Day 06

Wan Fang
Southern University of Science and Technology

Day 6

- 8:00 – 8:50 Lecture: Machine Learning Model Deployment
- 9:00 – 9:50 Review of designs
- 10:20 – 12:10 Exercise: Refining design diagrams

- 2:00 – 5:00 Lecture and Exercise: Design + Implementation
- 5:00 – 6:00 Review of the day

Machine Learning Model Deployment

What is Model Deployment?

- Deployment is the method by which you integrate a machine learning model into an existing production environment to make practical business decisions based on data.
- It is one of the last stages in the machine learning life cycle and can be one of the most cumbersome.

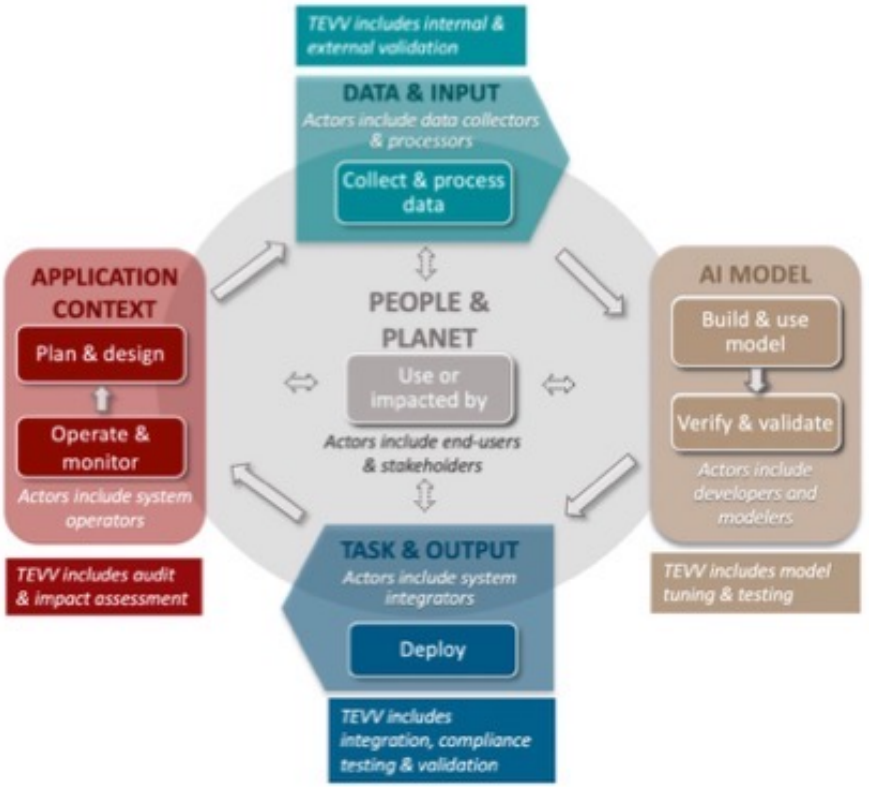

Cloud v.s. Edge

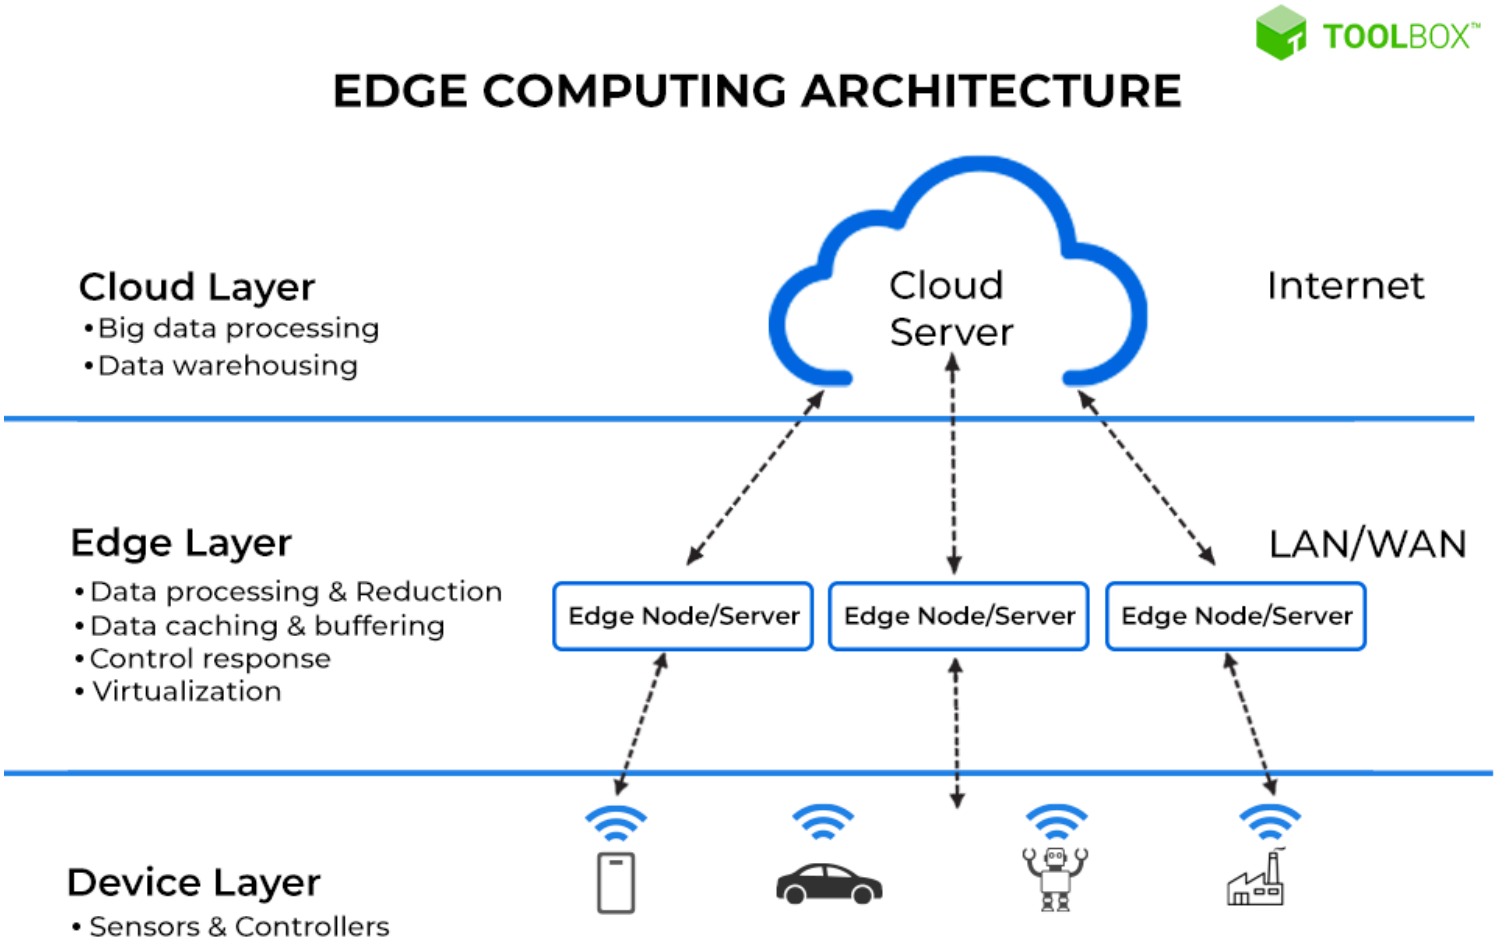

Source

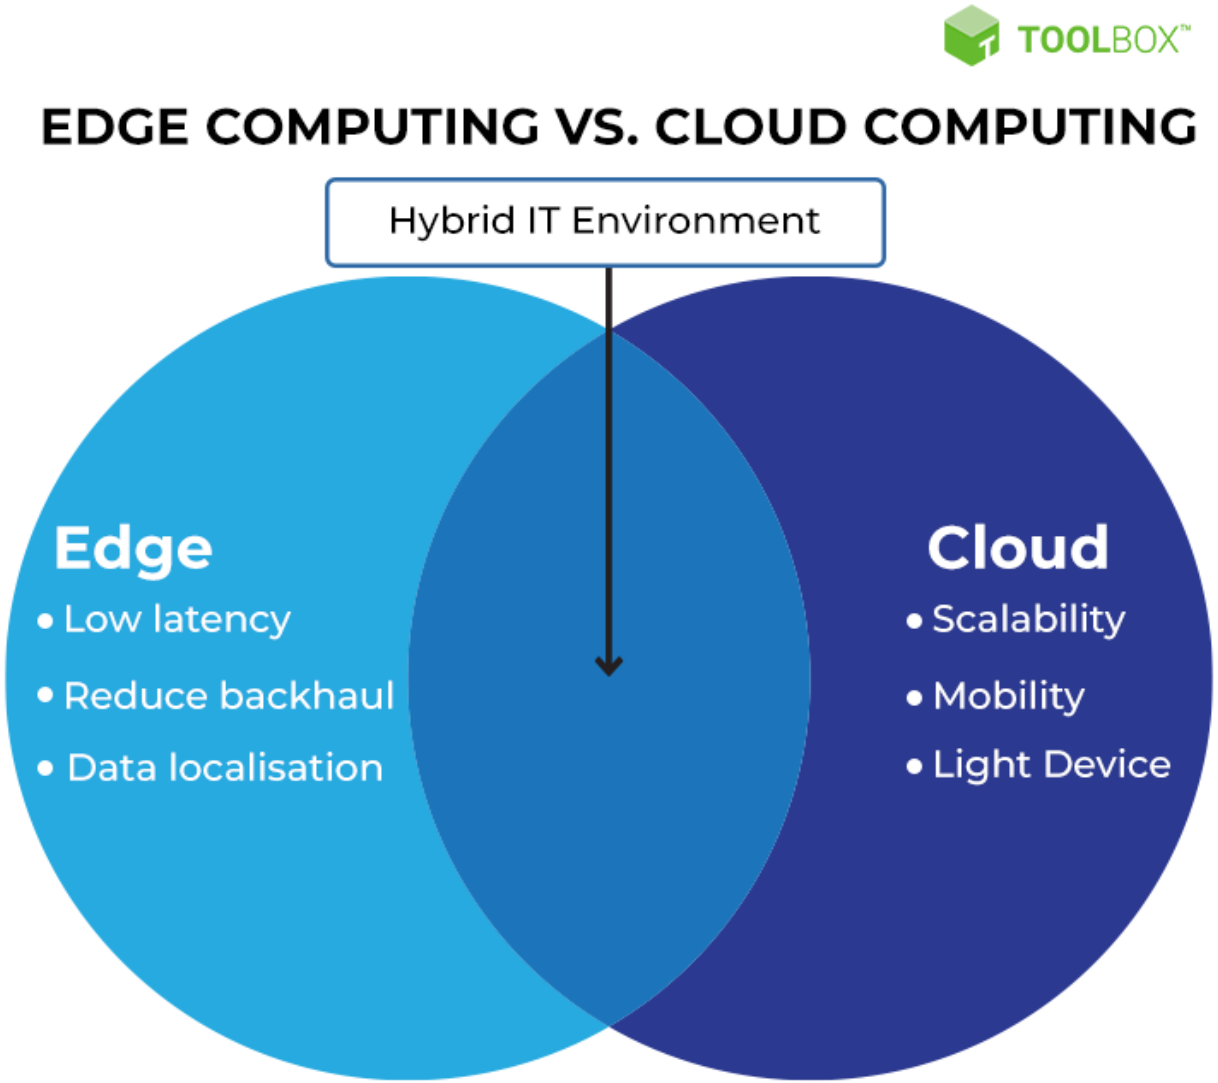
TOOLBOX™
EDGE COMPUTING VS. CLOUD COMPUTING
Hybrid IT Environment
Edge
• Low latency
• Reduce backhaul
• Data localisation
Cloud
• Scalability
• Mobility
• Light Device

Source


Top 8 Application of SaaS in Cloud Computing: Examples That Transformed World
Netflix
Dropbox
Hubspot
Zendesk
Slack
Canva
Zoom
Shopify
www.cloudconc.com

Examples of Edge Computing

Grid Edge Control and Analytics

Edge grid computing technologies are enabling utilities with advanced real-time monitoring and analytics capabilities.

Oil and Gas Remote Monitoring

Computing resources at the edge allow data to be analyzed, processed and delivered to end-users in real-time.

Using edge computing the gigabytes of sensory and special data is analyzed, filtered and compressed before being transmitted on IoT edge Gateways to several systems for further use.

Traffic Management

Autonomous Vehicles

Edge and distributed computing techniques increase safety, spatial awareness and interoperability with current-generation hardware.

