

2021-07

Many-objective Population Visualisation with Geons

Varga, Marius Nicolae

<http://hdl.handle.net/10026.1/17259>

10.1145/3449726.3463162

Proceedings of the Genetic and Evolutionary Computation Conference Companion

ACM

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Many-objective Population Visualisation with Geons

Marius N. Varga
marius.varga@plymouth.ac.uk
School of Engineering, Computing
and Mathematics,
University of Plymouth
Plymouth, UK

Swen E. Gaudl
swen.gaudl@gmail.com
School of Engineering, Computing
and Mathematics,
University of Plymouth
Plymouth, UK

David J. Walker
david.walker@plymouth.ac.uk
School of Engineering, Computing
and Mathematics,
University of Plymouth
Plymouth, UK

ABSTRACT

This paper explores the use of geometrical ions (called *geons*) to represent solutions in the approximated Pareto front generated by multi- and many-objective optimisers. The construction of geons based objects (*GBOs*) for solutions to a 3- and 5-objective problem is outlined, and the visualisation is embedded in a tool that has been tested with expert users. The findings suggest that our approach is promising, with all users successfully engaging with the given tasks and 4 out of 6 managing to complete some of the tasks they were assigned. Results indicate that the use of geometry, rather than colour as is often used to convey properties of Pareto front approximations, is a useful way of embedding multi-objective data.

CCS CONCEPTS

• **Human-centered computing** → *Information visualization; Visual analytics; User centered design;*

KEYWORDS

geons, information visualisation, visual analytics, Pareto Front, multi-dimensional dataset

ACM Reference Format:

Marius N. Varga, Swen E. Gaudl, and David J. Walker. 2021. Many-objective Population Visualisation with Geons. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3449726.3463162>

1 INTRODUCTION

As many-objective evolutionary algorithms (MaOEA) have matured, it has become clear that identifying tools that can be used to visualise their high-dimensional solution sets and support decision making is an important endeavour. In the last decade a range of methods have been proposed that enable visualisation of solutions in terms of the whole set of objectives, and which use dimension reduction to embed the solutions into two or three objectives [19, 24].

In this work we propose a novel way of visualising many-objective populations using geometrical ions (*geons*). Biederman in his theory,

Recognition By Components (RBC) [2], argues the case for a set of theoretical primitives, called geons that can be combined to form the structural description of a generic visual object. We built 36 unique geons as 3D models, using 4 attributes: edge, symmetry, size and axis (for a more in depth explanation please see Section 2.2). A combination of two or more geons are being used in a 3-dimensional virtual environment to represent a more complex visual object. For the rest of this paper, this visual object will be referred to as Geon Based Object (GBO). Representing each dimension of a multidimensional dataset by a geon, our approach can visualise a Pareto front solution with five objectives with a single GBO. This enables the observer to assess all five objectives simultaneously.

The remainder of this paper is organised as follows: Section 2 presents an overview of many-objectives visualisation followed by an introduction to the RBC theory and existing work with geons. Next, we present our approach to mapping the data to a geon, including the impact on the application design. Section 3 gives an overview of the methodology employed for the experiment including a deep dive into the applied experimental design. Section 4, is dedicated to the results and Section 5 (Discussion) focuses on the meaning behind the results. Finally in Conclusion and Future work, we summarize our work and discuss possible directions for our approach.

2 BACKGROUND

2.1 Many-objective visualisation

This work considers solutions to *multi-* and *many-*objective optimisation problems. Such problems are evaluated according to multiple measures of fitness, the objectives, between which a trade-off exists. A solution optimising one objective incurs a poor score according to another objective and it is not possible to simultaneously optimise all of the problem objectives, and the role of an EA is to locate a good approximation to the *Pareto front*, the set of feasible solutions that offer the best trade-off between the problem objectives.

A multi-objective problem comprises $M = 2$ objectives, while for a many-objective problem $M > 4$. In addition to the increased difficulty in optimising many-objective problems, part of the reason for this distinction is that solution sets to many-objective problems cannot be visualised using standard two or three dimensional visualisation methods such as scatter plots. Various approaches have been proposed to visualise the Pareto front approximations arising from many-objective optimisers (see [19] for a comparative review). Broadly, methods are grouped according to those that visualise solutions based on the full set of objectives, and those that use dimension reduction to project solutions into two or three dimensions for visualisation with a standard tool such as a scatter

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO '21 Companion, July 10–14, 2021, Lille, France

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-8351-6/21/07...\$15.00
<https://doi.org/10.1145/3449726.3463162>

plot. Of those relying on the full set of objectives, some enable the recovery of the individual objective values by inspection, and some do not.

The use of various dimension reduction methods has been proposed for visualising Pareto front approximations. Their purpose is to project the solution into two dimensions, and methods used to accomplish this include *principal component analysis* (PCA) [13], *multidimensional scaling* [24] and self-organising maps [16].

