Sensory and choice responses in MT distinct from motion encoding

Aaron J Levi^{1†}, Yuan Zhao^{2†}, Il Memming Park², Alexander C Huk^{1*}

*For correspondence:
huk@utexas.edu (FM)

[†]These authors contributed
equally to this work

¹ Center for Perceptual Systems, Departments of Neuroscience & Psychology, The University of Texas at Austin, Austin, Texas, United States of America; ² Department of Neurobiology and Behavior, Stony Brook University, Stony Brook, New York, United States of America

Abstract Macaque area MT is well known for its visual motion selectivity and relevance to motion perception, but the possibility of it also reflecting non-sensory functions has largely been ignored. Manipulating subjects' temporal evidence weighting revealed multiple components of MT responses that were, surprisingly, not interpretable as behaviorally-relevant modulations of motion encoding, nor as consequences of readout of motion direction. MT's time-varying motion-driven responses were starkly changed by our strategic manipulation, but with timecourses opposite the subjects' temporal weighting strategies. Furthermore, large choice-correlated signals were represented in population activity distinctly from motion responses (even after the stimulus) with multiple phases that both lagged psychophysical readout and preceded motor responses. These results reveal multiple cognitive contributions to MT responses that are task-related but not functionally relevant to encoding or decoding of motion for psychophysical direction discrimination, calling into question its nature as a simple sensory area.

Introduction

Primate area MT plays a critical role in the perception of visual motion. A long line of study has established that MT's encoding of motion direction is quantitatively consistent with perceptual sensitivity, that noise in its responses is correlated with behavioral variability, and that causal perturbations of its activity affect motion perception in lawful and substantial ways (*Newsome and Pare, 1988; Britten et al., 1992, 1996; Salzman et al., 1992*). Owing to this powerfully integrated set of findings, many models and experiments can safely assume that MT is the key place that the brain looks to for information about visual motion. However, these successes do not logically imply that MT only carries sensory information, leaving our understanding of MT at risk of overlooking additional signals and computations that are not aligned with representing motion for the sake of motion perception. In this work, we show that a manipulation of temporal strategy during motion discrimination reveals large signals in MT that are precisely related to components of performing the task, but which neither directly impact psychophysical performance nor reflect straightforward links between perceptual decisions and the sensory responses which informed them.

In addition to the large, classic body of work describing the form and fidelity of MT's representation of visual motion (*Born and Bradley, 2005; Cormack et al., 2017*), some prior work has identified cognitive modulations of MT's sensory-driven activity. Such modulations are still interpretable with respect to MT's representation of visual motion direction, however. Most notably, attention can modify the sensory-driven responses of MT neurons, typically boosting the gain of responses (*Treue and Maunsell, 1996; Seidemann and Newsome, 1999; Cook and Maunsell, 2004*).

43 These modulations of stimulus-driven activity modify MT's representation of motion, and thus play
44 out in behavior as if the visual motion itself had been modified. In contrast, recent work has shown
45 that MT's choice-correlated activity is distinguishable at the population level from its sensory-driven
46 responses, and follows a different time course than the read-out of motion, as inferred from the
47 psychophysical behavior (*Zhao et al., 2020*). While this intriguing initial observation suggests the
48 existence of task-related signals not directly related to motion encoding, interpretation of this
49 choice-related activity is constrained by the lack of any direct experimental manipulation of the
50 decision-making process.

51 To directly test for and characterize non-sensory signals in MT, we manipulated the time course
52 of psychophysical weighting while monkeys performed a direction-discrimination task, coupled
53 with ensemble recordings of multiple neurons in MT analyzed via population-coding techniques.
54 We explicitly manipulated whether early or late parts of the stimulus had stronger or weaker mo-
55 tion evidence on average, which affected the time course of how the visual motion stimulus was
56 weighted for task performance, as assessed via psychophysical reverse correlation. This manipu-
57 lation of temporal weighting strategy provided critical interpretive leverage for distinguishing the
58 time courses of decision formation and choice-correlated activity, and also caused a surprising
59 and strong modulation of the sensory responses themselves that was also not directly related to
60 forming decisions about motion.

61 When perceptual weighting was unconstrained, direction-discrimination behavior was based
62 primarily on early portions of the stimulus, the sensory representation showed a standard and
63 modest falloff over the course of stimulus presentation, and a distinct and substantial choice-
64 correlated response emerged during late portions of stimulus viewing. When we shifted the tem-
65 poral readout strategy to favor late portions of the stimulus, behavior relied preferentially on later
66 portions of the stimulus, but *later* portions of the sensory response were *decreased*, as opposed to
67 increased. Choice-correlated activity was significantly muted during the late-weighting condition.
68 However, choice-correlated activity was present after the stimulus, leading up to the response (a
69 novel phenomenon evident across all strategic conditions, in fact). When subjects' temporal weight-
70 ing strategy was then manipulated to preferentially rely on earlier portions of the stimulus, later
71 portions of the sensory response were increased, and choice-correlated activity was again evident
72 during the late portions of the stimulus. This last condition's effects were most striking, as a steep
73 falloff in perceptual weighting over time was accompanied by an increase in late sensory-driven
74 activity that led to a non-monotonic time course of motion-driven response.

75 The opposite effects of our experimental manipulations on temporal weighting strategy and
76 the timecourse of sensory gain run counter to any standard encoding model of MT simply repre-
77 senting behaviorally-relevant motion: In that framework, motion responses ought to mirror the
78 psychophysical weighting. Choice-correlated activity during the stimulus was also controlled by
79 changes in the psychophysical weighting, and across these psychophysical time courses, was al-
80 ways lagged relative to the periods when the subjects were "reading out" MT activity. But this
81 decision-lagged choice-related signal was not simple feedback linking a sensory response and a
82 subsequent, corresponding decision, not just because the choice signals affected MT population
83 activity differently than visual motion did; we also observed a distinct choice-related signal after
84 stimulus offset that was linked to impending response, and which was distinct from simple premo-
85 tor activity.

86 Together, these multiple components of the MT response, revealed while manipulating the
87 temporal weighting strategy, could be seen as lawful functions of the time course of decision for-
88 mation and the anticipation of the response. However, these response components could not be
89 interpreted as either modulations of the encoding that played out in perceptual reports, nor to the
90 effects of read out mechanisms that would either correlationally (via feed-forward mechanisms) or
91 causally (via straightforward feedback mechanisms) align with the sensory response. Thus, there
92 appear to be multiple, large components of MT activity that affect both its stimulus-driven activity
93 and which are separable from it— even during a well-studied direction-discrimination task with tight

94 control over motion readout strategy– that are inconsistent with its conventional designation as a
95 simple, low-dimensional, sensory encoding area.

96 Results

97 We measured the timecourse of sensory and choice-correlated responses from simultaneously
98 recorded groups of MT neurons using linear and nonlinear decoding approaches while rhesus
99 monkeys performed a motion direction discrimination task. We manipulated the time course of
100 stimulus evidence, and the subjects shifted their temporal weighting strategy to rely preferentially
101 on the stronger periods of stimulus motion. We began recordings in each subject with a baseline
102 “flat” stimulus phase for several experimental sessions, in which stimuli had a constant average
103 motion strength over time within a trial, as is the case in most related experiments (*Gold and*
104 *Shadlen, 2007*). We then shifted to several sessions in a “late” regime, in which the stronger motion
105 was present in the second half of the stimulus. Finally, we performed several sessions in an “early”
106 regime, in which the stronger motion was present in the first half.

107 Observers change temporal weighting strategies according to stimulus statistics

108 Two trained rhesus macaques (one male, one female) viewed sequences of seven motion pulses
109 and indicated perceived net motion with a saccade to one of two response targets (Figure 1A).
110 We measured traditional psychometric performance (i.e., accuracy as a function of net motion
111 strength on each trial), and the time course of weighting within each trial (i.e., using logistic regres-
112 sion between motion strength at each pulse and the binary choices, see Methods). We refer to the
113 resulting set of regression coefficients, or weights, as the temporal weighting strategy.

