

The AI Guidance Gap: Teaching, Learning, and Trust When Answers Are Everywhere

The AI Guidance Gap: Teaching, Learning, and Trust When Answers Are Everywhere

EVELYN CHAN

TORONTO, ONTARIO



The AI Guidance Gap: Teaching, Learning, and Trust When Answers Are Everywhere Copyright © 2026 by Evelyn Chan is licensed under a Creative Commons Attribution 4.0 International License, except where otherwise noted.

© 2026 Evelyn Chan. This work is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

Contents

Preface	vii
Introduction	1
Part I. The Student Gap	
1. Answers Without Confidence	11
2. Capability Is Not Competence	16
3. When Guidance Is Missing	22
4. The Hidden AI Conversation	27
5. Learning in the Age of Assistance	31
Part II. The Faculty Gap	
6. Trust Under Pressure	41
7. Rethinking Evidence of Learning	45
8. The Changing Role of Faculty	51
9. Teaching Through Uncertainty	57
Part III. The Institutional Gap	
10. When Institutional Guidance Is Not Enough	65
11. Preparing for the Next Phase	70
12. Building a Guidance Ecosystem	75

Part IV. Closing the Guidance Gap

13. A Framework for Guidance	81
14. From Guidance to Practice	88
Acknowledgements	93
OER Licensing	96
References	97
About the Author	104
Related Books by the Author	106

Preface

When I started thinking about my next book, I knew I wanted to continue writing about AI. But I also asked myself: **What else can I say about AI that would actually help educators, faculty, learning designers, and others working in teaching and learning?**

There are already so many books about AI.

What could I contribute from my own expertise and working experience?

I have pushed myself to grow tremendously over the past year. I attended more than five conferences and delivered five public presentations related to AI in education. Some involved students. Some focused on faculty practice. Others allowed me to share the early stages of my research. I also continued using Pressbooks and OER to openly share what I was learning.

Through all of this, I realized something about myself.

I have worked in higher education for two decades, and deep down, the work I do has always been very practical. I see a problem and think about how we can solve it. I work with faculty on real courses, real students, real barriers, and real constraints.

But when I went to interviews, conferences, or conversations where someone asked me to describe my “framework,” sometimes my mind suddenly went blank.

What framework did I use?

Then I realized that I was already using these frameworks and approaches every day. I just did not always put names around what I was doing.

ADDIE. Bloom's Taxonomy. Quality Matters. Universal Design for Learning. Accessibility. Interactive design. AI integration. Learning design.

These were not simply theories or approaches I knew about. They had become part of how I looked at a course, identified a problem, and worked with faculty to make something better.

That realization led me to write my first set of five books, sharing what I knew from the perspective of a developer and specialist working directly with faculty. The purpose was never to tell educators that there was one perfect way to design a course. It was to help them get started, make progress, and think about what might work in their own courses.

Then AI changed the conversation.

I began submitting conference proposals about AI in education, presenting my work, listening to feedback, and, most importantly, listening to other people.

My first *AI in Education* book had reached more than 1,400 visitors by the time I began writing this second book. Seeing people access an openly available resource reminded me why I wanted to keep sharing this work.

But the conferences changed my thinking even more.

Evolution of Practice and Inquiry

My work with AI did not develop through one project or one presentation. It evolved through the year. I began by thinking about how course review practices could help faculty establish clearer AI guidance and transparency.

That led to conversations about practical AI guidance.

Those conversations led me toward questions about trust, transparency, and the conditions we create for learning. Then students entered the conversation more directly. Listening to their experiences changed the questions again.

The conversation was no longer simply about what faculty should do with AI. It became a conversation about student voice, implementation, ethics, AI literacy, reflection, human judgment, and what happens when institutional expectations meet the reality of the classroom.

Through conferences and faculty learning opportunities, I continued exploring these ideas: practical AI guidance, trust and transparency, student experiences, ethical and reflective instructional design, clear AI expectations, and the importance of maintaining human judgment.

Eventually, those individual conversations began developing into something more scalable through *AI Literacy for All* and faculty- and student-facing AI literacy resources.

Looking back, I can see that these were never isolated presentations or initiatives. They were part of an evolving exploration of AI, teaching, and learning. And somewhere along that journey, another question began to emerge.

If students already have access to AI, why are they still asking for guidance?

Learning to Become a Researcher

My Scholarship of Teaching and Learning research brought another kind of growth. I started by learning how to prepare a Research Ethics Board (REB) application, and from there I went through the research ethics process, conducted the research,

analyzed what I was seeing, and began writing for publication. Going through the whole process helped me understand research differently. More importantly, I discovered how much I enjoyed it.

I also discovered rejection.

One of my papers moved through initial review and into the journal review process. Another paper, focused on UDL, was rejected.

I read the editor's feedback and realized that the topic needed a much deeper exploration of the literature. So I went back, read more, and began to understand the topic differently. Through that process, I started to see another gap—one that I had not clearly seen when I wrote the first paper.

That gap led me to write a different paper and submit it.

Unfortunately, it was rejected too.

Was I frustrated? **My answer is no.**

I read the feedback again. This time, I started to see more clearly what I was missing. But I also want to be honest: I still need time to fully understand and digest some of the feedback. I can see the direction the editor is pointing me toward, but I am not yet completely sure how to get there.

What I appreciate is that the feedback did not simply point out weaknesses. It helped me see that there may be more potential in the idea if I am willing to go deeper. For example, I had approached different areas of the literature somewhat separately. The feedback made me realize that I may need to think more carefully about how those areas connect, interact, and can be synthesized rather than simply examining them one by one.

I understand that direction better now, but understanding what I need to improve and knowing exactly how to do it are two different things. I still need time to read, think, and learn how to approach the literature differently.

And I think that is part of becoming a researcher too.

For me, this rejection process has become an important part of learning how research works. You try something, receive feedback, question your thinking, read more, and explore again. Sometimes that process leads you to realize that the original question was not the question you should have been asking. Sometimes you are accepted, and sometimes you are rejected. Then you have to decide whether the idea matters enough to keep working on it.

For me, it does.

I also increasingly recognize how much I want scholarly mentorship as I continue developing as a researcher. There are things I can learn independently by reading, writing, trying, and responding to feedback, but there are also moments when I want an experienced supervisor to challenge my thinking, help me see what I cannot yet see, and guide me toward becoming a stronger researcher.

Perhaps recognizing that I still need guidance is part of my growth too.

Still Learning

As I begin my EdD journey while continuing my professional work, I know there is still much more for me to learn and perhaps that is exactly why I want to write this book now. I am not writing because I have all the answers. I am writing because I

want to continue learning, questioning, and contributing to the teaching and learning community.

There is something fitting about that.

I am writing a book about the importance of guidance while recognizing that **I need guidance too**.

When answers are everywhere, perhaps what we need most is not another answer.

We need guidance.

And that is where *The AI Guidance Gap* begins.

Introduction

The Era When Answers Became Free

Several years into the widespread use of generative AI, “AI” is still a keyword that attracts attention.

Include “AI” in the title of a conference session, workshop, article, or professional development event, and people are likely to notice. At conferences, AI sessions continue to fill rooms. Educators want to know what is changing, what other institutions are doing, what tools are emerging, and, perhaps most urgently, what they are supposed to do next.

After attending several conferences, speaking with educators, listening to students, and having informal conversations with graduate students from different institutions, I have started to wonder whether we are asking the wrong question.

We have spent a great deal of time asking what AI can do and how we can prevent students from using AI to skip the learning process.

Perhaps we need to spend more time asking what students need from us now that AI can do so much.

The Learning Has Already Changed

The conversation around AI in education continues to circle around the same problem: How do we know whether students are using AI?

One emerging discussion, for example, is whether AI-gener-

ated content can contain invisible markers or other signals that could help identify how the content was created (Cao et al., 2026). Then another tool will inevitably emerge to detect, remove, or work around those signals. We find ourselves circling the same issue again: trying to determine whether students used AI rather than asking what happened in their learning process.

But AI use itself does not tell us whether learning happened.

I am using AI to help me develop this book.

After completing my first *AI in Education* book, I began thinking about the practical, day-to-day problems that faculty were still facing. I used ChatGPT Edu to help me organize those ideas and think about the structure of this second book. I have also used AI to challenge my writing. I have asked it to read from the perspective of a faculty member: Is this actually helpful? I have asked it to act as a book editor: What would you criticize? Is the argument clear? Does this chapter help faculty do something differently?

Most importantly, I keep returning to the same question: **Am I delivering the message I actually want educators to hear?**

Don't panic.

We do not need a quick-fix solution every time AI develops another capability. We can look at the learning design itself. Sometimes a small change to an activity, an assessment, or the way expectations are communicated can make a meaningful difference.

AI is a tool. Instead of being afraid that students will use it, we can teach them how, when, and why to use it to support their learning, and when they need to think, decide, create, or struggle through something themselves.

This is where AI has been unusual for education. It did not simply introduce another technology into the classroom. It changed students' access to assistance.

Across the conferences I attended, I kept hearing variations of one message from students:

Trust us.

Students understand that their professors have expertise. They are there to learn from them. What they are asking for is not unlimited permission to use AI. They want to understand the boundaries. What can AI help with? What should remain their own judgment? When does assistance support learning, and when does it begin to replace the learning they are supposed to do?

Faculty are trying to work through many of the same questions. What should remain human judgment? What can appropriately be supported by AI? How do we know whether learning has happened?

At the same time, learning behaviours are evolving.

Students are used to getting information quickly. When something does not capture their attention, they scroll to the next thing. When they do not understand something, they search for another explanation. Increasingly, they expect answers almost immediately.

And now those answers are only a click or a prompt away.

Students no longer have to wait until the next class to ask a question. They may not wait in line after class. They might message an instructor through Teams, search online, ask AI, or sometimes not ask the instructor at all.

I heard a very practical example of this during one of my con-

ference presentations that included student perspectives. Vraj, one of the student panelists, explained that he cannot ask his professor every question that comes up while he is learning. There are times when he needs an answer or explanation and turns to AI for help instead.

That stayed with me because it reflects the reality students are learning in. It does not necessarily mean that they value their professors less or no longer want faculty support. It means that assistance is now available in places and at times when the professor cannot always be there.

The question, then, is not simply whether students will ask AI. It is whether they have enough guidance to know what to do with the assistance they receive.

They can ask an AI system at two o'clock in the morning. AI can explain a difficult concept. It can provide another example. It can give instant feedback on an idea, help organize thoughts, summarize a reading, suggest ways to improve a paragraph, or try a different explanation when the first one does not make sense.

And everyone can keep asking, including students, educators, and myself.

There is no embarrassment about asking the same question again. There is no need to raise a hand in front of a class. There are no office hours to wait for.

When you think about all of this, do you realize how much things have fundamentally changed?

Think about the old days of formal education. Access to knowledge was much more limited. Books, libraries, teachers, professors, and educational institutions were where learners went to find information and expertise that were not easily available elsewhere.

The internet already dramatically expanded our access to information. With a few clicks, we could search for almost anything. But having access to information did not mean that everything we found was reliable.

We started teaching students how to evaluate information. What is a reliable source? Where did this information come from? Should I trust this website? Is this a content farm? Where can I find more credible information to verify what I am reading?

When I think about it, the challenge with AI is not completely new.

We still need students to question the information they receive. They need to think about where it comes from, whether it is accurate, whether it can be verified, and whether they understand it well enough to use it responsibly.

The technology has changed, but the need to teach students how to evaluate information has not.

What has changed is that students are no longer only searching for information. They are interacting with systems that can explain, suggest, generate, respond, and assist throughout the learning process. Knowing how to evaluate information is still important, but students also need to make decisions about **when, why, and how that assistance should be used** (Laupichler et al., 2022; Long & Magerko, 2020).

The AI Guidance Gap

I call the space between access and understanding **the AI Guidance Gap**.

