Raman Scattering Lab

Ben Nachumi

November 18, 2014

1 What is Raman Scattering?

Raman scattering is the inelastic scattering of light from vibrational or other excitations in a piece of matter. An incident photon excites (or absorbs) a phonon (or magnon, or...) and comes out a different color. This scattering occurs over a very short timescale, and is a coherent process over virtual intermediate states. This means that Raman scattered light is shifted from and tied to the incident photon frequency. If you illuminate a piece of matter with monochromatic light and analyze the light that comes out, you will detect new colors in the collected spectrum corresponding to the Raman-excitable modes of the sample. The *Raman shift* in the frequency of the collected light, measured from the incoming photon's frequency, corresponds to how much energy the incident photons deposit (or take from) the sample. [1]

Raman scattering can produce excitations when the combinaton of $\langle \vec{out}|$ (scattered) and $|\vec{in}\rangle$ (incident) polarizations of the photon field match (transform like) the symmetry of the excitation in the sample. Given a molecular or crystal structure, you can predict which excitations you ought to see, depending on how you select the polarization of the incoming and outgoing light. If you're interested, I expand upon this mathematics in section 5.

In a research context, a Raman experiment tells you stuff about the excitations in your sample–their shapes and energies. Add temperature dependence and you can often pick up evidence of structural and electronic phase transitions for much less than the cost of a neutron scattering or X-ray spectrometer. Raman scattering is widely used in chemistry for the rapid detection of various chemical groups from their signature vibrational frequencies. From the larger philosophical perspective, Raman scattering presents an accessible exploration of the consequences of symmetry in the interaction of light with matter.

---

[1] By contrast, in fluorescence (for example), the intermediate excited state of the sample is long lived and light of a different color comes out much later, as the relatively stable intermediate state decays. In fluorescence, the visible emission comes from a transition between two well-defined energy levels of the sample, and so is fixed in frequency.

2 The Big Pieces

Suppose you have a sample that should support Raman-excitable modes. To see Raman scattering, you need monochromatic incident light–a laser. In fact, to confirm that you are looking at a *Raman shift* from the laser line, you really ought to try to see the same peaks with the same shift from different laser colors. You need a dark room, a reasonable spectrometer, a quiet detector, and a rejection filter for your excitation, because as a higher-order process, Raman scattering is much weaker than the elastic scattering of the excitation.

The spectrometer is a research-grade single-grating Czerny-Turner (slit-mirror-grating-mirror-out) spectrometer (JY HR-640) that fell upon some hard times. It's been fixed up so that the grating moves as it should. There are lasers of at least two colors for excitation, and enough filters, lenses, and mirrors to make a good transfer of laser light to your sample, to focus the scattered light on the entrance slit of the spectrometer. For detection, there are, as of this writing, two usable low-dark-current CCDs. One is a liquid nitrogen-cooled camera (Photometrics CH250) and the other is a Peltier-cooled camera (SBIG ST-8300M).

3 Follow the Light Path!

3.1 Beam Transport: How Light gets to the Sample

A source of more-or-less monochromatic light is cleaned by bouncing it off a preliminary grating, then sent (via mirrors) over to the sample. Along the way it passes through one or more lenses and polarizers, so that it arrives at the sample in a custom-sized, custom-polarized spot. Note that the shorter the focal length and the larger the beam on the lens surface, the smaller the focused spot.

Now recall that you don't really want to see the laser light–you want the other colors emitted when the laser light is shifted. This is a much weaker signal; in fact **even a weak-looking, less than 1 mW direct laser beam quickly oversaturates, and can even damage, the detection CCD's**. You almost never want the specular beam, and you certainly don't want it at full strength. You want the laser to hit the sample such that the spectral (mirror-reflected) laser beam misses the collection lens. To hit the sample around the collection lens while the specular reflection escapes, you can bounce the laser beam off of one or more mirrors.

3.2 How Light Gets into the Spectrometer

Next, you want to be able to get the light from the sample into the spectrometer. The usual way to do this is with at least two lenses. The first lens collects light from the laser spot. The second lens refocuses the collected light on the entrance slit. For higher

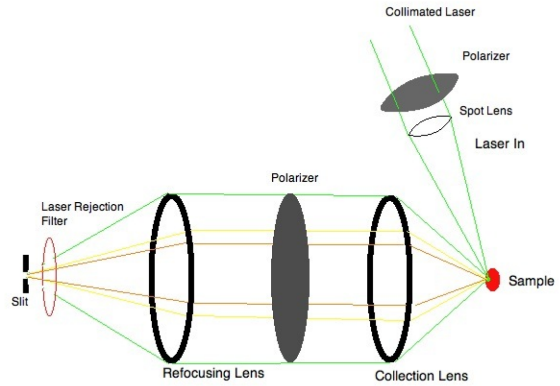

Figure 1: A simple input train.

signal, you want to collect lots of light. So for a collection lens, you should use a nice fast (large diameter/focal length, or low f-number) lens, and you want the sample as close to the lens as possible. A simple way to do this and make a real image is to put the sample at the focus of the fast collection lens. In ray optics, this configuration would turn an ideal point into a parallel beam. From this parallel beam, the refocusing lens produces a point at its focus–so the refocusing lens should have its focus right on the entrance slits.

While the basic idea is simple, there's a lot to think about. Call the focal length of the collection lens F, and that of the refocusing lens F'. If the size of the spot on the sample is r, then the size of the spot on the slits is $F'/F \times r$. You want the spot on the slit to be as small as possible so that, while collecting full signal, the slit can made as narrow as possible.

On the other hand, the spectrometer itself has an f-number: the size of its grating/the focal length. This is larger than the f-number of the collection lens. So as not to throw collected light away and to maximize your use of the grating, you want the refocusing lens to be as big as the collection lens, and yet to have the same f-number as the spectrometer. This means that there will have to be some magnification of the image spot at the slits, which means that the slits have to be bigger to admit the entire image spot. But widening the slit reduces resolution.

