

Dubroll — Privacy Policy

Last updated: 14 September 2026

Kind of change: substantial — the Terms of Use change; this Policy is re-issued together with them, because the two documents are accepted as one and carry a single date. Its own text does not change. Clause 17.2 of the Terms.

Effective from: 14 September 2026

Applies to: the Dubroll mobile app for Android, distributed through Google Play with the technical package `com.torvyxlabs.memefactory`, published by Torvyx Labs Ltd. („Торвикс Лабс“ ЕООД). This text describes the Android app; if we publish Dubroll on another platform, this document will be updated, or a separate one published, before we do.

How to read it: in the app — from the notice on first launch and from the "About & Legal" screen; a copy as a PDF in the "Downloads" folder on your phone.

Published at: <https://dubroll.app/legal/privacy/bg/> and <https://dubroll.app/legal/privacy/en/>; short address <https://dubroll.app/privacy>. The permanent address of every version is in section "Changes and languages".

In short

This section is a summary. The binding text is the whole document below; where the summary and the text diverge, the reading that is in your favour applies.

- **There are no accounts.** Dubroll asks for no name, email, telephone or sign-in in order to work; you give a name or an email only if you choose to — on withdrawal from a contract, or when you write to us. There are no adverts and no advertising libraries (SDKs), and no cross-app tracking. The notifications are local — they do not come from a server.
- **The most important sentence in this document:** in order for the commentary to be written and voiced, the video you record or pick — including the sound, that is the voices of everyone in the frame and behind the camera — is sent to Google Gemini through Firebase AI Logic. Google is our processor of personal data — with one exception, stated in section "Our role and Google's role". One and the same video may go up to three times for one clip — and once more on every repeated attempt after a network error or on "Try again".
- **Our alignment server** in the European Union receives the synthetic voice and the text of the commentary in order to fit the captions; it does not receive your video, frames from it or your voice, and it keeps nothing of the content.
- **The clip is made on your phone** and does not come to us. The finished clips (the "My clips" library), the temporary files (the cache) and the settings are on the phone. Only the settings file is copied into your Google account through the Android backup.
- **Statistics and error reports — only if you switch them on.** Two checkboxes, both empty by default; switching them off takes effect immediately for the future, and what has already been collected with us is deleted within 24 hours.

- **The reports of AI content** carry a reason, which commentator and language, the time, the model and the version of the app, a fingerprint of the text and an identifier of the clip — without the text itself, without video, without an installation identifier.
- **Purchases go through Google Play.** Google processes your payment as a separate controller under its own terms towards you and sees your card; we do not.
- **The integrity check.** Before a request to Google, the app asks Google to confirm that this is a genuine, unmodified copy of Dubroll on a genuine device. Apart from the IP address, which Google sees on every call, this is the only identifier that always travels — and it is an identifier of the installation, not of you.
- **Dubroll is for persons aged 18 and over.** We ask once and you confirm; **we do not ask for and do not keep your date of birth** — only the answer, the threshold and the time at which it was given (section "What is kept on your phone").
- **The people in the frame have a section of their own.** They are not our users; what reaches Google about them, what we have told the model and what you owe them when you film them and publish, is in section "People in the frame".
- **You can object.** Against any processing that rests on a legitimate interest — which processing that is, is stated in section "Legal bases by purpose" — you can object at any time by writing to contact@dubroll.app with the subject "Objection" (section "Objection and withdrawal of consent").
- **Deletion on request without an account.** The data that may exist about you with us is deleted on a request made from the app itself or by email with your order number from Google Play — except the documents that the law or the contract obliges us to keep: those are restricted, not deleted (section "Retention periods and deletion").

The rest of the document explains all of this in detail.

Contents

1. Who we are
2. What Dubroll does
3. How Dubroll is built
4. What goes to Google Gemini and why
5. Which Google terms apply
6. The voice
7. Our alignment server
8. What is kept on your phone
9. Usage statistics and error reports — only with your consent
10. The app integrity check
11. Identifiers, advertising and tracking
12. Android backup and transfer between phones

13. Things you send out
 14. The permissions
 15. Subscriptions, purchases and the rights
 16. The reports
 17. Our role and Google's role
 18. Legal bases by purpose
 19. Transfers of data to third countries
 20. Retention periods and deletion
 21. Your rights and the honest limits
 22. Objection and withdrawal of consent
 23. Automated decision-making and artificial intelligence
 24. People in the frame
 25. Age — Dubroll is 18+
 26. Security and personal data breaches
 27. Changes and languages
 28. Contact and complaints
-

1. Who we are

The controller of the personal data described in this document is **Torvyx Labs Ltd.** („Торвикс Лабс“ ЕООД), a company registered in the Republic of Bulgaria.

- Registered seat and address of management: 46 Tsar Asen I Street, Burgas 8000, Bulgaria
- Company number (ЕИК/UIC): 208895745 · VAT number: BG208895745
- Contact for everything in this document: **contact@dubroll.app** · **+359 877 001 885**
- Responsible for the protection of personal data and for compliance with this Policy: the manager of the company, at the address of management and at contact@dubroll.app.

We are a controller within the meaning of Article 4(7) of the General Data Protection Regulation (Regulation (EU) 2016/679, hereinafter "the Regulation"). We are established in the European Union, and therefore the Regulation applies to our processing in relation to every user of Dubroll, wherever they live.

We have not designated a data protection officer. Here is why. The criterion under Article 37(1) of the Regulation is not the size of the company, but whether our core activity consists in regular and systematic monitoring of people on a large scale, or in large-scale processing of special categories of data. The core activity of Dubroll is the processing of video with people's faces and voices — including children and people who are not our users — so that a commentary about them can be written and voiced. That is not "monitoring" of people within the meaning of point (b): we do not follow anyone over time, there is no file and there is no profile (section "Automated decision-making and artificial

intelligence"); we do not seek special categories of data (section "People in the frame"). Nor is the "large scale" present, measured by the number of persons concerned, the volume of data, the duration of storage and the geographical scope: we hold no copy of anyone's content, nothing of it is kept with us, and the processing is transient. The threshold at which the answer changes is written down in one place — in our internal assessment under Article 37 — and not in several documents with different figures. We review the assessment with every release of the app, with every new processing operation and when the number of active installations passes the recorded threshold. If the answer changes, we will designate an officer, we will notify the Commission for Personal Data Protection (Article 25b of the Personal Data Protection Act) and we will name them here. Until then, privacy questions are handled directly by the company at the contact details given above — with no automated request system and no external agency.

2. What Dubroll does

You film a short video with the camera — you press the red record button — or you pick a part of a video in the phone's gallery with the "Gallery" button. Dubroll sends the video, including the sound, to a Google Gemini model together with our instructions and the language of the clip, and receives back a commentary in the style of the commentator you have chosen. The commentary is voiced with a synthetic voice, captions are fitted onto it, and on your phone a finished clip is assembled (rendered), which goes into the "My clips" library.

The commentators are nine invented characters — "Sports Commentator", "Political Chronicler", "Angry Gamer", "Nature Documentarian", "Horror Narrator", "Village Gossip", "Drama Diva", "Femme Fatale", "Snob Critic" — and all of them are available to everyone. None of them is a real person; the voices are Google's synthetic voices (Gemini TTS). The recording has a shortest possible length — the app does not make a shorter clip — and a cap in "Cinema" mode; both numbers are on the "Your plan" screen in the app. Before the recording you choose a mode: "Short" — of a fixed length, at which the cap and the floor are the same number, or "Cinema" — from the shortest possible length to the cap of "Cinema" mode; the kind of the clip (a "short clip" or a "cinematic clip") follows the chosen mode, not the measured length.

The language of the clip is the language in which the commentary is written and voiced; you choose it in "Clip language", and there you see which languages are currently offered. The list can change, and that is why it is not written out here.

Every clip carries an AI marking: metadata in the MP4 file itself and a visible notice that the voice is generated by artificial intelligence (AI) — the notice is on every clip, whatever the tier; for its language see the section stated in section "Automated decision-making and artificial intelligence". The brand watermark depends on who pays for the clip: without it are the clips with PRO and the cinematic clip paid with a purchased clip; a clip paid with the daily quota or with the cinematic welcome clip carries the brand watermark. The AI metadata and the visible notice stay on all clips. What is included in free use, in PRO and in "Clips by the piece" is governed by the Terms of Use, section "What Dubroll does"; here we are concerned only with what happens to your data.

That is the whole product. There is no public feed of clips, no messages between users, no profile, no social layer. Another user of Dubroll has no way of seeing anything of yours through Dubroll — you publish, outside the app, if you decide to (section "Things you send out").

3. How Dubroll is built

This section describes how Dubroll is built; most of what follows comes from it. Dubroll has no accounts — here is what reaches whom.

Your phone. The recording, the cache during production, the "My clips" library with the finished clips and the full text of every commentary, the settings and the local event log are on your phone, in the app's private storage. The assembly (the rendering) of the clip happens on the phone.

Google. We use three different Google services, in three different roles. The first is Gemini and Gemini TTS through Firebase AI Logic (Agent Platform) — in two cloud projects of ours at Google, one for PRO and one for free use; both go through one and the same path in the code. The second is the integrity check (Firebase App Check with Google Play Integrity). The third is Google Play for the purchases. Which role Google has for each of them is in section "Our role and Google's role".

Our alignment server. A service of Torvyx Labs Ltd. in Google Cloud Run, region `europa-west1` (Belgium), in our main cloud project. It receives the synthetic voice and the text of the commentary in order to return the timings of the words for the captions; it does not receive the video. The details are in section "Our alignment server".

Our server for the rights. In the same project and region sits the server that checks your purchases with Google Play, receives Google Play's real-time messages about renewal and refunds, and keeps the record of the two statements made on a purchase. It is described in section "Subscriptions, purchases and the rights".

The receiver of the reports. The reports of AI content go to a receiver of ours in the same project — section "The reports".

The receiver of the statistics. If you switch on the statistics or the error reports, the events go to a receiver of ours in the same project and region, not to Google Analytics or Crashlytics — section "Usage statistics and error reports — only with your consent".

The other side of it, said in advance. "Stored only on your phone" is a statement about storage, not about transmission: in order for the commentary to be written, the video has to be transmitted to Google, and it is transmitted — sometimes more than once. The app is an Android app with network access: it connects to Google and, for the components described, to us. We do not claim that we run no servers or that we never receive a copy — with us that would not be true.

What never reaches a server of ours: your video and sound, the finished clip, the answer about age or the verdict about your age, the fingerprint of your device identifier, your contacts, your precise location, your Google account and your email — unless you enter it yourself in order to receive a confirmation of withdrawal, or write to us yourself.

The text of the commentary — which describes your video and the people in it — passes in transit through our alignment server, in memory, with no record and no logs carrying content. We treat it as personal data, not as impersonal technical text.

4. What goes to Google Gemini and why

The commentary is written by a Google Gemini model, which the app reaches through **Firebase AI Logic** via Google Cloud's Agent Platform service — at the time of writing the model is `gemini-3.6-flash`. We may move to another Gemini model; the recipient, Google, and everything else in this

section are not changed by that, and a change of model is for us an occasion to check afresh what is described below about recording at Google. A change of provider would be a material change and we will tell you about it in the app first (section "Changes and languages").

