On Biological Inspirations for Computer Science

Dr. Andrzej (AJ) Bieszczad

California State University

Channel Islands

aj.bieszczad@csuci.edu

Outline

- Biological inspirations
- Human and machine problem solving
- An artificial cortical column
- Architecture of the Neurosolver
- Learning in the Neurosolver
- Running rats in mazes with the Neurosolver
- Rat’s dilemma: multiple choices
- The context-based guided search
- Some conclusions
- TBDitF

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms
- Ants behavior → Swarm intelligence

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms
- Ants behavior → Swarm intelligence
- Nervous system → Neural Networks

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms
- Ants behavior → Swarm intelligence
- Nervous system → Neural Networks
- Neuron → Adaline

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms
- Ants behavior → Swarm intelligence
- Nervous system → Neural Networks
- Neuron → Adaline
- Immune system → Self-healing networks

...

Biological inspirations for computer science

- Evolution/Genetics → Genetic Algorithms
- Ants behavior → Swarm intelligence
- Nervous system → Neural Networks
- Neuron → Adaline
- Immune system → Self-healing networks

. . .

- Network of hypercolumns → Neurosolver

Human problem solving

Human problem solving

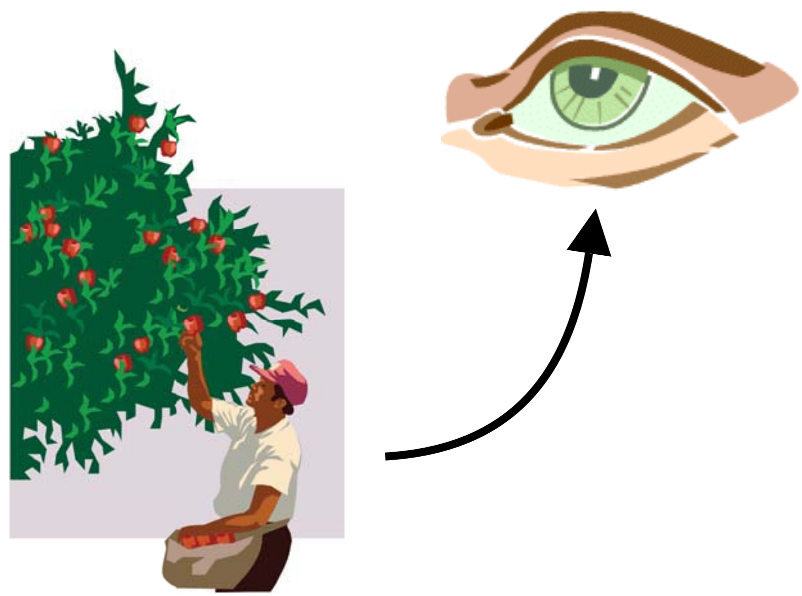

Human problem solving

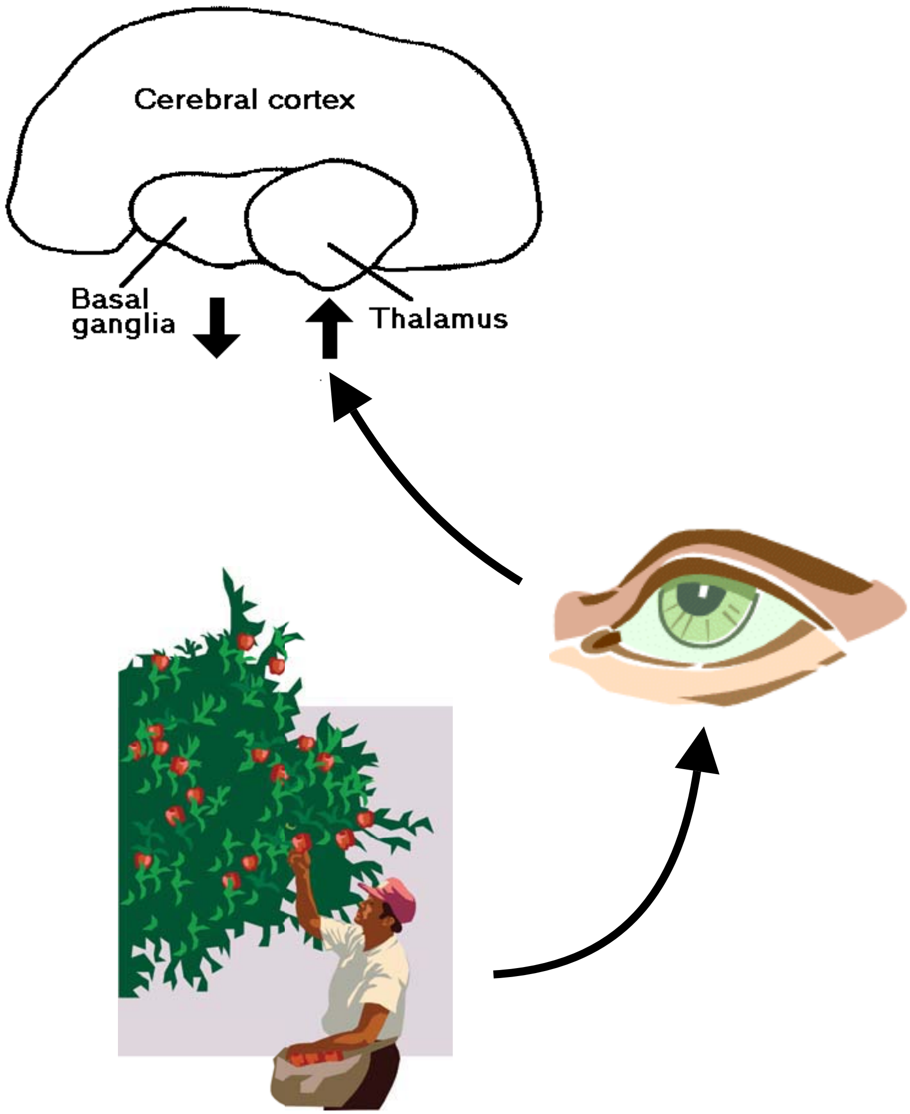

Human problem solving

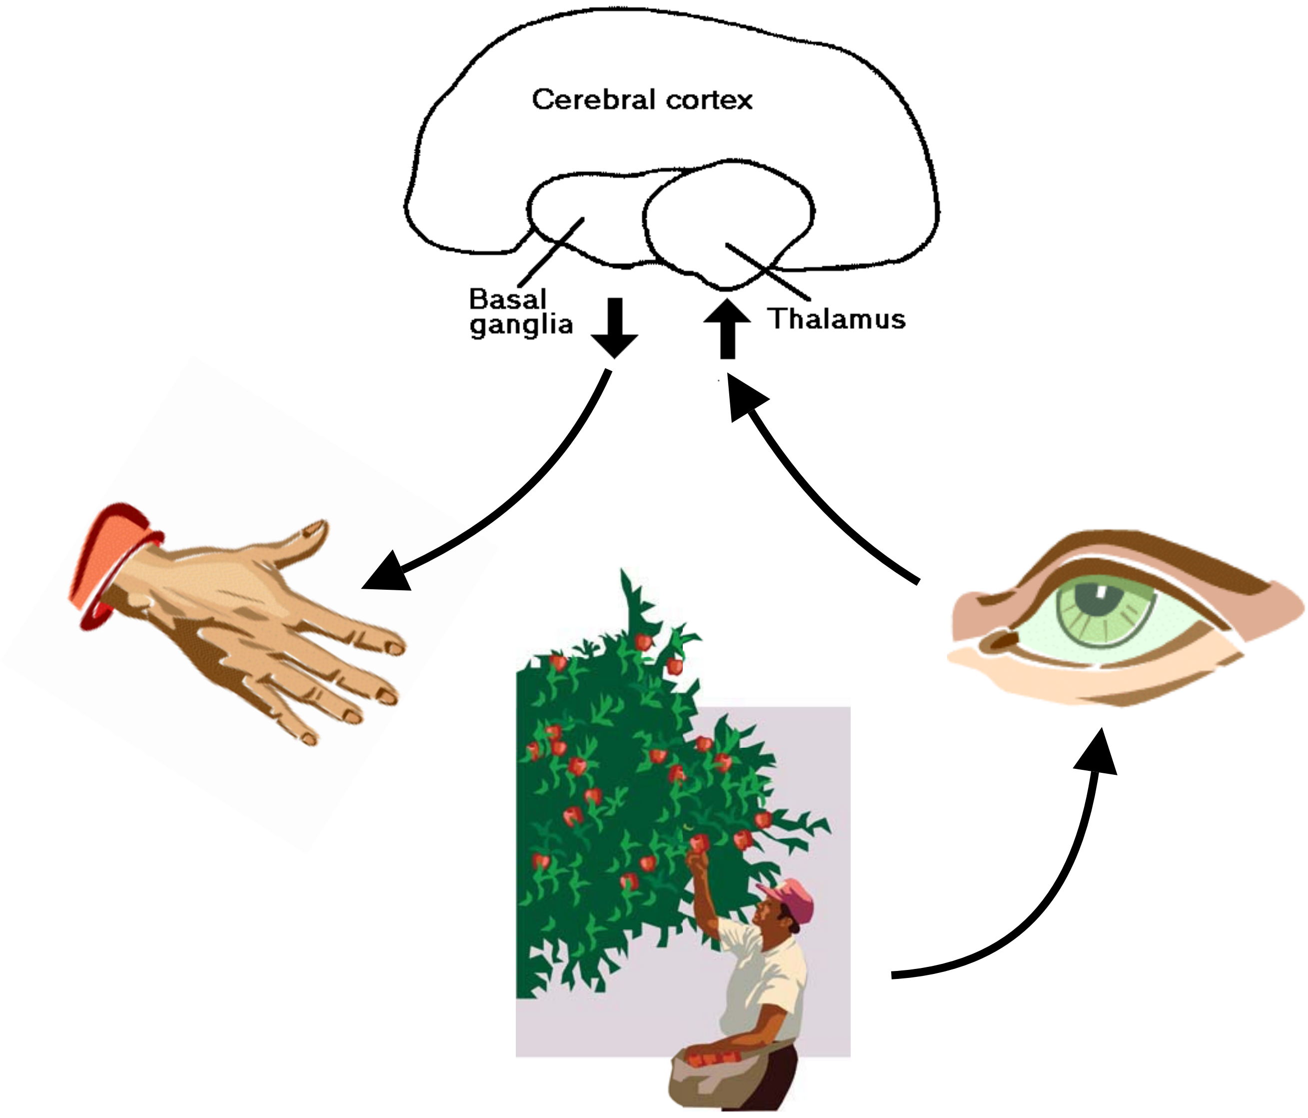

Human problem solving

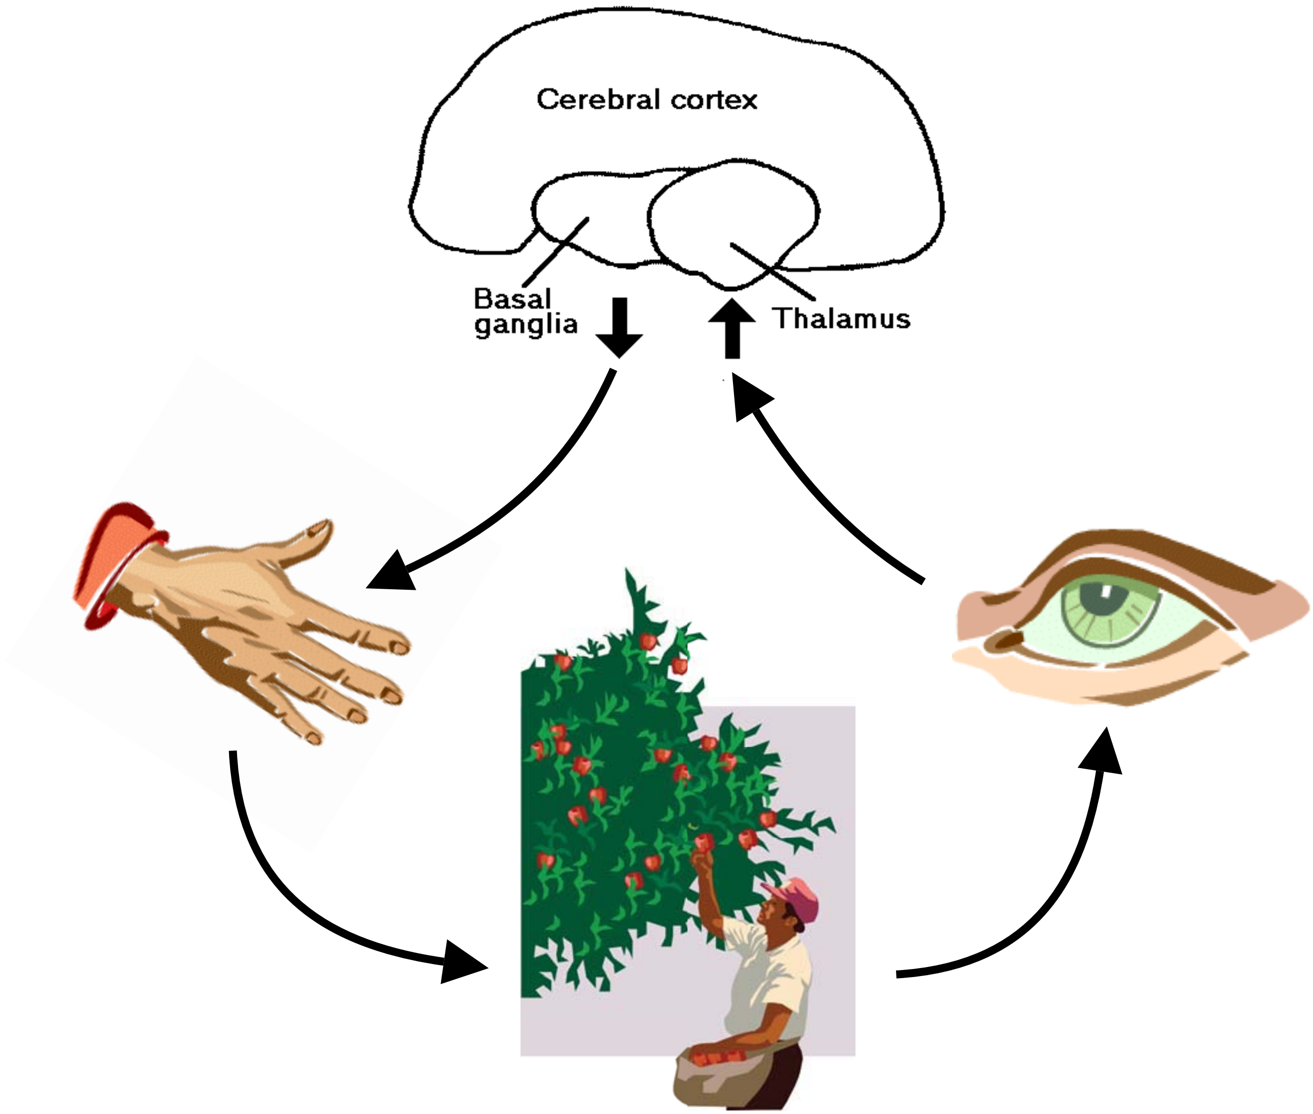

Modeling the real world

picking an apple

Modeling the real world

picking an apple

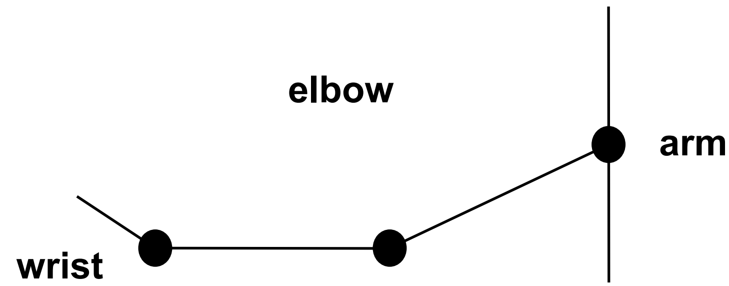

Modeling the real world

picking an apple

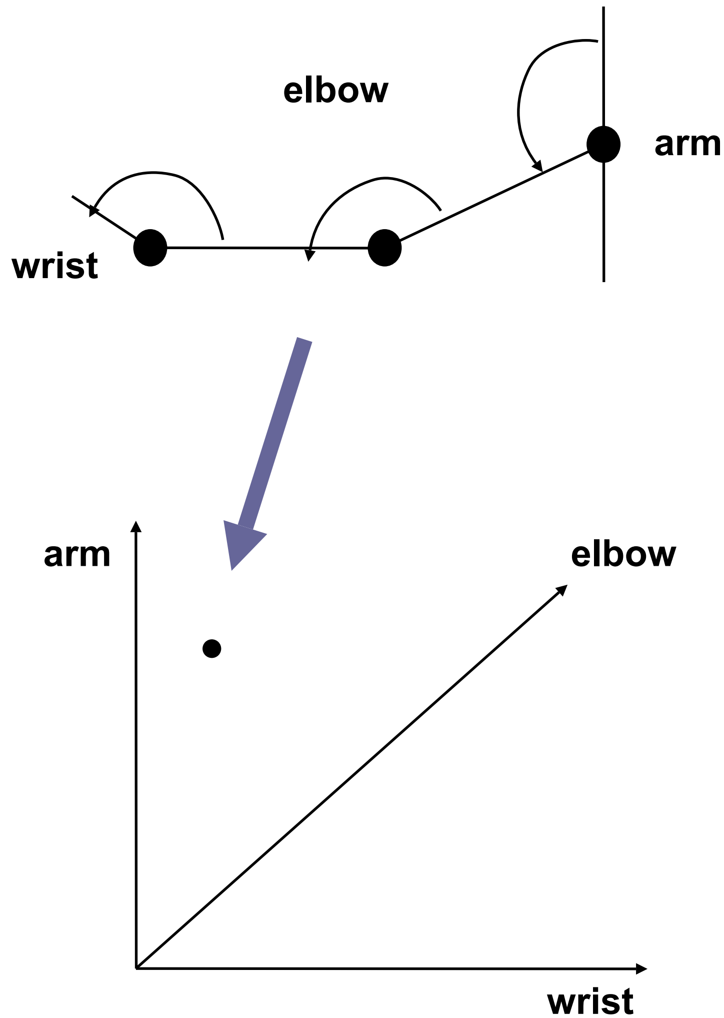

Modeling the real world

picking an apple

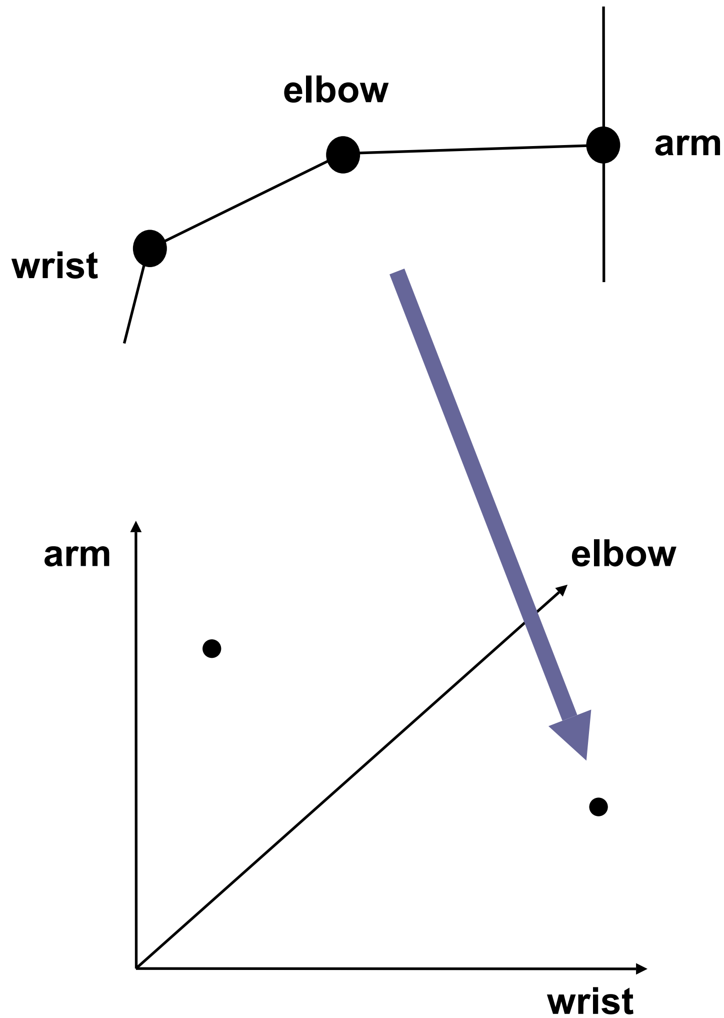

Modeling the real world

picking an apple

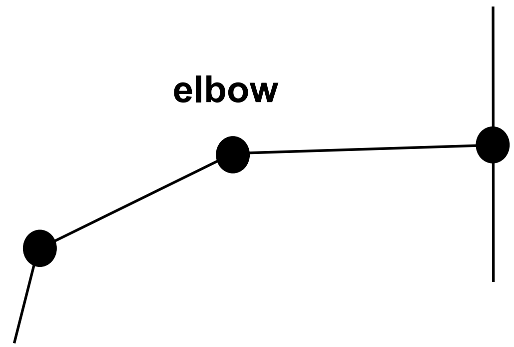

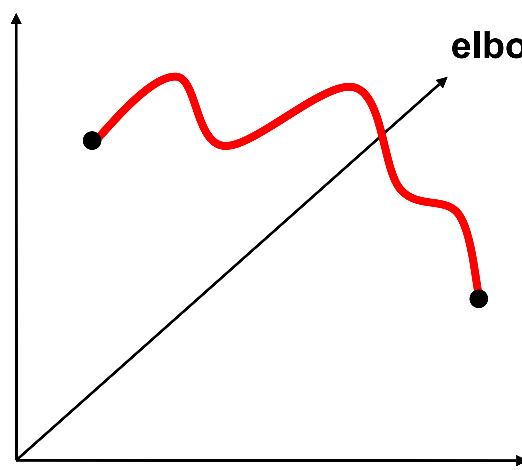

Solving a puzzle problem

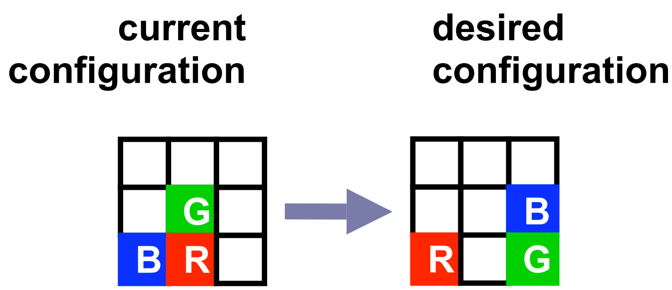

Solving a puzzle problem

current configuration

desired configuration

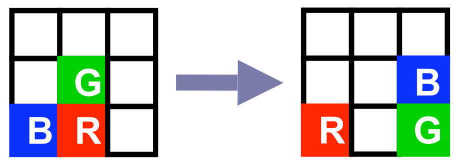

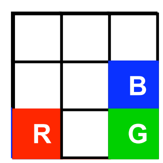

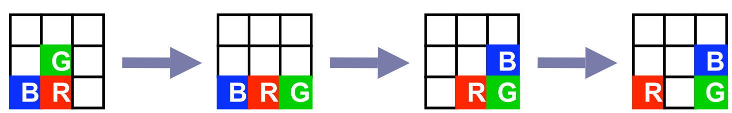

A search tree for a block rearrangement problem

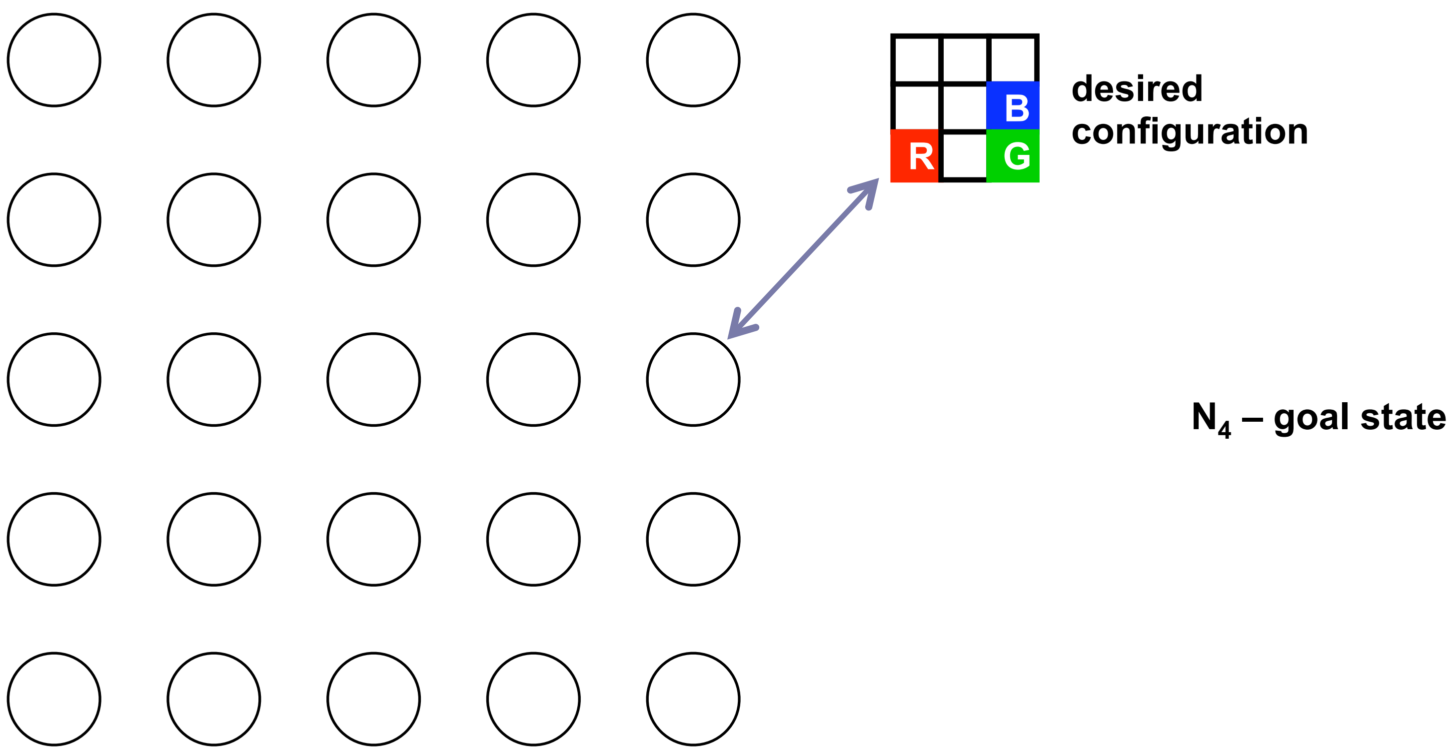

A search tree for a block rearrangement problem

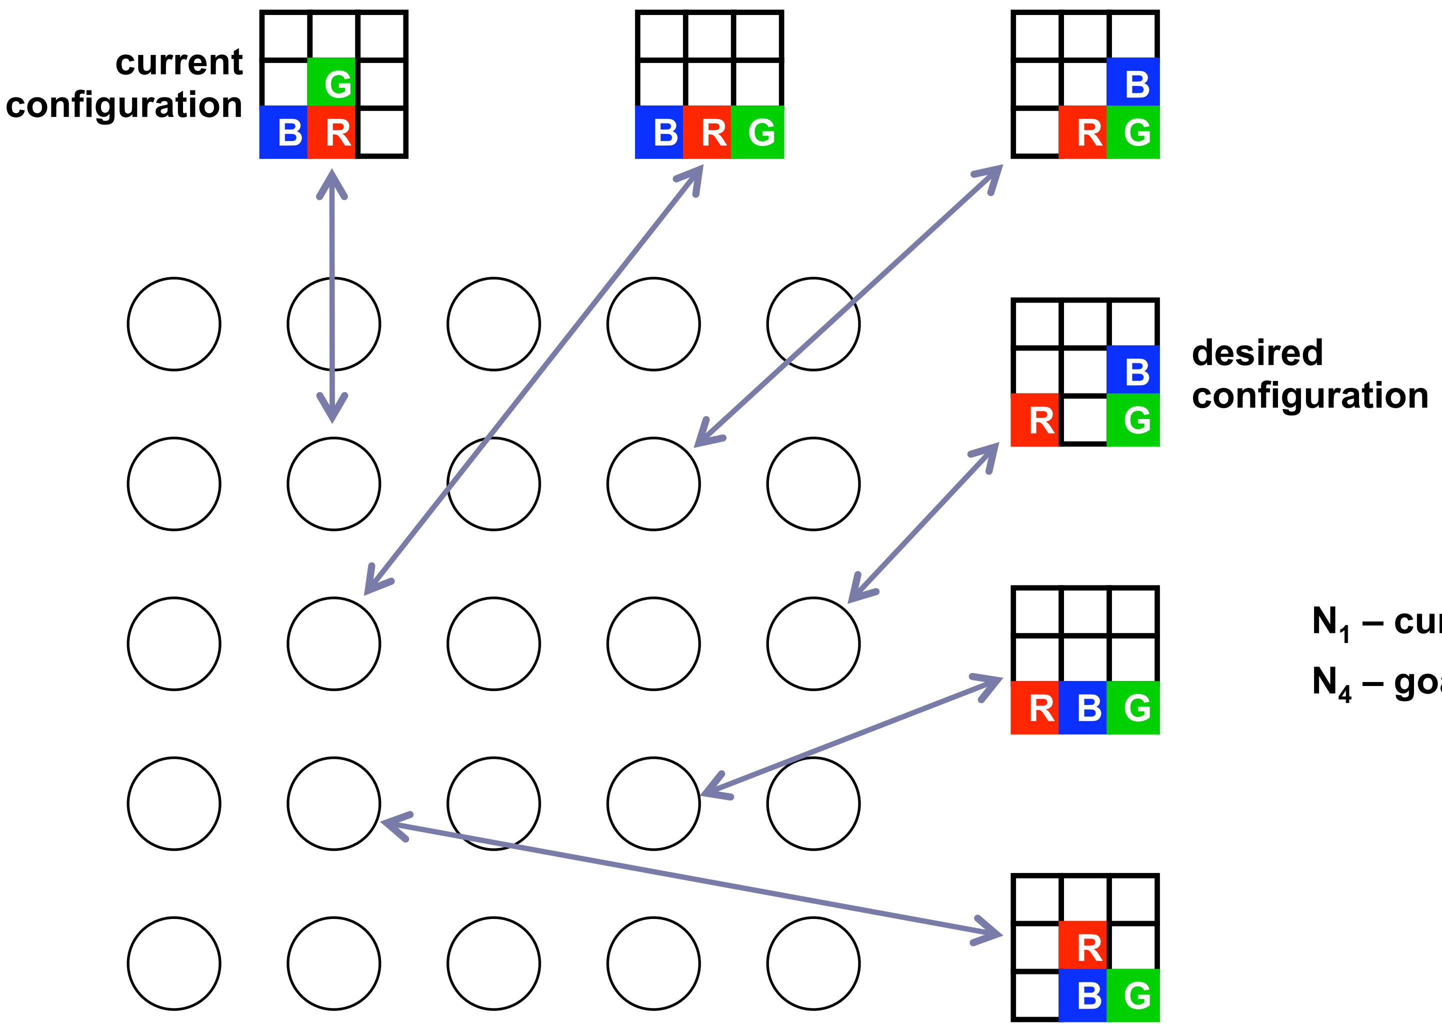

N_1 – current state

N_4 – goal state

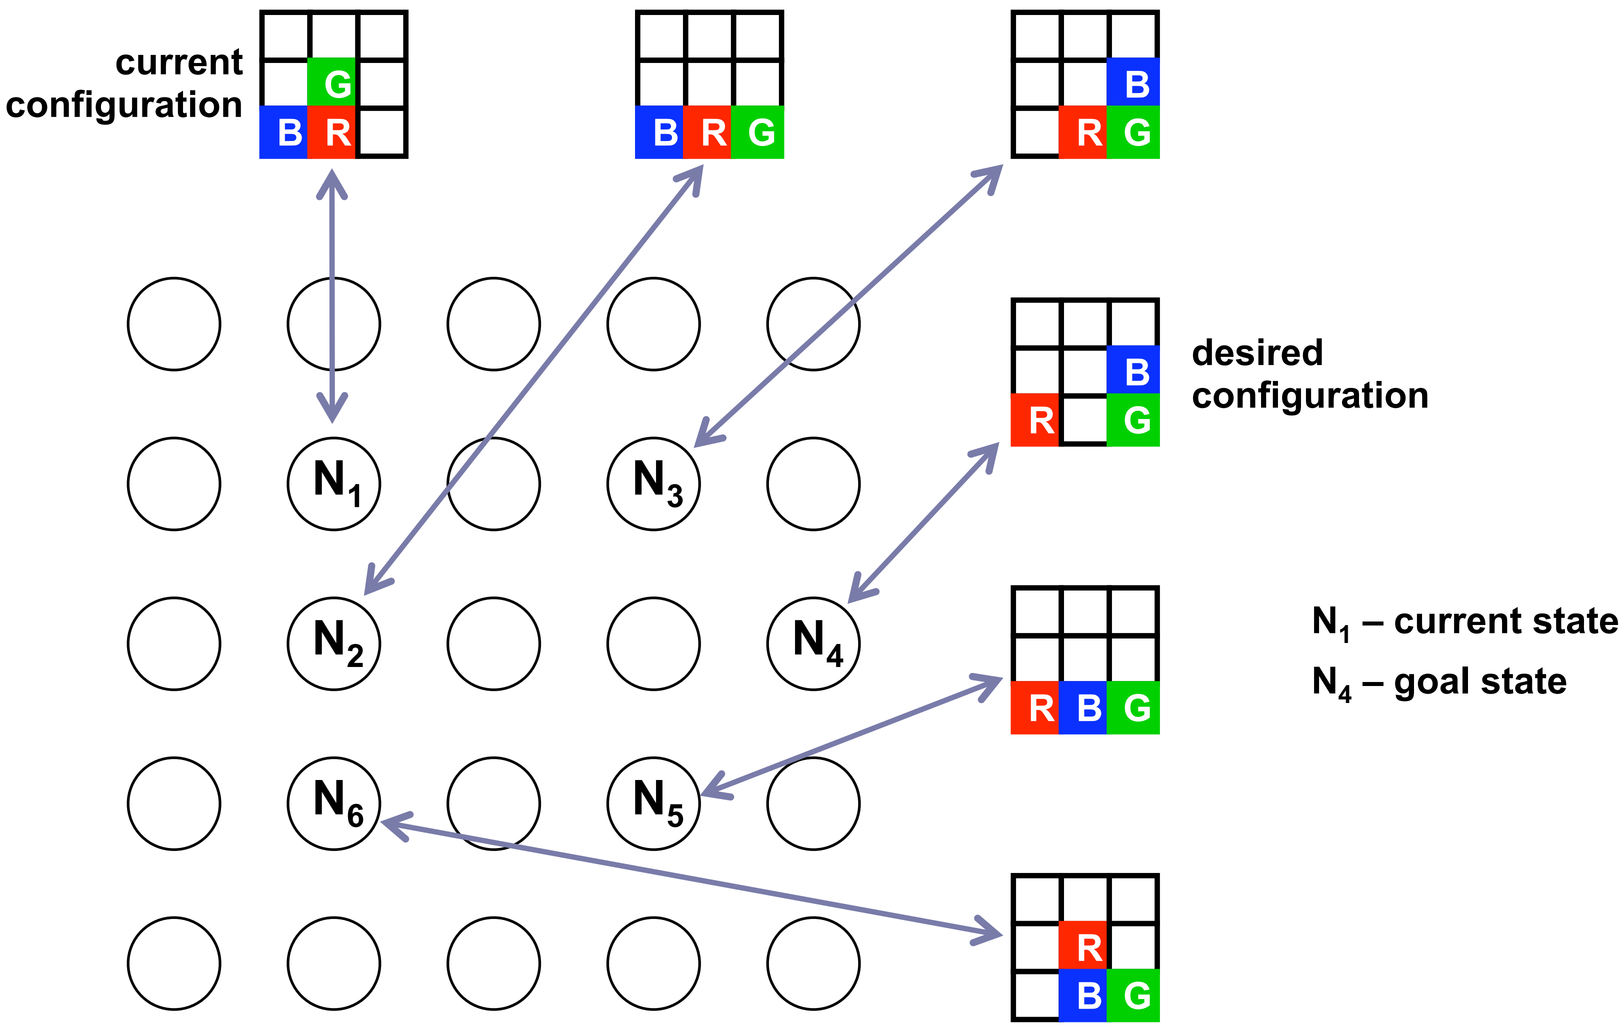
current
configuration
G
B
R
B
R
G
B
R
G
desired
configuration
B
R
G
R
B
G
R
B
G
N1
N2
N3
N4
N5
N6
N1 – current state
N4 – goal state

