

Considering the Ethics of Large Machine Learning Models in the Chemical Sciences

Evan Walter Clark Spotte-Smith^{1,2*} 

¹ Department of Chemical Engineering, Carnegie Mellon University

² Department of Materials Science and Engineering, Carnegie Mellon University

Foundation models, including large language models, vision-language models, and similar large-scale machine learning tools, are quickly becoming ubiquitous in society and in the professional world. Chemical practitioners are not immune to the appeal of foundation models, nor are they immune to the many risks and harms that these models introduce. In this work, I present the first analysis of foundation models using the lens of scientific ethics and chemical professional ethics. I find that common general-purpose foundation models are essentially incompatible with the moral practice of chemistry, though there are fewer ethical problems with chemistry-specific foundation models. My discussion, which includes environmental harm, epistemological risk, labor ethics, and more, concludes with an examination of how the harm associated with foundation models can be minimized and further poses a set of serious lingering questions for chemical practitioners and scientific ethicists.

1 Introduction

In recent years, machine learning (ML) models have grown in complexity and scale (in terms of e.g., number of learnable parameters and training data size) at a rapid pace, culminating in the development of so-called “foundation models” or “large models”. These models, which are intended to either be directly applicable to problems in a wide domain space or to be easily generalizable through fine-tuning and transfer learning, include generative models like large language models (LLMs, e.g., GPT-3,¹ DeepSeek-r1),² image and video diffusion models (e.g., DALL-E,³ Stable Diffusion),⁴ and vision-language models (VLMs, e.g., LLaVA),⁵ among other, less common examples. Foundation models have infiltrated nearly every aspect of the digital world, from standalone applications like the popular chat interface ChatGPT to digital search, code generation, automated writing, and much more.

The chemical sciences have not been immune to the proliferation of foundation models. There have been a number of recent efforts seeking to evaluate the capabilities of generative models like LLMs for chemical tasks, from answering simple questions⁶ and mining text⁷ to performing complex experiments by interfacing with laboratory hardware.^{8,9} Efforts to create chemistry-specific foundation models have also appeared, including protein-folding models^{10,11} and “universal” machine-learned interatomic potentials.^{12,13}

In the recent rush towards ever larger and more general ML models, comparatively little focus has been placed on the ethics of these models, though a growing field of “artificial intelligence” or “AI” ethics has emerged.^{14–16} Even less attention has been paid to the ethics of foundation models when applied in scientific domains, including chemistry. This work aims to fill that gap and initiate a

conversation about responsible development in applied data science and scientific ML.

I begin (Section 2) by providing a theoretical and practical basis for ethical discussion. Drawing from the field of science ethics as well as from several prominent ethical codes for chemists and chemical engineers, I distill some basic principles that (should) guide the practice of chemical science. While I will discuss alternative perspectives, my discussion favors a deontological ethic. I discuss the general risks and harms associated with foundation models (Section 3), focusing on generative models, and then (Section 4) examine how these risks are largely unacceptable within established scientific and chemical ethics. In Section 5, with these more general considerations at hand, I shift focus to examine ethical concerns that are unique to applications of foundation models in the chemical sciences. Finally (Section 6), I provide a brief discussion and conclusion. I suggest how some of the risks and harms mentioned might be mitigated, using analogies from efforts in other areas of chemistry, and I emphasize the need for further research into chemical “AI” ethics.

2 A Basis for Chemical Ethics

2.1 The Ideals of the (Chemical) Scientist

There have been many attempts to distill the moral ideals of science and scientists. For instance, Robert Merton famously laid out four principles now known as the “Mertonian Norms”:¹⁷ communism (sometimes instead called “communalism”), universalism, disinterestedness, and organized skepticism. Mostly, the different sets of ideals overlap considerably. In the interest of simplicity, I will focus on the six ideals described by the National Academies of Science, Engineering, and Medicine in their report *Fostering Integrity In Research*.¹⁸

*Corresponding author: ewcspottesmith@cmu.edu

1. *Objectivity*: the responsibility of a scientist to avoid letting their personal feelings or beliefs introduce bias into their findings
2. *Honesty*: a core responsibility which includes both truth-telling and transparency
3. *Openness*: a scientist's responsibility to present all relevant findings without obfuscation
4. *Accountability*: a scientist's obligation to be able to explain and justify their work, and to stand behind what they've done
5. *Fairness*: consistently making judgments based on clear, equitable criteria
6. *Stewardship*: an overarching responsibility to maintain relationships within scientific organizations and between scientists and broader communities

