Bayesian Decision Theory Tutorial: Applying Bayesian Statistics to Data

Adam Ahmed

April 18, 2024

Sponsors:
Dr. Sandra Hobson, DOT&E, Deputy Director, SIPET
Dr. Jeremy Werner, DOT&E, SIPET, Chief Scientist

COR:
Mr. Chris Dodson, DOT&E JT&E

Setup

- You are part of the T&E Working Integrated Product Team to test the Mock-Cannon (based on the 155 mm Howitzer).
- You are working alongside Bayesian trained SMEs brought on by the Product Manager to assist the Test Manager on inference for OT&E on this system.
- You and Bayesian SMEs have assisted in formulating the Detailed Test Plan (DTP).
- Given the requirements as defined in the Testing & Evaluation Master Plan (TEMP), evaluate the viability of the Mock-Cannon as a potential capability.

System Under Test: Mock-Cannon

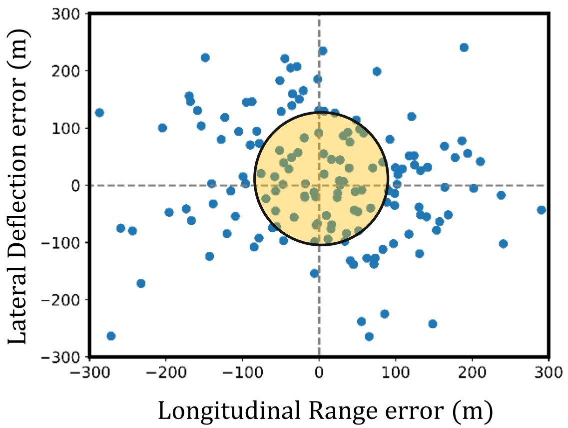

- TEMP defines an accuracy metric defined as a circular error probable (CEP):
  - Cover proportion, $p \in [0,1]$ of any shot group to fall within a predefined radius
  - CEPZ corresponds to $p = Z/100$ with $Z \in [0,100]$
- We will use a condition that R_{CEP50} = 100 m, i.e., 50% of all shots fall within a radius of 100 m.
- This metric determines if the Mock-Cannon is/is not an acceptable system.

https://en.wikipedia.org/wiki/M119_howitzer

Tutorial outline

Part 1 (Applying Bayes to the Mock-Cannon):

"Given that you've conducted a test/experiment, how do I apply Bayes Theorem?"

Model→ Likelihood→ Prior→ Hyperparameters→ Posterior

Part 2 (Bayesian Decision Theory):

"Should I conduct a test?"

State → Action→ Outcome→ Utility→ Projection

Part I:

Applying Bayes to the Mock-Cannon

Setup of the Mock-Cannon

- Munition is based off the M107(HE) 155 mm artillery projectile.
- Simulation is for *unguided* munitions.

- Independent variables:
  - Range: distance from gun-to-impact point
  - Quadrant Elevation: firing angle w.r.t. the i - j plane.
- Dependent variable is point of impact (POI) error

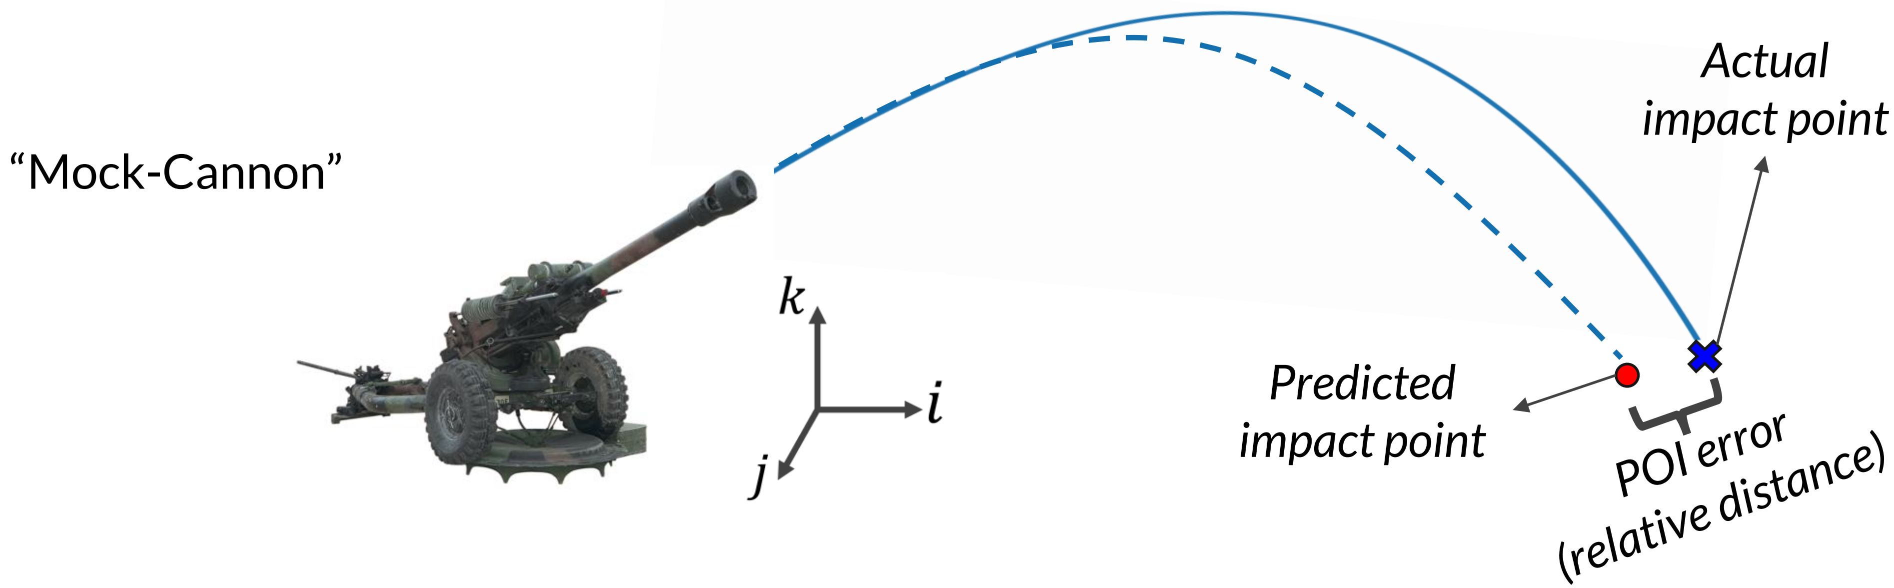

Contributions to projectile motion

- Main forces and torques
  - $\vec{F}_g$: gravitational force; attraction between earth and munition
  - $\vec{F}_d$: drag force; resistive force of an object travelling through a fluid
  - $\vec{F}_l$: lift force; responsible for lateral drift
  - Overturning moment: associated with lift force
  - Spin damping moment: opposes spin of projectile due to aerodynamic skin friction
- Additional contributions:
  - Coriolis Force
  - Magnus Force
  - Mach number dependent aerodynamic coefficients
  - Wind velocity field
  - Pressure
  - Temperature
  - Air density
  - Etc…

$\left.\begin{array}{l}\text{Wind velocity field}\\ \text{Pressure}\\ \text{Temperature}\\ \text{Air density}\end{array}\right\}$ Environmental conditions

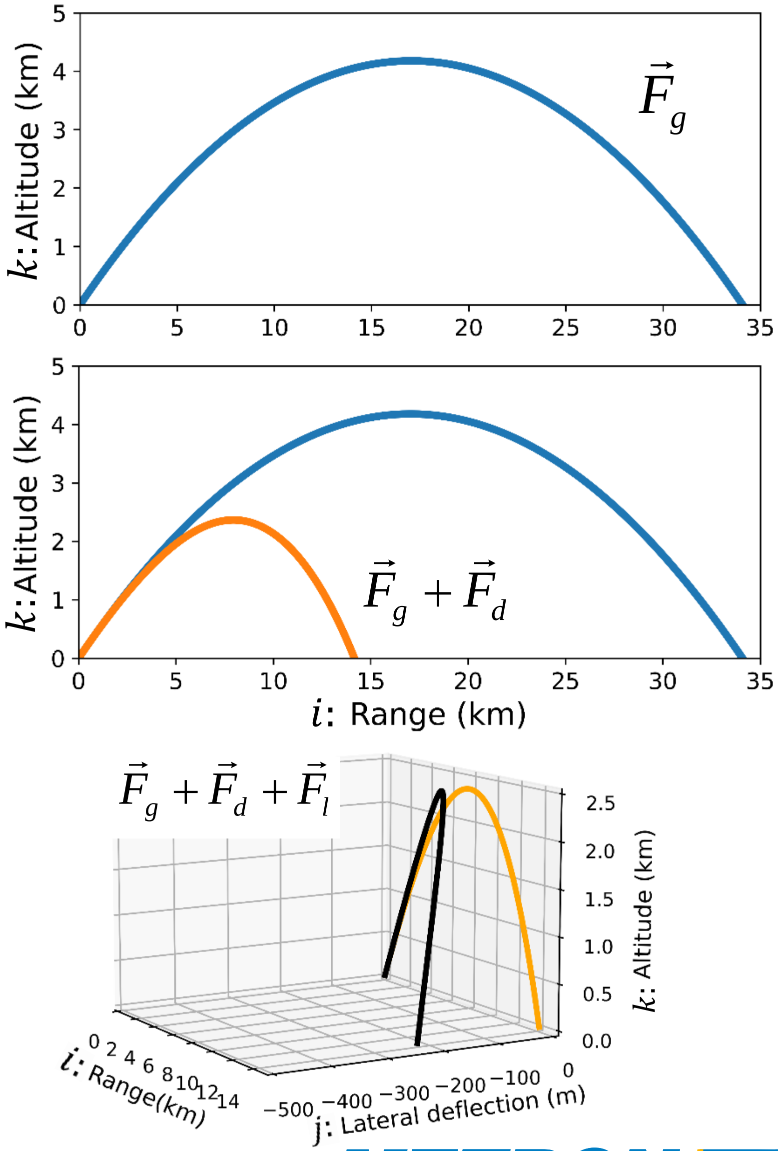

The Mock-Cannon simulator transforms experiments to outcomes

"Black Box" Transformation

Tester conducts experiment, $\boldsymbol{x}$:

1. Fire munition
2. Measure range and quadrant elevation

Quadrant Elevation	Range (km)
45°	15.4
63°	6.3
34°	9.4

? ? ?

$$f: x \rightarrow y$$

? ? ?

Tester observes outcome, y:

Calculate POI error

POI error
151 m
27 m
62 m

What is the relationship between *any* experiment and its outcome?

Model the outcomes

- A ***model*** is a specification between the *parameters* and the *observed data*. The *outcome*/*response* (POI error) variable can be written as

$$y = f(\boldsymbol{x}, \boldsymbol{\theta}) + \varepsilon$$

- Deterministic component, $f(\boldsymbol{x}, \boldsymbol{\theta})$: exact relationship between variables
- Stochastic component, ε: randomness/error inherent in the outcome

<u>*Mock-Cannon*</u> (assume a linear model)

$$f(\boldsymbol{x}, \boldsymbol{\theta}) = \boldsymbol{x}\boldsymbol{\theta}^T = \text{context} \times \text{parameters}$$

$$= \theta_{intercept} + \theta_R \underbrace{x_R}_{\text{Log(Range)}} + \theta_{QE} \underbrace{x_{QE}}_{\text{Quadrant Elevation (radians)}} + \theta_{RQE} \underbrace{x_R x_{QE}}_{\text{First-order cross term}}$$

$\varepsilon \sim \mathcal{N}(0, \sigma^2 \mathbb{I}) \rightarrow$ Assume errors i.i.d. and normally distributed

$\boldsymbol{x} : 1 \times d$
$X : n \times d$
$\boldsymbol{\theta} : 1 \times d$
$\vec{y} : n \times 1$
$\vec{\varepsilon} : n \times 1$
Number of trials, n
Number of features, d
Identity matrix, $\mathbb{I}$

Likelihood distribution of outcomes

- With the model fully specified, the statistical distribution of outcomes, a.k.a. the ***likelihood***, is now:

$$\mathcal{L}(y|\boldsymbol{\theta}, \boldsymbol{x}) = \mathcal{N}(y; \boldsymbol{x}\boldsymbol{\theta}^T, \sigma^2)$$

- From the Mock-Cannon data, we find that $\log(\text{POI error})$ better fits a normal distribution.

