AI Classical and Non-deterministic Planning: Model-based Autonomous Behavior

Sebastian Sardiña

School of Computing Technologies
RMIT University

Julio 28 - Agosto 1 2025

Part I

Introduction, Motivation, and AI Search

Part 1: Introduction, Motivation, and AI Search

Part 1: Introduction, Motivation, and AI Search

⚠️ Disclaimer / Descargo

Mixed-language warning

The talk will be in **Spanish**, but the slides are in **English**.
Sometimes I'll switch languages mid-sentence sin darme cuenta.

¿Por qué?

Soy argentino 🇦🇷, pero vivo hace muchos años afuera 🌏.
Enseño en inglés, pienso en pseudo-español, y me expreso en Spanglish.
Básicamente, no hablo bien **ninguno** de los dos idiomas 😅.
Pero tranqui, ¡igual nos vamos a entender!

Survival tips 🧠

- Don't worry, the concepts are the same in any idioma.
- Ask if you get lost (en cualquiera de los dos idiomas).

References

- S. Russell and P. Norvig. *Artificial Intelligence : A Modern approach*, Pearson. 4th, 2021.
- H. Geffner, B. Bonet. *A Concise Introduction to Models and Methods for Automated Planning*. Morgan & Claypool. 2013.
- Ghallab, M., Nau, D. & Traverso, P. 2004. *Automated Planning: Theory and Practice*. Elsevier.
- Patrik Haslum, Nir Lipovetzky, Daniele Magazzeni, Christian Muise: *An Introduction to the Planning Domain Definition Language*. Synthesis Lectures on Artificial Intelligence and Machine Learning, Morgan & Claypool Publishers 2019.
- **Other:** papers referenced in slides (slides available in Moodle)

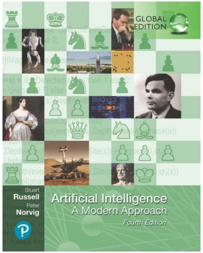

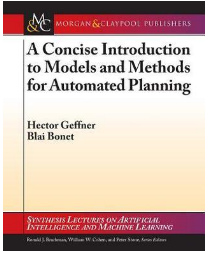

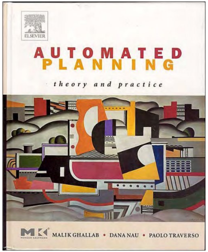

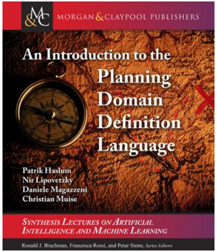

AI Classical and Non-deterministic Planning

*This course will survey **Automated Planning** as a model-based AI approach to sequential decision making, from the classical formulation to the more general variant with non-determinism that relates to SE formal methods.*

Special thanks to (and others!):

Hector Geffner @ RWTH Aachen University

Nir Lipovetzky @ Uni. of Melbourne

Course site: https://ssardina.github.io/courses/eci25/

Course Structure: 4 parts in 5 days

- **Part 1: Introduction, Motivation, and AI Search**
  - Introduction & Motivation: State of AI research.
  - AI Search: Uninformed Methods.
- **Part 2: Classical Planning: Languages**
  - Informed Search and Heuristics.
  - The Classical Model.
  - Planning languages: STRIPS and PDDL.
- **Part 3: Classical Planning: Methods and Algorithms**
  - Complexity of Planning.
  - Heuristic-based methods.
  - SAT-based solvers for planning.
- **Part 4: Non-deterministic Planning**
  - FOND Planning & solution concepts.
  - Methods for FOND Planning.
  - Conditional Fairness.

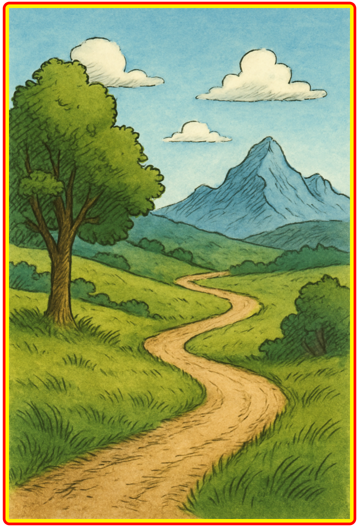

Course Structure: 4 parts in 5 days

- **Part 1: Introduction, Motivation, and AI Search**
  - Introduction & Motivation: State of AI research.
  - AI Search: Uninformed Methods.
- **Part 2: Classical Planning: Languages**
  - Informed Search and Heuristics.
  - The Classical Model.
  - Planning languages: STRIPS and PDDL.
- **Part 3: Classical Planning: Methods and Algorithms**
  - Complexity of Planning.
  - Heuristic-based methods.
  - SAT-based solvers for planning.
- **Part 4: Non-deterministic Planning**
  - FOND Planning & solution concepts.
  - Methods for FOND Planning.
  - Conditional Fairness.

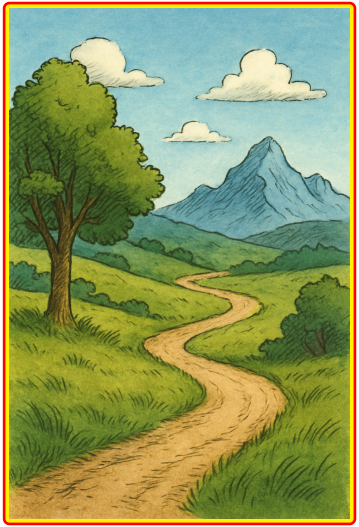

Course Structure: 4 parts in 5 days

- **Part 1: Introduction, Motivation, and AI Search**
  - Introduction & Motivation: State of AI research.
  - AI Search: Uninformed Methods.
- **Part 2: Classical Planning: Languages**
  - Informed Search and Heuristics.
  - The Classical Model.
  - Planning languages: STRIPS and PDDL.
- **Part 3: Classical Planning: Methods and Algorithms**
  - Complexity of Planning.
  - Heuristic-based methods.
  - SAT-based solvers for planning.
- **Part 4: Non-deterministic Planning**
  - FOND Planning & solution concepts.
  - Methods for FOND Planning.
  - Conditional Fairness.