114 The motion discrimination task was performed in three contexts (Figure 1B-D). First, in the
115 flat-stimulus condition (Figure 1B), average motion over time was equal within a trial. Many tra-
116 ditional sensory decision-making studies use stimuli with uniform information over time, and thus
117 the flat-stimulus condition served as a baseline in our experiments. Subjects’ temporal weighting
118 strategies were biased to have higher weight on early stimulus periods, despite uniform motion
119 expectation over time (Figure 1E). This default early weighting strategy is consistent with many
120 other findings (*Huk and Shadlen, 2005; Kiani et al., 2008; Nienborg and Cumming, 2009; Yates*
121 *et al., 2017; Levi et al., 2018; Kawaguchi et al., 2018*) and likely reflects a combination (*Levi and*
122 *Huk, 2020; Okazawa et al., 2018*) of improved sensory encoding at stimulus onset (*Osborne et al.,*
123 *2004; Churchland et al., 2010*), and the consequences of early termination of the decision process,
124 due to mechanisms like bounded accumulation (*Kiani et al., 2008*).

125 Next, we performed a series of experimental sessions in which the stimulus statistics were
126 manipulated such that the average motion strength was high for the last three pulses, while the first
127 three were near zero. We refer to this as the late-stimulus condition (Figure 1C). Although the first
128 3 pulses had motion strength near zero on average (regardless of full-trial, net motion strength),
129 on individual trials there was still variable nonzero motion possible for any pulse. Subjects were
130 rewarded based on the actual net motion direction presented on that particular trial, as opposed
131 to the average or expected value based on the condition from which the trial was generated. This
132 produced robust behavioral changes that tracked motion expectation in the stimulus design, such
133 that weight on the first three pulses decreased substantially, and the highest psychophysical weight
134 was placed on the later pulses (Figure 1F).

135 Finally, we performed a series of sessions in which the stimulus statistics were changed such
136 that the average motion strength was now high in the early half of the stimulus, and was near zero
137 for the last half of the stimulus; we refer to this as the early-stimulus condition (Figure 1D). This
138 successfully changed temporal weighting behavior back to pronounced early weighting, in which
139 the first pulses received drastically higher weight than the remainder of the stimulus (Figure 1G), in
140 a manner overall similar to the default strategy during the flat-stimulus (flat: $-0.091 [-0.113, 0.069]$,
141 late: $0.083 [0.015, 0.151]$, early: $-0.091 [-0.136, -0.081]$; slope of linear fit to the psychophysical ker-
142 nel [95% CIs]). In summary, the temporal weighting strategy shifted in concert with the time course

Figure 1. Sequence of trial events, temporal stimulus statistics, and successful manipulation of behavioral weighting strategy. **A**, Subjects fixated on a central point through the appearance of targets and motion stimulus until the disappearance of the fixation point (“go”). Choices were made with saccades to the target corresponding to the perceived net direction of motion. Initial fixation time, target-on duration, and time until fixation point disappearance were randomly varied. **B-D**, Average stimulus strength per pulse (bold lines) and individual trial examples (semi-transparent lines) for trials of different strength and direction (denoted by sign). In the flat-stimulus (**B**), motion strength is constant over time on average. In the late-stimulus (**C**) motion strength is reduced on average in the first three pulses such that the highest motion expectation is late. In the early-stimulus (**D**) motion strength is reduced in the last three pulses such that the highest motion expectation is early. Motion pulse values in individual trials (semitransparent traces) vary considerably (see Methods for detail). **E-G**, Temporal weighting behavior across conditions. **E**. Subjects preferentially weight the early pulses despite uniform motion expectation over time. **F**. Temporal weighting shifts during the late-stimulus condition to preferentially weight late pulses. **G**. Behavior reverts back to early-weighting when the stimulus statistics are biased towards high motion strength early.

of expected motion strength, placing higher weight on portions of the stimulus when higher motion strength was expected based on the experimental phase. This confirms that our manipulation of stimulus statistics affected the time course of psychophysical readout, allowing us to better interpret the time scale of neural responses relative to the timing of when the subject was "reading out" MT for the purpose of forming a decision about motion direction.

Choice-correlated activity in MT is large but does not align with stimulus encoding or behavioral readout

We recorded ensembles of single and multi-unit activity from area MT while monkeys performed the direction discrimination task, across the manipulation of temporal weighting strategy described in the previous section. We used both linear and nonlinear ensemble decoding frameworks to extract information about direction and choice from groups of simultaneously recorded MT neurons (Figure 2A). As a simple starting point, we used logistic regression (logReg) between the raw trial spike count vectors and either the stimulus direction (the "direction" axis) or the psychophysical choice (the "choice" axis) to find a reweighted population response that best mapped neural activity to the binary stimulus or choice (Figure 2A, left). Such linear models are likely easy for the brain to implement, but are limited in how they can capture relations between neurons and between neural activity and experimental factors. We therefore also used a more advanced nonlinear dimensionality reduction technique (variational latent Gaussian process model, vLGP) to extract smooth low-dimensional latent factors that explain correlations within the population spike trains (*Zhao and Park, 2017; Zhao et al., 2020*) (Figure 2A, right). It functions in a conceptually analogous manner to the simpler logistic regression approach (i.e., mapping ensemble activity to the stimulus or the choice), but has the ability to more effectively capture the complex joint statistics of the neural population while also providing access to a more concise representation of the relations between neural activity, stimulus direction, and psychophysical choices (by virtue of an intervening dimensionality reduction step to identify latent factors making up the ensemble activity).

Both analytic approaches revealed the presence of substantial choice-correlated activity in the MT population response, achieving large peak magnitudes (> 0.6 as measured by choice probability, CP; although we use this conventional metric in this paper, we emphasize that by calculating it on various dimensions of the ensemble response, we have generalized it beyond the classical approach of only looking at choice-correlated activity defined along the stimulus axis) (*Britten et al., 1996*). The largest choice-correlated activity was present in the population activity in a manner distinct from how the stimulus drove the ensemble of MT neurons. Via logReg, this was evident in significantly larger CP along the choice axis over the direction axis (Figure 2B-D), stemming from a weak correspondence between a neuron's weight in one model compared to the other ($r = 0.146$). The vLGP analysis showed that stimulus encoding was well described by a single dimension (termed the stimulus axis), but the stimulus axis had relatively small choice information when compared to the combined choice information in the top four latent factors altogether (*Zhao et al., 2020*) (Figure 2B-D).

Importantly, both analysis methods revealed that across pronounced changes in temporal weighting strategy, the time course of choice-correlated activities never mirrored the time course of psychophysical readout (Figure 2B-D, 1E-G). Instead, choice-correlated activity was consistently highest after the stimulus periods that were weighted the highest in the behavior. In the flat condition, both analysis approaches demonstrated increased choice probability during the last half of the stimulus, despite early weighting in the behavior. In the late condition, when behavior exhibited the strongest dependence on later portions of the stimulus, the strongest choice-correlated activity was still distinct from the stimulus-driven activity, and exhibited a more muted and flatter time course, though still characterized by an even later peak relative to the flat-stimulus condition. Finally, when subjects returned to an early weighting strategy in the early stimulus condition, the time course of choice probability returned to a similar rising profile, as originally measured during the flat condition. These observations are inconsistent both with classical interpretations that

choice probabilities reflect the feedforward consequences of sensory noise being read out as information about the stimulus (because the bulk of the choice-correlated activity arose after the psychophysical readout of MT was likely happening), as well as more recent interpretations that choice probabilities reflect feedback, because differential MT responses correlated with choice were not strongly aligned with the motion responses that gave rise to those decisions.

