Consent to data processing in the Lunar Oracle

Edition 5 of 25 September 2026, clarified on 2 November 2026: the name has been removed from what is passed to the model; less data is passed — no new consent is required. Consent to the Lunar Oracle consists of three documents, one per legal basis: this one (processing and transfer to the model), the consent to international transfer (GDPR Chapter V / Art. 49(1)(a)) and the consent to processing of special categories (Art. 9 — health and sex life). All three are accepted by a single tick, and the consent is indivisible: either you agree and use the Lunar Oracle, or the feature is unavailable. These documents are issued by the international entity because your registration is outside Russia; women registered in Russia accept a different entity's documents under a different law.

1. Controller

Controller: PT. TECHNO CAPITAL INDONESIA.

Registered address: Jalan Sunset Road No. 99D, Desa/Kelurahan Kuta, Kec. Kuta, Kab. Badung, Provinsi Bali, 80361, Indonesia.

Tax identification number (NPWP): 40.174.113.7-905.000.

Contact for data protection matters — e-mail: support@woomoon.app, Telegram: @lunnaya_care, or by post to the registered address above.

2. Who answers in the chat

The Lunar Oracle is a chat with a language model. The reply is written by a language model, not by a person: primarily DeepSeek (via kie.ai), with Google Gemini as the reserve when it is unavailable; the recipients are listed in section 7. The Oracle is not a doctor, makes no diagnoses and does not replace professional advice. In a crisis it shows helpline contacts.

3. The consent is indivisible

One tick covers all three documents. The box is not pre-ticked: silence and inaction are never consent. Withdrawal is all-or-nothing — there is no «I use it, but do not read that» setting. Without consent the Lunar Oracle is unavailable; everything else — cycle, moon, recommendations, calendar — works independently of it and of whatever you decide about the Oracle.

4. What goes to the model with every message

- the text of your message, the reply to it and the recent turns of this same conversation — up to the newest fifty, from your device;
- a short memory digest of the conversation (up to 600 characters), stored on your device and sent from it with the message;
- what makes the answer yours rather than generic: age in years, a note on which birth details are filled in on your profile (date, time, place — without the values themselves), lunar day, your archetype (natal nakshatra, natal rashi and natal lunar day), special days, the week ahead, library article titles, and the kind and date of what we sent you over the last few days;
- the positions of your natal chart from the calculation stored on our side: the ascendant (lagna); the sign of each of the twelve houses and its lord; the sign and house of each of the nine planets with its dignity, retrograde status and combustion; aspects between planets; the Moon's nakshatra with its pada, lord, deity, symbol and its span in signs;

the yogas in the chart with a description of how they formed, and the classical yogas that are not in the chart, with the reason; the yoga karaka; natural and functional (for your ascendant) benefic, malefic and neutral planets, and the planets that, per the classical text (BPHS), give a yoga for your ascendant, planets of mixed or conditional nature, and pairs of planets that do not form a yoga, with the explanation from the text; the current mahadasha and antardasha — without their start and end dates. Your date, time and place of birth and the degrees are not passed on, but together with your age these positions make it possible to work out your date and approximate time of birth fairly precisely;

- only if you yourself ask for a reading of your chart — further calculations made by our own service on our own server: dasha periods with their start and end dates; current transits relative to your natal Moon; a chart cast for the moment of your question (prashna) with its conclusions; the navamsha. The model receives only the result of the calculation; your birth data is not passed to it.

Your name and your account identifiers are not passed to the model; if you write your name in a message yourself, it goes as part of that message.

5. Context layers

This data is a special category and is described by its own document — the consent to processing of special categories. It is named here so you can see the whole picture in one place.

- Cycle: cycle phase and day, interval history, delays, the start date of your last period, the text of today's recommendation.
- Body signals from Apple Health: sleep, heart rate variability, resting heart rate, wrist temperature during sleep, cycle symptoms. They are computed on your phone, travel only with the message you sent, and are not stored on our servers. They are sent only if you allowed the iPhone app to read them in the Apple Health system prompt, which appears once after you give consent; declining that prompt closes nothing, and access is revoked in the iPhone settings. Menstrual flow is not part of this layer: it syncs with the calendar separately, in Apple Health settings.
- Intimacy marks you place in the calendar.

Health data and data concerning your sex life are special categories of personal data (GDPR Art. 9) and are processed on the basis of the separate consent linked above.

6. What never goes anywhere

- your full date and time of birth;
- e-mail;
- phone number;
- payment details.

