Transformation of value signaling in a striatopallidal circuit

Donghyung Lee[1], Lillian Liu[1], and Cory M. Root[1]

[1] University of California San Diego, Department of Neurobiology, School of Biological Sciences, San Diego, California

1 Abstract

The ways in which sensory stimuli acquire motivational valence through association with other stimuli is one of the simplest forms of learning. Though we have identified many brain nuclei that play various roles in reward processing, a significant gap remains in understanding how value encoding transforms through the layers of sensory processing. To address this gap, we carried out a comparative investigation of the olfactory tubercle (OT), and the ventral pallidum (VP) - 2 connected nuclei of the basal ganglia which have both been implicated in reward processing. First, using anterograde and retrograde tracing, we show that both D1 and D2 neurons of the OT project primarily to the VP and minimally elsewhere. Using 2-photon calcium imaging, we then investigated how the identity of the odor and reward contingency of the odor are differently encoded by neurons in either structure during a classical conditioning paradigm. We find that VP neurons robustly encode value, but not identity, in low-dimensional space. In contrast, OT neurons primarily encode odor identity in high-dimensional space. Though D1 OT neurons showed larger response vectors to rewarded odors than other odors, we propose this is better interpreted as identity encoding with enhanced contrast rather than as value encoding. Finally, using a novel conditioning paradigm that decouples reward contingency and licking vigor, we show that both features are encoded by non-overlapping VP neurons. These results provide a novel framework for the striatopallidal circuit in which a high-dimensional encoding of stimulus identity is collapsed onto a low-dimensional encoding of motivational valence.

Introduction

Animals exhibit an impressive ability to change how sensory inputs map onto behavioral outputs. Understanding how animals learn to output different behaviors through experience is one of the fundamental problems in neuroscience. Over the last half century, the field has developed compelling frameworks to tackle this problem at both the algorithmic level (Rescorla-Wagner models, Q-learning models) (Rescorla, 1972; Sutton, 1988) and the mechanistic level (Hebbian learning, STDP, neuromodulation) (Dan and Poo, 2004). By comparison, we lack frameworks through which to understand how the brain might implement learning algorithms through the updating of synaptic weights. One strategy has been to identify neural correlates of latent features assumed to be required for these algorithms (e.g. dopamine as a neural substrate for reward-prediction-error) (Hollerman and Schultz, 1998; Schultz et al., 1997). These results, however, can often be difficult to interpret because reward related signals are found globally throughout the brain (Allen et al., 2019), and are likely multiplexed with signals about motor output and/or stimulus identity. We propose that a more powerful approach is one that compares 1) how the encoding of reward cues changes from one brain nucleus to its downstream target and 2) how much of the encoding can be explained by value vs. other features such as identity or motor output. In this present work, we implement this comparative approach to the investigation of how encoding of olfactory reward cues is transformed between the olfactory tubercle (OT) and the ventral pallidum (VP) in the context of classical conditioning.

The OT, also known as the tubular striatum (Wesson, 2020), is a 3-layered striatal nucleus situated at the bottom of the forebrain. As with other striatal structures, the OT is composed primarily of Spiny Projection Neurons (SPN's) which express either the *Drd1* or *Drd2* DA receptors (abbreviated as OT_{D1} and OT_{D2}, respectively) (Tritsch and Sabatini, 2012). In addition to receiving a wide range of inputs from cortical and amygdalar areas (e.g. AI, OFC, BLA, PlCoA, Pir) (Zhang et al., 2017b), it receives dense DAergic input from the midbrain (Ikemoto, 2007) and direct input from the mitral and tufted cells of the olfactory bulb (Haberly and Price, 1977; Igarashi et al., 2012; Scott, 1981), a unimodal and primary sensory area. There is a range of experiments that suggest that the OT's DAergic innervation is involved in reward processing. Coincident stimulation of the lateral olfactory tract and DAergic midbrain afferents supports LTP of excitatory current (Wieland et al., 2015) and rats self-administer cocaine, a DAergic drug, into the medial OT more vigorously than to any other striatal nuclei (Ikemoto,

2003). And while the OT neurons are known to respond to a wide range of odorants (Wesson

and Wilson, 2010), pairing stimulation of midbrain DAergic neurons with an odor drives

appetitive behavior towards the paired odor (Zhang et al., 2017a) and enhances the contrast of

the odor-evoked activity (Oettl et al., 2020). Lastly, a number of recent publications report

varying degrees of value signals recorded from neurons in the OT (Gadziola et al., 2020, 2015;

Martiros et al., 2022; Millman and Murthy, 2020; Oettl et al., 2020).

The most well-established target of the OT is the VP (Newman and Winans, 1980; Zahm

and Heimer, 1987), a pallidal structure that lies immediately dorsal to the OT. In addition to OT

input, VP receives strong input from the nucleus accumbens (Jones and Mogenson, 1980) and

the subthalamic nucleus (Ricardo, 1980; Turner et al., 2001). More recently, it was reported that

VP also receives inputs from several cortical and amygdalar areas that also project to the OT

(e.g. Pir, BLA, OFC) (Stephenson-Jones et al., 2020). The VP contains GABAergic neurons,

which respond to positive value cues, and glutamatergic neurons, which respond to negative

value cues. Consistent with their responsiveness, the GABAergic and glutamatergic neurons

drive real time place preference and avoidance, respectively (Faget et al., 2018). Though it is

well-established that the VP plays a critical role in reward processing, there has been ongoing

disagreement on what specific latent features are encoded by VP neurons. Interpretations have

included value (Ottenheimer et al., 2018, 2020b; Richard et al., 2016; Tachibana and Hikosaka,

2012), hedonics (Smith et al., 2009; Tindell et al., 2006), motivation (Faget et al., 2018; Fujimoto

et al., 2019; Lederman et al., 2021; Tindell et al., 2005), and reward-prediction error

(Ottenheimer et al., 2020a). This ongoing discussion highlights the need to adopt a more

comparative approach outlined above.

Here, we investigated the transformation of value encoding between the OT and the VP.

We began by refining our understanding of OT's efferents to reveal that, contrary to a previous

report, both OT_{D1} and OT_{D2} neurons project primarily to the ventrolateral portion of the VP and

minimally elsewhere. Given this finding that VP may be the only robust output of the OT, we

proposed that the OT to VP circuit is an ideal model system for examining how the encoding of

reward cues is transformed between connected brain areas. Comparing the stimulus-evoked

activity in OT_{D1}, OT_{D2}, and VP neurons with 2-photon Ca^{2+} imaging, we found that VP neurons

encode reward-contingency in low-dimensional space with good generalizability. In contrast,

activity in both OT_{D1} and OT_{D2} neurons was high-dimensional and primarily contained

information about odor identity. By examining the same neurons across multiple days of pairing,

we propose a putative cellular mechanism for reward-cue responsiveness in VP wherein reward

responsive VP neurons gradually become reward-cue responsive. Finally, using a novel

classical conditioning paradigm, we provide evidence that non-overlapping sets of VP neurons

contain information about the vigor of licking and reward-contingency, but not both.

Results

In order to compare odor-evoked activity in connected brain nuclei, we first characterized
which specific subregions of the VP receive input from the anteromedial portion of the OT. While
considerable effort has been made to unravel the anatomy and function of the NAc, much less
attention has been directed at the OT. Though multiple studies have characterized its
anatomical connectivity (Zahm and Heimer, 1987; Zhang et al., 2017b; Zhou et al., 2003), there
is inconsistency regarding whether or not OT projects to areas other than the VP. We therefore
aimed to clarify previously reported OT connectivity by independently conducting anterograde
viral tracing experiments in OT_{D1} and OT_{D2} neurons. To this end, we injected AAVDJ-hSyn-
FLEX-mRuby-T2A-syn-eGFP in the anterior OT of *Drd1*-Cre (labels D1+ SPN's) and *Adora2a*-
Cre (labels D2+ SPN's) animals (Fig1A, FigS1C-E). Because viral contamination of areas dorsal
to the target site can lead to difficulties in interpretation of tracing data, we also injected the
same virus to the AcbSh immediately dorsal to the OT for comparison. Consistent with past
findings (Kupchik et al., 2015), we observed robust projections VP, LH, and VTA from $AcbSh_{D1}$
neurons and primarily VP projections $AcbSh_{D2}$ neurons (Fig1B-C). We also observed dense
labeling of the VP in D1-Cre and A2A-Cre animals injected at the OT. Contrary to one report
(Zhang et al., 2017b) but consistent with another (Zhou et al., 2003), we observed minimal
labeling in LH and VTA, or anywhere else in the brain, for both OT_{D1} and OT_{D2} experiments
(Fig1B-C, FigS1A-B), suggesting that neither OT subpopulation projects strongly outside the
VP. As previously reported (Groenewegen and Russchen, 1984). It is also notable that OT
projections were restricted to the lateral portions of the VP.
To corroborate and more precisely describe the OT to VP projection, we conducted
retrograde tracing by injecting CTB::488 and CTB::543 to the lateral and medial portion of
caudal VP, respectively (Fig1D, FigS1F-G). We found strong labeling of soma by both CTB::488
and CTB::543 in the Acb, AI, and Pir (Fig1E-F). By comparison, we found predominantly
CTB:488, but not CTB::543, labeling in OT soma, indicating OT neurons are more likely to
project to the lateral portion of the VP than to the medial. Similarly, to corroborate the lack of OT
to VTA projection, we injected CTB::647 into the VTA (Fig1G, FigS1H-I). Consistent with
previous findings (Beier et al., 2015; Faget et al., 2016; Watabe-Uchida et al., 2012), we found
dense labeling of soma in various areas of the striatum such as AcbSh, AcbC, and CPu (Fig1H-
I). We also found some labeling of soma in some frontal cortical regions such as PrL, AI and IL

cortices. In stark contrast, we found that hardly any OT neurons were labeled. The rare OT

neurons that did have CTB labeling were exclusively localized to the dorsalmost portion of layer

III, closely bordering the VP. Taken together, we conclude that both D1 and D2 SPN's of the OT

project primarily to the lateral portion of the VP and negligibly to other brain areas, including the

VTA.

Once we had identified that OT has extremely constrained outputs to the lateral VP, we

set out to comparatively characterize the encoding of reward cue in this striatopallidal circuit.

Past analysis of value encoding is confounded by not accounting for the difficult-to-avoid

overlaps among identity, salience, and reward contingency. To address this, we carefully

designed a 6-odor conditioning paradigm where these factors could be decoupled (Fig2A).

During each trial, the animal is exposed to 1 of 6 odors for 2 seconds. At the end of odor

delivery, the animal either receives: 2 μl of a 10% sucrose solution (S), 50 ms of airpuff at 70 psi

(P) or nothing (X). 3 of the odors are ketones (hexanone, heptanone, octanone) and the rest are

terpenes (terpinene, pinene, limonene), but the pairing contingencies are chosen such that each

contingency group (S, P, or X) includes 1 ketone and 1 terpene. In a value-encoding population,

but not in an identity-encoding population, we should see that odor pairs of different reward-

contingency (e.g. S_K, a sucrose-paired ketone vs. S_P, an airpuff-paired ketone) are more

different than odor pairs of same reward-contingency (e.g. S_K, a sucrose-paired ketone vs. S_T, a

sucrose-paired terpene). Additionally, because both sucrose-pairing and airpuff-pairing should

make the associated odor more salient, we can disambiguate between increased

discriminability due to salience vs. valence by comparing neural activity in response to sucrose-

cues or airpuff-cues.

To record the activity of the OT and VP neurons across multiple days of pairing, we

injected C57BL/6 mice with AAV9-hSyn-jGCaMP7s-WPRE (lateral VP) and *Drd1*-Cre or

Adora2a-Cre animals with AAV9-hSyn-FLEX-jGCaMP7s-WPRE (anteromedial OT) (Fig2B,

FigS2A-F). Additionally, we implanted a 600μm Gradient Refractive Index (GRIN) lens 150μm

dorsal to the virus injection site and cemented a head-fixation plate to the skull. 6-8 weeks after

surgery, animals were water-restricted and habituated for 3-5 days in the head-fixation setup

(Fig2C). We processed the acquired time-series images using Constrained Nonnegative Matrix

Factorization (Pnevmatikakis et al., 2016) to obtain fluorescence traces from each putative

neuron (Fig2D,E). In total, we recorded Ca^{2+} signals from 231 OT_{D2} neurons from 6 *Adora2a*-

Cre animals ([FigS3](#)), 288 OT$_{D1}$ neurons from 6 *Drd1*-Cre animals ([FigS4](#)) and 130 VP neurons

from 5 C57BL6/J animals ([FigS5](#)).

After 3 days of odor-sucrose associations, the animals displayed anticipatory licking

behavior primarily during sucrose-paired odors ([Fig2F,G](#)). Starting on day 4, the sucrose and

airpuff contingencies were switched such that every odor had a reassigned contingency. By day

6, animals had adapted their anticipatory licking behavior to match the new sucrose-contingency

([Fig2G](#)). Quantification of the animal's licking behavior showed that the accuracy of animals'

licks during odor increased across time and was not different across lens-placement groups

([Fig2H](#); $F_{day}=27.64$, $p_{day}=2.29e-16$, $F_{lens\ location}=2.30$, $p_{lens\ location}=0.11$). These results show that the

animals learn to associate S odors with reward in a flexible manner in our paradigm. Because

we saw the strongest behavioral evidence that animals learned odor-sucrose associations by

day 6, we focused our analysis on how reward cues are encoded on the last day of imaging.

The animals also showed trends of behavioral changes in response to airpuff-cues, though they

were not significant: during airpuff-cues, animals walked less and closed their eyes more than

during other odors ([FigS6D-G](#)). These behavioral changes for aversive cues were less robust

than that for reward association, however, animals show clear responses to the US indicating

that they perceive the aversive stimulus.

