AEXIA™ FRAMEWORK V3.1

User guide

Implementing AEXIA™ step by step: inventory, levels, gates, measures, steering

Mohamadou SIDIBE
AI BIAIS INITIATIVE
Chief Data Officer, Data and AI governance expert
Aligned with AEXIA™ Framework V3.1 · CC BY 4.0 licence

How to read this guide

This guide explains **how to use AEXIA™ Framework V3.1 day to day**: who does what, in which order, with which documents and evidence. It is written for CDOs, DPOs, Chief AI Officers, data stewards, risk managers and auditors. It does not replace the Framework: every definition and rule points back to its section.

The document kit

Document	What it is for	When to use it
AEXIA™ Framework V3.1	Reference: concepts, rules, measures	When a rule is in doubt
User guide (this document)	Step-by-step instructions	At start-up and each cycle
AEXIA-O evaluation grid	Measure the organisation's maturity, with evidence	Once or twice a year
Self-assessment sheet	Triptych summary, alerts, indicators, RAD-lite	Quarterly or semi-annually
AEXIA Card (template)	A system's identity sheet	One per system, updated at every gate and incident
AEXIA Card (Meridian example)	Completed example, fictitious system	To train teams

> **Status to keep in mind**
> AEXIA™ is an author's proposal, published in open access. The DCME, ECS and RAD measures have only had **preliminary validation on synthetic data and small public datasets**: they have not yet been validated on real systems or peer reviewed. Thresholds quoted in the documents are **policy benchmarks**, not standards. AEXIA is neither a certification nor a label.

1. AEXIA™ in ten minutes

AEXIA™ answers a question that existing frameworks (NIST AI RMF, ISO/IEC 42001, MLTRL, AI Act) do not address head-on: **who is accountable for a decision produced by a learning system, and on what evidence?** It orchestrates them rather than replacing them.

Two axes, one crossing

- **Axis O (organisation)**: six pillars (Foundations, Architecture, Governance, Compliance, Operations, Measurement and improvement) rated N0 to N5. The published level is that of the weakest pillar.
- **Axis S (system)**: each system receives a level from S0 to S9 (ten levels), aligned with MLTRL. It is reassessed at every gate and every incident.
- **The O × S crossing** says what the organisation is allowed to deploy, and with which controls.

The four mechanisms

Mechanism	What it requires
Weakest link	The published level is that of the weakest component, never an average.
Evidentiary gates	No level change by declaration: a closed set of artefacts, examined in review.
Downgrade	A qualified system can step back down: drift, incident, population change, regulatory change, retraining.
Proportionality	Effort is proportionate to the AI Act risk classification.

What is published: the triptych

Value	Definition	Use
AEXIA-O level	Level of the weakest pillar	Published externally
AEXIA-O profile	Levels of the six pillars (radar)	Internal steering: where to invest
AEXIA-S distribution	Fleet breakdown by level S0–S9	Operational trade-offs

The decision matrix

Org.	S0-S5 systems	S6-S7 systems	S8-S9 systems
N0-N1	Tolerated: exploratory phase	**Prohibited**: no controls	**Major alert**: production without governance
N2	Tolerated with traceability	Authorised under condition: gate passed in committee	Authorised for limited risk; high risk: no
N3	Normal	Normal, gate passed on evidence	Authorised, including high risk, under measured monitoring
N4-N5	Normal	Normal, automated	Normal, automated downgrade on drift

2. The six-step path

The AEXIA™ lifecycle (MAP, DEFINE, BUILD, GOVERN, OPERATE, IMPROVE) translates into six working steps. Durations are those given by the Framework for an organisation starting out; they vary with fleet size.

Phase	What you do	Deliverables	Indicative duration
MAP	Inventory systems, classify them (AI Act), place each on S0-S9	Inventory, classification, initial level, first AEXIA Cards	2 to 4 weeks
DEFINE	Set policy, roles (RACI), methods, policy thresholds	Charter, RACI, SHAP/LIME/DCME selection matrix	4 to 8 weeks
BUILD	Deploy tools, registries, dashboard, train	Augmented MLOps pipeline, catalogue, dashboard	2 to 6 months
GOVERN	Set up gates, the DGC, the audit trail	Gate decisions, DGC report, compliance report	Continuous
OPERATE	Manage explanations, incidents, service levels, RAD	Procedures, service levels, tickets	Continuous
IMPROVE	Reassess maturity, close gaps	Updated score, roadmap, executive report	Quarterly

The first 90 days

1. Appoint minimal roles (Data Owner, XAI Data Steward, DPO, CDO sponsor). 2. Inventory all AI systems, even “test” ones. 3. Classify each system under the AI Act and record the rationale. 4. Fill in a first AEXIA Card for high-risk systems. 5. Assess the organisation with the AEXIA-O grid. 6. Present the triptych to the executive committee.

