








## Show Me Your Best Side: Characteristics of User-Preferred Perspectives for 3D Graph Drawings

Lucas Joos<sup>1</sup>  Gavin J. Mooney<sup>2</sup>  Maximilian T. Fischer<sup>1</sup>  Daniel A. Keim<sup>1</sup>   
Falk Schreiber<sup>1,2</sup>  Helen C. Purchase<sup>2</sup>  Karsten Klein<sup>1</sup> 

<sup>1</sup>University of Konstanz, Germany

<sup>2</sup>Monash University, Australia

Submitted: October 2025    Accepted: April 2026    Published: April 2026

Article type: Regular paper

Communicated by:  
V. Dujmović, F. Montecchiani

**Abstract.** The visual analysis of graphs in 3D has become increasingly popular, accelerated by the rise of immersive technology, such as augmented and virtual reality. Unlike 2D drawings, 3D graph layouts are highly viewpoint-dependent, making perspective selection critical for revealing structural and relational patterns. Despite its importance, there is limited empirical evidence guiding what constitutes an effective or preferred viewpoint from the user’s perspective. In this paper, we present a systematic investigation into user-preferred viewpoints in 3D graph visualisations. We conducted a controlled study with 23 participants in a virtual reality environment, where users selected their most and least preferred viewpoints for 36 different graphs varying in size and layout. From this data, enriched by qualitative feedback, we distil common strategies underlying viewpoint choice. Using classical 2D aesthetic criteria (e.g., *Crossings*), 3D-specific measures (e.g., *Node-Node Occlusion*), and a novel measure capturing the perceivability of a graph’s principal axes (*Isometric Viewpoint Deviation*), we analyse the alignment of user preferences with graph aesthetics and explore their value landscapes across all perspectives. Our data-driven analysis indicates that *Stress*, *Crossings*, *Gabriel Ratio*, *Edge-Node Overlap*, and *Isometric Viewpoint Deviation* are key indicators of viewpoint preference. Beyond our findings, we contribute a publicly available dataset consisting of the graphs and computed aesthetic measures, supporting further research and the development of viewpoint evaluation measures for 3D graph drawing.

*Special issue on Selected papers from the Thirty-third International Symposium on Graph Drawing and Network Visualization, GD 2025*

The authors gratefully acknowledge financial support by the Federal Ministry for Economic Affairs and Climate Action (BMWK, grant No. 03EI1048D) and the Deutsche Forschungsgemeinschaft (DFG) – Project-ID 251654672 – TRR 161. A preliminary version of this paper appeared in [22].

*E-mail addresses:* [lucas.joos@uni-konstanz.de](mailto:lucas.joos@uni-konstanz.de) (Lucas Joos) [gavin.mooney@monash.edu](mailto:gavin.mooney@monash.edu) (Gavin J. Mooney) [max.fischer@uni-konstanz.de](mailto:max.fischer@uni-konstanz.de) (Maximilian T. Fischer) [keim@uni-konstanz.de](mailto:keim@uni-konstanz.de) (Daniel A. Keim) [falk.schreiber@uni-konstanz.de](mailto:falk.schreiber@uni-konstanz.de) (Falk Schreiber) [helen.purchase@monash.edu](mailto:helen.purchase@monash.edu) (Helen C. Purchase) [karsten.klein@uni-konstanz.de](mailto:karsten.klein@uni-konstanz.de) (Karsten Klein)



This work is licensed under the terms of the [CC-BY](https://creativecommons.org/licenses/by/4.0/) license.

## 1 Introduction

Beyond traditional graph drawing in two-dimensional space, the rise of immersive technology has enabled numerous applications and studies where graphs are arranged and viewed in three-dimensional space using devices that support stereoscopic 3D (S3D) vision [20]. Virtual (VR) and augmented reality (AR) devices are increasingly employed for this purpose, due to their growing availability, quality, and affordability, across fields such as biology [28], neuroscience [39], and the social sciences [11]. Although rare approaches visualise graphs as matrices in 3D space [38], the common method remains node-link diagrams [21] where node positions are typically derived from layout algorithms or reflect semantic 3D positions (such as brain regions [18]). Under certain conditions, S3D environments have been shown to enhance graph-related tasks, such as community detection [15] and path tracing [31].

However, due to the three-dimensional representation and the impact of depth and occlusion, viewpoints might either facilitate or prohibit the identification of important aspects and structures, thereby supporting or hindering the analysis process. Therefore, 3D visualisations are often viewed from several viewpoints for interactive graph analysis, which is mostly done by the analyst in an ad-hoc fashion. Investigating the impact of different viewpoints and their potential relation to layout characteristics and quality measures consequently deserves increased attention. One pathway to derive insight on this impact is to investigate the variability of established 2D quality measures for a sample of viewpoints without considering human perception at all. However, these measures might not be a good fit for the quality of perceived structures in a 3D environment. A complementary approach is to investigate the characteristics of viewpoints chosen by humans.

It is well-known that preference for a visualisation does not necessarily strongly correlate with better task performance or lower cognitive load [16]. Still, preferences that are not solely based on aesthetic appearance, but motivated by a context in graph analysis, can guide the evaluation of quality measures and facilitate the search for new ones. Even if humans would sometimes prefer less advantageous viewpoints, it is an important goal to avoid viewpoints that are reasoned to be subpar, in particular if objective guidance for improvement can be given. In interactive analysis, the choice of what to explore next can be influenced, or even misled, by the previous choice, especially if that earlier choice did not provide enough information to make reasonable decisions going forward. Hence, optimised initial viewpoints can be especially valuable. Further, identifying characteristics that lead to specific preferences allows us to gain further insight into human graph analysis behaviour. In the literature, measures intended to capture the quality of a graph drawing are described as either *aesthetic* or *readability* measures, terms whose meanings partly overlap but do not always coincide. As our study centres on measures grounded in viewers' preferences, we adopt the term *aesthetic* throughout this paper. Our investigation aims to better understand human preference for 3D viewpoint selection, investigating the following research question:

**RQ:** *What are users' preferences regarding viewpoints in 3D graph drawings, and to what extent can individual aesthetic measures (e.g., edge crossings, stress) or combinations of these characterise preferred or unfavoured perspectives?*

Thereby, we make the following contributions:

- A **VR user experiment** ( $n = 23$ ) collecting a set of user-preferred and -disregarded perspectives for a broad set of 36 3D graph drawings along with qualitative user feedback.
- A detailed analysis of the study results, consisting of a **qualitative and quantitative analysis** evaluating the expressiveness of individual aesthetic measures (21 in total, including

one introduced in this work) and **combinations** of these. We further analyse the **landscape** and local stability of the measures across all perspectives.

- We publish our **study data** along with all results and analyses as **open-access**, providing a valuable dataset for future research on this topic: [10.18419/DARUS-5108](https://doi.org/10.18419/DARUS-5108) [23].

## 2 Related Work

There is some evidence that people, when presented with a static graph drawing, tend to prefer those with lower values for quality measures such as edge crossings or stress [4]. However, in interactive settings, they do not necessarily optimise for the lowest values [29]. With the additional degree of freedom in 3D visualisations, this tendency may be even more pronounced due to both the limited applicability of classical quality measures to viewpoint evaluation and the difficulty humans face in judging layout quality for analytical purposes. Over the years, many quality measures have been proposed [3, 8], some of which are better supported by empirical studies than others. A relevant question is to what extent viewers can actually perceive and distinguish different values of such measures, i.e., whether this information can help them identify better viewpoints. Mooney et al. investigated the perception of stress [36], asking participants to determine which of two drawings had higher stress. They found that the task was feasible above a certain threshold difference. Mooney et al. further found that users prefer graphs with less stress, which also improved task performance [34].

In 2D settings, Huang et al. demonstrated that optimising a combination of criteria can improve task performance [17]. Several methods have been developed to optimise multiple quality measures simultaneously, or to support this through interactive user involvement [45, 43, 6]. Since layouts can either facilitate or hinder task performance, they also influence cognitive load [54, 16, 51], which is another consideration for future work in our context. Closely related to our approach of exploring quality measures through user preferences are those that invert the classical ‘algorithm first, evaluation second’ paradigm, deriving insight from human-generated layouts instead [24, 41].

In 3D graph visualisation, Joos et al. computed aesthetic measures for a large number of sampled viewpoints and visualised these in a VR environment, helping users identify potentially beneficial perspectives, which received positive feedback from domain experts [19]. Similarly, Wageningen et al. investigated how 2D quality measures vary across viewpoints sampled on a sphere surrounding 3D graph layouts, and applied gradient descent methods to find views with the highest aesthetic values [49, 50]. The investigation of the impact of such measures in stereoscopic 3D is more complex than in 2D, as factors such as depth, distance, and perspective might influence the perception significantly, as e.g., for edge crossings [56]. Drogemuller et al. [10] studied edge curvature aesthetics in Lombardi-inspired 3D layouts within stereoscopic 3D, finding that straight edges outperformed curved ones. Bennett et al. and Ware et al. have also examined the perceptual foundations of graph visualisation, including links to Gestalt principles [3, 51]. Zhang et al. [55] investigated gaze behaviour of users when solving graph analysis tasks in S3D, finding that gaze behaviour analysis can be used to measure visual complexity and cognitive load when analysing graphs in 3D.

Although considerable research has addressed drawing principles for 2D graph visualisations, and some initial work has explored advantageous viewpoints in 3D drawings, there remains no empirical evaluation of user-preferred perspectives in 3D and whether these preferences can be characterised by existing aesthetic measures or require new ones.

### 3 Methodology

To examine how viewers choose preferred viewpoints in 3D graph drawings, and to what extent these can be explained by aesthetic measures (see Section 1), we conducted an empirical user study and analysed the data through a structured pipeline (see Section 4). Participants were asked to rotate a series of 3D graphs, and for each, to identify best and worst viewpoints. In the following, we describe the dataset, the VR prototype, the study procedure, the aesthetic measures, and participant details.

#### 3.1 Stimuli

Drawing on results from earlier VR studies, we aimed to keep headset time below 45 minutes [42, 27]. With an average of one minute per graph plus setup and calibration, we decided on a sample of 36 graphs. To support general and class-specific conclusions, we included graphs from four size classes (S-S (20 nodes), S-M (50), S-L (100), and S-XL (200)) and three layout types (energy-based (L-E), layered (L-L), and semantic (L-S) types), aligned with those typically used in immersive settings [20]. With three graphs per size-layout combination, this yields 36 graph drawings. Edge density varied between 0.75% and 33.3%, computed as  $\frac{2m}{n(n-1)}$  for a graph  $G = (V, E)$ ,  $|V| = n$ ,  $|E| = m$ . Figure 3 shows a sample of these graphs, while a full documentation can be found in Appendix A. Following, we describe the generation process for each class.

**Energy-Based Layouts (L-E)** Energy-based algorithms are widely used for immersive graph visualisations [20, 15, 26], particularly when node positions lack semantic meaning, as many drawing algorithms natively support 3D. For this layout type, we selected 12 graphs (three per size) and computed 3D layouts using the *Stress Minimisation* algorithm from OGDF [5]. Graphs were drawn from the *Rome* [9] dataset and *SuiteSparse Matrix Collection* [7], both standard benchmarks. To include sparser, structured samples, we created graphs with the OGDF Random Tree Generator.

**Layered Layouts (L-L)** Layered graphs are relevant in cases with hierarchical structures or distinct node types [30]. We followed a method from a recent VR study on layered graphs for generation [14]. We used two layers for S-S and S-M and three for S-L and S-XL. Given a target size  $n$ , nodes were partitioned into  $l$  layers with a maximum deviation of  $\pm 10\%$  from  $\frac{n}{l}$ , ensuring a roughly even split with some perturbation. Within each layer, edges were added using OGDF’s *Simple Connected Graph Generator*, followed by the *FMMM* layout algorithm for 2D node placement. Between adjacent layers, 5 to  $0.7n$  inter-layer edges were added at random to produce varying densities.

**Semantic Layouts (L-S)** A key advantage of 3D visualisations is the ability to place nodes at meaningful spatial, semantic locations. A prominent case is brain activity networks, where nodes represent brain regions (centroids) and edges indicate functional correlations [18, 39]. These networks are undirected, fully connected, and weighted. Applying a threshold removes weaker edges, and discarding the weights yields an unweighted graph that highlights strong correlations. Region granularity is adjustable via brain atlases, offering flexibility in node and edge counts. Given their frequent use in immersive contexts, meaningful spatial structure, and adaptability, we included such networks in the evaluation. For S-L and S-XL, we used human fMRI data from the Human Connectome Project [47, 48] with the Schaefer (100 regions) and Brainnetome (246 downsampled to 200) atlases [44, 12]. For S-S and S-M, we used rat fMRI data [1, 2] (53 regions,




Figure 1: The VR study application, perceived and controlled using the Apple Vision Pro. The 3D graph drawing (A) can be rotated by a gaze and pinch interaction, or sliders (B). Users can store up to three preferred or disregarded perspectives (C).

downsampled to 20 or 50). Based on prior experience, edge thresholds were tuned to yield realistic connectivity for fMRI analysis (see Table 2).

### 3.2 Application

We developed a VR application for the Apple Vision Pro (2024 release) using Unity. The headset offered high resolution, substantial processing capabilities, and native hand- and gaze-based interaction, making it ideal for this evaluation. The virtual environment consisted of a minimal room designed to provide spatial orientation while minimising distractions. Participants experienced the environment while seated at a fixed position (only head movement was allowed). Graphs with pre-computed 3D layouts were presented in stereoscopic space, using blue spheres for nodes and black tubes for edges – a typical representation for graphs in immersive settings [20] (see Figure 1). Users rotated the graph by directing their gaze at it and performing a pinch gesture (thumb and index finger together), applying the hand rotation to the graph. Three sliders provided fine-grained rotation control around the X, Y, and Z axes. An interactive panel allowed participants to store up to three viewpoints per category (best or worst), review them, and select one final view for each category using toggles. Submissions were only enabled after both selections were made. Upon submission, the selected views and associated metadata (e.g., task duration and graph position) were logged in a database. Both user interface panels were movable within the 3D environment to avoid occlusion and support ergonomic interaction. While graphs remained fixed during each task, participants could adjust the default placement of the visualisation between tasks by moving a bounding sphere. Graph size was fixed, and full visibility was ensured at all times.

### 3.3 Study Procedure

The study took place in a controlled university lab and lasted about one hour. Participants received 15 € compensation. After providing informed consent, participants watched a four-minute video

explaining the setup, application, and task. They were instructed to rotate each graph and identify both the *best* and *worst* viewing perspectives from their point of view. Participants were informed that perspectives were not only to be evaluated based on visual appeal, but also on their ability to convey information relevant to task-solving. After addressing any questions, the headset was fitted and calibrated. A screen-mirroring setup allowed the instructor to monitor the session and offer guidance if needed. Participants completed two trial tasks to familiarise themselves with the interface and determine a default graph position. In the main study, they evaluated 36 graphs shown in random order. Submissions were only possible after rotating the graph at least 180° along each axis ensuring a thorough exploration. A soft time limit of one minute per task was set, which was not enforced, but participants were reminded if they exceeded it significantly. Breaks were allowed at any time, and questions were addressed as they arose. Audio was recorded, and participants were encouraged to verbalise their strategies. At the end, participants completed a brief questionnaire on demographics, strategies, and their experience before leaving.

### 3.4 Aesthetic Measures

Our analysis uses 21 aesthetic measures, mostly established 2D graph drawing metrics grounded in longstanding principles and empirical studies [35, 33, 3, 40]. Moreover, we include 3D-specific overlap measures and a novel metric for structural perceptibility in 3D.

**2D Measures** To obtain 2D positions for aesthetic evaluation, we applied perspective projection to the 3D graph from a given viewpoint. We used the Python library *GdMetriX* [37], which implements all relevant 2D measures found in the literature and required for our analysis. The following measures were calculated: *edge crossings* CR, *stress* ST [53] (discrepancy of geometric and topological distances), *crossing angular resolution* CAR (how much deviate angles of crossing edges from an optimum), *bounding box area* AR, *aspect ratio* ASP, *node concentration* CON [46] (estimating how (un)evenly distributed nodes are), *node orthogonality* NO (how ‘packed’ are nodes on the grid), *Gabriel ratio* GR [35] (measuring whether nodes are too close to edges), *angular resolution* ANGR [40] (the angle deviation of incident edges from an optimum), *edge orthogonality* EO [40] (measuring how horizontal/vertical edges are), and *edge length deviation* ELD [35] (the mean deviation of edge lengths, aiming for uniformity). To capture symmetry [32], we included three edge-based measures [25]: *reflective* ESR, *rotational* ESO, and *translational* EST symmetry, measuring whether a symmetry axis exists and quantifying rotational or translational invariance.

**3D Overlap** Unlike in 2D, where edge crossings are the main form of occlusion, 3D graph drawings are more susceptible to various overlaps. Prior work highlights their relevance for users [19], yet most studies only consider node-node or edge-node overlaps, without accounting for the order or visual depth. User feedback (see Section 4.1) suggests that the direction of occlusion matters. When a node appears in front of an edge, it can falsely imply a connection where none exists, making this type of overlap particularly misleading. Conversely, when an edge passes in front of a node, the node often remains partially visible because edges are typically thinner. However, this can still obscure important parts of the node and lead to misinterpretation, especially in dense regions. Users also noted that the degree of overlap influences perception, not only their pure existence. To reflect this, we include six overlap measures: *node-node* NNO, *edge-node* ENO, and *node-edge* NEO count occurrences, while *node-node overlap area* NNOA, *edge-node overlap area* ENOA, and *node-edge overlap area* NEOA quantify extent. Node-node areas (circle-circle intersections) can be computed in closed form [52]. Overlaps involving rectangles (edges) and circles

(nodes) are more complicated and approximated via fine-grained rasterisation. For two nodes (i.e., circles) with radii  $r_1, r_2$  and distance  $d$ , the overlapping area can be calculated as

$$A = r_1^2 \cos^{-1} \left( \frac{d^2 + r_1^2 - r_2^2}{2dr_1} \right) + r_2^2 \cos^{-1} \left( \frac{d^2 + r_2^2 - r_1^2}{2dr_2} \right) - \frac{1}{2} \sqrt{(-d + r_1 + r_2)(d + r_1 - r_2)(d - r_1 + r_2)(d + r_1 + r_2)}$$

while for complete overlap, it reduces to the area of the smaller circle  $A = \pi \cdot \min(r_1, r_2)^2$  [52].

**Isometric Viewpoint Deviation** While previous measures capture various aspects of viewpoint quality, they do not assess whether a graph’s global structure is perceivable from a given perspective. This aspect, although difficult to quantify, is crucial for understanding user preferences (see Section 4.1). It is not always clear what the overall structure of a graph is, or how to determine whether a viewpoint supports its perception, especially for arbitrary graphs with no prior semantic knowledge. However, user feedback in our study provided valuable insight into how participants perceive structure. Many described preferring a ‘diagonal view’ that offers a balanced look along the graph’s main axes, avoiding alignment between those axes and the view direction. For most graphs, the perceived axes correspond to directions with high node spread, which can be estimated using principal component analysis (PCA). Inspired by Farish’s concept of *isometric projections* [13], established in the 19th century, which are views in which the axes appear equally foreshortened (120° angle between each pair), we define a measure to quantify deviation from this ideal, taking into account whether PCA-derived axes meaningfully describe the graph’s structure (*anisotropy*). Let the normalised view vector be  $\mathbf{v} \in \mathbb{R}^3$ , and the PCA of the node positions yield orthonormal eigenvectors  $\mathbf{E} = [\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3] \in \mathbb{R}^{3 \times 3}$  and corresponding eigenvalues  $\boldsymbol{\lambda} = [\lambda_1, \lambda_2, \lambda_3]^\top$ . We define the *Isometric Viewpoint Deviation* score as

$$\text{ISO} = 1 - \alpha \cdot \frac{\sigma_w}{\sigma_{\max}}$$

- $\mathbf{a} = \left| \mathbf{E}^\top \mathbf{v} \right|$ : absolute projections on principal axes,
- $\mathbf{a}_{\text{norm}} = \mathbf{a} / \sum a_i$ : normalised projection weights,
- $w = \lambda / \sum \lambda_i$ : weights as normalised eigenvalues,
- $\mu = \sum w_i a_{\text{norm},i}$ : weighted mean,
- $\sigma_w = \sqrt{\sum w_i (a_{\text{norm},i} - \mu)^2}$ : weighted standard deviation of projections,
- $\alpha = \sigma(\mathbf{w}) / \sigma_{\max}$ : normalised anisotropy factor,
- $\sigma_{\max} = 1 / \sqrt{3}$ : normalisation constant.

The anisotropy factor  $\alpha$  ensures the score only penalises unbalanced projections when the graph has meaningful, deviating main axes (i.e., uneven eigenvalues). It is zero for isotropic graphs and one for maximally anisotropic ones. This is combined with  $\sigma_w$ , which captures how unevenly the principal axes are represented in the current view, normalised by the worst-case imbalance. The resulting score ISO ranges from 0 to 1, where 0 indicates highly imbalanced views of strongly structured graphs, and 1 corresponds to either balanced views or graphs lacking dominant directions

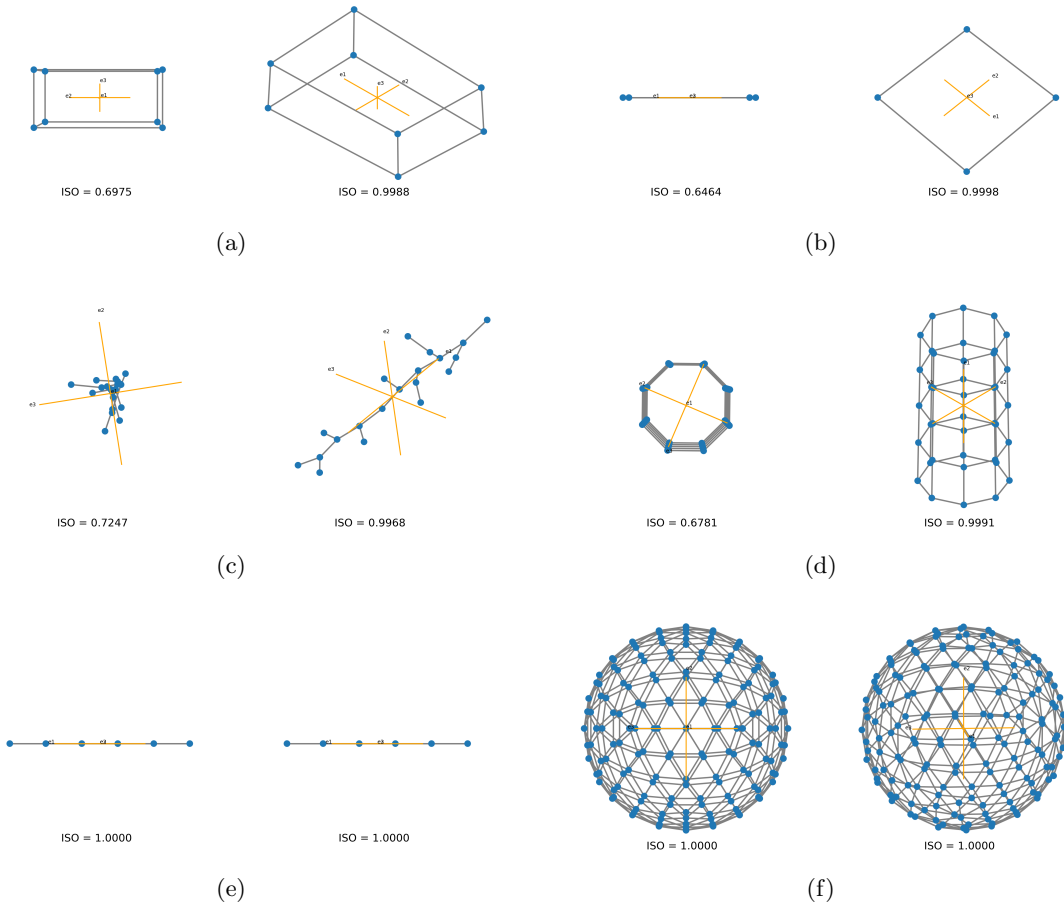


Figure 2: Examples of the ISO measure for six graphs and perspectives with lowest (left) and highest (right) ISO scores. (a) and (b) depict structured cuboid and plane like graphs. Low ISO views hide the dominant eigenvector along the view direction, while high values provide a balanced view of the eigenvectors (orange). This also holds for (c) and (d), where eigenvectors align with node distributions. The tree (c) and the cylindrical structures (d) are hidden at low values but visible in a balanced view. For (e) with only one meaningful eigenvector and (f) with uninformative eigenvectors, the ISO score is not meaningful, highlighting the need for aesthetics combinations.

(see Figure 2). While we acknowledge that there may be other, potentially better or more sophisticated measures to quantify the perception of main axes or to define *structure*, both the tests and user-reported strategies suggest that this approach is worth incorporating into the analysis and can serve as a baseline for future methods.

### 3.5 Participants

We recruited 23 participants (9 female, 14 male) from among students and associates at our university to take part in the study. The average age of participants was 25.96 years (SD =

3.94), with a median age of 25. All participants reported normal or corrected-to-normal vision. Corrected vision was achieved through contact lenses, as the VR device can not be used with glasses. Participants rated their experience with graphs on a five-point Likert scale ranging from 1 (No Experience) to 5 (Expert Level Experience), resulting in a mean rating of 2.52. Their experience with immersive technologies (such as augmented or virtual reality) was assessed using the same scale, with a mean rating of 2.27. After completing the study, participants rated the task difficulty on a five-point Likert scale from 1 (Very Easy) to 5 (Very Difficult), yielding a mean difficulty rating of 2.17.

## 4 Results

In the following, we present the qualitative feedback provided by participants, examine how their reported strategies align with their selected viewpoints, and analyse how these preferences vary. We also evaluate the aesthetic measures applied in the study and assess the extent to which individual or combined measures explain users’ choices. Lastly, we investigate the landscape of aesthetic measures and their combinations, offering insights into their correspondence with user selections and the robustness of these measures to minor perspective variations. Numbers in parentheses indicate how participants mentioned a strategy or observation.

### 4.1 Qualitative Results

During and after the study, participants described their strategies for identifying the most and least preferred viewpoints. Across all graph types, a central theme was ensuring the overall structure was visible ( $n = 17$ ), even at the cost of some overlap or edge crossings. ‘Diagonal’ views on the graph were considered beneficial for this purpose ( $n = 3$ ). Minimising node and edge overlap was frequently mentioned ( $n = 5$ ), with some participants focusing specifically on node-edge overlap ( $n = 5$ ), and one on edge-node overlap. Others stressed the importance of distinguishing connected from unconnected nodes ( $n = 1$ ), and warned against heavy node overlap ( $n = 3$ ), especially in depth ( $n = 1$ ), where structure became hard to discern otherwise ( $n = 2$ ). An observed strategy involved first finding a view revealing the global structure, then refining the view to reduce overlaps and crossings ( $n = 3$ ). Generally, global understanding was prioritised over resolving all local ambiguities. Participants identified several detrimental factors: edge crossings, compressed visualisation (especially into depth), misleading projections, and excessive flattening into 2D. Wide layouts were preferred over tall ones, as they better matched the VR field of view. Some noted that misleading visual cues, such as nodes occluding edges, suggesting false connections ( $n = 4$ ), were problematic, as were edges overlapping nodes ( $n = 2$ ). However, one participant mentioned that partially occluded long edges were mostly acceptable, as their course could still be inferred (Gestalt principle of continuity). Partial visibility of nodes or edges was sometimes sufficient for comprehension, while a larger projected area appeared to support understanding ( $n = 2$ ).

For layered graphs, a ‘tilted side view’ was frequently communicated as most effective ( $n = 9$ ), while ‘orthogonal’ views often led to overlap and were less helpful ( $n = 7$ ). Similar comments applied to semantically laid-out graphs, especially where parallel edges aligned with the view axis ( $n = 3$ ). Energy-based layouts prompted more diverse strategies. For tree-like structures, views that aligned the main branch with the view axis were discouraged ( $n = 2$ ), with diagonal views again preferred ( $n = 1$ ). In multi-component graphs (e.g., E-9, see [Figure A34](#)), overlap among components was a concern ( $n = 4$ ), and placing the more complex part in front improved clarity ( $n = 3$ ). The trade-off between revealing structure and reducing overlap was especially apparent

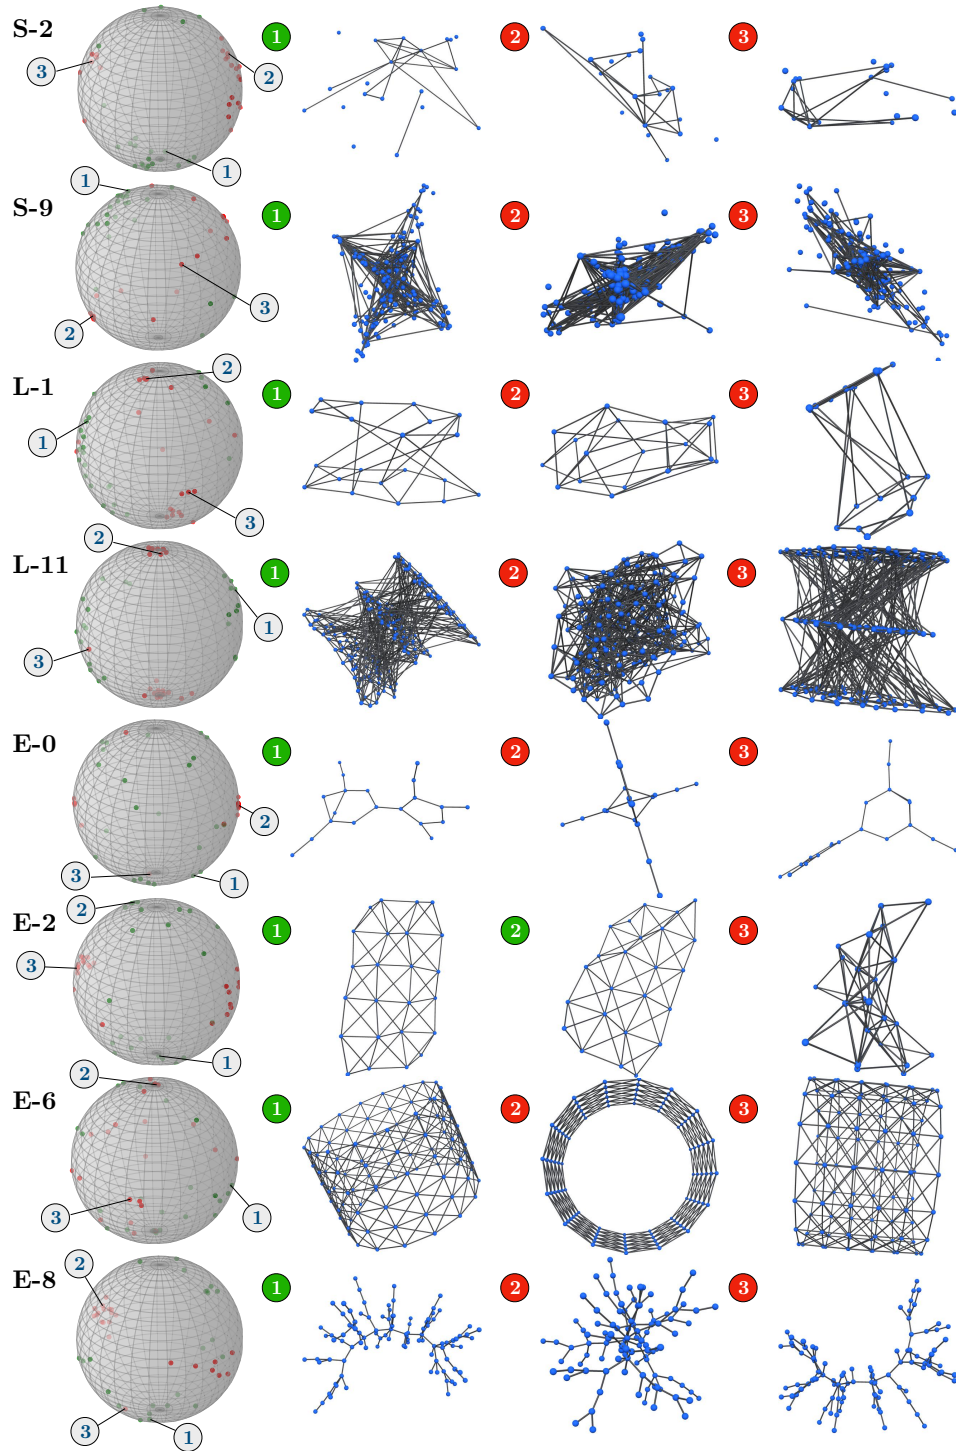


Figure 3: Example of eight graphs (of 36, reflecting the different types) showing the distribution of selected perspectives (sphere surface) along with three perspective projections (either *best* (green) or *worst* (red)), discussed in Section 4.2. The complete set of perspectives is shown in Appendix A.

for graph E-2. Participants noted that a slightly diagonal, non-planar view allowed the structure to be perceived, even at the cost of increased overlap ( $n = 3$ ). For E-6, the optimal view appeared task-dependent: diagonal views helped expose internal structure ( $n = 3$ ), while frontal or side views reduced overlap but obscured the overall form ( $n = 2$ ). Individual users sought a view that would reveal the full structure in a way that could be reconstructed by others ( $n = 1$ ), attempted to identify principal components ( $n = 1$ ), or aimed to balance visibility with occlusion ( $n = 1$ ). While selecting the best view was considered challenging in some cases, participants found the 3D environment intuitive and helpful for structural insight. Three participants compared the experience to viewing molecular models, and two noted that 3D perspectives revealed structures that would be difficult to interpret in 2D without stereoscopic cues.

