The Convergence of Parallel Medical Technologies

A philosophical analysis of the question of what happens when two parallel medical imaging technologies do not converge in their diagnostic conclusions

Wouter Ruben van Dijk

University of Twente

PSTS

Supervisor: Prof. Dr. Ir. Peter-Paul Verbeek

Second Reader: Prof. Dr. Ir. Mieke Boon

Contents

Abstract

This thesis compares and contrasts two different medical imaging technologies, x-ray mammography and photoacoustic mammography, in order to analyse the problems that appear when two different technologies are used to study the same object.

The main question answered in this thesis is: "What happens when two parallel medical imaging technologies do not converge?" In discussing this topic and the technologies involved, several sub-questions appear, which I also attempt to answer.

"How do Imaging technologies influence the way we see the patient?"

"How are medical images interpreted, especially when there are multiple seemingly valid interpretations?"

"Is there such a thing as a disease itself, if our access to the disease is always mediated?"

The thesis answers these questions by going through three different topics, each building on the ones that came previously. The second chapter, after the introduction, deals with technological mediation of the physician's access to his patient, with the help of the works of postphenomenological philosophers such as Peter-Paul Verbeek. The third chapter deals with expertise and tacit knowledge, which is important for the interpretation of medical images. This chapter heavily depends on the works of Collins and Dreyfus. The fourth chapter deals with epistemological issues, first through Don Ihde's Instrumental Realism, and eventually settling on Boon and van Baalen's concept of Epistemological Responsibility.

The thesis ends by answering the questions asked in the introduction, giving attention to the way in which the medical of the context demands solutions which would not appear in a purely philosophical discussion.

Chapter 1: Introduction

1.1: Introduction

In 1901, Wilhelm Röntgen was awarded the first Nobel Prize in Physics, for his discovery of the X-ray, a particle that allowed people to see through the human body. Still quite well known is the famous picture of his wife's hand, bones and wedding ring clearly visible.

This technology made what was seemingly impossible, possible. It allowed a physician to look inside his patient's body, without having to cut it open. Suddenly, a world of possibilities opened up. Today, X-rays are joined by a host of different imaging modalities, each allowing the user to look inside the body of his patient in a different way. The subject of this thesis is the similarities, differences, and interactions between these differing imaging modalities. Different technologies give us different ways to look inside of the body, so what happens when these technologies come into conflict? What happens if one technology tells us one thing, but another technology disagrees? What happens if one technology tells us that a given patient is sick, while another technologies' diagnosis is that the patient is healthy? Will the end result be that one technology takes precedence over the other, or will a third technology, perhaps a more invasive or expensive one, have the final say over the reality of the patient?

Which brings us to the question: "What happens when two parallel medical imaging technologies do not converge?" Do people look for a technological explanation behind the differences? Has one of the results been misinterpreted? Are the two technologies not looking at the same thing after all? What happens in a hospital? In a laboratory?

1.2: The case

For this thesis, I will focus on two different technologies, both intended for the same purpose, doing drastically different things. However, that is not to say that the ideas discussed here are only applicable to the comparison of those technologies. Neither are they specific to the situation described here.

The specific situation I will focus on is that of breast cancer screening. The idea behind the process of breast cancer screening is that early detection of breast cancer increases the probability of cure, and that mammography has the ability to reduce breast cancer mortality when implemented in a population-based screening program[1]. Thus, many countries have instituted these screening programs, using x-ray mammography to screen at-risk populations.

However, these programs have their problems, socially, medically, and financially. X-ray mammography brings with it a minimal amount of risk, that is almost negligible compared to the positive effects of screening, but it is still a risk. There are also other problems, such as over-diagnosis and overtreatment. That is to say, the screening detected breast cancer in a healthy patient, and the following treatment damaged the patient.[2] In addition to this, there are issues of cost effectiveness. A screening program costs money, money that could also be used for different measures. How effective is breast cancer screening compared to what it costs? Then, there is the last issue, the social effects of breast cancer screening. Besides calculus regarding the medical benefits of screening, there is also a social dimension to it. True positives, false positives, true negatives and false negatives have consequences on the lives of the participants in the screening, and with a wide screening program, this means that the technology used in the screening has a large effect.

In other words, breast cancer screening is a hot topic, and a lot of research is done into the area. This research is not limited to x-ray mammography and it effects, but also includes possible alternative technologies that do not have some of the problems mentioned above. There are technologies such as, low-dose mammography, Magnetic Resonance Imaging, contrast agents, Ultrasound Imaging, and Molecular imaging.[1] Besides using familiar techniques in this new context, there are also research trying to use new and upcoming technologies for the purpose of breast cancer screening.

Researchers at the University of Twente have been working on one such technology, called photoacoustic imaging. They are currently using their Twente Photoacoustic Mammoscope (PAM) in clinical trials, to study how their device holds up in vivo. They have found that they can visualize breast malignancies with high imaging contrast, and that the photoacoustic contrast is independent of mammographically estimated breast density.[3] That is, a breast which is dense to x-ray mammography, and difficult to image, is not necessarily difficult to diagnose using photoacoustics.

One of the main advantages of the technology however, is its non-invasiveness. Where x-ray mammography is minimally invasive because of the low amounts of radiation involved, photoacoustic measurements do not use harmful radiations, and do not require painful breast compression for measurements either. In other words, if and when the technology completes its development, it is a possible candidate to replace the role of x-ray mammography in breast cancer screening.

1.3: The main question

As such, we have two different technologies, for the same purpose, they can both diagnose breast cancer. So what happens when one technology tells us that a patient is sick, while the other technology tells us that the patient is healthy? The question is: 'What happens when two parallel medical imaging technologies do not converge?"

To answer this question however, we must first answer several subquestions. I will answer these subquestions, as well as the main question, in three steps. The first of these steps is to look at technological mediation, and answer the question of: "How do Imaging technologies influence the way we see the patient?" The second step is to look at the interpretation of medical images, or rather, at the expertise of the person doing the interpretation. The question then, is "How does the expertise of the interpreter influence the way medical images are interpreted?" The third step is more ontological, and has to do with the disease that is being imaged itself. The question is: "Is there such a thing as a disease itself, if our access to the disease is always mediated?" These three steps and questions will be handled in order, in chapters two, three, and four.

1.4: Step one, Mediation

The first step to answering our main question about parallel technologies is to look at the technologies themselves. What is it that makes them different? First, and most obviously, the two technologies are based on very different principles. Photoacoustics works by calculating the location of chromophores such as haemoglobin, and finding the cancer through that[3], while x-ray mammography uses differences in the breasts absorption and scattering of x-rays. Thus, we have two very different technologies that are used to do the same thing, breast cancer detection. How can this be? The answer to this question lies in the fact that the two technologies are aimed at different facets of the tumour. Photoacoustic measurements detect the changes in vascularization created by malignant tissue, while x-ray mammography can detect structural changes, as well as micro-calcifications, which can be early signs of breast cancer.[4] While the two technologies seem to do different things, there is an overlapping theory that explains how they both detect features of the same object, a tumor.

This brings us to a different problem, one that is important in answering the question. The problem is that, through technology, we do not see this tumor as a whole. Instead, we only specific features of the tumor. In other words, not only can we only access the tumor through medical technology, we can only access part of the tumor. The question then, is what the tumor actually is, because it seems that what a tumor is, is highly dependent on the technology used to look at it, in more ways than one. This is not just limited to scientific definitions, but goes into the social context as well. To illustrate, a tumor detected through x-ray mammography is a possibility, a likelihood that someone may have a disease. There is a high false positive rate, and early detection means that the chances of survival are high. The specifics of the technology influences the meaning of the diagnosis. In contrast, a biopsy has far more concrete results, and can be far more damning for the patient. A positive test-result gathered by biopsy changes the reality of the patient. No longer is she going to the hospital to check the result of a screening, she is now a sick patient.

That is not to say, however, that technology decides what the tumor means, and how we see it. Different medical imaging technologies give access to the tumor in different ways. That is to say, a physician sees different things through different technologies. Most medical technologies do not offer a diagnosis, instead, they give their user information with which to create such a diagnosis. We cannot reduce the process of diagnosis to the technology used in it. The physician role matters as well, and must therefore also be discussed.

In 'What Things Do: Philosophical Reflections on Technology, Agency, and Design' [5] philosopher Peter-Paul Verbeek introduces the concept of mediation. The argument is that technology is neither *Instrumentalist* nor *substantivist*. To elaborate, the instrumentalist view is that technology is neutral, a tool to achieve human ends. The substantivist view is that technology is not neutral, but rather a controlling feature of a society. To elaborate, think of the saying 'Guns don't kill people, people kill people.' Something often said by opponents of gun control in the United States. This is an Instrumentalist view. The gun is neutral, a means to an end. It is the person holding the gun that is killing people. The substantivist view then, is that the technology is responsible for the killing of people. Had the technology not been there, deaths would have been avoided, and it is the technology that is responsible for them.

Mediation theory then, is a postphenomenological theory that places itself somewhere in the middle. It holds that the relationship between the human and the world is never direct, and always mediated. Technology is one of those mediators, and in this case the important one. X-ray mammography mediates between the physician and the body of his patient. It shapes the relationship, but does not control it. The technology is not neutral. Had the physician chosen a different technology to mediate the relation between him and his patient, he would have seen something else. However, the technology is not the determining factor either. The patient, the physician, the world around them, these all play roles.

Thus, if we look at different medical imaging technologies, we can see that they mediate in different ways, but it is clear that this is not the end of the story. Apart from the technology, there is also the physician, the observer using the technology. This medical expert must interpret the information given to him through the technology. The problem here, is that technology is multistable. What that means, is that there is more than one way to interpret the technology. Or in this case, specifically, that there is more than one possible diagnosis for a medical image. A medical image can be interpreted in different ways, without one of them being obviously superior.

To interpret the image, and to choose which possible interpretation to believe in, is not a neutral process. It is something that is influenced by the knowledge and expertise of the person doing the interpretation. Two observers can see different things, and there is not always an obviously correct answer.

Here, one of our conundrums begins to appear. Technology is not neutral, it mediates between human and world, and different technologies mediating between a physician and a tumor can do so in different ways. At the same time, the use of the technology is dependent upon the user. In other words, technology has an influence on the human, and the human then has an influence on the technology. The question is, if all these things influence each other, how do we get real access to the tumor? Is it possible to see the object in a way unmediated by technology, unaffected by the specifics of the observer? To get to the essence of a tumor itself?

And, on a more philosophical level, is there such a thing as the essence of a tumor? And if so, is it even possible for us to have information about such a thing? Or is the tumor itself wholly constructed by the technology and the user of the technology?

1.5: Step two, Expertise

The first step, as explained above, is looking at technological mediation. However, this only tells half the story. Technology not only mediates, its results are also interpreted. As said above, medical images are multistable, and what they mean is not immediately obvious. The question then, is how does this interpretation work, and how do we make sure we have the correct interpretation. The answer, at least to the latter question, is that we leave it to the expert, which leads us to the question: what is an expert? The role of the physician in the interpretation of medical images is critical.

One of the things I explain in this thesis is a hypothetical quantified version of [6]. Ultimately, the idea behind this technology is to use it for local quantification of chromophores such as haemoglobin as a result of angiogenic processes, which can be correlated to tumor progression. In other words, the idea is that the technology will supply a number, which, cross-referenced with literature values, can supply a diagnosis.

This technology is of interest to our case, because it directly contrasts the other technologies discussed in how the results are interpreted, or rather, whether or not it is interpreted. Unlike the other two technologies, this interpretation is done beforehand. It is not the physician gaining mediated access to the tumor. Instead, it seems like the technology is supplying information. Yet it is not neutral, the processes are simply hidden from view. The expertise has been codified into the technology, made explicit, in order to have a more efficient device.

But is that all there is to expertise? And how exactly does expertise change the way in which people interpret medical technology? That is the question I wish to answer in the second step, in chapter 3, Expertise. The chapter will handle different viewpoints on expertise, and explain how expertise can change someone's mesoperceptual framework, the way in which he looks at the world, as influenced by his knowledge, skills and experience. The idea being that an expert will see different things, look at things through different lenses, and have different theories about how things work.

1.6: Step Three, Epistemology.

The third step then, is to return to the issue of looking at the tumor through technology, and how knowledge about the tumor can be gained. I first do this by taking a look at Don Ihde's theory of Instrumental Realism[7], the idea that true knowledge about the essence of a scientific object can be attained by seeing what remains the same when looking at the object with different technological variations.

The problem that Ihde is trying to solve with this theory is the question of how we know that what we see through scientific instruments is really there. His solution to this problem, is to look at the same object through many different instruments. On some points, those instruments will disagree, while on other points, their views of reality converge. Those things that remain the same then, we can assume to be features of the object itself, and are not caused by the mediating effect of the technology. This theory at least partially answers the question, as it holds that knowledge about the object is gained by agreements between different technologies. As such, disagreements between technologies mean that knowledge is not gained.

Yet, as I explained above, technological data must also be interpreted, and this process of interpretation is not neutral. Ihde's claim that we can look at different technological variations in order to find truth about the world falls apart when we realize that our mesoperceptual famework influences the interpretation of the scientific data, and while Instrumental Realism still holds value, it does not seem to properly answer our question.

The final step then, is to look on a higher level, not at the individual technologies but at the patient and his situation. The question is not what happens when the two technologies do not converge, but rather how to handle this situation. This is something that, in the medical field, is the responsibility of the physician, who takes responsibility for the treatment of his patient, and makes decisions regarding it. By changing the way the question is asked, we open a new solution. It is not necessary in any way to gain access to the true nature of the tumor which seems to be impossible. However, the physician can reflect on the ways in which technological mediation and his own mesoperceptual framework influence his knowledge of his patient, and he can use this to better help that patient. When parallel technologies do not converge, this is still a problem, but one that can be solved by the physician, who pragmatically picks one technology to believe, using all the knowledge and information available to him. His expertise, the nature of the technology, the specifics of the patient and countless other things

1.6: Audiences

I have written this thesis with several audiences in mind, who I believe will enjoy reading it, and may benefit from what I have to say here. The main audience is that of people who practice medicine, the people who interpret medical images and treat their patients accordingly. To this audience, I wish to explain the differences between imaging technologies, and the way they mediate between them and their patients. The importance of how technological details influence the way they work, their expertise, and ultimately how to handle the issues of mediation and interpretation in such a way as to do right to their patients.

I also have two secondary audiences in mind, one consisting of philosophers, and one consisting of engineers with an interest in the philosophical side of their work. For the philosophers, especially those already familiar with the themes discussed here, I think the technological perspective will be of interest, that is to say, the direct application of theory to technology, going beyond the superficial. As to the engineers, I discuss the effect of imaging technologies on the people that use them, technologies that have, of course, been designed. Thus, I hope that I can help this audience understand the role that the technology they create will play.

Chapter 2: Mediation

2.1: Introduction

Going by its name, a medical image is an image that is used in the medical field. It is at the same time an image, and an object of medical interest. Unlike a piece of art, say, a photograph of a landscape of a painting, its purpose is clearly defined. A medical image's purpose is, roughly speaking, to help a physician in the treatment of a patient. It can do so in many different ways, but the one that is most relevant here is diagnostics. The image tells us about the state of a patient, it can say whether someone is healthy or ill, what kind of disease someone has, and it can give us specific information about the disease. It can answer questions about where in the body the disease is located, about the progression of the disease, and it can help in giving a prognosis.

But the idea that an image says something is just a metaphor. An image, in and of itself, cannot speak. In order for it to be useful, the image must be interpreted. Someone, in the case of medical images a physician, must look at the image and say "this is what this means". The problem here is that this interpretation process is not neutral. It is affected by many different factors, and different interpreters can have different opinions about the image.

In addition to this, an image is not the object of the image. This is especially clear when we look at medical images. Once we move into this field, it becomes important that there are many different types of images, that all seem to show the same thing. They are images of the thing, but they are not the thing itself, only one way to look at it. The images created by an x-ray machine are very different from those created using MRI technology, or Ultrasound scanners.

This chapter will talk about the twin topics of image constitution, and image interpretation. How the image is created by technology, and how this image is then changed into something like a diagnosis. In the interpretation of a medical image, the reader of that image must relate the object, the patient's body, to the image. The question he must answer is how features in the image relates to features of the patient. To do this, he needs knowledge about how the image was created. To give an example, one must have knowledge of how X-rays work, and how an X-ray machine creates the images on the screen, in order to be able to correlate white forms on the screen with bones in the body, and thus say that the broken shape in the image means that the bone in the body is broken. Thus, while they are separate steps, the constitution of the image and its interpretation are linked.

I will start this chapter by talking about images and photography, explaining what photography does, and how it mediates between the observer and the world behind the photograph. This all to illustrate the differences between such photographs and medical images, to show how much more x-ray photography or photoacoustics mediate the relation between observer and world.

Then, I will explain the two technologies talked about in the introduction. X-ray photography and photoacoustic imaging. I will compare the two of them, and show the similarities as well as the differences. After this, I will explain the process of quantification involved in spectroscopic quantified photoacoustics.

After this, having set the stage, I will go into further detail on the philosophical side of things, further explaining how technology mediates human interactions with the world around them, not just in general, but in the specific case of the technologies introduced. Amongst other things, I will explain

Don Ihde's Human-Technology-World relations, as well as the way technology influences the way we see reality, while the way we see reality is also influenced by the technology in turn. [5]

Then, the last part of the chapter will end in an explanation of the motion of multistability, and how this issue relates to the topic at hand. This all leads to a conclusion about the importance of expertise in the process of 'reading' an image.

2.2: Images of the world

Everyone knows photography, the art of producing photographs. These photographs are images of the world around us, they show us things. A photograph of a person looks just like that person, a photograph of a building properly represents how that building looks. This means that, when we show someone a photograph we have taken, we show them what we have seen. At least, that is the underlying assumption.

Normal photography is a technology that is, in a way, transparent. Not in the sense that it is easy to see what happens, but in the sense of a window. A photograph does not obviously change what is seen, or obstruct our view. We can, metaphorically speaking, see through the photograph itself, looking straight at the thing that was photographed. It transports the observer to the time and place the photograph was taken, even if it is somewhere we've never been, before we were born. We trust the thing on the picture to have actually happened. This leads to an important point about photographs. We generally assume that they show the truth. If something is shown on a photograph, we expect that situation to have happened. Even if we know that it is possible to perform image manipulation, the general assumption is one of truthful depiction. The photograph shows something we could have seen, if we were there at that time and place. The first time I went to Paris, I was not surprised upon seeing the Eiffel tower. I'd seen pictures of it many times. Those pictures depended upon the existence of the Eiffel tower. If there was never such a thing as an Eiffel tower, then there would not be any photographs of it.[8]

That does not mean that there are no differences between a photo and reality. A photograph does not move, it is locked in time. It is 2-Dimensional in its existence. We can only see the object in it from one perspective, we can't move around to get a better view. The lens distortions, while small, are different from our eyes, colours can change in the process, and those on old photographs can fade. Going beyond vision, there is much more we miss. We cannot hear the sounds the photographer was hearing, smell the scents, and feel the temperature, etcetera. Our sensed and perspective are limited, but in return, the technology shows us another time and place in a fixed image that does not change.

