Impacts and Implications of Generative AI on Course Materials

David Wiley, PhD

Times

Are

Changing

This

AI

Contraption

Textbooks

Assumptions

Challenged

Take

A

Couple

Think

About what's

Coming

Everything Old is New

Wild Claims about Efficacy, 1500s Edition

- In the 1500s, a group of humanists includuing Joachimus Fortius Ringelbergius claimed that the **special “methods”** that underpinned their textbooks would enable the young to **learn “Latin in eight months, Greek in twenty days, astronomy in eight or ten days, philosophy and music in a month or less”.**

The Textbook Authoring Treadmill, 1500s Edition

- Rudolph Goclenius complained about the quality of his textbook on metaphysics, which he finished in the late 1500s.

- He acknowledged “the messiness and incoherence of his own book, no doubt produced in haste from his teaching notes. **This was textbook writing as we know it**: deadline-driven, compilatory, and not so much intellectually ambitious as haplessly apologetic.”

Undercutting the Used Market, 1600s Edition

- Judah Monis wrote a Hebrew grammar that was prohibitively expensive to print because of its Hebrew font.

- “But the expense was borne by the College [Harvard], **which then required every student to buy a new copy** (rather than a used one from another student).”

Look Back to Look Ahead

Before the Printing Press

- Books were ***incredibly expensive*** (had to be copied by hand)
- Faculty had to assume no student had the text
- Lead faculty to adopt "**dictation**" pedagogy
- Dictation was widely believed to have pedagogical merit, as "the act of copying out a text was often considered an essential part of mastering it," going back to Demosthenes and St. Jerome.

Before the Printing Press

- Students saw dictation as "**a cheaper way of procuring oneself a classroom text**."

- Consequently, a ban on dictations by Arts Faculty at the University of Paris in 1355 "anticipated **vehement student resistance** to the ban."

- "In Paris the ban on dictation of 1355 coincided with the development of commercial stationers who rented out exemplars of texts for classroom use and thus offered a reasonably priced alternative to taking down the text under dictation."

After the Printing Press (and the Bans)

- **Faculty continued dictations anyway**.

- But… how do you continue dictation when students are bringing copies of the textbooks with them to class?

- “Around 1500, German universities developed the practice of producing inexpensive printed copies of texts that were to be treated in lectures. They were printed… with **larger interlinear spacing and wide margins**.”

After the Printing Press (and the Bans)

- Leonhardt calls these “lecture texts… an inexpensive copy of the text under study, mass-produced, as it were, **with the ‘blanks’ to be filled in**.”

- “In many cases, several different [modern] libraries conserve copies of the same printed text, whereby the handwritten notes found in each copy are **literally identical or very nearly so**.”

After the Printing Press (and the Bans)

- Instead of dictating the text itself, **faculty dictated notes for students to write in the margins of the text.**

- (Fast forward 400 years, and all too often faculty are still standing at the front of the room essentially reading out notes for students to take down in the margins of the book they should have read before coming to class.)

After the Printing Press (and the Bans)

- The transition from not having course materials to having course materials didn't really change pedagogy.

- If not even the **printing press** could meaningfully change faculty behavior, can **generative AI**?

Why Do We Have Course Materials?

Access to Expertise

- Imagine you lived before the public availability of books, and wanted to learn something new.
- Your only real option was to find an expert to teach you.
- However, experts are **rare**, their time is **expensive**, and they're **busy**.

Access to Expertise

- Course materials are built on three assumptions:

- Experts are **scarce**
- Experts are **expensive**
- Experts are **slow**

Access to Expertise

- Printed textbooks **capture** a SME's expertise so another person can access it at another time and a different location.

- Online resources **capture** a SME's expertise so everyone can access it from everywhere, simultaneously.

- Elaborate methodologies for eliciting expertise from SMEs are also **expensive** and **slow**, but eventually provide access to a **cheaper** and **faster** "snapshot" of the SME's expertise.