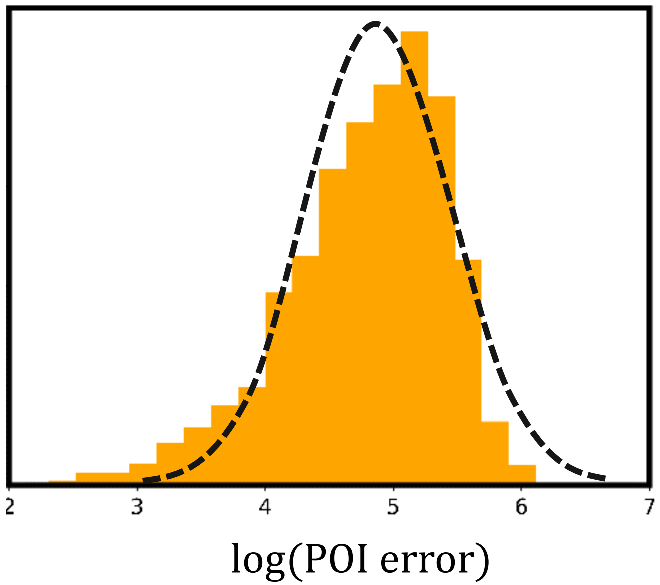

- For the rest of the talk, our outcome y is $\log(\text{POI error})$.

Prior distribution of parameters

- The ***prior*** defines the parameter probability distribution.
- For the mock-cannon, we'll assume:
  - $\boldsymbol{\theta}$ → random variable.
  - σ → known (set to 1 for $\log(\text{POI error})$)
- This give a simple Bayesian conjugate prior structure taking the form of a Normal-Normal distribution:

$$P(\boldsymbol{\theta}|y,X) = \frac{\mathcal{L}(y|\boldsymbol{\theta},X)P(\boldsymbol{\theta})}{P(y|X)}$$

$$\text{Posterior} \propto \text{Likelihood} \times \text{Prior}$$

$$\text{Normal} \propto \text{Normal} \times \text{Normal}$$

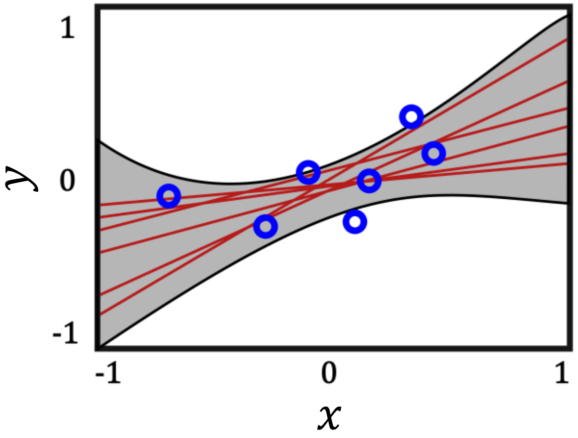

Distribution of "fits" /slopes for our linear model

Prior distribution defines hyperparameter vector, $\boldsymbol{\kappa}$

- The prior distribution on $\boldsymbol{\theta}$ takes the form of a multivariate normal:

$$P(\boldsymbol{\theta}) = \mathcal{N}\left(\boldsymbol{\theta}; \boldsymbol{\mu}, \sigma^2 V\right)$$

- The parameters that characterize the parameter distribution are called "hyperparameters". They define our knowledge about the Mock-Cannon system.

$$\kappa = (\boldsymbol{\mu}, V)$$

Posterior distribution: hyperparameter updates

- The ***posterior*** distribution on $\boldsymbol{\theta}$ takes the form of a multivariate normal:

$$P(\boldsymbol{\theta}|\vec{y}, X) = \mathcal{N}(\boldsymbol{\theta}; \boldsymbol{\mu}^+, \sigma^2 V^+)$$

- As the prior and posterior belong to the same parametric family (Normal), the prior is said to be *conjugate* to the likelihood.

- If you conduct n trials in environments X and measure outcomes $\vec{y}$, a <u>knowledge update</u> $(\kappa \to \kappa^+)$ is given by:

$$\kappa \stackrel{\text{def}}{=} (\boldsymbol{\mu}, V)$$

$$V^+ = \left(V^{-1} + X^T X\right)^{-1}$$

$$\boldsymbol{\mu}^+ = \left(\boldsymbol{\mu} V^{-1} + \vec{y}^T X\right) V^+$$

~OR~

Change of variables

$W = V^{-1}; \boldsymbol{\nu} = \boldsymbol{\mu} W$

$$\kappa \stackrel{\text{def}}{=} (\boldsymbol{\nu}, W)$$

$$\boxed{\begin{aligned} W^+ &= W + X^T X \\ \boldsymbol{\nu}^+ &= \boldsymbol{\nu} + \vec{y}^T X \end{aligned}}$$

Data adds to the knowledge!

$W = V^{-1}$
$\boldsymbol{\nu} = \boldsymbol{\mu} W$

Knowledge update example

Let prior knowledge be $\kappa = \left(\boldsymbol{\mu} = [7,1,-1,1], V = 10 \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}\right)$

1

$$W = V^{-1} = \begin{bmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.1 \end{bmatrix}$$

$$\boldsymbol{\nu} = \boldsymbol{\mu} W = [7,1,-1,1] W = [0.7, 0.1, -0.1, 0.1]$$

X (columns Intercept through QE×Log(Range)); $\vec{y}$ (column Log(POI error))

Intercept	QE	Log(Range)	QE×Log(Range)	Log(POI error)
1	0.8	9.46	7.56	4.52
1	1.1	8.69	9.56	4.21
1	0.6	8.85	5.31	2.64

Knowledge update example

$$W = V^{-1}$$
$$\boldsymbol{\nu} = \boldsymbol{\mu} W$$

Let prior knowledge be $\kappa = \left(\boldsymbol{\mu} = [7,1,-1,1], V = 10\begin{bmatrix}1 & 0 & 0 & 0\\ 0 & 1 & 0 & 0\\ 0 & 0 & 1 & 0\\ 0 & 0 & 0 & 1\end{bmatrix}\right)$

1

$$W = V^{-1} = \begin{bmatrix}0.1 & 0 & 0 & 0\\ 0 & 0.1 & 0 & 0\\ 0 & 0 & 0.1 & 0\\ 0 & 0 & 0 & 0.1\end{bmatrix}$$
$$\boldsymbol{\nu} = \boldsymbol{\mu} W = [7,1,-1,1]W = [0.7, 0.1, -0.1, 0.1]$$

2

$$W^+ = W + X^T X$$
$$W^+ = W + \begin{bmatrix}1 & 1 & 1\\ 0.8 & 1.1 & 0.6\\ 9.46 & 8.69 & 8.85\\ 7.56 & 9.56 & 5.31\end{bmatrix}\begin{bmatrix}1 & 0.8 & 9.46 & 7.56\\ 1 & 1.1 & 8.69 & 9.56\\ 1 & 0.6 & 8.85 & 5.31\end{bmatrix}$$
$$W^+ = \begin{bmatrix}3.1 & 2.5 & 27 & 22.4\\ 2.5 & 2.31 & 22.4 & 19.8\\ 27 & 22.4 & 243.1 & 201\\ 22.4 & 19.8 & 201 & 176.1\end{bmatrix}$$
$$\boldsymbol{\nu}^+ = \boldsymbol{\nu} + \vec{y}^T X$$
$$\boldsymbol{\nu}^+ = \boldsymbol{\nu} + [4.52, 4.21, 2.64]\begin{bmatrix}1 & 0.8 & 9.46 & 7.56\\ 1 & 1.1 & 8.69 & 9.56\\ 1 & 0.6 & 8.85 & 5.31\end{bmatrix}$$
$$\boldsymbol{\nu}^+ = [12.07, 9.93, 102, 88.61]$$

X				$\vec{y}$
Intercept	QE	Log(Range)	QE×Log(Range)	Log(POI error)
1	0.8	9.46	7.56	4.52
1	1.1	8.69	9.56	4.21
1	0.6	8.85	5.31	2.64

Knowledge update example

$$W = V^{-1}$$
$$\boldsymbol{\nu} = \boldsymbol{\mu} W$$

Let prior knowledge be $\kappa = \left(\boldsymbol{\mu} = [7,1,-1,1], V = 10 \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix}\right)$

1

$$W = V^{-1} = \begin{bmatrix} 0.1 & 0 & 0 & 0 \\ 0 & 0.1 & 0 & 0 \\ 0 & 0 & 0.1 & 0 \\ 0 & 0 & 0 & 0.1 \end{bmatrix}$$

$$\boldsymbol{\nu} = \boldsymbol{\mu} W = [7,1,-1,1] W = [0.7, 0.1, -0.1, 0.1]$$

2

$$W^+ = W + X^T X$$

$$W^+ = W + \begin{bmatrix} 1 & 1 & 1 \\ 0.8 & 1.1 & 0.6 \\ 9.46 & 8.69 & 8.85 \\ 7.56 & 9.56 & 5.31 \end{bmatrix} \begin{bmatrix} 1 & 0.8 & 9.46 & 7.56 \\ 1 & 1.1 & 8.69 & 9.56 \\ 1 & 0.6 & 8.85 & 5.31 \end{bmatrix}$$

$$W^+ = \begin{bmatrix} 3.1 & 2.5 & 27 & 22.4 \\ 2.5 & 2.31 & 22.4 & 19.8 \\ 27 & 22.4 & 243.1 & 201 \\ 22.4 & 19.8 & 201 & 176.1 \end{bmatrix}$$

$$\boldsymbol{\nu}^+ = \boldsymbol{\nu} + \vec{y}^T X$$

$$\boldsymbol{\nu}^+ = \boldsymbol{\nu} + [4.52, 4.21, 2.64] \begin{bmatrix} 1 & 0.8 & 9.46 & 7.56 \\ 1 & 1.1 & 8.69 & 9.56 \\ 1 & 0.6 & 8.85 & 5.31 \end{bmatrix}$$

$$\boldsymbol{\nu}^+ = [12.07, 9.93, 102, 88.61]$$

3

$$V^+ = (W^+)^{-1} = \begin{bmatrix} 8.29 & -1.11 & -0.89 & 0.13 \\ -1.11 & 8.84 & 0.12 & -0.93 \\ -0.89 & 0.12 & 0.10 & -0.02 \\ 0.13 & -0.93 & -0.02 & 0.11 \end{bmatrix}$$

$$\boldsymbol{\mu}^+ = \boldsymbol{\nu}^+ (W^+)^{-1} = [6.33, 0.31, -0.54, 0.28]$$

$$\underbrace{\text{Intercept} \quad \text{QE} \quad \text{Log(Range)} \quad \text{QE×Log(Range)}}_{X} \quad \underbrace{\text{Log(POI error)}}_{\vec{y}}$$

Intercept	QE	Log(Range)	QE×Log(Range)	Log(POI error)
1	0.8	9.46	7.56	4.52
1	1.1	8.69	9.56	4.21
1	0.6	8.85	5.31	2.64

