

AI tools in education: what is already there, and where Dify fits



Appendix to the AI Agents workshop

LeX Consultancy B.V. · 20 September 2026

Almost everyone in the room already uses an AI tool: Copilot because the school has Microsoft, Gemini because the school has Google, ChatGPT because it is the best known. This appendix puts the tools side by side in plain words: what it is, where it already sits in school, what a school should check, and when to use it or not. At the end is where Dify fits in this landscape, because Dify is not yet another chatbot but a building layer on top of these tools.

Everything below is the state of play on **20 September 2026**. This market changes every quarter. For decisions about use at school you always look at the current advice of your national education bodies and data protection authority.

The most important sentence of this appendix: **the difference is not the model but the account**. The same chatbot with a free private account is something quite different from the same chatbot with a school account under a data processing agreement. In the first case your conversations can be used to train the model and you have no rights; in the second case there are agreements about that. So never ask “may we use ChatGPT?” but “under which account, with which data, for which task?”

Three layers

It helps to keep three layers apart. Most discussions at school are about layer 1; most of the gain is in layers 2 and 3.

Layer	What it is	Examples	Who decides
1. Chatbot	You type a question, you get an answer. Every time you explain again what you want.	ChatGPT, the Gemini app, Claude, Le Chat, Copilot Chat	The user themselves, often without the school knowing
2. Assistant in your office suite	AI in the tools you already have: summarise mail, draft a presentation, explain a table.	Microsoft 365 Copilot, Gemini in Workspace for Education	The school, through licences and settings
3. Building platform	You build an app with fixed steps, your own rules and your own documents, and share it with colleagues.	Dify, Copilot Studio, Gems (Google), GPTs and Projects (OpenAI), Projects (Claude)	The school, per app; this is what this workshop is about

A chatbot (layer 1) is handy for someone who can steer well themselves. An app (layer 3) takes over that steering: the rules are built in, so a colleague who never wrote a prompt still gets a good result. That is exactly why the guides build apps rather than hand out prompts.

The tools side by side

Microsoft Copilot

From	Microsoft (US), models from OpenAI and its own models
Where it already sits in school	In Microsoft 365: Copilot Chat with a school account is often already there; Microsoft 365 Copilot (in Word, Outlook, Teams) is a separate licence
State of play	Several European education bodies have carried out data protection impact assessments on Microsoft 365 Copilot; the general line in 2025-2026 is cautious use with school agreements in place first. Check the assessment of your own national body
Suited for	Working in your own documents and mail, summaries of Teams meetings, first drafts of texts
Watch out	Copilot Chat with a school account (data stays in the tenant) is something different from Copilot in the browser with a private account. Check which one you have. Copilot Studio is Microsoft's building platform; comparable to Dify, but tied to Microsoft

Google Gemini

From	Google (US)
Where it already sits in school	In Workspace for Education: Gemini in Docs, Gmail and Classroom, the Gemini app, NotebookLM, Gems
State of play	Assessments by education bodies are ongoing in several countries; Google meanwhile switches Gemini on in more and more components, and administrators can switch it off per component
Suited for	Working in Google documents, NotebookLM for learning from your own sources, Gems as a simple assistant of your own
Watch out	Wait with use by students until an assessment is done. Check what your administrator has switched on; defaults change with updates

ChatGPT (OpenAI)

From	OpenAI (US)
Where it already sits in school	Usually not through school; massively in private use by teachers and students
State of play	With a free or Plus account conversations can by default be used for training (can be switched off in the settings); with Team, Enterprise and Edu not, and a data processing agreement is then possible
Suited for	The broadest ecosystem: GPTs, Projects, voice, images. For personal use by the teacher without school data
Watch out	Never student data or unpublished test material in a private account. Through the API (as in Dify) different conditions apply than in the app: by default no training on API traffic

Claude (Anthropic)

From	Anthropic (US)
Where it already sits in school	Rarely through school; available as a model through Dify
State of play	Claude for Education and a learning mode exist, aimed mainly at higher education. Business accounts do not train on conversations
Suited for	Reading and rewriting long documents, careful language, Projects with your own sources. In Dify's Agent Console one of the compatible models
Watch out	Same rule as for ChatGPT: the account determines what is allowed. More expensive per answer than gpt-5-mini

Mistral (Le Chat)

From	Mistral AI (France)
Where it already sits in school	Rarely through school
State of play	European provider, falls under EU law and hosting in the EU is possible; some models are open (you can run them yourself)
Suited for	Anyone who wants a European provider or wants to self-host; strong in French and other European languages
Watch out	Smaller ecosystem; fewer ready-made education features. Here too: a free account is not the same as a business account

