AllergyTrace — Statistical Methodology

Transparent documentation of the calculation methods, with a fully worked example.

SilvaGIS GmbH · Dumeni Cavegn · as of 1 June 2026

1. What the app does — and what it does not

AllergyTrace is an **offline allergy diary with statistical correlation analysis**. Users record daily:

- **Custom symptom fields** on a 0–10 scale (e.g. itchy eyes, sneezing, shortness of breath — freely defined by the user)
- **Allergen exposures** with date, intensity, optionally duration (pollen, foods, animals, etc.)
- **Relief-medication intakes** with date and effectiveness

After at least 6 daily entries, the app computes on-device:

- **Quick analysis**: per-symptom-field mean comparison with time-lag optimization (Welch's t-test, Cohen's d)
- **Scientific analysis**: ordinal logistic regression with 5 input variables to adjust for confounders
- **Relief-medication correlation**: comparison of symptom values under exposure with and without relief medication

> **Transparent methodology:** all calculations are implemented in JavaScript / TypeScript and documented in the source files `analysisEngine.ts`, `quickAnalysis.ts`, `ordinalLogit.ts`. Source code available on request.

> **What AllergyTrace is NOT:** not a medical device in the legal sense (MDR/IVDR), not a diagnosis, not a substitute for an allergological work-up. Statistical correlation is not causation. Results are hypotheses that patients can bring to a consultation — not medical confirmations.

2. Statistical formulas

1 Arithmetic mean

Standard formula:

$$\bar{x} = (1/n) \cdot \Sigma\, x_i$$

2 Sample standard deviation (Bessel correction)

With n−1 in the denominator (unbiased for samples):

$$s = \sqrt{(\Sigma(x_i - \bar{x})^2 / (n - 1))}$$

3 Cohen's d (effect size)

Effect size standardized in SD units:

$$d = | \bar{x}\_A - \bar{x}\_B | / s\_pooled$$

$$s\_pooled = \sqrt{((s\_A^2 + s\_B^2) / 2)}$$

Interpretation thresholds after Cohen (1988): **d ≈ 0.2 small, d ≈ 0.5 medium, d ≈ 0.8 large**. Sawilowsky (2009) adds: **d ≈ 1.2 very large, d ≈ 2.0 huge**. AllergyTrace uses the simple pooled SD √((s_A² + s_B²) / 2); the alternative Hedges variant d_s additionally weights by sample size — for strongly unequal n the values differ slightly.

4 Welch's t-test (independent samples, unequal variances)

t statistic:

$$t = (\bar{x}\_A - \bar{x}\_B) / \sqrt{(s\_A^2/n\_A + s\_B^2/n\_B)}$$

Welch–Satterthwaite degrees of freedom:

$$df = (s\_A^2/n\_A + s\_B^2/n\_B)^2 / ((s\_A^2/n\_A)^2/(n\_A - 1) + (s\_B^2/n\_B)^2/(n\_B - 1))$$

The app uses **Welch's t-test** instead of the classic Student's t-test, because in practice samples almost always have unequal variances. The Welch–Satterthwaite df approximation yields more accurate critical values under unequal variances.

5 Two-sided p-value via the Student-t distribution

> Relation of the t distribution to the regularized incomplete beta function (Abramowitz & Stegun §26.7.1):
>
> $$p_{(2\text{-sided})} = I_x(df/2, 1/2) \quad \text{with} \quad x = df / (df + t^2)$$

The regularized incomplete beta function $I_x(a, b)$ is computed numerically via a continued-fraction expansion (Numerical Recipes §6.4, Lentz's algorithm); log(Gamma(x)) via the Lanczos approximation.

