Public Summary of Training Content for Kimi K3

Version of the Summary:	*Version No.1*
Last update:	*October 22, 2026*

1. General information

1.1. Provider identification

Provider name and contact details:	BEIJING MOONSHOT TECHNOLOGY CO., LTD. Email: governance@moonshot.ai
Authorised representative name and contact details:	European Data Protection Office Email: info@edpo.com

1.2. Model identification

Versioned model name(s):	Kimi K3
Model dependencies:	None. Kimi K3 is natively pretrained rather than a modification of a previously marketed general-purpose AI model. Models used to generate synthetic training data are addressed in Section 2.5 (Synthetic data) of this Summary.
Date of placement of the model on the Union market:	16/07/2026

1.3 Modalities, overall training data size and other characteristics

This Section requires general information about the overall training data after pre-processing and before the training of the model.

Modality *Select the modalities present in the training data, to the extent that they are identifiable*	**Training data size** *For each selected modality, select the range within which the estimated total training data size for that modality falls. Dynamic datasets may be excluded from the estimation.*	**Types of content** *For each selected modality, provide a general description of the type of content that has been included in the training data.*
☒ Text	☐ Less than 1 billion tokens ☐ 1billion to 10 trillions tokens ☒ More than 10 trillions tokens Alternatively, specify the approximate size in a different measurement unit: *Replace this with your response…*	Web text, source code, mathematical materials and knowledge-oriented content, together with instruction-following, reasoning and task-oriented training examples.
☒ Image	☐ Less than 1 million images ☐ 1Million to1 billion images ☒ More than 1 billion images	Images and associated text; interleaved image–text documents; materials for optical character recognition, visual understanding and visual programming; and rendered visual materials.

☐ Audio[1]	☐ Less than 10 000 hours ☐ 10 000 to1 million hours ☐ More than 1 million hours	
☒ Video	☐ Less than 10 000 hours ☐ 10 000 to1 million hours ☒ More than 1 million hours	Video content, including long-form video and associated multimodal training materials.
☐ Other	*Specify the modality and for each one indicate approximate size and unit of measurement*	

Latest date of data acquisition/collection for model training: 01/07/2026

Description of the linguistic characteristics of the overall training data: Multilingual data, with Chinese and English coverage and other languages from around the world

Other relevant characteristics of the overall training data: N/A

Additional comments (optional): N/A

2. List of data sources

This Section requires information about specific sources of data used to train the general-purpose AI model. In this section "dataset" should be understood as a single, pre-packaged collection of data. The filtering and pre-processing of data collected from the same pre-packaged collection should not be considered a new dataset to be disclosed separately in the sections below. If a particular dataset can be assigned to more than one of the categories below, providers should select the most relevant category and only report the dataset in that category, except in the case of synthetic data (see Section 2.5).

2.1. Publicly available datasets

This Section requires information about datasets that were used to train the model and which have been compiled by a third party, are made available publicly for free, and are readily downloadable as a whole or in predefined chunks, such as datasets and collections available on public repositories and online platforms, specialised websites, or snapshots of common crawl. The public availability of the datasets for free does not mean that the content at issue is necessarily free of rights since it may be subject to licensing arrangements or conditions of use (e.g., certain free and/or open licenses may determine the scope of the uses, including prohibiting uses relating to model training).

A dataset is considered to be "large" if the total data size for any one of the modalities contained in the dataset exceeds 3% of the size of all publicly available datasets for that modality used for training. The size of the dataset should be based on its size after pre-processing (for example filtering), and without splitting the dataset to prevent reporting circumvention.

Have you used publicly available datasets to train the model? ☒ Yes ☐ No

[1] Excluding audio that is part of video, as this should be reported under the "video" modality instead. Furthermore, the Commission understands the modality of 'audio' to include 'speech'.

If yes, specify the modality(ies) of the content covered by the datasets concerned:	☒ Text ☒ Image ☒ Video ☐ Audio ☐ Other *If so, please specify…*
List of large publicly available datasets:	The training data for Kimi K3 includes text from Common Crawl.
General description of other publicly available datasets not listed above:	The visual training corpus includes open-source datasets alongside in-house data preparation pipelines. Across the visual corpus, the content includes image-caption pairs, interleaved image-text documents, optical character recognition (OCR), visual perception, video, and visual programming data. Data preparation includes filtering, synthetic data generation, and deduplication.
Additional comments (optional):	When processing public web datasets, Moonshot AI identifies and respects copyright reservations and opt-out signals such as robots.txt rules for our crawler on all listed domain.

