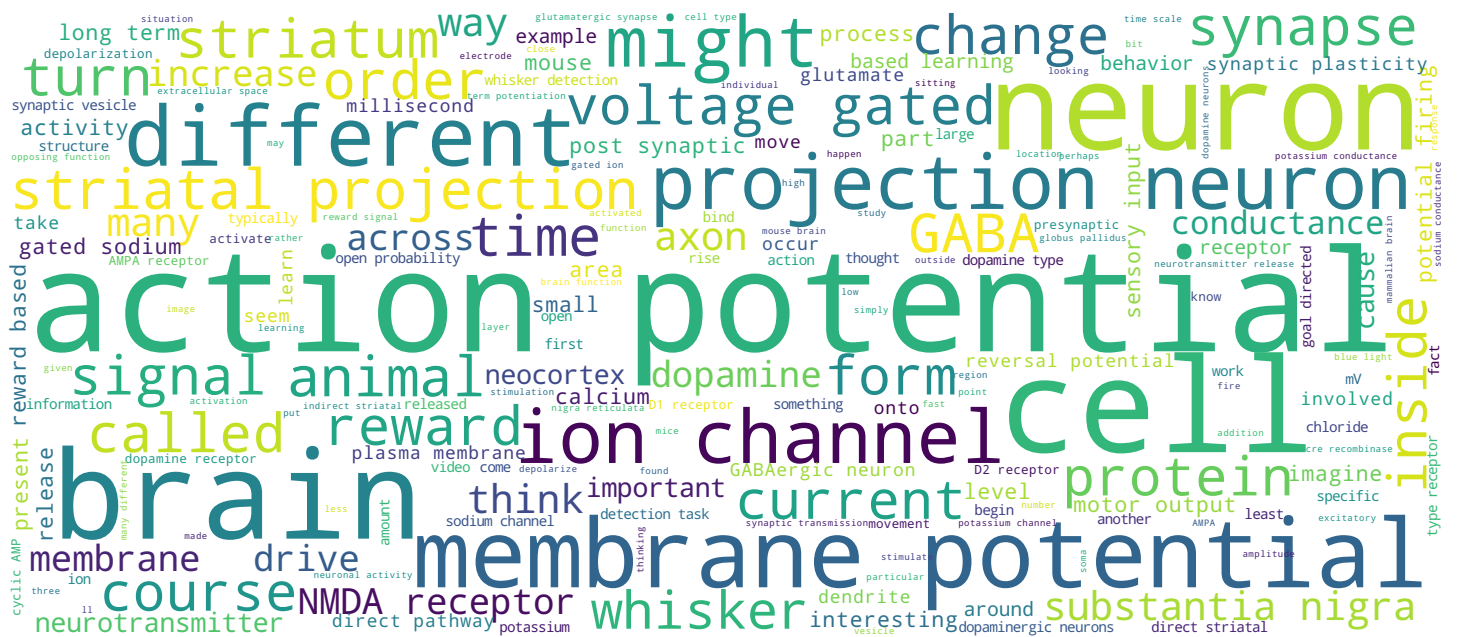
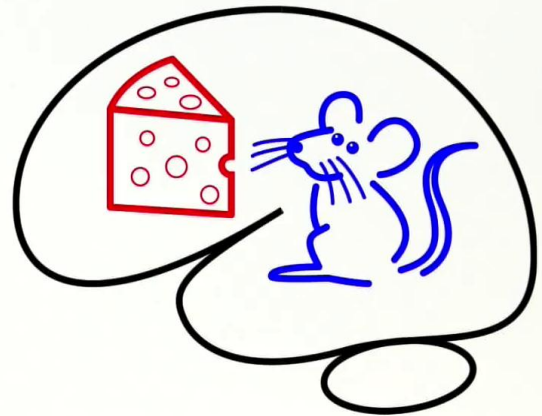


Cellular Mechanisms of Brain Function

Prof. Carl Petersen



Reward-based learning



Cellular Mechanisms of Brain Function

One of the most important features of animal behavior is our ability to learn from experience. Such adaptive behavior is important for survival in our ever-changing environment. The brain has evolved many mechanisms of learning which likely contribute differentially to distinct types of learning. Thus declarative learning, procedural learning, as well as fear learning, spatial learning and motor learning, may involve different changes in distinct neuronal circuits in different parts of the brain. In a previous video, video 6.2, we briefly discussed evidence that spatial learning through NMDA receptors in the CA1 region of the hippocampus might play an important role. Indeed, the hippocampus and neocortex are likely to be involved in many forms of learning, and this is supported by increasing amounts of evidence from experiments both correlating and manipulating neuronal activity in hippocampus and neocortex. However, in this video, we'll focus our attention on other aspects of learning that might be perhaps more ancient and thus perhaps more fundamental — a form of learning that likely existed long before the neocortex evolved. Here we will specifically focus on reward-based learning.

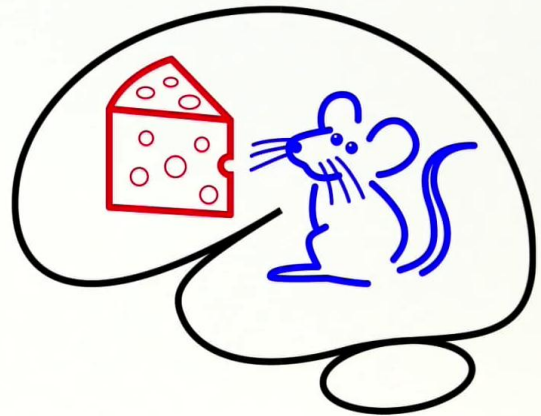
Notes

Summary



0m 05s

Reward-based learning



Cellular Mechanisms of Brain Function

Presumably an important motivation for carrying out various actions is a desire to achieve positive outcomes, accompanied by immediate or future rewards. Sometimes we're successful, either by design or chance, and it's important for us to learn from the outcome of our actions so that we can maximize our reward. In this video, we will discuss one attractive hypothesis for how reward-based learning might occur in the mammalian brain. However, it's good to bear in mind that there are many other hypotheses about how learning takes place and this is a very active field of research.