In my own institution, I have conducted Scholarship of Teach-

ing and Learning research related to students' use of AI (Chan, 2026a). I have also been involved in developing *AI Literacy for All*, including an *AI Faculty Learning Guide* and *A Practical Guide for Students*, with the resources launching at the beginning of the academic term.

Through this work, one thing has become increasingly clear to me.

Whenever I hear faculty asking questions about AI, such as what students should be allowed to do, how assessment should change, how we should talk to students about AI, and what should remain human, AI guidance is often a practical place to start.

But guidance is not simply about telling people what they can and cannot do.

I define **the AI Guidance Gap** as the distance between having access to AI and having the judgment, clarity, support, and trust needed to use it meaningfully for learning.

The AI Guidance Gap appears when students have powerful tools but do not understand what their instructors expect. It appears when students know how to ask AI for an answer but do not yet have enough knowledge or judgment to evaluate that answer.

It appears when faculty are expected to respond to AI but do not feel sufficiently supported to rethink their teaching and assessment. It also appears when institutions create policies and guidance, but the people using them still do not know what responsible AI use looks like in a particular classroom, assignment, discipline, or learning situation.

The gap cannot be closed simply by providing another AI tool. It cannot be closed by detection software, and it cannot be

closed by writing increasingly detailed rules for every possible use of a technology that will continue to change.

What we need is guidance.

That means conversations, clear expectations, opportunities to make choices, reflection on those choices, feedback, and trust. It also means educators who are willing to help learners develop judgment rather than simply provide or police answers.

This book, therefore, is not primarily about artificial intelligence.

AI is the condition that makes the problem visible.

The book is about what happens to teaching and learning when answers are everywhere, but understanding is not. It is about students who need guidance even when assistance is available almost anytime they need it. It is about faculty trying to determine what their role becomes when they are no longer the only source of explanation. It is also about institutions discovering that policies and high-level guidance can establish a foundation, but still need to become meaningful practice.

Ultimately, this book is about a question that may matter much more than what the next generation of AI will be able to do:

When answers are everywhere, what do learners still need from us?

PART I

THE STUDENT GAP

1. Answers Without Confidence

Key Question

Why do students with unlimited access to answers still need help?

Why Students Still Feel Uncertain Despite Having AI

As I mentioned in the introduction, I attended more than six conferences over the past year. The topics, presentations, institutions, and audiences were different, but I kept hearing variations of one message from students.

Trust us.

Students understand that they are there to learn. They respect that their professors have expertise and are there to teach them. What I heard was not students asking educators to let them use AI for everything.

Instead, the message sounded more like this:

Teach me how.

Students know that AI is available. Answers are everywhere. It is difficult to expect someone never to ask a tool that can provide an answer almost instantly. But students also do not want to be accused of cheating simply because they used AI. They want to

understand how to use it responsibly, what is appropriate, what is not, and where the boundaries are.

I saw this from several directions.

In my EvolvED presentation, I intentionally brought student perspectives into the conversation (Chan & Gaidosch, 2026). Through my other AI presentations and my own Scholarship of Teaching and Learning research, I continued exploring questions about expectations, responsibility, and student decision-making (Chan, 2026b). At the STLHE conference, similar messages appeared again in conversations and presentations.

Different settings. Different people.

I kept coming back to the same three words that I had been presenting in my own work:

Trust. Clarity. Transparency.

At the STLHE conference, I presented *AI in Education: Ethics, Literacy, and Interactive Design* on June 17 (Chan, 2026a). During that presentation, I emphasized these three ideas as important conditions for working with students around AI.

Then something happened that I did not expect.

The following day, Maheux-Pelletier et al. (2026) presented the keynote *Small Teaching Decisions, Big Shifts: Student Voices on Adapting Teaching and Learning in the Age of GenAI*. During the keynote, my presentation was mentioned by name, along with the three words I had emphasized: **“Trust. Clarity. Transparency.”**

I was genuinely surprised and honoured to have my work recognized during the keynote. It also reinforced the same message I had shared the day before: **trust, clarity, and transparency.**

More importantly, it reinforced something I had been noticing across my own work, research, presentations, and conversations with students. These were not simply three words that sounded good together. They were beginning to describe something students and educators actually needed.

If educators and students can build from trust, clarity, and transparency, perhaps AI does not have to become a tool for avoiding learning.

It can become a tool used within learning.

That distinction matters.

The Difference Between Receiving an Answer and Understanding It

Perhaps the easiest way for me to explain this is to show you how I am using AI to write the book you are reading right now.

I mentioned in the preface that my first AI book, *AI in Education: Ethics, Literacy, and Interactive Design*, focused on a message I wanted educators to hear: we do not need to be afraid of AI.

We need to understand it.

We need to think about how it can be used ethically and meaningfully.

And we need to recognize where human judgment should remain.

After completing that book, I began wondering what I could contribute next.

I knew I did not want to write another book simply explaining what AI is or introducing another collection of AI tools. There is already plenty of that information available.

More importantly, my day-to-day work was showing me that the problem had moved somewhere else.

Both faculty and students continue to ask questions.

Students are asking: *Can I use AI for this? How much can I use? Will my professor think I cheated? What do I need to disclose?*

A recent example came from a Travel Journal Writing course. The course had a clear policy stating that AI use was not permitted. However, students asked whether they could use generative AI to create images to accompany a writing assignment.

The initial answer was no because the course policy said no AI use.

But when the faculty member asked me for advice, I asked a different question:

What are you actually assessing?

If the learning being assessed is the students' writing, and the accompanying image is not being graded, should the use of AI to create that image be treated in the same way as using AI to generate the writing itself?

I was not suggesting that the course policy should simply be ignored. Rather, the student request created an opportunity to look more closely at the purpose of the assignment, what students were expected to learn, and where AI use might—or might not—interfere with that learning.

A rule can give us an answer.

No AI.

But understanding requires us to ask why that rule exists and how it applies to a particular learning situation.

Faculty are asking their own questions: *What should I allow? How do I know whether students are learning? What should remain human judgment? How do I redesign my assessment?*

The questions may be different, but both sides are looking for something similar: guidance.

The more I listened, the more I realized that the issue was not simply whether AI should or should not be used.

The issue was whether we had created enough **trust, clarity, and transparency** for faculty and students to understand how AI could be used within the learning process.

Resources existed here and there, but I could see a continuing need for clearer, more centralized AI literacy guidance that could help people make decisions about responsible AI use.

At the same time, I was thinking about future conference presentations. What could I share next? What had I learned from the conversations I had already been part of? What was still missing?

2. Capability Is Not Competence

Key Question

If AI helps us produce more, does that mean we have learned more?

How I Am Using AI to Write This Book

So I started talking to ChatGPT and used it to help me generate some bullet-point ideas for this section.

And this is where the distinction between **receiving an answer** and **understanding an answer** becomes important.

From those bullet points, I started drafting. As I wrote, I thought about practical examples I had encountered while working with faculty and began bringing those experiences into the chapter.

I did not type:

Write my second book about AI.

And then receive a whole book.

The process did not work like that.

I brought my experiences into the conversation. I started drafting this first chapter using some of the bullet points I had developed through my conversation with AI. Then I added

what I had heard from students. I added the questions faculty had been asking me. I brought in my conference experiences, my research, my first book, the AI literacy resources I had been developing, and practical problems I had encountered through my work.

AI helped me organize those thoughts. It suggested possibilities. I responded to them. Some ideas did not fit. Others made me think differently. Sometimes I changed direction. I added experiences and context that AI could not possibly know unless I provided them.

Eventually, a larger problem became clearer to me.

There was a gap between **having access to AI and knowing what to do with it.**

That became the idea behind *The AI Guidance Gap*.

But even then, I did not have a book.

I had an idea and some text to start with, and now I am doing the work of turning that idea into the draft you are reading.

I still need to decide what this book is actually trying to say. I need to determine the structure. What belongs in the student section? What belongs in the faculty section? What belongs at the institutional level? And just as importantly, what does not belong in the book at all?

Then I need to write.

When I finally had time to sit down and work seriously on the introduction, I wrote from my own experiences and ideas. But the page kept getting longer. At one point, I started wondering, **Did I just write the whole book in the introduction?**

I used AI to help me review, organize, challenge, and refine what I was trying to communicate. But I also noticed some-

thing about working with AI. It can easily follow the direction I am already leaning toward. If I tell it that I think a section should stay, it can give me reasons to keep it. If I tell it that I think the section is too long, it can give me reasons to cut it.

That means I cannot simply ask AI for an opinion and treat the response as the answer. Sometimes I need to challenge it from another perspective, ask the question differently, or simply walk away from the writing and come back later with my own judgment.

That is another reason I need to remain the one making the final decision.

Human Judgment Still Matters

This is where human judgment becomes important.

AI can suggest a structure, identify repetition, challenge an argument, or respond from the perspective of an editor or educator. Those suggestions can be useful, but AI does not automatically know what I am trying to contribute or which version best represents what I want my readers to understand.

I still need to make the final decision. I can take a suggestion and make a change, or I can reject it because I have decided that another direction better represents my purpose.

I have to remain in charge of the writing and take responsibility for the final result.

This is my book. It needs to represent my voice, my experience, and what I want to share with my readers. I bring the context and understand the people I am writing for. AI can help me think through possibilities, but I decide whether those possibilities belong in the book.

If I do not understand my own topic, I may not know whether an AI suggestion is useful at all.

That is the difference between receiving an answer and understanding it, and it leads to another important question:

If AI can make us more capable of producing something, does that necessarily mean we have become more competent at doing it?

Confidence, Competence, and Learning

That question brings us to the next part of the problem.

AI can make us feel capable very quickly.

But does that mean we have actually become more capable?

Think about brainstorming. Give AI a rough idea and it can return a list of possibilities within seconds. Some of those ideas may be useful. Some may sound polished. Some may even give us a direction we had not considered.

But will the final idea necessarily be one of those suggestions?

In my experience, often it is not.

When I work on a book, I sometimes begin with a suggested structure developed through a conversation with AI. It gives me somewhere to start. But by the time I actually begin writing, some of those original bullet points may no longer fit. I remove them. I move things around. New ideas emerge as I write. My own experiences remind me of examples I had not considered before.

Sometimes the final structure looks very different from where I started.

That is why I would say AI inspires and supports my thinking, but it does not do the thinking for me.

The same applies to writing.

Give AI a question and it can return a confident answer. Give it a paragraph and it can suggest improvements. Ask it to brainstorm and it can generate possibilities almost instantly.

But the quality of the output and the quality of our learning are not necessarily the same thing.

For example, I use AI to help me notice things I might overlook in my own writing. I may ask whether I am repeating the same idea, circling around the same point in different sections, or leaving something unclear. Sometimes I also ask AI to review my writing from an editor's perspective. I look at the suggestions, decide which ones are useful, reject the ones that do not fit, and continue editing until the writing represents what I actually want to say.

But being able to use AI to help structure and review this book does not mean that I suddenly understand learning theory better. The same thing can happen with students. A student can receive a well-written explanation without fully understanding the concept. They can generate a polished paragraph without knowing whether the argument is correct. They can also receive a confident AI-generated answer without having enough knowledge to decide whether that answer is useful, accurate, or appropriate.

This is where I started thinking about the difference between **confidence, competence, and learning**. AI can give us something to work with very quickly, and it may help us produce something that we could not produce as easily on our own. But being able to produce something does not necessarily mean

that we have developed the knowledge or skills to do it ourselves (Anders & Dux Speltz, 2025; Wang et al., 2026).

That is why I keep coming back to the key question of this chapter: **If AI helps us produce more, does that mean we have learned more?**

I do not think we can assume that it does. AI can support the process, but we still need to think about what the learner actually understands and what they can do because of what they learned.

3. When Guidance Is Missing

Key Question

What happens when expectations are unclear?

