Striking the Balance

Safeguarding data privacy while empowering employees

Emily Greene
Security & Privacy Engineer, Moveworks

Large models = LOTS of data

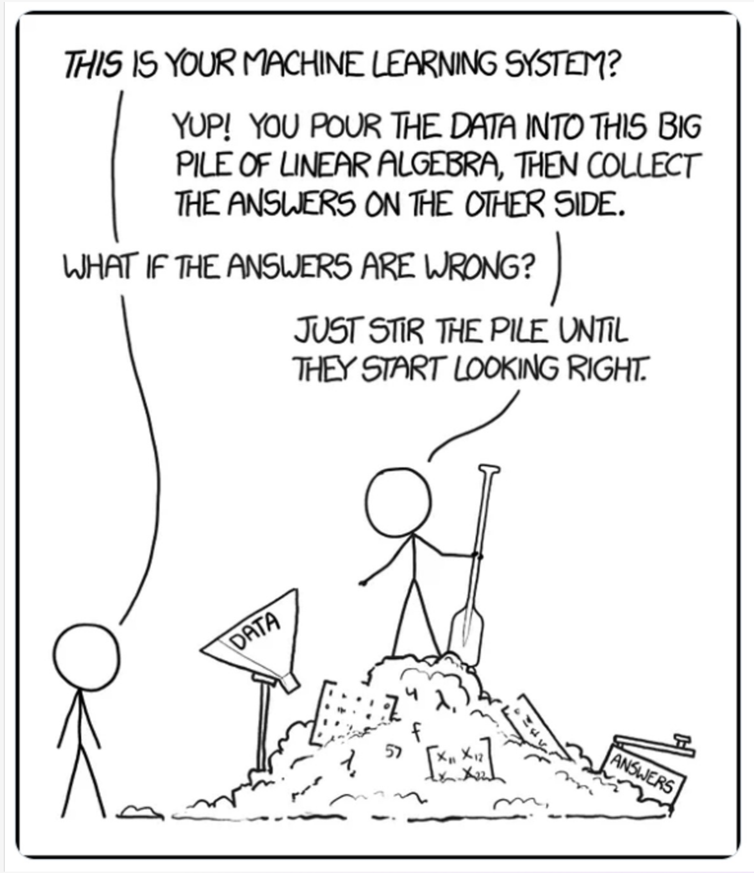
THIS IS YOUR MACHINE LEARNING SYSTEM?
YUP! YOU POUR THE DATA INTO THIS BIG PILE OF LINEAR ALGEBRA, THEN COLLECT THE ANSWERS ON THE OTHER SIDE.
WHAT IF THE ANSWERS ARE WRONG?
JUST STIR THE PILE UNTIL THEY START LOOKING RIGHT.
DATA
ANSWERS

Humans in the loop

“Need to deactivate Jon Snow’s employee access”

