

Designing a 3D visualization for exploring relations between semantically similar brain diseases in Augmented Reality

A Master's thesis by Thijs Hendrickx

**Supervised by Lynda Hardman
Second examiner: Wolfgang Hürst**

Submitted in partial fulfillment of the requirements
for the degree of Master of Science in Game & Media Technology



Utrecht University

Department of Information and Computing Sciences
Utrecht University
The Netherlands
20 July 2021

Designing a 3D visualization for exploring relations between semantically similar brain diseases in Augmented Reality

Thijs Hendrickx

Utrecht University

ABSTRACT

As more papers are published and the body of knowledge increases, it becomes difficult to maintain a sense of overview and recognize relations between discussed topics. In order to assist neuroscientists with their task of finding unexplored relations between brain diseases, the DatAR project provides a topic model which allows neuroscientists to see at a glance which brain diseases are semantically similar and thus likely correlated. However, this topic model lacks the ability to explain **why** and **how** these brain diseases are semantically similar. We set out to design the Brain Disease Co-occurrence Explorer (BDCE), a 3D visualization which visualizes the number of co-occurrences between two brain diseases and any other brain topic in literature, with co-occurrence defined as two topics mentioned in the same sentence or body of text. We design two iterations of the BDCE, the first of which is evaluated by interviewing two experts in the field of 3D data visualization and immersive analytics. Based on the results from these interviews, we improve upon the design and perform user evaluations of the final iteration with participants from a neuroscience background or different related biology field. The user evaluations show that the BDCE is capable of explaining the topic model and encouraging further investigation into potentially interesting relations between brain diseases. Finally, we propose some avenues for future research regarding additional BDCE features, different visualization designs for visualizing co-occurrences in 3D space and improvements to our number visualization style.

1 Introduction

As more papers are published and the body of knowledge increases, it becomes difficult to maintain a sense of overview and recognize relations between discussed topics. The DatAR project (*DATA exploration in Augmented Reality*) aims to assist researchers with their literature research by providing various visuospatial data representations of topics discussed in literature, in an Augmented




Figure 1: Topic model for showing the semantic similarity between brain diseases in the DatAR environment [1]. Each sphere in the model represents a brain disease, with the disease name written on a white label in front of the sphere.

Reality (AR) environment [3, 1]. Neuroscience is one such research field where understanding complex relations is increasingly assisted by linked, text-mined data exploration [4]. Neuroscientists often conduct literature studies to find unexplored connections between brain topics in order to direct their research focus. The DatAR project aims to assist neuroscientists with this task by providing a topic model visualization that maps brain diseases in 3D space based on their semantic similarity (see Figure 1). This semantic similarity is calculated based on the total number of direct and indirect co-occurrences in literature, with co-occurrence defined as two topics mentioned in the same sentence or body of text. The topic model allows neuroscientists to see at a glance which brain diseases are semantically similar and thus likely correlated, assisting them with their task of finding unexplored connections between brain diseases. However, the topic model lacks the ability to explain **why** and **how** these brain diseases

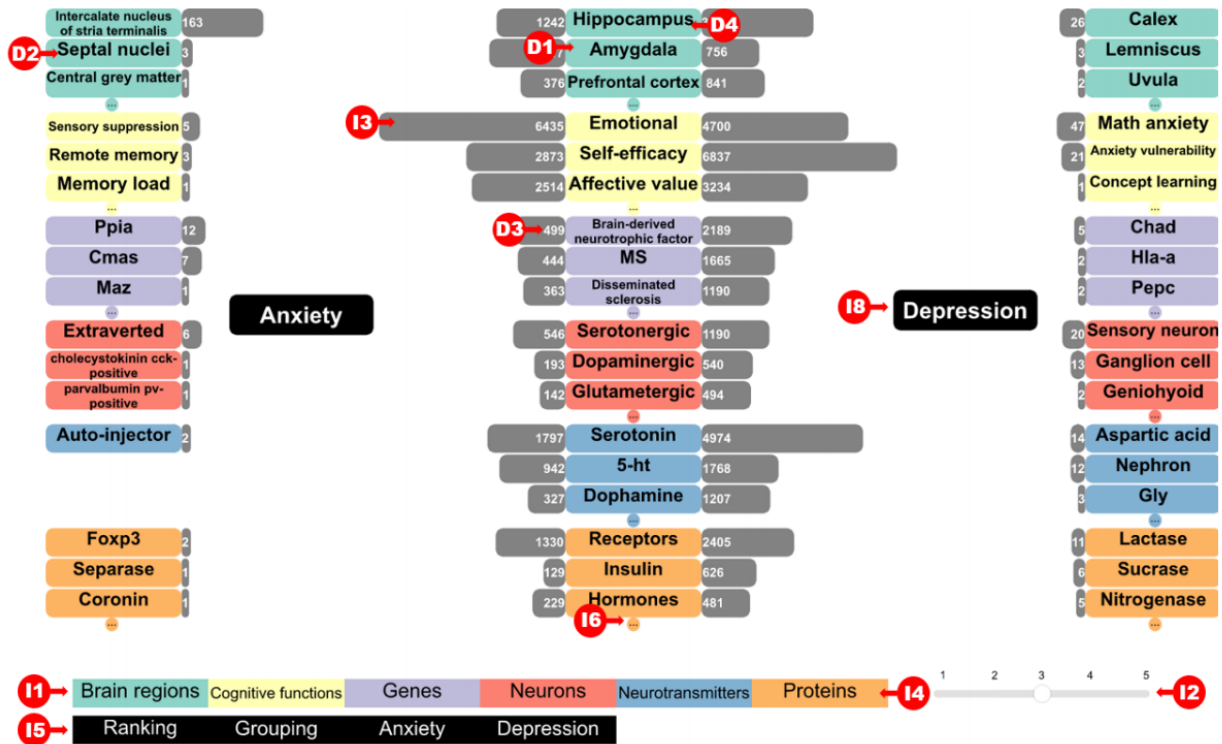


Figure 2: The Buiuter2D visualization design for exploring the semantic similarity between brain diseases [2].

are semantically similar. The details of the direct and indirect co-occurrences are not visualized in the topic model, which obscures the connections in literature that make the diseases similar. These connections are essential for understanding how brain diseases are correlated.

