Risk assessment with AILA:

Automated and Intelligent Likelihood Assignment

20th Workshop on Security Frameworks
"Security Testing"

20/12/2022 - Catania

Giampaolo Bella, *Università di Catania*
Cristian Daniele, *Radboud University*
Mario Raciti, *Scuola IMT Alti Studi di Lucca*

Risk Management in a Nutshell

"If you don't invest in risk management, it doesn't matter what business you're in, it's a risky business." - Gary Cohn

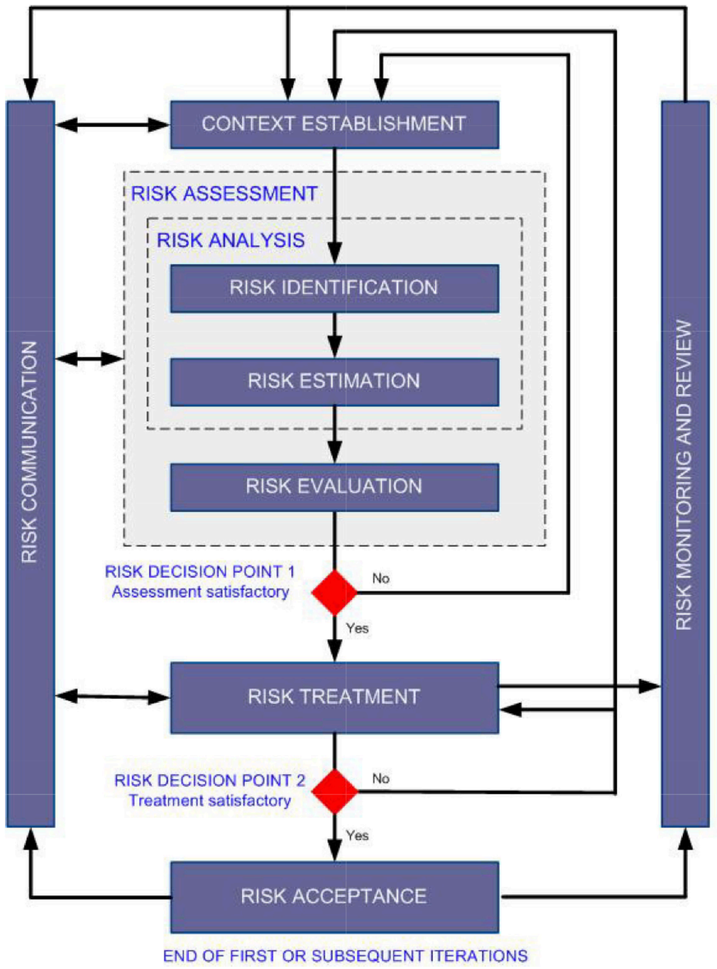

ISO 27005

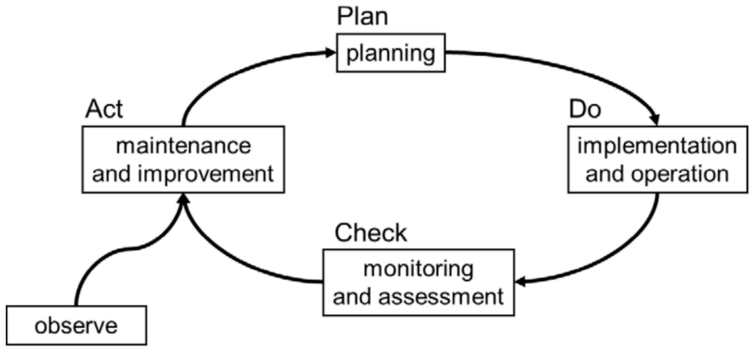

ISMS PDCA Cycle [ISO 27001]

Risk Assessment Concepts

RA inputs:

- Assets
- Threats
- Safeguards

RA outputs:

- Impact
- Risk

Risk		*Likelihood*				
		VL	L	M	H	VH
Impact	VH	H	VH	VH	VH	VH
	H	M	H	H	VH	VH
	M	L	M	M	H	H
	L	VL	L	L	M	M
	VL	VL	VL	VL	L	L

Other factors:

- Security dimensions
- Likelihood

Risk for dummies $R = L \times I$

Actual risk $R = \ldots ?$

where R is the risk, L the likelihood and I the impact.

Privacy Risks

A variety of **personal data** is collected by services to "*improve*" the user's experience.

OWASP Top 10 Privacy Risks

P1 - Web Application Vulnerabilities
P2 - Operator-sided Data Leakage
P3 - Insufficient Data Breach Response
P4 - Consent on Everything
P5 - Non-transparent Policies, Terms and Conditions
P6 - Insufficient Deletion of User Data
P7 - Insufficient Data Quality
P8 - Missing or Insufficient Session Expiration
P9 - Inability of Users to Access and Modify Data
P10 - Collection of Data Not Required for the User-Consented Purpose

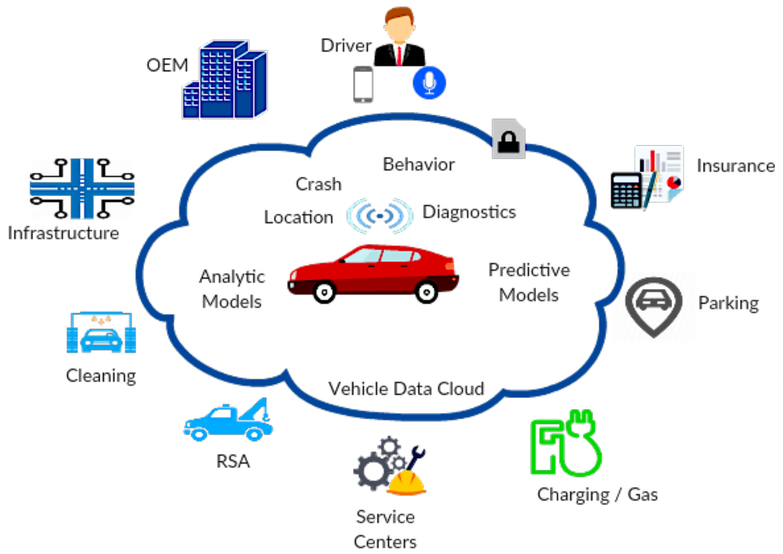

He's making a list
He's checking it twice
He's gonna find out who's naughty or nice
Santa Claus is in contravention of the GDPR (EU) 2016/679

