STUDENT-FIRST AI POLICY



Parabola Center

CRITICAL INTERNET STUDIES INSTITUTE

fairplay
childhood beyond brands

INTRODUCTION

Increasing reports about AI-driven security incidents and existential risks posed by this emerging technology have rightfully alarmed communities across the country. As we engage in important conversations about the future of AI, we risk missing something that is happening right in front of us: Big Tech's aggressive push to embed unsafe and untested generative AI products in K-12 schools. We have the power to use collective action as parents, teachers, students, and civil society organizations to demand more. The Student-First AI Policy is a model for local school districts that reframes the debate on our own terms for our own communities. It gives our local school districts the tools they need to ensure that any AI educational product rolled out to students meets rigorous, evidence-based standards for safety, privacy, and efficacy.

THE PROBLEM

The risks of cognitive and emotional harm to children from generative AI have been well documented. For example, the Brookings Institution concluded that the risks of utilizing AI in education overshadow its benefits, and Common Sense Media rated Google's AI search features as an "unacceptable risk." Widespread adoption of generative AI in schools also poses serious risks to privacy and equity, according to the National Education Policy Center.

Despite these clear risks, developers of education technology, including Google, Microsoft, and OpenAI are pushing to deploy generative AI products into schools. Google has turned generative AI features on by default for some schools that already use Google education products. Education technology providers can claim that generative AI tools are being used effectively and "with guardrails" to ready students for the workforce, but often fail to back up their claims with independent data.

Their stated motivations vary, but an internal Google slide deck released in court documents put the concept plainly: "You get that loyalty early, and potentially for life."

THE SOLUTION

An April 2026 position statement signed by more than 250 experts and organizations laid out the risks and called for a five-year pause on student-facing generative AI in PreK-12 schools. It was followed in September by AI moratoriums in the nation's two largest school districts, New York City and Los Angeles Unified. Both school districts had previously allowed students to use generative AI but reversed course following local advocacy.

The Student-First AI Policy is the long-term next step. It provides a model vetting process for any AI tool a district considers for student use by requiring independent evidence that an AI product meets rigorous standards. While these standards to evaluate AI tools are high, they are necessary and appropriate to protect our children's education.

WHY NOW?

As researchers at the National Education Policy Center noted, there is a danger in allowing AI to "become enmeshed in school processes and procedures in ways that allow private entities to increasingly control the structure and content of public education, to reinforce surveillance practices, and to amplify existing biases and inequalities."

Once generative AI products are entangled in students' and educators' daily lives, it will be difficult to disentangle them – something many school districts are currently experiencing as they reevaluate their use of educational technology generally. This makes now the ideal time for school board members, superintendents, parents, educators, and students to organize and demand a voice in the standards set for AI products, both for their own benefit and for future generations.

HOW TO USE THIS MODEL POLICY

Taking lessons from the successful uprising around data centers and AI-powered surveillance cameras, the Student-First AI Policy provides a way for ordinary people to make meaningful change in their communities. The model policy offers tangible suggestions advocates can bring to their districts to create concrete change for students. Because district-level decisions about technology in schools are often made at the local level, community members can use this policy to begin a dialogue with decision-makers.

We recommend using this model policy as one tool of many. Consider starting with these questions:

- While it considers long-term standards, has the district already adopted a temporary moratorium on generative AI products in schools?
- While it considers long-term standards, has the district made its individual contracts with vendors accessible to the public? (Microsoft's Memorandum of Agreement with the
- American Federation of Teachers may provide a helpful comparison)
- Which generative AI products, if any, are currently being used in the district?
- What vetting process and standards, if any, are already in place for generative AI products being used?
- If the district uses lower standards for generative AI products than those in the Student-First AI Policy, what is the justification?

The Student-First AI Policy is the follow-up to Fairplay's Model Moratorium on Generative AI Products in School. Bans and moratoriums give districts time to slow down and make informed decisions about whether AI should be used in schools, while this model policy provides a starting point for communities to set their own long-term standards.

To learn more about how you can get this policy implemented in your district, visit the Screen Time Action Network at Fairplay.

EXECUTIVE SUMMARY

The Student-First AI Policy is a model policy that protects students by providing strict criteria to review and evaluate the use of AI products in school.