OT and VP neurons showed heterogeneous responses to 6 odors across all 6 days of

imaging ([FigS3](#), [FigS4](#), [FigS5](#)). To unbiasedly describe the difference between regions, we

performed hierarchical clustering on the pooled trial-averaged responses to the 6 odors on the

6th day of imaging ([Fig3A](#)). We observed both inhibitory (clusters I, II) and excitatory (clusters

III-VI) responses to odors as well as broad (clusters II, VI) and narrow (clusters IV, V) odor-

tuning ([Fig3B](#)). Cluster I and cluster III most closely fit our description of putative valence-

encoding neurons, i.e. neurons that had similar responses to 2 sucrose-cues (S_K vs. S_T) but

different responses to a sucrose-cue and a puff-cue or control odor (S_K vs. P_K or X_K). Although

all clusters included neurons from all subpopulations, cluster I and cluster III, which showed

larger responses to odors predicting sucrose, were enriched for VP neurons ([Fig3C](#)), leading us

to hypothesize that neurons in the VP were more likely to be valence encoding neurons than in

either OT subpopulation.

To assess this hypothesis, we quantified the number of neurons that had statistically

significant responses to each of the 6 odors on the last day of imaging. We found that more VP

neurons were either excited (29.8±4.1%, 36.6±4.0% for S_K, S_T) or inhibited (24.5±3.0%,

29.4±3.8% for S_K, S_T) to either sucrose-paired odor than to control or puff-paired odors (7.6-

11.1% excited, 8.1-12.9% inhibited) (Fig3D,E). A 2-way ANOVA revealed that sucrose-

contingency, but not puff-contingency, affects the percentage of responsiveness among VP

neurons on day 3 and day 6 (FigS8; F=18.6, p=2.04x10^{-4} and F=32.2, p=5.78x10^{-6},

respectively), but not on day 1. By comparison, there was no significant effect of sucrose pairing

on the number of responsive neurons among OT_{D2} or OT_{D1} animals on day 6 (p>0.05). OT_{D2}

neurons were slightly more responsive to puff-cues on day 3 (FigS8; F=3.03, p=3.99x10^{-2}) and

OT_{D1} neurons were more responsive to sucrose-cues on day 3 (FigS8; F=7.33, p=1.08x10^{-2}),

but these effects were small when compared to the effect of sucrose pairing observed among

VP animals. When compared across days, we found that the percentage of VP neurons that

respond to both S odors increases from 6.1±2.2% on day 1 to 34.1±5.1% by the 6th day of

imaging (Fig3E). By comparison, the percentage of OT neurons that respond to both S odors in

the same direction (i.e. excited by both S odors or inhibited by both S odors) did not increase

through training. Furthermore, whereas OT_{D1} and OT_{D2} neurons were more likely to respond to

a single odor than they were to respond to both S odors (12.6 vs 31.3% in OT_{D1}, 11.8 vs 21.7%

in OT_{D2}), VP neurons were more likely to respond to both S odors than to a single odor (34.1 vs

23.3%).

Similarly, we found that the magnitude of trial-averaged odor responses in the VP were

significantly higher for S odors than X or P odors on the last day of imaging (FigS9; F=59.5,

p=3.83x10^{-14}). By comparison, neither sucrose-pairing nor airpuff-pairing had any impact on the

magnitude of odor responses in OT_{D2} neurons on day 6. And though we did observe a

significant effect of sucrose-pairing on response magnitudes in OT_{D1} neurons (FigS9; F=6.34,

p=7.97x10^{-3}), both the effect size and significance were weaker than observed in VP. Given that

an ideal value-encoding neuron should respond similarly to 2 odors of equal reward-contingency

but disparate molecular structure, we looked at the correlation between each neuron's response

to the sucrose-paired ketone (S_K) and to the sucrose-paired terpene (S_T). Consistent with our

hypothesis, we saw a high correlation between a VP neuron's responses to S_K and S_T (Fig3F;

R^2=0.89). This similarity was much higher than between the sucrose-paired ketone (S_K) and the

control ketone (X_K) despite the greater structural similarity between S_K and X_K (Fig3G; R^2=0.33).

In contrast, for both OT_{D1} and OT_{D2} neurons, there was a higher correlation between responses

to similar molecular structure (S_k and $X_{K,}$) than between responses to similar contingency (S_K

and S_T) (OT_{D2}: S_K vs. S_T R^2=0.04, S_K vs. X_K R^2=0.58; OT_{D1}: S_K vs. S_T R^2=0.13, S_K vs. X_K

R^2=0.40). Moreover, most VP neurons (76.5%), had a smaller absolute difference in the

response magnitude to the 2 S odors ($|S_K-S_T|$) than the absolute difference between the

sucrose-paired ketone and the control ketone ($|S_K-X_K|$) (Fig3H). By comparison, only half of

OT_{D2} and OT_{D1} neurons showed smaller $|S_K-S_T|$ than $|S_K-X_K|$, as would be expected if response

magnitude to an odor did not depend on reward-contingency. This trend was not due to the fact

that VP neurons were more likely to respond to both S odors than OT neurons were since it was

consistent across various thresholds for odor response magnitude. This trend was consistent for

other pairwise odor comparisons where one odor was a sucrose-cue and the other was not (e.g.

S_K vs. P_T, FigS10A-B).

Finally, we reasoned that the activity of reward-contingency encoding neurons would

support good decoding of odor pairs which have different valence but not of odor pairs that have

the same valence. To do this, we trained binary logistic classifiers from each neuron's response

to all 15 odor pairs and quantified the area under their receiver operating characteristic

(auROC). Because auROC values were non-normal with large spread, we quantified what

percentage of neurons had an auROC of at least 0.75, halfway between ideal and at-chance

decoding. We also note that all classifiers with auROC>0.75 showed bootstrapped p-values less

than 10^{-3} (FigS10C-D). To assess whether neurons from each region were encoding valence,

we compared a neuron's {S_K vs. X_K} decoder performance (*inter*valence classification) against

its {S_K vs. S_T} decoder performance (*intra*valence classification) (Fig3I-J, FigS10E-G). Across

multiple days of imaging, we found that the percentage of neurons that support intervalence

classification increased regardless of region but that this effect was markedly more pronounced

among VP neurons than among OT_{D1} or OT_{D2} neurons (Fig3I, FigS10F). Intravalence

classification, however, did not depend on days or region (Fig3J). By day 6, there were thrice as

many VP neurons with good intervalence decoding than with intravalence decoding (51.8±5.0%

vs. 14.4±5.8% for {S_K vs X_K} and {S_K vs S_T}, respectively). In contrast, a similar number of OT

neurons displayed good intervalence decoding as did intravalence decoding (20.8% vs 19.9% of

OT_{D1}; 12.8% and 21.0% of OT_{D2} for {S_K vs X_K} and {S_K vs S_T}, respectively). The pattern of

better intervalence decoding than intravalence decoding among VP neurons was observed

across all 15 pairwise classifiers (FigS10G). Whereas 10.2% of all day 6 VP neurons had

auROC>0.75 for {S_K vs. S_T}, 46.9-57.8% had auROC>0.75 for any classification between a

sucrose-cue and a control odor or airpuff-cue. By comparison, there were few neurons with

auROC>0.75 for any classification between a puff-cue and a control odor (2.3-10.9%),

suggesting that negative valence is not encoded in these VP neurons.

Plotting a neuron's $\{S_K$ vs. $S_T\}$ auROC against its $\{S_K$ vs. $X_K\}$ auROC, we can categorize

a neuron into the 4 categories ([Fig3K-L](#)). 1) a valence encoding neuron ($\{S_K$ vs. $S_T\}<0.75$ and

$\{S_K$ vs. $X_K\}>0.75$), 2) an identity encoding neuron (both auROC>0.75), 3) an identity encoding

neuron that does better with S odors ($\{S_K$ vs. $S_T\}>0.75$ and $\{S_K$ vs. $X_K\}<0.75$), and 4) an

uninformative neuron (both auROC<0.75). According to this categorization, half of VP neurons

were valence encoding by day 6, followed by OT_{D1} then OT_{D2} ([Fig3M](#); 47.7, 16.2, 7.3% for VP,

OT_{D1}, and OT_{D2}, respectively). The opposite was true for identity encoding. VP had a smaller

percentage of identity encoding neurons than either OT_{D1} or OT_{D2} (14.8, 21.1, 22.9% for VP,

OT_{D1}, and OT_{D2}, respectively). We note that these conclusions can also be replicated when

analyzing multinomial regression (MNR) classifiers trained on single neuron activities ([FigS11](#)).

Namely, the rates of confusion between the 2 sucrose cues are highest in VP and lowest in

OT_{D2} whereas the rates of confusion across all ketones (S_K, X_K, P_K) are highest in OT_{D2} and

lowest in VP. These single-neuron classifier analyses further consolidate our hypothesis that VP

neurons, more than either OT_{D2} or OT_{D1} neurons, were encoding reward contingency. However,

the most striking observation was that while only a subset (37.5%) of VP neurons had

auROC<0.75 for both $\{S_K$ vs. $X_K\}$ and $\{S_K$ vs. $X_K\}$, a majority of OT_{D2} and OT_{D1} neurons (69.7%

and 62.6%, respectively) showed auROC<0.75 for both $\{S_K$ vs. $S_T\}$ or $\{S_K$ vs. $P_K\}$. This led us to

conclude that, in comparison to the VP, most OT neurons have little discriminatory information

about olfactory stimuli regardless of valence at the single-neuron level and are rather used in a

population code.

Our data indicated that valence encoding emerges in VP neurons over the course of

learning. To explore the potential mechanisms at the cellular level, we compared the activity of a

subset of neurons we could observe on both day 1 and day 3 ([Fig4A-F](#)). We noticed there were

neurons that responded to the sucrose delivery on day 1 that responded to the sucrose cue on

day 3 ([Fig4C,F](#)), reminiscent of models of Hebbian plasticity. When quantified, we found that

17.9, 20.9% of VP neurons were responsive to sucrose on day 1 and S_K and S_T on day 3,

respectively ([Fig4G](#)). We specifically considered neurons that had the same direction of

response (excitation or inhibition) to both cues on separate days. This figure was much lower

among OT subpopulations (11.5, 8.2% for S_K and S_T in OT_{D1}; 10, 2.5% for S_K and S_T in OT_{D2}).

Consistent with above observations, we also found that the odor responses to sucrose-cues

were larger on day 3 than day 1 in 85% of tracked VP neurons, but only in 65% and 57% of

OT_{D1} and OT_{D2} neurons, respectively (Fig4H,I). We did not see the same effect in VP neurons'

responses to control or puff-paired odors. Together, our data suggest that sucrose pairing

causes sucrose-responsive VP neurons to increase their responses to the sucrose-predictive

odors.

Olfactory brain areas are known to use population codes to encode sensory information.

According to this model, single neurons are broadly tuned and their activities have weak

discriminatory information. But the activity of the population allows for an efficient encoding of

high-dimensional data. To assess if there is discriminatory information about the odorants within

the population-level activity, we compared the pairwise euclidean distance of trial-averaged odor

responses for all 15 odor pairs (Fig5A,B). We saw that, in general, the pairwise euclidean

distance for all odor pairs examined increases quickly after the onset of odor, reaches peak

distance towards the end of the 2 second odor delivery, and slowly decays after odor ends

(Fig5A). When examining the average pairwise distance during the last second of odor, there

was a relatively unstructured distribution of pairwise distance in OT_{D2} odor-response such that

$||S_K-X_K||$, $||S_K-P_K||$, $||S_K-S_T||$, and $||X_K-X_T||$ were all similar (Fig5B). By comparison, in VP

populations, the distribution was structured such that intervalence pairwise comparisons

between sucrose-paired and not sucrose-paired odors (e.g. $||S_K-P_K||$ and $||S_K-X_K||$) were larger

than intravalence pairwise comparisons (e.g. $||S_K-S_T||$, or $||X_K-X_T||$). OT_{D1} populations showed an

intermediate trend where most intravalence pairwise distances were smaller than intervalence

pairwise distances with the exception of $||S_K-S_T||$. Thus, at the population level VP

representations appear to encode valence but not identity, whereas OT representations is best

suited for identity encoding.

In parallel, we also performed decoding analysis using linear classifiers to assess how

reliably a given pair of odors could be decoded from population-level activity (Fig5C-D). To

quantify this, we extracted the average $\Delta F_{i,k}/F$ values for each trial $i \in [1,m]$ and each neuron $k \in$

$[1,n]$. The resulting matrix of size m x n was used to train a binary linear classifier with a logistic

learner. For each classifier, we looked at the average accuracy across 5-fold cross-validation

(CV accuracy). Classifiers were trained on neurons that were recorded together (i.e. neurons

from the same animal recorded on the same day) to capture biological variability. A total of 765

pairwise linear classifiers were trained (15 pairwise comparisons, 17 animals and 3 days). When

compared against 10,000 shuffles, 569 of these classifiers showed bootstrapped p-value less

than 0.001 ([FigS12A](#)). Importantly, all classifiers with CV accuracy higher than 0.75 had p-value
less than 0.001.

Linear classifiers trained on day 6 OT_{D2} population data had similar ranges of accuracy
regardless of valence ([Fig5D](#)). For example, the intravalence classification {S_K vs. S_T} was more
accurate (86.6±3.9**%**) than some and intervalence classifications (e.g. {S_K vs. X_K}, 72.8±5.4**%**)
but less accurate than others (e.g. {S_T vs. P_K}, 88.2±3.1**%**). Classifiers trained on VP population
activity, however, always showed more accurate intervalence decoding (range: 89.5-96.1%)
than intravalence decoding ({S_K vs. S_T}, 79.9±6.1**%**). Additionally, whereas OT_{D2} population
classifiers could decode the 2 control cues {X_K vs. X_T} at accuracy (85.8±4.2**%**) comparable to
sucrose-cue vs. non-sucrose-cue, VP population classifiers were consistently less accurate
(76.8±3.7**%**) at {X_K vs. X_T} than the aforementioned intervalence classifiers. This suggests that
whereas OT_{D2} encodes odor identity agnostic to the valence, VP does not encode identity at all
but rather encodes reward contingency or positive value. OT_{D1} pairwise classification was a
mixture of the other 2 regions: sucrose-cue vs. non-sucrose-cue classification was more
accurate than most other pairwise classifications (range: 86.4-94.3%), but the {S_K vs. S_T}
classification was comparably accurate (90.9±4.7%). This rules out the interpretation that OT_{D1}
strictly encodes value since the identity of 2 sucrose-cues can be decoded well. To address the
possibility that our results are due to the limitations of linear classification, we repeated the
analysis using support vector machines (SVM's) with a radial basis function kernel and found
we could draw the same conclusions ([FigS12E](#)). Similarly, to verify our results are not
epiphenomena of forcing the data into binary classification, we looked at population-level MNR
classifiers trained on day 6 data. Importantly, we observe high confusion between 2 sucrose
cues in MNR classifiers trained on VP data, but not those trained on OT_{D2} or OT_{D1} data
([FigS12F](#)), corroborating through an alternate analysis method that VP population activity
encodes reward contingency whereas either OT subpopulations likely encode identity.

