Vision, Modeling, and Visualization (2025)
B. Egger and T. Günther (Eds.)

Uncovering Relations in High-Dimensional Behavioural Data of Drosophila Melanogaster

Michael Thane[1], Kai M. Blum[1] & Dirk J. Lehmann[1]

[1]Ostfalia University of Applied Sciences, Wolfenbüttel, Germany

Abstract

Understanding how behaviour changes under genetic or experimental conditions is a key challenge in behavioural neuroscience. High-throughput tracking enables the collection of high-dimensional datasets describing locomotion, posture, and stimulus orientation in Drosophila melanogaster larvae (fruit fly). However, exploring relations across numerous dimensions remains challenging. We present a Visual Analytics system that integrates coordinated views, type-aware relation metrics, and hierarchical clustering to support relation discovery and validation in behavioural data. The system was initially developed based on prior experience and refined through evaluation with domain experts to address key analysis tasks, including grouping dimensions, exploring behavioural patterns, and validating hypotheses. We demonstrate how it supports both confirmatory and exploratory workflows, enabling users to confirm known effects and uncover novel patterns—such as an unexpected correlation between head-casting behaviour and locomotion speed. This work highlights how tailored visual analysis can advance behavioural research.

1. Introduction

A central challenge in behavioural neuroscience is understanding how behaviour adapts to neural, genetic, or environmental manipulations [PRPS17, TVM*19, TPS*23]. The fruit fly larva (*Drosophila melanogaster*) is a widely used model for studying how neural activity and memory formation shape behaviour.

Existing practices combine univariate or bivariate plots (boxplots, scatterplots, summary statistics) with heatmaps or correlation matrices for overview. Since these components are not integrated, users must manually connect findings and relation analysis across mixed numerical and categorical data remains poorly supported.

We present a visual analytics (VA) system that integrates coordinated views, type-aware relation metrics, and interactive clustering to support both hypothesis-driven and exploratory analysis of behavioural data. Following a design study methodology [SMM12], we worked closely with behavioural neurobiologists through interviews, task observations and iterative feedback to ensure domain alignment.

The remainder of this paper first presents the domain background and data (*Section* 2) to provide context for our design goals, then reviews related work (*Section* 3), describes the Visual Analytics system (*Section* 4), demonstrates use cases and expert feedback (*Sections* 5–6), and concludes with discussion (*Section* 7) and conclusion (*Section* 8).

This paper contributes:

- A task analysis of relation exploration in high-dimensional behavioural data.
- A VA system that unifies type-aware metrics, interactive clustering, and coordinated views.
- A qualitative expert evaluation showing support for both hypothesis validation and novel pattern discovery, including an unexpected correlation between head-casting and locomotion speed.

2. Domain Background and Requirements

According to Munzner's nested model of visualization design [Mun09], effective systems should address both high-level domain goals—what users ultimately want to understand—and lower-level tasks that support these goals. In this section, we outline both aspects for behavioural neuroscience, focusing on *Drosophila melanogaster* larvae.

2.1. Higher-level Domain Goals

Behavioural neuroscience seeks to understand how animals adapt their behaviour in response to sensory input, learning, or internal states. The larva of the fruit fly (*Drosophila melanogaster*) is a widely used model, offering genetic accessibility and relatively simple yet quantifiable behaviour [PRPS17, SRP*15, TPS*23].

A common paradigm is associative learning. Larvae are placed on a Petri dish where an odour (e.g., amylacetate, a banana scent component) is paired with a reward such as fructose. After a short exposure, they are transferred to a neutral Petri dish, and this sequence is repeated several times. In the subsequent test, larvae are

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.
This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

DOI: 10.2312/vmv.20251234

delivered by EUROGRAPHICS DIGITAL LIBRARY www.eg.org diglib.eg.org

released onto a Petri dish with the odour on one side and no odour on the other. The proportion orienting toward the odour provides a measure of memory strength. Such paradigms illustrate the overarching goal: to link experience and neural manipulations to behavioural strategies.

2.2. Data Collection

In a typical experiment (*Fig. 1*), a research question (**a**) guides experimental design (**b**), larval behaviour is recorded via computer vision (**c**), producing time series of locomotion, posture, and orientation across thousands of frames. These are aggregated into high-dimensional profiles (**d**), which are then explored through visualization to generate and refine hypotheses (**e**).

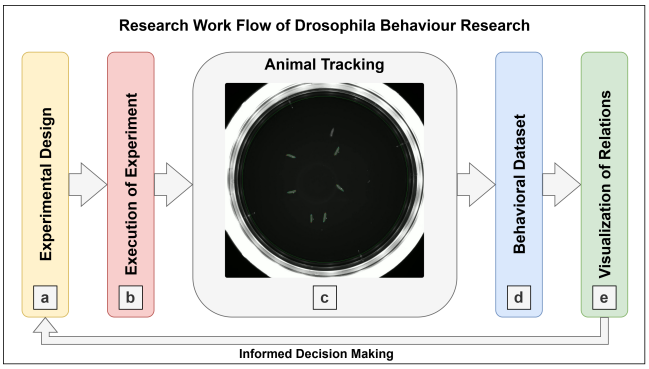


Figure 1: *Research Work Flow for Drosophila Larvae behaviour: (a) A research question or hypothesis guides the study. (b) Experiments are conducted in the lab. (c) Larvae are tracked using specialized software. (d) The resulting high-dimensional dataset is processed. (e) Data interpretation leads to new hypotheses, returning to (A).*

2.3. Data Sets

We analysed five behavioural datasets (*Fig. 1d*), each with $d = 103$ dimensions but differing in sample size and aggregation level:

D_1: Age-dependent locomotion, group-level ($n = 249$). Behaviour recorded from larvae aged four to six days [Sch24].

