

Important Considerations When Using AI at UNM

Data Privacy & Security:

Anything you share with a publicly available AI tool could be stored, reused, and incorporated into training of the model. This is especially risky with sensitive data like student records, health information, and proprietary research. Always use university-approved tools over public platforms, verify data policies before sharing anything, and when in doubt, leave it out.

Understanding how AI handles your data is critical. Individual subscriptions to public platforms like ChatGPT, Google Gemini, and other platforms that have not completed an enterprise agreement with UNM, often store your prompts and the data you provide and use them to train future models. They may also share your data with third parties. While this helps improve the vendor's AI technology, it means that anything you type into these tools could be retained, analyzed, and potentially surfaced in ways you didn't intend. This applies to text, uploaded documents, images, and any other content you share with the model.

The risks are especially significant in a university setting, where we regularly work with sensitive student records, proprietary research, personal health information, and other confidential data. Once that information is entered into a public AI platform, you may have no control over how it's stored, who can access it, or how long it's retained. In some cases, such disclosures may constitute disclosable breaches and require formal incident responses, as well as reports to regulatory bodies.

A few guidelines to protect yourself and others:

1. **Never input sensitive or confidential information into any AI tool** until a thorough UNM risk assessment has been completed, and you have received data steward approval. This includes student information, financial data, medical records, human subject research, other proprietary research, and anything covered by Controlled Unclassified Information (CUI), FERPA, HIPAA, or other regulatory requirements.
2. **Read the fine print.** Data policies vary widely between platforms and can change without notice. What was private yesterday may not be tomorrow. Don't assume your data is safe, always verify. Please also be attentive to updates to applicable terms and conditions.
3. **Understand the difference between public and enterprise tools.** There is an important difference between public AI tools and tools that have been vetted and

approved by the university. A paid subscription does not guarantee your data is protected. University-approved tools have been specifically reviewed to meet our security and privacy standards and applicable regulations for the data that will be used with them. Enterprise tools should be chosen for university tasks whenever possible.

4. **When in doubt, leave it out.** If you're unsure whether something is safe to share with an AI tool, don't. It's always better to err on the side of caution.

Your data and the data of those UNM serves is a responsibility, not just a convenience. Treat it accordingly.

Intellectual Property, Copyright, and Integrity:

Generative AI technologies are based on the development of sophisticated models that are trained using digital content that was originally created by humans. This model training was often completed without the consent or awareness of the original creators. When using generative AI, be aware of issues related to Intellectual Property, Copyright, and Integrity.

Case law related to AI-generated content and the intellectual property rights of the original creators of content that AI models were trained on is not settled. It is always a good idea to pay attention to the Intellectual Property and Copyright protections that may be related to content that you intend to share with an AI model or that the AI model produced for you.

Similarly, disclosure of your use of AI tools increases transparency and supports the academic and intellectual integrity of your work. In some cases, rights to your own intellectual property may be diluted if AI tools have been used to create content.

A few guidelines for using AI with integrity:

1. **Understand requirements and restrictions for AI usage.** Academic Journals, research projects, creative exhibitions, and academic courses may have specific limitations on how/if AI may be used. Check with the journal, funding agency, instructor, supervisor, etc. to understand how AI may be used before you use it. You are personally responsible for ensuring that you are not misusing AI in each context.
2. **Acknowledge and cite your use of AI.** Different situations may have different citation requirements, but it is common to include a citation that describes what AI models or services were used and how they contributed to the final product.
3. **Respect copyright and Intellectual Property.** Consider how your use of AI protects the copyright and intellectual property (including your own) for content you share

with AI models as well as whether generated results may infringe on the rights of creators whose content was used to train models.

4. **Promote trustworthy use of AI.** Whether through biased training data, inaccurate data sources, model error (hallucinations), or intentional prompting, AI can technically produce content that is not factual, applies biased decision-making frameworks, or promotes misinformation. Be sure to verify that AI content is accurate and truthful.
5. **Accept responsibility for the final product.** Whether used to brainstorm, generate a final product, or automate a task, the human directing the AI tool bears responsibility for the actions AI performs on their behalf, and for the accuracy of the content that is produced.

AI and the Environment:

AI can require significant amounts of energy and water to operate. As a university in the desert, we have extra reason to care. A single prompt has a small footprint, but it adds up at scale. Use AI intentionally, choose smaller models when possible, and skip unnecessary prompts.

AI systems require large amounts of energy and water to train and operate, which can carry a real, negative environmental cost and impact. The data centers powering these tools consume significant electricity and require vast amounts of water for cooling, and as AI adoption accelerates globally, that demand is growing rapidly. As a university in the desert, this is a particularly important issue for us to be aware of.

