Multiarrangement: A Plug & Play Geometric Data Collection Package For Video Stimuli

Umur Yıldız[1] and Burcu A. Ürgen[1,2,3]

[1]Department of Neuroscience, Bilkent University, Ankara, Turkey
[2]Aysel Sabuncu Brain Research Center and National Magnetic Resonance Research Center (UMRAM), Bilkent University, Ankara, Turkey
[3]Department of Psychology, Bilkent University, Ankara, Turkey

Abstract

We present Multiarrangement, an offline, open-source Python toolkit for collecting human similarity judgments for video stimuli through multi-arrangement tasks. Participants arrange subsets of stimuli in a 2D arena such that Euclidean distances reflect perceived dissimilarity. The toolkit supports two experiment paradigms: a set-cover scheduling system that uses combinatorial cover designs to guarantee pair coverage with minimal trials, aiming to avoid overwhelming the participant for stimulus-rich settings, and an adaptive Lift-the-Weakest scheduling that focuses each new trial on the globally least certain pair and informative neighbors. Across trials, partial distance evidence is integrated into a representational dissimilarity matrix using normalized weighted averaging, with an optional inverse MDS refinement that minimizes cross-trial prediction error. We document task design, algorithms, evaluations, and provide practical guidance for easy and reliable usage.

Keywords: multiarrangement, video stimuli, representational dissimilarity matrix (RDM)

Introduction

Video has become a first-class stimulus in many behavioral and neural research paradigms. Naturalistic movies capture rich, time-varying perception and cognition and are now common in neuroimaging (Hanke et al., 2014; Liu et al., 2022), where the shared narrative time-locks cortical responses across viewers (Hasson et al., 2004) and yields more reliable, behavior-predictive functional connectivity than rest (Finn et al., 2020). At the behavioral level, validated film-clip corpora reliably elicit target emotions and support ecologically valid laboratory experiments (Gilman et al., 2017; Gross & Levenson, 1995; Schaefer et al., 2010). Naturalistic videos also strengthen behavioral paradigms by preserving multimodal dynamics, narrative structure, and social context (Baldassano et al., 2017; Nastase et al., 2020; Sonkusare et al., 2019). This results in measurements that generalize better to everyday cognition and exposes processes such as event segmentation and predictive processing that are often not well captured by static stimuli (Baldassano et al., 2017; Nastase et al., 2020; Sonkusare et al., 2019). Furthermore, videos differ from their static counterparts in that they carry a manipulable temporal and audiovisual structure. Researchers can reorder shots or remove sound while keeping content realistic, and these manipulations have measurable effects on reliability and cognition (Hasson et al., 2008; Jääskeläinen et al., 2021; Lerner et al., 2011; Uhrig et al., 2016).

Although the value of video stimuli is better appreciated, most video-rating workflows are sequential, where participants watch a clip and provide a rat-

ing on a scale. Sequential methods are well matched to scalar targets (MOS, valence), but they are inefficient for capturing relational structure among many clips, and paired-comparison protocols or triplet designs scale poorly ($O(N^2)$ or $O(N^3)$). Contemporary continuous-annotation toolchains for video provide frame-level valence/arousal traces, yet still yield unidimensional trajectories per pass (Baveye et al., 2015; Koelstra et al., 2012). What is missing for video-centric studies is a method that allows participants to see and compare multiple clips at once, while scaling to dozens of stimuli without overwhelming either the display or the participant. No off-the-shelf software package currently exists to facilitate this type of data collection for video stimuli.

Multiarrangement addresses this gap. Participants arrange small subsets on a two-dimensional canvas, and we aggregate distances across trials into a single representational dissimilarity matrix (Kriegeskorte et al., 2008; Nili et al., 2014). Each placement yields many informative pairwise constraints, which reduces the quadratic burden of exhaustive pairwise ratings while preserving fine-grained structure. Balanced subset schedules give broad coverage with modest trial counts and reduce scale-use and anchoring artifacts that often appear in scalar ratings. For controlled studies, the resulting stimulus space supports targeted sampling, pre- to post-training comparisons of category structure, and explicit tests of dimensional organization using a single dataset. For cognitive neuroscience the same dissimilarity matrix fits directly into representational similarity analysis and model-based encoding or decoding, enabling comparisons between behavior and brain across images, audio and video, with straightforward reliability checks via split halves and cross-validation. As a library, it provides reproducible scheduling, streamlined interfaces for static and time-varying stimuli, detailed trial logs, and easy integration for downstream analyses.