These principles are far from universally accepted, and there is certainly room for argument and further development. For instance, the debate around whether or not human beings can ever be objective or act objectively (and under what conditions) continues.^{19–21} If objectivity is forever out of reach, it may not make sense as a guiding principle for human behavior. Other principles are more certainly attainable but are nonetheless rarely realized in practice. What scientist presents *all* of their findings without omission? How often can a researcher explain and justify *every* element of *every* article and project that they are a part of, particularly in an age when collaborations are becoming larger^{22,23} and researchers more specialized?²⁴ While these are questions worth asking, and while the principles enumerated above are perhaps flawed, for the purposes of this discussion I will act on the assumption that they are appropriate ideals for (chemical) scientists and that they are generally worth aspiring to.

It is worth pointing out that the principles of the National Academies speak to the social element of science. While science is an epistemology that has been called a pursuit of “reliable knowledge”,²⁵ as a profession and practice, it is also a social activity. Essentially all practicing scientists “know” or “believe” ideas about the natural world on the basis of the theories, simulations, and experiments of other scientists.²⁶ In this way, scientific epistemology and scientific ethics cannot be disentangled; our knowledge is only as good as the scientists who we are reading, listening to, and interacting with. By upholding honesty, we seek to improve the reliability of our knowledge and that of those around us, while stewardship more directly preserves the social connections which enables science's decentralized knowledge generation.

In this discussion, I am assuming that, once a set of principles are agreed upon, all chemical scientists will and should try to follow them. It is reasonable to challenge this assumption and ask “why should they”? Clearly, it would be insufficient to draw on legal or professional punishment. Laws are the morals of the State, not of people, and fear of retribution, however valid a motivation, is not a moral one. While the motivations for ethical behavior are outside of the scope of the present work, I should point out the work of Kovac,^{26,27} who has invoked virtue ethics and specifically the virtue of *reverence* as a possible answer.

2.2 Chemical Ethics in Practice

One would hope that most practicing scientists agree with the National Academies' principles, the Mertonian Norms, and other ethical ideals of science, but most scientists have not explicitly agreed to live and work according to these ideals. On the other hand, scientists, including chemists and chemical engineers, regularly agree to follow professional codes of ethics. It is thus perhaps more appropriate to ground ethical discussions in professional codes, rather than ideals.

Moreover, professional codes of ethics are typically somewhat more specific in what they require of the members of the profession. Here, I will discuss practical professional ethics in the form of the American Chemical Society's (ACS's) *The Chemical Professional's Code of Conduct*,²⁸ the *Global Chemists' Code of Ethics* (also developed through the ACS),²⁹ the American Institute of Chemical Engineers (AIChE) *Code of Ethics*,³⁰ and *The Hague Ethical Guidelines*.³¹ I discuss each source, focusing on specific text that is most relevant to the present discussion. After discussing each individual code, I consider where the codes share common ground and what might be the foundation of a general professional ethic for chemistry. While the four sources that I have selected are somewhat diverse in the types of professionals considered in their rules, they are mostly based on and written from an American context, which could potentially bias this and the following discussion.

I note that most chemical codes of ethics, including those discussed here, do not address what happens when a chemical scientist violates their agreed-upon norms. In the present work, I am mostly unconcerned with the repercussions of unethical behavior and wrongdoing. Taking a deontological perspective, I aim to analyze which behaviors are ethical or unethical. Rather than motivating ethical behavior (Section 2.1), I assume that chemical practitioners should (and will aim to) do what is ethical because it is good, and should avoid what is unethical because it is not good. I briefly discuss repercussions and regulations later in this work, e.g., in Section 6.

2.2.1 The Chemical Professional's Code of Conduct

The Chemical Professional's Code of Conduct, which was approved in its current form in late 2019, is motivated by a desire to “advance chemical science while striving for the highest standards of scientific integrity”. While much of the *Code* is quite general, some specific standards are listed, including sharing ideas, maintaining accurate records, and properly distributing credit. Similarly, some specific unethical behaviors are included, e.g., undisclosed conflicts of interest and plagiarism.

Chemical scientists under the *Code* are charged with maintaining public health and safety and “[serving] the public interest”. While non-human health and safety is not directly listed, there is a separate clause imploring chemists to strive for environmental sustainability which “includes continuing to work to develop sustainable products and processes that protect the health, safety, and prosperity of future generations”. This means that chemical scientists operating under ACS' *Code* must make ethical decisions considering individuals in the future under considerable uncertainty.

