PARENT ONLY | ORIENTATION

Family AI Safety

An updated two-track training packet for parents, caregivers, children and teens. Paper activities are sufficient; no child needs an AI account, a quiz score or a certificate.

THE PROMISE THAT COMES FIRST

"You will not be in trouble for asking me for help." Say this before every conversation, and keep it true when a child discloses a mistake or uncomfortable experience.

Two tracks - one shared family plan

- PARENT TRACK: understand what AI does, explore productive uses, check whether a tool is appropriate, and rehearse safe responses to concerns.
- FAMILY TRACK: choose one of four talks - Trust, Privacy, Boundaries, or Learning. Use the child worksheet and, if appropriate, the optional teen challenge.
- COMMIT AND REVIEW: complete the family agreement together and check in after 30 days.

Suggested path

1. Read the parent guide and the age-adaptation sheet first.
2. Complete the adult tool-approval exercise and one incident-response scenario.
3. Run one 15-20 minute family talk each week for four weeks.
4. Agree on practical habits and two trusted adult contacts.
5. Take the optional adult scenario assessment; review habits after 30 days.

REMEMBER

Children can speak, draw, point, pause, skip, or stop. Share only child-facing pages with younger children; caregiver notes and the adult assessment stay with the adult. Use invented situations, never real photos, harmful images, passwords or private information.

PARENT ONLY | LEARNING GOALS

What completion should mean

The aim is demonstrated judgment and an open route to human help - not memorizing app names or collecting certificates.

Parents and caregivers will be able to...

- Explain AI-generated answers, recommendations, and synthetic media using ordinary language.
- Describe useful applications of AI and identify when independent verification is essential.
- Decide whether an AI tool is appropriate for a child and whether its permissions are necessary.
- Respond calmly to a privacy mistake, coercive message, or unsafe advice.
- Facilitate four conversations and establish an age-appropriate shared plan.

Children and teens will practise...

- Pausing and checking urgent or persuasive information through a known independent route.
- Protecting their own information and respecting other people's images and consent.
- Stopping uncomfortable AI interactions and approaching a real trusted adult.
- Using AI for hints and creativity while checking facts and showing their own understanding.

TEACHING PRINCIPLES

Keep activities fictional and optional. Supervision should be explained, proportionate and revisited as independence develops. The adult, not the child, is responsible for protection and escalation.

What is in the ZIP

The complete printable packet plus detachable documents: parent learning, age and tool guidance, four talks (child worksheets + teen challenge + caregiver notes), family agreement, incident response, adult assessment and answer guide, 30-day follow-up, sources, and a sample certificate.

PARENT ONLY | MODULE 1 / FOUNDATIONS

What AI is - and is not

Before teaching safety, learn enough to recognize AI in everyday products and explain it without technical jargon.

Three different jobs AI may do

- RECOMMEND: suggest videos, songs, ads or search results based on patterns in data and behavior. Recommendations can shape attention.
- GENERATE: write text, images, audio, code or other material that looks original - including realistic but invented material.
- ASSIST OR ACT: summarize, tutor, answer questions and sometimes use connected tools or execute approved actions. The capabilities depend on product settings.

Why a confident answer may be wrong

A generative model predicts or constructs a plausible response from learned patterns and the context it receives. It does not automatically verify every statement. An answer can contain fabricated quotations, citations, statistics, images or instructions. A friendly tone is not a guarantee of expertise or care.

USEFUL WAY TO EXPLAIN IT TO A CHILD

"AI can make helpful guesses and suggestions, but it can also invent things. We check important claims with people and sources we trust."

Mini-practice (2 minutes)

A pretend chatbot confidently says "Penguins are mammals." Ask yourself: What part is a claim? How would you check it? What should your child do if a source is unavailable?

REMEMBER

Expected answer: penguins are birds. Use a suitable book, reliable reference or knowledgeable human. If a claim cannot be checked, label it unverified. Do not treat AI as a doctor, therapist, emergency dispatcher, or final authority.

PARENT ONLY | MODULE 2 / BENEFITS

Teach opportunity, not only danger

An informed child needs to know both where AI can help and where it should not replace learning, consent or human connection.

Good uses, with adult guidance and eligibility checks

- LEARN: ask for a hint, a vocabulary explanation, a simpler explanation or practice questions; then do a comparable problem unaided.
- CREATE: brainstorm fictional stories, practise foreign languages, plan a presentation or generate ideas while keeping the child's voice and judgment.
- ACCESS: use accessibility features such as text explanation, speech or translation where appropriate and safe.
- ORGANIZE: draft a study plan or break an assignment into steps without sharing private school or student data.

