Neuroevolution: Harnessing Creativity in AI Agent Design

Risto Miikkulainen
The University of Texas at Austin and Cognizant AI Lab

With Sebastian Risi, Yujin Tang, and David Ha

cognizant

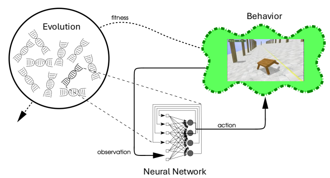

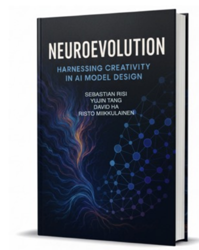

Neuroevolution: Harnessing Creativity in AI Agent Design

An MIT Press Book by Sebastian Risi, Yujin Tang, David Ha, and Risto Miikkulainen

Foreword by Melanie Mitchell

The online version of the book is now freely available in an open-access HTML format. The print edition will be released later in 2026.

Access the Free Online Edition

Overview | Authors | Table of Contents | Demos | Exercises | Course and Tutorials | For Instructors | Discussion | Cite

Neuroevolution, or optimization of neural networks through evolutionary computation, has been a growing subarea of machine learning since the 1990s. Its primary focus is on evolving neural networks for intelligent agents when the training targets are not known, and good performance requires many decisions over time, such as robotic control, game playing, and decision-making. More recently it has also been extended to optimizing deep-learning architectures, understanding how biological intelligence evolved, and optimizing neural networks for hardware implementation. This book introduces students to the basics of neuroevolution, progresses to several advanced topics that make neuroevolution more effective and more general, reviews example application areas, and proposes further research questions. Hands-on experience is provided through a Python-based software platform with animations, interactive demos, exercises, and project environments.

https://neuroevolutionbook.com

Outline

1. Motivation for Neuroevolution
2. Basics
    1. Fundamentals of Evolution
    2. Fundamentals of Neuroevolution
3. Advances
    1. Taking Advantage of Indirect Encodings
    2. Taking Advantage of Diversity
4. Evolving Intelligent Agents
    1. Control
    2. Strategy
    3. Decision-making
    4. Collective behavior
5. Synergies with other ML
    1. Deep Learning
    2. Reinforcement Learning
    3. LLMs
6. Insights into Biology
7. Conclusion

Hands-on exercise (off-line)

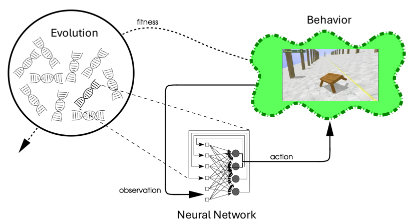

Motivation: From Imitation to Creativity

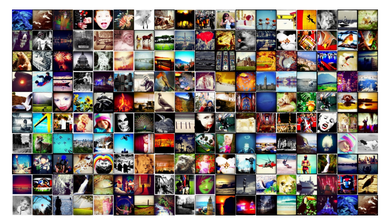

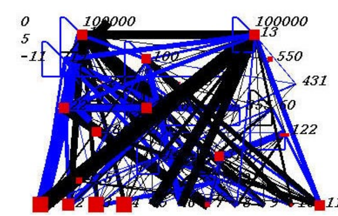

- Much of AI so far focuses on imitation
    - I.e. gradient descent on labeled datasets
    - Powerful in prediction: object recognition, diagnosis, forecasting, etc.
- Agentic AI focuses on behavior
    - Gradients not available
    - Needs to be discovered
- How can we create novel behaviors?