To avoid that penalty, you have to make the spot on the sample as small as possible. That means the input lens should have as short a focal length as possible, and should be as fully illuminated as possible. So the laser beam ought to be expanded on its way to the input lens. This means at least two lenses before the input lens. But because of intrinsic inefficiencies, such as back reflections from the lenses themselves, the more lenses there are in the input train, the less laser light gets to the sample. For an undergraduate experiment like this, safety considerations limit the laser power. So the actual excitation and input train is something of a compromise.

3.3 How to Get a Good Image on the Slit

Go through these steps to put a focused image on the entrance slits.

- Turn on the little focusing camera and the monitor, and swing around the flip mirror. This mirror sends light from the slits through an internal lens that makes a real image of the slit on the camera. Check that the focus camera is really imaging the front (vertical) slits–these determine the spectral resolution. The slits are focused if you see their sharp edges move around on the monitor when you adjust the slit width.

- Next, open the slits and center a refocusing lens in front of the shutter. By moving the lens forward and back, try to focus the image of the most distant object you can see in the monitor. The focal point of this lens is now just a bit in front of the slits (after all, the object isn't infinitely far away), but this gets you near where you eventually want to be.

- Now center the collection lens on the refocusing lens. Leave some space between them so you can insert a polarizer or a filter later. Move a needle or some other tiny pointy thing around until you see its focused image appear in the monitor along with the slits.

- The image of the needle is now falling on the slits, but you can improve things a bit by just touching the refocusing lens back toward the slits, and finding the image of the needle as before. If you can't get a focus, that means that the refocusing lens is too close to the slits. Move it a bit forward from them and try again.

- Now turn on the laser and steer the unfocused beam so you see it hit the needle in the monitor. Finally, focus the beam into a spot on the needle.

Your sample goes where the needle is, at the focus of the collection lens. When you swap in your sample, you'll have to tweak the beam steering and the sample position to get everything just right. Eventually, you'll want to close down the slits while keeping the spot visible.

3.4 Inside the Spectrometer

To begin to understand how the spectrometer works, refer to the following figure. The

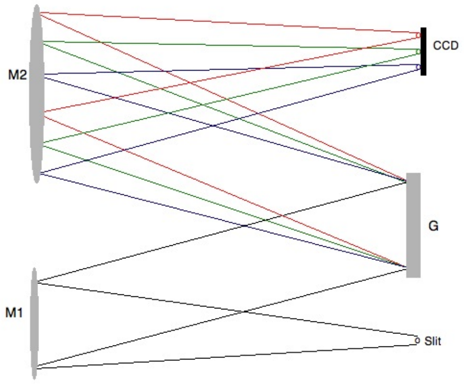

Figure 2: The light path inside the spectrometer.

basic idea is to make separate copies (one in each color present) of the illuminated entrance slit, and put them on a detector. Now the slit is in the focal plane of the concave mirror (**M1**), so light from the slit that hits **M1** gets sent parallel, to the grating (**G**). The grating splits the multicolored parallel beam into a fan of copies. Those copies that hit the second concave mirror (**M2**) are refocused onto the plane of the detector. So for each color that impinges on **M2**, you get a separate spot on the detector. Which section of spectrum you see depends on how the grating is turned, and on the pitch (grooves per mm) of **G**. The more densely grooved is **G**, the more rapidly the spectrum fans out, while **M2** of course stays the same size. So for a fixed grating angle, you'll catch less of the spectrum (but have better resolution) with a finer grating.

4 Making Sense of the Spectrum

4.1 The Central Wavelength

For all of this to be useful, you have to know how position on the detector corresponds to color. The task is to calculate this spectral range from the dimensions of the spectrometer and from the formulas for diffraction from a grating. You'll also learn how this particular commercial spectrometer associates a central wavelength wth a particular grating position. First, consider the undercarriage of the spectrometer. You'll see a big drive screw. Upon this screw rides a shuttle, which in turn moves a long arm. For reasons that should soon become obvious, this arm is called the *sine bar*. Upon the sine bar rides a small wheel, attached to another arm, which appears to turn on a hinge that

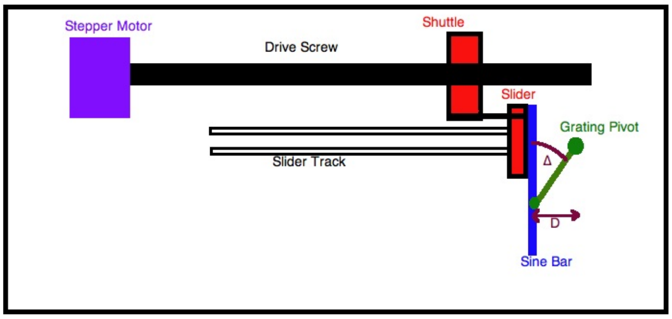

Figure 3: Schematic of the drive mechanism.

goes up into the spectrometer. In fact, this hinge turns the grating. When the screw turns, the sine bar rides up or down parallel to its axis, which pushes the arm that turns the grating.

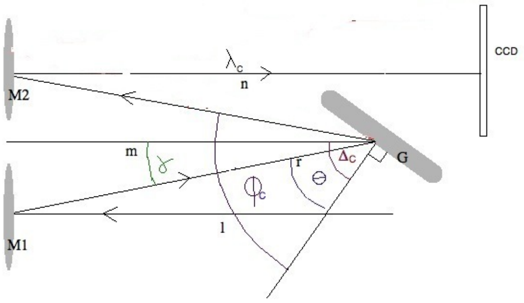

Figure 4: Determining λ_c.