2.2 Geon Based Objects (GBOs)

According to the Recognition By Components (RBC) theory, object recognition happens in a set of hierarchical steps. It starts with an early edge extraction stage that considers surface characteristics (luminance, texture or colour). The next stage is the identification of non-accidental properties [2] (collinearity, curvilinearity, symmetry, parallel curves and vertices termination at common point) together with parsing regions of concavity. The identified regions combined with the non-accidental properties form a structural description of the perceived object. This structural description is then matched against an object representation in the memory. Partial representation is possible at this stage if enough features are activated to create an overlap between the viewed image and the object representation in the memory.

Biederman theorised that each identified concave region can be represented by a set of simple primitives called geometrical ions (geons) [2]. Geons are typically symmetrical volumes without sharp concavities characterised by an axis and cross section typically at a right angle of the axis. The variation of non-accidental properties of 4 attributes of generalised cylinders (edge, symmetry, size and axis), generate 36 uniquely identifiable geons. The unique combination of geons and their positional relation to each other creates a GBO that can be easily recognised. Generally, the recognition process of the GBOs is fast (under 100ms exposure [2]) and invariable over viewing position. Low image quality has little impact on the recognition process with relatively strong recognition even from novel points of view or if the object is partially obscured.

The RBC theory was used by [11] to train a neural network to recognise CAD models in a database. The network uses the basic shape of geons and their relationship to each other as an input. A similar approach was used by [18] to train an unsupervised neural network in recognising a range of visual objects. The geons were used as an intermediary step in the training process. The experiments run by Biederman [2] used line drawing examples of visual objects composed of geons to test his theory. Both experiments [11], [18] used generated three-dimensional versions of the geons, but the focus was on neural network training using geons.

Irani [8], [9] had a more human focused approach in data representation using the RBC theory in a three-dimensional environment. The goal of the research was to enhance the semantic content in diagrams using perceptual syntax, a comparison between traditional UML diagrams and a geon based version of UML diagrams. A limited set of three-dimensional geons were generated and used in order to represent a set of UML diagrams. Although the luminance property of a material is enough to identify the shape [2], additional textures and colours were used in order to display the complex relationship of the UML diagrams. The results [9] were

encouraging, with both experts and novices displaying a smaller error value in the geon based version of diagrams, as opposed to traditional UML diagram approach.

3 METHODOLOGY

To understand how and why people interact with a novel tool that allows the user to visualise complex data relationships, we employed a mixed approach integrating a quantitative post-experiment tool to reflect upon the experiment with a qualitative approach during the experiment. Our Geon Visualisation tool uses a novel technique to visualize data and due to that, it is important to understand the usage process and keenly guide the development from an early barely usable stage towards a tool fit for purpose. Qualitative research [6] allows the researcher to look into specific usages and reasons why a participant decided to do a certain action, instead of looking into trends and quantifiable findings over larger groups of users which would be less useful in an early stage. A key part in this process to interpret users actions and answers to investigate the reasons behind their choices.

For our qualitative tool we are using the THINK ALOUD method [3, 22] where users are initially promoted to verbalise their thought process during a given set of tasks. The benefit of THINK ALOUD is that users are not prompted additional questions that might distract and disturb their attention from the task they performing and it is, according to Van Someren et al., often the case that participants come into a natural rhythm of verbalising their thoughts without much prompting, making the technique moderately easy to employ. To reinforce verbalisation "Why?" questions might be beneficial such as "Why did you move from this space in the dataset to the other one?" or "Why do you think, this is the best solution?". When employing THINK ALOUD we try to minimize disturbing the participant from their task while gaining some information on the user's introspection in the moment. Thus, we are aiming to reduce retrospective memory errors and justifications for certain actions.

To complement THINK ALOUD and to give participants some time to think and reflect we added a post experiment questionnaire that includes an adapted version of the Post-Study System Usability Questionnaire (PSSUQ) [12] with an additional 4 open questions to gather some qualitative feedback to conclude the session. The PSSUQ is a validated and established usability questionnaire developed by IBM to measure general usability of their developed software. It uses a Likert scale to quantify the overall usability of a given software artefact based on the users current experience. As we are developing a novel visualisation tool it is important to understand if the tool at any point in time is fit for purpose and if users can interact with the software. For an early prototype still undergoing massive changes, the PSSUQ by itself is not sufficient for identifying why the software does not perform well or what elements need be improved. To address those questions the previously discussed qualitative approach can give more detailed information and identify reasons for why participants had issues with the software or approached a task in a given way. To allow the participants to reflect on their experience after our PSSUQ four additional open questions were added employing the idea of the "START-STOP-CONTINUE" method [4] used in education.

