

利用 MATLAB/SIMULINK 開發強化式學習應用

Sarah Hung Application Engineer 2019/09/27





What is Reinforcement Learning?

• What is Reinforcement Learning?

- Type of machine learning that trains an
 'agent' through repeated interactions with an environment
- How does it work?
 - Through a trial & error process that maximizes success
- Why should you care about Reinforcement Learning?
 - It enables the use of deep learning for controls and decision-making applications



Controls



Autonomous driving



Robotics

Game Play



















Reinforcement learning:

- Learning through trial & error [*interaction*]
- Complex problems typically need deep learning [Deep Reinforcement Learning]
- It's about learning a behavior or accomplishing a task



A Practical Example of Reinforcement Learning Training a Self-Driving Car





A Practical Example of Reinforcement Learning A <u>Trained</u> Self-Driving Car Only Needs A Policy To Operate



- Vehicle's computer uses the final state-to-action mapping... (policy)
- to generate steering, braking, throttle commands,...
 (action)
- based on sensor readings from LIDAR, cameras,...
 (state)
- that represent road conditions, vehicle position,...
 (environment)

By definition, this trained policy is optimizing driver comfort & fuel efficiency



Deep Networks are commonly found in the agent, because they can model complex problems.





- Turn left
- Turn right
- Brake
- Accelerate



Reinforcement Learning Workflow

Prepare Data



Data access and preprocessing



Ground truth labeling

Train Model

Reinforcement learning



Training agent to perform task



Developing reward system to optimize performance

Simulink – generate data for dynamic systems (planes, cars, robots, etc.)



I	_
L	
<u>⊡</u>	

Enterprise Deployment



Why MATLAB and Simulink for Reinforcement Learning?

Virtual models allow you to simulate conditions hard to emulate in the real world.





Using MATLAB and Simulink for Reinforcement Learning

- Reinforcement learning is a dynamic process
- Decision making problems

 Financial trading, calibration, etc.
- Controls-based problems
 - Lane-keep assist, adaptive cruise control, robotics, etc.





Teach a robot to follow a straight line





Let's try to solve this problem the traditional way





What is the alternative approach?





Walking Robot Example

Reinforcement Learning:

Use past experiences to maximize expected reward





Environment: model a biped robot in Simscape[™] Multibody[™]

📣 Mechanics Explorers - Mechanics Explorer-rlWalkingBipedRobot		– 0 ×
<u>File Explorer Simulation View Tools Window Help</u>		X 5 K
🖳 🖳 💿 🌰 💩 💢 🗊 🗇 🗇 🗊 🗊 🗊 😭 🔛 🛄 뒢 View convention: Z up (X	(Top) 🔷 🍓 🔄 🖶 🗣 🕂 🖵 📜 🥵	
Image: Second		Mechanics Explorer X
Solid Properties Geometry Shape Brick Dimensions [torso_x torso_y torso_z] inertia		
Type Calculate from Geometry		
Density density ka/m^3		
Derived Values Update		
= Graphic		
Type From Geometry		
Erames		
Calid Durainting		
Solid Description *		

17



How Is the Agent Trained?



Decides which action to take



Deep Network Designer



```
];
```

HOME PLOTS APPS	LIVE EDITOR	INSERT			144	1546	0	Search	Documentation	🔎 🌲 Shou
Image: Save open save Image: Print Files Image: Pr		M Code &	Control ♥ 🧏 🕍 Refactor ♥ 🛐 📑 💽	Run Section Section Break Run to End	n De Nance Run	Step	Stop			
FRE NAVIGATE	TEXT		CODE	SECTION		RUN				
🖻 🗭 💽 🌄 📜 🔸 C: 🕨 Users 🕨 shmitra 🕨 Work	Deep_Learnin	g 🖡 Seminar 🕨 18b	 FoodDataTransferLea 	rning-master 🕨						
urrent Folder	🖲 🖬 Live Edit	or - C:\Users\shmitra\	Work\Deep_Learning\Se	minar\18b\FoodDataTra	nsferLearning-r	naster\Fo	odDataTra	nsfer 💿 🗙	Workspace	
Name +	FoodD	ataTransferLearning.ml	x × +						Name -	Value
i 🧍 foodData	53	layer	<pre>rs(end) = classif</pre>	icationLayer('Nam	e','classo	utput'));	A (5)	ans	5x2 table
HelperFunctions	54								doTest	0
i 🦲 results	55	layer	rs_train = layers	;					doTrain	0
exclude.xml	56							1	H II	5
FoodDataTransferLearning.mlx	57	case 'goo	oglenet' % DAG ne	twork				4	imagepath	'foodData\train'
T ReadMe.m	58		w 10 10 100						imagesize	[224,224,3]
	59	lgrag	oh = layerGraph(n	et);					imds	1x1 ImageDatastore
	60	lgrap	<pre>ph = removeLayers</pre>	(lgraph, {'prob',	'output'})	;			img abel	1x1 cotenorical
	61								abels	5x1 categorical
	62	larray = [fullyConnectedLayer(numClasses, Name', 'fc', WeightLear						layers	144x1 Layer	
	64		sortmaxta	stical auco ('Name'	dx J				nchoices	3
	64	classificationLayer(Name , classourput)];							📧 net	1x1 DAGNetwork
	66								netName .	'googlenet'
	67	1 gray	s topin = length	(TBuabu) Tossore	TOPPTITE!	, 10//0	sy),	=	nlabel	5
	68	2090	aller ann - aller abr	,					outputsize	[229,229] 1x1000 double
	69	end						_	prediction	1x1000 sinale
								*	- psorted	1x1000 single
									scores	[0.2859;0.1935;0.157
oodDataTransferLearning.mtx (Live Script)	Command	Window						C	topclasses	3x1 cell
	New to MA	TLAB? See resources f	or Getting Started.					,	trainDS	1x1 ImageDatastore
fx		x >>						vaiDS	1x1 ImageDatastore	
No details available		1							visimos	1x1 imageDatastore
									<	
.										