The next diagram shows how the grating angle determines the central wavelength. The lines l, m, and n are mutually parallel. l connects the center of the slit to the center of **M1**, m goes through the center of **G**, and the normal to the grating makes an angle (the grating angle) Δ with m. n passes through the center of the exit port. Also note for your later convenience the line r, which connects the centers of **M1** and **G**, and makes the angle γ with m.

The mirrors **M1** and **M2** are positioned so that when $\Delta = 0$, the specular reflection of light moving parallel to r from the grating of light from **M1** is ultimately sent out the center of the exit port (via **M2**) along n. For any position of the grating, one color of light will still follow this path. The color of the light that follws this path is called the *central wavelength* λ_c. For a perfectly calibrated system, this wavelength would correspond to the readout of the odometer on the side of the spectrometer which converts turns of the screw directly to λ_c.

This works because the grating inside the spectrometer is parallel to the arm that rides along the sine bar, and the sine bar moves along the main screw (see the undercarriage diagram), which makes the distance D along the screw proportional to $\sin\Delta$. Here's the proof:

Let θ be the angle between the normal to **G** and r, so that $\theta = \Delta - \gamma$. Call the overall deviation between the ray approaching **G** and its outgoing partner ϕ. For the central wavelength λ_c, $\phi_c = \Delta + \gamma$. If the spacing between neighboring grooves of the grating is ϵ, then the path length difference between two rays reflecting from neighboring grooves will be $\epsilon(\sin\phi + \sin\theta)$, and you'll see constructive interference along the λ_c's path for

$$x\lambda_c = \epsilon(\sin\phi + \sin\theta)$$

(x, the order of the interference maximum, is an integer). Since $\phi = \Delta + \gamma$ and $\theta = \Delta - \gamma$, this becomes

$$x\lambda_c = 2\epsilon \sin\Delta \cos\gamma.$$

Since γ is fixed, $\lambda_c \propto \sin\Delta \propto D$. So if you set the initial grating angle to zero, turns of the screw are directly proportional to λ_c.

By drawing a similar diagram, you can show that this same formula applies when the grating is turned so that the normal to **G** is on the other side of r and $x = -1$. As another exercise, you can calculate Δ_{532}, the angle appropriate for $\lambda_c = 5320\mathring{A}$, for a 1200 groove/mm grating (for the HR640, use $\gamma = 7.67936°$).

4.2 Beyond the Central Wavelength: Calculating Dispersion

If this were **BCCD** (**B**efore **CCD**'s–olden times), you would be taking your data by driving the grating through some angle, and recording the current readout of a photomultiplier tube sitting on the center of an exit slit. From the previous characterization of the tube, you would calculate the intensity for each specific grating angle and so reconstruct the spectrum point by point. In yet older times, you'd be recording a broad spectrum in one go on a photographic plate–but this would require you to chemically develop the image, and would also be very difficult to calibrate. The CCD is a pixel detector array. It acts like a nearly instantly readable film, with a very linear intensity versus illumination–so it combines the best of both older methods. However to understand the CCD spectrum, you need to modify the previous discussion for colors that are sent off of the central path.

This isn't too difficult to do. In the next diagram, I've added a new color λ_ν, coming off **G** at the angle ϕ_ν from **G**'s normal, with $\phi_\nu = \phi_c + \nu$. Going through essentially the same steps, you can show that $x\lambda_\nu = 2\epsilon \sin(\Delta + \nu/2)\cos(\gamma + \nu/2)$.

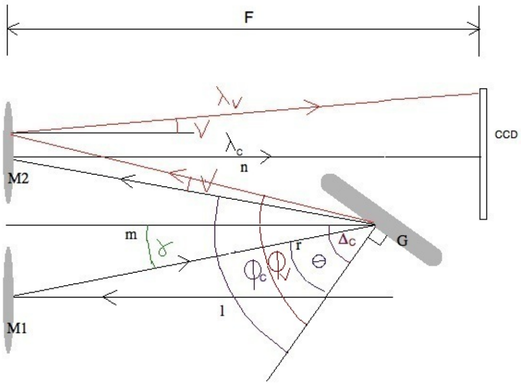

Figure 5: Dispersion geometry.

This extra angular deviation just carries through the rest of the system. Coming off **M2**, λ_ν will be deflected from the central path of λ_c by the angle ν. Call the focal length of **M2** F. In the focal plane of **M2**, at the CCD, λ_ν will make a bright spot a distance $y = F \tan \nu$ from the spot made by λ_c.

Using the dimensions of the CCD and the known Δ_c, you can determine the largest angluar deviation $\bar{\nu}$ from the center of the CCD. The complete range of wavelengths that you can see in one shot is then

$$\lambda_{Max} - \lambda_{Min} = 2\epsilon \sin \bar{\nu} \cos(\Delta + \gamma).$$

Since $\epsilon/F \ll 1$, $\tan\nu \approx \nu$, so the *linear dispersion* in units of (wave)length/length along the CCD is

$$\frac{d\lambda_\nu}{dy} \approx \frac{\epsilon}{F}\cos(\Delta+\gamma+\nu/2) \approx \frac{\epsilon}{F}\cos(\Delta+\gamma) \approx \frac{\epsilon}{F}\cos\Delta.$$

The penultimate, ν-independent expression is usually good as the width of the CCD is always much less than F, and if you are looking at visible light in a long focal-length instrument with a finely-pitched grating, you can also ignore γ in comparison with Δ. As an exercise, find the wavelength coverage that you expect for a 1200 groove/mm grating with $\lambda_c = 5320\mathring{A}$ (use 24.4 mm for the width of the CCD chip, 640mm for F, and $\gamma = 7.67936°$). Dividing the range by the width of the CCD, get the average linear dispersion for this setting. Compare this to the ν-independent and γ-independent formulas. Generate a plot of the true linear dispersion as a function of position (from -12.2 mm to +12.2 mm) across the CCD. On the same axes, plot the % deviation from the average linear dispersion.