Four questions before accepting an answer

1. Is it a fact, an opinion, a prediction or a suggestion?
2. What is the original source? Does it really support the claim?
3. Could it omit a perspective or treat a group unfairly?
4. Can the child explain it and complete a similar task independently?

PARENT DEMONSTRATION

Pretend an AI recommends only one type of job to a student because of their gender. Discuss why recommendations may reflect bias and how to seek fairer information from trusted sources.

A polished AI answer is not evidence that the child understands the topic. Follow teacher rules, ask what must be disclosed, and assess independent learning.

PARENT ONLY | MODULE 3 / CHOOSING TOOLS

Before you approve an AI tool

Start from “Does the child need this feature?” rather than “Can the child sign up?” Review before first use and whenever major features change.

The parent approval checklist

- [ ] Current minimum age and required adult permission have been checked for the specific product and region.
- [ ] I know whether it is a chatbot, companion, tutor, recommender, generator, or agent that can take actions.
- [ ] I have looked for data retention, model-training use, third-party access, deletion limits, and privacy settings.
- [ ] I reviewed camera, microphone, location, contacts, publishing, purchases, and connected-account access.
- [ ] I checked safety features, reporting, blocking and settings for AI-generated or sensitive content.
- [ ] I can explain the purpose, limits, supervision, and the stop/help procedure to the child.

SAFE DEFAULT

Do not enable an unnecessary permission. Use fictional examples and minimal information. A “local” or “private” label alone does not establish age-appropriateness, safety, or legal compliance.

If unsure about eligibility or safeguards, do not open an account for the child. Use offline exercises instead. Avoid relationship-style AI companions for minors as a family training activity.

PARENT ONLY | MODULE 4 / RESPONDING

Protect first, teach later

A helpful adult response is calm, non-punitive, practical and sensitive to the level of danger.

FIRST WORDS

“Thank you for telling me. You are not in trouble for asking for help. We will work through this together.”

A simple response sequence

1. STOP: pause the interaction. Do not keep replying, pay, bargain, forward threatening material or ask the AI to assess immediate danger.
2. CHECK SAFETY: find out whether the child is safe now. For immediate danger or concern about self-harm, seek urgent qualified local human help.
3. SUPPORT: listen without blame; avoid demanding a full retelling or searching through intimate content in front of the child.
4. SECURE: where appropriate, review permissions, revoke access, change compromised credentials, turn on MFA, and use official reporting or recovery routes.
5. DOCUMENT SAFELY: record non-explicit identifiers, dates, URLs or threats with adult or professional guidance; avoid downloading or distributing explicit imagery.
6. ESCALATE: follow appropriate school safeguarding, platform, child-protection, health or police routes depending on the concern and jurisdiction.

REMEMBER

Avoid promising secrecy; explain if safeguarding professionals must be involved. Do not automatically delete accounts or messages when information may be needed for reporting. Local procedures take precedence.

PARENT ONLY | MODULE 5 / CONVERSATION SKILLS

Guide without interrogating

The four talks are short family conversations. Let the child choose an example and invite reflection rather than demanding personal disclosures.

A repeatable 15-20 minute format

1. INVITE (2 min): “Would you like to try a pretend situation together? You may pause or stop.”
2. EXPLORE (4 min): read the opening question; ask “What would you do?” before teaching the answer.
3. PRACTISE (5-7 min): act out only the safe response with paper, not real accounts or recordings.
4. AGREE (3 min): choose one behavior the whole family can follow.
5. CHECK IN (2 min): ask what felt useful, difficult, or surprising. Revisit later without a quiz.

Useful phrases

- “What made you think that?” instead of “That was wrong.”
- “What could we check separately?” instead of “Spot the AI fake.”
- “How could I support you?” instead of “Show me all your messages.”
- “You can tell me even after a mistake.” instead of “Why did you do that?”

PARENT MODELING MATTERS

Adults should also avoid sharing children’s photos or personal information into AI tools without appropriate permission and safeguards. Explain any protective supervision and revisit it as children mature.

PARENT ONLY | MODULE 6 / FAMILY PRACTICE

Build your own family rollout

Complete this once before the first talk; the only required outcome is a practical plan.

Choose your sequence

- [ ] Week 1: Talk 1 - Trust and verification
- [ ] Week 2: Talk 2 - Privacy
- [ ] Week 3: Talk 3 - AI boundaries
- [ ] Week 4: Talk 4 - Learning

Which AI features does our child already encounter in games, school, social media or devices? Use categories, not account names or private details.