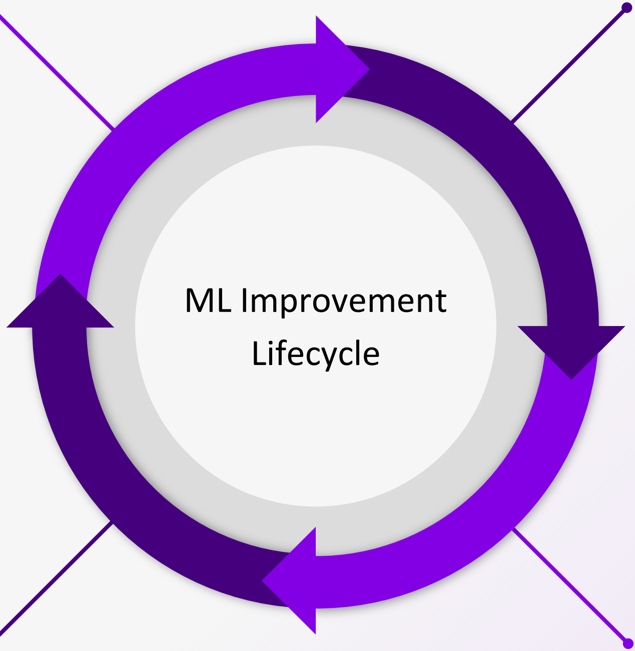

Escalation
On-call engineer

Employee deactivations aren’t working

Analytics
Data scientist

There have been 1000 failed deactivation requests in the last 3 months

Annotate
Data annotator

Need to deactivate [Action] Jon Snow [Employee Name]’s employee access

Re-train
ML engineer

New annotated data examples

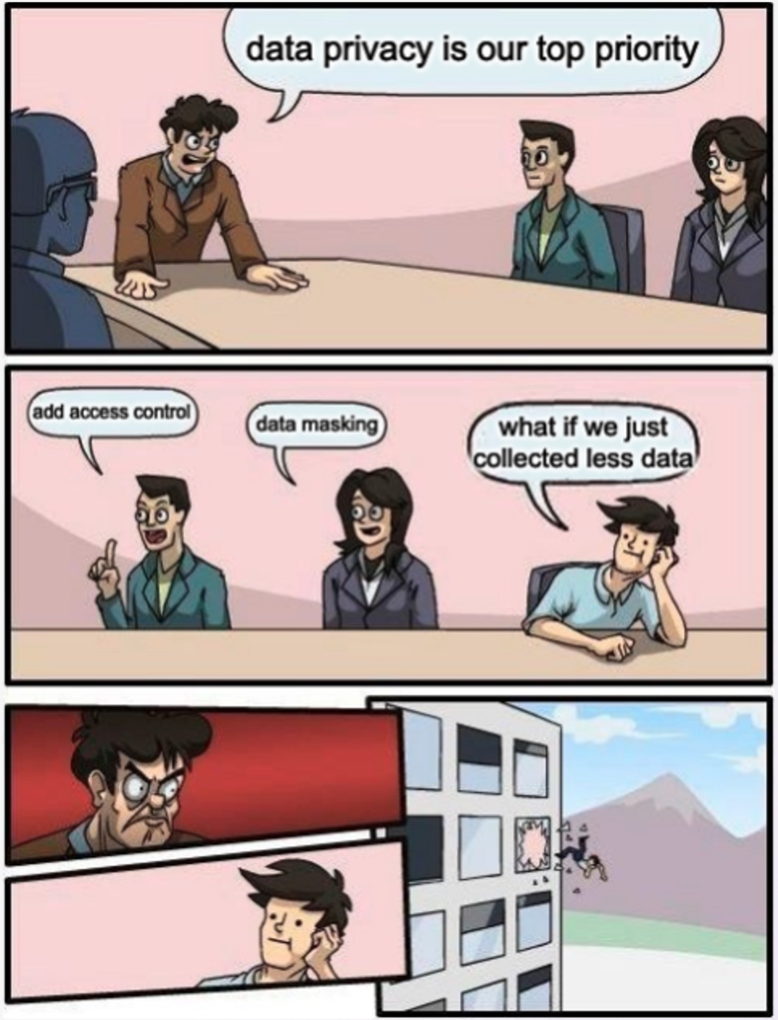
data privacy is our top priority
add access control
data masking
what if we just collected less data

Data Masking

What's the goal? Protect sensitive information from unauthorized exposure.