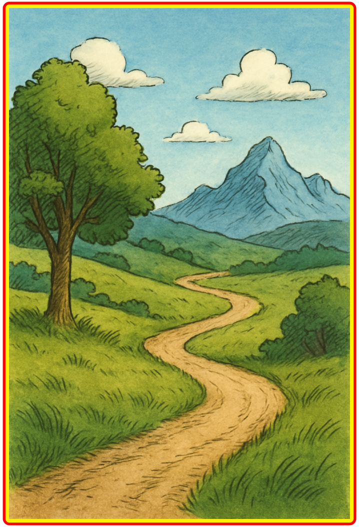

Course Structure: 4 parts in 5 days

- **Part 1: Introduction, Motivation, and AI Search**
  - Introduction & Motivation: State of AI research.
  - AI Search: Uninformed Methods.
- **Part 2: Classical Planning: Languages**
  - Informed Search and Heuristics.
  - The Classical Model.
  - Planning languages: STRIPS and PDDL.
- **Part 3: Classical Planning: Methods and Algorithms**
  - Complexity of Planning.
  - Heuristic-based methods.
  - SAT-based solvers for planning.
- **Part 4: Non-deterministic Planning**
  - FOND Planning & solution concepts.
  - Methods for FOND Planning.
  - Conditional Fairness.

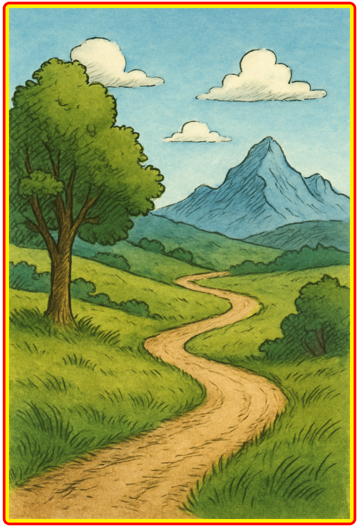

Plan for the rest of today

1. About me & us
2. State of AI research
3. AI search for sequential decision making

Course Structure: 4 parts in 5 days

- **Part 1: Introduction, Motivation, and AI Search**
  - Introduction & Motivation: State of AI research.
  - AI Search: Uninformed Methods.
- **Part 2: Classical Planning: Languages**
  - Informed Search and Heuristics.
  - The Classical Model.
  - Planning languages: STRIPS and PDDL.
- **Part 3: Classical Planning: Methods and Algorithms**
  - Complexity of Planning.
  - Heuristic-based methods.
  - SAT-based solvers for planning.
- **Part 4: Non-deterministic Planning**
  - FOND Planning & solution concepts.
  - Methods for FOND Planning.
  - Conditional Fairness.

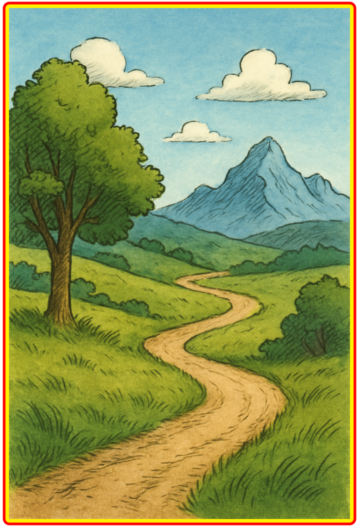

Part 1: Introduction, Motivation, and AI Search

Part 1: Introduction, Motivation, and AI Search


LET ME
INTRODUCE
MYSELF


LET ME
INTRODUCE
MYSELF

My CS journey started here!

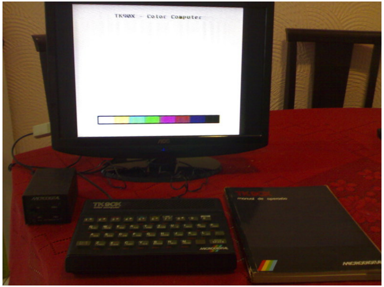

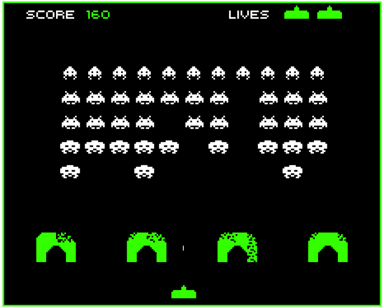

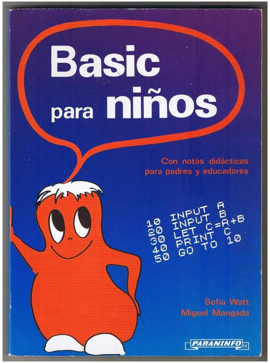

UNS (ARG) $\Longrightarrow$ UofT (CAN) $\Longrightarrow$ RMIT (AUS)

