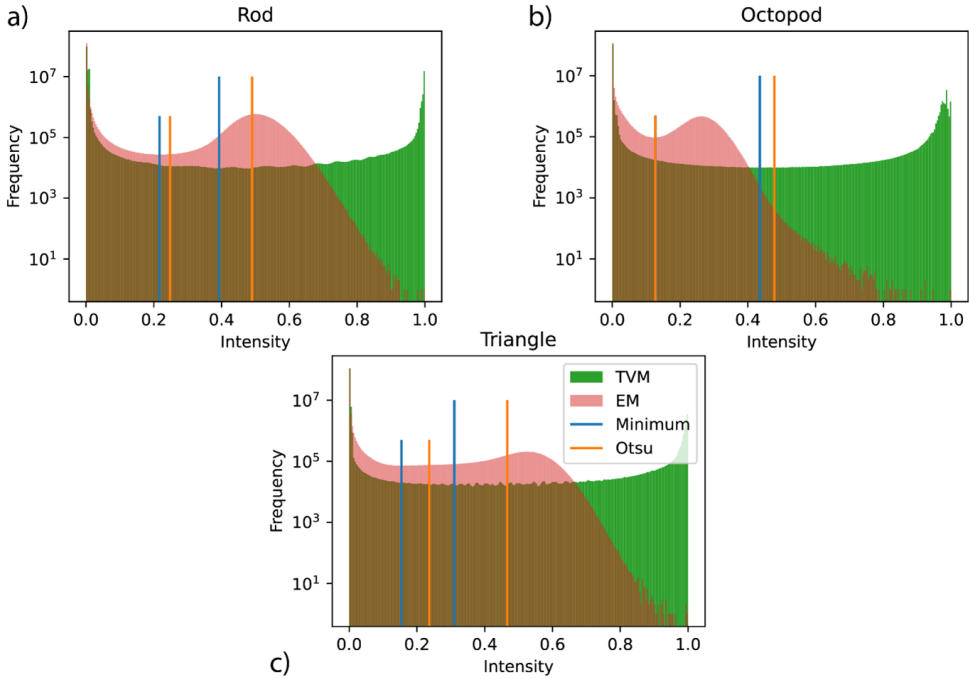


Figure 3.A3: Intensity distributions for the data in Figure 3.4 in the main text with their corresponding threshold values as vertical lines using either the Minimum or the Otsu method. The taller lines correspond to TVM and the shorter lines correspond to EM.

However, it is also possible to smooth the 3D reconstruction itself, i.e. without pre-processing the data. Figure 3.A4 shows that the shape error went down to a similar value when smoothing the 3D reconstruction compared to smoothing the projection images before reconstruction. It should be noted, however, that the level of smoothing and the optimal number of iterations for EM are coupled. Therefore, for noisy data it can be beneficial to modify both in an iterative manner, which is beyond the scope of this work [121].

It should be noted that there are more pre- and post-processing image modification options than explored here. The ultimate choice is highly connected to the input data. For our high signal to noise data, pre- and post-processing did not significantly influence the shape error. However, for noisier data, it can be expected that the shape error can be significantly more reduced when smoothing the input or reconstruction data or when using the other pre-processing methods.




Figure 3.A4: Shape error as a function of the standard deviation used for Gaussian smoothing after reconstructing the originally unprocessed simulated data.

### 3.A.2.5 Fitting algorithms for shape parametrization

For simple structures it is possible to fit their respective geometrical objects to the data by parametrization. This can be done for both 2D and 3D data by carefully selecting the outline of the thresholded data using binary erosion. Binary erosion removes the outer layer of pixels/voxels from segmented data. By subtracting this from the original segmented data, the outline is obtained. Subsequently, the residuals between an initial guess and the outline can be computed. This allows a least squares algorithm to establish a loss function and to find a local minimum of the related cost function. The result is a parametrized prediction of a geometrical object that matches the measured structure. Such a fitting procedure is an alternative to meshing the segmented data, but it is based on prior knowledge since the underlying geometrical shape needs to be describable. In our work, the shape fitting served as a comparative method to evaluate how large the shape error was when a shape was assumed and was used to define the ground truth shape for our tomography data simulations based on the experimentally fitted shape.

In this paper this method was applied to a rod-like object, which can be approximated with a stadium in 2D and a cylinder with hemispherical caps in 3D. The

parametrization for a stadium is given by:

$$F(x, y) = \begin{cases} y^2 + \left(x - \frac{l-r}{2}\right)^2 - r^2, & x > \frac{l-r}{2}, \\ y^2 + \left(x + \frac{l-r}{2}\right)^2 - r^2, & x < -\frac{l-r}{2}, \\ y^2 - r^2, & -\frac{l-r}{2} \leq x \leq \frac{l-r}{2}. \end{cases}$$

where  $r$  and  $l$  are the length and radius of the stadium, respectively, and  $x$  and  $y$  are the pixel coordinates. Therefore, the error function becomes:

$$e(x, y) = \begin{cases} y^2 + \left(x - \frac{l-r}{2}\right)^2 - r^2, & x > \frac{l-r}{2}, \\ y^2 + \left(x + \frac{l-r}{2}\right)^2 - r^2, & x < -\frac{l-r}{2}, \\ y^2 - r^2, & -\frac{l-r}{2} \leq x \leq \frac{l-r}{2}. \end{cases}$$

where  $x$  and  $y$  are the predicted pixel coordinates. The minimization of these residuals only works properly if the input data is centered around the origin. This can be accomplished by simple translations with mean  $x$  and  $y$  values. However, also rotation freedom should be implemented during optimization. For that we use the 2D rotation matrix:

$$R(\theta) = \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix}.$$

where  $\theta$  is the angle of rotation.

Similar functions can be defined for a 3D capsule where we take  $z$  as the long axis of the rod:

$$F(x, y, z) = \begin{cases} y^2 + z^2 + \left(x - \frac{l-r}{2}\right)^2 - r^2, & x > \frac{l-r}{2}, \\ y^2 + z^2 + \left(x + \frac{l-r}{2}\right)^2 - r^2, & x < -\frac{l-r}{2}, \\ y^2 + z^2 - r^2, & -\frac{l-r}{2} \leq x \leq \frac{l-r}{2}. \end{cases}$$

$$e(x, y, z) = \begin{cases} y^2 + z^2 + \left(x - \frac{l-r}{2}\right)^2 - r^2, & x > \frac{l-r}{2}, \\ y^2 + z^2 + \left(x + \frac{l-r}{2}\right)^2 - r^2, & x < -\frac{l-r}{2}, \\ y^2 + z^2 - r^2, & -\frac{l-r}{2} \leq x \leq \frac{l-r}{2}. \end{cases}$$

Therefore, it is only needed to provide rotational freedom along the other two axes:

$$R_x(\theta) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos \theta & -\sin \theta \\ 0 & \sin \theta & \cos \theta \end{bmatrix}$$

$$R_y(\theta) = \begin{bmatrix} \cos \theta & 0 & \sin \theta \\ 0 & 1 & 0 \\ -\sin \theta & 0 & \cos \theta \end{bmatrix}$$

We performed the fitting procedure on our experimental data in the main text. After reconstructing the experimental data with TVM using 1000 iterations and thresholding the result with the Minimum threshold the above-described 3D fit was performed. The result was used to rotate the experimental data before calculating the shape error using the ground truth as a reference. Figure 3.A5 shows a slice of the reconstruction with the corresponding fit in red. Slight local shape mismatches could already be observed within this slice which can be an explanation for the increased shape error of almost 5%. This is expected since actual synthesized nanorods follow crystallographic constraints and kinetic restrictions during the synthesis process and are therefore not perfect spherocylinders.

### 3.A.2.6 Meshing segmented reconstructions

As described in the main text, the segmented 3D data needed to be transformed into a surface mesh for the BEM simulations. Here, we elaborate on the effect on the shape error for the different steps. We limit ourselves to the unprocessed but reconstructed (15 iterations of EM) and segmented (Minimum threshold) simulated data. For surface meshing, we needed to ensure that the segmented 3D data were fully connected and unconnected fragments (normally single pixels stemming from background noise) were therefore removed from the data set. We did this by selecting the largest connected component in the voxelized data. This slightly lowered the shape error from 1.72% to 1.71%. On these data, the marching cubes algorithm to




Figure 3.A5: Slice of reconstruction from the experimental data with TVM using 1000 iterations and segmentation using the Minimum threshold in blue and the fit result from the 3D fitting algorithm in red.

create a triangular surface mesh from the “measure” module from the scikit-image library in Python was performed using Lewiner’s method. To compare the introduced shape error after meshing, the meshed output needed to be transformed back to a voxel grid. To do that, we applied an algorithm based on the VTK library, which creates a voxel grid and finds the voxels falling inside the specified mesh. We checked that this revoxelization resulted in a negligible shape error of 0.04% by meshing the ground truth and revoxelizing it. Revoxelizing the unprocessed but reconstructed (15 iterations of EM), segmented (Minimum threshold), and meshed (Lewiner’s method) simulated data, increased the shape error to 1.95%. This is an intermediate shape error, as the meshes still needed to be simplified (Section 3.A.2.7). We compared the marching cubes algorithm to the built-in isosurface function in the PyVista library, which is a wrapper of `vtkMarchingCubes` that uses Lorensen’s method. The difference between these two methods is well explained by Chen and Zhang [122]. The marching cubes algorithm from scikit-image also allows to use Lorensen’s method but gave slightly different final shape errors than the PyVista library (Figure 3.A7). All three algorithms took similar times to run. For these reasons, we chose the algorithm that gave the lowest shape error per dataset to create the meshes for BEM.

### 3.A.2.7 Simplifying meshed reconstructions

Subsequently, mesh simplification was performed using different simplification algorithms. On the one hand, existing implementations are based on the quadric error metrics algorithm described by Garland and Heckbert [119]:

the `meshing_decimation_quadric_edge_collapse` function from PyMeshLab, the Fast-Quadric-Mesh-Simplification Wrapper, which claims to be an optimized version of the `decimate` function in PyVista, which by itself is a wrapper function of `vtkQuadricDecimation`. And on the other hand, there is an implementation of the algorithm described by Schroeder, Zarge, and Lorensen [123]: the `decimate_pro` function in PyVista. Again, an intermediate shape error was obtained but now there was also a significant difference in running time between the algorithms. Since the fast simplification algorithm had both the lowest shape error (1.8%) and running time (1.79 s), we decided to continue with this algorithm for the remainder of this section. This was also the algorithm that was used to create the surface meshes that were simulated with BEM to obtain the scattering cross sections in Figure 3.5 of the main text. It is remarkable that applying the fast simplification algorithm decreased the shape error from 1.95% to 1.8%. This might be a result of the effective smoothing applied by the mesh simplification, which compensates for the noise in the mesh obtained from the original reconstruction.

An important parameter to set for the fast simplification algorithm is the aggression, which controls how aggressively the mesh is decimated. The above-mentioned reduced shape error of 1.8% was obtained with an aggression of 7. The higher the aggression, the more surface faces are removed from the mesh in a single step. Figure 3.A6 shows the effect of the aggression parameter on the shape error for simplifying the unprocessed but reconstructed (15 iterations of EM), segmented (Minimum threshold) and meshed (marching cubes) simulated data to different numbers of faces. It is clear that the aggression became increasingly more important when simplifying the mesh to a low number of faces. An aggression of 7 seemed to give the most accurate and robust results.

Unfortunately, we noticed that mesh simplification of the unprocessed reconstructions sometimes resulted in non-watertight surface meshes. Although the holes in the meshes were small, BEM simulations could not be performed for some of the meshes with holes. Therefore, smoothing with a Gaussian of 1 px was applied on the reconstruction before creating the mesh which prevented the subsequent creation of holes in the mesh.

Figure 3.A7 displays the effect on the shape error of the meshing pipeline discussed

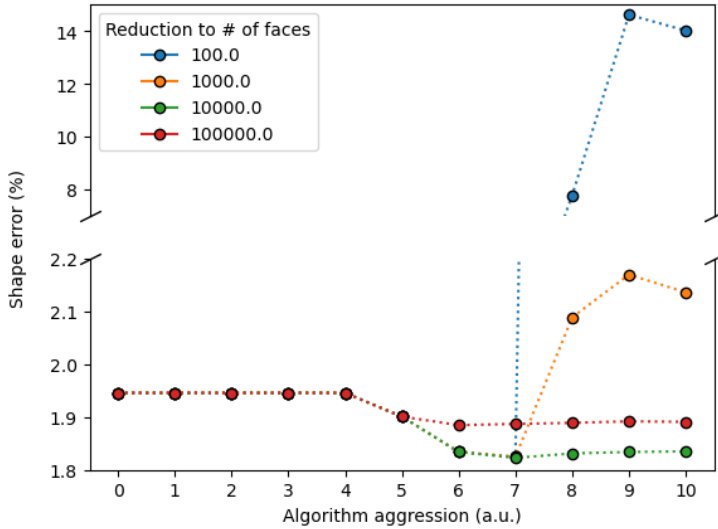


Figure 3.A6: The effect of the aggression parameter in the fast-simplification algorithm on the shape error of the unprocessed but reconstructed (15 iterations of EM), segmented (Minimum threshold) and meshed (marching cubes) simulated data for simplifying to different number of faces. The original mesh contained 995606 faces.

above. Note that the titles of the plots in Figure 3.A7 match the x labels of Figure 3.3 in the main text and that the legend refers to the explored marching cubes algorithms. Moreover, the final shape error is displayed, i.e. after smoothing, meshing, and simplifying. No significant shape error difference was obtained when using different meshing algorithms and in general the shape error remained more or less constant. However, it is remarkable that for the data reconstructed with 100 iterations of EM (Figure 3.A7d) or with SIRT (Figure 3.A7h) a significant drop in shape error was obtained by either smoothing or simplification. To decouple this effect, the shape errors were tracked for each step for the data using the meshing algorithm that resulted in the lowest shape error for each dataset (Figure 3.A8). It becomes clear that smoothing impacted the shape errors for both datasets most (Figures 3.A8d and h). Surprisingly, for the data with 100 iterations of EM this is not reflected in the spectrum error when a BEM simulation was performed with the obtained mesh (Figure 3.5 in the main text).

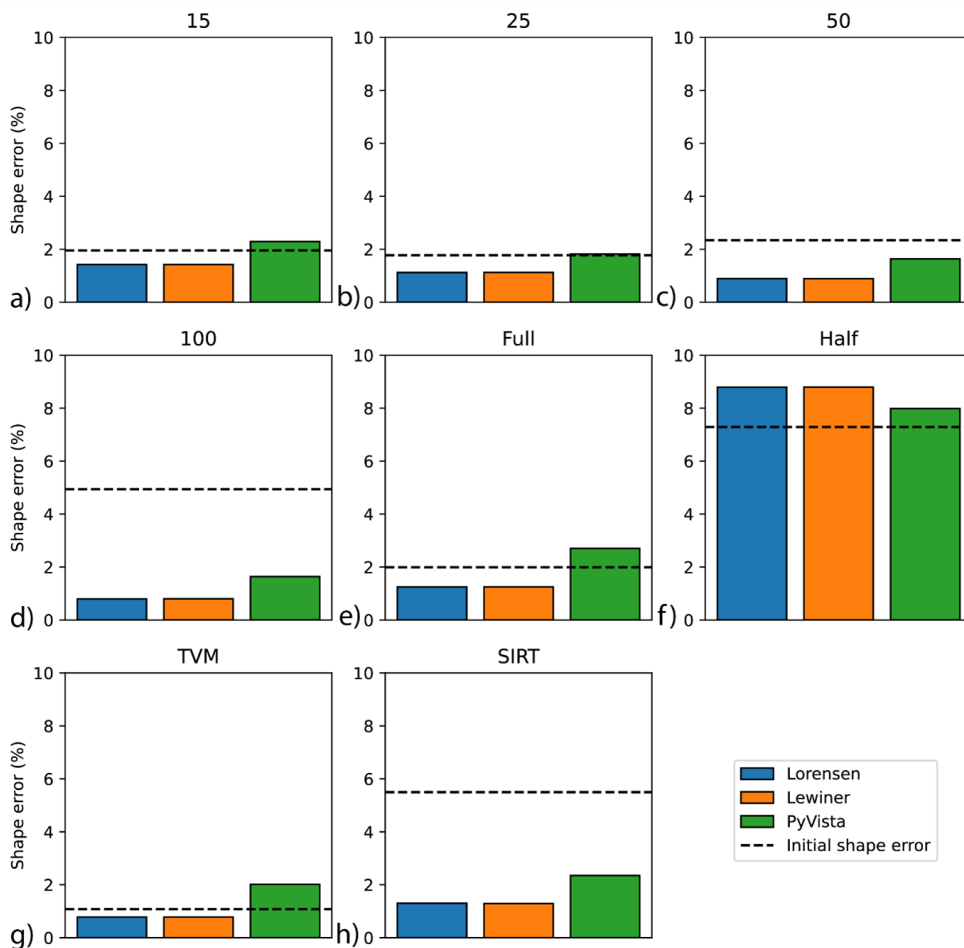


Figure 3.A7: The effect of the marching cubes algorithm on the final shape error for the data reported in Figure 3.3 of the main text. The titles of the subplots match with the x labels of Figure 3.3 in the main text. The labels Lewiner and Lorenson refer to different methods of the marching cubes function of the measure module from the scikit-image library in Python. PyVista refers to the `contour` function from that library, that should also use Lorenson’s method as it is a wrapper of `vtkMarchingCubes`.

### 3.A.2.8 Initial correlation of optical measurement to 2D HAADF-STEM data

The final quantitative correlation is a two-sided problem: unravelling the exact morphology (tackled here) and dielectric environment (now retrievable with our workflow). To demonstrate that these two unknowns need to be tackled independently, we show preliminary data on a darkfield scattering measurement correlated to 2D

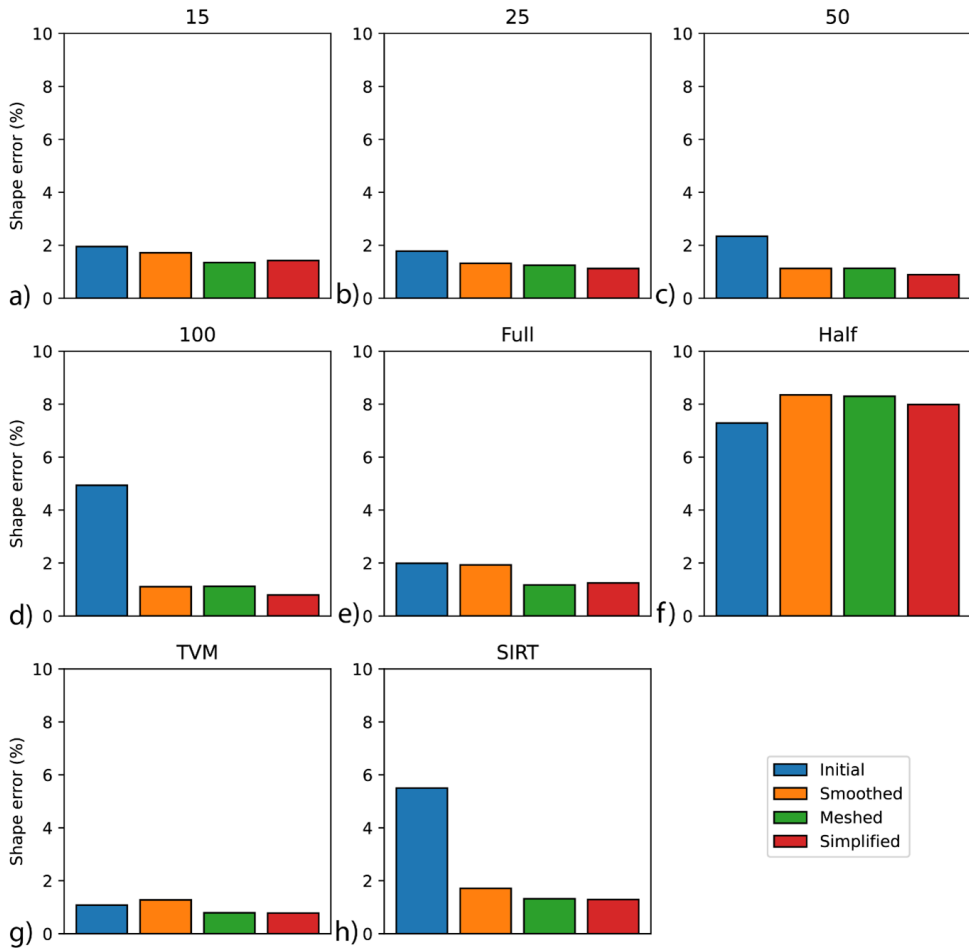


Figure 3.A8: The change of shape error throughout the processing pipeline for the data reported in Figure 3.3 of the main text. The label *initial* refers to the segmentation errors from Section 3.A.2.3. The remaining curves show smoothing, meshing, and simplification steps.

HAADF-STEM data of an individual Au nanorod (Figure 3.A10).

To show the effect of both thresholding and dielectric environment, we performed simulations with different combinations of them based on the experimental setup and STEM image. To retrieve the NR shape, a relative fraction of the Otsu threshold was swept. All of those thresholds look like reasonable values for segmentation based on the inset and histogram (not shown). In addition, an effective medium approach was used in which air/vacuum ( $n = 1$ ) and  $\text{SiO}_2$  ( $n = 1.5$ ) were used in different fractions. Figure 3.A11 shows that the measurement could in principle be matched

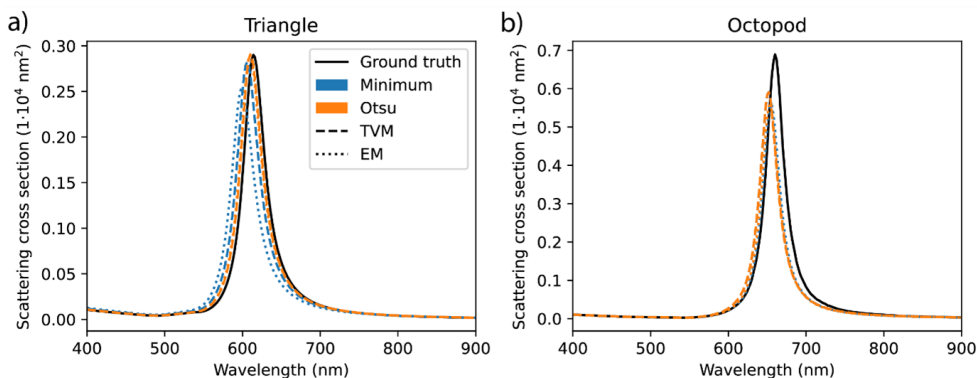


Figure 3.A9: Simulated BEM scattering spectra of meshes created by reconstructing, segmenting, smoothing, meshing, and simplifying simulated HAADF-STEM data of the triangle (a) and octopod (b). These spectra were used to calculate the spectrum errors in Figure 3.5c and d of the main text.




Figure 3.A10: 2D HAADF-STEM image of a Au nanorod.

with simulated data by tuning either the threshold (Otsu fraction) or the dielectric environment (effective medium approach), strengthening our point that they need to be tackled independently.

Although we were not able to obtain a full tomography data set for this Au nanorod, we can make a more accurate estimate of the shape by using our tomography simulations, such that we are only left with tuning the dielectric environment with the effective medium approach. For that we have to make the link between 2D and 3D data. We first aimed to find the optimal threshold that keeps the 2D size of the

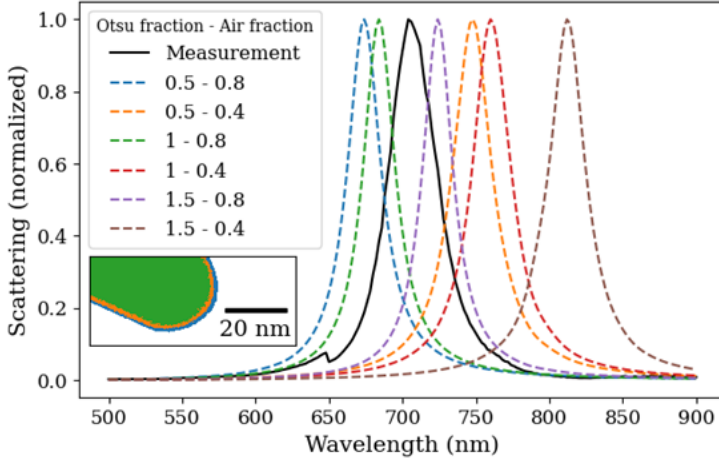


Figure 3.A11: Scattering measurement (solid line) of the Au nanorod in Figure 3.A10 and simulations with different Otsu and air fractions. The inset shows the segmentation results of different Otsu fractions.

rod as close as possible to the 3D ground truth. To this extent, we applied different threshold values on a simulated projection from our ground truth (see Figure 3.2b in the main text). After that we applied a 2D fitting algorithm to the data to determine the length and the diameter of the rod (Section 3.A.2.5), from which we can calculate the aspect ratio (Figure 3.A12).

It can be observed that the aspect ratio matches that of the ground truth ( $AR = 3.21$ ) most for an Otsu fraction of 0.36 ( $AR = 3.24$ ). To confirm that the aspect ratio is a good measure for this study, we performed BEM simulations with the rods obtained from thresholding with the different Otsu fraction values and compared them to the ground truth by means of calculating the spectrum error (Equation 2.7 in Chapter 2). Figure 3.A13 shows that indeed a lower aspect ratio results in a lower spectrum error.

Therefore, we chose to use an Otsu fraction of 0.36 on the measured 2D HAADF-STEM data as well (Figure 3.A10). Subsequently, we obtain a good match with the measurement when we used an air fraction of 0.65 (Figure 3.A14). This is a realistic value, which can be explained by describing the experimental setup: The AuNR that was measured was dropcasted on a  $\text{SiO}_2$  TEM grid, which was inverted onto a glass cover slip such that the rod was sandwiched between two glass surfaces. Our initial guesses in Figure 3.A11 were accounting for either one side of glass (air fraction = 0.4), or two sides of glass (air fraction = 0.8). Since we do not know the exact distance of the rod to the cover slip it could very well be something in between our initial guesses, like the air fraction of 0.65 we found matching our measurement.




Figure 3.A12: Correlation between the Otsu fraction used for segmentation and obtained aspect ratio of the simulated Au nanorod. The inset shows a zoomed part of the graph.




Figure 3.A13: Correlation between the aspect ratio and the spectrum error obtained by performing BEM simulations. The inset shows a zoomed part of the graph.




Figure 3.A14: Scattering measurement (solid line) of the Au nanorod in Figure S11 and a simulation (dashed line) using an Otsu fraction of 0.36 and an air fraction of 0.65.



## Chapter 4

---

# Cast Shadow – Quasi-Darkfield Microspectroscopy

*Cast Shadow – to send light or shadow in a particular direction*

**Abstract:** This chapter develops the optical framework required for quantitative single-particle scattering spectroscopy by introducing a quasi-darkfield microspectroscopy setup that selectively broadcasts the scattered light of individual nanostructures while suppressing background illumination. Building on a reflectance-mode geometry with a pupil-plane central stop, we treat illumination and detection as coupled angular filters that can be optimized jointly through beam shaping, back-focal-plane masking, and spatial filtering at the image plane. By modelling nanoparticle radiation patterns and propagating real beam parameters through the system, we establish design rules for maximizing collection efficiency and signal-to-noise ratio (SNR) across the visible spectrum. The chapter further outlines procedures for wavelength-dependent instrument-response calibration, alignment and focusing strategies, and substrate-index matching, all of which are essential for obtaining reproducible, quantitative spectra. Together, these elements define a robust and highly tunable measurement platform for the single-particle optical studies presented in the following chapter.

The content of this chapter is in preparation in M. Dieperink, P. Spaeth, W. Albrecht, "Decoupling Geometry and Environment in Plasmonic Nanorods via Correlative Optical–Electron Microscopy".

## 4.1 Introduction

The preceding chapters established the simulation half of the correlative framework: a validated Boundary Element Method (BEM) solver, a certified morphology-reconstruction pipeline, and a quantitative understanding of how uncertainties in geometry and dielectric environment propagate into computed spectra. The other half of the framework, the optical measurement, places equally demanding requirements on experimental fidelity. Quantitative single-particle spectroscopy requires not only detecting the faint scattered signal from an individual nanoparticle, but doing so with sufficient signal-to-noise ratio (SNR), spectral accuracy, and reproducibility to enable meaningful comparison with simulation.

Dark-field scattering spectroscopy has become a standard tool for interrogating single plasmonic nanoparticles because it suppresses the bright illumination beam while transmitting the scattered field. However, conventional dark-field configurations suffer from several limitations that hinder quantitative analysis: residual stray light, sensitivity to alignment, wavelength-dependent illumination/collection imbalance, and geometric constraints imposed by fixed dark-field condensers [30], [31]. In inverted microscopes, these constraints become even more restrictive, as specular reflections and substrate-mediated background often leak into the detection path (see cover).

This chapter describes the quasi-darkfield reflectance microspectroscopy setup developed to overcome these challenges. Starting from the principles of angular filtering and background suppression, we analyze each component of the detection chain: illumination pupil fill, partial-reflector geometry, pinhole filtering, instrument-response calibration, and autofocus strategy. Then, we show how their combined optimization yields single-particle spectra of sufficient quality for quantitative morphology–property correlations. The measurements presented in Chapter 5 rely directly on this setup.