3. Step 1 : Inventory and classify

3.1 Inventory

An AI system enters the inventory as soon as it produces or influences a decision affecting people: scoring, sorting, recommendation, detection, generation. Include purchased systems, supplier systems and prototypes reachable by real users. For each, record an internal code, the business owner, the purpose, the people affected and the first deployment date.

3.2 Classify under the AI Act

1. Is the system a prohibited practice (art. 5)? If so: **unacceptable**, stop.
2. Is it a safety component of a regulated product (Annex I) or does it fall in an Annex III domain (credit, employment, education, access to essential services, justice, etc.)? If so: **high risk**.
3. Does it interact with people or generate content subject to transparency (art. 50)? If so: **limited risk**.
4. Otherwise: **minimal risk**.

Record the rationale and the date: the card needs it and the regulator will ask for it. According to section 9 of the Framework (state of the law at 22 October 2026), obligations for Annex III high-risk systems apply from 2 December 2027 and for systems embedded in products from 2 August 2028; this section is reviewed every six months.

3.3 Proportion the effort

The Framework sets the proportionality principle (mechanism 4) without detailing the content. The following table is a **recommendation of this guide**, to be adapted by policy.

Class	Effort	Minimum content
Minimal	Light	Inventory; card sections 1 and 2; annual review
Limited (art. 50)	Standard	Card sections 1–4 and 6; explanation method chosen; semi-annual review
High risk / prohibited	Full	Full card; gates S5→S6, S7→S8, S8→S9; DCME if protected variables or proxies; ECS and RAD; monitoring (art. 72)

4. Step 2 : Place each system on the S axis

A system's S level is not a judgement: it is **the highest level whose evidence exists**, that of its least mature component. When in doubt, take the lower level.

S	Stage	Typical evidence	Environment
S0	First principles	Idea, literature review, mathematical basis	None
S1	Directed research	Recorded exploratory experiments	Laboratory
S2	Proof of principle	Working concept on a restricted case	Laboratory
S3	System development	Structured code, reproducible pipeline	Laboratory
S4	Proof of concept	Validation on real data	Representative
S5	Demonstrated capability	Identifiable, measured business capability	Representative · critical gate
S6	Application development	Production software quality; explanation registry	Representative
S7	Integration	Integration with production systems	Restricted operational
S8	Qualified	Complete compliance file; formal DGC decision	Operational
S9	In monitored service	Operation, maintenance, active monitoring (art. 72)	Operational

4.1 The four gates

Gate	Required artefacts	Decision-makers	Where to record it
S2 → S3	AI Act risk sheet, classification, data-readiness review	Data Owner + Data Steward	Card, section 3
S5 → S6	XAI method justification, quality thresholds, risk quantification	DGC (critical gate)	Card, sections 3 and 5
S7 → S8	Complete compliance file, frozen reference dataset, monitoring procedure	DGC + DPO	Card, sections 3, 6 and 8
S8 → S9	Active monitoring pipelines, drift thresholds, tested downgrade procedure	DGC + DSI	Card, sections 3 and 7

Running a gate. (1) The Data Steward assembles the artefacts and references the file in the card. (2) Decision-makers examine them in review, never by isolated signature. (3) The decision (passed or failed) is recorded with date, body and comment. (4) A failed gate stays failed: a remediation plan and a new review are needed. (5) The level is updated only after the decision.

4.2 Setting the expected level S_att

RAD compares expected and actual levels. The Framework does not fix S_att: it must be **declared by policy and published**. A benchmark suggested by this guide: S9 for a high-risk system in operation, S8 before its go-live, and the level matching the planned stage for a system in development.

5. Step 3 : Assess the organisation on the O axis

The AEXIA-O assessment uses the grid (6 pillars, numbered criteria, evidence per criterion). A first pass takes one to two days.