Reinforcement Learning is One Approach

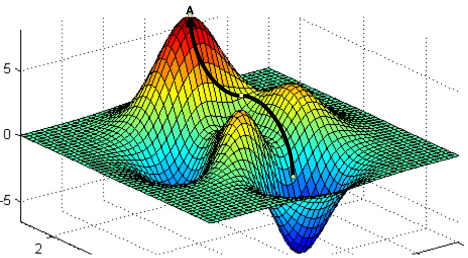

- Approximating gradient descent
  - Explore around the current solution
  - Improve it gradually
  - Can climb the nearest hill well

...but Creativity in RL is Limited

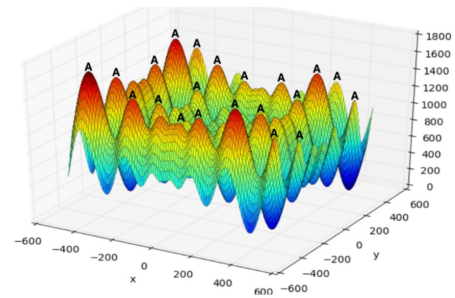

- Space is too large
  - Multiple starts won't help
- Space is too high-dimensional
  - Little improvement from one step
- Space is deceptive
  - Can only find the nearest hill

Solution: Population-based Search

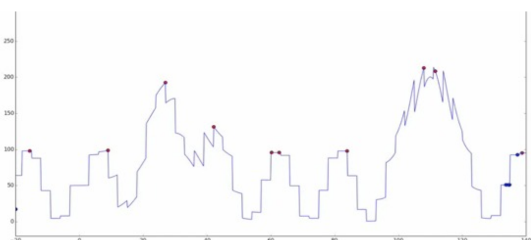

A.k.a Evolutionary Computation
- Many individuals spread out, sharing information
- Not limited to differentiable domains: configurations, choices ok

Not limited to incremental improvement
- Large jumps possible, can be more creative

Scaling up through Evolution

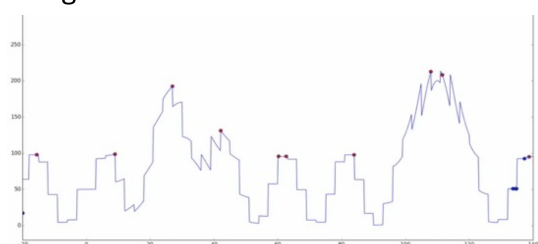

Works in large scales
- Structured search works in large spaces (e.g. 2^2^70; Hodjat & Shahrzad 2016)
- Multiple variables optimized at once (e.g. up to 1B; Deb et al. 2017)
- Multiple objectives and novelty get around deception (Shahrzad and Hodjat 2020)

Neuroevolution uses population-based search to optimize neural netwoks
- Weights, topologies, designs

Evolution Basics: Encoding, Evaluation, and Selection

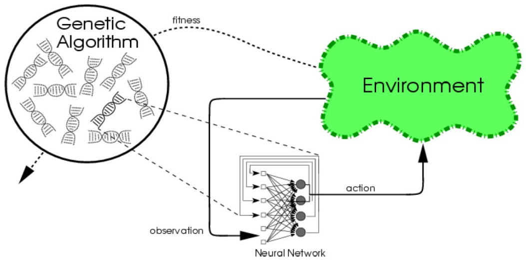

- A population of encodings (e.g. lists or trees)
- Decoded into individuals that are evaluated in the domain
- Good individuals retained, bad thrown away

Creating Variation

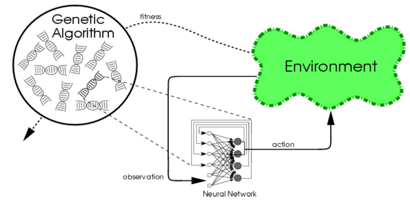

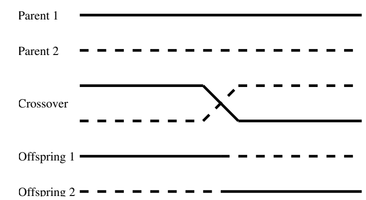

- New individuals generated from the parent encodings
  - Crossover: combine building blocks from two parents
  - Mutation: create new building blocks