Knowledge update example: Marginal distributions

$P(\boldsymbol{\theta}) \rightarrow$ **prior**

$P(\boldsymbol{\theta}|\vec{y}, X) \rightarrow$ **posterior**

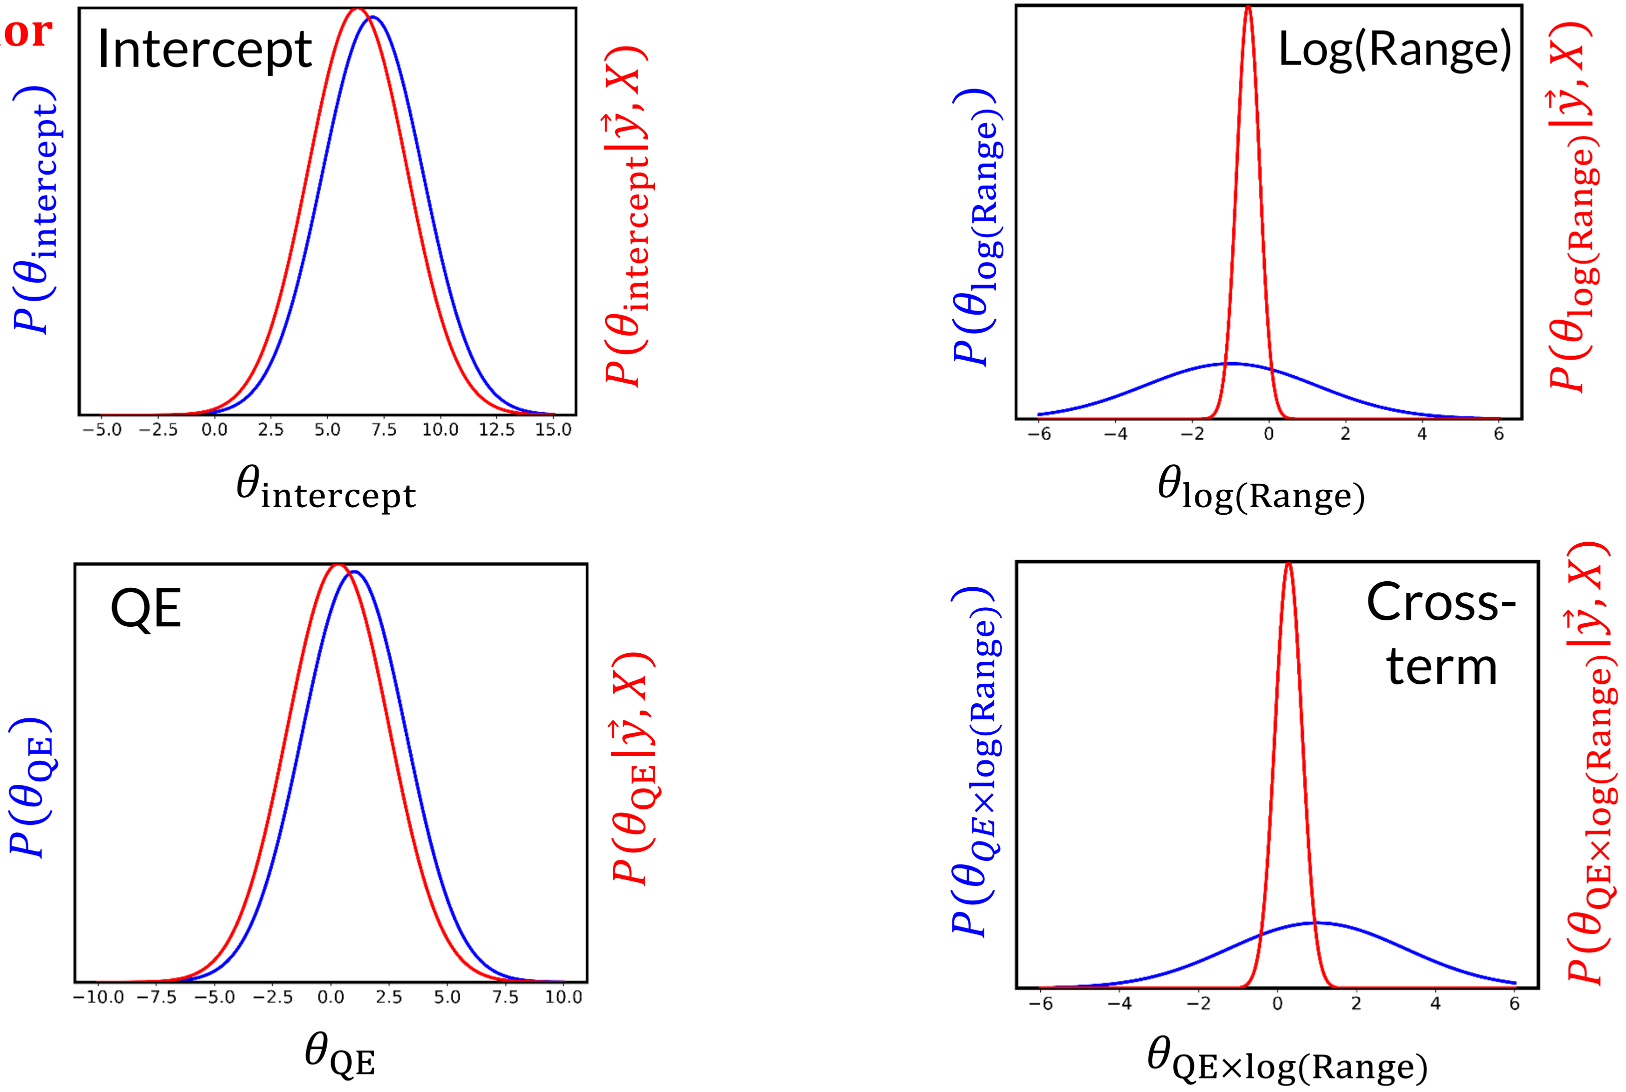

Setting the prior

- Incorporates subject matter expert (SME) knowledge.
- When little is known about $\boldsymbol{\theta}$, it's common to set $\boldsymbol{\mu} = 0$ and $\mathrm{V} = \text{large diagonal matrix.}$
- We'll do one better:
    - The intercept shares the same units as the POI error.
    - Set intercept to > 100 m (req. thresh).
        - We chose POI error intercept to be 500 m $\Longrightarrow \mu_{intercept} \approx 6.2$
    - Let $V = \lambda\mathbb{I}$. Vary λ and observe effect.
- When setting the prior, we can observe the effects on the *prior predictive distribution*:

$$P(\vec{y}|X,\kappa) = \int \mathcal{L}(\vec{y}|\boldsymbol{\theta}, X)P(\boldsymbol{\theta})d\boldsymbol{\theta} = \mathcal{N}\left(\vec{y}; X\boldsymbol{\mu}^T, \sigma^2(\mathbb{I} + XVX^T)\right)$$

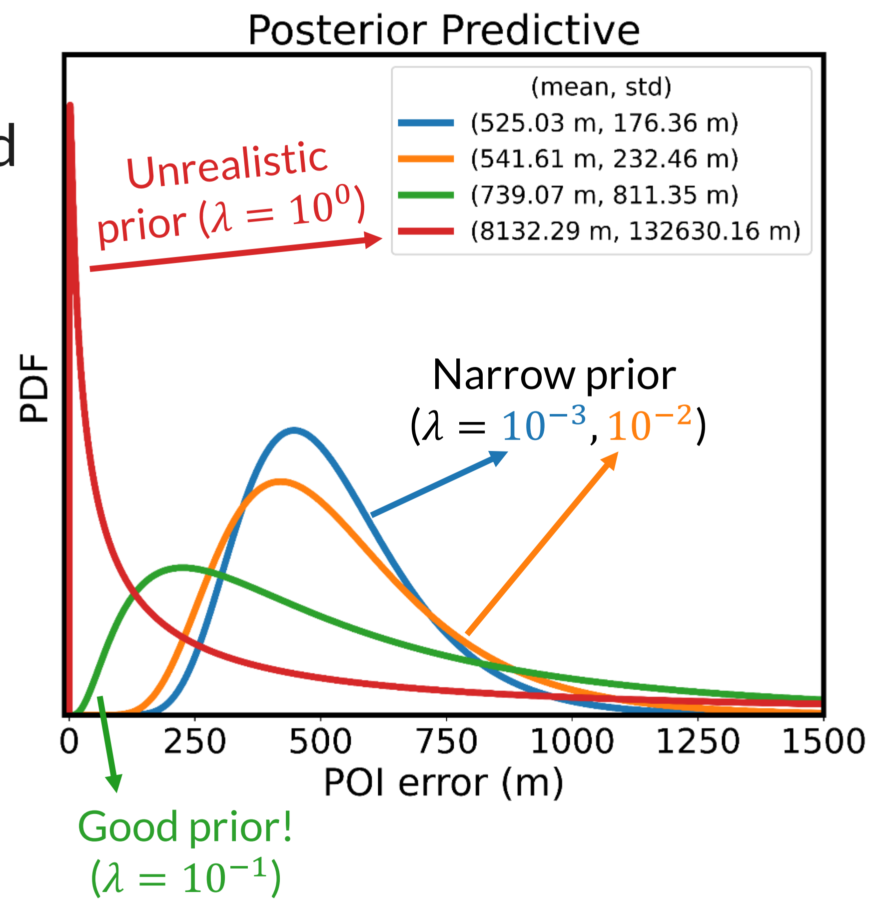

Tip: Convert normal to log-normal distribution to view POI error (instead of log(POI error))

Mock-cannon simulator: response surface

- Represent outcome randomness in POI through Gaussian variability in muzzle velocity and quadrant elevation to represent variation in repeated testing.
- This model shows curvature in response surface and non-trivial relationship between contexts and outcomes.

Mach number dependent aerodynamic coefficients
Gravitational force
Drag force
Lift force
Coriolis Force
Magnus Force
Overturning moment
Spin damping moment
Wind velocity field
Pressure
Temperature
Air density
Etc…

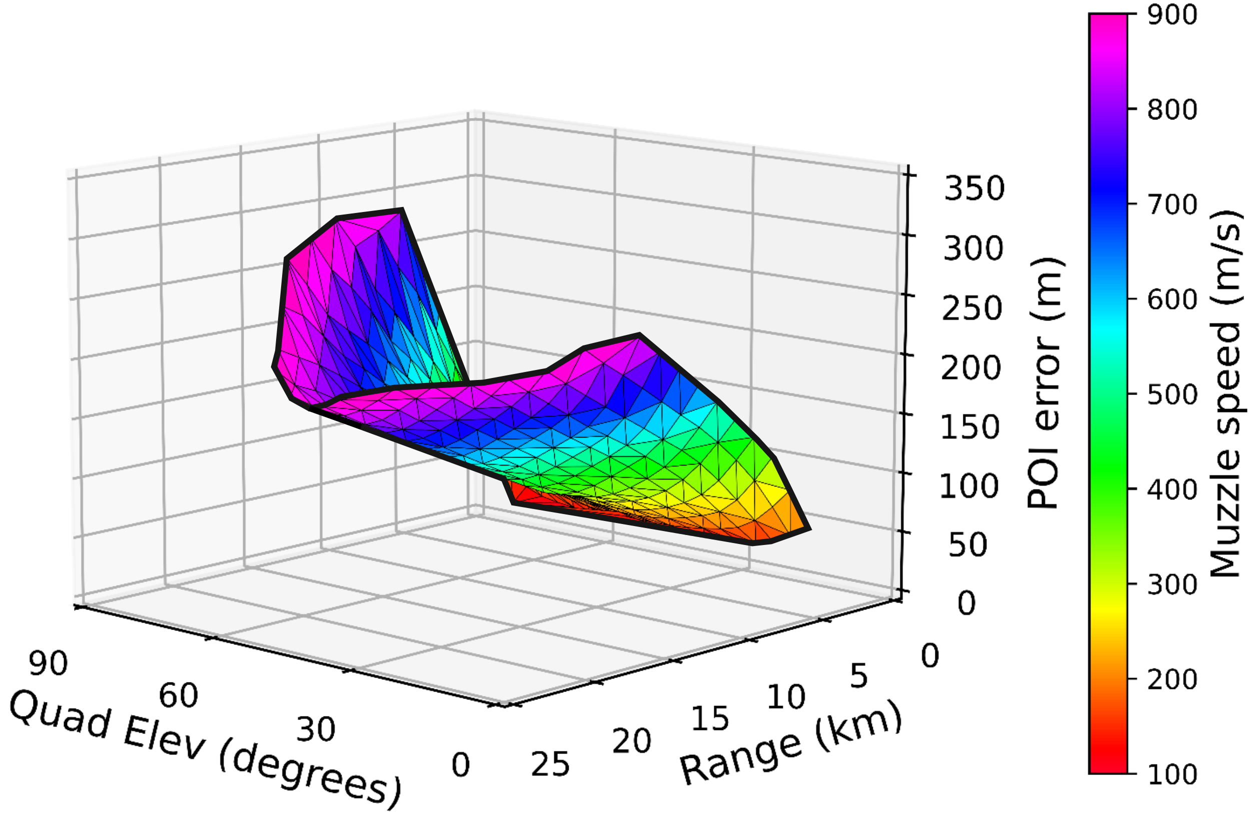

*Plot generated by simulating 50 firings per point and averaging POI error

Part II:

(Bayesian) Decision Theory

Decision Theory Basics

- Decision theory: a branch of probability that relates beliefs and preferences to making choices between alternatives.
- We concern ourself with *normative* decision theory:
  - Identifies optimal decisions assuming agent is fully rational
- Given two outcomes of the universe, A and B, express preferences via:
  - $A \succ B$ if we prefer A over B.
  - $A \sim B$ if we are indifferent between A and B.
  - $A \succcurlyeq$ B if we prefer A over B or are indifferent.
- Define real-valued function U (*utility*) that maintains the ordinal relationship between preferences:
  - $U(A) > U(B)$ iff $A \succ B$.
  - $U(A) = U(B)$ iff $A{\sim}B$.

Defining the decision space

- $\mathcal{S}$, state space: realizations/configurations of the world
- $\mathcal{A}$, action space: set of all available actions, $a: \mathcal{S} \rightarrow \mathcal{O}$
- $\mathcal{O}$, outcome space: consequence/reward of a joint state-action

- The preference of an outcome is defined as the ***utility***, $u_a(o|s)$ with $u: \mathcal{O} \rightarrow \mathbb{R}$.
  - Utilities are elicited that encode agent's preferences.

Example: should you take an umbrella?

