

Anyone Can Cook (AI Edition)

Machine Learning Singapore #MSLG

- Martin Andrews = DSTA-BrainHack-2026 @ mdda.net

10-June-2026



About Me

- ML / Startups / Finance
 - NYC → Singapore in Sep-2013
- 2014 = 'fun' :
 - Machine Learning, DL, NLP
 - Robots, drones
- Since 2015 = 'serious'
 - NLP + deep learning
 - Including Papers...
 - & GDE ML
 - ML-Singapore co-organiser...
 - & Red Dragon AI...



Quick Poll

= Show of hands

- Who has seen the 2007 movie "Ratatouille"?
- Who has been to a Machine Learning Singapore MeetUp (at Google)?



Outline

- A bit about the movie
 - one minor spoiler!
- Technical content
- "Anyone can Cook!"
- Wrap-up & QR-code

Characters

Main Characters

- Alfredo Linguini
 - Human 'Chef'
- Remy
 - Rat that loves food



Is this a Story about AI being in Control?

- In this analogy, is AI the Rat that uses the human as a meat puppet? :: NO
- This talk is not about how AI is coming to take all the jobs...
 - ... it's about how we can be the Rat with Taste



Remy

- Loves food
- Has a heightened sense of taste
 - has dreams
 - might feel threatened by AI



Émile (Remy's Brother)

- Loves to eat
 - likes as much food as possible
- Enjoyer of Slop
 - AI enables creation of Slop ...
- So, what do we mean by slop?
 - Why do models do this?

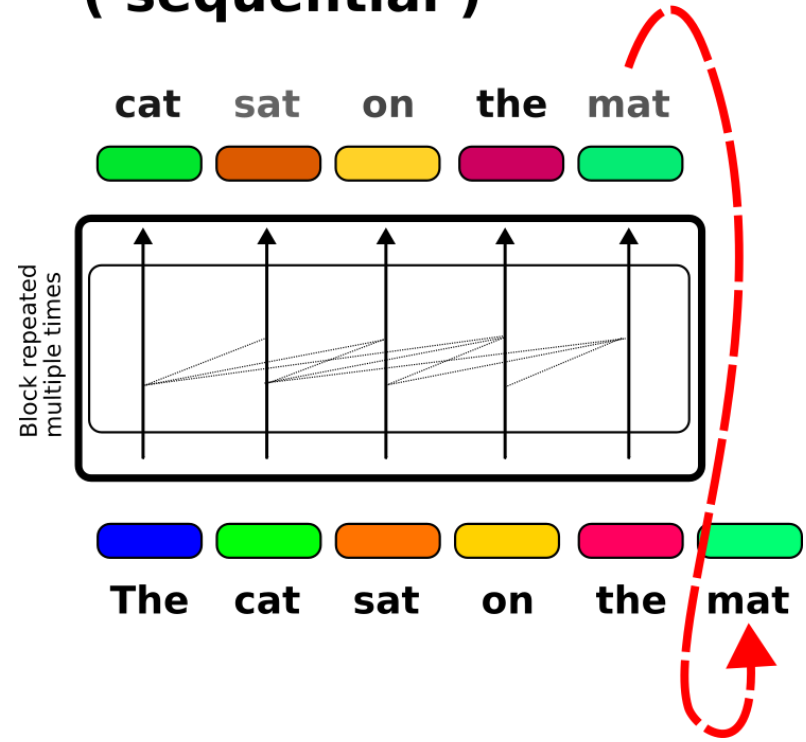


Getting Technical...

Large Language Models

- An LLM simply predicts the next token
 - given the preceding text
 - and then adds it on the end
- In practice, they predict a *distribution* over all possibilities
 - and then choose one
 - depends on the temperature (and top_p , etc)
- Amazing that it works so well!


Inference Mode (sequential)



LLM Pre-Training

- LLMs train on huge datasets
 - Multiple-Trillion tokens
 - Huge compute
- To create a :
 - Base Large Language Model
 - Predict 'what the internet would say'
 - ... token by token
- But this can be very wild
 - how do we guarantee that an LLM remains tame?

