

INTELLIGENT CLINICAL DECISION SUPPORT AND CURRENT MEDICAL PRACTICE

An Exploration of Compatibility and Epistemological Issues

Jan-Yme de Boer (s1013815)

*Philosophy of Science, Technology and Society
University of Twente
2016-2017*

*Supervisor: Dr. M.A.J. MacLeod
Second Reader: Dr. M. Boenink*

Acknowledgements

I have actually no idea if people write acknowledgements for theses anymore. Because of not taking the time to do some empirical research on the theses web page, I'll run the risk of doing an odd thing. That is a risk, however, that I gladly accept as doing an odd thing appears completely unrestraining to me, as I feel compelled to address some of the gratitude I experience in the context of this thesis.

I consider myself rather lucky in the sense that few of what I achieve, if they count as achievements at all, I could have achieved in solitude. This thesis is no exception to this and there are many people whose efforts to support me in delivering this thesis I could acknowledge. In fact—a bit depending on the degree of externality of contributing factors I would be willing to accept—I could probably fill an entire chapter with thank-yous.

As implied, it is not up to me to determine whether this is an achievement, that is a role reserved for the graduation committee. As a consequence, I cannot thank anyone yet for the achievement this thesis may or may not be in terms of quality. What I can establish myself however, is the fact that a certain sense of satisfaction has been achieved on my side. Satisfaction with the direct thesis-writing process that led to the delivery of this thesis, with the PSTS programme as a whole and with everything that the University of Twente and its people offered me before and besides. It is exactly this satisfaction for which I acknowledge the valuable contributions of all people involved.

However, to be a bit more specific in context of this thesis, I first of all thank my daily supervisor Miles McLeod. Miles managed to make and keep me motivated and committed from the start by example, and managed to continue this even after my explicit announcement to have been forced to initially aim for the bare minimum. This is something I believe one cannot reasonably expect from teachers, but of which I am confident enabled me to keep the project focussed and finished in time. Something that came to mean a lot for me.

I similarly thank Marianne Boenink for accepting a position in the examination committee and being flexible with regard to my schedule from the start. The feedback I received on the previous version of this thesis impeccably revealed the thesis' analytical weaknesses, at least anecdotally proving the quality of PSTS's teachers from which I had the opportunity to learn—for which gratitude would be in place—and helping me make the last improvements.

Thanks are furthermore in place for Melle, Jeroen, Joke, Marijke and Rianne, who helped me shape my research, and with whom I experienced Masterlab's countless hours as rather enjoyable.

Slightly more external to this thesis, I thank Iris for starting PSTS in 2014 and studying together for little over three quarters of a year, and all that followed. Without it PSTS would have been significantly less enjoyable. I am certain it has had a profound impact on my life as it continues to present itself, for the better I am sure.

Finally I thank anyone taking the time and the effort to read this thesis; although I must admit that I am quite certain this means that I am merely thanking the graduation committee twice.

TABLE OF CONTENTS

<i>INTRODUCTION</i>	<i>7</i>
---------------------	----------

<i>CHAPTER I</i> <i>Clinical Decision-Making & Knowledge Production</i>	<i>12</i>
--	-----------

<i>CHAPTER II</i> <i>Intelligent Clinical Decision Support System</i>	<i>37</i>
--	-----------

<i>CHAPTER III</i> <i>A brief reflection on methods of intelligence and expectations</i>	<i>57</i>
---	-----------

<i>CHAPTER IV</i> <i>Analysis of Differences and Issues</i>	<i>66</i>
--	-----------

<i>DISCUSSION & CONCLUSION</i>	<i>79</i>
--	-----------

INTRODUCTION

Having a background in biomedical engineering I have had a fascination for the intersection of medical practices and technological developments for quite a while. The fluid and soft organic context of the human body are starkly contrasted by the precise calculations and solid materials of technology. But this opposition is merely paradoxical, as the integration of technology and medicine has become ubiquitous in contemporary medicine, and has so often been shown to provide revolutionary improvements of the processes and outcomes of clinical practice. But what for now has remained the same is that medical practice remains a process in which human actors are faced with a great amount of choices about diagnostic and therapeutic options, choices they ultimately make themselves.

However, whether the role of human decision-makers will continue to remain a constant factor in the clinical decision-making process is possibly uncertain. The technological developments that were briefly mentioned have already been permeating into the decision-making process, for example by providing new knowledge and creating new options to choose from. This thesis project aims to investigate the direct involvement of technology into decision-making and knowledge production itself, quite possibly a future step in the hybrid between the medical and the technological. In more specific terms this thesis project will aim to come to an understanding of the issues that relate to the compatibility of the current modes of clinical decision-making and knowledge productions with artificially intelligent decision support systems.

Optimising Decisions

If we can agree that medicine is so heavily dependent of decision-making, it can go without question that the goal of medicine is to make these choices as good as possible. How good decisions should be made in medical practice is not self-evident. It is not self-evident what criteria should constitute *the good*, nor is it self-evident how decisions need to be made to meet these criteria. However, whatever good is precisely supposed to entail, it is generally true that the outcomes of decisions are supposed to provide the foundation for accurate, effective and reliable clinical results. Medicine as a general discipline has a long standing tradition of providing ways in which clinicians are enabled in this task. Perhaps in its most fundamental form by training practitioners in attitude, skill and knowledge, but also in more elaborate and external assistance to diagnostic and therapeutic decision-making; such as guidelines or diagnostic manuals. The process of doing medicine will later in this thesis be shown to be rather diverse in its methods of diagnosis and decision-making, but one particularly influential method of current medical practice has been, and continues to be, a principle named Evidence Based Medicine (EBM). EBM was announced as a "new paradigm" for teaching and practising medicine by the Evidence Based Medicine Working Group in the early nineties (Evidence-Based Medicine Working Group, 1992). EBM requires that clinical decisions are made as rationally as the most recent and relevant evidence from clinically relevant research allows (Sackett, Rosenberg, Gray, Haynes, & Richardson, 1996).

The ambitions of using EBM are easily understood as a normative ideal for what good decisions should be based on and to enable clinicians to make these decisions. However its implementation has, for reasons this thesis will come to return to later on, not been as straightforward as some may have hoped (Greenhalgh, Howick, & Maskrey, 2014), with actual application in clinical practice failing to really catch on in all aspects of clinical practice (El-Khatib, Zeineldine, Ayoub, Husari, & Bou-Khalil, 2010; Gurses et al., 2010; Shafi et al., 2014).

Though EBM was introduced as a solution to sub-optimal and varying treatment standards and outcomes throughout the medical process, it is apparently not free from shortcomings itself. This, and the fact that EBM is partially being replaced by new methods, such as narrative medicine, shows how this process of decision-making continuous to be a developing field with room for improvement.

This room for improvement is where technology, in this case information technology, comes in and aims to provide solutions for these shortcomings. The issue here is that researchers on the intersection of technology development and clinical practice are coming up with ways to contribute to, and integrate into, the decision-making process. One discipline of research and innovation that seems highly ambitious, promising and to attract a lot of attention is computer science and their implementation of artificially intelligent information technologies into the clinical decision-making processes. This, among other things, is supposed to potentially enable the integration of automatic generation of diagnostic and therapeutic decision-suggestions; based on knowledge autonomously derived from real-time monitoring of patients as well as from readily available knowledge from previous scientific research. This can be understood to be an evolution of existing *clinical decision support systems* (CDSS). CDSS has existed for quite a few decennia, generally in the form of a group of pre-programmed software tools that clinicians have the ability to consult when they are in doubt (Musen, Shahar, & Shortliffe, 2006). The evolution of CDSS as it will be characterised in this thesis is not a passive tool, nor are its diagnostic and therapeutic options necessarily pre-programmed. It may actively be taking part in clinical processes and may draw conclusions of optimal treatment options autonomously, and on the basis of large collections of existing data and knowledge; rather than on pre-programmed guidelines. In other words, this form of CDSS may be considered to be intelligent in the same way as for example self-driving cars, personal smart phone assistants and Youtube's recommendations are intelligent; arguably allowing for more accurate and more individualised treatment, allowing the application of knowledge independent of highly schooled specialists, and for a medicine that can go beyond the limits of the diagnostic and therapeutic categories it has become to rely on. We shall refer to this particular category of computer technology implementations as Intelligent Clinical Decision Support Systems (ICDSS). This provides an image of the future that is not just the starting point of this thesis proposal, but that is envisioned, advocated and initiated by others as well (Bennett & Hauser, 2013; Ghahramani, 2015; Hauskrecht & Fraser, 2000; Patel et al., 2009; Peek, Combi, Marin, & Bellazzi, 2015; Ramesh, Kambhampati, Monson, & Drew, 2004; Rowley, 2016). For the sceptical reader I will announce that, though many parts of this are still rather speculative, this thesis will convincingly argue that there are reasons to take these promises seriously.

Knowledge Production

Intertwined with these ideals for good clinical decisions and the decision-making process is the production of knowledge on the basis of which these decisions should be made. We can, even at this stage in the thesis, safely argue that any form of judgement is informed by some form of knowledge. Regardless of what decisions may be based on, knowledge is always part of the context in which a clinician evaluates information about the patient's status and especially the patient's options. In the most rudimentary forms of medicine this knowledge may be limited to unfounded beliefs about the human body; but in more meticulous and present day clinical methods, such as has already been briefly mentioned for EBM, it consists mainly and preferably of the outcomes of scientific research and extensive training.

There shall be ample opportunity in later chapters of this thesis to discuss the role and practice of knowledge production under current medical practice. What is important for now is to emphasise that ICDSS does not limit itself to changing the way in which the decision is made on an executive level. It is argued that its data-processing abilities provide the potential to change the way in which the necessary knowledge is produced on which decisions are based. Not through doing organised experiments such as randomised control trials, but through autonomously analysing large sets of already existing and real-time incoming data in order to construct knowledge on the basis of correlations in datasets.

Problem statement

With a bit of imagination, which suffices for now, ICDSS can easily be seen to be the announcement of a future practice in which clinical decisions will increasingly be based on the knowledge that is gathered, produced and used by increasingly autonomous and possibly increasingly opaque systems. This seemingly places ICDSS in a frame that implies a warning for the implicit risks of having to blindly trust machines, a frame that might inspire the assumption that this thesis aims to argue against an implementation of ICDSS. However, this thesis will assume the permeation of artificial intelligence into medical practice as inevitable, and assumes it taking the form of ICDSS as very likely (as will be supported later on). It is important to note that it is not the point of this thesis project to naively go along with the promises and expectations of entrepreneurs and visionaries, applaud their predictions and accept their images of the course of events and the future. Proposing the implementation of ICDSS into the current practice of making knowledge and decisions in medicine entails a great amount of complexities. Great differences may exist between the epistemic, technological and clinical cultures of current medicine and ICDSS; human activities, attitudes and beliefs may need to be rearranged and adjusted for ICDSS to properly function in medical organisations.

So, instead of solely focussing on an argument for or against implementations, this thesis mainly aims to come to an understanding of what medical decision making and knowledge production currently entail and what ICDSS may entail in these respects with a focus on epistemic and social processes; how different, mutually exclusive, or compatible they are, and how they may or may not assimilate into a new form of medical decision making and ultimately what may be required for that.

In more direct wording, this leads to the following main and sub-questions to be addressed in this thesis project:

What is the difference between current medical practice and Intelligent Clinical Decision Support Systems in terms of the way in which knowledge and guidelines are produced and clinical decisions are being made, and what are the issues with regard to their compatibility in case of integrating Intelligent Clinical Decision Support Systems into current medical practice.

In attempting to provide an answer to the main question, the first chapter explores what methods of knowledge production and decision-making are used in current medical practice, and—drawing from Miriam Solomon's work on medical epistemology—how current medicine can be characterised in general in terms of knowledge production and decision-making. It argues that a tidy methodological pluralism of Narrative Medicine, Evidence-Based Medicine, Consensus Conferences and Translational Medicine provides an accurate summary of current medical practice.

The second chapter will provide an overview of the promises and expectations that surround ICDSS and big data technologies in general as well as a methodological and functional overview of the technologies that underlie these promises and expectations in order to determine how ICDSS can be meaningfully characterised in terms of reasonably realistic clinically relevant capabilities. In doing so it argues collections of separate autonomously intelligent functionalities on the basis of big data technologies will be integrated with comprehensive patient monitoring systems and databases of written sources of medical knowledge into an ICDSS, allowing it—among other things—to assist diagnoses and therapy by providing suggestions on the basis of existing literature, generate new medical knowledge on the basis of correlations and patterns in available patient data and possibly provide predictive diagnoses as well.

Drawing from the first and the second chapter, the third chapter will aim to identify relevant (epistemic) issues that emerge from the hypothetical integration of ICDSS into current medical practice. This ultimately identifies ICDSS's implications on transparency, the differences in resolving scientific controversy between ICDSS and current clinical practice, ICDSS's need for standardising record keeping and the discontinuity ICDSS causes with regard to the role of scientific realism and causal-mechanistic reasoning as significant issues. On a methodological level this chapter identifies that the role of Narrative Medicine in current medical practice is problematic for ICDSS's capabilities, it is unclear how knowledge produced with ICDSS's can be integrated into EBM's evidential hierarchy, and that ICDSS's potential in autonomously executing translational processes of Translational Medicine as very limited.

Ultimately this thesis will conclude by structuring and reiterating these findings in more detail in order to identify themes and broader trends. Among which it shall discuss the role of trust and its implications on responsibility. Finally it shall reflect on the desirability of compatibility as a quality of ICDSS and current medical practice, concluding that though they can be discussed their applicability is limited.

CHAPTER I

Clinical Decision-Making & Knowledge Production

Everyday life of most people is intertwined with medicine to a significant extent. Many humans draw their first and their last breaths in hospitals or under medical supervision, and in the time between these key moments in their brief existence, they and the people around them may be vaccinated, screened, diagnosed, cured or be hearing disconcerting news from doctors. Though many people intuitively prefer not to require the services that medicine provides, medicine is rather ubiquitously present in most people's lives. Even if you are among the very few people whose lives have never been directly influenced by medicine, chances are slim that you have never seen one of the countless medically themed television programmes. From light-hearted comedy soap series to real-life documentary dramas, medicine seems to be a great place to tell a story. Perhaps medicine is so broadly represented in our lives because it seems to inform us about some of the most fundamental aspects of our existence; when it starts, when it ends and of the quality we may experience in between.

But this thesis project's title reveals that we will delve into decision-making, and not go into the relevance of medicine to existential philosophy. The ubiquity of medicine in our lives is a significant detail regardless, because it informs us that the average reader of this thesis is most likely to have an intuitive understanding of what decision-making entails in a clinical context. Readers may know that making a diagnosis is not always a logical and objective deduction, but that it often is rather a decision a clinician makes between varying likely options. We learn this from seeing how our general practitioners undecidedly provide us several possible explanations for the complaints we present, and subsequently make an explicit decision whether to refer us to a specialist or not. And if we haven't had the personal experience, we may have seen how Dr. House constantly makes decisions about what the most likely explanation of symptoms is, what additional diagnostic tests to order, what kind of treatments to apply and, especially in the case of Dr. House, how much risk to take in the process and when to throw everything over board and start again from the start. I think that, though we often conveniently rely on the authority and objectivity of clinicians, we know that medicine is being held together by individual decisions, made by doctors and nurses.

That does of course not mean that decisions are made out of nowhere based on nothing more but what feels right to the clinician. These decisions may be based on a wide variety of methods, assumed facts, rules, guidelines, advice from other, standard practices and many things more. But regardless how standardised and objective medicine has ideally become, there is still a major role reserved for the clinician's decision-making. Because it is crucial to understand the current process of clinical decision-making for the compatibility relation with ICDSS to be investigated, this first chapter aims to provide a

characterisation of current medical decision-making and knowledge production as a starting-point for the compatibility analysis.

In service of this characterisation, this chapter will first of all explore how medical decision-making can at all be characterised. The art-vs-science distinction will be argued to be a dominant heuristic for understanding decision-making. The information-processing model and the intuitive-humanist model will be presented as examples of this distinction. It will subsequently be presented that this thesis shall subscribe to Miriam Solomon's rejection of the art-vs-science distinction and use her proposed messy methodological pluralism as a starting point for coming to a characterisation of current medical practice in terms of decision-making and knowledge production. In service of this, a large selection of the methods Solomon identifies are translated to knowledge production and decision making, in order for them to be explained and analysed. The pluralism will be dealt with pragmatically, ultimately arguing that a tidy methodological pluralism is required for later analysis, which is presented to be partially found in a hybrid of Solomon's 'current methods'.

Clinical Decision-making - An Introduction

As has been mentioned, everyday understandings of medicine often conceive hospitals just as the place where the ill and the suffering go to be healed or relieved. In this view many patients expect their doctors to have the ability to objectively identify what is or is not wrong with them and to know what optimal treatment would entail. Actual medical practice deviates from this in many ways. Hospitals are first of all very complex organisations that do much more than merely treating patients. They are also sites where a lot of knowledge is created, where clinicians are being educated and trained and where all sorts of experiments are being performed in service of this. Another significant simplification is that clinicians have the ability to objectively establish a patient's condition and determine the best course of action. Of course clinicians do indeed come with conclusions on what is or is not wrong with a patient and they do indeed prescribe or perform treatments; ideally to satisfying results. But instead of these conclusions and prescriptions being self-evident, or in other ways entirely logically following from whatever is presented to a clinician, these conclusions and prescriptions are very often fully conscious decisions. Decisions that are made by clinicians as the result of decision-making processes that are complex enough to leave room for uncertainty, judgement calls and variability between identical patients; both in diagnosis as well as in treatment (Grimshaw & Russell, 1993; Shafi et al., 2014; Timmermans & Mauck, 2005).

In investigating the compatibility of ICDSS in current medical practice it is very important to understand how current processes of decision-making and knowledge production are taking place. The logical first step would be to come to an analytical understanding of what decision-making is. Doing this is not something in which this thesis will be the first. Many authors have made attempts to do this; and overwhelmingly many of them arrive at somewhat dichotomous characterisations, such as the art-versus-science distinction (Banning, 2008; Battista, Hodge, & Vineis, 1995; DiMatteo, 1979; Heymann, Swift Jr, & Ritter, 2014; Malterud, 2001; Meador, 1965; Peabody, 2015; Piquette - Miller & Grant, 2007).

Though this thesis will ultimately not come to use these types of characterisations, they are because of their ubiquity highly influential in the field of clinical decision-making. Addressing these

dichotomies will therefore be helpful to distinguish the approach followed by this thesis from these characterisations. In order to explain what is meant by the dichotomous characterisation', one example will be reviewed, of which many similar ones exist. Banning's (2008) description of the *Information Processing Model* and *The Intuitive-Humanist model* will be briefly presented.

Information Processing Model

The Information Processing Model poses that decisions are made by applying the hypothetico-deductive method as the basis for conscious and judicious reasoning (Gordon & Franklin, 2003; Graber, 2003). According to Kovacs and Croskerry, this model suggests that clinicians generate, evaluate, refine and verify hypotheses of possible explanations (1999). Generation entails that clinicians formulate possible explanations for the available observations and information on the patient (Barrows, Norman, Neufeld, & Feightner, 1982; Elstein, Shulman, & Sprafka, 1978; Kassirer & Gorry, 1985). Evaluation generally involves confirmation and elimination on the basis of gathering additional information on the patient (Kassirer & Gorry, 1985; Sackett et al., 1996). Refinement takes place parallel to evaluation of the hypothesis, and involves specifying certain hypotheses and dropping others on the basis of the outcomes of evaluation (Kovacs & Croskerry, 1999). Verification of the hypothesis ultimately involves a final check before accepting a certain diagnosis. Kassirer and Copelman argue that the hypothesis is tested for the consistency of the representation with the hypothesis (adequacy), the appropriateness of causal or pathophysiologic links (coherency), and whether the hypothesis involves the simplest possible explanation for the observed phenomena (parsimonious nature).

However, different interpretations of the hypothetico-deductive method in the context of the information processing model exist next to the one outlined by Kovacs and Croskerry. For example Tanner et al., (1987) propose that it involves the stages of: cue recognition or cue acquisition, hypothesis generation, cue interpretation and hypothesis evaluation. It features hypothesis generation and evaluation in a manner similar to Kovacs and Croskerry, but it focusses on the recognition, acquisition and interpretations of cues as distinct steps of the clinician's reasoning. According to Banning (2008) additional stages are proposed by Carnevali and Thomas (1993), namely: entry to the data search field and shaping the direction of data gathering, hypothesis and data directed search of data field and diagnosis. The point of outlining this here is not necessarily to provide a complete picture of how the hypothetico-deductive method is incorporated into clinician's decision-making. It rather aims to point out that different opinions exist on how the hypothetico-deductive method should be modelled.

It is furthermore interesting to mention that a major assumption behind the information processing model, and the assumed role for the hypothetico-deductive method, is that the decision-making of clinicians can be modelled as a process of rational and analytical logical reasoning. It is worth pointing this presupposed rationality out, because it is in stark contrast with the second model of clinical decision-making that Banning proposes, the Intuitive-Humanist model.

The Intuitive-Humanist model

The Intuitive-Humanist model, as it appears in literature, focusses on clinical decision-making in the context of nursing. Though this section drew from interpretations in nursing literature, its core concepts should translate to other areas of medicine as well. The model focusses on experiential knowledge and intuition as important pillars of decision-making, leaving logic and rule-following aside. The model is criticised for its lack of room for scientific and logical reasoning. To overcome this critique Benner (1982) argues that procedures and guidelines as part of logical reasoning are only used

by novice nursing clinicians, and that as experience increases decision-making starts to rely on intuition instead. According to Benner and Tanner (1987) this intuition can be defined as 'understanding without a rationale'. Banning argues that other definitions can be proposed, all of which come down to the absence of conscious reasoning and independence of linear analytical processes of reasoning. In many ways it seems to be understood in a way that is very compatible with the concept of tacit knowledge. A concept that was put forward by Karl Polanyi (1886-1964) as part of a comparable critique of rationality and algorithmic rules as the basis for knowing in general. The concept of tacit knowledge entails that there is knowledge that is not easily transferred or used by means of verbal communication or conscious logical reasoning, but that it is visible only in the context of applying the tacit knowledge in practice and that it is transferred by extensive personal contact, regular interaction and trust (Goffin & Koners, 2011; Schmidt & Hunter, 1993). Like tacit knowledge, intuition is learned through experience (Lam, 2000). King and Clark (2002) argue that it is also a condition for adequate medical practice because it improves the quality of decision-making, much in the same way that tacit knowledge is necessary for mastering a craft.

The Art vs. Science Distinction

As announced, the information processing model and the intuitive-humanist model were shown not to be a unique pair opposites in attempting to explain the way in which processes of clinical practice can be understood. It in fact, they relate back to the all-too common general distinction in medicine, the art-versus-science distinction. The art-versus-science distinction creates an image of medicine where the essence of successful clinical practice lies either in implicit understandings that are honed by experience and apprenticeship, or in logical reasoning, rule-following and calculated decisions and interactions. Especially for philosophers of science and technology this art-versus-science distinction has a familiar ring to it, and is argued by some to have originated from the Greek distinction between *technê* (craft) and *epistêmê* (knowledge) (Solomon, 2015, p. 6).

The examples that were outlined above on the basis of nursing literature provide only one crystallisation of this distinction, many other but ultimately similar distinct models of reasoning could have been constructed or described. Perhaps with different names, but all of which would most likely have come down to the art-versus-science distinction.

This distinction is not being addressed because of its accuracy or because of its helpfulness in understanding clinical decision-making. It is however a distinction that has so widely been adopted in the past, and that it may intuitively feel so useful and right, that it is necessary to explicitly establish the position that the art-versus-science distinction is not a helpful one if we aim to understand how current medical practice deals with making decisions.

The art-versus-science distinction is unhelpful in the context of this thesis, given the goal to analyse the compatibility of current medical practice with ICDSS. The problem of relying on the art-versus-science distinction is that it provides very little insight into the kinds and types of knowledge that are being used within clinical practice. It forces the assumingly rich variety of epistemic characteristics, processes and values to be reduced to falling into an epistemology of 'art' or epistemology of 'science'. Though that may be fine in some cases, it would result in a lack of detail for being able to adequately investigate the compatibility of medical practice with ICDSS. A lacking level

of detail is problematic, justifiably assuming that the epistemic compatibility of medical practice with ICDSS depends on the epistemic characteristics of the actual epistemic processes clinicians go through in decision-making and knowledge generation.

Instead of the art-versus-science distinction, a much more fine-grained account of the methods of decision-making and knowledge production that are employed in medical practice is needed. Such an account can be constructed from works of Miriam Solomon, an authority on the philosophy of medicine that has attempted to provide insight into the variety of methods that medical practice employs. In particular her perspective from "making medical knowledge" (2015), a book that shall be increasingly referred to in this thesis, will be shown to be useful inspiration. In line with the reasoning above, and the requirements for a characterisation of current medical practice, Solomon argues that the art-versus-science distinction is not only redundant in classifying medicine as either one of them, but that it is also unproductive as a heuristic to characterise different aspects of medical practice as hybrids of art and science. In other words, Solomon proposes to discard of the categories of art and science as tools to understanding medical practice.

Apart from the ones outlined above, Solomon provides some additional reasons why art-versus-science fails to do justice to understanding medicine. The first reason Solomon brings forward is that neither art nor science is apt to characterise the role of group judgement or consensus conferences in medicine, a role that Solomon describes to be essential to medicine. She describes group consensus as often thought of as an epistemic ideal, and not necessarily a scientific one; nor from art, for that matter (2015, p. 5). The second shortcoming Solomon identifies is that using unified categories obscures the fact that even within the methods that may be understood as 'science' or 'art', the differences between the methods may be epistemically significant. In more specific words, the differences between, for example, causal reasoning, evidence-based medicine, medical decision-making and translational medicine, are relevant in understanding the epistemic processes of medicine; even if they could be categorised as 'science' there is still more than one way of being scientific (2015, p. 5).

Solomon argues to replace this redundant dichotomy with what is seemingly needed by this thesis, an account of the actual methods that clinicians employ. In providing this account, Solomon sets herself the goal of providing a new and general way of conceptualising and evaluating epistemological techniques in medicine. Or in her own words, Solomon aims to solve the dichotomy by replacing it "with developing an untidy, methodological pluralism" (2015, p. 225). Solomon's account, as will be shown, particularly satisfies the requirements for the methodological account, because of the comprehensiveness of the account; not merely outlining the most important methods of today; but also addressing modern medical movements and those that are starting to belong to the past.

A collection of methods

Part of Solomon's messy methodological pluralism is that Solomon is very specific and distinct with regard to the different methods or epistemological techniques that she identifies within medicine. In this section we shall set out to briefly summarise the palette of methods that Solomon identifies. Describe and summarise them, and make some general remarks about how they relate to each other and current medical practice. From this description we can construct a rather complete characterisation of decision-making and knowledge production under current medical practice. It is

important to note here, that Solomon's perspective is explicitly about the epistemic aspects of medicine in general; and that she subsequently tends to focus her analysis on the construction of knowledge rather than on the application of knowledge in clinical decision-making. The dimensions of decision-making are rarely left unaddressed, but for solving situations where Solomon is insufficiently clear about the decision-making it has been argued in the previous section that it seems sensible to aim at a characterisation of decision-making on the basis of Solomon's descriptions of the methods of knowledge production as it provides a strong indication of the kind of knowledge the methods employ.

With regard to these description, Solomon proposes to put these methods into three categories. She explicitly mentions *traditional methods* and *new methods* as being marginally epistemically interesting and to be discussed only briefly; and implicitly announces *current methods* as interesting and to be described in length (2015, p. 16). Though this categorisation may come as a slight surprise, given that Solomon decidedly aims for untidiness in her overview, we shall follow this categorisation for reasons of structuring the collection. This should however be under the explicit remark that the implied hierarchical relation between the methods should be carefully dealt with, preventing us to follow it without proper questioning.

With regard to carefully dealing with Solomon's work, this thesis may at this stage in the process provide the implicit suggestion that Solomon's work is adequately applicable as it is; not requiring analytical work or interpretation from our side. This will however be shown to merely be a suggestion. It seems first of all clear that Solomon's attempt to construct a methodological pluralism is not aimed to providing a starting point for analysis. It is instead intended to provide insight into the processes of decision making in order to elucidate the epistemological role of the methods and to point out that many of these methods have strongly social components rather than merely epistemological ones.

Solomon's focus on contributing to understanding ultimately requires this thesis to reduce the collection of methods to a tidy methodological pluralism. Taking the insight from Solomon's messy pluralism into account, but presenting it in a 'tidy' way that allows for further investigation of the compatibility with ICDSS.