Privacy Policy RA

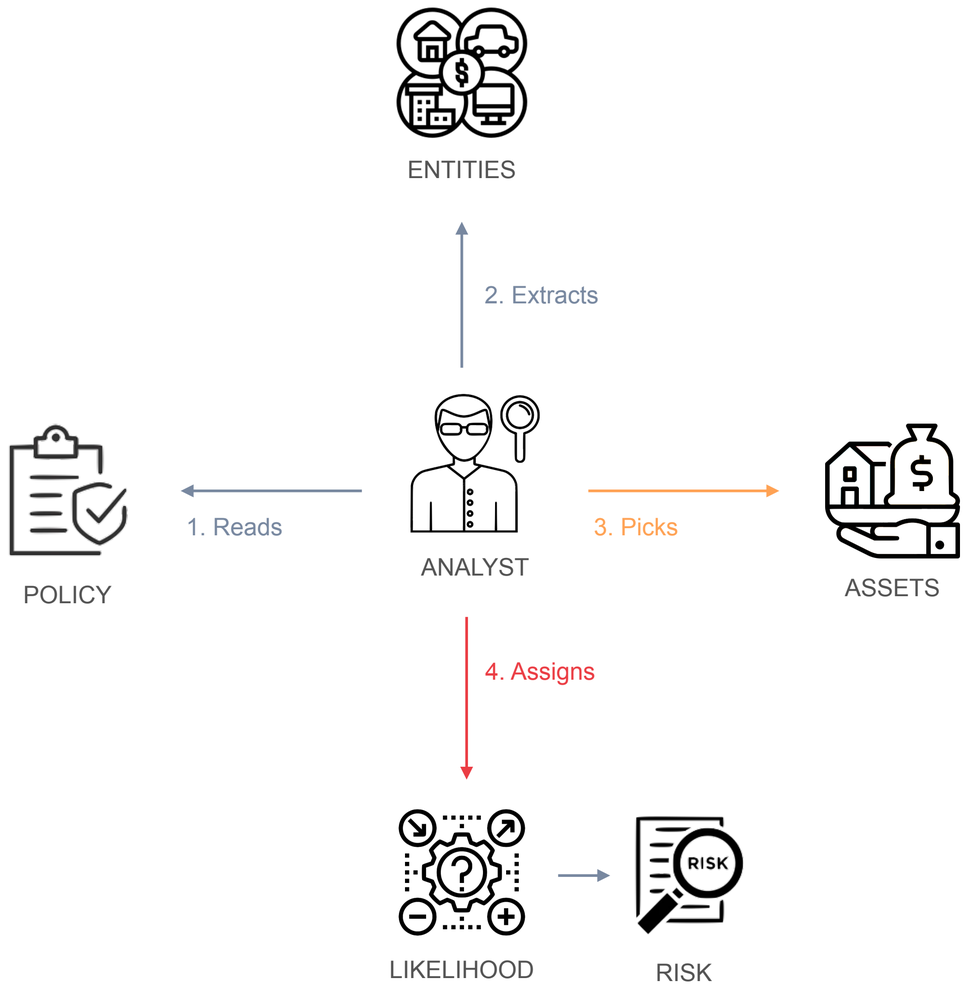

AILA - Automated and Intelligent Likelihood Assignment

AILA aims at reducing the influence of ***subjectivity*** and ***distraction***.

AILA uses *Natural Language Processing* and *Machine Learning*.

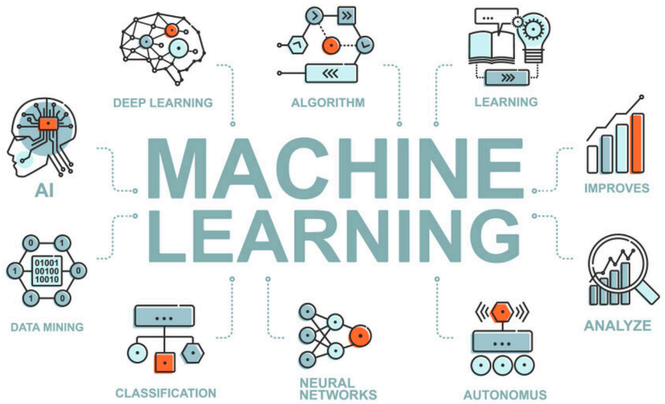

The process is also integrated with a *RA tool*.

