


PARENTS



STOP THE CARNAGE

— UNDERSTANDING • PROTECTING • TALKING —

PAUL NEUMAN

PARENTS: STOP THE CARNAGE

**UNDERSTANDING AI, PROTECTING YOUR CHILDREN,
TALKING ABOUT IT AS A FAMILY**

PAUL NEUMAN

© 2026 Paul Neuman

Published by IT Services Demystified Editions

All rights reserved.

No part of this publication may be reproduced, adapted, translated or distributed, in whole or in part, by any means, without the prior written permission of the author or publisher, except as permitted by applicable copyright law.

This book provides general guidance for parents and other adults supporting children and teenagers. It is not a substitute for medical, psychological, educational or legal advice tailored to an individual situation.

IN THE EVENT OF IMMEDIATE DANGER OR A RISK TO A CHILD'S SAFETY, CONTACT THE APPROPRIATE PROFESSIONALS OR RELEVANT SERVICES WITHOUT DELAY.

Laws, platforms and digital services may change after publication. For any action you may need to take, consult current official sources in your country of residence. Further resources and updates are available online; details are provided at the end of the book.

First edition, October 2026

ISBN: 979-8-1741-6964-7

Table of Contents

Introduction 1

How to Use This Guide 5

Chapter 1 – But First, What Exactly Is AI? 7

Chapter 2 – When AI Does the Homework.....13

Chapter 3 – When AI Becomes a Confidant.....21

Chapter 4 – Photos, Digital Traces & Privacy.....29

Chapter 5 – When an Image Is No Longer Proof.....37

Chapter 6 – Who Decides What Comes Next?..... 45

Chapter 7 – When AI Imitates Someone 53

Chapter 8 – What the Law Won’t Do.....59

Chapter 9 – The Family AI Agreement 65

Conclusion 75

AI Survival Dictionary 77

Resources and Updates.....83

Acknowledgements..... 85

About the Author..... 87



Introduction

Artificial intelligence did not knock on your door before entering your child's life.

One evening, it may be behind a surprisingly good school assignment. The sentences are flawless and the reasoning looks perfectly tidy. But when you ask your child to explain what they wrote, everything suddenly becomes less clear.

Another day, an image circulates in the class group chat and shows a scene that never happened. The image is fake, but the laughter, shame or fear it triggers are all very real.

Then your child tells you they sometimes talk to an AI. It answers quickly, does not make fun of them and always seems to be available. Perhaps, they tell you, it even understands them better than adults do.

None of this really looks like a technological revolution. These are now ordinary scenes of family life, and that is precisely how AI gradually finds its way in: through a school assignment, a video, an image, a search or a conversation held behind a closed door...

The title *Parents: Stop the Carnage* does not cast artificial intelligence as an enemy to be driven out of the home, nor does it suggest that parents should have to solve all the societal, environmental or economic consequences of AI.

It is simply a call to act where they can have a real influence: in their own family.

At heart, this book is about the place AI takes, or is allowed to take, in our children's lives, because difficulties arise when its use develops faster than our ability to set boundaries, exercise judgement and have the conversations needed to guide it.

And this is where the problems begin:

- help replacing the effort needed to learn;

- a convincing likeness passing itself off as proof;
- the ease of a click on a screen standing in for informed consent;
- a recommendation that, through repetition, starts to feel like a conscious choice;
- an always-available machine gradually taking the place of someone who can listen and act;
- above all, your child being left alone with what they have seen, shared, believed or confided, with no one to talk to about it.

Children often discover these tools before their parents do.

They know where to click, which service to use and how to quickly find what they are looking for, and that technical head start can easily make adults feel they have arrived too late.

Some respond by banning AI altogether, or at least some uses of it. But a ban your child does not understand will not necessarily stop them using it; above all, it may simply drive that use out of sight.

Others end up thinking: “They know more about this than I do. I have nothing to teach them.”

But that is wrong.

Knowing where to click does not mean knowing where to draw the line. Your child may know the tool better than you do without necessarily being better equipped to judge what is true, safe, private or acceptable.