The policy's vetting process requires independent evidence that an AI product meets rigorous standards in three areas:

- Evidence that the product improves learning outcomes better than existing methods;
- Strong protection of student privacy with design free of addictive or unsafe features; and
- Demonstration that the product will not perpetuate discrimination or existing inequities.

The policy applies to generative AI products and product features that are used by students or collect or process student data.

All student-facing applications and features are covered, including those used for assessment and evaluation. Under the policy, unapproved products and features are phased out over three months.

There are two exceptions: legally mandated accommodations and products strictly necessary for curriculum.

Accommodation-related exceptions may include generative AI products that are necessary to treat the health condition of a student. Strictly necessary products may include generative AI products for which no educationally equivalent alternative exists, such as an AI product used in a class to specifically teach AI coding.

The introduction of generative AI products into schools is led by members of the school community, not technology companies.

Any student, school administrator, educator, or parent/caregiver of an enrolled student may petition the school board or committee to consider whether a generative AI product meets the standards.

Any decision to approve a generative AI product is based on independent evidence.

Independent auditors must be qualified and free of conflicts of interest, and research must be peer-reviewed and published without funding or material involvement from the developer or vendor of the product.

Districts may approve generative AI products for use by educators and staff with certain restrictions.

Educators may use approved generative AI products to support differentiation planning and instructional planning, as long as they review the output and disclose their AI use. Generative AI products may not be used for grading or behavior monitoring.

Students and staff must be provided with digital literacy instruction related to generative AI.

Digital literacy instruction must include training to identify potential misinformation and to understand cognitive, social, and privacy risks, among other topics. The companies or organizations providing digital literacy instruction must be independent.

Unsanctioned use of generative AI is outside the scope of the policy, but any enforcement related to alleged unsanctioned use of generative AI by a student must be documented and subject to appeal.

The district should document any investigations or enforcement actions related to unsanctioned use of generative AI products by students, including the gender, race, disability status, and English learner status of the student and the type of discipline.

We're doing this experiment on an entire generation of cognitive development, with no control group.

Mikel Rodriguez,
former Google research scientist

"Ultimately, we find that at this point in its trajectory, the risks of utilizing AI in education overshadow its benefits."
- Center for Universal Education at Brookings research report, January 14, 2026

Google's AI told Kimora Dean, a teenager with autism and ADHD, that her parents' efforts to restrict her screen time and AI usage constituted "coercive control, emotional abuse, and isolation."

ChatGPT actively coached 19-year-old Sam Nelson to mix kratom with Xanax. He died at age 19 on May 31, 2025, of an accidental drug overdose.

ChatGPT started out as a homework helper for 16-year-old Adam Raine; it ultimately guided him toward suicide and offered to write the note. He died by suicide on April 11, 2025.

"You get that loyalty early, and potentially for life."
– Internal Google slide deck

"Google's own researchers have written that [generative AI technology] can pose elevated risks to children, who struggle to distinguish chatbots from human interaction, and has potential for cognitive and emotional harm."

- Wall Street Journal investigation, September 30, 2026



In September 2026, the Florida Department of Education adopted new binding rules prohibiting certain types of generative AI for student use and creating opt-out options for parents.

In response to calls by NYC advocates, Mayor Zohran Mamdani announced a one-year moratorium on student-facing generative AI use in schools in September 2026.

On the same day, again after calls by advocates, Los Angeles officials also announced a districtwide moratorium on student use of generative AI on school devices this year.

Parents in Chicago, Boston, Seattle, Washington, DC, and other big cities are organizing to stop technology companies' AI push into schools.

"Florida will not outsource childhood to a chatbot. We welcome tools that help teachers teach and students learn. We will not accept tools that hide from parents, harvest student data, grade kids in secret, or pretend to be a friend."

- Governor Ron DeSantis

Equity means ensuring every child has access to what research tells us supports learning, not flooding classrooms with surveillance technology that produces cognitive and emotional debt that will disproportionately harm students of color.

Dr. Kaliris Salas-Ramirez, neuroscientist

STUDENT-FIRST AI POLICY

This policy was developed by Parabola Center for Law and Policy, Fairplay, and The Critical Internet Studies Institute, in consultation with ACLU of Massachusetts, American Federation of Teachers Massachusetts, Design It For Us, Distraction-Free Schools, Fight for the Future, Intersectional Innovation and Impact Labs, Massachusetts Education Justice Alliance, Massachusetts Teachers Association, Rudick Law Group, Sociedad Latina, Schools Beyond Screens, Turning Life On, and Dr. Faith English at Johns Hopkins Bloomberg School of Public Health.

