

CLAIMS

1. A machine-learning model-based combination chunker/planner system, the chunker/planner system comprising:
 - i. a machine learning component (“Sequencer”) configured for receiving sequential input, for dividing the sequential input into one or more chunks, and for generating a plan corresponding to each chunk; and
 - ii. a second machine learning component (“Planner”) configured for pursuing a reward, for selecting from plans generated by the Sequencer those plans most closely associated with reaping the reward in a current state, and for activating a selected plan.
2. The chunker/planner system of claim 1 wherein the Sequencer is further configured for dividing the sequential input based on an element selected from the group consisting of: receiving an explicit end-of-sequence input, reaching a maximum size for a current chunk, receiving a reward to associate with a current chunk, and receiving an input whose value differs from expected values by more than a set threshold.
3. The chunker/planner system of claim 1 wherein generating a plan corresponding to a chunk comprises:
 - i. generating a declarative representation (“tonic”) associated with the entire chunk; and
 - ii. as each element of the chunk is examined in input sequence:
 1. querying the Planner for a complete plan consistent with the chunk as examined so far; and
 2. using a complete plan returned by the Planner, the tonic, a time-decaying context, and a most recently examined element in the chunk to predict a next element in the chunk.
4. The chunker/planner system of claim 1:
 - i. wherein the Planner is further configured for associating with a plan a change of state produced when the plan is activated; and
 - ii. wherein the Planner is further configured for pursuing a goal state, for selecting from plans generated by the Sequencer those plans most closely associated with accomplishing a change of state from a current state to a state closer to the goal

state, and for activating a selected plan.

5. The chunker/planner system of claim 1 further comprising an input buffer for the Sequencer; wherein the Sequencer is further configured to:
 - i. receive the sequential input into the input buffer;
 - ii. respond to a user command to discard the contents of the input buffer; and
 - iii. respond to a user command to train the Sequencer to turn the contents of the input buffer into a plan and to record it as a chunk.
6. The chunker/planner system of claim 1 further configured to:
 - i. receive a partial input;
 - ii. select a best match among existing plans that are consistent with the partial input;
 - iii. predict a result and reward from activating the selected plan; and
 - iv. activate the selected plan.
7. The chunker/planner system of claim 1 further configured to:
 - i. receive a partial input;
 - ii. infer a probability distribution of existing plans consistent with the partial input;
 - iii. predict a probability distribution of results and rewards from activating the consistent plans; and
 - iv. predict a next element in the input, the prediction based, at least in part, on the probability distribution.
8. The chunker/planner system of claim 1 wherein the Sequencer is a self organizing map.
9. The chunker/planner system of claim 1 or claim 2 wherein the Planner is a self organizing map.
10. In a computer-implemented system, a method for directing behaviour, the method comprising:
 - i. receiving, by a first machine learning component (“Sequencer”), sequential input;
 - ii. dividing the sequential input into one or more chunks;
 - iii. generating a plan corresponding to each chunk; and

- iv. pursuing, by a second machine learning component (“Planner”), a reward, wherein pursuing a reward comprises selecting from plans generated by the Sequencer those plans most closely associated with reaping the reward in a current state and activating a selected plan.

11. The method for directing behaviour of claim 10 further comprising:

- i. associating, by the Sequencer, with a plan a change of state produced when the plan is activated; and
- ii. pursuing, by the Planner, a goal state, wherein pursuing a goal state comprises selecting from plans generated by the Sequencer those plans most closely associated with accomplishing a change of state from a current state to a state closer to the goal state and activating a selected plan.

12. A system for controlling an application, the system comprising:

- i. a machine-learning model-based combination chunker/planner system comprising:

a first machine learning component (“Sequencer”) configured for receiving sequential input from the application, for dividing the sequential input into one or more chunks, and for generating a plan corresponding to each chunk; and

a second machine learning component (“Planner”) configured for pursuing a reward, for selecting from plans generated by the Sequencer those plans most closely associated with reaping the reward in a current state, and for activating a selected plan by communicating with the application;

- ii. wherein the Sequencer is further configured for associating with a plan a change of state produced when the plan is activated; and
- iii. wherein the Planner is further configured for pursuing a goal state, for selecting from plans generated by the Sequencer those plans most closely associated with accomplishing a change of state from a current state to a state closer to the goal state, and for activating a selected plan by communicating with the application.

13. The system of claim 12 wherein the controlled application is selected from the group consisting of: an industrial process, a manufacturing process, an online planning/collaboration application, and an online service avatar.

14. The system of claim 12 wherein the Sequencer is a self organizing map.
15. The system of claim 12 wherein the Planner is a self organizing map.
16. A machine-learning model-based system configured for detecting frequently occurring subsequences in sequential input and representing the subsequences as a whole, the system comprising: a neural-network self-organizing map (“Sequencer”) configured for receiving the sequential input, for dividing the sequential input into one or more subsequences, and for generating a plan corresponding to each subsequence.
17. The system of claim 16 wherein the Sequencer is further configured for dividing the sequential input based on an element selected from the group consisting of: receiving an explicit end-of-sequence input, and receiving an input whose value differs from expected values by more than a set threshold.
18. The system of claim 16 wherein generating a plan corresponding to a chunk comprises:
generating a declarative representation (“tonic”) associated with the entire chunk; and
as each element of the chunk is examined in input sequence:
using the tonic, a time-decaying context, and a most recently examined element in the chunk to predict a next element in the chunk.
19. The system of claim 16 further configured to:
receive as a partial input a fragment of a sequence;
select a best match among existing plans that are consistent with the partial input;
predict a likely next element in the sequential input from activating the selected plan; and
activate the selected plan.
20. The system of claim 16 further configured to:
receive as a partial input a fragment of a sequence;
infer a probability distribution of existing plans consistent with the partial input;
predict a probability distribution of likely next elements in the sequential input from activating the consistent plans; and
predict a next element in the input, the prediction based, at least in part, on the probability distribution.