A search tree for a block rearrangement problem

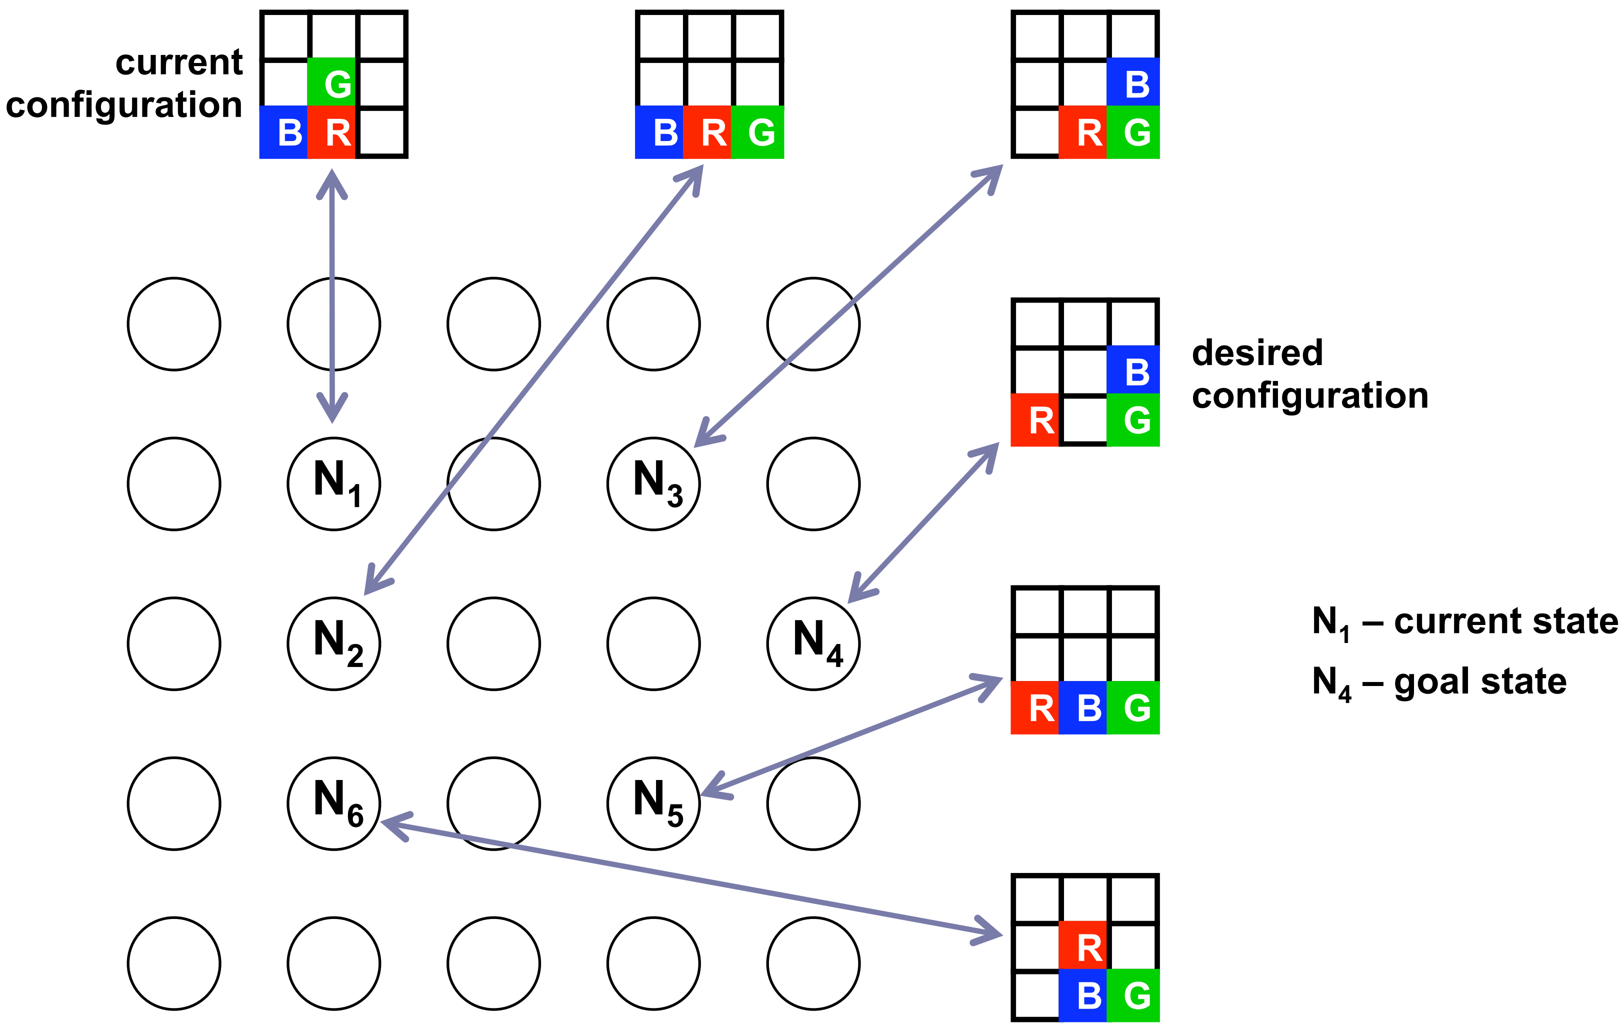

A search tree for a block rearrangement problem

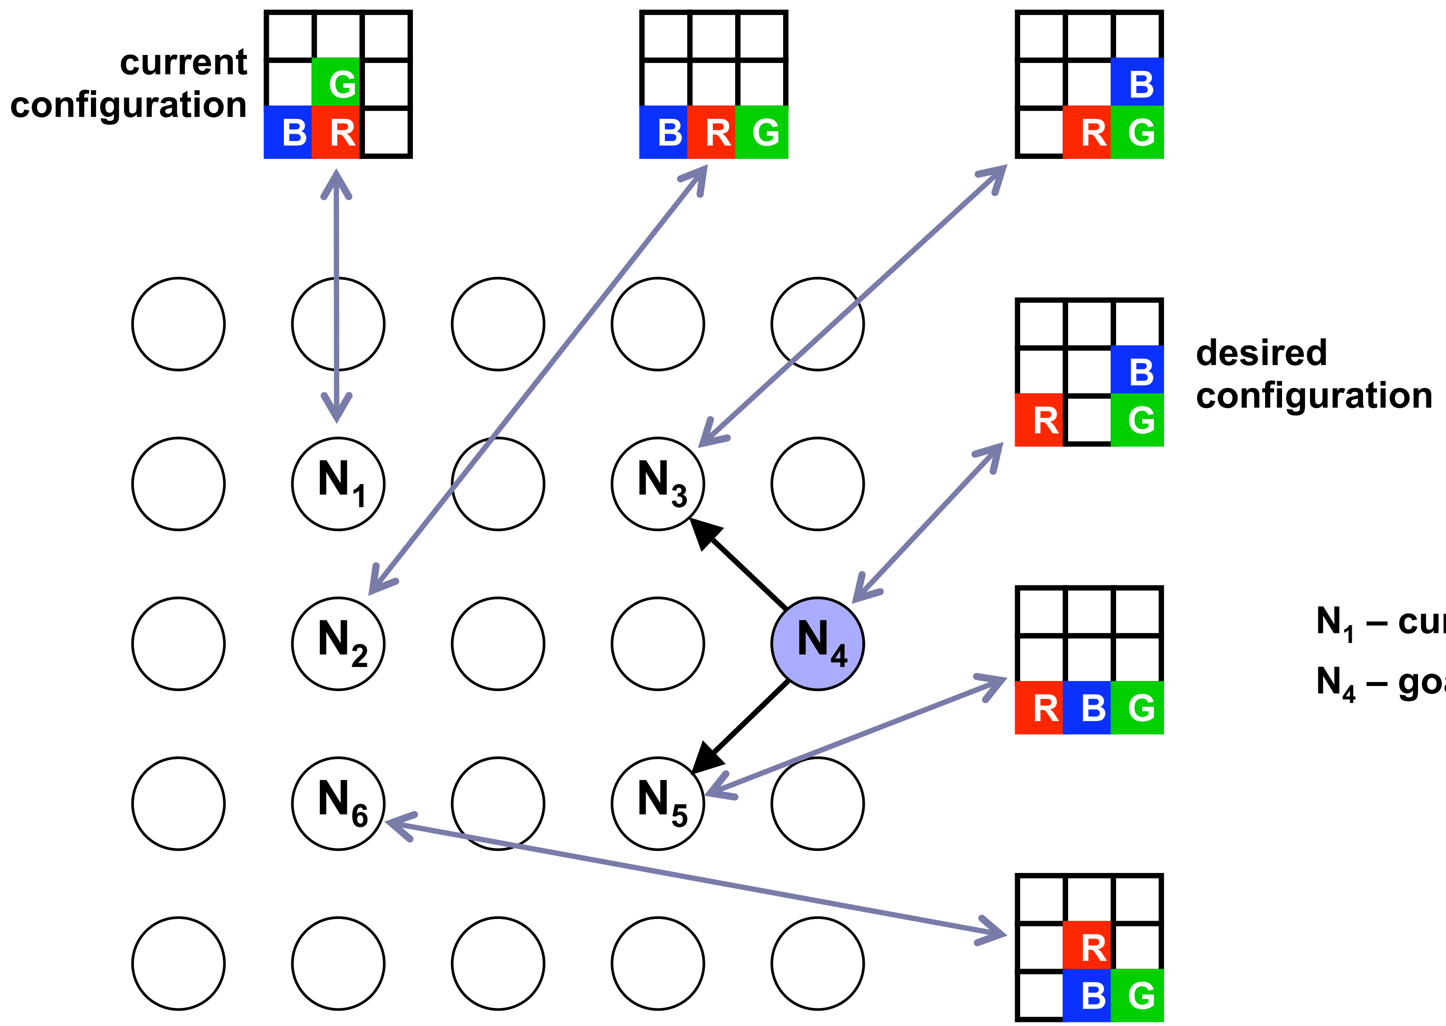

N_1 – current state

N_4 – goal state

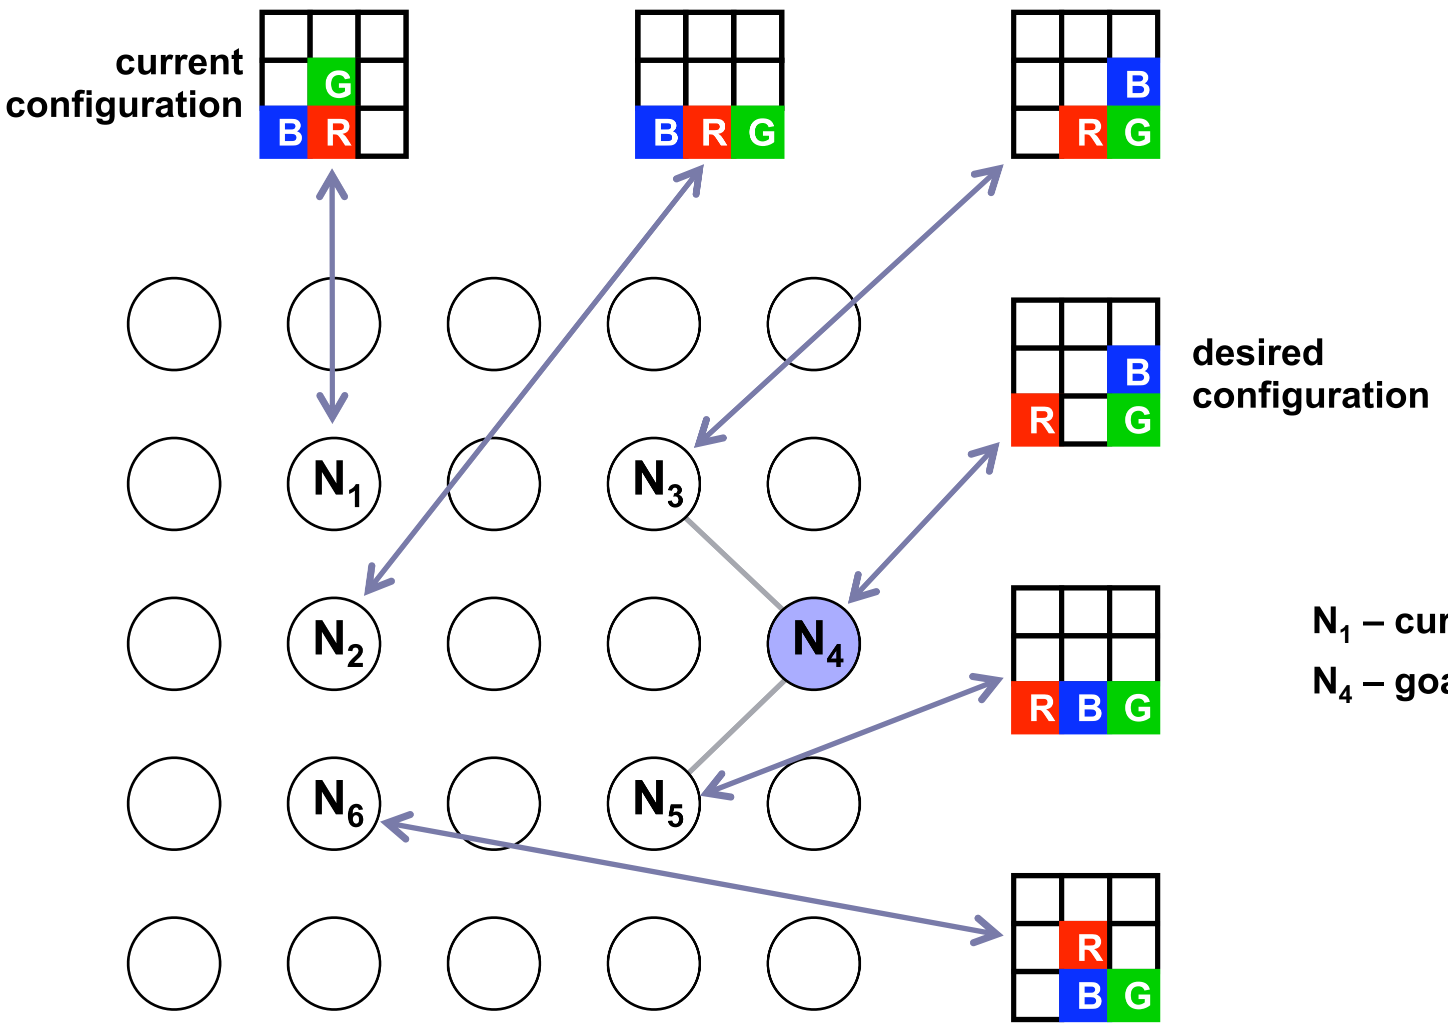

N_1 – current state

N_4 – goal state

A search tree for a block rearrangement problem

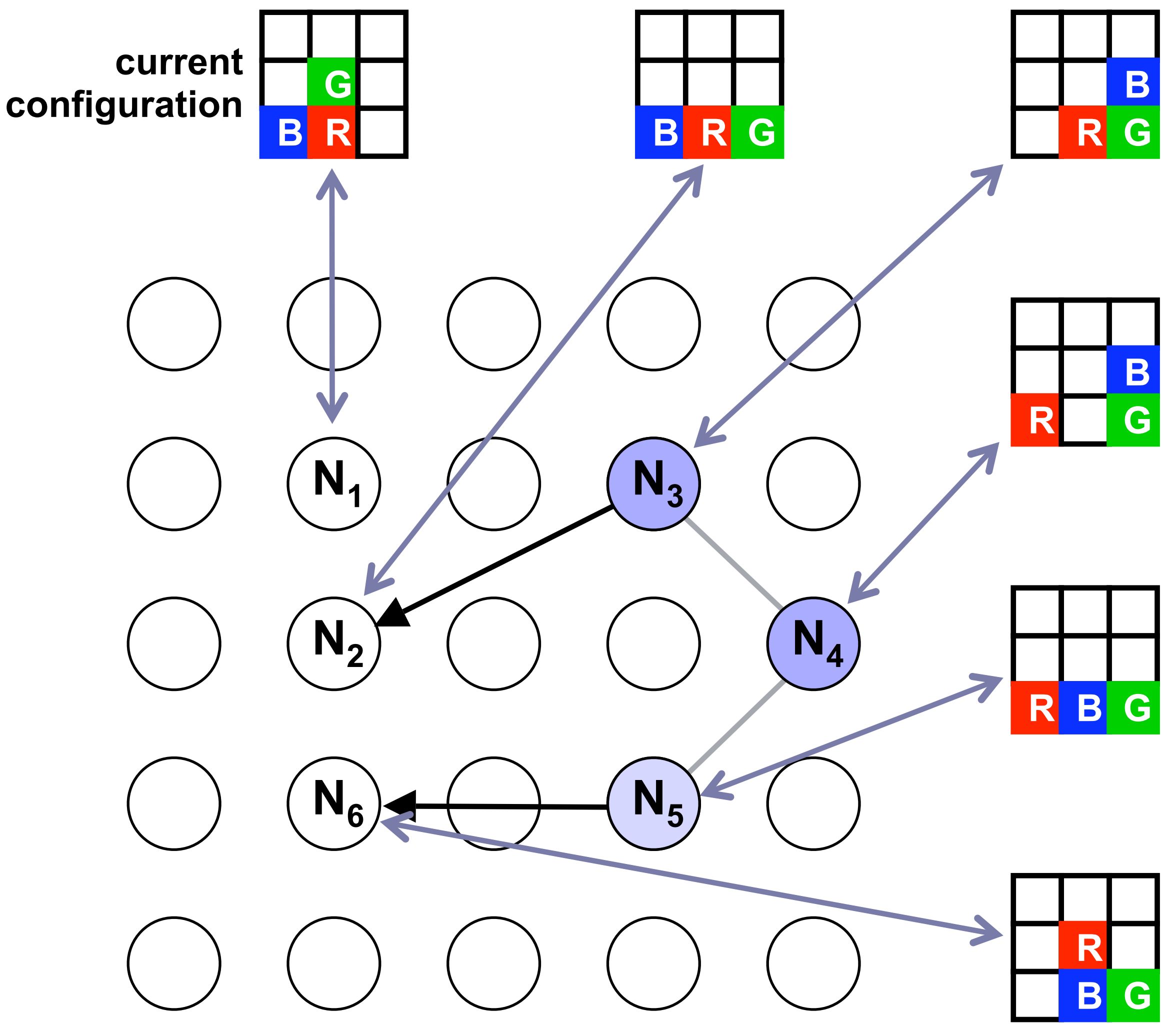

N_1 – current state

N_4 – goal state

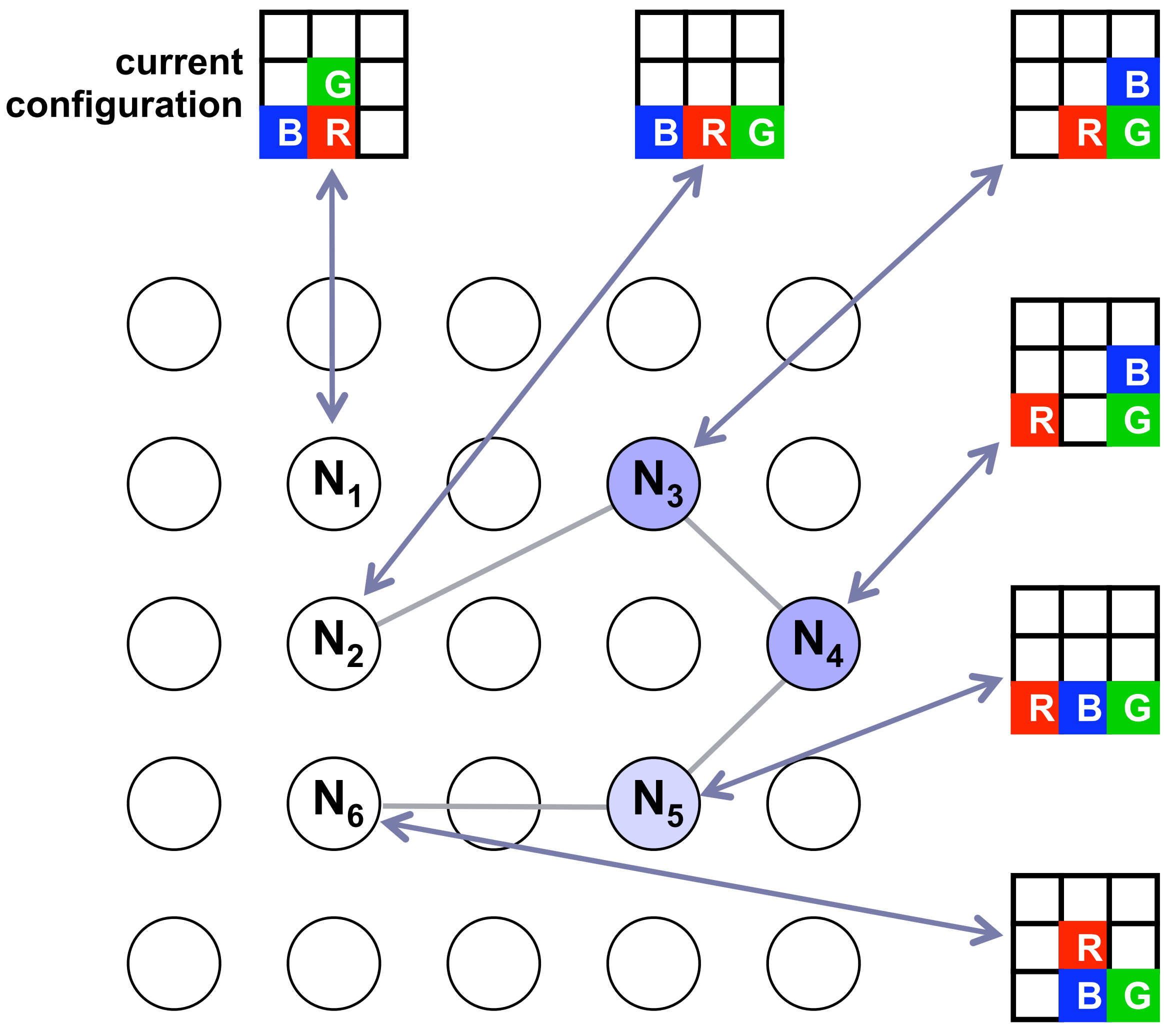

N_1 – current state

N_4 – goal state

A search tree for a block rearrangement problem

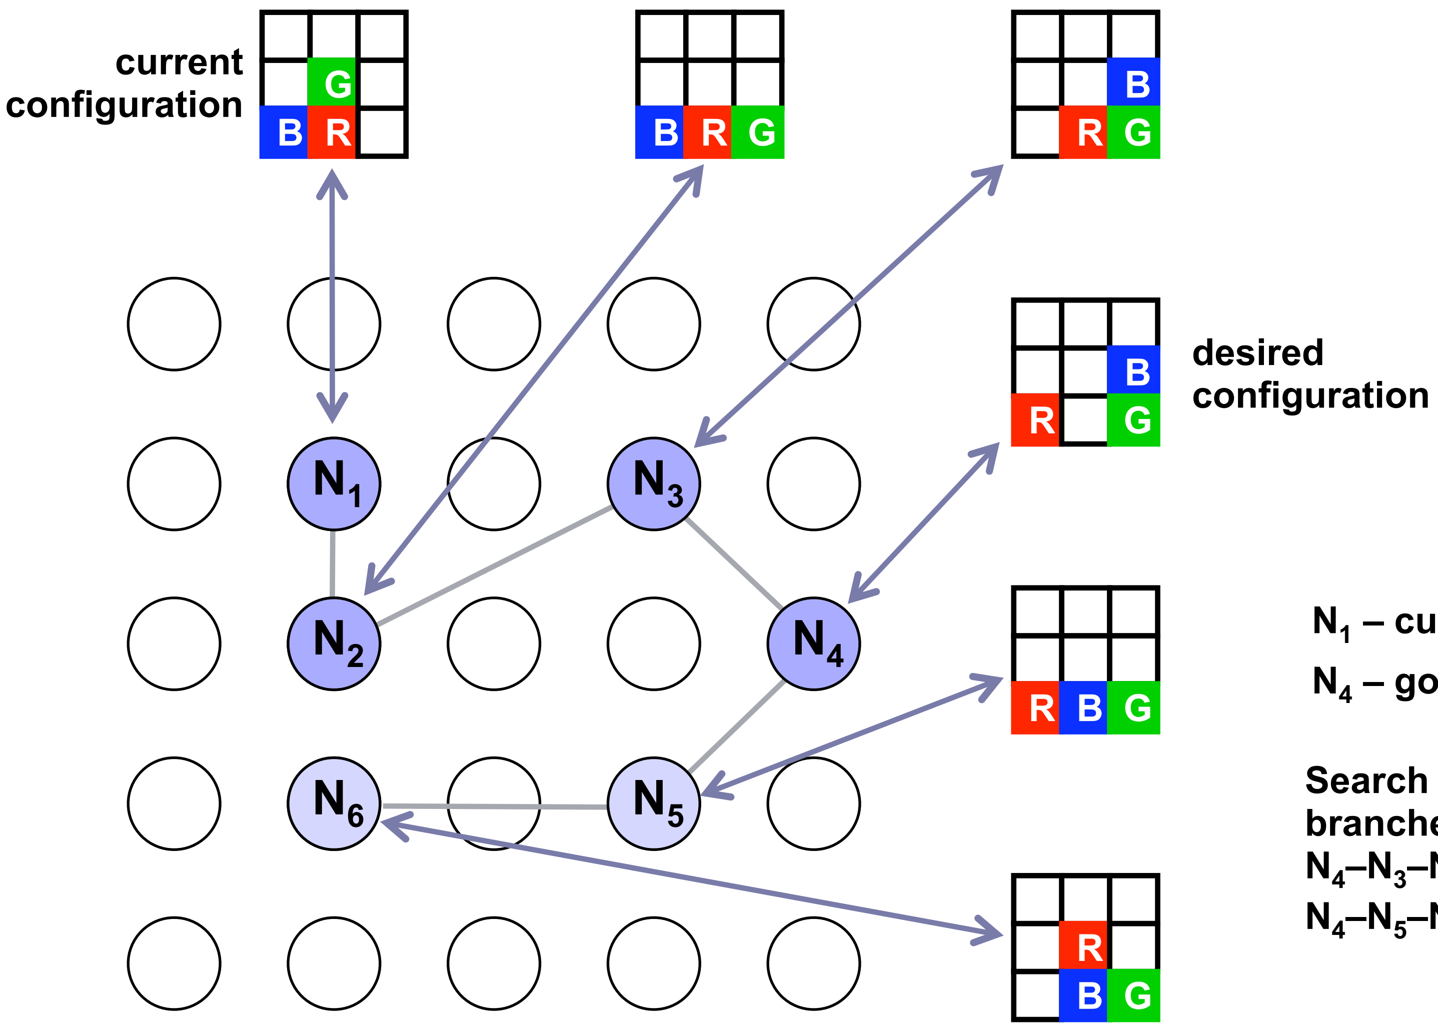

N_1 – current state

N_4 – goal state

Search Tree with two branches:
N_4–N_3–N_2–N_1 – branch
N_4–N_5–N_6 – branch

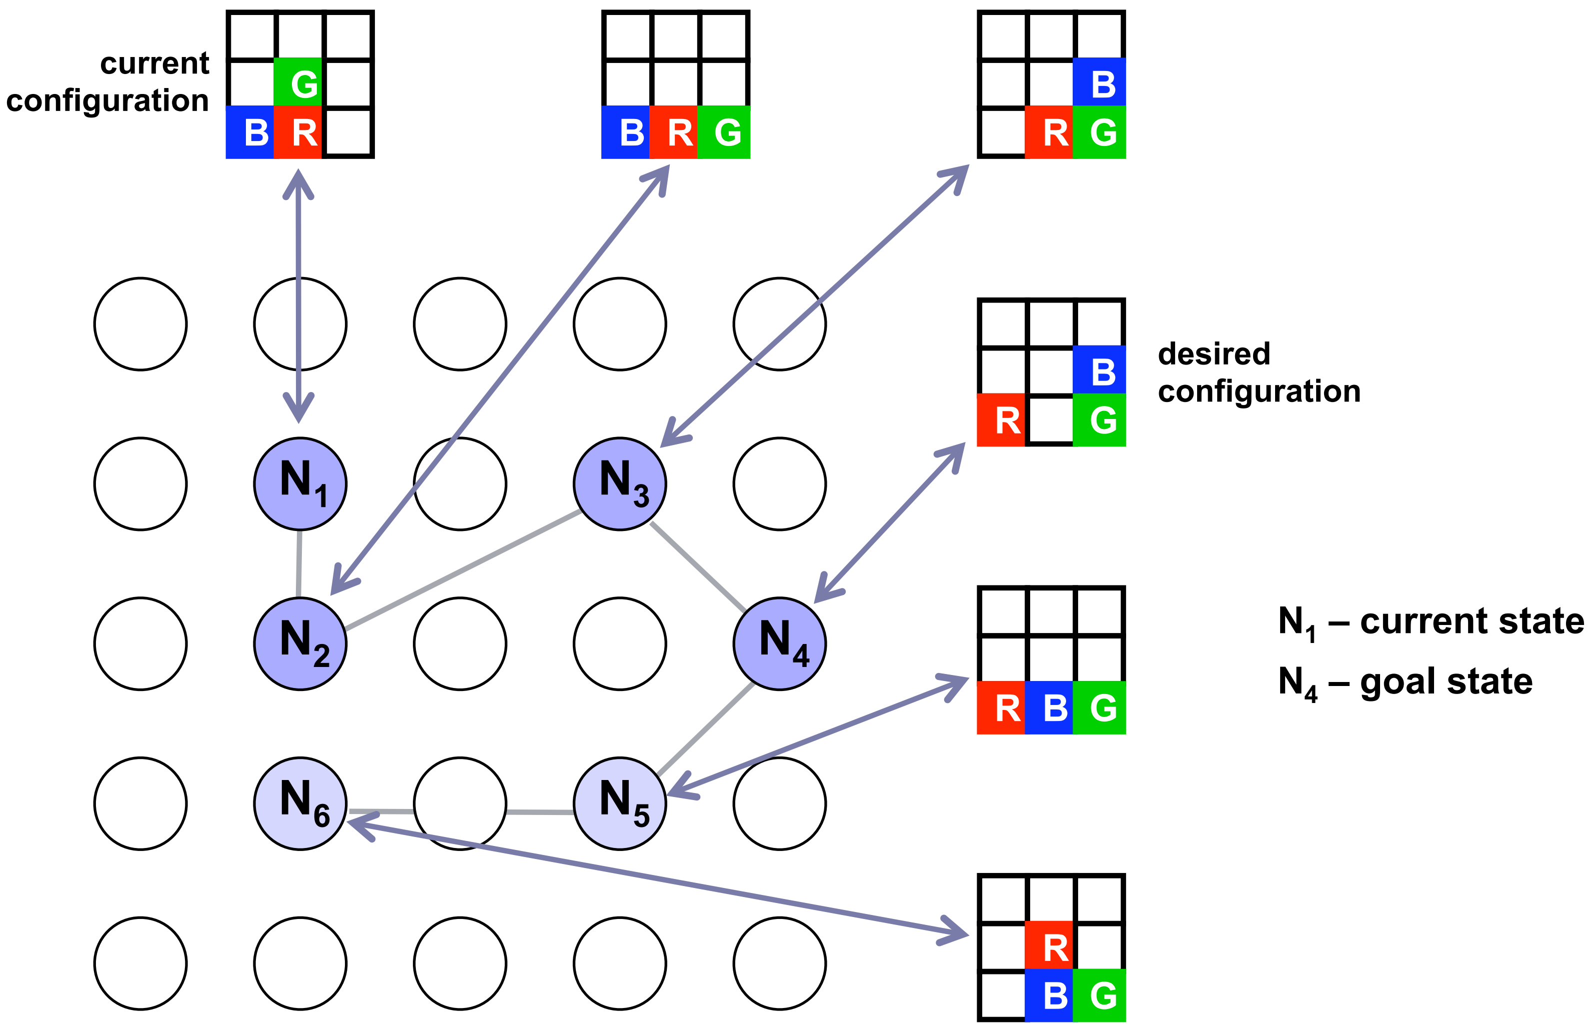