Aside from a more general clause about respect, which includes avoiding behaviors like bullying, harassment, and coercion, ACS specifically instructs its members to “avoid all bias based on race, gender, age, religion, ethnicity, nationality, sexual orientation, gender expression, gender identity, presence of disabilities, educational background, or other personal attributes”. Finally, *The Chemical Professional's Code of Conduct* makes it clear that it is not enough for an individual chemical professional to behave ethically; rather, one has a responsibility “to act or intercede where possible to prevent misconduct”.

2.2.2 Global Chemists' Code of Ethics

The Global Chemists' Code of Ethics begins by declaring that chemical scientists should “promote a positive perception and public

understanding and appreciation of chemistry”. This is not a negative or harmful idea, but it is curious in that this introductory motivation has to do with *perception*, rather than reality. In the framing of the *Global Chemists’ Code of Ethics*, it is important that the public perceive chemistry as being positive but not as important that chemistry actually is positive and beneficial for that same public.

Environmental concerns are central to the *Global Chemists’ Code of Ethics*, with a full section titled “Environment” immediately following the introduction. Under this code, chemical practitioners should work to make sustainability central in research and in education and should protect the environment for the wellbeing of those who come in the future. This call for environmentalism is not only individual; global chemists are called upon to make their organizations sustainable and promote public understanding of the environment and sustainability.

The *Global Chemists’ Code of Ethics* section on “Research” emphasizes that chemistry is and should be a public service, that “Research in the chemical sciences should benefit humankind and improve quality of life”. This central drive for public benefit calls back to the earlier environmental statements and includes future generations in the definition of “humankind”. Transparency and avoiding conflicts of interest are highlighted, as is the somewhat vague order to “practice collegiality in the best way”.

In “Scientific Writing and Publishing”, openness, honesty, integrity, reproducibility, and correctness are all mentioned as centering principles. Chemists have a personal responsibility and a responsibility for those they supervise to ensure that there are no errors in published work, to avoid plagiarism, and to “promote peaceful, beneficial applications and uses of science and technology through a variety of media”.

Finally, chemical scientists have obligations to ensure safety and security, with the latter including minimizing possible dual use risk.

2.2.3 AIChE Code of Ethics

The core requirement for chemical engineers, according to AIChE, is to “uphold and advance the integrity, honor and dignity of the engineering profession”. Like the *Global Chemists’ Code of Ethics*, this seems to emphasize public perception and social acceptability as opposed to making positive change in the world. Nonetheless, within this call for integrity and honor are some additional charges, specifically “being honest and impartial and serving with fidelity their employers, their clients, and the public; striving to increase the competence and prestige of the engineering profession; and using their knowledge and skill for the enhancement of human welfare”. Notably, employers are listed before clients and the public, perhaps implying that, when in conflict, one’s responsibility to an employer comes first. This is deeply problematic given the history of intentional and unintentional chemical disasters driven by government and corporate interests,³²⁻³⁶ but I digress.

While the environment is not mentioned in AIChE’s main ethical goals, environmental protection is mentioned along with protecting public welfare in a subsequent clause. Chemical engineers must also accept responsibility and actively seek out critiques of their work, avoid working in areas outside of their competence, and be fair and respectful. In the vein of respect, there are separate and specific clauses in the *AIChE Code of Ethics* stating that chemical engineers should promote diversity, equity, and inclusion and “never tolerate harassment”.

2.2.4 The Hague Ethical Guidelines

The *Hague Ethical Guidelines* were written not for a particular professional body but rather to support the 1993 Chemical Weapons Convention,³⁷ an international agreement binding 193 countries. It is a truly global collection of norms rooted in an understanding of the immense danger and destructive power of chemistry when practiced unethically. Chemical weapons and the deaths that chemistry can bring are central to the *Hague Guidelines*, which reiterate the agreement not to “develop, produce, otherwise acquire, stockpile or retain chemical weapons”, among others. At the same time, the authors argue that, behaving ethically as described by the *Hague Guidelines* will “[ensure] high quality results in science”.

The *Hague Guidelines* inextricably link ethics, public benefit, and environmental protection, stating, “The responsible practice of chemistry improves the quality of life of humankind and the environment”. Environmental protection is further emphasized, and under these guidelines, promoting sustainability is a “special responsibility”, requiring practitioners to “[meet] the needs of the present without compromising the ability of future generations to meet their own needs”.