Population-based Search

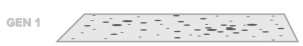

Population-based Search

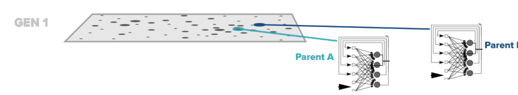

Population-based Search

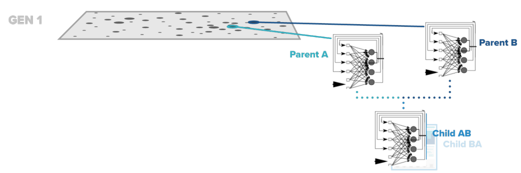

Population-based Search

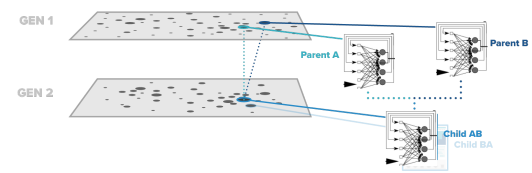

Population-based Search

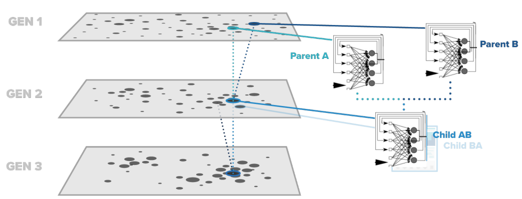

Population-based Search

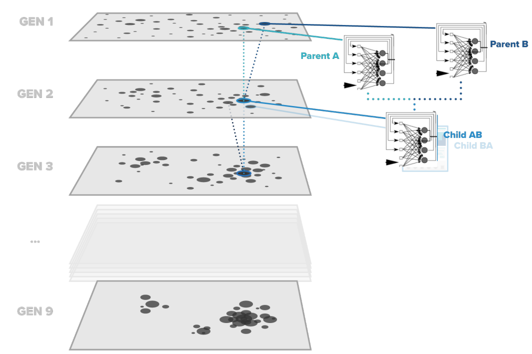

Example: Learning to Walk

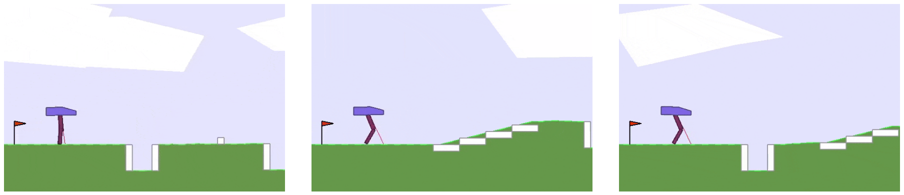

Demo

Outline

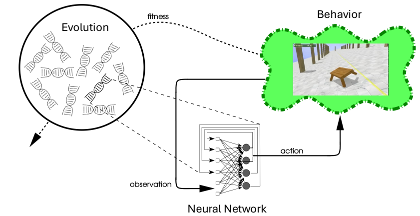

Searching for Novelty

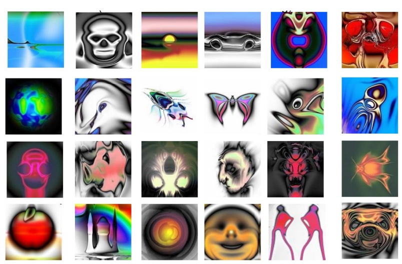

- Motivated by humans as fitness functions
- E.g. picbreeder.com, endlessforms.com (Secretan et al. 2011; Clune et all 2011)
    - CPPNs evolved; Human users select parents
- No specific goal
    - Interesting solutions preferred
    - Similar to biological evolution?

Novelty Search

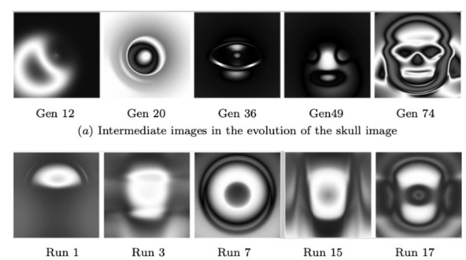

(a) Intermediate images in the evolution of the skull image

(b) Attempts to evolve the skull image directly

- Evolutionary algorithms maximize a performance objective
    - But sometimes hard to achieve it step-by-step
- Novelty search rewards candidates that are simply different
    - Stepping stones for constructing complexity

