

Notes on wavefunctions

The double slit experiment

In the double slit experiment, a beam of light is sent through a pair of slits, and then observed on a screen behind the slits. At first, we might expect to see a pair of bright lines on the screen, behind the two slits. In reality (assuming the separation between the slits is similar to the wavelength of the light), the brightest region is exactly in the middle, between the two slits. Surrounding this bright region are a series of bright and dark areas. If we define θ to be the angle of deflection of the light (see the diagram in the notes), the bright spots occur for $\sin(\theta) = n\lambda/d$ with $(n = 0, 1, 2, \dots)$ and the dark regions occur for $\sin(\theta) = (n + \frac{1}{2})\lambda/d$ with $(n = 0, 1, 2, \dots)$.

The explanation for this pattern of bright and dark spots is based on the interference of light from the two slits. At the location of the slits, the light waves are in phase with each other: the electric field oscillates up at one slit when it oscillates up at the other slit. The light waves then propagate outward from these two slits toward the screen. In the middle of the screen between the two slits, the distance to each slit is the same, so the light waves from these two slits are still in phase with each other. The two light waves constructively interfere, and we see a bright spot. As we move away from the center, we come to a place where the distance between the two slits differs by half a wavelength. Here, the two light waves are out of phase with each other; thus, we have destructive interference (the waves cancel each other out) and we have a dark spot.

The points on the screen at angles $\sin(\theta) = n\lambda/d$ are exactly the ones for which the distance to the two slits differs by an integer number of wavelengths (so that we get constructive interference) while the angles $\sin(\theta) = (n + \frac{1}{2})\lambda/d$ are the ones for which the distance to the two slits differs by $n + \frac{1}{2}$ wavelengths (so that we get destructive interference). The details of this may be found in the text, but for our purposes, there are two essential points to take away:

- The pattern on the screen is explained by understanding how light from the two slits combines.
- The pattern on the screen is explained using the fact that light is a wave with a particular wavelength.

The double slit experiment: photon explanation

What happens if we reduce the light intensity so much that only one photon at a time passes through the apparatus? In this case, we'll need to replace the screen with a more sensitive detector that can tell when and where a single photon hits (such a detector could be built from an array of photomultipliers or with a device similar to the "film" in digital cameras). In this case, we find that each photon hits a specific spot on the screen, but after many photons have hit, the distribution of hits reproduces the pattern of bright and dark regions that we had before.

How could we possibly get the same pattern as before if the pattern was caused by light from the two slits combining and interfering? If we just send individual photons through, don't they have to go through just one slit or the other? If so, how can there be anything like interference? The answer is that the individual photons don't just go through one slit or the other. Somehow, *each photon is able to go through both slits and interfere with itself!* To convince ourselves that this is true, we only need to repeat the experiment, but randomly cover up one slit or the other each time we send a photon through. If individual photons can only go through one slit or the other, this should not affect the pattern of hits, except that we should have half as many photons hitting the screen. But in fact, when we cover slits in this way, the pattern of hits completely changes. We conclude that the single photons must know about both slits; rather than being pointlike particles, the photons must be spread out in some way, enough so that they can pass through both of the slits.

But wait! didn't we say that each photon hits a specific spot on the screen? Now it sounds like we have a contradiction: we said each photon must be spread out enough to go through both slits, but also that each photon is detected at a single location on the screen.

Quantum superpositions

To resolve this, we need to bring in the idea of a QUANTUM SUPERPOSITION that we used to understand the polarization experiments. Here are the basic rules:

- We can have photons with specific positions, but these are special states, which we call POSITION EIGENSTATES.

- General states are combinations, or “quantum superpositions” of these position eigenstates. They don’t have a definite location.
- When a measurement of position is made, the photon becomes one of the eigenstates. It randomly “chooses” a position, with the relative probability for each possible position related to the amount of that eigenstate present in the quantum superposition.

How does this apply to our photons? We can say that the photons approaching the double slits are in spread-out quantum superpositions involving many different position eigenstates. It is possible for the photon to go through both slits, since it can be in a state which is the superposition of being at one slit and being at the other slit. As the photon approaches the screen, it is still in a superposition, so there is no predetermined location where it will hit. But hitting the screen is like making a measurement of the photon’s position, and at that time (according to the third rule) the photon will turn into one of the position eigenstates, and therefore hit the screen at a definite location.

Same story for electrons and other particles

It turns out (in accord with de Broglie’s prediction that particles like electrons will exhibit a wavelength) that the double slit experiment produces exactly the same results for electrons and other particles, and requires exactly the same type of explanation in terms of quantum superposition. Thus, going forward, the discussion will apply equally well to electrons, neutrons and other particles as it does to photons.

Mathematical description of quantum superpositions

Review of the polarization story

How do we describe these quantum superpositions precisely? Let’s first recall the polarization story. There, we represented polarization states by unit vectors $|\theta\rangle$. We said that for a given polarizer experiment, there are special states for which we can predict for sure whether or not the photon will pass through. These are the polarization eigenstates. For example, for a polarizer oriented vertically, we have two polarization eigenstates, which we call $|0^\circ\rangle$ and $|90^\circ\rangle$ (representing the unit vectors pointing vertically and horizontally).