D_2: Odour conditioning, group-level ($n = 358$). Classical conditioning with fructose (CS) paired with amylacetate (US) [RW72].

D_3: Odour conditioning, individual-level ($n = 869$). Equivalent to D_2 at single-larva resolution.

D_4: Training trials, group-level ($n = 321$). Behaviour across different numbers of training trials.

D_5: Training trials, individual-level ($n = 818$). Equivalent to D_4 at single-larva resolution.

The 103 dimensions span five categories: odour orientation (PREF), head-casting metrics (HC), run statistics (RUN), posture and motion features (TS), and tracking features (TRACK). Detailed formulas are provided in the supplementary material.

2.4. Lower-level Domain Goals

Most behavioural researchers are familiar with standard statistical methods, but current workflows rely on custom code and static plots, limiting flexible multivariate exploration. To inform our system design, we conducted a structured feedback session with four experts at Hokkaido University, who routinely perform larval learning experiments and data analysis (see *Section 6*). This process led to four key requirements:

R1: Dimensions exploration. Inspect distributions, variance, and group-wise differences of behavioural dimensions, supported by clear attribute descriptions to capture inter-individual variability [TPS*23].

R2: Relation discovery. Identify relations across categorical, numerical, and mixed dimensions (e.g., whether turning depends on stimulus history).

R3: Relation validation. Judge relations for robustness and biological plausibility through outlier inspection and subgroup comparisons.

R4: Dimension grouping. Manage complexity by clustering correlated or dependent dimensions to reduce redundancy.

These requirements directly informed our system design, which combines linked views, type-sensitive relation metrics, and interactive filtering to support both hypothesis-driven and exploratory workflows.

3. Related Work

This work intersects VA for relation exploration in heterogeneous high-dimensional data and design studies in scientific domains. Following the design study methodology [BLI*11, SMM12, Sed16], our approach addresses real-world challenges faced by behavioural neuroscientists, with a system tailored to their analytical workflows (**R1–R4**). Design studies contribute through problem abstraction, iterative refinement, and domain-specific insights, balancing practical utility with methodological rigor.

3.1. Relation Visualization Techniques

Relation visualization typically relies on matrix-based heatmaps or node-link diagrams. Seo and Shneiderman's rank-by-feature framework systematically explores multidimensional data by ranking 1D distributions and 2D scatterplots according to statistical criteria [SS05]. Our system follows a similar idea of supporting systematic relation discovery, but extends it beyond axis-parallel 1D/2D projections to heterogeneous numerical and categorical dimensions. Heatmaps are common in biological sciences [EWJP17, GES16], using color gradients and hierarchical clustering to reveal dense relation patterns. However, scalability suffers as axis referencing becomes cumbersome with increasing dimensionality. Variants like Meshmap [KAB*20] add a third encoding (height) but at higher cognitive cost. Layouts such as force-directed [FR91], or projections such as MDS [Tor58], t-SNE [vdMH08], and UMAP [MHSG18] emphasize clusters. They scale better for sparse matrices but suffer from clutter when densely connected. Heatmaps remain dominant in confirmatory tasks [GES16].

Our system combines network-based relation views and dendrogram clustering [War63] to support relation exploration and validation in heterogeneous high-dimensional data.

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

3.2. Visual Analytics for Heterogeneous High-Dimensional Data

Visual analysis of heterogeneous high-dimensional data is complicated by the need to preserve both numerical and categorical semantics. Early solutions like interactive quantification [JFJJ08] and MixSOM [HL12] converted categorical data into numerical proxies to enable joint analysis, as also seen in Correlation Map [ZMZM15] and Zhou et al.'s dual-analysis framework [MGU*21]. Bernard et. al [BRS*12] combine numerical and categorical data by enriching clusters with metadata information.

Network-based systems [BSW*14, ZCYY19] use type-sensitive similarity metrics and interactive filtering to manage sparsity in relation networks. CatNetVis [TBL23] focuses on categorical-categorical relations but lacks integration of numerical variables and validation tools. Muhammad et al. [SMM*20] combine star coordinates for numerical data with enhanced mosaic displays and pairwise contingency plots for categorical data, providing tightly linked exploration of heterogeneous datasets. DimLift [GMS*21] extends this space by combining dimensional bundling with embedding-driven exploration, allowing analysts to navigate heterogeneous data at multiple semantic levels. Efforts like the Parallel Assemblies Plot [CMADDF23] and evaluations by Linsen et al. [ML22] further underline that naive mappings of categorical data onto numerical axes increase cognitive load, reinforcing the need for type-aware relation visualization.

Our solution extends network-based relation analysis with type-sensitive metrics, hierarchical grouping, and integrated statistical validation, directly supporting scientific workflows in behavioural research without requiring full data-type transformation.

4. Visual Analytics System

Our system supports the visual exploration of relations through six coordinated views (*Fig. 2*): the **Dimension View**, **Network View**, **Inspection View**, **Validation View**, **Dendrogram View**, and **Settings View**. Following Shneiderman's Visual Information-Seeking Mantra [Shn96], the system provides an overview of global relations (via the **Network View**), supports zoom and filtering (through the **Selection View**), and offers details-on-demand with statistical plots and distributions (in the **Validation View**). The **Dimension View** and **Dendrogram View** help users structure and group dimensions to reduce complexity and enhance interpretability.