Buiter [2] contributed a design for a 2D data visualization that allows neuroscientists to better understand the relations between similar brain diseases (see Figure 2). While this 2D design is capable of explaining the semantic similarity of brain diseases, it lacks a 3D design and corresponding evaluation in the immersive 3D DatAR environment which neuroscientists use to inspect the topic model. We set out to design the Brain Disease Co-occurrence Explorer (BDCE), a 3D data visualization which visualizes the number of co-occurrences between two brain diseases and other brain topics mentioned in literature. The design for the visualization is largely on Buiter’s 2D design since it takes the same approach on how to explain the semantic similarity between brain diseases. The BDCE allows us to evaluate our ability to explain the similarities of brain diseases as shown in the topic model and encourage further investigation into these similarities. In summary, we aim to answer the following research question: **How do we explain and visualize brain disease similarities as shown in the topic model in an immersive 3D DatAR environment?**

¹Buiter [2] refers to *brain topics* as *brain concepts*. This was the previous term used throughout the DatAR project. They mean the same.

For sake of clarity, throughout this document we will refer to Buiter’s 2D data visualization as *Buiter2D*, the first BDCE iteration as *BDCE-1* and the second (and final) BDCE iteration as *BDCE-2*.

2 Related work

Since our BDCE design is largely based on the Buiter2D design, we explore the design requirements set out for the Buiter2D design. We additionally explore the topic of immersive analytics and the challenges of moving a 2D design over to 3D space. Finally, we present some 3D data visualizations used as inspiration for our design.

2.1 Retrospect of the Buiter2D design requirements

Buiter set out to formulate design requirements for an interface that would allow neuroscientists to better understand the semantic similarity of brain diseases as seen in the topic model. By visualizing the number of co-occurrences between brain diseases and brain topics¹ in literature, Buiter hypothesized that a neuroscientist can compare the co-occurrence rates to find unexplored connections between the brain diseases. The design require-

ments and visualization sections for Buiuter2D (see Figure 2) were refined and evaluated by design and visualization experts, as well as neuroscientists. We summarize the design requirements as formulated by Buiuter in the following three categories.

The desired information for comparing similar diseases should be available

The core requirement of the visualization is to *show the number of co-occurrences of any given brain topic with two selected diseases*. The visualization should in some way *highlight brain topics which co-occur with only one disease*. Users should be aware that *co-occurrences shown in the visualisations represent correlation, not causation*. Additionally, the visualized co-occurrences do not indicate the nature of any correlation (positive or negative). Lastly, users should be able to *investigate the articles that feature the presented co-occurrences in an intuitive manner*.

The visuals should be clear and devoid of clutter

There are various key aspects to keeping the visualization readable and clutter free. *Text should at all times be readable without clutter; and number of co-occurrences needs to be shown through visual cues apart from numbers*. *There should also be a clear visual distinction between brain classes², the brain topics and the selected diseases*. Lastly, *hidden parts of the visualization need to visually relay the fact that they are hidden*.

Interactions should conform to the expectations of the user

The interactions and tools available to the user are an integral part of the visualization. Since the target audience using this data visualization are experts in their research fields, they must be provided with the tools that allow them to retrieve the information that they want based on their expertise. The design requirements specify that *users must be able to rank the relevance of related topics (sorting), temporarily hide and show individual brain classes (filtering) and be able to visualize initially hidden co-occurrences between a topic and any topics it relates to*. This last requirement indicated the ability to explore co-occurrences spanning multiple brain topics.

2.2 Immersive analytics

Traditionally, information exploration was primarily done through manual searches with a desktop setup on a flat 2D screen. The field of visual analytics was introduced

a decade ago and developed techniques to represent data, primarily big data, in a more visual way. This had the benefit of building upon the human mind's ability to understand complex information better visually [5]. Recently, a new sub-field has developed called immersive analytics. This field takes advantage of the significant advances in the development of new technologies for Human-Computer Interaction (HCI) by investigating how new interaction and display technologies can be used to create a more immersive kind of data analysis and exploration [6]. One example of new technologies that allow immersive analytics are VR and AR Head Mounted Displays (HMDs) which allow the user to completely immerse themselves in a 3D data visualization. AR also supports the anchoring of these visualizations to the real world. DatAR takes advantage of this with some of its design ideas, such as real-time interactions between the virtual elements of DatAR and the user's desktop setup.

Since the 3D BDCE design is based on the Buiuter2D design, the simplest approach to make the transition between the two dimensional representations is by designing the BDCE as a 2.5D³ visualization in the immersive 3D DatAR environment. Marriott et al. note that one of the main benefits of immersive visualizations is the utilization of depth (the third dimension) [8]. Designing the BDCE as a 2.5D design adds this utilization of depth, allowing for a more immersive kind of experience while analysing data.

2.3 3D data visualizations of multivariate data

Visualizing information in an immersive 3D environment can be achieved in a multitude of ways. Cordeil et al. introduce *ImAxes*, an interactive multi-dimensional visualization tool that relies on the arrangement of data axes in space [9]. Users can freely move around available data axes and connect them in multiple ways to construct new types of graphs to view the data. Supported graphs include histograms (see Figure 3), 2D/3D scatterplots, parallel coordinates and more. The way *ImAxes* represents histograms is similar to how the horizontal bars are designed in Buiuter2D, and can be similarly utilized in the 2.5D BDCE design.

²Brain class refers to categories for brain topics, e.g. brain regions, neurons and cognitive functions.

³The term 2.5D is quite uncommon and ambiguous in its definition. We use the definition of Ware [7] which defines it as 2D with depth cues which provide a sense of 3D. The additional dimension is not used for higher dimensional data visualization and purely serves to provide a sense of depth and increase sense of orientation.




Figure 3: Histograms in the ImAxes environment [9].

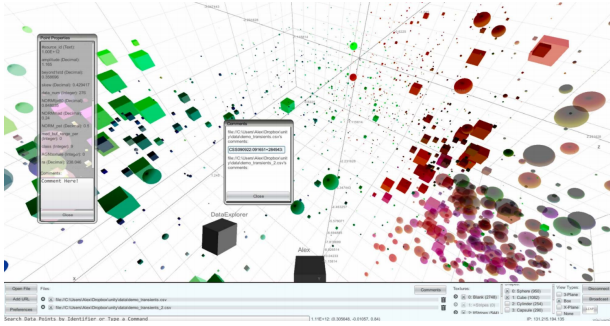


Figure 4: The iViz environment, showing hyperdimensional data encoded as 3D points with with various parameters such as shape, color and size [10]

Donalek et al. present iViz, a 3D data visualization tool for visualizing hyperdimensional data [10]. In iViz, data is visualized not only through XYZ spatial coordinates but also with additional parameters such as shape, color and size (see Figure 4). Additionally, iViz offers collaborative data exploration by allowing multiple users to view and manipulate the same environment. Their use of color and shape to encode hyperdimensional data serves as an inspiration for the BDCE design.

3 Methods

We followed an iterative design process, split into two cycles. During the first cycle, we designed BDCE-1 and evaluated our design approach by way of interviews with experts in the field of 3D data visualization and immersive analytics. The design for BDCE-1 was based on Buitier2D, with various alterations for better accommodation to a 3D immersive environment. Due to the lack of access to appropriate AR hardware, the implementation of BDCE-1 was built upon the existing DatAR implementation for VR.