## 4.2 Quasi-darkfield microspectroscopy

### 4.2.1 Introduction and motivation

Quantitative single-particle scattering spectroscopy demands high signal-to-noise ratio (SNR) and robust background suppression. Conventional dark-field microscopy is powerful, but in inverted geometries it is often limited by residual substrate reflections, parasitic stray light, and restricted control over the balance between illumination and collection angles [30], [31]. Standard dark-field condensers eliminate low-numerical aperture (NA) illumination but do not prevent all low-NA background from entering

the detection path, especially in spectroscopic measurements where the collected angular range must remain broad.

Building on the partial-reflector concept introduced by Kukura and co-workers for interferometric and extinction-based contrast enhancement [124], we adopt a reflectance-mode configuration in which a physical beam stop in the detection path suppresses the on-axis, low-NA background. This geometry provides fine control over the angular components reaching the detector while remaining compatible with a compact, inverted microscope.

In this chapter we show how optimizing beam parameters (illumination pupil fill, reflector size and placement), implementing wavelength-dependent focusing at the sample plane, and adding stray-light suppression and background-correction strategies improves the SNR relative to standard dark-field configurations. The central design idea is to treat the illumination beam diameter and the partial-reflector diameter as coupled parameters in the pupil plane: together they set the trade-off between spatial resolution, collection efficiency, and low-NA background rejection.

We refer to the method as *quasi-darkfield* because, unlike classical dark-field objectives that angularly separate illumination and collection to exclude the unscattered beam entirely, we operate in a common-path reflectance geometry. A partial reflector (central stop) in the detection arm physically blocks the specular reflection and low-NA illumination leakage while still transmitting mid- to high-NA scattered light. Because this stop also unavoidably removes some low-NA scattering from the nanoparticle, the configuration is not strictly dark-field. However, compared to conventional dark-field condensers with fixed hardware constraints, our approach allows the illumination fill and stop size to be tuned independently. This enables direct optimization of the trade-off between background suppression and signal collection, yielding a signal-to-noise ratio that exceeds that of related inverted dark-field configurations. We further show that, when combined with refractive-index matching of the substrate, this configuration enables quantitative and reproducible single-particle scattering spectra.

### 4.2.2 Setup overview

The quasi-darkfield geometry is implemented in an inverted reflectance-mode microscope, where a high-NA oil-immersion objective (NA 1.3) collects a large fraction of the scattered field from plasmonic nanoparticles (PNPs), which radiate into a broad angular range [30], [31]. Reflectance operation with a beamsplitter enables a compact, stable configuration in which the same high-NA objective is used for both illumination and collection, ensuring mechanical robustness and straightforward alignment for

spectroscopy. A collimated beam is cleaned with an iris/spatial filter, spectrally conditioned, and adjusted in diameter before the objective pupil (see Figure 4.A1):

- *Beam generation:* A SuperK continuum laser is fiber-connected to an LLTF filter that allows for wavelength selection with an average bandwidth of 2.5 nm.
- *Beam cleanup:* a pinhole and an iris suppress stray wings and higher-order structure prior to the objective pupil.
- *Pupil fill:* a beam reducer (or telescope) sets the collimated beam diameter at the objective pupil, thereby defining the illumination NA (Section 4.2.3).

The back-reflected illumination and scattered light are collected by the same objective and directed through the beamsplitter to the detection arm. A *partial reflector* (central stop) is placed at a plane conjugate to the objective back focal plane (Fourier plane) to block the on-axis/low-angle components dominated by specular reflection and residual illumination. At an image plane, a *pinhole* further suppresses out-of-focus and parasitic stray light prior to the avalanche photon detector (APD).

Figure 4.1 shows a simplified schematic of the optical layout. Light is delivered as a *collimated* beam whose diameter is set by a beam reducer (4) to define the illumination numerical aperture. The beam is reflected by a beamsplitter (3) into the objective (2) and onto the sample (1). The back-reflected illumination and scattered light are collected by the same objective and pass through the beamsplitter towards the detection arm. A *partial reflector* (5), placed at a plane conjugate to the objective back focal plane (Fourier plane), serves as a physical *central stop* that blocks the on-axis and low-angle components (dominated by specular reflection and residual illumination). A pinhole (6) at an image plane further suppresses out-of-focus and parasitic stray light before the spectrograph/camera.

### 4.2.3 Angular filtering and coupled design parameters

To quantify the angular filtering imposed by the central stop and the illumination beam diameter, we define two dimensionless design parameters that capture the core trade-off of the quasi-darkfield geometry. In the objective pupil (a pupil-conjugate Fourier plane), radial position maps directly to propagation angle and therefore to numerical aperture. Consequently, the collimated beam diameter at the objective pupil sets the illumination NA, and the diameter of the central stop sets the minimum NA that reaches the detector.

Let  $D_{\text{pupil}}$  denote the objective pupil diameter,  $D_{\text{beam}}$  the collimated illumination




Figure 4.1: Schematic of the inverted quasi-darkfield microscope: (1) gold nanorod on a glass slide, (2) objective lens, (3) beamsplitter, (4) beam reducer, (5) partial reflector (central stop in the detection path), and (6) pinhole. The reflector acts as a physical beam block at a pupil-conjugate plane to suppress on-axis/low-NA background while transmitting mid- and high-NA scattering.

diameter at the pupil, and  $D_{\text{stop}}$  the diameter of the central stop in the detection path. We define:

$$\alpha \equiv \frac{D_{\text{beam}}}{D_{\text{pupil}}}, \quad \beta \equiv \frac{D_{\text{stop}}}{D_{\text{pupil}}}.$$

These normalized diameters set the effective NA bounds:

$$\text{NA}_{\text{ill}} \approx \alpha \text{NA}_{\text{obj}}, \quad \text{NA}_{\text{min}} \approx \beta \text{NA}_{\text{obj}},$$

where  $\text{NA}_{\text{obj}} = 1.3$  in our setup (the hardware limit). Because the specular reflection retraces the illumination cone, effective suppression of the unscattered beam requires  $\beta > \alpha$ . In practice, a small safety margin is added so that  $\beta$  remains just above  $\alpha$  even under minor alignment drift or beam pointing fluctuations (typical margin  $\Delta\beta \sim 0.05\text{--}0.10$  in NA-normalized units).

This coupling expresses the central design trade-off:

- **Increasing**  $\alpha$  (filling more of the objective pupil) increases the illumination NA, improving spatial resolution and overlap with mid/high-NA scattered light, but it *also* demands a larger stop ( $\beta$ ) to block the specular core, which removes more low-NA scattered signal.

- **Decreasing**  $\alpha$  (underfilling the pupil) allows a smaller stop ( $\beta$ ) and thus rejects less scattered light, which can improve SNR when background is limiting. The trade-off is reduced spatial resolution and increased sensitivity to beam divergence over long paths.

Because PNPs scatter over a broad angular range [30], [31], the retained mid-/high-NA content is sufficient for sensitive single-particle measurements when  $\{\alpha, \beta\}$  are chosen appropriately. Qualitatively, background suppression improves rapidly as  $\beta$  is increased toward the illumination-cone edge, while useful signal decreases more gradually; therefore, the practical optimum is straightforward: set  $\beta$  just large enough to block the specular reflection cone, i.e.,  $\beta > \alpha$ , with a small margin for alignment tolerance. Any further increase in  $\beta$  suppresses useful low-NA scattered light without additional benefit to background rejection. In our implementation, we deliberately *underfill* the objective ( $\alpha < 1$ ) to keep the required stop only slightly larger than the illumination cone, minimizing lost scattered power while ensuring robust suppression of the specular core.

To make this quantitative in later sections, we define the relative collection efficiency  $\eta(\lambda)$  as the fraction of the nanoparticle’s backscattered power within the acceptance window  $[\theta_{\min}, \theta_{\max}]$  (set by  $\beta$  and the objective NA), integrated over the upper hemisphere. Computing  $\eta(\lambda)$  requires (i) the angular acceptance window determined by  $\{\alpha, \beta\}$  and (ii) the nanoparticle angular scattering distribution  $I_{\text{sca}}(\theta, \phi; \lambda)$ . These ingredients are established in Sections 4.2.3 and 4.2.5, respectively, and then combined to guide the choice of  $\{\alpha, \beta\}$ .

#### 4.2.4 Angular collection efficiency and optimal reflector size

To guide the choice of illumination fill and reflector size, we first quantify what fraction of a nanoparticle’s backscattered power falls within the angular collection window of the setup. This collection efficiency represents the *signal*-side of the SNR trade-off; the background contribution and empirical SNR optimisation are addressed in Section 4.3.1. The two ingredients required for this calculation are: (i) the nanoparticle’s angular scattering distribution, and (ii) the setup’s angular detection range  $[\theta_{\min}, \theta_{\max}]$  set by the illumination underfill and the objective NA.

We use a  $30 \times 96.5$  nm gold nanorod (AuNR), consistent with the particles studied in Chapters 2–3, as a representative scatterer. Its backscattering angular distribution  $I_{\text{sca}}(\theta, \phi; \lambda)$  is obtained from Finite-Difference Time-Domain (FDTD) simulations, which directly provide the radiated power per solid angle. We employ FDTD here (but BEM elsewhere) because our BEM implementation does not natively project the

far-field angular pattern.<sup>1</sup>

The first ingredient required for the collection efficiency  $\eta(\lambda)$  is the angular detection window  $[\theta_{\min}, \theta_{\max}]$  of the microscope. The *outer edge* is fixed by hardware:

$$\theta_{\max} = \arcsin\left(\frac{\text{NA}_{\text{obj}}}{n}\right),$$

with  $\text{NA}_{\text{obj}} = 1.3$  and immersion oil index  $n = 1.51$ , yielding  $\theta_{\max} = 59.43^\circ$ .

The *inner edge*  $\theta_{\min}$  is determined by the illumination NA, which is set not by hardware but by the beam diameter at the objective pupil. Underfilling the pupil reduces the illumination NA and therefore reduces  $\theta_{\min}$ .

To determine  $\text{NA}_{\text{illum}}(\lambda)$ , we must know the beam diameter at the objective entrance pupil. Because the laser output exhibits both wavelength-dependent beam size and finite divergence, we propagate the datasheet-provided beam diameters (Fig. 4.2a) through the optical path using the standard Gaussian-beam expression:

$$w(z) = w_0 \sqrt{1 + \left(\frac{\lambda z}{\pi w_0^2}\right)^2}.$$

The resulting beam diameters at successive optical planes are shown in Fig. 4.2b. Smaller initial waist radii  $w_0$  generate stronger divergence, causing the beam to expand more rapidly over the same propagation distances.

Let the objective entrance pupil diameter be  $D_{\text{pupil}} = 7$  mm (manufacturer specification). If  $D_{\text{beam}}(\lambda)$  is the propagated beam diameter at the pupil plane, the illumination NA is

$$\text{NA}_{\text{illum}}(\lambda) = \left(\frac{D_{\text{beam}}(\lambda)}{D_{\text{pupil}}}\right) \text{NA}_{\text{obj}}.$$

Full pupil fill ( $D_{\text{beam}} = D_{\text{pupil}}$ ) gives  $\text{NA}_{\text{illum}} = \text{NA}_{\text{obj}}$ , while underfilling reduces  $\text{NA}_{\text{illum}}$  proportionally. Figure 4.3a shows the resulting  $\text{NA}_{\text{illum}}(\lambda)$  and Fig. 4.3b shows the corresponding  $\theta_{\min}(\lambda) = \arcsin(\text{NA}_{\text{illum}}(\lambda)/1.51)$ .

With  $\theta_{\min}(\lambda)$  established, the required reflector size follows directly: placed at a pupil-conjugate (Fourier) plane, the reflector must be large enough to block the full illumination cone. In the notation of Section 4.2.3, this requires  $\beta > \alpha$ , where  $\alpha = D_{\text{beam}}/D_{\text{pupil}}$  and  $\beta = D_{\text{stop}}/D_{\text{pupil}}$ . Because small alignment drifts effectively increase the transmitted low-NA light, we choose a small safety margin above equality. Anticipating the experimental optimisation in Section 4.3.1, a beam diameter of

---

<sup>1</sup>FDTD permits direct far-field projection; BEM simulations in earlier chapters focused on cross sections, not angular scattering.




Figure 4.2: (a) Laser beam diameter from the specification sheet (three wavelengths) with a second-degree polynomial fit used to interpolate across the working spectral range. (b) Calculated beam diameter at successive planes in the setup (see Fig. 4.1), including the beam reducer and the objective pupil. Different colors indicate distinct initial waist conditions  $w_0$  propagated over the same distances.




Figure 4.3: (a) Illumination numerical aperture  $NA_{\text{illum}}(\lambda)$  obtained from the beam diameter at the objective pupil relative to the 7 mm entrance pupil (scaled by  $NA_{\text{obj}} = 1.3$ ). (b) Minimum collection angle  $\theta_{\text{min}}(\lambda)$  obtained from  $\theta_{\text{min}} = \arcsin(NA_{\text{illum}}/1.51)$ . The outer cutoff  $\theta_{\text{max}} = 59.43^\circ$  is fixed by objective hardware.

approximately 2 mm at the pupil provided near-optimal SNR in our setup.

With  $[\theta_{\text{min}}, \theta_{\text{max}}]$  and  $I_{\text{sca}}$  specified, the fraction of collected backscattered power is

$$\eta(\lambda) = \frac{\int_{\theta_{\text{min}}}^{\theta_{\text{max}}} d\theta \int_0^{2\pi} d\phi I_{\text{sca}}(\theta, \phi; \lambda) \sin \theta}{\int_0^{\pi/2} d\theta \int_0^{2\pi} d\phi I_{\text{sca}}(\theta, \phi; \lambda) \sin \theta} \quad (4.1)$$

i.e.  $\eta = I_{\text{collected}}/I_{\text{backscattered}}$ . This quantity represents only the *signal* side of the

SNR balance; background and stray-light contributions depend on  $\beta$  and are treated experimentally in Section 5.1.6.

### 4.2.5 Nanoparticle radiation pattern and collection efficiency

In the previous section, we established how the angular collection window  $[\theta_{\min}, \theta_{\max}]$  follows from the illumination beam diameter at the objective entrance pupil, which sets  $\text{NA}_{\text{illum}}$  and thereby the minimum detectable angle  $\theta_{\min}(\lambda)$ . For a given illumination fill (and thus  $\theta_{\min}$ ), the resulting collection efficiency depends on how much of the nanoparticle’s wavelength-dependent angular scattering distribution lies within this acceptance window. We therefore compute the spectral collection efficiency by combining the calculated angular limits with the nanoparticle’s backscattering profile.

Figure 4.4a shows the FDTD-calculated backscattered far-field radiation pattern of a  $30 \times 96.5$  nm AuNR in immersion oil, the same dimensions used in Chapters 2–3. The intensity is normalized such that the integral over the upper hemisphere equals unity. The plot displays the angular dependence of  $I_{\text{sca}}(\theta, \phi)$  itself. For a dipolar longitudinal plasmon mode, the physical radiated power per solid angle is proportional to  $I_{\text{sca}}(\theta, \phi) \sin \theta$ , and therefore peaks at intermediate angles, forming a doughnut-like distribution. Because Fig. 4.4a shows  $I_{\text{sca}}$  rather than  $\sin \theta I_{\text{sca}}$ , the maximum appears closer to the optical axis. The  $\sin \theta$  weighting is included in the final efficiency integration (Eq. 4.1), ensuring that the collected power is correctly weighted by solid angle.

This distinction is important: although much of the amplitude is concentrated near lower polar angles, the majority of the radiated power detectable in reflectance geometry lies at intermediate angles, which are transmitted by our quasi-darkfield geometry even when low-NA components are blocked by the central stop.

Because both  $\theta_{\min}(\lambda)$  and the nanoparticle radiation pattern  $I_{\text{sca}}(\theta, \phi; \lambda)$  vary with wavelength, the collection efficiency must be evaluated spectrally. Figure 4.4b reports the resulting  $\eta(\lambda)$  for an illumination beam diameter of approximately 2 mm at the objective pupil, the value that later proves optimal for maximizing SNR. Across the visible range, the model predicts a relative collection efficiency of about  $\sim 30\%$ . The step-like fine structure originates from the discrete angular sampling of the FDTD far-field monitor.

For clarity:

- $\theta_{\max} = 59.43^\circ$  is the *outer*, hardware-limited collection angle set by the NA 1.3 oil-immersion objective in immersion oil.




Figure 4.4: (a) Simulated backscattered far-field radiation profile (with FDTD) of a  $30 \times 96.5$  nm AuNR in immersion oil at its longitudinal resonance. The intensity is normalized such that the integral over the upper hemisphere equals unity. Polar angle  $\theta$  is measured from the surface normal, while the outer ring represents the azimuthal angle  $\phi$ . (b) Theoretical collection efficiency for an illumination beam diameter of approximately 2 mm at the objective pupil. The efficiency reflects the fraction of the scattered power falling between the wavelength-dependent inner cutoff  $\theta_{\min}(\lambda)$  and the hardware-limited outer angle  $\theta_{\max} = 59.43^\circ$ . Step-like features originate from the angular discretization of the FDTD monitor.

- $\theta_{\min}(\lambda)$  is the *inner* cutoff, determined by the illumination NA (i.e. by  $D_{\text{beam}}/D_{\text{pupil}}$ ). All scattering with  $\theta < \theta_{\min}$  is blocked by the central stop to suppress background.

Thus the detected scattering is that lying in the angular annulus  $\theta_{\min}(\lambda) < \theta < \theta_{\max}$ .

Importantly, this  $\sim 30\%$  collection efficiency is significantly higher than what is typically obtained in conventional dark-field microscopes that use high-NA condensers and annular illumination. In such systems, only a narrow angular annulus is transmitted, and the nanoparticle’s radiation pattern is only partially captured. In contrast, the quasi-darkfield geometry retains a much broader mid- and high-NA range while still rejecting low-NA background, which directly improves SNR by delivering a larger portion of the scattered power to the detector.

Besides its diameter, the reflector’s opacity critically determines background suppression. Kukura *et al.* demonstrated that an 82 nm Ag coating yields visible-range transmittance of  $10^{-3}$ – $10^{-4}$ , enabling interferometric contrast [124], [125]. Our application requires even stronger background suppression rather than interference, so we selected a thicker ( $\sim 120$  nm) Al film, which exhibits transmittance on the order of  $10^{-6}$ – $10^{-7}$  across the visible spectrum [126].

In practice, occasionally higher than expected low-NA leakage was observed. This

likely arises from microscopic defects in the deposited film such as nanometer-scale pinholes, incomplete grain coalescence, or local oxidation, none of which need to approach the full film thickness to introduce detectable transmission when placed at a Fourier plane. To eliminate such variability, we now use opaque neodymium micro-magnets (diameter 2 mm) mounted on a coverslip as the central stop. Their high opacity, rigidity, and well-defined diameter provide reproducible angular blocking, and the 2 mm size corresponds to the cutoff used in the collection-efficiency calculations. Further optimisation of reflector diameter can be performed following the procedure of Kukura *et al.*, where stop size is varied systematically while measuring the resulting background and signal levels; an analogous approach can be implemented for the quasi-darkfield geometry used here.

### 4.2.6 Pinhole design and stray-light suppression

Reflections within the objective and downstream optical components can introduce stray laser light at relatively high numerical apertures. In the configuration of Kukura *et al.*, such high-angle background is suppressed by placing the partial reflector directly at the objective back focal plane [124]. In our inverted layout this placement is not mechanically feasible. Instead, we suppress these unwanted high-NA components by inserting an additional spatial filter, a pinhole, at an image plane in the detection arm (element 6 in Fig. 4.1).

The pinhole serves two complementary purposes. First, it blocks high-angle stray light originating from internal reflections, microscopic surface imperfections within the objective, and clipping at intermediate apertures in the detection path. Second, it acts as a spatial-frequency filter that transmits only the *zeroth diffraction order* of the nanoparticle's image while rejecting higher-order diffraction rings. This is advantageous for spectroscopic measurements because the integrated scattering cross section is encoded entirely in the diffraction-limited central spot. Higher-order rings contain spatial structure but carry only a small fraction of the total power and would complicate quantitative integration.

To choose an appropriate pinhole diameter, we must determine the physical size of the zeroth-order Airy disk at the pinhole plane. Figure 4.A1 in the Appendix shows a qualitative diffraction pattern generated when a 532 nm beam is focused onto the pinhole plane: a bright central disk corresponding to the zeroth order surrounded by faint rings containing higher spatial frequencies. The pinhole must be chosen such that only this central disk reaches the detector.

To estimate the Airy-disk diameter, we use the standard expression for a focused

collimated beam of radius  $r$  incident on a lens of focal length  $f$ ,

$$D = \frac{\lambda f}{r},$$

which is accurate when the beam fills the lens aperture. In our collection path, stray light is not restricted to the  $\sim 2$  mm illumination beam used to set the quasi-darkfield angular window. Both the nanoparticle signal and all high-angle stray light fill the entire objective back aperture, because the nanoparticle scatters over a broad angular range and the objective collects up to its hardware limit. Consequently, the relevant beam diameter for estimating the Airy disk is the full 7 mm entrance pupil ( $r = 3.5$  mm). Using this value yields the wavelength-dependent Airy-disk diameters shown in Fig. 4.5a, which remain below approximately  $15 \mu\text{m}$  across the visible range. A pinhole of this diameter would therefore transmit the entire zeroth-order spot for both signal and stray light while blocking higher-order diffraction features.

To experimentally validate this prediction, we measured the diffraction-limited spot produced by a  $20 \times 75$  nm AuNR under 532 nm excitation. For this measurement the pinhole was deliberately enlarged to  $100 \mu\text{m}$  and slightly misaligned so that the full Airy pattern was visible on the camera. Figure 4.5b shows the resulting image. From the recorded pattern we extract a central-disk diameter of approximately  $8 \mu\text{m}$ , in excellent agreement with the theoretical prediction of  $7.6 \mu\text{m}$  at 532 nm. This agreement confirms that the theoretical Airy-disk curve in Fig. 4.5a accurately captures the wavelength dependence across the full spectral range.

Taken together, these results indicate that a maximum pinhole diameter of approximately  $15 \mu\text{m}$  would fully transmit the zeroth-order spot at all relevant wavelengths. To ensure robust suppression of higher-order spatial frequencies and residual high-angle background, we choose a more conservative pinhole diameter of  $10 \mu\text{m}$ . This value reliably transmits the diffraction-limited core across the spectrum while minimizing the probability of passing unwanted spatial content. As a consequence, some mild wavelength-dependent clipping occurs at the long-wavelength end, an effect that contributes to the spectral instrument response and is corrected for in the calibration procedure described in the next section.

### 4.3 Considerations for spectroscopy

The preceding sections established the optical geometry of the quasi-darkfield setup: illumination underfill, angular filtering by the partial reflector, and stray-light suppression via the pinhole. Having defined these elements of the optical pathway, we




Figure 4.5: (a) Theoretical diameter of the diffraction-limited Airy disk at the pinhole plane as a function of wavelength, assuming a 7 mm beam diameter at the focusing lens. (b) Measured diffraction pattern of a  $20 \times 75$  nm AuNR under 532 nm excitation with a deliberately oversized ( $100 \mu\text{m}$ ) and slightly misaligned pinhole, revealing the full Airy structure. Differences in magnification between (a) and (b) arise because the images correspond to distinct planes in the optical path.

now turn to the measurement protocol and calibration steps required to convert raw detector counts into *quantitative* single-particle scattering spectra. This includes (i) characterising the wavelength-dependent response of the detection chain, (ii) controlling the local optical environment to maximise signal quality, and (iii) ensuring stable focus and alignment across the full spectral range. Quantitative single-particle spectroscopy requires that the recorded signal reflects the true spectral shape of the nanoparticle rather than the wavelength-dependent transfer function of the detection hardware. In what follows, we outline how this instrument response is determined and corrected.

### 4.3.1 Instrument response function

Any optical measurement system has a wavelength-dependent response that distorts the recorded spectrum relative to the true scattered signal. To recover quantitative spectra, this *instrument response function* (IRF) must be characterised and accounted for. In our measurements, scattering spectra are acquired with an avalanche photodiode (APD) while scanning the excitation wavelength; for each wavelength step, the APD integrates the scattered signal transmitted through the quasi-darkfield detection path.

The overall wavelength-dependent response of this detection chain arises from multiple factors: the tunable filter throughput (including internal grating configuration), the transmission of relay optics, the beamsplitter, the wavelength-dependent coatings

of the objective, the APD quantum efficiency, and wavelength-dependent clipping at the pinhole. We denote the combined transfer function by the instrument response function  $\text{IRF}(\lambda)$ .

A direct “background” measurement is not suitable for determining  $\text{IRF}(\lambda)$ , because background is intentionally suppressed in quasi-darkfield and does not probe the throughput of the detection chain in a controlled manner. Instead, we determine the IRF using a *reference scatterer* with a known (or reliably simulated) spectral response. We employ 100 nm polystyrene (PS) beads, whose single-particle scattering spectra on glass can be accurately computed in the visible range using MNPBEM.

Let  $C_{\text{ref}}(\lambda)$  be the APD counts measured from a single PS bead under the same optical conditions used for our nanoparticle measurements, and let  $\sigma_{\text{ref}}^{(\text{sim})}(\lambda)$  be the simulated bead-scattering spectrum on glass (Fig. 4.6c). The instrument response function is then obtained as the ratio of simulated to measured spectra, normalised to unit mean:

$$\text{IRF}_{\text{norm}}(\lambda) = \frac{\sigma_{\text{ref}}^{(\text{sim})}(\lambda)}{C_{\text{ref}}(\lambda)} / \left\langle \frac{\sigma_{\text{ref}}^{(\text{sim})}(\lambda)}{C_{\text{ref}}(\lambda)} \right\rangle_{\lambda}.$$

For any measured single-particle spectrum  $C_{\text{meas}}(\lambda)$ , the corrected spectrum is obtained by multiplication:

$$S_{\text{corr}}(\lambda) = C_{\text{meas}}(\lambda) \text{IRF}_{\text{norm}}(\lambda).$$

Absolute scaling to obtain scattering cross sections is treated separately in Chapter 5; here we address only the spectral shape.

Figure 4.6 summarises the IRF determination procedure. Panel (a) shows an SEM image of the PS beads (100 nm). Panel (b) displays APD scattering spectra from three single beads and their average, which probe the full detection chain. Panel (c) shows the corresponding MNPBEM simulation (inset: bead-on-glass geometry). Panel (d) presents the resulting normalised IRF. For completeness, the APD quantum-efficiency curve is provided in Appendix Fig. 4.A3; the reduced QE at longer wavelengths partly explains the downward trend in  $\text{IRF}_{\text{norm}}$  and contributes to the reduced sensitivity observed in corrected spectra above  $\sim 700$  nm.

Our tunable wavelength filter internally switches gratings near  $\sim 650$  nm, producing a small but distinct throughput discontinuity visible in raw spectra. The IRF captures this feature and removes the resulting kink upon correction. Figure 4.7 compares a raw and an IRF-corrected spectrum of a single  $20 \times 75$  nm AuNR on glass. The

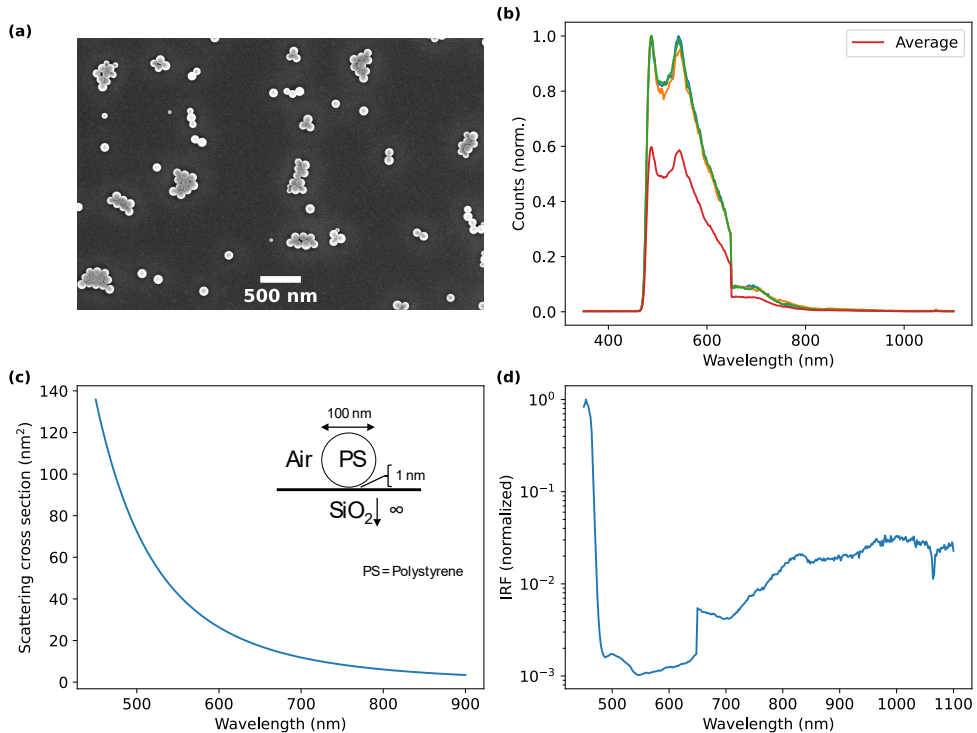


Figure 4.6: (a) SEM of polystyrene beads (mean diameter 100 nm). (b) APD scattering spectra of three single beads and their average. (c) MNPBEM simulation of the single-bead scattering on glass (inset: geometry). (d) Normalized instrument response function (IRF), obtained as the ratio between simulated and measured spectra and normalized to unit mean.

grating-switch artifact is largely eliminated. Residual asymmetry at longer wavelengths likely originates from reduced detection efficiency (Appendix Fig. 4.A3) and minor mismatches between the reference and sample alignment. Conversion from APD counts to absolute cross sections is presented in Chapter 5.