MARSHALL

Dynamic Expertise

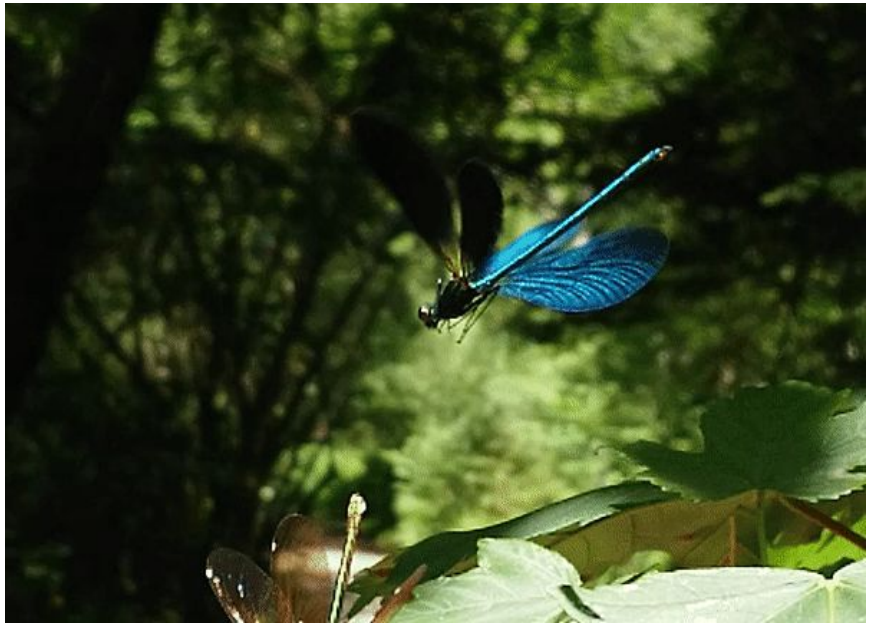

Static Content

Access to Expertise

- Access to static content is far better than no access to expertise at all.
- But when you're doing your homework at 12:30am and get stuck on a problem, you need access to dynamic expertise.
- You need to be able to ask your specific questions and get helpful answers.
- Generative AI provides access to this level of dynamic expertise - and a lot more.

Access to Expertise

- Publishers' first impulse has been to **bolt on** generative AI tutors to their existing static content.

- That's because they're don't realize **their fundamental assumptions are wrong now**.

- With generative AI:

    - Expertise is **abundant**
    - Expertise is **cheap**
    - Expertise is **fast**

Future Course Materials

Structure of Course Materials

- **Now**: Textbooks with chapters, sections, etc.
- **Soon**: Conversations with topics
- Both are structured around a standard scope and sequence and a detailed set of learning objectives that guarantee coverage
- Generative AI allows for meaningful personalization, questions of curiosity, additional explanations and examples, additional practice with feedback, etc.

Authoring of Course Materials

- **Now**: Authors create:
  - scope and sequence
  - detailed learning objectives
  - content
  - supplemental materials

- Author name and institution as primary signals of credibility and authority

Authoring of Course Materials

- **Soon**: Authors will create:
  - scope and sequence
  - detailed learning objectives
  - context and prompts that help models generate conversational content
  - context and prompts that help models generate supplemental materials

- AI model / architecture / provider as primary signal of credibility and authority

Prototype

- The learner chooses a book, a chapter, and an activity to study
- Book and chapter = **context**
- Activity = **prompt**

- Context and prompt are passed to an LLM, which initiates a conversation

Assessments with Course Materials

- **Now**:
  - Quiz banks
  - Assignments with rubrics

- **Soon**:
  - “Stealth assessment” within conversations
  - Automated, rubric-based grading of multimodal assignments

Stealth Assessment

- “Stealth assessment refers to ECD-based assessments that are woven **directly** and **invisibly** into the fabric of the gaming environment. During gameplay, students naturally produce rich sequences of actions while performing complex tasks, drawing on the very skills or competencies that we want to assess… Evidence needed to assess the skills is thus provided by the players’ interactions with the game itself (i.e., the processes of play).”

