

DSC291: Machine Learning with Few Labels

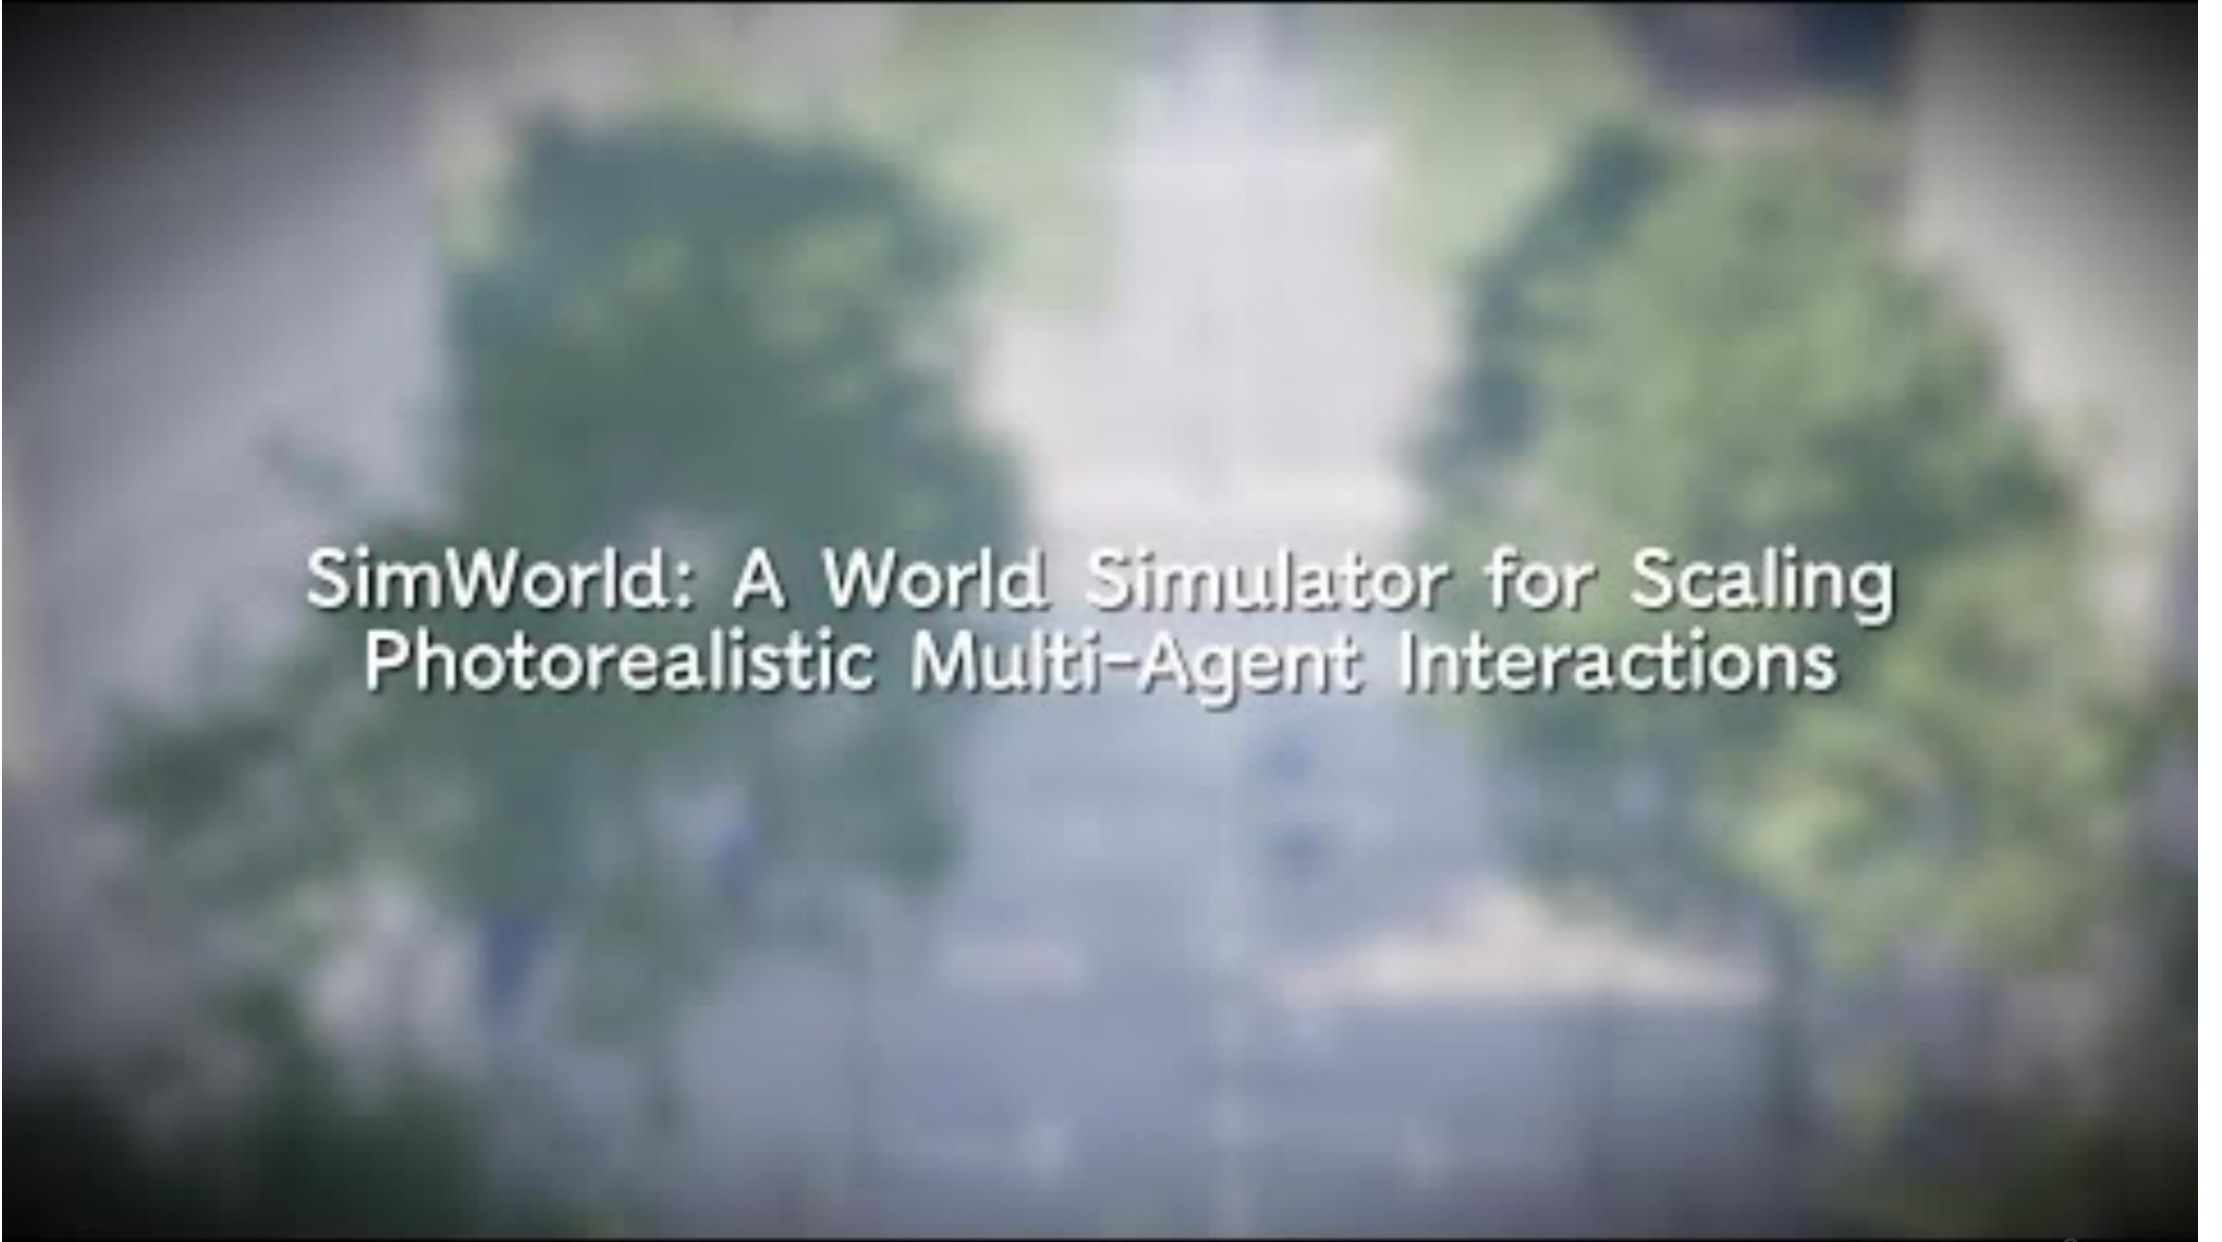
Overview

Zhiting Hu

Lecture 2, April 3rd, 2025

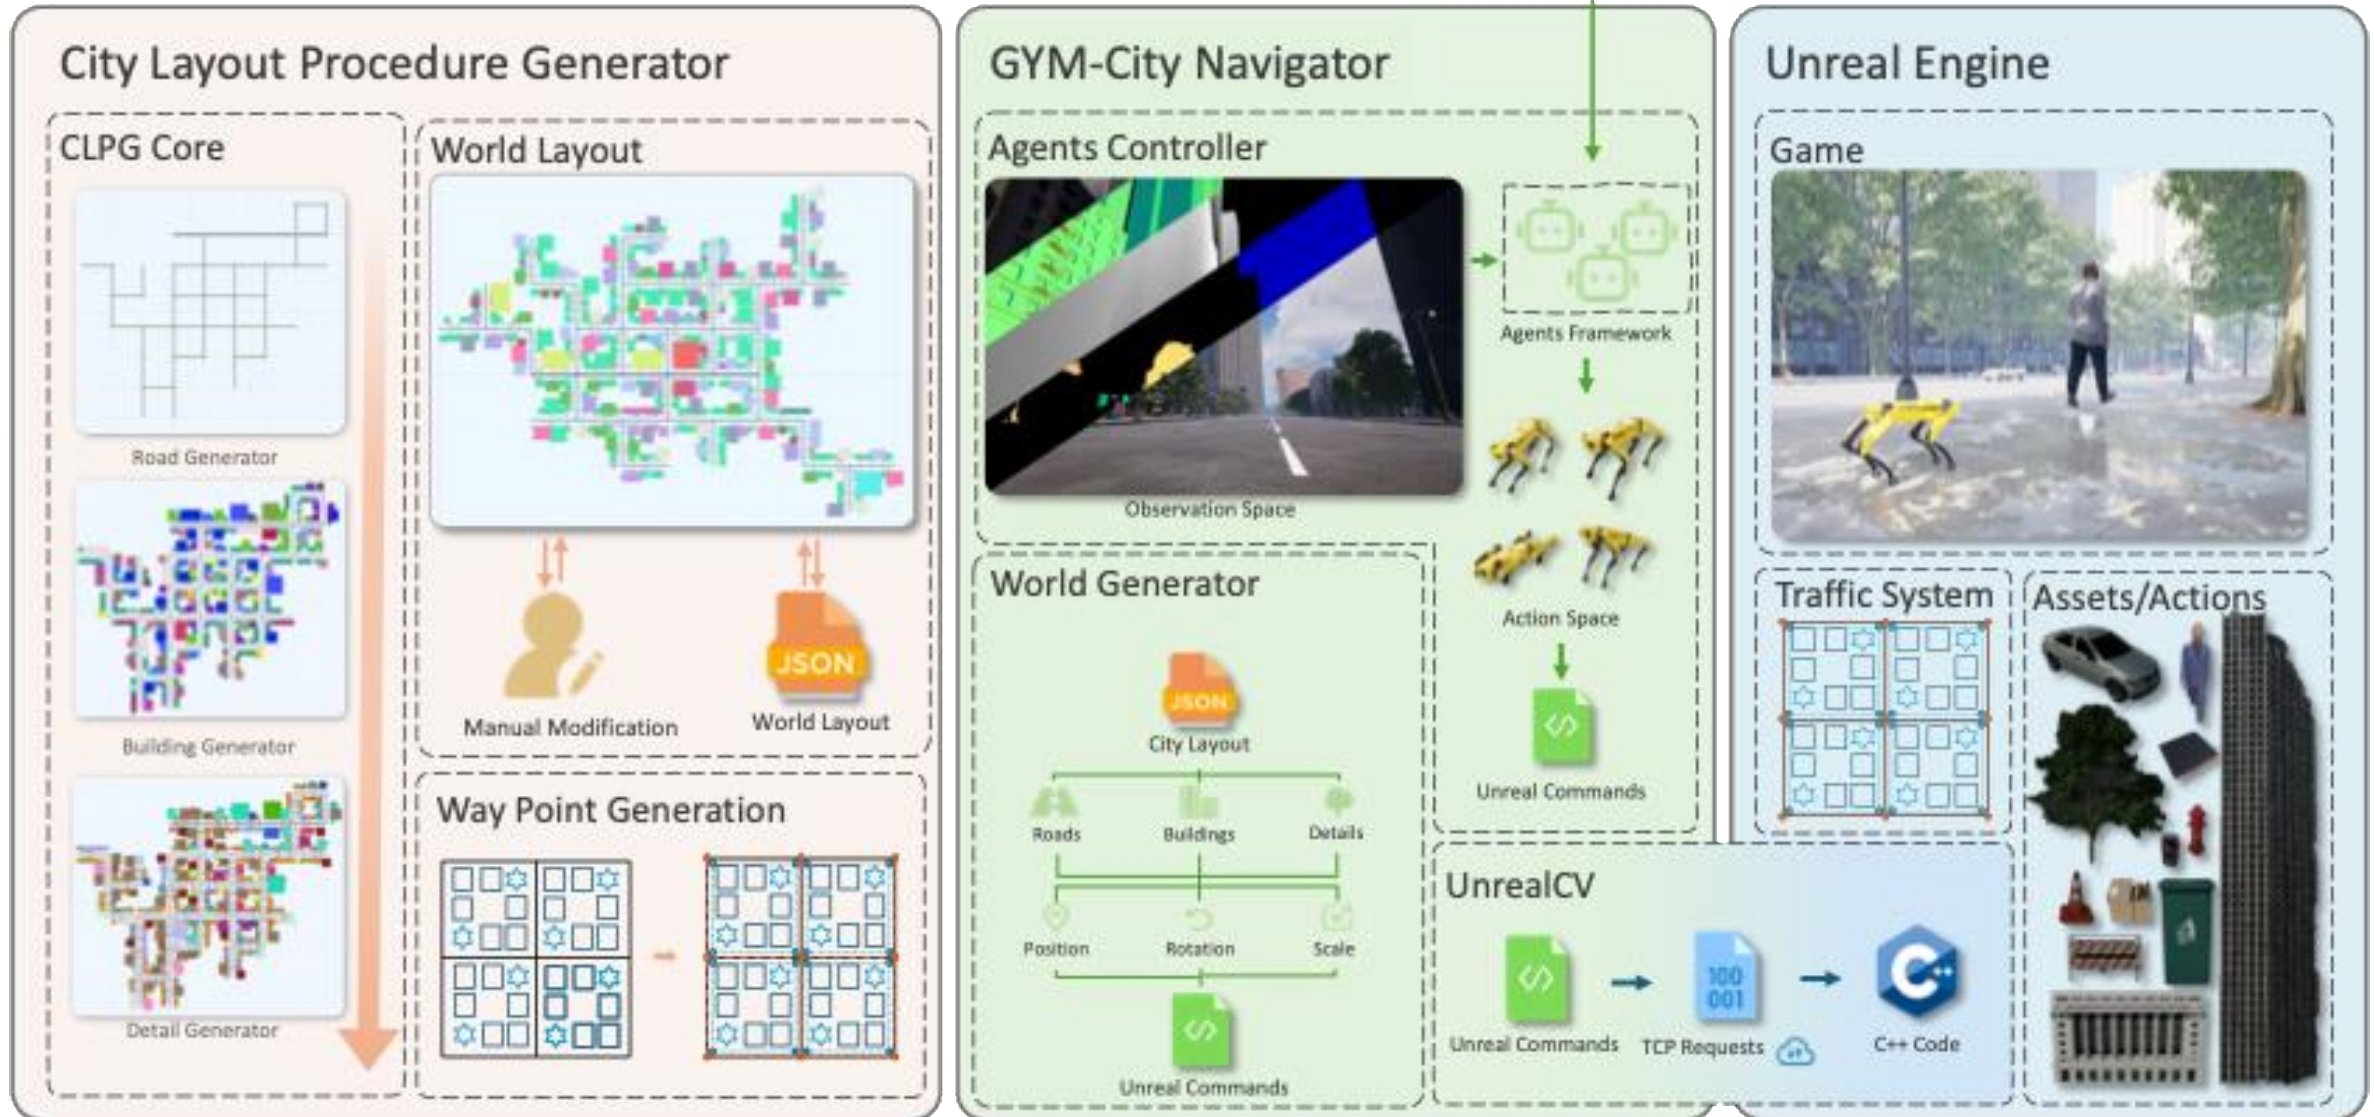
Possible Ideas of Course Project

SimWorld: Open-ended world simulation with tens to millions of agents



SimWorld: A World Simulator for Scaling
Photorealistic Multi-Agent Interactions

SimWorld: Open-ended world simulation with tens to millions of agents



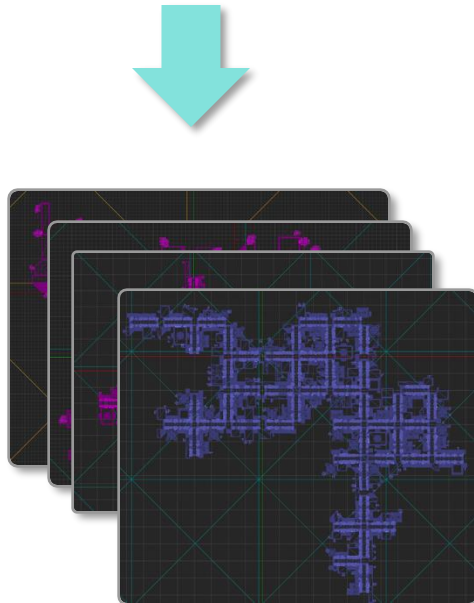
SimWorld: Open-ended world simulation with tens to millions of agents