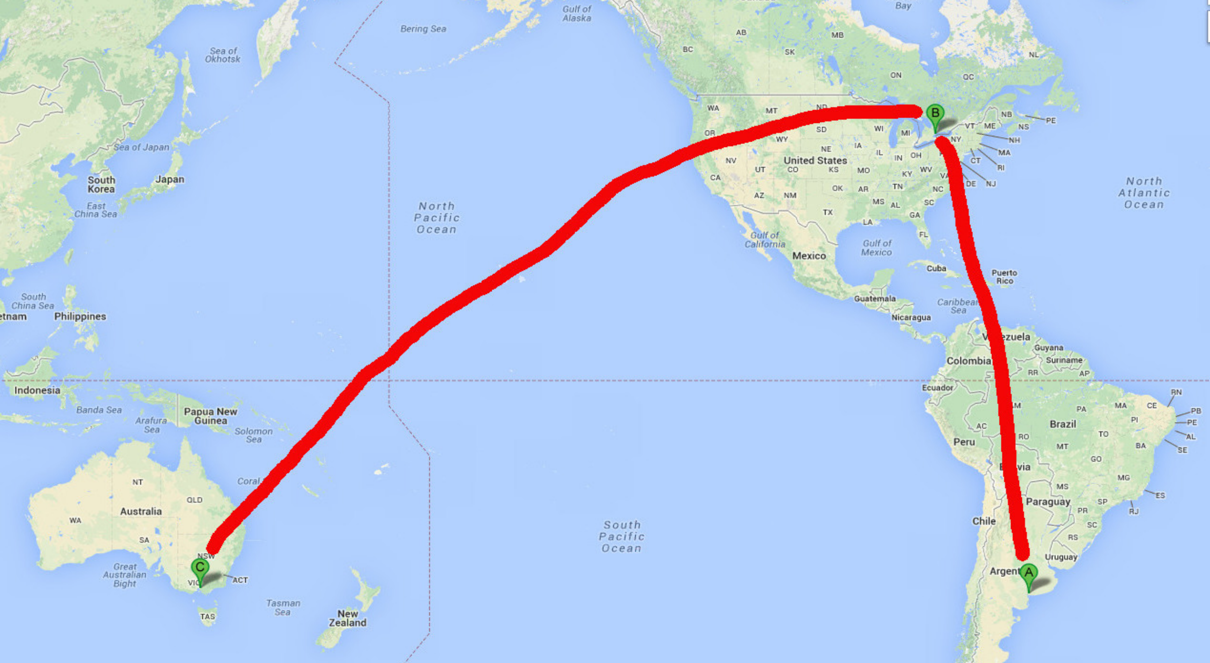

UNS (ARG) $\Longrightarrow$ UofT (CAN) $\Longrightarrow$ RMIT (AUS)

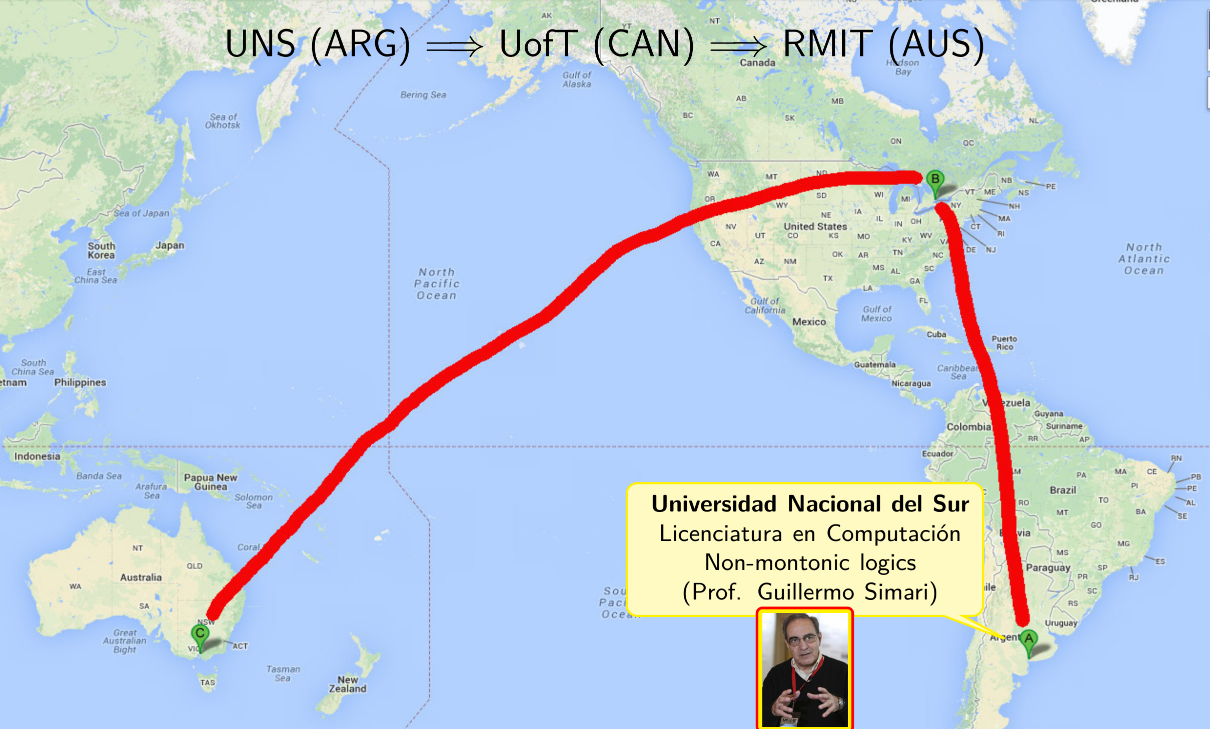

Universidad Nacional del Sur

Licenciatura en Computación

Non-montonic logics

(Prof. Guillermo Simari)

UNS (ARG) $\Longrightarrow$ UofT (CAN) $\Longrightarrow$ RMIT (AUS)