Changes in sensory encoding run opposite changes in temporal weighting strategy

Most surprisingly, we observed large changes to MT's time-varying sensory response that were incommensurate with perceptual readout. Here, we relied on the vLGP analysis to describe the temporal dependence of the population response on the motion pulses by looking at the directional response along the stimulus axis. We calculated a pulse-triggered average (PTA) to summarize the regression coefficients that predict the change in latent states (*Yates et al., 2017*). Each "bump" in Figure 3 represents the modulation of the stimulus-axis latent factor by a pulse of unit motion strength (i.e., a single Gabor drifting in one direction) for each pulse in the stimulus presentation (Figure 1). As temporal weighting strategy shifted across conditions, one might expect nothing to change in MT, consistent with a constant (and thus largely veridical) representation of visual information despite changes in readout/weighting strategy. An alternative hypothesis based on temporal attention would predict gain modulation congruent with behaviorally up-weighted and down-weighted stimulus epochs (*Ghose and Maunsell, 2002*). Instead, to our surprise, we observed changes to sensory encoding with an unintuitive, if almost paradoxical, link to psychophysical direction discrimination.

In the flat stimulus condition there was a modest decrease in the sensory response over time (i.e., PTA magnitude fell across the 7 pulse epochs; Figure 3A). Such a gradually-declining time course is consistent with known adaptation phenomena in many visual brain areas, and has been observed in MT during viewing of this same stimulus (*Yates et al., 2017*). However, during the late-stimulus condition, the sensory response decreased for the late pulses relative to the flat condition time course (Figure 3B). The behavioral profile shows precisely the opposite: relative down-weighting of early pulses and up-weighting of later pulses. And most strikingly, when subjects switched to the early-stimulus condition, the sensory response showed a stark up-weighting of later pulses, resulting in a dramatically non-monotonic, U-shaped profile (Figure 3C). Once again, this is directly at odds with the temporal weighting of behavior, which sharply favors the first 2-3 pulses over the rest. This modulation is counterintuitive from standard perspectives, which would predict that if any changes in sensory response are evident, they would be reflected by increases in response to stimulus portions that were weighted more strongly for decision making.

Instead of gain changes that reflect behavioral readout strategy, the sensory response modulations we observed make more sense viewed as compensating for "missing" signal relative to a time-stationary motion expectation. In our experiments, both animals were trained extensively on the flat condition before undergoing temporal manipulation. The change in gain thus manifested as a function of the mismatch between this apparently "default" temporally-uniform expectation of motion and the statistics of the currently-encountered condition. In more detail, during the late condition motion strength was decreased in the early portions of the stimulus, but the PTA revealed decreased gain on later pulses instead (Figure 3B). During the early condition, the motion strength on later pulses was decreased, but the PTA revealed a striking gain increase on these portions of the stimulus for which the expected motion was quite weak (Figure 3C). Thus, while the temporal weighting evident in behavior changed across conditions in a way that tracked changes in stimulus statistics (i.e., weighting the stronger periods of motion more, and weaker periods of motion less), MT's response to motion was changed inversely to those patterns.

Large choice-correlated activity also exists in the absence of the motion stimulus

We also observed another choice-related signal in MT of substantial magnitude. The vLGP analysis revealed significant choice-correlated activity after the offset of the motion stimulus, in anticipation

Figure 3. Time course of motion-driven MT response changes opposite that of changes in temporal weighting strategy. A-C. The pulse-triggered average (PTA) describes the modulation of the stimulus-axis latent factor by a pulse of unit motion strength for each of the seven pulses in the visual motion stimulus. **A.** The PTA for the flat-stimulus condition reflects the expected transient-to-sustained response, where a pulse at the beginning of the stimulus affects the MT response more than a pulse closer to the end of the stimulus. **B.** In the late-stimulus condition, the relative drop from early pulses to later ones is even more exaggerated than in the flat, despite highest motion strength occurring late in the trial. **C.** The PTA during the early-stimulus condition exhibits substantial increase on later pulses, despite a lack of high motion signal in the stimulus during those pulses.

of an upcoming saccade. There was a minimum 500 ms window between the stimulus offset and the disappearance of the fixation point which signaled the monkey could move their eyes to make their choice, and during this window we saw choice probabilities up to > 0.7 (Figure 4A).

The magnitude of post-stimulus choice probability is comparable to, and often greater than, what we observed from our decoders during the stimulus period, and is quite high compared to traditional measures of choice probability based on single neuron measurements. Most importantly, the finding of large amounts of choice-correlated activity without the presence of a visual stimulus in MT strengthens the case for such signals being non-sensory in origin. The choice signal measured during the delay period is present when there is no sensory drive whatsoever, further ruling out interpretations of choice probabilities as a product of noise in sensory representations. Instead, its full magnitude (revealed by "looking" off the stimulus axis), late time course, and presence up to the response are more similar to choice-related activity seen in a multitude of areas that are often considered much more cognitive or associative in nature, such as LIP and PFC (Roitman and Shadlen, 2002; Mante et al., 2013).

Interestingly, the onset of CP during the delay period varies with reaction time (RT) in a way that suggests the choice signal is not simple premotor activity. If this were the case, we would expect that CP would increase later on trials with longer RTs compared to trials with shorter RTs. Instead, when reaction times were longer than the median RT, the saccade-aligned CP increased noticeably earlier than on trials with reaction times in the shorter half of the RT distribution (Figure 4B). This was true of all three temporal stimulus conditions. The result is striking, especially given the fixed-stimulus experimental design and the coarse division of "short" and "long" RTs by median. Temporally divorced from stimulus processing and not tightly linked to motor behavior timing, this delay-period choice signal appears to have a more cognitive origin reflecting the maintenance of choice information between stimulus and response.

Time-varying readout of population activity confirms the dynamics of choice-related signals

In all analyses leading up to this point, the weights used to decode the stimulus or the choice were calculated using the neural responses and/or the derived latent factors from the entire stimulus

Figure 4. Presence of large choice-related signals in MT during post-stimulus delay period. A. CP along the choice (dashed lines) and direction (solid lines) axes, aligned to the time of the saccade. In all three conditions, there is high CP in the choice axis for the entire 500ms preceding the saccade, without any stimulus drive. CP increased over the last 200ms leading up to the saccade. There was relatively little CP along the stimulus axis. **B.** Saccade-aligned CP along the choice axis only, separated by median reaction time (RT). CP for longer RT trials (dashed lines) increased earlier than that of shorter RT trials (dotted lines). This was true in all three conditions.

period. Even with this fixed temporal readout scheme, we saw nuanced temporal dynamics in both sensory- and choice-related activity that differed from the time course of temporal weighting evident in the psychophysical behavior. Although, from a decoding perspective, using temporal fixed weights makes for a readout process that the brain might find easier to implement, we know very little about how sophisticated the brain's decoding machinery might be (and indeed, our ability to manipulate the timecourse of motion weighting suggests that temporally-static decoding is not a hard limit). Furthermore, from a purely statistical perspective, we were also motivated to consider decoding with a temporally dynamic readout scheme to confirm that the rich dynamics we observed were neither constrained nor distorted by the assumption of constant read-out weights. We therefore performed further latent factor analyses in which weights were fitted and applied based on the activity within individual 100 ms bins for both the delay and motion periods (Figure 5).

The timecourse of choice-correlated activity was quite similar from fixed to dynamic readout models. With temporally varied readout weights, the same pattern persisted: high CP late in the stimulus period regardless of temporal stimulus condition. This is strong support for CP as a top-down signal that arrives in MT mostly after decisions have been made. That is, after the pulses with the highest weight in the psychophysical kernel. In this interpretation, during the late condition we have in essence delayed the decision and thus further delayed the decision-correlated activity that follows. The time-varying readout schemes also confirmed the dynamics in the post-stimulus, delay period. In all three conditions, CP was high throughout the delay period, but increased over the last 200ms. Along the stimulus axis, CP was flatter and closer to chance. Altogether, the similarity in CP timecourse between fixed and dynamic readout models suggests that a fixed weighting scheme is sufficient to describe the temporal patterns of choice information in MT during motion information both during and after the stimulus.

Discussion

By manipulating the temporal weighting strategy of subjects while they performed a direction discrimination task, aided by ensemble recordings and population-level decoding analyses, we discovered multiple signals in MT that are distinct from its representation of motion direction, solidly established to be used by later decision stages for perceptual reports and behavior. Striking changes in sensory response were associated with the mismatch between the current strength of sensory evidence and prior, learned time courses of sensory evidence. Although these large modulations affected the sensory encoding, they appear not to have affected the psychophysical behavior. Choice-correlated activity was also surprisingly strong, but was delayed relative to temporal weighting behavior, even when the latter was under direct experimenter control. Furthermore, the choice-correlated activity was evident at the population level in a manner that was distinct from stimulus-driven responses in MT, and was "readout-irrelevant" as well, in that it was largest when the subjects were not primarily reading out the stimulus, or even viewing a stimulus at all.

