a

b

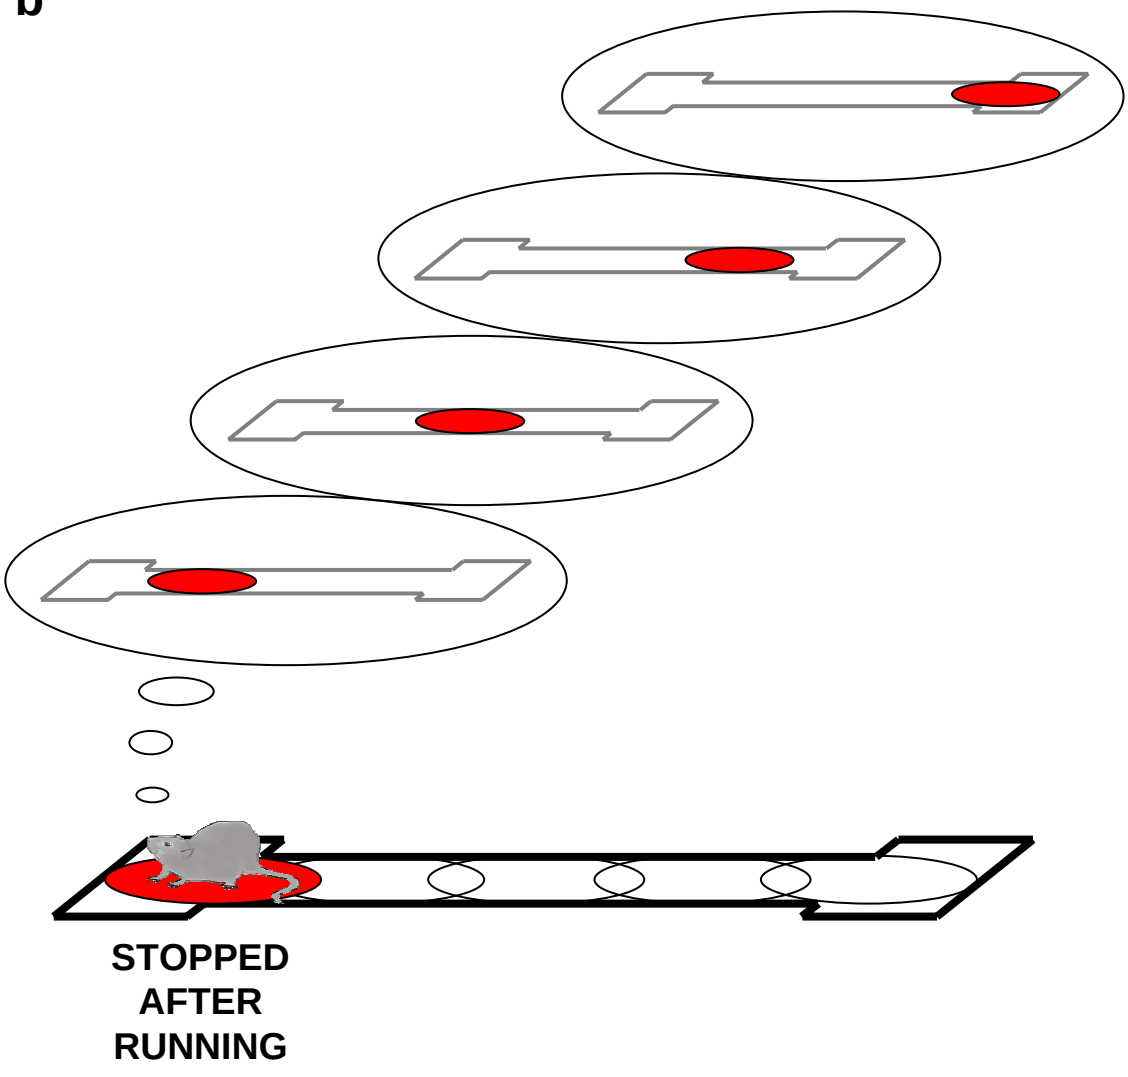

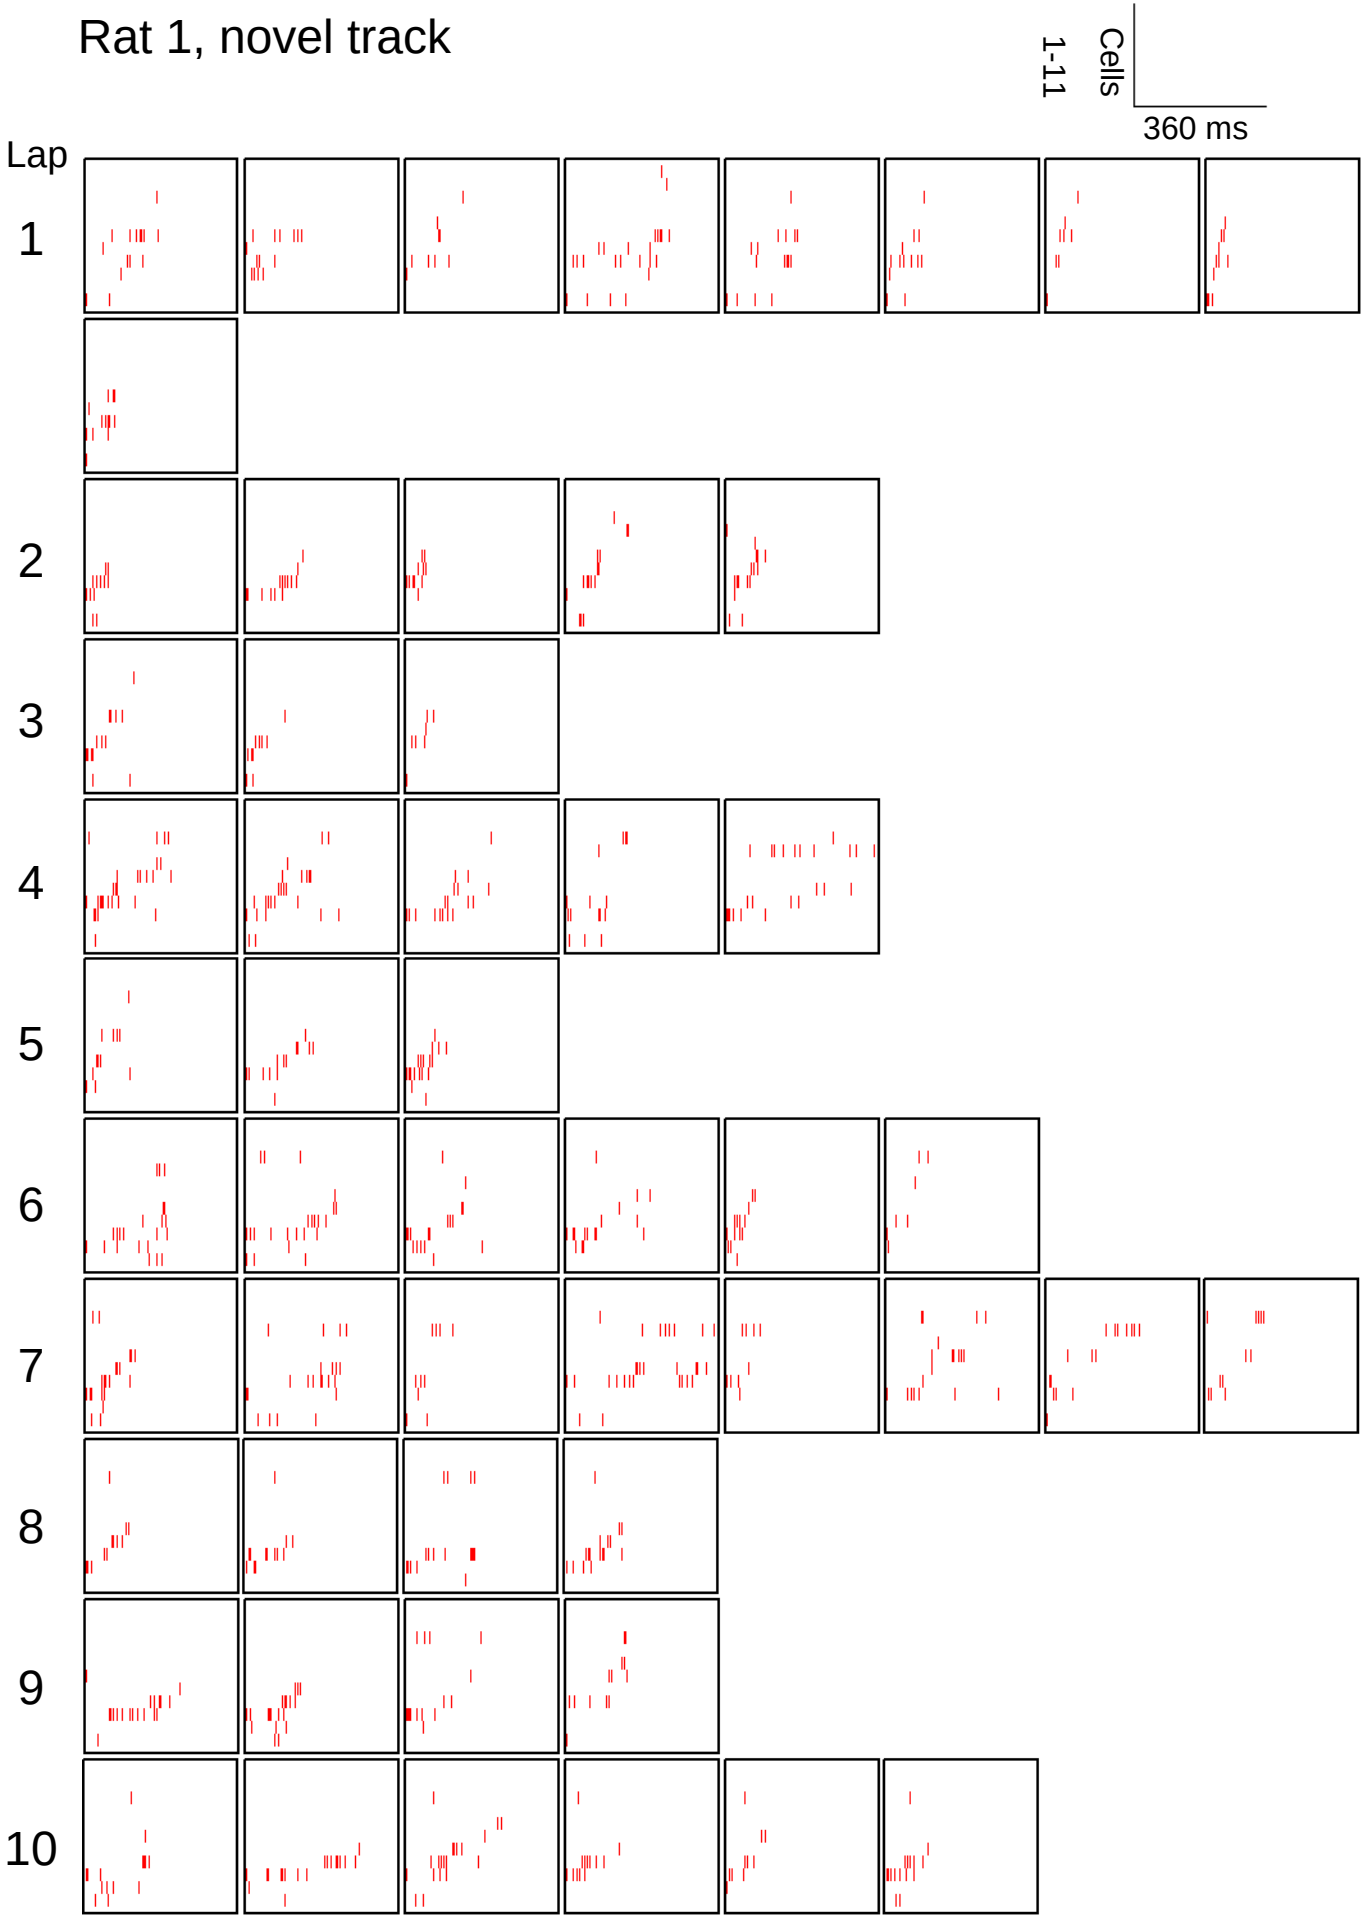
Rat 1, novel track
Cells
1-11
360 ms
Lap
1
2
3
4
5
6
7
8
9
10