AILA in a Nutshell

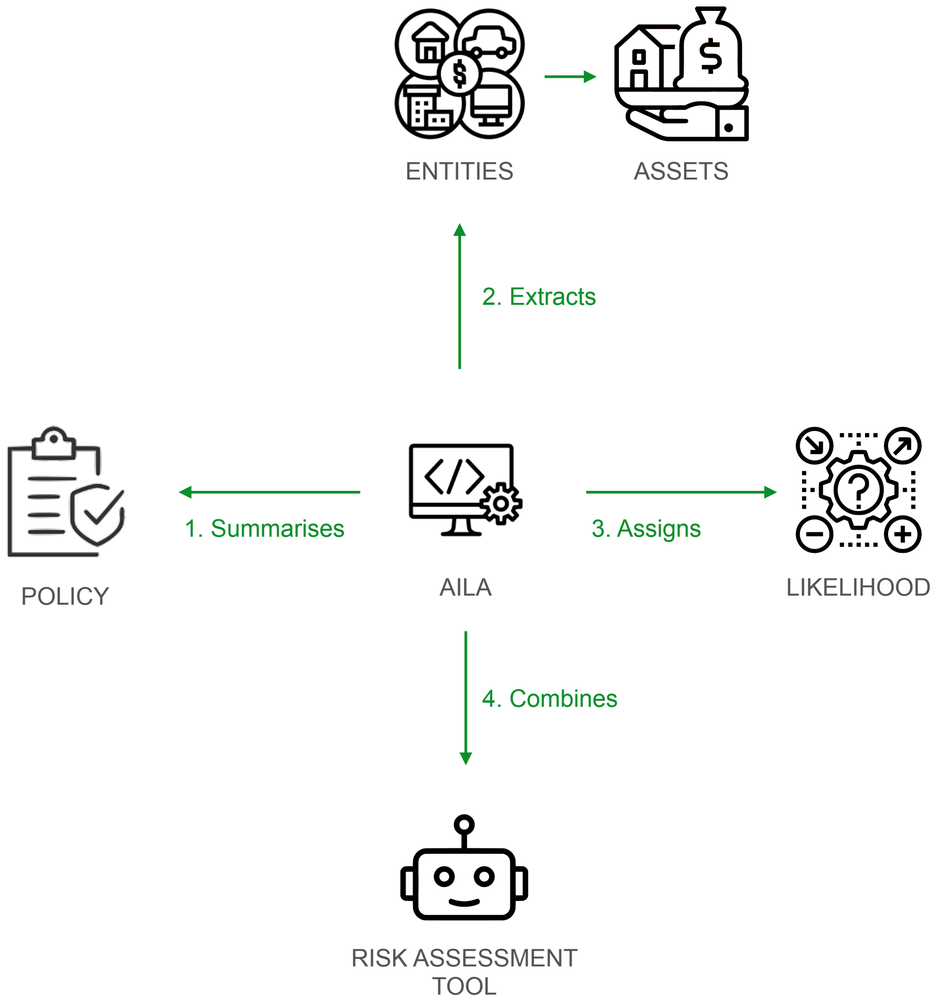

AILA Methodology

1. AUTOMATED ASSET EXTRACTION

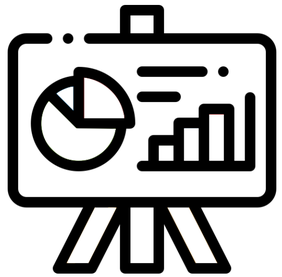

2. LIKELIHOOD DETERMINATION THROUGH AILA

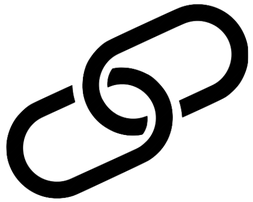

3. COMBINED LIKELIHOOD DETERMINATION

AILA Methodology

1. AUTOMATED ASSET EXTRACTION

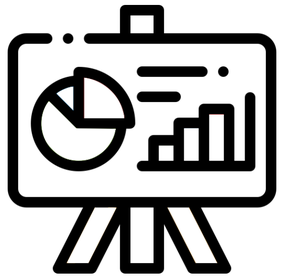

2. LIKELIHOOD DETERMINATION THROUGH AILA

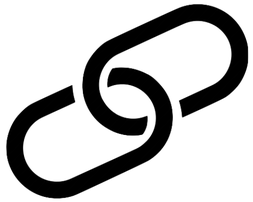

3. COMBINED LIKELIHOOD DETERMINATION

1. Automated Asset Extraction

This step is carried out by the **AILA Entity Extractor (AILAEE)**.

SUMMARISATION
USING N-GRAMS

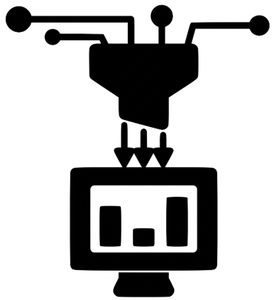

ENTITY RECOGNITION
ALGORITHM

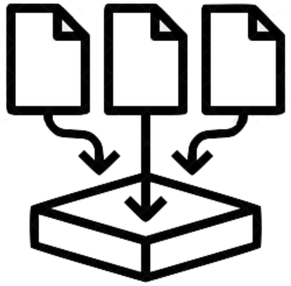

COLLECTION OF
POLICY SENTENCES

1. AILAEE - Preprocessing

> Bigrams Identification

This is a sentence. → [This, is], [is, a], [a, sentence], [sentence, .]

> Sentence Extraction

*Sentences containing **verbs** are the coolest ones!* → POS TAGGING → *I* (PRON) *love* (VERB) *AILA* (NOUN)

> Text Summarisation

The winners are those sentences containing the most frequent words.

> Entropy Measurement

Is the loss of information negligible? → **Shannon's Entropy**

NLTK

spaCy

1. AILAEE - Preprocessing

> Bigrams Identification

This is a sentence.

[This, is],
[is, a],
[a, sentence],
[sentence, .]

The aim is to **remove irrelevant text** and only keep the most relevant components (*nouns, verbs, adjectives*).

This is a sentence.

[This],
[is],
[a],
[sentence],
~~[.]~~

NLTK.WORD_TOKENIZE

[This],
[is],
[a],
[sentence]

[This],
[is],
~~[a]~~,
[sentence]

NLTK.BIGRAMCOLLOCATIONFINDER

[This, is],
[is, sentence]

1. AILAEE - Preprocessing

> Sentence Extraction

*Sentences containing **verbs** are the coolest ones!* ⇨ POS TAGGING

I *love* *AILA*
PRON *VERB* *NOUN*

The aim is to **keep sentences containing a verb**.

*Sentences containing **verbs** are the coolest ones!* ⇨ NLTK.POS_TAG

Sentences *containing* *verbs* *are* *the* *coolest* *ones* *!*
NOUN *VERB* *NOUN* *VERB* *DT* *ADJ* *NOUN*

For each bigram, it extract the original sentence containing it and the adjacent sentences.

1. AILAEE - Preprocessing

> Text Summarisation

The winners are those sentences containing the most frequent words.

The aim is to **summarise the text to improve entity recognition**.

For each sentence linked to a certain bigram:

1. *it tokenises all sentences and calculates the frequencies of each word;*
2. *it calculates the score of each sentence by adding up the frequencies of the words in the sentence;*
3. *it extracts the sentence with the greatest score.*

1. AILAEE - Preprocessing

> Entropy Measurement *Is the loss of information negligible?* → **Shannon's Entropy**

The aim is to **avoid loss of information** during summarisation.

It evaluates whether the loss of information between the original set of sentences and the chosen sentence is negligible.

1. AILAEE - Named Entity Recognition

> Sentence Tokenisation

I am pretending to be a great long text. This is a sentence. Another sentence here. Are you enjoying the talk?