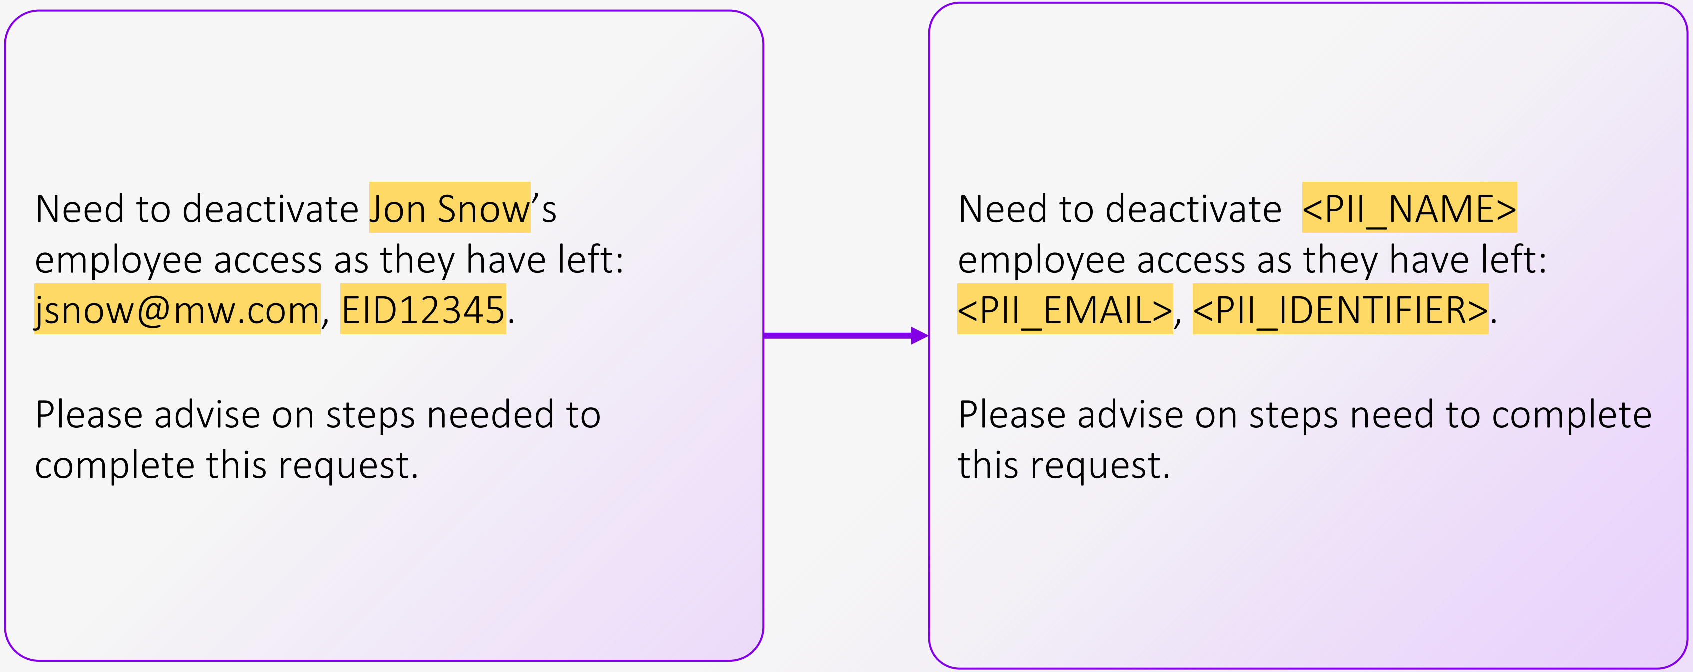

Need to deactivate Jon Snow's employee access as they have left: jsnow@mw.com, EID12345.

Please advise on steps needed to complete this request.

Need to deactivate <PII_NAME> employee access as they have left: <PII_EMAIL>, <PII_IDENTIFIER>.

Please advise on steps need to complete this request.