The fact that VP populations showed higher decoding for odor pairs of unequal sucrose-
contingency provides strong evidence that VP encodes reward-contingency more than identity.
Results from OT decoder analyses, however, are less intelligible: all 15 odor pairs, regardless of
sucrose-contingency, could be decoded with above-chance success. Though this result is
consistent with OT populations encoding identity rather than valence, it does not rule out the
possibility that valence and identity are both encoded. In the context of cue-association, 2 cues
of different valence cannot have the same identity, meaning that good decoding of {S_K vs. P_K}

can be extracted from either valence encoding or identity encoding populations. To

disambiguate these 2 possibilities, we looked at the generalizability of pairwise decoders.

Briefly, linear classifiers were trained on each of the 15 possible odor pairs. Afterwards, the

resulting classifier was tested on every other odor pair (Fig5E). We reasoned that if neural

populations encode valence in addition to identity, classifiers trained on any odor pair of unequal

sucrose-contingency should consistently perform above chance on a different odor pair of

unequal sucrose-contingency (e.g. train on $\{S_K$ vs. $P_K\}$, test on $\{S_T$ vs. $X_T\}$). In other words, given

valence encoding, $\{S_K$ vs. $P_K\}$ should be discriminable in a way that can also discriminate $\{S_T$ vs.

$X_T\}$. As expected, VP population decoders were consistently generalizable when trained on odor

pairs of unequal sucrose-contingency then tested on other odor pairs of unequal sucrose-

contingency (Fig5F,G). OT_{D2} population decoders, on the other hand, showed negligible

generalizability across pairs of unequal sucrose-contingency. Similarly to other metrics of

valence encoding, we found that OT_{D1} displayed a generalizability in between that of VP and

OT_{D1}, suggesting that OT_{D1} could encode some valence in addition to identity. However, we

note that the VP population, on average, vastly outperforms OT_{D1} at generalized valence

decoding (95.0±2.0% vs 78.5±3.9%).

After performing these population-level analyses, we noticed a discrepancy: although

single-neuron intervalence decoding was worse in OT than in VP (Fig3M), population-level

intervalence decoding was comparable between either OT subpopulations and the VP (Fig5A-

D). This led us to speculate that the encoding of odor information had a higher dimensionality in

OT than in VP. To explicitly compare the dimensionality of VP and OT population activities, we

looked at the extent to which the population vector is spread across multiple axes using

principal component analysis (PCA). Dimensionality can further be quantified using the

participation ratio (PR) of a population, which is the square of the sum of eigenvalues of its

covariance matrix divided by the sum of the squares of its eigenvalues (Litwin-Kumar et al.,

2017; Recanatesi et al., 2019). This value will have a range of 1 to n, where n is the total

number of features. If a single principal component can describe all of the total population

variance (i.e. the data is low-dimensional), the population will have PR equal to 1. Conversely, if

every principal component equally describes n^{th} of the total variance (i.e. the data is high-

dimensional), the population will have PR equal to n. Because the number of total neurons

recorded was different between OT and VP experiments, we first assessed if and how the

normalized PR would vary with the number of total neurons through random sampling (Fig5H).

After observing a consistent decrease in PR with increasing n, we compared the PR of OT and

VP animals by repeatedly subsampling a fixed number of neurons (k=15) and found that VP

animals had lower PR (PR_{VP} =5.83±0.80) than either OT_{D2} (PR_{D2}=9.61±0.37) or OT_{D1}

(PR_{D1}=9.24±0.44) animals after training (Fig5I). There was also a difference, however, in how

valence information vs. identity information was encoded by VP populations. Though the first

PC of each VP population was sufficient to train adequate {S_K vs. P_K} decoders (CV accuracy$_{PC1}$

= 85.5±2.7%), all 15 PC's were required for comparable {S_K vs. S_T} decoding (CV accuracy$_{PC1:15}$

= 75.1±13.4%) (Fig5J-L). In either OT populations, the first PC did not support good decoding of

either {S_K vs. P_K} or {S_K vs. S_T}. All together, our population-level analysis shows that VP

encodes valence, but not identity, in low-dimensional space, OT_{D2} encodes identity but not

valence in high-dimensional space, and OT_{D1}, though potentially has valence information,

primarily encodes identity in high-dimensional space.

Analyses at the single-neuron and population levels showed that VP activity encodes

reward contingency, rather than the identity, of the olfactory stimulus. However, due to the task

design, the reward-contingency of a stimulus was highly correlated with the vigor of licking

(Fig2F). This raised concerns that some neurons classified as robust reward-contingency

encoders were potentially encoding motor-related information. Indeed, many VP neurons

showed consistent increases in fluorescence time-locked to the onset of a licking-bout

(FigS13A-B), and could be used to train distributed lag models to predict onset of licking bouts

(FigS13C). Across all VP neurons, we observed a positive and significant correlation between a

neuron's valence decoding ability and licking decoding ability (FigS13D; slope=0.41, p=2.2x10^{-}

^{10}, R^2=0.28). This motivated us to develop a new conditioning paradigm that could decouple

reward-contingency of an odor cue from the behavioral output.

Initially, we attempted to train animals on a symmetric Go/No-Go operant task where

reward delivery was contingent on licking or withholding licks during odor. However, consistent

with previous findings (Gubner et al., 2010), we found that animals struggled to learn the No-Go

behavior in comparison to the Go behavior (data not shown). In an operant paradigm, this leads

to a problematic difference in valence of Go/No-Go cues. Consequently, we opted to develop a

classical conditioning paradigm whereby licks were encouraged/discouraged by physically

moving the lick spout before odor presentation (Fig6A).

Briefly, headfixed animals were presented with 1 of 6 odors in pseudorandomized order.

During the presentation of 3 of these odors, the lick spout was moved away from the mouse

with a linear stepper motor. These odors are denoted as N odors (N for No-lick spout). During

the presentation of the other 3 odors, the lick spout remained within licking distance of the

mouse's tongue. These odors are referred to as L odors (L for lick spout). 1 odor from each

group served as a control odor that had 0% reward-contingency (L_X, N_X). The other 2 odors in

each group were paired with sucrose at low (50%) or high (100%) probability (L_{lo}, L_{hi}, N_{lo}, N_{hi}).

We reasoned that this contingency could allow us to make pairwise comparisons where one

odor has a higher value but lower anticipatory licking than the other (e.g. N_{hi} vs. L_{lo}).

To monitor anticipatory licking in the absence of the lick spout, we trained a distributed lag

model (DLM) using features of the mouse's face tracked using DeepLabcut (Mathis et al., 2018)

(FigS13E-G). We chose to pool data across all mice from the day of the highest licking

differential between L_{hi} and N_{hi} odors (FigS13H) to maximize the decoupling of value and motor

output in our analysis. Anticipatory licking during L_{lo} or L_{hi} began during the last second of the

odor and increased gradually until sucrose delivery whereas licking during N_{lo} or N_{hi} was

delayed by about one second (Fig6B). When quantified across animals on their days of highest

lick differential, we found that mice consistently licked most during L_{hi}, followed by L_{lo} and N_{hi},

then N_{lo} (Fig6C). Mice showed little to no licking during either control odors. Thus, this

behavioral assay affords us the opportunity to assess the decoupled effects of reward-

contingency and licking vigor on neural activity. To begin to characterize the presence of

reward-contingency and/or licking vigor encoding in the VP, we first pooled and clustered the

neural activity taken from 5 animals on their days of highest lick-differential (Fig6D). When

clustering VP neurons into 3 clusters, we found that one cluster (I) showed a largely similar

inhibitory response to the 4 sucrose-paired odors, but not control odors- much like cluster I from

the previous conditioning experiment (Fig3A-B, Fig6D-E). Another cluster (II) by comparison,

showed a varied excitatory response to each of the 4 sucrose-paired odors, much like cluster III

from the previous experiment. Cluster II neurons seemed to have a particularly strong response

to L_{hi} for which there was most anticipatory licking. This led us to speculate the existence of both

reward-contingency encoding and vigor encoding neurons in the VP.

To test this directly, we quantified single neuron decodability of odor pairs and examined

how correlated decoding along the reward-contingency axis is to decoding along vigor axis

(Fig6F-G). We reasoned that auROC values for {L_{hi} vs. N_{hi}} would be high for vigor encoding

neurons but not value encoding neurons given these 2 odors have the same reward-

contingency but disparate licking behaviors. Similarly, we reasoned that auROC values for {N_{hi}

vs. N_X} would be high for reward-contingency encoding neurons but not vigor encoding neurons

given there is a large difference in value but small difference in licking between these 2 odors.

First, we saw that while single neuron decodability along the reward-contingency axis (e.g. {N_{hi}

vs N_X}) was higher than along the lick axis (e.g. {L_{hi} vs N_{hi}}), there were more neurons that could

decode {N_{hi} vs N_X} than could decode {L_X vs N_X} at auROC>0.75 (Fig6F-G). Furthermore, we

saw a lack of significant correlation between the single-neuron decodability of 2 odors that had

similar licking but different reward-contingency ({N_{hi} vs N_X}) and the decodability of 2 odors that

had different licking but same reward-contingency ({L_{hi} vs. N_{hi}}) (Fig6H; slope=0.038, p=0.61,

R^2=0.0039). This decoupling suggests that reward-contingency and vigor information are both

encoded in the VP but by different populations. As a control, we saw a significant correlation

between 2 pairwise comparisons that both had high difference in reward-contingency ({N_{hi} vs

N_X} and {L_{hi} vs L_X}) (Fig6I; slope=0.60, p=3.8x10^{-10}, R^2=0.30).

We also performed the converse experiment where the $\Delta\Delta F/F_{baseline}$ values of each

neuron were linearly fitted to either 1) the reward contingency, 2) the anticipatory licking or 3)

both values and the interaction term. Then, we compared the ΔR^2 when either variable was

omitted in the model and plotted the $\Delta R^2_{-valence}$ against the $\Delta R^2_{-licking}$ (Fig6J). We reasoned that, if

a typical VP neuron's activity could be well-explained by either reward-contingency or vigor but

not both, we would see points along either x or y-intercepts. On the other hand, if a typical

neuron's activity could be well-explained by a linear combination of the 2 variables, we would

see data fall along a line of positive slope. We found that most neurons tended to have large

$\Delta R^2_{-valence}$ *or* large $\Delta R^2_{-licking}$ values but not both, supporting the idea that 2 largely non-

overlapping sets of VP neurons encode reward-contingency *or* vigor but not both.

Lastly, we trained linear classifiers of pairwise odor comparisons using population-level

activity to assess if both reward-contingency and vigor information were present in the

population-level activity. Consistent with single-neuron decoder analysis, we found that {L_{hi} vs

L_X} and {N_{hi} vs N_X} could both be decoded better than {L_X vs N_X} (Fig6K-L). Because we train

each classifier using simultaneously recorded neural activity (i.e. from a single animal), we had

a total of 75 classifiers (15 pairwise classifiers for 5 animals). The cross validated accuracies of

these classifiers were then fitted to a linear model of pairwise differences in either 1) reward-

contingency, 2) anticipatory licking, or 3) both. If the population VP activity encodes either

reward-contingency or vigor but not both, we expect to see one of the single-variable models

outperform the other greatly. But if the population VP activity encodes both variables, we expect

the multivariable model would outperform either single model. We found that Δlicking has a

weak and not significant relationship with pairwise CV accuracy (Fig6M; slope=0.076, p=0.079,

R^2=0.042). By comparison Δreward-contingency had a larger and statistically significant

correlation with pairwise accuracy (Fig6N; slope=0.15, p=3.7×10^{-4}, R^2=0.16). The combined

model, however, showed larger coefficients and larger R^2 than either single variable model,

suggesting an additive effect of both features on CV accuracy (Fig6O; accuracy = 0.18Δlick +

0.24ΔP(S) - 0.22Δlick*ΔP(S) +0.65, R^2=0.23). Thus, we conclude that both reward-contingency

and licking vigor are encoded in the population-level activity of VP neurons.

Discussion

Our anatomical investigations demonstrate that the primary output of the OT is to the VP, with minimal connections to the VTA. Given its constrained connectivity, we propose that the OT to VP circuit is an ideal model system for examining how the encoding of reward cues is transformed across brain circuits. Utilizing comparative longitudinal imaging, we found that VP, but not OT_{D2}, robustly encodes the sucrose-contingency of odors. Although our analyses revealed that sucrose-contingency influences odor-evoked responses in OT_{D1} neurons more so than in OT_{D2} neurons, other evidence suggests value encoding is not the appropriate framework for interpreting OT_{D1} activity. Specifically, information about sucrose-contingency in OT_{D1} resides in a high-dimensional space and generalizes poorly, whereas VP encodes reward-contingency robustly in a low-dimensional and generalizable manner. Thus, we suggest that the changes in OT_{D1} activity are more likely to reflect increased contrast than value. Finally, using a novel classical conditioning paradigm, we assigned motor-related signals and expected-value signals to non-overlapping VP subpopulations.