4.1. Data Processing

For each pair of dimensions, we compute a type-aware association measure: Spearman's rank correlation coefficient (ρ_s) for numerical–numerical pairs, Cramér's V for categorical–categorical pairs, and the coefficient of determination (R^2) from a linear regression model with the categorical variable as predictor and the numerical variable as response for mixed-type pairs. These measures yield values in $[0,1]$ (or $[-1,1]$ for ρ_s, which we take in absolute value to quantify strength independent of direction). We interpret these association strengths a_{ij} as *similarities* between dimensions i and j. To obtain a metric suitable for projection and clustering, we convert them into dissimilarities $d_{ij} = 1 - a_{ij}$, forming a symmetric dissimilarity matrix $\mathbf{D}$. This matrix serves as input to dimensionality reduction (force-directed layout or multidimensional scaling) to embed dimensions in 2D, and to hierarchical clustering (average linkage) to group related dimensions. Previous approaches, such as [BSW*14], rely on binning and compute mutual information, which requires selecting binning parameters. In contrast, our type-aware approach avoids binning entirely, eliminating the need for parameter tuning while ensuring appropriate association measures for each data type.

4.2. Visual Design

The **Dimension View** (*Fig.2a*) presents all dimensions in a compact grid. Stroke color indicates cluster membership; fill color encodes centrality. Users can sort by name, cluster, type, or centrality. Tooltips provide descriptions on hover to aid interpretation. Selecting a dimension (highlighted in red in *Fig.2 (S1)*) links it across all views. This view supports initial exploration and grouping for hypothesis generation or validation (**R1**, **R4**).

The **Network View** (*Fig. 2b*) provides an overview of relations between all dimensions in the dataset. Nodes represent dimensions and edges indicate the strength of pairwise relations. Node shapes indicate dimension type: circles for numerical and squares for categorical dimensions. Edges are double-encoded by color and stroke width according to relation strength. Node color reflects cluster membership from hierarchical clustering using Ward's method. Users can interactively filter edges by relation strength using sliders, and choose to exclude outliers from relation calculations based on an adjustable IQR threshold. Brushing into a region activates the **Selection View** (*Fig. 2c*), which magnifies the selected subset while preserving the layout and encodings of the **Network View**. This facilitates focused analysis without losing context. We can see the selected dimensions are highlighted in red (*Fig. 2 (S2)*).

To explore hierarchical structure among dimensions, we include a **Dendrogram View** (*Fig. 2e*). Users can select linkage methods and adjust a cutoff threshold to control cluster granularity. This supports grouping dimensions based on relational patterns and complements the spatial layout of the **Network View** with a structural perspective (**R4**). The colour coding in both the **Dimension View** and **Network View** updates accordingly.

The **Validation View** (*Fig. 2f*) offers detailed visualizations tailored to the types and number of selected dimensions. For univariate analysis, numerical dimensions are displayed as histograms to show their distribution, while categorical dimensions are represented with bar charts indicating category frequencies. For bivariate analysis, pairs of numerical dimensions are shown as scatter plots annotated with Spearman's rank correlation coefficient and corresponding p-values. Categorical pairs are visualized as contingency table heatmaps, accompanied by results from a chi-square test. When analyzing one numerical and one categorical dimension, boxplots grouped by category are used, with additional statistical annotations such as the coefficient of determination (R^2) and p-values to support interpretation. For analysis of three numerical dimensions, the third variable is color encoded. When all three selected dimensions are categorical, the system generates faceted heatmaps grouped by the third dimension.

To support flexible analysis, users can apply custom binning strategies, exclude outliers from the plots, or interactively rearrange

