Psychological Review

The Format of Representation for Objects in Orbit: Composite and Internally Referenced

Qihan Wu and Jonathan I. Flombaum

Online First Publication, May 18, 2026. https://dx.doi.org/10.1037/rev0000629

CITATION
Wu, Q., & Flombaum, J. I. (2026). The format of representation for objects in orbit: Composite and internally referenced. *Psychological Review*. Advance online publication. https://dx.doi.org/10.1037/rev0000629

AMERICAN PSYCHOLOGICAL ASSOCIATION

Psychological Review

© 2026 American Psychological Association
ISSN: 0033-295X

https://doi.org/10.1037/rev0000629

The Format of Representation for Objects in Orbit: Composite and Internally Referenced

Qihan Wu and Jonathan I. Flombaum
Department of Psychological and Brain Sciences, Johns Hopkins University

When objects move on circular trajectories that could be described as orbiting, are they represented independently, as when translating linearly, or as parts of a composite system, akin to elements visibly affixed to a rotating structure? This fundamental question matters because representing a satellite as part of a rotating system implies that perception constructs composite representations, including internal references to anchor the system to the external present. We tested this composite representation hypothesis using a tracking task, predicting greater difficulty for rotational compared to translational motion along similar trajectories. Composite representations require additional computational steps to integrate internal and external references, unnecessary steps for independent tracking through translation. Two identical disks were placed on opposite corners of an invisible square, with another pair arranged similarly on a larger concentric square. Participants tracked one designated disk from each pair as they rotated or translated. Performance was significantly worse for rotation. A second experiment showed that visible lines connecting paired disks impaired translational tracking, demonstrating that composite representations can be induced both by visible connections and by the yoking inherent to rotation. Three additional experiments replicated these findings and revealed increased error rates on trials which demanded greater reexpression of relational variables. A sixth experiment extended the reasoning to the freezing illusion, introducing a novel method to quantify illusion strength. Rotation produced a stronger illusion than matched translation. Our findings indicate that principles of object-centered representation apply broadly in visual perception, including axes and relational variables not physically present in the stimulus.

Keywords: representation, motion perception, visual object, grouping

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

We investigate the hypothesis that a circular orbit is perceived like an object, meaning that it automatically erects a representation with its own internal reference frame. An object that moves along the orbit ceases to be an independent entity, transforming instead into a satellite that is part of the larger structure. For a satellite, the primary descriptors of position are supplied relative to an internal reference frame. For example, in Figure 1A, the blue satellite would be represented as occupying the 90° position *of the orbit*. In such a scheme, externally referenced position is derivable from stored and updated descriptors of how the whole orbit is situated relative to an external reference. In the same figure, descriptions of the orbit could include the kinds of information shown in Figure 1B: The orbit's center point is 20 m east of the external origin and 20 m south. The orbit radius is 10 m long, and the orbit's internal axis is parallel to the external north–south axis with its own positive side pointing directly north. Combining all of this with the blue satellite's location *within* the orbit, one can derive that the satellite currently occupies a position 30 m directly east and 20 m south of the external origin. But operations are required to reach the conclusion.

In such a scheme, the motion of a satellite relative to an external reference frame is assumed by the reorientation of the orbit's axes relative to external ones. As shown in Figure 1C and 1D, for example, one could describe an orbit as a circle with an internal axis that begins in parallel alignment to the external vertical axis. Then the alignment changes so that it takes on a 45° clockwise rotation relative to the external vertical, later a 90° rotation, and so on. If the blue satellite remains in the 90° position of the orbit—*that is, relative to the internal axis*—then it must be moving relative to the external frame, in lockstep with the description offered for the orbit as a whole. (An animated demonstration of an orbit like this together with its composite representation can be viewed at: https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_1.html.)

In theory, the representational scheme could be inverted (Tadin et al., 2002). Priority could be awarded to externally referenced positions and movements of what we are calling "a satellite." In Figure 1A, the primary description would be something more like, the blue object is at (30, 10). As the object moves, the representation of its coordinates can be updated accordingly, now it is at (20, 30), now at (10, 20), and so on. That "in orbit" correctly describes the object's motion could be derived secondarily from observed positional changes that trace a circle. Further descriptions of the orbit's properties—that the radius of the orbit is 10 m, for example—may emerge automatically or as needed, but crucially, through derivation from data whose native reference is external to the orbit.

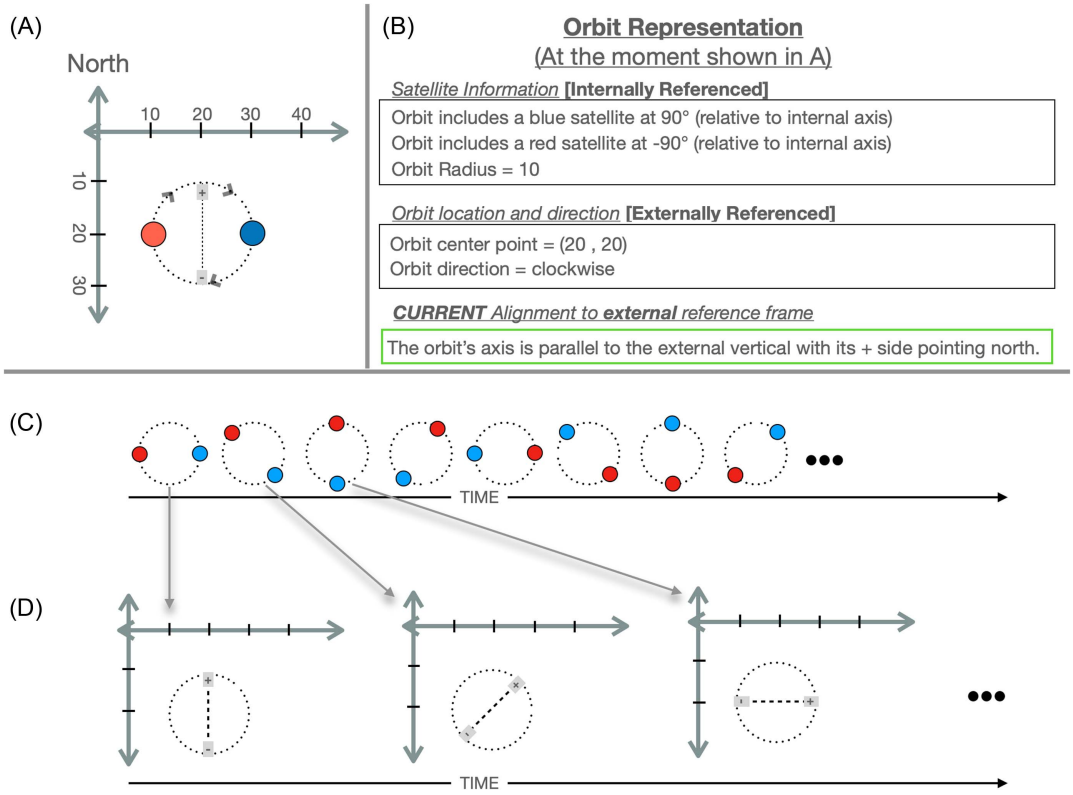
Hongjing Lu served as action editor.

Qihan Wu played a lead role in formal analysis, investigation, visualization, and writing–original draft and an equal role in conceptualization and methodology. Jonathan I. Flombaum played a lead role in writing–review and editing and an equal role in conceptualization and methodology.

Correspondence concerning this article should be addressed to Jonathan I. Flombaum, Department of Psychological and Brain Sciences, Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218, United States. Email: flombaum@jhu.edu

Figure 1

A Schematic Account of an Orbit Representation



Note. (A) An orbit that includes two satellites (blue and red) moving around a shared center point in a clockwise direction. (B) An orbit-centered description of the moment shown in (A). As time progresses (shown in C), continuing to represent the complete state of the system only requires updates to the information in the green box, which describes the alignment between the orbit-centered axis and the external reference frame. A subset of these alignment changes is represented pictorially in Panel D, beginning with the moment in time from Panel A, followed by a 45° rotation of the internal axis from the external north, and then a 90° rotation. An animated demonstration of an orbit like this together with its composite representation can be viewed at: https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_1.html. See the online article for the color version of this figure.

The hypothesis we investigate therefore includes two components: First, that moving around a circle automatically induces a representation of a higher order structure, an orbit. Any satellite thereby becomes a part of the system. And second, that this representation supplies the native reference frame to describe the motions and positions of satellites. Representing any satellite as an independent object with external reference is secondary. This is what we mean when we say that "a circular orbit is perceived like an object."

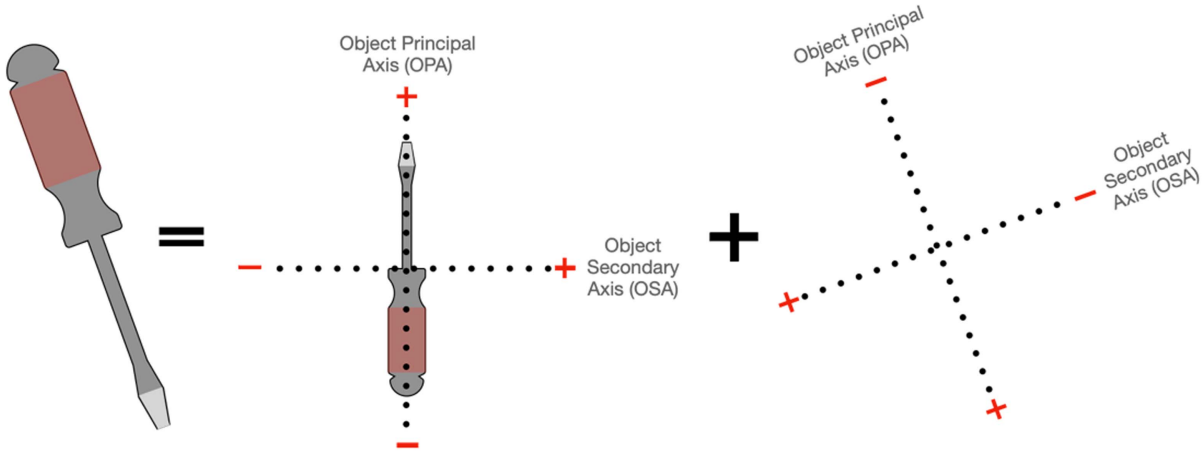
Research has suggested that object representations can reduce access to individual parts, for example, in faces (Murphy et al., 2017), or for an upright walking person compared to an upside down one (Poljac et al., 2012; see also Wu & Flombaum, 2024). An aspect our claim will be that in an orbit system, what becomes less accessible are the coordinates of its parts with respect to an external reference frame, owing to prioritization of internally referenced descriptions. As we have argued with respect to a related study (Wu & Flombaum, 2024), this could be a proximate cause for how and why wholes generally reduce access to parts.

In a proper object, parts occupy inherently relational roles, both conceptually and in visual representation. Research employing neuropsychological case studies (McCloskey et al., 2006), working memory paradigms (Gregory & McCloskey, 2010), and object-based attention paradigms (McCloskey & Rapp, 2000) has shown that the positions and orientations of parts in an object are represented with priority assigned to internal reference. For an elongated object such as a screwdriver, for example, the straight extent of the tool supplies an obvious internal axis, an object principal axis (OPA). It has been suggested that representing a tool like this further includes a secondary axis (object secondary axis) which perpendicularly bisects the OPA (Figure 2). Note that the object secondary axis is virtual, imposed in the representation and not present in the visual stimulus. But it allows for a wholly coherent description of the parts of the tool without reference to anything external to itself: The handle resides at the end of the principal axis *on the negative side*, with the "business end" residing on the positive end of the principal axis—all regardless of how the tool may sit currently. That a screwdriver can be held upside down makes sense because its internally referenced upside can be oriented in different ways relative to an external reference (Chaisilprungraung et al., 2019; see also McCloskey, 2009). Another way to state the

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 2

A Composite Representation of a Screwdriver



Note. A complete representation of a screwdriver at a given moment is composed from an object-centered representation (here with the handle on the negative end of the object's principal axis), combined with a description of how the object axes are oriented relative to an external reference in the present moment. One feature of such a representation is that it allows for the conception of an object as currently "upside down." (See McCloskey, 2009; McCloskey et al., 2006, for a detailed articulation of this representational framework.) See the online article for the color version of this figure.

hypothesis that an orbit is represented like an object is, therefore, that an orbit is a sort of thing that can be upside down.

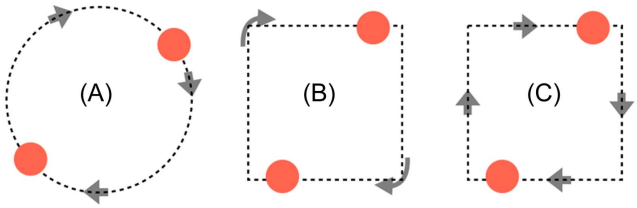
The hypothesis that an orbit is like an object in these ways may seem counterintuitive. Sure, a satellite's trajectory can be described as an orbit, but why subordinate the satellite to the trajectory, and why treat the trajectory as an underlying structure? We assume, in general, that efficiency considerations should guide the prioritization of one representational format or another. An advantage for orbit-centered representation is the ease of comparing relative satellite motions and positions. Returning to Figure 1, the red satellite and the blue one never move relative to one another, although they are both in motion relative to the external world. With an orbit-centered representation, like the one in Figure 1C, their relative stasis is obvious, and their externally referenced movements are assumed under the reorientation of the internal axis relative to the external ones. More broadly, by imposing an internal reference on a system, the relative arrangements of the parts become explicit, a paradigmatic instance of a point made frequently and clearly by Marr (1982) that "how information is represented can greatly affect how easy it is to do different things with it" (p. 21). Without internal reference, keeping track of relative positioning requires additional steps. Prioritizing internal reference is a way of prioritizing access to relative information about the parts of a system, at the expense of external relationships, and vice versa. This is perhaps why visual judgments about the alignment of independent objects can be difficult: for example, when placing nails in a wall for hanging pictures. Imposing internal reference on a system can simplify computation when relative information is desired. (Try tying a string tautly between the nails, and now a single judgment of parallelism will reveal whether they are aligned.)

Consider as well that an internally referenced representation may seem less counterintuitive in the case of a rotating circular object—a wheel, a gear, or a record, for example—because in these cases the represented axes have material manifestations. We propose that this same kind of representation—one that is wheel-like—will emerge and take precedence during the perception of objects in a circular orbit. Crucially, this happens absent the visible rotation of an underlying surface, and unlike in elongated objects, absent a physical manifestation for any of the represented internal axes. Orbiting, on this view, is perhaps a paradigmatic case of perception that is not direct (Gibson, 1979), with the representation of the stimulus imposing layers of structure over and above the raw input data.