The changes we observed in sensory responses may seem paradoxical at first, as the gain was increased for periods of the stimulus during which the subjects applied the smallest amount of weight in forming decisions. This is opposite the notion of attention affecting gain for parts of a stimulus that are more relevant for decisions (*Treue and Maunsell, 1996; Seidemann and Newsome, 1999*). But, these modulations appear more sensible when viewed as resulting from a mismatch between trained statistics and the current ones. The hypo-responsivity to late pulses in the late condition, and the hyper-responsivity to those same late pulses during the early condition, could both reflect a compensatory response to motion in the current condition compared to the expectation of the temporally uniform stimulus on which animals were trained. Indeed, potentially-related homeostatic mechanisms have been observed in sensory cortex (*Benucci et al., 2013*). Through this lens, the temporal changes in the PTA reflect a recalibration of incoming information to meet the expectation of a temporally-flat stimulus. Thus, even MT's sensory responses are

Figure 5. Time course of choice-related activity in MT is similar when time-varying decoding weights are used. Choice probabilities calculated with time-varying readout weights aligned to motion (left), and the saccade (right) for the flat (A), late (B), and early (C) conditions. CP along the choice axis is represented by dashed lines, while CP along the stimulus axis is represented by solid lines. Choice-axis CP was significantly higher in both the motion- and saccade-aligned time frames. During motion (left), we confirmed that CP was highest during later stimulus epochs, after those with highest psychophysical weight (Figure 1E-G.) During the post-stimulus period (right), we confirmed that CP increased primarily over the last 200ms preceding the saccade to levels even higher than motion-aligned CP.

319 strongly affected by cognitive factors in ways that are dissociable from its well-established, but no
320 longer sole role of representing retinal motion for the sake of perception and/or behavior.

321 Our findings regarding choice-related activity also add to the case for MT carrying substantial
322 non-sensory signals. Having previously used ensemble recordings and population decoding to
323 show that stimulus- and choice-related activity in MT are distinguishable (*Zhao et al., 2020*), our
324 findings in this study add several important facets. First, we exerted explicit control over the time
325 course of psychophysical weighting, which allowed us to experimentally dissociate the psychophys-
326 ical weighting from the time course of choice-correlated activity. By shifting the temporal weight-
327 ing strategy, we effectively changed the average time of the decision, allowing us to confirm that
328 choice signals followed primary decision formation when under explicit experimenter control. Sec-
329 ond, we saw choice activity of substantial magnitude during the post-stimulus delay period. This
330 result rejects virtually any stimulus-based interpretation, as the choice signal was present when the
331 sensory stimulus was not. These results also rule out straightforward forms of feedback creating
332 choice-related activity, as those explanations require the decision-related feedback to be aligned
333 with the sensory responses that gave rise to the corresponding choice. Furthermore, the delay
334 period choice signal was not entirely explainable as premotor. Given all these distinctions, the
335 oddly-parsimonious interpretation is that choice-related activity in MT is a distinct cognitive signal
336 (or set of signals) that are best understood outside of MT's encoding of visual motion. Although the
337 presence of large choice-related signals in macaque MT may be surprising at first, recent work in
338 other species (but also using ensemble recordings and analyses) has revealed widespread repre-
339 sentations of choice and other task-related signals (*Musall et al., 2019; Stringer et al., 2018; Grün-*
340 *demann et al., 2018*).

341 These findings provide new connections between MT function and well-established conceptual
342 and empirical frameworks. The sensory modulations associated with mismatches between ex-
343 pected and observed timecourses of motion aligns with both predictive coding and reinforcement
344 learning models, both of which are abstractly based on errors between expected and encountered
345 elements within a task (*Rescorla and Wagner, 1972; Engel et al., 2015*). Although our findings run
346 opposite known effects of temporal attention (*Ghose and Maunsell, 2002*) or attention-related gat-
347 ing of sensory responses (*Seidemann et al., 1998*), some recent work has decoupled attentional
348 modulations in MT and MST from task performance (*Recanzone and Wurtz, 2000*). Our dissociation
349 between MT modulations and task performance may be related, although in our case, their depen-
350 dence on the strategic history of the subjects revealed signals that are not wholly irrelevant to
351 the task, but are just not related to the formation of decisions on a trial-by-trial basis. This opens
352 up the possibility that some attention-like phenomena may arise from expectations of stimulus
353 statistics, instead of being modulations of sensory data per se. The post-stimulus choice signals
354 we observed in MT may be related to prior observations of small-amplitude, but tuned, persistent
355 activity in MT (*Bisley et al., 2004*); our findings suggest that those initial observations of relatively
356 small changes in spike rate may have simply caught a glimpse of larger non-sensory signals pre-
357 ceding the saccadic decisions mostly missed by single unit recordings that cannot see alternate
358 effects on population activity across diversely-tuned neurons. Finally, related work using a motion
359 categorization task has revealed strong non-sensory, category-related activity in area MST, but not
360 area MT (*Freedman and Assad, 2006; Zhou et al., 2020*). Such category-related activity can also be
361 thought of as "choice-correlated", as distinct from purely sensory-driven. Although the tasks, train-
362 ing histories, and analytic approaches differ between that work ours, our findings suggest that the
363 apparent distinction between MT and MST regarding the presence of such category/choice activity
364 might be less strict than previously observed. Again, the potential for ensemble recordings and
365 corresponding ensemble analyses may have been critical for not just observing these non-sensory
366 signals in MT, but for appreciating their substantial magnitude.

367 To conclude, our manipulation of temporal weighting strategy revealed a dissociation of sen-
368 sory response gain from decision formation. Likewise, our approach of using ensemble recordings
369 and population decoding allowed us to see large choice-related signals that were not just tempo-

rally dissociated from psychophysical weighting (or even stimulus viewing), but that were large in magnitude and distributed across the population in a manner distinct from how visual motion direction is represented. Together, these signals and modulations call for consideration of MT well beyond its role in encoding of retinal motion. Understanding the population coding structure and functional roles of such task-related but non-sensory computations are new open questions.

Methods and Materials

Stimulus presentation and design

Stimuli were presented using the Psychophysics Toolbox with Matlab (Math-works) using a Datapixx I/O box (Vpixx) for precise temporal registration (*Eastman and Huk, 2012*). Sample stimulus presentation code is available on request. Eye position was tracked using an Eyelink eye tracker (SR Research), sampled at 1 kHz. M1 was seated 57 cm away from a 150 cm × 86 cm rear-projection screen (IRUS; Draper Inc.) covering the central 106° × 73° of visual angle. Images were projected onto the screen by a PROPixx projector (VPixx Technologies Inc.) driven at a resolution of 1920 × 1080 pixels at 120 Hz. M2 viewed stimuli on a 55-inch LCD (LG) display (resolution = 1920 × 1080p, refresh rate = 60 Hz, background luminance = 26.49 cd/m²) that was corrected to have a linear gamma function. M2 viewed the stimulus from a distance of 118 cm (such that the screen width subtended 54° of visual angle, and each pixel subtended 0.0282° of visual angle). Auditory feedback was played at the end of every trial, and fluid reward was delivered through a computer-controlled solenoid.

Subjects were required to discriminate the net direction of a motion stimulus and communicate their decision with an eye movement to one of two targets, placed on either side of the motion stimulus. The sequence of task events is presented in Figure 1A. A trial began with the appearance of a fixation point. Once the subject acquired fixation and held for 750–1300 ms (uniform distribution), two targets appeared and remained visible until the end of the trial. 500–1000 ms after target onset, the motion stimulus was presented at a range of eccentricities from 4° to 12° for a duration of 1050 ms. The fixation point was extinguished 500–1000 ms after motion offset, and the subject was then required to shift their gaze toward one of the two targets within 600 ms (saccade endpoints within 3° of the target location were accepted). The timing of each event was randomly and independently jittered from trial to trial.