## 4.2 User-Chosen Perspectives

In the following, we discuss representative user-selected perspectives and their distributions for a sample of eight graphs (see Figure 3). All selected perspectives across users and graph visualisations are shown in Appendix A.

For L-S graphs, users made relatively consistent selections. In smaller, less dense graphs (e.g., S-2), the ‘optimal’ perspectives spread out the graph, minimising occlusion. In ‘worst’ views, the graph appears compact with heavy overlap. In more complex graphs (e.g., S-9), ‘positive’ perspectives reveal structure along distinct axes. In contrast, ‘poor’ views align these axes with the viewing direction, causing occlusion, which is consistent with reported strategies. L-L graphs (e.g., L-1, L-11) show similar patterns. ‘Best’ perspectives expose the layered structure, typically via a slightly angled view, matching participant strategies. In ‘worst’ views, a front-on view of one layer hides others, leading to strong overlap. Some users instead viewed layers orthogonally from the side, causing intra-layer node overlap. This was more common in simpler graphs (e.g., L-1) and rare for more complex ones (e.g., L-11). In L-E graphs, perspective selections varied more due to differing topologies despite identical layout algorithms. Structure remained a key factor. In E-0, diverse ‘good’ perspectives still conveyed the structure well. For ‘negative’ views, participants often selected misleading angles, e.g., cross-shaped (2) or partially obscured (3) structures. In E-2, a clear 2D projection (1) was rarely chosen. Instead, users preferred a less optimal projection (2) in terms of crossings but clearer 3D form, as supported by verbal feedback. ‘Worst’ perspectives failed to show the structure and caused major overlap. In E-6, resembling a cylinder, users varied more. Many chose a top-down, angled view (1) as ‘optimal’. Others preferred (2), which had fewer crossings but poorer structural clarity. Those who did not select it as ‘worst’ chose a side-on orthogonal view with severe occlusion. In tree-like graphs (e.g., E-8), ‘good’ perspectives revealed the trunk clearly (1), while ‘bad’ views rotated it into the image plane (2). A few alternatives showed overlapping side branches (3).

Overall, the strategies communicated by the participants are clearly reflected in the results. Some graphs prompted highly consistent choices, while others showed broader variation.

## 4.3 Aesthetic Measures

Before calculating the aesthetic measures, we converted each chosen perspective’s graph position and rotation, along with the user position and rotation, into a comparable view vector representation. We then computed 2D graph projections using a perspective projection matching the VR headset’s camera parameters. In addition to node and edge positions, the projected dimensions of the representation objects (circles and rectangles) were calculated. To ensure comparability,

all measures were adapted such that 0 indicates lowest and 1 highest quality. For each selected perspective (best and worst), we computed the aesthetic measures introduced in Section 3.4. Additionally, we uniformly sampled 5,000 perspectives per graph using *Fibonacci Lattice* (see Figure 6) and calculated all aesthetic measures for these as well. This enabled us to determine the full range of achievable values and normalise the selected perspectives accordingly. The resulting scores indicate where each selected perspective falls within this range. Figure 5 presents score distributions across all measures and selected perspectives, and Table 1 summarises mean scores across all data, layout types, and sizes (see Table 2 for graph-level results). The findings suggest that several measures, i.e., CAR, ASP, CON, ANGR, EO, ELD, ESR, ESO, and EST, do not characterise user choices. In contrast, the six overlap measures, along with CR, ST, NO, GR, and ISO, appear more relevant, as they differentiate best and worst perspectives to a certain extent. A user-centric analysis of the aesthetic measure results did not reveal any outliers or substantial user-specific anomalies. A correlation analysis (Figure 4, right) shows that the overlap measures are highly correlated (particularly overlap counts with corresponding areas, which is not surprising). Other notable correlations include those between NO and AR, and among CR, ST, and GR. In Section 4.4, we further examine how well individual and combined measures explain user preferences.

#### 4.4 Combinations of Aesthetics

Based on the calculated aesthetic measures, we investigate how well they align with user judgements and whether combined measures can better reflect their importance. A combined measure approximating user choices could be valuable in supporting the automated selection of *optimal* viewpoints. To get an initial impression of the data interdependences, we performed a PCA on the study results (see Figure 4, left). Interestingly, although not designed for this purpose, the first component already provides a good separation between best and worst perspectives. Since each PCA component represents a weighted linear combination of the original features (here, aes-

Table 1: Mean aesthetic values (*best*  $\uparrow$ , *worst*  $\downarrow$ ) for all graphs (All), layout classes L-S (*semantic*), L-L (*layered*), and L-E (*energy-based*), and size classes S-S (20 nodes), S-M (50), S-L (100), and S-XL (200). In addition to 21 aesthetics, combined scores (RdYlGr colour scale) using logistic regression (C-LR) and sequential quadratic programming (C-SQP) are shown (5 aesthetics each).

Data	NNO	NNOA	ENO	ENOA	NBO	NEOA	CR	ST	CAR	AR	ASP	CON	NO	GR	ANGR	EO	ELD	ESR	ESO	EST	ISO	C-LR	C-SQP
All $\uparrow$	.89	.91	.89	.89	.85	.83	.80	.79	.44	.49	.51	.72	.74	.66	.71	.38	.59	.46	.39	.33	.57	.77	.75
All $\downarrow$	.63	.69	.55	.62	.59	.62	.45	.39	.41	.28	.49	.64	.53	.28	.64	.41	.61	.49	.53	.46	.30	.47	.39
L-S $\uparrow$	.90	.90	.89	.89	.81	.79	.76	.81	.51	.55	.52	.73	.81	.63	.71	.43	.76	.51	.38	.33	.57	.76	.74
L-S $\downarrow$	.58	.61	.51	.59	.62	.64	.55	.43	.39	.25	.43	.58	.51	.28	.62	.45	.58	.53	.49	.51	.32	.49	.42
L-L $\uparrow$	.91	.94	.90	.90	.88	.87	.77	.70	.46	.45	.60	.72	.70	.60	.77	.36	.35	.43	.39	.28	.55	.77	.71
L-L $\downarrow$	.78	.81	.68	.69	.66	.67	.44	.47	.45	.32	.48	.71	.56	.30	.73	.37	.71	.44	.48	.39	.27	.52	.44
L-E $\uparrow$	.88	.88	.87	.87	.86	.83	.86	.86	.36	.47	.43	.69	.72	.75	.65	.37	.67	.45	.40	.39	.59	.78	.80
L-E $\downarrow$	.54	.64	.47	.58	.49	.54	.36	.28	.39	.29	.56	.62	.52	.27	.56	.39	.55	.52	.64	.47	.30	.41	.33
S-S $\uparrow$	.98	.99	.96	.96	.92	.92	.84	.82	.42	.50	.51	.75	.80	.63	.78	.37	.72	.33	.28	.26	.59	.82	.77
S-S $\downarrow$	.57	.68	.53	.65	.50	.56	.51	.38	.42	.23	.46	.57	.51	.23	.49	.42	.50	.41	.44	.40	.28	.46	.39
S-M $\uparrow$	.91	.93	.89	.90	.86	.84	.78	.71	.56	.47	.55	.75	.74	.65	.76	.39	.64	.53	.41	.30	.55	.77	.72
S-M $\downarrow$	.64	.68	.56	.61	.55	.60	.52	.41	.38	.29	.46	.61	.55	.37	.65	.43	.59	.51	.51	.43	.36	.50	.45
S-L $\uparrow$	.82	.84	.85	.85	.81	.77	.76	.78	.43	.48	.56	.71	.70	.68	.64	.40	.51	.50	.42	.34	.56	.74	.73
S-L $\downarrow$	.67	.71	.58	.62	.70	.68	.41	.34	.36	.33	.56	.69	.53	.28	.70	.39	.71	.52	.55	.41	.29	.49	.38
S-XL $\uparrow$	.88	.87	.84	.83	.82	.79	.81	.85	.37	.50	.44	.66	.73	.70	.66	.38	.49	.47	.46	.42	.57	.75	.77
S-XL $\downarrow$	.65	.68	.55	.61	.59	.64	.35	.43	.49	.29	.47	.67	.52	.24	.71	.39	.65	.54	.64	.58	.25	.44	.37

Table 2: For all 36 graphs, we computed the mean values of 21 aesthetic measures and their linear combinations weighted by logistic regression (C-LR) and sequential quadratic programming (C-SQP). Values are shown for participant-favoured (↑) and -disliked (↓) perspectives. The table also lists the number of nodes  $|V|$  and edges  $|E|$  per graph and is grouped by graph type: *semantic* (S), *layered* (L), and *energy-based* (E).

Graph	$ V $	$ E $	NNO	NNOA	ENO	ENOA	NEO	NEOA	CR	ST	CAR	AR	ASP	CON	NO	CR	ANCHR	EO	ELD	ESR	ESO	EST	ISO	C-LR	C-SQP	
S-0	20	16	↑	.92	.95	.96	.96	.88	.86	.74	.77	.74	.51	.56	.78	.81	.64	.67	.32	.76	.32	.29	.09	.46	.76	.72
			↓	.67	.71	.71	.75	.62	.49	.77	.39	.46	.26	.37	.49	.52	.28	.68	.37	.54	.18	.31	.20	.34	.60	.50
S-1	20	22	↑	.96	.98	.94	.95	.84	.88	.80	.75	.43	.68	.58	.75	.91	.37	.80	.39	.82	.30	.45	.38	.64	.80	.70
			↓	.59	.69	.57	.63	.38	.58	.58	.49	.27	.24	.36	.61	.50	.16	.41	.37	.71	.38	.62	.67	.38	.49	.44
S-2	20	19	↑	.96	.98	.93	.92	.82	.81	.76	.90	.41	.49	.60	.76	.86	.56	.74	.42	.83	.32	.36	.16	.36	.72	.72
			↓	.66	.72	.56	.55	.60	.57	.72	.48	.40	.17	.24	.61	.53	.23	.61	.50	.62	.31	.36	.37	.31	.53	.46
S-3	50	100	↑	.92	.89	.93	.90	.91	.91	.83	.82	.50	.62	.67	.82	.92	.69	.78	.48	.73	.68	.28	.32	.65	.83	.79
			↓	.56	.52	.51	.52	.62	.67	.43	.40	.24	.23	.41	.57	.60	.39	.59	.55	.39	.71	.36	.50	.32	.47	.41
S-4	50	100	↑	.90	.93	.94	.93	.87	.82	.92	.79	.73	.59	.56	.64	.87	.77	.81	.46	.85	.47	.26	.18	.69	.85	.83
			↓	.43	.33	.42	.50	.40	.44	.34	.29	.36	.14	.30	.41	.43	.19	.67	.53	.57	.49	.32	.38	.33	.38	.31
S-5	50	100	↑	.86	.88	.87	.89	.80	.86	.61	.80	.68	.59	.64	.80	.89	.63	.85	.41	.92	.67	.55	.28	.70	.75	.72
			↓	.66	.68	.35	.49	.52	.67	.67	.34	.29	.25	.38	.63	.62	.28	.59	.47	.46	.34	.67	.42	.40	.46	.41
S-6	100	144	↑	.77	.81	.83	.84	.84	.72	.79	.76	.57	.45	.55	.61	.64	.53	.76	.42	.63	.62	.30	.45	.56	.75	.70
			↓	.70	.73	.57	.64	.82	.69	.56	.52	.31	.32	.67	.47	.50	.34	.65	.44	.63	.69	.41	.64	.39	.57	.48
S-7	100	137	↑	.77	.84	.87	.89	.83	.75	.77	.79	.45	.47	.47	.77	.70	.72	.48	.60	.80	.26	.47	.45	.63	.76	.76
			↓	.60	.73	.47	.60	.81	.71	.45	.31	.33	.27	.63	.65	.49	.31	.67	.53	.80	.44	.55	.51	.26	.48	.36
S-8	100	131	↑	.83	.75	.89	.89	.63	.59	.82	.75	.51	.58	.38	.75	.73	.75	.68	.38	.71	.68	.40	.30	.58	.74	.76
			↓	.69	.56	.41	.50	.89	.76	.74	.35	.31	.32	.60	.59	.49	.19	.68	.35	.65	.67	.46	.28	.31	.56	.40
S-9	200	194	↑	.95	.92	.91	.87	.73	.77	.71	.95	.52	.56	.36	.60	.79	.74	.68	.38	.76	.82	.78	.38	.51	.73	.78
			↓	.57	.58	.48	.56	.58	.70	.49	.47	.33	.23	.35	.63	.45	.31	.56	.43	.56	.87	.87	.66	.28	.47	.41
S-10	200	202	↑	.96	.93	.82	.85	.81	.75	.53	.77	.22	.53	.50	.80	.78	.42	.56	.51	.40	.76	.19	.48	.48	.67	.61
			↓	.51	.58	.56	.66	.69	.71	.56	.81	.62	.31	.41	.68	.52	.43	.73	.51	.70	.91	.35	.64	.27	.53	.55
S-11	200	150	↑	.94	.93	.82	.83	.81	.81	.89	.91	.35	.52	.34	.73	.79	.77	.73	.37	.91	.18	.27	.47	.58	.75	.81
			↓	.36	.53	.50	.63	.52	.71	.24	.25	.71	.19	.41	.60	.43	.20	.66	.38	.35	.34	.58	.87	.22	.37	.28
L-0	20	36	↑	1.00	1.00	.97	.97	.96	.96	.85	.78	.35	.42	.58	.76	.75	.54	.77	.47	.60	.32	.13	.21	.61	.83	.75
			↓	.73	.78	.65	.74	.65	.72	.56	.38	.32	.20	.61	.51	.48	.19	.56	.49	.48	.38	.27	.32	.21	.51	.40
L-1	20	38	↑	.98	.99	.94	.97	.92	.92	.85	.82	.39	.52	.57	.76	.81	.49	.75	.44	.52	.37	.26	.29	.66	.83	.75
			↓	.71	.75	.64	.73	.62	.63	.62	.61	.35	.28	.45	.58	.53	.31	.51	.44	.63	.42	.33	.48	.34	.55	.51
L-2	20	40	↑	.99	1.00	.93	.95	.90	.93	.65	.63	.49	.42	.65	.75	.68	.53	.78	.35	.43	.39	.26	.21	.49	.74	.65
			↓	.79	.84	.65	.64	.56	.62	.58	.62	.48	.36	.57	.66	.62	.43	.54	.39	.52	.44	.27	.37	.44	.56	.55
L-3	50	101	↑	.94	.97	.90	.92	.90	.89	.72	.53	.59	.43	.52	.72	.67	.56	.82	.38	.36	.43	.39	.32	.54	.76	.65
			↓	.79	.83	.78	.79	.71	.73	.71	.71	.44	.34	.39	.68	.58	.68	.74	.46	.72	.33	.38	.38	.23	.60	.64
L-4	50	102	↑	.96	.98	.90	.94	.91	.91	.69	.44	.51	.40	.67	.80	.72	.59	.74	.20	.39	.48	.51	.32	.50	.74	.62
			↓	.74	.76	.66	.71	.60	.63	.67	.54	.44	.36	.66	.71	.67	.53	.58	.24	.55	.46	.52	.46	.47	.60	.58
L-5	50	93	↑	.98	.99	.93	.93	.94	.92	.65	.54	.70	.38	.54	.82	.62	.60	.81	.48	.38	.58	.30	.24	.28	.70	.61
			↓	.71	.75	.69	.71	.70	.71	.56	.62	.47	.30	.34	.67	.56	.54	.68	.48	.64	.64	.31	.38	.35	.58	.56
L-6	100	207	↑	.85	.89	.86	.85	.85	.77	.76	.77	.11	.40	.55	.72	.60	.51	.78	.27	.26	.66	.61	.19	.49	.74	.68
			↓	.77	.79	.69	.68	.69	.63	.20	.47	.08	.25	.47	.83	.42	.19	.89	.22	.85	.64	.71	.26	.18	.46	.35
L-7	100	253	↑	.86	.92	.88	.88	.86	.86	.63	.66	.70	.44	.73	.75	.67	.65	.80	.17	.31	.14	.38	.15	.56	.72	.68
			↓	.80	.87	.68	.73	.66	.70	.28	.22	.58	.32	.48	.84	.54	.17	.89	.17	.79	.16	.49	.24	.23	.46	.31
L-8	100	284	↑	.86	.92	.93	.86	.86	.84	.86	.90	.32	.60	.67	.66	.79	.75	.68	.49	.18	.73	.39	.38	.58	.81	.81
			↓	.74	.80	.69	.66	.66	.71	.29	.33	.44	.48	.48	.70	.69	.15	.68	.48	.76	.74	.54	.54	.20	.48	.33
L-9	200	431	↑	.90	.92	.87	.84	.89	.82	.81	.78	.33	.48	.54	.74	.68	.70	.86	.27	.30	.07	.23	.28	.68	.79	.77
			↓	.88	.88	.68	.61	.79	.72	.18	.35	.50	.36	.37	.82	.56	.15	.89	.31	.87	.07	.34	.40	.13	.44	.30
L-10	200	487	↑	.81	.86	.82	.84	.81	.80	.88	.73	.55	.49	.64	.59	.73	.74	.72	.37	.21	.82	.62	.24	.55	.76	.75
			↓	.84	.85	.59	.64	.63	.62	.32	.43	.78	.27	.44	.76	.53	.24	.85	.37	.80	.89	.82	.36	.23	.46	.37
L-11	200	508	↑	.77	.83	.87	.82	.81	.82	.86	.84	.48	.42	.52	.60	.69	.57	.76	.42	.19	.10	.60	.49	.61	.77	.75
			↓	.85	.86	.75	.70	.60	.64	.31	.34	.58	.30	.55	.79	.58	.07	.90	.43	.89	.10	.72	.51	.24	.48	.34
E-0	20	22	↑	1.00	1.00	1.00	1.00	.99	.99	.95	.89	.32	.50	.38	.75	.76	.93	.87	.32	.87	.33	.15	.26	.77	.91	.92
			↓	.31	.55	.19	.50	.37	.40	.25	.14	.75	.28	.46	.43	.54	.18	.34	.42	.18	.77	.69	.49	.14	.24	.18
E-1	20	19	↑	.99	1.00	1.00	1.00	1.00	1.00	.92	.00	.42	.41	.63	.82	.65	.81	.33	.83	.24	.26	.35	.70	.90	.86	
			↓	.16	.41	.39	.73	.33	.50	.40	.26	.57	.14	.63	.73	.46	.16	.47	.35	.34	.41	.65	.24	.13	.32	.27
E-2	24	68	↑	.98	1.00	.95	.96	.97	.95	.97	.90	.62	.50	.28	.80	.84	.92	.80	.33	.81	.35	.31	.40	.64	.86	.89
			↓	.51	.68	.39	.57	.42	.52	.14	.10	.16	.17	.47	.50	.44	.14	.26	.42	.43	.37	.49	.48	.25	.31	.20
E-3	48	129	↑	.90	.93	.89	.87	.88	.87	.82	.73	.64	.44	.52	.81	.61	.77	.70	.37	.75	.69	.56	.42	.32	.73	.72

thetic measures), we were motivated to explore interpretable, linear combinations of aesthetics that match the collected data. Given these first insights, and considering both the interpretability and the typically lower risk of overfitting associated with linear models – particularly in settings with limited data – we aimed to find ‘optimal’ linear aesthetic combinations corresponding to the collected data. We pursued two different approaches for finding such combinations. The first treats the task as a binary classification problem, using *logistic regression* (LR) to distinguish best from worst views. The resulting coefficients can be interpreted as importance weights. However, this approach aims only to classify scores as above or below 0.5, not necessarily to maximise separation. For clearer interpretability, assigning best views scores close to 1 and worst close to 0, we pursued an approach based on *sequential quadratic programming* (SQP): Given aesthetic measures  $\{A_1, \dots, A_m\}$  with  $A_i : p \rightarrow v$ , where  $p \in P$  is a perspective and  $0 \leq v \leq 1$ , and equally sized classes  $P_{best}$  and  $P_{worst}$ , we optimise:  $\left(\sum_p \sum_a w_a * A_a(P_{best_p})\right) - \left(\sum_p \sum_a w_a * A_a(P_{worst_p})\right)$ . This could be posed as a linear programme with a convex constraint ( $w_a \geq 0, \sum w_a = 1$ ), but such formulations tend to create degenerate solutions with all weights assigned to a single measure. Hence, we instead constrain the weight vector to satisfy  $\|\vec{w}\|_2 = 1$  and  $w_a \geq 0$ , allowing more balanced distributions. This constraint requires a non-linear solving approach, such as SQP, which handles the constraints efficiently and yields optimal weights maximising the separation between best and worst perspectives.

The resulting importance coefficients are shown in Table 3. In addition to computing importances across the entire dataset, we also report results per layout and graph size, and under constraints limiting combinations to five or three measures, to identify the most influential ones. Across conditions, the most relevant measures include the overlap measures, CR, ST, GR, and ISO. Although we expected overlap area to be more important than count, this is not reflected in the results. Interestingly, ENO is typically more relevant than NEO, while NNO appears less expressive. For L-S graphs, overlaps, particularly NNO, along with CR, ST, GR, and AR are most

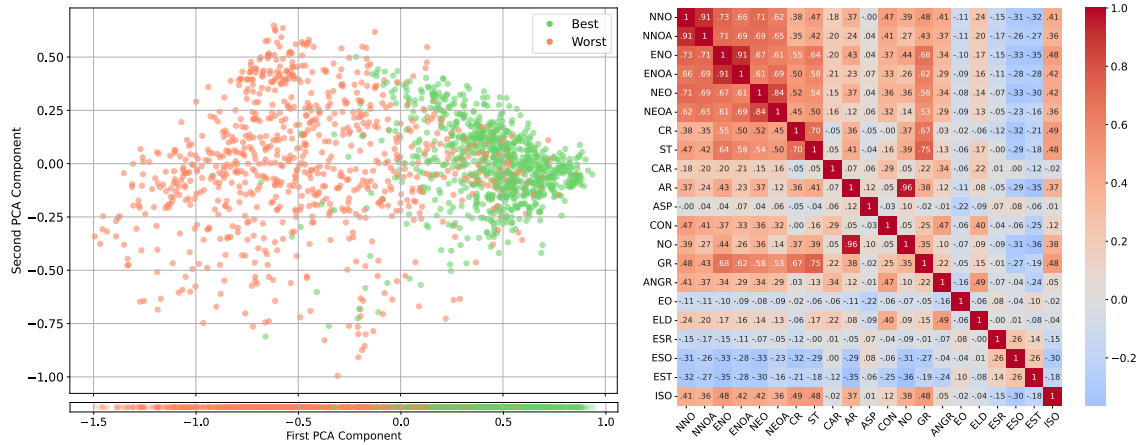


Figure 4: The first two components of a PCA of the *best* (green) and *worst* (red) user-chosen perspectives (left), in total  $2 \times 36 \times 23 = 1.656$  samples with 21 measures each. The projection on the first component (left, below) already shows a good separation of the data points. A correlation analysis indicates that some aesthetic measures are highly correlated (especially the overlap measures NN\*, EN\*, and NE\*), while others are not (right).

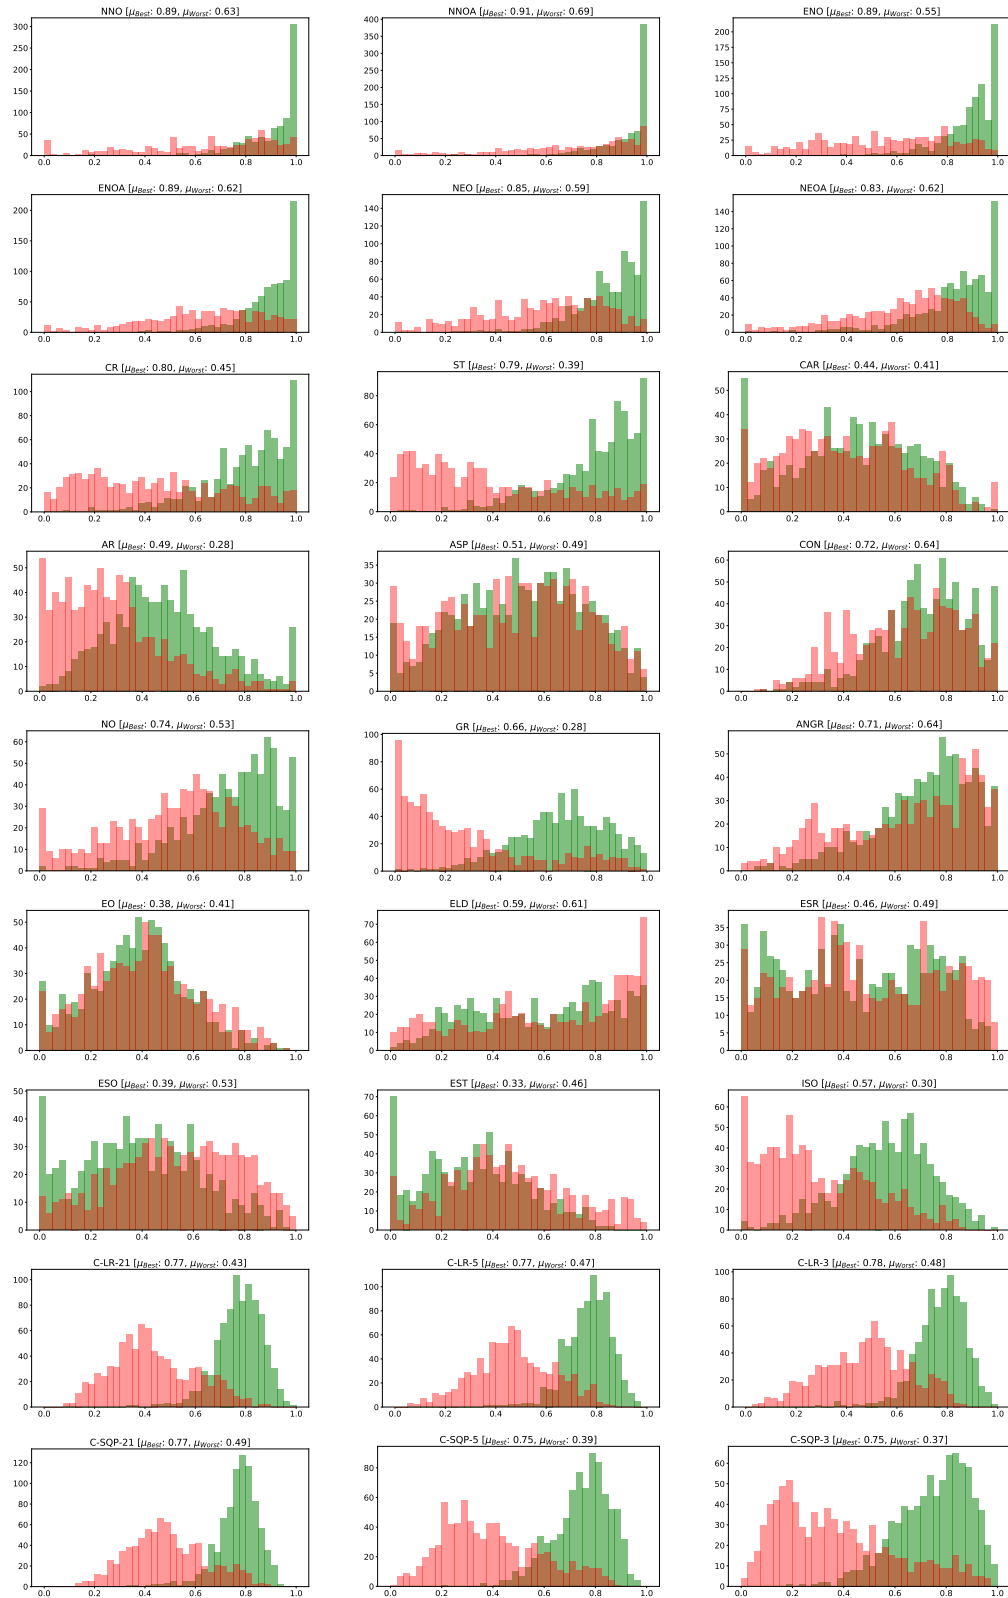


Figure 5: The distributions of values (*best*: green, *worst*: red) for 21 aesthetic measures (top seven rows) and combinations of these (LR and SQP) with 21, 5, or 3 measures each (last two rows).

relevant, while ISO is less. In L-L graphs, NEO dominates among overlap measures, with the usual others and especially ISO being important. L-E graphs show similar trends, though ENO is more prominent, while ISO is often excluded. Regarding graph size, smaller graphs appear more affected by overlaps (including NNO), whereas larger ones place more importance on ST and GR. Notably, the SQP method tends to favour measures like ST and GR, where best and worst scores are spread out more, even if some worst views receive high scores. In contrast, LR prefers measures that more clearly separate distributions, such as ISO and the overlap measures, resulting in better classifications (see Figure 5). For comparison with individual aesthetics, we computed linear combinations using weights from both C-LR (logistic regression) and C-SQP (sequential quadratic programming). The results appear to the right in all tables and in Figure 5.

## 4.5 Aesthetic Measure Landscapes

Besides investigating the aesthetic measure values for the user-selected viewpoints, we also analysed the entire landscape of aesthetics for all viewpoints of a graph structure. This allows us to further explore how user decisions interrelate with aesthetic measures and also provides insight into the distribution and local stability of values. In particular, in VR systems, where the perspective is never perfectly stable due to small head and body movements, it is desirable that selected perspectives are locally stable with regard to the visual appearance of the graph. To analyse the aesthetic measure landscapes, we used the *Fibonacci Lattice* method for sampling 5,000 uniformly distributed viewpoints (see Figure 6), for which we calculated all measures and their optimal combinations as described in Section 4.4 of five measures (SQP-5 and LR-5). The final results for all graphs can be found in Appendix A. For the 3D-specific overlap measures, we observe that the landscape appears highly scattered, and perspectives with a very high score are often located directly next to those with very low scores. This is especially the case for the count-only measures (NNO, NEO, ENO), while the area-based measures (NNOA, NEOA, ENOA) appear smoother. This is unsurprising, since the overlap area between two nodes changes gradually with small angular changes, whereas for count-only measures, there is no change until the overlap appears or disappears, leading to a more discontinuous landscape. Measures incorporating the area of the projection (AR, ASP) or the orthogonality of nodes or edges (NO, EO) typically lead to rather continuous value landscapes, as the area and orthogonality characteristics smoothly increase or decrease with small angular changes. Similarly, the deviation of edge lengths (ELD) and angular resolution (ANGR) changes in a continuous manner, resulting in a rather smooth value landscape. Analogously, for fixed main dimensions, the evaluation of how much a perspective deviates from an isometric view changes smoothly. Therefore, the ISO measure appears very similar across the graphs, with local maxima around the different isometric viewpoints and continuously decreasing when shifting towards axis-aligned views. We further find that ST and GR tend to lead to similar value landscapes across the viewpoints, but the results for ST appear much smoother and more continuous. While ST is based on the relation between topological and geometric distances between nodes, which changes continuously for small angular changes, GR measures whether nodes are located within the radius of edges, which can change abruptly for a small angular difference. This is also the case for CR, which relies on a count-based metric, namely the number of crossings. Furthermore, the CAR measure and node concentration (CON) yield rather diffuse values. Interestingly, the symmetry measures (ESR, ESO, EST) also lead to highly scattered values across the different viewpoints for all graphs and often do not provide a clear indication of where advantageous and disadvantageous perspectives are located. As they also do not align with any other measure, this might explain why the symmetry measures are almost never included in optimised combinations.

