

Machine Learning in 15:

A journey from beginner to advanced ML builder with SageMaker Studio Lab

Vadim Omeltchenko
Sr. AI/ML Solutions Architect
Amazon Web Services




What is Amazon SageMaker Studio Lab?

A free development environment for anyone to learn and experiment with ML

Amazon SageMaker Studio Lab

A no-charge, no-configuration environment that enables users to learn and experiment with machine learning



Create an account with an email address for free

No setup or configuration required

15 GBs to save your work projects

As many compute sessions as you need –
CPU (12 hours)/GPU (4 hours)

Access any notebook on GitHub

Graduate to SageMaker Studio when ready

Who are the Studio Lab users?



Academics

I want the right skills for a great career

Basic theory and learn Python/R



Developers

I want to expand my technical skills with data science

Learn Python/R
corporate data



Environment to practice

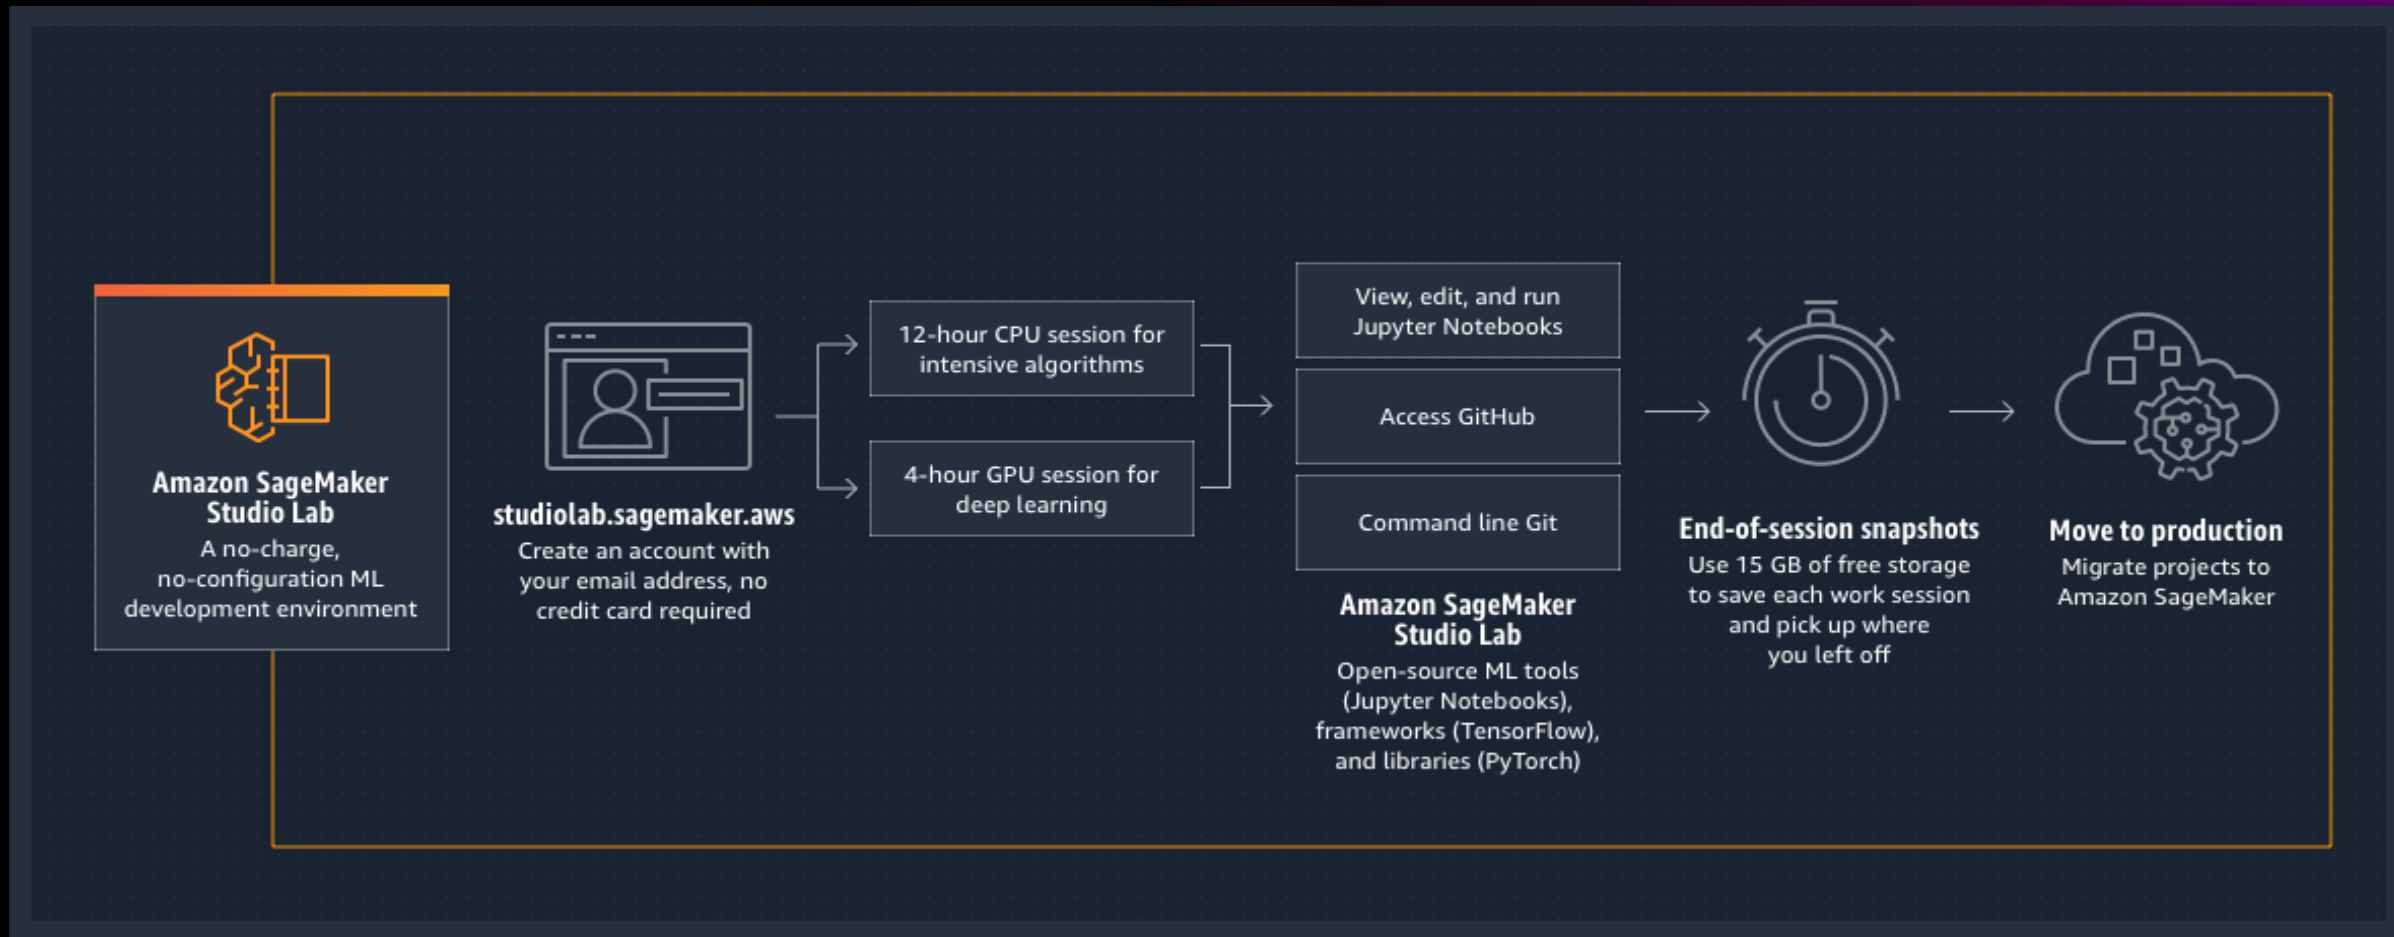


Data scientists

I want experiment ML and move them into production

Data science communities

How does it work?



