What does REALISTIC Look Like?

by Anthony Waichulis

What does Realistic Look Like?

By Anthony Waichulis

v1.2

CONTENTS

anthonywaichulis.com

“Imagine if a technology existed that could allow us to push neural activity toward a particular brain state, yielding any perceptual experience we like. While to some this may seem like an impossible dream of science fiction, I would argue that we have had access to such devices for thousands and thousands of years--and they are currently populating the walls of most museums and galleries today.

Keep this in mind the next time you pick up that paintbrush or pencil. You’re building an incredible machine with each mark.”

About this text:

Due to the nature of my work I am often asked, "*How do you get your paintings to look so realistic?*" It's a great question in the sense that it usually opens the door to an enjoyable conversation about art, visual perception, and a host of related topics that I find utterly fascinating. However, much of that ensuing conversation is determined by how the inquirer responds to a question that I pose in response to theirs,

"***What does realistic look like?***"

The text I present here is my framework for tackling this latter question. The answer to the former is the result of methodologies that I have developed from my efforts to answer the latter. While many may find this writing to be somewhat "sciencey," I have worked to simplify everything as much as I could without significantly diminishing the accuracy or usefulness of the content. I apologize in advance if this effort to simplify misrepresents or miscommunicates any of the science. If you would like to pursue a more robust understanding of any of the scientific ideas presented here, I will offer a list of resources at the end of this text so that you may pick up where this text leaves off.

I would also like to mention that in writing these articles, I tend to include the questions that most often pop up during my many conversations on these topics. I hope that the inclusion of these questions is not too annoying.

Lastly, if you have any questions, feedback, or would like to notify me of any errors that I made, please feel free to contact me via email at: yychuls@gmail.com

Best wishes,

Anthony Waichulis

Part I - Not-So-Terrible Lizards

Boy did I love dinosaurs. I mean I LOVED them. By the age of five, I knew more about the age of the "terrible lizard" than I did about the world that existed outside my front door. Don't get me wrong, I indeed climbed my fair share of trees and learned to corner with a big-wheel like I stole it, but the vast majority of my early years were spent staring at captivating images of elaborate Tyrannosaurus-Triceratops battles, learning about the armor that protected the heavily fortified Ankylosaurus, and marveling at the size of my favorite sauropods like the Diplodocus and the Brontosaurus. Each image-laden page of my favorite How and Why, Golden Book and Golden Guide was a gateway to a world of giant monsters that I would anthropomorphize into a full cast of beloved characters.

However, as with many of the new worlds or landscapes that we are fortunate enough to explore, the boundaries of this realm seemed to constrict with time. As such, it wasn't long before I felt the urge to take my familiar characters into new lands for fresh adventures. And so I did. I put aside the old gateways for a new one that was built from vehicles that seemed to hold a greater potential for exploration. Sometimes these vehicles were in the form of plastic toy reptiles, but most often they were formed amid a vast sea of drawing paper and a giant butter cookie tin of weathered crayons.

And then one day something happened. While drawing one of the tens of thousands of adventures with my "dino-pals," I made an incredible discovery. Now while this discovery might not seem as impressive to some as William Buckland discovering his Megalosaurus fossils in 1819, but to my young mind, it was just as profound. So much so that it would begin something that would shape the way my life would unfold. Interestingly enough, much like a fossil hunter, my discovery was to be found with the leg of one dinosaur and the jaw of another.

Let me explain.

Up until the day of my discovery, the vast majority of my drawings resembled the images that one might expect to find scrawled on the walls of certain caves in France or Spain. Even though my visual experience with dinosaurs had come from robust illustrations that were rife with representations of depth and dimension, my early representations were quite flat and seemingly one dimensional. My characters were almost always in profile, contained little-to-no depth cues, and seemed to reflect what some psychologists might describe as "split-type" characteristics (i.e., a type of representation that aims to communicate the important features of an object in a manner that would not normally be seen from a single perspective.)

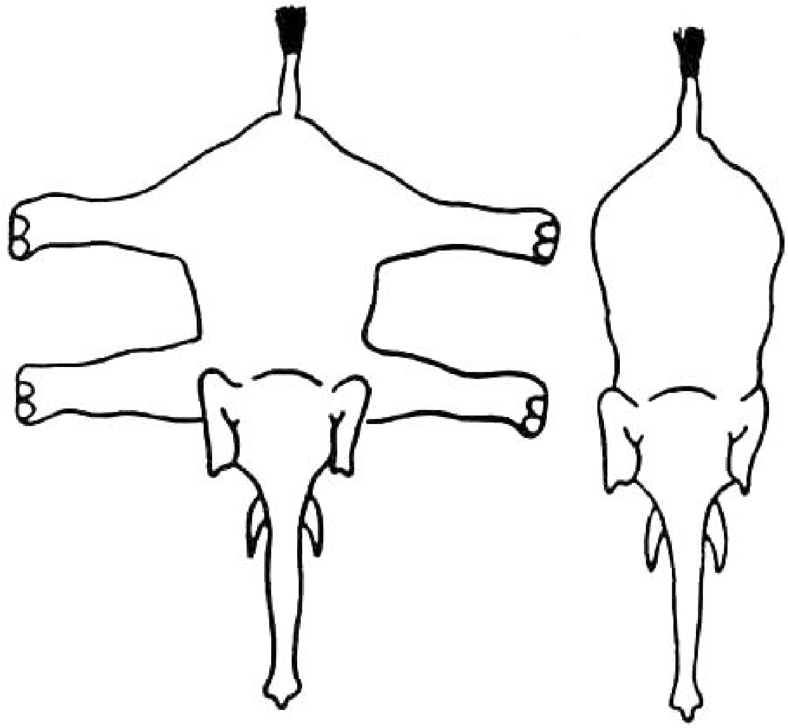

Elephant drawing split-view and top-view perspective. Such illustrations were used by psychologists William Hudson and Jan B. Deregowski in their efforts to better understand the impact of culture on visual perception.

One day, while drawing one of my favorite dinosaurs, the Dimetrodon (the one with a big sail on its back), I remember that I stopped drawing as I approached the underside. I had confidently drawn this fellow countless times, but something seemed different this time. I remember pausing for a great while to study the lines defining the bottom of the dinosaur. And then, as if by accident, I added a line that I never did before.

The line that I added separated the leg opposite the viewer from the belly of a Dimetrodon. It immediately put the leg BEHIND the mass of the body. This was the discovery--a simple line that fundamentally changed drawing for me. The “flat” giants that I had become so adept at drawing evolved instantly before my eyes. The familiar sailback was immediately something more than it was before.

It was not long before all of my characters had this dimensional augment. And while this newly found update to my representations would have been enough to keep me content for some time, it quickly happened again. This time though it was with the jaw of the infamous Tyrannosaurus Rex. It was actually a line that I made too long. However, working in crayon did not offer much in terms of erasing, so I just continued it to the nearest contour--and voila--another transformative improvement was uncovered. I simply added a small u-turn-esque curve to the T. Rex jawline that instantly communicated something more dimensional. Like the legs and torso of the Dimetrodon, the altering of my routine profile for a more depth-rich, near-three-quarter representation was like stumbling across another great secret of the universe. I was elated beyond words.

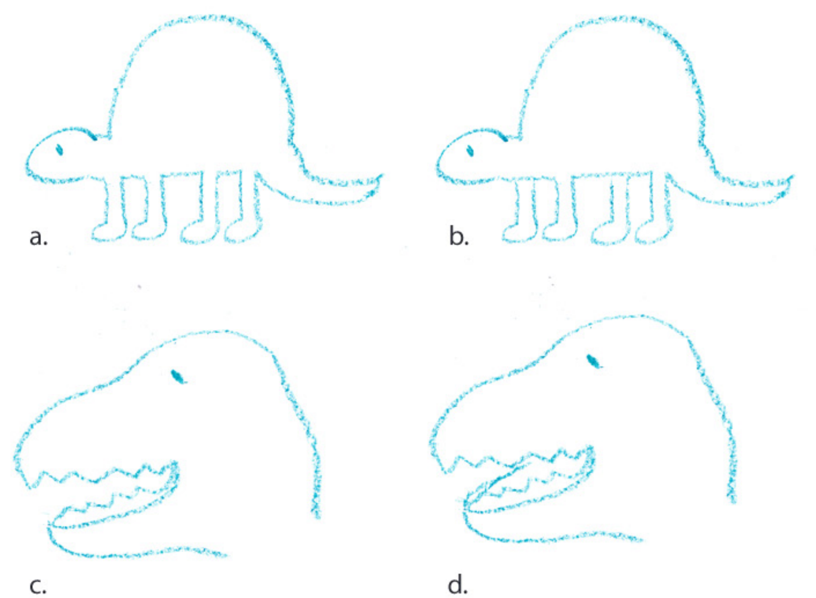

Illustrations simulating the described line additions/augments that made the representations appear more "realistic" to me. (a) Dimetrodon prior to leg occlusion line. (b) Dimetrodon after leg occlusion add. (c) T. Rex drawing prior to curved jawline augment. (d) T. Rex drawing after curved jawline augment.

But now, many years and seemingly thousands of artistic endeavors later, I find myself wondering why I was so elated by these developments in my early drawing efforts. I remember an immense feeling of pride in the fact that my representations began to look more "realistic" due to these augments, but in hindsight--it seems absurd that I would have thought that. I mean, it's not like I had ever seen real dinosaurs to compare my drawings with. I had only ever seen illustrations in books that I also thought looked VERY realistic. But why did I think that those illustrations were realistic? It's not like I thought that the artists that created those striking images had ever seen a real dinosaur either as just about every book opened with the fact that the dinosaurs had been long gone before any humans showed up with their pencils and sketch pads.

So what do we mean when we describe a visual representation as "realistic?" If we look to the current dictionary resources available on the web we can find some fairly straightforward definitions regarding realistic depiction:

Dictionary.com defines realistic as "pertaining to, characterized by, or given to the representation in literature or art of things as they really are," while the Oxford dictionary offers, "representing things in a way that is accurate and true to life." Merriam Webster defines the term as the following, “the theory or practice of fidelity in art and literature to nature or to real life and to accurate representation without idealization; accurately representing what is natural or real; convincingly rendered to appear natural.”

realistic
adjective • **US** /ˌri-əˈlɪs-tɪk/

- **having or showing a practical awareness of things as they are:**
  She is realistic about her chances of winning.
- **realistic also means appearing to be existing or happening in fact:**
  the scene in the movie where the dinosaur hatches from the egg is incredibly realistic.

Screen capture from the online Cambridge Dictionary (https://dictionary.cambridge.org/us). Again, how does the speaker in the second example know what a dinosaur egg hatching would look like enough to deem the representation realistic?

Unfortunately, these definitions, while useful in many general contexts, come up short for those visual artists looking to gain a deeper understanding of the nature of realistic representation. To understand why this is the case we will have to begin with an examination of the fascinating and often counterintuitive biological process of visual perception.

Part II - Visual Perception

Visual perception can be defined as the ability to interpret the surrounding environment by processing information that is contained in visible light. It is a fascinating process that takes up about 30% of our brain's cortex. While there is much that we have come to understand about this remarkable ability, there is a great deal of mystery still to be solved. Unfortunately, very little about our vision system is intuitive. In fact, it is not uncommon to find people that intuitively think of the eyes as tiny cameras that record accurate, clear pictures that are eventually sent on to some inner region of the brain where a small "inner self" or homunculus (a Greek term meaning "little man") awaits to review the image. What's more is that even though some can push past the "homunculus fallacy," they still manage to get mired in the idea that vision is (at least in some part) veridical (objectively truthful, corresponding to objective measurement or fact.) Current evidence about the nature of visual perception refutes this idea, but that does not stop many from arguing for the veridical nature of vision regardless. One of the most recent arguments I have encountered regarding this misconception was from a colleague who insisted that it is very likely that vision may become more accurate/veridical when percept properties reduce in complexity. Unfortunately, this argument would be akin to one which claims that the accuracy of one's "clairvoyance" increases among pick-a-card prediction tasks when the number of cards in play is reduced. With this reasoning, if we limit the pile to one card– one would be likely to hit 100% accuracy.

The truth is that a growing body of evidence indicates that our visual percepts are generated according to the empirical significance of light stimuli (empirical information derived from past experience), rather than the characteristics of the stimuli as such. In other words, modern research points to the idea that the function of the visual system is not to provide a veridical window on the environment, but rather, has evolved to provide adaptive, neural responses that are useful in achieving successful behavior amid a dynamic, and sometimes hostile, environment. The vision system did not evolve for veridicality--but rather for evolutionary fitness. As such, the properties of the visual system should not be confused with devices that can garner reasonably accurate, objective measurements of the physical world (e.g. calipers, light meter, spectrophotometer, etc.) To this point, Neuroscientist Dale Purves states,"*...vision works by having patterns of light on the retina trigger reflex patterns of neural activity that have been shaped entirely by the past consequences of visually guided behavior over evolutionary and individual life time. Using the only information available on the retina (i.e., frequencies of occurrence of visual stimuli, light intensities), this strategy gives rise to percepts which incorporate experience from trial and error behaviors in the past. Percepts generated on this basis thus correspond only coincidentally with the measured properties of the stimulus or the underlying objects.*"

Support for this idea comes not only from fields like perceptual neuroscience and modern vision science but from other realms of inquiry. Donald D. Hoffman, a professor of cognitive science at the University of California, Irvine has spent the past three decades studying perception, artificial intelligence, evolutionary game theory, and the brain. His findings indeed lend significant support to the idea that visual perception is indeed non-veridical. In a 2016 article with Quanta Magazine, Dr. Hoffman states, "*The classic argument is that those of our ancestors who saw more accurately had a competitive advantage over those who saw less accurately and thus were more likely to pass on their genes that coded for those more accurate perceptions, so after thousands of generations we can be quite confident that we're the offspring of those who saw accurately, and so we see accurately. That sounds very plausible. But I think it is utterly false. It misunderstands the fundamental fact about evolution, which is that it's about fitness functions — mathematical functions that describe how well a given strategy achieves the goals of survival and reproduction. The mathematical physicist Chetan Prakash proved a theorem that I devised that says: According to evolution by natural selection, an organism that sees reality as it is will never be more fit than an organism of equal complexity that sees none of reality but is just tuned to fitness. Never.*" Furthermore, in regards to the hundreds of thousands of computer simulations run by Dr. Hoffman and his research team, he states, "*Some of the [virtual] organisms see all of the reality, others see just part of the reality, and some see none of the reality, only fitness. Who wins? Well, I hate to break it to you, but perception of reality goes extinct. In almost every simulation, organisms that see none of reality but are just tuned to fitness drive to extinction all the organisms that perceive reality as it is. So the bottom line is, evolution does not favor veridical, or accurate perceptions. Those perceptions of reality go extinct.*"

But the counterintuitiveness of the vision system is not limited to the big picture (pardon the pun.) Even at the earliest steps through the vision process, we find one counterintuitive factor or scenario after another. For example, did you know that when light energy interacts with our specialized light-sensitive receptors, they respond with less activity instead of more? (in other words, our light-sensitive cells (photoreceptors) are far more active in the absence of light!) Or how about the fact that all of the biological machinery that effect processes downstream from those receptors is found upstream (getting in the way of the light)? And while many of you might be aware that the incoming light patterns are inverted and reversed on the retina, did you know that those patterns must first pass through a dense web of blood vessels that you have perceptually adapted to? It's true. So all in all, in the light of all of the interdisciplinary evidence that we have at present, visual perception is a counterintuitive, experience-driven, non-veridical sensory process that is tuned to fitness. It is not a clear window on the environment.