We tested the hypothesis via a prediction: That it should be difficult to track the individual identities of satellites in the same orbit compared to tracking the identities of objects that move around a noncircular trajectory. Figure 3A helps to explain the prediction. (Animated demonstration two is also provided to assist: https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_2.html). The dotted line marks an orbit shared by the two red disks. Each of the disks is a satellite that moves around the same origin. Suppose an observer is asked to track one of the two disks (a target) and to identify the tracked individual after a period of motion. Because they are identical in terms of shape and color, tracking the disks must rely on continually updated representations of their positions (Pylyshyn & Storm, 1988). If the primary representation of the target and nontarget marks them as static relative to one another—each of them permanently affixed to a location on the perimeter of an orbit—then tracking a specific satellite must follow updates about how the whole system reorients relative to the external origin. The secondary and derived nature of target identity under these constraints should add difficulty to accurate tracking. If, however, the two targets do not rotate around a shared orbit, but move clearly as independent objects, tracking one should be relatively easy, a task that amounts to iterated updates of the object's coordinates relative to an external origin (Upadhyayula & Flombaum, 2020). Our prediction therefore follows from the observation that a simple tracking task does not demand a relative assessment of target locations compared to nontargets. It requires only that an observer accurately update the absolute

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 3
Stimuli for Comparison of Tracking During Rotation Versus Translation



Note. (A) Two featurally identical satellites (red disks) share an orbit. Before motion onset, one was marked (in some way) as a target. The primary experimental prediction is that tracking in this case should be difficult compared to a comparable condition where the motion would not be described as a rotational orbit. Placing the target and nontarget on the perimeter of an invisible square, as in (B) and (C), allows for a controlled comparison of the desired nature. In (B), the whole square rotates around its center, including the tracking target and nontarget that are affixed to its perimeter. When the square perimeter is not shown, the percept is of satellites on a circular orbit, a condition that should challenge accurate tracking compared to the condition shown in (C), where the disks translate around the square perimeter. See also: https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_2.html. See the online article for the color version of this figure.

(i.e., externally referenced) position of the target. But if the stimulus somehow biases representation to employ internal reference nonetheless, then tracking performance should suffer.

To compare such conditions with better control, we placed targets and nontargets on the perimeters of invisible squares (Figure 3B and 3C). When the invisible squares rotated—spun around the center point of the square (Figure 3B)—we could ask participants to track a rotating target in a shared orbit with a nontarget. A square also allows for nonrotational motion around a perimeter. The experiments therefore included conditions in which a target and nontarget translated around the perimeter of an invisible square (Figure 3B). The key prediction was that tracking performance should be better when targets translate (around a square) compared to when they orbit (affixed to locations on the same virtual square), owing to the cost of orbit-centered representation for the localization of an individual.

Experiment 1: Deficient Tracking of Identity in a Rotational Orbit

We sought to test the prediction that it should be difficult to track the individual identities of satellites in the same orbit compared to tracking the individual identities of objects that move around a noncircular shared trajectory. Because tracking experiments often involve rotational motions and orbits, we began by identifying an extant paradigm to modify. Holcombe and Chen (2012) employed concentric orbits in a basic tracking task. We modeled our experiments closely to theirs by defining two homocentric virtual squares to build trials upon. To each square, we affixed two red disks opposite one another. At the start of a trial, one disk from each pair turned white, identifying it as a target. In half the trials, the virtual squares rotated. This created the impression of concentric orbits. In half the trials, the disks translated along the virtual square perimeters. We measured tracking accuracy with a probe-one-task (as described below; cf. Holcombe & Chen, 2012). We predicted better performance in the translation motion trials compared to the rotation motion trials.

Method

Preregistration

We report how we determined our sample size, all data exclusions (if any), all manipulations, and all measures in the study. The methods for this experiment were preregistered. Registration can be accessed through the following link: https://aspredicted.org/yzfd-snmz.pdf. Links to all preregistrations, online experiments, and to open data and code repositories can be found in Appendix A (Wu & Flombaum, 2026). Appendix B includes links and descriptions for demos provided for explanatory purposes.

Experiment Link

Experiment 1 can be viewed online with a web browser at: https://qihan.visualthinkingresearch.com/tracking_COR/exp_1/tracking_exp_1.html.

Participants

We aimed to collect data from 30 participants. The sample size was determined based on pilot experiments conducted before preregistration. Any participant with a correct response rate below 60%, overall, was excluded and replaced. Such data were discarded automatically and never viewed. The same criterion is applied to all subsequent experiments. We tested a total of 35 undergraduates for Experiment 1, with 5 excluded as described.

Participants were recruited through an online system for students to enroll in studies and receive course-related credit. Demographic information was not collected since it is not related to the hypothesis. Participants completed the experiment on a personal device. The instructions and sign-up page asked that they use a laptop or desktop computer, not a mobile device. All reported experiments were approved by the Johns Hopkins University Homewood Institutional Review Board.

Stimuli and Procedure

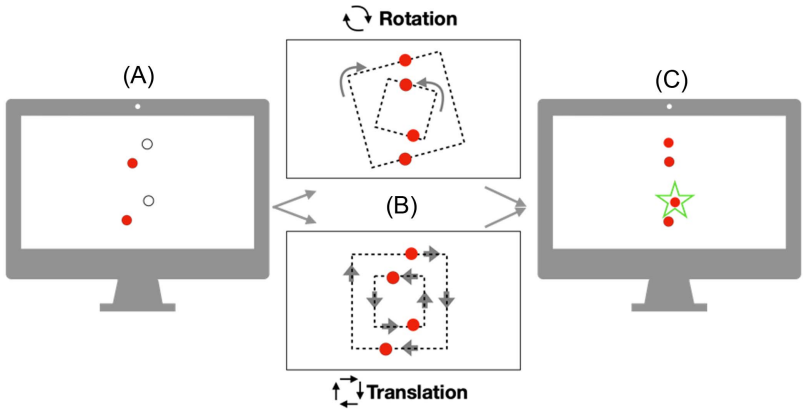
In this and all subsequent experiments, parameters are reported in pixels, because the experiments were run online. The task and stimuli were derived from a study by Holcombe and Chen (2012). Each trial included two invisible squares that shared one origin. The squares are designated in Figure 4B by dotted lines *which were not visible in the actual experiment*. The smaller (inside) square had a side length of 140 pixels. The side length of the larger (outer) square measured 380 pixels.

Each participant completed 160 trials. To begin a trial of any kind, two pairs of disks appeared in the display, along with a fixation point in the center. The disks appeared on a gray background. One pair of disks always occupied the perimeter of the smaller square, always appearing at the start of the trial on the opposite upper and lower walls of the smaller square. The position of the disk on the top edge was selected randomly along the length of the edge, with the lower disk occupying the position of a reflection of the upper one across the center point.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 4

Schematic Depiction (Not to Scale) of Methods Used in Experiment 1 and Forming the Basis for Subsequent Experiments



Note. (A) Every trial starts in the same way, with four disks each affixed in opposite-location-pairs to the perimeter of one of two virtual squares (shown in B). One disk in each pair appears white to mark it as a target for tracking. The white disks then turn red, so that all the disks look identical, and then (B) a period of motion and tracking begins. In half the trials, the motion occurs through the rotation of the virtual squares (each square in the opposite direction of the other). In the remaining half of trials, the disks on each square translate in synchrony along the virtual perimeters (the disks on one square moving in the opposite direction from the disks on the other). Finally, all trials end in the same way (C): motion ceases, and one of the four disks is highlighted (here with a green star for clarity but with a simple circular outline in practice). Participants indicate by keypress if the highlighted disk was a tracked target or not. See the online article for the color version of this figure.

The second pair of disks always occupied the perimeter of the outer and larger invisible square. For these disks, the position of one was selected randomly along any point on the square's perimeter. The position of the second was always a reflection of the first across the square's center.

At the beginning of a trial, all four disks were stationary for 1,500 ms. Two disks (one from each pair) were white to designate them as the tracking targets. After the 1,500-ms stationary period, the two white disks turned red so that all the disks were visually identical. Now the disks moved for a duration of 2,300 ms. When motion ceased, one disk was probed through the appearance of a white circle around its perimeter. Half of the time the probed item was a former target, and half of the time it was not. Half the time, the circle probed was on the outer square, and half the time it was on the inner one. Participants were instructed to respond whether the probed disk was a target or not.

Trials included two motion types, randomly and equally distributed throughout a testing session. In the rotation condition, the disks rotated as pairs, in synchrony, around the screen center point, defining concentric circular orbits with diameters equal to the distances between the two disks in each pair. The initial rotation direction for the inner pair and the outer pair was always in the opposite direction. Direction changes took place simultaneously for the pairs of disks every 600 ms during the motion period. To make this concrete: if the inner pair started by rotating clockwise, the outer pair would rotate counter. They switched every 600 ms.

In the translation condition, the disks in each pair moved along the perimeters of their respective squares, turning corners as necessary. The top disk of the inner pair always began by moving to the left (equivalently, counterclockwise), with the bottom disk moving in the opposite direction. For the disks on the outer square motion direction was randomly up or down, or left or right, depending on the edge the disks were on, but always in directions that could be described as counter to the motion of the inner square. Again, direction changes took place every 600 ms.

Motion speeds were set as follows. The angular speed for all disks in rotation trials was 576°/s. This meant that the number of radians a disk traveled per second varied somewhat as a function of where the disk happened to start along a square's edge. (i.e., A pair of disks that started near corners traced a larger circle than a pair that started near the middle of an edge.) The minimum radian speed for a rotating pair was equal to half the length of the respective (inner or outer) square multiplied by 576. The maximum was the same multiplied by the square root of two. We took the average of these two values (respectively, for each virtual square) to set the linear motion speed for disks in the translation condition.

Reaction Time and Accuracy Validation

Results were computed while excluding any individual trial with a reaction time less than 200 ms or greater than 3,000 ms. 200 ms is a typical cutoff time in such experiments, close to the time required to execute a keypress. 3,000-ms (or greater) trials were excluded based on pilot experiments before preregistration, which suggested that some subjects might step away from the task for long durations, a challenge of online experiments. In addition to these criteria, a programming glitch resulted in a small number of trials for which responses were not recorded. This affected eight participants as

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

follows: 26 trials were lost for one of the eight; two trials for two of the eight; three trials for three of the eight; and one trial for two of the eight. All together we analyzed results from 4,603 trials in Experiment 1.

Results

Tracking Accuracy

A planned one-tailed *t* test showed significantly better performance in the translation condition ($M = 87\%$) compared to rotation (84%), $t(29) = 2.47$, $p = 9.66 \times 10^{-3}$, $d = 0.45$ (Figure 5A). Figure 5B plots the accuracy difference between translation and rotation trials on an individual participant basis.

Discussion

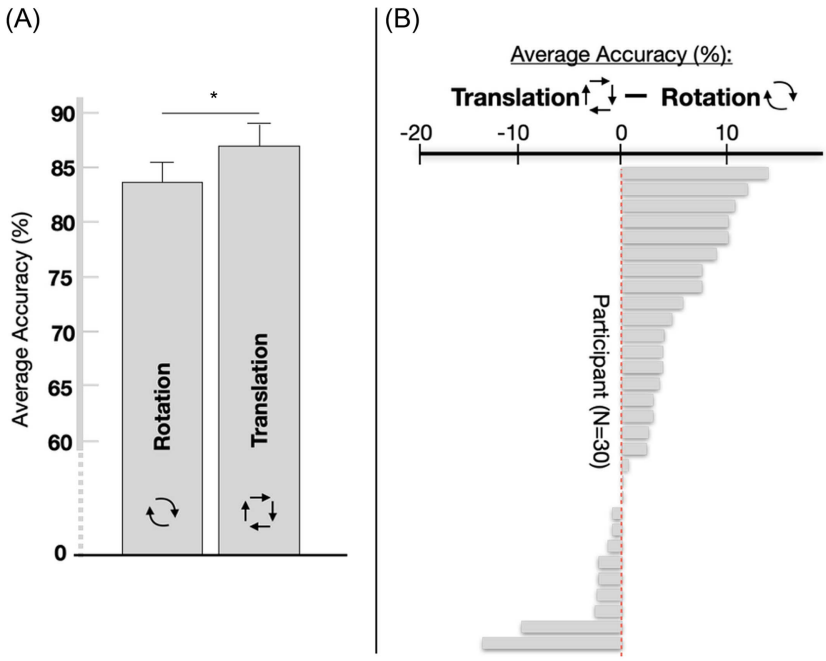
Results confirmed the prediction that tracking objects within circular orbits is more difficult than tracking in an equivalent but nonrotational condition. We interpret the deficiency for orbit tracking as evidence of an object-like representation, with a prioritized internal reference frame.

Experiment 2: Deficient Tracking of Proper Objects

Our reasoning demands that if an object-like representation impairs the tracking of an element, then an object representation should do the same. Experiments with a more typical multiple-object paradigm are consistent with this expectation, having shown deficient tracking when a visible connection conjoins target and distractor pairs as parts of the same object (Scholl et al., 2001). In that study, participants tracked four squares in a group of eight, all of them identical, and each moving independently on unpredictable trajectories occupying the entire display (Figure 6). In some trials, each target was conjoined with a nontarget via a thin line that compressed or elongated as the items moved, so that the same exact trajectories could be used in connected and unconnected trials. In theory, adding such connections should have improved tracking performance. Being certain about the identity of any one target should exclude its conjoined pair as a target candidate. But conjoining items in this way made tracking performance *worse*.

The authors explain that they sought to determine if the "undifferentiated parts or ends of objects can be independently selected and tracked over time." The answer was clearly that parts are difficult to track independently. We interpret the results of Experiment 1 similarly, on the logic that satellites are represented as parts of an object-like system, thereby rendered less accessible to tracking as independent entities. Our results may at the same time shade some light on the previous ones. *Why* is it difficult to *independently* track the undifferentiated parts of objects? We have suggested that it is because the locations of the parts are addressed with object-internal reference. With respect to task demands such representations are defiantly *dependent*, requiring mediation to uncover the location of a target in externally referenced terms.