Table 3: The relative importance of the aesthetic measures when combined to express the perspective choices of users. We applied logistic regression (LR) and sequential quadratic programming (SQP) to find the most relevant aesthetic measures when using 21, 5, and 3 measures ( $|A|$ , i.e. shown as indicated measure cells per each row). We investigated the importance for all data at once (All), the three layout classes (L-S, L-L, L-E), and the four graph sizes (S-S, S-M, S-L, S-XL).

$ A $	Data	Method	NNO	NNOA	ENO	ENOA	NEO	NEOA	CR	ST	CAR	AR	ASP	CON	NO	GR	ANGR	EO	ELD	ESR	ESO	EST	ISO	
21	All	LR	.09	.03	.14	.11	.08	.04	.12	.11	-.02	.11	-.00	.01	.04	.08	.02	.04	-.17	.01	.00	-.02	.19	
		SQP	.07	.06	.09	.07	.07	.06	.10	.11	.01	.06	.01	.02	.06	.11	.02							.08
	L-S	LR	.08	.08	.10	.10	.03	.03	.05	.08	.00	.07	.03	.04	.04	.04	.09	.06	.01	.05	.04	-.01	-.05	.08
		SQP	.08	.07	.09	.08	.05	.04	.05	.09	.03	.07	.02	.04	.07	.09	.02			.04				.06
	L-L	LR	-.01	.02	.25	.21	.37	.30	.17	.13	-.05	.12	-.00	-.01	.07	.18	.09	-.04	-.82	-.04	-.01	-.18	.26	
		SQP	.05	.05	.08	.08	.09	.07	.12	.09	.00	.05	.04	.00	.05	.11	.02							.10
	L-E	LR	.13	.02	.13	.11	.03	-.02	.10	.19	-.03	.01	-.05	.08	.06	.14	-.04	.06	.02	.02	-.07	.00	.10	
		SQP	.08	.06	.09	.06	.08	.06	.11	.13		.04	.02	.04	.11	.02			.03				.06	
	S-S	LR	.08	.07	.11	.08	.09	.07	.13	.07	.00	.06	.01	.01	.08	.08	.08	.01	-.06	-.02	.01	-.03	.06	
		SQP	.08	.06	.09	.06	.08	.07	.07	.09		.05	.01	.04	.06	.08	.06		.04				.06	
	S-M	LR	.06	.09	.17	.16	.14	.04	.13	.06	.08	.11	-.01	-.01	.03	-.06	-.02	.05	-.19	.02	-.02	-.01	.18	
		SQP	.07	.07	.09	.08	.08	.07	.07	.08	.05	.05	.02	.04	.05	.07	.03		.01	.01			.05	
	S-L	LR	.07	.05	.06	.09	-.03	.01	.11	.20	.05	.04	.01	.05	.05	.11	.03	.04	-.06	.05	-.01	-.09	.17	
		SQP	.05	.05	.10	.08	.04	.03	.13	.15	.02	.05	.01	.06	.14				.00				.09	
	S-XL	LR	.08	.07	.19	.16	.07	.09	.23	.20	-.13	.07	.04	-.06	.11	.18	-.10	-.02	-.17	-.02	-.08	-.04	.12	
		SQP	.07	.05	.09	.07	.07	.04	.13	.12		.06			.06	.13							.09	
	5	All	LR			.32		.19		.22											.03		.24	
			SQP			.19				.20	.23							.22						.16
L-S		LR		.21	.29				.11			.24											.15	
		SQP	.18		.22	.18			.22								.20							
L-L		LR			.13	.36						.10							.05				.35	
		SQP				.17			.24	.17							.22						.20	
L-E		LR			.23	.13	.16		.22	.25							.24		.07					
		SQP			.17		.16		.22	.25							.21							
S-S		LR			.32	.29			.22		.16								.01					
		SQP	.20		.21	.20			.21		.21						.19							
S-M		LR	.21		.37	.20					.13						.02						.28	
		SQP			.23	.19	.20		.20								.18							
S-L		LR			.26				.34		.06	.03					.31							
		SQP			.16				.21	.25							.23						.15	
S-XL		LR	.12		.20				.29	.19						.20								
		SQP			.15				.23	.22							.23						.16	
3		All	LR			.41		.27															.32	
			SQP							.31	.35							.34						
	L-S	LR		.37								.31											.32	
		SQP			.34						.34												.32	
	L-L	LR				.42			.31														.27	
		SQP							.36														.31	
	L-E	LR			.29				.36														.35	
		SQP							.32	.37													.31	
	S-S	LR	.30		.42				.28															
		SQP			.34		.33			.34														
	S-M	LR			.52						.17												.31	
		SQP			.36		.32			.32														
	S-L	LR				.30			.35														.35	
		SQP							.30	.37													.33	
	S-XL	LR			.38				.42							.20								
		SQP							.34	.32													.34	

For the combined measures SQP-5 and LR-5, we observe a rather continuous landscape of values. This is especially the case for SQP-5, which puts more weight on highly continuous measures such as ST. This stability is valuable for evaluating perspective qualities and also for finding optimal perspectives (i.e., finding the global maximum of the combined measure) using methods such as gradient descent. The exemplary case of graph E-5 (see Figure 7) illustrates the benefits of combining measures with different objectives. The essential measures ST, GR, CR, and the overlap




Figure 6: For measure normalisation and aesthetic measure landscape analysis across all perspectives, we sampled 5,000 perspectives using *Fibonacci Lattice*. The orange spheres illustrate a subset of the sampled perspectives.

aesthetics all achieve the best scores for the perspective  $P3$ . However, the ISO aesthetic penalises this perspective, as it directly aligns with one of the main axes, making the variances in this direction disappear. Thus, the combined measures achieve the best score for a perspective that is slightly shifted away from the pole, which is in accordance with the user choices.

While the chosen perspectives are generally in accordance with the combined measures (i.e., high scores for *best*, low scores for *worst* perspectives), the perspectives chosen as *best* are sometimes not located at the global maximum. This may be due to the need for the combined measure to generalise across all graph types and sizes and to the restriction to five measures in this case, but it could also indicate that further measures are required to fully explain the user decisions.

## 5 Discussion

Our analysis of user-selected viewpoints revealed several user strategies and demonstrated how their choices interrelate with aesthetic measures. Notably, preferences varied across graph types and sizes, often reflecting trade-offs between structural clarity and local features. This supports the view that users do not optimise for a single aesthetic, but rather negotiate among multiple perceptual cues, many of which are not captured by traditional 2D aesthetics.

The strongest 2D aesthetic predictors of viewpoint preference were CR, ST, and GR, aligning with prior findings in 2D [4, 36], and extending their validity into S3D. However, 3D-specific measures, such as ISO and the overlap aesthetics, proved highly relevant. The prominence of ENO over NNO is particularly noteworthy. While one might expect node-node occlusions to be more perceptually disruptive, the results indicate that edge-node intersections may be more misleading. Node-edge occlusions appeared less impactful, possibly due to the Gestalt principle of continuity. Although overlap area was computed as a potentially more accurate representation, it showed no advantage over overlap counts. This may reflect that users prefer views in which objects are clearly distinguishable, as indicated by the relevance of GR penalising nodes and edges in close proximity. Another key insight lies in the differing impact of layout type. Layered and semantic layouts led to more consistent viewpoint choices than energy-based ones. In semantic graphs, where node positions carry intrinsic meaning, participants prioritised preserving perceived structure. This was reflected in the strong relevance of ISO for semantic and layered layouts, where users

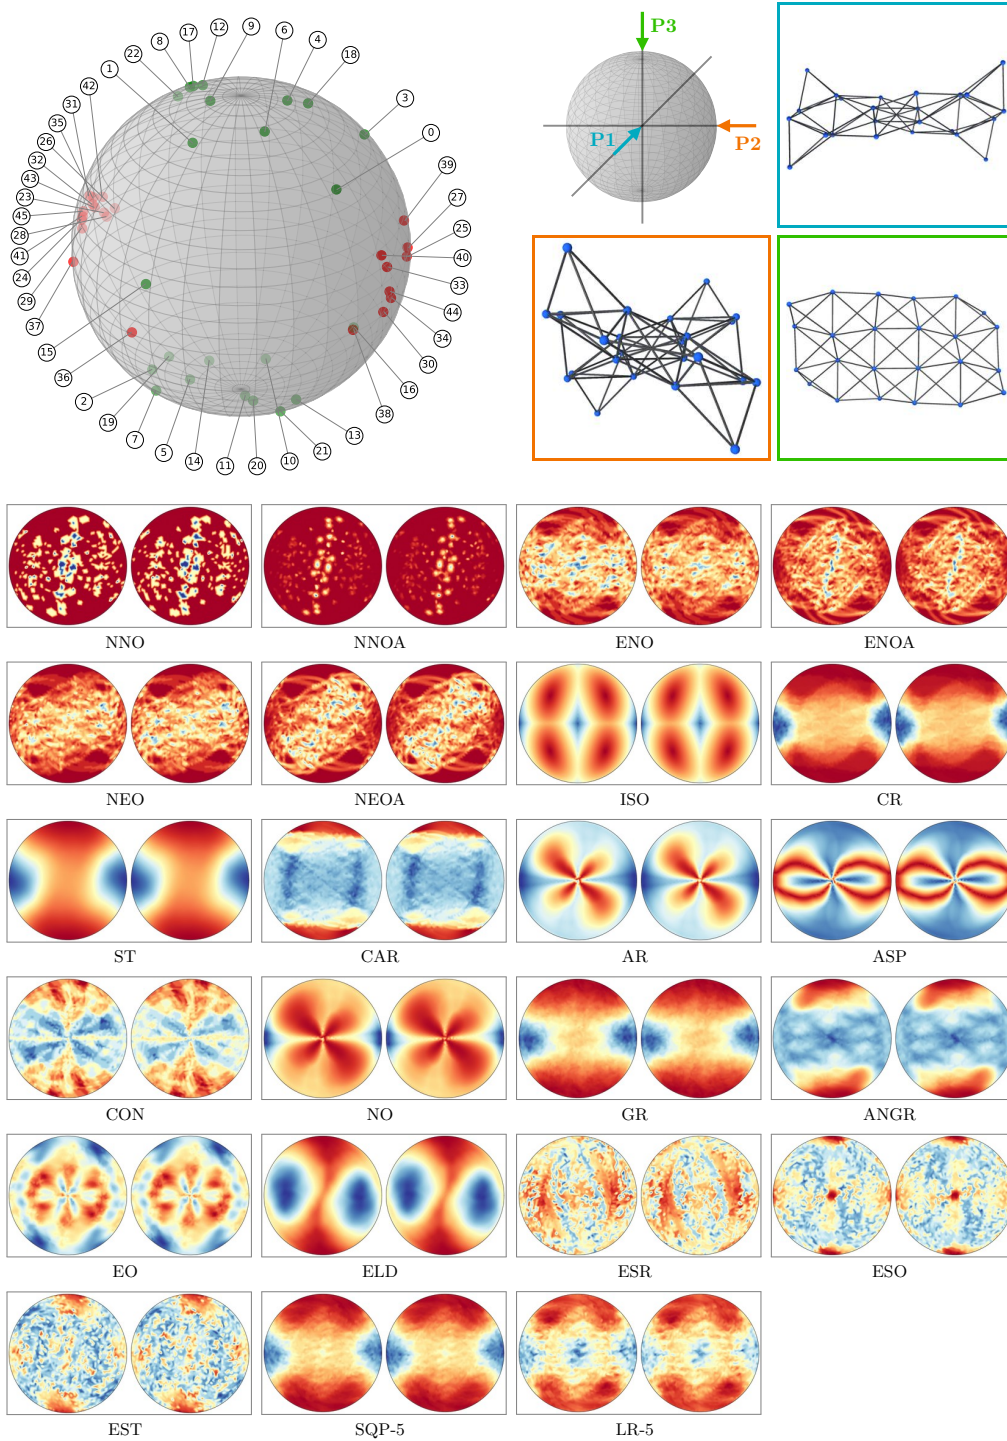


Figure 7: The landscapes of all normalised aesthetic measures and combinations (LR and SQP, five measures only) for graph **E-2** (can.24) as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red for optimal scores). See [Appendix A](#) for all graphs.

appeared to prefer perspectives maintaining orthogonal balance or offering diagonal overviews. In contrast, energy-based graphs, which showed greater topological variation in our dataset, elicited more diverse and individual strategies, with user preferences often diverging even within the same graph. These cases suggest that, in the absence of salient or meaningful structure, users default to evaluating local features such as component separation, edge crossings, or occlusions.

One somewhat surprising result was the limited explanatory power of several aesthetics such as ANGR, ASP, and symmetry measures. Although considered beneficial in 2D, these did not significantly distinguish best and worst views in our study. One explanation may be that such aesthetics target fine-grained aspects of local visual quality, while participants were more influenced by coarse, global cues, such as whether the overall structure ‘popped out’ in space, or whether the orientation supported mental model construction. Another possibility is that these aesthetics are harder to judge consistently in stereoscopic views, where depth cues from parallax and motion influence perception in ways not captured by 2D projections.

PCA and correlation analysis revealed that many aesthetic measures are indeed redundant, especially among the overlap measures. Thus, a small number of well-chosen features can explain much of the variance in user preferences. Beyond redundancy, our analysis of aesthetic landscapes highlights that not all measures are equally suitable for optimisation. Smooth landscapes such as those of ST or ISO support stable local maxima, making them robust for viewpoint selection and gradient-based optimisation. In contrast, highly scattered landscapes, especially among count-based measures, exhibit abrupt discontinuities that complicate optimisation and may lead to unstable viewpoint predictions. This distinction suggests that landscape stability should be considered alongside predictive power when selecting measures for automated viewpoint recommendation. Our comparison of linear combination methods (logistic regression and sequential quadratic programming) confirms that different optimisation objectives yield subtly different rankings, though both approaches consistently identify the most predictive measures. Combinations incorporating ST, GR, CR, ENO, and ISO consistently achieved the best separation between preferred and unfavoured viewpoints.

Beyond the numerical results, qualitative feedback reinforced the value of 3D interaction and stereoscopic cues in disambiguating structure, particularly in dense or ambiguous graphs. Interestingly, user choices were largely consistent, with no substantial differences between novices and more experienced users. This reinforces the value of identifying general principles for viewpoint preference, as we aimed to do in this work.

Despite careful design, this study has limitations. Although we included a broad sample of graphs and layouts, results may differ for other or more specific datasets. Moreover, our participants were primarily novices and generalisation to expert users or larger populations may require further study. While we focused on user *preference*, the relationship to task-solving performance remains unclear. Thus, future work could build on previously introduced methods to identify aesthetic-optimal viewpoints, incorporating our findings on preference-driven combinations to investigate their effects on task performance. Moreover, our analysis has so far treated each viewpoint as static and independent, whereas in real-world use, users navigate through sequences of views. Studying viewpoint transitions and the temporal dynamics of understanding could offer insights into how users mentally assemble spatial structures over time. Another open question concerns how aesthetic measure landscapes align with user navigation in immersive environments. Since users rarely remain at a static viewpoint, locally smooth measures may better capture the perceptual continuity experienced during small viewpoint changes. Conversely, scattered measures may explain why some perspectives felt unstable or less satisfying despite achieving high scores. Future work could therefore examine how landscape properties relate to dynamic viewpoint exploration and whether

new measures can be designed to combine predictive power with local stability. Furthermore, based on our tests with the ISO measure, we see potential for developing new, 3D-specific aesthetic measures that reflect how users perceive graphs in a S3D setup. Using our openly accessible dataset, such measures can be evaluated against the user preferences collected in this study.

## 6 Conclusion

In this work, we conducted a comprehensive study with 23 participants on viewpoint preference in 3D graph visualisation. Participants selected their most and least preferred viewpoints across 36 graphs varying in layout, size, and topology, and shared their selection strategies. Our findings show that, in addition to classical 2D aesthetic measures such as *Stress*, *Crossings*, and the *Gabriel Ratio*, 3D-specific measures, especially *Edge-Node Overlap* and overall structure-capturing methods like our proposed *Isometric Viewpoint Deviation*, are highly reflecting user preference. Participants were often willing to tolerate more overlap and crossings when a viewpoint conveyed the global structure more clearly. Our analysis of combined aesthetics using logistic regression and sequential quadratic programming underlines these results, showing that a small set of measures characterises user choices to a high degree. However, the results indicate the need to adapt and extend traditional quality measures for the 3D context. In addition, our analysis of aesthetic measure landscapes revealed that some measures yield smooth and stable distributions across viewpoints, while others are highly scattered and discontinuous. This distinction is crucial for practical applications, as stable landscapes enable reliable optimisation and support perceptual continuity when users explore viewpoints dynamically. Beyond our empirical findings, the open-access dataset provides a foundation for future research on viewpoint evaluation and optimisation in immersive graph analysis.

## References

- [1] S. Achard, J.-F. Coeurjolly, P. L. de Micheaux, H. Lbath, and J. Richiardi. Inter-regional correlation estimators for functional magnetic resonance imaging. *NeuroImage*, 282:120388, 2023. doi:10.1016/j.neuroimage.2023.120388.
- [2] G. J.-P. C. Becq, E. L. Barbier, S. Achard, and H. Lbath. Preprocessed rat brain voxel time series, Nov. 2022. Zenodo. doi:10.5281/zenodo.7254133.
- [3] C. Bennett, J. Ryall, L. Spalteholz, and A. Gooch. The Aesthetics of Graph Visualization. In D. W. Cunningham, G. Meyer, and L. Neumann, editors, *Computational Aesthetics in Graphics, Visualization, and Imaging*. The Eurographics Association, 2007. doi:10.2312/COMPAESTH/COMPAESTH07/057-064.
- [4] M. Chimani, P. Eades, S.-H. Hong, W. Huang, K. Klein, M. Marner, R. Smith, and B. Thomas. People prefer less stress and fewer crossings. In *Graph Drawing: 22nd International Symposium*, volume 8871 of *LNCS*, pages 523–524, 2014.
- [5] M. Chimani, C. Gutwenger, M. Jünger, G. W. Klau, K. Klein, and P. Mutzel. The open graph drawing framework (ogdf). *Handbook of graph drawing and visualization*, 2011:543–569, 2013.
- [6] M. K. Coleman and D. S. Parker. Aesthetics-based graph layout for human consumption. *Software: Practice and Experience*, 26(12):1415–1438, 1996. doi:10.1002/(SICI)1097-024X(199612)26:12<1415::AID-SPE69>3.0.CO;2-P.

- [7] T. A. Davis and Y. Hu. The university of florida sparse matrix collection. *ACM Transactions on Mathematical Software*, 38(1), 2011. doi:10.1145/2049662.2049663.
- [8] S. Di Bartolomeo, T. Crnovrsanin, D. Saffo, E. Puerta, C. Wilson, and C. Dunne. Evaluating graph layout algorithms: A systematic review of methods and best practices. *Computer Graphics Forum*, 43(6):e15073, 2024. doi:10.1111/cgf.15073.
- [9] G. Di Battista, A. Garg, G. Liotta, R. Tamassia, E. Tassinari, and F. Vargiu. An experimental comparison of four graph drawing algorithms. *Computational Geometry*, 7(5):303–325, 1997. 11th ACM Symposium on Computational Geometry. doi:10.1016/S0925-7721(96)00005-3.
- [10] A. Drogemuller, A. Cunningham, J. Walsh, J. Baumeister, R. T. Smith, and B. H. Thomas. Aesthetics and comprehension of curved 3D graphs in virtual reality. *Journal of Computer Languages*, 75:101198, 2023. doi:10.1016/j.cola.2023.101198.
- [11] A. Drogemuller, A. Cunningham, J. Walsh, W. Ross, and B. H. Thomas. Vrige: exploring social network interactions in immersive virtual environments. In *Proceedings of the international symposium on big data visual analytics (BDVA)*, volume 5, 2017. doi:10.13140/RG.2.2.34201.67689.
- [12] L. Fan, H. Li, J. Zhuo, Y. Zhang, J. Wang, L. Chen, Z. Yang, C. Chu, S. Xie, A. R. Laird, P. T. Fox, S. B. Eickhoff, C. Yu, and T. Jiang. The human brainnetome atlas: A new brain atlas based on connectional architecture. *Cerebral Cortex*, 26(8):3508–3526, 2016. doi:10.1093/cercor/bhw157.
- [13] W. Farish. On isometrical perspective. *Cambridge Philosophical Transactions*, 1822.
- [14] S. P. Feyer, B. Pinaud, S. Kobourov, N. Brich, M. Krone, A. Kerren, M. Behrisch, F. Schreiber, and K. Klein. 2D, 2.5D, or 3D? An exploratory study on multilayer network visualisations in virtual reality. *IEEE Transactions on Visualization & Computer Graphics*, 30(1):469–479, 2024. doi:10.1109/TVCG.2023.3327402.
- [15] N. Greffard, F. Picarougne, and P. Kuntz. Beyond the classical monoscopic 3D in graph analytics: An experimental study of the impact of stereoscopy. In *2014 IEEE VIS International Workshop on 3DVis*, pages 19–24, 2014. doi:10.1109/3DVis.2014.7160095.
- [16] W. Huang, P. Eades, and S.-H. Hong. Measuring effectiveness of graph visualizations: A cognitive load perspective. *Information Visualization*, 8(3):139–152, 2009. doi:10.1057/ivs.2009.10.
- [17] W. Huang, P. Eades, S.-H. Hong, and C.-C. Lin. Improving multiple aesthetics produces better graph drawings. *Journal of Visual Languages & Computing*, 24(4):262–272, 2013. doi:10.1016/j.jvlc.2011.12.002.
- [18] S. Jaeger, K. Klein, L. Joos, J. Zagermann, M. de Ridder, J. Kim, J. Yang, U. Pfeil, H. Reiterer, and F. Schreiber. Challenges for brain data analysis in VR environments. In *2019 IEEE Pacific Visualization Symposium (PacificVis)*, pages 42–46, 2019. doi:10.1109/PacificVis.2019.00013.
- [19] L. Joos, M. T. Fischer, D. A. Keim, and J. Fuchs. Aesthetic-driven navigation for node-link diagrams in vr. In *Proceedings of the 2023 ACM Symposium on Spatial User Interaction, SUI '23*. ACM, 2023. doi:10.1145/3607822.3614537.

- [20] L. Joos, M. T. Fischer, J. Rauscher, D. A. Keim, T. Dwyer, F. Schreiber, and K. Klein. Visual network analysis in immersive environments: A survey, 2025. [arXiv:2501.08500](https://arxiv.org/abs/2501.08500).
- [21] L. Joos, S. Jaeger-Honz, F. Schreiber, D. A. Keim, and K. Klein. Visual comparison of networks in vr. *IEEE Transactions on Visualization and Computer Graphics*, 28(11):3651–3661, 2022. [doi:10.1109/TVCG.2022.3203001](https://doi.org/10.1109/TVCG.2022.3203001).
- [22] L. Joos, G. J. Mooney, M. T. Fischer, D. A. Keim, F. Schreiber, H. Purchase, and K. Klein. Show me your best side: Characteristics of user-preferred perspectives for 3D graph drawings. In *33rd International Symposium on Graph Drawing and Network Visualization (GD 2025)*, 2025. [doi:10.4230/LIPIcs.GD.2025.35](https://doi.org/10.4230/LIPIcs.GD.2025.35).
- [23] L. Joos, G. J. Mooney, M. T. Fischer, D. A. Keim, F. Schreiber, H. C. Purchase, and K. Klein. Replication Data for: “Show Me Your Best Side: Characteristics of User-Preferred Perspectives for 3D Graph Drawings”, 2025. [doi:10.18419/DARUS-5108](https://doi.org/10.18419/DARUS-5108).
- [24] S. Kieffer, T. Dwyer, K. Marriott, and M. Wybrow. HOLA: Human-like orthogonal network layout. *IEEE Transactions on Visualization & Computer Graphics*, 22(01):349–358, 2016. [doi:10.1109/TVCG.2015.2467451](https://doi.org/10.1109/TVCG.2015.2467451).
- [25] R. Klapaukh, S. Marshall, and D. Pearce. A symmetry metric for graphs and line diagrams. In *Diagrammatic Representation and Inference*, pages 739–742. Springer International Publishing, 2018. [doi:10.1007/978-3-319-91376-6\\_71](https://doi.org/10.1007/978-3-319-91376-6_71).
- [26] J. Kotlarek, O.-H. Kwon, K.-L. Ma, P. Eades, A. Kerren, K. Klein, and F. Schreiber. A study of mental maps in immersive network visualization. In *2020 IEEE Pacific Visualization Symposium (PacificVis)*, pages 1–10, 2020. [doi:10.1109/PacificVis48177.2020.4722](https://doi.org/10.1109/PacificVis48177.2020.4722).
- [27] P. Kourtesis, S. Collina, L. A. A. Dumas, and S. E. MacPherson. Validation of the virtual reality neuroscience questionnaire: Maximum duration of immersive virtual reality sessions without the presence of pertinent adverse symptomatology. *Frontiers in Human Neuroscience*, 13, 2019. [doi:10.3389/fnhum.2019.00417](https://doi.org/10.3389/fnhum.2019.00417).
- [28] M. Kuznetsov, A. Elor, S. Kurniawan, C. Bosworth, Y. Rosen, N. Heyer, M. Teodorescu, B. Paten, and D. Haussler. The immersive graph genome explorer: Navigating genomics in immersive virtual reality. In *2021 IEEE 9th International Conference on Serious Games and Applications for Health (SeGAH)*, pages 1–8, 2021. [doi:10.1109/SEGAH52098.2021.9551857](https://doi.org/10.1109/SEGAH52098.2021.9551857).
- [29] M. R. Marner, R. T. Smith, B. H. Thomas, K. Klein, P. Eades, and S.-H. Hong. GION: Interactively untangling large graphs on wall-sized displays. In *Graph Drawing: 22nd International Symposium*, volume 8871 of *LNCS*, pages 113–124. Springer, 2014. [doi:10.1007/978-3-662-45803-7\\_10](https://doi.org/10.1007/978-3-662-45803-7_10).
- [30] F. McGee, M. Ghoniem, G. Melançon, B. Otjacques, and B. Pinaud. The state of the art in multilayer network visualization. *Computer Graphics Forum*, 38(6):125–149, 2019. [doi:10.1111/cgf.13610](https://doi.org/10.1111/cgf.13610).
- [31] M. J. McGuffin, R. Servera, and M. Forest. Path tracing in 2d, 3d, and physicalized networks. *IEEE TVCG*, 30(7):3564–3577, 2024. [doi:10.1109/TVCG.2023.3238989](https://doi.org/10.1109/TVCG.2023.3238989).

- [32] A. Meidiana, S.-H. Hong, P. Eades, and D. Keim. Quality metrics for symmetric graph drawings. In *2020 IEEE Pacific Visualization Symposium (PacificVis)*, pages 11–15, 2020. doi:10.1109/PacificVis48177.2020.1022.
- [33] G. J. Mooney, T. Hegemann, A. Wolff, M. Wybrow, and H. C. Purchase. Universal quality metrics for graph drawings: Which graphs excite us most? In *33rd International Symposium on Graph Drawing and Network Visualization (GD 2025)*, pages 1–21, 2025. doi:10.4230/LIPIcs.GD.2025.28.
- [34] G. J. Mooney, J. Miller, M. Wybrow, S. Kobourov, and H. C. Purchase. Stress in graph drawings: Perception, preference, and performance. In *33rd International Symposium on Graph Drawing and Network Visualization (GD 2025)*, pages 1–25, 2025. doi:10.4230/LIPIcs.GD.2025.36.
- [35] G. J. Mooney, H. C. Purchase, M. Wybrow, and S. G. Kobourov. The multi-dimensional landscape of graph drawing metrics. In *2024 IEEE 17th Pacific Visualization Conference (PacificVis)*, pages 122–131, 2024. doi:10.1109/PacificVis60374.2024.00022.
- [36] G. J. Mooney, H. C. Purchase, M. Wybrow, S. G. Kobourov, and J. Miller. The perception of stress in graph drawings. In *32nd International Symposium on Graph Drawing and Network Visualization (GD 2024)*, volume 320 of *LIPIcs*, pages 21:1–21:17. Schloss Dagstuhl, 2024. doi:10.4230/LIPIcs.GD.2024.21.
- [37] M. Nöllenburg, S. Röder, and M. Wallinger. GdMetriX - A NetworkX extension for graph drawing metrics. In *32nd International Symposium on Graph Drawing and Network Visualization (GD 2024)*, volume 320 of *LIPIcs*, pages 45:1–45:3. Schloss Dagstuhl, 2024. doi:10.4230/LIPIcs.GD.2024.45.
- [38] R. Pan, H. C. Purchase, T. Dwyer, and W. Chen. Extending adjacency matrices to 3d with triangles. In *2024 IEEE 17th Pacific Visualization Conference (PacificVis)*, pages 232–241, 2024. doi:10.1109/PacificVis60374.2024.00033.
- [39] S. Pirch, F. Müller, E. Iofinova, J. Pazmandi, C. V. R. Hütter, M. Chiettoni, C. Sin, K. Boztug, I. Podkosova, H. Kaufmann, and J. Menche. The vrnetzer platform enables interactive network analysis in virtual reality. *Nature Communications*, 12(1):2432, 2021. doi:10.1038/s41467-021-22570-w.
- [40] H. C. Purchase. Metrics for graph drawing aesthetics. *Journal of Visual Languages & Computing*, 13(5):501–516, 2002. doi:10.1006/jv1c.2002.0232.
- [41] H. C. Purchase, C. Pilcher, and B. Plimmer. Graph drawing aesthetics—created by users, not algorithms. *IEEE Transactions on Visualization & Computer Graphics*, 18:81–92, 12 2012. doi:10.1109/TVCG.2010.269.
- [42] R. Ratan, Q. Lin, C. Lim, R. Park, A. Lover, E. Han, D. Jang, A. P. Leith, and J. N. Bailenson. Time matters in vr: Students benefit from longer vr class duration, but certain outcomes decline after 45 minutes, with large individual variance. *Computers & Education*, 235:105328, 2025. doi:10.1016/j.compedu.2025.105328.
- [43] K. Ryall, J. Marks, and S. Shieber. An interactive constraint-based system for drawing graphs. In *Proceedings of the 10th Annual ACM Symposium on User Interface Software and Technology*, UIST '97, pages 97–104. ACM, 1997. doi:10.1145/263407.263521.

- [44] A. Schaefer, R. Kong, E. M. Gordon, T. O. Laumann, X.-N. Zuo, A. J. Holmes, S. B. Eickhoff, and B. T. T. Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral Cortex*, 28(9):3095–3114, 2017. doi:10.1093/cercor/bhx179.
- [45] M. Spönemann, B. Duderstadt, and R. von Hanxleden. Evolutionary meta layout of graphs. In *Diagrammatic Representation and Inference*, pages 16–30. Springer, 2014. doi:10.1007/978-3-662-44043-8\_3.
- [46] M. Taylor and P. Rodgers. Applying graphical design techniques to graph visualisation. In *Ninth International Conference on Information Visualisation (IV'05)*, pages 651–656, 2005. doi:10.1109/IV.2005.19.
- [47] F. Tijhuis, M. Schepers, E. Centeno, B. de Maciel, L. Douw, and F. Nobrega Santos. Human connectome project rfMRI repository. <https://github.com/floristijhuis/HCP-rfMRI-repository>, 2025. San Francisco (CA): GitHub; [Last Accessed 30.09.2025].
- [48] F. B. Tijhuis, M. Schepers, E. Centeno, B. d. A. Maciel, L. Douw, and F. Nobrega Santos. Human connectome project resting-state fmri connectivity matrices (young adult + aging), June 2022. Zenodo. doi:10.5281/zenodo.6770120.
- [49] S. v. Wageningen, T. Mchedlidze, and A. Telea. An experimental evaluation of viewpoint-based 3D graph drawing. *Computer Graphics Forum*, 43(3), 2024. doi:10.1111/cgf.15077.
- [50] S. v. Wageningen, T. Mchedlidze, and A. Telea. Viewpoint Optimization for 3D Graph Drawings. *Computer Graphics Forum*, 44(3), 2025. doi:10.1111/cgf.70127.
- [51] C. Ware, H. Purchase, L. Colpoys, and M. McGill. Cognitive measurements of graph aesthetics. *Information Visualization*, 1(2):103–110, 2002. doi:10.1057/palgrave.ivs.9500013.
- [52] E. W. Weisstein. Circle-circle intersection. <https://mathworld.wolfram.com/Circle-CircleIntersection.html>. From MathWorld—A Wolfram Web Resource. [Last Accessed 30.09.2025].
- [53] E. Welch and S. Kobourov. Measuring symmetry in drawings of graphs. *Computer Graphics Forum*, 36(3):341–351, 2017. doi:10.1111/cgf.13192.
- [54] V. Yoghourdjian, Y. Yang, T. Dwyer, L. Lawrence, M. Wybrow, and K. Marriott. Scalability of network visualisation from a cognitive load perspective. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1677–1687, 2021. doi:10.1109/TVCG.2020.3030459.
- [55] L. Zhang, Y. Wang, Y. Zhang, W. Kerle-Malcharek, K. Klein, F. Schreiber, and A. Bulling. Towards a better understanding of graph perception in immersive environments. In *33rd International Symposium on Graph Drawing and Network Visualization (GD 2025)*, pages 1–19, 2025. doi:10.4230/LIPIcs.GD.2025.9.
- [56] Y. Zhang, N. Groene, K. Klein, G. Liotta, and F. Schreiber. Investigating crossing perception in 3D graph visualisation, 2025. arXiv:2508.00950.