The reverse-correlation motion stimulus contained motion toward one direction or the opposite, with varying motion strength. Spatially, the stimulus consisted of a hexagonal grid of 19 Gabor elements, where individual Gabor elements were set to approximate the receptive field (RF) size of a V1 neuron, and the entire motion stimulus approximated the RF size of an MT neuron, which scaled based on eccentricity from fixation (*Van Essen et al., 1981*). Motion was presented by varying the phase of the sine-wave carrier of the Gabors. Each Gabor underwent a sinusoidal contrast modulation over time with independent random phase. Gabor spatial frequency (0.8 cycles/° sigma = 0.1 × eccentricity) and temporal frequency 5–6 Hz, yielding velocities of 5.55–6.66°/s, respectively) were selected to match the approximate sensitivity of MT neurons (*Bair and Movshon, 2004*).

Each motion stimulus presentation consisted of seven consecutive motion pulses lasting 150 ms each (9 frames on the 60 Hz display, 18 on the 120 Hz display), producing a motion sequence of 1050 ms in duration in total. On any given pulse, a number of Gabor elements would have their carrier sine waves drift in unison to produce motion (“signal elements”), and the remaining would counter-phase flicker (“noise elements”). Within any given pulse, signal elements were spatially assigned at random within the grid, and all signal element drifted in the same direction.

Motion strength on pulse i was defined as the proportion of signal elements out of the total number of elements, the value of which was drawn from a Gaussian distribution, $X_i \sim N(\mu_k, s)$ and rounded to the nearest integer, where k is the distribution index for the five trial types (strong left, weak left, zero-mean, weak right, strong right). Thus, while each pulse within a sequence could take on any value (and either sign/direction) from distribution $N(\mu_k, s)$, the expectation of a sequence

would be μ_k (Figure 1B-D). The subjects were rewarded for selecting the target consistent with the sign of the motion pulse sequence sum (i.e., the net direction), independent of the distribution μ_k from which the pulses were drawn.

Subjects performed the motion-discrimination task with three variations of temporal stimulus statistics (Levi et al., 2018). First was the flat-stimulus, in which expected motion strength was uniform over time within a trial. In other words, the mean of the motion strength distribution $N(\mu_k, s)$ would be held constant throughout a stimulus presentation. In other words, the mean of the distribution from which X_i was drawn was fixed at (μ_k) , for pulses 1–7 (Figure 1B).

Next, subjects encountered the late-stimulus, where motion strength was reduced substantially in early pulses, but not late. In this condition, μ_k is set to 0 for the first pulse ($i = 1$), and reaches its expected value (μ_k) by pulse 7. Finally, the opposite is done for the “early-stimulus” condition (Figure 1D), in which the first pulses maintain mean motion strength equal to μ_k and later pulses have a mean near zero. In the late- and early-stimulus conditions, the transition from μ_k at pulse 1 to μ_k at pulse 7 is governed by a logistic function with parameters chosen to result in a smooth transition between the first 3 and last 3 pulses (midpoint = 4, slope = 0.3).

All subjects began the experiments with the flat-stimulus condition (Monkey L: 13; Monkey N: 10 sessions). After multiple sessions of stable psychophysical performance, the stimulus was changed to the late-stimulus conditions (Monkey L: 11; Monkey N: 11 sessions). Finally, after multiple sessions of stable psychophysical performance the stimulus was changed to the early-stimulus condition (Monkey L: 11 sessions; Monkey N: 15 sessions). Subjects were exposed to only one stimulus condition per session and were not cued as to which stimulus condition they were viewing before or during any given session (other than the stimulus statistics themselves).

Throughout all conditions, there existed a subset of “zero-mean” trials in which $\mu_k = 0$ for all 7 pulses, regardless of whether the stimulus condition is flat, late, or early. Sessions also contained 5-10% frozen seed trials, which were identical stimulus displays. The “frozen noise” stimulus always summed to zero, had the same temporal structure across sessions, and was completely identical within sessions. Subjects were rewarded at random on frozen noise trials.

Behavioral analysis

Subject choices in the direction-discrimination task were analyzed with a maximum likelihood fit of a three-parameter logistic function (Wichmann and Hill, 2001) assuming a Bernoulli distribution of binary choices, in which the probability of a rightward choice is p and leftward choice is $1 - p$, where p is given by

$$p = \gamma + (1 - 2\gamma) \frac{1}{1 + e^{-\beta(x-a)}} \quad (1)$$

where x is the net motion strength value (z-scored over all sessions for each subject separately), a is the bias parameter (reflecting the midpoint of the function in units of motion strength), b is the slope (i.e., sensitivity, in units of log-odds per motion strength), and g captures the lapse rate as the offset from the 0 and 1 bounds. Error estimates on the parameters were obtained from the square root of the diagonal of the inverse Hessian (2nd derivative matrix) of the negative log-likelihood. The temporal weighting kernel (which we also refer to as “temporal weighting strategy” or “temporal weighting profile”) was computed using ridge regression via maximum likelihood. The log posterior of the psychophysical weights is given by

$$L(w) = \sum_{i=1}^N [Y_i w^T X_i - \log(1 + \exp(w^T X_i))] + \lambda \|w\|^2 \quad (2)$$

where $Y \in \{0, 1\}$ is a vector of choice on every trial and X is a matrix of the seven pulses on each trial, augmented by a column of ones (to capture bias). λ was estimated using evidence optimization (Sahani and Linden, 2003). Psychophysical weights are normalized by the Euclidean norm of the vector of weights. The seven temporal weights assigned to the seven motion pulses, w , were

463 computed by using all trials within a session. These include trials where μ_k was set to zero (i.e.
464 “zero-mean trials”, where motion on a given pulse is temporally independent of all other pulses
465 in the sequence) and trials where μ_k was set to a non-zero value (“signal trials”, where motion is
466 correlated over pulses)

467 Electrophysiology

468 A custom titanium chamber was fabricated and placed over the superior temporal sulcus and in-
469 traparietal sulcus to allow for a dorsal approach to access area MT. Chamber placement was as
470 guided by structural MRI and cranial landmarks. Extracellular recordings were performed using
471 linear electrode arrays from Plexon (U-Probe, V-Probe, or S-Probe; 24 or 32 channels; 50-100 mi-
472 crometer spacing).

473 MT was identified using electrode depths and paths (i.e., sulcal anatomy), and functional map-
474 ping. Functionally, MT was identified via size and location of RFs, and preponderance of direction
475 selective neurons. MT units were hand-mapped using a field of moving dots with experimenter
476 control of stimulus location, aperture size, dot speed, dot size, and dot density. Upon choosing
477 the stimulus location that maximally drove the highest number of neurons, direction tuning was
478 measured by 500ms presentations of a randomly drawn direction of motion from one of 12 di-
479 rections from 0 to 330 degrees. A total of 71 recording sessions were performed; 23 during the
480 flat-stimulus condition (Monkey L: 13; Monkey N: 10), 22 during the late-stimulus condition (Mon-
481 key L: 11; Monkey N: 11) and 26 during the early-stimulus condition (Monkey L: 11; Monkey N:
482 15).

483 Spike sorting was performed using KiloSort (*Pachitariu et al., 2016*) followed by manual merging
484 and splitting of clusters as necessary. A total of 583 units were identified; 161 during the flat-
485 stimulus condition, 219 during the late-stimulus, and 203 during the early stimulus.

486 Logistic regression neural decoder

487 To interrogate the roles and relationship of direction and decision-related signals, we used various
488 decoding methods to approximate how information may be gleaned from groups of MT neurons.
489 The first method we employed was logistic regression directly between spike counts and the binary
490 direction or choice on each trial (*Kiani et al., 2014; Yates et al., 2020*). The regression is done for
491 each session such that each neuron is a feature in the model, where each neuron received a weight
492 according to how well it predicts the binary outcome of interest. The result is a linear readout model
493 that allows for maximal prediction of the stimulus direction or the animal's choice.