The essential parenting skill is to stay the adult in the room, even when you do not master every technology.

It also means knowing how to set boundaries. Some can be discussed with your child; others cannot, particularly when it comes to protecting them.

Sometimes, you need to slow things down when the tool pushes for speed, protect what is private and remember that a well-written answer or a convincing image can be false.

A machine that responds gently is not, by doing so, caring for your child.



None of this requires monitoring every screen or reading every conversation. It is mainly about staying present and knowing when the situation calls for you to step in.

AI can be in the room all the time without becoming the “adult”.

It can explain things differently, translate, correct, stimulate creativity or make some forms of learning more accessible. Your child can also use it to look for information, prepare a question or put words to what they are feeling.

But depending on how it is used, the same tool can write for them, make things up without ever sounding uncertain, retain personal information, imitate someone or make them feel understood.

Asking whether “AI” is good or bad is too vague a question to be genuinely useful.

Ultimately, what matters is exactly what your child shares with AI, what they get back from it and the place it eventually comes to occupy.

You do not guide a nine-year-old, a thirteen-year-old teenager and a seventeen-year-old on the verge of adulthood in the same way. As your child grows, they should of course gain more autonomy, but that does not mean you become a spectator.

The need for dialogue and protection does not disappear. What changes is how you set boundaries and support your child.

This guide grew out of the situations in which AI stops being an abstract subject and enters family life: homework, confidences, digital traces, fabricated images, recommendations and impersonation.

And while some of these issues are clearly not new, particularly social media, digital traces, recommendations or impersonation, AI is now woven into them: it ranks, recommends, generates, transforms and imitates.

Applications will change their names, appearance and promises, and some will simply have disappeared before you have even got to grips with their latest menu, so there is little point chasing them.

Three habits, however, should outlast all these applications:

- understand what the tool actually does and check what it produces;
- protect what should remain personal;
- talk before a problem takes hold.

You will not always know what to say, and that is fine because being a parent in the age of AI is not a test of your knowledge of new technology.

You can ask a question, check something with your child, set a boundary and, above all, make sure they are not left alone with a machine when the situation becomes important or worrying.

This book will help you recognise when AI is genuinely helping your child, when it starts to take up too much space and when you need to step back in.

Knowing when to intervene matters more than knowing everything about AI.



Chapter 2

When AI Does the Homework

“Help me to do it myself.”

Phrase associated with Montessori education

Impeccable homework. Well-constructed sentences. A clear structure. Vocabulary that suddenly sounds more “grown-up”. Yet your child cannot even explain the second sentence.

The carnage does not always look like a disaster. Sometimes, it even gets full marks.

A good grade can now hide learning that never actually happened.

Your child may not have asked AI to do everything for them. Perhaps they started by asking for an explanation, then a structure, a rewording and finally a conclusion. Layer by layer, the assignment becomes lasagne homework and, by the end, it carries their name but no longer quite belongs to them.

Asking your child “Did you use AI?” is therefore not enough, because the answer may be “no”, “just a little” or “everyone does it”. Your child may even be completely sincere because, if AI is involved at every stage of the work, they may eventually no longer know very clearly themselves where the help stopped and the replacement began.

The useful question needs to be more precise: “What did you ask AI to do, and what did you do yourself?”

That question asks your child to describe how they worked and makes clear what still belongs to them.

Generative AI can be a very good research assistant because it can help your child formulate a question, find keywords, compare different explanations or identify what they have not yet understood.

But it can just as easily invent a source, attribute a quotation to the wrong person or present something uncertain as fact. And all in polished prose, just to make the mistake look more convincing. Facts, figures, quotations and references therefore need to be identifiable in trusted sources.

The real trap lies elsewhere, however, because homework is not designed to produce text but to produce learning.

If an exercise asks your child to search, select, verify, reason or develop a structure, handing those operations over to AI removes precisely what your child was supposed to practise. So even if the homework progresses, the learning is likely to stall.

Correcting spelling may be acceptable in a history assignment, but in a spelling exercise, it amounts to letting AI sit the test. Similarly, asking AI to suggest a structure may help with revision, but if the assignment is about building an argument, then part of the essential work has already been delegated.