Thus, a photograph does not, in fact, show us the world. It shows us specific things about the world, but not others. Some such things are enhanced, using lenses, or flash photography, to see what normally cannot be seen. Things change in the taking of the image, but this does not necessarily make an image 'less' than reality. While some information is thrown away, other things are gained, expanded. One can even see things the eye could never see. The easiest example of this can be seen by looking at an active TV-remote through a camera. The infra-red light is invisible to the naked eye, but the camera can see and display it, translating it to another colour, usually white.

In other words, while photographs link us to another time and place, they mediate this relationship. The way we look at an object trough a photograph is different from the way we look at an object directly. Even a seemingly transparent technology has a mediating role between human and world. Think, for example, of photographs of famous paintings. To see Rembrandt's 'De Nachtwacht' on a poster, in a book or on the internet is one thing. To see it in person, to be there, is different, even if the image is the same.

Some medical imaging techniques are very similar to photographs, especially X-ray photography. It enhances, translates wavelengths we cannot see to those we can. But in those changes, information is both lost and gained. They do not purely represent reality, but rather show us a part of it, mediated by the technology.

2.3: Medical Imaging

As already explained, medical imaging regards the creation of medical images. A medical image being an image that gives medical information to its user. Unlike photographs, these images usually seek to show us what we usually cannot see, the inside of the body. The most well-known type of medical image is probably the X-ray scan, something which most people can immediately identify. On the archetypical x-ray photograph, the human body is transparent, with bones being the only thing that can be clearly seen. We recognize this, and can figure out what things mean. There are also however, other types of imaging techniques, these imaging techniques are usually called imaging modalities in the field, and they come in many different shapes and forms.

X-ray photographs, MRI scans, CT-scans, Ultrasound images, shear-wave imaging, Optical Coherence Tomography, and photoacoustic imaging, all different imaging modalities. While their ultimate purpose might be the same, (supplying medical information about the patient), they achieve this purpose in different ways. While an X-ray image and an MRI-scan can both be used to detect broken bones, they work in a different manner, and show their results in different way.

In this thesis, I will focus on two different imaging modalities: X-ray imaging and photoacoustic imaging. Photoacoustic imaging is a technology that has been proposed as an alternative for breast-cancer screening, possibly replacing X-ray mammography. However, while both techniques can detect breast cancer, or at least markers for breast cancer, they do so in drastically different ways. The way in which they mediate between the patient's body and the radiologists is different.

2.3.1: X-Rays

X-rays are a form of electromagnetic radiation, much like the light we see with our eyes. The difference, is its wavelength, which lies between 0.01 and 10 nanometres, smaller than that of light. This also means it has more energy, and is more capable of ionizing atoms, which is what makes them harmful to humans. However, while harmful, they are also useful. They are able to penetrate far deeper into human tissue than light, without being scattered. But while they can penetrate the human body, they cannot penetrate all its parts equally. Bones for example, are very dense for x-rays. This is why they are clearly visible in X-ray images. When using lower-energy waves, and sensitive equipment, more subtle differences, such as those involved in breast cancer, also become visible.

X-ray imaging uses a vacuum tube with a cathode and an anode, to create x-rays. These x-rays are then directed through the body of a patient. Different structures in the body block these rays at varying rates, letting only a part of the original photons through. The rays are then picked up by a detector. In the past, this detector was roughly analogous to a negatives used in photography, which was then developed into a picture. Modern systems are almost always digital, using a digital detector. There are several different types, and the question of which type is used can have an effect on the quality of the image, but the principle remains the same.[9] This data is then shown in an image. In these images, areas of high absorption such as the bones are usually coloured white, while low absorption areas are black, with a grey-values in between. This is slightly counterintuitive. In a photograph, and image of a lamp will show the light source in white. However, the white parts on an X-ray image are those parts where there is little to no “light”.

2.3.2: The photoacoustic effect

The other imaging technique is Photoacoustic (PA) imaging. Much like X-ray imaging, PA looks inside the body, but it does so in a very different way. Photoacoustic imaging works by allowing optically absorbing chromophores to emit ultrasound (US) waves. This occurs through thermo-elastic expansion at the absorption site, in response to short pulses of light.[10] On a fundamental level is uses the principle that, when things heat up, they tend to expand. A short beam of light, generally in the infrared, is sent into the tissue, and starts scattering. This light is absorbed by chromophores, in the case of near-infrared light, blood. The absorbed light creates an increase in temperature, which makes the chromophore larger, sending a tiny shockwave through the tissue around it because of the sudden expansion. That shockwave can then be detected using Ultrasound detectors at the surface of the tissue, and it can be used to reconstruct the location of its source. This allows PA to use optical imaging to detect the vascularization-driven optical absorption contrast from cancer, while using the high resolution and imaging depth from Ultrasound Imaging. In addition to that, through the calculations used to reconstruct the tissue from Ultrasonic data, a three-dimensional region can be reconstructed.

2.3.3: The comparison

As seen in the short overview above, both x-ray mammography and photoacoustic imaging are imaging techniques that can be used to image breast tissue for the purpose of identifying breast cancer. X-ray Mammography is currently being used in many countries around the world for the national screening of breast cancer, for early detection. Scientists at the University of Twente are currently working on photoacoustic breast-cancer screening technologies, to either supplement or replace x-ray mammography[3, 11]. On the surface, these technologies seem similar. Both technologies are medical, highly advanced, and can look inside the human body. Furthermore, they are, or at least can be, used for the same goals. They mediate the relationship between a physician and the tissue he wants to look inside of. The images created by the technology give a view of the inside of the tissue, allowing a radiologist to “look inside” the body. However, when you open up the technology and look ‘under the hood’, there is a drastic difference. X-ray imaging and photographs are isomorphic, the difference being the wavelength of the photons, and that x-ray imaging always requires a flash, while most of the rest is pretty similar. Photoacoustics images are different. In X-ray imaging the tissue plays a passive role, while photoacoustics make the tissue an active participant in the imaging process. Photons still enter the body, but their different wavelength means that they are immediately scattered, making them seemingly useless for locating structures in the body, after all, you don’t know where a detected photon comes from. However, where in X-ray imaging the particles that don't get absorbed are measured, photoacoustics does the opposite. The absorbed particles create ultrasonic waves that are picked up. These ultrasonic waves are then reconstructed into an image, but these are waves of pressure, not photons.

Where an x-ray image can be seen as analogous to a photograph, photoacoustic images are wholly constructed. To explain, an x-ray image can be said to translate vision. Unlike a microscope, it is not an amplification of an object, but a shift in wavelength. The wavelengths picked up by the technology are translated into those that can be picked up by the eye, the invisible is made visible. As implied before, x-ray images take the form of a negative. Where you would expect darkness/black (the bones, which absorb the "light"), light/white is seen, and vice versa. The way in which photoacoustic images work is very different. Rather than being analogous to a photograph, they are constructed in a way that is far more artificial. The way in which the images are created is slightly arbitrary, there seem to be more choices made in how to represent the data, which must be known in order to read the images. They do not naturally come from the data. After all, ultrasonic waves are received by a sensor, but an image, thus, light, is displayed on a screen. Computer analysis is necessary for the creation of the image, and not just a more efficient replacement for traditional picture development.

Another important difference between the two technologies is the type of image they create. For X-rays, the tissue becomes transparent, the user sees all of it at the same time. Three-dimensional data about the body is seen in a two-dimensional image because the body is transparent to X-rays. If an X-ray scan is made from the front of the patient to the back, it is difficult to figure out how deep in the body a structure is located through a single image, because all depths are shown at once. In mathematical terms, an X-ray image has an X-axis and a Y-axis in the height and width of the image, and shows the entire Z-axis at once.

In the case of photoacoustics, something different happens. A threedimensional area can be reconstructed, by comparing the signals picked up by different detectors. Then, from this three-dimensional area, a slice can be shown as a two-dimensional picture. This slice is more akin to what you see when have a slice of fruit, you can see the inside in two dimensions. But this single picture only gives you information about a single slice. To return to the mathematical terminology, we see along the X-axis and the Y-axis, but we only see a specific location on the Z-axis.

A final important difference is that photoacoustics can be used for functional imaging[12], unlike x-ray imaging, which is limited to structural imaging. What this means is that X-ray images can show the structure of the tissue that is imaged. What is located where, and how dense to X-rays are the images. photoacoustic images however, can also detect blood, and blood oxygen concentrations, this is related to what is happening inside the body, meaning that its images contain information about the functioning of the tissue as well as its structure.

2.3.4: Quantifying measurements

While different, X-ray Imaging and photoacoustic imaging are also quite similar in that they create images whose data is not quantified. The data in them is relational instead. If one place in the image is white, and the other is black, then one has a higher value than the other, but we cannot translate that to quantified numbers about the tissue.

One thing that researchers try to do is figure out a way to quantify this data. This is important for several reasons, most importantly, that it is far more objective. Quantified data, if done correctly, is not dependent upon the observer, the place of measurement, or the instruments. Two different people using a thermometer will still get the same temperature number in the same place. While if they try to describe it, one might say it's a nice temperature, while another finds it too cold. This quantification makes it far easier to relate a body of research to an image. For example, if the level of blood oxygenation in a body part is quantified, and the literature says that that level of oxygenation is correlated with tumorous tissues, then it becomes very simple to identify the tumor.

PA is generally unable to create quantified images of absorption coefficients, because under the conditions of heat and stress confinement, the initial stress distribution that creates the ultrasonic waves depends on the absorbed energy density. That is to say, the stress distribution depends on the Grüneiesen parameter, the local absorption coefficient, the local fluence, the concentration of a chromophore, and its molar absorption. All these quantities vary with position in the tissue, and the quantification problem in photoacoustics is the problem of how to decompose this initial stress distribution into the fluence, (that is, the amount of light in an area) and the absorption coefficient.[13]

One way to solve this problem in the specific case of measuring absolute blood oxygen saturation, is to use multi-wavelengths measurements with combined photoacoustics and acousto-optics. It uses acousto-optics to determine the relative differences in fluence at two different wavelengths, and then uses that information to compensate the relative difference in fluence in photoacoustics. Normally, this would only give relative values, but we are interested specifically in blood oxygenation, and through the difference in absorption spectra between haemoglobin with and without oxygen, it is possible to mathematically calculate the amount of oxygenated blood.[6]

This is relevant, because such quantified measurements will greatly decrease the role of the interpreter in the analysis of data. Instead of having to 'read' a medical image, numbers can be read directly from the data, and compared to literature values in order to make a diagnosis.

2.3.5: Comparing quantified measurements

The differences between X-ray and photoacoustic images are, as explained above, largely related to what exactly they measure. The presentation of both is roughly the same, as a 2-Dimensional image. Their purpose, too, seems to be the same, giving information to a doctor in order to allow for diagnosis. The quantified information gathered by spectroscopic fluence-compensated photoacoustics is different in other ways. While its way of gathering (some of its) information from the tissue is the same as in normal photoacoustics, the big difference is in what it does. It allows for numerical analysis of the body, instead of the relative measurements of x-ray imaging and photoacoustics. This has consequences for how physicians use the data. Rather than the type of interpretation found in analysing x-ray images, they now see a number that can be directly related to a diagnosis. Therefore, the quantified technology is, or at least can be, drastically different. On one hand it is possible to create a 3-D map, and then a 2-Dimensional image of a plane in that map, of blood oxygen saturation levels. This would mimic the other two technologies. However, blood oxygenation levels in the body are generally rather high, around 95% or higher, while haemoglobin in a malignant tumor has oxygenation levels that are far lower. Given the disparity between normal levels and those found in a tumor, as well as the quantified nature of the measurements, something very interesting happens. Rather than the image showing us characteristics of a human body to be interpreted, the image can now tell us the location of malignant cancers. Taking this to its (il?)logical conclusion, breast cancer screening could become rather similar to an at-home pregnancy test, with a computer giving either a positive or negative result. As such, the increased precision of this technology (specifically the quantified nature), drastically changes the way people interact with it.

2.4: Not a photograph.

Before, I have shortly explained how images, photographs specifically, are at the same time like, and unlike, reality. Now, after going through the more technical details behind X-ray photography and photoacoustic imaging, we can do a more detailed analysis. What is it that is changed in the creation of the image, and what points that people may assume images to have, are not as simple as they might seem?

In the article 'Are Neuroimages Like Photographs of the Brain?'[8] Adina L. Roskies talks about how people often look at neuroimages, images created using functional MRI measurements of the brain, as if they are photographs. To the layman, or in fact the philosopher, it is an attractive option to assume that scientifically produced images have the same features that photographs have. For example, if an image from a functional MRI scanner shows us that brain region A is active in mental activity B, then brain region A and activity B being related is as sure as a picture on a postcard telling me that the Eiffel Tower is in Paris. After all, we can see the brain region light up in the image. Roskies criticises this by looking at several features that are classically ascribed to photographs, and shows how these do not apply to neuroimages, leaving out the question of whether or not photographs actually have these features. In the following section, I will use some of these features to describe how X-ray and photoacoustic images are, like fMRI images, not photographs, hopefully illustrating to the reader why we should take a better look at them. The features I will discuss are as follows

- Visual Dependence
- Causal and counterfactual dependence
- Belief Independence

2.4.1: Visual dependence

One of the features of photographs is their Visual Dependence on the original object. A photograph looks like the thing it is photographing. When speaking about x-rays, on first glance, they do depend on the visual properties of the object. Bones appear in white on the image, and bones are themselves also white when we look at them. This makes it clear that what we are seeing are bones. But the colour of the bones in the x-ray image is a mere coincidence. There is, in fact, a good argument to be made for them to be black, and everything else white, instead. Still, while there are some differences, when we shift the meaning of visual to a different wavelength, and accept the fact that we are looking at a negative, it becomes possible to see that x-ray images hold the visual properties of their subject. At least for whole body images. When chest, or breast, x-ray images are concerned, the story is wholly different. While some outlines are usually recognizable to the lay-man, there is very little stereotypically visual about them.

Things get worse when we look at photoacoustics, or Roskies' fMRI images. In her treatment of the fMRI images, Roskies notes that what an fMRI image shows is a so-called BOLD, or Blood Oxygen

Level Dependent, signal. This is a signal that is created by changes in the process of dephasing of protons in water molecules, caused by a change in the oxygenation level, and the concurrent change in magnetic properties of blood in the brain. As such, this is very clearly something that is not visual information. However, this is not all. The BOLD signal is in some way related to brain activity, this much is known. What is not known, however, is how the two are related, or what the delay is.

Photoacoustics also shows information that is not necessarily visual. First of all, one of the differences between photoacoustics and X-ray images are that x-ray images rely on x-ray transparency of the object. The image is as if we look through something. A Photo-Acoustic image is more like taking a slice out of the body and looking at that instead. Photo-Acoustic images are based, not upon light going through the body, but rather light being absorbed by the body. The signal is sonic rather than photonic, and the image created is computer-reconstructed rather then direct. Although one of the things they depend on are the "visual" property of absorption, photoacoustic images are dependent upon this in a very non-visual way.

Quantified photoacoustics are very much like photoacoustics in this respect, but one of the things that should not be forgotten is its quantified nature. fMRI and PA images have in common that, although they are reconstructed using mathematical techniques, and are in fact based upon numbers, these numbers are not "hard" numbers, in that they do not specifically refer to an absolute value, especially in photoacoustics. By putting a number that is actually known in an image (blood oxygenation level), the image becomes less like a photograph, and more like a mathematical graph.

2.4.2: Causal and Counterfactual Dependence

The second factor discussed is that of "causal and counterfactual dependence". The existence of a photograph is dependent upon the object of the photograph. If the object does not exist, the photograph cannot exist. Furthermore, had the object been different, so too would the photograph have been. In the case of fMRI images, Roskies claims that this counterfactual dependence is not achieved. In fMRI, it is not known what events or changes in the brain could occur without changing the image, where we do know this in photography. Because fMRI's work indirectly through BOLD signals that only somewhat refer to brain activity in not entirely understood ways, there is a large degree of change possible without a change in the created image. For example, it is impossible to know if activity in an fMRI image is inhibitory or excitatory. Two complete opposites, of which we cannot know which is what really happens. For X-rays, the phenomenon is well understood, and, at least to the expert, there is knowledge about what changes in the tissue lead to what changes in the image.

Of course, by their very nature, diagnostic images deal with this topic. If a change in tissue (the existence of a tumor) does not lead to a change in the image (seeing micro-calcifications , lesions, etc.), then the image would be useless for diagnostics. For photoacoustics, the situation is different. While generally not a mystery, there is less knowledge about photoacoustics then there is about X-ray, due to the relative ages and popularities of the fields. Furthermore, the fluence problem discussed earlier in this chapter is a good example of changes in the tissue not creating changes in the image. It is possible for, say, the absorption rate to double, and the local fluence to halve, without a corresponding change in signal. The risk of this having a large effect is not very high. After all, the images created are relative, and sharp changes will probably be to changes in absorption, rather than light distribution, but the possibility is there. This is where quantification plays an interesting role, in that it attempts to perfectly capture what happens, at least regarding blood oxygenation. Due to the quantified nature of the measurements, it is impossible for there to be a change in blood oxygenation without a corresponding change in the image (within the context of normal diagnostics). Of course, changes in the body that do not have a corresponding change in blood oxygenation level are invisible to the technology. This affects x-ray imaging too, if a change occurs outside of the wavelength of x-rays, it is invisible to the technology.

2.4.3: Belief-Independence

The third point Roskies makes is about belief-independence of photographs. She talks about photographs being belief-independent because they can be interpreted without knowing the beliefs of the creator. For the average photograph, simply seeing it is enough to understand what it is. A building, a car, a group of friends. However, when looking at highly scientific images such as those created by an fMRI machine, it becomes clear that we need theories and skills from those fields to properly interpret them. The image itself does not directly make clear what it means in itself. One needs theories about how the brain works, and how to colour that on an image, to understand what it means.

An x-ray image of the human body usually does not have this problem, at least for our society. It seems ingrained in our culture that we know, sort of, what our skeletons look like. Thus, when we see an image of a body with a clearly broken arm, we can interpret this quite easily. When looking at more subtle bone fractures, or images created for the analysis of things other than broken bones, more understanding of the human physiology is needed.

In photoacoustics, this problem is exacerbated, since the way the image is created is less analogous to normal vision. Much like in fMRI's, statistical analysis and reconstruction are used in the creation of an image, and in order to understand what the image means, these techniques need to be understood. When we look at quantified imaging, the image becomes even more belief-dependent, the quantified aspect can only be understood when the viewer knows what the numbers mean, and this is not limited to a contrast between low and high values which could be easily translated to a colour-map. Theories about what specific oxygenation levels mean, for example, are needed to fully appreciate the power of these quantified images.

2.4.4: Hermeneutics and mediation.

To summarize this section, images created by technologies such as x-ray imaging or photoacoustic imaging are not simply a way to see the inside of the body. They influence the way in which the body is seen.

In more philosophical terminology, x-ray machines of PA set-ups are mediators, mediating between their user and the world around them. They have their own features, which influence the way the world is seen through them. At the same time, however, their results must also be interpreted, the technology is not in charge, and it is a human that does the work. There is an influence but it is not absolute.

This then, is mediation theory, the idea that the technology mediates between human and world. Different aspects of the world are highlighted by different technologies, while other parts are neglected.[5]

The other side of the coin is the interpretation of a thing. While two people looking at a scan will have their access to the world mediated in a different way, they interpret the image differently. This process of 'reading' the meaning of an image, or anything else, is a hermeneutic process. The image is, metaphorically speaking, a text, that has to be read. As such, what it means is not self-evident. In this hermeneutic process, many things play a role. The mind-set and cultural background of the observer, his experience and knowledge, and other factors.[14]

2.4.5: Ontology

I have talked about technology mediating between human and world, and in doing so, I have made an assumption. Namely, that of the existence of the world.