3.1 Data mapping

In our approach, the goal was to map data containing a set of solutions for multi- and many-objective problems to a set of GBOs that would enable the viewer to visually explore and understand the individual solutions. Data was generated for two instances of the benchmark test problem DTLZ1 [5] – one for an instance comprising three objectives and one comprising five objectives. The algorithm executed for 50,000 function evaluations in both cases, and returned 100 non-dominated solutions representing the final approximated Pareto fronts. We note that while better approximation sets can be achieved with bespoke many-objective optimisers for the $M = 5$ case, demonstrating the efficacy of the proposed visualisation approach did not require highly optimal solutions to the given problems.

In his RBC theory, Biederman argues that the representation of an object would be a structural description of the components that form the object and their relationship to each other. These relationships include the relative size of the components, orientation and the locus of their attachment. The typical configuration of a geon based object (GBO) is with a main geon that acts as the body of the object and the rest of the geons are attached to the main geon or to other geons (already attached to the main geon). The main geon is the largest in size, all the geons attached to it are smaller in size (medium size) and the geons attached to the medium size geons will be even smaller. Typically the largest geon sits at the center of the GBO and the medium size geons can connect left, right, top and bottom of it. The next layer can be the smallest geons that connect at the top and the sides of the medium size geons. With this approach we could technically display up to 17 dimensions into one GBO. This can easily be extended by adding more sizes to geons. An important factor to take into consideration is the fact that humans are slow at recognising quantitative differences for size or curvature [14] [7], therefore we need to have a significant difference between relative sizes in order to ensure fast recognition.

For the datasets used in the visualisation, we used the central (largest) geon to represent the average rank [1] of all the objectives for that particular solution. Based on the number of objectives (3 or 5) we used medium geons and small geons respectively to represent each objective. The location of the geon representing the objective stays constant for the entire dataset, therefore a viewer will always know that the Objective 2 for example is on the left hand side of the main geon. For a more detailed explanation of how we mapped the data for the 5 objective solution, please see Figure 1. For the 3 objective solution the mapping is similar except is missing 2 of the dimensions.

In order to visualise each objective's value as a geon, we first normalise the data and then we split the data into 4 quantiles. Once each objective has a range assigned to it, we generate a geon for that particular range and we place it at the location corresponding to each objective's location in the GBO. For Figure 1, objective 1, 2, 3 and 5 are in the same range, therefore represented by the same shape, squashed lemon. Objective 4 has a different range and is represented by the wedge shape geon. The central geon (largest in size) has yet a different range from the previously mentioned geons and is represented by the lemon shape geon.

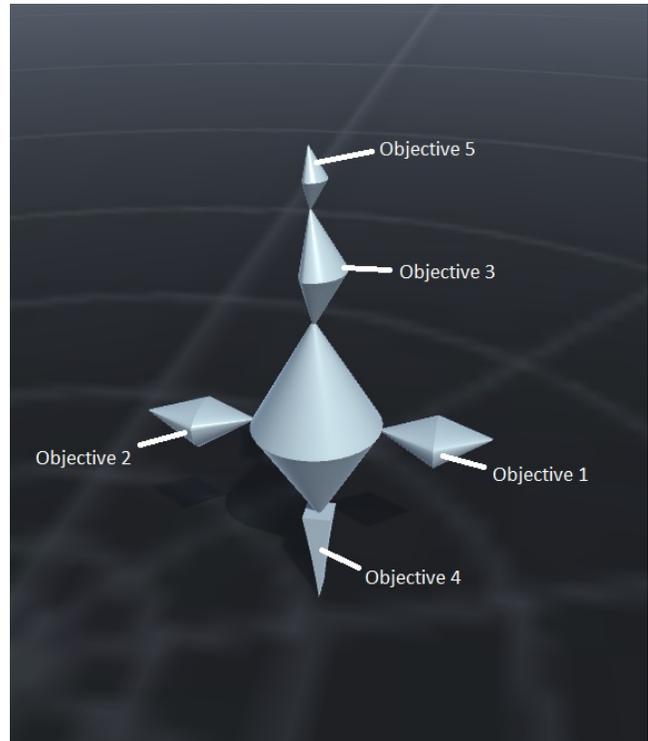


Figure 1: Geon placement in a 5 objective GBO. The main geon has the largest size and sits at the centre of the GBO. The rest of the objects that represent Objective 1, 2, 3 and 4 are all medium size and they attach right, left top and bottom respectively. The fifth objective is represented by the smallest size geon.

The data used for this study were the solutions to a many-objective problem using five objectives. Each solution was represented through a single GBO and positioned onto a 2D plane using PCA to reduce the dimensionality of the data. The Data set was centered around a Polar Coordinate system offsetting the X-Y-axis by $[-0.5, -0.5]$, see Figure 2. As all solution coordinates are in the range of $[0.0 - 1.0]$ the solutions center around the polar coordinate system's origin giving the user a visual guide. The positioning of the data has no correspondent to the values displayed, for example we cannot find the best value for objective 2 around central point but a pattern is visible in the arrangement of the GBO with clusters of similar shape being observed in the data. The distribution of solutions follows the expected pattern for the DTLZ1 problem. Pareto optimal solutions for this problem lie on the hyperplane $\sum_{m=1}^M = 0.5$. As can be seen, the solutions have been arranged in a triangular shape, with a corner corresponding to each of the objectives.