The useful measure is therefore not how much text AI produced, but how much of the intellectual work your child actually did.

To move beyond the same old debate about “help” versus “cheating”, it is better to return to two very concrete questions: what was the exercise intended to teach your child, and what did they actually do?



USAGE

GUIDANCE

CAN HELP

Explaining a concept, suggesting research directions or keywords, asking questions, giving an example or pointing out mistakes after your child has made a first attempt.

Any information or sources it suggests must still be checked.

AI lights the path, but your child is the one who walks it.

NEEDS BOUNDARIES

Suggesting a structure, summarising a document or rewording a passage.

The tool is already touching the structure of the work. Your child therefore needs to check the result, make their own choices, rewrite the content in their own words and follow the teacher's instructions.

Otherwise, the help provided eventually starts to replace the exercise itself.

TAKES OVER

Building the reasoning, selecting the arguments, writing the answers or producing the entire assignment instead of your child.

And at this point, it is no longer assistance.

The homework then has your child's name on it, but not their reasoning.

AI has also taken over when your child uses an answer they cannot explain or defend.

This framework is not a parallel set of school rules, and the teacher's instructions remain the priority. If AI use is prohibited, hiding it certainly does not turn breaking the rules into educational innovation.

But even when it is allowed, its use can cause more subtle forms of harm.

AI responds quickly, comes up with a structure, examples and often sentences more polished than your child's first attempts. Over time, your child may start asking the machine before even trying, convinced that their own words will never be good enough.

The risk then goes beyond an assignment that no longer really belongs to them, because your child may lose the habit of starting anything without AI.

Learning also means spending some time alone with a question that has no immediate answer. Searching, hesitating, getting something wrong, correcting it and trying again are not failures, but an integral part of the process.

An imperfect assignment your child understands is worth more than a perfect one written by a chatbot.



TRY THIS TOGETHER

Ask your child to explain their assignment to you for three minutes without reading from it.

They should be able to:

- summarise the main idea or the topic covered;
- explain their reasoning;
- say what they checked and how;
- acknowledge what they are unsure about.

The three-minute test is not at all designed to detect AI use, but it checks something even more useful for your child: have they genuinely made the work they are about to submit their own?

If they understand the ideas, can defend their choices, explain what they checked and recognise what they are unsure about, AI has been useful by supporting their work without replacing it.

If your child seems to be discovering their own assignment as they read it back, something has gone wrong and the tool has probably taken up



too much space. Or perhaps the work was simply not understood. In either case, it is better to go back through the reasoning together than to put the machine on trial.

Banning AI from all homework may sound like a brilliant idea, but it is more likely to push its use out of your sight than make it disappear.

Conversely, accepting everything in the name of progress means confusing technical sophistication with educational quality. An answer produced in two seconds is unlikely, in the eyes of teachers, to count as educational progress.

The question to ask then is: “At what point did AI become involved, and what did you use it for?”

The logic would therefore be to proceed as follows:

1. your child tries first;
2. they ask for help with a specific point;
3. they check the response they receive;
4. they make their own choices and put the result into their own words;
5. they are able to explain and defend their work.

Your child can also get help without asking for the answer. For example: “Ask me questions to help me work it out for myself.” or “Explain where I’m getting stuck without solving the exercise for me.”

They can then work on the content of the exercise while learning to use AI without handing over their reasoning.

But if AI steps in before the very first attempt and directly provides a result that is then reused without any checking, all the steps that mattered have just been skipped.

The teacher’s instructions remain the number one priority. When in doubt, your child should always ask what is allowed and clearly explain how AI helped them. Asking AI afterwards to disguise that help, for example with a prompt such as “Add a few mistakes so the teacher doesn’t notice”, is prompt gaming.

The way you support your child changes with age. A younger child will often need to use the tool with an adult. With a teenager, supervision can decrease, but responsibility never does.