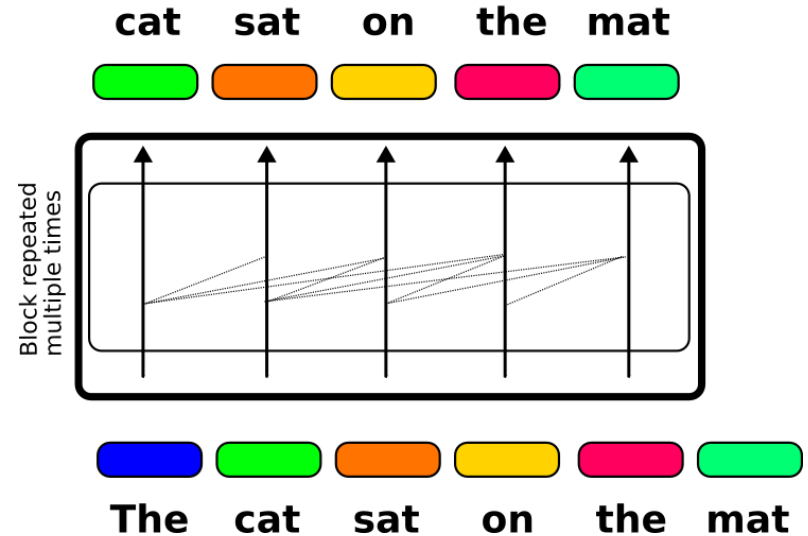
Training Mode (parallel)



Supervised Fine-tuning

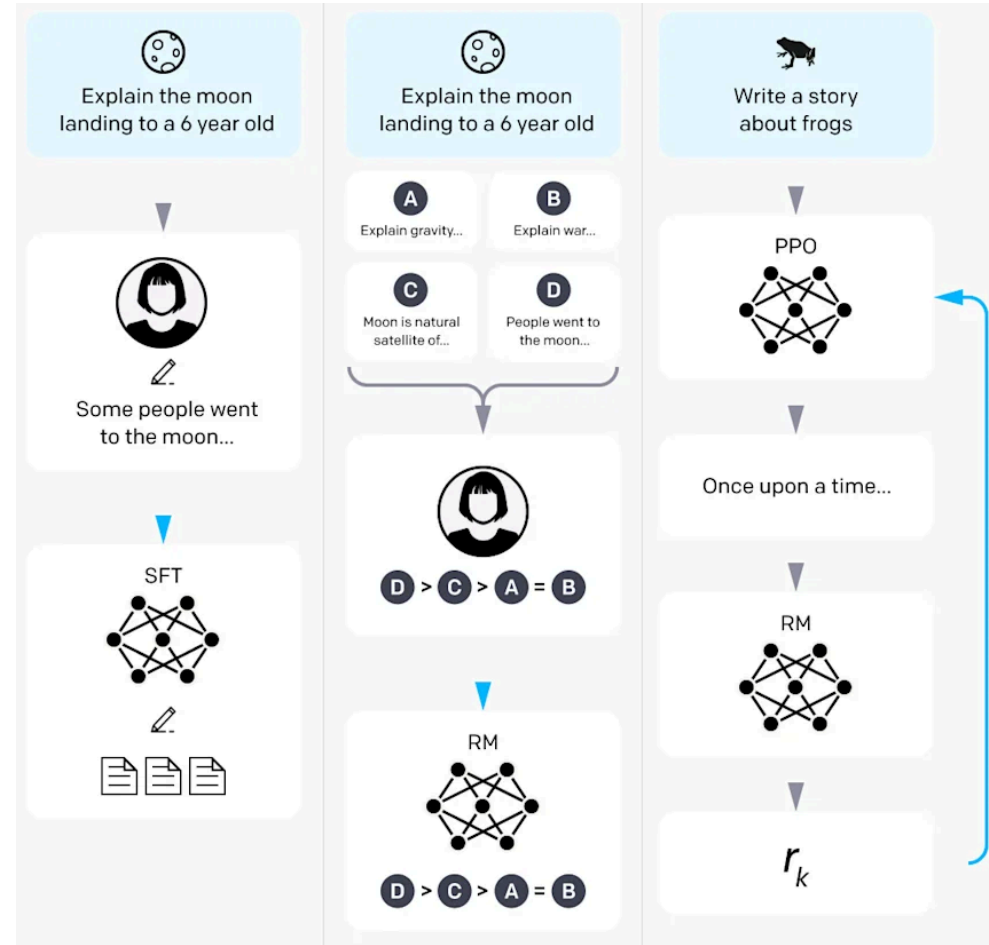
- Models for answering questions
 - trained by seeing lots pairs of Questions and Answers
- This makes them 'more useful'
 - ... ideally
 - but at what cost?
- Questions in the training data *have answers*
 - so models will always answer
 - even if they don't have the knowledge
 - ... and we don't know what knowledge they have
 - this gives rise to *hallucinations*