1. **Gather the evidence**: charters, RACI, DGC minutes, catalogues, audit reports, dashboards.
2. **Rate pillar by pillar**: tick a criterion only if evidenced; a partial criterion counts as not met.
3. **Assign the level**: the highest level whose criteria, and those of all lower levels, are met.
4. **Have a second assessor review** if the result is communicated externally; note disagreements.
5. **Publish the weakest pillar's level** and use the profile to decide investments.

> **Why not an average score?**
> With weights F 20 %, A 15 %, G 20 %, C 20 %, O 15 %, M 10 %, an organisation whose Compliance is at N0 and whose other five pillars are at N4 scores 3.2, i.e. "Managed", although it faces direct legal exposure. The weakest link shows N0 and makes the exposure visible. The grid offers the weighted score as an optional internal benchmark, never published alone.

Usual evidence by pillar

Pillar	Examples of acceptable evidence
F · Foundations	Validated glossary, method benchmark, selection matrix, DCME report with published graph, explanation quality standards
A · Architecture	MLOps pipeline diagram, data catalogue, explanation registry, API documentation, automatic gate configuration
G · Governance	Signed charter, RACI, quarterly DGC minutes, gate decisions, RAD-lite report
C · Compliance	AI risk register, classification of each system, GDPR / AI Act mapping, annual audit report
O · Operations	XAI runbook, service-level agreements, incident tickets, operational indicators
M · Measurement	Annual score, quarterly dashboard, ECS reports compliant with publication rules, improvement plans

6. Step 4 : Fill in the AEXIA Card

A card describes one system. It is the document a regulator will ask for first: it must stand on its own and point to evidence files. The pre-registered study aims at filling in a first card in **four hours or less** (median); this is a target to verify, not an established result.

Card section	What to enter	Pitfall to avoid
Header	ID AEXIA-XXXXXXXX, version x.y.z, status	Reusing an ID; not versioning
1 Identification	Name, code, purpose, people affected, roles (functions rather than names)	Vague purpose; forgetting the supported decision
2 Classification	Category, annex, rationale, application date	Classifying without rationale
3 S level	Current level, S_att, gate history with files	Declared level without file
4 Weakest link	Determining component, plan, deadline; O level and matrix cell	Plan with no owner or date
5 Method and measures	Method; DCME, ECS and RAD blocks with their publication rules	ECS without ε, controls or fidelity
6 Data	Readiness level, frozen set, lineage, known biases (TriAxis reference)	Reference set not frozen
7 Monitoring	Thresholds for each of the five triggers, frequency, last control	Monitoring without a downgrade threshold
8 Compliance	Documents, audit, gaps, deadline, information to individuals	Open gaps with no date
9 to 11	Card history, approvals, notes	Signatures without role

6.1 Versioning the card

- **Major (x)**: change of S level, classification or explainability method.
- **Minor (y)**: update of measures (ECS, RAD), data or monitoring.
- **Patch (z)**: formal correction with no effect on content.

6.2 Validate the format

The Framework provides a JSON schema (aexia_card_schema.json, Draft-07) and a Python validator. Required fields are: id, version, system_name, level_s, xai_method, data_readiness, compliance.

```
{ "id": "AEXIA-MERSCR01", "version": "0.3.0",
  "system_name": "Meridian consumer-credit scoring",
  "classification_ai_act": "haut_risque", "level_s": 5,
  "xai_method": "SHAP",
  "data_readiness": {"level": 4}, "surveillance": {},
  "compliance": {"status": "partial"} }
```

Limitation of the current schema: it accepts a single xai_method value, whereas a system may combine SHAP and DCME, and it does not contain the ECS and RAD fields of the V3.1 card.

7. Step 5 : Apply the matrix and decide

1. Take the organisation's **published AEXIA-O level**.
2. Take the system's **S level** (card, section 3).
3. Read the matrix cell and record it in card section 4.
4. If the cell is **Prohibited** or **Major alert**, escalate to the DGC and CDO without waiting for the next committee.
5. If the cell is **Authorised under condition**, check the gate was passed in committee and the condition is written down.

Three typical readings

Situation	Cell	Expected decision
N2 organisation, high-risk system at S7	Authorised under condition	Gate passed in committee, written condition (restricted scope, plan towards S8)
N1 organisation, system at S6	Prohibited	Do not deploy: raise the organisation to N2 or bring the system back to the lab
N3 organisation, high-risk system at S9	Authorised under measured monitoring	Keep monitoring, thresholds and downgrade procedure up to date

8. Step 6 · Monitor, downgrade, return

An AI system does not stay qualified: data drift, populations change, models are retrained. AEXIA provides five downgrade triggers.