Avoid opening with: “Did you use AI for your homework?” It is a closed question that mainly invites a “yes”, a “no”, or an impressive display of verbal gymnastics to avoid answering. What’s more, it tells you almost nothing about the work that was actually done.

Ask your child to talk you through how they approached the assignment instead. You can simply add: “What was the hardest part for you? I’m not trying to catch you out. I want to understand what you did, what AI did and what you can explain today.”

You can also reaffirm your support by reminding them: “I would rather see an imperfect assignment that you genuinely understand than a perfect one you cannot explain.”

This discussion should make clear how much space the tool has taken up, without turning into a police interrogation.

These phrases can open the conversation, while the family rule sets the line that must not be crossed.



FAMILY RULE

Your child tries first without AI.

AI can then help them search, understand or check.

It does not do for them what the assignment is meant to teach them to do.

School is not meant to produce immaculate homework, but to teach children how to search, understand, write, reason and explain.

AI can be useful when it helps your child practise those skills. But as soon as it does that work in your child’s place, it is no longer really helping them learn; it is mainly giving them the illusion of success.



Your role is to stay the adult in the room by looking at what your child has actually learned, without tracking every prompt or being impressed by the perfection of the assignment.



Chapter 5

When an Image Is No Longer Proof

“This is not a pipe.”

René Magritte, *The Treachery of Images*, 1929

An image appears in a class group chat.

In it, a student is recognisable in a ridiculous, shocking or humiliating situation. Some people laugh. Others comment. Someone is already forwarding the image to another group.

Nobody yet knows where it came from.

The image may show something that really happened. But it may also have been cropped, altered or accompanied by a false explanation. Or it may have been created entirely by AI.

None of that stops the laughter, the accusations or the sharing.

The image may be fake. Its effects are already real.

Image generators can nevertheless be remarkable creative tools. Your child can illustrate a story, visualise a historical place, imagine a character, experiment with a visual style or give form to an idea they could not draw themselves.

The problem begins when a creation is presented as evidence, or when it depicts a real person in a situation they never experienced in order to deceive, humiliate or harm them.

For a long time, a film photograph could seem to offer simple proof. Someone was there, the camera captured the scene, so it must have happened.

That reasoning was never entirely reliable, since a photograph could be staged, cropped, altered or simply accompanied by a misleading caption.

But with AI, the problem changes scale because a scene no longer even needs to have existed in front of a camera. An imaginary place, a fabricated event or a person in an invented situation can now be created in moments.

I have a particularly clear sense of this change in scale because, in the mid-1990s, at engineering school, I set up a student television association called EPTv. Our speciality was producing news bulletins at events organised by other grandes écoles, whether sports competitions or parties. For each event we covered, we would adapt our logo using computer-generated imagery. Back then, a few seconds of animation could take several hours of processing. Today, the same kind of animation can be produced in seconds.

An image can therefore be fakethentic, convincing enough to look authentic without actually being so. It can just as easily be authentic while telling a completely false story.



IN PLAIN TERMS

An image can be:

- authentic and accurately presented, if it really shows the scene, at the stated place and time;
- authentic but taken out of context, if the photograph is real but the date, location or explanation is false;
- altered, if elements have been added, removed or changed;
- entirely generated, if the scene never existed in front of a camera.

A fabricated image can mislead, but a genuine photograph accompanied by a false story can do just as much.



Checking an image therefore means examining what it shows, where it came from and what is being claimed about it.

We have heard plenty about six-fingered hands, strange eyes and impossible shadows. That used to be reassuring because all you had to do was spot a flaw.

That reflex is no longer enough.

Today, a fabricated image may contain no obvious anomaly at all, while a perfectly genuine photograph may look strange because of movement, a reflection, poor lighting or simply lower quality.

Even spotting a suspicious detail may raise doubts, but it certainly should not be enough to reach a conclusion.

A few basic questions can already make things clearer:

- Who published the image?
- What claim accompanies it?
- Can you find an earlier post or the original version?
- Can the date, location and context be confirmed?
- Do other reliable sources show the same scene?
- If a real person is involved, can the information be checked with them or with a source directly connected to what happened?