DeepSeek

From	DeepSeek (China)
Where it already sits in school	Not through school
State of play	The app and the website process data on servers in China; several European data protection authorities have warned about it. The open models can be hosted by European parties
Suited for	Only as an open model, hosted by a party in the EU, for those who want to experiment technically
Watch out	Do not use with school data through the app. The model is not the problem, the processing is

Grok (xAI)

From	xAI (US), intertwined with X
Where it already sits in school	Not through school
State of play	No meaningful use in education; repeated controversies about unfiltered output and the use of X posts for training
Suited for	Nothing in a school context
Watch out	Do not use. Do mention it in media literacy lessons: students meet it on X

Perplexity

From	Perplexity (US)
Where it already sits in school	Not through school; popular with students as a “search engine”
State of play	No education assessment. Searches the web and gives sources
Suited for	Exploratory searching with sources; check the sources yourself afterwards
Watch out	The sources are not always the best, and the summary is not always what the source says

NotebookLM

From	Google (US)
Where it already sits in school	Part of Workspace for Education (the administrator can switch it on or off)
State of play	Falls under the Google assessments of your education body; check whether your administrator has switched it on
Suited for	Learning from your own sources: summary, questions, an audio overview of a chapter
Watch out	Only with material you may upload. The same rules as for a knowledge base in Dify

European public models

From	National and European public initiatives (for example the Dutch GPT-NL, the EU's OpenEuroLLM and similar projects)
Where it already sits in school	Not yet in schools
State of play	Language models in development with attention to lawful training data and European values; meant for public organisations
Suited for	In time: a model whose data provenance is known
Watch out	Not yet a product you use today; do follow it

Four questions for every tool

Whether it is Copilot or an app you build yourself in Dify, the same four questions determine whether it is allowed and whether it is wise.

1. **Under which account?** Private, free, paid, or a school account with a data processing agreement. Only in the last case does the school have a say in what happens to the data.
2. **Which data goes in?** No personal data of students, parents or colleagues, unless the tool has been assessed for that. Unpublished test material and confidential documents count too.
3. **Is it trained on what you enter?** With consumer accounts often yes by default (sometimes you can switch it off), with business accounts and API use usually no. Check it; do not take a colleague's word for it.
4. **Who decides and who checks?** A tool one teacher uses for themselves is something different from a tool students use. For the second a school decision is needed, and someone who checks the output.

Aids that already exist:

- **The assessment framework of your national education body** for AI systems under the EU AI Act: most countries now have a decision tree and checklists for school and supplier. Use it at procurement and for every tool that touches students.
 - **Team agreements about generative AI:** many education bodies publish a toolkit to make agreements as a team. Do that before the tools, not after.
 - **The rule of thumb on age:** several education bodies advise that children under 13 do not work with generative AI independently.
 - **The impact assessments of your national bodies** on Microsoft and Google: the only independent assessments for the education sector. Other providers have no such assessment; that does not mean they are unsafe, but it does mean the school has to find out for itself.
-

GDPR and the AI Act in brief

Two laws determine what is allowed. The **GDPR** is about personal data, the **EU AI Act** about AI systems. They overlap, but ask different questions.

GDPR: who is responsible for the data?

- The school is the controller. Dify (cloud, outside the EU) and the model provider (OpenAI, Anthropic, Google) are processors. Without a data processing agreement no personal data may go in, and a sandbox account has no such agreement. That is why the guides work with fictional material.
- As soon as an app processes data of students, parents or colleagues (a name in a parent letter, a personal message to the Study coach that goes to the teacher), you need a legal basis, data minimisation, a retention period for the logs and, for high risk, a data protection impact assessment. With students that is almost always high risk.
- Transfer outside the EU needs a legal basis (for example the EU-US Data Privacy Framework for US providers that fall under it); check that per provider, it changes.

AI Act: what kind of system is it, and what do you have to do?

Part	Since	What it means for a school
Art. 4, AI literacy	2 February 2025	Whoever uses AI makes sure staff know enough about it. A workshop like this is such a measure; record who attended
Art. 5, prohibited practices	2 February 2025	Among others emotion recognition in education. Do not use a tool that measures the mood or attention of students
Annex III, high risk in education	planned August 2026; in 2025-2026 a delay was discussed, check the current state	Admission and placement, assessing learning outcomes or level, proctoring tests. Then heavy requirements apply to supplier and school: risk management, human oversight, documentation, registration
Art. 50, transparency	August 2026	A chatbot lets people know they are talking to AI; AI-generated content is made recognisable as such where relevant

The apps in this series (rewriting text, answering questions from a school guide, drafting messages, summarising a transcript) take no decisions about people and do not fall under high risk. The line is close: a Study coach that explains and gives hints is not high risk; the same coach that would record who does not understand the material and report that to the teacher as an assessment would be. That is a design choice, and the design card asks about it in box 4.