The Assignment Says Nothing About AI

At my institution, faculty are provided with sample course-outline language that includes three broad approaches to generative AI: AI use is not permitted, AI use is permitted for specific components of the course, or AI use is permitted in the course. Students are also encouraged to ask their instructors about the rules for using generative AI in each course (George Brown Polytechnic, n.d.).

That gives students a starting point, but think about it from their perspective. A student opens an assignment and starts reading the instructions. They may know the overall position on AI in the course, but they still need to understand what that means for this particular assignment.

But even asking a question can be a challenge. Where do they ask? How do they ask? And what exactly should they ask the instructor?

The assignment may already include the learning outcomes, requirements, rubric, due date, and everything else the student needs to complete the work. But there may be very little infor-

mation about what AI use actually means for this particular assignment.

Students may already be using AI in many different ways to support their learning and work on assignments. In another course, for example, a professor may have encouraged students to use AI to support their learning, perhaps beginning with brainstorming. In another assignment, students may have been allowed to use AI as long as they disclosed how they used it. Another professor may have said absolutely **NO to AI**.

From the student's perspective, this can become confusing. When can they use AI, and when can't they?

A whole list of questions may come up:

- Why is it okay here but not there?
- Why can I use AI for one assignment but not another?
- What exactly does "AI use permitted" mean?
- Can I use AI to brainstorm?
- Can I ask it to explain the assignment?
- Can I use it to organize my ideas?
- What about grammar?
- What if I write everything myself but ask AI for feedback?
- Do I need to tell my professor?

If the assignment does not answer these questions, the student may have to make a guess and hope that their interpretation matches what this particular professor considers acceptable.

The problem is not necessarily that the answers are different. The problem is that students may not understand why they are different.

When guidance is missing, students do not stop making decisions about AI. They still have to decide what is appropriate,

what is not appropriate, and where the boundary might be. Sometimes they may also be worrying about what will happen if their interpretation is wrong.

This is what happens when guidance is missing.

Students Fill in the Blank

When we leave something unclear, students still have to decide what to do. They are already making judgments, but they may be making those judgments without knowing whether their interpretation matches what the educator intended.

A professor may believe the expectation is already clear because students were told that AI is allowed or not allowed. The student may also believe that their own interpretation of that statement is reasonable.

But “AI is allowed” does not necessarily explain how AI should be used, and “AI is not allowed” may not answer every question about the AI-enabled tools students encounter in their everyday work.

This is where I see the gap.

The faculty member may believe the expectation has been communicated, while the student may believe they understand it. But they may actually have two different interpretations, and neither side may realize that until something goes wrong.

Sometimes that does not happen until after the assignment has already been submitted.

There is another assumption we need to think about here. As educators, we may think that if students are unsure, they can simply ask their professor. But will they?

There could be many reasons why a student does not approach their professor before submitting an assignment. A student taking an asynchronous online course may rarely interact directly with the professor. A student in a hybrid course may not attend every in-person opportunity. Even in an in-person class, some students may feel uncomfortable or afraid to approach their professor. Others may search for the answer themselves, ask their classmates, or now ask AI instead.

I have seen and heard examples of these situations through my own work and conversations with students. That is why I do not think we can simply assume that unclear guidance will always lead to a student asking us for clarification.

As educators, we may expect the conversation before submission to be:

How should I use AI for this assignment?

But if that conversation never happens, the question after submission may become:

Why did you use AI for this assignment?

Those are very different conversations.

This is why I think we need to focus more on **preventing the misunderstanding rather than policing it after it has already happened**. Before submission, there is still an opportunity to clarify expectations and provide guidance. After submission,

both the student and faculty member may already be trying to explain what they thought the expectation meant.

If we want students to develop responsible ways of using AI, **the more useful conversation needs to happen before they click Submit, not after.**

4. The Hidden AI Conversation

When faculty receive an assignment, we see the final product. But how much do we actually know about what happened before the student clicked **Submit**?

With AI, there may be many decisions behind that final product. The student may have considered using AI and decided not to. They may have used it to brainstorm, asked for feedback, rejected what it suggested, or used part of a response and changed direction.

Most of those decisions are hidden unless we create a reason for students to tell us about them.

My own Scholarship of Teaching and Learning (SoTL) research gave me another opportunity to see part of this hidden conversation.

In the study, students were provided with AI guidance, but using AI was optional. They were not graded based on whether they chose to use or not use AI. Instead, students were asked to explain whether they used AI, how they used it, and, importantly, why they made that decision (Chan, 2026b).

That last question matters to me.

Asking students *why* gave me a different window into their decision-making process.

Across the student statements and reflections, I found recurring patterns related to deliberate decision-making, using AI as support rather than substitution, maintaining creative ownership, and reflecting on learning and authorship (Chan, 2026b).

I will not reproduce the full study here. What is important for this discussion is what those reflections helped make visible:

Students were already making decisions about AI.

The transparency statement gave them a place to explain those decisions.

That made me think differently about transparency. Perhaps transparency should not only be something we ask students to provide after they use AI. It can also become part of the learning process by encouraging students to think about what they used, why they used it, and what remained their own responsibility (Ratnayake et al., 2024; Weaver, 2024).

As I was preparing to release this book, I came across a newly published UNESCO collection that connected closely with this thinking. Desai (2026) argues that as final artefacts become less reliable indicators of student learning, greater attention can be given to the learning process, including students' thinking, iteration, dialogue, and metacognitive reflection.

This does not mean that educators need to see every prompt a student enters into an AI system.

That is not the point.

What may be more valuable is creating opportunities for students to explain the decisions that matter.

- *Why did you use AI here?*
- *Why did you choose not to use it?*
- *What did AI contribute?*
- *What did you change or reject?*
- *What remained your responsibility?*

Perhaps we do not need to uncover every AI interaction.

But if students are already having this hidden conversation—making choices, questioning boundaries, accepting some suggestions and rejecting others—we can create learning environments where they feel able to bring some of that thinking into the open.

The hidden AI conversation is already happening.

The question is whether we create a reason for students to share it with us.

Related Scholarship and Ongoing Inquiry

This chapter draws in part from my ongoing Scholarship of Teaching and Learning inquiry into AI transparency, reflection, and student decision-making. Readers interested in exploring the research and practice behind this work can find more through the following resources.

SoTL Project: *From Compliance to Reflection: AI Transparency Statements as a Pedagogical Strategy Supporting Student Agency in Higher Education*

This SoTL inquiry examines how students describe and reflect on their AI-related decisions when AI guidance and transparency are intentionally incorporated into the learning process. The study explores transparency not simply as disclosure or compliance, but as a potential opportunity for reflection, decision-making, and student agency.

SoTL Spotlight: *Watch the SoTL Spotlight video*

This short research spotlight introduces the development of the inquiry, emerging insights, and its connection to ongoing work on responsible AI integration and student decision-making.

Related Conference Presentation:

Chan, E., & Gaidosch, B. (2026). *From Guidance to Practice: Starting Small with AI Integration—A Film Studies Case Study*. EvolvED Conference, George Brown Polytechnic, Toronto, ON.

This presentation shares the Film Studies implementation that helped inform the ongoing SoTL inquiry, including the use of AI guidance, transparency, student reflection, and student perspectives.

Chan, E. (2026, October 30–31). *Learning across states of awareness: Reflective inquiry and educational leadership* [Accepted paper presentation]. Centre for Leadership and Diversity 7th Annual International Conference, Ontario Institute for Studies in Education, University of Toronto, Toronto, ON, Canada.

This paper presentation extends the SoTL inquiry into a broader examination of reflective inquiry and educational leadership, adapting the work to the context of the OISE conference.

5. Learning in the Age of Assistance

Learning in the Age of Assistance

Key Question

How do students learn when assistance is always available?

At my institution, there are certain elements that we consistently include in our courses. Whether a course is delivered in person, online, or in a hybrid format, the learning management system provides students with a central place to access important course information and support.

My own background also shapes how I look at this. Much of my course design and development work has been informed by Quality Matters (QM) and Universal Design for Learning (UDL) (CAST, 2024; Quality Matters, n.d.). I naturally look for clear course structure, alignment, accessibility, learner choice, and whether students can easily find the information and support they need.

That is why an area such as **Essential Information (Start Here)** matters to me.

It is not simply a place to store documents. It is part of how we guide students into the learning environment and help them understand where to begin, what is expected, and where they can find support.

Students may find the course introduction, learning outcomes,

professor information, course outline, textbook and resource information, evaluation details, academic policies, accessibility information, and directions for where to get help.

As AI became part of teaching and learning, we also needed to think about what additional information students needed, including AI usage policies, citation and reflection guidance, and information about responsible AI use.

More recently, through our *AI Literacy for All* work, we have been developing a shared foundation for AI literacy, along with learning guides designed specifically for faculty and students.

When I look at how these resources have evolved, I am reminded of something important.

Students have always needed assistance with learning.

That part is not new.

Before generative AI, students asked their professors questions when they needed help. They collaborated with classmates. Group work has long been one way for students to learn with and from others. Students visited tutoring centres, writing centres, libraries, and other academic support services.

Then the internet changed how we looked for help.

Students searched Google. They watched YouTube videos to learn how to do something. They used online examples and digital textbooks. They asked questions in online communities such as Reddit. And, of course, they still asked classmates, professors, friends, and family.

The pandemic accelerated another change.

When physical campuses around the world suddenly closed, many forms of support that had traditionally been connected to physical locations had to move online. Students and edu-

cators became increasingly accustomed to accessing learning, resources, communication, and support without necessarily being in the same place.

These are changes I have observed through my own experience working in higher education and watching learning environments evolve.

Including information about where to get help in a course is not new.

But something about getting help has changed.

For students who grew up with the internet, information was already relatively easy to search for. We talked about digital citizenship and helping students learn how to participate responsibly in increasingly digital environments.

When generative AI emerged, I initially began thinking about something similar: **AI citizenship**.

My practical thinking was to begin with faculty. If we want students to develop responsible ways of using AI, faculty also need opportunities to understand AI, experiment with it, and think about how it fits within their own teaching.

Over time, my language shifted from *AI citizenship* toward **AI literacy**. The word *citizenship* did not always connect with people and could sound more serious than the practical approach I was trying to introduce.

But the idea behind it continued.

That thinking eventually developed into faculty learning opportunities such as the **AI Literacy in the Classroom Series: *Simple Ways to Set Clear AI Expectations in Your Course***.

The purpose is intentionally practical and low-overload. Rather than asking faculty to become AI experts, the workshop helps

them think about something they can immediately apply in their own courses: **How can I make my AI expectations clearer for students?**

Faculty work with ready-to-use AI guidance templates, reflection prompts, and reusable learning materials that they can adapt for their own teaching. The approach has continued to evolve through implementation and iteration in actual courses.

For me, this is an important part of AI literacy.

We cannot simply tell students to **“use AI responsibly”** and expect them to know what responsible use means in every learning situation.

Faculty need guidance too.

And when faculty become more confident about explaining when, why, and how AI can be used within their courses, they are in a better position to guide students in making those decisions themselves.

The terminology evolved from AI citizenship to AI literacy. But the need remained the same: help people make informed decisions about AI rather than simply giving them access to it (Hackl et al., 2026; Laupichler et al., 2022; Long & Magerko, 2020).

Dependency Versus Support

I do not think the problem is that students ask AI for help. We want students to reach out for help when they need it, whether that help comes from their professor, a librarian, another academic support service, or an appropriate use of technology.

I thought about this a lot while working on ***A Practical Guide for Students*** as part of our *AI Literacy for All* work. Portions of

the instructional content were inspired by and adapted from the University of Toronto's *Coursework and GenAI: A Practical Guide for Students* (University of Toronto, 2025a). From that starting point, we developed and adapted the resource for our own George Brown context.

We also wanted student voices to be part of the resource. With the support of a faculty member, we invited two students from a public speaking class to record introduction and summary videos for the modules. I prepared the scripts for them, but even that process taught me something. There were terms that those of us working with AI may use regularly that were not necessarily familiar to students. For example, when we talked about **LLMs**, we realized that we needed to explain what that actually meant rather than assume students already knew.