2.2 Private non-publicly available datasets obtained from third parties

This Section requires information about private non-publicly available datasets of third parties that are not publicly available and not disclosed under Section 2.1. These include:

1) *datasets for which transactional commercial licensing agreements were concluded between the provider and the rightsholders or their representatives, including by collective management organisations and legitimate content aggregators who have the right to collectively license works on behalf of rightsholders (Section 2.2.1);*

other private datasets obtained through data intermediaries, non-publicly available databases and datasets of third parties for which transactional commercial licenses have not been concluded with rightsholders or their representatives (Section 2.2.2).

2.2.1. Datasets commercially licensed by rightsholders or their representatives

Have you concluded transactional commercial licensing agreement(s) with rightsholder(s) or with their representatives?	☒ Yes ☐ No
If yes, specify the modality(ies) of the content covered by the datasets concerned:	☒ Text ☒ Image ☒ Video ☐ Audio ☐ Other *If so, please specify…*

2.2.2. Private datasets obtained from other third parties

Have you obtained private datasets from third parties that are not licensed as described in Section 2.2.1, such as data obtained from providers of private databases, or data intermediaries?	☒ Yes ☐ No
If yes, specify the modality(ies) of the content covered by the datasets concerned:	☒ Text ☒ Image ☒ Video ☐ Audio ☐ Other *If so, please specify…*
If publicly known, list private datasets obtained from other third parties:	N/A
General description of non-publicly known private datasets obtained from third parties	Moonshot AI collaborates with third parties to obtain multi-domain, multi-scenario data for model improvement, with all data access conducted in full compliance with relevant laws.

Additional comments (optional):	N/A

2.3 Data crawled and scraped from online sources

This Section requires information about crawled, scraped data, or otherwise compiled from online sources directly by the provider of the model or on their behalf (i.e. excluding publicly available datasets already compiled by third parties and made available on platforms such as common crawl that are covered under Section 2.1).

Were crawlers used by the provider or on behalf of?	☒ Yes ☐ No
If yes, specify crawler name(s)/identifier(s):	KimiBot (User-Agent identifier: KimiBot/1.0)
Purposes of the crawler(s):	KimiBot collects web content that may be used to train Kimi models.
General description of crawler behaviour:	KimiBot follows industry-standard robots.txt directives, allowing website administrators to block the crawler, restrict specific paths, or set a crawl delay.
Period of data collection:	01/01/2024-01/03/2026
Comprehensive description of the type of content and online sources crawled:	Online content crawled for the training of Kimi K3 includes publicly available web text, code, mathematical and knowledge-oriented materials. The crawled content comprises text, images and videos. Before training, it is subject to filtering and processing measures.
Type of modality covered:	☒ Text ☒ Image ☒ Video ☐ Audio ☐ Other *If so, please specify…*
Summary of the most relevant domain names crawled:	The crawled content includes a variety of publicly accessible online materials from various fields such as academic, educational, scientific, technical, research, and general-interest content.
Additional comments (optional):	N/A

2.4 User data

This Section requires information about user data collected by all services and products of the provider, including through mail services, social media platforms, content platforms or interaction with the providers' AI models and/or systems. This does not cover data licensed by users based on commercial transactional agreements described in Section 2.2.1., or customer data to fine-tune models for specific purposes.

Was data from user interactions with the AI model (e.g. user input and prompts) used to train the model?	☐ Yes ☒ No
Was data collected from user interactions with the provider's other services or products used to train the model?	☐ Yes ☒ No
If yes, provide a general description of the provider's services or products that were used to collect the user data:	N/A

Type of modality covered:	☐ Text ☐ Image ☐ Video ☐ Audio ☐ Other *If so, please specify…*
Additional comments (optional):	N/A

2.5 Synthetic data

This Section requires information about synthetic data created by or on behalf of the provider for training the model directly on the outputs of another AI model, in particular through model distillation or model alignment (e.g. AI feedback through reinforcement learning). This does not include the use of AI models to clean or enrich data (e.g. AI-generated metadata to enrich or modify a dataset, such as creating depth maps or text descriptions of images). In case this concerns publicly available datasets as described in Section 2.1, these should be reported in that Section of the Template. In case this concerns synthetic datasets created by third parties on behalf of the provider, these should be reported in this Section of the Template instead of in Section 2.2.2.