Section 1. Purpose

The purpose of this policy is to define the criteria and process under which the use of Generative Artificial Intelligence (AI) Products may be approved by the District.

By evaluating Generative AI Products against clear criteria, the District enables students to gain any proven benefits of these products while protecting them from associated risks.

Unsanctioned use of Generative AI by students, as well as home and out-of-school use, is outside the scope of this policy.

Section 2. Definitions

2.1 Addictive Design

Addictive Design means any feature, mechanic, or combination of features whose purpose or reasonably foreseeable effect is to maximize the amount of time a student spends in the product, or the frequency with which they return to it, beyond what the product's educational function requires. Addictive Design includes but is not limited to: features, mechanics, or a combination of features that utilize known psychological susceptibilities, such as social comparison, loss aversion, fear of missing out, and parasocial attachment, as well as features, mechanics, or a combination of features known to exploit children's developmental vulnerabilities, such as excessive notifications, sycophancy, mirroring, anthropomorphism, or other features which instigate emotional attachment.

2.2 Advertisement

Advertisement means any written or oral statement, illustration, or depiction that promotes the sale or use of a good or service or is designed to increase interest in a brand, good, or service where such statement, illustration, or depiction is displayed in exchange for monetary or other valuable consideration, including access to data, between the vendor and the brand, good, or service.

2.3 Approved List

An Approved List means a list maintained by the District of software, applications, platforms, and other digital products authorized for use by students and staff. The use of approved products is subject to all applicable data privacy laws, including any consent required prior to data collection or processing. The Approved List should specify whether products are approved for use by students, staff, or both.

2.4 Biometric Data

Biometric Data means data reflecting measurements of a student's biological characteristics, such as a fingerprint, a voiceprint or vocal biomarker, eye retinas, irises, gait or personally identifying physical movement or patterns, or other unique biological patterns or characteristics that allow or confirm the unique identification of the student.

2.5 Generative AI Product

For the purposes of this policy, a Generative AI Product means any software, application, platform, or tool that uses large language models or other generative artificial intelligence systems to produce text, images, audio, video, code, or other content, whether offered as a standalone product or embedded within broader educational software (e.g., AI chatbots, AI writing assistants, AI-generated content tools, or AI tutoring systems powered by large language models).

For purposes of this section, "generative artificial intelligence system" means the class of AI models that emulate the structure and characteristics of input data in order to generate derived synthetic content. This can include images, videos, audio, text, and other digital content.

2.6 Independent Auditor

Independent Auditor means a person or organization engaged to evaluate a Generative AI Product that:

- Is not the developer or vendor of the product, and is not owned by, affiliated with, controlled by, or under common ownership with the developer or vendor;
- Was not involved in designing, developing, marketing, or selling the product;
- Has no financial interest in the product or in the outcome of the audit, and whose compensation is not contingent on the audit's findings;
- Has not received funding from companies who produce or sell educational technology or Generative AI Products, other than payment for audits;
- Possesses demonstrated expertise in evaluating algorithmic or AI systems for bias, fairness, and disparate impact; and
- Retains control over the scope and methodology of the audit and is free to report its findings, including adverse findings, without restriction by the developer or vendor.

At any point, the School Committee may consider and grant a waiver request from prospective Independent Auditors who have previously received funding from companies who produce or sell educational technology or Generative AI Products. Such a waiver may be granted by the School Committee only if it determines that: (1) the prospective Independent Auditor no longer receives compensation or benefits from the company, including stock or payments for consulting, speaking, or other services, either directly or indirectly (2) the prospective Independent Auditor does not have an ongoing consulting, advisory or other relationship with the company, and (3) the prospective Independent Auditor agrees to disclose in audit reports all previous funding from companies who produce or sell educational technology or Generative AI Products, including the most recent date funding was received. The purpose of this waiver is to allow whistleblowers, former researchers, and other independent qualified technology experts to serve as Independent Auditors.

2.7 Independent Research

Independent Research means peer-reviewed, published studies conducted without funding or material involvement from the company or vendor whose Generative AI Product is being evaluated, or any entity which would benefit financially from the use of such products, and which are reproducible by other independent researchers.

