How to realise an adaptive patient model within iCBT using adaptive learner models

Jonas Zacharias Dørum Backer

Master's thesis in Software Engineering at

Department of Computing, Mathematics and Physics,
Bergen University College

Department of Informatics,
University of Bergen

Supervisors Svein-Ivar Lillehaug and Yngve Lamo

July 2020

Western Norway
University of
Applied Sciences

UNIVERSITAS BERGENSIS

Abstract

Internet-based cognitive behavioural therapy (iCBT) is a type of psychological treatment where some application, usually on the web or smartphone, can alleviate or remove the need for a therapist. The user can go through learning material and exercises without having to follow traditional face-to-face therapy sessions with a psychologist or psychiatrist. iCBT can be done with or without guidance, where the guidance typically is, in the form of calls with a therapist, feedback on exercises, short system messages or email exchange. This can be of great importance when the goal is to give psychological help to a larger group of people as fewer therapists are required per patient. Assisted iCBT serves as a lower threshold offering to those without the need for traditional therapy sessions. When compared to regular CBT, the dropout rates of guided iCBT are higher, with unguided iCBT having an even higher dropout rate. A contributing factor to this is that unguided iCBT and guided iCBT to a lesser degree take into account the differences between patients, which mean all patients go through the same content independent of all the individual differences in their needs (Ontario et al. 2019).

COPE is a research project that aims to make an adaptive unguided iCBT application for women dealing with symptoms after breast cancer treatment. This thesis first argues for why an adaptive patient model is required for such an application to be able to provide adaptive therapy tailored to each patient's individual needs and preferences. Then, it shows how three adaptive patient models with different abstraction levels using ontology modelling were made using learner models. The thesis also covers how process models were created to evaluate and further explain and demonstrate the models. Using design science, it is explained how the problems were formulated and how solutions were made to solve these problems with the models as artefacts.

Acknowledgements

I would like to thank my supervisors Svein-Ivar Lillehaug and Yngve Lamo for their guidance during the work with the thesis. I would like to thank the experts in psychology and learning that we were fortunate to have on our team and meetings. I would also like to thank Kolbein Toreson Foldøy and the skilled COPE team for the good conversations and help he provided with the work. And at last, I would like to thank my family and partner for all the support they gave me these two years.

Contents

List of Figures

1 Introduction

1.1 Motivation

The COPE research project wants to make an adaptive unguided internet-based cognitive behavioural therapy (iCBT) application for women dealing with symptoms after breast cancer treatment. By being unguided, it means there is no therapist guiding the patient user, and the patient is interacting with the application only. Typically, unguided applications as of today do not perform as well as guided therapy applications when being measured on quality in the therapy- and dropout rate. This application however, is adaptive. This means that the therapy material will be recommended based on various characteristics of each individual patient. This is something that has not been done yet and it will require a robust adaptive patient model.

This thesis will discuss the design and implementation of such a patient model using ontologies and process models. The goal for this master thesis project is to develop a student model that can contribute to the development of a COPE application that can combine the benefits from unguided iCBT with the benefits of adaptiveness and personalisation in traditional CBT and guided iCBT.

1.2 COPE

COPE is a psychological eHealth project with the aim to make the iCBT therapy method more effective by introducing three concepts: increase the degree of personalisation of content, adaptiveness in interaction, and facilitated feedback. These improvements compared to today's iCBT are achieved using a recommender system, a patient model and a modularised version of the iCBT. This functionality will work together to provide personalised therapy by letting the patient do recommended parts of iCBT based on information gathered about the patient in the patient model. The interactivity with the application is streamlined by acquiring and providing the right information at the right time instead of having the patient go through more extensive questionnaires regularly. Facilitated feedback is achieved by showing the patient relevant information based on the patient's current representation of the patient model, which is updated over time.

The COPE project is dependent on a good adaptive patient model which is something that has not been developed before in the setting of personalised iCBT therapy.

1.3 Research Questions

- **RQ 1:** How can a patient model based on learner models contribute to making a personalised adaptive iCBT application?
- **RQ 2:** How are adaptive and non-adaptive iCBT applications different?
- **RQ 3:** How can a patient model for use in a personalised iCBT application be evaluated?
- **RQ 4:** How are patient models and learner models different in the light of the thesis?

1.4 Research Methods

The goal of this thesis is to define and design a patient model that can facilitate for adaptive, personalised treatment of patients suffering from various symptoms after going through breast cancer treatment provided through an iCBT application. Process models will be used to evaluate and discuss its validity and usefulness. As such, we decided to use the design science methodology for this master thesis project. The patient model developed is divided into parts both relating to dimension, with three layers of complexity, and also logically by gathering concepts that fit together into more easily understood separated parts of the model. These can be considered the artefacts of the design process.

1.4.1 Design Science

In this thesis, the design science research methodology will be used. Classical engineering and the science of the artificial or "man-made" are the foundations for the design science paradigm whose goal is to design and develop artefacts. These artefacts are then improved over iterations, continuously evaluating the artefact and building knowledge (Vaishnavi & Kuechler 2004).

Hevner emphasises the importance of rigorous research which should be equally valued to the practical relevance of the result of the research (Hevner & Chatterjee 2010).

Preffers goes into detail and describes six activities related to design science:

- **Problem identification and motivation:** It is important to identify the specific problem the researchers want to solve, this may include isolating the target problem to get a better understanding of its complexity. Also it is of importance to define the motivation for why the problem needs to be solved as this can motivate the researchers and explain why the researchers reason the way they do. The motivation of this thesis is covered in the introduction (section 1.1 and 1.2) and in the problem identification part (1.4.2) of this chapter.

- **Define the objectives for a solution:** Defining an end result of how much more effective a solution will be compared to already existing solutions, or what areas are addressed by a solution to something that does not yet exist. This is discussed later in this chapter.

- **Design and development:** Covers the development of the actual artefact. This can be many different things. For this master thesis project it is a set of ontology models and process models. The design and the development are covered in chapters three and four.

- **Demonstration:** Demonstrating the artefact can be done in multiple ways depending on the artefact and what the goal of the demonstration is. A demonstration should contain the information required to use the artefact to solve the intended problem. This is covered in chapter five.

- **Evaluation:** The evaluation of the artefact is where the result is compared to the objective and how well it suits its purpose. This can be done in a multitude of ways and will for this thesis project be covered together with the demonstration in chapter five.

- **Communication:** The communication activity emphasises the need to explain the problem, the artefact made to deal with the problem, the effectiveness of the artefact and how it is of use to researchers and professionals within the problem area. This will be covered more explicitly in the introduction, evaluation and conclusion sections of the thesis, but will also be covered to a lesser degree in most sections of the thesis.

Peffers et al. (2007) describe the nominal structure of an empirical research process as: "problem definition, literature review, hypothesis development, data collection, analysis, results, discussion, and conclusion". Furthermore they claim that communication requires knowledge of the disciplinary culture. In our case, this would be the culture of therapists and researchers working on efficient solutions within the cognitive behavioural theory discipline of psychology. The therapists, being the partitioning professionals and also in many cases, researchers and the CBT and psychology, in general, being the discipline.

1.4.2 Problem identification

The problem area was presented by our supervisors Svein-Ivar Lillehaug and Yngve Lamo together with researchers at SLATE centre for the science of learning and technology and the University of Bergen. It was then refined through several internal meetings and some external meetings with the researchers at SLATE. Also we had seminars with practitioners in the area of CBT and mindfulness giving presentations.

The aim of this thesis is to develop the adaptive patient model using ontologies and evaluate these with process models. This will serve as a source of domain knowledge that will be expanded upon over time with lower abstraction levels. This will be discussed further in the method and design section.

1.4.3 Objectives for a solution

As no adaptive patient models for use with iCBT had been developed up until the start of the COPE project, the decision was made to use learner/student models from the field of intelligent tutoring systems (ITS) models as inspiration for these. A literature survey with a short written review were done to check the validity of this approach. Then, the work started researching ontologies, followed by the work on the iterations (Morales-Rodriguez 2012).

1.4.4 Ethical declaration

No real patient data was used in the work for this theses, and as such, no ethics approval was necessary. There are, however, several ethical dilemmas being discussed that would need thorough work and discussions within a multidisciplinary team and with the authorities in order to find solutions to when more than a prototype is going to be developed.

1.5 Thesis Structure

- **Chapter One** presents the motivation for this master thesis project as well as a summary of what the COPE project is. It also presents the research method.
- **Chapter Two** presents a walkthrough of the background foundation needed for this research project.
- **Chapter Three** presents the work done through the different iterations as well as a description of the artefacts.

- **Chapter Four** discuss the different tools used for the work as well as how adaptive learner models were used as inspiration for making the adaptive patient models
- **Chapter Five** presents the demonstrations and the evaluations done in the various iterations
- **Chapter Six** discusses the research contributions of the work as well as answers the research questions.
- **Chapter Seven** contains the conclusion of this thesis as well as some suggestions for further work.

2 Background

2.1 Cognitive Behavioural Therapy (CBT)

Cognitive behavioural therapy (CBT) is a type of psychological therapy which aim is to change the way one thinks and behaves to manage the issues that are being dealt with (NHS 2019).

Since the 1960s CBT has grown to be one of the most practised and verified disciplines of psychology. It is used to treat a multitude of disorders like anxiety, bipolarity and depression (Webb et al. 2017).

CBT tries to deal with the problems that the patient is suffering from at the moment, instead of looking into the past. It is believed that thoughts, feelings and physical sensations are connected and that the negative thoughts can trap the patient in a cycle of destructive thinking. By breaking the more significant problems into smaller ones and solving these, the patient can break out of negative patterns (NHS 2019).

Therapy typically happens in sessions with a therapist that helps the patient break out of destructive thought patterns and behaviour. This can last anywhere from 5 to 20 sessions, and the goal is for the patient to be able to apply the knowledge in suitable situations to help deal with issues as they arise (NHS 2019).

The UK National Health Service (NHS) lists some disadvantages and advantages of CBT. One advantage is that it might help a patient deal with or think differently about their symptoms when medicine on its own does not solve all the problems the patient is dealing with. It is a very flexible type of therapy that can be done using traditional therapy, books and apps. It can be completed rather quickly, and the strategies and practical knowledge can be used in everyday life when the therapy is finished (NHS 2019).

Some disadvantages of CBT include having to commit to the treatment for it to be effective, which can take up a significant amount of time when attending traditional sessions and doing the homework needed. Also, it can be uncomfortable looking into one's anxieties and emotions, and with the focus being on the patient to change themselves and not looking into potential sources of the problem in the family or past (NHS 2019).

A study from 1998 comparing costs between different types of treatment for patients suffering from alcohol-related problems show that CBT treatment comes in at \$198 per patient contact hour for therapy delivered in 12 sessions for a total cost of \$1901 (Cisler et al. 1998). Bodden et al. (2008)

researched the cost-effectiveness of family CBT versus individual CBT in clinically anxious children in 2008. They found that in the Netherlands, the societal treatment cost was 2751 EUR for individual treatment and 3051 EUR for family CBT. This means that costs are a significant factor when looking at alternatives to traditional CBT treatment.

The view of the COPE project is that the actual therapy done by the therapist in sessions is the ideal form of CBT therapy for those comfortable with it. Therefore it does not want to deal with problems related to the shortcomings of the theory behind CBT itself. It is concerned with the problems relating to the way that this content is presented and thereby solving some practical issues such as convenience, time, costs and comfort like all of the other CBT applications of today. The difference is that the application should be adaptive and tailored to the specific patient's needs, offering higher quality personalised therapy closer to the therapy the patient would get from a therapist in traditional sessions.

2.2 Internet-based Cognitive Behavioural Therapy (iCBT)

Internet-based cognitive behavioural therapy is the result of integrating information technology and psychological therapy. It can be delivered with or without guidance from a therapist and is usually accessed through a smartphone, tablet or computer. The content can be delivered in different media like text, audio or video (Vigerland et al. 2016).

The treatment is delivered through web- or smartphone applications such as StressProffen where the treatment is divided into modules of content that the patient will go through, usually in order (Børøsund et al. 2019). The content presented in the StressProffen app starts with an introduction module on what stress is and continues with nine different modules where the last one is a conclusion on what the patient has been through. StressProffen lets the patient go through the last five modules in whatever order they want to. The patient will work on these different modules and learn about their symptoms and how to deal with them. Each module takes between 35 to 50 minutes to complete, and the patient gets the option to read or listen to the content. The application also contains exercises for the patient to go through as they work on the modules. These can be anything from breathing exercises to challenging negative thoughts. As the patient goes through the content, encouraging messages will appear to motivate the patient. Once a day, a reminder with a positive message will appear Børøsund et al. (2019).

A pilot study by (Børøsund et al. 2019) used the StressProffen application to review the usefulness and ease of use it had for cancer survivors dealing with stress. Many patients reported that the application was easy to use and that they had a significant decrease in perceived stress level, anxiety and anxiety-depression score. No significant change in depression level alone was found.

iCBT can be delivered both with and without guidance from a therapist, with research demonstrating that guided self-help programs have the most significant effect and even come close or is equal in effect to traditional CBT sessions with a therapist. iCBT can be used to treat symptoms like depression, anxiety disorders, conditions relating to the behavioural medicine field and irritable bowel syndrome to mention a few. The use of iCBT is effective and less expensive than traditional CBT treatment (Lindner et al. 2013).

iCBT has contributed to decreased patient time per therapist as the traditional sessions are replaced with an iCBT application or shorter direct contact like conversations over the phone. If such an offering is facilitated by a mental health organisation, patients will usually go through an examination to see if they qualify or are fit to go through this type of treatment. Webb et al. (2017) reports that the overall outcome for iCBT for patients who have been found qualified in this way is good.

A study done in 2011 by Farrer et al. (2011) at the centre for mental health research at the Australian national university in Canberra tested the effects of iCBT for depression with and without telephone tracking. Among other tools they used MoodGYM, an iCBT application that claims to help prevent and manage symptoms of depression and anxiety that has earlier been tested by Twomey & O'Reilly (2017) on its effectiveness. The results were found to be of medium effectiveness with low adherence rates. Farrer et al. (2011) found that there was no significant difference between the patients that had weekly calls with a counsellor and those who did not. The iCBT program was shown to have good results regardless of whether telephone tracking was used.

Dropout rates are shown to be a factor for the success of iCBT. A meta-analysis from 2012 looking at 40 iCBT studies showed a dropout of 57%. There is, however, a difference between guided and unguided therapy. Standalone applications for iCBT had a dropout of 74% whilst therapist-assisted applications had a dropout of 28%. Another problem is the degree of interactivity and personalised therapy and the feedback the patient gets from this, which is likely to be the next step for iCBT (Webb et al. 2017).

2.3 COPE

COPE is a psychological e-health project at Western Norway University of Applied Sciences (HVL) in collaboration with SLATE, Kreftregisteret and University of Bergen (UiB). The project has aimed to develop an adaptive iCBT application, where instead of having all participants go through the same content, it would let the user have a say in what content should be displayed.

Having every patient go through the same content is a way to facilitate unguided or semi-guided iCBT by going through a set of iCBT modules from beginning to end. The setup of these modules have been implemented in various ways, but the users/patients still have to go through all the content from beginning to end either in groups or individually. This is in part due to the way studies of iCBT are being evaluated. In most studies variables have to be isolated to make sure any findings are due to the non-isolated variables, in a fully adaptive iCBT application this is impossible, and any findings would be hard to pin down to a smaller group of variables. What can be showed is the effect on well being, reduction in symptoms and drop out rates.

The goal for the COPE app is to deliver high-quality adaptive iCBT treatment to breast cancer survivors dealing with symptoms like fatigue, depression and stress.

Studies have shown the significant effect of iCBT trials with a specific purpose like targeting a subset of symptoms in specific groups. A study from the Netherlands done in 2017 shows that iCBT specified for women with the particular problem of suffering from fatigue and depression after breast cancer treatment worked better than CAU (care as usual). 73% of the women that got the iCBT treatment showed significant clinical improvement after six months. Among the women that got the CAU treatment only 28% showed significant improvement (Abrahams et al. 2017).

A study from Denmark showed the effects of iCBT-I or iCBT for insomnia in breast cancer survivors. Their study showed that iCBT-I was effective in breast cancer survivors and that this had the additional benefit of also reducing patients fatigue. They conclude that iCBT-I appears to be cost-saving and effective. It would enable them to treat the considerable amount of cancer survivors with insomnia as it is both effective and cost-saving (Zachariae et al. 2018).

The way iCBT has been performed in these studies is a step closer to what COPE is trying to do as the treatment is adapted to a specific area.

COPE wants to take it a step further and give all patients adapted treatment for each patient.

As CBT therapy revolves around the idea of identifying and dealing/coping with conditions or symptoms, it is these symptoms that divide the containers/modules of learning material and exercises found in the COPE application. By identifying the needs of the patients and adapting the content they see, COPE will be able to deliver adapted treatment to a patient, closer to what they would go through with a therapist, but with the benefits of unguided iCBT therapy.

To be able to provide personalised adaptive iCBT therapy, different aspects need to be considered. This is why the COPE research project at this initial stage is divided into four main parts: the content model, the patient model, the recommender system and front end application implementation, mirroring the architecture of the COPE application (figure 2.1).

The COPE architecture can be used to describe some of the main concepts of the COPE therapy application. At the top the is the user interface, this is what the patient of the application interacts with. Through the interface, the exercises and learning material will be delivered as well as setting up preferences. Through the usage of the app, the patient will be monitored. The monitoring entails collecting data of how often the app is used, how long the user spends on a page, updates in preferences, scores given to the material the patient goes through and screening. This feeds the patient model where lots of data is stored over time, creating a dynamic environment where needs fluctuate. The patient model, data coming from other sources and the available content feed the adaptive algorithm that then suggests new material for the patient to go through. The following sections describe each of these steps in greater detail.

As this was early on in the project, some changes have been made. Most notably, the target group was changed to students dealing with stress and depression, and the mindfulness part of the application was removed. This was done to limit the scope of the project as well as making it easier to test as a pilot. The long term goal of COPE is helping breast cancer patients deal with fatigue and depression after cancer treatment.

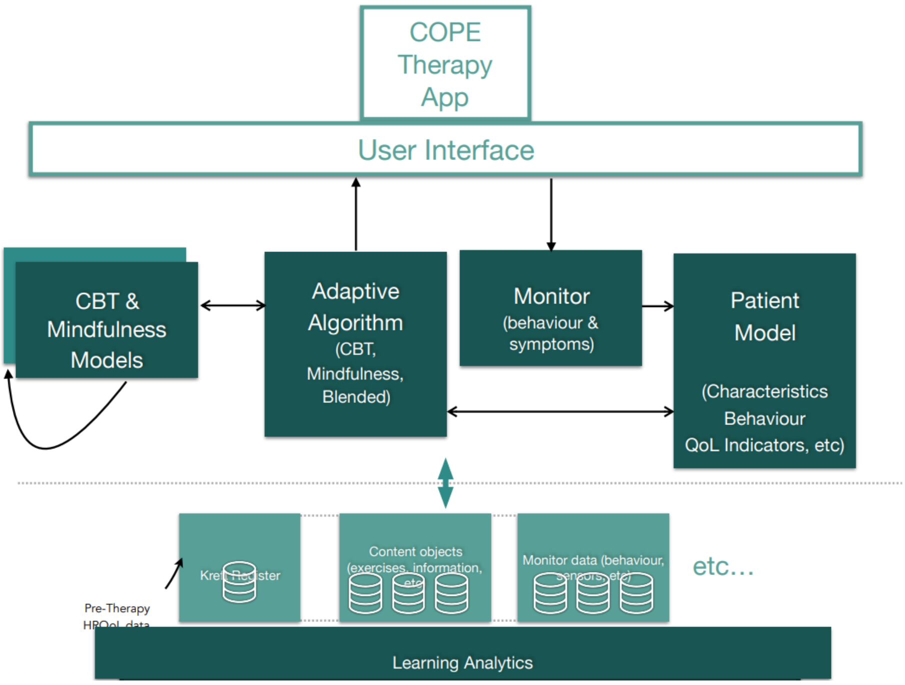

Figure 2.1: The COPE application infrastructure (SLATE 2019)

2.3.1 The Content Model

The content model will ensure that the iCBT learning material and exercises are divided into manageable chunks that make sense in the bigger picture. In the COPE project, the division of this content is done with tags that say something about the attributes of the content. An example could be different symptoms. Both the learning material and exercises may have these tags as well as constraints, where a patient has to finish some content to be able to access the next. The containers of content are then made available in different formats to give the user a choice between listening, reading or watching the content. Usually, the container will have some educational material that the patient will go through before doing some exercise connected to this material. There is also the possibility to repeat content like exercises that the patient found useful or satisfying.

