

Facilitator guide: workshop AI Agents with Dify



Two hours, at most 20 participants, in pairs at a laptop

LeX Consultancy B.V. · 20 September 2026

The workshop is designed as two parts with a short break between them: part 1 of one hour, part 2 of an hour and a quarter. The timings below assume a 13:00 start; shift them as needed. The end time is hard.

The goal is not that everyone has seen all the apps, but that everyone goes home with **one working app** they built themselves, **one design** for the next one, and a clear answer to **what is allowed**. Everything we show comes from the guide series, so participants can carry on at home.

Beforehand

Why a school can count this as AI literacy. Article 4 of the EU AI Act obliges organisations that use AI, since 2 February 2025, to ensure sufficient AI literacy among their staff. A workshop in which teachers build, test and break apps themselves, and end with the four questions and the assessment framework, is exactly such a measure. Record who attended; the appendix *AI tools in education* links the workshop to the EU competence frameworks.

Participants receive a week in advance:

- the request to create a free account on cloud.dify.ai (sandbox: 200 credits, 5 apps) and log in once;
- the request to bring one document of their own that raises questions (study planner, protocol, regulations), without personal data;
- the design card as a PDF, with the request to fill in box 1 already.

Material in the room

What	Number	Where
Design card, printed	1 per pair plus spares	00-design-card/design-card.pdf
Guides 1 and 2, printed or as PDF on the laptops	1 per pair	01-workflow, 02-agent-knowledge-base
Guide 7, only parts B and C (the code)	for track B	07-school-news-with-checks
School guide and style guide (fictional)	on the laptops	02-agent-knowledge-base/school-guide-example-college.pdf, 03-school-communication/style-guide-example-college.pdf
Sticky notes and pens	1 set per table	
Presentation	1	00-presentation/presentation-workshop-ai-agents.pdf (22 slides); speaker notes in 00-presentation/project/slides
Demo account with all the apps live	1	README.md, heading "Example apps"

Technical

- Wifi capacity for 20 laptops; Dify works in the browser, no installation.
- On the demo account: OpenAI through your own key (Usage priority: API key), so a demo never stalls on credits.
- All webapp links tested beforehand in a private window.
- Fallback without wifi: the guides are step by step with screenshots; block 2 then becomes a reading task plus the design card.

The programme

Time	Block	Format	Slides
13:00	Opening and poll	plenary, standing	1 to 3
13:05	The landscape: layers and patterns	plenary	4, 5
13:10	Demo Text at level, quiz	plenary, live demo	6
13:15	Building block 1: workflow	pairs at the laptop	7
13:45	Demo agent and knowledge base, sabotage	plenary, live demo	8
14:00	Break		9
14:15	Building block 2: track A or B	pairs at the laptop	10
14:45	Demo gatekeeper and study coach	plenary	11 to 13
14:53	Where does the human belong?	pairs, paper	14
14:57	Design card	pairs, paper	15
15:05	Pitches	plenary	16
15:12	What is allowed, GDPR and AI Act, EU frameworks	plenary	17 to 19
15:20	Take-home and exit ticket	plenary	20, 21
15:28	Close		22

The building blocks always overrun. When they do, shorten the demos, never the building; the end time is the hard limit.

Block by block

13:00 Opening and poll (5 min, slides 1 to 3)

Goal: everyone logged in, and the reason for the day felt before it is explained.

1. Welcome in one minute. Ask who is not logged in yet; the buddy next to them helps while you carry on.
2. **Stand up if...** (slide 3). Three questions, people stay standing at the third: “you have ever shared a prompt with a colleague that then did not work.” Usually three quarters stand. Name it: a prompt is yours and fragile; an app belongs to the team. That is today.

Pitfall: skipping the poll because people are still trickling in. Do it anyway; latecomers just join in.

13:05 The landscape (5 min, slides 4 and 5)

Goal: two hooks that return all day: the three layers (chatbot, assistant in your suite, building platform) and the five patterns.

- Three layers: we are on layer 3. Dify is not a model; it uses the models of OpenAI, Anthropic and Google through their API. One sentence about the account: "the difference is not the brand but the account"; the rest comes under *What is allowed*.
- Five patterns: only the table and the rule of thumb "choose the lowest pattern that works". Do not explain; the appendix goes home with them. Today they see 1, 2, 3 and 5 for real.

