Modeling secrets

Mário S. Alvim
(msalvim@dcc.ufmg.br)

Quantitative Information Flow

DCC-UFMG
(2021/1)

Introduction

- We begin our detailed study of quantitative information flow by asking:

  What is really a secret?

- We can cite some non-controversial examples:

  (a) a user's password,

  (b) social security number,

  (c) current location,

  (d) etc. . .

- Here we'll deal with what does it mean for us to treat them as secrets.

Secrets and probability distributions

Secrets and probability distributions

- Let's start by defining secrets.

Definition (2.1)

A **secret**, called X, has a set $\mathcal{X}$ of possible values, which we call the **type** of X and which we assume to be **finite**.

Moreover we assume that the knowledge that some adversary has about X is given by a **probability distribution** π on $\mathcal{X}$ that specifies the probability π_x of each possible value x of X.

(Thus the secrecy of X is defined relative to a **particular** adversary, leaving open the possibility that a **different** adversary might have different knowledge of X, and so require a different distribution to reflect that.)

Secrets and probability distributions: Terminology and notation

- A **probability distribution** $\delta: \mathcal{X} \to [0,1]$ on a finite set $\mathcal{X}$ is a function from $\mathcal{X}$ to the interval $[0,1]$ s.t.:

$$\sum_{x \in \mathcal{X}} \delta_x = 1.$$

  We write $\mathbb{D}\mathcal{X}$ for the **set of all distributions on** $\mathcal{X}$.

- We sometimes denote a distribution simply as a tuple of its probabilities.
  (a) E.g., with $\mathcal{X} = \{x_1, x_2, x_3, x_4\}$, we might write $\delta = (1/2, 1/6, 0, 1/3)$.

- The **support** $\lceil \delta \rceil$ of a distribution is the set of values to which it gives nonzero probability.
  (a) The δ above has support $\{x_1, x_2, x_4\}$.

  A distribution δ in $\mathbb{D}\mathcal{X}$ is said to be **full support** if $\lceil \delta \rceil = \mathcal{X}$.

Secrets and probability distributions: Prior

- We refer to π as the **prior distribution**, because it represents the adversary's knowledge about the system before any output is observed.

  (Later we will consider **posterior distributions**, which represent the adversary's knowledge after observing the output of a system.)

- We now consider some examples of prior distributions.

Secrets and probability distributions: Prior

- If X is generated by a known probabilistic process, then π is clear.
- Example 1 (**3 coin flips**) If X is generated by 3 flips of a fair coin, then X has 8 equally probable values, giving a **uniform** prior π^{coin}:

HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$

◁

Secrets and probability distributions: Prior

- Example 2 (**Sum of 2 fair dice**) If X is generated as the sum of a roll of two fair dice, then we get a **non-uniform** prior π^{dice}:

2	3	4	5	6	7	8	9	10	11	12
$1/36$	$2/36$	$3/36$	$4/36$	$5/36$	$6/36$	$5/36$	$4/36$	$3/36$	$2/36$	$1/36$

◁

Secrets and probability distributions: Prior

- The appropriate prior distribution for secrets in general can be more doubtful.

- Example 3 (**Mother's maiden name**) If X is the mother's maiden name of a certain user, then X is not generated by a clear random process.

  In that case, it seems more appropriate for the prior distribution to reflect the adversary's knowledge of the population that the user comes from.

  E.g., if the adversary knows only that the user is an American born in the second half of the 20$^{\text{th}}$ century, then the prior distribution π^{mother} can be inferred on the 2000 Census data from the USA.

  In this case the space $\mathcal{X}$ of possible values is very large: there are 151,671 different surnames occurring at least 100 times.

  This leads to a prior π^{mother} as follows.

Smith	Johnson	Williams	Brown	Jones	Miller	Davis	Garcia	...
0.00881	0.00688	0.00569	0.00512	0.00505	0.00418	0.00398	0.00318	...

Secrets and probability distributions: Prior

- Example 4 (**Location on Saturday night**) If X is a certain young person's location on Saturday night, then an appropriate prior distribution would be quite hard to determine.

  But it could in principle be based on the popularity of various bars, clubs, and other activities in the population the person comes from. ◁

Secrets and probability distributions: Secrecy

- We have illustrated various **kinds** of secrets.

  Now let's consider the **quantification of secrecy**.

- Given a secret X with (prior) distribution π, it is natural to say that the "amount" of secrecy X has is determined by π.

  Note that, intuitively:

  - A uniform π should mean "more" secrecy.
  - A skewed π should mean "less" secrecy.

    In particular, if π is a **point distribution** $[x]$ on some x (meaning that $\pi_x = 1$), then there is no secrecy at all.

- To quantify secrecy, we can measure:

  - the adversary's "uncertainty" about the secret (the lower the better for the adversary); or
  - the 'threat" to the secret (the higher the better for the adversary).

  These perspectives are complementary, and we'll alternate between them.