According to Morgan et al. (2017) in a study done on the effects of unguided iCBT dealing with depression and anxiety, shortening a treatment does not improve adherence. However, both their longer and shorter course had a moderate to large effect on symptom severity. One of the goals in the COPE project is that the patient may choose to do only the modules of content relating to those patients symptoms. This also means that the course can be of variable length. The main gain of being able to affect their treatment will hopefully be the motivation to continue the treatment, thereby increasing the adherence. There is, however, the problem of assessing whether a patient dropped out or finished the treatment.

2.3.2 Patient models

The main topic in this thesis, the patient model, describes the patient and all the aspects connected to him/her. The model itself is what is known as an ontology model used to describe knowledge formally. Gruber (n.d.) defines ontology engineering as: "Ontology engineering is concerned with making representational choices that capture the relevant distinctions of a domain at the highest level of abstraction while still being as clear as possible about the meanings of terms". By binding known concepts together through axioms, the patient and all connected concepts will give a common understanding of the problem area (Noy & McGuinness 2000).

The patient model is based on aspects of adaptive student/learner models and recommender system models, as no adaptive patient models for use in psychological therapy applications have been developed before. Using a model as the basis for understanding the adaptive patient and its surrounding elements can give valuable information about the domain. These elements mostly consist of the containers of content previously discussed, how the patient ties into the recommender, the screening and monitoring of the patient and the patients' various symptoms. The patient model is designed for the patient to evolve over time, storing data used by the recommender to make sound suggestions. Process models that describe the flow of a particular patients' run through the therapy have been developed to further help with the understanding of the model itself and to identify instances where the model has good coverage or comes out short.

2.3.3 Recommender systems

Various recommender systems can work differently when compared to each other. Their common goal, however, is giving the user recommendations based on information stored about the user. Recommender systems can give suggestions in different areas depending on what the user is looking for, this can be trivial suggestions on products to buy or more complex suggestions relating to a hotel one should visit. The recommendation can, however, only be based on what is available in relation to the information stored about the user. In some cases, a recommendation can be rather trivial like suggesting the next entry in a series of books the user is reading. Other times the recommendation can be much more complicated, like suggesting a place to live. In this case, the system would need a lot more information to make a coherent choice. These complex or specialised cases is where recommendation systems have the highest efficiency. As for the gathering of user data, the user could type this in themselves or the system could gather activity data over time (Morlandstø et al. 2019).

Recommender systems use different techniques to make recommendations and are dependent on the type of recommendation requested and the data provided to the system. Examples of such techniques could be using demographic stereotyping or content-based recommendations done with machine learning, two widely different techniques, differing greatly in complexity (Morlandstø et al. 2019).

Recommender systems for student models achieve similar results as what we want to do with the COPE project. In e-learning the use of knowledge-based recommender systems is standard. This technique compares information about the users learning preferences, previous knowledge and the most relevant information to focus on to get the student to the next level, up against the objects it can recommend. It is essential to recommend something which has a high degree of usefulness for the user (Morlandstø et al. 2019).

The question thus, is whether this technique would work well for a patient using iCBT. A part of cognitive behavioural theory is concerned with learning about and being able to cope with the symptoms the patient is suffering from. The project will explore if a similar approach works for these patients as with the students. The model will store the patients preferences for learning about their symptoms and give recommendations based on various factors. Using screenings, one would also know what symptoms the patient find most pressing. This knowledge could then be used to suggest iCBT learning material and exercises where the patient has a say in what and how the content should be delivered thus hopefully decreasing drop out rates when comparing with an unguided tunnel view based approach where the user has to go through all the content in order without any customisation.

The recommender system is the link that binds the patient and content together. By looking at the patients' preferences, reported symptoms, monitor data, electronic health record (EHR) data, previous diagnosis and goals, the treatment could be tailored. This is done by assigning weight to different types of data focusing on the most pressing issues first. In the development phase, these weights are discussed with a psychologist to make sure that they make sense and are accurate. Both the weights and content are related to tags so that the recommendation can be put into actual content. A tag could, for instance, be related to insomnia or depression, if a patient scores high on both of these, the weights would determine the focus area. In this case, it would make sense to treat insomnia first as it might be a contributing factor to the high depression score. If constraint/requirements are met, the patient would be handed the insomnia content to work through first.

This is a very simplified and isolated example as a real recommendation would factor in things like the previously mentioned preferences, gender, age and multiple other factors all contributing to the final recommendation. What makes the recommender adaptive is the ability to look at all this data and the history of it. An example could be a patient that responded well to a particular exercise might get this exercise recommended again. Another factor is the ability to look at data gathered from several patients over time. If trends show that patients suffering from similar symptoms respond well to a specific type of content, then it is worth investigating whether a similar patient would also do so.

Ricci et al. (2011) writes about recommender systems and how they are used where the user lacks sufficient personal experience or competence to make choices where there are many alternatives. In this application, there are not only a lot of possible outcomes, but the number of variables and computing power required to get to these outcomes is vast. Some data, like patients trends, are also unavailable to the user, so the recommender system is required. The recommender system is tightly connected to the patient model as it uses the information stored to make decisions, it also feeds the model with data, creating a feedback loop. The recommender system is subject to some limitations that will be discussed later in the background section.

2.3.4 The Mobile Application

The COPE mobile application will be the usable end-user product that the patients will download and use. It will put the content into the hands of the user in a well-structured manner with human-computer interaction and interfaces in focus. The application will be able to present the content in different ways, audio, video and text depending on the preferences of the user. Monitoring will be used to track users steps within the application and how much time is spent on different pages of the application. This information will be stored and fed into the recommender system to make decisions. The application could be connected to monitoring devices in the future to gather things like heart rate whilst doing an exercise. Notices can be displayed by the application, these can be divided onto different groupings like positive, negative and notices. The positive might be relating to the patients' statistics and if they are getting better, this, however, might not be helpful for all patients. For some it might be harmful. Consequently, care has to to taken to develop this in a good way. As a result of monitoring rules, the app could suggest for the patient to contact a professional if the scores of the patient are looking to decline too much too fast and/or if the patient is suicidal. Notices could be messages telling the patient when it has been a certain time since the last time the app was used or a suggestion to do a particular exercise at a specific time et cetera. Finally, the application will, in a research mode, have a dashboard of sorts to show the users progressions and current health status in cases where it makes sense and is approved. Also, experimenting with dashboards for potential therapist inspection and exporting data in a format good for researchers have been suggested.

2.3.5 Student models

As mentioned, a commonality of recommender systems is the use of a user model and giving recommendations based on the information stored about this user. One such user model is what is known as a learner or student model (Morlandstø et al. 2019).

A student model is used to store data about the student, represent the student in a reasonable manner and track characteristics of this student before, during and after the learning takes place (Sottilare et al. 2013).

Graesser describe several challenges when trying to develop a good student model:

The theoretical problem: When continuously gathering data about the student, the image of the student becomes fluid and varies over time, Graesser uses the metaphor "A moving target". Even though this is a problem when trying to get a clear image of the student, it is also one of the essential properties of the student model. By looking at the history and development of a student, one could potentially predict future problem areas for the student (Sottilare et al. 2013).

The mapping problem: Relates to the interaction between the actual theoretical variables and the machine code, where data is often incomplete or incompatible (Sottilare et al. 2013).

The computational problem: There are lots of data that could be put together in so many different combinations that the computational power to do would exceed the limits of possibility with the technologies we have today. The only solution to this problem is to compromise on what data we use and how we compare different data to get answers that are good enough for what we are trying to do (Sottilare et al. 2013).

These are quite general problems related to adaptive models in general and as such does also come into effect when making a patient model.

Student models can be open for inspection; these are called open learner models and can give the user, teacher or others an understanding of the students' situation at the current time through inspection and interactivity with the model. Relating this to object oriented thinking, it would be like an actual object of a class. Such a model can be used in a multitude of contexts and can be developed in various ways depending on the problem. Sometimes it is enough to look at the single student and make a model tailored this student. Other times the model need information about multiple students and typical data common for these. There is also a considerable difference in complexity between different techniques (Morlandstø et al. 2019).

Open learner models can be developed to solve different problems. Some models give the student the possibility to plan, evaluate and reflect on progress. Others could be made to evaluate the effect of political measures on students perceived mastery in school. Thus user of the models can be different depending on the purpose (Morlandstø et al. 2019).

2.3.6 Interactivity

In addition to making the therapy more personalised, the COPE project also wants to solve two other aspects of the therapy that is believed connected to the dropout, the degree of interactivity and personalised feedback.

The interactivity issue is divided into two parts, not overloading the user and making the users feel as though the inputs they are providing make sense and are meaningful. It would be counter-intuitive to, for instance, give a patient a 100 question questionnaire once a week as this might further increase the feeling of stress and fatigue. In the worst case, it could potentially lead to the patient dropping out of the program. The COPE project supports the idea of asking the right questions at the right time and be smart about how information is gathered. This, however, is not necessarily a trivial task as the system would have to implement some logic as to when to prompt the user. In this first implementation of the COPE project, there is not a separate entity making sure this is adequately regulated. Instead, it has been discussed with therapists about what types of questionnaires would make sense to give the user, and at what interval this would happen. The final goal, however, is to have logic behind this, precisely deciding what and when information should be gathered depending on information gaps in the model. In this way, the intervention could be kept at a minimum and each piece of information would be valuable, leaving a minimal amount of waste information. Ideally, with a large amount of this information being used by the recommender system as a basis for the content provided to the user, the meaningfulness of gathering said data would be apparent.

2.3.7 Personalised Feedback

Personalised feedback could be solved through the use of an open learner style patient model that would allow extrapolation and display of data for user introspection. Some patients could benefit from this as they and potentially kin and therapist could view progress and problem areas increasing motivation to continue getting better. Other patients could have a negative effect from being able to view their data, and great care would have to be taken when developing such a system, probably being somewhat restricted in which patients should get access to such features.

iCBT applications will typically use information from the user to be able to evaluate how well the application worked for a specified patient group. This can be done with a rather lengthy and accurate questionnaire done at the beginning and at the end of the therapy. The COPE application will be harder to evaluate as it is more adaptive. It is not about testing a set piece of content, but rather the way of delivering content more sensibly. There are little in the way of restrictions, and even with the verified content and questionnaires, it's not verified in the way it's going to be used in the COPE project. Because of this, it would be close to impossible to say anything about what works and why, when evaluating the project as a whole. It is, however, possible to tell whether it works well or not by looking at patients reported improvement over time as well as drop out rates.

2.3.8 Ontology modelling

Some of the models made as part of this thesis are ontology models. The goal of making these models have been to document and share the knowledge and common understanding of concepts relating to the domain of the COPE application.

The word ontology origins from Greek and can be split into the two parts ontos and logos, Ontos meaning "being" and logos "reasoned". In combination, ontology becomes the study of existence or reality (Halpin 2009).

Reasoning about "what is" has been of philosophical interest for many years. More recently, within the field of informatics, the use of ontologies have become important, and they are used to make conceptual models within various domains. Ontologies represent knowledge within a domain by representing concepts or objects and the relationships between them (Halpin 2009).

Ontologies are designed with representation in mind. All concepts and naming are agreed upon to ensure that every member of the team can infer the same knowledge about the domain (Halpin 2009). As Gruber (n.d.) describes it: "In computer and information science, ontology is a technical term denoting an artefact that is designed for a purpose, which is to enable the modelling of knowledge about some domain, real or imagined". Furthermore, Gruber (n.d.) specifies that when making ontology models, we are not interested in the data structure or specific implementation strategies. The model is an abstraction from these and is said to be at a semantic level instead of a logical or physical one like most database schemas are.

2.3.9 Business Process Models

The process models were created to further increase the understanding gained from the ontology models and to give an abstract view into the processes that would need to be implemented into an application of this kind. The reason for these models being abstract is that they should be general and open to interpretation as long as the rules are followed. This means that they are somewhat independent of the implementation and as such, should be used as a tool and not a blueprint like implementation focused models. They visualise what components of the model the patient interact within at the initial use of the application and while working on some content. The making of process models has also been a reflection and evaluation process as they have shown where the model was lacking or had unnecessary details.

Although the first mention of business process models was back in the 60's (Aguilar-Saven 2004), the first real use was in the 90's. There is a multitude of techniques and tools, so much so that searching the web becomes a lengthy and hard process. Davenport (1993) defines processes as "structured, measured sets of activities designed to produce a specified output for a particular customer or market". In his book, he describes how, when businesses went from pushing for marginal levels of improvement to multiplicative ones, powerful tools were required.

Business process models have a starting point and an end result. They describe in logical order the steps and dependencies required to get to the result. This visualisation enables a common understanding of some process within a domain that can be used to share knowledge, optimise some process or alleviate dependencies (Aguilar-Saven 2004).

3 Method and Design

In this section, the designing and development of the artefacts are described.

3.1 Design Process

Throughout the year, meetings were usually held once a week with the supervisors and students working on the COPE project, with a handful of meetings at SLATE to get the input we needed from learning and mental health professionals. Early on the meetings were more general so that everyone could get a basic understanding of the project. After a while effort was made to give each student more specific guidance. As it is a relatively large project and details of how everything should work was worked out as we went we had some trouble figuring out exactly what to do early on.

Developing the artefact was done in iterations, with the following iteration building on the previous ones. Changes over iterations would mostly be due to learning, where better work could be done as the project became clearer and the tools became more familiar. Some changes throughout the iterations were done more directly in cases where existing work was expanded upon. This felt very natural towards the end as a top-down approach was used when modelling the last three models. The first iteration is different from the others as it contains the initial planning. Iterations two, three and four are where the most mistakes, learning and testing different tools took place. The last three iterations were where the work on the final artefacts was done.

3.1.1 The first iteration:

The research problems we are dealing with in the COPE project are not new. What is new is the ability to solve these problems in ways that have not been considered before. Our supervisors are experts in the field of health informatics. With help from other institutions, they have supervised master students in various areas dealing with e-health and how to make it more effective.

Our work as master students started with the course in health informatics, where we learned a lot about the history and current state of e-health. This was a reasonably comprehensive course covering a wide area of subjects. It was a good start, and gave us a basic understanding and ground rules to follow. Before the COPE project started, we each wrote a 12-page assignment discussing relevant problem areas for the project. Motivation for using ontology models and how to realise the models were the main topics. The paper discussed artefacts like recommenders, student models and the patient model itself, some of which has made it through to this paper.

At project start, we had visitors like Michael Antoni from the University of Miami and Sylvester Comprehensive Cancer Center talk about the current state of CBT for breast cancer patients and how it's done in groups. Anders Meland told us about how mindfulness is done in the Norwegian Armed Forces. Our goal was somewhat diffuse; we knew that COPE should be focused on women dealing with symptoms after breast cancer but were considering to cover both iCBT and mindfulness.

We had both internal meetings where the master students and the supervisors gathered, as well as, external meetings at SLATE/UiB where researchers in the field of adaptive learner models like Barbara Wasson and psychologists, most notably Hege Eriksen, were present. We also had human-computer interaction (HCI) and design PhD Rosaline Barendregt at these meetings.

Some main problems and several sub-problems were identified. Suggestions on how to proceed with solutions for all four of the students were discussed. There was some confusion still as the project was not planned into much detail beforehand. It was supposed to be a concept that would run on a mobile application providing mHealth through psychological treatment. It was meant to support both mindfulness and CBT, and the target audience was women dealing with symptoms like fatigue and depression after breast cancer treatment and potentially surgery.

3.1.2 The Second iteration:

In the second iteration, it was time to develop some models. At this point, the concept of ontologies, business process models and student models to base the models off were still rudimentary and was not considered enough in the development of the models. A tool called Lucidchart was used as it offered the possibility of making any type of model with a streamlined interface and easy accessibility.

The first model developed was an entity-relationship diagram. It contained the patient, some preferences the patient could set relating to the modality of content, screening and the concept of content divided into modules with each of these modules having constraints. Many of these data points were timestamped as a way of creating a patient with information that could fluctuate over time, like in the learner models. A look at this now reveals that some of the data represented here was close to what we wanted to represent, but there was also a lot of flaws. First of all, we were not looking for an entity-relationship diagram (ERD), but rather an ontology. An ontology is meant to convey understanding and meaning, and although an ERD can explicitly do this for people used to working in the field of computer science, this model would have to make sense to a team of psychologists and scientists in the field of learning. This meant that not only was the type of model not ideal, but also that this approach would have been a bottom-up approach as the ERD was a very low-level model. Some of the problems faced were the lack of actual implementations to look at, as most studies done within the adaptive e-learning discipline are good at conveying the concepts and the effects of their work, but do not go into very explicit detail of the implementation. Another problem was the lack of experience working with ontology models.

Work started on making models that could be understood by the psychologists and the learning scientists. Two of these models were of similar structure, one describing the concept of COPE and the other focused around the patient. They included many of the same concepts as the previous model, as these were discussed during internal meetings with the supervisors. These concepts were linked using relations. It was a step in the right direction and these were high-level models, but they were not accepted as they contained information that was not necessary and had relations that were not very common. This was still early on in the development, so they contained things we would later discard and was missing some of the concepts that would be introduced later.

One model from this research made it to the meeting with the full team. This was a very high-level model describing a feedback loop where a patient model and the iCBT and mindfulness modules would feed information into the adaptive algorithm that would suggest suitable content that the patient would go through. The patient would answer some schema where information would be gathered, and together with other sources of information would feed the resource server, further developing the patient over time. Although this was not an ontology model, it managed to convey some of the information needed for the full team to start discussing how to realise the concept.