City Generation

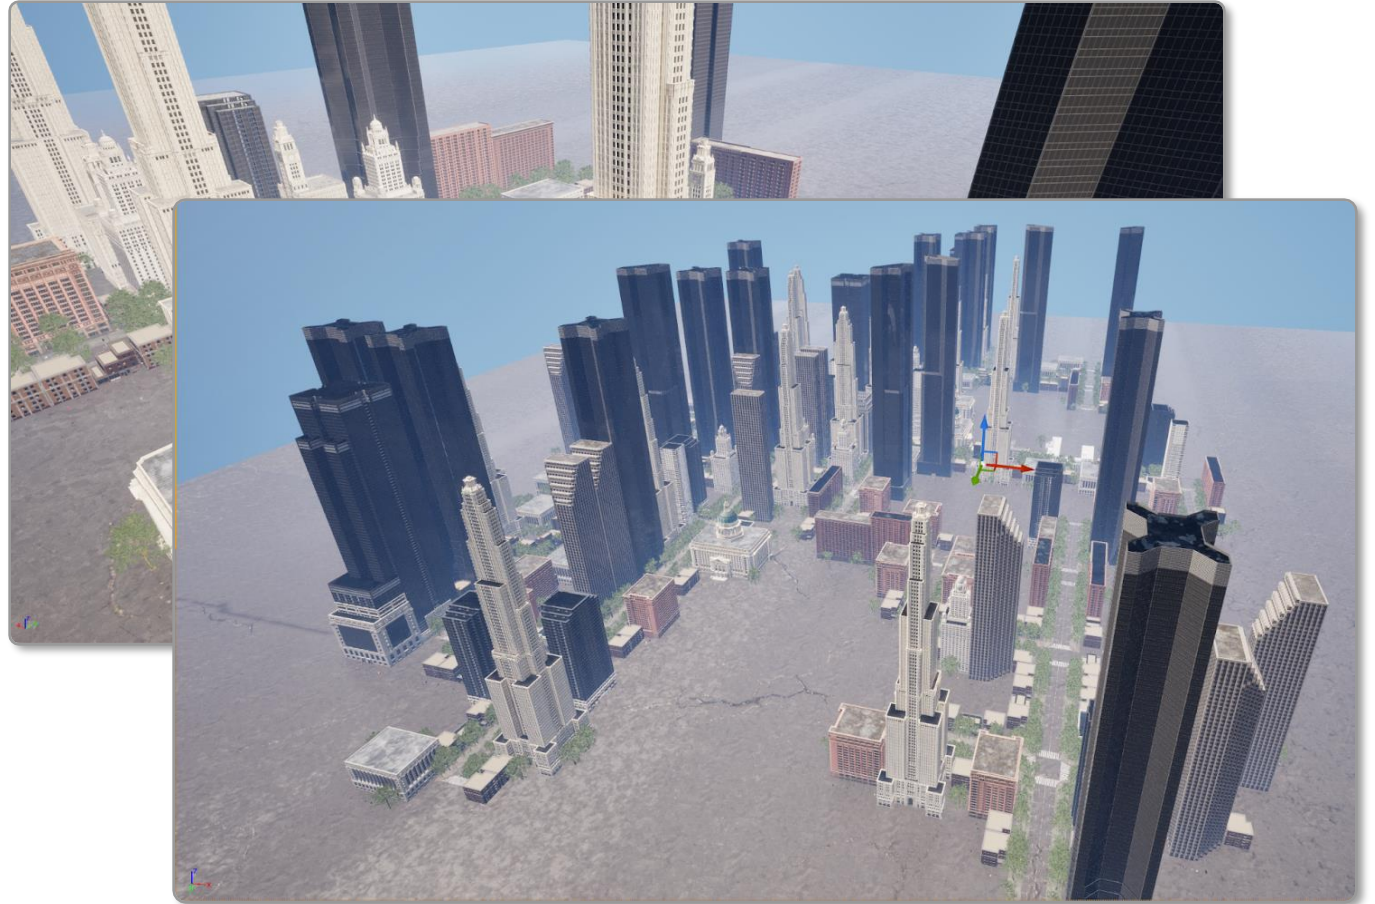
- **Example projects (I): language-guided city/scene generation**

```
import clpg  
  
cfg = clpg.function_call.CityFunctionCall()  
cfg.generate_city()  
cfg.export_city("output_test")
```

4 lines of python code using CLPG



Generated layouts of cities



Rendered cities according to layouts in Unreal Engine

SimWorld: Open-ended world simulation with tens to millions of agents

City Generation

- Example projects (I): language-guided city/scene generation

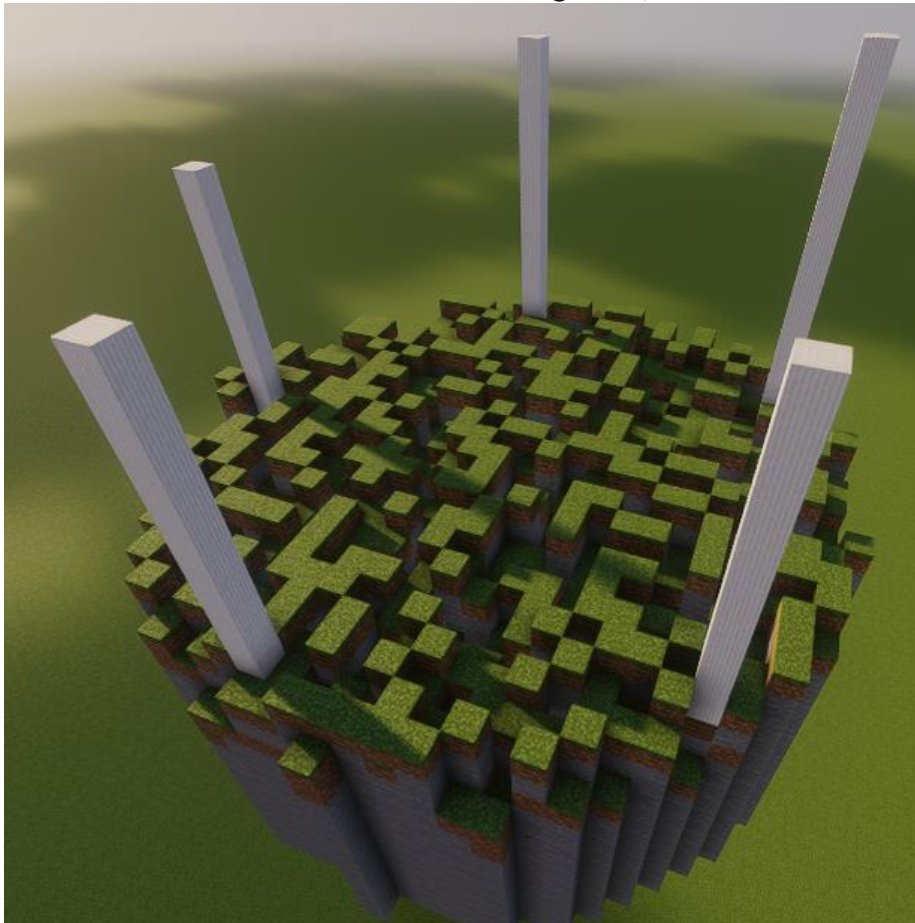


SimWorld: Open-ended world simulation with tens to millions of agents

City Generation

- **Example projects (I): language-guided city/scene generation**

“Build me an amazing, large, organic and epic floating island city right above you with a ton of detail. Make it a goal, iterate.”