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

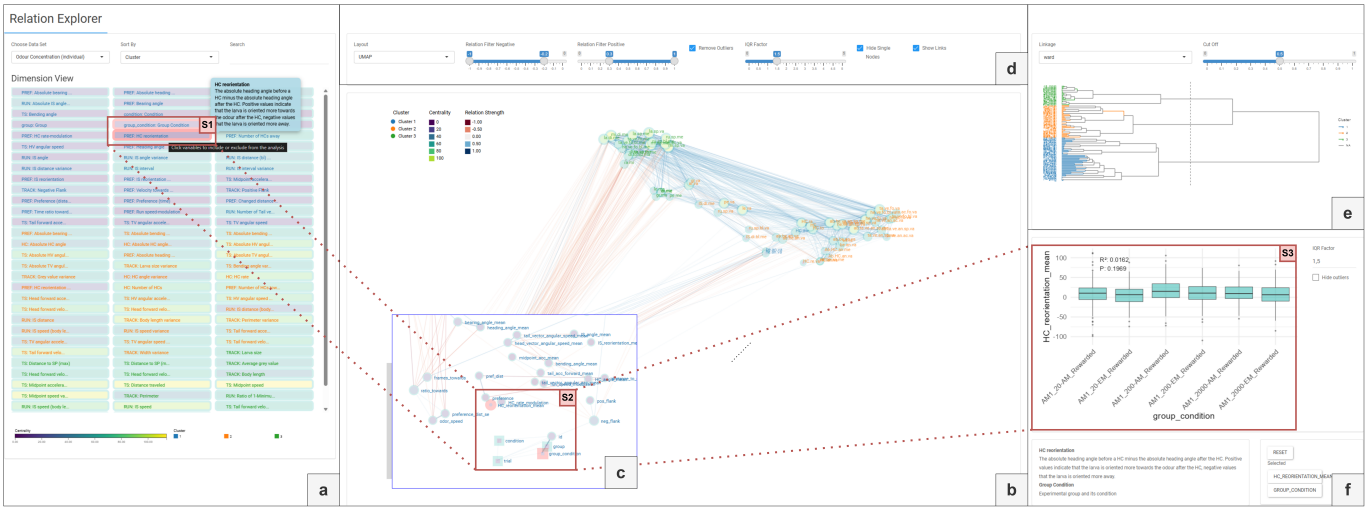


Figure 2: *System using coordinated views: (a)* ***Dimension View*** *displays all attributes grouped by cluster, with color and layout encodings for type, centrality, and cluster membership. (b)* ***Network View*** *encodes type-aware relations (Spearman, Cramér's V, R²) using colored edges and clustered layout. (c)* ***Selection View*** *magnifies selected subnetworks for closer inspection. (d)* ***Settings Panel*** *allows users to control relation strength thresholds, layout options (e.g. UMAP, t-SNE) and outlier filtering. (e)* ***Dendrogram View*** *visualizes hierarchical relations based on Ward linkage, enabling dimension grouping. (f)* ***Validation View*** *shows statistical plots for selected dimensions with interactive filters and tooltips. (S1) Selection in the* ***Dimension View*** *(S2) Linking in the* ***Selection View****, (S3) Linking in the* ***Validation View***

axes to explore different dimension combinations. All interactions are coordinated with the rest of the system (*Fig. 2 (S3)*), enabling users to link observations in the **Validation View** back to other visual contexts, thus supporting deeper investigation of discovered relations (cf. **R2**, **R3**).

The **Settings View** (*Fig. 2d*) provides interactive controls for adjusting the layout and filtering parameters of the relation network. Users can switch between different layout algorithms, including force-directed, multidimensional scaling (MDS), t-distributed stochastic neighbor embedding (t-SNE), and uniform manifold approximation and projection (UMAP). Correlation thresholds for positive and negative relations are controlled via separate sliders to refine which links are displayed. Additional options allow users to hide singleton nodes and disconnected edges for a cleaner and more interpretable view. An interquartile range (IQR)-based filter is applied to remove outliers from the relation matrix calculation. These controls enable users to dynamically adapt the network to their current analytical focus.

4.3. Implementation

The system is implemented as a web-based application using `R` and `Shiny` [CCA*21]. The server-side logic handles data preprocessing, computation of relation matrices, clustering, and statistical validation using native `R` functions and the `ggplot2` [Wic16] library for generating statistical plots. The client-side interface is built with `D3.js` [BOH11] to render the interactive network view.

5. Use Cases

We demonstrate how the system supports relation exploration, dimension grouping, and validation through use cases on real behavioural datasets. These examples show how users can reveal global patterns, inspect clusters, and investigate correlations in line with our domain requirements (**R1–R4**).

5.1. Overview and Dimension Grouping

This use case demonstrates how the system supports an overview of behavioural dimensions, their clustering, and their relational structure, as well as detailed inspection of specific relations (*Fig. 3*). The **Dimension View** *(Fig.3a)* provides an overview of all attributes, grouped by cluster and colored according to the dendrogram cutoff. The cluster assignments are derived from the **Dendrogram View** (*Fig. 3c*), where the cutoff threshold determines the cluster boundaries.

The color encoding for clusters, centrality, and relation strength is explained in the **Legend** (*Fig. 3b*). The **Network View** (*Fig. 3d*) reveals a global structure, with three main clusters. The orange cluster (**C3**) consists primarily of speed- and morphology-related metrics, such as *Midpoint Speed*, *Tail Forward Velocity*, and *Body Length*. In contrast, the green cluster (**C2**) contains variance-related metrics like *IS Speed Variance* and *Head Velocity Variance*. A strong negative correlation between **C3** and **C2** is clearly visible as the dominant pattern of red edges between them. This indicates that individuals with higher average speed tend to exhibit less variability in their speed—an intuitively plausible relation.

The blue cluster (**C1**) contains head-movement-related metrics including *Head Cast Rate*, *Bending Angle*, and *Head Vector Angular Speed*, alongside odour orientation features (*PREF*), indicating a close relation between navigation behaviour and head movement dynamics.

To confirm the negative correlation observed between speed and speed variance in the **Network View** (*Fig. 3d*), the relation between