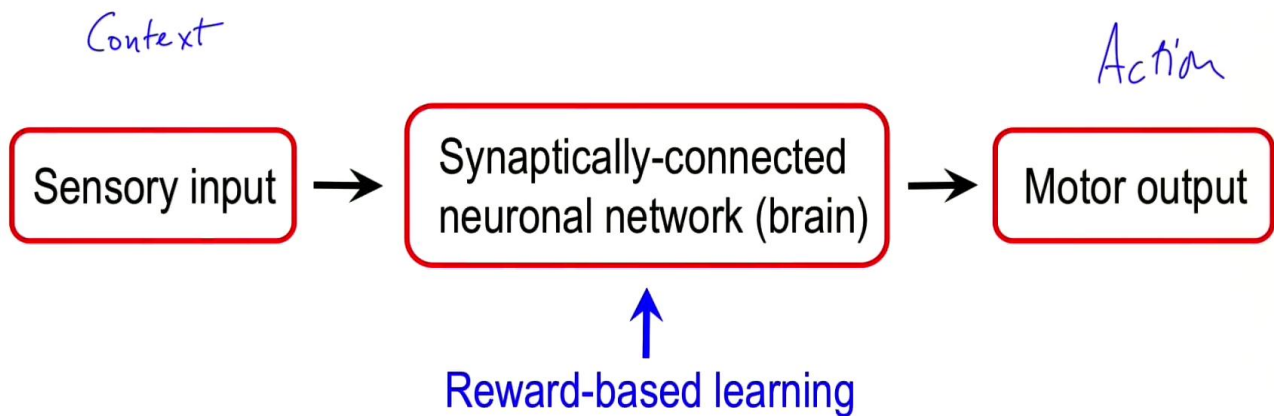
Notes

Summary



1m 30s

Reward-based learning



Cellular Mechanisms of Brain Function

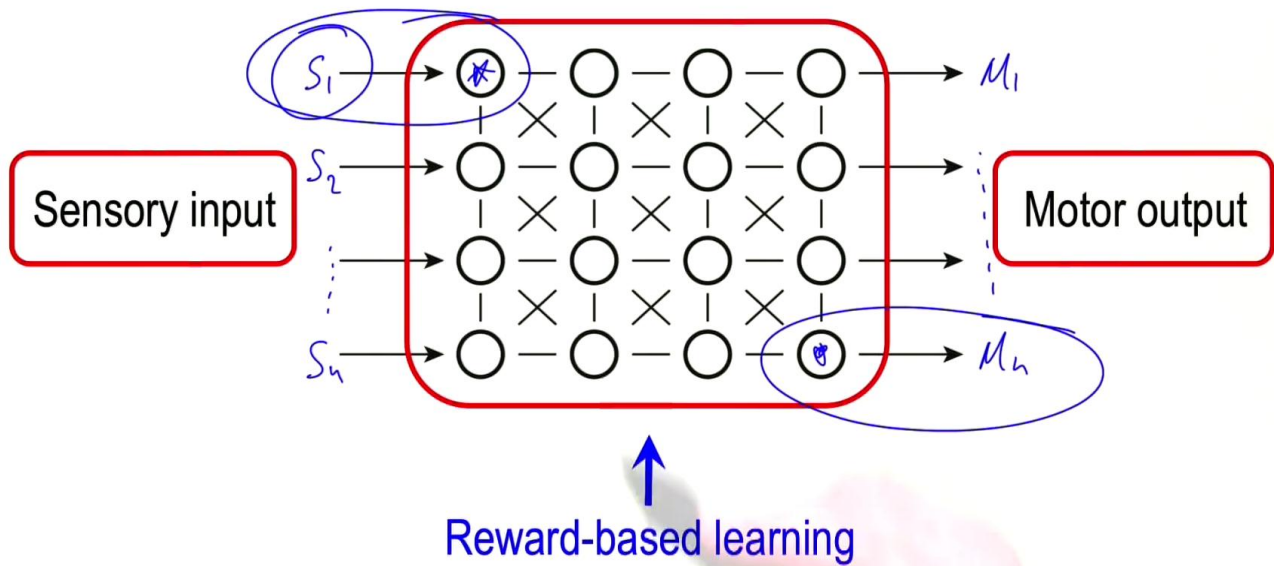
Let's think a bit about how reward-based learning might work. We imagine there's some sensory input, there's some sort of context that's provided by sensory input — that's how we know where we are. That's processed by the brain, the synaptically connected neuronal networks that govern brain function, and that then gives rise to some form of motor output — an action. That action, that motor output, might give rise to reward. It may be a successful action and we may be rewarded for the performance of that action in the context of where that sensory input arrived. So the basic idea behind reward-based learning is to learn that in a specific stimulus context we should perform a specific motor output and that will then lead us to reward.

Notes

Summary



Reward-based learning in neuronal networks



Cellular Mechanisms of Brain Function

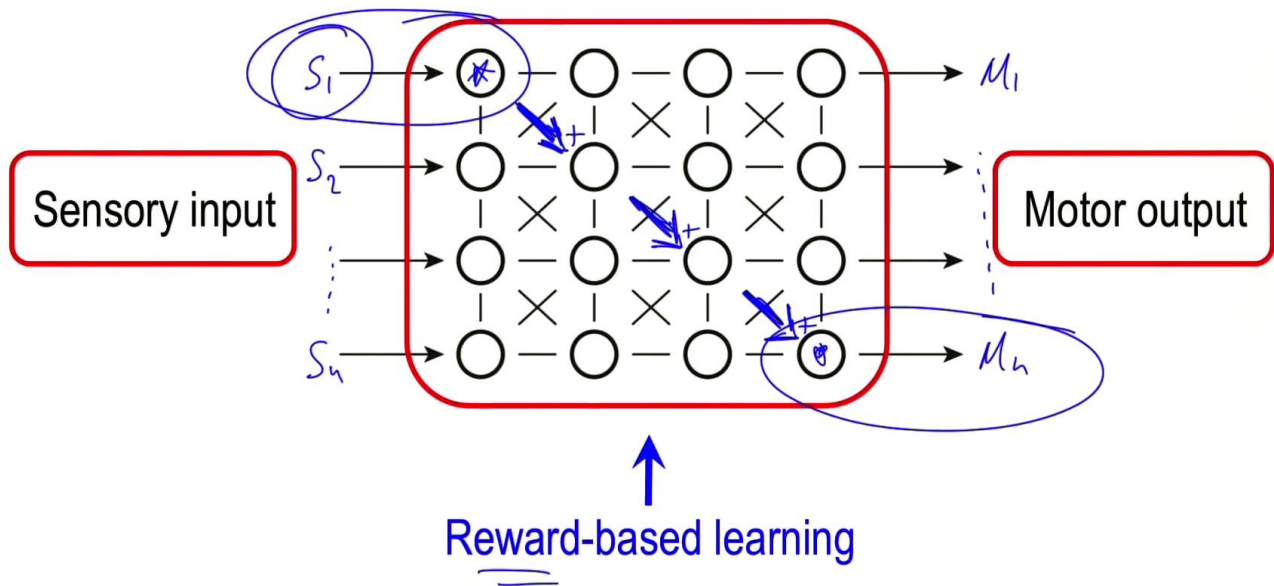
At the level of neuronal networks, we can begin to think a bit more specifically about what types of learning ought to occur in the brain. We again imagine our context provided by the sensory input and of course there are many different channels and types of sensory input that we could have. That comes in, excites neurons or groups of neurons inside the brain and that then forms these connected neuronal networks where different neurons in different parts of the brain talk to each other. They process the information and give rise to motor output, and of course there are many different motor outputs that we can imagine. In one particular behavior, we might imagine that there's a given sensory stimulus that comes in, a given context that's provided to the animal that generates activity in one node of the network, and in order for the animal to receive reward, it needs to carry out a specific action that it turns out is driven by activity in this motor node of the network. So what the animal needs to learn is to convert this sensory input into this motor output. In terms of how one might do this at the level of individual neurons, you might want to strengthen this most direct pathway that links the sensory stimulation to the motor output.