4.3 Resolution

Colloquially speaking, the *resolution* is the smallest separation between two bright spots that you can just barely distinguish with your CCD. The spectral resolution is the corresponding minimal difference in λ. For a working definition of "barely distinguishable", the common choice is the *Rayleigh criterion*: two peaks are barely distinguishable if the maximum of one falls on the minimum of the other. To apply it, you need to know the width of the bright spot–the distance along the CCD from maximum intensity to zero. For a given grating of line spacing ϵ, with N lines illuminated, with a bright beam going off at at angle Θ the angular width of the xth maximum is

$$\delta(\sin\Theta) = \frac{\lambda}{N\epsilon}.$$

This just describes the change in angle necessary to get a phase shift of π over half the width of the grating, so that reflections in lines cancel in pairs (as in the standard analysis for diffraction from a single slit). On the other hand, from the diffraction maximum condition $\epsilon\sin\Theta = x\lambda$ (x integer), the next distinguishable peak, $\delta\lambda$ away, will deviate from Θ by $\delta(\sin\Theta) = \frac{x\delta\lambda}{\epsilon}$. Setting these two $\delta(\sin\Theta)$'s equal gives the *grating resolution*

$$\frac{\delta\lambda}{\lambda} = \frac{1}{Nx}.$$

Note that the resolution improves as the illuminated area of the grating increases. Using the dispersion to convert this information to a width on the CCD is straightforward. To find the spot size δs, just divide the wavelength width of the peak by the dispersion; you'll get

$$\delta s = \frac{\lambda}{Nx} \times \frac{F}{\epsilon\cos\Delta} = \frac{\lambda F}{xG\cos\Delta},$$

where $G = N\epsilon$ is the illuminated width of the grating. For a first-order ($x = 1$) peak, with a typical $\lambda = 5 \times 10^{-7}$m, with $F = 0.64$m, and assuming the entire grating illuminated with $G = 0.1$m, the intrinsic spot size (that is, the spot size just due to the grating resolution) will be $3 \times 10^{-6}/\cos\Delta$m.

When the pixel size on the CCD is larger than the intrinsic spot size, the effective smallest possible spot size is just the pixel size. However the width of the entrance slit, which is almost always larger than the size of the pixel, usually controls the *experimental resolution.* For example, if the pixel size is 4μm, and the intrinsic spot size is 1μm, and the slit itself is 100μm, then the slit spot size will be about about 25 pixels [2]. Hence for a large range of slit values, you can predict your experimental resolution pretty well by multiplying the dispersion by the slit width.

4.4 CCD Wavelength Calibration and Binning

The wavelength readout on the side of the spectrometer is not as reliable as you might like it to be. To really know what you are looking at, you can take a reference spectrum from a known source and use that to calibrate your CCD.

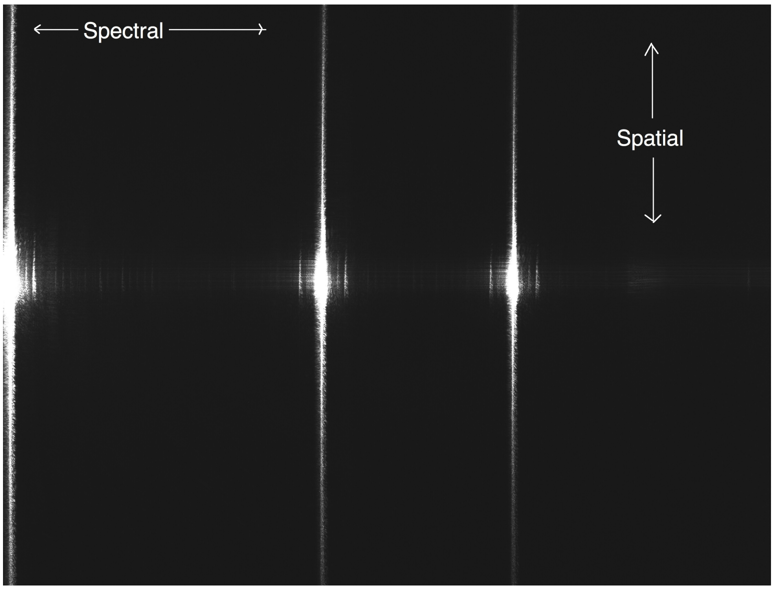

Figure 6: A reference spectrum from the unfiltered emission of the multiline argon laser. Visible are the $4879\AA, 4972\AA$, and $5015\AA$ lines.

If you are taking a spectrum for a weak signal, it's a good idea to bin the CCD in the spatial (vertical) direction. Binning combines several pixels into a single superpixel. You generally want to bin quite coarsely in the spatial direction, because you are treating all

---

[2] for $F = 640$mm and $\lambda = 5000\AA$, can you ignore diffraction from the slit?

illuminated parts of your sample as equivalent. You can also enhance signal (at the cost of resolution) by binning in the spectral (horizontal) direction. The next figure is a CCD image from a piece of silicon wafer, binned by a factor of 128 spatially and spectrally by a factor of 3. (Granted, the spectrum in the binned image is a bit difficult to make out–but a plot of intensity versus bin number makes things obvious.)

Figure 7: A binned CCD image from a piece of silicon wafer. There are three bright spots: two laser lines and one phonon. Note that to avoid saturating and possibly damaging the CCD in this longer exposure, the $4879\mathring{A}$ excitation has, by a slight turn of the grating, been moved out of the frame. The two remaining laser lines are residue or "contamination" of the filtered beam.