current configuration
G
B R
B R G
B
R G
desired configuration
B
R G
R B G
R
B G
N1
N3
N2
N4
N6
N5
N1 – current state
N4 – goal state

A solution path to a problem of rearranging blocks

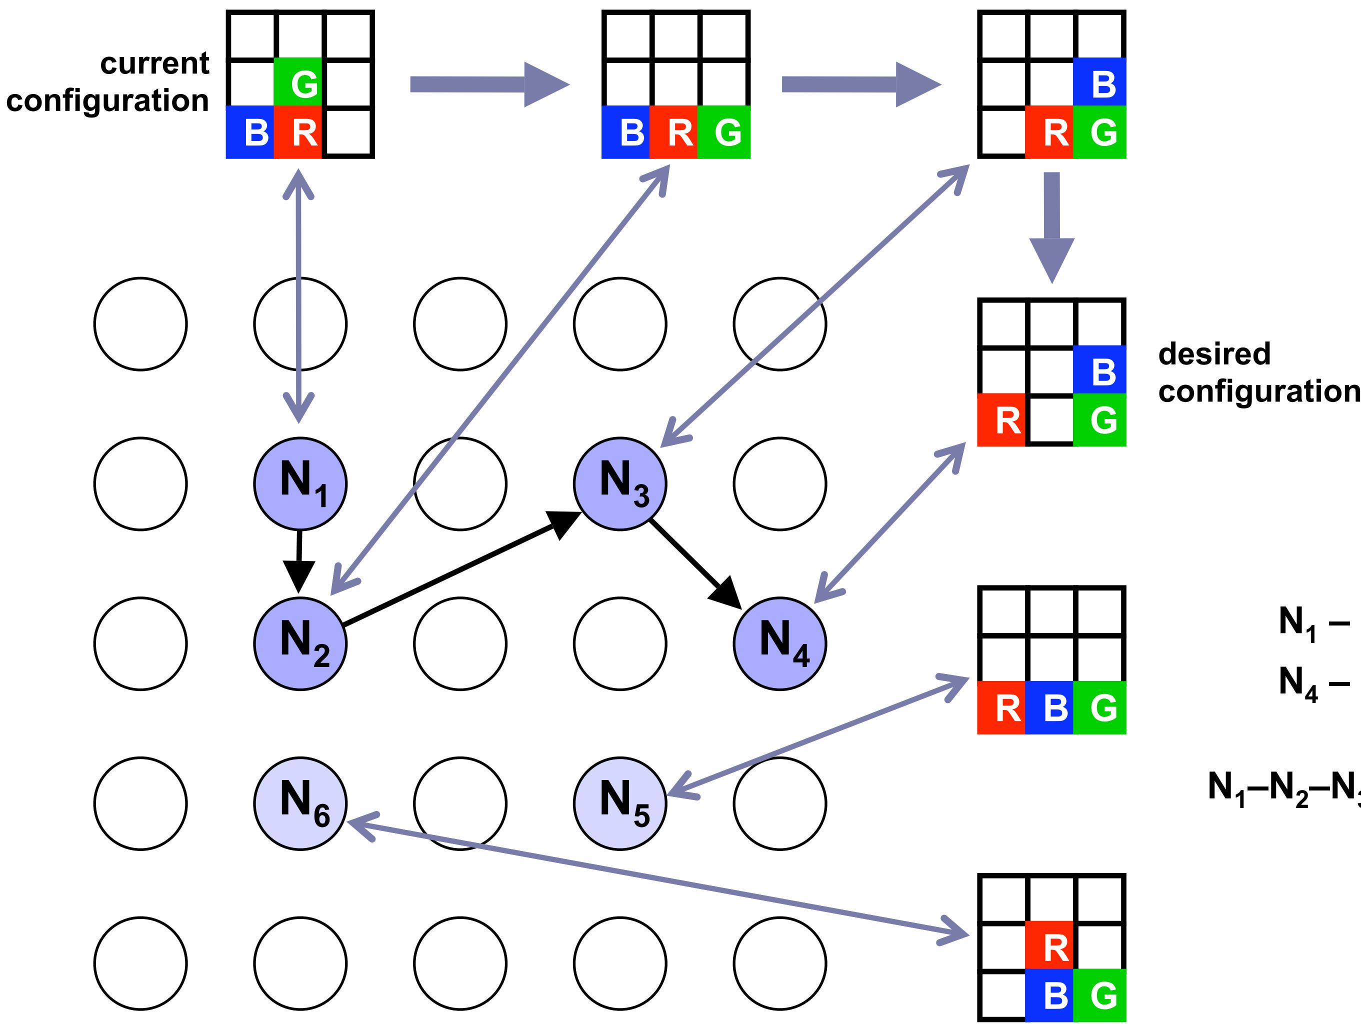

N_1 – current state

N_4 – goal state

N_1–N_2–N_3–N_4 – solution path

How to construct the map?

- In the human brain

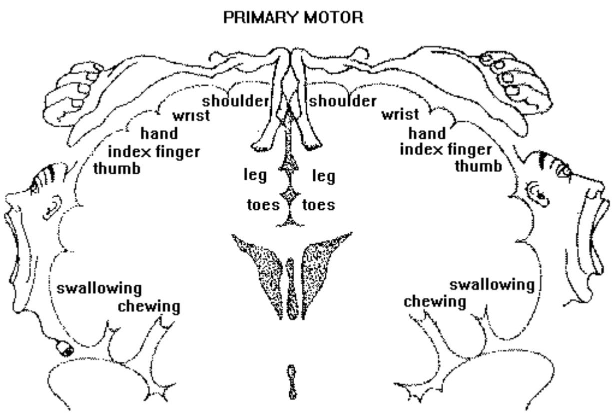

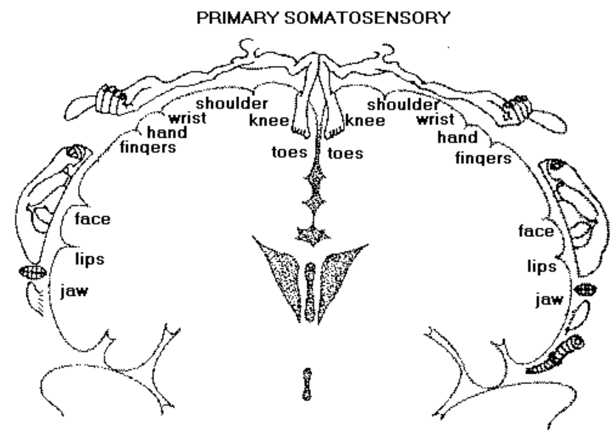

How to construct the map?

- In the human brain

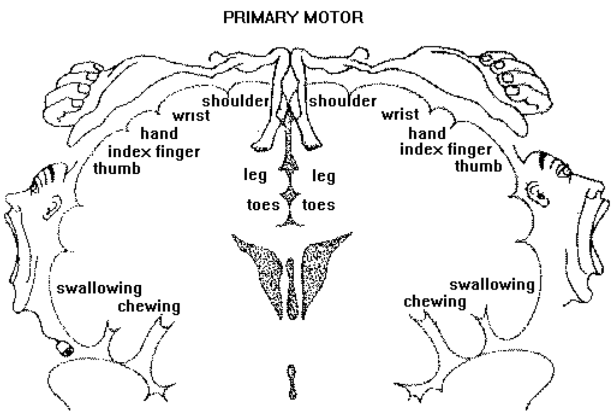

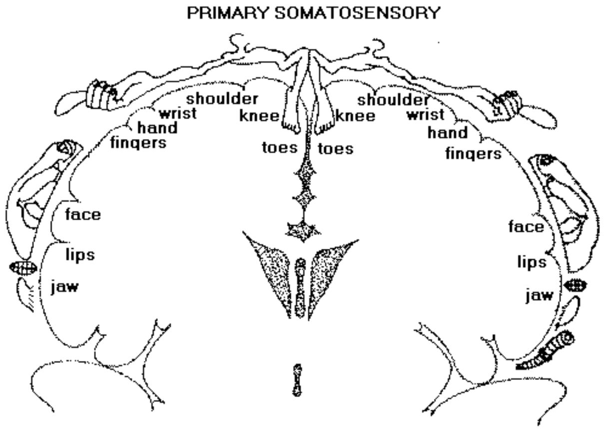

- In an artificial network
  - backpropagation
  - Kohonen
  - Hopfield
  - other classification techniques

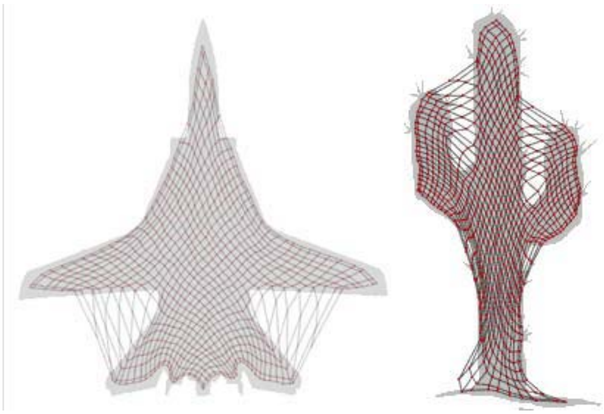

What is a node?