That reminded me that if we are developing AI literacy resources for students, we cannot design them only from the perspective of people who already work with AI every day. Bringing student voices into the modules helped us think again about the language we were using and how other students might experience the material.

The Practical Guide also does not simply tell students to use AI responsibly and leave them to figure out what that means. Along the way, we identify places where students can get help. For example, if they are unsure about sources or information generated by AI, they can reach out to library support. If they are unsure whether AI is permitted for an assignment, they need to check the assignment guidance or ask their instructor.

We also adapted an open **Generative AI Decision Flowchart** from The Learning Portal, College Libraries Ontario (n.d.). The flowchart gives students a practical way to work through questions about permission, assignment instructions, evaluating AI-generated information, and acknowledging or disclosing AI use.

That last question connects directly with what I have been thinking about in this chapter.

A student may ask AI to explain a difficult concept in simpler language and then use that explanation to continue working through the material. Another student may ask AI to produce the answer and move on without really understanding it. Both students used AI, but what they did with the assistance was different.

This is why I do not think we can simply categorize AI as good or bad for learning. We need to look more closely at what the student is doing with the assistance and whether that assistance is helping them continue learning or beginning to replace something they need to learn themselves.

Students can also learn to ask themselves these questions. Did AI help me understand something, or did it simply help me finish faster? Could I explain what I learned without AI beside me? Did I use the response to continue thinking, or did I allow AI to do the thinking I was supposed to do?

A policy cannot answer all of these questions for students. Telling them to **“use AI responsibly”** is also not enough if we do not help them understand what responsible use looks like in an actual learning situation.

This is where guidance becomes part of AI literacy. Students need opportunities to make decisions about AI, think about what happened, receive feedback, and use what they learned when they make the next decision (Anders & Dux Speltz, 2025; Ratnayake et al., 2024).

Learning in an age of assistance does not mean that students should never ask for help, and it does not mean that AI should become the answer to everything. What matters is helping stu-

dents recognize when assistance is supporting their learning and when they need to do more of the thinking themselves.

Closing the Student Gap

Across this first part of the book, I have looked at the Student Gap from several different directions. Students have access to answers, but having an answer does not mean they understand it. AI may help them produce something they could not produce as easily before, but that does not automatically mean they have developed the knowledge or skills behind it.

When guidance is unclear, students still have to make decisions. They have to interpret what their professor expects, decide what kind of AI use may be appropriate, and sometimes make those decisions before faculty ever see the final work. Now they are also learning in an environment where assistance can be available almost anytime they want it.

I do not see this as a problem with students. Students need opportunities to make decisions, reflect on what happened, receive feedback, and make different decisions the next time. But if we want them to use AI responsibly, we also need to give them enough guidance to understand what responsible use means within their learning.

That is what I mean by the **Student Gap**. It is the space between the access students already have and the guidance they need to develop judgment around that access.

This leads naturally to the next part of the book.

**If students need this kind of guidance, are faculty
being given enough support to provide it?**

That brings us to the Faculty Gap.

PART II

THE FACULTY GAP

6. Trust Under Pressure

Why Focusing on Catching Students Does Not Solve the Problem

In the *Silver Lining for Learning* podcast discussion, Dr. Lydia Cao also challenged approaches that focus primarily on policing students' AI use (Cao et al., 2026).

This connects closely with my own thinking.

What I appreciate about her framework is that responsibility is built into the learning process. Instead of simply telling students whether AI is allowed or prohibited, the framework creates an opportunity to think about the learning objective, how AI may affect that learning, what role AI should play, and when it should be used (Cao, 2025).

This is closely aligned with an approach I have been developing through my own work since my first *AI in Education* book: **teach students how to make responsible decisions about AI rather than focusing primarily on catching them using it.**

I want to be clear that I am not saying academic integrity no longer matters. Of course it matters. Faculty still have a responsibility to assess whether students have demonstrated the required learning, and trusting students does not mean

that we simply accept every submitted assignment without asking questions.

But I do not think detection is where we should spend most of our energy.

Even if technology becomes much better at identifying AI-generated content, knowing that AI was involved does not automatically tell us what the student did with it. A student may have used AI to help understand a concept, receive feedback, or improve something they had already developed. Another student may have asked AI to do most of the thinking for them. These are very different learning situations, even though AI was involved in both.

This is why I keep coming back to the learning itself.

I recently worked with a faculty member who was reconsidering the discussion boards in a course. Students were expected to read course materials and then choose three questions to answer from a set of five. There were approximately ten discussion boards across the course, with the discussion activities together accounting for a meaningful portion of the final grade.

The original purpose was to give students an opportunity to think about the course material and discuss the topics. But the faculty member had started noticing that some of the responses were becoming extremely polished.

That raised a concern.

If a student can put the discussion question into AI, receive a strong response, and post something that sounds excellent, what are we actually learning from that discussion-board response?

We could focus on trying to determine whether AI generated

the response. But even if we could detect that perfectly, we would still need to return to the more important question: **What was the discussion board supposed to help students learn?**

If the purpose was for students to read, think about the material, develop their own position, and interact with the ideas of their classmates, then perhaps the issue is not simply whether AI should be prohibited. We may need to look again at the design of the activity and ask what would make the student's thinking more visible.

Could students connect the discussion to a specific example from the course, perhaps by watching a short video clip and responding to something they observed? They could explain why they chose a particular position, respond to something another student said and consider whether it changed their thinking, or reflect briefly on how they arrived at their response.

Small Changes

These are small changes, but they move our attention away from trying to catch AI and back toward the purpose of the learning activity.

This is where I think we need to shift our thinking. Instead of focusing mainly on **“How do I catch students using AI?”**, we can ask what students are supposed to learn, where AI can

appropriately help, where it may interfere with that learning, and what students should still be able to demonstrate themselves.

Throughout this book, I keep returning to the same three ideas: **Trust. Clarity. Transparency.** Trust does not mean ignoring inappropriate AI use. Clarity means helping students understand what is expected and why. Transparency gives students an opportunity to explain how they made their decisions rather than leaving faculty to guess what happened from the final product alone.

For me, the more useful question is therefore not simply whether AI was used. I want to know **what happened to the learning when AI was used.**

Detection may give us some information about the final product, but it cannot fully explain the learning process behind it. This is where guidance and human judgment become much more important.

If we want students to be transparent with us about their AI use, we also need to create an environment where they feel that they can have that conversation with us.

Transparency also does not automatically create trust. In a recent UNESCO collection, Muellauer (2026) describes being open with students about using AI to support feedback, while still finding that some students interpreted the presence of AI as the absence of the teacher's personal attention. This reminds me that transparency matters, but trust also depends on how students experience the relationship behind the use of AI.

7. Rethinking Evidence of Learning

Key Question

What should assessment measure now?

When the Final Product Is No Longer Enough

Assessment was already evolving before generative AI.

From my own experience working in the polytechnic and college sector, we have long used formative and summative assessments throughout a course rather than relying only on one final product. My own course design practice has also been influenced by Quality Matters (QM), Universal Design for Learning (UDL), accessibility, alignment, and student agency.

I have explored many of these ideas in my earlier work, so I will not repeat them here.

What generative AI adds is a new urgency to a familiar question:

What actually counts as evidence that learning happened?

I have seen this question emerge through my own course development work. In one project that I later presented at *D2L Fusion – Adapting Film Studies Courses for Diverse Learning Environments Using D2L Brightspace*, we moved away from

relying primarily on a traditional final assessment. Students began with a technical proposal and then worked through learning activities, feedback, and developmental steps leading toward their final project (Chan, 2025).

The final project still mattered, but it was not the only place where learning became visible.

Now bring generative AI into that picture.

If AI can help write an essay, organize a report, create presentation slides, generate images, summarize information, write code, or provide answers to questions, what exactly does the final product tell us?

A polished assignment may represent significant student learning or it may not.

That is the assessment dilemma.

The problem is not that our existing assessments suddenly became useless. The problem is that we may no longer be able to assume that the final product alone tells us everything we need to know about the learning that happened before it (Gulikers et al., 2004; Perkins, 2026).

Earlier in this book, I asked a faculty member a question I often return to:

What are you actually assessing?

If AI performs the skill we intended the student to demon-

strate, we have a problem. If AI supports a part of the task that is not actually being assessed, the educational question may be different.

This is why I do not think the solution is simply to label every assessment **“AI allowed”** or **“AI prohibited.”** I discussed this approach in my first AI book. The issue goes deeper than permission.

We need to understand what evidence of learning we actually need from the student.

Rethinking Evidence of Learning

Imagine two students submitting equally strong final assignments.

The first student develops the idea, researches the topic, drafts the work, uses AI for feedback, rejects some of its suggestions, revises the assignment, and submits the final version.

The second student asks AI to generate most of the assignment and then makes minor edits before submitting it.

The final products might look similar.

But the learning processes are very different.

If we assess only the final product, how much of that difference can we actually see?

This connects directly with the previous chapter.

Detection may tell us something about whether AI was involved. But assessment has a different purpose.

Assessment should help us understand what students know, what they can do, and how they are developing.

That does not mean faculty need to monitor every step students take. It does not mean asking students to submit hundreds of AI prompts. And it certainly does not mean adding so many new requirements that assessment becomes unmanageable for both faculty and students.

Instead, perhaps we need small opportunities for students to make important parts of their learning process visible.

Making Learning More Visible

Sometimes the change can be surprisingly small.

A short reflection, an early draft, a checkpoint, a rationale, a conversation, or a brief explanation of an important decision may give us another form of evidence (Perkins, 2026; Ratnayake et al., 2024).

Instead of asking students only what they produced, we can also ask how they got there. What decisions did they make?

Where did they need assistance? If they used AI, why did they use it? What did they accept or reject? What remained their responsibility? And, perhaps most importantly, what did they learn?

This is something I also began seeing through my own SoTL research. Asking students not only whether and how they used AI, but why, created an opportunity for students to make some of their reasoning visible (Chan, 2026b).

For me, the reflection was not simply about documenting AI use. It became part of the learning process.

This does not mean we need to redesign every assessment.

Sometimes adding one meaningful reflection question, one checkpoint, or one opportunity for students to explain a decision may tell us something that the final product alone cannot.

None of these approaches makes an assessment “AI-proof.”

I do not think that should be the goal.

The goal is to create better opportunities for students to demonstrate learning.

Instead of beginning with **“How do I stop students from using AI?”**, perhaps we should begin with **“What evidence would help me know that my students learned?”**

Once we understand that, we can make a much more informed decision about what role AI should or should not play (Tai et al., 2018).

Knowledge, skills, application, critical thinking, creativity, and judgment still matter. What may need to change is how much we rely on a polished final product as evidence that all of those things happened.

When polished outputs are increasingly easy to generate, we may need to make selected parts of the learning process more visible—not every step, not every prompt, and not through constant surveillance, but enough for students to demonstrate that they understand what they did, can explain the decisions they made, and remain responsible for the work they submit.

Perhaps assessment in an AI-enabled world should therefore not only ask what students produced. It should also help us understand **what they understood, what decisions they made, and what they can now do because of what they learned.**

8. The Changing Role of Faculty

Key Question

What becomes more important when AI can explain almost anything?

Expert, Coach, Mentor, Guide

I have worked with faculty who moved courses from in-person teaching into asynchronous online environments. That transition already changed the role of the professor.

Instead of standing in a classroom and delivering most of the course content directly, faculty may use OER, videos, interactive activities, online discussions, and other digital resources to support learning. The faculty member does not disappear, but the role shifts. In many ways, the professor becomes more of a facilitator of learning.

Generative AI pushes that shift even further.

If students can ask AI to explain a theory, define a term, summarize information, or provide another example, what becomes more important about the professor?

I do not think faculty need to compete with AI by proving that they know more information. Their expertise is much deeper than that.