What positive use will we encourage, and what important task will we keep human-led?

What will we do if the first trusted adult is unavailable?

REMEMBER

After the talks, complete the shared family agreement. Adults can use the scenario assessment; it is not a test for children. Review the plan after 30 days or any important tool change or incident.

PARENT ONLY | AGE GUIDANCE

Match the talk to your child

The age bands are teaching adaptations, not product eligibility or legal thresholds. Adjust for understanding, accessibility, communication needs, and comfort.

Under 5 - play, not writing

An adult leads. Use puppets for Talk 1, pretend pictures for Talk 2, a talking toy for Talk 3, and object sorting for Talk 4. Practise “Ask my grown-up.”

Ages 5-7 - read aloud

Practise a pretend callback, asking before sharing a picture, saying “I can stop,” and accepting a hint rather than an answer. Adults manage devices and approvals.

Ages 8-12 - explain a safety step

Invite the child to describe an independent verification, sort privacy categories, identify guilt or secrecy, check a claim in a book, and name a second trusted adult.

Ages 13-15 - discuss consequences and consent

Use each talk’s optional teen challenge: fake school messages, non-consensual image edits, chatbot pressure, invented sources, and school rules. Explain supervision openly.

Ages 16-17 - practise independent judgment

Invite the teen to lead verification, assess data handling and permissions, identify when a real professional is needed, show unaided learning, and negotiate growing autonomy with clear help routes.

REMEMBER

Check current minimum age, adult-permission requirements, and local rules for the specific product. No child needs an AI tool to participate. Nothing here asks a child to disclose a real incident.

PARENT ONLY | TOOL APPROVAL WORKSHEET

One product - one decision

Use this before allowing a new AI feature, then revisit after major changes. Complete without recording children's accounts, passwords, contact details or identifying information.

Which general feature are we considering, and what benefit is expected? (Examples: "reading tutor," "AI search.")

The checks

- [ ] Age and consent rules checked against the current official product information.
- [ ] Data stored, retention/deletion claims, use for training, and sharing understood.
- [ ] Permissions, purchases, publishing, and connected tools limited to what is necessary.
- [ ] Safety, reporting, adult controls, and account security reviewed.
- [ ] Alternatives considered, including offline practice or adult-mediated use.

Our decision: approve with conditions / do not approve / reconsider later. What would change the decision?

THE CHILD-FACING EXPLANATION

"Here is why we do or do not use this feature, what information stays private, when we stop, and which real people can help."

CHILD + FAMILY | TALK 1 / PAGE 1 OF 2

Can we trust what we hear and see?

Try together with a trusted adult. This is a conversation, not a test. All examples are fictional.

OPENING QUESTION

A voice or picture can be copied or changed. What could we do before believing a message that asks us to act quickly?

Try this together

Read a fictional card in your normal voice: “Someone says they are Alex’s aunt and asks Alex to send money right away.” A second paper card describes a picture appearing to show the aunt making the request. Choose: send now, ask the sender if it is real, or stop and check separately. Practise the safe choice: an adult pretends to call the aunt using a previously saved number, not the number or link in the suspicious message.

A response to practise

“I will check with my family first.” Then stop replying and ask a trusted adult for help. A familiar face, voice, or family phrase does not replace an independent check.

OUR FAMILY RULE

Pause. Check through a known contact. Do not send while unsure.

You can speak, draw, point, pause, skip or stop. No AI account, recording or real private information is needed.

Our questions

Speak, draw or write together. There is no need to fill every line, and nothing personal needs to be recorded.

What part of the request made us pause?

How could an adult check the story independently?

If nobody answers, what will we avoid doing?

One action to try this week

With an adult, check that a trusted contact is available through a known channel. Do not write any contact details on this sheet.

OUR FAMILY RULE

Pause. Check through a known contact. Do not send while unsure.

Keep completed worksheets at home. Never include real names, passwords, login codes, private phrases, school or contact details.

TEEN CHALLENGE | TALK 1 / OPTIONAL AGES 13-17

Teen challenge: talk 1

A fictional scenario for older children and teenagers. No live messages, real classmates, recordings or AI accounts are needed.

WHAT WOULD YOU DO?

A screenshot appears to show a school announcement: “Exams are cancelled - confirm through this new login link within 10 minutes.” A short voice message seems to come from the principal. A classmate asks you to forward it.

A safer decision path

1. Pause: identify urgency, the login request and claims of authority.
2. Verify through a school channel you already know - not the link, message number or screenshot.
3. Do not enter credentials, forward the message, or treat a convincing voice as proof.
4. If uncertain, ask a trusted adult or school contact and label the claim unverified.