### 4.3.2 Signal quality and alignment requirements

To ensure that the recorded spectra faithfully represent the nanoparticle's scattering response, it is important to understand how the optical system forms the nanoparticle image and how deviations from ideal alignment affect the detected signal. With collimated illumination underfilling the objective, the nanoparticle image on the detector corresponds to the diffraction-limited point-spread function (PSF) of the high-NA objective, scaled by the particle's scattering strength.




Figure 4.7: Raw (blue) and IRF-corrected (orange) APD scattering spectra of a single  $20 \times 75$  nm AuNR on a glass coverslip. The discontinuity at  $\sim 650$  nm originates from the grating switch inside the tunable filter and is removed by IRF correction.

The 3D PSF of a high-NA system is intrinsically elongated along the optical axis. Approximate lateral and axial extents are given by

$$\Delta x \approx \Delta y \approx 0.61 \frac{\lambda}{\text{NA}}, \quad \Delta z \approx 2 \frac{\lambda n}{\text{NA}^2},$$

with refractive index  $n$  set by the immersion medium. For  $\text{NA} = 1.3$  and oil ( $n = 1.51$ ), this yields  $\Delta x, \Delta y \sim 250\text{--}300$  nm and  $\Delta z \sim 1.4 \mu\text{m}$  at visible wavelengths. The latter estimate agrees with the measured PSF shown in Fig. 4.8a, but the axial extent appears larger  $\sim 500$  nm in practice because the illumination underfills the objective, effectively reducing the operational NA for axial confinement.

When the sample plane is even slightly tilted or defocused, the camera records a *slanted projection* of this 3D PSF, producing an elongated, asymmetric spot. Figure 4.8 illustrates this: panel (a) shows a confocal 3D measurement of a single  $20 \times 75$  nm AuNR, and panel (b) shows a slice highlighting the slant generated by a small tilt or focus error. Such distortions are detrimental for spectroscopy because a slanted PSF couples less stably through the  $10 \mu\text{m}$  pinhole, leading to wavelength-dependent throughput variations and degraded SNR.

In practice, we minimise PSF slant by (i) leveling the sample, (ii) setting focus by maximising the APD signal transmitted through the pinhole, and (iii) employing the

autofocus strategy described in the next section.

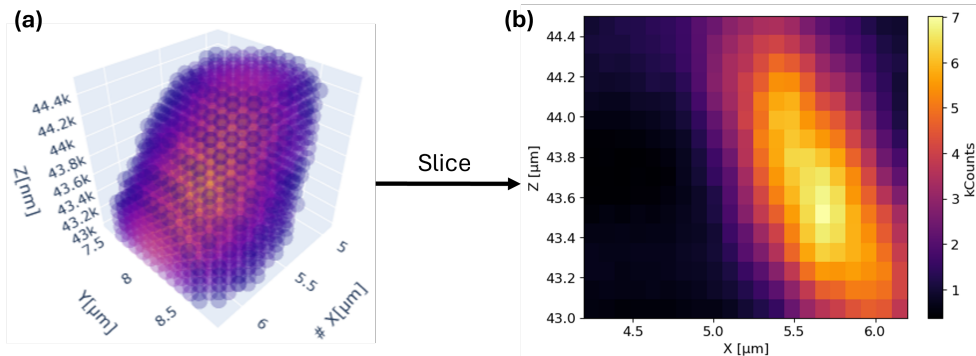


Figure 4.8: (a) Confocal 3D measurement of a single  $20 \times 75$  nm AuNR on glass showing the elongated 3D PSF of the high-NA system. (b) Slice view highlighting the slanted projection that appears when the sample plane is slightly tilted or defocused relative to the optical axis.

Throughout this chapter, unless otherwise stated, we quantify signal quality using

$$\text{SNR} \equiv \frac{S_{\text{peak}}}{\langle B \rangle},$$

the ratio of the peak nanoparticle signal to the mean background level in a particle-free region of the same image or spectrum. This definition is robust and directly reflects the practical detectability of single particles in our setup.

### 4.3.3 Refractive-index matching and background suppression

Background suppression in single-particle spectroscopy depends not only on the angular filtering provided by the quasi-darkfield geometry, but also on the refractive index of the immersion medium relative to the substrate. A large index contrast (e.g., air/glass) enhances scattering amplitude but also increases Fresnel reflections, substrate-mediated stray light, and background. As a result, the SNR may deteriorate even when the absolute signal appears higher.

By *index matching* the immersion medium to the substrate, we suppress Fresnel reflections and reduce background without compromising the mid- to high-NA scattered light admitted by the quasi-darkfield stop.

To demonstrate this experimentally, Fig. 4.9 compares 2D APD scan maps of  $20 \times 75$  nm AuNRs on a glass coverslip in air, water, and immersion oil (indices 1.0, 1.33, and 1.51). As the immersion index approaches that of the glass substrate,

background diminishes markedly and SNR increases. In immersion oil, we routinely achieve  $\text{SNR} \approx 65$ , compared to  $\text{SNR} \approx 15$  reported by Kukura *et al.* for related interferometric measurements [124], even though their measurements also used an oil-immersion objective. This underscores the combined effect of angular filtering, pinhole suppression, and index matching in our setup. Notably, relying on absolute signal amplitude alone is misleading: in air the analytic scattering is higher due to the larger refractive-index contrast, but the background increases even more strongly, leading to inferior SNR.




Figure 4.9: 2D microspectroscopy scans of  $20 \times 75$  nm AuNRs on a glass coverslip in (a) air, (b) water, and (c) immersion oil. As the immersion index approaches that of glass, background is reduced and SNR increases.

Index matching is even more important when performing spectroscopy on *TEM window* substrates. As discussed earlier, correlative single-particle studies require that optical measurements be carried out directly on TEM grids so that structural characterisation of the same particle can be obtained afterward by electron microscopy. TEM windows introduce additional background pathways due to surface roughness, stress in the thin membrane, and internal reflections not present in coverslips.

Figure 4.10 compares 2D APD scan maps on glass coverslips (air and oil) and on a glass TEM window in oil. Panels (d–f) display the same maps on a logarithmic scale to emphasise single-particle signals. Somewhat surprisingly, the SNR on the glass TEM window in oil is comparable to, and in our sample even higher than, that on standard glass. This is encouraging for correlative measurements: the extremely thin TEM membrane contains very little material to generate bulk scattering or internal reflections, and its effective refractive index may be closely matched to the immersion oil. Small wrinkle-like patterns observed in panel (e) likely arise from minor local thickness variations at the oil–glass interface. With careful sample preparation, the quasi-darkfield geometry therefore yields excellent SNR on TEM grids, enabling reliable single-particle spectroscopy directly on substrates compatible with electron microscopy.




Figure 4.10: 2D scans of  $20 \times 75$  nm AuNRs on (a) glass in air, (b) glass in oil, and (c) a glass TEM window in oil. (d–f) show the same maps on a natural logarithmic scale to enhance particle visibility. Despite the thin-membrane substrate, the TEM window exhibits comparable or even slightly higher SNR, demonstrating the suitability of this geometry for correlative optical–electron measurements.

#### 4.3.4 Wavelength-focusing

While the static alignment procedure in Section 4.3.2 establishes a well-focused starting point, chromatic aberrations in the optical system shift the best-focus position with excitation wavelength. Maintaining stable coupling through the pinhole therefore requires *active refocusing at each wavelength step*. Without this correction, wavelength-dependent focus errors translate directly into wavelength-dependent throughput variations, biasing the measured spectra (see also Fig. 4.8).

Accurate focusing is especially critical in quasi-darkfield microscopy, where the pinhole defines a tight spatial-frequency filter. Even small focus errors alter the shape and position of the PSF relative to the pinhole aperture, reducing transmitted signal and introducing artificial spectral modulations. Unlike broadband autofocus approaches that determine focus once and assume achromaticity [127], [128], our method scans the *excitation wavelength* and determines focus *at each wavelength*.

We compare two autofocus strategies. The first, the *max-counts* method, performs a 3D grid scan around the nominal focus at each wavelength using the piezo stage. For

the first wavelength the grid is centered on the manually-defined focus; at subsequent wavelengths it is centered on the previously found focus to minimise scan volume. At each voxel the APD counts are recorded, and focus is defined as the voxel with the maximum intensity.

Although simple and fast, this method is vulnerable to several failure modes: (i) residual high-NA background can create spurious bright voxels; (ii) sample drift during the scan can shift the apparent peak; (iii) rare detector spikes (e.g., cosmic rays) can produce artificial maxima; (iv) in asymmetric or tilted PSFs, the brightest voxel need not coincide with the geometric PSF center.

The practical consequence is shown in Fig. 4.11: rather than following a smooth trajectory dominated by chromatic focal shift, the sequence of focus coordinates appears as a *random 3D walk* even though the particle is fixed. At many wavelengths the system is therefore effectively out of focus, introducing precisely the kind of wavelength-dependent intensity errors that autofocus is meant to prevent.

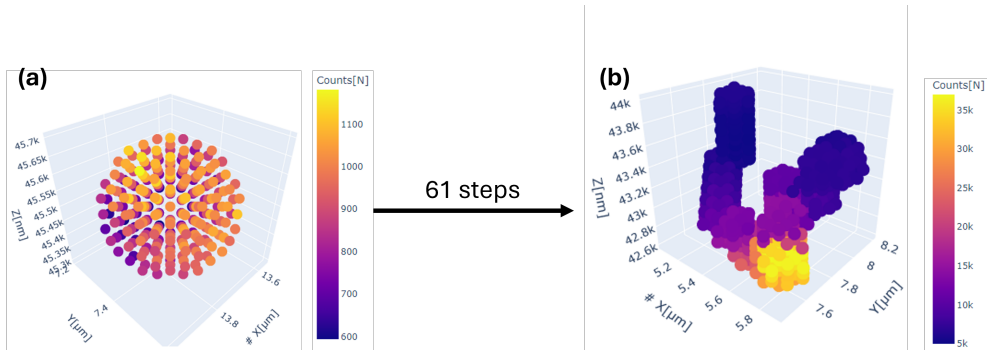


Figure 4.11: (a) Max-counts autofocus applied to a single  $20 \times 75$  nm AuNR on glass during an excitation-wavelength scan. (b) Retrieved focus coordinates (one point per wavelength) form a random 3D path, demonstrating that the brightest-voxel criterion is perturbed by background and outliers rather than tracing a smooth chromatic focus shift.

To increase *wavelength-focusing robustness*, we exploit the fact that the high-NA 3D PSF has an approximately ellipsoidal shape. For each wavelength, we acquire a 3D grid scan and perform a *3D ellipsoid fit* to the measured intensity distribution. The fitted center gives the focus; the fitted orientation accounts for sample tilt; the fitted axes capture lateral and axial PSF extents. Unlike max-counts, which relies on a single voxel, the ellipsoid fit aggregates information across the entire 3D PSF. This makes it highly resistant to: spurious bright voxels, smooth background gradients, substrate structure (e.g., TEM window texture), and low SNR per voxel.

Before fitting, we suppress outliers using a two-stage rejection: (i) a percentile-based upper cutoff to remove saturated or cosmic-ray voxels, and (ii) a uniform background subtraction. While such filtering could be added to max-counts, aggregating across all voxels remains the core reason the ellipsoid fit performs better under challenging conditions. Figure 4.12 demonstrates the difference: for the same particle on a SiO<sub>2</sub> TEM grid, max-counts selects a bright voxel near the PSF edge (a local artifact), whereas the ellipsoid fit identifies the true geometric center.

Figure 4.13 compares spectra recorded without refocusing, with max-counts autofocus, and with ellipsoid-fit autofocus for the same  $20 \times 75$  nm AuNR on glass in air. The initial focus was set at  $\lambda = 700$  nm; accordingly, all spectra coincide near this wavelength. Autofocus strongly improves spectral quality relative to no refocus. Under these benign conditions, high signal, low background, max-counts performs slightly better than ellipsoid-fit. This is expected: when SNR is high and background is minimal, the brightest voxel genuinely coincides with the PSF maximum (the regime where max-counts works as intended). The small advantage may also reflect the fact that max-counts directly maximises detected intensity, whereas the geometric PSF center found by the ellipsoid fit need not exactly coincide with the intensity peak for an elongated PSF. With further optimisation of scan volume and step size, ellipsoid-fit would likely match or exceed max-counts even under these benign conditions, while retaining the robustness advantage demonstrated below.

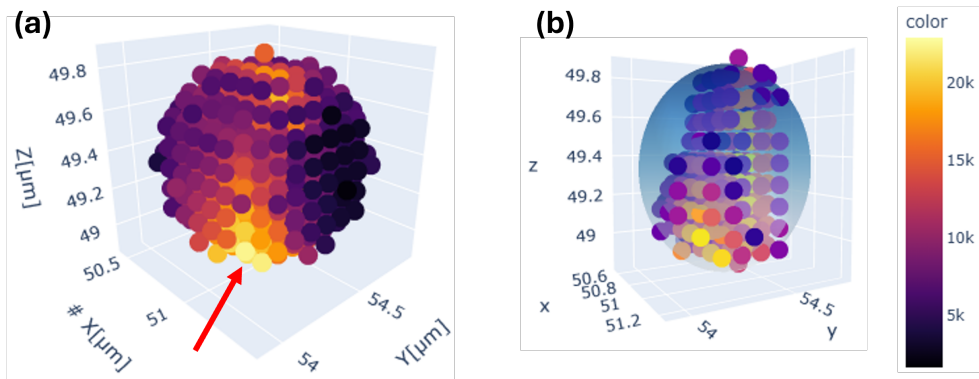


Figure 4.12: Autofocus comparison on the *same*  $20 \times 75$  nm AuNR on a SiO<sub>2</sub> TEM grid. (a) Max-counts: the brightest voxel (red arrow) lies near the PSF edge, indicating influence from local background rather than true focus. (b) Ellipsoid-fit: the fitted center (crosshair) coincides with the geometric PSF center, reflecting the true focus position.




Figure 4.13: Spectral comparison of no refocus, max-counts, and ellipsoid-fit autofocus for a  $20 \times 75$  nm AuNR on glass in air. Numbered labels denote ellipsoid-fit grid configurations (`radius_step` in  $\mu\text{m}$ ). Autofocus improves SNR relative to no refocus; under benign conditions max-counts slightly exceeds ellipsoid-fit, while the latter becomes advantageous under low-SNR or structured backgrounds.

As anticipated from the failure modes described above, the advantage of ellipsoid-fit becomes most pronounced when background is elevated or per-voxel signal is limited, precisely the case on TEM windows and for small particles.

In Fig. 4.14a (AuNR on TEM window in oil), max-counts loses focus beyond the tunable-filter grating switch ( $< 650$  nm) and again at long wavelengths ( $> 800$  nm), where the scattered signal is weakest. These focus failures introduce artificial spectral features unrelated to the nanoparticle’s physical response. Ellipsoid-fit, by contrast, yields a smoother baseline outside resonance and modest improvements near the LSPR peak ( $\sim 700$  nm), indicating partial stabilisation in this challenging regime.

Although Section 4.3.3 showed that AuNRs on oil-immersed TEM windows can produce high SNR in *2D scans at fixed wavelength*, autofocus distributes the signal across many voxels in a 3D grid. Far from resonance, where scattering is weak, the per-voxel signal drops near the detection limit, making TEM-window autofocus intrinsically more demanding. The benefit of ellipsoid-fit is even clearer in Fig. 4.14b for a 13 nm Au nanosphere on a TEM window in oil, a particularly demanding case because its scattering cross section is roughly two orders of magnitude smaller than that of the

AuNR. Both no-refocus and max-counts yield broad, poorly defined spectra with no discernible physical lineshape, whereas ellipsoid-fit recovers a credible resonance feature. In such low-SNR conditions, ellipsoid-fit is not merely incrementally better than max-counts, it makes the difference between a usable and an unusable spectrum.

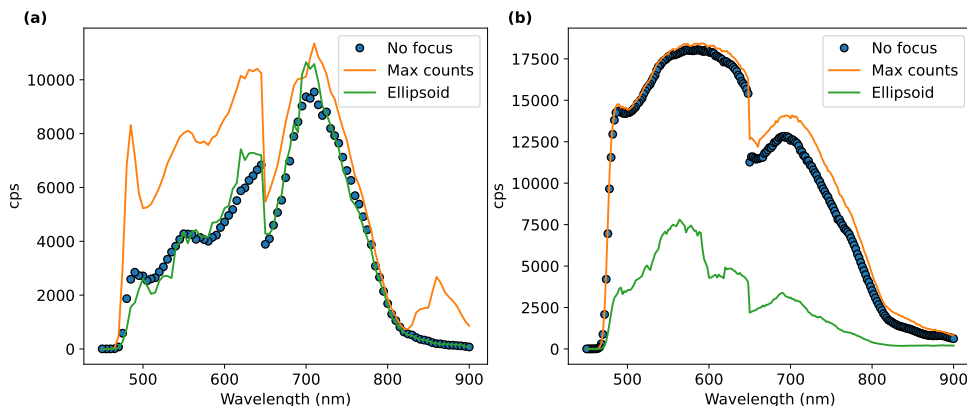


Figure 4.14: Autofocus comparison on (a) a  $20 \times 75$  nm AuNR and (b) a 13 nm Au nanosphere, both on TEM windows in oil. Max-counts loses focus at wavelengths where the per-voxel signal is lowest, introducing artificial spectral features. Ellipsoid-fit provides smoother baselines and recovers the expected lineshape for the lower-SNR nanoparticle in (b).

## 4.4 Conclusion

In this chapter we established the optical half of the quantitative structure–property framework developed throughout this thesis. Building on the quasi-darkfield geometry introduced in Section 4.2, we analysed how illumination underfill, partial-reflector diameter, and pinhole filtering together define the angular window through which single-particle scattering is detected. By explicitly modelling nanoparticle radiation patterns and propagating realistic beam parameters to the objective pupil, we derived a set of coupled design rules that maximise collection efficiency while suppressing low-NA background. This rationalizes the choice of a  $\sim 2$  mm illumination beam and a central stop just larger than the illumination cone, as well as the need for a  $10 \mu\text{m}$  pinhole to reject high-angle stray light without compromising the diffraction-limited signal.

Beyond the optical geometry, we developed a measurement workflow that enables reliable quantitative spectra. We showed how refractive-index matching between the immersion medium and substrate reduces Fresnel background and stabilises

the measured SNR, even on TEM window substrates required for correlative optical–electron microscopy. We also introduced a robust procedure for determining the instrument response function (IRF) using well-characterised reference scatterers, allowing wavelength-dependent system throughput, including grating switches, detector efficiency, and minor pinhole clipping, to be corrected systematically.

Finally, we addressed the practical challenge of maintaining stable coupling through the pinhole during spectral scans by introducing a wavelength-by-wavelength autofocus scheme. While a simple max-counts criterion can work under high-SNR conditions, we demonstrated that a 3D ellipsoid-fit method provides superior stability in low-SNR scenarios, on structured substrates, and for small particles. This approach uses the full 3D intensity distribution and is therefore resistant to background gradients, spurious bright voxels, and chromatic focal drift.

Taken together, these elements define a robust and fully calibrated quasi-darkfield microspectroscopy platform. The system achieves high signal-to-noise ratios, precise wavelength-dependent throughput correction, and stable focus across the visible range. With this optical framework in place, Chapter 5 leverages these capabilities to extract quantitative single-particle scattering spectra and correlate them directly with the simulated responses based on electron-microscopy-derived morphologies.

## 4.A Appendix




Figure 4.A1: Diffraction pattern of a 532 nm beam focused onto a pinhole. The bright central disk corresponds to the zeroth diffraction order (Airy disk), while the surrounding rings arise from higher-order spatial frequencies. The pinhole diameter is chosen such that only the central disk is transmitted to the detector.




Figure 4.A2: Detailed schematic of our home-built quasi-darkfield microscope.




Figure 4.A3: Spectral quantum efficiency of our APD (green)

## Chapter 5

---

# Typecast – Quantitative Optical–Morphological Correlation: Extracting Local Dielectric Environments

*Typecast – to always give an actor the same type of character to play, usually because he, she or they are physically suited to that type of part*

**Abstract:** This chapter brings together the experimental and computational frameworks developed throughout the thesis by establishing quantitative, single-particle correlations between measured scattering spectra and electromagnetic simulations. Using TEM-determined morphologies as fixed geometric input and instrument response function (IRF)-corrected dark-field spectroscopy acquired directly on TEM grids, we isolate the remaining unknown in the inverse problem: the local dielectric environment. By fitting measured spectra with physical models and comparing them to BEM simulations across discrete effective-medium conditions, we extract a local refractive index per particle and characterise its variation across a population. Measuring the same particles before and after TEM exposure, introducing a controlled perturbation of the dielectric environment through electron-beam-induced ligand modification and carbonaceous deposition, demonstrates that the framework can detect subtle environmental changes at the single-particle level.

The content of this chapter is in preparation in M. Dieperink, P. Spaeth, W. Albrecht, “Decoupling Geometry and Environment in Plasmonic Nanorods via Correlative Optical–Electron Microscopy”.

## 5.1 Introduction

Chapters 2 and 3 established the electromagnetic simulation framework and the morphology-extraction pipeline required for quantitative single-particle correlations, and Chapter 4 developed the optical measurement setup with the signal quality needed to support them. This chapter brings these components together into a complete correlative workflow, applied to a population of individual gold nanorods measured by single-particle dark-field spectroscopy and subsequently characterised by electron microscopy.

The central scientific question addressed here is not merely whether simulation and measurement can be brought into agreement; it is what physical information can be extracted once they are. With particle geometry fixed from TEM and the simulation framework validated in Chapters 2–3, the dominant remaining free parameter is the *local dielectric environment* surrounding each nanoparticle. The environment is nominally set by immersion oil and the substrate but is, in practice, influenced by surface ligands, local chemical variations, and substrate inhomogeneities that are difficult to characterise independently. By finding the effective-medium fraction that brings simulated and measured spectra into quantitative agreement for each particle, we extract the local dielectric environment on a single-particle basis.

## 5.2 Optical measurements and spectral analysis

For morphology–property correlations on the *same* particles, optical measurements must be performed directly on TEM grids to enable subsequent structural characterisation. Chapter 4 showed that glass coverslips in oil offer the highest reliable single-particle signal-to-noise ratio (SNR) for spectroscopy; however, correlative measurements require TEM-compatible substrates. We therefore used SiO<sub>2</sub> TEM grids as the best available compromise: their refractive index is close to that of the immersion-oil objective (minimising index mismatch relative to SiN or carbon-based membranes), and as shown in Sec. 4.2.2 they can achieve sufficient SNR for single-particle spectroscopy despite additional background challenges. Practically, the SNR on such grids is high enough for the particle sizes studied here but can vary from grid to grid because strain, height variations, and thickness inhomogeneities randomly modulate the background (see Figure 4.10).

Large-area microscopy scans were recorded to locate individual particles and select suitable targets for spectroscopy. We used an illumination wavelength of 700 nm for the overview because it is close to the expected longitudinal LSPR of the 20 × 75 nm AuNRs

and thus enhances contrast for reliable localisation. Figure 5.1a shows an overview scan of a  $55 \times 55 \mu\text{m}$   $\text{SiO}_2$  TEM window containing  $20 \times 75 \text{ nm}$  AuNRs (supplier-specified nominal dimensions in solution). The marked particles are precisely those analysed in this chapter. Each nanoparticle appears as a compact, near-diffraction-limited spot against a suppressed background due to the quasi-dark-field geometry, enabling reliable localisation and target selection across the full field of view. The edges of the TEM window are clearly visible, and the remaining background originates from substrate roughness and surface strain.

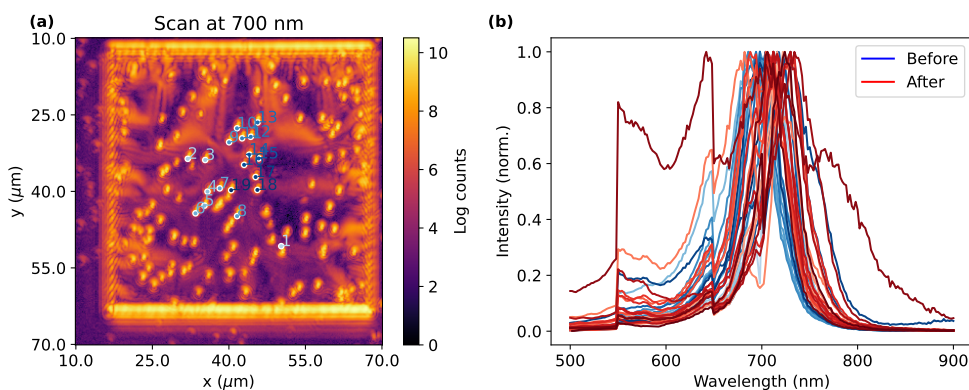


Figure 5.1: (a) Overview microscopy scan at 700 nm of a  $55 \times 55 \mu\text{m}$   $\text{SiO}_2$  TEM window with  $20 \times 75 \text{ nm}$  AuNRs (nominal supplier dimensions). Markers indicate the 23 positions selected for spectroscopy; these are the particles analysed in this chapter. (b) Instrument response function (IRF)-corrected APD scattering spectra at the marked positions. The compact, near-diffraction-limited spots in (a) enable precise target selection, while the visible TEM window boundaries and residual background illustrate the additional challenges on  $\text{SiO}_2$  TEM grids (SNR remains sufficient but varies from grid to grid).

Within this window, we performed 23 dark-field spectroscopy measurements using the max-counts autofocus, which Chapter 4 showed to be robust for the particle sizes studied here and faster than the ellipsoid-based method. Of these 23 particles, 15 were retained after quality selection (criteria detailed in Sec. 5.4). The corresponding IRF-corrected spectra are shown in Fig. 5.1b. Across multiple windows, we measured 41 positions, obtaining 26 successful single-particle spectra; from these, 20 were retained for the pre-TEM population analysis, and 10 of those had successful paired measurements after TEM exposure (Sec. 5.4).

### 5.2.1 Spectral model and fit strategy

To enable quantitative comparison between measured and simulated spectra, the IRF-corrected spectra are fitted with an analytical model that accounts for both the intrinsic plasmonic lineshape and the instrumental broadening of the tunable filter. For longitudinal modes dominated by a single resonance, the intrinsic lineshape is well described by a Lorentzian (radiative and nonradiative damping). The tunable wavelength filter in our setup contributes an instrumental Gaussian broadening of approximately 2.5 nm (see Sec. 4.2.2), so the measured spectrum is modelled by the convolution of a Lorentzian with this Gaussian, i.e. a Voigt-like model.

To ensure consistent comparison, the convolved model is evaluated at the same discrete wavelength points as the measured spectrum rather than on an interpolated grid. To avoid artifacts unrelated to the plasmon response (tunable-filter grating switch around 650 nm and increased noise in the outer tails), we restrict the fit to the resonance region, operationally defined as the interval between the full width at half maximum (FWHM) bounds of the peak. The fitted model is then extended beyond the fit range to define the smooth reference lineshape used for the error metric. The spectral fit serves two purposes: (i) to extract peak position and linewidth as parameters for comparison with simulation, and (ii) to provide a clean fitted lineshape as the reference for the spectrum-error metric from Chapter 2, thereby avoiding the influence of noise and filter artifacts when making quantitative comparisons.

Figure 5.2a illustrates a corrected spectrum (blue markers) fitted with the convolved Voigt-like model (red solid line); the red dashed line shows the fitted model extrapolated outside the fit range to illustrate the quality of the lineshape description in the spectral tails. The corresponding HAADF-STEM image for morphology extraction is shown in Fig. 5.2b.

## 5.3 Morphology extraction and environment determination

After optical measurements, the TEM grid was transferred to the TEM. We re-located the same region and the same particles using the grid coordinates, low-magnification overview images, and the particle localisation map from the optical overview scan. This allowed one-to-one correlation between optical spectra and structural images for each selected nanoparticle and also allowed us to discard any non-single-particle measurements in this analysis.

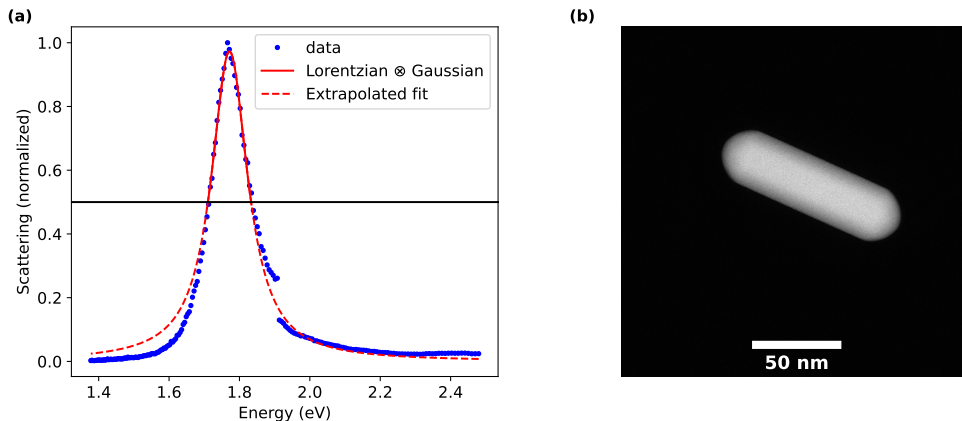


Figure 5.2: (a) IRF-corrected APD spectrum of a single AuNR (blue markers) fitted with a Lorentzian convolved with a Gaussian of standard deviation 2.5 nm (red solid line). The red dashed line shows the extrapolated fit outside the fit range; the fit is restricted to the full width at half max (FWHM)-bounded resonance region to avoid the filter transition and noisy tails. (b) HAADF-STEM image of the same AuNR (particle 2 in Figure 5.1) used to extract geometric parameters for simulation (see Sec. 5.3).