Some of our findings were unexpected. For example, we found no evidence that either OT_{D1} or OT_{D2} have significant extrapallidal outputs. This is in direct contrast to a previous study which reported that OT_{D1} neurons, and to a lesser extent, OT_{D2} neurons, project to the LH and VTA (Zhang et al., 2017b). We suspect that at least some of the VTA labeling Zhang and colleagues observed from anterograde viral tracing experiments could be due to backflow of the tracer virus in nuclei immediately dorsal to the OT (e.g. AcbSh). As a critical control, we provide evidence that retrograde tracing from VTA robustly labels AcbSh neurons but hardly any OT neurons. And the few VTA projecting OT neurons we did observe were restricted to the distal portions of layer III bordering the VP. Consistent with this, quantification of OT afferents is glaringly absent from 2 independent characterizations of brainwide inputs onto VTA (Beier et al., 2015; Faget et al., 2016). In contrast, OT has been reported to be one of the most prominent inputs to both GAD2+ and Vglut2+ VP neurons (Stephenson-Jones et al., 2020). We note that another retrograde tracing study (In 't Zandt et al., 2019) was unable to corroborate reports of OT to OB projections made by Zhang and colleagues from their anterograde viral tracing. This highlights the importance of verifying anterograde tracing experiments with retrograde tracing, especially when the identified output area is also a well-known target of nuclei dorsal to the injection site. It is difficult, however, to completely rule out the existence of OT to midbrain

projections due to the limitations of our experiments: we primarily targeted layer II in the

anteromedial portion of the OT for anterograde tracing and only tested the VTA with retrograde

tracers. More posterior and/or lateral portions of the OT could have extrapallidal outputs

posterior to the VTA. Despite these caveats, the evidence strongly suggests that *Drd1+* neurons

in the anteromedial portion of the OT, unlike in the AcbSh, have minimal extrapallidal

projections.

Interestingly, there is a parallel discrepancy between the anatomical organization of

dorsal striatum (DS) vs. ventral striatum (VS): SPN's of the DS project exclusively to either the

substantia nigra pars reticulata (SNr) of the midbrain (*Drd1+*) or the exterior portion of the

globus pallidus (GPe) (*Drd2+*), but *Drd1+* neurons in the VS (Acb) project to both the VTA of the

midbrain and the ventral pallidum (Kupchik et al., 2015). Our results hint at a gradient of

anatomical connectivity where the most dorsal *Drd1+* SPN's project primarily to the midbrain

and the most ventral *Drd1+* SPN's (i.e. OT_{D1} neurons) project primarily to the pallidum.

Functionally, the lack of evidence for OT_{D1} to midbrain connectivity challenges the dichotomy of

direct vs. indirect pathways in the ventral basal ganglia. In this model, DA orchestrates motor

initiation by oppositely modulating *Drd1+* and *Drd2+* SPN's, which have differential downstream

targets. Given the lack of clear differences in OT_{D1} and OT_{D2} projections, we think this canonical

model of basal ganglia connectivity inadequately explains the functional consequences of DA

modulation in the OT.

Though we found little difference in the output patterns of OT_{D1} and OT_{D2} neurons, we

observed differences in how these 2 subpopulations encode odor valence. Consistent with a

previous report (Martiros et al., 2022), we found that OT_{D1} activity, more than OT_{D2} activity, is

modulated by reward contingency. For example, OT_{D1} neurons, but not OT_{D2} neurons, were

more likely to respond to sucrose-paired odors than other odors. And the magnitude of

responses in OT_{D1} but not OT_{D2} neurons were significantly larger to sucrose-paired odors than

to other odors. We refrain, however, from concluding that the primary feature encoded in OT_{D1}

neurons is value or reward contingency, for the following reasons. First, the above-mentioned

effects of sucrose-contingency on neural activity are much stronger for VP than for OT_{D1}.

Additionally, whereas more than 50% of VP neurons could be categorized as reward-

contingency encoders, this figure was less than 20% for OT_{D1}. Lastly, population-level decoders

trained on odor pairs of different valence can generalize in the case of VP populations, but not

OT_{D1} populations. While we acknowledge that there is poor standardization when it comes to

defining value encoding, it is unlikely that discrepancies between our conclusions and those of

Martiros et al. stem from differences in interpretation alone. Comparative examination of our

analyses reveals clear dissimilarities in the effect-size of shared metrics (e.g. % odor

responsive). Given the high Z-resolution afforded by 2-photon microscopy, it is probable that we

recorded from different layers of the OT, which should not be assumed to have identical

physiology. We note that the lens placements in our experiments are considerably more ventral

than those reported in Martiros et al. This raises the possibility that some of the neurons

recorded in Martiros et al. could be from the rostral portion of the VP which lies immediately

dorsal to layer III of the OT. Though *Adora2a* and *Drd1* are not expressed as mRNA in the VP,

the BAC-transgenic lines used for both the present work and work by Martiros et al. labels

neurons in the VP.

Our comparison of OT and VP is reminiscent of previous comparisons made between

value encoding in VP and NAc (Ottenheimer et al., 2018; Richard et al., 2016). These

publications showed that VP encodes incentive value more robustly than the NAc. Given that

OT and NAc share many anatomical, physiological, and molecular traits, it is tempting to

speculate that the encoding schemes, too, would be similar between the 2 areas. Optogenetic

activation of OT_{D1} supports RTPP (Murata et al., 2019), as does activation of D1 or D2 neurons

of the NAc (Soares-Cunha et al., 2020). While we acknowledge stimulation experiments provide

unique insights that cannot be obtained from recordings alone, we note that SPN's have

extensive inhibitory collaterals and exhibit high-dimensional activity. Given these peculiarities of

the striatum, we predict that bulk stimulation leads to activity patterns so far outside the

physiologically relevant range that it warrants conservative extrapolations regarding their

endogenous role.

An exciting conclusion from our work is that, within the context of our conditioning

paradigm, the dimensionality of neural activity was much lower in VP than in OT. Furthermore,

the dimensionality of the imaged subpopulations were anti-correlated with the robustness of

sucrose-contingency encoding: OT_{D2} displayed the highest dimensionality and lowest value

encoding whereas VP displayed the lowest dimensionality and highest value encoding. As

discussed elegantly by others (Chu et al., 2016; Shannon, 1948), there is generally a tradeoff

between the efficiency of a neural population (i.e. its total information capacity) and the

robustness of its encoding scheme (i.e. redundancy of encoding). Consistently, it is likely that

VP neurons display such robust encoding of valence, in large part, due to the loss of odor

identity information. By comparison, OT populations may be able to encode information about
the large olfactory identity space due to their high dimensionality. We speculate that the
extensive inhibitory collaterals among SPN's play a role in enforcing the high dimensionality of
OT activity. Though it is entirely unknown what anatomical or physiological strategies are used
to reduce VP dimensionality, we consider this an important piece of the puzzle in understanding
VP computations.

We saw little evidence of negative valence neurons in any of the 3 populations that
were imaged. This was surprising given previous reports of negative valence neurons in the VP
(Stephenson-Jones et al., 2020). We consider 2 potential explanations for this discrepancy.
First, it is possible that our conditioning paradigm was not sufficiently aversive for the animals.
Although our behavioral evidence for aversive association is significant, it is less robust than
sucrose association raising the possibility that the learning was insufficient. This could be due to
the fact that we targeted the airpuff to the animal's hindquarters rather than to the face. But we
note that in a previous report, airpuff delivery to the snout and to hindquarters elicited similar
ingress response in a burrowing assay (Fink et al., 2019). Additionally, we observed clear
unconditioned responses to the airpuff itself. For these reasons we think it is more probable
that, while negative valence neurons do exist in the VP, as has been reported, they were
outside of our field-of-view. Previous work in the VP supports positive and negative valence as
being encoded by *Vgat+* and *Vglut2+* neurons, respectively (Faget et al., 2018; Stephenson-
Jones et al., 2020). Most *Vglut2+* neurons are found in the dorsomedial portion of the VP,
whereas our lenses were specifically targeted to the ventrolateral portion where we found the
most OT afferents. Given this distinction, our results are not inconsistent with previous reports of
negative valence neurons in the VP.

In our work, we successfully described key differences in how reward cues are encoded
in 2 synaptically connected nuclei. But what insights can we infer about the role of OT on
shaping VP activity through this comparison? The most salient observation of VP activity is the
large and widespread excitatory responses to sucrose-cues. Though the effect size is smaller,
OT_{D1} neurons also showed larger excitatory responses to sucrose-cue when compared to other
odors. Given that these neurons are GABAergic and their primary target are the VP neurons, it
is difficult to explain how these 2 responses are related. We consider 3 possible explanations for
this paradox. First, in addition to large excitatory responses that were specific to the sucrose-
cues, we also observed inhibitory responses that were specific to the sucrose-cues. It is

possible that the excitatory VP activity during sucrose-cue presentation is driven mainly by the

numerous excitatory afferents (Pir, BLA, etc.) while the inhibitory VP activity is driven mainly by

OT_{D1} and OT_{D2} afferents. In a second model, there could be mechanisms downstream of

somatic activity that could explain the discrepancy. For example, though brief optical stimulation

of D2 neurons in Acb leads to a decrease of VP activity, prolonged activation causes an

increase in VP activity via the δ-opioid receptor (Soares-Cunha et al., 2020). Our experiments

do not provide any information on how neuropeptide release from OT neurons is different during

presentation of sucrose-cue vs. control odor. Similarly, we cannot measure if and how positively

valent stimuli change the input-output-function of OT neurons. Previous reports have found that

Drd2 agonism in Acb neurons leads to a decrease in collateral inhibition through a presynaptic

mechanism (Dobbs et al., 2016). Given that more DA is expected to be released during

presentation of sucrose-cues, it is plausible that the probability of GABA release from OT

boutons onto VP dendrites is affected. In a third and perhaps the most parsimonious model,

endogenous OT activity does not contribute significantly to explaining the bulk excitatory activity

in VP. This goes against the prevailing working model in Acb to VP circuit which assumes that

Acb excitation leads to VTA disinhibition by inhibiting the VP. And while there is evidence

supporting from bulk stimulation of D1 or D2 neurons in the Acb (Soares-Cunha et al., 2020),

under endogenous conditions, both Acb neurons and VP neurons are excited in response to

reward-cues (Lederman et al., 2021; Ottenheimer et al., 2018). Furthermore, given that

GABAergic synapses from SPN's to VP neurons is likely dendritic (Bolam et al., 1986), we think

it is unlikely that OT to VP drives large-scale shunting of action potential in the presence of

excitatory drive from other areas known to respond preferentially to reward cues such as the

BLA (Beyeler et al., 2018) or the OFC (Wang et al., 2020). We consequently propose an

alternate framework in which the mechanistic role of the OT in this circuit is to provide

spatiotemporally precise inhibition to coordinate the integration of excitatory inputs onto VP.

This form of inhibition could powerfully gate which excitatory synapses go through Hebbian

potentiation vs. anti-Hebbian depression. Under such a framework, OT would function as a

high-dimensional filter for VP neurons to adaptively scale its various excitatory afferents.

In this work, we present evidence that contradict previous anatomical and physiological

characterizations of the OT. We conclude that the anteromedial portion of the OT sends high-

dimensional information about odor identity primarily to the VP and not the VTA. By directly

comparing OT and VP population-level activity in the same paradigm, we bridge together, for

the first time, the fields of OT and VP. This provides valuable context which not only helps us

evaluate past conclusions about valence encoding in the OT but also consider the implications

of the stimulus-evoked activity in the OT. Lastly, we propose a novel framework for

understanding the contributions of striatopallidal projections in reward processing. In our

framework, OT neurons guide VP's integration of its excitatory afferents rather than being the

primary drivers of VP activity. We believe this framework effectively addresses the observed

paradox that both OT_{D1} and VP show increased excitatory activity to reward cues despite being

connected through inhibitory projections. Furthermore, we speculate this framework will be

useful for understanding striatopallidal projections at large.

Methods

Stereotaxic Surgery:

All procedures were approved by the UCSD Institutional Animal Care and Use

Committee. Animals were anesthetized with isoflurane (3% for induction, 1.5-2.0% afterward)

and placed in a stereotaxic frame (Kopf Model 1900). Mouse blood oxygenation, heart rate and

breathing were monitored throughout surgery, and body temperature was regulated using a

heating pad (Physio Suite, Kent Scientific). A small craniotomy above the injection site was

made using standard aseptic technique. Virus was injected with needles pulled from capillary

glass (3-000-203-G/X, Drummond Scientific) at a flow rate of 2nl/s using a micropump (Nanoject

III, Drummond Scientific). For OT anterograde tracing experiments, 50µl of AAV9-phSyn1-

FLEX-tdTomato-T2A-SypEGFP-WPRE diluted to 10^{12} vg/ml was injected into the rostral portion

of the medial OT (AP: 1.6mm, ML: -1.0mm, DV: -5.375mm) in *Drd1*-Cre or *Adora2a*-Cre mice.

For VP retrograde tracing experiments, 100 nl of Cholera Toxin Subunit B CF 488A (Biotium)

was injected at into the caudal portion of the ventrolateral VP (AP: 0.75mm, ML: -1.4mm, DV: -

5.4mm) and 100 nl of Cholera Toxin Subunit B CF 543 (Biotium) was injected into the

dorsomedial VP (AP: 0.75mm, ML: -1.0mm, DV: -5.35mm) in C57BL6/J mice. For VTA

retrograde tracing experiments, 100 nl of Cholera Toxin Subunit B CF 647 (Biotium) was

injected to the rostral portion of the VTA (AP: -3.1mm, ML: 0.8mm, DV: -4.5mm) in C57BL6/J

mice. CTB injections were done at 1 mg/ml dilution in PBS. In some cases, tracers were

injected bilaterally and each hemisphere was analyzed independently. Following each injection,

the injection needle was left at the injection site for 10 minutes then slowly withdrawn.

For imaging experiments, the skull was prepared with OptiBond^{TM} XTR primer and

adhesive (KaVo Kerr) prior to the craniotomy. After performing a craniotomy 800 um in diameter

centered around the virus injection site, a 27G blunt needle was used to aspirate 1.5 mm below

the brain surface. For OT imaging experiments, 500 ul of AAV9-syn-FLEX-jGCaMP7s-WPRE

(Addgene viral prep #104491-AAV9) was diluted to 10^{12} vg/ml and injected into the left and

rostral portion of the medial OT in D1-Cre or A2A-Cre mice. For VP imaging experiments, 300 ul

of AAV9-syn-jGCaMP7s-WPRE (Addgene viral prep #104487-AAV9) was diluted to 10^{12} vg/ml

and injected into the left and caudal portion of the ventrolateral VP in C57BL6/J mice. Following

the viral injection, a head-plate (Model 4, Neurotar) was secured to the mouse's skull using

light-curing glue (Tetric Evoflow, Ivoclar Group). At least 30 minutes after viral injection, a

600um GRIN lens (NA, ~1.9 pitch, GrinTech) was sterilized with Peridox-RTU then slowly

lowered at a rate of 500 um/min into the craniotomy until it was 200 um dorsal to the injection

coordinate. The lens was adhered to the surface of the skull using Tetric Evoflow. We then

placed a hollow threaded post (AE825ES, Thorlabs) to act as a housing for the lens and

adhered it using Tetric Evoflow. Any part of the skull that was still visible was covered using

dental cement (Lang Dental). Finally, the housing was covered with a Nylon cap nut

(94922A325, McMaster-Carr) screwed onto the thread post to protect the lens in between

imaging. Animals were left on the heating pad until they fully recovered from anesthesia.

