Comparing Forgetting Algorithms for Artificial Episodic Memory Systems[1]

A. Nuxoll[i], D. Tecuci[ii], W.C. Ho[iii], & N. Wang[i]

University of Portland
[i]Dept. of Electrical Engineering and Computer Science
University of Texas at Austin
[ii]Dept. of Computer Science
University of Hertfordshire
[iii]School of Computer Science

March 2010

[1]Andrew Nuxoll et al. "Comparing Forgetting Algorithms for Artificial Episodic Memory Systems". In: *Proceedings of the Society for the Study of Artificial Intelligence and Simulation of Behavior (AISB) Convention* (2010).

Outline

Outline

Agent

An Agent, or Intelligent Agent,

- exists in an Environment,

Agent

An Agent, or Intelligent Agent,

- exists in an Environment,
- perceives through Sensors,

Agent

An Agent, or Intelligent Agent,

- exists in an Environment,
- perceives through Sensors,
- effects through Actuators.

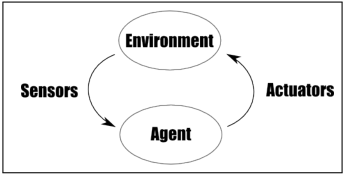

Figure: An agent.

Reflex Agents

Simple Reflex agents

- if-then rules,

Reflex Agents

Simple Reflex agents

- if-then rules,
- perform actions when a rule matches.

Reflex Agents

Simple Reflex agents

- if-then rules,
- perform actions when a rule matches.

Model-Based agents

- track the state of the environment.

Reflex Agents

Simple Reflex agents

- if-then rules,
- perform actions when a rule matches.

Model-Based agents

- track the state of the environment.

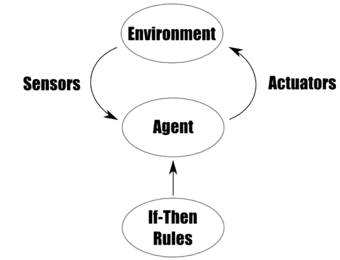

(a) A simple reflex agent.

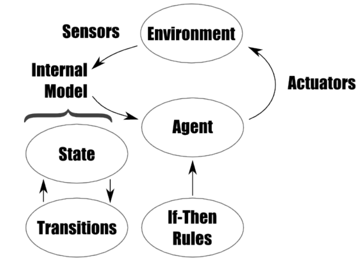

(b) A model-based reflex agent.

Figure: Depiction of simple reflex agent and model-based reflex agent.

Planning Agents

Planning agents

- model-based agents,

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Goal-Based planning agents

- have goals,

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Goal-Based planning agents

- have goals,
- use a knowledge base to construct a plan.

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Goal-Based planning agents

- have goals,
- use a knowledge base to construct a plan.

Utility-Based planning agents

- assign utility (preference) to actions and states,

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Goal-Based planning agents

- have goals,
- use a knowledge base to construct a plan.

Utility-Based planning agents

- assign utility (preference) to actions and states,
- maximize expected utility,

Planning Agents

Planning agents

- model-based agents,
- execute sequences of actions (a plan).

Goal-Based planning agents

- have goals,
- use a knowledge base to construct a plan.

Utility-Based planning agents

- assign utility (preference) to actions and states,
- maximize expected utility,
- may choose a plan with short-term low utility.

Outline

Multi-Agent Systems

Planning agents

- good in static environments,
- poor in dynamic environments.

Multi-Agent Systems

Planning agents

- good in static environments,
- poor in dynamic environments.
  - e.g. An environment with other agents (people, pets, etc.).

Multi-Agent Systems

Planning agents

- good in static environments,
- poor in dynamic environments.
  - e.g. An environment with other agents (people, pets, etc.).

In a Multi-Agent System, multiple agents interact within an environment.

Multi-Agent Systems

Planning agents

- good in static environments,
- poor in dynamic environments.
  - e.g. An environment with other agents (people, pets, etc.).

In a Multi-Agent System, multiple agents interact within an environment.

Coping

- Use reasoning.