- Great work underway at ETS on “Designing and Evaluating Evidence-Centered Design based Conversations for Assessment with LLMs”

MARSHALL

Unit Economics of Digital Course Materials

- **Now**:
  - Author royalties

- **Soon**:
  - Author royalties
  - Tokens*

Unit Economics of Open Educational Resources

- **Now**:
  - No author royalties

- **Soon**:
  - No author royalties
  - Tokens*

Legal Considerations

- **Now**:
  - Still navigating how publishers can or can't integrate TCM into generative AI tools and how royalties would work

- **Soon**:
  - OER's 5R permissions mean OER can be integrated into generative AI applications

- For both instructors, students, and publishers, there's never been a better time to choose to use OER over TCM

Now: How Gen AI is Impacting OER Course Materials

Traditional OER

- “Teaching, learning, and research materials that reside in the public domain or have been released under an open license that permits their free use and re-purposing by others.”

- Articles, chapters, essays, textbooks, images, videos, audio, etc.

- Generative AI can profound affect the way traditional OER are **authored** and the ways they are **revised** and **remixed**.

Authoring Traditional OER

- Before generative AI, traditional OER were **hand-crafted**.
- With generative AI, traditional OER can now be **AI-drafted**.
- Using generative AI can reduce the amount of time and resources necessary to reach a first draft by an **order of magnitude** (literally divide by 10) or more.
- Philanthropists, and others who want to be **effective stewards** of funding, will likely begin requiring new traditional OER to be AI-drafted in future grants.

MARSHALL

Revising and Remixing Traditional OER

- Open licenses make it **legal** to revise and remix resources.

- But open licenses don't grant users the **resources** and **expertise** they need to engage in high-demand revise and remix activities:

    - Translating a resource into another language
    - Creating interactive resources based on static resources
    - Adjusting the reading level of a text

- **Generative AI can perform tasks like these** for users that don't have the time, resources, or expertise themselves.

Time and Effort as Obstacles

- Research has shown that **the more time and effort are required** to engage in a specific revise or remix activity, **the less often users will engage** in that activity.

- For example, users are far more likely to delete a chapter from an open textbook than make extensive revisions to an existing chapter.

- Inasmuch as generative AI makes many tasks faster and easier, we should expect to see an **increase in revising and remixing in the future**.

Quality of Revise and Remix

- Research has also shown that the **productivity gains** associated with using generative AI are highest among **lower-skilled** workers.

- For example, on translation tasks, GenAI will be far more helpful to a person who doesn't speak a second language than someone who does.

- Inasmuch as generative AI makes many tasks possible that were previously impossible or impractical, we should expect to see an **increase in the variety and quality of revising and remixing** in the future.

MARSHALL

Soon: Future OER Course Materials

Generative AI + OER = “Generative OER”

- Generative OER are not designed to be used directly for teaching, learning, or research. They are designed to be used as input to a generative AI system and include:

- Openly licensed **context**

- Openly licensed **prompts**

- Openly licensed **model weights**

Open Context

- Looks a lot like **traditional OER**
- **NOT** designed to help students learn
- Designed to **help LLMs give accurate responses**

Open Prompts

- Many of the prompts written by novice GenAI users are **relatively simple** (e.g., “write an essay about the causes World War II”).

- Prompts eliciting more complex behavior, like an extended tutoring session, can be **hundreds or thousands of words long**.

- Consequently, these more useful prompts are **automatically copyrighted** to the full extent of the law (thanks, Berne!), meaning that sharing them for legal revising and remixing will **require open licensing**.

Open Prompts

“Greet me enthusiastically and ask my name. Wait for my answer.

Show me the topics in the outcomes tag as a numbered list. Ask me which topic I would like to learn more about. Wait for my answer.

Once I select a topic, ask me to share several interests or hobbies you can use to explain the topic in more detail as we talk about it. Wait for my answer.

Once I provide a list of interests, explain the topic to me using one of my interests. Ground your examples in the information in the content tag so that your responses will be accurate…”

Why Open Prompts Matter

- The open education community understand the need to localize traditional OER for **linguistic, cultural, and other reasons**. Open prompts need to be revisable and remixable for these same reasons.

- Users will choose to use different GenAI models for a range of reasons. Because different models respond in subtly different ways to the same prompt, another critically important aspect of “localizing” prompts is **refining them to work effectively with different models**.