In order to evaluate the design approach for BDCE-1, we made use of the expert interview methodology. Expert interviews have a long history in social research, as an efficient and concentrated method of gathering information compared to other methods such as quantitative surveys [11]. This methodology lends itself well to evaluating systems in both a qualitative and quantitative manner, which is why we decided to utilize this methodology. Our interview procedure consisted of two parts. During the first

part, the interviewee is unaware of BDCE-1 and Buitier2D. This allows us to ask questions without their answers being skewed towards the existing design. During the second part, the interviewee is shown BDCE-1 which then allows us to discuss the design direction.

The second design cycle iterates on BDCE-1 by incorporating the feedback and advice from the expert interviews. We port the system from the VR version of DatAR to an AR version supported by appropriate hardware, resulting in final BDCE design iteration: BDCE-2. In order to answer the research question and evaluate the BDCE-2, we perform user evaluations of BDCE-2 in AR with neuroscientists and researchers with a research background in biology. The user evaluation consisted of two parts: a questionnaire and a set of guided tasks using the AR device. Users would answer questions on the usability and data-interpretation aspects of BDCE-2 while performing the tasks. With the results from the expert interviews and user evaluations we answered the research question.

4 BDCE-1 visualization design

We describe the design rationale behind the first iteration of the BDCE, and discuss its implementation and data set.

4.1 BDCE-1 design rationale

BDCE-1 was designed to visualize co-occurrences between brain diseases and other brain topics in an immersive 3D environment. While BDCE-1 was based on the Buitier2D design, we made various alterations to better accommodate the design to an immersive 3D environment.

One major concern was the scalability of the design, and by extension its vertical structure. Buitier2D shows at most five brain topics per class and does not scale well with more topics, as this would increase the verticality of the visualization. In immersive AR, vertical field of view is limited so it is important that the vertical bounds of BDCE-1 are kept to a minimum. We opted to go for a hierarchical design structure in the form of a "brain class view" (top of the hierarchy) and a "brain topic view" (bottom of the hierarchy) (see Figure 5). The initial view showed each brain class as a singular sphere with the total number of co-occurrences that the topics within that class had with the diseases on either side. The user could then interact with the class spheres by "grabbing" them with their hand. This would change the view from a class overview to a topic overview of the class that had been selected. This view is similar to the overview demonstrated in Buitier2D, with a filter applied so only topics from the selected class would be displayed. By separating the overview seen in Buitier2D into two views, we introduced potential for scalability since the topic view had

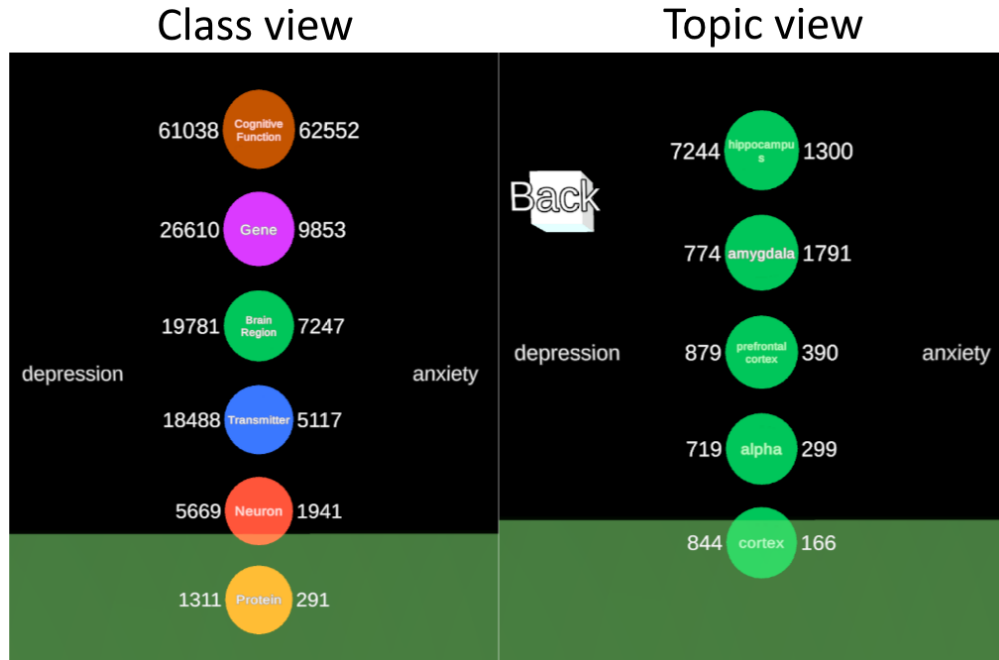


Figure 5: Implementation of BDCE-1 in the VR DatAR environment. The two selected diseases are on either side of a central vertical list. The numbers on either side of the spheres represent the number of co-occurrences that those brain classes or topics have with either disease. **Left:** The brain class view. Users can interact with a sphere to go to the brain topic view of that class. **Right:** The brain topic view of brain regions. Users can interact with the "Back" cube to return to the brain class view.

fewer topics shown at any given time. Users could return to the class overview by "grabbing" the *Back* cube on the top-left of BDCE-1.

In order to reduce design complexity, we also opted to not include the discriminatory topics as designed in Buiuter2D (see D2 in Figure 2). The discriminatory topics are topics that only co-occur with one of the diseases, rather than both. Previous evaluations of Buiuter2D concluded that the participants remained largely confused as to the purpose of the discriminatory topics, or thought they were insignificant and not that interesting [2]. For BDCE-1 we decided to only show the topics that have the most co-occurrences with the diseases, which allowed us to focus the expert interviews on the brain class and topic representations.

4.2 BDCE-1 implementation

As we lacked access to AR hardware at this point in the study, BDCE-1 was implemented as a widget in the established VR DatAR framework in Unity⁴ [1]. The widget accepts two brain diseases and subsequently generates a BDCE-1 visualization with the two diseases on either side. Because BDCE-1 visualizes co-occurrences between brain diseases and other brain topics, we needed a dataset that contained these co-occurrences. The Linked

Brain Data (LBD)⁵ repository contains co-occurrences of known topics discussed in neuroscience literature from various catalogs, including PubMed. This data can be queried through a SPARQL endpoint. Slotting two diseases in the widget sends a SPARQL query to the LBD SPARQL endpoint, in order to retrieve the number of co-occurrences that the selected two diseases have with any brain topic stored in the LBD. This data then gets populated across BDCE-1, which appears above the widget cube after the query has completed.

5 Expert interviews

In order to evaluate our design approach for the BDCE-1, we interviewed two experts in the field of information visualization and immersive analytics. In this section we describe the interview setup and discuss the results from the two interviews, after which we summarize the changes that we made to the design direction.