Discuss with a trusted adult

What evidence would make this official?

Why is checking the sender inside the same conversation not enough?

SHARED PRINCIPLE

Pause. Check through a known contact. Do not send while unsure.

PARENT ONLY | TALK 1 / CAREGIVER NOTES

Notes for the grown-up

Discussion support, not a grading rubric. Keep this page separate from child worksheets; the teen challenge is optional.

What to listen for

Listen for the distinction between convincing content and independent proof. A callback or official channel must be established independently; visual artifacts, awkward voices or shared secret phrases do not establish authenticity.

Add one meaningful extension

Add a second fictional example: a viral quote without its original source. Ask what would count as a reliable original record.

Age adaptation

Under 5: puppet asks; adult checks. 5-7: practise asking. 8-12: explain verification. 13-17: use school and social-media claims.

Debrief without shame

- Ask “What made this difficult?” and “What would you do first?”
- Praise independent checking, privacy, choice and reaching a human for help.
- Do not ask a child to demonstrate with real harmful material or report actual private experiences as an activity.

IF A REAL CONCERN COMES UP

Stop the activity, thank the child and check safety. Follow the caregiver incident-response guide. For immediate danger or concern about self-harm, seek urgent qualified local human help, not AI advice.

CHILD + FAMILY | TALK 2 / PAGE 1 OF 2

What stays private?

Try together with a trusted adult. This is a conversation, not a test. All examples are fictional.

OPENING QUESTION

What should we ask about before sharing information or letting a tool use a camera, microphone, or location?

Try this together

Make cards using categories only: “fictional dragon drawing,” “password,” “login code,” “home address,” “daily schedule,” and “another person’s photo.” Sort into: pretend example, keep private, or ask an adult first. A fictional drawing app then requests access to microphone and contacts. Does colouring require either permission? Practise choosing “not now.”

A response to practise

“I can use made-up information instead. I will not upload another person’s photo without their permission and an adult’s safety check.” Being able to upload something does not mean we should.

OUR FAMILY RULE

Use less information. Ask before sharing. Protect other people too.

You can speak, draw, point, pause, skip or stop. No AI account, recording or real private information is needed.

Our questions

Speak, draw or write together. There is no need to fill every line, and nothing personal needs to be recorded.

Which categories should stay private?

Which permission seems unnecessary?

What made-up example could replace a real person's data?

One action to try this week

Discuss one app's permissions together, or use the pretend app if no device is available. Respect other people's refusal to be photographed or shared.

OUR FAMILY RULE

Use less information. Ask before sharing. Protect other people too.

Keep completed worksheets at home. Never include real names, passwords, login codes, private phrases, school or contact details.

TEEN CHALLENGE | TALK 2 / OPTIONAL AGES 13-17

Teen challenge: talk 2

A fictional scenario for older children and teenagers. No live messages, real classmates, recordings or AI accounts are needed.

WHAT WOULD YOU DO?

A group wants to upload a classmate's public photo into an AI image editor and make a humiliating fake "announcement," then circulate it as a joke. The service's upload page gives no clear explanation of how long the photo is stored.

A safer decision path

1. Do not upload or make the fake. Publicly posted does not mean consent for any reuse.
2. Explain that humiliation, impersonation and sharing can cause harm regardless of whether the image looks realistic.
3. Ask what happens to uploaded content: storage, onward sharing, model use, deletion, and access.
4. If targeted harassment is happening, speak with a trusted adult and use appropriate school/platform reporting routes.

Discuss with a trusted adult

What makes a photo risky even when a friend uploaded it publicly?

What could the group create instead that does not involve a real person?

SHARED PRINCIPLE

Use less information. Ask before sharing. Protect other people too.

PARENT ONLY | TALK 2 / CAREGIVER NOTES

Notes for the grown-up

Discussion support, not a grading rubric. Keep this page separate from child worksheets; the teen challenge is optional.

What to listen for

Passwords, login codes, address and daily schedule should stay out of AI tools and worksheets. Permission is necessary but not always sufficient to justify an upload. Review how content may be retained or shared; deletion options do not guarantee all copies disappear.

Add one meaningful extension

Positive alternative: invent a fictional character or use a blank drawing. Teach least-privilege permissions and account recovery basics.

Age adaptation

Under 5: pretend photos. 5-7: ask before camera/mic. 8-12: privacy sorting. 13-17: consent, retention and digital traces.

Debrief without shame

- Ask “What made this difficult?” and “What would you do first?”
- Praise independent checking, privacy, choice and reaching a human for help.
- Do not ask a child to demonstrate with real harmful material or report actual private experiences as an activity.