Open Weights

- (I'm purposefully skipping over the debate about "**open models**.")
- GenAI model weights are conceptually similar to the beta weights calculated in a linear regression. **They're numbers**, like 0.0078183742.
- When we say "model weights" we mean the thousands of matrices containing the **billions of individual weights** that comprise a model.
- "Open weights" are generative AI model weights that are openly licensed so that users can **retain, revise, remix, reuse, and redistribute** them. What does that look like?

MARSHALL

Open Weights

- oss-gpt (OpenAI)
- Gemma (Google)
- Llama (Meta)
- Mistral (Mistral)
- Qwen (Alibaba)

Revising and Remixing Open Weights

- **Revising** is about editing or adapting a resource.
  - **Quantization** is conceptually like rounding (0.42563 to 0.43). Reduces precision but makes model's memory footprint smaller.

- **Remixing**
  - **Fine-tuning** updates model weights through additional training
  - **Distillation** fine-tunes a smaller model on data generated by a larger model

Revising: Quantization

- **Revising** is about editing or adapting a resource.

- **Quantization is conceptually similar to rounding** (0.42563 to 0.43). It reduces the precision of a model's weights, making the model's memory footprint smaller, but also making the model "dumber."

Remixing: Fine-tuning

- **Remixing** is about combining multiple resources into a new one.

- **Fine-tuning**: a process by which model weights are updated through additional training on **new data curated specifically to change model behavior in a specific way**.

- For example, if the fine-tuning data include 10,000 examples of interactions between students and expert tutors, after fine-tuning the model will behave more like an expert tutor.

Remixing: Distillation

- **Remixing** is about combining multiple resources into a new one.

- The process whereby a **smaller model** is fine-tuned on data generated by a **larger model**, in order to **transfer knowledge and behavior from the larger model to the smaller model**.

Why Open Weights Matter

- Foundation models are not designed to **behave pedagogically**.
- Foundation models can **lack** the disciplinary knowledge, cultural knowledge, and other information necessary for specific teaching and learning situations.
- While prompting and context engineering can temporarily improve a model's knowledge or behavior, **fine-tuning** is required to change them permanently.
- Quantization, pruning, and other changes are often necessary to **run models on consumer hardware**.

MARSHALL

Why Open Weights Matter

- Customizing and running these models locally **improves privacy**, **decreases energy consumption**, and addresses other concerns.

- The ability to run models locally is key **improving equity and access** (c.f. MIT OCW's Mirror Site program).

Future Course Materials

- Openly licensed **context** and **prompts**
- Openly licensed, **pedagogically aligned open weights**
- Open source **UI layer** and **backend orchestration system**
- **Running locally** on users' devices
- Syncing performance and other analytics data **if desired**

MARSHALL

Other Implications

GenAI-enabled Pedagogy

- **GenAI-enabled pedagogy** is the set of teaching and learning practices that are only possible or practical when teachers and students are able to use generative AI.

- **‘The things you can do with GenAI that you can’t do otherwise.’**

- For example, we typically don’t assign synchronous collaborative activities to students in asynchronous online courses. But when generative AI can play the partner, **activities that require a collaborator can be done at any place and time**.

Reifying Our Pedagogy in Course Materials

- Generative OER **reify our pedagogies**, enabling us to **share, scrutinize, replicate, and evaluate** our teaching as never before.

- Course materials efficacy research is about to **get real.**

- (**I am not claiming** that all pedagogies can be implemented through generative OER. But many can be.)

“Productizing” EBPs in Course Materials

- **Most people don’t understand** the technical differences between the 4G and 5G wireless standards. You don’t have to - 5G is built into your phone so that you can benefit from higher speeds without needing a masters degree in information theory.

- **Most instructors and students are unaware** of the large body of research into evidence-based teaching and learning practices. But they don’t have to be - if these practices are built into generative OER, everyone can benefit from **doing more effective things** without a masters in teaching and learning.

Conclusion

Conclusion

- The **internet** dramatically changed how we deliver course materials, but resulted in only incremental changes to their design.

- **Generative AI** could dramatically change how we design course materials because their **underlying assumptions are no longer true** - expertise is now abundant, cheap, and fast.

- The market can / will produce these new materials **only if faculty will adopt them**.

MARSHALL

Thank You!

david.wiley@marshall.edu

bit.ly/2025nacs