## A Selected Perspectives & Aesthetic Landscapes

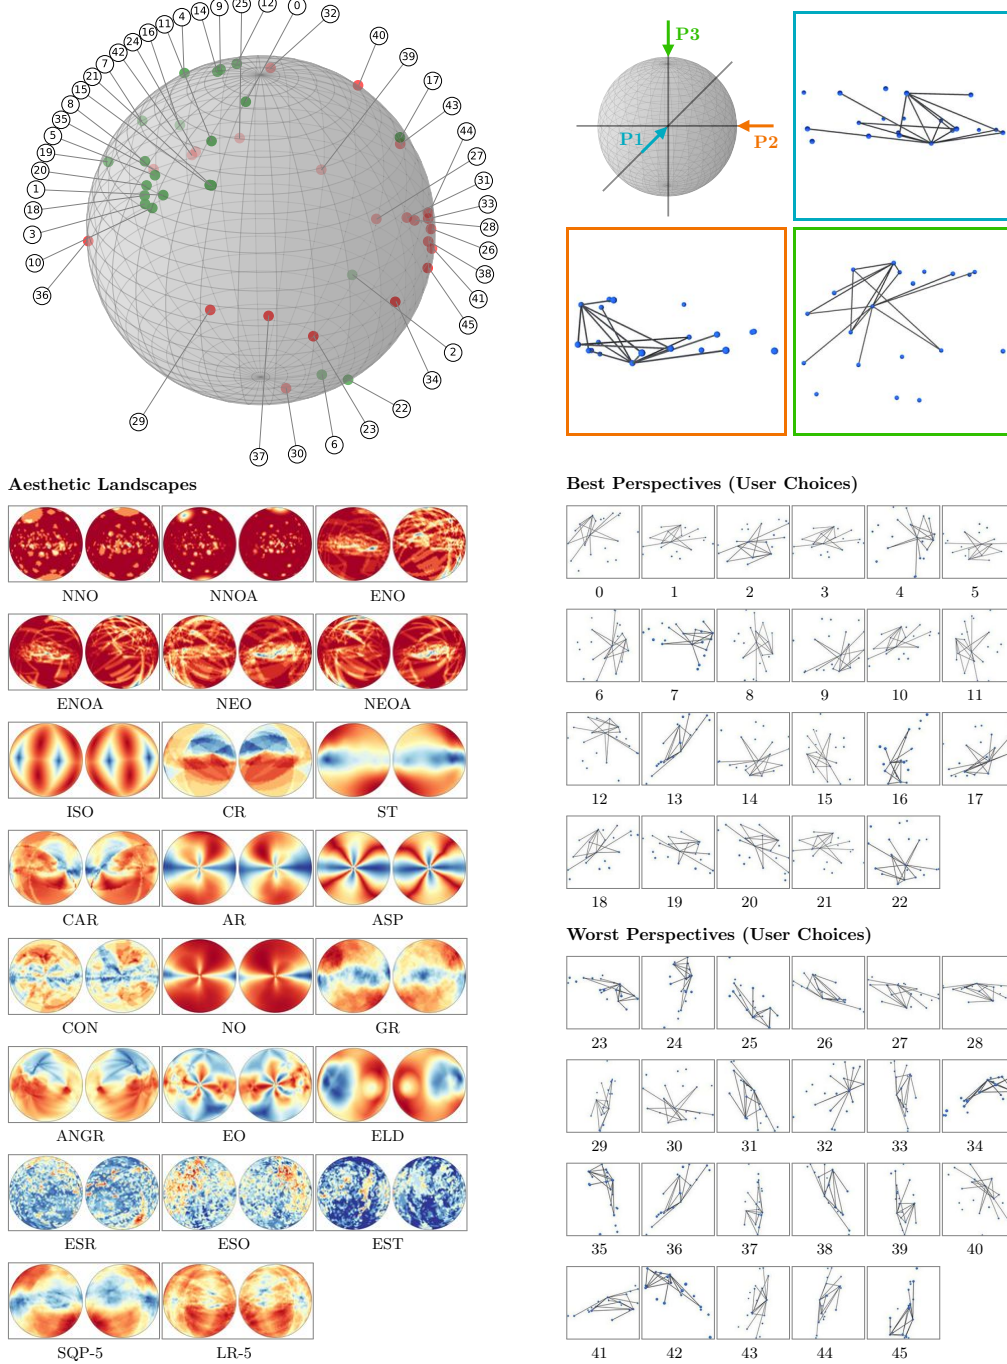


Figure A1: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-0** (sem.0\_20\_16) with  $|V| = 20$ ,  $|E| = 16$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

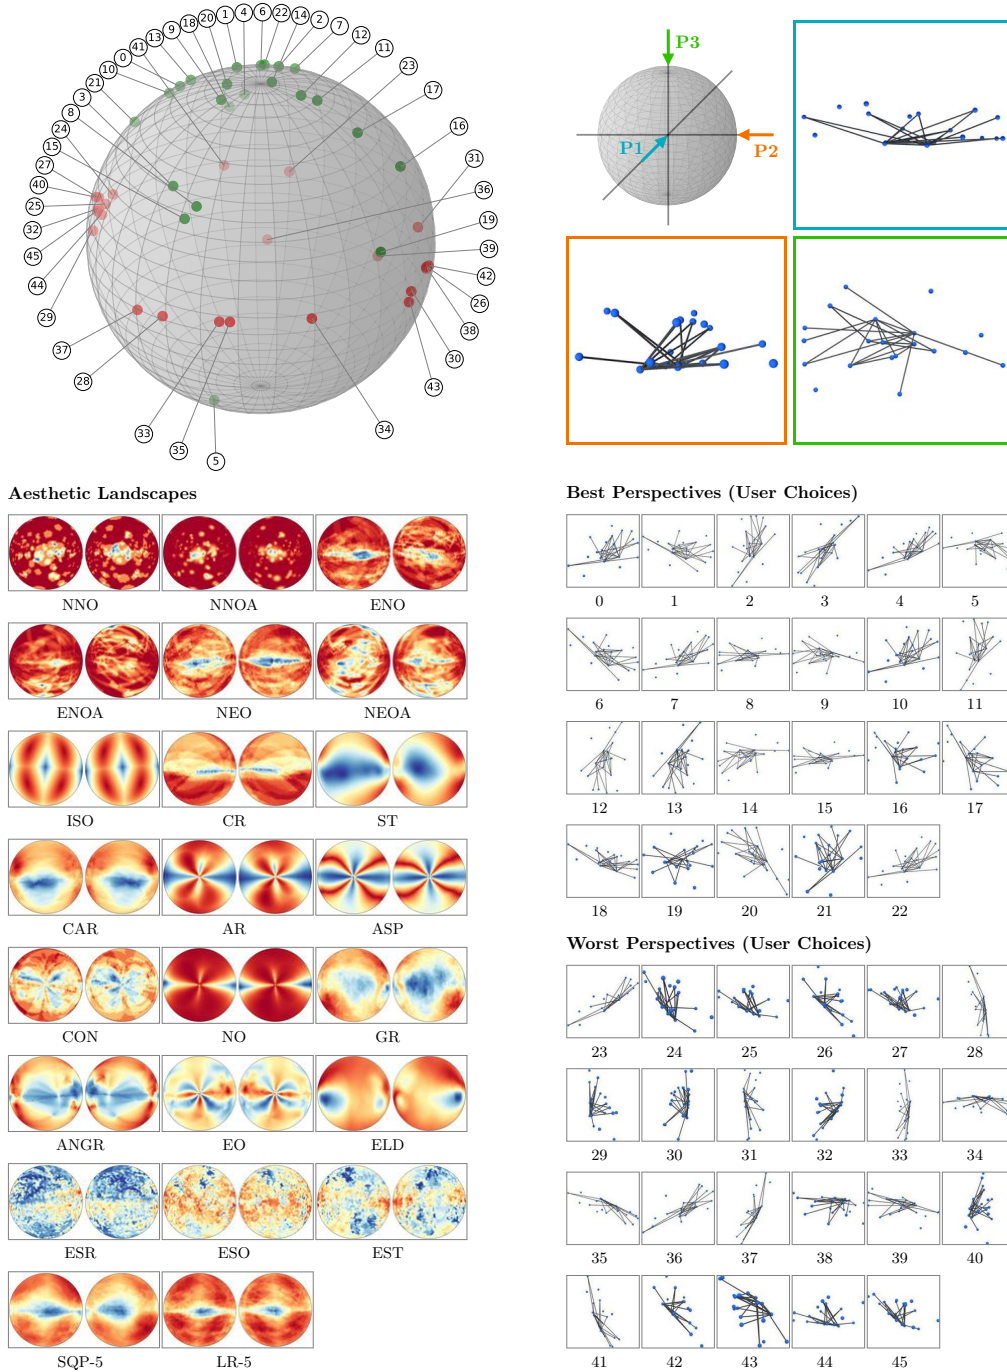


Figure A2: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-1** (sem.1.20.22) with  $|V| = 20$ ,  $|E| = 22$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

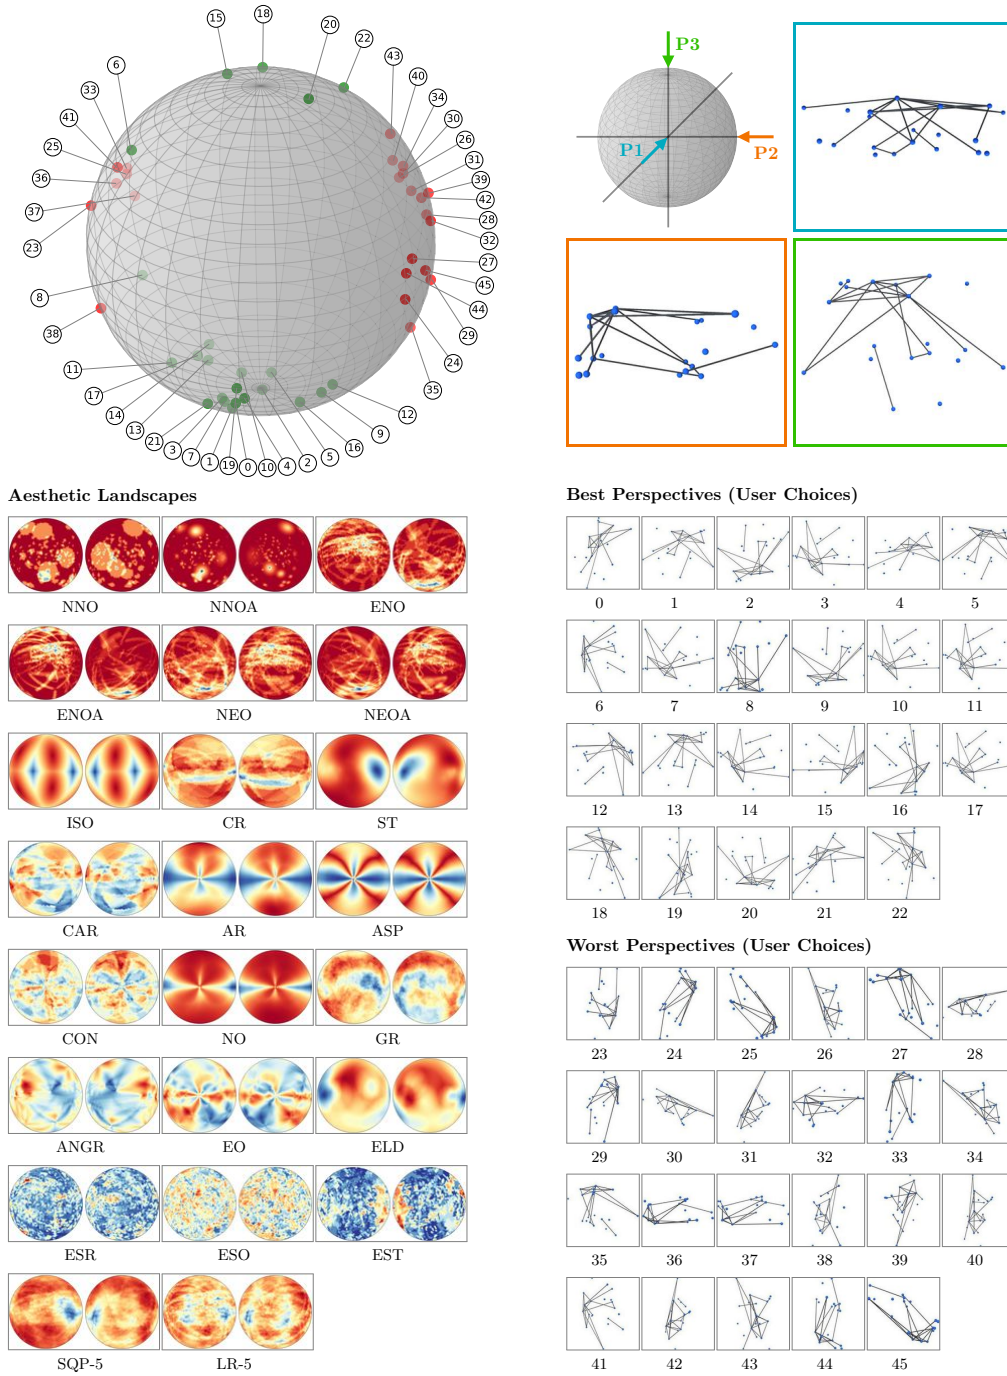


Figure A3: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-2** (sem.2.20.19) with  $|V| = 20$ ,  $|E| = 19$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

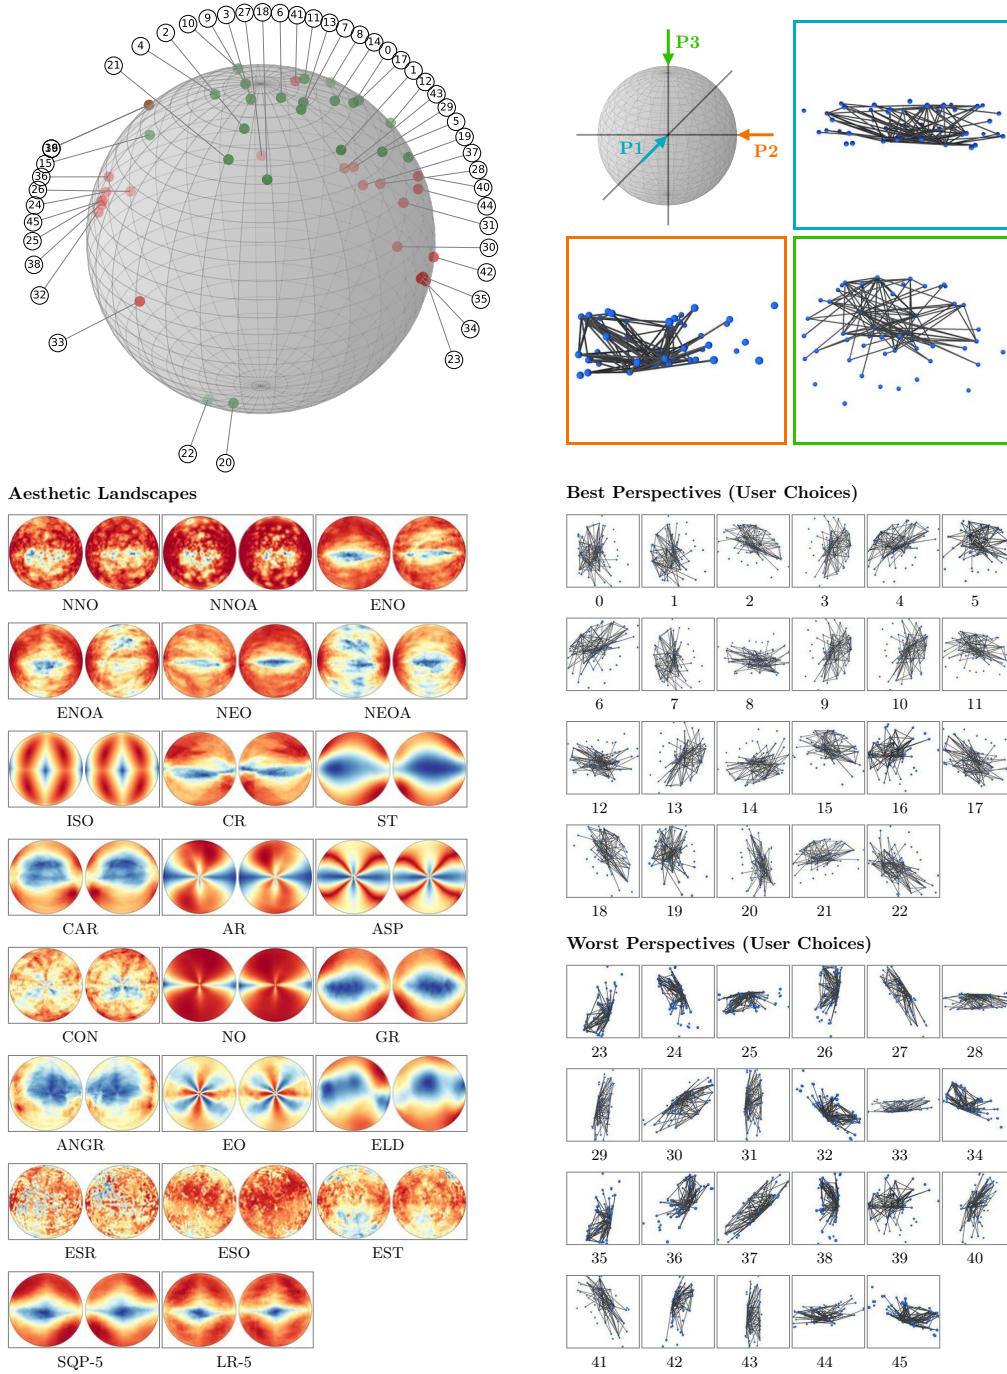


Figure A4: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-3** (sem\_0.50\_100) with  $|V| = 50$ ,  $|E| = 100$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

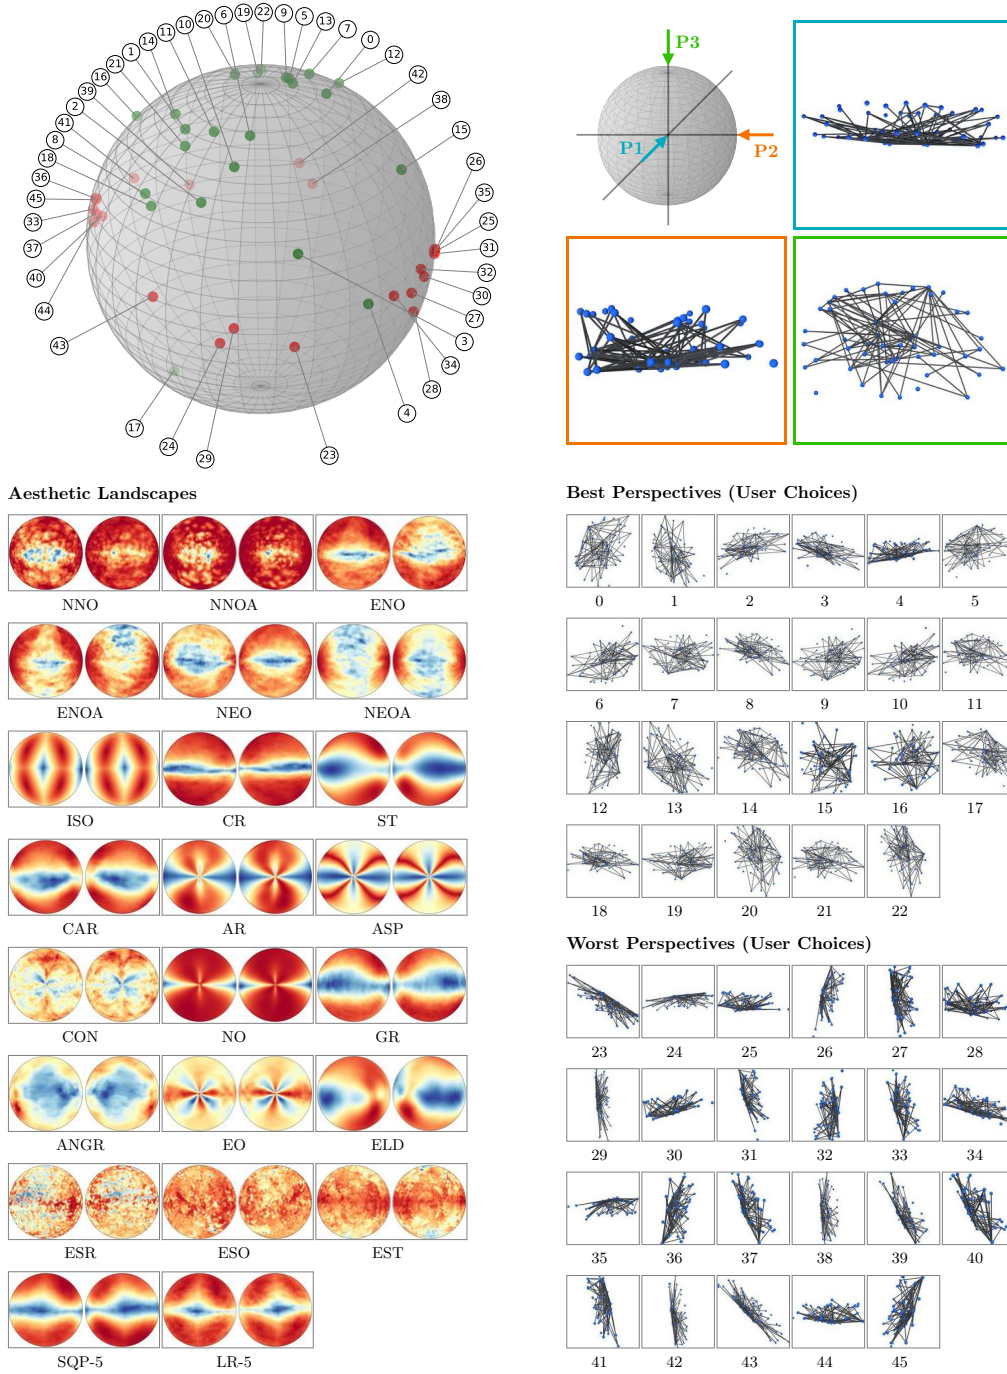


Figure A5: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-4** (sem\_1.50\_100) with  $|V| = 50$ ,  $|E| = 100$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

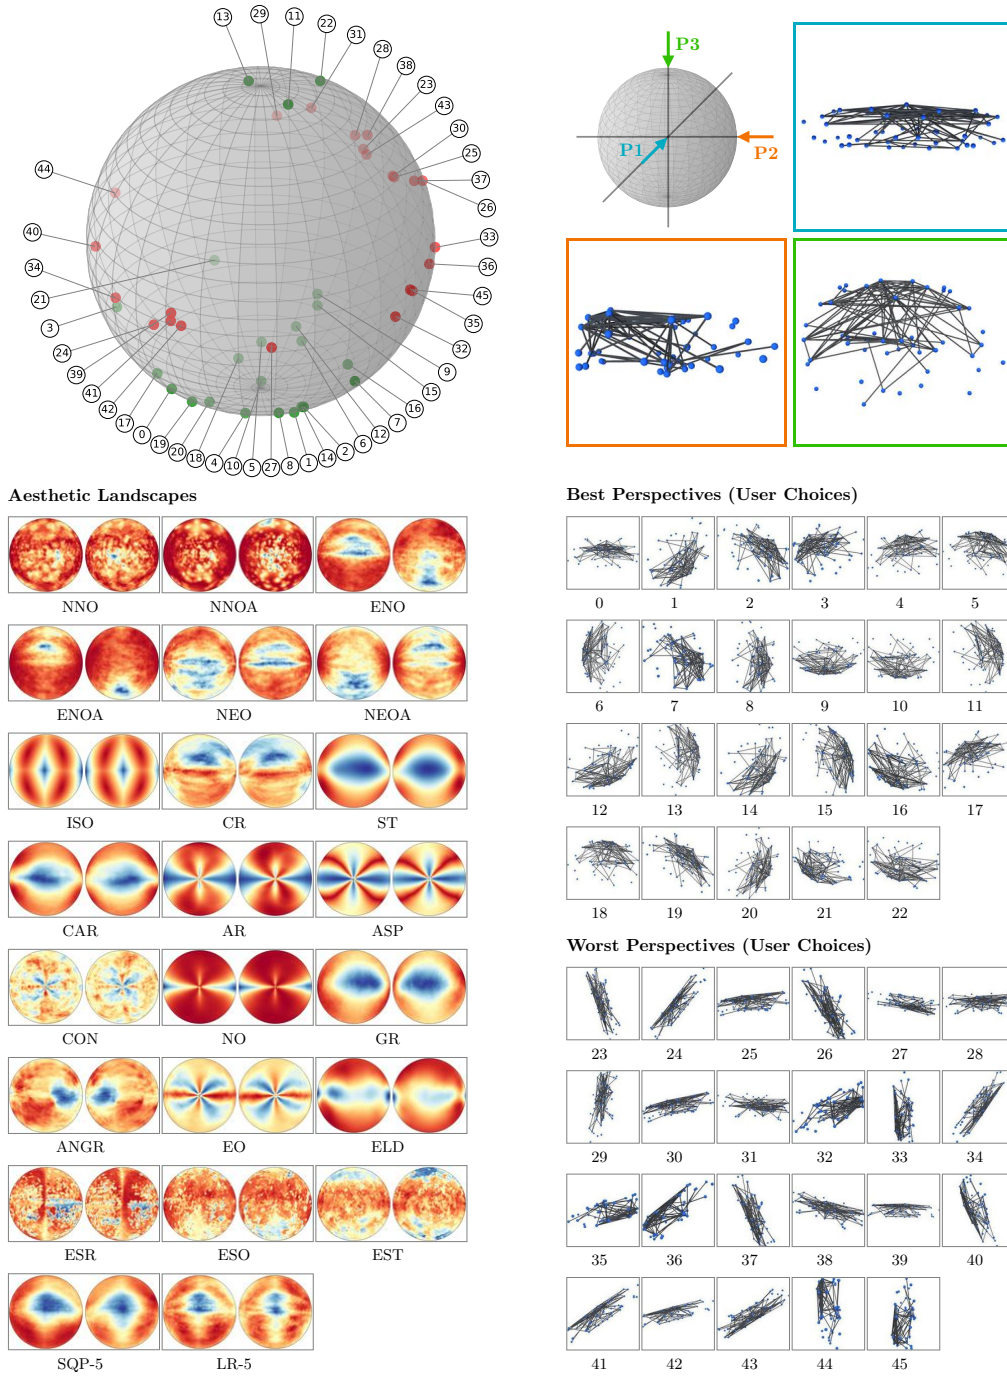


Figure A6: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-5** (sem\_2.50.100) with  $|V| = 50$ ,  $|E| = 100$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

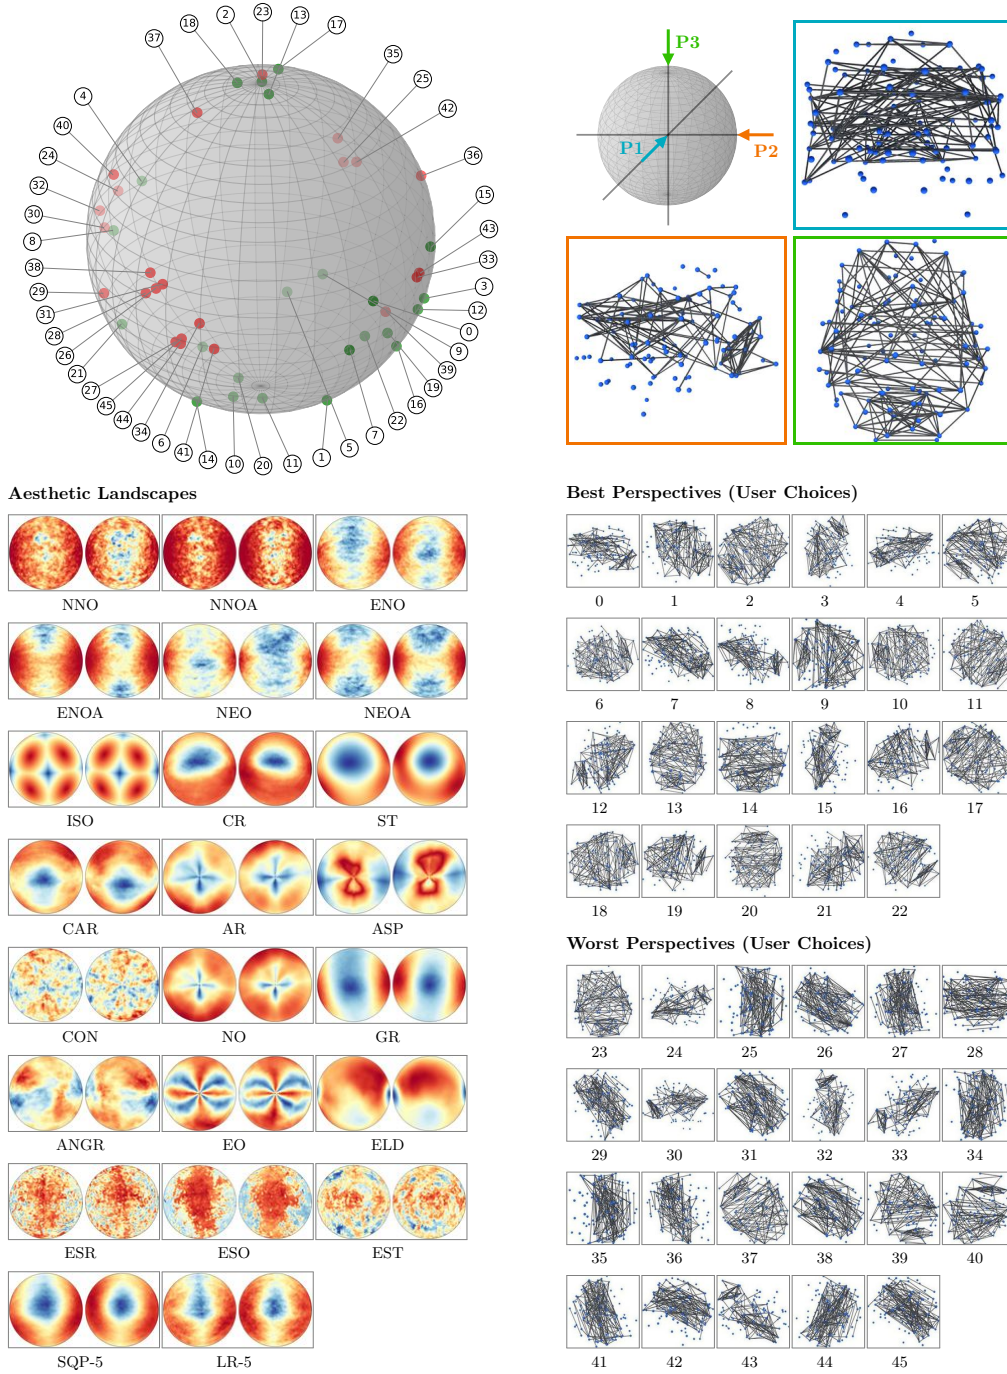


Figure A7: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-6** (sem.2\_100\_144) with  $|V| = 100$ ,  $|E| = 144$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