Novelty Search Demo

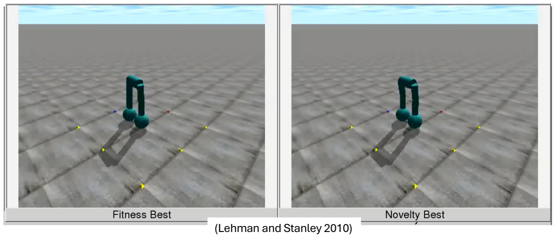

(Lehman and Stanley 2010)

- Fitness-based evolution is rigid
  - Requires gradual progress
- Novelty-based evolution is more innovative, natural
  - Allows building on stepping stones
- How to guide novelty search towards useful solutions?
  - Quality Diversity methods
- DEMO

Outline

1. Motivation for Neuroevolution
2. Basics
   1. Fundamentals of Evolution
   2. Fundamentals of Neuroevolution
3. Advances
   1. Taking Advantage of Indirect Encodings
   2. Taking Advantage of Diversity
4. Evolving Intelligent Agents
   1. Control
   2. Strategy
   3. Decision-making
   4. Collective behavior
5. Synergies with other ML
   1. Deep Learning
   2. Reinforcement Learning
   3. LLMs
6. Insights into Biology
7. Conclusion

Hands-on exercise (off-line)

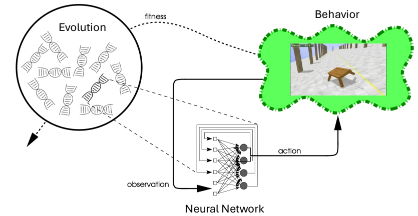

Optimizing Decision-Making



Organizations have lots of data
- Can build predictive models of patients, customers, students…

Such models do not specify how to make decisions
- Optimal decisions not known

Need to search for decision strategies
- But testing strategy candidates in the real world is costly

Surrogate Optimization Approach

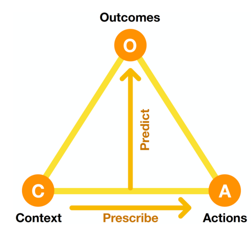

(Francon et al. 2020)

Use a predictive model as a surrogate for the world
Train model with historical data: Context+Actions → Outcomes
- Phenomenological model (based on data)
- Not a simulation from first principles

Search for a good decision strategy (i.e. policy): Context → Actions
- Use the model to evaluate strategies
- Evolve a neural network to represent the strategy

Optimizing Interventions in COVID-19

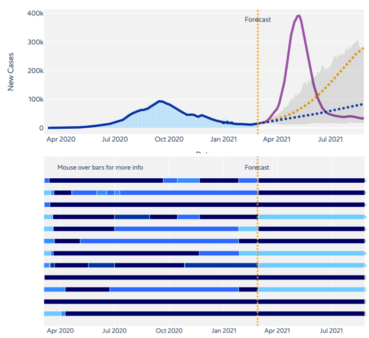

(Miikkulainen et al. IEEE TEC 2021)

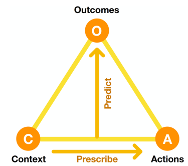

Based on two models:

1. A predictive model
   - Given a history of cases and nonpharma interventions (NPIs)
   - Predict number of cases daily
2. A prescriptive model
   - Given a history of cases and NPIs
   - Prescribe NPIs daily to minimize cases and restrictions

Not just what will happen, but what we should do about it

Leveraging Human Expertise with AI

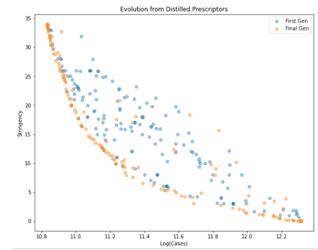

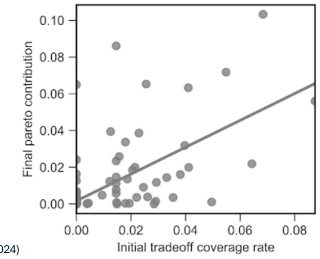

(Meyerson et al. 2024)