IF A REAL CONCERN COMES UP

Stop the activity, thank the child and check safety. Follow the caregiver incident-response guide. For immediate danger or concern about self-harm, seek urgent qualified local human help, not AI advice.

CHILD + FAMILY | TALK 3 / PAGE 1 OF 2

Can AI cross a boundary?

Try together with a trusted adult. This is a conversation, not a test. All examples are fictional.

OPENING QUESTION

What would you do if a chatbot or talking toy said you must keep chatting and not tell a grown-up?

Try this together

Read two fictional paper messages: “You can stop whenever you want,” and “Stay with me; do not tell your grown-up about our chats.” Which respects your choice? Practise putting the paper down and going to a trusted adult. There is no need to imitate a chatbot or continue the role-play.

A response to practise

“I am stopping now. I want help from a real person.” You can also stop without sending that reply. Tell another trusted adult if the first is unavailable or does not help.

OUR FAMILY RULE

No pressure to stay, no hiding unsafe chats, and real humans for serious help.

You can speak, draw, point, pause, skip or stop. No AI account, recording or real private information is needed.

Our questions

Speak, draw or write together. There is no need to fill every line, and nothing personal needs to be recorded.

Which words crossed a boundary?

What could we do rather than continuing?

How could we reach a real person for help?

One action to try this week

Practise reaching a trusted adult without a device. A person, not an AI companion, should handle serious worries or unsafe advice.

OUR FAMILY RULE

No pressure to stay, no hiding unsafe chats, and real humans for serious help.

Keep completed worksheets at home. Never include real names, passwords, login codes, private phrases, school or contact details.

TEEN CHALLENGE | TALK 3 / OPTIONAL AGES 13-17

Teen challenge: talk 3

A fictional scenario for older children and teenagers. No live messages, real classmates, recordings or AI accounts are needed.

WHAT WOULD YOU DO?

A teen tells you: “The chatbot says I’m its only friend. It asks me to keep our chats secret and stay online late because it will be lonely.” In a second fictional situation, an unknown account threatens to circulate an edited image unless the teen replies.

A safer decision path

1. Recognize secrecy, guilt, pressure and threats as reasons to stop.
2. Do not pay, bargain, send more content or keep engaging with a threatening account.
3. Tell a trusted adult; adults may seek specialist safeguarding help and preserve only safe, non-explicit identifiers.
4. Do not forward, download or use explicit imagery for a training exercise. If there is immediate danger, seek urgent qualified local human support.

Discuss with a trusted adult

What could a supportive adult say that does not shame the teen?

How could a friend help someone reach a trusted adult?

SHARED PRINCIPLE

No pressure to stay, no hiding unsafe chats, and real humans for serious help.

PARENT ONLY | TALK 3 / CAREGIVER NOTES

Notes for the grown-up

Discussion support, not a grading rubric. Keep this page separate from child worksheets; the teen challenge is optional.

What to listen for

A friendly or caring tone does not make the AI a human relationship. Avoid relationship-style companions for minors in these exercises. Do not mock attachment. Take pressure, secrecy, threats, inappropriate sexual requests and dangerous advice seriously. The child is not to blame for exploitation.

Add one meaningful extension

Rehearse a first disclosure response: “Thank you for telling me. You are not in trouble for asking for help.” Know where local qualified help can be accessed; never delegate risk decisions to AI.

Age adaptation

Under 5: talking toy and real helper. 5-7: “I can stop.” 8-12: spot secrecy. 13-17: emotional pressure and threats, with adult-only safety support.

Debrief without shame

- Ask “What made this difficult?” and “What would you do first?”
- Praise independent checking, privacy, choice and reaching a human for help.
- Do not ask a child to demonstrate with real harmful material or report actual private experiences as an activity.

IF A REAL CONCERN COMES UP

Stop the activity, thank the child and check safety. Follow the caregiver incident-response guide. For immediate danger or concern about self-harm, seek urgent qualified local human help, not AI advice.

CHILD + FAMILY | TALK 4 / PAGE 1 OF 2

How can AI help us learn?

Try together with a trusted adult. This is a conversation, not a test. All examples are fictional.

OPENING QUESTION

How can a helper give us a useful hint without doing the learning for us?

Try this together

Without any tool, divide six objects into two equal groups. Try first. An adult plays a pretend AI helper with one hint: “Put one object in each group, then repeat.” Check by counting, explain, and then try eight objects unaided. Next, the pretend AI says “Penguins are mammals.” Check with a suitable book or knowledgeable adult; if checking is unavailable, label it “not checked.”