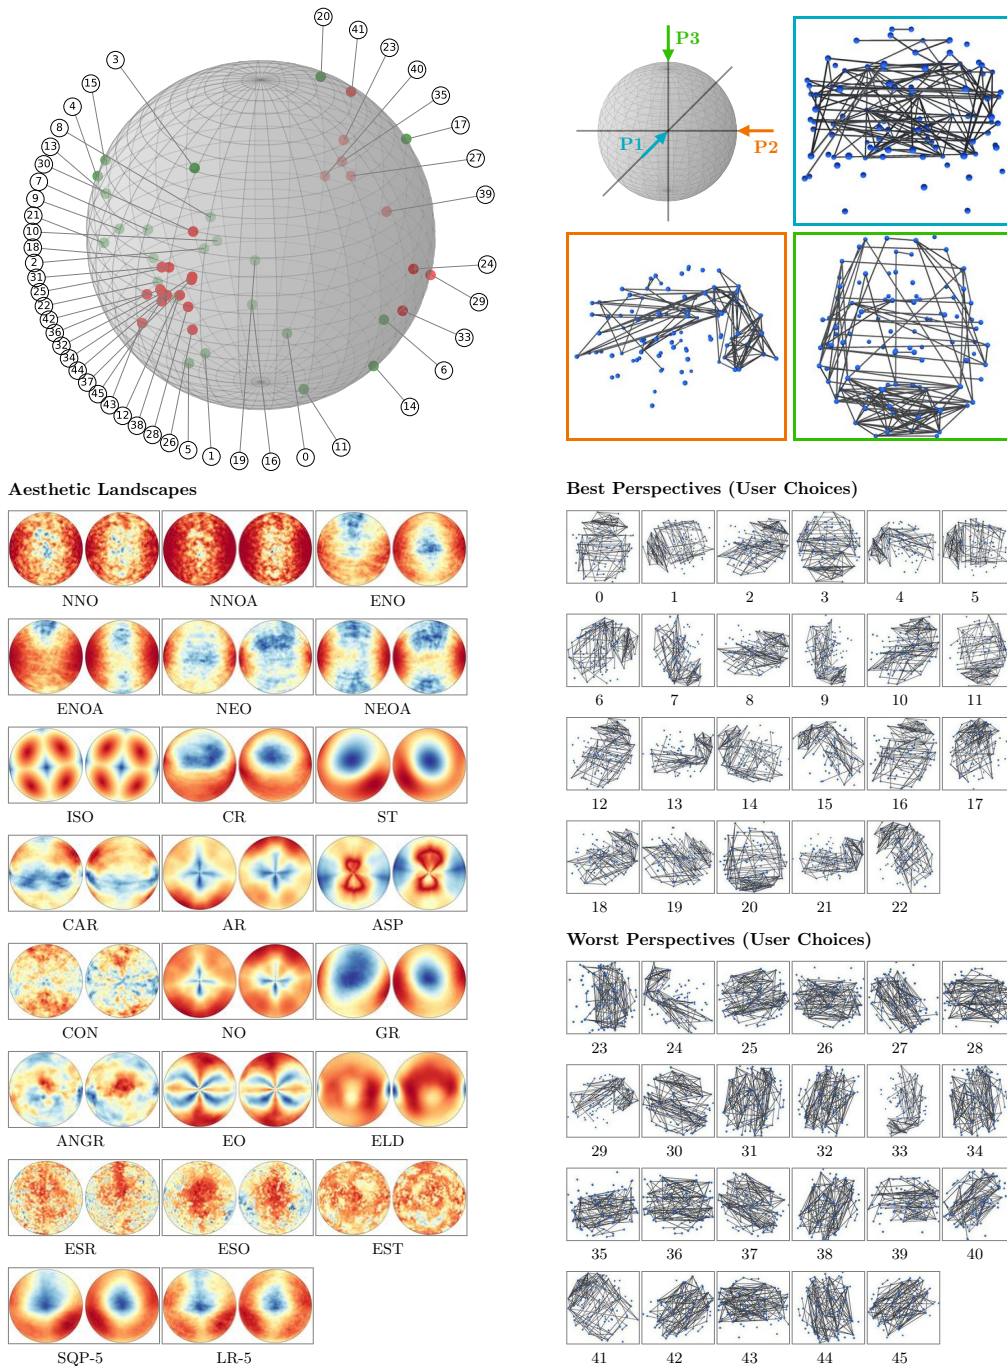


Figure A8: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-7** (sem.3\_100\_137) with  $|V| = 100$ ,  $|E| = 137$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

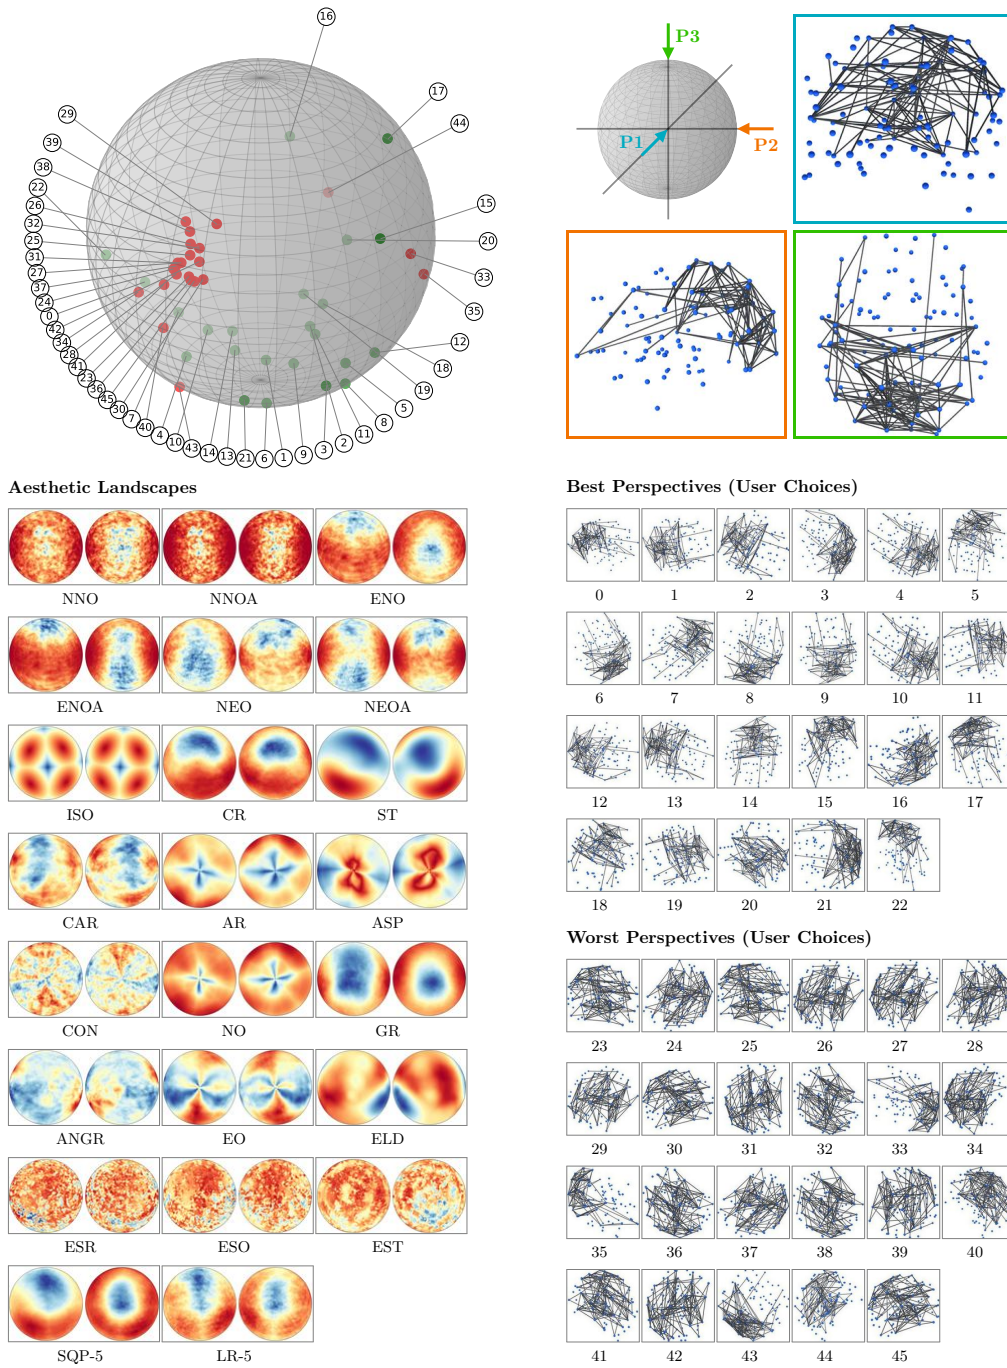


Figure A9: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-8** (sem.9\_100\_131) with  $|V| = 100$ ,  $|E| = 131$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

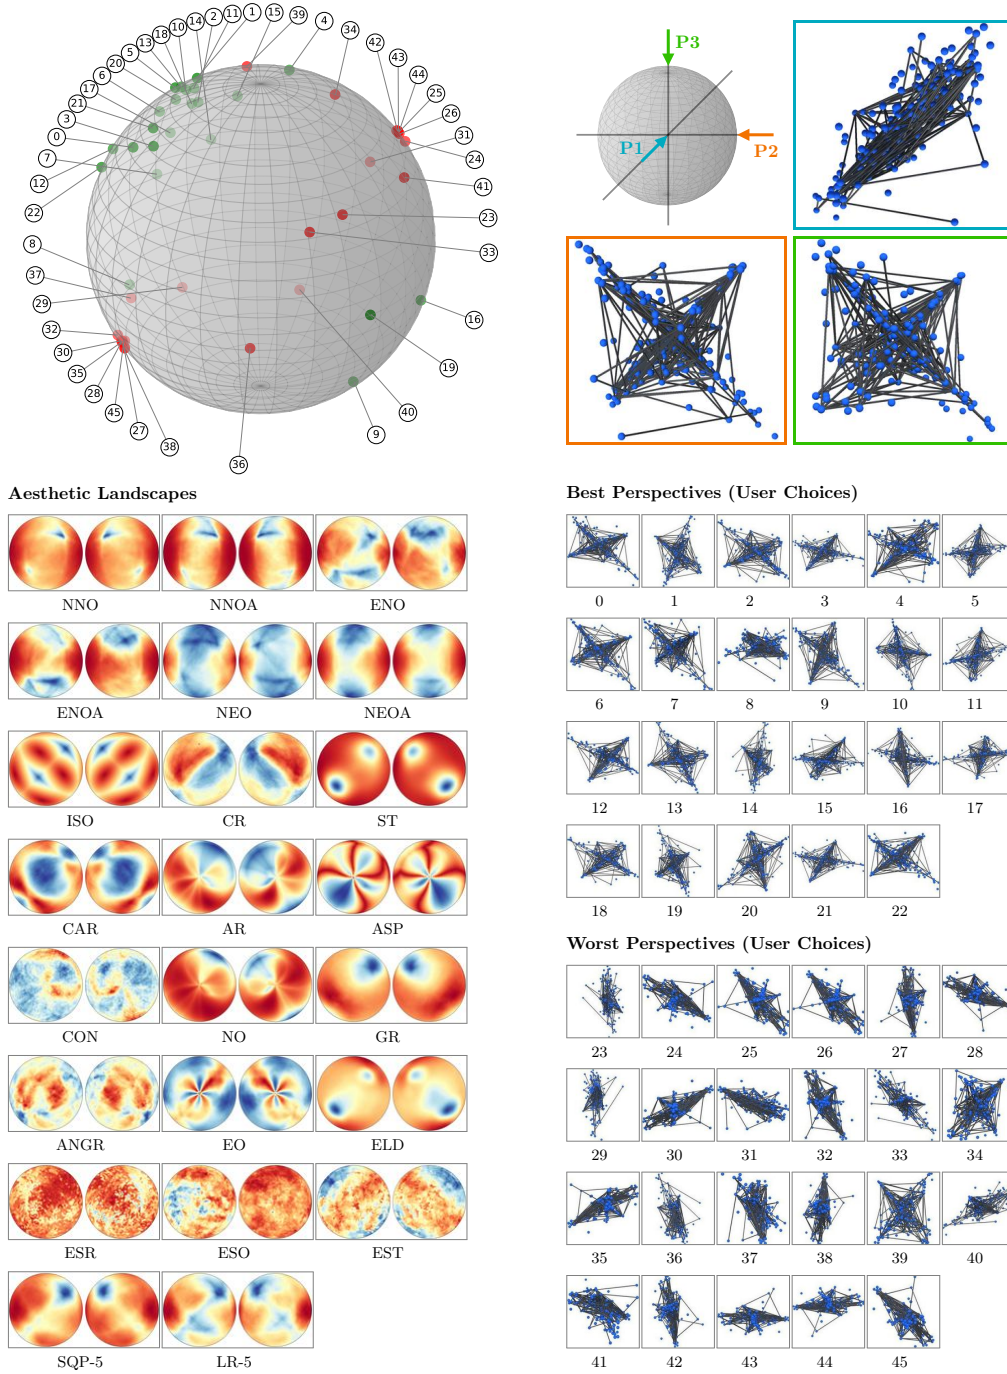


Figure A10: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-9** (sem.5\_200\_194) with  $|V| = 200$ ,  $|E| = 194$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

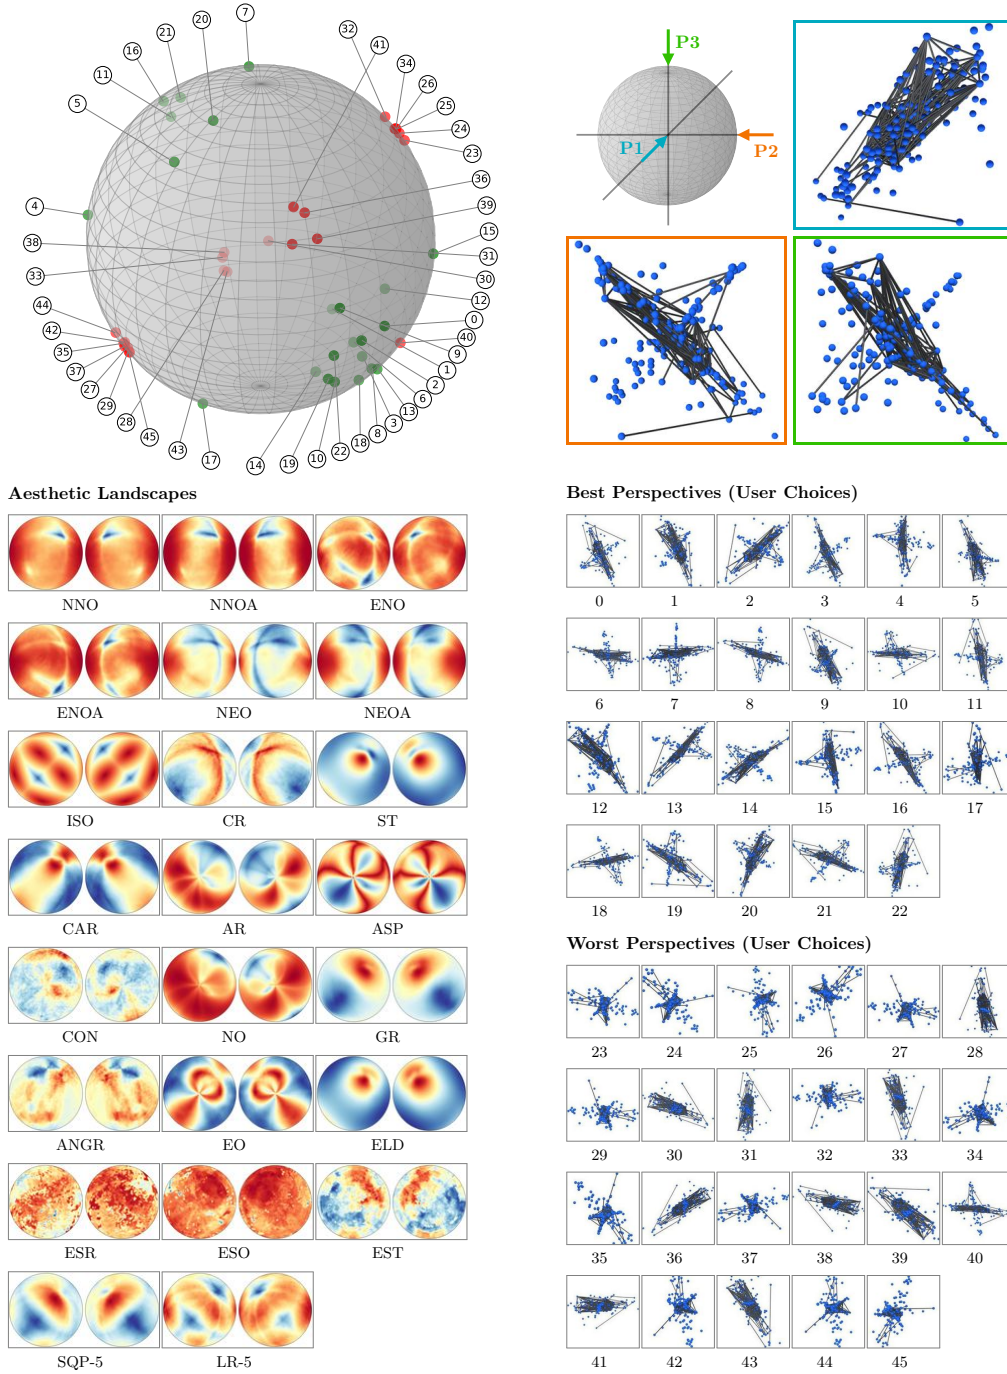


Figure A11: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-10** (sem\_6.200.202) with  $|V| = 200$ ,  $|E| = 202$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

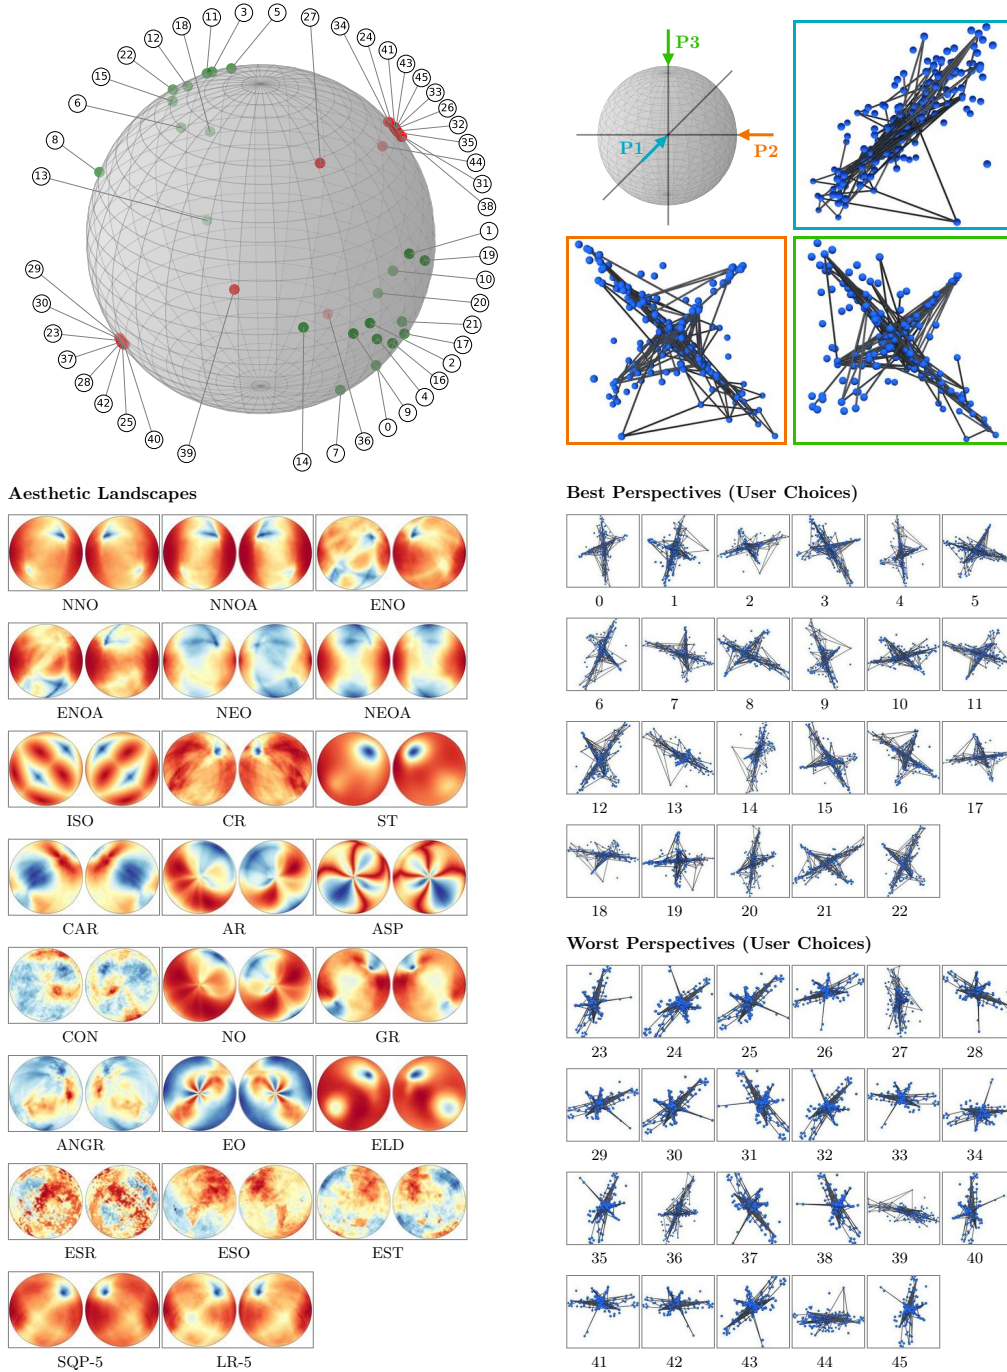


Figure A12: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **S-11** (sem\_9.200.150) with  $|V| = 200$ ,  $|E| = 150$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

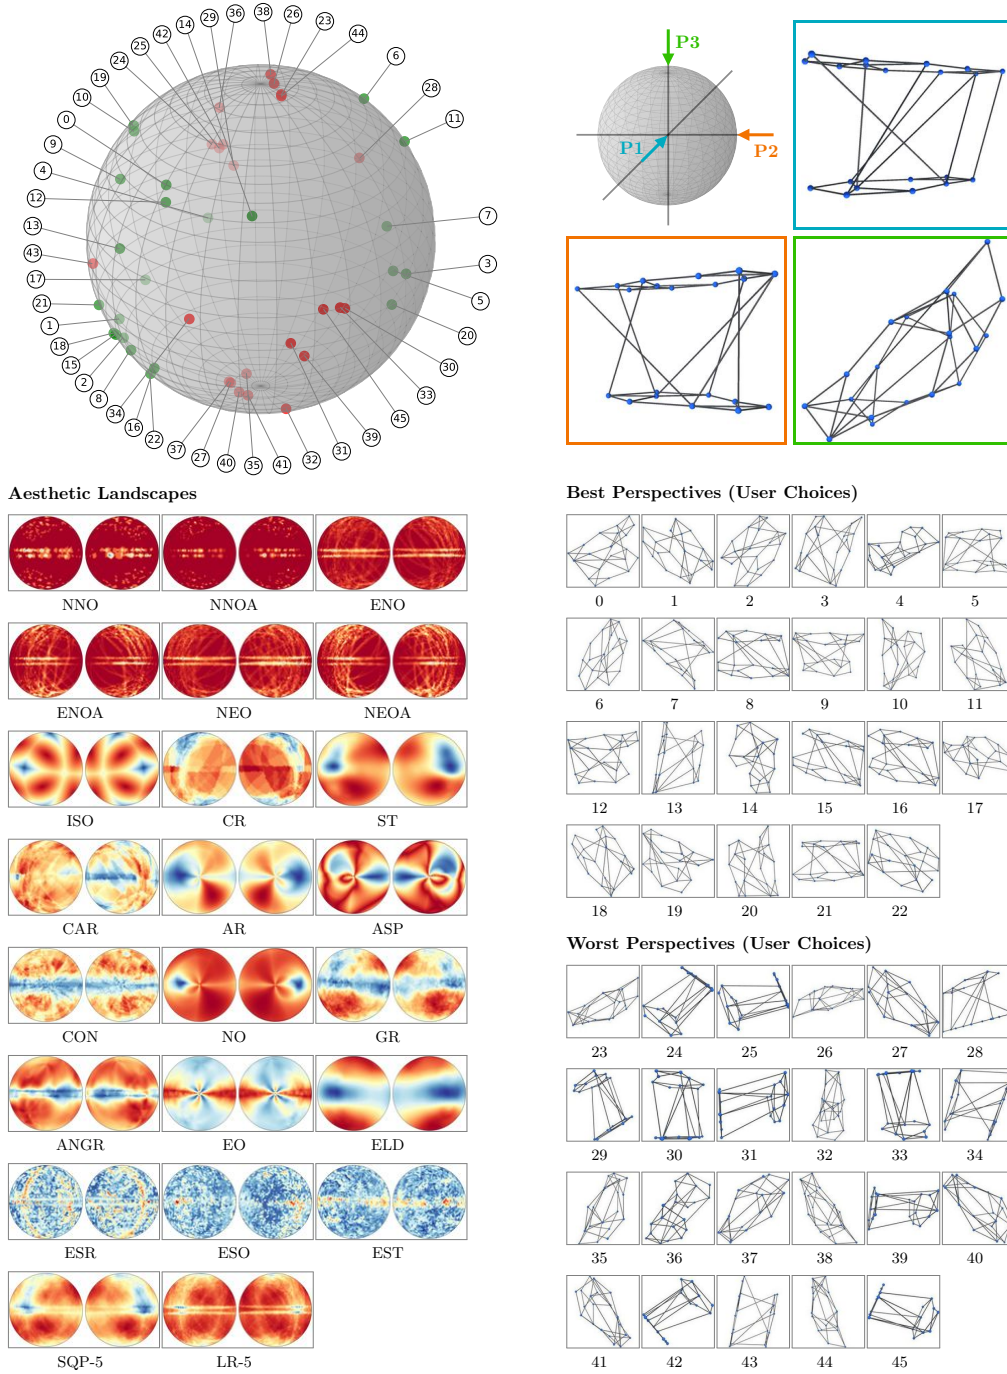


Figure A13: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-0** (layered\_2.20.36) with  $|V| = 20$ ,  $|E| = 36$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

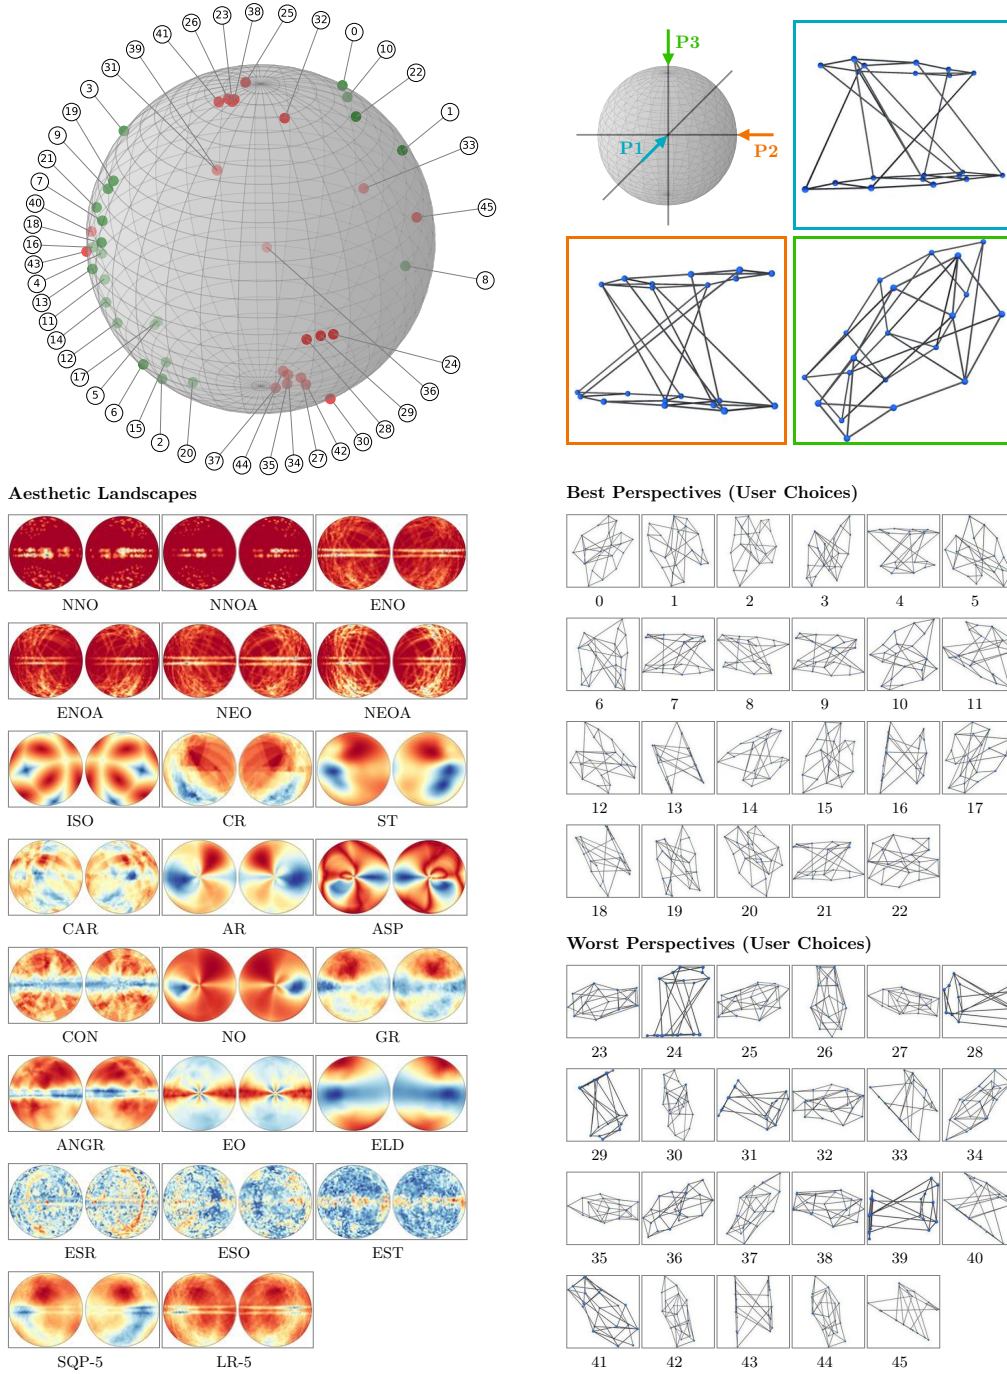


Figure A14: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-1** (layered\_2.20.38) with  $|V| = 20$ ,  $|E| = 38$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