Was synthetic AI-generated data created by the provider or on their behalf to train the model?	☒ Yes ☐ No
If yes, modality of the synthetic data:	☒ Text ☒ Image ☒ Video ☐ Audio ☐ Other *If so, please specify…*
If yes, specify the general-purpose AI model(s) used to generate the synthetic data if available on the market:	Kimi K2, Kimi K2.5, and Kimi K2.6, were used to generate synthetic training data for Kimi K3, including domain-specific training trajectories.
Information about other AI models, including provider's own AI model(s) not available on the market, used to generate synthetic data to train the model to which this Summary applies:	Internally developed expert models were used to generate synthetic training environments and trajectories. The resulting training content covered instruction following, reasoning, coding, knowledge-intensive and agentic tasks.
Additional comments (optional):	N/A

2.6 Other sources of data

This Section requires information about data that does not fall under any of the categories in the previous Sections, for example data collected from offline sources, self-digitised media (e.g., digitised analog text context, images), datasets labelled by humans commissioned by the provider, or human generated data through reinforcement learning.

Have data sources other than those described in Sections 2.1 to 2.5 been used to train the model?	☒ Yes ☐ No
If yes, provide a narrative description of these data sources and the data:	An internal team of professionals undertakes the creation, labeling, and auditing of data across certain domains.
Additional comments (optional):	N/A

3. Data processing aspects

3.1. Respect of reservation of rights from text and data mining exception or limitation

This Section concerns measures implemented by the provider to identify and comply with the reservation of rights from the text and data mining (TDM) exception or limitation expressed pursuant to Article 4(3) of Directive (EU) 2019/790, as outlined in the copyright policy put in place by the provider in accordance with Article 53(1)(c) AI Act.

Are you a Signatory to the Code of Practice for general-purpose AI models that includes commitments to respect reservations of rights from the TDM exception or limitation?	☐ Yes ☒ No
Describe the measures implemented before model training to respect reservations of rights from the TDM exception or limitation before and during data collection, including the opt-out protocols and solutions honoured by the provider or, as applicable, by third parties from which datasets have been obtained:	Moonshot AI attaches great importance to complying with laws and regulations concerning intellectual property, trade secrets, and personal privacy. For publicly accessible web data, Kimi takes steps prior to and throughout the data collection process to honor applicable rights reservations and opt-out signals related to text and data mining; where applicable: (1) its crawler is built to follow robots.txt instructions and other standard web protocols; (2) it is not intended to bypass captchas, paywalls, or password-protected content; (3) it does not source content from websites identified by EU or Member State authorities as persistently infringing copyright on a commercial scale; (4) rightsholders may contact kimi at support@kimi.ai to request opt-out or lodge complaints, which are processed diligently within a reasonable time.
Additional comments (optional):	N/A

3.2 Removal of illegal content

This Section concerns measures taken to avoid or remove illegal content under Union law from the training data (such as blacklists, keywords, and model-based classifiers), without requiring disclosure of specific details about the provider's internal business practices or trade secrets. Such measures are advisable if the training data is likely to include illegal or unlawful content under Union law, in particular child sexual abuse material and terrorist content and the non-authorised use of material protected by intellectual property rights. Such measures do not include data selection practices, for example to increase the capability of the model.

General description of measures taken:	Moonshot AI carries out preprocessing and screening procedures on its training data to identify any proprietary

	information protected by intellectual property rights or trade secrets, as well as any content related to pornography, gambling, drugs, prejudice, discrimination, and the like, so as to prevent or eliminate illegal material from the training dataset. Such procedures may involve automated filtering, manual review, keyword-based checks, and classification model detection.

3.3. Other information (optional)

Other relevant information about data processing (optional):	N/A