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

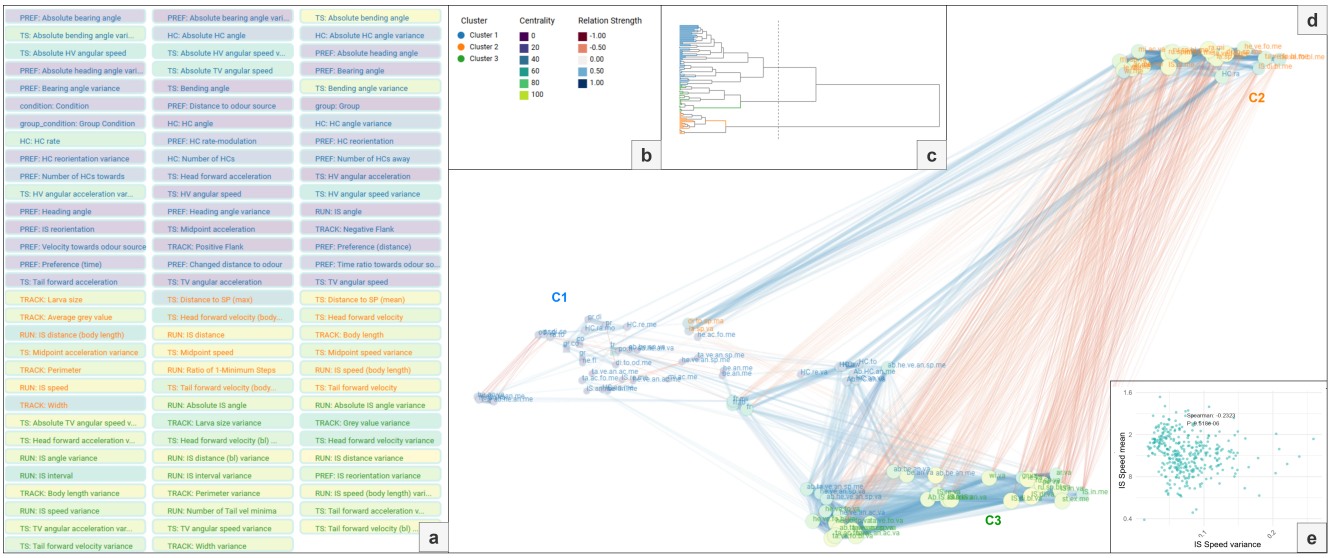


Figure 3: *Dimension grouping (using D_3): (a) Dimension View, (b) Legend for cluster, centrality, and relation strength, (c) Dendogram View, (d) Network View showing negative correlation between C3 and C2, (e) Validation View showing negative correlation between IS speed mean and IS speed variance*

IS Speed Mean and *IS Speed Variance* was further inspected in the **Validation View** (*Fig. 3e*). The scatterplot confirms the negative correlation, with a Spearman coefficient of –0.23 ($p < 0.001$), supporting the global network pattern with statistical validation.

This use case demonstrates how the coordinated views facilitate grouping behavioural dimensions based on their interdependencies and examining both the composition of each group and their relations to others.

5.2. Relation Validation and Interpretation

After obtaining an overview of all dimensions, their clusters, and relationships, the VA system enables users to explore specific relations in greater detail. For this, dimensions can be selected using the **Dimension View** or the **Selection View**. We proceed to give examples where such validation confirmed previous research and present a surprising discovery that was enabled with our approach.

Different Stages of Development The first dataset, D_1 [Sch24], examines larvae at different ages: four, five, and six days old. Initially a subset of nodes from the **Network View** is selected for inspection (*Fig. 4a*). The analysis reveals a clear relation between larval width and age (*Fig. 4b*). Further exploration highlights a cluster of interconnected dimensions. Selecting these dimensions shows that length, perimeter and run speed are all correlated (*Fig. 4c*). This correlation is plausible, as these dimensions describe physical characteristics that change as the larvae grow. Interestingly the correlation between the length and the run speed is weaker for 4 days old larvae (0.54) than for 5 day old larvae (0.38) (*Fig. 4d*). The correlation of these dimensions is strongest for the 6 day old larvae (0.69).

Unexpected Relation between Head Cast Rate and Midpoint Speed Using a force-directed layout and a correlation filter of 0.4–1, domain experts observed a positive correlation between HC rate and midpoint speed. *Fig. 5* shows this for D_2, but the pattern persists in both learned behaviour datasets (D_2 and D_3). This is surprising for our domain experts as HC events are known to reduce the speed - which would have predicted the opposite correlation, if any. Upon closer investigation, our domain experts found that there is also a positive correlation between the HC rate and the run speed which explicitly excludes HC events in the calculation. This correlation could be found in all groups of both experiments, but it seemed to be more pronounced with more training trials.

In addition, in the free locomotion data set D_1 that compared larvae of different age, no such correlation could be observed. Notably, these animals were of the same genotype as in the learning experiments, and the main differences between the experiments were that in the learning experiments a training procedure occurred before the recording and an odour source was present during the recording. Based on the observation that the correlation becomes stronger with the number of training trials (D_4, D_5), our domain experts hypothesize that the training procedure causes some individuals to become generally more active, leading to both an increased HC rate and higher speed. This surprising and novel hypothesis was made possible by the interactive and exploratory nature of our system. It has not been reported before and will be further investigated through targeted experiments by our collaborators.

6. Domain Expert Feedback

To assess usability and effectiveness, we interviewed four domain experts in Drosophila behavioural research at Hokkaido University: a PI, a PhD candidate, and two Master's students. None had been involved in system development, and this was their first exposure. The PI was most familiar with data processing, while the students and PhD candidate regularly perform larval learning experiments

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

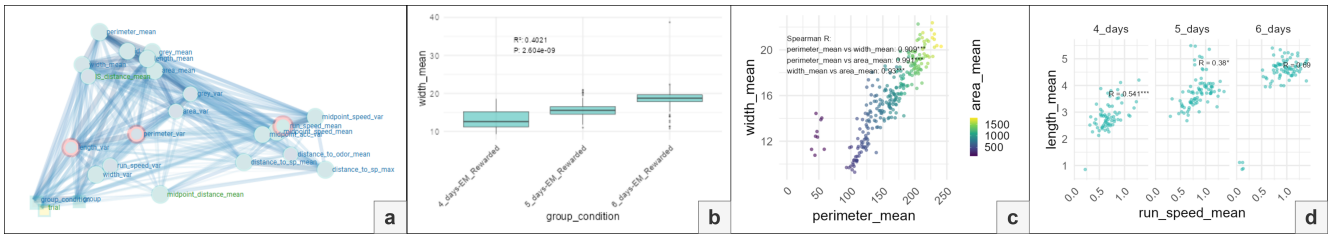


Figure 4: *Use Case for D_1 (a) **Selection View** of a cluster, (b) Boxplot of relation between the categorical dimension group_condition (The age in days) and the Width mean, (c) Use case for D_2 Relation between Midpoint speed mean and HC Rate*