Notes

Summary



Reward-based learning in neuronal networks



Cellular Mechanisms of Brain Function

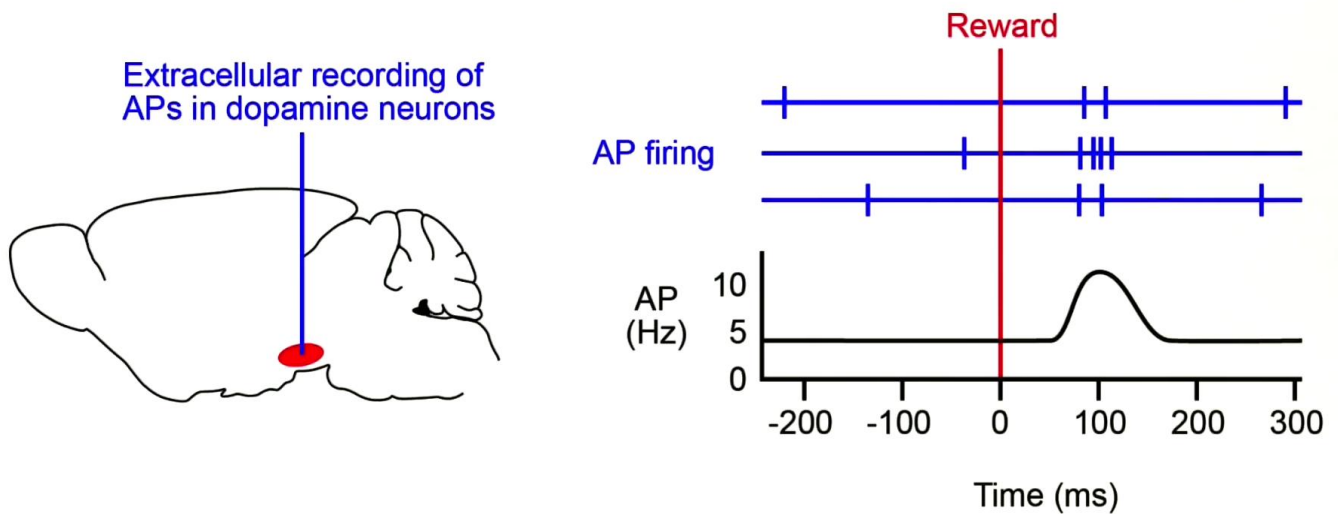
So we might want to — if we imagine that these are individual neurons and this might be glutamatergic synapses, what we'd want to do is to increase the efficacy of these synapses, perhaps through the process of synaptic plasticity and long-term potentiation. If we could strengthen the link between this sensory input and that motor output, then that would give rise to an enhanced sensorimotor transformation that would then give rise to reward for the animal. So what reward should do, then, is to drive this form of plasticity in the sensorimotor networks of the brain. So a key question, obviously, is how reward is represented in the brain.

Notes

Summary



Dopamine reward signal



Cellular Mechanisms of Brain Function

In the now classic work of Wolfram Schultz, who recorded dopaminergic neurons of the midbrain in the monkey, it seems that there's an interesting reward signal that's present there in the dopamine neurons. What Wolfram Schultz observed is that when he gave reward to his monkeys, there would be a brief transient increase in the action potential firing of the dopamine neurons, and that appears, then, to encode a reward signal. More recently, these experiments have been repeated in mice by the work of Naoshige Uchida and in mice, in addition, we have the ability to verify that these really are dopaminergic neurons through detailed examination of the properties of these cells. The work of Naoshige Uchida and Wolfram Schultz are in close agreement, and what both of them found is that when reward is delivered to an animal, and typically this is in the form of a juice, for example, that's squirted into the mouth of a thirsty monkey, some 100 milliseconds or so afterwards there's a phasic increase in the action potential firing of the dopaminergic neurons.

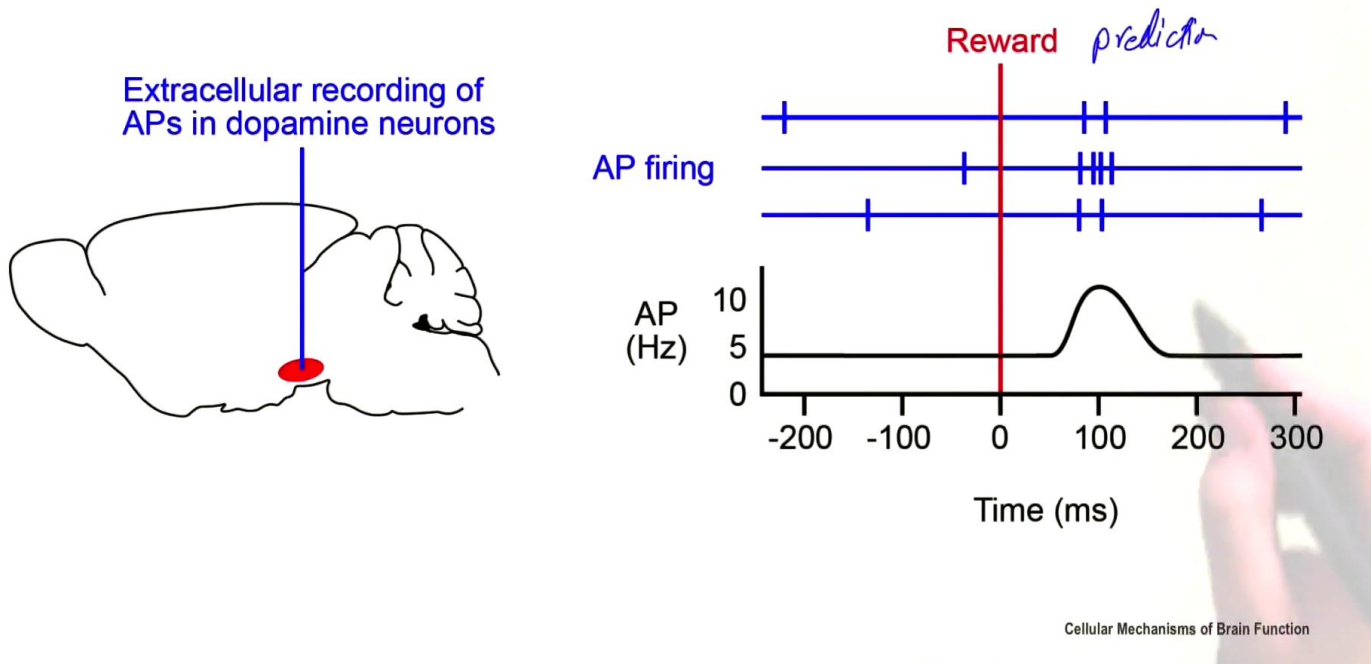
Notes

Summary



5m 31s

Dopamine reward signal



So there's an extracellular recording electrode that's been lowered to where the midbrain dopamine neurons are— these are spontaneously active neurons and so they will fire action potentials every few hundred milliseconds; maybe they have a basal firing rate of around five Hertz, but when reward is delivered, some hundred milliseconds later there's an increase in the action potential firing of these dopaminergic neurons and that's a transient signal that again lasts about 100 milliseconds or so. So reward is associated with a brief increase in dopaminergic signalling. Interestingly, if a cue is presented before the reward and it's done so on a regular basis — say it's a one-second interval between the cue — this could be a sound stimulus or a visual stimulus that predicts reward delivery — then the dopaminergic neurons fire not so much to the reward but rather to the cue that predicts the reward. So in fact the signal here that drives dopamine signaling is not so much reward but reward prediction. So this phasic dopamine signal seems to move to the earliest predictor of reward coming the way of the animal.