Signup

- **Fill out account request form**
- **Provide referral code (if available)**
- **Wait for the account to be approved**

Request account

Request a free Amazon SageMaker Studio Lab account.

Enter your email*

Enter your first name

Enter your last name

Select your country

Enter your company or organization name

Select your occupation

Why are you interested in Amazon SageMaker Studio Lab?

Enter referral code

Submit request



Project page

Status of your project

- Time remaining
- Select runtime type
- Start or stop instance

Start a user session

- CPU — 12 hours
- GPU — 4 hours

Open project (in a new browser tab)

Assets to get you started

- Dive into Deep Learning
- AWS Machine Learning University
- Hugging Face
- Popular blogs
- Community links

amazon SageMaker Studio Lab michele

My project

GPU runtime limits have changed. You can use GPU for up to 4 hours at a time and up to 8 hours in a 24-hour period.

Runtime status: **Stopped** Runtime remaining: — Compute type: CPU GPU [▶ Start runtime](#) [Open project](#)

Learn and experiment

Dive into Deep Learning (D2L)

Level up your understanding of machine learning with a free interactive book (150 Jupyter notebooks) that teaches the ideas, the math, and the code. Adopted at 300 universities from 55 countries including Stanford, MIT, Harvard, and Cambridge.

[Open D2L notebooks](#)

AWS Machine Learning University

Get access to the same machine learning courses used to train Amazon's own developers on machine learning. Learn how to use ML with the learn-at-your-own-pace MLU Accelerator learning series.

[Open MLU notebooks](#)

Resources and community

Hugging Face

Hugging Face is the home of the Transformers library and state-of-the-art natural language processing, speech, and computer vision models.