We deliberately use the 2D morphology workflow from Chapter 3 for this chapter because: (i) for rod-like particles, the 2D protocol (thresholding + 2D geometric fitting) provides sufficient morphological information to fix the dominant geometric degrees of freedom (length and diameter); (ii) electron tomography is time-consuming, making it impractical for the population-level statistics that are one of our goals; and (iii) current commercially available  $\text{SiO}_2$  TEM grids do not support the large tilt range (approximately  $\pm 75^\circ$  or higher) needed for high-quality tomographic reconstruction, membrane geometry and support bars limit tilt and introduce shadowing that compromises reconstruction quality. This practical limitation motivates ongoing development of custom grids in the group. We therefore frame the 2D approach as the best currently available option for population studies, with 3D reconstructions reserved for representative subsets in future work.

Using the thresholding and 2D fitting protocol developed in Chapter 3 (see Appendix 3.A.2.8 and Sec. 3.A.2.5), we extracted rod length and diameter from HAADF-STEM images (e.g. Fig. 5.2b). These dimensions were then used as fixed geometric input for electromagnetic simulations using the MNPBEM toolbox (Chapters 2–3). To account for the local environment, we employed a minimal one-parameter

effective-medium model,

$$\epsilon_{\text{eff}} = f \epsilon_1 + (1 - f) \epsilon_2, \quad \epsilon_1 = 1 \text{ (air)}, \quad \epsilon_2 = 2.28 \text{ (oil)}, \quad f \in [0, 1]. \quad (5.1)$$

For each  $f$ , we computed the simulated spectrum and convolved it with the same 2.5 nm Gaussian as in the measurement model to ensure consistent instrumental broadening.

This linear mixing is the simplest heterogeneous-environment description with only one free parameter  $f$ , which keeps the inverse problem well posed with the available data. Physically,  $f$  is an *effective* parameter that captures the combined influence of all environmental contributions (ligands, substrate contact fraction, nanoscale gaps, local chemistry) projected onto a single scalar; it is not solely a literal oil-versus-air fraction at the nanoscale. Using oil as the upper bound is consistent with typical refractive indices for cetrimonium bromide (CTAB) ( $n \approx 1.45$ – $1.50$ ) and  $\text{SiO}_2$  ( $n \approx 1.46$ ). In a small number of cases (6 out of 26 particles), unconstrained optimisations suggested values slightly outside  $[0, 1]$ ; since such fractions are unphysical within the linear-mixing model, these particles were excluded from further analysis. Practically, these exclusions tended to coincide with resonances falling partly outside the simulated wavelength sweep or spectra affected by residual artifacts, thus they were also removed on quality grounds. The model bound and how to extend it to higher-index or lossy surroundings is revisited in Secs. 5.4 and 5.5.

### 5.3.1 Determining the effective environment for a single particle

With the measured spectrum, TEM-extracted geometry, and effective-medium model in hand, we determine the optimal  $f^*$  for each particle by comparing simulated and measured spectra across the discrete set of  $f$  values.

We evaluated three candidate metrics for extracting  $f^*$ : (i) the spectrum error defined in Chapter 2 (Eq. 2.7), evaluated with the fitted experimental lineshape as the reference; (ii) the difference in peak position  $\Delta\lambda$  between simulation and measurement; and (iii) the difference in FWHM. Figure 5.3 summarises these comparisons for a representative particle (particle 2; the same as in Fig. 5.2). Panel (a) shows that the spectrum-error metric can be non-monotonic and somewhat noisy across discrete  $f$ . This is partly a consequence of the global nature of the metric (it integrates spectral differences across the full wavelength range including noise-dominated regions away from the resonance peak) and partly reflects an irreducible shape mismatch when only the real part of the environment is varied. Panel (c) shows that  $\Delta\text{FWHM}$  varies only

weakly with  $f$ , because linewidth is governed by damping channels that depend on losses (imaginary permittivity) in both the particle and the environment, which are not captured by the current real-valued model.




Figure 5.3: Representative single-particle optimisation (particle 2). (a) Spectrum error (Eq. 2.7 in Chapter 2) comparing convolved simulations to the *fitted* experimental lineshape for different air fractions  $f$ . (b) Difference in peak position  $\Delta\lambda$  (blue markers) with linear regression (red dashed line); the green marker indicates the x-intercept that defines  $f^*$ . (c) Difference in FWHM between simulation and measurement vs.  $f$ , showing relatively weak variation. The approximately linear  $\Delta\lambda(f)$  makes the x-intercept a robust estimator compared to (a) and (c).

By contrast, panel (b) shows that  $\Delta\lambda(f)$  is approximately linear, physically expected since peak position is primarily governed by the real part of the effective permittivity, which is the quantity directly varied in the model. We therefore fit a straight line to  $\Delta\lambda$  vs.  $f$  and take the x-intercept as the optimal fraction  $f^*$  for which the simulated peak aligns with the measured one ( $\Delta\lambda = 0$ ). This choice yields a robust and well-defined estimator for  $f^*$  for each particle.

## 5.4 Sensitivity and reproducibility of the dielectric-environment extraction

We applied the workflow to the full set of measured particles and retained 20 of 26 after quality selection (exclusions for: insufficient spectrum quality, non-single particles, or non-convergent  $\Delta\lambda(f)$  regressions). Across the retained set, the per-particle optimisations yielded an average effective-medium fraction

$$\bar{f}^* = 0.71 \quad (\text{std} = 0.08),$$

corresponding in the linear model  $\epsilon_{\text{eff}} = f\epsilon_{\text{air}} + (1-f)\epsilon_{\text{oil}}$  to an average  $\epsilon_{\text{eff}} \approx 0.71 \cdot 1 + 0.29 \cdot 2.28 \approx 1.37$  with a standard deviation of  $\sim 0.10$  as shown in Figure 5.4a.

Thus, the average effective environment lies appreciably closer to air than to oil. Although counter-intuitive given that the particles are immersed in oil, this is physically plausible because the LSPR samples the environment only within a few nanometers of the surface. The CTAB ligands form a bilayer on the gold surface with hydrophobic tails pointing inward, creating a low-density hydrophobic interior that resists wetting by oil, meaning the immediate near-field environment is dominated by the low-index CTAB interior rather than bulk oil. Additionally, partial de-wetting at the particle–substrate interface or nanoscale air gaps above the rod could further reduce the effective permittivity sensed by the plasmon.

At the same time, the single-parameter linear-mixing model conflates multiple environmental contributions (CTAB interior, substrate-contact geometry, trapped air/gaps) into a single scalar  $f^*$ , making it difficult to attribute the extracted value to any one physical contribution independently. More detailed models (e.g. ligand–substrate/oil mixing, or core–shell representations with explicit ligand layers) would likely yield more physically interpretable parameters.

An important practical implication of the observed spread (std  $\approx 0.08$  in  $f^*$ ) is that assuming a single uniform environment for all particles, as is typically done in ensemble or non-correlative measurements, introduces a systematic error in the predicted peak positions whose magnitude is set by this spread. To quantify this, we simulated all particles using the population mean  $\bar{f}^*$  as a common environment and compared the predicted peak positions to the measurements (Fig. 5.4d, blue). The resulting deviations directly benchmark the error introduced by not performing per-particle environment extraction, and demonstrate the value of the correlative workflow for quantitative single-particle plasmonics.

Optical measurements were deliberately performed before TEM characterisation, since electron-beam exposure is known to modify the local dielectric environment through ligand modification and carbonaceous deposition. We exploit this perturbation as a sensitivity test: by measuring the same particle before and after TEM exposure, we compare the optical response under two slightly different environmental conditions on an identical geometry, providing a stringent test of the framework’s ability to detect subtle environmental changes.

Figure 5.4b shows representative spectra *before* and *after* TEM for a single particle (particle 2), together with simulations using the effective-medium model. A clear red-shift is observed after TEM, consistent with an increased  $\epsilon_{\text{eff}}$  due to a carbon-rich overlayer. Importantly, carbonaceous deposition increases the local refractive index and could in principle push the post-TEM environment beyond the nominal oil bound

#### 5.4. Sensitivity and reproducibility of the dielectric-environment extraction

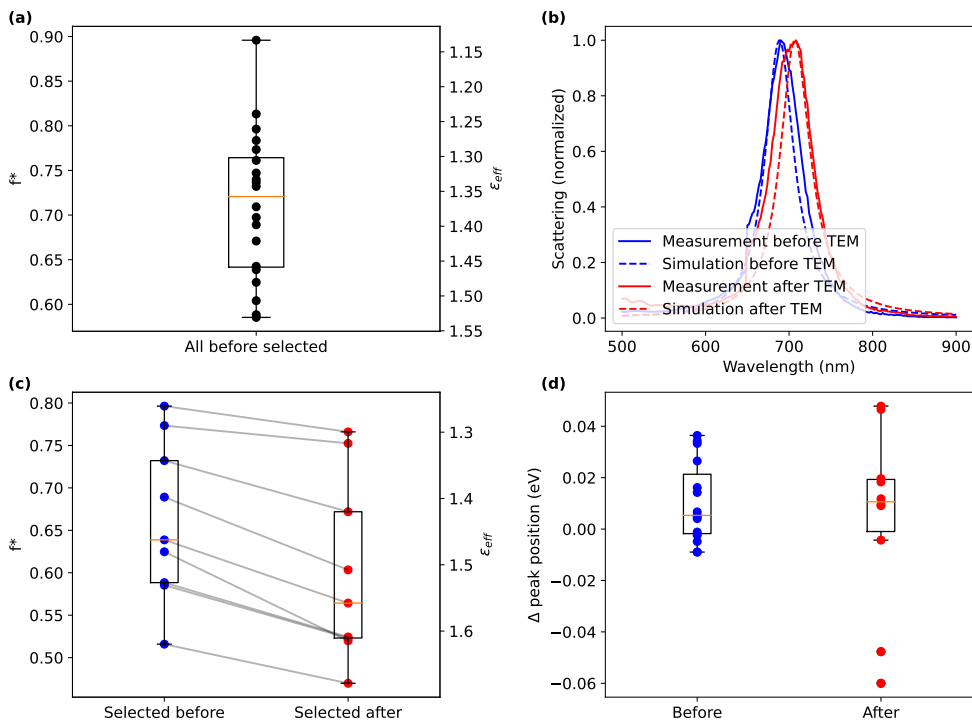


Figure 5.4: Population statistics and TEM-exposure effect. (a) Boxplot of extracted  $f^*$  across the pre-TEM population (**20** particles); the secondary axis maps  $f^*$  to  $\epsilon_{\text{eff}}$ . (b) Representative scattering spectra of a single AuNR before (blue) and after (red) TEM, together with simulations using the effective-medium model (blue and red dashed lines, respectively). A clear red-shift after TEM is evident, consistent with an increased effective environmental permittivity due to beam-induced carbonaceous deposition. (c) Paired comparison of  $f^*$  (or  $\epsilon_{\text{eff}}$ ) before and after TEM for the same **10** particles, showing a systematic shift across the paired subset. (d) Boxplots of peak-position differences (simulation minus measurement) across particles when using a common-environment model, before and after TEM. The larger spread after TEM demonstrates that a single common environment becomes increasingly inadequate once beam-induced modifications introduce particle-specific environmental heterogeneity.

of the current model, which would correspond to  $f^* < 0$  and is unphysical within Eq. 5.1. In practice, given the relatively high air-leaning  $\bar{f}^*$  before TEM and the subtle nature of the beam-induced perturbation, we do not expect this to occur systematically, but it is a limitation to keep in mind when applying the pre-exposure model to post-exposure data.

Panels (a), (c), and (d) present the population-level statistics. Panel (a) shows a boxplot of extracted  $f^*$  values (and corresponding  $\epsilon_{\text{eff}}$  on a secondary axis) *before*

TEM for the 20 retained particles. Panel (c) shows a paired comparison of  $f^*$  before vs. after TEM for the same 10 particles for which both measurements were successful (the pre-TEM mean within this paired subset is slightly different from the full-population mean due to selection, not a systematic effect). Quantitatively, across this paired subset,

- the mean fraction shifts from  $f_{\text{before}}^* = 0.66$  to  $f_{\text{after}}^* = 0.59$ ;
- correspondingly,  $\epsilon_{\text{eff}}$  shifts from  $\approx 1.43$  to  $\approx 1.51$ .

Although small in absolute terms, this shift is systematic across the paired population, demonstrating that the framework is sensitive to subtle changes in the local dielectric environment at the single-particle level. What makes the correlative approach uniquely powerful is not just absolute refractive-index sensitivity (which is well known for plasmonic sensors), but the ability to resolve *individual* environmental responses at the single-particle level with geometry held fixed by TEM.

Panel (d) shows peak-position differences (simulation minus measurement) under a common-environment model before and after TEM. The substantially larger spread after TEM reflects the highly individual nature of beam-induced modifications: each particle accumulates a different amount of carbonaceous deposition depending on its local geometry and position on the grid, producing particle-specific environmental changes that would be averaged out and invisible in ensemble measurements. Because particle geometry is fixed from TEM, environmental changes can be unambiguously separated from morphological ones, a decoupling that is impossible in standard plasmonic sensing.

## 5.5 Linewidths and amplitudes: observations and model limitations

Peak position, linewidth, and amplitude are independent observables that carry complementary physical information: the peak position probes the real part of the surrounding dielectric response; the linewidth encodes damping channels (radiative, bulk nonradiative, surface-scattering, and interface-related losses); and the amplitude reflects both the scattering cross-section and local-field factors, and is sensitive to loss in the environment.

We expected that carbonaceous contamination and ligand cross-linking could broaden the plasmon linewidth after TEM exposure through additional absorption

losses in the surrounding layer. However, the measured FWHM changes are not systematically positive (Fig. 5.5a). A plausible explanation is that the TEM-induced red-shift moves the resonance to a spectral region where the imaginary part of the gold dielectric function is smaller (which is indeed the case in the measured LSPR range), thereby reducing the nonradiative damping contribution from interband absorption and partially offsetting the broadening from the carbonaceous shell. The net linewidth change therefore depends on the balance between these competing contributions, which varies from particle to particle depending on the extent of carbonaceous deposition, yielding changes of varying sign across the population.

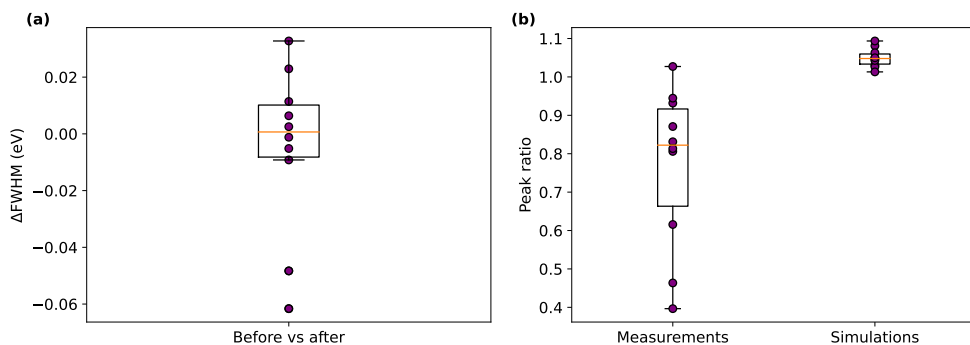


Figure 5.5: (a) Change in fitted FWHM (after minus before TEM) across particles shows no systematic broadening, consistent with competing effects from added environmental absorption and changes in Au nonradiative losses due to the resonance red-shift. (b) Ratios of peak heights (after/before) indicate a systematic amplitude reduction after exposure; discrepancies with the purely real effective-medium simulations highlight the need to include a lossy environmental contribution. The reduced spread of simulated amplitudes arises because a common effective environment was used for those simulations, while measurements reflect rod-specific variation.

By contrast, the peak amplitude is frequently reduced after TEM (Fig. 5.5b). This behaviour is not captured by simulations that vary only the real part of the environment via  $\epsilon_{\text{eff}}$ . An absorbing overlayer introduces an imaginary component to the environment, reducing the net scattering amplitude without a simple one-to-one mapping to Lorentzian width within the current model. Incorporating a lossy effective medium or a thin absorbing shell would likely improve quantitative agreement for both amplitude and linewidth. Practically, the amplitude reduction is a potentially useful observable in its own right: if combined with position and width, it could enable extraction of the imaginary part of the effective environment, providing additional information about lossy surroundings (not only carbon from TEM exposure) beyond what the peak position alone reveals. Finally, we note that the simulated amplitude

spread in Fig. 5.5b is narrower than in the measurements. This is expected because a common effective environment was used for the simulations shown there, while the measurements reflect true rod-to-rod variability (both before and after TEM).

## 5.6 Conclusions and outlook

This chapter demonstrates *quantitative, single-particle extraction of the local dielectric environment* by correlating IRF-corrected dark-field spectroscopy on TEM grids with electromagnetic simulations using geometry fixed from TEM. First, we extracted a per-particle effective environmental parameter  $f^*$  by bringing simulated and measured spectra into quantitative agreement. Across 20 particles, we obtained a consistent population mean  $\bar{f}^* = 0.71$  (std = 0.08), corresponding (in the linear model) to an average  $\epsilon_{\text{eff}}$  closer to air than oil. We provided a physical interpretation (dominant ligand-shell influence within a few nanometres and possible nanoscale de-wetting) and identified the limitations of a single-parameter linear model for full physical attribution. Second, we validated sensitivity by deliberately perturbing the local dielectric environment via TEM exposure and observed a systematic red-shift attributable to carbonaceous deposition. On a paired-population level, the average effective parameter shifts from  $f_{\text{before}}^* = 0.66$  to  $f_{\text{after}}^* = 0.59$ , corresponding to an  $\epsilon_{\text{eff}}$  change from  $\approx 1.43$  to 1.51. The large particle-to-particle variation in the magnitude of this shift demonstrates that beam-induced environmental modifications are highly individual, a heterogeneity that is averaged out in ensemble measurements and only accessible through single-particle correlation with geometry fixed from TEM. Third, by analysing linewidths and amplitudes we showed where the minimal model falls short: non-systematic FWHM changes and systematic amplitude suppression after TEM indicate lossy environmental contributions and interface-related damping not captured by a purely real  $\epsilon_{\text{eff}}$ . These observations highlight that isolating specific damping channels requires models beyond a single real-valued environmental parameter, motivating explicit inclusion of a lossy shell or complex effective medium as the natural next step toward a quantitative damping analysis.

Looking forward, the correlative framework established here provides a controlled setting for testing specific physical models, the overarching goal highlighted in Chapter 2. Extending the environment model to include lossy contributions (thin absorbing shells, complex effective media) and upgrading the morphology pipeline to 3D for selected particles will enable isolation and quantification of distinct damping channels (surface scattering vs. interface-related losses) and broaden applicability to catalysis and sensing, where small environmental changes are the signal of interest.

## Chapter 6

---

# Melt-cast – Reshaping Dynamics of Gold Nanorods

*Melt-cast – pouring hot liquid into a container and leaving it to become solid*

**Abstract:** This chapter investigates how gold nanorods dynamically reshape under pulsed optical excitation, transforming in real time much like molten metal cast into new forms. Using synchronized laser pulses within an in situ TEM, we track structural evolution on a pulse-by-pulse basis and reveal highly non-equilibrium pathways, including reversible, liquid-like reshaping, defect formation, long-range mass redistribution, and metastable immobilization within clusters. Heat-diffusion modelling shows that each pulse drives the system into a regime of rapid atomic mobility before ultrafast cooling, enabling repeated exploration of kinetically accessible configurations rather than equilibration to a single morphology. These experiments establish pulsed-illumination TEM as a powerful framework for visualizing and controlling nanoscale transformations, offering mechanistic insight into non-equilibrium structural dynamics relevant to reconfigurable materials and time-dependent catalytic functionality.

The content of this chapter is submitted as M. Dieperink, L. Monin, I. Ohnishi, T. Sasaki, Y. Jimbo, E. Garnett, W. Albrecht, “In-Situ Laser Reshaping Dynamics of Gold Nanoparticles in the TEM”.

## 6.1 Introduction

Having established in Chapter 3 how nanoscale morphology can be extracted, quantified, and incorporated into reliable optical simulations, we now move from static structure toward structural dynamics. The underlying motivation remains continuous with the rest of the thesis: a quantitative understanding of plasmonic nanoparticles ultimately requires not only knowing how morphology shapes the optical response, but also how optical excitation reshapes the morphology itself. Here we therefore treat morphology as a dynamic variable, exploring how gold nanorods reorganize, redistribute mass, and undergo structural transitions when driven far from equilibrium. By combining pulsed optical excitation with in situ TEM, we directly observe these transformations with single-pulse temporal resolution, enabling mechanistic insight into reshaping pathways that cannot be inferred from static measurements alone.

The structure and function of crystalline materials are traditionally understood through the lens of their equilibrium lattice parameters, which are stabilized by strong interatomic bonds. At the nanoscale, however, this thermodynamic rigidity gives way to a rich landscape of metastable configurations whose energetics are dominated by surfaces, interfaces, and defects. In such small systems, kinetics can outweigh thermodynamics, enabling the synthesis of nanocrystals with tailored, yet still imperfectly controlled, morphologies and surface terminations. Because catalytic [129], [130], optical, magnetic [131], [132], and mechanical properties are acutely sensitive to the arrangement of atoms at facets, edges, and defects, the ability to manipulate atomic configurations with precision, and ideally in real time, remains one of the major frontiers in nanoscience.

Among the different routes to reshape nanoscale matter, pulsed optical excitation has emerged as an especially powerful driver of atomic motion. Metal nanoparticles, in particular, exhibit strong coupling between light, electrons, and phonons through their localized surface plasmon resonances (LSPRs). Absorption of ultrafast optical pulses generates steep, spatially confined temperature gradients and transient stress fields on sub-picosecond timescales [133], [134]. These impulsive excitations can launch rapid heating–cooling cycles and acoustic strain waves, pushing the lattice into highly non-equilibrium regimes that are inaccessible under conventional thermal processing [135]. As a result, pulsed illumination has been used to reshape nanoparticles, redistribute alloy components, and generate unusual defect motifs, all while preserving the overall nanoscale architecture [136], [137], [138], [139], [140], [141], [142], [143], [144], [145], [146], [147], [148]. Yet despite these advances, the atomic pathways by which nanocrystals melt, diffuse, reform, twin, or exchange mass under such extreme and

transient conditions remain inferred primarily from before-and-after characterization and molecular dynamics simulations [149], [150].

Real-time visualization of laser-driven atomic motion is essential for building a mechanistic foundation for non-equilibrium nanoscale transformations. Such insight is also becoming increasingly important in the context of dynamic and resonant catalysis, where time-dependent modulation of surface structure, strain, and energetics is predicted to overcome conventional scaling relations and enable new modes of reaction control. Theoretical work suggests that even small, rapidly imposed structural perturbations, such as transient lattice strain or facet reshaping, could dramatically alter adsorbate binding and reaction pathways. However, experimental access to these processes at the relevant length and time scales remains extremely limited.

Transmission electron microscopy (TEM) offers a unique window into these ultrafast-to-millisecond structural dynamics by providing atomic-resolution imaging of individual nanocrystals while they are stimulated by optical pulses. Recent in situ TEM studies have revealed that laser excitation can generate crystal defects, induce partial melting, and reshape metal nanoparticles, yet the temporal evolution of these processes and their interplay across multiple particles remain largely unexplored.

Here, we directly image the structural evolution of gold nanorods under controlled pulsed illumination inside a TEM. We uncover a highly dynamic regime in which atoms migrate across several nanorods, defects nucleate and rotate, and particle morphologies reorganize on millisecond timescales. These observations provide unprecedented experimental access to the atomic mechanisms of laser-driven restructuring, revealing pathways of mass transport, strain relaxation, and transient defect stabilization that have not been previously visualized. Our findings establish in situ pulsed-illumination TEM as a powerful platform for interrogating non-equilibrium nanoscale processes and open new avenues for designing materials whose functionalities emerge from dynamically reconfigurable atomic structures.

## 6.2 Experimental setup and imaging scheme

Gold nanorods (Au NRs) with average dimensions of  $40 \times 68$  nm (NanoSeedz Ltd.) were deposited onto TEM grids for in situ optical excitation. Their structural evolution was monitored using a JEOL GrandARM TEM (300 kV) equipped with an IDES Luminary Micro laser-incoupling system. A 577 nm modulated CW laser was focused to a  $40 \mu\text{m}$  spot on the sample to excite the NRs near their simulated longitudinal LSPR (580 nm, Figure 6.A1). Continuous illumination rapidly damaged the thin support film, necessitating the use of short  $400 \mu\text{s}$  pulses at an average power of

50 mW. To temporally isolate individual excitation events, the pulse repetition rate was synchronized with the 40 ms camera integration time so that each frame contained exactly one optical pulse. Because the laser is on for only about 1% of the frame duration, the remaining time provide stable imaging conditions. This scheme allows us to record high-quality snapshots of the nanorods per each pulse, effectively capturing the structural consequences of a single photothermal heating event on a frame-by-frame basis.

### 6.3 Initial morphology and pre-illumination characterization

Before initiating optical excitation, we recorded several reference frames to establish the initial morphology of the gold nanorods. Upon switching on the laser, we immediately observed a slight mechanical drift of the support film, likely due to rapid photothermal expansion, which caused some particles to move out of the field of view. Figure 6.1a therefore corresponds to the last stable frame acquired six frames prior to the first laser pulse, and represents the closest pre-excitation configuration in which all nanorods remained visible. The colloidal Au NRs are known to be single-crystalline with an fcc lattice structure [136], a characteristic confirmed by their homogeneous contrast in the initial frame. Variations in the apparent intensity among different nanorods arise from thickness differences and from diffraction contrast. In bright-field TEM, diffraction contrast occurs when the crystal is oriented such that certain lattice planes satisfy the Bragg condition, leading to enhanced scattering of electrons out of the direct beam. Nanorods oriented closer to a zone axis therefore appear darker, whereas those tilted away from strong diffracting conditions appear brighter. Following the sequence of laser pulses, pronounced morphological changes became evident. Several nanorods reshaped substantially (cyan arrows), and some appeared to have vanished from the field of view (blue arrows), as shown in Figure 6.1b. In addition to overall morphological evolution, twinning defects emerged (white arrows), visible through characteristic diffraction-contrast features. Such laser-induced defect formation is well aligned with previous reports of plasmon-driven restructuring in gold nanocrystals, consistent with the behavior described in earlier studies as detailed in the introduction.

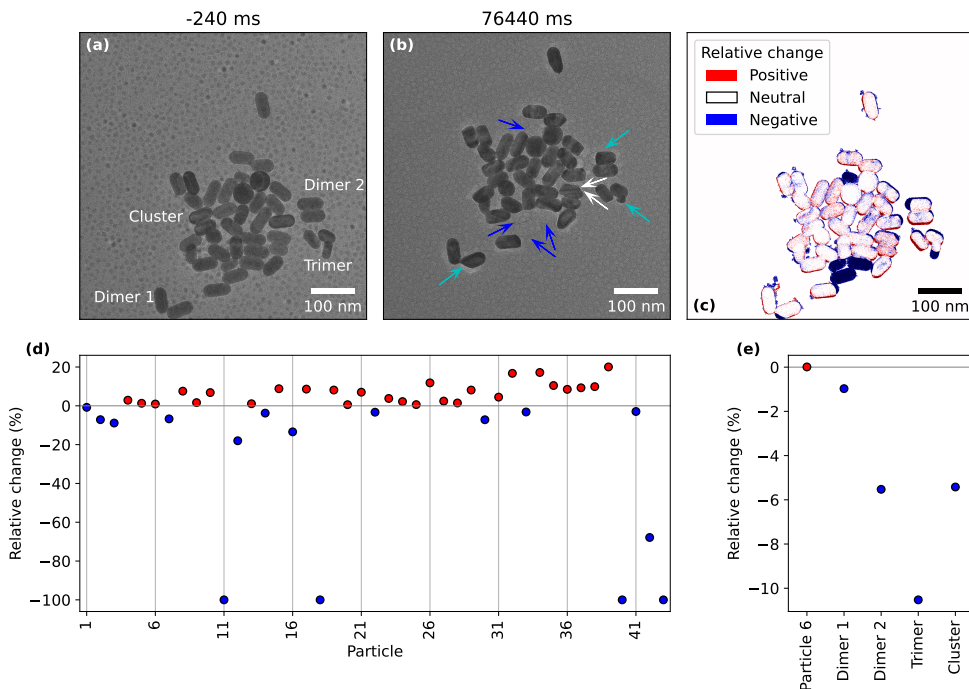


Figure 6.1: The total structural change over the course of the laser excitation experiment. 6 frames before (a) and the last frame (b) of the laser excitation experiment with highlighted changes and twinning defects in (b). The aligned subtraction of the binarized versions (c) (see Figure 6.A2a,b). The relative change in pixels observed for each particle (d) and grouped isolated islands (e) from our manual segmentation analysis (see Figure 6.A3a).