The problem of mediation, is that everything is always mediated. It is impossible to look at, say, a tumor, without the relationship being mediated in some way. Even if we cut open the patient and look with the naked eye, the eye itself mediates between us and the object we are studying. The question of course, is, if we are only capable of accessing a mediated world, how do we know there is such a thing as the world itself? Is it not possible that the world is generated through the mediation technologies we use to view it? After all, we have never seen the world directly, it is impossible to see.

Is there such a thing as a tumor that we study different facets of through different technologies, or would the tumor not be there without the technology? How can we know anything about the world if it is impossible to observe it directly?

I handle this topic in chapter four, epistemology, using philosopher Don Ihde's theory of instrumental realism.[7] Until that point, I ask of the reader to keep this conundrum in mind when reading this chapter and the next. While I speak of technology allowing access to the world, there is always that tiny voice in the back of your head, asking "what world?"

It is in chapter four that I give arguments regarding this topic. Until that point, when I talk about access to the world, it is mostly a manner of speaking, a way of talking about things without the sentences becoming overly complicated, it is not a point about the structure of the world.

2.5: Human-Technology-World

One of the people talking about the way technology mediates our access to the world is Don Ihde. He has categorized several different ways in which technology mediates our access to the world, with different technologies having different effect. He identifies four different ways in which technologies can mediate access between the Human and the World. In this section, I will explain these different categories, and explain where the technologies under discussion fall, as well as expand upon the ideas of Ihde by changing them from categories into a spectrum. The categories that Ihde identified are: Embodiment relations, Hermeneutic relations, Alterity relations and Backround relations. The following then, is an explanation of these relations according to Verbeek's interpretation of Ihde's technology relations in 'What things do.'[5]

2.5.1: Embodiment relation: (I-Technology) -> World

An embodiment relation is a human-technology relation in which the person 'becomes one' with the technology, the technology becoming transparent to the user It is still there, but the user is not actively thinking about it. A good example of this is eye-glasses. When someone wears glasses, a lens before his or her eyes changes the path of the light coming towards his eyes; the very photons themselves are modified by the lenses in the glasses. This seems drastic; the way in which the entire world is seen is changed by the technology. However, when someone wearing glasses is asked to describe what he sees, he will describe the object in front of him, he will not tell you about curved glass and a frame directly in front of his eyes. While obviously a part of the human-world interaction, the glasses are invisible unless attention is called to them.

However, this example calls attention to something. When someone who never wears glasses puts on a pair, the frame and distortion will be very apparent. There is more involved than just the technology itself. Another example is the photograph. We say that a photograph becomes transparent, since we can see the object in the image without continuously being aware that it is actually a picture, chemicals on a piece of paper.

However, this reading of a photograph is not a natural thing, it is a skill, if an ubiquitous one. Take, for example, the bicycle. Through it, a rider can feel the road below him. But, only if he is familiar with the feeling of sitting on a bike. Other skills such as these are less ubiquitous. Take a blind man's cane. In the hands of the lay-man, it is an object that can be used to poke stuff, a stick of sorts. To the sufficiently experienced, it is an extension of reach, and gives information about the ground in the same way stepping on it would. The cane is embodied. Another good example of this is the computer mouse. In the hands of an experienced user, it is the cursor on the screen that is moved, not the mouse, the mouse moving seems secondary to a movement on the screen, even though the process is entirely artificial. Thus, embodiment relations are almost always skill-based, even if these skills are ubiquitous in our society.

To explain the term used in the title of this paragraph, in an embodiment relation, human and technology become one in their interaction with the world.

2.5.2: Hermeneutic relation: I -> (Technology-world)

An embodiment relation disappears; the technology is embodied, and in so doing is forgotten by the user. The hermeneutic relation is different. Hermeneutic technologies extend the senses of the user in a more abstract way, requiring interpretation in order to be useful. Hermeneutic technologies have a meaning, but this meaning is not as directly clear as the one found in an embodiment relation.

Take the example of a thermometer. A thermometer allows the user to measure the temperature of something, even at a distance. However, it does not so by heating up or growing cold on the display, it does so by displaying a number, a number that is interpreted as relating to the temperature of the original. The user needs to read this number, interpret it, in order for the number to say anything about the temperature. It is a theoretical process, rather than a practical one.

A good example of the contrast between Hermeneutic and Embodied relations can be found in music. Right now, my laptop is playing me a song. I hear a singing voice, a guitar, some bass-lines. The technology is giving me access to sounds that were recorded in a studio somewhere, even though all that is happening is that s a membrane vibrating in my speakers. The technology is embodied. However, I could open a sound analysis program, get an overview of frequencies and amplitudes. With enough knowledge on how to interpret this, I could read the music, a hermeneutic relationship. More traditionally, I could read sheet music.

At first glance hermeneutic and embodied relations are very different. However, I posit that they are actually on a spectrum with each other, and that it is user dependent what role a technology plays. Think of Beethoven for example. He famously became deaf late in his career, but he kept composing new music. Even though he could no longer hear music, he could still interact with it fully. To him, music written down on a sheet was equivalent to music that was heard. Can we say, then, that he still had to interpret the sheet music in a hermeneutic relationship? Or was this process automatic, did the paper create sounds in his mind, just like vibrating air would?

To explain the schematic, a hermeneutic relation is one in which the human, the I, interprets the world as the technology represents it.

2.5.3: Alterity relations: I -> Technology (-world)

An alterity relation is a relation in which the human interacts with the technology itself, rather than interacting with the world through the technology. An object in an alterity relation has some seeming autonomy, which allows for interaction between man and technology. The level of autonomy in play can, of course, vary. On one side of the spectrum there is the full artificial intelligence that interacts, while on the other side there is something as simple as a spinning top. In an alterity relation, the technology takes on the role of the “other”, as such, one can speak of interaction *with* technology rather than *through* technology. Think of playing against a chess computer, rather than a human player. Thus, in an alterity relation, the focus is the technology itself, the world being involved only as an afterthought.

2.5.4: Background relations.

A technology in a background technology is a technology that seemingly disappears. This might sound like the same thing as an embodiment relation. However, an embodiment technology disappears in use, while a background technology the technology disappears because it is not “in use”. The thermostat in your home or office is an example of a background technology. The technology mediates the way in which we experience the world, but does so from the background.

2.5.5: Human-Technology-World relationships of Imaging technologies

As already explained, this list is not exhaustive, and merely a general overview of different types of human-technology relations. As seen in the previous sub-section, both x-ray images and photoacoustic images should not be seen as photographs of the inside of the body, they are different in what they detect and how they are constituted. However, this does not mean that they might not count as technologies involved in an embodied human-technology relationship. X-ray imaging seems to be involved in such a relationship, at least on a superficial level. However, when you go deeper into an analysis, the features of the technology become very important, and equally so the skill that the user has in using the technology.

Photoacoustic imaging is an interesting case. At first glance it seems to create an image in much the same way x-ray imaging does, and, being similar to a photograph, an easy candidate for an embodiment relation. However this changes when looking into the technology. When you look under the metaphorical hood of photoacoustics, you find out that the image is based upon a 3-Dimensional reconstruction of ultrasonic waves created by active interrogation of the tissue. Given this, it seems to be a hermeneutic relationship, one reads the numbers in the computer in order to gather information about the object being studied. However, by creating an image, there is an attempt to shift the technology back to being embodied, showing the inside of the body.

To give an example using a more general technology, imagine a digital thermometer. This is a hermeneutic relationship, we read the number on the thermometer and interpret that as telling us something about the temperature of a room or object. However, imagine a heating/cooling element that is attached to this thermometer, that automatically changes to the temperature of the thermometer. If the sensor of the thermometer is placed in a different place from the heating/cooling element, it is possible to feel, rather than read, a temperature in a different place. As such, a hermeneutic relation seems to change to an embodied one. This is roughly equivalent to what happens during the creation of a photoacoustic image. The hermeneutic I –> (Technology-world) relationship becomes something closer to (I-Technology)-> (Technology-world). Technology is used both to analyse the world, and to help in the interpretation of that analysis.

However, as explained before, embodiment relations require some sort of skill. It is with this skill that the relationship becomes embodied. If this skill is not present, the information must be interpreted in a more conscious, hermeneutic process. Thus, I propose that we see the relation between embodiment and hermeneutic relationships not as a binary divide, but rather as a sliding scale. This means that technologies can be involved in a relationship that is somewhere in between hermeneutic and embodied. In addition to this, its position on the scale can shift depending on the skills of the user of the technology.

Medical images are, on one hand, hermeneutic. They have to be interpreted as to their relationship to the body of the patient and how they represent it. Understanding what colours in a picture mean in regards to a patient is a hermeneutic process. However, they can also take on an embodiment relationship, for example during computer assisted surgery, in which a surgeon uses medical imaging to keep track of the exact position of his instruments, and where they need to go. This is a delicate procedure in which the visual information provided by imaging techniques is directly used by the surgeon to guide his handheld instruments, an embodiment relationship.

Furthermore, where a novice might need to consciously think about what an image means, an expert can sometimes read an image intuitively. Thus, expertise turns a hermeneutic relationship into an embodied one in some cases. To make matters more complicated, technological interference can change the exact data to be more embodied, or more hermeneutic. To make a dataset more embodied, numbers and other mathematical information can be used to construct an image. To make it more hermeneutic, something like a photograph can be distilled into a set of Red, Blue, and Green values.

While x-ray imaging and photoacoustic imaging seem to lie somewhere on the spectrum between embodied and hermeneutic technologies, quantified photoacoustics take on a much more hermeneutic relationship, and is even capable of taking an alterity relationship. If only the numbers are shown, it is a clear hermeneutic technology relation. However, if analysis is done by a computer, then the technology acts as an “other” in that it performs a diagnosis itself, rather than leaving interpretation up to the doctor in charge. This means it becomes an alterity relation. As such, small changes in the type of imaging done can drastically change the type of human-world relation the technology is involved in. Furthermore, where a fluence-compensated image is only of use in an embodied relationship when the user doesn't know how to relate explicit numbers to the image, understanding medical theories about this makes the user capable of making it into a hermeneutic relationship.

To conclude, where a technology lies on the scale between Embodied and Hermeneutic is based on both the technology and the user. The user's skillset, cultural background, background knowledge, even his mind-set, these all influence the way in which the technology mediates between observer and world. A photograph, even if it seems an embodied way of looking at the world to us, has a different role in the eyes of someone from a drastically different culture without photography. To someone analysing the paper on which it is printed, the image is more hermeneutic.

2.6: Multistability

As explained, technology mediates the relationship between its user and the world. The way in which the world is experienced is shaped by the technology in play. However, technology does not have a single purpose, and a single object can have different effects in different situations. Technological objects are multistable, in that they have different, stable, interpretation. There seems to be no single, objectively "best" interpretation in many of these. This regards the use of the technology, but also its interpretation. To give an example of the first, multistability in use, many objects can serve as either a tool or a symbol. A military officer wielding a sword in the twenty-first century does so not as a weapon, but as a symbol of authority. A picture can serve as a piece of art, or as a representation of an object. The second type of multistability, regarding interpretation, is more relevant to this thesis.

A good example of this type of multistability is found in the Necker cube, a drawing of a translucent cube seen from either the upper-right side of the bottom-left side. Both interpretations of the image are equally valid, there is no 'correct' interpretation that is better than the other, both are stable. In the same way, a medical image can represent two things, a healthy patient, or a sick patient. A patient without breast cancer, or one with breast cancer. As such, when confronted with a medical image, one must interpret it as either of these. The problem, however, is that unlike the Necker cube, there is a 'correct' interpretation in this case. After all, the patient is either healthy or unhealthy.

However, while there might be only one interpretations that corresponds to medical reality, the image itself still seems to be multistable. Some images may clearly depict healthy patients. Others can clearly depict sick patients. The third category, is images that can be interpreted as depicting either a healthy, or a sick patient, both interpretations being stable. Are those points an anomaly, or the mark of a tumor? Is the inflammation caused by cancerous tissue, or simply stress? If a medical image has more than one stable interpretation, more than one valid story that can be told about what it depicts, then how do we choose which is better?

In Robert Rosenberger's 'Perceiving Other Planets: Bodily Experience, Interpretation, and the Mars Orbiter Camera"[15], he talks about the multistability of images create by a camera in orbit around the planet mars. He talks about multistable images with the term 'hermeneutic strategy'. A term he uses to describe the ways in which the images can be interpreted. In order to see an image in a specific way, Rosenberger says that the watcher has to know about the necessary hermeneutic strategy. One has to be aware of the existence of nails and how to use them in order to be able to interpret a hammer as a nail-hammering device. In much the same way, we need to know about paper and paper-weights if we want to interpret the hammer as an impromptu paper-weight. In the more scientific context Rosenberger talks about, It is the existence of different theoretical frameworks and strategies that explains how there can be equally valid interpretations of an image. A hermeneutic strategy for the interpretation contains everything necessary for understanding how the image came to be the way it is. This means that it not only relates to how an image is constituted, (How does an X-ray machine work, how does it interact with tissue) but also what the object being imaged exactly is. As such, a difference in opinion about the nature of the imaged breast can lead to a different opinion in much the same way as a difference in opinion on just how the image is constituted. A hermeneutic interpretation can walk the path backwards, so to speak, starting the story with neurons firing in the brain, to photons reaching the eyes from a computer screen, to an image file being processed and displayed by a computer, to arriving X-rays being translated into an image file, to x-rays interacting with tissue, and possible tumor markers, to the creation of those possible markers by a tumor. A hermeneutic interpretation at least implicitly contains all of these transformations of the signal in order to explain how to read the image.

These transformations of data can be rather complex. For example, disagreements over how to interpret an image can be caused by disagreement over how an image is constituted. For example, photoacoustic measurements can be done a single time to create an image, but due to their speed and non-invasive nature, they can also quite easily be done 50 times in a row and then averaged. As such, an object in the image can be interpreted as either an actual object, or a random artefact that should have averaged out if more samples were taken. There are other mathematical processes as well, image compression for example, or simply pixel sizes.

Rosenberger discusses the case of the Eberwalde delta, seemingly an ancient river delta on Mars. This is a formation on mars in the shape of an earth fiver delta, but “fossilized”, meaning it is inverted, with rocks forming the delta, rather than gaps in the ground, due to once being underground, but now having the wind blow the soil around it away. The debate around the issue is not about whether or not these formations are evidence of water on Mars, but rather around how long water persistently flowed. Different hermeneutic strategies involved in the interpretation of these MOC (Mars Orbiter Camera) images do not stop with the images themselves. They often use additional information from other cameras, laser altitude measurements and competing geological theories. This with the idea that the additional information will collapse some of the stable interpretations of the MOC-image. The different hermeneutic strategies interpret the water-flow as having taken either hundreds of thousands of years, or closer to about a hundred years. A difference of several orders of magnitude, yet both interpretations seem to be stable with regards to the images.

Thus, outside information can be used to limit the amount of viable interpretations somewhat, at least narrowing down the multistability. In much the same way, other tests can be done on a patent in order to try and collapse some interpretations of the medical images, but even then, sometimes the debate on the “correct” interpretation remains. Where this is a question that seemingly has an actual answer based in historical fact, this is not the fact in all cases. We do not have time travel to check, we may never know which interpretation is true, but we do accept that only one of them can be. It is, at least theoretically, possible to ‘solve’ some multistable objects, although in other cases it is not. An image like the Necker cube that is intended to be multistable does not have this trait

What is interesting in Rosenberger’s analysis is where he places the topic of scientific controversy. Most people would say that the controversy is about how long water has flowed through the Eberswalde delta. Rosenberger’s postphenomenological perspective however, places the topic of the controversy at the images of the delta. It is the images that are multistable and subject to interpretation, not the delta itself. Since these images and the MOC are our main method of access to the delta, what is happening regards hermeneutic strategies for the interpretation of images, not theories about the history of mars. In other words, while the images themselves could depict a water-flow of a hundred, or of a hundred-thousand years, and are thus multistable, the surface of mars itself supports only one truth. In the same way, medical images of breasts are multistable, but the breasts themselves are not.

Different hermeneutic strategies can form a problem in the medical field as much as in Martian Geology. One of the criticisms of breast cancer screening, for example, is the high false-positive rate, images of healthy patients being interpreted as images of sick patients. In other words, the meaning ascribed to the image is not the one that corresponds to the reality of the breast tissue. However, we are talking about a screening process here. The question is not whether or not there is a tumor in the breast. The question is whether or not the image should be interpreted as one of a breast with a tumor in it. If it is interpreted as an image of a breast with cancer, then there is cause for further testing. So when the image is multistable, and the radiologist interprets the image “incorrectly”, this is not immediately a large problem. Further diagnostic procedures will result in the “incorrect” interpretation becoming invalid, and no longer stable. The problem being, that these further tests are expensive, time-consuming, inconvenient and often invasive.

Both x-ray images and photoacoustic images are multistable. They have different possible interpretations, and there is not always a single stable interpretation in the medical context. Furthermore, unlike in the case of the mars pictures, there is usually little extra information available. In a screening context for example, the entire point is to make a judgement without asking for much more than a single scan. As such, the choice made by the radiologist is not between “breast with tumor” or “breast without tumor”. Rather, the choice is between “image that warrants no further attention” and “image that has a valid hermeneutic interpretation of the image as one of that of a breast with cancer”. In the first case, the screening is seen as negative, no further attention is required, because the radiologist does not believe there is a good hermeneutic explanation that could be used to explain the image as one of a diseased breast. The second option is vaguer, and relates to the possibility of breast cancer. There are one or more viable hermeneutic interpretations of the image as the image of a breast with cancer. Further testing can narrow down what the correct interpretation is. However, given what is required to do this additional testing, and the large scale of population screening, a radiologist also has to keep this false-positive rate down. Thus, in some cases he may also choose to ignore a possible hermeneutic interpretation, given that it has a very small chance of being the “correct” one. After all, if he were to only dismiss those images that were truly, 100%, cancer free, the medical system would be overloaded with false positives, doing more harm than good.

2.6.1: What is the tumor itself?

As explained, the difference between X-ray and photoacoustic imaging lies not only in the way the image is constituted, but also in what is measured. X-ray imaging measures the absorption of x-rays, while photoacoustics measures the absorption of (usually) visible light, by a different route. This means that different structures become visible. In the case of x-rays, breast cancer takes the form of micro-calcifications or other structural issues like lesions, whereas photoacoustic mammography looks at the distribution of blood in the tissue instead. The same object, a tumor, is ultimately being imaged, but in a very different way.

So if the object, a tumor, is being detected in different ways, then what is the tumor itself? It is not the micro-calcifications detected with X-rays, nor is it the chromophores that absorb the light used in photoacoustics. Is it the change in the DNA of the cell then? It is the physical deformation that actually kills the patient in most cases, not the DNA itself. If a tumor is all these things combined, then what it seems to be depends on how we look at it, on the instruments we use. Different instruments give a different view, but seem to all be valid. The way we see it is dependent on our instruments. One way of viewing the tumor might lead to multistability, with 'tumor' and 'not-tumor' both being possible interpretations, while another way may not have this problem, Thus, when combined, at least part of this problem of multistability is solved.

