Second Writing

- Writing this after some ideas about handling orientation.

This system builds on the earlier one, which was limited to 1-d positions:

1. Each point on the sensor has a corresponding unit.
2. There are two points on the sensor receiving input. The corresponding units activate.
3. Propagation: Each time step, the activity moves one unit in a fixed direction along the array of units.
4. Single unit oscillations: When a unit first activates, it starts oscillating, such as 3 time steps in the oscillatory down state and 1 time step in the oscillatory up state.
5. Grid response: When a unit is active due to the propagation (or direct sensory input, which is just the first time step of the propagation) at the same time as it is in its oscillatory up state, it represents a relative distance between two points which is a multiple of a fixed distance.
6. With multiple sets of units, each with different oscillation frequencies, the exact distance between the two points is unambiguous.
7. To identify the relative position between a point sensed before a movement and a point sensed after a movement, the grid cells oscillate at a slightly different rate such that, if you were to extend their grid fields beyond the sensor's surface, the grid fields would still have grid fields on the same points for the entire course of the movement and after it. The propagating signal also speeds up or slows down to account for the movement.

That system is a foundation for determining allocentric locations in more than one dimensions. The exact number of dimensions cannot be fixed or limited to three dimensions because the brain might use more than three dimensions in some cases. For example, it might represent color as a location because colors are continuous attributes, not binary categories. Also, the same mechanisms probably apply to egocentric regions, not to mention mixed what/where regions, which might have a separate dimension for things like the angle of each joint.

One way to extend the 1-d system to more dimensions is to use an additional propagation and an additional single cell oscillation for each additional dimension. However, that introduces some potential problems and the neuroscience gets much more shaky. Also, with more than one dimension, rotations exist, which would be difficult to account for during movement since the pre-movement position in one dimension could influence the post-movement position in another dimension. Instead, it might be better to use separate units for each dimension.

◆

First Writing

- Didn’t update this doc with anything from 7/7/18-7/28/18. This section is from 5/31/18-7/7/18.
- This writing is to organize my thoughts for writing an article. I didn’t always check the facts.
- After this, mostly just stuff pasted over from a notepad I write ideas in.

Neocortical layer 5 is an enigma. As the cortical output to subcortical structures, it produces motor commands. However, the same cells that target subcortical structures often also project up the cortical hierarchy via the thalamus, suggesting a role in sensory processing. Exactly why the same cell can both generate behavior and contribute to sensory responses will clarify the mechanisms of sensorimotor integration.

One possibility is that layer 5 is actually just a motor layer. It projects up the hierarchy via thalamus in the same way as primary sensory thalamus relays sensory input, but nonprimary regions have other sources of sensory input, and every region probably needs input about motor output to understand the resulting sensory input. That would seem to resolve the apparent contradiction.

That explanation does not give the cortex enough credit. As a rule of thumb, we should expect the cortex to do a lot with a small number of characteristics and solve problems in unexpected ways. Otherwise, it would have hundreds of unique modules rather than the same six layers for all functions. We should expect cortical motor processing to utilize a large fraction of the same mechanisms as sensory processing or sensory and motor to be incorrect ways to categorize the functions of the cortex.

Layer 5 cannot only be a motor output. The shared circuitry of sensory and motor regions is the first clue, and sensory regions, even primary ones, have minor motor outputs. Despite the weakness of these motor outputs, however, not a single region has a particularly thin layer 5 compared to its other layers. Furthermore, layer 5 projects to subcortical sensory structures, such as S1 to the trigeminal principal nucleus or A1 to parts of the inferior colliculus, suggesting a role in sensory processing. Projections from cortex to subcortical sensory structures could simply convey a motor copy signal to account for the effects of behavior on sensory input, but that adds an additional constraint to a large list required for L5 to be purely motor. Additionally, if layer 5 were simply motor output, it would have to directly convert sensory input into behavior because layer 5 receives a direct sensory input from thalamus which evokes firing. That is not a broadly recognized input, so that thalamic input might only exist in some regions, but the fact that the input exists in any region contradicts a pure motor role. Even if there are alternative explanations to the directly driven response to the thalamocortical sensory input, as well as the apparently shared circuitry of sensory and motor regions, the prominence of layer 5 in sensory regions, and the projections from layer 5 to subcortical sensory structures, layer 5 being a purely motor layer would be too restrictive. Most brain regions connect to a large number of other brain regions, although often weakly, as if seeking any additional information or interpretation that helps whatsoever. Layer 5 is the only way for most subcortical regions to receive information from the cortex, as it is the only cortical projection to most subcortical structures, with the most notable exceptions being thalamus and possibly striatum, which also receive layer 6 and possibly L2/3, respectively. Thus, confining layer 5 to motor processing would prohibit subcortical sensory structures from making use of the massive amount of highly processed sensory information available in the cortex. A pure motor role for layer 5 would violate that principle.

Layer 5 is of course also a motor layer, not just a sensory layer, since it projects to the spinal cord and subcortical motor centers. That much is clear, but so many factors strongly indicate a sensory role that layer 5 must be a sensory layer. The apparent contradiction remains. How can the same cell produce a twitch of the shoulder and contribute towards recognizing features of the world?

Even though layer 5 has a sensory role, its characteristics seem in opposition to a sensory role, especially when compared to other layers. It is the only layer which projects to the vast majority of subcortical structures, implying that cortical sensory processing cannot be used in subcortical sensory processing. Furthermore, the cortical structures which receive from other layers are the thalamus and possibly the striatum, which notably also receives input from a class of layer 5 cell without other subcortical targets, and these two structures are repeated for various modalities so they are more closely related to the cortex than other subcortical structures. Other subcortical structures are not repeated for each modality. Thus, it would seem that these subcortical structures, which are unique to each modality, cannot utilize the type of sensory processing and representation utilized by the cortex because it is applicable to any modality. Since this cortical sensory processing appears incompatible with general subcortex and layer 5 is repeated for each modality just like the other layers, it seems that layer 5 cannot have a sensory role. Just as the code utilized by other layers appears incompatible with general subcortex, layer 5 appears to use a different code which seems more effective for encoding motor commands than encoding sensory items. L5 has generally high firing rates, in contrast to the sparse firing in other layers which is useful for coding complex discrete things but less useful for coding non-discrete values such as distance or force of movement. By being suited to code non-discrete things, L5 is suited to code behaviors since it can reuse most of the same representation for slightly different distances, forces, speeds, angles, and so forth. Layer 5 even fires when there is no sensory stimulation, suggesting its firing does not code sensory items. That is not to say that L5 cells always tick along firing at the same elevated firing rate when there is no sensory input to the region or behavior coming out of the region. Rather, L5 cells have noisy spontaneous input, which adds a degree of randomized firing. This random activity has been proposed to contribute to experimental behavior. Further exaggerating their firing patterns compared to other layers, L5 cells are by far the most bursty, which seems better for encoding a motor command, which occurs at a precise time, than sensory items, which occur for longer periods of time. Besides exaggerated firing, layer 5 cells have exaggerated dendritic trees with many synapses, like giant funnels out of the cortex. Perhaps as a result of their apparently low selectivity for synaptic input sources, they have low selectivity for sensory details and the positioning of those sensory details, as L5 cells have large receptive fields. These characteristics of L5 are shared with the deep layers of the superior colliculus, where neurons are more concerned with identifying potential saccadic targets than understanding those targets. So long as the feature is a saccadic target, these neurons in the superior colliculus respond to visual stimuli with little concern for sensory details and limited concern for position in visual space, helping them find targets. They form a bump of activity which averages out to drive a saccade to the target's exact location, suggesting that layer 5 cells have large receptive fields because a large number of L5 cells are required for a sufficiently strong signal to produce movements. Like the deep layers of the superior colliculus, it is likely that the movement generated by L5 reaches the exact position of the target in the receptive fields because the resulting motor target commands average out to that exact point. As a whole, it appears impossible for layer 5 to serve any notable role in sensory processing.

Yet, L5 must have a role in sensory processing, so it is worth reconsidering some assumptions about sensory processing. A widespread notion is that, to recognize something, the brain must recognize color, texture, pressure, and the like at each point on the sensor, and the spatial pattern of those points. As an additional part of this notion, since there are many ways to sense the same thing, the brain labels many instances of this type of spatial pattern as the same sensory item. This notion is flawed. Information like color describes a single point on the physical surface being sensed, not the sensory surface. The brain does not directly receive this information, of course. The physically arranged point descriptors on the physical surface project onto the sensory surface, producing a sensation which depends on how that physical surface is sensed. However, the brain should not directly group these practically random patterns by physical sensory item. Doing so would be difficult because there are a multitude of ways to view the same sensory item which produce completely different sensory inputs. Additionally, directly labelling the projection onto the sensor ignores the actual nature of the physical sensory item, so effective generalization is far more difficult. By first representing the point descriptors arranged on the physical surface, the cortex can produce similar representations of similar things. With that representation, the cortex need only recognize one particular arrangement of point descriptors on one particular surface shape in order to recognize the sensed feature.