Now even at this point, I think that one can begin to see the glaring issue that exists with the aforementioned definitions for "realistic." Nonetheless, I would like to look at a some of the observable machinery and processes that we find in our "visual pathway" so that you may better appreciate, or dissect, some or all of the concepts that will be presented a bit later on.

Let's begin with light. As I put forth at the beginning of this section, visual perception can be defined as the ability to interpret the surrounding environment by processing information that is contained in visible light. This term, visible light, describes a portion of the electromagnetic spectrum that is visible to the human eye. The electromagnetic radiation in this range of wavelengths is called visible light or more simply, light. A typical human eye will respond to wavelengths ranging from about 390 to 700 nanometers.

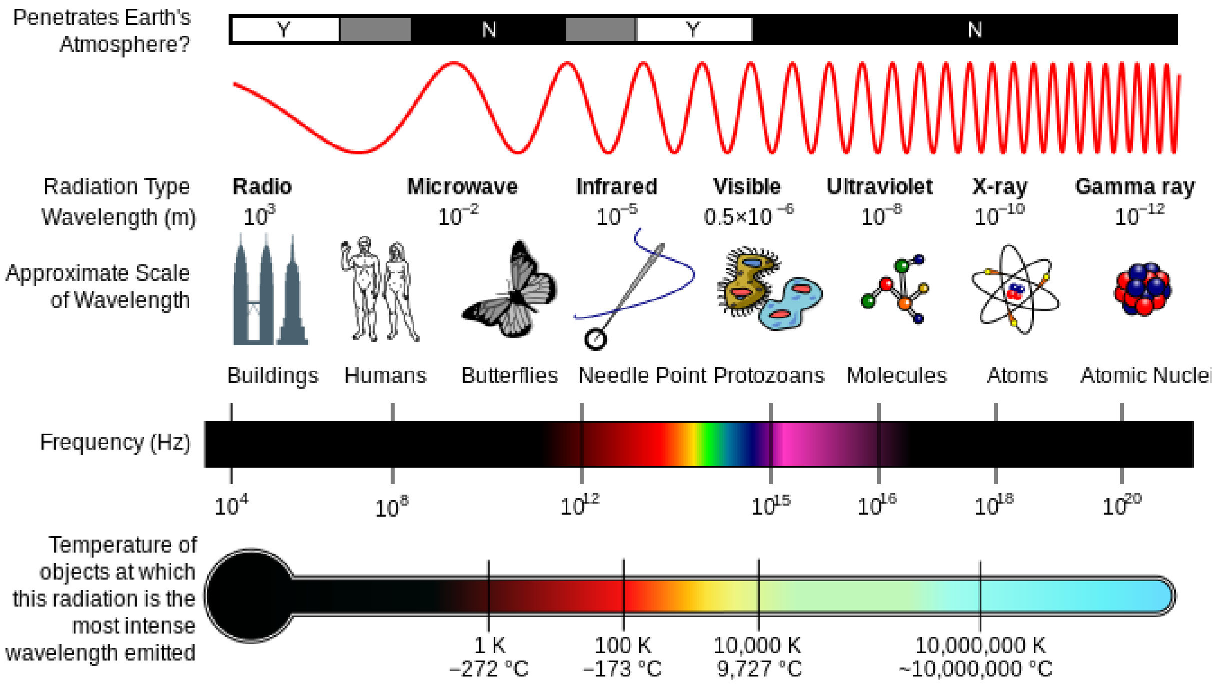

A diagram of the electromagnetic spectrum, showing various properties across the range of frequencies and wavelengths.

Visible light rays enter the eye through a small aperture called the pupil. A dome-shaped transparent structure covering the pupil (called the cornea) assists in this entrance, and with the help of a biconvex lens right inside of the pupil, guides the light rays into a focused light pattern onto a region at the back of the inner eyeball called the retina. This tissue is the neural component of the eye that contains specialized light-sensitive cells called photoreceptors. It is these specialized photosensitive receptors that will convert the environmental energy (light) into electrical signals our biological machinery can use in a process called phototransduction.

Now upon reading that last paragraph, some might be quick to state that thus far, the eye sure sounds like a camera. I mean, the basic idea with photography is to record a projected light pattern with an electronic sensor or light-sensitive plate that can be used to generate a percept surrogate. Generally speaking, in the case of the digital camera sensors (something often compared with the retina), each pixel in the sensor's array absorbs photons and generates electrons. These electrons are stored as an electrical charge proportional to the light intensity at a location on the sensor called a potential well. The electric charge is then converted to an analog voltage that is then amplified and digitized (turned into numbers.) The composite pattern of data from this process (stored as binary) represents the pattern of light that the sensor was exposed to and may be used to create an image of the light pattern recorded during the exposure event. Additionally, a Bayer mask or Bayer filter (a color filter array) is placed over the sensor so as to collect wavelength information at each pixel in addition to information regarding light intensity.

So is this how the eye or vision works? Do we simply use objective recordings of wavelength and intensity responses at each photoreceptor to produce a clear percept?

Hardly. There are indeed some similarities between the eye and a camera in terms of optics and photosensitive materials--but that is where any significant similarities end. Let's take a look at some of the differences:

First, it is important to understand that our photoreceptor landscape is not anything like the uniform array of pixels found with a digital image sensor. The 32mm retina (ora to ora) has a very uneven distribution of photoreceptors. Two main types, known as rods and cones, differ significantly in number, morphology, and function as well as in their manner of synaptic connection. Rods are far more numerous than cones (about 120 million rods to 8 million cones), are far more light sensitive, provide very low spatial resolution, and hold only one photopigment. Conversely, cones are far fewer in number, less sensitive to light, but provide very high spatial resolution, and come in three types with each type carrying a photopigment that is differentially sensitive to specific wavelengths of light (thus facilitating what we understand as color vision.)

Cones are present at a low density throughout the retina with a sharp peak within a 1.5mm central region known as the fovea. Rods, however, have a high-density distribution throughout the retina but have a sharp decline in the fovea, being completely absent at the absolute center of the fovea (a .35mm central region called the foveola.) To better appreciate the size of our high acuity window resulting from this photoreceptor distribution you need only look to your thumbnail at arm's length.

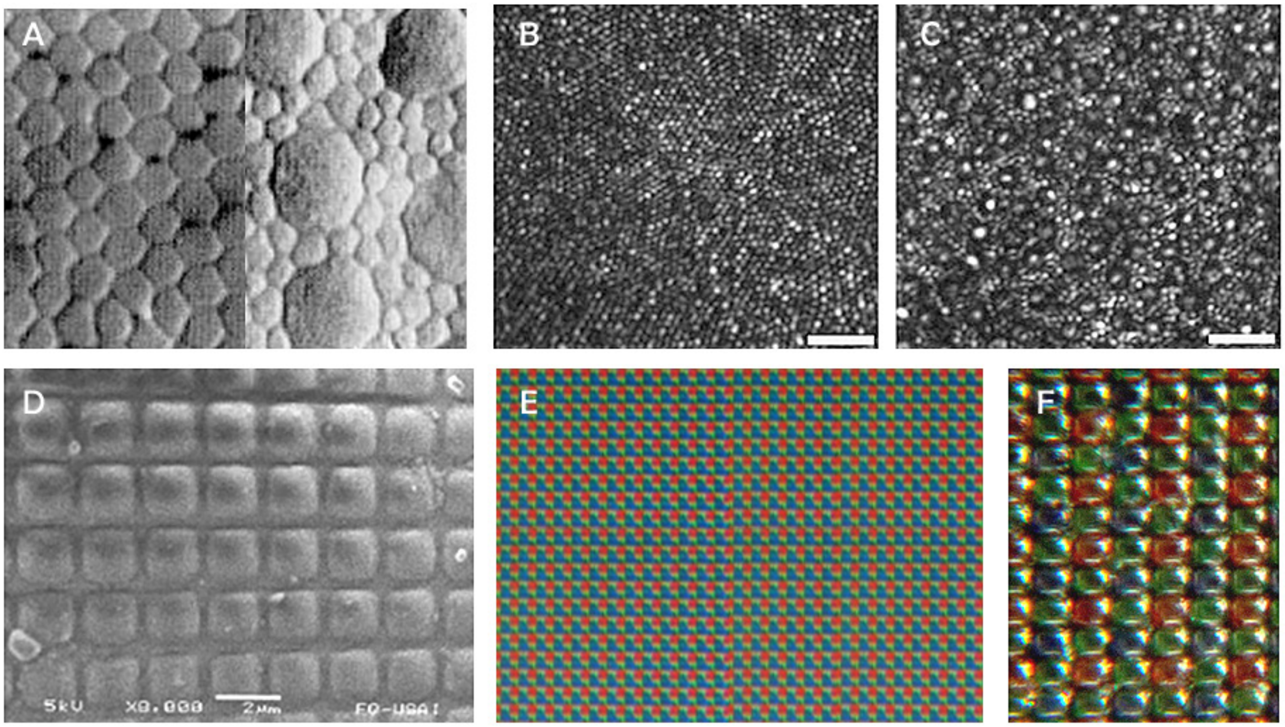

Here we can see some of the uniformity and distribution differences between the photosensitive cells of the retina and the photosensitive pixels in a digital camera sensor. (A) Distribution of cone photoreceptors in the fovea (left) and the cone/ rod distribution the periphery (right). (B) High-resolution image of the foveal cone mosaic obtained with the Rochester AOSLO (adaptive optics scanning laser ophthalmoscopy). (C) Peripheral photoreceptor mosaic showing both rods and cones imaged at 10° temporal and 1° inferior. (obtained with the Rochester AOSLO.) Scale bars are 20 microns. (D) Micrograph of a CMOS sensor at 2 microns. (E) Micrograph of a digital sensor with Bayer mask/filter. (F) A closer look at a Bayer mask/filter.

Another intuitive misconception worth mentioning here, in relation to acuity, is the idea that information gets more "blurry" as we move outward from the fovea to the periphery. The truth is, this lower acuity does not yield image blur, but rather a spatial imprecision. Neuroscientist Margaret Livingstone offers a great demonstration of this point in her book Vision and Art: The Biology of Seeing. Here is a recreation of that demonstration:

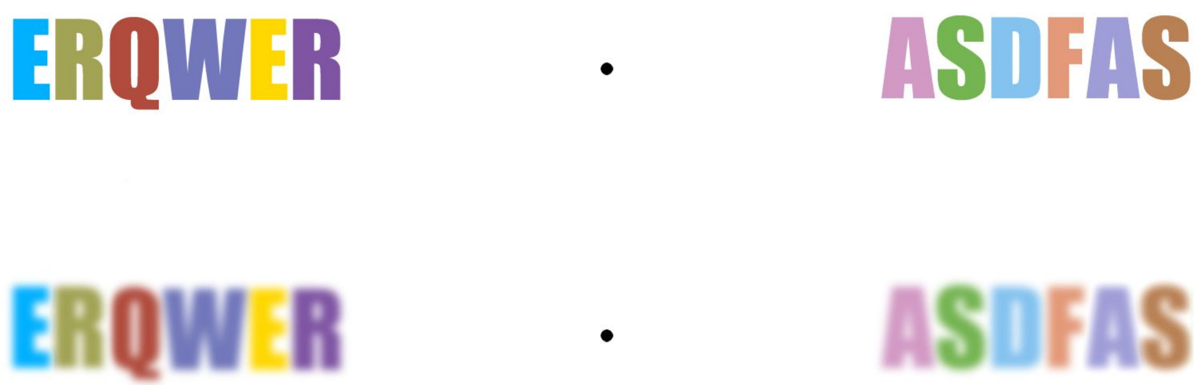

As one stares at the top black dot between the letter strings, we find difficulty in identifying the individual letters. The same holds for the bottom example, however, in staring at the bottom dot we can perceive the blur even though the letters continue to evade identification. This is because our periphery is not "out of focus" or blurred, but is rather spatially imprecise.

Second, we need to look at what happens with the output of our photoreceptors. It is very important to understand here that our rods and cones do not register some absolute measurement of light intensity and wavelength as seen with the image sensor. Rather, responses from these specialized receptors trigger a cascade of highly dynamic, complex processing through multiple cell layers and a myriad of receptive fields. The resulting signals from this retinal activity are then ushered off to our next stop on the visual pathway--the thalamus. But before we head over to this well known "relay station" in the brain, I would like to present another issue at the level of the retina worth consideration when comparing a camera with the eye--the quality of that initial light pattern projection.

Do you remember the last time you took a picture with your camera when you had some debris or smudge on the lens? Did it ruin the shot you were trying to take? How about the last time you had to deal with a piece of tape keeping light from entering a part of the lens? Or the last time you pulled up a pile of plant roots to suspend in front of your camera before taking that nice portrait shot?

Do those last two questions seem a bit ridiculous? Well, those seemingly ridiculous factors represent some real issues that our vision system has to contend with early on in the perception process. As I mentioned when first describing the counterintuitiveness of the visual system, the human eye has evolved with all of the biological "machinery" used to process the photoreceptors output in FRONT of the photoreceptors themselves. That's right--all of the incoming light has to pass through all of the cell layers that will process the output of the photoreceptors. Now while these cell layers themselves are not too much of a problem (as they are relatively transparent), a problem indeed arises when the signals from all of that machinery need to exit the eye.

After the signals that arise from the photoreceptors make their way through all of the other cell layers, they will eventually reach the cells (ganglion cells) whose axons (long slender projections of nerve cells, or neurons, that usually conduct electrical impulses away from the neuron's cell body or soma) will need to leave the eye in the form of a nerve bundle that we call the optic nerve. The region where this nerve exits the eye is called the optic disk, and it indeed creates a deficit in our receptor array. This exit, or “blind spot,” is about 1.86 × 1.75 mm. Oddly enough, this deficit in our visual field measures slightly LARGER than our rod-free region of highest visual acuity.

Here we see an image of the back of the eye (which is known as a fundus photograph). The darker region in the center is an area known as the macula (about 5.5 mm) which contains the fovea. To the right of this darker region, we can see a lighter region which is, in fact, the optic disk. IMAGE CREDIT: Häggström, Mikael, "Medical gallery of Mikael Häggström 2014". WikiJournal of Medicine 1.