2.6.2: Quantification and multistability

The problem behind quantification of photoacoustic measurements can also be explained using the concept of multistability. A photoacoustic signal is dependent upon both fluence and local rate of absorption. Thus, any change in the image can be interpreted as being due to a change in fluence, a change in absorption rate, or a combination of these two factors. To solve this problem outside information is brought in, in the form of Acousto-Optic measurements. If only one wavelength is used, then the fluence distribution could easily be used to compensate for differences caused by difference in fluence, leading to a better quality image. When information from another wavelength is brought in, then it becomes possible to form a hermeneutic strategy that relates the different signals to blood oxygenation. If scientific theories about the oxygen saturation in tumors are added to this, then the result becomes a hermeneutic strategy that interprets numbers as diseases, seemingly eliminating all other now no longer viable hermeneutic strategies and automatically selecting the "correct" one. This leads to a reduction in the level of multistability of the results, to the point that, theoretically, a computer could pick the only "valid" interpretation of the image according to the theories behind it. The task of choosing a hermeneutic interpretation is thus left to the machine, an alterity relation instead of a hermeneutic one.

2.7: The Physicians role

The topic of this chapter has been the mediating role of technology. I have explained how different medical images are constituted, and how the way in which they are constituted has an influence on how they give us access to the world. In this, one important aspect has been the role of the physician in the interpretation of medical images. Since medical images are multistable, that is to say, there seem to be several different, stable interpretations for medical images, it is important to properly interpret them. The assumption being, of course, that the state of the patient is independent on the interpretation of the medical image. This leads us to the question of what the role of that physician is. What is it that the physician needs to do? To understand what a medical image means, the physician must understand how it is created. In addition to that, he must know how to read the image. As we have seen in paragraph 2.5, the skill of the physician has an influence on the way he interacts with his patient through the technology, changing between human-technology-world relations.

To make this more complicated, there is the issue of quantified measurements and other automated techniques, the idea that it is possible to take this process of the interpretation of a medical image, and to then outsource this to a machine. In a way, this is something you would expect on an episode of Star-Trek, where the captain is beamed to the sickbay and the medical AI tells everyone what is wrong with him, but this kind of technology is in use right now, and often quite widespread. A prime example of this process of technological interpretation of the results is in the at-home pregnancy test. It is a simple device with simple instructions, which gives its users a clean binary answer to the question of whether or not they are pregnant. The technology is not one that is used in an embodiment relationship, and on first sight, it does not seem that a large amount of background knowledge or experience is necessary to correctly interpret the test result. However, even in this relatively simple case, whether a woman does the test herself or lets a trained professional do it has a great effect on the end result [16]. Whereas professional use of such tests were about 97.4 % accurate, consumer testing was far less accurate, with false-negative reports reaching as high as 24.3%. As such, though it might seem that the role of expertise can be relegated to technology, this does not seem to be the case in every situation

So what is the role of the physician? Is it to interpret medical images and create hermeneutic strategies that can be used to analyse the patient? Or is there more to it? And what is it that allows the physician to do this in a way that normal people cannot? The answer we seem to have is that the physician is an expert. He has the knowledge and the skills necessary to properly apply tests and interpret the test results, as well as all the other odds and ends that make up his job. The problem now, is that the question has not been answered, or rather, that the answer automatically begets another question. What is expertise? How does this expertise allow the physician to make *'better´* interpretations of a medical image? Is it the fact that he has been trained to do so, that he has the skills to make the process embodied? Is it his background knowledge about how the images are constituted, and what cancer does to the body of the patient? Or is it the combination of all of these? In the case of the pregnancy test, knowledge about how the device works may not help much in the interpretation of the results, but it will help in proper application. In other situations, other factors are important, and expertise is not supplanted by the technology, but complemented.

A good example of this can be found in the results of blood tests. Many chemicals present in blood have a set of reference values; values that correspond to those found in healthy people. A deviation from the reference values means that something is "wrong", it something that, generally speaking, does not happen, at least not without reason. Seeing whether or not the values are within acceptable bounds is often rather simple, for things like blood acidity, there are upper and lower bounds for acceptable levels that can be found in the scientific literature of the field, and when the pH values of a patient fall outside of these bounds, something is generally wrong. As such, a test result brought to a physician can automatically state whether or not the results are acceptable. However, every patient is different, and it is difficult to translate blood levels into a diagnosis. To do so requires knowledge about the patient, as well as the diseases the patient has, or is suspected to have.. If a patient has a known thyroid problem, but is in the hospital for cancer treatment, then the oncologist has to keep the thyroid problem in mind as being unrelated to the issue at hand. A quantification and literature regarding the normal values for a quantified measure do not replace expertise, but rather give the physician another tool. A tool that, once again, the physician has to learn to use, and has to acquire expertise in.

2.8: Conclusion

To summarize, this chapter has discussed the way in which medical technology mediates between the user and the world. How the technology interrogates the body of the patient, and how it presents this information to the physician. This brought attention to three different, and important, things. The first of these is the way in which the technology is not neutral, and that different technologies used for the same purpose can do so in very different ways. It mediates and influences. The second of these is the fact that the medical images created by the technologies need to be interpreted. They are not self-evident. The third of these is the relevance of the interpreter, the person looking through the technology. It is the observer that partially determines how a technology is interpreted, and in what ways it mediates.

The observer then, goes through two different layers in order to gain access to the world. First, that of technological mediation, with his medical imaging devices showing only certain aspects of the patient while ignoring others, and constituting the medical image in a specific way. The second of these, is the process of interpretation which, in a strange parallel to technological mediation, means that we cannot directly access the medical image either, we need a way to interpret it, since it is not self-evident but multistable.

One of the things we discussed is how the different technologies can be used in different ways, and how a physician can have a different type of relation with the two technologies, depending on the characteristics of the technology as well as his skill with the technology. We also discussed how the technologies accessed the tumor in different ways. Both of these answer one of our sub-questions: “How do Imaging technologies influence the way we see the patient?” The different imaging technologies grant us access to different facets of the tumor, and the way in which they are interpreted also depends on the technology, as well as the physician performing the diagnosis.

Of course, this leaves us with the other questions. We touched upon the issue of multistability, of what happened when there is more than one possible interpretation of an image. The answer that came up was to use expertise in the interpretation of the images. That then, is the subject of the following chapter, which deals with the nature of expertise and tacit knowledge. What is it that makes an expert and expert? How is an expert different from a layman, and how does he interpret the medical images given to him. I will take into account different viewpoints about expertise, especially regarding the role of tacit knowledge, and context awareness, and talk about how these relate to the interpretation of medical images.

Chapter 3: Expertise

3.1: Introduction.

The previous chapter has discussed several different medical imaging techniques, where they are similar, and how they differ from each other. It also introduced a philosophical framework for discussing these technologies. The postphenomenological approach I used places a lot of emphasis on the technology, and how the technology mediates the relationship between a radiologist, and the patient under diagnosis. However, as explained before, the role of the human in Ihde's technology relationships is not a neutral one. Terms like 'Multistability' and 'Hermeneutic Strategy' show that there is a role for the observer in this technological mediation. Postphenomenological analysis shows that medical images can have many different stable interpretations. At the beginning of the development of x-ray technology, many people saw taking a picture of their bones as a fun or romantic experience. If you gave your lover an image of your innermost self, you were showing your devotion. X-ray images only gained their explicitly medical characteristic once scientists figured out how dangerous they could be. Another good example of a widely multistable imaging modality is that of echoscopy of pregnant women. In Dutch society, so-called 'pret-echos'[17] refer to the idea of performing an echoscopy on a pregnant woman, not to check on the health and growth of the unborn child, but rather to take pictures for the baby-album, and get to know the new addition to the family before he or she is even born.

But even if there is agreement on the purpose of the image, namely that it is to be used for a diagnosis, it is still a multistable object. After deciding to interpret the image as being one of medical significance, rather than a social happening or romantic gesture, the image must still be interpreted in order to decide what it means. A diagnosis, or meaning, must be attached to the image, and this interpretation must then lead to a proper response.

What makes this interpretation of the image so important, is that while an image is multistable, this does not necessarily have to be the case for the medical realities of the patient. Generally speaking, a potential patient is either sick, or not sick. Thus, while the medical image is multistable, there seems to be a correct interpretation, that corresponds to the condition of the patient. When an image is created, it is possible for it to be interpreted as either an image of a body with cancerous tissue, or an image of a body without cancerous tissue. At the moment of testing, these two interpretations can both be equally "correct", so to speak. However, eventually, one will collapse, no longer being a stable interpretation of the image. If a patient dies of cancer, the "non-cancer" diagnosis ceases to be stable.

The problem, then, is how to figure out which of two seemingly stable interpretations of an image is the "right" one, without waiting for the patient to either live or die. One way of doing this is to use additional medical tests to gather more data, but these tests too are multistable, and subject to interpretation. The way in which this is solved in medical practice, is to have an expert judge the image, in order to find the best interpretation. In order to do this, radiologists and other medical professionals go through a long process of training. They learn the anatomical structure of the human body, They learn how to "properly" interpret medical images. That is, interpret them according to the discipline of medicine. This is done, for example, with the help of images where the eventual outcome of a biopsy or something similar is known. In other words, they acquire expertise in the reading of images, and this expertise we connect with authority. We trust and experts interpretation, because he trained for it. If a lay-person interprets one way, and an expert another, then the expert opinion is generally more accepted. The experts are responsible for determining what we know from the image, an epistemological responsibility. Chapter 4 will discuss this idea of epistemological responsibility in more detail, but to do this, I must first talk about expertise itself in this chapter.

First, I will introduce two different frameworks for expertise from the literature. Then, I will compare and contrast these frameworks, and see how they complement each other. Once I have done this, I will discuss this new framework, and the notion of taciticity in it, in regards to the notions of micro and macro-perception. After doing this, I will go through some critiques of expertise, and take a short empirical detour through medical expertise, relating practice to the theories discussed.

3.2: Expertise and socialization

3.2.1: The periodic table of expertise

An expert is a person who is, in some way, better at something then a person who is not an expert, a lay-person. What the expert has that the lay-person does not, is expertise. A simple explanation, but one that does not bring us very far when we actually want to talk about expertise. To give an example, A person that knows that Fat Man was dropped on Nagasaki, and Little Boy on Hiroshima, knows more about nuclear warfare then a person who does not, but this does not seem to fit the idea of expertise in a real sense. In their ' periodic table' of expertises [18], Harry Collins and Rob Evans have created a framework with which to categorize and identify different types of expertise, and their relations to each other. In this framework, they delineate expertise into five rough categories.

The first of these categories is what they call beer-mat knowledge. Much like my nuclear example, this type of expertise relates to the kind of knowledge that can be found on a beer-mat. It is the type of "expertise" that is useful during trivia night, but not in more situations. This is the lowest level of expertise, a simple fact that can be parroted without any understanding of the subject matter.

Beyond this most basic type of expertise comes popular understanding. Where beer-mat knowledge relates to trivia about a subject, popular understanding is a step above it. Popular understanding of a topic can be gained through the watching of a well-made documentary, or reading an in-depth article in a magazine. This type of expertise relates to knowledge that is more advanced than beer-mat knowledge, but the type of information that can be gained in this way is still very limited. The subject is simplified, made abstract, and otherwise different from a real expert working in the field. To give an example, when I was a young child, I had some books on geology, They contained information about plate tectonics, the earth's crust, the formation of mountain ranges, volcanoes and earthquakes, and the movement of continents. I have a level of understanding of these things that is higher than that of someone who has read the phrase "A volcano is a rupture on the crust of a planetary-mass object, such as Earth, that allows hot lava, volcanic ash, and gases to escape from a magma chamber below the surface."(The first sentence in the Wikipedia article about volcanoes) This is a level of understanding that goes beyond knowing such trivia, but is quite obviously still very different from the type of understanding that one gains from reading scientific texts about the topic.

The third category Collins and Evans mention is the having of primary source knowledge. The type of expertise gained through reading the primary literature used by the "real" experts in a field. Someone who interested by a documentary about a topic of research might go online and try to find papers being published in the field, search for videos of conversations between experts, for experts, or read standard text-books used in the field. However, Collins and Evans argue that simply reading the literature, or even performing standard experiments in a university, do not make one a full expert. The primary source reader might know all the theoretical knowledge put into the literature, but lacks the tacit knowledge required to work in a field, the type of knowledge that is far more difficult to transfer via text. The other problem with primary source expertise they identify between a primary source reader and full experts is the matter of what texts to read in the matter of disputed research. In order to properly understand what a scientific field agrees upon, one needs some way to communicate with the experts on what should be taken seriously, and what should not be.

Above these three steps in the ladder of expertise, Collins and Evans come into the types of expertise that require tacit knowledge specific to the expertise. Someone may read everything that is written in and about a subject, but this knowledge remains limited to explicit knowledge if that is all that happens. Tacit knowledge, refers to the type of knowledge that, by definition, cannot be written down. Thus, it cannot be read either. Collins and Evans claim that tacit knowledge can only be gained through interactive immersion. Without being immersed, one cannot fully grasp the nuances of the "language" of an expertise.

Within expertise that requires tacit knowledge, Collins and Evans identify two different types of expertise. Interactional expertise, and contributory expertise. To simplify, contributory expertise in a (scientific) field is the type expertise necessary to work in that field, while interactional expertise is, roughly speaking, the expertise necessary to talk about the field with a contributory expert, and not be immediately unmasked as someone who cannot work in that field. As such, an interactional expert in biochemistry will be able to have a conversation about biochemistry with a contributory expert in the field, and it will not be obvious to this contributory expert that the interactional expert would have no idea what to do once in a lab environment. Later work, by Collins and Evans, as well as others, has further refined the concept of interactional expertise, but more on this later.

There are a lot of texts about the difference between interactional expertise and contributory expertise, and when something counts as interactional expertise. One oft-given criteria for when someone's level of expertise counts as true interactional expertise is the one that Collins and Evans give in later articles on the subject, where the interactional expert can pass a modified form of the Turing-test. In this modified version of the test, the interactional expert counts as such if he is able to pass as a contributory expert in a conversation. In the theoretical test, a tester (an expert in the field) has two different conversations, one with a real contributory expert, and one with an interactional expert. If he cannot reliably distinguish which is the contributory expert and which is the interaction expert during a conversation, the interactional expert counts as a real interactional expert.

Other authors have argued that this criteria is too strict. They claim that being able to pretend to be a contributory expert is not necessary, as long as one is able to have an "interesting" conversation on the topic, on a high level. One of the examples of interactional experts often given in the literature is that of AIDS activists[19], who, although not trained doctors, or able to work in a lab, knew what they were talking about after doing their research. They probably would not pass the proposed Turing-test, but this group was rather different from normal laymen. They knew what the experts were talking about, and could still speak in the scientific language.

Furthermore, one of the primary values of interactional experts is that they are different from normal experts. They have a different perspective from, and often challenge assumptions held by, the normal experts. While an interactional expert needs to be informed about a topic in order to “contribute” to it, being immersed to such an extent that one is able to pass an interactional Turing-test means being able to fake that which can make an interactional expert valuable. This is also where the idea of a linear ladder of expertise, as proposed by Collins and Evans, starts to break down. If Interactional Expertise is seen as Contributory Expertise minus practical ability, then Interactional expertise is indeed lesser, falling somewhere between the contributory expert and the primary source reader. But if interactional expertise is taken more seriously, then interactional experts are interesting exactly because they bring something new to the table. As such, interactional expertise should be seen as a separate category from, but not necessarily beneath, contributory expertise.

3.2.2: Tacticity

Harry Collins’ notion of expertise and the necessity of tacit knowledge mostly comes from the field of Studies of Expertise and Experience, or SEE. As such, while his discussion of different kinds of expertise, as well as the necessity of tacit knowledge in expertise are important, they are only one of many different types of analysis. The classification of types of expertise, and the following discussion of expertise and tacit knowledge, is, however, a good starting point into the discussion.

Collins’ interpretation of the notion of tacit knowledge is based upon its relation to explicit knowledge. Explicit knowledge is that which can be directly expressed, put on paper, shown directly. Tacit knowledge, then, is the knowledge which is not explicit. However, this does not mean that there is only one type of tacit knowledge. Collins identifies several different kinds of tacit knowledge.[20] One of the most important examples of tacit knowledge to Collins is language. Acquiring the ability to speak a language requires socialization in a community that already speaks this language, whether this is a traditional language (in order to correctly speak English, one must socialize with English-speaking people to learn the tacit rules of natural language) or an expert language/jargon (in order to correctly speak in the terminology of gravitational wave physicists, one must learn the tacit rules of their expert language). This is the type of tacit knowledge that Collins calls collective tacit knowledge. However, there is also tacit knowledge that is not explicitly collective, and does not require socialization. A good example for this is driving a bicycle. The driving itself, keeping balance, steering by shifting your body weight, etc. is something that is decidedly tacit. Collins calls this kind of knowledge ‘*somatic-limit tacit knowledge’*. How to then drive your bike through the street, in interaction with other cyclers, cars, pedestrians, and other participants, is part of *‘collective tacit knowledge’*. The greatest difference between these two, according to Collins, is that somatic-limit tacit knowledge is knowledge that, while tacit, can be (theoretically) explicated. Bike riding, for example, can be explained in a set of anatomical positions, muscle movements, and other specific and technical terms once correctly analysed. However, this information wouldn’t help someone learn how to ride a bicycle, because we cannot translate this explicit knowledge into a set of instructions for our body to follow. Applying a torque of 5.42 Newton Meter to the left knee in clockwise direction is easy to write down, but not quite as easy to do. Somatic-limit tacit knowledge is thus tacit knowledge that is tacit because of the human body, not because of any inherent feature of itself.

It is possible for some strange variant of human, (for example, a cyborg) to be capable of translating explicit information to proper body movements, in much the same way as we can use it as a set of instructions for a robot. However, according to Collins, it will never be possible to properly learn how to behave in traffic using only explicit knowledge. There are tacit rules of the road that are taught through socialization and immersion. Participation In traffic is required to learn the rules of the road. Thus, the knowledge required to participate in traffic is tacit because of its own properties, rather than because of biological limitations. Other examples of this type of tacit knowledge is the knowledge required to speak a language. Any explicit rules about the application of language must necessarily already be expressed in a language, thus, some sort of tacit knowledge is required to make sense of it. Furthermore, the social rules of language, or traffic, are always changing, with respect to both location and time, they are a form of 'collective tacit knowledge'. This social type of knowledge is not contained within an individual either; it belongs to groups of people together. If only one man spoke a language, that would be no language at all.

There is a third type of expertise that Collins talks about, called relational tacit knowledge. This type concerns tacit knowledge that is tacit for reasons contained in a culture, there is no deep reason for this knowledge to be tacit, it simply happens to be. This third category of weak tacit knowledge is of little interest to us here, but is included in this paragraph for the sake of completion.

3.2.3: Dimensions of expertise.

In his article "three dimensions of Expertise"[21], Collins discusses different dimensions of expertise, and how they relate to the expert. The three different dimensions he identifies are the level of esotericity, the level of accomplishment within the domain, and the exposure to the tacit knowledge of the domain. The first of the three dimensions regards the ubiquity of the expertise. Although generally speaking, expertise is usually regarded as something rare, this is a strange categorization. Think for example of reading and writing. We do not see someone able to do that as an expert in western countries, but if you go back in time a thousand years, someone with these skills would be highly regarded as an expert. Furthermore, although driving is a skill that is highly ubiquitous, there is still an exam that has to be passed as proof of expertise.