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Types of reasoning:

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Types of reasoning:

- Speculative reasoning: contemplative and certain thinking; the use of reason to decide what to believe.

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Types of reasoning:

- Speculative reasoning: contemplative and certain thinking; the use of reason to decide what to believe.
- Practical reasoning: immediate and situational thinking; the use of reason to decide how to act.

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Types of reasoning:

- Speculative reasoning: contemplative and certain thinking; the use of reason to decide what to believe.
- Practical reasoning: immediate and situational thinking; the use of reason to decide how to act.
- Reactive reasoning: the absence of reasoning.

Deliberative Agents

Deliberative, or intentional, agents

- planning agents,
- use reasoning to choose actions.

Types of reasoning:

- Speculative reasoning: contemplative and certain thinking; the use of reason to decide what to believe.
- Practical reasoning: immediate and situational thinking; the use of reason to decide how to act.
- Reactive reasoning: the absence of reasoning.
- Hybrid reasoning: a combination of the above.

The BDI Model

The BDI (Belief, Desire, Intention) model[2] employs practical reasoning.

[2] Michael Bratman. *Intention, Plans, and Practical Reason*. 1987.

The BDI Model

The BDI (Belief, Desire, Intention) model[2] employs practical reasoning.

- Beliefs: (fallible) facts that represent the state of the environment.

[2] Michael Bratman. *Intention, Plans, and Practical Reason*. 1987.

The BDI Model

The BDI (Belief, Desire, Intention) model[2] employs practical reasoning.

- Beliefs: (fallible) facts that represent the state of the environment.
- Desires: aims the model is to accomplish, sometime in the future.

[2] Michael Bratman. *Intention, Plans, and Practical Reason*. 1987.

The BDI Model

The BDI (Belief, Desire, Intention) model[2] employs practical reasoning.

- Beliefs: (fallible) facts that represent the state of the environment.
- Desires: aims the model is to accomplish, sometime in the future.
- Intentions: *committed* aims the model is to accomplish, in the *near* future.

---

[2] Michael Bratman. *Intention, Plans, and Practical Reason*. 1987.

BDI Agents

BDI Agents

- deliberative agents,

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,
- aims as desires,

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,
- aims as desires,
- current action as intention,

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,
- aims as desires,
- current action as intention,
- multiple (non-conflicting) intentions,

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,
- aims as desires,
- current action as intention,
- multiple (non-conflicting) intentions,
- practical reasoning to choose next intention.

BDI Agents

BDI Agents

- deliberative agents,
- state as beliefs,
- aims as desires,
- current action as intention,
- multiple (non-conflicting) intentions,
- practical reasoning to choose next intention.

In a BDI agent, planning and doing are interleaved.

From the Literature

Some examples of BDI agents:

- Robots exhibiting flexible teamwork[3].
- Simulation of human crowds[4].
- Automated tourist guide & planner[5].

---

[3] G A Kaminka and I Frenkel. "Flexible Teamwork in Behavior-Based Robots". In: *Proceedings of the Twentieth National Conference on Artificial Intelligence (AAAI-05)* (2005).

[4] Ameya Shendarkar et al. "Crowd Simulation for Emergency Response Using BDI Agent Based on Virtual Reality". In: *Proceedings of the 38th Conference on Winter Simulation* (2006).

[5] Juan M Corchado et al. "Development of CBR-BDI Agents: A Tourist Guide Application". In: *International Journal of Computer Science and Applications* (2005).

Outline

Long-Term Memory

Long-term Memories are typically[6] divided[7] into two categories: procedural and declarative memories.

[6] J R Anderson. *Language, Memory and Thought*. 1976.

[7] Brenda Milner. "Les Troubles de la Mémoire Accompagnant des Lesions Hippocampiques Bilatérales". In: *Physiologie de l'Hippocampe* (1962).

Long-Term Memory

Long-term Memories are typically[6] divided[7] into two categories: procedural and declarative memories.

Procedural Memory: things one knows how to do; cognitive or motor skills.

[6] J R Anderson. *Language, Memory and Thought*. 1976.