While a single prompt uses relatively little energy and water, the impact adds up across millions of users over time. This isn't unique to AI – many of the technologies we rely on daily consume similar resources. In some cases, AI can actually reduce overall resource use by streamlining workflows that were previously more energy or time intensive.

The goal isn't to avoid AI altogether, but to learn to use it intentionally and responsibly.

A few ways you can help:

1. **Use smaller models when possible.** AI platforms often offer a range of models, from lightweight, efficient options to larger, more powerful ones. For everyday tasks like drafting an email, summarizing a document, or brainstorming ideas, a smaller model will often get the job done just as well while using a fraction of the resources. Save the more powerful models for complex tasks that genuinely require them.
2. **Keep your prompts concise** while still including all the necessary information.

3. **Make minor edits yourself** rather than sending follow-up prompts for small tweaks to a response.
4. **Perform a final review of the work** to help ensure that such work is complete and accurate, especially for work that may have consequences for employees, patients, research subjects, and/ or students. You are responsible for the work that AI does on your behalf.
5. **Be purposeful about when you use AI.** Ask yourself whether AI is the best tool for the task, or whether you're reaching for it out of habit. Asking an AI about the weather or a simple math problem, for instance, uses far more resources than using a tool that was designed for that purpose.

Thoughtful, intentional use of AI allows us to harness its benefits while minimizing its environmental impacts, and other potential harms.

Appropriate Use Cases:

AI should enhance your work, not replace the thinking behind it or shortcut the cognitive processes needed to thoroughly evaluate ideas. The end result should be of the same quality or better than what you'd produce without it. Always verify AI output, be transparent about your use, and make sure the final product reflects your own expertise, intentions, and responsibilities.

AI is a powerful tool, but its value depends on how you use it. Used well, it can help brainstorm ideas, streamline repetitive tasks, and accelerate research. Used poorly, it can lead to inaccurate, generic, or low-quality work, and can create avoidable harms to the communities UNM serves.

The key is treating AI as a collaborator, not a replacement for your own thinking. When used properly, the result should be the same quality or better than what you would have produced without it. If someone can tell you used AI, that's a sign to put more of yourself into the work. A few guidelines to keep in mind:

1. **Use AI to enhance your work, not bypass it.** Let it help you get started, organize your thoughts, or work through a problem, but make sure the final product reflects your own expertise and judgment. Don't forget that sometimes the value of the activity is in the brainstorming and applied thought processes.
2. **Always verify AI-generated content.** AI models can produce information that sounds convincing but is inaccurate or fabricated. Never take outputs at face value and always do your own research to validate.

3. **Be transparent about your use of AI.** If AI played a meaningful role in your work, disclose it. For students, always follow your instructor's specific policies on AI use. If you have created a quality product, you shouldn't be embarrassed to say AI assisted you.
4. **Ask yourself: is AI the right tool for this task?** Sometimes the process of doing the work yourself by thinking critically, writing from scratch, and problem-solving, is more valuable than the end result.
5. **Follow university policies.** AI use should align with existing university policies and guidelines. For students, this includes the UNM Student Code of Conduct (Pathfinder), the Faculty Handbook policy on Dishonesty in Academic Matters, and any guidance provided by your instructor in the course syllabus. For faculty and staff, this includes applicable data governance, procurement, and information security policies.

AI works best when it amplifies human effort and promotes human-centered experiences rather than replacing it. By using it thoughtfully, we maintain the quality and integrity of our work while still benefiting from what this technology has to offer.

Loss of Critical Thinking and Creativity:

AI can erode critical thinking and creativity or sharpen them. The difference is in the approach. Use it to challenge your ideas, break through blocks, and build new skills, not to skip the thinking entirely. AI should make you better at what you do, not do it for you.

One of the most common concerns about AI is that it will erode critical thinking, stifle creativity, or create dependency. These are valid concerns and they can absolutely become reality when AI is used as a shortcut rather than a tool. If you're asking for AI to do your thinking for you, you will lose those skills over time.

When used with intention, however, AI can actually sharpen your thinking, assist your learning, expand your creativity, and open doors that previously felt out of reach. This doesn't diminish anyone's expertise; it gives more people the ability to participate, contribute, and grow in areas that might have felt inaccessible before. The difference isn't the technology, it's the approach.

Consider using AI to:

1. **Stress-test your ideas.** Ask it to poke holes in your argument or present counterpoints, then strengthen your work based on what you learn. AI tools often

tend to be overly agreeable, so you may need to clearly tell it that you want honest pushback and scrutiny of your assumptions.