If you like, you can even “experience” your blind spot with a simple exercise:

Cover your left eye and look at the dot on the left in this image. You will notice the cross in your periphery, but don't look at it - just keep your eye on the dot. Move your face closer to the page *or monitor (depending on how you are reading this), and farther away. At some point, you should notice the cross disappear. Now, stay at that point, cover your right eye, and look to the cross. You should see that the dot has disappeared.*

Another consideration worth mentioning here is in regards to something that you might have already noticed with the fundus photograph above--a rich web of blood vessels that populates the eye. As you might suspect, these little buggers are also in the way of the incoming light that is heading toward the retina. Now due to what we call sensory or neural adaptation (when a sensory stimulus is unchanging, we tend to stop "processing" it—like the way you don't feel the clothes on your body after a bit), we don't normally perceive these vessels, but by influencing the incoming light we can force this network to reveal itself. To do this, you'll need an index card, a pencil or something to poke a small hole in the index card, and a bright surface.

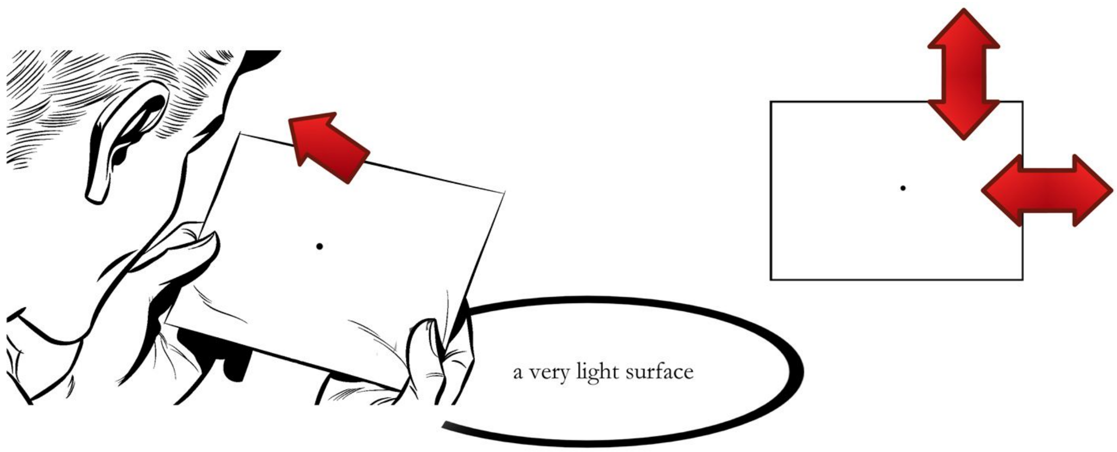

Close one eye, and closely look through the hole at a plain (homogenous) brightly illuminated or light surface. (The card should be right up to your eye). Carefully jitter the card horizontally or vertically and, almost immediately, you will begin to see a grayish web of blood vessels appear. The hole in the index card changes the way that light is entering the eye and thus begins to change the way the blood vessels cast shadows onto the retina. This change allows us to perceive them for a short time.

So if you want to buy into the idea that your eye is some clear window on the world, you need to contend with the fact that this window initially yields a light pattern on the retina that is inverted, left-right reversed, diminished by passage through ill-placed machinery (which ultimately generates a significant receptor deficit), faces partial occlusion by a rich web of blood vessels, only to fall onto a remarkably uneven distribution of varied photoreceptors. So what might that look like?

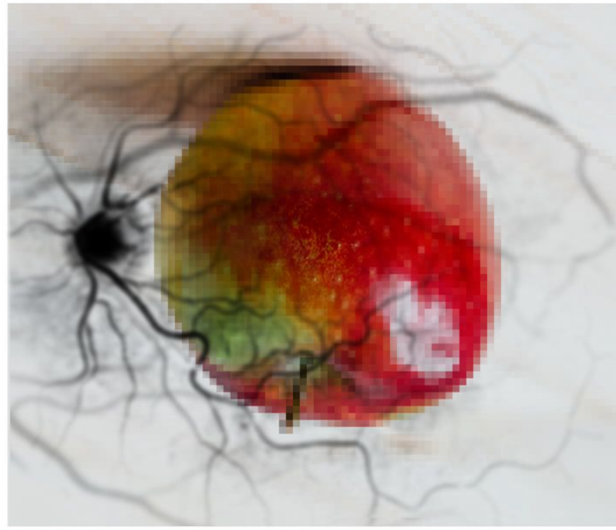

■ infomration defecit

On the left, we see a surrogate representing a live percept while the right simulates how a light pattern that is inverted, left-right reversed, affected by acuity variation due to photoreceptor distribution, diminished by blood vessel interference and a receptor deficit (blind spot), may appear on the retina.

Now many might be quick to intuitively think that this just can't be true. The image I "see" is so rich and detailed, how could all of this "stuff" be in the way? Again, because visual perception does not involve any sort of clear window on the world. Take a moment to consider macular degeneration. This unfortunate, incurable condition is the leading cause of vision loss, affecting more than 10 million Americans–more than cataracts and glaucoma combined. It is the deterioration of the central portion of the retina, known as the macula (which contains the fovea and foveola.) It just so happens that because of the way our brain uses sensory data, the deterioration may simply go unnoticed, especially in cases with spotty macular cell damage or dysfunction, thus leaving many to visit their ophthalmologist only when the disease is fairly advanced. Our brain takes what sensory data it encounters and reflexively responds in a manner that we have evolved to find useful. I can't stress enough how much visual perception is NOT like a camera taking snapshots via light intensity/wavelength measurements… It is just nothing like it.

Oh, and I almost forgot! There really isn't any "color" in this initial projection as color is not a property of the environment. Contrary to what some may believe, we do not sense color. While that may also sound counterintuitive--it is indeed true. Color is the visual experience that arises from our biology interacting with the spectral composition of the light (electromagnetic radiation) that is emitted, transmitted, or reflected by the environment. Our various photoreceptors DO respond differently to different wavelengths of light, thus resulting in an ability to discriminate different wavelengths, resulting in an experience of color vision, but it is just not a component of the early projection. So, perhaps a more "realistic" representation of what is falling on the retina might be this:


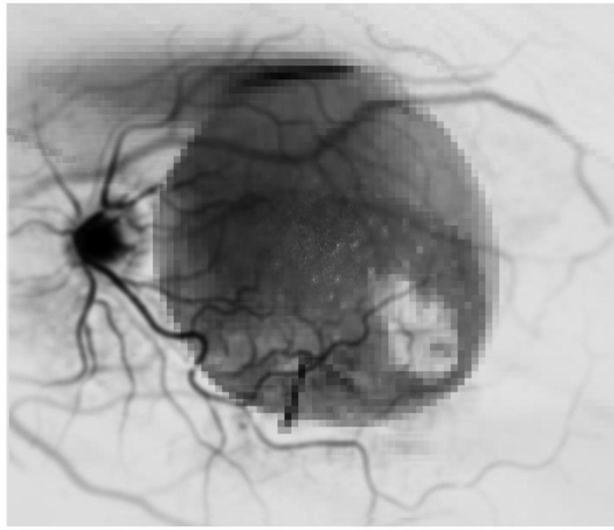

While we are on this color issue, exploring the counterintuitiveness of our vision system, and examining the comparison between a camera and the eye, I'd like to take a moment to share an interesting argument that I encountered not long ago from, believe it or not, a professional artist who was absolutely convinced that color is indeed part of the environment. No matter what evidence I was able to put forward to demonstrate that this was NOT the case, he would not budge from his position. When I asked him to present the evidence that justified his position, he stated that color MUST be a physical property of the environment "because a camera can record it."

Now I am sure that a good number of you reading this have already realized the glaring fatal flaw in this argument, but for many individuals unfamiliar with the fundamentals of color vision and color photography, the argument might seem to have some teeth. However, like most intuitive arguments for a flat-earth, a young earth, or intelligent design, they quickly deteriorate with an increase in scientific literacy and critical thinking. Just to be clear though--with color photography, electronic sensors or light-sensitive chemicals respond to specific aspects of electromagnetic radiation of at the time of exposure. The recorded information is then used to create a percept surrogate by mixing various proportions of specific light wavelengths ("additive color", used for video displays, digital projectors and some historical photographic processes), or by using dyes or pigments to remove various proportions of the particular wavelengths in white light ("subtractive color", used for prints on paper and transparencies on film). The color that you perceive "in" a surrogate, like a traditional photograph or digital image, is being generated by the wavelengths of light emitting from the surrogate-not because the camera "captured" color like some fairy in the garden.

What might make wrapping our head around this even easier is if we take a moment to clearly define two basic terms concerning the manner with which we interact with the environment, sensation and perception. These terms are often used synonymously–but they indeed describe two different aspects of what we normally understand as "experiences." The term sensation describes a low-level process during which particular receptor cells respond to particular stimuli. At the level of sensation, our sensory organs are engaging in what is known as transduction (or in our case, as we mentioned earlier, phototransduction), or the conversion of energy from the environment into a form of energy that our nervous system can use

Perception, on the other hand, can be simply defined as the assignment of “meaning” to a sensation.

So what is actually happening when we visually encounter something that we might understand or describe as “blue”?

Putting aside scenarios in which the object may be an actual source of light or a structural configuration that is bending light, it is likely that the surface of the object is absorbing all of the available wavelengths of the visible spectrum except for some that are relatively short. Now the standard human observer has a specific type of photoreceptor cell in the retina that is “tuned” for such short wavelengths. That means that when this particular cell type encounters this type of wavelength, it responds by initiating a complex cascade of electrochemical events that will eventually lead to more and more complex processes along a particular route we are currently exploring--the “visual pathway.” This low-level cascade initiation is what we define as a sensation.

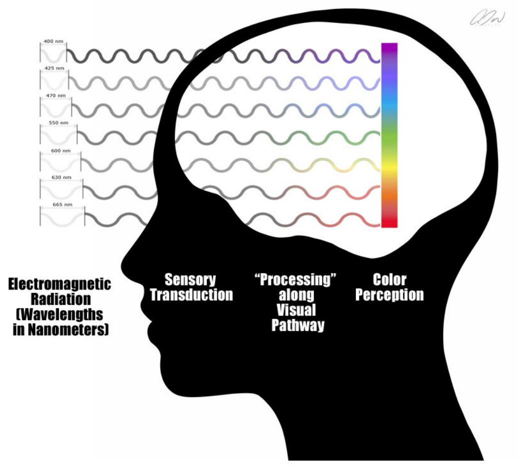

The cascade of events initiated in the retina will eventually lead to specific activity in other, “higher”, or more complex processing regions of the brain such as, but not limited to, the lateral geniculate nucleus of the thalamus, the striate and extrastriate cortex of the occipital lobe, and the temporal lobe (regions which we will be heading to soon in our visual pathway when our walkthrough continues). It is through the aggregate activity of these brain regions that we find a perceptual response–in our case here–”blue.”

So as you might already be starting to suspect, in this example the object in the environment is NOT physically blue, the wavelengths reflected off the surface of the object is not physically blue, nor do the cells referenced above “sense” blue. Blue is not a sensation–rather, it is a perception. We assign “blue” to a particular sensation that is itself a reflexive response to an exposure to a certain wavelength of visible light.

Hopefully, that makes some sense, clears up a few more counterintuitive factors, and allows us to continue our glimpse at a few key points on the visual pathway.

From the retina, we will jump to what is often referred to as the main relay station of the brain known as the thalamus, or more specifically, a region of the thalamus called the lateral geniculate nucleus, or LGN. Information from both motor and sensory systems (with the exception of the olfactory system) relay signals through the thalamus where they are processed before being sent off to a myriad of cortical regions. Yet, The LGN does seem to be much more than a mere relay station or gateway to the visual cortex. It is a multi-layered array of sophisticated microcircuits that would seem to suggest a region of very complex processing. What makes this area even more fascinating is that only 10% of the inputs to this region are coming from the retinas. The other 90% are inputs from a number of cortical and subcortical regions including significant input from the primary visual cortex itself. With what we can observe anatomically, it would seem that this region might be a major site of top-down influence. In other words, it appears that the visual cortex may play a very large role here role in controlling what is actually sent on to--the visual cortex. (Hopefully, at this point, many of you are starting to develop a better appreciation of just how unlike a camera this system really is.)

From here we will exit the LGN via a large fan of axons radiating outward, appropriately dubbed optic radiations, and land on the banks of (and somewhat within) the calcarine sulcus (or fissure) of the occipital lobe. This noticeably striated region of tissue is home to what call the primary visual cortex. Here we find the most studied part of the visual brain.

The part of the occipital lobe that receives the projections from the LGN is known as the primary visual cortex (also referred to as visual area 1 (V1), as well as the striate cortex.) Immediately surrounding this region, above and below the calcarine sulcus are what have been dubbed extrastriate regions. These regions consist of visual areas 2 (V2), 3 (V3), 4 (V4), 5 (V5 or MT) and 6 (V6 or DM). Now before you shudder expecting me to go into a large series of complex processes here--don't worry. I am not. Rather I am going to jump forward to the routes that information takes from here as I believe understanding those paths will be more relevant to our exploration of what "realistic" looks like.

So where do things go from here? Each V1 (remember that we have two halves here) transmits information to two distinct pathways or processing "streams," called the ventral or "what" stream, and the dorsal or "where" stream. The information that is relevant to these two pathways is separate but remains physically integrated through much of the visual pathway. It is when this integrated information reaches "higher" levels that we see a physical separation.
Anatomically, the ventral stream begins with V1, goes through V2, through V4, and then on to regions in the inferior temporal cortex (IT). This stream is often referred to as the "what" pathway as it is associated with form recognition and object representation. As such we find neurons in this stream that respond selectively to signals that might represent particular wavelengths of light, shapes, textures, and at the "highest" levels of this pathway, faces and entire objects. There are a number of regions within the inferior temporal cortex (ITC) that work together for processing and recognizing neural activity relevant to "what" something is. In fact, discrete categories of objects are even associated with different regions. For example, the fusiform gyrus or fusiform face area (FFA) exhibits selectivity for incoming activity patterns linked to faces and bodies while activity in the parahippocampal place area (PPA) helps us to differentiate between scenes and objects. The extrastriate body area (EBA) aids in distinguishing body parts from other objects while the lateral occipital complex (LOC) assists in discrimination tasks regarding the separation of shapes and "scrambled" stimuli. These areas all work together, along with the hippocampus, a dynamic memory region that is believed to be significantly involved in object "compare and contrast" tasks, in order to create an array of understanding of the physical world. This pathway also holds strong connections to the medial temporal lobe (long-term memories), the limbic system (emotion), and the aforementioned dorsal stream. This stream has a lower contrast sensitivity compared to the dorsal and is somewhat slower to respond. However, it does have a slightly higher acuity than its counterpart and carries all information about what will eventually yield an experience of color.