A faculty member brings disciplinary knowledge, professional

and research experience, context, judgment, and an understanding of how knowledge is actually applied in the field. A student may be able to ask AI for an explanation of a theory, but the professor can explain what that theory looks like in practice, where it becomes complicated, what happened when they applied it professionally, or what they discovered through their own research.

Perhaps this is why the faculty role increasingly includes several roles at the same time: **expert, facilitator, coach, mentor, and guide.**

Technology Still Needs Human Interpretation

Faculty may also use intelligent technologies themselves. Learning management systems, for example, can provide analytics and automated tools that help faculty identify patterns in student activity.

A system might show that a student has not logged in recently. Does that mean the student is struggling academically? Perhaps. But the student may also be experiencing an access problem, dealing with work or family responsibilities, or only having certain times during the week when they can study.

Similarly, a student who appears to spend several hours on one page may be struggling with the content, or they may simply have left the browser open.

This is why I think we need to be careful about what technology tells us. The data can provide a signal, but the educator still needs to interpret what that signal means.

An intelligent agent in an LMS can send a reminder, and an

announcement can communicate with the entire class. These tools can help faculty reach students more consistently, but they do not necessarily understand why a student has stopped participating or what kind of support that student may need.

The technology can assist us. **Human judgment still matters.**

Facilitating Learning and Designing Connection

If information is increasingly easy to access, facilitating learning becomes even more important. Faculty still teach, explain concepts, and share disciplinary knowledge, but teaching has never been only about delivering information. It also involves feedback, questions, encouragement, challenge, conversation, and helping students connect what they are learning with something meaningful.

There is another part of this that I have been thinking about through my conference experiences.

Sometimes students do not approach faculty even when they need help. That does not necessarily mean they do not want support. Sometimes they may simply be uncomfortable or afraid to ask.

I encountered an interesting perspective through Dr. Samantha Chang's session, *Designing for Connection: Required One-on-One Meetings as a Condition for Learning* (Chang, 2026). One question from that work stayed with me:

What changes when connection is designed into a course rather than left to chance?

The session brought attention to presence, trust, belonging, access, listening, visibility, support, responsiveness, and care

(Chang, 2026). It made me think about how much time we spend designing content, assessments, activities, navigation, and access to resources while perhaps giving less attention to intentionally designing opportunities for human connection.

This does not mean every course needs mandatory individual meetings. Different courses, class sizes, faculty workloads, and student needs require different approaches. But simply writing “contact me if you need help” may not always be enough. Sometimes we need to create an opportunity for connection rather than wait for students to initiate it.

My conference experiences reinforced this in another way. Much of the information I received during presentations could have been delivered online. The slides could have been shared, presentations recorded, and many of the concepts could probably have been explained by AI.

But that would not completely reproduce the learning experience.

What I valued was also the spontaneous conversation, the questions that changed the direction of a discussion, the different perspectives in the room, and the conversations that happened between sessions. The information mattered, but so did the interaction.

Students are not simply numbers appearing in learning analytics or users accessing course content. They are real people with different experiences, responsibilities, challenges, and reasons for engaging with a course in different ways.

That is why human connection still matters.

The Enduring Value of Human Educators

So what becomes more important when AI can explain almost anything?

I do not think the answer is that faculty need to provide even more information. What becomes more valuable is what educators can do with that information and how they use their expertise to support the people who are learning it.

Faculty can bring context to a theory, challenge an assumption, notice confusion, ask a better question, share professional and research experience, interpret evidence, and provide meaningful feedback. They can listen to students, recognize when someone may be struggling, create a sense of belonging, and help students make decisions when there is no single correct answer.

AI may support faculty in some of these areas, and I do not think educators should be afraid to use that support either. But having AI available does not remove the need for educators. Instead, it may make us think more carefully about which parts of teaching become more valuable when information and explanations are already easy to access.

The faculty role has been evolving for a long time. Generative AI pushes that evolution further. Instead of being the only source of answers, faculty can help students evaluate the many answers now available to them. Instead of simply telling students what to do, they can help students develop the judgment to make decisions themselves. And instead of designing only content and assessments, we can also think about how we intentionally design opportunities for connection.

Perhaps that is the enduring value of the human educator in the age of AI. It is not that we will always have an answer that

AI cannot provide. It is that educators understand the purpose of the learning, bring disciplinary and lived experience into that learning, and develop relationships with the students they are there to teach.

AI can provide an explanation, but education has never been only about providing explanations.

9. Teaching Through Uncertainty

Key Question

How do educators lead through change?

Faculty Learning Alongside Students

When you have worked with faculty for a long time, you see many different approaches to teaching.

Some faculty have taught the same course for many years and have materials that have worked well for them over time. There is nothing necessarily wrong with that. If something has worked, there is naturally a reason to keep using it.

But I have also worked with faculty who continuously update their courses. They try a new technology, change an activity, add a new example, adjust an assessment, respond to student feedback, and learn from what happened the last time they taught the course.

They are learning alongside their students.

Generative AI makes this willingness to continue learning even more important because faculty cannot wait until AI stops changing before deciding how to respond.

It will continue changing.

New tools will appear. Existing tools will gain new capabilities.

Students will find new ways to use them. Institutional guidance will continue to evolve.

That means educators may sometimes need to teach while they are still learning themselves.

I do not think that makes someone less of an expert.

Perhaps part of expertise is being able to recognize when the environment has changed, continue learning, and make informed adjustments rather than pretending that we already know everything.

Experimentation and Adaptation

This is also one reason Scholarship of Teaching and Learning has become increasingly meaningful to me.

SoTL gives educators an opportunity to look systematically at what is happening in their own teaching. We can introduce something intentionally, observe what happens, listen to students, examine the evidence, reflect on what we are seeing, and use what we learn to inform the next iteration (Hutchings & Shulman, 1999).

As an educational technology and digital content specialist, I often work in partnership with faculty when we introduce new learning activities, technologies, or approaches. The change does not always need to be large. We can start with one small change, see how students respond, collect feedback, and look at what happened. From there, we can decide what worked, what needs to be adjusted, and what we want to try next.

My recent SoTL work developed in much the same way. We did not begin by redesigning an entire course around AI. We introduced clearer AI guidance and opportunities for transparency

and reflection into an existing learning environment. Students could choose whether or not to use AI, and the reflections gave us another way to understand some of the decisions they were making (Chan, 2026b).

The results gave me new questions as much as they gave me answers.

That is an important part of experimentation. We do not always need to prove immediately that an intervention “works.” Sometimes, we need to understand what is happening, learn from it, and use what we learn to make the next decision.

This is particularly important with AI because everything is still changing so quickly. We may implement something that works well now, but later the technology changes, students start using AI differently, or another issue comes up that we did not expect. We need to learn from what happens and continue adjusting our practice as things change.

Building Confidence Without Certainty

Faculty often want to know the right answer.

Should I allow AI?

How much should I allow?

Should I change this assessment?

What if students use AI differently from what I expected?

I understand these questions because I hear them regularly in my work. The difficulty is that with AI changing so quickly, we may not always have one clear answer.

Faculty do not need to know everything about the latest AI

tool before they can make a decision about their teaching. They know the course, the learning outcomes, the assessment, and what students are supposed to demonstrate. They can use that knowledge to establish expectations, explain those expectations to students, and make changes later if something does not work as expected.

This brings me back again to **trust, clarity, and transparency**.

When expectations are clear and faculty and students can talk openly about AI use, students have a better opportunity to explain the decisions they made. Faculty can then look more closely at whether those decisions supported the learning rather than focusing only on whether AI was used.

There is still one part of this that I continue to think about. In my own work, I sometimes find that conversations about trusting students can be more difficult than conversations about asking students to trust their faculty.

Faculty may say that they want to build trust with students. But then the next comment may be: *The student's writing suddenly looks different. The answer is too perfect. Did they use AI?*

This is where trust becomes difficult in practice. It is relatively easy to say that we want a relationship built on trust, but when something makes us question whether the work really represents the student's learning, that trust can quickly be challenged.

I do not present this as a research finding. It is something I continue to observe in my work, and it makes me think about what trust actually means when AI is involved.

If we are asking students to be transparent about their AI use, we also need to think about what happens when they are transparent with us. If students believe that admitting any AI

use will immediately make their work questionable, why would they feel comfortable telling us?

This does not mean that faculty should ignore concerns about student work or academic integrity. It means that building trust is much easier to say than to put into practice.

Closing the Faculty Gap

Across this part of the book, we have looked at several challenges faculty are facing with AI. How do we know whether student work represents their own learning? What counts as evidence of learning when AI can help produce a polished final product? What becomes more important about the faculty role when students can get explanations almost instantly? And how do we make decisions about AI when the technology keeps changing?

These are not easy questions, and I do not think the Faculty Gap means that faculty do not have enough expertise.

The gap comes from how quickly the teaching and learning environment has changed. Faculty are being asked to make decisions about AI, assessment, academic integrity, student learning, and new technologies while still teaching and supporting their students. They also need practical guidance, time, support, and opportunities to learn and try different approaches.

One thing I want faculty to remember is that they do not need to become AI experts. They need enough understanding to make educational decisions about AI based on their course, their students, their learning outcomes, and what they are actually assessing.

Sometimes it is also okay to say, *“I don’t know yet. Let me learn more.”*

We are all still learning in this changing environment. Faculty do not need to have every answer before their students do. What matters is being willing to keep learning, try something when it makes sense, make changes when something does not work, and continue helping students learn along the way.

Perhaps that is also part of the changing role of faculty. **We are not only helping students learn. Sometimes, we are learning alongside them.**

PART III

THE INSTITUTIONAL GAP

10. When Institutional Guidance Is Not Enough

Key Question

Why do institutions still struggle even when AI policies and guidance exist?

When I wrote my first *AI in Education* book, one of the issues I kept returning to was that policy alone was not enough. At that stage of the generative AI conversation, what I was seeing across higher education was often high-level institutional direction around AI, academic integrity, privacy, and responsible use. When we moved from the institutional level into an actual course, however, much of the decision-making was still left to individual faculty.

A high-level policy could tell us the institution's position on AI, but it could not always answer the practical questions students and faculty were facing. Can students use AI for this assignment? How much can they use? What should they disclose? What does responsible AI use actually look like in this particular learning activity?

That was why I kept coming back to the need for **guidance**. Students needed more than a policy, and faculty needed practical support to help translate institutional principles into teaching and learning decisions.

But AI moves quickly, and institutions are moving too.

While developing this second book, I went back to look at some of the institutions I had encountered through my conference work. Within only a few months, I could already see changes. Institutions that I had previously looked to for AI policies or high-level information were now providing faculty resources, student resources, graduate student guidance, assessment guides, course-outline suggestions, citation resources, and examples of responsible AI use.

That made me reconsider the argument I originally planned for this chapter. I was going to call it **Policy Is Not Guidance**, but that no longer fully described what I was seeing. Many institutions were already moving beyond policy and developing more guidance and practical resources. This made me realize that the gap had also changed. If policy and institutional guidance are already available, the next question is how that guidance becomes meaningful practice in an actual program, course, or assignment.

Moving Beyond Policy

The University of Waterloo is one example of this evolution. Its guidance for graduate students and supervisors does not simply say whether AI is allowed or prohibited. Students are directed to check whether AI is permitted, discuss expectations with instructors or supervisors, document AI use when appropriate, remain accountable for the accuracy and integrity of their work, and be able to explain material produced with AI (University of Waterloo, n.d.-b).

Waterloo also provides additional resources related to assessment, course outlines, citation, academic integrity, privacy, writing, and conversations between students and instructors (University of Waterloo, n.d.-a, n.d.-b).

One statement particularly caught my attention. In its academic integrity guidance, **Waterloo recommends that instructors focus on evidence of learning rather than focusing on finding misuse of artificial intelligence** (University of Waterloo, n.d.-a).

That is closely connected with the argument I have been developing throughout this book.