4.1. Everything that travels in the request

The video — including the sound. The app makes a reduced copy of the recording for the model: a lower resolution (480p), a few frames per second (6 or 8 depending on the length) and a mono audio track. The sound travels — that is, the voices of everyone in the frame and behind the camera, the music and the ambient noise. The copy goes, if it is up to 10 MB. If the copy cannot be made and the raw recording is up to 10 MB, the raw recording goes. Above that cap, instead of video, between eight and twenty-four separate frames go as 768 px images. On an import from the gallery a normalised copy of the selected part goes, without the metadata of the original file; the original in your gallery is not changed.

The metadata of the recording. Nothing that goes to the model carries the metadata of the recording (device, time of recording and similar fields from the file container) — they are removed before the file departs. That holds for the reduced copy, for the individual frames, and for the case where the recording itself goes instead of a copy: it then passes through a copy without metadata, and if even such a copy cannot be made, the recording is not sent and frames go instead.

The emergency "two frames" mode. When the video recording fails technically, the app sends two still frames (JPEG) instead of video. This is an emergency path, not a feature. Those images cannot carry a location — the app does not ask for access to it.

Our instructions to the model and the language of the clip. Together with the video travels the system instruction: the style of the chosen commentator, the rules about the people in the frame, about the language, about the grammar and about what the model must not do, and the maximum number of words. The instructions are ours, not yours, and they are described in substance in section "People in the frame".

Text-only sendings. Once the commentary has been written, the same text — and it describes your scene and the people in it — may go to Gemini several more times without video: to choose between several candidates, for a proof-reading check, for pronunciation notes, to tighten it to fewer words, to extend it and to translate the captions into another language (PRO). Each of these is a separate request, which carries the text.

To Gemini TTS. For the voicing, the commentary, a style instruction (how it is to be read) and the name of the voice go to Gemini TTS; only audio comes back.

4.2. How many times the video travels

One and the same video may be sent up to three times for one clip. The first time is for the commentary. The second is for a clip over 20 seconds — as a "probe", with which the model finds when what appears in the frame. The third is when all the candidate commentaries breach our rules and the text has to be corrected afresh. Separately, every network error or temporary refusal by Google leads to a repeated attempt — up to three times for one request — and every repeated attempt uploads the same video again. "Try again" on the screen runs the whole pipeline afresh with the same recording, that is, a new upload. That is why we do not write that the video is sent "once" — it is not once.

4.3. Which cloud project and in which region

The request goes to one of our two cloud projects at Google — the one for PRO or the one for free use — according to whether you have PRO at the moment of the production; the path in the code is one and the same for both. The video is processed in Google's multi-region zone for the European Union ("eu") — the zone is set in the code. The voice is synthesised by Google where the voice model is offered: first in the "eu" zone, and when the model is not offered there — as has been the case since 11 September 2026 — in Google's global infrastructure, including outside the European Union, under the safeguards in section "Transfers of data to third countries"; only the text of the commentary goes to the voice, not the video. Both paths are under the same contract and under the same rules on recording at Google described below.

4.4. Google's purpose, which is not ours

Our purpose is one: that the commentary you have asked for is written and voiced. Google has a purpose of its own, which we cannot set aside by instruction: under Google Cloud's terms, automated safety classifiers review the requests to Gemini. A request that the classifiers do not flag is not recorded as content — for the model we use today; the technical record of the request ("Service Data", section "Which Google terms apply") is a separate thing and stays with Google. A flagged request may be recorded — only that one — for up to 90 days, solely in order to check whether Google's prohibited use policy has been breached, and authorised Google staff may review it. For a certain class of models ("Advanced AI") Google documents that it records all requests and responses for 30 days — that is not the model we use as at the date of this document, and a change of model is the moment at which we check this afresh. Google also maintains a short-term cache of the requests (up to 24 hours), which can be switched off at project level; until it is switched off for our two projects, your requests may sit in it for up to 24 hours. Google does not use your content to train its models ("Training Restriction" in the Service Specific Terms). We do not rely on the exception for short-term processing before Google Play, because the period at Google is not under our control.

An exemption from recording of a flagged request ("zero retention") can be asked of Google by a separate application; if we ask for it and it is granted, the row for Gemini in section "Retention periods and deletion" will become "not kept" and we will say so here.

4.5. What we see

Nothing of the content is kept on a server. The request to Gemini goes from your phone to Google and does not pass through us; we receive no copy of the video, and the text of the commentary is not recorded with us. On your phone the app keeps a local event log; what it contains is described in section "Usage statistics and error reports — only with your consent".

5. Which Google terms apply

The service we use is a paid Google Cloud service — Gemini through Firebase AI Logic in Agent Platform. We do not use the Gemini Developer API (the service for developers with a separate data regime) and the app published on Google Play has no way of switching to it: the mode is chosen once, when the package is assembled. The pre-publication check refuses a package with that mode. That is why only the Google Cloud documents are described here, by title. "Paid service" is a property of our contract with Google, not of your subscription: free use of Dubroll goes down the same paid path.

The Google documents that govern your data in this service (by title; Google renumbers its sections often, which is why we do not cite numbers):

- **Cloud Data Processing Addendum** — the contract under which Google processes personal data as our processor (Article 28 of the Regulation).
- **Google Cloud Platform Terms of Service** — the "Generative AI Safety and Abuse" clause (the abuse monitoring described in section "What goes to Google Gemini and why").
- **Service Specific Terms** — the clauses on generative artificial intelligence ("Generative AI Services"): "Training Restriction" (no training on your data) and "Age Restrictions" (why Dubroll is 18+; section "Age — Dubroll is 18+").
- **Google Cloud Privacy Notice** — "Service Data": the technical record of every request (number of tokens — units of text, status, result of the filters, IP address), for which Google is a separate controller.
- **Firebase Data Processing and Security Terms** — for the integrity check (Firebase App Check).
- The Agent Platform documentation on abuse monitoring ("Abuse monitoring") and on the location of the data ("Data residency").

The notice in the app — the "Before your first clip" screen — says the same thing in fewer words; where it is stricter than this section, it applies. Content that has already been sent to Google we cannot pull back, check or delete — that is true of us as well.

6. The voice

6.1. The synthetic voice

The voice you hear in the clip is generated by Google (Gemini TTS) from the text of the commentary and from a style instruction describing how it is to be read. The synthesis is in Google's cloud, not on the phone. The nine commentators are invented; none of them is a real person and no voice is a recording of a human being. If the current model is not available, the app tries a backup Google model — only in that case.

6.2. The backup voice

When Gemini TTS refuses or does not answer, the commentary is read by your phone's own speech synthesiser (Google TTS, Samsung TTS or another one your device uses). That is a third-party component with terms of its own. The app passes it the text of the commentary and the language code. Whether that synthesiser uses the network depends on it and on your phone's settings — we cannot promise it one way or the other. The app tells you when this has happened ("The voice actor did not show up - a backup voice was used."), and such a clip spends nothing from your daily quota.

6.3. People's voices

The voices of the people in the frame and behind the camera travel to Google as the sound of the video — section "What goes to Google Gemini and why". Dubroll does not identify anyone by voice and makes no voiceprints. More — in section "People in the frame".

6.4. The voice showcase

The samples you listen to when choosing a commentator ("Hear the voice") are files embedded in the app; when you listen to them nothing leaves the phone.

6.5. The disclosure

The voice and the commentary are artificially generated and the clip says so — what every clip carries today and what is planned is stated in section "Automated decision-making and artificial intelligence".

7. Our alignment server

The captions are fitted to the voice by the timings of the individual words. When our alignment server is not used or is not reachable, those timings are estimated on the phone itself and nothing leaves the device for that purpose.

The alignment server is a service of Torvyx Labs Ltd. running in Google Cloud Run, region `europa-west1` (Belgium), in our main cloud project. Google is our processor for the infrastructure under the Cloud Data Processing Addendum; the controller for the processing is us. The service is used by every published version of the app.

What it receives. The synthetic voice as an audio file (WAV), the text of the commentary (up to 4000 characters) and the language of the clip — and only when the captions are not "Off". The request is authenticated by the token from the integrity check before a single byte of its body is read; the token is always from our main project, whichever project the clip went through. It returns the timings of the words and the text that the recognition "heard" — and that is our own synthetic commentary.

What it does NOT receive. Your video, frames from it, your voice or the voices of the people in the frame, an identifier of a user or of an installation.

Another processor for six languages. For Hindi, Marathi, Tamil, Bengali, Telugu and Thai the server forwards the sound of the synthetic voice to Google Speech-to-Text v2, region `eu`, in order to recognise the words; Google is a processor under the Cloud Data Processing Addendum and data logging in that service is disabled — that is Google's default setting and we do not switch it on. For the other languages the recognition is on the server itself.

What it keeps. Nothing of the content: the server is stateless — nothing from the request is written to disk after it; only the logs below remain. Its application logs carry the language, the recogniser used, sizes, timings and coverage — no content at all. The Cloud Run request logs, which carry an IP address, are switched off for all our services; the application logs without an IP address are kept for 30 days in a regional log store in the EU. Google as a processor maintains infrastructure logs of its own under the Cloud Data Processing Addendum.

The availability check. The app sends the server an empty request to check availability — with no content and no identifier; only the technical connection travels, whose IP address we do not record. No network request to us starts before you have accepted the Terms of Use and this Policy — including the sending of the queue of reports, which waits for the acceptance.

When the server is not reachable. The clip is made with captions estimated on the phone. The text of the commentary describes your private video — people, children, home — which is why we treat it as personal data processed by us in transit, and not as impersonal text.

8. What is kept on your phone

We list what the app keeps; we do not give a number — the number of records changes with every version, and a change in their composition is a change of this document (section "Changes and languages"). Everything below is in the app's private storage on your phone. The app adds no encryption of its own on top of those files; it relies on Android's isolated storage (the sandbox) and on the encryption of your device.

8.1. The settings file (travels in the Android backup)

One settings file, which Android copies into your Google account and carries over when you change phone (section "Android backup and transfer between phones"). In it are:

- **The answer about age** — your confirmation that you are 18 or older, the threshold against which it was given, and the time. **A date of birth is not collected and not kept, and there is no "blocked" state** (decisions of 6 September 2026). An installation from before the change that had a stored date of birth loses it at the first launch of the new version: the date is deleted and the confirmation is asked for again under the new model.
- **The record of the acceptance** of the Terms of Use and of this Policy — which version you accepted (the "Effective from" date), when and in which language; the record is append-only and is not overwritten — a second acceptance does not erase the first.
- **The two statements on a purchase** (section "Subscriptions, purchases and the rights") — a unique number (nonce), the time, the version of the text, the language, the plan, the letter of the price tier and the price shown; an append-only list.
- **The order number** in Google Play (it starts with "GPA."), the product and the time of the last confirmation of PRO, the time of the first check at which Google Play did not return a subscription.
- **The balance of the purchased cinematic clips** and the numbers of the last 50 purchases, so that one purchase is not counted twice.
- **The daily counters** — short clips used from the daily quota for the day, successful renders, cinematic clips, whether the cinematic welcome clip has been used; the streak of consecutive days.
- **The times of the refused clips** over the last 24 hours — the counter behind the one-hour pause (section "Automated decision-making and artificial intelligence").
- **Your preferences** — the language of the clip, the chosen commentator, the caption settings (mode, language, letters, font).
- **The learned tempo of each voice** — a measure of how fast a given voice speaks in a given language, so that the clips fit better; that is data about the voice, not about you.
- **The counter of gallery imports** for the day — for measurement only, it limits nothing.
- **Two housekeeping flags** — that the introduction has been shown and that your first clip has been made; zeros and ones, with no other content.