Validation against tabulated values (standard reference):

t	df	Table p	App p	Deviation
2.228	10	0.0500	0.0500	< 0.0001
1.812	10	0.1000	0.1001	< 0.0001
3.169	10	0.0100	0.0100	< 0.0001
4.000	20	0.0007	0.0007	< 0.0001

6 Significance criterion

AllergyTrace only labels an effect as "significant" when **both** conditions are met:

1. **Effect size**: |Δ| > 1.0 points on the 0–10 scale (≈ 10% of the scale)
2. **Statistics**: p < 0.05 in Welch's t-test

This prevents trivial effects from appearing "significant" merely because of a large sample.

3. Worked example: birch pollen vs. itchy eyes

3.1 Dataset

30-day simulated diary (1–30 April 2026), with realistic noise: some exposure days with a mild reaction, some non-exposure days with an elevated value (other triggers such as stress, sleep, other allergens). "Itchy eyes" as a custom symptom field, 0–10 scale, lower is better. Birch-pollen exposure on 14 days.

Date	Pollen exposure	Itchy eyes	Date	Pollen exposure	Itchy eyes
2026-04-01	yes	6	2026-04-16	yes	6
2026-04-02	no	2	2026-04-17	no	5
2026-04-03	yes	5	2026-04-18	yes	6
2026-04-04	yes	7	2026-04-19	yes	7
2026-04-05	no	1	2026-04-20	no	2
2026-04-06	yes	6	2026-04-21	no	1
2026-04-07	no	4	2026-04-22	yes	5
2026-04-08	yes	4	2026-04-23	no	3
2026-04-09	yes	8	2026-04-24	yes	6
2026-04-10	no	2	2026-04-25	no	2
2026-04-11	yes	5	2026-04-26	no	1
2026-04-12	no	3	2026-04-27	no	3
2026-04-13	yes	7	2026-04-28	no	2
2026-04-14	yes	3	2026-04-29	no	4
2026-04-15	no	2	2026-04-30	no	2

Group A (exposure): n=14, values = {6, 5, 7, 6, 4, 8, 5, 7, 3, 6, 6, 7, 5, 6}, Σ = 81. Note: two exposure days with a mild symptom (4, 3) — typical of days with low pollen counts or an additional protective measure.

Group B (no exposure): n=16, values = {2, 1, 4, 2, 3, 2, 5, 2, 1, 3, 2, 1, 3, 2, 4, 2}, Σ = 39. Note: three "no exposure" days with an elevated symptom (4, 5, 4) — presumably other triggers (stress, lack of sleep, other allergens).

3.2 Step-by-step calculation

Step 1 — Arithmetic mean

$$\bar{x}_A = 81 / 14 = \mathbf{5.7857}$$
$$\bar{x}_B = 39 / 16 = \mathbf{2.4375}$$
$$\Delta = \bar{x}_A - \bar{x}_B = \mathbf{3.3482}$$

Step 2 — Sample SD

$$\Sigma(x_A - \bar{x}_A)^2 = 22.3571 \rightarrow s_A = \sqrt{(22.3571 / 13)} = \mathbf{1.3114}$$
$$\Sigma(x_B - \bar{x}_B)^2 = 19.9375 \rightarrow s_B = \sqrt{(19.9375 / 15)} = \mathbf{1.1529}$$

The SDs show spread in both groups — a plausible reality for self-reported symptom scales.

Step 3 — Cohen's d

$$s_{pooled} = \sqrt{((1.3114^2 + 1.1529^2) / 2)} = \sqrt{(1.5239)} = \mathbf{1.2347}$$
$$d = |3.3482| / 1.2347 = \mathbf{2.7118}$$

Interpretation: d = 2.71 lies far above the threshold for "large" introduced by Cohen (1988) (d ≈ 0.8). Sawilowsky (2009) extends the scale with "very large" (d ≈ 1.2) and "huge" (d ≈ 2.0). For self-report data from a diary, an effect of this magnitude is a clinically relevant signal.