- $\mathcal{S}$, state space: the weather forecast (sunny or raining)
- $\mathcal{A}$, action space: either taking an umbrella or not
- $\mathcal{O}$, outcome space: the consequence of (not) having an umbrella in different weather conditions

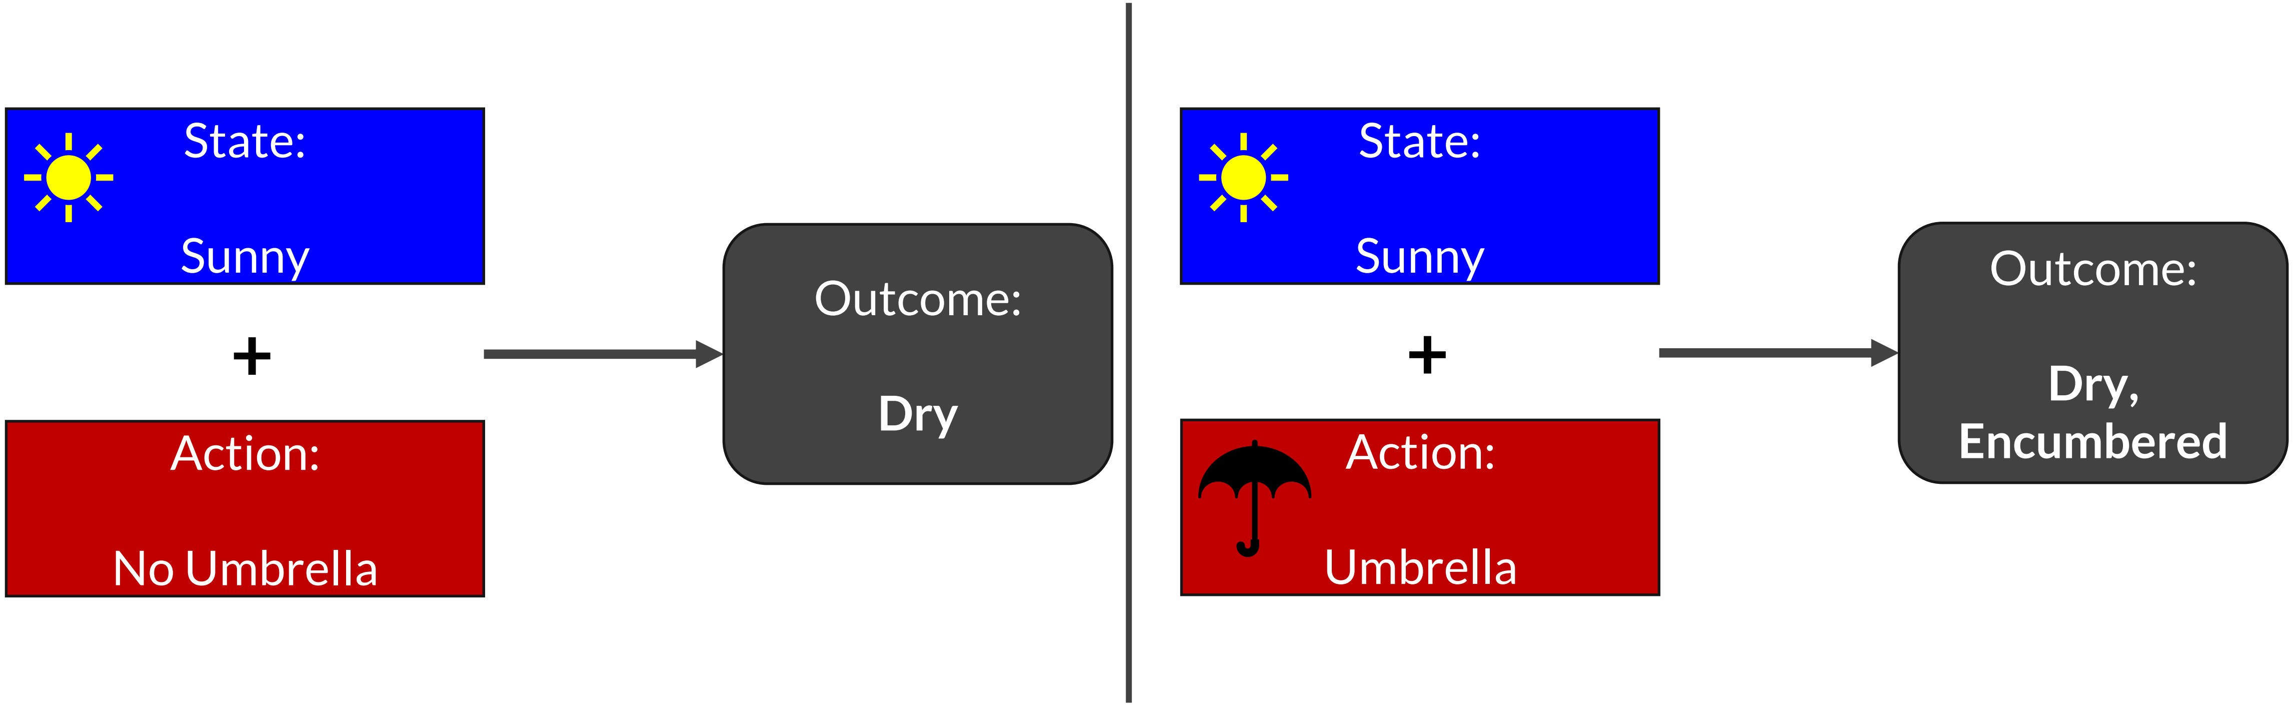

Example: should you take an umbrella?

- $\mathcal{S}$, state space: the weather forecast (sunny or raining)
- $\mathcal{A}$, action space: either taking an umbrella or not
- $\mathcal{O}$, outcome space: the consequence of (not) having an umbrella in different weather conditions

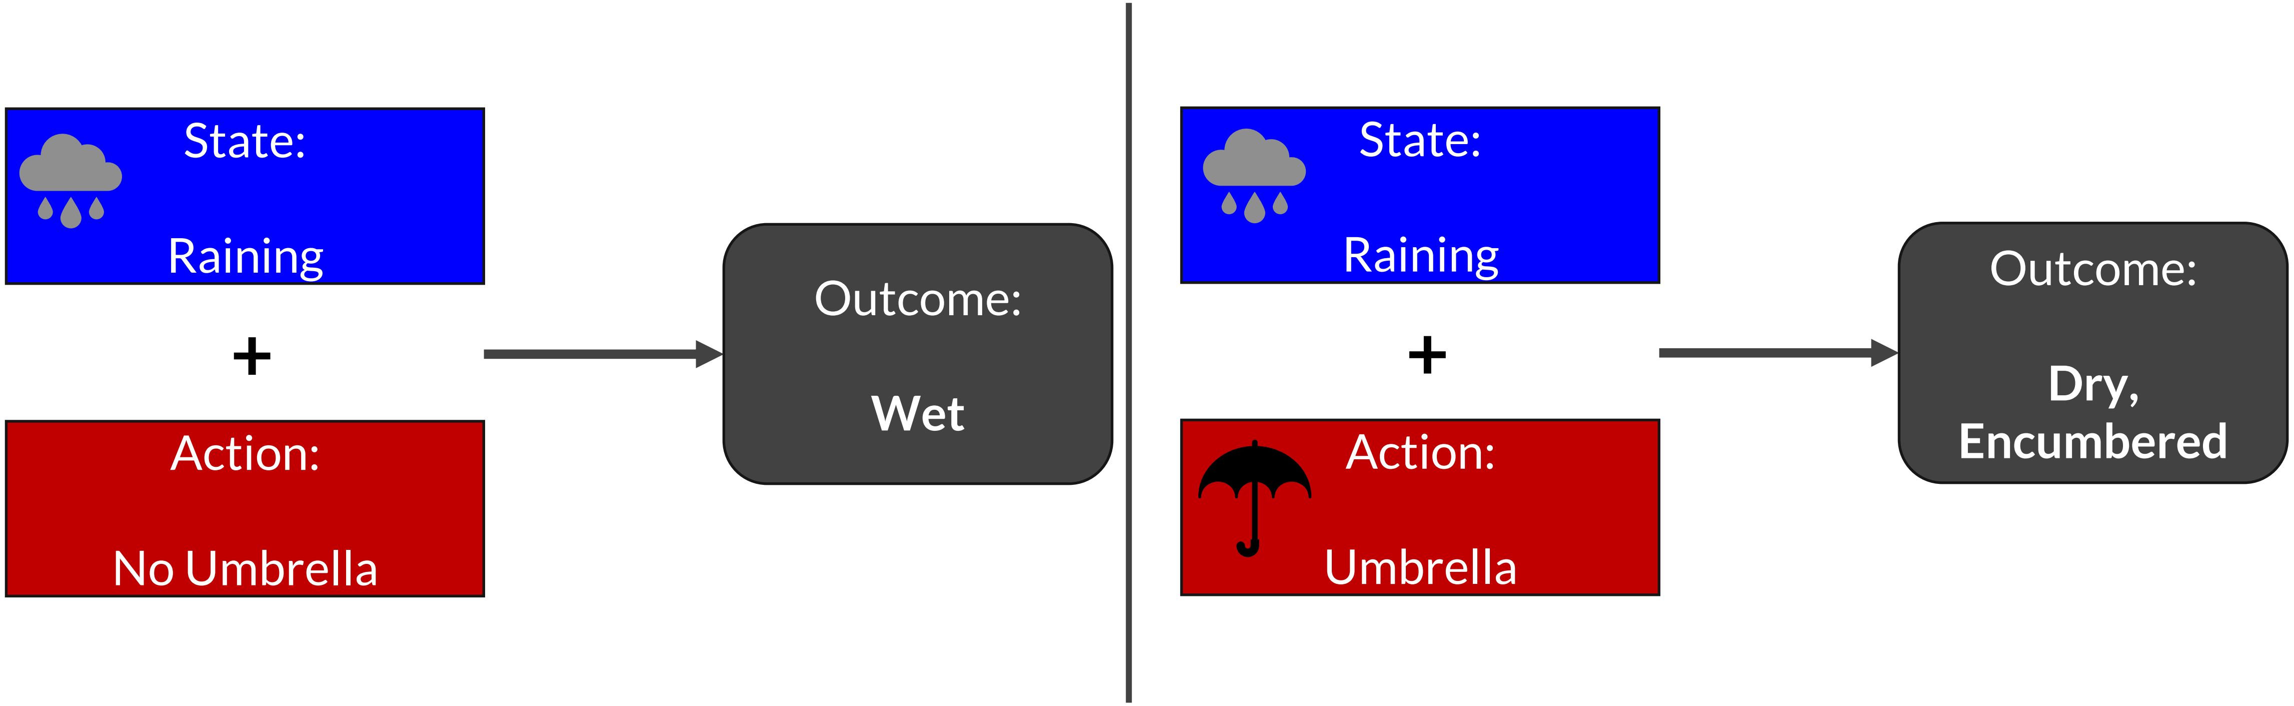

Rank preferences using a utility function

- A utility function takes our qualitative ranking of outcomes and represents them as ordinal, real numbers.

Preference

$$\text{Dry} \quad \succ \quad \text{Dry, Encumbered} \quad \succ \quad \text{Wet}$$

$$(\mathrm{D}|S, NU) \succ (DE|S, U) \sim (DE|R, U) \succ (\mathrm{W}|\mathrm{R}, \mathrm{NU})$$

Utility

$$u_{NU}(\mathrm{D}|S) > u_U(DE|S) = u_U(DE|R) > u_{\mathrm{NU}}(\mathrm{W}|\mathrm{R})$$

$$1 \quad > \quad 0.80 \quad = \quad 0.80 \quad > \quad 0$$

States:
- Raining, R
- Sunny, S

Actions:
- No Umbrella, NU
- Umbrella, U

Outcomes:
- Dry, D
- Dry, Encumbered, DE
- Wet, W

Some actions have costs! Can be thought of as:

$$u_U(DE|S) = u_{NU}(D|S) - C(U)$$
$$= \quad 1 \quad - \quad 0.2$$

$C(U)$: cost of carrying umbrella

Decisions under uncertainty

Principle of maximum expected utility: Under state uncertainty, the optimal action maximizes the expected utility:

$$a^* \in \underset{a \in \mathcal{A}}{\operatorname{argmax}}\ \mathbb{E}_{\mathcal{S}}[u_a(o|s)]$$

"Decision Table"		
States / Actions	*Sunny* $P(s_1) = 70\%$	*Raining* $P(s_2) = 30\%$
No umbrella, a_1	Dry $u_{a_1}(o_{11}\|s_1) = 1$	Wet $u_{a_1}(o_{12}\|s_2) = 0$
Umbrella, a_2	Encumbered, dry $u_{a_2}(o_{21}\|s_1) = 0.8$	Encumbered, dry $u_{a_2}(o_{22}\|s_2) = 0.8$

"Irrational" action

$$\bar{u}(a_1) = 1 \times 0.7 + 0 \times 0.3 = 0.7$$

$$\boxed{\bar{u}(a_2) = 0.8 \times 0.7 + 0.8 \times 0.3 = 0.8}$$

"Rational" action

How do we apply decisions to the Mock-Cannon?

- A tester's decision process involves a *sequence* of decisions (sequential decision theory).