The dorsal stream also begins with V1 and goes through V2, but then travels into the dorsomedial area (V6/DM), middle temporal region (V5 (MT)), and onto the posterior parietal cortex. The dorsal stream, often referred to as the "where" or "how" pathway, is associated with motion and the representation of object locations. As such, within this stream, one would find neurons that show selectivity, not for signals representing shape or light wavelength, but rather for signals that represent location, direction and speed. This stream is essentially "colorblind," but has greater sensitivity to contrast, is quicker to respond (albeit more transient), and has a slightly lower acuity.

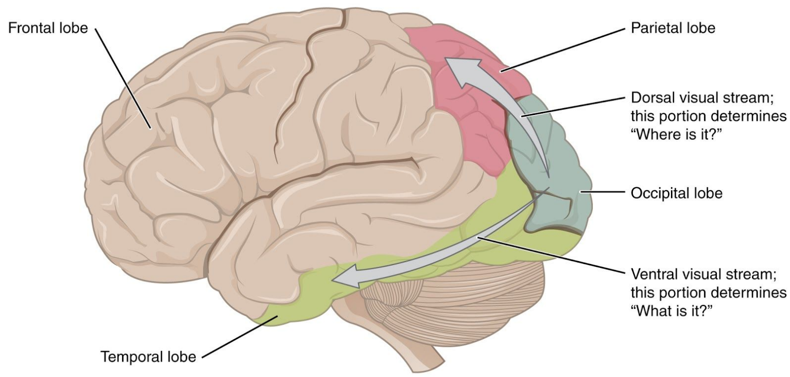

Illustration of the ventral processing stream, or "what" pathway, and the dorsal processing stream, or "where" pathway.

So, as with our blindspot, blood vessels, as well as the spatial imprecision of our periphery, can we see some demonstration of how these two streams process things differently?

You betcha.

One of my favorite demonstrations for this is to introduce a stimulus (object or image) that is visible to only ONE of the streams-thus upsetting their normal operational "balance."

How might we do this? **With color.** Take a look at this image:

Do you notice an odd visual shimmer, jitter, or vibration when trying to read the letters here? This is because that while your "what" system can easily process the color contrast, the colorblind "where" system cannot. In other words, your "where" system is having a heck of a time trying to determine where the boundaries of those letters actually are.

One of the most famous examples of an artist exploiting this phenomenon is Claude Monet's Impression Sunrise. In the piece, many reported that the sun within the image appeared to "vibrate."

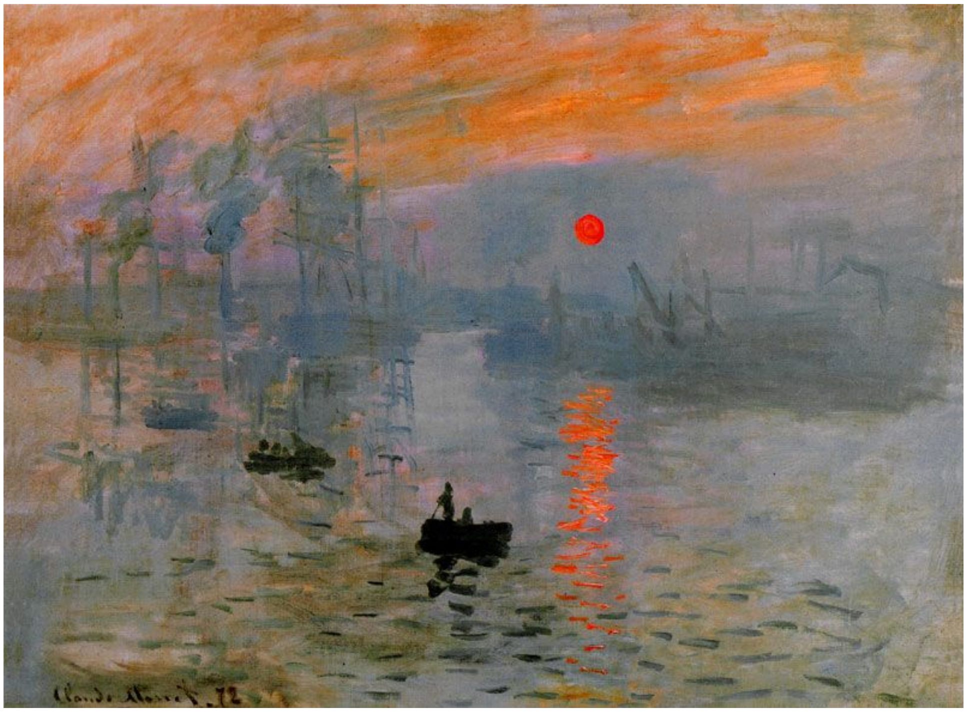

Impression, Sunrise by Claude Monet 48 cm × 63 cm (18.9 in × 24.8 in)
Oil on canvas, 1872

Grayscale version of Impression, Sunrise by Claude Monet.
Notice how the sun is nearly invisible due to its equiluminance
with the surrounding clouds.

Neurobiologist Margaret Livingstone explains the peculiarity of Monet's equiluminant sun in her book Vision and Art: The Biology of Seeing. She writes, "The sun in this painting seems both hot and cold, light and dark. It appears so brilliant that it seems to pulsate. But the sun is actually no lighter than the background clouds, as we can see in the grayscale version. It is precisely equiluminant with–that is, it has the same lightness as–the gray of the background clouds. This lack of luminance contrast may explain the sun's eerie quality: to the more primitive subdivision of the visual system (which is concerned with movement and position) the painting appears as it does in the grayscale version; the sun almost invisible. But the primate-specific part of the visual system sees it clearly. The inconsistency in perception of the sun in the different part of the visual system gives it this weird quality. The fact that the sun is invisible to the part of the visual system that carries information about position and movement means that its position and motionlessness are poorly defined, so it may seem to vibrate or pulsate. Monet's sun really is both light and dark, hot and cold."

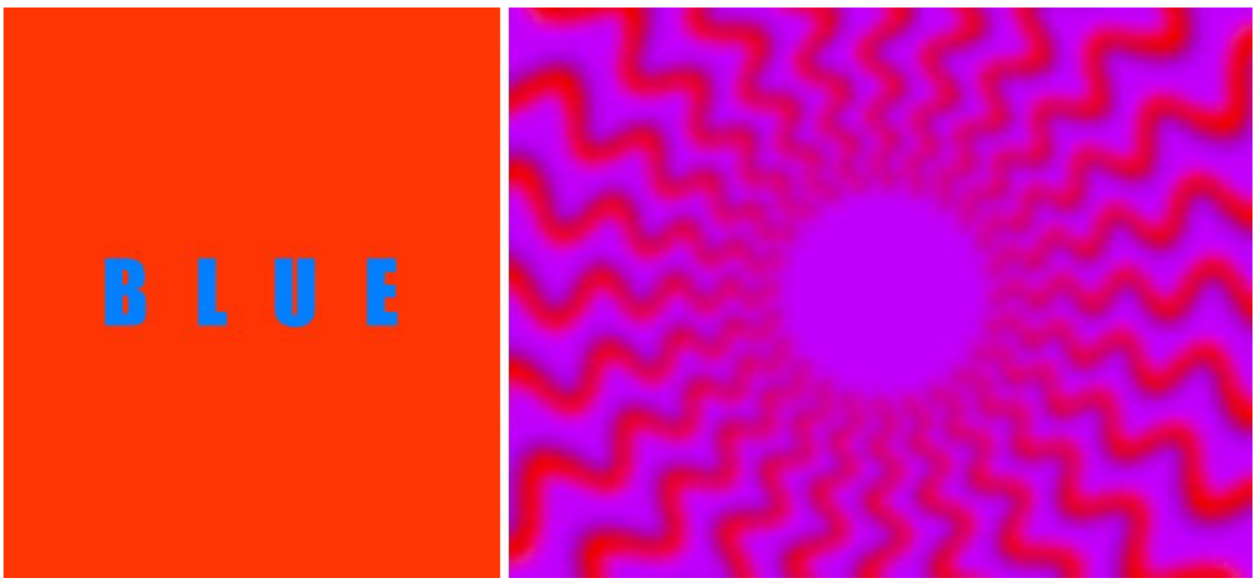

Equiluminance (left) can indeed be a powerful device for achieving a sensation of vibration or pulsation–however, specific variations in contrast, color and element orientation can give rise to even more powerful perceptions of motion such as those created by psychologist Akiyoshi Kitaoka to demonstrate the effect of "Perceptual Drift" (right).

So at this point, I hope that you can better appreciate the fact that our visual system is indeed non-veridical and quite counterintuitive. It does not use any system of objective measurements to yield a percept. Unlike a camera which uses digitized measurements of light intensity and wavelength to allow for the recorded light pattern to be recreated, the visual system uses the adaptive, tuned responses within a neural sheet of photosensitive cells to generate a long and complex cascade of activity, modulated by experience-driven, dynamic biological machinery, which results in a brain state (or series of states) that yields the experience of perception. There is no objective measurement whatsoever contained in the latter. And while sometimes we may find some kinship between what we see and what can be measured objectively, this, as neuroscientist Dale Purves has pointed out, is merely coincidence. In fact, it is important to keep in mind that our visual system has no direct access to the physical world at all and yet, it continues to map on to successful behavior remarkably well.

And while you can find countless demonstrations in textbooks, classrooms, and websites that attest to this fact of non-veridical vision, I could not move on here without sharing at least one of my favorite demonstrations--the simultaneous brightness/lightness contrast demonstration.

Demonstration of simultaneous lightness/brightness contrast. A target (grey square) with a lighter surround (left) is perceived as being darker than an identical target with a darker surround (right). If our visual system operated via objective samplings, one might suspect that there would be no such perceived disparity here.

Now with all of this all in mind, can you understand how the definitions for “realistic” presented at the close of the last section might be, at least in part, problematic? How can something that is deemed realistic “represent things in a way that is accurate and true to life” or be “accurately representing what is natural or real” when our vision system is non-veridical?

Therefore it is my position that when we describe something as looking “realistic,” we are not describing any aspect of real-world properties or even the kinship between our perception of those properties and the properties themselves--rather, we are describing something that I find far more fascinating.

What I believe “realistic” describes is a degree of relative similarity between a perceptual response to a surrogate, simulacrum, simulation, or other representation and the past perceptual responses to the stimulus, stimulus components, or experience that is being represented.

In other words, a realistic image looks like the past.

Part III - The Match Game

In a 2016 essay for Aeon titled The Empty Brain, psychologist Robert Epstein writes, “No matter how hard they try, brain scientists and cognitive psychologists will never find a copy of Beethoven’s 5th Symphony in the brain–or copies of words, pictures, grammatical rules or any other kinds of environmental stimuli. The human brain isn’t really empty, of course. But it does not contain most of the things people think it does–not even simple things such as ‘memories’...We don’t store words or the rules that tell us how to manipulate them. We don’t create representations of visual stimuli, store them in a short-term memory buffer, and then transfer the representation into a long-term memory device. We don’t retrieve information or images or words from memory registers. Computers do all of these things, but organisms do not.”

To further his argument, Mr. Epstein put forward a great exercise that you can try for yourself. For years, Mr. Epstein has been asking his students to try and draw a representation of a dollar bill from memory with as much detail as possible. When the student has finished, he covers the drawing and then has the student repeat the exercise with an actual dollar bill as a reference. To get a general idea of the typical disparities found between the two representations, Mr. Epstein asked a student intern at the time (Jinny Hyun) to share her attempts:

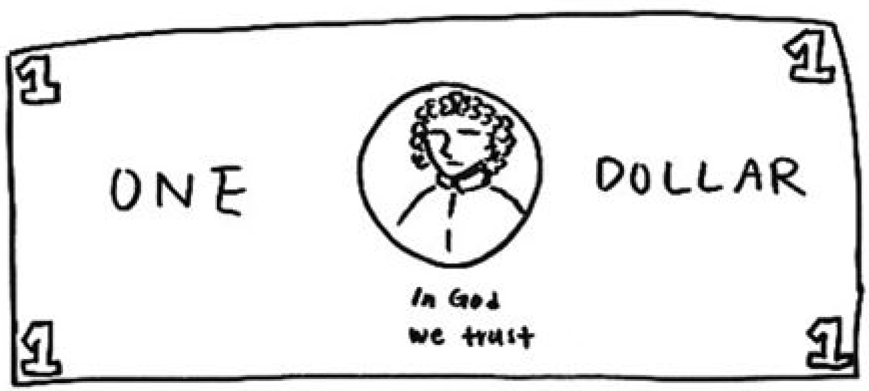

Dollar bill “memory” drawing by Jinny Hyun

Dollar bill "observational" drawing by Jinny Hyun

Mr. Epstein explains, "*As you can see, the drawing made in the absence of the dollar bill is horrible compared with the drawing made from an exemplar, even though Jinny has seen a dollar bill thousands of times. What is the problem? Don't we have a 'representation' of the dollar bill 'stored' in a 'memory register' in our brains? Can't we just 'retrieve' it and use it to make our drawing? Obviously not, and a thousand years of neuroscience will never locate a representation of a dollar bill stored inside the human brain for the simple reason that it is not there to be found.*"

Now while I would indeed agree that we do not carry an objectively recorded, representation of a dollar in some memory register in our brain, that does not mean that there is nothing that we might define or interpret as a representation. In fact, if there was nothing to draw from, why would Ms. Hyun's first attempt be populated with anything? If there was no endogenous information source, shouldn't the first image be completely blank? But the first image is not blank. Ms. Hyun is populating the paper with something that represents, at least to her, a dollar bill.

To understand where the content for the first drawing is coming from you need to understand a little bit about the malleability or plasticity of the brain as well a thing or two about memory.

The brain is incredibly malleable or "plastic" in a great many ways. It is often busy forging new connections, pruning old ones, strengthening those circuits that show great promise and coordination, and diminishing those that are deemed a hot mess. It is this neural maintenance plan that allows us to adapt to a highly dynamic environment. For example, every time you visually encounter an apple, your brain responds with an array of neural activity. Not only does this activity determine what you visually perceive, but it also uses the experience to shape future responses. You see, with continued experience, the perceptual responses that are demonstrated to be useful during an apple encounter are reinforced, while others that are deemed less useful, are not. This will result in a strengthened connection between those neural circuits that have facilitated successful, visually guided interactions with an apple (often across various viewing conditions) in the past. It is here, with these bound and reinforced connections that we find the substrate for a mental representation (an internal cognitive symbol that represents a perception of an external reality.)

Now when someone asks me to picture an apple, I don't perceive an apple magically popping into existence (that would be rather terrifying and not very useful)--rather, I respond with the cultivated neural activity (minus the activity that would result from the presence of the actual stimulus of course) that I have cognitively bound to the concept of "apple." I experience this neural activity as the mental representation. It was this mental representation that Ms. Hyun pulled from for her first drawing.