## 6.4 Quantification of reshaping through projection difference mapping

To quantify the structural evolution, we analyzed the changes in the projected area of each nanorod by segmenting their 2D bright-field TEM projections from the background (Figure 6.A2). The difference map shown in Figure 6.1c represents the pixel-by-pixel change between the final and initial segmented images, and therefore reflects modifications in the projected footprint of the nanorods. All analysis is necessarily based on 2D projection images; thus, the measured changes correspond to lateral redistribution of projected mass and do not capture possible variations in particle thickness. Slight apparent changes can arise from sample drift or small mechanical contractions of the support film during pulsed illumination (Figure 6.A4), which we corrected for as described in Appendix 6.A.2. However, the dominant

signals in Figure 6.1c clearly originate from significant reshaping of the nanorods. The four nanorods that “disappear” from the field of view (blue arrows) manifest as extended blue regions, indicating a substantial reduction of their projected area. The remaining nanorods, in contrast, show systematic increases in their projected footprint (red regions), indicating that material redistributed within the cluster. We emphasize that the complex contrast mechanisms in bright-field TEM, combined with projection limitations, preclude direct inference of absolute particle thickness or total volume. Nevertheless, the difference maps reveal clear and substantial rearrangements of projected mass within the ensemble, consistent with extensive nanoscale material redistribution under pulsed optical excitation.

To clarify whether the apparent loss of projected area reflects true mass loss or positional redistribution, we quantified the relative change in segmented pixel area for each nanorod (Figure 6.1d; particle labels in Figure 6.A3) and then examined how these changes depend on particle proximity by comparing isolated rods, dimers, trimers, and the main cluster (Figure 6.1e). Four nanorods that disappear from projection are clearly identified by a relative change of -100%, while the remaining particles show more modest relative changes consistent with the redistribution patterns in Figure 6.1c. Importantly, these values are derived from two-dimensional projection images and therefore reflect changes in lateral footprint rather than direct measurements of three-dimensional volume. A systematic trend emerges when grouping the particles by local configuration. The fully isolated nanorod (particle 6) shows no measurable loss of projected area, whereas particles in close proximity exhibit progressively larger relative decreases: approximately 2% and 6% for the two dimers, 10% for the trimer, and 6% for the main cluster. When summing over all particles, the ensemble exhibits an overall projected-area loss of about 11%, which, assuming a comparable reduction along the thickness direction, corresponds to an estimated volume loss on the order of 15%. This indicates that some material loss, likely via sublimation or evaporation, does occur under the applied pulsed illumination conditions. Strikingly, however, this loss does not correlate with particle isolation. If sublimation were governed primarily by surface exposure, the isolated nanorod would be expected to lose material most readily, which is not observed. Instead, particles in close proximity exhibit larger relative losses, suggesting that collective effects, such as enhanced local heating, transient melting, and interparticle mass redistribution, play a dominant role. In this context, the apparent loss of projected area within dimers, trimers, and the cluster likely reflects a combination of modest material loss and extensive three-dimensional reshaping, including out-of-plane mass redistribution and compaction into geometries with a smaller lateral footprint. Although projection-based analysis cannot fully disentangle

these contributions, the observed proximity-dependent trends demonstrate that pulsed illumination drives substantial, collective restructuring across the ensemble. In the following, we examine the time-resolved dynamics underlying these transformations.

## 6.5 Pulse-resolved dynamics and automated tracking

Building on the pulse-synchronized imaging scheme described above, each 40 ms TEM frame captures the structural state following a single 400  $\mu$ s laser pulse, allowing us to track reshaping dynamics one excitation event at a time. To analyze these frame-to-frame changes across long sequences and many particles, we therefore switch from the manual segmentation used for Figure 6.1 to an automated difference-mapping workflow (see Section 6.A.1.2). Using this analysis, we observed three distinct categories of reshaping dynamics: (i) no reshaping, (ii) isolated reshaping, and (iii) coupled reshaping. Figure 6.2 shows representative consecutive frames for each case. In Figure 6.2a, nanorod 4 located near the center of the cluster remains essentially unchanged over the displayed time window and thus provides a useful internal control for validating the automated analysis; any apparent fluctuations in its difference maps arise from noise and minor diffraction-contrast contributions, as detailed in Section 6.A.1.2. Figure 6.2b shows nanorod 6, the isolated single rod, which undergoes clear shape changes between consecutive frames, indicating that substantial atomic reconfiguration occurs within a single pulse-to-pulse interval. Reshaping in such isolated rods typically initiates at the tips, consistent with prior observations. Finally, Figure 6.2c highlights coupled reshaping within an interacting pair: nanorods 12 and 19 exchange projected area over several consecutive pulses before the dimer abruptly becomes immobilized after frame (iv), suggesting entry into a metastable configuration that we analyze further in Figure 6.6. The specific intervals in Figure 6.2 are representative examples; notably, some rods (including nanorod 6) remain dynamically active over much longer portions of the experiment.

Dynamical reshaping was frequently accompanied by the formation and evolution of twin boundaries. We observed twinning within single nanorods (Figure 6.3a), spanning neighboring nanorods (Figure 6.3b), and in some cases extending across large parts of the cluster (Figure 6.3e,f). Notably, the spatial extent of twinning increased over the course of the experiment: at early stages twin boundaries appeared in individual or adjacent particles, whereas at later times they propagated through multiple nanorods with a common orientation. In fcc Au, coherent twin boundaries are constrained to lie parallel to  $\{111\}$  planes; therefore, for a twin boundary to appear continuous across

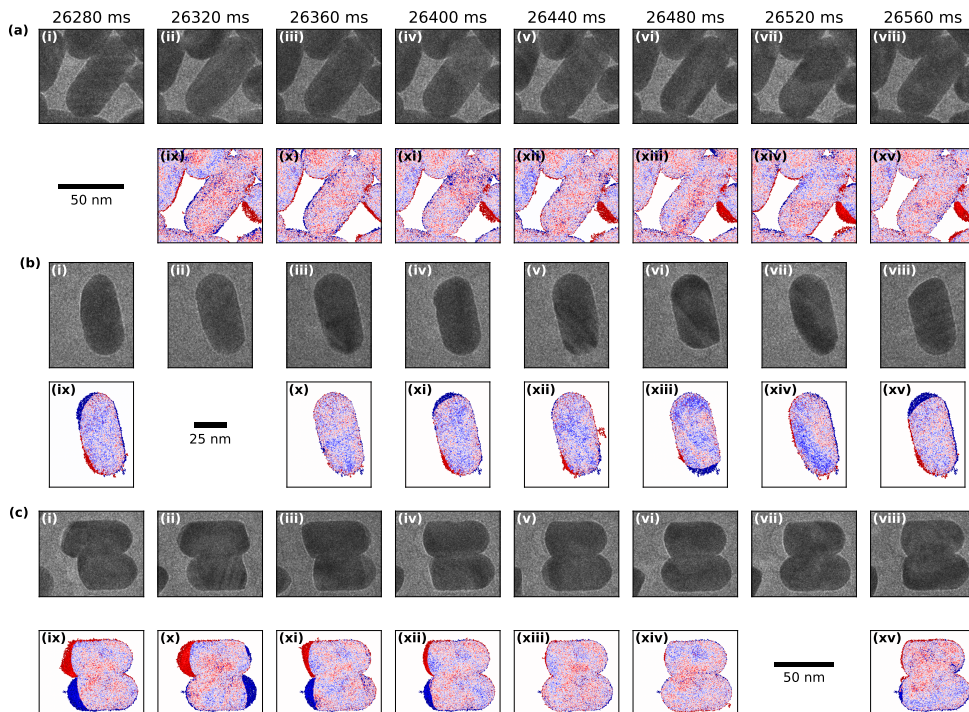


Figure 6.2: 3 categories of dynamic reshaping: no reshaping (a), isolated reshaping (b), and coupled reshaping (c). A zoom in on NR 4 (a), NR 6 (b), and NRs 12 and 19 over the course of 8 consecutive frames (i-viii) and their change throughout time (ix-xv). For (a) the first frame is used as minuend and the consecutive frames as subtrahends to show the change throughout time rather than the difference between consecutive frames. The minuend is changed for (b and c) where they are the frames that have the scalebar reported below them (ii and vii, respectively). The colors correspond to gain (red), neutral (white), and loss (blue) which is the same as for Figure 6.1.

several particles, those nanorods must share a crystallographic alignment in the [111] viewing direction. In addition to this cumulative alignment, twin boundaries within individual nanorods were highly dynamic. As shown in Figure 6.3c,d, twin-related diffraction contrast can reorient completely between consecutive frames, implying rapid restructuring of the underlying defect geometry within a single pulse-to-pulse interval. Together, the progressive emergence of cluster-wide  $\{111\}$ -aligned twins and their frame-to-frame reorientation point to a highly labile lattice under pulsed excitation. A natural hypothesis consistent with these observations is that the nanorods transiently enter a liquid-like state and recrystallize between pulses; without partial melting and resolidification, it would be difficult for many particles to develop a common [111] alignment and for twin orientations to reset on sub-frame timescales within our

experiment.




Figure 6.3: The formation of twin boundaries in individual particles (a), throughout multiple particles (b), along parts of the cluster (e), and along the whole cluster (f). A double threshold and inverted colormap was used to make the twinning defects more visible, creating the images with a white background. The original frames for (e) and (f) can be found in Figure 6.A5. (c) and (d) show two consecutive frames in which the twin boundaries completely change orientation.

## 6.6 Heat-diffusion modelling and temperature analysis

To evaluate whether the heat generated under our pulsed illumination conditions is sufficient to account for the observed lattice lability and possible transient melting, we performed heat-diffusion simulations in COMSOL (see Section 6.A.2 for details). Because the optical pulse duration ( $400 \mu\text{s}$ ) is many orders of magnitude longer than both the electron–phonon coupling time in gold (a few picoseconds) and the characteristic heat-dissipation time of the nanorod–support system, the thermal response during each pulse can be treated as quasi–steady state. The resulting steady-state temperature map is shown in Figure 6.4. Although excitation is near the longitudinal plasmon resonance of the nanorods, the simulations indicate that the dominant heat source is photothermal absorption in the carbon support, with an additional but smaller contribution from direct nanorod absorption. Taking both effects into account, we obtain a steady-state temperature increase of about  $820 \text{ }^\circ\text{C}$ . Consistent with spatial heat accumulation in a dense ensemble, nanorods located in the interior of the cluster reach slightly higher temperatures than those at the periphery [151], [152], [153]. This temperature gradient is further reinforced by the

Gaussian intensity profile of the laser beam, which results in higher optical power density near the beam center and therefore enhanced heating of particles located closer to the center of the illuminated area. The simulated temperatures are below the bulk melting point of gold (1064 °C), yet they are high enough to render surface atoms highly mobile at the nanoscale and to drive substantial reshaping. Indeed, laser- or thermally induced restructuring of Au nanoparticles has been reported at much lower nominal temperatures under continuous heating, including 223 °C [154], 250 °C [155], and 400 °C [156]. Such literature benchmarks, together with the strong size- and defect-dependent melting point depression in nanocrystals, indicate that the temperatures reached here are fully compatible with rapid surface diffusion and likely approach regimes of partial melting and resolidification under repeated pulsing. Because the pulse duration is 400  $\mu$ s and the thermal response during each pulse is quasi-steady state, the nanorods experience an extended high-temperature window for atomic diffusion and defect reconfiguration. Notably, 400  $\mu$ s is orders of magnitude longer than the nanoscale melting and structural rearrangement times reported for gold nanoparticles (typically in the ps–ns regime), so substantial reshaping can readily occur within a single excitation event. Time-dependent heat simulations shown in Figure 6.A8 further show that the system cools back toward room temperature within about 100 ns after each pulse, which is far shorter than the 39.6 ms interval between pulses, indicating that restructuring is driven primarily during the excitation window of each pulse rather than by cumulative heating between pulses.




Figure 6.4: Simulated temperature increase reached during the experiment

## 6.7 Quantification of reshaping through tip-difference mapping

To quantify the pulse-resolved reshaping dynamics revealed qualitatively in Figure 6.2, we track each nanorod in every frame and monitor where projected area is lost along its long axis. Specifically, for each rod we take the segmented outline from the first frame as a fixed reference mask, split this initial outline into two halves (upper and lower tip regions), and then compute in every subsequent frame how much projected area remains within each half of that original mask. We define a tip-asymmetry signal by subtracting the remaining projected area in the upper half from that in the lower half. In the absence of reshaping, both halves retain the same projected area within the reference outline and the difference stays close to zero; any residual fluctuations reflect segmentation uncertainty and small tracking/orientation errors. The noise level of this asymmetry signal is determined from time intervals in which the laser is off and the particles remain static (see Section 6.A.1.2). These noise boundaries are indicated by black lines in Figure 6.5. Values within the boundaries correspond to a neutral state with no detectable reshaping, whereas excursions beyond the boundaries indicate a discrete structural change in which one half of the rod loses more projected area than the other relative to the initial shape. We first validate the method on the stationary reference nanorod (NR 4 in Figure 6.2a): its asymmetry signal stays mostly within the noise limits throughout the experiment (Figure 6.5a), confirming that the analysis works as expected. In contrast, the isolated mobile nanorod (NR 6 in Figure 6.2b) shows large, repeated positive and negative excursions far outside the noise limits over nearly the entire series (Figure 6.5b), demonstrating sustained, pulse-to-pulse reshaping. A zoom-in of the NR 6 signal (Figure 6.5c), aligned with the frames shown in Figure 6.2b (see also Figure 6.A6), reveals that the asymmetry can switch sign between consecutive frames, meaning that the dominant loss alternates between the two tips on a single-pulse timescale, evidence of highly dynamic and, at times, reversible reshaping.

This asymmetry analysis also allows us to extract a pulse-resolved switching rate. We identify a “switch” whenever the asymmetry signal changes sign between consecutive frames while remaining outside the noise band, i.e., when the dominant loss shifts from one tip to the other. The number of such sign changes is counted in 1 s bins (25 frames per bin), yielding the instantaneous switching rate shown in Figure 6.5d,e for NR 4 and NR 6, respectively. The stationary reference rod (NR 4) exhibits a small, finite apparent switching rate of a few events per second (Figure 6.5d), which reflects residual segmentation/tracking noise and the deliberately

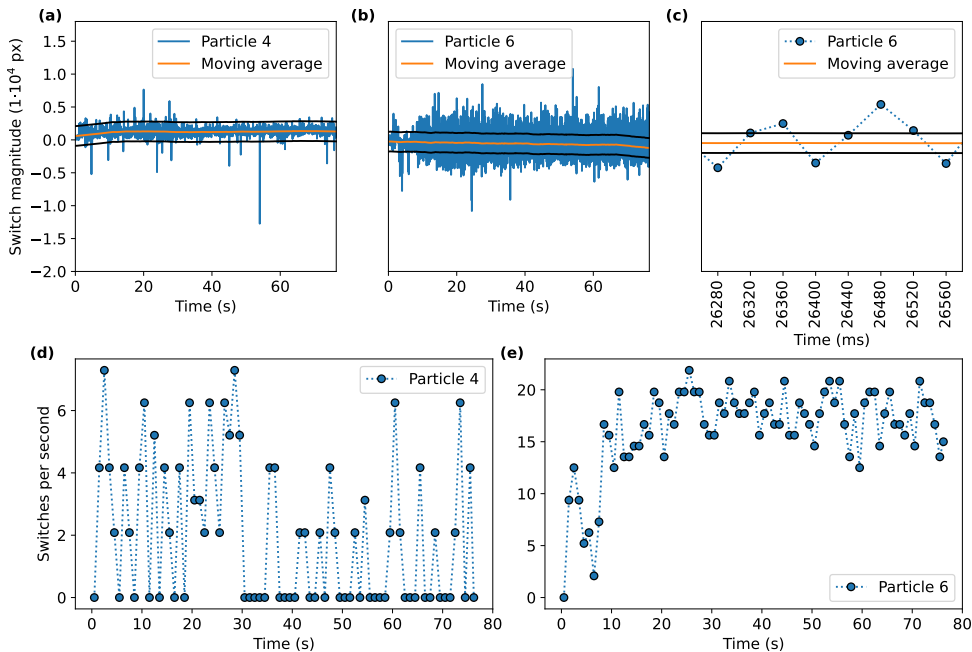


Figure 6.5: Quantification of dynamical reshaping in NRs 4 and 6. The switch magnitude of NR 4 (a) and 6 (b) and a zoom (c) that matches the frames in Figure 6.2b (displayed together in Figure 6.A6). The switch rate with 1 s binning for NR 4 (d) and 6 (e).

tight choice of the noise boundaries rather than genuine reshaping. In contrast, the isolated mobile rod (NR 6) shows a pronounced increase in switching activity during the first 20 s of illumination, after which the rate stabilizes at roughly 20 switches per second (Figure 6.5e). This temporal “acceleration” of dynamics mirrors the progressive emergence and alignment of twin defects across the ensemble described above, indicating that the system becomes increasingly susceptible to pulse-induced restructuring as the experiment proceeds. While cumulative heating might naively be expected to contribute, our heat-dissipation simulations indicate that thermal relaxation occurs on about 100 ns timescales, far shorter than the 39.6 ms laser-off interval between pulses. Thus, significant heat accumulation between pulses is unlikely to be the sole driver. The origin of the increasing switching rate therefore remains open and may involve progressive structural conditioning of the nanorods (defect and facet evolution) that enhances their susceptibility to subsequent pulse-induced reshaping.

## 6.8 Mass redistribution and metastable locking in dimers

When nanorods approach each other closely enough to exchange material, the reshaping dynamics become qualitatively richer, because structural evolution in one particle can directly reshape the other. The pulse-resolved analysis allows us to track such events by monitoring the total projected area within each nanorod’s initial outline. Figure 6.6 focuses on a strongly coupled pair (NRs 12 and 19) in close proximity, where inter-particle mass exchange is possible. Using the same projection-based metric as above, we follow the number of segmented pixels retained within the original masks for each rod. At  $t = 26.44$  s (red vertical line), NR 12 shows an abrupt decrease in projected area (Figure 6.6a), providing clear evidence for a directed transfer of projected mass away from this particle. NR 19 is the apparent recipient; however, because our automated method only detects losses within fixed reference masks, any gain in projected area by NR 19 cannot be quantified and therefore does not appear as a corresponding increase in Figure 6.6d. This limitation is detailed in Appendix 6.A.1.2 and does not affect the qualitative identification of the transfer event, which is also evident in the raw image series (image (v) in Figure 6.2c) and confirmed by manual segmentation of these frames. Strikingly, this single exchange event coincides with a sharp transition in dynamics. Before the transfer, both nanorods exhibit frequent pulse-to-pulse switching, similar to the isolated mobile rod discussed above. Immediately after the transfer, reshaping activity collapses: the switch magnitude remains within the noise band for extended periods, indicating dynamical arrest (Figure 6.6b,e). This immobilization is most pronounced for NR 19, whose asymmetry signal becomes essentially dormant after the red line. NR 12, in contrast, shows sporadic residual switching events that remain negative in sign, which we identified as a combination of segmentation artefacts and directionally biased reshaping rather than alternating between tips. Consistent with this behavioural change, the switching rate for the pair drops sharply after the exchange (Figure 6.6c,f), demonstrating that close-range coupling can drive an abrupt, metastable locking of the reshaping dynamics following material transfer.

## 6.9 Cluster-mediated reshaping and long-range correlations

The dynamics become even richer when nanorods are embedded in the dense cluster, where multiple neighbors can participate in, or mediate, structural reconfiguration. A

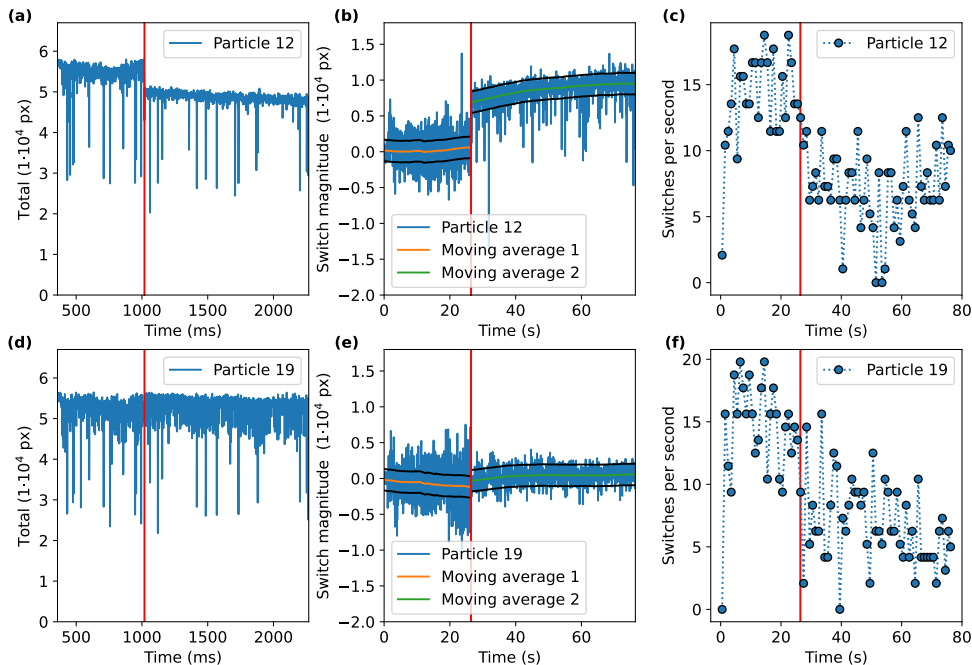


Figure 6.6: Quantification of dynamical reshaping in NRs 12 and 19. The total particle areas (a,d), switch magnitudes (b,e), and switch rates (c,f) for NRs 12 and 19, respectively. The red vertical lines at 26.44s indicate the moment of mass transfer and immobilization of the dimer. A zoom-in of the switch magnitude is plotted together with Figure 6.2c in Figure 6.A7

striking example is shown in Figure 6.7a,b for NR 11, which exhibits highly reversible behavior over extended time windows. The corresponding frame series in Figure 6.7a shows that the nanorod can completely disappear from projection and later reappear, with the transition proceeding through intermediate, liquid-like morphologies. Importantly, this disappearance-reappearance cycle is not tied to a single pulse; instead it unfolds over several consecutive pulses, indicating that the underlying mass redistribution occurs on comparatively long (multi-pulse) diffusion timescales. Despite these large excursions, the nanorod repeatedly recovers a projected footprint closely resembling its original morphology, demonstrating a remarkable degree of reversibility even after substantial restructuring. The liquid-like evolution is consistent with the high temperatures predicted by our heat-diffusion simulations and supports the view that, under pulsed excitation, nanorods can intermittently access a highly labile state that enables rapid reshaping and recovery.

A second example of cluster-mediated dynamics is provided in Figure 6.7d, where

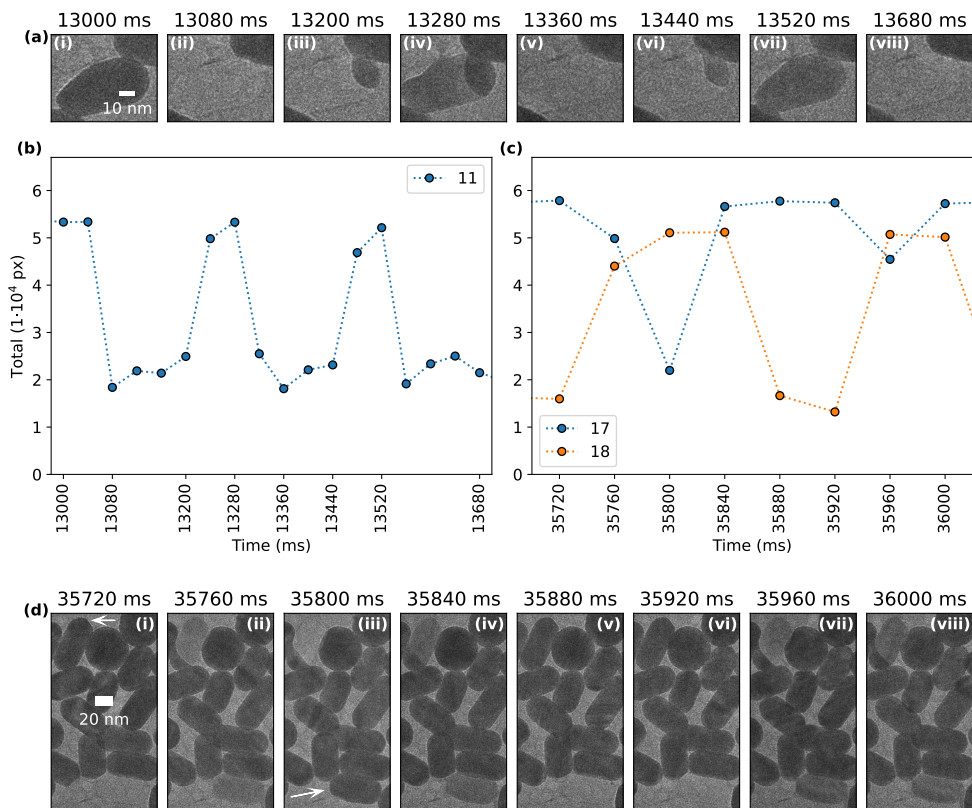


Figure 6.7: Mobility and reversibility examples of dynamic reshaping. Zoom in of cluster frames (a) and total particle areas for NR 17 and 18 (b) and reversed for NR 11 (c,d).

two nanorods (NRs 17 and 18) undergo coupled reshaping despite being separated by several intervening rods, corresponding to a center-to-center distance of hundreds of nanometres. Across eight consecutive frames, projected area appears to transfer between these two particles, again via intermediate, liquid-like morphologies in which partial signatures of both rods coexist. Notably, the nanorods located between NRs 17 and 18 exhibit comparatively small changes over the same interval. This suggests that the correlated evolution of NRs 17 and 18 does not require concerted restructuring of every intermediate particle, and may instead be mediated by a localized, transient redistribution pathway (e.g., via an intermittent surface/wetting layer or out-of-plane mass transport) that is not fully captured in projection. The complementary time traces of projected area within the fixed reference masks (Figure 6.7c) corroborate this coupled behavior: when one rod loses projected area or vanishes from projection, the other exhibits a corresponding gain in retained footprint, within the limits of the

automated segmentation (which can overestimate partially visible rods). Together, these examples demonstrate that in the clustered environment, pulsed illumination can drive both long-range, coupled mass redistribution and strongly reversible, liquid-like reshaping events that extend beyond the behavior of isolated nanorods.

## 6.10 Mechanistic interpretation: diffusion, carbon Shell, and kinetics

The rich and highly reversible reshaping dynamics observed under pulsed illumination reflect the interplay of rapid atomic mobility, interfacial interactions, and kinetic constraints imposed by the experimental environment. A central factor is the transformation of surface ligands into an amorphous or partially graphitized carbon shell under combined electron-beam exposure and laser heating. Such carbonaceous layers are well known to form in TEM experiments on ligand-stabilized nanocrystals and profoundly modify surface energetics and diffusion pathways. Rather than acting as a rigid encapsulant, this carbon shell likely serves as a kinetic moderator: it suppresses unconstrained sublimation while still permitting extensive atomic rearrangement along the metal-carbon interface.

Atomic-scale studies have shown that gold atoms exhibit exceptionally high mobility on amorphous and graphitic carbon substrates, even at room temperature. Early aberration-corrected STEM experiments directly visualized gold atoms migrating over nanometer distances on carbon films [157], while *ab initio* and molecular-dynamics simulations revealed that such jumps can occur over 1-4 nm within tens of picoseconds [158]. Diffusion is favored on graphitic terraces, whereas step edges, defects, and thickness variations act as stronger binding sites [159]. Under the significantly elevated temperatures reached during pulsed illumination in the present experiment, these diffusion processes are expected to be greatly accelerated, enabling rapid redistribution of atoms over distances far exceeding the dimensions of individual nanorods within the 400  $\mu$ s excitation window.

At the same time, gold exhibits poor wettability on graphitic carbon. Experimental measurements and simulations consistently report contact angles around 130°, reflecting the dominance of Au-Au bonding over Au-C interactions [160], [161]. Density functional theory calculations place the Au-C interaction energy in the range of only 0.01-0.02 eV for such contact geometries [162], substantially weaker than metallic bonding within gold [159]. As a result, while gold atoms and clusters can transiently spread along carbon surfaces during high-temperature excitation, thermodynamic

driving forces favor dewetting and re-compaction into more compact, gold-dominated morphologies upon cooling. This balance between high atomic mobility and weak interfacial bonding provides a natural explanation for the strongly reversible, liquid-like reshaping events observed here, including the repeated collapse and re-formation of nanorod-like geometries following each pulse.

Time-dependent heat simulations further clarify the kinetic regime of the experiment (Figure 6.A8). Although the system cools back toward room temperature within about 100 ns after each pulse, ruling out significant heat accumulation between pulses, the 400  $\mu$ s excitation duration is orders of magnitude longer than the picosecond-to-nanosecond timescales associated with nanoscale melting, surface diffusion, and structural rearrangement in gold nanoparticles. Consequently, each pulse provides an extended high-temperature window in which atoms can explore a wide range of configurations before rapidly quenching into metastable states during post-pulse cooling. The modest net mass loss observed across the ensemble suggests that sublimation does occur, but is strongly limited by carbon confinement and does not dominate the observed dynamics.

Beyond these general considerations, the experiment reveals clear evidence of history-dependent behavior. Most prominently, the switching rate of an isolated nanorod increases progressively over the first 20 s of pulsed excitation (Figure 6.5e), despite the absence of cumulative heating. This behavior is inconsistent with a purely stationary energy landscape and instead suggests that repeated excitation gradually modifies the local kinetic environment. Plausible mechanisms include progressive carbonization or annealing of the ligand-derived shell, evolution of defect-mediated pinning sites, and subtle morphology hysteresis of the nanorod itself. In this view, atomic motion during each pulse remains stochastic and extremely fast compared to the imaging rate, but the effective activation barrier for entering the mobile, reshaping state decreases with repeated cycling.