- Simple units used in most neural networks are not adequate
  - need extended functionality

What is a node?

- Simple units used in most neural networks are not adequate
  - need extended functionality
- Hyper-column suggested to posses needed features
  - Y Cajal (Noble Price, 1906), Szentágothai, Hubel/Wiesel (Noble Price, 1981)
  - Burnod (brain modeling)
  - See BlueBrain Project:
    - http://bluebrainproject.epfl.ch/
    - first phase finished on Nov. 26th, 2007!

Cortical hyper-column

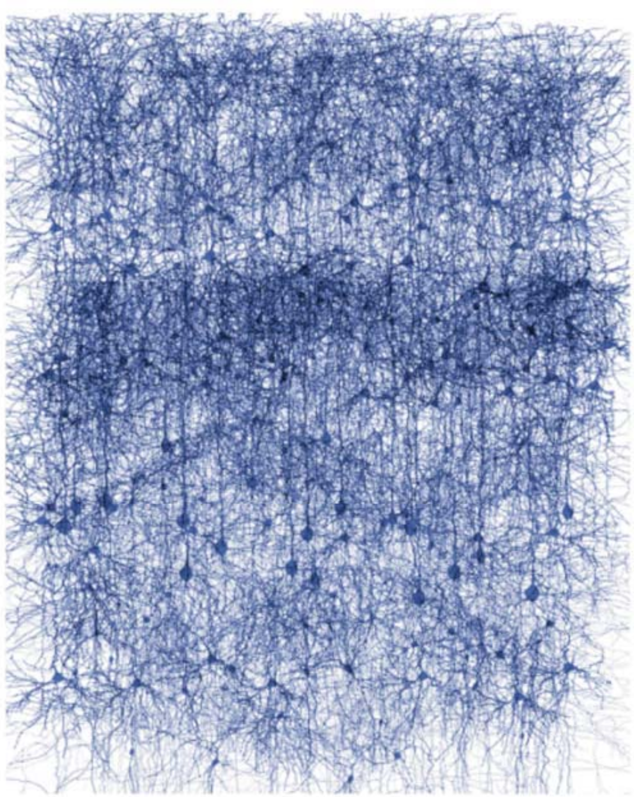

A forest of neurons. A dye is injected into each neuron and then developed inorder to reveal the morhology. This image shows a minute fraction of the cells and connections within the microcircuitry of the Neocortex

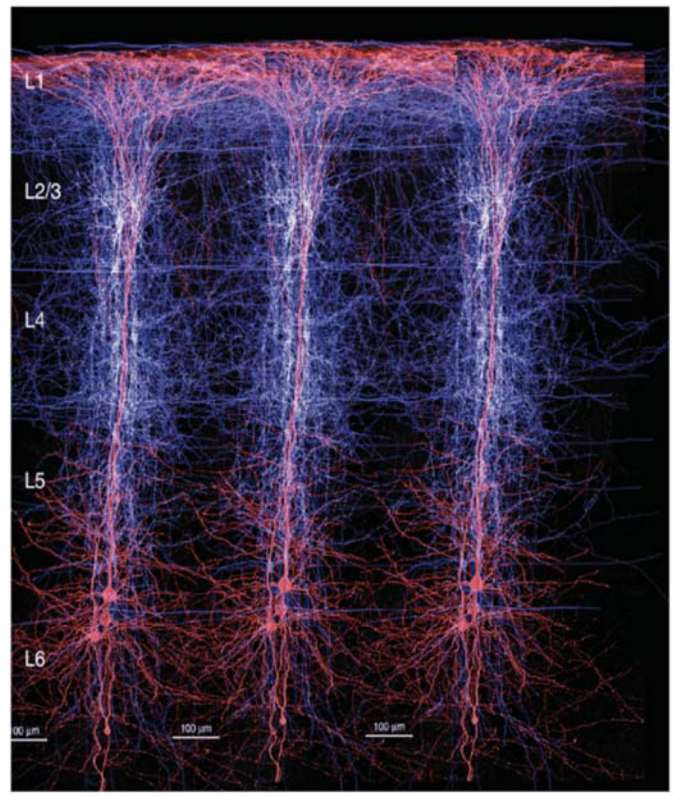

The activity in the Neocortex is tightly control by inhibitory neurons. Shown here are the inhibitory fiberse in blue that wrap around the pyramidal neurons, in red, in order to control their activity and prevent epilepsy.

Cortical Hypercolumn

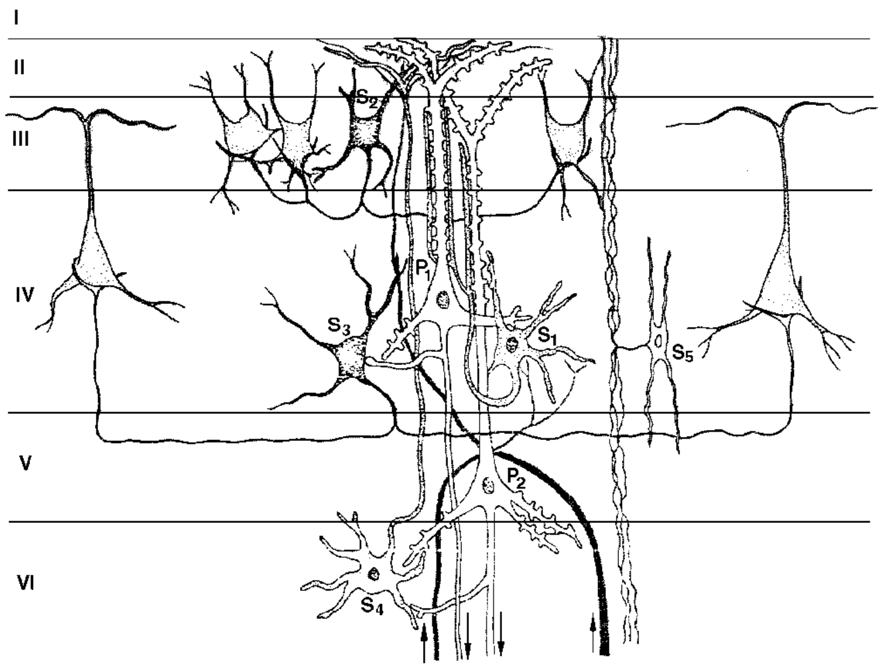

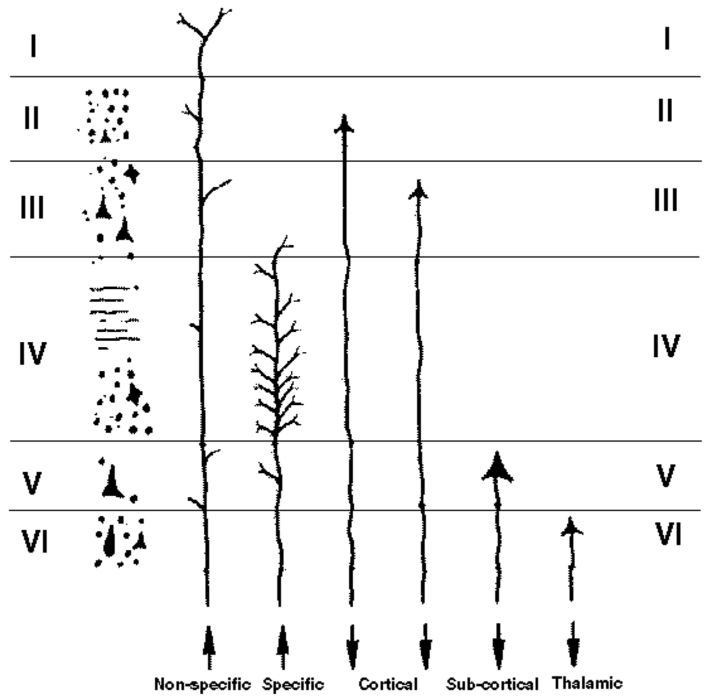

Architecture of Neurosolver

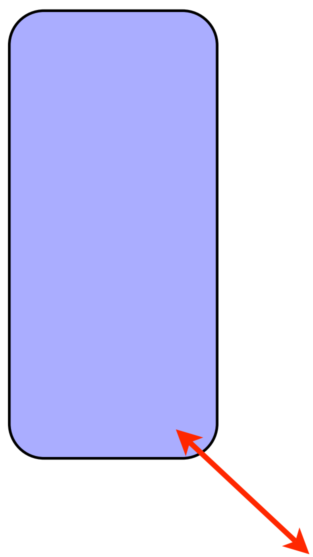

node
artificial hypercolumn

Architecture of Neurosolver

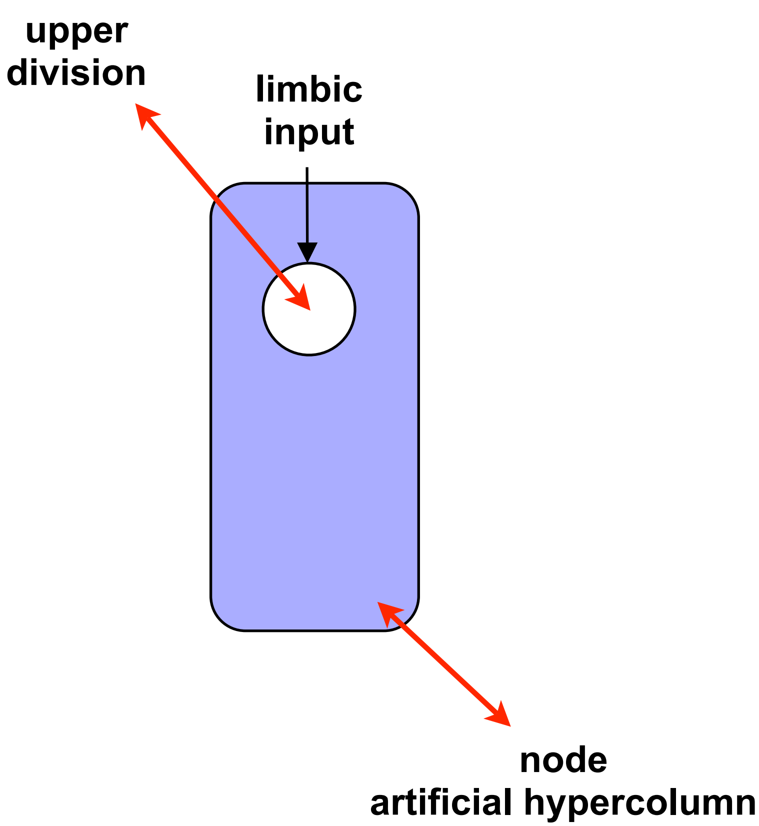

Architecture of Neurosolver

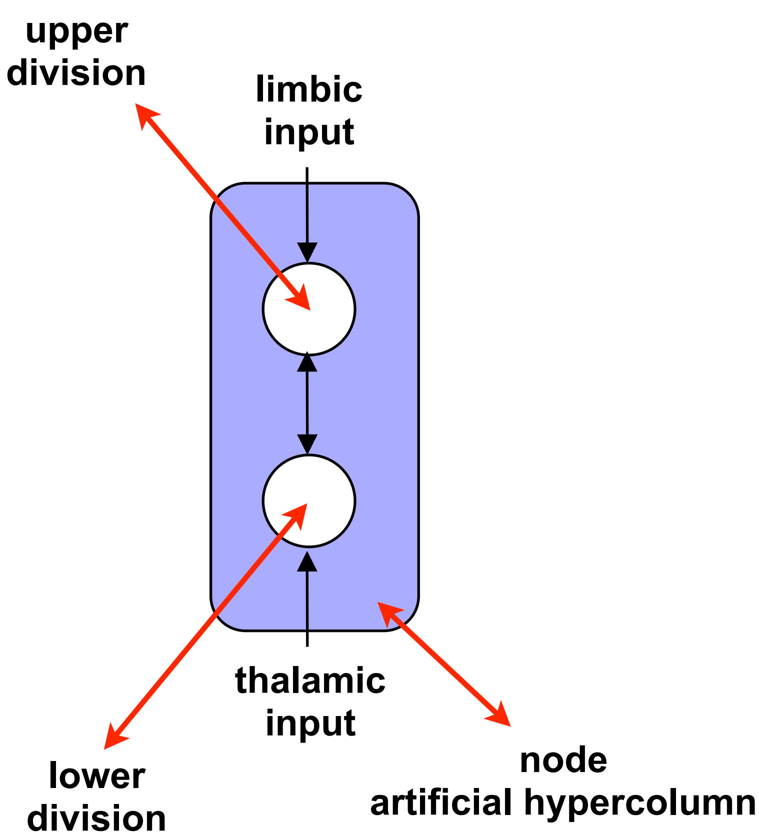

Architecture of Neurosolver

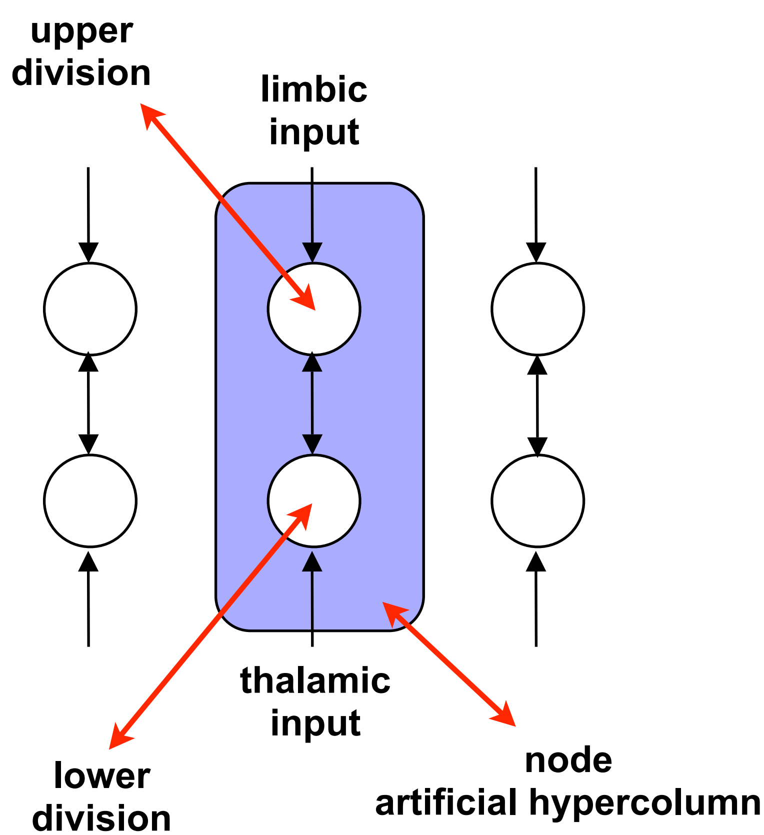

Architecture of Neurosolver

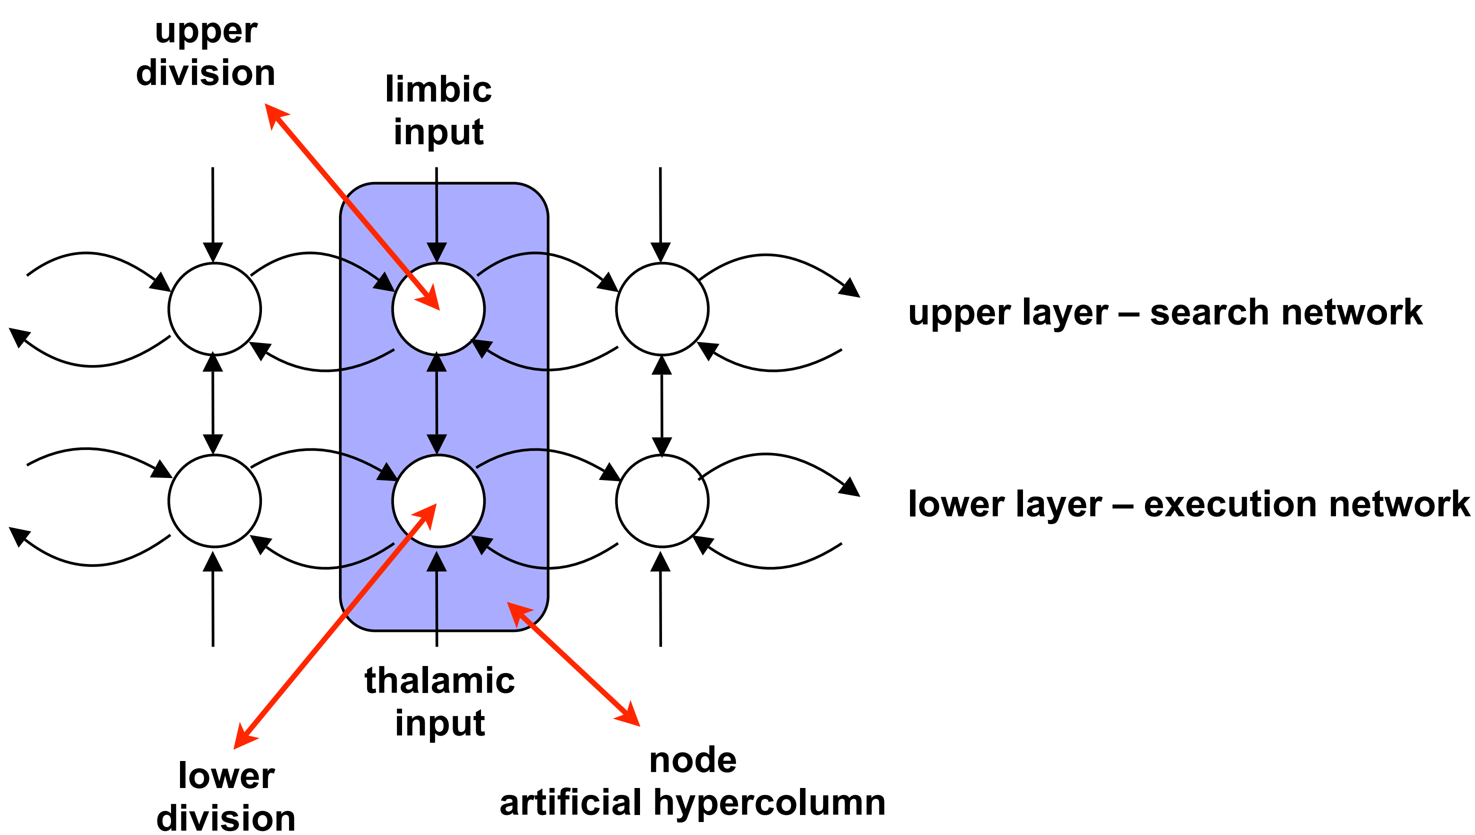

Adaptation in the Neurosolver

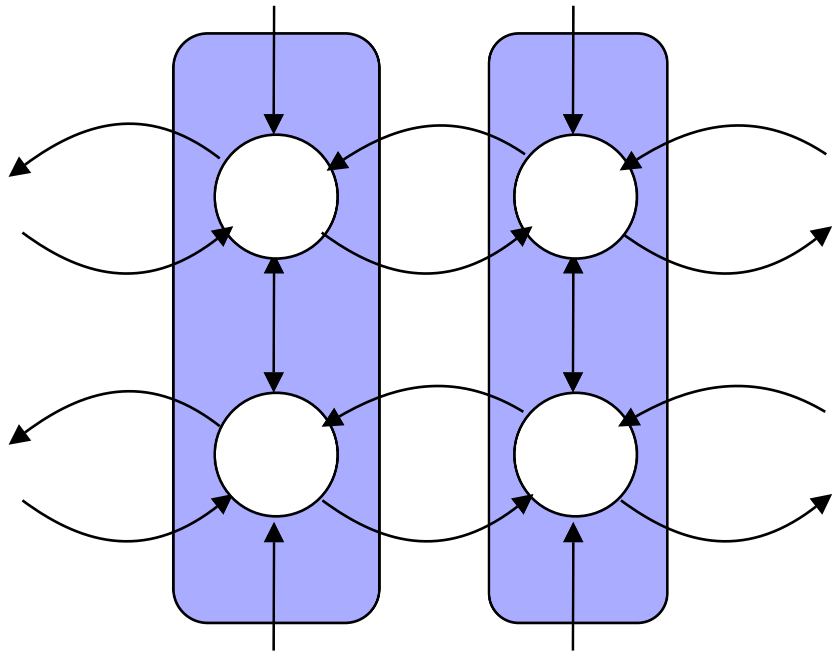

If N_2 fires after N_1 fired, then the connection from the upper division of N_2 and to the upper division of N_1 is strengthen. At the same time, the connection from the lower division of N_1 to the lower division of N_2 is strengthen as well.

Adaptation in the Neurosolver

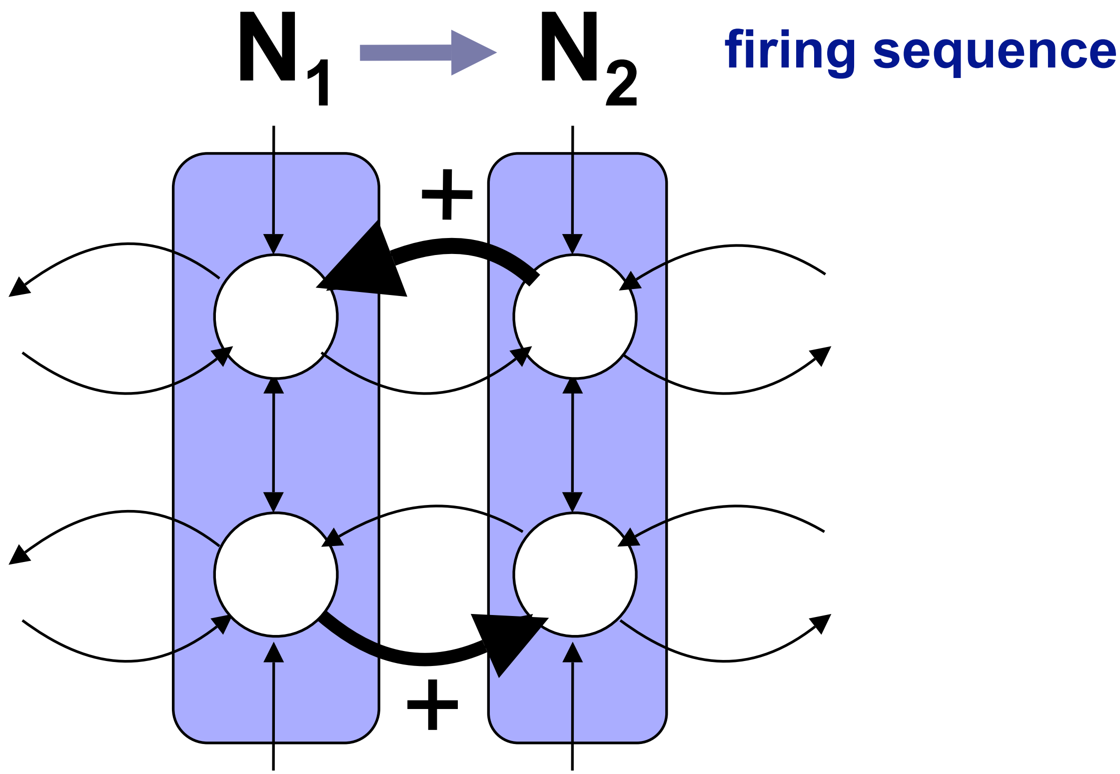

If N_2 fires after N_1 fired, then the connection from the upper division of N_2 and to the upper division of N_1 is strengthen. At the same time, the connection from the lower division of N_1 to the lower division of N_2 is strengthen as well.