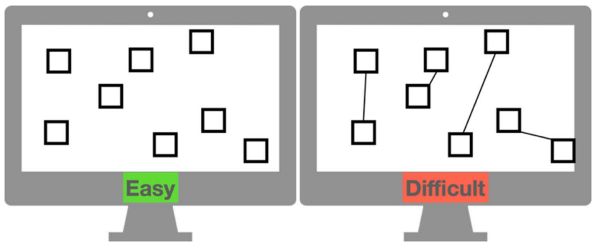
Figure 5
Results of Experiment 1



Note. (A) Average tracking accuracy was significantly worse in the rotation trials compared to the translation trials. Error bars show standard errors. (B) Average accuracy in translation trials minus rotation trials for each participant ($N = 30$). See the online article for the color version of this figure.
* $p < .05$ (for a planned comparison).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 6
Schematic Depiction of Two Conditions From a Study by Scholl et al. (2001)



Note. Participants were asked to track four out of eight identical squares, all of which moved independently through the display. At the end of a trail, participants identified what they believed were the original targets. In one condition, each target was conjoined to a nontarget, creating pairs, via a thin line. With everything else left unchanged in terms of trajectories, tracking performance declined compared to the unpaired condition. See the online article for the color version of this figure.

To allow for more direct comparison of orbit-based and object-based tracking constraints, Experiment 2 therefore sought to confirm that object parts would be difficult to track individually in the exact paradigm devised for Experiment 1. To meet this objective, we added visible connections between the disks in each tracked pair (Figure 6), creating dumbbells, but otherwise adopting the same paradigm as in Experiment 1. A uniform and continuous connection will typically elicit a representation of parts belonging to the same object (Behrmann & Tipper, 1994, 1999; Watson & Kramer, 1999). We predicted that visible connections between disk pairs should impair tracking when the disks translate, compared to translation without visible connections (i.e., the translation condition of Experiment 1). Put slightly differently, it should be harder to track the translating target end of a dumbbell than to track under the same conditions, but without a visible line to pair disks as dumbbells. Experiment 2 included trials with connected pairs, as shown in Figure 7, as well as unconnected trials, replicating those from Experiment 1 (and shown in Figure 4).

Method

Preregistration

Registration can be accessed through the following link: https://aspredicted.org/29jf-y338.pdf.

Experiment Link

Experiment 2 can be viewed online with a web browser at: https://qihan.visualthinkingresearch.com/tracking_COR/exp_2/tracking_exp_2.html.

Participants

We aimed to collect data from 30 participants. We tested a total of 48 undergraduates for Experiment 2. Eighteen were excluded automatically because of low overall performance, employing our preregistered exclusion criteria exactly as in Experiment 1.

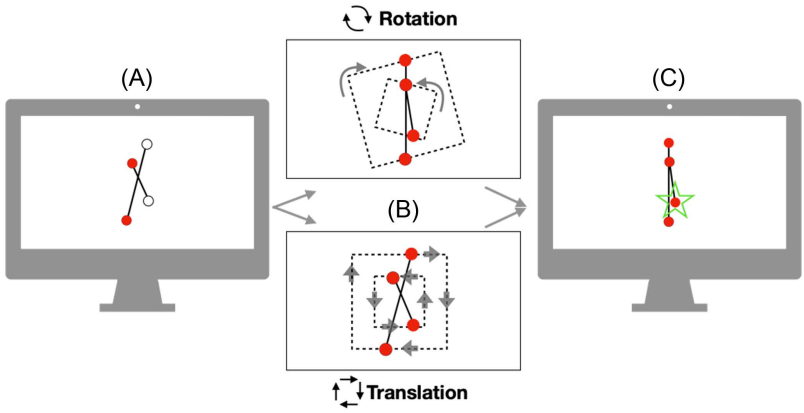
Stimuli and Procedure

Each participant completed 180 trials. Half of the trials were identical to the trials in Experiment 1. The remaining half differed from Experiment 1 in only one way. Each pair of disks was connected by a 3-pixel width black line, joining them from center to center during each frame of a given trial.

Reaction Time and Accuracy Validation

We analyzed 5,189 trials that met our inclusion criteria.

Figure 7
Schematic Depiction (Not to Scale) of Methods Used in Experiment 2



Note. Methods were identical to Experiment 1, including a target marking phase (A) and a motion period (B) followed by target identification (C). The only deviation involved the presence of a visible line connecting each pair of disks. Experiment 2 also included unconnected trials, exactly like those from Experiment 1. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Results

Tracking Accuracy

Replicating the results of Experiment 1, a planned one-tailed t test showed significantly better performance in the translation condition (90%) compared to the rotation condition (88%) when the disks were not visibly connected, $t(29) = 2.00$, $p = 2.73 \times 10^{-2}$, $d = 0.37$. But when the disks were connected, the same test was not significant (both means were 85%), $t(29) = 0.32$, $p = .38$, $d = 5.87 \times 10^{-2}$. Performance in each of the unconnected conditions was significantly better than in each connected counterpart, for rotation: $t(29) = 2.41$, $p = 1.11 \times 10^{-2}$, $d = 0.44$. For translation: $t(29) = 3.13$, $p = 2.00 \times 10^{-3}$, $d = 0.57$. Results are summarized in Figure 8.

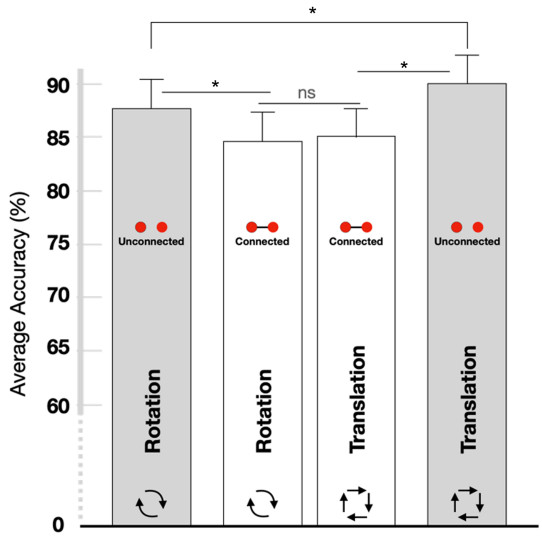
A repeated measure two-way analysis of variance between motion type and connection type did not show a significant main effect of motion type, $F(1, 29) = 2.24$, $p = .15$, but it did show a significant main effect of connection type, $F(1, 29) = 12.53$, $p = 1.38 \times 10^{-3}$, without a significant interaction between the two factors, $F(1, 29) = 1.16$, $p = .29$.

Figure 9 plots, on an individual participant basis, what we see as the key comparison of Experiment 2: average tracking performance for unconnected translation trials minus connected translation trials. 22/30 participants showed a positive difference—better performance for unconnected compared to connected pairs with translating motion.

Discussion

Results confirmed the prediction that elements in a visibly connected pair are more difficult to track than elements engaged in identical motion but without a visible connection between them. Especially noteworthy is the fact that a visible connection impaired performance for translating objects not just for rotating ones. This is consistent with the hypothesis that visible connection induces a whole object representation, one that demands reference frame management regardless of motion. A consolidated interpretation of Experiments 1 and 2 is that what it means—in terms of visual representation—to be a part of something is to have an address that is linked to internally referenced coordinates so that uncovering a part's independent location in external space demands computation. Both orbits and objects are systems with internal reference for their parts.

Figure 8
Results of Experiment 2



Note. Average tracking accuracy was significantly better for Unconnected compared to connected stimuli. Error bars show standard errors. ns = not signficiant. See the online article for the color version of this figure.
* $p < .05$ (for a planned comparison).

Supplementary Data Analysis: Relative Starting Positions in Experiments 1 and 2

We conducted secondary analyses of the results of Experiment 1 to explore the hypothesis that the cause of deficient tracking under rotation resides in the primacy of internally referenced, orbit-centered representation. The analyses derive from our understanding of how the targets, pairs, and orbits are represented. We therefore review in brief.

For the orbiting pairs in Experiment 1, our main proposal is that the representation of each target amounts to a description with the following pieces of information included.

(a) Which item within each pair is the target?

One could answer the question by saying something like, "It is the top one *in the pair*," "the one on the positive side of the objects' principal axis" (see e.g., McCloskey, 2009), or "the one located at the orbit's 0° position." The point is that the target is specified in terms that are independent of the current placement of the orbit in the world.

(b) Where is the orbit located as a whole?

Answering this question requires some reference to position in an external frame—for example, the center point of the orbit is at (x, y) on the computer screen. It also needs to describe how the orbit is aligned to the external frame of reference. For example, one could describe the orbit's top as currently pointing toward the top end of an external vertical axis, or as in Figure 1, one could describe an orbit-internal axis as parallel, perpendicular to, or angled relative to an external axis.

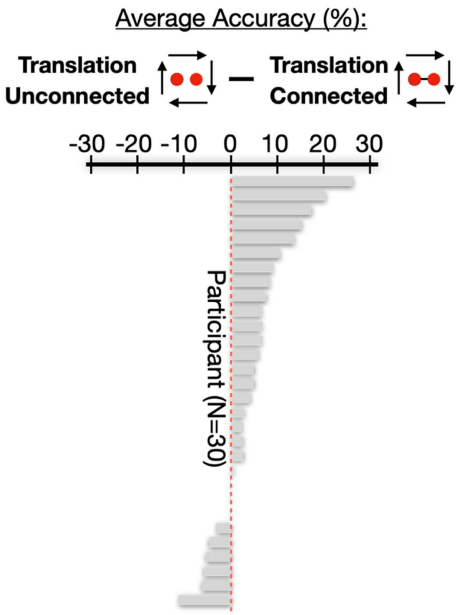
Notice that (b) included no mention of targets or nontargets, and that (a) lacked any reference to the external world. But putting them together does allow for inference about where a target resides in the external world. For example, if the target is the one on the + side of the principal axis, and that axis is parallel to the external vertical with top (+) sides aligned, then currently, the target must be the disk that is located above the other disk in the image. As the orbit unfolds, one would need to update the description, for example, to identify that the orbit's top now points to the left end of an external horizontal axis—indicating that the target is the leftmost object in the current image. Eventually, there will even be a mismatch, where the inference that the target is the one on the bottom (in the image) derives from the knowledge that the orbit's + end is aligned with the − side of an external vertical. When prioritizing orbit-centric terms, the target can remain "on top" within the orbit, and it can still appear on the bottom in a current scene because the orbit itself is "upside down" (see also Figure 2).

If this seems complicated, that is the point! Starting with an internally referenced representation requires that externally referenced relations are constantly updated, and it demands an inferential

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 9

Individual Participant Results for Experiment 2



Note. Average accuracy in unconnected translation trials minus connected translation trials for each participant ($N = 30$). See the online article for the color version of this figure.

step to describe a current scene in practice. It should therefore be error-prone under some conditions.

Specifically, when tracking two targets in two different orbits, we expect *more* errors under conditions where there is a mismatch between how the two orbits are aligned relative to the external frame of reference. Figure 10 shows three types of target arrangements that could be present at a given moment in Experiment 1. For simplicity, we can rely on a basic version of an internally referenced description of the disk pairs: Each pair has a target side and a nontarget side (as with the dumbbells in Experiment 2). The task is to keep track of which side is the target side at every moment. In Figure 9A (left example), because the two targets are in the top half of the image, a complete representation for localizing the targets is as simple as noting that the target side of each pair is above the horizontal meridian or on the + side of the external *y*-axis.

However, when one target is on top in one pair and the other target is on the bottom in its pair, as in Figure 9B, a separate description is required to localize the individual target in each pair. For example, "The target for the inner pair is on top. The target for the outer pair is on the bottom." Two sentences are needed instead of one. And these sentences have key parts that will produce tracking errors if they become confused. Mistakenly switching "inner" and "outer" or "top" and on "the bottom" in these sentences will identify nontargets as targets. We suspect that tracking errors should be more common when targets are "anti-oriented" in the sense of Figure 9B compared to when they are co-oriented, as in Figure 9A.

There is also a third option, where the disk pairs are arranged orthogonally relative to one another (see Figure 9C). We are unsure what to predict about these cases. Presently, we seek an analysis that allows us to consider whether the relative orientation of the disk pairs has any effect on tracking performance, and especially, whether anti-orientation aggravates performance compared to co-orientation.

Overall, we expect that as two targets spend more time co-oriented, it will be easier to track them. As they spend more time anti-oriented, it should make tracking harder. To analyze our data, we took a streamlined approach, comparing tracking performance as a function of relative *starting* orientations. We used relative starting arrangements as a proxy for time spent in each configuration, labeling trials based on how the disk pairs arranged when the targets were revealed, and therefore, for at least the first few moments of tracking.

Figure 11 plots the results from Experiment 1. In the rotation condition, performance was significantly worse when disks commenced with anti-orientations compared to co-orientation, $t(29) = 2.77$, $p = 4.89 \times 10^{-3}$, $d = 0.41$. Initial co-orientation advantaged tracking in rotation trials.

The same was not true in translation trials. Here, the comparison between co-orientation and anti-orientation was not significant, $t(29) = 1.04$, $p = .15$, $d = 0.12$.

We reanalyzed the results of Experiment 2 in the same way. For disconnected objects (Figure 12, left panel)—trials which were identical to those in Experiment 1—the effect of starting orientation replicated what we found in the secondary analysis of Experiment 1: significantly better performance for co-oriented trials compared to anti-oriented trials when the motion condition was rotation, but not when it was translation. Rotation: $t(29) = 2.34$, $p = 1.32 \times 10^{-2}$, $d = 0.24$; translation: $t(29) = 0.70$, $p = .25$, $d = 0.12$. Interestingly, for the connected objects we now found a significant effect of starting orientation in both rotation and translation trials. Rotation: $t(29) = 2.10$, $p = 2.23 \times 10^{-2}$, $d = 0.29$; translation: $t(29) = 1.87$, $p = 3.60 \times 10^{-2}$, $d = 0.25$.

Secondary analyses showed that targets are more difficult to track when they start a trial with opposite relative orientations, compared to when they start with the same relative orientations—for *connected* items regardless of motion, and for *disconnected* items, when the motion condition is rotation, but not when it is translation. A similar pattern appears true of co-oriented stimuli compared to ortho-oriented ones, although these arrangements were not the focus of our analyses.