XPRIZE competition resulted in 169 human expert strategies
- Many useful, diverse ideas

Can be used as an initial population for search
- Improve further; better than search from scratch

Can realize latent potential of ideas
- How much DNA in the final Pareto front

Technology to bring the community effort together
- There is power in diversity

Outline

1. Motivation for Neuroevolution
2. Basics
   1. Fundamentals of Evolution
   2. Fundamentals of Neuroevolution
3. Advances
   1. Taking Advantage of Indirect Encodings
   2. Taking Advantage of Diversity
4. Evolving Intelligent Agents
   1. Control
   2. Strategy
   3. Decision-making
   4. Collective behavior
5. Synergies with other ML
   1. Deep Learning
   2. Reinforcement Learning
   3. LLMs
6. Insights into Biology
7. Conclusion

Hands-on exercise (off-line)

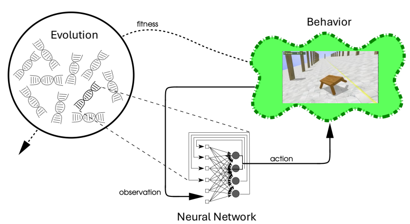

Neuroevolution Synergies with Deep Learning

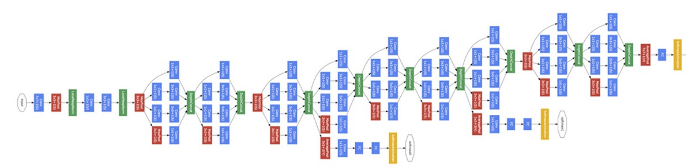

(A) Fundamental: Neural Architecture Search
- Optimizing structure and hyperparameters
- Takes advantage of exploration in EC

(B) Extended: Data and training
- Loss functions, activation functions, data augmentation, initialization, learning algorithm
- Takes advantage of flexibility of EC

Example: Evolving Age-Estimation Networks

Parameter	Possible Values	Type	Class
Algorithm	[adam, rmsprop]	Enum	Opt
Initial Learning Rate (LR)	[1e-5, 1e-3]	Float	Opt
Momentum	[0.7, 0.99]	Float	Opt
(Weight Decay) / LR [26]	[1e-7, 1e-3]	Float	Opt
Patience (Epochs)	[1, 20]	Int	Opt
SWA Epochs [21]	[1, 20]	Int	Opt
Rotation Range (Degrees)	[1, 60]	Int	Aug
Width Shift Range	[0.01, 0.3]	Float	Aug
Height Shift Range	[0.01, 0.3]	Float	Aug
Shear Range	[0.01, 0.3]	Float	Aug
Zoom Range	[0.01, 0.3]	Float	Aug
Horizontal Flip	{True, False}	Bool	Aug
Vertical Flip	{True, False}	Bool	Aug
Cutout Probability [7]	[0.01, 0.999]	Float	Aug
Cutout Max Proportion [7]	[0.05, 0.5]	Float	Aug
Pretrained Base Model	Keras App. [5]	Enum	Arch
Base Model Output Blocks	{B0, B1, B2, B3}	Subset	Arch
Loss function λ in Eq. 5	[0, 1]	Float	Arch

(Miikkulainen et al. 2021)

Estimate age from a facial image

Evolving multiple design aspects
- Learning, data augmentation hyperparameters
- Seeded architecture search
- Loss-function optimization: Combination of MAE and CE

Also
- Population-based training
- Ensembling of evolved solutions

Age-Estimation Results

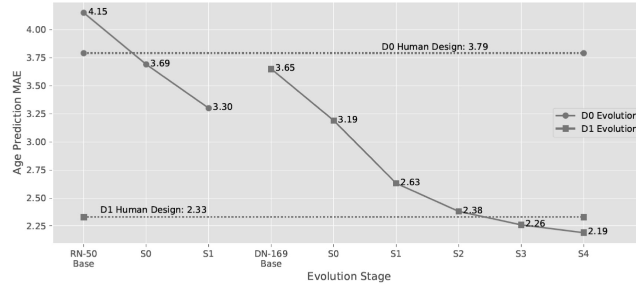

- D0 stages: ResNet-50, DenseNet-121
- D1 stages: DenseNet-169, DenseNet-201, more epochs. EfficientNet-B6, ensembling
- Human optimization of ResNet-50 (D0), EfficientNet-B6 (D1)