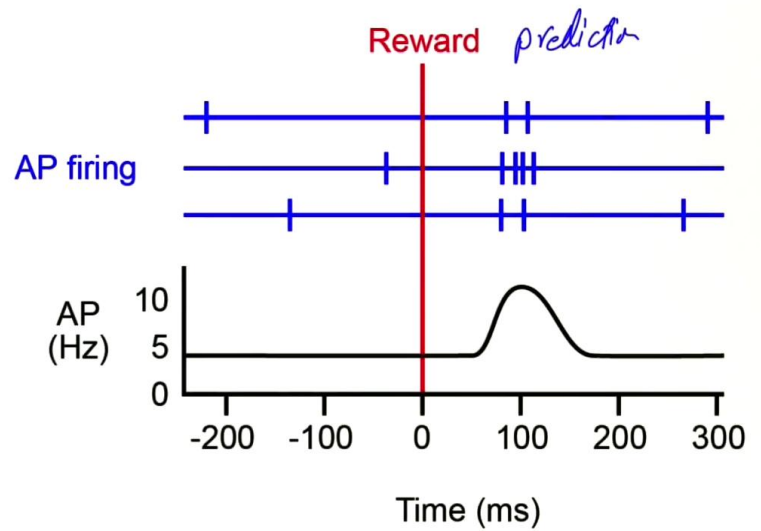
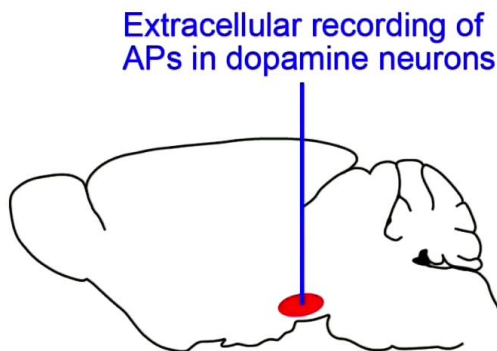
Notes

Summary



6m 45s

Dopamine reward signal



Cellular Mechanisms of Brain Function

That's clearly a very interesting signal in terms of learning how the animal should figure out how it should best behave in order to achieve more reward. That seems to be then given by dopaminergic signals that drive the sensorimotor loops for learning how to behave and how to get more reward. Having found this reward signal in the dopamine neurons, it's obviously interesting to see where the signal goes in the brain.

Notes

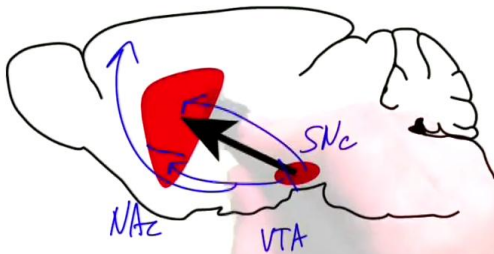
Summary



8m 04s

Dopaminergic innervation of striatum

Midbrain dopaminergic neurons prominently innervate striatum



Cellular Mechanisms of Brain Function

The midbrain dopaminergic neurons project most strongly to the striatum. So that's the location where they release dopamine most prominently in the brain. The striatum is a relatively large brain area. It can be divided into the dorsal and the *ventral striatum*, and the ventral striatum is sometimes called the *nucleus accumbens*. The midbrain dopaminergic neurons can also be divided into two groups. There's the so-called *ventral tegmental area* that heavily innervates the nucleus accumbens and also signals to the frontal cortex and there's the *substantia nigra compacta* that signals to the dorsal striatum. The heavy innervation of striatum by dopamine then suggests that the dopamine is probably carrying out some important functions in the striatum. So if you now zoom in on a small bit of the striatum, we'll see that there are axons, dopaminergic axons that have swellings in them, and at these swellings, there are many synaptic vesicles that contain dopamine and it's thought that dopamine is released from these swellings very much like they are in other synapses of the brain.

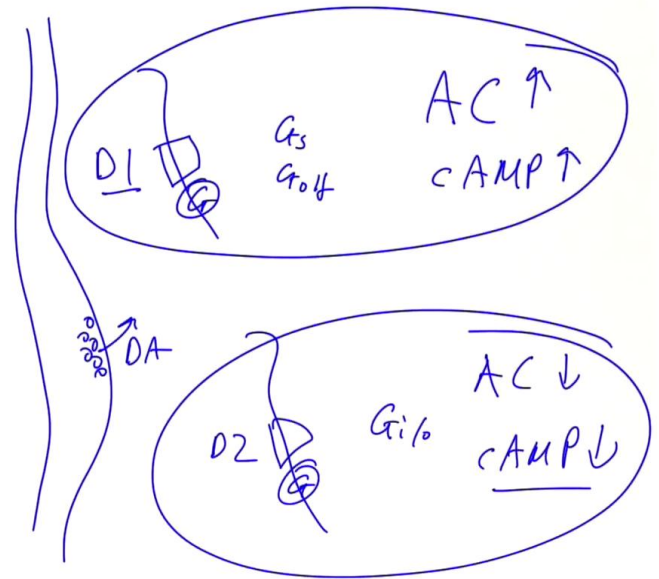
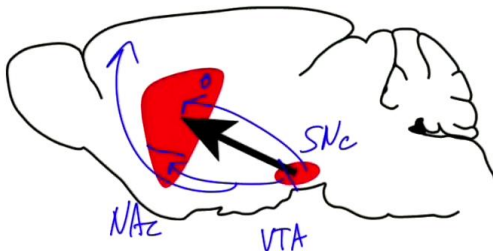
Notes

Summary



Dopaminergic innervation of striatum

Midbrain dopaminergic neurons prominently innervate striatum



Cellular Mechanisms of Brain Function

There are however no immediate post-synaptic membranes that are tightly opposed to these release sites, so the dopamine is thought to diffuse in the extracellular space. There it can bind to dopamine receptors that are present on membranes around it. There are two important types of dopamine receptors. There are the dopamine type 1 receptors and there are the dopamine type 2 receptors and both of them are set in trans-membrane receptors that couple through G proteins, so the next downstream signal here is G protein activation after dopamine binding to the receptor. The dopamine type 1 receptor signals the G(s) and G(off) — these are two different types of G proteins, and they signal through a variety of different pathways including the stimulation of adenylate cyclase, that leads to increased production of the second messenger, cyclic AMP. The D(2) receptors signal through other G proteins, G(i) or G(o), G proteins that have a variety of signalling targets including Phospholipase C but in addition they also inhibit adenylate cyclase and thus reduce cyclic AMP concentrations. Although they probably have many other functions, there at least is, at some level, the dopamine type 1 receptors increase cyclic AMP and the dopamine type 2 receptors decrease cyclic AMP, so they have in some respects opposing functions.