Step 4 — Welch's t statistic

$$SE_A = s_A^2 / n_A = 1.7198 / 14 = 0.12284$$
$$SE_B = s_B^2 / n_B = 1.3292 / 16 = 0.08307$$
$$SE = \sqrt{(SE_A + SE_B)} = \sqrt{(0.20591)} = \mathbf{0.4538}$$
$$t = 3.3482 / 0.4538 = \mathbf{7.3785}$$

Step 5 — Welch–Satterthwaite df

$$\text{Numerator} = (0.12284 + 0.08307)^2 = 4.2399 \cdot 10^{-2}$$
$$\text{Denominator} = (0.12284)^2 / 13 + (0.08307)^2 / 15 = 1.6209 \cdot 10^{-3}$$
$$df = 0.042399 / 0.0016209 = \mathbf{26.1596}$$

Step 6 — Two-sided p-value

> $x = df / (df + t^2) = 26.1596 / 80.6921 = 0.32419$
> $p = I_x(df/2, 1/2) = I_{0.3242}(13.08, 0.5) = \mathbf{7.4941 \cdot 10^{-8}}$
> $\rightarrow p < 0.001$

3.3 Verification: code output vs. hand calculation

```
App code output (direct call of tTest() and cohensD()):
  mean(A)   = 5.7857     ✓ identical
  mean(B)   = 2.4375     ✓ identical
  std(A)    = 1.3114     ✓ identical
  std(B)    = 1.1529     ✓ identical
  cohensD   = 2.7118     ✓ identical
  t         = 7.3785     ✓ identical
  df        = 26.1596    ✓ identical
  p         = 7.4941e-8  ✓ identical

Significance check:
  |Δ| = 3.35 > 1.0 (minimum-effect threshold)   : true
  p = 7.49e-8 < 0.05                            : true
  → SIGNIFICANT (both conditions met)
```

The app implementation reproduces the hand calculation to machine precision (deviation < 10^{-4}). Reproducible via `node scripts/verify-statistics.mjs`.

3.4 What the app shows users

From the same calculation, AllergyTrace produces three display levels (disclosure levels). Users choose for themselves how much detail they want to see:

Level 1 — Simple

🌳 **Birch pollen**

Likely triggers symptoms

Confidence: usable (84%)

Plain-language interpretation without statistical vocabulary. For first-time users and everyday use.

Level 2 — Detailed

🌳 **Birch pollen → itchy eyes**

With exposure: **5.8 / 10** (n=14)
Without exposure: **2.4 / 10** (n=16)
Difference: **+3.3 points**

Based on 30 days · 84% confidence

Concrete numbers with bar visualization. For engaged users.

Level 3 — Scientific

🌳 Birch pollen → itchy eyes (quick analysis)

Metric	Value	Interpretation
Mean with exposure	5.79	n_A = 14, s_A = 1.31
Mean without exposure	2.44	n_B = 16, s_B = 1.15
Δ (mean difference)	+3.35	on the 0–10 scale
Cohen's d	2.71	very large effect
Welch's t statistic	7.38	df = 26.16
p-value (two-sided)	< 0.001	(exact: $7.5 \cdot 10^{-8}$)
Data basis	30 days	Best lag: 0 h (immediate reaction)
Confidence level	usable (0.835 / 1.000) — another season improves the statement	

Methodology: Welch's t-test with Welch–Satterthwaite df, p via the Student-t CDF (regularized incomplete beta function). Statistical correlation, not a diagnosis.

Full statistics with a methodology reference. For patients who want to discuss the finding with an allergist.