Shannon entropy

Shannon entropy

- From the "uncertainty", it's tempting to measure the secrecy of an X having distribution π as the **Shannon entropy** $H(\pi)$, defined by

$$H(\pi) \;:=\; -\sum_{x \in \mathcal{X}} \pi_x \log_2 \pi_x \quad .$$

- Note that if $|\mathcal{X}| = n$ then the possible values of $H(\pi)$ range:
  - from 0, in the case of a point distribution,
  - up to $\log_2 n$, in the case of a uniform distribution.

  But this alone does not imply that different values of Shannon entropy are associated with definite security guarantees.

- Every measure needs an **operational significance**, which is an interpretation of what the numbers it gives means in real life.

Any measure is only as relevant as its operational significance!

Shannon entropy: Operational significance

- **Shannon's source-coding theorem** gives an operational significance to Shannon entropy:

  "$H(\pi)$ is the average number of bits required to transmit X under an optimal encoding."

- Example 5 Suppose that $\mathcal{X} = \{x_1, x_2, x_3, x_4\}$ and $\pi = (1/2, 1/4, 1/8, 1/8)$.

  We can encode X using a **Huffman code**, which uses shorter codes for more likely values, and longer codes for less likely values:

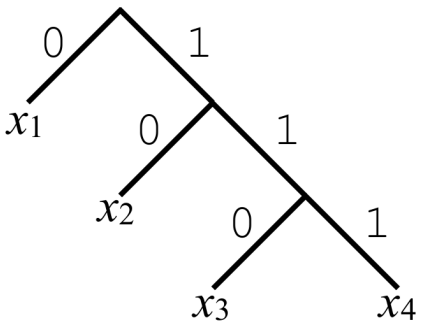

$x_1 : 0$
$x_2 : 10$
$x_3 : 110$
$x_4 : 111$

Under this code, the expected number of bits required to transmit X is

$$1/2 \cdot 1 + 1/4 \cdot 2 + 1/8 \cdot 3 + 1/8 \cdot 3 = 7/4,$$

whereas

$$H(\pi) = 1/2 \log_2 2 + 1/4 \log_2 4 + 1/8 \log_2 8 + 1/8 \log_2 8 = 7/4.$$

Shannon entropy: Operational significance

- When $H(\pi)$ is large, many bits are required on average to transmit X.

  We might then think that the adversary has a great deal of "uncertainty" about X, and that there is little "threat" to X.

- More precisely, if we recall that

Distribution on X	**Shannon entropy $H(X)$**	**Adversary's chance of guessing X in one try**
Point: $[x]$	0	1
Uniform: $(1/n, \ldots, 1/n)$	$\log_2 n$	$1/n$

  then we might conjecture that Shannon entropy has the following operational significance for confidentiality.

Conjecture (2.2)

If a secret X has distribution π, then an optimal adversary's probability of guessing the value of X correctly in one try is at most $2^{-H(\pi)}$.

Shannon entropy: Operational significance

- But this conjecture is false!

  The following example illustrates that.

- Example 6 Consider a secret X with $\mathcal{X} = \{x_1, x_2, \ldots, x_n\}$.

  With a uniform π we have

$$H(\pi) = \log_2 n \ .$$

  But a smart adversary's probability of correctly guessing X in one try is

$$1/n = 2^{-\log_2 n} = 2^{-H(\pi)} \ .$$

Shannon entropy: Operational significance

- Example 6 (Continued)

But now suppose that we shift half of the mass from x_2 through x_n to x_1.

The result is a distribution π' where $\pi'_{x_1} = {}^{(n+1)}/_{2n}$ and $\pi'_{x_2} = \cdots \pi'_{x_n} = {}^{1}/_{2n}$.

This change increases the adversary's chance of correctly guessing X in one try to over ${}^{1}/_{2}$.

But we find that $H(\pi')$ is still about half as big as $H(\pi)$ is, as for large n:

$$\begin{aligned} H(\pi') &= \frac{n+1}{2n}\log_2\left(\frac{2n}{n+1}\right) + \frac{n-1}{2n}\log_2(2n) \\ &= \left(\frac{n+1}{2n} + \frac{n-1}{2n}\right)\log_2(2n) - \frac{n+1}{2n}\log_2(n+1) \\ &= \log_2 n - \frac{n+1}{2n}\log_2(n+1) + 1 \\ &\approx \frac{1}{2}\log_2 n + 1 \end{aligned}$$

Shannon entropy: Operational significance

- Example 6 (Continued)

  Concretely, if X is a 128-bit secret, so that $n = 2^{128}$, then we find that

$$H(\pi') \approx \frac{1}{2}\log_2 2^{128} + 1 = 65 ,$$

  so the conjecture would predict a guessing probability of at most 2^{-65} .

  But the adversary's guessing probability is indeed

$$\frac{n+1}{2n} \approx \frac{1}{2} = 2^{-1} ,$$

  which is far above the prediction. ◁

- This example shows that Shannon entropy can be quite misleading wrt. confidentiality!

  We'll consider other measures of secrecy with better operational significance.

Bayes vulnerability

Bayes vulnerability

- Let's refocus on a more basic security question:

  "What is the adversary's probability of guessing X correctly in one try?"

  Note that we are now switching our perspective from the adversary's "uncertainty" about X to the "threat" to X.

