InViTAG: a web application for AI-assisted exploration and grouping of health images and data

Ala Abuthawabeh*
Amman Arab University

Michael Aupetit†
Qatar Computing Research Institute,
Hamad Bin Khalifa University

Abstract

In health and clinical research, as in many other scientific domains, important discoveries stem from carefully exploring hundreds of cases, sifting through their graphical representations or images. Once the researcher has identified and validated typical cases, Artificial Intelligence techniques like clustering and classification, can be leveraged to classify the remaining cases at scale. `InViTAG` is a web application designed to support exactly this initial exploratory phase, where machines can barely help. We present the main features of this system on a use case where it helps discover categories of sleep and activity patterns in wearable data.

Index Terms: Interactive Voronoi Treemap; Augmented Intelligence; Arrangement and Grouping; Images Categorization

1 Introduction

An essential step before using data images in machine learning pipelines is data and image curation [34]. For instance, systems like `ImageJ` [30] help users enhance image quality using computer vision techniques, and detect, identify, and annotate regions of interest, enriching image data with expert knowledge. Another aspect of this data preparation relies on identifying and grouping data into meaningful categories that will serve the training of supervised classifiers [13, 25]. A critical part of this categorization process is arrangement and grouping (A&G) actions, which are key [22, 12] for domain experts to build trust and knowledge [32] while exploring new data and generating hypotheses about groups. Enhanced file managers like `Tropy` [6] make it easier to organize scientific images but lack tools to scale up the grouping process. Human-machine hybrid approaches have been proposed in the Visualization community to support these tasks [13, 36, 14, 25]. We propose an Augmented Intelligence system [16] using interactive visualization aided with semi-automatic clustering and classification called `InViTAG`, which stands for interactive Voronoi treemaps for arrangement and grouping. The design rationale of this tool has been presented in more detail in previous publications [8, 9, 10]. While `InViTAG` is still a prototype, we make it publicly available to serve the scientific community and to collect feedback on the related GitHub page. The system can be accessed from the following URL: `https://github.com/michaelaupetit/invitag` where comments are welcome. In the sequel, we remind the main design rationale and present the features of this system.

2 Context, design rationale, and technical details

This system was initially designed to support a clinician-researcher studying the physical activity and sleep of 264 patients suffering various degrees of obesity and diabetes [17, 28, 18]. They were equipped with wearable sensors capturing their level of physical activity every minute during one week, which was encoded for each patient as a bar chart (Figure 1(Z)) with sleep (blue), sedentary, moderate, and vigorous (orange to red) segments.

The clinician needed to visualize each patient's data to allow the visual discovery and categorization of activity and sleep patterns that were likely unknown by the clinician and the health community at large given the novelty of the technology at that time [7, 18, 11, 21]. This process typically resorts to exploratory data analysis tools like dimensionality reduction (DR) [27, 35, 26] to convert multivariate measurements into scatterplots, or clustering techniques [36, 20, 31, 23, 15, 29] to create groups semi-automatically. In contrast to these approaches, we opted for an interactive visualization that can better capture the early stages of the analytic process, where the external knowledge and intuition used by the analyst are crucial to determine patterns of interest. Such a system [35] records this knowledge-in-the-making through the initial category seeds visually identified and manually grouped, and learns from them the features' importance and the multidimensional metrics to use for further automation and assistance. As this initial stage is fully manual and visual, grouping images requires optimizing space to avoid clutter and overlap while keeping individual images readable. Grouping also requires optimizing time and interactions to reduce the burden of manual and visual exploration and to ease the discovery of interesting patterns and populate the corresponding groups [10]. `Piling.js` [24] proposes a piling metaphor to support the arrangement and grouping of images. `InViTAG` enriches the piling metaphor with area-based grouping and interactions [9] to create groups semi-automatically, replacing the snap-to-grid or other DR layout arrangements with more flexible, space-efficient, and expressive power diagrams [8, 9].

`InViTAG` is a ready-to-use `Flask` application [3] for data scientists and domain experts, using `Scikit-learn` [5] for automatic classification and clustering [10] to scale the categorization process by exploiting the initial groups the analyst forms with only a few images. It uses `d3.js` [2] and `Weighted Voronoi` [1] with additional custom logic to realize the novel interactive Voronoi treemap component. It embeds `LineUp.js` [19, 4] in the `Features` tab.

3 Structure, data, and group generation

`InViTAG` interface comprises `Data`, `Features`, and `Groups` tabs.

