Hot Nanoparticles

Jollans, T.G.W.

5 Time-resolved measurement of electronic temperatures in a single gold nanoparticle

Anti-Stokes nanoparticle photoluminescence, in which light extracts energy from the thermal bath in the nanoparticle, is an indicator of the temperature prevalent within a nanoparticle. We use it to measure the electron temperature in a single gold nanoparticle under pulsed illumination (~10^3 K) and its strikingly nonlinear dependence on pulse energy. Pump-probe anti-Stokes spectroscopy allows us to measure ultrafast dynamics of a pulse-excited hot electron population with sub-picosecond time resolution. We measure the heating and cooling, all within picoseconds, of the electrons and find that the highest apparent temperature state is reached 0.6 ps before the maximally excited state.

5.1 Introduction

5.1.1 Background

Plasmonic nanoparticles are remarkable in their ability to couple to light whose wavelength is significantly larger than their size. This property is primarily due to their localized surface plasmon resonance — a resonance of collective electron oscillations enabled by the spatial confinement of the electron gas.

Leaving aside the rich physics created by various nanoparticle geometries, from tunable spectroscopic properties to powerful near-field effects, absorption of light by a plasmonic nanosphere can be understood in the following simplified way:

If the dielectric permittivity of the material is known, the key parameters of scattering and absorption can be calculated using Mie theory. For far-field interactions, these are conveniently given in terms of cross sections: the scattering cross section σ_{sca}, the absorption cross section σ_{abs}, and the extinction cross section $\sigma_{\text{ext}} = \sigma_{\text{sca}} + \sigma_{\text{abs}}$. Through σ_{abs}, we have access to the rate at which energy is absorbed by the nanoparticle, and thus to the amount of energy that must (eventually) leave the nanoparticle when it equilibrates with its environment.

The vast majority of the energy absorbed is converted to heat, which the (now hot) nanoparticle will dissipate. A vanishingly small fraction of the energy absorbed goes towards radiative emission channels referred to as photoluminescence; while this emission is measurable and useful — and indeed central to the technique used in this chapter — for the purposes of the energy balance, it can safely be neglected ($P_{\text{pl}} \ll P_{\text{heat}}$). This approached based on scalar cross-sections is the approach taken in the previous chapters of this thesis.

While this simple picture works well for low-electric-field near-steady-state continuous-wave excitation, it ignores the dynamics of light absorption, and the ultrafast heat transfer dynamics which occur in the metal. These come to the fore in particular when a plasmonic nanoparticle is excited with a pulsed laser with a pulse width below about a picosecond.

Ultrafast laser absorption in a metal, be it in a nanoparticle, a film, or bulk metal, is typically modelled using a two-temperature model; the electron gas and the metal lattice are treated as two distinct coupled subsystems with independent temperatures. In bulk metal, both temperatures must be treated

as temperature fields varying in space. In nanoparticles that are sufficiently small compared to the wavelength and focus size of the heating laser, we can assume uniform heating. The temperatures of the electrons and of the lattice can then be taken to be constant across the entire volume of the particle, and may be treated as scalars.

In the two-temperature model, under ultrafast laser excitation, we typically assume the electrons are heated to some temperature $T_e(t = 0)$ instantaneously. This requires that the temperature be well-defined, i.e. that the electrons thermalize to a Fermi–Dirac distribution with no significant delay. The characteristic thermalization time is of order ~0.5 ps in gold [22, 111]; however, the electron distribution being instantaneously non-thermal barely has an impact on most measurements, especially if the excitation pulse width is of the same order. Assuming instantaneous thermalization, we model the evolution of the temperatures using the following coupled ordinary differential equations [105]:

$$C_e(T_e)\frac{\mathrm{d}T_e}{\mathrm{d}t} = -G(T_e - T_l) \tag{5.1}$$

$$C_l\frac{\mathrm{d}T_l}{\mathrm{d}t} = G(T_e - T_l) - Q_{\mathrm{bath}}(T_l), \tag{5.2}$$

where T_e and T_l are the temperatures of the electrons and the lattice, C_e and C_l are the corresponding heat capacities, G is an electron–phonon coupling constant, and Q_{bath} represents heat transfer from the lattice to the environment. Note that the electronic heat capacity depends on the electronic temperature. The electronic excitation generally relaxes within a few picoseconds, while the lattice cools down over a longer period of time, which is limited by the thermal properties of the environment, especially the thermal conductivity and, for small particles in particular, the interfacial thermal resistance. Under typical conditions, relaxation of the lattice takes place over a period of hundreds of picoseconds.

Because of the low heat capacity of the electrons (viz. about $2 \times 10^{-2}\,\mathrm{J\,cm^{-3}\,K^{-1}}$ [105], or 0.8 % of that of the gold lattice, at STP), the instantaneous temperature of the electrons upon pulsed excitation is very high. It has been estimated under various conditions to reach into the thousands of kelvin under modest illumination, both indirectly from extinction and/or scattering measurements [105, 112, 113], and more directly, using anti-Stokes emission (see § 5.1.2).

