RYAN A. BIZZARRO, CHAIRMAN
116 MAIN CAPITOL BUILDING
P.O. BOX 202003
HARRISBURG, PA 17120-2003
(717) 772-2297

HOUSE MAJORITY POLICY COMMITTEE
PAHOUSE.COM/POLICY
POLICY@PAHOUSE.NET
@PADEMPOLICY

HOUSE OF REPRESENTATIVES
COMMONWEALTH *of* PENNSYLVANIA

House Democratic Policy Committee Hearing

Artificial Intelligence + Child Safety

Wednesday, August 5, 2026 | 10:00 a.m.

Representative Chris Pielli

OPENING REMARKS
10:00 a.m. **Rep. Chris Pielli**, D-Chester
Rep. Tarik Khan, D-Philadelphia
Rep. Greg Scott, D-Montgomery

ROUNDTABLE DISCUSSION
10:10 a.m. **Dr. David Christopher,** Superintendent
West Chester Area School District

Melissa Kleiman, Director of Pupil Services
West Chester Area School District

Chris de-Barrena-Sarobe, District Attorney
Chester County

Jess Boyle, Senior Director of Digital Care Transformation
Children's Hospital of Philadelphia

Hannah Diem, Student
East High School, West Chester Area School District

CLOSING REMARKS
11:25 p.m. **Rep. Chris Pielli,** D-Chester

Q & A with Legislators

Pennsylvania House Majority Policy Committee Hearing on Artificial Intelligence and Student Safety

August 5, 2026
10:00 AM West Chester East High School

Prepared by: David E. Christopher, EdD, Superintendent
West Chester Area School District.

Members of the Committee, thank you for the opportunity to co-host today's hearing about the impact of artificial intelligence on Pennsylvania's students and schools.

As a public-school superintendent, I believe artificial intelligence offers tremendous opportunities to enhance learning, improve efficiency, and better prepare students for a rapidly changing workforce. I also believe that each school must consider important safeguards to ensure student safety and to ensure that AI does not disrupt student learning. Additionally, as policymakers consider the role of AI in education, I encourage you to focus on both its potential and safeguards necessary to ensure that students can benefit from these tools safely and equitably.

The first issue I would highlight is **equitable access**.

Throughout the history of public education, technology has often amplified existing inequities before it helped reduce them. Artificial intelligence presents a similar challenge. Students from well-resourced communities may have access to advanced AI tools, subscriptions, devices, and connectivity, while others may have limited or no access at all.

If AI becomes an essential tool for learning, research, writing, and problem solving, then unequal access may create a new educational divide. We should be asking not only whether students have access to AI, but whether all students have access to high-quality AI tools and the instruction necessary to use them responsibly. If Pennsylvania is serious about preparing students for the future, we must ensure that opportunity is not determined by a student's ZIP code.

A second area of concern is **student data privacy**.

Many AI tools rely on large amounts of user-generated information. Students may enter essays, personal reflections, academic work, or other sensitive information into AI platforms without understanding how that information is stored, shared, or used.

School leaders want to ensure that student information remains protected, but the rapid pace of AI development has often outpaced the guidance available to districts. Clear expectations regarding data privacy, vendor accountability, and the use of student information would provide important protections for families while helping school districts make responsible decisions regarding AI adoption.

A third issue involves the growing threat of **AI-generated deepfakes and synthetic media**.

Artificial intelligence now makes it possible to create realistic images, audio recordings, and videos that never actually occurred. While these technologies can be used creatively and productively, they can also be used to impersonate students and staff, harass peers, damage reputations, and undermine trust.

School districts across the country are increasingly confronting concerns related to cyberbullying, misinformation, and the misuse of synthetic media. Existing policies were largely written before these technologies became widely accessible. As policymakers consider student safety, this is an area where thoughtful guidance could help schools respond effectively while protecting students from emerging forms of digital harm.

Finally, I would encourage the Committee to think about the long-term impact of AI on **student well-being and human relationships**.

At present, most school districts are not witnessing widespread concerns regarding students forming unhealthy attachments to AI systems. However, as AI tools become increasingly personalized, conversational, and human-like, it is appropriate to monitor their potential impact on young people.

Schools play a critical role in helping students build relationships with peers, teachers, counselors, coaches, and other trusted adults. We should be cautious about a future in

which technology begins to substitute for, rather than supplement, those important human connections. While we should embrace innovation, we should also ensure that AI remains a tool that supports human development rather than replacing it.

In closing, I believe the greatest challenge facing Pennsylvania schools is not whether artificial intelligence will become part of students' lives. It already is. The challenge is ensuring that AI is viewed as a tool and introduced responsibly, equitably, and safely. I believe it is also critical for policymakers to ensure that any policies they adopt allow educators and staff to continue using AI to support their work and to personalize learning for students.