Trigger	Effect	Obligation to return
Data or concept drift beyond threshold	S9 → S6	Method review, recalibration, new S7 → S8 gate
Explainability incident	S9 → S7	Root-cause analysis, correction, DGC review
Significant population change	S9 → S6	Data-readiness review, fairness check, compliance re-examination
Applicable regulatory change	S9 → S8	Updated compliance file, new formal decision
Model modification (retraining)	S9 → S7	Comparison with the reference set, validation of explanations

Six-beat procedure

6. **Detect**: an indicator crosses its policy threshold, or an incident is declared.
7. **Qualify**: identify which of the five triggers applies.
8. **Downgrade**: update the S level in the card (major version) and re-read the matrix cell.
9. **Decide**: if the new cell is Prohibited or the risk requires it, suspend or restrict use.

10. **Correct**: carry out the return obligation for the trigger.
11. **Pass the gate again**: new decision in review, recorded in the card.

Detection thresholds (for example a population stability index, a performance drop, an ECS degradation on recent data) are **set by policy** and recorded in card section 7. Post-market monitoring (art. 72) requires a plan for high-risk systems.

9. Using DCME, ECS and RAD without mistakes

The three measures are the author's original contribution; they are **preliminary**. Their use follows strict publication rules: without them, a figure can mislead. Numerical examples come from the Meridian case study (synthetic data).

9.1 DCME: where influence flows

- **When**: high-risk system where a protected characteristic (or proxy) may act on variables the model uses. Otherwise SHAP is enough.
- **Procedure**: build and publish the causal graph with its rationale; estimate ϕD (direct) and ϕC (causal); compute ψ = ϕC − ϕD; check efficiency and Σψ = 0 (tolerance 10^{-4}); run the graph sensitivity analysis.
- **Reading**: ϕD answers "does the model use the variable?", ϕC "does the variable influence the score, along all paths?". ψ shows where influence is redistributed. On Meridian: the group's direct contribution is zero, its causal contribution +0.95 points.
- **If identifiability is not defensible** (possible unobserved confounding): publish ϕD only and declare the limit; do not speak of causality.

9.2 ECS: is the explanation stable?

- **When**: for any explanation meant for a customer, a regulator or an audit.
- **Publication rules**: report ε, N, the normalisation and a fidelity measure; include a random control (floor near 0) and a constant control (misleading ceiling of 1); give several radii (0.05; 0.1; 0.25 σ); publish the bootstrap confidence interval; set a minimum threshold by policy.
- **Reading**: ECS measures stability, never truth or quality. On Meridian: 0.90 at ε = 0.1 σ, fidelity 0.71, controls 0.00 and 1.00. Also compare ECS across subgroups.

9.3 RAD: where and since when are we behind?

- **When**: to steer a system portfolio in front of the executive committee.
- **Calculation**: RAD-lite = Σ β × max(0, S_att − S_eff) / Σ β × S_att, between 0 and 1. Report **separately** the age of the oldest gap and the average opening duration.
- **Reading**: a weighted share of the remaining path. Look at each system's contribution: on a four-system portfolio (Meridian example), RAD-lite = 36 / 81 = 0.444, 89 % of it carried by two systems.
- **Precautions**: no normative threshold before calibration; ordinal S levels (read the gap as missing gates); a variation of a few hundredths may be measurement noise; cumulative variant optional, with T½ declared.

10. Steer and report

Steering fits on one page: the triptych, RAD-lite with gap age, the matrix alerts, the period's gates and the decisions requested from the committee.

Item	Cadence	Body
A system's S level	At every gate and incident	DGC (per gate)
AEXIA-O assessment	Once or twice a year	CDO, DGC
Indicators, RAD-lite, alerts	Quarterly (dashboard)	DGC then executive committee
Regulatory section	Semi-annual	DPO, CDO
Improvement roadmap	Quarterly (IMPROVE phase)	CDO

Typical content of the committee report

- Published AEXIA-O level, six-pillar profile and change since the last assessment.
- AEXIA-S distribution, including high-risk systems, and matrix alerts.
- RAD-lite, each system's contribution, gap age.
- Gates passed, failed, downgrades and incidents over the period.
- Decisions requested: investment trade-offs, conditional authorisations, AI Act dates.