5.1 Interview setup

The aim of the expert interviews was to acquire creative ideas which could contribute to the current design or alter the design direction, as well as feedback on the cur-

⁴<https://unity.com/>

⁵<http://www.linked-brain-data.org/>

rent design. The interviewees were first presented with a slideshow, detailing the aims of the DatAR project and our role within. They were then shown a sample of the data from the LBD repository which we were working with, after which the participants were asked how they would visualize this sort of data in 3D space. This allowed us to acquire new design insights and potential BDCE-1 adjustments from an unbiased perspective. Afterwards, the experts were shown the BDCE-1 in a video demonstration of a probable use-case where the visualization could be used. They were then asked to have a conversation about the design direction of BDCE-1.

We specifically directed the conversation to the following limitations that we had identified with the design prior to the interviews:

- While BDCE-1 addressed some of the concerns regarding vertical scaling by splitting the class and topic views, the topic view still had the potential to show too many topics vertically from the same class. We asked the experts for their opinions on how to approach this issue.
- Since BDCE-1 lacked a visual representation of the co-occurrence numbers, we took this opportunity to ask the experts for their opinion on how they would visualize these numbers in our design.

The interviews were recorded and subsequently summarized into specific points of feedback.

5.2 Interview 1

The first interviewee commented during the presentation on DatAR and the LBD why brain diseases are not considered a brain class together with the 6 other classes. This is a valid observation, as in normal circumstances they would be considered a class. The primary reason that brain diseases are not being included as a class for comparison with the two selected diseases is that there are simply not many direct co-occurrences between diseases. While not negligible, for ease of implementation we excluded brain diseases as a class in our design. The interviewee also noted that brain classes have inherent differences in the number of literature available, as this is field dependent. For example, there is a larger body of work on Cognitive Functions compared to Neurons. According to them, a comparison across classes becomes less interesting in that case and in-class comparisons are the most interesting way to view this data. Though if there were some way to show the global distribution within each class, the topics could be comparable across classes. When asked how they would visualize data from the LBD repository in 3D, the interviewee suggested that there were two possibilities. You can show several diseases and their co-occurrences with different colors, and some way to indi-

cate the strength of the relations. This would show the relative strength of each co-occurrence. Or you can single out one specific disease, and show in a glyph or spider chart a number of co-occurring topics, with again ways to indicate the strength of those relations. Since these suggestions are vastly different from our design, they can instead serve as inspiration for future work which can look at different ways of visualizing co-occurrences in 3D.

After showing BDCE-1, the interviewee suggests to make the data an "extension" of the visualized topics, with the ability to see more detail on a specific topic. Having the ability to store a previous search query/selection would be helpful as well. Utilizing the third dimension for this purpose could be an idea. When asked how to handle showing more than five topics in a vertical list, the interviewee asked if we were sure that neuroscientists needed to view more than five topics at the same time. With proper sorting and organization tools allowing neuroscientists to view the data in the way they want, it might not be necessary. Lastly, they commented on the hierarchical manner in which the data is displayed. According to them, the design lacked a clear connection between the class and topic views. They noted that there are visualization methods in 2D which are designed to compare hierarchical data structures, and these methods clearly signal the connection between data points.

5.3 Interview 2

After asking the interviewee how they would visualize the data from the LBD repository in 3D, the interviewee mentioned some kind of 3D scatter plot with disease, topic and class on the three axes. Each point on the scatter plot could indicate the number of co-occurrences with their size. After showing them the BDCE-1, they immediately suggested using visual encoding of the numbers to show the magnitude. Perhaps with horizontal bars. Furthermore, they suggested looking into ways to better accomplish hierarchical visualization instead of two separated views, such as Sankey diagrams. BDCE-1 lacked visual connection between the classes and the topics, which hierarchical visualizations solve. They further noted a desire to store topics of particular interest after navigating the hierarchy, which then allows the user to navigate freely through the visualization again without losing the previous topic of interest. After being asked how they would handle showing more than five topics in a vertical list, the interviewee suggested a movable selector similar to a fish-eye lens that the user can move past a list of topics which are represented as a chart. The selector could then show details of the topics besides the chart, perhaps including some sentences from literature where these co-occurrences happen directly.



Figure 6: The BDCE-2 design implemented in AR. Lights in the room were turned off for better visibility of the visualization in the photo. **Left:** Total view, showing the total number of co-occurrences on either side of the central brain topic spheres. **Middle:** Ratio view, which shows the percentage of co-occurrences between either disease for each topic. **Right:** Ratio view with the Brain Region filter applied, showing the top five brain regions in terms of total number of co-occurrences with both diseases.

5.4 Interview summary & Design decisions

We compile the results from the interviews in the following points, and discuss our design decision on each point.

Hierarchical design

One of the interviewees stated that they had doubts about our hierarchical design approach which lacked clear connection between the class view and the topic view. Since Buitier2D also did not incorporate this hierarchical structure, we decided to step away from this hierarchical approach and change the design to the "overview" style from Buitier2D (see Figure 2).

Displayed topic limit

With the design moving away from the hierarchical style, our initial concern of the vertical limit topic limit was once again relevant. However, the interviews made it clear that we needed to re-evaluate the need for displaying more than five topics at once. We decided to not change our design to support more than five topics, and instead ask neuroscientists during the user evaluations of BDCE-2 if they desired the ability to view more than five topics at once. If the results showed that they would, we could suggest the fish-eye lens proposed by the second interviewee as a possible solution to be evaluated in future work.

Number visualization

The comments given by the interviewees on the lack of a number visualization for the co-occurrences made it clear that this was an area of priority to be implemented in BDCE-2. One interviewee suggested a standard histogram visualization, also seen in Buitier2D. We decided to go with this approach.

Normalized ratio view

The comment on the inability to compare topics across classes showed us the need for some kind of normalized ratio view to make topics comparable regardless of research scale. We decided to add a ratio view to the design which instead of showing the raw number of co-occurrences, would show the relative percentages of co-occurrences between the two brain diseases.

Annotated history of search queries

Both interviewees expressed their desire to view an annotated history of search queries performed with BDCE-1. We decided to not include this in BDCE-2 since it had no precedence in Buitier2D, but this could be explored in future work.

6 BDCE-2 visualization design

We proceeded with the next iteration of the BDCE design, including the design decisions formulated from the expert interviews. In this section we describe the design rationale behind BDCE-2 and briefly discuss the process of porting the DatAR VR environment to support AR, namely the HoloLens 2 HMD.