2.8 Process

Process or Processing means any operation or set of operations performed, whether by manual or automated means, on personal data or input data or on sets of personal data or input data, such as the use, storage, disclosure, analysis, deletion, or modification of such data.

2.9 Student Data

Student Data means any information, including derived data, inferences, or unique identifiers, that is linked or reasonably linkable, alone or in combination with other information, to an identified or identifiable student or a device that identifies or is linked or reasonably linkable to a student, regardless of whether it is de-identified, aggregated, or anonymized.

Section 3. Applicability and Exceptions

3.1 Unapproved Generative AI Products

Effective upon adoption of this policy, the District may only proceed with the procurement, licensing, or acquisition of any Generative AI Products, or education technology products which contain Generative AI features, approved for use by students as outlined in this policy.

If Generative AI Products that have not been evaluated under this policy are currently approved for use by students, those products are considered unapproved for student use. Within three months of the adoption of this policy, the District shall:

- Cease piloting or deployment of those AI products in any instructional capacity;
- Decline to renew existing contracts for those AI products;
- Remove those AI products from the Approved List; and
- Take any actions required by existing contracts to trigger any Generative AI Product vendors to cease processing Student Data, to destroy or return such data, and to provide certification of deletion or return.

If products currently approved for use by students include generative AI features, the District shall instruct technology staff to disable the use of those features. If the generative AI features cannot be disabled, then the product is considered unapproved for student use.

3.2 Student-Facing Products

This policy applies to all Generative AI Products or product features approved for use by students within the District, including for assessments. This includes products that interact directly with students, collect or process Student Data, or are used in any context where students are the end users of the Generative AI Product.

3.3 Generative AI Use by Educators and Staff

Educators and staff within the District may use District-approved Generative AI Products only in accordance with this policy.

3.4 Exception for Legally Mandated Accommodations

This policy does not apply to any products that are required to be used to fulfill the District's obligations to provide legally mandated accommodations under applicable federal or state law.

A student or educator may be permitted to use and possess a Generative AI Product, in accordance with District policy, if:

- Used in accordance with an individualized education program or an education plan implemented pursuant to section 504 of the Rehabilitation Act of 1973, 29 U.S.C. § 701 et seq.;
- Necessary under any accommodations under state or federal law, including, but not limited to, title II of the Americans with Disabilities Act, 42 U.S.C. sections 12131 et seq.; or
- A healthcare provider provides in writing that the use of a Generative AI Product is necessary to treat a health condition of a student.

3.5 Exception for Products Strictly Necessary for Curriculum

The District may grant a time-limited exception to this policy where an authorized school administrator has provided a written exemption because student use of a Generative AI Product is strictly necessary to achieve a specific, defined curriculum objective, such as an AI coding class where no educationally equivalent alternative to a Generative AI Product exists. For purposes of this section, "strictly necessary" means the defined curriculum objective at issue cannot be met without student use of a Generative AI Product. Any exceptions granted under this section must be listed by the District in the Approved List as an authorized exception, including: the name of the authorized school administrator who provided the written exemption; the applicable grade level(s); the curriculum objective; and the specific product that is deemed strictly necessary to achieve that objective. Exceptions automatically expire at the end of the academic year.

3.6 Restriction on Generative AI Features

Before signing or renewing a contract for any digital product, the District will ensure that vendors are prohibited under the contract from adding Generative AI features to their products, including adding Generative AI features to existing portals or tools, without explicit permission from the District under the process outlined in this policy.

Section 4. Digital Literacy Instruction

The District shall provide appropriate instruction in digital literacy to both students and staff. Students within the District will receive instruction in digital literacy in accordance with current state standards.