Images and features The user starts from the `Data` tab, uploading a set of *image* files paired with a *feature* file containing a vector of numerical features for each image. Images can come in standard formats (`PNG`, `JPG`, `GIF`, or `SVG`) with various aspect ratios and sizes. The feature file can be in `CSV` or `JSON` format and must contain an `image_id` column with image filenames. The same file will be output, enriched with `group_name` and `group_color` columns identifying the group name and color assignment of each image to enable statistical analysis of the groups or training of machine learning models, and to allow reusing the same colors in external tools. `InViTAG` can be initialized with pre-existing group names and colors by providing these columns in the input feature file. Only the features are used by the semi-automatic clustering and classification processes. The size and number of images and

*e-mail: a.abuthawabeh@aau.edu.jo
†e-mail: maupetit@hbku.edu.qa (Corresponding author)

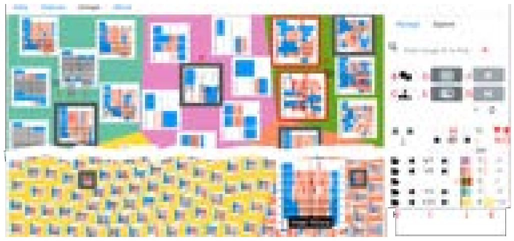

Figure 1: `InViTAG` grouping interface showing global (top vs bottom) and local on-demand (Z) visual scaling. The Explore panel (right) provides multiple actions (Section 3).

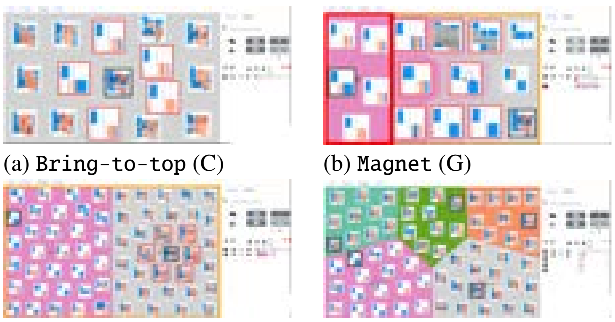

(a) `Bring-to-top` (C) (b) `Magnet` (G)

(c) `Clusterize` (F) seeds & (d) groups generation

Figure 2: (a) All the images manually selected (S)(red frame) across pages of the *Uncategorized* group (grey) have been moved (C) to the first page. (b) The user generated (N) the *Outlier* group (pink) with them and made it a target (red border). The `Magnet` action (G) attracts closer to the target group the images of the *Uncategorized* group with similar feature values. (c) `Clusterize` (F) shows six seed images (centroids of K-means with K=6). (d) The user kept three of them, which became the representative images (R) of the generated groups (teal, green, orange), filled with the top-9 nearest images to the centroids. The user always decides on selections and transfers.

(a) `Show inliers & outliers` (b) `Features` (`LineUp.js`)

Figure 3: (a) The user collapsed all groups (L) except the *Outlier* group (pink), maximizing the visual space to clean it up. The `Inliers`(D)/`Outliers`(E) actions show near the center, the most feature-based typical (top)/outlying (bottom) images of that group. (b) The `Features` tab uses `LineUp` to analyze group features.

features are only limited by the client browser's memory and by the server capacity for the automatic clustering and classification part.`InViTAG` does not implement *features-to-image* nor *image-to-features* transformations, as each application domain has specialized processing pipelines and visualizations. For instance, when exploring natural images, features can be extracted from the images with computer vision techniques [33]; images can be computed from the features, for instance, visualizing multivariate data as radar plots; or features and images can be extracted from other sources as in the proposed use case, where wearable data have been processed to extract 16 features like *average wake-up time* or *average duration of vigorous activity*, and visualized as bar charts.

Once the data have been loaded, the user can proceed with the `Features` or the `Groups` tab. The `Features` tab uses the `LineUp` interface [19] for feature-based exploration (Figure 3b). The `Groups` tab (Figure 1) is the core component of `InViTAG` that allows group creation, exploration, and management. A typical user proceeds with the following steps (All capital letters refer to Figure 1; see the GitHub project page for a complete documentation).

Optimize the visual space Users first set the size and number of images to fit the screen without clutter (M). They can page through a group (I) when not all images are visible simultaneously. The `Explore` side panel (A-O) shows the statistics of total (K) and visible (J) images in each group, and enables access to semi-automatic actions (A,B,C,D,E,F,G).