Training Mode (parallel)



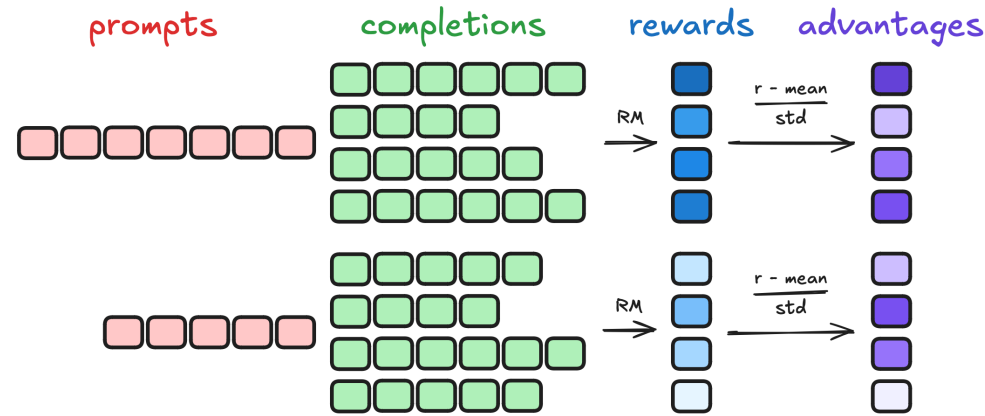
RLHF

- Reinforcement Learning from Human Feedback
- Models learn what humans prefer
 - by seeing how humans rank different answers produced by the model
- But this is where Sycophancy creeps in...
- Fighting sycophancy:
 - say 'my rival says...'
 - get the model create its own alternative answer(s)
 - then do Pros/Cons on both