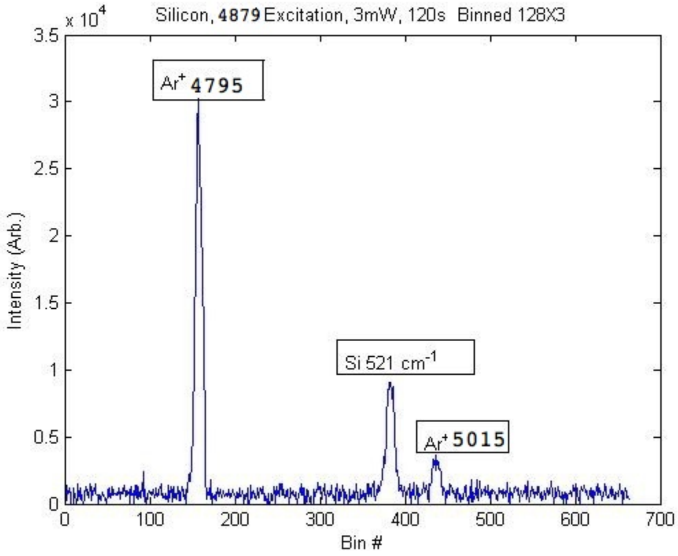

Figure 8: A histogram of the charge in each CCD bin for the center line of the previous image. The two residual Ar^+ laser lines serve as a simultaneous calibration of the dispersion.

4.5 Plotting the Raman Shift

Your data comes out naturally as a histogram of intensity or number of photons versus the absolute wavelengths of the collected light. But the deposited energy is the fundamental quantity of interest. The energy of any photon is proportional to its frequency, so you really want to plot the number of scattered photons versus the Raman shift. The peaks in this latter plot will correspond to the energies of the excitations in the sample.

In spectroscopy, the usual unit of frequency is the *inverse centimeter* (cm^{-1}), which is just the usual frequency (in s^{-1}) divided by the speed of light in cm/s. The conversion is $1\mathrm{cm}^{-1} \leftrightarrow 30\mathrm{GHz}$. With $\lambda_c = 5000\mathring{A}$, the spectral coverage of the spectrometer with the 1200 groove/mm grating installed, is about 1000 cm^{-1}.

A positive energy deposited in the sample means that you collect less energetic (longer λ) photons. So to convert your histogram from counts versus wavelength to counts versus Raman shift, you

1. Convert the laser wavelength into absolute cm^{-1}. For example, the $\lambda = 4879\mathring{A}$ line of the argon laser is the $f = 20496\mathrm{cm}^{-1}$ excitation frquency.
2. Convert your CCD spectrum bin numbers to wavelengths using the calculated linear dispersion.
3. Convert your wavelength scale to absolute cm^{-1}.
4. Convert your absolute cm^{-1} scale to Raman shift by subtracting off the laser frequency.

You can easily convince yourself that the intrinsic resolution of the grating is unchanged by the conversion to frequency units: $\delta f/f = \delta\lambda/\lambda$. On the other hand, the dispersion does change. When you convert dispersion from wavelength/mm to cm^{-1}/mm, the transformation is nonlinear. Since $\lambda = 1/f$, $\delta f = -\delta\lambda/\lambda^2$, and a constant linear dispersion $d\lambda/dx$ in wavelength units becomes

$$\frac{df}{dx} = \frac{-1}{\lambda^2}\frac{d\lambda}{dx}$$

when converted to frequency units. Since the pixel size stays constant, the frequency increment per pixel decreases with increasing λ.

5 Origin and Use of Raman Selection Rules

This is an introduction to the symmetry aspects of Raman scattering. It assumes a notional and notational familiarity with transition amplitudes in quantum mechanics. For the sake of definiteness, this treatment assumes that the scattering occurs from phonons–atomic vibrational modes–although the same ideas apply to any sort of excitation in any kind of sample.

The term "selection rule" is shorthand for the set of conditions under which a given transition amplitude is switched on or off. A Raman scattering experiment determines the polarization selection rules for light-scattering from a sample, in the event that the incident and scattered light have different frequencies. These selection rules depend on how the excitation in the sample changes under the symmetry operations of the experiment.

5.1 Simple Examples of Selection Rules: the Polarizer and Parity

Consider the absorption of light by a polarizer. The polarizer seems to split the light into two parts, one with the electric oscillation $\vec{E}$ perpendicular, the other parallel, to the direction of the little absorbing wires in the polarizer. Current density $\vec{J}$ is induced along the wires. In the usual parlance, we say that the power absorbed is proportional to $\vec{J} \cdot \vec{E}$, and this dot product is zero for $\vec{E} \perp \vec{J}$.

Here's how you could use symmetry arguments to analyze the same phenomenon. Imagine that you know that light has a polarization, and that the absorption of light in a polarizer has something to do with the induced current $\vec{J}$ in the little wires. Suppose that the polarizer lies in the xy plane, with the little wires in the $+x$ direction, and that $\vec{E}$ points wherever it wants in the xy plane. If absorption occurs, how should $\vec{J}$ go?

Suppose that we were to rotate the whole experiment ($\vec{E}$ plus wires), around some arbitrary axis. The underlying assumption of any symmetry argument is the principle of *objectivity*: that under a symmetry operation, nothing measurable should change, because the outcome of an experiment shouldn't depend on our choice of coordinates. Therefore the power developed by the interaction of $\vec{J}$ and $\vec{E}$ must be a scalar function of $\vec{J}$ and $\vec{E}$. Next, suppose we twirl just the polarizer by π around the z axis, while keeping $\vec{E}$ the same, and then repeat the experiment. If $\vec{J}$ followed the polarizer, this would introduce a measurable angle between $\vec{J}$ and $\vec{E}$. But unless there is a preferred orientation for our wires, that can't be. In a theoretical picture using ideal wires, this would violate the principle of objectivity.