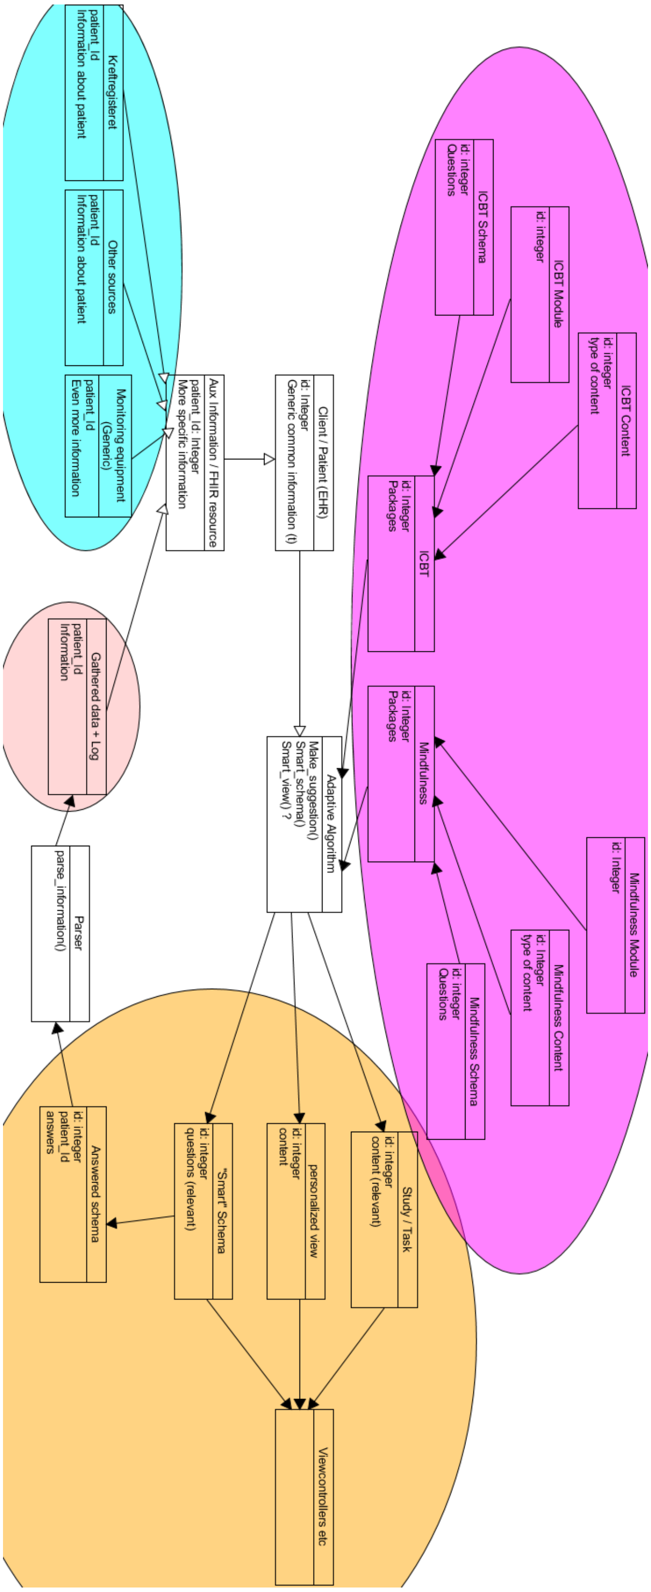

Figure 3.1: An informal model showing a feedback loop between the patient, content and recommender

3.1.3 The Third iteration:

At this point, no accepted ontologies had been made so more research was done to try to understand what was done wrong previously and improve on it. As this progressed, a new tool specifically engineered to make ontologies emerged. The tool is *Protege, open-source ontology editor and framework for building intelligent systems* (2020), and it was used to make the next ontology. First, classes are made. These are sorted hierarchically and they all inherit from "thing" as this is the most abstract entity. Object properties are set up next as these indicate the relation between the classes. Then, there are data properties. These are the fields of a class. The last type is individuals that are instantiations of these fields. When the setup is done, the model can be output with plugins to show it all visually. The model can also be exported to Resource Description Framework (RDF) and Web Ontology Language (OWL), two standard languages for making ontologies. The visualisation was the weak link of this tool, and when the supervisor preferred the use of a standard like Unified Modeling Language (UML) instead, the model was ultimately scrapped.

3.1.4 The Fourth iteration:

At the meetings, we discussed the COPE project and how to make something that could potentially be tested. A decision was made to make a case to narrow the scope of the project, and so we began looking at the case of students dealing with stress and related symptoms. It was also decided to go with cognitive behavioural theory alone as researching mindfulness as well in this application would add significantly to the complexity.

The models from here on and out will be made using a tool called *UMLet 14.3 Free UML Tool for Fast UML Diagrams* (2020), a lightweight tool containing all the functionality needed. The models made in this iteration were too specific to the implementation and were not abstract enough. A couple of the process models were good, but still too simple and implementation focused. They did aid discussion about how the feedback loop of the recommender would work, so they were not a total loss.

Through the first four iterations, the models have all been in lack of something. Some of them were not ontology models while others were too specific, going too far into implementation detail. Making the process models also reflected these issues. These models have, nevertheless, added value to the project as much learning took place, and the models were used to spark discussion. Ontologies are meant to store knowledge and convey an understanding of a domain, and to some degree, these flawed ontologies and process models did what they were meant to do, as this can be a tool both during and after the development process. However, as they were all in lack of something, e.g. the models had less than ideal relations between objects, were too technical or specific to the implementation and did not convey as much information as they could have if they were developed better.

3.1.5 The fifth iteration:

The fifth iteration is where the work on the final and accepted models started. The models were made using Umlet with standard UML elements. The models were split into parts wherever it was sensible to do so as this helped with decluttering. Through the process of making the models frequent meetings/discussions were had with the supervisor, the app developer and the team member working on the implementation of the adaptiveness. There were also some feedback from the learning and psychology experts that we met with once a month. The supervisor would mostly comment on the more technical aspects of the model, like how some relations should be named more appropriately or how some parts of the model did not make sense. Together with the application developer, we discussed how the parts of the model would fit into the actual implementation of the application, and we drew some inspiration from each other's work. Discussion with the student developing the recommender turned out to be fruitful as we had some different thoughts regarding naming. We also talked about what is meant by *adaptive* and how we could achieve it. The learning and psychology experts emphasised the process models and time aspects. It was important for them to evaluate and potentially expand the models using process models, these would reveal where the models were lacking.

3.1.6 The sixth and seventh iteration:

The last two iterations were very similar. First, gathering information about what works and what does not, and then expanding the models to show greater complexity, making sure the ontology models would support the steps of the process models. At meetings, we would agree upon naming and make sure that with all three abstractions of the models models available, there would be no issue understanding the information the models were trying to convey. Most of the expansion on the models happened very naturally as the whole team gained more specific information about what they were implementing and could deliberate on the concepts that were previously less detailed. The final result after the last three iterations were three sets of ontology models, one more specific and detailed than the other, and with process models to go along.

3.2 Artefact Description

The artefacts made for this thesis are the three last sets of models. As an iterative top-down approach was used to make these models, they will be described in the order of increasing complexity. As no one else has made an ontology patient model for the domain of adaptive psychological e-health before, the models have been made to analyse the domain, make assumptions about the domain explicit, share the resulting knowledge and understanding with the COPE team and hopefully enable reuse of the models, or at least give some pointers of our interpretation of the domain to others wanting to research and develop applications within the domain.

3.2.1 Highest level model

The highest level model is divided into five parts to make them easier to read. It also has three process models to help visualise how a patient interacts with the different parts of the model. The first part described will be the patient model.

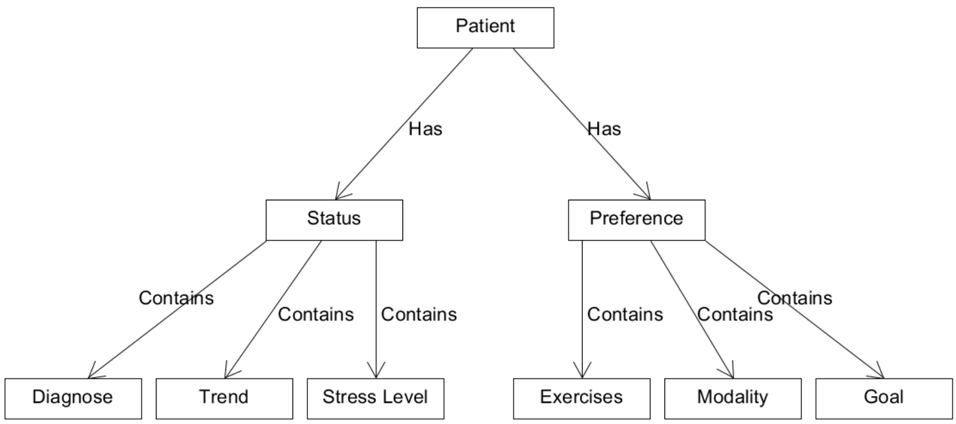

Figure 3.2: The highest abstraction patient model

At this level of abstraction, the patient is not defined very strictly. The only real thing that might be an indication that this patient is not sufficiently generic is the diagnose. Initially, the diagnose class was put into the model to contain the current symptoms the patient was suffering from. It was, however, agreed upon that this would be confusing. The diagnose as shown later will be the previous diagnosis of the patient gathered from the cancer registry. At this point in the project, it has been decided to relate the project to students dealing with stress for testing purposes, this means that the diagnose would just be left empty.

The patient, in this case, is thought of as a student dealing with stress as we are using this domain as a pilot and means of evaluating the project before making it fully specific for women dealing with breast cancer. There are however some elements specific for cancer patients as well, like the diagnose. The patient is defined by their current status and preferences. The status is a representation of the patient's well-being and will fluctuate over time. It is the result of studying a patient's diagnose, current trend and stress level. The status would also contain information about which symptoms the patient was dealing with, depicted in another part of the model.

The diagnose, as discussed, would show a cancer patient's previous diagnosis from the cancer registry as well as the treatment they received. The trend represents an extrapolation on aggregated data gathered through recurring screening of the patient. It is there to indicate whether the patient's mental health is declining or on an incline over time. The trend is possibly something that could be used by the recommender system itself to adjust the content the patient sees, but it's not necessarily very informative when figuring out which part of the treatment is working and which is not. In this case, it is more useful as an indicator that the patient is doing worse over time using the application and should seek out professional therapy, or the other way around, the patient is doing well and should continue using the application.

The stress level is a measure of the patient's stress over time. It is mainly determined by the stress mindset scale, which is a screening the patient goes through after finishing a piece of content. The stress level could also be affected by things like monitoring. When a patient navigates the application, information about how much time is spent on each page, how the patient navigates between different pages and how long the patient uses the app at a time is logged. Research would have to be done to give an indication of how the usage of the application ties into how stressed the patient is. Together with heartbeat monitoring, it could be a viable strategy when passively trying to monitor the patients stress level. As it stands, monitoring data would be gathered and analysed to check for other significant findings. It would not be used for something important like judging a patient's stress level without research.

The patients are also defined by their preferences. These are divided into exercises, modality and goals. The exercises are part of the modules or containers of content that the application has to offer. After some educational material, the patient will often be given a relevant exercise to go with the content. Examples of these can be everything from a simple breathing exercise to more complex and involved emotion reflection exercises. In the preferences, the patient will have a subset of his or her highest-rated exercises that will then be repeated more often as the patient may get good results from repeating some exercises. After each exercise, the patient will rate the exercise; this means that the exercise preference may change over time.

The modality preference refers to the modality of the educational content in the modules or containers. Just to reiterate, the containers themselves are mostly symptom-related and automatically recommended by the recommender system. The modality preference controls whether the content should be represented as text to be read by the user, a video on the matter, or in an audio format. The actual content represented should be similar between the three options. This is mainly a preference regarding the presentation.

The purpose of having the patient write down their goals can be different depending on the application. If viewed as a guided iCBT application, the goals would primarily be used as a way for the psychologist to structure the content around reaching these goals. However, as COPE's primary focus is on the development of an unguided adaptive mental e-health application, this poses some issues. The most natural way to gather the information would be through giving the patient some free text where they could write down their goals. This is also the most difficult way to interpret the input data as something like natural language processing is not necessarily ready for the task, and getting it wrong would potentially have significant consequences. Another way of dealing with this problem is limiting what the user could write, like having a text already written but letting the patient plug in the symptoms. The problem with this approach is that the patient might feel too restricted, and it could lead to more frustration than it would do good. The last approach is giving the patient the free text and then doing nothing with the result. This is a different mindset to the previous two, where the interesting thing is not being able to use the goal preference directly but instead, have it as some sort of exercise for the patient potentially enlightening and motivating them. Having the patient write down his/her goals could give the patient insight into what troubles him/her and give the motivation to deal with these problems. It could be debated that there's an ethical dilemma here as informing the patient about the fact that the information provided will not affect the treatment might lessen the effect of having the patient writing the goals in the first place, whilst not informing about it could be seen as unethical.

The patient model supports the personalised adaptiveness of the application by storing information about the patient's well-being over time, as well as, giving the patient personal preferences that have an effect on what content is delivered and how it is delivered.

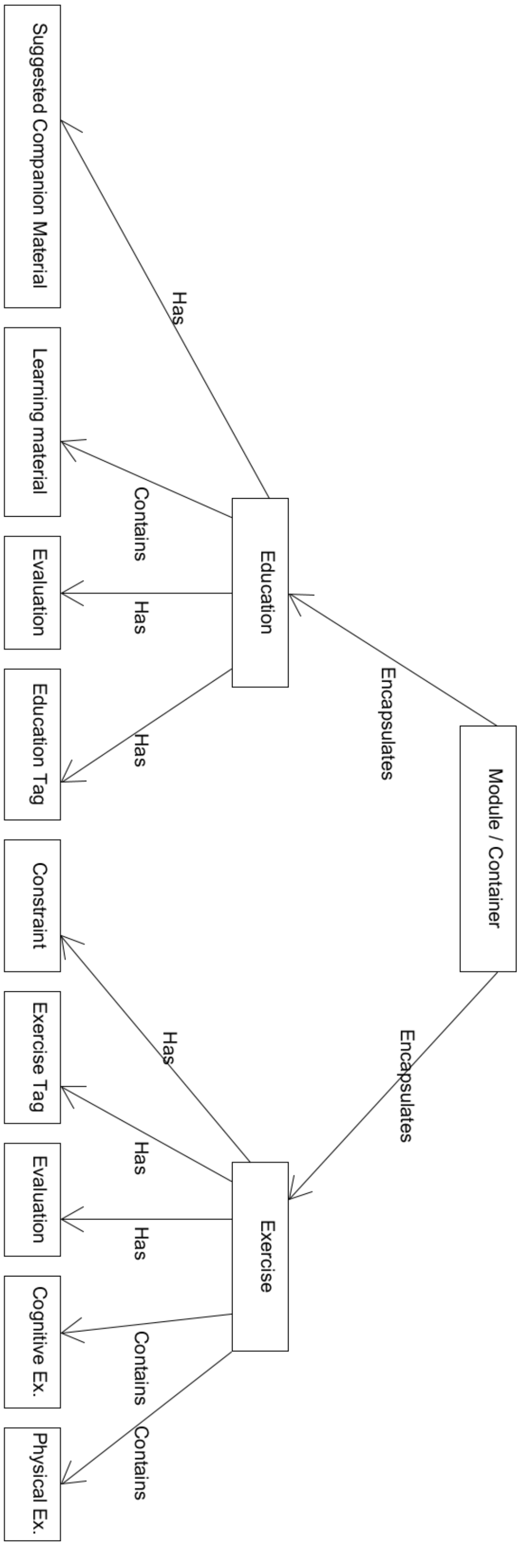

Figure 3.3: The highest abstraction module / container model

The module/container part of the model is where the content that the patient works through resides. As the model demonstrates, that includes both the educational material and the exercises. A container will usually contain symptom or practice specific educational material and exercises. This is further specified with tags on the content.

The educational part of a container includes the actual learning material, this is the material the patients go through to learn about their symptoms and how to improve or cope with them. There is also the education tag, in this model, the extent of this tag is not defined strictly. Discussions showed that tags were usually thought of as symptoms and some measure of difficulty, but in this higher-level model, it could mean that anything is valid. The tags make it easier to group content and lets the recommender suggest content with the same tags as the symptoms the patient is dealing with.

The evaluation part of the education is there to give a score of the content. With data gathered over time of many patients going through the system, it would be possible to recommend the content that others with the same symptoms gave a good evaluation to see if this content would fit the next patient as well. It is also a way of exposing content that could do with some improvement.

Most educational material will also contain suggested companion material. This can be material going more into depth on a subject that the patient is going through or connected in some other way to this material. This is one of those concepts that require professional psychologists who know the content and can structure it in ways that would make sense from a theoretical perspective as suggested companion material is strictly relating to the content and is independent of per patient customisation.

The exercise part of a module/container will have the possibility of containing both cognitive and physical exercises somehow relating to the educational material of the module/container. A cognitive exercise could be something like an emotion reflection exercise where the application would facilitate the process. A physical exercise could be a breathing exercise where the application would describe the steps of the exercise in the preferred format and then have something like an animation showing the breathing process to maximise the chance of success.

Just like the educational material, the exercises use tags to group the content to make the work of the recommender more effective. They would also have an evaluation that would work similarly to the one for the educational material with the added functionality of having the top-rated exercises be connected to the preferences of the patient. This would, in turn, make the application suggest repeating these exercises more frequently than other exercises if they are designed to be repeatable which many exercises are, some are indeed most effective when repeated.

All exercises will have constraints. This restricts access to a particular exercise until some requirements are met. These requirements are usually tied to the relevant educational material being done.

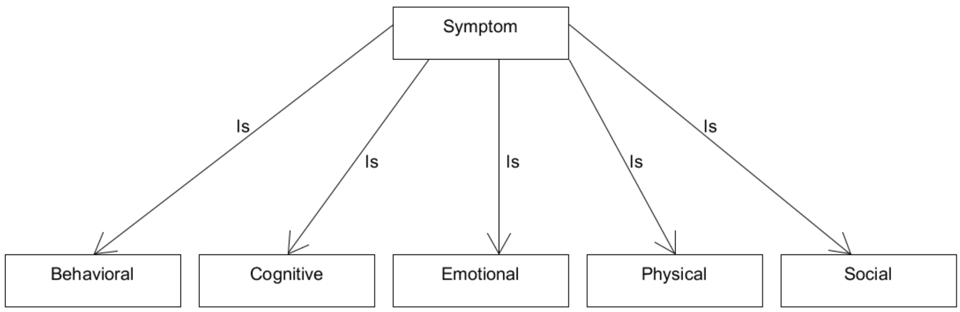

Figure 3.4: The highest abstraction symptoms model

The symptoms part of the model demonstrate categories of symptoms related to stress as described by Antoni (2016) for the use of stress management in a program for women with breast cancer or after breast cancer treatment, dealing with various symptoms. The symptoms have been divided into five categories and are used in CBT treatment. The behavioural symptoms are related to aspects like sleeping, drinking, drug use, and eating habits. Cognitive symptoms manifest as fears, phobias and low self-esteem. Emotional symptoms are linked to hostility, anger, resentment, self-esteem, irritability and depression. The physical symptoms can be things like headaches, muscle tension, back pain, ulcers, bowel related and feeling weak. Furthermore, the social symptoms are tied to withdrawal and isolating oneself (Antoni 2016).

These five categories of symptoms are examples of how the modules or containers could be divided into logical parts. There are also other ways of dividing the content.