The fact that Solomon's description aims to focus on the kinds of knowledge that are used in medical practice, also means that the methods she identifies are primarily described in terms of their role in the processes of knowledge production rather than decision making. That requires us to perform some interpretative work, and making combinations with existing descriptive empirical work on decision-making, for the tidy methodological pluralism to find its focus in the combination of decision-making and knowledge production.

It is furthermore seemingly unclear what Solomon in general identifies as a method. Though it seems that Solomon largely follows an existing ontology of methods, as can be found throughout literature on medical epistemology and decision-making, and that this seemingly secures an adequately complete overview of the methods, it should be taken into consideration that this may make this thesis susceptible to oversight of methods and potential other ways of understanding clinical practice in terms of decision-making and knowledge production.

Traditional methods in Medicine

Chronologically the most logical category to start with are the traditional forms of medicine. The name of the category may come to suggest that they have gotten out of fashion and may be used no longer in current medical practice. Solomon, however, argues that, though methods of clinical

decision-making and their corresponding epistemic cultures are often understood as paradigms, they fail to comply with a key characteristic of Kuhnian understanding of paradigms (Kuhn, 1962). In particular Kuhn's understanding of a paradigm implies that the 'new' paradigm fully replaces the old. Making the old indeed the previous and the new indeed the current. This is precisely where these clinical decision-making methods are different, because they are not constantly being replaced by the 'new'. Instead 'the new' becomes part of the current just by becoming integrated into a hybrid of methods and principles, Solomon argues (2015, p. 5). An argument that will be shown to be sensible as soon as several methods have been described, allowing the reader to see that they may indeed coexist without problems.

In other words, the methods that will be discussed under traditional methods shall still be a significant part of what current medical practice consists of; even if another paradigm (in a more relaxed non-Kuhnian interpretation) appear to be dominant. Additionally, traditional in this context means that these concern methods or mechanisms that have traditionally been relied on. Solomon provides case study reasoning, appeal to authority, clinical experience and causal mechanistic knowledge as primary examples (2015, p. 5). Though she announces to explain how these methods are to be understood, it unfortunately appears that she doesn't always do so explicitly.

In order to overcome this mild shortcoming we shall briefly construct plausible descriptions of how Solomon's intends the examples to be understood. Case study reasoning, the first example Solomon provided, is named in such a way that it suggests that the reasoning in a decision-making process of a newly presented and not immediately familiar problem is significantly inspired by what the decision-maker knows from previous cases, described and analysed in case studies. In a brief literature review of what case study research entails, Sarah Crowe et al (2011) conclude on the basis of influential publications by Stake (1995), Yin (2003), Miles and Huberman (1994), Green and Thorogood (2009) and George and Bennet (2013) that an essential aspect of case study research is that it offers a solution for *"the need to explore an event or phenomenon in depth and in its natural context"* by generating a *"multi-faceted understanding of a complex issue"*. In the medical context this seems to translate to a mode of decision-making on the basis of in-depth understanding of descriptions and analyses of previous clinical interactions.

Medical reasoning on the basis of authority, or authority-based medicine, is described by Solomon as always having been a foundation for medical knowledge by means of establishing and transferring knowledge based on trust of authorities (2015, p. 225). Decision-making on the basis of authority could consequently involve one of two different mechanisms. Firstly this may involve that a decision-maker in a specific situation follows a particular suggestion or instruction that has been put forward by someone of authority, mostly, or even entirely, because of the proponent's epistemic authority. A second, and more 'remote' mechanism, is when a decision-maker takes any kind of decision belonging to any kind of method, on the basis of at least partially authority-derived knowledge.

Reasoning on the basis of clinical experience, the third method of knowledge production that Solomon identifies as traditional, shows a lot of overlap with the example of the intuitive-humanist decision-making model. The essence of this aspect reasoning on the basis of clinical experience is that the experiential knowledge that a practitioner possesses, is key to the decision-making process. In other words; clinical experience based decision-making employs intuitions and tacit knowledge that the practitioner develops on the basis of their own personal experiences with clinical practice.

However, Solomon also identifies that reasoning on the basis of clinical experience can also be based on more articulated; verbalised or explicit forms of knowledge. Solomon writes that, though clinical experience "*is not as rigorous as clinical trials, [it]... is capable of detecting reproducibly strong evidence*" (2015, p. 118). This means that reasoning on the clinical experience not only includes judgement on the basis of tacit forms of knowledge, but also on explicit forms of knowledge. To explicate this in the context of an example, both treating a patient a certain way because the practitioner explicitly knows he or she has seen him or herself that it has worked in previous cases, is a form of experience-based decision-making; as well as a practitioner 'intuitively' ordering additional diagnostic tests because he or she vaguely has 'a bad feeling about it on the basis of unarticulated previous experiences'.

Causal-Mechanistic Knowledge, the last traditional method that Solomon identifies clearly relates to the rather well described method of causal-mechanistic reasoning. Causal-mechanistic knowledge, in the context of clinical decision-making, entails knowledge on the basis of a pathophysiologic rationale (Howick, Glasziou, & Aronson, 2010). In simpler terms this means that causal-mechanistic knowledge consists of explicit models of how the individual parts of the human body function in relation to each other and how that relates to health and disease. This may entail models of the 'function' of individual organs, but also on individual cells within the organs; or even individual molecules within these cells. The causal element means that the method presupposes that there exists causal relations between different 'actions' of the individual parts that can be known and understood through the model. The combination of dividing the functioning of the human body into different parts, and 'knowing' the modelled causalities between the different parts ultimately gives the decision-maker the ability to understand how disease occurs and how and whether there should be intervened with these mechanisms for them in order to produce health instead of disease¹. In causal-mechanistic reasoning the decision-maker can subsequently use that model to come to rational explanations for symptoms the patient provides or to come up with solutions for these symptoms and make decisions about this. Causal-mechanistic reasoning, though traditional by Solomon's standards and controversial in terms of accuracy, is still widely developed and used in current medical practice (Howick et al., 2010).

Though Solomon's categorisation as 'traditional' may suggest that this method may be somewhat insignificant in comparison to other methods, it is useful at this point to explicitly mention that this thesis shall later on show that the principles of causal-mechanistic reasoning are, even as a tradition, extremely deeply rooted into medical practice.

New methods

The second group of methods that Solomon identifies, and labels as epistemically significant to a limited sense in current medicine is the group of *new methods*. Whereas the group of traditional methods is characterised by methods that have been around for a long time throughout the medical tradition, new methods are characterised by their recent emergence into medical practice. This also means that, whereas influences of the traditional methods can be found to almost ubiquitously play

¹ Please note that adherents of scientific realism might argue that the knowledge that feeds into causal-mechanistic reasoning is not merely knowledge of models proposing an understanding of the human body; but is actual knowledge of the real functioning of the human body itself.

part in decision-making, similar influences of the new methods might lie in the future and currently be limited in scope.

Because this thesis aims to investigate the compatibility between ICDSS and current medical practice, one might argue that the influence of the new methods is currently too insignificant to be relevant for a characterisation of current medical decision-making. However, because ICDSS is a technology whose ultimate impact would take shape in the future rather than the present; and that it would most likely have to be integrated into an evolution of the current, it is probably informative to characterise these methods nonetheless.

It should nevertheless be taken into consideration that some of the new methods may actually be in early conceptual phases. This means that the accounts that Solomon provides may in many cases be normative ideals, describing how a certain method should take shape, rather than derivatives from decades of empirical research and insight into the actual shaping of the traditional and the current methods.

The first new method that Solomon addresses is *medical decision making*. A confusing name, given that this thesis has used these and very similar words throughout this piece. In order to prevent confusion on the reader's part, it shall be referred to as *medical decision making (method)*². It is important to prevent this confusion because according to Solomon, medical decision making (method) refers to a specific sub-category of decision-making methods. Solomon describes that medical decision making (method) "makes use of Bayesian probability theory and an investigation of harms, benefits and individual variation on those harms and benefits" (2015, p. 115). This is explicated when Solomon states that it is a movement that is derived from important advances in the field of the cognitive psychology of choice, among others by Tversky and Kahneman (Elstein, 2004). As a result, Solomon writes: "Medical decision making (method) seeks to avoid common errors of judgement, such as overestimation or underestimation of risk, framing biases, and availability and salience biases" (2015, p. 215). Put in terms of decision-making this means that knowledge of statistical analysis and risk assessment in combination with specific statistics about the relation between symptoms, diseases, therapies, beneficence, etc. feed into the decision-making process. In other words, medical decision making (method) may help a clinician to take decisions on the basis of statistics; for example of what the most likely diagnosis is, or on what the best possible treatment is, in relation to the chances of harm and benefit due to success or failure.

Though Solomon places this method into the 'new' methods, Bayesian reasoning has a considerable history in the field of decision-making. To be more precise, as will be shown in the next chapter, Bayesian reasoning has been used in some of the early attempts at providing algorithmical clinical decision support.

Shared decision making, the second new method Solomon mentions, refers to a method that is exceptional in the sense that the clinician does not make a decision on his or her own for the patient, but instead makes a decision based on mutual agreement with the patient. This actually seems a lot like the principle of informed consent; a requirement for medical interventions that states that the patient is informed of the chances and types of harm and benefit of a particular procedure, and consciously agrees to its execution. However, under informed consent the patient is figuratively speaking merely the one who has a final say in whether to treat or not. Under shared decision making,

² Please note the absence of a joining hyphen

however, the patient is invited to weigh into the decision-making process itself by expressing and identifying values they may have with regard to the decision, desires for avoiding or taking on particular outcomes, acceptance of risks and chances, or other medical questions where uncertainty and normativity plays a significant role and clinicians have no clear and self-evident standards for what is 'best' for the patient (Frosch & Kaplan, 1999). It should furthermore be noted that shared decision is not a stand-alone method for decision-making. The clinician, after all, needs to go through his or her own processes of decision-making in order to decide what the possibilities are at all. Shared decision-making is furthermore ideal for uncertainty with regard to treatment, but stops to be useful with regard to diagnostic decision-making or acute medical emergencies.

The third new method Solomon identifies is systems analysis and is again a method that Solomon does not explicitly address. It is however most likely that Solomon refers to the field of systems medicine, a fairly recent development spun off from the field of systems biology. Systems biology is a holistic approach to understanding complex biological systems, mostly on the cell-level, through understanding the roles of the individual sub-systems in relation to each other and to the entire system's functioning (Kitano, 2002). In the context of medicine that means that the functioning of the human body is being studied as the outcome of the interrelations between ideally all of the smallest parts and processes of the human body. It should then come as no surprise that systems medicine is heavily engaged with the field of genomics, given that genes are generally considered to play a crucial role in how the human body as a complex system is regulated (Auffray, Chen, & Hood, 2009). This may seem a lot like causal-mechanistic reasoning in terms of the dependence of knowledge of a system of interacting parts. However, the method of causal-mechanistic reasoning is often not used as a holistic approach; meaning that they permit parts and sub-systems to be isolated from rest of the parts and systems of the human body, and to be reasoned with. This often means that causal-mechanistic reasoning presupposes that the decision-maker can actually use a functional model him or herself to explain certain symptoms or to plan for interventions accordingly, systems medicine considers these models as highly complex and not reasonably comprehensible by human actors. This means that models for causal-mechanistic reasoning are significantly simpler and more macroscopic than models for systems medicine. Systems medicine subsequently relies heavily on computational models that incorporate the countless relations and interactions between the many parts and processes the human body consist of. Systems medicine is currently perhaps more of a scientific enterprise than it is incorporated directly into the daily practice of clinicians. Knowledge produced by systems medicine on the other hand might actually currently feed into other decision-making methods. On the other hand, in order to actually fully operationalise systems medicine as a method for decision-making would require a computational model to be fed a lot of knowledge about the patient, and for a clinician to consult a computational model.

Personalised medicine is probably the best known method of the new methods that Solomon provides. This is assumingly partially due to a number of relatively successful commercial businesses bringing some of the early promises of personalised medicine to the consumer markets, such as how 23andMe provides customers personalised health advice on the basis of DNA-sequencing. The idea behind personalised medicine is that individual patients differ among each other in many ways, and that one of these may include the way they may react to treatment, or what kind of symptoms they present under different circumstances on the basis of their *personal* biochemistry. Solomon describes that one of the early hopes of personalised medicine was that "*new genetic therapies would be*

tailored to individual mutations", but that the meaning of the term has come to exceed the genetic domain, and include *"all ways of tailoring medical knowledge in application to specific individuals and to signify that "cookbook medicine" is not being used"* (2015, p. 247). In other words, in order to be able to take adequately effective decisions about a patient's diagnosis or treatment, the clinician needs to know the specific characteristics of that particular type of patient. In terms of decision-making it should be mentioned that personalised medicine is not as much a method for decision-making as it is an ideal for optimal clinical care. The parallels can furthermore be observed with systems medicine, suggest that— depending on the level of personalisation—decision-making in the context of personalised medicine is very complex and may not be executed by a fully autonomous clinician. He or she would most likely need to consult some computational model that incorporates the influence of the variations in biochemistry (or other characteristics) on the processes of diagnosis and therapy into classifying and understanding the conditions of the individual patient.

"Integrative medicine represents a higher-order system of systems of care that emphasizes wellness and healing of the entire person (bio-psycho-socio-spiritual dimensions) as primary goals, drawing on both conventional and CAM [complementary and alternative medicine] approaches in the context of a supportive and effective physician-patient relationship." (Bell et al., 2002).

Integrative medicine, in other words is to some extent a holistic treatment approach, in which decision-making aims to achieve the optimal composition of biological, psychological, social and spiritual factors. Part of this is that 'health' as the ideal outcome of care, is defined broader than merely functional and mechanical aspects of the human body; for example by incorporating the patients experience of spirituality. This also implies that the methods that are used to achieve outcomes with regard to the non-biological aspects of health lie outside the domain of what we often consider as conventional medicine. So in addition to conventional medicine integrative medicine requires the, unsurprisingly, integration of so-called complementary and alternative medicine. The presupposition is not necessarily that these complementary and alternative forms of medicine contribute to the biological aspects of health, but that they may very well be helpful in the spiritual aspects of health. Because these non-biological elements, and perhaps the biological elements as well, are highly personal, there must be a strong element of shared decision making in the process of integrative medicine. Integrative medicine in general relies on other methods of decision-making to explore options and to make decision within the different aspects of health. It then still requires decision-making by clinicians within the framework of integrative medicine to determine how possibly conflicting insights and options from the different frameworks relate to each other; especially between different medical paradigms (Bell et al., 2002).

The last new method that Solomon refers to is hermeneutical medicine. Solomon mentions it only briefly when she states that *"a physician with hermeneutic goals is likelier to be empathic"* (Solomon, 2015) and that *"Recent works on empathy has roots in the hermeneutics tradition"* (2015, p. 178). Solomon later on defines empathy as the quality of a practitioner as not merely the ability to listen to a patient, but to know what a patient is going through at an experiential level (2015, p. 183). Searches through medical literature did not reveal clear and explicit understanding of what Solomon might understand to be hermeneutical medicine. However, drawing from the citations above it seems that Solomon means to refer to a form of medical decision-making that reserves a crucial role for clinicians to know the experiential implications of their decisions for their patients, and take them as knowledge into their considerations. In terms of knowledge production this means that clinicians

require experiential knowledge of undergoing all kinds of clinical scenarios, and have an understanding of how this knowledge should be factored into the decision-making process; something that does not seem self-explanatory at first sight.

Current Methods

Current methods is a description for the group of methods that Solomon identifies as being of significant influence to understanding the social epistemology of clinical practices under current medicine. Whereas Solomon only briefly addresses, if at all, the traditional and new methods of knowledge production and decision-making, she describes the current methods of consensus conference, evidence-based medicine, translational medicine and narrative medicine in great detail in her critical analyses of these methods. In this section we shall follow Solomon's example and characterise them as constitutive of current medical practice. To be clear, it should be emphasised once more that Solomon's work focusses on the knowledge producing aspects of the four methods, and the characterisation aims to come to a more complete understanding of the epistemic processes by deriving and characterising aspects of knowledge use and decision-making as well.

Consensus Conferences

Consensus conferences, the first current method that Solomon mention, are very briefly put meetings where experts gather to resolve the inhibiting effect of scientific controversy by seeking consensus (Solomon, 2015, p. 26). Solomon describes how the consensus conferences started as the product of the United States National Institute of Health's (NIH) Consensus Development Conference Program in the late seventies of the previous century. The foundation for the idea of organised consensus seeking can be found in Kantrowitz's model of the 'science court' (1967, 1977). The science court is a proposal based on legal courts, in the sense that it would feature independent judges evaluating and judging scientific evidence. The goal of the court model was not to arrive at truths or at certain facts, but to evaluate the existing scientific knowledge objectively and free from biases, to enable collective action (such as implementing technologies) while there would still be scientific controversy over particular facts (Solomon, 2015, p. 26)

It seems that especially this last feature inspired the NIH, because none of the models describing the NIH's consensus conferences include judges; however all of them incorporate the aim for an instrument to deal with uncertainty and controversy in medical scientific knowledge. In the court model this instrument is a judge's, or a panel of judges', verdict. In the case of consensus conferences this instrument is the consensus found by a panel of a dozen or more panellists. The consensus could lie in one of two 'dimensions' of consensus; in *technical* consensus and in *interface consensus*. Technical consensus related to scientific matters whereas interface consensus related to broader societal, ethical, economical, etc. matters (Solomon, 2015, p. 27)

Solomon explains that this meant in practice that after topics for contemplation had been successfully suggested by anyone, a planning committee consisting of a variety of experts decided how to frame the questions, after which panellists would be chosen from unbiased clinicians, researchers, research methodologists and patient representatives (2015, p. 29). The actual conferences, open to

public audiences, took place over a period of three days of presentations, discussions and several iterations of drafting and adapting the consensus statement that would be announced in a press conference by the end of the third day; after which a final dissemination to the medical and research communities followed with a month's delay after a final round of revisions.

Solomon describes how other organisations followed the NIH's example, and started to organise consensus-oriented meetings both within the US and across Canada and Europe. Precise implementations varied among time and place, but remained recognisable as gatherings of people assessing knowledge claims and seeking for consensus. Though consensus conferences are often assumed to have been fading out into something of the past since the turn of the millennium, Solomon argues that mostly applies to the name but not to the practice: "*...the same basic ritual of consensus conferences continues. Face-to-face meetings lasting 2–3 days with 10-20 invited people [producing] guidelines, recommendations, and standards of care*" (Solomon, 2015, p. 78)

Especially because of the continuing relevance that Solomon describes, it is important to delve into the epistemic role that the consensus conferences of the past and the nameless gatherings of the present play in the production of knowledge and decision-making. A few of Solomon's important observations are pointed out when she writes that "... consensus development conferences were choreographed social epistemic rituals, designed to produce conclusions that were perceived as objective by its intended audience..." (2015, p. 61). Though she writes this in the context of the NIH variety, it seems that this applies reasonably well to other consensus conferences as well because they are all equally social, and can be just as well understood as ritualistic.

Solomon's choice for the word 'ritual' is provocative. On the one hand the word ritual resonates well with the connotation of ritual as an often repeated and standardised sequence of activities, which these conferences seem to have been judging from Solomon's historiography. However on the other hand, the word ritual inspires the suggestion that the standardised sequence is more like an unquestioned product of tradition and arbitrary culture, then it is like the continuously evolving and rational outcome of a justifiable and rational (as the NIH intended it to be) process of evaluation and development. This suggests that there is no clear justification for the epistemic outcomes of these conferences. In other words, if there is no leading rationale behind the way in which consensus conferences arrive at resolutions of knowledge debates, there is no obvious purely epistemic argument for its validity or usefulness.

It is suggested in the choice for the word 'social', as an adjective to epistemic ritual, that Solomon largely shares the previous conclusion. Social in this context suggests that the justification or the origin of the knowledge that the consensus conferences produced did indeed not lie in clear epistemological rules that could have been objectively applied by one single individual; but rather in the social processes performed in definition by groups.

Apart from this social dimension (regardless of whether it should be seen as desirable or not) there are other interesting epistemic considerations to be made about consensus conferences. First of all there is the fundamental question of the epistemic value of consensus as an instrument of epistemic goals in general. A position that is critically supported by a variety of historical examples where consensus existed on issues that are now considered entirely wrong, such as the motion of the earth, alchemic practices and miasma theory (Solomon, 2015, p. 37). Other critiques involved scepticism about deviating from the scientific tradition and ideal of resolving controversy by doing "rigorous

testing and critical evaluation", and the claim that scientific controversy is useful for scientific productivity (Gunby, 1980; Solomon, 2003). It was furthermore argued that consensus conferences failed to facilitate nuanced judgements with the too-general guidelines the conferences produced with great authority (Holleb, 1980).

Solomon adds to this herself a firm critique on consensus conferences through a critique of group judgement in general. Solomon argues that there are reasons to believe that group processes may be less objective, and thus epistemologically generally speaking less desirable, than what heterogeneous individuals might achieve. Reasons for this, she supposes, may be found in psychological factors causing biases; such groupthink, peer pressure, anchoring phenomena, etc. All these lead the group deliberation to make sub-optimal use of the available knowledge and attitudes that may exist and ultimately contribute to the process (2003, pp. 97–98).

Evidence-Based Medicine

The second current method that Solomon announces is the method of Evidence Based Medicine (EBM). Already in the introduction of this thesis attention has been explicitly attracted to evidence-based medicine, already providing an early suggestion of its position as an important method of knowledge-production and decision-making. The standard story about the inception of EBM starts with the Evidence Based Medicine Working Group publishing in which EBM is announced as a 'new paradigm' for medical practice (1992). Much of the standard story of EBM focusses on how it has been, and continues to be, a normative ideal for how clinical decision-making should be taking place on the basis of scientific evidence. An influential formulation of this ideal was published in an editorial piece in the British Medical Journal, and state that: "*Evidence based medicine is the conscientious, explicit, and judicious use of current best evidence in making decisions about the care of individual patients.*" (Sackett et al., 1996). 'Current best evidence' is to be understood as the outcomes of "*...clinically relevant research, often from the basic sciences of medicine, but especially from patient centred clinical research into the accuracy and precision of diagnostic tests (including the clinical examination), the power of prognostic markers, and the efficacy and safety of therapeutic, rehabilitative, and preventive regimens.*" (Sackett et al., 1996)

This formulation confirms, and perhaps inspired, the conceptualisation of the novelty and relevance of EBM as lying in the crucial position of evidence. Solomon, however, argues that would be a misrepresentation, and that evidence has always played an important role in medicine (2015, p. 106). This is a strong point. Drawing from the countless methods that have been described earlier in this chapter it can easily be seen how so many of them rely on one form of evidence or another. But if it is not the cruciality of evidence that sets it apart from the other methods, what does?

Whereas the ordinary narration of the history of EBM starts with the working group's publication, Solomon presents how the working group was in fact inspired by *Effectiveness and efficiency: random reflections on health services*, in which Archibald Cochrane powerfully and ultimately influentially argues for the widespread use of randomised controlled trials (RCT) and for the existence of accessible databases of their results (Cochrane, 1972). Instead of putting evidence in general centre stage, Cochrane praises one very specific kind of evidence above all others; namely evidence following from RCTs.

In other words; for Cochrane to be able to put RCTs on top of the list; the available forms of scientific evidence had to be put in some order of desirability. This evaluation of the epistemic quality of evidence, and the resulting pecking order, is precisely what EBM incorporated. It has actually been very formally included into EBM in the form of the so-called *Hierarchy of Evidence*. A hierarchy of evidence can of course be constructed in several ways, and subsequently there are a few different attempts to provide tools to hierarchise evidence. The most influential attempt to do this is the *Grading of Recommendations Assessment, Development and Evaluation (GRADE)* system, which orders evidence into four categories of quality on the basis of the type and quality of the study producing the evidence: High, moderate, low and very low quality (Guyatt et al., 2008; Solomon, 2003, p. 109). Ironically, given that EBM is a method to partially overcome the epistemic weaknesses of consensus, as both speaks from the title of Guyatt et al.'s 2008 article as well as is identified by Solomon explicitly; deciding on the proper hierarchy, and determining how different types of research should be qualified within it are again a highly social and consensus-oriented process (2015, p. 5).

It is according to Solomon especially this crucial role for the hierarchy of evidence in EBM that sets it apart from other methods of knowledge production and decision making. Not is merely stating that one method is the best, but by hierarchising all the possible forms of evidence—from RCTs of proper quality and size to case-based evidence and expert opinions—and using that to systematically and rationally draw conclusions from large collections of evidence. It is for this reason that Solomon makes the playful remark that EBM should actually have been called "*evidence-hierarchy medicine*" (2015, p. 106). This actually makes a lot of sense, realising that of only few methods of decision-making could be argued that they do not base themselves on some form of evidence.

Finally, before going into the epistemological implications and aspects of EBM (regardless of how the method should be named) Solomon argues that Davidoff et al. (1995) provide a useful 'canonical' definition by stating that: "*In essence evidence based medicine is rooted in five linked ideas: firstly clinical decisions should be based on the best available scientific evidence; secondly the clinical problem—rather than habits or protocols—should determine the type of evidence to be thought; thirdly, identifying the best evidence means using epidemiological and biostatistical ways of thinking; fourthly, conclusions derived from identifying and critically appraising evidence are useful only if put into action in managing patients or making health care decisions; and finally, performance should be constantly evaluated.*"

In identifying and hierarchising the quality of evidence also lies a major epistemological controversy. Namely: can evidence at all be hierarchised, or is the collection of knowledge of evidence so diffuse and heterogeneous—even within certain categories of experimental methods—that trying to do so discards of valuable scientific input. This problem may lie in the impossibility of establishing reliable criteria for assessing the quality, something Michael Rawlins claims it to be (2008, p. 579). But the fact that meta-analyses rank high in many hierarchies, also means that the selections that these meta-analyses make on which trials and studies to include in their analysis have an aggregating influence, putting pressure on the diversity of available studies because the meta-analyses are weighed heavier in the hierarchy than the individual studies (Solomon, 2015, p. 113).

This connects to a similar general criticism on the authority that many of these hierarchy systems, such as the GRADE system, attach to RCTs. An RCT is a research method that tests for the influence of certain interventions on patients by separating the research population into an intervention group and a control group. Both groups are treated as similarly as possible, apart from

the fact that the control group does not actually receive the intervention whereas the intervention group does. The authority of RCTs can become problematic because not every research questions can be answered through RCTs. For example because there are first of all medical scientific questions that cannot be answered without treating both groups differently, for example because some interventions are hard to simulate for the control group. But there are also medical scientific questions that can only be answered longitudinally with too many variables to call it an RCT, and of course questions that are not clearly related interventions. Not being able to use the most authoritative form of evidence in answering these questions has an influence on the epistemic authority that the possible answers to these questions may attract under the paradigm of EBM; even though they may entail empirically perfectly adequate outcomes, or provide knowledge that could be highly valuable for other areas of research. In other words, discriminating between evidence on the basis of criteria that are not perfectly equipped to do justice to the value of evidence inescapably leads to the marginalisation and exclusion of perfectly useful evidence, and the inflation and inclusion of possibly inadequate evidence.

It might seem that decision-making within EBM requires specific knowledge about the available study outcomes on diagnosis, therapy and prevention and knowledge of how the different 'pieces' of knowledge are supposed to be put together and weighed on the basis of the hierarchy of evidence. However, because it would be unrealistic to expect from clinicians that they would hierarchise the available evidence and draw adequate conclusions from them every time they encounter a new patient with new decisions to be made; EBM relies on the construction of standardised clinical guidelines. These guidelines are constructed on the basis of the evidence and their respective hierarchies, and are then communicated to clinicians to be applied to problems (El-Khatib et al., 2010; Gabbay & le May, 2004). This means that the individual clinician only requires knowledge of the guideline in to make clinical decisions. Though this allows clinicians to act according to the normative ideals of EBM, it also requires patients and problems to be standardised in order for the guidelines to be applied. Opponents of this dependence on guidelines often refer to it as 'cookbook medicine' (Dans, 1994; Delamothe, 1993). Solomon explains that many opponents argue that the statistics that EBM relies on to construct these guidelines do not apply to individual patients. Ignoring this shortcoming and, unreflective of the individual patient, applying these guidelines as rigid rules is what cookbook medicine refers to (Solomon, 2015, p. 146).

Translational Medicine

Solomon argues that the second current method that she describes, translational medicine, "*is a response to the combined shortfalls of consensus conferences, evidence-based medicine, and basic science research in the 1990s*" (2003, p. 156). Though translational medicine is attracting a lot of attention and though its ideals are being operationalised in medical practice, exactly what translational medicine is supposed to refer to remains somewhat fugitive, and simple definitions—such as were available for the previous ones—seem to be hard to find. But regardless of the difficulty, Solomon argues that it should be understood as referring to the translation from "*applied research from bench to bedside (and back)*" and "*Moving successful new therapies from research to clinical contexts*", with wider usages of them being opportunistic band-wagon jumping rather than being substantive (2015, p. 159). The *bench to bedside and back* step is referred to as T1, *new therapies to clinical contexts* as T2.