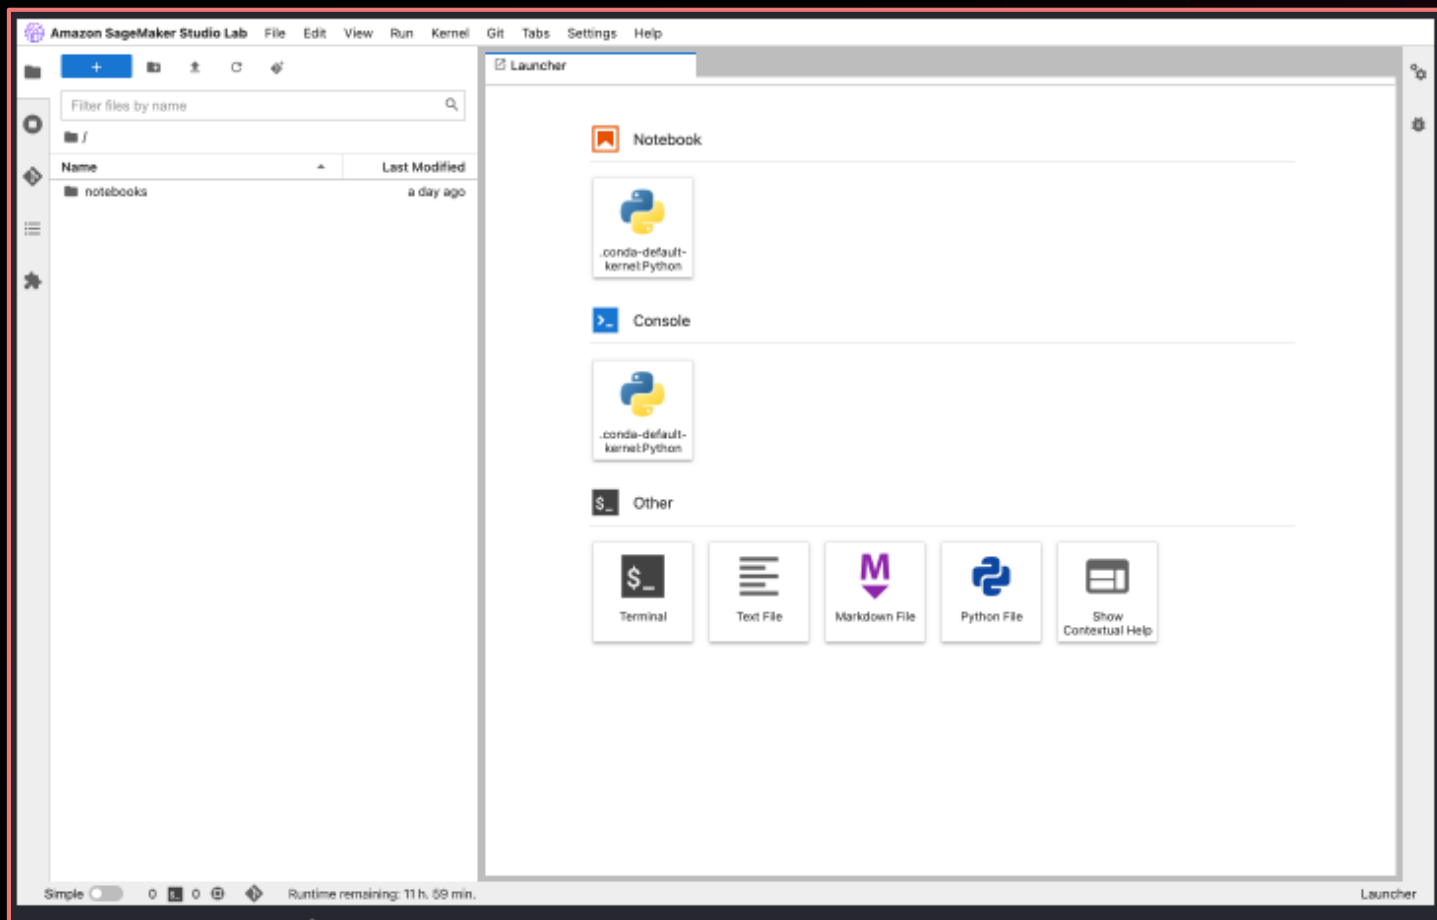
[Open Hugging Face notebooks](#)

Machine Learning Blog

Stay up-to-date with the latest developments, research, and techniques in the AI and machine learning space.

[AWS Machine Learning Blogs](#)

Notebook development environment



Familiar JupyterLab experience

- Terminal access
- Git/GitHub

Your ML environment on AWS

- Compute dedicated to you
- 12 hours CPU/4 hours GPU
- Install the libraries you want
- Dedicated 15 GB for your project
- Unlimited user sessions
- Pick up where you left off

Studio Lab Environment

Base image

Python 3.9
bzip2
build-essential
curl
git
libgl1-mesa-glx
nano
rsync
unzip
wget
ca-certificates
pip
ipykernel-6.4

Supported frameworks

PyTorch 1.9
TensorFlow 1.15 and 2.6
MxNet 1.8
Hugging Face
AutoGluon 0.3.1
Scikit-learn 0.24
PyTorch ecosystem
OpenCV
scipy
numpy