Figure 3 illustrates the corresponding plot for the Pareto front approximation of the 5-objective problem. This illustrates a similar arrangement of solutions, with a region optimising each of the five objectives being clear.



Figure 2: Geon visualisation of the three objective DTLZ1 solutions in Polar Coordinate System

3.2 Application Design

For the development of the application we used Unity3D [20] version 2020.3.0f1. The initial build was for PC but later on we switched to a WebGL [10] version deployed to a server, due to participants' difficulty to control the application through a streaming service. All the geons' meshes have been generated programmatically in Unity3D based on a set of four distinct attributes for surface definition: Edge, Symmetry, Size and Axis. These attributes are at the core of Biederman's RBC theory and geon's generation and result in 36 distinct primitives.

At the beginning of the application the user has an option to load any dataset that we added to the application, in the future we hope to have a system where the user can load their own data. The data is being processed at runtime, by normalising all the values, calculating the average rank for each solution, calculating the coordinates using the PCA dimensionality reduction algorithm and the application generates the geons shapes based on the values that they represent and finally they are assembled into GBOs based on the assigned location for each objective. Please see Figure 4.

Once the GBOs are being displayed, the user has an option to left click on a geon and display for 3 seconds information about that particular geon. The information is the attribute the geon represents, the range and the normalised value. This information is useful for geon comparison and for users to identify the particular attribute that geon represents. Please see Figure 5.

Depending on the similarities between various data elements, the coordinates generated by the PCA algorithm sometimes result in GBOs overlapping. This creates a problem, as the user struggles

to identify different GBOs, especially when they similar structural representation. In order to overcome this issue, we introduced a value that controls the distance between the GBOs. We keep the GBO's relative position to each other but the space between them can be adjusted by the value mentioned above. The GBOs can be crunched into a small area or expanded over a wider area, while maintaining the relative distance between the visual objects. The value adjustment is done through a UI slider that is part of the menu. The menu's position is on the left side of the screen and it covers approximately 20% of the screen. The menu can be toggled on and off by clicking on the right mouse button. Toggling the menu on also freezes the camera rotation, enabling the users to click on the geons and explore their values.

The flying model [25] is used as the navigation metaphor inside the virtual environment. This is typical of the first person controllers found in gaming environments but with the additional benefit that the user can explore the space freely by flying and not being limited to ground level locomotion. The forward/backward movement are mapped to W and S keys respectively and they are also mirrored onto the Up/Down arrow keys. The side movement is mapped to A and D key respectively and mirrored onto Left/Right arrow keys. Mouse movement allows the user to rotate the camera inside the virtual environment.

3.3 Experimental Design

For the first experiment only experts or participants familiar with complex datasets were selected.



Figure 3: Geon visualisation of the five objective DTLZ1 solutions.

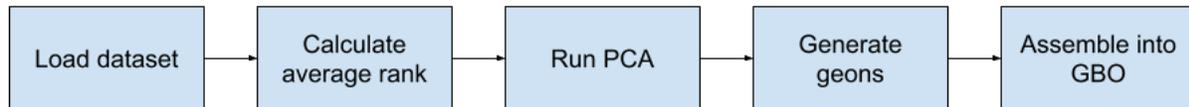


Figure 4: The steps taken at runtime to process the data and transform it into GBOs.

The experiment was set up in the following way. The participants were invited to virtual meeting call where one participant was meeting an experimenter. The entire meeting was recorded. To start the experiment a short brief was given to the participant detailing the structure. Next, an initial short background questionnaire was filled in by the participant using an online survey tool and the participant was asked to stop and respond after completing the general part about. The collected items in this part were age, education, first language as well as current job role to gain background information on the participants expertise.

Next, the experimenter read out the experiment protocol describing the next steps and explaining the software controls and what the users will see once they see the software. A brief description of what Geons are and what type of data is represented was given. The experimenter gave a quick example of how the controls work and showed a reduced dataset to the participant.

Next, the experimenter gave the participants ten minutes to complete five tasks. Stating that completion was not important but that the participants should verbalize how they approach the

task and what their thought process was. The tasks were finding good/bad solutions for individual objectives as well as globally good solutions.

During those 10 minutes the experimenter occasionally reminded the participants to verbalize. They also asked the participants why they performed certain actions or what they intended to do. In some cases the initial explanation of how the Geons are constructed was repeated as well as details on the controls.