[7] Brenda Milner. "Les Troubles de la Mémoire Accompagnant des Lesions Hippocampiques Bilatérales". In: *Physiologie de l'Hippocampe* (1962).

Long-Term Memory

Long-term Memories are typically[6] divided[7] into two categories: procedural and declarative memories.

Procedural Memory: things one knows how to do; cognitive or motor skills. Examples:

- Using a pencil.
- Riding a bicycle.

[6] J R Anderson. *Language, Memory and Thought*. 1976.

[7] Brenda Milner. "Les Troubles de la Mémoire Accompagnant des Lesions Hippocampiques Bilatérales". In: *Physiologie de l'Hippocampe* (1962).

Long-Term Memory

Long-term Memories are typically[6] divided[7] into two categories: procedural and declarative memories.

Procedural Memory: things one knows how to do; cognitive or motor skills. Examples:

- Using a pencil.
- Riding a bicycle.

Procedural memory is much like the if-then knowledge base of a reflex agent.

[6] J R Anderson. *Language, Memory and Thought*. 1976.

[7] Brenda Milner. "Les Troubles de la Mémoire Accompagnant des Lesions Hippocampiques Bilatérales". In: *Physiologie de l'Hippocampe* (1962).

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Types of declarative memory:

- Semantic Memory: factual information.

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Types of declarative memory:

- Semantic Memory: factual information. Example:
    - Recalling the definition of a word.

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Types of declarative memory:

- Semantic Memory: factual information. Example:
  - Recalling the definition of a word.
- Episodic Memory: specific events in time.

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Types of declarative memory:

- Semantic Memory: factual information. Example:
  - Recalling the definition of a word.
- Episodic Memory: specific events in time. Example:
  - Remembering that traffic is busy around 16:30.

Declarative Memory

Declarative Memory: things one knows; facts or knowledge.

Types of declarative memory:

- Semantic Memory: factual information. Example:
  - Recalling the definition of a word.
- Episodic Memory: specific events in time. Example:
  - Remembering that traffic is busy around 16:30.

Declarative memory is much like the knowledge base of a model-based agent.

HOMER

HOMER[8] is a BDI agent with episodic memory.

[8] S Vere and T Bickmore. "A Basic Agent". In: *Computational Intelligence* (1990).

HOMER

HOMER[8] is a BDI agent with episodic memory.

The agent is a simulated robotic submarine in a virtual two-dimensional "Seaworld".

[8] S Vere and T Bickmore. "A Basic Agent". In: *Computational Intelligence* (1990).

HOMER

HOMER[8] is a BDI agent with episodic memory.
The agent is a simulated robotic submarine in a virtual two-dimensional “Seaworld”.

HOMER is able to:

- Plan tasks around current/conflicting tasks.

[8] S Vere and T Bickmore. “A Basic Agent”. In: *Computational Intelligence* (1990).

HOMER

HOMER[8] is a BDI agent with episodic memory.
The agent is a simulated robotic submarine in a virtual two-dimensional "Seaworld".

HOMER is able to:

- Plan tasks around current/conflicting tasks.
- Examine its own episodic memory to answer user prompts.

[8] S Vere and T Bickmore. "A Basic Agent". In: *Computational Intelligence* (1990).

HOMER

HOMER[8] is a BDI agent with episodic memory.
The agent is a simulated robotic submarine in a virtual two-dimensional "Seaworld".

HOMER is able to:

- Plan tasks around current/conflicting tasks.
- Examine its own episodic memory to answer user prompts.

```
Example Dialogue
User:  Drop the package at the barge next Saturday at 9pm.
HOMER: Okay.
[HOMER Generates Plan]
User:  Are you going to the barge next Saturday?
HOMER: Yes.
User:  What time?
HOMER: 8:56pm.
```

[8] S Vere and T Bickmore. "A Basic Agent". In: *Computational Intelligence* (1990).

Outline

Neuromimetism

Amnesiac case-studies show that a lack of episodic memory is a severe handicap.