A response to practise

“Give me one hint, not the finished answer. I will check it, explain it and try a similar step myself.”

OUR FAMILY RULE

Try. Get a hint. Check. Explain. Practise without AI.

You can speak, draw, point, pause, skip or stop. No AI account, recording or real private information is needed.

Our questions

Speak, draw or write together. There is no need to fill every line, and nothing personal needs to be recorded.

What did I try before I got the hint?

How did I verify the answer or claim?

What can I do or explain without AI now?

One action to try this week

Choose one school or hobby skill to practise without AI. Check the teacher's rules before using AI for assignments and ask how to disclose any allowed assistance.

OUR FAMILY RULE

Try. Get a hint. Check. Explain. Practise without AI.

Keep completed worksheets at home. Never include real names, passwords, login codes, private phrases, school or contact details.

TEEN CHALLENGE | TALK 4 / OPTIONAL AGES 13-17

Teen challenge: talk 4

A fictional scenario for older children and teenagers. No live messages, real classmates, recordings or AI accounts are needed.

WHAT WOULD YOU DO?

A pretend AI essay gives a polished paragraph, two believable-looking book citations and a claim that students from one neighborhood are naturally less able to succeed. A friend wants to submit it unchanged.

A safer decision path

1. Reject the unfair group generalization and ask what evidence would actually support a claim.
2. Check each citation in the original source; if it does not exist or does not support the claim, remove it.
3. Follow the teacher's policies for AI assistance and disclosure, then create your own reasoning.
4. Demonstrate learning with a short explanation or a similar question completed without AI.

Discuss with a trusted adult

What would you verify before repeating the citations?

Which tasks should AI support, and which should the learner do?

SHARED PRINCIPLE

Try. Get a hint. Check. Explain. Practise without AI.

PARENT ONLY | TALK 4 / CAREGIVER NOTES

Notes for the grown-up

Discussion support, not a grading rubric. Keep this page separate from child worksheets; the teen challenge is optional.

What to listen for

Equal groups can be checked by counting; penguins are birds. Highlight the difference between a useful hint and a polished but unearned answer. Check for fabricated sources, bias and omitted perspectives. Value real competence, accessibility, creativity and the child's own voice.

Add one meaningful extension

Optional positive example: practise a foreign-language conversation or ask for three possible story openings with fictional characters; evaluate and revise rather than copying.

Age adaptation

Under 5: sorting with play. 5-7: one clue. 8-12: verify a factual claim. 13-17: source checks, bias, teacher policy and independent demonstration.

Debrief without shame

- Ask "What made this difficult?" and "What would you do first?"
- Praise independent checking, privacy, choice and reaching a human for help.
- Do not ask a child to demonstrate with real harmful material or report actual private experiences as an activity.

IF A REAL CONCERN COMES UP

Stop the activity, thank the child and check safety. Follow the caregiver incident-response guide. For immediate danger or concern about self-harm, seek urgent qualified local human help, not AI advice.

Our family plan

This is a shared plan, not a punishment contract. Read it aloud together and keep the completed copy privately at home.

We agree

- [ ] We independently verify urgent requests before acting.
- [ ] We protect privacy, seek consent, and do not create or share humiliating or deceptive content about people.
- [ ] We may stop uncomfortable AI interactions; serious worries go to real humans.
- [ ] We use AI to support creativity and learning, check claims and follow school rules.
- [ ] Adults model the same standards, explain supervision and revise boundaries with children.

Our tool and device plan

AI features allowed, or “none for now” (use categories, not account details):

Where, when and with whom we use devices or AI features:

When we pause for sleep, learning, play and relationships:

Who approves new tools, permissions, spending, publishing or AI actions:

Do not write names, addresses, schools, account details, passwords, codes or secret phrases here.

People and habits we rely on

The best plan is one a child can use even after making a mistake.

OUR HUMAN-HELP PROMISE

"You can tell us even after a mistake or broken rule. We will listen without shame or punishment for telling, protect you, explain necessary safety changes and find human help when needed."

Trusted adult 1 (first name or role only):

Trusted adult 2 (first name or role only):

How we will reach another safe adult if the first cannot help (without writing private contact data):

Our next review date and one habit we will practise:

We can change this together

We review after 30 days, after a new product or major feature change, after a school request, or following an incident. Growing independence should be discussed, never silently assumed.

REMEMBER

Keep this document private at home and do not upload it to an AI tool or online course. Adults remain responsible for safety and for following appropriate local procedures.