Chemical practitioners, including chemistry teachers, should promote beneficial (*i.e.*, peaceful) applications of science and technology while preventing the misuse of existing technologies and the research and development of new harmful technologies. Preventing harm includes promoting safety and ensuring security, *e.g.*, preventing theft and harmful or destructive applications of chemical supplies.

2.2.5 Synthesizing Professional Chemical Ethics

While there are some significant and some more subtle differences between the four codes studied, they concur on many of the responsibilities of chemical scientists. With broad agreement, chemists, chemical engineers, and other chemical practitioners should:

1. Work to ensure the wellbeing of current populations, future generations of humans, and the environment;
2. Share ideas openly while attributing ideas appropriately;
3. Avoid a set of intrinsically harmful activities, including plagiarism, harassment, dishonesty, misusing technology, and developing technologies with destructive applications;
4. Increase knowledge, both within the scientific community and in the general public, by speaking truthfully and accurately;
5. Protect health, safety, and security, broadly defined.

Several codes emphasize that chemists should not only meet these obligations as individuals but should drive other individuals and communities to also behave ethically.

The responsibilities listed here are compatible with and extend the general principles outlined in Section 2.1. While there are certainly other requirements for ethical practice of (chemical) science not listed, if one follows these guidelines, one does not risk going against any foundational ideals.

2.3 Considering Chemical Utilitarianism

It is worth taking a brief detour at this point to consider alternative ethical schools of thought. Many of the ideas discussed in Section 2.2 are concerned with particular outcomes, such as the development of chemical weapons (dangerous and likely to lead to mass death) or an increase in the prestige of chemical engineering as a field (potentially beneficial to science). It is therefore attractive to apply a consequentialist or utilitarian ethic.³⁸ I note that these areas of ethics

are diverse and not always in agreement, and for reasons of space and simplicity I will refrain from an in-depth analysis. At the risk of generalizing, a consequentialist might say something like, “If the net benefit is positive (*i.e.*, more benefits comes out of the action than harms), then the action is morally acceptable”. However, consequentialist and utilitarian views are particularly ill-suited for application in the field of chemistry.

A general problem with consequentialist and utilitarian ethics is *measurement*. To know if an action is good or just as a consequentialist, one must wait to see what happens after the action. But when does one stop “counting” benefits and harms? How widely must one consider impacts? The problem of measurement is particularly thorny in chemistry, where practitioners often do not immediately know how a new compound, material, or process could affect human or environmental wellbeing and where negative impacts are often obscured, *e.g.*, by chemical manufacturers.³⁹ In research contexts, there is perhaps even greater uncertainty about impacts. A researcher could argue that their work will improve sustainability (*i.e.*, it will have a positive net effect) as a way to justify ethically questionable behavior, but what happens if the researcher’s ideas don’t work out, if their experiments fail? Even worse, what if some results, though initially intended to benefit humankind, are used to a nefarious end? In such cases, an expected net benefit becomes a net harm.

Another problem with consequentialism and utilitarianism concerns the distribution of impacts. Most utilitarian frameworks aim to maximize wellbeing in a global sense.⁴⁰ A global optimum (if such a thing could be identified), or the point which creates “the greatest amount of good for the greatest number”, could include severe suffering in certain populations but still be considered moral and worth pursuing. This is dangerous when one considers that chemistry’s benefits tend to be diffuse, while its negative impacts are often concentrated.⁴¹ Industrial facilities such as combustion-based power plants and chemical factories are frequently located near marginalized communities,⁴² (*e.g.*, in the United States, they are more often located near Black and other racialized communities).⁴³ While the chemicals produced may help millions of people, the local impacts are often starkly different, with proximity to certain types of facilities being correlated with negative health outcomes.⁴⁴ Such a disparate outcome would directly go against the principles laid out in Section 2.1, including fairness and stewardship, and subjecting marginalized communities to disproportionate harm is inherently immoral outside of any chemistry-specific considerations.

For these reasons, chemical scientists must resist the urge to fall back on consequentialist and utilitarian arguments, either directly or in arguments based on related concepts like “net good”. In the remainder of this piece, I will refrain from such positions, grounding my arguments in moral ideals, rules, and responsibilities.