Neuromimetism

Amnesiac case-studies show that a lack of episodic memory is a severe handicap.

If an agent is given

- episodic storage and
- recall strategies,

what advantages are afforded?

Neuromimetism

Amnesiac case-studies show that a lack of episodic memory is a severe handicap.

If an agent is given

- episodic storage and
- recall strategies,

what advantages are afforded?

Such an agent could

- make better decisions,

Neuromimetism

Amnesiac case-studies show that a lack of episodic memory is a severe handicap.

If an agent is given

- episodic storage and
- recall strategies,

what advantages are afforded?

Such an agent could

- make better decisions,
- detect problems in plans,

Neuromimetism

Amnesiac case-studies show that a lack of episodic memory is a severe handicap.

If an agent is given

- episodic storage and
- recall strategies,

what advantages are afforded?

Such an agent could

- make better decisions,
- detect problems in plans,
- monitor long-term goals.

Forgetting

We note the following from observation:

- Humans routinely forget. Is the information
  - lost?
  - unretrievable?

[9] J T Wixted and E B Ebbesen. "Genuine Power Curves in Forgetting: A Quantitative Analysis of Individual Subject Forgetting Functions". In: *Memory & Cognition* (1997).

[10] J W Alba and L Hasher. "Is Memory Schematic?" In: *Psychological Bulletin 93* (1983).

Forgetting

We note the following from observation:

- Humans routinely forget. Is the information
  - lost?
  - unretrievable?
- Memory performance declines over time[9]
  - decay of unused information,
  - interference of new and old information.

---

[9] J T Wixted and E B Ebbesen. "Genuine Power Curves in Forgetting: A Quantitative Analysis of Individual Subject Forgetting Functions". In: *Memory & Cognition* (1997).

[10] J W Alba and L Hasher. "Is Memory Schematic?" In: *Psychological Bulletin 93* (1983).

Forgetting

We note the following from observation:

- Humans routinely forget. Is the information
  - lost?
  - unretrievable?
- Memory performance declines over time[9]
  - decay of unused information,
  - interference of new and old information.
- Memories are stored as schemata[10].

---

[9] J T Wixted and E B Ebbesen. "Genuine Power Curves in Forgetting: A Quantitative Analysis of Individual Subject Forgetting Functions". In: *Memory & Cognition* (1997).

[10] J W Alba and L Hasher. "Is Memory Schematic?" In: *Psychological Bulletin 93* (1983).

Anticipated Problems

- Growth is unbounded.

Anticipated Problems

- Growth is unbounded.
- Some episodes must be discarded.

Anticipated Problems

- Growth is unbounded.
- Some episodes must be discarded.
- Which features of an episode are important?

Outline

Methods of Forgetting

Methods for removing episodes:

Methods of Forgetting

Methods for removing episodes:

- remove,

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,
- decay.

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,
- decay.

Episode selection criteria:

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,
- decay.

Episode selection criteria:

- oldest,

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,
- decay.

Episode selection criteria:

- oldest,
- least used,

Methods of Forgetting

Methods for removing episodes:

- remove,
- combine,
- decay.

Episode selection criteria:

- oldest,
- least used,
- redundant.

Remove Only

Only episodic removal is investigated, not decay or combine.

Three Architectures

Different episodic memory systems are considered:

- Computational Autobiographic[11],
- Generic Memory[12],
- Soar Episodic[13].

[11] W C Ho, K Dautenhahn, and C L Nehaniv. "Computational Memory Architectures for Autobiographic Agents Interacting in a Complex Environment: A Working Model". In: *Connection Science* (2008).

[12] Dan Tecuci. *A Generic Memory Model for Events*. 2007.

[13] Andrew M Nuxoll and John E Laird. "Extending Cognitive Architecture with Episodic Memory". In: *Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-07)* (2009).

The Algorithms

Three algorithms are described:

The Algorithms

Three algorithms are described:

- random,

The Algorithms

Three algorithms are described:

- random,
- oldest,

The Algorithms

Three algorithms are described:

- random,
- oldest,
- activation (frequency and recency).