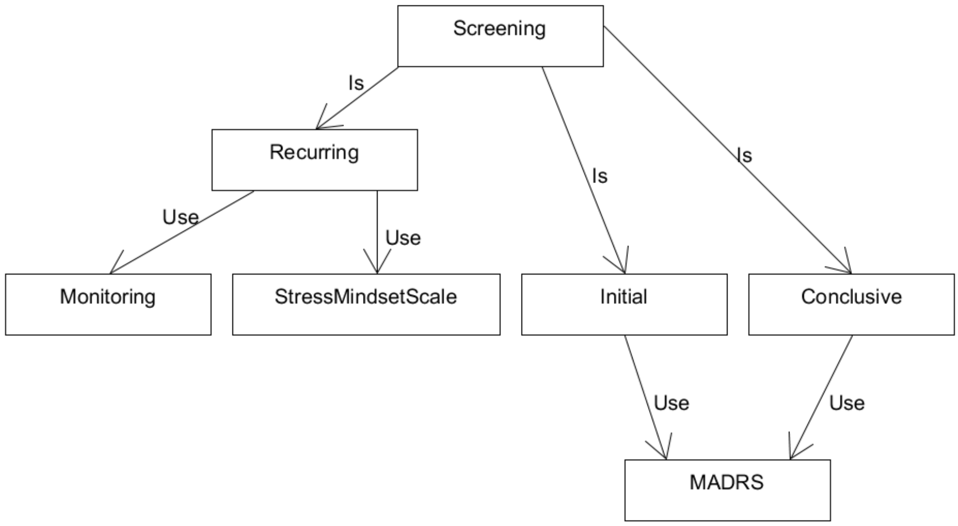

Figure 3.5: The highest abstraction screening model

The screening part of the model covers how screening could be done in the application when the area of interest is stress, as with our application used for demonstration and evaluation. The screening itself is divided into three kinds, initial, conclusive and recurring. The reasoning behind this, was finding a middle ground between what is currently legitimised and accepted within the research communities of mental health, in other words using a standardised test like the MADRS, and what would be convenient for the user, the stress mindset scale. As a baseline, every patient should go through the MADRS before starting treatment and again after treatment to show how the patient's mental health has been improved or worsened. However, this test would be cumbersome to have to do regularly because of the scale of the test. This is why the stress mindset scale is used in recurring screening. Its lightweight and specific results in the stress domain makes it a good fit for everyday patient testing with the use of the application.

Monitoring is both a way of surveillance and screening in the application. The application could potentially use heartbeat monitoring while the patient is doing some exercise to see the effects that might occur. This would have to be based on some actual research, and therapists would have to legitimise the use of it. The way monitoring is used for surveillance is that it captures activity data and stores it. This way, it is possible to tell if a user has had time to read the content of the learning material and other things.

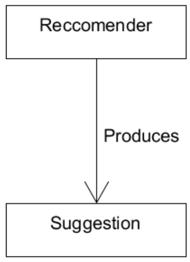

Figure 3.6: The highest abstraction recommender model

At this point in the model development **the recommender** is not portrayed in detail. This is because it is there to spark conversation, meanwhile not setting any limits. It is also because the development process of the recommender was at an early stage where details had not yet been sorted out.

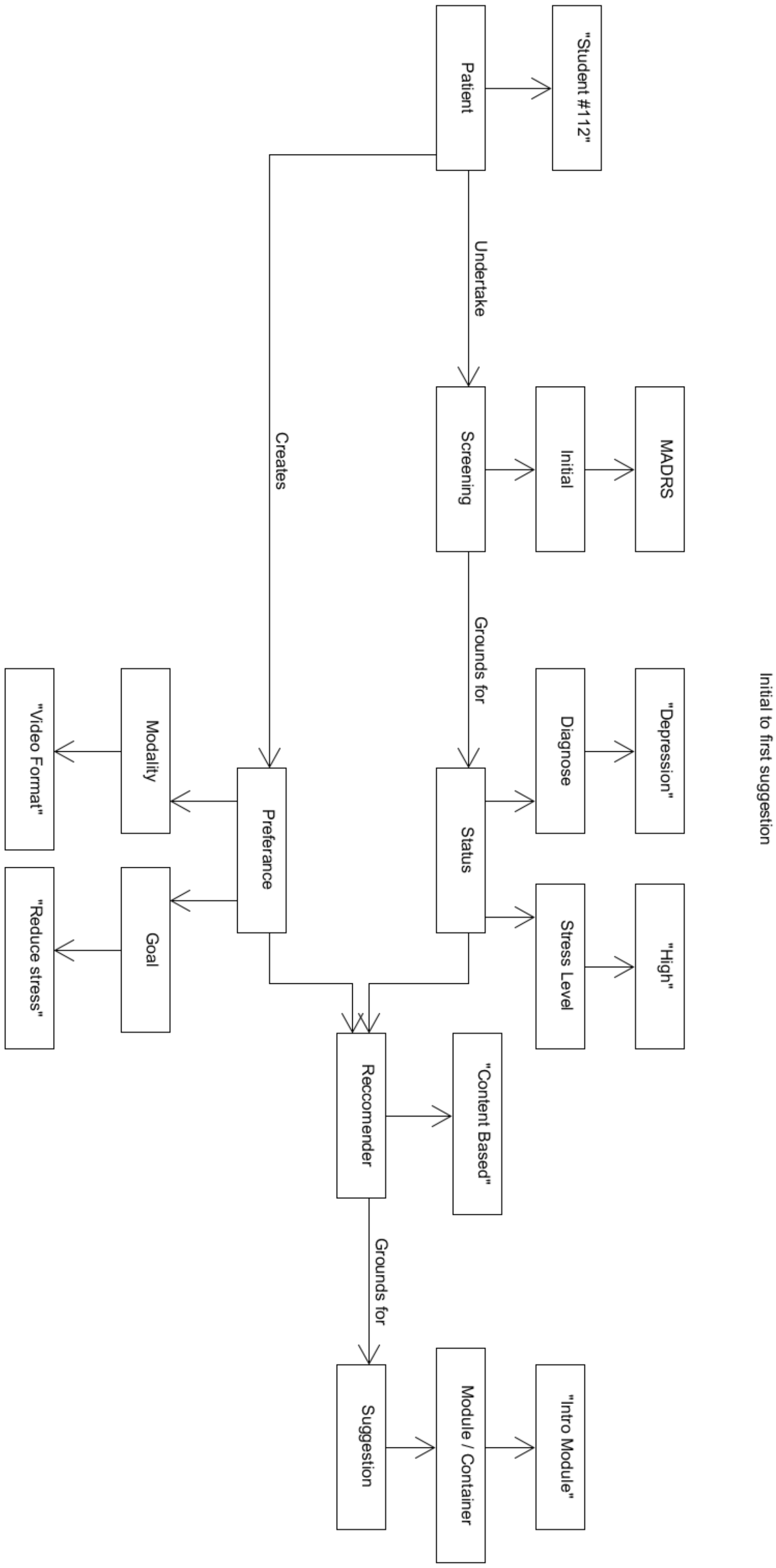

Figure 3.7: Process model depicting the patient's initial use of the application based on the high level models

The process models illustrate the flow of the application based on the high-level models with examples. This is done to give all participants on the project a better understanding of the models and to discuss where changes need to be made in the models. This first process model shows the patient's initial use of the application.

First off, we have the patient, in this case, a student who is going to partake in the iCBT application focused on students dealing with stress. The patient starts by doing the initial screening and setting up his/her preferences, this sets up the first information the recommender can use for its recommendation. As we can see, the patient chose the video format modality and wrote down a goal of reducing stress. The screening determined the initial status of the patient, as mentioned earlier, the diagnose will not be determined by this screening but will be something the patient brought with it from a registry like Kreftregisteret. This has been specified in the lower-level models. The screening did, however provide some information about the patient having a high stress level.

At this point the student working on the recommender was choosing a type of recommender suitable for the job, a content-based one has been set as an example here, but this is also something that was discussed further. A hybrid approach between a content-based and knowledge-based one was later used. The recommender goes through its various steps, but at this point, the constraints of the modules will make the recommender recommend the intro module as this is a requirement.

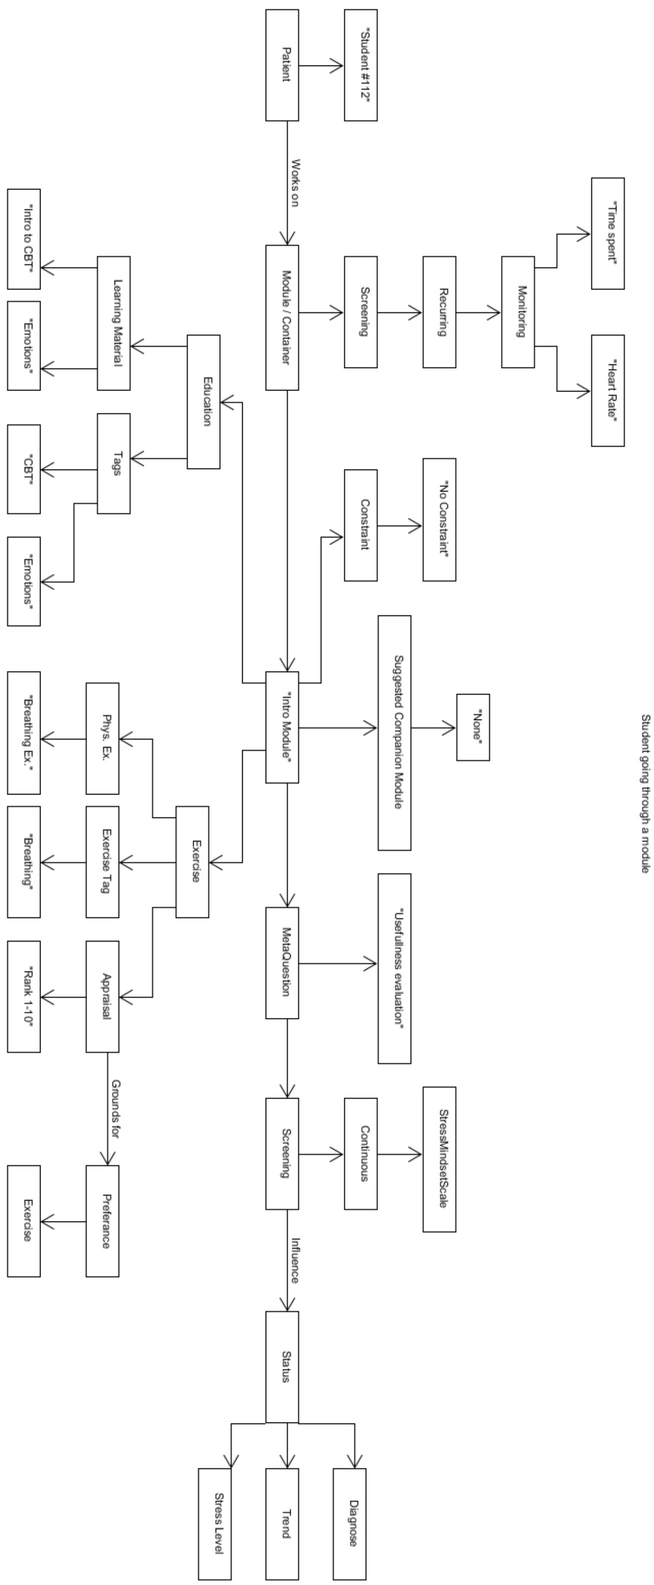

Figure 3.8: Process model depicting the patient working on a module based on the high level models

Now that the patient is working on the module/container that was recommended, the process is monitored using user data collection about how long the patient spends on a page et cetera. It would also be possible to use monitoring like heart rate monitors. The content being worked on in this case is the "intro module", it is an intro to CBT treatment and emotion and as such is also tagged similarly. The tags, in this case, are simplified and could be expanded to fit the needs of the application, this also controls how much information the recommender gets about the modules/containers.

The same goes for the exercise connected to this module, its a breathing exercise and if appraised high enough by the user, it could affect the exercise preferences in the patient's profile. The intro module does not have any constraints and also no suggested companion modules as it is very general and not specific to any symptoms.

Next, the patient scores the learning material and exercise before doing the stress mindset scale screening. This then updates the status to reflect any changes in the trend and stress level.

3.2.2 Lower level models

After the work on the higher-level models, it was clear that the model of the recommender would be the next step. At this point, the recommender was discussed in more detail, and a feedback loop was conceptualised.

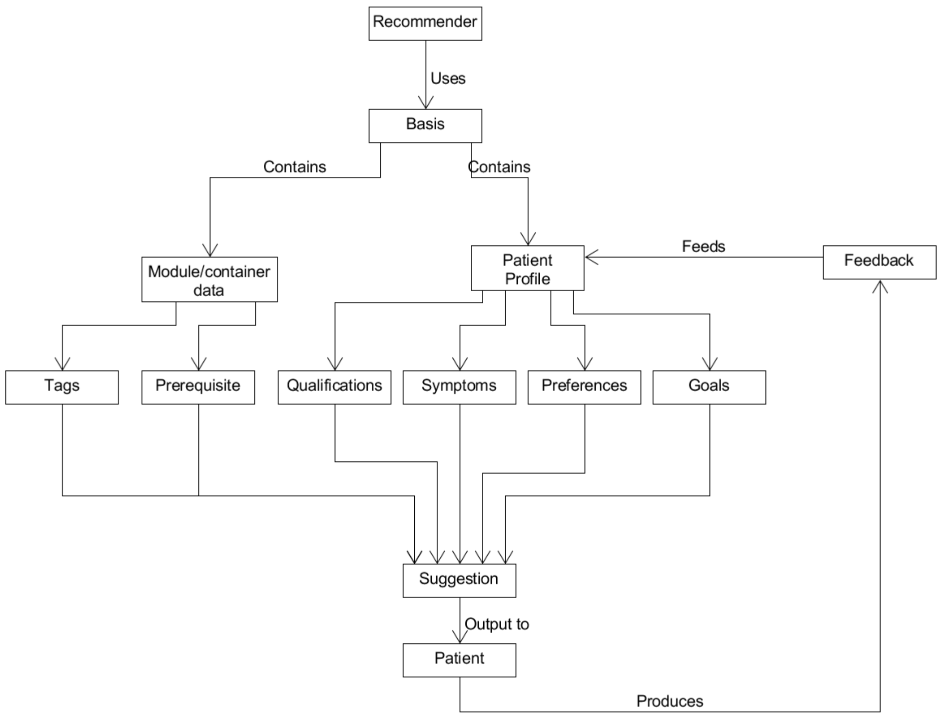

Figure 3.9: A model showing the feedback loop of the recommender

The model shows how **the recommender** uses the module/container data and the patient profile data as a basis for a suggestion. The tags are still the tags shown earlier in the high-level module/container model that really can be anything describing the content as not to go into specific implementation detail. In the application at this point, it has mostly been used to tie the content to a specific symptom to allow grouping of modules that contain learning material and exercises related to the same symptoms. The prerequisites earlier talked about as constraints, are hard limiters in place to give the content some order so that the patient is given content that would give meaning when previous material is done.

The patient data is divided into four main parts at this point. The qualifications of the patient are the modules/containers the patient has already been through and is tied to the prerequisite. The symptoms are the symptoms the patient is struggling with and severity of these. Preferences at this point are the preferred modality of content and current favourite exercises. The goals are some goals the patient has set like reducing stress.

This data all goes into the recommender that then outputs a suggestion to the patient. The patient works through the recommended content and gives feedback through screening and rating of the content.

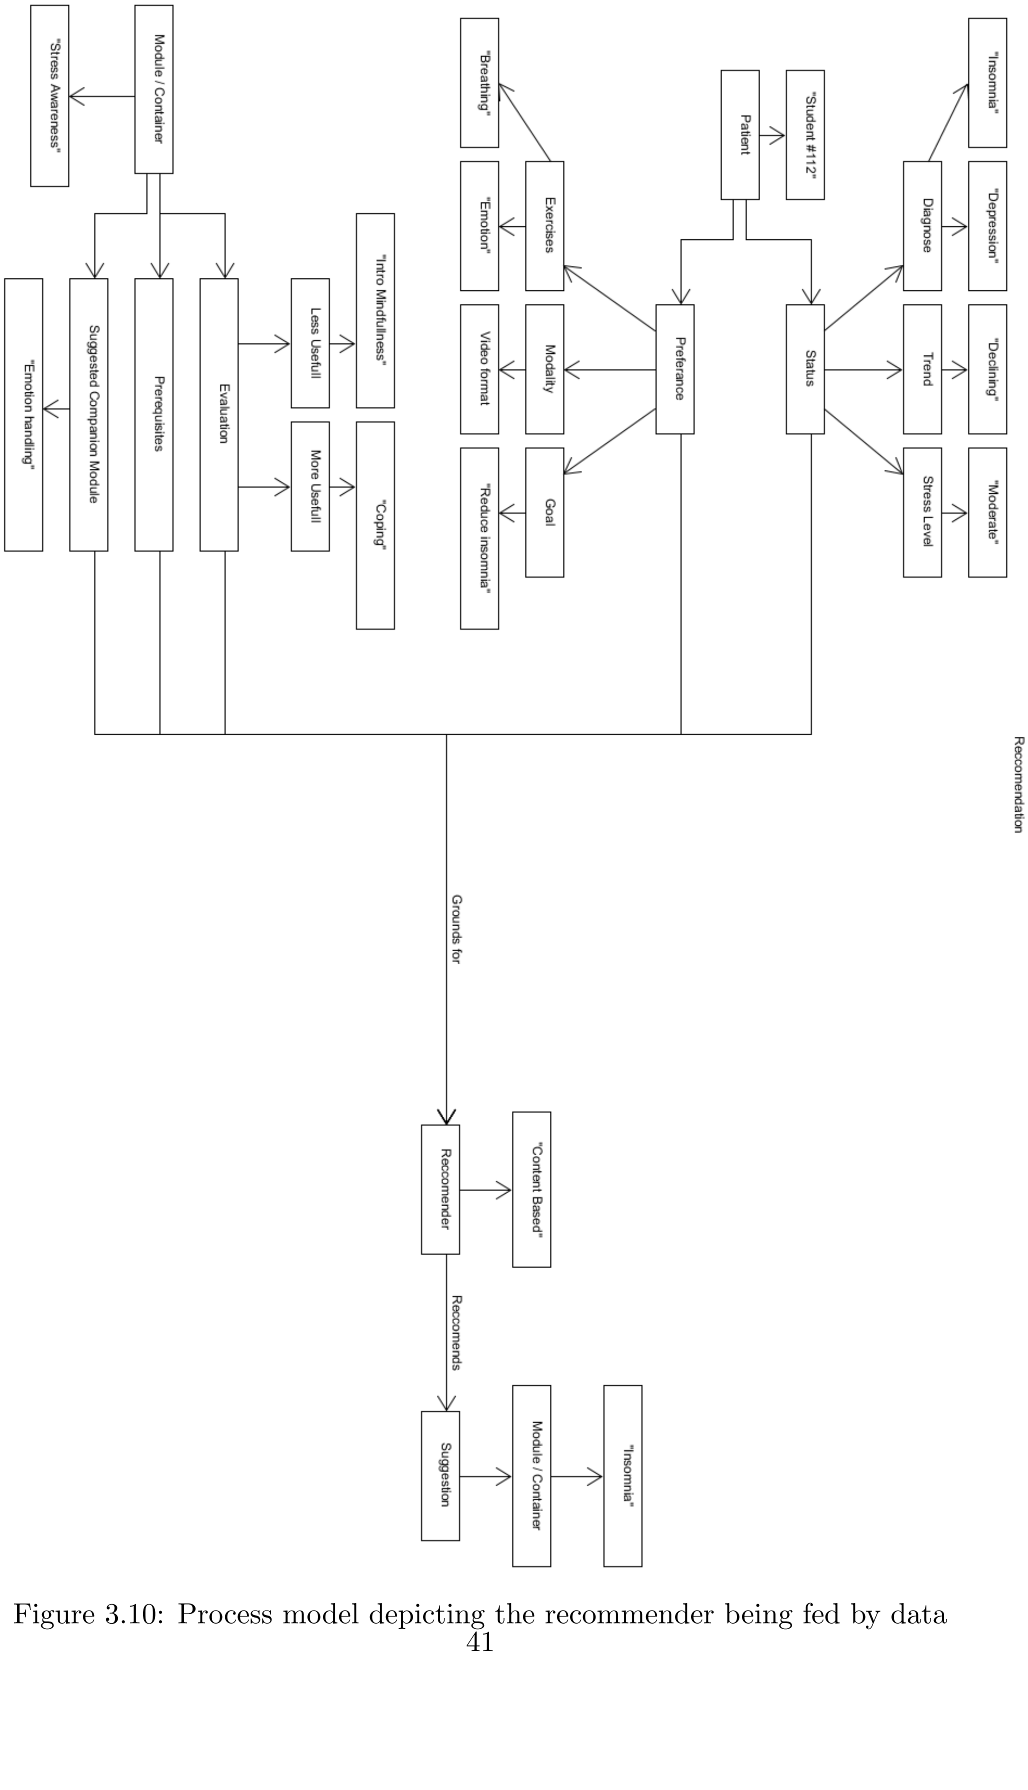

Figure 3.10: Process model depicting the recommender being fed by data

This **process model** shows how the status and preferences of the patient combine with the evaluation of earlier content as well as suggested companion modules to feed the recommender the information needed to make a recommendation. This model was made around the same time as the recommender model above. Notice how the model above is lacking the input of suggested companion module from the module/container data. This is something that became very clear when making the process model and as such, is featured in the next iteration of the process model. The inputs are weighted and count differently depending on importance. The importance factor of the inputs was discussed with a professional psychologist.