After completing the ten minutes, even if not all tasks were completed the experimenter asked the participants to complete modified PSSUQ in the same online survey tool and stop before completing the remaining 4 open questions. After the participants reported back they thanked the experimenter thus far and did a short debrief after which they asked the participants to take some time for completing the last four open questions.

To reduce the time spent on the questionnaire and not overload the participants we reduced the granularity of the Likert scale from seven to five points removing the two elements around the Neutral axis. This reduces the complexity and expressiveness of the PSSUQ

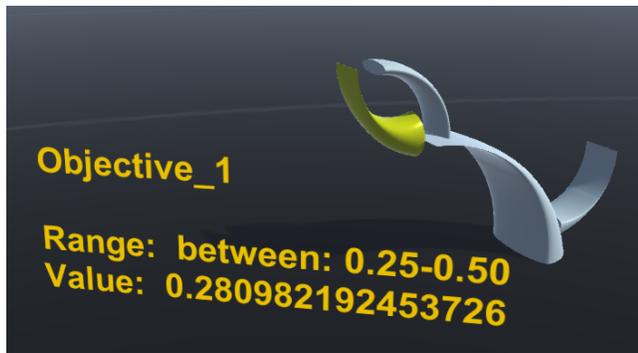


Figure 5: Floating text being displayed for 3 seconds when the user clicks on a specific geon. In this case the geon represents objective 1 and we can observe the name of the information displayed, the range of the value that the geon represents and the normalised value of that particular entry in the dataset.

which is only used for indicating the general usability of an early prototype.

4 RESULTS

A total of six participants were recruited for the pilot study and the participants were not paid. The participants were split in two groups: the first group of three users interacted with the visualisation tool through screen sharing, thus video compression and lag in action were present and it made participant's interaction with the geons very difficult. We decided to offer the second group composed of the other three participants to use the tool through a web-interface, thus having no lag or image compression, resulting in better interaction with the visual objects.

The pilot study had two main goals, first to establish the cognitive process needed for understanding, evaluating and exploring multi dimensional data and second goal was to run a usability study of our application design in order to facilitate a better interaction metaphor. The two goals are closely linked as the user interaction with the application is crucial at building the mental model needed for solving the given tasks.

4.1 Cognitive process

In order to understand the cognitive process, five tasks were given to each participant. The experimenter was revealing each task in turn upon completion of the previous task. Examples of tasks were: "The solution with the best score for objective two" or "The solution with the worst score for objective four" or "A solution that provides a good trade-off between all the objectives".

All the participants attempted at least two tasks and no participant managed to attempt all five tasks. For details on the participants' tasks attempts and successes please see Figure 6. For the 6 participants, the overall average value for task attempts is 2.83 tasks and the average value for completed tasks is 1.83 tasks.

The two participants who managed to complete successfully all the attempted tasks (participant 2 and 5), showed good understanding of the dataset mapping to geons. The initial start was slow due to

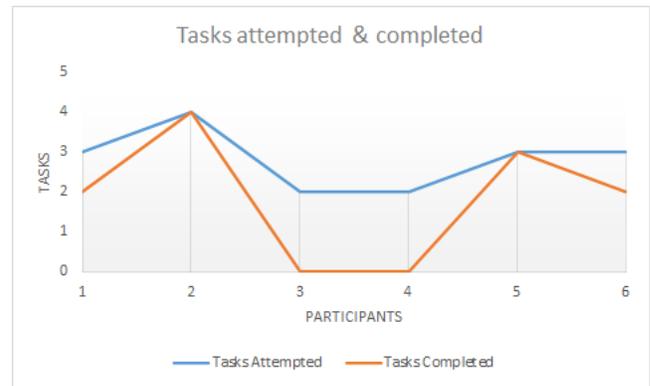


Figure 6: The chart shows the number of attempted tasks and the number of completed tasks for each participant. No participant have attempted all 5 tasks. Two participants manage to complete all their attempted tasks successfully and two participants attempted two task with no successful completion.

the novelty of the interface and the participants not being familiar with the cognitive process. Once the mental model was formed, the participants managed to complete tasks in relatively short amount of time compared to the first task and they verbalised quite effortlessly the steps needed for task completion. One of the major hurdles in completing all five tasks was the application controls, especially positioning of the camera in the virtual environment, which majority of the participants found it quite difficult.

For the two participants who did not manage to complete any of the tasks successfully, the evaluation of the task analysis shows partial understanding of the mapping process from data to geon. These participants were able to find the lowest values for a geon (which tells us that value to geon transformation was understood) but the location of the geon in the overall structure of the GBO was wrong, basically they were looking at the wrong geon in the GBO. This clearly shows participants' lack of understanding of the GBO construction process.