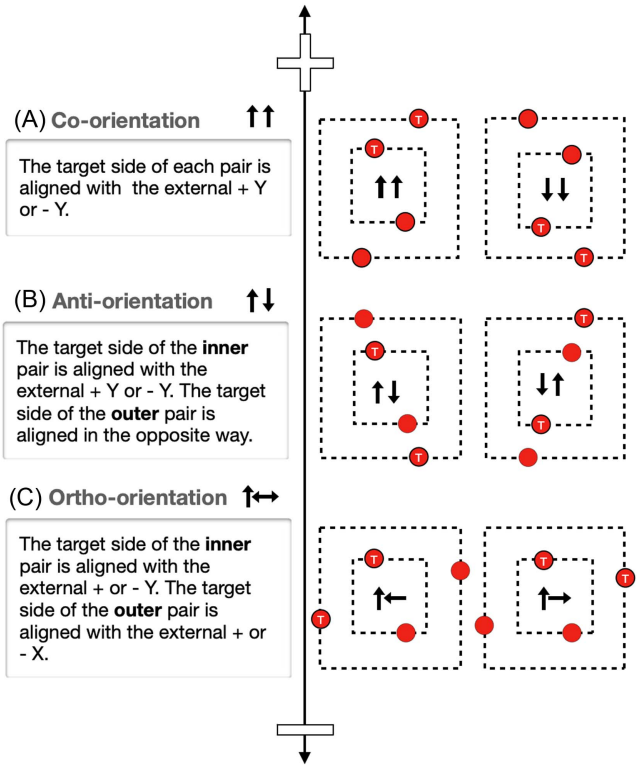
That *relative arrangement* of the disk pairs impacts performance at all is evidence that they are represented as pairs when connected or when rotating, though not when they translate disconnected. This is because *arrangement* as we have defined it is a property of disk pairs. And *relative arrangement* is a property of pairs of pairs. Saying that "both targets are on top" inherently refers to where each target resides relative to a paired nontarget. Whereas saying "that is a target and that is also a target," makes neither implicit nor explicit reference to tandems. The effect of arrangement thereby supports the first portion of the hypothesis that we originally sought to test: That moving around a circle automatically induces a representation of a higher order structure that places satellites inherently in relation to one another.

That *starting* arrangement, specifically, has predictive power, in our view stems from the fact that relative arrangement should constrain performance in general. Co-arrangement should be easier to manage, because it allows for shared descriptors to apply to more than one set of items, for example, "they are both upside down." This means that co-orientation creates a lesser need to keep pieces of

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 10

Separate Orbits Defined Relative to a Vertical External Axis



Note. Two orbits in a scene can be (A) co-oriented relative to an external frame of reference; they can be anti-oriented, with opposite alignments to an external frame (B); or their orientations relative to external references can be orthogonal to one another, ortho-orientation (C). Co-orientation affords efficient descriptions of targets' locations, for example, "The target side of each pair is aligned with the + side of an external *y*-axis." Anti- and ortho-arrangements require separate descriptions for each target. See the online article for the color version of this figure.

representations properly assigned to their respective items, compared to anti-orientation and probably ortho-orientation as well. What we do not know from these analyses is why exactly the starting position is predictive: Is it because it is correlated with the amount of time the items spend in each relative arrangement, how frequently arrangement switches take place, or moments when errors tend to take place (or a combination thereof).

Experiment 3: Replication of Experiment 1

The methods in these experiments are new, developed to ask a specific question, and the results are novel. In a tracking task like this one, floor performance is inherently capped at 50% accuracy (guessing), and effects can be quite small in percentage terms. We therefore replicated Experiment 1 in several ways. Experiment 2 included a direct replication, alongside the connected objects that extended the results. Experiment 3, which we report now, was a direct replication. Experiments 4 and 5 replicate the results again, with the orbits arranged adjacent to one another.

Method

This was a direct replication of Experiment 1, with the only differences as follow. Working from the data of Experiment 1, we conducted a power analysis, determining that to observe an effect of 0.4, we should test 20 participants. We implemented the same inclusion criteria as before. A new group of 28 participants completed the study, and we analyzed data from 22. An additional two participants, beyond the 20 sought, were unplanned, caused by two subjects who signed up for the experiment, neglected to complete it, and then later sought to complete it to receive the attendant course credit.

Experiment 3 differed from Experiment 1 in one more important way: Target pairs always began a trial in an anti-orientation (50% of trials) or in co-orientation (50% of trials). To implement this, the position of one target was randomly chosen from all the points on the perimeter of the respective virtual square. The position of the second target, on the other square, was then chosen from among the points that would produce the desired orientation. Each participant completed 256 trials.

Results

Tracking Accuracy

As shown in Figure 13, a planned one-tailed t test showed significantly better performance in the translation condition (86%) compared to rotation (84%), $t(21) = 1.89$, $p = 3.65 \times 10^{-2}$, $d = 0.40$. 14/22 participants showed better performance in the translation condition compared to rotation (Figure 13).

Relative Starting Positions

Figure 14 plots the results of a relative orientation analysis. In the rotation condition, performance was significantly worse when discs commenced with anti-orientations (80%) compared to co-orientation (87%), $t(21) = 3.27$, $p = 1.82 \times 10^{-3}$, $d = 0.47$. Thus, initial co-orientation advantaged tracking in rotation trials. The same was true in translation trials. Here, the comparison between co-orientation (88%) and anti-orientation (83%) was significant as well, $t(21) = 2.14$, $p = 2.22 \times 10^{-2}$, $d = 0.42$.

Discussion

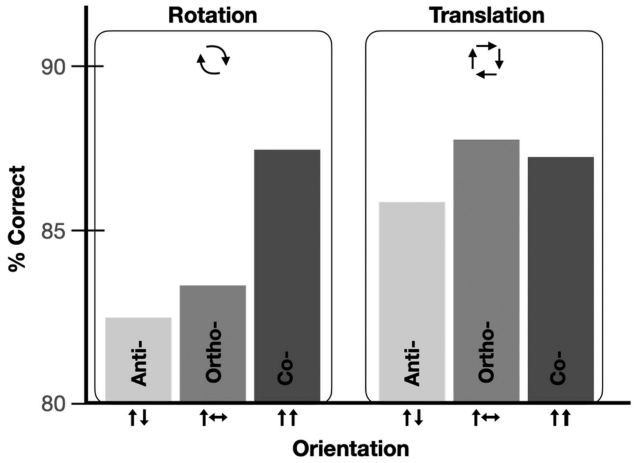
We successfully replicated the primary result of Experiment 1: a disadvantage for tracking through rotation compared to translation. This experiment also replicated the effects of starting arrangement reported as supplementary analyses with respect to Experiment 1. One difference here was that we saw an advantage for co-orientation compared to anti-orientation not only for rotation but also for translation (whereas in Experiment 1 the effect was isolated to rotation.) A practical reason for this could be that Experiment 3 did not include orthogonally oriented starting arrangements, limiting the possible arrangements to co- and anti-orientations. This meant that there were more trials associated with these cases compared to previously. Perhaps this allowed the significant effect in the translation case to emerge.

More research will be needed to understand exactly when and why moving items are represented relative to one another—with internal reference—as opposed to independently, and also whether both forms of representations may be present simultaneously, or if

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 11

Results of Secondary Analysis of Experiment 1



Note. Performance by motion type and the relative starting orientations of the targets in each trial.

instead, they are exclusive at a given time. Our suspicion is that in the current experiments, participants sometimes represent the discs as dyads—with internal reference—and that they sometimes represent them independently—that is, representing the targets as individuals with locations referenced externally. We think so because of the repeated (and frankly, contrived) nature of any experimental setup. Participants do many trials that are indistinguishable from one another at the start. Given that half the trials then rotate and half translate makes it reasonable that participants might represent pairs as such—with pair-internal reference—at the beginnings of some translation trials, and inversely, that they might represent items independently at the beginning of some rotation ones. Their presence at the start of a trial could also be a reason for why visible object connections appear to be strong inducers of internal reference, though we are hesitant to conclude that they are stronger than rotation based on the results in the current report.

Finally, a related and especially interesting possibility is that internal and external priority can be imposed intentionally, a kind of willful grouping (or ungrouping). Ongchoco and Scholl (2019) reported that classic same-object advantages can be observed in purely imagined objects, a demonstration of how object representations can be imposed with intention on otherwise unstructured displays to create "attention-based objects." If, as we have argued, the mechanisms of object-based representation apply broadly, inclusive rotational systems, then this specific mechanism should apply broadly.

These are topics for future research. For current purposes, the fact of an arrangement effect in the translation trials is consistent with our general theorizing. As we said in the context of the supplementary analyses, that relative arrangement of the disk pairs impacts performance at all is evidence that they are sometimes represented as pairs with internal reference.

Experiment 4: Side-by-Side Horizontal Placement

To test the generalizability of the new effects reported, we replicated the prior experiments with a different response measure and with the stimuli arranged differently on the screen. Rather than the homocentric arrangement of the two (virtual) squares that anchored disks and their motions, we arranged the squares horizontally side by side (Figure 15). Among other differences, this meant that the two virtual squares could be equally sized. In this experiment, we also asked participants to click on each of the two targets at the end of a trial, rather than probing one item and asking participants to report whether it was a target or a nontarget.

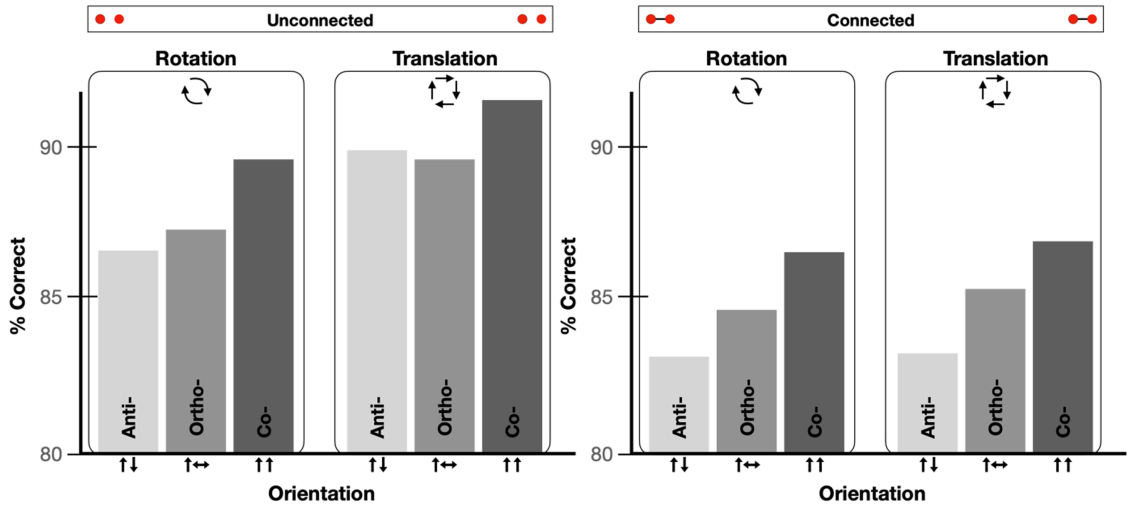
Method

Preregistration

This experiment was not preregistered.

Figure 12

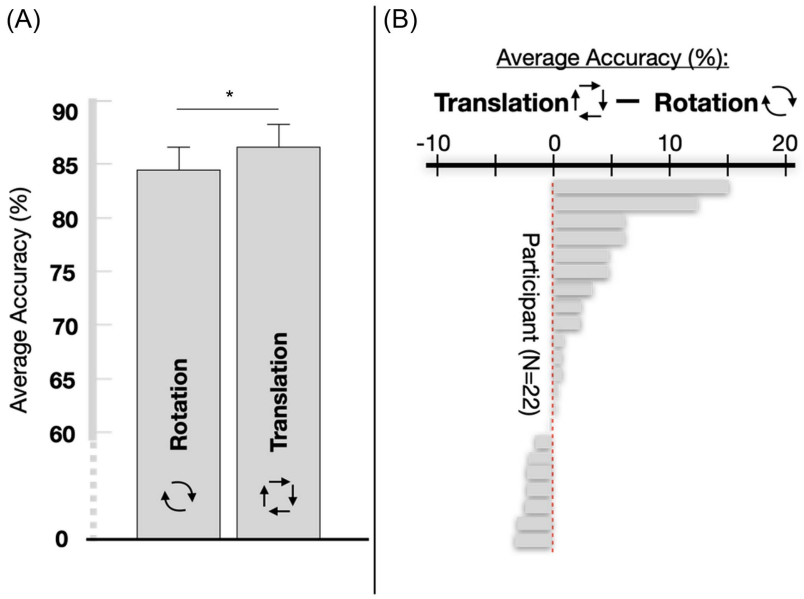
Results of Secondary Analysis of Experiment 2



Note. Performance by motion type, connection and the relative starting orientations of the targets in each trial. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 13
Results of Experiment 3



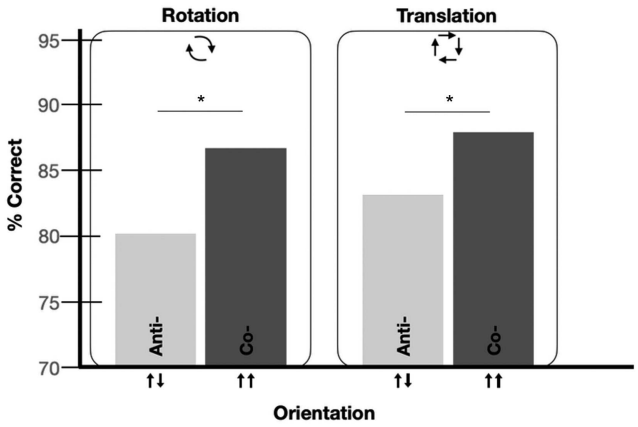
Note. (A) Average tracking accuracy was significantly better in the rotation trials compared to the translation trials. Error bars show standard errors. (B) Average accuracy in translation trials minus rotation trials for each participant ($N = 22$). See the online article for the color version of this figure.
* $p < .05$ (for a planned comparison).

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Experiment Link

Experiment 4 can be viewed online with a web browser at: https://qihan.visualthinkingresearch.com/tracking_COR/exp_4/tracking_exp_4.html.

Figure 14
Relative Orientation Results of Experiment 3



Note. Performance by motion type and the relative orientations of the targets at the start of a trial ($N = 22$).
* $p < .05$ (for a planned comparison).

Participants

We tested a total of 40 undergraduates, 10 excluded from analysis using the same criteria as in all the prior experiments.

Stimuli and Procedure

Each trial included two invisible squares placed 55 pixels to the left and right of the screen center (which was marked by a small fixation square in black). The squares are designated in Figure 15B by dotted lines which were not visible in the actual experiment. The squares each had a side length of 190 pixels.