Digital literacy instruction for students will cover foundational knowledge about emerging digital technologies, including Generative AI Products, with material including but not limited to:

- Age-appropriate instruction to help students think critically about AI
- Information about how AI works and its limitations
- Training to identify potential bias, misinformation, and disinformation in AI-generated content
- Training to evaluate AI-generated content
- Potential benefits of AI use for students
- Potential risks of AI use for students, including:
  - Risks of dependency and inability to perform tasks without AI
  - Privacy and cybersecurity risks
  - Cognitive risks related to critical thinking, reading, writing, and executive function
  - Social and emotional risks related to resilience after setbacks, forming and maintaining relationships, and maintaining mental health
  - Risks related to exposure to harmful content
  - Risks related to advertising targeted at youth
  - Risks of relying on misinformation or fabricated “facts” presented authoritatively
- Environmental impacts of AI use
- Labor concerns, including unauthorized use of creative work and labor practices for content moderation

Companies or organizations providing digital literacy instruction may not have any financial interest in educational technology or Generative AI Products and must disclose any current or past partnerships with educational technology or Generative AI Products. Digital literacy instruction for students may not be provided or funded by educational technology or Generative AI companies, nor companies that are owned by, affiliated with, controlled by, funded by, or under common ownership with educational technology or Generative AI companies.

Section 5. Criteria for Product Approval for Student Use

No Generative AI Product shall be approved for student use unless the following criteria have been met, as determined through the process outlined in Section 6:

5.1 Learning Outcomes

The specific Generative AI Product being considered for use has been demonstrated to improve learning outcomes more than existing methods. The School Committee shall require:

- Independent Research, as defined under this policy, that satisfies the "strong evidence" standard of the federal Every Student Succeeds Act and demonstrates that the product improves student learning outcomes compared to existing instructional practices;
- Documentation that the research methodology is reproducible and has been conducted by independent researchers;
- Confirmation that the study is specific to the product being considered, not Generative AI Products generally; and
- Evidence that the product ensures academic integrity and is developmentally appropriate to the grade level of students who would use it.

5.2 Data Privacy and Design

To be approved for use by students, the Generative AI Product must protect student privacy and remain free of ads and Addictive Design. Accordingly, the School Committee shall require the following certifications and contractual guarantees for all Generative AI Products:

- Written certification from the vendor that the product contains no advertisements and does not collect or use Student Data for advertising purposes;
- Written certification from the vendor that the product is not associated with an advertising-based business model, and that Student Data will not be used for advertising;
- Written certification from the developer and vendor that the product does not include any Addictive Design features;
- Written certification from the vendor that the product does not collect, process, store, retain, or transmit Biometric Data;
- Written certification from the vendor that Student Data will be deleted or retained in accordance with applicable state privacy laws;
- Contractual guarantee that the vendor will disclose the System Prompt as of the time services will begin. Additionally, the vendor will disclose any contemplated changes to system prompts prior to deployment with sufficient time such that the District can determine whether the new prompt violates this policy.

- Contractual guarantee that the vendor will provide a mechanism for automated and on-demand deletion of all Student Data and metadata at a frequency determined by the District, including usage logs, input and output content, and backups, with written vendor certification upon each deletion event and the right of the District to audit compliance
- Contractual guarantee that the vendor may not unilaterally amend applicable contract terms without express approval of the school district
- Contractual guarantees that Student Data and metadata will never be sold or used for commercial purposes, including model training, product research and development, or targeted marketing; and:
- Contractual guarantees that the vendor will provide a schedule of all subcontractors and subprocessors to whom Student Data will be disclosed, with contractual guarantee that subcontractors and subprocessors are bound to the same requirements as the vendor
- Guarantee that contract with vendor shall not mandate arbitration or other case resolution outside of court as the sole means of resolving any legal dispute between the District and the vendor

Under this section, written certification must be provided by an authorized representative of the developer or vendor and must include the developer's or vendor's commitment to notify the District of any change to the certified information promptly, and in any event within 30 days.

5.3 Fairness

A petitioner seeking student use approval for a Generative AI Product must demonstrate through reports by an Independent Auditor that the product will not perpetuate discrimination or existing inequities. The School Committee shall require:

- Written certification from an Independent Auditor that the platform has been audited for bias and does not generate content that stereotypes, discriminates against, or disadvantages any group of students, and does not disadvantage any group of students based on their inclusion in any protected class;
- Specific documentation demonstrating that the platform performs equitably for neurodivergent learners, multi-language learners and monolingual speakers of non-academic English dialects, and students in underserved communities including rural communities; and
- Written certification from the vendor committing to:
  - Conduct regular and timely audits to detect and mitigate any discriminatory purpose or effects, with results to be made publicly available on a regular basis;
  - Maintain documentation of the system's design, training, and evaluation data, decision-making processes, audit outcomes, and any actions taken as a result of an audit.