Amazon SageMaker Studio Lab vs Studio

Amazon SageMaker

PREPARE

SageMaker Ground Truth

Label training data for machine learning

SageMaker Data Wrangler

Aggregate and prepare data for machine learning

SageMaker Processing

Built-in Python, BYO R/Spark

SageMaker Feature Store

Store, update, retrieve, and share features

SageMaker Clarify

Detect bias and understand model predictions

BUILD

SageMaker Studio Notebooks & Notebook Instances

Jupyter notebooks with elastic compute and sharing

Built-in and bring-your-own algorithms

Dozens of optimized algorithms or bring your own

Local mode

Test and prototype on your local machine

SageMaker Autopilot

Automatically create machine learning models with full visibility

SageMaker JumpStart

Pre-built solutions for common use cases

TRAIN & TUNE

Managed training

Distributed infrastructure management

SageMaker Experiments

Capture, organize, and compare every step

Automatic model tuning

Hyperparameter optimization

Distributed training libraries

Training for large datasets and models

SageMaker Debugger

Debug and profile training runs

Managed Spot training

Reduce training cost by 90%

DEPLOY & MANAGE

Managed deployment

Fully managed, ultra low latency, high throughput

Kubernetes and Kubeflow Integration

Simplify Kubernetes-based machine learning

Multi-model endpoints

Reduce cost by hosting multiple models per instance

SageMaker Model Monitor

Maintain accuracy of deployed models

SageMaker Edge Manager

Manage and monitor models on edge devices

SageMaker Pipelines

Workflow orchestration and automation



Amazon SageMaker Studio Lab

A no charge, no setup ML development environment

Scheduled notebooks

The screenshot displays the Amazon SageMaker Studio Lab interface. On the left, a file explorer shows a directory structure with 'my-notebook.ipynb' selected. A context menu is open over this file, with 'Create Notebook Job' highlighted by an orange circle. The main editor shows Python code for training a model, including a loop for printing statistics and a 'Finished Training' message. Below the code, there are sections for 'Save the model' and 'Save to S3'. A 'Schedule' dialog box is overlaid on the right, with 'Run on a schedule' selected. The dialog includes fields for 'Interval' (set to 'Weekday') and 'Time' (set to '00:00'). A note specifies that times are in UTC and affected by daylight saving time. 'Cancel' and 'Create' buttons are at the bottom right of the dialog.

```
optimizer.step()

# print statistics
running_loss += loss.item()
if i % 2000 == 1999: # print every 2000 mini-batches
    print(f'[{epoch + 1}, {i + 1:5d}] loss: {running_loss:.3f}')
    running_loss = 0.0

print('Finished Training')
```

PATH = './model.pth'
torch.save(net.state_dict(), PATH)

! aws s3 cp ./model.pth s3://sagemaker-my-model-bucket/model

Interval: Weekday
Time: 00:00
12:00 AM

Specify time in UTC (subtract 10 hours from local time)
Schedules in UTC are affected by daylight saving time or summer time changes

Cancel Create

<https://aws.amazon.com/blogs/machine-learning/run-notebooks-as-batch-jobs-in-amazon-sagemaker-studio-lab/>

Export Studio Lab environment to SageMaker

Step 1: Export your Studio Lab conda environment

Step 2: Save your Studio Lab artifacts

Step 3: Import your Studio Lab artifacts to Studio

Step 4: Install your Studio Lab conda environments in Studio

<https://docs.aws.amazon.com/sagemaker/latest/dg/studio-lab-use-migrate.html>



How to get started?

1. Login at <https://studiolab.sagemaker.aws/>
2. Resources <https://aws.amazon.com/sagemaker/studio-lab/>

Amazon SageMaker Studio Lab

Learn and experiment with ML using a no-setup, free development environment

Get started with SageMaker Studio Lab

- Free machine learning development environment that provides the compute, storage, and security to learn and experiment with ML
- Get started with a valid email address —no need to configure infrastructure or manage identity and access or even sign up for an AWS account
- GitHub integration and preconfigured with the most popular ML tools, frameworks, and libraries so you can get started immediately

BLOG

How to get started with SageMaker Studio Lab

[Read more »](#)

VIDEO

Deep Dive presentation on SageMaker Studio Lab

[Read more »](#)

VIDEO

Use SageMaker Studio Lab to improve disaster response

[Read more >>](#)

Q & A

