



MindSpark Artificial Intelligence (AI) Framework

Safety and Responsibility

MindSpark Ventures Inc.

Version 1 | October 2026

Table of Contents

Table of Contents.....	2
Executive Summary.....	3
Our Public Commitments.....	3
1 Our Approach to Artificial Intelligence.....	3
AI Supports Learning Rather Than Replacing It.....	3
“Quiet” AI.....	4
Controlled Expansion.....	4
2 What AI May and May Not Do.....	4
Appropriate Uses.....	4
Prohibited Uses.....	4
Human Review.....	5
3 Safety and Responsibility Assessment.....	5
4 How MindSpark Reviews AI.....	6
Before an AI Feature Is Released.....	6
During Use.....	6
Review Schedule.....	6
What Happens When a Test Fails.....	7
5 How Approved Source Libraries Are Used.....	7
Why MindSpark Uses Approved Source Libraries.....	7
Safeguards.....	7
6 Student Teacher and Parent Oversight.....	7
What Students and Adults May See.....	7
What MindSpark May Monitor Internally.....	8
Important Limits.....	8
7 MindSpark AI Agents.....	8
Standards for Every Agent.....	9
8 AI Advisory Committee.....	9
Membership.....	9
Responsibilities.....	10
9 Concerns Incidents and Continuous Improvement.....	10
10 Public Commitments.....	10
References and Further Reading.....	11
Document Governance.....	12

Executive Summary

MindSpark uses artificial intelligence to support learning, not to replace students, teachers, or parents. Our goal is to help students plan meaningful projects, receive useful guidance, reflect on their progress, and show what they have learned. Adults remain responsible for important educational decisions.

MindSpark is a personalized learning platform for students in grades 6 through 12. Students complete projects through a series of stages, tasks, reviews, and reflections. Teachers and parents serve as mentors and evaluators. Artificial intelligence, or AI, works quietly behind these experiences. It may organize information, suggest resources, help students and educators design a project, or identify information that deserves human attention.

MindSpark does not give AI final authority over grades, mastery, discipline, placement, access to services, or other decisions that can significantly affect a student. When an AI-supported recommendation matters, an authorized adult must review it and may change or reject it.

MindSpark has established important safeguards, including secure sign-in, role-based access, encrypted connections, system monitoring, adult review, and the ability to turn off AI services when necessary. These include regular safety testing, clearer records of how each AI feature works, testing of deletion procedures, accessibility reviews, fairness reviews, and better measurement of whether an AI feature is actually helping (MindSpark Ventures Inc., 2026).

This framework explains our current direction and the standards we expect MindSpark to meet. It is a public accountability document. It does not claim that AI is error-free, and it does not treat planned improvements as completed work.

Our Public Commitments

Commitment	What It Means
Learning comes first	AI must serve a clear educational purpose. A feature should help students learn, create, revise, reflect, or demonstrate understanding.
People remain responsible	Students do their own thinking. Teachers and parents guide learning. Authorized adults make important decisions.
Data use is limited	MindSpark uses only the information needed for an approved purpose and limits who can access it.
Users receive clear information	Families, educators, and students should understand when AI is involved, what it does, and where it may be wrong.
Safety is checked repeatedly	AI features are reviewed before release and again as models, content, users, or risks change.
Problems lead to action	Concerns are investigated, features may be paused, and software changes must be tested before normal use resumes.

1 Our Approach to Artificial Intelligence

AI Supports Learning Rather Than Replacing It

MindSpark is designed around student ownership. Students choose, plan, create, revise, and reflect. AI may help them move through that process, but it should not complete assessed work for them or remove the productive struggle that helps learning happen.

Teachers and parents remain central. They understand the learner, the classroom or family context, and the meaning of the student's work in ways that an automated system cannot. AI can provide information or a draft, but adults decide how that information should be used.

This approach is consistent with responsible-AI guidance that emphasizes educational value, child wellbeing, human oversight, privacy, fairness, and transparency (Khan Academy, 2025; The Learning Agency & Renaissance Philanthropy, 2026).