Computing the confidence level (84% in the example)

$$\text{dataFactor} = \min(\text{N\_days} / 30, 1) \quad \text{(saturates at 30 diary entries)}$$
$$\text{effectFactor} = \min(|\Delta| / 5, 1) \quad \text{(saturates at 5 points of mean difference)}$$
$$\text{confidence} = 0.5 \cdot \text{dataFactor} + 0.5 \cdot \text{effectFactor}$$

Heuristic measure — combines data quantity and effect size into a single 0–1 number. Thresholds for the display:

Confidence	Display label
< 0.30	not evaluable
0.30 – 0.50	too little data
0.50 – 0.70	first tendency
0.70 – 0.85	usable
≥ 0.85	reliable

For the example: dataFactor = 30/30 = 1.00 (full saturation), effectFactor = 3.35/5 = 0.67, confidence = 0.5·1.00 + 0.5·0.67 = **0.835** → label "usable". The measure is not statistically interpretable (no power analysis); it serves UX communication: when is there enough data, when is the effect clear enough?

3.5 Accounting for relief medication

Relief medications (antihistamines, nasal sprays, eye drops, corticosteroids, etc.) change symptoms — and therefore also the means within the allergen groups. AllergyTrace accounts for them in two independent ways:

Way 1 — Relief medication as a confounder in the OLR model

The OLR model (see §4) includes `reliefTaken` as one of its 5 features. This lets the model estimate the isolated allergen effect even when relief medication was taken on some exposure days.

Why this matters: if a user took cetirizine on a birch day and the itchy eyes were milder as a result, a naive model would underestimate the birch effect. With `reliefTaken` as a covariate, OLR removes this bias automatically.

Way 2 — Relief-medication correlation as its own sub-analysis

In addition, the app runs a separate analysis per relief medication × per symptom field: *"Does this medication help with symptoms when this allergen is active?"*

Procedure:

1. Restrict the data basis to **exposure days** (only there can the relief medication take effect)
2. Split those days into two groups:
   - Group A': exposure days **with** relief-medication intake
   - Group B': exposure days **without** relief-medication intake
3. Same statistics as the main analysis: Welch's t, Cohen's d, p-value
4. Effect-size threshold: |Δ| > 0.5 points (instead of 1.0 in the main analysis, because relief-medication effects are typically smaller than allergen effects)

> $\Delta\_\text{relief} = \bar{x}\_A' - \bar{x}\_B' \quad (\text{negative} = \text{symptoms are relieved})$

3.5.1 Worked example: cetirizine against birch-pollen-related itchy eyes

Hypothetical illustration scenario — independent of the main dataset in §3.1. Here we show what the relief-medication sub-analysis would look like for a typical antihistamine effect: 14 pollen-exposure days, on 8 of which 10 mg of cetirizine was taken, on 6 not. The itchy-eyes values below are chosen to show a plausible cetirizine effect — they do *not* match the values in the main dataset §3.1, because the main example contains no antihistamine signal.

Date	Pollen	Cetirizine	Itchy eyes	Note
2026-04-01	yes	yes	5	cetirizine taken after waking
2026-04-03	yes	no	5	forgotten
2026-04-04	yes	yes	6	cetirizine, heavy pollen count

2026-04-06	yes	yes	4	—
2026-04-08	yes	no	4	mild day
2026-04-09	yes	no	8	forgotten, heavy pollen count
2026-04-11	yes	yes	4	—
2026-04-13	yes	no	7	forgotten
2026-04-14	yes	yes	3	—
2026-04-16	yes	yes	5	—
2026-04-18	yes	no	6	forgotten
2026-04-19	yes	yes	5	—
2026-04-22	yes	yes	4	—
2026-04-24	yes	no	6	—

Group A' (pollen + cetirizine), n=8: {5, 6, 4, 4, 3, 5, 5, 4} — mean 4.50, SD 0.93

Group B' (pollen without cetirizine), n=6: {5, 4, 8, 7, 6, 6} — mean 6.00, SD 1.41