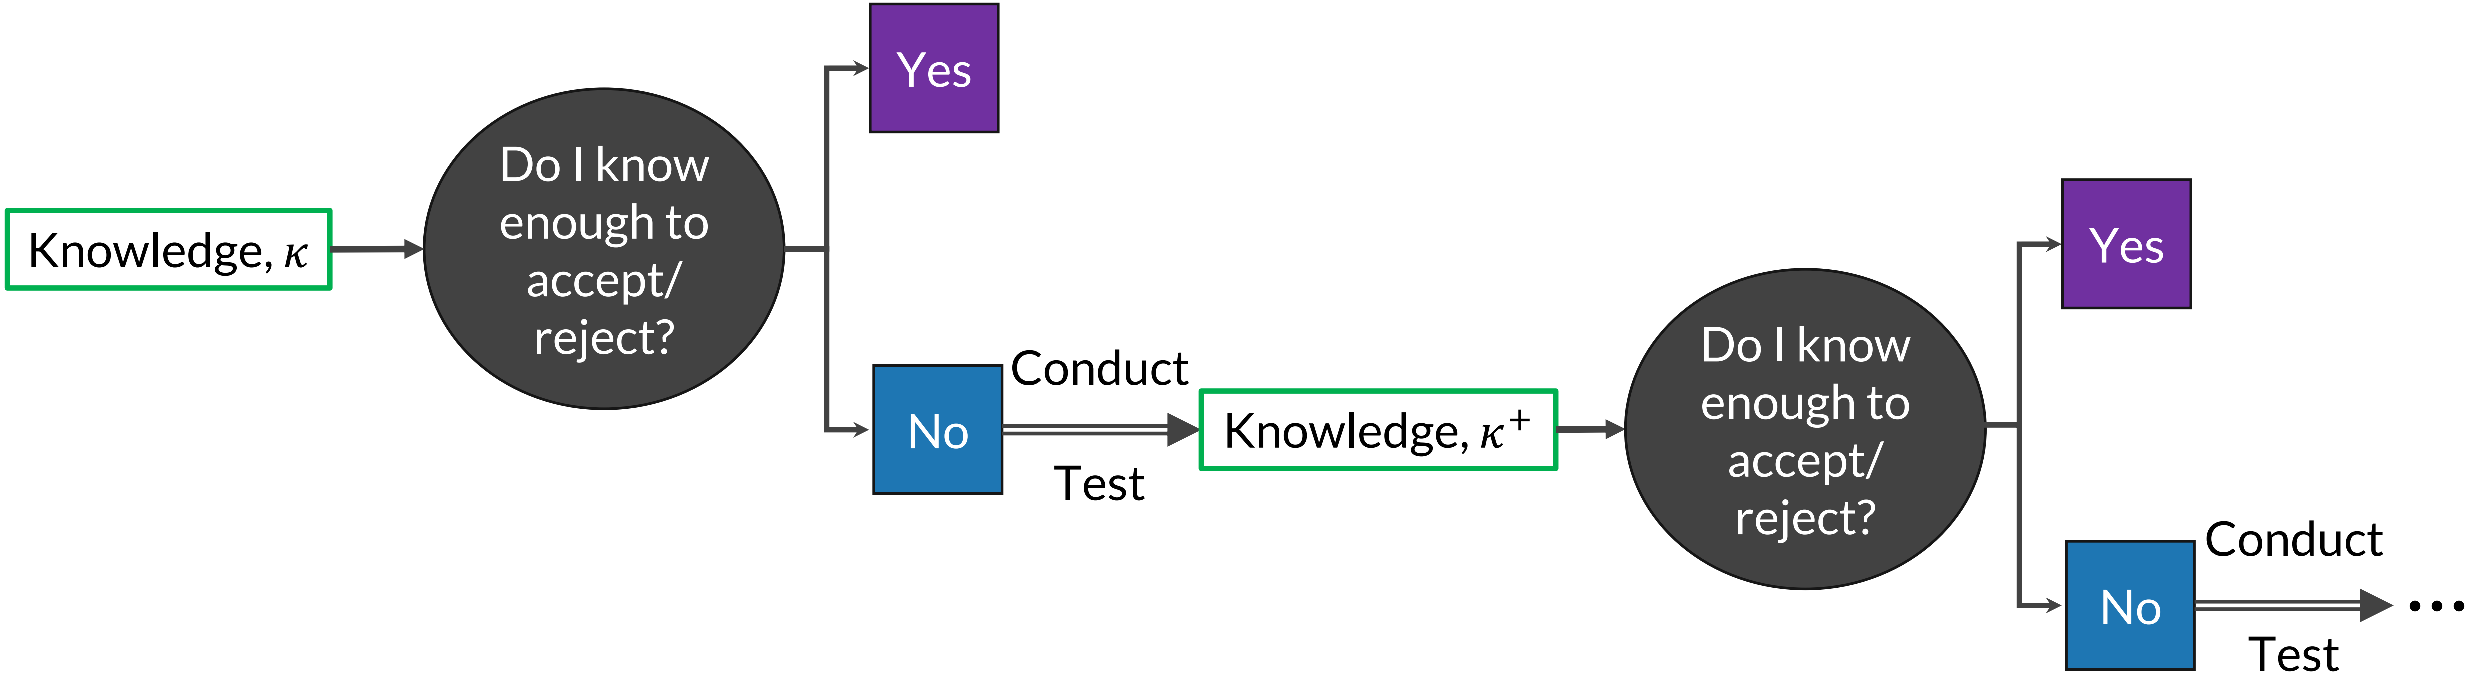

- *Terminal decisions* (**Yes** nodes) are decisions that end the sequence.
- *Intermediate (non-terminal) decisions* (**No** nodes) influence future terminal decisions.

Bayesian Decision Theory handles state uncertainty

- Extend the discussion to the continuum, e.g., the mock-cannon (more difficult to enumerate in practice; integrate over states)
  - $\mathcal{S}$, state space: uncertain parameters $\boldsymbol{\theta}$ as defined by κ, or $\boldsymbol{\theta}|\kappa$
    - Bayes Rule: probability distributions defined by the prior
  - $\mathcal{A}$, action space: $\{\text{Accept, Reject, Continue Testing}\} = \{A, R, T\}$
    - Terminal decisions are defined as D: $\{A, R\}$
  - $\mathcal{O}$, outcome space: fraction of points contained within CEP50 radius
    - No longer discrete outcomes
    - "Containment Probability"

State: $\theta|\kappa$

+

Action:
D:{Accept, Reject}
T:{Continue Test}

→

Outcome:
Containment Probability

Assessing the utility of terminal decisions

Reject utility, $\boldsymbol{u_R(\kappa)}$

- If you reject the Mock-Cannon, i.e., send it back to the manufacturer, then it adds no value to you.

$$\boxed{u_R(\kappa) = 0}$$

Accept option: representing preferences

- We define our preference to be the *containment probability*, p_c, representing the probability of fires that land within $R_{CEP50} = 100\ m$ ($\log(100) \approx 4.61$).
- Integrate area under the likelihood distribution:

$$p_c(\boldsymbol{x}, \boldsymbol{\theta}) = \int_{-\infty}^{\log(R_{CEP50})} \mathcal{L}(y | \boldsymbol{x}\boldsymbol{\theta}^T, \sigma^2) dy$$

- Our *preference* of any $\boldsymbol{\theta}$ is averaged over all operational environments:

$$\bar{p}_c(\boldsymbol{\theta}) = \mathbb{E}_{\boldsymbol{x}}[p_c(\boldsymbol{x}, \boldsymbol{\theta})]$$

- The Mock-Cannon is more preferable with increasing containment probability.

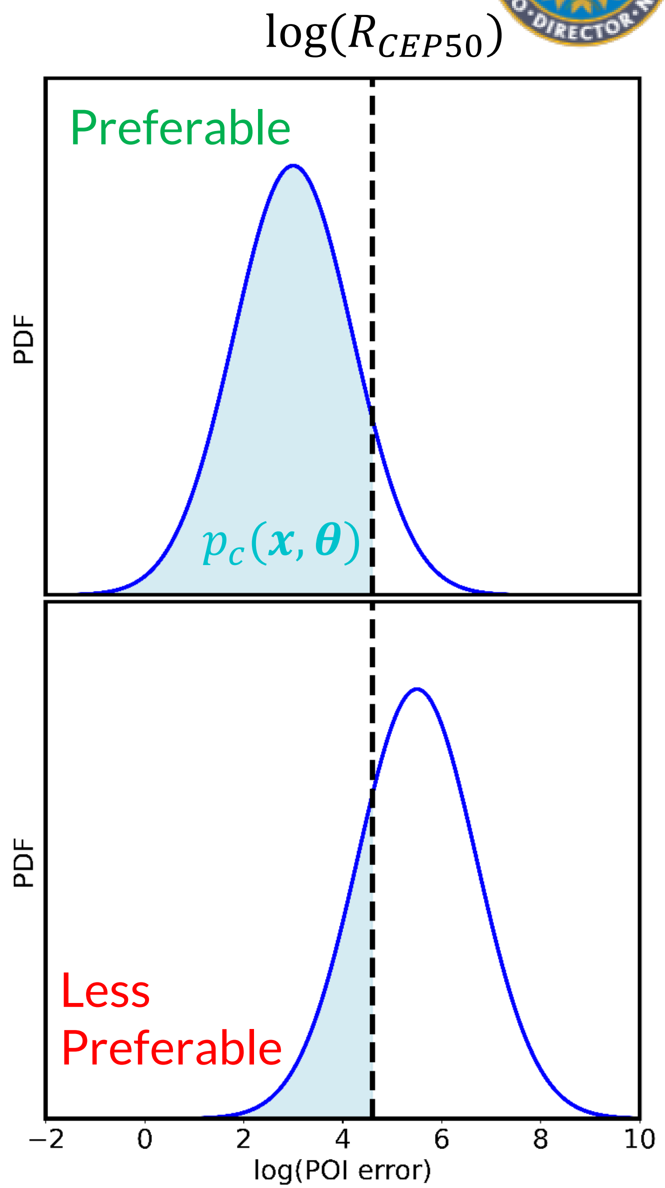

Accept utility, $\boldsymbol{u_A(\kappa)}$

- The utility of any state/parameter $\boldsymbol{\theta}$ encodes our preference relative to the requirement threshold:

$$u(\boldsymbol{\theta}) = \text{sigmoid}(\bar{p}_c(\boldsymbol{\theta}) - 0.5)$$

- The ***accept utility*** is the expected utility over all parameters *minus* some fixed acquisition cost (e.g. cost to productionalize the system) (C_A):

$$\boxed{u_A(\kappa) = \mathbb{E}_{\boldsymbol{\theta}|\kappa}[u(\boldsymbol{\theta})] - C_A}$$

Utility ($C_A = 0$)

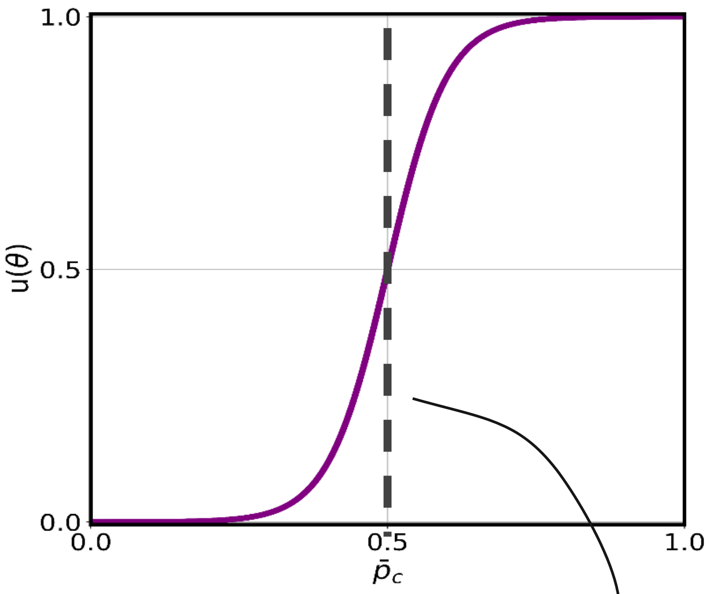

Requirement threshold (CEP50)

Terminal utility, $\boldsymbol{u_D(\kappa)}$

- Terminal decisions are actions that end the sequential decision process, $d \in D: \{A, R\}$.