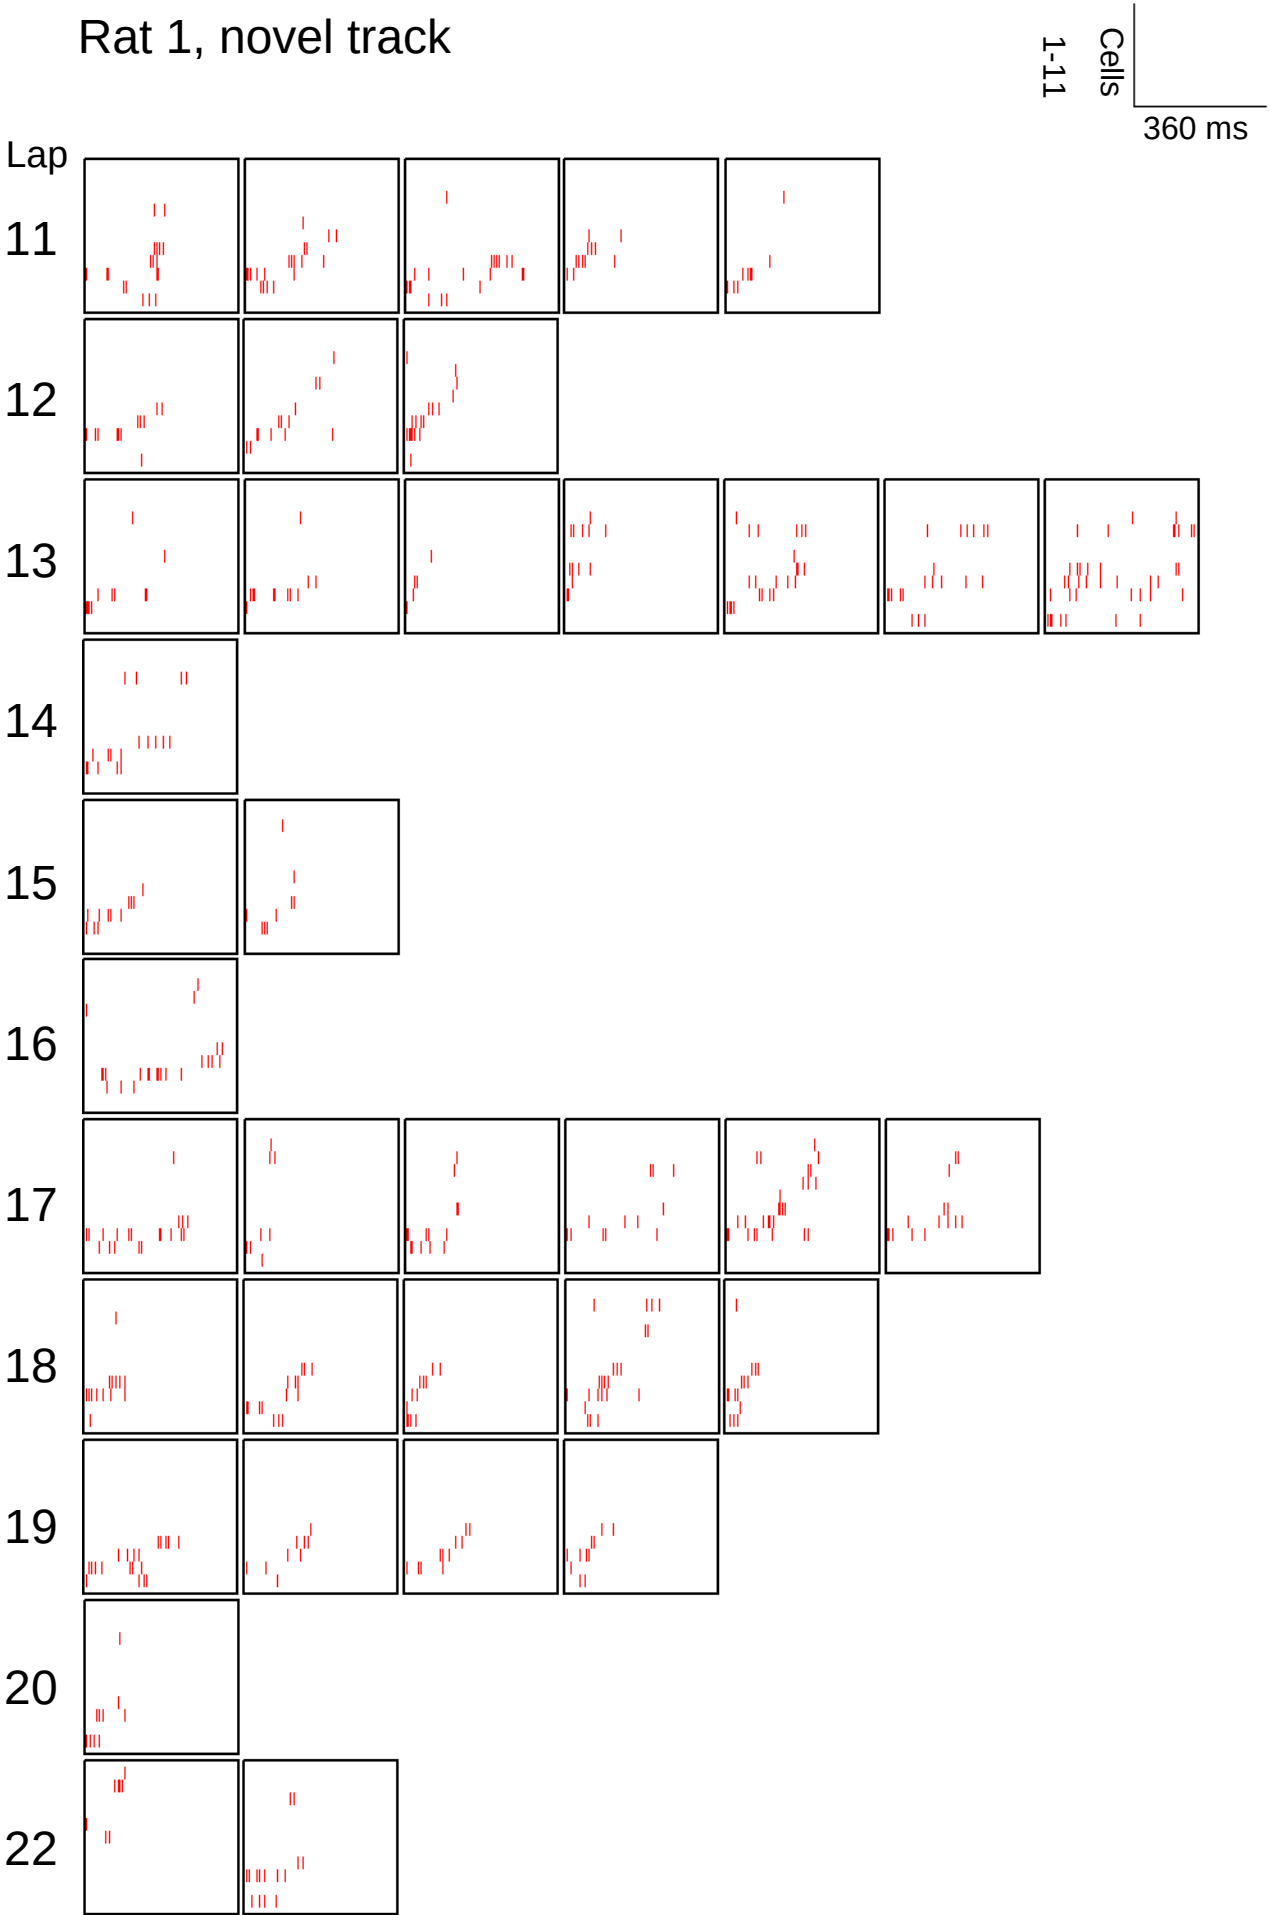
Rat 1, novel track
Cells
1-11
360 ms
Lap
11
12
13
14
15
16
17
18
19
20
22