3 Ethical Risks and Harms of Foundation Models

In this section, I consider foundation models designed for general public use, mainly generative models like LLMs, VLMs, and diffusion models. I assume that the person using such models may or may not be a (chemical) scientist but that the application is not specialized to chemistry. Risks that are heightened in chemical applications or are unique to chemistry are discussed in Section 5 below.

3.1 Hegemonic Values and Prejudices

Generative foundation models like LLMs, VLMs, and diffusion models are trained on human-generated media, introducing human biases.

Alarming, it has been well documented that LLMs and diffusion models reproduce harmful prejudices and stereotypes, including sexism,^{45,46} racism,^{47,48} ableism,^{49,50} anti-LGBTQIA+ prejudice,⁵¹⁻⁵³ and more,⁵⁴ reflecting hegemonic positions. Though VLMs have been developed and deployed more recently, it appears that hegemonic biases are present in these models as well.^{55,56} For instance, VLMs produce starkly different responses to the same prompt based on the perceived gender and/or race of the people depicted in provided images.^{57,58}

Several techniques have been identified to mitigate foundation model prejudice, including direct preference modeling,^{59,60} but model bias remains a persistent challenge. Human preferences provide limited protection against prejudice. Indeed, human evaluators often cannot even agree on what statements are harmful,⁶¹ introducing significant noise. Even if evaluators were in perfect agreement, model bias is not always overt. LLMs can produce prejudicial responses with ostensibly non-prejudicial explanations,⁶² making the task of even identifying social biases challenging. Concept editing,⁶³ a recent development in which foundation models are selectively altered to remove harmful associations, is perhaps a more promising direction for bias mitigation, and further investigation in that direction should be conducted.

3.2 Plagiarism and Privacy

While the complete training corpora of many foundation models are unreported, it has been shown that LLMs repeat published materials verbatim without attribution.^{64,65} The legal status of foundation models is still in question, and it is possible that courts will not interpret foundation model repetition as copyright infringement. Regardless, reproducing text without attribution constitutes plagiarism in academic contexts. The ethical question for chemical ethics is: are individuals responsible for plagiarizing text *via* foundation model if they were unaware that the text is plagiarized?

LLMs provide a convenient excuse for unscrupulous researchers and writers, who can claim that they did not know or even that they could not know that their text was plagiarized. At the same time, conventional publication ethics would suggest that even an ignorant author is morally in the wrong. LLMs and other foundation models lack agency and complete tasks only in response to user instructions. If we treat foundation models as tools, then a writer is obviously responsible for their tool use. Regardless of whether LLMs are considered as agents or tools, all authors of a publication are responsible for the text.^{66,67} Even a writer who did not themselves use an LLM is ethically responsible if some of their text is plagiarized.

In the most egregious cases, LLMs do not only repeat another author’s words but “leak” private information contained in *e.g.*, training data.⁶⁸ Like efforts to reduce bias, research into improving privacy protections and mitigating leaking have shown limited success;⁶⁹ careful attack strategies can even overcome model “hallucination” (see Section 3.5) and tendency to “catastrophically forget” training data to extract target private information.⁷⁰ For those who have an ethical responsibility to maintain privacy, foundation model use with *e.g.*, medical, genetic, or patent data is highly suspect.

3.3 Environmental and Resource Considerations

The energy and water required to train and use LLMs are immense. The increased use of graphics processing units (GPUs) for ML model training and inference has directly been implicated in rapidly increasing energy consumption by data centers. In the just five years between

2018 to 2023 (a time period that notably saw the development and widespread use of generative foundation models like ChatGPT), the portion of United States energy consumption attributed to data centers increased from 1.9% to 4.4%.⁷¹

In addition to straining electrical grids, foundation models significantly contribute to greenhouse gas emissions and, thus, anthropogenic climate change. In principle, the energy to power data centers could be provided by renewable sources, but in practice this is not often the case, and many data centers rely on fossil fuels for power.⁷² In the recent explosion of interest in “AI”, many companies in the data and tech sectors are actually becoming less sustainable in terms of greenhouse gas emissions.⁷³

It is difficult to precisely estimate the environmental and climate impact of foundation models, and making such an estimate is outside of the scope of this work. Though exact figures will vary, training a single model can easily release dozens of tons of carbon dioxide into the atmosphere,⁷⁴ and this is only a small part of the model’s impact. Foundation model training is a fixed cost that takes place only once or perhaps a handful of times (depending on hyperparameter tuning), while the costs of fine-tuning and model inference are ever increasing.⁷⁵ ChatGPT alone receives billions of visits per month and responds to roughly 10,000,000 user queries daily,⁷⁶ providing a sense of the scale of this add-on cost.