Wait a minute--- doesn't the system that you describe present a glaring issue for your definition of realistic? For example, If your definition is the degree of relative similarity between a perceptual response to a surrogate, simulacrum, simulation, or other representation and the past perceptual responses to the stimulus, stimulus components, or experience that is being represented---and as Ms. Hyun's drawings clearly demonstrates, we do not hold a complete or accurate recording of our perceptual experiences--how then can we make any such comparisons?

No, this is not a problem whatsoever. We **CAN** interpret the configurations of neural circuits to be a form of "information" storage as it is the manner in which these circuits function that provides us with the content of our perceptions and the basis for our memories. But to understand why our lack of a clear image database is not a problem for my definition of realistic, you need to appreciate that the brain is the source of all human experience. There is no vision, hearing, olfaction, touch, proprioception, taste, emotion, or other personal experience without it. Each experience that we have is the result of complex activities taking place across billions of neurons and trillions of synapses.

With that said, let's return to our example of an apple sitting on a table. The visual encounter with the apple elicits a pattern of neural activity that we will call brain state A. We experience brain state A as "seeing an apple." As stated earlier, if someone asks me to visualize that apple, it does not mean that my brain returns to brain state A (the state that results in an experience of actually viewing the apple) even though my ensuing neural activity would most likely share some of the same patterns of activity found with brain state A. If however, we had the capability to artificially return the brain to a particular state, I would argue that the resulting experience would return as well. In a sense, you would essentially be experiencing "seeing an apple" again rather than merely recalling or visualizing some stored mental representation of it. **And that's the point--we wouldn't need to access some stored photograph in some colloquial sense to return to the experience of the apple--we would just need to return the brain to a particular state.**

Now through experience, my cultivated responses to visually guided interactions with an “apple” have become bound to higher-level responses (something we might look at like the “tags” used in social media) so as to facilitate the continuing assignments of useful meaning to sensations (e.g., recognition, categorization). Therefore, even though I haven’t got that clear photograph of an apple in my head for comparisons, I can experience the activity of those “tags’” during perceptual events to arrive at perceived degrees of similarity.

With all of that that said, I would like to drive home the fact that memory **DOES** play a significant role here. It is important to understand that **BOTH** of Ms. Hyun’s drawings involved different types of memory resources. As I touched upon earlier, the first drawing was created from a mental representation. This representation was comprised of reinforced neural activity that developed over time, with experience. These reinforced representations are often experienced as abstract, spatially imprecise amalgams with the most strengthened attributes or components “appearing” most prominent. These representations can be remarkably long-lasting, and as such, are often associated with a memory resource that is referred to as long-term memory (LTM).

The second drawing was also developed from long-term memory resources but differs in that the contributions were coupled with transient, high capacity memory resources known as iconic and visual short-term memory (VSTM.) Iconic memory is an extremely short-lived form of visual memory which allows us to briefly point our attention to aspects of the experience of a visual response after exposure to a stimulus has ceased (some aspects of the afterimage phenomenon can be attributed to iconic memory.) As one might suspect, the transient nature of this resource tends to limit its content to more low-level features (e.g., size, shape, and color.) Access to the information provided by this visual memory component lasts for a few hundred milliseconds and quickly deteriorates with the introduction of any subsequent activity. The VSTM is considered by some to be an intermediate bridge between Iconic memory and LTM. Lasting for up to a few seconds, it can allow for additional attention to, and maintenance of, residual visual response information, albeit more abstracted, in the absence of an external stimulus.

Now in the context of discussing the “content” of these memory resources, we might be tempted to conclude that we are likely to determine how realistic a representation is by the degree to which it is able to elicit the same (or very similar) low-level responses that we would find appropriate to iconic and VSTM resources during, or immediately/briefly following exposure to, the percept being represented. Now, this does make a great deal of sense. I mean, it is true that we are far more likely to find a representation created solely from LTM resources to be far less realistic, or far more abstract, than a representation created from LTM, iconic and VSTM resources together--which is EXACTLY what Ms. Hyun's drawings demonstrate. However, with that said, to limit our definition to the recreation of low-level responses appropriate to iconic and VSTM resources may lead us to a significant under-appreciation of the higher-level cognitive contributions to our judgments about “realistic.”

Once again, to be clear--my working definition for realistic is wholly compatible with the nature and function of these biological systems as we currently understand them. Again, my aim is not to communicate that the assignment of realistic is determined by how well a perceptual experience aligns with the mental representations that we hold (even though such activity may be required to facilitate the perception of the similarity), nor do I think it is adequate to focus only on low-level responses, rather it is my assertion that the assignment is determined by how close we find the perceptual experience resulting from a surrogate or other representation to be to the perceptual experience elicited by past exposure to the stimulus that is being represented.

In any case, all of this talk of dollar bills reminds me of a popular riddle from the 1930s that is often referred to as “The Missing Dollar”. It goes something like this, “Three people check into a hotel room. The clerk charges $30 for the room, so each guest contributes $10. It is only later that the clerk realizes that there is a special room rate in effect so the guests should have only been charged $25. To rectify this, he gives the bellhop $5 to return to the guests. On the way to the room, the bellhop realizes that he cannot divide the money equally, and as the guests didn't know the total of the revised bill, the bellhop decides to just give each guest $1 and keep the remaining $2 as a tip for himself. So if each guest got $1 back, they ultimately paid $9 each for the room.

But if each guest paid 9$ (totaling $27) and we add the $2 taken by the bellhop--we have $29 accounted for. What happened to the last dollar?”

I’ll give you a hint: ***It’s in the same drawer that your brain keeps all the images in.***

Part IV - Novelty and Other Cognitive Stuff

So let's recap some important points thus far:

1. Vision is non-veridical. Unlike a camera, our visual system does not "produce images" via absolute measurements of any kind. Rather, our visual system works by having patterns of light on the retina trigger reflex patterns of neural activity that have been shaped entirely by the past consequences of visually guided behavior over evolutionary and individual lifetimes.

2. Configurations of neural circuits can be regarded as a type of information storage which is responsible for our perceptions as well as our memories. However, while mental representations indeed shape our perceptions, they are not "stored" snapshots of reality, nor do they allow us access to the neural activity that is experienced during an actual perceptual event.

3. Point 1 & 2 demonstrate that many existing definitions for realistic are fundamentally at odds with the nature of perception.

4 .The visual information available to us from different memory resources is quite different. (e.g., in general, a drawing created from LTM resources (mental representations) will differ greatly from a drawing created from LTM, Iconic and VSTM resources (less abstracted information that holds more kinship with the brain state of an actual perceptual event.)

5. My usage of the term realistic does not describe how well an experience aligns with the mental representations that we hold, but rather describes how closely (relatively) the brain state/perceptual experience resulting from a surrogate or other representation is to the past brain states/ perceptual experiences elicited by the stimulus that is being represented.

While we have indeed covered a good amount of ground in our exploration of realistic so far, we still have yet to tackle one of the most interesting aspects of the assignment--specifically, our tendency to apply the descriptor to seemingly novel percepts or objects. Now some of you reading this might already be thinking that my definition's reliance on past experience gets me into a rather large pickle here.

For example, in regards to the Cambridge Dictionary definition speaker mentioned in the first section who determined that a representation of a hatching dinosaur egg was realistic---how can he/she determine whether or not the representation was realistic if there was no past visual experience of such eggs hatching?

To understand how we navigate novelty in this context, and why the aforementioned lack of direct dino-hatching experience is not a big problem for us here, we need to appreciate how novelty interacts with our penchant for prediction and the HUGE difference between the manner in which percepts are shaped by past experience and the way we use past experience to arrive at conclusions about those percepts.

First, it is important to understand that massive amounts of neuronal resources in the human brain are devoted to predicting what will happen from moment to moment. This fact has led many to regard the brain as a dynamic prediction machine. Some would even go so far as to state that the entire neocortex has evolved to do just that. In his book On Intelligence, Jeff Hawkins writes "*Your brain receives patterns from the outside world, stores them as memories, and makes predictions by combining what it has seen before and what is happening now. Prediction is not just one of the things your brain does. It is the primary function of the neocortex, and the foundation of intelligence.*" As such, we seem primed to actually seek out novelty so as to become better predictors. Some parts of our brain even seem to become quite active in the presence of novelty. Take our substantia nigra/ventral tegmental area or SN/VTA for example. This midbrain region, responsible for regulating our motivation and reward-processing, is closely linked to areas regarded as heavy hitters in processes like learning and memory (e.g., hippocampus and the amygdala) and has been shown to respond far better to novelty than to the familiar.

Second, past experience in regards to visual perception refers to the past consequences of visually guided behavior over **evolutionary and individual lifetimes**. (*And no, before you even think it, this does not mean that some humanoid ancestor was running around with dinosaurs, watching their eggs hatch and storing a biological MPEG for us to access to better appreciate certain scenes in movies.*) What this means is that over time, selective pressures have given rise to certain neural configurations and behaviors that shaped the perceptions of our ancestors, and continues to shape how we perceive the world today. Simply put, visual perceptions are reflexive manifestations shaped not solely by our personal history or development--but also by our evolution. On the other hand, our decision making routines are governed by cognitive mechanisms that have evolved to generate successful behavior in spite of limited data. We refer to such mechanisms as cognitive heuristics and cognitive biases. Let's take a look at these cognitive mechanisms and see how they might help us deal with novelty.

At the very end of Part II, I presented you with two sentences describing what I believe the term realistic actually describes. The first sentence there is very specific but incredibly wordy, while the second is far more succinct but incredibly vague. Now while this may initially seem like a silly question, of the two, which did you prefer? Which one do you think you would be more likely remember?

I would venture to say that most would prefer (and remember) the latter. Aside from the fact that it is almost always easier to remember less information than more, a big factor at play here is what psychologists have dubbed a fluency heuristic.

A heuristic is a sort of mental shortcut for use with problem-solving, learning, or discovery that employs a practical method with seemingly efficient rules not guaranteed to be optimal or perfect but are relatively sufficient for immediate goals. These shortcuts, which often involve focusing on one aspect of an issue, object, or concept, and ignoring all others, can work well in many circumstances, but they can also lead to systematic deviations from logic, probability, or rational thinking. Such deviations often result in the development or promotion of irrational influences on our behaviors called “cognitive bias.”

The errors in logic and judgment that can result from heuristics are called "cognitive biases.” And even though biases are often labeled as irrational, carrying negative connotations in some contexts, they can also prove to be very useful in others (e.g., allowing us to reach decisions quickly in the presence of danger.)

In the case of the fluency heuristic, efforts to ease cognitive workload lead to preferences for that which can be processed or understood more fluently, faster, or more smoothly compared to more complex alternatives. The heuristic promotes a method of evaluation that increases focus on the aspect of processing ease while pushing us to ignore all else. The resulting bias is the belief that material that carries a lower cognitive cost carries a higher intrinsic value.

A very similar heuristic, dubbed the “Keats” heuristic, seems to share traits with the fluency heuristic as well as the resulting bias. With Keats though, we use the perceived aesthetic properties (e.g., eloquence) of a statement as a measure of its truthfulness or accuracy. One of the resulting cognitive biases from this heuristic is what is called the “rhyme as reason” bias. This bias is the belief that a statement written to rhyme is inherently more truthful or accurate. For those of you thinking that something this illogical cannot possibly carry significant influence, consider the effectiveness of Johnnie Cochran’s mantra of “If it doesn’t fit, you must acquit!” during the 1995 murder trial of O.J. Simpson.

Now while I would argue that the fluency heuristic indeed plays a role in the higher-level cognitive contributions to visual perception, others may hold the key to understanding how our brains deal with more novel stimuli.

Two heuristics worth our time and effort here are the known as the **availability heuristic** and the **representativeness heuristic**. The availability heuristic is a process in which we use or adapt readily available, specific information to form beliefs about distant or more general concepts that we deem comparable. Additionally, in a fashion similar to our fluency heuristic, the description of the availability heuristic process are sometimes found to include a weighted consideration of available information based on the ease with which certain information is brought to mind. Examples of this heuristic can be found everywhere in our lives, from calculating the likelihood of a train derailment to buying a lottery ticket. For example, many people that are subjected to news stories about train derailments may be likely to believe that such events are far more typical than they truly are. Another example can be found with individuals purchasing lottery tickets as it is far more likely that the purchaser's mind is filled images of celebrating a colossal win as opposed to any actual statistics regarding the likelihood of winning. As such, you can see how the availability heuristic may account for people being swayed by a big, flashy story instead of a large body of scientific evidence.

The representativeness heuristic is a process in which we apply the perceived properties of a group, class or category prototype to each assigned member. (The term prototype here is used to describe an example or standard that exhibits the essential features or greatest representativeness of a particular category.) While availability involves specific instances, representativeness has much more to do with the prototypes, stereotypes, or averages. Psychologists have argued that many judgments relating to likelihood, or to cause and effect, are based on how representative one thing is of another, or of a category. We tend to feel far more confident about this process when a target in question appears similar to our prototype. While this shortcut can be useful when evaluating a relatively novel target, it can also lead to a significant amount of irrational stereotyping.

So how might examining these heuristics help us better appreciate how we arrive at assignments of realistic?

Let's start with the dinosaur eggs that I mentioned earlier. When I typed "realistic dinosaur egg hatching scene" into Google, I was treated to a wide variety of images including this one from the movie Jurassic World.

Still image from Universal Studios Home Entertainment (2015). Jurassic World.

Even though I have never seen an actual dinosaur hatching from an egg, I would be very comfortable in describing this representation as a realistic depiction of just that. Why? Because the perceptual experiences elicited by this surrogate hold a significant kinship with the perceptual experiences I have had from encountering similar stimuli (e.g., forms and surfaces) in the real world. In other words, the low-level components (and a number of their configurations) comprising this image hold significant kinship to many past experiences even though the perceived “whole” of the image may not. Furthermore, the aforementioned cognitive heuristics (e.g., availability and representativeness) pull from experience-cultivated, higher-level cognitive resources (neural activity associated with eggs, animals hatching from eggs, categorical information about lizards/reptiles, etc.,) to yield an array of assumptions and/or predictions about what a dinosaur hatching from an egg WOULD or SHOULD look like. It is this magnificent synergy between seemingly disparate neural operations that allows us to simultaneously apply the familiar to the novel while using the novel to build future familiarity. It is an amazingly adaptive system that can allow something utterly foreign to appear quite realistic.

But how about the turning this around? Can something real look unrealistic?

Absolutely.

Again, we need to keep in mind (pardon the pun) that our assignment of realistic has nothing to do with the actual properties or even existence of real-world objects. It’s all about neural activity.

For example, here is a post that I borrowed from a colleague’s (Katherine Nelson’s) social media page.