Features and Functionality

The design of Multiarrangement focuses on providing a turn-key solution for geometric data collection by combining a straightforward user interface with an efficient back-end for trial management. The goal is to offer a tool that is both accessible for the researcher and effective for data collection. This is accomplished through two main components, which we detail in the following sections:

User Interface

The user interface of Multiarrangement adheres to a design similar to prior multi-arrangement tools (Kriegeskorte & Mur, 2012), but is tailored for the specific demands of dynamic media. Multiarrangement offers a workspace in which a single circular arrangement area is present. Stimulus tokens initially appear in a seating area outside the circle, from which

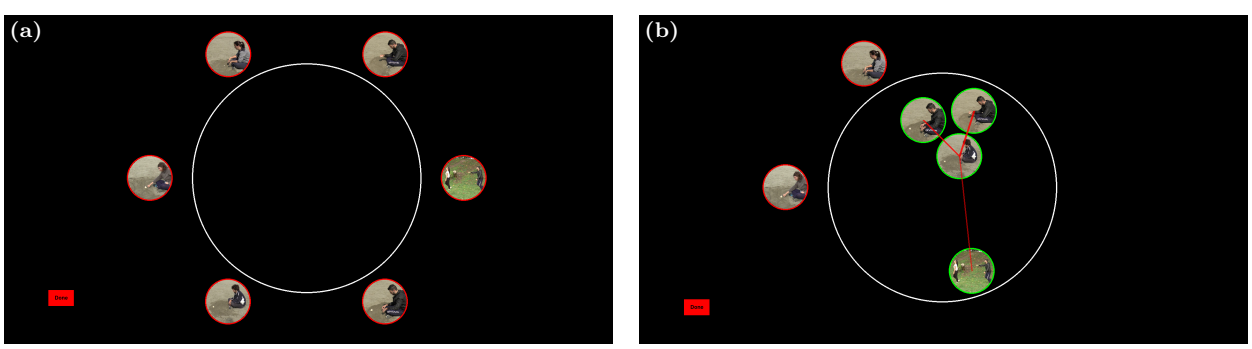

Figure 1. *Multiarrangement user interface.* The workspace at the start of a trial (a) shows the central circular arrangement area and the peripheral seating area. During placement (b), dragging a token displays transient guidance lines whose thickness/opacity scale with proximity.

the user must move them. This two-region design is a deliberate choice, adhering to the paradigm established in seminal multi-arrangement research. The circular workspace is specifically employed to minimize layout biases, as it lacks the corners or implicit axes of a square that might otherwise influence participant placements (Figure 1, panel a). For large-subset trials, an optional, center-locked zoom allows local magnification for precise placement.

The primary mode of interaction is direct manipulation via drag-and-drop. Participants use the mouse to move stimulus tokens (typically thumbnails for videos and images, or icons for audio) from the seating area into the circular workspace. A critical feature, designed specifically for dynamic media, is the ability to inspect stimuli on demand. By double-clicking any stimulus token, the participant can play the associated video in a pop-up window. This allows for repeated review and fine-grained comparisons, a necessary function when judging the similarity of complex, time-varying stimuli.

While a stimulus token is being dragged, transient guidance lines connect it to every other token currently in the circle. Line thickness/opacity scale with proximity, providing a quick visual sense of relative distances (Figure 1, panel b). By default, a trial can be submitted only after every token has been inspected at least once and all tokens are inside the circle.

The interface also uses colored rims around stimulus tokens to signal compliance status: red indicates the requirements for submission pertaining to the specific token have not yet been fulfilled, and green confirms the token satisfies the per-trial requirement. Default instruction sets are provided in Turkish and English, and the interface supports user-specified instructional text and labels via simple templates to enable customization.