Even if foundation models like ChatGPT were powered entirely by renewable energy, model water consumption is unsustainable. While the exact water consumption required to train and employ foundation models is not reported by “AI” companies like OpenAI, Anthropic, and Microsoft, Li et al.⁷⁷ note that data centers currently evaporate quantities of water that are comparable to the amounts used in the beverage industry, and, just as with energy, data center water usage is accelerating.⁷⁸ Though perhaps obvious, it bears mentioning that while beverages sustain life, data centers do not in any way directly promote human wellbeing. Placing foundation models within this broader context, Li et al. estimate that training a single GPT-3 model (a relatively small foundation model at the time of this writing) can consume 700,000 liters of water and that the same model conservatively requires on the order of one liter of water per 20-100 queries.

It is comforting for a chemical practitioner (particularly one concerned with sustainability) to take the long view. It has been suggested that foundation models could dramatically accelerate human progress, including in science.^{79,80} Yes, right now foundation models are (somewhat literally) bleeding the Earth dry, and yes, they’re worsening the ongoing climate catastrophe, but what if these models eventually lead to breakthroughs in climate technologies and sustainable energy? What if the current environmental harms are a necessary cost for a one-day utopia?

This argument falls into one of the traps of consequentialism that I outlined in Section 2.3 (*i.e.*, the problem of measurement), but as this is a particularly common line of reasoning, it deserves a direct refutation.

The hope that the deaths and destruction that “AI” is causing today may lead to a better future is a dangerous one. The future that this hope rests on is dubious at best, depending on the development of scientific breakthroughs (in particular, breakthroughs brought about by foundation models that could not be discovered at an acceptable rate without foundation models) and new technologies as well as massive changes in energy production. On the other hand, climate change, drought, famine, and other associated disasters are definite, existing now with as much certainty as empirical science

allows.^{81,82} Climate change and environmental degradation are existential threats to human society and entire ecosystems, and it is irresponsible and morally unacceptable to reserve judgment on environmental sustainability in the face of such severe risk.

3.4 Labor

There has been considerable discussion both in the scientific literature and in the popular press concerning the effect of foundation models and “AI” on labor. Most of this discussion centers on the question, “Will ‘AI’ take jobs away from humans?”

Foundation models remain a relatively new development in machine learning, and the full impact of these models on society has yet to be seen. Preliminary data suggests that there is real risk to human labor, which under Capitalism amounts to a risk to human security and wellbeing. Hui et al. found that generative models such as ChatGPT and Midjourney cause the employment and economic wellbeing of freelance workers to suffer.⁸³ Skilled freelancers who produce high-quality work consistently appear to be disproportionately affected, undermining any argument that “AI” will only affect “unskilled” labor. Through a literature meta-analysis, Zarifhonarvar⁸⁴ found that roughly two-thirds of occupations will be affected by LLMs like ChatGPT (other foundation models were not considered in the analysis). Zarifhonarvar suggests that roughly one-third of occupations face “full impact” from LLMs, meaning that all tasks and skills involved in the occupation can be automated. These “full impact” occupations face mass job loss. While Guliyev takes a contrary view,⁸⁵ using an analysis of 24 developed countries to suggest that “AI” increases employment and worker wellbeing, his analysis concludes in 2021, early in the current explosion of foundation models, and therefore may no longer apply.

Foundation model development presents its own troubling labor ethics. Training a model on the scale of a modern LLM or generative diffusion model is challenging, involving data collection, data cleaning, and model testing and fine-tuning, and labor abuses happen throughout this process.

It can be inferred from the models’ ability to plagiarize (Section 3.2) that works under copyright are included in model training. Artists, writers, and other human creators whose work is incorporated into foundation models are uncompensated. Thus foundation models represent theft of intellectual property, labor (particularly creative, artistic, and intellectual labor), and wages/compensation on an unprecedented scale.

To avoid producing undesirable and illegal material (*e.g.*, hate speech or images depicting child pornography), foundation model datasets are cleansed of offending materials. The labor of labeling offending materials, along with model testing, is often conducted by underpaid workers in the Global South.⁸⁶ The conditions that these workers face has been described as exploitative,⁸⁷ and even if the workers were treated and compensated fairly, the work of weeding out some of the worst content that humanity has created is inherently traumatizing. While some may see generative foundation models and “AI” as bringing about a future where humans no longer need to work, this dream ignores the widespread damage that the same models are producing for human workers today.