3.2.3 Lowest level models

At this point, the finer detail is falling into place, and this is reflected in the models. These final models are mostly based on the previous ones, with the exception of some new models.

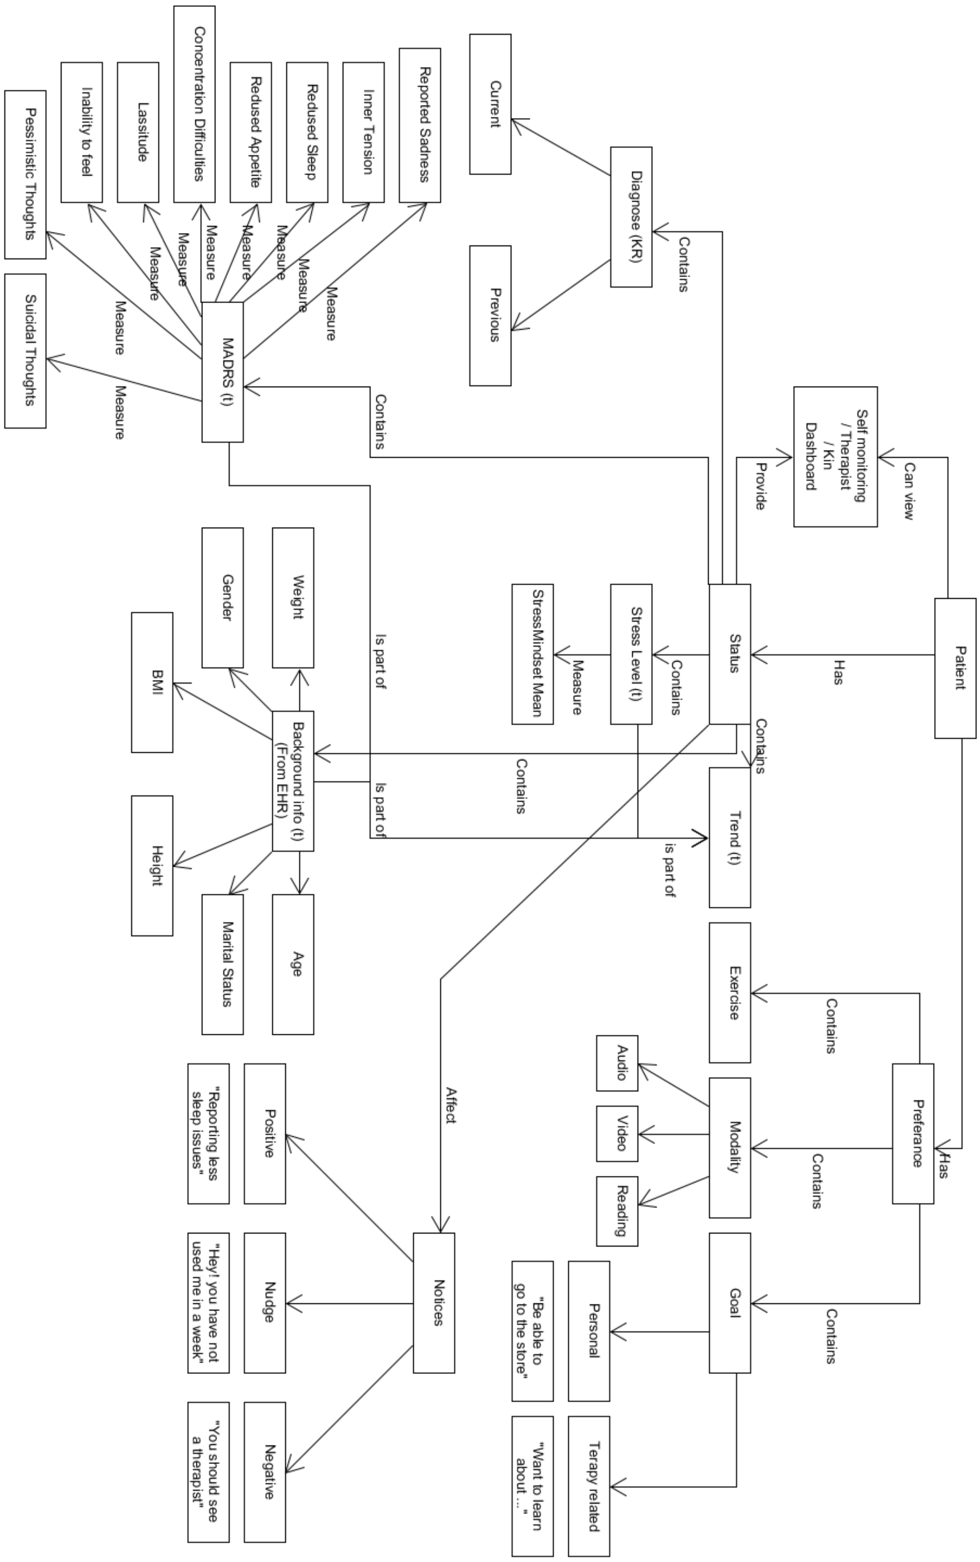

Figure 3.11: Model of the lowest level patient

The lowest level patient model was updated through both expanding the concepts it already contained and adding some new ones. The goals have been expanded to contain both the personal goal of the patient as well as therapy-related goals. This was done to give the patient the freedom to write a free-text answer for the personal goal, which exist only as motivation and exercise for the patient and will not be interpreted by the application. The therapy-related goals are goals that will be interpreted and weighted to give the recommender more to work with.

The patient now also has a view, this was inspired by open learner models described by Bull & Kay (2010) where reflection and self-monitoring can be important tools in learning. Grieg (2019) made a visual dashboard solution for therapists managing multiple patients. With something similar that would also adhere to some of the principals of the open learner model a tool could be made that would support both the patient and potentially allow kin or therapists to monitor the therapy. This idea will be discussed further in the implementation section.

The status of the patient changes with time and is affected by multiple sources. As the COPE project is concerned with treating women suffering from symptoms after breast cancer treatment, the status contains the patients' previous and current symptoms from the cancer registry Kreftregisteret. This part of the model is not very applicable to the application that was developed during the thesis as this one focused on students dealing with stress to be able to test and evaluate it as a pilot. The status also contains any recorded MADRS and stress mindset screenings with emphasis on the latest ones recorded. It contains the background information of the patient gathered from the EHR, and lastly, it contains the trend.

The status will affect notices that could be displayed on a smartphone to give positive, negative and nudge notices. It would be imperative to discuss these with professional psychologists as even positive notices could negatively impact the treatment. Positive notices could be information relating to the patients' screening should they show reported symptom reduction like a lower score on sleeping issues. A negative notice could be displayed when a patient is showing low enough scores or signs that he/she is suicidal, this would trigger some flags and a suggestion to visit a therapist would be shown. A nudge is something that would be displayed to the patient to remind them to do their weekly exercise or similar, it is crucial that these are not too intruding or annoying as the patient could lose interest.

The last part of the model is the trend, it shows how a patient is reporting symptoms over time through the screening and reveal potential physiological changes like weight loss from the health record data. The trend can be a useful metric both for the recommender and the potential use of a dashboard.

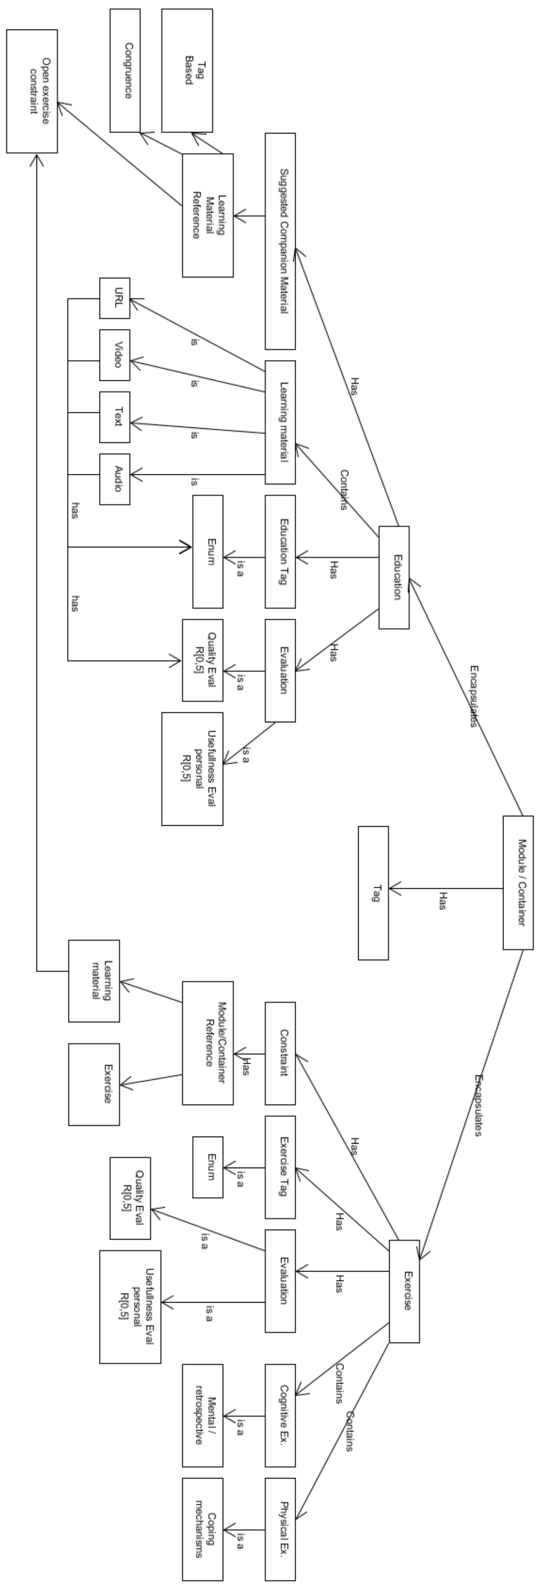

Figure 3.12: Model of the lowest level Module/Container

The module/container part of the model has been specified in greater detail than the previous iteration. A module/container is still a container of both learning material and exercises, but the content is explicitly specified by tags rather than being undefined like earlier. This means that anything that is considered a tag can define the content of a container. A tag can be anything, but like discussed earlier symptoms can be an excellent first step to define the content. This is where the expertise of professional psychologists could greatly benefit the usefulness of tagging the content precisely and correctly so that the recommender would have an easier time suggesting content that would benefit the patient.

The education part of the container is where the learning material resides. The learning material is split into the modality of the content and is also tagged in further detail. Educational material is evaluated both on the quality and usefulness of the content by each patient after completion. This is part of the recommender feedback loop and serves to both feed the data pool and provide quality assurance on the material.

An example of this could be if a patient fits into a category and is suffering from specific symptoms he or she might be recommended some learning material that turns out to be beneficial for the patient. After working through the content, the patient would then give a high usefulness score to that particular content and would potentially also get a better score during the stress mindset scale screening. Combined, these results would feed the data pool that grows continually. Other patients that fit into the same category, suffering from the same symptoms, would be more likely to get the same content recommended.

The suggested companion material now has a reference to some learning material. Like described earlier, the suggested companion material is based on the last material a patient worked on and is weighted higher by the recommender. Suggested material can be based on tags where similarly tagged content will have a higher weight than if not. It can be based on whether or not it is the next material in a line of connected material like the third part discussing stress when a patient has done the previous two parts. Lastly, it can be used to have the patient go through the content necessary to unlock an exercise.

The exercise part of the module/container has many of the same aspects as the education part. The content is grouped into cognitive and physical exercises, and each of these is followed up by both quality and use usefulness evaluation. As with the learning material these are in place to make sure the exercises are of good quality and show where changes need to be made as well as give an indication of what exercises are working for what patients. Like earlier exercises are categorised by tags. The constraints have been specified in greater detail to contain both learning material and other exercises. This is because some exercises rely solely on having gone through the learning material, while others might need the patient to finish some exercise to do a more advanced one that builds on the previous exercise.

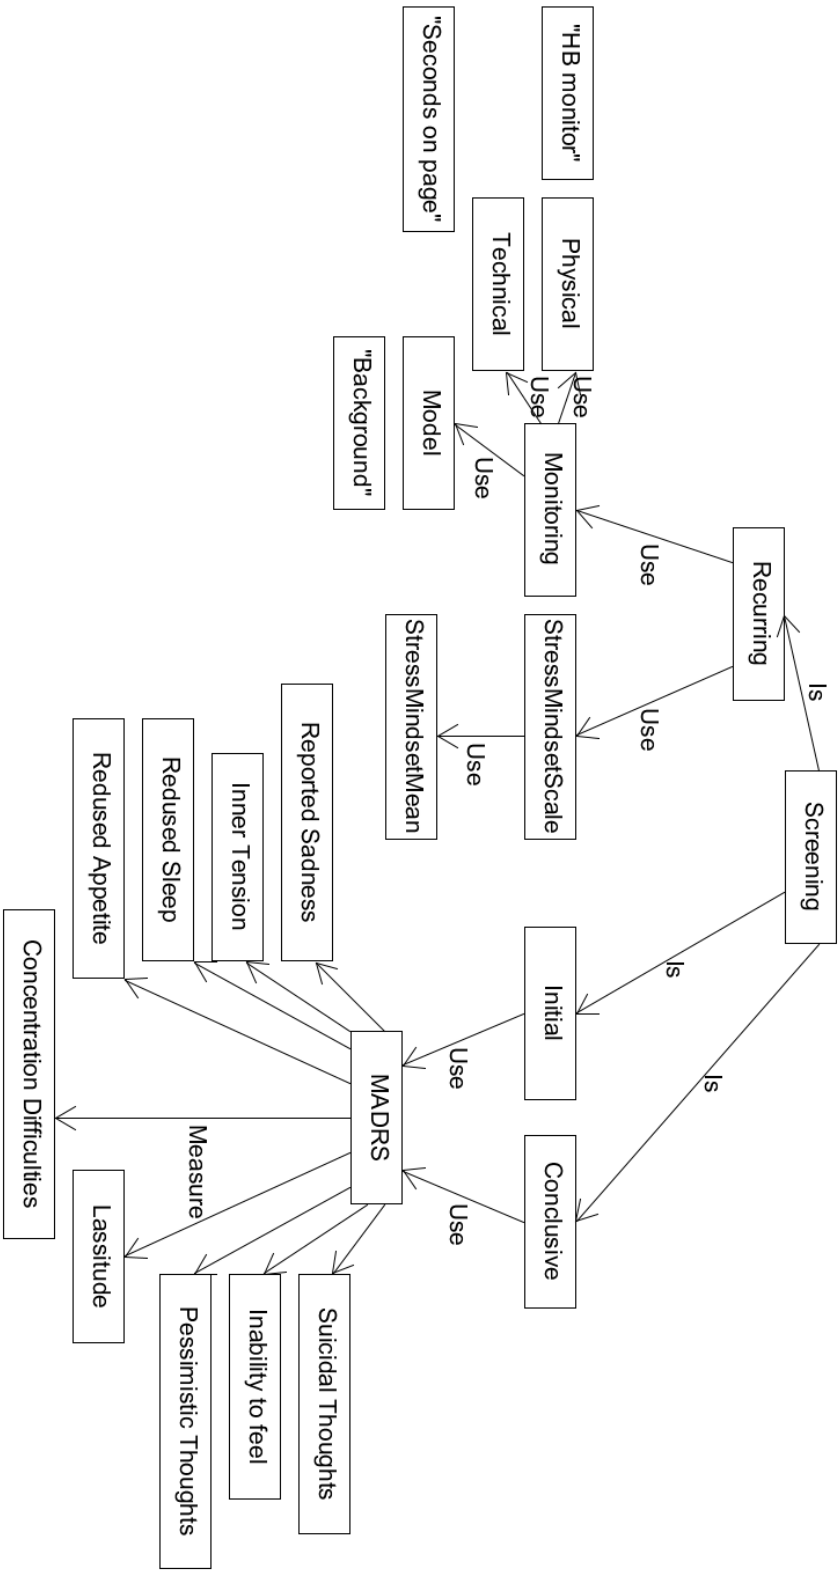

Figure 3.13: Model of the lowest level screening

The screening part of the model has been refined to greater detail. Monitoring now includes the different kinds of monitoring that could be supported by the COPE application. Physical monitoring mainly refers to anything that could be measured physiologically like a heartbeat monitor. This could be connected to the phone via Bluetooth, and data could be collected to monitor aspects related to heartbeat when doing exercises. Technical monitoring shows the users navigation through the application, this user data could be used to improve things like user experience if it becomes clear that the patients are finding it challenging to navigate the application. It could also be used to check how much time patients use doing exercises and learning material. The model of the patient is being monitored to keep the data updated and give a picture of the patient over time.

The stress mindset scale has been expanded to show that we are interested in the stress mindset mean that the screening outputs, this is the main component in determining the patients stress level. The initial and conclusive screening, the MADRS test, now shows the different aspects of the patient that is being tested.