Upon first inspection this may seem to be hardly different than what should be expected of EBM. However, even though translational medicine is a fugitive concept, there is an important difference.

What translational medicine ideally allows is for therapies that are thought of and developed in laboratory environments as the result of rather fundamental research, that are supposed to work but have never been verified, to be applied directly into clinical practice (Solomon, 2015, p. 159). Under EBM this would arguably not be possible as the individual clinician can ideally hardly use fundamental research in combination with mechanistic reasoning because it would rank very low in the evidential hierarchy. Instead a longer path of increasingly advancing clinical trials, initiated from the research domain, would be required. This in turn would reduce the speed with which new insights can be brought to the clinic and with which experimental research and follow-up questions can be asked.

However, this may be more of a stigmatisation of the methods of the past than that it is an adequate reflection of what is 'new' and 'current' about translational medicine. Solomon even argues that there is no new methodology to translational medicine, but that there remains substance to the concept (2015, p. 164). Solomon makes a long case showing that, especially in the pharmaceutical industry, there is a very wide gap between basic science and clinical practice. This gap results in many drugs ending up being discarded before entering serious trials because they supposedly fail to work in the clinical, human, context. However, discarding of drugs because they fail to live up to the expectations is the problem, perhaps, it is argued, there used to be an underestimation of the translation from basic science to clinical practice, easily leading a failed trial to perform as evidence of the impossibility of the approach, rather than a sign that making the translation may be more complex than thought of. Solomon follows Maienschein et al. (2008) in their conclusion that what is new is that there is the explicit recognition of the difficulties that translation entail. Taking these difficulties seriously may lead to improved anticipation of difficulties of translation, and perhaps lead to new and more effective insights with regard to this in terms of clinical outcomes (2015, p. 157).

So, rather than being a mode of clinical decision-making, it is a conceptual development aiming to contribute to the way in which knowledge of basic science can be translated to clinical applications and integrated into clinical practice in general; decision-making included. This also urges us to re-emphasise that this is somewhat of a normative ideal rather than a description of actual practices taking place; also showing from the lack of empirical sources used. Purely for the sake of completeness it should be mentioned that, perhaps, Solomon has been somewhat optimistic in that regard in grouping translational medicine into the group of current methods. Perhaps the ambition is what belongs to the current, while its practice belongs to the future. Maintaining this category seems harmless as long as these nuances are kept in mind.

Narrative Medicine

Solomon addresses narrative medicine as the last and most recent of the current methods of knowledge production and decision-making. Solomon identifies Rita Charon as a key proponent in the field of narrative medicine, who characterises it as "*medicine practiced with narrative competences*" (2001, p. 1897). These competences ensure the clinician's ability to "... *acknowledge, absorb, interpret and act on the stories and plights of others*". Upon first inspection this may seem to be a more articulate version of the device '*to listen to the patient*'. However, Solomon explicitly remarks that this falls short of the actual 'newness' and relevance of the narrative method. Instead narrative medicine shares goals and presuppositions with integrative and hermeneutical medicine in the sense that biochemical 'health' is not the only useful parameter to judge the desirability of clinical outcomes.

There should be attention in the clinical process for the story of the patient to be able to decide what the best course of action or inaction is. The turn towards this more narrative understanding is not merely relevant in clinical decision-making, but also presented as contributing to the relationship between the clinician and the patient.

What 'narrative' precisely is, is hard to grasp. Solomon even argues that a clear and useful definition cannot be provided and that the field of narrative medicine tends to define narrative, and medical narrative, as broadly as possible (2015, p. 181). Though the concept of narrative receives little explicit features, the narrative competences that should lead to the '*acknowledgement, absorption, interpretation and acting on the stories and plights of others*' mostly are explicated by Solomon.

These competences ultimately all contribute to the knowledge that is needed in the decision-making process. The first narrative competence that Solomon identifies is '*listening and witnessing*', and includes more than merely recording the words the patient speaks and registering how he or she moves, but requires attentiveness to "mood, tone and desire". It furthermore requires the clinician's ability to be moved by what the patient expresses (2015, p. 183). '*Empathy*', the second competence, is an extension of listening and witnessing, and requires the clinician to know at an experiential level what the patient experiences. Thirdly '*Narrative Detective Work*' requires the clinician to be able to extract meaningful insight from the narrative of the patient. The diffuse and non-obvious expressions of the patients can be understood as having concrete meanings through detailed knowledge of narrative (2015, p. 188). Finally Solomon identifies the competence of '*meaning giving*', which entails contributing to intimacy between patient and clinician through giving emotional and existential meaning to the results of the narrative detective work.

But the free and open nature of narrative medicine does not mean that there are no epistemological issues. First of all the premise that personal narratives are singular seems to be untenable; larger public narratives sometimes emerge and become part of the personal narrative possibly becoming taught to the patient rather than felt. Furthermore there is the dependence on the truthfulness of the patient. This causes problems when a patient "*misremember, fabricate, exaggerate, lie, distort, selectively tell, or otherwise intentionally or unintentionally, explicitly or inexplicitly report falsehoods...*" (Solomon, 2015, p. 200). When the clinician picks up on the falsehoods it may be of important input to the narrative analysis. However, when the clinician wrongly assumes the presented narrative, or parts of it, to be true; it limits the applicability and effectiveness of the outcomes it may have (assuming there is a positive relation between accurate narratives and outcomes of the clinical process (Solomon, 2015, p. 200).

Finally it should be mentioned that Solomon explicitly sees narrative medicine as a possible solution for the desire for a more humanistic medicine. Though she is not omitting the fact that narrative medicine may be flawed, or that there may be better alternatives; she does not question the need for a more humanist medicine. Though I myself do not have a definitive position with regard to this issue, it seems relevant to express that this position may not be universally shared. It can for example be easily imagined that by some people economic efficiency may outweigh humanist medical ideals.

Methodological Pluralism

In the past sections a detailed overview of available methods of clinical decision-making and knowledge-production in medicine has been provided. Though some of them have explicitly been linked for reasons of similarity, integrability or overlap, the majority of methods have been described independent of the wider context of clinical practice. But in service of coming to a productive characterisation of clinical practice it is important to come to a conclusion on how they relate to each other.

Some of the methods that have been described have been positioned as paradigms, replacements of the old, a pair of goggles through which medical practice should be seen from then on. If this revolutionary picture of the development of medical decision-making is adopted then the characterisation of current medicine should be based on only one of the recent methods of this revolutionary model; such as EBM or translational medicine.

However, this revolutionary picture is not self-evident and, as has been previously mentioned, is contrasted by a more evolutionary approach in which medical practice is constituted of an diversity of methods that changes in constitution; and of which the individual methods are not stable but changing themselves. In this view methods do not continuously replace each other but are added to the collection of methods that are applied. Solomon, as has been previously mentioned, described conceptualising the collection of methods in this evolutionary frame as a messy methodological pluralism. This means that there is an alternative to characterising medical decision-making by one big method, by adopting Solomon's pluralism.

The question then remains which position to assume. Without going through too much trouble, the revolutionary model can easily be discarded as inadequate. First of all because some methods formally include other methods; such as the way in which causal-mechanistic reasoning, case-based reasoning and expert opinion and consensus are included into the decision-making processes of EBM. How hermeneutical medicine and shared decision making seem to be inherent parts of narrative medicine, or how medical decision-making method plays a role within personalised medicine. Second, because methods, even if they are not integrated into each other, may still very well exist next to each other. Within the same countries, same institutions, even employed by the same clinician. There is no reason why a clinician could not employ methods of narrative medicine to develop an intimate relation with the patient and gain narrative knowledge of the patients values and needs, to subsequently formulate and compare best possible options for treatment in accordance with these values and needs on the basis of the best available evidence, and to ultimately ask an authoritative colleague and make an appeal to his or hers intuitions with regard to these options and, in accordance with his patient, base his ultimate decision on it.

This second argument against the 'tidy methodological singularity' of the revolutionary model is at the same time an argument for the messy evolutionary methodological pluralism, given that combinations of methods, as demonstrated above, can be made in countless different configurations and with many more methods included.

However, as useless as the revolutionary model is for reasons of its inadequacy to describe clinical decision-making, is the messy methodological pluralism difficult to use in the context of a concise analysis of issues with regard to compatibility. This difficulty is caused by the complexity of

analysing and making comparisons with the totality of all methods in all possible combinations that the 'messiness' entails.

However, perhaps more importantly, representing processes of decision-making and knowledge production as a messy methodological pluralism seemingly overestimates the actual diversity and independence of the individual methods. The messiness in Solomon's account implies all methods to be completely unique and independently existing in the hybrid of methods. As has already been mentioned, some methods can be seen as complementary to each other, to integrate each other, or at least to be constituted of rather similar elements. In other words, Solomon shows there is diversity among methods, but in doing so obscures the fact that there is structure and similarity between them.

So for both it being difficult to work with Solomon's messy pluralism, as well as for it being misleading to perceive the pluralism as fully diverse and independent, some structuring should take place on the methodological pluralism. In other words, Solomon's should be tidied up slightly in order to make it more useful and more accurately depicting the internal structure. It should be pluralistically enough to do justice to the complexities of actual current medical practice, but tidy enough to enable analysis and reflect the relatedness of methods.

Tidy Methodological Pluralism

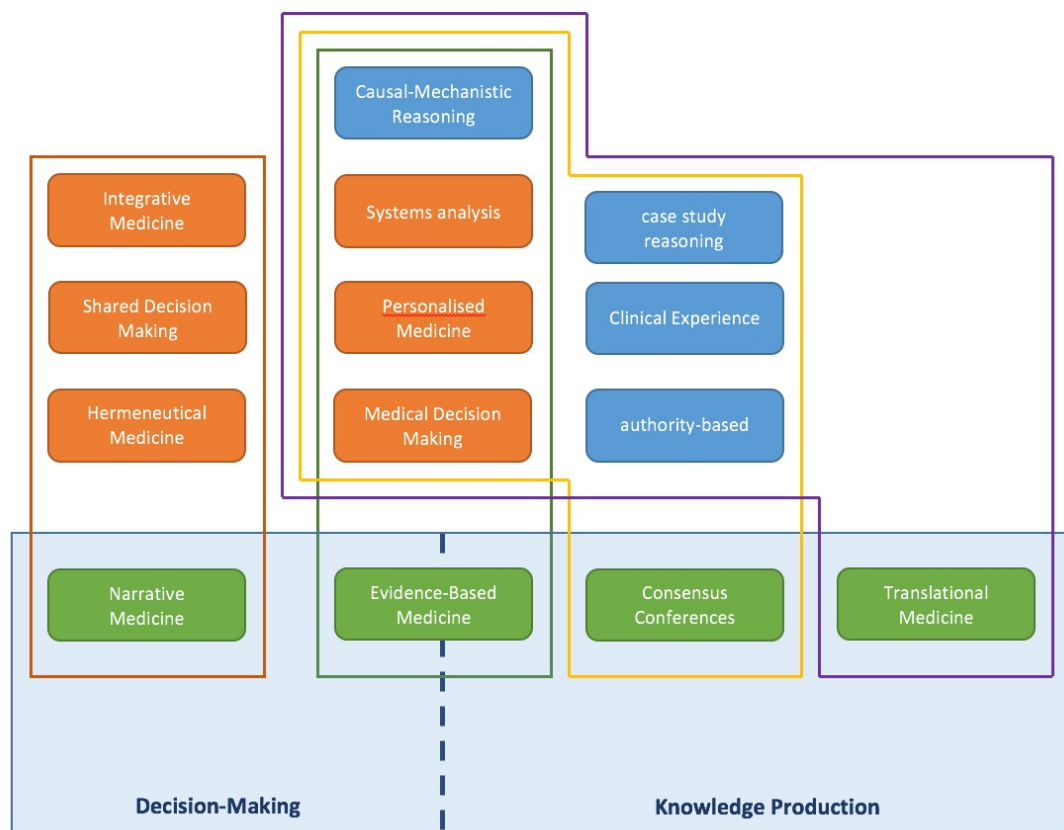


Figure 1 - Schematic representation of tidy methodological pluralism

The question then is of course how methods relate to each other, and which methods are more prominent than others. In order to come to an answer to this question Solomon's characterisation of the 'current methods' will be used as a starting point for the prominent methods. It may seem that we are blindly adopting Solomon's interpretation that these are the most prominent methods constituting current medicine, and though this thesis has certainly come to the conclusion that these current methods do indeed provide an adequate overview of current medicine, it is not without reason. This reason has been found in the fact that all of the traditional methods (In Figure 1 depicted in blue) and future methods (In Figure 1 depicted in orange) can be roughly subsumed under the current methods (In Figure 1 depicted in green).

Current medical practice has, in line with how the distinction has been made throughout this thesis, has been divided into processes of decision-making and processes of knowledge production. It has subsequently been identified that some methods are of more use to one of both processes. To be specific, it has been identified that narrative medicine plays a seemingly exclusive role in the decision-making process rather than the knowledge production process. One may argue that narrative medicine produces knowledge of a patient's 'broad' needs, and that it would therewith also fall within the category of knowledge production. With the category of knowledge-production, however, a more traditional understanding of knowledge production is followed, which rather refers to the production of more generalisable and verbalised knowledge. Intuitively one could argue that it is more 'scientific' in nature (though this undoubtedly opens up a whole new discussion). In the same light, consensus conferences and translational medicine have been placed under knowledge production, as their mutual goal is to distil this type of broadly applicable and explicit knowledge. It could be argued that consensus conferences could be organised as part of an individual's clinical decision, but that would be highly unique, and perhaps even disturbing the relation between clinician and patient (if it occurred ever at all). More seriously one could argue that translational medicine's knowledge producing processes would often take place in direct relation to individual patient care, therewith inevitably being involved in decision-making. However, the purpose of translational medicine is not decision-making, it is instead aimed at generating knowledge from the application to patient care, arguably legitimising translational medicine exclusively being placed under knowledge production.

With regard to EBM this distinction could not be—and therewith should not be—made as hard. EBM plays an important methodological role in the decision-making process, requiring clinicians to adhere to certain standards with regard to the evaluation of available evidence regarding clinical cases. At the same time EBM provides standards and an influential methodological framework for medical scientific research, making it just as much a method of knowledge production. EBM has subsequently been placed on the interface between both processes, serving both knowledge production and decision making.

Please note how the different traditional and future methods are placed under the four current methods. Integrative medicine, shared decision making and hermeneutical medicine are all placed under narrative medicine, as all four of them extend the domain of relevant clinical knowledge to include more socially holistically oriented aspects of their needs. They are not subsumed under EBM, consensus conferences or translational medicine as they are all very much focussed on the individual patient, being much less relevant to knowledge production.

Mechanistic reasoning, systems analysis, personalised medicine and medical decision making (method) are all subsumed under both EBM, consensus conferences and translational medicine. They

first of all fall under EBM as they are all types of knowledge production that feed into the continually influential standards of EBM. In other words, they all have the ability to produce evidence strong enough to be meaningful within the evidential hierarchy that EBM maintains.

Similarly these methods fall into consensus conferences, as all of them may be used in the epistemic rituals by which consensus conferences produce knowledge. It should be unsurprising as consensus conferences may also be used to resolve scientific controversy within EBM; it should then be able to consider its forms of evidence. Additionally it is important to note that consensus conferences, as opposed to EBM, may reasonably integrate case study reasoning, clinical experience and authority-based reasoning into its knowledge producing processes. These forms of knowledge production do not produce the kind of evidence that ranks high enough into EBM to be reasonably considered, but are—following Solomon's insight—often meaningful within the social processes that drive and characterise consensus conference.

Translational medicine furthermore integrates the same methods as consensus conferences do. The purpose of translational medicine, to bring scientific laboratory knowledge into clinical practice, facilitates EBM's methods to be subsumed under translational medicine, as a lot of the (early) scientific knowledge translational medicine aims to implement is produced under EBM. As the knowledge translational medicine aims to implement is not always fully verified or completely refined, it requires clinicians to be able to work outside the rigid methods that EBM subsumes, but also draw from case study reasoning, clinical experience and authority based reasoning in order to deal with the scientific uncertainty that translational processes involve.

The representation above does not imply that all methods are of equal significance to current medical practice. Translational medicine is obviously applicable to a narrower selection of ailments than EBM is, as not every treatment requires uncertain solutions from the fundamental sciences. Furthermore, though narrative medicine may be growing with time, it is identified by Solomon as much less prominent than EBM and translational medicine. Similarly Solomon identified that EBM is currently a much more dominant paradigm than consensus conferences in knowledge production.

This means that the characterisation assumes a dominant role for EBM in both the domains of decision-making and knowledge-production. Translational medicine and narrative medicine add to EBM in the domain of decision-making in, respectively, the practices of implementing new knowledge and shared decision-making. In the domain of knowledge-production consensus conferences, according to Solomon, add to EBM by dealing with scientific controversy and uncertainty, and translational medicine increasingly becomes a platform for experimental treatment; which EBM cannot always provide.

Reflections on current medicine

In this chapter the concept of medical decision-making and knowledge production have been introduced as interesting places of epistemological activity within the clinical environment. It has been shown how clinical decision-making can be characterised using versions of the art-vs-science distinction; but that the start of a more complete understanding could be derived from assuming parts

of Solomon's messy methodological pluralism. Solomon's pluralism should be understood as an evolutionary representation of developments in medical decision-making and knowledge production in which methods are added to the collection of existing methods rather than replacing them by revolutions.

It has furthermore been shown that knowledge plays a crucial role in decision-making, in the sense that decisions always rely on some sort of knowledge. The kind of knowledge differs greatly among the different methods of decision-making, but it is evident that knowledge and its production are a necessary condition in the decision-making process. This implies that it is necessary to address knowledge production as part of the ultimate analysis of the compatibility of ICDSS and clinical decision-making.

With regard to Solomon's messy pluralism an extensive overview has been provided with descriptions of the of available methods. It has subsequently been argued that Solomon's pluralism is unnecessarily too messy to come to a characterisation of medical decision-making that can be meaningfully used in service of assessing the epistemological compatibility of ICDSS with current medical practice. Instead a tidy methodological pluralism has been proposed in which Solomon's 'current' methods are put in the centre of the characterisation of clinical decision-making and knowledge production under current medicine. Though this perhaps limits the prominence of the 'new' and 'traditional' methods in the analysis, they will remain relevant regardless. On the one hand, some are integrated into current methods. On the other hand, the 'new' and 'traditional' methods may continue to function as frames of reference by which to verify or reject preliminary outcomes of the analysis.

CHAPTER II

Intelligent Clinical Decision Support Systems

The previous chapter of this thesis introduced Solomon's pluralism on a conceptual level. It explained the principles on which the methods of decision-making were based; but it went into little detail with regard to how they are being applied in the clinical practice. One significant aspect of clinical practice that has been mostly omitted is the role that technologies play within the different methods of decision making.

This chapter will not aim to address the role of technology in the broad sense for all of the presented methods; that would go far beyond the scope of this project and the available time and space. Instead it will focus on one specific area of technological development, that of clinical decision support systems. Briefly put, this concerns computer science and artificial intelligence with the goal of assisting in the decision-making process and the production of knowledge. As will be shown in this chapter, there is quite some history to these systems and an equally long list of promises and expectations for the future. What there is currently not, however, is a fully operational technology that qualifies as intelligent clinical decision-support. This poses a challenge given the aims of this thesis to evaluate the compatibility of ICDSS with current clinical decision-making. This shortcoming can be overcome, in a manner analogous to the previous chapter, by constructing a characterisation of what could be qualified as ICDSS. Drawing from academic literature and more public statements, this chapter aims at constructing such a characterisation by analysing the promises and expectations that are put forward by influential opinions in the field, and by evaluating their likeliness on the basis of critical reflection on data analysis methods and computational technologies.

Ultimately this should lead to a characterisation of what ICDSS may turn out to require and to deliver. It will be aimed at making it broad enough for it to overcome the inherent uncertainty of anticipation, and to facilitate the hypothetical implementation in a wide spectrum of biomedical technological environments within the characterisation of current medicine—while at the same time being specific enough for it to be adequately substantive and meaningful about the epistemological characteristics and implications of ICDSS.

In pursuit of this goal the chapter will first go into the question of how ICDSS should be understood and what its history entails; what promises and expectations are being made and raised, and how ICDSS both enables decision-support and knowledge production. The chapter shall subsequently go into the underlying technologies and methodologies that enable ICDSS's capacities, and evaluate and characterise their epistemological properties. This should enable the assessment of what ICDSS may realistically be capable of, and enable epistemological issues to be identified for later

analysis. On the basis of these outcomes, this thesis aims to formulate a characterisation of ICDSS in the context of clinical decision-making and knowledge production.

What Is ICDSS

Three types of CDSS

Answering the question of what intelligent clinical decision support systems (ICDSS) are, begins with an answer to what clinical decision support systems (CDSS) in general are. CDSS, as has been briefly mentioned, have a substantive history in medicine. A definition to CDSS is provided by Mark Musen, a distinguished scholar in the field of bio-informatics, and director of the Stanford Center for Biomedical Informatics Research. He defines CDSS as: "*A clinical decision-support system is any computer program designed to help healthcare professionals to make clinical decisions.*" (2006, p. 700). Several other definitions of CDSS can be found in academic literature, but they are all ultimately very similar to Musen's definition (Miller, 1991; Sim et al., 2001)

The first thing that stands out, is that Musen emphasises the immaterial aspects of the technology by focussing on the programme; rather on the device on which it runs. The second thing that is clear is that Musen provides an incredibly broad definition that includes a truly vast selection of imaginable and actual forms of data processing that takes place in a clinical environment. To create some structure in this overwhelmingly large category of systems; Musen proposes to identify three different types of functions CDSS could perform as a tool. As a tool for: information management, focussing attention, providing patient-specific recommendations (2006, p. 701).

As a tool for 'information management', CDSS functions as an environment for storing and retrieving clinical data. The systems helps the clinician with organising and providing data and knowledge, but does not decide itself what information should be presented to the clinician, nor does it draw conclusions on the basis of it. Systems that aim at 'focussing attention', are designed to alert the clinician of abnormalities in data or suggest possible diagnoses that may commonly be overlooked. Musen states that these systems are often based on simple logic, applying simple pre-programmed rules (2006, p. 700). As an example Musen provides how clinical laboratories use software to alert laboratory analysts of known drug interactions (2006, p. 701). 'Providing patient-specific recommendations' concerns the formulation of custom-tailored advice on the basis of patient-specific data. This may be advice about diagnoses, treatment options or suggestions for what additional information could help the diagnostic process (2006, p. 701). This thesis project aims at a concept of ICDSS that, though it may incorporate all of these functions, ultimately focusses on the role of providing patient-specific information.

A bit of history

Before going further into what it means for clinical decision support systems to be intelligent in providing patient-specific recommendations, it would be useful to establish a more complete understanding of what CDSS is and how it has come to be, guided by a few historical examples. CDSS as has been briefly mentioned in the introduction to this thesis, has been around for quite a while, perhaps longer than one might come to expect given the suggested dependence of computers.

Though this thesis will follow Musen's definition, equating CDSS to computer programmes, the idea of developing heuristics to aid medical decision-making also existed before outside the scope of computer programmes. One interesting example is the 'diagnostic slide rule', or 'logoscope' by Dr. Firmin Nash; a pre-computer aid to draw diagnostic conclusions. The ruler had a long list of possible conditions printed along the side, and different sticks for the presented symptoms could be inserted into the ruler. The sticks featured lines at the height of the conditions the symptoms corresponded with. Lines lining up with each other and a condition then indicate that the presented symptoms could be explained by that condition. The idea was that the more lines aligned the more reasonably it would be assume the condition to be the reason (Nash, 1954). Nash had tried to emulate the decision-making that clinicians did according to his account by mechanising the logic they applied to determine conditions into the ruler. As should be clear by now; medical decision-making involves much more than mere logical reasoning on the basis of the incidence of symptoms, and it should come as no surprise that the device was rather unsuccessful in its aims. Though it did indeed provide a convenient summary of 'book knowledge' (Nash, 1960) and the developed product made it to the market, it failed to be fully effective in clinical practice. One instance of a critical review by R.W. Pain showed it was only of any use in just 23% of presented cases (Pain, 1975).



Figure 2 - Image of Nash Logoscope indicating tuberculosis as plausible explanation for the symptoms of swelling of a bone, hiccup and anaemia - Image by Science Museum London

Though Nash's solution did not really work out, his quest to understand decision-making and produce universal aids for choosing from different possible explanations outlasted the Logoscope.

Towards the end of the 1950's R.S Ledley and L.B Lusted, saw the potential for computer science to pick up the glove, and were among the first to describe in a paper how "*With the use of computers several mathematical techniques can be applied to certain aspects of medical diagnosis*" (1959, p. 1970). Though the authors identify that '*much work remains to be accomplished*' for computers to actually be able to contribute, they feel confident enough to put an extensive list forward of functions that a computer should be able to perform.

1. *Produce a list of possible diagnoses for a hospital case by analysing the symptoms' presented with respect to data characteristic of certain diseases.*
2. *Indicate further diagnostic tests which best differentiate between remaining disease possibilities.*
3. *Calculate probabilities for the alternate diagnostic possibilities.*
4. *Aid in an analysis, based on hospital case data, of value decisions which lead to treatment planning.*
5. *Since the above functions must be based on extensive medical data, it is possible that computers could simultaneously compile statistics that relate symptom combinations to disease states, and treatment to prognosis*
6. *In addition, computers can aid data recording and analysis of certain diagnostic procedures. For example. Tabulate quantitative criteria derived from electrocardiograms and electroencephalogram, and perform calculations based on such data.*
7. *Retrieve current information relative to the above functions.*
8. *Record and recall desired aspects of a particular patient's total medical record (such as total radiation dosage received, previous allergic reactions, individual biochemical and physiologic norms and deviations, etc.) which might be useful in a current evaluation of the patient's status.*

(Ledley & Lusted, 1959)

Though some of the functions computers could have according to Ledley and Lusted are rather technical in nature, the collection of these eight functions reveals that many of the features they put forward are very similar to the three types of decision support that Musen identified. Not only that, but an early interpretation of a key feature of ICDSS that has been referred to throughout this thesis project, can be read in Ledley and Lusted's fifth and sixth function: the ability of the programme to learn from raw input data and produce knowledge.

Much of the rest of the article focusses on the actual mathematical methods that could be used by computers to perform the task of making differential diagnoses. A differential diagnostic process involves the diagnostic choice between several different conditions that all explain or align with the presented symptoms. One method that Ledley and Lusted propose to use and present is Bayesian statistics; a statistical method that provides means to compare between the different conditional probabilities for the different possible differential diagnoses. In other words, it provides the probability of *condition A, given the presented symptoms; condition B, given the presented symptoms*, etc. and subsequently allows the decision-maker to order the different explanations on the basis of likeliness. In a brief and simplified example: this method provides statistical support for a patient presenting just a runny nose having more likely caught a cold or suffering from hay fever than having contracted HIV. In actual differential diagnoses this may be a lot more complex given a much wider spectrum of present and absent symptoms and possible conditions.

Bayesian statistical logic became the foundation for many attempts to build computerised CDSS. One of these attempts was developed at the *department of surgery and computational science and the electronic computing laboratory* at the University of Leeds. The Leeds Abdominal Pain System (LAPS) was developed towards the end of the 1960s as a means to aid in the differential diagnosis of seven possible diagnoses of acute abdominal pain in the emergency room environment (de Dombal, Leaper, Staniland, McCann, & Horrocks, 1972). Developing LAPS, like any other Bayesian-based system, required large amounts of high quality data about the prevalence of symptoms and seven different conditions related to abdominal pain (Adams et al., 1986). And though simplifications had to be made for reasons of computability, the system turned out to be very effective in correctly diagnosing patients. In a 1972 study by De Dombal et al. on more than 300 cases of acute abdominal pain found that LAPS accurately diagnosed patients in 91.8% of all the presented case; as opposed to an accuracy between 65 to 80% for clinicians depending on their amount of expertise (de Dombal et al., 1972). These results were so convincing that LAPS was exported outside the UK with the emergence of personal computing; only never to produce results as accurate as they had been in Leeds.

Bayesian statistics, however, was (and of course still is) not the only possible foundation for computers to assist the diagnostic process. An example of an alternative is the MYCIN programme. A rule-based system designed to support with the management (therapy) of infectious disease rather than with diagnosis (Shortliffe et al., 1975). Without going into too much detail, the programme allowed for, and relied on, conditional statements to be put into the programme. Conditional statements, or rules, "indicate what conclusions can be reached or actions taken if a specified set of conditions is found to be true" (Musen et al., 2006, p. 704). In other words, they are if-then statements; with the possibility of having several ifs and thens in one rule. When the user asks a question to the MYCIN programme, it needs to determine which rules apply and how they relate to each other. The idea of relying on input from clinicians was to make use of the expert knowledge that they possessed and make it available to the wider community. Though research suggested MYCINs advice was slightly more effective than human clinicians, MYCIN never became widely used in clinical practice (Victor et al., 1984; Yu et al., 1979).