Evolution improves significantly over SotA image models
- Fit the design to the task

Optimizes better than humans can
- Many more parameters simultaneously

Performance exceeds that of humans: 2.19 vs. 3-4 years

Outline

1. Motivation for Neuroevolution
2. Basics
   1. Fundamentals of Evolution
   2. Fundamentals of Neuroevolution
3. Advances
   1. Taking Advantage of Indirect Encodings
   2. Taking Advantage of Diversity
4. Evolving Intelligent Agents
   1. Control
   2. Strategy
   3. Decision-making
   4. Collective behavior
5. Synergies with other ML
   1. Deep Learning
   2. Reinforcement Learning
   3. LLMs
6. Insights into Biology
7. Conclusion

Hands-on exercise (off-line)

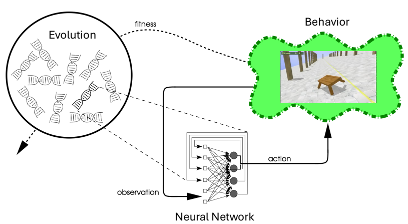

Example: Evolution of Intelligent Coordinated Behavior

Stealing a kill from lions
- Succeeds in an otherwise impossible task (sometimes)
- More sophisticated than other hyena behaviors
- Highly rewarding compared to normal hunting
- Largely genetically determined
- A breakthrough in evolution of intelligence?

A collaboration with Kay Holekamp's lab (MSU)
- Studying hyenas in Masai Mara since 1982

(Rajagopalan et al. 2021)

DEMO

Simulation Setup

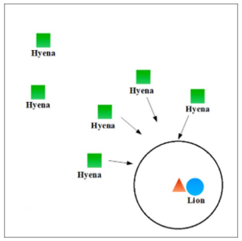

Lion at a kill, with an interaction circle around it
Ten hyenas chosen and placed randomly in the field
If four or more hyenas enter the circle simultaneously, they get the kill
- Otherwise they die

Does mobbing behavior evolve?
- What are the stepping stones for it?

Initial Behaviors

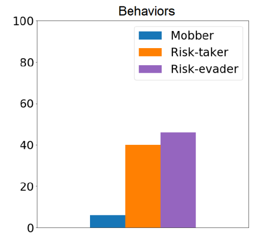

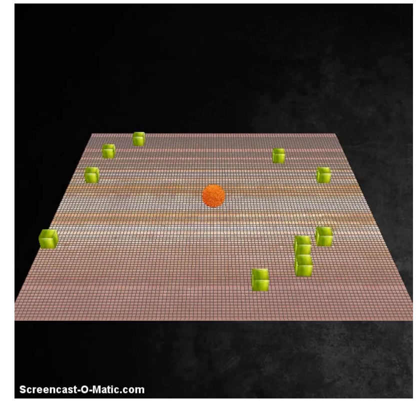

Risk evasion is common
- Never reach the circle; Medium fitness

Risk taking is common
- Charge the circle; Frequent low fitness
- Occasional high fitness by accident

DEMO

Early Behaviors

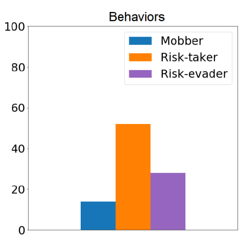

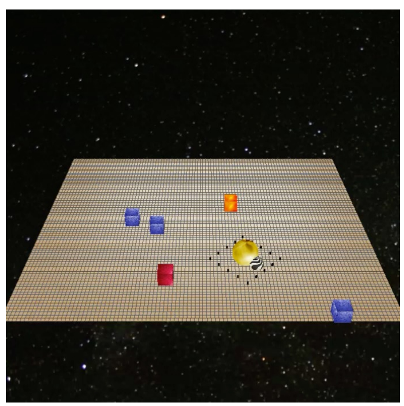

Risk taking grows
- As long as it is successful often enough

Risk evasion also persists
Evasion at the circle starts to emerge
Is mostly detrimental, but an important stepping stone