Verification therefore begins by tracing the image back, certainly not by becoming an expert in pixels.

Above all, resist the urge to react immediately to a shocking, funny or outrageous image, because responding, accusing or sharing before even knowing what it really shows is bound to be a bad idea.

The stronger the emotion, the further your hand should move away from the “forward” button.

Adding “Is this real?” will never turn sharing into verification, nor will it absolve your child if they help circulate the image, potentially making the humiliation of the person shown even worse.

Before even asking whether the image is authentic, one question matters: “Do we really need to circulate it?”



WHAT TO WATCH FOR

Stop before sharing when an image:

- immediately provokes fear, anger or outrage;
- accuses someone without identifying a clear and verifiable source;
- ridicules or humiliates a person;
- shows an intimate or sexualised situation;
- explicitly urges people to share it quickly;
- comes from an unknown account or one impersonating a familiar source;
- has no clear origin or verifiable context.

A strong emotional reaction is not evidence. It is one more reason to slow down.

We actually need to separate two problems: on the one hand, whether the image is authentic and the context in which it is presented, and on the other, what it may do to the person concerned if their image has been used or altered without their permission to deceive, humiliate or harm them.

However, for some images, tracing the source becomes almost secondary, particularly when a real person’s face or body has been used to create fake nudity, a sexualised scene or any other humiliating situation.

Even if the image is fake, the humiliation is not.

Saying “it isn’t real” is not enough when the person targeted is suffering, receiving threats or seeing the image circulate among the people around them.



So if your child receives this kind of content, it is essential that they know what to do:

- do not forward or share it, and do not make another copy;
- do not confront the person threatening or pressuring them on their own;
- note any useful information, such as the sender's account, a link, a date, a time or an associated message;
- tell a trusted adult immediately, whether or not they are a family member;
- ask for help to report the situation and find out how to preserve any necessary evidence.

Your child should not have to face this situation alone.

Your first reaction will be decisive because, if your child is afraid first and foremost of being punished after receiving, viewing or even forwarding the image, there is a strong chance that the problem will remain hidden and any intervention will be delayed.

They may have made a mistake, but that can always be discussed afterwards.

The immediate priority is to stop the circulation, protect the person targeted and get help.

With a younger child, you will generally need to take charge of the necessary steps yourself. With a teenager, involve them in decisions and respect their privacy as far as possible, without leaving them to manage a humiliating, intimate or threatening situation on their own.



WHAT TO SAY TO YOUR CHILD

“A convincing image isn’t necessarily real. Before believing what you see, look carefully for its source and context.”

“Adding ‘Is this real?’ doesn’t justify forwarding an image that humiliates someone.”

“An intimate or sexualised image can be fabricated and still cause real harm. If you receive one, don’t share it. Tell an adult immediately.”

Your child can invent a character, illustrate a poem or depict a scene that could never be photographed, but that creation becomes deceptive when it is presented as a photograph of a real event.

What distinguishes creation from deception lies solely in how the image is presented and used, not in its pixels.

Some platforms add a label when an image appears to have been generated or substantially altered by AI. Although useful, that indication remains limited because it may be absent, lost when the image is shared or fail to tell the whole story.

The absence of a label proves nothing.

And when a real person is depicted, transparency is no longer enough, because writing “AI-generated” underneath a humiliating fake does not make it acceptable.

Technology may make the act much easier, but it gives us no additional rights over someone else’s image or body.



TRY THIS TOGETHER

Choose a non-sensitive image that claims to show a real fact or event.

Ask yourselves:

- What story is this image telling?
- Who published it first?
- Can we find an original or earlier version?
- Can the date, location and context be confirmed?
- Does another reliable source describe the same event?
- What could happen if we shared it and the story turned out to be false?

There is no need to hunt for the sixth finger, because the point of the exercise here is to trace the image back to its source, check the story it tells and consider the consequences of sharing it.

To stay the adult in the room here means resisting the three-second verdict: “It’s AI.”

When someone is being humiliated, the priority is to stop the circulation and protect the person targeted, then look for reliable information before reaching a conclusion.