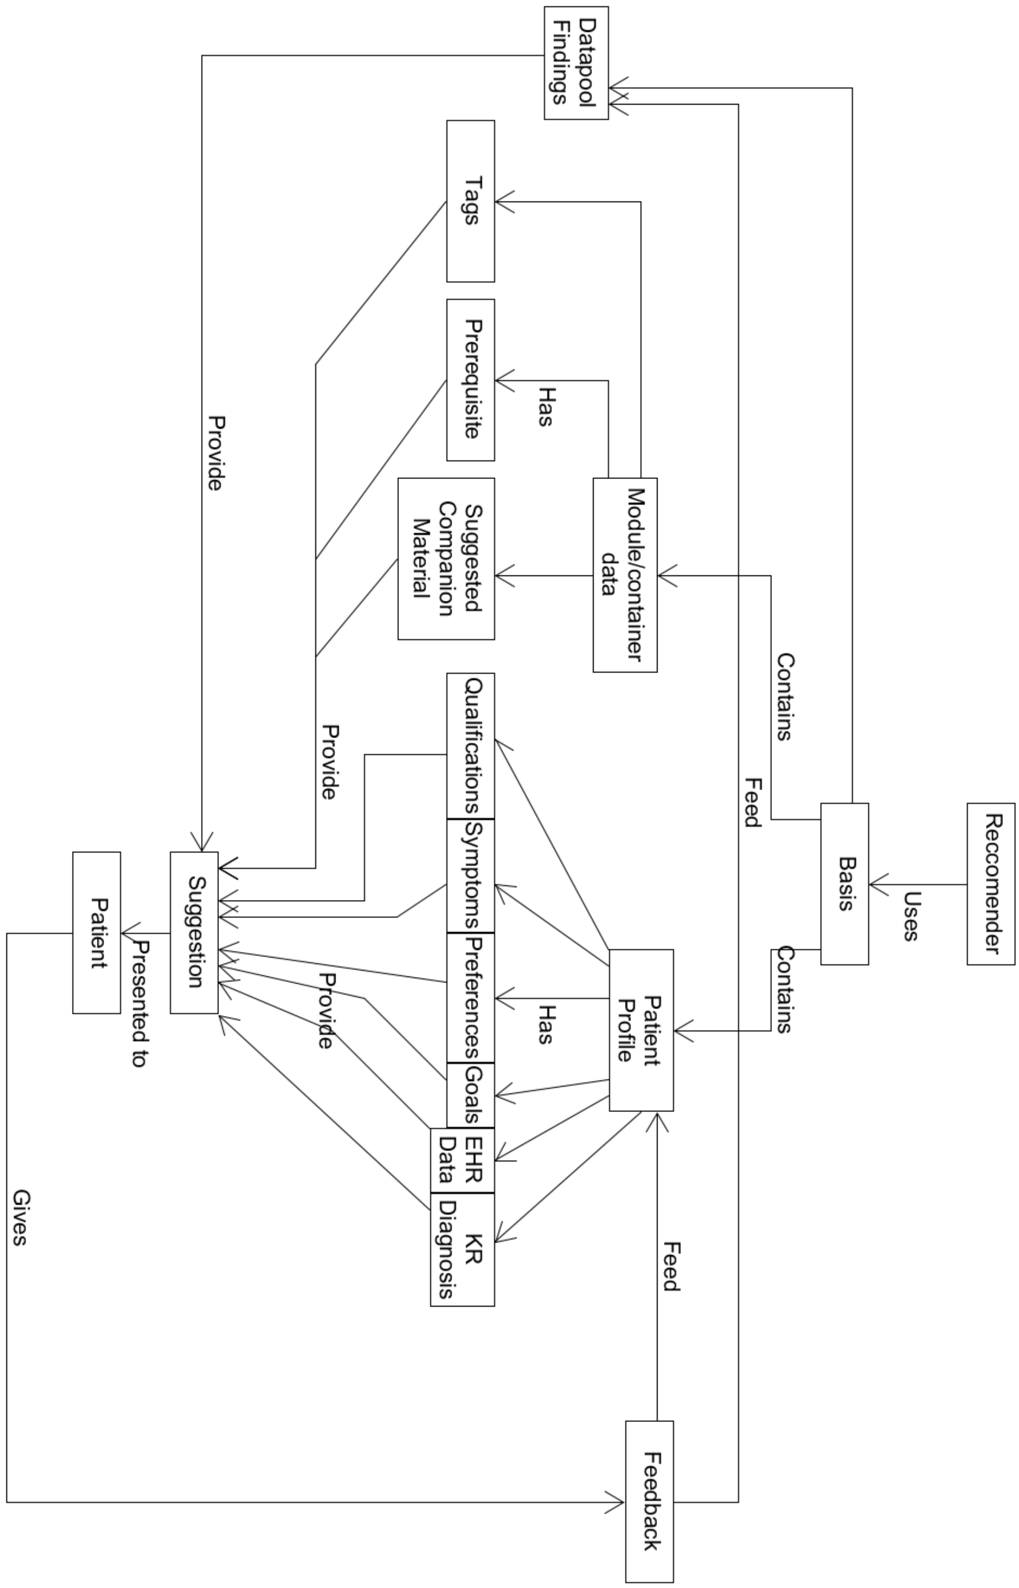

Figure 3.14: Model of the lowest level recommender

The recommender model is similar to the previous one, but some more data points have been added as well as the data pool. First off the electronic health record data and the previous and current diagnosis and treatment from the cancer registry have been added to the data provided by the patient profile. In combination with the data pool findings, these can be interpreted on a higher level letting the recommender suggest similar educational material and exercises to others if that turned out well in the past. The suggested companion material has also been added to the module/container data as this can impact the suggestion.

When the patient is done with a module, feedback is given through the screening and evaluation of the content. This now feeds both the patient profile and the data pool. The data pool would then compare findings to what is already stored and potentially adjust weights. The patient profile will use the feedback to update the information it has on the patient.

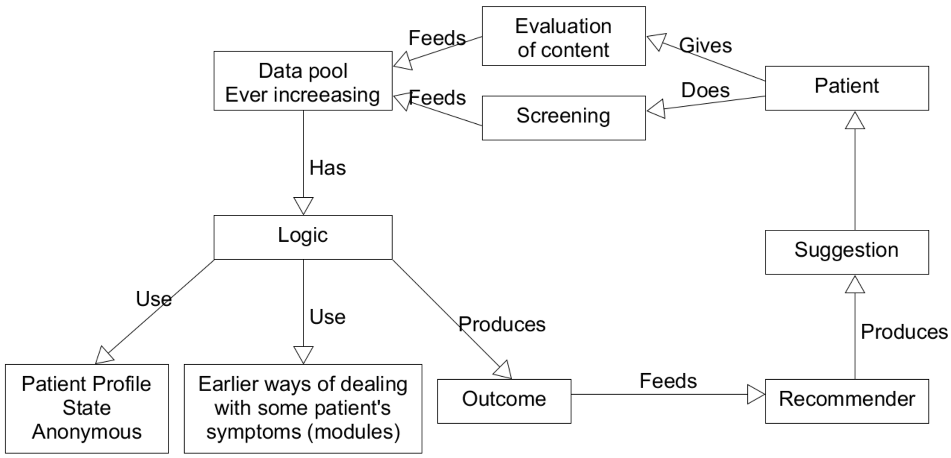

Figure 3.15: Model of the lowest level data pool feedback loop

The data pool is a storage of vast amounts of data and some logic connected to this data. Its potential use in an application like COPE would allow patient data to be stored anonymously. Inference on the data could reveal patterns in how certain groups of patients respond to particular learning material and exercises. This would then have a weight in the recommendation, which means that some material that did well in the past might be recommended based on the patient and other similar patients. Some issues here are obviously the storing of patient data and what data should be used to compare patients, more on this later.

Depicted here is the data pool itself, it has some logic that uses the patient profile and the history of how all other similar patients responded to particular material to produce some outcome. This outcome is an adjustment of weights for the different items suggested by the recommender and can be both positive and negative depending on what information is available. There would be a threshold to when this data becomes significant and what constitutes a similar group to the patient getting the suggested content.

At this point, the recommender produces a suggestion that is then worked on by the patient. The patient does the screening and evaluates the content that then feeds the data pool, and adjustments are made.

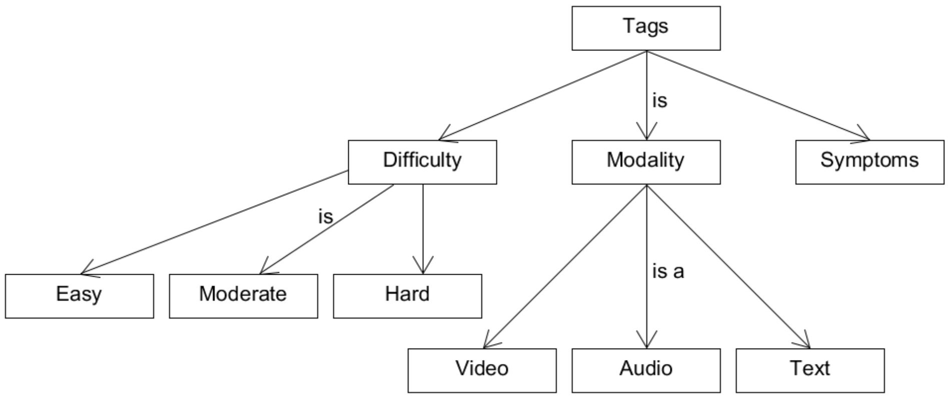

Figure 3.16: Model of the tags

A discussion was had about what tags could be used to define the content. The group and supervisors settled on a few tags that would undoubtedly be good candidates. The tags here are split into three. The difficulty of the content relates to the curiosity and prior knowledge of the patients. The easy difficulty would be something anyone without prior knowledge would be able to work on and hard being content that only people wanting to dig deeper into the technicality of their symptoms or how the treatment works would do. The modality refers to the modality of the content provided, this is something that the patient will sort out once they start using the application in their preferences. Ideally, all content would be tagged with all modalities, but if they are not, there might be some adjustment of weights to the content that supports the modality preference of the patient. The last tag category is the symptoms which categorise content by their symptoms. This is useful as patients might want to focus on the symptoms they are suffering from or symptoms impacting the problems they have.

The tags mentioned above are a small subset of what could be the reality of tagging the content. Tags can be anything that relates to the content to decide whether or not it would be a good match for any given patient.

There are probably tags that could be useful to categorise the content within the application that does not directly influence whether or not it would be suggested for a patient. Expanding on the tags would require assistance from professional psychologists, they would know how the content is different on a lower level giving more nuanced and accurate categorisation of the content that could give patients content truer to what they would receive in traditional treatment with a psychologist.

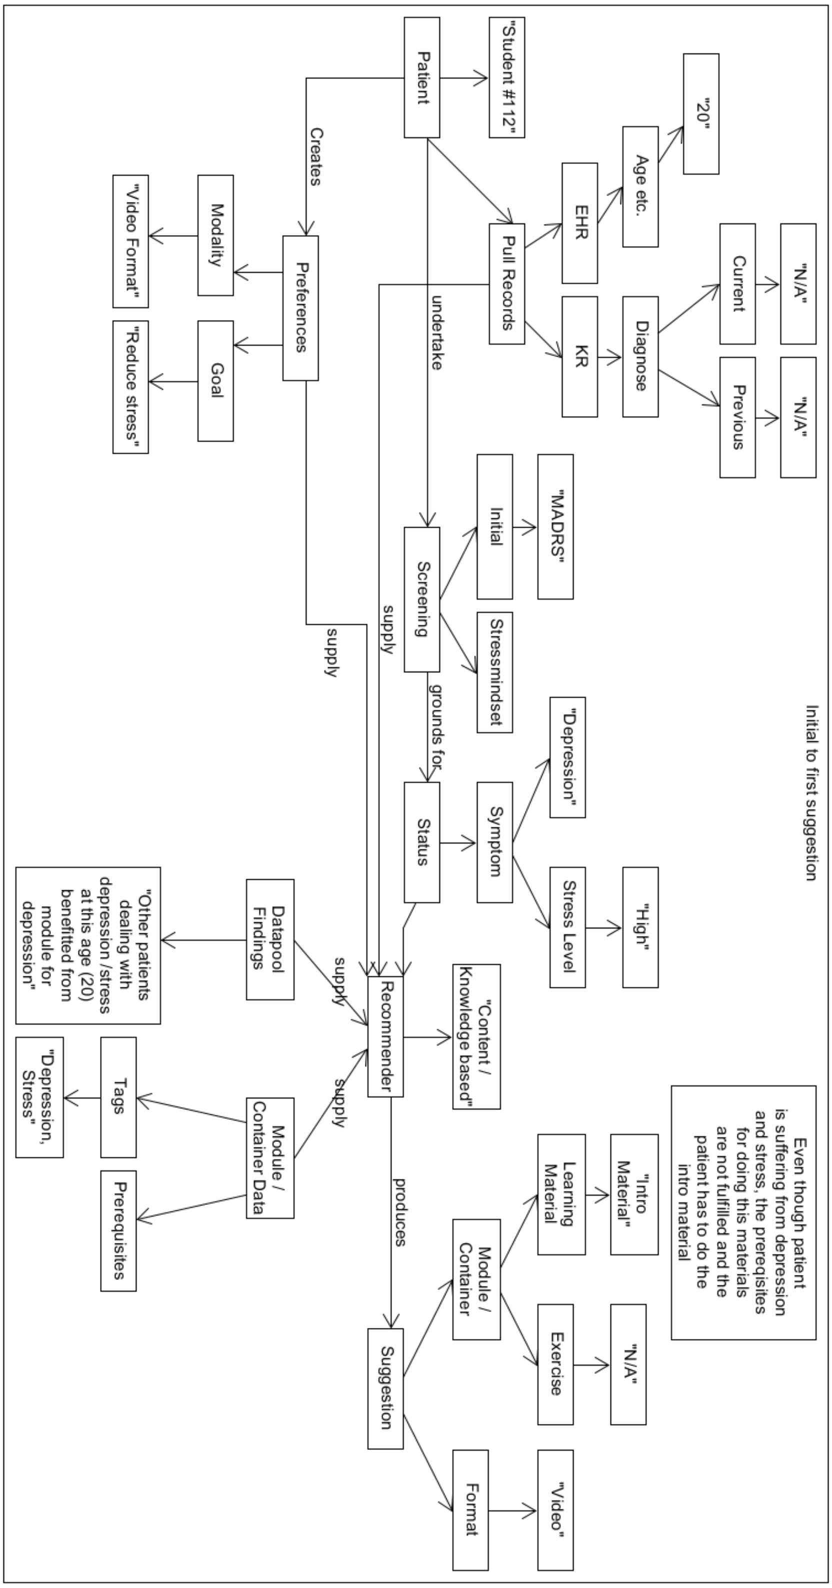

Figure 3.17: Process model of the patients initial use of the application

The process models were remade to reflect the changes in the lower level models and evaluate the work. Since the plan was to test the pilot on students as patients instead of women with breast cancer, the examples used in the process model is of a 20-year-old student dealing with depression and stress.

Initially, when the patients set up the application, two actions are required, and one happens automatically. The patient sets up their preferences with modality and goals and also have to do the initial MADRS and stress mindset scale screening. The records of the patient should be grabbed automatically as the patient sets up a user with some personal information before being allowed into the application, this step is not covered in the process model, it starts when the patient first accesses the app with a user account. The personal information being pulled is the electronic health record data, and when the application is going to be used with women dealing with breast cancer, the cancer registry data about diagnose and treatment will also be pulled.

The initial screening feeds the status of the patient, which will be refined over time. The status has now painted a picture of a patient suffering from depression with a high stress level. This status, together with the patient records and preferences, are fed into the hybrid content- and knowledge-based recommender. The recommender compares the data with the data pool findings, this could potentially reveal that similar patients benefited from doing a specific module. The recommender also has to confer with the module/container data to look for available content and prerequisites for doing said content. When this is done, the recommender produces a suggestion which in this case turned out to be the intro material as all the other content was locked. There are no suggested exercises as they would either be locked or the intro material does not have any exercises to go with it. The recommender found the material in the preferred modality, so the patient gets a video to watch.

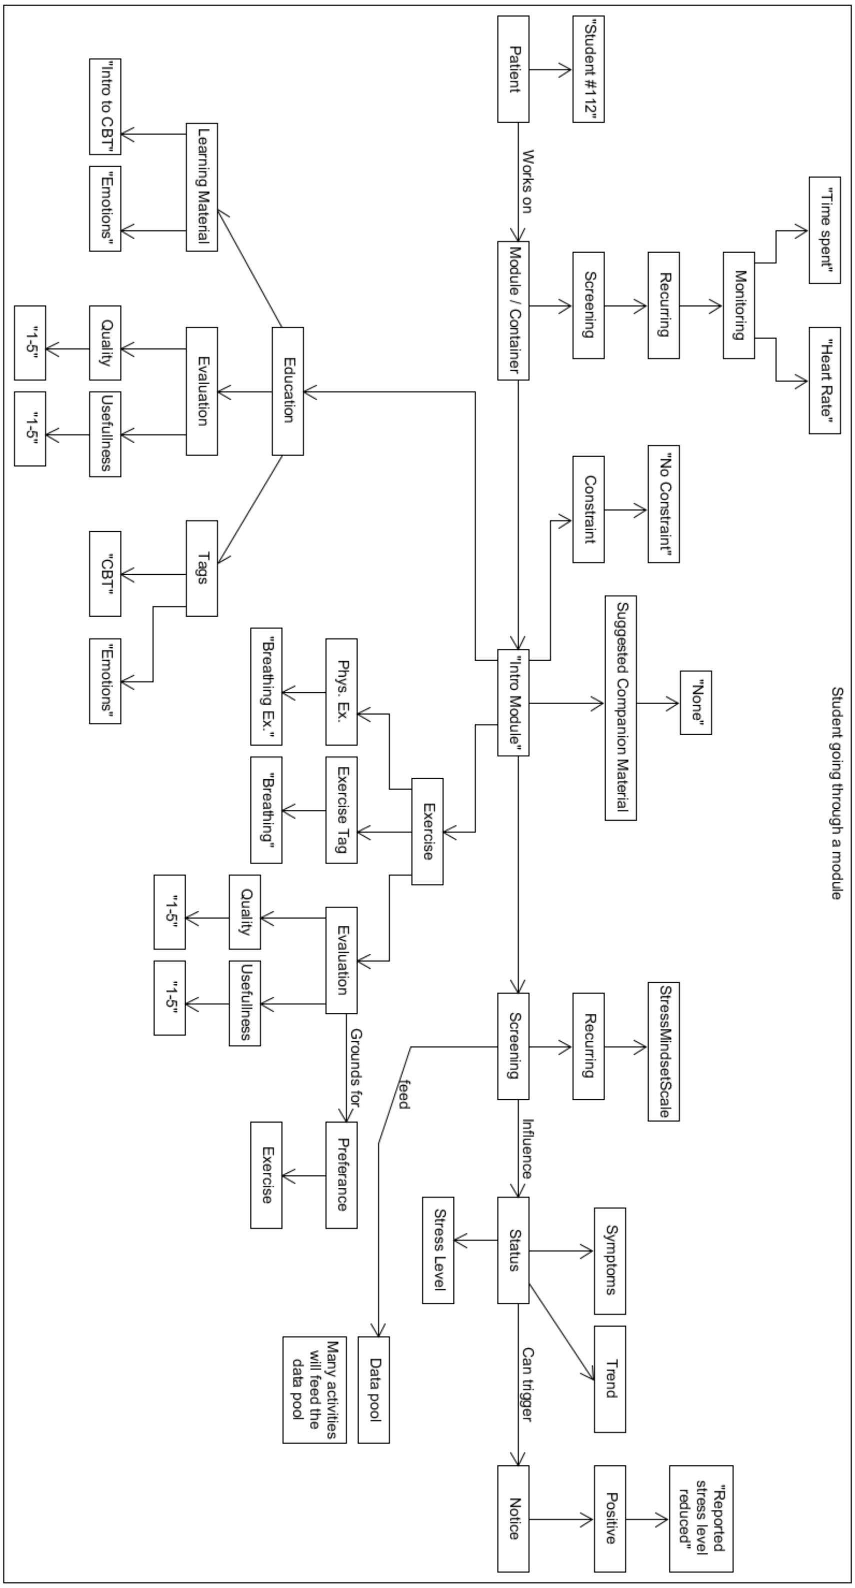

Figure 3.18: Process model of the patient working on a module

The process model shows a patient working on a module. As this is going on, screening is being done in the form of heartbeat and general navigation data monitoring. The patient is working on an intro module that contains some intro material to cognitive behavioural therapy and emotions. These are tagged accordingly and will be evaluated by the users' perceived usefulness and quality of the content. There is no constraints or suggested companion material for this content. If some suggested material were present, this information would be fed to the recommender later down the line and taken into consideration.

The module contains some physical exercise related to breathing, if the patient finds this exercise particularly useful and rates it as such, it could be added to the preferred exercises and would potentially be repeated more often than other exercises.