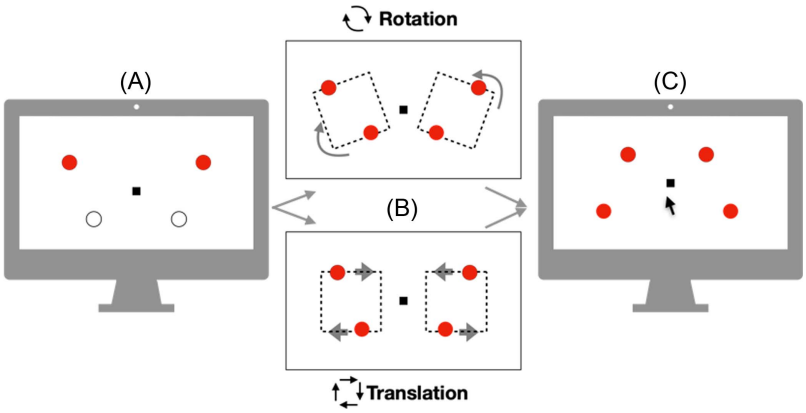
Each participant completed 240 trials. To begin a trial of any kind, two pairs of disks appeared in the display, each pair occupying the perimeter of one of the virtual squares. For each square, the position of one disk was chosen randomly to occupy any point on the square's perimeter with the restriction that the point be at least 30 pixels away from the nearest corner of the respective virtual square. The second disk on a given square always occupied the reflected position of the first.

At the beginning of a trial, all four disks remained stationary for 1,500 ms. Two disks (one from each pair) were white to designate them as the tracking targets. After the 1,500 ms stationary period, the two white disks turned red so that all the disks were visually identical. Now the disks moved for a duration of 2,300 ms. When motion ceased, an arrow cursor appeared in the display. Participants were instructed to use the mouse to click on the two disks that they believed to be targets.

Trials included two motion types, randomly and equally distributed throughout a testing session. In the rotation condition, the

Figure 15

Schematic Depiction (Not to Scale) of Methods Used in Experiment 4



Note. (A) Every trial starts in the same way, with four disks each affixed in opposite-location-pairs to the perimeter of one of two virtual squares (shown in B). One disk in each pair appears white to mark it as a target for tracking. The white disks then turn red, so that all the disks look identical, and then (B) a period of motion and tracking begins. In half the trials, the motion occurs through the rotation of the virtual squares (each square in the opposite direction of the other). In the remaining half of trials, the disks on each square translate in synchrony along the virtual perimeters (the disks on one square moving in the opposite direction from the disks on the other). Finally, all trials end in the same way (C): motion ceases, and an arrow cursor appears. Participants clicked to select the disk in each pair that they believe was the target. See the online article for the color version of this figure.

disks rotated as pairs around the center points of their respective virtual squares, defining orbits with diameters equal to the distances between them. The initial rotation direction for the right pair and the left pair was always in the opposite direction. Direction changes took place asynchronously. For the left disk pair changes took place every 300 ms and 600 ms (alternating), while direction changes for the right disk pair took place every 600 ms and 300 ms (alternating). In the translation condition, the disks in each pair moved along the perimeters of their respective squares, turning corners as necessary.

Motion speeds were set with same rules but different parameters as in Experiment 1. The angular speed for all disks in rotation trials was 576°/s. This meant that the number of radians a disk traveled per second varied as a function of where the disks happened to start along a square's edge. (i.e., A pair of disks that started near corners traced a larger circle than a pair that started near the middle of an edge). The minimum radian speed for a rotating pair was equal to half the square length (190 pixels) multiplied by 576. The maximum was the same multiplied by the square root of two. We took the average of these two values to set the linear motion speed for disks in the translation condition.

Reaction Time and Accuracy Validation

We analyzed 6,050 trials with correct responses and latencies that were between 200 ms and 3,000 ms.

Results

Tracking Accuracy

The responses in this experiment constituted two clicks, one target selection from among each pair of disks on either side of the display. We therefore analyzed responses by comparing first-click accuracy for translation and rotation trials separately from second-click accuracy (for the same comparison). A planned one-tailed t test showed significantly better performance on the first click in the translation condition (90%) compared to rotation (89%), $t(29) = 2.19$, $p = 1.83 \times 10^{-2}$, $d = 0.40$. Performance was also significantly better for the second click in the translation condition (89%) compared to rotation (87%), $t(29) = 2.72$, $p = 5.45 \times 10^{-3}$, $d = 0.50$. Figure 16 plots these results, along with individual participant performance differences (here first and second clicks are averaged together).

Discussion

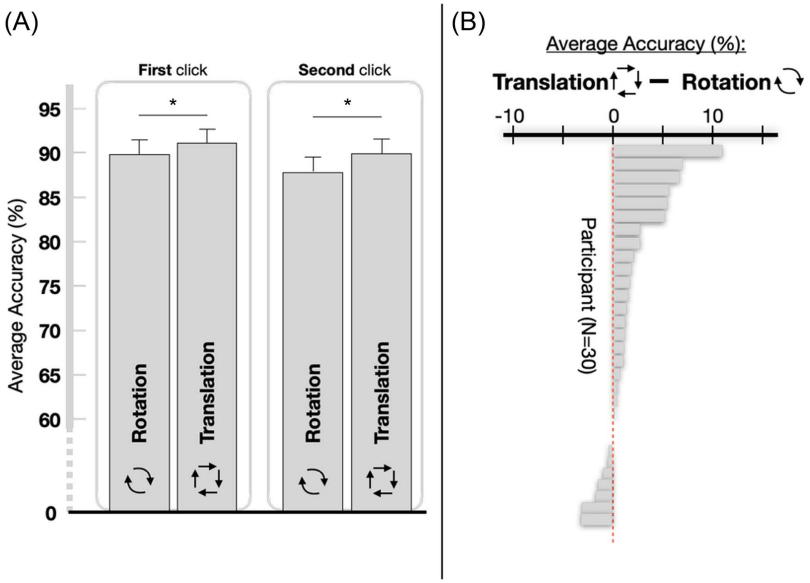
The experiment replicated the main result of Experiment 1 with a different configuration of stimuli and with a different measure of tracking accuracy. A further value of this specific replication is that it demonstrates that tracking challenges observed in this paradigm are not caused by hemifield-specific resource constraints. It has been suggested that each visual hemifield operates under independent resource constraints for tracking, generally making tracking easier when parallel demands are hemifield segregated (Alvarez & Cavanagh, 2005). In the previous three experiments, the central and homocentric arrangement of the tracking stimuli could, in theory, draw on the same pool of resources. In this replication, we observed a significant cost for tracking under rotation compared to translation despite a horizontal arrangement.

Experiment 5: Replication With Vertical Placement

We sought to replicate our main findings one more time while changing theoretically nonessential parameters. This experiment

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 16
Results of Experiment 4



Note. (A) Average tracking accuracy was significantly worse in the rotation trials compared to the translation trials, both for the first target choice made and the second. Error bars show standard errors. (B) Average accuracy in translation trials minus rotation trials (averaged across clicks) for each participant ($N = 30$). See the online article for the color version of this figure.
* $p < .05$ (for a planned comparison).

was identical to Experiment 4 but with stimuli arranged vertically (Figure 17) rather than horizontally.

Method

Preregistration

This experiment was not preregistered.

Experiment Link

Experiment 5 can be viewed online with a web browser at: https://qihan.visualthinkingresearch.com/tracking_COR/exp_5/tracking_exp_5.html.

Participants

We tested a total of 41 undergraduates, 11 excluded from analysis using the same criteria as in all the prior experiments.

Stimuli and Procedure

Each participant completed 240 trials. Each trial included two invisible squares placed 55 pixels above and below the screen center (which was marked by a small fixation square in black). The squares are designated in Figure 17B by dotted lines which were not visible in the actual experiment. The squares each had a side length of 190 pixels. In all other respects, the stimuli and procedures were identical to those of Experiment 4.

Reaction Time and Accuracy Validation

We analyzed 6,246 trials with correct responses and latencies between 200 ms and 3,000 ms.

Results

Tracking Accuracy

The responses in this experiment constituted two clicks, reflecting a target selection from among each pair of disks on the top and bottom of the display. We therefore analyzed responses by comparing first-click accuracy for translation and rotation trials separately from second-click accuracy (for the same comparison). A planned one-tailed t test showed significantly better performance on the first click in the translation condition (95%) compared to rotation (93%), $t(29) = 2.14$, $p = 2.05 \times 10^{-2}$, $d = 0.39$. Performance was also significantly better for the second click in the translation condition (92%) compared to rotation (88%), $t(29) = 2.77$, $p = 4.83 \times 10^{-3}$, $d = 0.51$. Figure 18 plots these results, along with individual participant performance differences (here, first and second clicks are averaged together).

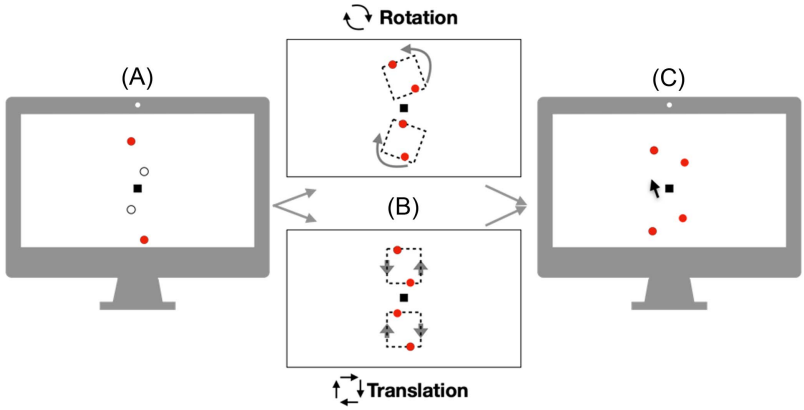
Discussion

The experiment once again replicated the main result of Experiment 1, with a different configuration of stimuli and with a different measure of tracking accuracy. As with Experiment 4, this experiment makes it difficult to interpret the results of the

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 17

Schematic Depiction (Not to Scale) of Methods Used in Experiment 5



Note. (A) every trial starts in the same way, with four disks each affixed in opposite-location-pairs to the perimeter of one of two virtual squares (shown in B). One disk in each pair appears white to mark it as a target for tracking. The white disks then turn red, so that all the disks look identical, and then (B) a period of motion and tracking begins. In half the trials, the motion occurs through the rotation of the virtual squares (each square in the opposite direction of the other). In the remaining half of trials, the disks on each square translate along the virtual perimeters (the disks on one square moving in the opposite direction from the disks on the other). Finally, all trials end in the same way (C): motion ceases, and an arrow cursor appears. Participants clicked to select the disk in each pair that they believe was the target. See the online article for the color version of this figure.

previous three as originating with idiosyncrasies of homocentric arrangement.

Experiment 6: Explaining the Freezing Illusion Through Orbit-Centered Reference

Here, we report a rather different experiment designed to provide converging evidence for internally referenced representations of orbiting. A feature of an internal reference is that it can be compared to multiple other reference frames. To the extent that an orbit is "like an object" therefore, one can lack motion relative to one reference frame while moving relative to another. Anything that is object-like can move relative to one reference frame while remaining static relative to another, if that frame moves with it synchronously. In such a case, to answer the question, "is it moving?" the observer needs to know, "relative to what?"

Consider again the pairs of rotating disk-pairs in Experiment 1. Our secondary analysis of those data suggested that observers naturally relate the internal reference frames of the two orbits to one another. We therefore hypothesized that when considering how (and even whether) an orbit is in motion, priority is assigned to nearby systems with their own internal reference frames: that the answer to the question "relative to what" is first, relative to the internal reference frame of an adjacent system, and only after, relative to a more global, external reference.

To test the hypothesis, we turned to a visual illusion called *the freezing rotation illusion*, characterized by a "decrease in perceived speed of a continuously rotating central region when a swaying surround corotates" (Dürsteler, 2008, p. 283). The reader can view a demonstration movie with the following link: https://qihan.visualthinkingresearch.com/dots2/demo/displayVideo.html.

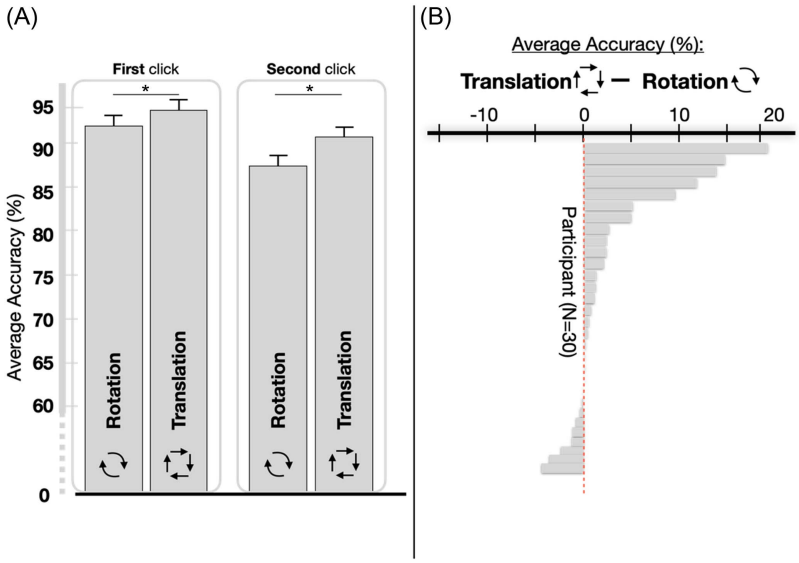
This demo is one that we developed for the purposes of the current experiment. It produces the same effect as more typical versions. In our demo, the viewer will see a square made up from many white disks. Homocentric with the square are two yellow disks. The two yellow disks rotate continuously as a tandem in one direction. The square on the outside rotates as well, but it changes direction every so often from clockwise to counter. The illusion: When the outer square changes direction so that it goes from moving opposite to moving with the disk pair, then the disks appear to freeze momentarily, even though they never stop moving or change speed in practice.

We reserve interpretation of the illusion to follow the presentation of results. The experiment we designed asked observers to answer the question "is it moving?" with the intended meaning applied relative to an external reference frame. Is it moving, period? We anticipated that answering the question should be difficult—error prone and slow in a specific case: when the question applies to an orbiting system that does not move in relation to another rotating system although it does relative to a global frame. The prediction follows if visual processing prioritizes a notation of the relationship between internal reference frames.