The second dimension, the level of accomplishment, regards what is traditionally seen as expertise, the amount of skill in something a person, or community, has. This is something that generally increases with time and experience. This is related to, but not equal to, the amount of tacit knowledge of the domain. By reading a lot about a topic, the individual accomplishment increases. However, as explained in the beginning of this chapter, Collins believes that only reading about a subject will not lead to "real" expertise. Tacit knowledge is required for this type of expertise. As such, while gaining tacit knowledge, one's accomplishment typically increases, but not necessarily so. One of the examples of different kinds of expertise within these three domains that Collins gives is that of chess.

First of all, the skills of a chess-master are highly esoteric, very few people are chess masters, even though most people understand the basic rules of chess. What makes chess interesting for our subject is that is an expertise in which human players have been irrevocably beaten by computers. As such, the highest level of accomplishment in this field is given to something that is unable to hold tacit knowledge (since any programming is explicit in nature). This means that, if we define winning matches as the measurement for expertise in chess, tacit knowledge is not required for expertise. This raises up two important questions. First of all, what is it that expertise is in? While a computer may have superior expertise in winning chess matches, it lacks the tacit knowledge required to teach others how to play chess the way a chess player can, or to hold a high-level discussion of the game. These are things that require knowledge about other things than the purely explicit rules about game-piece movement and win conditions. The second question is what type of tacit knowledge is required for chess. An experienced chess player sees patterns and lines that are difficult to describe in words, something that falls under tacit knowledge. The question is whether this is somatic-limit tacit knowledge, or collective tacit knowledge. If it takes the form of somatic-limit tacit knowledge, then it is possible to explicate the knowledge required for doing this and giving that expertise to a computer program. For example, a chess player may "feel" that two rooks for a queen is a good trade, and gain this understanding tacitly. If you gave him the information in a mathematical equation or some other explicit format, he would not gain that understanding. In contrast, the computer can only know it explicitly through its programming. However, there are other situations in chess where more tacit "if I do this he'll think I'll follow up with that" situations come in where a player gains a 'feel' for the opponent, something far more social, which can fall under collective tacit knowledge.

3.2.4: Application to imaging

To Collins, the reason we confuse somatic-limit tacit knowledge and collective tacit knowledge is because humans acquire them in the same way, by interaction, instead of explicitly taking in information. What is interesting here, is where medical diagnostic skills would fall. Can they be categorized as somatic-limit tacit knowledge, or is it collective tacit knowledge? Or, to make matters more complicated, is it dependent on the type of imaging technology used?

If it is purely somatic-limit tacit knowledge, then that would mean that it is perfectly possible to write a computer program that is able to diagnose in the same way (or better) as a medical professional. If the tacit knowledge involved is collective, then this becomes more difficult. And to return to the example of chess, it is also possible for an expert and a program to diagnose in different ways. And, above all, we should not forget that the job of a radiologist is not simply to diagnose, but rather to help his patients, something in which diagnosis only plays a small part, and a process which is far more likely to require collective tacit knowledge

Furthermore, tacit knowledge also forms a problem when looking at differences between imaging techniques. X-Ray Imaging and photoacoustic imaging give the user a different type of image, they show different things in a different manner. In order to form a hermeneutic strategy to interpret a medical image, the user needs to know how the technology works 'under the hood' so to speak. However, reading a textbook explaining the difference between the two techniques, and information on how to recognize specific structures on an image, only conveys explicit knowledge. Even if the type of tacit knowledge necessary for interpretation is only somatic-limit tacit knowledge, simply describing the new technology and how it is different is not a guarantee that this is enough for the expert to properly use the new technology. Here we can return to the bicycle example, while an explanation of muscle movements and torque is interesting, it is not very useful for biking, and an explanation of the differences in riding a bicycle and a motorbike will not be enough to teach someone how to ride one or the other. A human radiologist requires tacit knowledge/experience, in order to be able to form a good hermeneutic strategy to interpret an image. And while a large part of that tacit knowledge might be transferrable, (how to talk to a patient does not change), there is still new tacit knowledge necessary.

Tacit knowledge is also an important concept when we talk about why we should trust experts. In a perfect world, a radiologist would be able to perfectly explain to a layman how he reached his conclusion regarding the diagnosis of an image. As such, we could listen to this explanation, would come to the same conclusion, and thus trust in the expertise of the expert. However, the diagnosis requires tacit knowledge, making it impossible for the expert to explain everything about the diagnosis to someone not immersed in the practice of radiology. This creates an epistemological problem. Should we accept the diagnosis of an expert as true of he cannot explain how he reached this conclusion, because he cannot provide the tacit knowledge upon which it is based.

3.3: The Dreyfusian Model

3.3.1: The model itself

A different notion of expertise is introduced by Stuart and Hubert Dreyfus.[22] Their theory about expertise is a different one then the theory used by Collins and Evans. The differences can be found in several places. First of all there is the type of activity in which expertise is found. The Dreyfusian model is based in specific skills, rather than the broader fields often found in Collins' work. It locates expertise in changes in understanding and the internalization of skills, rather than increases in explicit and tacit knowledge. The model works with five different stages of expertise, one logically following from the previous stage. This is also different from the periodic table proposed by Collins and Evans, in which people need to do different things for different levels of expertise, and may skip certain stages as well. The five stages in the Dreyfusian model are: Novice, Competence, Proficiency, Expertise and Mastery.

The first of these stages, the Novice, concerns an individual with an instruction manual, someone who follows explicit rules in order to solve a problem. These are rules concerning context-free features. For example, a chess novice will not only know the rules involved in playing a game of chess, but will also be able to understand the pieces as having a specific value, as lined out by a chess guide. He will be able to move his pieces, and make a few decisions based upon the rules (for example, always trade a rook for a queen if possible), but he will not have a deeper understanding of why these rules apply, or when they don't apply. Context-dependent decisions based upon the flow of the game and overall positioning on the board will not apply.

After this stage, the competent practitioner arrives with a large amount of experience with real situations, rather than merely rules and instructions. Recurrent patterns are picked out by the competent person, these are called aspects. These recurrent patterns can be called to by an instructor, in the giving of 'guidelines'. Something more efficient than simple rules, they are dependent on experience instead of merely explicit language. Patterns, instead of specific situations, are what guidelines refer to.

At the third stage, someone is said to be proficient. At this stage patterns are related to goals, and judged upon their relevancy. Rather than merely recognizing recurrent patterns, the practitioner is able to judge whether or not the patterns are relevant, and what to do in case of which patterns and goals. Two objectively equal situations with different goals are treated differently. The way in which the relevance of patterns relate to goals are called maxims.

The fourth stage, the expert level, is the stage at which the earlier rules, guidelines, and maxims, become less relevant, the expert has internalized his experiences and does what he does intuitively. Every situation has an associated response, and because the repertoire is so vast, this counts for all situations, even if that specific situation has not yet been encountered, it is similar enough. Instead of being analytic, experience becomes intuitive. In the terms used at the start of this chapter, explicit knowledge becomes, at least partly, tacit.

The fifth stage, mastery, is not truly a category, but rather a sub state of expertise. Mastery happens when someone at the fourth stage, an expert, is in a state in which he ceases to pay conscious attention to the performance. As such, attention can be given to other tasks instead.

There are several interesting things about this model of expertise. First of all there is the intuitive nature of the work of an expert. An expert makes his decisions not based upon rules, guidelines or maxims, but rather based upon an intuitive understanding of the situation. Things like rules and guidelines are useful during the acquisition of expertise, but once someone becomes an expert, they are no longer necessary. In fact, almost the opposite is true. One significant factor is choking,[23] Choking is a phenomenon where, when asked to think about the activity, the skill of an expert decreases. The idea behind this being that, when attention is called to the activity, the expert will think about the explicit rules, guidelines and maxims he has been taught, rather than use his intuitive understanding of the issue. This can be seen in increased chances of failure in high-pressure environments, for example, high-stake moments in sports, in which thinking about the movement decreases the skill of the sportsman, and the pressure makes it difficult not to think explicitly about what you are doing.

The main similarity between the two models I have explained is their focus on the place of tacit knowledge in expertise. Both models agree that true expertise requires tacit knowledge, knowledge that cannot be made explicit. However, they identify this tacit knowledge in very different places, and it is acquired in different manners. Where the tacit knowledge described by Collins is acquired through socialization and immersion, Dreyfusian tacit knowledge is gained through experience and internalization.

3.3.2: Dreyfusian expertise and medical imaging

Much like the previously explained model, the Dreyfusian account of expertise can be related to the topic of medical imaging, and the difference between x-ray and photoacoustic imaging, as well as diagnosis using quantified measurements. In regards to using a medical imaging technology, a novice learns to recognise anatomic structures, she learns how to relate these structures to each other as a competent reader, and then learns how to use them to diagnose. However, when the level of expert is reached, these once explicit rules found in text-books are no longer applicable. Understanding how a technology creates an image is no longer something that is explicitly thought about, but is present tacitly in the analysis. This makes changing between technologies difficult. While it is possible to explicitly state the difference in image formation between X-ray imaging and photoacoustics imaging, this cannot be done tacitly. A physician might explicitly know the difference between the technologies, but without experience, this information is not very useful. However, this does not mean that changing between types of imaging means that a physician has to completely learn a new technique. First of all, his anatomical knowledge remains applicable. Furthermore, expertise in the Dreyfusian model is marked by the ability to apply skills, rules and guidelines in new situations. As long as two imaging techniques are relatively similar, some of the mental techniques used in interpreting one image can, with the necessary expertise, be used for the other type of image.

The other interesting topic is quantified diagnosis, the interpreting of an image by a machine. The diagnosis of a medical image is a skill, one in which a practitioner becomes better as he gains more expertise. However, take the analogy of chess again. A chess grandmaster is an expert, he has experience and knowledge about the game, and an intuitive understanding through his tacit knowledge. However, that does not mean that he is automatically better at chess. A computer, perfectly following explicit rules, can beat him. In the same way, a computer program using quantified (or unquantified) data could, at least theoretically, outdo a human in the diagnosis of cancer. If the explicit rules are good enough, precisely following them may lead to better results than an intuitive understanding.

3.4: Unifying views of expertise

Of note in discussing Dreyfus and Collins is that, while they both discuss the topic of expertise, they approach the subject from different angles. Collins is mainly focused around academic expertise, with experts being professionals in a community of other experts. The Dreyfusian model approach expertise from the topic of skills instead. Both of these seem to be applicable to radiologists. On one hand, a radiologist will learn specific things about reading medical images, which are then, with experience, internalized in the way described by Dreyfus. On the other hand, radiology is very much an academic field, and in order to truly be a radiologist, one must learn the ways of radiology. The profession cannot be reduced to one specific activity.

Collins view of expertise includes knowledge and skills on different levels. There is the level of overarching theories, and there are the skills necessary for scientific experimentation, both of which are necessary for a contributory expert. But the lynchpin is placed in the social context of an expertise, the language of it. Not just in talking to other experts, but in understanding the literature and what theories mean. The Dreyfusian model is based around the increase of skill in an activity after practice. Taciticity is located within the forgetting of explicit rules. Once someone is an expert, the brain works in a way that is better than merely following explicit rules, by seeing situations and judging them intuitively.

What makes combining the two interesting, is that someone who is an expert in the Dreyfusian sense need not be an expert in the way Collins describes expertise. Someone who is very good at a skill, but not part of a greater community of practitioners, not skilled in a 'language' belonging to the expertise, is not necessarily an "expert". Someone could be an amazing physicist, but if he does not, cannot, communicate with the greater community of experts, then there is neither interactional nor contributory expertise. Eventually, he needs to leave his laboratory to discuss his findings and publish them. The question then, is whether this is also true the other way around. Is it possible to be a contributory or interactional expert without being an expert in the Dreyfusian sense. In this, I suggest that we look at the difference between contributional and interactional expertise when trying to find Dreyfusian expertise. It is not the social part of the expertise, but rather the specific skills, that the interactional expert is missing. As such, a large part of the difference between the interactional and contributory experts can be found in the Dreyfusian expertise that contributory experts have. As such, the models of Dreyfus and Collins complement each other. Collins' socialization is what makes a physician a physician, but Dreyfusian skill acquisition allows the physician to actually perform tasks beyond interaction with other experts.

Combining these two different viewpoints has interesting applications when we look at medical imaging. One of the observations about different imaging technologies given in the previous chapter is that they look at the same object, for the same purpose, but do so in drastically different ways. As such, the process of interpreting the image differs between different technologies. The skills necessary for the interpretation of the images created are different, and from our discussion of tacit knowledge, we know that simply knowing explicitly what the exact differences are in the constitution of the images is not enough to transfer expertise between the diagnosis of the two types of images. However, the context of possible patients, cancer and treatment options remains largely the same between different technologies. As such, the tacit expertise found in Collins' model does not change overly much when changing imaging technology.

This also links back to an observation back in paragraph 2.5, regarding the sliding scale between Embodiment relationships and Hermeneutic relationships. The process of the use of a technology becoming like a second nature to the user, and therefore the shirt from Hermeneutic to Embodied relationship, is very similar to the Dreyfusian expertise, and the role of tacit knowledge in this framework. At the start, the use of a technology is explicit, and more Hermeneutic. Data gathered from technology is read and interpreted through a strict framework. With experience, this process becomes more tacit, and at the same time more embodied. As such, we can say that the two concepts are related, with explicit knowledge becoming tacit knowledge possibly leading to a more embodied human-technology relationship.

Quantification and computer diagnosis play a role here too. The technology is able to diagnose automatically based upon patient data, thus taking over the role of the radiologist in this way. However, they lack the wider societal context, and a lot of the nuances of a skilled human expert. Statistically, they may be better and more efficient than a human expert, but they do so in a different way. As such, an automated system does not count as an expert in the Dreyfusian model (no skills are internalized and made tacit), but is able to replace it. If we zoom out, we see that the contribution that makes the contributory expert into what he is seems to be partly taken over by a machine. However, the wider societal context of the problem remains, some form of expertise is still necessary to perform the role of the expert. More on this topic will follow in Chapter four, regarding the epistemological processes involved the interpretation of an image, and diagnosis of a patient. This chapter will now continue in a discussion of perception, and the role of tacit knowledge, expertise, and culture within this, to see how the viewpoint of an expert is different from a layman.

3.5: Micro-perception, Meso-perception and Macro-perception

3.5.1: Perceptions

Our observations, our ideas about the world, the diagnosis of a patient. These things are based upon the senses. Our eyes see light coming from the world, our ears hear the sounds of the world, our fingers feel the world. We interact with the world with our bodies. The direct perception of this is what Don Ihde calls Micro-Perception. The warmth of the sun on your skin, the colour of the pixels on a screen, these are Micro-Perceptions. Complementary to this is Macro-Perception. Macro-perception is what comes into play when we look beyond the Micro-perceptions, to what we perceive. The pixels on a screen become images that we recognize, a symbol gains meaning, a sound becomes a speech about hope or hatred. Macro-perception is about the way we perceive the world, the way that our ideas, the concepts we know and our understanding of things colours what we perceive. Heideggers notion of Gestell, or Kuhn's notion of the scientific paradigm, are examples of macro-perception. Overarching ideas that influence the way we see things. Where primeval man sees a barren desert with black goop, the modern man sees a source of oil. To apply these concepts to medical imaging, micro-perception regards the way the light from the pixels on the screen reach the eyes of the beholder. However, when this information reaches the brain, the physician has to interpret the image, a process of macro-perception. One is seeing, the other is understanding what is seen.

What is interesting is that they are not truly separate things. There can be no macro-perception without the micro-perceptions of the body to base its ideas in. In order for the 'higher order' of macro-perception to be possible, there must first be micro-perceptions. However, Don Ihde claims that macro-perceptions also influence micro-perception. Our culture, the ideas we hold, influences the way we see the world. This is an easy claim to be sceptical about. After all, light illuminating cones and rods on our retina seems to be independent of larger cultural frameworks, physics doesn't seem to care about things. But we do not actually see activated cones and rods in our retina, we see things. If someone does not know the symbol for radioactivity, then he will not see it if it stands amongst many other similar patterns, but if someone does know the symbol, and its significance, then his sight will be drawn to it through simple pattern recognition. We can also look at the context of the medical image. A lay-person and an expert will, on one hand, see the exact same thing. However, the experts macro-perceptual frameworks will include what is seen as the relevant markers for a disease, micro-calcifications in breast tissue for example. Where the eyes of the Lay-person will look right past these specific details, the macro-perceptual framework of the expert will be drawn to them, a seemingly irrelevant feature taking central stage.

The term macro-perception is, however, overly broad. Therefore, for this thesis, I will make the distinction between two subcategories in macro-perception. The first of these is the large, overarching cultural type of macro-perception. One that is shared by all people in a civilisation, based upon the culture of those people. In the introduction, I gave the example of a photograph. What we see when we look at a photograph, our micro-perception, is based upon our macro-perceptual framework. We see it as an image of the world, rather than a collection of ink colours on paper. This type of macro-perception is on the level of Heidegger's Gestell, a culture-wide type of interpretation that influences the specifics of what we see. Something that is almost inescapable, part of who we are as people. The other type of macro-perception is, instead of Heideggerian, more Kuhnian in nature. A scientific paradigm is something that influences the way we see things, with useless noise becoming important scientific principles during a paradigm change, but it is on a decidedly smaller scale then Heideggers Gestell. Kuhn's paradigm allows for the existence of several different paradigms, although not shared by the same person. In this sense of macro-perception, there is a difference in interpretation within cultures. After all, a doctor works under a different paradigm then a mathematician.

The first of these categories then, I will keep calling macro-perception. The second however, the one concerning factors that can differ between people in a culture, of even change over relatively short amounts of time in a person's lifespan, is meso-perception.

As said before, imaging technologies such as PA require mathematical reconstruction before an image is created. This requires certain choices to be made by the engineer responsible for the creation of this reconstruction. However, the engineer has a different background than the medical personnel that will be using the mathematical reconstruction. His meso-perception is different. Thus, he will notice different things in the images he creates, and he will create an image that fits what he sees in the data. What is dangerous here, is that instead of simply seeing different things, the engineer might remove those things that he does not perceive in the creation of the image. Thus, instead of simply not seeing these details in the image himself, he makes them invisible, the details are fully removed, and cannot be seen by the medically trained user of the device.

3.5.2: Tacit knowledge, Heidegger, Macroperception

Ben Trubody has analysed Collins' theories about tacit knowledge from the perspective of Martin Heidegger, with a focus on the idea of collective tacit knowledge [24]. According to Trubody, the problem with Collins' ideas about tacit knowledge is that he places the hard problem of tacit knowledge in the socialization problem. Collective tacit knowledge is seen as being the cause of the problems computers have with natural language, the information can only be conveyed by socializing with a group, and thus, a computer can never be an expert quantum wave physicist, because it cannot socialize with quantum wave physicists. To Trubody, the secret to inexplicable tacit knowledge should not be sought in the process of socialization with others; instead, he finds it in Martin Heideggers work, and the notion of "Being-in-the World". In Trubody's interpretation, Heidegger talks about "beings" (objects) and "Being", with beings being the objects and conditions around us. Being, however, refers to the possibility for things to show themselves as things, beings can only be because of our "Being", humans realise we are objects ("being") but also reflect on ourself and the fact that we are. To Heidegger, this awareness of ourselves and our condition, this "Being" is primary to "being". We cannot explain Being in the terms of being, and to try this is the primary sin of western philosophy. The human condition of "Being" cannot be described in the terms of being. Trubody connects Collins' notions of explicit knowledge and tacit knowledge with Heideggers present-at-hand and ready-to-hand. Explicit knowledge has the characteristic of that which is present-at-hand. It is about an object that is observed, that can be thought and talked about. It is a 'being'. However, once such an object is used, if it is well-made and well-used, the object disappears and becomes ready-to-hand. A hammer, for example, disappears in the hands of the skilled carpenter. It is not thought about, this Trubody thus connects to Collins' taciticity. To Heidegger, science is concerned with beings and beings alone, not "Being". "Being-in-the World" refers to the idea that we belong to our world, and it is only through this that we are able to experience things. Trubody claims that the social tacit of Collins is not tacit enough. Explicit accounts of knowledge are present-at-hand accounts, and tacit knowledge is ready-to-hand, but both of these require "Being-in-the World". Trubody claims that we should instead locate the tacit in "Being" and the "World", the fore-structure of understanding the world, instead of in a specific non-explicable understanding of the world. The taciticity that Collins talks about should be seen as ready-to-hand, not as true tacit-ness.