“Quiet” AI

MindSpark describes its approach as **quiet AI**. This means AI usually works within a defined learning activity instead of becoming the center of the experience. For example, it may help an educator check whether a project addresses selected learning standards or help organize information for a parent to review.

Quiet AI does not mean hidden AI. Users should still be told when AI plays a meaningful role. The term describes the place of AI in the learning experience: helpful, limited, and accountable to people.

Controlled Expansion

MindSpark uses a controlled-expansion approach. New AI features begin with a limited purpose and a limited group of users. Before a feature is used broadly, MindSpark should be able to explain:

- What educational need the feature addresses.
- Who will use it and who will review its work.
- What information it uses.
- What the feature is not allowed to do.
- How accuracy, safety, privacy, accessibility, and fairness were checked.
- How the feature can be corrected, limited, or turned off.

2 What AI May and May Not Do

Appropriate Uses

Depending on the product/feature and the user's role, MindSpark AI may:

- Help educators design project plans, learning tasks, questions, rubrics, and assessment materials.
- Suggest learning resources connected to an approved project or standard.
- Help organize project information for students, teachers, or parents.
- Draft summaries that an authorized person reviews before use.

- Help identify patterns in student progress, revision, reflection, or support needs.
- Assist technical staff with approved system, development, or operational work.

These uses remain subject to access controls, human review, and the limits described in this framework.

Prohibited Uses

MindSpark does not authorize AI to:

- Make final decisions about grades, mastery, discipline, placement, eligibility, or access to educational services.
- Present a guess as a verified fact.
- Complete assessed student work in place of the student.
- Use protected or sensitive personal characteristics to make inappropriate decisions about a learner.
- Use information for a new purpose without review and authorization.
- Share personal information with people or systems that are not authorized to receive it.
- Train an outside company's AI model on MindSpark customer information unless a specific agreement and governance decision allow it.

Human Review

Human review (referred to as “human in the loop”) means more than placing a person at the end of an automated process. The reviewer must have enough information and authority to question the result, correct it, reject it, and explain what happened.

3 Safety and Responsibility Assessment

MindSpark's company direction is generally aligned with responsible-AI guidance. The table below presents that conclusion in public-facing language. These are management findings, not independent certifications.

Area	What is Working	Ongoing Work
Educational value	Projects, review gates, revision, reflection, and adult oversight keep the focus on learning.	Measure the quality and educational value of each AI feature, not just overall product activity.
Student safety and wellbeing	AI is limited to defined activities, adults review important work, and AI services can be turned off.	Conduct regular safety challenges, set response expectations, and make concern reporting easier.
Privacy and security	Secure sign-in, role-based access, encrypted connections, managed credentials, and monitoring are in place.	Regularly review privacy and security policies with customers and our AI advisory board.
Accessibility and fairness	Human review can identify some problems, and MindSpark has committed to equitable access.	Complete accessibility testing and examine whether AI quality differs across groups of learners.
Transparency	MindSpark describes AI as advisory and recognizes the need for human review.	Give users clearer notices about data, limitations, sources, review, and appeals.

Area	What is Working	Ongoing Work
User control	Schools, educators, and parents have oversight roles.	Establish regular review cycles.

The assessment supports one clear conclusion: MindSpark focuses on evidence, testing, and transparency before adding more AI features (MindSpark Ventures Inc., 2026).

4 How MindSpark Reviews AI

MindSpark reviews AI throughout its life, from the first idea through use, improvement, and retirement. The level of review depends on the age of the users, the sensitivity of the information, and the possible harm if the AI is wrong.

Before an AI Feature Is Released

MindSpark conducts a thorough review of any AI-related interaction. Our developers test all prompts using a variety of available AI models to ensure the accuracy and performance of that designated feature. Testing includes normal use as well as attempts to make the system behave in unsafe or unintended ways.

Examples of safety tests include:

- Does the AI invent facts or sources?
- Can a user trick it into ignoring its rules?
- Does it reveal information that should remain private?
- Does it provide harmful or age-inappropriate content?
- Does it give students completed answers when it should support their thinking?
- Does it treat similar learners or requests inconsistently?
- Can people using keyboards, screen readers, captions, or other assistive technology use the feature?