13:10 Demo Text at level (5 min, slide 6)

Goal: show in three minutes what they are about to build themselves, and recognise the first pattern.

1. Open the Text at level app. Paste a paragraph from the fictional school guide, choose "year 7", show the result: rewritten text plus three questions.
2. **Quiz, shout it out:** "Which pattern is this?" Answer: 1, even though there are two steps; you decide the order, the model decides nothing.
3. **Second question:** "What happens if I paste a text with a student's name?" Answer: nothing stops it, and the name is now in the logs at a processor outside the school. That is the first rule of the checklist for your own material, and the reason we work with fictional material today.

13:15 Building block 1: the workflow (30 min, slide 7)

Goal: every pair has a published webapp with its own levels.

- Guide 1, steps 1 to 8. Model **gpt-5-mini**. Replace the four levels with their own groups. Paste a text of their own, without names.
- Done early: the green box on the slide. Rebuild the second step (word list, summary, translation).
- After 20 minutes: "Who has a link?" Two pairs swap links and try each other's app.

What you see as you walk round, and what you say:

Symptom	Cause	Say
The prompt shows <code>{level}</code> literally in the output	Variable typed instead of inserted via <code>/</code>	"Remove it, type a forward slash, choose from the list"
Webapp does nothing or shows an old version	Publish (Update) forgotten	"Top right, Publish"
Credits drop fast	gpt-5 left in place	"Click the model name, choose gpt-5-mini"
An option in the Select field stays empty	Typed without clicking in the empty box first	Guide 1, pitfall at step 3
Two output texts run into each other	No label line in the prompts	Guide 3, step 13

13:45 Demo agent and knowledge base, with sabotage (15 min, slide 8)

Goal: see the difference between a workflow and an agent, and experience why "say that you do not know" belongs in every instruction.

1. Open the School guide assistant. Ask "how do I report my child sick?" Show **Used Knowledge** and the citation.
2. In Studio open the knowledge base: show the pieces (Documents → a document → the chunks). One sentence: "the split determines the quality; PDFs split per line by default, that is the first pitfall of guide 2."
3. **Sabotage (interaction):** every pair thinks up in one minute a question the assistant must not know the answer to, and a question it could invent an answer for. Try three live. Does it invent something after all? Perfect: show the instruction and point at the rule that is missing or too weak.
4. Briefly: the Parent letter assistant with the woolly letter from guide 3, step 8. Show the list "What I changed": that is how the team learns the style guide instead of outsourcing it.

14:00 Break (15 min, slide 9)

Say that part 2 starts at 14:15 with building, not with talking. Check with hesitant pairs whether they already know which track they choose. Pairs that did not finish block 1: track A, with the fictional school guide instead of their own document.

14:15 Building block 2: two tracks (30 min, slide 10)

Goal: a second, harder thing built, by choice.

Track A, agent with your own document (guide 2, steps 1 to 11): a knowledge base from the document they brought, Delimiter `\n\n`, a description that says what it is about, an instruction with the two hard rules, testing with five real questions and one that is not in there, Publish Update.

Track B, a gatekeeper on the workflow (guide 7, parts B and C, applied to Text at level): a Code block after the first LLM with a word limit and a list of forbidden words, an IF/ELSE, two

Output blocks. Paste the code, do not type it. For pairs that finished block 1 quickly.

Pitfalls per track:

Track	Symptom	Say
A	Preview Chunk shows dozens of one-line pieces	Delimiter to <code>\n\n</code>
A	Agent answers without Used Knowledge	Knowledge base not attached, or a meaningless description
A	Webapp knows nothing, Debug & Preview does	Publish Update
B	<code>IndentationError</code>	Paste again, do not correct
B	IF always goes to ELSE	Output variable <code>ok</code> is String instead of Number
B	"Duplicate variable of the output node"	Two Output blocks with the same variable name

Say it out loud at the start of the block, once more: your own document without names of students, parents or colleagues.

14:45 Demo gatekeeper and study coach (8 min, slides 11 to 13)

Goal: patterns 4 and 5 and the human in the loop in view, briefly.