Simulator:
Minecraft



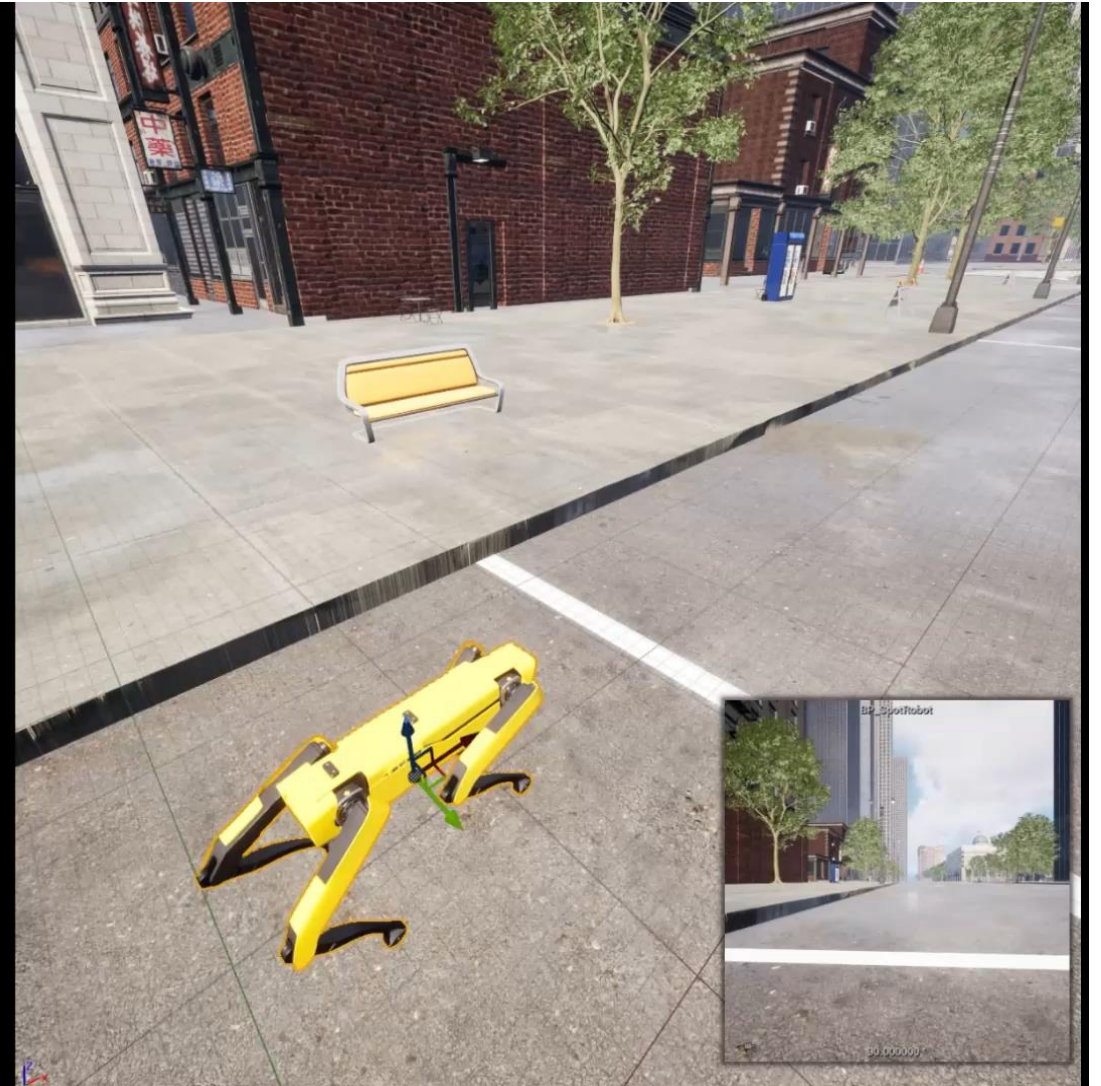
SimWorld: Open-ended world simulation with tens to millions of agents

- Example projects (II): embodied agent for navigation or other tasks

Robot dog
controlled by
GPT-4o

speed of video: 5x
target: blue vending machine
model: GPT-4o (with simple reasoning)
step:
1: rotation(duration=5, angle=15, direction=-1)

planner:
- The blue vending machine is in the field of view.
- The relative direction of the blue vending machine is slightly to the left.
- Suggestion: Slightly rotate left.



SimWorld: Open-ended world simulation with tens to millions of agents

- Example projects (III): multi-agent interactions



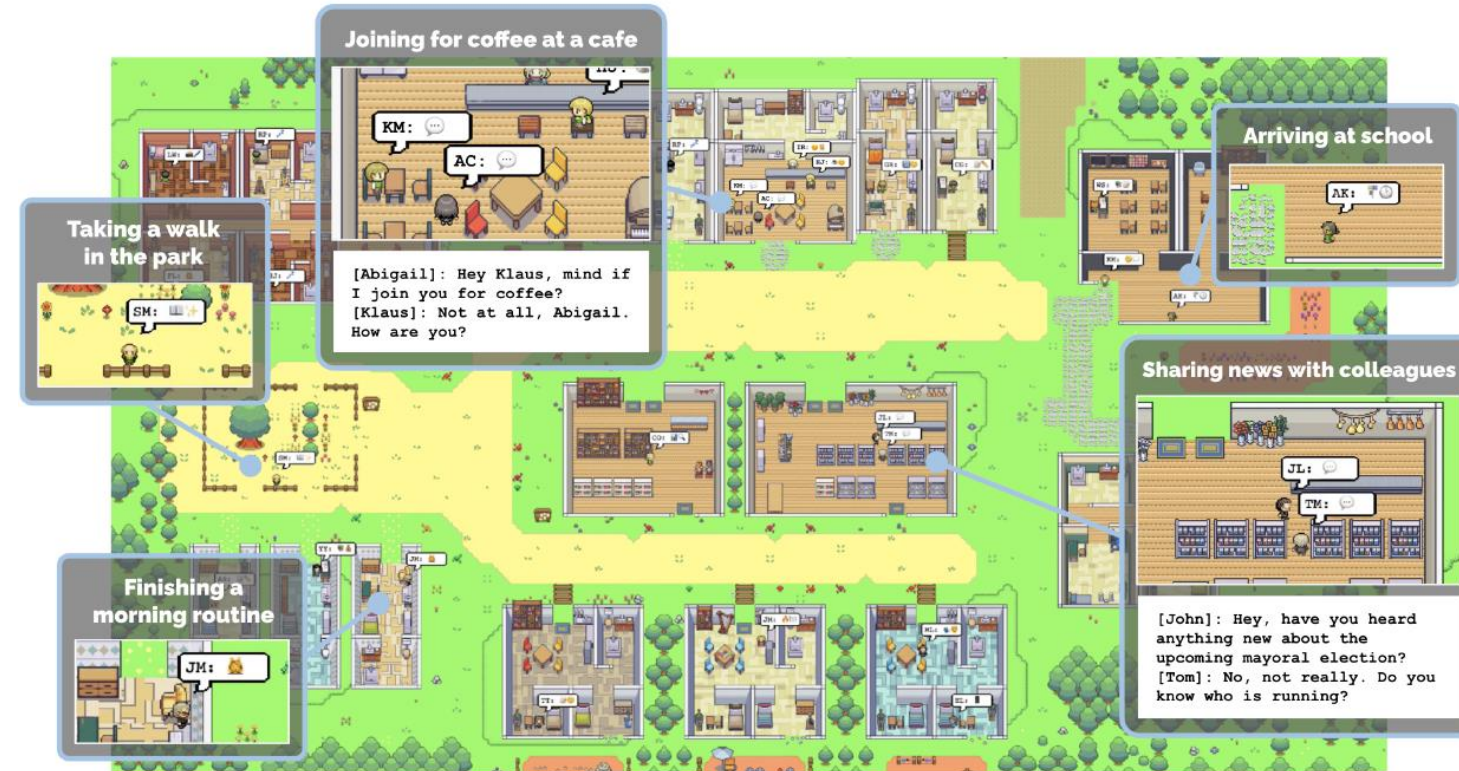
Two agents moving forward concurrently, controlled by our APIs



84 agents moving forward concurrently, no need to manually define them because we supply APIs to generate and control these agents

SimWorld: Open-ended world simulation with tens to millions of agents

- Example projects (III): multi-agent interactions



25 agents, each controlled by individual LLM, converse with each other

- For studying emerging communication behaviors

[Park et al., 2023]

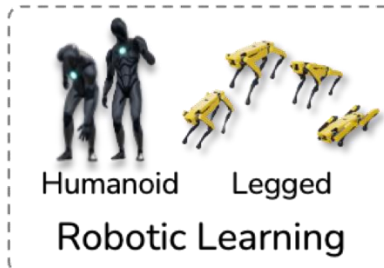
SimWorld: Open-ended world simulation with tens to millions of agents

- Example projects (III): multi-agent interactions



SimWorld: Open-ended world simulation with tens to millions of agents

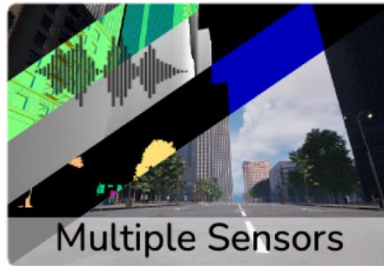
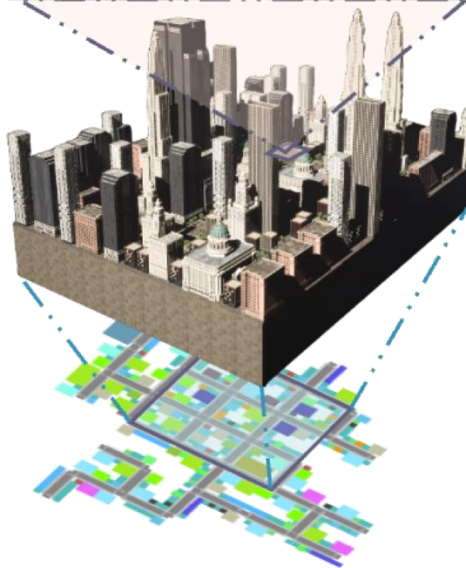
Potential Applications