Another significant programme that combined the role of 'information management' with 'focussing attention' was the HELP system, a system that continues to be of great influence (Kuperman, Maack, Bauer, & Gardner, 1991; Musen et al., 2006). The HELP system was a programme that was developed to perform four functions: To accommodate an ever-expanding medical database of clinical records, support decision logic to facilitate diagnostic, therapeutic and alarm functions, serve medical and clinical needs of the hospital, and provide effective research sub-systems to enable clinical research to be performed on large datasets (Pryor, Gardner, Clayton, & Warner, 1983). As opposed to the examples that have just been discussed, or the many similar CDSSs of the fifties and sixties, the HELP system did not aim at providing a single-focus heuristic for making domain-specific decisions on the explicit request of the clinician. Instead it aimed at a centralised system; containing and integrating the records of all patients with knowledge and information of decision-making, diagnoses, therapies and standards and deviations. Also of interest to this thesis, not only providing decision-support, but also tools for monitoring of patients, and for knowledge production (Musen et al., 2006, p. 705).

Intelligent CDSS?

Now that we have come to explain how CDSS is to be understood, and have complemented that with some examples, the question of what it means for CDSS to be intelligent emerges. Given that this is a thesis project in philosophy of science, technology and society one might expect descriptions and reflections on the philosophical discourse on what intelligence is, or whether it can be possessed by machines. Though these questions are very interesting indeed, it is not the point of this project to come to a philosophically justifiable account of what intelligence means in the context of CDSS. This project, instead, uses the word intelligence as an adjective indicating that it concerns the evolution of CDSS. In this regard, word intelligent could just as well have been substituted by another word that similarly expresses the current trend of making known and existing objects and concepts, such as CDSS, 'smart'. Therefor the question is not what may or may not be intelligence from a philosophical point of view. The question is what is added in terms of functionality to CDSS that makes it ICDSS, more or less regardless of whether it should indeed be classified as intelligent from an analytical point of view.

The short answer to this questions of what intelligence in the context of this, would involve some statement about the implementation of artificial intelligence (AI). Though this does indeed address part of what it means for CDSS to be intelligent within the context of this thesis, it does not provide enough insight for it to be used in the characterisation of ICDSS. This is first of all due to the ambiguity of the term artificial intelligence. Apart from referring to the abstract principle of, indeed, artificial intelligence, it has come to refer to a form of technology instead. In other words, what was once referred to as AI may now just be considered a simple programme. Merely characterising ICDSS by the implementation of AI furthermore obscures precisely what functions it is supposed to have and on which methods the system relies to satisfies these functions.

In order to actually come to an understanding of what these functions may entail, and what methods ICDSS may employ, this section shall aim to gather an understanding from a diverse selection of sources that express expectations, visions and arguments regarding this. These materials will be drawn from a variety of sources, promotional texts, scientific publications, interviews and opinion pieces.

There is no opportunity in the context of this thesis project to perform a complete discourse analysis, or exhaustive analysis of the expectations. That does however not mean that an exploration of expectations and visions of ICDSS will produce meaningless results. The sources shall be selected on their embeddedness into the scientific and industrial discourse; warranted by using relatively highly cited academic articles, articles making claims of which similar claims are conventionally made in academic literature, and projects that attract significant public and corporate attention. Again not to produce an exhaustive analysis, but to gain meaningful and reliable insight into the visions, expectations and perhaps trends that are present and steer developments; ultimately to come to a sensible characterisation of what ICDSS may turn out to be.

Promises and Expectations

IBM Watson

Visions of what ICDSSs are, turn out to often be formulated in the context of research projects. There is a large number of research project aiming at some form of intelligent or autonomous functions in the clinical context, and many of them come with, or can be interpreted as, visions of the future. One particularly prominent project of developing artificial intelligence, is a programme named Watson. Watson is a computer programme developed by IBM for the specific purpose of natural language processing and Question Answering; in other words, to be capable of answering questions asked in natural language, such as English, rather than abstract codes or programming languages (Ferrucci et al., 2009, 2010). Watson became known for outperforming 'champion' human competitors in a well-known television quiz show, Jeopardy—and continues to be very present in the popular discourse on the future of AI with regular appearances in the news media. In playing Jeopardy Watson dealt adequately with, indeed, questions asked in English, it dealt with the extraction of knowledge from written sources (such as Wikipedia), gathering textual evidence for possible answers, identifying implicit relations between knowledge, and with inferences needed to come to accurate answers to these questions (Chu-Carroll, Brown, Lally, & Murdock, 2012; Fan, Kalyanpur, Gondek, & Ferrucci, 2012; Kalyanpur et al., 2012; Lally & Fodor, 2011; Murdock, Fan, Lally, Shima, & Boguraev, 2012).

Soon after demonstrating Watson's ability to accurately process knowledge and evidence in answering trivia questions, IBM announced that Watson³ would be developed and trained to go beyond Jeopardy and be of use in the field of medicine (Simonite, 2011). A logical step, given that this thesis has shown that decision-making is a process that at least partially involves answering questions on the basis of readily available knowledge. Especially viewed in the context of EBM, which the previous chapter identified as still being a very influential method of decision-making, this step makes sense given the clear integrability of clinical evidence into the operations of Watson.

In an interview with computer world, Dr. Eliot Siegel, a professor of the University of Maryland's School of Medicine—IBM's early partner in developing 'Dr. Watson'—paints a picture of what these capabilities should be able to do in medical practice.

"There is a major challenge in medicine today. There's an incredible amount of information in a patient's medical record. It's in the form of abbreviations and short text. There's a tremendous amount of redundancy, and a lot of it is written in a free-form fashion like a blog or text. As a physician or radiologist, it might take me 10 or 20 or 60 minutes or more just to understand what's in a patient's medical record,"
"If all Dr. Watson did was allow me to organize electronic medical records and bring to my attention what's most important and summarize it, that would be incredibly valuable to me."

(Siegel 2011, as cyted by Gaudin, 2011)

Employing Watson's abilities merely to structure the heterogeneity of knowledge that may be contained in localised electronic health records, is actually still a rather modest picture of what can be expected from Watson's capabilities compared to other people in the field. For example Jennifer Chu-Carroll, one of the lead researchers on the Watson project, and arguably a more authoritative expert on the technical capabilities of Watson, argues:

³ It is important at this point to identify that 'Watson' is not as singular as IBM's marketing prefers to refer to it. Watson should be understood as a collection of methods and technologies, that can be tuned, trained and developed for particular purposes; but ultimately being different programmes for different applications. In other words, the Watson playing trivia quizzes is not exactly the same as the medical Watson.

"Think of some version of Watson being a physician's assistant, in its spare time, Watson can read all the latest medical journals and get updated. Then it can go with the doctor into exam rooms and listen in as patients tell doctors about their symptoms. It can start coming up with hypotheses about what ails the patient." "The physician will make the decisions, but Watson can help."

(Chu-Carroll, as cited by Gaudin, 2011)

This is a picture in which Watson not only has knowledge of the patient records, it absorbs the full body of medical literature that is available at any point in time and provides the ability to understand it and relate it to the clinical circumstances and needs of the individual patient and provide decision support in the form of suggestions on diagnoses and possible treatments. It is interesting to note how Chu-Carroll adds that the physician will remain in charge of the actual decision.

The statements above should be put in a little context by emphasising that they were made in the early stages of Watson's integration into medicine, and that they describe very broad visions of what the technology in general terms could become to be. However, judging from the five projects that IBM is currently visibly working on with Watson in medicine, it seems that though these projects are considerably more modest in scale, many of these ambitions are visible in the actual projects. This is very interesting in coming to an understanding of ICDSS, coming to a useful understanding of what Watson can mean for ICDSS.

To describe a few, mainly focussing on their functionalities, without going too deeply into the technical details, 'Watson for Genomics' for example aims to provide clinicians with advice on cancer treatment on the basis of the outcomes of genomic tumour analysis; a technique that provides genomic information of a particular tumour. In a promotional text, IBM states that. *"Watson employs advanced cognitive computing to automatically extract and analyze relevant and validated data from established guidelines, medical texts, and clinical trials to help enhance your patient's results"* (IBM Watson Health, 2017b). Though this suggests that Watson only performs a very specific task and only answers one very specific question⁴, and it is oriented on a confined set of domain-specific guidelines and literature, it seemingly does feature autonomy in the interpretation of both clinical data and written informational sources. Watson for Drug Discovery, as the name suggests, is tuned to aid in the discovery of new therapeutic drugs, currently for example for ALS. IBM states that Watson helps to *"search research documents for potential connections and help identify new RNA-binding proteins linked to ALS"* (IBM Watson Health, 2016). In other words, it is quite similar to what Watson Genomics is currently doing; though not in a clinical setting; but in a research setting instead.

'Watson for Oncology' is a project that is increasingly becoming available for hospitals to be actually used clinically, and is described in a bit more detail in a video released by IBM as part of the product page. The video showcases *"Watson's unique capability to analyze a patient's medical record to help identify for the clinician evidence-based and personalized treatment options"* (IBM Watson Health, 2017a). It shows how Watson takes the form of a clinical archive of patient records to which the clinician logs on and sees a list of patients for the day. Upon selecting a patient *"Watson analyzes relevant portions of her electronic medical record including her family history, notes from prior office visits, and test results"* (IBM Watson Health, 2017a). This means that the programme summarises the patient's entire record and presents it as a consumable amount of structured and

⁴ For example: 'What would be the most effective treatments for patient X based on the genetic characteristics of the patient's tumour'

relevant information, and shows the information that is missing and could be of relevance. Upon request, by the click of a button, *"Watson analyzes the case information, identifies a prioritized list of treatment options... and provides links to supporting evidence"* (IBM Watson Health, 2017a). The video continues to explain that Watson draws from *"curated literature and rationales from leading oncologists, as well as from 300 medical journals, over 200 textbooks and almost 15 million pages of text"* (IBM Watson Health, 2017a). The 'prioritised' list of treatment options is presented in such a way that the clinician can see on what sections from what sources Watson bases its suggestion and that the clinician can provide feedback to the sources used. Watson for oncology ultimately provides the possibility to see the compatibility of a patient with certain medications, and allow the clinician to share the treatment plan directly with the patient.

Reiterating what can be learnt from IBM's description of Watson for Oncology, it is interesting to see how IBM sees Watson both in the context of evidence-based decision-making as well as making it personalised. It also shows how Watson is indeed integrated with the electronic patient records, like the MYCIN programme. What is furthermore interesting to learn is that the system aims to communicate to the clinician about the foundations for the suggestions and that it allows the clinician to provide feedback on this.

DeepMind

IBM is not the only leading high-performance computing company that tries to make its way into developing systems for intelligent clinical decision support. Quite recently a few different collaborations between Google Deepmind and health care institutions were announced that should bring the power of artificial intelligence into clinical practice (Baraniuk, 2016; Hern, 2016; Rigg, 2016; Stevens, 2017; Turk, 2016). Deepmind, being a Google subsidiary since 2014, specialises in artificial intelligence and has especially employed machine learning in solving complex tasks before (Shead, 2016). Though this section will inevitably show some similarities between both Watson's and DeepMind's projects, they will be shown to be somewhat different in focus, goals and approach. One difference is that Watson is already a platform-like product that businesses and hospitals can make use of, whereas DeepMind is much more experimental. However, just like Watson, DeepMind is not merely being employed in the field of healthcare, and analogously DeepMind concerns several distinct projects within their healthcare related activities. In order to come to an understanding of how ICDSS is being interpreted in the context of these projects, and to draw input for characterising ICDSS from that, we shall again describe and inspect some examples of projects.

The most advanced and prominent example of such a project is DeepMind Streams. A name that closely relates to a system that tasks itself with *"Streaming the right information to the right clinician at the right time"* (DeepMind Health, 2017). The idea behind it is that patients suffer or die because clinicians *"don't have real-time information about who urgently needs their care"* (DeepMind Health, 2017).

In practice, they envision, *"every patient receives the right care from the right clinician at the right time, made possible through cutting-edge mobile technology that pushes patient alerts to nurses and doctors, enables them to securely assign and communicate about clinical tasks, and gives them all the information they need to make the right diagnoses and decisions."* (DeepMind Health, 2017). Like the way Watson for oncology, Streams takes the shape of a MYCIN-like platform;

integrating the patients' medical history and logistical software with providing its autonomous functions—in this case providing clinicians push-notifications of alerts in a mobile app.

As is the case with many of these projects, this is a vision that the current width at which it is being experimented with. The Streams project is currently being applied to acute kidney injury, a potentially life-threatening condition that requires early diagnosis and immediate care for optimal clinical outcomes. However, the condition often occurs as a consequence of other conditions (for which patients are often hospitalised) and its causes remain unknown; making prevention challenging (Bellomo, Kellum, & Ronco, 2012). It is furthermore difficult to make a direct diagnosis on the basis of one simple test, it requires the clinician to make inferences of the presence of injury; though often not with optimal outcomes (Kellum, 2008).

However surprisingly, this support is currently being provided without the use of, what DeepMind itself refers to as, AI. Instead, they are performed by pre-programmed rules on the basis of an established body of knowledge and only serve to bring test results to the clinician as soon as they are available. But the point of discussing Streams is not to describe a 'mere' logistical innovation. DeepMind is a company that specialises in a method referred they refer to as machine learning, and unsurprisingly their ambitions in the Streams project are about implementing these methods; hence the reason for reflecting on it.

About machine learning, a general statement of DeepMind's involvement in healthcare shows that DeepMind envisions a role for itself in improving clinical decision-making as well as medical research: *"We think that machine learning technology, a type of artificial intelligence, can bring huge benefits to medical research. By using this technology to analyse medical data, we want to find ways to improve how illnesses are diagnosed and treated."* (DeepMind Health, 2017). Given DeepMind's focus on data science, this statement most plausibly refers to a system that learns from these streams of data and identifies relations between certain developments and parameters and clinically relevant events. This benefits medical practice both in service of making diagnoses of conditions that would have gone unnoticed with current diagnostic methods, as well as in service of creating new data-based knowledge of clinical conditions⁵.

In the context of the Streams programme, this may mean that it continuously analyses incoming data from patient monitoring devices and the electronic record and searches for the detection of known and unknown patterns; for example between sets certain lab-results in combination with fluctuations in blood pressure and heart-rate over time and the risk of acute kidney injury.

So, precisely what AI would mean for Streams remains unknown at this stage, but if DeepMind's general statement provides an accurate means to understand its direction, there are a few things about Streams that stand out. The first thing is how Streams and Watson differ currently see different devices as the locus of the ICDSS. Watson uses a web-page running on a conventional computer, and Streams a mobile device such as a phone or a tablet. However, a much more important difference, can be found in their relation to knowledge. Watson had a strong focus on knowledge retrieval from written sources, both in giving diagnostic advice as well as research-oriented advice. Streams' ultimate vision, however, seems to find knowledge in the patterns (that are assumed) to occur within the

⁵ For a reader without any comprehension of machine learning (and related methods) this may seem like an unfounded inference. Please refer to the next chapter for a general introduction of these methods.

patient-specific flow of continuously generated data. On the one hand relying on actionable correlations and patterns, and on the other hand continuously searching for new correlations.

Scientific Literature

Apart from Streams and Watson, that are both being presented by their owners and promoters as cutting edge and future-proof, there is a field of academic research that concerns itself with ICDSS. In this section a look shall be taken into some of their influential expectations for the future of (I)CDSS.

Mark Musen, a leading authority on CDSS that has previously been of service to this thesis by now, dedicates a section of his book chapter to expected future developments in the field of CDSS. Musen first of all identifies how the growing ubiquity of the internet will enable a convenient sharing of 'libraries' containing re-usable functional modules. This enables convenient integration of functionality into existing CDSSs, for example certain diagnostic functionalities for specific sub-domains (Musen et al., 2006, p. 733). This connects to the 'specialised' approaches that both IBM and DeepMind use, not developing general systems capable of everything in medical practice, but building their methods and software around particular features, such as cancer treatment planning, or the screening for acute kidney failure.

Musen furthermore forecasts that this convenient integrability of systems makes it unnecessary for ICDSS to choose one particular data analysis method. In other words, the collection of functions that the ICDSS features, may be Bayesian reasoning, mathematical models, pattern recognition or whatever method. What is, however, certain according to Musen, is that *"automated decision support will take place with every practitioner's routine access to clinical data in a manner that is unobtrusive, transparent, and tailored to the specific patient situation"* (2006, p. 733). According to Musen the deep integration of intelligent assistants into clinical practice will lead the concept of decision-support system to 'fade away'. This prediction is interesting in two ways; first of all it confirms the relevance of this thesis project. More importantly however, it again focusses our attention of the personalisation of care and diagnosis as a crucial aspect of ICDSS; something of which both Watson and Streams provide examples of.

Jaspers, another well-cited leading CDSS scholar and professor of bio informatics at the Amsterdam Medical Centre, agrees with Musen's expectation that CDSS may come to play a significant role in all aspects of medical practice. However, in the concluding remarks of a literature review of the efficacy of CDSS in clinical practice, Jaspers expresses some important challenges for building CDSSs in general that may actually work, challenges that apply to ICDSS equally well (Jaspers, Smeulers, Vermeulen, & Peute, 2011). Jaspers comes to these challenges after concluding what is needed for CDSSs to gain the efficacy her study identifies it lacks. She identifies that for improving its efficacy, the specificity and sensitivity of the advice CDSSs produce need to be increased (2011, p. 333), meaning that both the portion of correctly identified positives as well as correctly identified negatives need to be increased. Jaspers furthermore argues that the need for manual input by clinicians into the system needs to be taken away, and that the advice the CDSS produces is created at the time the clinician actually makes the decision, not interrupting the already established workflow (2011, p. 333). This leads Jaspers, as we have seen with other authors, to argue for a merging of electronic medical records with the CDSS. It furthermore leads her to draw the following conclusion:

This suggests that these systems should work in the background and continuously monitor and check whether the care (to be) delivered to individual patients is in accordance with applicable guidelines. The CDSS should then only deliver its advice in situations where clinicians do not follow these guideline recommendations or when unforeseen patient outcomes occur.

(Jaspers et al., 2011, p. 333)

Apart from clarifying that especially the ability to autonomously provide decision support requires continuous monitoring, it elucidates a very important aspect of the purpose of and role of ICDSS that Jaspers subscribes to. Namely, that CDSSs should be tasked with making sure that clinicians stick to the standards and guidelines, or are at least made aware when they choose to deviate. This is confirmed by the way in which she positions CDSS as a possible solution to the lacking adherence of clinicians to guidelines and recommendations within EBM (Jaspers et al., 2011, p. 327). Though this is seemingly a logical train of thought, it is not necessarily self-evident that CDSS should be used for the purpose of monitoring clinician performance, making it valuable insight into the envisioned role of EBM.

Another interesting issue regarding standardisation to which attention is raised, though of a more practical nature, is the lack of standardisation. Both in terms of 'commonly accepted' terminologies and ontologies currently in medicine, as well as in 'interfacing standards'; respectively making integration of the variety of differently standardised patient records and the integration of different ICDSS features across different systems a major challenge (2011, p. 333).

Reflecting on the views of Musen and Jaspers, it stands out that they both argue for a form of CDSS that is integrated into a system of patient record management. It is furthermore implied that they expect the intelligence of ICDSS to predominantly lie in individually artificially intelligent and autonomous functions that can be developed, implemented and exchanged across ICDSSs, rather than having one all-encompassing form of AI that is responsible for assisting the entire clinical process. Musen and Jaspers furthermore agree with each other that CDSS shall be integrated into patient monitoring and record keeping systems, though worth noting is that this is an expectation for Musen and a requirement for Jaspers.

Musen and Jaspers are of course not the only people involved into the academic discourse on the role and form of intelligent features of clinical decision support. As far as other authors in the field go into similar details, their views do however seem to be quite representative, and at least reconcilable, with the views of a wide selection of other authors in the business of investigating and analysing CDSS in clinical care (Jackson, Bolen, Brancati, Batts Turner, & Gary, 2006; Kaplan, 2001; Mack, Wheeler, & Embi, 2009; Mollon et al., 2009; Walton, Dovey, Harvey, & Freemantle, 1999)

With regard to these autonomous and Intelligent Sub-Systems that Jaspers and Musen propose, some recent examples shall be briefly inspected. The first of which is a so-called 'fuzzy rule-based' risk level prediction system for heart disease (Kunjunainair, 2011). Without going too deeply into the technologies underlying it at this point, it aims to find predictive correlations between certain phenomena within a patient's data and heart disease by employing a form of machine learning. The philosophy behind machine learning is that, when the machine is presented enough historical data about a large enough group of patients, including a (verified) diagnosis, it is capable of finding

'existing' correlations between different aspects of the patient data without requiring explicit instructions of where to find these correlations. Kunjunninair's work is just one well-cited example of this category of work, but many similar projects involving prediction in clinical practice exist (De, Biswas, & Roy, 2001; Esfandiari, Babavalian, Moghadam, & Tabar, 2014; Khatibi & Montazer, 2010; Palaniappan & Awang, 2008; Tsipouras et al., 2008; Warren, Beliakov, & Van Der Zwaag, 2000).

In Kunjunninair's research, the aim for the system is to 'learn' from a large dataset of patients. In his case, the data sets contained information on "*Age, sex, chest pain type, resting blood pressure, serum cholesterol in mg/dl, fasting blood sugar, resting electrocardiographic results, maximum heart rate achieved, exercise induced angina, ST depression, slope of the peak exercise ST segment, number of major vessels, thal and diagnosis of heart disease*" (2011, p. 30) The system then learns about the correlation between different aspects of their records and the corresponding diagnoses; so that the machine (after 'practice') ultimately gains the ability to infer a correct diagnoses for records without the diagnoses included. More specifically Kunjunninair's system aims to look at aspects of the record 'long' before the actual incidents and symptoms that led to the diagnosis in order to not just be able to infer a diagnosis on historical records, but to be able to predict heart disease in pre-symptomatic stages on the basis of their already existing data (2011).

To a large extent Kunjunninair's research looks a lot like any other machine learning project in making diagnoses on the basis of clinical data. What is interesting however, is that he aims his systems to be predictive as well. This is especially interesting because this strongly deviates from current medical practise. In most aspects of current conventional medical practice a patient gets to see a doctor when as he or she experiences problems. In medical terminology, apart from screening programmes and occasional periodical check-ups, the majority of patients only shows up while or after they experience the symptoms of underlying conditions. These symptoms are then in turn also being used to make clinical decisions, for example for making a diagnoses or for ordering additional tests and examinations. In Kununninair's predictive system, and other ones like it, however, allows for patients to be diagnosed without the conventional symptoms already being present. In turn causing an interesting problem in the verifications of the prediction. Though in some cases additional tests may be able to confirm or disconfirm a predictive diagnosis, there may be circumstances under which diagnosis cannot be checked by any of the currently available diagnostic tools; leading the grounds for the prediction and the machine's self-discovered correlations being the only verifiable aspects of the predictive diagnosis.

This 'turn' towards the predictive and preventative is fuelled by the conviction of data scientists that meaningful and actionable correlations will exist within collections of data that are too large for human cognitive abilities to adequately perceive. This is a position that this chapter will come to reflect on in further sections, but in this conviction it is generally speaking that the larger and more diverse the dataset the more likely meaningful correlations are. It should then come as no surprise that there is also research taking place aiming to make use of the ever-increasing amount of data that is being generated; also real-time, and also outside the scope of the clinic itself.

A recent review paper by Banaee, Ahmed & Loutfi (2013) about trends and challenges of 'data mining' in data from continuous health monitoring systems. Health monitoring systems are devices (often) that track certain parameters of patients, such as heart rate, blood oxygenation, blood pressure, electrocardiograms, blood glucose level, etc. in order for clinicians to monitor their patients outside of the hospital, possibly even real time. A domain of CDSS that lies outside of the clinic but within

medicine, that is highly interesting and relevant given the growing amount of health-oriented consumer wearable's, and the increasing capabilities of these devices. Banae et al. show that much of the research in this field focusses on three tasks: "Prediction, anomaly detection... and diagnosis" (2013, p. 17475). Anomaly detection is a self-explanatory concept, and requires a CDSS to have some sort of reference of what normal patterns of behaviour are. Prediction involves the ability to identify events that have not yet occurred, for example predicting blood-glucose levels in case of diabetics or predicting cardiac arrests on the basis of ECG data (2013, p. 17477). Making diagnoses can be seen as similar to anomaly detection, with the only difference being that diagnostics requires the system to have a reference of what healthy is and what not, rather than what normal is and what not (Banae et al., 2013, p. 17477).

What can be learnt from both Kunnunair's research and Banae et al.'s review is that future developments in the field of wearable technologies as well the in the field AI techniques such as data mining and machine learning, causes ICDSS to extend beyond traditional diagnosis (traditional as: on the basis of physical phenomena/symptoms rather than phenomena within data patterns) and beyond the clinic. As mentioned this is a significant addition to the image of intelligent features of CDSS we have drawn so far, given the growing (consumer) interest in self-tracking health devices.

Characterising the I.

The examples that chapter section has reviewed above, drawing from commercial parties, meta-academics and perspectives from executive researchers, provide perhaps just a limited overview of what intelligence may entail and how it would be implemented into CDSS. Regardless, it has provided useful insight into the forms, methods, functions and roles of future ICDSS. This allows for two things. First of all, to explicate an understanding of what intelligence means in the context of CDSS. Second, it allows for the construction of a characterisation of ICDSS, that shall be used in the following chapter for the analysis of the epistemological compatibility with the characterisation of current medical practice.

As was announced at the beginning of this section, the obvious answer to the question of what intelligence entails in the context of CDSS has something to do with artificial intelligence. To be more specific, the sources that have been reviewed suggest that it refers to a system having certain functions that are being performed autonomously on the basis of artificial intelligence. In that same introduction it was also declared that this knowledge is not the focus in characterising ICDSS. The truly relevant question is what this intelligence implies in terms of capabilities and deficiencies, something to which the analysis of the previous section provides interesting insight.

The first main element of ICDSS that surfaces from the analysis is that ICDSS will most likely be integrated into patient management systems. Apart from the benefits in terms of work-flow and the acceptability of the technology, as both stressed by Musen and Jaspers, it is also crucial for many of the intelligent capabilities of ICDSSs to have access to both historical and incoming patient data to perform many of the proposed real-time assistance functions.

This implicitly already partially addressed the second main technical specification for ICDSS, namely that the intelligent capabilities of ICDSSs are not part of one large general artificial intelligence, but are more likely modules instead. Apart from both Musen and Jaspers arguing for this on the basis of respectively possibility and desirability, this connects well to the overall impression that both IBM, DeepMind and the individual research projects revolve around making use of AI technologies to perform relatively 'small' and delineated tasks within the healthcare process. As the next chapter will identify, it is very likely that developing one general AI lies outside the possibilities in the near future; making modular systems the ideal platform for slowly expanding the capabilities of ICDSS systems.

Another interesting insight of a more technical nature is the fact that the variety of sources drawn from above, confirm the idea that an ICDSS fulfils two rather distinct roles. On the one hand applying existing knowledge—in case of making diagnoses or giving therapeutic suggestions as outlined by all of the sources—and on the other hand generating new knowledge on the basis the available and incoming data that is being generated, that can then later be used to draw diagnostic and therapeutic conclusions from.

The question of how ICDSSs arrive at knowledge is a very interesting aspect of their intelligent features, and a diversity of visions exist with regard to this. IBM Watson first of all proposed, and are arguably already using, a feature that allows knowledge extraction from written sources. This ideally enables the ICDSS to gather knowledge on the basis of analysing medical literature, both educational as well as scientific. DeepMind's projects do however seem to remain powered by explicit rules and guidelines that are implemented into intelligent features by incorporating medical experts into the developmental process. At the same time the analysis of DeepMind assumes an anticipated future in which this knowledge can as well be generated on the basis of the patient data, something of which the intentions and technological possibilities are confirmed by Banae et al. and Kununninair. This generation of knowledge on the basis of patient data can take many different forms, depending on the methods that are being used; something the next chapter will reflect further on.

Following the assumption that ICDSS needs to be considered to be modular, these different sources of knowledge should pose no problem. There is then no demand for one perfect foundation for decision-making in all domains. That means it is possible for certain decisions to be made on the basis of knowledge from written sources, whereas other modules base themselves on pre-formulated guidelines or one of the many forms of 'machine learnt' knowledge.

Together with the idea of creating knowledge from data, as follows from the analysis of work by Banae et al. and Kununninair, comes the potential to perform predictive forms of diagnosis. Both on the basis of existing medical records, but also on the basis of real-time data collection. If this function proves to be feasible in the next chapter, this would imply a major extension of current medical diagnostic and preventative activities, and presents an ICDSS that goes beyond providing support to clinicians in their current tasks. In other words, it would help clinicians to make a new kind of decision, one that is rarely practiced in current medicine.