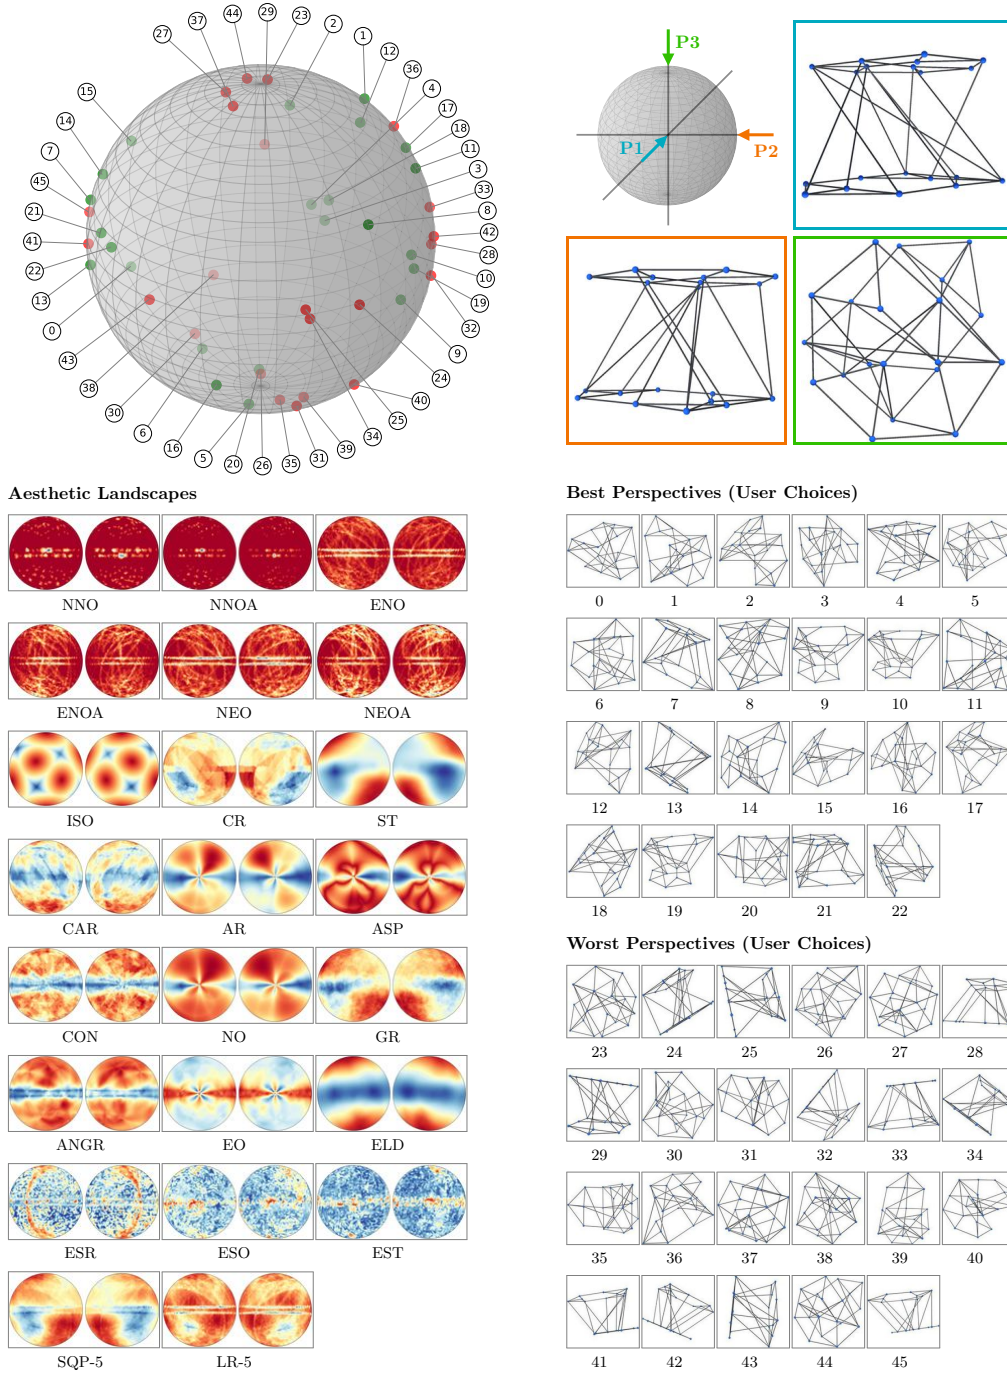


Figure A15: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-2** (layered\_2\_20\_40) with  $|V| = 20$ ,  $|E| = 40$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

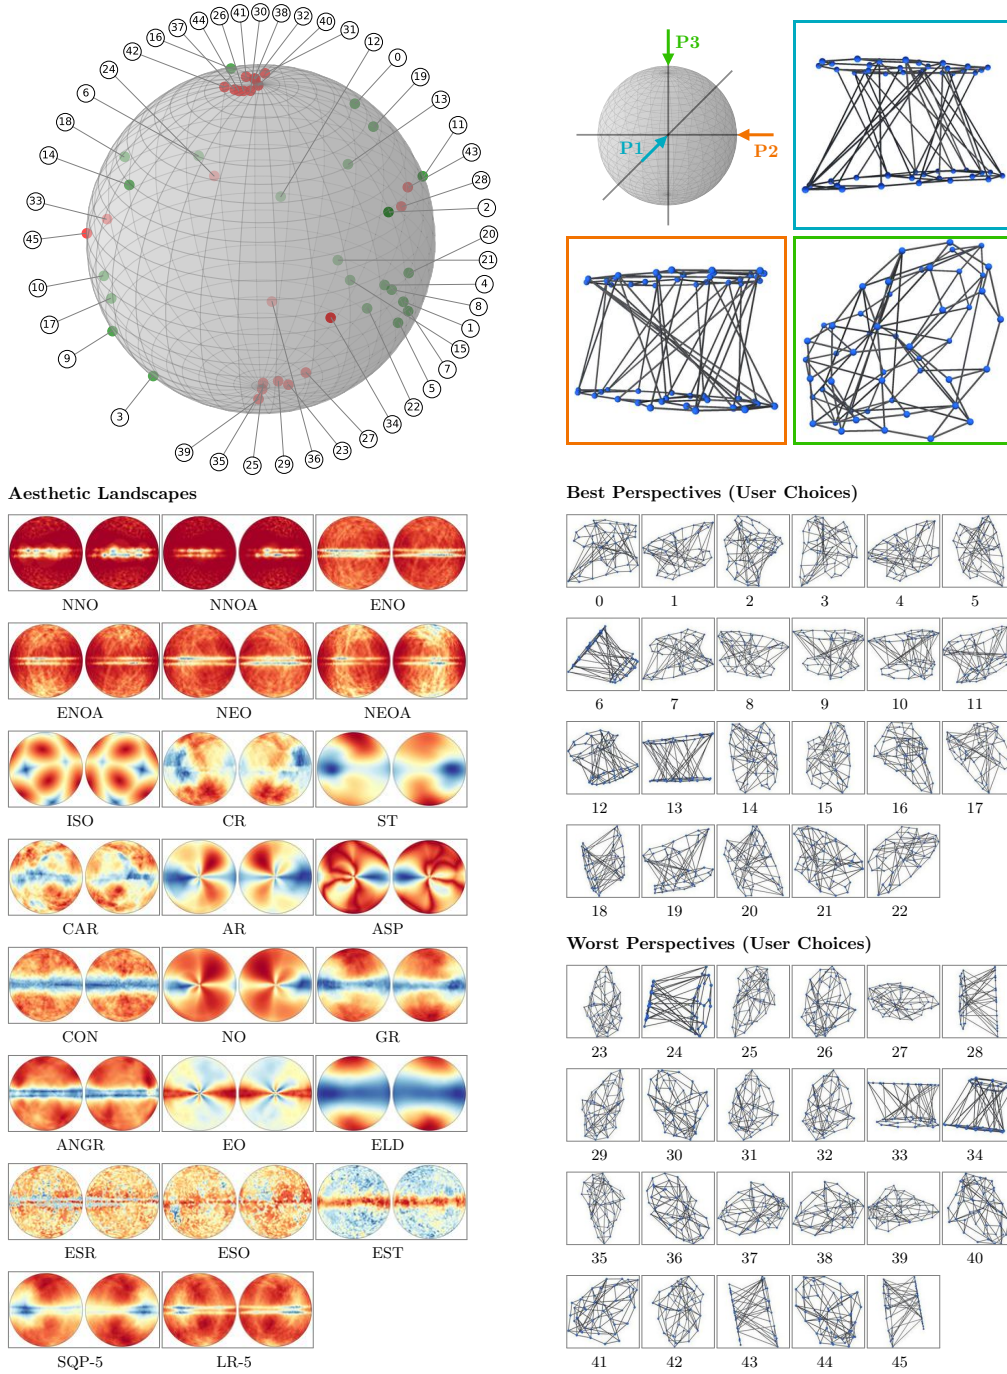


Figure A16: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-3** (layered\_2.50\_101) with  $|V| = 50$ ,  $|E| = 101$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

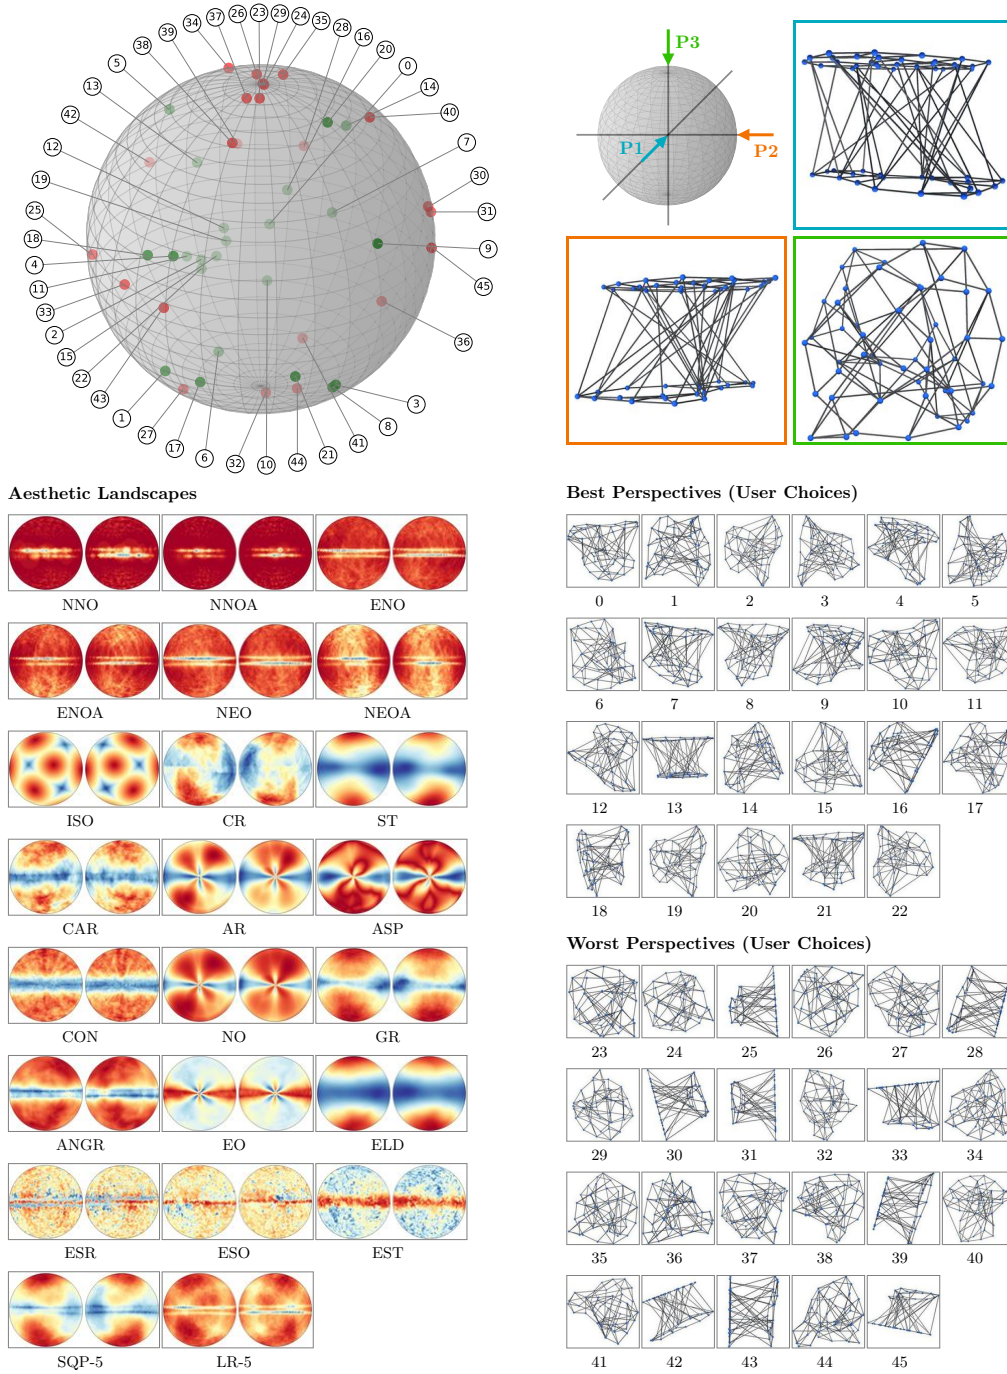


Figure A17: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-4** (layered\_2.50\_102) with  $|V| = 50$ ,  $|E| = 102$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

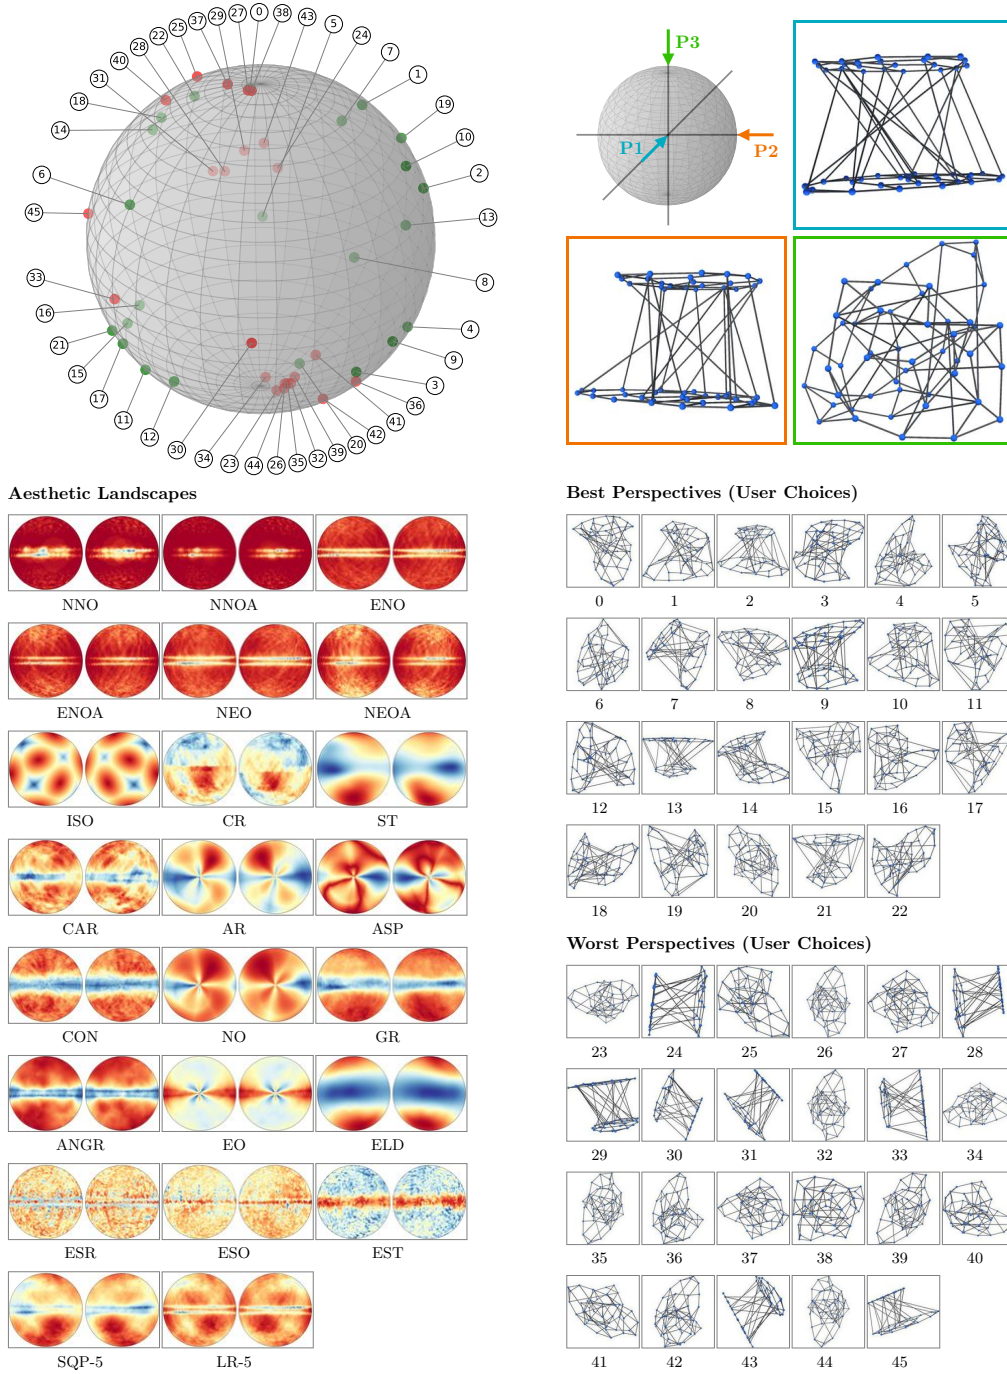


Figure A18: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-5** (layered\_2.50\_93) with  $|V| = 50$ ,  $|E| = 93$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

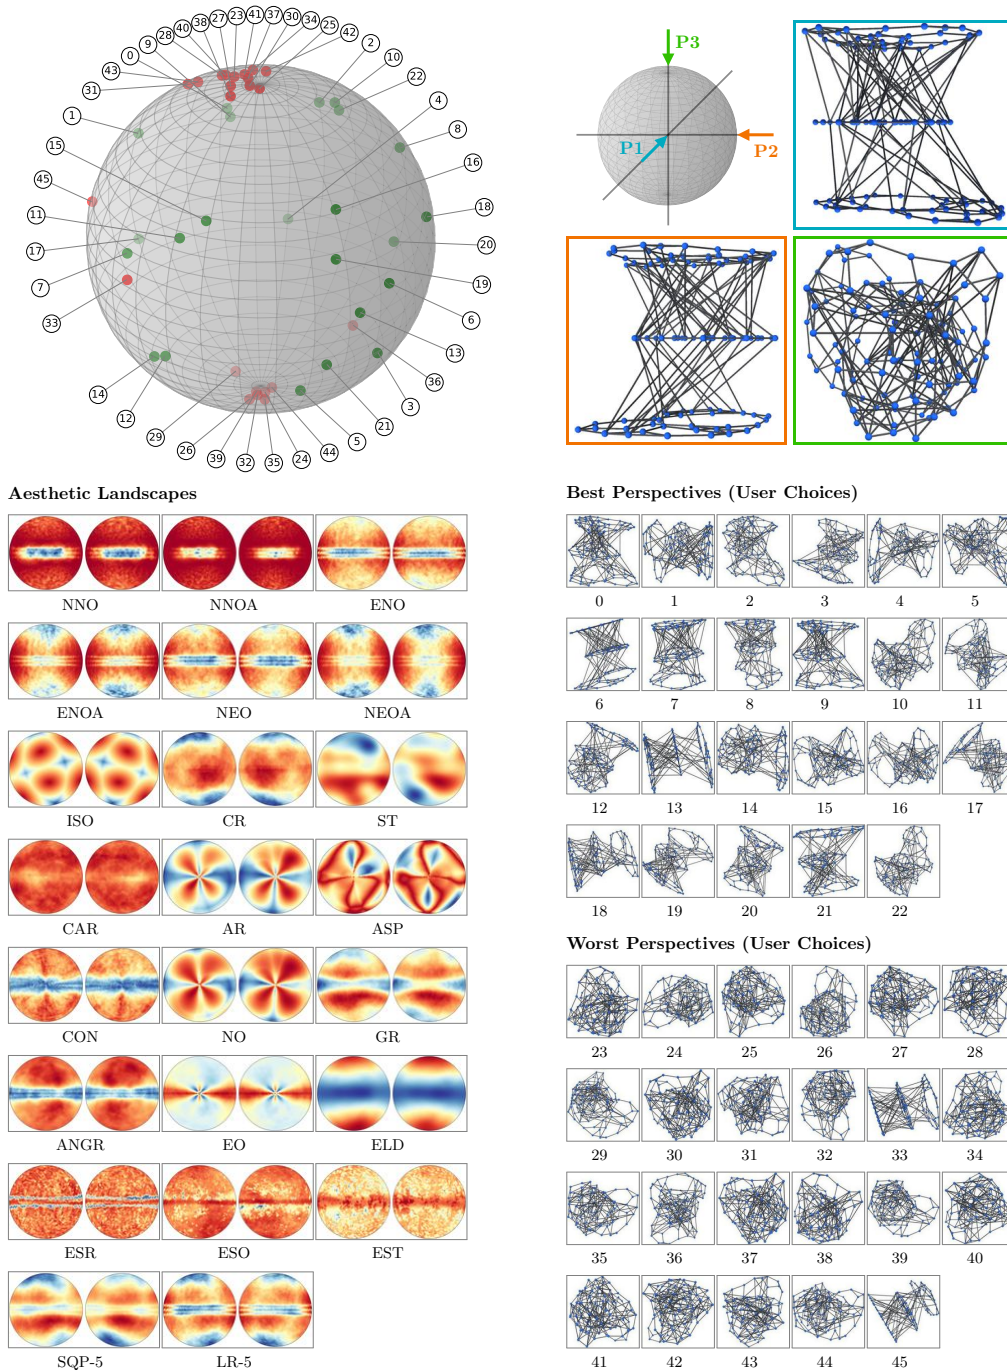


Figure A19: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-6** (layered.3\_100\_207) with  $|V| = 100$ ,  $|E| = 207$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

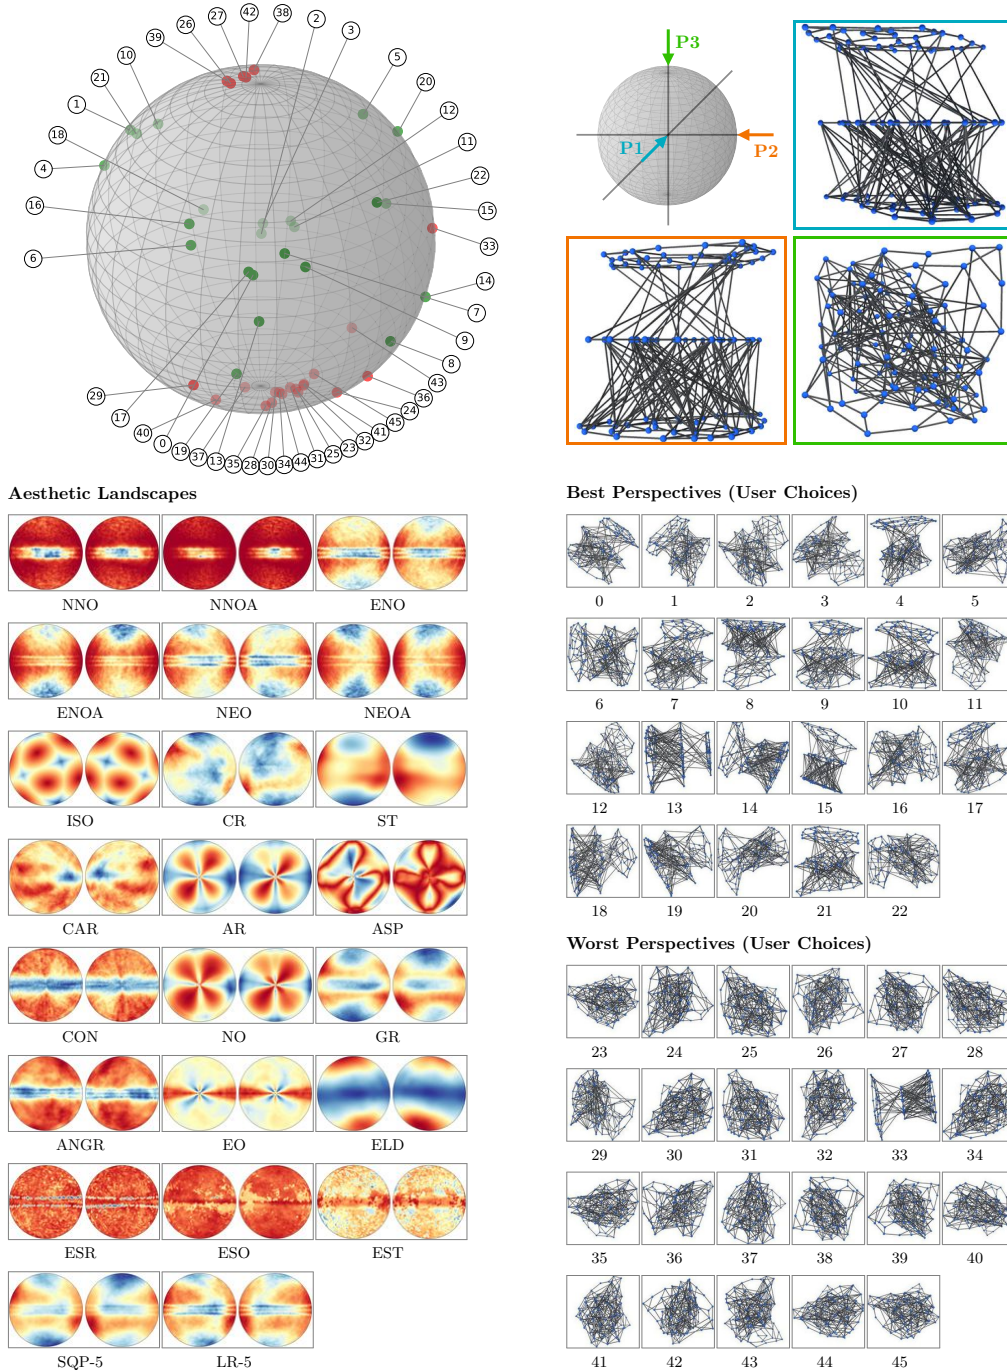


Figure A20: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-7** (layered.3\_100\_253) with  $|V| = 100$ ,  $|E| = 253$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

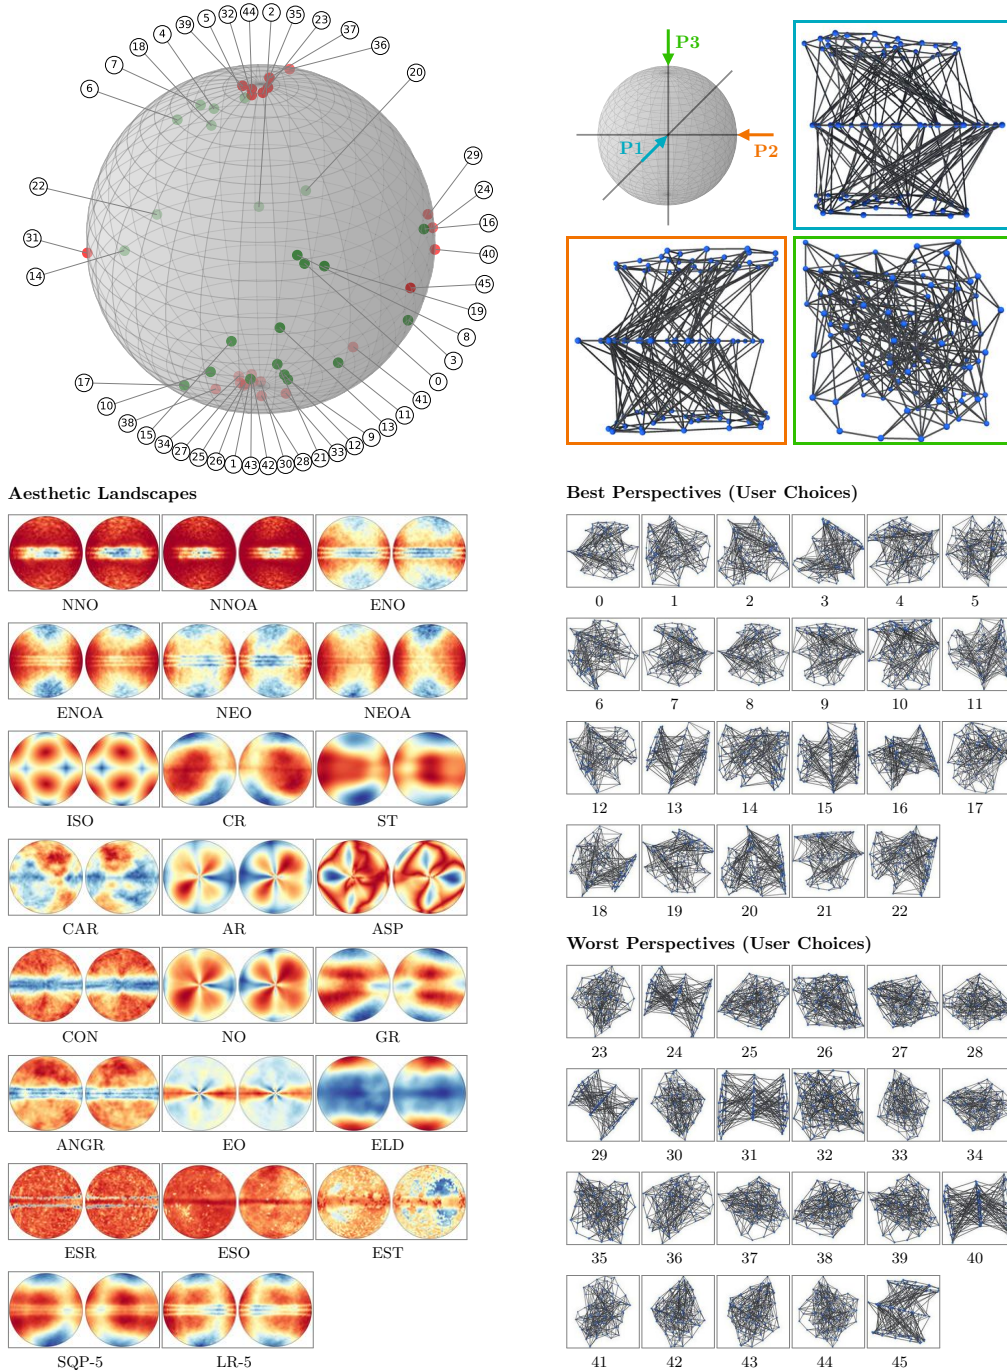


Figure A21: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-8** (layered.3\_100\_284) with  $|V| = 100$ ,  $|E| = 284$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