Histology:

Mice were administered ketamine (100 mg/kg) and xylazine (10 mg/kg) and euthanized

by transcardial perfusion with 10 ml of cold PBS followed by 10 ml of cold 4% paraformaldehyde

in PBS. Brains were extracted and left in a 4% PFA solution in PBS overnight. 50 um coronal

sections were cut on a vibratome (VT1000, Leica). A subset of tissue was labeled with the

following antibodies: Rb α-substance P (1:1,000 dilution; 20064, Immunostar), Rb α-TH (1:1,000

dilution; AB152, Millipore), Dk α-Rb Alexa Fluor$_{TM}$ 488 (1:2,000 dilution; A-210206, Thermo

Fisher Scientific), Dk α-Rb Alexa Fluor$_{TM}$ 647 (1:2,000 dilution; A-31573, Thermo Fisher

Scientific). Slices were mounted using Fluoromount with a DAPI counterstain (SouthernBiotech)

and imaged on an Olympus BX61 VS120 Virtual Slide Scanner and 10x objective (Olympus).

Brains were harvested 21-30 days or 5-7 days after surgery for anterograde and retrograde

tracing experiments, respectively. Brains injected for Ca^{2+} imaging were harvested within a week

of the last imaging session.

For anterograde tracing quantification, 4-6 slices containing each of the brain regions of

interest (VP, LH, and VTA) were analyzed per animal. To quantify the relative abundance of OT

axons in a given brain region, boundaries for the region were drawn on ImageJ Fiji (National

Institutes of Health) with reference to the Paxinos and Franklin Mouse Brain Atlas. Afterwards,

the percentage of the 16-bit pixels within the boundary that had intensity above 200 was

quantified. For retrograde tracing experiments, cells were counted manually every 4th slice.

For anterograde tracing quantification, 4-6 slices containing each of the brain regions of

interest (VP, LH, and VTA) were analyzed per animal. To quantify the relative abundance of OT

axons in a given brain region, boundaries for the region were drawn on ImageJ Fiji (National

Institutes of Health) with reference to the Paxinos and Franklin Mouse Brain Atlas. Afterwards,

the percentage of the 16-bit pixels within the boundary that had intensity above 200 was

quantified. For retrograde tracing experiments, cells were counted manually every 4th slice.

Behavior

Mice were water restricted to reach 85-90% of their initial body weight and given access

to water for 5 minutes a day in order to maintain desired weight. Prior to imaging, mice were

habituated to the head fixation device (Neurotar) and treadmill for 3-5 days, 15-30 minutes per

session. The treadmill parts were 3D printed using a LCD printer (X1-N, EPAX) from publicly

available designs (Jackson et al., 2018). During habituation, mice were provided 10% sucrose

from the water spout. Walking and licking behaviors were measured using a quadrature encoder

(HEDR-5420-es214, Broadcom) and a capacitance sensor (1129_1, Phidgets), respectively. A

video feed of the animal's face was also recorded using a camera (acA1300-30um, Basler) with

a 8-50mm zoom lens (C2308ZM50, Arducam) at 20 Hz with infrared illumination (VQ2121,

Lorex Technology).

Odor was delivered to the mouse using a custom-built olfactometer. Compressed

medical air was split into 2 gas-mass flow controllers (GFC17, Aalborg). One flow controller

directed a constant rate of 1.5 L/min to a hollowed out teflon cylinder. The other flow regulator

was connected to a 3-way solenoid valve (LHDB1223418H, The Lee Co.). Prior to odor delivery,

the 3-way valve directs clean air at 0.5 L/min to the teflon cylinder. During odor delivery, the 3-

way valve directs air to an odor manifold, which consists of an array of 2-way solenoid valves

(LHDB1242115H, The Lee Co.), each connected to a different odor bottle. Depending on the

trial type, the appropriate 2-way valve opens, directing 0.5 L/min of air flow through the odor

bottle containing a kimwipe blotted with 50 ul of diluted odor. All odors were diluted in mineral oil

(M5310, Sigma-Aldrich) to 1.5 mmHg. The kinetics and consistency of odor delivery were

characterized for 30 trials of terpinene delivery using a miniature Photoionization Detector

(mPID) (Aurora Scientific, Inc).

During classical conditioning, animals were exposed to the following odors for 2

seconds: 3-hexanone, 3-heptanone, 3-octanone, α-terpinene, α-pinene, and (R)-(+)-limonene

(all odors were purchased from Sigma with the highest available purity). In days 1-3 of training,

each of the 6 odors and associated outcomes were provided 30 times with 12-18 seconds of

inter-trial interval. Hexanone and terpinene were not associated with any outcome, heptanone

and pinene were associated with 2 ul of 10% sucrose, and octanone and limonene were

associated with a 75 psi airpuff delivered to their hindquarters. Sucrose or airpuff was delivered

100-300 ms after the end of odor delivery. Trials were organized into 30 blocks, each of which

consisted of 1 trial of each of the 6 odors in randomized order. In days 4-6 of training, the

outcome contingencies were switched such that heptanone and limonene were not associated

with any outcome, octanone and terpinene were associated with 2 ul of 10% sucrose, and

hexanone and pinene were associated with 75 psi airpuff.

In the lick-no-lick paradigm, trials were also structured into 30 blocks, each of which

consisted of 1 trial of each of the 6 odors in randomized order. Hexanone and terpinene were

not associated with any outcome, heptanone and pinene were paired with 2 ul of 10% sucrose

at 50% chance, and octanone and limonene were paired with 2 ul of 10% sucrose at 100%

chance. 200 ms prior to the onset of 3 of the odors (terpinene, octanone, and limonene), the lick

spout was retracted 30 mm away from the animal's mouth using a linear stepper motor (BE073-

1, Befenybay) and driver (A4988, BIQU). The lick spout would return to its original position 100

ms prior to the earliest possible time of sucrose delivery.

DeepLabCut

DeepLabCut2.3.3 with Tensorflow 2.12 was used to track 4 points on the periphery of the eye during 2-photon Ca^{2+} imaging. The mini-batch k-means clustering method was used to extract a total of 100 frames (20 frames from 5 animals). These frames were labeled and used to train a Deep Neural Network (DNN) model for 100,000 iterations. After the first training session, 20 outlier frames were picked up from each video and added to the training data for a second training session. The area of the eye at a given time point was estimated as an ellipse. For the lick-no-lick paradigm, we used DeepLabCut to track the tip of the tongue, the corner of the mouth, the upper lip and the lower lip. To record licking in the absence of the lick spout, we trained a linear classifier using logistic regression of the following metrics: 1) the confidence score for the tip of the tongue, 2) the confidence score for the corner of the mouth and 3) the Euclidean distance between the upper and lower lip. Data collected from the capacitive lick sensor was used as ground truth for the classifier.

2-photon Ca^{2+} imaging in head-fixed, behaving mice

Mice were habituated to the head-fixation setup for 3 days beginning 8-10 weeks after surgery. Ca^{2+} imaging data was acquired using an Olympus FV-MPE-RS Multiphoton microscope with Spectra Physics MaiTai HPDS laser, tuned to 920 nm with 100 fs pulse width at 80 MHz. Each 128x128 pixel scan was acquired with a 20x air objective (LCPLN20XIR, Olympus), using a Galvo-Galvo scanner at 5Hz. Stimulus delivery and behavioral measurements were controlled through a custom software written in LabVIEW (National Instruments) and operated through a DAQ (USB-6008, National Instruments). Each imaging session lasted between 30-45 minutes and was synchronized with the stimulus delivery software through a TTL pulse. The imaging depth was manually adjusted to closely match that of the first imaging day such that we recorded from overlapping populations across days of imaging. Animals were excluded from analysis if a) histology showed that either the GRIN lens or the jGCaMP7s virus was mistargeted or b) the motion during imaging was too severe for successful motion-correction. 2 animals were excluded due to mistargeting and 2 animals were excluded due to excessive motion.

Image Processing:

Ca^{2+} imaging data were first motion-corrected using the non-rigid motion correction
algorithm NoRMCorre (Pnevmatikakis and Giovannucci, 2017). Afterwards, neural traces were
extracted from the motion-corrected data using constrained nonnegative matrix factorization
(CNMF) (Giovannucci et al., 2019; Pnevmatikakis et al., 2016). Spatial components identified by
CNMF were inspected by eye to ensure they were not artifacts. A Gaussian Mixture Model
(GMM) was used to estimate the baseline fluorescence of each neuron. To account for potential
low-frequency drift in the baseline, the GMM was applied along a moving window of 2,500
frames (500 seconds). The fluorescence of each neuron at each time point t was then
normalized to the moving baseline to calculate $\Delta F/F = F_t - F_{baseline}/F_{baseline}$. All subsequent
analysis was performed using custom code written in MATLAB (R2022b).

Hierarchical clustering of pooled averaged responses

796    $\Delta F/F$ in response to all 6 odors on day 6 were averaged across trials then Z-scored. The
797   resulting trial-average values from the following timebins were averaged across time: 1) the first
798   second during each odor, 2) the last second during each odor, and 3) the first second after each
799   odor. The resulting 18-element vectors were sorted into 6 clusters after agglomerative
800   hierarchical clustering using euclidean distance and ward linkage.

Responsiveness criteria

To determine how many neurons were responsive to a given odor, we compared $\Delta F/F$ at
each frame during the 2 second odor period against a pooled distribution of $\Delta F/F$ values from
the 2-seconds prior to odor onset using a Wilcoxon rank sum test. The resulting p-values were
evaluated with Holm-Bonferroni correction to ensure that familywise error rate (FWER) was
below 0.05. We then calculated the percentage of responsive neurons for each animal to show
the mean and the standard error as a function of time. We also counted the number of neurons
that were significantly responsive for at least 4 frames during the odor period to report the total
percentage of responsive neurons during odor.

Single neuron logistic classifiers

To test how reliably a single neuron's fluorescence could discriminate between 2 odors,

we assessed the performance of binary logistic classifiers trained on a single neuron's

responses to 2 odors. For each neuron and odor pair, we averaged the ΔF/F during the last

second of the odor exposure for each trial then Z-scored across all trials. The resulting 60-

element vector was used to train a linear classifier using logistic regression. The receiver

operator characteristic (ROC) was evaluated for each single neuron pairwise classifier and the

area under the curve (AUC) reported. To test if a given pairwise classifier performed

significantly better than chance, we compared the accuracy of each classifier against a

distribution of 10,000 classifiers trained on shuffled labels.

Normalized ΔΔF/F correlations

To compare the average response of a neuron to each odor, the trial-averaged ΔF/F

during the last second of odor exposure from each trial was averaged and then subtracted from

the trial-averaged ΔF/F during the 2 seconds prior to odor delivery. This ΔΔF/F value was

scaled to the largest positive ΔΔF/F value of each neuron for all odors. To assess the similarity

of the average response to a given pair of odors i and j, we looked at the null linear model in

which all neurons respond identically to both odors, i.e. $\Delta\Delta F_j/F = \Delta\Delta F_i/F$. To assess how well

this describes the data, we report the R^2 value of the fit.

Pairwise euclidean distance

To quantify the differences among population-level responses to the 6 odors, we

quantified the pairwise Euclidean distance between the trajectories of odor responses. First, we

subtracted the ΔF/F values during the 2 seconds prior to odor delivery from each frame then

averaged these values across trials for each odor. The pairwise Euclidean distance at each

frame was computed for each odor pair and normalized to the maximum pairwise distance

measured in all odor pairs at any time bin. These calculations were carried out separately for

each animal and then averaged across biological replicates to report the mean and the standard

error.

Population pairwise classifiers

To assess the discriminability of odor responses in high-dimensional space, we measured the accuracy of binary classifiers for a given odor pair. At each time point relative to odor delivery, we pooled ΔF/F values from all trials during which either odor was presented. These values were then normalized and used to train a linear classifier using either a logistic regression or a Support Vector Machine (SVM). The accuracy of the classifier was evaluated via 5-fold cross-validation. To test if a given pairwise decoder performed significantly better than chance, we compared the accuracy of each classifier against a distribution of 10,000 classifiers trained on shuffled labels. All classifiers were trained on populations of neurons simultaneously recorded from individual mice. The resulting cross-validated accuracies were averaged across biological replicates to report the mean and the standard error.

Dimensionality analysis

To quantify the dimensionality of each simultaneously recorded neural population, we calculated its participation ratio (PR). First, we performed principal component analysis of the whole dataset using the singular value decomposition algorithm. The PR was calculated as the square of the sum of the eigenvalues of the covariance matrix divided by the sum of the square of its eigenvalues (Litwin-Kumar et al., 2017; Recanatesi et al., 2019). To account for the differences in number of recorded neurons across individuals, we bootstrapped the PR by randomly sampling n neurons from each dataset 1,000 times and reported the average PR value.

Author contributions
D.L. and C.M.R. conceived of the project, participated in its development and wrote the manuscript. L.L. assisted with anatomy, histology and behavioral analysis. D.L. performed all imaging experiments and analyzed the data.

Acknowledgments
We thank members of Root lab for discussions, M. Aoi for discussions on data analysis, and T. Komiyama and for comments on the manuscript. This research was supported by grants from the NIH (R00DC014516, R01DC018313), and C.M.R. was a Hellman Fellow.