6.1 BDCE-2 design rationale

With our decision to step away from a hierarchical design, BDCE-2 shifted to a single overview which lists the top three brain topics of each class in terms of co-occurrences

(see Figure 6). With this change, the filters that were integrated in the BDCE-1 now became separate options. The filter options were displayed on their own separate interaction panel on the left of the widget cube as buttons, and when pressed would hide the overview and display five topics of the selected class in a larger font. While more topics could be displayed while these filters were enabled, we proceeded with five since our aim was to ask neuroscientists participating in the user evaluations if showing more than five is desirable. Finally, we introduced a ratio view which shows the relative percentages of co-occurrences between the two diseases. This view allows for a normalized comparison between the co-occurrences of the two diseases across brain classes. The ratio view could be toggled through the interaction panel on the right of the widget cube.

6.2 Porting DatAR from VR to AR

While the DatAR system had been tested in the past using a Meta 2 HMD, this hardware proved lacking in hand tracking capabilities and the position of virtual objects did not remain static while the user moved their head. Furthermore, support had been dropped since 2018 with the company behind it no longer existing. We therefore performed a short study to compare current available HMD options on various criteria such as Field of View, tracking capabilities and development support. The Hololens 2⁶ fulfilled the selected criteria and was accessible to us, thus we chose to develop for that platform. The full comparison study can be found in Appendix A.

In order to develop for the Hololens 2, the Microsoft Mixed Reality Toolkit (MRTK)⁷ for Unity can be used. This is an officially supported library for Unity, providing all the tools for developing for Microsoft Mixed Reality HMDs such as the Hololens 2. Due to its vastly different build requirements compared to SteamVR⁸ in which the VR DatAR framework was built, we created a separate Unity project for the Hololens 2 version instead of adding onto the existing framework. This would save time, and the VR version would become obsolete anyway once all the core framework functionalities had been ported over to support the Hololens 2 since the team will have regular access to a Hololens 2 for performing evaluations.

With the core functionalities of the DatAR framework ported to the MRTK, BDCE-2 was implemented using the same process as BDCE-1 (see Section 4.2).

7 User evaluations

In order to evaluate BDCE-2 and answer our research question, we performed user evaluations with participants from a neuroscience background or different related biology field. In this section we describe the evaluation setup and discuss the results.

7.1 Setup of the user evaluations

We invited twelve participants to take part in an evaluation of BDCE-2 in AR, using the Hololens 2. These participants were selected on their background in neuroscience or a different related biology field. We first asked participants to fill in a questionnaire to collect information on their personal background and VR/AR experience. The participants were then shortly briefed on the project with a 5-minute video demonstrating the DatAR system in VR, and an image of the BDCE-2 visualization which they would shortly be seeing in AR. Afterwards, the participants were asked to put the Hololens 2 on their heads. We streamed the view from the Hololens 2 to a laptop, allowing us to watch along and record their experience. To allow the participants to become familiar with the AR environment at their own pace, they began their Hololens 2 experience with a tutorial session on how to use hand gestures to interact with the AR environment.

After the tutorial session, the participants opened the DatAR application and began the BDCE-2 evaluation. They were provided with the following scenario:

You are a neuroscientist looking for potentially interesting relations between brain diseases to research further. You decide to use the DatAR system to help you find new relations.

The participants were given tasks to complete while using the system and were asked to answer questions during their experience. We decide to provide the instructions and questions orally, which is identified by Satkowski et al. as a viable method for AR studies, especially for systems that require complete use of the interaction set in AR [12]. They do note that this method can make it harder for the participants to think about the question and is susceptible to distracting noises from the surroundings. By conducting the evaluation in a silent room, noise concerns can be eliminated but participants might still have difficulty thinking about the question while being distracted by the AR environment. To minimize this distraction, we made sure to ask questions in-between tasks to allow the participants to focus on the questions.

⁶<https://docs.microsoft.com/en-us/hololens/hololens2-hardware>

⁷<https://docs.microsoft.com/en-us/windows/mixed-reality/>

⁸https://valvesoftware.github.io/steamvr_unity_plugin/

The tasks can be summarized in the following objectives:

1. Generate a topic model for brain diseases
2. Find the *depression* and *anxiety* disease spheres
3. Generate a BDCE-2 visualization to compare *depression* and *anxiety*
4. Apply filters & observe the two different views

The participants were guided step by step through the tasks, and asked questions at specific points throughout the experience. The questions can be split in two types: **Usability questions** and **Data-interpretation questions**. The usability questions aimed to get information regarding the HCI aspects of the system; how usable is BDCE-2 in AR? The data-interpretation questions aimed to see if BDCE-2 succeeded in explaining the topic model, and if it conveyed its visualized data in an understandable manner. For a full breakdown of the tasks and the questions, see Appendix B.

7.2 Results of the user evaluations

7.2.1 Participant's background & VR/AR experience

The participants came from a wide variety of backgrounds, including five different nationalities and eight different research fields. All but one participant (who had a background in data science for healthcare) had a background in neuroscience or different related biology field with close to no computer science overlap. We saw this reflected in the VR/AR experiences, as only three participants indicated that they used VR once or twice before and no participant had used AR before. Nine of the participants were female, three were male. Two of the participants held the title of Principal Investigator (PI), while the others held a Master's degree or lower. Three of the participants were already familiar with DatAR, with two of these participants having taken part in a DatAR study before.

7.2.2 Usability of the BDCE-2 visualization

In this section we discuss the results from the usability study on BDCE-2. We have to note that one of the participants seemed to be unable to accommodate their eyes to the virtual objects, causing the objects to remain blurred and out of focus for them. Chakravarthula et al. investigate this issue and propose auto-focus AR glasses that adjust optical focus based on gaze fixation [13]. Since the participant happened to be highly familiar with the project and the BDCE designs, they were able to perform the tasks and answer the questions without being able to read the text.

Readability of the topic model & BDCE-2

The participants were first asked to generate and observe a brain disease topic model. Most thought the model was clear and readable in AR, with a few comments on how they can move their head closer to read it clearer. We also asked the participants about the clarity and readability of the BDCE-2. Most participants thought it was clear and readable, however a few noted the text in the topic spheres was too small even when they moved their head closer.

Meaning of the topic sphere colors

When asked what the colors of the topic spheres represented, we got a mix of responses. Seven participants instantly made the connection with the colors on the filter panel and recognized them as brain classes. One of the participants was colorblind so they did not see much of a difference between the colors to begin with, this was not something we had taken into consideration. The other four started guessing the brain classes based on the topics that they saw. Most of them guessed correctly, but our aim was for the filters in the filter panel to be the clear identifier of what the colors represented. Once we pointed this out to the participants, they instantly recognized the connection.

Usability of the filter panel

When it came to the filter panel, the buttons were quite intuitive to use with nine participants immediately succeeding in enabling and disabling the filters. One participant required an explanation on how to disable the filter after enabling one. Some participants tried enabling multiple filters at the same time, as they expected this to be a possibility.