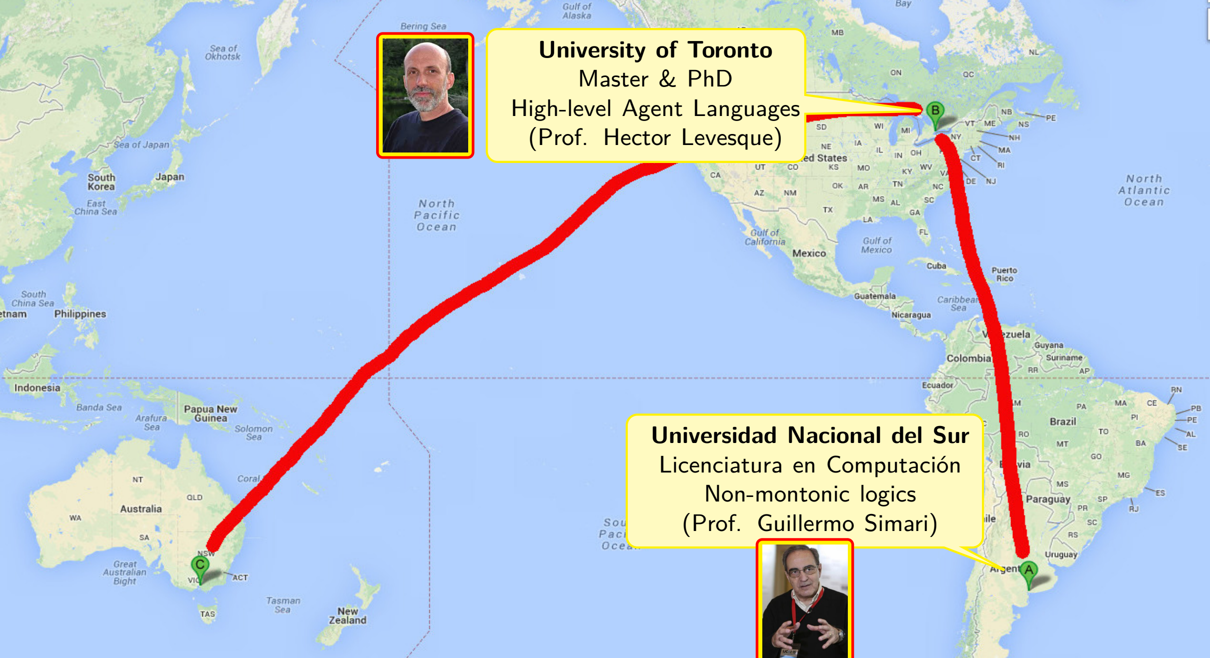

University of Toronto
Master & PhD
High-level Agent Languages
(Prof. Hector Levesque)

Universidad Nacional del Sur
Licenciatura en Computación
Non-montonic logics
(Prof. Guillermo Simari)

UNS (ARG) $\Longrightarrow$ UofT (CAN) $\Longrightarrow$ RMIT (AUS)

University of Toronto
Master & PhD
High-level Agent Languages
(Prof. Hector Levesque)

RMIT Unviersity
Postdoctoral
Planning in BDI Systems
(Prof. Lin Padgham)

Universidad Nacional del Sur
Licenciatura en Computación
Non-montonic logics
(Prof. Guillermo Simari)

Gracias.... totales!

- Founded in 1946 - 1956 (seventh national university created in the country).
- Structured in "Departments" (not Faculties!) 30,000+ students.

Gracias.... totales!

- Founded in 1946 - 1956 (seventh national university created in the country).
- Structured in "Departments" (not Faculties!) 30,000+ students.

- Started Computer Science in 1993 in the Math Department - *CS Dept. created in 1994!*
  - ➠ Graduated in 1997 (Thesis on Non-monotonic Logics).
- Tutor ("ayudate") 1994-1997 and head tutor ("JTP") 1997-1998.
- President of CeCom - Centro de Estudiantes de Computacion 1997-1998.
- Member of Departmental Council & University Assembly.

En defensa de la universidad pública...

En defensa de la universidad pública...

March 8th, 2024

Engineer Mr. Nicolás Posse, Chief of the Cabinet of Ministers
c.c. Dr. Daniel Salamone, President of CONICET
c.c. Members of the Board of Directors of CONICET

As members of the Computer Science international scientific community, we write to express our strong support to the Argentine scientific community in these difficult times. We are deeply concerned by the recent developments in Argentina in regards to how the prestigious Argentine national science and technology system has been brought to a standstill that undermines the country's science and technology sector due to the actions and inactions of your government.

We believe that decisions such as cutting PhD fellowships and promotions, withdrawing already committed funds to ongoing research projects, laying off administrative employees in research institutions, and freezing the investment in science in the context of high inflation levels have a short and long term devastating effect on the national scientific and technology system of Argentina.

Neglecting the role of the *state* in supporting science and technology is myopic and detrimentally affects the development possibilities of the country. There is extensive evidence that the *state*, by actively investing in science and technology when private investors found it too risky to do so, is a lead investor and key enabler of innovative knowledge and technologies that promote economic growth. In fact, the state has been behind the most significant technological advancements in our field that we see and enjoy today, from the search algorithms behind Google to the many technologies packed inside an iPhone to the wireless technology and the Internet itself, and to today AI-based technologies running Machine Learning algorithms.

Ignoring and disregarding the role of science and technology in modern society and the role of the state in promoting and fostering them is something a country cannot afford.

We ask you to listen to the Argentine scientific community's demands and actively work with their members towards preserving and improving a system that fosters the progress of the country's science and technology for the benefit of the nation.