11. Roles and responsibilities

Proposed allocation, adapted from section 4.3 of the Framework (A = accountable, R = responsible, C = consulted, I = informed). To be validated in the organisation's charter.

Activity	CDO	DSI	DPO	Data Owner	Data Steward	DGC
System inventory	A	C	C	R	R	I
AI Act classification	A	C	R	R	C	I
S level and card completion	I	C	C	A	R	I
Gate S2 → S3	I	I	C	A/R	R	I
Gate S5 → S6	I	C	C	C	R	A
Gate S7 → S8	I	C	R	C	R	A
Gate S8 → S9	I	R	C	C	R	A
DCME, ECS, RAD measures	I	C	C	A	R	I
Monitoring and downgrade	I	R	C	A	R	I
AEXIA-O assessment	A	C	C	C	R	C
Executive committee report	A/R	C	C	I	C	C

12. Worked example: Meridian Bank

Meridian is a fictitious bank. Its credit score does not use the sensitive variable, yet applications from the protected group are accepted less often (rate ratio about 0.83). Here is the full AEXIA path, on synthetic data; the corresponding completed card is in the kit.

Step	Result
1. Inventory	Consumer-credit score MER-SCR-01, in shadow mode since 1 September 2026
2. Classification	High risk, Annex III (creditworthiness of natural persons); deadline 2 December 2027; full effort
3. S level	S5; S_att = S8; S5 → S6 gate failed (quality thresholds and risk quantification not finalised)
4. Matrix	N2 organisation, S5 system: tolerated with traceability (no decision applied to customers)
5. DCME	Group's direct contribution: 0. Causal contribution: +0.95 points. Without the group's edges, 28 % of the effect is recovered: the graph must be justified
6. ECS	0.90 at ε = 0.1 σ; controls 0.00 and 1.00; fidelity 0.71; 0.03 gap between groups (less stable for the protected group)
7. RAD	System: 12 of 32; four-system portfolio: 36 / 81 = 0.444; age 24 months
8. Decision	Authorise only under conditions: documented decision on the legitimacy of influence paths, upgrade plan, new review of the S5 → S6 gate

The figures above come from the synthetic-data case study; they illustrate the method and do not constitute validation on real systems.

13. Mistakes to avoid

Mistake	Consequence	Correction
Publishing an average score	Hides a legal exposure	Publish the weakest-link level
Declaring a level without a file	Indefensible level in audit	Reference every piece of evidence
Adding indirect effects to causal Shapley values	Double counting (10.32 instead of 5 on the toy example)	Use $\phi C = \phi D + \psi$, $\Sigma\psi = 0$
Speaking of causality without a defensible graph	Biased conclusions	Publish ϕD only and declare the limit
Presenting an ECS without ε or controls	Uninterpretable score	Apply the five publication rules
Reading ECS = 1 as a good explanation	A constant explainer scores 1	Require fidelity alongside ECS
Using the v3.2 RAD thresholds 0.10 and 0.25	Thresholds withdrawn as not reproducible	No normative reading before calibration
Comparing two organisations by RAD	Premature comparison	Wait for calibration on real systems
Treating S levels as intervals	Misread gaps	Count missing gates
Not downgrading after retraining	System not requalified	Apply S9 → S7 and compare with the reference set
Not updating the card after an incident	Broken traceability	Version the card at every incident

14. Frequently asked questions

Is AEXIA a certification?

No. It is an open-access governance framework. Levels are documented positionings, not labels issued by a third party.

Must DCME be used everywhere?

No. It is useful when a protected characteristic or proxy can act through intermediate variables. Elsewhere SHAP and ECS suffice.

How to choose the criticality β?

The Framework ties it to AI Act classification (1 to 4) without fixing the mapping. Benchmark of this guide, to be declared: 1 minimal, 2 limited, 3 high risk, 4 high risk with strong impact on people (credit, employment, justice).

Is the ECS ≥ 0.8 threshold mandatory?

No. It is an example policy threshold: to be chosen, justified and published with ε.

How often to assess?

Organisation: once or twice a year. System: at every gate and incident. Indicators: quarterly.

Can it be used to assess a supplier?

Yes, by requesting its card and evidence files. Without access to evidence, the level stays “not established”.