Notes

Summary

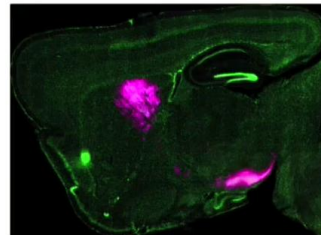
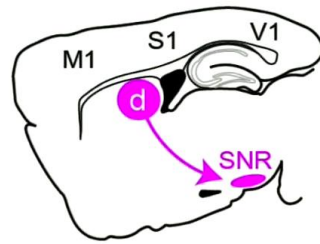


Striatal projection neurons

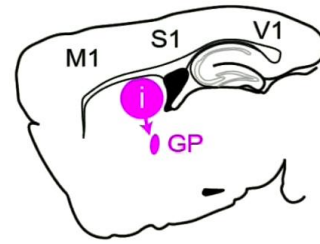
Direct pathway
striatal projection
neurons (dSPNs)

Express D1Rs

D1-Cre



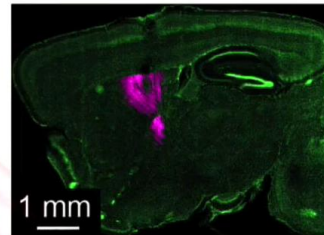
Sippy, Lapray, Crochet & Petersen



Indirect pathway
striatal projection
neurons (iSPNs)

Express D2Rs

D2-Cre



Cellular Mechanisms of Brain Function

Interestingly, these dopamine receptors are expressed on distinct neurons in the striatum. So there are two types of striatal projection neurons. There are the ones that express the dopamine type 1 receptors and there is another type of neuron that expresses the dopamine type 2 receptor. We can look at these specifically through mouse genetics. We can express *cre recombinase* downstream of the promoter of the D1 receptor or we can express *cre recombinase* downstream of the promoter for the D2 receptor. If we now inject a virus that's activated by *cre recombinase* to express green fluorescent protein, then we can see where the axons of these neurons go, and so that's exactly what we've done here. So here's the striatum. This is a sagittal section of the mouse brain, it's a D1 *cre* mouse, we've injected virus here into the striatum, so we've infected cells and expressed green fluorescent protein in the D1 *cre* animals, and the GFB here is shown in pink and so here we see the cells that are infected by the virus and further away we see the axonal projection of these neurons. So they send axons that go all the way over here to the *substantia nigra reticulata*.

Notes

Summary



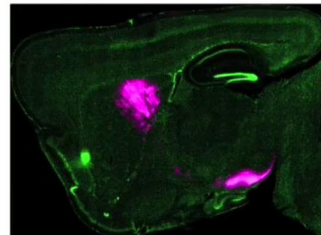
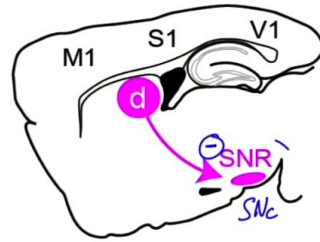
11m 37s

Striatal projection neurons

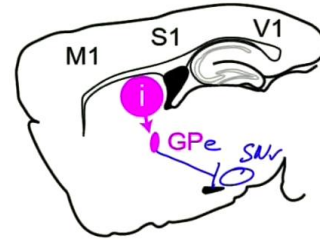
Direct pathway
striatal projection
neurons (dSPNs)

Express D1Rs

D1-Cre



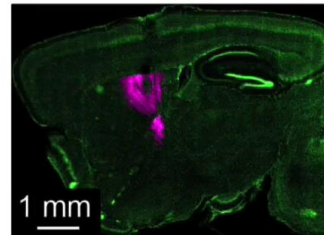
Sippy, Lapray, Crochet & Petersen



Indirect pathway
striatal projection
neurons (iSPNs)

Express D2Rs

D2-Cre



Cellular Mechanisms of Brain Function

So that's actually immediately next to where the dopamine neurons are of the substantia nigra compacta that signal back here to the striatum and release dopamine. Here we're looking at the reverse projection, the striatal projection neurons, the ones that express D1 receptors, they send axons to the substantia nigra pars reticulata and all the projection neurons here in the striatum are GABAergic neurons and so this direct pathway inhibits substantia nigra reticulata. The indirect pathway we can study in these D2 cre expressing mice. Again, here we have a sagittal section, we've injected a virus that gets activated by cre recombinase, we infect cells in the striatum, and now we see that there's a very different axonal projection that targets an area here much closer to the striatum, the so-called *globus pallidus* external part. It's a GABAergic projection from these D2 expressing neurons that inhibits the globus pallidus. The globus pallidus, in turn, contains GABAergic neurons that then inhibits substantia nigra reticulata. So not only do these two pathways express dopamine receptors that at least at some level have opposing functions, where this increases cyclic AMP concentration and this decreases cyclic AMP concentration but in terms of the impact upon the substantia nigra reticulata, they have opposing functions.

Notes

Summary



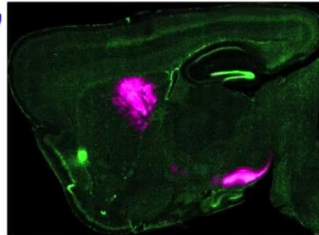
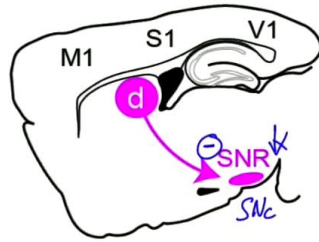
13m 06s

Striatal projection neurons

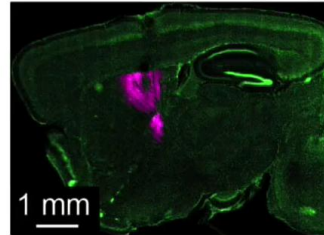
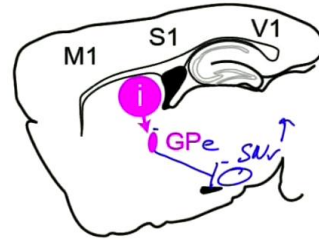
Direct pathway
striatal projection
neurons (dSPNs)

Express D1Rs

D1-Cre



Sippy, Lapray, Crochet & Petersen



Indirect pathway
striatal projection
neurons (iSPNs)

Express D2Rs

D2-Cre

Cellular Mechanisms of Brain Function

The direct pathway directly inhibits substantia nigra reticulata whereas the indirect pathway inhibits globus pallidus that then inhibits substantia nigra reticulata, so that's a disinhibitory pathway that in the end leads to an increase of reticulata function and here a decrease of reticulata function. These two pathways are rather interesting in the sense that they're very clearly distinct in their projections, their functions, and also in terms of the dopamine receptors that they express.