Public AI risk-assessment and product-safety resources inform this approach (UK Department for Education, 2026; Youth AI Safety Institute, n.d.).

During Use

MindSpark reviews reports from users, examples of AI output, system errors, corrections made by educators, and signs that a feature is operating outside its approved purpose. Technical records may be used to understand what happened, but access to those records is limited to authorized staff. Development advisors and systems managers are responsible for implementing the following review schedule.

Review Schedule

When	What MindSpark Reviews
Ongoing	Serious errors, safety concerns, unusual access, service failures, and reports from users.
Before a release	Purpose, data use, safety tests, accessibility, adult review, user notice, and the ability to reverse the change.

When	What MindSpark Reviews
Monthly	Samples of AI work, repeated problems, educator corrections, and unresolved concerns.
Quarterly	AI features, outside providers, access permissions, approved information sources, incidents, and overdue improvements.
Annually	Overall privacy, security, accessibility, fairness, effectiveness, and governance practices.
After a major change	The feature is reviewed again when its AI model, information sources, purpose, users, or possible consequences change.

What Happens When a Test Fails

Not every error has the same level of risk. MindSpark considers the possible harm, the number of people affected, the sensitivity of the information, and whether the problem could influence an important decision. Serious failures can block a release or require MindSpark to pause the feature until the problem is corrected and tested again.

5 How Approved Source Libraries Are Used

Some MindSpark AI features use an **approved source library**. The technical term for part of this process is a *vector store*. In plain language, it is a searchable collection of approved materials that helps AI find passages related to a user's question.

For example, an approved source library might contain learning standards, project instructions, school-approved resources, or MindSpark support documents. The AI searches that collection for relevant information and uses it to prepare a response. This process is often called *retrieval-augmented generation*, but users do not need to know that term to understand the safeguard: the AI is being directed toward approved sources instead of relying only on its general training.

Why MindSpark Uses Approved Source Libraries

- To make responses more relevant to the project or task.
- To connect answers to materials that MindSpark or an authorized customer has approved.
- To make it easier to identify the source behind a response.
- To reduce unsupported answers.

Safeguards

MindSpark reviews who owns each source, why it is included, who may retrieve it, and when it should be updated or removed. Personal or sensitive information is never included within our source library.

MindSpark also checks for outdated sources, duplicated information, conflicting guidance, and hidden instructions that attempt to change how the AI behaves. When information is deleted, the related searchable copies and indexes should also be removed according to MindSpark's retention practices and technical capabilities.

An approved source library improves grounding, but it does not guarantee that an answer is true or complete. The AI can retrieve the wrong passage, miss important context, or misunderstand a good source. Human review remains necessary.

6 Student Teacher and Parent Oversight

MindSpark uses monitoring to support learning, safety, and reliable service. Monitoring should be limited to a legitimate educational, support, safety, or operational purpose.

What Students and Adults May See

- The student's current project stage and tasks.
- Work that has been submitted or needs revision.
- Upcoming review gates or deadlines.
- Evidence connected to learning standards.
- Feedback, reflection, and progress information.
- Items waiting for teacher or parent review.

What MindSpark May Monitor Internally

Authorized staff may review whether users can sign in, whether services are working, whether information is being retrieved correctly, and whether an AI feature is producing repeated errors or unsafe results. MindSpark may also review when educators accept, change, or reject an AI suggestion so the company can identify weak features.

This approach reflects guidance that school technology should have a clear purpose, visible governance, and safeguards for children (Common Sense Education, n.d.; Fazlullah et al., 2026).

Important Limits

An automated warning is not proof that a student or educator did something wrong. If a warning could affect a person, an authorized adult must review the context. The person should have an appropriate opportunity to explain, correct, or challenge the information.

Access to monitoring information is based on role. MindSpark limits retention according to educational, contractual, business, and legal needs.

7 MindSpark AI Agents

MindSpark gives several AI features names so users and staff can understand their different jobs. We call these features **AI agents**. An agent is not a person and does not have independent authority. It is a named AI helper with a defined purpose, approved information, limited permissions, and a responsible human owner.