All other polarization states can be written as superpositions of these two eigenstates:

$$|\theta\rangle = a_{0^\circ}|0^\circ\rangle + a_{90^\circ}|90^\circ\rangle$$

where a_{0° and a_{90° are two numbers that determine which state we are talking about (for example $a_{0^\circ} = a_{90^\circ} = 1/\sqrt{2}$ gives us a photon polarized at 45 degrees). Thus, the eigenstates act like basis vectors.

For general states, we can't predict with certainty whether or not they will pass through the polarizer, but we can say that with probability $|a_{0^\circ}|^2$ the photon will become the $|0^\circ\rangle$ eigenstate when it hits the polarizer and pass through, and with probability $|a_{90^\circ}|^2$ it will become the $|90^\circ\rangle$ eigenstate and be absorbed.

Superpositions of position eigenstates and wavefunctions.

Now on to the position eigenstates. In this case, we have a lot more eigenstates: in an experiment where we measure a particle's location, the result could be anywhere, so we have one eigenstate for every position in space. We'll label these states $|\mathbf{x}\rangle = |(x, y, z)\rangle$, so for example, $|(0, 0, 0)\rangle$ represents a state where we have a particle with a definite location at the origin. In the model of quantum superposition, these again behave like basis vectors.¹ So we can consider a more general state which is a superposition of these. For example

$$a_{(0,0,0)}|(0, 0, 0)\rangle + a_{(1,2,3)}|(1, 2, 3)\rangle$$

represents a state of a particle that doesn't have a definite location. For this particular state, if we make a measurement of the electrons position, we will find it at $\vec{x} = (0, 0, 0)$ or at $\vec{x} = (1, 2, 3)$ with probability $|a_{(0,0,0)}|^2$ and $|a_{(1,2,3)}|^2$ respectively.

The most general state we could imagine is a superposition of not just two, but all the possible eigenstates. To describe such a state, we need a number for each eigenstate that tells us how much of that state we have in the superposition (like the numbers $a_{(0,0,0)}$ and $a_{(1,2,3)}$ in the previous example). We'll denote this number by $\psi(x, y, z)$. To write this general state, we want to add up all the eigenstates $|(x, y, z)\rangle$ with $\psi(x, y, z)$ as the coefficient. Since

¹Note: these are no longer vectors in a two-dimensional space. Since we have an infinite number of basis vectors, we should just think about them as living in some kind of abstract space of all the possible states; in particular it is not correct to think of the basis vector $|(x, y, z)\rangle$ as pointing in the three-dimensional where the particles live.

the states are labeled by continuous parameters, we need an integral rather than a sum, so we write

$$\int \int \int \psi(x, y, z) |(x, y, z)\rangle dx dy dz$$

or in a more compact notation

$$\int d^3x \psi(\mathbf{x}) |\mathbf{x}\rangle .$$

To summarize, the integral here is just a sum over all the possible positions, and in the sum we are adding up eigenstates with $\psi(\mathbf{x})$ telling us how much of each to add.

In practice, the function $\psi(\mathbf{x})$ tells us everything we can know about the state of the particle. This is known as the WAVEFUNCTION for the particle. If we make a measurement of position, $|\psi(\mathbf{x})|^2$ tells us the probability density for finding the particle at location \mathbf{x} . For example, to find the probability that the electron will be measured in a spatial region V , we integrate $|\psi(\mathbf{x})|^2$ over this region. Since the net probability of finding the electron somewhere in space must be 1, we require that any physical wavefunction must have $|\psi(\mathbf{x})|^2$ integrate to 1 over all space.

Measurement

The wavefunctions for particle states are generally spread out (non-zero for a wide range of positions). However, when we measure the location of a particle, we always find it in one place. Furthermore, an immediately repeated measurement of the particle's location will give the same result (otherwise, it would be somewhat meaningless to say that the electron is at the location where we measure it). This implies that the wavefunction generally changes in a measurement: after the location is measured, the particle does have a definite location so the state is now (approximately) one of the position eigenstates. The wavefunction changes from being spread out (non-zero for a wide range of positions) to being localized (non-zero only for one location or a small range of locations). This is often referred to as the COLLAPSE of the wavefunction.

Why should a measurement have any effect on the state of the particle? Can't we measure the particle without disturbing it? The answer is that

in order to learn something about the particle, we have to interact with it somehow (e.g. by scattering a photon off it or having it hit a detector). In doing this, we always affect the state of the particle, apparently in such a way that the wavefunction collapses to the wavefunction for an eigenstate.

Classical versus quantum

The wavefunction (or the idea of quantum superposition) represents an entirely new description of physical systems. In the classical (pre-quantum) picture, the state of a particle at a given time is completely specified by giving its position and velocity. From these, the position and velocity at any later time can be predicted using Newton's Laws. Now, instead of these six numbers (three position coordinates and three velocity components), we need to give a whole function $\psi(\mathbf{x})$ to describe the state of a particle at some time.

Just as position and velocity evolve with time in the classical picture, the wavefunction for a particle evolves with time (as we saw in the double-slit simulation). We therefore need something to replace Newton's laws; something that will allow us to predict the wavefunction at some later time given the wavefunction at some initial time. This evolution equation for wavefunctions is known as the SCHRÖDINGER EQUATION. Our next task will be to "derive" this based on some simple observations and principles.