This same history dependence also provides a natural framework for understanding the apparent long-range coupling observed within clustered regions. Correlated reshaping and mass redistribution between nanorods separated by hundreds of nanometers does not require continuous material bridges or uniform participation of all intermediate particles. Instead, it can emerge from biased diffusion through a structured and evolving energy landscape, in which fast atomic motion during the excitation window is followed by repeated trapping at preferred sites upon cooling. Because these trapping sites are spatially fixed, e.g. carbon defects, thickness variations, or local curvature extrema, the system can repeatedly return to similar configurations, giving rise to both temporal “memory” in switching behavior and spatially correlated

evolution across extended distances.

A remaining question is why the nanostructures repeatedly return to rod-like morphologies rather than relaxing toward the thermodynamically stable truncated-octahedral shape expected for free gold nanoparticles. The key distinction is that the present system never approaches the equilibrium of an unconstrained particle. Carbon confinement, weak but spatially heterogeneous Au–C interactions, finite excitation times, and repeated quenching impose strong kinetic constraints that bias diffusion and trapping along anisotropic pathways. As a result, the system repeatedly relaxes into metastable, history-dependent configurations that retain memory of the original rod geometry, rather than accessing the global minimum predicted by Wulff constructions.

Taken together, these observations indicate that pulsed illumination drives a kinetically controlled, non-equilibrium process in which rapid diffusion, weak wettability, and evolving interfacial constraints combine to produce reversible reshaping, long-range correlation, and history-dependent dynamics. Rather than reflecting simple melting or evaporation, the behavior arises from repeated exploration and selective stabilization of metastable configurations within a carbon-mediated energy landscape, a regime that is inaccessible under static heating or in isolated nanostructures. More broadly, the ability to repeatedly access, bias, and stabilize specific structural configurations through time-dependent excitation points toward a route for dynamically controlling nanoscale structure, and thereby function, in operando. In particular, such stimulus-driven restructuring under pulsed or resonant excitation may provide a physical basis for dynamically tuned catalytic systems, where transient geometries, defect populations, or electronic states are selectively populated during operation, enabling catalytic behavior that cannot be achieved under steady-state conditions.

## 6.11 Conclusion

In conclusion, by performing time-resolved pulsed-laser irradiation of Au nanorods in situ inside the TEM, we uncovered highly dynamic structural processes that remain hidden in static experiments. Under pulsed excitation, the nanorods undergo pronounced reshaping and mass redistribution, ranging from reversible, liquid-like transformations to correlated behavior within clusters, on timescales far shorter than the 40 ms frame time. The 400  $\mu$ s excitation window enables rapid atomic rearrangement followed by recrystallization upon cooling, leading to pulse-to-pulse changes in crystallographic orientation and defect structure. Strikingly, the particles repeatedly return to rod-like configurations, revealing a clear “memory” of their initial morphology despite entering highly mobile regimes, consistent with kinetic

constraints imposed by the ligand-derived amorphous carbon shell and spatially heterogeneous trapping that biases diffusion and reshaping pathways. Because pulse width and repetition can be tuned, these experiments highlight a route to actively steer nanoscale structure with time-dependent stimuli, with potential relevance to dynamically controlled functionality and resonant or dynamic catalysis.

## 6.A Appendix




Figure 6.A1: BEM simulation of a single nanorod with the average size from all fitted rods in vacuum and in an effective medium where 15% is carbon.

### 6.A.1 Segmentation

#### 6.A.1.1 Manual

Literature provides us with many automated segmentation methods. However, to study static change few images have to be segmented. In that case, manual segmentation suffices. Figure 6.A2a,b show binarized images of Figure 6.1a,b, respectively. These were created by using Otsu's threshold to remove the background, followed by island size based segmentation to remove any residual noise. In ImageJ these images were segmented further by removing pixels in between particles, resulting in Figure 6.A2c,d. In the case of Figure 6.A2a, also residual noise around particles was removed. Since a 2-based (diagonal inclusion) connectivity was used for island labeling, a brush with a pixel size of 2 was used. This prevents particles from remaining connected when a diagonal line is drawn in between them.




Figure 6.A2: Thresholding and segmentation. Binarized images of Figure 6.1a,b (a,b respectively) and their manually segmented versions from ImageJ (c,d).

### 6.A.1.2 Automatic

Subsequently, we can label each individual particle and fit a shape to them using our earlier described method (Figure 6.A3a). Note that for this process it is important to remove the artefacts from the diffraction contrast by filling any holes in the particles. Figure 6.A3b shows a zoom-in of the filled version of nanorod (NR) 6 where the diffraction contrast is no longer visible. Since we do not have the time to manually segment all 2398 frames, a mask was created for each NR using the parameters from the fit. Therefore, the dimensions for each NR remained fixed, while the translation

and rotation were altered in a way that is detailed later. This was necessary to correct for sample drift during the laser excitation experiment. Shortcomings of this method are that there were no corrections for magnification changes and that no size increases could be captured because of the fixed size of the mask. Fortunately, these two shortcomings somewhat cancel each other as there was vertical shrinking of the frames (Figure 6.A4). This is likely also the reason that we obtained some positive values for the automatic segmentation method in Figure 6.1d, e.g. particle 39. A benefit of this method is that cross-correlation of the mask resulted in improved quantification of state switches. To explain this, we need to look at the analysis pipeline in more detail.

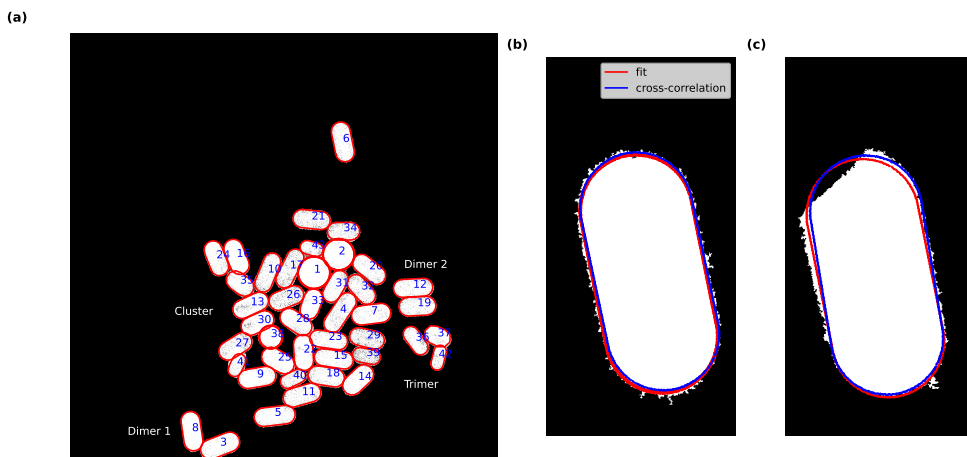


Figure 6.A3: Labeling, fitting, and grouping of NRs (a) within Figure 6.A2c. Zoom-in of NR 6 with fit and mask from automatic analysis before experiment (b) and during experiment (c). When pixels are missing the mask cross-correlation matches and maybe outperforms the fit for quantification.

To follow the sample drift a reference point is required in the frame. The center of the cluster could be used for this. However, to account for rotation it is easiest to track the orientation of an elongated structure. Since the cluster’s aspect ratio is close to 1, we chose to use NR 6 as a reference for rotation and translation. Since the labeling is based on size and NR 6 is highly dynamic throughout the experiment it could not be identified only by labeling. Therefore, we used the cluster center as an initial reference point to locate NR 6 in each frame. From this we computed a relative shift and rotation list for each frame that we used to track all other individual NRs throughout the experiment and to line up each mask per NR per frame. Note that it is likely that the quality of the analysis decreases throughout the experiment as the vertical shrinking (Figure 6.A4) might induce misalignments, especially for particles close to the center of the frame.




Figure 6.A4: Shrink factor along the vertical direction of the frames over the course of the experiment determined from manual segmentation.

Finally, the pixels inside each mask were summed to obtain a total per NR per frame. Subtracting the halves of each mask resulted in the switch magnitude. To obtain the noise boundaries we first computed the moving average for each dynamic range. Then, we chose NR 4 as reference and made sure that  $> 95\%$  of the signal falls in between the noise boundaries, such that the boundaries were placed 1500 pixels away from the moving average. These noise boundaries seem to give accurate switching identification for all NRs, e.g. comparing Figure 6.2 with Figures 6.5 and 6.6.




Figure 6.A5: Original frames of Figure 6.3e,f left and right, respectively.

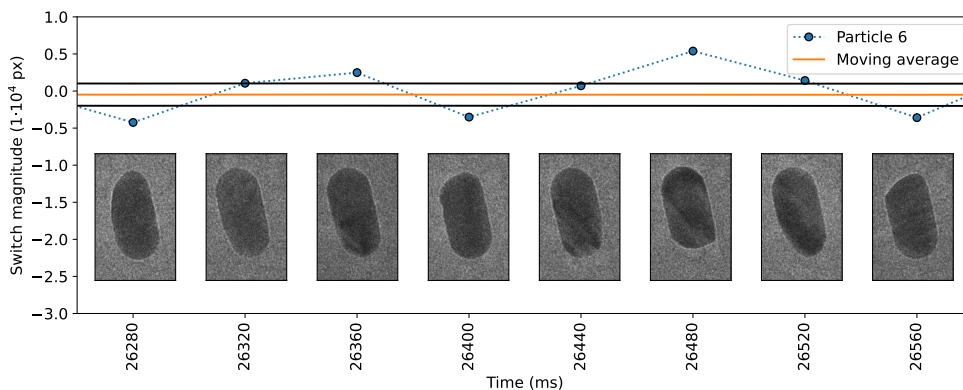


Figure 6.A6: Combination of Figure 6.2b and Figure 6.5c to facilitate comparison. Note that the limits of the y axis have been changed compared to the main text.

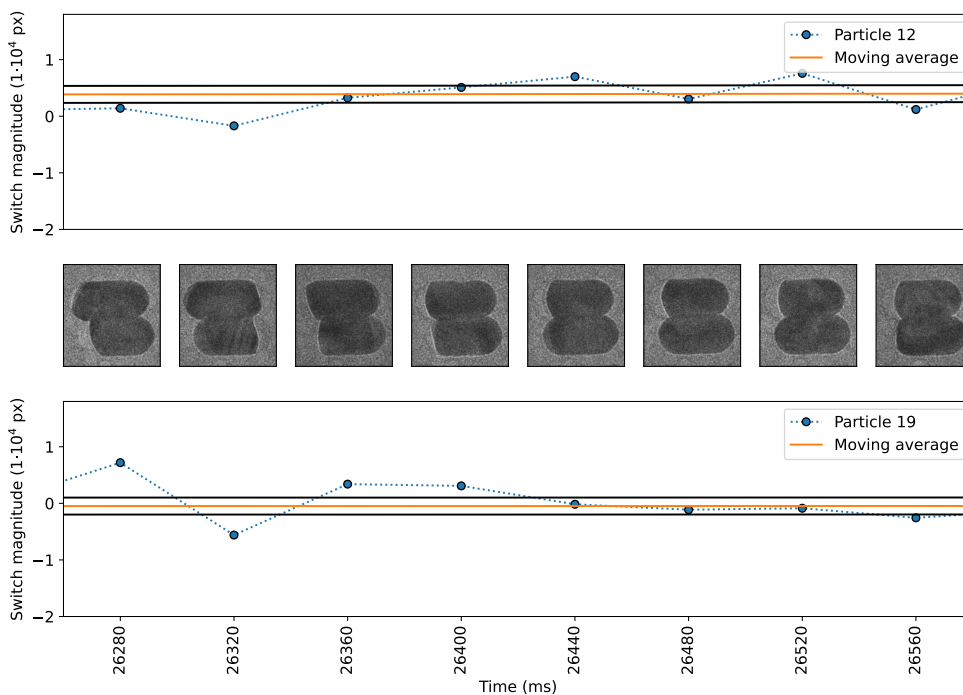


Figure 6.A7: Combination of Figure 6.2c and switch magnitude of NR 12 and 19.

## 6.A.2 Simulations

### Absorbed power by the carbon substrate

First, we simulate the absorption of the incoming laser light by the carbon thin film. For this, we use the electromagnetic wave frequency domain of COMSOL v5.5. The simulation domain is a 20 nm thin film suspended in vacuum. We assume the absorption to be homogeneous, and therefore we only simulate  $8 \times 8 \mu\text{m}$  to reduce the computation time. We set the thickness of the perfectly matched layer to  $2 \mu\text{m}$  and we take the refractive index from Larruquet *et al.* [163]. The absorption of the thin film is computed for an incoming plane wave linearly polarized along  $z$ . The resulting absorbed power averaged over the simulated thin film volume is  $3.17 \times 10^{13} \text{ W/m}^3$  by the thin film at 577 nm for the laser fluence used in the experiment.

### Heat simulations

We simulate a full  $100 \times 100 \mu\text{m}$  carbon TEM grid window on which the 43 gold nanoparticles are laying in vacuum. The dimensions and position of the particles were extracted from Figure 6.A3. The heat diffusion equations are solved with the Heat Transfer in Solids module from COMSOL v5.5. The material properties are reported in Table 6.A1. If at the start of the experiment, the substrate is an amorphous carbon thin film, due to the high temperatures reached during the experiment, we are crystallizing this amorphous thin film and forming some graphite. Graphite is known to have large in plane thermal conductivities on the order of  $1000\text{-}2500 \text{ W/m}^3\text{K}$  while it has a relatively low out of plane conductivity  $5\text{-}10 \text{ W/m}^3\text{K}$ . However, it has also been shown that the in plane thermal conductivity can significantly decrease when the twisting angle between the layers is high. Taking all of these considerations, we use a value of  $10 \text{ W/m}^3\text{K}$  for the thermal conductivity of graphite. We define heat sources for the nanoparticles as the product of their absorption cross section with the laser fluence. To the carbon substrate, we add a Gaussian heat source from the absorbed power computed as described previously. The width of the Gaussian is  $40 \mu\text{m}$  corresponding to the laser beam size. A fixed temperature boundary condition is applied to the edges of the TEM grid which act as a perfect heat sink, representative of the TEM grid copper frame. The domains are meshed with a tetrahedral mesh with a minimum and maximum element size of 1 and 10 nm respectively. Finally, we conduct a static study with a fully coupled GMRES solver to compute the temperature increase of the nanoparticles and TEM membrane during the experiment.

Thermal conductivity gold	315 [W/m*K]
Heat capacity gold	129 [J/kg*K]
Gold density	19320 [kg/m <sup>3</sup> ]
Thermal conductivity carbon	8 [W/m*K]
Heat capacity carbon	710 [J/kg*K]
Density carbon	2330 [kg/m <sup>3</sup> ]
Laser beam radius	20 [ $\mu$ m]
Laser power	50 [mW]

Table 6.A1: Parameters used for the heat the simulations

As mentioned in the main text, we conduct time resolved heat simulations to confirm that the system has enough time to cool down between the pulses. This is simply done by using the solution of the static simulations shown in Figure 6.4 as initial condition. We then let the system go back to equilibrium and we find that the particles cool down to reach room temperature within 100 ns.




Figure 6.A8: Cooling dynamics of the nanoparticles

---

# Bibliography

- [1] World Economics Forum, *The Global Risks Report 2023 (18th ed.)* 2023. [Online]. Available: <https://www.weforum.org/publications/global-risks-report-2023/in-full/>.
- [2] World Economics Forum, *The Global Risks Report 2024 (19th ed.)* 2024. [Online]. Available: <https://www.weforum.org/publications/global-risks-report-2024/in-full/>.
- [3] World Economics Forum, *The Global Risks Report 2025 (20th ed.)* 2025. [Online]. Available: <https://www.weforum.org/publications/global-risks-report-2025/in-full/>.
- [4] I. P. on Climate Change (IPCC), *Climate Change 2021 – The Physical Science Basis*. Cambridge University Press, 2023, ISBN: 9781009157896.
- [5] Deloitte, *Global risk management survey (12 ed.)* 2020. [Online]. Available: <https://iapp.org/resources/article/deloitte-global-risk-management-survey/>.
- [6] Deloitte, *Integrated Annual Report Coöperatief Deloitte U. A.* 2024. [Online]. Available: <https://annualreport.deloitte.nl/downloads>.
- [7] EY Global, *Taskforce on Climate-Related Financial Disclosures Report*, 2024. [Online]. Available: [https://www.ey.com/en\\_nl/insights/sustainability/tcfd-annual-report](https://www.ey.com/en_nl/insights/sustainability/tcfd-annual-report).
- [8] EY Global, *Risk management's strategic opportunity in a time of turbulence*, 2025. [Online]. Available: <https://www.iif.com/Publications/ID/6160/EYIIF-Second-Annual-Global-Insurance-Risk-Management-Survey>.
- [9] KPMG, *Climate Risk Report*, 2024. [Online]. Available: <https://kpmg.com/xx/en/what-we-do/services/audit/corporate-reporting-institute/ifrs/climate-hub.html>.
- [10] KPMG, *Top geopolitical risks 2025*, 2025. [Online]. Available: <https://kpmg.com/xx/en/our-insights/risk-and-regulation/top-risks-forecast-2025.html>.
- [11] PwC, *PwC's Global Risk Survey 2023*, 2023. [Online]. Available: <https://www.pwc.com/gx/en/issues/risk-regulation/global-risk-survey.html>.
- [12] PwC, *PwC's Global Compliance Survey 2025*, 2025.

- [13] M. Draghi, *Report on the Future of European Competitiveness*, 2024. [Online]. Available: [https://commission.europa.eu/topics/eu-competitiveness/draghi-report\\_en](https://commission.europa.eu/topics/eu-competitiveness/draghi-report_en).
- [14] H. H. Thorp, “Here comes the Sun,” *Science*, vol. 390, pp. 1195–1195, 2025.
- [15] K. Calvin et al., *Climate Change 2023: Synthesis Report, Summary for Policy-makers*. 2023.
- [16] IEA, IRENA, UNSD, W. Bank, and WHO, *Tracking SDG7: The Energy Progress Report*, 2024.
- [17] L. Novotny and B. Hecht, *Principles of Nano-Optics*. Cambridge University Press, 2012, ISBN: 9781107005464.
- [18] M. L. Brongersma, N. J. Halas, and P. Nordlander, “Plasmon-induced hot carrier science and technology,” *Nature Nanotechnology*, vol. 10, pp. 25–34, 2015.
- [19] E. Cortés et al., “Challenges in Plasmonic Catalysis,” *ACS Nano*, vol. 14, pp. 16 202–16 219, 2020.
- [20] J. Li, S. K. Cushing, F. Meng, T. R. Senty, A. D. Bristow, and N. Wu, “Plasmon-induced resonance energy transfer for solar energy conversion,” *Nature Photonics*, vol. 9, pp. 601–607, 2015.
- [21] B. Ostovar et al., “The role of the plasmon in interfacial charge transfer,” *Science Advances*, vol. 10, 2024.
- [22] M. Green, *Third Generation Photovoltaics*. Springer Berlin Heidelberg, 2006, vol. 12, ISBN: 978-3-540-26562-7.
- [23] A. LaPotin et al., “Thermophotovoltaic efficiency of 40%,” *Nature*, vol. 604, pp. 287–291, 2022.
- [24] S. B. Ramakrishnan et al., “Photoinduced Electron and Energy Transfer Pathways and Photocatalytic Mechanisms in Hybrid Plasmonic Photocatalysis,” *Advanced Optical Materials*, vol. 9, p. 2 101 128, 2021.
- [25] K. Sytwu, M. Vadai, and J. A. Dionne, “Bimetallic nanostructures: combining plasmonic and catalytic metals for photocatalysis,” *Advances in Physics: X*, vol. 4, p. 1 619 480, 2019.
- [26] K. V. Gordon et al., “Tuning the Growth of Chiral Gold Nanoparticles Through Rational Design of a Chiral Molecular Inducer,” *Nano Letters*, vol. 23, pp. 9880–9886, 2023.
- [27] G. González-Rubio et al., “Micelle-directed chiral seeded growth on anisotropic gold nanocrystals,” *Science*, vol. 368, pp. 1472–1477, 2020.

- [28] J. Rodríguez-Fernández et al., “Spectroscopy, Imaging, and Modeling of Individual Gold Decahedra,” *The Journal of Physical Chemistry C*, vol. 113, pp. 18 623–18 631, 2009.
- [29] A. Al-Zubeidi, L. A. McCarthy, A. Rafiei-Miandashti, T. S. Heiderscheid, and S. Link, “Single-particle scattering spectroscopy: fundamentals and applications,” *Nanophotonics*, vol. 10, pp. 1621–1655, 2021.
- [30] M. W. Knight, J. Fan, F. Capasso, and N. J. Halas, “Influence of excitation and collection geometry on the dark field spectra of individual plasmonic nanostructures,” *Optics Express*, vol. 18, p. 2579, 2010.
- [31] R. E. Armstrong, J. C. van Liempt, and P. Zijlstra, “Effect of Film Thickness on the Far- and Near-Field Optical Response of Nanoparticle-on-Film Systems,” *The Journal of Physical Chemistry C*, vol. 123, pp. 25 801–25 808, 2019.
- [32] M. Dieperink, F. Scalerandi, and W. Albrecht, “Correlating structure, morphology and properties of metal nanostructures by combining single-particle optical spectroscopy and electron microscopy,” *Nanoscale*, vol. 14, pp. 7460–7472, 2022.
- [33] E. M. Perassi et al., “Quantitative Understanding of the Optical Properties of a Single, Complex-Shaped Gold Nanoparticle from Experiment and Theory,” *ACS Nano*, vol. 8, pp. 4395–4402, 2014.
- [34] V. Myroshnychenko, E. Carbó-Argibay, I. Pastoriza-Santos, J. Pérez-Juste, L. M. Liz-Marzán, and F. J. G. de Abajo, “Modeling the Optical Response of Highly Faceted Metal Nanoparticles with a Fully 3D Boundary Element Method,” *Advanced Materials*, vol. 20, pp. 4288–4293, 2008.
- [35] P. C. Angelomé et al., “Seedless Synthesis of Single Crystalline Au Nanoparticles with Unusual Shapes and Tunable LSPR in the near-IR,” *Chemistry of Materials*, vol. 24, pp. 1393–1399, 2012.
- [36] P. Midgley and M. Weyland, “3D electron microscopy in the physical sciences: the development of Z-contrast and EFTEM tomography,” *Ultramicroscopy*, vol. 96, pp. 413–431, 2003.
- [37] W. Albrecht and S. Bals, “Fast Electron Tomography for Nanomaterials,” *The Journal of Physical Chemistry C*, vol. 124, pp. 27 276–27 286, 2020.
- [38] P. Spaeth et al., “Photothermal Circular Dichroism Measurements of Single Chiral Gold Nanoparticles Correlated with Electron Tomography,” *ACS Photonics*, vol. 9, pp. 3995–4004, 2022.
- [39] K. Shiratori et al., “Machine-Learned Decision Trees for Predicting Gold Nanorod Sizes from Spectra,” *The Journal of Physical Chemistry C*, vol. 125, pp. 19 353–19 361, 2021.

- [40] S. Alekseeva, I. I. Nedrygailov, and C. Langhammer, "Single Particle Plasmonics for Materials Science and Single Particle Catalysis," *ACS Photonics*, vol. 6, pp. 1319–1330, 2019.
- [41] A. Zilli, W. Langbein, and P. Borri, "Quantitative Measurement of the Optical Cross Sections of Single Nano-objects by Correlative Transmission and Scattering Microspectroscopy," *ACS Photonics*, vol. 6, pp. 2149–2160, 2019.
- [42] L. M. Payne, W. Albrecht, W. Langbein, and P. Borri, "The optical nanosizer – quantitative size and shape analysis of individual nanoparticles by high-throughput widefield extinction microscopy," *Nanoscale*, vol. 12, pp. 16 215–16 228, 2020.
- [43] L. M. Payne, F. Masia, A. Zilli, W. Albrecht, P. Borri, and W. Langbein, "Quantitative morphometric analysis of single gold nanoparticles by optical extinction microscopy: Material permittivity and surface damping effects," *The Journal of Chemical Physics*, vol. 154, p. 044 702, 2021.
- [44] Y. Wang et al., "Quantitatively linking morphology and optical response of individual silver nanohedra," *Nanoscale*, vol. 14, pp. 11 028–11 037, 2022.
- [45] K. Jenkinson, L. M. Liz-Marzán, and S. Bals, "Multimode Electron Tomography Sheds Light on Synthesis, Structure, and Properties of Complex Metal-Based Nanoparticles," *Advanced Materials*, vol. 34, 2022.
- [46] W. Albrecht et al., "3D Atomic-Scale Dynamics of Laser-Light-Induced Restructuring of Nanoparticles Unraveled by Electron Tomography," *Advanced Materials*, vol. 33, p. 2 100 972, 2021.
- [47] A. Skorikov et al., "Quantitative 3D Characterization of Elemental Diffusion Dynamics in Individual Ag@Au Nanoparticles with Different Shapes," *ACS Nano*, vol. 13, pp. 13 421–13 429, 2019.
- [48] T. M. de Oliveira et al., "3D Characterization and Plasmon Mapping of Gold Nanorods Welded by Femtosecond Laser Irradiation," *ACS Nano*, vol. 14, pp. 12 558–12 570, 2020.
- [49] W. Albrecht, E. Bladt, H. Vanrompay, J. D. Smith, S. E. Skrabalak, and S. Bals, "Thermal Stability of Gold/Palladium Octopods Studied *in situ* in 3D: Understanding Design Rules for Thermally Stable Metal Nanoparticles," *ACS Nano*, vol. 13, pp. 6522–6530, 2019.
- [50] H. Vanrompay et al., "3D characterization of heat-induced morphological changes of Au nanostars by fast *in situ* electron tomography," *Nanoscale*, vol. 10, pp. 22 792–22 801, 2018.

- [51] Y. Wang, A. B. Serrano, K. Sentosun, S. Bals, and L. M. Liz-Marzán, “Stabilization and Encapsulation of Gold Nanostars Mediated by Dithiols,” *Small*, vol. 11, pp. 4314–4320, 2015.
- [52] W. Albrecht, S. V. Aert, and S. Bals, “Three-Dimensional Nanoparticle Transformations Captured by an Electron Microscope,” *Accounts of Chemical Research*, vol. 54, pp. 1189–1199, 2021.
- [53] J. S. Googasian and S. E. Skrabalak, “Practical Considerations for Simulating the Plasmonic Properties of Metal Nanoparticles,” *ACS Physical Chemistry Au*, vol. 3, pp. 252–262, 2023.
- [54] A. Trügler, *Optical Properties of Metallic Nanoparticles*. Springer International Publishing, 2016, vol. 232, ISBN: 978-3-319-25072-4.
- [55] A. Amirjani and S. K. Sadrnezhaad, “Computational electromagnetics in plasmonic nanostructures,” *Journal of Materials Chemistry C*, vol. 9, pp. 9791–9819, 2021.
- [56] J. E. Ortiz-Castillo, R. C. Gallo-Villanueva, M. J. Madou, and V. H. Perez-Gonzalez, “Anisotropic gold nanoparticles: A survey of recent synthetic methodologies,” *Coordination Chemistry Reviews*, vol. 425, p. 213 489, 2020.
- [57] S. E. Lohse, N. D. Burrows, L. Scarabelli, L. M. Liz-Marzán, and C. J. Murphy, “Anisotropic Noble Metal Nanocrystal Growth: The Role of Halides,” *Chemistry of Materials*, vol. 26, pp. 34–43, 2014.
- [58] G. Lin, W. Lu, W. Cui, and L. Jiang, “A Simple Synthesis Method for Gold Nano- and Microplate Fabrication Using a Tree-Type Multiple-Amine Head Surfactant,” *Crystal Growth Design*, vol. 10, pp. 1118–1123, 2010.
- [59] F. Kim, S. Connor, H. Song, T. Kuykendall, and P. Yang, “Platonic Gold Nanocrystals,” *Angewandte Chemie International Edition*, vol. 43, pp. 3673–3677, 2004.
- [60] S. Barbosa et al., “Tuning Size and Sensing Properties in Colloidal Gold Nanostars,” *Langmuir*, vol. 26, pp. 14 943–14 950, 2010.
- [61] A. Guerrero-Martínez, S. Barbosa, I. Pastoriza-Santos, and L. M. Liz-Marzán, “Nanostars shine bright for you,” *Current Opinion in Colloid Interface Science*, vol. 16, pp. 118–127, 2011.
- [62] S. Adhikari, P. Spaeth, A. Kar, M. D. Baaske, S. Khatua, and M. Orrit, “Photothermal Microscopy: Imaging the Optical Absorption of Single Nanoparticles and Single Molecules,” *ACS Nano*, vol. 14, pp. 16 414–16 445, 2020.