NLTK.SENT_TOKENIZE

[I am pretending to be a great long text],
[This is a sentence],
[Another sentence here],
[Are you enjoying the talk]

> Entity Recognition

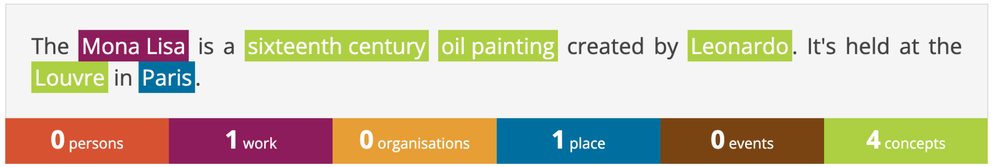

DANDELION API

RETURNS

Entities and synonyms!

> Sentence Gathering

For each entity, we gather all the sentences containing the entity or its synonyms in the original text.

AILA Methodology

1. AUTOMATED ASSET EXTRACTION

2. LIKELIHOOD DETERMINATION THROUGH AILA

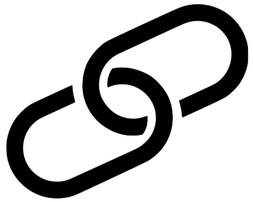

3. COMBINED LIKELIHOOD DETERMINATION

2. Likelihood Determination Through AILA

The **AILA Classifier (AILAC)** addresses this second challenge.

2. AILAC - Fairness

It indicates how **fair**, proper and clean a text is, regarding the *users' privacy concerns*.

Fairness per asset	AILA Likelihood	
0 - 0.20	5	**VH**
0.21 - 0.40	4	**H**
0.41- 0.60	3	**M**
0.61- 0.80	2	**L**
0.81- 1	1	**VL**

AILA LIKELIHOOD DEFINITION

AILA Likelihood for dummies

$$L = 1 - F$$

2. AILAC - Dataset

We got 500 sentences enriched with *text augmentation* and *synonyms*.

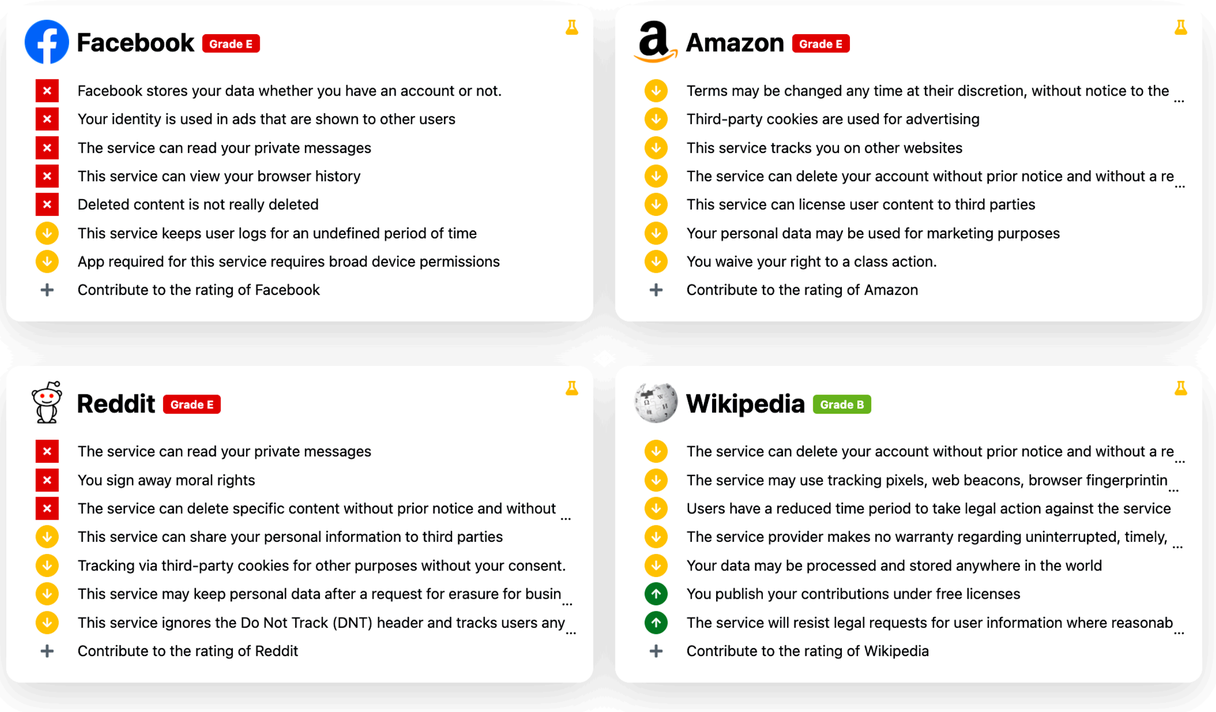

Terms of Service Didn't Read

AILAC corpus consists of over 100.000 labelled sentences.

2. AILAC - Model

The dataset is split into 2 parts:
- *75% training*
- *25% testing*

The labelled sentences are transformed into a *2-D feature matrix*.

The *Relu function* is chosen as activation function of the first layer, the *Sigmoid function* for the second layer and the *Adam function* for the optimisation.

The model gets trained for *15 epochs* using the *binary cross-entropy function* as loss function, 0.0001 as learning rate, and 50 as batch size.

The model has an **accuracy of 96%**.

2. AILAC - Model

The model is used to evaluate the **fairness** of the sentences extracted in the previous step.

For each entity, it calculates the fairness of the related sentences and assign the **mean fairness** to the entity.

E.g.

Asset: ***data***

Sentences: ***"****We will sell all your* ***data****", "We will save your* ***data*** *at any cost"*