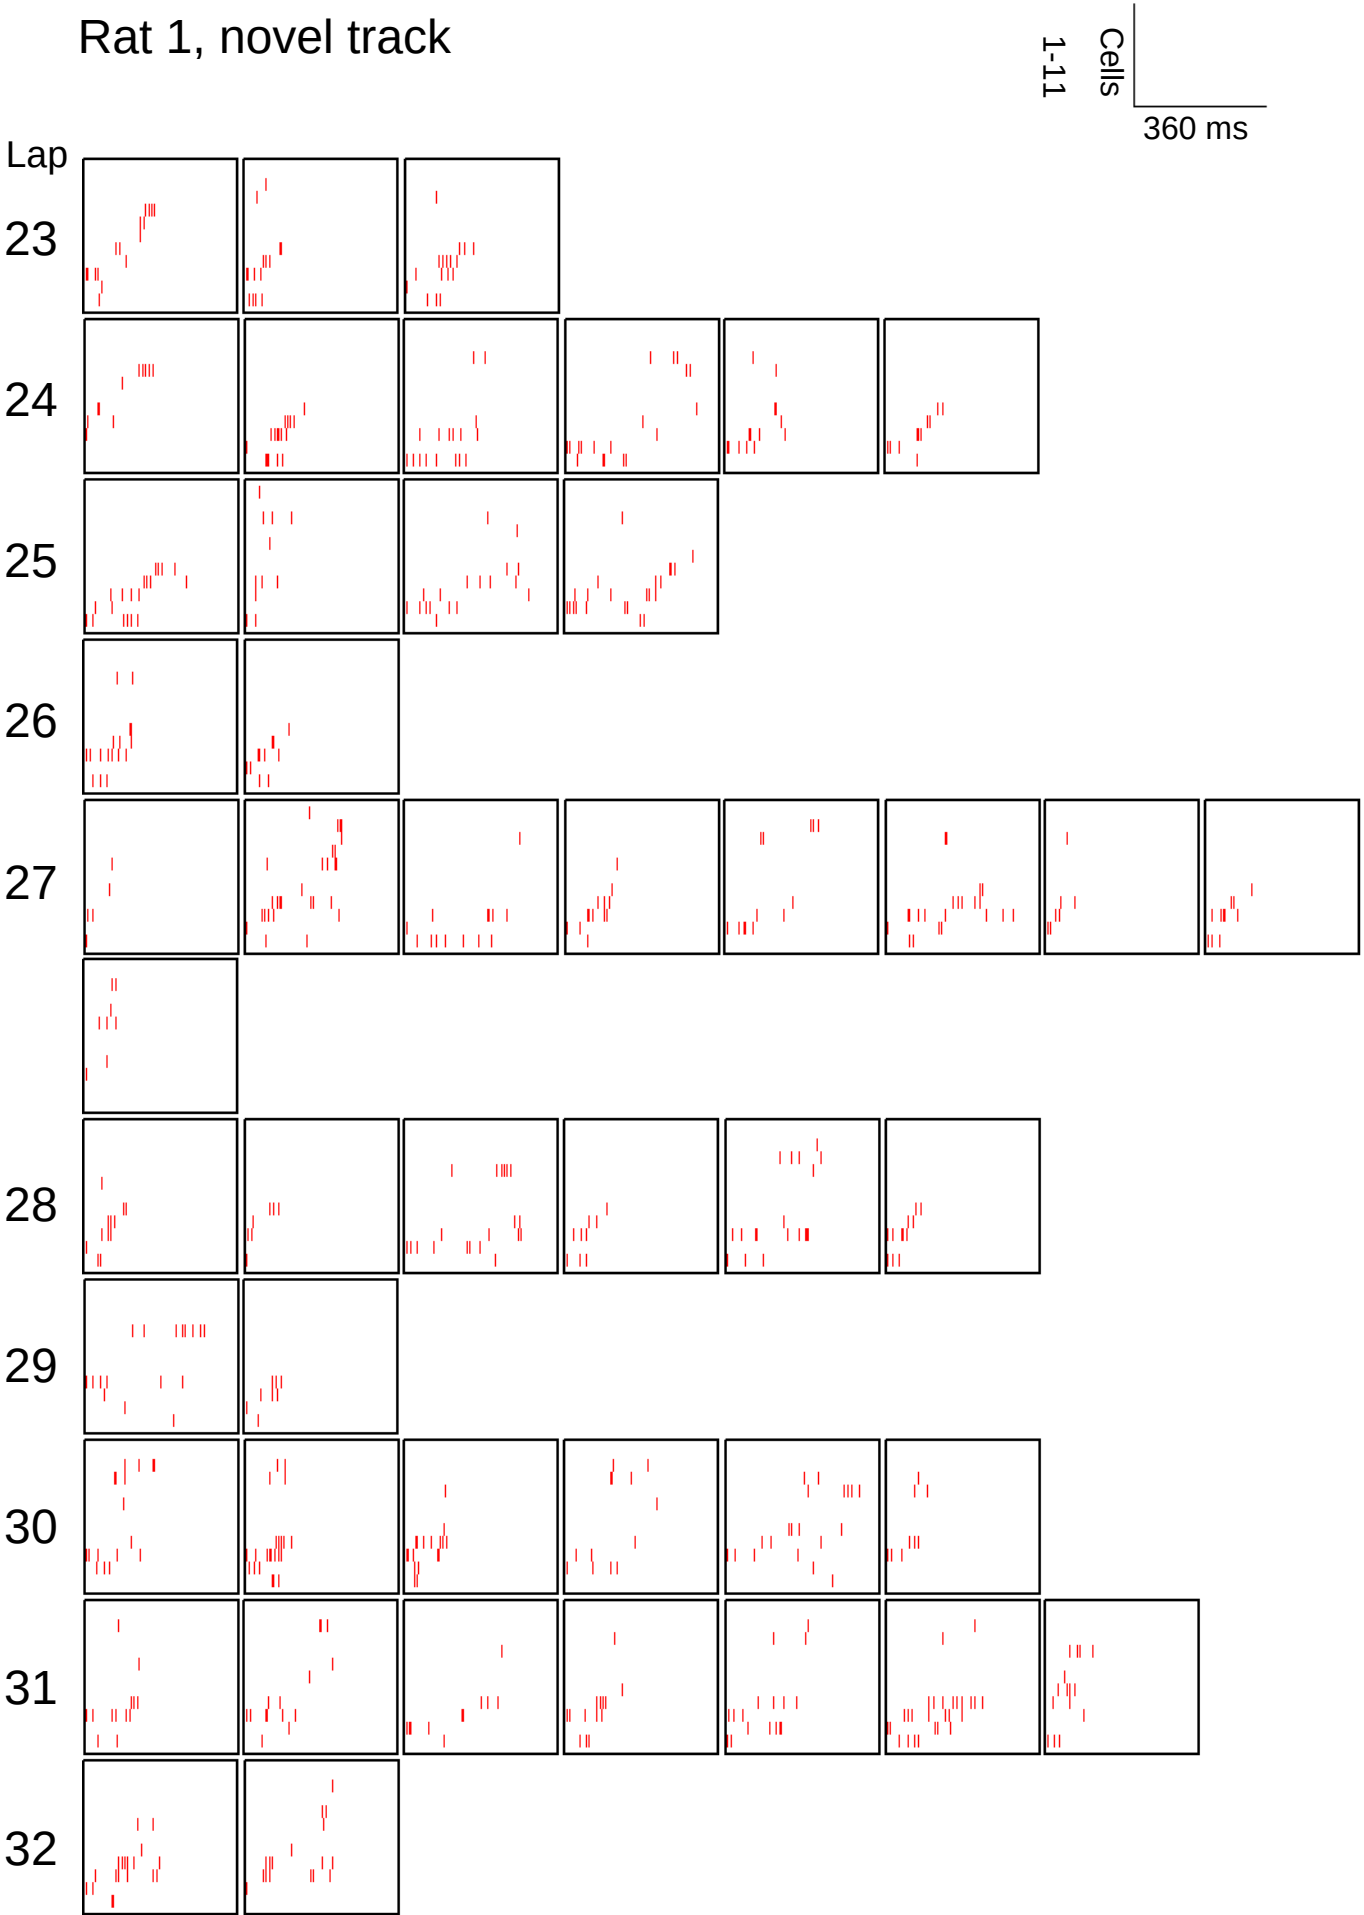
Rat 1, novel track
Cells
1-11
360 ms
Lap
23
24
25
26
27
28
29
30
31
32