- [63] B. Ni et al., “Chiral Seeded Growth of Gold Nanorods Into Fourfold Twisted Nanoparticles with Plasmonic Optical Activity,” *Advanced Materials*, vol. 35, 2023.
- [64] L. Wang, M. H. Kafshgari, and M. Meunier, “Optical Properties and Applications of Plasmonic-Metal Nanoparticles,” *Advanced Functional Materials*, vol. 30, p. 2005400, 2020.
- [65] P. Zijlstra, J. W. M. Chon, and M. Gu, “Five-dimensional optical recording mediated by surface plasmons in gold nanorods,” *Nature*, vol. 459, pp. 410–413, 2009.
- [66] W. Li et al., “Bidirectional plasmonic coloration with gold nanoparticles by wavelength-switched photoredox reaction,” *Nanoscale*, vol. 10, pp. 21910–21917, 2018.
- [67] J. N. Anker, W. P. Hall, O. Lyandres, N. C. Shah, J. Zhao, and R. P. V. Duyne, “Biosensing with plasmonic nanosensors,” *Nature Materials*, vol. 7, pp. 442–453, 2008.
- [68] B. S. Hoener et al., “Plasmonic Sensing and Control of Single-Nanoparticle Electrochemistry,” *Chem*, vol. 4, pp. 1560–1585, 2018.
- [69] J. B. Sambur and P. Chen, “Approaches to Single-Nanoparticle Catalysis,” *Annual Review of Physical Chemistry*, vol. 65, pp. 395–422, 2014.
- [70] S. Yu, A. J. Wilson, G. Kumari, X. Zhang, and P. K. Jain, “Opportunities and Challenges of Solar-Energy-Driven Carbon Dioxide to Fuel Conversion with Plasmonic Catalysts,” *ACS Energy Letters*, vol. 2, pp. 2058–2070, 2017.
- [71] C. C. Carlin et al., “Nanoscale and ultrafast *in situ* techniques to probe plasmon photocatalysis,” *Chemical Physics Reviews*, vol. 4, 2023.
- [72] A. Baldi and S. H. C. Askes, “Pulsed Photothermal Heterogeneous Catalysis,” *ACS Catalysis*, vol. 13, pp. 3419–3432, 2023.
- [73] K. L. Kelly, E. Coronado, L. L. Zhao, and G. C. Schatz, “The Optical Properties of Metal Nanoparticles: The Influence of Size, Shape, and Dielectric Environment,” *The Journal of Physical Chemistry B*, vol. 107, pp. 668–677, 2003.
- [74] L. M. Liz-Marzán, C. R. Kagan, and J. E. Millstone, “Reproducibility in Nanocrystal Synthesis? Watch Out for Impurities!” *ACS Nano*, vol. 14, pp. 6359–6361, 2020.

- [75] A. M. Brown, R. Sundararaman, P. Narang, W. A. Goddard, and H. A. Atwater, “Nonradiative Plasmon Decay and Hot Carrier Dynamics: Effects of Phonons, Surfaces, and Geometry,” *ACS Nano*, vol. 10, pp. 957–966, 2016.
- [76] J. Zheng et al., “Gold Nanorods: The Most Versatile Plasmonic Nanoparticles,” *Chemical Reviews*, vol. 121, pp. 13 342–13 453, 2021.
- [77] K. A. Dahan, Y. Li, J. Xu, and C. Kan, “Recent progress of gold nanostructures and their applications,” *Physical Chemistry Chemical Physics*, vol. 25, pp. 18 545–18 576, 2023.
- [78] P. Drude, “Zur Elektronentheorie der Metalle,” *Annalen der Physik*, vol. 306, pp. 566–613, 1900.
- [79] A. Sommerfeld, “Zur Elektronentheorie der Metalle auf Grund der Fermischen Statistik,” *Zeitschrift fr Physik*, vol. 47, pp. 1–32, 1928.
- [80] K. M. McPeak et al., “Plasmonic Films Can Easily Be Better: Rules and Recipes,” *ACS Photonics*, vol. 2, pp. 326–333, 2015.
- [81] P. B. Johnson and R. W. Christy, “Optical Constants of the Noble Metals,” *Physical Review B*, vol. 6, pp. 4370–4379, 1972.
- [82] E. D. Palik, *Handbook of Optical Constants of Solids*, E. D. Palik, Ed. Academic Press, 1998, vol. 3.
- [83] R. L. Olmon et al., “Optical dielectric function of gold,” *Physical Review B*, vol. 86, p. 235 147, 2012.
- [84] C. Novo et al., “Contributions from radiation damping and surface scattering to the linewidth of the longitudinal plasmon band of gold nanorods: a single particle study,” *Physical Chemistry Chemical Physics*, vol. 8, p. 3540, 2006.
- [85] B. Foerster, A. Joplin, K. Kaefer, S. Celiksoy, S. Link, and C. Sönnichsen, “Chemical Interface Damping Depends on Electrons Reaching the Surface,” *ACS Nano*, vol. 11, pp. 2886–2893, 2017.
- [86] S. A. Lee and S. Link, “Chemical Interface Damping of Surface Plasmon Resonances,” *Accounts of Chemical Research*, vol. 54, pp. 1950–1960, 2021.
- [87] J. Strutt, “On the light from the sky, its polarization and colour,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 41, pp. 107–120, 1871.
- [88] G. Mie, “Beiträge zur Optik trüber Medien, speziell kolloidaler Metallösungen,” *Annalen der Physik*, vol. 330, pp. 377–445, 1908.
- [89] M. R. Gonçalves, “Plasmonic nanoparticles: fabrication, simulation and experiments,” *Journal of Physics D: Applied Physics*, vol. 47, p. 213 001, 2014.

- [90] U. Hohenester and A. Trügler, “MNPBEM - A Matlab toolbox for the simulation of plasmonic nanoparticles,” 2011.
- [91] H. Sun, D. L. Darmofal, and R. Haimes, “On the impact of triangle shapes for boundary layer problems using high-order finite element discretization,” *Journal of Computational Physics*, vol. 231, pp. 541–557, 2012.
- [92] P. M. Knupp, “Achieving finite element mesh quality via optimization of the Jacobian matrix norm and associated quantities. Part I-a framework for surface mesh optimization,” *International Journal for Numerical Methods in Engineering*, vol. 48, pp. 401–420, 2000.
- [93] P. Stoller, V. Jacobsen, and V. Sandoghdar, “Measurement of the complex dielectric constant of a single gold nanoparticle,” *Optics Letters*, vol. 31, p. 2474, 2006.
- [94] L. V. R.-d. Marcos, J. I. Larruquert, J. A. Méndez, and J. A. Aznárez, “Self-consistent optical constants of SiO<sub>2</sub> and Ta<sub>2</sub>O<sub>5</sub> films,” *Optical Materials Express*, vol. 6, p. 3622, 2016.
- [95] K. Luke, Y. Okawachi, M. R. E. Lamont, A. L. Gaeta, and M. Lipson, “Broadband mid-infrared frequency comb generation in a Si<sub>3</sub>N<sub>4</sub> microresonator,” *Optics Letters*, vol. 40, p. 4823, 2015.
- [96] J. I. Larruquert, L. V. R.-d. Marcos, J. A. Méndez, P. J. Martin, and A. Bendavid, “High reflectance ta-C coatings in the extreme ultraviolet,” *Optics Express*, vol. 21, p. 27 537, 2013.
- [97] B. T. Draine and P. J. Flatau, “Discrete-Dipole Approximation For Scattering Calculations,” *Journal of the Optical Society of America A*, vol. 11, p. 1491, 1994.
- [98] L. Wang, M. H. Kafshgari, and M. Meunier, “Optical Properties and Applications of Plasmonic-Metal Nanoparticles,” *Advanced Functional Materials*, vol. 30, p. 2005 400, 2020.
- [99] L. Scarabelli, A. Sánchez-Iglesias, J. Pérez-Juste, and L. M. Liz-Marzán, “A “Tips and Tricks” Practical Guide to the Synthesis of Gold Nanorods,” *The Journal of Physical Chemistry Letters*, vol. 6, pp. 4270–4279, 2015.
- [100] J. Rodríguez-Fernández et al., “Spectroscopy, Imaging, and Modeling of Individual Gold Decahedra,” *The Journal of Physical Chemistry C*, vol. 113, pp. 18 623–18 631, 2009.
- [101] E. Ringe, B. Sharma, A.-I. Henry, L. D. Marks, and R. P. V. Duyne, “Single nanoparticle plasmonics,” *Physical Chemistry Chemical Physics*, vol. 15, p. 4110, 2013.

- [102] R. F. Hamans, R. Kamarudheen, and A. Baldi, “Single Particle Approaches to Plasmon-Driven Catalysis,” *Nanomaterials*, vol. 10, p. 2377, 2020.
- [103] P. Zijlstra and M. Orrit, “Single metal nanoparticles: optical detection, spectroscopy and applications,” *Reports on Progress in Physics*, vol. 74, p. 106 401, 2011.
- [104] J. Olson, S. Dominguez-Medina, A. Hoggard, L.-Y. Wang, W.-S. Chang, and S. Link, “Optical characterization of single plasmonic nanoparticles,” *Chemical Society Reviews*, vol. 44, pp. 40–57, 2015.
- [105] A.-I. Henry, J. M. Bingham, E. Ringe, L. D. Marks, G. C. Schatz, and R. P. V. Duyne, “Correlated Structure and Optical Property Studies of Plasmonic Nanoparticles,” *The Journal of Physical Chemistry C*, vol. 115, pp. 9291–9305, 2011.
- [106] T. Jollans, M. D. Baaske, and M. Orrit, “Nonfluorescent Optical Probing of Single Molecules and Nanoparticles,” *The Journal of Physical Chemistry C*, vol. 123, pp. 14 107–14 117, 2019.
- [107] Y. Wang, A. Zilli, Z. Sztranyovszky, W. Langbein, and P. Borri, “Quantitative optical microspectroscopy, electron microscopy, and modelling of individual silver nanocubes reveal surface compositional changes at the nanoscale,” *Nanoscale Advances*, vol. 2, pp. 2485–2496, 2020.
- [108] T. M. de Oliveira et al., “3D Characterization and Plasmon Mapping of Gold Nanorods Welded by Femtosecond Laser Irradiation,” *ACS Nano*, vol. 14, pp. 12 558–12 570, 2020.
- [109] B. Goris, T. Roelandts, K. Batenburg, H. H. Mezerji, and S. Bals, “Advanced reconstruction algorithms for electron tomography: From comparison to combination,” *Ultramicroscopy*, vol. 127, pp. 40–47, 2013.
- [110] C. Glasbey, “An Analysis of Histogram-Based Thresholding Algorithms,” *CVGIP: Graphical Models and Image Processing*, vol. 55, pp. 532–537, 1993.
- [111] T. MacDonald, *Isosurface Extraction*, 2011.
- [112] E. J. Kirkland, *Advanced Computing in Electron Microscopy*. Springer International Publishing, 2020, ISBN: 978-3-030-33259-4.
- [113] H. Vanrompay et al., “Fast versus conventional HAADF-STEM tomography of nanoparticles: advantages and challenges,” *Ultramicroscopy*, vol. 221, p. 113 191, 2021.
- [114] N. Otsu, “A Threshold Selection Method from Gray-Level Histograms,” *IEEE Transactions on Systems, Man, and Cybernetics*, vol. 9, pp. 62–66, 1979.

- [115] W. van Aarle et al., “The ASTRA Toolbox: A platform for advanced algorithm development in electron tomography,” *Ultramicroscopy*, vol. 157, pp. 35–47, 2015.
- [116] H. Wiecek, “The image quality of FBP and MLEM reconstruction,” *Physics in Medicine and Biology*, vol. 55, pp. 3161–3176, 2010.
- [117] W. E. Lorensen and H. E. Cline, “Marching cubes,” in ACM, 1998, pp. 347–353.
- [118] T. Lewiner, H. Lopes, A. W. Vieira, and G. Tavares, “Efficient Implementation of Marching Cubes’ Cases with Topological Guarantees,” *Journal of Graphics Tools*, vol. 8, pp. 1–15, 2003.
- [119] M. Garland and P. S. Heckbert, “Surface simplification using quadric error metrics,” ACM Press, 1997, pp. 209–216, ISBN: 0897918967.
- [120] J. Koo, A. B. Dahl, J. A. Bærentzen, Q. Chen, S. Bals, and V. A. Dahl, “Shape from projections via differentiable forward projector for computed tomography,” *Ultramicroscopy*, vol. 224, p. 113 239, 2021.
- [121] A. Skorikov, W. Heyvaert, W. Albecht, D. M. Pelt, and S. Bals, “Deep learning-based denoising for improved dose efficiency in EDX tomography of nanoparticles,” *Nanoscale*, vol. 13, pp. 12 242–12 249, 2021.
- [122] Z. Chen and H. Zhang, “Neural marching cubes,” *ACM Transactions on Graphics*, vol. 40, pp. 1–15, 2021.
- [123] W. J. Schroeder, J. A. Zarge, and W. E. Lorensen, “Decimation of triangle meshes,” *Proceedings of the 19th annual conference on Computer graphics and interactive techniques*, pp. 65–70, 1992.
- [124] A. Weigel, A. Sebesta, and P. Kukura, “Dark Field Microspectroscopy with Single Molecule Fluorescence Sensitivity,” *ACS Photonics*, vol. 1, pp. 848–856, 2014.
- [125] A. Ciesielski, L. Skowronski, M. Trzcinski, and T. Szoplik, “Controlling the optical parameters of self-assembled silver films with wetting layers and annealing,” *Applied Surface Science*, vol. 421, pp. 349–356, 2017.
- [126] F. Cheng, P.-H. Su, J. Choi, S. Gwo, X. Li, and C.-K. Shih, “Epitaxial Growth of Atomically Smooth Aluminum on Silicon and Its Intrinsic Optical Properties,” *ACS Nano*, vol. 10, pp. 9852–9860, 2016.
- [127] Z. Bian et al., *Autofocusing technologies for whole slide imaging and automated microscopy*, 2020. arXiv: 2006.08610 [physics.med-ph].

- [128] J. Knapper, J. T. Collins, J. Stirling, S. McDermott, W. Wadsworth, and R. W. Bowman, “Fast, high-precision autofocus on a motorised microscope: Automating blood sample imaging on the OpenFlexure Microscope,” *Journal of Microscopy*, vol. 285, pp. 29–39, 2022.
- [129] L. Dubau et al., “Defects do catalysis: CO monolayer oxidation and oxygen reduction reaction on hollow PtNi/C nanoparticles,” *ACS Catalysis*, vol. 6, no. 7, pp. 4673–4684, 2016.
- [130] Y. Jia, K. Jiang, H. Wang, and X. Yao, “The role of defect sites in nanomaterials for electrocatalytic energy conversion,” *Chem*, vol. 5, no. 6, pp. 1371–1397, 2019.
- [131] A. Lak, S. Disch, and P. Bender, “Embracing defects and disorder in magnetic nanoparticles,” *Advanced Science*, vol. 8, no. 7, p. 2002682, 2021.
- [132] M. Oliveira et al., “3D Characterization and plasmon mapping of gold nanorods welded by femtosecond laser irradiation,” *ACS Nano*, vol. 14, no. 10, pp. 12558–12570, 2020.
- [133] Y. U. Staechelin, D. Hoeing, F. Schulz, and H. Lange, “Size-dependent electron–phonon coupling in monocrystalline gold nanoparticles,” *ACS Photonics*, vol. 8, no. 3, pp. 752–757, 2021.
- [134] W.-Y. Chiang et al., “Electron–phonon relaxation dynamics of hot electrons in gold nanoparticles are independent of excitation pathway,” *The Journal of Physical Chemistry C*, vol. 127, no. 43, pp. 21176–21185, 2023.
- [135] G. González-Rubio and W. Albrecht, “Engineering of plasmonic gold nanocrystals through pulsed laser irradiation,” *Applied Physics Letters*, vol. 121, no. 20, 2022.
- [136] W. Albrecht et al., “3D Atomic-Scale Dynamics of Laser-Light-Induced Restructuring of Nanoparticles Unraveled by Electron Tomography,” *Advanced Materials*, vol. 33, no. 33, p. 2100972, 2021.
- [137] G. González-Rubio et al., “Controlled Alloying of Au@Ag Core–Shell Nanorods Induced by Femtosecond Laser Irradiation,” *Advanced Optical Materials*, vol. 9, no. 10, p. 2002134, 2021.
- [138] G. González-Rubio et al., “Femtosecond laser reshaping yields gold nanorods with ultranarrow surface plasmon resonances,” *Science*, vol. 358, no. 6363, pp. 640–644, 2017.
- [139] V. Manzaneda-González et al., “From multi- to single-hollow trimetallic nanocrystals by ultrafast heating,” *Chemistry of Materials*, vol. 35, no. 22, pp. 9603–9612, 2023.

- [140] T. Milagres de Oliveira et al., “3D characterization and plasmon mapping of gold nanorods welded by femtosecond laser irradiation,” *ACS nano*, vol. 14, no. 10, pp. 12 558–12 570, 2020.
- [141] G. González-Rubio et al., “Formation of hollow gold nanocrystals by nanosecond laser irradiation,” *The Journal of Physical Chemistry Letters*, vol. 11, no. 3, pp. 670–677, 2020.
- [142] G. González-Rubio, A. Guerrero-Martínez, and L. M. Liz-Marzán, “Reshaping, fragmentation, and assembly of gold nanoparticles assisted by pulse lasers,” *Accounts of Chemical Research*, vol. 49, no. 4, pp. 678–686, 2016.
- [143] G. González-Rubio et al., “Femtosecond laser-controlled tip-to-tip assembly and welding of gold nanorods,” *Nano letters*, vol. 15, no. 12, pp. 8282–8288, 2015.
- [144] J. C. Castro-Palacio et al., “Hollow gold nanoparticles produced by femtosecond laser irradiation,” *The Journal of Physical Chemistry Letters*, vol. 11, no. 13, pp. 5108–5114, 2020.
- [145] P. Díaz-Núñez et al., “Using femtosecond laser irradiation to grow the belly of gold nanorods,” *The Journal of Physical Chemistry C*, vol. 122, no. 34, pp. 19 816–19 822, 2018.
- [146] S. Link, Z. L. Wang, and M. A. El-Sayed, “How does a gold nanorod melt?” *The Journal of Physical Chemistry B*, vol. 104, no. 33, pp. 7867–7870, 2000.
- [147] J. Yan, D. Zhu, J. Xie, Y. Shao, and W. Xiao, “Light tailoring of internal atomic structure of gold nanorods,” *Small*, vol. 16, no. 22, p. 2 001 101, 2020.
- [148] N. Sumimoto, K. Nakao, T. Yamamoto, K. Yasuda, S. Matsumura, and Y. Niidome, “In situ observation of structural transformation of gold nanorods under pulsed laser irradiation in an HVEM,” *Journal of Electron Microscopy*, vol. 63, no. 4, pp. 261–268, 2014.
- [149] G. Opletal, G. Grochola, Y. H. Chui, I. Snook, and S. Russo, “Stability and transformations of heated gold nanorods,” *The Journal of Physical Chemistry C*, vol. 115, no. 11, pp. 4375–4380, 2011.
- [150] Y. Wang, S. Teitel, and C. Dellago, “Surface-driven bulk reorganization of gold nanorods,” *Nano letters*, vol. 5, no. 11, pp. 2174–2178, 2005.
- [151] G. Baffou, R. Quidant, and F. J. García de Abajo, “Nanoscale control of optical heating in complex plasmonic systems,” *ACS nano*, vol. 4, no. 2, pp. 709–716, 2010.

- [152] A. Baldi and S. H. Askes, "Pulsed photothermal heterogeneous catalysis," *ACS catalysis*, vol. 13, no. 5, pp. 3419–3432, 2023.
- [153] P. Li et al., "Nanoscale Thermometry of Plasmonic Structures via Raman Shifts in Copper Phthalocyanine," *The Journal of Physical Chemistry C*, vol. 127, no. 20, pp. 9690–9698, 2023.
- [154] K. Imura, N. Hayashi, and K. Kawashima, "Ultrafine Modulation of Plasmon Resonance in Single Gold Nanorods via Photothermal Effects," *The Journal of Physical Chemistry Letters*, vol. 16, pp. 11 037–11 042, 2025.
- [155] H. Petrova, J. P. Juste, I. Pastoriza-Santos, G. V. Hartland, L. M. Liz-Marzán, and P. Mulvaney, "On the temperature stability of gold nanorods: comparison between thermal and ultrafast laser-induced heating," *Physical Chemistry Chemical Physics*, vol. 8, no. 7, pp. 814–821, 2006.
- [156] H. Cho, J. W. Shin, and R. Ryoo, "Atomic scale mechanisms underlying thermal reshaping of anisotropic gold nanocrystals revealed by in situ electron microscopy," *The Journal of Physical Chemistry C*, vol. 124, no. 23, pp. 12 855–12 863, 2020.
- [157] P. E. Batson, "Motion of gold atoms on carbon in the aberration-corrected STEM," *Microscopy and Microanalysis*, vol. 14, no. 1, pp. 89–97, 2008.
- [158] V. Gervilla, M. Zarshenas, D. G. Sangiovanni, and K. Sarakinos, "Anomalous versus normal room-temperature diffusion of metal adatoms on graphene," *The Journal of Physical Chemistry Letters*, vol. 11, no. 21, pp. 8930–8936, 2020.
- [159] D. Appy et al., "Transition metals on the (0 0 0 1) surface of graphite: Fundamental aspects of adsorption, diffusion, and morphology," *Progress in surface science*, vol. 89, no. 3-4, pp. 219–238, 2014.
- [160] B. Wu et al., "A Simple Method to Measure the Contact Angle of Metal Droplets on Graphite," *Nanomanufacturing and Metrology*, vol. 6, no. 1, p. 31, 2023.
- [161] A. Habenicht, M. Olapinski, F. Burmeister, P. Leiderer, and J. Boneberg, "Jumping nanodroplets," *Science*, vol. 309, no. 5743, pp. 2043–2045, 2005.
- [162] S. Kumar, "Wettability and structural evolution of gold over a single-walled carbon nanotube: an atomistic investigation," *The Journal of Physical Chemistry C*, vol. 122, no. 28, pp. 16 346–16 355, 2018.
- [163] J. I. Larruquert, L. Rodríguez-de Marcos, J. A. Méndez, P. Martín, and A. Bendavid, "High reflectance ta-C coatings in the extreme ultraviolet," *Optics express*, vol. 21, no. 23, pp. 27 537–27 549, 2013.



---

# List of publications and author contributions

This thesis is based on the following publications.

1. M. Dieperink\*, F. Scalerandi\*, and W. Albrecht, "Correlating structure, morphology and properties of metal nanostructures by combining single-particle optical spectroscopy and electron microscopy," *Nanoscale*, vol. 14, pp. 7460–7472, 20 2022, doi: 10.1039/D1NR08130F. **Chapter 1**

WA initiated the project. MD and FS performed an extensive literature study under the supervision of WA. All authors co-wrote the manuscript.

2. M. Dieperink, A. Skorikov, N. Claes, S. Bals, and W. Albrecht, "Considerations for electromagnetic simulations for a quantitative correlation of optical spectroscopy and electron tomography of plasmonic nanoparticles," *Nanophotonics*, vol. 13, pp. 4647–4665, 25 Dec. 2024, doi: 10.1515/nanoph-2024-0238. **Chapter 2 and 3**

WA initiated the project. MD performed the research, analyzed the data and wrote the original draft of the manuscript. WA and AS edited the draft. NC performed the electron tomography experiments under the supervision of SB. AS performed the TVM reconstructions and helped with the data interpretation. All authors have accepted responsibility for the entire content of this manuscript and approved its submission.

3. M. Dieperink, P. Spaeth, W. Albrecht, "Decoupling Geometry and Environment in Plasmonic Nanorods via Correlative Optical–Electron Microscopy". (*in preparation*) **Chapter 4 and 5**

WA initiated the project. PS designed the concept of the setup. MD built the setup and refined the design. MD performed the research, analyzed the data and wrote the original draft of the manuscript. WA edited the draft.

4. M. Dieperink, L. Monin, I. Ohnishi, T. Sasaki, Y. Jimbo, E. Garnett, W. Albrecht, "In-Situ Laser Reshaping Dynamics of Gold Nanoparticles in the TEM". (*submitted*) **Chapter 6**

WA initiated the project. MD prepared the sample. IO, TS, YJ, EG, and WA performed the experiment. MD performed the data analysis. LM performed the heat simulations. MD and LM wrote the original draft of the manuscript. WA and EG edited the draft.

\*equal contributions

## Other publications

1. Y. Bleiji, M. Dieperink, I. Schuringa, H. Sun, and E. Alarcon-Llado, "Influence of the crystallographic texture of ITO on the electrodeposition of silver nanoparticles," *RSC Advances*, vol. 13, pp. 6490–6497, 10 2023, doi: 10.1039/D3RA00577A.
2. T. Bouwens et al., "Using supramolecular machinery to engineer directional charge propagation in photoelectrochemical devices," *Nature Chemistry*, Oct. 2022, doi: 10.1038/s41557-022-01068-y.

The author acknowledges the use of text-based artificial intelligence tools, which were employed to provide sentence-level suggestions related to grammar and vocabulary.

---

# Rodcast – Summary

*Rodcast – The comprehensive broadcast of everything one could possibly learn from a single nanorod, including how to cast light on it, screen-cast it, cast shadow on it, typecast it, and eventually melt-cast it.*

This thesis develops a complete, quantitative framework for understanding, simulating, and experimentally probing the optical response of individual plasmonic nanoparticles, with a particular focus on gold nanorods. It links the electron-scale physics of plasmonic excitations to the real, experimentally observed behaviour of nanoparticles, both in static configurations and under strong non-equilibrium perturbations. The work integrates multiple methodological layers: global context and physical motivation, accurate electrodynamic simulation, reliable three-dimensional morphology extraction, high-SNR quasi-darkfield single-particle spectroscopy, per-particle optical-simulation correlation on TEM grids, and, finally, *in situ* structural dynamics under pulsed excitation. Across these layers, the thesis builds a unified picture in which morphology is not merely an assumed input to simulation but an experimentally validated, dynamically evolving quantity with measurable consequences for plasmonic resonances.

The thesis begins by situating nanoscale plasmonics within broader technological and societal challenges. Global assessments by organisations such as the World Economic Forum, the IPCC, and major industrial consultancies repeatedly identify the energy transition, climate mitigation, and sustainable industrial processes as long-term priorities. Many technologies central to these goals, solar photovoltaics, catalytic CO<sub>2</sub> conversion, and materials for energy storage, depend on efficient nanoscale interactions between light, charge, and matter. Plasmonic nanoparticles, which support intense and tunable optical resonances, play a key role in these processes through their ability to scatter, absorb, and concentrate electromagnetic fields in sub-wavelength volumes. The first chapter therefore introduces the fundamental electrodynamics of plasmons, examines the material conditions under which metals support resonances, and explains the relevance of radiative and non-radiative damping pathways for photovoltaics, catalysis, and sensing. This establishes why accurate, single-particle understanding of plasmonic behaviour is both scientifically essential and technologically relevant.

Building on this context, Chapter 2 develops the foundations of electromagnetic simulations. Gold’s dielectric function is dissected into Drude-like free-electron behaviour

and Lorentzian interband contributions, and these analytical models are compared to experimentally measured dielectric datasets, which differ due to microstructure, grain boundaries, and ellipsometry modelling. The chapter quantifies the breakdown of the quasistatic approximation, showing that particles larger than roughly 20 nm, or elongated nanorods with long axes well beyond this threshold, require fully retarded electrostatics. A benchmark across major full-wave solvers (boundary element, finite-difference time-domain, discontinuous Galerkin, and discrete dipole) against the exact Mie solution demonstrates that the boundary element method uniquely combines accuracy, predictable convergence, and computational efficiency for metallic nanoparticles. Practical meshing rules, environmental modelling strategies, and dataset considerations are derived to ensure reliable simulations. These results constitute the computational backbone for all later quantitative structure–property correlations.

Chapter 3 turns to morphology acquisition. Electromagnetic simulations require watertight, quantitative surface geometries, yet the morphology of real nanoparticles is inherently three-dimensional, with subtle distortions introduced by synthesis, ligands, and sample preparation. Two-dimensional electron micrographs suffer from projection ambiguity: even small segmentation changes or projection effects can shift extracted lengths and widths enough to move plasmonic resonances by tens of nanometres. To overcome these limitations, the chapter develops a benchmarked tomography pipeline. Using simulated HAADF–STEM tilt series with known ground truth, including noise, blur, background drift, and a missing wedge, the performance of reconstruction algorithms, segmentation strategies, and meshing/simplification methods is systematically evaluated. Total-variation minimization consistently produces the lowest shape error, even for complex geometries such as nanotriangles and octopods. Otsu’s method proves robust for segmentation across noise levels, while minimum-threshold methods are only reliable for high-SNR datasets. Marching-cubes meshing with carefully controlled simplification preserves curvature and facet structure while keeping meshes tractable. The chapter shows how small morphological inaccuracies propagate into spectral errors, and how these errors scale strongly with geometric complexity and local curvature. The resulting tomography-to-mesh workflow forms a validated and reproducible route to simulation-grade morphology.

With the simulation and morphology tools established, Chapter 4 constructs the high-SNR quasi-darkfield platform used to measure single-particle spectra directly on TEM windows. A central stop placed at a pupil plane suppresses low-NA background while transmitting mid- to high-NA scattered light, ensuring strong contrast without a dedicated dark-field condenser. The chapter derives the angular illumination and collection windows from the optical geometry and shows how these windows determine

collection efficiency when combined with the angular radiation pattern of a nanorod. A pinhole sized from the measured and simulated Airy spots passes the diffraction-limited core of the scattered field while rejecting stray reflections. Careful index matching, stray-light management, and alignment yield reproducible measurements on both glass and SiO<sub>2</sub> TEM supports. To ensure spectral fidelity, the chapter introduces an instrument-response calibration using polystyrene scatterers and BEM simulations, along with a per-wavelength autofocus method based on fitting the 3D PSF ellipsoid. This establishes the experimental foundation for quantitative comparison between measured and simulated spectra.