Sincerely,

Prof. Sebastian Sardiña
RMIT University, AUSTRALIA

Prof. Hector Geffner
RWTH Aachen University, GERMANY
Alexander von Humboldt Professor in AI, AAAI and EurAI fellow

Senior Lecturer Dr. Damiano Spina
RMIT University, AUSTRALIA

Prof. Diego Calvanese
Free University of Bozen-Bolzano, ITALY
ACM Fellow, EurAI Fellow, AAIA Fellow

RMIT University

What does "RMIT" stand for? *What about the "R"?*

- Public university.
- Founded 1887 (training institute for workers).
- 80,000+ students.
- 3 campuses in Melbourne
  - 1 campus in Vietnam.
  - 1 center in Barcelona!
- Known for Art & Design, and Architecture.
- Also very strong in Engineering, Business and IT.

(click to see 1 min video)

RMIT University

RMIT = Royal Melbourne Institute of Technology 😲

- Public university.
- Founded 1887 (training institute for workers).
- 80,000+ students.
- 3 campuses in Melbourne
  - 1 campus in Vietnam.
  - 1 center in Barcelona!
- Known for Art & Design, and Architecture.
- Also very strong in Engineering, Business and IT.

(click to see 1 min video)

AI Innovation Lab

We are AI researchers who develop and extend solutions to further and enhance human capabilities, improving our quality of life through the use of artificial intelligence.

We target problems that have a direct impact, focusing on solving practical real-world problems by bringing the cutting-edge of AI to Industries including Transport, Food & Agriculture, and Advanced Manufacturing.

Research capabilities

Robotics & Human Collaboration

We focus on software technologies for intelligent collaboration between humans and robots, applying this to problems which are too dangerous or tedious for humans to complete themselves. We

Optimisation and Planning

We develop algorithms that find the optimal solutions and plans of action for complex problems. Our expertise includes nature-inspired and large-scale optimisation, operational research, machine

Autonomous Decision Systems

Many real-world problems are far too complex for a single human mind to handle, instead our capacity to solve these complex problems can only be enhanced by augmenting it with automated

My research/work

- Did my PhD at University of Toronto, 1998-2005.
  - *Supervised by Hector Levesque; Winograd schema challenge*
- Started at RMIT in July 2025 as postdoc; permanent academic since 2010
- **Teach** "foundational" CS courses:
  - Maths for CS (1st year)
  - Theory of Computation
  - Artificial Intelligence
  - Constraint Programming / Answer Set Programming
- **Research** areas/topics = KR $\cap$ Agents $\cap$ Planning
  - Cognitive Robotics / Agent programming
  - AI Planning
  - Goal/intention recognition
  - Behavior Composition

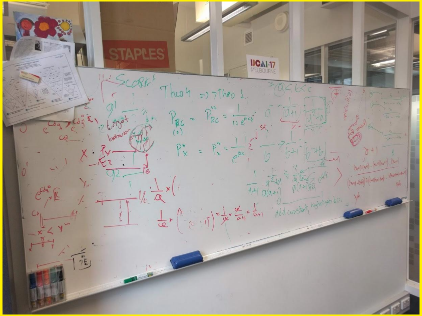

- Also contribute to Computational Thinking in the **community** (schools & centers, Victorian Curriculum, school teachers' professional development, etc.)

Who are we? Your turn!

2552 6250 @ menti.com

https://www.menti.com/al89ktgno9yf

Part 1: Introduction, Motivation, and AI Search

Part 1: Introduction, Motivation, and AI Search

A Bit of History: AI Programming and Problem of Generality

There was a time (60s, 70s, 80s) when AI was done mostly by **programming**:

1. pick up a challenging task and domain X (humor, story understanding, ...)
2. analyze/introspect/find out how task is solved
3. capture this reasoning in a program (usually knowledge base + rules)

A Bit of History: AI Programming and Problem of Generality

There was a time (60s, 70s, 80s) when AI was done mostly by **programming**:

1. pick up a challenging task and domain X (humor, story understanding, ...)
2. analyze/introspect/find out how task is solved
3. capture this reasoning in a program (usually knowledge base + rules)

Great ideas on programming and **AI programming**, but **methodological problem:** 👎

- ✖ Programs written by hand were clever but **not robust or general**.
- ✖ They worked on scenarios envisioned by programmer but **failed** on others.
- ✖ Difficult to **understand/debug** when failing: far from the actual problem/task.

AI Winter: the 80's

The rule+knowledge-based approach reached an **impasse** in the 80's, a time also of debates and controversies:

- **Good Old Fashioned AI** is 'rule application' but intelligence is not (J. Haugeland)

⚠️ Many criticisms of mainstream AI partially valid then; less valid now.

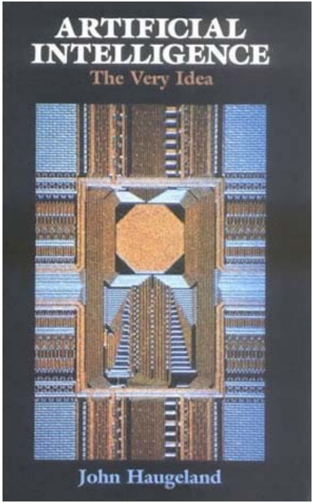

AI 90's - 2020

Formalization of AI techniques and increased use of mathematics. Recent issues of AIJ, JAIR, AAAI or IJCAI shows papers on:

1. SAT and Constraints
2. **Search and Planning** 👉
3. Probabilistic Reasoning
4. Probabilistic Planning
5. Inference in First-Order Logic
6. Machine Learning
7. Natural Language
8. Vision and Robotics
9. Multi-Agent Systems

✼ Areas 1 to 4 often deemed about techniques, but more accurate to regard them as models and solvers.