Not every agent or listed capability is available in every MindSpark product. Each agent must be approved for the specific product and user group in which it operates.

Agent	What the Agent Does	What MindSpark Checks
Ari	Supports Huddles and collaboration by helping connect participants, organize agendas, and identify next steps. Ari may support teacher-facing Ignite Insights. Ari is not the Family Studio agent.	Whether the right people can access the Huddle; whether summaries and next steps are accurate; whether personal information is handled appropriately; and whether Ari stays out of Family Studio workflows.
Satori	Provides help and guidance to students and educators without completing assessed work for the student. It is primarily a knowledge agent assisting students and educators with answers to questions.	Whether guidance is useful and age appropriate; whether Satori invents facts, gives unsafe advice, provides completed answers, or ignores the limits of the request; and whether unresolved concerns reach a person.
Kaizen	Helps students and educators understand growth, reflection, revision, feedback use, and Learning Velocity. Kaizen is a growth coach, not a grader.	Whether summaries match the student's actual work; whether growth measures can be checked; whether recommendations avoid unsupported conclusions; and whether Kaizen stays out of grading and mastery decisions.
Akira	Helps authorized staff monitor system health and approved safety signals.	Whether warnings are accurate; how often false warnings occur; whether only approved events are flagged; who can see monitoring information; and whether a warning is reviewed by a person before action is taken.
Kage	Helps authorized technical staff run and monitor approved background processes, such as scheduled reports or large processing jobs.	Whether each job was approved; whether it ran once and completed correctly; whether failures and repeated attempts are controlled; and whether staff receive clear notice when something goes wrong.
Ren	Serves as the Family Studio parent guide. Ren can answer parent questions and prepare draft rubrics, questions, evidence prompts, and short assessments for parent review.	Whether the parent is authorized to view the child's information; whether conversation memory stays within the approved family context; whether drafts are accurate and age appropriate; and whether parents review materials before using them.
Takumi	Supports Project Design Studio by helping educators develop project plans and check quality and alignment.	Whether projects match the educator's goals and selected standards; whether suggested activities are feasible, accessible, and fair; whether sources are valid; and whether an educator approves the project before classroom use.

Standards for Every Agent

MindSpark keeps a record for each agent that explains its purpose, owner, users, information sources, permissions, limits, adult reviewer, known weaknesses, test results, and current release status. If an agent changes models, gains a new source of information, serves a new group, or takes on a new job, it must be reviewed again.

Every agent is tested for inaccurate information, unsafe content, attempts to bypass rules, privacy problems, unfair results, accessibility barriers, and failure recovery. Serious problems block release or require the agent to be paused until the problem is corrected.

8 AI Advisory Committee

MindSpark's AI Advisory Committee provides outside perspective and internal accountability. It reviews evidence, raises concerns, and recommends improvements. The committee does not replace the responsibilities of company leadership, educators, schools, security staff, or legal advisers.

Membership

- Teachers and school leaders.
- Students and family representatives.
- AI and data specialists.
- Privacy and security specialists.
- Accessibility and learner-support specialists.
- People with expertise in ethics or child development.
- MindSpark leaders and product owners who can act on recommendations.

Student and family participation should be voluntary, age appropriate, supported, and protected from retaliation.

Responsibilities

The committee reviews higher-risk AI features, major changes, serious incidents, accessibility findings, fairness concerns, user complaints, and overdue corrective actions. It may recommend more testing, a smaller pilot, changes to a feature, a temporary pause, or retirement of a feature.

The committee should meet at least quarterly and more often when an urgent issue arises. Its records should identify the evidence reviewed, recommendations made, people responsible for follow-up, and expected completion dates. When leadership does not follow a significant recommendation, the reason should be documented.

9 Concerns Incidents and Continuous Improvement

Students, families, educators, schools, and partners should have a clear way to report AI behavior that appears inaccurate, unsafe, biased, inaccessible, or invasive of privacy.