As a consequence of our assumption about the objectivity of the experiment, we conclude that, for the polarizer to affect the light, $\vec{J}$ must be able to follow $\vec{E}$. And for ideal wires, that can happen only to the extent that the wires and $\vec{E}$ are parallel–which leads to the usual rules for how ideal polarizers work.

Perhaps the simplest example of a selection rule in quantum mechanics comes from parity. The parity operation $\mathbf{P}$ flips spatial coordinates: $\mathbf{P}((x,y,z)) = (-x,-y,-z)$. For wavefunctions, the parity operation flips the arguments $\mathbf{P}(\psi(x,y,z)) = \psi(-x,-y,-z)$. Clearly, $\mathbf{P}^2(\psi(x,y,z)) = \psi(x,y,z)$, so that $\mathbf{P}$ has eigenvalues of ± 1, and an arbitrary state $|\phi\rangle$ can be broken up into pieces that are “even” (eigenvalue 1) or “odd” (eigenvalue -1) under $\mathbf{P}$.

In an introductory quantum mechanics class, one often sees the matrix element of an operator like the electric dipole operator, $\mathcal{D} = e\vec{r}$, between two parity-labeled states: $\langle a'|\mathcal{D}|a\rangle$. The question is: which pairs of states can be connected through $\mathcal{D}$, and which cannot?

$\mathcal{D}$ is odd under the parity operation ($\mathbf{P}(\mathcal{D}) = -\mathcal{D}$). Meanwhile, $\langle a'|\mathcal{D}|a\rangle$ represents a physical process with a classical analogue. We assume the objectivity of $\langle a'|\mathcal{D}|a\rangle$, so it *shouldn’t care* whether we number our axes left-to-right or right-to-left:

$$\mathbf{P}(\langle a'|\mathcal{D}|a\rangle) = +\langle a'|\mathcal{D}|a\rangle.$$

Hence the electric dipole operator can only connect states of opposite parity–one changes sign, and one doesn’t, so that the overall sign of the matrix element doesn’t change under $\mathbf{P}$. This is the “parity selection rule” for the electric dipole operator.

In a slightly more advanced class, one considers the combined transformations of the original object and the electric field. The states of the world are then $|a\rangle|\gamma\rangle$, where γ stands for the field coordinates, and it seems only fair to use $e\vec{E} \cdot \vec{r}$ for $\mathcal{D}$. Now $\mathbf{P}(\mathcal{D}) = +\mathcal{D}$, as befits a dot product of two vectors. But the selection rule follows anyway because now in the statement of objectivity

$$\mathbf{P}(\langle \gamma \pm 1|\langle a'|\mathcal{D}|a\rangle|\gamma\rangle) = +\langle \gamma \pm 1|\langle a'|\mathcal{D}|a\rangle|\gamma\rangle,$$

$\mathbf{P}(\langle \gamma \pm 1|...|\gamma\rangle)$ introduces a phase of -1, which $\mathbf{P}(\langle a'|...|a\rangle)$ must match.

5.2 A Slightly More Advanced Picture of Raman Scattering

Consider the following matrix element as a representative Raman scattering event:

$$\langle \mathbf{OUT}|\mathcal{S}|\mathbf{IN}\rangle = \langle n'\gamma'|\mathcal{S}|n\gamma\rangle.$$

This matrix element corresponds to the absorption of a photon (a γ) by the sample, an internal transition of the phonon field (the n’s), and the reemission of a (different frequency) photon. The scattering operator $\mathcal{S}$ represents the effect of the sample. The number of photons before and after the scattering is the same, but the disposition of photons changes. On the other hand, after the scattering, the number of phonons or other excitations in the sample has changed. By examining the way the different parts of the matrix element (phonons, light, $\mathcal{S}$ operator) transform under the symmetry operations of the experiment, we'll see how selection rules work for Raman scattering.

5.3 The Raman Tensor

Think of $\mathcal{S}$ as the product of two operators, one that makes transitions of the sample, and one that takes care of the photon fields:

$$\mathcal{S} = \mathcal{S}_{\{n\}} \otimes \mathcal{S}_{\{\gamma\}}.$$

$\mathcal{S}_{\{\gamma\}}$ must be invariant with respect to general three-dimensional rotations and inversion (because Maxwell's equations are invariant under general rotations, translations, and spatial inversion). $\mathcal{S}_{\{n\}}$ has the symmetry of the sample, that is, it is unchanged under the rotations, reflections, etc., that leave the idealized lattice the same. All of these are part of the larger group for Maxwell's equations. Thus the total $\mathcal{S}$ always has a phase of $+1$ after transformations that leave the sample invariant.

Suppose now that we do a scattering experiment, and then repeat the same experiment after rotating the sample through the symmetry operations of the sample's group, while keeping the rest of the apparatus (laser polarizations, lab room) fixed. We assume that the result of the experiment should not change: the sample should behave the same way when considered in any equivalent orientation. In other words, we assume the sample should behave "objectively". In comparing the two experiments we are effectively rotating the part matrix element/part operator ("**PMPO**")

$$\langle n'|\mathcal{S}|n\rangle = \langle n'|\mathcal{S}_{\{n\}}|n\rangle \mathcal{S}_{\{\gamma\}}$$

against a fixed background

$$\langle \gamma'|...|\gamma\rangle.$$

Since we are assuming that the $\mathcal{S}$ operators don't change, the sandwiching phonon states will have to change in a way that, after the transformation, they again match the outer photon sandwich.