Motivation: Models and Solvers

$$\textit{Problem} \Longrightarrow \boxed{\textsc{Solver}} \Longrightarrow \textit{Solution}$$

Example

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?
- **Expressed as:** $J = 3P \;\; ; \;\; J + 10 = 2(P + 10)$
- **Solver:** Gauss-Jordan (Variable Elimination)
- **Solution:** $P = 10 \;\; ; \;\; J = 30$

Motivation: Models and Solvers

Problem $\Longrightarrow$ $\boxed{\textsc{Solver}}$ $\Longrightarrow$ *Solution*

Example

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?
- **Expressed as:** $J = 3P \;\; ; \;\; J + 10 = 2(P + 10)$
- **Solver:** Gauss-Jordan (Variable Elimination)
- **Solution:** $P = 10 \;\; ; \;\; J = 30$

➠ Solver is **general**: deals with any problem expressed as an instance of **model**.

➠ **Linear equations** model is too simple; **AI models** more challenging.

Motivation: Models and Solvers

Problem $\Longrightarrow$ $\boxed{\textsc{Solver}}$ $\Longrightarrow$ *Solution*

Example

- **Problem:** The age of John is 3 times the age of Peter. In 10 years, it will be only 2 times. How old are John and Peter?
- **Expressed as:** $J = 3P \;\; ; \;\; J + 10 = 2(P + 10)$
- **Solver:** Gauss-Jordan (Variable Elimination)
- **Solution:** $P = 10 \;\; ; \;\; J = 30$

➠ Solver is **general**: deals with any problem expressed as an instance of **model**.
➠ **Linear equations** model is too simple; **AI models** more challenging.

👍 **Models** good not just for **solving** but also for **understanding** problems.

From Programs to Solvers and Learners

- Generality problem increasingly led to methodological shift in 80s-90s:
  - from **programs** for **ill-defined problems** ...
  - to **algorithms** for **well-defined mathematical tasks**.
- New programs, **solvers** and **learners**, have a **crisp functionality**, and both can be seen as computing **functions** that map inputs into outputs

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f} \Longrightarrow \textit{Output } f(x)$$

- The algorithms are general: not tied to particular examples but to classes of **models** and **tasks** expressed in **mathematical form**.

Solvers (Reasoners)

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f} \Longrightarrow \textit{Output } f(x)$$

- **Solvers** derive output $f(x)$ for **given input** x from **model**:
  - **SAT:** x is a formula in CNF, $f(x) = 1$ if x satisfiable, else $f(x) = 0$.
  - **Classical planner:** x is a planning problem P, and $f(x)$ is plan that solves P. 👈
  - **Bayesian net:** x is a query over Bayes Net and $f(x)$ is the answer.
  - **Constraint satisfaction, Markov decision processes, POMDPs, ...**

✓ **Generality:** Solvers not tailored to particular examples.

✓ **Expressivity:** Some models very expressive; e.g., POMDPs.

✖ **Challenges:**
  - Scalability; computation of $f(x)$ is NP-hard (or more!).
  - Models must be known.

Solvers (Reasoners)

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f} \Longrightarrow \textit{Output } f(x)$$

- **Solvers** derive output $f(x)$ for **given input** x from **model**:
  - **SAT:** x is a formula in CNF, $f(x) = 1$ if x satisfiable, else $f(x) = 0$.
  - **Classical planner:** x is a planning problem P, and $f(x)$ is plan that solves P. 👈
  - **Bayesian net:** x is a query over Bayes Net and $f(x)$ is the answer.
  - **Constraint satisfaction, Markov decision processes, POMDPs, ...**
- ✓ **Generality:** Solvers not tailored to particular examples.
- ✓ **Expressivity:** Some models very expressive; e.g., POMDPs.
- ✖ **Challenges:**
  - Scalability; computation of $f(x)$ is NP-hard (or more!).
  - Models must be known.

✲ **Learners are solvers too:** $\mathrm{argmin}_w \sum_{x \in D} L(x, f_w(x))$ (Differentiable programming)

Learners

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f_\theta} \Longrightarrow \textit{Output } f_\theta(x)$$

- In **deep learning (DL)** and **deep reinforcement learning (DRL)**, training results (the "model") in function $f_\theta(\cdot)$.

Learners

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f_\theta} \Longrightarrow \textit{Output } f_\theta(x)$$

- In **deep learning (DL)** and **deep reinforcement learning (DRL)**, training results (the "model") in function $f_\theta(\cdot)$.
- $f_\theta(\cdot)$ given by structure of **neural network** and adjustable parameters θ.
  - In DL, **input** x may be an image and **output** $f_\theta(x)$ a classification label.
  - In DRL, **input** x may be state of game, and **output** $f_\theta(x)$, value of state.
- Parameters θ learned by **minimizing error function** by stochastic gradient descent.
  - In DL, error depends on inputs and target outputs in training set.
  - In DRL, error depends on value of states and successor states.

Learners

Input $x \Longrightarrow$ $\boxed{\textsc{Function}\ f_\theta}$ $\Longrightarrow$ *Output* $f_\theta(x)$

- In **deep learning (DL)** and **deep reinforcement learning (DRL)**, training results (the "model") in function $f_\theta(\cdot)$.
- $f_\theta(\cdot)$ given by structure of **neural network** and adjustable parameters θ.
  - In DL, **input** x may be an image and **output** $f_\theta(x)$ a classification label.
  - In DRL, **input** x may be state of game, and **output** $f_\theta(x)$, value of state.
- Parameters θ learned by **minimizing error function** by stochastic gradient descent.
  - In DL, error depends on inputs and target outputs in training set.
  - In DRL, error depends on value of states and successor states.

✓ A true revolution in AI still **unfolding**...