Notes

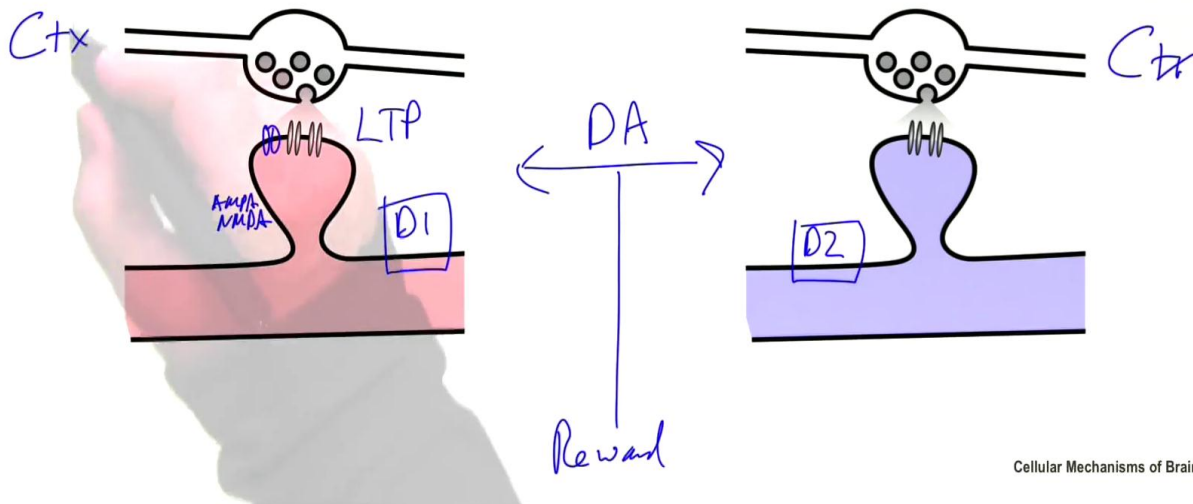
Summary



Dopamine modulation of synaptic plasticity

D1 receptor expressing
striatal projection neuron

D2 receptor expressing
striatal projection neuron



Cellular Mechanisms of Brain Function

How does that help us understand reward-based learning? Let's have a look and see what happens if we imagine that there's some reward that causes activation of the dopamine pathway, release of dopamine into the striatum, the dopamine diffuses in the extracellular space, and it then binds to D2 receptors or D1 receptors, as we've seen expressed in different cell types. The action of the dopamine receptors seems to be particularly involved in regulating synaptic plasticity at the glutamatergic synapses that arrive onto the striatal projection neurons. It turns out that a large amount of the excitation that arrives in the striatum comes from the neocortex, and so cortical projections come into the striatum and release glutamate onto the striatal projection neurons. That's true of both pathways, both the direct and the indirect pathways. The action of dopamine in the context of the glutamatergic synapse here is on the D1 receptors to promote the insertion of both AMPA and NMDA receptors onto the postsynaptic membranes and D1 activation appears to be essential for promoting long-term potentiation at this synapse here. So the cortical input here onto the direct striatal projection neurons appears to be potentiated by the action of dopamine.

Notes

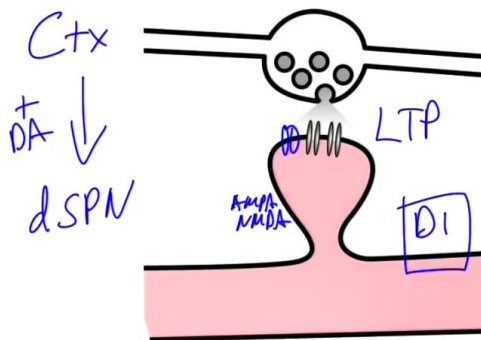
Summary



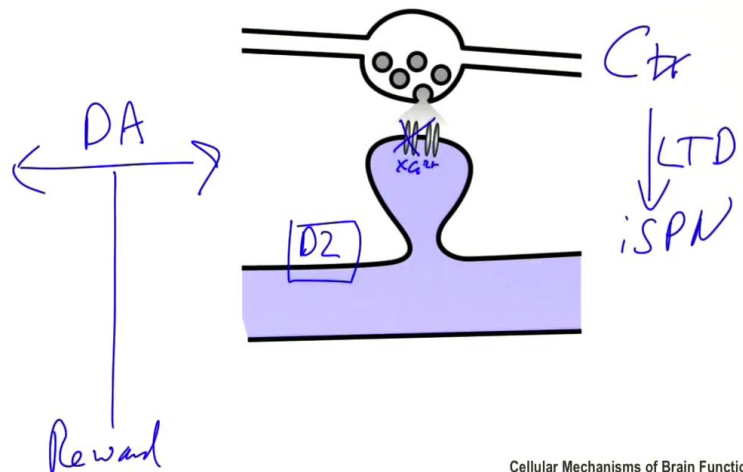
15m 14s

Dopamine modulation of synaptic plasticity

D1 receptor expressing striatal projection neuron



D2 receptor expressing striatal projection neuron



Cellular Mechanisms of Brain Function

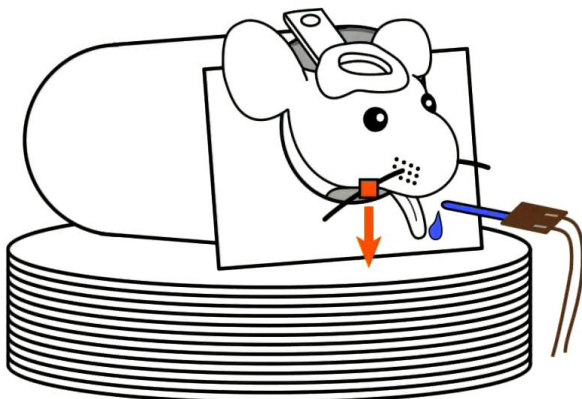
So that's a strengthening of that pathway in the presence of reward. In contrast, the activation of D2 receptors seems to be involved in just the opposite type of synaptic plasticity. Activation of D2 receptors prevents calcium entry through the NMDA receptor, and also seems to be involved in potentiating the long-term depression at this synapse. So D2 activation gives rise to long-term depression, here onto the indirect striatal projection neurons. So dopamine seems to have two opposing functions. It strengthens the direct pathway, cortex to direct striatal projection neurons, and it weakens the pathway onto the indirect striatal projection neurons.