Trials included two yellow disks on either side of a display. Around each pair of disks we also placed two squares (made up from groups of white dots), as shown in Figure 19. In each trial, *the yellow disks on one and only one side of the screen rotated* continuously in one direction. The pair on the other side was static. The two surrounding squares in a trial always moved in exactly the same way, either rotating as an object, or through the translation of the constituent white dotes, and either in the same direction as the moving yellow disks in that trial, or in the opposite direction.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 18
Results of Experiment 5



Note. (A) Average tracking accuracy was significantly worse in the rotation trials compared to the translation trials, both for the first target choice made and the second. Error bars show standard errors. (B) Average accuracy in translation trials minus rotation trials (averaged across clicks) for each participant ($N = 30$). See the online article for the color version of this figure.
* $p < .05$ (for a planned comparison).

The task for participants was to report by keypress which side included the moving disks, producing a measure of accuracy and also latency (see also Wu & Flombaum, 2024). We predicted that when the outer squares rotated in the same direction as the rotating disks, that the task would be more difficult (elevated error rates and reaction times) compared to when the rotating squares moved opposite the disks. This would provide, to our knowledge, the first quantitative measure of the strength of the freezing illusion, employing the logic that if a synchronously moving surround causes the internal disks to appear to "freeze," then detecting their motion should be error prone and delayed. We further predicted that the translation of the constituent dots would produce less or no "freezing," even when the dots move in the same direction as the disk rotation, because translation should reduce reliance on an internal reference frame in the representation of the surround. Quantitatively, this effect should appear as lower error rates and faster response latencies for the same-direction translation trials compared to the same-direction rotation trials.

Method

Experiment Link

Experiment 6 can be viewed online with a web browser at: https://qihan.visualthinkingresearch.com/dots2/exp4/Dots_exp4.html.

Participants

We sought to collect data from 20 participants. We planned that any participant with a correct response rate below 75% overall would be excluded and replaced automatically. Ultimately, we tested a total of 20 participants.

Participants were recruited through an online system for students to enroll in studies and receive course-related credit. Demographic information was not collected since it is not related to the hypothesis. Participants completed the experiment on a personal device. The instructions and sign-up page asked that they use a laptop or desktop computer not a mobile device.

Stimuli and Procedure

Each trial included two square surrounds separated from each other by 200 pixels between their right and left edges and centered in the display vertically. The surrounds comprised one hundred white dots randomly arranged in the space between two homocentric virtual squares. The outer square had a length of 450 pixels, and the inner square had a length of 250 pixels. Each individual dot was 10 pixels in radius.

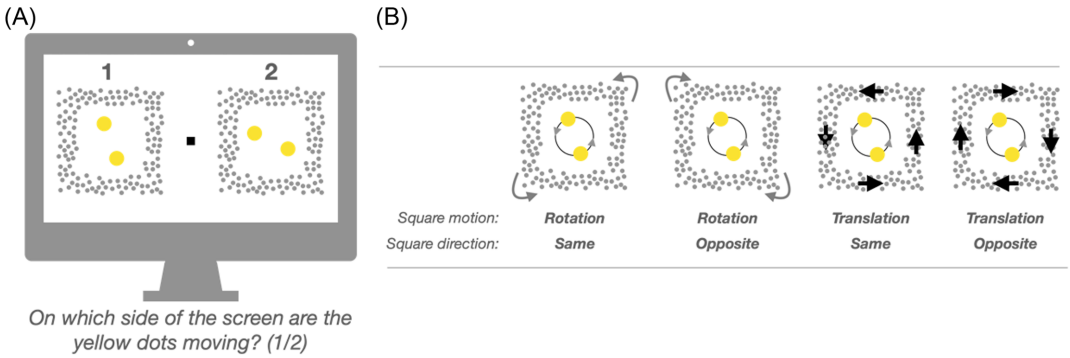
Each trial also included two pairs of yellow disks (radius of 15 pixels, RGB value of 255, 255, 0). Each pair of disks was placed on an invisible square of length of 160 pixels, centered within one of the surround squares. One disk in each pair was placed randomly on the perimeter of the respective virtual square (excluding points within 30 pixels of a corner). The second disk in the pair was always reflected to the opposite position on the virtual square perimeter

One of the two orbit systems always remained stationary, while the other rotated clockwise or counterclockwise continuously. The square backgrounds always moved in the same exact way in a given trial, either translating or rotating, to produce four motion

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 19

Schematic Depiction (Not to Scale) of the Methods in Experiment 6



Note. Each trial included side-by-side stimuli including a pair of yellow disks surrounded by a square comprising a collection of white dots (A). Here, the dots are shown in gray for print clarity. The yellow disks on only one side of the display rotated continuously, either clockwise or counterclockwise. The task was always the same: to report which side of the screen included the rotating pair of disks. The motion of white dots varied by trial as shown in (B). The square could rotate continuously as a whole object (in the same or opposite direction from the yellow disks in the given trial). Or the constituent white dots could translate continuously (again, in the same direction as, or opposite to the yellow disks in a trial). The two squares in a given trial always moved in exactly the same way. The task was to indicate on which side of the screen the yellow dots were moving. See the online article for the color version of this figure.

conditions: same direction rotation, opposite direction rotation, same direction translation, opposite direction translation.

The disk pairs and the square surrounds rotated at a rate of 90°/s. During translation trials the constituent dots moved at a speed of 194 pixels/s. Each motion condition was repeated 60 times so that an experimental session included a total of 240 (randomly distributed and counterbalanced) trials per participant.

The RGB values for the webpage background were RGB (128, 128, 128). A black fixation square (10 pixels) was placed in the middle of the screen between the two square surrounds. Instructions to participants included the request that they fixate the center of the display during the times between trials, prior to the onset of stimuli. The task for this experiment was to press a key to indicate which set of stimuli includes the rotating yellow disks.

Reaction Time and Accuracy Validation

We analyzed 4,776 trials with response latencies between 200 ms and 2,000 ms. Among them 4,667 trials had correct responses.

Results

Accuracy

Figure 20 plots average accuracy as a function of the motion conditions. A repeated measures analysis of variance showed a significant main effect of relative motion on average accuracy, $F(1, 19) = 18.15$, $p = 4.23 \times 10^{-4}$; no significant main effect of background motion type $F(1, 19) = 1.10$, $p = .31$; and a significant interaction between these two factors $F(1, 19) = 7.36$, $p = 1.38 \times 10^{-2}$.

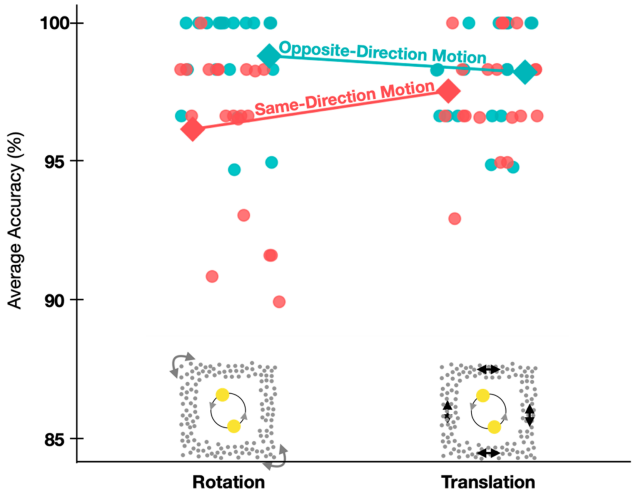
A planned one-tailed t test showed a significant accuracy difference for rotation trials in which the surround rotated in the same direction as the moving disk pair compared to trials when it moved in the opposite direction, $t(19) = 4.37$, $p = 1.65 \times 10^{-4}$, $d = 1.02$. The same planned comparison for translation trials was marginal, $t(19) = 1.52$, $p = .07$, $d = 0.36$. Finally, accuracy was significantly better in same motion trials under translation compared to rotation, $t(19) = 2.08$, $p = 2.58 \times 10^{-2}$, $d = 0.53$. But the comparison between translation and rotation was not significant in opposite direction trials, $t(19) = -1.78$, $p = 9.10 \times 10^{-2}$, $d = -0.33$.

Reaction Time

Figure 21 plots average response latency as a function of motion conditions. A repeated measure analysis of variance showed

Figure 20

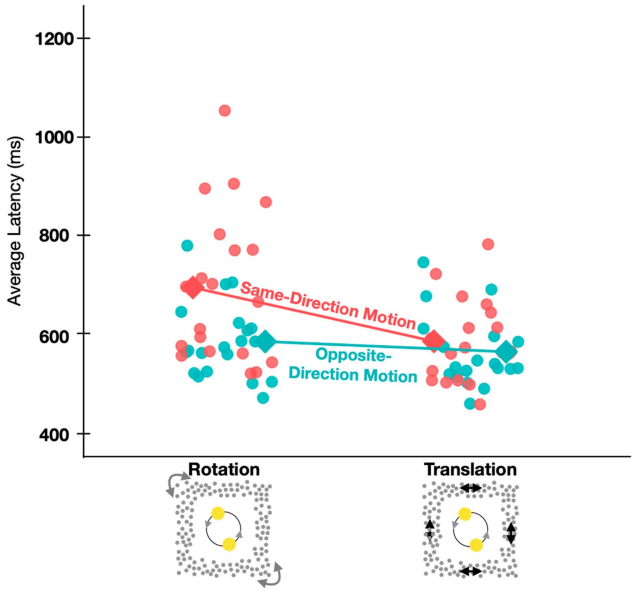
Accuracy Results of Experiment 6



Note. Average accuracy as a function of motion type (rotation vs. translation) and motion direction (same vs. opposite). Dots show individual participant responses, and diamonds mark group averages by motion direction. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Figure 21
Latency Results of Experiment 6



Note. Average response latency (ms) as a function of motion type (rotation vs. translation) and motion direction (same vs. opposite). Dots show individual participant responses, and diamonds mark group averages by motion direction. See the online article for the color version of this figure.

significant main effect of relative motion on average latency, $F(1, 19) = 26.97$, $p = 5.17 \times 10^{-5}$; a significant main effect of background motion type, $F(1, 19) = 38.36$, $p = 5.97 \times 10^{-6}$; and a significant interaction between these two factors, $F(1, 19) = 16.72$, $p = 6.25 \times 10^{-4}$.

A planned one-tailed t test showed a significant accuracy difference for rotation trials in which the surround rotated in the same direction as the moving disk pair compared to trials when it moved in the opposite, $t(19) = 4.86$, $p = 5.39 \times 10^{-5}$, $d = 0.69$. The same planned comparison for translation trials was also significant $t(19) = 3.23$, $p = 2.23 \times 10^{-3}$, $d = 0.26$. Finally, latency was significantly faster in same direction trials under translation compared to rotation, $t(19) = 5.30$, $p = 2.03 \times 10^{-5}$, $d = 0.64$. The comparison between translation and rotation was also significant in opposite direction trials (but with a smaller effect size), $t(19) = 3.80$, $p = 6.10 \times 10^{-4}$, $d = 0.27$.

Discussion

Our interpretation of the freezing illusion is that it arises when the stimulus surround and its internal components rotate in the same direction because in these moments, the internal reference axes of the two systems do not move relative to one another. In fact, it is possible that visual processing goes a step further. As became clear in the context of our supplementary relative arrangement analyses of the tracking results, a single representation of an internal axis can occasionally support the representation of parts of multiple objects or systems. In our freezing illusion stimuli, there are no targets, and especially with the homocentric arrangement of the stimuli, a single internal reference axis may be sufficient to describe the motions of both the yellow dot system and the surrounding square system, *but only when they rotate in the same direction.* When they rotate opposite one another, they require independent internal reference axes.

The results of Experiment 6 support this interpretation by demonstrating that when a surround translates rather than rotates, the illusion weakens or disappears because translation does not strongly invoke an internal reference frame. Consequently, observers can easily detect the disks' rotation relative to translating surrounds, as the internal reference of the disks remains distinct and clearly in motion. Our findings confirm that orbit-centered reference frames are integral to motion perception, with synchronous rotation uniquely diminishing perceived motion through the alignment of internal frames.

General Discussion

The key finding across the experiments is that tracking objects that move in a circular orbit is more difficult compared to tracking objects that move along a noncircular trajectory, even when the motion paths are otherwise matched. This deficit for tracking orbiting objects was observed both when the objects were visibly connected (Experiment 2) and when they were unconnected (Experiments 1, 3–5). We also discovered relative difficulty detecting the motion of a rotating orbit when an external surround rotates in the same direction (Experiment 6).

Secondary analyses revealed that the relative starting orientation of the tracked objects impacted performance. Specifically, when the two tracked objects started in a co-oriented arrangement (both on the top or bottom or left or right in their respective orbits), tracking was easier compared to when they started in an anti-oriented arrangement (one on top, one on bottom, etc.). This effect was most pronounced in the rotation conditions but also observed in some of the translation conditions.

Because the objects in half the conditions underwent rotational motion and the objects in the other half did not, the experiments necessitated performance comparisons across speeds that are best described as linear versus ones best described as angular. To address this issue, we fixed angular speed in rotation trials at a single value (576°/s), with no acceleration or deceleration so that motions onset at top velocity and objects came to immediate stops at the ends of a trials. Because the tracked disks and foils could appear in different places along the sides of virtual squares, this meant that linear speeds in rotation trials varied: When the disks were placed close to the midpoint of a side they moved around a smaller circle (linearly, more slowly) than when they were placed near the corners of their virtual sides. We fixed the linear speed of translation trials at the average of the available linear speeds in rotation trials (again without acceleration or deceleration), so that motion speed in translation trials was at most 20% faster than the slowest rotation trials and 15% slower than the fastest ones, but equal to all the rotation trials on average. Speed differences cannot therefore explain the observed effects, because translation trials moved *faster* than rotation trials as often as they moved more *slowly*. Moreover, speed cannot explain the arrangement effects observed nor the effects of object connections which impaired performance.

A related concern could be that in the translation trials, moving disks underwent direction changes at the (invisible) corners of the

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

virtual squares. It is not obvious, however, why this would make tracking *easier* in the translation trials, and in general, research on multiple object tracking suggests no advantage or disadvantages for curvilinear versus rectilinear trajectories. (See e.g., Keane & Pylyshyn, 2006, where curvilinear trajectories were directly compared to rectilinear paths, and performance was either equal across conditions or *better* for the curved trajectories.) Like with speed, these kinematic differences between conditions fail to explain arrangement and connectivity effects coherently.