- The terminal utility is maximized over the terminal action utilities in accordance with the *principle of maximum expected utility*:

$$\boxed{u_D(\kappa) = \max_{d \in D}(u_A(\kappa), u_R(\kappa) = 0)}$$

Assessing the utility of the intermediate decision

Utility function for Continue Testing

- The continue testing action, T, addresses the questions:
  - "What is the value in conducting another set of tests in environment X_T?"
  - "Would a **future** test X_T yield a higher overall terminal utility?"
- We will project out future states by sampling outcomes.
- As we do not know the outcomes before conducting X_T, we can simulate future outcomes from the posterior predictive distribution:

$$P(\vec{y}_{pp}|X,\kappa) = \int \mathcal{L}(\vec{y}_{pp}|\boldsymbol{\theta},X)P(\boldsymbol{\theta}|\vec{y},X)d\boldsymbol{\theta} = \mathcal{N}\left(\vec{y}_{pp}; X\boldsymbol{\mu}^T, \sigma^2(\mathbb{I}+XVX^T)\right)$$

- $\vec{y}_{pp}$: *simulated* outcome vectors from the ***p**osterior **p**redictive*.

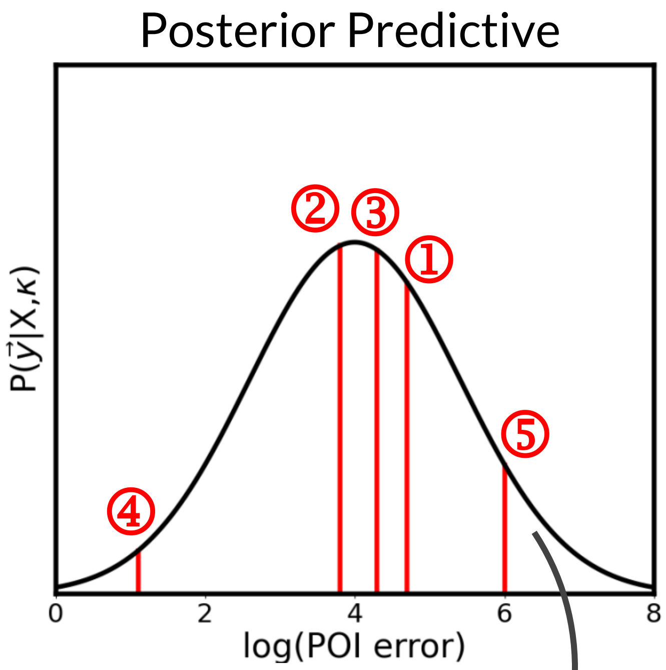

Example of 5 random draws that represent potential outcomes from conducting a test

X_T

Intercept	QE	Log(Range)	QE×Log(Range)	$\vec{y}_{pp,1}$	$\vec{y}_{pp,2}$	$\vec{y}_{pp,3}$
1	0.2	8.08	1.67	4.52	7.14	5.80
1	1.1	8.69	9.56	3.43	8.55	6.01
1	0.6	8.85	5.31	2.34	7.28	5.44

Sequential decision theory projects future knowledge states with a decision tree

- Instead of a decision table, we can visualize the $(\mathcal{A}, \mathcal{S}, \mathcal{O})$ with a *decision tree.*
- The branching possibilities are in principle endless for continuous states. Visualize for a *1-step lookahead.*
- Project out M knowledge representations (κ^+) for test option X_T.

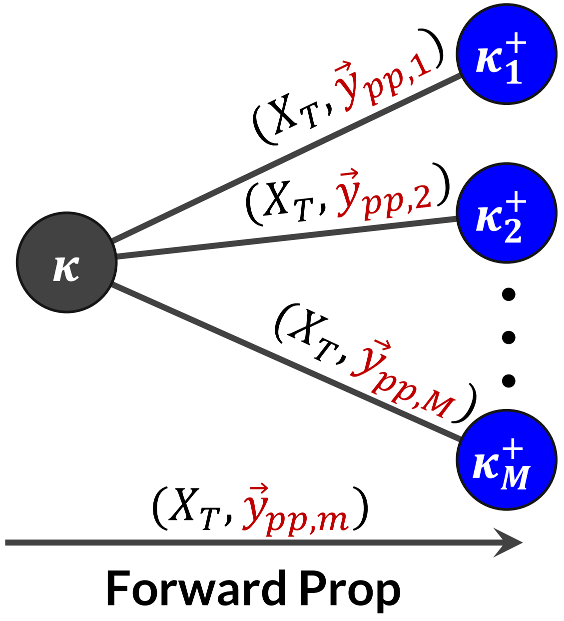

<u>Update Equations</u>

$$V_m^+ = \left(V^{-1} + X_T^T X_T\right)^{-1}$$
$$\boldsymbol{\mu}_m^+ = (\boldsymbol{\mu} V^{-1} + \vec{y}_{pp,m}^{\,T} X_T) V_m^+$$

Continue Testing utility is the expected terminal utility

- The continue testing utility is the expectation over all future terminal utilities *minus* the cost of performing the test (C_T):

$$\boxed{u_T(\kappa) = \mathbb{E}_{\kappa^+|\kappa}[u_D(\kappa^+)] - C_T}$$

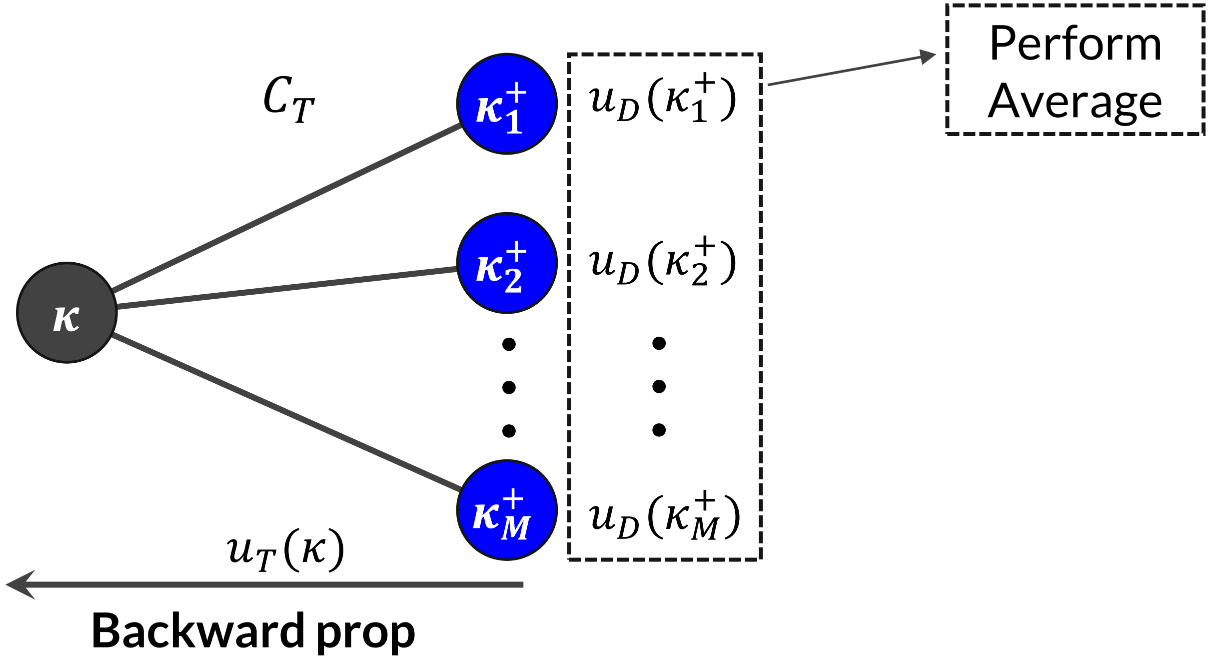

Note: this is an inductive process. Additional future projections follow a similar pattern, e.g.,

$$u_T(\kappa^+) = \mathbb{E}_{\kappa^{++}|\kappa^+}[u_D(\kappa^{++})] - C_T$$

The optimal action

- We arrive at the optimal action over the entire action space:

$$a^* = \underset{a\in\mathcal{A}}{\operatorname{argmax}}\left(u_a(\kappa)\right) = \underset{a\in\mathcal{A}}{\operatorname{argmax}}(u_D(\kappa), u_T(\kappa)) = \underset{a\in\mathcal{A}}{\operatorname{argmax}}(u_A(\kappa), u_R(\kappa) = 0, u_T(\kappa))$$

- This process repeats sequentially with every test conducted until the intermediate action (Continue Testing) is no longer optimal/maximal.

- X_T consists of 30 test points.
  - $a^* = \underset{a\in\mathcal{A}}{\operatorname{argmax}}(\underset{A}{0.10}, \underset{R}{0}, \underset{T}{0.14}) = T$

Intercept	QE	Log(Range)	QE×Log(Range)	Log(POI error) (*true outcome*)
1	0.8	9.46	7.56	4.52
1	1.1	8.69	9.56	4.21
1	0.6	8.85	5.31	2.64
~ 27 more rows ~				