Being-in-the-World, then, seems to hold a level of taciticity that goes above the type that is discussed by Collins or Dreyfus. It regards the necessary knowledge to even begin understanding the world, and what it is to live in it. This, then, makes it very obvious why there is a difference between meso-perception and macro-perception. When we talk about tacit knowledge involved in macro-perception, it is nearly impossible to even talk about it, whereas when we look at meso-perception, we can talk about how the knowledge is acquired, and what it does, just not the explicitly the contents. This also makes it clear why people can have different meso-percepetions, but to have truly different macro-perception seems unlikely, our mode of Being in the world is relatively similar. That is, as long as we remain human. The Heidegerrian tacit knowledge that Trubody talks about is something that all humans hold, but disappears when we look outside of the human experience. A computer does now know what it is like to Be, it can only be. As such, even if we teach a computer to talk in the language of the expert, and to make judgements better and faster than the expert, it still does not have the tacit knowledge required to be an expert.

3.6: Critique of expertise

Earlier in this chapter, I have explained two different theories about what expertise is, and how to classify it. A repeating theme in this was that an expert was someone who was very good at something. Leaving something to the experts was seen as a good thing. The expert has more knowledge, both explicit and tacit, to base his decisions upon, thus the decisions have a higher quality than the decisions made by lay-people. However, not everyone agrees with this interpretation of the role of experts. There are people who criticise experts, and their role in things. In his article “Feyerabend's democratic critique of expertise”[25] philosopher Evan Selinger brings together some of these criticisms, focussed around the viewpoint of Paul Feyerabend. The criticism that Selinger describes comes from several different angles.

One of these angles is related to the issue of meso-perception described above. This angle of attack is based upon the idea that an expert not only has strengths, but also weaknesses that make flaws in his reasoning invisible. The claim is that there is a dogmatization within scientific expertises, because experts are unable to critically evaluate the underpinnings of their expertise. Their worldview means that, in some cases, non-expert are better at locating prejudices in the work of experts. In dreyfusian terms, the underpinnings of a discipline become tacit knowledge, making it difficult for the expert to give them an explicit and critical look, while a lay-person would have an easier task of this. The meso-perceptual framework is a burden instead of a blessing.

The second angle of criticism is found in the case of David Sackett. A researcher that did work in therapeutic regime compliance, evidence-based medicine, and randomized trials. Sackett noticed that his work had an authority beyond what his work deserved on its scientific grounds. Thus, because he was seen as an expert, he found that people placed too much confidence in his work. More generally, the authority of experts means that new ideas, that are not accepted by the current experts, have a much harder time being accepted, slowing the march towards truth. To avoid this problem, Sackett changed fields from compliance to evidence-based medicine, and later from evidence-based medicine to randomized trials after seeing his influence and authority.

Feyerabend’s democratic critique, then, considers how to properly control experts. According to Selinger, Feyerabend believes that the knowledge experts hold is not as hard to acquire as it seems to be, and that experts are okay with this because it increases their prestige. As example, Feyerabend gives medical training during the World Wars for soldiers as an example. A medic could be trained in half a year, while a doctor requires four years or more, thus, when necessary, the expertise is easier to acquire. This, of course, leaves out the difference in materials covered and the level of expertise reached during the training. In order to control experts, who would otherwise rule their area of expertise, Feyerabend proposes democratic oversight, with experts explaining their theories and methods to a jury of lay-people, and these lay-people deciding upon the validity of the theory. Besides the problem of the subject material being too difficult for the average jury of lay-persons, there is another problem here having to do with tacit knowledge. It is simply impossible for an expert to explicate all the knowledge that went into his decisions, thus it is impossible for the jury to make a decision based upon the facts presented by the expert. By forcing the expert to explicitly give all of the knowledge that went into his decision-making process, Feyerabend's proposal throws away the very thing that makes an expert an expert, his tacit knowledge.

All in all, almost all of Feyerabend's critiques break down under inspection, only the meso-perceptual problem remaining. Much like how a technology mediates the users view of the world, enhancing some parts and making others invisible, so too does expertise hide some of its own details. Something that can be circumvented in much the same way as the problem in mediating technologies, which will be discussed in the next chapter.

3.7: Medical expertise in practice

In addition to the above, which has been a more philosophical look, in the following section, I will show the expertise involved in the reading of medical images from a more practical, scientific perspective, while linking this to previously discussed concepts.

3.7.1: Gathering information

In order to gain the information that is necessary for a proper diagnosis, a doctor will often need to talk with a patient. In such a doctor-patient interview, the patient will usually tell the doctor a story about their situation. It is the doctor's responsibility to take this story and gather the relevant facts from it, this information can range from what kind of problem the patient has been experiencing, to medical history, to general lifestyle. Such a conversation can be rather broad, because the patient does not know what facts that may be relevant to the case. The doctor has to ask for both questions and clarifications in order to obtain the information that he needs. Furthermore, the doctor then has to translate the patients answers into medically relevant information. On one hand, this can be a literal translation. For example, when a patient points to a location on his body and explains that "this bone here hurts", a physician will then translate that as something like "proximal ulna", using the language specific to his expertise. On the other hand, this can also be more interpretative, with the doctor realizing that a specific pain is usually described as a specific disease profile [26].

The doctor has to be able to navigate a social situation in which medically relevant information is gathered. To do this, the doctor not relies not only on his explicit medical knowledge, but also on the tacit knowledge necessary to interpret subjective statements and talk with the patient. "people skills" are involved. And while there is a difference in expertise between the patient and the doctor, they have the shared experience of being human, of existing in the world. When a patient describes being tired, or dizzy, or a specific type of pain, the medical expert is able to understand this because they are based upon a shared experience in a way that would not be possible of the doctor did not have a body just like the one the patient has. In addition to that, if a patient describes a subjective phenomenon like pain, it is up to the physician to interpret what that means. Is the patient a strong, though outdoorsman who says that he feels 'a tiny sting' when his arm is broken, or is he secretly a rather meek individual that starts crying at the slightest needle-prick? In this, the necessity of being able to manoeuvre in a social setting becomes clear, and with it the importance of expertise in navigating this situation. There is more to this aspect than simply good bedside manner.

3.7.2: Prewhitening.

To borrow from the terminology of the field of image processing, an image consists of two different things. There is the signal, which is what the observer wants to see, and the noise, which is everything else in the image. An image can be seen as perfect when it consists wholly of signal, without any noise to speak of. Of course, how the signal is defined depends upon what the observer wants to see. When looking for tumors, the tumor can be seen as the signal, while the noise is everything else. Part of this is like the noise you hear from a badly tuned radio, or an old tv with a slightly lose cable. Something happening in the background that slightly changes every part of the signal stochastically, making it impossible to retrieve some of the information. Noise caused by faulty, or less than perfect, equipment or simply a feature of the system that is being studied. However, there is also the rest of the body, which can also be seen as noise. In the perfect imaging technique, you would see the tumor, and only the tumor, not the tissue surrounding it.

Thus, in breast-cancer imaging, the signal is the tumor, or rather, the marker for the tumor. Noise, then, is filtered out in different ways. When comparing older x-ray images to newer ones, a drastic reduction in noise is almost immediately visible, because of a variety of techniques and new or improved technologies. One of the ways to reduce this kind of noise is with mathematics, and one of these ways is by pre-whitening, a term that comes from the field of signal compression. In this field, prewhitening is a process that subtracts a predicted value from a signal, to minimize data. To give a simplified example, think of the tide. There is well known pattern to I, it generally goes up and down. To "whiten" the noise in this case is to make every data point independent of the previous one, by removing the general trend. For example, by taking average data about water and tidal movements, and subtracting that, all that is left is the variation in water height not dependent upon tidal movements. The signal (small changes in water height unrelated to the movement of the tides) remains, while the noise (changes in water height due to tidal movements) has been removed, the ratio of signal to noise is greatly improved.

In the analysis of diagnostic images, something similar is happening. The observer that looks for something filters out structured noise in his head, thus, humans are able to 'prewhiten' the image in their head. This can be done to artefacts creating by image processing, but also to the anatomical background. When looking at an image of a patient, an expert is able to partly remove the underlying, expected, structure in his mind, seeing only what is different from the norm. One of the ways this happens is by doing comparing scans. By tracking eye movements, scientists have shown that radiologists sometimes do this by quickly moving their eyes between areas on the image. A radiologist can do this to compare left and right breasts, looking for differences that might indicate a tumor (mentally "subtracting" one from the other, leaving only the possible signal) but also by looking at areas of similar densities within the same breast [27]. This is a skill honed through practice, tacit knowledge that allow the physician to know what to look for, and what not to look for. Through practice and smart thinking, the physician knows what is signal, and what is noise. But this does not seem to be purely tacit, there is a conscious element to it, and it can be easily explained what is being done when, for example, comparing images of two different breasts. This then, seems to be similar to Collins called somatic-limit tacit knowledge. The physician knows what he is doing, and he can explain it, but the expertise inherent to doing it is tacit. It is perfectly possible for a skilled programmer to listen to the physician's explanation of what he is doing, and write a computer program that is able to 'prewhiten' the image, but the physician does not use mathematical algorithms but his own skill and expertise.

3.7.3: Scanning.

The human eye does not see everything equally, and most of its power is focussed in one specific, and small, area, the fovea. An area that contains far more receptors then the rest of the eyy. This central area of the eye is responsible for sharp vision, with the rest of the eye providing the peripheral vision, something noticeably less powerful. What this means in the context of medical diagnosis, is that not all of the image can be observed at the same time, instead, the image has to be scanned. The eye has to "move over" the image in order to see every part of it. These movements can be tracked by a camera, tracking the eyes. Visual search patterns are thus not only an important part of the expertise of a physician, but also something that can be researched. There are several different visual tasks involved in the work of a radiologist. The Handbook of medical imaging. Vol 1, physics and psychophysics[28], classifies them as detection, comparison, location, classification and estimation. The ones that are most important for our purpose are detection and classification. First, a possible tumor must be detected in the image, then, it must be classified as being a tumor or not.

By doing tests, it can be seen that the scanpaths of experts and laymen differ drastically when looking at an image. The expert has experience in looking at a type of image, and knows what to look for. The expert's expertise influences not only the search pattern, but also when he fixates on a part of the image, taking a closer look at an area. However, when research was done in order to understand how large the area of focus has to be by artificially limiting it, they found that radiologists actually preferred having visual information beyond only the area focussed on by the eye. Furthermore, experts are able to scan for abnormalities much faster than can be explained by just their foveal cone, the main area of focus in their vision. They are able to use their peripheral vision much better than less experienced observers. Where beginners will follow hard lines in the images in their search patterns, those of trained radiologists will focus on specific areas and locate abnormalities faster. This is very reminiscent of the Dreyfusian model, with the beginners still following rules and guidelines, trying to work through systematically, the expert is able to take in the wider picture, and looks at what is important. His mastery allows him to spread his attention, taking in information outside of the foveal cone.

When looking at the difference between beginning radiologists (residents with little real experience) and full radiologists, a study[27] showed that experienced readers were easily able to shift between areas of interest while searching. They would locate one lesion and look in the area for other signs, then immediately move on to the next. Less experienced readers were also able to quickly find the first lesion, but needed a systematic search of the image in order to find further lesions. When no areas of interest were found immediately, both experienced and less experienced radiologists went through the entire image part by part, in order to make sure they didn't miss anything. This again is reminiscent of the Dreyfusian model, with the more systematic approach being rule-based. The expert no longer needs this for a large portion of the work, but given the nature of the activity, still uses the rules afterwards to make sure.

What is interesting here is that it shows that the expertise of a radiologist not only means that he is, statistically speaking, better at his job, but also that he works quantitatively different from the less-experienced. In chapter two I discussed the difference between Hermeneutic and Embodiment human-technology relations, and explained how differences in expertise could create qualitatively different relations. Here then, we have a piece of practical evidence for this, showing how the way the medical images are interacted with changes for experts. Systematically working your way through an image looking for regions of interest can be said to be more Hermeneutic in nature, while organically moving through the image is more Embodied, no longer requiring the strictness. Interesting then, is that while the expert radiologists are capable of this more embodied process, they afterwards shift to the more hermeneutic style of image analysis, in order to double-check their work. The expert, in this case, does not just have a different approach, but is able to follow two different approaches, depending on what is necessary.

3.7.4: Scanning and recognition.

The process of diagnosing an image is not limited purely to searching the image for possible tumor markers. Once something is detected, it must then be classified (mentally). As such, even if the eyes focus upon a suspect area in the image, this does not mean that it is automatically regarded as being dangerous tissue, which would result in a large amount of false-positives. For example, a study [29] about chest x-ray images has shown that about 65% of missed nodules (the visual objects that radiologists are looking for) are fixated upon by the foveal cone. In other words, even though the radiologist's vision focussed upon the object, it was not recognized for what it was. In another study about breast cancer images specifically[27], when dividing into search errors, recognition errors and decision errors, with search errors being those errors where the eye did not fixate on a location, recognition errors being those errors where gaze duration was less than one second, and recognition errors those where fixations took longer than a second. The study found that, of the experienced radiologists, two false-negative results were search errors, two were recognition errors, and four were decision errors. The inexperienced image readers had three and a half times as many false-negatives at 21, with six search errors, nine recognition errors, and six decision errors. As such, the difference experience, expertise, brought, was not only a reduction in amount of false-negatives, but also a change in where the errors were made. This is interesting when we return to our descriptions expertise, as the three different errors seem to be in different categories. Search errors are almost completely tacit, depending on skills honed by experience, while recognition errors and decision errors are more explicit in nature. Especially the decision errors, in which the radiologist decides whether a given area of interest is a tumor or not, can be made explicit. These errors too, rely on background knowledge from the field and how the images are constituted, unlike the experience of searching through an image. The expert, when compared to the novice, does better in both these areas, but makes the least improvement in the area of decision errors. It seems than, that the difference in explicit knowledge is relatively small when comparing the novice (who has gone through medical school) and the veteran.

3.8: Conclusion.

Again, we return to the questions asked in the introduction, and partially answered at the end of chapter two. Here, we have answered the question of how the correct interpretation of an image is decided upon. Through expertise, which gives the physician the skills, knowledge and mesoperceptual framework he needs to make the correct decision. This expertise contains several components, and knowledge both explicit and tacit. In addition to that, what is also important is the physician himself, the fact that he has a body, and is located in the world, which gives him tacit knowledge he requires which could not be programmed or otherwise placed in technology. Furthermore, the expert is embedded in a context, a group of people sharing the same language of expertise.

When we go back to the question of how different imaging technologies influence the way diseases are diagnosed, we are now also aware of the ways in which they do not. While Dreyfusian expertise is very much connected to the specific technologies, the wider, more social expertise that Collins talks about is separate from the specific, and embedded in the expertise itself, medicine. Thus, while some things change, other things remain the same. That is, of course, in the case that the technology does not cause the field itself to change.

Now that we have answered this part of the question, what remain are the main question of what happens when two technologies do not converge, and the question of the existence of disease itself. The next chapter will talk about epistemology, the theory of knowledge and when we can hold a belief to be justified. Ultimately, that discussion will build on this one, and shift to the topic of the epistemological responsibility of the expert.

Chapter 4: Epistemology

4.1: Introduction

In chapter two, we have discussed how different technologies mediate out relationship to the world in different ways, as well as how the way in which this is done depend on the user. The different medical technologies discussed seemed to look at the same object in drastically different ways. Both in how they interrogated the object, and in how that information was then read by the user. The question that was asked was how to interpret what was seen, if a medical image is by its very nature multistable. This was partially handled in the third chapter, dealing with expertise.

The question that is central to this chapter then, is whether or not what we see through technology is truly there in the way we see it. Is there a way to see the world itself through technology, or do we only see a mediated version that always has something of the technology itself in it? Or is the world itself constituted through technology, and does it make no sense to talk about unmediated access? Was the tumor that is detected by the instruments there before it was detected, or is it only there because it has been detected?

What makes our specific case especially interesting is that it lies in the medical field. Unlike philosophers, who can write treatises, publish books, and debate each other for decades, physicians must act, and their actions have drastic consequences. This brings a level of pragmatism to their decisions, they simply must decide.

This chapter will begin by looking at Don Ihde's theories of instrumental realism, and try to see how he attempts to answer of whether or not what is observed through scientific instruments is truly there. Following that, I will introduce some points of critique related to the earlier chapters, and the role of expertise in the interpretation of data.

4.2: Instrumental Realism

Scientific instruments allow us access to parts of the world that cannot normally be seen. They are the technology in the man-technology-world schematic. Entire fields like microbiology, particle physics, astrophysics, and many more, depend entirely upon instruments. The objects of their study are almost exclusively detected through instrumentation.

In, for example, a thermometer, we have an instrument capable of measuring something we also have direct access to. A thermometer can give us a more precise measurement, but it measures something that we know is there. However, many scientific instruments allow us to measure things that cannot be seen. We do not have access to the raw phenomenon, we only know it by its measurement. Even in our specific case, the tumor is not directly observed during screening, it can only be "directly" observed by cutting open the patient, and even then microscopic evidence is usually preferred.

Historically speaking, science has often posited things that were incorrect, and which did not actually exist. Temperature has been theorized as a type of particle, the substance of phlogiston was thought to play a role in combustion, and the sun revolves around the earth. This brings us to the question of scientific realism, whether or not the objects studied by science truly exist in the way that scientific theories claim they do. Are science and its instruments a tool that bring us true knowledge about the world, or are they simply models that, although they may be useful, have no real relation to the actual state of the world?

In other words, a scientist's access to the object he studies are always mediated. He can only see the object through his instruments. But, as we have seen, scientific instruments both mediate between scientists and the objects they study, and they create multistable hermeneutic "texts" that must first be interpreted. Thus, the way the scientist thinks, what he believes about the world, and his expertise, influence how he interprets the data from his instruments, which, in turn, influences his theories about the world. This becomes more complicated when we look at different instruments used to observe the same thing. X-ray mammography and photoacoustics both claim to be able to study the same object, in our case breast tumors. But one of them observes how many x-rays are not absorbed by the tissue, while the other uses ultrasonic pulses created through heating by laser light as a measurement. How do we know if the two allow us access to the same object from a different perspective, or are completely unrelated? Is their spatial position enough to correlate them? Even if this position is wholly dependent upon how the data from the instruments is interpreted?

in his book "Instrumental Realism" [7] Ihde shows how various authors from different fields have arrived at a similar conclusion about the question scientific realism, and he groups these people together as "instrumental realists", himself being one of them. What these people share, according to Ihde, is a belief in the embodied and instrumental nature of science. The idea is that science cannot be separated from the instruments involved in it, and thus the power of a scientific claim can depend on the instruments used in its creation.