Source

Deploy with Cloud and IoT Edge

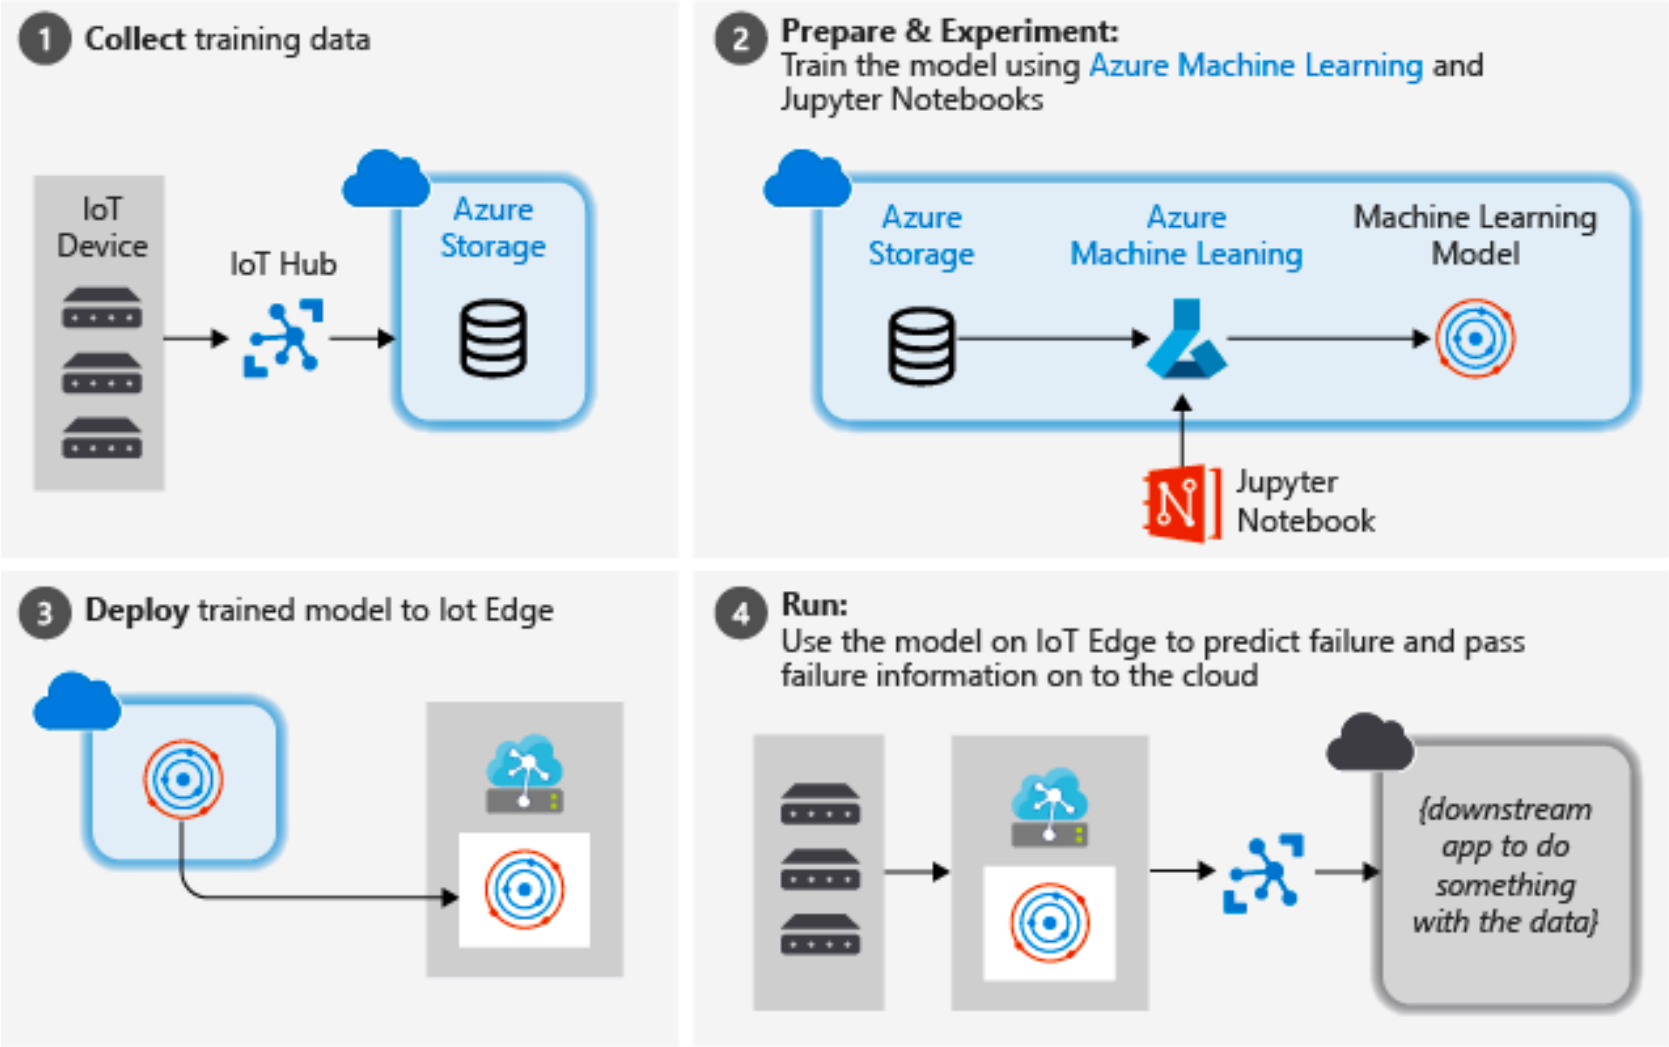

Source

Deploy with Cloud and IoT Edge

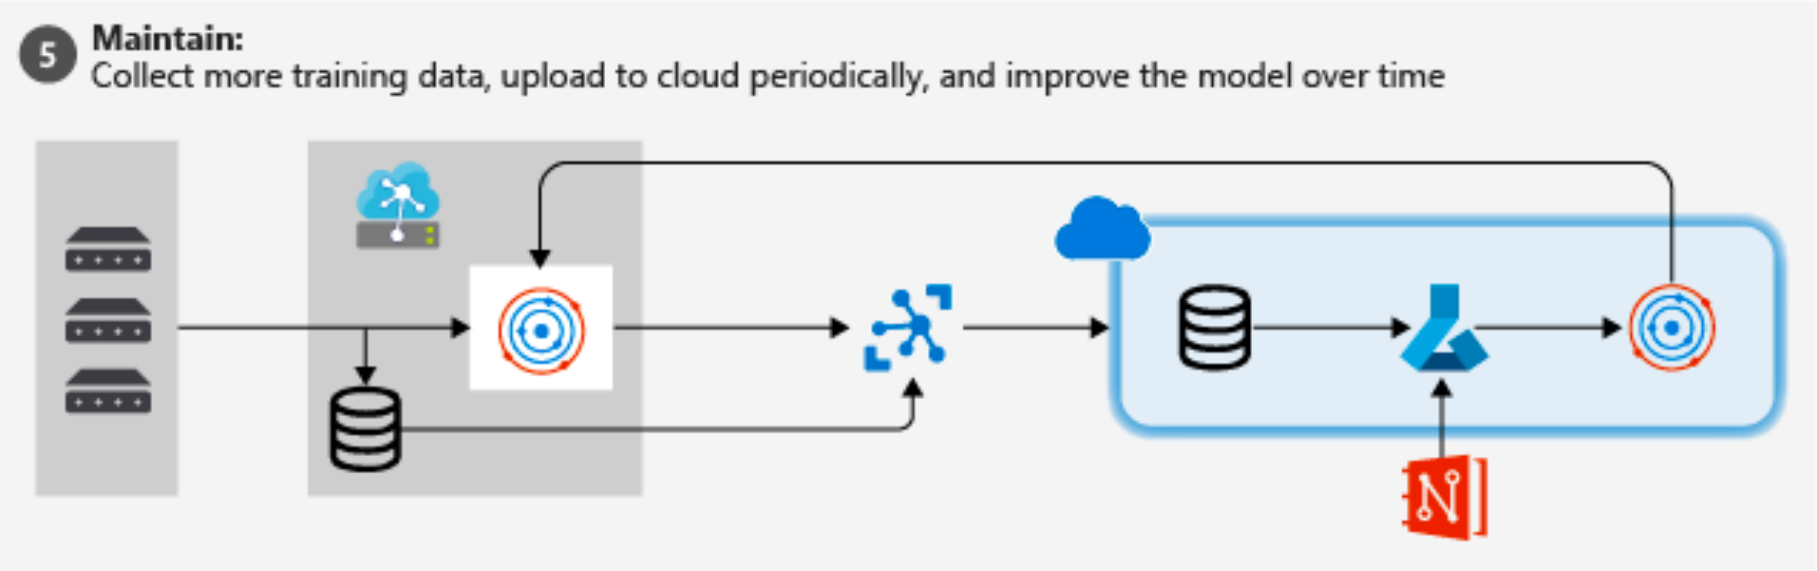

Source

Edge Computing Devices

- https://www.waveshare.net/
- https://www.nvidia.com/en-us/autonomous-machines/embedded-systems/