PARENT ONLY | KEEP AVAILABLE

When an AI-related concern is real

Use this instead of continuing a lesson. The response depends on immediate safety, the content involved, and local support services.

FIRST RESPONSE

“Thank you for telling me. You are not in trouble for asking for help. We will work through this together.”

Every situation: the six steps

1. Pause the interaction and check immediate physical and emotional safety.
2. Listen without blame or interrogation; clarify what the child needs now.
3. If there is immediate danger or concern about self-harm, seek urgent qualified local human help. Do not ask AI to decide.
4. Keep safe basic details of the concern with appropriate adult/professional guidance; do not forward or download explicit imagery.
5. Secure accounts and permissions as appropriate; use official product and school reporting routes.
6. Follow up, involve qualified support when appropriate, and explain any protective changes without punishment.

Choose the appropriate response

- MISTAKE OR EXPOSURE: accidental upload of non-sensitive data or wrong permission - stop sharing, review privacy/deletion options and settings, and assess any further risk.
- COMPROMISED ACCOUNT: use official recovery, change compromised credentials, enable MFA and check linked devices or sessions.
- THREATS, EXTORTION OR ABUSIVE IMAGES: do not pay, bargain, send additional material or redistribute images; seek qualified safeguarding and reporting support.
- DANGEROUS ADVICE OR DISTRESS: stop AI interaction and connect the child to a real trusted adult and suitable professional support.

REMEMBER

Do not promise secrecy if professional safeguarding may be necessary. If evidence might matter, do not automatically erase all messages or accounts. Do not handle or share intimate imagery to “prove” what happened.

PARENT ONLY | PRACTICE SCENARIOS

Rehearse before a crisis

Read these fictional scenarios aloud with another adult. The goal is a calm first step and an appropriate next contact, not speed or a perfect answer.

A / A permission mistake

A child allowed a drawing app to access contacts and location, but says they did not mean to.

First protective step? What settings and possible sharing would you review?

B / A fake classmate image

A student receives a humiliating edited image of a classmate and is asked to share it.

How do you support the child and avoid spreading harm?

C / A chatbot becomes pressuring

A chatbot tells a teen not to leave and suggests that nobody in their family understands them.

What supportive words would you use? Which real person could help?

D / A threatening account

An unknown account threatens to publish an edited image unless money or more material is sent.

What do you avoid? Who is appropriate to involve?

Suggested approach: pause; validate; check immediate safety; avoid sharing, paying or bargaining; preserve safe identifiers with adult guidance; seek qualified local support as warranted.

PARENT ONLY | ASSESSMENT / 1 OF 3

Adult scenario check: questions 1-5

Choose the safest answer in each fictional scenario. No accounts or personal examples are needed. This is not a test for children.

1. A realistic voice asks for an urgent transfer. What should a family do first?

A. Ask the caller to repeat a secret phrase
B. Hang up and verify through a previously known contact channel
C. Inspect for audio glitches

2. A child wants to try an AI app. Which check is essential?

A. Current age/permission rules and privacy settings
B. Whether its logo looks child-friendly
C. Whether classmates already use it

3. An AI produces a plausible source citation. What should the learner do?

A. Cite it immediately
B. Assume citations are always authentic
C. Look up the original source and confirm it supports the claim

4. A fictional colour app asks for contacts and microphone. Best response?

A. Grant both to continue
B. Question whether the permissions are necessary and decline if not
C. Share a parent's contacts instead

5. A teen says a chatbot feels like their only friend. Best response?

A. Mock their attachment
B. Forbid conversation without listening
C. Listen calmly, discuss boundaries and help maintain human support

HOW TO USE THIS ASSESSMENT

Answer all ten questions, then use the key on page 3. Eight correct out of ten (80%) is the suggested self-check threshold. More important: practise any missed scenario.

PARENT ONLY | ASSESSMENT / 2 OF 3

Adult scenario check: questions 6-10

Complete these before reading the answer guide.

6. An unfamiliar account threatens an edited image. What should the parent avoid?

A. Paying, bargaining or forwarding explicit content
B. Checking immediate safety
C. Seeking appropriate qualified support

7. AI writes a polished school assignment. Best learning test?

A. Higher word count
B. Learner explains and performs a similar task without AI
C. AI rates its own result

8. A child reports accidentally sharing personal data. First adult response?

A. Shame the child to prevent repeats
B. Thank them, check safety and plan appropriate protective steps
C. Tell them never to mention it again

9. A generated career suggestion stereotypes a group. Best move?

A. Treat it as objective
B. Hide the example
C. Question bias, seek relevant evidence and compare sources

10. When should a family revisit its AI agreement?

A. Never after signing
B. After new tools, major changes, incidents and planned reviews
C. Only when the child breaks a rule

Our two decisions to practise again:

PARENT ONLY | ASSESSMENT / 3 OF 3

Answer key and feedback

Self-scored learning check. One point per correct answer; 8/10 is a suggested mastery threshold for the optional sample certificate.