The two-temperature model, by allocating the electrons to a Fermi distribution with a certain elevated temperature T_e at time zero, does not take into account how the electrons got there. Considering this is important for two main reasons: Firstly, if the excitation laser pulse width is comparable to the electron–phonon coupling time, then they must, in principle, be considered together. Secondly, there is no reason to believe that optical excitation of the electron gas yields a thermal distribution of electrons, and this must be considered if the excitation laser pulse approaches the electron–electron coupling time. Indeed, several studies have considered non-thermal electron distributions and how they thermalize [22, 114–116].

5.1.2 Anti-Stokes emission as a measure of temperature

When a nanoparticle is illuminated at a suitable wavelength, it can be detected optically, either through elastic (Rayleigh) scattering, in which the detected photons have the same energy as the illumination, or through inelastic processes, in which the detected photons have either gained energy from or lost energy to the nanoparticle. If the emitted photons have a higher energy (blue-shift), this is known as an anti-Stokes process, if they have a lower energy (red-shift), a Stokes process.

The simplest such inelastic process would be a single scattering event, in which a single incoming photon interacts once with a single particle (such as an electron, a hole, or a phonon) in the material, exchanges some energy and momentum with it, and is reemitted. Such lowest-order processes are collectively known as Raman scattering. A wide variety of higher-order processes may also occur: an interaction may involve two or more photons, or two or more electrons. The photon or photons may create an excited state, and this exciton may then decay radiatively some time later after some number of interactions with electrons, holes, and phonons.

All this is vague enough to be true for any sample which interacts with light, but which types of interactions will dominate varies considerably. The main variable here is the availability of excited states: in insulators or semiconductors with a band gap greater than the photon energy, any excited states are negligible and the inelastic interactions are dominated by the lowest order process, Raman scattering. In molecules and semiconductors which happen to have long-lived excited states resonant with the incoming light, these, instead, dominate the inelastic emission. This is known as fluorescence.

In metal nanoparticles specifically, excited states (electron-hole pairs, and, collectively, localized surface plasmon polaritons) are available, but they have relatively short lifetimes on the order of tens of femtoseconds [117]. This means that they do not dominate and that the excitons only experience a limited number of interactions with the thermal bath. In other words, the excited states decay radiatively before they have fully thermalized with their environment — this is known as 'hot' luminescence.

In general, certainly in metals, Raman scattering and hot luminescence will always occur together [118], and cannot easily be distinguished from one another in practice. In any case, energy is gained and lost through interactions with the thermal bath. Thus, anti-Stokes emission, in which the emitted photons have an energy $\hbar(\omega_L + \delta\omega)$, which is larger than the energy $\hbar\omega_L$ of the laser illumination, can only occur to the extent that occupied states with an energy of $\epsilon = \hbar\delta\omega$ from which energy can be extracted are present in the metal. Both phononic and electronic states may, in principle, contribute, to these processes[1].

The emission spectrum will depend on on the energy distribution of occupied 'hot' states which might donate energy, $n(T, \epsilon)$. Broadly speaking, this will follow a Boltzmann-type distribution: If hot electrons dominate the interactions, which is likely at large $\delta\omega$ due to the higher temperature achieved by the electrons, it should obey Fermi–Dirac statistics. If, on the other hand, interactions with phonons dominate, it should follow Bose–Einstein statistics. In practice, the differences between the predictions of the three distributions are marginal for range of $\delta\omega$ which we have access to experimentally[2]. The emission can then be used to deduce an electronic temperature which matches up well with *a priori* predictions [119, 120] and, by extrapolating to zero laser heating, to measure the temperature of the environment [121].

In the following, the term 'photoluminescence' will be used for inelastic emission in a broad sense, without prejudice as to the mechanism which causes it.

[1]The terminology of Raman scattering is a case in point: 'traditional' Raman scattering exchanges energy with vibrational states, while electronic Raman scattering is universally couched in terms of scattering with quasi-free electrons, which dominate the behaviour of metals

[2]Nota bene, other authors have made the same observation about their measurements. [119]

5.2 Method

5.2.1 Premise

We probe the dynamics which occur in response to pulsed irradiation of a gold nanosphere using time-resolved anti-Stokes spectroscopy; to achieve ultrafast time resolution, we use a two-colour pump-probe technique:

Two ~ 350 fs laser pulses are sent to the sample with a particular delay τ between one and the other. The first, the 'pump' pulse, with a central wavelength of $\lambda = 785$ nm, is far to the red of the localized surface plasmon resonance of the nanosphere. The second, the 'probe' pulse, at $\lambda = 594$ nm, is near the resonance. We acquire emission spectra in the neighbourhood of 594 nm using spectral filters to remove the signal from the lasers and from elastic scattering.

Both colours are absorbed by the nanosphere, though the pump beam is absorbed significantly less efficiently. What's crucial is that since the pump laser is ca. 0.5 eV to the red of both our observation range and the plasmon resonance, (i) any possible emission is not enhanced by the plasmon, and (ii) whatever inelastic emission may be detectable without plasmon enhancement is spectrally separated from the probe signal. In other words, the pump is, on its own, effectively invisible.