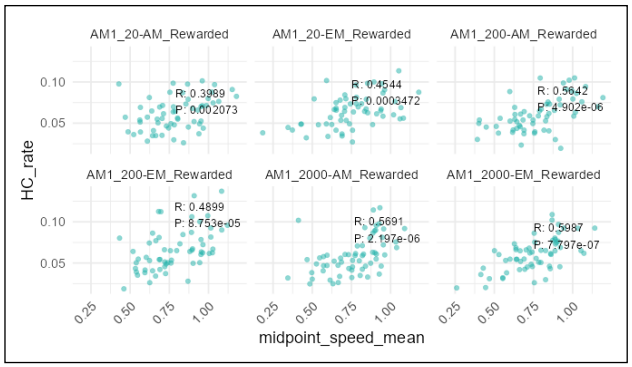


Figure 5: *Relation between midpoint speed and HC rate in D_2. Odour conditions are `AM_Rewarded` (Amylacetate rewarded) and `EM_Rewarded` (Empty rewarded), with odour concentrations of 20, 200, and 2000 mM.*

and data analysis. Feedback was gathered in an earlier development stage and led to key improvements, including tooltips, a Dendrogram view, and enhanced filtering.

Each 30-minute interview followed a structured format: 5 minutes of dataset introduction, 5 minutes of system presentation, 15 minutes of hands-on exploration, and 10 minutes for qualitative feedback and the System Usability Score (SUS) questionnaire. Participants are referred to as U1–U4 in the analysis.

6.1. Exploration Tasks

In this part of the interview the users where asked to use the system to solve three tasks using D_2.

T1: What is the strongest relation that you could find in the data set?
T2: Use the cutoff, filtering and layout options to determine how many groups (clusters) of dimensions are related in the data.
T3: In which dimensions do the experimental groups (1 trial, 2 trials and 3 trials) show a difference and why?

The goal was to find out strengths and weaknesses in the exploration process by using a *Think-Aloud-Protocol* [ES80]. Participants were instructed to verbalize their thoughts as they worked through the tasks, providing insight into their interactions with the system.

Across all users, there were common themes related to identifying relations between dimensions and navigating the system. Strong relations were found between various dimensions, such as *Tail acceleration*, *Head forward velocity*, and *Absolute IS angle mean* for several users, while clusters were identified, though their interpretation varied. U1 and U4 noted the discovery of multiple clusters, with U4 identifying 4-5 larger clusters and U1 observing clusters linked to angular speeds and velocities. U3 also saw differences in experimental groups, particularly *Preference* and *Experimental Condition* (In this case the difference in training trials), a relation similarly noted by U2.

Users found the system's filtering, cutoff, and layout options helpful for identifying relations, though clustering was challenging, particularly in defining the number of groups and their relations. U2 and U3 found clustering slightly confusing, with U3 struggling to derive specific insights despite using the options.

Users suggested improvements for navigation. U1 recommended refining the sliders and dimension selection tool, while U2 suggested better visibility of dimension names in the plots for easier interpretation. U4 had difficulty navigating between panels. Despite these challenges, users appreciated the tool's effectiveness, especially its descriptive features, color coding, and ability to reveal meaningful relations.

6.2. Usability

Statement	U1	U2	U3	U4
I think that I would like to use this system frequently.	4	4	5	4
I found the system unnecessarily complex.	1	2	3	1
I thought the system was easy to use.	5	3	5	4
I would need the support of a technical person to use this system.	1	2	2	2
I found the functions were well integrated.	4	5	5	5
I thought there was too much inconsistency.	1	1	1	1
Most people would learn to use this quickly.	5	4	5	5
I found the system cumbersome.	2	1	1	1
I felt confident using the system.	3	2	2	4
I needed to learn a lot before getting started.	3	3	3	3
Overall Score	**82.5**	**72.5**	**80**	**85**

Table 1: *User responses with System Usability Scale (SUS) results.*

To evaluate the system's usability, we combined a structured

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

questionnaire with qualitative feedback. Specifically, we administered the standardized System Usability Scale (SUS) [Bro96] and asked open-ended questions about user experience.

The SUS is a widely used 10-item Likert-scale questionnaire for assessing system usability. While typically applied in larger-scale studies, SUS has also been used effectively in small-sample qualitative design studies to capture subjective impressions of usability. Despite its simplicity, it is considered reliable even with expert groups of 5–10 participants, with scores above 70 indicating good usability. In our study, all participants rated the system positively, with scores ranging from 72.5 to 85 (Table 1).

In addition to the SUS, participants answered three open-ended questions:

Q1: Did the system help to better understand the underlying research question?
Q2: Which features supported understanding the relations in the data? What features were missing?
Q3: Would you consider using this software in the future?

Participants highlighted the system's intuitive interaction, clarity of visual encodings, and its ability to support both hypothesis-driven and exploratory workflows. Network and scatter plots were praised for effectively conveying relations and clusters (U1–U3), though one user noted visual density (U3). Suggested improvements included larger viewports, tooltips, and more consistent filtering (U1–U3).