Another interesting expectation that is closely related to the predictive capabilities of ICDSS is the real-time monitoring of patients outside the walls of the hospitals. In order to to exploit the full potential that the increasing amount of available health data presents, with regard to these predictive capabilities, it requires ICDSS to incorporate also these forms of knowledge that are continuously generated. This both includes monitoring and data collection after a clinician's orders, but also data

generation initiated and performed by patient-customers that initiate and perform their own monitoring with the growing selection of consumer-oriented wearable devices with the capabilities to monitor health-related parameters. For making optimal use of the meaning of the relation of different pieces of data, patient records would need to be centralised so that every patient has one 'personalised' record containing all health-related information and to allow for continuous updating of the contents from different sources; including raw data.

Finally this brings in the issue of personalisation. Many of the sources that have been used to characterise ICDSS proposed this as an ideal for developing intelligent features, either explicitly or implicitly. Musen, as has been shown, makes it explicit, but personalisation is also an issue in Watson Oncology—providing therapy tailored to the individual patient—and also generating knowledge from patient-data may be employed as a tool to get a better grasp of the individual patient in order to personalise treatment.

Whether this characterisation of ICDSS is exhaustive is hard to argue. The sources that were chosen for review were selected on the basis of how widely they have been cited, as an indication of their influence, so there are most certainly visions and research projects that have not been taken into account. Taking these more 'marginal' perspectives into account or going into more technical detail may have resulted in a more fine-grained picture of what ICDSS is in terms of features. However, the point is to arrive with a some degree of reliability at a characterisation of something that is—by nature of being something of the future—surrounded by uncertainty. The fact that the sources that were used (with exception of whether the system is of general intelligence, or a collection of intelligent modules) were in implicitly or explicitly in agreement or reconcilable with each other, can be understood as a signal of confirmation these are indeed part of a generally widely shared expectation of the future of ICDSS; and with it a reliable enough starting-point for our later analysis. Additionally the characterisation shares a lot of overlap with the insight drawn from the historical examples of ICDSS.

Illustrative Hypothetical Scenario

What remains, is to show how these characteristics would fit into a vision of what ICDSS would, according to these sources, ideally entail. In order to provide such a vision, drawing from the characteristics and their ideal uses, a hypothetical case-study shall be constructed that aims to elucidate how both patients and clinicians would come to interact with ICDSS and how it relates to both decision-making and knowledge production.

Imagine a fifty-year-old male person, the patient. He has received medical aid a few times in his life, but never for anything serious. The patient is an avid runner, and is busy training for a marathon while his sports watch—measuring and tracking his heart rate—streams the heart-rate data into the patient's patient-centred medical file. Back home, showering, the patient notices he feels unusually tired and wonders why when the phone rings. The person on the other side of the line calls on behalf of a clinician—could be a GP or some other specialised organisation—to make an urgent appointment at the clinic. The caller explains that the computer [the interface of the CDSS in this case] detected an anomalous pattern in the patient's heart rate and that it alerted the patient's clinician of this.

In the years preceding this incident, the use of sports trackers and the sharing of their data has grown substantively. This had given researchers the ability to data mine for, and ultimately find, correlations between specific data-patterns and medical conditions. A self-learning algorithm had been trained to recognise and classify these patterns, and had recently been implemented into the clinic's patient management/decisions support system.

The patient, instructed not to drive himself, takes a taxi to see the clinician, where he is informed that the system predicted some acute heart disease. After a brief physical examination and questioning of the patient, the clinician is uncertain why the system alarmed, how to treat the underlying causes, and how prevent it from happening. The clinician makes a decision about what tests to order and perform to narrow down his options. The clinician registers the outcomes of the consultation and orders additional tests in the system. In the meantime the system checks if these tests are indeed in line with the recommended guidelines for this type of patient. The system verifies and signals that it is, employing a module that processes the patient's updated medical file, and compares it with an analysis of both written sources as well as existing pre-constructed guidelines for his type of patient.

The clinician continues working on other patients while he awaits the results of the additional tests when an alarm on his phone rings. The results of some of the additional tests had been added to the patient's medical file and the system, drawing from literature and guidelines and historical cases, proposes a preliminary diagnosis and recommends to acutely perform a certain preventative measure. The clinician—not understanding why—requests an explanation and is provided with passages from literature, guidelines, previous cases and with statistics about the system's certainty and the risks involved in treatment. The clinician reviews the test results, the system's 'reasoning' and decides upon a certain course of action. The patient consents to the treatment and for his medical record and the results to be used in scientific research.

The patient is one of the many patients coming through the hospital. Patients with similar stories, and with radically different ones. Modules of the ICDSS may be performing certain types of knowledge generation on the basis of their medical files. When new predictive, diagnostic or therapeutic capabilities result, their underlying mechanisms may be reviewed, evaluated and validated. Ultimately to be standardised and shared across other ICDSSs across clinics and other health institutions.

Endless variations and boundary conditions

There are many things that could have happened differently in this story. The patient may not have been predictively diagnosed on the basis of a consumer wearable, but just as well because of being a long-time heart patient under constant electrocardiographic monitoring. And instead of coming to the clinic himself, the whole story could have played within an ambulance on the way from where he collapsed towards the hospital. His predictive and urgent diagnosis could have been a non-urgent diagnosis; not based on data-mining heart rate patterns, but on new insight from artificially intelligent epidemiological research involving the relation between some environmental conditions the patient the patient had been subjected to. Ultimately, the wearable hearth-tracking device may have been part of the ICDSS, independently alerting the patient to seek urgent help itself.

Furthermore the clinician in charge could have failed to choose the diagnostic tests in line with the guidelines, in which case the system may have alerted him of deviation. The clinician would in that

case have had to make a decision, to agree or disagree with the system. In fact, the clinician could have based his entire actions on the system's judgements.

The system may also have had no reason for alarm; patiently waiting for all test results to come back. Not needing to provide decision support because the clinician had seen and treated countless similar cases. Similarly the clinician may have recognised the system's reasoning to be flawed, and sent feedback to the system for improvement.

This story also identifies some boundary conditions that need to be met in order for this story to play out, both social as well as technical. A few of the technical criteria have already been discussed, such as the existence of a patient-centred medical file that can be shared across medical facilities, and the technical convenience with which the file is being filled out to completeness. Also the integrability of different modules seems critical for the possibility of sharing. Furthermore, but this will be reviewed in the next chapter, all the artificially intelligent features as are being envisioned and expected by the previously discussed sources are assumed to be feasible and adequately performing.

Not only are they presupposed to be performing, they also need to be seen by both patients and clinicians are reliable. If not, the clinic would perhaps not have called the patient and not to take the system's alarm seriously later on, or the patient may have chosen not to rely on these kinds of predictive capabilities. Also the acceptance of societies of to the relatively intense data sharing and analysis that these systems would require is of great importance for these sorts of systems to work as 'advertised'. If the patient would not have shared his heart rate data into his medical file, the anomalies would not have been detected; and if the many patients for him would not have allowed for the analysis of their data a correlation may never have been found to begin with.

Critical remarks have been made about the necessary conditions for these characteristics with regard to the technical and social aspects of its implementation. The next chapter will go deep into the technological possibilities—and most importantly the epistemological implications—of the artificially intelligent functionalities of ICDSS.

Conclusion

In this chapter it has been aimed to come to an understanding of what ICDSS may be and how it can be characterised in order to be analysed in terms of feasibility and epistemological implications; ultimately to be used for the assessment of issues emerging with regard to the compatibility with current medical practice.

It has been shown how early attempts at constructing decision support had difficulty with the complexity of medical decision-making; but also how much of their methods and aims has remained intact under current definitions, visions and expectations of clinical decisions support.

It has been explored how intelligence can be understood in the context of CDSS, and it has been argued that this intelligence lies in the integration of, often, autonomous forms of artificial intelligence into systems of clinical decision support. It has subsequently been argued that therefor ICDSS should be understood in terms of the types of functions these systems entail rather than by an all-encompassing definition of intelligence in this context.

It has been chosen for these functions to be derived from analysing visions, promises and expectations in statements and publications by leading scientists and representatives involved in the development of and reflection on (I)CDSS.

This analysis has created a picture of ICDSS that involves a few distinct characteristics. ICDSS is first of all not a general form of intelligence, it is rather a collection of intelligent functions integrated into hybrids of CDSS and patient management systems. These functions can be based on many methods, from Bayesian reasoning to Machine learning; and enable ICDSS to perform knowledge generation as well as application. It, among other things, allows CDSS to provide diagnostic support, give therapeutic advice, and provide predictive forms of diagnosis, all on the basis of AI generated understandings, drawn from the analysis of patient data and written sources. These self-learning capabilities may provide the possibility of a more personalised medicine, in which diagnosis, therapy and prevention can be tailored to smaller groups of patients within the entire population of patients. It has ultimately been identified that the generation of clinically relevant data increasingly continues outside the hospital walls; making ICDSS, at least materially, exceed the walls of the clinic.

CHAPTER III

A brief reflection on methods of intelligence and expectations

Now that it has been established how ICDSS should be characterised, and with it what is technologically required for it to work, it remains an important question on what underlying methods of data-analysis the features and functions that researchers and corporations have put forward are based on. Though this is not strictly necessary for establishing a characterisation of ICDSS, it should provide some additional insight into epistemic properties of its outcomes and provide some handles for later analysis.

It should, perhaps unnecessarily, be emphasised that this descriptions of the underlying methods shall only be in a general and brief way. This will most certainly mean that it will fail to do justice to the actual technological and methodological complexities that are involved, and that it will overlook nuances in their respective fields that some may argue are essential for understanding it. Though care shall be taken to be as complete and as nuanced as the limited extent allows, it should be noted that completeness is not the primary goal of going into the methods. It is rather to provide some grasp of how these features may work. On the one hand verifying the possibility and reasonability of the characterisation of ICDSS, and on the other hand, as mentioned, providing extra handles for reflecting on the related epistemological aspects.

Big Data

One of the underlying concepts is the concept of Big Data. Over the past few years the concept of big data evolved from something few people had heard of, to something few people don't know about. We can read and hear about big data in the context of many knowledge-intensive disciplines, for example in service of optimising business processes (Bughin, Chui, & Manyika, 2010; H. Chen, Chiang, & Storey, 2012), doing scientific research (Agarwal & Dhar, 2014; Marx, 2013; Provost & Fawcett, 2013; Swan, 2013), improving and personalising education (Picciano, 2012; Siemens & Long, 2011), and unsurprisingly, for knowledge production and decision-support in medicine (Bennett & Hauser, 2013; Johnson et al., 2016; Swan, 2012).

There is controversy over what big data precisely refers to, or where it originated, as different interpretations exist and different stakeholders have different interests in it. However, most accepted understandings of what big data is refer to the three underlying concepts of 'volume, velocity and variety' (Philip Chen & Zhang, 2014; Zikopoulos & Eaton, 2011). More precisely this refers to, large volumes of data, a high velocity of data production and processing and a wide variety of types of data and sources part of the data set.

Precisely when these dimensions are large enough for the process to be qualified as big data is undetermined, which makes it hard to identify the border between what is big data and what is not (Gandomi & Haider, 2015). In case of the ICDSS this discussion shouldn't be a problem, as it draws its data from medical files; a category of data certainly large and diverse enough to qualify as big.

The reason why the big data trend has been emerging over the past few years is most likely partially due to the very rapid expansion of the amount of available data available for analysis that has coincided with the massive increase of connectivity and data-recording devices and services (Lynch, 2008; Szalay & Gray, 2006). However, the mere availability of volume, variety and velocity has no explanatory power over the very wide interest in it. What makes big data so highly anticipated and investigated are the many possibilities that volume, velocity and variety unlock are expected and promised to unlock. The promises and expectations that have been integrated into the characterisation of ICDSS can largely be seen to stem from these very promises and expectations.

In its most modest form, big data promises the analysis of data that would be too complex to process under traditional methods of statistical analysis. More concretely this leads advocates to conclude that big data would be able to find patterns and correlations within the data that would not be visible from smaller and less diverse data sets, and that would be unable to be observed without the methods that big data is comprised of (Gandomi & Haider, 2015; Ghahramani, 2015; Kitchin, 2014a; Sabina Leonelli, 2014; Sagirolu & Sinanc, 2013; Shin & Markey, 2006; Witten, Frank, Hall, & Pal, 2016). These patterns and correlations would then allow the beholder a unprecedented insight into the processes that the involved data describe. These promises therefor ideally ultimately lead to the insight needed to improve certain processes. For example granting competitive advantages to businesses and identifying hypothetical causal relations between processes of scientific domains. However, most importantly for this thesis, it promises insight into yet unknown relations between physical and behavioural characteristics and health and disease. In the characterisation of ICDSS we can see this in its function to learn from patient data stored in electronic patient files, both on a population-level—for example in evaluating a certain therapy's effectiveness in relation to any of all other characteristics of patients—and the individual level—leading to a personalisation of care on the basis of structuring individual patients that can be clinically distinguished, again on the basis of any combination of all of the available types of data. Perhaps counterintuitively, big data's promises also enable Watson's ability to 'read' and understand text, something that shall be briefly explained below.

Machine Learning and Data Mining

Until this point, we have allowed a vague and broad understanding of big data that provides little insight into how the value that the promises and expectations entail are supposed to come to exist. That means that though it has been briefly shown that ICDSS's functions align with broadly shared expectations of big data, it has not been shown how these functions are supposed to work. Two general categories of methodologies, machine learning and data mining are particularly prominent, and cover many of the functions that big data entails—reading text, learning from patient records and subsequent predictive diagnoses and personalisation included.

Data Mining, is a concept that is quite broadly understood in the big data research communities, but that is summarised by many slightly differing definitions. Fayyad, Piatetsky-Shapiro and Smyth define data mining as "*... the application of specific algorithms for extracting patterns from data.*"

and "... a step in the KDD⁶ process that consists of applying data analysis and discovery algorithms that, under acceptable computational efficiency limitations, produce a particular enumeration of patterns (or models) over the data." (1996). In a more recent publication, Fayyad and Uthurusamy add to this that data mining "... is defined as the identification of interesting structure in data. Structure designates patterns, statistical or predictive models of the data, and relationships among parts of the data." (2002). In addition they explain that "Each of these terms—patterns, models, and relationships—has a concrete definition in the context of data mining". They explain that a pattern is a parsimonious summary of a subset of the data and that a model has predictive power over future scenarios (2002). In other words a pattern is the most frugal, but yet adequate, description of some regularity of some elements within the data; and a model can be predictive on the basis of historical results. Fayyad and Uthurusamy furthermore make an interesting remark when they state that data mining still requires the patterns and models to be comprehensible by the human users of the system (2002, p. 30). If a data mining algorithm would uncover a pattern or a model that would be too complex to verify or at all to comprehend, it becomes seemingly impossible to use it. However, humans are necessarily required in connecting consequences to the findings of data mining. The patterns and models could for example be, as complex as they may be, used in other algorithms, to connect consequences to them without humans having to understand and act upon them. Take for example a scenario in which a data mining algorithm would uncover a complex pattern within some aspect of consumer behaviour depending on the placement of products and sections in the supermarket. Instead of requiring humans to understand the assumingly complex pattern and to rearrange the positioning according to the desired shopping outcomes, it would perhaps be more efficient to programme the pattern into an algorithm or a model to render possible arrangements with their corresponding chances of success.

Machine Learning

This somewhat 'manual' solution is however not the only solution to this problem. In line with it lies the practice of machine learning; the second prominent big data analysis method. A very commonly used definition of machine learning is one put forward by Tom M. Mitchell, stating that "A computer program is said to learn from experience *E* with respect to some class of tasks *T* and performance measure *P* if its performance at tasks in *T*, as measured by *P*, improves with experience *E*." (1997) In a similar but slightly less technical definition, Ethem Alpaydin adds to Mitchell's definition by describing machine learning as "... programming computers to optimize a performance criterion using data or past experience" (Alpaydin, 2014, p. 3).

In other words machine learning can be understood as a form of data analysis that aims to 'learn' to solve problems on the basis of data of historical examples of actual solutions. So like data mining, machine learning involves the identification and construction of patterns, and models, but ultimately all aimed at the relations between controllable circumstances, uncontrollable circumstances and producing eventual successful outcomes. Mitchell's notion of 'performance measure' and Alpaydin's mentioning of 'a performance criterion' furthermore add that the performance of the algorithm employing the patterns and models that were found and constructed is quantified; a quantification that is subsequently used to continuously improve and optimise the patterns and models it uses.

⁶ Knowledge discovery from databases

The way in which machine learning systems actually learn, is commonly divided into three types of learning: supervised learning, unsupervised learning and reinforcement learning (Chapelle, Scholkopf, & Zien, 2009; Dougherty, Kohavi, & Sahami, 1995; Huang, Song, Gupta, & Wu, 2014; Li, 1994; Schölkopf & Smola, 2002; Smith, 2002; Tesauro, 1994)

In a publication on the potential of machine learning in clinical decision support systems Shin and Markey summarise unsupervised learning as "*... the computer attempting to identify natural groupings within a dataset based on criteria that define how "similar" items are and what makes a "good" group, but without being provided examples of the feature values of items and associated "correct" class membership*" (2006, p. 238). In other words, unsupervised learning requires the system to find patterns of which merely has requirements, but no explicit examples. In the example of ICDSS, unsupervised learning might for example take place by providing an autonomous sub-function with an instruction of how to 'recognise' an etiological relation between external factors and health and disease and subsequently instruct the function to analyse a large dataset including descriptions of external factors and diagnoses for particular individual, or groups of, patients.

In supervised learning, Shin and Markey state, "*...the computer is provided with examples of the feature values of items and associated "correct" class membership.*" (2006, p. 238). So in supervised learning the system is provided with examples of the patterns it is supposed to find instead of clear criteria for the kind of relation. That is a mode of learning that is particularly useful when it is complicated, or even impossible, to adequately describe what it is that the system must detect. In the context of ICDSS this could for example take the form of providing the system with a large set of accurately diagnosed mammograms, to identify how certain visual properties relate to the diagnoses, and then apply these identified relations between visual properties and diagnoses to undiagnosed mammograms in order to diagnose them. Under supervised learning the algorithms receives specifically quantified scores, such as a percentage of accurately diagnosed images that it then needs to optimise.

Reinforcement learning, on the other hand, the system is neither unsupervised, nor supervised, instead it receives "*...less specific feedback that indicates if the system is on the right track.*" (2006, p. 238). This 'less specific feedback' then often consists of positive and negative stimulation (reward and punishment). The system then does not precisely know on what aspect it is being rewarded or punished, it is merely motivated to estimate what could be done in pursuit of more positive and less negative feedback. In the context of ICDSS this could be applied by allowing clinicians to give a 'thumbs up' or 'thumbs down' to any given form of advice or summary the system may provide; for example like Watson Oncology allows a clinician to give 'feedback' on it's diagnostic reasoning.

hype cycle

It is clear that these three different methods of machine learning enable different applications. It is however not immediately clear where their respective limitations lie with regard to precisely what these methods enable. If we were to follow Gartner's 2016 hype cycle, a debatable but also authoritative indicator of technological development, we come to conclude that the technology is generally most likely overestimated. The cycle identifies that machine learning is at the 'peak of inflated expectations' in 2016, which implies that the expectations are currently beyond the capabilities of the technology and impact it eventually may have (The Gartner Group, 2016).

This is an important remark, since we have characterised ICDSS at least partially on the basis of expectations—including those expressed by parties potentially benefitting from inflated expectations on the short term (for example in terms of business interest or attracting research funding. Though great care has been taken to select credible sources and to verify and balance the expectations from commercial businesses with the outcomes of empirical work, a brief investigation of what this inflation entails and how it relates to our characterisation of ICDSS is in place.

Precisely what inflation is is unclear, and no full-scale investigation of the inflated expectations of machine learning can be found among its large body of literature. There are however some separate sources that together provide an argument for what these expectations entail, and why they are inflated. Please note that this does not focus on all the expectations of machine learning, just on the tenets that are clear subject to the inflation.

Assumingly the first and most fundamental of these tenets, contain inflated expectations on the time scale for machine learning to come to fruition and on the ultimate ubiquity it will come to have (The Gartner Group, 2016). In other words, one may expect machine learning to come to realise its full technological potential faster than it reasonably can, and expect the machine learning to be applied quicker to more aspects of the world than it often turns out to be.

There are however also more specific expectations with regard to the underlying principles of machine learning, that also relate to the expected ubiquity of the technology. The first of which is the assumption that all, or at least most, aspects of the world are structured and bound by rules, and that these rules inherently cause identifiable patterns to exist within the data describing the world. These patterns are identified using correlations, that are then in turn assumed to correspond with causal relations existing in the 'real' world (Domingos, 2012). In other words it is assumed that data describing the world is necessarily full of correlations, and correlation can be enough to suppose causation.

Specifically relevant to this thesis, these two core assumptions behind much of big data expectations lead to a particularly appealing picture of knowledge production, that is not often stated explicitly, but that can be recognised in many of the optimistic statements of machine learning's potential with regard to knowledge producing data analysis tasks. One, intentionally, rather provocative article fully embodying these assumptions to great expectations was written by Chris Anderson, at that time Wired Magazine's editor in chief (2008). Though it was written almost a decade before Gartner's 'peak of inflated expectations', Anderson's expectations are quite high. In the article he argues how modelling is soon to become an unnecessary activity in scientific inquiry. Instead the patterns within big data allow objective empirical facts to be discovered that conventional scientific methods would not be able to.

More precisely, Anderson opposes the concept of modelling by arguing that: *"There is now a better way. Petabytes allow us to say: "Correlation is enough." We can stop looking for models. We can analyze the data without hypotheses about what it might show. We can throw the numbers into the biggest computing clusters the world has ever seen and let statistical algorithms find patterns where science cannot"* Anderson adds to this that *"The new availability of huge amounts of data, along with the statistical tools to crunch these numbers, offers a whole new way of understanding the world. Correlation supersedes causation, and science can advance even without coherent models, unified theories, or really any mechanistic explanation at all"* (Anderson, 2008).

Anderson draws a very explicit parallel to scientific research, and makes the claim that knowledge production no longer requires the scientific method, the formulation of hypotheses or causal-mechanistic models and explanations. In other words, Anderson argues for the irrelevance of theory. Anderson is just one very accessible example of these expectations, but very similar positions can be found throughout works by more 'visionary' authors (Dyche, 2012; Prensky, 2009)

Robert Kitchin refers to their shared position of theory-free science with correlation as ultimate epistemic authority, as proposing a 'new era of empiricism' (Kitchin, 2014a). More specifically Kitchin identifies that the following four "powerful and attractive" ideas are at work that 'enable' these inflated expectations.

1. *Big Data can capture a whole domain and provide full resolution;*
 2. *There is no need for a priori theory, models or hypotheses;*
 3. *Through the application of agnostic data analytics the data can speak for themselves free of human bias or framing, and any patterns and relationships within Big Data are inherently meaningful and truthful;*
 4. *Meaning transcends context or domain-specific knowledge, thus can be interpreted by anyone who can decode a statistic or data visualization.*
- (Kitchin, 2014b)

But, regardless of how powerful and attractive these ideas are, they remain ideas. Demonstrably wrong, Kitchin shows. Very briefly summarised, Kitchin argues that: **(1)** Data can never capture a whole domain, it is furthermore subject to sampling bias and its always a view from somewhere rather than an "infallible gods-eye view" here (Amin & Thrift, 2003; Crawford, 2013; Kitchin, 2014a, p. 4). **(2)** Big data never comes out of nowhere. Instead Kitchin argues certain kinds of data are being captured within the scope of scientific frameworks and that they are subsequently processed on the basis of scientific reasoning (S Leonelli, 2012). He subsequently concludes that: "*New analytics might present the illusion of automatically discovering insights without asking questions, but the algorithms used most certainly did arise and were tested scientifically for validity and veracity*" (Kitchin, 2014a, p. 5). **(3)** Data can never speak for themselves, instead they only receive meaning upon observation through a certain lens. Even if this lens is situated in an autonomous algorithm, it is still "*imbued with particular values and contextualized within a particular scientific approach*" (Kitchin, 2014a, p. 5). He furthermore makes the point that correlations can be random in nature, and are subsequently not inherently meaningful, nor useful. **(4)** Though data and correlations can perhaps be interpreted without knowing anything about the domain from which they were distilled; it is unlikely that these correlations will be particularly useful. Or as Kitchin states slightly more eloquently: "*Put simply, whilst data can be interpreted free of context and domain-specific expertise, such an epistemological interpretation is likely to be anaemic or unhelpful as it lacks embedding in wider debates and knowledge*" (2014b, p. 5).

Kitchin convincingly shows that the expectations of a fully autonomous empiricism independent of, and disconnected from, current and traditional scientific disciplines and knowledge are indeed

inflated. Not merely beyond what it may turn out to be, but inflated beyond what is logically possible. It has been mentioned that the discourse surrounding a 'new empiricism' is a relatively radical one that may be considered to be part of the most optimistic explicit expectations that can be found. But as has been argued as well, elements—such as the epistemic authority of correlation, the objectivity of data and their subsequent conclusions and the independence of theory—of this optimism have found their way into more modest expectations within the domain of data science and business. Parts of it have even made their way into our characterisation of ICDSS. The urgent question here, is whether that poses a crucial problem.

Arguably not. Though the characterisation incorporates some elements of the optimism that was described—such as the convenient levels of autonomy and the ability to identify patterns within complex and seemingly incoherent data—none of these functions require the data to describe the whole domain, be free of biases or underlying methods and models. Nor do the correlations need to speak for themselves. There is no reason why all of these machine learning activities would not take place under intensive human involvement, understanding the domain, interpreting correlations, supporting the whole knowledge creation process and verifying its outcomes.

In fact, on somewhat of a side note for now, as long as human clinicians remain the ultimate decision makers, knowledge generated using machine learning (and data mining for that matter) is required to be embedded in the clinician's understanding of the medical context; if ICDSS is to be transparent to some extent. It is therefore seemingly not useful at all to have some free-of-theory type of correlation based knowledge; as it would be impossible to implement it in a practice that heuristically relies in its very functioning on biased theories and understandings the human clinicians hold in order to make decisions and perform procedures.

Kitchin agrees with the position that the 'optimistic' features of autonomy and the ability to detect invisible patterns do not necessitate the four 'ideas' behind the empiricist view to be complied with; nor are they necessary for machine learning to be of value to the knowledge creation process. Not in conventional science, and arguably neither in clinical practice. Instead Kitchin shows 'data-driven science' to be a more likely method of using big data analytics in knowledge creation processes. With a few statements Kitchin argues that data-driven science "*... is guided in the sense that existing theory is used to direct the process of knowledge discovery, rather than simply hoping to identify all relationships within a dataset and assuming they are meaningful in some way.*"

Furthermore, "*Data are not generated by every means possible, using very kind of available technology or every kind of sampling framework; rather, strategies of data generation and repurposing are carefully thought out, with strategic decisions made to harvest certain kinds of data and not others.*"

"The data are not subject to every ontological framing possible, or every form of data-mining technique in the hope that they reveal some hidden truth. Rather, theoretically informed decisions are made as to how best to tackle a data set such that it will reveal information which will be of potential interest and is worthy of further research." (Kitchin, 2014b, p. 6)

In other words, Kitchin distances himself from a highly autonomous knowledge generation process in which machines process data by applying endless methods, frameworks and interpretations, and proposes a mode of generating knowledge instead, that—though originating its findings within data—intensively requires human agents in the process to interpret, theorise, hypothesise and verify. A picture of knowledge generation that is highly compatible with our previous analysis identifying that

there is no need for a theory-free fully autonomous science for ICDSS to perform its intelligent functionalities.

CHAPTER IV

Analysis of Differences and Issues

Now that this thesis has presented a substantive amount of background on clinical practice, and has provided a workable and reasonably feasible characterisation of ICDSS, the requirements for being able to explicate how they differ and to delve into issues of epistemology and compatibility of ICDSS with current medical practice have largely been met. But before starting to identify epistemological issues and aspects of the compatibility of ICDSS and clinical practice, it is necessary to explicate what is meant by this and how this task will be interpreted.

Throughout this thesis it has become more than clear that ICDSS is very much a technology of the future. Though certain limited functionalities may be operational within certain experimental or commercial contexts (which partially confirms their feasibility); a system that incorporates functions that were described as part of the characterisation is nowhere near available to clinical practice.

This also highlights again the fact that it is in fact a rather explorative exercise to come to an understanding of the differences between ICDSS and current medical practice and to come to conclusions about the issues regarding their compatibility. It is difficult to consistently do justice to these uncertainties in the following analysis, and some claims that will be made could be accompanied by countless nuances. Though this thesis aims to ultimately provide useful insights, it is not feasible within the scope of this thesis to provide that level of completeness.