Scheduling Algorithms

The scheduler selects, on each trial, a subset of stimuli that balances simultaneous comparison against per-trial effort while accumulating pairwise evidence toward a reliable RDM (Representational Dissimilarity Matrix) within a fixed session budget. We provide two complementary strategies. First is a Set-Cover mode that precomputes small batches that together cover all unordered pairs and yields a fixed, between-subject-comparable schedule. In this approach we seed from published $(v, k, t{=}2)$ coverings in the La Jolla Covering Repository (Gordon, Kuperberg & Patashnik, 1995; La Jolla Covering Repository). The second strategy is an adaptive mode that maintains an evidence matrix, selects the globally weakest pair as anchors, and expands to a bounded subset until an evidence threshold or time limit is reached. We follow the "Lift-the-Weakest" formulation of Kriegeskorte and Mur (2012) with some modifications.

Set-Cover Implementation

Problem setup Let $\mathcal{I} = \{1, \ldots, N\}$ index the stimuli. A trial presents a batch $B \subset \mathcal{I}$ of size k, covering the intra–batch pairs $\mathcal{P}(B) = \{(i, j) : i < j,\ i, j \in B\}$. We seek a small family $\mathcal{B} = \{B_m\}_{m=1}^{M}$ such that $\bigcup_{m=1}^{M} \mathcal{P}(B_m) = \{(i, j) : i < j\}$ while bounding k to avoid on-screen overload and keeping M small enough to fit the session budget.

Construction When available, we seed from published $(v{=}N, k, t{=}2)$ coverings in the La Jolla Covering Repository, otherwise we initialize greedily by iteratively adding a batch that covers the largest number of yet-uncovered pairs. We then apply a light refinement pipeline consistent with the released code: *repair/prune* to guarantee full coverage while removing batches that add no unique pairs, followed by *local search + group DFS* to smooth per-item load via item swaps and small group moves without breaking coverage. A shrink-only variant (*flex*) may reduce some late batches to $k_{\min}$ when this closes residual gaps or trims M without increasing per-trial burden.

Execution controls for video We show one batch per trial. For video, we keep k in a narrow range (e.g., $6 \leq k \leq 10$) to bound playback and interaction time, randomize initial token positions within the arena to reduce placement bias, and interleave batches so that no stimulus appears in adjacent trials beyond a user-specified limit. The schedule, RNG

```
# C_ij: pair-coverage counts; deg(i): per-item load
Input: N, k, k_min, seed, flex
rng <- init_rng(seed)

if LJCR_has(N, k) then
   B <- LJCR_seed(N, k) # published (v,k,2) blocks
else
   B <- greedy_seed(N, k, rng) # add batches
        covering most uncovered pairs
end if

if flex then
   B <- flex_shrink(B, k_min) # shrink late
        batches to k_min if helpful
end if

B <- repair_prune(B) # repair: cover all uncovered
     pairs; prune: drop batches with no unique
     pairs (keep C_ij >= 1)

B <- local_search_groupDFS(B, rng) # swaps/one-out-
     one-in/small-group moves; accept only if
     coverage preserved (C_ij >= 1) and var(deg)
     decreases

return B # final batch schedule
```

Figure 2. Set-cover scheduler.[1] Initialization uses a published $(v, k, 2)$ covering when available or a greedy seed otherwise. `flex_shrink` can trim late batches to $k_{\min}$. `repair_prune` enforces full pair coverage and removes batches contributing no unique pairs, and `local_search_groupDFS` preserves coverage while reducing exposure variance via small swaps/group moves. Output: ordered batches and coverage diagnostics.

seed, and coverage diagnostics are recorded for reproducibility.

Set-Cover for video Set-Cover fixes the per-trial set size k and reduces total trials compared with naive pairwise enumeration as N grows. Participants compare a small, consistent number of clips per trial while every unordered pair is co-presented at least once. To control duplication, we refine the schedule to remove redundant batches and balance per-item exposure $\deg(i)$ via local swaps and prune & repair iterations, which minimizes repeated co-occurrences and distributes any necessary repeats evenly across items and pairs. For remaining repeats, we aggregate per-trial distances after max-normalization using reliability weights r_t and we optionally apply a limited inverse MDS refinement to reconcile cross-trial scale and placement differences. Together, these steps preserve the time efficiency and manageable on-screen complexity of a fixed-batch design while yielding a smoother, globally consistent RDM.