4.2 PSSUQ - Usability study

The PSSUQ questionnaire is widely used to measure users' perceived satisfaction with an application. It uses a set of 16 standard questions and calculates a set of scores using a Likert Scale of 7 points, for our questionnaire we used a 5 point Likert scale. The overall score is calculated as the average of questions 1 to 16. System Usefulness (SYSUSE): the average scores of questions 1 to 6. Information Quality (INFOQUAL): the average scores of questions 7 to 12. Interface Quality (INTERQUAL): the average scores of questions 13 to 15.

The scores for our visualisation can be seen in Figure 7. The smaller the value the better. An ideal value would be 1 corresponding to Strongly Agree. The neutral option is value 4 and a value of 7 will indicate a Strongly Disagree.

Interestingly, one of the questions "I believe I could become productive quickly using this system" has achieved an average score of 1.83, which is significantly better than any average score in

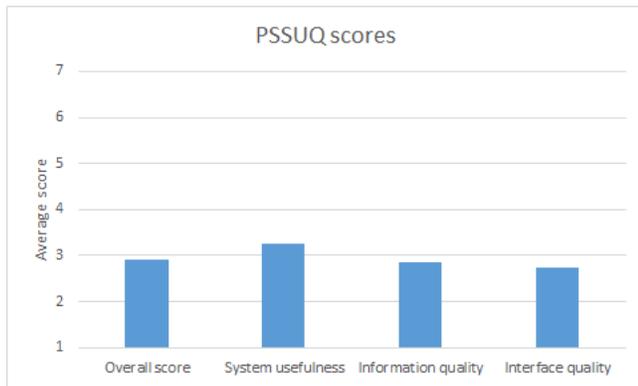


Figure 7: The average PSSUQ scores from all six participants with one exception the Information quality score. One of the participants did not answer enough questions to be able to calculate a score. The information quality score is the average of only 5 participants.

Figure 7. This score is significant as the participants seem to see the potential of the geon based visualisation. This was also confirmed by the additional qualitative feedback given in the four additional questions.

5 DISCUSSION

In a novel visualisation, simply presenting the data to a user is not enough, the understanding of the mapping process from data to geon together with the ability to interact with the tool are vital steps in building a coherent mental model for the user. This was evident from the study results, no participant attempted or completed all five tasks (see Figure 6). Some of the participants had partially formed mental model of the visualisation paradigm and others taking longer to learn the controls and simply running out of time. First part of this section is dedicated to the cognitive model and its importance in the visualisation and the second part discusses the user's engagement with application and its significance.

5.1 Cognitive model

In our overall goal of deconstructing the expert's cognitive model used for multi-dimensional data analysis, we decided to employ a bottom-up approach [21], where the focus is on the task analysis. The goal of this pilot study was to observe the experts' cognitive process during the execution of a set of simple tasks and their interaction with the tool during that process. These findings are the building blocks of the cognitive process that enables users to explore, understand and evaluate a dataset using our approach.

Our observations suggest that those participants that manage to build a coherent mental model of geon's representation had no issues performing the given tasks. The participants who lacked full understanding of the mapping process from data to geon made some assumptions based on their previous experiences with other visualisation applications.

One of the assumptions was that size is a key characteristic in the data mapping to a geon. The expectation was that two geons that look the same in shape but have different values should have

different sizes. It is a fair assumption, but our approach mapped the range value of the data to the geon's shape and the size is simply applied according to the RBC theory in order to facilitate easy shape recognition. The actual value of the data entry was only made available as "on demand" additional information simply to be used as a comparison between geons of the same shape. Our hypothesis was that participants will use their gaze to gain instant understanding of the value range and for deeper understanding they need to interact with the geon by clicking on it and displaying additional information. Therefore, a better way of explaining the mapping process is needed. Even though the experimenter had brief each individual, in depth, about the mapping system and what the geons actually represent, the novelty of the system and the sheer volume of new information that users had to absorb in short amount of time lead to incomplete mental models. In addition, the participants found the controls and navigation quite challenging further reducing the cognitive capacity to remember and apply all the new knowledge.

This was evident from feedback we received from some participants, they found it difficult to remember the range each geon represents. Even though they had the ability to click on a geon and quickly see the range value, the limited space in working memory meant that as soon as the participant was looking for a new range they completely forgot the shape of the geon that represents that shape. Most of the working memory was preoccupied with exploring the virtual space and the application controls, leaving little room for additional information to be retained. Suggestions from some of the participants was to include a key table for geons and the range that they represent in order to deal with this particular issue. Probably a more hands-on approach in introducing the mapping process, will enable the user to form a better mental model of the relationship. Enabling the participants to assign a geon to each range value before data is being displayed, we believe can help with building a better mental model of how geons work. At the moment, the geons assignment to a value range is done randomly, by software, during the loading process. This means that participants need to first gain understanding what value range each geon represents and only after that they can start the data exploration process. Enabling the users to actively choose the geons, before the data is displayed, will ensure that a model of that the geons represent is formed before the data exploration stage starts, leaving more room for the cognitive process needed to explore the represented data.