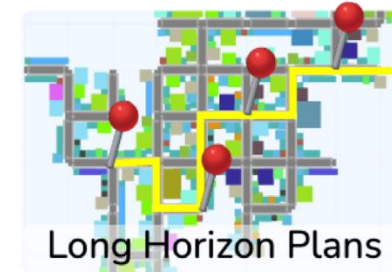
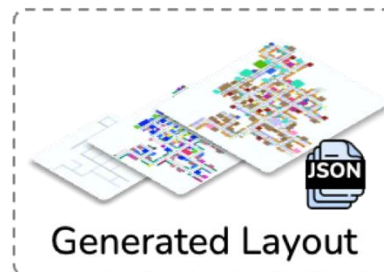
Multi-agent Interactions



Physical Simulation



World Layout & Social Rules



LLM Reasoning and Agent



LLM Reasoners

A library for advanced reasoning with large language models

<https://github.com/matrix-org/llm-reasoners>

- Example projects (IV): web-agent
 - <https://github.com/matrix-org/llm-reasoners/tree/main/examples/browsergym>

LLM Reasoning and Agent



LLM Reasoners

A library for advanced reasoning with large language models

<https://github.com/LLM-Reasoners/LLM-Reasoners>

- Example
 - <https://llm-reasoners.net>

The screenshot displays the LLM Reasoners website, which features a variety of reasoning tasks and a prominent banner for 'Inference-Time Planning for Web Browsing'.

LLM Reasoners
A library for advanced reasoning with large language models

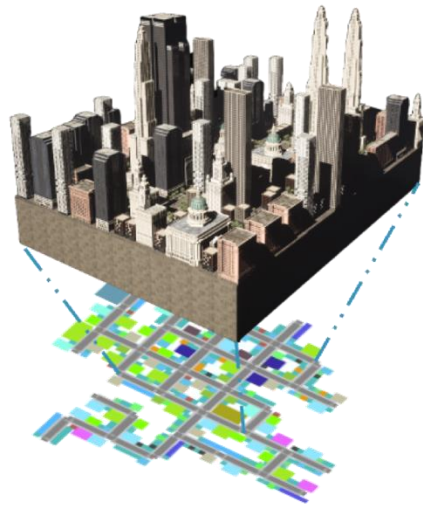
Inference-Time Planning for Web Browsing
NOW SUPPORTED !!

The website includes several sections:

- Search Algorithm:** A diagram showing a search tree with nodes labeled s_0, s_1, s_2, s_3 and a search configuration.
- Model:** A list of supported models: GPT-4 turbo, Claude-3 Opus, Gemini Pro, InternLM-2 7B, and Mixtral 8x7B.
- Reasoning Chains:** A diagram showing a reasoning chain with nodes labeled s_0, s_1, s_2, s_3 and a search configuration.
- World Model:** A diagram showing a world model with nodes labeled s_0, s_1, s_2, s_3 and a search configuration.
- Reasoners:** A section titled 'Reasoners' showing a list of reasoning chains and a diagram of a reasoning chain.
- GSMBk (MCTS):** A section titled 'GSMBk (MCTS)' showing a diagram of a search tree and a list of reasoning chains.
- What mistakes did the student make?:** A section titled 'What mistakes did the student make?' showing a list of reasoning chains and a diagram of a reasoning chain.
- II: Detecting the error:** A section titled 'II: Detecting the error' showing a list of reasoning chains and a diagram of a reasoning chain.
- Dataset:** A section titled 'Dataset' showing a bar chart with four bars labeled 'AQuA*', 'Game24', 'PrOntoQA', and 'Str'.

Possible Ideas of Course Project

- (I): language-guided city/scene generation
- (II): embodied agent for navigation or other tasks
- (III): multi-agent interactions
- (IV): web-agent



LLM Reasoners

A library for advanced reasoning with large language models

Overview

Recap: What is Machine Learning?

- Computational methods that enable machines to learn concepts and improve performance from **experience**.

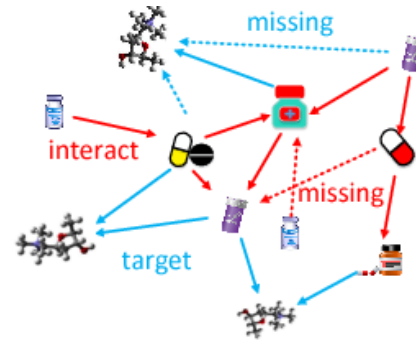
Recap: Experience of all kinds



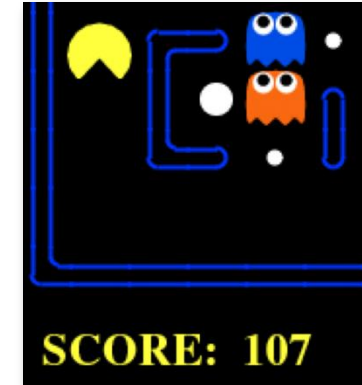
Data examples

Type-2
diabetes is 90%
more common
than type-1

Rules/Constraints



Knowledge graphs



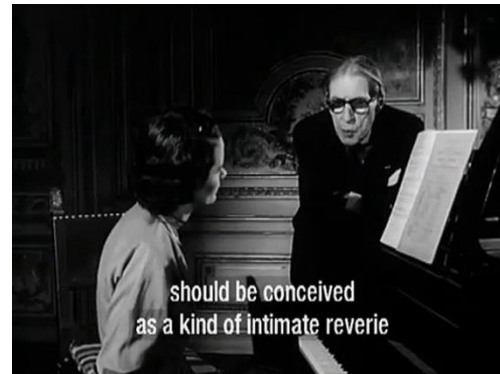
Rewards



Auxiliary agents



Adversaries



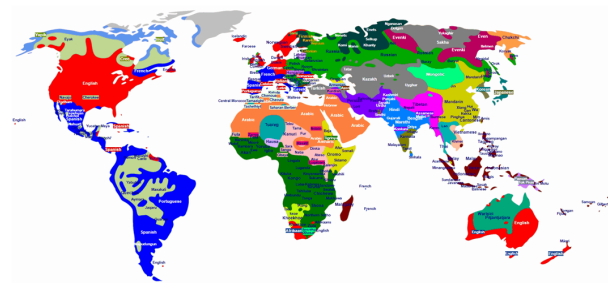
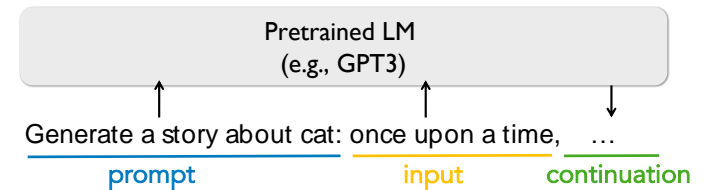
Master classes

...

And all combinations thereof

Recap: Problems with few data (labels)

- Privacy, security issues
- Expensive to collect/annotate
- Difficult / expertise-demanding to annotate
- Specific domain

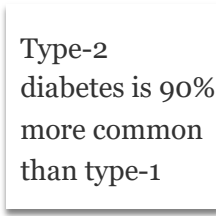


Machine learning solutions given few data (labels)

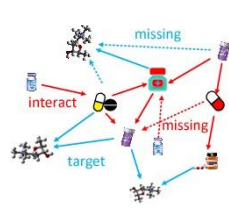
- How can we make more efficient use of **data**?
 - Clean but small-size
 - Noisy
 - Out-of-domain
- Can we incorporate **other types of experience** in learning?