Notes

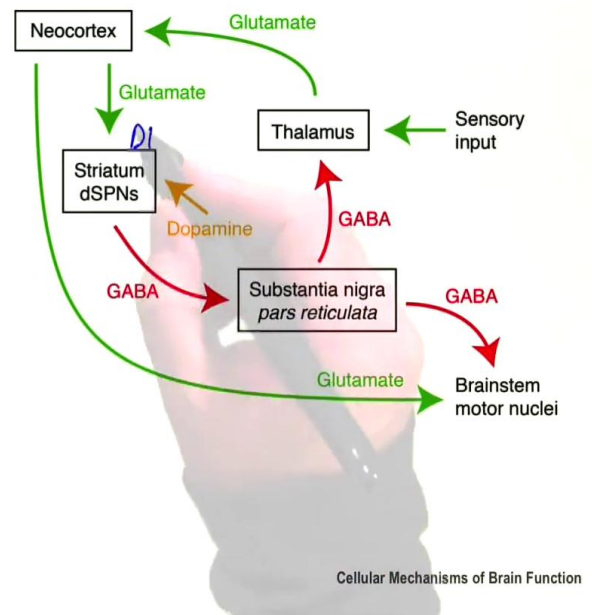
Summary



Whisker detection task



Sachidhanandam, Sreenivasan, Kyriakatos, Kremer & Petersen, 2013



Cellular Mechanisms of Brain Function

Let's look at that in the context of one goal-directed sensorimotor transformation that we've already been thinking about. Let's take the whisker detection task where we stimulate a whisker and the mouse needs to learn to do a goal-directed motor output involving licking a reward spout in order to obtain reward. At the level of a circuit diagram, this is how it might work. The sensory input arrives at the periphery, drives excitation of somatosensory thalamus that in turn excites the somatosensory neocortex and we already saw that activity in the neocortex is essential in order to carry out this whisker detection task. The neocortex in turn sends glutamatergic projections to the striatum, and one interesting place that could be reinforced by the reward signal by dopamine is the glutamatergic input onto the striatal direct projection neurons, the D1 receptor-expressing pathway. Dopamine might potentiate that, drive more activity in the direct pathway downstream of sensory activity, the activity in these direct pathway neurons might then inhibit the substantia nigra pars reticulata and that contains GABAergic neurons that are tonically active and normally they inhibit the brain stem motor nuclei and prevent the animal from moving in some sort of random way.

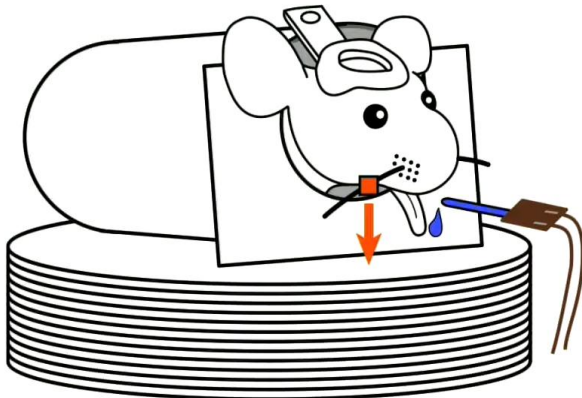
Notes

Summary

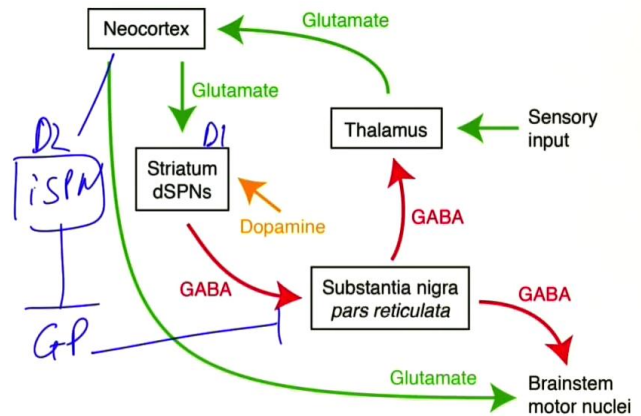


17m 54s

Whisker detection task



Sachidhanandam, Sreenivasan, Kyriakatos, Kremer & Petersen, 2013



Cellular Mechanisms of Brain Function

So there's a tonic inhibition of brain stem motor nuclei and that tonic activity in the substantia nigra might be what's been directly inhibited by the striatal direct projection neurons. The D2 receptor indirect pathway spinal projection neurons might be weakened by the dopamine synaptic plasticity. There would then be less inhibition onto the globus pallidus neurons and they would then in turn inhibit the substantia nigra pars reticulata more, and so we would have two ways in which the basal ganglia circuit via striatum and substantia nigra might contribute to the licking motor output in this goal-directed sensorimotor transformation of the whisker detection task. So that's a nice hypothesis. Let's see if there's any data that might support it.

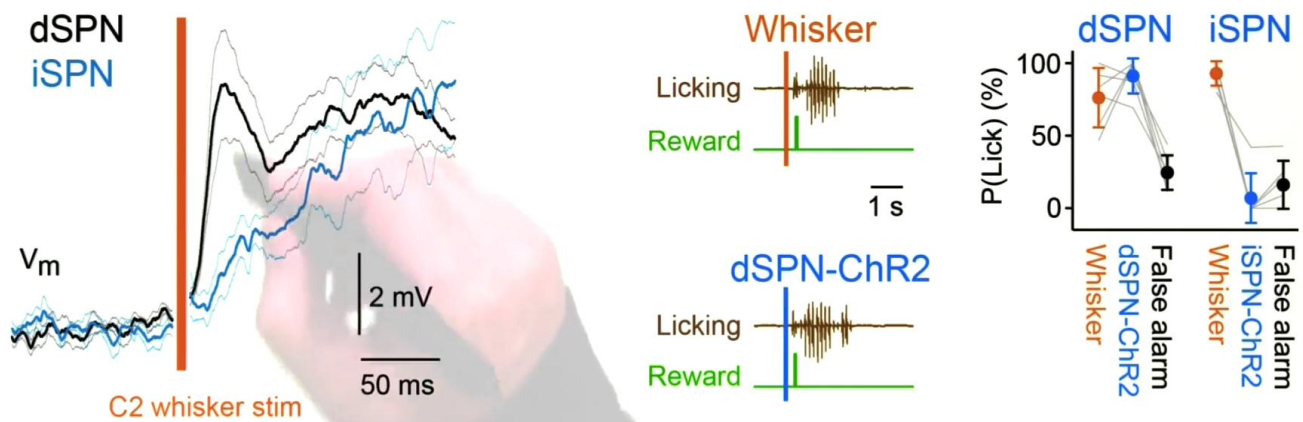
Notes

Summary



19m 25s

Cell-type-specific function of striatal neurons



Sippy, Lapray, Crochet & Petersen

Cellular Mechanisms of Brain Function

So in the first membrane potential recordings of striatal projection neurons during goal-directed behavior, we found that there's in fact a brief, transient sensory response in the direct striatal projection neurons and that seems to be missing or at least much weaker in the indirect striatal projection neurons. So this transient excitation of the direct pathway could contribute to a go signal initiating the licking motor response during the whisker detection task. To see if that brief activation of the direct pathway might be sufficient in order to drive licking behavior, we tried an optogenetic substitution experiment. So we trained the animals as before in the whisker detection task, we stimulate the whisker, the animal licks, it gets a reward, and now, once the animal's learned the whisker task we then start stimulating with blue light the direct striatal projection neurons that have been infected by our virus-expressing channelrhodopsin. So we can activate the striatal projection neurons, these dSPNs that are normally activated by the sensory stimulus during the whisker detection task. Now we bypass the sensory stimulus and we directly stimulate those neurons.