494 Specifically, the decoding weights are calculated as coefficients in a logistic regression between
495 trial spike counts (summed over a window starting at stimulus onset and ending 150ms after stim-
496 ulus onset) and one of two binary variables (the stimulus direction, or the observer's choice) using
497 MATLAB's glmfit. The choice decoder weights were calculated using only the zero-sum, frozen
498 noise trials, while the direction decoder used all other trials.

499 The probability of a trial's stimulus direction or choice given each neuron's firing rate is given
500 by:

$$p(X|Y, \beta) = \frac{\exp(Yb)}{1 + \exp(Yb)} \quad (3)$$

501 Where $b = \beta_0 + \sum_{i=1}^N \beta_i X_i$ for N neurons present during a session. X is a vector of spike counts per
502 neuron, and the choice or direction is $Y \in \{0, 1\}$. The weights are then applied to their respective
503 neuron's temporally binned trial spike rates. Spikes were counted in 10ms bins and smoothed
504 with a 50ms boxcar. This was expressed in terms of rates by dividing by the bin size. The result is a
505 population response that best represented stimulus or choice information present in a recording
506 session.

507 The resulting decoder output was then used to calculate population-level choice probability (CP)
508 for each session. We measured CP over the course of stimulus presentation as a metric of trial-by-

trial correlation between firing rate and choice, given a fixed stimulus. CP was calculated as the area under the ROC curve generated from choice-conditioned distributions of the reweighted activity in each temporal bin. CP time course traces were smoothed with a 100ms boxcar for visualization.

Latent factor analysis

To understand how stimulus and perceptual choice are encoded across the population, we employed the variational latent Gaussian process (vLGP) method (Zhao and Park, 2017) to extract single-trial low-dimensional latent factors from population recordings in area MT. We used the recording between target onset and reward. The spike counts were binned at 10 ms. Let $\mathbf{x}_k$ denote the k -th dimension of the latent factors. We assumed that the spatial dimensions of latent factors are independent and imposed a Gaussian Process (GP) prior to the temporal correlation of each dimension,

$$\mathbf{x}_k \sim \mathcal{N}(0, \mathbf{K}). \quad (4)$$

To obtain smoothness, we used the squared exponential covariance function and respective covariance matrix $\mathbf{K}$ in the case of discrete time. Let y_{tn} denote the occurrence of a spike of the n th neuron at time t , $y_{tn} = 1$ if there was a spike at time t and $y_{tn} = 0$ otherwise at this time resolution. Then $\mathbf{y}_t$ is the vector of length N , total number of neurons in a session, that concatenates all neurons at time t . The spikes $\mathbf{y}_t$ are assumed to be a point-process generated by the latent state $\mathbf{x}_t$ at that time via a linear-nonlinear model,

$$\mathbf{y}_t \sim \text{Poisson}(\exp(\mathbf{A}\mathbf{x}_t + \mathbf{b})). \quad (5)$$

To infer the latent factors ($\mathbf{x}_t$ for each trial) and the model parameters ($\mathbf{A}$ and $\mathbf{b}$), we used variational inference technique, as the pair of prior and likelihood do not have an tractable posterior. We assumed parametric variational posterior distribution of the latent factors,

$$q(\mathbf{x}_k) = \mathcal{N}(\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k). \quad (6)$$

We analyzed the mean $\{\boldsymbol{\mu}_k\}$ as the latent factors in this study. The dimensionality of the latent factors was determined to be 4 by leave-one-neuron-out cross-validation on the session with the largest population. All the sessions with at least 4 simultaneously recorded units were included in this analysis (Monkey N: 13 sessions, Monkey L: 28 sessions).

Pulse-triggered average

To measure the relationship between the time-varying pulse strength and the inferred latent factors, we measured the contribution of pulses to the latent factors. The pulse-triggered average (PTA) measures the change in latent factors resulting from an additional pulse at a particular time of unit strength. To calculate the PTA, we used the pulse stimulus and latent response at 1 ms resolution. For each session, let s_i denote the value of the i -th motion stimulus, and let x_{ik} denote the k -th dimension of the latent factors at time t . All trials were concatenated such that the latent factors $\mathbf{X}$ is a matrix of length $T \times 4$, where T is the total time. For the i -th pulse, s_i is the number of Gabors pulsing, with $s_i > 0$ for pulses in one direction and $s_i < 0$ for pulses in the other direction. To calculate the temporal lags of the PTA, we built design matrices, $\mathbf{D} = [\mathbf{D}_1, \mathbf{D}_2, \dots, \mathbf{D}_7]$. For the i -th pulse, the design matrix $\mathbf{D}_i$ is a $T \times 28$ matrix that consists of 4 cosine basis functions at the $4i + 1, 4i + 2, \dots, 4i + 4$ -th columns and 0 elsewhere. These basis functions starts at 0 ms, 50 ms, 100 ms and 150 ms after the onset, lasts 100 ms each and spans the rows of $\mathbf{D}_i$. The magnitude of the bases is equal to the corresponding pulse value s_i . We calculated a separate $\mathbf{D}_i$ for each of the seven pulses and concatenated them to obtain a design matrix for all seven pulses and estimated

the weights with ℓ_2 -regularization,

$$\begin{aligned} \mathbf{X} &= \mathbf{D}\mathbf{W} + \mathbf{E} \\ \mathbf{W} &= \arg \min_{\mathbf{W}} \|\mathbf{X} - \mathbf{D}\mathbf{W}\|_2^2 + \gamma \|\mathbf{W}\|_2^2 \end{aligned} \quad (7)$$

where $\mathbf{W}$ is the weight matrix to estimate and $\mathbf{E}$ is the Gaussian noise matrix and the regularization hyperparameter γ was chosen by the generalized cross-validation (GCV) (Golub et al., 1979). The PTA was calculated with the design matrices of unit-strength pulse and the estimated weights $\mathbf{W}$. We smoothed the PTA with a temporal Gaussian kernel (40 ms kernel width).

Subject to arbitrary rotations, a latent trajectory forms an equivalence class of which the members have the same explanatory power in the vLGP model. We seek a particular rotation for each session that makes the encoded task signal concentrate in the first few dimensions. By singular value decomposition, $\mathbf{W}^T = \mathbf{U}\mathbf{S}\mathbf{V}^T$, we rotate the factors $\mathbf{x}$ to $\mathbf{U}^T\mathbf{x}$.

Choice decoder

Since there were some recording sessions with less than ideal number of frozen trials (identical visual motion trials) for the calculation of choice probability, we instead analyzed the “weak” trials of which the monkeys’ correct rate was below a threshold (65%). We started at the trials of zero pulse coherence and gradually increased the magnitude of coherence (absolute value) until the correct rate reached the threshold. One of the sessions containing less than 100 weak trials was excluded in this analysis.

We removed the stimulus information that is encoded in the latent factors of weak trials by regressing out the pulses and analyzed the residuals. The latent factors were re-binned at 100 ms resolution where the value of each bin is the sum of latent state $\mathbf{x}_t$ or spike counts $\mathbf{y}_t$ over the bin for $t = 1, 2, \dots, T$. For each t , we assumed a linear model

$$\mathbf{x}_t = \sum_{i=1}^7 \mathbf{w}_{ti} s_i + \mathbf{e}, \quad (8)$$

where s_i denote the strength of the i -th pulse, $\mathbf{w}_{ti}$ is the weight vector corresponding to the bin and pulse, and $\mathbf{e}$ is the homogeneous Gaussian noise across all bins. We estimated the weight vector by least-squares with ℓ_2 -regularization to prevent over-fitting,

$$\mathbf{w}_{ti} = \arg \min_{\mathbf{w}_{ti}} \|\mathbf{x}_t - \sum_{i=1}^7 \mathbf{w}_{ti} s_i\|_2^2 + \gamma \|\mathbf{w}_{ti}\|_2^2. \quad (9)$$

Again, the hyperparameter of regularization was chosen by GCV. We then analyzed the contribution of behavioral choice on the residuals

$$\mathbf{r}_t = \mathbf{x}_t - \sum_{i=1}^7 \mathbf{w}_{ti} s_i. \quad (10)$$

For the whole trial we used the sum residual of the windows $\mathbf{r} = \sum_t \mathbf{r}_t$. The range of t depends on the period of interest.