How do phonons transform? Phonons are identified by their pattern of atomic displacements, which are periodic in time. So, for example, a phonon consisting of just an upward atomic displacement will acquire a phase of -1 under a rotation of π about an axis perpendicular to the displacement:

$$\Bigg\Uparrow = -1 \times \Bigg\Downarrow.$$

Under a rotation of $\pi/2$ around the y axis, a displacement in the $+z$ direction will turn into a displacement in the $+x$ direction. Hence it is clear that our **PMPO** in fact has components. What is it? A vector?

Make the assumption of overall objectivity: if we we rotate the entire world of the experiment, crystal plus apparatus, corresponding to the entire matrix element, then there can be no change at all, not even a phase. The change of the sandwiching $\langle\gamma'|...|\gamma\rangle$ of the photon states must exactly compensate for the change we get by rotating the crystal alone, e.g., the change due to the phonons in the **PMPO**. But photon states can be described by their vector polarizations, and vectors transform like vectors, so the **PMPO** connecting them must be a rank-two tensor.

To emphasize: In order to connect two vectors into an invariant amplitude, the **PMPO** must be a rank-two tensor. This **PMPO** tensor is called the *Raman tensor* (for that particular phonon), $\mathcal{R}(ph)$. The three possible values (x, y or z) of the two tensor indices (out/row, in/column) of $\mathcal{R}(ph)$ match up with the polarization directions of the incoming and outgoing light. What you see in an experiment is the total Raman tensor, $\mathcal{R}$, which is a sum of the $\mathcal{R}(ph)$'s and similar things for other kinds of modes that can live in a sample. In a typical experiment, one controls the polarization of the incident light, and selects different polarizations of the scattered light. In this manner one probes different components of $\mathcal{R}$.

5.4 Displacement Patterns and Polarization Selection Rules

Suppose you are interested in a particular Raman phonon. From the sample structure, you have guessed at the displacement pattern of the atoms involved in the phonon. You can use this information to predict which incoming/outgoing combination of light polarizations could create one of these. Pick in/out $\langle\gamma'|...|\gamma\rangle$ polarization conditions, and note how those change. If the phases these pick up match the phase the displacement pattern picks up, then the phonon might be seen in those polarization conditions. And if they don't match, then in order to have a chance to see it, you need to rotate your polarizers. This is a Raman selection rule.

Now consider again the actions of the sample's symmetry group on $\langle n'|\mathcal{S}|n\rangle = \mathcal{R}(ph)$. This is a sandwich of phonon states around an invariant operator $\mathcal{S}$. We know that $\mathcal{R}(ph)$ is a 3×3 tensor. Because $\langle n|n\rangle = 1$, the sandwiching states that differ by a single phonon $\langle n+1|...|n\rangle$ must transform like the single extra phonon. It's the difference between the two states that matters. This means that under all symmetry operations, $\mathcal{R}(ph)$ must behave the same way as the displacement pattern of the single phonon.

In stronger terms: For the purposes of the scattering experiment, $\mathcal{R}(ph)$ *is* the phonon. Because of this fact, even without knowing the displacement pattern, you can see what sorts of $\mathcal{R}(ph)$ are "allowed" to exist, and so which $\langle\gamma'|...|\gamma\rangle$ polarizations might be interesting to try. You do this by comparing $\mathcal{R}(ph)$ with a symmetry-transformed version of itself.

As an example, consider a flat (x, y) plane square molecule confined to move in two dimensions, with symmetry C_4:

$$\begin{array}{cc} \uparrow & \rightarrow \\ \leftarrow & \downarrow \end{array} .$$

Assume the excitation comes in along the z axis, and the scattered light comes out in the opposite direction, so we can represent our $\mathcal{R}$ with a $2{\times}2$ matrix and not worry about polarizations in z. Let's also denote the experimental polarization conditions by $[\vec{in}, \vec{out}]$.

The square molecule is invariant under the identity, and turns by $\pi/2$, π, and $3\pi/2$ about a z axis through its center. Under the same operations, the polarization combination$[\vec{x}, \vec{y}]$ rotates to $[\vec{x}, \vec{y}]$, $[\vec{y}, -\vec{x}]$, $[-\vec{x}, -\vec{y}]$, $[-\vec{x}, -\vec{y}]$, and $[\vec{x}, -\vec{y}]$. Hence the phonon states that you should see in $[\vec{x}, \vec{y}]$ and $[\vec{y}, \vec{x}]$ must pick up phases of $1, -1, 1, -1$. A phonon visible in these conditions should have the same transformation properties. For such a 2×2 tensor,

$$\begin{pmatrix} 0 & 1 \\ -1 & 0 \end{pmatrix} \begin{pmatrix} A & B \\ C & D \end{pmatrix} \begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} -A & -B \\ -C & -D \end{pmatrix} .$$

By evaluating the product,

$$\begin{pmatrix} D & -C \\ -B & A \end{pmatrix} = \begin{pmatrix} -A & -B \\ -C & -D \end{pmatrix} ,$$

which implies

$$\mathcal{R}(ph) = \begin{pmatrix} A & B \\ B & -A \end{pmatrix} .$$

If we got rid of the arrows, giving the molecule a center of inversion, then we could also exchange the diagonal elements. This would imply that, for this particular phonon, the diagonal A's would be zero. Anyhow, it's easy to find a displacement pattern which fits the bill: a diamond stretch/squish distortion.