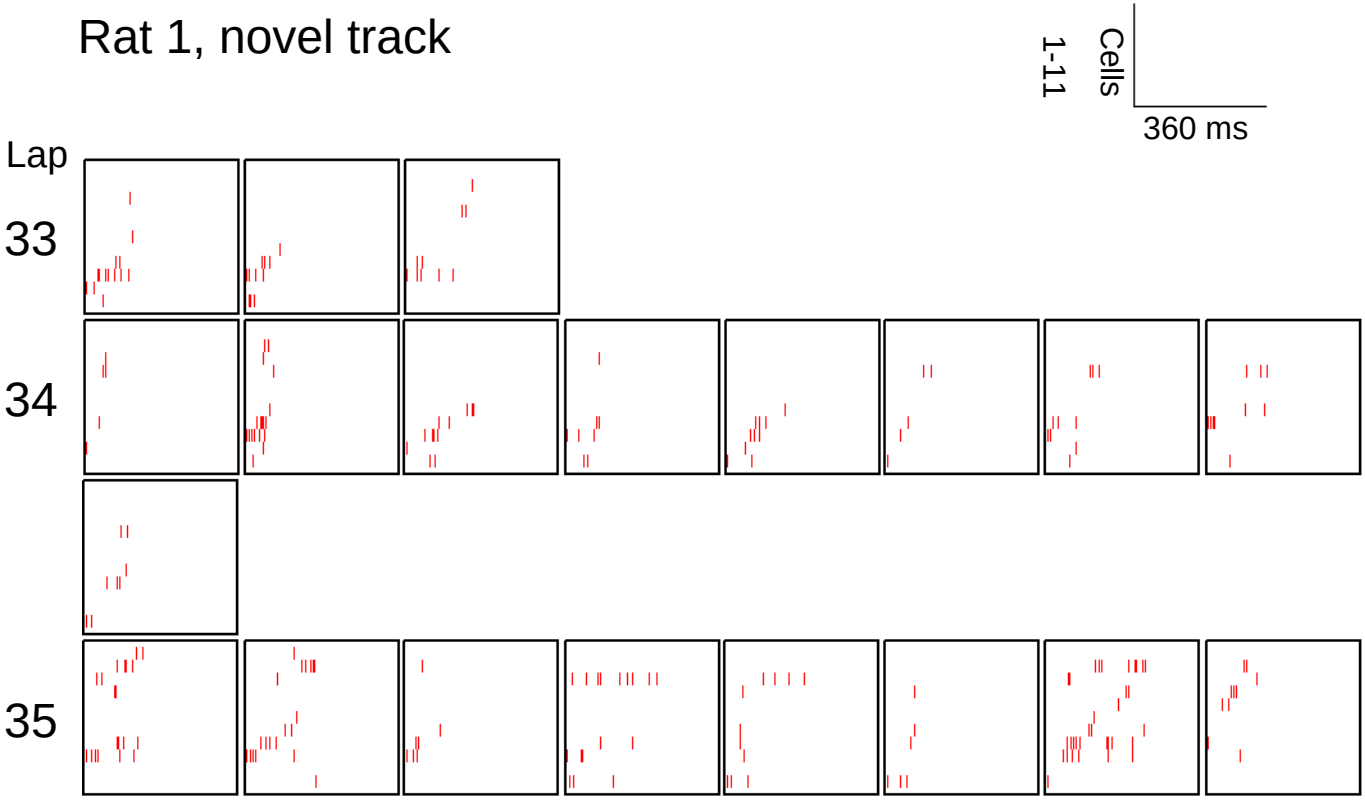
Rat 1, novel track
Cells
1-11
360 ms
Lap
33
34
35

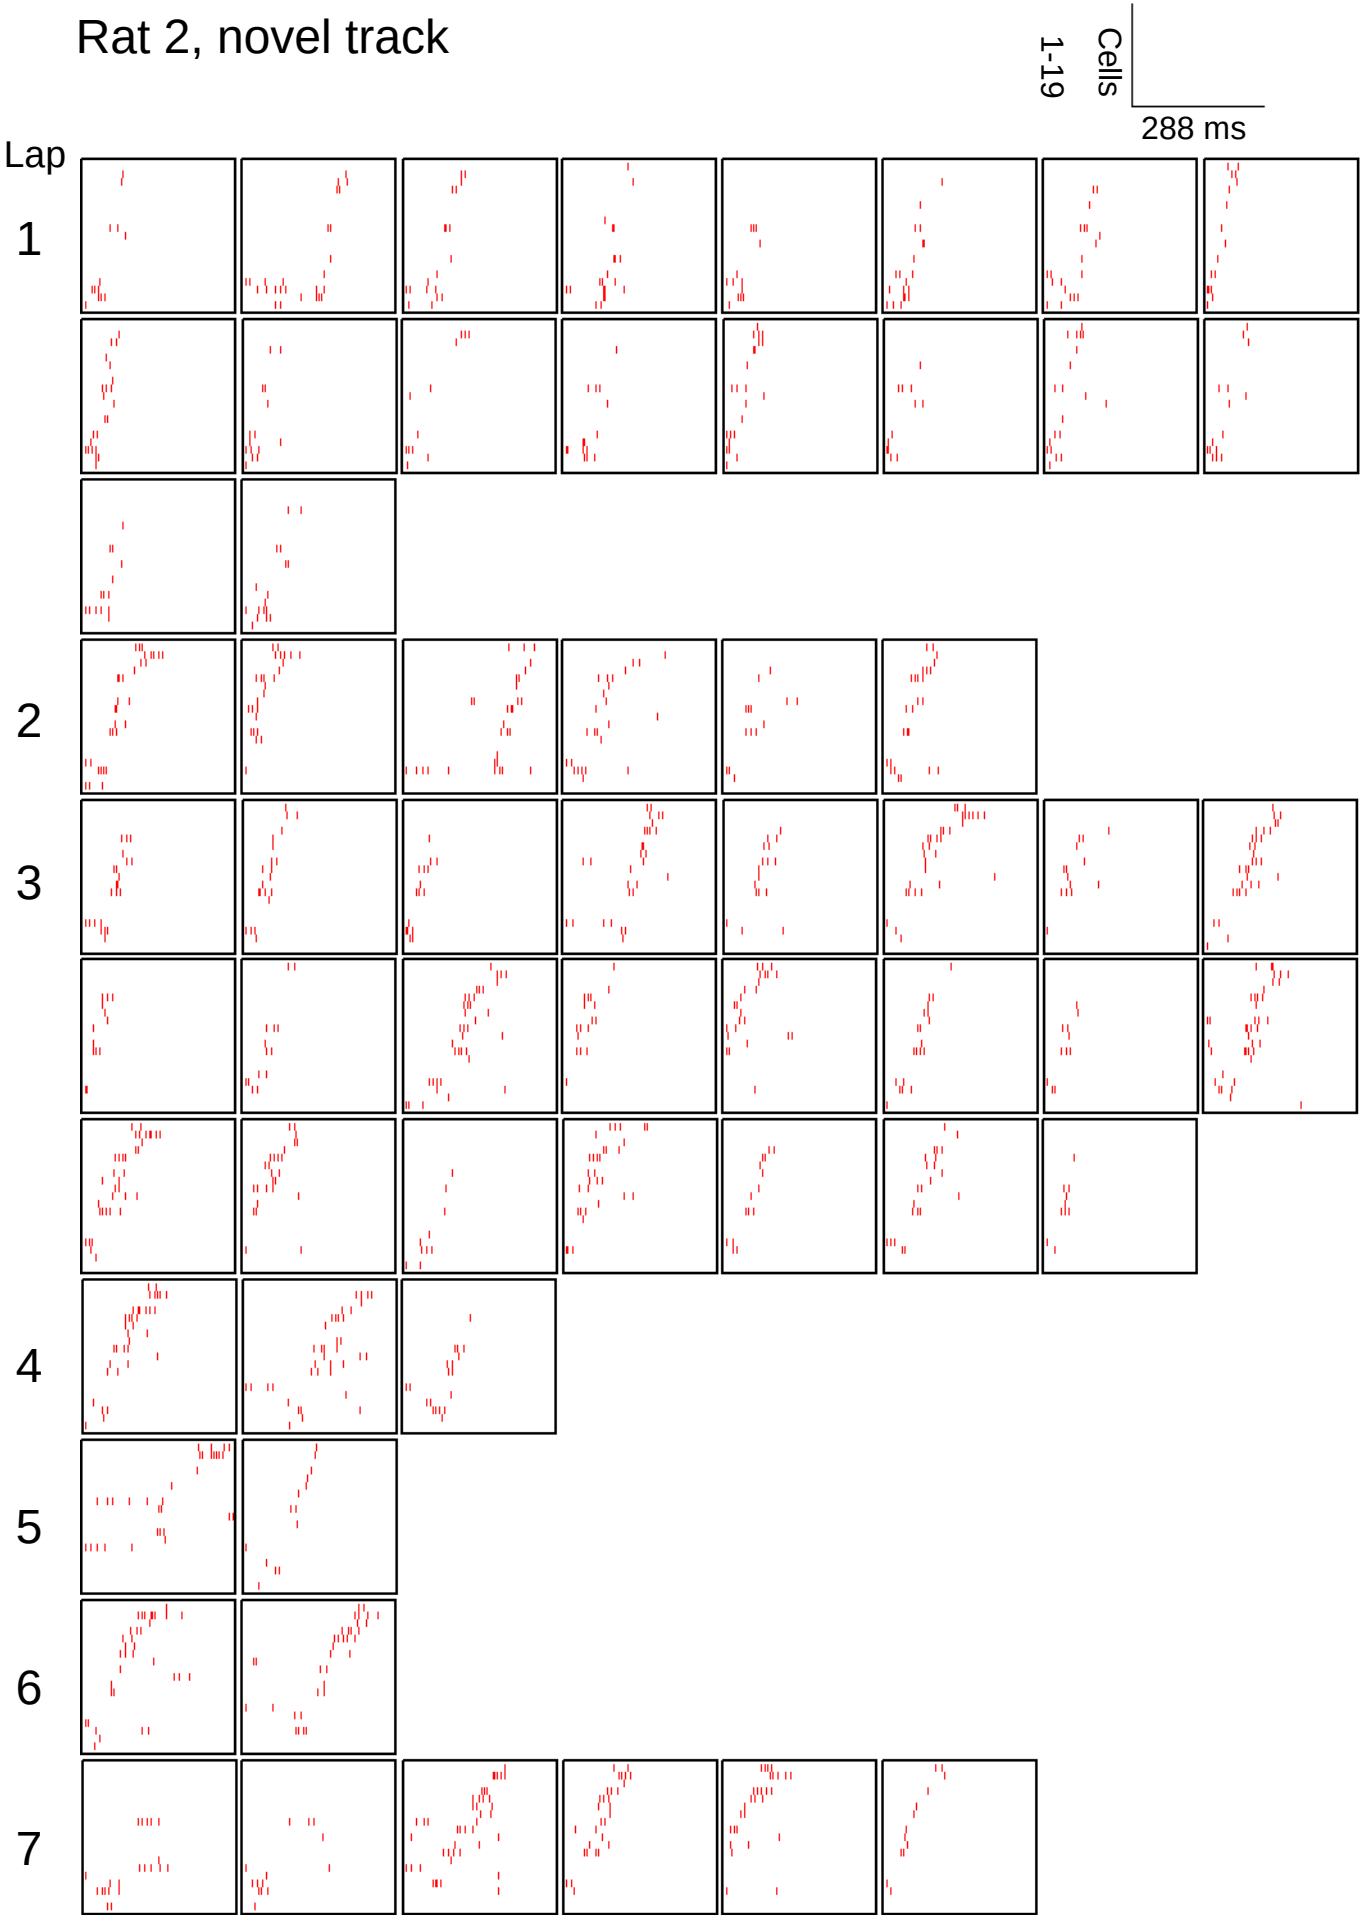
Rat 2, novel track
Cells
1-19
288 ms
Lap
1
2
3
4
5
6
7