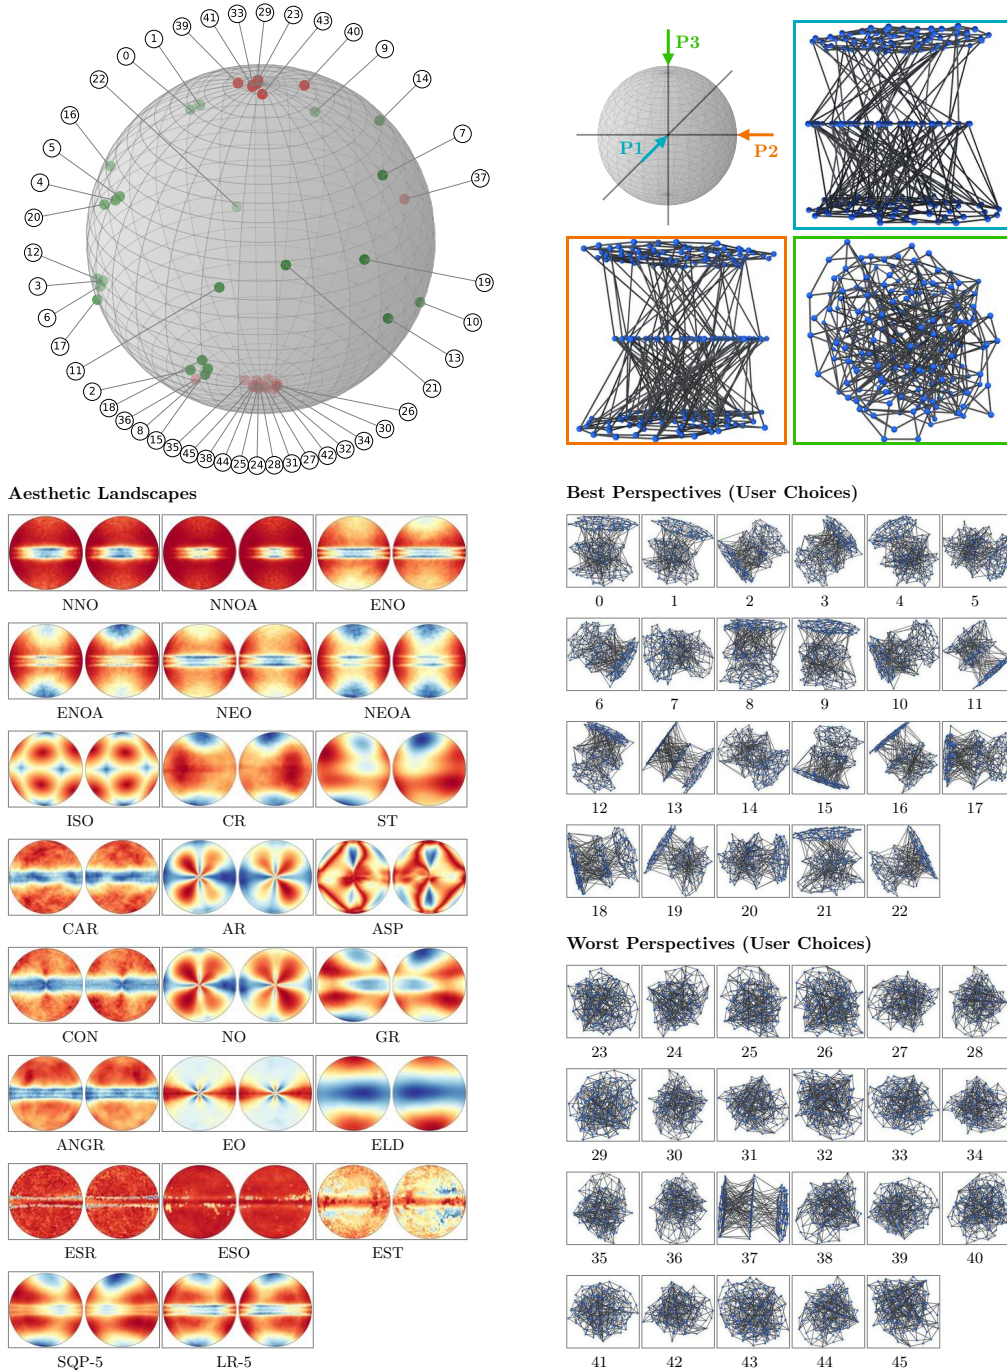


Figure A22: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-9** (layered.3\_200\_431) with  $|V| = 200$ ,  $|E| = 431$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

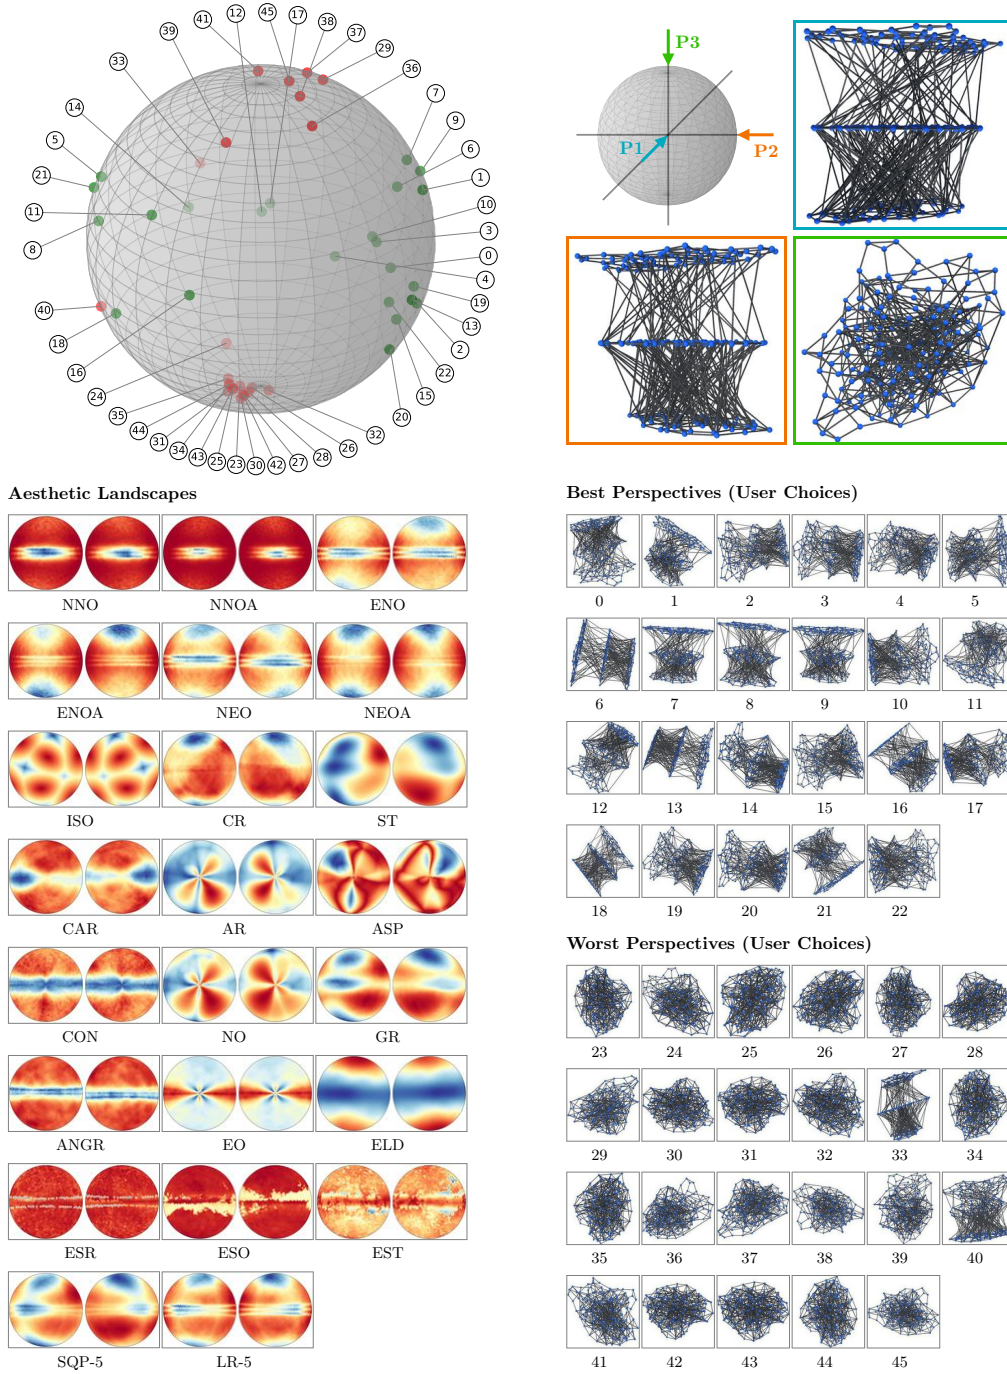


Figure A23: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-10** (layered.3.200.487) with  $|V| = 200$ ,  $|E| = 487$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

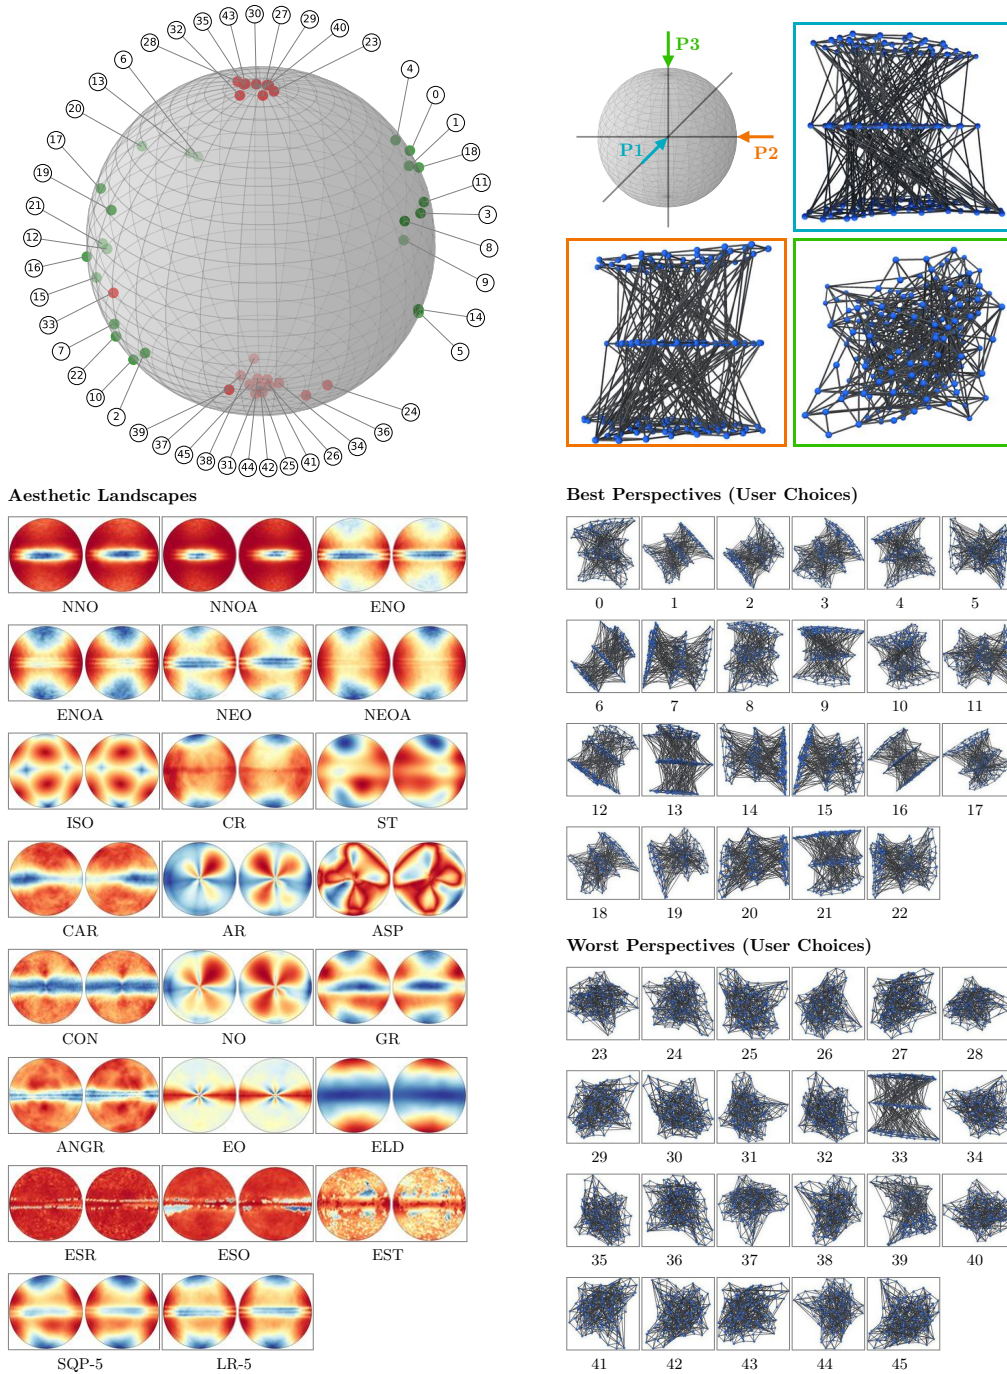


Figure A24: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **L-11** (layered\_3\_200\_508) with  $|V| = 200$ ,  $|E| = 508$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

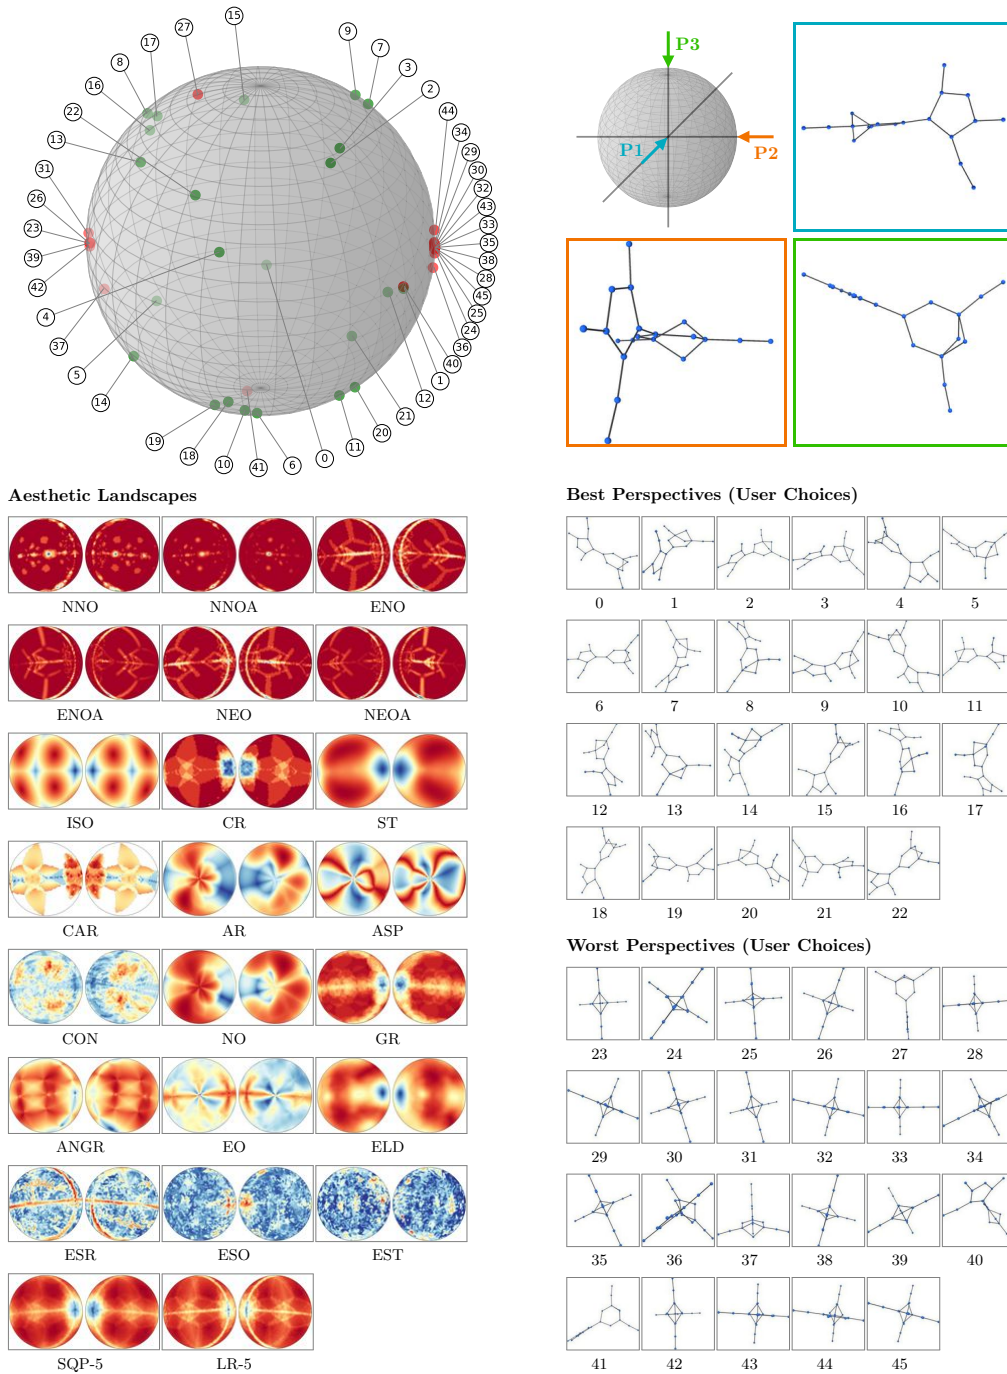


Figure A25: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-0** (grafo2347.20) with  $|V| = 20$ ,  $|E| = 22$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

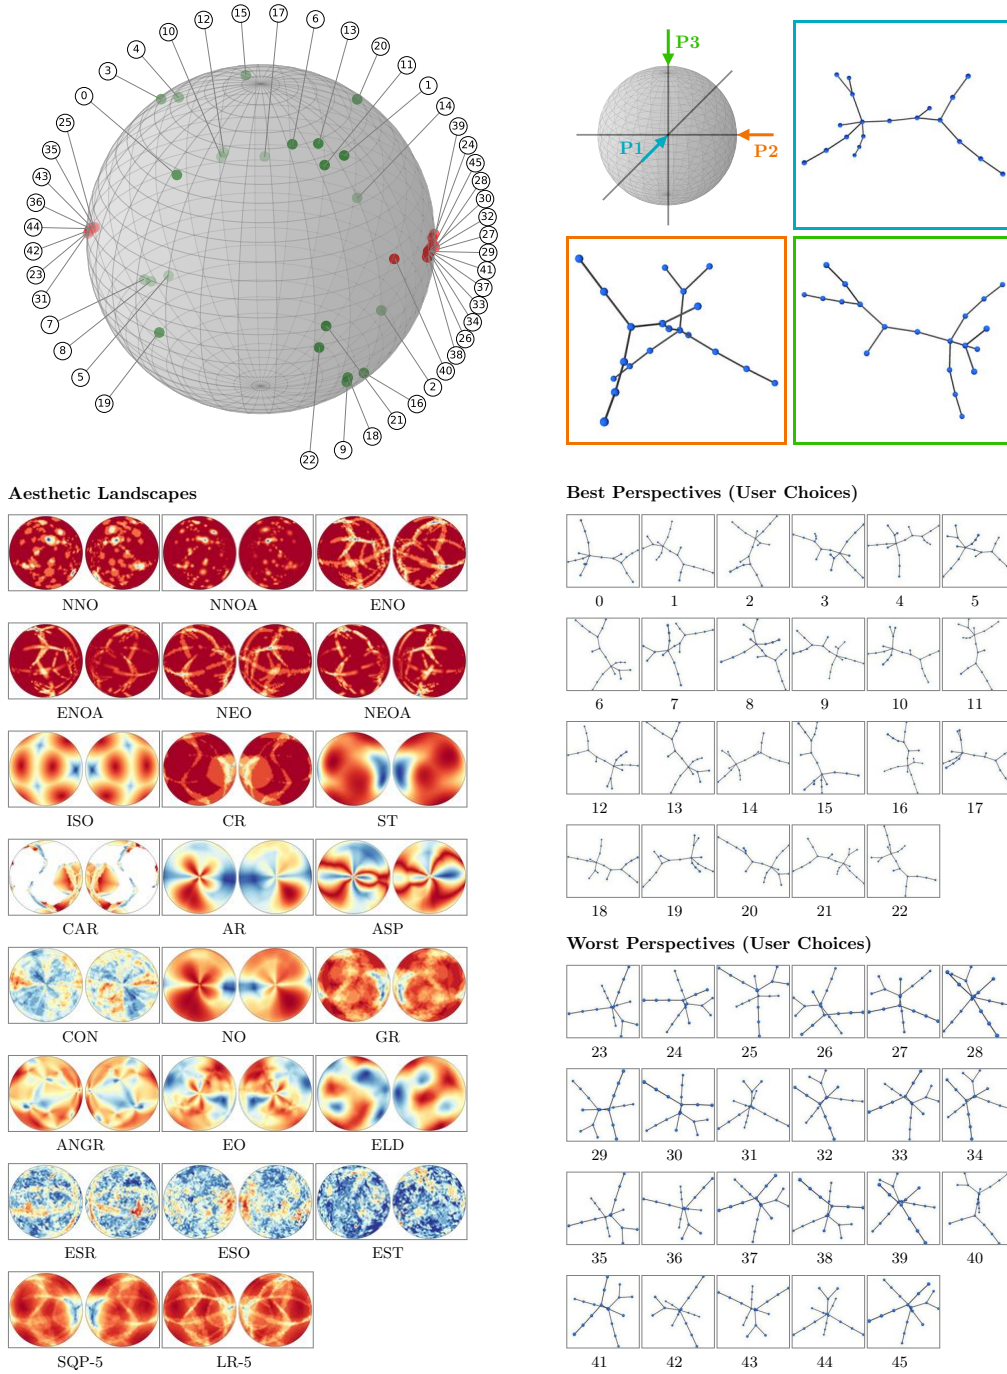


Figure A26: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-1** (tree.20.0) with  $|V| = 20$ ,  $|E| = 19$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

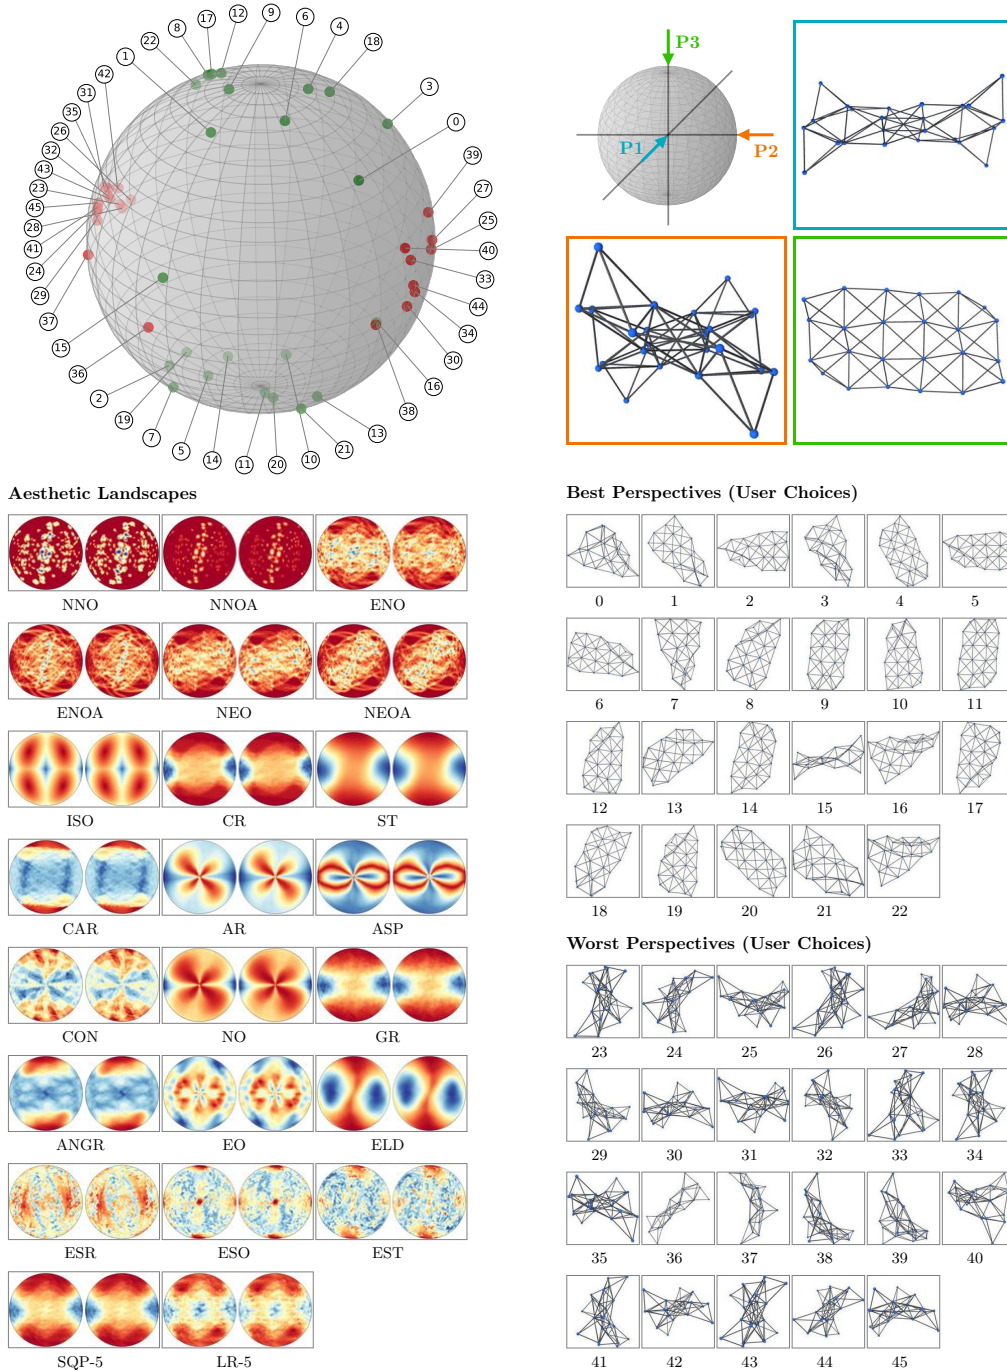


Figure A27: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-2** (can\_24) with  $|V| = 24$ ,  $|E| = 92$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

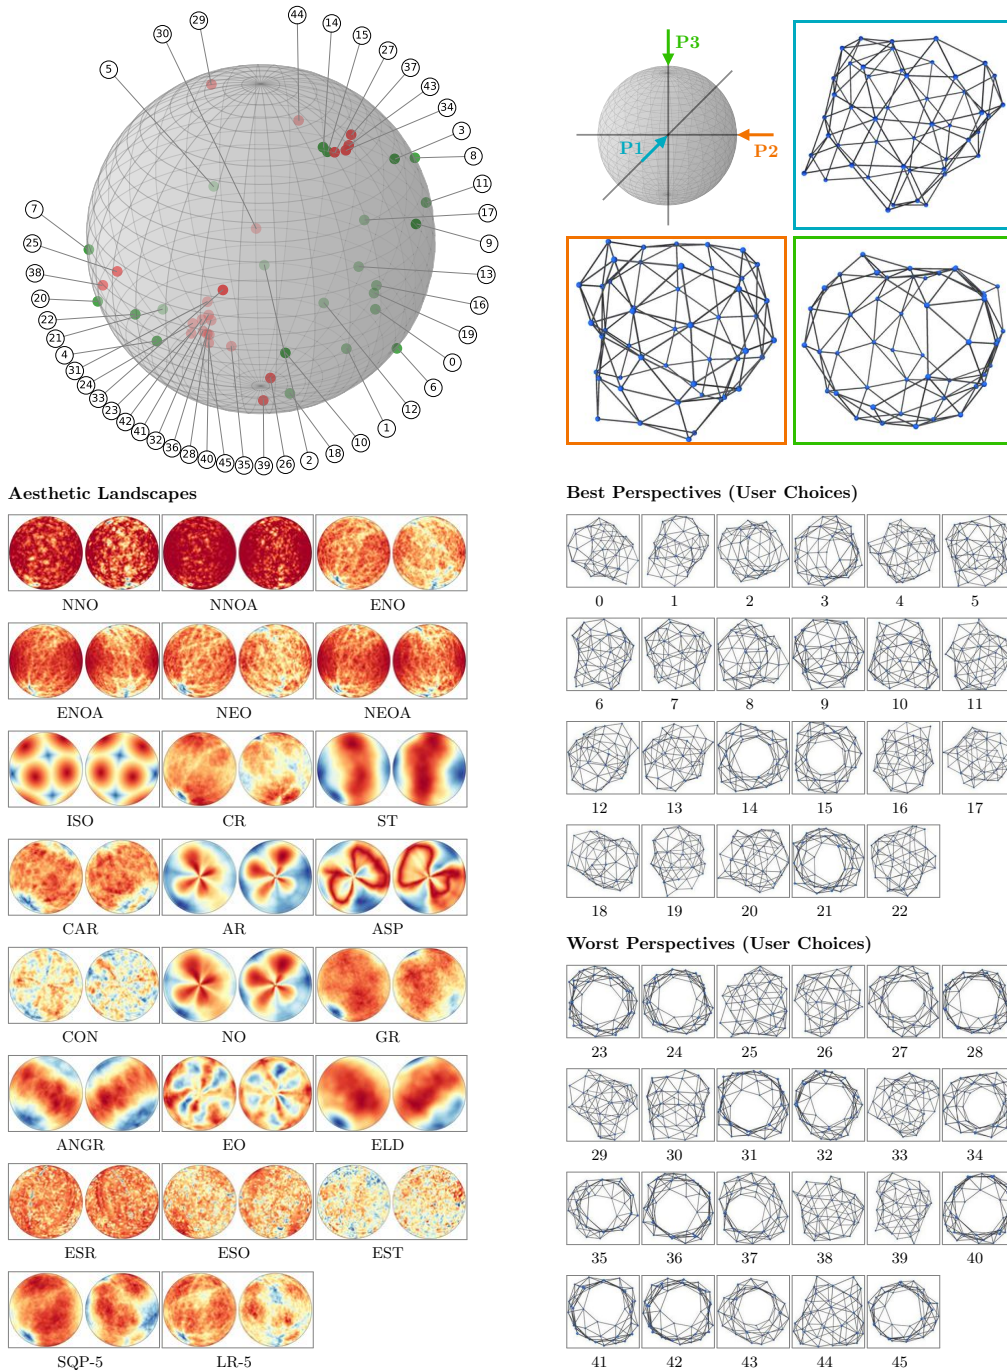


Figure A28: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-3** (mesh1em1) with  $|V| = 48$ ,  $|E| = 177$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

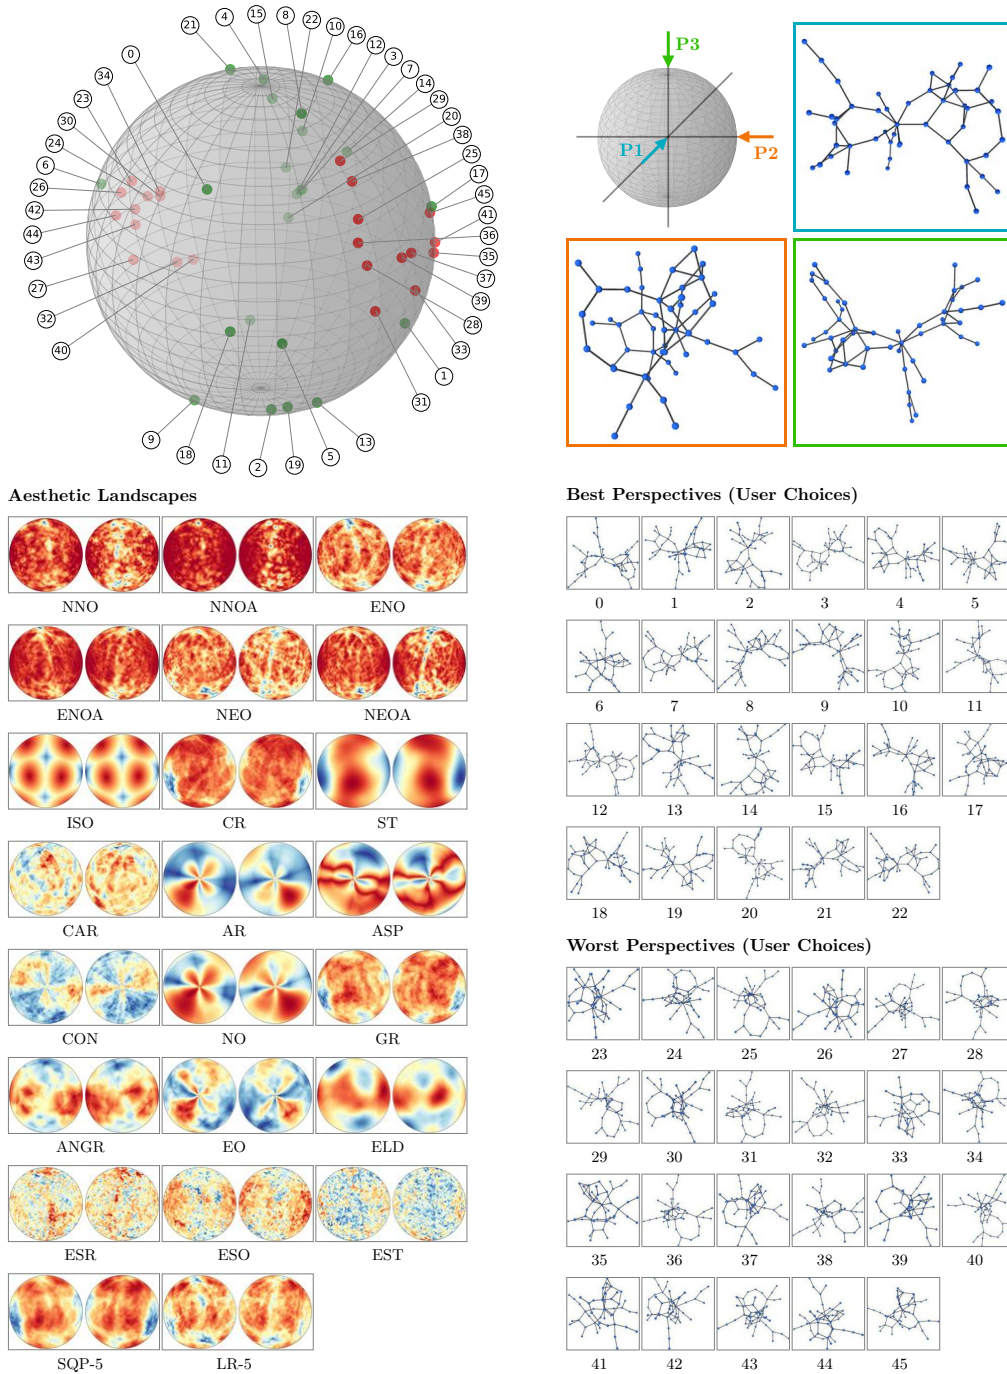


Figure A29: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-4** (grafo5640.50) with  $|V| = 50$ ,  $|E| = 58$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