In order to prevent us from subsequently writing one irrelevant general analysis, this chapter discusses the differences and issues separately for the different methods of current medical practice in our tidy methodological pluralism. It shall subsequently address overlapping and general issues ordered into sub-themes. The themes that will be discussed in this chapter are issues of Big Data transparency, autonomous dealing with scientific controversy, dealing with personalised knowledge, predictive diagnoses, the challenge of data as a diagnostic instrument, clinician responsibility and issues with standardisation.

Compatibility of ICDSS with the methods of current medical practice

Narrative Medicine

Decision-making has been characterised as being predominantly composed of EBM and narrative medicine. Narrative medicine has been explained to focus on the patient's story, by 'acknowledging, absorbing, interpreting and acting on the stories and plights of others' as Rita Charon defines it. A more humanistic mode of medicine, in which health should not be defined in mere biochemical or functionalistic terms.

In providing an answer to whether there are epistemic issues with the role of ICDSS in narrative medicine we shall start by exploring an answer whether ICDSS would be able to present the abilities that would be required by narrative medicine. In the first chapter it has been shown how Solomon identified and defined that there were 4 competences that clinicians would need to have. 'Listening and Witnessing', 'Empathy', 'Narrative Detective Work' and 'Meaning Giving'. Though an argument could be made that ICDSS would be capable of listening and witnessing and narrative detective work, using algorithms to classify moods and narrative patterns, and that meaning giving would be dependent on how patients perceive artificial intelligence in social relations, it seems well beyond the big data capabilities of ICDSS to 'know at an experiential level what a patient goes through'.

If the full set of competences is indeed necessary to be able to actually perform narrative medicine autonomously, this means that ICDSS will not be capable to perform narrative medicine. The remaining question then is whether it is still possible for ICDSS to provide support to decision-making given that narrative medicine is an inherent part of this.

It seems that there is a fundamental epistemic issue involved in this question. After all, the fact is that what narrative defines as 'good' clinical outcomes, is difficult or even impossible to quantify given the unstructured integration of subjective narrative components into the decision-making. Narrative medicine has no definitive criterion for what constitutes ideal clinical outcomes, other than perhaps the patient narrative having been perfectly incorporated into the decision-making process—in turn again being unquantifiable.

The reason that this is a problem is that ICDSS's intelligent features are all based on the processing of large amounts of data. Data of the patient narrative could then certainly exist, for example in the form of transcripts of the interactions between clinician and patient. But what cannot be turned into data is the extent to which narrative has been optimally incorporated, something machine learning so desperately requires as a performance indicator. In other words, ICDSS would not be able to teach itself as there is no way in which it could know how well its functioning aligns with narrative medicine's ideals.

Of course one could find alternative factors, that could be quantified, such as the degree to which the patient is satisfied about the integration of his or her narrative into the clinical process. But that introduces methodological issues. Because this would mean that good narrative medicine would lead to satisfied patients—with satisfaction not necessarily connected to what the patient requires. This also causes issues with regard to the continuity of patient desires. A patient narrative may very clearly suggest a particular course of action at the beginning of some therapy, perhaps even with the patient being highly satisfied about the decision that has been made, while the patient may be deeply unsatisfied with it afterwards. Not meaning that narrative medicine has been sub-optimally implemented, but merely that the patient is not satisfied.

However, this is problematic only for scenarios in which the ICDSS provides clinicians with only one best available course of action to a patient. Though that would be the most 'powerful' picture of decision support, it may very well not be the most realistic picture and certainly not the only one. A scenario that presents itself as a plausible alternative here, is one in which the ICDSS provides the clinician with a selection of possible courses of action that can then be taken into consideration into the narrative process; discussing options with patients, weighing the consequences of options against patients' secondary values and needs, etc.

Evidence-Based Medicine

EBM has in terms of decision-making initially been introduced in this thesis as a being proposed as a technological solution to the limitations human clinicians have in making decisions in accordance with the strict criteria EBM imposes. Throughout this thesis, this idea may have been somewhat nuanced, even though the ability to process the full body of available written sources has partially been left intact.

The impossibility of acting completely in accordance with EBM, lies in the idea that EBM requires a perfect evaluation of scientific evidence in order to come to the best possible treatment. This is already highly problematic for human clinicians, given that there are little rational grounds to prefer one study when confronted with some conflicting studies of equal quality and evidential hierarchy. This problem is not at all being solved by moving towards an algorithmic epistemology. The underlying problem here is that it is hard to establish a priori what, in clinical terms, the best possible conclusion from a conflicting body of literature is, a problem that strongly relates to the issues of controversy resolution.

As a partial solution, one could though argue that ICDSS would be able to use its 'knowledge' of historical records to provide historical empirical support for certain conclusions of the controversy, assuming that a given interpretation of some controversy has been assumed in historical cases of which clinical outcomes have been recorded. In other words, ICDSS may be capable to perform additional research on the basis of its large collection of records with regard to the effectiveness of certain therapeutic decisions, simply because guideline-following is not flawless in medical practice and variations among treatments for similar conditions exist.

Many reservations should be made about this strategy as the ICDSS would be unable to verify whether what has been recorded is indeed how treatment was applied and subsequently have no control over the quality of the virtual experiment. It would subsequently be absolutely impossible to meet the requirements of randomised controlled trials.

Another very different problem that emerges with the knowledge production capabilities of EBM in ICDSS, is that it is unclear how the machine learning and data mining capabilities rank into the evidential hierarchy that EBM maintains. It is clear that under current medical practice the randomised control trial ranks highest, and that an ICDSS can't autonomously simulate it, but it is unclear how knowledge of a particular pattern or correlation would rank into the evidential hierarchy. Solution could be found to this question, simply by developing a system for this, but the question of how evidence ranks is an interesting one.

However, it is important to note that this problem only exists if ICDSS would indeed make use of the evidential hierarchy to order the significance of different pieces of knowledge. Given that ICDSS is a big data driven type of knowledge production, it is not unlikely that ICDSS would use its own machine learnt way of ordering different pieces of data rather than using the rigid pre-determined categories of EBM. Questions could be asked whether this would then still be EBM, or whether it should be understood as a new method of decision-making.

Consensus Conferences

We have argued that ICDSS is faced with a number of challenges with regard to the resolution of controversy. The conclusion of this is that it is unclear whether and how ICDSS would resolve scientific controversies.

Whether there is a role for consensus conferences with ICDSS in this respect, perhaps depends on how authoritative and determinant ICDSS's functionalities would be considered. In a scenario in which ICDSS is completely trusted in its autonomy, there may not be room for human clinicians and experts to intervene with the knowledge that has been generated by the machine, nor would there be a need for them to convene and determine the way in which there should be decided upon some scientific controversy.

However, trust must generally be earned, and it is unlikely that trust would be widespread in the early phases of its implementation. So even in a scenario in which a perfectly adequate ICDSS exists that functions in accordance with our characterisation, it would be unlikely that it would be blindly trusted; perhaps in none of its autonomous features.

In the situation that trust would not be yet in place, it makes sense for ICDSS to be programmed to acknowledge some of its assumed weaknesses and signal human clinicians when these weaknesses may influence the ICDSS's functioning. It could for example request a consensus-seeking effort from human clinicians when it detects unresolved scientific controversies, or when it finds a seemingly controversial pattern or correlation itself. Instead of trusting and forcing the ICDSS to make a decision about their value and act on them immediately, there would be room for human clinicians and experts to assist the knowledge production capabilities of the ICDSS and help it determine how to deal with a controversy or newly found piece of knowledge by organising a consensus conference. If not because the ICDSS would actually require it, then perhaps because human clinicians demand to remain in charge over controversial aspects of the ICDSS's capabilities.

Translational Medicine

The second chapter showed that translational medicine, though difficult to precisely characterise, involves a continuous effort to make laboratory knowledge available to actual clinical practice by applying it to patient care.

With regard to the compatibility of ICDSS with Translational Medicine there are both challenging issues and opportunities. Under current medical practice clinicians can impossibly be aware of all the available knowledge that is relevant to their discipline, for them to also be aware of all the knowledge that has been produced in highly controlled laboratory environments would be even less feasible. Assuming that ICDSS would indeed be able to link the available written sources to actual clinical cases in at least some cases, scenarios can be imagined in which ICDSS notifies clinicians that certain experimental treatments may deserve consideration for certain groups of patients, even across hospitals. A scenario in which the ICDSS merely makes the clinicians aware of existing knowledge.

This may seem like a promising possibility, however, the problem is that laboratory knowledge is not always explicitly linked to medicine, human physiology or clinical practice in general—even if it would actually be applicable and productive there. Laboratory knowledge involves fundamental findings from basic research in some domain within the sciences, describing abstract phenomena within controlled environments. This makes it potentially extremely complex for an ICDSS to identify that some of the described knowledge has a potential application in clinical practice, as its intelligence is so dependent on verbalised knowledge and relations.

Furthermore, much of the knowledge that is being produced in laboratories, or at least the most recent interpretations of that knowledge, is most possibly not published yet; meaning that ICDSS's text mining capabilities would not have access to them. Clinicians working in close contact with these researchers however would, and would be enabled to bring knowledge into clinical practice at an even earlier stage.

This is however all very much a theoretical perspective on the translation of knowledge; seemingly presupposing that merely the availability of laboratory knowledge is enough to enable the translation of it to clinical practice. With this it assumes that 'random' clinicians would be able to apply the very specific laboratory knowledge in a way that actually contributes to effectively bringing knowledge from the bench to bedside. The process of actually translating the laboratory knowledge to relevant and effective clinical knowledge involves much more than simply being made aware of the possibility of applying it to a particular clinical case. Much rather it may involve clinicians with detailed understandings of the laboratory knowledge and with well-developed translational and experimental skills to be able to perform the difficult and manual process translation from basic science to clinical trials and from clinical trials to clinical practice (Mankoff, Brander, Ferrone, & Marincola, 2004). After all, applying unrefined laboratory knowledge is seemingly certainly not a process of following guidelines or otherwise straight-forward application, but of tailoring instead. It should however be considered that ICDSS may be able to aid doing the research, by data analysis, by finding similar patients, by facilitating collaboration across clinics, etc.

Additional Issues

Though this thesis has largely followed a somewhat analytical approach by deducing issues from the characterisations of ICDSS and current medical practice; additional issues came up in the process of writing and contemplating the contents of this thesis project. Precisely how these issues were identified or from which sections they developed is therefore not entirely clear. The fact that clarity in this respect would probably have strengthened the analytical qualities of this thesis, should however not be a reason to leave them unaddressed, neither to simulate some clear origination. These additional issues will be shown to be valuable, also without needing a clear connection to our analysis.

Issues of Big Data Transparency

Big Data transparency can metaphorically be understood as the property of being able to see into the container in which ICDSS's processes take place; allowing someone with some knowledge of the system to understand why the system does what it does. This would for example allow someone to come to a meaningful understanding of why an ICDSS provides certain answers to certain questions, or makes the predictions and diagnoses it makes on the basis of how it processes what data (Lisboa & Taktak, 2006, p. 1).

Providing ICDSS with this type of transparency is not without challenges, and it is unlikely that transparency can be provided in all aspects of the system's functioning. This becomes clear first of all in ICDSS's text mining capabilities that were discussed in the second chapter. According to IBM, Watson has the ability to 'read' large amounts of written text and to bring them into meaningful relations to each other.

In the example of Watson, the ICDSS communicate the assisted diagnoses and therapies that with brief sections from sources from which Watson arguably draws its conclusions. Though this certainly enables the clinician to some extent determine for him or herself whether the sources that the ICDSS provides and the conclusion it draws are reasonable, and one may argue that this may be enough to work with Watson in a responsible and accountable way, it does not really lead to an explanation of why Watson precisely selected those sources and ignored others. Neither is being able to 'verify' a certain advice on the basis of the provided literature enough to understand why the machine drew a certain conclusion. After all, since the ICDSS's understanding is machine-based and lies in information on probabilistic relations between words and sentences rather than in a clear verbalised conceptual understanding, it may very well be impossible for the average clinician to retrace the underlying logic of the machine's understanding and deduce why it produced the outputs it did (Abdelfattah & Fuji, 2008; Aggarwal, 2015).

On a more methodological level transparency is an issue with machine learning and data mining, methods that have been shown to be crucial for ICDSS's functioning. With these types of data processing transparency is a rather common challenge, as the whole point is to exploit patterns that may be hard to find and conceive by humans (especially in the case of artificial neural networks). Though it is of course possible that a pattern is hard to find, but reasonably easy to recognise by a human, this would still require the human clinician to have a good enough understanding of data in general and the relevant data in particular to be able to do so.

However, for a pattern to be reasonably recognisable by humans, it must not be too complex. It can though easily be imagined that these algorithms would detect statistically significant relations between a large and complex number of variables that may make it beyond reasonably possible to see the pattern from the raw data, even if it is perfectly described right in front of the clinician. This would make it impossible for the clinician to verify whether the conclusion the ICDSS draws indeed follows from the presented data.

As was explained in the previous chapter, machine learning aims to improve itself in a particular task by finding relations between the environment and the successful outcome of the task. In practice this often means that a machine learning algorithm goes through a process of trial and error of implementing and trying out the influence of taking into account every single correlation it can find in a given dataset in order to improve the performance measure. Because the point is not to provide intelligible relations or conclusions to human decision makers, but for the machine to make its decisions itself (as is the case in for example self-driving cars, or automatic mammogram analysis) there is no limit to how far-fetched a correlation can be as long as it appears effective in relation to maximising the performance measure (i.e. driving as safe as possible or being as accurate as possible in detecting tumours).

In effect this results in the possibility of machine learning algorithms to function outside of the conventional understanding of medical practice, which causes severe problems for the relation between input and output to be understood by the ordinary humans that most clinicians seem to be. In other words, being transparent is truly a challenge for the functionalities of ICDSS that are based on machine learning.

Whether transparency, or rather lacking transparency, should be considered as problematic depends on the necessity of transparency under current medical practice. The positive connotation that transparency carries throughout many aspects of society makes it convenient to deem it desirable or even crucial to the process of decision making under current medical practice. However, drawing from the tidy methodological pluralism by which this thesis characterises current medical practice, it seems that current medical practice is not necessarily so dependent on transparency. As Solomon describes, the outcomes of EBM's guideline production are very often outcomes of (partially undocumented) social processes, leading to a social epistemology of limited transparency rather than a fully retraceable decision on the basis of perfectly rational and logical evaluations of facts (if such a thing would even exist conceptually). The same thing counts for narrative medicine; after all, how transparent can the relation be between the clinician's perceived patient narrative and the therapeutic conclusions he or she connects to them.

Furthermore, as far as the transparency is concerned, it seems that the medical technologies that clinicians routinely rely on provide little transparency with regard to how they produce their information. After all, assumingly few clinicians would be able to verify that, for example, an MRI scanner or an automated laboratory test produces accurate results on the basis of inspecting the internal operations within these machines. Nor is it common practice to verify an MRI scan by cutting a patient open and manually inspect a patient's internals. Instead, it seems that these technologies are considered to have been adequately verified to the point where trust replaces the need for transparency. To draw from Latour's vocabulary⁷ on the opacity of science and technology, one could say that current medical practice does not seem to be in a serious conflict with technologies being black-boxed as long as the level of trust in its operation is high enough for the clinician's responsibility to accurately perceive to be outsourced to the machine.

Though this makes it seem reasonable that also the machine learning capabilities of ICDSS may become black-boxed without much concern as well as long as it lives up to its promises and provides good enough reasons for it to be trusted. The thing with machine learning is that they are ideally continually trained in order to be further improved. This means that black-boxing a functionality that is supported by machine learning does not merely necessitate the current state of the algorithm to be trusted, but also requires trusting the self-learning capabilities that underlie the functionalities to be trusted, and therewith the future states of the algorithm, as well.

Dealing With Scientific Controversy

Another important epistemological issue that is strongly related to the automated analysis of written sources is the way in which it deals with scientific controversy. Scientific controversy fundamentally comes down to the fact that science is not a straight-forward process of knowledge production. Different explanations may exist for similar results, and different theories, each with their own implications, may seemingly effectively explain the same phenomena.

This continuous controversy is ultimately reflected in the fact that scientific literature is not always in perfect agreement with each other about the best way to diagnose or treat a patient or the

⁷ "the way scientific and technical work is made invisible by its own success. When a machine runs efficiently, when a matter of fact is settled, one need focus only on its inputs and outputs and not on its internal complexity. Thus, paradoxically, the more science and technology succeed, the more opaque and obscure they become." (Latour, 1999)

ideal therapeutic course of action. This seems to be obvious at this point, given that we have discussed in some detail how EBM has continuously struggled with the problem of distilling evidence-based guidelines from complex collections of unaligned and often conflicting scientific knowledge. As we have also drawn from Solomon's work, much of the resolution of the controversies that underlie the problem of finding guidelines, has been found in social epistemic rituals, such as consensus conferences where diversely opinionated people collectively come to some consensus on how best to interpret the available data and come to one actionable conclusion.

If ICDSS is to autonomously perform a continuous analysis of all the written literature that is being published. It would, assuming that it would indeed be capable of accurately understanding the content of the literature, inevitably encounter the incommensurability of certain works. Given that disagreement, though of varying degree, is so widespread within the medical scientific communities the system would not be able to exploit its text analysis capabilities if it would not have some way to draw authoritative conclusions from the very different and conflicting views that literature may provide—at least as long as the system would be supposed to provide one, or a limited number, of best possible answers instead of identifying all collections of the available mutually commensurable views.

An algorithm would most certainly be able to construct any of the many possible conclusions from the controversy in one of many ways. In a least refined scenario it could for example randomly, without any epistemological support, pick one. Or, just count how often a certain claim is supported and choose the most 'popular' one. But it could also derive epistemic authority from external information, such as the scientific status of certain authors, articles and journals and build a system in which those external factors can be structured in relation to each other, and factor that into how certain claims and results are weighed into the controversy.

However, realist interpretations of scientific knowledge within medical practice are not uncommon. The idea then is that an understandable reality exists, and the question is what theory provides a more likely and more accurate description of that reality. The most truthful theory should then be privileged. Scientific realism, is commonly understood as a matter of belief, something ICDSS has not been characterised to be capable of. This raises an interesting issue in terms of the role for scientific realism in controversy resolution by ICDSSs. After all, unless the ICDSS would have direct access to the truth, something we shall assume not to be the case (if at all possible, the system has no definitive way to measure or determine the truthfulness of theories, claims and results. In other words, the ICDSS cannot provide its machine learning algorithms with a truthfulness performance measure, and subsequently can't hierarchise possible outcomes accordingly, nor train itself to deal with it.

Apart from the issue of ICDSS being unable to judge the veracity of claims and theories, there is another potentially significant difference with current medical practice, regardless of the method it employs. As mentioned, current medical practice relies on a social epistemology in the resolution of controversy. Solomon has been shown to explain, this means that knowledge is constructed through social processes; something this thesis ascribes to. The ICDSS can however not perform these social processes autonomously, all it can is exploit certain performance indicators or pre-programmed methods or rules to construct knowledge. It would be an algorithmic epistemology rather than a social one. The issue here is that this changes the nature of medical knowledge.

What the consequences of such a change from a social to an algorithmical nature are, cannot be immediately clear. In that respect the nature of knowledge is hard to grasp. What is clear nonetheless, is that it again emphasises the question of where trust is to be placed. Under current medical practice

this trust is placed in authorities individuals and institutions and the 'power' of social processes of consensus seeking (such as conferences), whereas ICDSS would require trust to be placed in the machines capabilities to perform an analogous operation based on logic and rule-following.

The issue of personalised knowledge

The uncomplicated part of understanding ICDSS in terms of personalised medicine lies in the general idea behind personalised medicine. This general idea is that "*... to optimize medical care and outcomes for each individual, to include treatments, medication types and dosages, and/or prevention strategies that may differ from person to person—resulting in an unprecedented customization of patient care.*" (Ginsburg & Willard, 2009, p. 279). Though this is merely a single quote, it provides a nice description of how personalisation as a general concept can be understood as aiming to come to better understandings of patient diversity in order to be able to tailor diagnosis, treatment and prevention to their personal clinical needs. The underlying presupposition is of course that differences between patients imply different clinical needs; the question then remains what 'differences' are relevant and should be seen as having predictive value for patients' clinical needs.

Perhaps not quite surprisingly, these differences are very often considered to be most relevant in patient's genomes; following the established belief that genomes provide the natural basis for human (patho)physiology (Alberts et al., 2013). Concretely this means that much of personalised medicine focusses on patient's genomes in order to determine what their clinical needs are (Burke & Psaty, 2007; Ginsburg & McCarthy, 2001; Hamburg & Collins, 2010). The first way in which the genome may be used, which is often the case, is through a rather common process of 'stratification'. In stratification patients are grouped by genetic markers of which evidence exists that they correlate with relevant clinical characteristics (Cardon & Palmer, 2003; Lin, Yu, & Yang, 2010; Lindon & Nicholson, 2014). This could for example mean that, if knowledge exists of a relation between a certain gene and higher effectiveness of therapy A over B, a clinician can use that knowledge to help to give the best possible treatment to a patient that turns out to have that gene in his or her genome.

This method of using genomic knowledge for personalisation is very much dependent on correlations between genetic markers and clinical needs. Though not entirely separate from this, systems biology and similar disciplines aim to provide personalised medicine by aiming more for simulating the processes that are derived from a patient's genome, and how they relate to therapy, diagnosis and prevention (R. Chen & Snyder, 2012).

Under these understandings of personalised medicine, however, it seems as if ICDSS does not bring in anything dramatically different from what can already be seen with regard to personalisation under current medicine. After all, knowledge for increasingly specific categories and clinical needs are being developed under evidence-based medicine already (Altman & Royston, 2000; Cardon & Palmer, 2003; Oliver, Britton, Seed, Martin, & Hopper, 1997; Oliver, Daly, Martin, & McMurdo, 2004; Schoene et al., 2013).

The work that this thesis has done so far on personalisation has however not at all explicitly gone into the genomic side. In fact, this thesis has consistently maintained an understanding in which patient data has been treated in a general sense. Referring to patient data in this thesis subsumed genomic and genetic data as well, but is explicitly not limited to it. This is largely following the dominant machine learning discourse in which helpful correlation does not necessarily need to be

derived from parameters that would be considered meaningful under current medicine, but may be found and productive within any type of data.

With this, there is a significant difference between current forms of personalisation and the one that ICDSS entails. A difference that unsurprisingly lies in the kind of data from which correlations may be drawn. Under current medical practice, these correlations are all being based on aspects of which widely shared theories predict or presuppose clear causal mechanistic significance with a patient's physiology. A person's genome, as mentioned, is for example often understood to be directly linked to the molecular physiology of a patient (Alberts et al., 2013). Similar mechanistic linkages are being presupposed for other quantifiable properties of the human body, such as for blood perfusion of body parts, electrocardiograms, imaging technologies, cell counts, biochemical markers, sex, ethnicity and race, and many more (Cheng, Chan, Cembrowski, & Van Assendelft, 2003; C. Mitchell et al., 2009; Monin et al., 2001; Nyman et al., 1993; Owens et al., 2007; Roberts & Fromm, 1998; Sharir et al., 2001). ICDSS, as characterised, would instead be able to use all these patient properties just as well. However, where it grows apart from current medical practice is that it may be using patient properties that under current medical practice are understood to be of little causal mechanistic meaning. It should be mentioned that this argument is subject to the possibly inflated expectations that machine learning entails, namely that productive correlations may be found within the most trivial and unanticipated types of data.

Assuming that it may indeed turn out to be realistic, it would mean that medicine would need to accept seemingly trivial and unconnected parameters of patients as potentially meaningful. Apart from possibly extending the understanding of clinically relevant data, it would also require clinicians to develop a trust in ICDSS and personalised medicine that is not founded on causal mechanistic understandings of human physiology and health, but on the empirical adequacy of its clinical outcomes.

Predictive Diagnosis

The concept of having to gather clinically relevant data about patients in order to be able to classify them for purpose of personalised medicine, may bring up the idea that under this system people are increasingly permanently being perceived as patients, or at least patients-to-be. This shift towards permanent patient-ship would be brought even further by the notion of predictive diagnosis. A functionality that would not merely require data in order to put patients into slightly more refined categories, but a functionality that requires a continuous data collection to be in place from which inferences of near future states and conditions can be made.

It must be mentioned however that also under current medical practice, this shift towards permanently being a (potential) patient, is already taking place with small steps in the form of screening programmes for people that fall, on the basis of relatively simple criteria such as sex and age, into groups of higher risks for certain diseases. This ranges from cancer screening to infectious disease and from relatively small groups to the wider population, but in all cases it departs from the idea that they are potentially patients, regardless of the absence of clear symptoms or complaints.

Furthermore, predictions and anticipating future states of patients seems to be common practice already under current medical practice. Clinicians routinely make estimates of how a patients conditions shall improve or deteriorate in the absence or presence of some treatment, and make treatment choices accordingly. Some of these may merely be based on hunches, fed by a clinician's

experience, but others are based on evidence-based correlations between certain stages or characteristics and conditions at a later stage. Take for example the way in which genetic screening programmes make estimates of certain patients being at risk of ultimately conceiving a particular condition (Daly et al., 2016; Lyssenko & Laakso, 2013; Moyer, 2014; Tikkanen, Havulinna, Palotie, Salomaa, & Ripatti, 2013). More direct forms can be found as well, such as how an HIV infection, if untreated, will quite possibly lead to the condition of AIDS, and there is relatively little uncertainty that cancer moves through its stages without adequate treatment.

Neither the fact that current medicine already involves a fair bit of permanent patient-ship nor that it already involves aspects of anticipation and prediction would legitimise the idea that nothing much changes with the introduction of ICDSS. Though the basic elements may be already to some extent be in place, a fully predictive ICDSS functionality would have more fundamental issues attached to it.

What sets an ordinary diagnosis apart from a predictive diagnosis is the fact that the predictive diagnosis' adequacy can only be verified with the passing of time. There are after all no medical diagnostic instruments that can perceive or visualise the future. It is then the clinician's knowledge of the relation between the current state and the passing of time that allows for predictions to be made about the future states of the patients under current medicine, and the ICDSS's insight into that same relation analogously. In this regard there does not seem to be much of a difference between current medical practice and ICDSS when it concerns predictive diagnoses. After all, it seems that the only way to really be sure that HIV develops into AIDS, and a stage 1 cancer developing into stage 2, is by simply waiting for it to become the patient's current state, and confirming it then with the available diagnostic instruments.

However, what does set them apart, however, is the fact that predictions based on ICDSS may be much less clearly linked to the patient's current clinical state. Having HIV or stage 1 cancer is, reasoning from current medicine, ultimately worlds apart from having an inconspicuous heart rate pattern during a morning run, eating certain food in combination with certain genes, or any other pattern within some quantified aspect of a person's life; as hypothesised in the previous chapter. What is so different about it is that HIV leading to AIDS and stage 1 leading to stage 2 cancer can perfectly well be supported and predicted using the dominant causal-mechanistic theories explaining both pathophysiologies; facilitating causal reasoning to legitimise and support the clinician's experience with the relation between current and future states and subsequent predictions and anticipation. In other words, under current medical practice, predictions are mostly dependent on existing understandings of the influence of time on a causal mechanistic processes that may often have been verified using scientific methods, and that may be relatively broadly shared, whereas ICDSS exploits clinically seemingly unrelated aspects of patient's lives that may not be 'verifiable' using dominant causal-mechanistic theories.

What may furthermore be very different is the extent in which the patterns that are being exploited for prediction can be generalised for the larger patient population. With the increasing role for the personalisation of medicine, it seems reasonable that these predictions will come to be made for smaller groups of patients under ICDSS, also within current medicine's classification systems. Making predictions being more personalised, and therewith assumingly less comprehensible and transparent.

Need for standardisation

An important requirement that Jaspers in the second chapter of this thesis was shown to have expressed, was that there is a need for standardisation, in service of making the most of the available data; and for ICDSS to work with data across sources and systems. The reason for Jaspers for bringing this up is the fact that patient records are written down by clinicians in a very wide variety of ways. Every clinician may have his or her own way of keeping patient records. The problem Jaspers identifies with this, is that this makes it hard for ICDSS to recognise what types of information stored in these records can be compared across records. In other words, writing similar information down in different ways may obscure the fact that they should be supposed to be similar enough for them to be compared with each other.

This is seemingly unproblematic when it concerns data that is quantified using the same instruments, but recorded in different formats. Under such circumstances it matters little whether a record shows an average of *110BPM*, *HR110*, *Heart Rate 110*, *A heart rate of one hundred and ten* or any other way of writing down the same quantitative determination. As long as its use is consistent across records this all means precisely the same thing, especially when the same methods have been used to determine them.