Usability of the views panel

Most participants also understood the purpose behind the views panel, with seven participants correctly identifying the two views. Some however misunderstood the Total view as showing the number of publications or citations, rather than the number of co-occurrences. Two participants misunderstood the Ratio view, depicting it as a view which showed the relative proportions (e.g. for X number of co-occurrences with the first disease there are Y with the second disease) rather than a percentage. One participant commented that a percentage sign behind the numbers in the ratio view would make this clearer.

7.2.3 Interpreting the data in BDCE-2

In this section we discuss if participants were able to extract and interpret meaningful data from BDCE-2

Connection between BDCE-2 and the topic model

After the participants generated and explored the BDCE-2 for depression and anxiety, they were asked if they found it logical that depression and anxiety were close together in the topic model based on the numbers visualized in BDCE-2. Most thought it was, with nine participants noting that both diseases share a similar number and ratio of co-occurrences with the same topics. Most of these participants based their logic primarily on the co-occurrences with cognitive functions, with one participant noting that in other topics the relation was less obvious as there were larger differences between the numbers on either side. Three participants were of the opinion that this larger difference in numbers on either side of most brain topics made it difficult to draw logical connections between BDCE-2 and the positions in the topic model. One participant noted that since they had no idea how the visualization would look with diseases that are far apart in the topic model, they could not immediately draw connections between the topic model and BDCE-2.

Six of the participants who saw logic in the connection between BDCE-2 and the topic model were also asked if they could guess how the BDCE-2 visualization would look with diseases that are far apart in the topic model. Five of them predicted (correctly) that the bars would be a lot less symmetrical, with larger differences between the numbers on either side. One participant instead expected the numbers to be a lot smaller overall.

Usefulness of the Total and Ratio views

When the participants were asked about their thoughts on the usefulness of the Total and Ratio views for understanding the similarity between brain diseases, they almost unanimously answered that it was useful. Eleven participants argued that both views are necessary; the Total view to get a sense of scale of the research done on the relations between diseases and any topic, and the Ratio view in order to better compare the two diseases indiscriminately from the number of co-occurrences found. One participant was unsure about the uses for a ratio view.

Motivation to further investigate interesting relations

As a final question while still wearing the HoloLens 2, the participants were asked if anything stood out to them from the data that they saw in BDCE-2 that might warrant further research. The intention behind this question was seeing whether the participants were becoming inspired to pursue further investigation into a relation.

Ten participants did see interesting aspects to the data, mostly in the brain regions (the Amygdala has more co-occurrences with anxiety than depression as the only topic in the list) but also with proteins since there were so few co-occurrences in that area. Two participants did not see anything of interest in BDCE-2, with one commenting that they were not a neuroscientist thus lacked field knowledge to make any interesting connections.

Opinion on the five displayed topics limit

In accordance with our conclusions from the expert interviews, we asked the participants if they thought viewing a maximum of five topics per class was enough or if they would have liked to see more. Responses were mixed, with seven participants indicating that they would like to have the option to see more if it was properly supported with the ability to scroll through the list. One of these participants only wanted to see a total of around ten topics, as they argued that more topics loses a sense of overview and adds a chance that you fall into a rabbit hole of topics that do not matter. The other five participants thought five topics was enough, arguing the same points of losing a sense of overview. One participant expressed that we are already facing more literature than we know what to do with, so by seeing more you lose focus on the most relevant topics. Most participants did express a desire to be able to sort and filter the list of topics, and be able to search for specific topics.

Sentence explorer feature

The participants were then asked on their thoughts of a sentence explorer feature being added to the BDCE-2. This feature, which already existed in DatAR as a separate widget, would be able to show sentences from publications where a disease and a different brain topic co-occur. The participants unanimously thought a sentence explorer would be a great addition to the system, with the most common argument being that this would be essential in understanding what type of relation a co-occurrence indicates. Co-occurrences in literature are always multifactorial, so having the ability to read the sentences where they occur is essential in understanding the relation between the topic and a disease. Some participants raised concerns about the selection criteria for showing the sentences however, how do you determine which sentences to show? One participant added that they would like the ability to see the number of publications that have topic co-occurrences with both diseases.

Using BDCE-2 for their own literature research

Finally, the participants were asked if they would use an AR system such as DatAR and BDCE-2 for their own literature research. They unanimously stated that they would when it came to orienting on what research to perform.

Multiple participants stated that it would make the (boring) process of literature research a lot more fun, and as visual thinkers would be more productive than reading through multiple bodies of literature to manually find connections. Other participants stated that traditional search engines do not easily allow them to find co-occurrences between topics in literature and this tool would automate this process to some degree. One participant stated that the tool is similar to entering a search query with multiple topics in PubMed, but the AR application worked better and could be integrated with their real computer for easy access to the publications.

8 Discussion

8.1 Discussing the user evaluation results

The results from the user evaluations show that BDCE-2 is beneficial for exploring similarities between brain diseases. All participants were positive about the design direction and could see themselves using the tool for their own literature research, with the primary criticisms revolving around lack of features. While this degree of positivity could be biased due to the participants trying out a new technology, as noted by Andersen et al. [14], the detailed answers, constructive feedback and interest in BDCE-2 that each participant showed suggest that regardless of trying out a new technology the overall opinion of BDCE-2 was positive.

Usability wise, five of the participants had trouble recognizing the meaning of the colors. Future changes to the design should address this by making the connection between the colors and the brain classes more apparent. Support for color blindness is also paramount, by for example providing color schemes suitable for those with color blindness. Since multiple participants wrongly identified what the numbers in BDCE-2 meant, this could be an indication that future design should make it clearer that the numbers represent co-occurrences in literature. The question is however if it is the task of the visualization to explain this, since we designed BDCE-2 for expert users that have expertise in the applicable research field and received information about BDCE-2 before viewing it in AR. We argue that it might not be necessary for BDCE-2 to indicate the meaning of the numbers, if pre-usage instructions and documentation are available and clear to the user.

The majority of participants saw a connection between the disease positions in the topic model and the data shown by the BDCE-2 visualization. This indicates that the BDCE-2 is successful at explaining the topic model, however more could be done to make this link more apparent since three participants could not recognize a connection. There was also a clear indication that the Total and Ratio view were essential in understanding the similarities.

When it came to the limit of five displayed topics, mixed responses indicated that some users desired the ability to see more while others would prefer to only see the most relevant topics. To conform to both sides, an optional setting could be added to expand on the total number of viewed topics if the user so desires. Overall, the participants felt encouraged to research brain disease similarities and desired more features to better narrow down their searches.