Explore patterns diversity Users explore the image patterns diversity by paging the groups up and down (I). They can use the `Randomize` action (B) to make a new random subset of images visible in each group. They can drag images side-by-side to identify common or distinct patterns, zooming in on an image transiently (Z) to better see its details.

Generate groups Users start grouping manually, by selecting individual images of interest (red frame) (S) across different pages. They move all the selected images to the first page with `Bring-to-top` action (C) (Figure 2a) to check if they could form a group, unselecting individual images that do not fit. They can use line or lasso selection to select multiple images, or unselect all images (O). They create a new group (N) from the whole selection, or create multiple groups, one selected image in each. They can also use the `Clusterize` action (F) (Figure 2cd) to generate K clusters of 10 images, using K-means in the feature space over the *Uncategorized* group (grey). Each time a group is created, the seed or an arbitrary image of this group is made its representative (bold frame (R)). This image can be changed by users, and stays visible on any page of that group to summarize its content.

Populate groups Once several groups have been created, users can manually transfer single or multiple selected images from one group to another by designing a target group (red border, Figure 2b). As the visual space gets crowded with multiple groups, users can regain focus by collapsing some of them (H,L), leaving only their representative image visible, and allocating freed slots to images of the other uncollapsed groups in proportion to their size (Figure 3a). Whole groups can be re-arranged or merged, too. Once a group is big enough, users can run the `Magnet` action (G). It uses multinomial logistic regression to order all images of the other groups by their similarity in the feature space to the ones of the target group. It arranges images closer to the target group axially (Figure 2b) emphasized with a color overlay. Users can check and select up to the page showing the last image deemed acceptable, aided with a cVIL-like probability line chart (J) [25].

Check groups Users run the `Show inliers/outliers` actions (D/E), to visualize the most/least (feature-based) typical images in a group, bringing them upfront, arranged in a radial fashion emphasized by color overlay (Figure 3a). Users can search specific images by their names (`image_id`) using the `Search` action (A). At anytime, users can switch to the `Features` tab to visualize the feature statistics of each group (Figure 3b). Any feature vector selected in the `Features` tab (orange rows) corresponds to a selected image in the `Groups` tab (S) and *vice-versa*.

Name and save groups The Manage tab of the side panel allows renaming groups and downloading the group-enriched feature file as `CSV` or `JSON`, and the entire set of images organized into group folders as a `ZIP` file.

Acknowledgments

This contribution was made possible by `NPRP11C-0115-180010/NPRP11C-30752-SP-32` Grant from the Qatar National Research Fund member of Qatar Foundation. The statements made herein are the sole authors' responsibility.

References

[1] d3 Weighted Voronoi Treemap. https://github.com/Kcnarf/d3-voronoi-map. Acc: 2020-04-08. 1

[2] D3.js, Acc [2025-07-01]. https://d3js.org/. 1

[3] Flask, Acc [2025-07-01]. https://flask.palletsprojects.com/. 1

[4] Line-Up, Acc [2025-07-01]. https://lineup.js.org/. 1

[5] scikit-learn, Acc [2025-07-01]. https://scikit-learn.org/. 1

[6] Tropy, Acc [2025-07-01]. https://tropy.org/. 1

[7] Y. Abdelaal, M. Aupetit, A. Baggag, M. Bashir, and D. Al-Thani. How much wearable data is enough for the utility and trust of augmented artificial intelligence systems? A scenario-based interview with medical professionals. *Int. J. Hum. Comput. Interact.*, 41(12):7684–7710, 2025. doi: 10.1080/10447318.2024.2400388 1

[8] A. Abuthawabeh and M. Aupetit. A force-directed power diagram approach for interactive voronoi treemaps. In A. Kerren, C. Garth, and G. E. Marai, eds., *EuroVis, Norrköping, Sweden [online only]*, pp. 109–113. Eurographics Association, 2020. doi: 10.2312/evs.20201057 1

[9] A. Abuthawabeh and M. Aupetit. Toward an Interactive Voronoi Treemap for Manual Arrangement and Grouping. In M. Agus, C. Garth, and A. Kerren, eds., *EuroVis - Short Papers*, pp. 97–101, 2021. doi: 10.2312/evs.20211062 1