8.2. The files (they do not travel in the backup)

- **The "My clips" library** — every finished clip (MP4), a still frame for its cover and an index in which, for each clip, sit the full text of the commentary, the commentator, the language, the time of creation, whether it carries a brand watermark, the duration and whether it has been saved to the gallery.

- **The event log** — the local file with records of the app's operation: at 1 MB the current file is archived, the previous archive is deleted and recording continues in a new file (section "Usage statistics and error reports — only with your consent").
- **The queue of the reports** of AI content — up to 90 days and up to 200 rows (section "The reports").
- **The deletion marker** — a single file that tells the app on the next launch that the data has been deleted with the row "Delete my local data"; it is reported once locally and deleted.
- **The cache during production** — the raw recording, the copy of the part picked from the gallery and the folders of the system picker, the copy for the model, the audio file of the voice, frames, captions, the intermediate render. On a successful production the inputs and the audio are deleted; on the next launch of the app the leftovers are cleaned up; on a failure the raw recording stays so that "Try again" can use it, and on a final error it is deleted at once.
- **Fonts and brand files** — they are not personal data.
- **The internal stores of the Firebase library** — the installation identifier and the token of the integrity check are files of Google's library (section "The app integrity check"); the row "Delete my local data" deletes our files, not those.

The retention periods for each of them are in section "Retention periods and deletion".

9. Usage statistics and error reports — only with your consent

The app keeps an event log on your phone: the name of the event, the time and a few fields — numbers, zeros and ones, labels from a fixed list (language code, commentator, plan, the letter of the price tier, class of error). The file is overwritten after 1 MB and stays on the phone; today's log also contains a few fields with free text: words and pairs of lines from the commentary, on two errors the message of the exception, and on one of them the first lines of the internal trace of the error, as well as messages from Google's third-party libraries; they do not leave the phone and are to be removed by a fix in the code. The log also carries your order number in Google Play and the number of the two statements on a purchase — on the phone, not outside. The same rows are repeated in the Android system log (logcat) as well, including in the published version of the app — readable with the developer tools when you have enabled debugging on your phone, not over the network. We do not use Google Analytics, Firebase Crashlytics or Firebase Performance — there is no such library in the app.

9.1. The two checkboxes

We send usage statistics and error reports only if you have switched on the relevant checkbox — "Usage statistics" and "Error reports". Both are empty by default. They stand on the notice in the app on first launch and permanently in the "Your data" section of the "About & Legal" screen. The text next to each checkbox says what is recorded on your device for it and what categories of data are sent.

Your consent is tied to the version of the text you saw: if we change the text of the checkbox, it is switched off and we ask you afresh. On a new phone the checkboxes are also asked afresh. Only events that occurred after the moment of consent are uploaded — what was recorded in the local log before it is not read for sending.

Withdrawal is as easy as giving: you switch the checkbox off. Switching it off stops the sending immediately for the future, deletes from the phone the random number for this installation (section

"Identifiers, advertising and tracking") and sends a request to delete what has already been collected with us — that is deleted within 24 hours of its receipt. No backups of this data are made.

The legal basis is your consent under Article 6(1)(a) of the Regulation, and the writing of the random number onto your device — Article 4a of the Electronic Commerce Act (the transposition of Article 5(3) of Directive 2002/58/EC). The service does not depend on these checkboxes: the clips are made in the same way without them.

9.2. Where it goes and who sees it

The events go to a receiver of ours in Google Cloud Storage, region `europa-west1`, in our main cloud project; Google is a processor under the Cloud Data Processing Addendum. They do not go to Google Analytics, Crashlytics or any other analytics product. Access is held by the manager of the company (and an authorised deputy in their absence) and the service accounts of the services themselves.

9.3. What every event carries and what it cannot carry

Every event carries a name from a fixed list, the time, the version of the app, the letter of the price tier (A or B), the language of the app and the random number of the installation, plus fields that are numbers, zeros and ones or labels from a fixed list. A double list of the permitted fields — in the app and on the server, with a test that fails on every new field — throws out everything that is not on the list. What cannot get through: text from the commentary, an order number, the number of the statements (nonce), a purchase token, file names; the IP address of the request is not recorded, because the request logs are switched off. The error reports carry the type of the exception, our list of the calls (stack trace) without file paths, the version of the app, the version of Android and the model of the device; the message of the error is not sent — only its type. The full list of the events and their fields is produced from the allow-list in the code and is published together with the version of the app that sends them.

9.4. Retention periods

The raw events are kept for 90 days against the random number of the installation; after that only daily counts by event, version and price tier remain, without any identifier. The error reports are kept for 90 days. An installation with no events for 90 days is deleted: by uninstalling you reach the same result as by withdrawing consent, but after up to 90 days, not up to 24 hours. The price of refusing: without consent we do not see that the app has stopped with an error on you, and we cannot fix it for you.

10. The app integrity check

Before a request to Google, the app asks Google to confirm that this is a genuine, unmodified copy of Dubroll on a genuine device — Firebase App Check with Google Play Integrity, for our two cloud projects. Google returns a token, which travels with the request; the token carries signals about the device and about the app, not content. No identifier of a user of ours travels — the request to Gemini carries only that token.

For that purpose Google's library keeps on your phone an installation identifier (Firebase installation ID) — one for each of the two projects — and the token. Those identifiers are of the installation, not of you; at Google they live under the Firebase terms (Firebase Data Processing and Security Terms). We do not record them anywhere.

The check is activated when you accept the Terms of Use and this Policy — not earlier. Before the acceptance the app does not attest itself to Google and does not ask for a token, and the attestation itself

carries nothing of yours. The basis is our legitimate interest in protecting from abuse the service we pay Google for (Article 6(1)(f) of the Regulation); the access to your device for the token is strictly necessary for the service you have asked for — the exception under Article 4a(4)(2) of the Electronic Commerce Act (Article 5(3) of Directive 2002/58/EC). There is no switch: without the check the service cannot be protected.

If the check does not pass — for example on a device without Google Play services or with a modified app — a new clip cannot be made and the app tells you so ("The studio didn't recognise this copy of the app. Install it from Google Play or try again later."). Your finished clips remain readable.

The same token also authenticates the requests to our alignment server — always from our main project, whichever project the clip went through.

The token also authenticates the requests to our server for the rights, to the receiver of the reports and to the receiver of the statistics.

11. Identifiers, advertising and tracking

There are no adverts, no advertising library (SDK), no install attribution and no notifications from a server — the Android notifications you see during a production are local. The Android advertising identifier is not asked for: there is no such permission in the assembled package of the app. We do not ask for location, contacts or biometrics. We do not sell your data and do not provide it to third parties for their own purposes; the two exceptions at Google — the technical record of the request and the copy for abuse monitoring — are described in section "Our role and Google's role".

Put in the language of the US privacy laws as well, including the California Consumer Privacy Act: we do not sell personal data and we do not share it with third parties for targeted (behavioural) advertising — the app has no adverts, no advertising libraries and no advertising identifier — and we have not done so in the past either. Nobody pays us for the data that reaches Google; it goes to Google as our processor or on the terms in section "Our role and Google's role". Should this ever change, the change is material and we will show it to you in the app before it takes effect (section "Changes and languages").

The device identifier and the price tier. The app reads your device's Android system identifier, makes a fingerprint of it (SHA-256) and, from the first byte of the fingerprint, one letter: A or B. The letter chooses which of the two sets of PRO prices you see (the price tier). The fingerprint does not leave the phone and is not kept; it is computed afresh at every launch of the app, but the result is the same each time — your letter does not change. The letter goes into the local event log and into the record of the two statements on a purchase. The allocation of the letter is one-off and random and does not depend on your behaviour or your purchases — otherwise it would be profiling. Which plans exist for the letter in the country of your Google Play profile is decided by our setting in Google Play; where the two plans for the letter B are missing, you see the prices for the letter A. Why it exists and how it is disclosed on a purchase is in the Terms of Use, section "Subscriptions, clips by the piece, billing, withdrawal and termination"; the disclosure in the pre-contractual information — that the price has been set automatically by a device identifier, within a test with two price tiers (Article 47(1)(21) of the Consumer Protection Act) — is a planned fix in the app, which goes into the code before this text is published. The basis is our legitimate interest (Article 6(1)(f) of the Regulation). The reading of the identifier is access to information in your device under Article 4a(1) of the Electronic Commerce Act (Article 5(3) of Directive 2002/58/EC): the law requires us to give you the possibility of refusing it, and the exception for what is strictly necessary for the service you have asked for does not apply here — your

clip does not depend on the price tier. There is no switch for refusing it in the app today; you can object (section "Objection and withdrawal of consent"). Whether that is enough is being checked by our lawyer; if the answer is "no", the test with two tiers falls before it affects anyone.

The number of the two statements (nonce). On a purchase the app creates on the phone a unique number for the two statements (section "Subscriptions, purchases and the rights") and passes it to Google Play as an opaque identifier of the order (`obfuscatedAccountId`). It links the order to the statements you have ticked and is our proof that they were given. It contains nothing about you other than the moment of its creation.

The Firebase installation identifiers and the tokens of the integrity check — section "The app integrity check".

Your Google account. Google Play knows it — for the purchases and for the backup. We do not receive it.

Your IP address. Google sees it on every call the app makes to Google's services — that is the technical record of the request ("Service Data"), for which Google is a separate controller.

Our alignment server does not record your IP address: the request logs that carry it are switched off, and the application logs are without an IP address.

The model of the phone. The app reads the model, the processor and the memory of the device in order to choose a capture mode; that stays on the phone — only a "class" of the device (high, medium or low) goes into the event log.

The model of the device travels solely in an error report, if you have switched that checkbox on (section "Usage statistics and error reports — only with your consent").

The random number of the installation (`installKey`). It is created only when you switch on the statistics or the error reports, lives in a folder that is not included in the Android backup and is expressly excluded from the backup rules — it does not travel when you change phone. It lives until the checkbox is switched off, until 90 days without events, or until deletion on request. It is shown to you in "About & Legal" and it is the key with which a deletion request finds your statistics with us (section "Retention periods and deletion").

12. Android backup and transfer between phones

Only the settings file is copied — into Google Drive through the automatic Android backup and when the data is transferred to a new phone. The "My clips" library, the event log and the queue of the reports are not copied. What the file contains is listed in section "What is kept on your phone": the answer about age, the record of the acceptance, the two statements, the order number, the balance of the purchased clips, the counters, the refused clips, the preferences.