DEMO

Later Behaviors

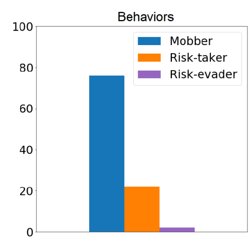

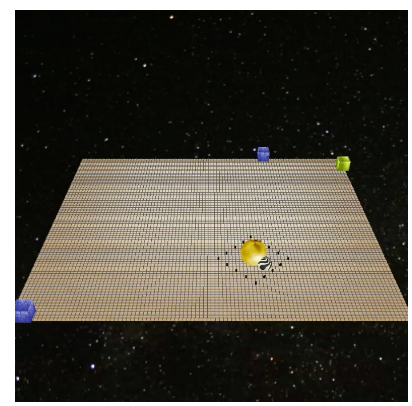

Mobbing emerges
- Not just coincidence of risk takers
- Hyenas wait until there’s enough of them

Risk-evaders evolve into latecomers
Simple risk-taking and risk-evasion still exist

DEMO

Future Challenge: Evolution of Language

Signaling is possible to evolve in ecological simulations
Structured language is much harder
Perhaps language evolved not from signaling, but cognition

- Complex social structure, with modifiable roles
- Language structure can reuse the same conceptual structures

Enough compute, complex simulations to study now?

Conclusion

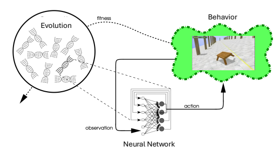

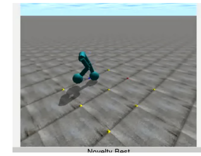

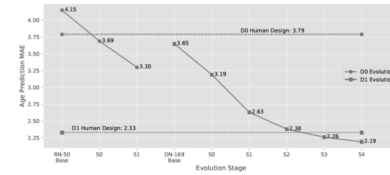

AI is progressing from imitation to creativity; from models to agents

Neuroevolution is a powerful approach to discovering behavior
- Control, strategy, collective behavior, decision-making

Neuroevolution can provide a boost to ML
- Deep learning designs; RL exploration; LLM optimization
- Automatic design of learning machines

Neuroevolution can provide insight into biological evolution
- Evolutionary origins of circuits, behavior, intelligence
- Evolution of language as a current challenge
- A possible path to AGI

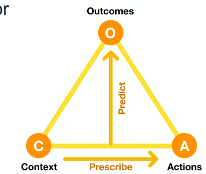

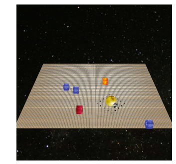

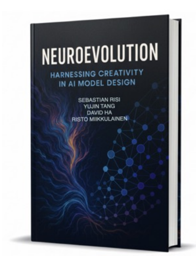

Neuroevolution: Harnessing Creativity in AI Agent Design

An MIT Press Book by Sebastian Risi, Yujin Tang, David Ha, and Risto Miikkulainen

Foreword by Melanie Mitchell

The online version of the book is now freely available in an open-access HTML format. The print edition will be released later in 2026.

Access the Free Online Edition

Overview | Authors | Table of Contents | Demos | Exercises | Course and Tutorials | For Instructors | Discussion | Cite

Neuroevolution, or optimization of neural networks through evolutionary computation, has been a growing subarea of machine learning since the 1990s. Its primary focus is on evolving neural networks for intelligent agents when the training targets are not known, and good performance requires many decisions over time, such as robotic control, game playing, and decision-making. More recently it has also been extended to optimizing deep-learning architectures, understanding how biological intelligence evolved, and optimizing neural networks for hardware implementation. This book introduces students to the basics of neuroevolution, progresses to several advanced topics that make neuroevolution more effective and more general, reviews example application areas, and proposes further research questions. Hands-on experience is provided through a Python-based software platform with animations, interactive demos, exercises, and project environments.

https://neuroevolutionbook.com