The pump pulse creates some excited population of hot electrons $\Delta n_{\mathrm{ir}}(t, \epsilon)$. The probe pulse, arriving a delay of τ later, also creates an excited population of hot electrons, $\Delta n_{\mathrm{vis}}(t - \tau, \epsilon)$. However, its anti-Stokes emission depends on the total populations of electrons, which we may write as

$$n(t, \epsilon)\Big|_{t=\tau} = \Big[n_0(\epsilon) + \Delta n_{\mathrm{vis}}(t - \tau, \epsilon) + \Delta n_{\mathrm{ir}}(t, \epsilon) \Big]_{t=\tau} \tag{5.3}$$

if higher-order non-linear effects are small enough for the contributions of the two pulses to be additive. By varying the delay τ between the pump and probe pulses, as sketched in fig. 5.1, and examining the anti-Stokes spectra due to the probe, we can then elucidate the dynamics of $\Delta n_{\mathrm{ir}}(t, \epsilon)$ and the electron temperature with sub-picosecond time resolution.

5.2.2 Experimental details

Two correlated laser pulses are prepared using a titanium-sapphire (Ti:Sapph) laser, which produces a 75.8 MHz train of near-transform-limited pulses with

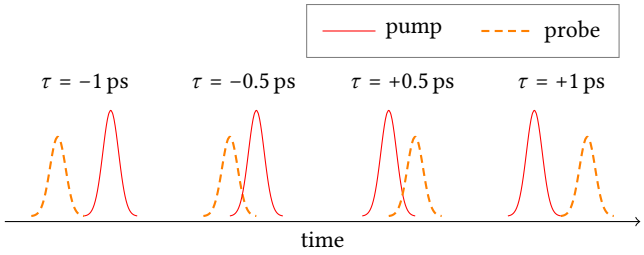

Figure 5.1: Visual representation of different pulse delays τ. Note that, typically, $\tau \geq 0$ for the pump pulse to have an effect on the measurement with the probe.

a central wavelength of 785 nm (the 'pump' pulses), and with a frequency-doubled optical parametric oscillator (OPO), which is pumped by the Ti:Sapph laser. The fact that one laser pumps the other means that every pulse from the OPO corresponds to a Ti:Sapph pulse — the two are locked together. The visible output of the OPO is tuned to 594 nm (the 'probe'). Both pulses individually pass suitable dielectric band-pass filters which are well-matched to the notch filters in the detection. Additionally, we use a solid state continuous-wave 532 nm laser to identify particles, for fine adjustments, and for CW photoluminescence spectra.

The pulse width is measured to be equal to approximately 350 fs after an acousto-optic modulator in the beam path, which considerably lengthens the pulse. This component is needed only for the acquisition of transient extinction time traces (cf. chapter 4). Further optical components other than the objective should have little to no further effect on the pulse width; we can estimate the effect of the objective by approximating it as a solid block of glass with a length of ~5 cm; for BK7 glass, this would result in a negligible lengthening of the pulse by 0.06 fs (at 785 nm) to 0.14 fs (at 594 nm). The equivalent values for other glasses are similar.

The delay τ between the pulses is adjusted using an optomechanical delay line with a length of up to 1 ns in the path of the near-infrared (pump) beam. All three are then carefully overlapped and tightly focussed on the sample with an oil-immersion objective (Olympus, NA = 1.4). A second objective on the far side of the sample (Olympus, NA = 0.75) collects the transmitted light, which passes a spectral filter to remove the NIR component before

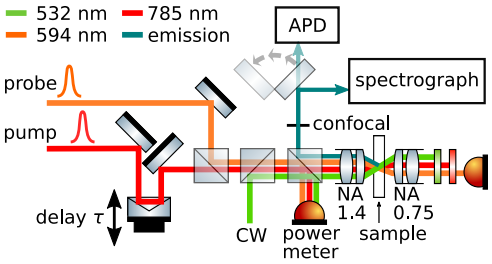

Figure 5.2: Highly simplified sketch of the experiment showing the logical beam paths.

being focussed on a fast photodiode (FEMTO Messtechnik). This arrangement allows measuring the change in extinction of the visible pulse as a function of inter-pulse delay and the acquisition of transient extinction time traces (as in chapter 4), but is further not required for the work discussed here. The pulse overlap is optimized using pump-probe extinction contrast (for the two pulses) and photothermal contrast (to overlap the pulses with the CW laser).

Meanwhile on the near side of the sample, the light (back-scattering and photoluminescence) collected by the NA = 1.4 objective is split off into the confocal detection path with a 50:50 beam splitter. After passing through a confocal pinhole (50 µm) and a series of notch filters for both colours, the emitted light is sent either to an avalanche photodiode, used for focussing, or to a liquid N_2-cooled spectrograph (Acton Research SpectraPro-500i). Fig. 5.2 shows a rough logical sketch of the beam paths from source to detection, disregarding most optical components.

Photoluminescence spectra are acquired with an integration time of 240 s, with the exception of some spectra recorded at very low or very high intensity.[3] All spectra shown here have been normalized by the integration time. Before the acquisition of each spectrum, the particle is brought into focus by maximizing photoluminescence with an automatic focussing routine. This is done to compensate for any slight drift in the system that may occur over the course of a long measurement. During the acquisition, the inter-pulse delay τ is kept constant.

Quoted pulse energies are measured as average powers in the back focal plane, before the objective, as indicated in fig. 5.2. They are not normalized by the transmission of the objective or by the absorptivity of the sample.

[3] The very long integration time is necessitated by the low intensity of the signal from a single nanoparticle. In chapter 4, the need for such long integration times is obviated by the use of a lock-in scheme.