The backup is under your Google account and under Google's terms towards you. The app does not encrypt the file itself; the copy may reach Google Drive without encryption on the part of the app, and on a phone without a screen lock. Whether and how you can switch off the backup and the transfer on your side is decided by Android and by the maker of your phone; which of our files go into the backup is decided by us — only the settings file.

The restore happens without any action by you: on a reinstall Android brings the settings back without asking you. With them come back the daily counters as well, and the record that the cinematic welcome

clip has been used. That is why "once per installation" in practice means "once for the settings file" — and Android brings that back on a reinstall. The record of the acceptance also comes back — which is why the notice is not shown again on the new phone. Asking again on a new version of the text happens at the next launch (section "What is kept on your phone").

The row "Delete my local data" asks Android to replace the cloud copy at the next automatic backup; that is a request, not a guarantee of a moment, and the screen tells you so: "The backup of your settings in your own Google account is replaced at the next automatic backup."

The random number for the statistics does not travel in the backup and is not transferred — on the new phone the checkboxes are asked afresh.

13. Things you send out

Everything in this section happens only by an action of yours.

"Save to gallery". It saves a copy of the clip in the phone's public media library — there it is visible to every app with access to the gallery and, if you have a cloud backup of your photos switched on (Google Photos, for example), it goes into your account there. Uninstalling Dubroll does not remove it; the row "Delete my local data" does not delete it — the copies in the gallery are yours and we do not change them.

"Share". The Android system share sheet passes the clip to the app you have chosen, together with a text with hashtags: "#Dubroll #AIParody" and up to 3 words from the commentary itself. Words that describe your video become public text if you publish it. The same text is copied to the phone's clipboard on every share — and the clipboard is readable by other apps too. The shared file carries the AI marking (the metadata and the visible notice), the brand watermark if the clip is branded, and the inherited metadata of the recording (section "What goes to Google Gemini and why" for the fix concerning the copy for the model). You publish — Dubroll publishes nowhere.

Import from the gallery. The "Gallery" button opens the Android system picker, which gives the app a copy of the selected video; the original is not changed, and the normalised copy for the model is without the metadata of the original.

The files in the "Downloads" folder. On acceptance of the Terms of Use and of this Policy the app saves the two documents as PDF files in the "Downloads" folder of your phone — a medium that you hold and that the app never changes, so that it serves as a durable medium within the meaning of § 13(32) of the Consumer Protection Act. The app never changes and never deletes a file it has saved in "Downloads" — that is a rule in the code, with a test. The files are yours; after an uninstall Android takes our access to them away.

The confirmation on withdrawal and the file with your data. On withdrawal from a contract the confirmation is saved in the same folder, and on a request for your data — the file with it.

14. The permissions

The list comes from the assembled manifest of the app, not from an intention. The permissions for camera and microphone are asked for only after you have accepted the Terms of Use and this Policy.

Permission	Why it exists
Camera, Microphone	The core: the recording is video with sound. They are asked for after the acceptance, when the camera is first opened.
Internet	The requests to Google and to our components, described in section "How Dubroll is built".
Foreground service of type "data synchronisation"	It keeps the network connection while the clip is being made, even if you leave the app; the progress notification is the price of it.
Notifications	Asked for at the first production, for the progress notification ("Making your clip", "Your clip is ready 🎬"); refusing breaks nothing.
Vibration	Haptic response of the buttons.
Writing to storage (only up to API 29)	For "Save to gallery" on older versions of Android. Reading from storage has been expressly removed.
Brought in by libraries	Network state, keeping the processor awake (wake lock), Google Play Billing and Google services — from the libraries for purchases and Firebase.

What we do not ask for: location, contacts, biometrics, an advertising identifier, reading of media files, Bluetooth.

"Visibility to other apps" (queries in the manifest): the system action "process text", the speech synthesisers on the phone (for the backup voice) and Google Play Billing.

Components open to the outside: the library through which we reach Gemini indirectly includes, as a dependency of a dependency, a sign-in library (Firebase Auth) that the app does not use — it brings in two activities accessible to other apps, even though there is no sign-in in Dubroll. A planned fix: the target state is for those activities not to be accessible from outside, or for the dependency to fall away.

What the video itself may carry in it (the metadata of the recording) — section "What goes to Google Gemini and why".

15. Subscriptions, purchases and the rights

Google Play Billing. The purchase of PRO or of clips by the piece is concluded and paid for in Google Play. The app asks Google Play about your purchases at launch and when "Restore purchases" is pressed, and shows only prices that have come from Google Play. The asking at launch happens after the acceptance of the Terms of Use and of this Policy and carries nothing of yours. Who is the merchant of record for the transaction, which the countries are and how the subscription is terminated is in the Terms of Use, section "Subscriptions, clips by the piece, billing, withdrawal and termination". How you learn of a change in the price for renewal is in the same section of the Terms of Use; with us, of the prices only the price shown to you remains, in the record of the two statements (section "What is kept on your phone").

What Google receives and what we do not see. Google processes your payment data — card, account, billing address — as a separate controller, under its own terms towards you. We do not see your card. Google Play gives the app a purchase token (an opaque identifier of the purchase itself), the order number (it starts with "GPA."), the product and the plan and the time of the purchase.

The two statements. Before a purchase with a right of withdrawal you tick the two statements under the Consumer Protection Act — the express request that the service should start at once and the

confirmation of the consequence (Article 49(9)). Today's code asks for them both on the subscription and on the clips by the piece; the text of the two statements differs, because the pack is a service with prepaid units, not a subscription with renewal, and the two records are kept separately on your phone. Their record is made on your phone before the purchase (section "What is kept on your phone"), and their unique number goes to Google Play in the order itself (section "Identifiers, advertising and tracking"). The record and the number are our proof that the statements were given — the burden of proof is on us (Article 49(10) of the same Act). The basis is our legal obligation (Article 6(1)(c) of the Regulation).

PRO on the device. The app keeps on your phone when PRO was last confirmed. PRO stops on the device only after a proven online check with Google Play at which an active subscription was not returned; where there is no connection PRO keeps working for a short period after the last successful check; how long that period is, is in the Terms of Use, section "Subscriptions, clips by the piece, billing, withdrawal and termination".

15.1. Our server for the rights

Whether you have PRO and how many purchased clips are left to you is checked with Google Play through a server of ours, and not against a record on the phone, which can be forged. The server runs in Google Cloud Run, region `europa-west1`, in our main cloud project; Google is a processor under the Cloud Data Processing Addendum.

What it does. After the acceptance of the Terms, the app sends it the tokens of the purchases that Google Play has returned for this installation, and the records of the two statements for which a purchase already exists. The server checks every token with Google Play (Play Developer API), acknowledges the purchase on our side and returns the state: whether PRO is active, which plan, until when, whether it renews automatically, how many purchased clips are left, how many cinematic clips have been made today. From Google Play the server also receives the time of expiry, the automatic renewal, the state of the subscription according to Google Play and the country of your Google Play profile — the last of these is a country of the profile, not a location. Separately, Google Play sends us real-time messages on renewal, termination and refund (RTDN); the raw message is kept for 90 days. Once a day the server asks Google Play about voided purchases. If Google Play refunds an amount or voids a purchase at its own discretion, we learn the fact, the date and the reason code as Google Play gives them — not your correspondence with Google; from that fact PRO stops and the remainder of the purchased clips becomes zero.

What it keeps and for how long. For every purchase — a document by a fingerprint of the token: the token, always encrypted at rest and decipherable when needed for up to 90 days after the end of the purchase, and replaced by its fingerprint alone after that, the kind, the product and the plan, the letter of the price tier, the order numbers, the country of the Google Play profile, the price and the currency as Google Play gives them, the times, the state according to Google Play, the remainder of purchased clips, the daily counter of cinematic clips for the last 7 days, the number of the two statements. For every purchase — also a full record of the two statements (time, version of the text, language, plan, price tier, price shown, version of the app, fingerprint of the accepted Terms) — that is our proof and it remains even if you delete the data on your phone. A record of statements without a purchase is deleted after 30 days. The documents are kept for 3 years after the end of the last purchase — expiry, revocation or refund — the limitation period for periodic payments and for compensation and penalties from a contract that has not been performed under Article 111(b) and (c) of the Obligations and Contracts Act;

the sales reports from Google Play (order number, country, product, currency) are kept in our accounts for 10 years under Article 12 of the Accountancy Act.

The daily amounts of PRO are on the subscription, not on the device. The server counts the cinematic clips for the day against the purchase — two phones with one Google account share the same included clips (the cinematic clips and the short clips are counted separately). The local counter remains for working without a connection.

Clips by the piece. The purchase is acknowledged immediately by the server, the remainder is held with us against the purchase, and the marking of the purchase as consumed in Google Play happens when the remainder is zero; so on a reinstall Google Play returns the unspent purchase and the server returns the remainder — the purchased clips do not expire and are restored. While you have a remainder from "Clips by the piece" of a given size, the same size cannot be bought again; the Terms of Use say so word for word.

Bases. The contract with you (Article 6(1)(b) of the Regulation) for the check, the access and the counters; the legal obligation (point (c)) for the accounting reports; our legitimate interest (point (f)) for keeping them after the end of the purchase — a defence in a dispute about withdrawal, refund or lack of conformity. You can object to the keeping under Article 21 — the document is marked as restricted at once, without argument, and remains only for keeping until the end of the period.

Access. The service account of the server — only to those documents; Google Play — through a service account with the smallest rights, invited into the Play Console; the manager and an authorised deputy. We do not use the Google Play service that returns the buyer's address.

Recipients by law. The company's accountant receives the sales reports from Google Play; a lawyer — in a dispute.

15.2. The withdrawal function and the page of the contract

For every contract — the subscription and every purchase of clips — the app shows a withdrawal function labelled with the words "withdraw from contract here", visible and available from conclusion until the end of the 14-day withdrawal period, and a separate confirmation function labelled only with the words "confirm withdrawal" (Article 11a of Directive 2011/83/EU, introduced by Directive (EU) 2023/2673). What data is processed on a withdrawal: the text of your statement, your name — only if you enter it, the field is optional — the date and the time (of the device and of the server), the order number, the indicative proportional part for the days elapsed, the channel you have chosen for the confirmation and — only if you choose "email" — the email you have entered, which is deleted 30 days after a delivered confirmation. The confirmation is recorded on a durable medium of your choice: a file in "Downloads", an email or sharing to an app of your choice. The basis is our legal obligation to accept the withdrawal and to confirm to you that we have received it (Articles 52 and 54 of the Consumer Protection Act; Article 6(1)(c) of the Regulation); the function itself is under Article 11a of Directive 2011/83/EU, which we implement whether or not it has been introduced into Bulgarian law. The conditions of withdrawal and the money are in the Terms of Use.

The page of the contract in the app shows the product, the price and the currency as Google Play gives them, the period, the start and the next renewal, the state according to Google Play, the accepted version of the Terms and of this Policy — the "Effective from" date — and a link to the PDF file in "Downloads", and the text of the two statements. There is no web page with a sign-in: without accounts there is no way for us to recognise you in one.