[10] A. Abuthawabeh, A. Baggag, and M. Aupetit. Augmented intelligence with interactive voronoi treemap for scalable grouping: a usage scenario with wearable data. In M. Agus, W. Aigner, and T. Höllt, eds., *EuroVis - Short Papers*, pp. 43–47. Eurographics Association, 2022. doi: 10.2312/EVS.20221091 1

[11] M. Aupetit, L. Fernández-Luque, M. Singh, and J. Srivastava. Visualization of wearable data and biometrics for analysis and recommendations in childhood obesity. In P. D. Bamidis, S. T. Konstantinidis, and P. P. Rodrigues, eds., *30th IEEE International Symposium on Computer-Based Medical Systems, CBMS 2017, Thessaloniki, Greece, June 22-24, 2017*, pp. 678–679. IEEE Computer Society, 2017. doi: 10.1109/CBMS.2017.120 1

[12] E. Beauxis-Aussalet, M. Behrisch, R. Borgo, D. H. Chau, C. Collins, D. Ebert, M. El-Assady, A. Endert, D. A. Keim, J. Kohlhammer, D. Oelke, J. Peltonen, M. Riveiro, T. Schreck, H. Strobelt, and J. J. van Wijk. The role of interactive visualization in fostering trust in ai. *IEEE Computer Graphics and Applications*, 41(6):7–12, 2021. doi: 10.1109/MCG.2021.3107875 1

[13] J. Bernard, M. Zeppelzauer, M. Sedlmair, and W. Aigner. VIAL: a unified process for visual interactive labeling. *The Visual Computer*, 34(9):1189–1207, 2018. doi: 10.1007/s00371-018-1500-3 1

[14] S. Bonakala, M. Aupetit, H. Bensmail, and F. El-Mellouhi. A human-in-the-loop approach for visual clustering of overlapping materials science data. *Digital Discovery*, 3:502–513, 2024. doi: 10.1039/D3DD00179B 1

[15] M. Cavallo and C. Demiralp. Clustrophile 2: Guided visual clustering analysis. *IEEE TVCG*, 25(1):267–276, 2019. doi: 10.1109/TVCG.2018.2864477 1

[16] F. D. Felice, A. Petrillo, C. D. Luca, and I. Baffo. Artificial intelligence or augmented intelligence? impact on our lives, rights and ethics. *Procedia Computer Science*, 200:1846–1856, 2022. 3rd International Conference on Industry 4.0 and Smart Manufacturing. doi: 10.1016/j.procs.2022.01.385 1

[17] L. Fernandez-Luque, M. Aupetit, J. Palotti, M. Singh, A. Fadlelbari, A. Baggag, K. Khowaja, and D. Al-Thani. *Health Lifestyle Data-Driven Applications Using Pervasive Computing*, pp. 115–126. Springer International Publishing, Cham, 2019. doi: 10.1007/978-3-030-06109-8_10 1

[18] L. Fernández-Luque, M. Singh, F. Ofli, Y. Mejova, I. Weber, M. Aupetit, S. K. Jreige, A. K. Elmagarmid, J. Srivastava, and M. Ahmedna. Implementing 360° quantified self for childhood obesity: feasibility study and experiences from a weight loss camp in qatar. *BMC Medical Informatics Decis. Mak.*, 17(1):37:1–37:13, 2017. doi: 10.1186/S12911-017-0432-6 1

[19] S. Gratzl, A. Lex, N. Gehlenborg, H. Pfister, and M. Streit. Lineup: Visual analysis of multi-attribute rankings. *IEEE TVCG*, 19(12):2277–2286, 2013. doi: 10.1109/TVCG.2013.173 1, 2

[20] H. Jeon, M. Aupetit, D. Shin, A. Cho, S. Park, and J. Seo. Measuring the validity of clustering validation datasets. *IEEE TPAMI*, 47(6):5045–5058, 2025. doi: 10.1109/TPAMI.2025.3548011 1

[21] K. Khowaja, W. W. Syed, M. Singh, S. Taheri, O. Chagoury, D. Al-Thani, and M. Aupetit. A participatory design approach to develop visualization of wearable actigraphy data for health care professionals: Case study in qatar. *JMIR Hum Factors*, 9(2):e25880, 2022. doi: 10.2196/25880 1

[22] D. Kirsh. The intelligent use of space. *Artif. Intell.*, 73(1-2):31–68, 1995. doi: 10.1016/0004-3702(94)00017-U 1