5.2 Points of view

An interesting observation in our study was the GBOs' positioning in the virtual environment enabling users to look at them from a dual point of view macroscopic and microscopic, as seen in the visualisation of a complex neuronal pattern [23]. In our approach we refer to a macroscopic point of view as a top down approach with the camera situated at the top of the displayed data, visually capturing most of the GBOs. This point of view combined with the coordinates generated by the PCA algorithm results in GBOs forming patterns or clusters. These clusters presented points of interests for the participants in their exploration. Some of the participants verbalised their visual search patterns with phrases like "central point" or "

“this group”, indicating that the macroscopic point of view is part of the cognitive process for data exploration.

The majority of the geons that form the clustered GBOs tend to have close values, visually the GBOs look similar to each other with some small variations of one or two geons. From our observations, the users tend to explore the grouped GBOs as a group, selecting geons from one GBO and moving on to another for comparison. This type of exploration can be confusing without a point of reference. In order to help the participants with the orientation and navigation at macroscopic level, we used a textured image of concentric circles (see Figure 3 or Figure 2). The central point is centered at 0 on both x and z axis, similar to the central point coordinates for the PCA algorithm. This is important in the mental mapping of the GBOs in the virtual environment and it is crucial in building a “mental history” of the places visited by the user when exploring the data.

Generally, the best view to observe an individual geon is the canonical orientation [17] [2] [26], generally a three-quarters front view, maximising the geon’s features and enabling quick recognition. We need to keep in mind that the Biederman’s study on geon recognition was done on a single GBO at the time with no other shapes in the background or other type of noise that might interfere with the recognition process. In our study, at microscopic level (close up look of the GBOs), the participants place the view instinctively on a similar canonical orientation but with the camera being slightly up, looking down on the GBOs. This orientation basically minimises the background noise which are other GBOs. Simply pointing the camera slightly down gives a user the best view of the geons’ features and minimises the interference from other objects.

The controls should be intuitive and seamless to enable the users to place the view in the best possible spot in order to observe the geon’s features. There is strong evidence of this in the participants’ feedback asking for better controls and suggesting that the camera should pivot around points of interests when the users selected a geon on the screen.

The textured pattern on the floor continues to play an important role even at microscopic level. According to [15] a strong sense of depth can be achieved from the texture floor, as the only points of reference in the virtual environment are the other GBOs. The textured floor can also help with the “mental history” of visited places through an egocentric view, placing the user amongst the data itself rather than observing it from outside.

6 CONCLUSION & FUTURE WORK

The visualisation of many-objective solutions is not a trivial task and remains a challenge overall. In this paper we presented a novel way to visualise a many-objective population using geons and implicitly GBOs. We ran a small study with the main focus on better understanding the cognitive process employed to multi-dimensional data exploration. We explored the process through a bottom-up approach of task analysis executed by experts and observing their behaviour. The results were encouraging with some participants achieving a full working mental model being able to complete the majority of the tasks successfully. A secondary goal of the study was to analyse the application’s suitability for the task at hand by

asking the participants to complete a PSSUQ usability study with encouraging results for our approach.

For future work we plan to extend our pilot study, that was mainly focused on experts, to include a higher number of participants including non-experts. We also plan to iterate through the interaction metaphor in order to allow the participants for better controls and minimise the cognitive load when exploring the virtual space. Which in turn should enable us to refine the cognitive model for exploring a set of many-objective solutions through the geon based objects. The GBOs positioning in the virtual space remains an interesting problem, with current application employing a PCA algorithm to generate a set of two dimensional coordinates. The question remains if a more suitable approach exists? Maybe the GBOs positioning should be focused more around the user, in order to enhance the egocentric perception of the visual elements.

Another possible area to explore will be a more complex set of objectives (maybe 10 objectives). The solution number can be an interesting challenge, currently in our study we used a dataset with 100 solutions. How would we display a dataset with a higher number of solutions, let’s say 1000? What’s the interaction metaphor for a more complex data set, both in objectives and in solutions? How is the cognitive process affected by the sheer volume of information that the user needs to deal with? All very interesting questions that we plan to answer in the future.