[1] Algorithm pseudocode listings are framed in blue, and runnable code examples are framed in black.

Lift-the-Weakest Implementation

Problem setup Adaptive scheduling maintains an evidence matrix $E \in \mathbb{R}^{N \times N}$ that accumulates support for pairwise dissimilarities across trials. At each step the scheduler targets the least supported region while keeping per-trial effort bounded, following the inverse MDS Lift-the-Weakest principle (Kriegeskorte & Mur, 2012).

Policy The next subset anchors on the globally weakest pair $(a, b) = \arg\min_{i<j} E_{ij}$. Starting with $S = \{a, b\}$, items are added greedily up to a size bound to increase expected information, preferring candidates that create low-evidence links to members of S. After a layout is collected, within-trial distances are normalized by the trial maximum and accumulated into E with weights $w_{ij} = (d_{ij}^{(t)}/d_{\max}^{(t)})^{\alpha}$ with a small $\alpha \in [1, 2]$, and the run stops when $\min_{i<j} E_{ij} \geq \tau$ or a trial or time cap is reached (Kriegeskorte & Mur, 2012).

Execution controls for video Subset size is bounded $k_{\min} \leq |S| \leq k_{\max}$ to keep playback and interaction time practical. Light diversity constraints limit consecutive reuse of the same anchors and discourage high Jaccard overlap between successive subsets. A limited-iteration inverse MDS refinement can be applied intermittently or at the end to improve cross-trial consistency without large runtime (Kriegeskorte & Mur, 2012).

Video-specific adaptations We follow Lift-the-Weakest and introduce pragmatic modifications for

```
# Inputs: N, min_subset_size, max_subset_size,
     target_trial_time, alpha, tau, max_trials,
     time_budget, seed
rng <- init_rng(seed); E,W,Num,D_hat <- zeros(N,N);
     subsets <- []; t0 <- now();
viewed <- zeros(N) # 1 if inspected at least once
recent <- zeros(N) # light fatigue counter

for t in 1..max_trials:
  if min(E_ij) >= tau or now()-t0 >= time_budget:
       break

  (a,b) <- argmin_{i<j} E_ij
  S <- {a,b}
  while |S| < max_subset_size and est_cost(S) <=
       target_trial_time:
    x <- argmax_x ( gain(x|S,E) + eta*stress(x|S) +
         beta*(viewed[x]==0) - lambda*recent[x] ) /
         review_cost(x)
    if x == None: break
    S <- S  {x}
  if |S| < min_subset_size: S <- pad_with_cheapest(S)

  layout, meta <- present_and_collect(S) # meta: dmax
       , r_t, etc.
  dmax <- max_pairwise_distance(layout)
  for (i,j) in pairs(S):
    d <- clamp( d(i,j)/max(dmax,eps), 0, 1 )
    w <- meta.r_t * d^alpha
    E_ij += w; W_ij += w; Num_ij += w*d
    D_hat_ij <- Num_ij / max(W_ij,eps)

  mark_viewed(viewed, S)
  decay(recent); bump(recent, S)
  subsets.append(S)

return { E, D_hat, W, subsets }
```

Figure 3. Lift-the-Weakest scheduler. Each iteration anchors on the weakest pair $\arg\min E_{ij}$, expands with a cost-aware utility (gain, boundary stress, unseen boost, recency penalty), then updates evidence by max-normalizing within-trial distances and accumulating reliability-weighted contributions into E, W, and Num to maintain $\hat{D} = \text{Num}/W$. Returns E, $\hat{D}$, W, and the subset history.

dynamic video stimuli. Expansion uses a cost-aware utility that divides expected evidence gain by an estimated per-token review cost based on clip duration and recent behavior to keep per-trial time near a target. Subset size is chosen within bounds to satisfy a per-trial time target rather than a fixed k. A compliance-aware cold start boosts items not yet inspected and requires early subsets to include at least one unseen token. Light diversity constraints limit anchor reuse and cap the Jaccard overlap between successive subsets to reduce fatigue. Boundary-focused expansion prioritizes candidates that increase evidence around locally uncertain regions by combining low-evidence links with a simple local stress score. Evidence updating uses per-trial max normalization and a squared weight by default, with optional winsorization or Huber reweighting to reduce the influence of outlier placements, and a limited-iteration inverse MDS refinement is applied intermittently or at the end for stability (Kriegeskorte & Mur, 2012).