At my own institution, George Brown Polytechnic, I can see another version of this evolution. Institutional AI guiding principles address areas such as responsible and ethical AI, educational enhancement, transparency, data protection, AI literacy and accountability, and professional discretion for educators. These principles are also connected with AI tools, learning and training resources, policies, and additional guidance for the institutional community (George Brown Polytechnic, 2026a, 2026b).

Many institutions already have AI policies, and they are also developing more guidance and practical resources.

The next question is how we bring that guidance into actual teaching and learning practice.

From Policy to Guidance

Policies are still important. Institutions need clear expectations around academic integrity, privacy, security, ethical AI use, appropriate technology use, and responsibility.

But a policy cannot tell us exactly what to do in every course or assignment.

An institution may allow faculty to decide how AI fits within their teaching, but the faculty member still needs to make that decision. Should students use AI for this assignment? If yes, what can they use it for? What should they still do themselves? What should they disclose?

Students face the same challenge. They may be told to follow the AI expectations set by their instructor, but they still need to understand what those expectations mean for the work they are actually doing.

This is where I see the difference between policy and guidance. Policy gives us the overall expectations and boundaries. Guidance helps us understand what those expectations mean when we need to make a decision.

Even with good guidance, there will still be situations that are not covered. An institutional website can provide examples, templates, resources, workshops, and recommendations, but eventually someone still needs to decide what responsible AI use looks like in a particular learning situation.

That decision may happen within a program, a course, or an individual assignment. Faculty may need to make the decision with their students, and sometimes students will need to make their own decisions based on the guidance they have received.

This is where I see the next part of the Institutional Gap. Having

AI policy and guidance is important, but the work does not stop there.

We still need to bring that guidance into actual teaching and learning practice.

11. Preparing for the Next Phase

Key Question

How do institutions move from having AI guidance to making that guidance easier to use in practice?

From Having Guidance to Using Guidance

Having institutional AI guidance is an important step, but publishing guidance does not mean that the work is finished.

When I started looking more closely at institutional AI resources for this book, I noticed how quickly things had changed. Earlier in the generative AI era, one of my concerns was that institutions often had high-level policies but not enough practical guidance. Now, many institutions have developed much more.

It is actually a positive change but it also creates another challenge.

The University of Toronto is one example that made me think about this differently. When I looked through its AI resources, there was a lot of information available. Graduate students and supervisors can find guidance related to research, writing, editing, disclosure, academic integrity, privacy, copyright, publication, and other uses of generative AI. There are also FAQs, classroom resources, syllabus and assignment guidance, and

links to additional policies and supports (University of Toronto, n.d.; University of Toronto, School of Graduate Studies, 2026).

U of T also has detailed AI guidance for administrative staff. That guidance covers responsible AI use, acceptable and prohibited uses, privacy and data security, transparency, bias, procurement, training, and other areas (University of Toronto, 2025b).

Looking at all of these guidance materials now, I do not think the problem is that there is not enough information.

My next question is: How much information can we expect someone to read before they know what to do when a particular situation comes up?

A faculty member may be preparing an assignment and simply want to know how to communicate appropriate AI use to students. A student may be working on an assignment and want to know whether AI can be used for a particular part of the work. They may not know which institutional page to visit, which guideline applies, or how a high-level principle connects with the decision they need to make at that moment.

This is where my learning design background influences how I look at the problem.

Having information available is not the same as making that information easy to use.

Perhaps the next phase of institutional AI guidance is not simply creating more guidance. It is thinking about how we make the guidance we already have more actionable.

Making Guidance Actionable

What would actionable guidance look like?

I do not think the solution is another long document that summarizes all the other documents. That may simply give people one more resource to read.

Instead, we can think about where people actually need the guidance.

For faculty, that may mean practical templates they can adapt directly within their courses or assignments. It may mean examples from different disciplines, short decision guides, professional learning opportunities, or someone they can contact when they are not sure what to do.

For students, guidance may need to appear closer to the learning activity itself. If an assignment has particular expectations around AI, those expectations should be available where students are actually reading the assignment rather than requiring them to search through several institutional webpages to figure out what applies.

This does not mean that detailed institutional resources are unnecessary. We still need them. Privacy, security, academic integrity, research ethics, copyright, accessibility, and other areas require more detailed information.

But not everyone needs every piece of information at the same time. The real challenge is helping people get the right level of guidance at the time they need it.

From Institutional Guidance to Classroom Practice

This is where different parts of the institution can work together. The institution can establish the overall principles and expectations, but those principles still need to reach faculty and students in ways they can actually use.

Because I work within Teaching and Learning, this is where I see an important role for our work. Our focus is not to create another institutional AI policy. We help faculty and students understand what those broader expectations mean for teaching and learning and provide practical ways to apply them.

One example from my own work is **AI Literacy for All**, which was developed to provide a shared foundation for both faculty and students. It includes an **AI Faculty Learning Guide** and **A Practical Guide for Students**, along with interactive learning activities, guidance around responsible AI use and transparency, and resources that faculty can bring into their existing courses.

Having both faculty and student resources is important. If we want faculty and students to have clearer conversations about AI, we cannot focus only on helping faculty understand responsible AI use and then expect students to already know the same things. Both groups need opportunities to develop a shared foundation.

I have also been developing **Practical AI Guidance** for teaching and learning. This is a growing collection of faculty resources designed to help move AI guidance closer to day-to-day teaching practice. The resources include adaptable prompts, implementation examples, and practical guidance related to areas such as assessment, accessibility, student engagement, transparency, privacy, AI detection, and responsible AI use.

The purpose is not to replace faculty judgment. Faculty still need to decide what makes sense for their discipline, course, learning outcomes, assessments, and students. The resources give them somewhere to start rather than expecting every faculty member to interpret all of the available institutional information and develop everything from the beginning.

This connects directly with what I noticed when looking at the University of Toronto example. There may already be a great deal of useful AI information available. The next challenge is helping people move from finding information to knowing what to do with it.

The institution can establish the foundation. Teaching and Learning can help turn that foundation into practical learning opportunities and resources. Faculty can adapt those resources to their own teaching context, and students can use clearer guidance when making decisions about AI.

Each level has a different role, but they need to connect.

This is where I see another part of the AI Guidance Gap. The guidance may already exist, but if people cannot easily find it, understand it, or apply it when they need to make a decision, the gap remains.

12. Building a Guidance Ecosystem

Next Step

We can tell everyone to **use AI responsibly**, but responsible AI use may look different depending on what students are learning, what faculty are teaching, and what students are being asked to demonstrate. This connects with Cao's (2025) work on responsive AI policies, which emphasizes the importance of considering the learning context and learning goals when making decisions about AI use.

My next question is how that contextual approach connects with the different levels of guidance across an institution. Policy, institutional guidance, teaching and learning support, faculty practice, and student learning all have different purposes, but they need to work together.

This leads me to the idea of building a **guidance ecosystem**.

Building a Guidance Ecosystem

In a guidance ecosystem, the institution establishes the overall foundation and expectations. Guidance helps turn that foundation into something faculty and students can understand and use. Faculty then apply and adapt that guidance within their courses, while students use it to help make decisions about their own learning.

But the process should not stop after the guidance is imple-

mented. What faculty and students experience in practice can help us understand what works, what remains confusing, and what new questions are appearing. We can then use what we learn to improve the guidance again.

I see this becoming an ongoing cycle:

Policy → Institutional Guidance → Contextualized Guidance → Practice → Reflection → Better Guidance

Policy = institutional policies, principles, responsibilities, and boundaries.

Institutional Guidance = the broader AI guidance available on institutional websites, including responsible use, privacy, academic integrity, transparency, approved tools, etc.

Contextualized Guidance = bringing that broader guidance into a particular program, course, assignment, or assessment.

For example: *What does responsible AI use mean for this specific discussion board or writing assignment?*

Practice = faculty and students actually applying the guidance.

Reflection = looking at what happened. What worked? What remained confusing? What did students do? What did faculty learn?

Better Guidance = using what was learned to revise the guidance for the next time.

Then the cycle starts again.

This is particularly important with AI because things continue to change quickly. A new AI capability may appear, students may begin using AI in ways we did not expect, or faculty may try an approach and discover that it does not work the way they intended. New questions about privacy, accessibility, assessment, ethics, and learning will also continue to appear.

That means we cannot simply publish AI guidance and consider the work finished.

We need to continue listening to the people who are actually using it. Faculty can tell us where guidance becomes difficult to apply in a real course. Students can tell us where expectations are still confusing. Teaching and Learning teams can identify areas where faculty continue to need practical support. Research can also help us understand what is happening as these practices develop.

What we learn from those experiences can then help us improve the next version.

This is why I think the next phase of institutional AI work is not simply about having more guidance. We already have a lot of information available in many institutions.

The next step is to make that guidance easier to find, easier to understand, easier to apply, and easier to improve as we continue learning.

That is how guidance moves from information into practice,

and how the different parts of an institution can begin working together as a guidance ecosystem.

PART IV

CLOSING THE GUIDANCE GAP

13. A Framework for Guidance

Key Question

What does meaningful AI guidance actually need?

Throughout this book, I have looked at the AI Guidance Gap from three different perspectives: students, faculty, and institutions. The challenges are different at each level, but they continue to connect with one another.

Students have access to AI but may not know what is expected, whether their use is appropriate, or whether using AI is actually helping them learn. Faculty are being asked to make decisions about authenticity, assessment, AI use, and changing teaching practices while the technology continues to evolve. Institutions are developing policies, principles, resources, and guidance, but that guidance still needs to become meaningful practice within individual programs, courses, and assignments.

After looking across these different experiences, I keep returning to five ideas:

Clarity. Transparency. Reflection. Judgment. Trust.

These ideas have appeared throughout my own practice, conference presentations, conversations with students and faculty, AI literacy work, and ongoing Scholarship of Teaching and Learning inquiry. I do not present them as a finished or validated model. Instead, I see them as an emerging framework for thinking about what meaningful AI guidance may require.

For me, the relationship between them is important:

Clarity → Transparency → Reflection →
Judgment → Trust

Each part supports the next, but they are also connected and can continue to influence one another.

Clarity: What Is Expected?

Guidance begins with clarity.

Students should not have to guess whether AI can be used, what kinds of use are appropriate, what needs to remain their own work, or what they are expected to disclose. Throughout this book, I have returned to the question: **What are students actually supposed to learn?**

When the learning purpose is clear, we are in a better position to explain what role AI should have. AI may support one part of an activity while interfering with another. Different assignments may therefore have different expectations, even within the same course.

Clarity does not mean creating a rule for every possible situation. It means giving people enough information to understand the purpose, the boundaries, and where to ask when something is unclear (Winkelmes, 2023; Winkelmes et al., 2019).

Transparency: What Happened?

Once expectations are clear, transparency gives students and faculty a way to talk about what actually happened.

This is an area that has become particularly important through my own SoTL inquiry. In that work, students were not simply asked whether they used AI. They were asked whether they used it, how they used it, and importantly, **why they made that decision**.

That small difference gave me another way to see parts of the learning process that may otherwise have remained hidden.

Transparency, therefore, does not need to mean surveillance. It does not mean collecting every prompt or trying to monitor everything a student does with AI.

Instead, transparency can create an opportunity to make relevant decisions visible. What role did AI play? Why was it used? What did the student do with the output? What remained the student's responsibility? (Weaver, 2024).

My Phase I research does not establish that transparency causes student agency or better learning. What it did show me was that students' transparency statements and reflections could reveal deliberate decision-making, AI use as support rather than substitution, attention to creative ownership, and reflection on learning and authorship (Chan, 2026b).

That made me think about transparency differently. It can be more than disclosure and it can create an opportunity for learning.

Reflection: What Did This Mean for My Learning?

Transparency can tell us what happened. Reflection allows us to ask what that experience meant.