While all participants were positive about BDCE-2 in AR, the design could work just as well in a desktop setup due to its 2.5D nature. The participants did not comment much on the specific uses of AR for BDCE-2, aside from how it made literature exploration more fun and enticing. Bach et al. identify that 3D visualization exploration was about equally as precise and efficient on a desktop setup compared to using immersive AR [15]. They also note that immersive AR was better for tasks which required precise actions in 3D. The current BDCE-2 design does not require precise actions to access its features, which begs the question if AR is a meaningful addition to the BDCE-2 design.

8.2 Future work

When it comes to future work, we propose multiple avenues for further research.

Additional BDCE-2 features

The participants of the user evaluation showed a clear desire for more features in BDCE-2. Some of these features include:

- The ability to sort topics in an ascending and descending manner, with various criteria (total number of co-occurrences with one disease, co-occurrence ratios, total number of publications etc.)
- A sentence explorer, allowing users to read some of the sentences where co-occurrences between brain diseases and topics occur.
- The ability to see an annotated history of your search queries in BDCE-2. This could be encoded in the depth dimension as suggested by the interviewees during the expert interviews.
- An optional setting for displaying more than five topics at once in the filtered views. This can be accomplished using the fish-eye design style proposed by one of the interviewees during the expert interview (see Section 5.3).
- An additional view, that shows the number of publications in which co-occurrences between brain diseases and topics occur. This could be expanded by also showing the number of publications that con-

tain co-occurrences between a brain topic and **both** diseases.

We propose a deeper look at the inclusion of these features, and design studies on further iterations of the BDCE design.

Novel BDCE designs in 3D & use of AR

As proposed by one of the interviewees during the expert interviews (see Section 5.2), co-occurrences could be visualized in a multitude of different ways in 3D space. They gave two examples. The first example had individual brain diseases connected to individual topics, with this connection somehow indicating the strength of the co-occurrence. The second example was a 3D spider chart with a brain disease in the center and connections to co-occurring brain topics spread around the center, with again some way of indicating the strength of the co-occurrence through the connection. These are but some examples of different ways that co-occurrences between brain diseases and other topics can be visualized, which might warrant further exploration. Additionally, the benefits of AR for the BDCE-2 design are questioned in the previous section. Further research into the benefits of AR when it comes to 2.5D visualization designs might be worthwhile.

Improved BDCE-2 visualization design

The manner in which numbers are currently visualized in the BDCE-2 design could perhaps be improved. Nacenta et al. introduce FatFonts, a numeric typeface design which bridges the gap between numeric and visual representations [16]. With this font, numbers are represented as single glyphs where the amount of ink in the glyph is proportional to its quantitative value. This allows viewers of a visualization to immediately get a sense of the quantitative value of a number, without having to process the number itself. Displaying the numbers in BDCE-2 in this manner can add another way for users to see the co-occurrence strength at a glance. Currently the histogram serves this purpose, but it might be beneficial to have multiple indicators of co-occurrence strength. Further ways of encoding co-occurrence strength can also be explored, such as sphere size, shape etc.

9 Conclusion

Recognizing connections between topics discussed in literature gets more difficult as the body of knowledge increases with each publication. The topic model in DatAR aims to assist neuroscientists with their task of finding unexplored connections between brain diseases by seeing at a glance which brain diseases might be semantically similar. However, this topic model lacks the ability to

explain **why** and **how** these brain diseases are semantically similar which is essential for neuroscientists to understand the type of connections that brain disease similarity indicates. We designed BDCE-2, a 3D visualization which allows neuroscientists to investigate these brain disease similarities in more detail, by showing the number of co-occurrences between two brain diseases and any brain topic discussed in literature. We evaluated our first design iteration, BDCE-1, through expert interviews which gave us insight in our design approach and encouraged us to move away from a hierarchical design. Instead we would focus on a single overview with tools to manipulate the topics displayed in the overview. The second and final design iteration, BDCE-2, was evaluated through user evaluations. These clearly indicated that BDCE-2 has potential to successfully explain the topic model and encourage further investigation into potentially interesting relations between similar diseases, answering our research question. Finally, we propose avenues for future work, including additional features for BDCE-2, an exploration into additional 3D visualization designs for visualizing the co-occurrences between brain diseases and brain topics, and a study on the potential for improving the number visualization style in BDCE-2.

Conceptual search is a powerful tool, and as the academic world slowly moves away from traditional text-based search we believe that our design will make a great contribution by making literature research more efficient, intuitive and enjoyable.

Acknowledgements

Special thanks to Marie Denoulet & Géry Laurent for the constructive discussions and help with the recruitment of evaluation participants. Also thanks to Lynda for the high levels of patience, tolerance and flexibility that she has shown me during this project. Her constant uplifting attitude, especially when I did not feel capable and motivated, was essential for the conclusion of this project!