We trained logistic models, to which we refer to as *choice decoders*, to predict the choice on each trial using latent factors. The weights β and bias β_0 were estimated by maximum likelihood with ℓ_2 -regularization,

$$\beta, \beta_0 = \arg \max_{\beta, \beta_0} \log L(\text{choice} \mid \mathbf{r}; \beta, \beta_0) - \gamma \|\beta, \beta_0\|_2^2. \quad (11)$$

The hyperparameter of regularization was chosen via 5-fold (balanced classes in test set) cross-validation for every session individually.

Choice mapping

The conventional choice probability only applies to univariate variables. However, the latent factors and population activity are multivariate. We transformed the multivariate variables mentioned above onto a one-dimensional subspace that has the same direction as the choice through the choice decoders,

$$c = \frac{1}{1 + e^{-\beta^T \mathbf{r} - \beta_0}} \quad (12)$$

We refer to the transform as the *choice mapping*. The quantity c is a normalized value within [0, 1] that maps the residual onto the choice direction (Lueckmann et al., 2018), and enables pooling across sessions.

In order to prevent potential inflation of choice probability due to multidimensionality (3D), we regularized the choice decoder and used only the choice mapping on the test set (pooled samples held-out by cross-validation). This approach guarantees that choice probability will not be overestimated.

We pooled these mappings across all sessions. Using different subsets of latent factors as $\mathbf{r}$ in the mapping, we obtained the choice-mapping of the stimulus-dimension and non-stimulus-dimensions of latent factors. Then we calculated the choice probability of the corresponding dimensions based on the values. To investigate the time course of choice probabilities, we used choice decoders to perform choice-mapping on the whole dataset with non-overlapping moving windows. For fixed readout, we estimated the weights using mean value of 0-1.2s for the stimulus period and -0.5-0s for the delay period. We use the weights to obtain readout and CP values with 10ms moving window, and smooth the CP values with a 100ms boxcar. Finally; for dynamic readout, we estimated the weights and calculated the CP values within 100ms moving windows individually.

Acknowledgments

Thanks to A Laudano, K Mitchell, and C Carter for animal support and C Badillo for technical support. I.M.P and A.C.H were supported by NSF IIS-1734910. A.J.L was supported by NIH/NEI R01-EY017366.

References

- Bair W, Movshon JA. Adaptive Temporal Integration of Motion in Direction-Selective Neurons in Macaque Visual Cortex. *Journal of Neuroscience*. 2004; 24(33):7305–7323. doi: 10.1523/jneurosci.0554-04.2004.
- Benucci A, Saleem AB, Carandini M. Adaptation maintains population homeostasis in primary visual cortex. *Nature Neuroscience*. 2013; 16(6):724–729. doi: 10.1038/nn.3382.
- Bisley JW, Zaksas D, Droll JA, Pasternak T. Activity of Neurons in Cortical Area MT During a Memory for Motion Task. *Journal of Neurophysiology*. 2004; 91(1):286–300. doi: 10.1152/jn.00870.2003.
- Born RT, Bradley DC. Structure And Function Of Visual Area MT. *Annual Review of Neuroscience*. 2005; 28(1):157–189. doi: 10.1146/annurev.neuro.26.041002.131052.
- Britten KH, Newsome WT, Shadlen MN, Celebrini S, Movshon JA. A relationship between behavioral choice and the visual responses of neurons in macaque MT. *Visual Neuroscience*. 1996; 13(1):87–100. doi: 10.1017/s095252380000715x.
- Britten KH, Shadlen MN, Newsome WT, Movshon JA. The analysis of visual motion: a comparison of neuronal and psychophysical performance. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 1992 Dec; 12(12):4745–4765. <https://www.ncbi.nlm.nih.gov/pubmed/1464765>.
- Churchland MM, Yu BM, Cunningham JP, Sugrue LP, Cohen MR, Corrado GS, Newsome WT, Clark AM, Hosseini P, Scott BB, et al. Stimulus onset quenches neural variability: a widespread cortical phenomenon. *Nature Neuroscience*. 2010; 13(3):369–378. doi: 10.1038/nn.2501.
- Cook EP, Maunsell JH. Attentional Modulation of Motion Integration of Individual Neurons in the Middle Temporal Visual Area. *Journal of Neuroscience*. 2004; 24(36):7964–7977. doi: 10.1523/jneurosci.5102-03.2004.

- 625 **Cormack LK**, Czuba TB, Knöll J, Huk AC. Binocular Mechanisms of 3D Motion Processing. *Annual Review of*
626 *Vision Science*. 2017; 3(1):297–318. doi: 10.1146/annurev-vision-102016-061259.
- 627 **Eastman KM**, Huk AC. PLDAPS: A Hardware Architecture and Software Toolbox for Neurophysiology Requir-
628 ing Complex Visual Stimuli and Online Behavioral Control. *Frontiers in Neuroinformatics*. 2012; 6. doi:
629 [10.3389/fninf.2012.00001](https://doi.org/10.3389/fninf.2012.00001).
- 630 **Engel TA**, Chaisangmongkon W, Freedman DJ, Wang XJ. Choice-correlated activity fluctuations underlie learning
631 of neuronal category representation. *Nature Communications*. 2015; 6(1). doi: 10.1038/ncomms7454.
- 632 **Freedman DJ**, Assad JA. Experience-dependent representation of visual categories in parietal cortex. *Nature*.
633 2006; 443(7107):85–88. doi: 10.1038/nature05078.
- 634 **Ghose GM**, Maunsell JHR. Attentional modulation in visual cortex depends on task timing. *Nature*. 2002;
635 419(6907):616–620. doi: 10.1038/nature01057.
- 636 **Gold JI**, Shadlen MN. The Neural Basis of Decision Making. *Annual Review of Neuroscience*. 2007; 30(1):535–574.
637 doi: [10.1146/annurev.neuro.29.051605.113038](https://doi.org/10.1146/annurev.neuro.29.051605.113038).
- 638 **Golub GH**, Heath M, Wahba G. Generalized Cross-Validation as a Method for Choosing a Good Ridge
639 Parameter. *Technometrics: a journal of statistics for the physical, chemical, and engineering sciences*.
640 1979 May; 21(2):215–223. <https://amstat.tandfonline.com/doi/abs/10.1080/00401706.1979.10489751>, doi:
641 [10.1080/00401706.1979.10489751](https://doi.org/10.1080/00401706.1979.10489751).
- 642 **Gründemann J**, Bitterman Y, Lu T, Krabbe S, Grewe BF, Schnitzer MJ, Lüthi A. Amygdala neuronal ensembles
643 dynamically encode behavioral states. *Science*. 2018; doi: 10.1101/425736.
- 644 **Huk AC**, Shadlen MN. Neural Activity in Macaque Parietal Cortex Reflects Temporal Integration of Visual Mo-
645 tion Signals during Perceptual Decision Making. *Journal of Neuroscience*. 2005; 25(45):10420–10436. doi:
646 [10.1523/jneurosci.4684-04.2005](https://doi.org/10.1523/jneurosci.4684-04.2005).
- 647 **Kawaguchi K**, Clery S, Pourriahi P, Seillier L, Haefner RM, Nienborg H. Differentiating between Models of
648 Perceptual Decision Making Using Pupil Size Inferred Confidence. *The Journal of Neuroscience*. 2018;
649 38(41):8874–8888. doi: [10.1523/jneurosci.0735-18.2018](https://doi.org/10.1523/jneurosci.0735-18.2018).
- 650 **Kiani R**, Cueva CJ, Reppas JB, Newsome WT. Dynamics of Neural Population Responses in Prefrontal
651 Cortex Indicate Changes of Mind on Single Trials. *Current Biology*. 2014; 24(13):1542–1547. doi:
652 [10.1016/j.cub.2014.05.049](https://doi.org/10.1016/j.cub.2014.05.049).
- 653 **Kiani R**, Hanks TD, Shadlen MN. Bounded integration in parietal cortex underlies decisions even when viewing
654 duration is dictated by the environment. *The Journal of neuroscience: the official journal of the Society*
655 *for Neuroscience*. 2008 Mar; 28(12):3017–3029. <http://dx.doi.org/10.1523/JNEUROSCI.4761-07.2008>, doi:
656 [10.1523/JNEUROSCI.4761-07.2008](https://doi.org/10.1523/JNEUROSCI.4761-07.2008).
- 657 **Levi AJ**, Huk AC. Interpreting temporal dynamics during sensory decision-making. *Current opinion in physiol-*
658 *ogy*. 2020 Aug; 16:27–32. <http://dx.doi.org/10.1016/j.cophys.2020.04.006>, doi: [10.1016/j.cophys.2020.04.006](https://doi.org/10.1016/j.cophys.2020.04.006).
- 659 **Levi AJ**, Yates JL, Huk AC, Katz LN. Strategic and Dynamic Temporal Weighting for Perceptual Decisions in
660 Humans and Macaques. *eneuro*. 2018; 5(5). doi: [10.1523/eneuro.0169-18.2018](https://doi.org/10.1523/eneuro.0169-18.2018).
- 661 **Lueckmann JM**, Macke JH, Nienborg H. Can Serial Dependencies in Choices and Neural Activity Explain
662 Choice Probabilities? *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 2018
663 Apr; 38(14):3495–3506. <http://dx.doi.org/10.1523/JNEUROSCI.2225-17.2018>, doi: [10.1523/JNEUROSCI.2225-17.2018](https://doi.org/10.1523/JNEUROSCI.2225-17.2018).
- 665 **Mante V**, Sussillo D, Shenoy KV, Newsome WT. Context-dependent computation by recurrent dynamics in
666 prefrontal cortex. *Nature*. 2013; 503(7474):78–84. doi: 10.1038/nature12742.
- 667 **Musall S**, Kaufman MT, Juavinett AL, Gluf S, Churchland AK. Single-trial neural dynamics are dominated by
668 richly varied movements. *Nature neuroscience*. 2019; doi: [10.14224/1.38599](https://doi.org/10.14224/1.38599).
- 669 **Newsome WT**, Pare EB. A selective impairment of motion perception following lesions of the middle tem-
670 poral visual area (MT). *The Journal of Neuroscience*. 1988; 8(6):2201–2211. doi: [10.1523/jneurosci.08-06-02201.1988](https://doi.org/10.1523/jneurosci.08-06-02201.1988).
- 672 **Nienborg H**, Cumming BG. Decision-related activity in sensory neurons reflects more than a neuron's causal
673 effect. *Nature*. 2009; 459(7243):89–92. doi: 10.1038/nature07821.