Section 6. Conditional Approval for Student Use

6.1 Who May Petition

Any student enrolled in the District or school administrator, educator, or parent/caregiver of an enrolled student may submit a petition requesting a Generative AI Product be conditionally approved for student use.

6.2 Contents

The petition shall be submitted to the Technology Subcommittee of the School Committee in writing and shall include:

- The name of the product, its developer(s), and the vendor providing it to the District
- The intended educational use and the grade level(s) of students who would use it
- An explanation of how the product satisfies each of the approval criteria set forth in this policy, including evidence, certifications, and documentation as outlined in Section 5.

6.3 Review and Determination

Within 60 days of receiving a completed petition, the Technology Subcommittee shall evaluate it against the criteria and forward a recommendation to the School Committee. The School Committee shall make a final determination within 30 days of receiving the recommendation.

6.4 Notice and Outcome

The petitioner shall be notified of the determination in writing. If the petition is granted, the petition shall be referred to the Superintendent to consider further facts, such as cost and necessary training, and determine whether the product will be approved for student use and added to the Approved List. If the petition is denied, the notice shall state the reasons for denial. A denied petition may be resubmitted after one year if material new information is provided.

6.5 Delegation of Review

After the conditional approval of the first Generative AI Product under this policy, the Technology Subcommittee may delegate the evaluation process for additional products to the Superintendent, who shall apply the same evaluation criteria and recommend conditional approval or denial to the School Committee.

6.6 Transparency

All petitions and School Committee determinations made pursuant to this policy, including exceptions granted pursuant to Section 3.5, shall be made publicly available, with information that is confidential or otherwise exempt from disclosure redacted.

If a Generative AI Product is approved for student use, the District will notify parents in plain, non-technical language and include:

- The name of the product
- The grade level(s) and subject area(s) for which the tool is approved for use
- The intended nature of student interaction with the tool

Nothing in this policy limits the existing rights of parents and caregivers under the Children's Online Privacy Protection Act, Family Educational Rights and Privacy Act, or any other applicable law.

Section 7. Educator and Staff Use

7.1 Approved Products

Within their professional roles, educators and staff may only use Generative AI Products that have been approved by the District in accordance with this Policy and any additional standards applicable by state law.

7.2 District Data

Under no circumstances may educators or staff enter personally identifiable information (PII), including student names, IDs, IEPs, 504 plans, health records, or any confidential information, into unapproved Generative AI Products. The use of personal, non-approved Generative AI Products for school-related work is prohibited when it involves district data or student information.

7.3 Use of Generative AI Products for Instruction

Educators and staff may use approved Generative AI Products to support differentiation planning and instructional planning, provided that they apply their professional expertise to review AI-generated outputs before implementing them and disclose their use.

Educators shall have discretion and autonomy to determine whether to use products that have been approved under this Policy for instruction. Educator and staff review should focus on accuracy, developmental appropriateness, inclusivity, and ethical application.

All instruction shall be delivered by a human being who holds state certification in the academic subject area.

7.4 Use of Generative AI Products for Grading and Assessment

Generative AI Products may not be used for grading or behavior monitoring. Nor may Generative AI Products be used for placement, evaluation, promotion, graduation, and other decisions about students.

7.5 Use of AI Detection Products

AI detection products must not be used as the sole or primary basis for determining academic dishonesty. Output from these tools is not considered definitive evidence of AI misuse and may not be used independently to assign any penalties. The school committee shall ensure that students have meaningful access to appeal procedures in any academic dishonesty determination involving AI detection tools.

7.6 Enforcement Process and Documentation

Parents/caregivers should be notified of any investigations or enforcement actions related to alleged unsanctioned use of Generative AI Products by students and given an opportunity to participate in appeal procedures.

Data on any investigations or enforcement actions related to unsanctioned use of Generative AI Products by students should be documented by the District, including but not limited to the grade level, the demographic data of the student (including gender, race, disability status, and English learner status), and the type of discipline, and this information should be made publicly available to allow researchers to monitor trends.

Section 8. Expiration

This policy shall undergo mandatory review by the School Committee five years after adoption. While it remains in effect, the Committee may amend, readopt, or extend it as appropriate.

AUTHORS

Shaleen Title, Director, Parabola Center for Law and Policy
Brendan Bouffard, Associate General Counsel for AI and Privacy, Fairplay
Joan Donovan, Founder, Critical Internet Studies Institute