The **Dimension View** and the search function were valued for narrowing down relevant dimensions (U2, U3), while the **Validation View** was seen as effective for detailed inspection and outlier exclusion (U1, U2). Familiar visual idioms like boxplots, heatmaps, and histograms were well-received, although SPLOMs were sometimes mentioned as hard to interpret, this is why we opted f. One user recommended improved font scaling and onboarding (U3). Overall, participants expressed that the system would be useful for exploring behavioural datasets, including in new research contexts such as decision-making studies in Drosophila larvae (U3).

7. Discussion

This design study demonstrates how type-aware relation visualization supports exploratory analysis in heterogeneous high-dimensional behavioural data. A key finding is that combining network layouts with hierarchical clustering and interactive selection addresses typical scalability challenges. While node-link diagrams are known to suffer from clutter in dense datasets, the integration of the **Selection View** mitigates this issue by allowing users to focus on subsets without losing the global context.

A limitation lies in the expressiveness of the relation measures. The system focuses on pairwise relations—Spearman correlation, Cramér's V, and R^2—which capture linear or monotonic dependencies but cannot represent more complex or conditional relationships. Extending the framework to detect nonlinear patterns or conditional dependencies remains an open challenge.

A major lesson learned concerns the importance of integrating semantic information directly into the visual interface. Features like dimension descriptions and linked views were repeatedly mentioned as critical for reducing cognitive load. Without these, users struggled to interpret dimension names and relations.

The design process also revealed how visual overviews shift analytical workflows. While domain experts traditionally relied on hypothesis-driven comparisons, access to a global relation view encouraged open-ended exploration. This shift directly contributed to the discovery of a previously unknown correlation between head casting and locomotion speed, illustrating how the tool supports not only validation but also hypothesis generation.

The system currently relies on pairwise relation metrics, which do not capture higher-order or conditional dependencies. Future work should extend the framework to support non-linear or multi-way relations. Another limitation is the focus on aggregated features rather than direct integration of raw time series or video data. Linking relational patterns back to temporal dynamics would further enhance the system's analytical power.

8. Conclusion

We presented a VA system that helps behavioural researchers explore, group, and validate relations in heterogeneous high-dimensional data. Combining type-aware relation metrics, hierarchical clustering, and coordinated views, it supports both hypothesis-driven and exploratory workflows. Semantic metadata and validation plots bridge global structure and detailed analysis.

Validated with domain experts, the system confirmed known patterns and revealed new insights—such as a surprising correlation between head casting and locomotion speed. This highlights how task-driven visual tools can extend traditional analysis. Future work will generalize this approach to other domains and extend the framework beyond pairwise analysis to capture higher-order and nonlinear relationships.

Acknowledgements

This research was supported by a stipend from the Goho Life-science Foundation, whose financial assistance is gratefully acknowledged.

References

[BLI*11] BERTINI E., LAM H., ISENBERG P., PLAISANT C., CARPENDALE S.: Empirical studies in information visualization: Seven scenarios. *IEEE Transactions on Visualization and Computer Graphics* (2011). 2

[BOH11] BOSTOCK M., OGIEVETSKY V., HEER J.: D3: Data-driven documents. *IEEE Transactions on Visualization and Computer Graphics* (2011). 4

[Bro96] BROOKE J.: Sus: A "quick and dirty" usability scale. *Usability Evaluation in Industry* (1996). 7

[BRS*12] BERNARD J., RUPPERT T., SCHERER M., SCHRECK T., KOHLHAMMER J.: Guided discovery of interesting relationships between time series clusters and metadata properties. In *Proceedings of the 12th International Conference on Knowledge Management and Knowledge Technologies* (2012), i-KNOW '12, Association for Computing Machinery. 3

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.

[BSW*14] BERNARD J., STEIGER M., WIDMER S., LÜCKE-TIEKE H., MAY T., KOHLHAMMER J.: Visual-interactive exploration of interesting multivariate relations in mixed research data sets. *Computer Graphics Forum* (2014). 3

[CCA*21] CHANG W., CHENG J., ALLAIRE J., XIE Y., MCPHERSON J.: *shiny: Web Application Framework for R*, 2021. 4

[CMADDF23] CANTU A., MICÓ-AMIGO M. E., DEL DIN S., FERNSTAD S. J.: Parallel assemblies plot, a visualization tool to explore categorical and quantitative data: application to digital mobility outcomes. In *2023 IEEE 16th Pacific Visualization Symposium (PacificVis)* (2023). 3

[ES80] ERICSSON K. A., SIMON H. A.: Verbal reports as data. *Psychological Review* (1980). 6

[EWJP17] ENGLE S., WHALEN S., JOSHI A., POLLARD K. S.: Unboxing cluster heatmaps. *BMC Bioinformatics* (2017). 2

[FR91] FRUCHTERMAN T. M., REINGOLD E. M.: Graph drawing by force-directed placement. *Software: Practice and Experience 21*, 11 (1991), 1129–1164. 2

[GES16] GU Z., EILS R., SCHLESNER M.: Complex heatmaps reveal patterns and correlations in multidimensional genomic data. *Bioinformatics* (2016). 2

[GMS*21] GARRISON L., MÜLLER J., SCHREIBER S., OELTZE-JAFRA S., HAUSER H., BRUCKNER S.: Dimlift: Interactive hierarchical data exploration through dimensional bundling. *IEEE Transactions on Visualization and Computer Graphics* (2021). 3