16. The reports

The "Report" button opens the "Report this clip" screen with reasons from a fixed list: "Harassment or bullying", "Sexual content", "Hate or discrimination", "Child safety", "Third-party rights" (a breach of rights over image or voice), "Other". The report carries: the time, the reason, the commentator and the language of the reported clip (those of the REPORTED clip, not the ones currently selected on the screen), the model, the version of the app, an identifier of the clip (formed from the moment of its creation, to the millisecond; known only to your phone), a fingerprint of the text of the commentary and its length. It does not carry the text, it does not carry the video, it does not carry an identifier of a user or of an installation.

The report is first written into a queue on your phone — up to 90 days and up to 200 rows — and is sent only when there is a receiver to receive it; without a receiver it stays on the phone.

The receiver. The reports go to a receiver of ours in Google Cloud (Firestore), region `europa-west1`, in our main cloud project. The day is recorded, not the hour. The sixth reason above — "Third-party rights" — is for a clip that affects a person who did not want to be in it. No installation identifier is recorded and no IP address is recorded (the request logs are switched off). The reports are pseudonymous — the identifier of the clip and the fingerprint of the text are keys that exist only on your phone — and we have no key with which to find a report by person. That is why we cannot carry out the deletion of a report on request, and we say so (Article 11(2) of the Regulation) — unless you yourself give us the identifier of the clip from your phone.

Period and review. The reports are kept for 12 months, after which only aggregate counts without identifiers remain. The manager reviews them weekly under a written procedure; the note from the review carries the reason, the commentator and the action taken — with no free text about the clip. The person reporting is not notified — there is no channel by which we could reach them.

Purpose and basis. Built-in reporting of AI content is a Google Play requirement for apps with generative artificial intelligence and a safety measure of ours; the reports also show us when which commentator or language crosses the line. The basis is our legitimate interest (Article 6(1)(f) of the Regulation) — the safety of the product and of the people in the frame.

17. Our role and Google's role

Torvyx Labs Ltd. is the controller for the collection of your video and sound, for their transmission to Google, for the writing and voicing of the commentary, for the integrity check and for everything the app keeps on your phone. We decided what the app collects, which service receives it, how it is presented and what is asked of it. The fact that the content does not touch a server of ours does not change that.

Role by service:

- **Gemini and Gemini TTS through Firebase AI Logic (Google Cloud)** — Google is our **processor of personal data** under the Cloud Data Processing Addendum. One named exception: the copy of a flagged request for abuse monitoring (section "What goes to Google Gemini and why") is a purpose of Google's. Google states that there too it stands in the role of a processor, as a security measure for the service; a supervisory authority may see it as a separate controller. We do not present one of the two readings as settled — we state both.

- **The technical record of the request** (number of tokens, status, result of the filters, IP address) — "Service Data" under the Google Cloud Privacy Notice; Google is a **separate controller** for it.
- **Firebase App Check / Google Play Integrity** — Google processes under the Firebase Data Processing and Security Terms; the signals of Play Integrity are under the terms of the Google Play Integrity API. Which is under the Cloud Data Processing Addendum and which under the Firebase Data Processing and Security Terms for each service we are checking with our lawyer and will state here, on every row of this list.
- **Google Play (purchases)** — Google is a **separate controller** for your payment and for your Google Play profile.
- **Google Drive (the backup of the settings)** — under your Google account, under Google's terms towards you.
- **The phone's own speech synthesiser** (the backup voice) — its provider (Google, Samsung or another) is a separate provider with terms of its own.
- **Google Cloud Run and Cloud Logging for our alignment server** — Google is a **processor**; **Google Speech-to-Text v2** — **another processor** (a sub-processor) through our server, for six languages.
- **Firestore, Cloud Storage, Cloud KMS, Cloud Run, Pub/Sub and Cloud Logging for our server components** (rights, reports, statistics, rights requests) — Google is a **processor**.
- **Firebase Hosting for dubroll.app** — Google is a **processor**; the Hosting request logs (with an IP address) stay with Google "for a few months" under the Firebase documentation. The site is static, with no cookies, no analytics, no forms.

Which Google company. Our counterparty for Google Cloud and for Firebase is a company from the Google group; we will name it here once we have checked the contract and the invoice.

Other recipients. Google — as the provider of our contact@dubroll.app mailbox (Google Workspace), for the letters you send us yourself; professional advisers (a lawyer, an accountant) under an obligation of confidentiality; a public authority, where the law obliges us. We have no other recipients. We do not sell data and do not provide it for other people's purposes beyond the two exceptions at Google named above.

You as a controller. When you publish a clip with other people in it before an indefinite circle of people, for the data of the people in the frame the controller is you — section "People in the frame". Whether in some cases we are joint controllers with you under Article 26 of the Regulation is a question this document does not settle on its own: we do not assert it and we do not deny it.

The Google documents referred to in this section and in section "Which Google terms apply" are at the following addresses: Google Cloud Platform Terms of Service — <https://cloud.google.com/terms/>; Service Specific Terms — <https://cloud.google.com/terms/service-terms>; Cloud Data Processing Addendum — <https://cloud.google.com/terms/data-processing-addendum>; Google Cloud Privacy Notice — <https://cloud.google.com/terms/cloud-privacy-notice>; Firebase Data Processing and Security Terms — <https://firebase.google.com/terms/data-processing-terms>; the Agent Platform documentation on abuse monitoring — <https://docs.cloud.google.com/gemini-enterprise-agent-platform/models/abuse-monitoring> — and on the location of the data — <https://docs.cloud.google.com/gemini-enterprise-agent-platform/resources/data-residency>. We do not point to the documents of the Gemini Developer API — they describe another service with another data regime, which we do not use.

18. Legal bases by purpose

For every purpose — its legal basis. We do not rely on consent for the service itself: consent that you cannot meaningfully refuse is not consent.

Purpose	Legal basis
Sending your video and sound to Google Gemini and writing the commentary at your request	The contract with you — Article 6(1)(b) of the Regulation. The acceptance of the Terms and of this Policy on first launch is information and acceptance of a contract, not "consent" under point (a).
The data of the people in the frame — image, voice, behaviour, setting — in the same sending, including children	Legitimate interest — Article 6(1)(f) (ours and yours: that the clip you have asked for is made), with a recorded balancing assessment. We do not seek special categories of data and do not rely on a condition under Article 9(2); Google's safety filters and our instructions to the model are a partial safeguard, not a guarantee.
Synthesising the voice (Gemini TTS)	The contract — point (b).
Fitting the captions on our alignment server	The contract — point (b).
The app integrity check (App Check / Play Integrity)	Legitimate interest — point (f): protection of the service from abuse; the access to the device is strictly necessary for the service asked for — Article 4a(4)(2) of the Electronic Commerce Act.
A fingerprint of the device identifier for the price tier	Legitimate interest — point (f) (a price test with a one-off random allocation); the access to the identifier on the device — Article 4a(1) of the Electronic Commerce Act, without the exception under paragraph 4(2) and without a switch for refusing it in the app today (section "Identifiers, advertising and tracking"); the right to object is live.
The number of the two statements, passed to Google Play with the order	Legal obligation — point (c): proving the statements (Article 49(10) of the Consumer Protection Act).
The answer about age (the verdict)	The contract with you — point (b): a step before its conclusion, with which we check legal capacity under Articles 3 and 4 of the Persons and Family Act. A reason for the threshold of 18 is also our contract with Google, which prohibits access by persons under 18 (Service Specific Terms → "Generative AI Services" → "Age Restrictions") — but that is a contract, not a law, and it is not a basis under point (c).
The record of the acceptance of the Terms and of this Policy	Legal obligation — point (c): the burden of proving that you received and accepted them is on us (Article 147a(4) of the Consumer Protection Act).
The daily counters, the streak of days, the pause after refusals — on the device	The contract — point (b): these are the rules of free use or of PRO.
The "My clips" library, the cache during production and the local event log — on the device	The contract — point (b): this is the service itself and its diagnostics on your phone; none of it leaves the device by this route.
The backup voice — the text of the commentary to the phone's own synthesiser	The contract — point (b): without it a clip without Google's voice would not come out.
The inclusion of the settings file in the Android backup	The contract — point (b): it carries the acceptance, the verdict about age and the purchases over when you change phone; the copy is under your Google account and under Android's rules.

Purpose	Legal basis
Checking the purchases with Google Play on your phone — purchase token, order number, product, plan and time; access to PRO and to the purchased clips	The contract — point (b).
Checking the purchases through our server for the rights and the daily counter of the subscription	The contract — point (b).
Keeping the documents on purchases, statements and withdrawals after the end of the purchase	Legitimate interest — point (f): a defence in a dispute about withdrawal, refund or lack of conformity, within the limitation period; your objection restricts it at once.
The sales reports in the accounts	Legal obligation — point (c) (Article 12 of the Accountancy Act).
The withdrawal function, the confirmation and the performance of the withdrawal	Legal obligation — point (c) (Articles 52 and 54 of the Consumer Protection Act; the function — Article 11a of Directive 2011/83/EU).
The reports of AI content	Legitimate interest — point (f): safety of the product and of the people in the frame; the Google Play requirement.
Usage statistics and error reports	Consent — point (a), withdrawable at any time; the writing of the random number onto the device — Article 4a of the Electronic Commerce Act.
Requests for access, deletion and objection; the record of the requests themselves	Legal obligation — point (c) (Articles 12 and 15–21 of the Regulation; accountability under Article 5(2)).
Answering a letter you have sent us yourself	Legal obligation — point (c), where you are exercising a right; otherwise legitimate interest — point (f): answering you.
Notifying the supervisory authority and the persons concerned of a personal data breach	Legal obligation — point (c) (Articles 33 and 34 of the Regulation).

What is compulsory and what is not. The video and the sound are necessary for the contract itself: without them there is no commentary and no clip. The answer about age is compulsory — without the confirmation the app does not start; whoever does not confirm simply does not continue and can try again: the app does not lock devices and there is no "blocked" record. The integrity check is compulsory for every request to Google — without it a new clip cannot be made. On a purchase the purchase token and the order number come from Google Play — without them the purchase cannot be checked and credited, and the two statements are a condition only for a purchase with a right of withdrawal: without the tick the purchase button stays disabled, and free use does not depend on them. Everything else is your choice and refusing takes nothing of the service away from you: reporting a clip and writing to us.

Optional too are the statistics and the error reports — the two checkboxes; the service does not depend on them.

There is no row for direct marketing: we send no marketing messages and do not ask for your email for such.

19. Transfers of data to third countries

Google is a global company and Google LLC (USA) may process data under the Cloud Data Processing Addendum. The safeguards are two, one on top of the other: the European Commission's adequacy decision of 10 July 2023 (the EU–US Data Privacy Framework), under which Google LLC is certified, and the standard contractual clauses under Article 46(2)(c) of the Regulation as a fallback, should the decision fall away; the clauses are part of the Cloud Data Processing Addendum, which is at the address in section "Our role and Google's role" — from there you obtain a copy. We do not claim that the adequacy decision is permanent — it is being challenged and may be annulled; if that happens, the standard clauses remain. We do not rely on the exception (the derogation) under Article 49(1)(a) (explicit consent to a transfer) — we do not ask you for such consent. A copy of the safeguards: the standard contractual clauses are part of the Cloud Data Processing Addendum at the address given in section "Our role and Google's role"; on request at contact@dubroll.app we will send you the link and a copy.