Validation

Protocol We compared multi-arrangement RDMs against one-by-one comparison RDMs within-subject on the 58-item set of natural human action videos, taken from a validated corpus for cognitive neuroscience (Urgen et al., 2023). Distances were normalized to [0,1] before analysis. Multi-arrangement trials followed a fixed set-cover schedule with $N = 58$ and batch size $k = 8$, guaranteeing at least one co-presentation of every unordered pair while keeping per-trial load bounded. Agreement was quantified using Pearson and Spearman correlations, the concordance correlation coefficient, Deming regression with 95% confidence intervals, RMSE and MAE, and Bland–Altman bias and limits of agreement. Summary metrics are reported in Table 1.

Interpretation Subject 1 shows high correlation and substantial concordance. Subject 2 shows moderate correlation and concordance. Deming slopes below 1.0 with small negative intercepts indicate that multi-arrangement yields slightly compressed dissimilarities relative to one-by-one at higher values, consistent with a mild range contraction in simultaneous layouts. Bland–Altman analysis shows a small positive bias for one-by-one over multi-arrangement with reasonably narrow limits of agreement. Overall, multi-arrangement reproduces the structure of

	Correlation and error metrics			
Subject	Pearson	Spearman	CCC	$\mathrm{RMSE}_{\mathrm{norm}}$
Subject 1	0.846	0.744	0.778	0.207
Subject 2	0.694	0.593	0.578	0.200
Mean	0.770	0.669	0.678	0.203
	Deming regression			
Subject	Slope	$\mathrm{CI}_{\mathrm{low}}$	$\mathrm{CI}_{\mathrm{high}}$	Intercept
Subject 1	0.934	0.914	0.956	−0.073
Subject 2	0.872	0.827	0.911	−0.024
	Bland–Altman			
Subject	Mean diff	$\mathrm{LoA}_{\mathrm{low}}$	$\mathrm{LoA}_{\mathrm{high}}$	
Subject 1	0.123	−0.204	0.450	
Subject 2	0.120	−0.175	0.415	

Table 1: Agreement between one-by-one and multi-arrangement dissimilarities.

the one-by-one RDM while using far fewer trials, enabling more context-aware, simultaneously referenced judgments through side-by-side comparison of multiple clips, and doing so without overwhelming participants (see Table 1 and Figure 4).

Usage of the Package

Installation and setup

Install from PyPI (Python 3.8+) and ensure common multimedia dependencies are available:

```
pip install multiarrangement
# or: uv pip install multiarrangement
```

The package bundles demo media (videos, images, audio), default instruction clips, and a cache of covering designs (LJCR) to enable offline use.

Preparing stimuli

Place your files (videos, audio, or images) in a single input folder. Type of modality is detected automatically, and requirements for submissions are adaptively set and mixed modalities are supported. The UI will show a confirmation if different types are detected. Filenames become token labels so that entries in the resulting RDM and csv/xlsx file can be tracked.

Fixed-batch (Set-Cover) Usage

The following example demonstrates how to use the Set-Cover mode in our API: first, precompute a fixed-k schedule of covering batches, seeding for reproducibility and using La Jolla coverings when available, then launch the session with that schedule to collect placements and automatically save the resulting RDM and logs. Batch size k can be tuned and shrink-only refinement can be enabled with `flex=True`, while `seed` guarantees the same schedule across runs. The run writes per-trial logs, coverage diagnostics, and the final RDM to the designated output folder to be used in visualization or downstream analysis.

```
import multiarrangement as ma

input_dir = "./videos" # your stimuli
output_dir = "./results"

# Build batches: uses LJCR cache when available, else
     greedy seed + refinement
batches = ma.create_batches(ma.auto_detect_stimuli(
     input_dir), k=8, seed=42, flex=False)
# flex=True enables shrink-only late blocks to k_min
     where helpful

res = ma.multiarrangement(
   input_dir=input_dir,
   batches=batches,
   output_dir=output_dir,
   fullscreen=False,
   language="en", # "tr" for Turkish or provide custom
         list below
   instructions="default", # or None, or a list of
        strings
   show_first_frames=True # show thumbnails for
        videos
)

res.vis(title="Set-Cover RDM")
res.savefig(f"{output_dir}/rdm_setcover.png", title="
     Set-Cover RDM")
```