3.5 Epistemological Risk

Foundation models, particularly LLMs, VLMs, and related generative methods, challenge the very notion of truth. Whereas it has long been said that “seeing is believing”, now it is possible, even easy, to generate as much text as one wants about any topic and to gen-

erate photorealistic images and videos depicting events that never happened. “Deepfake” images and videos⁸⁸ depicting real people can damage individual reputations⁸⁹ and manipulate public opinion.^{90,91} LLMs can generate outputs that is increasingly difficult to distinguish from human-generated text.^{92,93} On a societal level, foundation models represent serious threats to how we make sense of the world and form our beliefs.

Over the last several years, an enormous effort has been put into reducing erroneous output from LLMs (so-called “hallucination”),^{94–96} The basic premise is that, with enough data, fine-tuning, and specific instructions, LLMs will cease contributing to misinformation, and users seeking to use LLMs to produce technically accurate text will be able to do so without worry. Even if erroneous LLM output has decreased over time, this project will always necessarily be incomplete. LLMs are stochastic text generators;^{97,98} they are trained (and, barring a dramatic change in model architecture and training approach, will continue to be trained) to produce *statistically likely* text. Given that human beings, who produced the vast majority of training data for LLMs and related models, are frequently wrong and frequently lie using natural language, a statistically likely response is not and can never be guaranteed to be accurate.

It is not accurate to say that LLMs lie (which would imply that they know the truth and choose to say otherwise) or that they hallucinate (which would imply knowing the truth but somehow being deceived by an altered “mental” state to produce erroneous data). Rather, to use the philosophical term of Frankfurt,⁹⁹ they bullshit.¹⁰⁰ On a basic level, they have no concern for truth and whether their outputs are true or false, and thus they are harmful to truth-seeking enterprises.

But what of those who seek to intentionally deceive through deepfakes and generative, human-like text? There have been two major approaches to address this risk. A number of adversarial ML models have been trained to distinguish between foundation model-generated and human-generated/real media.^{101,102} Outside of data-driven methods, there have been a number of guides published in the popular and scientific press attempting to train human consumers of media to distinguish real from model-generated images and videos.^{103,104} Both approaches seem promising, though as foundation model capacity increases, it is difficult to see a resolution to this challenge of distinguishing reality from stochastic fiction. Rather, it is likely that the arms race between generative model developers and truth-seeking individuals will continue.

The epistemological risks and harms associated with foundation models are an existential threat to science. How can one trust another scientist’s article or book if they do not know if that scientist actually wrote it or if an LLM generated the text? How can one trust an image that could easily have been generated by the authors to tell a convenient or desirable narrative? Note that this line of inquiry does not require one to assume that any researcher intends to harm others. It only requires that researchers be willing to cut corners to publish, which, given increasing demands to “publish or perish”^{105,106} in an increasingly tight academic job market,¹⁰⁷ is plausible.

If the mutual trust that undergirds scientific inquiry is disrupted, then “every researcher for themselves” could become the norm, with more and more time being spent trying to verify claims rather than using established studies to move one’s own research forward. Indeed, in the most extreme (though unlikely) case of widespread distrust, science could cease altogether to be a community effort if researchers feel that they could more effectively produce more solid results in isolation than by consulting a fundamentally flawed literature.

Finally, while my analysis here has been mainly structural, focusing on knowledge on a community and societal level, it is important not to ignore the risk of foundation models on individual knowledge. A recent study found that those who place confidence in generative “AI” tools like LLMs engage in less critical thinking overall.¹⁰⁸ The authors also found that the use foundation models resulted in a perceived qualitative change in how critical thinking is used; in terms of Bloom’s taxonomy, comprehension (*i.e.*, knowledge organization and summarizing) and synthesis (*i.e.*, generating new ideas by combining different concepts, or devising new meanings and interpretations) are most negatively impacted by foundation model use, while evaluation (*i.e.*, judging information) was least negatively impacted. This study is limited, for instance because it relies entirely on self-reported survey data, but it nonetheless points to the chilling idea that many are treating statistical models as a principal source of knowledge and not their own minds.