2. **Break through creative blocks.** Use it as a brainstorming partner to generate possibilities you might not have considered, then build on them with your own voice and vision.
3. **Deepen your understanding.** Ask it to explain a concept in a different way, quiz you on what you've learned, or walk you through a topic step by step. Then engage your critical thinking skills by challenging yourself to evaluate whether its output is accurate and well-reasoned.
4. **Bridge skill gaps.** Use it to learn skills that might have previously felt out of reach, not to skip the learning entirely. The goal is to come out the other side more capable, not more dependent.

The people who will thrive alongside AI aren't those who hand over the wheel, they're the ones who use it to see further, think deeper, and create with more confidence. AI should make you better at what you do, not do it for you.

Hallucinations/Inaccuracies:

Even the most advanced AI models hallucinate. They can produce convincing but entirely wrong information, and they have no way of knowing the difference. Your job is to verify and treat every output as a draft, not a fact.

AI models don't "know" things the way people do. They generate responses by predicting what comes next based on patterns in their training data. This means that AI models can produce information that sounds confident and convincing but is completely wrong. These are often called "hallucinations," and they're one of the most important risks to understand when working with AI.

Inaccuracies can also stem from the data the models were trained on. AI models learn from massive amounts of existing content, and not all of that content is accurate. Misinformation, outdated articles, satire, and even April Fools' jokes can all end up in the training data. Even when the model isn't "making things up," it may be confidently repeating something it learned that is false.

While this improves over time, this isn't a flaw that will simply be "fixed in the next update." Even the most advanced models are prone to hallucinations. It is a fundamental part of how these models work. In fact, large language models do not "know" facts and knowledge

in the same way that humans or even a database would “know” it. They are statistical engines that are designed to predict the most likely response to a question. The more technical or specialized the topic, the harder these errors can be to spot.

The good news is that with the right habits, you can significantly reduce the risk:

1. **Treat AI output as a starting point, not a final answer.** Always verify facts, data, and sources through your own research before relying on them, especially those that are novel or surprising.
2. **Provide it with background information.** When AI outputs are reliant on domain expertise or specialized knowledge, your results will be better if you can tell the AI where to look for that kind of information. Also known as Retrieval Augmented Generation (RAG), telling the AI what documents and websites to review in generating the answer will ensure that the answer is more closely matched to your needs and expectations.
3. **Be especially cautious with citations and references.** AI is particularly prone to fabricating sources that don't exist or applying a source as a piece of evidence for a claim that the original article did not make. If it cites something, look it up.
4. **Ask it to explain its reasoning.** If something doesn't feel right, push back. Ask the model how it arrived at its answer. This can surface weak logic or unsupported claims.
5. **Know the limits of your model.** Most AI models have a knowledge cutoff date and don't have access to real-time information unless specifically equipped with search tools. Keep this in mind when asking about recent events or current data.

AI can be an incredibly useful research and productivity tool but only when you bring your own judgment to the table.

Bias/Lack of Diversity:

AI can reinforce the very biases we're working institutionally to address. Its training data is shaped by the world's existing gaps and imbalances, and by those of the model's authors and its outputs reflect that. Don't mistake confident delivery for objectivity and always evaluate AI output with a critical lens.

AI models are shaped by the data they're trained on and that data reflects the biases, gaps, and imbalances of the world it came from – including the biases of those who created the models or, selected the data on which the models are trained. If certain perspectives, cultures, or communities are underrepresented in the training data, the model's outputs

will often reflect that. This can show up in subtle ways: defaulting to certain assumptions, reinforcing stereotypes, or consistently overlooking viewpoints that weren't well-represented in the data.

Left unchecked, biased outputs can lead to real harm. Across industries, unchecked AI bias has led to real consequences, from reinforcing stereotypes in content to disadvantaging certain groups in hiring and policy decisions. Viewing AI outputs with a critical lens isn't optional, **it's essential**.

Since AI presents its responses with confidence and authority, these biases can easily go unnoticed and become baked into our work – especially if we share in some of those biases. If we accept AI outputs without questioning whose perspective is present, and whose is missing, we risk not just reflecting existing biases but actively amplifying them.