> $\Delta = 4.50 - 6.00 = \mathbf{-1.50}$ (negative → relief)
> $\text{s\_pooled} = \sqrt{((0.93^2 + 1.41^2) / 2)} = \mathbf{1.20}$
> $\text{Cohen's d} = |-1.50| / 1.20 = \mathbf{1.25}$ (large effect)
> $\text{SE} = \sqrt{(0.93^2/8 + 1.41^2/6)} = \sqrt{(0.108 + 0.331)} = \mathbf{0.663}$
> $\text{t} = -1.50 / 0.663 = \mathbf{-2.260}$, $\text{df} = 8.13$ (Welch–Satterthwaite)
> $\text{p} = \mathbf{0.053}$

Statistical picture in this example: a **large effect size** (d = 1.25) with **borderline significance** (p = 0.053), due to the small subgroups (n_A' = 8, n_B' = 6). The app presents Δ, d, t, df and p directly; the clinical interpretation and any therapeutic conclusions rest with the treating professional.

> **Important:** relief-medication correlation is *only available in the Pro tier*. It serves as a hypothesis for the consultation — not a recommendation on dosage, switching, or discontinuation.

3.6 Multiple allergens tracked in parallel

In practice, the patient usually does not know what triggers the symptoms — otherwise the app would not be needed. Real diaries therefore contain **several allergens tracked in parallel** (pollen from various plants, animal hair, foods, dust mites, etc.). The question is: how does AllergyTrace isolate the effect of a single allergen when several are active on some days?

Approach: one analysis per allergen × symptom field

AllergyTrace computes a *separate* model per allergen for each symptom field. With three allergens (e.g. birch pollen, cat hair, gluten) and three symptom fields (itchy eyes, sneezing, stomach ache), up to 9 analysis results are produced — each with its own means, Cohen's d and p-value. The effects are thus directly comparable.

Confounder adjustment via the OLR (feature `otherExposureCount`)

So that the birch effect is not distorted by cat days that happen to coincide, the OLR model includes an `otherExposureCount` feature (the number of other allergen exposures on the same day, normalized to 0–1). The model thereby estimates the isolated allergen effect while accounting for parallel triggers. The *quick analysis* (mean comparison) without OLR does not contain this adjustment — it is therefore intended as an orienting first step, while the OLR result is methodologically more robust.

3.6.1 Worked example: birch pollen + cat hair + gluten in parallel

Assumption: the same patient as in the main example additionally tracks cat hair (visits to relatives with a cat) and gluten (eating out). Over the 30 days:

- **Birch pollen**: 14 exposure days (as before)
- **Cat hair**: 6 exposure days (weekend visits)
- **Gluten**: 9 exposure days (lunch at a restaurant)
- **Overlap birch + cat**: 3 days
- **Overlap birch + gluten**: 4 days
- **Overlap cat + gluten**: 2 days

AllergyTrace computes three separate per-symptom-field analyses. Results for the symptom *itchy eyes*:

Allergen	n exposed	$\bar{x}$ with	$\bar{x}$ without	Δ	Cohen's d	p (Welch)	App finding
Birch pollen	14	5.79	2.44	+3.35	2.71	< 0.001	Triggers symptoms (significant)
Cat hair	6	5.83	3.50	+2.33	1.42	0.012	Triggers symptoms (significant)
Gluten	9	3.78	3.62	+0.16	0.09	0.812	No clear relationship

Note: the birch-pollen row corresponds directly to the worked example in §3.1 / §3.2. The cat-hair and gluten rows are illustrative figures that depict typical multi-allergen findings — one clearly triggering allergen (cat), one null effect (gluten). They are not derived from a specific 30-day dataset.

Interpretation:

- Birch pollen **and** cat hair each show a significant effect on itchy eyes (both with $|\Delta| > 1.0$ and $p < 0.05$).
- Gluten shows only a 0.16-point difference at $p \approx 0.81$ — no indication of a contribution to itchy eyes (which fits clinically: celiac disease / gluten intolerance manifests differently).

- For sneezing and stomach ache, the same analysis would run separately and may identify different allergens as triggers.