Lecture
AI Meets Design

Wan Fang
Southern University of Science and Technology

Activity:
Plotting your model for concept development

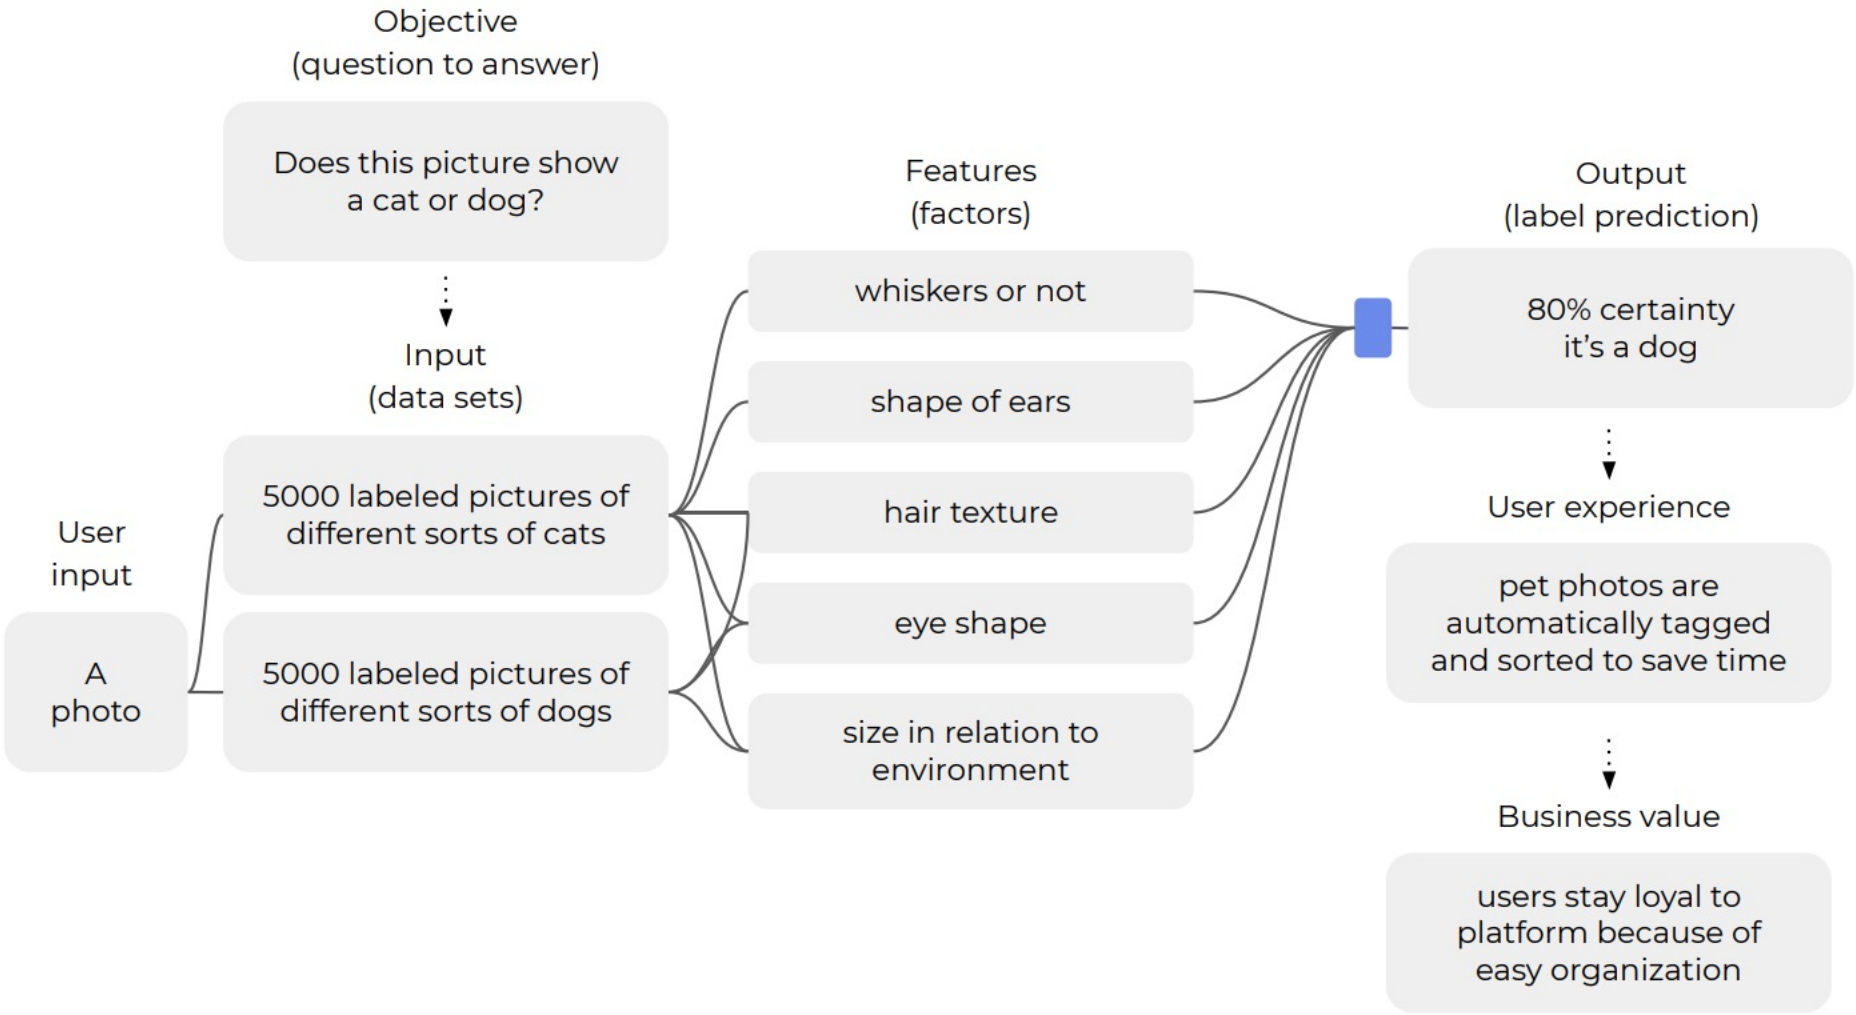

Worksheet:
Plotting your model

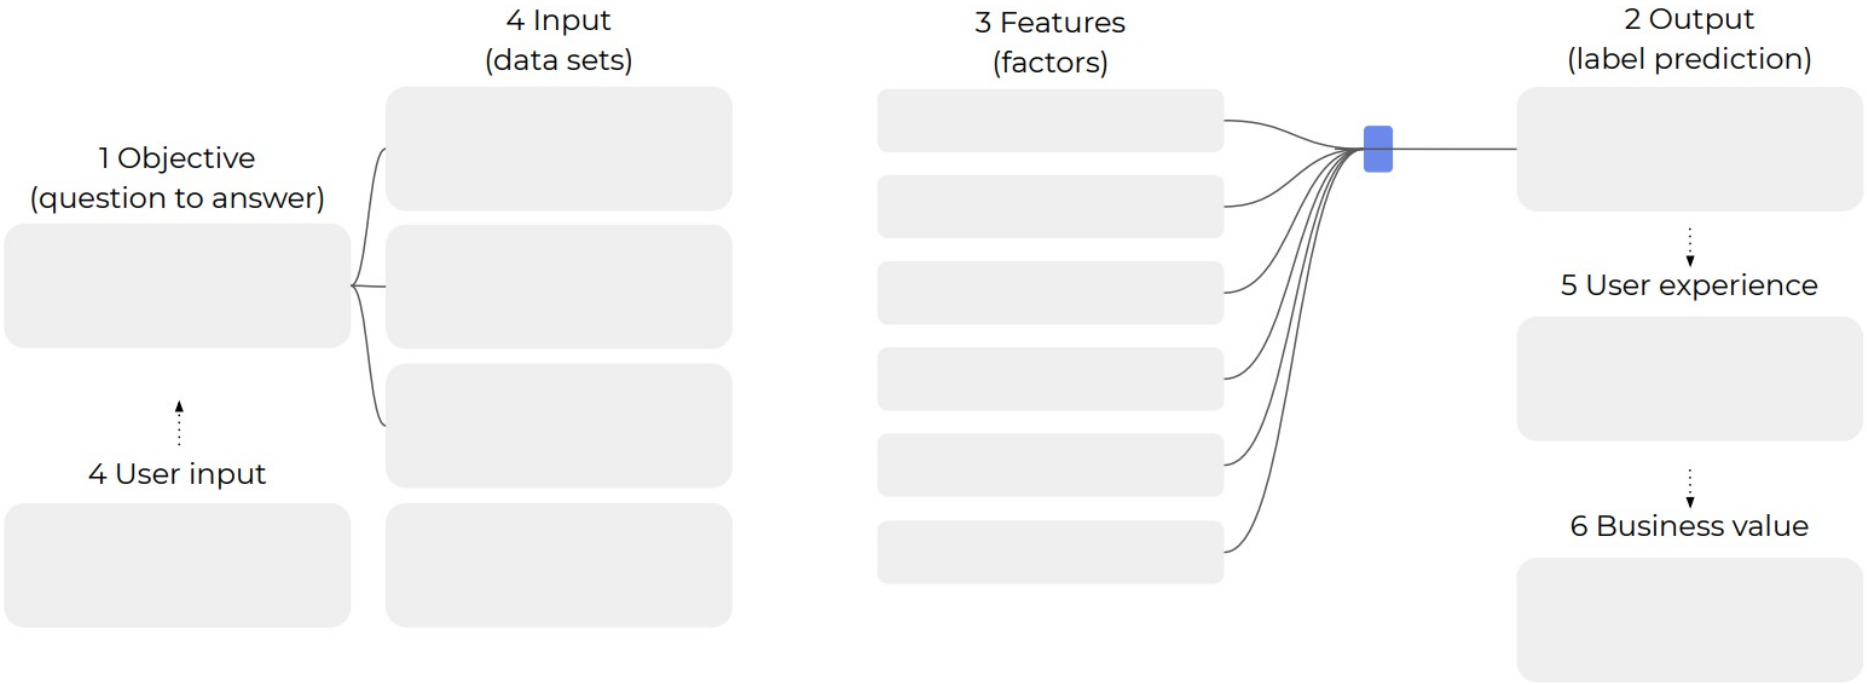

1 Objective
What is the question we're trying to answer and asking the machine?

2 Output
How is the machine's answer presented and interpreted?

Formulate your output as a probability.

3 Features
What data points do you need or are important factors in answering the question?

Do you know which features go into the answer? Think about the variables and patterns humans look at when performing this task or answering this question.

4 Input
Which data sets does that data reside in? What data will the model be trained on? What data does the user input?

Do you have this data to input? If not, how do you acquire it?

+ Connect
Draw connections between the assumed features and data sets they reside in.

5 User experience
How does the outcome get presented to and help the user?

6 Business value
How does the solution return value to the organization?

Design
+ Implementation

You've developed the concept, validated with your users,
and are ready to start building and bringing your idea into the world.

How to build it? How do you translate user needs to algorithmic parameters? What considerations do you need to make during design and implementation? How will you align needs across departments and stakeholders while keeping your user at the core?

In this chapter you will find:

Defining success and failure *
to understand confusion matrix and cost of falses

Mapping user needs to models
to translate subjective user needs into model trade-offs

Metrics to evaluate by *
to agree on when your model will be good enough

UX and design challenges of AI *
to learn about the unique design challenges of the material

Capturing design tensions *
to explore the complexity of designing for human values

Consequence wheel *
to anticipate (un)intended consequences and impact

Activity:
Defining success and accounting for failure

To define what success looks like, it can help to imagine what failure looks like. Understanding cost of false, we can make conscious decisions about what to optimize for and what to trade-off.

1

Confusion matrix

The confusion matrix (right) helps you map the impact of correct and false predictions.

True & false - The impact of misdiagnosing illness is much greater than the impact of misdiagnosing whether I like a song.

Positive & negative - In addition, the impact of diagnosing me with illness when I'm healthy, or failing to diagnose me when I am sick, is also different.

		Machine prediction: Positive	Machine prediction: Negative
User reality	Positive	:) True positive	:(False negative
User reality	Negative	:(False positive	:) True negative

Think through and discuss each of the 4 states and how they affect the your user and stakeholders. Re-clarify your objective. Which of the 4 states is most and least desired?

2

Precision vs recall trade-off