References

Allen WE, Chen MZ, Pichamoorthy N, Tien RH, Pachitariu M, Luo L, Deisseroth K. 2019. Thirst
regulates motivated behavior through modulation of brainwide neural population dynamics.
Science **364**:253.
Beier KT, Steinberg EE, DeLoach KE, Xie S, Miyamichi K, Schwarz L, Gao XJ, Kremer EJ,
Malenka RC, Luo L. 2015. Circuit Architecture of VTA Dopamine Neurons Revealed by
Systematic Input-Output Mapping. *Cell* **162**:622–634.
Beyeler A, Chang C-J, Silvestre M, Lévêque C, Namburi P, Wildes CP, Tye KM. 2018.
Organization of Valence-Encoding and Projection-Defined Neurons in the Basolateral
Amygdala. *Cell Rep* **22**:905–918.
Bolam JP, Ingham CA, Izzo PN, Levey AI, Rye DB, Smith AD, Wainer BH. 1986. Substance P-
containing terminals in synaptic contact with cholinergic neurons in the neostriatum and
basal forebrain: a double immunocytochemical study in the rat. *Brain Res* **397**:279–289.
Chu MW, Li WL, Komiyama T. 2016. Balancing the Robustness and Efficiency of Odor
Representations during Learning. *Neuron* **92**:174–186.
Dan Y, Poo M-M. 2004. Spike timing-dependent plasticity of neural circuits. *Neuron* **44**:23–30.
Dobbs LK, Kaplan AR, Lemos JC, Matsui A, Rubinstein M, Alvarez VA. 2016. Dopamine
Regulation of Lateral Inhibition between Striatal Neurons Gates the Stimulant Actions of
Cocaine. *Neuron* **90**:1100–1113.
Faget L, Osakada F, Duan J, Ressler R, Johnson AB, Proudfoot JA, Yoo JH, Callaway EM,
Hnasko TS. 2016. Afferent Inputs to Neurotransmitter-Defined Cell Types in the Ventral
Tegmental Area. *Cell Rep* **15**:2796–2808.
Faget L, Zell V, Souter E, McPherson A, Ressler R, Gutierrez-Reed N, Yoo JH, Dulcis D,
Hnasko TS. 2018. Opponent control of behavioral reinforcement by inhibitory and excitatory
projections from the ventral pallidum. *Nat Commun* **9**:849.
Fink AJP, Axel R, Schoonover CE. 2019. A virtual burrow assay for head–fixed mice measures
habituation, discrimination, exploration and avoidance without training. *Elife* **8**:e45658.
Fujimoto A, Hori Y, Nagai Y, Kikuchi E, Oyama K, Suhara T, Minamimoto T. 2019. Signaling
Incentive and Drive in the Primate Ventral Pallidum for Motivational Control of Goal-
Directed Action. *J Neurosci* **39**:1793–1804.
Gadziola MA, Stetzik LA, Wright KN, Milton AJ, Arakawa K, Del Mar Cortijo M, Wesson DW.
2020. A Neural System that Represents the Association of Odors with Rewarded Outcomes
and Promotes Behavioral Engagement. *Cell Rep* **32**:107919.
Gadziola MA, Tylicki KA, Christian DL, Wesson DW. 2015. The olfactory tubercle encodes odor
valence in behaving mice. *J Neurosci* **35**:4515–4527.
Giovannucci A, Friedrich J, Gunn P, Kalfon J, Brown BL, Koay SA, Taxidis J, Najafi F, Gauthier
JL, Zhou P, Khakh BS, Tank DW, Chklovskii DB, Pnevmatikakis EA. 2019. CaImAn an
open source tool for scalable calcium imaging data analysis. *Elife* **8**.
doi:10.7554/eLife.38173
Groenewegen HJ, Russchen FT. 1984. Organization of the efferent projections of the nucleus
accumbens to pallidal, hypothalamic, and mesencephalic structures: a tracing and
immunohistochemical study in the cat. *J Comp Neurol* **223**:347–367.
Gubner NR, Wilhelm CJ, Phillips TJ, Mitchell SH. 2010. Strain differences in behavioral
inhibition in a Go/No-go task demonstrated using 15 inbred mouse strains. *Alcohol Clin Exp*
Res **34**:1353–1362.
Haberly LB, Price JL. 1977. The axonal projection patterns of the mitral and tufted cells of the

olfactory bulb in the rat. *Brain Res* **129**:152–157.

Hollerman JR, Schultz W. 1998. Dopamine neurons report an error in the temporal prediction of
reward during learning. *Nat Neurosci* **1**:304–309.

Igarashi KM, Ieki N, An M, Yamaguchi Y, Nagayama S, Kobayakawa K, Kobayakawa R, Tanifuji
M, Sakano H, Chen WR, Mori K. 2012. Parallel mitral and tufted cell pathways route distinct
odor information to different targets in the olfactory cortex. *J Neurosci* **32**:7970–7985.

Ikemoto S. 2007. Dopamine reward circuitry: two projection systems from the ventral midbrain
to the nucleus accumbens-olfactory tubercle complex. *Brain Res Rev* **56**:27–78.

Ikemoto S. 2003. Involvement of the olfactory tubercle in cocaine reward: intracranial self-
administration studies. *J Neurosci* **23**:9305–9311.

In 't Zandt EE, Cansler HL, Denson HB, Wesson DW. 2019. Centrifugal Innervation of the
Olfactory Bulb: A Reappraisal. *eNeuro* **6**. doi:10.1523/ENEURO.0390-18.2019

Jackson J, Karnani MM, Zemelman BV, Burdakov D, Lee AK. 2018. Inhibitory Control of
Prefrontal Cortex by the Claustrum. *Neuron* **99**:1029–1039.e4.

Jones DL, Mogenson GJ. 1980. Nucleus accumbens to globus pallidus GABA projection:
electrophysiological and iontophoretic investigations. *Brain Res* **188**:93–105.

Kupchik YM, Brown RM, Heinsbroek JA, Lobo MK, Schwartz DJ, Kalivas PW. 2015. Coding the
direct/indirect pathways by D1 and D2 receptors is not valid for accumbens projections. *Nat*
Neurosci **18**:1230–1232.

Lederman J, Lardeux S, Nicola SM. 2021. Vigor Encoding in the Ventral Pallidum. *eNeuro* **8**.
doi:10.1523/ENEURO.0064-21.2021

Litwin-Kumar A, Harris KD, Axel R, Sompolinsky H, Abbott LF. 2017. Optimal Degrees of
Synaptic Connectivity. *Neuron* **93**:1153–1164.e7.

Martiros N, Kapoor V, Kim SE, Murthy VN. 2022. Distinct representation of cue-outcome
association by D1 and D2 neurons in the ventral striatum's olfactory tubercle. *Elife* **11**.
doi:10.7554/eLife.75463

Mathis A, Mamidanna P, Cury KM, Abe T, Murthy VN, Mathis MW, Bethge M. 2018.
DeepLabCut: markerless pose estimation of user-defined body parts with deep learning.
Nat Neurosci **21**:1281–1289.

Millman DJ, Murthy VN. 2020. Rapid Learning of Odor–Value Association in the Olfactory
Striatum. *J Neurosci* **40**:4335–4347.

Murata K, Kinoshita T, Fukazawa Y, Kobayashi K, Yamanaka A, Hikida T, Manabe H,
Yamaguchi M. 2019. Opposing Roles of Dopamine Receptor D1- and D2-Expressing
Neurons in the Anteromedial Olfactory Tubercle in Acquisition of Place Preference in Mice.
Front Behav Neurosci **13**:50.

Newman R, Winans SS. 1980. An experimental study of the ventral striatum of the golden
hamster. II. Neuronal connections of the olfactory tubercle. *J Comp Neurol* **191**:193–212.

Oettl L-L, Scheller M, Filosa C, Wieland S, Haag F, Loeb C, Durstewitz D, Shusterman R,
Russo E, Kelsch W. 2020. Phasic dopamine reinforces distinct striatal stimulus encoding in
the olfactory tubercle driving dopaminergic reward prediction. *Nat Commun* **11**:3460.

Ottenheimer DJ, Bari BA, Sutlief E, Fraser KM, Kim TH, Richard JM, Cohen JY, Janak PH.
2020a. A quantitative reward prediction error signal in the ventral pallidum. *Nat Neurosci*
23:1267–1276.

Ottenheimer DJ, Wang K, Tong X, Fraser KM, Richard JM, Janak PH. 2020b. Reward activity in
ventral pallidum tracks satiety-sensitive preference and drives choice behavior. *Sci Adv* **6**.
doi:10.1126/sciadv.abc9321

Ottenheimer D, Richard JM, Janak PH. 2018. Ventral pallidum encodes relative reward value
earlier and more robustly than nucleus accumbens. *Nat Commun* **9**:4350.

Pnevmatikakis EA, Giovannucci A. 2017. NoRMCorre: An online algorithm for piecewise rigid
motion correction of calcium imaging data. *J Neurosci Methods* **291**:83–94.
Pnevmatikakis EA, Soudry D, Gao Y, Machado TA, Merel J, Pfau D, Reardon T, Mu Y,
Lacefield C, Yang W, Ahrens M, Bruno R, Jessell TM, Peterka DS, Yuste R, Paninski L.
2016. Simultaneous Denoising, Deconvolution, and Demixing of Calcium Imaging Data.
Neuron **89**:285–299.
Recanatesi S, Ocker GK, Buice MA, Shea-Brown E. 2019. Dimensionality in recurrent spiking
networks: Global trends in activity and local origins in connectivity. *PLoS Comput Biol*
15:e1006446.
Rescorla RA. 1972. A theory of Pavlovian conditioning : Variations in the effectiveness of
reinforcement and non-reinforcement. *Classical conditioning, Current research and theory*
2:64–69.
Ricardo JA. 1980. Efferent connections of the subthalamic region in the rat. I. The subthalamic
nucleus of Luys. *Brain Res* **202**:257–271.
Richard JM, Ambroggi F, Janak PH, Fields HL. 2016. Ventral Pallidum Neurons Encode
Incentive Value and Promote Cue-Elicited Instrumental Actions. *Neuron* **90**:1165–1173.
Schultz W, Dayan P, Montague PR. 1997. A neural substrate of prediction and reward. *Science*
275:1593–1599.
Scott JW. 1981. Electrophysiological identification of mitral and tufted cells and distributions of
their axons in olfactory system of the rat. *J Neurophysiol* **46**:918–931.
Shannon CE. 1948. A mathematical theory of communication. *The Bell System Technical*
Journal **27**:379–423.
Smith KS, Tindell AJ, Aldridge JW, Berridge KC. 2009. Ventral pallidum roles in reward and
motivation. *Behav Brain Res* **196**:155–167.
Soares-Cunha C, de Vasconcelos NAP, Coimbra B, Domingues AV, Silva JM, Loureiro-Campos
E, Gaspar R, Sotiropoulos I, Sousa N, Rodrigues AJ. 2020. Nucleus accumbens medium
spiny neurons subtypes signal both reward and aversion. *Mol Psychiatry* **25**:3241–3255.
Stephenson-Jones M, Bravo-Rivera C, Ahrens S, Furlan A, Xiao X, Fernandes-Henriques C, Li
B. 2020. Opposing Contributions of GABAergic and Glutamatergic Ventral Pallidal Neurons
to Motivational Behaviors. *Neuron* **105**:921–933.e5.
Sutton RS. 1988. Learning to predict by the methods of temporal differences. *Mach Learn* **3**:9–
44.
Tachibana Y, Hikosaka O. 2012. The primate ventral pallidum encodes expected reward value
and regulates motor action. *Neuron* **76**:826–837.
Tindell AJ, Smith KS, Berridge KC, Aldridge JW. 2005. VP neurons integrate learning and
physiological signals to code incentive salience of conditioned cues. *Soc Neurosci Abstr*.
Tindell AJ, Smith KS, Peciña S, Berridge KC, Aldridge JW. 2006. Ventral pallidum firing codes
hedonic reward: when a bad taste turns good. *J Neurophysiol* **96**:2399–2409.
Tritsch NX, Sabatini BL. 2012. Dopaminergic modulation of synaptic transmission in cortex and
striatum. *Neuron* **76**:33–50.
Turner MS, Lavin A, Grace AA, Napier TC. 2001. Regulation of limbic information outflow by the
subthalamic nucleus: excitatory amino acid projections to the ventral pallidum. *J Neurosci*
21:2820–2832.
Wang PY, Boboila C, Chin M, Higashi-Howard A, Shamash P, Wu Z, Stein NP, Abbott LF, Axel
R. 2020. Transient and Persistent Representations of Odor Value in Prefrontal Cortex.
Neuron **108**:209–224.e6.
Watabe-Uchida M, Zhu L, Ogawa SK, Vamanrao A, Uchida N. 2012. Whole-brain mapping of
direct inputs to midbrain dopamine neurons. *Neuron* **74**:858–873.

Wesson DW. 2020. The Tubular Striatum. *J Neurosci* **40**:7379–7386.
Wesson DW, Wilson DA. 2010. Smelling Sounds: Olfactory–Auditory Sensory Convergence in
the Olfactory Tubercle. *J Neurosci* **30**:3013–3021.
Wieland S, Schindler S, Huber C, Köhr G, Oswald MJ, Kelsch W. 2015. Phasic Dopamine
Modifies Sensory-Driven Output of Striatal Neurons through Synaptic Plasticity. *J Neurosci*
35:9946–9956.
Zahm DS, Heimer L. 1987. The ventral striatopallidothalamic projection. III. Striatal cells of the
olfactory tubercle establish direct synaptic contact with ventral pallidal cells projecting to
mediodorsal thalamus. *Brain Res* **404**:327–331.
Zhang Z, Liu Q, Wen P, Zhang J, Rao X, Zhou Z, Zhang H, He X, Li J, Zhou Z, Xu X, Zhang X,
Luo R, Lv G, Li H, Cao P, Wang L, Xu F. 2017a. Activation of the dopaminergic pathway
from VTA to the medial olfactory tubercle generates odor-preference and reward. *Elife* **6**.
doi:10.7554/eLife.25423
Zhang Z, Zhang H, Wen P, Zhu X, Wang L, Liu Q, Wang J, He X, Wang H, Xu F. 2017b. Whole-
Brain Mapping of the Inputs and Outputs of the Medial Part of the Olfactory Tubercle. *Front*
Neural Circuits **11**:52.
Zhou L, Furuta T, Kaneko T. 2003. Chemical organization of projection neurons in the rat
accumbens nucleus and olfactory tubercle. *Neuroscience* **120**:783–798.