Where your content is processed. The requests to Gemini go to Google's multi-region zone for the European Union ("eu"), for which Google documents that the processing by the model and the copy for abuse monitoring stay in Member States; under Google's documentation on the location of the data, the United Kingdom and Switzerland are not part of "eu". The voice is the exception: the voice model is sought first in the "eu" zone, and when it is not offered there — as has been the case since 11 September 2026 — the text of the commentary is synthesised in Google's global infrastructure and may be processed outside the European Union, at Google LLC, under the two safeguards at the start of this section. The video and the frames do not take that path — the vision stays in the "eu" zone. This happens under the same contract and under the same rules on recording (section "What goes to Google Gemini and why"). The "eu" zone and the global path for the voice are set in the code of the app; should Google bring the voice model back to "eu", the app returns there on its own. Administrative access by Google from outside the EU for support remains possible under the Cloud Data Processing Addendum.

Our alignment server is in `europa-west1` (Belgium); Google Speech-to-Text v2 for the six languages — in the "eu" region.

Our server stores (rights, reports, statistics, rights requests, logs) are in `europa-west1`; the key for encrypting the tokens — as well.

Google Play and the integrity check are global Google services and may be processed outside the European Economic Area — Google Play as a separate controller, under Google's terms towards you; the integrity check — by Google as our processor under the Firebase Data Processing and Security Terms. Which transfer safeguards under Article 46 those terms provide we are checking with our lawyer and will state here.

Wherever you live: data stored at Google may be requested by a public authority under the law of the country in which it is stored. That is not something we can change by contract.

20. Retention periods and deletion

20.1. For how long

What	Where	Period
The "My clips" library — clips, cover frames, index with the full text of the commentary	your phone	until you delete the clip or the local data, or uninstall the app
Raw recording, copy for the model, audio of the voice, frames, intermediate render	your phone (cache)	until the end of the production; leftovers — until the next launch of the app: our files are deleted at once, and the raw recordings when they are older than about 1 hour; on a failure the raw recording stays for "Try again"
The event log	your phone	at 1 MB the current file is archived, the previous archive is deleted
The queue of the reports	your phone	up to 90 days and up to 200 rows
The settings file (answer about age, acceptance, statements, order, counters, preferences)	your phone and the backup in your Google account	until uninstall or deletion; in the backup — under Android's rules
Your video and sound, the text to the voice	Google Gemini and Gemini TTS (processor)	transient; a short-term cache of up to 24 hours, until it is switched off for our two projects; a request flagged by a classifier — up to 90 days in the region; for the "Advanced AI" class of models — 30 days for all requests (not the model we use)
The token of the integrity check and the installation identifier	Google (Firebase)	under the Firebase terms; we do not record them
Your payment data and your Google Play profile	Google Play (separate controller)	under Google's terms towards you
The backup of the settings	Google Drive, your account	under the rules of Android and of your account
The text of the commentary passed to the phone's own synthesiser (backup voice)	the provider of the synthesiser (Google, Samsung or another)	under its terms
The synthetic voice and the text of the commentary for alignment	our alignment server	transit — not kept; application logs without content and without an IP address — 30 days
Documents on purchases, statements, withdrawal requests	our server (Firestore)	3 years after the end of the last purchase (expiry, revocation or refund); the token in clear form — 90 days after the end; raw messages from Google Play (RTDN) — 90 days; an email for confirming a withdrawal — 30 days after a delivered confirmation; a record of statements without a purchase — 30 days

What	Where	Period
Sales reports from Google Play (order, country, product, currency)	our accounts	10 years from the first of January of the year following the reporting year (Article 12 of the Accountancy Act)
Reports of AI content	our receiver (Firestore)	12 months; after that only aggregate counts
Usage statistics (with consent)	our receiver (Cloud Storage)	raw events 90 days; aggregate data without an identifier — no period, because without an identifier it is no longer personal data and the rule under the table does not apply to it
Error reports (with consent)	our receiver (Cloud Storage)	90 days
Rights requests (pseudonymous records)	our server (Firestore)	12 months
Letters to support	the mailbox	12 months after the matter is closed
Backups of the server databases (point-in-time recovery)	Google (processor)	7 days
Register of security incidents	with us	5 years
The Firebase Hosting request logs for dubroll.app (with an IP address)	Google (processor)	"for a few months" under the Firebase documentation

The rule behind the table: every field has a reason and a period; "just in case" is not a reason (Article 5(1)(e) of the Regulation).

20.2. The ways to remove data

First — "Delete clip" in the "My clips" library: it deletes the clip, the cover frame and the entry in the index, including the text of the commentary. A copy saved to the gallery stays there — it is yours.

Second — the row "Delete my local data" in the "Your data" section of the "About & Legal" screen: it deletes the "My clips" library, the event log, the queue of the reports, the cache and your settings, writes a deletion marker and asks Android to replace the cloud copy at the next automatic backup. WHAT STAYS: the record that you accepted the Terms and this Policy; your answer about age (it is given with the same single press as the acceptance — had it fallen alone, you would never have been asked again); the statements made at purchase; today's counters (the daily quota, fair use, the cinematic cap) and the pause after clips blocked by moderation. The first three are proof we are required to keep; the counters are there so that this row does not become a way to give yourself a new day, and the pause so that it does not become a way to clear a penalty. The confirmation dialog lists exactly this. It also deletes the record on the phone of the purchased cinematic clips — which the confirmation dialog does not name either: their balance is part of the settings (section "What is kept on your phone"); the Terms of Use say so in section "Your data, in one paragraph". What it does NOT delete: the copies in the gallery; PRO — it lives in your Google account and is restored with "Restore purchases"; the internal stores of Google's library (installation identifier, token of the check); the copy in Google Drive until the next backup.

The remainder of the purchased clips is also held with us, against the purchase itself, and is returned to the phone at the next check of the purchases — deleting the local data does not take it away from you (section "Subscriptions, purchases and the rights").

For the data with us there is the fourth way below.

Third — uninstalling. Android deletes the app's private storage, that is, everything from the first two ways. Three things are not removed by uninstalling: the copies in the gallery, PRO in your Google account and the backup of the settings in Google Drive, which Android brings back on a reinstall — and with it comes back the balance of the purchased clips.

The purchased clips are restored even without a backup — from our server, against the purchase itself (section "Subscriptions, purchases and the rights").

Fourth — a deletion request without an account — for the data with us. In "About & Legal" the button "Delete my data at Dubroll" sends to our server the random number of the installation (if you have ever switched the statistics on) and the tokens of the purchases that Google Play returns for this installation. The statistics and the error reports are deleted within 24 hours. The documents on purchases, statements and withdrawals cannot be deleted while they are necessary for the contract or for a defence in a dispute (Article 17(3)(e) of the Regulation) — they are marked as restricted and remain only for keeping until the end of the period; the receipt you get in "Downloads" says exactly that, the reason and your rights of complaint and of judicial remedy. Without the app — by email to contact@dubroll.app with the order number from Google Play (it starts with "GPA.") and the date of the purchase or the product as a second check; where there is doubt about identity we only restrict, we do not delete. Statistics cannot be found without the app — their key is only on your phone; for a person with no purchase and with the statistics not switched on there is nothing to be found with us. If you use Dubroll on two phones, the random number for the statistics is different on each: a request from one phone deletes the statistics only on it, and from the other you make a separate one; the documents on purchases are common to your Google account and one request is enough for them. Period: up to one month.

Content already sent to Google Gemini we cannot recall, check or delete — it has been processed and discarded as described in section "What goes to Google Gemini and why", and the flagged requests stay with Google under its terms.

21. Your rights and the honest limits

You have the rights under Articles 15–22 of the Regulation: access, rectification, erasure, restriction, portability, objection, and the right not to be subject to a decision based solely on automated processing. Here is what they mean at Dubroll — and where they end.

Access. The "My clips" library shows you every clip and an extract from its commentary; the full text of the commentary, the settings file and the event log are in the app's private storage and have no screen that displays them — what they contain is described in section "What is kept on your phone". If you want an extract from them, write to us at contact@dubroll.app: we will describe to you what each of them contains and how to get it off your phone, and when the app has a way out for it, we will say so here.

Portability. "Save to gallery" gives you the clip as an MP4 file; the text of the commentary has no separate export from the app other than the words in the share text.

Rectification. The commentary is generated text; rectifying it means deleting the clip and making a new one.

Erasure. The ways are in section "Retention periods and deletion".

Restriction and objection — section "Objection and withdrawal of consent".

The limit. For content sent to Google Gemini we cannot carry out access, rectification, erasure or portability: we do not hold it and we cannot identify you in it — the only identifier in the request that comes from the app is the token of the installation, which is not linked to you (Article 11(2) of the Regulation). That is a description of the architecture and it applies to us as well: we too cannot take it out.

For the data with us the limit is a different one and Article 11 does not apply to it: the documents on purchases can be found by the purchase token or by the order number, and the statistics — by the random number of the installation from your phone. For them access, portability, restriction and erasure are feasible — the button "My data at Dubroll" in "About & Legal" writes into "Downloads" a file with all the documents about you.

Period and form. We answer within one month of receiving the request; where the matter is complex the period may be extended by two further months, which we tell you within the first month, with the reason. No fee, no form — a letter to contact@dubroll.app is enough. If we have a reasonable doubt about your identity, we may ask for additional information in order to confirm it — and only that much.

Outside the European Union, including the United States and California. If you are outside the European Union, your rights of access, rectification, erasure and objection are the same as those described in this section and are exercised by the same route, at the same address and within the same period.

United Kingdom. If you are in the United Kingdom, in addition to the Regulation — which applies to us because we are established in the European Union — your processing is also subject to the law of the United Kingdom: the UK GDPR (Regulation (EU) 2016/679 as it forms part of the law of the United Kingdom) and the Data Protection Act 2018. Your rights of access, rectification, erasure, restriction, portability and objection (Articles 15–18, 20 and 21 of the UK GDPR) are the same as those described in this section and are exercised in the same way, at contact@dubroll.app, within the same period. The supervisory authority for the United Kingdom is the Information Commissioner's Office (ICO); you have the right to make a complaint to it under section 165 of the Data Protection Act 2018 at <https://ico.org.uk/make-a-complaint/>, and you may also complain to us first. Your data is processed in the European Union, as described in section "Transfers of data to third countries".

A complaint to us. If you consider that we process your personal data in breach of the law, you may complain directly to us — at contact@dubroll.app, with no form and no fee; for users in the United Kingdom this is also the route under section 164A of the Data Protection Act 2018, and for everyone else it is the same route. We acknowledge receipt within 30 days and respond on the substance without undue delay, within the period in paragraph "Period and form", telling you what we did. If our answer does not satisfy you, you keep your right to complain to the supervisory authority — the ICO for the United Kingdom, and for the European Union the authority in section "Contact and complaints".