Adjusted birch effect in the OLR model

Because 3 of the 14 birch days also had cat exposure, the question is: would the birch effect be the same without this co-exposure? The OLR model estimates the adjusted β for `exposedBinary` while simultaneously controlling for `otherExposureCount`. With low confounder overlap (as here, 3 of 14 days), the estimated β usually stays close to the unadjusted value. With strongly overlapping seasons (e.g. hazel: January–March, alder: February–April; about two months of co-exposure), the confounder adjustment is more impactful, since the effects would otherwise not be separable.

3.6.2 Limits of the multi-allergen analysis

- **Strongly correlated allergens**: if two triggers almost always occur together (e.g. two early-blooming pollens), the statistics cannot cleanly separate the effects. In such cases the app reports both as probable triggers, without a causal hierarchy.
- **Rare exposures**: with fewer than 3 exposure days, no model is computed (minimum sample size).
- **Undetected triggers**: allergens the patient does not record in the diary cannot be analyzed — so it is important to track broadly in the diary, even when unsure.
- **No multivariate correction across multiple symptom fields**: when an allergen is tested against several symptoms, the family of tests is not adjusted for multiple testing (e.g. no Bonferroni correction). The threshold $|\Delta| > 1.0$ in addition to $p < 0.05$ mitigates this risk in practice.

4. Ordinal logistic regression (advanced analysis)

For the *scientific analysis level*, AllergyTrace complements the quick analysis with an ordinal threshold model (ordinal logistic regression, OLR) with confounder adjustment. Advantage over the pure mean comparison: simultaneous consideration of several influencing factors.

4.1 Model structure

Ordinal threshold model ("cumulative link model" in the English literature) with 5 classes — the 0–10 scale is compressed to {0–2, 3–4, 5, 6–7, 8–10} (justified in §5.2). A separate model is trained per symptom field.

> $P(Y \le k \mid x) = \sigma(\tau_k - \beta^T x)$

With the logistic sigmoid $\sigma(z) = 1 / (1 + e^{-z})$, ascending thresholds $\tau_0 < \tau_1 < \tau_2 < \tau_3$ (monotonicity enforced during training), and coefficient vector β.

4.2 Input variables (confounder adjustment)

1. `exposedBinary` — 0/1 indicator for allergen exposure on the relevant day (with a lag if applicable, see §4.4)
2. `intensity` — normalized exposure intensity: minimal = 0.2, low = 0.4, moderate = 0.6, high = 0.8, severe = 1.0
3. `otherExposureCount` — number of other allergen exposures on the same day, scaled to 0–1 (saturates at 5)
4. `reliefTaken` — 0/1 indicator for relief-medication intake on the symptom day
5. `dayOfWeek` — weekday (Sunday=0 … Saturday=6) scaled to 0–1; captures weekly activity patterns

4.3 Loss and optimization

> Negative log-likelihood:
> $L(\tau, \beta) = -(1/N) \cdot \Sigma_i \log P(Y_i = y_i \mid x_i)$

Minimized via gradient descent with a finite-difference approximation of the gradient. Hyperparameters (step size, number of epochs, stopping criterion) are chosen so that the optimization converges in < 1 s on a mid-range smartphone. The monotonicity of the thresholds $\tau_0 < \tau_1 < \tau_2 < \tau_3$ is restored after each update. Validation: per-field convergence is shown visually in the loss-curve plot (in scientific mode).

4.4 Time lag (lag analysis)

Some allergens act immediately (e.g. inhalation of animal hair → itchy eyes within minutes), others with a delay (e.g. food intolerance → digestive complaints after 1–4 h). The quick analysis tests lags derived from the reaction profile (typically 5–8 values between 0 h and durationHours) and selects the one with the largest cumulative |Δ| across all symptom fields. The OLR (more compute-intensive) tests only three lag points for performance reasons:

- **0 h** — immediate reaction
- **mean** of the lag list (typically near peak_h)
- **longest lag** from the list (typically durationHours)

Reaction profiles are values predefined per allergen for `onsetMinutes` (reaction onset), `peakHours` (symptom maximum) and `durationHours` (decay time), based on the allergological literature. Examples from the app database: *birch* 15 min / 2 h / 12 h, *hazel* 15 min / 2 h / 10 h, *peanuts* 5 min / 0.5 h / 4 h, *lactose* 30 min / 2 h / 8 h, *milk* 30 min / 2 h / 6 h. Full values in `src/data/allergens-db.ts`.