The EU competence frameworks: where building fits

The AI Act says that staff must be AI literate; two European frameworks say what that means.

- The **AI Literacy Framework for Primary and Secondary Education** (OECD and European Commission, final in 2026) describes four competence domains for students, each with knowledge, skills and attitudes: **Engage with AI** (understanding how it works and dealing with it responsibly), **Create with AI** (making something with it), **Manage AI** (using and monitoring it responsibly) and **Shape AI** (making choices and exerting influence).
- **DigCompEdu** is the European framework for the digital competences of teachers (six areas, from professional engagement to fostering students' digital competence); the Commission has added supplements for AI and data.

The guide series and the workshop touch all four domains, and that can be used in a professional development plan:

Domain	In the series
Engage with AI	Three layers, five patterns, sabotaging the agent, the knowledge base that splits wrongly
Create with AI	The building blocks: workflow, agent, chatflow, gatekeeper, agent console
Manage AI	Gatekeeper and self-critique, the human in the loop, Logs, the checklists for your own material, the four questions
Shape AI	The design card, the pitch, the choice of where the human belongs and what an app must never decide

Sources: ailiteracyframework.org and education.ec.europa.eu (AI Literacy Framework); joint-research-centre.ec.europa.eu (DigCompEdu).

Where Dify fits

Dify is not a model and not a chatbot. It is a building platform (layer 3) that calls the models of the providers above through their API. In the guides you chose a model per app: usually gpt-5-mini from OpenAI, in the Agent Console also Claude or Gemini. Three consequences:

1. You choose the model per task, and can switch. Rewriting a parent letter can be done with a small, cheap model. Analysing a long policy document asks for a bigger model. In Dify you change that in one menu, without the user of the app noticing. If a European model good enough arrives tomorrow, you put that underneath.

2. API traffic is not the same as the chatbot. What goes through the API of OpenAI, Anthropic or Google is by default not used to train models, and there is a business agreement. That is a better position than a free chatbot account. But the data still goes to a processor outside the school, usually outside the EU. The rules from the four questions still apply.

3. Dify cloud itself is a processor too. The sandbox account from the workshop runs on Dify's servers, outside the EU. For practising with fictional material that is fine. For real use with school data there are two routes: host Dify yourself (it is open source) or run it in an environment of your school or trust with central model keys. Guide 5 (Transcription tool) shows how to move a prototype out of the cloud.

When to use what

Situation	Take
A quick question, a first draft, something for yourself	The chatbot your school has an account for (Copilot Chat or Gemini with a school account)
A task that comes back every week and has to follow rules (house style, format, sources)	An app in Dify, or in the building platform of your suite (Copilot Studio, Gems)
Working in your own documents and mail	Microsoft 365 Copilot or Gemini in Workspace, once your school leadership has taken a decision based on the assessments
Something for students	First the school decision, the assessment framework and the age limit; only then the tool
Experimenting with what is possible	Dify with fictional material, as in this workshop

To remember

- The account determines more than the brand.
- GDPR: no personal data without a data processing agreement. AI Act: literacy is mandatory, assessing students is high risk, a chatbot says it is AI.
- Copilot and Gemini are already in school; whether they may be switched on is a leadership decision based on the assessments of your education bodies, not the choice of an individual teacher.
- ChatGPT, Claude and Mistral are fine tools for the teacher themselves, without school data, and as the model under an app.
- The DeepSeek app and Grok stay out of school.
- You put Dify on top of these models: you build the rules in, choose the model per task and can switch when the market changes. For real use the building layer belongs with the school or trust, not in a free cloud account.

Licence: CC BY-SA 4.0 (<https://creativecommons.org/licenses/by-sa/4.0/>). You may copy, adapt and reuse this document, including at your own school, provided you credit the source (Guido van Dijk, LeX Consultancy B.V., <https://git.lexconsultancy.nl/guido/dify-guides>) and share your adaptation under the same licence. The LeX Consultancy logo is excluded; product names are trademarks of their owners. Schools, people and documents in the examples are fictional.

Owner and contact: Guido van Dijk, LeX Consultancy B.V. · guido@l-e-x.nl