Location is therefore crucial. It is generally accepted that coordinate transforms are important, such as for recognizing an object based on the relative locations of its features, but a coordinate transform is necessary even before the cortex can recognize the feature being sensed in a coordinate system with respect to the sensor. Before recognizing the contacted feature, the brain must convert the sensation of point descriptors arranged on the sensory surface into a sensation of point descriptors arranged on the surface being sensed. This coordinate transform is akin to warping the sensor to match the shape of the surface. In order to produce this effect, the cortex must treat surface shapes differently from the point descriptors arranged on those surfaces, as point descriptors add unneeded context to representations of surface shape. The cortex needs some sort of pure location processing.

Layer 5 is a good fit for this location processing, thus resolving the apparent contradiction between its roles in sensory processing and motor processing. Location is important for both functions, and the properties of L5 which might seem inconsistent with a role in sensory processing are useful for location processing. Weak selectivity is useful for processing location because the feature at each location does not matter for location. Likewise, large receptive fields are necessary for coordinate transforms, as a coordinate transform converts one location into another in different ways depending on the circumstances so each L5 cell must be able to respond to many locations on the sensor. High firing rates are useful for coding similar locations similarly, even if the firing rates themselves do not code anything, as the high firing rates reflect a membrane potential close to firing threshold and so a similar location is more likely to activate the same neuron. Likewise, the pseudorandom firing is useful for learning continuous variables like location because it tests various timings of various states and learns the pattern when it happens to get it right, evoking the same representation again when the pattern occurs again. By coding location, L5 is also a viable source of general subcortical projections because every subcortical structure, regardless of whether or not it specializes to a particular modality and regardless of whether it has a sensory or motor role, likely represents location in the same way. Unlike sensory details, which have different characteristics in different modalities such as vibrations indicating texture or on-off center receptive fields in the retina, location on the sensor is the same property for all modalities. In retrospect, it makes sense for layer 5 to perform sensory and motor functions because location is a key aspect of both brain functions.

The coordinate transform done by layer 5 produces a map of the sensed surface, but that only requires small shifts in locations. The map is similar because, effectively, both arrangements are topographical with respect to the cortical sheet. The sensory surface has a topographical correspondence with the cortical sheet, and adjacent points on a sensed surface are usually sensed by adjacent points on the sensor. Although nonadjacent points on a sensed surface can be adjacent on the retina, that is the same problem as separating overlapping objects in vision. The only difference between the two is that the former problem leads to incorrectly linking two points to the same feature, whereas the latter problem leads to incorrectly linking two features to the same object. Distinguishing nonadjacent points on the surface which are sensed with adjacent points on the sensor is a problem for all senses, not just vision. Multiple auditory objects can occur at the same time, and a small area on the skin can touch two objects at once or touch different objects over time. There must be mechanisms which solve this issue in every modality. Most of the time, adjacent points on the sensor correspond to adjacent points on the sensed surface, so this problem can be ignored for now.

Further simplifying the coordinate transform from sensory surface to sensed surface, there is no need to remove information about the sensory surface. Since this coordinate transform occurs before features are even recognized, the resulting map should not be allocentric, nor even account for different orientations of the sensor relative to the surface. Therefore, the impacts of the sensory surface’s shape (on the mapping from sensed surface to cortical receptive fields) are always the same. Layer 5 only has to add context about the sensed surface’s shape relative to the sensory surface, and each point on any given sensory surface will correspond to one point in the new coordinate system. At this early stage of location processing, the purpose is only to add back in the information lost by projecting the sensed surface onto the sensory surface. Since it adds context rather than removing it like egocentric to allocentric location transforms, this coordinate transform can be fairly straightforward.

The transform is not so simple that a direct conversion based on spatial patterns would work. At any moment in time, the sensed surface is projected onto the sensory surface, making the coordinate transform impossible without additional information. For example, upon contacting a surface with a fingertip, the skin distorts to effectively project the surface onto the sensor. Points on a bump will match up with points on the skin based on the two dimensional position on the surface, without regards for the height of that point. Although there is sufficient sensory information to determine the height of that point, such as based on pressure, that information is encoded as point descriptors, not as locations, so it is difficult to operate on. For vision and hearing, this problem is not obvious. The visual input projects onto the retina, but binocular vision helps, although it is not sufficient in some situations. The first iteration of the coordinate transform will only apply to vision in certain circumstances, but the second iteration is applicable because it [not complete, just groundwork right now] performs a transform to allocentric coordinates. In hearing, the location of the sound is resolved by using two sensors like in vision, but the map on the cortical sheet is of sound frequencies rather than locations. This suggests that the auditory cortex codes frequency in a similar way to how other regions encode location, with locations of sounds acting as context, like point descriptors in other senses. This is like graphing the frequency over time and treating that as a shape. Since the auditory shape is conveyed as a sequence of stimuli, only one dimension of the shape is present at any given moment, which is the frequency, whereas the other dimension, time, is conveyed over time.

The projection problem is the exact same for touch and hearing. Hearing clearly requires temporal processing, but the solution for touch is also temporal processing, and both operate on precise time scales. When the skin makes contact with a surface, it does so sequentially, making contact at initial points and propagating contact from those points. That spread of contact is analogous to shifts in sound frequency. Rough estimates based on slowed video suggest when only moving the finger, the fingertip makes full contact with another fingertip in around 40 milliseconds if moving very quickly. That seems to have a similar timescale to sounds and it is a reasonable timescale for neural computations. Since sensory input spreads on the sensory surface on neural timescales, the sensory input is reminiscent of propagating oscillations in the cortical sheet.

By utilizing propagating oscillations to process sequences of the sort described, the cortex can add information about surface shape to the neurons which receive direct sensory input in L5. When the first point of the sensor makes contact, it initiates a propagating oscillation which produces grid cell-like responses with respect to the sensory surface. When a point on the cortical sheet activates, interneurons activate in response to contact slightly afterwards, ending the response, and inhibiting pyramidal cells in adjacent points on the sheet to not respond to contact. Since interneurons target other interneurons, those in the adjacent points are inhibited as well, so the neighbors of the neighbors of the first point of contact are disinhibited, allowing them to respond to contact. The pattern continues from there, producing rings of responsive points around the initial point of contact. This periodicity may be useful for allocentric coordinate transforms. Additionally, the sequence of contact matters, with new separate points of contact acting as new sources of the pattern, colliding with other sources of the pattern, and the sequence of the shape of the spreading contact also changes the pattern of cells which respond. This coordinate transform also [hypothetically] produces similar patterns for similar surface shapes, improving generalization. It is applicable to hearing, not just touch, and it is applicable to vision during gaze tracking, such as during head movements associated with saccades. Although it is likely insufficient, this coordinate transform might help produce responses which are invariant to the orientation of the sensor relative to the surface because it operates on local scales and the localized sequence of contacts is similar for different orientations of the fingertip. However, this was not tested and is meant as a starting point for describing a coordinate transform to allocentric space.

Layer 5 is a good fit for the described coordinate transform. It has been proposed to handle precise timing both for behavior and sensory processing, which is important for recognizing surface shape in this manner. It also explains an apparent contradiction. Layer 5 processes location, which seems mostly spatial, but it generates behavior, which causes a change in location but is mostly sequential. The mechanisms are unknown at this point in this paper, but there seems to be some shared mechanisms for behavior and location processing. This coordinate transform also explains the various aspects of layer 5 which imply a sensory role, such as direct thalamic input, and those which imply a motor role, such as weakly selective receptive fields. The mechanisms of this coordinate transform are utilized in the second iteration.

Oscillations are thought to play a key role in coordinate transforms. They can produce periodic responses like grid fields, which can be allocentric. It is easier to add information than remove it, and periodic responses reduce the amount of information that must be removed from egocentric representations. There are numerous egocentric representations of the same allocentric object because there are numerous ways to contact it, so directly grouping them into allocentric objects would be difficult or inefficient. Instead, the solution may be to respond to periodic egocentric locations, shifting these locations reasonably small distances as necessary.