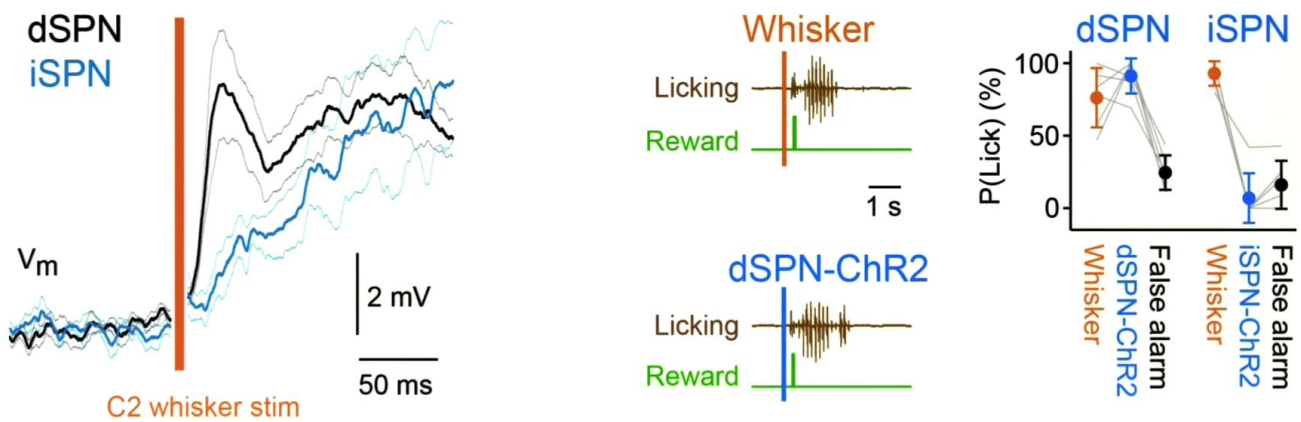
Notes

Summary



20m 22s

Cell-type-specific function of striatal neurons



Sippy, Lapray, Crochet & Petersen

Cellular Mechanisms of Brain Function

We give them a brief few-millisecond blue light flash, and that turns out to drive the licking motor output. So apparently this activity in the direct striatal projection neurons is sufficient to drive the licking motor response. Here you can see the group statistics: the stimulation of the direct striatal projection neurons is in fact even better than the whisker stimulus at driving licking, and of course, false alarm rates are low under these conditions, whereas stimulating the indirect striatal projection neurons in no way drives further licking. If anything, it prevents the animal from licking. So there's a highly specific pathway through the direct striatal projection neurons that drives licking and the indirect neurons seem to have the opposing function.

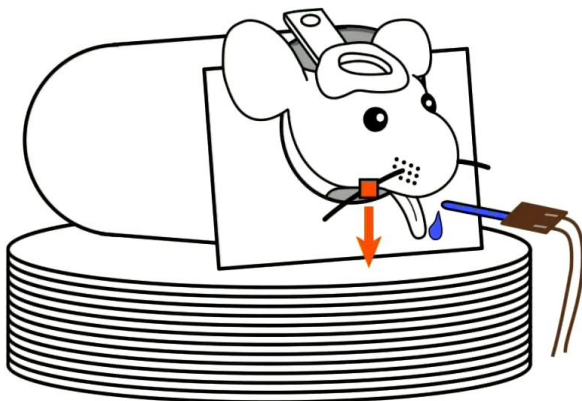
Notes

Summary

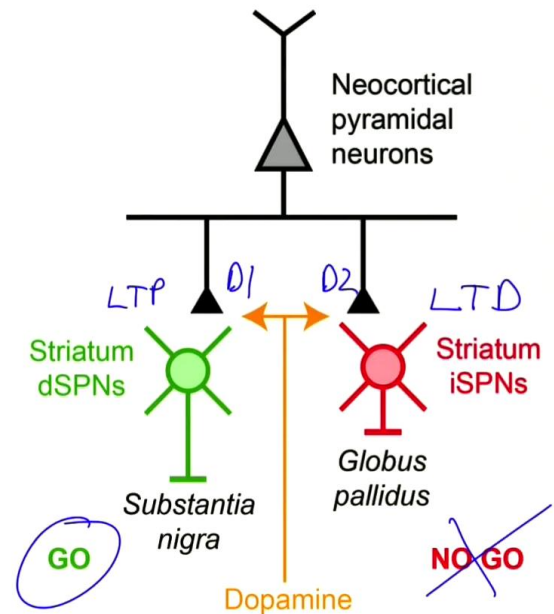
21m 40s



Dopamine learning in the whisker detection task ?



Sachidhanandam, Sreenivasan, Kyriakatos, Kremer & Petersen, 2013



Cellular Mechanisms of Brain Function

So if you now put this at the level of a circuit and synaptic wiring diagram, we think that the neocortex is involved in the performance of this whisker detection task, and one of the outputs of the neocortex is the glutamatergic innervation of the striatum. The dopamine reward signal that presumably is involved in the learning of this task might contribute to long-term potentiation at the neocortical synapse onto the direct pathway striatal projection neurons through the activation of D1 receptors. That would then strengthen the go pathway, whereas the glutamatergic input onto the indirect striatal projection neurons might be decreased in efficacy through long-term depression driven by the D2 receptors and the depression of this pathway here might then prevent the no-go signal. So altogether what we see is that this pathway could be involved in the reward-based learning of this very simple goal-directed sensorimotor loop downstream of the whisker stimulus.

Notes

Summary



Mechanisms underlying reward-based learning



- Reward-based learning likely reinforces the synaptic circuits driving goal-directed behavior.
- Transient action potential firing in dopamine neurons signals unexpected reward, an important learning signal. The released dopamine could drive synaptic plasticity underlying learning.

Cellular Mechanisms of Brain Function

In this video, we've discussed one specific hypothesis that attempts to provide a mechanistic neuronal circuit and synaptic plasticity mechanisms underlying reward-based learning. The basic concept is simple: reward-based learning should reinforce the synaptic circuits driving goal-directed sensorimotor transformations. That way, actions leading to rewards should be carried out in the appropriate context. Mechanistically, the transient action potential firing of dopaminergic neurons in the midbrain appears to signal reward, or cues that predict reward. In fact, more formally it's thought to be reward prediction errors. This could form the signal driving synaptic plasticity at corticostriatal synapses, which might participate importantly in reward-based learning. These same circuits that are so useful for learning goal-directed behavior also go wrong in some brain diseases. We'll discuss that in the next video.

Notes

Summary



23m 48s