22. Objection and withdrawal of consent

Objection (Article 21 of the Regulation). You can object at any time, on grounds relating to your particular situation, to any processing that rests on a legitimate interest (section "Legal bases by purpose"). Those are: the data of the people in the frame — including if you are the person in the frame, without our identifying you; the integrity check; the fingerprint for the price tier; the reports of AI content; and the answering of a letter you have sent us yourself, where by it you are not exercising a right. Write to contact@dubroll.app with the subject "Objection". We stop the processing, unless we demonstrate compelling legal grounds which override, or it is needed for legal claims. The objection does not go through a switch in the app — we say so because we do not want you to look for one there.

You can also object to the keeping of the documents on purchases after their end; the objection also has a button — "I object to the keeping" in "About & Legal": the document is marked as restricted at once, without argument.

Withdrawal of consent (Article 7(3) of the Regulation). The only consent we ask of you is for the statistics and the error reports.

You withdraw it with the two checkboxes in the "Your data" section of the "About & Legal" screen — as easily as you gave it, with no uninstalling. It takes effect immediately for the future; what has already been collected is deleted within 24 hours. The withdrawal does not affect the lawfulness of the processing before it.

23. Automated decision-making and artificial intelligence

There are no decisions under Article 22 of the Regulation. Dubroll takes no decision based solely on automated processing which produces legal effects concerning you or similarly significantly affects you. Three things that could look like one, and why they are not:

- **The price.** The price tier is determined by a fingerprint of your device identifier — one-off, random and the same for all devices with that letter; it does not depend on your behaviour or your purchases (which plans are available for the letter in the country of your Google Play profile is decided by our setting there) and it is disclosed in the pre-contractual information — the disclosure itself in the app is a planned fix (section "Identifiers, advertising and tracking"). This is an automated setting of a price — which is why we disclose it under Article 47(1)(21) of the Consumer Protection Act — but it is not a decision under Article 22: it produces no legal effects and does not significantly affect you — it is a matter of two sets of prices for one and the same service, your price is shown to you before the purchase, access is not refused and neither your personality nor your creditworthiness is assessed.
- **The one-hour pause.** After three clips refused by the filters within 24 hours the app stops new clips for one hour ("Too many blocked clips in a row — the studio takes an hour's break."). That is a counter of refusals with a described effect, and not a decision about you: it does not assess your personality, it does not refuse you access permanently and it resets by itself.
- **A refusal of the generation.** Google's safety filters and our thresholds may stop an individual clip ("This clip isn't one our comedy AI can roast. Try a different scene."). The refusal relates to the particular scene, not to you; with us it leaves no trace other than in the counter above and in the local event log, and at Google the refused request may also be "flagged" (section "What goes to Google Gemini and why").

We do not profile people over time. There is no file about you and there are no inferred characteristics; what history of events and purchases we keep, when and for how long, is in sections "Usage statistics and error reports — only with your consent" and "Retention periods and deletion"; the learned tempo in the settings is about the voice, not about you.

Artificial intelligence by name. The commentary and the voice in the clip are artificially generated by an artificial intelligence system — Google Gemini and Gemini TTS. Regulation (EU) 2024/1689 (the Artificial Intelligence Act) applies to us: the transparency obligations under Article 50 apply from 2 August 2026, and the transitional period under Article 111(4) does not help us, because Dubroll was not on the market before that date.

Our role. In relation to Google's model we are a deployer — we use a system over which we exercise control; in relation to Dubroll as a product that puts out a synthetic voice under its own name, we are a provider. The exact classification will be confirmed by our lawyer; the transparency obligations under Article 50 we fulfil in both roles.

Emotions and biometric categorisation. We do not ask the model to recognise emotions or to categorise people by features of the body and the face. Our instruction is to comment on the action, the moment and the situation, and, in a scene that looks delicate or with a person who looks hurt, not to joke — that is a safeguard in favour of the person in the frame, not a function for recognising emotions; the commentary may nevertheless describe an expression or a mood in the way a viewer would describe it. Whether such a safeguard falls under Article 50(3) of the Artificial Intelligence Act is a question on which our lawyer will pronounce — we do not assert it and do not deny it on our own authority. The use of Dubroll for monitoring, assessing, ranking or punishing people in an employment, official or educational relationship, and for inferences about their emotions, is prohibited by contract as well — the Terms of Use, section "Acceptable use".

How the clip says it. Every clip carries metadata in the MP4 file saying that the commentary and the voice are artificially generated. The visible notice that the voice is AI is on every clip and in all cases: with PRO, without PRO and with a purchased cinematic clip. PRO and the purchased clip remove only the brand watermark; the AI notice is not part of the brand and is not for sale. It is in the language of the clip, except where that language's script cannot be drawn by the device's system font — there it is in Latin script ("AI voice"). That is a choice, not an omission: a notice shown as empty boxes looks fulfilled without being so. The metadata is the machine-readable layer under Article 50(2); it does not necessarily survive the re-encoding that the social platforms perform on upload — the visible notice is there for exactly that. Whether the audio from Gemini TTS also carries an embedded Google watermark (SynthID) and whether it survives re-encoding, we are measuring; we will state it here when we have the result. A signature under a content provenance standard (C2PA) we do not apply today; a server-side signature would require the clip to pass through a server of ours, which this document promises will not happen.

Is the clip a deep fake? As a rule, no. A "deep fake" under Article 3(60) of the Artificial Intelligence Act is generated or manipulated content that resembles existing persons or events and creates a false impression of being authentic. The commentators are invented and do not imitate real persons; the voice is not the voice of the person in the frame and does not present itself as such; the clip is evidently comedic. That is why the particular disclosure obligation under Article 50(4) of the same regulation is as a rule not triggered; if an individual clip nevertheless is such content, the obligation — for a work of an evidently satirical character — is disclosure in a clear and visible manner that does not hamper the viewing, and that is exactly what is planned for every clip — with the metadata every clip carries and

with the visible notice, whose scope today is limited (above, "How the clip says it"). A "deployer" under Article 3(4) of the same regulation is a person who uses an artificial intelligence system outside a personal non-professional activity: if you publish clips as part of a professional or commercial activity (a monetised channel, for example), on publication the deployer is you; in personal use you are not. In both cases you must not remove the notice — the Terms of Use govern this in section "The licence for the app and the rights in the clip".

The system can be confidently wrong. The commentary is a joke written by a model from frames and sound: it may be untrue, exaggerated or offensive, including towards the people in the frame, and it may describe something that is not in the frame. The instructions we give the model are instructions, not guarantees. The "Report" button exists for exactly that.

24. People in the frame

This section is the centre of the Policy. When you film, in the frame and behind the camera there are people — adults, children, passers-by, voices. That person is not our user. They have not heard of us and have not been asked.

What reaches Google about them. Their image, voice, behaviour and setting — the video with the sound, up to three times for one clip and once more on every repeated attempt after an error (section "What goes to Google Gemini and why"). Then — the text of the commentary that describes them, on every text-only sending.

The text of the commentary that describes them also passes in transit through our alignment server (section "Our alignment server").

What we do not do. We do not identify anyone. There is no facial recognition and there are no voiceprints: the faces and the voices in the frame are not biometric data under Article 4(14) of the Regulation until they are subjected to technical processing that identifies the person — and we have no such processing and do not want it. There is no file about anyone — neither about you, nor about the people you film.

What we have told the model. Our instructions to Google Gemini say — here translated from their English original — that people are welcome as a subject, but the joke is about what they are doing — the action, the moment, the situation — and never about the body, the weight, the skin, the face, the disability, the age, the ethnicity, the gender, the sexuality, the religion or the poverty of a real person; never about clothing or appearance, not even as a compliment; with children in the frame — warm and harmless, never at their expense; in a delicate scene, or with a person who looks hurt — no joke. A separate instruction prohibits guessing a person's gender from their occupation, the setting or their clothes. Five of the commentators carry a limit of their own as well — for example the anger of the "Angry Gamer" is aimed at the universe and at objects, never an insult to the person in the frame. These are instructions to a model, not guarantees: the model may breach them. Google's safety filters and our thresholds on top of them (sexual content — at generation, at the strictest threshold Google offers: it stops even at low probability; hate, harassment, dangerous content — at a medium threshold) are a partial safeguard, not a guarantee.

What you owe. You accepted it on the "Before your first clip" screen: "Only film people who are OK with it — no nudity, no violence, no mocking people without their consent, no filming copyrighted screens." Nobody may be filmed or recorded without their knowledge or against their express disagreement, except in the cases provided for by law (Article 32(2) of the Constitution of the Republic

of Bulgaria). Do not film children at their expense; do not film delicate scenes. While you film your family for yourself, the Regulation does not apply to you (Article 2(2)(c)); once you film in a public place or publish the clip before an indefinite circle of people, that is no longer a purely personal or household activity (Court of Justice of the EU, C-25/17 *Jehovan todistajat*, paragraph 42; C-101/01 *Lindqvist*, paragraph 47) and for the data of the people in the frame the controller is you — with everything that follows. The Terms of Use govern this in section "Dubroll is comedy, not fact. People in the frame".

The information we owe the person in the frame (Article 14 of the Regulation). Here it is, in one place: controller — Torvyx Labs Ltd., contact@dubroll.app; source of the data — the user who filmed them or picked the video from the gallery of their own phone (where such a video comes from and whether it was publicly available, we do not know); categories — image, voice, behaviour, setting, and the generated commentary that describes them; purpose — writing and voicing a comedy commentary on the video at the user's request; basis — legitimate interest (Article 6(1)(f)), with a recorded balancing assessment; recipients — Google as our processor and, if Google's voice refuses, the phone's own synthesiser on the user's phone — for the text of the commentary; transfers to third countries — section "Transfers of data to third countries"; period — transient, section "Retention periods and deletion"; rights — objection, access and erasure, section "Your rights and the honest limits", and a complaint, section "Contact and complaints". We cannot provide this information to the person in the frame personally: we do not know who they are and cannot identify them — the impossibility under Article 14(5)(b) comes from the fact that we cannot identify them, and not from the fact that we have no server. The measure the law requires in that case is public access to the information — this section — and the text below, which the user can show.

A recipient of the text of the commentary is also our alignment server, and for six languages — Google Speech-to-Text as another processor through it (section "Our alignment server").

A ready text you can show the person you are filming: "I am filming you with an app that sends the video and the sound to Google Gemini so that it can write and voice a comedy commentary. The clip stays on my phone until I publish it. What your rights are is written in the "People in the frame" section of dubroll.app/privacy."

Special categories. We do not seek and do not ask the model to infer from the frame racial or ethnic origin, religion, health, sexual orientation or other special categories of data under Article 9, and we do not rely on a condition under Article 9(2) for the service. Such data may be visible in the frame — the instructions to the model prohibit it from commenting on exactly those things, and that is a measure, not a guarantee. For children the particular protection under the Regulation applies; they appear here as people in the frame, not as users — Dubroll has no users under 18 (section "Age — Dubroll is 18+").