Medical practice however, as should be clear by now, is not merely a discipline of objective facts and quantitative data. For a large part concerns subjective interpretations and observations that clinicians do. These more subjective pieces of information are also routinely recorded in patient records. The position that standardisation would be helpful in general is seemingly under the assumption that the way of writing things down is irrelevant with regard to what is being described.

Instead one could argue that the way in which something is written down is of great significance to the epistemic content of written information. Limiting the variety of the ways things can to be described by enforcing a strict regime of standardisation may also limit the clinician's ability to write things down in the way that best reflects their subjective experiences of a patient. Think for example about the influence of standardising clinical language and the influence that would have on the processes of narrative medicine, that are so dependent on verbal language exchange. Or clinicians having to write the impression a patient makes, the way they enter the room, shake hands, speak, the posture a patient presents or the pitch of their voice down in standardised categories or using a pre-determined vocabulary. The issue here is that standardisation on the one hand indeed helps to enable functionalities of ICDSS, but on the other hand paradoxically implies that the recorded information will provide poorer information about patients.

DISCUSSION & CONCLUSION

Introduction:

This thesis has been an attempt to provide insight into what the difference are between current medical practice and Intelligent Clinical Decision Support Systems in terms of the way in which knowledge and guidelines are produced and clinical decisions are being made, and what the issues are with regard to their compatibility in case of integrating Intelligent Clinical Decision Support Systems into current medical practice.

In doing so, this thesis has produced a characterisation as a heuristic to understand and analyse current medical practice and similarly characterised the feasible and probable functionalities of ICDSS. In the discussion of the previous chapter these were used to identify the differences and (epistemological) issues this thesis had been aiming to identify.

In this conclusion we shall briefly review these issues and explore whether broader themes can be identified. As mentioned at the beginning of this thesis, the discussion will also make a brief excursion to discuss the sensibility of the assumption that compatibility should be the goal here, or whether alternative outcomes may be more desirable.

Providing answers and conclusion to the sub questions

The first chapter, aiming to provide a characterisation of current medical practice, argued that the art-versus-science distinction, that is so omnipresent among characterisations of clinical practice, fails to provide insight into the different types and kinds of knowledge that are being produced and used in clinical practice and fails to do justice to the variety of methods in which art and science cannot be separated. Instead, this thesis constructed a characterisation of clinical practice drawing from Miriam Solomon's very detailed and elaborate collection of methods by 'summarising' it to a small number of primary methods constituting the tidy methodological pluralism. This placed Narrative Medicine and Evidence-Based Medicine at the centre of decision making under current medical practice, and placed Evidence-Based Medicine, Consensus Conferences and Translational Medicine at the centre of knowledge production under current medical practice.

The second and third chapter characterised ICDSS on the basis of historical examples and perspectives, expectations from businesses and scholars and scientists on the macro and micro level. Assuming its technological feasibility, the characterisation described an ICDSS as a collection of autonomous functionalities integrated into an elaborate patient management system, using big data technologies and machine learning to enable knowledge production and application on the basis of patterns and correlations, extraction of knowledge from written sources, using real-time monitoring both within and outside of the hospital walls in order to provide clinicians real-time assistance, both responsively as well as predictively.

General Overview

Then, most importantly to this thesis, these characterisations and the underlying understanding of the methods of current medicine and ICDSS's functionalities were used to come to an understanding of the differences between ICDSS and current medical practice in terms of knowledge production and decision-making and to identify issues with regard to their compatibility, both on a general level as well as on a method-specific level.

This led us first of all to argue that ICDSS introduces issues with regard to the transparency for a number of its functionalities. On the one hand this issue emerges from the assumed capability of ICDSS to analyse and effectively process written sources. This thesis argued for the impossibility of making transparent how suggestions on the basis of these written sources are formed by the ICDSS in the decision-making process.

On the other hand issues emerged in the knowledge production process due to ICDSS's dependence of on data mining and machine learning and their exploitation of complex correlations and pattern. Within both knowledge production and decision-making the enormous amount of available data is a crucial component of the improbability of seeing through the system's processes.

At the same time it has been argued that medicine has a long tradition of black-boxing technologies and methods, and that transparency is not necessarily a necessary condition for current medical practice to function. In other words, though current medical practice may very well be more transparent at this point than ICDSS suggests to be, it is not necessary a reason to presuppose incompatibility between the two. It was argued instead that generally trust seems to replace the need for transparency. This, and the inherent opaqueness of big data processing, makes us aware that for ICDSS to be useful within medical practice according to our characterisation its autonomous functionalities need to be trusted by its user, whether it deserves it or not.

Trust does however not eliminate the existence of epistemological issues. Dealing with scientific controversy provides a few of these. First of all it has been shown that ICDSS does not possess the potential of having some logical operation for resolving scientific controversy. It could be programmed to do so anyway, on the basis of some performance measure, but that provides an algorithmic foundation for the resolution of controversy, rather than the social foundation it has been shown to have under current medical practice. The unexplored question here is how the method of resolving controversy entails particular epistemic values to be prioritised over others; possibly deeply influencing how and what knowledge is established. Shifting values could thus have a major impact on how controversy is resolved, and subsequently demanding medical practice to adapt to new epistemic values and products.

It has been argued that, with regard to epistemic values, that veracity, an epistemic value so highly valued among scientific realists, cannot be determined by ICDSSs due to its inability to access reality. Instead much of ICDSS's intelligent features are founded in mere empirical adequacy. This thesis has left unanswered how significant scientific realism is for medical practice, but the prominent place of causal-mechanistic models within medicine's, as has been shown in this thesis, suggests that it may at least be quite widely adopted.

It has however also been shown that this problem could theoretically perhaps be avoided, were ICDSS and human clinicians to cooperate in the resolution of controversy by the system suggesting a consensus conference if a social epistemology is desired. The problem with this is that much scientific controversy may be resolved informally, between clinicians, outside the walls of a conference room. Implying that there is a lot more controversy among literature then the frequency of consensus conferences suggests.

This thesis furthermore argued that the possible need for standardisation may cause issues with regard to the clinician's ability to express their clinically relevant observations about patients, leading to an impoverishment of patient records. On the other hand, the discussion of issues concerning personalised knowledge showed that a wider variety of information about patients may become relevant clinical data; depending on whether they will be or become part of meaningful correlations.

This discussion in the previous chapter made numerous references to causal-mechanistic models and reasoning in explaining some of the differences between current medical practice and ICDSS. For example when arguing for the difference between current predictive diagnosis and under ICDSS being positioned in the legitimisation using causal-mechanistic understanding of conditions. A similar case was made to explain how personalised medicine under ICDSS would be different from current practices; by stratification on the basis of seemingly causally unrelated aspects of patients. This is an interesting fact, given that this thesis largely followed Miriam Solomon's classification as a traditional method; suggesting that its relevance for current medicine would perhaps be smaller than the current methods. This suggestion seems to be unjustified, given that the role of causal-mechanistic it has provided us with ability to explain the differences between some aspects of ICDSS and current medical practice.

Reflection on Methods

In the second part of the previous chapter this thesis identified some additional differences and issues on the methodological level. It first of all identified that ICDSS would be unable to integrate the full set of competences that Narrative Medicine requires in order to actually consider the patient narrative in treatment considerations. As a consequence this would limit ICDSS's to diagnostic suggestions, and when assisting with treatment options, to providing a complete range of theoretically possible options; leaving out any attempt to come to a definitive conclusion because that would require the integration of narrative aspects.

ICDSS has been shown to necessarily be an imperfect medium for EBM, mostly because of methodological shortcomings of EBM itself. Though EBM provides an evidential hierarchy it does not provide a logical solution for weighing large collections of evidence arguing slightly different things. Further epistemic issues have been raised with regard to how the knowledge generated by EBM's knowledge production abilities should be placed into the evidential hierarchy. An issue that can certainly be resolved in one way or another. The question then remains what the relevance of the evidential hierarchy is to the ICDSS as it may not be able to apply it to its knowledge discovery processes, but that needs attention in order to be able to employ the principles of EBM, as current medical practice has been characterised to demand.

As a consequence of not having a definitive way to choose between scientific explanations or results in the resolution of controversy, ICDSS may be developed in such a way to keep human clinicians and experts in charge of the resolution of scientific controversy.

The previous chapter has furthermore been reserved with regard to the usefulness of ICDSS in translational medicine. Translational processes require skill and expertise rather than merely being notified of the existence of some possibly relevant knowledge. ICDSS could however be supportive by facilitating efforts in the planning and execution of clinical trials and other additional clinical research.

Trust and Responsibility

The role of trust has been discussed briefly in this thesis. These discussions have however mostly focussed on the individual clinician's trust in the outcomes of ICDSS's processes. Though this is definitely a crucial aspect of the integration of ICDSS into medical practice, the clinician functions within a larger system. The broader medical practice, the public opinion, laws and legislations etcetera. What this may mean is illustrated by what Jasper's (2011) was shown to foresee as an interaction between the clinician's decision making and ICDSS in the second chapter of this thesis. Jasper's argues that CDSS, apart from merely providing suggestions and assistance, could be used to monitor the clinician's decision making, notifying the clinician when he or she deviates from the recommended guidelines and protocols.

This means that the question of whether an individual clinician trusts the ICDSS's outcomes is seemingly subordinate in relation to the question whether the relevant powers within the system trust the ICDSS's outcomes. Given that ICDSS (as it has been characterised) would be able to keep track of how clinicians diagnose and provide care to patients, it would become seemingly unattractive for clinicians to base their decisions on factor other than the one's provided by the ICDSS's trusted functionalities. Depending on the precise level of trust that authorities place in particular functionalities of ICDSS, scenarios can be imagined in which not acting in accordance with ICDSS's outcomes may be equal to a violation of the standard of medical care and perhaps malpractice. This can in fact already be the case with current medicine's guidelines, the difference however, is that ICDSS may keep a much more detailed record of a clinician's acting and be leaving less room for discussion with regard to what guideline applies if it prompts an immediate answer (Hurwitz, 1999; Mello, 2001; Woolf, Grol, Hutchinson, Eccles, & Grimshaw, 1999).

This would create a clinical practice in which clinicians would need to be very certain that the ICDSS is wrong in order for it to make sense for them to act against the ICDSS's suggestions . This level of certainty may however be difficult to reach. The recurring issues of transparency may make it difficult to falsify an ICDSS's conclusions; making it hard to be certain about not following the ICDSS. In this form of clinical practice, one could argue, the clinician is no longer actually in charge; regardless of whether he or she is formally responsible for the decision-making, as long as authoritative powers are convinced that trust should be placed in ICDSS's capabilities.

This goes against the continuous reassurances of commercial parties that clinicians remain in charge of patient care and that ICDSS is only there to help them. These assurances hold as long as the human clinician remains the site where trust is placed. Most likely as long as the ICDSS's abilities fail to be more accurate and reliable than human clinicians.

There is little certainty about ICDSS's precise potential in medicine, but even if a few of the functionalities that this thesis proposed would only be half-way met it will most likely imply redistributions of trust and subsequently responsibility.

Compatibility, and then?

Finally there is one question that has yet been left unanswered. The question of whether compatibility with current medical practice should be a desired quality for ICDSS, or whether alternatives exist and should be preferred. Positioning compatibility as a value is an odd thing to do, as few would recognise the intrinsic value of things being compatible. Perhaps it is more of an aesthetic value, something that one may pursue without clear functional or consequential reasons. However, in the context of this conclusion, it shows how throughout this thesis an implicit assumption has been made; namely that compatibility would make for a painless integration of the one into the other, without either one having the need to adjust.

Reflecting on the thesis, a naive position, as the hybrid between ICDSS and current medical practice would always entail a completely different mode of medical practice, regardless of whether the methods of the tidy methodological pluralism would remain uninfluenced.

This thesis has however shown that these methods would most likely not remain uninfluenced. The implementation of ICDSS into medical practice would change the significance of the methods of decision-making and knowledge production in relation to each other, and may actually change the way in which the methods are applied and understood.

It has first of all, for example, been argued how narrative medicine would be unable to uphold its normative ideals if ICDSS were to take over aspects of its tasks, as narrative medicine's criteria cannot be quantified. As an example it was shown how using patient satisfaction, a seemingly narrative aspect of medical practice, fails to solve this. However, this merely fails to solve this in terms of the normative ideal. What could happen is that the understanding, or at least the performance, of narrative medicine may actually come to be shaped by measuring and weighing factors such as patient satisfaction into clinical decision-making. This would then indeed lead to a more hermeneutical medicine, albeit not perfectly in accordance with the current ideals that narrative medicine upholds. This may in fact even make narrative more prominently featured in relation to EBM in terms of methods of decision-making, as it may be experienced to be much easier than acting in accordance with the actual competences that have been put forward for performing narrative medicine.

Furthermore, it has been shown that aspects of what belongs to EBM in the decision-making practice may also be strengthened when integrated with ICDSS. In particular the focus on standardised guidelines and rule-following may become increasingly prominently featured within EBM as ICDSS would be perfectly capable of keeping track of the extent to which clinicians stick to their guidelines. Assuming that ICDSS will formulate increasingly many rules on the basis of 'best available evidence' it can easily be imagined that the actual decision-making would largely be delegated to applying the right guideline; rather than weighing the best available evidence by the clinician his or herself. As a final example, it has been briefly mentioned how consensus conferences could be used within a hybrid with ICDSS in order to deal with the unresolvable controversies that ICDSS will discover in its search to aggregate and process all sorts of written literature. If this is a solution that would work, it could become a crucial aspect of allowing ICDSS to properly draw conclusions from medical literature.

So perhaps the question of compatibility is not as much about whether ICDSS and current medical practice can be integrated into each other, it is instead about whether they should. So instead

of asking whether compatibility is a desirable quality, it rather asks whether integration is desirable in general. So assuming that ICDSS and current medical practice are at least partially compatible, as this thesis has suggested, the question of alternatives to integration is reduced to two options. Not making use of ICDSS at all, or replacing current medical practice for a medical practice that is completely built around ICDSS.

It seems sensible to suppose that the reasons for preferring either scenario, or arguing why they ought to be preferred, may be very diverse. Though this thesis has not aimed at elucidating these reasons, it seems that some key elements have been identified that may play crucial roles in these preferences.

First of all the role of trust can be a crucial component in either scenario, seemingly as a necessary condition. After all, if one does not trust ICDSS's clinical adequacy it makes sense not to want ICDSS to be implemented into current clinical practice and no sense to want it to replace current medical practice.

Second, one may ascribe intrinsic value to aspects of current medical practice that would degrade or disappear in cases where ICDSS would be integrated into clinical practice, or replace it altogether. Depending on how high these aspects are valued one may come to conclude that there should be no place for ICDSS in clinical practice. Opposite scenarios may equally well exist, in which value is ascribed to aspects of ICDSS that that would be satisfied less in scenarios of exclusion or compatibility.

One of these aspects in both scenarios could for example be their ability to provide good clinical care. One may argue that one provides better care over the other, and better care over a hybrid of both methods. Reasons for this may include highly valuing social epistemic processes or the value of patient narrative, or oppositely algorithmic epistemology and the power of machine learning.

Another possibly quite influential factor could simply be cost. The cost of medical care is of great concern in many societies where investment in medical care is always a trade-off with other values a society maintains. Depending on whether ICDSS corresponds with a decrease or an increase in cost one may choose to respectively wish to stimulate or avoid implementation.

In addition to this it is relevant to reflect on the fact that the characterisation of current medical practice as it has been used in this thesis may be of limited value to medical care in some parts of the world. It has been mentioned somewhere early in this theses that the focus predominantly lies on medical practice in Western European, North American and countries, regions or institutions with similarly organised medical care. There are however many regions in the world where medical care is organised radically different. Severe scarcity of money and materials play crucial roles in some of these sites. But perhaps more importantly for the potential of ICDSS, so does scarcity of expertise. Training clinicians and enabling them to form the professional community that current medical practice requires them to be is not always among the possibilities on the short term.

Though this would not really involve a replacement or compatibility of current medical practice, as what medical practice is in those cases differs greatly from how it has been characterised in this thesis, a scenario in which ICDSS would be brought in to aid localised medical practices would be involve a very different distribution of values and interests and a very different perspective on compatibility could be seen as desirable.

It may finally be relevant to reflect on the implicit suggestion that discussing the desire for current medical practice to be replaced by ICDSS gives off. Primarily that it seemingly gives off is that it is reasonable to think that it is indeed possible for current medical practice to be replaced by ICDSS.

Though I have been hesitant writing strong conclusions about this down, and though it is not precisely clear what 'replacement' entails; the findings that this thesis made quite clearly suggest that ICDSS—as it has been characterised in this thesis— does not have the potential to be a complete replacement of medical practices, current medical practice included.

There are assumingly many aspects of medical practice that could be taken away without changing the essence of health-care. Decision-making could for example easily be done without narrative medicine or EBM. It would be a different type of decision-making, probably with a great impact on the effectiveness of healthcare, but it would be medical decision-making nonetheless. And similarly could consensus conferences and translational medicine be removed. Assuming the effectiveness of ICDSS, it could be imagined that algorithmic alternatives could be used to take over the processes of decision making and knowledge production.

However, ultimately medical practice consists of more than merely knowledge production and decision-making. Though it has not been part of the characterisation of current medical practice, medicine has a very large component of manual labour. An ICDSS would perhaps be capable of making a therapeutic decision, but would not be capable of performing surgical interventions. And it would be capable of determining what additional blood tests need to be taken, but nor can it take the blood itself. Similarly is it unable to wash patients, transfer them to a wheel chair or treat minor injuries.

It is of course more than possible to develop machines and robots for all these physical aspects of clinical care. But that would require a fully integrated system of material and cognitive functionalities; something that seems much further into the future than the collection of intelligent functionalities that ICDSS currently seems to become.

At the same time increasing implementation of intelligent features into medical practice seems inevitable as some of the examples at the beginning of this thesis has shown. Implying that if replacement is not an option, and preventing implementation seeming an option of the past, that integration is the only option there truly is. The question is then not whether to desire compatibility and integration or not, but on how to integrate ICDSS and clinical practice in such a way that important values of both domains are identified, weighed against each other where they may be conflicting, and aiming to assure the materialisation of these values built into the design of the technologies and the protocols and rules surrounding them. The aim of this thesis has after all not been to defend or reject compatibility. It has aimed to provide a starting point in understanding the complex balancing by which this must ultimately take place.

REFERENCES

- Abdelfattah, M., & Fuji, R. (2008). Automatic text summarization. *World Academy of Science, Engineering and Technology*, 37(17), 192–195.
- Adams, I. D., Chan, M., Clifford, P. C., Cooke, W. M., Dallos, V., de Dombal, F. T., ... McIntyre, N. (1986). Computer aided diagnosis of acute abdominal pain: a multicentre study. *British Medical Journal (Clinical Research Ed.)*, 293(6550), 800–4. <http://doi.org/10.1136/bmj.293.6550.800>
- Agarwal, R., & Dhar, V. (2014). Editorial—Big data, data science, and analytics: The opportunity and challenge for IS research. *INFORMS*.
- Aggarwal, C. C. (2015). Mining Text Data. In *Data Mining: The Textbook* (pp. 429–455). Cham: Springer International Publishing. http://doi.org/10.1007/978-3-319-14142-8_13
- Alberts, B., Bray, D., Hopkin, K., Johnson, A., Lewis, J., Raff, M., ... Walter, P. (2013). *Essential cell biology*. Garland Science.
- Alpaydin, E. (2014). *Introduction to Machine Learning* (Third). MIT Press.
- Altman, D. G., & Royston, P. (2000). What do we mean by validating a prognostic model? *Statistics in Medicine*, 19(4), 453–473.
- Amin, A., & Thrift, N. (2003). Cities: Reimagining the urban. *Urban Policy and Research*, 21(2), 217.
- Anderson, C. (2008). The End of Theory: The Data Deluge Makes the Scientific Method Obsolete | WIRED. Retrieved November 8, 2016, from <https://www.wired.com/2008/06/pb-theory/>
- Auffray, C., Chen, Z., & Hood, L. (2009). Systems medicine: the future of medical genomics and healthcare. *Genome Medicine*, 1(1), 2. <http://doi.org/10.1186/gm2>
- Banaee, H., Ahmed, M. U., & Loutfi, A. (2013). Data mining for wearable sensors in health monitoring systems: a review of recent trends and challenges. *Sensors*, 13(12), 17472–17500. <http://doi.org/10.3390/s131217472>
- Banning, M. (2008). A review of clinical decision- making : models and current research . *Journal of Clinical Nursing*, 17(2), 187–195. <http://doi.org/10.1111/j.1365-2702.2006.01791.x>
- Baraniuk, C. (2016). Google DeepMind targets NHS head and neck cancer treatment - BBC News. Retrieved April 22, 2017, from <http://www.bbc.com/news/technology-37230806>
- Barrows, H. S., Norman, G. R., Neufeld, V. R., & Feightner, J. W. (1982). The clinical reasoning of randomly selected physicians in general medical practice. *Clinical and Investigative Medicine*, 5(1), 49–55.
- Battista, R. N., Hodge, M. J., & Vineis, P. (1995). Medicine, practice and guidelines: the uneasy juncture of science and art. *Journal of Clinical Epidemiology*, 48(7), 875–880.
- Bell, I. R., Caspi, O., Schwartz, G. E. R., Grant, K. L., Gaudet, T. W., Rychener, D., ... Weil, A. (2002). Integrative Medicine and Systemic Outcomes Research Issues in the Emergence of a New Model for Primary Health Care. *Arch Intern Med*, 162, 133–140. <http://doi.org/10.1001/archinte.162.2.133>
- Bellomo, R., Kellum, J. A., & Ronco, C. (2012). Acute kidney injury. *The Lancet*, 380(9843), 756–766. [http://doi.org/10.1016/S0140-6736\(11\)61454-2](http://doi.org/10.1016/S0140-6736(11)61454-2)

- Benner, P. (1982). From novice to expert. *AJN The American Journal of Nursing*, 82(3), 402–407.
- Benner, P., & Tanner, C. (1987). How expert nurses use intuition. *AJN The American Journal of Nursing*, 87(1), 23–34.
- Bennett, C. C., & Hauser, K. (2013). Artificial intelligence framework for simulating clinical decision-making: A Markov decision process approach. *Artificial Intelligence in Medicine*, 57(1), 9–19. <http://doi.org/10.1016/j.artmed.2012.12.003>
- Bughin, J., Chui, M., & Manyika, J. (2010). Clouds, big data, and smart assets: Ten tech-enabled business trends to watch. *McKinsey Quarterly*, 56(1), 75–86.
- Burke, W., & Psaty, B. M. (2007). Personalized medicine in the era of genomics. *Jama*, 298(14), 1682–1684.
- Cardon, L. R., & Palmer, L. J. (2003). Population stratification and spurious allelic association. *The Lancet*, 361(9357), 598–604.
- Carnevali, D., & Thomas, M. (1993). *Diagnostic Reasoning and Treatment Decision Making in Nursing*. Philadelphia: Lippincott.
- Chapelle, O., Scholkopf, B., & Zien, A. (2009). Semi-supervised learning (chapelle, o. et al., eds.; 2006)[book reviews]. *IEEE Transactions on Neural Networks*, 20(3), 542.
- Charon, R. (2001). Narrative Medicine. *Jama The Journal Of The American Medical Association*, 286(15), 1897–1902. <http://doi.org/10.1001/jama.286.15.1897>
- Chen, H., Chiang, R. H. L., & Storey, V. C. (2012). Business intelligence and analytics: From big data to big impact. *MIS Quarterly*, 36(4), 1165–1188.
- Chen, R., & Snyder, M. (2012). Systems biology: Personalized medicine for the future? *Current Opinion in Pharmacology*, 12(5), 623–628. <http://doi.org/10.1016/j.coph.2012.07.011>
- Cheng, C. K., Chan, J., Cembrowski, G. S., & Van Assendelft, O. W. (2003). Complete blood count reference interval diagrams derived from NHANES III: stratification by age, sex, and race. *Laboratory Hematology: Official Publication of the International Society for Laboratory Hematology*, 10(1), 42–53.
- Chu-Carroll, J., Brown, E. W., Lally, A., & Murdock, J. W. (2012). Identifying implicit relationships. *IBM Journal of Research and Development*, 56(3.4), 11–12.
- Cochrane, A. L. (1972). Effectiveness and efficiency: random reflections on health services. *London: England*.
- Crawford, K. (2013). The hidden biases in big data. *HBR Blog Network*, 1.
- Crowe, S., Cresswell, K., Robertson, A., Huby, G., Avery, A., & Sheikh, A. (2011). The case study approach. *BMC Medical Research Methodology*, 11, 100. <http://doi.org/10.1186/1471-2288-11-100>
- Daly, M. B., Pilarski, R., Axilbund, J. E., Berry, M., Buys, S. S., Crawford, B., ... Khan, S. (2016). Genetic/familial high-risk assessment: breast and ovarian, version 2.2015. *Journal of the National Comprehensive Cancer Network*, 14(2), 153–162.
- Dans, P. E. (1994). Credibility, cookbook medicine, and common sense: guidelines and the college. *Annals of Internal Medicine*, 120(11), 966–968.
- Davidoff, F., Haynes, B., Sackett, D., & Smith, R. (1995). Evidence based medicine. *BMJ: British Medical Journal*, 310(6987), 1085.
- De, S. K., Biswas, R., & Roy, A. R. (2001). An application of intuitionistic fuzzy sets in medical diagnosis. *Fuzzy Sets and Systems*, 117(2), 209–213.