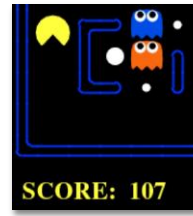
Data examples



Rules/Constraints



Knowledge graphs



Rewards



Auxiliary agents



Adversaries

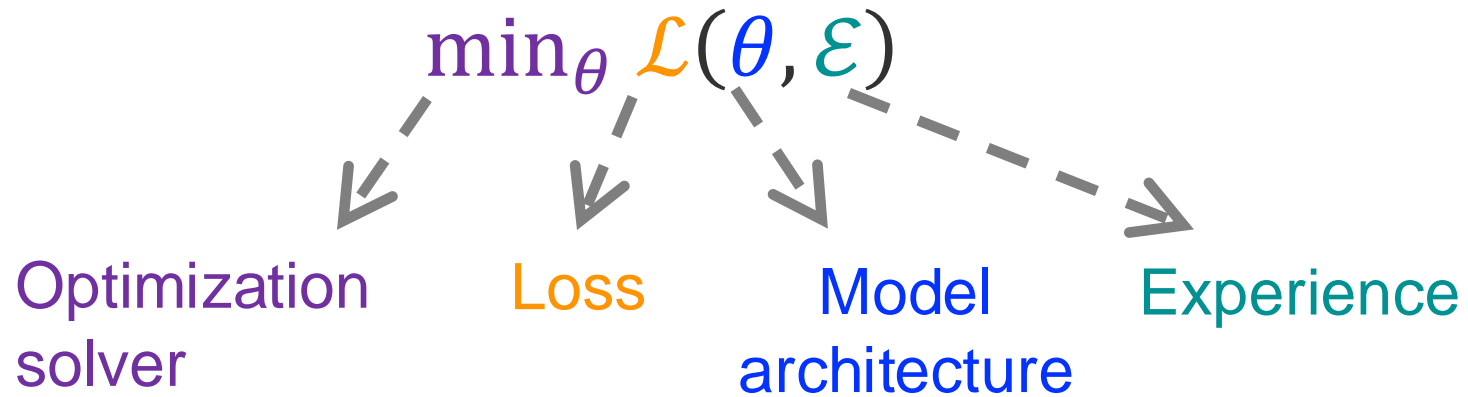


Master classes

... *And all combinations thereof*

Components of a ML solution (roughly)

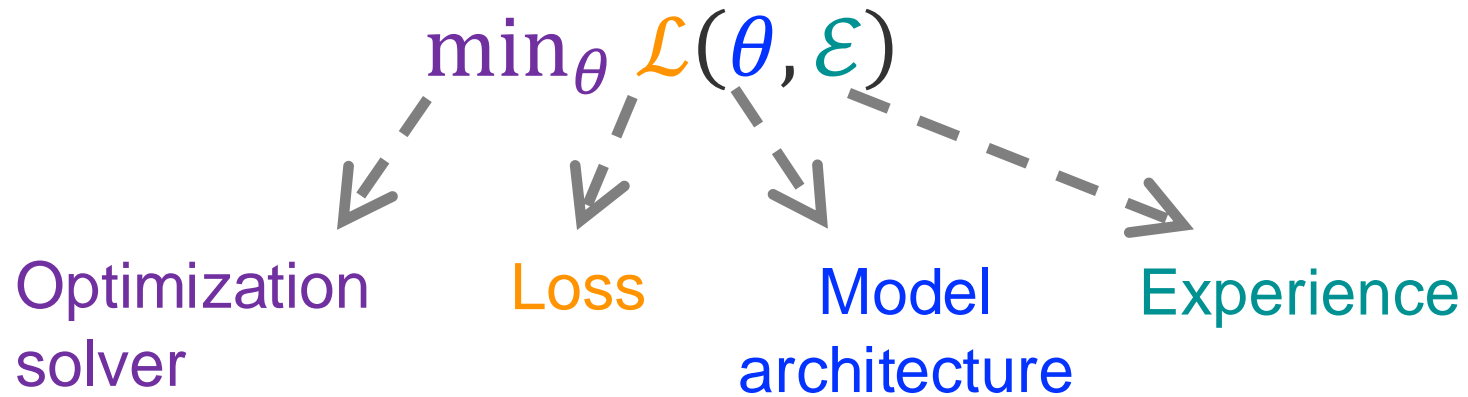
- Loss
- Experience
- Optimization solver
- Model architecture



Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course does **not** discuss model architecture



Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(x, y)$ or $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures

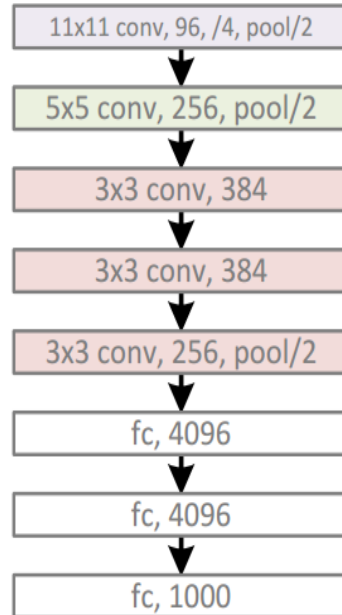
Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

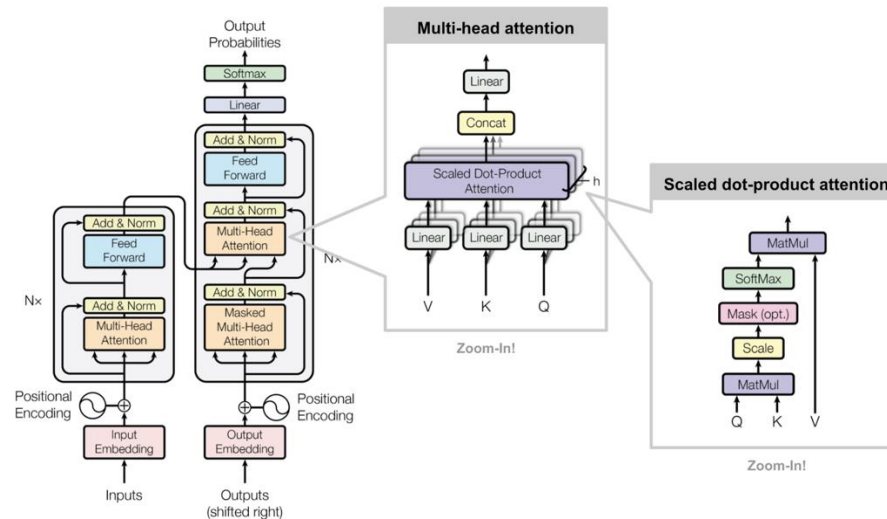
This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(x, y)$ or $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures



Convolutional networks



Transformers

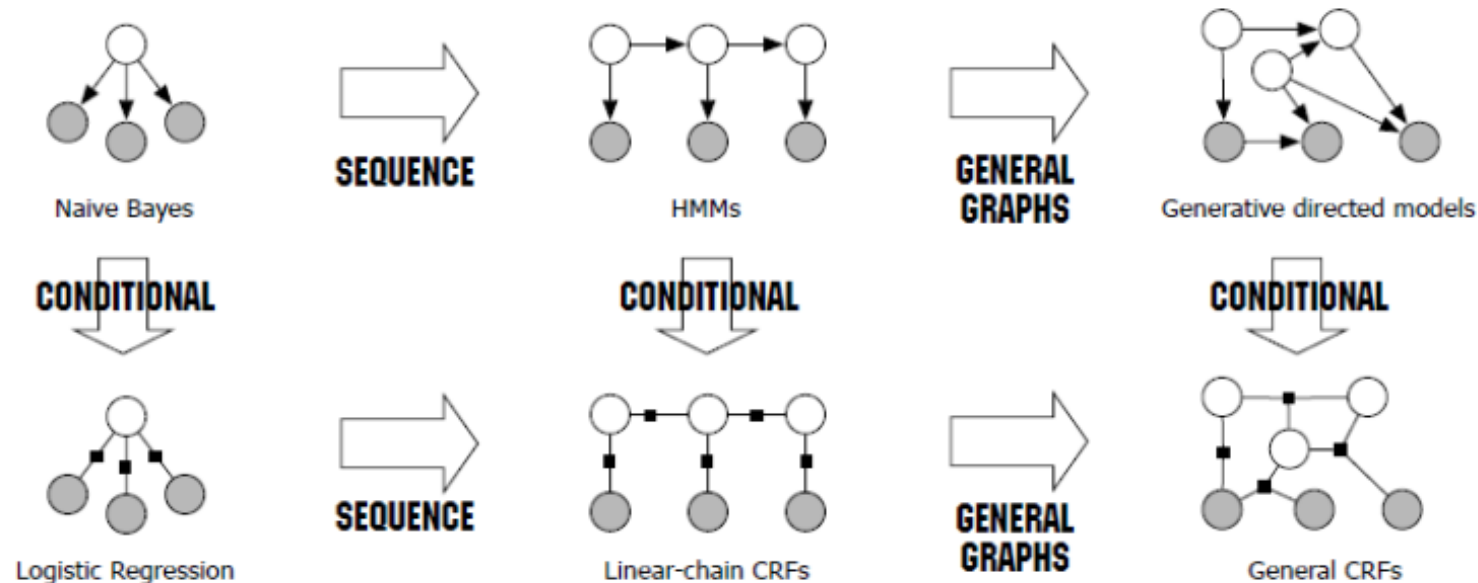
Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- **Model architecture**

This course does **not** discuss model architecture

Model of certain architecture whose parameters are the subject to be learned, $p_{\theta}(x, y)$ or $p_{\theta}(y|x)$