I was captivated when I first saw this image and doubly so when I read the caption. This is a great example of what I mean when I state that something real may not necessarily appear realistic. As opposed to the image of the dinosaur hatching above, with this picture both my colleague and I (and possibly you) are finding it difficult to establish significant kinships between the perceptual experiences elicited by this surrogate and the perceptual experiences resulting from encountering relatively similar stimuli in the real world. Some low-level components (and a number of their configurations) comprising this image are indeed familiar, as with the hatching dinosaur image, but not quite enough to warrant the same assignment of realistic (to the stimulus being represented). Both Katherine and I do not have enough information to deem what this object **WOULD** or **SHOULD** look like in reality, thus her comment “..it hardly looks real.”

Now you need to be careful here in that you do not quickly conclude from the example of this fascinating sea creature that realistic is somehow a result of object recognition (i.e., the creature does not look realistic because I cannot recognize/categorize the object.) I would passionately argue against this idea as we can find many things realistic that we just cannot be quick to recognize or categorize. Just like the dinosaur egg example above, as long as perceptual experiences elicited by a surrogate hold a significant kinship with the perceptual experiences I have had from encountering similar stimuli (e.g., forms and surfaces) in the real world, I can experience the representation as realistic. Even if my higher-level cognitive mechanisms are engaged in a series of speculations about how something WOULD or SHOULD look like due to the novel nature of the subject, I can still find it quite realistic.

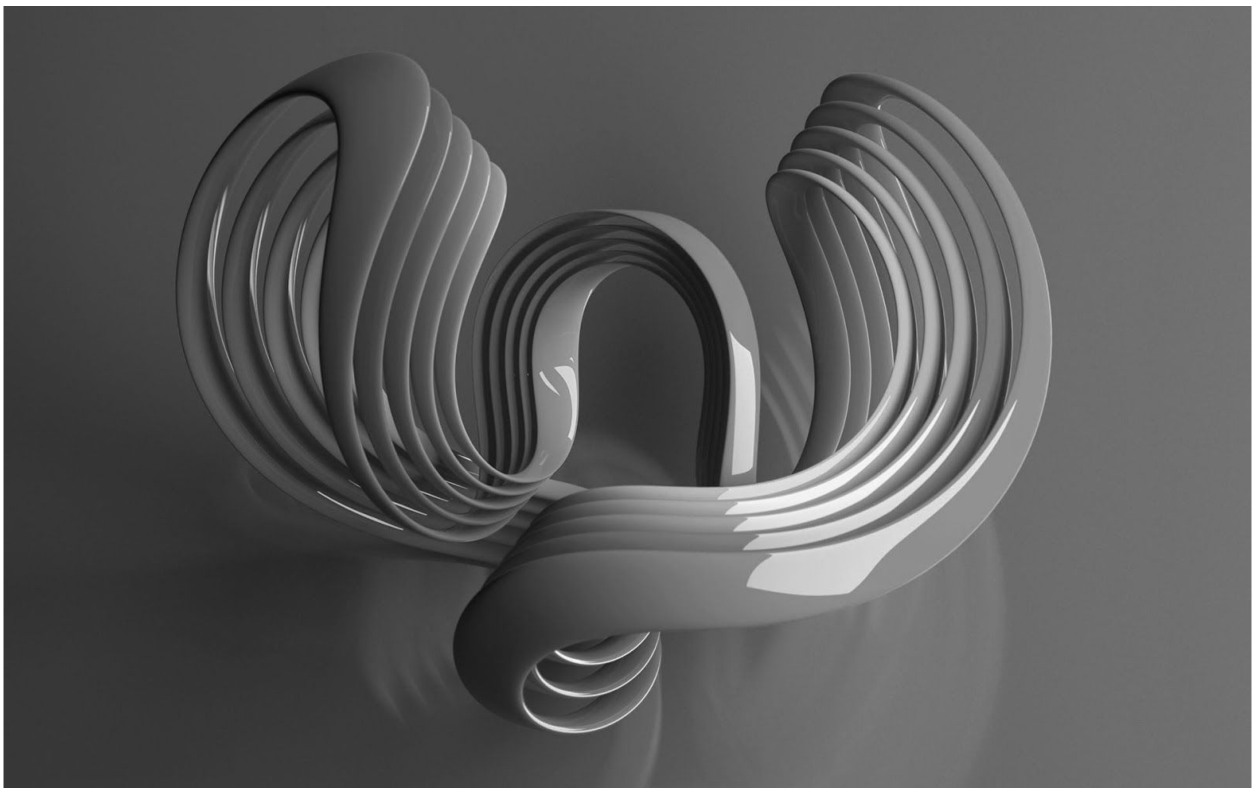

Take this representation of an abstract, biomorphic form for example. I don't recognize the subject matter as anything specific other than an abstract, biomorphic form, but I immediately found the configurations/relationships of the visual components holding significant kinship with past experiences of particular surfaces, forms, etc. Enough of these kinships exists so that I find it to be a realistic representation of something I simply don't recognize.

An interesting aside to this would be an interesting phenomenon connected to object recognition known as canonical perspective. In 1981, researchers Stephen E. Palmer, Eleanor Rosch, and Paul Chase began to systematically investigate into the influence of perspective effects on object characterization. These studies involved rating tasks in which participants rated how much a representation (e.g., a photograph) looked like the actual object Here is what they found:

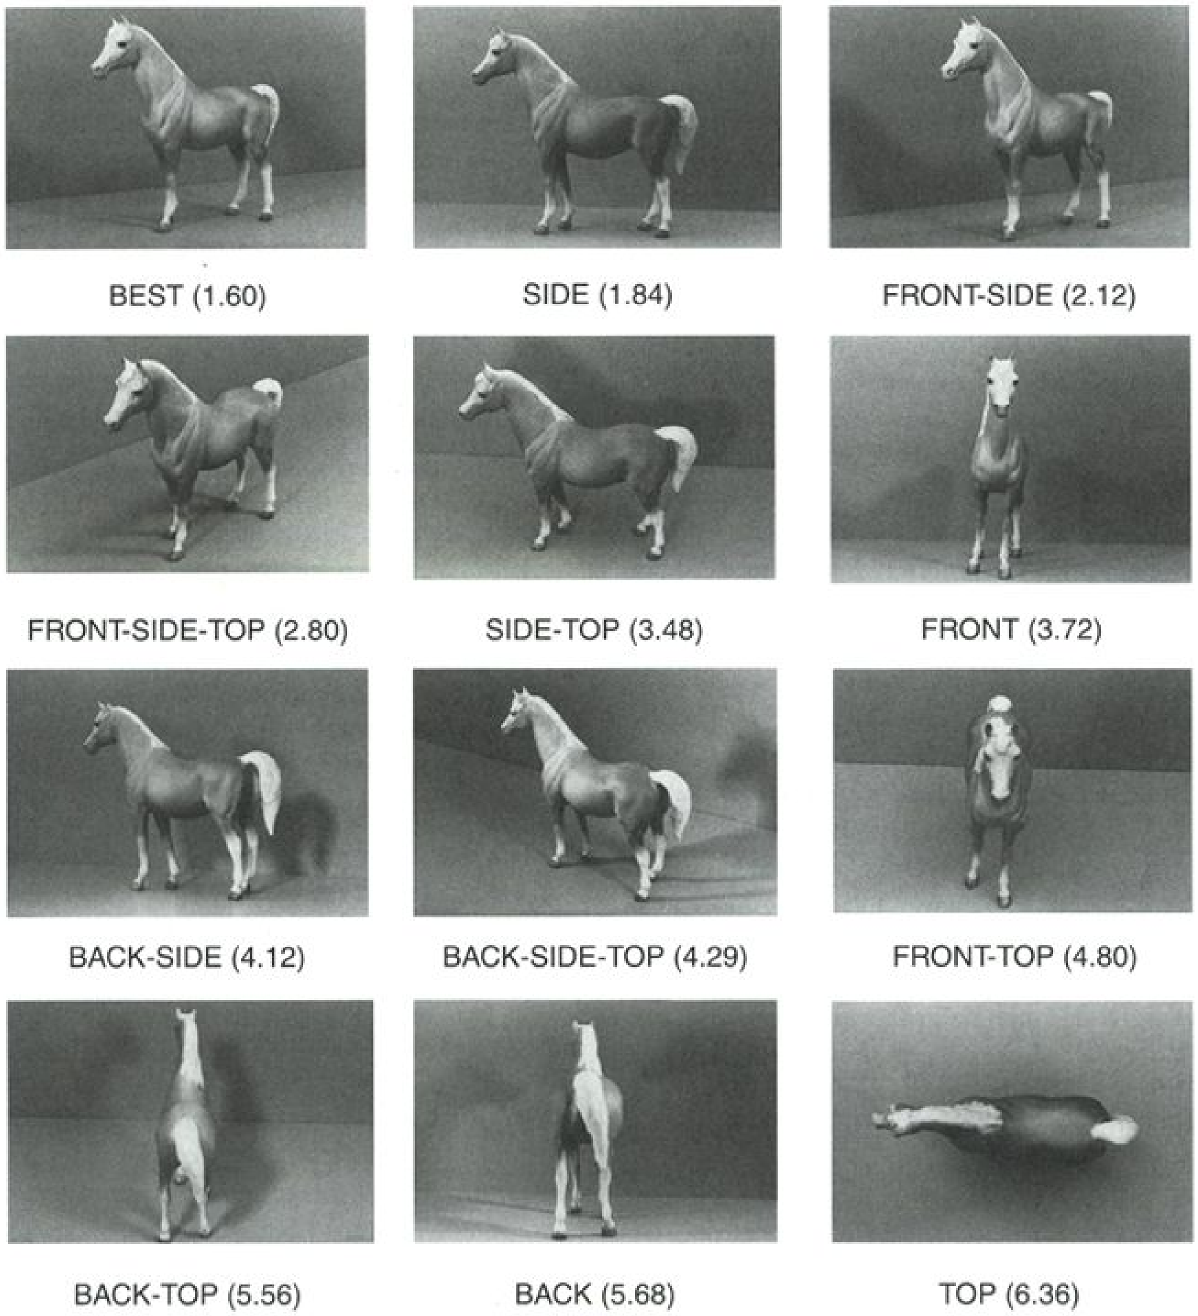

Canonical Perspective results from Palmer, Stephen E. Vision science: Photons to phenomenology. MIT press, 1999. (fig 9.2.4)

Perspective views of a horse differ significantly in how much they look like the object they depict. The numbers in the image above represent average ratings by study subjects on a scale from 1 (looks very much like the object) to 7 (looks very unlike the object.)

Palmer, Rosch, and Chase then had subjects name the entry-level categories for represented objects as quickly as possible with the same array of photographs (taken from various perspectives.) Results from that stage of the study showed that the pictures that were rated as “best” representations earlier were successfully identified/categorized fastest and that naming latencies increased as ratings from the previous stage declined. In other words, the representations rated “worst” were named far more slowly than the ones rated “best.”

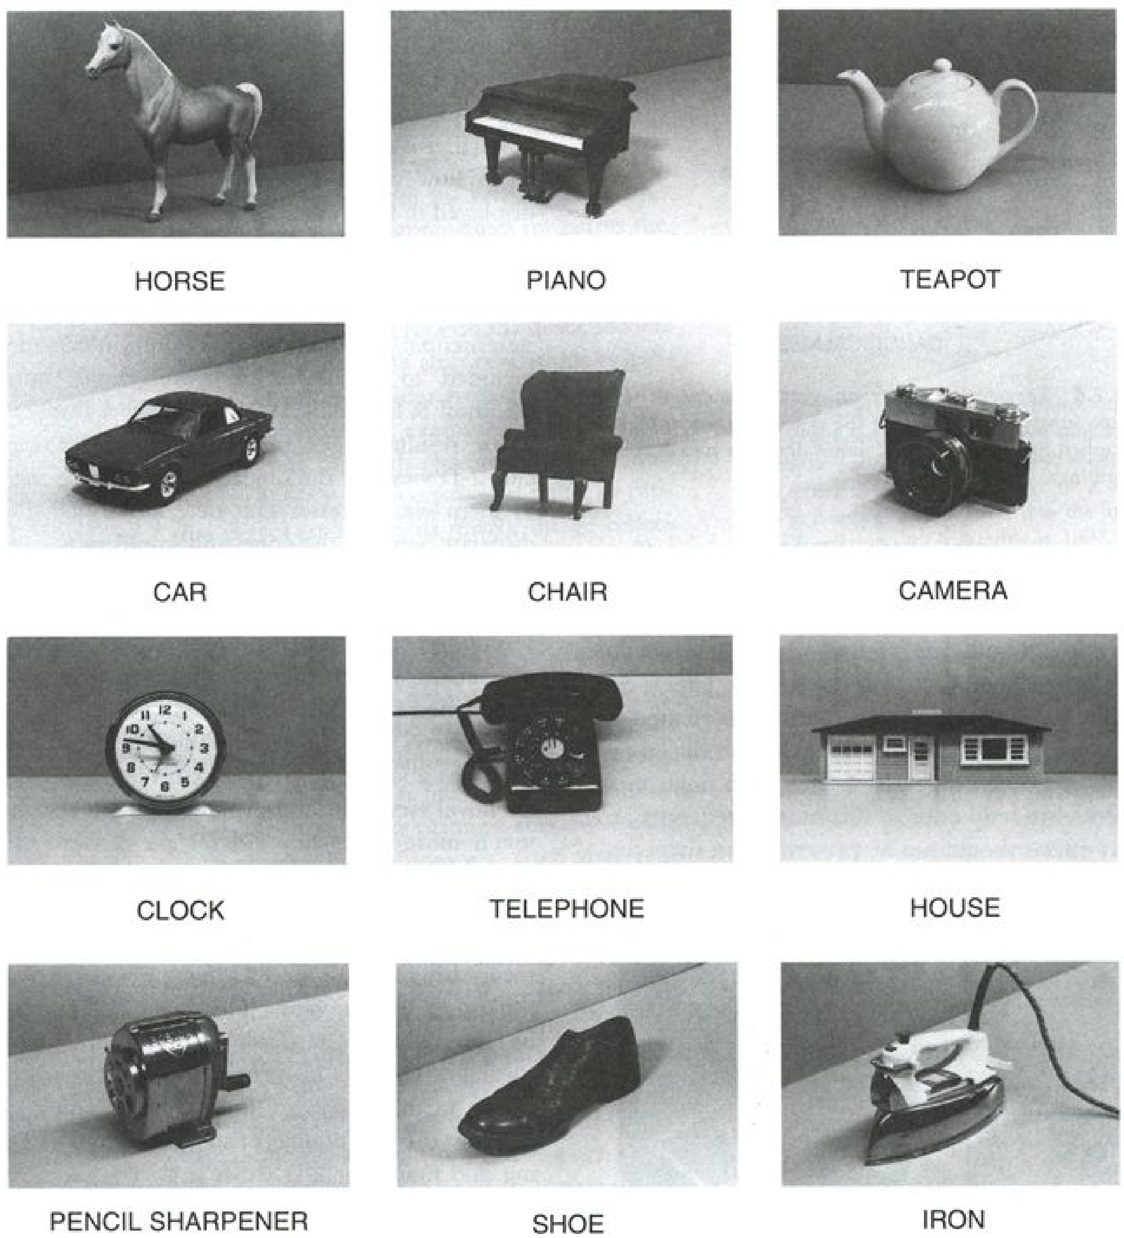

Canonical Perspective results for 12 objects from the from Palmer, Rosch, and Chase study. Image taken from Palmer, Stephen E. Vision science: Photons to phenomenology. MIT press, 1999. (fig 9.2.4)

Palmer, Rosch, and Chase dubbed these “preferred” perspective views canonical perspectives. While there are a number of hypotheses put forward to explain the latencies and preferences found with this study. Two of the most obvious were explained in Stephen E. Palmer’s book, Vision science: Photons to phenomenology. He writes, “*1.Frequency Hypothesis. One possibility is that the speed of naming is simply a function of the number of times we see the objects from those viewpoints. By itself, this explanation is unlikely because there are several cases in which it gives the wrong predictions. Cups, for example, are very often seen from directly above, especially as we guide them to our mouths in the act of drinking, yet this view was rated as quite a poor view and was slow to be identified...2.Maximal Information Hypothesis. A second possibility is that the perspective effects simply reflect the amount of information different views reveal about the shape and use of the object. [This] seems contradicted, however, by the fact that there are some objects for which the view of a single surface is the best view, such as the purely frontal view of the clock.*” He goes on to state, “*It is likely that both hypotheses contain some measure of truth and that the perspective effects that Palmer et al. reported depend jointly on both."*

Since the original 1981 study, experiments have continued to explore this phenomenon. While many continue to find strong preferences for a general perspective view that can be described as slightly above and at an approximately 30 to 45-degree angle from the front of the object--possibly offering the most surface information, other experiments continue to demonstrate that people will just get used to whatever perspective they're shown most often--and will, in turn, come to be most familiar with the object from that perspective. In my opinion, these variations in results simply reflect the dynamic nature of what we might deem “useful.”