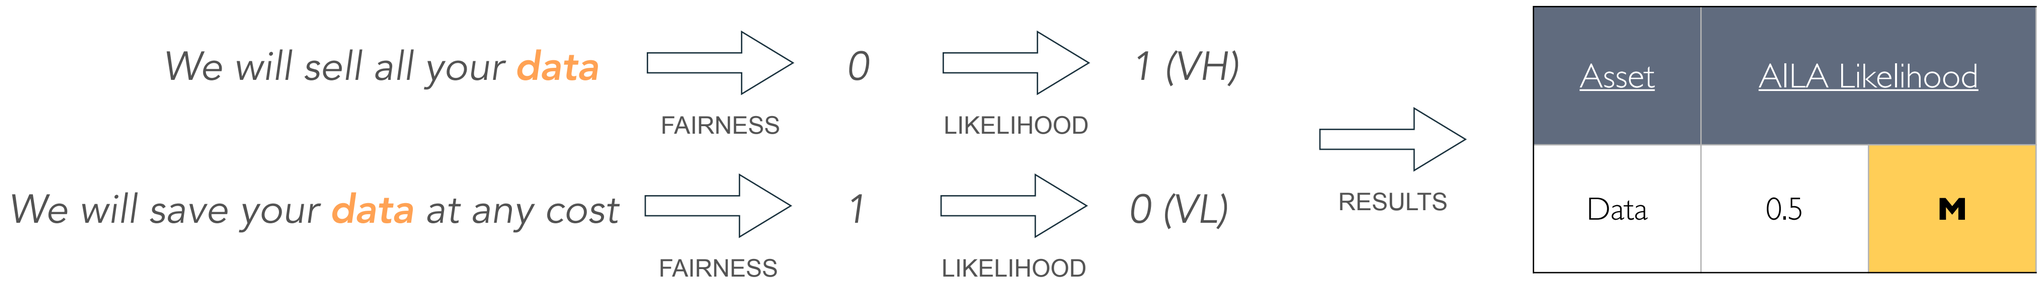

2. AILAC - Outputs Sample

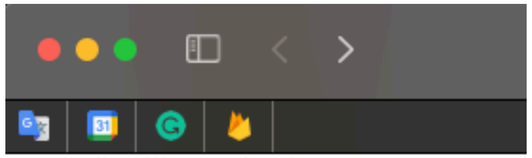

- Satellite navigation
  - Fairness mean: 0.39
  - Likelihood: M
- Web portal
  - Fairness mean: 0.93
  - Likelihood: VL
- Vehicular communication systems
  - Fairness mean: 0.05
  - Likelihood: VH
- Vehicle tracking system
  - Fairness mean: 0.47
  - Likelihood: M
- Mobile app
  - Fairness mean: 0.44
  - Likelihood: M
- Data collection
  - Fairness mean: 0.49
  - Likelihood: M

LIST OF ENTITIES

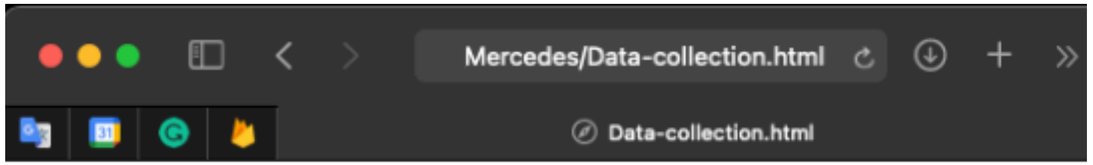

	Data collection	Label
0	Your Choices Opt Out of data collection You can opt out of the collection of data via certain Mercedes me connect services by deactivating those specific services through the Mercedes me connect portal.	1.000000
1	Contact Mercedes me connect Support for additional information on how to opt out of data collection or deactivating services: 888- 628-7232 or me-connect.usa@cac.mercedes-benz.com.	0.980000
2	This Privacy Notice does not address the collection, use, or sharing of information regarding how you use or interact with our websites and mobile applications.	0.090000
3	Live Traffic, Navigation, Concierge, Car-to-X communications, Assist Services, Parked Vehicle Locator, Vehicle Tracker, Geofencing, Route Planning, Mercedes-Benz Apps, and Product Improvement services, for example, involve the collection of Geolocation Information to determine the location of your vehicle.	0.110000
4	In addition, some specific Mercedes me connect services may involve the collection and use of Driving Behavior Information.	0.260000

LIST OF SENTENCES FOR A SELECTED ENTITY

AILA Methodology

1. AUTOMATED ASSET EXTRACTION

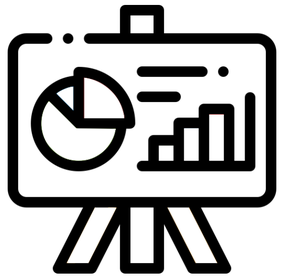

2. LIKELIHOOD DETERMINATION THROUGH AILA

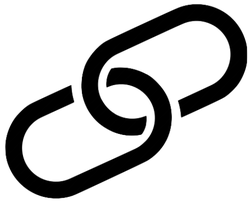

3. COMBINED LIKELIHOOD DETERMINATION

3. Combined Likelihood Determination

AILA Likelihood can be used to *sculpt* the Likelihood outputted by a standard tool on a specific privacy policy.

We combined AILA with *PILAR*.

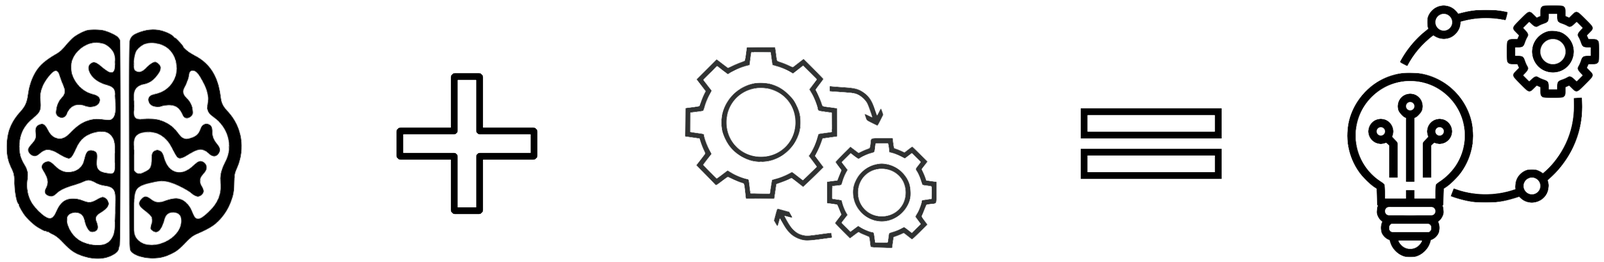

PILAR Reverse Engineering

Impact $I = V \times d$

where I is the impact, V the asset value and d the degradation.