References

- [1] Ivar Troost. “Supporting Relation-Finding in Neuroscientific Text Collections using Augmented Reality: A Design Exploration”. MA thesis. Utrecht University, 2020. URL: <https://dspace.library.uu.nl/handle/1874/397219>.
- [2] Freek Buitter. “Designing a user interface for the discovery of relations between semantically similar brain diseases”. 2021.
- [3] Ghazaleh Tanhaei, Lynda Hardman, and Wolfgang Huerst. “DatAR: Your Brain, Your Data, On Your Desk - A Research Proposal”. In: *2019 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR)*. IEEE, Dec. 2019. DOI: 10.1109/aivr46125.2019.00029. URL: <https://doi.org/10.1109/aivr46125.2019.00029>.
- [4] B Nolan Nichols et al. “Linked Data in Neuroscience: Applications, Benefits, and Challenges”. In: (May 2016). DOI: 10.1101/053934. URL: <https://doi.org/10.1101/053934>.
- [5] J.J. Thomas and K.A. Cook. “A visual analytics agenda”. In: *IEEE Computer Graphics and Applications* 26.1 (Jan. 2006), pp. 10–13. DOI: 10.1109/mcg.2006.5. URL: <https://doi.org/10.1109/mcg.2006.5>.
- [6] Tom Chandler et al. “Immersive Analytics”. In: *2015 Big Data Visual Analytics (BDVA)*. IEEE, Sept. 2015. DOI: 10.1109/bdva.2015.7314296. URL: <https://doi.org/10.1109/bdva.2015.7314296>.
- [7] Colin Ware. “Designing with a 2½D attitude”. In: *Theme: Information landscapes* 10.3 (Dec. 2001), pp. 258–265. DOI: 10.1075/idj.10.3.07war. URL: <https://doi.org/10.1075/idj.10.3.07war>.
- [8] Kim Marriott et al. “Immersive Analytics: Time to Reconsider the Value of 3D for Information Visualisation”. In: *Immersive Analytics*. Springer International Publishing, 2018, pp. 25–55. DOI: 10.1007/978-3-030-01388-2_2. URL: https://doi.org/10.1007/978-3-030-01388-2_2.
- [9] Maxime Cordeil et al. “ImAxes: Immersive Axes as Embodied Affordances for Interactive Multivariate Data Visualisation”. In: *Proceedings of the 30th Annual ACM Symposium on User Interface Software and Technology*. ACM, Oct. 2017. DOI: 10.1145/3126594.3126613. URL: <https://doi.org/10.1145/3126594.3126613>.
- [10] Ciro Donalek et al. “Immersive and collaborative data visualization using virtual reality platforms”. In: *2014 IEEE International Conference on Big Data (Big Data)*. IEEE, Oct. 2014. DOI: 10.1109/bigdata.2014.7004282. URL: <https://doi.org/10.1109/bigdata.2014.7004282>.
- [11] Alexander Bogner, Beate Littig, and Wolfgang Menz. “Introduction: Expert Interviews — An Introduction to a New Methodological Debate”. In: *Interviewing Experts*. Palgrave Macmillan UK, 2009, pp. 1–13. DOI: 10.1057/9780230244276_1. URL: https://doi.org/10.1057/9780230244276_1.
- [12] Marc Satkowski, Wolfgang Büschel, and Raimund Dachsel. *Experiences with User Studies in Augmented Reality*. 2021. eprint: arXiv:2104.03795.
- [13] Praneeth Chakravarthula et al. “FocusAR: Autofocus Augmented Reality Eyeglasses for both Real World and Virtual Imagery”. In: *IEEE Transactions on Visualization and Computer Graphics* 24.11 (Nov. 2018), pp. 2906–2916. DOI: 10.1109/tvcg.2018.2868532. URL: <https://doi.org/10.1109/tvcg.2018.2868532>.
- [14] Benjamin J. H. Andersen et al. “Immersion or Diversion: Does Virtual Reality Make Data Visualisation More Effective?” In: *2019 International Conference on Electronics, Information, and Communication (ICEIC)*. IEEE, Jan. 2019. DOI: 10.23919/elinfocom.2019.8706403. URL: <https://doi.org/10.23919/elinfocom.2019.8706403>.
- [15] Benjamin Bach et al. “The Hologram in My Hand: How Effective is Interactive Exploration of 3D Visualizations in Immersive Tangible Augmented Reality?” In: *IEEE Transactions on Visualization and Computer Graphics* 24.1 (Jan. 2018), pp. 457–467. DOI: 10.1109/tvcg.2017.2745941. URL: <https://doi.org/10.1109/tvcg.2017.2745941>.
- [16] Miguel Nacenta, Uta Hinrichs, and Sheelagh Carpendale. “FatFonts”. In: *Proceedings of the International Working Conference on Advanced Visual Interfaces - AVI '12*. ACM Press, 2012. DOI: 10.1145/2254556.2254636. URL: <https://doi.org/10.1145/2254556.2254636>.

Appendix A - Comparison of AR Head Mounted Displays (HMDs)

In this study, we conduct a short comparison of two HMDs for AR on the following criteria:

- The ability for virtual objects to remain static
- Standalone vs PC-tethered
- Tracking capabilities
- Ongoing support from the manufacturer
- Accessibility

The two HMDs being compared are the Meta 2 and the Hololens 2. The Meta 2 was chosen since the VR version of DatAR had existing support for it, while the Hololens 2 was chosen as it is regarded as the current best AR HMD available for consumers. These are the comparison results.

Static virtual objects

With the Hololens 2, thanks to its 240hz refresh rate display, virtual objects remain stable to the point of it seeming like there is no movement at all. Virtual objects viewed using the Meta 2 move around a lot while the user moves their head.

Standalone vs PC-tethered

The Hololens 2 is standalone, making it more suitable for demo's in conferences, classrooms etc. where access to a powerful enough PC might not be practical. The Meta 2 is tethered and requires a separate PC to function.

Tracking capabilities

Hand tracking is more refined with the Hololens 2 compared to the Meta 2. Since DatAR has frequent user interactions using hand gestures, this is an important aspect.

Ongoing support from the manufacturer

Meta, the company that developed the Meta 2, went bankrupt, which means that no further support is provided for Meta 2 HMDs. The Hololens 2 on the contrary is still supported by Microsoft, as it is a newer device.

Accessibility

The Meta 2 is cheaper and since existing code in the VR DatAR implementation supports it it would take minimum effort to get the system working with the Meta 2. On the contrary, the Hololens 2 is much more expensive and more difficult to get access to. Evaluating the DatAR system using the Hololens 2 would also require a partial rewrite of the DatAR VR version in order to support it.

Conclusion

Due to the straight up better specifications, and ongoing longer term support, we decide to continue the project on the Hololens 2. We are required to get access to one, and it does require that we rewrite parts of the DatAR codebase, but since future studies on the DatAR project will also require evaluations in AR this move benefits the whole project.

Appendix B - User evaluations tasks & questions breakdown

Here we describe the tasks that the participants of the user evaluations undertook while using the Hololens 2 to evaluate BDCE-2. We also include the questions that were asked at specific points during the experience.

1. Move the Topic Model cube center in your field of view.
2. Drag a "Disease" orb into the Topic Model socket.
3. Explore the topic model.
 - (Usability) Is the topic model clear and readable to you?
4. Find "Depression" and "Anxiety" in the topic model.
5. Drag the Depression and Anxiety spheres out of the topic model.
6. Place the Depression and Anxiety spheres in the other cube on your left (which is the BDCE-2 widget).
7. Explore the visualization
 - (Usability) Is the BDCE-2 visualization clear and readable to you?
 - (Usability) The topic spheres all have different colors. Do you have an idea what they mean?
 - (Data interpretation) By just looking at the data in the BDCE-2 visualization, does the fact that depression and anxiety are close to each other in the topic model make sense to you? Why do you think so?
8. Look at the bottom-left filter panel. Try applying some of these filters, and then try to go back to the overview.
 - (Usability) Are the filters intuitive to use?

9. Now look at the bottom-right panel.

- (Usability) What do you think this panel is for? Feel free to press the buttons and see what happens.
- (Data interpretation) The ratio view shows the percentage of the total number of co-occurrences that both diseases have with a specific topic. Do you think this way of representing co-occurrences is beneficial?

10. Look back at the BDCE-2 visualization.

- (Data interpretation) Does anything stand out to you about the data shown in this visualization? Something that would make you interested in investigating that connection in more detail?
- (Data interpretation) One core feature that we would like to implement is the ability for you to select a brain topic, and then directly view some of the sentences from literature where the co-occurrences between that topic and the diseases occur. Do you think this would be beneficial to make the model even better at explaining the similarities in the topic model?
- (Usability) Do you see yourself using a system like this in AR for your own literature research? Why?