You can also start with a displacement pattern and predict polarization conditions. A displacement pattern that preserves the square would be like the displacement patterns mentioned above, but with all displacements in phase. The matching polarization conditions are $[\vec{x}, \vec{x}]$ and $[\vec{y}, \vec{y}]$, which each have a phase of $+1$ under all rotations. Repeating the same analysis with these phases yields the partial Raman tensor

$$\mathcal{R}(ph) = \begin{pmatrix} A & 0 \\ 0 & A \end{pmatrix} .$$

6 Enumerating the Modes (requires some group theory)

An irreducible representation of a group is defined by the list of its traces. In one-dimensional representations, these are just phases. From the foregoing discussion of displacement patterns, it should be no surprise that the phonon modes can be classified by irreducible representation (*ie* the set of phonon phases) in the group.[3]

It's easy to construct a general representation for "any" phonon displacement pattern. Each irreducible representation's contribution to this general representation can be counted by the standard procedure for reducing representations, in much the same way one projects out vector components using a dot product. The method hinges on the orthogonality of the characters of irreducible representations. Character tables of the point groups and for crystal space groups at representative points in the Brillouin zone are available in books and online.
[4]

6.1 Constructing the General Phonon Representation

With crystals, Raman phonons live near the center of the Brillouin zone (zero wavevector) because the wavelength of visible light is large compared to the lattice spacing. We can therefore concentrate on the subgroup of the crystal space group corresponding to zero wavevector. This is just the point group of the unit cell, with symmetry elements distrubuted according to the particular space group (e.g., for D_{4h}, it could be D_{4h}^{14}). For molecules, of course, there is only the point group.

For the point group, the general representation of a phonon is the one which simultaneously depicts the vector displacements of all of the atoms in the molecule or the unit cell–it is DN dimensional, if N is the number of atoms in the molecule/unit cell and D is the dimension. Under a symmetry operation, two tiers of transformations occur.

Firstly, sets of equivalent ions are shuttled among themselves. If, as a basis of the DN dimensional representation, we use the displacement coordinates of each ion, then for a given element of the group, only those ions whose identities are unchanged in the transformation can contribute to the character. For crystals, these are ions that reside on the rotational axes, reflection planes, etc. The others will be by definition off-diagonal, and so cannot contribute to the trace.

---

[3] What might surprise you is that some phonons acquire phases that are incompatible with representation by a rank-two tensor: if you try to make a $\mathcal{R}(ph)$ out of these, you get all zeros! These phonons are called *Raman silent modes.*

[4] For crystals, you need to identify the symmetry elements (e.g., rotational axes, mirror planes, etc.) in the unit cell. For this purpose, you might use the excellent diagrams in *The Analytical Expression of the Results of the Theory of Space Groups*, by Ralph W. G. Wyckoff (Carnegie Institute of Washington, 1930).

For each one of these special “diagonal” ions, the contribution to the trace is just the contribution of a vector displacement of unit length under the given transformation. The sum of the vector-trace contributions for the diagonal ions gives the trace of the DN-dimensional representation for the transformation; the collection of these traces for the entire group gives the character of this representation for the group.

6.2 Reducing the General Phonon Representation

Once you have the traces of the DN dimensional representation, you can decompose it into irreducible components by using the orthogonality of characters of irreducible representations. Each allowed phonon will appear once, and so the number of times $m(\alpha)$ a given irreducible representation α comes up gives the number of allowed phonons of that symmetry species. If d_α stands for the dimension of representation α, the sum rule

$$\sum_\alpha d_\alpha m(\alpha) = DN$$

must of course hold. This is a good check of your arithmetic.

For an example, consider our square molecule, now allowed to move in three dimensions. The character table of C_{4h} is

	e	C_4	C_2	C_4^3	σ_h	S_4	I	S_4^3
A	1	1	1	1	1	1	1	1
A'	1	1	1	1	-1	-1	-1	-1
B	1	-1	1	-1	1	-1	1	-1
B'	1	1	-1	1	-1	1	-1	1
E_1	1	i	-1	$-i$	1	i	-1	$-i$
E_2	1	$-i$	-1	i	1	$-i$	-1	i
E_1'	1	i	-1	$-i$	-1	$-i$	1	i
E_2'	1	$-i$	-1	i	-1	i	1	$-i$

Furthermore, the only operations that leave atoms in place are the identity e, and the reflection in the molecular plane, σ_h, which leave all four. In three dimensions, the vector character of the identity is obviously 3, and therefore the number of atoms × the vector character, $N_A\chi^v = 12$. For σ_h, the vector trace is 1 (because z changes sign but x, y don't). So for σ_h, $N_A\chi^v = 4$. Using the orthogonality relation

$$\sum_{i\epsilon g} \chi_i^{\alpha *} \chi_i^{\alpha'} = |g|\delta_{\alpha,\alpha'},$$

to decompose the 3N-dimensional representation (here $|g|$=8), compute

$$\begin{gathered}
8m(A) = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^A = 16 \Rightarrow 2A \\
8m(A') = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{A'} = 8 \Rightarrow 1A' \\
8m(B) = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^B = 16 \Rightarrow 2B \\
8m(B') = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{B'} = 8 \Rightarrow 1B' \\
8m(E_1) = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{E_1} = 16 \Rightarrow 2E_1 \\
8m(E_2) = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{E_2} = 16 \Rightarrow 2E_2 \\
8m(E_1') = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{E_1'} = 8 \Rightarrow 1E_1' \\
8m(E_2') = \textstyle\sum_{i\epsilon g} N_A(i)\chi_i^v\chi_i^{E_2'} = 8 \Rightarrow 1E_2'
\end{gathered}$$

There are $N_A \times D = 12$ degrees of freedom, breaking down into the modes

$$2A \oplus A' \oplus 2B \oplus B' \oplus 2E_{1,2} \oplus E'_{1,2} (= 12),$$

where the usual thing is to collect the sets of complex conjugate E_1, E_2 representations into 2-dimensional representations.

Of course three of these modes correspond to translations of the whole molecule, and three to rigid rotations. Because displacements are just vectors, the trace for translations is just χ^v again. A quick calculation using χ^v for all the elements gives the translational modes as E_1, E_2, A'. Rotations transform like pseudovectors, and you could figure out the pseudovector χ^p's, etc. How should the three rotations transform?