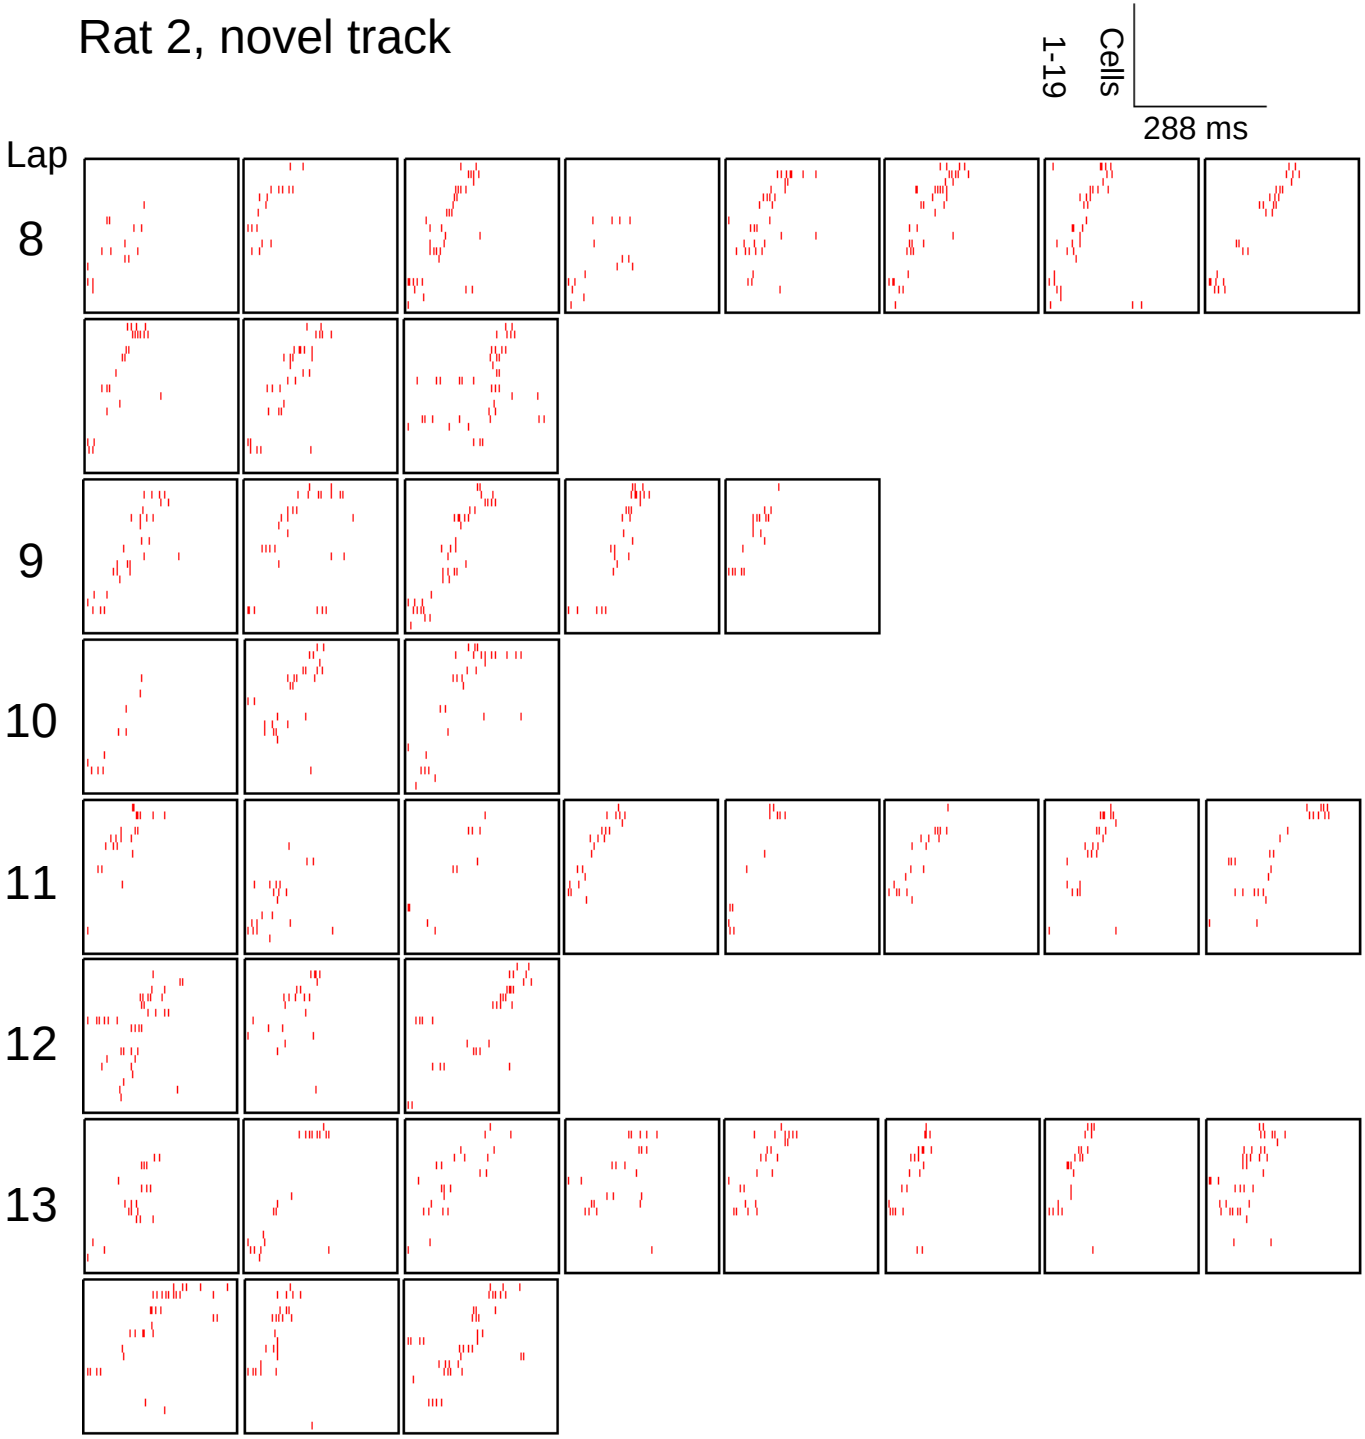
Rat 2, novel track
Cells
1-19
288 ms
Lap
8
9
10
11
12
13