✖ **Limitations:** transparency, amounts of data, generalization, understanding

Learners *vs* Solvers

$$\textit{Input } x \Longrightarrow \boxed{\textsc{Function } f} \Longrightarrow \textit{Output } f(x)$$

- **Learners** require **experience over related problems** x but then fast!
  - They compute function f from training, then apply it.
- **Solvers** deal with **new problems** x but need **models**, and need to **"think"** hard.
  - They compute $f(x)$ for each input x from scratch; out of the box.

Learners and Solvers: System 1 and System 2?

Dual process accounts of the human mind assume two processes
(D. Kahneman: Thinking, Fast and Slow, 2011; K. Stanovich: The Robot's Rebellion, 2005)

(Intuitive Mind) *(Analytical Mind)*

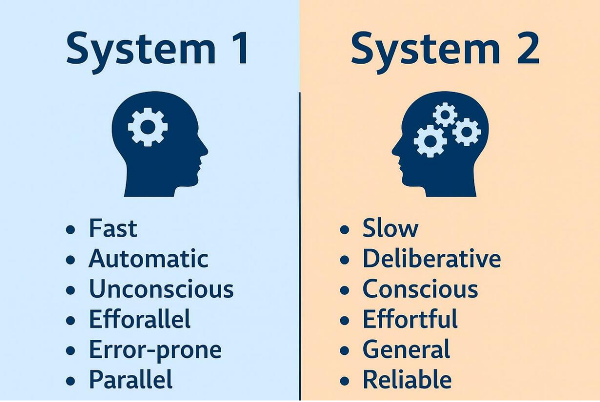

Learners? *Solvers?*

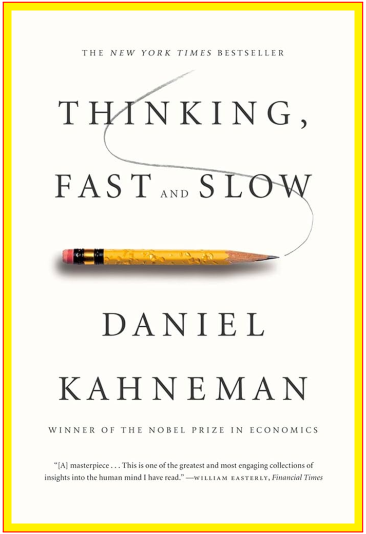

SAT and CSPs

- SAT: determine if there is a truth assignment that satisfies a set of clauses:

$$(x \vee \neg y \vee \neg z) \wedge (\neg x \vee y) \wedge (y \vee z) \wedge \ldots$$

- Problem is NP-Complete, which in practice means worst-case behavior of SAT algorithms is **exponential** in number of variables ($2^{100} = 10^{30}$).
- Yet current SAT solvers manage to solve problems with thousands of variables and clauses, and used widely (circuit design, verification, planning, etc).
- Constraint Satisfaction Problems (CSPs) generalize SAT by accommodating non-boolean variables as well, and constraints that are not clauses.
- Key is **efficient (poly-time) inference** in every node of search tree: **unit resolution, conflict-based learning**, ...

Classical Planning Model

- Planning is the **model-based approach** to autonomous behavior.
- A system can be in one of many **states**.
- States assign **values** to a set of **variables**.
- **Actions** change the values of certain variables.
- **Basic task:** find **action sequence** to drive **initial** state into **goal** state:

$$\textit{Model World } x \Longrightarrow \boxed{\textsc{Planner}\ f} \Longrightarrow \textit{Action Sequence } f(x)$$

- **Complexity**: NP-hard+; i.e., exponential in number of vars in **worst case**.
- Planner is generic: should work on any domain no matter what variables are about.

Why do we need such AI Planning?

Settings where greater autonomy required:

- **Space Exploration**: (RAX) first artificial intelligence control system to control a spacecraft without human supervision (1998)
- **Business Process Management**
- **First Person Shooters & Games:** classical planners playing Atari Games
- **Interactive Storytelling**
- **Network Security**
- **Logistics/Transportation/Manufacturing**: Multi-model Transportation, forest fire fighting, PARC printer
- **Wherehouse Automation**: Multi-Agent Path Finding, Post China, Amazon
- Automation of Industrial Operations (Schlumberger)
- Self Driving Cars ...

✲ Find out more at ICAPS in Action (right panel)


Poster by
AAAI
Association for the Advancement of Artificial Intelligence
AI Magazine
Poster development supported in part by
Microsoft Research
YAHOO! RESEARCH
Optimizing Paths & Flows
AI in Art
AI in Music
AI & Creative Expression
Planning
Plan Recognition
Autonomous Vehicles & Safety
Robot Guides & Assistants
Humanoid Robots
Augmenting Cognition
Intelligent Tutoring
Ubiquitous Computing
Social Computing
AI HISTORY TOUR
Gesture Recognition
Multimodal Interfaces
Mixed-Initiative Collaboration
Autonomous Space Exploration
AI and Preferences, Media & Entertainment
Home Robotics
Descartes
Aristotle
Leibniz
Whitehead
Lovelace
Russell
Turing
Robots For Education
Robotic Surgery
Diagnosis
Drug Design
Scientific Discovery
AI LAB
Assistive Technology
MEDICAL CENTER
Search & Retrieval
Security & Privacy
User Modeling
Vehicle Navigation
Handwriting & Sketch Recognition
BOOKS & COFFEE
Recommender Systems & Question Answering
Communications Triage
Machine Translation
Ecocomputing
See the AI timeline and more at www.aaai.org/AILandscape
The AI Landscape
David Leake, Indiana University, Poster Development Committee Chair
Poster Design: Giacomo Marchesi, www.GiacomoMarchesi.com

Summary: AI and Automated Problem Solving

- A research agenda emerged in last 20 years: **solvers** for a range of **intractable models**.
- **Solvers** unlike other programs are **general** as they do not target individual problems but families of problems (**models**).