[HL12] HSU C.-C., LIN S.-H.: Visualized analysis of mixed numeric and categorical data via extended self-organizing map. *IEEE Transactions on Neural Networks and Learning Systems 23* (2012), 72–86. 3

[JFJJ08] JOHANSSON FERNSTAD S., JERN M., JOHANSSON J.: Interactive quantification of categorical variables in mixed data sets. *2008 12th International Conference Information Visualisation* (2008). 3

[KAB*20] KRAUS M., ANGERBAUER K., BUCHMÜLLER J., SCHWEITZER D., KEIM D. A., SEDLMAIR M., FUCHS J.: Assessing 2d and 3d heatmaps for comparative analysis: An empirical study. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (2020). 2

[MGU*21] MÜLLER J., GARRISON L., ULBRICH P., SCHREIBER S., BRUCKNER S., HAUSER H., OELTZE-JAFRA S.: Integrated dual analysis of quantitative and qualitative high-dimensional data. *IEEE Transactions on Visualization and Computer Graphics* (2021). 3

[MHSG18] MCINNES L., HEALY J., SAUL N., GROSSBERGER L.: Umap: Uniform manifold approximation and projection. *Journal of Open Source Software* (2018). 2

[ML22] MATUTE J., LINSEN L.: Evaluating data-type heterogeneity in interactive visual analyses with parallel axes. *Computer Graphics Forum* (2022). 3

[Mun09] MUNZNER T.: A nested model for visualization design and validation. *IEEE Transactions on Visualization and Computer Graphics* (2009). 1

[PRPS17] PAISIOS E., RJOSK A., PAMIR E., SCHLEYER M.: Common microbehavioral footprint of two distinct classes of conditioned aversion. *Learning & Memory 24*, 5 (Apr. 2017), 191–198. 1

[RW72] RESCORLA R. A., WAGNER A. R.: A theory of Pavlovian conditioning: Variations on the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory*, Black A. H., Prokasy W. F., (Eds.). Appleton-Century-Crofts, New York, 1972, pp. 64–99. 2

[Sch24] SCHLEYER M.: Locomotion of drosophila melanogaster larvae depending on age. DOI: https://doi.org/10.12751/g-node.vlg6pi, 2024. Accessed: 2024-12-04. 2, 5

[Sed16] SEDLMAIR M.: Design study contributions come in different guises: Seven guiding scenarios. In *Proceedings of the Sixth Workshop on Beyond Time and Errors on Novel Evaluation Methods for Visualization* (2016), Association for Computing Machinery. 2

[Shn96] SHNEIDERMAN B.: The eyes have it: A task by data type taxonomy for information visualizations. *Proceedings 1996 IEEE Symposium on Visual Languages* (1996), 336–343. 3

[SMM12] SEDLMAIR M., MEYER M., MUNZNER T.: Design study methodology: Reflections from the trenches and the stacks. *IEEE Transactions on Visualization and Computer Graphics 18* (12 2012), 2431–2440. 1, 2

[SMM*20] SHAHID M. L. U. R., MOLCHANOV V., MIR J., SHAUKAT F., LINSEN L.: Interactive visual analytics tool for multidimensional quantitative and categorical data analysis. *Information Visualization* (2020). 3

[SRP*15] SCHLEYER M., REID S. F., PAMIR E., SAUMWEBER T., PAISIOS E., DAVIES A., GERBER B., LOUIS M.: The impact of odor–reward memory on chemotaxis in larval idrosophila/i. *Learning &amp Memory 22*, 5 (Apr. 2015), 267–277. 1

[SS05] SEO J., SHNEIDERMAN B.: A rank-by-feature framework for interactive exploration of multidimensional data. *Information Visualization* (2005). 2

[TBL23] THANE M., BLUM K. M., LEHMANN D. J.: CatNetVis: Semantic Visual Exploration of Categorical High-Dimensional Data with Force-Directed Graph Layouts. In *EuroVis 2023 - Short Papers* (2023). 3

[Tor58] TORGERSON W. S.: *Theory and methods of scaling*. Wiley, 1958. 2

[TPS*23] THANE M., PAISIOS E., STÖTER T., KRÜGER A.-R., GLÄSS S., DAHSE A.-K., SCHOLZ N., GERBER B., LEHMANN D. J., SCHLEYER M.: High-resolution analysis of individual drosophila melanogaster larvae uncovers individual variability in locomotion and its neurogenetic modulation. *Open Biology* (Apr. 2023). 1, 2

[TVM*19] THANE M., VISWANATHAN V., MEYER T. C., PAISIOS E., SCHLEYER M.: Modulations of microbehaviour by associative memory strength in drosophila larvae. *PLOS ONE* (2019). 1

[vdMH08] VAN DER MAATEN L., HINTON G.: Visualizing data using t-sne. *Journal of Machine Learning Research 9* (2008). 2

[War63] WARD J. H.: Hierarchical grouping to optimize an objective function. *Journal of the American Statistical Association* (1963). 2

[Wic16] WICKHAM H.: *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2016. 4

[ZCYY19] ZHANGA C., CHENB Y., YANGA J., YINC Z.: An association rule based approach to reducing visual clutter in parallel setsi. 3

[ZMZM15] ZHANG Z., MCDONNELL K. T., ZADOK E., MUELLER K.: Visual correlation analysis of numerical and categorical data on the correlation map. *IEEE Transactions on Visualization and Computer Graphics* (2015). 3

© 2025 The Author(s).
Proceedings published by Eurographics - The European Association for Computer Graphics.