By focusing on equitable access, strong privacy protections, safeguards against emerging harms such as deepfakes, and a continued commitment to healthy human development, Pennsylvania can lead with a balanced approach that embraces innovation while placing the needs of students first. The role of policymakers should be to establish clear guardrails that protect students while preserving the flexibility of local educators to innovate and meet the needs of their communities.

Thank you for the opportunity to share these thoughts, and I would be happy to answer any questions.

Hannah Diem Testimony

I am Hannah Diem, an incoming senior at West Chester East High School. I am honored to be given the chance to share my thoughts on the impacts of artificial intelligence and technology on children from the perspective of a student. When I was first given this opportunity, I began to consider what encapsulates my concerns the best; what are the things that I care the most about? I created a list of these things and I will take time to share my experience surrounding each item on the list. It is important to note that I do not claim to be an expert on these things, but I have done some research surrounding these topics and throughout my real-world exposure I have become increasingly aware of the presence of AI. I want my testimony to reflect my firsthand knowledge, because I believe the lived experiences of students are missing from the research and statistics surrounding the successes and pitfalls of AI. I cannot speak for all students, but I know a lot of people I am familiar with feel a similar way as I do.

Education

As most are aware of, AI has had severe impacts on education. AI was first introduced when I was in middle school. With little formal education on the topic, students quickly began to find different chatbots and websites to assist them with their assignments. The most common and direct form of AI could be found on Snapchat. MyAI, which automatically pinned itself to the top of the feed, looked identical to every other profile and was easily accessible for Snapchat users. It was nearly like the chatbot took the form of a real person. There were various harms to this, but as it relates to education, the answers were placed in the palm of our hands. Despite our teachers' best efforts, it was very difficult to keep this resource from being utilized.

I have grown up at a time in education where no one understands where to go next. Over the past three to four years, AI has been in flux. At times, teachers treat AI as the plague and are appalled at the idea of integrating this technology into our education system. But at other times, they encourage its integration and utilize it themselves. In truth, this has shown how little training and guidance they had on the technology. But no matter how much training they got on the topic, AI had already been released to students. AI was a cancer that took hold in education; treatment was not given soon enough and it spread to irreversible bounds. Everyone understood this to be true, yet no one has been willing to find a solution.

Midway through my sophomore year, the school district released online asynchronous AI training which everyone was required to complete. These trainings were meant to educate on the ethical use of AI and the capabilities that AI has. But instead, these stale and underresearched trainings were simply information that many students had already known. Ethical use of AI is not something that students do not understand, it is something that they do not care to apply. As I previously mentioned, we had experience with AI for two years prior at this point; if someone was not using it ethically that is not something that they were likely to change. Over the course of the remainder of the year and the entirety of my junior year, AI was not addressed by the district again.

An Output-Based Education

All of this is to say that AI has already been released into education and we cannot turn back time. Instead, I believe that we must focus on changing how we approach education as a whole. At the

moment, education is incredibly output based. Colleges, scholarships and awards are all chalked up to grades and academic achievement. Grades are on the basis of the quality of the work completed, and this plays a role in AI use. If students are given the capability to use AI for an assignment, and they know they will get a better grade if they do so, what is incentivizing them not to? Everyone can achieve the same output with modern technology, so in this case it matters more what students are putting into their work. But instead of rewarding the students who choose to put in the work, many teachers continuously miss or do not penalize the use of AI. This is not every teacher, but there are many of them. Those who do choose to look for AI, it is often biased and they are seemingly untrained to look for it. I have heard many cases in which teachers miss AI use or they penalize someone who did not use AI but had a similar writing style. I am not criticizing the teachers who are trying, but this all goes back to the idea that the output is still the center of their grading system. I believe that education needs to evolve to fit the new world we are living in. AI needs to be integrated but students need to be put in situations where the grade is not dependent on the final product they put forward. In tandem, education should foster the development of soft skills such as teamwork, practical analysis and presentation because the ability to do these things will likely take precedence over the ability to complete tasks. As a whole, this AI revolution is no different from the industrial revolution; in the same way that education evolved to fit those technologies, education must evolve to fit the technologies of the modern world.

The Modern Job Market

In a similar manner, it is uncertain what the job market will look like by the time we are ready to enter the workforce. There are many factors that play a role in this, but when choosing a career to go into I know it is in the forefront of many teens' minds. Throughout high school and college, students learn skills in an effort to become more competitive in the workforce. But when they don't know what they are preparing for, how does any of that education feel valuable? Specifically with college, there is a lot of anxiety surrounding what major to pick to ensure the most job security. College price tags are increasing, students understand they could experience years of debt by getting a college degree, but many don't even know if that will be worthwhile. This increased concern is a common experience, and it is something that changed as soon as AI was introduced and automation began to take hold in many industries.