Instead, the results overall are consistent with the hypothesis that moving objects in a circular orbit automatically induce an object-like representation, with an internal frame of reference, rather than a representation that simply tracks the external coordinates of the objects. An important set of open questions concerns the temporal dynamics of representation: when do object-centered versus externally referenced representations emerge or predominate during the time course of a trial? Future studies could investigate the time course of these different representational formats. A related issue is whether starting position is predictive because it is correlated with the amount of time the items spend in each relative arrangement, because it is correlated with how frequently arrangement switches take place, or because errors tend to take place early on in a trial (or a combination thereof). More broadly, when it comes to tracking individual objects in space (in the paradigm often called "multiple object tracking"), it has become clear that tracking errors emerge primarily when a target and a nontarget approach close enough to one another so that an observer confuses them, replacing the target in their mind with a nontarget (Bae & Flombaum, 2012; Upadhyayula & Flombaum, 2020). Closing distance between targets and nontargets is not the cause of errors in the current tracking experiments (as distance changes did not take place). Questions therefore remain about how and when targets and nontargets become confused with one another.

The key theoretical implication of the experiments is that the principles of object-centered representation may apply quite broadly to visual processing, outside of instances of what might be described as "proper" objects. This extension of the concept is especially important because it involves the representation of properties that are not physically present in the stimulus—axes and relational variables, in particular. The discussion that follows elaborates on these implications by analogy to basic elements of computer-assisted design and with reference to related neuropsychological research.

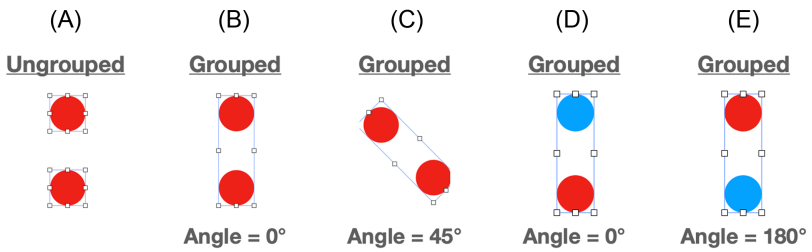
Grouping and Internal Reference in Visual Processing

In Computer-Assisted Design software grouping is a basic tool that allows the user to combine multiple objects into a single unit, making it easier to move, to resize, and especially to rotate them together. In other words, grouping allows the user to create a unit that behaves like an object out of elements that can be described as objects in their own rights. Once a pair of elements (or more) is grouped, they indeed meet the cardinal provision of objecthood: grab and pull on one end and the other end will come along for the ride. To achieve these outcomes, the grouping function represents structure that is not present in the individual elements, including axes that define the group's skeleton, circumscribing boundaries that include the elements and any shared negative space, and internally referenced variables to describe the positions of the elements in the group. Finally, the program describes the location and orientation of the group on the larger canvas.

Figure 22 shows five simple stimuli and associated variables as they would be assigned in basic slide making software (such as Apple's Keynote, where the variable values would appear in the *Arrange* window). The drawn image includes two red disks. In the first frame (A), the disks are represented by the system as independent objects. But in the second frame (B), the user has chosen to group them, and now the system bounds them in a shared rectangle to allow the user to understand the contours of the new object created. In the third frame (C), the user has rotated the grouped object by 45° counterclockwise. Note that the angle variable set to 45° is a descriptor that cannot apply to the elements individually and

Figure 22

Grouping and Rotation as Variables



Note. Five simple stimuli and associated variables as they would be assigned in basic slide making software (such as Apple's Keynote, where the variable values would appear in the *Arrange* window). The drawn image in frames A–C includes two red disks. In the first frame (A), the disks are represented by the system as independent objects. But in the second frame (B), the user has chosen to group them, and now the system bounds them in a shared virtual rectangle. In frame (C), the user has rotated the grouping by 45° counterclockwise. In (D), the user has grouped a blue disk, and a red disk arranged vertically on the canvas. Rotating the group by 180° (E) would yield the corresponding description in the arrange window—of a grouped set with an angle of 180°—although the pair of disks could just as easily be described as arranged vertically and lacking rotation. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

that inherently refers to an unrotated version of the group, a version that is no longer present on the canvas. In Frame D, the paired disks are now blue and red, with the blue one on top. Rotating the pair by 180°, as in E, would create a representation in the software (and in the associated menu) with clear internal reference: though still perfectly vertical in the canvas, the pair is represented as an upside-down version of what was on the canvas previously. Describing the pair as rotated by 180° refers to a past version of the stimulus, one that is no longer present.

Our proposal vis-à-vis the perception of orbiting motion in the reported experiments is that perception operates like a reverse Computer-Assisted Design. It represents the moving disks as satellites that are part of a larger system, one treated like an object in the sense of possessing internal reference. To accomplish the identification goal of the task assigned, we propose that observers represent the *relative* position of the target in a pair at the start of a trial. Note that it need not be so. It should be possible, in theory, to represent the target by tracking only its location, not caring whether a nontarget is present at all (see e.g., Pylyshyn, 2006), and certainly without representing the target location through reference to the nontarget. Indeed, we believe that this alternative is more frequently what occurs in translation trials. The difference between translation trials and rotation trials can be stated as something like a difference in how frequently perception describes a target as "that one" versus "the one on top."

In rotation trials, observers then hold onto their relative representation of a target's position as a canonical or platonic arrangement, so to speak. Tracking therefore depends on iteratively updating a representation of how an axis within the group rotates relative to the canvas, allowing for an inference about the position of the target at rest. For example, if the target is the one on the right when the internal axis is parallel to an external vertical axis, and the internal axis has rotated by 180° (on the canvas) since the start of the trial, then *currently* the target must be the one on the left on the canvas. See again Panels D and E of Figure 22.

As with any reverse graphics proposal, we intend our usage of software terms as analogies. The *arrange* dialogue and its variables are not to be taken as literal embodiments of mental events, but instead, as functional descriptions. But the terms "group" and "grouping" have played an important role in the discussion so far, and they have long-held status in perceptual psychology (Wagemans et al., 2012). What is the relationship, then, between our usage of these terms and their usages in perception research?

We can answer the first question confidently. In our usage "grouping" is a procedure that links together elements in an input image, minimally, by assigning them a shared (internal) axis of reference. It follows that a group, as we use the term, is a kind of representation, one that describes otherwise distinct elements with respect to shared axes that need not manifest physically in the raw stimulus. A group is further a representation in the sense that it becomes the input for routines, such as tracking, and the form of the representation dictates the kinds of routines available to perform a task. If the disks in our dyads were ungrouped, they would be tracked independently, through reference to external coordinates (in something like the way proposed, e.g., by Upadhyayula & Flombaum, 2020; see also Vul et al., 2009). But because they are grouped, they are tracked with an entirely different routine, one that relies on deductions from internal references to external ones.

To contextualize how the terms group/grouping are used in perception research, we borrow a minimal definition from a comprehensive review, where the terms are said to refer to "the fact that observers perceive some elements of the visual field as 'going together' more strongly than others (Wagemans et al., 2012, p. 1180)." We have made the case that two satellites on a rotational orbit are perceived as "going together more strongly" than the same objects on translational paths, and perhaps even as strongly as the two ends of a dumbbell go together. That does not necessarily mean that objects in rotation are a group in the same sense as two items that are grouped, for example, by common region, proximity, or any other cue. We are also not aware of research on grouping that offers a specific proposal about the underlying format of a representation of a group. Gestalt and related research describe in detail the conditions under which the perception of a group arises, and more computational research has been similarly engaged with models of how a system determines when and which elements go together more strongly than others. Perhaps, then, the key distinction between our usage of the term "grouping" and more traditional usages is that the latter refers to the fact that elements become assigned to a group, whereas we have made further claims about the underlying form of the new representation that emerges, and the inclusion of representational elements beyond the stimulus itself.

Where does that leave the relationship between perceiving an orbit and perceiving a group more generally? One option is that they are just different. The other possibility is that what grouping sometimes does is impose upon elements an internally referenced representational format, complete with structures that need not be present in the raw stimulus. To the extent that the principles of object-centered representation may apply quite broadly to visual processing, Gestalt groups may be another instance.

Object-Centered Reference: Neurodivergence, Working Memory, and Faces

To describe an object as upside down, the object must have an upside. Upside down, moreover, inherently refers to a frame of reference internal to an object, independent from any other frame of reference. That language can describe objects with terms that imply internal reference, and that people use these terms, does not necessarily mean that perception represents objects in the same way. Beyond our account of orbiting motion, what evidence suggests that perceptual representations can possess internal reference?

Faces, people, and living things are candidate stimuli, because they naturally possess internal axes that could be included profitably in perceptual representations. It is well known that these stimuli suffer processing disadvantages when displayed upside down. Inversion impairs the detection and discrimination of point-light walkers, for example, (e.g., Pavlova & Sokolov, 2003; Shipley, 2003; Troje, 2003). Of course, inversion causes severe impairments to face processing (e.g., Farah et al., 1995). Animals are easier to find in visual search when oriented upright (e.g., Ballaz et al., 2005). And there is also evidence that inversion constraints may arise very early in development (e.g., Schwarzer, 2000; Vallortigara & Regolin, 2006). However, inversion effects are not dispositive. Sliding a credit card into a gas pump with the chip side facing out

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

will not activate the pump, understandably, without dragging representations or internal references into the discussion.

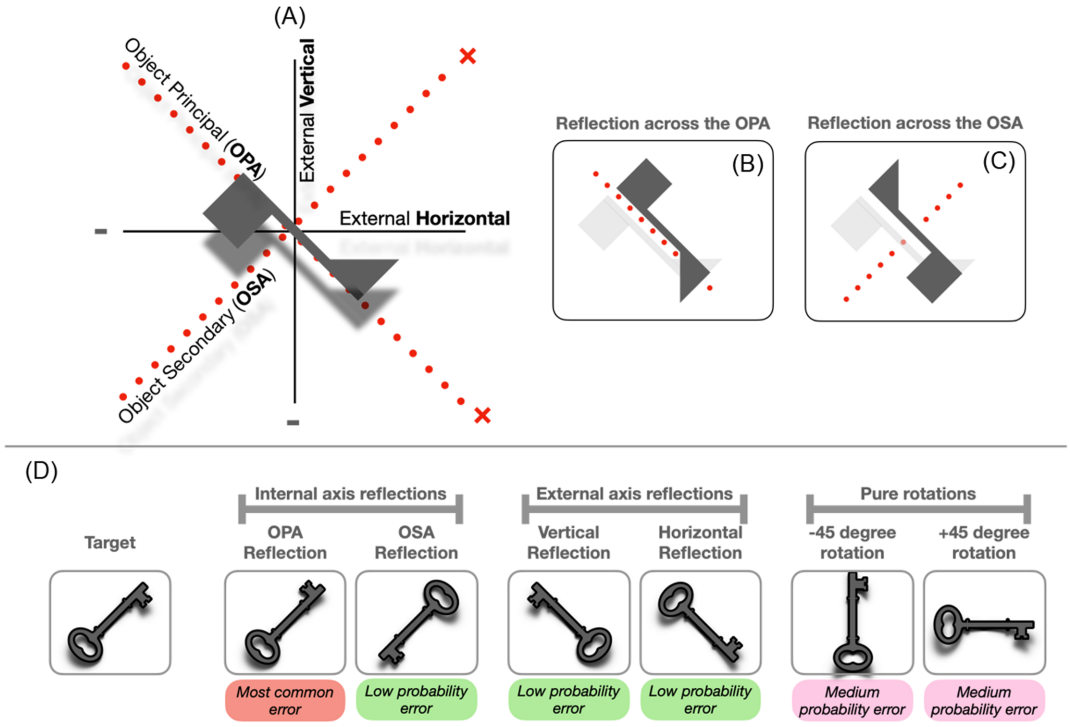
In the case of faces, though, the disorder prosopometamorphopsia suggests that internal reference may be deeply ingrained. In prosopometamorphopsia, an observer perceives faces as severely distorted (Blom et al., 2021). The disorder has high recovery rates, though it can be chronic. The distortions often appear to be unilateral, but in at least one patient, it seems possible that they could be side-of-face lateralized, that is, face-centered. The patient always sees the *right side of a face* as distorted: appropriately present or absent when a face is shown in profile, and on the left side of space when a face is upside down (Almeida et al., 2020). A face-centered representation—one that can suffer a distortion that is referenced with respect to an axis inside the face—is a specific case of what we have meant by object- or system-centered, of which an orbit-centered representation is an additional specific case.

Because faces are often described as "special," one might question the generality, pervasiveness, and flexibility of system-centered formats in visual perception. This would be unfortunate as well, because faces are not an ideal stimulus for investigating the general properties of such representations nor how they may place constraints on upstream processing. Research with neurodivergent individuals has filled the gap, illustrating how object-centered representations support and constrain the perception of elongated objects, both abstract and real world.

McCloskey reports extensively regarding an individual with the initials A.H., an individual with what appears to be a vision-specific developmental deficit (McCloskey, 2004, 2009; McCloskey et al., 2006). The key evidence emerging from her case is a preponderance of *object-centered* reflection errors in the perception of elongated objects. When shown a picture like the one in Figure 23A, A.H. frequently reports seeing the reflections

Figure 23

Asymmetric Objects and Reflection Errors



Note. (A) An object of the variety often used to investigate object-centered representation in neurodivergent individuals and working memory experiments. The object has an elongated axis and parts that elicit representation in an internal reference frame such as the one shown in red. A representation of the current state of the object in the world requires a description of how the internal reference frame is positioned and oriented relative to an external reference frame, such as the one shown in black. For example, one could describe the object as possessing an OPA that has its negative end on the negative side of the external horizontal axis, with a tilt of 45° counterclockwise relative to the external vertical axis. Similarly, the OSA has its positive end on the positive side of the external horizontal and it is tilted 45° clockwise relative to the external vertical. (B) and (C) show reflection errors that are common given the stimulus in A (the original is now shown in light gray for comparison). A reflection across the OPA (B) occurs if the polarity assignment between the OSA and the external reference were to become confused (here, if one labeled the positive side of the OSA as being on the negative side of the external horizontal). Likewise, a reflection across the OSA (C) reveals an error in the coding of how the OPA relates to the external axes. Note that neither (B) nor (C) can be described as pure rotation errors. (D) An example of a target stimulus and possible error types from working memory experiments. Stimuli like this one include asymmetry that allows for distinguishing between reflections across internal axes and reflections across external ones. Reflections across external axes tend to be rare in these experiments, especially compared to errors across internal axes (e.g., Gregory & McCloskey, 2010). OPA = object principal axis; OSA = object secondary axis. See the online article for the color version of this figure.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

shown in Figure 23B or 23C. Neither a noisy representation of orientation nor a reflection across an external axis will produce these specific errors. They will emerge naturally, however, through mislabeling of the relationship between an object-internal axis to an external axis. In addition to A.H., McCloskey et al. (2006) report on two individuals (with acquired deficits) whose systematic patterns of reflection errors constitute evidence for composite and internally referenced representations of orientation. More recently, Vannuscorps and colleagues reported on an individual named Davida with a developmental deficit strikingly like that of A.H., but not identical (Vannuscorps et al., 2021, 2022). They describe her patterns of systematic and independent reflection errors as evidence of mid-level shape-centered visual representations, where shape-centered refers to description relative to a shape-internal axis and attendant descriptions that situate the shape relative to external references at a given moment.