Two learned chains: backward search and

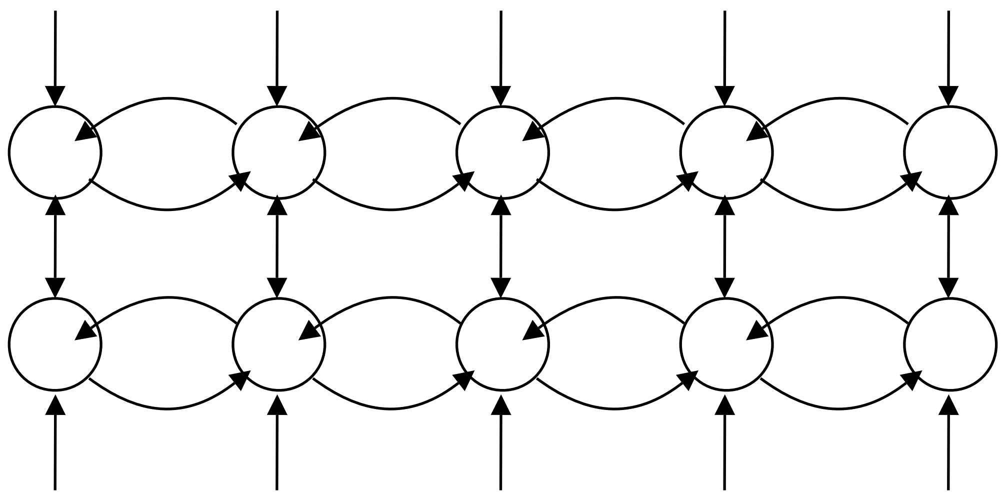

Two learned chains: backward search and

direction of learning sequence

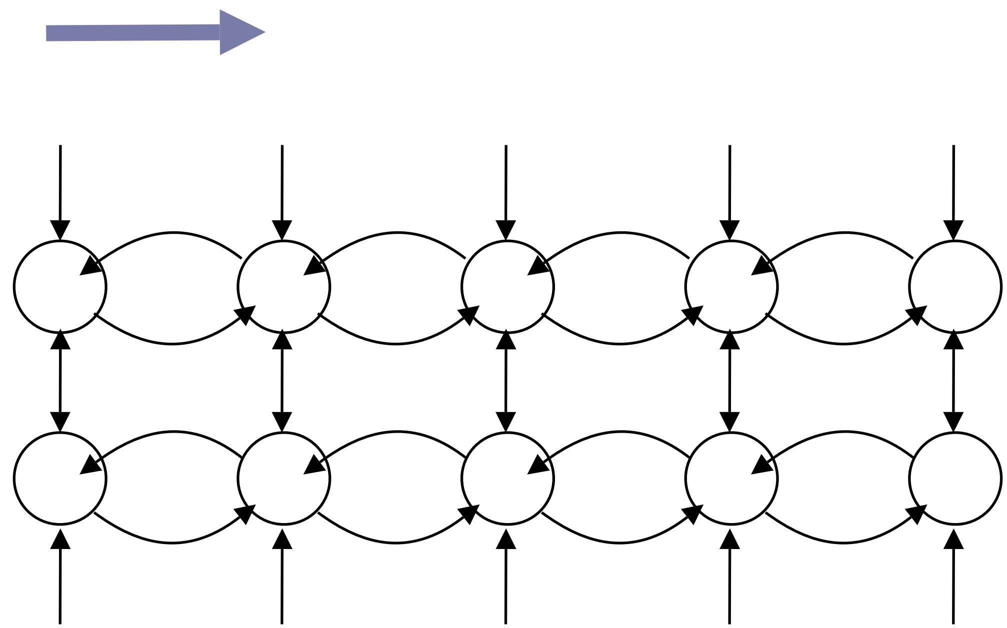

Two learned chains: backward search and

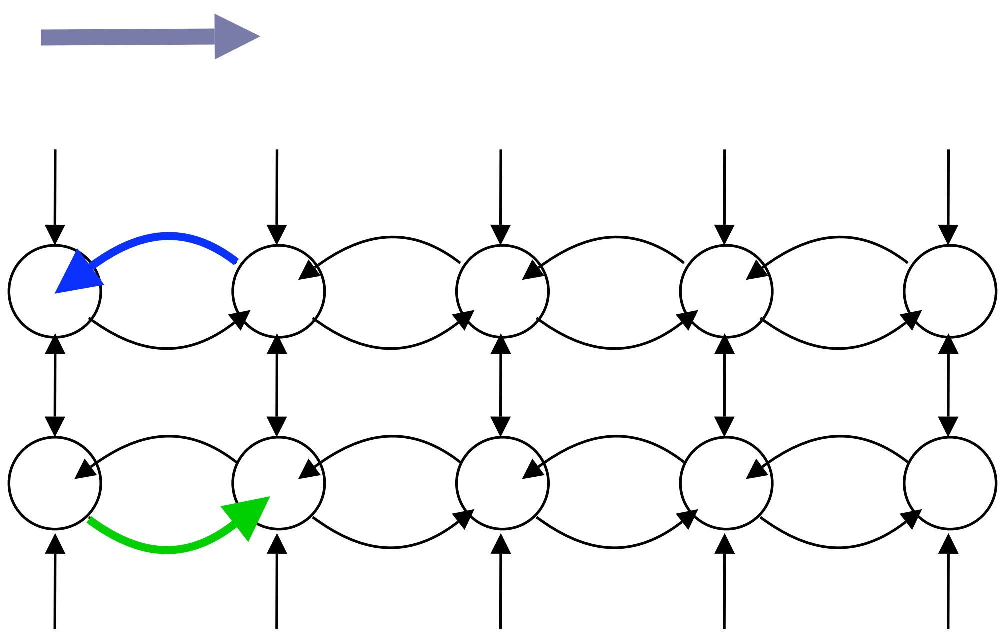

direction of learning sequence

backward search chain

forward execution chain

Two learned chains: backward search and

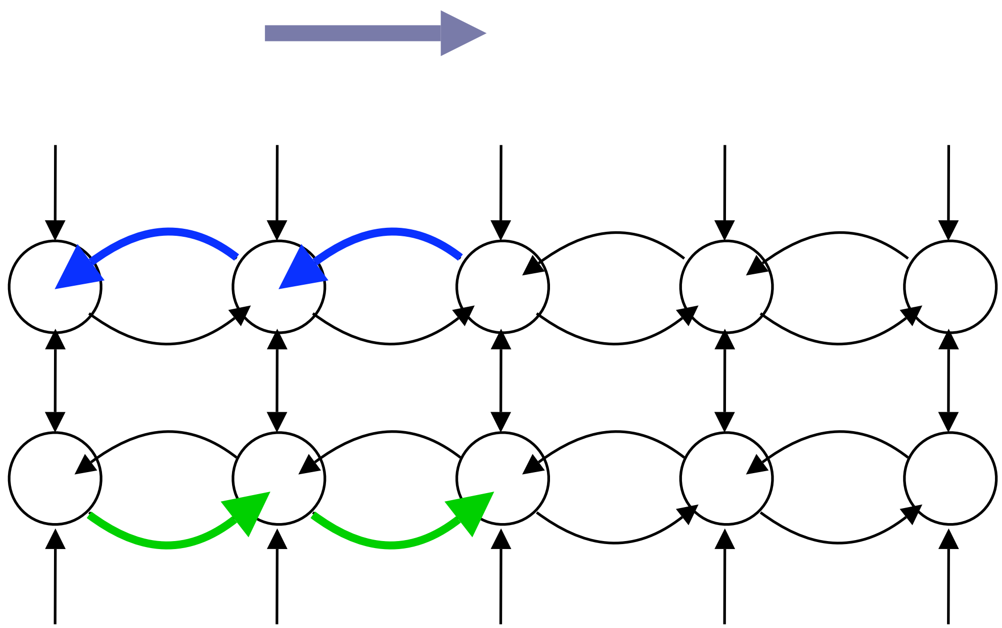

direction of learning sequence

backward search chain

forward execution chain

Two learned chains: backward search and

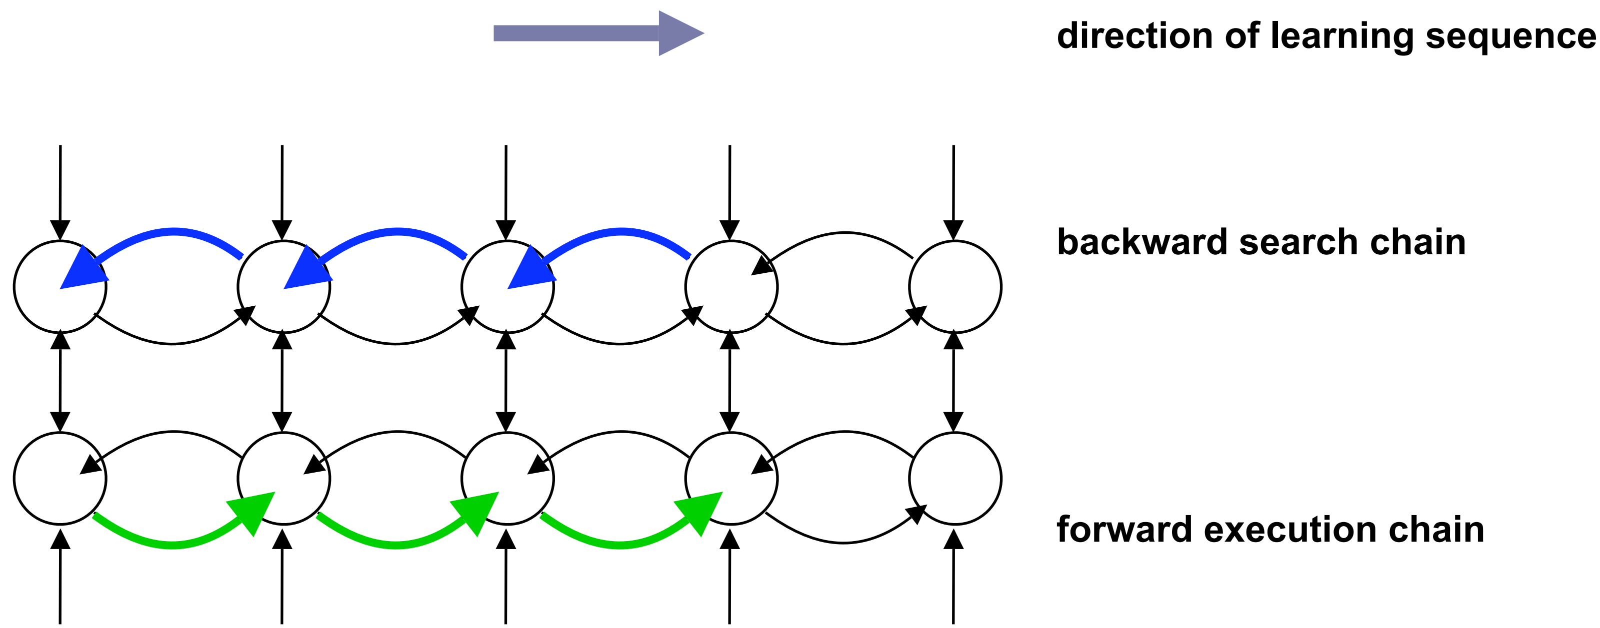

Two learned chains: backward search and

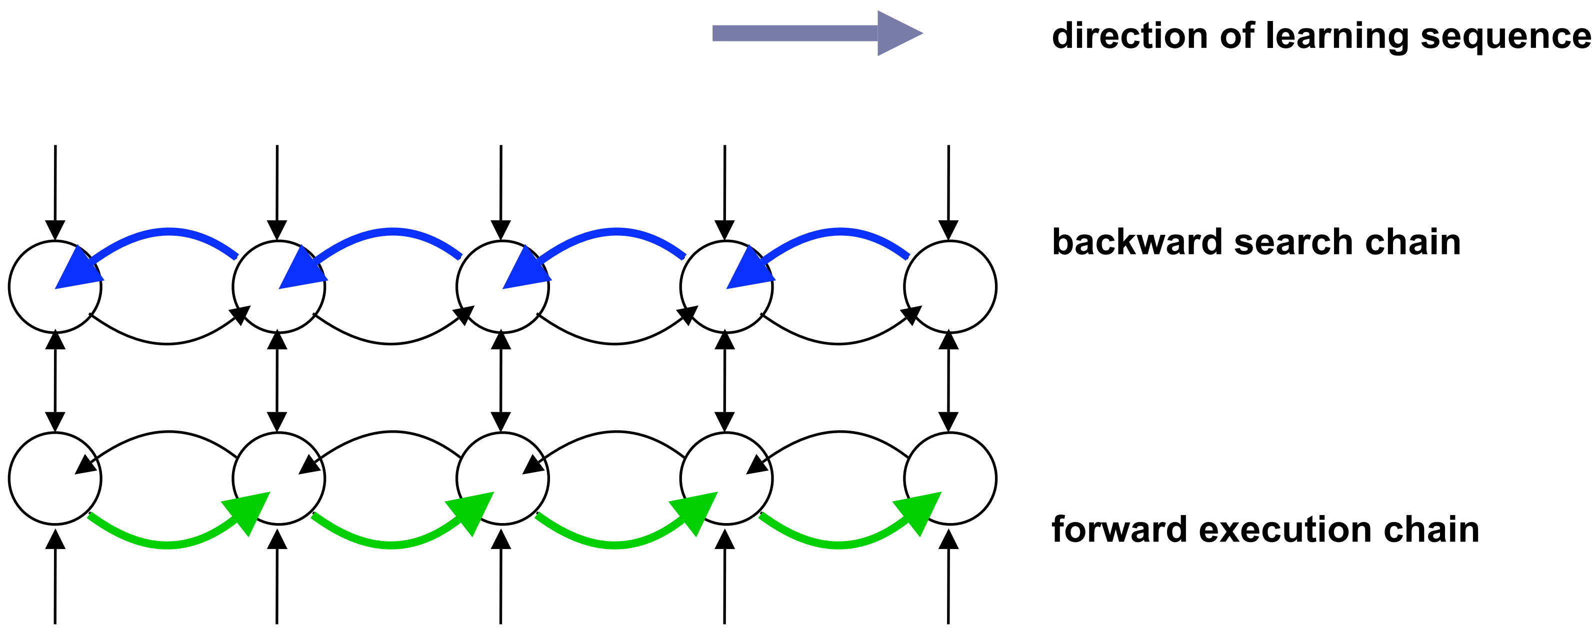

Two learned chains: backward search and

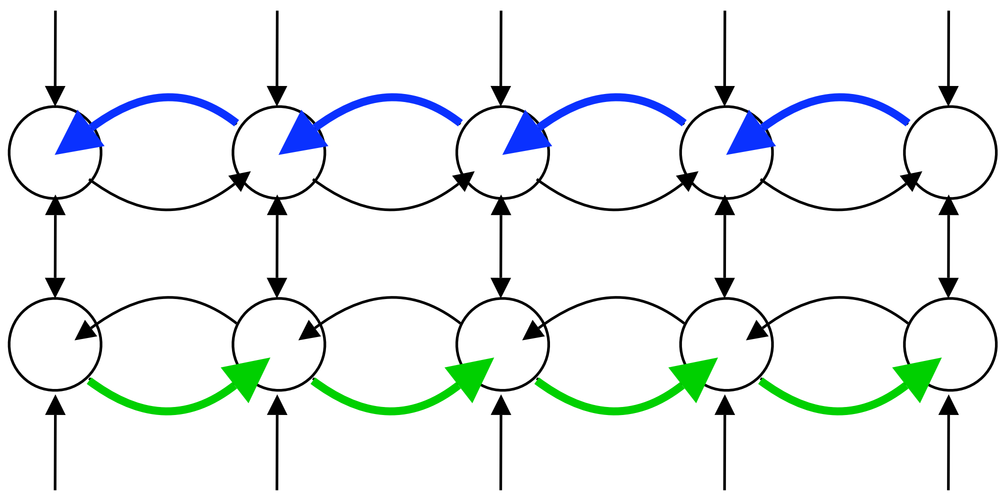

direction of learning sequence

backward search chain

forward execution chain

Search along the backward chain

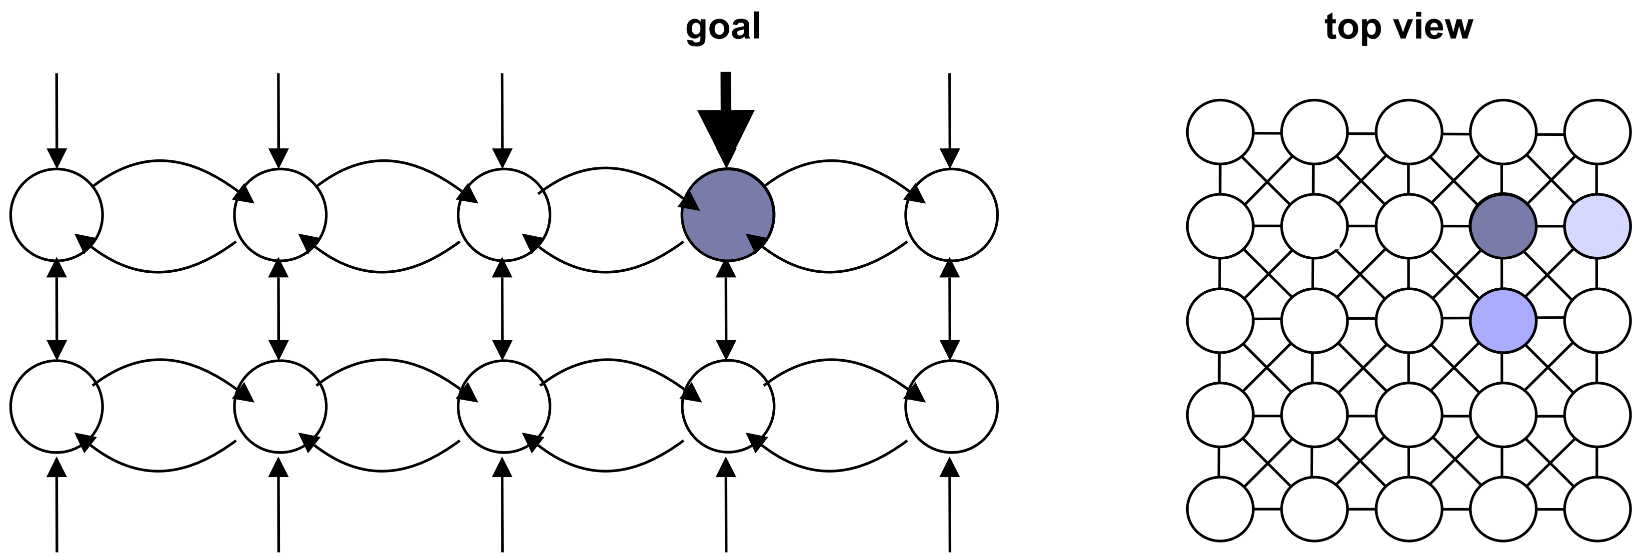

The top view shows a search tree with many branches; the cross-section is just for one of the branches (color intensity indicates search level of activity)

Search along the backward chain

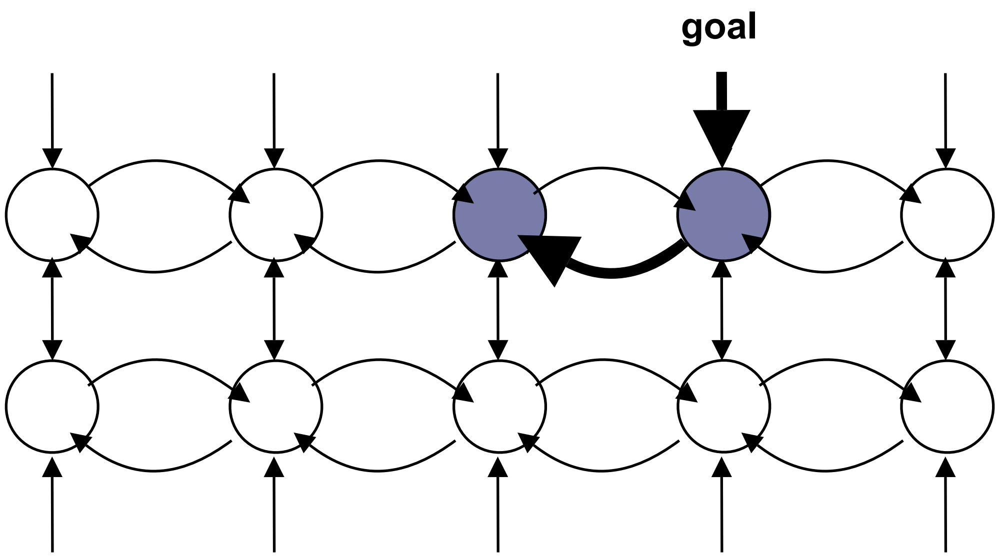

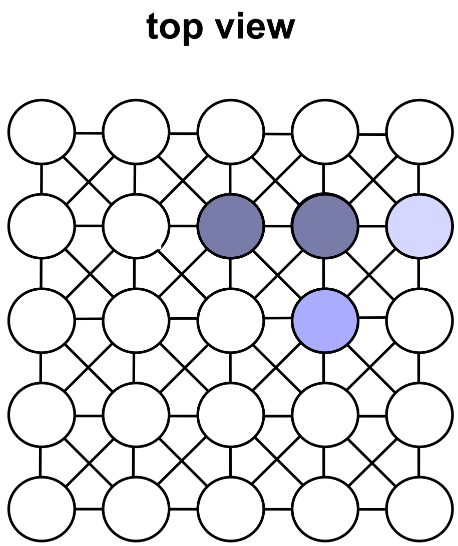

The top view shows a search tree with many branches; the cross-section is just for one of the branches (color intensity indicates search level of activity)

Search along the backward chain

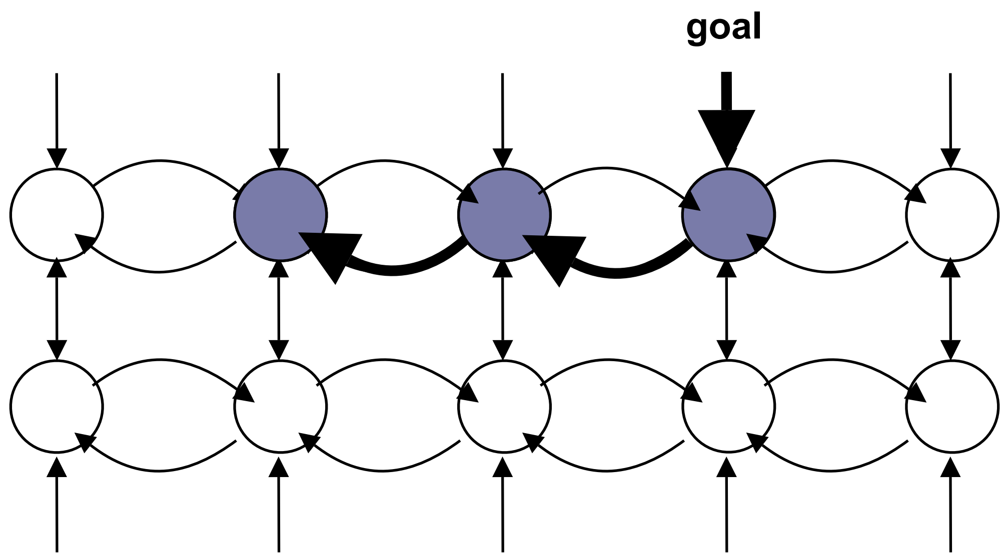

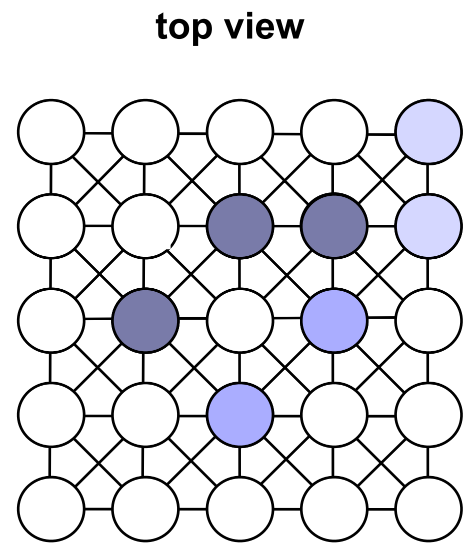

The top view shows a search tree with many branches; the cross-section is just for one of the branches (color intensity indicates search level of activity)

action potential = activity * P

goal

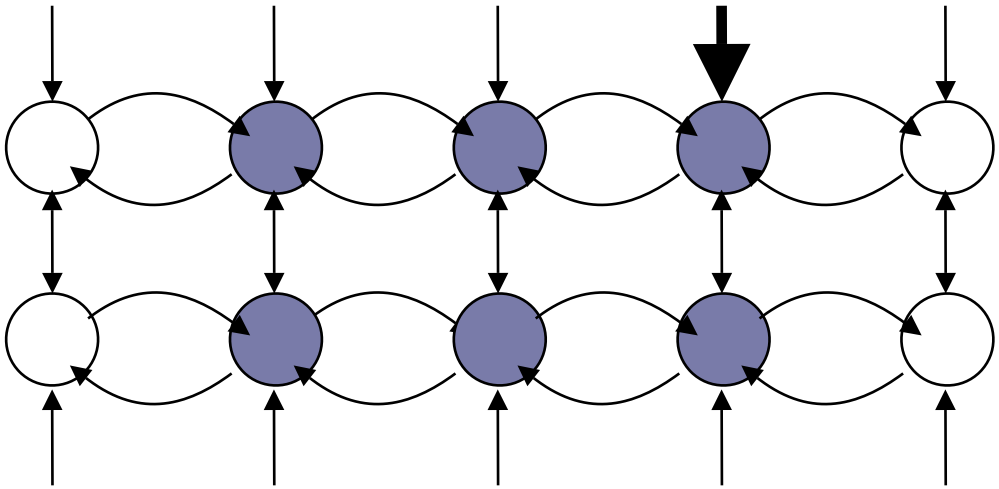

top view

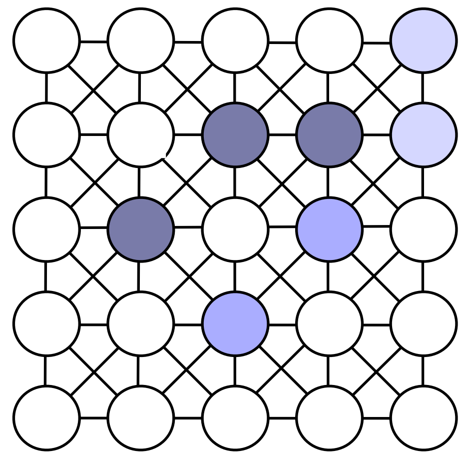

Triggering the solution

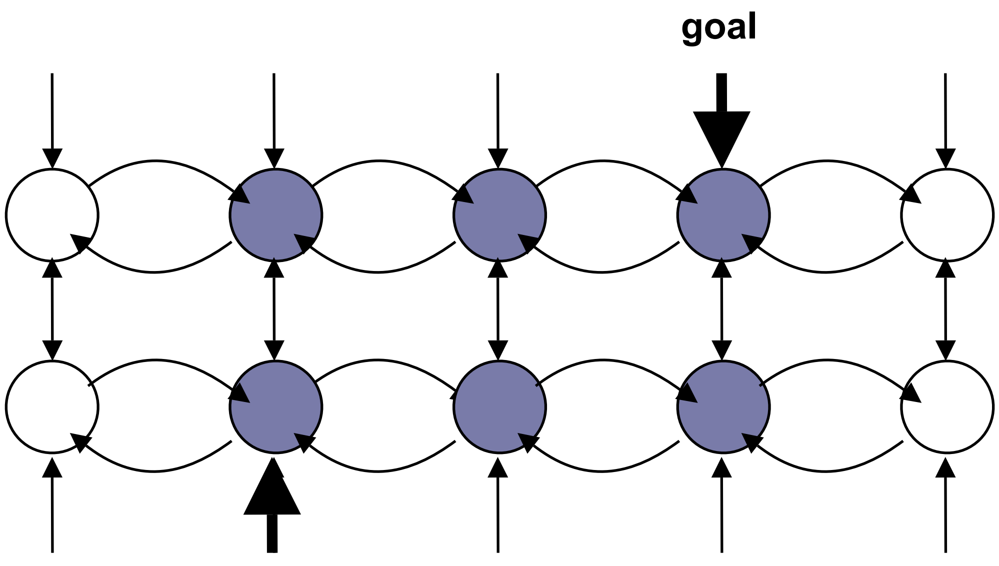

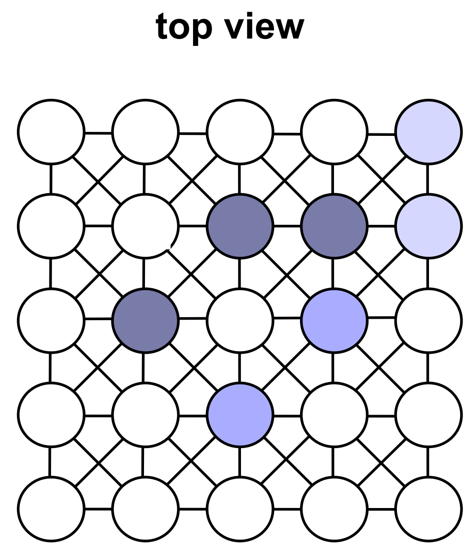

(dark navy color indicates a firing level of activity).