5. Conventions, limitations and disclaimer

5.1 Methodological conventions

- **Sample SD with n−1** (Bessel correction) — standard for inferential statistics
- **Cohen's d with the simple pooled SD** √((s_A² + s_B²) / 2) instead of the sample-size-weighted variant (Hedges' g / d_s). For roughly equal groups practically identical; for strongly unequal n minimal deviations — both conventions accepted in the literature.
- **Welch's t-test** instead of Student's t-test — more robust against variance heterogeneity
- **Welch–Satterthwaite df** — precise degrees-of-freedom approximation for Welch's t
- **p-value via the t-distribution CDF** (regularized incomplete beta function) — exact values for any df ≥ 1
- **Significance requires a minimum effect** (|Δ| > 1.0) AND p < 0.05

5.2 Limitations

- At least 6 daily entries for the quick analysis, at least 10 for the OLR
- Minimum group size: 2 observations per group (with/without exposure) in the quick analysis, 3 in the OLR
- OLR model compressed to 5 classes (performance on mobile devices) — slight loss of granularity compared to 11 classes
- Per-field analysis — no multivariate correction for simultaneously testing multiple fields (e.g. no Bonferroni correction)
- Time-lag optimization selects the lag with the largest |Δ| — possible optimistic bias from multiple selection
- Data are **self-reported** (recall and confirmation bias possible)

5.3 Disclaimer

> The app provides **statistical correlation indicators** — not a medical diagnosis. Correlation is not causation: a low p-value (< 0.05) and a large d mean that the symptoms differ statistically significantly on exposure days, not necessarily that the allergen is the cause (confounders such as weather, stress, other allergens are possible).
>
> The app is **not a regulated medical device** (no MDR/IVDR conformity assessment). Results are intended as hypotheses for a consultation with an allergist, not as a substitute for clinical diagnostics (skin-prick test, sIgE, provocation test).

6. Privacy and architecture

- **100% offline**: all calculations on the device, no cloud sync, no server communication
- **Local SQLite**: all diary data in an app-internal SQLite database, not accessible to other apps
- **PIN + biometrics**: optional access protection (Face ID / fingerprint via Expo SecureStore)
- **No analytics, no advertising, no trackers**: no third-party SDK that exports usage data
- **Swiss privacy**: developed in Graubünden, SilvaGIS GmbH. The offline architecture structurally minimizes the privacy risk: with no transmission to third parties, typical GDPR/revFADP processes (data processing agreements, cross-border transfers, external deletion requests) do not arise. Access, deletion and data portability rest solely with the user (local backup export; uninstalling the app deletes all data).

7. Source-code references

File	Contents
`src/services/analysisEngine.ts`	Mathematical helper functions: mean, standard deviation, Cohen's d, Welch's t, log-gamma, regularized incomplete beta function, Student-t CDF
`src/services/quickAnalysis.ts`	Quick analysis: per-symptom-field mean comparison with lag optimization and relief-medication effect
`src/services/ordinalLogit.ts`	OLR model: ordinal threshold model with 5 classes, 5 input variables, gradient descent with early stopping
`src/services/analysisFormatter.ts`	Output formatting for the 3 display levels (Simple / Detailed / Scientific)
`scripts/verify-statistics.mjs`	Standalone verification script: computes the worked example from §3 without app context, reproducible with `node scripts/verify-statistics.mjs`

Source code available on request: dumenicavegn@gmail.com