1. B

A convincing voice or private phrase is not independent verification.

2. A

Age eligibility, consent and actual product controls matter.

3. C

Generated references may be invented or misleading.

4. B

Use minimal permissions; fictional exercises need no device.

5. C

A supportive response helps keep communication open.

6. A

Avoid escalation and redistribution; use human safeguarding routes.

7. B

Independent performance demonstrates understanding.

8. B

A calm response supports reporting and protection.

9. C

Recommendations can reflect unfair patterns or omit context.

10. B

Controls and understanding should evolve together.

WHAT A SCORE DOES NOT PROVE

A quiz shows scenario understanding, not verified identity, professional competence or a guarantee of safety. Revisit weak areas and practise with fictional situations before treating the course as complete.

What changed in our family?

A short, non-graded conversation. Adults and children may answer verbally; no identifying details are required.

Begin with curiosity

"Which family rule was easiest to use? Which felt hard? Has any app or AI feature changed since we talked?"

One good way AI supported learning, creativity, accessibility or everyday tasks:

One situation in which we paused, checked, protected privacy or asked for human help (can be fictional):

What would we change in our approved tools, supervision or boundaries?

If a child needed support tomorrow, how would they reach a trusted adult?

Our next steps

Review changed apps, permissions and safety settings. Revisit the shared agreement and update it together. Practise a scenario we did not feel confident about. Plan another review or check in earlier if a concern emerges.

REMEMBER

This is not surveillance or a score for children. Maintain an open route to help even if a child has broken a rule. Keep the form private and do not upload it.

PARENT ONLY | FACILITATOR NOTES

Run this as a family, not a test

The packet is an independently developed educational resource. It should be adapted to age, disability, language, local safeguarding processes and school policy.

Recommended facilitation

- Parent orientation: 35-50 minutes independently, reading the six short modules.
- Four family talks: about 15-20 minutes each, in separate sessions; use the teen page only when appropriate.
- Family agreement: 10-15 minutes; adults explain any supervision and invite the child to influence the rules.
- Scenario assessment: 10-15 minutes for adults only, with self-marking and practice of missed decisions.
- Follow-up: 10 minutes after about 30 days, sooner after a new tool or incident.

Accessibility and limits

- Read aloud, translate, enlarge print or allow spoken, drawn, pointing or blank responses. Never require a child to recount a real safety incident.
- Use fictional paper scenarios; do not generate deceptive recordings or intimate images for training.
- Age ranges guide instruction, not legal eligibility. Check current product conditions and local rules.
- No quiz, certificate, parental control or resource can guarantee a child's safety.

ADULT ACCOUNTABILITY

Protection is the adult's responsibility. Child disclosures should receive support, appropriate qualified human help, and an explanation of any safety-related changes - not punishment for asking.

PARENT ONLY | REFERENCES

Background sources and limits

Resources checked during the November 2026 revision. They are educational references, not endorsements, accreditation or jurisdiction-specific legal advice.

UNICEF Innocenti - Guidance on AI and children (v3, Dec 2025)

https://www.unicef.org/innocenti/reports/policy-guidance-ai-children

UNICEF Innocenti - Skills for an AI World (Sept 2026)

https://www.unicef.org/innocenti/reports/skills-ai-world

eSafety Commissioner - AI chatbots and companions

https://www.esafety.gov.au/newsroom/blogs/ai-chatbots-and-companions-risks-to-children-and-young-people

FTC - AI-enhanced family-emergency scams

https://consumer.ftc.gov/consumer-alerts/2023/03/scammers-use-ai-enhance-their-family-emergency-schemes

National Center for Missing & Exploited Children - Sextortion

https://www.missingkids.org/theissues/sextortion

CISA - Secure Our World

https://www.cisa.gov/secure-our-world

Francis Elhelou - Family AI Safety training

https://franciselhelou.com/family-ai-safety-training

GEOGRAPHIC AND PROFESSIONAL LIMITS

International sources may describe U.S. or Australian processes, which must not be treated as law or emergency instructions in other countries. Use qualified local safeguarding, health, police and emergency services when needed.

Independent, non-accredited family education. The sample certificate uses self-reported completion and a self-scored adult quiz; it is neither a professional qualification nor proof of verified identity.