PILAR Impact $I = V - \delta$ where $\delta = \begin{cases} 6 & \text{if } d = 1\,\% \\ 3 & \text{if } d = 10\,\% \\ 2 & \text{if } d = 20\,\% \\ 1 & \text{if } d = 50\,\% \\ 0 & \text{if } d = 100\,\% \end{cases}$

Exponential fit $y = 1002.75e^{0.767241x}$ with $r = 0.99$

E.g. $V = 6\ (= 100000),\ d = 20\,\%$

$I = V - \delta = 6 - 2 = 4$

$I = V \times d = 100000 \times 20\,\% = 20000 \simeq_{(Exp\ fit)} 3.9 \simeq 4$

Level	Value
0	1000
1	2150
2	4650
3	10000
4	21500
5	46500
6	100000
7	215000
8	465000
9	1000000
10	2150000

PILAR Levels Map

PILAR Reverse Engineering

PILAR Conjectured Risk

$$R = 0.6I + \lambda$$

where R is the risk, I the impact and $\lambda = \begin{cases} -0.9 & \text{if } L = VL \\ 0 & \text{if } L = L \\ 0.9 & \text{if } L = M \\ 1.8 & \text{if } L = H \\ 2.7 & \text{if } L = VH \end{cases}$

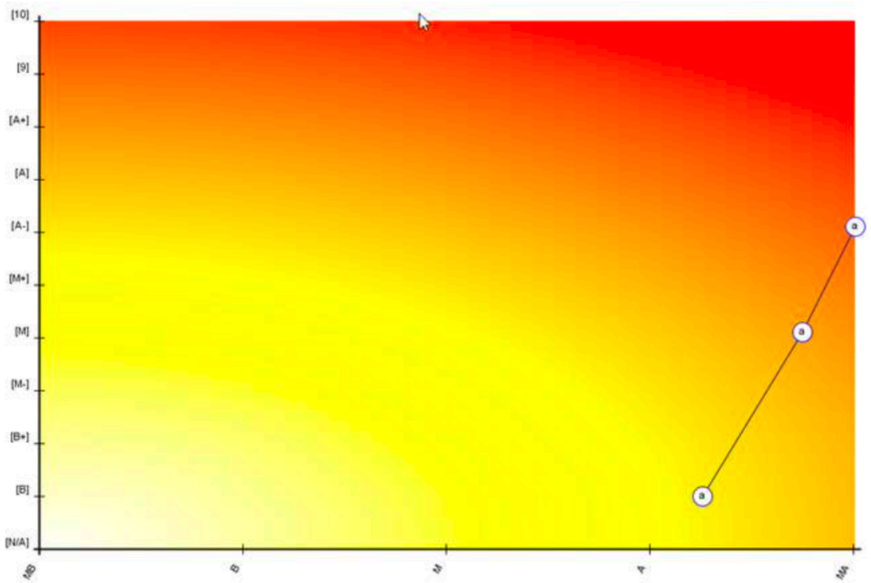

PILAR Heat Map

Risk	-0,9	0	0,9	1,8	2,7
10	5,1	6	6,9	7,8	8,7
9	4,5	5,4	6,3	7,2	8,1
8	3,9	4,8	5,7	6,6	7,5
7	3,3	4,2	5,1	6	6,9
6	2,7	3,6	4,5	5,4	6,3
5	2,1	3	3,9	4,8	5,7
4	1,5	2,4	3,3	4,2	5,1
3	0,9	1,8	2,7	3,6	4,5
2	0,3	1,2	2,1	3	3,9
1	0	0,6	1,5	2,4	3,3
0	0	0	0,9	1,8	2,7