The first step of this coordinate transform is a sequential shifting of receptive fields. A stimulus is in some large overlapping receptive fields, providing input to neighboring cells with thalamic input for a map which is topographical to the sensor. The size of the receptive fields determines the extent of the shift towards allocentric representation, as the cells cannot respond to anything outside their large receptive fields. Those receptive fields are only subthreshold without additional input, producing smaller, shiftable receptive fields; hierarchy would allow shifting shifted receptive fields to enhance allocentricity up the hierarchy. Until other mechanisms are added, the cells respond to a stimulus anywhere in their large receptive fields, but not all of them respond at the same time. Each cell's small receptive field travels to eventually cover the entire large receptive field, producing sequential responses. To do so, a low frequency oscillation is used. There is a gradual change in oscillatory phase, like a propagating wave although it is more likely subcortically generated than intracortically. As a result, an oscillatory peak travels across the cells. Since those cells are organized with receptive field centers topographical to the sensory surface, and the sensory input is stronger close to the center of the cell's receptive field, each cell shifts its smaller receptive field across the extent of its larger receptive field. [I still need to do some thinking about this but I'm not concerned with it right now because it's a straightforward effect].

So far, since the smaller receptive field is suprathreshold but not the larger receptive field, the response is like moving the sensor across the surface. At some point, disregarding the limits to the wide receptive fields, all cells will respond to a stimulus, and they do so sequentially. The initial position of the oscillatory peak does not matter because another will come anyway, but it will change the sequence of the response. That exact sequence is not important, although the fact that cells respond sequentially will be important. All that matters for now is which cells fire how many times[move this] [using the output layer removing states with each input, so all that is left is a representation of the set of cells which fired enough; details are murky on whether to represent the combo of every round of or just take the combo of cells which fired in the top x% times]. With additional mechanisms, when cells fire multiple times, they have found a relative location between two features in their large receptive fields.

A second oscillation adds one more level to the receptive fields to produce grid fields.It has a much faster frequency and smaller wavelength on the cortical sheet than the low frequency oscillation, and it is responsible for creating grid fields. Like the low frequency oscillation, the high frequency oscillation has a travelling peak, but because it is always aligned with the most recent spike. This peak is aligned with the low frequency peak because that is when cells can fire, meaning one of the high frequency peaks occurs at the same time as the low frequency peak each cycle, so it is aligned with the small near-threshold receptive fields, perhaps with an individual oscillation for each cell. When the cell's sliding receptive field encounters a stimulus at the same time as the cell's high frequency oscillation peaks, the cell fires. When a cell fires two low frequency oscillations in a row, it has found a relative location with a certain multiple of a distance away, producing a grid field. If it fires three times in a row, it has started to find an allocentric object. These mechanisms implement a two step process for one dimensional allocentric location. First, a receptive field sweeps the area. Then, after finding something, it tries to find something a multiple of a certain relative distance away.

The exact same oscillations control behavior, thereby accounting for movements of the sensor by behavior. If, for example, the map is one dimensional and the behavior moves the sensor to the left, the receptive fields of the grid cells must shift to the right. To do so, the cortex generates a movement to the left by increasing the frequency of the low frequency oscillation

during the movement. At the same rate as the sensor moves, the smaller grid cell receptive fields move, so they stick on the same place. However, this is only sufficient if the moved distance is a multiple of the grid cell wavelength. Therefore, the low frequency oscillation is shifted by the exact movement, but the single cell oscillations control precise distances of movement by speeding up or slowing down. This occurs indirectly. The cells firing at elevated rates is what generates the movement, so it always occurs a period of time aligned with the grid cell cyclic pattern. Through precise timing mechanisms, they change firing rates slightly on top of undergoing more rapid low frequency oscillations. This determines the precise distance to move because around when the low frequency oscillation returns to normal frequency, the time when they stop firing a little more rapidly than usual determines the distance moved. In the process, they shift their grid fields slightly to keep it aligned with their grid fields before the movement. Thus, behavior is generated by correcting for the resulting change in position.
-mechanisms for “behavior is generated by correcting for the resulting change in position” allow separation of firing rate (more exactly, FR during oscillatory peak) and behavior. So they can respond directly to thalamus.

-need to write/think about how during sequences of more than two relative locations, the fast oscillations are given consistent relative offsets to produces consistent grid fields/maybe path integration.

Other mechanisms likely contribute because this mechanism is fragile, but it will be useful for dealing with non-simultaneously contacted relative features. The second oscillation is faster than the first, and, unlike the low frequency oscillation, the high frequency oscillation does not have a travelling oscillatory peak. Instead, each cell has a different phase. The potentially fragile but later useful part is, these high frequency oscillations are generated by single-cell intrinsic mechanisms and aligned with most recent time the cell fired [this might be by
-just use organization of map -> firing time for this.

-the fast oscillation is initially centered on the position of the first contact but that doesn't matter because relative positions are still the same. Altho not sure about path integration.

-maybe for object ID + explaining why those that fire e.g. twice are the ones considered active after the xform, use disambiguation mechanisms.

It is easiest to first consider allocentric representations of simultaneous contact with two features on the same rigid, joint-free sensory surface.

To produce an allocentric representation, consider all possible relative locations between two points with a fixed distance between them. That may seem to require massive amounts of neural processing, and it might if done in a single instance, but it is much easier to test one at a time and then use the set of cells which fired as the resulting state.

Essentially, it shifts the responsive cell (or in other words, RFs) to cover all allocentric positions, although a complete allocentric transform requires hierarchy. This is by oscillating across the surface selecting cells with overlapping RFs sequentially, shifting them.

An oscillation, subcortically generating and with a peak which travels over the sheet, tests each by selecting subsections of large receptive fields with encompass the features, essentially shifting smaller receptive fields over time to test different relative locations. First cell initiates then it will respond to other feature if mod distance away.

[scrap this? The proposed mechanisms to produce a map of the sensed surface rely on the temporal properties of layer 5 receptive fields in barrel cortex. L5 cells receive direct thalamic

input, but only from one corresponding whisker. The rest of the suprathreshold receptive field is driven intracortically or via the loop between cortex and thalamus, and it extends out several whiskers. Further from the principal whisker, the response latency increases. These response properties probably result from propagating waves of activity. When a whisker makes contact, the thalamus drives activity in the corresponding barrel column. The activated neurons in turn activate interneurons, which inhibit neighboring neurons. Because some of those neighboring]
-talk about this about whiskers, not just fingertip and vision because that's the evidence for the mechanisms.

[It is easy to think of sensory processing like converting a picture of a feature into a label for that feature. However, sensory processing is not so straightforward, even leaving aside the problems of generalization and the multitude of ways of sensing the same thing. Sensory input is sequential. This property is most apparent in hearing because sounds are highly sequential, but every sense is sequential. In touch, fingertips often brush against surfaces rather than simply poking a surface and producing a static sensation. Even in vision, sensation is fundamentally sequential. As one saccades around an object, a sequence of images appear on the retina, not to mention the eye floaters that pass by every time the eye moves. Besides the auditory ones, the sequences mentioned so far are generated by behavior, so they are not informative about the world and the effects should get filtered out. At least, that is the common assumption when considering how the brain accounts for the results of behavior.]

[Behaviorally generated sequences of sensations are are actually *more* useful than static sensory inputs during active sensation. Even during active sensation, behavior is often considered an extra factor for which the brain must account in order to recognize features. While it is true that behavior can cause many different sensations for the same feature, and one purpose of sensory processing is to represent the same feature in the same way regardless of additional factors, that does not mean]

[For example, when contacting a surface with the skin, the sequence in which each point on the skin first makes contact depends on the shape of the surface feature. Although a static sensation after the skin sequentially makes contact also indicates the shape of the surface feature because of receptors like pressure sensors, the sensation encodes the feature differently. For a static

sensation, the entire pattern of pressures on the skin matters. To recognize a feature requires considering the entire pattern at once. In contrast, for a sequential sensation, which points on the skin just made contact is sufficient to recognize the feature when understood as a sequence. The simplest type of sensory input is whether or not a point on the sensor is receiving sensory stimulation; information about the stimulation at that point adds more complex information.]
[Even though most objects modern humans interact with are solid, this type of brief sequence is still useful. For example, the sequence of contact made with a surface depends on the shape of that surface. The importance of this kind of sequence is least clear for vision, and sequences might not be as important for vision as for touch or especially hearing, but sequential input still exists. For example, whenever the head moves, including during the period after a saccade, the eyes remain fixated but the sensory input nevertheless changes sequentially. For every sense, sequences are highly informative.]