To give an example, the first telescope was able to show its user objects that were, before its invention, invisible. These objects could only be seen through a telescope, and were invisible to the naked eye. Were these object truly there, but not normally visible, or were they merely instrumental artifacts created by the telescope? Amongst these objects were Galileo's Medician stars, the moons of Jupiter. The existence of these objects were once the topic of discussion, today we accept their existence. So what has changed? After all, they are still invisible to the naked eye. What happened is that instruments became better, and the scientific frontier moved on. Today we have orbital cameras that take pictures of water channels on Mars, and probes flung into deep-space. All of these can confirm the existence of the Medician stars. Through a multitude of different technological measurements, their existence has seemingly been confirmed.

Now let us return to the topic of breast-cancer diagnosis. After screening, there is a single measurement, one point of data which has a multistable hermeneutic interpretation. The question, then, is whether or not greater scientific testing through a multitude of instruments and methods can confirm whether a tumor is real, or just an illusion, an artefact.

Instrumental realism, according to Ihde, looks at the praxis of technological perception within science, the way in which it is done, rather than merely its theories and texts. This allows for conclusions surrounding the question of realism that attempt to solve the issue.

Ihde argues that science is in many ways embodied and perceptual, and therefore reliant upon its instruments. However, these instruments also rely on scientific principles, which strengthens the reliability of those principles. He uses several other authors which he places in his camp of instrumental realists to strengthen the position, even if he does not agree with them entirely.

To solve the problem of mediated perception, that what a scientist sees is always mediated, Ihde has an example in Patrick Heelan, who solves the epistemic issue phenomenologically. Rather than trying to argue the objective truth of instrumental scientific instrumentation, Heelan argues that bodily observations are always already subjective, dependent upon their observer, who is both positioned and finite. As such, the notion of objective observation is impossible. From this, he argues that there is therefore no deep distinction between perceiving through the body, and perceiving through the body and a technology. Thus, if we accept the world we see directly as true, then we should also accept the world we see indirectly as true. The problem that remains, of course, is whether or not we accept that the world we see directly is true. For Heelan, the notion of embodiment becomes important. If the technology is embodied, if it is a direct part of the observer, there is no deep distinction. If the technology is too hermeneutic in its man-technology-world relation, a distinction is created, and we can no longer say that the two are equivalent.

Similarly, in order to look at the existence of scientific entities, he uses Ian Hacking to strengthen his position. Hacking's argument being that technological manipulation of a theoretical entity in an experiment is how the object gains its epistemic weight. When an invisible entity changes from being a theoretical explanation of a phenomenon, into something that can be manipulated and used in experiments, it gains weight. For example, take the difference between the theoretical DNA being discussed in the past, with theories being constructed about what its structure was, and our current knowledge of DNA, where we can read, analyse and even manipulate it. The entity is no longer theoretical in nature, even though we cannot observe it with the naked eye, but only through instruments.. An invisible entity too small to see with the eye alone may have little epistemic weight on its own, but if this entity can be manipulated and used in experiments, rather than merely being theorized, it is no longer entirely hypothetical in nature. If the object is being manipulated on a regular basis, it is strange to say that it is entirely theoretical.

As such, science becomes based upon technological praxis, and the technology unveils the scientific reality. In "Expanding Hermeneutics" he goes further into the topic of science as a visual, phenomenological, hermeneutic activity. His own argument about how scientific instruments allow access to reality comes forth from the other thinkers he studies, and is related to that as well.

Ihde's solution to the problem is based upon what he calls "instrumental phenomenological variations.", or simply variations for short. The idea behind this is that, through variation of measurement, the object itself can speak. This solves two problems, the question of whether or not the observed object exists, or is an artefact created by the technology, and the question of what the technology truly is.

Having referred to others to empower his idea of the validity of observation, Ihde does not automatically accept the existence of what is seen using an instrument. Similar to Hacking, he holds that if drastically different techniques come to the same conclusion, it is unlikely that the conclusion is simply an artefact. Thus, if only one instrument sees an object, it could be a mistake. If many different objects see it, it is rather unlikely. For example, the moons of Jupiter were once thought of as an artefact of the telescope, when they were first seen. Now, however, we have many different subfields that all agree on their existence. In diagnosis this is also applicable. One measurement could be a mistake, while if there are many different ones, the possibility that something is an artefact goes down. Interesting too, is that he applies this to the theory itself. By taking different philosophers from different traditions, and observing how they reach the same conclusion of Instrumental Realism, Ihde gives weight to his findings

However, Ihde also argues that these variations give a form of true access to the object itself. While any individual point of access to the object is mediated, the collective of observations allows a more unmediated access. To illustrate this, Ihde uses the metaphor of sound. Every object in our life has a sound, it can speak. However, an object on its own cannot make a sound. In order for this to happen, it must interact. A bell must be struck by a hammer to speak with its voice. However, in striking this bell, the hammer itself is also a participant. Striking a bell with a single hammer will only yield data about the bell being struck with that hammer. We do not know what part of the sound comes from the bell, and what comes from the hammer. It is only by using a plethora of different hammers, or other bell-striking implements, that we can find the true tone of the bell. To bring this metaphor to our medical topic, the bell is a tumor, while the different striking implements are photoacoustics measurements and X-ray imaging. Both technologies show something of themselves, as well as something of the tumor, in the creation of an image. But by including different types of technology, we get a better view of the bell. The technologies mediate the way in which we access the tumor, but by looking at the issue from different sides, we can look at the nature of the tumor itself. By looking at variations upon a theme, we can see the true nature of the object.

Thus, if we accept this notion of variations, the fact that the two technologies work different is a good thing, since the chance that the two are in agreement through coincidence and artefacts of the technology is lower. That which is in agreement between the two is truly there. Thus, our medical imaging technologies serve as variations that increase the epistemic weight of the tumor they are studying.

The problem that remains however, is that unlike the analogy using a bell and a hammer, these are not simply different hammers creating slightly different sounds, and we cannot directly observe that they are all interacting with the same object either. So what is it that can hold them together as being variations upon the same theme, rather than separate observations? One way to do this is with an appeal to theory, the hermeneutic framework used to explain the "text" provided by the images. A theory about what is inside the body that predicts what you would see from different perspective. The most obvious point of convergence here is the shape and location of the object. Both X-ray imaging and photoacoustics are special techniques. They attach a location to what they measure. Thus, while the technologies may create images that correspond to very different features of the tumor, they still do this in the same location.

To illustrate why something like locality is such a powerful tool, let us compare it with another type of test, for example, something that tests for a biomarker in a bodily fluid. To simplify: A biomarker is a particle that is associated with a disease, or set of diseases. If the disease is present in the body, the biomarker will also be there. A well-known example is the pregnancy test. What is tested is whether or not the condition of pregnancy is in the body, not the location of the pregnancy. In the same vein, a diagnostic test can tell us that a tumor may be present, but it does not directly tell us where that tumor is located. In this case, the claim that the test and a medical image correspond to the same object is far weaker.

What then, if something shows up in one test, but not in another, and is most definitely a part of the object? Theory seems to be needed. A theory about the disease, and what measurements in one technique mean for another, can help connect the tests. This, of course, is where the divide between science and medicine becomes clear, and where analogy becomes more difficult. Science attempts to find these theories that connect different measurements, whereas medicine uses them to optimally treat a patient. If a correlation between oxygenation levels and tumor staging through x-ray imaging is found through scientific research, then we assume that this theory holds for other tumors, and that the hypothetical object we are studying corresponds to the theory. About the general class of such objects.

As such, Ihde places scientific praxis, and holds it to be embodied in the technology. Science is dependent upon its instruments. Those instruments are varied, and often embodied in their use. He brings forward examples of astronomers feeling heat from light when looking at the sun through filters. The instruments are embodied in their use, allowing scientists access to the world through them.

4.3: Critique of Instrumental Realism

In his talk of instrumental realism, Ihde talks about the role of scientific instruments in the process of science, and about the embodied nature of science that goes beyond the purely theoretical in nature. However, human-technology-world relationships also include a human factor. The instruments all create multistable objects to be interpreted, rather than objective information. The technology mediates the way in which the human sees the world, but does not decide it. In addition to that, the person using the technology also influences the technology in turn.

Technology influences man, and man in turn interprets technology different and influences it, on all levels. To illustrate this, I turn to both Thomas Kuhn, for the macro level, and the work of Annemarie Mol, for a more micro level.

4.3.1: Kuhn

Thomas Kuhn is probably most famous for his influential “The structure of Scientific Revolution”[30], published in 1962. In this book, Kuhn explains and provides evidence for his theory of scientific paradigms. A paradigm is, according to Kuhn, the way in which a scientific discipline is structured, the basic *rules of the game* so to speak. Scientists usually work within their paradigm in order to solve problems they find. To give examples of paradigm, think of Newtonian or Einsteinian physics. Both of these are different paradigms, working by different rules. Before Einstein, everyone worked under Newtonian rules, and did Newtonian physics. For the most part, this went very well, Newtonian physics described the world, and within its framework solutions to problems could be found, at least, in regular science. Sometimes irregularities were found. According to Kuhn, these irregularities are then ignored since they do not fit into the current scientific paradigm. Only when there are simply too many irregularities to ignore, and there is a new, competing theory/paradigm that better explains all the known facts, a “scientific revolution” occurs. Slowly but steadily, the scientific population in the field (usually starting with younger scientists) starts to convert to this new paradigm. For normal science, a paradigm is the presupposition, the framework upon which science is build. Microbiology for example, relies on there to be cells in order to test its theories. To bring it closer to medical imaging, in order to do research into a disease, one first needs a theory of disease. Thus, research presupposes a paradigm under which the research is done. Research into cancer, or HIV, or almost any other disease, would be drastically different if medical researchers still worked with the idea that disease is caused by an imbalance between the humours, and that some bloodletting should get rid of the problem. Kuhn calls these paradigms a way of structured seeing, In the terminology I introduced in chapter 3, this type of perception falls under the umbrella of meso-perception. A learned pattern that influences the way in which the world is perceived, but one that can change. Two scientists looking at the same object under different paradigms can see different things; one can see oxygen involved in a chemical process, where the other sees the tell-tale signs of phlogiston being released during combustion.

Medical diagnosis works under a medical paradigm. All diagnosis is based upon this prerequisite paradigm. One of the things that Ihde touches upon in 'instrumental realism' is that Kuhn noticed that changes in paradigms, scientific revolutions, often coincide with changes in the scientific instrumentation. This means that changes in instrumentation used for research can influence the paradigms under which science does its research. It is the telescope that changes scientific paradigms surrounding the heavens, rather than the change in paradigm introducing the telescope.

This clearly illustrates the effect that technology can have on science. And how different technologies can have an effect beyond the location where they're used. One of the things that Kuhn notes is that there are often small errors in a theory before a revolution occurs. However, because these are minor, and because there is no opposing theory that does explain these things, they are largely ignored. However, if there is a major change in used instrumentation, it would be foolish not to be wary of small errors suddenly becoming much harder to ignore. Thus, the paradigm itself may be influenced by the technology used in it, and a change in instrumentation changes the way the disease is observed.

4.3.2: Mol

To look more at a micro level, I turn to Annemarie Mol's "The Body Multiple: ontology in medical practice"[26]. In this book, she talks about her studies of the disease: "atherosclerosis of the lower limbs". She shows how one disease can take on many different forms. For example, there is atherosclerosis "under a microscope". Under a microscope, when looked at by a pathologist, atherosclerosis of the lower limbs is "a thickening of the intima". It is characterized by the changes in the blood vessel. However, when talking to a patient in the clinic, a physician characterizes it by a lack of pulsation in the lower legs. To the patient, the disease takes the form of pain as well as difficulties when walking. As such, we can see the relevance of Ihde's notion of variations, each perspective gives us a clear view of what artherosclerosis truly is. Different technologies mediate between the observer and the disease in different ways.

What is interesting in her book, is her description of the procedures of angiography and duplex scanning. Both of these are techniques used to observe the severity of the disease, and they do so in two different ways.

Angiography is an older technology that uses a dye that shows up on x-ray images, and makes images that show its spread. This allows the doctor to look at the degree of lumen loss, how much of an artery has been lost to the disease. The technology is, however, seen as invasive. A dye needs to be injected, and several x-rays must be taken. It is not as invasive as a biopsy, but not harmless either.

Duplex scanning is a newer technology that is replacing angiography, and uses ultrasound and the Doppler effect to look at the movement of blood through the arteries. Its velocity is measured, and expressed in the peak systolic velocity, the highest velocity during the contraction of the heart. Since it only uses ultrasound, it is completely non-invasive.

Since funding and time is not infinite, it is not accepted practice to do two tests when only one is necessary. Simply doing both tests and trying to find the true atherosclerosis is not an option here, while in the scientific process, a newly discovered entity can easily be analysed through many different processes. It can't be done because costs too much money, takes too much time, is too taxing for the patient, or for other reasons. Thus, it seems that Duplex scanning takes over the role of angiography as one of the main forms of technological mediation between physician and atherosclerosis. But what happens with the terminology is very interesting. What doctors did was relate the Peak Systolic Velocity to the degree of lumen loss that would've been detected using angiography. Rather than using a different type of technological mediation on its own, it was used mostly as a replacement. Research was done and a theory was constructed to link PSV to lumen loss, and this relation was then used to diagnose a patient and set up a treatment plan, rather than using the duplex measurements directly. The expertise of many of the doctors was in the interpretation of angiographical results, thus, they translated the duplex scanners results before interpreting them.

As such, it seemed that the doctors did not agree with Ihde's notion of phenomenological variations. Their expertise was in using lumen loss for measurements, not PSV measurements. Then, that lumen loss was related to the disease. To link it back to the original metaphor, think of the creation of a mathematical equation capable of transforming the tone of one hammer striking the bell, into another. Rather than using the power of new instruments to increase the knowledge of the studied object, knowledge of the object is used to transform new instruments into old ones. The expertise that the physicians already had greatly influenced how a new technology was used, which could not be predicted using theories that led out this type of expertise. On the micro-level, expertise influenced the way the technology was used.

To further go through Mol's work with our postphenomenological terminology, we can say that there are different technologies and methods that mediate access to the disease. A pathologist uses his medical framework to interpret slices of flesh under a microscope. In the clinic, physical touch or other tools to measure things like blood pressure are the medium through which access to the disease is mediated. And to the patient, a method is used that could not possibly be more embodied, as the patient has access through his own nervous system.

But this way of thinking still seems to presuppose that there is something which is being looked at through different technologies. What Mol tries to figure out, is what happens when they do not correspond. What if the patient has complaints about difficulties walking, when the pathologist cannot find anything, even though the doctors measurements of systolic pressure agree with the patient? In Mol's understanding, what happens when measurements clash, is that people seek an explanation. Does the patient proclaim pain, while there is no loss of systolic pressure? Then maybe there is a different disease, or the patient is incorrect. What is interesting is that, from the viewpoint of the patient, it may be more likely that the doctors measurements are incorrect. In the end, though, conflicting measurements must be solved, and the doctors start looking for explanations. Perhaps the pain in the legs is caused by some other disease, or perhaps the patient's blood vessels are so calcified that normal blood pressure measurements are impossible. If the results no longer match up, in other words, if the different technological variations no longer show us different facets of the same object, a closer look is taken at the processes involved. Perhaps the specifics of a test mean that they work incorrectly on a patient, perhaps the patient has a different disease that interferes with the expected results. The technology can longer be seen as a simple provider of data, but must be inspected, they must take a look at what it does, and how they can explain the divergent results.

The initial assumption is that all the ways of looking are at the same object. Only when this does not hold up, when the different variations do not work together to create an ever better view of the patient's condition, does it become an option to drop this assumption. Is the patient complaining about something which is similar, but different? Were some of the tests performed incorrectly? Was the equipment faulty? The assumption of Ihde's instrumental realist, that different perspectives look at an object in different ways, and therefore, combining them gives us information about the thing itself, does not seem to hold up. One cannot naively assume that the object to which access is mediated is always the same object. There is a human factor at play, in judging what is happening if the measurements do not match with each other. Especially in the medical field, where a lack of conclusions is not simple a scientific curiosity, but a matter of life and death.

4.4: Epistemological responsibility.

Thus, where Ihde's theory deals with the notion of mediation, and how to solve the problem of access to the world always being mediated by the technological instruments used to observe it, he leaves out the issue of appropriation. Of the way in which the data created by the instruments is interpreted, and made one's own. This process of interpretation is highly dependent upon the person doing the interpretation, his skills and expertise, on both the macro and micro levels. As I have argued in chapter two, and given evidence for in paragraph 3.7, the very way someone relates to a technology (hermeneutic or embodied) is dependent on their expertise in using that technology. More than that, not only does expertise and the paradigm under which someone works influence the way he interprets things such as multistable medical images, his expertise and the paradigm he works under are in turn also influenced by the instrumentation he is using.

Thus, our original problem remains, and is joined by another. We asked what happens when we have two competing technologies that say different things about the world. Ihde then, claims that th differences go away when we look through enough different variations, leading us to what is true, but as we have seen, this does not take into account the individual doing the interpretation, who influences the process on both a macro and micro level. So how do we handle this problem? Who is right? The answer, in this case, lies in medical pragmatism. It is not necessary to know which technology is correct or which interpretation is the right one, it is necessary to treat the patient. To handle this then, we need to look beyond the individual case, and look on a higher level, both in science in general, and the medical field specifically. I will do this by first introducing the subject through what I call second order interpretation, and then bring up the notion of Epistemological responsibility as proposed by van Baalen and Boon[31].

Both medicine and science are more than simply the interpretation of observations. A good example, and one that Ihde brings up, is Galileo's telescope. The telescope mediated Galileo's access to the stars above, the technology amplified the power of his vision, but at the same time added limitations based upon the instrumentation. Not just the limitations we see in the modern day, (Seeing only a small, circular field.) but also artefacts created by technological deficiency of the early telescope. Ihde explains how Galileo proclaimed that anyone could, with the use of his new instrument the telescope, see what the ancient astronomers could not. Given, of course, that Galileo teaches them how to use the instrument. As such, one can see that even one of the very first examples of scientific instrumentation required a skilled user and interpreter. In order to see a planet through a telescope, one needs to be able to both use a telescope, and to interpret the image as being that of a planet. In much the same way, in order to diagnose breast cancer, one first needs the expertise to create an image, and then the expertise to interpret what the image means.

But Galileo did more than just watch planets and identifying how they moved. He also interpreted what those movements meant, and constructed theories around these movements. There was a second order interpretation. The instrumentation supplied the hermeneutic text, but Galileo went beyond that, towards theories about the solar system. He did not limit himself to the observations themselves, but went beyond them. Thus, the role of the scientist. In much the same vein, doctors do not limit themselves to identifying a disease. Their diagnosis is a theory about the body of the patient, and they theorize on what to do to solve the problems of the patient.

In a 2015 article [31], van Baalen and Boon discuss the topic by introducing the concept of epistemological responsibility. With this concept, it is the professional responsibility of a medical doctor to create a mental “image” of a patient and the disease, using diagnostic tests, scientific-medical knowledge, and contextual information about the patient and the surrounding medical system. This picture then functions as an epistemic tool that can be used for further reasoning about the patient’s condition, such as predictions for treatments and the formation of hypothesis about the patient’s disease.