PILAR Conjectured Map

PILAR Reverse Engineering

Linear fit $y = 0.97x + 0.15$ with $r = 0.9909792073$

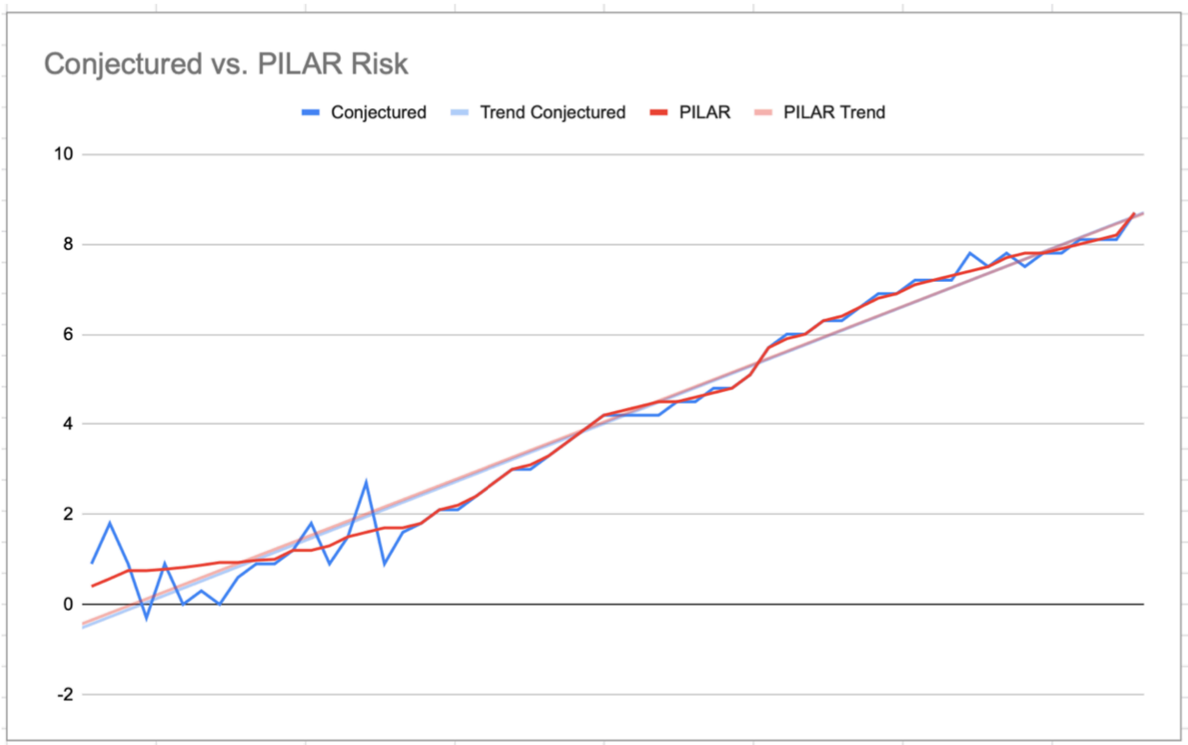

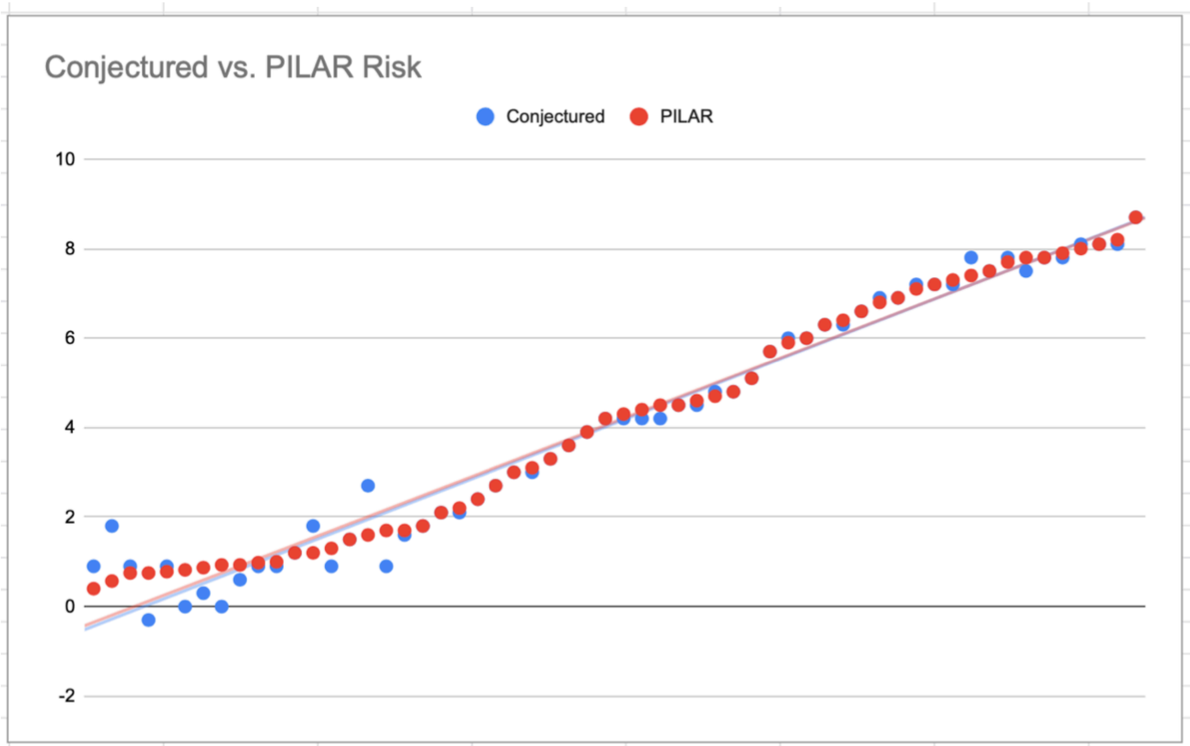

Case Study - Automotive

Mercedes-Benz

Toyota and **Mercedes** were the first two car brands in *Interbrand's 2020 Best Global Brands (BGB) Report*.

Tesla has a pioneer role on electric cars.

Case Study - Assets Extracted from Policies

Policy	Original words	Words after summarisation	Entities	Assets
Toyota	3526	768	**52**	**19**
Mercedes	1800	402	**57**	**17**
Tesla	6860	1164	**72**	**21**

Case Study - AILA Results

PILAR Class	AILA asset	PILAR Threat	Toyota Privacy Policy			Mercedes Privacy Policy			Tesla Privacy Policy		
			PILAR Likelihood	AILA Likelihood	Combined Likelihood	PILAR Likelihood	AILA Likelihood	Combined Likelihood	PILAR Likelihood	AILA Likelihood	Combined Likelihood
Software	Application	Hardware or software failure	**3**	**4**	**3.6**	**3**	**3**	**3.1**	**3**	**4**	**3.6**
		Software vulnerabilities	**3**			**3**			**3**		
		Defects in software maintenance / updating	**4**			**4**			**4**		
		Malware diffusion	**3**			**3**			**3**		
		Software manipulation	**3**			**3**			**3**		
Communication	Location	Accidental alteration of the information	**3**	**5**	**4**	**3**	**3**	**3**	**3**	**3**	**3**
		Information leaks	**3**			**3**			**3**		
		Unauthorised access	**3**			**3**			**3**		
		Traffic analysis	**3**			**3**			**3**		
		Deliberate alteration of information	**3**			**3**			**3**		
		Destruction of information	**3**			**3**			**3**		

RESULTS SAMPLE

Case Study - AILA Validation

We validated AILA with a tool promoted by ENISA. Mercedes's privacy policy was chosen as test data.

Asset	AILA Fairness	AILA Likelihood		ENISA Likelihood
Geolocation	0.23	0.77	**High**	**High**
Maintenance	0.38	0.62	**Medium**	**High**
Vehicle Tracking System	0.4	0.6	**Medium**	**Medium**
System	0.1	0.9	**Very High**	**High**
Mobile Application	0.44	0.56	**Medium**	**Medium**
Payment Information	0.05	0.95	**Very High**	**High**
Data Collection	0.49	0.51	**Medium**	**Medium**

AILA AND ENISA LIKELIHOOD SAMPLES

AILA rocks!

$$r = 0.93$$

$$r_s = 0.91$$

$$p - \text{value} = 0.00026$$

Case Study - The Winner (or the Loser) is...