[This type of sequence is not like a melody, but rather, it is more like a brief sequence during the onset of sensory stimulation. Although one could include sequences that last for seconds, briefer sequences are easier to recognize because neural timescales are around a fiftieth of a second, and that is a similar timescale to making contact with a surface.]
[is extremely brief, perhaps sometimes less than ten or twenty milliseconds such as when the fingertip rapidly pokes a surface, although this type of sequence can last . Layer 5 likely serves this role. It does not need to detect the information on the sensory surface, only which sensory surfaces are stimulated. Even if it did detect that, it would still appear as if it were not particularly selective because there are multiple sensory stimuli over the course of the sensory onset sequence. It would even appear that way to other neurons since the sequences are often brief enough for EPSPs, or at least NMDA spikes, to last as long as the sequences, so little is to be gained by recognizing things like texture or color in layer 5, and that would just cause confusion in some cases, so it is best left to other layers.]

[However, layer 5 does not seem to process sensory information in the same way as other layers. (only cortical motor output, high FRs, noisy firing, bursting, and nonselective responses. There are roles besides motor and sensory, e.g. attention).]

- The corticostriatal pathway is the real motor output, so there is no mystery why L5 TT projects to both motor and sensory subcortical structures such as the thalamus and superficial superior colliculus.
- L5 TT projects to the striatum, but since it projects to so many subcortical structures and projections to striatum are highly variant between species, that projection serves the same function as the projections to other subcortical structures.
- L5 TT doesn’t receive particularly strong driving thalamic input in V1 because, whereas in touch the sensory shape is clear based on the sensory input, it is not clear for vision until other layers divide objects.
- Motor L5 -> thalamus must be corollary discharge because otherwise, not all regions receive corollary discharge.
- Just as corollary discharge filters irrelevant results of behavior, the same signal from sensory L5 filters irrelevant sensory input, since higher order thalamus is involved in attention.
- L5 projections to subcortical sensory structures are predictive because L5 in at least some regions has longer latency than all subcortical sensory structures, so it must change the sensory response before L5 itself responds.
- L5 projections to subcortical sensory structures are not for attention. L6 already fits attention, and that is too large a limitation on sensory signals from cortex to subcortex.
- Since L5 TT cells are the only source of subcortical projections except to thalamus and striatum, they have some unique characteristic which makes them the only cells fit to do so. They cannot simply be the best fit for subcortical projections, or else other layers would project, even just rarely and sparsely, to provide additional information.

Misc

- To combine knowledge of different cortical columns, they just have to vote together.
- Columns can start voting together in multiple ways. They can start projecting to other cortical columns for voting over the course of development, or already existing connections between columns can learn voting over the course of learning about something in particular. They can also start voting together, in a sense, by their signals travelling up the hierarchy to regions with progressively greater degrees of columnar voting.
- There are multiple possible types of independent perspectives of the world a cortical column can have: not voting together with columns for a distant sensory patch, not voting together with columns for a different part of the region's map on the cortical sheet of object ID, restricted to small features rather than the whole object or objects rather than constellations of objects, or whatever other form of different or non-conjoined perspectives of the world.
- By starting with independent perspectives of the world and starting to vote together, columns become intelligent. By narrowing down possibilities with multiple ways of considering something, the possibilities left behind are generalized, invariant, and capable of concept representation.
- Since the EC sheet map is of grid cell scales rather than places, travelling up the hierarchy, the cortex produces location invariance by changing from a map of space to a map of periodicity scales.
- It is hard to pool a bunch of egocentric locations into the proper allocentric location, but easier to pool periodic egocentric locations into a grid cell.

Research-Oriented

- Since oscillations arise from the grid cell mechanism, some mechanisms and/or evidence utilized by oscillation-based grid cell models are applicable.

——

- The first point of contact matters, but that response is obscured by longer latency responses to more distance points in the surround RF because it is unclear whether the response is to the PW or to a SW. The prominent bursting and subsequent plateau-enhanced RS resolves this confusion.
- Burst -> plateau potential/RS helps detect onset/onset details and then continue responding depending on that.
- Even if bursting is not real, the briefly-sustained plateau potential clearly is real because of the prominent calcium spike initiation zone.
- Low prominence of direct lateral connectivity is because slice cuts them.

——

- The 50 ms L5 EPSPs following sensory input are actually calcium spikes.
- Although vS1 cells aren't very selective for deflection direction, they are selective for deflection amplitude. Therefore, they are selective for deflection azimuthal direction and/or deflection duration, since that correlates with deflection amplitude.
- As the fingertip makes contact, horizontal inputs arrive at a similar time as input from the fingertip. However, in vision, input is constant, which either masks primary thalamic input or makes it unnecessary

Grid Cells

Sensory Onset

- Bursts are for detecting sensory initiation events
- L5 ST and TT are both for sensory onset processing because they are the two most bursty pyramidal cell types.
- L5 TT is more directly concerned with sensory onset events because it has more prominent burstiness and shorter latency sensory responses.

——

- Since there is a map of objects in high allocentric regions, high level/broad concepts arise from representing objects like locations. Therefore, the point of hierarchy is to build up location-like abstraction where the location-like representation includes sensory information on a concept level.

——

Broad

- L5 ST and/or TT is for some sort of map-level function, not sensory contents processing.
- L5 is for sensory initiation events, as well as things associated with initiation events, such as contact order.
- The duel input/output layers are because of the distinction between sensory onset dynamics and sensory sustain dynamics.
- The corticostriatal pathway is the real motor output, so there is no mystery why L5 TT projects to both motor and sensory subcortical structures such as the thalamus and superficial superior colliculus.
- L5 TT subcortical outputs to motor and nonmotor structures serve the same function They fine tune responses by precise timing.
- Precise timing is a non-sensory contents processing roles of L5.
- L5 is sensitive to onset properties and has sustained responses (such as from apical metabotropic input), which might seem contradictory.
- However, sensory onset information must be sustained longer than the onset period, to utilize alongside sensory sustain information, to utilize alongside later sensory onset events, and to utilize alongside later periods of the overall sensory onset on different parts of the cortical sheet.

——

- If oscillations are how the brain accounts for speed, oscillations are a good fit for how the brain determines self-movement speed.

——

- L5 TT projects to the striatum, but since it projects to so many subcortical structures and projections to striatum are highly variant between species, that projection serves the same function as the projections to other subcortical structures.

——

- L5 TT doesn't receive particularly strong driving thalamic input in V1 because, whereas in touch the sensory shape is clear based on the sensory input, it is not clear for vision until other layers divide objects.
- L5 TT isn't really less selective than other layers for sensory details, but it appears so because during sensory onset there are many stimuli so cells are more likely to receive one which fits.
- To sustain information about onset dynamics, POm -> mGluRs. NMDARs are insufficient for sustained working memory-like responses.
- Rather than directly generating firing, the mGluRs change responses to later sensory input.
- To utilize the idea of objects, lower regions do not need to recognize an object. They just need to respond differently to the sensory input depending on the object. In fact, they should not represent the object, because the working memory may no longer apply.
- Motor L5 -> thalamus must be corollary discharge because otherwise, not all regions receive corollary discharge.
- Just because L5 projects to subcortical sensory structures doesn't mean it has a role in sensory processing, because the signal could be for things like coordinate transforms or other sensory processing which requires corollary discharge.
- Alternatively, just as corollary discharge filters irrelevant results of behavior, the same signal from sensory L5 filters irrelevant sensory input, since higher order thalamus is involved in attention.
- Since L5 ST receives proximal POm, its role is linked to the otherwise modulatory signals sent by POm.
- Corollary discharge predicts the result in sensory cortex. The oscillation rate or firing frequency of that corollary discharge indicates timing of each item in the motor-generated sequence because oscillation rate reflects movement speed.
- That is the signal sent by motor L5 TT. Therefore, sensory L5 TT has a similar role. It is for precise timing of sensory onset events.

- Self-movement produced sound filtering is temporally precise, so it is not filtered if it occurs at an unexpected time.
- The same mechanism contributes to disentangling sounds. Likewise, the cortex needs to disentable objects, whether represented simultaneously such as when exploring or utilizing multiple objects at once, or in when objects overlap like in vision or simultaneously touching multiple objects. It also has to disentangle possible object representations, which can be a massive union especially in higher regions.
- Precise timing and mechanisms for retaining onset information allow disentangling multiple objects.
- Mistiming of behavior is involved in motor learning. The same mechanisms are utilized for learning the precise timing of sensory onset.

——

- Never represent the entire object because that is not useful for higher regions. Rather, represent the place on the object, whether as the object in context of position on it or as a location signal in a coordinate system unique to the object.