The rights of the person in the frame and how to exercise them. Write to contact@dubroll.app — you need not be a user, nor have the app; describe what you know: when and where it was filmed, who filmed it, where the clip was published. We answer within one month. The objection under Article 21 is live even without our identifying you — the subject "Objection" is enough; the same route is open to a parent or guardian of a child in the frame; a complaint to the supervisory authority — section "Contact and complaints". The limits: deleting the clip is in the hands of the user who made it — the clip is on their phone and we do not hold it; we cannot delete it from someone else's phone, nor take it down from a platform on which it has been published (for a takedown, use that platform's own mechanism for reporting infringements); we have no identifier with which to stop a particular device or person — what we really can do is in section "Suspension and termination" of the Terms of Use; we do not know who the user is and cannot tell you; content already sent to Google we

cannot recall (section "What goes to Google Gemini and why"). What we can do: look at the particular case you describe to us, and change the instructions and the thresholds when the reports show it. The "Report" button is in the user's hands, not yours, and it has a "Child safety" category. The Terms of Use govern the same case from the contractual side — section "Reporting AI content".

It also has a category for infringement of a third party's rights — image or voice.

The worst case. The worst realistic case is a published clip in which an artificially written commentary mocks a real person — a child, an acquaintance, a stranger — in front of people who know them, without their having wanted it. Our measures against it, with their status: the instructions to the model (in force; not a guarantee); the caps on the length of the recording (in force); Google's filters and our thresholds (in force; partial); processing only by a processor under Article 28 and transient (in force); the clip is on the device until publication (in force); deletion from a single screen (in force); objection without identification (in force); the prohibition in the Terms on filming people without consent and on removing the notice (in force; contractual). The impact assessment under Article 35, the balancing assessment and the record of processing activities under Article 30 are internal documents of ours, reviewed with every release.

25. Age — Dubroll is 18+

Dubroll is for persons aged 18 and over. Three reasons, each separate from the others. The first is our contract with Google: the terms of the service through which we reach Gemini prohibit the app from being directed towards, or likely to be accessed by, persons under the age of 18 (Service Specific Terms → "Generative AI Services" → "Age Restrictions"). The second is the contract with you: the subscription is a contract which a person between fourteen and eighteen concludes only with the agreement of a parent or a guardian, while for a person under fourteen the parents or the guardian act (Articles 3 and 4 of the Persons and Family Act; on their own, minors conclude only ordinary small transactions for current needs) — we do not claim that such a contract would be "invalid", we claim that without that agreement it is not binding. The third is the subject matter: comedy about real people, which may be offensive.

A statement, not a check. On first launch the app asks once, on the screen before the first clip: by pressing you confirm that you are 18 or older. **A date of birth is not asked for.** We keep the confirmation, the threshold and the time. We know what this choice costs and we say it: a screen that shows the threshold is easier to get around than a neutral question about a date — but a date is just as unverifiable, and collecting it would be personal data without benefit. So we rely on your word, on the age rating in Google Play, and on the fact that we do not advertise Dubroll to minors.

The app does not lock devices. There is no screen telling you "that is the end", and no "blocked" record travelling in the backup. If you are not 18, simply do not confirm — Dubroll does not start without the confirmation, but your device stays free. The rule that "the first answer is final" was dropped on 6 September 2026 together with the date of birth: it locked out real people over a mistyped digit and gave not a grain of assurance about their age.

Children. We do not knowingly process the data of a user under the age of 18. Children in the frame are a different matter — they are people in the frame and section "People in the frame" is for them. The threshold under Article 25c of the Bulgarian Personal Data Protection Act — consent of a child under fourteen for information society services — does not apply to us, because we do not ask for consent from

children and do not offer the service to persons under 18. If we learn that a user is under 18, we stop the service as far as we can — the Terms of Use, section "Who may use Dubroll".

The documents on the purchase held with us in such a case are restricted at once and remain only for the refund of the amount and for a defence in a dispute, until the end of their period.

26. Security and personal data breaches

26.1. What protects your data

Mostly the architecture. There is no central store of your content — the video, the sound and the clips do not reach a server of ours, so there is no database of ours that could leak with all your clips. The requests to Google are protected by the integrity check, and the app's files — by the Android sandbox, in which another app cannot read them, and by the encryption of your device. What is missing: the app has no lock of its own with a PIN or biometrics, it does not prevent screenshots and it does not encrypt its files itself — whoever unlocks your phone sees your "My clips" library.

If your phone is lost or stolen. The "My clips" library, the text of every commentary and the settings are on the device itself, and they are protected by what protects the phone: the screen lock and the encryption of the device. We have no channel to your phone — we cannot delete anything remotely, nor find it; the means for that are those of Android and of your Google account. On the new phone PRO is restored with "Restore purchases", and the settings come back with the backup, if it was switched on (section "Android backup and transfer between phones").

The data with us does not depend on the phone: without the app you reach it and delete it by email with the order number from Google Play (section "Retention periods and deletion").

For our server components: an encrypted connection (TLS) for everything; the integrity check for every endpoint of ours; the database with no client access — only the service accounts of the services; the store with no public access; the purchase tokens — encrypted with a key that lives in the EU; no service account keys in files; logs without request bodies, without tokens and without an IP address; alarms on errors and on unusual volume; point-in-time recovery for the database — 7 days; every new endpoint is released with a test that is deliberately broken.

26.2. What could realistically be a breach

The places:

1. **The copy for abuse monitoring at Google** — content of flagged requests of many users, for up to 90 days. That is the largest real scenario for your content, and it is at our processor, not at us.
2. **Our developer account at Google Play and the process of signing the app** — if they were taken over, they would allow a substituted app to reach you.
3. **Our mailbox** — the letters you have sent us yourself.
4. **Our server components** — the documents on purchases and statements, the reports, the statistics. With a server a central incident at us becomes possible; that is what the measures in the previous subsection are for.

26.3. What we do on a breach

We record it in the register of incidents (the facts, the consequences, the measures) and assess the risk. Where there is a likely risk to people's rights and freedoms we notify the Commission for Personal Data Protection without undue delay and, where possible, within 72 hours of becoming aware (Article 33 of the Regulation). On an incident at Google we record both dates — when Google became aware and when it told us. A decision "we do not notify" is documented as well.

26.4. How we will tell you

We do not have your email in order to write to you. Where the risk is high we notify personally everyone with whom we have a channel — a person who has written to us or has entered an email for a confirmation of withdrawal. For everyone else we publish a communication that informs you in an equally effective manner (Article 34(3)(c) of the Regulation): in our listing on Google Play, which anyone can open without giving us data. When the in-app notice starts to be shown again with a new version, the communication will stand there too.

We also publish it at <https://dubroll.app/security> — the permanent place for security notices, which stays there also when there is nothing to report.

Where the process of releasing the app is compromised — there, where the attacker is not in control: in our listing on Google Play. Trust the source the attacker cannot substitute. Where in doubt, we publish.

27. Changes and languages

The dates. "Last updated" is the date of this text; "Effective from" is the date from which the version binds you.

Every version of this document has a permanent dated address on [dubroll.app: /legal/privacy/bg/<date>](https://dubroll.app/legal/privacy/bg/<date>) for the Bulgarian and [/legal/privacy/en/<date>](https://dubroll.app/legal/privacy/en/<date>) for the English version; the previous versions remain available in an archive and every page points to the other language.

The version of this Policy is the "Effective from" date — it changes on a material change and stays the same on an editorial correction, and the new version is accepted afresh at the next launch (the Terms of Use, section "Changes to the Terms"); the record of which version you accepted is kept on your phone (section "What is kept on your phone"). **A material change** takes effect only after you have seen it: the notice in the app is shown again with the new version and you accept afresh, before you make the next clip or purchase; what your rights are if you do not accept, including with a paid subscription, is in the Terms of Use, section "Availability, changes to the app and to the models". Material is a change that adds a new recipient or a new category of data, changes the region of processing, changes the model with an effect on recording at Google, adds a new purpose or changes what a subscription is needed for. Of a new purpose we tell you before we begin (Article 13(3) of the Regulation). Your "My clips" library is not locked on a change — it is on your phone.

Transfer of the business or discontinuation of Dubroll. If a company takes over our business, it takes over this Policy as well: the data with us is not passed to it on weaker conditions, we notify you before that in the app and here, and the new controller is entered in section "Who we are" of this document with a new version. If we discontinue Dubroll, we will tell you in advance in the app (the Terms of Use, section "Availability, changes to the app and to the models"); the "My clips" library stays

on your phone and the clips remain readable; the data with us is kept only until the periods in section "Retention periods and deletion" — the discontinuation extends not a single one of them.

On a new version of the notice your consent for the statistics and the error reports is switched off and asked afresh.

Languages. This document exists in Bulgarian and in English, with one and the same effective date; in no other language does it exist. The app also has an interface in other languages and may in time receive further ones — a language of the interface does not create a legal version of this document: with a Bulgarian interface you receive the Bulgarian document, with any other interface the English one, while the screens speak in the language of the interface. The languages of the clip are not languages of this document: they are the languages in which the commentary is written and voiced, and which of them are currently offered you see in "Clip language".

Which language prevails. Wherever the Bulgarian version gives you more than the English one, the Bulgarian version prevails: it is the first version of this document, because the information under a distance contract with a consumer is provided in Bulgarian (Article 47(7) and Article 5(1) of the Consumer Protection Act), and it is read under the rule that terms must be drafted in clear and unambiguous language and that, in case of doubt, they are interpreted in the way that is more favourable to you (Article 147(1) and (2) of the same Act). The English version is a translation of the Bulgarian one and prevails only in the case of an editorial divergence that does not affect your position — for example, a different word for one and the same concept; where the divergence changes a right, a period, an amount or an obligation, the Bulgarian text applies. The two versions are published together, with one and the same effective date, and this paragraph stands word for word in the "Miscellaneous" section of the Terms of Use and in the "Changes and languages" section of the Privacy Policy — if they ever diverge by even one letter, the version that gives you more applies.

28. Contact and complaints

Torvyx Labs Ltd. („Торвикс Лабс“ ЕООД) — 46 Tsar Asen I Street, Burgas 8000, Bulgaria · telephone +359 877 001 885 · **contact@dubroll.app**. We answer every request under this document within one month of receiving it (Article 12(3) of the Regulation).

The telephone and the call. The telephone is answered on working days between 10:00 and 18:00 Eastern European Time. A call is not recorded: we have no recording system and we keep no register of calls; the traffic data is kept by the telephone operator under its own law, not by us. What you tell us on the telephone we use only in order to answer. If the matter needs a written trace — a request under your rights, a withdrawal from the contract, a complaint — we will ask you to write to us at contact@dubroll.app: that way the evidence remains with both sides.

Supervisory authority. The lead supervisory authority for us is the Commission for Personal Data Protection (CPDP), Bulgaria — [address, telephone, email and website: to be filled in from the Commission's own details on the day of publication]. You have the right to lodge a complaint with it or with the supervisory authority in the Member State in which you live or work (Article 77 of the Regulation), as well as the right to an effective judicial remedy before the court of your habitual residence (Article 79). A complaint to the CPDP does not require you to have turned to us first; if you write to us first, most often we can put things right quickly.

Deletion and access without an account. The buttons "My data at Dubroll", "Delete my data at Dubroll" and "I object to the keeping" are in "About & Legal"; by email — with the order number from

Google Play (section "Retention periods and deletion").

For a discrepancy between this document and the app, or between the two language versions, write to us: we will correct the document so that it tells the truth, and we will answer you on the substance of your own case.