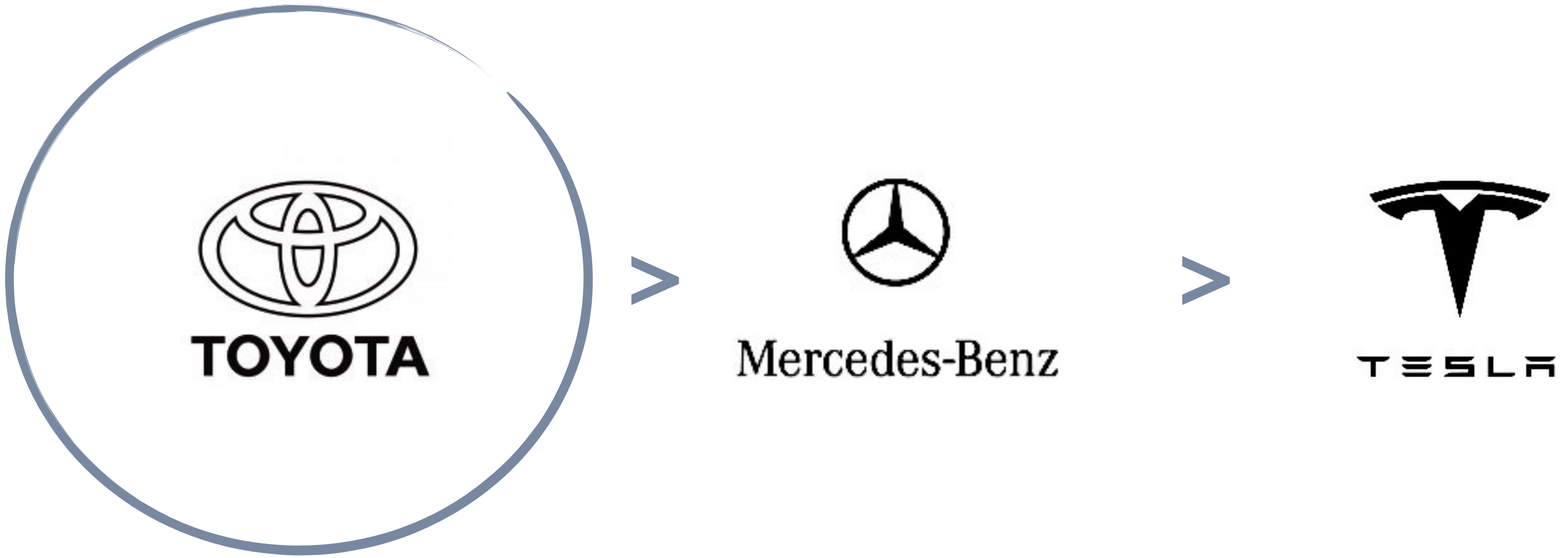

What's the Deal in Short?

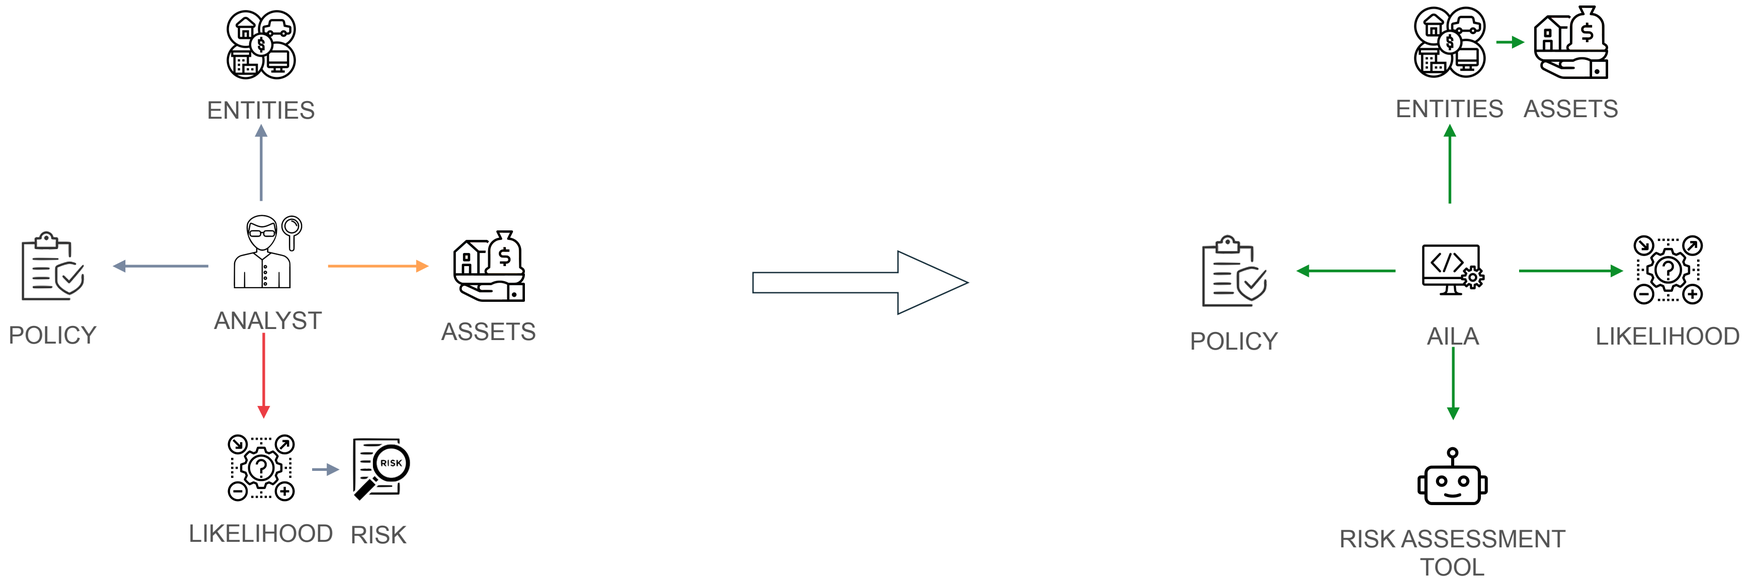

> AILA reduces **human subjectivity** through risk assessment.
> It facilitates **asset extraction** dramatically.
> It automates the **analyst's perception** of a policy.

Future work includes deeper semantic analysis and creation of a RA tool from scratch.

Q&A

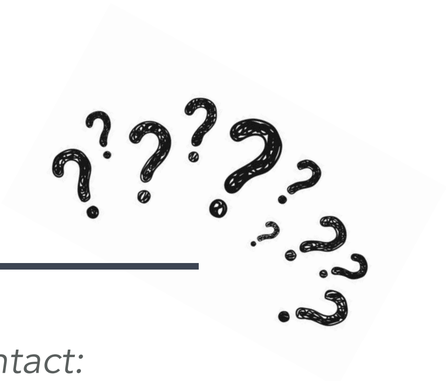

For more information or questions, please contact:

mario.raciti@imtlucca.it