——

- Since hierarchy is for egocentric scale, each region needs to represent multiple abstraction levels of a single object, and it needs to use a gradient of abstraction levels rather than discrete levels of abstraction.
- This gradient is implemented by incomplete disambiguation. Each feature removes some cells which represent objects that do not include that feature at that location, but leaves some behind. That way, some cells represent the whole combination of features as the exact object, whereas others do not care about some parts of the object.

——

- Onset dynamics processing or the mechanisms to do so must apply to higher levels of abstraction.
- L5 projections to subcortical sensory structures are predictive because L5 in at least some regions has longer latency than all subcortical sensory structures, so it must change the sensory response before L5 itself responds.

- L5 projections to subcortical sensory structures are not for attention. L6 already fits attention, and that is too large a limitation on sensory signals from cortex to subcortex.
- Since L5 TT cells are the only source of subcortical projections except to thalamus and striatum, they have some unique characteristic which makes them the only cells fit to do so. They cannot simply be the best fit for subcortical projections, or else other layers would project, even just rarely and sparsely, to provide additional information.

——

- Higher order thalamic cells represent objects.
- Wide range input to L1 causes synaptic depression for highly distinct sensory inputs. That prevents bursting for later sensory inputs that are part of the same object represented by the higher order thalamic cell, contributing to disambiguation.

- Related to L5 detecting contact order/sensory onset:
- Horizontal connectivity leads to longer latencies when further from the topographically aligned position of the sensory input on the cortical sheet. That propagating signal oscillates as it travels because of grid/mod-like inhibition, forming grid-like fields.
- The horizontal connectivity does not just propagate the sensory input. It also serves other functions. For example, sequence memory arises from incomplete propagation-driving excitatory connectivity because whether or not the propagation reaches a cell depends on which cells are active, and this effect chains.
- Cells have grid fields in terms of whether they can fire at all, and the set of cells which fire is further narrowed by sequence.
- The grid cell mechanism results in a simple response when there is only a single point of brief sensory input, but when the sensory input is a sequence of contacts, the response

gets more complex. Perhaps this and/or the sequence memory aspect of the propagating signal lead to path integration.

——

- The grid cell mechanism doesn’t have to depend on contact speed. The first point of contact inhibits adjacent points on the cortical sheet which therefore do not inhibit their adjacent points, and this chains. Contact order matters, but not contact speed.

——

- The grid cell mechanism operates on the level of bursting, not singlets, and uses martinotti cells as the inhibitory signal to nearby points on the cortical sheet.

——

- The grid cell mechanism leads to a grid scale gradient starting from the initial point of contact and expanding outwards on the cortical sheet, because of less excitatory drive to cells with center RFs more distant from the stimulus.

——

——

- Since oscillations in vS1 synchronize with whisking, each cell fires best once per whisking cycle, so it is selective to whisking phase. As the oscillation propagates, it moves activity along the cortical map, moving responsive cells, while meanwhile whisking phase progresses. Therefore, at some scale, the cortical sheet must represent something besides whisking phase, such as whisking space.
- Since the first contact likely initiates the propagating oscillation, this suggests a means of recognizing relative positions of features in whisking phase space. The cells at the initial point of contact are selective for the whisking phase at that time of contact, and the rest of the cortical sheet forms a map of whisking phase which extends from that initial point.
- Since propagating waves are generated by inhibitory-disinhibitory dynamics, the wavelength of the wave is fixed. In other words, since the activity propagates by a cycle of inhibition and disinhibition along the cortical sheet, and since inhibitory connections (both those that inhibit pyramidal cells and those that inhibit other inhibitory cells) have

fixed axon arbor radii, a single inhibition-disinhibition cycle occurs over a fixed distance along the cortical sheet, with at most only a little flexibility.

- Since oscillation is generated by the same inhibitory-disinhibitory dynamics which mediate propagating waves, with excitatory lateral connections mostly being distal synapses and therefore not driving, oscillation and propagation are the same things. Therefore, just as the propagating wave wavelength is fixed, the oscillation wavelength is also fixed.
- Since oscillation frequency is the same as whisking frequency, as a result of the fixed propagating oscillation wavelength, the propagation travels faster along the cortical sheet with greater whisking frequency. Therefore, the propagation travels faster with faster whisking speed.
- Since the oscillation is synchronized with the whisking cycle, the propagation must travel faster with faster whisking speed in a manner such that each point on the cortical sheet corresponds to a single place in whisked space. Therefore, there is a map of scanned space on the cortical sheet, and this map arises from propagating oscillation.
- From there, using oscillatory interference, grid fields emerge.
- To determine the coordinate system, different oscillating or oscillation-generating sources of corollary discharge are used. For example, a map of whisking angle would require only a signal from motor cortex cells which project to the whisking central pattern generator. To add a map of whisking angle up/down requires a signal from motor cortex cells which project to motor centers which control muscles for vertical whisker angle.

——

- Even though L5 TT doesn’t project to L5 ST, their proximity suggests shared inhibitory cells, leading to linked oscillations, linked competition, and linked minicolumnar states.

——

- L5 ST enhances TT bursting/subsequence plateau potentials via synapses on the tuft.
- Therefore, the grid cell mechanism might need two separate layers with similar roles (implied by burstiness).

- For example, one layer must not retain info about sensory onset dynamics because there needs to be a pure response to later sensory onset dynamics.
- L5 TT responds with short latency to sensory onset dynamics, and that response depends on prior sensory onset dynamics. Then, L5 ST responds purely to the current sensory onset dynamics, and the signal to L5 TT only matters if L5 TT could not understand the new sensory onset in light of prior sensory onsets, because there is not much impact if cells have already bursted and thereby generated apical inhibition by martinotti cells. [relates to martinotti cells being for burst SM not competition].
- Separating responses, into pure sensory onset dynamics and sensory onset in light of previous sensory onsets, contributes to interactions between sequential sensory onsets and makes them less chaotic.
- Later: Origins of cortical layer V surround receptive fields in the rat barrel cortex claims that L5a response latency (but not L5b) increases between the 1st and 3rd whisking cycles. So maybe L5 ST is serving that function, but also increases latency after understanding the stimulus so it aids L5 TT with the pure sensory onset info and then by increasing latency, aids only when L5 TT needs it.

——

- L5 vS1 cells have short principal whisker latencies, but also longer and longer latency responses to more and more distant surround whiskers. Therefore, they fit the longer latency responses to sensory stimuli further from the cell’s RF center in the grid cell mechanism.

——

-

——

——

-

——

——

——

-

——

——

——

——

——

——

——

——

- Since bursting is for onset dynamics detection and apical input enables bursting, the apical dendrite has a role in onset dynamics detection. The inputs to apical dendrites are topographically wide or not topographical at all, whereas onset dynamics are topographically narrow. Likewise, plateau potentials are long lasting while proximal input is brief.
- POm mGluR increases the ability to burst, so L5 TT cells process sensory onset in light of earlier sensory onsets.

——

- If L5 TT understands the sensory onset in light of earlier sensory onsets, it sends a signal through POm to ST cells, so they also process sensory onset in light of earlier sensory onsets. However, if the signal is not sent because L5 TT does not understand, L5 ST processes the pure sensory onset.

——

- The propagating m-type activity proposed by JH for precise timing results from propagating L5 CT activity.

——

- L5 fine-tunes things by predicting them/predictive firing.

——

- L5 TT detects onset dynamics:
- 1. Initial input causes L5 singlets, which evoke thalamic spikes.
- 2. Thalamus sends a signal to some L5 apical dendrites.
- 3. Sensory-evoked inhibition begins. L5 cells without apical input mostly cannot fire in all later steps.
- 4. One precise time step after sensory onset, cells with paired apical input and proximal sensory input respond by firing more rapidly.
- 5. Overcoming synaptic depression, the elevated firing rates are sufficient to evoke thalamic firing.
- 6. The elevated firing cells have fired several spikes rapidly. This activates martinotti cells which inhibit cells in that area of the cortical sheet. Intrinsic mechanisms further contribute to ending the elevated firing.
- 7. The process repeats.
- 8. After sensory onset, the L5 cells which have at some point had elevated firing rates have reduced synaptic input thresholds.
- Even without sensory input for brief periods, the activity still propagates because no sensory input means no sensory-evoked inhibition so lateral connections are effective, allowing precise timing during periods without sensory onsets. This requires an understanding of other layers.
- In motor areas, the thalamic input is about behavior, or at least partially about behavior. This feedback loop implements precise timing for behavior as well.