Building on that understanding, we can inform our position in common trade-off known as precision vs recall. While your engineers will work on both, at some point they'll need to make a decision on which to optimize for. In your use case, what is worse? Making a false prediction (optimize for precision), or missing true prediction (optimize for recall)?

How important is ..

Precision
% of predictions that are spot-on

VS

Recall
% of objects that are predicted

Worksheet:
Confusion matrix

Machine prediction

User reality	Positive	Negative
Positive	:) What does a **true positive** look like?	:(What does a **false negative** look like?
Negative	:(What does a **false positive** look like?	:) What does a **true negative** look like?

The confusion matrix is a diagram to help map the impact of correct and false predictions.

True & false
The impact of misdiagnosing illness is much greater than the impact of misdiagnosing whether I like a song.

Positive & negative
In addition, the impact of diagnosing me with illness when I'm healthy, or failing to diagnose me when I am sick, is also different.

Think through and discuss each of the 4 states and how they affect the your user and stakeholders. Re-clarify your objective. Which of the 4 states is most and least desired?

How important is ..

Accuracy
% of predictions are correct

VS

Transparency
ability to trace back why/how

1

Reward & penalize
In the same way we train pets, you can reward or penalize the model. Think about what behavior to reward or penalize to direct your model towards ideal behavior. It can help to imagine you're onboarding a person for this task. How would you respond to them in any outcome so they improve for next time?

2

Machine teaching
Your user can play an active role in teaching the machine what is desired behavior through explicit and implicit feedback.

This way you can leverage what AI practitioners call "the virtuous cycle of data" (on the right) where more data means a better product, a better product means more users, which in turn means more data. Think about how you might set up these feedback loops in your interface.

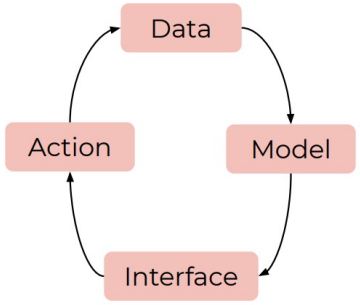

3

Accuracy vs transparency trade-off
Most models suffer from a trade-off between accuracy and transparency. How important is it for your user and your organization to understand why a certain prediction was made?

Activity:
Mapping user needs to models for machine teaching strategies

Optimizing for one thing always implies letting go of another. Once we have a clear objective in mind, we must train the model to understand and pursue such behavior.

By default, engineers will strive for maximum accuracy. In some cases, a lower score suffices at fulfilling a need, or the UX can make up for a range of errors. As designers, we can help define when models are or aren't 'good enough' to provide value to our user.

Activity:

Metrics to evaluate your model by

1

Benchmarking

Discuss and place on the spectrum between 1% and 99%:

100% accuracy

Current human benchmark based on how people currently perform the task (optional: expert & layman)

Baseline model based on research and industry standard

Minimum confidence level based on cost of falses. How confident does the model need to be before presenting an answer? At what point is making the wrong prediction more harmful than not making one at all?