Some practical steps to help spot and address bias:

1. **Question the defaults.** If an AI response feels one-dimensional or makes assumptions, ask yourself which perspectives are present and which might be missing.
2. **Be intentional with your prompts.** The way you frame a question influences the answer. If you want diverse or inclusive perspectives, ask for them explicitly. Unless you specifically prompt it to do otherwise, most models will look for evidence to support your position rather than objectively evaluating all evidence that is out there.
3. **Don't treat AI as a neutral source.** It can feel objective because it sounds confident and measured, but its output reflects its training data, not an unbiased truth. Some vendors intentionally build a certain level of agreeableness into their models because it helps keep users engaged with their platform.
4. **Beware of Sycophancy.** Many commercial models were intentionally designed to agree with the perspectives of the person prompting them.
5. **Bring your own perspective.** You know your audience, your community, and your context better than any model does. Use that knowledge to evaluate and shape what AI gives you.

AI can be a helpful tool for reaching broader audiences and exploring new ideas, but it requires a critical eye to make sure the output truly represents the diversity of thought and experience we value.

Emotional/Psychological Risks:

Treating AI like a companion or emotional support system can lead to real psychological harm, including unhealthy attachment and a distorted sense of reality. Never use AI as a substitute for mental health support or medical advice and be mindful of how your interactions with it may influence others.

AI is designed to generate responses that are helpful, engaging, and can feel personal, but it is important to remember that AI is not a person and in the current state of these technologies, is not intelligent. When we begin treating AI as a confidant, companion, or emotional support system, it can lead to unhealthy attachment, social isolation, and a distorted sense of reality. In extreme cases, prolonged and emotionally dependent interactions with AI have led to AI-associated psychosis, which involves delusional beliefs rooted in things that the model has said or validated.

AI also poses serious risks when it intersects with mental health. When someone is in crisis, they need genuine human understanding and resources. AI models have no ability to read between the lines, assess real danger, or provide the careful, trained support these moments require. Research has consistently shown that irresponsible representations of suicide, including detailed descriptions or methods, can increase risk for vulnerable individuals. AI models can and do produce this type of content when prompted, which is why appropriate guardrails and ongoing testing of model responses to sensitive topics are critical.

A few things everyone should keep in mind:

1. **Never use AI as a substitute for mental health support.** If you or someone you know is struggling, reach out to a qualified professional or contact the 988 Suicide & Crisis Lifeline.
2. **Don't personify AI models.** Even if it feels harmless or lighthearted to you, treating AI like a friend or companion normalizes a dynamic that can lead to real harm for others who may be more vulnerable.
3. **Be aware of how AI interactions affect you.** If you notice yourself turning to AI for emotional comfort, validation, or companionship, take a step back and reconnect with real people in your life.
4. **Speak up if something feels off.** If you encounter AI-generated content that feels harmful or inappropriate, report it. If you notice someone in your life developing an unhealthy relationship with AI, say something.
5. **If you need help, please access UNM's mental health resources for support:**

- [UNM Mental Health Resources](#)
- [UNM Agora Crisis Center](#)
- [UNM Student Health and Counseling](#)

We share these risks not to cause fear, but because understanding them is the first step in preventing them. How we interact with this technology, and how we model that behavior for others, matters.

Communicate a Clear Plan:

People are looking for clear guidance on AI. Set specific expectations, explain your reasoning, create space for questions, and don't force adoption on those who aren't ready. Clarity is the foundation for responsible AI use.

One of the most important things you can do as a leader, instructor, or manager is to set clear expectations around AI use. Whether it's in the classroom, the office, or a research lab, the people you work with are looking for guidance; not just permission to use the technology, but a clear understanding of when and how it's appropriate.

Without that clarity, people are left guessing. Some will avoid AI entirely out of fear of doing something wrong. Others will use it freely without considering the implications. Neither outcome is ideal.

A few steps to get started:

1. **Define what's acceptable.** Be specific about when AI use is required or encouraged, when it's acceptable with limitations, and when it's off the table entirely. Vague policies lead to vague compliance. It is critical that our employees and students clearly understand the boundaries of AI use as we are establishing the norms around AI use.
2. **Explain why.** People are more likely to follow guidelines when they understand the reasoning behind them. If you restrict AI use for certain assignments or tasks, share what you want them to gain from doing the work themselves – or the risks you see in the use of AI for certain tasks.
3. **Create space for questions.** AI is new territory for everyone. Foster an environment where people feel comfortable asking questions, raising concerns, or admitting they're unsure about the boundaries.

4. **Revisit and adapt.** This technology is evolving quickly, and your guidelines should too. Check in regularly with your team or students to see what's working and what needs to change.
5. **Don't force it.** Many people have real and valid concerns about AI, from privacy to ethics to job security. Requiring AI use when it isn't necessary dismisses those concerns and breeds resentment. Instead, meet people where they are. Address their questions honestly, give them time to get comfortable, and let them see the value for themselves.