The sample consists of 100 nm gold nanospheres (Nanopartz Inc.) spin-coated on a glass cover slip (Menzel). Spheres were chosen as other geometries are more liable to reshape under pulsed illumination. The nanoparticles are very dilute; in all cases the nearest neighbour of the nanoparticle being studied was more than ~2 μm away. The nanoparticles on the glass are further immersed in a reservoir of ultrapure water.

5.3 Results

5.3.1 Dependence of the electronic temperature on intensity

In order to establish clearly which parts of the resulting anti-Stokes spectra are due to heating by the probe pulse — in the language of eq. (5.3), in order to get an idea of $n_0(\epsilon) + \Delta n_{\text{vis}}(\epsilon)$ — we first measure Stokes and anti-Stokes spectra of the nanoparticle using the probe alone, at different excitation powers. The spectra are shown in fig. 5.3.

In general terms, if the anti-Stokes emission results from interactions with a Boltzmann-form thermal bath, we can say that the intensity should follow

$$I_{\text{AS}}(\delta\omega) \propto f(\hbar\omega)g(\hbar\omega)\exp\left(\frac{-\hbar\delta\omega}{k_B T}\right), \tag{5.4}$$

where $g(\hbar\omega)$ represents a density of states and $f(\hbar\omega)$ represents a probability of emission. The use of a Boltzmann factor here disguises the distinction between the electron and phonon thermal baths, and is only justifiable if $\hbar\delta\omega$ is sufficiently large, i.e. far from the laser. In our measurements, the spectral filter blocking out the laser largely obscures the energy range in which this approximation breaks down.

Carattino et al. [121] could make two simplifying assumptions: firstly, since they were using nanorods rather than nanospheres, they could identify $f(\hbar\omega)$ with a sharp plasmon resonance; this approximation is not valid for spheres as their resonance is much broader. Further, they were operating far enough from the main interband transitions of gold to assume $g(\hbar\omega)$ is constant. As the resonance of gold nanospheres is much closer to the interband transitions, this approximation is not valid, either [122]. Their use of the surface plasmon resonance as a normalizing factor can therefore not be replicated here; this is consistent with the lack of apparent sign of the plasmon in the spectra in fig. 5.3.

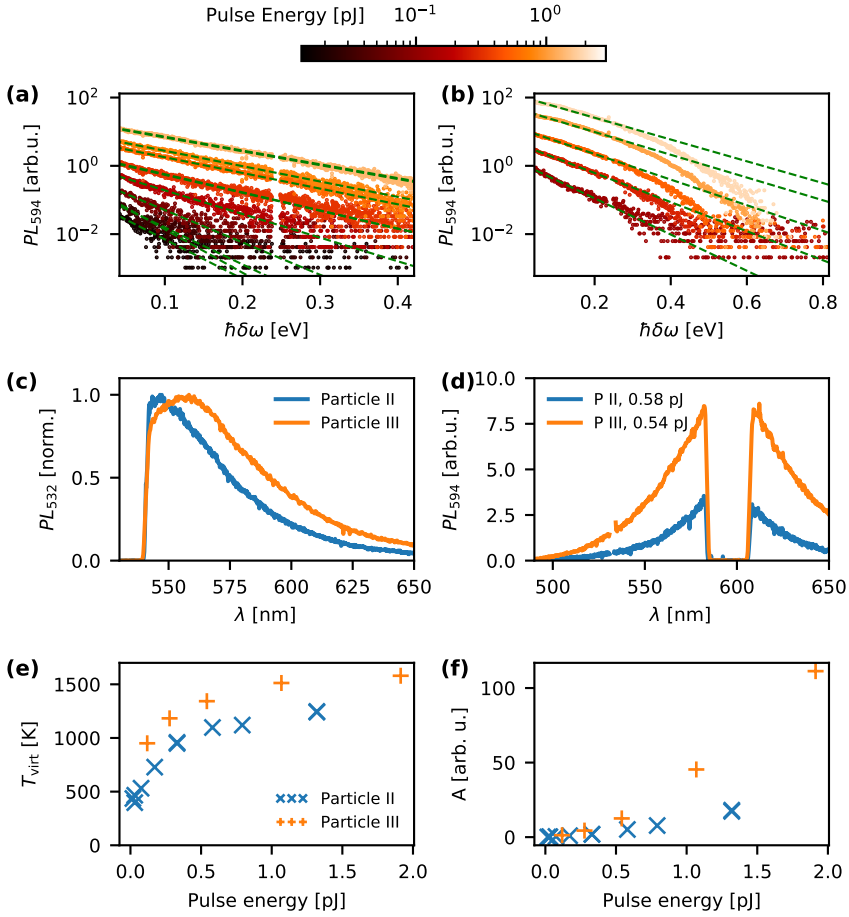

Figure 5.3: Anti-Stokes spectra of particles 'II' **(a)** and 'III' **(b)** excited by 594 nm pulses; different colours represent different pulse energies (colour bar above). Dashed lines are fits to the Boltzmann distribution, eq. (5.5). **(c)** Photoluminescence spectra of the same particles excited at 532 nm (CW). **(d)** Examples of spectra including the Stokes and anti-Stokes components, on a linear scale. **(e)** Temperatures derived from the fits. **(f)** Scaling factors A derived from the fits. The pulse energies quoted are measured in the back focal plane; note that the particles absorb only a fraction of the available energy.