The Literacy Crisis

The literacy crisis is a hot topic in many educational discussions, but it is worth mentioning in the context of the impacts of AI. The literacy crisis does not only encompass vocabulary or reading comprehension, but it relates to the ability to analyze. Whether it is a book for English class or discussing a movie with friends, in depth analysis requires one to go beyond the surface level ideas and look thematically or thoroughly. This way of thinking can be applied to relationships with others, such as an argument with friends or trying to reach a common understanding of a situation. Many analysis skills are developed in school, they often begin with the ability to analyze written works. The literacy crisis began before AI was introduced into education, it is not a new phenomenon. Because of this existing issue, it was exacerbated by the use of AI in English circumstances. AI has a phenomenal ability to analyze, at a much deeper level than many students believe they would be able to. But the impacts of using AI for an essay or to summarize a book goes beyond what the impacts would be in a math or science class solely because it slows the development of analysis skills. The impacts of this will go beyond the workplace, they will stretch throughout different aspects of our modern society.

Digital Literacy

One place that responses to AI have fallen short is the lack of digital literacy education. Children have been raised in an environment where they are surrounded by YouTube videos, phones and constant stimulation. Because they have grown up with these technologies, many educators seem to be under the impression that they understand how to navigate this media-based culture. There is minimal direct digital literacy education as a result of this conclusion. But if children were not taught how to consume digital media properly or use AI ethically, how can they be expected to develop that skill themselves? The amount of fake news, AI generated images or false information on the internet is staggering, but many still do not know how to discern what is real.

Until digital literacy is integrated into education, the impacts of this will continue to grow. As AI continues to improve, it becomes increasingly difficult to tell what to believe. This leads to skepticism and anxiety surrounding the quality of any information. Digital literacy is as necessary as math or reading in the modern world, as AI continues to spread on platforms that allow it to do so. Digital literacy needs to be included in education as young as elementary school, to build a foundation of critical thinking skills to take on these real world situations. In order to protect children's safety and create more informed citizens and individuals, digital skills are vital.

Chatbots and Deepfakes

Chatbots and deepfakes are some of the most concerning and dangerous aspects of AI, and they have had the most tragic consequences. Chatbots are AI bots which offer emotional support to the users who interact with them. They often offer advice or responses to their users as a way of interacting with them. The horror stories which revolve around emotional reliance on a fake, digital personality are terrifying. Cases of suicide or collapsing mental health due to these chatbots have spread in the news and exemplify the consequences that they can have. I cannot testify on these heavily because I have never used one, nor do I know anyone who has. What I can say is that these chatbots should not be allowed to exist in any capacity. AI should not be used in an emotional capacity, it is inherently an emotionless being. Regulations must be put in place in order to criminalize these actions, and that is slowly beginning to happen.

Deepfakes are AI generated images which can take an image of someone and create another fake image of them. In the worst cases, these are used in a way that can severely impact a person. Deepfakes can be used to create nudes or other emotionally damaging images. If these images are posted online, they could spread quickly around a school or the internet. These are incredibly damaging and harmful, but unfortunately it is not very difficult to do so. These have the ability to ruin people's lives and it is terrifying to realize that it could potentially happen to anyone.

AI can also be used to alter an image and many people utilize AI to alter the images they post on social media. This can have significant negative effects when it comes to body standards or expectations. When an image is edited to make an influencer look skinnier or their muscles look bigger, the person viewing that image would likely wonder why they cannot look the same. Little do they know that the image itself is not even real, that body is not real. This is incredibly damaging for those who struggle with

their appearance and it is very difficult to discern an edited image from an unedited one. Unrealistic standards are damaging for mental health and impact the confidence many feel about themselves. It is difficult to think of a way to end this editing, because it happens so seamlessly, but the repercussions of these unrealistic standards can be seen in the increasing use of weight loss drugs and the rise in eating disorders. For the rising generation, I worry how these standards will progress and the impacts it will have on their development.

Pennsylvania House Democratic Policy Committee
Children's Health and the Regulation of Artificial Intelligence
Prepared by Jess Boyle, Senior Director of Digital Care Transformation and Digital Health at Children's Hospital of Philadelphia
East High School, West Chester, PA
August 5, 2026

Good Morning, and thank you to the Pennsylvania House Democratic Policy Committee and to everyone gathered here at East High School today,

I am Jess Boyle, Senior Director of Digital Care Transformation and Digital Health at Children's Hospital of Philadelphia. Thank you for this opportunity to speak in support of meaningful, children's health-focused regulation of artificial intelligence.