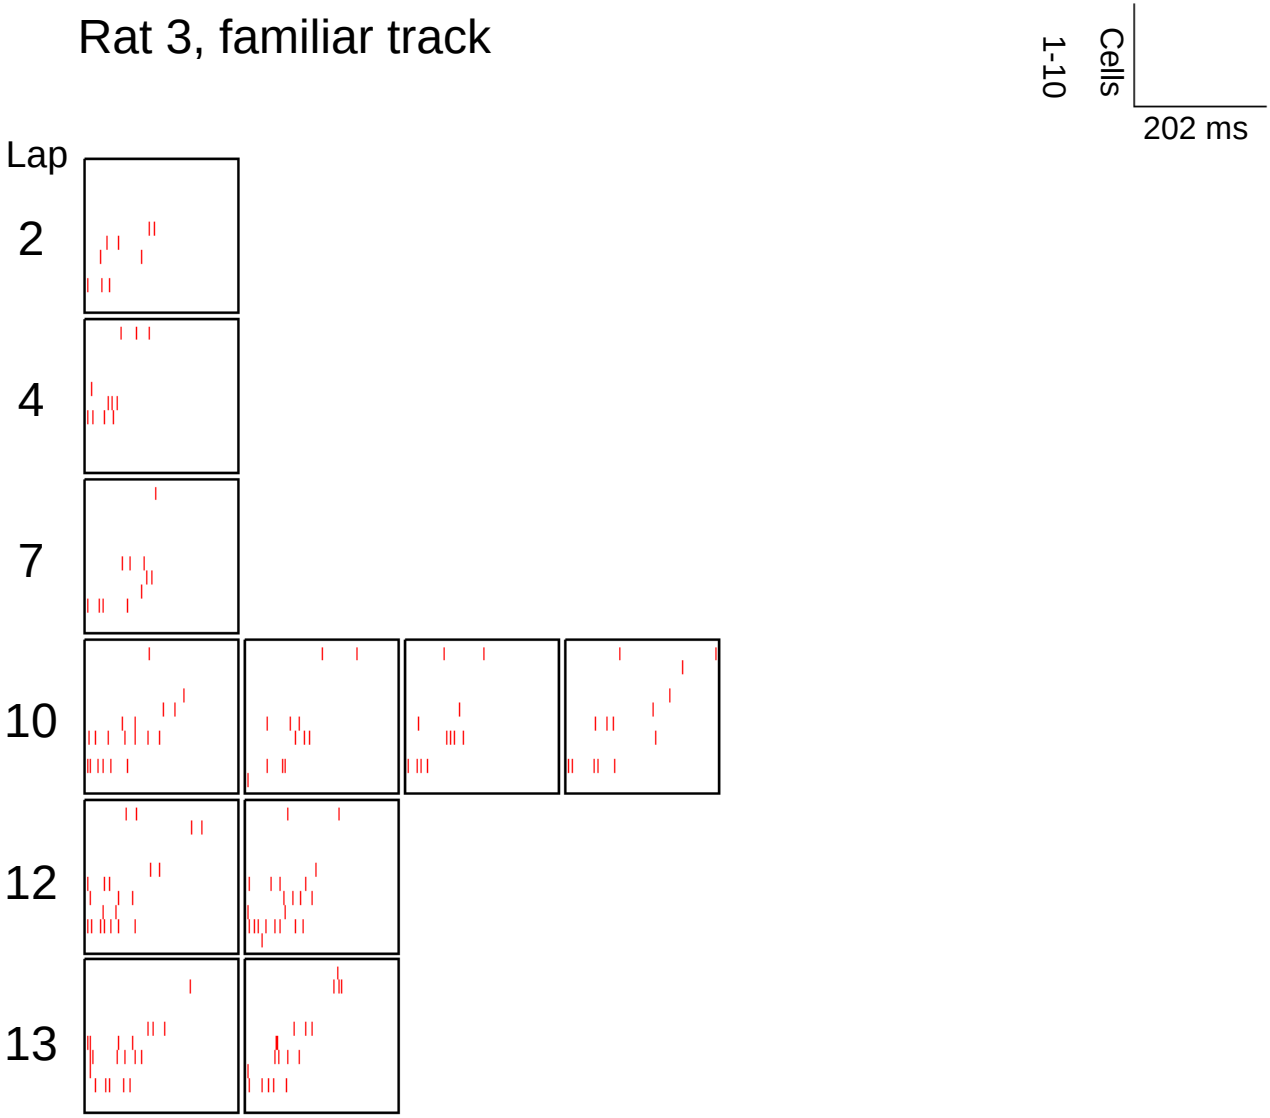
Rat 3, familiar track
Cells
1-10
202 ms
Lap
2
4
7
10
12
13