0% accuracy

Minimum benchmark to provide value to user based on your use case and user research

Targeted benchmark based on the above and cost of falses. At what point is the prediction accurate enough to provide consistent value without risking user trust?

2

Evaluating

You can use these metrics to track the performance of your model on training and new data - and know when you're ready to deploy your first version for users.

Worksheet:
Benchmarking

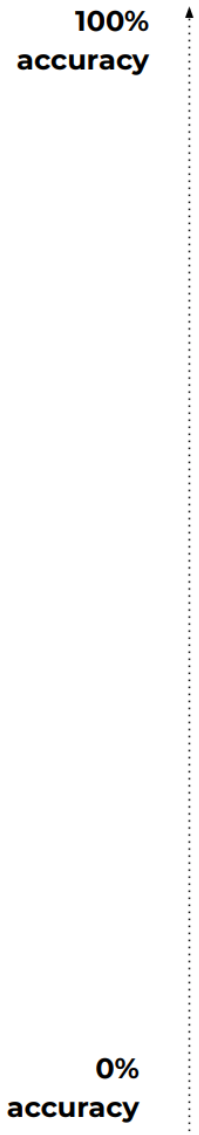

Discuss and place on the spectrum between 1% and 99%:

Current human benchmark based on how people currently perform the task (optional: expert & layman)

Baseline model based on research and industry standard

Minimum confidence level based on cost of falses. How confident does the model need to be before presenting an answer? At what point is making the wrong prediction more harmful than not making one at all?

Minimum benchmark to provide value to user based on your use case and user research

Targeted benchmark based on the above and cost of falses.

At what point is the prediction correct enough to provide value? How confident does the model need to present an answer? At what point is making the wrong prediction more harmful than not making one at all?

Below, we identify and introduce you to 9 challenges in designing intelligent, adaptive, and (semi-)autonomous systems:

Each design material comes with unique challenges. In the same way that designing a poster is different from designing a mobile app, designing AI-driven applications is different from designing apps.

User Trust & Transparency

1. **Explainability**
Making sense of the machine and communicating to the user why the system acts the way it does

2. **Managing expectations**
Helping the user understand what the system can and can not do (over time) by being transparent about abilities and limitations and building helpful mental models of it

3. **Graceful failure & accountability**
Assume failure and design graceful recoveries. Take accountability for mistakes and minimize cost of falses for your user

User Autonomy & Controls

4. **Machine teaching & user feedback**
Allowing the user to teach the machine with implicit and explicit feedback loops and collecting direct data input

5. **User controls & customization**
Giving users the controls to customize the system/algorithm to their needs and intervene with the course of a model if needed

6. **Data privacy & security**
Collect, handle, and store user data with care. Be transparent about who can access what data and why while acknowledging their ownership

Value Alignment

7. **Computational virtue**
Translating subjective human needs, values, and experiences into algorithmic parameters the model can optimize for

8. **Bias & inclusivity**
Mitigating bias and guarding inclusivity in data and models to ensure fair treatment for all

9. **Ethics & (un)intended consequences**
Unprecedented scale, speed and complexity call for a new level of thoughtfulness and responsibility in anticipating impact and (un)intended consequences

Resource:
UX and design challenges of AI as a design material

Worksheet: UX of AI challenges

1. Explainability - How will we help our user understand certain outcomes?	2. Managing expectations - How will we establish realistic expectations?	3. Graceful failure & accountability - How will we design for trust in case of failure?	**User Trust & Transparency**
4. User feedback - How will your user provide feedback to the system?	5. User autonomy - How will the user be able to customize their experience?	6. Data privacy & security - How will you collect, store, and handle data?	**User Autonomy & Control**
7. Computational translation - How will you turn needs into parameters?	8. Bias & inclusivity - How will you prevent bias and guard inclusivity?	9. Ethics & (un)intended consequences - How will you look out for negative and positive impact?	**Value Alignment**
10. Which other (design) challenges do you foresee?			

Design ethics, as general ethics, aren't as simple as wrong and right. More often, we are designing for a tension between different values.

Making these polarities explicit can help us design for them in a constructive way: deciding where on the spectrum we want our product or service to reside, on what to foster, and look out for.

Tool:
Capturing design tensions and value polarities

1. Name and write the polarizing values on the top and bottom. Iterate on the terms as you go to capture their essence.

2. For both values, think about all the positives and negatives. Write them into the boxes or on post-its first. Questions to get started:

   What opportunities and benefits are present? What challenges and threats? Why would users want this? Why not? What human values are at play? Human rights even?

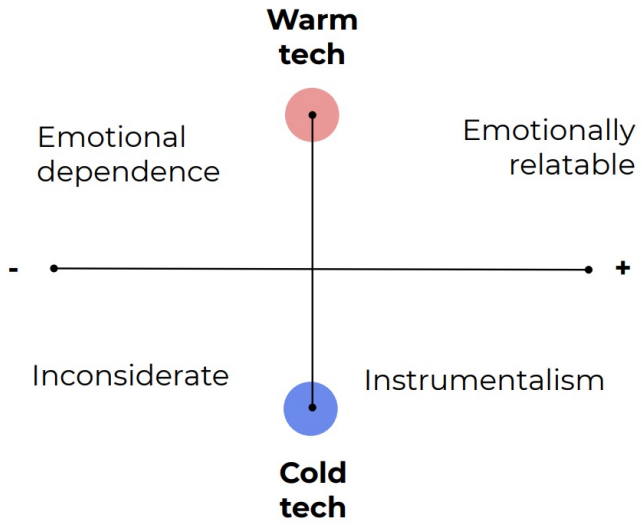

how might we relate to tech emotionally in a healthy way?

3. Give each of the dimensions a term to summarize its sentiment.

4. Considering positive and negative aspects, draw an x where you think your product or service should reside. If you're working in a team, let everyone do this individually and then discuss and agree on differences.

   Repeat making polarity maps as many times as needed and allowed depending on your context. Any product would have at least 3-5 polarities.

5. Think about how to harness the positive and how to limit the negative aspects. Keep this mental model front and centre as you continue to design, develop, and deploy your AI-driven application.