Now with all of that said, do these perspective effects affect our assignments of realistic?

Well let’s look at the “best” and “worst” horse representation from Palmer’s experiment:

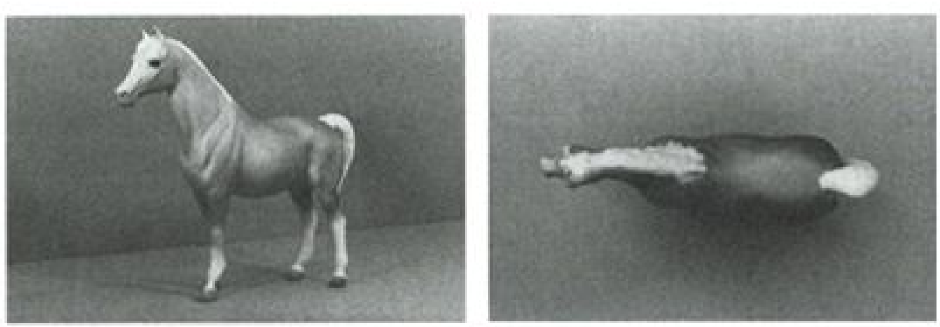

Would you deem one representation to be more realistic than another? I have a hard time saying yes to this one.

How about this?

While some of these skulls might be recognized as a skull faster than others, do you find any one to be any more realistic than another?

The takeaway here is that I cannot see very strong evidence that higher-level recognition or categorization systems significantly impact our judgements regarding what appears more or less realistic. While these systems absolutely help us to contend with novelty, it seems quite possible to create representations that yield perceptual experiences that hold significant kinship with past responses to real-world configurations of surface, form, texture, etc. The amalgam of these components may yield something novel, but that does not seem to greatly impact how realistic we might deem the representation.

Part V - Visiting the Valley

As with all jaunts into the hard sciences, explorers are often faced with the realization that all is not as it might seem. Does the sound of that make you feel a bit unsettled? If it has then I have successfully begun to set the stage for a fascinating phenomenon that seems to plague those chasing down more-realistic percept surrogacy. In fact, this rather bizarre phenomenon easily has the power to lead a developing percept surrogate right off of a cliff…

..and into a valley.

Take a look at this image:

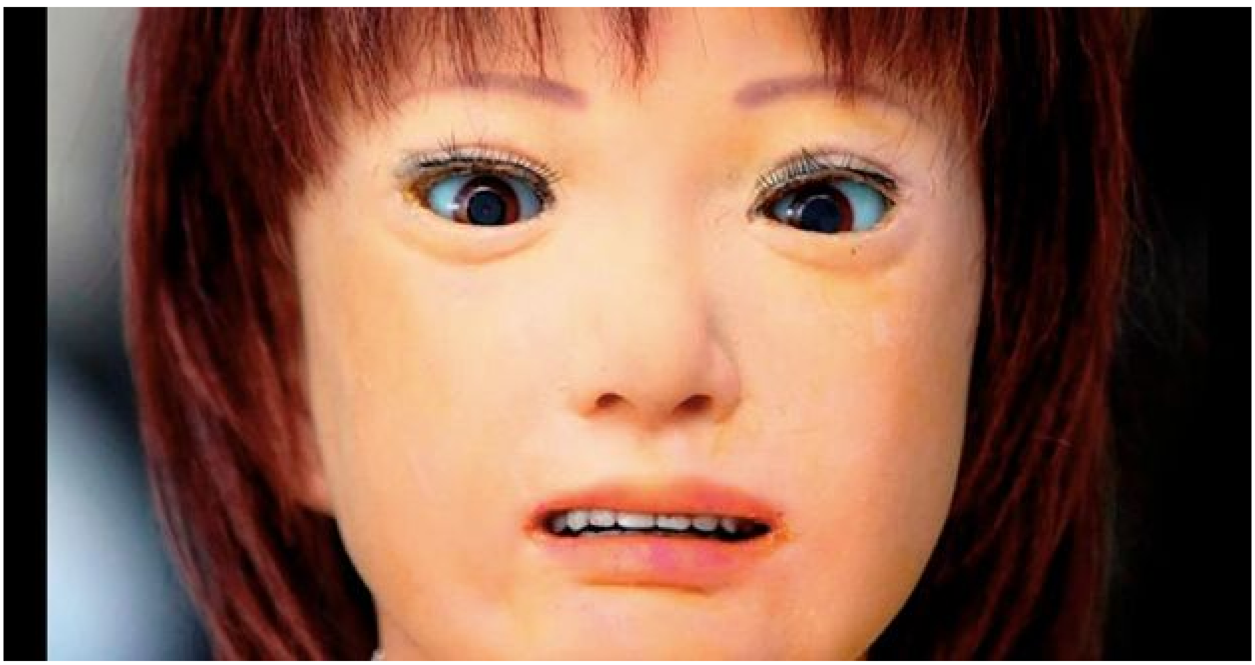

Does this creep you out at all? Do you find this somewhat unsettling? Well, hopefully, you aren't scared off just yet as it's about to get interesting…

The reason that many may find this image unsettling (or just downright creepy) is due to the phenomenon known as the "Uncanny Valley." In aesthetics, the uncanny valley is the hypothesis that human replicas which appear almost, but not exactly, like real human beings elicit feelings of eeriness and revulsion among some observers. "Valley" denotes a dip in the human observer's affinity for the replica, a relation that otherwise increases with the replica's human likeness. The concept was identified by the robotics professor Masahiro Mori in 1970. The term was first translated as uncanny valley in the 1978 book Robots: Fact, Fiction, and Prediction, written by Jasia Reichardt, thus forging an unintended link to Ernst Jentsch's concept of the uncanny, introduced in a 1906 essay "On the Psychology of the Uncanny."

In his essay, Jentsch explains that a person with the feeling of the uncanny is “not quite at home.” He goes on to describe a man sitting on an old tree trunk, only to find the trunk moving, and suddenly turning into a snake. The “uncanny” is that time in between the “trunk” and the “snake.”

So why should we find a near-but-not-quite human representation so potentially bothersome?

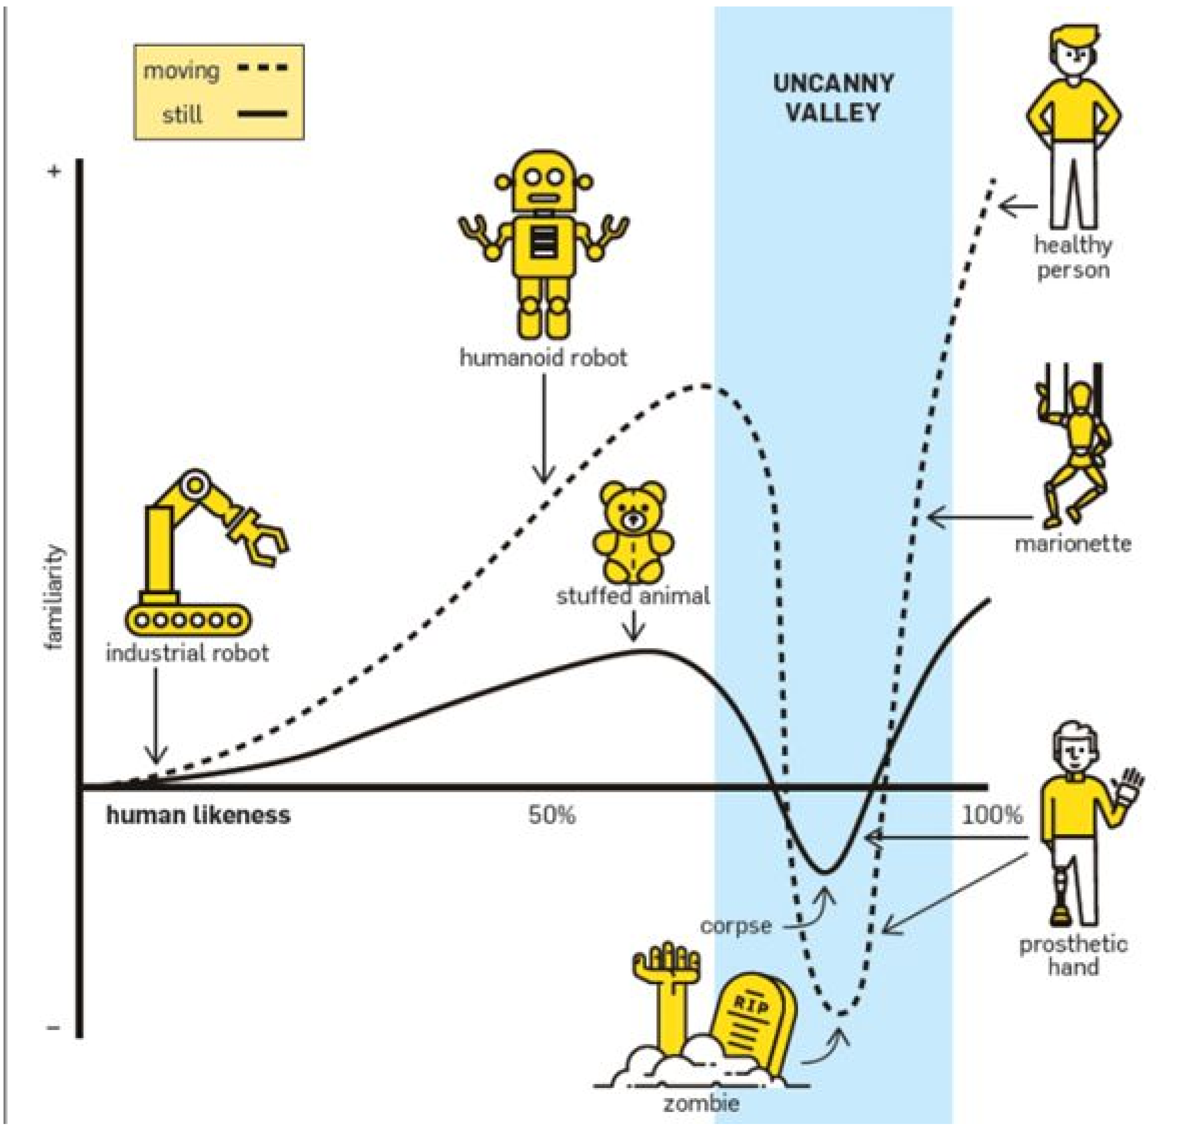

Source: ACM, 2016. Hypothesized emotional response of subjects is plotted against anthropomorphism of a robot, following Masahiro Mori’s statements. The uncanny valley is the region of negative emotional response towards robots that seem “almost” human. Movement amplifies the emotional response.

There are a number of hypotheses with some pretty good explanatory power. Here are a few of the most common:

1. Mate selection. Automatic, stimulus-driven appraisals of uncanny stimuli elicit aversion by activating an evolved cognitive mechanism for the avoidance of selecting mates with low fertility, poor hormonal health, or ineffective immune systems based on visible features of the face and body that are predictive of those traits.

2. Mortality salience. Viewing an "uncanny" robot elicits an innate fear of death and culturally-supported defenses for coping with death's inevitability.

3. Pathogen avoidance. Uncanny stimuli may activate a cognitive mechanism that originally evolved to motivate the avoidance of potential sources of pathogens by eliciting a disgust response. "The more human an organism looks, the stronger the aversion to its defects, because (1) defects indicate disease, (2) more human-looking organisms are more closely related to human beings genetically, and (3) the probability of contracting disease-causing bacteria, viruses, and other parasites increases with genetic similarity."

4. Sorites paradoxes. Stimuli with human and nonhuman traits undermine our sense of human identity by linking qualitatively different categories, human and nonhuman, by a quantitative metric, degree of human likeness.

5. Violation of human norms. The uncanny valley may "be symptomatic of entities that elicit a model of a human other but do not measure up to it". If an entity looks sufficiently nonhuman, its human characteristics are noticeable, generating empathy. However, if the entity looks almost human, it elicits our model of a human other and its detailed normative expectations. The nonhuman characteristics are noticeable, giving the human viewer a sense of strangeness. In other words, a robot stuck inside the uncanny valley is no longer judged by the standards of a robot doing a passable job at pretending to be human, but is instead judged by the standards of a human doing a terrible job at acting like a normal person. This has been linked to perceptual uncertainty and the theory of predictive coding.

6. Conflicting perceptual cues. The negative effect associated with uncanny stimuli is produced by the activation of conflicting cognitive representations. Perceptual tension occurs when an individual perceives conflicting cues to category membership, such as when a humanoid figure moves like a robot or has other visible robot features. This cognitive conflict is experienced as psychological discomfort (i.e., "eeriness"), much like the discomfort that is experienced with cognitive dissonance.

Now while it may seem that the Uncanny Valley is a concern limited to representations of the human form, you need first to appreciate the amazing tendency we have to anthropomorphize (to ascribe human form or attributes to) a vast multitude of objects. In much of my writings on pictorial composition, I reference this tendency often in regards to how prediction tasks may govern much of our experience with a picture.