- This leads to the definition of what we call Bayes vulnerability.

Definition (2.3)

Given a distribution π on a finite set $\mathcal{X}$ (viewed as the distribution of possible values of a secret X) the **Bayes vulnerability** of π is defined by

$$V_1(\pi) \ := \ \max_{x \in \mathcal{X}} \pi_x \quad .$$

Bayes vulnerability

- Example 7 Recall the examples from before.

  - **3 coin flips:** $V_1(\pi^{coin}) = 1/8$ (all guesses are equally good)

HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$

  - **Sum of 2 dice:** $V_1(\pi^{dice}) = 6/36$ (7 is the best guess)

2	3	4	5	6	7	8	9	10	11	12
$1/36$	$2/36$	$3/36$	$4/36$	$5/36$	$6/36$	$5/36$	$4/36$	$3/36$	$2/36$	$1/36$

  - **Mother's maiden name:** $V_1(\pi^{mother}) = 0.00881$ (Smith is the best guess).

Smith	Johnson	Williams	Brown	Jones	Miller	Davis	Garcia	...
0.00881	0.00688	0.00569	0.00512	0.00505	0.00418	0.00398	0.00318	...

Bayes vulnerability: Operational significance

- Bayes vulnerability has a clear operational significance wrt. confidentiality.

  $V_1(\pi)$ is adversary $\mathcal{A}$'s maximum probability of winning the following game:

> *-- Defender*
> $X :\in \pi$ *-- Choose X according to prior.*
>
> *-- and then Adversary*
> $W :\in \mathcal{A}(\pi)$ *-- Choose guess, and...*
> `if` $W{=}X$ `then` "win" `else` "lose" `--` *...win if correct.*

More informally,

"Bayes vulnerability is the probability of an optimal adversary guessing the secret in one try."

Bayes vulnerability: Adversary capabilities

- Note that we assume the adversary $\mathcal{A}$ has access to π.

  This means that, at a minimum, that $\mathcal{A}$ has the ability to compute π_x for any x in $\mathcal{X}$.

  But notice that to achieve the Bayes vulnerability in the game above, $\mathcal{A}$ actually needs to find a value x such that π_x is maximal.

  It is of course possible for $\mathcal{A}$ to do that by searching through all probabilities, but it might require "finding a needle in a haystack".

  In this course we ignore such questions of efficiency, focusing only on the **information** available to the adversary.

Throughout this course, we assume that adversaries are information theoretic — that is, without computational resource bounds.

A more general view

A more general view

- Bayes vulnerability's clear operational significance is a strength, but it can also be seen as a limitation.
- There are many possible operational scenarios, by which we mean not only the "rules" about what the adversary is allowed to do but also the adversary's particular goals.

  Depending on the scenario, Bayes vulnerability might or might not do a good job of measuring the threat to the secret.

A more general view

- A more general approach is to imagine an adversary's trying to decide among some possible actions to take in order to exploit her knowledge about X.

  For each possible action and each possible secret value, there will be some "reward" that she would achieve.

- Example 8 We can see that Bayes vulnerability corresponds to a scenario where the adversary is required to make a guess (perhaps by entering it into a keypad) and is rewarded only if that one guess is exactly correct. ◁

A more general view

- Example 9 (**Three tries**) Consider a login scenario where a user is allowed three tries to enter the correct password before she is locked out.

  An adversary's possible actions would consist of sets of three passwords to be tried, and she would be rewarded if any of the three were correct.

  Here it would be more accurate to model the threat using "three tries" Bayes vulnerability, defined as the sum of the 3 largest probabilities assigned by π.

A more general view

- Example 9 (Continued)

Denoting that by V_3, we have the following.

- **3 coin flips:** $V_3(\pi^{coin}) = \pi^{coin}_{HHH} + \pi^{coin}_{HHT} + \pi^{coin}_{HTH} = 3/8$,

HHH	HHT	HTH	HTT	THH	THT	TTH	TTT
$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$	$1/8$

- **Sum of 2 dice:** $V_3(\pi^{dice}) = \pi^{dice}_7 + \pi^{dice}_6 + \pi^{dice}_8 = 6/36 + 5/36 + 5/36 = 4/9$

2	3	4	5	6	7	8	9	10	11	12
$1/36$	$2/36$	$3/36$	$4/36$	$5/36$	$6/36$	$5/36$	$4/36$	$3/36$	$2/36$	$1/36$

- **Mother's maiden name:** $V_3(\pi^{mother}) = \pi^{mother}_{\text{Smith}} + \pi^{mother}_{\text{Johnson}} + \pi^{mother}_{\text{Williams}} = 0.02138$.

Smith	Johnson	Williams	Brown	Jones	Miller	Davis	Garcia	...
0.00881	0.00688	0.00569	0.00512	0.00505	0.00418	0.00398	0.00318	...

A more general view

- There are many ways to measure "uncertainty" or "threat" for a secret X, and they depend on the operational scenario.

  Is there a single framework that can encompass them all?

- Indeed there is!

  We'll see it in Chapter 3, when we develop the framework of g-**vulnerability**.