[23] B. C. Kwon, B. Eysenbach, J. Verma, K. Ng, C. De Filippi, W. F. Stewart, and A. Perer. Clustervision: Visual supervision of unsupervised clustering. *IEEE TVCG*, 24(1):142–151, 2018. doi: 10.1109/TVCG.2017.2745085 1

[24] F. Lekschas, X. Zhou, W. Chen, N. Gehlenborg, B. Bach, and H. Pfister. A generic framework and library for exploration of small multiples through interactive piling. *IEEE TVCG*, (01):1–1, oct 2020. doi: 10.1109/TVCG.2020.3028948 1

[25] M. Matt, M. Zeppelzauer, and M. Waldner. cVIL: Class-Centric Visual Interactive Labeling. In M. El-Assady and H.-J. Schulz, eds., *EuroVis Workshop on Visual Analytics (EuroVA)*. The Eurographics Association, 2024. doi: 10.2312/eurova.20241113 1, 2

[26] T. Metsalu and J. Vilo. Clustvis: a web tool for visualizing clustering of multivariate data using principal component analysis and heatmap. *Nucleic Acids Research*, 43(W1):W566–W570, 05 2015. doi: 10.1093/nar/gkv468 1

[27] L. G. Nonato and M. Aupetit. Multidimensional projection for visual analytics: Linking techniques with distortions, tasks, and layout enrichment. *IEEE TVCG*, 2018. doi: 10.1109/TVCG.2018.2846735 1

[28] I. Perez-Pozuelo, B. Zhai, J. R. M. Palotti, R. Mall, M. Aupetit, J. M. Garcia-Gomez, S. Taheri, Y. Guan, and L. Fernández-Luque. The future of sleep health: a data-driven revolution in sleep science and medicine. *npj Digit. Medicine*, 3, 2020. doi: 10.1038/S41746-020-0244-4 1

[29] A. Pister, P. Buono, J.-D. Fekete, C. Plaisant, and P. Valdivia. Integrating prior knowledge in mixed-initiative social network clustering. *IEEE TVCG*, 27(2):1775–1785, 2021. doi: 10.1109/TVCG.2020.3030347 1

[30] C. T. Rueden, J. Schindelin, M. C. Hiner, B. E. DeZonia, A. E. Walter, E. T. Arena, and K. W. Eliceiri. Imagej2: Imagej for the next generation of scientific image data. *BMC Bioinformatics*, 18(1):529, Nov 2017. doi: 10.1186/s12859-017-1934-z 1

[31] D. Sacha, M. Kraus, J. Bernard, M. Behrisch, T. Schreck, Y. Asano, and D. A. Keim. Somflow: Guided exploratory cluster analysis with self-organizing maps and analytic provenance. *IEEE TVCG*, 24(1):120–130, 2018. doi: 10.1109/TVCG.2017.2744805 1

[32] D. Sacha, A. Stoffel, F. Stoffel, B. C. Kwon, G. Ellis, and D. A. Keim. Knowledge generation model for visual analytics. *IEEE TVCG*, 20(12):1604–1613, 2014. 1

[33] M. A. Smith and T. Chen. 9.1 - image and video indexing and retrieval. In A. BOVIK, ed., *Handbook of Image and Video Processing (Second Edition)*, Communications, Networking and Multimedia, pp. 993–XXXI. Academic Press, Burlington, second edition ed., 2005. doi: 10.1016/B978-012119792-6/50121-2 2

[34] S. Vahdati, B. Khosravi, E. Mahmoudi, K. Zhang, P. Rouzrokh, S. Faghani, M. Moassefi, A. Tahmasebi, K. P. Andriole, P. Chang, K. Farahani, M. G. Flores, L. Folio, S. Houshmand, M. L. Giger, J. W. Gichoya, and B. J. Erickson. A guideline for Open-Source tools to make medical imaging data ready for artificial intelligence applications: A society of imaging informatics in medicine (SIIM) survey. *J Imaging Inform Med*, 37(5):2015–2024, Apr. 2024. 1

[35] J. Wenskovitch and C. North. Observation-level interaction with clustering and dimension reduction algorithms. In *Proc. of the 2nd Workshop HILDA'17*. Association for Computing Machinery, New York, NY, USA, 2017. doi: 10.1145/3077257.3077259 1

[36] J. Wenskovitch and C. North. An examination of grouping and spatial organization tasks for high-dimensional data exploration. *IEEE TVCG*, pp. 1–1, 2020. doi: 10.1109/TVCG.2020.3028890 1