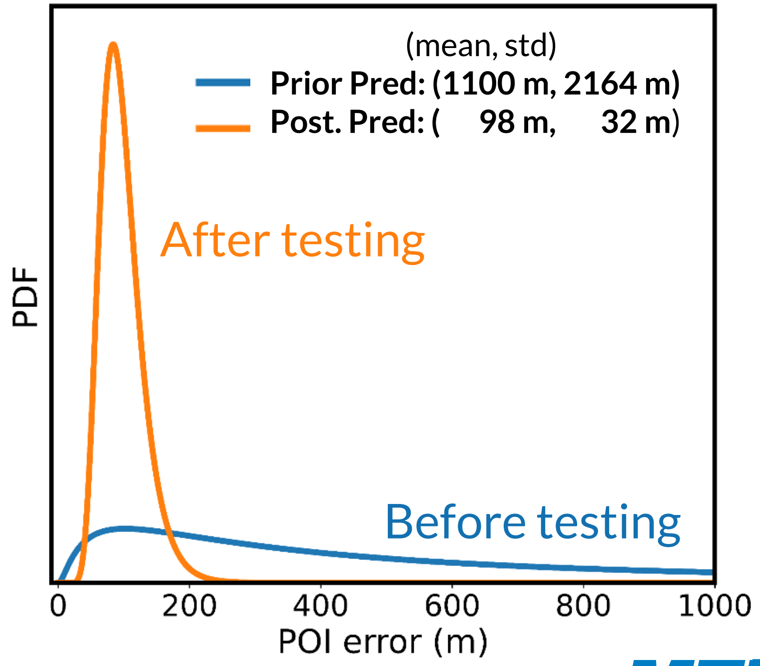

Summary, Part 1: Applying Bayes to the Mock-Cannon

Model, Likelihood, Evaluation

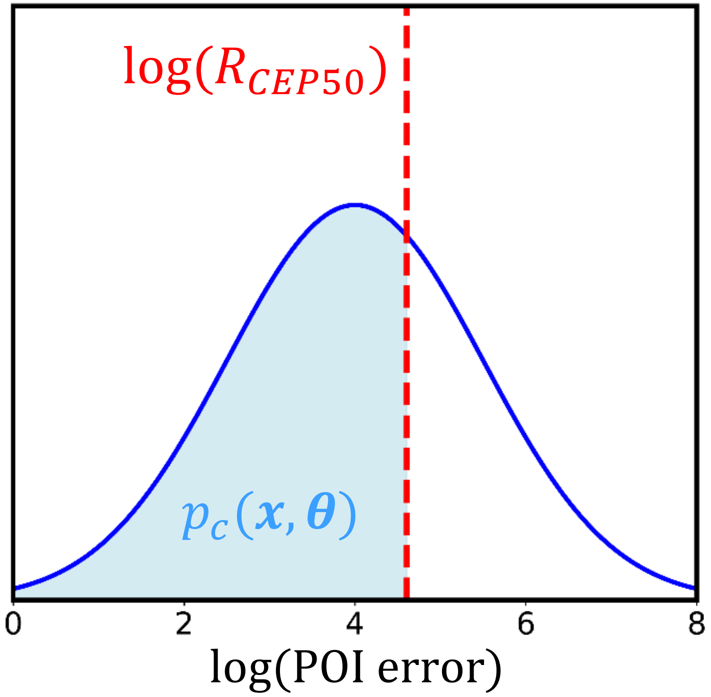

Prior, Posterior, Knowledge update

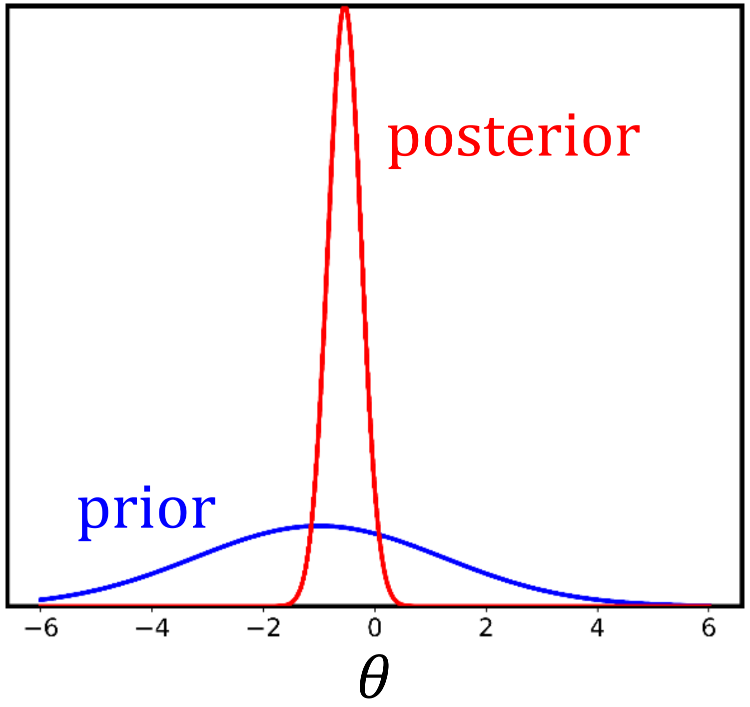

Setting a prior

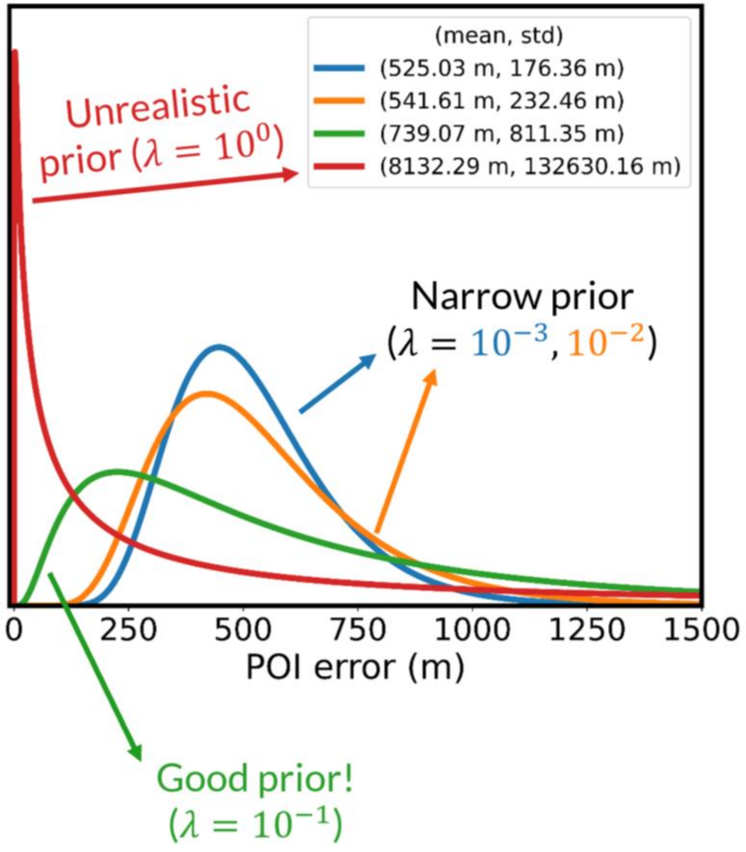

Summary, Part II: Bayesian Decision Theory

Decision space

State: $\theta|\kappa$

+

Action:
D:{Accept, Reject}
T: {Continue Test}

Outcome:
Containment
Prob, $\bar{p}_c(\theta)$

Preferences encoded
in utility function

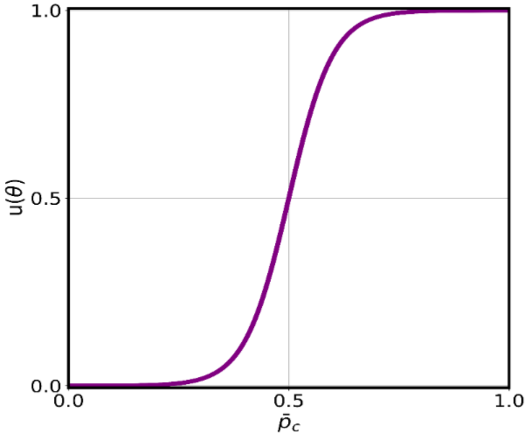

Principle of maximum expected utility

$$a^* \in \underset{a' \in \mathcal{A}}{\operatorname{argmax}} \, \mathbb{E}_S\left[u_{a'}(o|s)\right]$$

"Rational agents maximize expected utility."

Sequential decisions analyzed by projecting future states

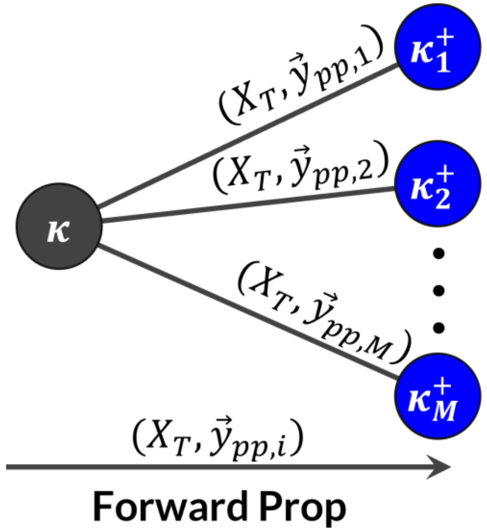

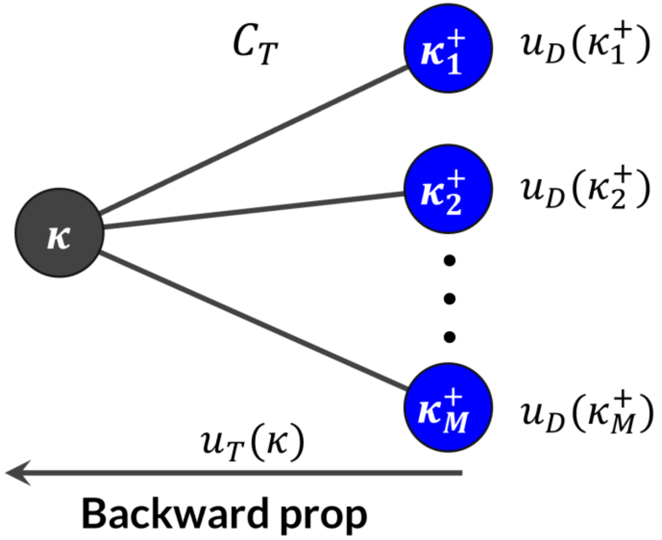

Questions?

References

- Bayesian Decision Theory
  - Berger, J. O. *Statistical Decision Theory and Bayesian Analysis 2nd Edition*. Springer, New York (2013).
  - Parmigiani, G. and Inoue, L. *Decision Theory: Principles and Applications*. Wiley, London, UK (2009).
  - Lindley, D. V. *Making Decisions 2nd Edition*. Wiley, London, UK (1971).
  - Lindley, D.V. "On a measure of the information provided by an experiment," *Ann. Math. Statistics*, vol. 27, pp. 986-1005 (1956).
  - H. Chernoff. *Sequential Analysis and Optimal Design*. Philadelphia, SIAM (1972).
- Munition Simulations
  - W. Y. Lim. "Predicting the Accuracy of Unguided Artillery Projectiles" M.S. thesis, Naval Postgraduate School, Montery, CA (2016). https://apps.dtic.mil/sti/tr/pdf/AD1029824.pdf
  - M. E. Wessam and Z. H. Chen. "Firing precision evaluation for unguided artillery projectile" in *Int. Conf. Artificial Intell. Ind. Eng.* (2015).
  - R. L. McCoy. "Modern Exterior Ballistics" Atglen, PA: Schiffer Publishing (2012).
  - M. Skande. "Numerical solution to a nonlinear external ballistics model for a direct fire control system" M.S. thesis, Dept. Mach. Design, KTH School of Ind. Eng. Manage., Stockholm, Sweden (2014).
  - M. Khalil, H. Abdalla, and O. Kamal. "Dispersion analysis for spinning artillery projectile". *13th Int. Conf. ASAT*, Cairo, Egypt, Mar 26-28 (2009).

Backup slides

Define the likelihood of outcomes

- The ***likelihood*** is a specification of the stochastic nature of outcomes.

2D Normal of positions → *1D Rayleigh of radii* $\xrightarrow{Y = \log(y)}$ *1D Log-Rayleigh of radii exponents* $\approx$ *1D Normal of radii exponents*

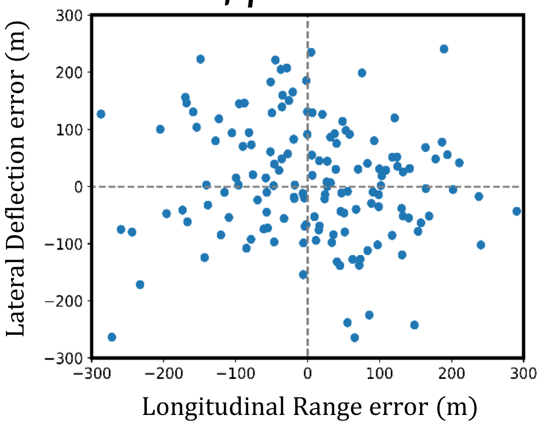

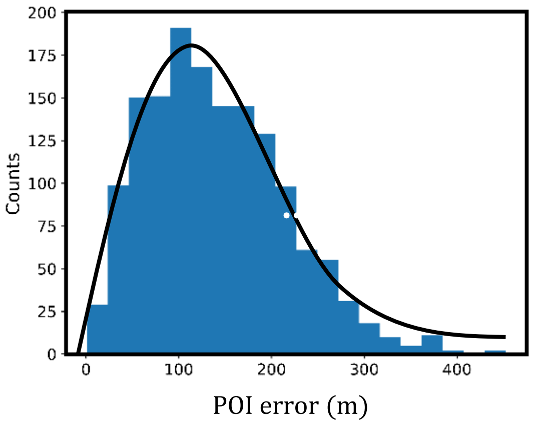

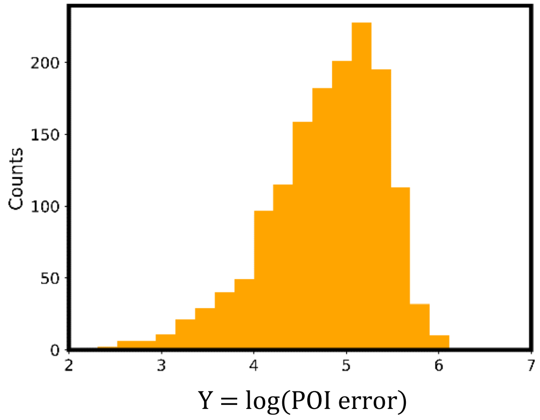

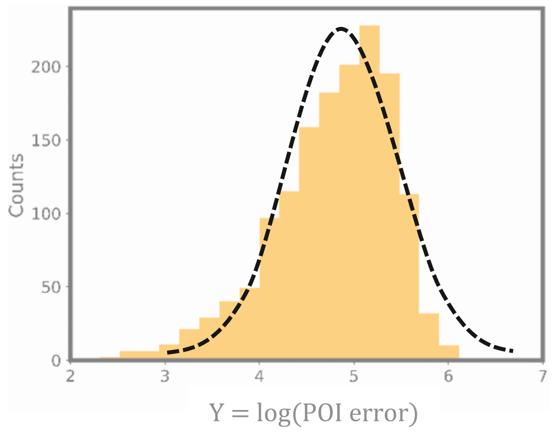

- To use the Bayesian conjugate prior construction, likelihood needs to be normally distributed.

<u>The likelihood</u>:

$$\mathcal{L}(y|\boldsymbol{\theta}, \boldsymbol{x}) = \mathcal{N}(y; \boldsymbol{x}\boldsymbol{\theta}^T, \sigma^2)$$

Trajectory model

- Linear motion can be described by Newton's 2nd law:
  - $\vec{F}_g$: gravitational force; attraction between earth and munition
  - $\vec{F}_d$: drag force; resistive force of an object travelling through a fluid
  - $\vec{F}_l$: lift force; responsible for lateral drift
- Rotational kinematics accounted for:
  - Overturning moment: associated with lift force
  - Spin damping moment: opposes spin of projectile due to aerodynamic skin friction
- We are implementing the Indirect Fires Delivery Accuracy Program (IFDAP) model.
  - Omits Coriolis and Magnus forces as they have negligible effects
  - Fast computation; maintains high fidelity compared to full models.