We resort instead to a rather simpler approximation: we assume that the function $f(\hbar\omega)g(\hbar\omega)$ varies much more slowly than the exponential Boltzmann factor. This leaves us with

$$I_{\mathrm{AS}}(\delta\omega) = A \exp\left(\frac{-\hbar\delta\omega}{k_B T}\right), \tag{5.5}$$

where A is a proportionality coefficient.

Due to the lack of a *sharp* resonance, this approximation is reasonable sufficiently far from the interband transitions. The nearby interband transition at 470 nm [123] corresponds to $\hbar\delta\omega \approx 0.5\,\mathrm{eV}$ (i.e. $6 \times 10^3\,\mathrm{K} \times k_B$), at which point we do indeed clearly see the fits to eq. (5.5) break down in fig. 5.3b.

The same approximation has been used successfully by other authors in the past, such as He et al. [124] and Cai et al. [119]. Authors inclined to interpret similar measurements as electronic Raman scattering, such as Crampton et al. [120], instead assume that $f(\hbar\omega)g(\hbar\omega) \propto \omega^3$, which is practically equivalent to our approximation since ω^3 varies much more slowly than the Boltzmann factor.

This simplistic Boltzmann description fits the data well up to $\hbar\delta\omega \approx 0.4\,\mathrm{eV}$, at which point the approximations start to break down as expected. Note, however, that the extracted temperatures reflect the slope of logarithm of the data at lower values of $\delta\omega$ than that, and are thus not particularly affected by the interband transitions.

The temperatures T and amplitudes A, shown in fig. 5.3e and 5.3f respectively, broadly show the expected features:

(i) The temperatures reach >1000 K, and increase monotonically with increased heating power. The high temperatures indicate that we are indeed measuring the temperatures of the electrons, and not of the gold lattice.

(ii) The intensity of the anti-Stokes emission also increases with heating power, both due to the increased number of photons at higher powers, and due to the larger anti-Stokes to Stokes ratio at elevated temperatures. The combination of the two leads to a faster-than-linear increase.

(iii) The rate of temperature change with heating power $\partial T/\partial P$ slows with increased power. Multiple factors contribute to this surprising effect: On the one hand, the heat capacity of the electrons $C_e(T_e)$ increases as the temperature goes up, meaning the same amount of heat corresponds to a smaller temperature increase [105]. On the other hand, the permittivity, and thus the amount of absorbed heat, also depends on the temperature, which has been

shown theoretically to lead to a similar slowing of the change in temperature [125, 126].

(iv) The temperature roughly approaches room temperature as $P \longrightarrow 0$. This is, however, far from exact for two main reasons. Firstly, extrapolating to zero power is not possible without greater knowledge of the nonlinear function $T(P)$, and secondly, the margins of error are substantial: on the one hand, the rather low signal-to-noise ratio at the lowest powers increases the uncertainty as the temperature decreases, and on the other, the approximations inherent in eq. (5.5) limit the reliability of exact quantitative conclusions in the first place.

Evidently, a more detailed model of the emission mechanism as well as the thermodynamic and photothermal effects in this system is needed to accurately calculate and understand the true electronic temperature, and its behaviour as a function of power.

Fig. 5.3 shows a remarkable heterogeneity between the two particles: particle III appears to be heated significantly more efficiently than is particle II, even though the particles have nearly the same size; this implicates other factors that may influence the absorption at 594 nm and elsewhere, such as the precise shape of the particles (including surface facets), interactions with the substrate, or crystal defects.

5.3.2 Hot electron dynamics

In the two-colour experiment, spectra are recorded while focussing two pulse trains on the particles, a probe pulse which causes the measured emission, and a pump pulse, whose direct effect is invisible. The delay between the pulses τ is varied from one spectrum acquisition to the next. τ is calibrated such that at positive τ the probe pulse arrives after the pump pulse, and such that $\tau = 0$ corresponds to the initial maximum of a 'traditional' pump-probe extinction time trace (such as those in chapter 4).

Fig. 5.4 shows such time-resolved spectra for three different 100 nm (nominal) gold nanospheres. The effect of the presence of the pump pulse on the spectra is limited to a period of ~1 ps; at longer delays the emission returns to its initial state. The change on the anti-Stokes side of the spectrum is significantly greater than the change on the Stokes side.