Outline

Warehouse Domain

Environment criteria:

- Very large number of possible states.

Warehouse Domain

Environment criteria:

- Very large number of possible states.
- Large degree of success/failure.

Warehouse Domain

Environment criteria:

- Very large number of possible states.
- Large degree of success/failure.
- Optimal solution is difficult, better-than-random easy.

Warehouse Domain

Environment criteria:

- Very large number of possible states.
- Large degree of success/failure.
- Optimal solution is difficult, better-than-random easy.
- Unlimited running time (no goal or equilibrium state).

Warehouse Domain

Environment criteria:

- Very large number of possible states.
- Large degree of success/failure.
- Optimal solution is difficult, better-than-random easy.
- Unlimited running time (no goal or equilibrium state).

The investigated environment is a combination of the knapsack problem and the scheduling problem; both NP-Hard.

Variables

The authors consider three variables:

- System;
  - Autobiographic,
  - Generic,
  - Soar.

Variables

The authors consider three variables:

- System;
  - Autobiographic,
  - Generic,
  - Soar.
- Algorithm;
  - Random,
  - Oldest,
  - Activation.

Variables

The authors consider three variables:

- System;
  - Autobiographic,
  - Generic,
  - Soar.
- Algorithm;
  - Random,
  - Oldest,
  - Activation.
- Episode memory size (number of episodes stored).
  - 10 - 10000 episodes.

Variables

The authors consider three variables:

- System;
  - Autobiographic,
  - Generic,
  - Soar.
- Algorithm;
  - Random,
  - Oldest,
  - Activation.
- Episode memory size (number of episodes stored).
  - 10 - 10000 episodes.

Measure

- Difference from the optimal solution.

Variables

The authors consider three variables:

- System;
  - Autobiographic,
  - Generic,
  - Soar.
- Algorithm;
  - Random,
  - Oldest,
  - Activation.
- Episode memory size (number of episodes stored).
  - 10 - 10000 episodes.

Measure

- Difference from the optimal solution.

Aggregate

- Data averaged over 10 runs.

Outline

Memory Size

Performance

- Inverse-power function.

Memory Size

Performance

- Inverse-power function.
- Poor for small memory size.

Memory Size

Performance

- Inverse-power function.
- Poor for small memory size.
- Good for large memory size.

Memory Size

Performance

- Inverse-power function.
- Poor for small memory size.
- Good for large memory size.
- Little gain for additional memory.

System 1: Autobiographical

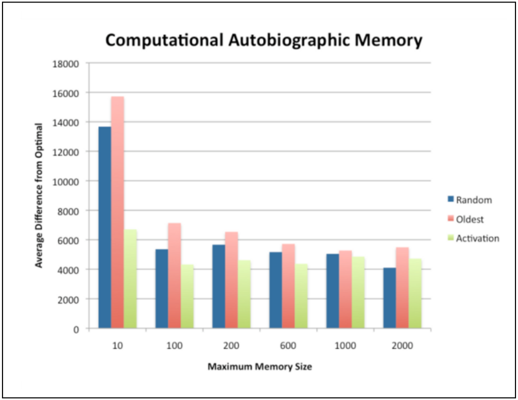

Figure: Results for the autobiographical system.

- Activation outperforms at low memory size.

System 1: Autobiographical

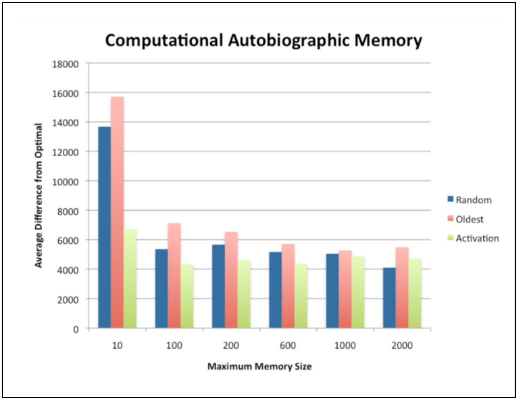

Figure: Results for the autobiographical system.