Clear communication and trustworthy relationships are not just helpful; they are the foundation for responsible AI use in any setting. The goal is to create an environment where people feel informed, supported, and empowered to engage with AI on their own terms.

Listen To and Address Concerns:

Worries about AI aren't obstacles to progress, they're essential to getting it right. Whether people are concerned about job security, overwhelmed by the pace of change, or skeptical of the hype, their perspective matters. What matters is that we keep the conversation going, take concern seriously, and make decisions about AI together.

As AI continues to evolve and become more embedded in how we work and learn, concerns will inevitably arise, and that's a good thing. The questions people ask and the hesitations they express aren't obstacles to progress; they're essential to getting this right.

Some people worry about being replaced. Others feel frustrated by the pressure to learn yet another new technology when they're already stretched thin. Some are skeptical that AI will live up to the hype and question whether it's worth the investment of time and energy. Others have real ethical concerns about the environmental or social impact. These are all valid concerns, and they deserve honest, thoughtful responses, not dismissal.

No single policy or set of guidelines will have all the answers. This technology is changing rapidly, and our understanding of its impact is changing with it. What matters most is that we approach this as an ongoing conversation where every voice is heard, every concern is taken seriously, and no one feels left behind or pressured into something they don't understand or trust yet.

If something doesn't feel right, say so. If you have questions, ask them. And if you're in a position of leadership, make it clear that this kind of feedback isn't just welcome, it's needed. The decisions we make about AI today will shape how this technology impacts our

community for years to come. We owe it to ourselves and each other to make those decisions thoughtfully, together.

Shadow AI

Using unapproved AI tools, even with good intentions, can create compliance risks, data privacy gaps, and redundant workflows. Not all AI tools are created equal, and a lack of a ban doesn't equal a green light. University-approved tools are vetted to protect you and the university.

Shadow AI, the use of unapproved or unvetted AI tools within an organization, is one of the fastest-growing risks in higher education and the workplace. It often happens with good intentions: someone finds a tool that helps them work faster, shares it with their team, and before long it's embedded in a workflow without ever being reviewed for risk, security, compliance, or data privacy.

While UNM does not currently ban any specific AI tools across the board, that doesn't mean every tool is safe to use for every purpose. The tools and platforms we approve go through a deliberate vetting process designed to identify and address the very risks outlined in this guide, from data privacy and regulatory compliance to bias mitigation, content safety, and appropriate guardrails. When AI tools are adopted outside of this process, those protections disappear.

Use of commercial AI products that have not been evaluated by UNM to ensure that they protect information security and privacy is strongly discouraged. **Individuals should exercise caution and refrain from inputting any data that is protected under CUI, FERPA, GLBA, IRB, PII, or other secured data categories.** Examples include but are not limited to:

- Student records,
- Personally Identifiable Information (PII) such as
 - o Social Security numbers,
 - o addresses,
 - o Banner IDs,
 - o phone numbers, and
 - o email addresses

Secured data categories also encompass information protected by the Health Insurance Portability and Accountability Act (HIPAA), Payment Card Industry Data Security Standard (PCI DSS), Federal Information Security Management Act (FISMA), and any other data that is considered sensitive or confidential.

Before using UNM Data, it is always advisable to consult with the appropriate **UNM Data Steward** for guidance on tools where third parties could gain access to that data.

Here's why this matters:

1. **Compliance and regulatory risk.** Unapproved tools may not meet FERPA, HIPAA, or other regulatory requirements the university is obligated to follow. Use of unapproved tools with this data may lead to data breaches and required regulatory reporting. In some cases, costly, substantial incident and breach response activities may be required.
2. **Inconsistent standards and oversight.** University-approved tools come with documented considerations, known limitations, and guidance for users on potential risks. Shadow AI tools are unlikely to offer the same level of transparency, and when teams are using different tools with different capabilities, it becomes harder to maintain quality and consistency across the university.
3. **Data stewardship gaps.** The university has established data stewardship processes to ensure that data is handled, stored, and accessed responsibly. Shadow AI often bypasses these critical steps entirely, meaning there's no way to track what data is being shared, where it's stored, or who has access to it.
4. **Redundancy and waste.** Shadow AI often leads to multiple teams independently adopting different tools for the same purpose, creating duplicated costs and fragmented workflows.

We encourage exploration, but within the process, not around it. If you've found a tool you think could benefit your work, bring it forward by emailing us at loboai@unm.edu and let us evaluate it properly. The goal isn't to slow innovation down; it's to make sure we're innovating responsibly.