This aligns with our expectations, as the anti-Stokes emission is far more dependent on the temperature than the Stokes emission; changes in Stokes

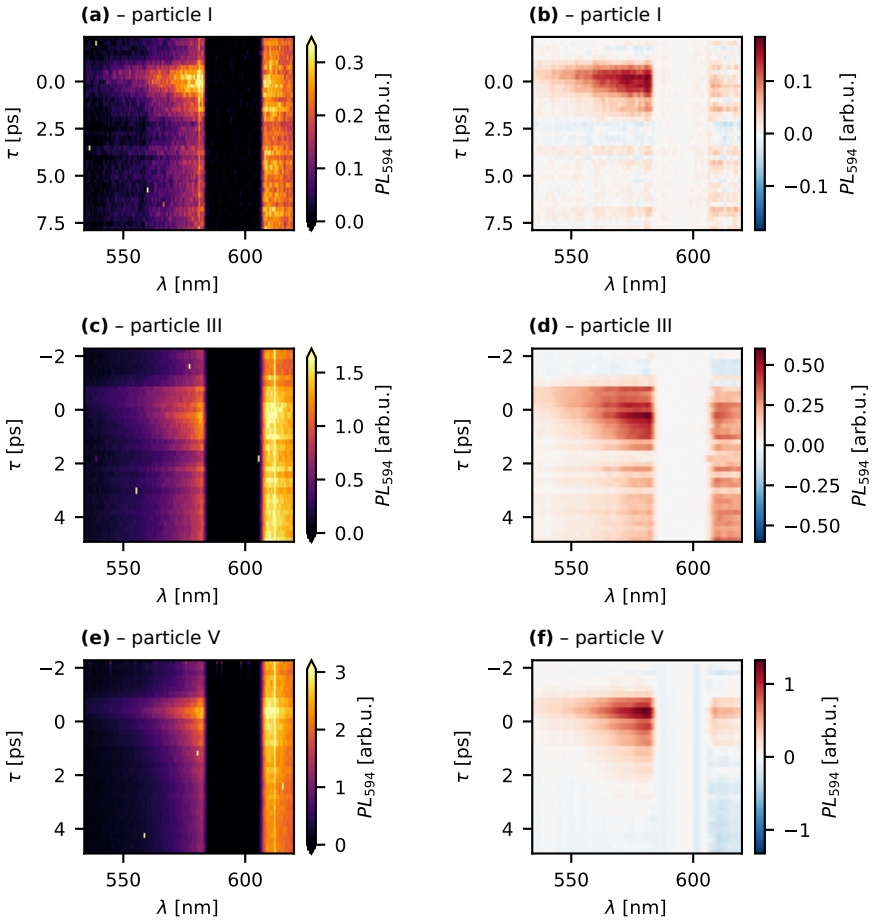

Figure 5.4: Time-resolved spectra of three different 100 nm (nominal) gold nanospheres, as a function of the wavelength λ and pump-probe interpulse delay τ. **(a,c,d)** The raw spectra on the left, **(b,d,f)** spectra with the τ-independent contribution of the probe pulse subtracted on the right, showing only the contribution of the pump pulse. Probe pulse energies were **(a,b)** (85 ± 3) fJ, **(c,d)** (119 ± 11) fJ and **(e,f)** (190 ± 11) fJ. The units on the colour scales are arbitrary but all equivalent to each other. $\tau = 0$ is defined as the peak of the pump-probe extinction spectra as shown in fig. 5.5.

emission due to elevated electron temperatures are due only to secondary effects, such as the change in the dielectric constant.

As stated in § 5.2.1, the measured anti-Stokes spectra are characteristic of the total electron distribution at the moment the probe pulse arrives, as given in eq. (5.3): they are characteristic of $n(t, T, \epsilon)|_{t=\tau}$. Here, two approaches for analysis present themselves:

For a straightforward parametrization, we can simply fit eq. (5.5) to the spectra for each time τ. The equation, now

$$I_{\mathrm{AS}}(\tau, \delta\omega) = A(\tau)\, \exp\left(\frac{-\hbar\delta\omega}{k_B\, T(\tau)}\right), \tag{5.6}$$

immediately gives us two time-dependent parameters that characterize the resulting emission quite well: a temperature $T(\tau)$ and a quasi-amplitude $A(\tau)$.

Alternatively, we can make use of the fact that $[n_0(\epsilon) + \Delta n_{\mathrm{vis}}(t - \tau, \epsilon)]_{t=\tau}$, i.e. the effect of the probe pulse alone, is independent of the inter-pulse delay τ to our advantage. Figs. 5.4b, 5.4d and 5.4f show the difference spectra

$$\Delta I_{\mathrm{AS}}(\tau, \delta\omega) = I_{\mathrm{AS}}(\tau, \delta\omega) - \hat{I}_{\mathrm{AS}}(\delta\omega) \tag{5.7}$$

if we subtract this τ-independent component, the baseline spectrum $\hat{I}_{\mathrm{AS}}(\delta\omega)$. For the purpose of parametrizing the entirety of the data, this approach is less practical since $\Delta I_{\mathrm{AS}}(\tau, \delta\omega)$ is overwhelmed by noise after a picosecond or two.

The parameters $A(\tau)$ and $T(\tau)$ resulting from a fit to eq. (5.6) are shown in fig. 5.5, along with the corresponding pump-probe extinction spectra from which $\tau = 0$ is calibrated. The magnitude of the change in the anti-Stokes spectra, represented by $A(\tau)$ follows the behaviour of the change in extinction well: Both have their maxima at the same inter-pulse delay, and both show the same asymmetric behaviour as a function of τ, with a steep rising edge as the pulse is absorbed, and a slower ~ ps decay as the absorbed energy is released into the metal lattice.

The highest apparent temperature state is, however, reached earlier. Across our measurements, temperature peaked around (0.66 ± 0.10) ps before the amplitude, which is almost twice our pulse width of 350 fs. The high-temperature state then rapidly decays to its initial value as the anti-Stokes intensity increases.