- 674 **Okazawa G**, Sha L, Purcell BA, Kiani R. Psychophysical reverse correlation reflects both sensory and
675 decision-making processes. *Nature communications*. 2018 Aug; 9(1):3479. <http://dx.doi.org/10.1038/s41467-018-05797-y>, doi: 10.1038/s41467-018-05797-y.
- 677 **Osborne LC**, Bialek W, Lisberger SG. Time course of information about motion direction in visual area MT of
678 macaque monkeys. *The Journal of neuroscience: the official journal of the Society for Neuroscience*. 2004
679 Mar; 24(13):3210–3222. <http://dx.doi.org/10.1523/JNEUROSCI.5305-03.2004>, doi: 10.1523/JNEUROSCI.5305-
680 03.2004.
- 681 **Pachitariu M**, Steinmetz N, Kadir S, Carandini M, D HK. Kilosort: realtime spike-sorting for extracellular elec-
682 trophysiology with hundreds of channels. *bioRxiv*. 2016; doi: 10.1101/061481.
- 683 **Recanzone GH**, Wurtz RH. Effects of Attention on MT and MST Neuronal Activity During Pursuit Initiation.
684 *Journal of Neurophysiology*. 2000; 83(2):777–790. doi: 10.1152/jn.2000.83.2.777.
- 685 **Rescorla RA**, Wagner AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement
686 and nonreinforcement. *Classical Conditioning II*. 1972; .
- 687 **Roitman JD**, Shadlen MN. Response of Neurons in the Lateral Intraparietal Area during a Combined Vi-
688 sual Discrimination Reaction Time Task. *The Journal of Neuroscience*. 2002; 22(21):9475–9489. doi:
689 10.1523/jneurosci.22-21-09475.2002.
- 690 **Sahani M**, Linden JF. Evidence optimization techniques for estimating stimulus-response functions. In: *Ad-
691 vances in Neural Information Processing Systems* MIT Press; 2003.p. 301–308. [http://citeseer.ist.psu.edu/703007.](http://citeseer.ist.psu.edu/703007.html)
692 [html](http://citeseer.ist.psu.edu/703007.html).
- 693 **Salzman CD**, Murasugi CM, Britten KH, Newsome WT. Microstimulation in visual area MT: effects on direction
694 discrimination performance. *The Journal of Neuroscience*. 1992; 12(6):2331–2355. doi: 10.1523/jneurosci.12-
695 06-02331.1992.
- 696 **Seidemann E**, Newsome WT. Effect of Spatial Attention on the Responses of Area MT Neurons. *Journal of*
697 *Neurophysiology*. 1999; 81(4):1783–1794. doi: 10.1152/jn.1999.81.4.1783.
- 698 **Seidemann E**, Zohary E, Newsome WT. Temporal gating of neural signals during performance of a visual
699 discrimination task. *Nature*. 1998; 394(6688):72–75. doi: 10.1038/27906.
- 700 **Stringer C**, Pachitariu M, Steinmetz N, Reddy CB, Carandini M, Harris KD. Spontaneous behaviors drive multi-
701 dimensional, brain-wide activity. *Science*. 2018; doi: 10.1101/306019.
- 702 **Treue S**, Maunsell JH. Attentional modulation of visual motion processing in cortical areas MT and MST. *Nature*.
703 1996 Aug; 382(6591):539–541. <http://dx.doi.org/10.1038/382539a0>, doi: 10.1038/382539a0.
- 704 **Van Essen DC**, Maunsell JHR, Bixby JL. The middle temporal visual area in the macaque: Myeloarchitecture,
705 connections, functional properties and topographic organization. *The Journal of Comparative Neurology*.
706 1981; 199(3):293–326. doi: 10.1002/cne.901990302.
- 707 **Wichmann FA**, Hill NJ. The psychometric function: I. Fitting, sampling, and goodness of fit. *Perception &
708 Psychophysics*. 2001; 63(8):1293–1313. doi: 10.3758/bf03194544.
- 709 **Yates JL**, Katz LN, Levi AJ, Pillow JW, Huk AC. A simple linear readout of MT supports motion direction-
710 discrimination performance. *Journal of Neurophysiology*. 2020; 123(2):682–694. doi: 10.1152/jn.00117.2019.
- 711 **Yates JL**, Park IM, Katz LN, Pillow JW, Huk AC. Functional dissection of signal and noise in MT and LIP during
712 decision-making. *Nature Neuroscience*. 2017 Jul; 20:1285–1292. doi: 10.1038/nn.4611.
- 713 **Zhao Y**, Park IM. Variational Latent Gaussian Process for Recovering Single-Trial Dynamics from Population
714 Spike Trains. *Neural Computation*. 2017 May; 29(5). doi: 10.1162/NECO_a_00953.
- 715 **Zhao Y**, Yates JL, Levi A, Huk A, Park IM. Stimulus-choice (mis)alignment in primate area MT. *PLOS Computa-
716 tional Biology*. 2020 May; <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1007614>, doi:
717 10.1371/journal.pcbi.1007614.
- 718 **Zhou Y**, Mohan K, Freedman DJ. Distributed Visual Category Processing Across Medial Superior Temporal and
719 Lateral Intraparietal Cortices. *bioRxiv*. 2020; doi: 10.1101/2020.08.25.266791.