- Neural networks
- Graphical models
- Compositional architectures



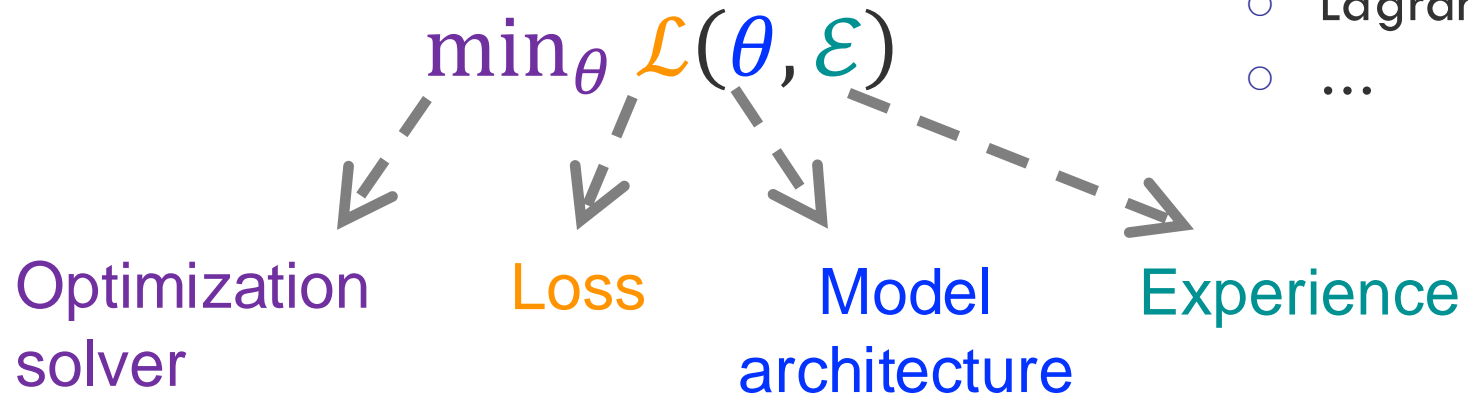
Components of a ML solution (roughly)

- Loss
- Experience
- Optimization solver
- Model architecture

This course discusses *a little* about optimization

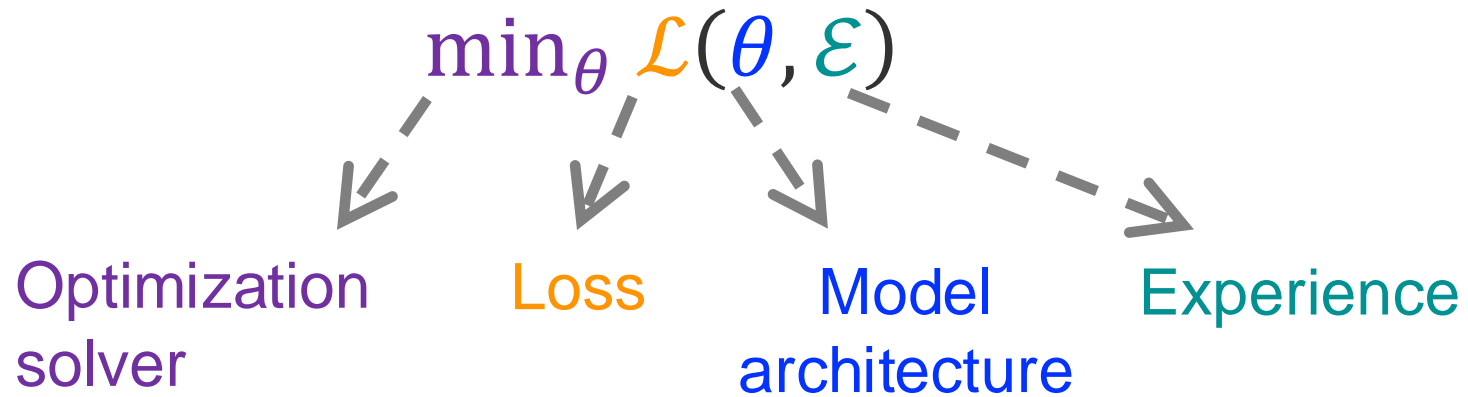
Assuming you know basic procedures:

- (Stochastic) gradient descent
- Backpropagation
- Lagrange multiplier
- ...



Components of a ML solution (roughly)

- Loss This course discusses *a lot* of loss & experience
- Experience
- Optimization solver Core of most learning algorithms
- Model architecture



Machine learning solutions given few data (labels)

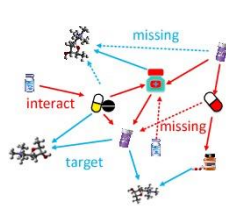
- (1) How can we make more efficient use of **data**?
 - Clean but small-size, Noisy, Out-of-domain
- (2) Can we incorporate **other types of experience** in learning?



Data examples

Type-2
diabetes is 90%
more common
than type-1

Rules/Constraints



Knowledge graphs



Rewards



Auxiliary agents



Adversaries



Master classes

...

And all combinations thereof

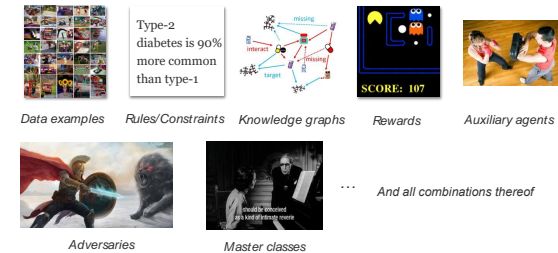
Machine learning solutions given few data (labels)

- (1) How can we make more efficient use of **data**?
 - Clean but small-size, Noisy, Out-of-domain, ...
- Algorithms
 - **Supervised learning**: MLE, maximum entropy principle
 - **Unsupervised learning**: EM, variational inference, VAEs
 - **Self-supervised learning**: successful instances, e.g., BERT, GPTs, contrastive learning, applications to downstream tasks
 - **Distant/weakly supervised learning**: successful instances
 - **Data manipulation**: augmentation, re-weighting, curriculum learning, ...
 - **Meta-learning**

Mostly first half of the course

Machine learning solutions given few data (labels)

- (2) Can we incorporate **other types of experience** in learning?
 - Learning from auxiliary models, e.g., adversarial models:
 - Generative adversarial learning (GANs and variants), co-training, ...
 - Learning from structured knowledge
 - Posterior regularization, constraint-driven learning, ...
 - Learning from rewards
 - Reinforcement learning: model-free vs model-based, policy-based vs value-based, on-policy vs off-policy, extrinsic reward vs intrinsic reward, ...
 - Learning in dynamic environment (**not covered**)
 - Online learning, lifelong/continual learning, ...



Algorithm marketplace

Designs driven by: experience, task, loss function, training procedure ...

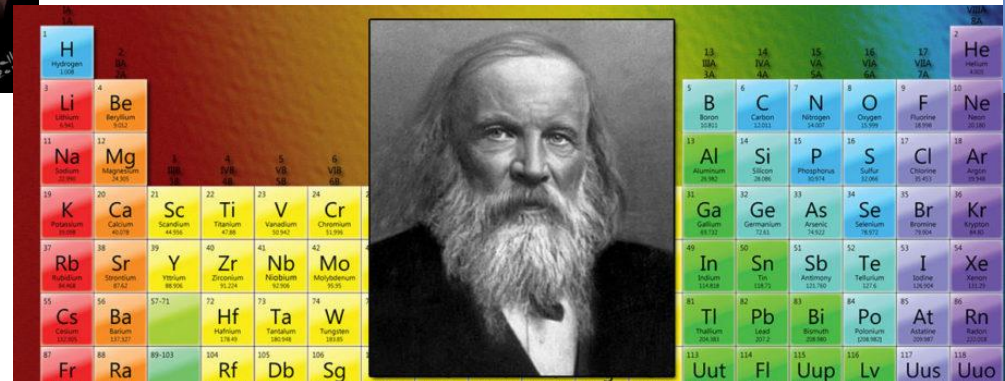
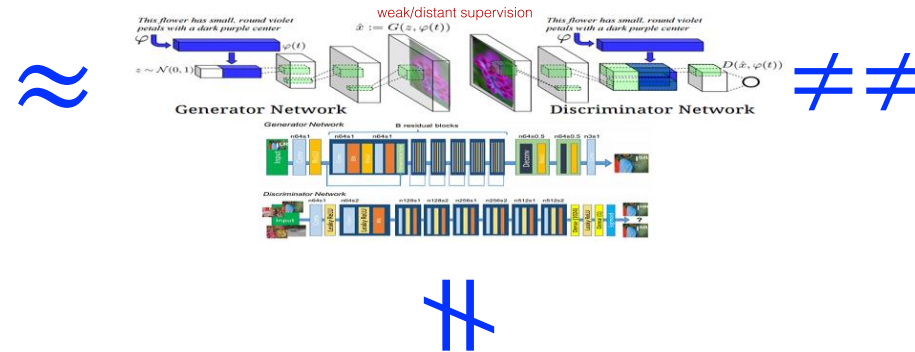


maximum likelihood estimation reinforcement learning as inference
data re-weighting inverse RL active learning
policy optimization
data augmentation reward-augmented maximum likelihood
label smoothing imitation learning softmax policy gradient
actor-critic adversarial domain adaptation
GANs posterior regularization
knowledge distillation intrinsic reward constraint-driven learning
prediction minimization generalized expectation
regularized Bayes
learning from measurements
energy-based GANs
weak/distant supervision

Where we are now? Where we want to be?