REFERENCES

- [1] P. J. Bentley and J. P. Wakefield. 1998. Finding acceptable solutions in the Pareto-optimal range using multiobjective genetic algorithms. In *Soft Computing in Engineering Design and Manufacturing*. 231–240.
- [2] I. Biederman. 1987. Recognition by components: A theory of human image understanding. *Psychological review* 94, 2 (1987), 115–147. <https://doi.org/10.1037/0033-295X.94.2.115>
- [3] E. Charters. 2003. The use of think-aloud methods in qualitative research an introduction to think-aloud methods. *Brock Education: A Journal of Educational Research and Practice* 12, 2 (2003).
- [4] C. M. Cunningham and T. L. White. 2020. What are they trying to tell me? Large-scale viability of the Start, Stop, Continue teaching evaluation method. *Innovations in Education and Teaching International* (2020), 1–10.
- [5] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. 2002. Scalable Multi-Objective Optimization Test Problems. In *Proceedings of IEEE Congress on Evolutionary Computation*, Vol. 1. 825–830.
- [6] E. Fossey, C. Harvey, F. McDermott, and L. Davidson. 2002. Understanding and evaluating qualitative research. *Australian & New Zealand Journal of Psychiatry* 36, 6 (2002), 717–732.
- [7] W.R. Garner. 1962. Uncertainty and Structure as Psychological Concepts. *Science* 140, 3568 (1962), 799.
- [8] P. Irani, M. Tingley, and C. Ware. 2001. Using perceptual syntax to enhance semantic content in diagrams. *IEEE Computer Graphics and Applications* 21, 4 (Sept. 2001), 76–84. <https://doi.org/10.1109/38.946634>
- [9] P. Irani and C. Ware. 204. The effect of a perceptual syntax on the learnability of novel concepts. In *Proceedings. Eighth International Conference on Information Visualisation, 2004. IV 2004*. IEEE, 308–314. <https://doi.org/10.1109/IV.2004.1320162>
- [10] The Khronos. 2021. *WebGL*. Retrieved April 10, 2021 from <https://www.khronos.org/webgl/>
- [11] S. Leonard, R. Lepage, and T. Redarce. 1999. Associative memory for geon-based object identification. In *IJCNN'99. International Joint Conference on Neural Networks. Proceedings (Cat. No.99CH36339)*, Vol. 5. IEEE, 3494–3499. <https://doi.org/10.1109/IJCNN.1999.836229>
- [12] J. R. Lewis. 1992. Psychometric evaluation of the post-study system usability questionnaire: The PSSUQ. In *Proceedings of the Human Factors Society Annual Meeting*, Vol. 36. Sage Publications Sage CA: Los Angeles, CA, 1259–1260.
- [13] Y. Masafumi, Y. Tomohiro, and F. Takeshi. 2010. Study on effect of MOGA with interactive island model using visualization. In *IEEE Congress on Evolutionary Computation*. IEEE, 1–6.
- [14] G.A. Miller. 1956. The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review* 63, 2 (1956), 81–97.

- [15] K. Nemire, R.H. Jacoby, and S.R. Ellis. 1994. Simulation Fidelity of a Virtual Environment Display. *Human Factors: The Journal of the Human Factors and Ergonomics Society* 36, 1 (1994), 79–93.
- [16] S. Obayashi. 2002. Pareto Solutions of Multipoint Design of Supersonic Wings using Evolutionary Algorithms. In *Adaptive Computing in Design and Manufacture V*, I.C. Parmee (Ed.). Springer-Verlag, London, 3–15.
- [17] S. Palmer, E. Rosch, and P. Chase. 1981. Canonical perspective and the perception of objects. *Attention and Performance IX Hillsdale, NJ: Erlbaum J. Long & A. Baddeley(Eds.)* (1981), 135–151.
- [18] L. Shams and C. der Malsburg. 2002. Acquisition of visual shape primitives. *Vision Research* 42, 17 (Aug. 2002), 2105–2122. [https://doi.org/10.1016/S0042-6989\(02\)00130-X](https://doi.org/10.1016/S0042-6989(02)00130-X)
- [19] T. Tušar and B. Filipič. 2014. Visualization of Pareto front approximations in evolutionary multiobjective optimization: A critical review and the projection method. *IEEE Transactions on Evolutionary Computation* 19, 2 (2014), 225–245.
- [20] unity.com. 2021. *Unity3D*. Retrieved April 7, 2021 from <https://unity.com/>
- [21] M.W. van Someren, Y.F. Barnard, and J.A.C. Sandberg. 1994. *A practical guide to modelling cognitive processes*. Academic Press, London.
- [22] M.W. Van Someren, Y.F. Barnard, and J.A.C. Sandberg. 1994. The think aloud method: a practical approach to modelling cognitive processes. *London: Academic Press* (1994).
- [23] M.N. Varga, R. Merrison-Hort, P. Watson, R. Borisyuk, and D. Livingstone. 2021. Tadpole VR: virtual reality visualization of a simulated tadpole spinal cord. *Virtual Reality* 25 (2021), 1–17.
- [24] D. J. Walker, R. M. Everson, and J. E. Fieldsend. 2013. Visualising Mutually Non-dominating Solution Sets in Many-objective Optimization. *IEEE Transactions on Evolutionary Computation* 17, 2 (April 2013).
- [25] C. Ware. 2012. *Information visualization : perception for design*. Morgan Kaufmann, New York. 358 pages.
- [26] A.T. Woods, A. Moore, and F.N. Newell. 2008. Canonical Views in Haptic Object Perception. *Perception* 37, 12 (2008), 1867–1878.