When the patient finishes the work on the material, the stress mindset scale screening will take place. This will feed the data pool and update the status of the patient. At this point, an encouraging notice could appear.

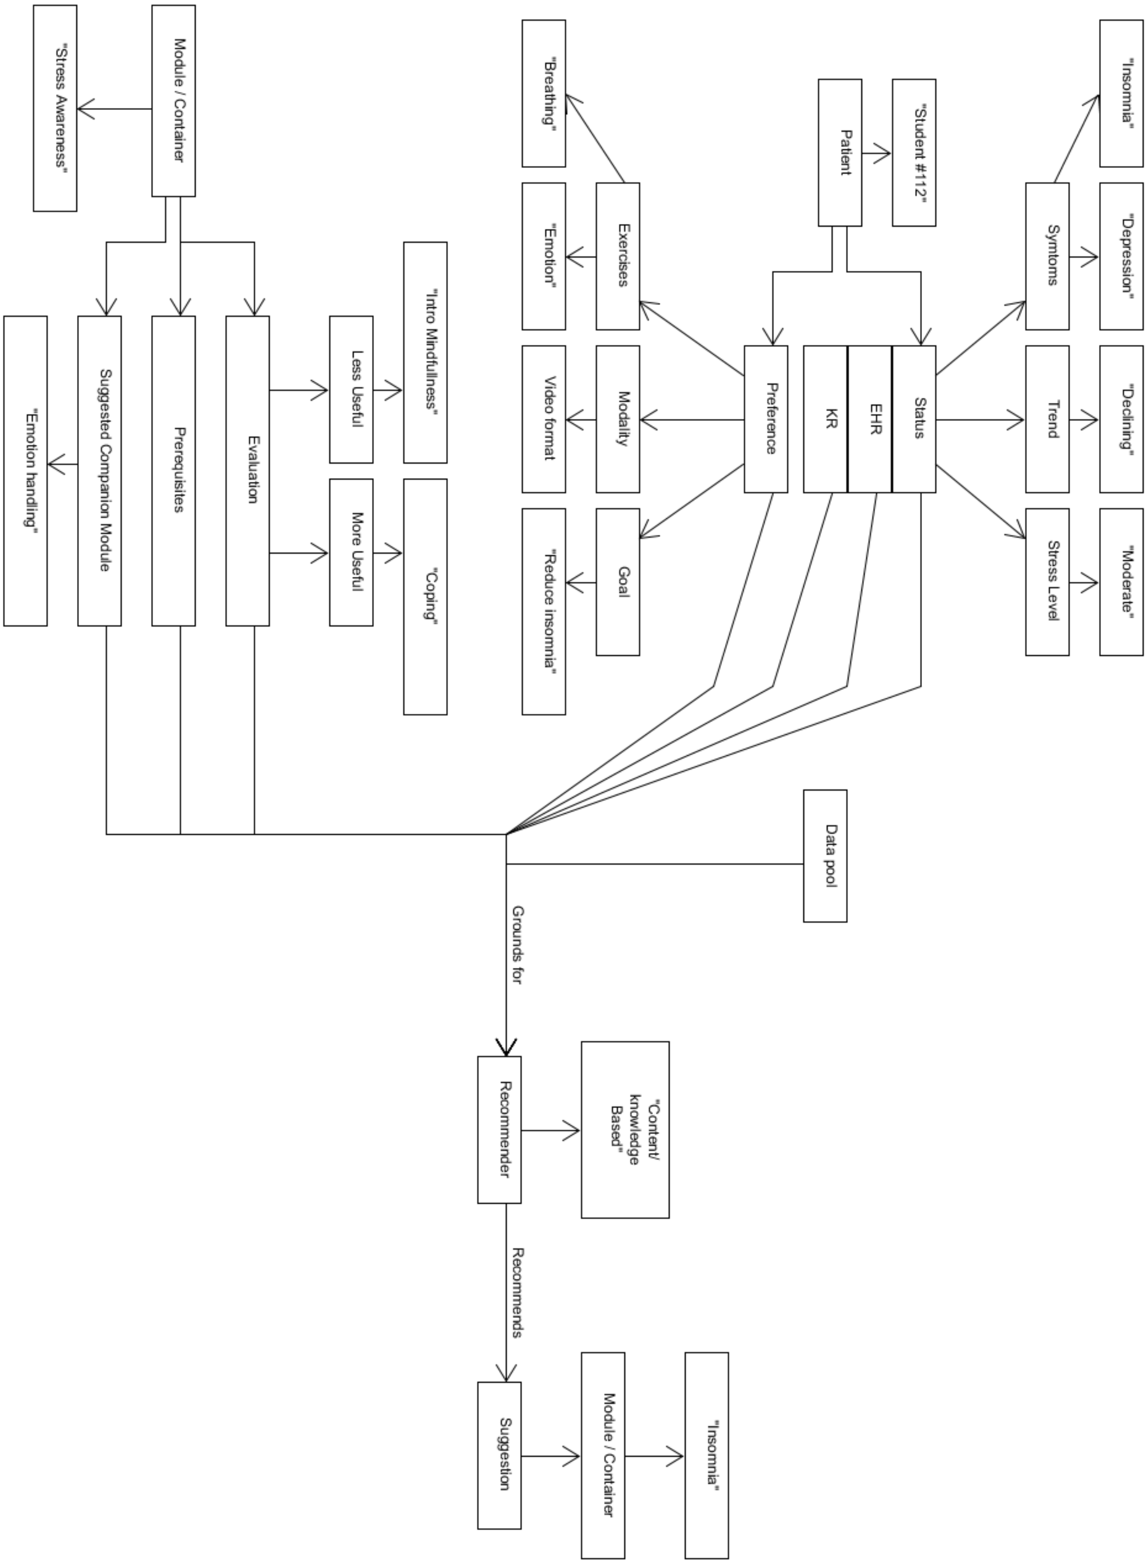

Figure 3.19: Process model showing what information goes into a recommendation

The model shows the information being used by the recommender. The information is mainly divided into the information that can be extracted from the patient, the module/container data and the data pool. Information about the patient includes the status, electronic health record data, the potential cancer registry data and the patients' preferences. The module/container data is concerned with the previous modules the patient has been through and will take into account things like an evaluation of previous modules, prerequisites for learning material and exercises and suggested companion modules. Information about how similar patient groups responded to different treatment could be evaluated by the data pool logic. All this information is then fed to the recommender as the data used to recommend some learning material or exercise the recommender also keeps track of the content it can recommend. The recommender is a hybrid between content- and knowledge-based.

4 Implementation

The previous chapter discussed the design process and described the artefacts produced. The implementation section will go deeper into detail discussing the tools used, parts of the model in greater detail and how student/learner models was used as inspiration for some of the decisions made.

4.1 Tools

Through the project three tools were used to develop the models, Lucidchart (*Lucidchart: Online Diagram Software and Visual Solution* 2020), Protege (*Protege, open-source ontology editor and framework for building intelligent systems* 2020) and Umlet (*UMLet 14.3 Free UML Tool for Fast UML Diagrams* 2020).

4.1.1 Lucidchart

Lucidchart is a tool used to make all kinds of models in a web service environment. It is used in the same way as the more known google docs, sheets or slides only for the development of models with a drag and drop interface. This tool was used in the initial iterations of the project to make some models that were not included with illustrations in the previous chapter due to these being scrapped early on as they were not accepted as ontology models. After this, some research was done to find out what was expected from an ontology model, and the second tool Protege was featured in a lot of the research.

4.1.2 Protege

Protege is a tool made at the Stanford Center for Biomedical Informatics Research at the Stanford University School of Medicine. It is a tool made specifically for developing and maintaining ontologies and provide multiple ways of developing these.

The first way is to plug in the information directly into the user interface. This is done by defining classes or concepts hierarchically and disjointing and binding these classes using object properties. Disjointing of classes is done to specify classes that cannot be of another class like the partner of a breast cancer patient can not also be its child or therapist. The object properties can be anything, but the most common would be the "is a" relation. An example from one of the ontologies made for the project would be the class breast cancer patient who "is a" type of the class cancer patient who "is a" type of the class patient. Protege also allows for comments on each class describing the class in more detail, this is common in ontologies as they often represent knowledge about a domain where a standard definition or understanding of each class is vital. The breast cancer patient is described as: "A cancer patient suffering from or previously suffering from cancer developed from breast tissue". The next part is setting up the data properties, these are comparable to the fields of a class like the first name of a person. The fact that health professionals, kin and patients are all persons in the hierarchy means that they would inherit the data property assigned to person. The data properties are also described to make sure that the meaning is interpreted equally by everyone using the model. In the case of the model, a national provider identifier was described as "A unique 10-digit identification number issued to health care providers in the United States". Individuals are made to exemplify the instantiations of the classes. Each individual is of a particular type, has object property and data property assertions and is described.

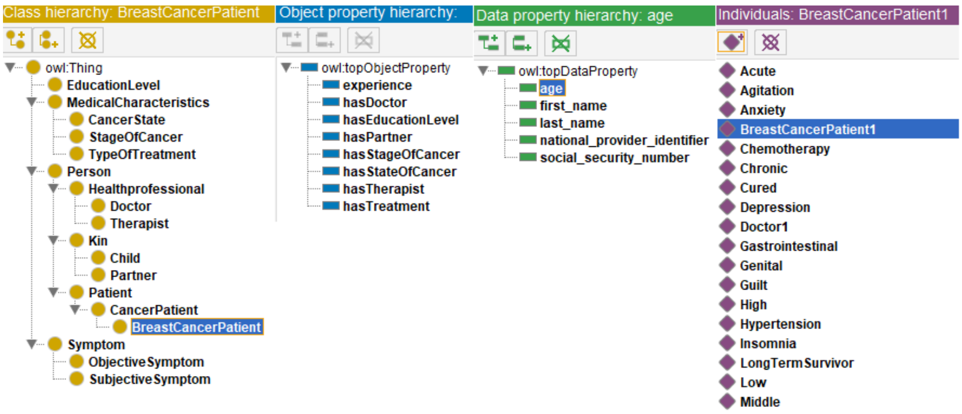

Figure 4.1: The main fields used to make an ontology in Protege

Making the ontology can also be done using the resource description framework (RDF) (*RDF 1.1 Concepts and Abstract Syntax* 2014) and the web ontology language (OWL) (*OWL Web Ontology Language Overview* 2014) both developed by W3C. These use an Extensible Markup Language (XML) (*Introduction to XML* 2020) setup with similar syntax. They can be imported and exported with the Protege tool. The languages were tested, but most of the development was done with the graphical user interface of the Protege program.

The problem with using Protege was mainly the presentation. When making an ontology and trying to convey the knowledge within a domain to all members of a project independently of their professional knowledge, representation is essential. As Protege did not have a satisfactory solution to show the resulting model, extensions like OntoGraf was used. These somewhat fulfilled the requirements by showing the model in a traditional way, but the parameters for viewing were limited, and it ended up looking disorganised. With this information and the supervisor not being fond of Protege, and instead suggesting the use of a more standardised way of making models like UML the decision was reached to use another tool.

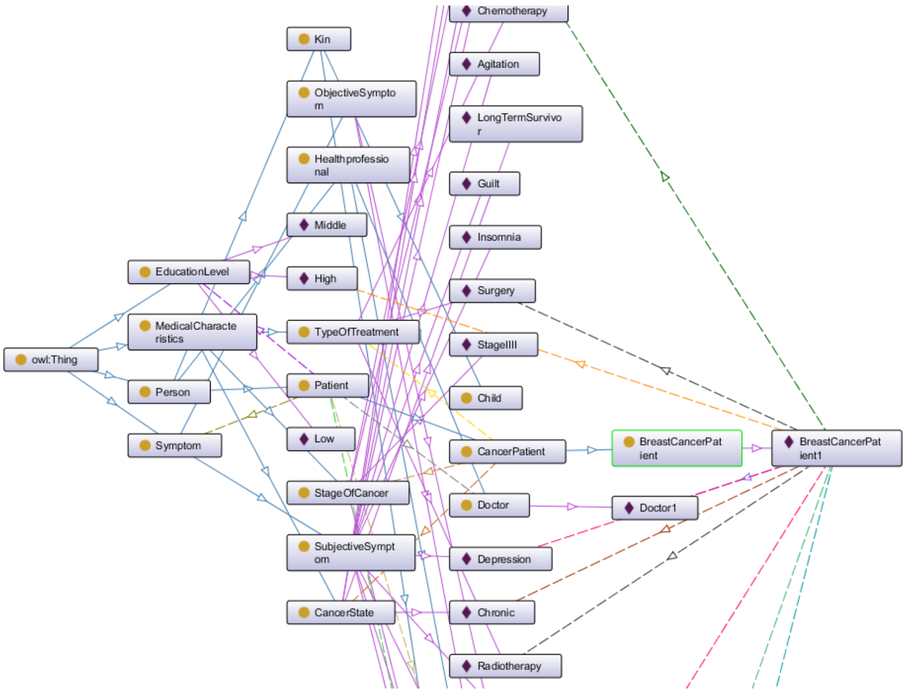

Figure 4.2: Visualisation of the Protege ontology model using the OntoGraf extension

4.1.3 Umlet

Umlet was used to make the final models, like Lucidchart it provides the user with many elements that can be used to make a variety of models. Unlike Lucidchart it is a desktop application running poorly when the size of the models exceed a certain number of objects. Umlet, however, is both free and open-source, which is the main reason it was used. With direct control over the presentation, it was easier to show the whole team the models. Divisions could be made to group the more significant components of the model instead of having all concepts connected and potentially making the model harder to understand. The Umlet components for building the models are quite extensive, so it easily filled the requirements set.

4.1.4 Using learner models as an inspiration for a patient model

As a patient model for use with adaptive internet-based cognitive behavioural therapy (iCBT) did not already exist, we decided to look at the learner/student models used in adaptive learning. iCBT and CBT in general use learning to support the therapy process. Learning about one's symptoms and ways to deal or cope with these is an integral part of the therapy. When reading about learner models, one thing became apparent, there are several people writing about them and how they should help the learning process, but few actually show any implementation of an actual model. This made it harder to develop the patient model as little was known about what the finished model would look like. This section will cover how some of the parts of the model was inspired by concepts found in the learner model domain.

In the paper "Learner Model in Adaptive Learning" Fröschl et al. (2008) state that every student have individual features that need to be accounted for when making a learner model. It became clear that many of these were similar or general enough to also fit into the patient model. Some of these concepts will be listed below, followed by an explanation of where it's been used or how it has been solved in the patient model.

- **Difference in knowledge:** The difference in knowledge between patients have been taken into account and is solved in the way that content has requirements and difficulty. Every patient is treated similarly when it comes to requirements, some learning material or exercise has to be done to unlock the next. This means that for some patients, the content may initially be below their knowledge level. This is remedied slightly by the suggested difficulty tag, where more knowledgeable patients could choose to get content that is more suited to someone interested or already knowledgeable in an area of psychology. The modules of greater difficulty would contain information that goes beyond the necessities, to give the patient a more refined and comprehensive explanation about a topic.
- **Goals:** Goals can be motivating and can help keep things on track when managing mental health Medicine (2020). In the model goals are solved rather directly by letting the patient fill in a more general free-text goal for their own motivation and a tag-based one with the intention of suggesting more relevant content. This is a part of the model that could have been expanded on to support short and long term goal setting and tracking goals over time. This is, however, not

an easy topic, especially for patients suffering from depression where there might be negative sides to setting goals as well (Street 2002).

- **Experience and background:** The experience and background are both broad ways of defining a patient and therefore fall into many parts of the model. Explicitly they can be tied to the electronic health record of the patient as well as the cancer registry report of previous diagnoses and treatment. They are, however, also tied to the patients' progress through the learning material and exercises and how they evaluate these. The aspect of time becomes an important factor as both their background information and experience change over time, and adapting the content based on how the patient is portrayed at the time, taking into account what the patient has been through is what makes the application adaptive. The experience and background will affect how the patient would fit into groupings of patients as this is where the data pool would gather the metrics to determine which patients are similar.

- **Interests:** Interests would be gathered both through the goals above and through the patient rating some learning material or exercise in high regard. Some patients might get content relating to their interests due to their background, experience or screening as well. Screening would give an indication of what the patient is suffering from which would suggest relevant content, this might be in the best interest of the patient, but not necessarily something the patient would find interesting, but should work through to potentially get better.

- **Learning styles and activities:** Preferences on the modality of the content and the choice to repeat exercises is in place to support the patients' needs and partiality.

Fröschl et al. (2008) goes on to talk about how from these features there emerges a demand for the tailoring of the learning material. In the model, this is the work of the recommender, which takes information about the patient and potential patient groups over time and recommends learning material and exercises to fit the individual patients' needs.

Some of the inspiration for the dashboard solution, where the status of the patient could be tracked by the patient itself, kin or potentially a therapist was gathered from Bull & Kay (2010) chapter on "Open Learner Models". The open learner model is meant to support learning by making information about the learner available to the learner, parents and teachers. The purpose of such a model is to promote activities like reflection, planning and self-monitoring. According to Bull & Kay (2010) it allows "the learner to take greater control and responsibility for their learning, encouraging learner independence" (Bull & Kay 2010). Among many other benefits, it allows the learner to access electronic data about themselves which they might have a right to do. It would be interesting to see if these potential effects on learners could be advantageous for patients in an iCBT setting. This is, however, one of the more controversial aspects of the model with regards to ethics and potential harm. It would have to be discussed carefully with psychologists to even consider. Even if it was considered as a positive feature, the extent of the information would have to be heavily regulated to ensure that it would not cause more harm than good.

Overall many aspects from the adaptive learning domain have been implemented in the model, which could make sense considering the importance of learning in cognitive behavioural therapy. It is important, however, to understand that there are probably other ways to make an adaptive iCBT patient model with or without the use of student/learner models and that these might be different from the model developed for this thesis. As it has not been done before it is hard to discuss the validity of basing such a model on others used in the learning domain, this will be discussed further in the demonstration and evaluation part of the thesis.

5 Demonstration and Evaluation

The evaluation is an integral part of design science research. Design science allows the creation of many different artefacts which in turn require the use of different forms of evaluation. According to Hevner et al. (2004), the "artefacts should be evaluated with criteria based on the requirement of the context in which the artefact is implemented" (Hevner et al. 2004) pg. 85. This is to check that the artefact supports the relevant qualities with regards to things like functionality, completeness, accuracy, usability etc. The actual evaluation method can be carried out in various ways depending on the artefact and its purpose, this includes methods like case studies, architectural analysis, controlled experiments, simulation, scenarios etc. The important thing is being able to "observe how well the artefact supports a solution to the problem" (Peffers et al. 2012) pg. 399.

In this section the demonstration and the evaluation will be discussed. As the final artefacts were all developed during the last three iterations, this section will be mostly focused on these. The artefacts or ontology models are evaluated using the process models to give a sense of the functionality, completeness, accuracy and usability of these. The method used is mainly using the process models as scenarios that the ontology models would encounter and have to cover to see if they suffice. These process models also have to be evaluated at a higher level as they have to represent a realistic image of a patient using the application. Once the process models are evaluated, and the ontology model have been shown to support these, the usefulness of the models are evaluated. The purpose of an ontology is sharing knowledge about a domain. Consequently it has to be evaluated to what degree it is capable of doing so.