Chapter 5 brings these components together to achieve quantitative morphology–property correlations on the same nanoparticles. Using TEM-based geometry (in two dimensions for scalability) and instrument-corrected quasi-darkfield spectra, the chapter compares measured spectra to BEM simulations across a discrete set of effective dielectric environments. For each particle, the best-matching environment is extracted by aligning simulated and measured peak positions using a simple regression procedure. Across a population of rods, the extracted effective-medium fraction clusters around a reproducible value of approximately 0.71, indicating a local environment dominated by immersion oil but with a significant air-like component caused by ligand shells, contamination layers, or small unfilled gaps at the substrate interface. TEM exposure predictably shifts plasmonic resonances, confirming the sensitivity of the method. While the effective-medium approach captures peak positions well, it does not fully reproduce linewidths or amplitudes, motivating the use of lossy, multilayer, or spatially varying environmental models, and highlighting the added value of 3D morphology for high-precision studies. The chapter thus provides a scalable two-tier workflow: rapid 2D-based correlations for population-level trends, and full 3D reconstructions for precision cases.

Finally, Chapter 6 examines how nanoparticle morphologies evolve when driven far from equilibrium. Using synchronised optical pulses delivered inside a TEM, the chapter records a frame after each pulse, producing pulse-by-pulse structural trajectories on millisecond timescales. The results reveal a wide range of reshaping behaviours: isolated rods that repeatedly melt and regenerate their tips; dimers exchanging material until reaching metastable locked states; long-range correlations between rods hundreds of nanometers apart; and reorientation of twin boundaries consistent with partial melting and recrystallisation. A tip-asymmetry metric quantifies the dynamics, showing an acceleration of structural change over time despite fast (<100 ns) thermal relaxation between pulses. Heat-diffusion simulations indicate that each 400 μs pulse drives the rods to a quasi-steady temperature of roughly 820 °C,

high enough for rapid surface diffusion and partial melting. Combined with the weak wettability of gold on carbonaceous supports, the chapter proposes a kinetic mechanism in which nanorods explore a landscape of metastable liquid-like configurations before recrystallising along biased pathways. This reveals a rich, history-dependent structural landscape fundamentally inaccessible to equilibrium TEM imaging.

Taken together, this thesis develops the conceptual, computational, and experimental foundations for quantitative single-particle plasmonics. It shows that accurate electrodynamic simulations require careful treatment of retardation, realistic dielectric functions, and mesh quality; that morphology extraction must be validated and precise, especially for complex shapes; that high-SNR quasi-darkfield spectroscopy enables reproducible, quantitative optical measurements on TEM grids; that per-particle environment extraction enables meaningful experiment–simulation alignment; and that nanoparticles under pulsed excitation explore highly dynamic and reversible structural states. By integrating these elements into a coherent framework, the thesis provides new tools and insights for understanding and ultimately controlling plasmonic behaviour at the level of individual nanoparticles, both in static and dynamically driven regimes.

---

# Samenvatting

Dit proefschrift ontwikkelt een volledig, kwantitatief ontwerp om de optische respons van individuele plasmonische nanodeeltjes te begrijpen, te simuleren en experimenteel te meten, met specifieke aandacht voor goudnanostaven. Het werk verbindt de elektron-schaalfysica van plasmonische excitatie met het reële, experimenteel waarneembare gedrag van nanodeeltjes, zowel in statische configuraties als onder sterke niet-evenwichtsexcitatie. Het combineert meerdere methodologische lagen: maatschappelijke en technologische motivatie, nauwkeurige elektrodynamische simulatie, betrouwbare driedimensionale morfologie-extractie, hoog-SNR quasi-donkerveldmicroscopie, per-deeltje correlatie tussen optica en simulatie op TEM-substraten, en, tenslotte, *in situ* structurele dynamica onder gepulste excitatie. Door deze lagen te integreren, laat het proefschrift zien dat morfologie niet slechts een aangenomen model voor simulaties is, maar een experimenteel verifieerbare en dynamisch veranderende grootheid met meetbare gevolgen voor plasmonresonanties.

Het proefschrift begint met het plaatsen van nanoschaalplasmonica in een bredere maatschappelijke context. Internationale technologische en risico-onderzoeken, onder meer van het World Economic Forum en het IPCC, benadrukken herhaaldelijk de noodzaak van koolstofarme energieoplossingen, duurzame productie en geavanceerde materiaalsystemen. Veel technologieën die hierin centraal staan, zoals fotonische elektriciteitsopwekking, CO<sub>2</sub>-conversie en geavanceerde elektrochemische materialen, zijn afhankelijk van efficiënte interacties tussen licht, lading en materie op de nanoschaal. Plasmonische nanodeeltjes kunnen licht op sterke en afstembare wijze verstrooien, absorberen en concentreren binnen subgolflengteschalen en spelen daardoor een belangrijke rol in dergelijke processen. Het proefschrift introduceert daarom eerst de basis van plasmonfysica, de voorwaarden waaronder metalen plasmonen ondersteunen, en de rol van radiatieve en niet-radiatieve dissipatie in toepassingen zoals fotonische katalyse. Dit vormt de fysieke motivatie voor gedetailleerde analyses op enkel-deeltjesniveau.

Daarna verschuift de focus naar de fundamentele elektrodynamische simulaties. Goud, het hoofdmateriaal van dit werk, vertoont een dielektrische respons die zowel vrije-elektronen als interbandovergangen omvat. Om deze respons nauwkeurig te modelleren, worden de Drude-achtige intrabandcomponenten en Lorentzianse interbandabsorptie ontleed en vergeleken met experimenteel gemeten dielektrische functies. Deze analyse toont aan dat geen enkel Drude-model de zichtbare en nabij-infrarode

respons kan beschrijven, wat het gebruik van gemeten datasets noodzakelijk maakt. Het proefschrift kwantificeert vervolgens de grenzen van de quasi-statische benadering: voor de hier bestudeerde nanodeeltjes is volledige, geretardeerde elektrodynamica nodig. Een systematische vergelijking tussen vier veelgebruikte simulatiemethoden (boundary element, finite-difference time-domain, discontinuous Galerkin en discrete dipole) en de exacte oplossing van Mie toont aan dat de boundary-elementmethode de beste combinatie van nauwkeurigheid, reproduceerbaarheid en rekentijd biedt. Praktische regels voor meshing, materiaalkeuze en omgevingsmodellering worden afgeleid en dienen als basis voor alle latere simulaties.

Met deze simulatiegrondslag verschuift het proefschrift naar de experimentele bepaling van morfologie. Elektrodynamische simulaties vereisen immers nauwkeurige, waterdichte oppervlakken, maar echte nanodeeltjes vertonen subtiele driedimensionale variaties als gevolg van synthese, liganden, kristaldefecten en meetcondities. Tweedimensionale TEM- of SEM-beelden zijn hierbij ontoereikend: kleine drempelwaarschommelingen bij segmentatie kunnen al tot merkbare verschuivingen in de optische respons leiden. Daarom ontwikkelt het proefschrift een volledige, gecalibreerde workflow voor elektronentomografie. Met gesimuleerde HAADF-STEM tiltseries met bekende grondwaarheid, inclusief realistische ruis, onscherpte, achtergrondvariaties en een ‘missing wedge’, worden reconstructiealgoritmen, segmentatiemethoden, meshing en mesh-simplificatie uitgebreid vergeleken. Total-variation minimization levert de meest nauwkeurige reconstructies, zelfs bij complexere geometrieën. Segmentatie met Otsu is robuust, terwijl ‘Minimum’-segmentatie enkel betrouwbaar werkt bij hoge signaal-ruisverhoudingen. De resultaten tonen hoe kleine morfologiefouten zich vertalen naar significante spectrale verschuivingen, vooral in hoog-gekromde gebieden. Hiermee ontstaat een gevalideerde, reproduceerbare pipeline van tomografie naar simuleerbare meshes.

Om plasmonresonanties op dezelfde deeltjes te meten die later zullen worden gecorreleerd aan hun morfologie, ontwikkelt het proefschrift vervolgens een quasi-donkerveldopstelling met een pupilvlak-stop die lage-NA-achtergrond onderdrukt terwijl midden- en hoog-NA-verstrooiing behouden blijven. De hoekacceptatie wordt analytisch bepaald en gekoppeld aan de gesimuleerde stralingspatronen van nanostaven. Een zorgvuldig gekozen pinhole laat alleen de primaire diffractie-lobe door, terwijl strooilicht van interne reflecties wordt onderdrukt. Door index-matching, verstrooid licht onderdrukking, en stabiele uitlijning wordt een hoge signaal-ruisverhouding bereikt, zelfs op SiO<sub>2</sub>-TEM-substraten. Vervolgens wordt een genormaliseerde instrumentresponsfunctie bepaald met polystyreensferen als referentie, en wordt een robuuste driedimensionale autofocus ontwikkeld op basis van ellipsoidale PSF-fitting

om chromatische defocus te vermijden. Deze optische set-up vormt de experimentele basis voor latere kwantitatieve correlatie.

Daarna brengt het proefschrift alle elementen samen om kwantitatieve morfologie-eigenschapcorrelaties te realiseren. Met TEM-geometrie (hier in 2D voor statistische schaalbaarheid) en IRF-gecorrigeerde spectra worden boundary-element-simulaties uitgevoerd voor verschillende effectieve omgevingen. De vergelijking tussen gemeten en gesimuleerde piekposities levert voor ieder deeltje een optimale effectieve brekingsindexfractie. Over een populatie van nanostaven resulteert dit in een consistent gemiddelde van circa 0.71, wat duidt op een lokale omgeving die voornamelijk door immersie-olie wordt bepaald, maar met een significante bijdrage van luchtachtige of koolstofrijke zones rond het deeltje. TEM-blootstelling veroorzaakt voorspelbare piekverschuivingen, wat de gevoeligheid van de methode bevestigt. Tegelijkertijd blijkt dat amplitude- en lijnbreedte-veranderingen niet volledig met een enkel effectieve-omgeving-model kunnen worden gereproduceerd, wat wijst op de noodzaak voor modellen met verlies of ruimtelijk variërende omgevingen, vooral op TEM-membranen. De analyse leidt tot een tweesporenstrategie: snelle 2D-correlaties voor grote datasets en trends, en volledige 3D-morfologie voor ultieme kwantitatieve precisie.

In het laatste deel wordt onderzocht hoe nanodeeltjes zich gedragen wanneer ze sterk uit evenwicht worden gebracht. In een *in situ* TEM worden goudnanostaven blootgesteld aan gesynchroniseerde optische pulsen; na elke puls wordt een beeld genomen. Deze aanpak onthult een rijke dynamiek van herstructurering op milliseconde-schaal: geïsoleerde staafjes die hun uiteinden telkens herconfigureren, dimeren die massa uitwisselen totdat een metastabiele toestand wordt bereikt, langeafstandsinteracties tussen deeltjes die tientallen tot honderden nanometers uit elkaar liggen, en de vorming en rotatie van tweelingsgrenzen van frame tot frame. Warmtesimulaties tonen aan dat elke puls een quasi-stationaire temperatuur van circa 820°C veroorzaakt: hoog genoeg voor snelle oppervlakdiffusie en gedeeltelijke smelting, gevolgd door snelle afkoeling op een tijdschaal van circa 100 ns. In combinatie met de zwakke bevochtiging van goud op koolstof en de hoge atoommobiliteit op amorfe of grafische oppervlakken ontstaat een kinetisch gebied waarin nanodeeltjes tijdelijk vloeibaar worden, een landschap van metastabiele configuraties verkennen, en via omgevingsgeïnduceerde voorkeurstrajecten weer kristalliseren. Deze reversibele, geschiedenisafhankelijke herstructurering is onzichtbaar in standaard, statische TEM-experimenten.

Gezamenlijk ontwikkelt dit proefschrift de conceptuele, computationele en experimentele fundamenteen voor kwantitatieve enkel-deeltje-plasmonica. Het toont aan dat nauwkeurige elektrodynamische simulaties zorgvuldig omgaan met retardatie, realistisch

che diëlektrische functies en numerieke convergentie vereisen; dat morfologie-extractie gevalideerd en precies moet zijn, vooral voor complexe vormen; dat nanodeeltjes onder gepulste excitatie zeer dynamische en omkeerbare structurele toestanden kunnen verkennen; en dat hoog-SNR quasi-donkerveld betrouwbare, kwantitatieve vergelijking tussen metingen en simulaties mogelijk maakt. Door deze componenten te integreren tot één coherent ontwerp, biedt het proefschrift nieuwe inzichten en gereedschappen om plasmonisch gedrag op het niveau van individuele nanodeeltjes te begrijpen en uiteindelijk te sturen, zowel in statische als in dynamisch aangedreven regimes.

---

# The Cast – Acknowledgements

*The Cast – the actors in a film, play, or show*

*A PhD does feel a bit like a movie where you are the main character going through an emotional rollercoaster. Luckily, there are other characters in the movie as well, that can help you get through it in the best way possible. I was lucky to have many people around me during my PhD, who all had their impact, big or small, and I would like to sincerely thank them for that.*

*I would like to start by thanking the director, who found funding for the movie, and guided me through it. **Wiebke**, bedankt voor het vertrouwen om mij als eerste PhD student in je groep te laten starten. Zoals tijdens elk promotietraject gebeuren er wel eens onverwachte dingen, maar ik kan altijd bij je terecht voor advies. Dat is echt zeer waardevol. Ook vind ik het super fijn dat we altijd eerlijk en efficiënt met elkaar kunnen communiceren. Daarnaast vind ik het geweldig dat je bij sollicitanten niet alleen maar op hun wetenschappelijke achtergrond let, maar ook of diegene binnen de onderzoeksgroep past. Het is net zoals tijdens die lasergame escape room: je geeft echt alles voor je team. *This movie also had a co-director, who was involved closely during some of its parts.* **Albert**, bedankt voor alle kansen en het vertrouwen de afgelopen jaren. De PV course is een van de redenen dat ik de scheikunde op de UvA achter me heb gelaten en bij AMOLF aan de slag ben gegaan. Toen je Francesca en mij uitnodigde om bij jullie groepsuitje naar Schiermonnikoog aan te sluiten wist ik het zeker: AMOLF is een plek waar mensen echt naar elkaar om kijken en dat is precies wat jij ook uitstraalt. Nog een komische uitkomst van de afgelopen jaren: ik heb je eigen interviewtactiek op jou mogen toepassen en volgens mij heeft het gewerkt. *And then, there was my recruiter, who allowed me a role in the cast of the prequel.* **Esther**, thank you for trusting me to start my master's project during the corona pandemic within your group. You and your group taught me how to create cohesion within a research group and I definitely tried to apply that to my 'new' research group as well. *Of course, movies also get reviewed by critics, who I would like to thank as well.* **Jorik**, bedankt voor alle discussies tijdens postersessies, je hebt een goed oog voor detail en stelt vaak vragen vanuit een compleet nieuwe invalshoek. **Stephan**, thank you and your wife for welcoming me *impromptu* at Rice university a couple of years ago. I showed up unannounced, got a lab tour from Stephen Lee, and even managed to chat to you for a bit on what already seemed like a busy day. **Peter**, bedankt voor de feedback tijdens mijn presentatie op Physics@Veldhoven. Ik heb mijn simulaties nog*

nauwkeuriger kunnen maken. **Paul** en de bovengenoemde commissieleden, bedankt voor het lezen en beoordelen van mijn thesis.

*Then, I would like to thank my lawyers, who help defend me against the critics, or in PhD terms, my paronyms.* **Igor**, mijn DJ leraar, PV bestuurslid, en the P&O fairy, bedankt voor het doorstaan van alle woordgrappen de afgelopen jaren. Ik heb geweldige herinneringen aan onze werkuitjes: van hotelontbijtjes en speakeasies tot aan nachttreinen en -boten. Maar daarnaast zagen we elkaar ook vaak op AMOLF zelf. Ik ben ontzettend dankbaar voor alle gezelligheid die we hebben gehad en hebben gecreëerd. *This brings me to my second paronym, who is currently in my research group.* **Quynh**, who would've thought that we end up in the same research group after all. Faith wants us to have some secret kind of competition: applying for the same PhD position, you already being a postdoc, but me defending my thesis a couple of days before you. However, instead of letting that drive us apart, I think it brought us closer (yes, cliché I know). But really, I think I've had my most vulnerable chats in life with you, and you always comforted me, either with excellent life advice, or with pictures of your majestically cute cats.

*Of course, there were more people in my research group. Someone who entered the movie only in the second half, but we did become each other's sidekick, is Michael Scott.* **Sergio**, you may have convinced me that the more politically incorrect you can think, the more open-minded and creative you can be as a person. At this point we have so many inside jokes that we could probably spend a whole Friday afternoon laughing our asses off without saying a word. I highly appreciate you for this. **Francesca**, I am grateful that you always stayed real and honest with me. People often hide and lie whenever a situation becomes uncomfortable. It's very brave to be confrontational and keep communicating, even in tough times. In the end, I think we did a great job in creating a nice atmosphere within our group. Our research group rocks! And that's also because of us. You should be proud, and not only because of this, but also because of the great researcher you have become! Really! **Ethan**, or shall I say lord Farquaad? Thank you for providing me with a lifetime worth of conspiracy theories that might actually not be conspiracies at all. This intrinsic curiosity and drive for discovery will bring you that *Science* paper, I'm sure about it. Just make sure that you don't have to travel when picking up your Nobel prize. I volunteer as stand in. **Anezka**, of we zijn allebei meer haveremelkelite dan we willen toegeven, of we hebben daadwerkelijk gedeelde interesses. Los daarvan kom ik je graag nog vaker random tegen in de stad om even heerlijk sarcastisch over werk te praten. Ook kom ik nog graag een keer bij je band kijken, invite me please! **Loriane**, if I ever need advice on chocolate I hope I can give you a call. Thank you for your modest but honest way

of communicating, I enjoyed working with you. **Nathan**, apart from your inspiring lifestyle, I think you could be a great teacher with these presentation and explanation skills. You are literally the ideal postdoc and I am happy that you joined our group. **Davide**, I am still impressed by how fast you picked up the knowledge on plasmonic dimers. This project has been going on since I started, but you showed the most promising results so far! Also thank you for supporting AFC AMOLF together with **Ece**; I spent years figuring out what meshes and simulation configurations are the best. It is insane to see that you picked it up so fast and I'm sure you will finish your project with a great result.

*It has been a long movie with many characters, and some of them are no longer on the screen. I would like to thank the people that have already left our research group.* **Patrick**, thank you for a crash course on optics and elaborately teaching me about nano optics. Without you, the setup would not have been there. You always remain positive and I cherish great memories of our kayak experience with Francesca in Het Twiske. **Jesper**, je bent de enige masterstudent die ik heb mogen begeleiden en dan ook nog eens als co-supervisor. Desalniettemin ben ik onder de indruk van je geduld: Ik bleef maar vragen of we nog meer berekeningen konden doen, en... je deed ze. Tegelijkertijd heb je misschien wel een nanowereldrecord gemeten, maar je was nog wel het meest waardevol voor AFC AMOLF. **Floor**, volgens mij ben ik er nog steeds niet helemaal achter hoe aardig en slim jij bent. Ook vraag ik me af waar jouw doorzettingsvermogen ophoudt. Toen jij vanwege een hockeybal met hechtingstape op je gezicht naar AMOLF kwam werd dit bevestigd. **Devin**, you are unhinged in the best way possible and your love for writing and language is amazing. I still have the handwritten Chinese: 'you are one of a million' on my desk and I try to still skedaddle everyday to go do some shenanigans. Thank you for all your joy! **Mareike**, you were not only very organized and prepared in your project, but also in your view on life. Thank you for the critical discussions we had. I hope you have a wonderful and impactful career! **Tjom**, in die paar maanden die je hier was, heb je meer bereikt dan 2 postdoc jaren die volgden. Ik wens je veel succes nog in je eigen PhD, dat komt vast helemaal goed! **Marco**, you reminded me that I would not be made for AUC, but you handled all this work perfectly. I wish you a great scientific career! **Stan**, jij dacht vast al: Waar blijf ik in dit rijtje? Dat komt omdat jij altijd alle details ziet, respect. Nou is het zo, dat jij nog op AMOLF bent, maar niet meer in onze groep zit. Toch zorgt jouw betrokkenheid ervoor dat het niet voelt alsof jij onze groep hebt verlaten en dat is echt veel waard. Bedankt voor alle fijne discussies!

*This brings me back to the rest of AMOLF, and here it becomes a bit messy as characters have played multiple roles sometimes. So don't get confused if you find your*

*name only at a certain part, it is just that I know you best from that.* Allereerst wil ik de directe technische support bedanken die ik heb gekregen tijdens het bouwen van de setup: **Marko, Dion, Xander**, en **Ivo**, super bedankt voor alle hulp bij de hardware en software kant van SParC.

Then, I would like to thank the staff association (PV) where I spent 3 years of my AMOLF time. **Evelijn**, iemand noemde ons ooit de Paula en Huib van de PV en dat klopt best wel goed. Wat zijn wij een goed organisatie-team zeg; we zitten (bijna) altijd op dezelfde lijn. Ik ben je ook super dankbaar voor elk moment dat je langskwam (vaak om half 5) om over van alles en nog wat te praten. Dat leek misschien op een afleiding maar had ook de functie om weer nieuwe energie voor het werk te creëren. **Daan**, wat ben jij een goede DJ zeg, als je ooit nog een Grilliams-achtig evenement doet dan hoop ik een VIP-pas te krijgen. **Masha**, you are the only person with whom I have a special greeting. I can always be weird with you but also have deep and inspiring chats. **Imme**, jij bent een van de meest creatieve mensen die ik ken en daarnaast ook nog eens super behulpzaam. **Elaina**, we have had quite some party fun, and I really appreciate your level of wokeness, it creates a perfect balance. I would also like to thank **Yorick, Sofija, Timo, Mels, Moritz, Wessel, Bas, Jente, Yvonne, Megan, Manuel, Margo**, and **Anne-Sophie** for our shared time at the PV.

Another amazing group of people was part of Amsterdam Football Club AMOLF, or AFCA. I would like to thank **Bernat, Agustin, Christiaan, Tom, Kasper, Jon, Giorgos, Oleg, Laurin, Achille, Jack, Tuoyu, Finn**, and **Juzhe** for all the great matches we played.

I am also deeply grateful for the diversity day I organized together with **Parisa, Bibiane, Cherin**, and **Ananya**, we had so many interesting discussions and above all, a lot of fun.

We hebben de sfeer op AMOLF ook met de buitenwereld weten te delen via ExpeditieNext. **Linde**, je bent een echte bron van gezelligheid en een geweldige violist. Ik ben fan en hoop nog vaak optredens van je te zien. **Falco**, met jou heb ik de snelste autorit van 3 uur ooit gehad. Wat een verhalenmachine en grappenmaker ben jij. **Petra**, bedankt voor alle outreach kansen, ik kom graag een keer mee vogelen. **Jorijn, Lidewei, Marloes**, zonder jullie had ExpeditieNext nooit zo kunnen slagen, super bedankt!

Another external endeavour was with Physics with Industry. **Amber, Parsa, Semanur, Jingmin, Rick**, thanks a lot for an intense but successful week.

This brings me to maybe the most fun days I had at AMOLF, which was the LMPV summer school. **Daphne D.**, please stay the silent whisperer you are, I love some good

gossip. **Larissa vdV.**, bedankt voor alle gezelligheid en wetenschappelijke discussies. **Rohit**, thanks for teaching me some new dance moves, you rock. **Jerome**, I am still impressed by your knowledge on optics, I might consult you in the future. **Antony**, good morning, sir. **Nika**, thanks for all the great laughs and series recommendations. **Saskia**, I envy your perseverance. **Bruno** and **Erik**, thank you for sharing some very personal stories during the cultural intelligence workshop. **Omolara**, I'm looking forward to a dinner with more than 10 pieces of cutlery with you. And thank you for all the fun chats, **Marcel**, **Robin**, **Luna**, **Shi Wei**, **Fanny**, **Francesco**, and **Alex**.

Then, there are people within LMPV I either didn't manage to chat with during the summer school or that weren't there: **Basita**, **Mike**, **Sarah**, **Hollie**, **Alvaro**, **Lars**, **Jaime**, **Larissa B.**, **Sam**, **Susanne**, **Lucienne**, and everyone I forgot to mention now, thank you all for the great time!

I am sure that I forget to thank some fellow AMOLF scientists, but I would like to thank **Susanna B.**, **Susanna E.**, **Nebojsa** for the nice chats we had over the years.

I also had a lot of support at AMOLF. **Isabelle**, ik heb helaas geen tijd gevonden om je terug te pakken met een sinterklaasgedicht. **Hinco**, bedankt voor alle enthousiasme en hulp bij de PV. **Bob K.**, bedankt voor alle schaduw PV activiteiten die je hebt verricht. **Remko**, bedankt voor alle ICT magie. **Berend**, bedankt voor de hulp met elektronica. **Roy**, bedankt voor alle gezelligheid in York en het doorstaan van Igor's en mijn humor. **Juliette**, bedankt voor alle goedemorgens en het interessante gesprek over zonnetjes. **Nik** bedankt voor je grappige winacties. En ook bedankt **Teressa**, **Max**, **Ulas**, **Bob D.**, **Esther**, **Marc**, **Gino**, **Richard**, **Rutger**, **Tarik**, **Ilja**, **Wilco**, en **Marco vdL.** voor de ondersteuning de afgelopen jaren.

Then, there are some people I would like to thank that already left AMOLF: **Susan**, **Sven**, **Andrea**, **Nick**, **Matthias**, **Stefan**, **Jeroen**, **Deba**, **Lucie**, **Nelson**, **Julia**, **Eitan**, and then I probably still forgot some people. I would like to give a special shoutout to **Dhawal**. You, your wife **Lisha**, and brother **Sawal** are the best hosts one can imagine. Next to that you are an amazing football captain and team player that cares about and sees everyone. Thank you for everything you organized!

Outside of AMOLF I would first like to thank my collaborators: **Alexander**, **Nathalie**, and **Sara**.

Next to the office hours, I had a lot of support from other friends and family. Thank you for organizing the occasional football games **Steven** and **Matteo**, where I got closer to **Paolo**, **Davide**, **Kyros**, and **Marijn**. I spent some nights with **Iso**, **Alan**, **Mars**, **Maurizio**, **Anna**, **Kai**, and **Max**. Thank you for making me feel safe and all the fun time! The first years of my PhD I lived together with **Thomas**, **Niek**,

and partially **Gal**. Thank you for all the positivity and cosy nights at home. En over gezelligheid gesproken, super bedankt ‘**Superfun**’ en in het bijzonder **Sylvia**, **Thomas** en **Jesper** voor alle tijd die we al met elkaar hebben doorgebracht. Ook wil ik graag mijn bachelorvrienden **Siza**, **Davita**, **Marit**, **Ninke**, **Eveline**, **Sophie**, **Marie-Lou**, **Tori**, en **Daan** bedanken.

Before I close these acknowledgments, I would like to thank my second homes. The first one is here in the Netherlands with **Sara** and **Lia**. Thank you so much for all the positivity, comforting conversations, and fun we had. The second one is in Italy. I would sincerely like to thank all these people for either inviting me to their homes or having a fun time with me: **Simonetta**, **Roberto**, **Luca**, **Grazia**, **Brynn**, **Gilberto**, **Filippo**, **Gloria A.**, **Gloria B.**, **Malano**, and **Laura**.

Als laatste wil ik mijn familie bedanken voor alle steun. **Papa** en **mama**, ik kan altijd bij jullie terecht en daar ben ik echt enorm dankbaar voor. Jullie hebben me geleerd wat doorzettingsvermogen betekent en me er altijd aan herinnert dat stress goed is, maar wel tot op zekere hoogte. Dit is echt wat mijn PhD tot een succes heeft gemaakt. **Donna** en **Loes**, ondanks dat we elkaar wat minder vaak zien dan vroeger, voel ik nog wel altijd jullie enorme steun. Alle Snapchats en berichtjes maken mijn dagen stukken leuker. En als we elkaar dan zien, dan is het net zo gek doen als vroeger, heerlijk. **Monique**, onze jaarlijkse concerten zijn een traditie geworden die ik enorm waardeer. Ook wil ik je enorm bedanken voor je oprechte interesse in mijn leven en PhD, dat heeft me enorm geholpen om mijn onderzoek uitlegbaar te maken. **Opa** en **oma**, jullie enthousiasme als ik uitlegde over mijn PhD was aanstekelijk, super bedankt hiervoor.

**Marco**, I don’t think I could have done my PhD without you. I honestly think you are the reason I was never too stressed or too sad about failed experiments. You give me so much energy that I sometimes don’t know what to do with it. And that is not because of the amazing food you cook but because of the person you are. I love you!

---

# About the author

Mees Dieperink was born on April 30, 1998 in Amsterdam. He grew up in Wijk aan Zee and obtained his gymnasium high-school diploma from the Kennemer College (now Castor college). He continued his studies at the University of Amsterdam and the Vrije Universiteit Amsterdam, where he graduated with a bachelor's degree in chemistry in 2019. The master Science for Energy and Sustainability (SfES) sparked his interest, but Mees decided to keep chemistry as his main subject, while following courses of the SfES track. Here, he became fascinated by sustainable energy materials, and particularly light-matter interactions. After graduating with a master thesis on the electrochemical growth of silver nanoparticles in 2021, Mees started his PhD at AMOLF in the group of dr. Wiebke Albrecht. His research focused on the interaction of light with gold nanoparticles, in particular how this interaction correlates quantitatively to morphology. The results of this work are presented in this thesis. Next to his scientific interests, Mees enjoys playing football, going to concerts, and hosting friends.