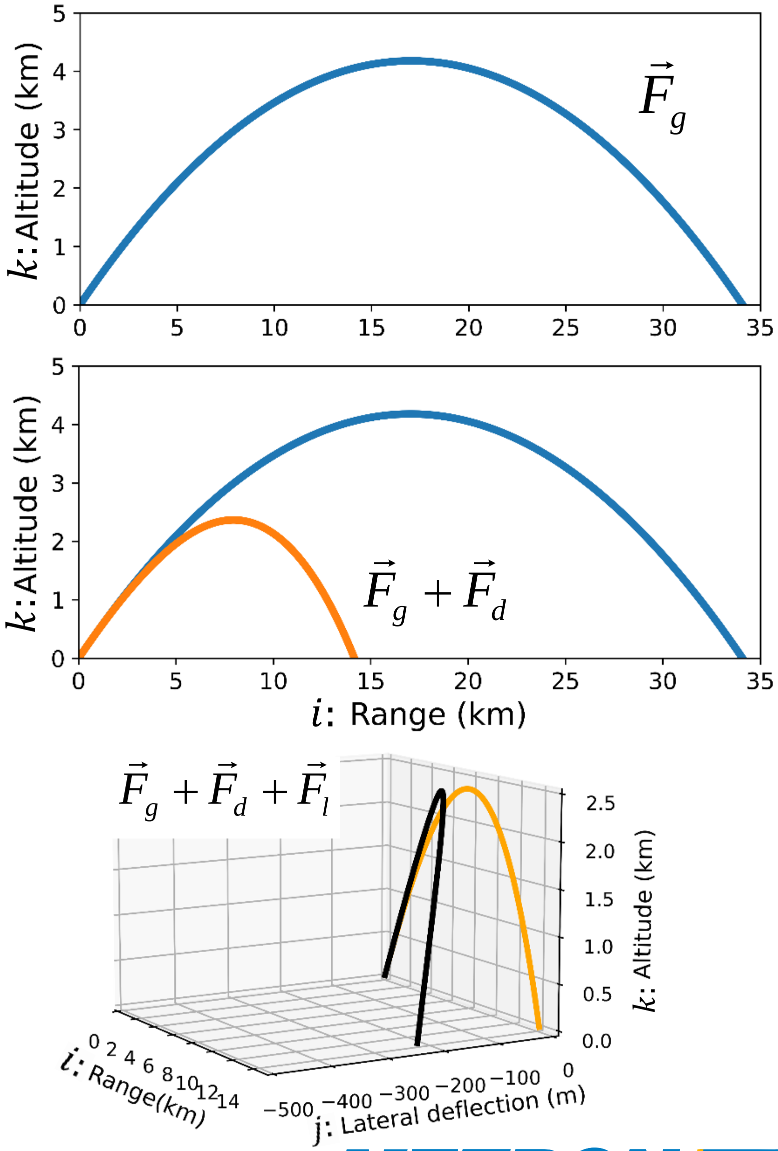

UNCLASSIFIED

M107(HE) 155 mm properties

General Parameters		Weapon Data	
Time Step	0.01 s	Twist rate	20 calibers/rev
Gravitational acceleration	9.81 m/s^2	I_{xx}	0.1461 kg m^3
Air C_p (specific heat at constant pressure)	1005 J/(kg K)	Mass	43.091 kg
Air C_V (specific heat at constant volume)	718 J/(kg K)	Projectile diameter	0.155 m

https://en.wikipedia.org/wiki/M107_projectile

Gravitational Force

$$\vec{F}_g = \frac{GMm}{\|\vec{r}\|^2}\hat{r} \approx -mg\,\hat{y}$$

- F_g : gravitational force: attraction between two objects with mass
- Induces symmetric parabolic motion
  - G : gravitational constant
  - M : mass of earth

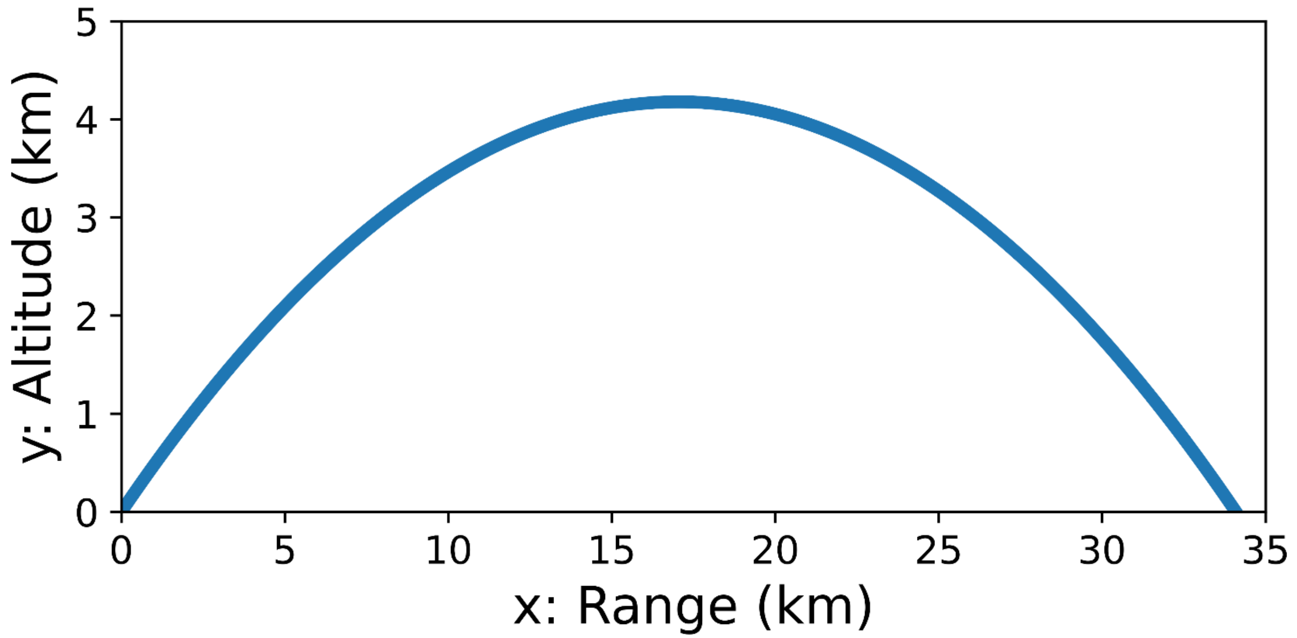

Drag force

$$\vec{F}_d = -\frac{1}{2} S C_d \rho V \vec{V}$$

- Drag force: air resistance; reduces projectile range and causes higher angle of impact
  - S : reference cross-sectional area perpendicular to axis of symmetry
  - C_d : drag force coefficient
  - ρ : density of air
  - $\vec{V}$: projectile velocity vector

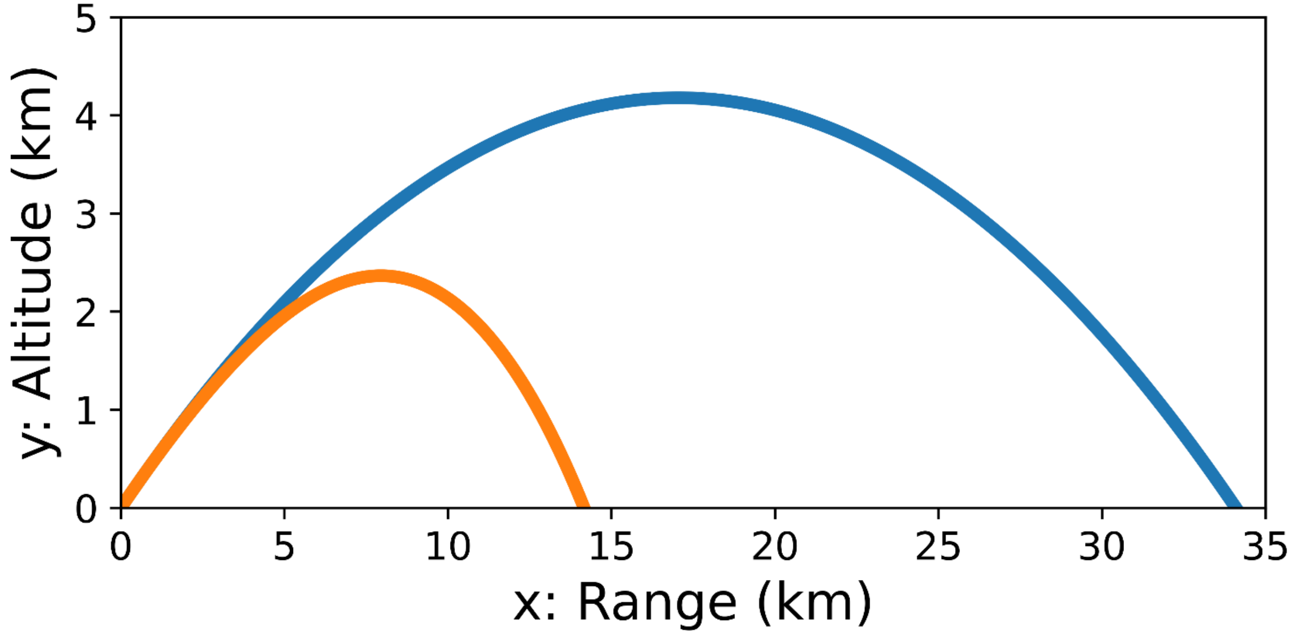

Lift force

$$\vec{F}_l = \frac{1}{2}\rho S C_{L_\alpha}[\vec{V}\times(\hat{j}\times\vec{V})]$$

- Lift force: perpendicular to the trajectory, tending to pull the projectile in the direction its nose is pointed.
  - E.g., if the nose is pointed above the trajectory, the lift force causes the projectile to climb.
  - S : reference cross-sectional area
  - C_{L_α} : drag force coefficient
  - ρ : density of air
  - $\vec{V}$: projectile velocity vector
  - $\hat{j}$: unit vector along projectile axis of symmetry
- The spinning munition precesses about $\hat{V}$, and the z-component of the lift force produces a lateral drift in range.

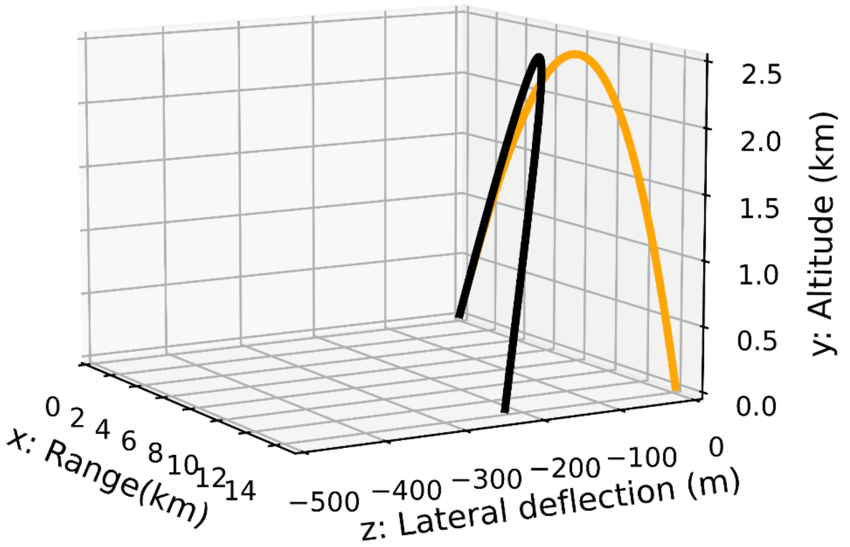

Rotational effects

- Forces acting on projectile induce rotations, which can be written down using Newton's second law for rotation:

$$\sum \vec{\tau} = I\vec{\alpha}$$

- Sum of the torques = (moment of inertia) × (angular acceleration)
- Two significant moments
  - Overturning moment: $\tau_o = \frac{1}{2}\rho S d C_{M_\alpha} V^2 (\hat{V} \times \hat{j})$
    - Associated with lift force
    - If projectile's nose lies above its trajectory, a positive overturning moment acts to increase yaw angle
  - Spin damping moment: $\tau_{sd} = -\frac{1}{2}\rho S d^2 V \omega C_{lp} \hat{j}$
    - ω : axial angular speed (radians/second)
    - d : reference diameter
    - Opposes spin of projectile due to aerodynamic skin friction
    - Always tends to reduce axial spin
- Rotational deflections from the x-y plane (precession) tend to result in lateral deflections.

Decision theory constraints

- Constraints are imposed upon preferences:
  - $Completeness$. Exactly on the following holds: $A \succ B$, $B \succ A$, or $A \sim B$.
  - $Transitivity$. If $A \succcurlyeq B$ and $B \succcurlyeq C$, then $A \succcurlyeq C$.
  - $Continuity$. If $A \succcurlyeq C \succcurlyeq B$, then there exists a probability p such that $\mathrm{p}A + (1-p)B \sim C$.
    - Implies you cannot have a discontinuous jump in preferences once you encounter uncertainty
    - Violated with lexicographic preferences
  - $Independence$. If $A \succ B$, then for any C and probability p, $[A\colon p; C\colon 1-p] \succcurlyeq [B\colon p; C\colon 1-p]$.
- These are the constraints on rational preferences. If follows from this that there exists a real-valued function U such that
  - $U(A) > U(B)$ iff $A \succ B$.
  - $U(A) = U(B)$ iff $A \sim B$.