In their article, van Baalen and Boon set this viewpoint of the epistemological responsibility of the doctor for the creation of this epistemic tool in opposition to the practice of Evidence-Based medicine. Evidence-Based Medicine being the modern way of thinking within medicine that all decision-making should, ultimately, be based upon medical research, in the form of Random Controlled Trials. Thus, is part of the argument, the doctor’s role is reduced to one of a flow-chart reader, doing exactly what the research tells him is optimal, without leaving room for the specific context of the patient, or the doctors own expertise. The focus with epistemological responsibility then, is not rule based reasoning, in the form of Picture X therefore disease Y therefore treatment Z, with the relations between these found in scientific literature based upon double-blind studies, but instead looking at how to help a specific individual in his or her own context.

The expertise of the physician here is not simply the application of scientific knowledge in a rule-based manner, nor is it limited to the interpretation of medical images. Instead, it is found in the creation of, and working with, the epistemic tool, the image of the patient. To do this, tacit knowledge about the art of being a doctor is necessary. Simple application of explicit rules is not enough.

The question then, is what kind of tacit knowledge is necessary here. The skill of interpreting medical imaging, of being able to properly handle the information, and even the skills of handling medical equipment and performing procedures are all Dreyfusian in nature. Skills where explicit rules are slowly turned into tacit knowledge on how to properly perform the actions.

When we look at a higher level, on decision-making within the context of the individual patient, the medical system, the hospital and other such factors, then we are looking on the level where Collins places tacit knowledge and expertise. To couple this back to the notion of variations, interpreting the different images created by different imaging technologies would fall squarely in the field of Dreyfusian expertise, while Collins' notion has more to do with combining different observations to find the core of the issue, and deciding what to do with that information.

What then, about Heideggerian taciticity. If Being is necessary to truly understand a patient's condition, it is no wonder that the worldview of Evidence-Based-Medicine is not enough to treat a patient. It is the physician that has the knowledge to understand the existence of the patient, through his Being-in-the World.

Thus, when we work through the concept of epistemological responsibility, a tumor is no longer a clump if cells. Instead, it is something that is a part of a patient. What a diagnosis using a medical image does, is add that information to the epistemic tool for thinking about the patient's situation. But when a medical image is interpreted as being one that corresponds to a body with a tumor, what that means depends on the patient.

A good example of the context dependency of disease can be found in seniors. If a 90-year old person is diagnosed with a form of cancer that will be deadly in ten or twenty years, when the person already has a different ailment that is likely to be fatal in half the time, then it is foolish to start an extensive regime of chemotherapy. The disease is not a disease at all, but simply something that is happening in the body.

Thus, two equal tumors, can be very different depending on the patients and the surrounding context. What is dangerous to a poor person living in a low-resource country may be solvable with a routine operation in a country with high-quality healthcare. What it is is the same thing, but what it means is something very different. The disease is more than simply the state of the body.

4.5: Conclusion

For the third time, we return to the questions posed at the start of this thesis, that have been discussed in this chapter.

I have discussed what happened when the technologies do not converge, using the work of Mol. The doctors start looking at the specifics, under the hood of the technology, to see which measurement they can discard, and which they must keep. Is it the patient's experience that is wrong? One of their measurement devices? In the same sense, when we move back to medical imaging, we can bring up the epistemological responsibility of the physician here. In the case of non-converging results, what happens is that the doctor must make a decision, must decide which test he believes in order to help his patient, or what extra information to gather to come to a decision.

The other issue is that of the existence of the disease itself. Don Ihde's Instrumental Realism gave us an in, a way of accessing the essence of something through variations upon it, keeping what remained through the different variations. Yet this only accounted for the technology, and not the perceptual frameworks and expertise of the person using the technology, and it seems unlikely if not impossible that we can have variations of these, since being an expert in the topic, and therefore able to correctly interpret, comes with a meso-perceptual framework. To vary upon this would be to ask the layman, who would have no idea how to interpret the image, how the medical image relates to the patient. After all, had he been capable of this, he would not have been a layman.

My answer here then, is that it seems very likely that there is indeed an essence, a disease that remains constant, but that this neutral essence is unattainable. Too many things colour the observations and interfere with each other to somehow observe it. But the fact that we can look at it means that we can say things about it, and performs actions such as a diagnosis.

The physician then, takes upon himself an epistemological responsibility, and creates the best picture of the patient he can, which he uses to treat that patient within his or her context. This physician has to be an expert, both in the technologies he uses, as well as the field of medicine. Even in the case of quantified, automatic measurements, he needs to understand what is happening. After all, if two technologies do not converge, the physician must be able to analyse why, and know which result to discard. If he only understands the result, and not the how, he cannot do this.

Chapter 5: Conclusion.

5.1: Answering questions

At the start of this thesis, I asked the question of what happens when two technologies, observing the same object, show different things. Specifically, I asked this about medical imaging techniques in the context of breast cancer screening. What is it that happens, if one technology says one thing, and another technology says something else, the two things being mutually exclusive? How do you decide which technology to trust? How do you find out what the truth is?

The first step in answering this question, was to look at the technology itself. For technology does not simply say something, and medical imaging technologies do not make conclusions. Instead, what they do is mediate between their user and the world. They magnify some aspects, while leaving away others, and influence what can be seen, and how it is seen. The technologies of photoacoustic imaging and x-ray mammography do the same things, but in a different way, by looking at a different feature of the tumor. They do not show a perfectly neutral picture of a tumor, but create an image of it in their own way. The technology is not a photograph that can be used to easily see what is inside a person, but had important mediating tendencies. The way in which they mediate is different, as is the way in which people use the technology, the human-technology-world relationship. In addition to that, there is the issue of multistability, and the interpretation of the medical images. The meaning of the images created by imaging technologies is not immediately clear. There are multiple, often equally valid, ways to interpret their meaning. As such, the person interpreting the images has an important role.

This then, makes the question we asked much more complicated. For when we as what happens when two different technologies say different things about the same object, we must now look with knowledge of both the different mediating effects of the technologies, and the way in which they are interpreted. If we have two medical imaging technologies, one of which claims that a patient is ill, while the other claims the patient is in perfect health, we must realize that what actually happens is that the doctor has interpreted the (mediated) medical image of the patient as being that of an ill person when using one technology, while interpreting the medical image created by the other technology as that of a healthy person.

To properly handle the question of interpretation, we turned towards the question of expertise. For if it is the expert opinion that makes a certain interpretation the correct one, we must know what that expertise is. Two different models of expertise were discussed, both looking at it from different perspectives. Collins focussed on expertise as belonging to a group of experts, as well as the nature of tacit knowledge. Most of his concept of expertise is broad, a field of study or a profession with a shared language. In contrast to this lies the Dreyfusian model, which is more focussed on the activity itself, becoming an expert in a thing rather than in a field. Expertise is, in many ways, a skill. Important to both however, is the notion of taciticity, tacit knowledge, which is knowledge that cannot be made explicit. Something that can, in theory, be problematic, as it is not something that can be easily explained. Because of the tacit nature of at least some of his knowledge and skills, the expert cannot perfectly explain why a medical image must be interpreted in a certain way.

What is important here, is that expertise leads to the expert interpreting in different ways. An expert is not just more effective, but he is quantitatively different than someone with less expertise, and the expert uses both highly tacit skills and explicit knowledge. What is important then, is that expertise changes the way we view the world. An expert perceives the world in a different way, he has a different mesoperceptual framework of the world, though which his microperceptions are filtered.

This then, tells us once again how important the observer, the interpreter, is in our question about medical imaging technologies. It is not that the two images are interpreted in different ways, but that they are interpreted that way by an observer with a specific mesoperceptual framework. Yet again, this does not give us an answer to the question of what is happening when two different technologies disagree.

The final move then, was a step towards the field of epistemology, and the things we hold to be true. Here, we seemed to have a partial answer in Don Ihde's Instrumental Realism. Instrumental Realism holds that when we have different technological variations, that which holds between all of them can be taken for truth, for if they were only artefacts of the technology, then they would not show up when looking through different technologies. This also answers our question somewhat. If two technologies show different things, then we cannot accept either as true, since what gives it epistemic weight is the way in which the phenomenon stays between different technologies.

Yet when we look at Instrumental Realism, we see that what misses is the topic of interpretation, the way in which the expertise, the mesoperceptual framework, influences the way in which the data provided by technology is interpreted, both on the micro and macro level. The solution then, is to move further into the realm of the expert, and looking at the pragmatic nature of the field of medicine in itself. By bringing in the concept of epistemological responsibility, and noting that what is important that the patient is helped, rather than that one of the two technologies is supreme. Unlike science, it does not matter whether or not we have true knowledge of a tumor, what matters is that the patient is helped. Thus, when we return to our question: 'What happens when two different medical imaging techniques disagree about the state of the patient?' we find our solution in the fact that the expert, who holds epistemological responsibility for determining the state of the patient, makes a decision, based upon his skills and knowledge about the patient, the context surround the patient and the hospital, and the technologies involved in creation of the medical images. The responsible expert may not be able to have true, unmediated knowledge of the state of the patient, but by reflecting upon the ways in which his diagnosis is created, he is still able to properly diagnose the patient, and create a treatment plan. If this is not possible, because the different interpretations of data available to the expert do not converge towards a single coherent disease, then it is up to the physician to make them converge, by reinterpreting, looking for possible sources of error, or even discounting one of the results. What is important is not truth, but rather helping the patient. If the interpretation of one type of medical image says the patient is sick, while an interpretation of another type of medical image says that the patient is healthy, then the doctor may decide to go with the interpretation that puts the patient at the least risk when it turns out to be incorrect, instead of going through an analysis of both technologies, and trying to make the results match up that way. The important question is not how two technologies can have different results, but rather the question of how to handle that situation. The question is how to take epistemological responsibility. In much the same way, the question of the disease itself, the essence not influenced by the mesoperceptual framework of the interpreter or the mediating effect of the technology, is irrelevant, as what is necessary is knowledge about how to help the patient, not true knowledge about the essence of the tumor.

To take his epistemological responsibility, the expert takes into account many things that have been talked about into thesis. The way in which technology mediates access to the world, and the user-dependency of the human-technology relationship, the way in which his own mesoperceptual framework influences his perceptions and interpretations, and the way in which different technological variations can give more insight into the essence of a tumor. Beyond this, other factors such as the situation of the patient also play a role, as does his knowledge of the disease, current treatment methods and the hospital itself. In the end, what is important here is reflection, the ability to understood how things like technological mediation and mesoperceptual frameworks influence access to the tumor.

5.2: Future perspectives

In the second chapter, I introduced the notion of the quantified measurement. Something that, rather than being interpreted, simply created a number that meant something. A result that seemingly does not require expertise in its interpretation. It is, in many ways, similar to the way pregnancy tests work, being capable of showing a result in a simple binary yes/no answer.

Thus, we seemingly have a technology that does not require any expertise to use, at least if we look at it naively.

This is because the expertise, and the process of interpretation of a multistable result, is instead done during the development of the technology. The decision of whether or not to interpret a specific result as one option or the other is still there. The entire process of meso-perception influencing what is seen is simply moved to the development of the device, and coded into the technology.

In addition to that, of course, there is the second level of interpretation, that of deciding what to do with the diagnosis. After all, to return to the pregnancy test, simply knowing whether or not you are pregnant is simply the beginning of a larger process, something that leads into further testing and meetings with medical experts.

A possible problem with this, is that the level of interpretation, and the mediating effect of the technology, is hidden. As we have seen, it is often vital that the physician is able to understand what a technology does, not just in order to understand how it mediates between him and the world, but also in order to be able to solve the problems of non-converging variations.

However, this change also has it positive sides. The expertise required to read the technology is no longer something that is required of every user of the technology. This allows for more efficiency, especially in fields like medicine, where the time and expertise of experts comes at a financial cost. The problem however, is that the physician cannot gain expertise in the technology without using it. Even if he reads through the documentation to understand what is under the hood of the technology, he only gains explicit knowledge about the technology, and not the tacit knowledge he would gain by using the technology. The solution to this problem may be found in division of labour, with several physicians combining their different expertises, able to properly communicate because they all hold the tacit knowledge of their field.

In the end, then, these are pros and cons that must be measured against each other, and given the nature of medicine, I believe that this equation will be found to be in favour of automatisation and quantification. Using such technology means that less time and expertise is required of a physician, allowing him or her to treat more patients. The other side of the coin, that physicians are no longer able to properly reflect upon how the technology they are using mediates between them and the world, and on what influences the interpretation of the results, may result in some problems, but then again, medicine will never be perfect. A great increase in the capacity of a medical system may be worth a slight loss in rigor, as it ultimately results in more lives saved.

One of the solutions then, is to return to the topic of screening. It can be argued that the goal of breast cancer screening is not actually to detect cancer. Rather, its purpose is to find people who are very likely to have cancer, that can then be referred to the medical system. In such a specific case, the individual context of the patient is not immediately important. Thus, while the use of quantifying or automatically diagnosing technologies may be problematic in some use cases, in others it holds little problem. After all, there is a big difference between something requiring a life-saving operation, and an advanced microchip that prescribes someone an over-the-counter remedy.

The danger, of course, is that such things would move into the theatre of science, where unlike in medicine, the focus is not on saving lives, but rather on finding knowledge. As long as the process of defining what a disease is, or creating new treatments, is not compromised by black-boxing technologies, it seems that we are partially safe. The problem being, of course, that the line between medical research, and medicine, is not a clear divide. Physicians do medical research, researchers work in hospitals. As such, the perceptual frameworks of one group influence those of the other.

In the end then, it comes back to the physician, and his ability to understand context. After all, if we demand the ability to self-reflect, we can assume that a good physician also has the knowledge to know when efficiency is more important than being able to reflect upon the influences of the technology you are using. Thus, a physician in an area that is understaffed, or sees a lot of patients with little resources, can choose to use these black-boxing technologies, knowing about the consequences, while a physician working with a smaller amount of patients, and a great amount of time to demand to each of them, may be able to afford the time and energy necessary to acquire expertise in the technologies that he is using.

References

1. Drukteinis, J.S., et al., *Beyond mammography: new frontiers in breast cancer screening.* Am J Med, 2013. **126**(6): p. 472-9.
2. Boenink, M., *Debating the desirability of new biomedical technologies: lessons from the introduction of breast cancer screening in the Netherlands.* Health Care Anal, 2012. **20**(1): p. 84-102.
3. Heijblom, M., et al., *The state of the art in breast imaging using the Twente Photoacoustic Mammoscope: results from 31 measurements on malignancies.* Eur Radiol, 2016. **26**(11): p. 3874-3887.
4. Kim, K.I., et al., *Changing patterns of microcalcification on screening mammography for prediction of breast cancer.* Breast Cancer, 2016. **23**(3): p. 471-8.
5. Verbeek, P.-P., *What things do : Philosophical reflections on technology, agency, and design*. 2005, University Park, Pa.: Pennsylvania State University Press. viii, 249 p.
6. Hussain, A., et al., *Quantitative blood oxygen saturation imaging using combined photoacoustics and acousto-optics.* Optics Letters, 2016. **41**(8): p. 1720-1723.
7. Ihde, D., *Instrumental realism : the interface between philosophy of science and philosophy of technology*. The Indiana series in the philosophy of technology. 1991, Bloomington: Indiana University Press. xiv, 159 p.
8. Roskies, A.L., *Are neuroimages like photographs of the brain?* Philosophy of Science, 2007. **74**(5): p. 860-872.
9. Mackenzie, A., et al., *Breast cancer detection rates using four different types of mammography detectors.* European Radiology, 2016. **26**(3): p. 874-883.
10. Xia, W., W. Steenbergen, and S. Manohar, *Photoacoustic mammography: prospects and promises.* Breast Cancer, 2014. **3**(5): p. 387-390.
11. Heijblom, M., et al., *Visualizing breast cancer using the Twente photoacoustic mammoscope: what do we learn from twelve new patient measurements?* Opt Express, 2012. **20**(11): p. 11582-97.
12. Zhang, H.F., et al., *Functional photoacoustic microscopy for high-resolution and noninvasive in vivo imaging.* Nat Biotech, 2006. **24**(7): p. 848-851.
13. Daoudi, K., et al., *Correcting photoacoustic signals for fluence variations using acousto-optic modulation.* Optics Express, 2012. **20**(13): p. 14117-14129.
14. Ihde, D., *Expanding hermeneutics : visualism in science*. Northwestern University studies in phenomenology and existential philosophy. 1998, Evanston, Ill.: Northwestern University Press. xii, 216 p.
15. Rosenberger, R., *Perceiving Other Planets: Bodily Experience, Interpretation, and the Mars Orbiter Camera.* Human Studies, 2008. **31**(1): p. 63-75.
16. Bastian, L.A., et al., *Diagnostic efficiency of home pregnancy test kits. A meta-analysis.* Arch Fam Med, 1998. **7**(5): p. 465-9.
17. van Dijck, J., *het transparante lichaam medische visualisering in media en cultuur*. 2001, Amsterdam: Amsterdam University Press.
18. COLLINS HARRY, A.E.R. *`PERIODIC TABLE OF EXPERTISES'*. 2004; Available from: WWW.CF.AC.UK/SOCSI/EXPERT.
19. Epstein, S., *Impure Science: AIDS, Activism, and the Politics of Knowledge*. 1996: University of California Press.
20. Collins, H., *Bicycling on the Moon: Collective Tacit Knowledge and Somatic-limit Tacit Knowledge.* Organization Studies, 2007. **28**(2): p. 257-262.
21. Collins, H., *Three dimensions of expertise.* Phenomenology and the Cognitive Sciences, 2013. **12**(2): p. 253-273.

22. Dreyfus, S.E. and H.L. Dreyfus, *A five-stage model of the mental activities involved in directed skill acquisition*, 1980, DTIC Document.
23. Hutto, D.D. and R. Sánchez-García, *Choking RECtified: embodied expertise beyond Dreyfus.* Phenomenology and the Cognitive Sciences, 2015. **14**(2): p. 309-331.
24. Trubody, B., *When Tacit is Not Tacit Enough: A Heideggerian Critique of Collins' "tacit" Knowledge.* Meta: Research in Hermeneutics, Phenomenology, and Practicap Philosophy, 2013(Vol. V, No. 2): p. 315-335.
25. Selinger, E.M., *Feyerabend's democratic critique of expertise.* Critical Review, 2003. **15**(3-4): p. 359-373.
26. Mol, A., *The Body Multiple: Ontology in Medical Practice*. 2002: Duke University Press.
27. Krupinski, E.A., *Visual scanning patterns of radiologists searching mammograms.* Acad Radiol, 1996. **3**(2): p. 137-44.
28. Beutel, J., H.L. Kundel, and R.L. Van Metter, *Handbook of Medical Imaging: Physics and psychophysics*. 2000: SPIE Press.
29. Kundel, H.L., C.F. Nodine, and E.A. Krupinski, *Searching for lung nodules. Visual dwell indicates locations of false-positive and false-negative decisions.* Invest Radiol, 1989. **24**(6): p. 472-8.
30. Kuhn, T.S., *The structure of scientific revolutions*. 2012: University of Chicago press.
31. van Baalen, S. and M. Boon, *An epistemological shift: from evidence-based medicine to epistemological responsibility.* J Eval Clin Pract, 2015. **21**(3): p. 433-9.