Worksheet:
Value polarities

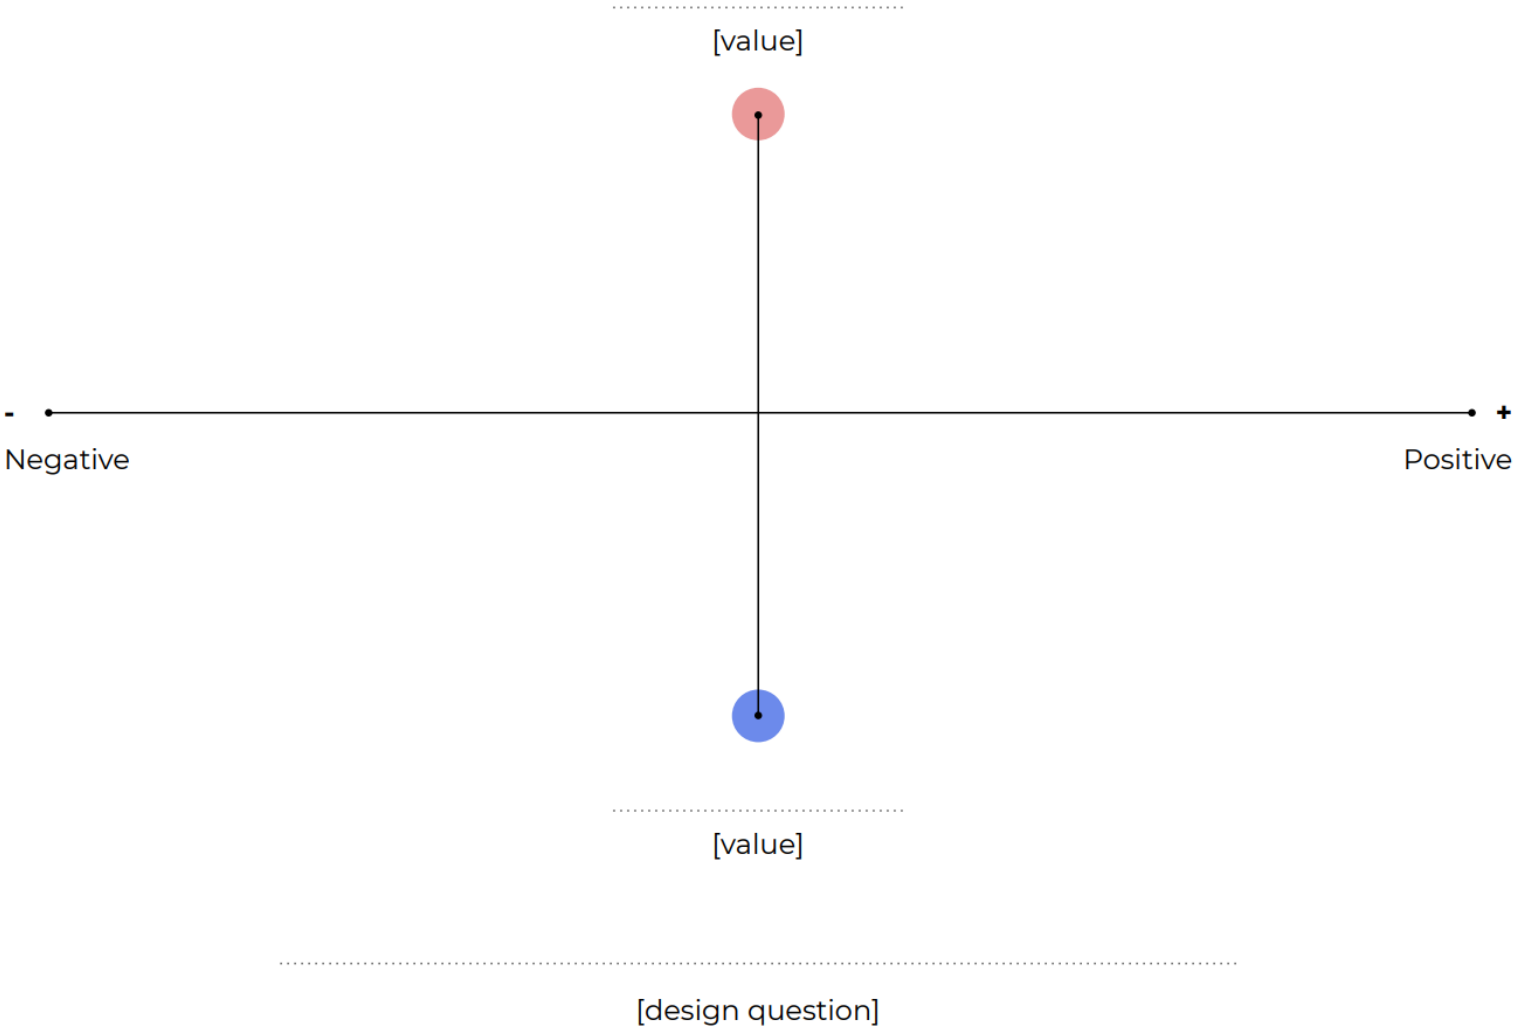

Name and write the polarizing values on the top and bottom. Iterate on the terms as you go to capture their essence.

For both values, think about all the positives and negatives. Write them into the boxes or on post-its first.

Questions to get started: What opportunities and benefits are present? What challenges and threats? Why would users want this? Why not? What human values are at play? Human rights even?

Give each of the dimensions a term to summarize its sentiment.

Considering positive and negative aspects, draw an x where you think your product or service should reside. If you're working in a team, let everyone do this individually and then discuss and agree on differences.

Repeat making polarity maps as many times as needed and allowed depending on your context. Any product would have at least 3-5 polarities.

Think about how to harness the positive and how to limit the negative aspects. Keep this mental model front and centre as you continue to design, develop, and deploy your AI-driven application.

Guiding questions:

- What consequences & impact might this have? Think of culture, politics, economics, etc.
- What's the best-case scenario? Worst case? Who benefits? Who suffers?
- Is this a world you'd like to live in? Why or why not?
- Which industry, services, or social rituals might be disrupted by this?
- Which new user pain points or opportunities exist in this world?

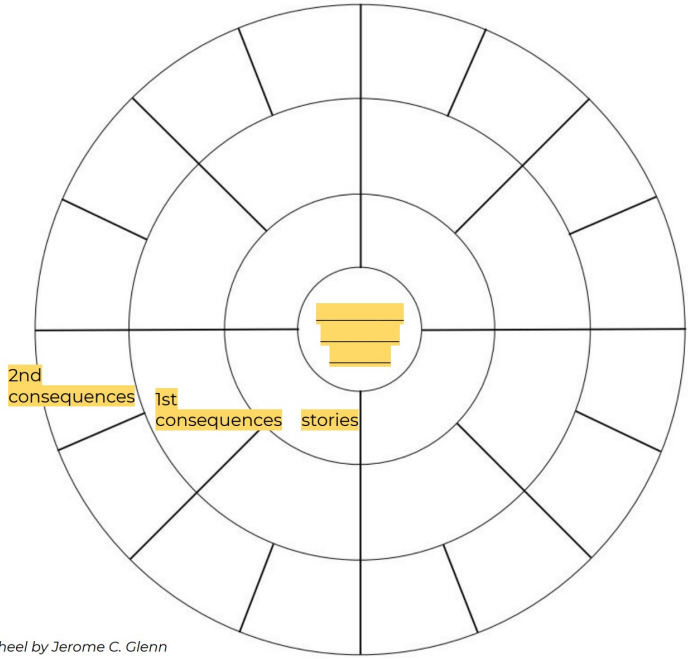

Futures Wheel by Jerome C. Glenn

1. Start with your product, service, interaction, or AI-driven application in the middle.
2. Write a handful of user stories and experiences in the second row.
3. Begin thinking through consequences by asking "if this then what?" as you work outward. Consider different groups, scenarios, and lenses. Do this for at least 3 and ideally more levels.

Tool:
Consequence wheel
to anticipate (unintended) consequences

Worksheet:
Consequence wheel

Start with your product, service, interaction, or AI-driven application in the middle.

Write a handful of user stories and experiences in the second row.

Begin thinking through consequences by asking “if this then what?” as you work outward. Consider different groups, scenarios, and lenses. Do this for at least 3 and ideally more levels.

Guiding questions:

- What consequences and impact might this have? Think of culture, politics, economics, ecology, etc.
- What’s the best case scenario? Worst case? Who benefits? Who suffers?
- Is this a world you’d like to live in? Why or why not?
- Which industry, services, or social rituals might be disrupted by this?
- Which new user pain points or opportunities exist in this world?