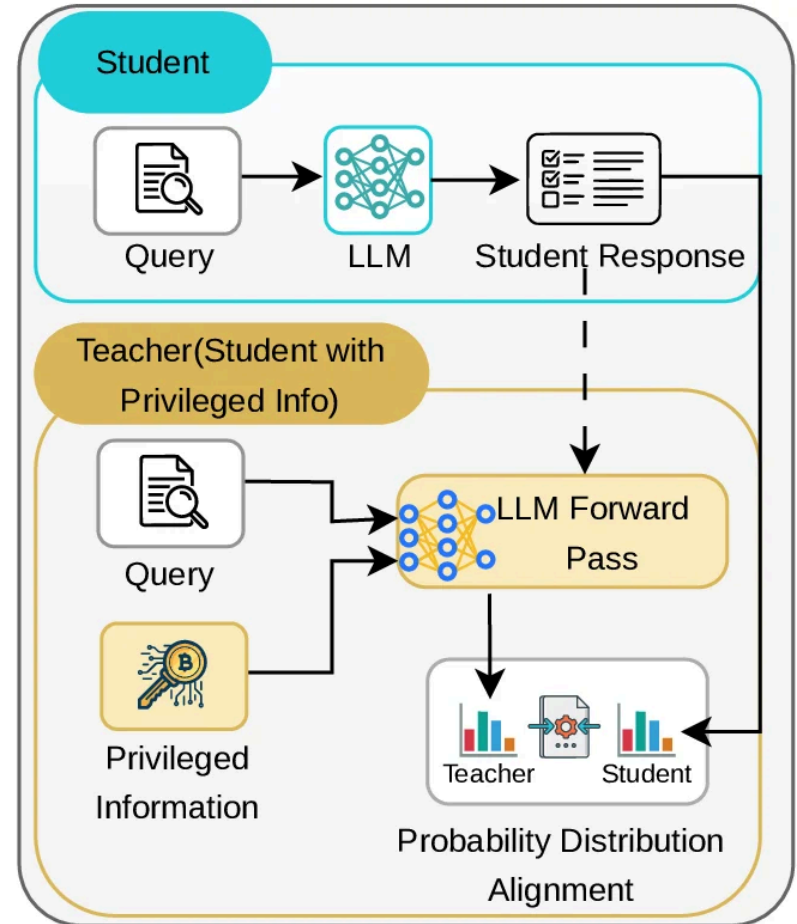
RLVR

- Reinforcement Learning from Verifiable Rewards
- Models learn to Reason
 - by 'rolling out' thinking traces
 - and getting rewards for correct answers
 - ... requires problems with Verifiable answers
 - has worked for ~18 months
 - DeepSeek published 'the trick' in Jan-2025
- Major labs hide the reasoning tokens
 - but Open Source models show us their thinking
 - super important for safety...



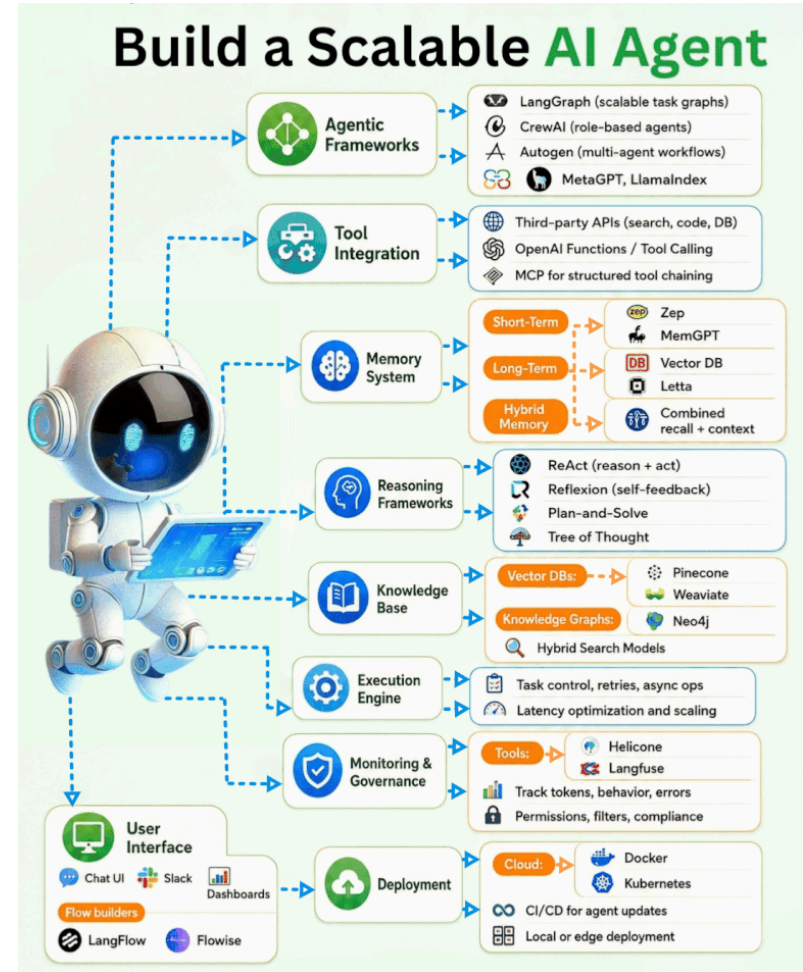
On-Policy Distillation

- Use a teacher to mark student roll-outs
 - can gain more information per *reasoning trace*
 - but need to calibrate the Teacher
- Should a good teacher:
 - Be smarter overall?
 - but what taught them?
 - Have the answer beforehand?
 - Adjust difficulty dynamically?
 - Compete with student somehow?
- Now we're getting to the forefront of research...



... so many things to explore!

- Agents = much wider scope
 - 'just' builds on what we've already covered
 - breakthroughs in performance in last 6 months...
- Next for Agents:
 - longer horizon tasks
 - self-checking
 - skill self-development
 - auto-research ...
- Next for research:
 - in-distribution vs creativity
 - ...