Summary: AI and Automated Problem Solving

- A research agenda emerged in last 20 years: **solvers** for a range of **intractable models**.
- **Solvers** unlike other programs are **general** as they do not target individual problems but families of problems (**models**).
- The challenge is **computational**: how to scale up.
- Sheer **size of problem** shouldn't be impediment to meaningful solution.
- **Structure** of given problem must recognized and **exploited**.

Summary: AI and Automated Problem Solving

- A research agenda emerged in last 20 years: **solvers** for a range of **intractable models**.
- **Solvers** unlike other programs are **general** as they do not target individual problems but families of problems (**models**).
- The challenge is **computational**: how to scale up.
- Sheer **size of problem** shouldn't be impediment to meaningful solution.
- **Structure** of given problem must recognized and **exploited**.
- Lots of room for **ideas** but methodology **empirical**.
- Consistent **progress**:
  - effective inference methods (derivation of h, conflict-learning)
  - islands of tractability (treewidth methods and relaxations)
  - transformations (compiling away incomplete info, extended goals, ...)

Course Aim

- **Not a full-fledge course on AI Planning**; too much for us...
  - *Full semester courses (12+ weeks) and still not complete*
- Focus is on **coherent research thread** that covers a lot of ground:
  - *Crisp and simple ideas and formulations for **stating**, **understanding**, and **addressing** key problems.*
- Many open problems; many opportunities for research

System 1 and 2 Intelligence: A Key Challenge in AI

- General **two-way integration** of System 1 and System 2 inference in AI systems:
  - **Learners** and **solvers** should inform, complement, and enhance each other.
- **Yoshua Bengio**'s challenges reflected in title of his IJCAI 2021 talk:
  - *System 2 Deep Learning: Higher-level cognition, agency, out-of-distribution generalization and causality*.
- **Yann LeCun**'s three challenges, AAAI 2020:
  - AI must learn to represent the world.
  - AI must think and plan in ways compatible with gradient-based learning.
  - AI must learn hierarchical representation of action plans.

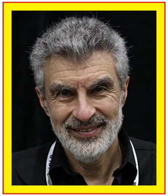

Research Challenge: Minigrid

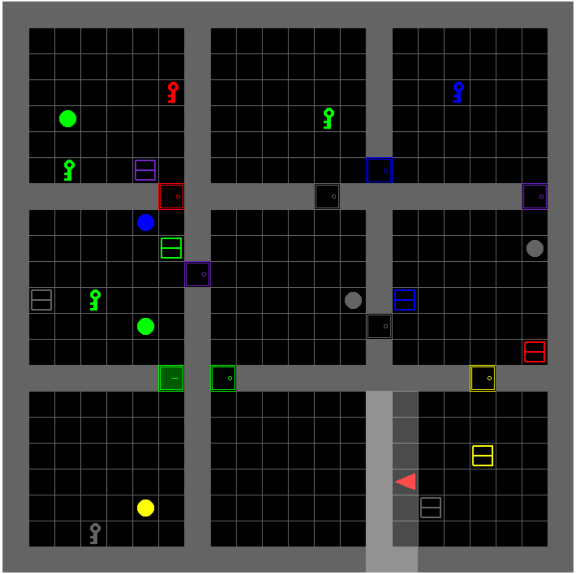

- **Task:** *Pick up grey box behind you, then go to grey key and open door*
- Agent is red triangle at bottom right. Light-grey is field of view.
- Learn **controller** that accepts **goals** and **observations**, and outputs **actions**.
- **How** to get such a controller? Action model and goal language **not known**, but can do **trial-and-error**.

Methodology: Bottom-Up vs. Top-Down Learning

- **Deep (reinforcement) learning methods** struggle in these problems,but manage to generate meaningful behavior after millions of trials (despite **so little prior knowledge**).

- Yet **methodology** largely **ad-hoc:** from intuitions to **architectures** and **experiments** using baselines; performance improvements but **no crisp understanding**.

Methodology: Bottom-Up vs. Top-Down Learning

- **Deep (reinforcement) learning methods** struggle in these problems,but manage to generate meaningful behavior after millions of trials (despite **so little prior knowledge**).

- Yet **methodology** largely **ad-hoc:** from intuitions to **architectures** and **experiments** using baselines; performance improvements but **no crisp understanding**.

Alternative: Top-Down

Alternative: **complementary, top-down** approach asks crisp questions like:

- What are the **domain-independent languages** for representing **dynamics**?
- What the **languages** for representing general reactive **policies**, **subgoals**?
- What are good **solvers** for those representations?
- How can **representations** over such languages be **learned**?

AI and Social Impact

- **System 2** not only necessary for AI systems; essential for people and **societies**.
- AI far from human-level intelligence, yet it can be used for **good** or **ill**.
- **Ethical committees** and **AI principles** good but not sufficient.
- **Markets and politics** play our **System 1**, focused on the **bottom line**. 😵

AI and Social Impact

- **System 2** not only necessary for AI systems; essential for people and **societies**.
- AI far from human-level intelligence, yet it can be used for **good** or **ill**.
- **Ethical committees** and **AI principles** good but not sufficient.
- **Markets and politics** play our **System 1**, focused on the **bottom line**. 😵

- If we want **good AI**, we need a **good and decent society**, that make use of our **System 2** and cares about truth, reason, knowledge, and the common good.
- Take courses on *Social and technological change*... 🙌

Part 1: Introduction, Motivation, and AI Search

Part 1: Introduction, Motivation, and AI Search

Ready to go?

AI's favourite trick