This early high-T, low-A state is easy to see in the spectra directly now that we know what we're looking for: Fig. 5.6 shows difference spectra from fig.

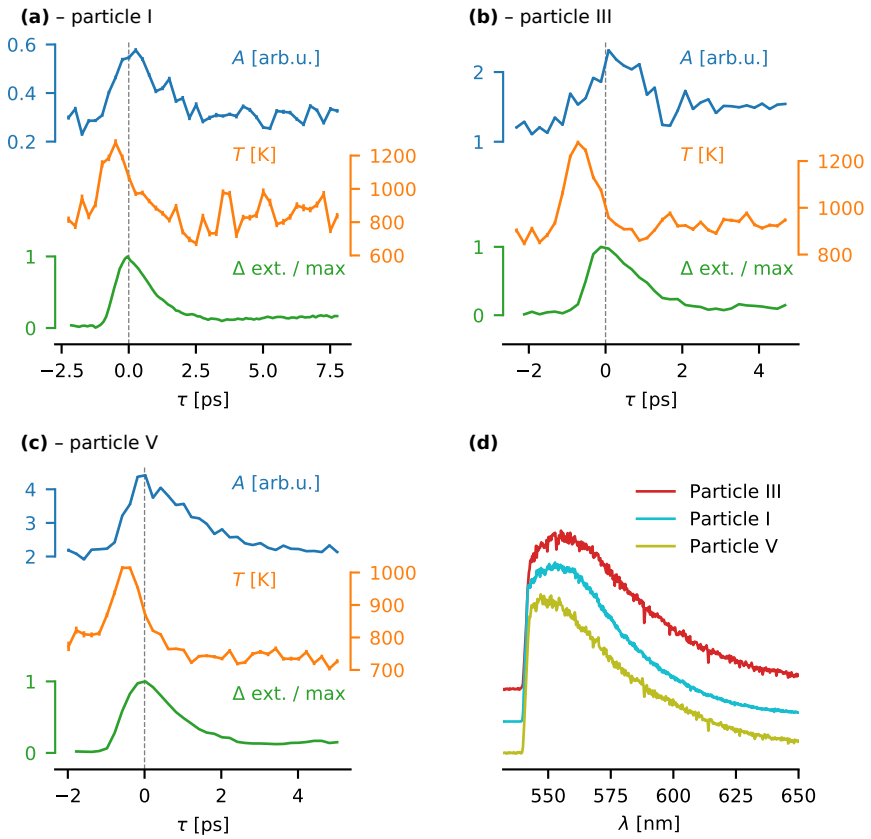

Figure 5.5: (a–c) Parametrization of the spectra in fig. 5.4 in terms of $A(\tau)$ [blue] and $T(\tau)$ [orange] according to eq. (5.6). The units for $A(\tau)$ are arbitrary but correspond to those used in fig. 5.4. Below: pump-probe extinction spectra [green] of the particles for comparison. **(d)** Photoluminescence spectra excited at 532 nm (CW) of the three particles, normalized to have the same maximum intensity and shifted for clarity.

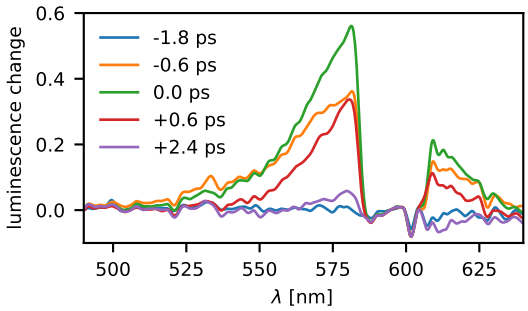

Figure 5.6: Difference spectra from fig. 5.4f: In order of increasing τ, the state before the pump pulse, the highest T state ($\tau = -0.6$ ps), the highest intensity state ($\tau = 0$), a slightly later state ($\tau = 0.6$ ps) and a late-τ state.

5.4f for selected delays τ. The highest-T spectrum, $\tau = -0.6$ ps, clearly decays much more slowly with decreasing λ than does either the tallest spectrum at $\tau = 0$ or indeed any later spectrum, such as the example from $\tau = +0.6$ ps.

The fact that the state of maximum temperature occurs so much earlier than the state of peak response in terms of both extinction and luminescence calls into question, at the very least, the implicit assumption that the peak of the pump-probe extinction spectrum, which we refer to as $\tau = 0$, corresponds to the pump and probe pulses arriving at the same moment.

The temperature changes due to the pump pulse as reported in fig. 5.5, some 300 K or so, appear quite low. This is an artefact caused by the underlying assumption of eq. (5.6): for the fit to give a good understanding of the temperatures, the entire hot electron gas would have to have a well-defined temperature, i.e., would have to be fully thermalized. In terms of eq. (5.3), it would require Δn_{ir} to have equilibrated with Δn_{vis} and n_0 before the emission of any photons.

As suggested above, we can get closer to that picture of the electron distribution by subtracting a baseline spectrum $\hat{I}_{\mathrm{AS}}(\delta\omega)$ from the full spectrum $I_{\mathrm{AS}}(\tau, \delta\omega)$ to arrive at the effect of the pump pulse alone, without the effect of the probe pulse. Just as with the full spectrum, we may then imagine it to be caused by a thermal distribution of electrons[4] and fit the spectra to a

[4]While we expect the initial distribution of hot electrons to be non-thermal [22, 114–116], and we see no *a priori* reason why the emission would be thermal, our data does not clearly show it to be otherwise.