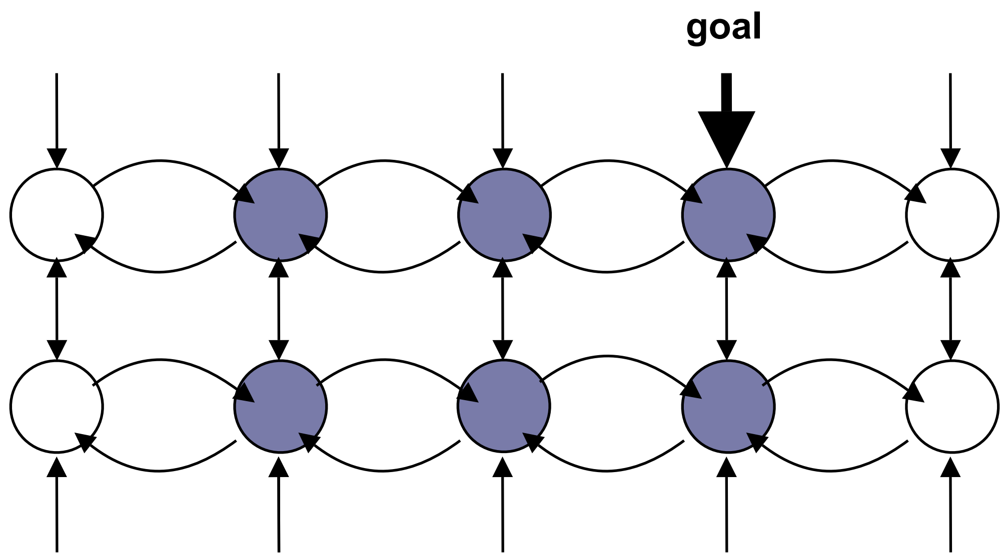

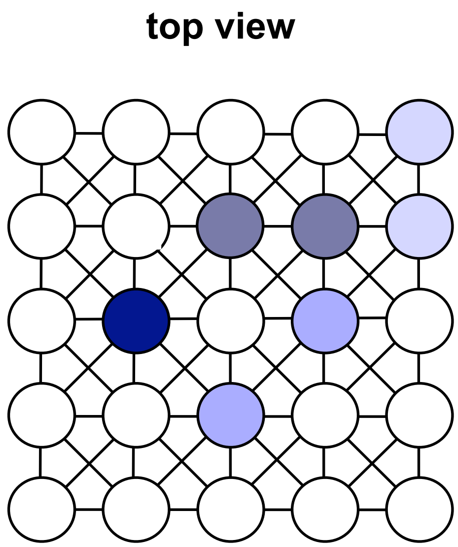

Note, that the node that has fired is shut down; i.e., it is not active anymore.

Stepping through the solution path

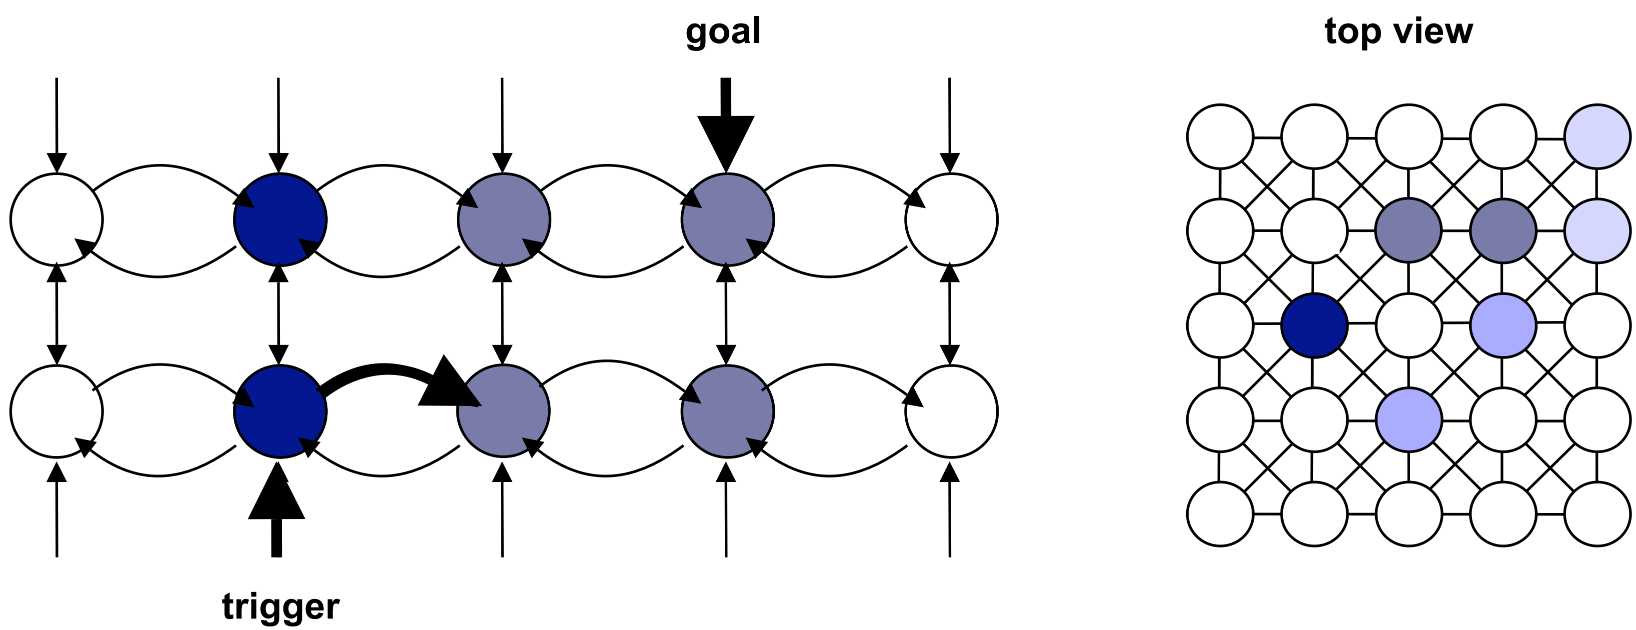

Note, that the node that has fired is shut down; i.e., it is not active anymore.

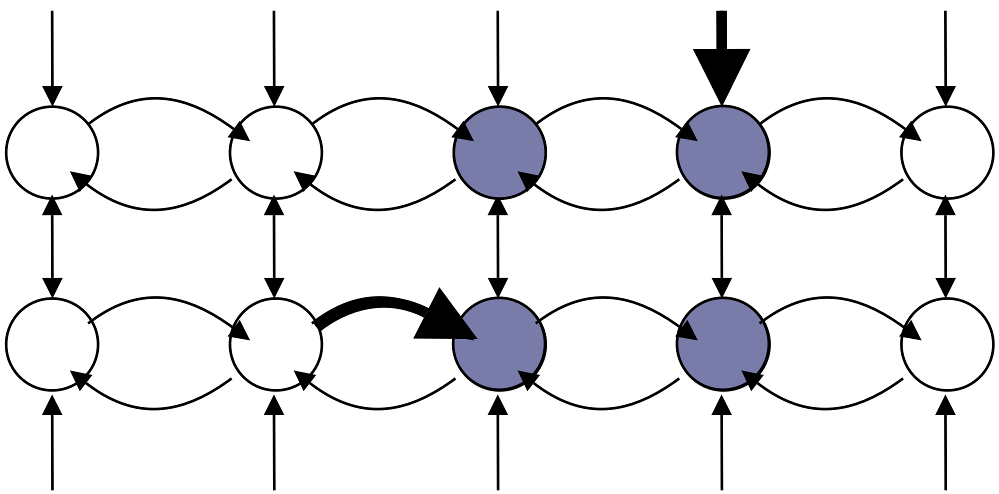
goal

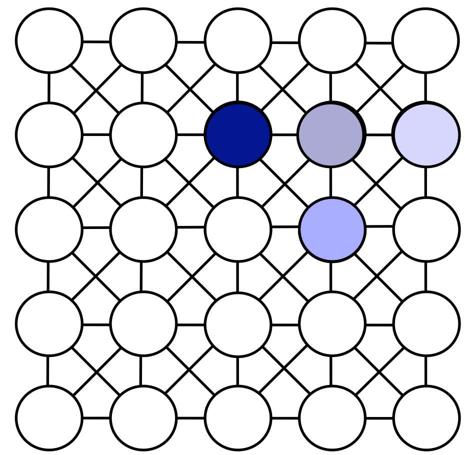
top view

Final step in the solution path

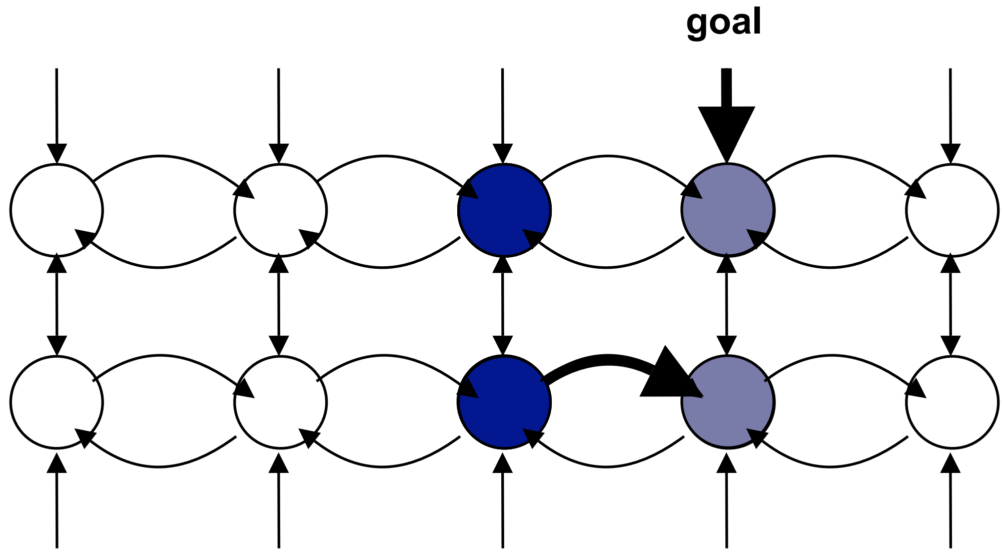

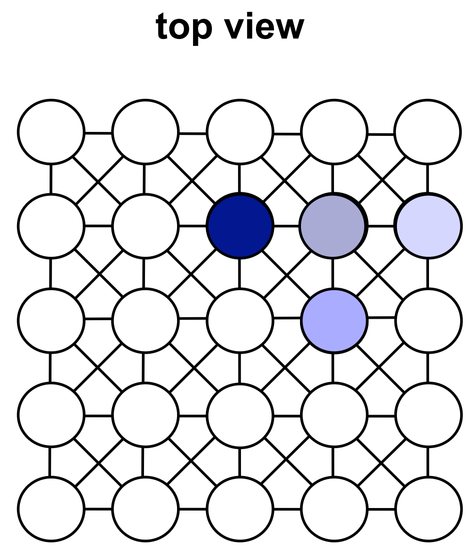

Final step in the solution path

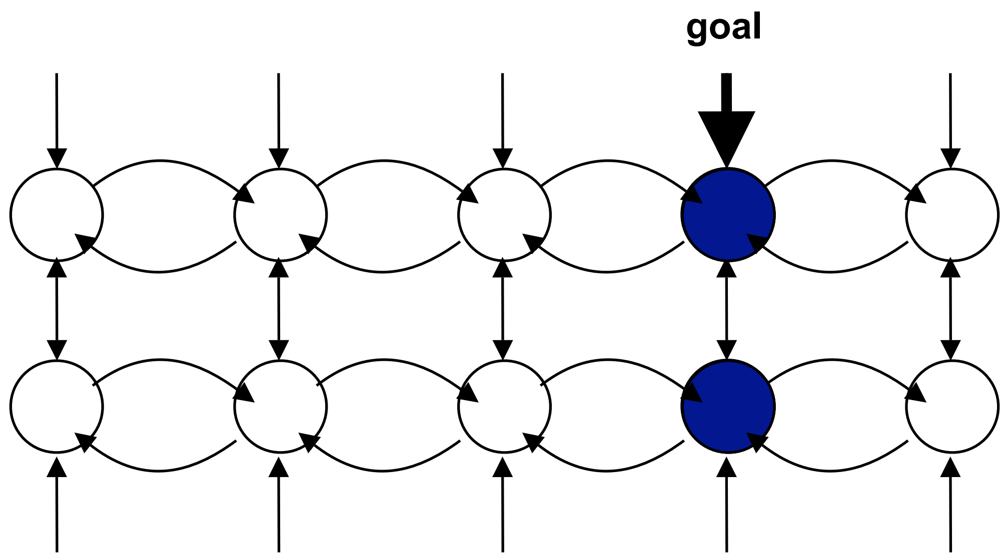

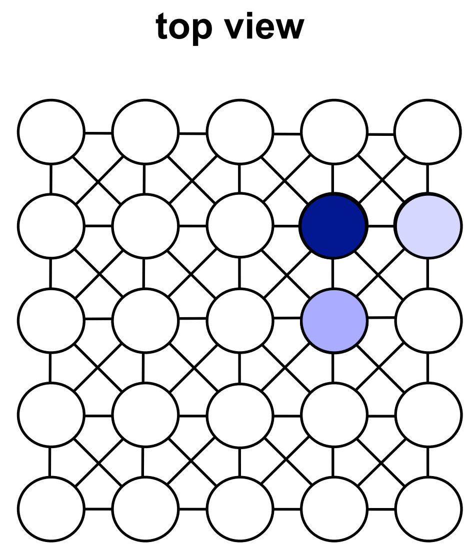

Final step in the solution path

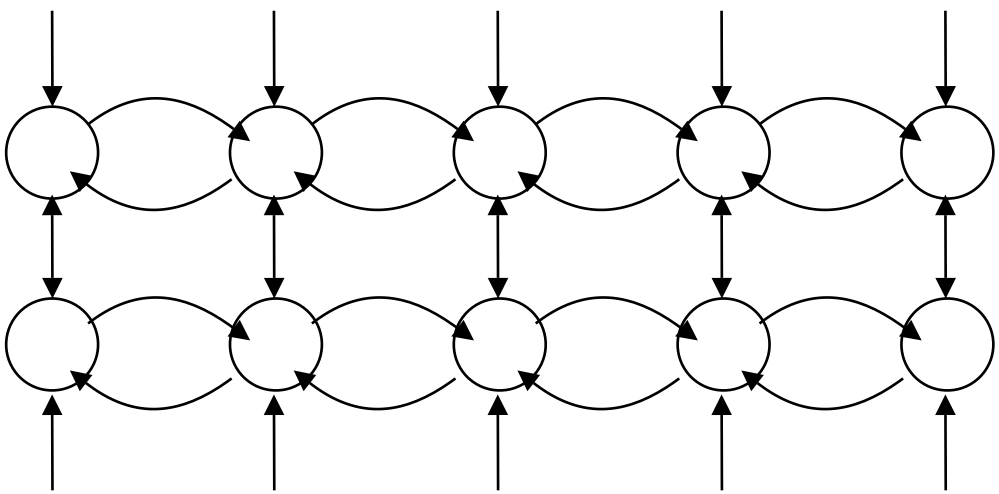

top view

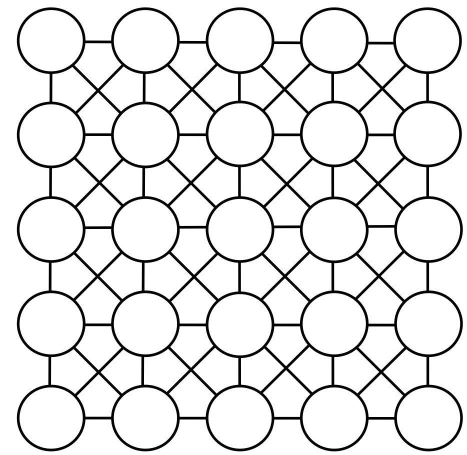

Goal attained – activity ceases

When the node corresponding to the goal fires, it is shut down as well. In that way, the source of the search activity disapears: an indication that the goal has been achieved.

Maze simulator

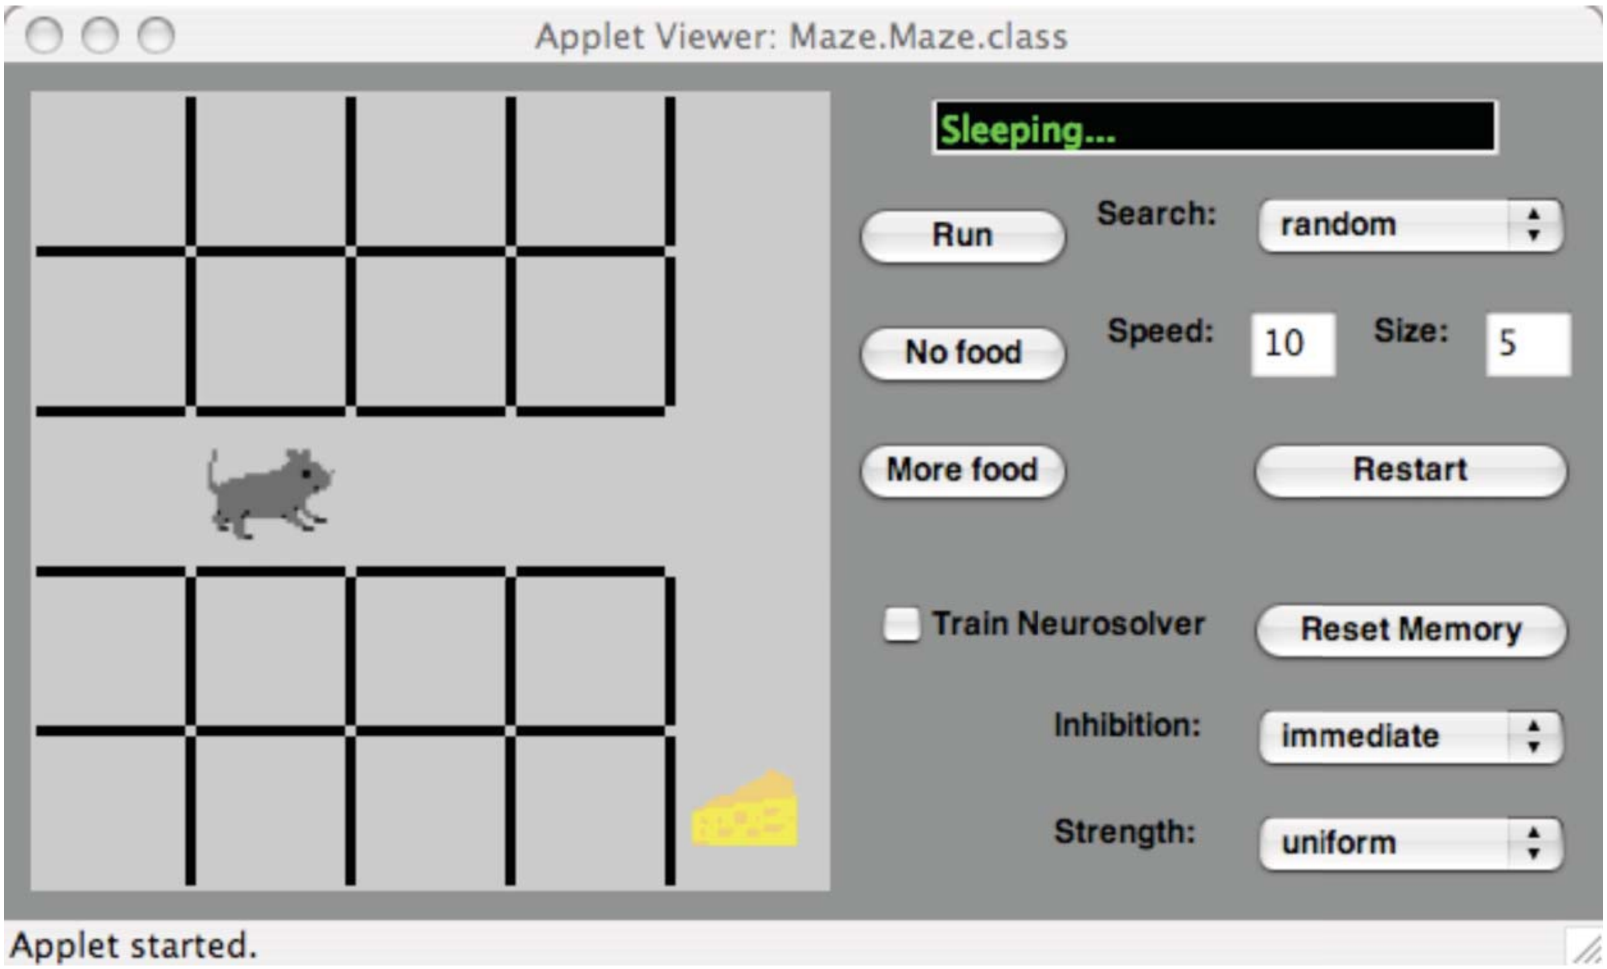

Experiments: simple T-maze

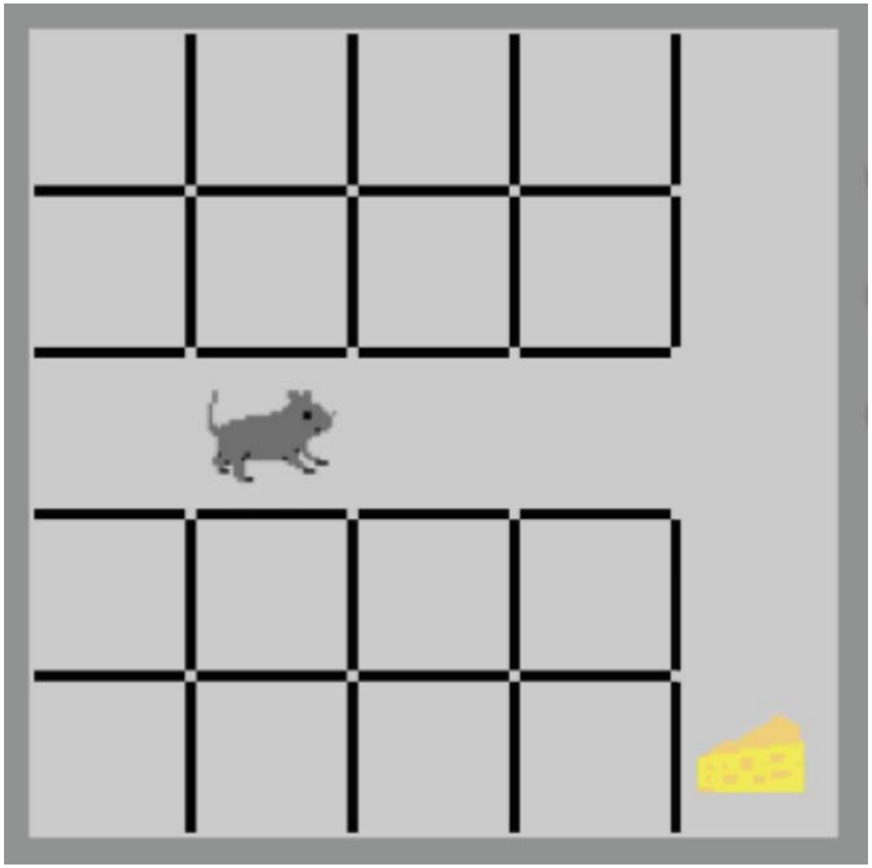

If food is placed consistently in one arm of the T, then this is the arm that will be selected by the rats in the subsequent runs. If the rat obtained food from both arms then it will choose the one that has a better trace in memory (higher probability).

Live rats may exhibit aberrant behavior under stress.