- Did AI help me understand something?
- Did it give me another perspective?
- Did I accept its suggestion because it was useful, or because it was easier?
- Did I reject something AI produced? Why?
- Did AI help me continue thinking, or did it remove the thinking I was supposed to do?
- Would I make the same decision next time?

These questions move AI literacy beyond knowing the capabilities and risks of a technology. They encourage learners to think about their own behaviour and the relationship between AI use and learning (Anders & Dux Speltz, 2025; Ratnayake et al., 2024).

Reflection also applies to faculty and institutions. Faculty can reflect on whether an AI expectation worked as intended. Institutions can look at where guidance is being used, where people remain confused, and what needs to change.

This connects with the cycle I introduced earlier:

Policy → Institutional Guidance →
Contextualized Guidance → Practice →
Reflection → Better Guidance

Reflection is what allows guidance to evolve rather than becoming a static document that gradually falls behind practice.

Judgment: What Should I Do?

Eventually, guidance cannot answer every question.

A new AI capability will appear. A student will use a tool in a way the faculty member did not anticipate. An assignment will create a situation that was not covered by the course guidance. Institutional resources will not have an example for every discipline and every learning activity.

Someone will still have to make a decision. That is where judgment becomes important.

Students need to develop the ability to decide when AI can support their learning, when they need to verify an answer, when they should try something themselves first, and when AI may be replacing something they need to learn. This connects with the broader educational idea of developing learners' capacity for evaluative judgment (Tai et al., 2018).

Faculty need judgment when deciding what role AI should play within an assessment, how much evidence of process is needed, and whether a student's AI use supported or interfered with the intended learning.

Institutions also need judgment when balancing consistency with professional discretion and disciplinary differences.

The goal of guidance, therefore, cannot be to remove every decision.

Perhaps good guidance should help people become better at making decisions.

This connects with UNESCO's (2026) recent call for educators to determine if, when, and how AI should be used according to educational purpose and context.

Trust: What Allows the Process to Work?

Trust has appeared throughout this book, but I intentionally place it at the end of this framework.

I do not think trust means simply telling faculty to trust students or asking students to trust their professors.

Trust needs conditions that allow it to develop.

When expectations are clear, students have less need to guess. When AI use can be discussed transparently, faculty have another window into the learning process. When students reflect on their decisions, they have an opportunity to explain their reasoning. When both students and faculty develop judgment, the conversation can move beyond simply asking whether AI was used.

That creates a stronger foundation for trust.

Trust still does not remove accountability. Students remain responsible for the work they submit. Faculty remain responsible for assessing learning. Institutions remain responsible for establishing appropriate principles and support.

But accountability and trust do not have to work against each other.

Students need to know that being transparent about responsible AI use will not automatically make their work suspicious. Faculty need confidence that students can learn to make responsible decisions. Institutions need to support faculty in exercising professional judgment while still providing a shared foundation.

For me, trust is therefore not the starting assumption that everything is fine.

It is something that can grow when clarity, transparency, reflection, and judgment are intentionally supported.

14. From Guidance to Practice

Key Question

What happens when we move from talking about AI guidance to actually designing it into learning?

Throughout this book, I have argued that guidance needs to move closer to the actual learning situation. Institutional principles are important, but eventually students and faculty need to make decisions within a particular course, assignment, or learning activity.

This is where guidance becomes practice.

From Phase I to Phase II

My Phase I Scholarship of Teaching and Learning inquiry gave me an opportunity to look at what happened when students were provided with AI guidance but still had the choice to use or not use AI.

Students were asked to explain whether they used AI, how they used it, and importantly, why they made that decision. Through their transparency statements and reflections, I started to see recurring patterns related to deliberate decision-making, using AI as support rather than substitution, maintaining creative ownership, and reflecting on learning and authorship (Chan, 2026b).

Phase I helped make some of the students' thinking and decision-making visible.

But it also left me with another question:

What happens when the guidance itself becomes more specific to the course and assignment?

That question leads into Phase II.

Phase II begins in September 2026, with implementation and data collection taking place from September to December across four online Film Studies sections. This time, the guidance will move from broader AI guidance toward more course- and assignment-specific guidance. The purpose is not simply to give students more information. I want to make the connection between AI use, the assignment, and the learning purpose clearer.

From January to April 2027, I will move into the next stage of the study: analyzing the Phase II data, comparing the patterns with Phase I, and developing the next paper from what emerges.

From General Guidance to Context-Specific Guidance

In Phase II, I am putting this into practice through a progression:

General Guidance → AI Guidance Designer → Context-Specific Guidance

Phase II Implementation	Comparative Inquiry
4 online Film Studies sections	Compare Phase I and Phase II patterns
~200 students anticipated	Examine how students justify AI decisions
Course- and assignment-specific AI guidance	Explore AI roles, authorship, and ownership
Guidance embedded in assessment	Examine changes in reasoning and reflection
Continued AI transparency statements and reflections	

Research question: How might moving from general to context-specific AI guidance shape students' reasoning and reflection about AI use?

The general student guidance provides a starting point. It explains that AI can support activities such as brainstorming, planning, clarifying concepts, and improving clarity, while students remain responsible for their own thinking, decisions, and final work. It also introduces expectations around transparency and reflection.

The next step is to bring that guidance closer to what students are actually learning.

To support this process, I developed an **AI Guidance Designer**, a custom AI agent that helps translate broader AI guidance into a particular course and assessment context. I have developed versions of the agent in both ChatGPT and Microsoft

Copilot so that the guidance-design approach is not limited to a single AI environment. The purpose is not to have AI decide what students should or should not be allowed to do. Instead, the agent supports the design process by helping consider the learning purpose, assessment requirements, appropriate roles for AI, and what students still need to demonstrate themselves.

For Film Studies, this means looking at the purpose of the course and the assignments before deciding where AI fits. Students need to demonstrate their own film viewing, use of film terminology, interpretation, selection of evidence, and creative judgment. AI may help them organize a Tech Sheet, review terminology, develop a checklist, or receive feedback on clarity, but it should not decide which evidence matters, invent evidence, perform the analysis, or make the creative decisions for them.

The underlying idea has not changed. AI can support learning without replacing the learning students need to demonstrate. What changes in Phase II is how explicitly the guidance is connected with the course, assignment, and learning purpose.

I will continue looking at students' transparency statements and reflections and compare what I see with Phase I. Will students make different decisions? Will they explain their decisions differently? Will clearer and more specific guidance make any difference at all? **At this point, I do not know what the results will show, and that is part of the research process.**

I am not conducting Phase II simply to prove the argument I have made in this book. I want to see what actually happens. The results may support some of my thinking, challenge it, or show me something I have not considered yet.

That is also part of moving from guidance to practice. We develop something, put it into practice, observe what happens,

reflect on what we learn, and then decide what should change next.

What I learn from Phase II will become part of the next stage of this work. Just as the guidance itself can change through practice and reflection, my understanding of the AI Guidance Gap may continue to change as I collect more evidence and learn from students.

This book captures where my thinking is now. The research will continue from here.

Acknowledgements

This book developed through many conversations, collaborations, presentations, and learning experiences across my work in teaching and learning. It is a practice-informed book grounded in professional experience and connected with relevant scholarship. The ideas grew from my work with faculty and students, conference conversations, AI literacy initiatives, and ongoing Scholarship of Teaching and Learning inquiry.

I would like to thank the faculty, students, colleagues, and conference participants who shared their questions, experiences, and perspectives with me. Many of the ideas explored in this book grew from listening to the questions people continue to ask about AI: what is appropriate, what supports learning, what should remain human, and how we can create clearer expectations without losing trust.

I am especially grateful to the students who have shared their perspectives through presentations, conversations, and my ongoing Scholarship of Teaching and Learning work. Their voices have continued to remind me that students are not simply asking whether they can use AI. They are also asking for clarity, guidance, and trust.

I would also like to acknowledge the faculty members and colleagues I have worked with while developing and implementing practical approaches to AI in teaching and learning. Working directly with faculty has helped me understand where institutional guidance meets the reality of an individual course, assignment, or learning activity.

The development of **AI Literacy for All**, including faculty- and student-facing learning resources, has also informed this book. I am grateful to Silvia Caicedo for collaborating with me on this

work and bringing a faculty perspective to the development of these resources. Bringing together faculty and educational technology perspectives helped strengthen the practical focus of the work and reinforced something I return to throughout this book: guidance needs to make sense to the people who will actually use it.

I am grateful to everyone who engaged with my first *AI in Education* book, attended any of my presentations, asked a question, challenged an idea, or continued a conversation afterward. Those interactions encouraged me to keep thinking, researching, and sharing openly.

Generative AI, primarily ChatGPT, was used during the development of this book to support brainstorming, organization, review, editing, and refinement of the writing. I used AI to help identify repetition, consider alternative ways of organizing ideas, review clarity, and challenge parts of the manuscript from different perspectives.

The ideas, professional experiences, practice-based examples, research interpretations, and final decisions presented in this book remain my own. AI-generated suggestions were reviewed, revised, rejected, or rewritten as appropriate, and I remained responsible for the final content of the book.

This process also became part of the reflection explored throughout the book: AI can support the work, but human judgment remains necessary in deciding what belongs, what needs to change, and what ultimately represents the author's intended meaning.

Finally, I want to acknowledge that this book was written while I am still learning. My research is continuing, my practice is evolving, and the ideas presented here may continue to develop as I learn from students, faculty, colleagues, research, and future practice.

The AI Guidance Gap is therefore not the result of one project or one conversation. It grew from many experiences that gradually led me back to the same question:

When answers are everywhere, what do learners still need from us?

OER Licensing

OER Licensing

© 2026 Evelyn Chan

This book is licensed under a Creative Commons Attribution 4.0 International License (CC BY 4.0).

You are free to:

- **Share** — copy and redistribute the material in any medium or format
- **Adapt** — remix, transform, and build upon the material for any purpose, even commercially

Under the following terms:

- **Attribution** — You must give appropriate credit, provide a link to the license, and indicate if changes were made.

To view the full license, visit: [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

If you reuse or adapt this material, please include the following attribution:

“Adapted from *The AI Guidance Gap: Teaching, Learning, and Trust When Answers Are Everywhere* by Evelyn Chan, licensed under CC BY 4.0.”

If you choose to adapt or reuse this book, you are welcome to share how you are using it. Hearing how these materials support teaching and learning helps inform future revisions and improvements.