Back to the movie




Anton Ego

- A critic
 - a.k.a The Establishment
- Believes that only a few people are capable of achieving greatness
 - essentially a gatekeeper
 - but not a *builder*
- Distainful of the book he's holding...

Auguste Gusteau

- Famous chef
 - Restaurant critically acclaimed
 - Author of "Anyone Can Cook!"
- Does this mean:
 - "Cooking is easy and skill-less" ?
 - NO!
- "Anyone Can Cook!" means:
 - "Great cooking can come from anyone"
 - there are no actual barriers
 - ... just passion and determination
- Machine Learning Singapore:
 - mantra since 2017
 - "You can do AI for yourself"



Wrap-Up

- AI progress has significantly sped up
 - ... since the beginning of the year
 - definitely time to re-calibrate
- And AI may seem daunting ...
 - but understanding the ideas is not impossible
 - and experimentation is open to all!
- Anyone Can Cook!



Machine Learning SG MeetUp Group

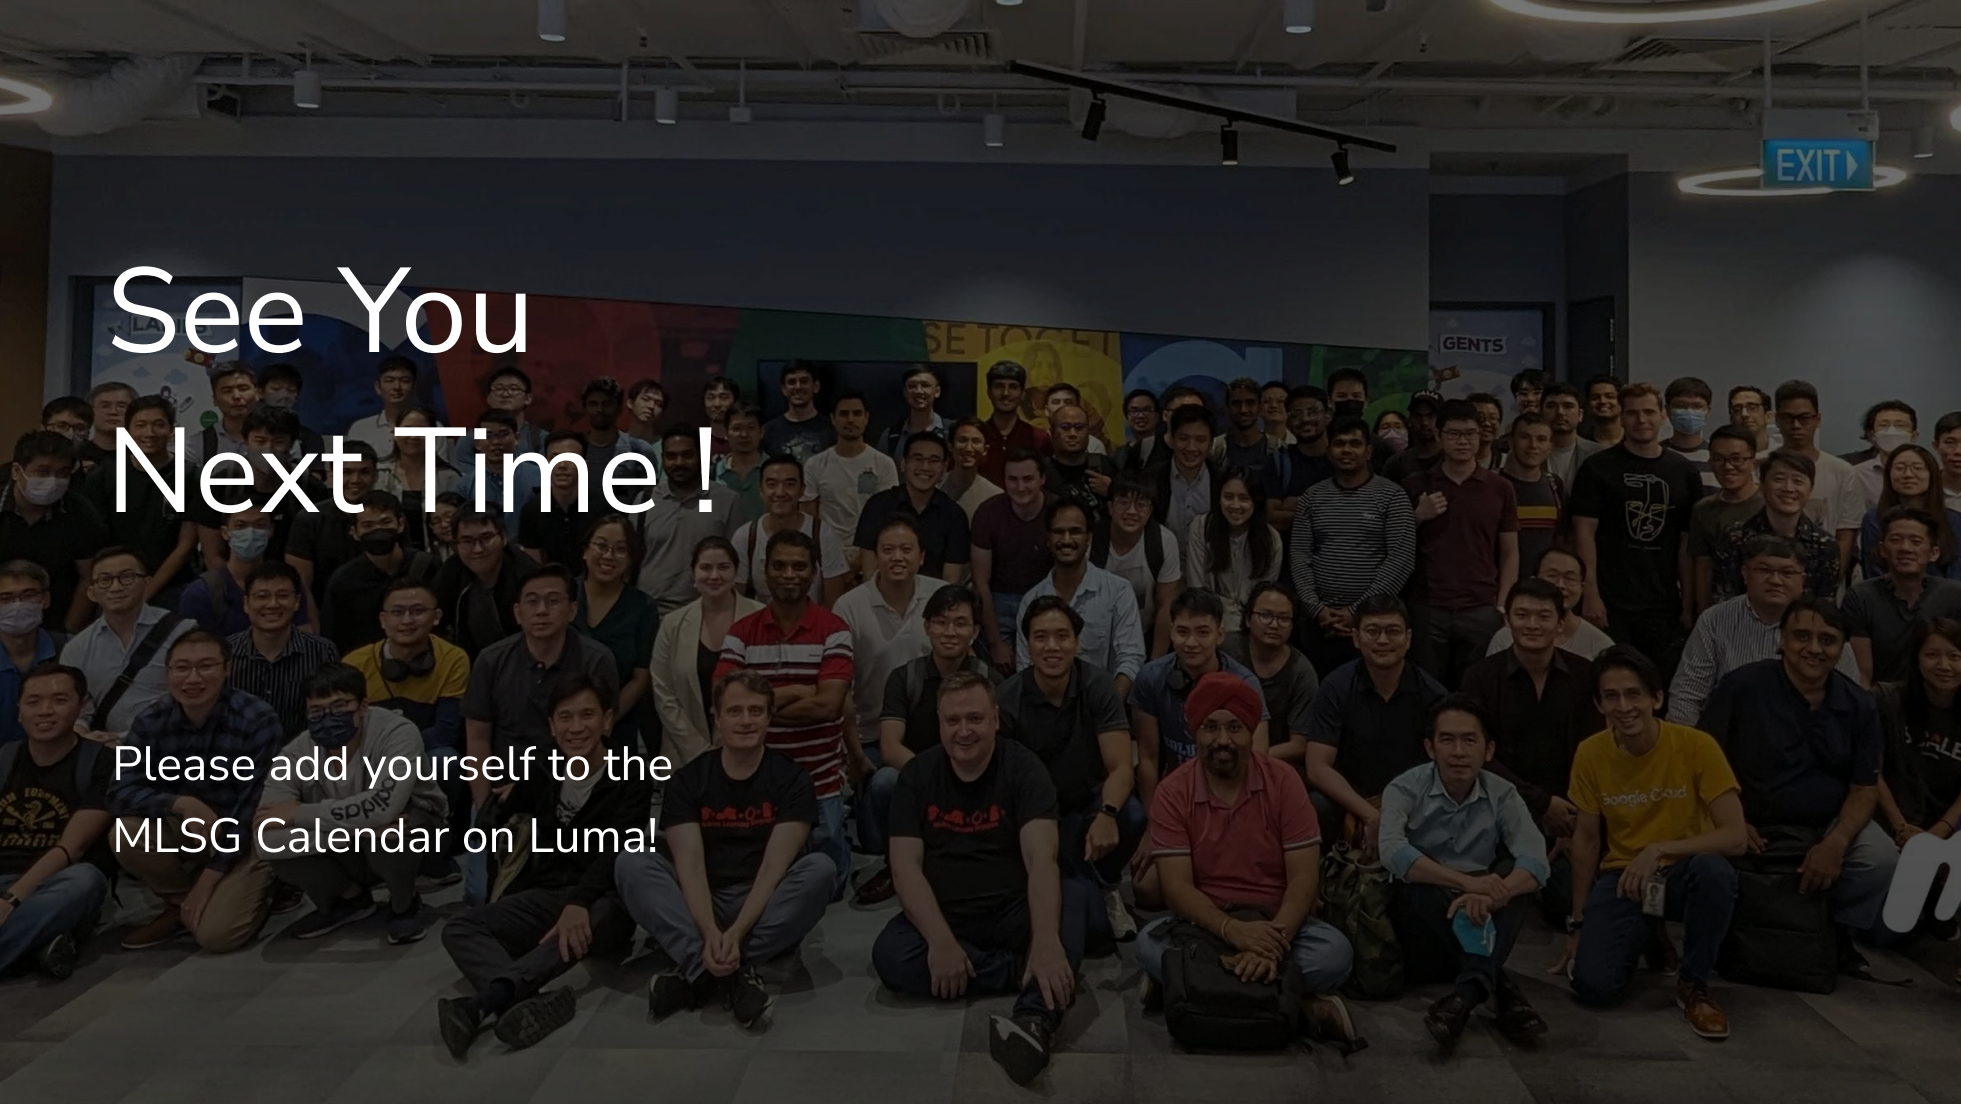
- Next Meeting = 25-June-2026
@ Google
- Topic : "TBD"
- Typical Contents :
 - Talk for people starting out
 - Something from the bleeding-edge
 - Lightning Talks
- MeetUp.com / Machine-Learning-Singapore



Link to Slides

- <https://bit.ly/4e4R7QG>





See You Next Time!

Please add yourself to the
MLSG Calendar on Luma!