7. Where it goes

The controller passes the data to the recipients listed below only to produce the reply: to the primary one — kie.ai with the DeepSeek model — or, when it is unavailable, to the reserve, Google. There is no data processing contract with kie.ai or DeepSeek, so they may process the data they receive for their own purposes as independent controllers; the terms

of each recipient are set out below. Processing takes place outside your country — in the USA and possibly in the People's Republic of China, and that is a separate legal basis with its own document: the consent to international transfer.

Recipients as of edition 5 (2026-09-25): primary recipient — kie.ai → DeepSeek; reserve, used when the primary is unavailable — Google. Within this list and on these terms the model provider may change without a new consent; a new recipient or new processing terms require a new consent.

- NEXUSAI SERVICES LLC (kie.ai) — the kie.ai API aggregator: receives the request, processes it with the DeepSeek model and returns the reply. Address: 118 Krameria St, Denver, CO 80220, USA. Country of processing: the United States of America — kie.ai does not disclose where exactly the model runs; by its public statements, on servers in the USA. Training on data: the kie.ai privacy policy does not say whether the content of API requests is used to train models. Retention: not disclosed — the kie.ai privacy policy does not state how long request content is kept. Data processing agreement: not confirmed — we have no separate agreement with kie.ai. Capacity: if it uses the data beyond relaying it to the model, it processes it for its own purposes as an independent controller, not on the controller's behalf; there is no contract with kie.ai.
- Hangzhou DeepSeek Artificial Intelligence Co., Ltd. (DeepSeek) — developer of the DeepSeek model that writes the primary recipient's reply; a possible recipient: kie.ai does not disclose where exactly the model runs, and it cannot be ruled out that request content is passed to DeepSeek in the PRC, so the controller assumes the worst case. Address: Room 1201, Building West 1, Huijin International Mansion, No. 169 Huancheng North Road, Gongshu District, Hangzhou, Zhejiang, People's Republic of China. Country of processing: possibly the People's Republic of China — if request content is passed to DeepSeek. Training on data: DeepSeek's Open Platform terms do not describe the use of data received for training; the controller assumes that DeepSeek may use it to train its models, and this cannot be opted out of through the app. Retention: not disclosed — DeepSeek's Open Platform terms do not describe how long data is kept. Data processing agreement: not confirmed — we have no separate agreement with DeepSeek. Capacity: may process the data it receives for its own purposes as an independent controller, not only on the controller's behalf; there is no contract with DeepSeek.
- Google LLC (Google) — owner of the Gemini model that writes the reply (paid Gemini API tier). Address: 1600 Amphitheatre Parkway, Mountain View, CA 94043, USA. Country of processing: the United States of America and any other country where Google operates capacity: Google processes and may temporarily cache data wherever its infrastructure runs. Training on data: under the paid Gemini API terms Google does not use prompts and replies to train models or improve its products, and does not have them reviewed by people for that purpose. Retention: prompts and replies are logged for a limited period — solely to detect policy violations and for legally required disclosures. Data processing agreement: confirmed by the controller on 19 September 2026 (Gemini API Paid Services terms). Capacity: processes the data on the controller's behalf under the paid Gemini API terms.

The history of this list is on the AI recipients page.

8. Purpose

To help you understand your own state better. The controller does not use this data for advertising or profiling, does not sell it, and passes it only to the recipients in section 7 and only to produce the reply. kie.ai and DeepSeek, however, have no data processing contract with the controller and may process what they receive for their own purposes as independent controllers under their own terms; DeepSeek's terms do not describe whether data received is used for training, so the controller assumes that DeepSeek may use it to train its models, and says so plainly. No decision producing legal effects concerning you is taken on the basis of automated processing alone.

9. Where the conversation lives, and how long the consent lasts

The conversation is not stored on our servers — it lives only on your device. The honest flip side: a browser can clear its storage (Safari does so after roughly a week without a visit), and then the conversation is gone, with nowhere for us to restore it from. The consent itself lasts while you use the feature and ends when you withdraw it.

10. How to withdraw

At any time: Settings → Consents (in the iPhone app and in the web app). Withdrawing is no harder than giving: one button. Withdrawal closes the Oracle entirely — the consent is indivisible — and touches nothing else: cycle, moon, recommendations and calendar keep working exactly as before. Withdrawal applies to future transfers — what has already left cannot be recalled, including where a recipient has already used it to train its model. You can erase all your data entirely in the same place, in account settings. You also have the right to lodge a complaint with a supervisory authority — in the EU/EEA, in the member state of your habitual residence, place of work or place of the alleged infringement (GDPR Art. 77); in the United Kingdom, with the Information Commissioner's Office.

The full description of processing is in the privacy policy. This document covers the Lunar Oracle only.