I remember when I first took notice of our collective propensity to anthropomorphize objects while teaching a college still life class in the late 90s. When soliciting direction from the class regarding the arrangement of subjects for the day, students would refer to an object without a salient front or face (like an orange) and say—"No, face it the other way."

Give it a try for yourself:

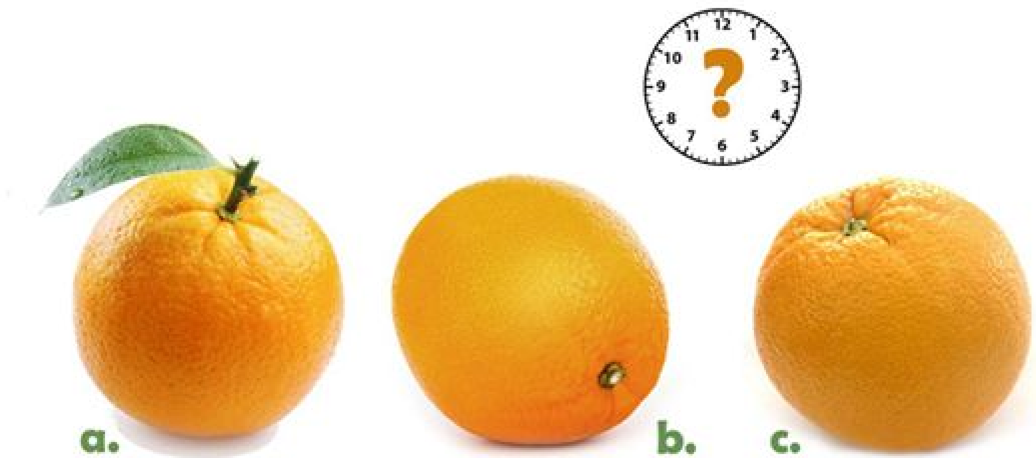

Take a quick look at these three oranges. Limiting yourself to answers via standard clock positions, which way do you think the following oranges (a, b, and c) are "facing?"

Did you come up with some directions? If so, you are anthropomorphizing these oranges. (You might see where this is going at this point.) So if we experience the "uncanny" in those representations of humanity that get close but just fall a hair short—could we experience it with objects that we are assigning human attributes to? I think so–at least to a degree.

If we substitute a simple subject like an apple into the uncanny valley graphic that we used earlier, we can see that much may still hold true. From the more utilitarian representations up to the more realistic percept surrogates, we seem fine until we hit the cliff with a rise in attributes that may seem overly artificial or generic—something that just takes us one step back from a natural appearance.

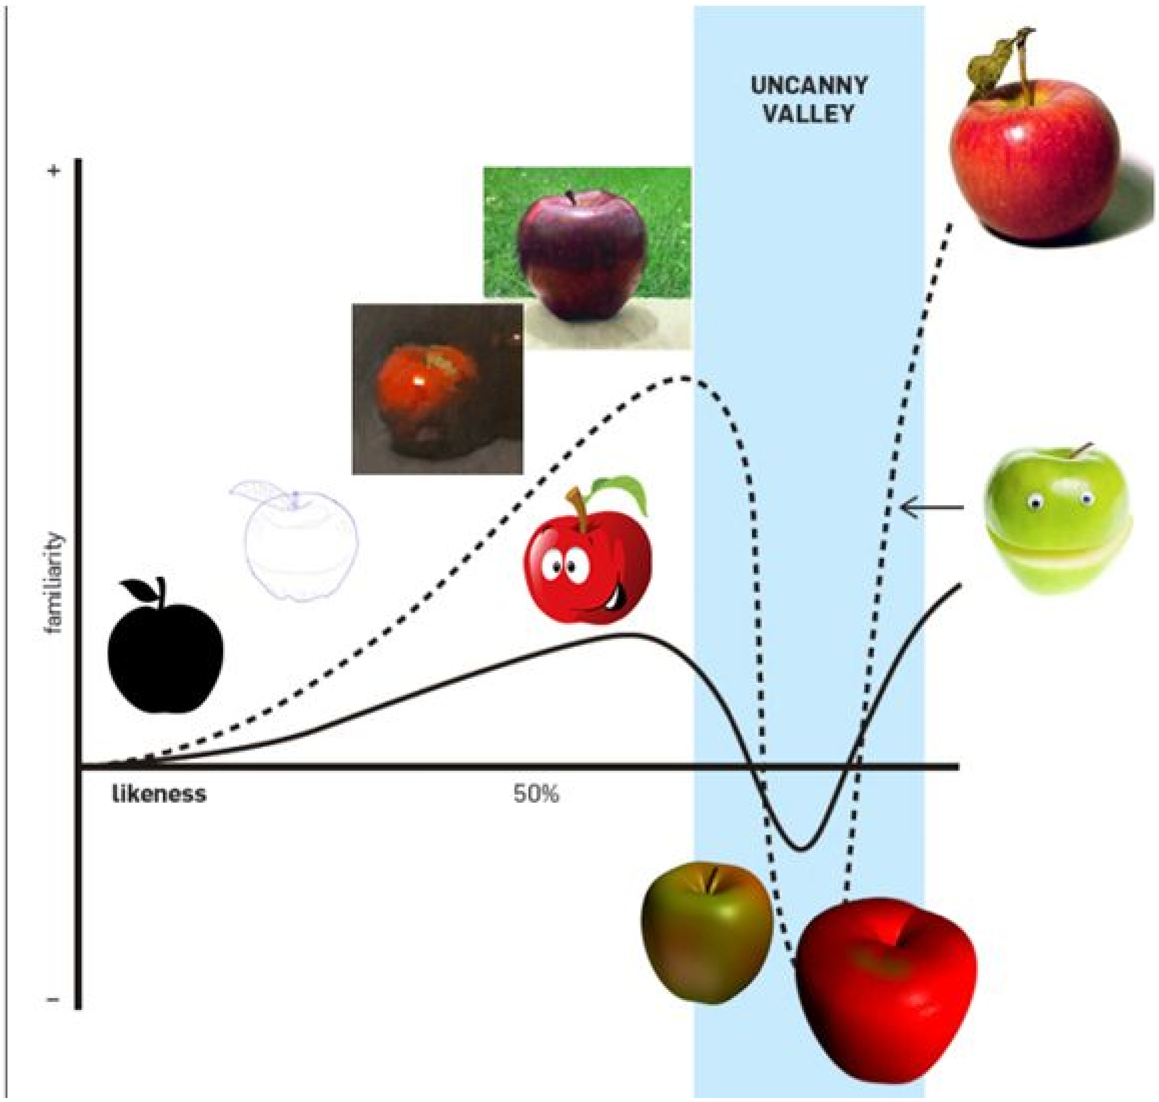

So what does this mean for our assignments of realistic?

I wanted to visit the valley here for a few reasons. First, it is a truly fascinating phenomenon that is often encountered (to some degree) when artists or designers are pursuing realistic surrogates.

Second, I think that the exploring the valley in this context is incredibly useful here to help establish or to better understand, the difference between an assignment of realistic and aesthetic preference. These concepts often get conflated, especially when discussing forms of art that involve realistic percept surrogacy (e.g., photorealism or hyperrealism.) Remember that aesthetic qualities are those characteristics of a stimulus that elicit adaptive responses that have evolved to reinforce or discourage specific behaviors. We may prefer one type of sensory experience over another—describing one as repulsive and the other beautiful. The realistic nature of a representation may indeed amplify or attenuate our responses to such characteristics, but I find it unhelpful to regard “realistic” as an independent channel here. For example. If we look at the above charts with representations of apples and the human figure, the valley does not necessarily involve a dip in our assessment of realistic, but rather demonstrates an aesthetic response to certain stimuli that are nearing realistic.

The third reason for my introducing of this phenomenon into our exploration is that my experience suggests that as representations grow more realistic, the gain or weight of visual component relationships is increased in such a way that any exaggeration along a single dimension can significantly alter the context of all dimensions--thus potentially steering the greater web of neural activity in a direction that one seeking a realistic surrogate may not want. As we discussed with novelty, it is not just the independent visual components themselves that are generating the brain state and resulting perceptual experience, but the relationships between them.

Part VI - Realism versus Realistic

In many of the conversations that I have about creating realistic percept surrogates, I notice that many are quick to use the terms realistic and realism interchangeably. Now in some contexts, this might be fine--but I often address it to make sure that there are not correlations being made that are unhelpful.

If we were to look up "realism" today we might find some very general definitions like this one: "*Realism in the arts is the attempt to represent subject matter truthfully, without artificiality and avoiding artistic conventions, implausible, exotic, and supernatural elements*." This seems not too far off from our concept of realistic, but as you might have noticed it includes that pesky term "**truthfully**" (which is another term for **veridical**) and I really don't think I need to address that any further.

However, there is an art **MOVEMENT** that is also referred to as Realism. ***Wikipedia*** provides the following description of the movement:

"*Realism was an artistic movement that began in France in the 1850s, after the 1848 Revolution. Realists rejected Romanticism, which had dominated French literature and art since the late 18th century. Realism revolted against the exotic subject matter and exaggerated emotionalism and drama of the Romantic movement. Instead, it sought to portray real and typical contemporary people and situations with truth and accuracy, and not avoiding unpleasant or sordid aspects of life. Realist works depicted people of all classes in situations that arise in ordinary life, and often reflected the changes brought by the Industrial and Commercial Revolutions. The popularity of such "realistic" works grew with the introduction of photography—a new visual source that created a desire for people to produce representations which look objectively real.*

IMAGE RIGHT: Jules Bastien-Lepage, October, 1878, National Gallery of Victoria

The Realists depicted everyday subjects and situations in contemporary settings, and attempted to depict individuals of all social classes in a similar manner. Classical idealism and Romantic emotionalism and drama were avoided equally, and often sordid or untidy elements of subjects were not smoothed over or omitted. Social realism emphasizes the depiction of the working class, and treating them with the same seriousness as other classes in art, but realism, as the avoidance of artificiality, in the treatment of human relations and emotions was also an aim of Realism. Treatments of subjects in a heroic or sentimental manner were equally rejected.

Realism as an art movement (and an aesthetic) was led by Gustave Courbet in France. It spread across Europe and was influential for the rest of the century and beyond, but as it became adopted into the mainstream of painting it becomes less common and useful as a term to define artistic style. After the arrival of Impressionism and later movements which downgraded the importance of precise illusionistic brushwork, it often came to refer simply to the use of a more traditional and tighter painting style. It has been used for a number of later movements and trends in art, some involving careful illusionistic representation, such as Photorealism, and others the depiction of "realist" subject matter in a social sense, or attempts at both."

I am sure that you can see some of the common threads between the general textbook definition of realism, my definition for realistic and the Realism movement described here. There ARE similarities, but there are also important differences that can make synonymous use highly problematic. For example, both the textbook definition for realism and the description of the Realist movement involve variables of "truth" and "accuracy" which makes sense in a general context--but as I hope I have clearly demonstrated, misrepresents the nature of human perception, and as such, misrepresents what is responsible for our assignments of realistic to a percept surrogate. Additionally, much writing on the Realist movement describes realism as an "aesthetic." Again, this can make sense in the context an artistic movement, however, as I mentioned in the section on the Uncanny Valley, conflating the nature of realistic with an aesthetic quality can often lead to more confusion than useful insight.

Closing - Final Thoughts

Albert Einstein once wrote that, “*The eternal mystery of the world is its comprehensibility…The fact that it is comprehensible is a miracle*.” This great line was later shortened to “*The most incomprehensible thing about the universe is that it is comprehensible*.” In either case, this sentiment echoes my own feelings in regards to our ability to recognize anything within a highly dynamic environment with a visual system that does not accord with the physical properties of sources, nor any features of the conflated light patterns that fall upon our retinas *(let alone our ability to use that system for communicating ideas and experiences to another mind across significant oceans of time and space.*)

Well, at this point I hope that I have been able to provide you with a useful framework for understanding what “realistic looks like.” It might sound counterintuitive at times, but I can assure you that it is highly compatible with what we currently understand about perception, cognition, and visual communication. In any case, I will continue to investigate and test the concepts that were put forward here as I hope you will. I look forward to the new insights that may reside on the horizon and the ensuing experiments that will be carried out at the easel. I will continue to try and make new marks in the hopes of discovering new ways to elicit the responses that I seek. In other words, I will continue to try and make more realistic dinosaurs every chance I get. In the end, I'm still that excited little kid wearing down countless crayons drawing dinosaur after dinosaur. I just wear bigger shoes now.

Happy art-making to you all!

<3

SELECT RESOURCES:

Deregowski, Jan B. “Pictorial perception and culture.” Scientific American(1972). Nov.:82-88.Hudson (1960)

Dictionary, Cambridge. "Cambridge dictionaries online." (2015). (https://dictionary.cambridge.org/us).

Dictionary, Merriam-Webster. "Merriam-Webster." ***On-line at http://www. mw. com/home. htm*** **(2002).**

Dictionary, Oxford English. "Oxford English dictionary online." ***Mount Royal College Lib., Calgary*** **14 (2004).**

Epstein, R. "The Empty Brain. Aeon." (2016).

Gefter, Amanda, and Donald D. Hoffman. "The Evolutionary Argument Against Reality." ***Quanta Magazine*** **(2016).**

Hawkins, Jeff, and Sandra Blakeslee. ***On intelligence: How a new understanding of the brain will lead to the creation of truly intelligent machines*****. Macmillan, 2007.**

Hoffman, Donald D. ***Visual intelligence: How we create what we see*****. WW Norton & Company, 2000.**

Jentsch, Ernst. "On the psychology of the uncanny (1906)." ***Angelaki: Journal of the Theoretical Humanities*** **2.1 (1997): 7-16.**

Ishii, Y., Watanabe, T., Akiyama, M., & Mori, T. (2016, March). Clone or relative?: Understanding the origins of similar android apps. In ***Proceedings of the 2016 ACM on International Workshop on Security And Privacy Analytics*** **(pp. 25-32). ACM.**

Livingstone, M. Foreword by David H. (2002). Vision and art : the biology of seeing. New York: Harry N. Abrams. ISBN 0-8109-0406-3.

Lotto, Beau. ***Deviate: The Science of Seeing Differently*****. Hachette Books, 2017.**

Palmer, Stephen. "Canonical perspective and the perception of objects." ***Attention and performance*** **(1981): 135-151.**

Palmer, Stephen E. *Vision science: Photons to phenomenology*. MIT press, 1999.

Purves, D. (2012). *Neuroscience*. Fifth Edition, Sunderland, Mass: Sinauer Associates.

Purves, Dale, and R. Beau Lotto. *Why we see what we do: An empirical theory of vision*. Sinauer Associates, 2003.

Reichardt, Jasia. *Robots: Fact, fiction, and prediction*. London: Thames and Hudson, 1978.

Rossi, E. A., et al. "Imaging retinal mosaics in the living eye." *Eye* 25.3 (2011): 301.