——

- Since bursts last longer than the onset event and longer than cells firing singlets, they associate to postsynaptic neurons which tend to fire afterwards, so they predict them.
- They do not associate postsynaptically to cells which fire beforehand because of temporally asymmetric STDP.
- Even when they do not burst, they still predict the target.

- Normally, STDP causes LTP when the cells fire at the same time. However, because the burst extends for a couple spikes afterwards, this leads to LTD. Furthermore,
- When in singlet mode, they associate postsynaptically when coincident, driving the target to fire longer than usual when bursting. This does not occur if they only associate postsynaptically when in burst mode, which depends on the target cell type.
- By predictively depolarizing cells, L5 TT biases subcortical targets including thalamus.
- Before sensory input-induced inhibition, L5 TT cells outside the current space of the sensory input on the sensor respond with singlets to the apical input from thalamus paired with proximal inputs from other L5 TT cells. The proximal inputs are predictive, but not the apical inputs because those inputs are not bursts and the learning rules are different. When the sensory input is unknown or too simple to narrow down the object pooler much, this manifests as larger RFs.
- The L5 TT predictive signal is not predictive in the same sense as in sequence memory because that occurs 50-100 ms ahead of time, whereas the burst predicts <30 ms ahead of time, brief enough that the last spike of burst itself is influenced by the predicted event. This allows predictions to lead to predictions while still constraining those predictions by what happens.
- For some targets, rather than causing predictive firing per se, the predictive input reduces response latency to sensory input, which is predictive relative to other cells.
- By sending a predictive signal, it causes pooling of sensory concepts on pseudoproximal segments, which are not SM-type predictive, as they lack NMDA spikes, but do not directly cause firing. They serve to identify simultaneous patterns which cause the cell to fire but are still independent patterns which do not have to co-occur.
- By the ~predictive firing, L5 TT cells signal multiple patterns simultaneously because they are predictions of multiple possibilities at multiple times. If the postsynaptic cell's represented pattern can occur after the presynaptic cell a few tens of ms later, the presynaptic cell predicts the postsynaptic cell, although not the same type of prediction as SM because it is anchored to sensory input.

- Even if the burst occurs immediately before the postsynaptic cell fires, the burst has more spikes afterwards, so it dissociates. Thus, predictions are of a window starting a bit afterwa

- Because of the duration of a burst, even if the burst occurs immediately before the postsynaptic spike, the burst will have more spikes after the postsynaptic singlet, so it will still dissociate. This means there is a window starting and also ending a little bit before the postsynaptic spike in which the presynaptic spike must occur for LTP to occur. Also, temporal summation produces a stronger postsynaptic depolarization towards spiking if more spikes in the burst occur before the driving input. This produces grid-like sensory onset sequence tracking. Since L5 TT cells have fairly unselective RFs, they can often respond to the center RF regardless of the surface information.
-
- The burst is sort of the opposite of prediction. Rather than predicting first and then trying to match what occurs afterwards, it sees what occurs first and then tries to predict what occurs afterwards, which is whether or not the cell bursts. Whether or not the cell bursts is decided after the first spike because of the bAP delay and because of how the precise timing loop works. Therefore, it helps decide whether to make the prediction by inputs to the apical dendrite. Since the calcium signal which evokes the burst is far larger than that evoked by a bAP, it essentially ignores the first spike in a burst (and singlets) for STDP at the distal apical dendrite and tuft, although the first spike still determines whether or not it can burst. Inputs to the apical tuft (non-primary for TT) produce NMDA spikes and therefore predict whether or not it will burst like SM prediction.
-
-
-

——

-

——

-

——

-

——

-

From notepad called "Hypothesis"

Layer 5 TT detects onset dynamics.

1. Initial input causes L5 singlets. Since the synapses are not depressed, they evoke thalamic spikes.
2. Thalamus sends a signal to some L5 apical dendrites.
3. Sensory-evoked inhibition begins. L5 cells without apical input cannot fire.
4. L5 cells with apical input respond with bursting to sensory input one precise timing time step after sensory onset begins.
5. Despite synaptic depression, the bursts are sufficient to evoke thalamic firing. They are not always obvious bursts, but simply elevations in firing rate for which synaptic depression has not yet accounted.
6. Since the cells fire for a longer period and more rapidly than those that are inhibited, they activate martinotti cells which inhibit cells in that area of the cortical sheet. Cells which happen to get out of control because of random potential synapses are inhibited by intrinsic mechanisms.
6. The process repeats.
7. By the end of the sensory onset, a set of L5 cells have bursted or at least had apical Ca2+ spikes or metabotropic events. These cells have enhanced firing rates for a little while.

Even without sensory input for brief periods, the activity still propagates because no sensory input means no sensory-evoked inhibition. This requires an understanding of other layers.

In motor areas, the thalamic input is about behavior, or at least partially about behavior. This feedback loop implements precise timing for behavior as well.

Sinces pseudobursts last ~30 ms even if proximal input ends, they associate to postsynaptic neurons which tend to fire afterwards. They would also associate to those which tend to fire beforehand, but asymmetric learning rules cause more LTD than LTP. Thus, they predictively depolarize postsynaptic targets, regardless of whether or not they (presynaptic cells) burst. They do not associate well to targets which fire at the same time because of the aforementioned temporally asymmetric learning rules, when in burst mode. When in singlet mode, it associates when coincident, and so drives the target to fire longer than usual when the population targetting that neuron is in burst mode.

By predictively depolarizing cells, L5 TT modulates or biases some of its subcortical targets, including thalamus. This is modulatory when it is insufficient to trigger firing on its own but contributes to firing evoked by slightly later additional depolarization.

Because there isn't yet sensory input-induced inhibition, L5 TT cells outside the current space of the sensory input on the sensor respond with single spikes to the apical input from thalamus + cortical inputs, including from other L5 TT cells, predictively because proximal input is necessary for associating to lateral connections, and because of the predictive results of asymmetric learning rules + bursting mentioned before. When the sensory input is unknown or too simple to narrow down the object pooler much, this manifests as larger receptive fields.

The L5 TT predictive signal is not predictive in the exact same sense as the predictively depolarized state in sequence memory because that occurs ~50-100 ms ahead of time, where the predictive firing of L5 TT occurs <30 ms ahead of time, brief enough that as a burst it is influenced by the predicted event itself towards the end of its predictive response. This allows predictions to lead to predictions while still contstraining those predictions by what happens, both in terms of learning and overall response. The overall response could theoretically create a premature incorrect predictive spike which propagates through the network, but the primary contribution is to the response to actual sensory input, although it does reduce response latency especially to precise timing sequences.

The predictive signal causes enhanced responses to the sensory input even when the predictive signal is gone because it associates better since it responds more rapidly and fires more reliably. Specifically, it causes pooling of sensory concepts. Pseudoproximal inputs have dendritic segment processing but are not predictive as they lack NMDA spikes. Since they drive the cell to fire themselves fairly weakly, their inputs are normally dissociated from because there are many others (not to mention the perisomatic zone) so they usually aren't what caused the cell to fire. However, they can serve to identify simultaneous patterns which cause the cell to fire but are still independent patterns, as one can fully not occur at the same time as another occurs a fraction of the time since they are summed independently and ~ORed by sublinear summation when the neuron's represented pattern is clearly present (well above spike threshold). By pseudopredictive firing, L5 TT cells signal multiple patterns simultaneously because they are predictive, not the single current sensory input. If the postsynaptic cell's represented pattern can (not always, because bAPs are weaker signals than perisomatic firing and because burst input causes strong association which outpaces LTD) occur after the previous cell just a few tens of millseconds later, then the presynaptic cell predicts the postsynaptic cell. This is only ever a pseudopredictive signal, not a driving signal, because if the driving patterns to the pre- and post-synaptic cells occur simultaneously then most of the spikes in the presynaptic burst are after the postsynaptic singlet.

It doesn't matter much in general whether the postsynaptic cell bursts for learning because the postsynaptic cell's firing is shifted forwards but the burst occurs after the driving input, especially given the time step of the precise timing loop of ~10 ms. Not learning as much during later spikes in the postsynaptic burst might be necessary.

Because of the duration of a burst, even if the burst occurs immediately before the postsynaptic spike, the burst will have more spikes after the postsynaptic singlet, so it will still dissociate. This means there is a window starting and also ending a little bit before the postsynaptic spike in which the presynaptic spike must occur for LTP to occur. Also, temporal summation produces a stronger postsynaptic depolarization towards spiking if more spikes in the burst occur before the driving input. This produces grid-like sensory onset sequence tracking. Since L5 TT cells have fairly unselective RFs, they can often respond to the center RF regardless of the surface information.