Adaptive (Lift-the-Weakest) Usage

The following example demonstrates how to run the adaptive Lift-the-Weakest mode within the Multiarrangement library: Set subset size bounds and a target per-trial time or evidence threshold, optionally enable inverse MDS, then run to produce the evolving RDM, evidence matrix, and logs.

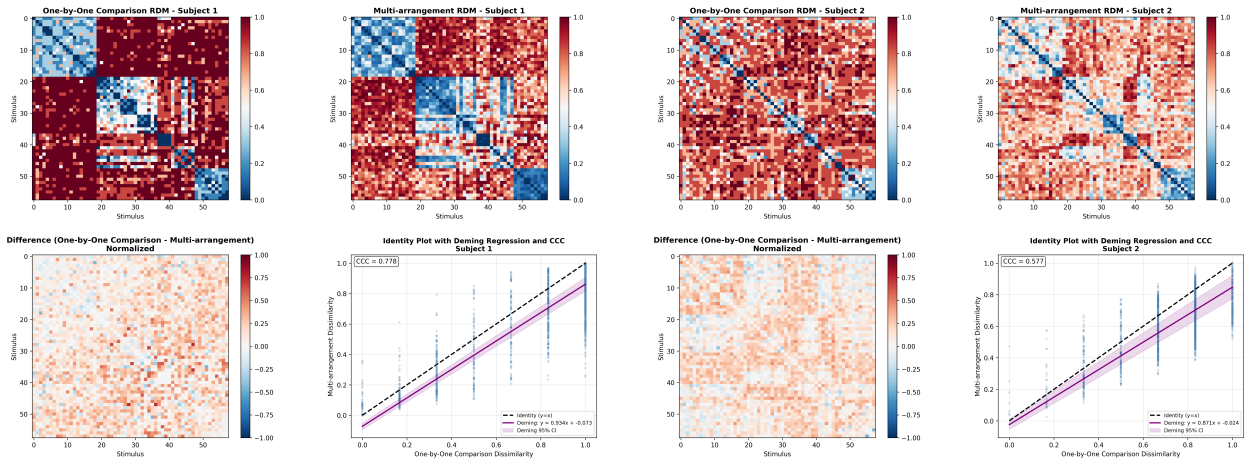

Figure 4. RDM agreement for two subjects. Panels correspond to Subject 1 (left) and Subject 2 (right). Within each panel: top row shows the one-by-one and multi-arrangement RDMs. Bottom left shows the normalized differences (one-by-one minus multi-arrangement) while the bottom right shows the identity plot with Deming regression and CCC.

```
import multiarrangement as ma

res = ma.multiarrangement_adaptive(
    input_dir="./videos",
    output_dir="./results",
    participant_id="S01",
    language="en",
    fullscreen=True,
    # key scheduling controls:
    evidence_threshold=0.35, # stop when min pair-
         evidence  threshold
    min_subset_size=4, max_subset_size=6,
    time_limit_minutes=None,
    # optional inverse MDS refinement:
    use_inverse_mds=True, inverse_mds_max_iter=15,
    inverse_mds_step_c=0.3, inverse_mds_tol=1e-4,
    instructions="default"
)

res.vis(title="Adaptive LTW RDM")
res.savefig("./results/rdm_adaptive.png", title="
     Adaptive LTW RDM")
```

Custom instructions and localization

Default instructions are provided in English and Turkish for convenience. To override, pass a list of strings to be shown as paginated screens before the task, per the following example:

```
custom = [
  "Welcome to the study.",
  "Drag each item into the white circle.",
  "Double-click a token to play/replay the video.",
  "Press SPACE to continue."
]
ma.multiarrangement(input_dir, batches, "./results",
     instructions=custom)
# ...or for adaptive:
ma.multiarrangement_adaptive("./videos", "./results",
     instructions=custom)
```

For further detailed documentation, see our GitHub page.[2] We release all code, experiment templates, and analysis scripts used in this paper.