At CHOP, my team's work sits at the intersection of digital care, clinical AI, and consumer health informatics. I lead programs that make children healthier with technology, and I believe deeply in what these tools can do.

We work with two very different kinds of artificial intelligence, and that difference is what I most want to leave with you today. The AI a health system like CHOP places near a child is narrow and purpose-focused: validated for a specific clinical task and pediatric population, run under clinical and operational oversight, analytical rigor, and used to support a trained professional accountable for that child's care. Consumer AI is nearly the opposite: general-purpose, optimized to maximize engagement, released to everyone including young children, and very broadly accountable for the outcome.

That difference is the product of governance. Before any AI tool is used in patient or clinical care at CHOP, it passes through rigorous review: whether it is appropriate for the use at all, whether it performs against our own pediatric population, whether it is safe and unbiased, and whether it protects patient privacy. Once it is live, we do not walk away, we monitor for accuracy, drift, and unintended effects, and keep a qualified clinician accountable for every decision. We ask not just whether we can use a tool, but whether we should. That is the standard any AI acting on a child's health should meet, and exactly the standard consumer AI reaching our children unfortunately does not.

Let me turn to what is happening outside of health system guardrails. In a nationally representative survey of nearly 3,500 American teenagers, more than six in ten reported using a conversational AI chatbot, and one in nine use one nearly every day. This is not homework help. Six in ten cited friendship and companionship as a reason, because these products are engineered to supply precisely that. The American Academy of Pediatrics warns that these companions use engagement-prolonging design to meet a child's emotional needs while posing an unregulated risk to their well-being. Children are forming attachments to systems built to hold their attention, not to protect their development.

And the harm is not rare. In that same survey, nearly half of teen users reported at least one concrete harmful experience. Roughly one in three were asked for uncomfortable personal information. Nearly one in four felt manipulated. One in seven received self-harm content, and one in eight received messages related to suicide. The youngest surveryed, thirteen-year-olds, reported the highest rates of nearly every category of harm. These are population-level exposures affecting millions of children.

Children's Hospital of Philadelphia℠
3401 Civic Center Blvd. • Philadelphia, PA 19104 • 215-590-1000 • chop.edu

And children are differentially vulnerable by design: in adolescence, the brain's reward circuitry is heightened while self-control is still maturing: the exact profile that engagement-maximizing design can exploit. Children are not small adults, and a product that may be safe for a forty-year-old is not, for that reason, safe for a thirteen-year-old.

The consumer safeguards we would hope protect them are currently failing where it matters most. When researchers tested leading chatbots across thousands of suicide-related questions, the systems reliably refused the most extreme and answered the most benign, but could not distinguish the intermediate band of risk, exactly where many at-risk adolescents sit.

The same holds for health information. In an audit of five major chatbots across 250 health questions, about half of the responses were problematic, and the systems declined to answer fewer than one percent of the time. They delivered wrong answers with complete confidence, at a college reading level, and could not produce an accurate list of sources. A child bringing a serious health question to one of these systems is likely to get an answer that sounds authoritative but may be dangerously wrong.

I want to be fair about the benefits, because there are some. The strongest evidence of benefit in young people shows a modest, short-term improvement in depressive symptoms in mild cases, alongside no measurable effect on anxiety, stress, or overall well-being. This is unfortunately narrow. It does not justify deploying these systems to minors at population scale, with few of the protections we require of anything else that affects a child's health.

So what am I asking of you? Not prohibition, but proportionate and measurable accountability, much of which already carries strong professional consensus. Safety and privacy on by default for any minor, with no targeted advertising to children. Platforms required to identify their child users and design for them, with dedicated child-safety teams, algorithmic transparency, and a clear prohibition of harmful content. Because these systems are probabilistic and cannot be fully tested before release, one-time approval is not enough; we need continuous post-market surveillance, adverse-event reporting, and periodic re-authorization. This is the clinical-grade accountability we demand of anything that acts on a child's health. Children deserve pediatric-specific standards: developmentally appropriate testing, explicit labeling of what is and is not suitable for a child, and data-protection frameworks built for minors.

Consumer AI is already a companion, a counselor, and a health advisor to millions of children, carrying very few of the safeguards we would demand of any other product in that role, safeguards we apply every day for the children in our care. We already know how to close that gap…we do it at CHOP everyday.

I stand ready to offer CHOP's assistance on any policy or technical questions as you take up this work. Thank you so much for your time today, and thank you for advocating for the health and safety of our children.

Children's Hospital of Philadelphia℠
3401 Civic Center Blvd. • Philadelphia, PA 19104 • 215-590-1000 • chop.edu