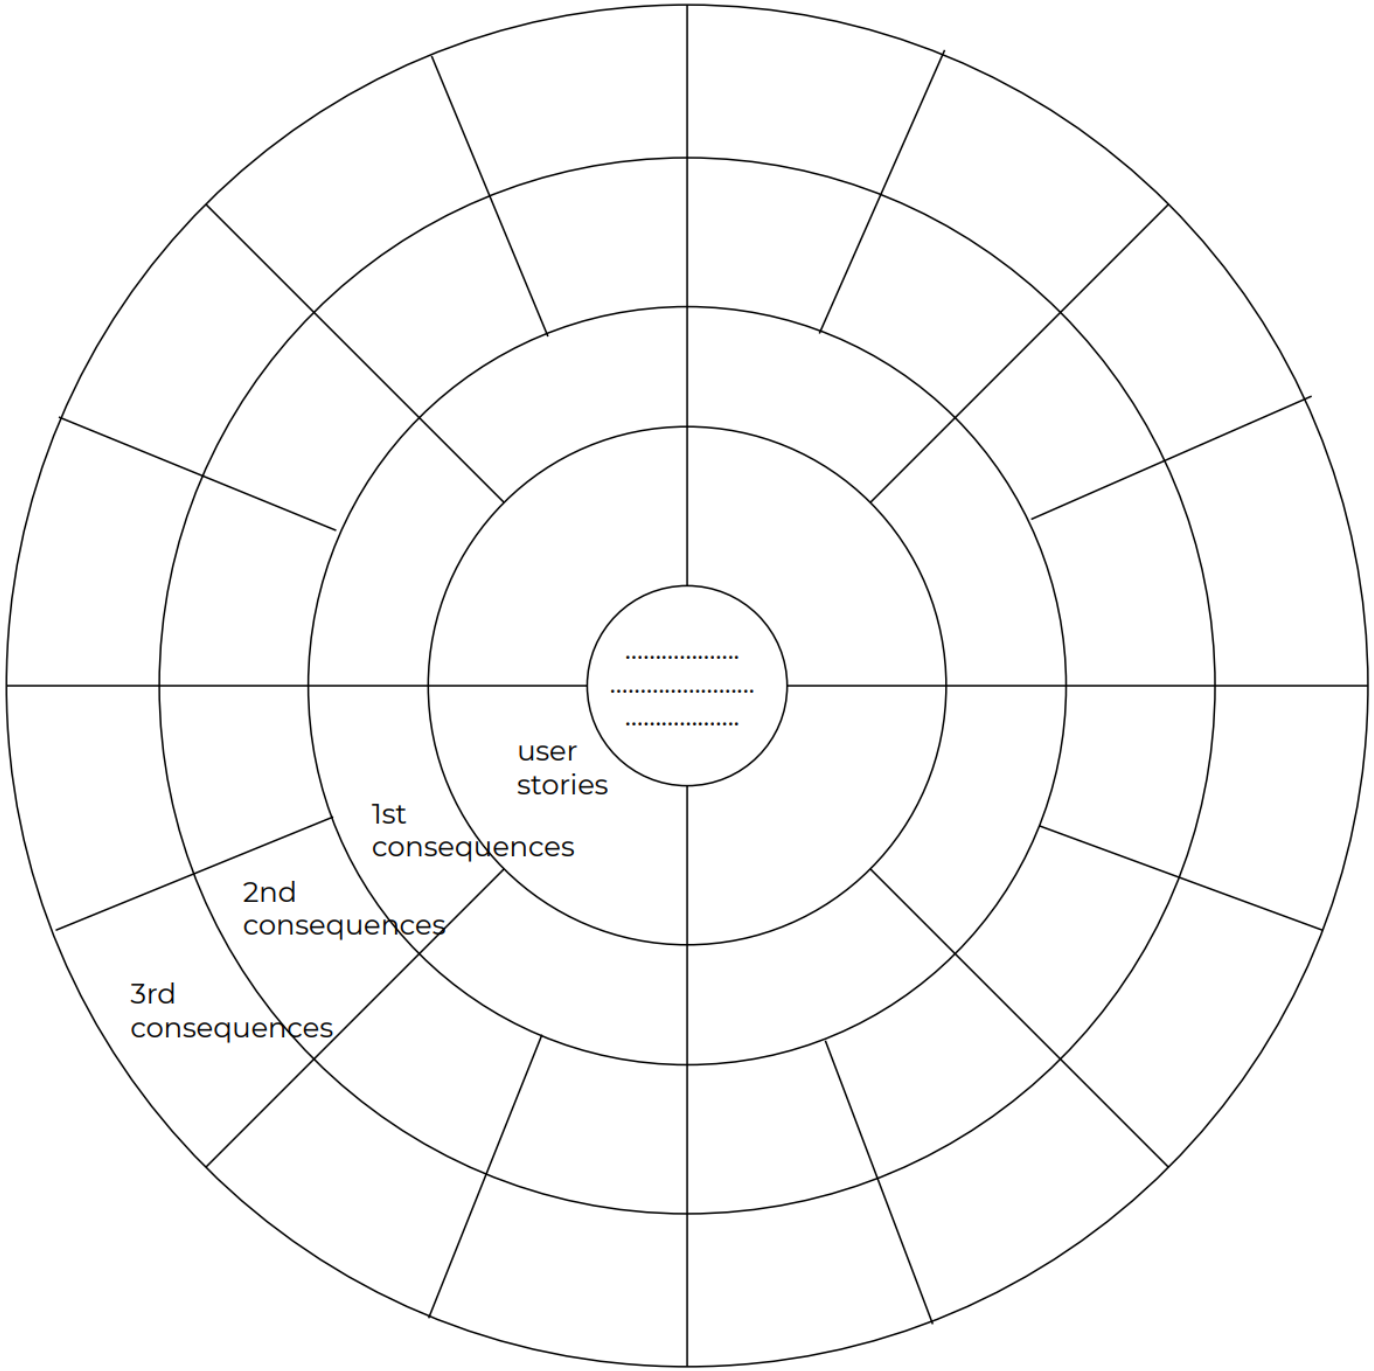

Additional resources

If this toolkit has succeeded at peeking your interest and you're hungry for more, these are additional resources to continue your learning journey.

To learn

If you want learn more about AI, here are some of my favorite learning resources:

- Elementsofai.com
- People + AI Guidebook by Google
- AI for Everyone on Coursera with Andrew Ng
- AI-driven design e-book series by Adyen & AWWWARDS
- Algorithms.design
- Machine Learning for Designers by Patrick Hebron for O'REILLY

To apply

If you want to explore more design tools, check out these projects:

- Intelligence Augmentation Design
- Toolkit by Futurice
- AI Ethics cards by IDEO
- Machine Ethics Toolkit
- IBM AI Camp DIY Guide

To play

If you're feeling a little overwhelmed, here are a few of my favorite AI games and experiments you can play with:

- Teachable Machine by Google
- Emoji Scavenger Hunt by Google
- Quick, Draw! by Google
- Runway ML

To build

Plenty MLaaS (Machine Learning as a Service), and cloud platform solutions have become available to cater to different goals and levels of technical expertise.

Type of dev	Pre-trained models on their data	Train models on own data	GUI for custom models	Libraries for custom models
Tools	APIs & SDKs. Mostly by Amazon, Microsoft, Google, IBM. Check rapidapi.com	AutoML by Google, Microsoft Azure ML Studio, Amazon ML	Amazon Sagemaker, IBM Watson ML Studio	Tensorflow, sci-kit, pytorch, Keras

Additional resources

If this toolkit has succeeded at peeking your interest and you're hungry for more, these are additional resources to continue your learning journey.

Day 02
AI Meets Design II

Thank you~

Wan Fang
Southern University of Science and Technology