How does it relate to ISO/IEC 42001?

The standard provides the management system; AEXIA adds per-system maturity, gates and explainability measures.

What if there is no second assessor?

Document the absence and have the DGC review the levels before any external communication.

15. Implementation roadmap

Horizon	Months	Priorities	O target	S target
Short term	M0-M3	Inventory, AI Act classification, first AEXIA Card, role assignment	N0 → N1	All placed
Medium term	M3-M12	Policy, S5 → S6 and S7 → S8 gates, explanation registry, monitoring pipelines	N1 → N2	No high-risk system beyond S7 without a gate
Long term 1	M12-M24	Internal audit, tested downgrade, automation, formalisation of DCME, ECS and RAD	N2 → N3	Stabilised distribution
Long term 2	M24-M36	Annex III compliance before 2 December 2027, ISO standardisation, publication	N3 → N4	Automated downgrade

Limits and status

- Preliminary validation on synthetic data and small public datasets; no validation on real systems.
- Inter-rater reliability of S levels not yet measured (study goal: Krippendorff's alpha ≥ 0.70).
- Thresholds (ECS, RAD, indicators): policy benchmarks, provisional.
- Regulatory section dated 22 October 2026; semi-annual review needed.
- This guide is not legal advice.

Annex A : Glossary

Term	Definition
AEXIA-O / AEXIA-S	Organisational (6 pillars, N0-N5) and system (S0-S9) axes
AEXIA Card	A system's identity sheet: level, method, data, monitoring, compliance
DGC	Data and AI governance committee that decides critical gates
Evidentiary gate	Level change subject to a closed set of artefacts examined in review
Weakest link	Rule: the level is that of the weakest component
Downgrade	Drop in S level after drift, incident, population change, regulatory change or retraining
DCME	Multi-scale presentation of causal Shapley values: $\phi C = \phi D + \psi$
ϕD, ϕC, ψ	Direct contribution, total causal contribution, mediated redistribution (zero sum)
ECS	Local stability score of an explanation over a neighbourhood of radius ε
Fidelity	Agreement between the explanation and the model's behaviour (by deleting or inserting variables)
Random / constant control	Reference explainers: expected floor near 0 and misleading ceiling equal to 1
RAD-lite	Weighted share of the portfolio's missing maturity levels, between 0 and 1
Gap age	Months since the oldest unresolved gap was opened
β, S_att, S_eff	Criticality (1-4), expected S level, actual S level
TriAxis	The initiative's bias reference: Origin × Lifecycle × Specific taxonomy

Annex B : Checklists

Before a gate decision

☐ All gate artefacts are assembled and referenced in the card

☐ The file has been reviewed by the prescribed decision-makers

☐ The AI Act classification is up to date

☐ The weakest link is identified and its plan is dated

☐ Decision, date and body are recorded

Before publishing an ECS

☐ ε, N, normalisation

☐ Fidelity measure and its value

☐ Random and constant controls

☐ Several radii and bootstrap confidence interval

☐ Justified policy threshold

Before publishing a DCME

☐ Graph published and justified

☐ Identifiability defended, otherwise ϕD only

☐ Efficiency and Σψ = 0 checked

☐ Graph sensitivity analysis

☐ Comparison of ϕD with interventional SHAP

Before publishing a RAD

☐ S_att and β declared by policy

☐ Gap read as missing gates

☐ Gap age reported separately

☐ No normative reading of values

☐ If cumulative variant: T½ declared and justified

Main references

Sidibe, M. (2026). AEXIA™ Framework, version 3.3. AI BIAIS INITIATIVE. CC BY 4.0.

Lavin, A. et al. (2022). Technology readiness levels for machine learning systems. Nature Communications, 13, 6039.

Regulation (EU) 2024/1689 (AI Act); Regulation (EU) 2026/1744 amending its application timetable.

NIST (2023). AI Risk Management Framework 1.0. ISO/IEC 42001:2023.

Heskes, T. et al. (2020). Causal Shapley Values. NeurIPS.

Alvarez-Melis, D., Jaakkola, T. (2018). On the Robustness of Interpretability Methods. ICML Workshop.

Stevens, S. S. (1946). On the Theory of Scales of Measurement. Science, 103.

Cunningham, W. (1992). The WyCash Portfolio Management System. OOPSLA.

Before distribution, check each reference against its original source.