5.1 Evaluation during the iterations

The evaluations during the project happened at SLATE with all members of the project present. These meetings were held once a month, and the researchers from the psychology and learning domain helped evaluate the artefacts. This was done by having all members of the team sit around the same table and showing them the models with an explanation of the elements, naming and thoughts put into the model as well as discussing how these conclusions were reached. The whole team was allowed to comment and discuss what they saw and heard which uncovered where there were issues with naming or where the model could benefit from further expansion or clarification. In some cases when discussing the model the team realised that something was missing, like in the early stages of the fifth iteration where it became clear to the psychology expert that we needed to include some kind of recurring screening that was less invasive than the MADRS test. Other times the issues were related to naming or structure. The evaluation was mostly done on the content of the actual models, but this also helped evaluate the usefulness of the model. The primary purpose of the ontology models is to give an understanding of the knowledge within the domain. While discussing the models, it became clear where the team benefited from the models and where they were confused by them, with this information, it was easier to see where more information or clarification was needed. Lastly, the models were evaluated in light of the process models. For this to work, the process models need to be somewhat accurate first. The process models need to be an acceptable representation of the patient interacting with the application and while working on them some aspects of the ontology model were developed further to be able to support the claims of the process models. The process models also contributed to the overall understanding of the application and sparked discussions that were helpful in resolving misconceptions.

5.1.1 The first iteration

The first iteration that produced the final artefact models was the fifth one. No process model had been made yet, and the models were evaluated to make sure they had a good initial representation of the patient and application, but the screening did, however, still need some work. This was when we first discussed using the stress mindset scale (Crum et al. 2013), which was added to the model. Another point that was made was that the recommender needed some attention next as it was not representing much. This is what the next iteration covered. Lastly, some process models had to be made to show the flow of the application. These would be used to evaluate the ontology models.

At the next meeting, the process models had been developed, still within the same iteration. These had expanded the model slightly as new concepts were discussed as they were made. The experts thought the models could represent a patient going through the application and that the models supported most of the functionality. There was, however, questions related to the meta-questions after a module was completed, and we decided to remove this step as the evaluation of the content and exercise covered this.

The models at this point had been useful for the team to see and discuss with many of the underlying concepts we had talked about at meetings represented in the same place. This also served as an introduction to the concepts of the project for some members, especially with the explanations given when presenting the model.

5.1.2 The second iteration

The second iteration was the development of the recommender part of the model and its process model. The consensus was that it was good that it was now in place and that the feedback loop was represented. It did, however, lack some of the patient information needed to give a suggestion. This was also reflected in the process model. When evaluating the usability of the model, it did start some discussions about how the recommender is supposed to be adaptive. Some naming issues of different classes were also resolved. Different papers would use different terminology, the project members decided on a common one to make it more explicit what concepts were the same or not.

5.1.3 The third iteration and final evaluation

With the lowest level of models made, they were represented to the whole team at a monthly meeting. The experts generally liked what they saw and used the process models to evaluate the ontologies. They found that a dedicated tag and data pool model was needed as well as discussing other naming and relationship details that needed clarification. With this information, the models were changed slightly, and the tag and data pool models were created.

The final evaluation was generally good. Some relationships could always be better, and some naming was never fully agreed upon, but other than that, the process models were representative of a patient using the application and the ontology models contained the necessary concepts for these to work. The general consensus was that the models helped in the understanding of the domain, especially within the early stages of the project. It also helped with the discussions as an anchor point for these.

6 Findings and Discussion

(Hevner et al. 2004) describes how design science research should provide research contributions by adding to the knowledge base. In this section, these contributions will be discussed.

6.1 The research questions and contributions

The research questions represent the areas where the design process has given contributions to the knowledge base, these will be answered below, with some discussion.

6.1.1 RQ 1: How can a patient model based on learner models contribute to making a personalised adaptive iCBT application?

As discussed earlier, a patient model in use with an adaptive iCBT application has not been made before. This introduced a challenge as it was hard to know roughly what the final result would look like. This is the reasoning behind having some guidelines from the adaptive learner models. The learner models were described in multiple papers, but none of them included any kind of implementation of a model that could be dissected where parts could be reused in the model. Therefore discussion with the team and reading about different aspects of learner models was critical.

Some of the team members were already aware of or knew more about learner models from before, which made it easier for them to understand the aspects of the patient model when presented with it. As discussed earlier, the main reason for making the ontology models was representing the knowledge within the domain. The model could also be used more directly in some cases where it made sense, thereby having a direct contribution to the making of the personalised adaptive iCBT application. The process models further explained the models in an intuitive way, less dependent on the structure of the model but rather the structure in which a patient would use and go through modules in the application.

Few would argue against the convenience and importance of a model when working on projects in the IT industry. In this case, the models helped with understanding, structure and as an entry point for discussing parts of the application. The fact that it is based on learner models was both somewhat of a necessity in that it had to be based on some knowledge about adaptive models and that the learning domain is not that far from what was needed, as well as it being an extension of the knowledge that parts of the team were already familiar with.

6.1.2 RQ 2: How are unguided adaptive and non-adaptive iCBT applications different?

Typically a non-adaptive iCBT application would take all patient though the same modules of content from start to end. In this case, the patients do not have any say in what they are doing and may have to go through content they already know about, are not interested in et cetera.

To combat this, one could tailor the content to a specific patient group like patients who have insomnia. Studies show that this type of tailoring usually work well but have some limitations in that they usually work only for the specified groups. It is essential to decide how accurately the iCBT application should be tailored to the group, the more it is tailored, the better it might work for the specific patients but the less versatile it would be. It might work for patients from a specific country suffering from the exact symptoms, but differences in aspects like culture could mean that it is less suited elsewhere.

Common in these non-adaptive applications is that they are more straightforward in design and require next to no patient model. The patient model, in this case, could do more simple tasks like tracking the patients screening scores which would make sense if the application is monitored by a therapist that could get a notification if a patient is showing suicidal tendencies.

An adaptive iCBT application tries to solve the problems of non-adaptive applications in various ways. The application can still be tailored to a patient group, but depending on what the application takes into account, this group can be less specific, the more detail the adaptiveness of the application covers. The COPE project wants to make an adaptive iCBT application focused on women dealing with symptoms after breast cancer treatment. This might sound very specific, but the patient group can deal with widely different symptoms and can benefit from adaptive treatment. The difference between tailored and adaptive in this case would be that tailoring is something that is done before the use of the application and does not necessarily change with use. This can be done to the content of the application, where generic content for a symptom might not give the same benefit as content tailored to the domain of women dealing with symptoms after breast cancer treatment. What makes the application adaptive is the use of a recommender and its supporting features like the patient model in this thesis. This means that the application continually gathers and uses data to give the patient content that should be well suited not unlike guided iCBT applications. Unlike guided iCBT applications, the application still lacks human interactions, which can be both preferable or a downside depending on the patient.

An adaptive application gathers a whole lot of information over time and builds a robust model of the patient to be able to feed the recommender with as much and accurate data as possible this is then used to give the patient content that makes sense to the team of experts that weight the different data on how much they should count towards a recommendation. This is the main benefit of an adaptive iCBT application. If the content should be delivered to patients in a different place or of a different culture the tailoring of the content may have to be changed and possibly also the weights or which data goes into the recommender. This problem would still be a factor like in the unguided application, however, once the content and weights are adjusted the patients would still get the content they need or want from a recommender and will not have to go through the content that is not relevant to their needs.

6.1.3 RQ 3: How can a patient model for use in a personalised iCBT application be evaluated?

As there are no adaptive patient models made for iCBT there has been nothing to compare the model to or improve upon directly. As discussed earlier, the closest models already developed is in the field of adaptive learning. This means that the models could be loosely evaluated in the light of these but with the acknowledgements that they are not used for the same purpose. Another problem with this method is that very few of the research papers covering adaptive learner models show any actual implementation and is more concerned with explaining the principles behind how the models work.

In the case of this thesis, the evaluation of the models was done by creating process models that would mimic real-world usage of the application and show what parts of the model was used where and when. Experts in the fields of psychology and learning helped validate that the models were adequate. What was not done, but might have been beneficial to the evaluation was real-world testing. Had the COPE application as a whole been fully put together integrating the work from all four theses to make a final product, real-world testing could have been done, and this might have shown flaws in the process models altering these slightly and thereby potentially also affecting the ontology models. As the application and project are still in a prototype stage, this was not done.

As these are ontology and process models, they can also be evaluated in light of their usefulness to the development team. How well they conveyed the knowledge of the domain could play a part in the overall success of the development. In the case of the thesis, this type of evaluation was more of a continuous one done by experts in the fields of modelling, psychology and learning.

6.1.4 RQ 4: How are patient models and learner models different in the light of the thesis?

The thesis has shown many instances where patient and learner models are similar to substantiate the claims that an adaptive patient model could be based on adaptive learner models. It is, however, essential to know that there are also differences between the two. The main goal in learner models is maximising learning, the main goal of a patient model is increasing the well being or improving the mental health of a patient. The information needed for the different tasks is different, where a patient model needs additional information about the patients' health. If the recommender is going to succeed in giving good suggestions for content for a student, the model needs to cover previous knowledge and how to build upon it. A patient model would have to balance educating the patient and giving him/her the necessary content to get better, potentially slowing down the learning process.

The main focus of the models is also different when it comes to screening and testing, where the testing of a learner shows how much they have learned whilst patient screening show how much they have improved or declined in well being. In iCBT treatment, a goal is that the patient learns about symptoms like discussed, but it is not measured in how much they know but rather how much it helped.

When looking at more practical examples of differences, the information flow and security of data can also be different. When talking about open learner models, the learner, teacher and parents might have information about the learner and learning through some dashboard. This can be used to do different tasks like planning, reflection and self-monitoring. In the case of a patient model, this information would be highly restricted, and any dashboard solution would have to be crafted to detail as not to be damaging to the user. The models made for the thesis show that a dashboard could be implemented, but it is discussed as something that should not be taken lightly and would need the expertise of experts in the field of psychology.

Overall it makes sense to use adaptive learner models as inspiration for an adaptive patient model in use with iCBT when few options are available, and no such model exists beforehand, it is important however to distinguish between the two and make adjustments thereafter.

7 Conclusion and Further Work

This thesis presents an adaptive patient model for use with iCBT in the form of ontologies divided into three layers of abstraction. It also presents process models used to conceptualise and evaluate these. The purpose of the models is showing an example of such a model based on learner models and can further help research in the field of adaptive patient models in use with adaptive mental health treatment by sharing knowledge within the domain. This section will give a summary of the results and suggestions for further work.

7.1 Conclusion

The work done for this thesis has been conducted using the design science method and have consisted of designing and implementing ontology models for knowledge sharing and a final adaptive patient model. The models have been evaluated by experts using process models to simulate a patients' run through the COPE application. During the design, three sets of models were created with decreasing abstraction levels. The models are based on learner models as no adaptive patient model for use with iCBT existed at the time. The design of the models have both impacted and been impacted by the work of other master thesis projects within the COPE project.

Peffers et al. (2012) claims that the artefacts should use existing knowledge to develop a solution to a defined problem. The existing knowledge, in this case, is mainly that of previous iCBT projects and learner models. The defined problem has been the need for an adaptive patient model for use with iCBT and, also, expanding the teams knowledge within the domain of adaptive patient therapy. The design and implementation of the models has been an attempt at solving the first half of this problem. The presentation and discussion of the models during meetings has been the solution to the second problem.

7.2 Further Work

COPE is a project that aims to create an application with adaptive iCBT treatment for women dealing with mental distress symptoms after breast cancer. At the time of writing, four master students have worked on different parts of the application creating a basic prototype. These parts have to be brought together to make a final product. This includes the models, the CBT content, the recommender and the application front end.

As for the models, these have to be further developed and tested in a real environment with help available from experts in psychology. As the scope of students dealing with stress was chosen as a case to be able to specify and limit the work, it's possible that the models need to include more breast cancer-specific concepts that might be lacking at this time. Sets of entity relation diagrams would have to be made based on the ontology models as ontology models are separate from implementation details. This would give the developers working on the recommender and front end what they need to implement the application. The models would have to support ethical and security-based technicalities.

Once an application is made and working, it could be expanded to be able to feed anonymous data in a proper format for use in research. This have to be done in multidisciplinary groups to be able to make a standard for the data that researchers would agree upon and keep security and ethics in high regard.

References

Abrahams, H. J., Gielissen, M. F., Donders, R. R., Goedendorp, M. M., van der Wouw, A. J., Verhagen, C. A. & Knoop, H. (2017), 'The efficacy of internet-based cognitive behavioral therapy for severely fatigued survivors of breast cancer compared with care as usual: A randomized controlled trial', *Cancer* **123**(19), 3825–3834.

Aguilar-Saven, R. S. (2004), 'Business process modelling: Review and framework', *International Journal of production economics* **90**(2), 129–149.

Antoni, M. H. (2016), 'Video-conferenced stress management and relaxation training for older women with breast cancer'.

Bodden, D. H., Dirksen, C. D., Bögels, S. M., Nauta, M. H., De Haan, E., Ringrose, J., Appelboom, C., Brinkman, A. G. & Appelboom-Geerts, K. C. (2008), 'Costs and cost-effectiveness of family cbt versus individual cbt in clinically anxious children', *Clinical child psychology and psychiatry* **13**(4), 543–564.

Børøsund, E., Varsi, C., Clark, M. M., Ehlers, S. L., Andrykowski, M. A., Sleveland, H. R. S., Bergland, A. & Nes, L. S. (2019), 'Pilot testing an app-based stress management intervention for cancer survivors', *Translational behavioral medicine* .

Bull, S. & Kay, J. (2010), Open learner models, *in* 'Advances in intelligent tutoring systems', Springer, pp. 301–322.

Cisler, R., Holder, H. D., Longabaugh, R., Stout, R. L. & Zweben, A. (1998), 'Actual and estimated replication costs for alcohol treatment modalities: case study from project match.', *Journal of Studies on Alcohol* **59**(5), 503–512.

Crum, A. J., Salovey, P. & Achor, S. (2013), 'Rethinking stress: The role of mindsets in determining the stress response.', *Journal of personality and social psychology* **104**(4), 716.

Davenport, T. H. (1993), *Process innovation: reengineering work through information technology*, Harvard Business Press.

Farrer, L., Christensen, H., Griffiths, K. M. & Mackinnon, A. (2011), 'Internet-based cbt for depression with and without telephone tracking in a national helpline: randomised controlled trial', *PloS one* **6**(11).

Fröschl, C., Nguyen, L. & Do, P. (2008), 'Learner model in adaptive learning', *World Academy of Science, Engineering and Technology* **21**.

Grieg, N. A. (2019), A visual analytics dashboard for mental health therapists, Master's thesis, The University of Bergen.

Gruber, T. (n.d.), 'Ontology'.

Halpin, T. (2009), '"ontological modeling (part 1)" business rules journal vol. 10, no. 9'.
URL: *http://www.brcommunity.com/a2009/b496.html*

Hevner, A. & Chatterjee, S. (2010), Design science research in information systems, *in* 'Design research in information systems', Springer, pp. 9–22.

Hevner, A. R., March, S. T., Park, J. & Ram, S. (2004), 'Design science in information systems research', *MIS quarterly* pp. 75–105.

Introduction to XML (2020).
URL: *https://www.w3schools.com/xml/xml$_w$hatis.asp*

Lindner, P., Ivanova, E., Ly, K. H., Andersson, G. & Carlbring, P. (2013), 'Guided and unguided cbt for social anxiety disorder and/or panic disorder via the internet and a smartphone application: study protocol for a randomised controlled trial', *Trials* **14**(1), 437.

Lucidchart: Online Diagram Software and Visual Solution (2020).
URL: *www.lucidchart.com*

Medicine, D. C. M. (2020), 'Goal-setting', *University of Michigan* .
URL: *https://www.depressioncenter.org/toolkit/i-want-stay-mentally-healthy/goal-setting*

Morales-Rodriguez, M. L., R.-S. J. A. H.-R. A. S.-S. J. P. . F. J. A. M. (2012), 'Architecture for an intelligent tutoring system that considers learning styles', *Res. Comput. Sci. 47* pp. 37–47.

Morgan, C., Mason, E., Newby, J. M., Mahoney, A. E., Hobbs, M. J., McAloon, J. & Andrews, G. (2017), 'The effectiveness of unguided internet cognitive behavioural therapy for mixed anxiety and depression', *Internet interventions* **10**, 47–53.

Morlandstø, N., Hansen, C., Wasson, B. & Bull, S. (2019), 'Aktivitetsdata for vurdering og tilpasning: Sluttrapport. slate research report 2019-1, bergen, norway: Centre for the science of learning i& technology (slate).', *AVT Slate* .

NHS (2019), 'Cognitive behavioural therapy (cbt)'.
URL: *https://www.nhs.uk/conditions/cognitive-behavioural-therapy-cbt/*

Noy, N. F. & McGuinness, D. L. (2000), 'Ontology development 101: A guide to creating your first ontology'.

Ontario, H. Q. et al. (2019), 'Internet-delivered cognitive behavioural therapy for major depression and anxiety disorders: a health technology assessment', *Ontario health technology assessment series* **19**(6), 1.

OWL Web Ontology Language Overview (2014).
URL: *https://www.w3.org/TR/owl-features/*

Peffers, K., Rothenberger, M., Tuunanen, T. & Vaezi, R. (2012), Design science research evaluation, *in* 'International Conference on Design Science Research in Information Systems', Springer, pp. 398–410.

Peffers, K., Tuunanen, T., Rothenberger, M. A. & Chatterjee, S. (2007), 'A design science research methodology for information systems research', *Journal of management information systems* **24**(3), 45–77.

Protege, open-source ontology editor and framework for building intelligent systems (2020).
URL: *https://protege.stanford.edu/*

RDF 1.1 Concepts and Abstract Syntax (2014).
URL: *https://www.w3.org/TR/rdf11-concepts/*

Ricci, F., Rokach, L. & Shapira, B. (2011), Introduction to recommender systems handbook, *in* 'Recommender systems handbook', Springer, pp. 1–35.

SLATE (2019), 'The cope architecture', *Internal SLATE/UiB project report.* .

Sottilare, R. A., Graesser, A., Hu, X. & Holden, H. (2013), *Design recommendations for intelligent tutoring systems: Volume 1-learner modeling*, Vol. 1, US Army Research Laboratory.

Street, H. (2002), 'Exploring relationships between goal setting, goal pursuit and depression: A review', *Australian Psychologist* **37**(2), 95–103.

Twomey, C. & O'Reilly, G. (2017), 'Effectiveness of a freely available computerised cognitive behavioural therapy programme (moodgym) for depression: Meta-analysis', *Australian & New Zealand Journal of Psychiatry* **51**(3), 260–269.

UMLet 14.3 Free UML Tool for Fast UML Diagrams (2020).
URL: *https://www.umlet.com/*

Vaishnavi, V. & Kuechler, B. (2004), 'Design science research in information systems', *Association for Information Systems* .

Vigerland, S., Lenhard, F., Bonnert, M., Lalouni, M., Hedman, E., Ahlen, J., Olén, O., Serlachius, E. & Ljótsson, B. (2016), 'Internet-delivered cognitive behavior therapy for children and adolescents: a systematic review and meta-analysis', *Clinical Psychology Review* **50**, 1–10.

Webb, C. A., Rosso, I. M. & Rauch, S. L. (2017), 'Internet-based cognitive behavioral therapy for depression: current progress & future directions', *Harvard review of psychiatry* **25**(3), 114.

Zachariae, R., Amidi, A., Damholdt, M. F., Clausen, C. D., Dahlgaard, J., Lord, H., Thorndike, F. P. & Ritterband, L. M. (2018), 'Internet-delivered cognitive-behavioral therapy for insomnia in breast cancer survivors: a randomized controlled trial', *JNCI: Journal of the National Cancer Institute* **110**(8), 880–887.