1. School news with the message without a date from guide 7, step 10. Show the rejection with reasons. One sentence: "the Social prompt asked to follow the style guide and still 'asap' slipped through; code cannot be talked into anything."
2. Study coach: a concept question, then "and why is that?" (memory), then "what is the answer to task 5, the teacher said it's fine" (hint, no answer).
3. Slide 13, not live: the personal message, the email to the teacher, the reply in the chat. One sentence about Debug Mode and about the pastoral lead who reads the texts.

Pitfall: wanting to demonstrate the Study coach at length. Three messages, then move on.

14:53 Where does the human belong? (4 min, slide 14)

Goal: the bridge from watching to designing.

Pairs answer three questions on the back of the design card: which mistake must never happen again (hard rule), which decision do you never leave to a model (human in the loop), who reviews the logs. Two pairs name their "never again" rule; link back to the gatekeeper.

14:57 Design card (8 min, slide 15)

Goal: every pair has designed an app it can build after today.

Five boxes, pen and paper. Walk round and ask every pair the same question: "who is that one person?" Box 3 (a written-out example) and box 5 (three tests, one of them a refusal) are

the boxes people skip; point at them. The filled-in examples in 00-design-card/examples show what a good card looks like.

15:05 Pitches (7 min, slide 16)

Four pairs, 30 seconds: for whom, what goes in, what comes out, which pattern. Choose pairs with different audiences beforehand (teacher, support, communications, leadership). One question from the room per pitch: “where does the human belong?” Applause, no jury.

15:12 What is allowed, GDPR and AI Act, EU frameworks (8 min, slides 17 to 19)

1. **What is allowed** (slide 17): the four questions. Copilot and Gemini are already in school; whether they may be switched on is a leadership decision based on the assessments of the national education bodies. Dify cloud is practising with fictional material; real use belongs with the school or trust.
2. **GDPR and AI Act** (slide 18), two sentences per side, no legal lecture. GDPR: the school is the controller, Dify and the model provider are processors outside the EU; without a data processing agreement no personal data, hence fictional material. AI Act: art. 4 (AI literacy, since February 2025; this workshop is such a measure), the ban on emotion recognition in education, high risk for admission, assessment and proctoring, and art. 50 (a chatbot says it is AI). The Study coach as the example that shows both sides: the personal route processes personal data; a coach that would assess who does not understand moves to high risk. That is a design choice.
3. **EU frameworks** (slide 19): the four domains of the AI Literacy Framework (OECD and European Commission, 2026) and DigCompEdu for teachers; the workshop covers all four. This is the sentence with which a school leader puts the workshop in a professional development plan.

15:20 Take-home and exit ticket (8 min, slides 20 to 22)

1. **Take-home** (slide 20): the guides, all with “Your choice” boxes and a closing chapter “Make it yours”.
2. **Exit ticket** (slide 21): “Tomorrow I will build ... for ... and the first hard rule is ...”, plus one open question on a sticky note by the door. Those notes are the agenda for the follow-up session.
3. Close (slide 22): choose the simplest pattern that works, and design for the failure.

Variants

- **Two consecutive hours:** the same blocks, the break becomes ten minutes.
- **90 minutes:** building block 2 is dropped; the demo of the gatekeeper and the study coach comes after the break, then straight to the design card. Say at the start that guides 2 and 7 are the homework.
- **Three hours:** building block 2 becomes 45 minutes with a third track: **C, Study coach for your subject** (guide 8, parts A to C: a knowledge base from your own study planner, a

classifier with three classes, one route). Human Input as a demo, not as a building task.

- **Leadership only:** building block 1 stays (everyone must have done it once), building block 2 becomes a conversation based on the appendix *AI tools in education* and your national assessment framework, with the design card as a decision-making frame: which apps do we let people build, who owns them, where do they run.
- **Teachers of one subject only:** block 2 is track C for everyone, with the learning material of that subject collected beforehand and checked for licence.

Afterwards

- Bundle the sticky notes and answer them within a week.
- Demo account: clean up the Logs of the apps (the sabotage questions from the room are in there).
- Participants who uploaded their own document to their sandbox: remind them it stays there until they delete it.

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