- Activation outperforms at low memory size.
- At higher memory sizes, all are comparable.

System 2: Generic

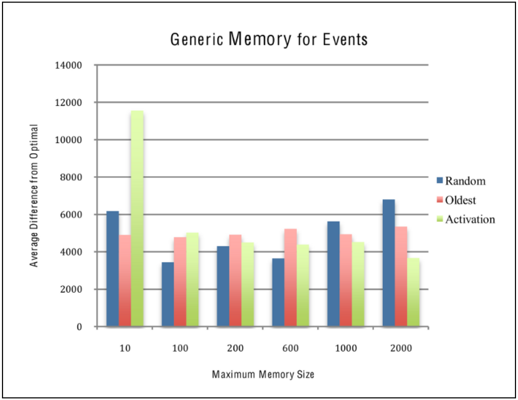

Figure: Results for the generic system.

- Activation algorithm was poor at low memory size.

System 2: Generic

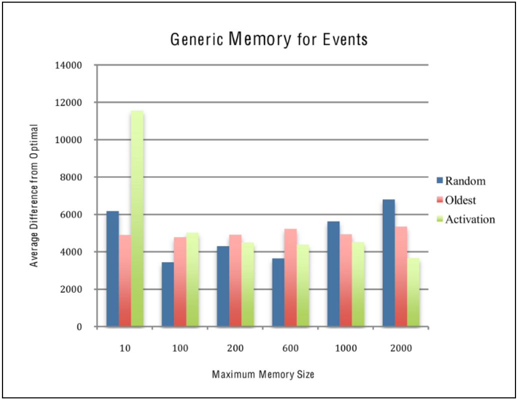

Figure: Results for the generic system.

- Activation algorithm was poor at low memory size.
- At higher memory sizes, all are comparable.

System 3: Soar

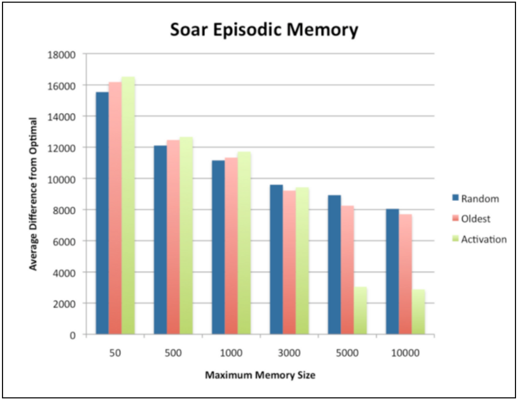

Figure: Results for the Soar system.

- All algorithms improve as memory increases.

System 3: Soar

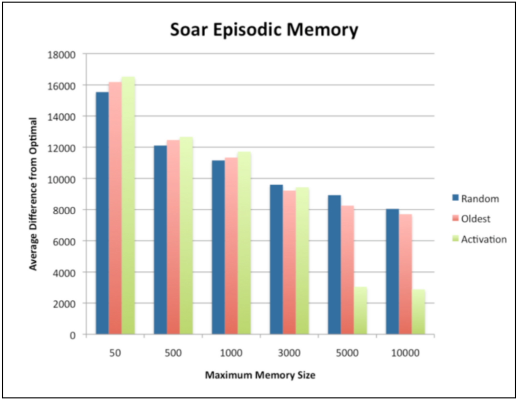

Figure: Results for the Soar system.

- All algorithms improve as memory increases.
- At higher memory size, activation *significantly* outperforms.

Summary

- Performance increases exponentially after a certain threshold in memory size.

Summary

- Performance increases exponentially after a certain threshold in memory size.
- This memory size threshold is different for each architecture.

Summary

- Performance increases exponentially after a certain threshold in memory size.
- This memory size threshold is different for each architecture.
  - Implies each domain needs a minimum number of stored episodes to be useful.

Summary

- Performance increases exponentially after a certain threshold in memory size.
- This memory size threshold is different for each architecture.
  - Implies each domain needs a minimum number of stored episodes to be useful.
- This memory size threshold is partially determined by the choice of algorithm.