- Alchemy vs chemistry

maximum likelihood estimation reinforcement learning as inference
 data re-weighting inverse RL active learning
 data augmentation policy optimization reward-augmented maximum likelihood
 label smoothing imitation learning softmax policy gradient
 actor-critic adversarial domain adaptation
 GANs posterior regularization
 knowledge distillation intrinsic reward constraint-driven learning
 prediction minimization generalized expectation
 regularized Bayes learning from measurements
 energy-based GANs



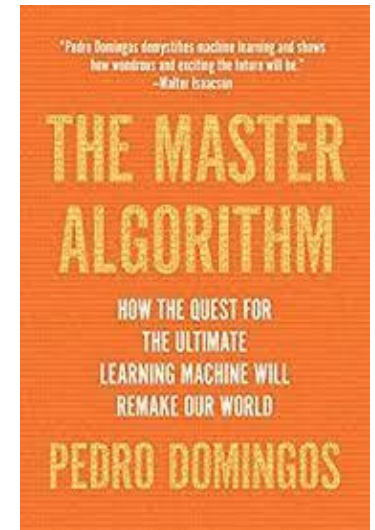
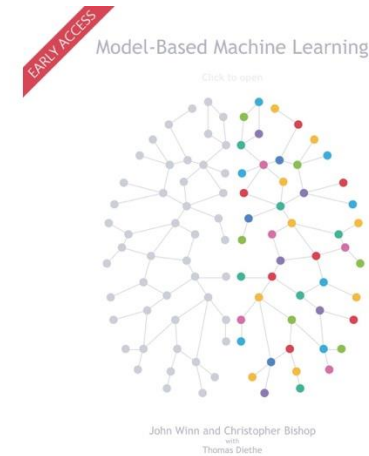
Quest for more standardized, unified ML principles

Machine Learning 3: 253–259, 1989
© 1989 Kluwer Academic Publishers – Manufactured in The Netherlands

EDITORIAL

Toward a Unified Science of Machine Learning

[P. Langley, 1989]



REVIEW

 Communicated by Steven Nowlan

A Unifying Review of Linear Gaussian Models

Sam Roweis*

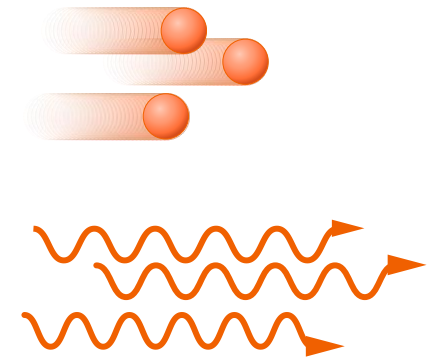
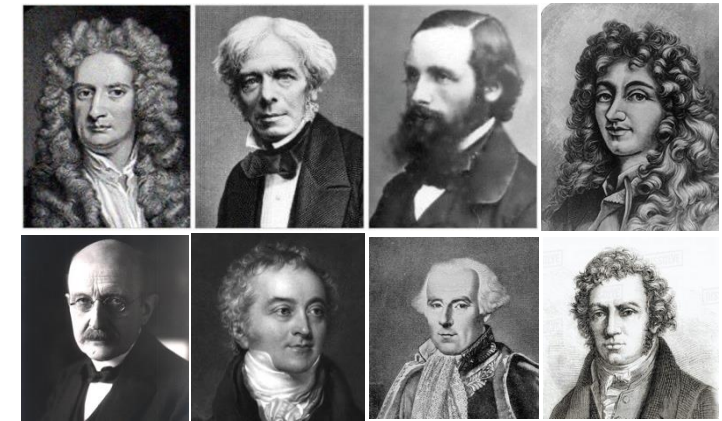
Computation and Neural Systems, California Institute of Technology, Pasadena, CA
91125, U.S.A.

Zoubin Ghahramani*

Department of Computer Science, University of Toronto, Toronto, Canada

Physics in the 1800's

- Electricity & magnetism:
 - Coulomb's law, Ampère, Faraday, ...
- Theory of light beams:
 - Particle theory: Isaac Newton, Laplace, Plank
 - Wave theory: Grimaldi, Chris Huygens, Thomas Young, Maxwell
- Law of gravity
 - Aristotle, Galileo, Newton, ...



“Standard equations” in Physics

Maxwell's Eqns:
original form

| | |
|---|---|
| $e + \frac{df}{dx} + \frac{dg}{dy} + \frac{dh}{dz} = 0$ | (1) Gauss' Law |
| $\mu\alpha = \frac{dH}{dy} - \frac{dG}{dz}$ $\mu\beta = \frac{dF}{dz} - \frac{dH}{dx}$ $\mu\gamma = \frac{dG}{dx} - \frac{dF}{dy}$ | (2) Equivalent to Gauss' Law for magnetism |
| $P = \mu \left(\gamma \frac{dy}{dt} - \beta \frac{dz}{dt} \right) - \frac{dF}{dt} - \frac{d\Psi}{dz}$ $Q = \mu \left(\alpha \frac{dz}{dt} - \gamma \frac{dx}{dt} \right) - \frac{dG}{dt} - \frac{d\Psi}{dy}$ $R = \mu \left(\beta \frac{dx}{dt} - \alpha \frac{dy}{dt} \right) - \frac{dH}{dt} - \frac{d\Psi}{dx}$ | (3) Faraday's Law (with the Lorentz Force and Poisson's Law) |
| $\frac{d\gamma}{dy} - \frac{d\beta}{dz} = 4\pi p'$ $\frac{d\alpha}{dz} - \frac{d\gamma}{dx} = 4\pi q'$ $\frac{d\beta}{dx} - \frac{d\alpha}{dy} = 4\pi r'$ $p' = p + \frac{df}{dt}$ $q' = q + \frac{dg}{dt}$ $r' = r + \frac{dh}{dt}$ | (4) Ampère-Maxwell Law |
| $P = -\xi p \quad Q = -\xi q \quad R = -\xi r$ | Ohm's Law |
| $P = kf \quad Q = kg \quad R = kh$ | The electric elasticity equation ($\mathbf{E} = \mathbf{D}/\epsilon$) |
| $\frac{de}{dt} + \frac{dp}{dx} + \frac{dq}{dy} + \frac{dr}{dz} = 0$ | Continuity of charge |

Diverse
electro-
magnetic
theories



Maxwell's Eqns
simplified w/
rotational
symmetry

$$\nabla \cdot \mathbf{D} = \rho_V$$

$$\nabla \cdot \mathbf{B} = 0$$

$$\nabla \times \mathbf{E} = -\frac{\partial \mathbf{B}}{\partial t}$$

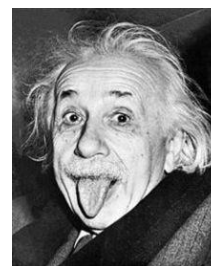
$$\nabla \times \mathbf{H} = \frac{\partial \mathbf{D}}{\partial t} + \mathbf{J}$$



Maxwell's Eqns
further simplified
w/ symmetry of
special relativity

$$\epsilon^{uvk\lambda} \partial_v F_{k\lambda} = 0$$

$$\partial_v F^{uv} = \frac{4\pi}{c} j^u$$



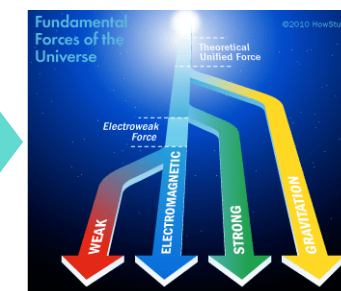
Standard Model
w/ Yang-Mills
theory and US(3)
symmetry

$$\mathcal{L}_{\text{gf}} = -\frac{1}{2} \text{Tr}(F^2)$$

$$= -\frac{1}{4} F^{a\mu\nu} F_{\mu\nu}^a$$



Unification of
fundamental
forces?



1861

1910s

1970s



A “standardized formalism” of ML



Data examples

Type-2 diabetes
is 90% more
common than
type-1

Constraints



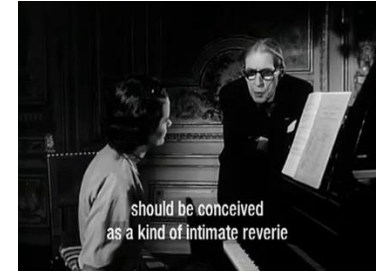
Rewards



Auxiliary agents



Adversaries



Imitation

$$\min_{q, \theta} - \mathbb{H} + \mathbb{D} - \mathbb{E}$$

Uncertainty Divergence Experience

The diagram shows the equation $\min_{q, \theta} - \mathbb{H} + \mathbb{D} - \mathbb{E}$ with three dashed arrows pointing from the terms to labels below: $-\mathbb{H}$ points to 'Uncertainty', $+\mathbb{D}$ points to 'Divergence', and $-\mathbb{E}$ points to 'Experience'.

- Panoramically learn from all types of experience
- Subsumes many existing algorithms as special cases

Will discuss in later in the class

Questions?