Access Control

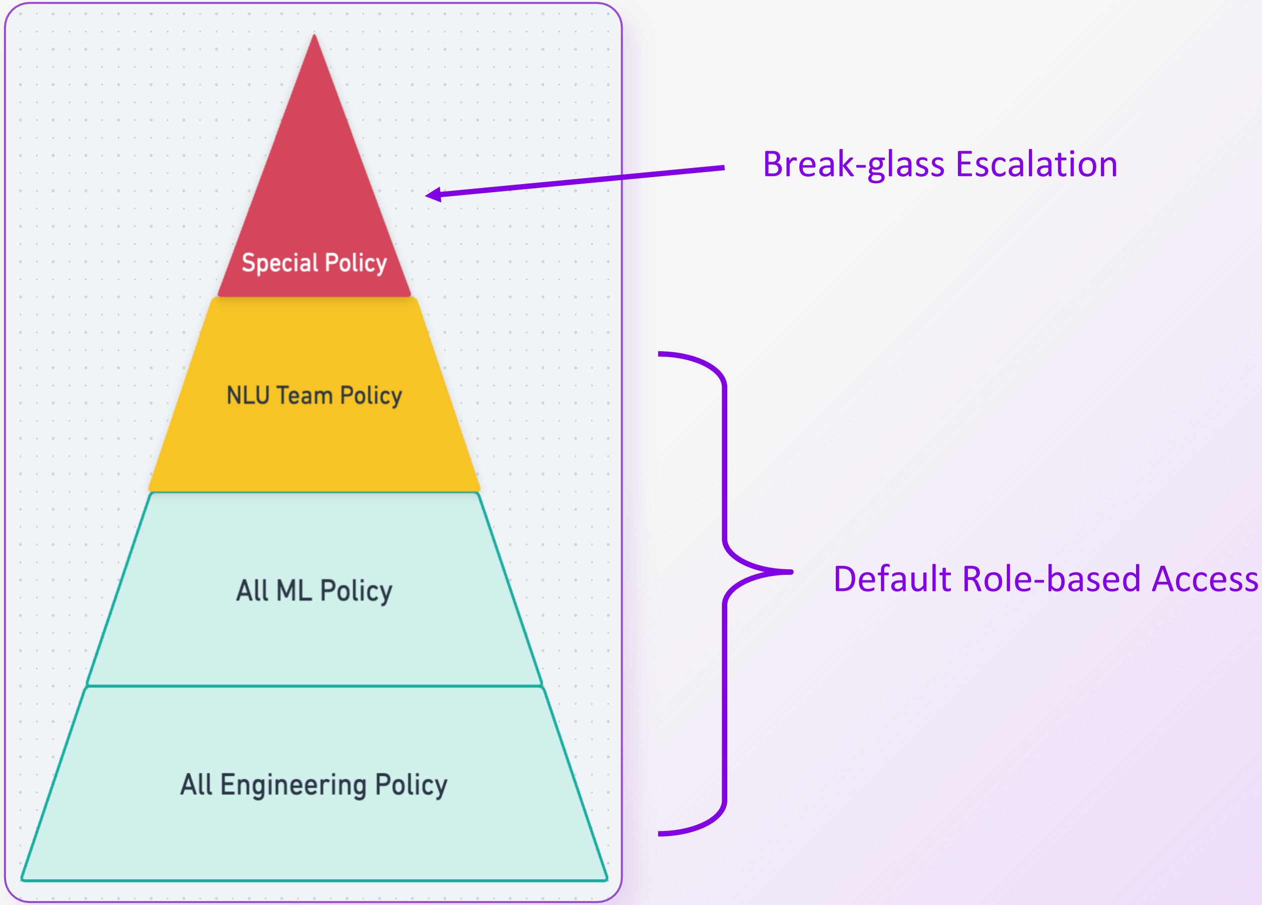

Special Policy

NLU Team Policy

All ML Policy

All Engineering Policy

Break-glass Escalation

Default Role-based Access

How do we apply these tools?

		Data Masking	Access Control	Details
1	**Escalation** Oncall Engineer	✓	✓	● Who can access the portal ● Mask utterance and metadata ● Break-glass mechanism
2	**Analytics** Data Scientist	?	✓	● Mask highly-sensitive PII ● Who can perform the analytics ● Restrict JOINs
3	**Annotate** Data Annotator	✓	✓	● Only annotate on masked data ● Break-glass mechanism
4	**Re-train** ML Engineer	?	✓	● Some models can be trained on masked data, but some can't ● Who can access the training data

Key Takeaways

What data do your employees need?

How can access to that data be limited by role?

How can your employees operate on masked data?

How can we identify sensitive data for masking?

Do we need a break-glass unmasking mechanism?

How can we provide time-bound access?

Thank you!

Emily Greene

egreene@moveworks.ai