Conclusions

In this paper, we introduce Multiarrangement, a turn-key, offline, open-source software package for collecting human similarity judgments with dynamic visual stimuli. The system unifies a simple, robust interface with a circular arena, on-demand playback, compliance cues, optional zoom, and bilingual de-

[2]GitHub: UYildiz12/Multiarrangement-for-videos

faults with scheduling back-ends that make large stimulus sets practical. A fixed Set-Cover mode produces deterministic, between-subject-comparable schedules with bounded per-trial load, while an adaptive Lift-the-Weakest mode targets the least-supported region of the space under explicit time and subset-size constraints. Across trials, distances are aggregated with per-trial max normalization and reliability weighting, and an optional inverse MDS refinement reconciles cross-trial scale and placement to yield a coherent RDM estimate. Our validation shows that the approach reproduces the structure obtained from one-by-one comparisons while using far fewer trials, enabling context-aware judgments without overwhelming participants.

Beyond video, the same workflow supports images and audio with no changes to the analysis path, and results are written in standard formats to facilitate downstream use in behavioral and neuroimaging pipelines. The package is designed to be easily extended. Researchers can customize instructional text, enforce additional compliance rules, adjust scheduler parameters, and integrate new quality-control or weighting schemes while preserving reproducibility via recorded seeds and schedules. We hope this plug-and-play framework lowers the barrier to deploying multi-arrangement studies at scale and stimulates new experimental designs that leverage simultaneous comparison to probe perceptual and cognitive structure in richer, more naturalistic settings.

Authors' contributions Both authors jointly conceptualized the study. U.Y. implemented the software, conducted testing & validation, and collected the data. Both authors co-wrote the first draft, critically revised subsequent drafts, and approved the final manuscript for submission.

Funding This research received no specific grant from any funding agency, commercial or not-for-profit sectors.

Availability of data and materials All datasets analyzed in this article are included in the reproducibility archive at 10.5281/zenodo.17329545.

Code availability All analysis and figure-generation code is archived at 10.5281/zenodo.17329545. Reproduce the figures/tables offline via: `python reproduce.py`.

Declarations

Competing interests The authors declare no competing interests.

Ethics approval The study was approved by, and conducted along the guidelines of, the Bilkent University Ethics Committee.

Consent to participate All participants provided written informed consent prior to participation. Data were analyzed anonymously.

Consent for publication All participants provided written informed consent for publication of anonymized data.

References

Baldassano, C., Chen, J., Zadbood, A., Pillow, J. W., Hasson, U., & Norman, K. A. (2017). Discovering event structure in continuous narrative perception and memory. *Neuron*, *95*(3), 709–721.e5. https://doi.org/10.1016/j.neuron.2017.06.041

Baveye, Y., Dellandréa, E., Chamaret, C., & Chen, L. (2015). LIRIS-ACCEDE: A video database for affective content analysis. *IEEE Transactions on Affective Computing*, *6*(1), 43–55. https://doi.org/10.1109/TAFFC.2015.2396531

Finn, E. S., Glerean, E., Khojandi, A. Y., Nielson, D., Molfese, P. J., Handwerker, D. A., & Bandettini, P. A. (2020). Idiosynchrony: From shared responses to individual differences during naturalistic neuroimaging. *NeuroImage*, *215*, 116828. https://doi.org/10.1016/j.neuroimage.2020.116828

Gilman, T. L., Shaheen, R., Nylocks, K. M., Halachoff, D., Chapman, J., Flynn, J. J., Matt, L. M., & Coifman, K. G. (2017). A film set for the elicitation of emotion in research: A comprehensive catalog derived from four decades of investigation. *Behavior Research Methods*, *49*(6), 2061–2082. https://doi.org/10.3758/s13428-016-0842-x

Gross, J. J., & Levenson, R. W. (1995). Emotion elicitation using films. *Cognition & Emotion*, *9*(1), 87–108.