Experiments: multiple paths

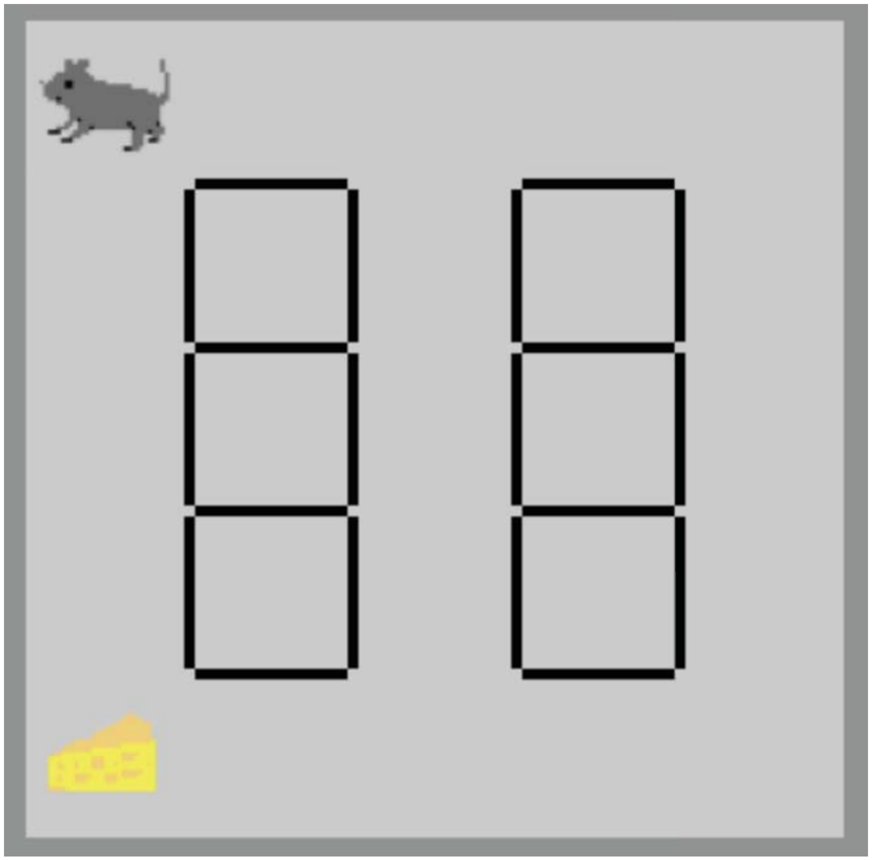

The rat selects the shortest path if the uniform learning is selected. If probabilistic learning is chosen, then the most probable path is taken: the path most often followed and rewarded in the past. This behavior comes from the fundamental characteristics of the Neurosolver. If a wall is created along the shortest path, then the rat reconsiders the plan and selects an alternate path backtracking as necessary.

Experiments: star-shaped T-maze

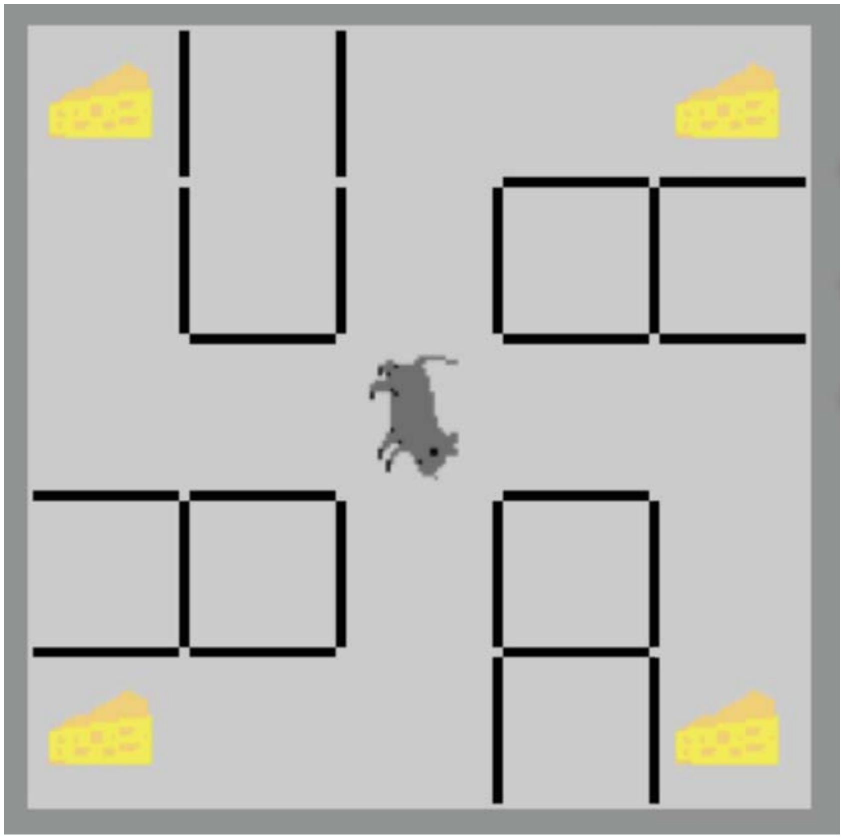

The simulated rat trying to get all food navigates to food along the minimal path.

If food is removed from certain locations, then the rat will tend to move to the branches that provided consistent food-reward.

Experiments: complex T-maze

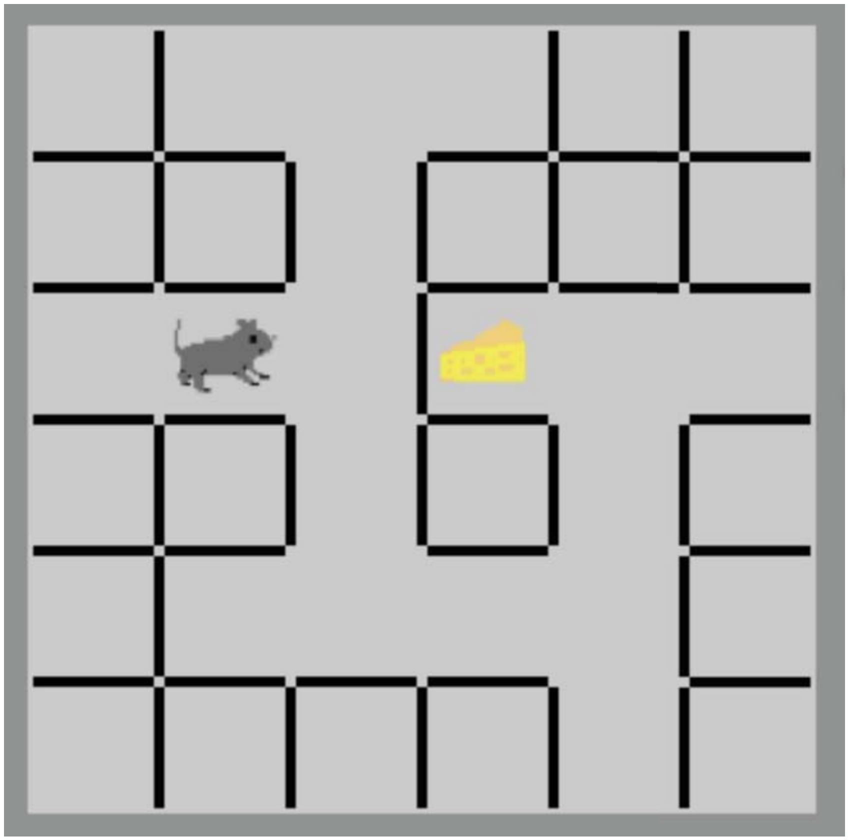

The rat is faced with multiple choices (T's) on their path to the food. This is a more challenging task to live rats. It also takes a longer training session for the artificial rat to build a map, and higher motivation to find a path to the food.

Live rats often fail to learn complex mazes

The rat's dilemma: Multiple choices

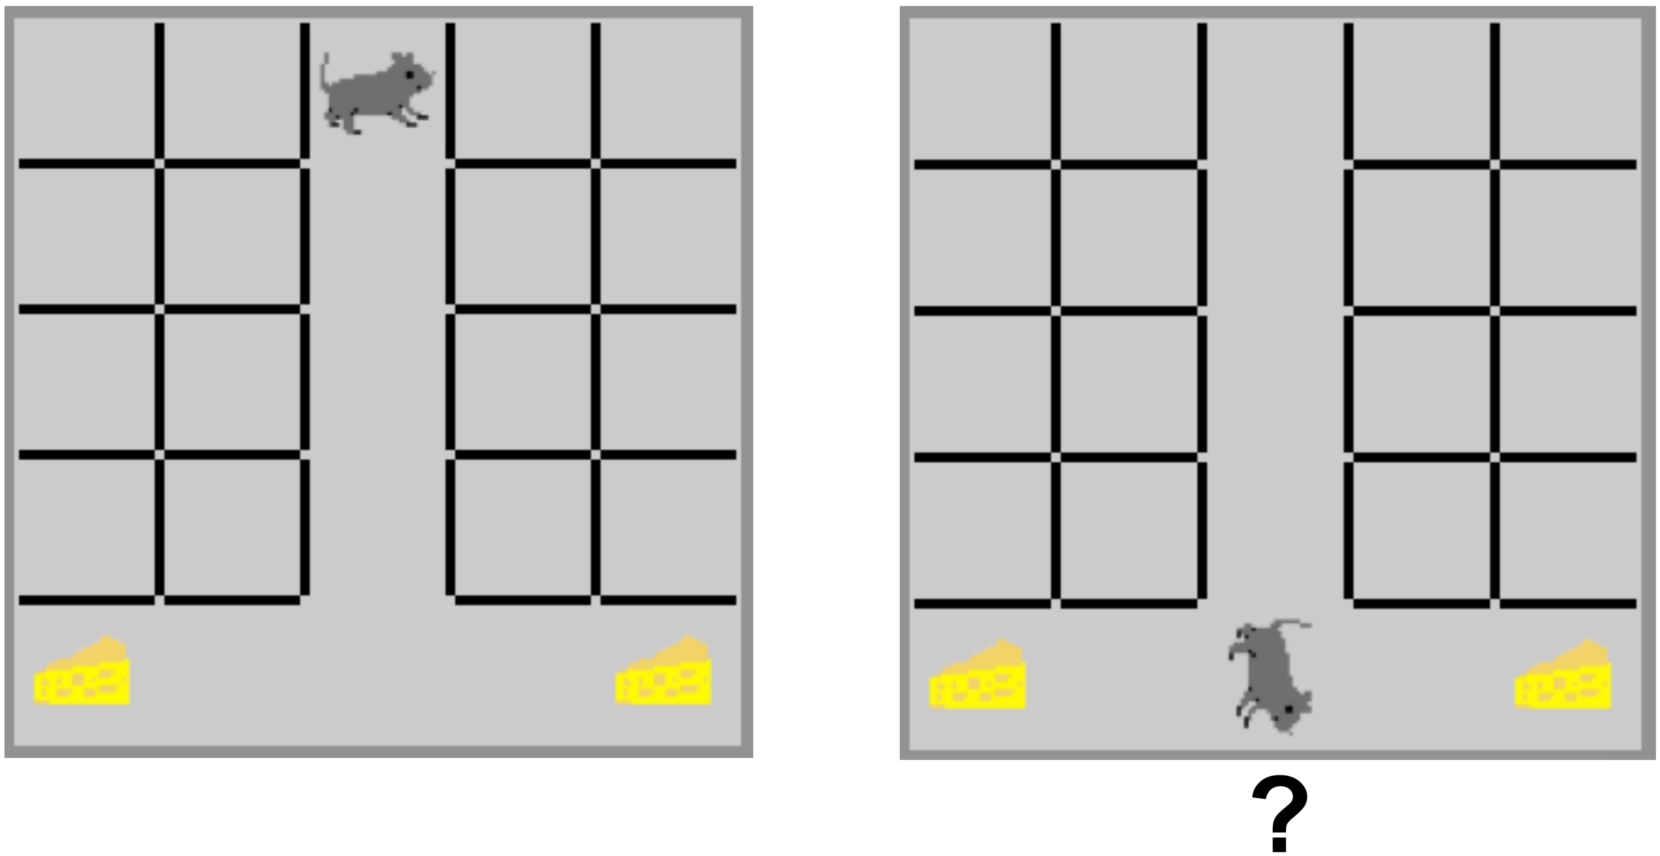

Problems with multiple choices

run 1

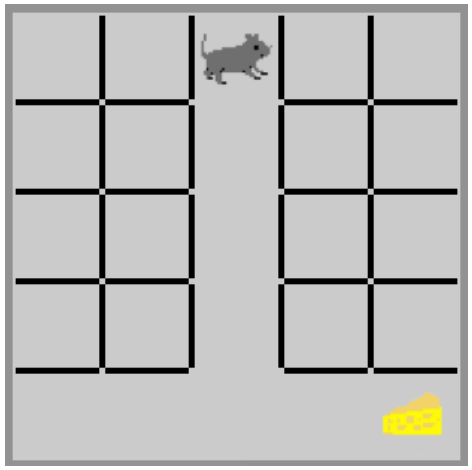

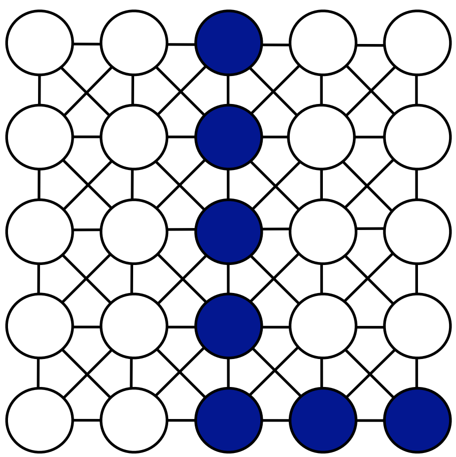

Problems with multiple choices

run 1

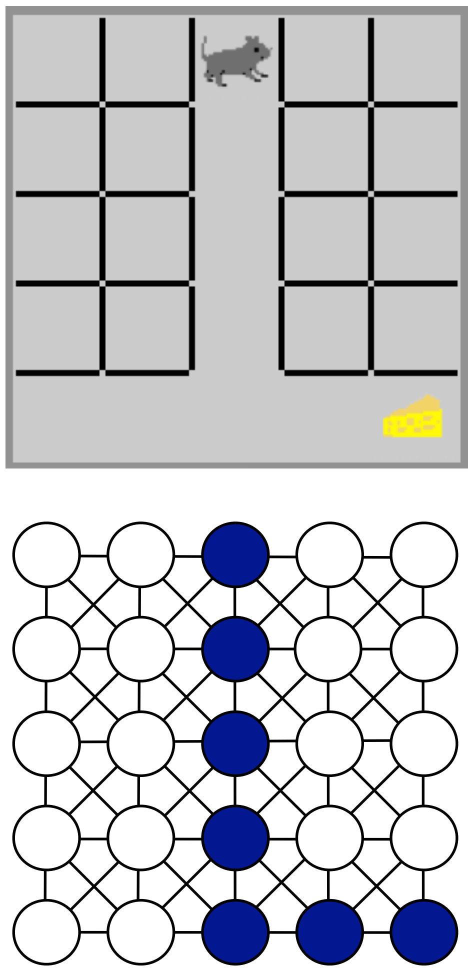

run 2

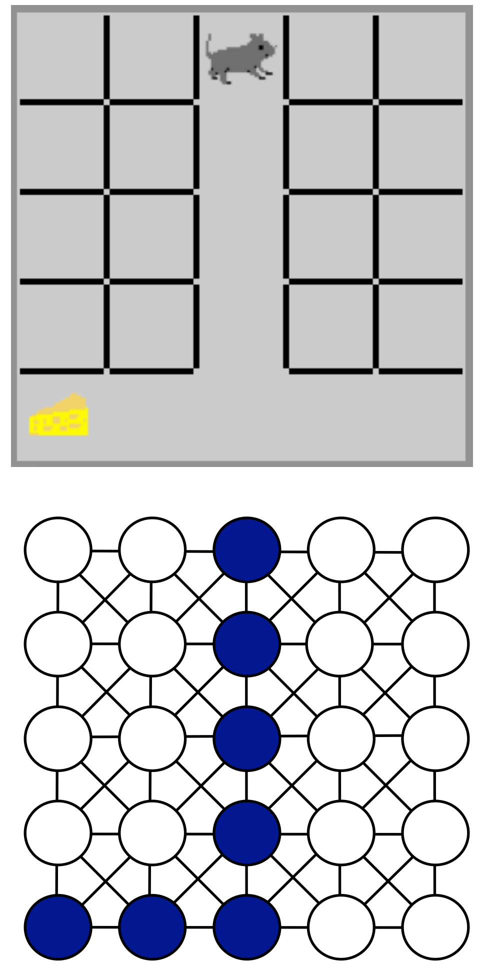

What's happening in the rat's brain?

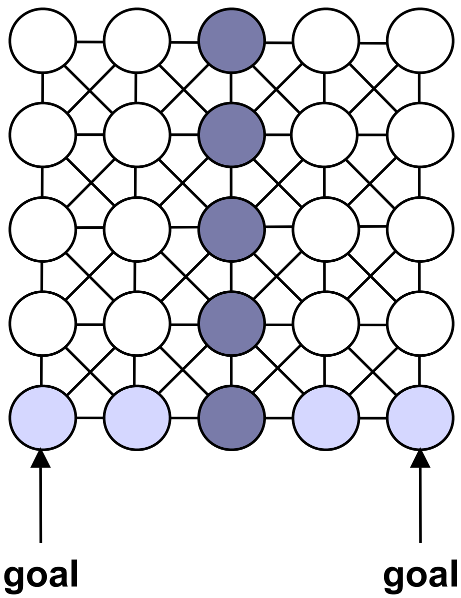

What's happening in the rat's brain?

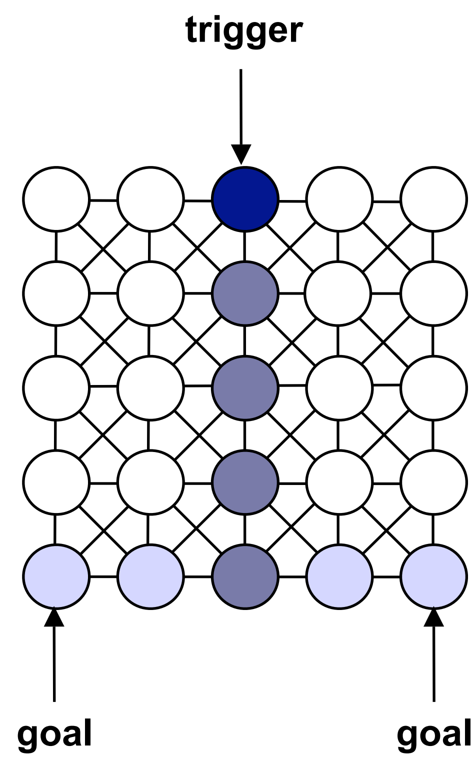

What's happening in the rat's brain?