Steps in the Reinforcement Learning Workflow



Environment



Reward



Policy



Agent







Reinforcement Learning Toolbox New Product in R2019a

- Built-in and custom reinforcement learning algorithms
- Environment modeling in MATLAB and Simulink
 - Existing scripts and models can be reused!
- Deep Learning Toolbox support for representing policies
- Training acceleration with PCT and MATLAB Parallel Server
- Deployment of trained policies with GPU Coder and MATLAB Coder
- Reference examples for getting started (control systems, automotive, robotics)



teinforcement Learning Toolbox.™ provides functions and blocks for training policies sing reinforcement learning algorithms including DQN, A2C, and DDPG. You can se these policies to implement controllers and decision-making algorithms for omplex systems such as robots and autonomous systems. You can implement the olicies using deep neural networks, polynomials, or look-up tables.

The toolbox lets you train policies by enabling them to interact with environments represented by MATLAB' or Simuline' models. You can evaluate algorithms, experiment with hyperparameter settings, and monitor training progress. To inprove training performance, you can run simulations in parallel on the cloud, computer clusters, and GPUs (with Parallel Computing ToolboxTM and MATLAB Parallel ServerTM).

Through the ONNX** model formal, existing policies can be imported from deep learning frameworks such as TensorFlow** Karas and PyTorch (with Deep Learning Toolbox***). You can generate optimized C, C++, and CUDA code to deploy trained policies on microcontrollers and GPUs.

The toolbox includes reference examples for using reinforcement learning to design controllers for robotics and automated driving applications.





Reinforcement Learning Toolbox Workflow



Reinforcement Learning Toolbox

Design and train policies using reinforcement learning

Reinforcement Learning Toolbox[™] provides functions and blocks for training policies using reinforcement learning algorithms including DQN, A2C, and DDPG. You can use these policies to implement controllers and decision-making algorithms for complex systems such as robots and autonomous systems. You can implement the policies using deep neural networks, polynomials, or look-up tables.

The toolbox lets you train policies by enabling them to interact with environments represented by MATLAB[®] or Simulink[®] models. You can evaluate algorithms, experiment with hyperparameter settings, and monitor training progress. To improve training performance, you can run simulations in parallel on the cloud, computer clusters, and GPUs (with Parallel Computing Toolbox[™] and MATLAB Parallel Server[™]).

Through the ONNX[™] model format, existing policies can be imported from deep learning frameworks such as TensorFlow[™] Keras and PyTorch (with Deep Learning Toolbox[™]). You can generate optimized C, C++, and CUDA code to deploy trained policies on microcontrollers and GPUs.

The toolbox includes reference examples for using reinforcement learning to design controllers for robotics and automated driving applications.

Getting Started

Learn the basics of Reinforcement Learning Toolbox

MATLAB Environments Model reinforcement learning environment dynamics using MATLAB

Simulink Environments Model reinforcement learning environment dynamics using Simulink models

Policies and Value Functions

Define policy and value function representations, such as deep neural networks and Q tables

Agents

Create and configure reinforcement learning agents using common algorithms, such as SARSA, DQN, DDPG, and A2C

Training and Validation

Train and simulate reinforcement learning agents

Policy Deployment

Code generation and deployment of trained policies

R2019a

Release Notes PDF Documentation

📣 MathWorks

Examples



MathWorks[®]

Key Takeaways

- Reinforcement learning can solve complicated (control) problems.
- Reinforcement learning can use *neural networks* to handle continuous or high-dimensional state and action spaces.
- Choosing good states, actions, and reward functions is extremely important!
- Reinforcement learning requires a lot of simulation data.
- DDPG is a *high variance* algorithm. Run several times and validate results!