The burst is sort of the opposite of prediction. Rather than predicting first and then trying to match what occurs afterwards, it sees what occurs first and then tries to predict what occurs afterwards, which is whether or not the cell bursts. Whether or not the cell bursts is decided after the first spike because of the bAP delay and because of how the precise timing loop works. Therefore, it helps decide whether to make the prediction by inputs to the apical dendrite. Since the calcium signal which evokes the burst is far larger than that evoked by a bAP, it essentially ignores the first spike in a burst (and singlets) for STDP at the distal apical dendrite and tuft, although the first spike still determines whether or not it can burst. Inputs to the apical tuft (non-primary for TT) produce NMDA spikes and therefore predict whether or not it will burst like SM prediction.

You could argue that L5 is just a motor output and signals to subcortical sensory structures are only corollary discharge. But then there's no way for the sensory processing in the cortex to signal to subcortical sensory processes. Also, the canonical primary thalamic input suggests otherwise, as well as the fact that primary thalamus can be sensory and higher order thalamus projects feedforward to L4 suggesting it is the same, just with more m-type cells. Also, the fact that vS1 L5 is reasonably thick but has only a minor motor output. And you don't need 1000 cells to control the behavior of a single whisker.

L5 has high spontaneous FRs, but don't think of that as a constant baseline frequency. Rather, it's evoked by noise and therefore variable, so it must be some sort of intentional randomness.

Maybe that variability in firing has to do with precise timing/flexible timing. Test different precise timings so find the one which matches. Otherwise, need a lot more cells to find one that matches the timing.
Maybe that also fits other things which are continuous and need flexibility. E.g. physical space. And maybe physical space on the sensor.
That relates to the idea that L5 fine tunes things by predicting them/predictive firing, since adjustable non-discrete values (motor/sensory precise timing, space, anything else) are tuned.

Just because L5 TT has a weaker response to POm input than L5 ST doesn't mean it receives less input on the tuft, because of the input to ST in L5 and because signals on distal apical are much stronger at the soma than signals at secondary+ tuft. So that's not necessarily an issue with the idea about ~loop (just signals back in L1) with m-type thalamus for precise timing recent post (posted a day or two before 6/28/18).

Timing is adjustable bc variable firing time -> learns when the right timing -> reinforces firing at that time.

Besides adjustable timing from noise-induced firing, being excitable means high firing rates for some cells. More like a wider range of FRs because it's more of a multiplier than just gain. That is useful for representing non-discrete values, although high firing rates might just be for uncertain timing, thereby being able to associate at more times (not just coincident with spiking but also a bit of a window, but again, don't think of it as high continuous firing rates, so it's more like being able to associate at more chunks of time rather than more intervals, since higher average FRs =/= higher constant FR) and also predicting things at more times/more chunks of time.

Maybe continuous space is dealt with in a way based on the previous note. Since there is a lot of variability in what % of time as cell is ~active (i.e. % of time there are postsynaptic EPSPs from a cell/in the window for presynaptic plasticity), this allows responding at different times based on continuous location. So since the precise timing of the sensory input varies continuously as a result of continuous variation in physical space, the two are the same problem in that way. And if it fires based on the precise timing that results from variation in location, it also associates to the continuous feature that co-occurs.

For that to work, it can't dissociate from the precise timing input when the complex feature it learns occurs with different timing. So e.g. broad shape/no surface information from m-type thalamus onto apical as the precise timing signal, because a bAP does not include plasticity in the tuft, only a calcium spike does. Just needs to be a distal segment so it requires that segment to fire, not just the neuron, for plasticity. But apical tuft makes the most sense. Also, because tuft segments probably don't undergo plasticity as a result of a calcium spike unless there is a local NMDA spike, it can learn multiple precise timings.

Sequence memory and object recognition are related in that both involve resolving ambiguity.

For SM, of sequence context, whereas for object recognition, of the object.

Is there a way to combine those two in a way which can also be split, explaining the gradient of thalamic input to L3a?

SM is on the feature level whereas object recognition is on the object level. So either sequences of objects or recognizing which object the feature is on.

The object can narrow down sequence contexts. Maybe that leads to object recognition based on sequence. Initially, recognize objects based on set of features (set because SM minicolumns are

all-cell firing bc not predicted). Then, learn feedforward connections from SM cells to object and feedback from object to distal segments of SM.
But if each cell in the SM narrows down the object, then learning the sequence will instead impair object disambiguation. So required: each cell in the SM cell contributes to some object recognition cells, and only those above a threshold remain active. That way, fewer sequence cells -> fewer object cells.
But that seems pretty inelegant. Since each feature is in location context, what it should be doing is infering the set of feature-location pairs based on sequence. Not narrowing it down to represent the combination of all sequences on the object.

Is there a way to narrow down the object and track the sequence using only a representation of the current feature in object/sequence gradient context? That could help explain the lack of sustained primary cortical cell firing.

If L5 CT thalamic synapses depress so strongly, then if the firing rate were to decline sufficiently (e.g. after a burst), the cell would become hyperpolarized, putting it in burst mode.

Since the synaptic depression takes a while to remove with L5 spontaneous FRs, that means the thalamic cell will not respond until the cell bursts again.

Whereas if the first input after silence was a singlet, the cell bursts, and then can respond in tonic mode to singlets. Since it is rare for there not to be any input (only e.g. during slow wave sleep), the real rule is respond to singlets with singlets.

So to summarize, singlets driven by sensory onset (specifically, onset of the CT cell increasing FR) drive singlets in thalamic cells, whereas a burst causes a singlet and then puts the thalamic cell in burst mode and it doesn't respond for a while except to bursts and is depressed longer by elevated firing rates from sensory (whether sensory onset or sustain; the elevated firing rate is all that matters because this is synaptic depression).

This could contribute to resolving object ambiguity. Using the sensory onset precise timing mechanism (maybe need to add a duration after bursting in which cannot burst), L5 CT cells burst in response to the first somatotopically aligned sensory input. That causes synaptic depression for a little bit on its own, which lasts even longer because of enhanced firing from additional somatotopically aligned sensory input (since either there is almost certainly more somatotopically aligned sensory input on a submacrocolumnar scale or it isn't examining the object with that fingertip anymore). As a result, it no longer drives thalamic firing unless the L5 CT cell bursts again. To burst again requires thalamic input to the tuft. Since that thalamic input depends on the previously responsive cells, this loop narrows down the object, or rather which feature-sensitive cells can respond to the object. But it's also a bit like SM and involves precise timing. Does this actually do all 3 and how to conceptualize that?

The article which describes CT depression also suggests that depression -> coincidence detection mode. That probably also contributes to the SM/precise timing/object recognition circuit, because the combination of different ~features -> object, or at least a more complex feature. But if all of those synapses burst, since a single synapse can't cause firing immediately after bursting even with another burst (same frequency), then it still isn't sufficient to change it to tonic mode. But if one of the convergent firing axons is firing a singlet at recently elevated FR, the switch to tonic mode following the burst evoked by the singlet input allows the inputs which were already depressed by bursting to contribute to reaching threshold.

Maybe a grid cell is a physically arranged patch on sheet, with individual non-sheet organized offsets locally.

If L5 TT is for fine-tuning responses, including sensory responses, why doesn't it project to primary sensory thalamus?

With a map on the cortical sheet of e.g. bar orientation, since there are longer latency responses to less prefered stimuli, there is a longer latency response further away from the input source on the submacrocolumnar level, not just between points on the sensor.

It is better to think of martinotti cells as inhibiting individual distal segments of the apical tuft, since the synapses are not on the apical trunk. Therefore, since apical tuft segments drive bursting, FDDI is not competition for bursting or the like. Rather, it is for sequence memory on the level of bursts rather than singlets.

L2/3 might be inhibited by whisker contact by the whisker ahead of it, but if there is a fine map, it is probably sent by L5 either directly from ST or via thalamus from TT
(remember, ST can burst, even if it might not burst for somatic current injection alone).

-later: I guess that's an explanation for dual input/output layers. Also, maybe contact order is like SM if it is based partially on latencies (as described in the post).

Hypothesis related to ramping activity/object pooler in post 6 24 18. Need to deal with inactivating/inhibiting behaviors after execution. Also other ideas in same post.

-to extend the sensory surface -> sensed surface xform to longer timescales, maybe larger cortical columns and/or longer range connections between CCs -> larger periods of propagation.