Studies of visual working memory in neurotypical adults and children demonstrate the presence of object-centered representations in everyday perception, and how these representations place constraints on cognitive performance (Chaisilprungraung et al., 2019; Gregory & McCloskey, 2010; Gregory et al., 2011). Figure 23D shows an example of the kind of stimuli used in these studies. Asymmetry serves a key role in the stimuli because it distinguishes between errors of reflection across *internal* axes, reflections across *external* axes, pure rotations, and combinations thereof. The punchline: Internal axis reflections, especially across a principal axis of elongation, are far more frequent than other error types. Recall that such errors arise from a mislabeling of how an internal axis is aligned at a given moment in relation to an external one. For example, believing that the positive end of an object-secondary axis is on the negative side of an external vertical axis, when it is actually on the positive side of the external—that will cause reflection across the object's principal axis.

The transition between representing an object in its own frame of reference to situating it in the present in relation to external reference frames, therefore, appears tenuous and error prone. Our claims with respect to tracking rotational motion hang on this point. Rotation erects an internally referenced representation like an object does. Tracking becomes a process of updating internal references in relation to external ones, a process that is inherently error prone.

Conclusion

The experiments support the hypothesis that orbiting motion induces perceptual representations that are composite and internally referenced, even in the absence of visible structural connections. The consistent tracking deficit for objects in circular orbits, compared to translational motion, suggests that perception treats orbits as structured systems, requiring additional computational steps to maintain separate identities for satellites within the system. This parallels how the features of faces and the parts of objects (such as tools) are represented in object-centered terms, in relation to internal references rather than independent coordinates. Although tracking is typically treated as a cognitive function that is constrained by limited resources (such as attention or working memory), tracking objects in orbit is a case where constraints are imposed by the underlying format of a perceptual representation.

References

Almeida, J., Freixo, A., Tábuas-Pereira, M., Herald, S. B., Valério, D., Schu, G., Duro, D., Cunha, G., Bukhari, Q., & Duchaine, B. (2020). Face-specific perceptual distortions reveal a view-and orientation-independent face template. *Current Biology*, *30*(20), 4071–4077.e4. https://doi.org/10.1016/j.cub.2020.07.067

Alvarez, G. A., & Cavanagh, P. (2005). Independent resources for attentional tracking in the left and right visual hemifields. *Psychological Science*, *16*(8), 637–643. https://doi.org/10.1111/j.1467-9280.2005.01587.x

Bae, G. Y., & Flombaum, J. I. (2012). Close encounters of the distracting kind: Identifying the cause of visual tracking errors. *Attention, Perception, & Psychophysics*, *74*(4), 703–715. https://doi.org/10.3758/s13414-011-0260-1

Ballaz, C., Boutsen, L., Peyrin, C., Humphreys, G. W., & Marendaz, C. (2005). Visual search for object orientation can be modulated by canonical orientation. *Journal of Experimental Psychology: Human Perception and Performance*, *31*(1), 20–39. https://doi.org/10.1037/0096-1523.31.1.20

Behrmann, M., & Tipper, S. P. (1994). Object-based attentional mechanisms: Evidence from patients with unilateral neglect. In C. Umiltà & M. Moscovitch (Eds.), *Attention and performance 15: Conscious and nonconscious information processing* (pp. 351–375). The MIT Press.

Behrmann, M., & Tipper, S. P. (1999). Attention accesses multiple reference frames: Evidence from visual neglect. *Journal of Experimental Psychology: Human Perception and Performance*, *25*(1), 83–101. https://doi.org/10.1037//0096-1523.25.1.83

Blom, J. D., Ter Meulen, B. C., & Dool, J. (2021). A century of prosopometamorphopsia studies. *Cortex*, *139*, 298–308. https://doi.org/10.1016/j.cortex.2021.03.001

Chaisilprungraung, T., German, J., & McCloskey, M. (2019). How are object shape axes defined? Evidence from mirror-image confusions. *Journal of Experimental Psychology: Human Perception and Performance*, *45*(1), 111–124. https://doi.org/10.1037/xhp0000592

Dürsteler, M. R. (2008). The freezing rotation illusion. *Progress in Brain Research*, *171*, 283–285. https://doi.org/10.1016/S0079-6123(08)00640-7

Farah, M. J., Tanaka, J. W., & Drain, H. M. (1995). What causes the face inversion effect? *Journal of Experimental Psychology: Human Perception and Performance*, *21*(3), 628–634. https://doi.org/10.1037/0096-1523.21.3.628

Gibson, J. J. (1979). *The ecological approach to visual perception*. Houghton Mifflin.

Gregory, E., Landau, B., & McCloskey, M. (2011). Representation of object orientation in children: Evidence from mirror-image confusions. *Visual Cognition*, *19*(8), 1035–1062. https://doi.org/10.1080/13506285.2011.610764

Gregory, E., & McCloskey, M. (2010). Mirror-image confusions: Implications for representation and processing of object orientation. *Cognition*, *116*(1), 110–129. https://doi.org/10.1016/j.cognition.2010.04.005

Holcombe, A. O., & Chen, W.-Y. (2012). Exhausting attentional tracking resources with a single fast-moving object. *Cognition*, *123*(2), 218–228. https://doi.org/10.1016/j.cognition.2011.10.003

Keane, B. P., & Pylyshyn, Z. W. (2006). Is motion extrapolation employed in multiple object tracking? Tracking as a low-level, non-predictive function. *Cognitive Psychology*, *52*(4), 346–368. https://doi.org/10.1016/j.cogpsych.2005.12.001

Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information*. W. H. Freeman and Company.

McCloskey, M. (2004). Spatial representations and multiple-visual-systems hypotheses: Evidence from a developmental deficit in visual location and orientation processing. *Cortex*, *40*(4–5), 677–694. https://doi.org/10.1016/S0010-9452(08)70164-3

McCloskey, M. (2009). *Visual reflections: A perceptual deficit and its implications*. Oxford University Press.

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

McCloskey, M., & Rapp, B. (2000). Attention-referenced visual representations: Evidence from impaired visual localization. *Journal of Experimental Psychology: Human Perception and Performance*, *26*(3), 917–933. https://doi.org/10.1037//0096-1523.26.3.917

McCloskey, M., Valtonen, J., & Cohen Sherman, J. (2006). Representing orientation: A coordinate-system hypothesis and evidence from developmental deficits. *Cognitive Neuropsychology*, *23*(5), 680–713. https://doi.org/10.1080/02643290500538356

Murphy, J., Gray, K. L., & Cook, R. (2017). The composite face illusion. *Psychonomic Bulletin & Review*, *24*(2), 245–261. https://doi.org/10.3758/s13423-016-1131-5

Ongchoco, J. D. K., & Scholl, B. J. (2019). How to create objects with your mind: From object-based attention to attention-based objects. *Psychological Science*, *30*(11), 1648–1655. https://doi.org/10.1177/0956797619863072

Pavlova, M., & Sokolov, A. (2003). Prior knowledge about display inversion in biological motion perception. *Perception*, *32*(8), 937–946. https://doi.org/10.1068/p3428

Poljac, E., de-Wit, L., & Wagemans, J. (2012). Perceptual wholes can reduce the conscious accessibility of their parts. *Cognition*, *123*(2), 308–312. https://doi.org/10.1016/j.cognition.2012.01.001

Pylyshyn, Z. W. (2006). Some puzzling findings in multiple object tracking (MOT): II. Inhibition of moving nontargets. *Visual Cognition*, *14*(2), 175–198. https://doi.org/10.1080/13506280544000200

Pylyshyn, Z. W., & Storm, R. W. (1988). Tracking multiple independent targets: Evidence for a parallel tracking mechanism. *Spatial Vision*, *3*(3), 179–197. https://doi.org/10.1163/156856888X00122

Scholl, B. J., Pylyshyn, Z. W., & Feldman, J. (2001). What is a visual object? Evidence from target merging in multiple object tracking. *Cognition*, *80*(1–2), 159–177. https://doi.org/10.1016/S0010-0277(00)00157-8

Schwarzer, G. (2000). Development of face processing: The effect of face inversion. *Child Development*, *71*(2), 391–401. https://doi.org/10.1111/1467-8624.00152

Shipley, T. F. (2003). The effect of object and event orientation on perception of biological motion. *Psychological Science*, *14*(4), 377–380. https://doi.org/10.1111/1467-9280.24471

Tadin, D., Lappin, J. S., Blake, R., & Grossman, E. D. (2002). What constitutes an efficient reference frame for vision? *Nature Neuroscience*, *5*(10), 1010–1015. https://doi.org/10.1038/nn914

Troje, N. F. (2003). Reference frames for orientation anisotropies in face recognition and biological-motion perception. *Perception*, *32*(2), 201–210. https://doi.org/10.1068/p3392

Upadhyayula, A., & Flombaum, J. (2020). A model that adopts human fixations explains individual differences in multiple object tracking. *Cognition*, *205*, Article 104418. https://doi.org/10.1016/j.cognition.2020.104418

Vallortigara, G., & Regolin, L. (2006). Gravity bias in the interpretation of biological motion by inexperienced chicks. *Current Biology*, *16*(8), R279–R280. https://doi.org/10.1016/j.cub.2006.03.052

Vannuscorps, G., Galaburda, A., & Caramazza, A. (2021). The form of reference frames in vision: The case of intermediate shape-centered representations. *Neuropsychologia*, *162*, Article 108053. https://doi.org/10.1016/j.neuropsychologia.2021.108053

Vannuscorps, G., Galaburda, A., & Caramazza, A. (2022). Shape-centered representations of bounded regions of space mediate the perception of objects. *Cognitive Neuropsychology*, *39*(1–2), 1–50. https://doi.org/10.1080/02643294.2021.1960495

Vul, E., Alvarez, G., Tenenbaum, J., & Black, M. (2009). Explaining human multiple object tracking as resource-constrained approximate inference in a dynamic probabilistic model. *Advances in Neural Information Processing Systems*, *22*, 1955–1963. https://proceedings.neurips.cc/paper/2009/hash/d79aac075930c83c2f1e369a511148fe-Abstract.html

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M., & Von der Heydt, R. (2012). A century of gestalt psychology in visual perception: I. Perceptual grouping and figure–ground organization. *Psychological Bulletin*, *138*(6), 1172–1217. https://doi.org/10.1037/a0029333

Watson, S. E., & Kramer, A. F. (1999). Object-based visual selective attention and perceptual organization. *Perception & Psychophysics*, *61*(1), 31–49. https://doi.org/10.3758/BF03211947

Wu, Q., & Flombaum, J. I. (2024). The motion-silencing illusion depends on object-centered representation. *Psychological Science*, *35*(5), 504–516. https://doi.org/10.1177/09567976241235104

Wu, Q., & Flombaum, J. I. (2026, March 26). *Tracking and OCR*. OSF. https://osf.io/zjhr9

(*Appendices follow*)

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.

Appendix A

Table of Preregistrations, Demonstrations, Data, and Code for All Experiments

Experiment	Preregistration link	Demo link	Code	Data
Experiment 1: Deficient tracking of identity in a rotational orbit	https://aspredicted.org/yzfd-snmz.pdf	https://qihan.visualthinkingresearch.com/tracking_COR/exp_1/tracking_exp_1.html.	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8
Experiment 2: Deficient tracking of proper objects	https://aspredicted.org/29jf-y338.pdf.	https://qihan.visualthinkingresearch.com/tracking_COR/exp_2/tracking_exp_2.html	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8
Experiment 3: Replication of Experiment 1	Not preregistered	Same as Experiment 1	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8
Experiment 4: Side-by-side horizontal placement	Not preregistered	https://qihan.visualthinkingresearch.com/tracking_COR/exp_4/tracking_exp_4.html.	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8
Experiment 5: Replication with vertical placement	Not preregistered	https://qihan.visualthinkingresearch.com/tracking_COR/exp_5/tracking_exp_5.html.	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8
Experiment 6: Explaining the freezing illusion through orbit-centered reference	Not preregistered	https://qihan.visualthinkingresearch.com/dots2/exp4/Dots_exp4.html.	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8	https://osf.io/zjhr9/overview?view_only=48c8d87f282c4989ba149ef71b912ea8

Appendix B

Table of Demonstrations Intended for Explanatory Purposes

Video name	Video link and description
Demo1—Wu&Flombaum—Basic-Composite-Orbit	https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_1.html A composite representation of a rotating pair. The blue and red satellites on the right rotate around a shared origin. The mental representation is composite: including a static description of the pair's initial arrangement (with respect to an axis of internal reference), together with a dynamic description of the changing orientation of the internal axis relative to the external environment.
Demo2—Wu&Flombaum—Composite-Tracking	https://qihan.visualthinkingresearch.com/tracking_COR/demo/trackingCOR_demo_2.html A composite representation of a rotating pair. The identical satellites on the right rotate around a shared origin. At the start of the trial, the one on the left is marked with a T to indicate that it is the Target. The composite mental representation marks the Target in the static description of the pair's initial arrangement (with respect to an axis of internal reference). The rotation is tracked through a dynamic description of the changing orientation of the internal axis relative to the external environment. Identifying the target at the end of a trial will rely on an inference across the parts of the composite.
Demo3—freezingRotationIllusion_demo	https://qihan.visualthinkingresearch.com/dots2/demo/displayVideo.html When the outer square changes direction so that it goes from moving opposite to moving with the disk pair, then the disks appear to freeze momentarily, even though they never stop moving or change speed in practice.

Received April 30, 2025
Revision received February 6, 2026
Accepted March 7, 2026 ■

This document is copyrighted by the American Psychological Association or one of its allied publishers.
This article is intended solely for the personal use of the individual user and is not to be disseminated broadly.
All rights, including for text and data mining, AI training, and similar technologies, are reserved.