Faced with a humiliating image, dignity comes before curiosity.

Faced with an image that makes an accusation, verification comes before reaction.

Faced with a creation, the pleasure of imagining does not remove the need for transparency, respect or consent.

**FAMILY RULE**

A creation must be presented as such.

An image that humiliates someone is never shared.

An image that accuses or shocks is always checked before we believe it, comment on it or forward it.

We do not place a real person's face or body into a generated or altered image without their permission.

Image generators give your child remarkable creative power. That power does not entitle them to use another person's body, identity or reputation, or to present an invented image as proof of an event that never happened.

An image can be true, false or misleading.

Whatever the case, sharing it should remain a choice.



AI Survival Dictionary

This dictionary brings together a few useful terms and memorable expressions for the main ideas and habits in this guide. It gives families simple language for recognising a situation and talking about it, without pretending to cover the entire vocabulary of artificial intelligence.

ADULT IN THE ROOM, *n. phr.*: Someone who may not know the latest tool, but can still ask a question, slow things down, check, protect or say no. The adult in the room may be clumsy, tired or technologically behind. But they can notice what is not being said, act when something happens and take responsibility for the consequences of their decisions.

AI, *abbr.*: Artificial intelligence. Two letters that make it sound like a single technology, when they actually cover very different systems: recognising, classifying, recommending, estimating, generating, conversing or chaining actions together. Saying “I used AI” therefore explains almost nothing. And sometimes, you may wonder whether you are dealing with Artificial Intelligence or something Artificially Intelligent: convincing enough to sound clever, not reliable enough to stop checking.

ALLOWED DOESN'T MEAN SUITABLE, *phr.*: A reminder that a tool may be available, comply with applicable rules or be accessible from the stated minimum age without necessarily being suitable for your child's maturity, intended use or the information they may share with it. The “Download” button is not a parental seal of approval.

BIAS, *n.*: A tendency or imbalance that may come from the data used to train a system, the way those data were selected, or choices made during its design. Bias can reproduce stereotypes, make some perspectives more visible than others or disadvantage certain people. A fluent answer is not necessarily a neutral one.



About the Author



EPITA graduate and INSEAD alumnus, Paul Neuman has spent nearly thirty years working in technology and services. He has supported several major transformations across the sector, with the same concern throughout: understanding what is genuinely changing behind the rhetoric and helping organisations make the most of it without losing sight of operational realities.

His approach starts with practical use: explaining complex subjects simply and distinguishing between what tools can actually do and what people assume they can do.

He is the author of several books, articles and white papers, including *Artificial Intelligence: Governing Through Clarity*. He advocates an approach to artificial intelligence based neither on blind fascination nor systematic rejection. Technology can be useful, creative and sometimes impressive. That does not remove the need to understand what it does or to decide what place we want to give it.

With *Parents: Stop the Carnage*, he writes for parents, teachers and other adults supporting children who are already encountering AI in their homework, images, videos, searches and conversations.

He offers them simple guidance to understand what these tools do and know when to step in, while keeping the dialogue open without trying to become specialists.

Children will sometimes discover new technologies before adults do. That does not diminish the role of those guiding them.

**A parent does not need to master everything to
remain the adult in the room.**

AI IS ALREADY IN YOUR CHILD'S LIFE.

**The real issue is not banning everything.
It is knowing when to step in.**

Flawless homework, a chatbot as confidant, fabricated images, endless recommendations, personal data shared too quickly, voices or faces imitated...

AI can help, explain and create.

It can also replace effort, blur the line between true and false, expose what should remain private, or take up a place that no one has consciously decided to give it.

This guide gives you practical reference points to take back control without becoming an expert in artificial intelligence.

UNDERSTAND

what the tool is actually
doing.

PROTECT

what should remain
personal.

TALK

before a problem
takes hold.

**You do not need to know everything about AI to
stay the adult in the room.**

Paul Neuman is an engineer by training and an author. For nearly thirty years, he has observed and supported technological change and the ways people use it. After *AI – Governing Through Clarity*, he turns here to what artificial intelligence is concretely changing in family life.