Rat 4, familiar track

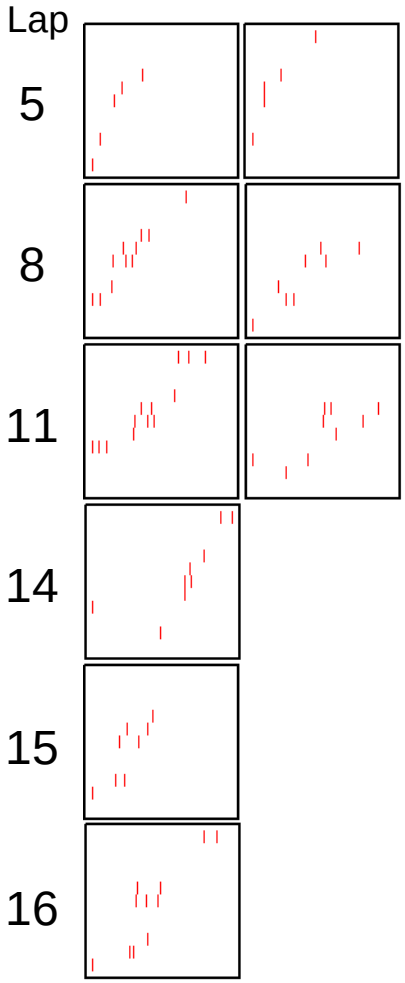

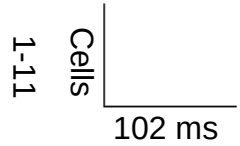

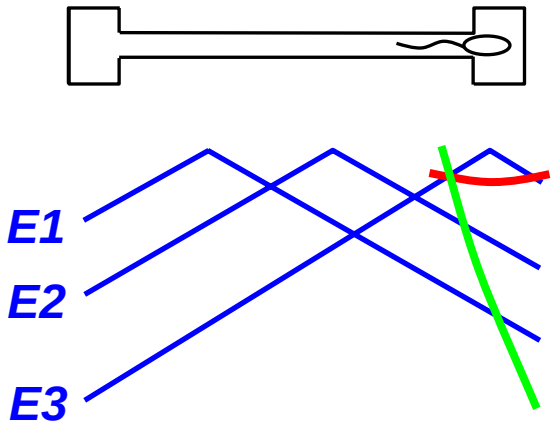

This figure represents a single moment in time, time t_1, when the rat is at the goal location at the end of the track, having just run along it. The blue cones (labelled E1, E2, E3) represent the rising and falling excitatory drives to three cells, as a function of position, **as they stand at time t_1**. At this time they incorporate two effects: a spatial tuning derived from their inputs, and an activity dependent bias, which reflects the fact that the cells have been activated during the running episode. Many possible mechanisms might provide this bias, including synaptic potentiation or intrinsic excitability increases. Cell activity in the model is then controlled by inhibition, which operates in two distinct modes. In one mode, during the non-sharp-wave state, inhibition is high (indicated by the red line), allowing only a small fraction of hippocampal place cells to fire (in this case only cell 3). In the other mode, during sharp waves, transient drops in inhibition (indicated by the green line) reveal the sub-threshold fields of the cells. Critically, as is made clear by the figure, the sub-threshold fields are revealed in reverse order, beginning with cell 3, and then cell 2, and then cell 1.

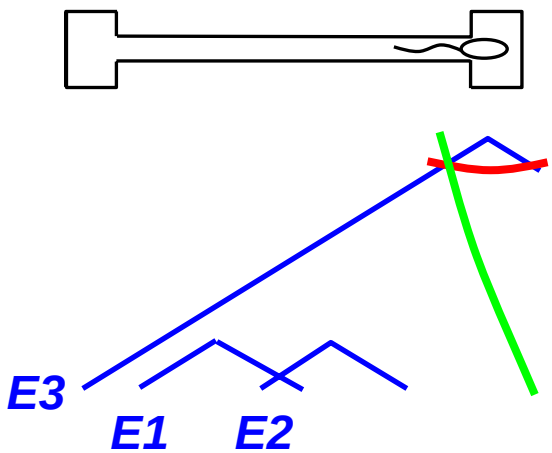

This second figure represents a different moment in time, time t_2, when the rat has been placed at the goal location, without having run through any of the locations preceding it. At this time, the excitatory drives to cells 1 and 2 reflect the spatial tuning of their inputs, but have not been biased because they have not yet been active. Cell 3 which is currently active is shown in the figure as having the bias. As before, high inhibition during the non-sharp-wave state (indicated by the red line) allows only cell 3 to fire. However, the transient drop in inhibition associated with the sharp wave (indicated by the green line) fails to reveal any activity from cells 2 and 3, because their excitatory input is too low.

a

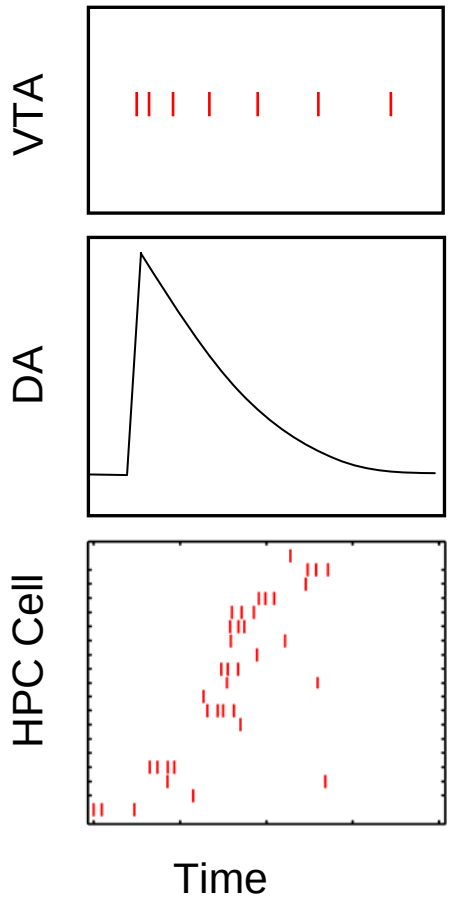

A simple demonstration of how reverse replay of hippocampal (HPC) place cells might support spatial learning. If the output of place cells is paired with a dopamine signal as shown in the second panel, those place cells with fields close to the goal will be maximally paired with reward, while the remaining cells will be paired with reward as a decreasing function of distance from the goal.

The model predicts that neurons responsible for the DA signal, such as neurons in the ventral tegmental area (VTA), would fire in bursts temporally synchronised with hippocampal sharp waves (as reported for ventral striatum by Pennarz et al, *J. Neurosci.* **24**(29):6446-6456. (2004)) and with an increasing spike interval profile, as shown.

b

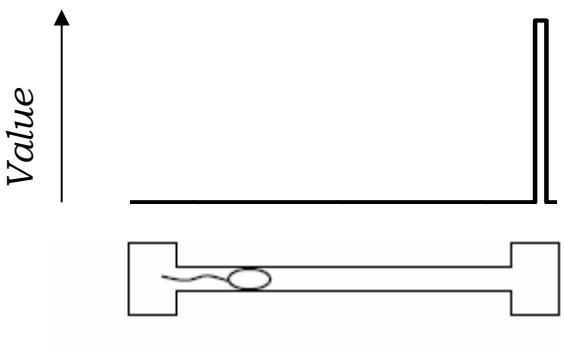

Before pairing, information about reward is available only at the goal location, and hence cannot guide action choice at locations far from the goal.

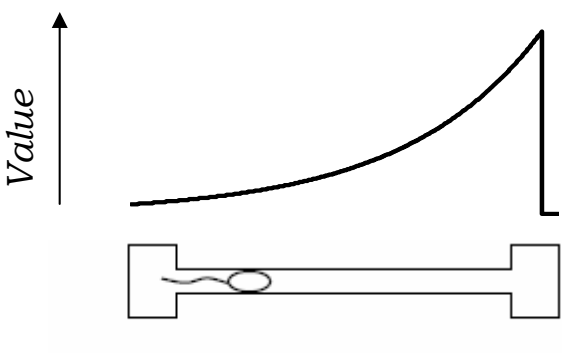

After pairing, gradient information about expected value is available at all track locations, since place cells are now associated with a value function.