- de Dombal, F. T., Leaper, D. J., Staniland, J. R., McCann, a P., & Horrocks, J. C. (1972). Computer-aided diagnosis of acute abdominal pain. *British Medical Journal*, 2(5804), 9–13.
<http://doi.org/10.1136/bmj.2.5804.9>
- DeepMind Health. (2017). Streams. Retrieved April 18, 2017, from
<https://deepmind.com/applied/deepmind-health/streams/>
- Delamothe, T. (1993). Wanted: guidelines that doctors will follow. *BMJ: British Medical Journal*, 307(6898), 218.
- DiMatteo, M. R. (1979). A social-psychological analysis of physician-patient rapport: Toward a science of the art of medicine. *Journal of Social Issues*, 35(1), 12–33.
- Domingos, P. (2012). A few useful things to know about machine learning. *Communications of the ACM*, 55(10), 78–87.
- Dougherty, J., Kohavi, R., & Sahami, M. (1995). Supervised and unsupervised discretization of continuous features. In *Machine learning: proceedings of the twelfth international conference* (Vol. 12, pp. 194–202).
- Dyche, J. (2012). Big data “Eurekas!” don’t just happen. *Harvard Business Review Blog*, 20.
- El-Khatib, M. F., Zeineldine, S., Ayoub, C., Husari, A., & Bou-Khalil, P. K. (2010). Critical care clinicians’ knowledge of evidence-based guidelines for preventing ventilator-associated pneumonia. *American Journal of Critical Care : An Official Publication, American Association of Critical-Care Nurses*, 19(3), 272–6. <http://doi.org/10.4037/ajcc2009131>
- Elstein, A. s. (2004). On the origins and development of evidence-based medicine and medical decision making. *Inflammation Research : Official Journal of the European Histamine Research Society ... [et Al.]*, 53 Suppl 2, S184–S189. <http://doi.org/10.1007/s00011-004-0357-2>
- Elstein, A., Shulman, L., & Sprafka, S. (1978). *Medical Problem Solving: An Analysis of Clinical Reasoning*. Harvard University Press (April 20, 1978).
- Esfandiari, N., Babavalian, M. R., Moghadam, A.-M. E., & Tabar, V. K. (2014). Knowledge discovery in medicine: Current issue and future trend. *Expert Systems with Applications*, 41(9), 4434–4463.
- Evidence-Based Medicine Working Group. (1992). Evidence-based medicine. A new approach to teaching the practice of medicine. *Jama The Journal Of The American Medical Association*.
<http://doi.org/10.1001/jama.1992.03490170092032>
- Fan, J., Kalyanpur, A., Gondek, D. C., & Ferrucci, D. A. (2012). Automatic knowledge extraction from documents. *IBM Journal of Research and Development*, 56(3.4), 1–5.
- Fayyad, U., Piatetsky-Shapiro, G., & Smyth, P. (1996). From Data Mining to Knowledge Discovery in Databases. *AI Magazine*, 17(3), 37. <http://doi.org/10.1609/aimag.v17i3.1230>
- Fayyad, U., & Uthurusamy, R. (2002). EVOLVING DATA MINING INTO SOLUTIONS FOR INSIGHTS. *Communications of the ACM*, 45(8), 28–31. <http://doi.org/10.1109/45.624335>
- Ferrucci, D., Brown, E., Chu-Carroll, J., Fan, J., Gondek, D., Kalyanpur, A. a., ... Welty, C. (2010). Building Watson: An Overview of the DeepQA Project. *AI Magazine*, 31(3), 59–79.
<http://doi.org/10.1609/aimag.v31i3.2303>
- Ferrucci, D., Nyberg, E., Allan, J., Barker, K., Brown, E., Chu-carroll, J., ... Heights, Y. (2009). Towards the Open Advancement of Question Answering Systems. *IBM Research Report*, 24789, RC24789.
- Frosch, D. L., & Kaplan, R. M. (1999). Shared decision making in clinical medicine: Past research and

- future directions. *American Journal of Preventive Medicine*, 17(4), 285–294.
[http://doi.org/10.1016/S0749-3797\(99\)00097-5](http://doi.org/10.1016/S0749-3797(99)00097-5)
- Gabbay, J., & le May, A. (2004). Evidence based guidelines or collectively constructed “mindlines?”
 Ethnographic study of knowledge management in primary care. *BMJ (Clinical Research Ed.)*,
 329(7473), 1013. <http://doi.org/10.1136/bmj.329.7473.1013>
- Gandomi, A., & Haider, M. (2015). Beyond the hype: Big data concepts, methods, and analytics.
International Journal of Information Management, 35(2), 137–144.
<http://doi.org/10.1016/j.ijinfomgt.2014.10.007>
- Gaudin, S. (2011). IBM’s Watson could usher in new era of medicine | Computerworld. Retrieved April
 17, 2017, from <http://www.computerworld.com/article/2513198/high-performance-computing/ibm-s-watson-could-usher-in-new-era-of-medicine.html?page=2>
- Ghahramani, Z. (2015). Probabilistic machine learning and artificial intelligence. *Nature*, 521(7553),
 452–459. Retrieved from <http://dx.doi.org/10.1038/nature14541>
- Ginsburg, G. S., & McCarthy, J. J. (2001). Personalized medicine: revolutionizing drug discovery and
 patient care. *TRENDS in Biotechnology*, 19(12), 491–496.
- Ginsburg, G. S., & Willard, H. F. (2009). Genomic and personalized medicine: foundations and
 applications. *Translational Research*, 154(6), 277–287.
<http://doi.org/10.1016/j.trsl.2009.09.005>
- Goffin, K., & Koners, U. (2011). Tacit knowledge, lessons learnt, and new product development.
Journal of Product Innovation Management, 28(2), 300–318. <http://doi.org/10.1111/j.1540-5885.2010.00798.x>
- Gordon, R., & Franklin, N. (2003). Cognitive underpinnings of diagnostic error. *Academic Medicine :
 Journal of the Association of American Medical Colleges*, 78(8), 782.
<http://doi.org/10.1097/00001888-200308000-00005>
- Graber, M. (2003). Metacognitive training to reduce diagnostic errors: ready for prime time?
Academic Medicine : Journal of the Association of American Medical Colleges, 78(8), 781.
<http://doi.org/10.1097/00001888-200308000-00004>
- Green, J., & Thorogood, N. (2009). *Qualitative Methods for Health Research. Introducing qualitative
 methods*. <http://doi.org/10.1177/1049732305285708>
- Greenhalgh, T., Howick, J., & Maskrey, N. (2014). Evidence based medicine: a movement in crisis?
Bmj, 348(jun13 4), g3725–g3725. <http://doi.org/10.1136/bmj.g3725>
- Grimshaw, J. M., & Russell, I. T. (1993). Effect of clinical guidelines on medical practice: a systematic
 review of rigorous evaluations. *The Lancet*, 342(8883), 1317–1322.
- Gunby, P. (1980). Pros and cons of NIH consensus conferences. *Journal of the American Medical
 Association*, 244(13), 1413–1414.
- Gurses, A. P., Marsteller, J. A., Ant Ozok, A., Xiao, Y., Owens, S., & Pronovost, P. J. (2010). Using an
 interdisciplinary approach to identify factors that affect clinicians’ compliance with evidence-
 based guidelines. *Critical Care Medicine*, 38(8 Suppl), S282–91.
<http://doi.org/10.1097/CCM.0b013e3181e69e02>
- Guyatt, G. H., Oxman, A. D., Vist, G. E., Kunz, R., Falck-, Y., Alonso-coello, P., & Schünemann, H. J.
 (2008). Rating Quality of Evidence and Strength of Recommendations : GRADE : An Emerging
 Consensus on Rating Quality of Evidence and Strength of Recommendations. *British Medical*

- Journal*, 336(April), 924–926. <http://doi.org/10.1136/bmj.39489.470347.AD>
- Hamburg, M. A., & Collins, F. S. (2010). The path to personalized medicine. *N Engl J Med*, 2010(363), 301–304.
- Hauskrecht, M., & Fraser, H. (2000). Planning treatment of ischemic heart disease with partially observable Markov decision processes. *Artificial Intelligence in Medicine*, 18, 221–244. [http://doi.org/10.1016/S0933-3657\(99\)00042-1](http://doi.org/10.1016/S0933-3657(99)00042-1)
- Hern, A. (2016). Google DeepMind pairs with NHS to use machine learning to fight blindness | Technology | The Guardian. Retrieved April 22, 2017, from <https://www.theguardian.com/technology/2016/jul/05/google-deepmind-nhs-machine-learning-blindness>
- Heymann, H. O., Swift Jr, E. J., & Ritter, A. V. (2014). *Sturdevant's art & science of operative dentistry* (6th ed.). St. Louis: Elsevier Health Sciences.
- Holleb, A. I. (1980). Editorial: Medicine by proxy? *CA: A Cancer Journal for Clinicians*, 30(3), 191–192.
- Howick, J., Glasziou, P., & Aronson, J. K. (2010). Evidence-based mechanistic reasoning. *Journal of the Royal Society of Medicine*, 103(11), 433–41. <http://doi.org/10.1258/jrsm.2010.100146>
- Huang, G., Song, S., Gupta, J. N. D., & Wu, C. (2014). Semi-supervised and unsupervised extreme learning machines. *IEEE Transactions on Cybernetics*, 44(12), 2405–2417.
- Hurwitz, B. (1999). Clinical guidelines: legal and political considerations of clinical practice guidelines. *BMJ: British Medical Journal*, 318(7184), 661.
- IBM Watson Health. (2016). *Helping Barrow Neurological Institute Discover Potential New Treatments for ALS - YouTube*. Retrieved from <https://www.youtube.com/watch?v=2oJNxvRvt6Y>
- IBM Watson Health. (2017a). *Product Vignette: IBM Watson for Oncology - YouTube*. Retrieved from https://www.youtube.com/watch?v=f_aSqh16jWo
- IBM Watson Health. (2017b). Watson Genomics from Quest Core - Helping your care team make informed decisions. Retrieved from <https://public.dhe.ibm.com/common/ssi/ecm/hl/en/hlb03015usen/HLB03015USEN.PDF>
- Jackson, C. L., Bolen, S., Brancati, F. L., Batts Turner, M. L., & Gary, T. L. (2006). A systematic review of interactive computer-assisted technology in diabetes care. *Journal of General Internal Medicine*, 21(2), 105–110.
- Jaspers, M. W. M., Smeulders, M., Vermeulen, H., & Peute, L. W. (2011). Effects of clinical decision-support systems on practitioner performance and patient outcomes: a synthesis of high-quality systematic review findings. *Journal of the American Medical Informatics Association : JAMIA*, 18(3), 327–34. <http://doi.org/10.1136/amiajnl-2011-000094>
- Johnson, A. E. W., Ghassemi, M. M., Nemati, S., Niehaus, K. E., Clifton, D., & Clifford, G. D. (2016). Machine Learning and Decision Support in Critical Care. *Proceedings of the IEEE*, 104(2), 444–466. <http://doi.org/10.1109/JPROC.2015.2501978>
- Kalyanpur, A., Boguraev, B. K., Patwardhan, S., Murdock, J. W., Lally, A., Welty, C., ... Zhang, L. (2012). Structured data and inference in DeepQA. *IBM Journal of Research and Development*, 56(3.4), 10–11.
- Kantrowitz, A. (1967). Proposal for an Institution for Scientific Judgment. *Science*, 156(3776), 763 LP–

764. Retrieved from <http://science.sciencemag.org/content/156/3776/763.abstract>
- Kantrowitz, A. (1977). THE SCIENCE COURT EXPERIMENT. *Jurimetrics Journal*, 17(4), 332–341.
Retrieved from <http://www.jstor.org/stable/29761595>
- Kaplan, B. (2001). Evaluating informatics applications—clinical decision support systems literature review. *International Journal of Medical Informatics*, 64(1), 15–37.
- Kassirer, J. P., & Gorry, A. G. (1985). Clinical Problem Solving: A Behavioral Analysis. In *Computer-Assisted Medical Decision Making* (pp. 85–107). New York: Springer-Verlag.
- Kellum, J. A. (2008). Acute kidney injury. *Critical Care Medicine*, 36(4), 141–145.
<http://doi.org/10.2310/7900.1164>
- Khatibi, V., & Montazer, G. A. (2010). A fuzzy-evidential hybrid inference engine for coronary heart disease risk assessment. *Expert Systems with Applications*, 37(12), 8536–8542.
- King, L., & Clark, J. M. (2002). Intuition and the development of expertise in surgical ward and intensive care nurses. *Journal of Advanced Nursing*, 37(4), 322–329.
- Kitano, H. (2002). Systems biology: a brief overview. *Science (New York, N.Y.)*, 295(5560), 1662–4.
<http://doi.org/10.1126/science.1069492>
- Kitchin, R. (2014a). Big Data, new epistemologies and paradigm shifts. *Big Data & Society*, 1(1), 2053951714528481. <http://doi.org/10.1177/2053951714528481>
- Kitchin, R. (2014b). *The Data Revolution*. (R. Rojek & K. Dickens, Eds.). Sage Publishing.
- Kovacs, G., & Croskerry, P. (1999). Clinical decision making: An emergency medicine perspective. *Academic Emergency Medicine*, 6(9), 947–952. <http://doi.org/10.1111/j.1553-2712.1999.tb01246.x>
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions* (Vol. I and II). Chicago: The University of Chicago Press. <http://doi.org/10.1119/1.1969660>
- Kunjunnair, A. P. (2011). Clinical decision support system: risk level prediction of heart disease using weighted fuzzy rules and decision tree rules. *Central European Journal of Computer Science*, 1(4), 27–40. <http://doi.org/10.2478/s13537-012-0007-7>
- Kuperman, G. J., Maack, B. B., Bauer, K., & Gardner, R. M. (1991). Innovations and research review: the impact of the HELP computer system on the LDS Hospital paper medical record. *Topics in Health Record Management*, 12(2), 76–85.
- Lally, A., & Fodor, P. (2011). Natural language processing with Prolog in the IBM Watson system. *The Association for Logic Programming (ALP)*, 1–4. Retrieved from <http://www.cs.nmsu.edu/ALP/wp-content/uploads/2011/03/PrologAndWatson1.pdf>
- Lam, A. (2000). Tacit knowledge, organizational learning and societal institutions: An integrated framework. *Organization Studies*, 21(3), 487–513.
- Latour, B. (1999). *Pandora's hope: essays on the reality of science studies*. Harvard university press.
- Ledley, R. S., & Lusted, L. B. (1959). The Use of Electronic Computers to Aid in Medical Diagnosis. *Proceedings of the IRE*, 47, 183–188. <http://doi.org/10.1109/JRPROC.1959.287213>
- Leonelli, S. (2012). Introduction: Making sense of data-driven research in the biological and biomedical sciences. *Studies in History and Philosophy of Biol & Biomed Sci*, 43(1), 1–3.
<http://doi.org/10.1016/j.shpsc.2011.10.001>
- Leonelli, S. (2014). What Difference Does Quantity Make? On the Epistemology of Big Data in Biology. *Big Data & Society*, 1(1), 2053951714534395. <http://doi.org/10.1177/2053951714534395>
- Li, E. Y. (1994). Artificial neural networks and their business applications. *Information &*

- Management*, 27(5), 303–313.
- Lin, P. Y., Yu, S. L., & Yang, P. C. (2010). MicroRNA in lung cancer. *British Journal of Cancer*, 103(8), 1144–1148.
- Lindon, J. C., & Nicholson, J. K. (2014). The emergent role of metabolic phenotyping in dynamic patient stratification. Taylor & Francis.
- Lisboa, P. J., & Taktak, A. F. G. (2006). The use of artificial neural networks in decision support in cancer: A systematic review. *Neural Networks*, 19(4), 408–415.
<http://doi.org/10.1016/j.neunet.2005.10.007>
- Lynch, C. (2008). Big data: How do your data grow? *Nature*, 455(7209), 28–29.
- Lyssenko, V., & Laakso, M. (2013). Genetic screening for the risk of type 2 diabetes. *Diabetes Care*, 36(Supplement 2), S120–S126.
- Mack, E. H., Wheeler, D. S., & Embi, P. J. (2009). Clinical decision support systems in the pediatric intensive care unit. *Pediatric Critical Care Medicine*, 10(1), 23–28.
- Maienschein, J., Ankeny, R. A., & Robert, J. S. (2008). The Ethos and Ethics of Translational Research. *The American Journal of Bioethics*, 8(3), 43–51.
<http://doi.org/10.1080/15265160802109314>
- Malterud, K. (2001). The art and science of clinical knowledge: evidence beyond measures and numbers. *The Lancet*, 358(9279), 397–400.
- Mankoff, S. P., Brander, C., Ferrone, S., & Marincola, F. M. (2004). Lost in Translation: Obstacles to Translational Medicine. *Journal of Translational Medicine*, 2, 14. <http://doi.org/10.1186/1479-5876-2-14>
- Marx, V. (2013). Biology: The big challenges of big data. *Nature*, 498(7453), 255–260.
- Meador, C. K. (1965). The art and science of nondisease. *New England Journal of Medicine*, 272, 92–95.
- Mello, M. M. (2001). Of swords and shields: the role of clinical practice guidelines in medical malpractice litigation. *University of Pennsylvania Law Review*, 149(3), 645–710.
- Miles, M. B., & Huberman, A. M. (1994). *Qualitative Data Analysis. An expanded sourcebook* (Vol. 3). <http://doi.org/10.1080/0140528790010406>
- Miller, R. (1991). Medical Diagnostic Decision Support Systems -Past , Present , and Future : Work on Human Diagnostic Problem Solving.
- Mitchell, C., Payne, J., Wade, R., Vora, A., Kinsey, S., Richards, S., & Eden, T. (2009). The impact of risk stratification by early bone-marrow response in childhood lymphoblastic leukaemia: results from the United Kingdom Medical Research Council trial ALL97 and ALL97/99. *British Journal of Haematology*, 146(4), 424–436.
- Mitchell, T. M. (1997). Machine learning. 1997. *Burr Ridge, IL: McGraw Hill*, 45(37), 2.
- Mollon, B., Chong, J. J. R., Holbrook, A. M., Sung, M., Thabane, L., & Foster, G. (2009). Features predicting the success of computerized decision support for prescribing: a systematic review of randomized controlled trials. *BMC Medical Informatics and Decision Making*, 9(1), 11.
- Monin, J.-L., Monchi, M., Gest, V., Duval-Moulin, A.-M., Dubois-Rande, J.-L., & Gueret, P. (2001). Aortic stenosis with severe left ventricular dysfunction and low transvalvular pressure gradients: risk stratification by low-dose dobutamine echocardiography. *Journal of the American College of Cardiology*, 37(8), 2101–2107.

- Moyer, V. A. (2014). Risk assessment, genetic counseling, and genetic testing for BRCA-related cancer in women: US Preventive Services Task Force recommendation statement. *Annals of Internal Medicine*, 160(4), 271–281.
- Murdock, J. W., Fan, J., Lally, A., Shima, H., & Boguraev, B. K. (2012). Textual evidence gathering and analysis. *IBM Journal of Research and Development*, 56(3.4), 1–8.
- Musen, M. A., Shahar, Y., & Shortliffe, E. H. (2006). Clinical Decision-Support Systems. In J. H. Van Bommel & M. Musen (Eds.), *Handbook of Medical Informatics* (pp. 698–736). New York: Springer. Retrieved from <http://www.springerlink.com/content/w43jt5449q7801n2/>
- Nash, F. A. (1954). Differential Diagnosis An apparatus to assist the logical faculties. *The Lancet*, 874–875.
- Nash, F. A. (1960). Diagnostic Reasoning and the Logoscope. *The Lancet*, (Nash), 1442–1446.
- Nyman, I., Wallentin, L., Areskog, M., Areskog, N.-H., Swahn, E., & Group, R. S. (1993). Risk stratification by early exercise testing after an episode of unstable coronary artery disease. *International Journal of Cardiology*, 39(2), 131–142.
- Oliver, D., Britton, M., Seed, P., Martin, F. C., & Hopper, A. H. (1997). Development and evaluation of evidence based risk assessment tool (STRATIFY) to predict which elderly inpatients will fall: case-control and cohort studies. *Bmj*, 315(7115), 1049–1053.
- Oliver, D., Daly, F., Martin, F. C., & McMurdo, M. E. T. (2004). Risk factors and risk assessment tools for falls in hospital in-patients: a systematic review. *Age and Ageing*, 33(2), 122–130.
- Owens, C. D., Ho, K. J., Kim, S., Schanzer, A., Lin, J., Matros, E., ... Conte, M. S. (2007). Refinement of survival prediction in patients undergoing lower extremity bypass surgery: stratification by chronic kidney disease classification. *Journal of Vascular Surgery*, 45(5), 944–952.
- Pain, R. W. (1975). Limitations of the Nash Logoscope or diagnostic slide rule. *The Medical Journal of Australia*, 2(18), 714.
- Palaniappan, S., & Awang, R. (2008). Intelligent heart disease prediction system using data mining techniques. *2008 IEEE/ACS International Conference on Computer Systems and Applications*, 108–115. <http://doi.org/10.1109/AICCSA.2008.4493524>
- Patel, V. L., Shortliffe, E. H., Stefanelli, M., Szolovits, P., Berthold, M. R., Bellazzi, R., & Abu-Hanna, A. (2009). The coming of age of artificial intelligence in medicine. *Artificial Intelligence in Medicine*, 46(1), 5–17. <http://doi.org/10.1016/j.artmed.2008.07.017>
- Peabody, F. W. (2015). The care of the patient. *Jama The Journal Of The American Medical Association*, 313(18), 1868.
- Peek, N., Combi, C., Marin, R., & Bellazzi, R. (2015, September 1). Thirty years of artificial intelligence in medicine (AIME) conferences: A review of research themes. *Artificial Intelligence in Medicine*. Elsevier.
- Philip Chen, C. L., & Zhang, C. Y. (2014). Data-intensive applications, challenges, techniques and technologies: A survey on Big Data. *Information Sciences*, 275, 314–347. <http://doi.org/10.1016/j.ins.2014.01.015>
- Picciano, A. G. (2012). The Evolution of Big Data and Learning Analytics in American Higher Education. *Journal of Asynchronous Learning Networks*, 16(3), 9–20.
- Piquette-Miller, M., & Grant, D. M. (2007). The art and science of personalized medicine. *Clinical Pharmacology & Therapeutics*, 81(3), 311–315.

- Prensky, M. (2009). H. sapiens digital: From digital immigrants and digital natives to digital wisdom. *Innovate: Journal of Online Education*, 5(3), 1.
- Provost, F., & Fawcett, T. (2013). Data science and its relationship to big data and data-driven decision making. *Big Data*, 1(1), 51–59.
- Pryor, T. A., Gardner, R. M., Clayton, P. D., & Warner, H. R. (1983). The HELP System. *Journal of Medical Systems*, 7(2), 87–102.
- Ramesh, A. N., Kambhampati, C., Monson, J. R. T., & Drew, P. J. (2004, September). Artificial intelligence in medicine. *Annals of the Royal College of Surgeons of England*.
- Rawlins, M. (2008). De Testimonio: On the evidence for decisions about the use of therapeutic interventions. *Clinical Medicine, Journal of the Royal College of Physicians of London*, 8(6), 579–588. [http://doi.org/10.1016/S0140-6736\(08\)61930-3](http://doi.org/10.1016/S0140-6736(08)61930-3)
- Rigg, J. (2016). NHS to use Google DeepMind AI app to help treat patients. Retrieved November 28, 2016, from <https://www.engadget.com/2016/11/22/nhs-deepmind-ai-app/>
- Roberts, R., & Fromm, R. E. (1998). Management of acute coronary syndromes based on risk stratification by biochemical markers. *Am Heart Assoc*.
- Rowley, R. (2016). Artificial Intelligence in healthcare | Flow Health. Retrieved December 6, 2016, from <https://flowhealth.com/blog/2016/01/laying-the-foundation-for-artificial-intelligence-in-health-care/>
- Sackett, D. L., Rosenberg, W. M. C., Gray, J. A. M., Haynes, R. B., & Richardson, W. S. (1996). Evidence based medicine: what it is and what it isn't. *BMJ*, 312(7023), 71–72. <http://doi.org/10.1136/bmj.312.7023.71>
- Sagiroglu, S., & Sinanc, D. (2013). Big data: A review. *2013 International Conference on Collaboration Technologies and Systems (CTS)*, 42–47. <http://doi.org/10.1109/CTS.2013.6567202>
- Schmidt, F. L., & Hunter, J. E. (1993). Tacit Knowledge, Practical Intelligence, General Mental Ability, and Job Knowledge. *Current Directions in Psychological Science*, 2(1), 8–9. <http://doi.org/10.1111/1467-8721.ep10770456>
- Schoene, D., Wu, S. M., Mikolaizak, A. S., Menant, J. C., Smith, S. T., Delbaere, K., & Lord, S. R. (2013). Discriminative ability and predictive validity of the timed Up and Go test in identifying older people who fall: systematic review and meta-analysis. *Journal of the American Geriatrics Society*, 61(2), 202–208.
- Schölkopf, B., & Smola, A. J. (2002). *Learning with kernels: support vector machines, regularization, optimization, and beyond*. MIT press.
- Shafi, S., Barnes, S. a, Millar, D., Sobrino, J., Kudyakov, R., Berryman, C., ... Nirula, R. (2014). Suboptimal compliance with evidence-based guidelines in patients with traumatic brain injuries. *Journal of Neurosurgery*, 120(3), 773–7. <http://doi.org/10.3171/2013.12.JNS132151>
- Sharir, T., Germano, G., Kang, X., Lewin, H. C., Miranda, R., Cohen, I., ... Berman, D. S. (2001). Prediction of myocardial infarction versus cardiac death by gated myocardial perfusion SPECT: risk stratification by the amount of stress-induced ischemia and the poststress ejection fraction. *Journal of Nuclear Medicine*, 42(6), 831–837.
- Shead, S. (2016). Google's £400 million acquisition of DeepMind is looking good - Business Insider. Retrieved November 28, 2016, from <http://uk.businessinsider.com/googles-400-million-acquisition-of-deepmind-is-looking-good-2016-7?international=true&r=UK&IR=T>

- Shin, H., & Markey, M. K. (2006). A machine learning perspective on the development of clinical decision support systems utilizing mass spectra of blood samples. *Journal of Biomedical Informatics*, 39(2), 227–248. <http://doi.org/10.1016/j.jbi.2005.04.002>
- Shortliffe, E. H., Davis, R., Axline, S. G., Buchanan, B. G., Green, C. C., & Cohen, S. N. (1975). Computer based consultations in clinical therapeutics: explanation and rule acquisition capabilities of the MYCIN system. *Computers and Biomedical Research*, 8(4), 303–320. [http://doi.org/10.1016/0010-4809\(75\)90011-7](http://doi.org/10.1016/0010-4809(75)90011-7)
- Siemens, G., & Long, P. (2011). Penetrating the fog: Analytics in learning and education. *EDUCAUSE Review*, 46(5), 30.
- Sim, I., Gorman, P., Greenes, R., Haynes, B., Kaplan, Bo., Lehmann, H., & Tang, P. (2001). Clinical decision support systems for the practice of evidence-based medicine. *Journal of the American Medical Informatics Association*, 8(6), 527–534.
- Simonite, T. (2011). Where Next for IBM's Watson? - MIT Technology Review. Retrieved April 17, 2017, from <https://www.technologyreview.com/s/422783/where-next-for-ibms-watson/>
- Smith, A. J. (2002). Applications of the self-organising map to reinforcement learning. *Neural Networks*, 15(8), 1107–1124.
- Solomon, M. (2003). Social Empiricism. *MIT Press*, 94(2), 422–423. <http://doi.org/10.1086/379487>
- Solomon, M. (2015). *Making Medical Knowledge*. Oxford: Oxford University Press.
- Stake, R. E. (1995). The art of case study research. *The Art of Case Study Research*. <http://doi.org/10.1108/eb024859>
- Stevens, L. (2017). Big Read: What does Google DeepMind want with the NHS? Retrieved April 22, 2017, from <https://www.digitalhealth.net/2017/03/deepmind-mustafa-suleyman-interview/>
- Swan, M. (2012). Health 2050: The Realization of Personalized Medicine through Crowdsourcing, the Quantified Self, and the Participatory Biocitizen. *Journal of Personalized Medicine*, 2(4), 93–118. <http://doi.org/10.3390/jpm2030093>
- Swan, M. (2013). The quantified self: Fundamental disruption in big data science and biological discovery. *Big Data*, 1(2), 85–99.
- Szalay, A., & Gray, J. (2006). 2020 Computing: Science in an Exponential World. *Nature*, 440(7083), 413–414. <http://doi.org/10.1038/440413a>
- Tanner, C. A., Padrick, K. P., Westfall, U. E., & Putzier, D. J. (1987). Diagnostic reasoning strategies of nurses and nursing students. *Nursing Research*, 36, 358–363. <http://doi.org/10.1097/00006199-198711000-00010>
- Tesauro, G. (1994). TD-Gammon, a self-teaching backgammon program, achieves master-level play. *Neural Computation*, 6(2), 215–219.
- The Gartner Group. (2016). *Gartner's 2016 Hype Cycle for Emerging Technologies Identifies Three Key Trends That Organizations Must Track to Gain Competitive Advantage*. New York. Retrieved from <http://www.gartner.com/newsroom/id/3412017>
- Tikkanen, E., Havulinna, A. S., Palotie, A., Salomaa, V., & Ripatti, S. (2013). Genetic risk prediction and a 2-stage risk screening strategy for coronary heart disease. *Arteriosclerosis, Thrombosis, and Vascular Biology*, 33(9), 2261–2266.
- Timmermans, S., & Mauck, A. (2005). The promises and pitfalls of evidence-based medicine. *Health Affairs*, 24(1), 18–28.
- Tsipouras, M. G., Exarchos, T. P., Fotiadis, D. I., Kotsia, A. P., Vakalis, K. V, Naka, K. K., & Michalis, L.

- K. (2008). Automated diagnosis of coronary artery disease based on data mining and fuzzy modeling. *IEEE Transactions on Information Technology in Biomedicine*, 12(4), 447–458.
- Turk, V. (2016). Google's DeepMind agrees new deal to share NHS patient data | New Scientist. Retrieved April 22, 2017, from <https://www.newscientist.com/article/2113701-googles-deepmind-agrees-new-deal-to-share-nhs-patient-data/>
- Victor, L. Y., Fagan, L. M., Bennett, S. W., Clancey, W. J., Scott, A. C., Hannigan, J. F., ... Cohen, S. N. (1984). An evaluation of MYCIN's advice. *Buchanan and Shortliffe [16]*, 589–596.
- Walton, R., Dovey, S., Harvey, E., & Freemantle, N. (1999). Computer support for determining drug dose: systematic review and meta-analysis. *Bmj*, 318(7189), 984–990.
- Warren, J., Beliakov, G., & Van Der Zwaag, B. (2000). Fuzzy logic in clinical practice decision support systems. In *System Sciences, 2000. Proceedings of the 33rd Annual Hawaii International Conference on* (p. 10–pp). IEEE.
- Witten, I. H., Frank, E., Hall, M. A., & Pal, C. J. (2016). *Data Mining: Practical machine learning tools and techniques*. Morgan Kaufmann.
- Wolf, S. H., Grol, R., Hutchinson, A., Eccles, M., & Grimshaw, J. (1999). Clinical guidelines: potential benefits, limitations, and harms of clinical guidelines. *BMJ: British Medical Journal*, 318(7182), 527.
- Yin, R. K. (2003). Case Study Research . Design and Methods. *SAGE Publications*.
<http://doi.org/10.1097/FCH.ob013e31822dda9e>
- Yu, V. L., Fagan, L. M., Wraith, S. M., Clancey, W. J., Scott, A. C., Hannigan, J., ... Cohen, S. N. (1979). Antimicrobial selection by a computer. A blinded evaluation by infectious diseases experts. *Jama*, 242(12), 1279–1282.
- Zikopoulos, P., & Eaton, C. (2011). *Understanding big data: Analytics for enterprise class hadoop and streaming data*. McGraw-Hill Osborne Media.