-how to make the xform invariant to the sequence of contact? E.g. changing orientation during contact. Maybe it doesn't matter because the oscillations have a sufficiently large wavelength that the tiny changes in sequence of contact don't matter, or maybe the mechanism should only be concerned with order of contact, not timing, and only be concerned with order of contact between adjacent points so the changes in order of contact induced by orientation don't matter.

-how does the xform being based on a brief sequence match up with the input being continuous, not a sequence, and L5 also doing precise timing? Maybe L5 adjusts the intervals of steps by precise timing mechanisms

-maybe think of colliding propagating waves in terms of amplitude/relative oscillation phases. Maybe instead of the coord xform being by which phase a point on the map is in when it makes contact, with the phase being constant and just being an inhibition-disinhibition spatial pattern on the sheet, everywhere oscillates on the sheet and the influence of making contact depends on the oscillatory phase at the time of contact. If in phase, it generates a large amplitude oscillation, which sends a large amplitude oscillation elsewhere. If out of phase, reduces the amplitude, reducing the magnitude and spatial extent of the oscillation from that point. And because it needs to oscillate with sufficient amplitude or else fewer cells reach threshold (with the oscillation left with few excitatory cells firing being inhibitory cell-only oscillations by inhibitory-inhibitory connections). So depending on the phase when contact begins, the oscillation at each point on the map grows or diminishes, and the amplitude and each point is like a hill because it contributes to amplitudes elsewhere less and less further away. The oscillation is either constant relative phases at each point, or maybe the relative phases at each point depend on the magnitude of it and nearby magnitudes bias the phase, so the hill-like influence both increases firing rates and shifts phases towards its local phase. That causes a warping of the map to increase the size of the parts of the map of the sensory surface around the peak of the hill because if other areas are in phase (accounting for latency though, but that latency is the same as that of lateral signals anyway) then they are enhanced by input nearer the peak. Also, during constant contact, the constant input drives oscillation at the amplitude because the negative-positive feedback loop has been established at a certain magnitude, involving more cells and therefore drawing more cells to reach threshold to take part in the oscillation, sustaining the amplitude.

phases are probably partly automatically determined/propagate because of subcortical oscillation-generating signals, relating to e.g. 100 ms sequence time steps. Maybe with wide RFs shifted to do coord xforms, the propagating oscillation means it responds at a given point on the cortical sheet where the stimulus and oscillatory peak line up, then at another point, it represents relative position (or rather, positions every n distance because modulus for grid cells) by shifting RFs (or more like selecting subsections of RFs) based on the oscillation needing to be peak. It doesn't have to pick just one phase, and instead could shift the source/relative cells as the oscillatory peak moves.

Oscillation -> during simultaneous touch on the same sensory surface (i.e. no joints) causes responses to represent relative positions, or actually shifted towards a mapping of relative positions.

Maybe a coord xform isn't even necessary for this same-sensory surface allo representation. To simplify things (bc it's normally just a shift towards allocentricity), let's say cells are responsive to that whole surface, more so towards a center of the RF (or in this case, each with a random point on the map bc they all have the same rfs in this thought experiment). Those RF centers are organized by the cortical sheet so organized the same as the propagating oscillation (i.e. the relative phase gradient). If there are two points generating input (which are the wavelength apart but would be handled by multiple wavelengths simultaneously and/or across a map and/or multiple wavelengths varying in influence magnitude across a map) and those are the wavelength apart, it generates two corresponding points firing with the same relative positions, but the positions also move sequentially as the oscillatory peak propagates. The set of cells which fired represent the relative location because they would've fired, although in a different order, for other relative positions.

But can't just look at the set to determine relative positions, need to look at the set of pairs which fired. Instead of doing that, add another level on top of the oscillation with higher frequency. This is to generate grid fields, while the lower frequency one is to shift RFs over time to create allocentricity. So cut the two point thing so far. If doesnt create grid fields, consider same wavelength but briefer peak/longer trough. Cells could still fire without the higher frequency and instead do grid fields by how many times the cell fires.

At this point, just need to create egocentric grid fields, which are enhanced when multiple points of stimulation are in the mod grid field. To do that, higher frequency oscillation with constant peak positions on the sheet, but different offsets in peak position for the cell. -> allocentric because that causes responses based on relative positions in egocentric RFs and the egocentric RFs shift over time so the shift causes them to be in egocentric RFs at some point if they are at the allocentric relative positions encoded. -> allocentric grid cells.

Maybe if these oscillations are in L5, L5 accounts for movement for the coord xform by generating movement by changing the oscillations in the exact same way necessary to account for movements.

Let's say the movement is encoded as a shift of the sensory surface to the right or left. To account for that shift, the low frequency oscillatory peak moves maybe. what about the higher frequency oscillation? Maybe the higher frequency oscillation is just by the firing rhythm, and it is adjusted by movement in a precise manner (smaller distances are all that matter beyond a movement of mod distance), using the same precise timing mechanisms which add precise distance beyond a mod distance motor command.

Without accounting for the movement, ignoring precise distance beyond mod distance with the grid cell spatial wavelength, a movement to the left effectively shifts the sensory surface to the right.

If the sensory surface shifts to the right, the low frequency oscillation needs to chase it (but that just means oscillating faster bc the peak does the chasing). This assumes the low frequency oscillatory peak travels from left to right. So to generate behavior which moves to the left, oscillate faster.

So to generate different behaviors, need different flows of propagation. They can be mostly the same flows for small numbers of behaviors in some regions, e.g. vS1 only has one behavior, retraction, so one flow of propagation. The motor map could explain that. Maybe the motor map is actually caused by randomly distributed by locally similar propagation sequences. Maybe the motor input to a section of the sheet (in e.g. particular layers; other motor input for other things) determines the motor output it produces, oscillating faster to generate the behavior. Or maybe it is different on a neuron by neuron basis because the motor map might not include movement direction.

To do precise movement distances, just use precise timing of firing since the firing rhythm determines the offset and timing/duration of movement determines movement distance. Whereas coarse distance is encoded by how rapidy low frequency oscillations occur, the exact distance is encoded by brief or not extreme elevations of firing of all the cells generating the movement. That causes the high frequency oscillation to adjust phase.

Could low frequency oscillatory peaks be randomly distruted, like independent relative phases?

The low frequency oscillation shifts RFs by only letting neurons fire when they have low frequency peak. Each point of sensory stimulation makes an area on the sheet weakly active, then the low frequency propagation makes cells fire in a sliding matter over the sheet. Maybe to shift RFs it requires first activating those aligned with center RF and then less and less aligned. Maybe that's the source of the grid offsets because firing rhythms begin based on distance of RF center from the feature.

Shouldn't the higher frequency oscillation also move? Without it, the sets of cells which fire at some moment represent relative location. With it, cells should fire for mod relative positions. But it's still mod relative to the sensor so if it shifts positions while still same relative positions, won't respond. So that needs to be a step after the lower frequency. Slide the RFs in one set of cells. Send that signal to another set of cells. They fire based on mod relative positions of those sliding RFs.

Actually, a separate layer isn't necessary. Instead, start the firing rhythm (=high frequency oscillation) when the cell first fires. Cells which fire multiple times have found a relative position in allocentric space.

Maybe can extend that idea of when the firing rhythm starts to allow path integration. When it contacts a feature, some cells start their firing rhythm. That firing rhythm can dictate the firing rhythms of other cells as part of path integration, by being mod distance with offset with respect to the first contacted feature.

Maybe synchronizing them is involved in behavior.

If the low frequency propagations are like the picture, with a sudden high oscillatory peak trailed by a mnuch longer lasting gradual decrease (more decrease towards the start of decrease), that means the low frequency oscillation is like a single point of responsive cells sliding along the sheet, because cells which respond a bit after that sudden peak get inhibited and nothing is active just before the sudden peak so it's not the same as a sharp peak, and then if nothing responds, no inhibition so uses a more lenient enhancement, with less spatial selectivity of that peak where cells don't fire.
Maybe can apply grid cells to motor by each representing a mod distance of different offset/spatial frequency and summing like place cells.

How to combine that with the sensory -> sensed surface xform? Maybe use the idea of creating a consistent coordinate frame by consistent firing rhythm offsets, but on the scale of a contacted surface.

Article: Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex.
Localized cluster of neurons have oscillation amplitude based on how well the stimulus matches their prefered stimulus.

If L5 cells generate behavior by oscillatory frequency, that helps explain sensory responses.

1st iteration xform:
-maybe instead of propagation spreading out, synchronize the oscillations sort of. Synchronize the oscillation at that moment at that position with the sensory input. Other points do not respond because the oscillatory peak moves/there is a gradient of relative phases.