References

References

Core scholarship

Anders, A. D., & Dux Speltz, E. (2025). Developing generative AI literacies through self-regulated learning: A human-centered approach. *Computers and Education: Artificial Intelligence*, 9, 100482. <https://doi.org/10.1016/j.caeai.2025.100482>

Brandt, W. C. (2024). *Measuring student success skills: A review of the literature on student agency*. National Center for the Improvement of Educational Assessment. <https://files.eric.ed.gov/fulltext/ED671663.pdf>

CAST. (2024). *CAST Universal Design for Learning Guidelines version 3.0*. <https://udlguidelines.cast.org>

Desai, H. (2026). What's worth measuring? The future of assessment in the AI age. In UNESCO, *The algorithm in the room: Seeing and confronting the implications of AI for the future of education* (pp. 64–68). UNESCO. <https://doi.org/10.54675/PMRO5023>

Fu, Y., & Weng, Z. (2024). Navigating the ethical terrain of AI in education: A systematic review on framing responsible human-centered AI practices. *Computers and Education: Artificial Intelligence*, 7, 100306. <https://doi.org/10.1016/j.caeai.2024.100306>

Gulikers, J. T. M., Bastiaens, T. J., & Kirschner, P. A. (2004). A five-dimensional framework for authentic assessment. *Educa-*

tional Technology Research and Development, 52(3), 67–86.
<https://doi.org/10.1007/BF02504676>

Hackl, V., Müller, A. E., & Sailer, M. (2026). The AI literacy heptagon: A structured approach to AI literacy in higher education. *Computers and Education: Artificial Intelligence*, 10, 100540.
<https://doi.org/10.1016/j.caeai.2026.100540>

Hutchings, P., & Shulman, L. S. (1999). The scholarship of teaching: New elaborations, new developments. *Change: The Magazine of Higher Learning*, 31(5), 10–15. <https://doi.org/10.1080/00091389909604218>

Laupichler, M. C., Aster, A., Schirch, J., & Raupach, T. (2022). Artificial intelligence literacy in higher and adult education: A scoping literature review. *Computers and Education: Artificial Intelligence*, 3, 100101. <https://doi.org/10.1016/j.caeai.2022.100101>

Long, D., & Magerko, B. (2020). What is AI literacy? Competencies and design considerations. In *Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems* (pp. 1–16). Association for Computing Machinery. <https://doi.org/10.1145/3313831.3376727>

Muellauer, C. (2026). When my students told me I ‘sound like AI’. In UNESCO, *The algorithm in the room: Seeing and confronting the implications of AI for the future of education* (pp. 69–72). UNESCO. <https://doi.org/10.54675/PMRO5023>

Perkins, G. (2026). Resilient assessment in the age of AI: Authentic design and the case for verbal examinations in business education. *Assessment & Evaluation in Higher Education*. Advance online publication. <https://doi.org/10.1080/02602938.2026.2644516>

Ratnayake, A., Bansal, A., Wong, N., Saseetharan, T., Prompiengchai, S., Jenne, A., Thiagavel, J., & Ashok, A. (2024). All “wrapped” up in reflection: Supporting metacognitive aware-

ness to promote students' self-regulated learning. *Journal of Microbiology & Biology Education*, 25(1), e00103-23. <https://doi.org/10.1128/jmbe.00103-23>

Tai, J., Ajjawi, R., Boud, D., Dawson, P., & Panadero, E. (2018). Developing evaluative judgement: Enabling students to make decisions about the quality of work. *Higher Education*, 76, 467–481. <https://doi.org/10.1007/s10734-017-0220-3>

Wang, G., Wang, W., Yang, D., & Ren, J. (2026). Generative AI, cognitive offloading, and learner agency in higher education: A scoping review. *Behavioral Sciences*, 16(7), 1150. <https://doi.org/10.3390/bs16071150>

Weaver, K. D. (2024). The Artificial Intelligence Disclosure (AID) Framework: An introduction. *College & Research Libraries News*, 85(10), 407–411. <https://doi.org/10.5860/crln.85.10.407>

Winkelmes, M. (2023). Introduction to transparency in learning and teaching. *Perspectives in Learning*, 20(1), 4–12. <https://csuepress.columbusstate.edu/pil/vol20/iss1/2>

Winkelmes, M. A., Boye, A., & Tapp, S. (Eds.). (2019). *Transparent design in higher education teaching and leadership: A guide to implementing the transparency framework institution-wide to improve learning and retention*. Stylus. <https://doi.org/10.4324/9781003448396>

UNESCO. (2026). *The algorithm in the room: Seeing and confronting the implications of AI for the future of education*. UNESCO. <https://doi.org/10.54675/PMRO5023>

Practice, Research, Presentations, and Professional Learning

Cao, L., Dede, C., Bonk, C., Zhao, Y., & Mishra, P. (2026, August

29). *Invisible watermarks, visible consequences: AI mandates and academic integrity* (Episode 277) [Video podcast episode]. *Silver Lining for Learning*. https://www.linkedin.com/posts/silver-lining-for-learning-podcast_aiineducation-academicintegrity-euaiaact-activity-7496322180910108673-U1Do

Cao, L. (2025, October 23). *Why universal AI rules don't work: How to build responsive AI policies with your students* [Webinar]. Contact North | Contact Nord. <https://teachonline.ca/webinars/why-universal-ai-rules-dont-work-how-to-build-responsive-ai-policies-with-your-students/>

Chan, E. (2025, July). *Adapting Film Studies courses for diverse learning environments using D2L Brightspace* [Online presentation]. D2L Fusion On Demand. <https://youtu.be/F9tIHk1jHAM>

Chan, E. (2026a, June 17). *AI in education: Ethics, literacy, and interactive design* [Conference presentation]. Society for Teaching and Learning in Higher Education (STLHE) Conference, Hamilton, ON, Canada.

Chan, E. (2026b). *From compliance to reflection: AI transparency statements as a pedagogical strategy supporting student agency in higher education*. George Brown Polytechnic, Teaching and Learning Exchange. <https://www.georgebrown.ca/teaching-and-learning-exchange/scholarship-of-teaching-and-learning-sotl/sotl-projects>

Chan, E. (2026c). *From compliance to reflection: Rethinking AI transparency as a learning strategy* [SoTL Spotlight]. George Brown Polytechnic, Teaching and Learning Exchange. <https://www.georgebrown.ca/teaching-and-learning-exchange/scholarship-of-teaching-and-learning-sotl/sotl-spotlight>

Chan, E. (2026d, October 30–31). *Learning across states of awareness: Reflective inquiry and educational leadership* [Accepted paper presentation]. Centre for Leadership and Diversity 7th Annual International Conference, Ontario Institute for Studies in Education, University of Toronto, Toronto, ON, Canada.

Chan, E., & Gaidosch, B. (2026). *From guidance to practice: Starting small with AI integration—A Film Studies case study* [Conference presentation]. Evolved Conference, George Brown Polytechnic, Toronto, ON, Canada. <https://ecampusontario.pressbooks.pub/evolved2026/chapter/from-guidance-to-practice-starting-small-with-ai-integration-a-film-studies-case-study/>

Chang, S. (2026). *Designing for connection: Required one-on-one meetings as a condition for learning* [Roundtable presentation]. Teaching and Learning Innovations Conference, University of Guelph.

Maheux-Pelletier, G., Akobi, I. O., & Kapasa, A. M. (2026, June 18). *Small teaching decisions, big shifts: Student voices on adapting teaching and learning in the age of GenAI* [Keynote address]. Society for Teaching and Learning in Higher Education (STLHE) Conference, Hamilton, ON, Canada.

Institutional and open resources

George Brown Polytechnic. (2026a). *George Brown Polytechnic's position on artificial intelligence*. <https://www.georgebrown.ca/ai-centre-of-excellence/guidance-policies-responsible-ai/george-brown-polytechnics-position-on-artificial-intelligence>

George Brown Polytechnic. (2026b). *Guidance, policies &*

responsible AI. <https://www.georgebrown.ca/ai-centre-of-excellence/guidance-policies-responsible-ai>

George Brown Polytechnic. (n.d.). *Academic integrity (Generative AI)*. <https://www.georgebrown.ca/teaching-and-learning-exchange/educational-technology/teaching-with-ai/generative-ai/academic-integrity>

Quality Matters. (n.d.). *Higher Ed Course Design Rubric*. <https://www.qualitymatters.org/qa-resources/rubric-standards/higher-ed-rubric>

The Learning Portal, College Libraries Ontario. (n.d.). *Generative AI decision flowchart*. The Learning Portal. <https://tlp-lpa.ca/learning-online/GenAI>

University of Toronto. (2025a). *Coursework and GenAI: A practical guide for students* [Open educational resource]. Open UToronto. <https://ocw.utoronto.ca/innovation-projects/genai-literacy-course-modules/>

University of Toronto. (2025b, October 27). *Guideline on the use of artificial intelligence by University administrative staff*. <https://people.utoronto.ca/employees/administrative-staff-and-ai/guideline-on-the-use-of-artificial-intelligence-by-university-administrative-staff/>

University of Toronto, School of Graduate Studies. (2026). *Guidance on the appropriate use of generative artificial intelligence for graduate academic milestones*. <https://www.sgs.utoronto.ca/about/guidance-on-the-use-of-generative-artificial-intelligence/>

University of Toronto. (n.d.). *Guidelines: Toward an AI-ready university*. <https://ai.utoronto.ca/guidelines/>

University of Waterloo. (n.d.-a). *Artificial intelligence and Chat-*

GPT. <https://uwaterloo.ca/academic-integrity/artificial-intelligence-and-chatgpt>

University of Waterloo. (n.d.-b). *Generative artificial intelligence (GenAI) guidance for graduate students and supervisors*. <https://uwaterloo.ca/current-graduate-students/academics/governing-rules-and-regulations/generative-artificial-intelligence-genai-guidance-graduate>

About the Author

Evelyn Chan works in educational technology and digital content in postsecondary education, where she supports faculty in designing clear, accessible, inclusive, and meaningful learning experiences across online and technology-enabled environments. Her work focuses on how educators use digital tools in practice, with particular attention to course design, Universal Design for Learning, AI literacy, responsible AI use, and human judgment.

Her recent work increasingly explores what happens when artificial intelligence becomes part of everyday teaching and learning. Through faculty development, course design, AI literacy initiatives, conference presentations, and Scholarship of Teaching and Learning research, Evelyn examines how students and educators make decisions about AI and what kinds of guidance can support those decisions.

Her ongoing SoTL research explores AI transparency statements, student decision-making, reflection, agency, and responsible AI use within authentic learning environments. This work has contributed to a broader interest in moving conversations about AI beyond detection and compliance toward **clarity, transparency, reflection, judgment, and trust.**

Through her role, Evelyn collaborates with faculty, students, accessibility specialists, and teaching and learning colleagues to develop practical approaches to AI integration. Her recent work includes **AI Literacy for All**, faculty- and student-facing AI literacy resources, practical AI guidance for teaching and learning, and professional learning initiatives that help educators translate broader AI principles into decisions they can apply within their own courses and assessments.

Evelyn is the author of a growing collection of open Pressbooks on course quality, inclusive design, digital learning, accessibility, and artificial intelligence. Her earlier work established a foundation in quality course design, Universal Design for Learning, and meaningful digital learning experiences. Her AI-focused books extend that work into questions of AI literacy, responsible use, student agency, assessment, human judgment, and the relationship between institutional guidance and classroom practice.

Her approach remains strongly practice-informed. She writes from the perspective of someone who works directly with faculty and students, experiments with new approaches, learns from implementation, and uses research and reflection to continue developing her practice.

Evelyn believes that AI should support human judgment rather than replace it. As AI makes answers and assistance increasingly easy to access, her work asks a question that continues to shape her research and practice:

When answers are everywhere, what do learners still need from us?

Evelyn lives and works in Ontario, Canada and continues her professional and scholarly development through collaboration, open educational practice, research, and ongoing doctoral study.

ORCID profile (Evelyn Chan)

Google Scholar profile (Evelyn Chan)

If you would like to connect or continue the conversation, you can find her on LinkedIn.

Additional information about publications, presentations, and current projects is available at **evelynchan.ca**.

Related Books by the Author

This book is part of a body of open Pressbooks by Evelyn Chan focused on quality assurance, inclusive design, interaction, and AI ethics in digital learning.



Selected titles:

- **The QM Learning Path: Faculty Development and Course Design for Quality Online Learning**
eCampusOntario Pressbooks
- **UDL in Action: Designing Inclusive and Interactive Digital Learning in Higher Education**
eCampusOntario Pressbooks
- **The Interactive Learning Path: Building Meaningful Digital Learning Experiences**
eCampusOntario Pressbooks
- **AI in Education: Ethics, Literacy, and Interactive Design**
eCampusOntario Pressbooks
- **Accessible Digital Learning: Practical Design for Equity and Inclusion**
eCampusOntario Pressbooks

Guides and Practice-Based Works

- **UDL Quick Start Guide for Faculty**
eCampusOntario Pressbooks
- **The QM Learning Path: A Faculty Quick Start Guide**
eCampusOntario Pressbooks
- **The QM Learning Path: Scholarly and Practice-Based Works**
eCampusOntario Pressbooks
- **Access, Equity, and Digital Learning: Scholarly and**

Practice-Based Works

eCampusOntario Pressbooks

Additional publications and resources are available on **Evelyn Chan's website**.



ecampus Ontario evelynchan.ca