1031

Figure 1:

Figure 1. OT$_{D1}$ and OT$_{D2}$ primarily project to the lateral portion of the VP. (**A**) Schematic representation of Cre-dependent anterograde axonal AAV tracing experiments used to characterize outputs of OT neurons. *Drd1+* and *Drd2+* neurons were separately labeled by using *Drd1*-Cre and *Adora2a*-Cre mouse lines, respectively. (**B**) Representative images from OT$_{D1}$ (top) vs. the AcbSh$_{D1}$ injection (bot). Target sites (far-left column) are stained with α-tyrosine hydroxylase antibodies to visualize the boundary between VP and OT. (**C**) Quantifying the % of output regions with fluorescence (n=3-4). (**D**) Schematic representation of 2-color retrograde CTB tracing experiment used to confirm OT to VP connectivity. CTB::488 and CTB::543 were injected to the lateral and medial portion of the VP, respectively. (**E**) Representative images of CTB labeled neurons in the OT and Acb. (**F**) The number of labeled cells was quantified (n=4). (**G**) Schematic representation of retrograde CTB tracing experiment used to test OT to VTA connectivity. CTB::647 was injected in the VTA. (**H**) Representative image shows robust AcbSh and AcbC labeling but no OT labeling. (**I**) Quantification of labeling in different nuclei (n=3). Two-way ANOVAs followed by post-hoc t-tests for each output region were performed to generate p-values, which were then corrected for FDR by Benjamini-Hocherg procedure. ***p<0.001, **p<0.01, *p<0.05.

Figure 2:

Figure 2. Head-fixed 2-photon Ca^{2+} imaging of OT$_{D1}$, OT$_{D2}$, or VP neurons during 6-odor conditioning paradigm. (**A**) State-diagram of odor conditioning paradigm. Each trial begins with 2 seconds of odor delivery. Odors are chosen in pseudorandomized order such that the same odor is not repeated more than twice in a row. At the end of odor delivery, there is a variable delay (100-300ms), after which the animal is given either a 10% sucrose solution (S$_K$ and S$_T$), a 70 psi airpuff (P$_K$ and P$_T$), or nothing (X$_K$ and X$_T$). Trials are separated by a variable intertrial interval (ITI; 12-18s). Schematic representation of (**B**) lens implant surgery and (**C**) headfix 2-photon microscopy setup. An example of spatial (**D**) and temporal (**E**) components extracted by CNMF from *Drd1*-Cre animal on day 3 of imaging. (**D**) The spatial footprints of 20 example neurons are shown on top of a maximum-correlation pixel image that was used to seed the factorization. The number displayed over each neuron matches the row number of the temporal components in (**E**). (**F**) An example raster plot (top) and averaged-across-trials trace (bottom) of the licking behavior recorded concurrently as (**D**) and (**E**). The timing of odor delivery is shown as shaded rectangles. The timing of US delivery is shown as arrowheads. (**G**) The mean total licks during each of the odors is shown averaged across all animals (n=17) after application of a moving-average filter with a window size of 10 trials. Red line marks the sucrose and airpuff contingency switch between day 3 and day 4. (**H**) Bar graph showing the ratio of licks during either S$_K$ or S$_T$ to licks during any odor. The effects of imaging day and lens target site were analyzed using a 2-way ANOVA. The imaging day, but not lens target site, was a significant factor on licking accuracy (F$_{day}$=27.64 vs. F$_{lens\ location}$=2.30).

Figure 3:

Figure 3. VP neurons encode reward-contingency more robustly than OT_{D1} or OT_{D2} neurons. (**A**) Heatmap of odor-evoked activities in OT_{D1}, OT_{D2}, and VP neurons from day 6 of imaging. The fluorescence measurements from each neuron were averaged over trials, Z-scored, then pooled for hierarchical clustering. Neurons are grouped by similarity, with the dendrogram shown on the right and a raster plot on the left indicating which region a given neuron is from. Horizontal white lines demarcate the boundaries between the 6 clusters. Odor delivered at 0-2 seconds marked by vertical red lines and US delivery is marked by arrowheads. From left to right, the columns represent neural responses to sucrose-paired ketone and terpene, control ketone and terpene, and airpuff-paired ketone and terpene (S_K, S_T, X_K, X_T, P_K, P_T). (**B**) Average Z-scored activity of each cluster to each of the 6 odors on day 6 of imaging. Yellow bar indicates 2-seconds of odor exposure. (**C**) The distribution of clusters by population. (**D**) Percentage of total neurons that were significantly excited or inhibited by each odor (Bonferroni-adjusted FDR < 0.05) as a function of time relative to odor. Bar represents the mean across biological replicates and x's mark individual animals. (**E**) Bar graph showing % of neurons from each population that are responsive to both sucrose-paired odors in the same direction (left), responsive to only a single odor (middle), or responsive to at least 3 odors (right). (**F**) Scatterplot comparing the magnitudes of S_K responses ($\Delta\Delta S_K$) to S_T responses ($\Delta\Delta S_T$). The dotted line represents the hypothetical scenario where $\Delta\Delta S_K = \Delta\Delta S_T$. For each population, the R^2 value of the 2-d distribution compared to the $\Delta\Delta S_K = \Delta\Delta S_T$ line is reported. (**G**) Same as F but comparing $\Delta\Delta S_K$ to $\Delta\Delta X_K$. (**H**) Lineplot showing the % of neurons from each population where the difference between $\Delta\Delta S_K$ and $\Delta\Delta X_K$ is lower than that between $\Delta\Delta S_K$ and $\Delta\Delta S_T$. (**I**) Bargraph showing % of neurons whose responses to {S_K vs. X_K} can be discriminated by a linear classifier with auROC>0.75. (**J**) Same as H but for {S_K vs S_T}. (**K**) Schematic representation of 4 possible categories for a joint-distribution of {S_K vs. X_K} and {S_K vs. S_T} auROC values. Identity-encoding neurons could be in any quadrant other than the bottom-left whereas valence-encoding neurons should be in the bottom-right quadrant. (**L**) Scatterplot of each neuron's auROC value for {S_K vs. X_K} on the x-axis and {S_K vs. S_T} on the y-axis on days 1, 3 and 6 of imaging. (**M**) Stacked bar graph showing the distribution of neurons from each population that fall into each of the 4 quadrants across the 3 different imaging days.

Figure 4:

Figure 4. Sucrose responsive VP neurons become sucrose-cue responsive after pairing.
(**A**) The spatial footprints of 15 neurons from day 1 are outlined over a max-correlation projection image. (**B**) Heatmap of averaged-over-trials ΔF/F in response to 6 odors on day 1. Odor delivery period is shown with 2 red vertical lines and sucrose/airpuff timing is shown with downward arrowhead. (**C**) An example neuron's responses on day 1 across 30 trials to 6 different odors. Individual trial traces are shown in light gray whereas the averaged-across trials trace is shown in black. Odor delivery period is depicted as shaded rectangles and US delivery is marked by arrowheads. (**D-F**) Same as (**A-C**), respectively, but for day 3. (**G**) Percentage of all tracked neurons that were both sucrose-responsive on day 1 and odor-responsive in the same direction on day 3. (**H**) Scatter plot of averaged-over-trials responses to S_K or S_T on day 1 (x-axis) and day 3 (y-axis). Each point is a neuron that was successfully matched from day 1 and day 3. Neurons from OT_{D2}, OT_{D1}, and VP are plotted as pink circles, blue crosses, and yellow squares, respectively. Neurons that have increased response magnitudes on day 3 would fall between the 2 dotted lines. (**I**) Violin plot showing the distributions of day 3 responsive magnitude - day 1 response magnitude.

Figure 5:

Figure 5. OT encodes odor identity in high-dimensional space and VP encodes reward-contingency in low-dimensional space. (**A**) Average normalized pairwise euclidean distance between odor-evoked population-level activity from day 6 of imaging shown as a function of time relative to odor delivery. Traces show the average value across biological replicates of the same population and the shaded areas represent the average ± SEM. (**B**) A heatmap of the average normalized pairwise distance during the odor delivery period. (**C**) Average CV accuracy of binary pairwise linear classifiers trained on population data plotted against time relative to odor delivery. (**D**). A heatmap of the average CV accuracy during the odor delivery period. (**E**) Schematic representation of generalized linear classification performance for an idealized valence encoder. Each row corresponds to the training odor-pair and each column corresponds to the testing odor-pair. For an idealized valence encoder, the decodability would generalize well across odor-pairs of the equal valence grouping outlined in red. Note that the elements along the diagonal are cases where training and testing odor-pairs are identical and do not reflect generalizability. (**F**) Heatmap representing the maximum generalized linear classification accuracy during odor delivery period averaged across biological replicates for each population. (**G**) Maximum cross validated linear classifier accuracy for S_K vs. X_K (trained on S_K vs. X_K) and the generalized accuracy for S_K vs. X_K after training on S_T vs X_T. (**H**) Average PR normalized to *n* calculated after randomly subsampling an increasing number of neurons. (**I**) Average PR calculated after subsampling 15 neurons. (**J**) Average CV accuracy of linear classifiers trained on {S_K vs. P_K}. For each simultaneously imaged group of neurons, 15 neurons were subsampled and classifiers were trained on an increasing number of principal components. The average across subsampling is reported. (**K**) Average CV accuracy of linear classifiers trained on {S_K vs. S_T}. (**L**) Comparison of the average accuracy of {S_K vs. P_K} classifiers trained on the 1st PC vs. {S_K vs. S_T} classifiers trained on all 15 PC's.

Figure 6:

Figure 6. Separate VP populations encode reward-contingency and licking vigor.
(**A**) State diagram for odor pairing paradigm where lick spout is removed during the presentation of half of the odors. The paradigm is similar to one described in Fig2A with the following key differences: 1) the lick spout is moved away from the animal's mouth during the presentation of half of the odors (N_{hi}, N_{lo}, N_X). 2) sucrose is delivered after a longer variable delay (1.1-1.3s). 3) 2 of the odors have 100% sucrose contingency (L_{hi}, N_{hi}), 2 of the odors have 50% sucrose contingency (L_{lo}, N_{lo}), and the other 2 have 0% sucrose contingency (L_X, N_X). (**B**) Licking behavior to 6 odors averaged across 30 trials from a representative animal. Duration of odor delivery is marked by the shaded rectangle and the average time of sucrose delivery is marked by the arrowhead. The time bin used for subsequent analysis (last 0.5s of odor and first 0.5s of delay) is outlined by square brackets (**C**) Average licks/s for each odor measured between the last 0.5s of odor and the first 0.5s of delay. Data were pooled from the day of highest difference between licks to L_{hi} and N_{hi}. The effect of lick spout presence and sucrose contingency were analyzed using a 2-way ANOVA followed by post-hoc t-tests were performed to generate p-values. These values were then corrected for FDR by the Benjamini-Hocherg procedure. $p<0.001$, $p<0.01$, $p<0.05$. (**D**) Heatmap of odor-evoked activity in VP neurons pooled from each animal's day of highest difference between licks to L_{hi} and N_{hi}. Neurons are grouped according to the clustering dendrogram, shown on the right. Horizontal white lines demarcate the boundaries between the 3 clusters. Odor delivery is marked by vertical red lines. (**E**) Average Z-scored activity of each cluster to each of the 6 odors. Yellow bar indicates 2-seconds of odor exposure. (**F**) The percentage of single-neuron linear classifiers with auROC>0.75 as a function of time relative to odor delivery. Shaded area represents the SEM across biological replicates (n=5). (**G**) Heatmap of the percentage of pooled VP neurons with auROC>0.75 during the last 0.5s of odor and first 0.5s of delay. (**H**) Scatterplot comparing the auROC for {L_{hi} vs N_{hi}} (y-axis) and {N_{hi} vs. N_X} (x-axis) for each neuron. The line of best fit is plotted as a dotted line, with the 95% confidence interval shaded in. (**I**) Same as (H) but comparing the auROC for {L_{hi} vs L_X} (y-axis) and {N_{hi} vs. N_X} (x-axis). (**J**) Scatterplot comparing regression models that explain each neuron's activity on a given trial as a function of anticipatory licking or sucrose contingency. The values plotted are the loss in R^2 in models without anticipatory licking (y-axis) or sucrose contingency (x-axis) when compared to a model with both variables and their interaction term. (**K**) CV accuracy for 5 different odor pairs as a function of time relative to odor delivery. (**L**) Heatmap of average pairwise CV accuracy trained on the last 0.5s of odor and the first 0.5s of delay. (**M**) Scatterplot of all pairwise classifier accuracies from all animals (y-axis) and the corresponding range-normalized average pairwise difference in anticipatory licking (x-axis). (**N**) Scatterplot of all pairwise classifier accuracies from all animals (y-axis) and the corresponding pairwise difference in reward-contingency (x-axis). (**O**) Scatterplot of all pairwise classifier accuracies (y-axis) and the adjusted combined model of ranged-normalized Δlick and Δreward-contingency (x-axis).

Figure S1:

Figure S1. OT$_{D1}$ and OT$_{D2}$ primarily project to the lateral portion of the VP.
(**A**) Serial coronal sections from a representative experiment where AAVDJ-hSyn-FLEX-mRuby-T2A-syn-eGFP virus was injected into the anterior OT of an *Adora2a*-Cre mouse. Sections are roughly 400μm apart from each other and range from +2.5mm to -3.0mm relative to bregma. (**B**) Same as (A) but in a *Drd1*-Cre mouse. (**C**) Schematic showing the centroids of 4 injection sites for OT and AcbSh anterograde tracing experiments. (**D**, **E**) Representative images of the injection sites shown in (C). Sections were counterstained with α-TH to delineate the boundary between the striatum and the rostral ventral pallidum. The centroids of these samples are marked as red x's in (C). (**F**) Schematic showing the centroids of 4 CTB injection sites for lateral and medial VP. CTB::488 injection to the lateral VP are marked by +'s whereas CTB::543 injection to the medial VP are marked by x's. (**G**) Representative image of the injection sites shown in (F). Sections were counterstained with α-Substance P to mark the boundary of the VP. The centroids of this sample are marked by the red + and red x in (F). (**H**) Schematic showing the centroids of 3 CTB injection sites for VTA. (**I**) Representative image of the injection sites shown in (H). Sections were counterstained with α-TH to mark the boundary of the VTA.

