

Twitter Thread by Goku Mohandas



Goku Mohandas

@GokuMohandas




■ Putting ML in Production! We're going to publicly develop @madewithml's first ML service. Here is the broad curriculum:


- ■ Product
- ■ Data
- ■ Modeling
- ■ Scripting
- ■ API
- ■ Production


More details (lessons, task, etc.) here: <https://t.co/xmMm9XGK9j>


Thread ■


Made With ML


Putting ML in Production


@GokuMohandas





Product


- Objective
- Solutions (UX, technical)
- Constraints (performance & compute)
- Evaluation


Data


- Collection
- Annotation
- Exploratory data analysis
- Splitting
- Preprocessing
- Augmentation
- Versioning


Modeling


- Baselines
- Pretraining
- Experiment tracking
- Evaluation
- Constraints
- Inspection
- Optimization
- Inference


Scripting

- OOPs
- Virtual envs
- Debugger
- Logging
- Testing (unit, e2e, regression, ML robustness)
- Styling
- Precommit
- Makefile
- Git


APIs

- FastAPI
- SQL + DB
- User management
- Authentication
- Docker
- Documentation


Production

- Serving
- Environments
- Monitoring (performance, drift)
- CI/CD (GitHub actions)
- Active learning
- Scaling

<https://madewithml.com/courses/putting-ml-in-production/>

Questions that this thread will answer:

- What is it?
- Who is this course for?
- What is the format?
- What makes this course unique?
- Why constrain to open source tools?
- What are my qualifications?
- Why is this free?
- What are the prerequisites?

<https://t.co/xmMm9XGK9j>

What is it?

Putting ML in Production: a guide and code-driven case study on MLOps. We will be developing and deploying Made With ML's first ML service, from Product → ML → Production, with open source tools.

This ML service will act as a foundation for all future ML features and subsequent iterations. The first feature is tagifai - multilabel classification of tags for a project. We'll discuss the need and utility of this feature in the first lesson.

Who is this course for?

- ML developers looking to become end-to-end ML developers.
- Software engineers looking to learn how to responsibly deploy and monitor ML systems.
- Product managers who want to have a comprehensive understanding of the different stages of ML dev.

What is the format of each lesson?

- Intuition: high level overview of the concepts.
- Code: simple code examples to illustrate the concept.
- Application: applying the concept to our specific task.
- Extensions: brief look at other tools and techniques that will be useful.

What makes this course unique?

1. Hands-on
2. Intuition-first
3. Software engineering
4. Focused yet holistic
5. Open source tools

1. Hands-on:

If you search production ML or MLOps online, you'll find great blog posts and tweets. But in order to really understand these concepts, you need to implement them.

2. Intuition-first:

We will never jump straight to code. In every lesson, we will develop intuition for the concepts and think about it from a product perspective.

3. Software engineering:

This course isn't just about ML. In fact, it's mostly about clean software engineering! We'll cover important concepts like versioning, testing, logging, etc. that really makes this a production-grade product.

4. Focused yet holistic:

For every concept, we'll not only cover what's most important for our specific task (this is the case study aspect) but we'll also cover related methods (this is the guide aspect) which may prove to be useful in other situations.

4. (cont.) For example, when we're serving our application, we'll expose our latest model as an API endpoint. However, there are several other popular ways to serving models and we'll briefly illustrate those and talk about advantages / disadvantages.

5. Open source tools:

We will be using only open source tools for this project, with the exception of [@googlecloud](#) for storage and compute (free credit will be plenty).

Why constrain to open source tools?

1. We can focus on the fundamentals, everyone can participate (single player mode as my friend [@eugeneyan](#) coined) and you will have much better understanding when you do finally use a paid tool at work (if you want to).
2. Large companies that deploy ML to production have complicated and scaled processes that don't make sense for the vast majority of companies / individuals.

Note: I will regularly make suggestions for tools (other open source, freemium and paid) as we progress because they each have their unique advantages.

For example, for data versioning and experiment tracking, we'll use [@DVCorg](#) + [@MLflow](#) but we'll also highlight why you may consider [@weights_biases](#) or [@Cometml](#) because it's important to know about them and what they each bring to the table.

My qualifications for teaching this:

1. I've deployed large scale ML systems at [@Apple](#) as well as smaller systems with constraints at startups and want to share the common principles I've learned along the way.
2. I created [@madewithml](#) so that the community can explore, learn and build ML and I learned how to build it into an end-to-end product that's currently used by over 5K daily active users.

You can learn more at my personal website or LinkedIn.

LinkedIn: <https://t.co/xWmPKz53vw>

Personal website: <https://t.co/NpLSczadPn>

Why is this free?

1. Personal reason: Every day, people explore the amazing work on [@madewithml](#) to learn from and contribute themselves. To stay consistent with this free spirit, I'm releasing this free course to pass on the lessons I've learned from my mentors and experiences.
2. Societal reason: This is especially targeted for people who don't have as much opportunity around the ■. I firmly believe that creativity and intelligence are randomly distributed but opportunity is siloed. I want to enable more people to create and contribute to innovation.

What are the prerequisites?

- You should have some familiarity with Python and basic ML algorithms. While we will be experimenting with deep learning (w.r.t compute/performance tradeoffs), you can easily apply the lessons to any class of ML models.

<https://t.co/V35zXocadQ>

The course hasn't even begun yet but there's already quite a few friends to thank for helping me thinking through some of this. We'll be referring to their work throughout the course. [@josh_tobin](#) [@jeremyjordan](#) [@eugeneyan](#) [@nlpguy](#) [@MLinProduction](#) [@FullStackML](#)

First lesson releases next week (■ Product) & subsequent lessons will follow a weekly cadence. Be sure to follow me or [@madewithml](#) for updates, discussions & feedback because I'll be creating the course content dynamically using the community's feedback.

<https://t.co/cmTTeALWz1>