When MindSpark receives a concern (via our support agent or support@mindsparkd.org), the company should:

1. Record and acknowledge the report.
2. Determine the possible harm, urgency, people affected, and information involved.
3. Limit or pause the feature when continued use could cause harm.
4. Review the AI model, instructions, source materials, permissions, and human process involved.
5. Correct the problem and notify affected parties when appropriate or required.
6. Test the correction before restoring normal use.
7. Document what happened and improve the related safeguards.

MindSpark maintains the ability to suspend AI-dependent functions when continued operation creates unacceptable risk. Restoring service requires confirmation that the immediate problem has been addressed and that appropriate monitoring is in place.

10 Public Commitments

MindSpark's next stage of responsible-AI work focuses on proving that safeguards work in practice.

Priority	What the public should expect...
Clear records for every AI feature	MindSpark documents the feature's purpose, owner, users, information, limits, tests, known weaknesses, and release status.
Repeatable safety testing	AI features are tested before release and after important changes, with failures corrected and tested again.
Meaningful measures	MindSpark measures the work the AI actually performs, including accuracy, usefulness, educator corrections, source quality, and learning value.
Accessibility and fairness	MindSpark tests with assistive technology, reviews generated materials for bias, and investigates meaningful differences in quality or outcomes.
Visible user choices	Users can find information about AI use and have practical ways to report concerns, request review, correct information, and seek deletion when applicable.
Public reporting	MindSpark provides understandable updates about governance, major improvements, known limitations, and the status of responsible-AI work.

This framework will be reviewed at least annually and after major changes to MindSpark products, AI models, information use, law, or risk. When MindSpark has not yet met a commitment, the gap should be recorded, assigned to an owner, and tracked until it is resolved.

References and Further Reading

Common Sense Education. (n.d.). *Tech in schools toolkit*.

<https://www.commonsense.org/education/tech-in-schools>

Fazlullah, A., Grosshans, H., Fox Johnson, A., Leasor, B., Weiss, D., & Barber, N. (2026). *Protecting kids online in the AI era: The path forward*. Common Sense Media.

<https://www.commonsensemedia.org/sites/default/files/featured-content/files/common-sense-media-protecting-kids-online-in-the-ai-era-the-path-forward.pdf>

Khan Academy. (2025, April 3). *Khan Academy's framework for responsible AI in education*.

<https://blog.khanacademy.org/khan-academys-framework-for-responsible-ai-in-education/>

The Learning Agency. (2026, October 7). *Building accessible and fair AI tools* [Webinar]. Ticket Tailor.

<https://www.tickettailor.com/events/thelearningagency/2417587>

The Learning Agency. (2026, September 30). *Mitigating the risks of Gen AI-powered tools* [Webinar recording]. Google Drive.

<https://drive.google.com/file/d/12cm9VZ-Z4S9-wcdfp0B378NO4-hjUgGA/view?usp=sharing>

- The Learning Agency, & Renaissance Philanthropy. (2026). *Guidelines for the safe and responsible use of AI in ed tech*.
<https://the-learning-agency.com/guides-resources/guidelines-for-the-safe-and-responsible-use-of-ai-in-ed-tech/>
- UK Department for Education. (2026, January 19). *Generative AI: Product safety standards*. GOV.UK.
<https://www.gov.uk/government/publications/generative-ai-product-safety-standards/generative-ai-product-safety-standards>
- Youth AI Safety Institute. (n.d.). *AI risk assessments*. Retrieved October 1, 2026, from
<https://institute.common sense media.org/risk-assessments>
- Zent, M., Smith, D., & Woodhead, S. (2025). *Question-anchored tutoring dialogues 2K* [Data set]. Hugging Face. <https://huggingface.co/datasets/Eedi/Question-Anchored-Tutoring-Dialogues-2k>

Document Governance

Owner	MindSpark Ventures Inc.
Audience	Students, families, educators, schools, districts, partners, and the general public.
Version	1
Effective date	October 2026
Review cycle	At least annually and after major changes.
Related materials	Privacy notices, terms, security practices, accessibility information, incident procedures, vendor agreements, and records for individual AI features.