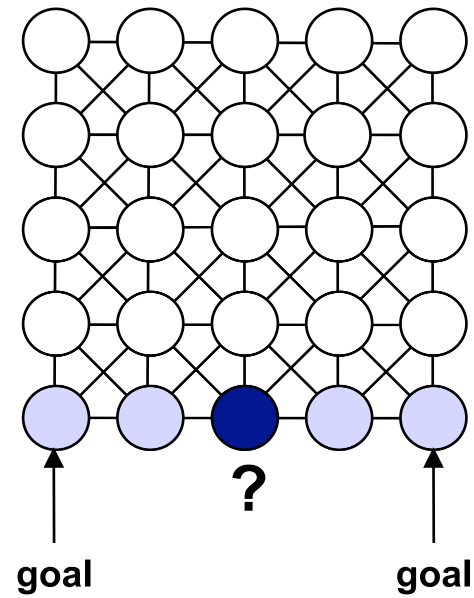

Adding context

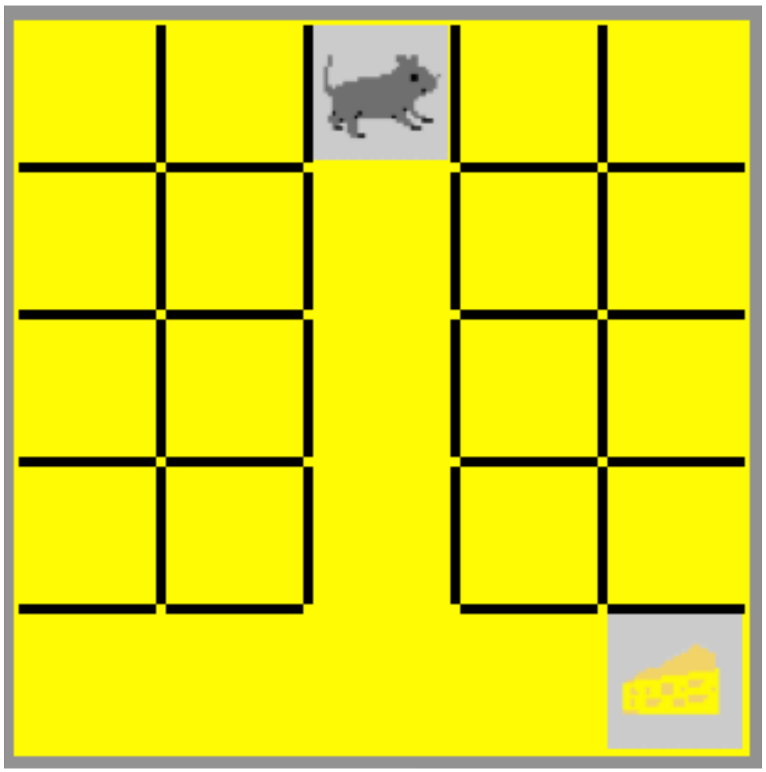

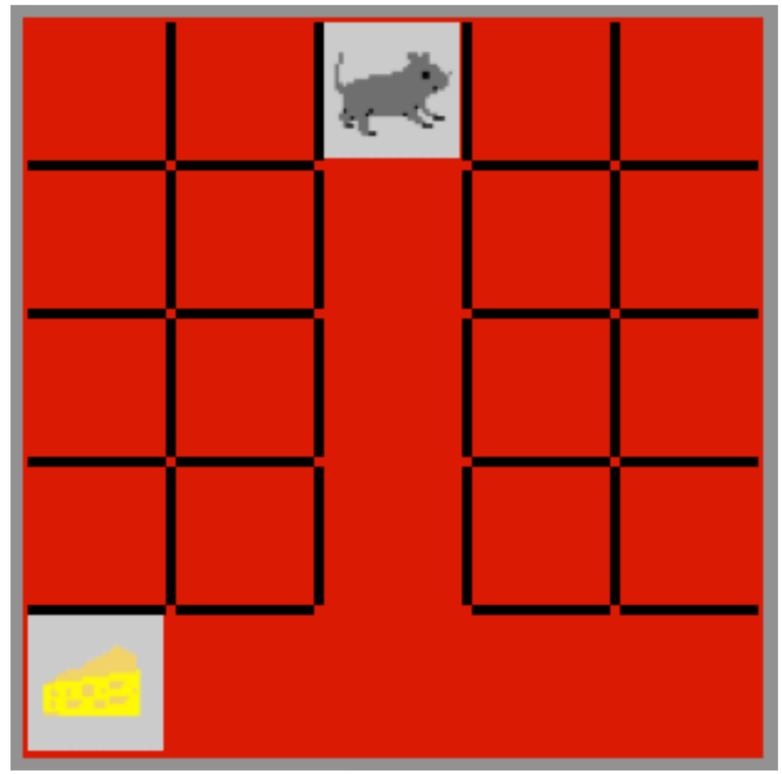

Learning with context

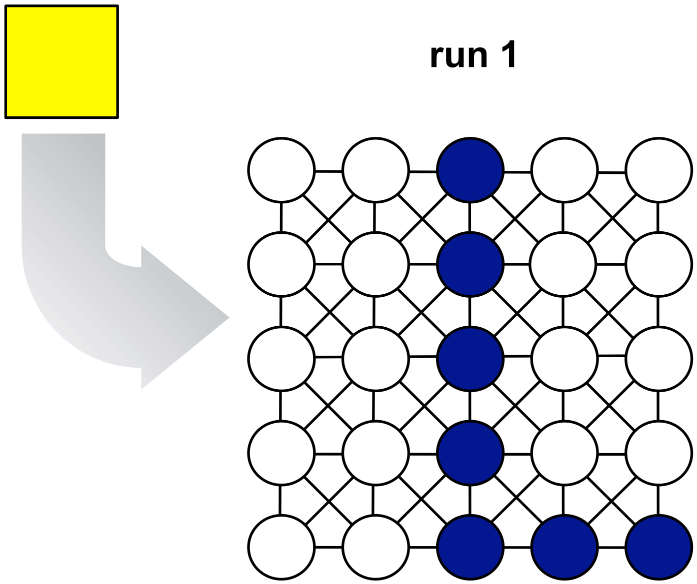

Learning with context

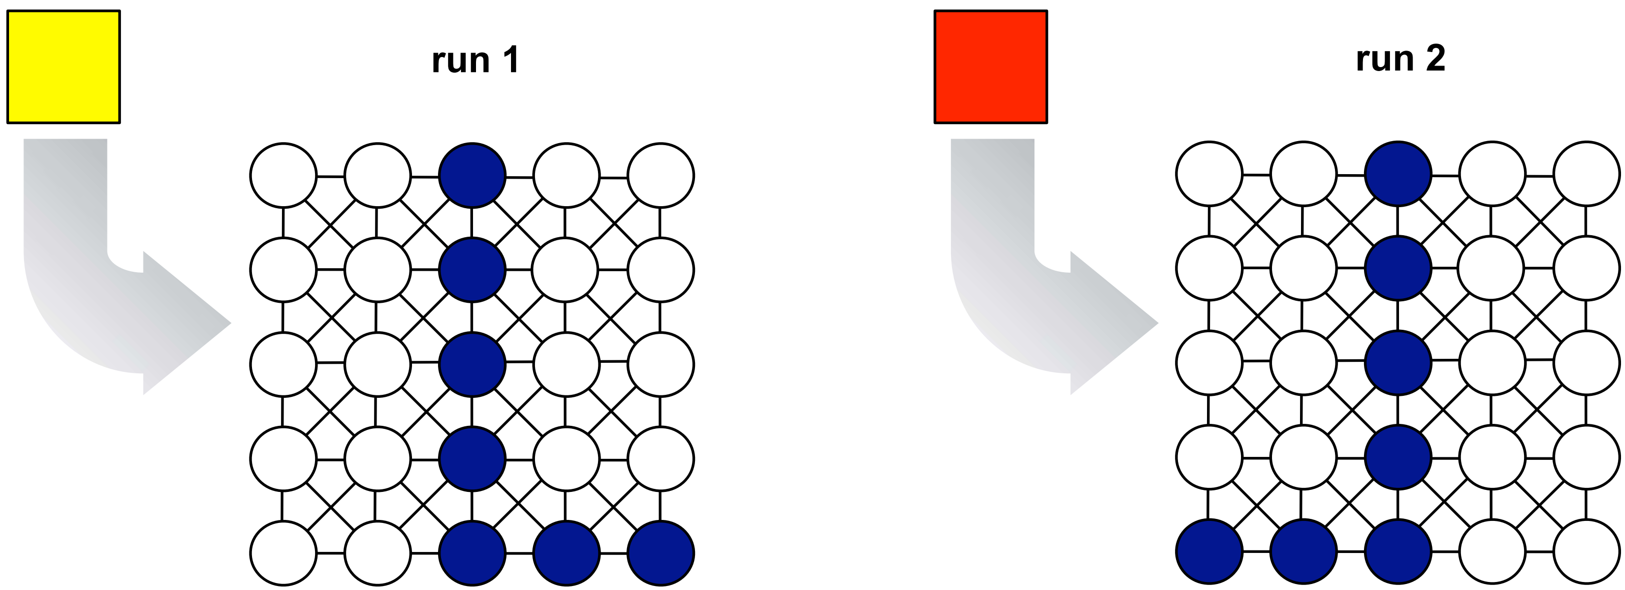

The context mechanism in the Neurosolver

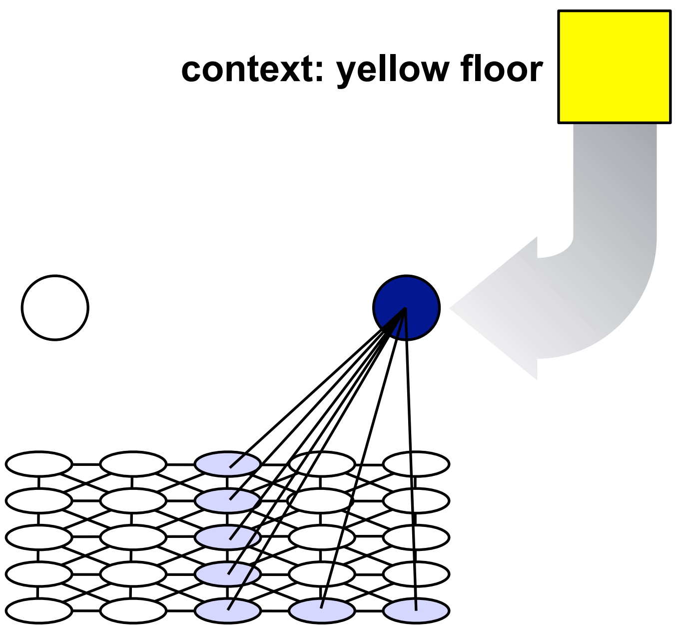

The context mechanism in the Neurosolver

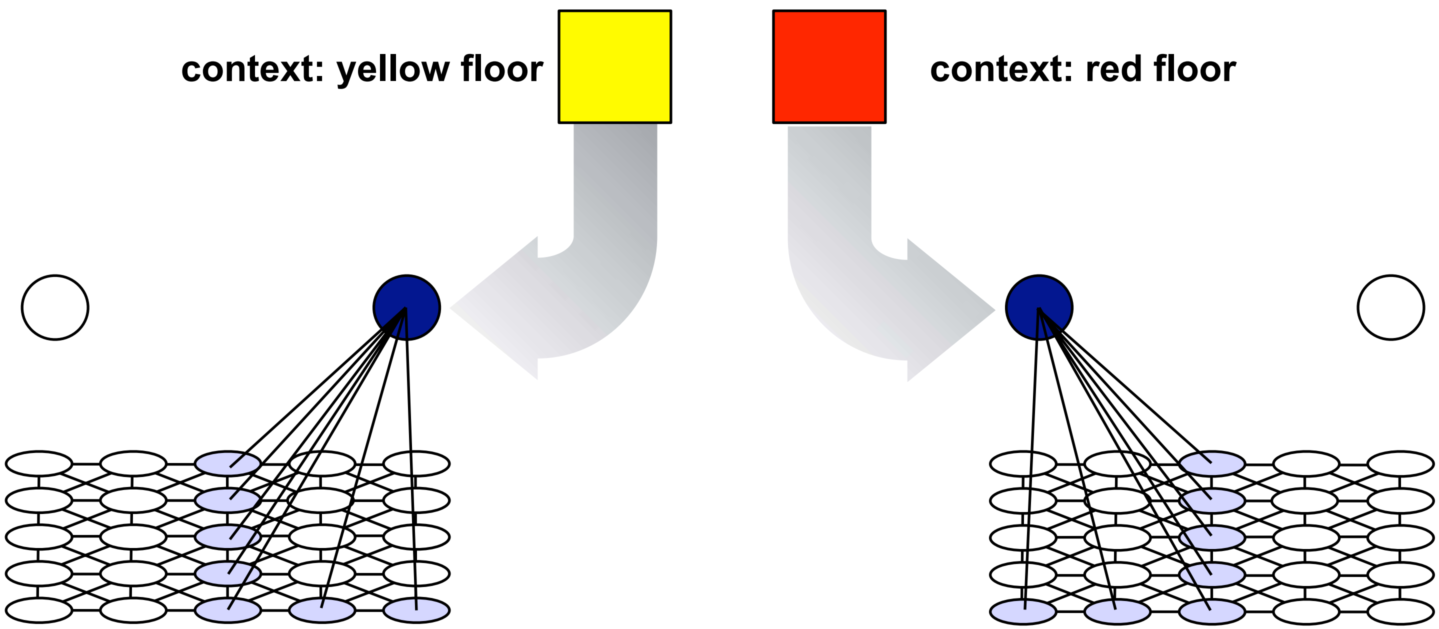

Context-modulated search

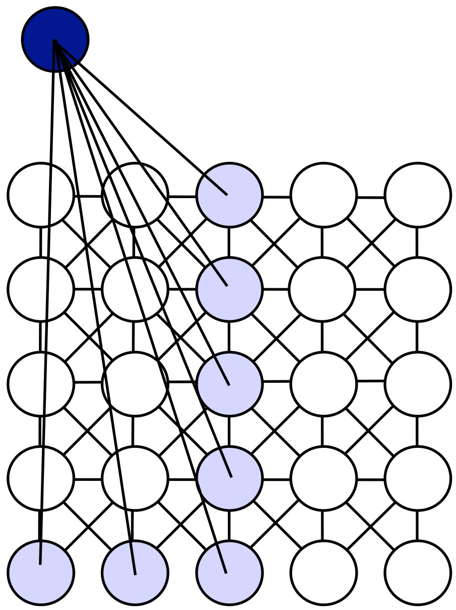

Context-modulated search

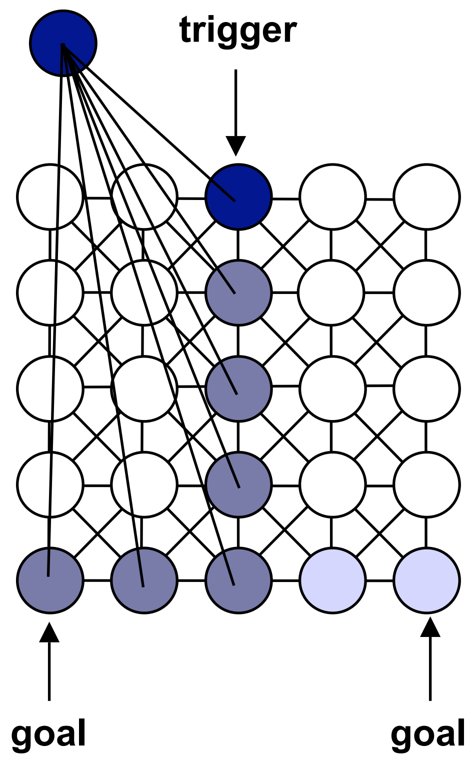

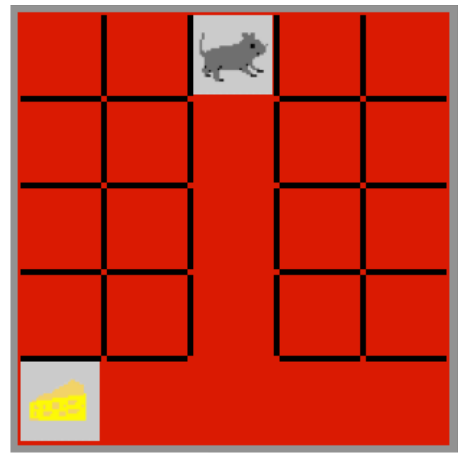

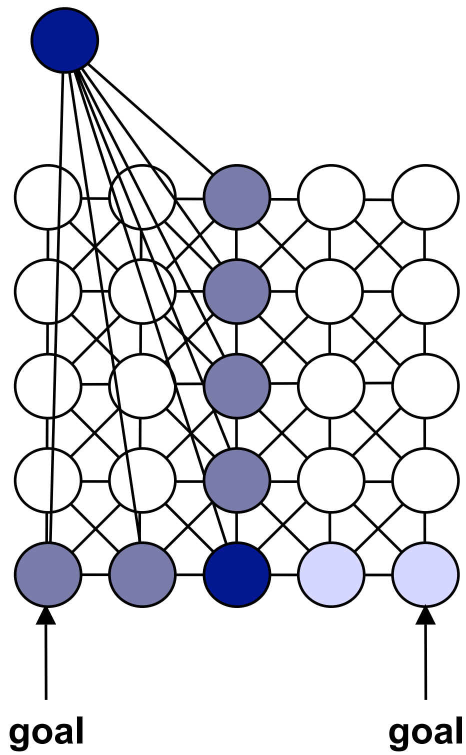
goal
goal

Running with the contextual activity

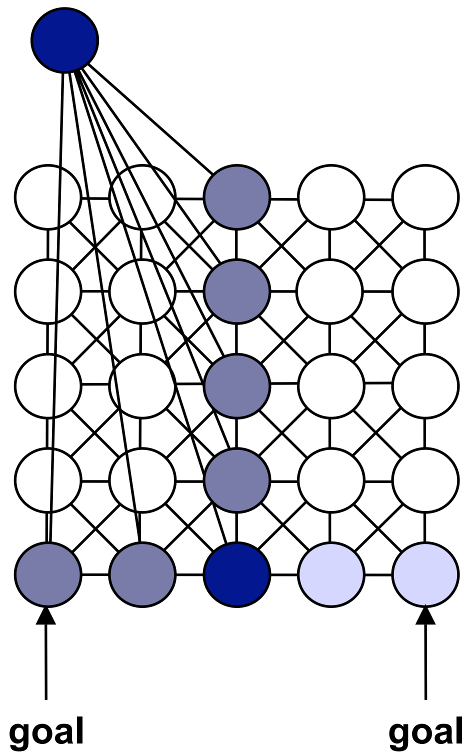

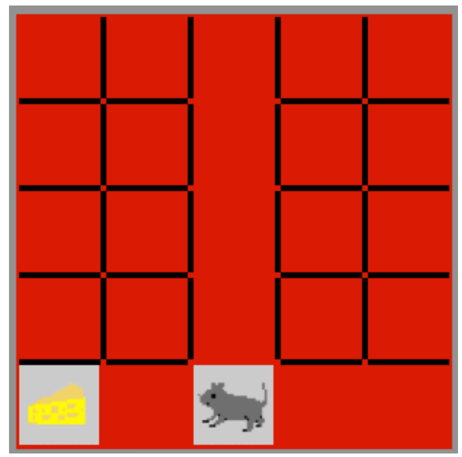

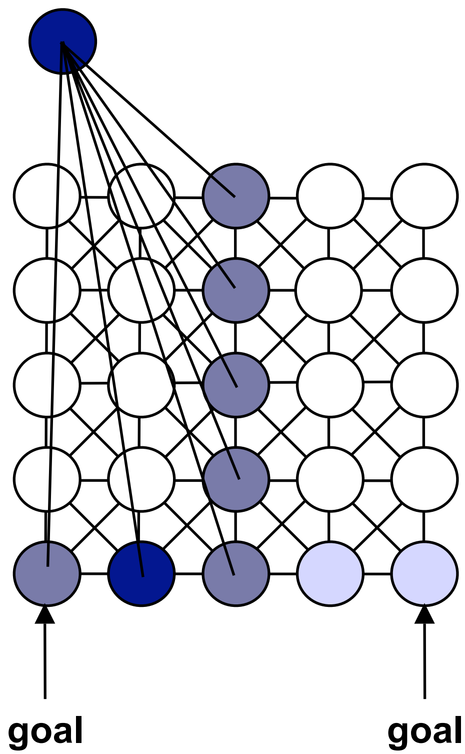
goal
goal

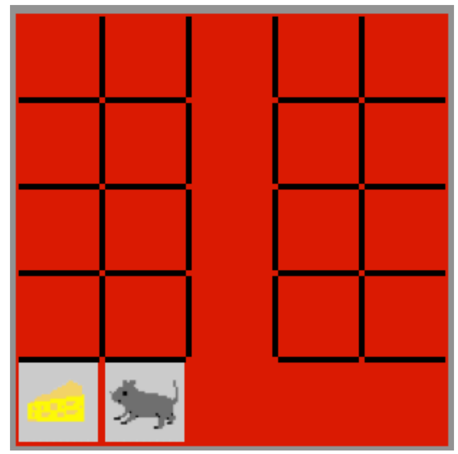

Running with the contextual activity

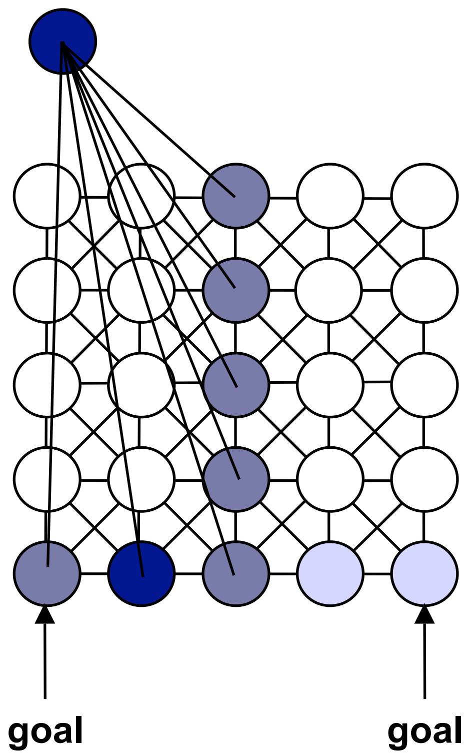

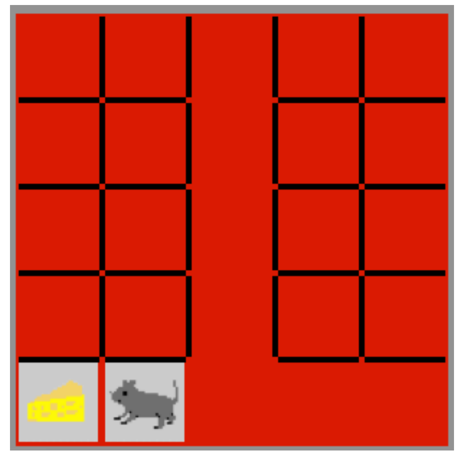

Rat maze simulator with context

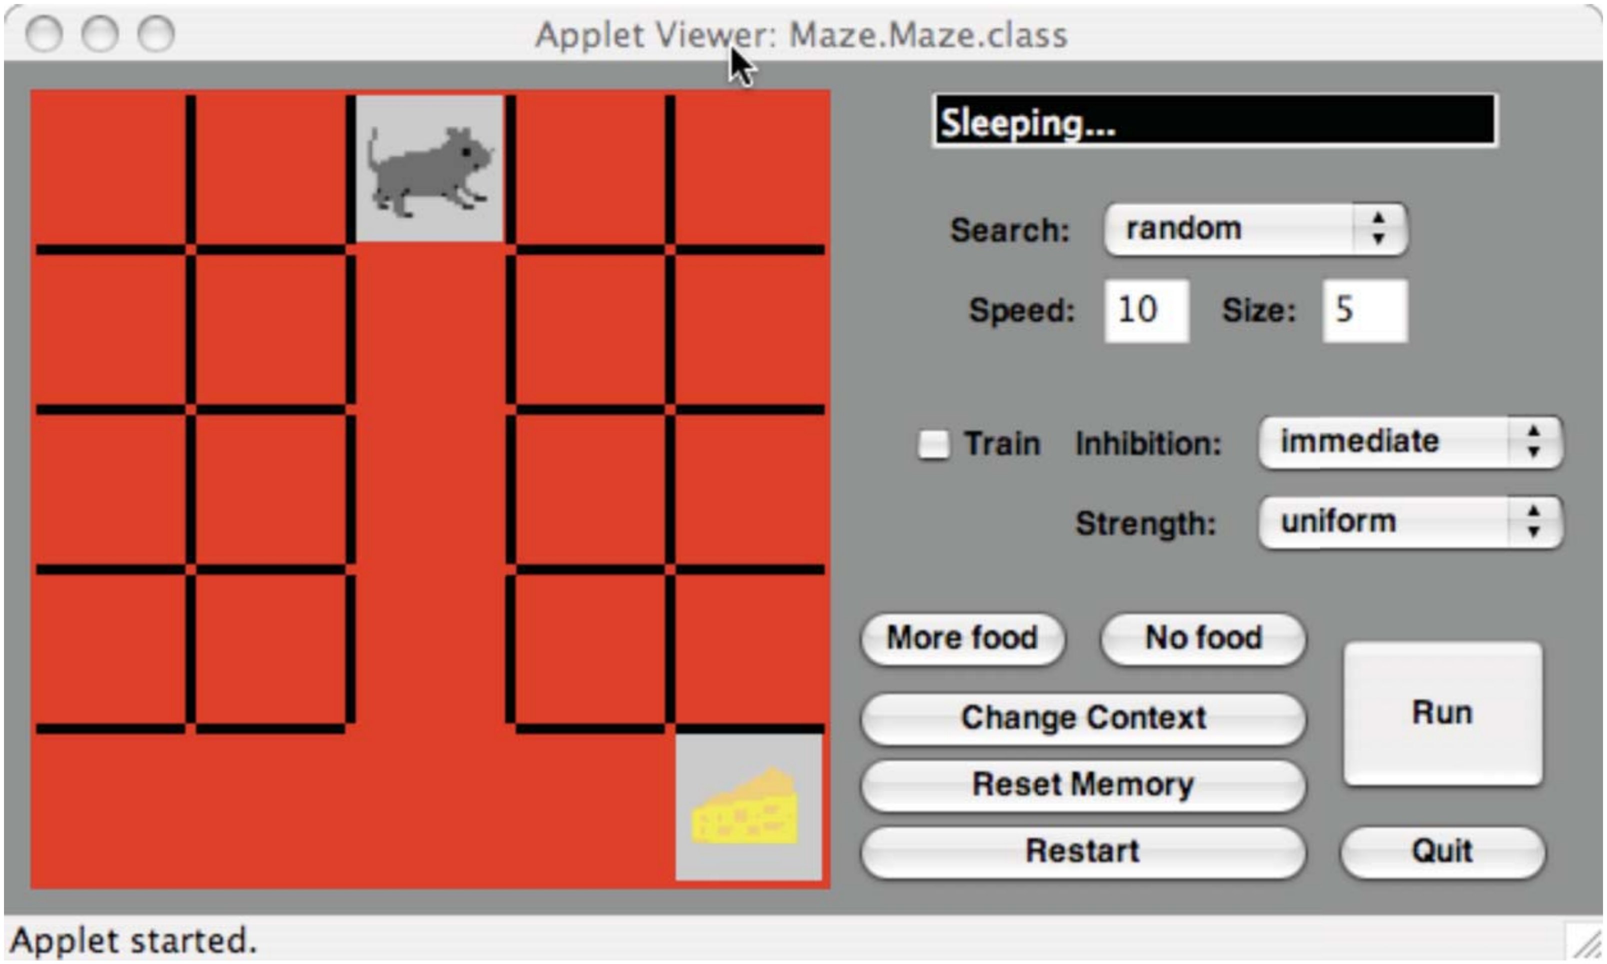

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable
    - rats may get frustrated - not the Neurosolver

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable
    - rats may get frustrated - not the Neurosolver
  - the Neurosolver always selects the best solution and performs

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable
    - rats may get frustrated - not the Neurosolver
  - the Neurosolver always selects the best solution and performs
    - uses the contextual cues if available

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable
    - rats may get frustrated - not the Neurosolver
  - the Neurosolver always selects the best solution and performs
    - uses the contextual cues if available
    - there is also a random selection mechanism built in to deal with race conditions

Conclusions

- In spite of dramatic simplifications in the model, similar behavior of the artificial rat to live rats running in mazes without stress
  - under stress, animal behavior becomes unpredictable
    - rats may get frustrated - not the Neurosolver
  - the Neurosolver always selects the best solution and performs
    - uses the contextual cues if available
    - there is also a random selection mechanism built in to deal with race conditions
- Neurosolver models the mechanism for path storage and recovery, but does not capture many other nuances that control the behavior of live rats.

TBDitF

- State space reduction
- Exploring application of functional areas
  - e.g., separate storage for place cells
    - place cells observed in hippocampus, but in the Neurosolver all is in one network mixed with hypercolumns
- Exploring new data from neuroscience
  - e.g., evident plasticity (programmability) in the hippocampus brings an idea of neuromorphic subroutines
    - the hippocampus produces maps (composed of place cells), but does not seem to store them
- Taxon navigation
  - real-time processing
- Continuous learning
- Goal management
- Alternative implementations
  - Software and hardware
- 3D model for the rat maze

On Biological Inspirations for Computer Science

Q&A

Dr. Andrzej (AJ) Bieszczad
California State University Channel Islands
aj.bieszczad@csuci.edu

EXTRAS

On Biological Inspirations for Computer Science

EXTRAS

Dr. Andrzej (AJ) Bieszczad
California State University Channel Islands
aj.bieszczad@csuci.edu

Probabilistic learning

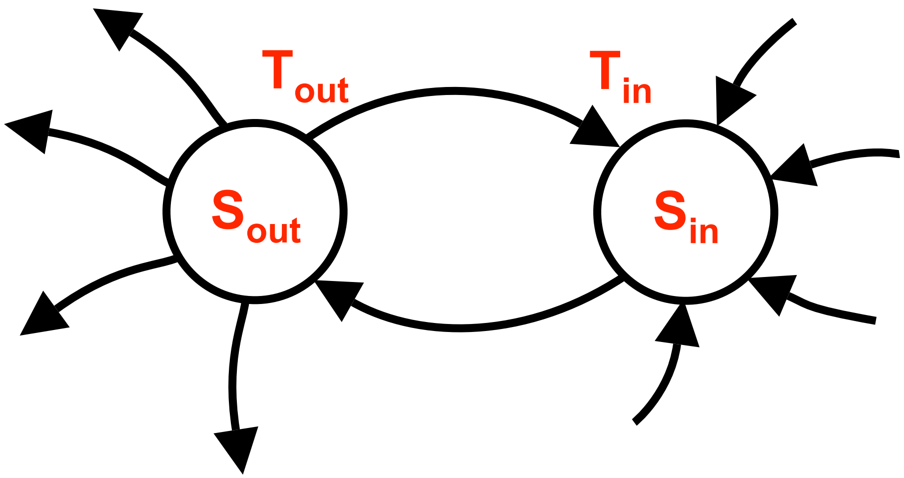

T_{out} - number of transmissions of an action potential
S_{out} - total number of cases when a division positively influenced other nodes
T_{in} - the number of times when an action potential transmitted over the connection contributed to the firing of the target node
S_{in} - the total number of times when any node positively influenced the node.

$$P = P_{out} \cdot P_{in} = (T_{out}/S_{out}) \cdot (T_{in}/$$

Probabilistic connection strength

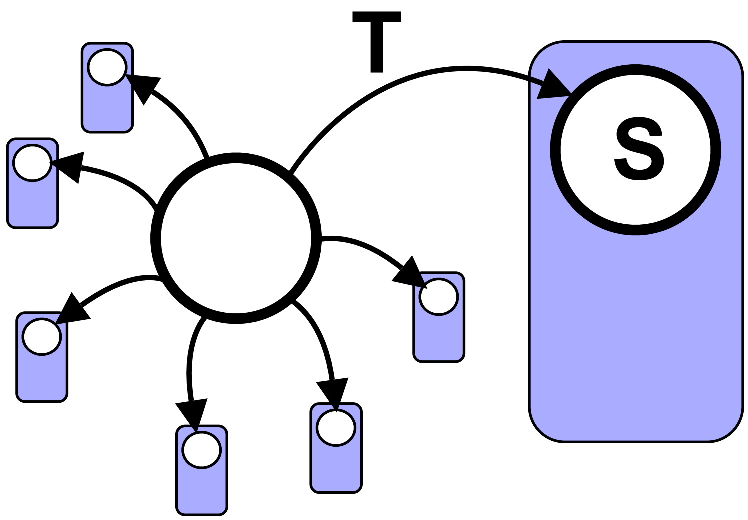

S - total number of columnar firings

T - number of contextual co-activations

$$P = T / S$$