	T from I_{AS}		T_Δ from ΔI_{AS}	
Particle	max	$\tau = 0$	max	$\tau = 0$
I	1278 K	1067 K	1661 K	1206 K
III	1279 K	958 K	1756 K	1004 K
V	1015 K	868 K	1306 K	936 K

Table 5.1: Apparent temperatures of the total spectra (figs. 5.4a, 5.4c, 5.4e and eq. (5.6)) and difference spectra (figs. 5.4b, 5.4d, 5.4f and eq. (5.8)) at maximum temperature and at $\tau = 0$. The values for the former correspond to those in fig. 5.5. Note the spectra for particles III and V were measured on the same day and with the same pump pulse energy, viz. ca. 2.2 pJ as measured in the back focal plane.

Boltzmann distribution:

$$\Delta I_{\mathrm{AS}}(\tau, \delta\omega) = I_{\mathrm{AS}}(\tau, \delta\omega) - \hat{I}_{\mathrm{AS}}(\delta\omega) \stackrel{!}{=} A_\Delta(\tau) \exp\left(\frac{-\hbar\delta\omega}{k_B\, T_\Delta(\tau)}\right). \tag{5.8}$$

The results of this fit for the high-T and high-A states are listed in table 5.1 alongside the corresponding values for the fit of the full spectra from fig. 5.5. It shows that the distributions created by the pump pulse initially have significantly higher characteristic temperatures than those present due to the probe pulse, which decays rapidly, as seen in the previous figures.

5.4 Discussion and conclusion

We have measured anti-Stokes photoluminescence of single gold nanoparticles due to pulsed illumination with characteristic temperatures of order 10^3 K. The temperatures are extracted using a simple Boltzmann approximation (eq. (5.5)) that fits well for intermediate anti-Stokes shifts that are not so small that the Boltzmann distribution would cease to apply, and that are not so large that the interband transitions of gold become significant.

The small-shift limit does not contribute to the measurement as the corresponding light is close enough to the laser to be rejected by our spectral filters. The large-shift limit barely contributes to the fit as the signal in that region is very weak in the first place. The extracted temperatures reflect the spectral region which the simple approximations are best suited to.

For a more exact and more precise extraction of the temperature from the photoluminescence spectra a more thorough model of the origin of gold nanoparticle photoluminescence is needed. Such a model may also provide more insight into the nature of the apparent temperature (i.e., is it entirely due to the electron temperature?). However, such a 'full' model should only amount to a modest correction to the deduced temperature.

In eq. (5.3) we assumed that the excited electron populations created by the two laser pulses are independent of one another. However, in fig. 5.3e, we see that dependence of the anti-Stokes spectra on pulse energy deviates significantly from linearity. This is not, in and of itself, terribly surprising: we know that the heat capacity of the electron gas and the dielectric permittivity both depend on temperature. The probe pulse energies used in the two-colour measurements are on the low side (~0.1 pJ, see fig. 5.4 caption), so the linear approximation inherent in eq. (5.3) may still be reasonable, but we cannot exclude interactions between the pulses that would affect the interpretation of these measurements.

It may be interesting to extend fig. 5.3e to higher pulse energies to establish whether the temperature saturates under strong illumination before melting of the particle becomes an issue.

We have shown in § 5.3.2 that the apparent temperature increase due to pulsed excitation decays on a ~ps timescale, which agrees with previous measurements of the electron-phonon coupling time. Surprisingly, the peak apparent temperature is reached early in the process (fig. 5.6), 0.6 ps before the peak amplitude. This time consistent with the thermalization times previously measured in bulk gold using time-resolved photoemission spectroscopy [111]. We note that what appears to be an early high-temperature state may in fact be a non-thermal state with a 'hot electron' contribution (which we measure) and a lower-temperature contribution (which is mostly obscured by the spectral filter near the laser).

This surprising observation may allow some deeper insight into the thermalization dynamics of the electron gas in a gold nanoparticle during the arrival of short laser pulses, if it were examined using a detailed model of the electron distribution and its evolution as a function of time.

All our results appear, however, to be consistent with a thermal distribution of electrons. We expect this is due to two main factors: Firstly, the probe pulse width is much longer than the electron scattering time (viz. <10 fs [22]) and of the same order of magnitude as the thermalization time. In short, there is

plenty of time during the probe pulse for the electron distribution to approach a thermal one. Secondly, the spectral filter removing the laser obscures small energy shifts, which means we can't compare the distributions at low and at high energies as extensively as we might like.

Either one of these shortcomings could be improved somewhat, but there is a trade-off between the two: If one wanted to work with a shorter pulse, this pulse would inevitably have a larger bandwidth (at least 0.2 eV for ~10 fs, for instance), which would obscure more of the anti-Stokes spectrum. Vice versa, a much narrower filter would soon require the use of longer laser pulses.

Rather, to measure electron temperature dynamics using anti-Stokes emission, our measurements have to be integrated over a longer period of time (as they are here), and we must rely on theory to clear up the details.