Hanke, M., Baumgartner, F. J., Ibe, P., Kaule, F. R., Pollmann, S., Speck, O., Zinke, W., & Stadler, J. (2014). A high-resolution 7-tesla fMRI dataset from complex natural stimulation with an audio movie. *Scientific Data*, *1*, 140003. https://doi.org/10.1038/sdata.2014.3

Hasson, U., Landesman, O., Knappmeyer, B., Vallines, I., Rubin, N., & Heeger, D. J. (2008). Neurocinematics: The neuroscience of film. *Projections*, *2*(1), 1–26. https://doi.org/10.3167/proj.2008.020102

Hasson, U., Nir, Y., Levy, I., Fuhrmann, G., & Malach, R. (2004). Intersubject synchronization of cortical activity during natural vision. *Science*, *303*(5664), 1634–1640.

Jääskeläinen, I. P., Sams, M., Glerean, E., & Ahveninen, J. (2021). Movies and narratives as naturalistic stimuli in neuroimaging. *NeuroImage*, *224*, 117445. https://doi.org/10.1016/j.neuroimage.2020.117445

Koelstra, S., Mühl, C., Soleymani, M., Lee, J.-S., Yazdani, A., Ebrahimi, T., Pun, T., Nijholt, A., & Patras, I. (2012). DEAP: A database for emotion analysis using physiological signals. *IEEE Transactions on Affective Computing*, *3*(1), 18–31. https://doi.org/10.1109/T-AFFC.2011.15

Kriegeskorte, N., & Mur, M. (2012). Inverse MDS: Inferring dissimilarity structure from multiple item arrangements. *Frontiers in Psychology*, *3*, 245. https://doi.org/10.3389/fpsyg.2012.00245

Kriegeskorte, N., Mur, M., & Bandettini, P. A. (2008). Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in Systems Neuroscience*, *2*, 4. https://doi.org/10.3389/neuro.06.004.2008

Lerner, Y., Honey, C. J., Silbert, L. J., & Hasson, U. (2011). Topographic mapping of a hierarchy of temporal receptive windows using a narrated story. *Journal of Neuroscience*, *31*(8), 2906–2915. https://doi.org/10.1523/JNEUROSCI.3684-10.2011

Liu, X., Dai, Y., Xie, H., & Zhen, Z. (2022). A studyforrest extension, MEG recordings while watching the audio-visual movie Forrest Gump. *Scientific Data*, *9*, 206. https://doi.org/10.1038/s41597-022-01299-1

Nastase, S. A., Goldstein, A., & Hasson, U. (2020). Keep it real: Rethinking the primacy of experimental control in cognitive neuroscience. *NeuroImage*, *222*, 117254. https://doi.org/10.1016/j.neuroimage.2020.117254

Nili, H., Wingfield, C., Walther, A., Su, L., Marslen-Wilson, W., & Kriegeskorte, N. (2014). A toolbox for representational similarity analysis. *PLOS Computational Biology*, *10*(4), e1003553. https://doi.org/10.1371/journal.pcbi.1003553

Schaefer, A., Nils, F., Sanchez, X., & Philippot, P. (2010). Assessing the effectiveness of a large database of emotion-eliciting films: A new tool for emotion researchers. *Cognition & Emotion*, *24*(7), 1153–1172. https://doi.org/10.1080/02699930903274322

Sonkusare, S., Breakspear, M., & Guo, C. (2019). Naturalistic stimuli in neuroscience: Critically acclaimed. *Trends in Cognitive Sciences*, *23*(8), 699–714. https://doi.org/10.1016/j.tics.2019.05.004

Uhrig, L., Janssen, D., Dehaene, S., & Jarraya, B. (2016). Cerebral responses to local and global auditory novelty under general anesthesia. *NeuroImage*, *141*, 326–340. https://doi.org/10.1016/j.neuroimage.2016.08.004

Urgen, B. A., Nizamoğlu, H., Eroğlu, A., & Orban, G. A. (2023). A large video set of natural human actions for visual and cognitive neuroscience studies and its validation with fmri. *Brain Sciences*, *13*(1), 61. https://doi.org/10.3390/brainsci13010061