Figure S2:

Figure S2 Histological verification of lens implant location.
(**A**) Schematic showing the center, along the AP axis, of the implanted GRIN lens in 6 OT_{D2} jGCaMP7s animals. (**B**) Representative image of lens implant sites shown in (A). Sections were counterstained with α-Substance P to delineate the boundary of the VP. This sample is marked by a red horizontal line in (A). (**C**) Schematic as in (A) for 6 OT_{D1} jGCaMP7s animals. (**D**) Representative image of lens implant sites shown in (C). (**E**) Schematic as in (A) for 5 VP jGCaMP7s animals. (**F**) Representative image of lens implant sites shown in (E). (**G**) Schematic as in (A) for 5 VP jGCaMP7s animals recorded during the lick spout retraction paradigm (Figure 6). (**H**) Representative image of lens implant sites shown in (G). Sections were counterstained with α-Substance P to delineate the boundary of the VP. This sample is marked by a red horizontal line in (G).

Figure S3:

Figure S3. Pooled averaged-over-trials neural activity of all neurons from OT$_{D2}$ animals across days. Heatmap of odor-evoked activity in OT$_{D2}$ neurons from day 1, day 3, and day 6 of imaging. The fluorescence measurements from each neuron were averaged over trials, Z-scored, then pooled for hierarchical clustering. Neurons are grouped by similarity, with the dendrogram shown on the right. Horizontal white lines demarcate the boundaries between the 6 clusters. Odor delivered at 0-2 seconds marked by vertical red lines. From left to right, the columns represent neural responses to sucrose-paired ketone and terpene, control ketone and terpene, and airpuff-paired ketone and terpene (S$_K$, S$_T$, X$_K$, X$_T$, P$_K$, P$_T$). Data is pooled from 6 animals.

Figure S4:

Figure S4. Pooled averaged-over-trials neural activity of all neurons from OT$_{D1}$ animals across days. Heatmap of odor-evoked activity in OT$_{D1}$ neurons from day 1, day 3, and day 6 of imaging.. The fluorescence measurements from each neuron were averaged over trials, Z-scored, then pooled for hierarchical clustering. Neurons are grouped by similarity, with the dendrogram shown on the right. Horizontal white lines demarcate the boundaries between the 6 clusters. Odor delivered at 0-2 seconds marked by vertical red lines. From left to right, the columns represent neural responses to sucrose-paired ketone and terpene, control ketone and terpene, and airpuff-paired ketone and terpene (S$_K$, S$_T$, X$_K$, X$_T$, P$_K$, P$_T$). Data is pooled from 6 animals.

Figure S5:

Figure S5. Pooled averaged-over-trials neural activity of all neurons from VP animals across days. Heatmap of odor-evoked activity in VP neurons from day 1, day 3, and day 6 of imaging. The fluorescence measurements from each neuron were averaged over trials, Z-scored, then pooled for hierarchical clustering. Neurons are grouped by similarity, with the dendrogram shown on the right. Horizontal white lines demarcate the boundaries between the 6 clusters. Odor delivered at 0-2 seconds marked by vertical red lines. From left to right, the columns represent neural responses to sucrose-paired ketone and terpene, control ketone and terpene, and airpuff-paired ketone and terpene (S_K, S_T, X_K, X_T, P_K, P_T). Data is pooled from 5 animals.

Figure S6:

Figure S6. Extended behavioral analysis from imaging period. (**A**) mPID voltage reading in response to 30 trials of odor (terpinene) delivery. The time period during which the odor valve was turned on is shown by the yellow rectangle. Individual recordings are shown in gray and the average is shown in black. (**B**) On-kinetics of odor delivery. On delay refers to the interval between the valve turning on and the mPID voltage increasing by more than 10% of baseline. (**C**) Off-kinetics of odor delivery. Off delay refers to the interval between the odor valve turning off and the mPID voltage decreasing by more than 10% of its maximum. (**D**) Representative velocity of the head-fixed mouse in response to the 6 different odors measured by a digital encoder. The lines represent the average across 30 trials and the shaded areas represent the SEM. The black arrowhead marks when the US is delivered. (**E**) The difference in walking velocity in response to odor (left) and US delivery (right). Differences are calculated between the last second before odor delivery and the last second before the odor exposure (left) or the first half second after US delivery (right) grouped by US pairing. Circles represent the average across animals and the error bars show SEM. (**F**) Representative changes in range-normalized eye-size in response to the 6 different odors. (**G**) The difference in eye-size in response to odor (left) and US delivery (right). Differences are calculated between the last second before odor delivery and the last second before the odor exposure (left) or the first half second after US delivery (right) grouped by US pairing. Circles represent the average across animals and the error bars show SEM.

Figure S7:

time to odor (s)

Figure S7. Traces of example neurons and their corresponding metrics. (**A**) Example traces from an OT_{D1} neuron recorded on day 3. Each column shows this neuron's response to a given odor across 30 trials (gray). The average across all trials is shown in black. For green and red arrowheads mark the time at which sucrose and airpuff were delivered, respectively. auROC values from single-neuron binary classifiers for discriminating {S_K vs. X_K}, {S_K vs. P_K}, and {S_K vs. S_T} are displayed on the right. Additionally, the results of statistical analysis to determine if this neuron reliably responded to each odor (in. = significant inhibitory response, exc. = significant excitatory response, ns= no significant difference between baseline and odor period). (**B**) Example traces from an OT_{D2} neuron recorded on day 1. (**C**) Example traces from a VP neuron recorded on day 3. (**D**) Example traces from a VP neuron recorded on day 6.

Figure S8:

Figure S8. Percentage of neurons responsive to each odor across days. Bar graphs showing percentage of neurons from OT_{D2} on day 1 that were significantly excited or inhibited by each odor and each US (S for sucrose, P for airpuff). The average is shown by the bar while the + and x markers show individual biological replicates. For each region and imaging day pair, 3 one-way ANOVA's were run to test for the effect of sucrose, airpuff, or ketone vs. terpene on percentage responsiveness. Only the significant F-statistics are reported with their corresponding p-value.

Figure S9:

Figure S9. Distribution of response magnitudes to each odor across days. Violin plots showing the averaged-over-trials response magnitudes to each odor during the last second of odor exposure. For each region and imaging day pair, 3 one-way ANOVA's were performed to test for the effect of sucrose, airpuff, or ketone vs. terpene on response magnitude. Only the significant F-statistics are reported with their corresponding p-value.

Figure S10:

Figure S10. Pairwise analysis of single neuron odor encoding.
(**A**) Scatterplot comparing the magnitudes of S_K responses ($\Delta\Delta S_K$) to P_T responses ($\Delta\Delta P_T$). The dotted line represents the hypothetical scenario where $\Delta\Delta S_K = \Delta\Delta P_T$. For each population, the R^2 value of the 2-d distribution compared to the $\Delta\Delta S_K = \Delta\Delta P_T$ line is reported. (**B**) The percentage of neurons from each population where the difference between $\Delta\Delta S_K$ and $\Delta\Delta P_T$ is lower than that between $\Delta\Delta S_K$ and $\Delta\Delta S_T$. (**C**) Bootstrapped FDR-adjusted p-values as a function of auROC values of single-neuron binary classifiers. In total, there are 27,900 single-neuron binary classifiers (15 pairwise classifiers for each of the 1860 recordings across 3 regions and days 1, 3 and 6 of imaging). Each classifier was compared against 10,000 shuffles. Horizontal magenta line marks FDR-adjusted p-value of 0.001 and the vertical magenta line marks auROC of 0.75. (**D**) The 25^{th}, 50^{th}, 75^{th}, 84^{th} and 95^{th} percentiles of auROC values and their corresponding unadjusted p-values. For auROC values that were greater than all 10,000 shuffles, a conservative p-value of 0.0001 was assigned. (**E**) The percentage of day 6 $\{S_K$ vs $P_K\}$, $\{S_K$ vs $X_K\}$, and $\{S_K$ vs $S_T\}$ auROC values greater than 0.75 as a function of time relative to odor, grouped by region. Lines represent the average across biological replicates and the shaded area shows the SEM. (**F**) Violin plot showing the distribution of pooled day 6 $\{S_K$ vs $X_K\}$ (left) and $\{S_K$ vs $S_T\}$ (right) auROC values grouped by region. Horizontal dotted line marks auROC = 0.75. (**G**) Heatmap of percentage of single-neuron pairwise classifiers with auROC>0.75. Classifiers were trained from neural activity recorded during the last second of odor exposure. Percent responsiveness was averaged across animals and grouped by region.

Figure S11:

Figure S11. Multinomial analysis of single neuron odor encoding (**A**) Confusion matrix of single-neuron MNR classifiers trained on neural activity during the last second of odor exposure on day 6 of imaging. Rows represent the true class while columns represent the predicted class. Each confusion matrix is averaged across 10 k-fold and across all neurons of a given region. ϕ represents data taken from 30 pre-odor bins randomly sampled from -1.5 to -0.5 seconds relative to odor delivery. (**B**) Violin plot of single-neuron MNR classifier accuracy, averaged across 10 k-fold, grouped by region. (**C**) Violin plot of the single-neuron MNR classifier accuracy trained on shuffled data. Each data point represents the average across 10 shuffles. (**D**) Violin plot of MNR S-cue confusion, i.e. confusion between S_K and S_T. This corresponds to 1) when the true class was S_K but predicted class was S_T and 2) when the true class was S_T but the predicted class was S_K. (**E**) Violin plot of MNR confusion among all ketones. This corresponds when the true class was a ketone and the predicted class was a different ketone (e.g. true class = X_K and predicted class = P_K). (**F**) Scatterplot of each neuron's ketone confusion on the x-axis and S-cue confusion on the y-axis on days 1, 3, and 6 of imaging. (**G**) Stacked bar \graph showing the distribution of neurons from each population that fall into each of the 4 quadrants across the 3 different imaging days.

Figure S12:

Figure S12. **Analysis of population-level odor encoding** (**A**) Scatterplot of CV-accuracy of linear classifiers trained on simultaneously-recorded neurons on the x-axis and their bootstrapped unadjusted p-values on the y-axis. Red horizontal line marks p = 0.001 and red red vertical line marks CV accuracy = 0.75. All classifiers with CV accuracy higher than 0.75 had p<0.001. In total, there are 765 binary classifiers (15 pairwise classifiers for each of the 51 recordings across 3 regions and days 1, 3, and 6 of imaging). Each classifier was compared against 10,000 shuffles. For auROC values that were greater than all 10,000 shuffles, a conservative p-value of 0.0001 was assigned. (**B**) The CV accuracy for $\{S_K \text{ vs } X_K\}$ binary classification trained on the last second of population-level activity. The bars represent the average across biological replicates. CV accuracy from individual animals are shown as x's. (**C**) Same as (B) but for $\{S_K \text{ vs. } P_T\}$ classification. (**D**) Same as (B) but for $\{S_K \text{ vs. } S_T\}$ classification. (**E**) Heatmap of CV accuracy from binary SVM's trained on day 6 of imaging with a radial basis function kernel. CV accuracy was averaged across biological replicates. (**F**) Confusion matrix of population-level MNR classifiers trained on neural activity during the last second of odor exposure on day 6 of imaging. Rows represent the true class while columns represent the predicted class. Each confusion matrix is averaged across biological replicates. ϕ represents data taken from 30 pre-odor bins randomly sampled from -1.5 to -0.5 seconds relative to odor delivery.

Figure S13:

Figure S13. Camera-based detection of licking in head-fixed animals (**A-C**) Metrics of a representative neurons with activity that predicts licking. (**A**) Representative neuron's Z-scored ΔF/F (dark green) and Z-scored dF/dt (light green) aligned to the onset of a lickbout. Each line represents the same neuron's activity during an individual lickbout. (**B**) Stemplot showing an example of the lagged correlation between the onset of licking and the fluorescence of a neuron across 9 frames (1 frame = 0.2s). Time bins of various lags are shown on the x-axis (negative number denotes frames that precede onset of licking) and the resulting R^2 is plotted on the y-axis. Red vertical line marks the case where the 9 frames are centered on the onset of licking. As an example, [-6:2] refers to fluorescence between 1.2 seconds prior to lickbout onset and 0.4 seconds after lickbout onset. (**C**) The output of a distributed lag model (DLM) that predicts the onset of a lickbout from ΔF/F of a single neuron. ΔF/F (dark green, top), DLM score (magenta, middle), and the licking recorded by a capacitive sensor (black, bottom) are shown in parallel. The DLM model was trained using 9 distributed frames ([-6:2]) of ΔF/F for each frame of lickbout onset. (**D**) Scatterplot of day 6 VP neurons' DLM lick classifier auROC on the y-axis plotted against their mean {S_K vs. X_K} auROC on the x-axis. The slope, intercept, R^2, and p-value of the slope are shown on the top left corner. (**E**) 3 example snapshots of the camera feed during moving lick spout paradigm with overlay of DeepLabCut labeling. The coordinates of top lip, bottom lip, base of tongue (tbase), and tip of tongue (ttip) are displayed with a probability cutoff of 0.4. (**F**) Range-normalized metrics from DLC labeling. P(ttip) (red, top) is the probability score assigned to the labeling of the tongue tip. P(tbase) (yellow, middle) is the probability score assigned to the labeling of the tongue base. Mgap (purple, bottom) is the Euclidean distance between the top lip and the bottom lip. The 3 vertical magenta lines represent the timing of the 3 snapshots shown in (E). (**G**) The same range-normalized metrics as in (F) plotted against the ground truth licking data from capacitive sensor (black, bottom) and DLC-based licking classifier score (magenta, second from bottom). (**H**) Lineplot showing the difference in total licking to L_{hi} nd N_{hi} during the time bin (1.5-2.5 seconds after odor onset) used for most analyses plotted against imaging day for individual animals. The time of peak difference is circled in black.