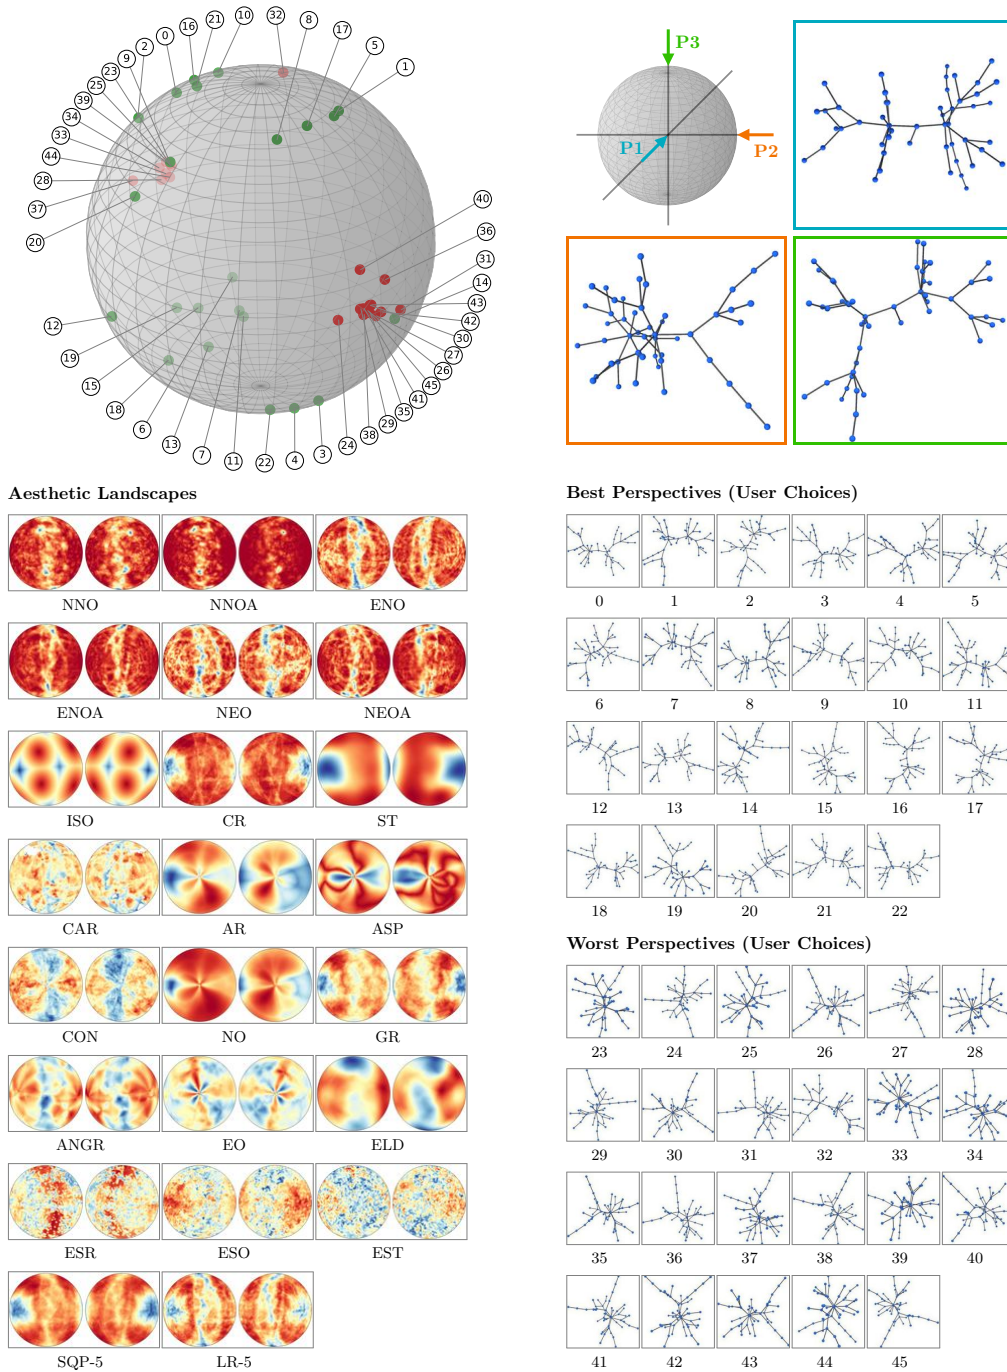


Figure A30: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-5** (tree.50.2) with  $|V| = 50$ ,  $|E| = 49$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

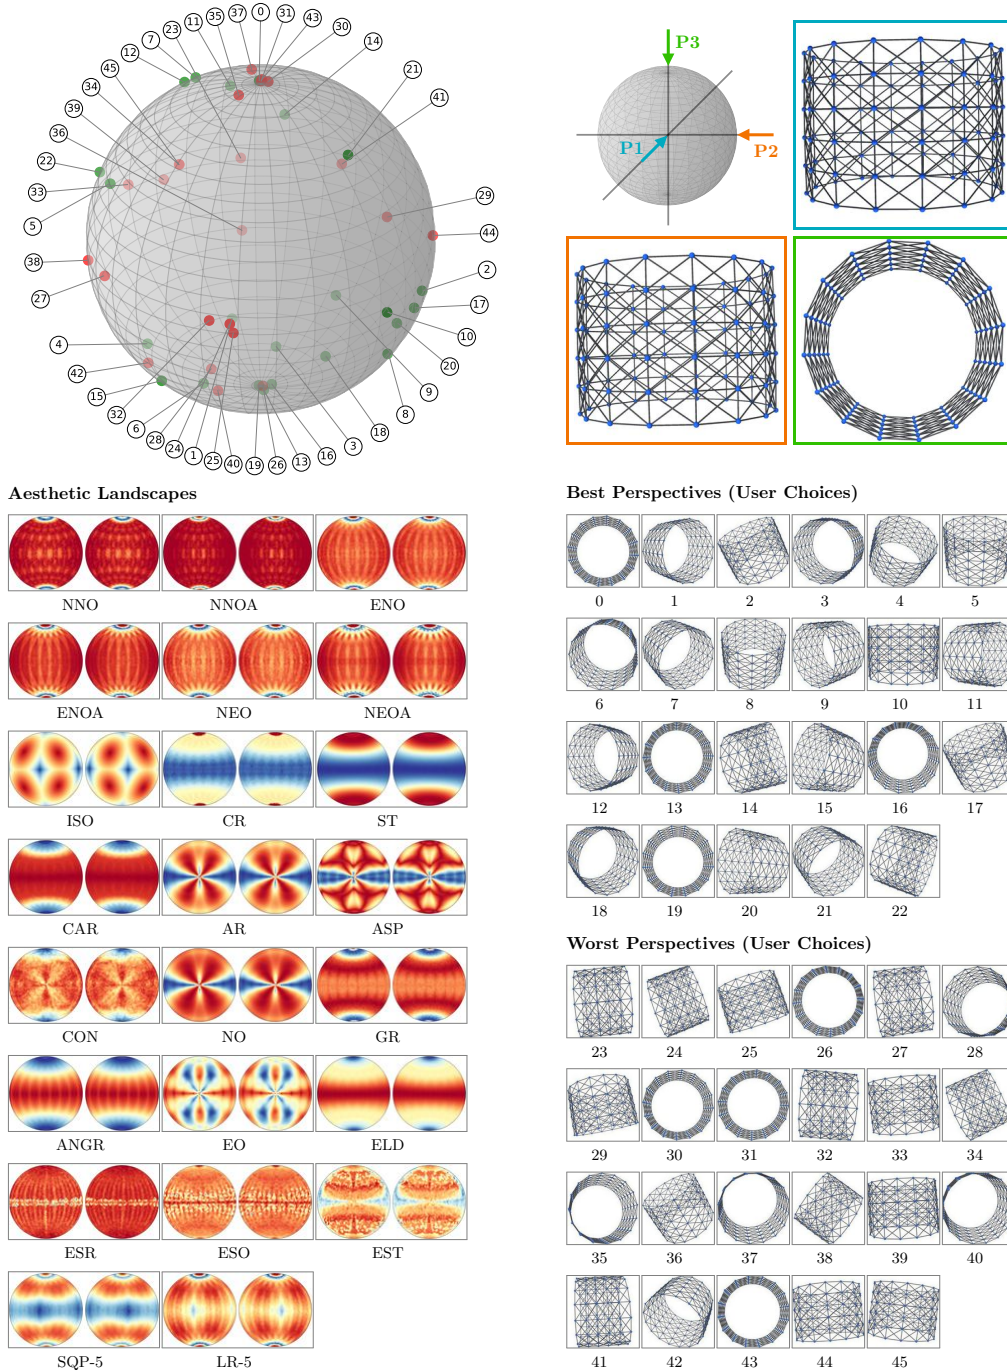


Figure A31: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-6** (can.96) with  $|V| = 96$ ,  $|E| = 432$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

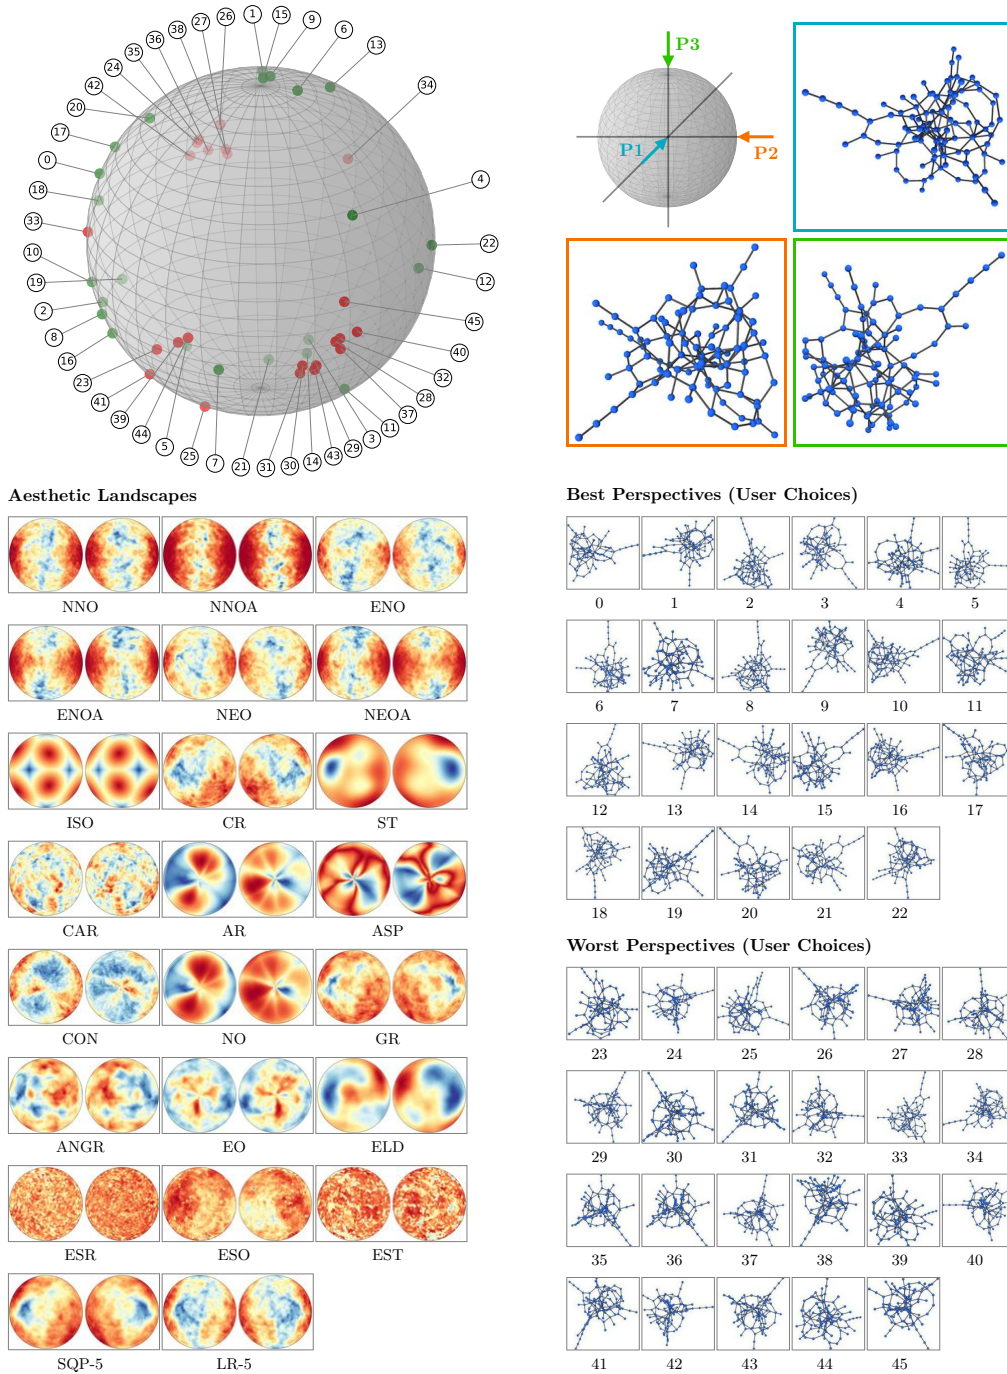


Figure A32: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-7** (grafo10183.100) with  $|V| = 100$ ,  $|E| = 132$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

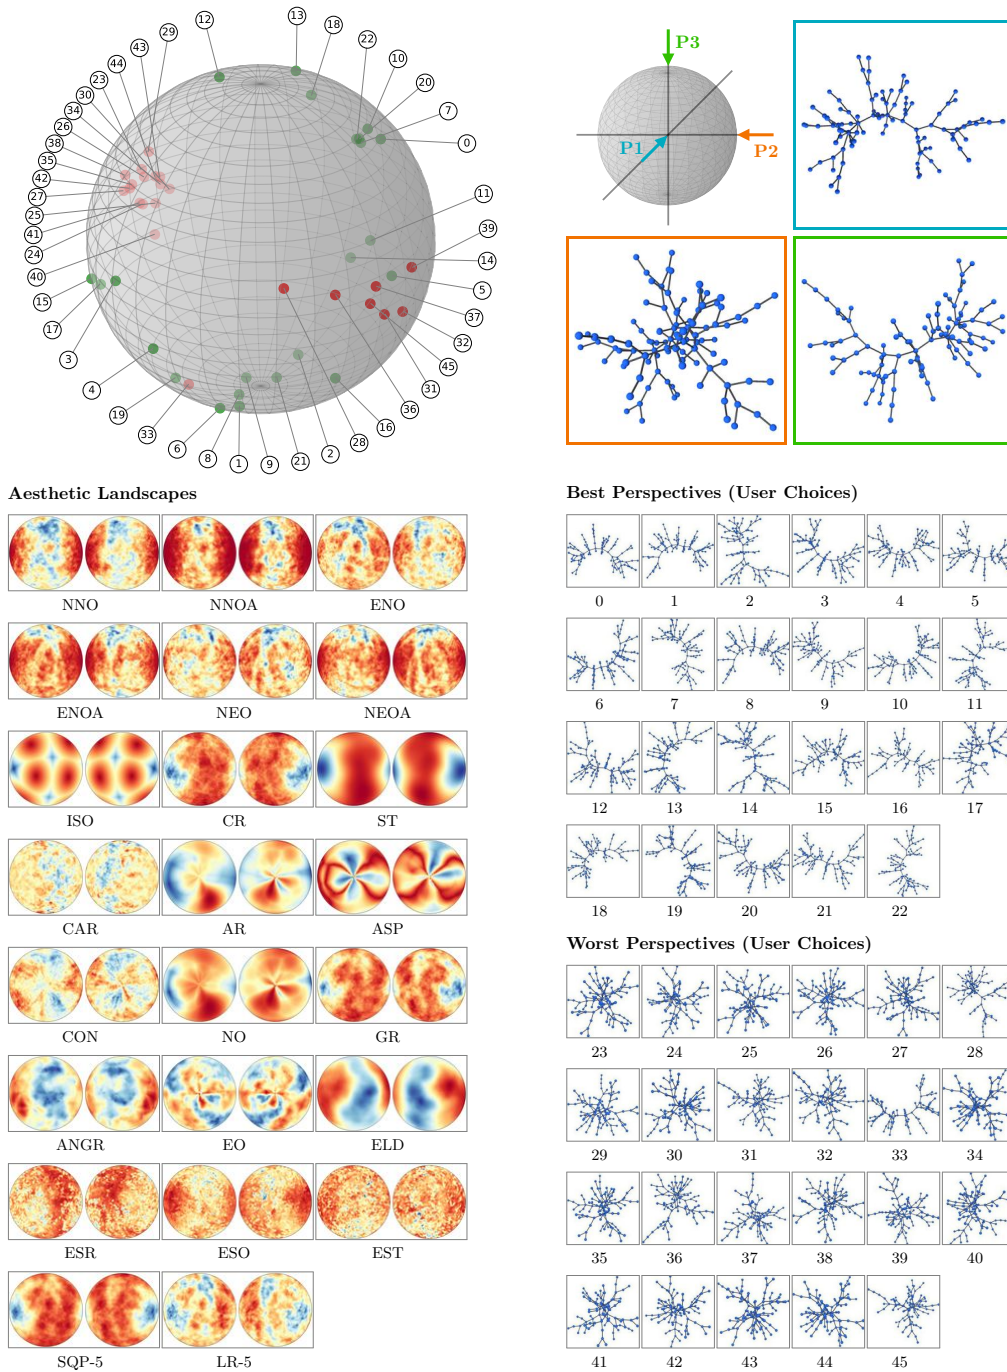


Figure A33: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-8** (tree\_100\_2) with  $|V| = 100$ ,  $|E| = 99$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

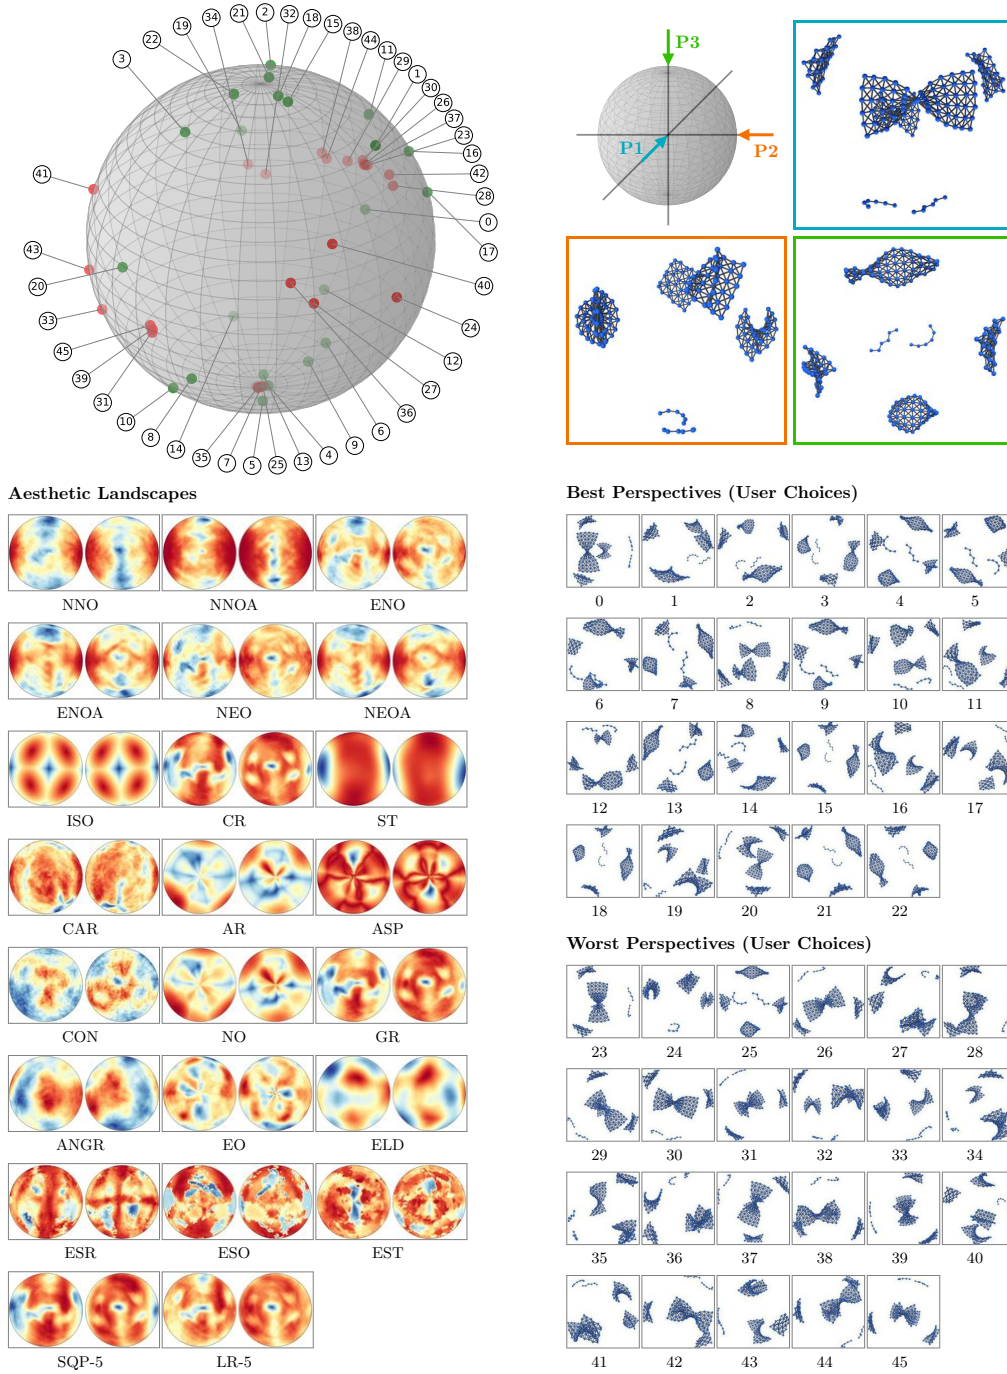


Figure A34: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-9** (dwt\_198) with  $|V| = 198$ ,  $|E| = 795$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

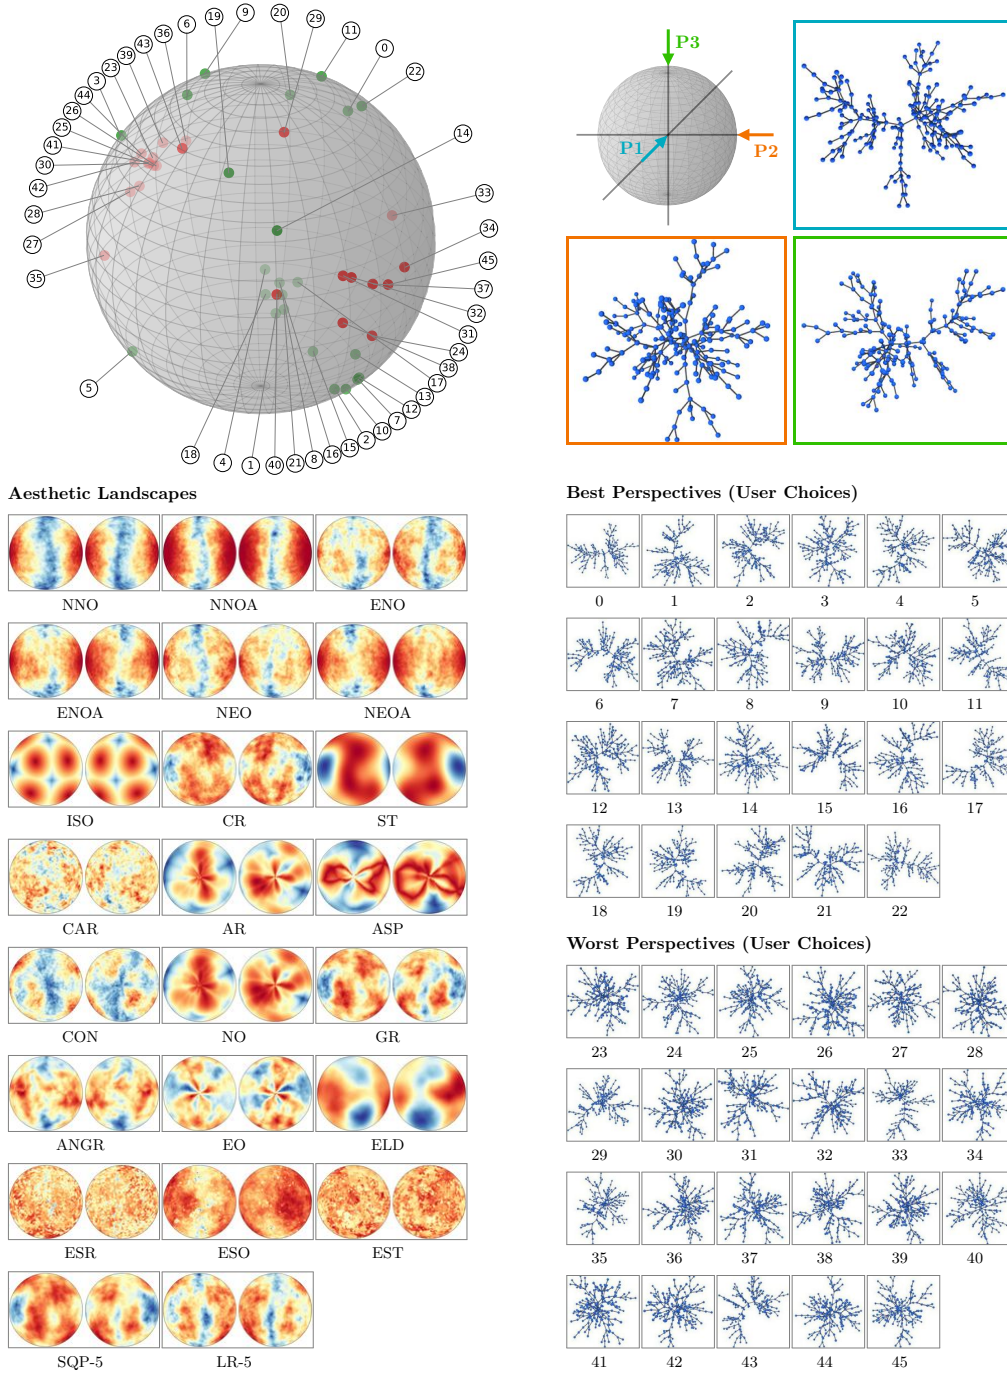


Figure A35: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-10** (tree\_200\_2) with  $|V| = 200$ ,  $|E| = 199$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P1$  (left) and  $-P1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).

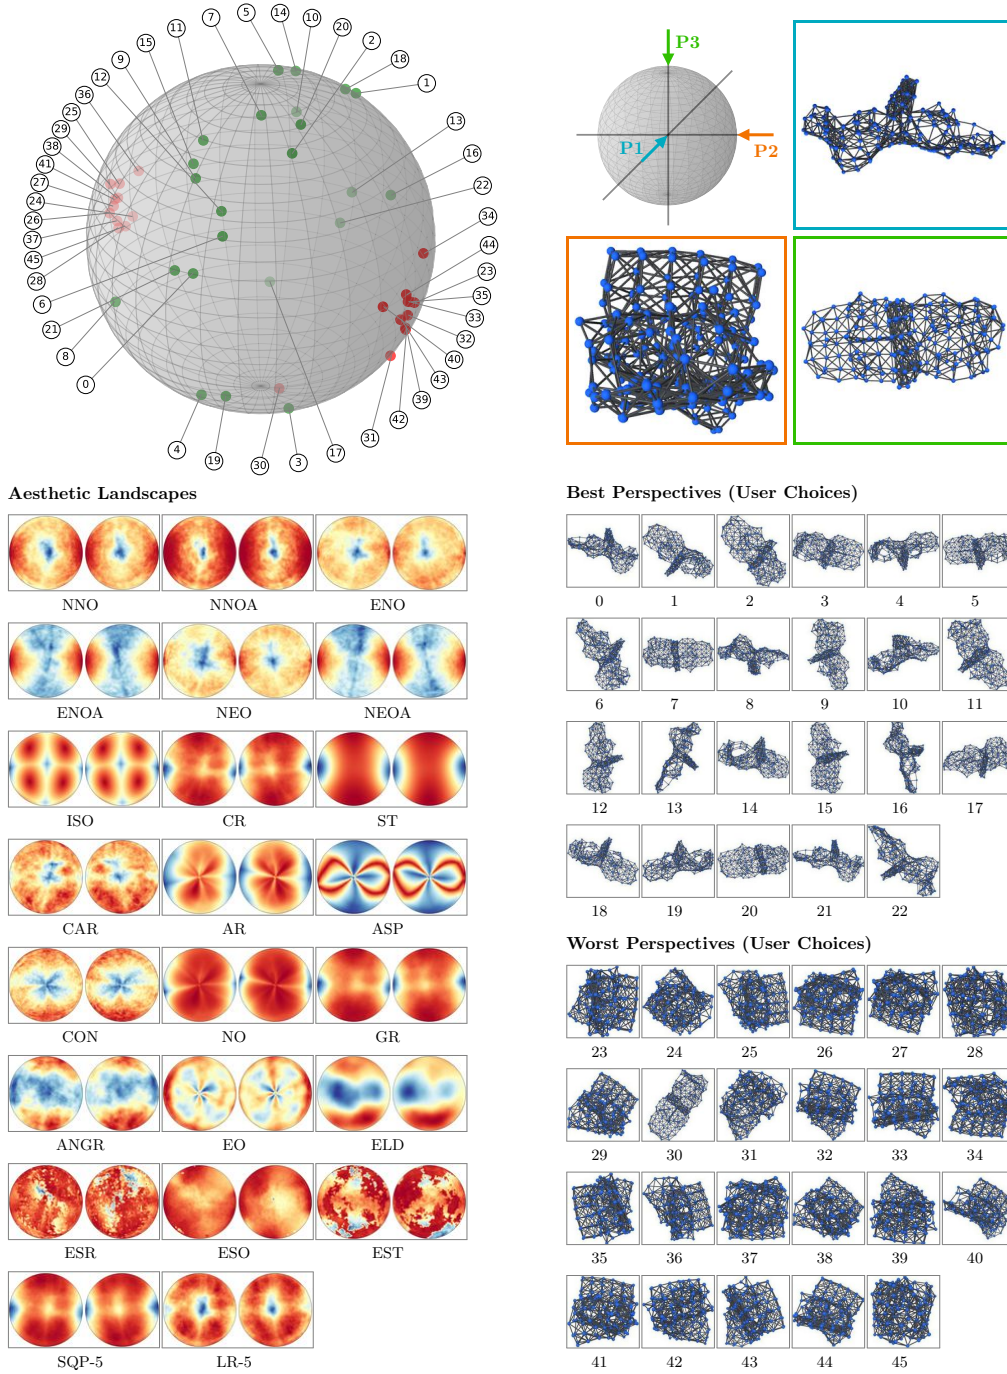


Figure A36: The distribution of all perspectives selected by users as *best* (green) and *worst* (red) mapped on a sphere surface for graph **E-11** (dwt.209) with  $|V| = 209$ ,  $|E| = 976$ . Further, the landscapes of all aesthetic measures and combinations (LR and SQP, restricted to 5 measures) are illustrated as seen from  $P_1$  (left) and  $-P_1$  (right). Colour maps values between 0 and 1 from blue to red (red corresponds to an optimal score).