4 Squaring Foundation Models with Chemical Ethics

Chemical scientists are already using generative foundation models in research contexts, and doubtless, these include non-chemistry-specific applications (*e.g.*, image generation, brainstorming, or code generation). But can this general foundation model use align with the ideals listed in Section 2.1 and the practical guidelines discussed in Section 2.2? That is, it is ethical for chemical scientists to lean on foundation models? In this section, I consider general-purpose applications of foundation models and associated ethical challenges, as discussed in Section 3; specific ethical conflicts related to chemical applications will be explored in Section 5 below.

4.1 Failing Science’s Ideals

The epistemological risks of foundation models are likely to lead to violations of several scientific ideals. As bullshit engines, foundation models have no notion of truth and frequently produce flawed or untrue outputs. If researchers using these models are not careful, they could develop flawed text or code and communicate falsehoods, ultimately undermining trust in scientific communications and thereby weakening scientific relationships and hampering scientific stewardship. It is also possible that researchers could use foundation models to avoid accountability, using excuses along the lines of, “It wasn’t me, it was the machine!” At least within the context of scientific publishing, accountability diffusion does not seem like a serious problem, as a number of journals and publishers have already provided ethical guidance reaffirming that authors must remain individually and collectively accountable.^{109,110} As long as publication outlets consistently apply and enforce their guidelines, adherence to the ideal of accountability can be maintained.

The ideals of science are even more fundamentally incompatible with foundation models in light of the models’ embedded prejudices. While computational and data-driven models are often praised for their objectivity,¹¹¹ foundation models are informed by human subjective biases and reproduce them in a range of tasks. The preferential, prejudicial behavior of many foundation models is also tautologically unfair and unjust.

Even aside from the epistemological damage that foundation models can cause as bullshit engines, foundation models risk the ideal of stewardship. By repeating ideas without appropriate attribution, foundation models can damage relationships within the scientific community, erasing researchers’ contributions and implicitly eliminating them from scientific conversations. Relations (especially

between scientists and the broader human community) are strained by the toxic labor relations baked into foundation model development and implementation, and foundation models are further worsening land relations through their unsustainable energy and water consumption.

Overall, it is difficult to imagine a consistent argument that foundation models are aligned with science's stated ideals. In fact, from this non-exhaustive analysis, it appears that the only ideal that is not obviously violated or threatened by the development and/or application of LLMs, VLMs, and diffusion-based image generators is openness. Even that principle can only be maintained provided that individuals explain transparently how they have used the foundation models.

4.2 Violating Professional Chemical Ethics

The practical chemist might roll their eyes at the preceding discussion. This would be fair. Even within this work (Section 2.1), I acknowledge that scientists rarely if ever live and work perfectly in line with their ideals. While it is deeply troubling in just how many ways and to what extent foundation models challenge science's ideals, it is hardly surprising that they are not in full alignment. However, foundation models are just as poorly aligned with the expectations of the chemical profession, which should discomfort even the least idealistic chemist.

At present, foundation models are causing active harm to the environment, and given the increasingly catastrophic effects of anthropogenic climate change and global warming, this also threatens the wellbeing of future generations. As discussed in Section 3.3, the environmental degradation associated with foundation models is a black mark against them, even if they are used towards ostensibly sustainable, environmentally benign, or beneficial ends.

In that foundation models serve to accelerate anthropogenic climate change, they are indirectly implicated in associated climate catastrophes that threaten human and non-human health, safety, and wellbeing. The chemists' duty to protect security is also compromised by foundation models' proclivity towards repeating sensitive information.

Most damningly, foundation models such as LLMs, VLMs, and generative image and video diffusion models go against many central and specific intellectual norms of chemical practice. Whereas chemical scientists are urged to share ideas with appropriate attribution while avoiding plagiarism, LLMs function as plagiarism machines. Chemists are required to be honest and increase knowledge, but foundation models produce information without care towards truth, spreading misinformation.

The implications of this analysis are clear: uncritical development, promotion, and use of the types of foundation models discussed above put chemical scientists at odds with their professional obligations. It could even be said that use of LLMs, VLMs, and related general-purpose foundation models inherently constitutes a misuse of the underlying technologies (e.g., statistical learning, transformer architectures, and/or reinforcement learning), further deepening the betrayal of professional duties.

5 Chemical Foundation Model Use

The purpose of this section is not to review either the uses of generative foundation models in the chemical sciences or chemistry-specific foundation models. A number of thorough reviews on both topics have already appeared in the literature.^{112–116} Rather, I will briefly and non-

exhaustively survey different areas where these models can be used, in doing so addressing potential ethical problems and how they align or fail to align with scientific and chemical ethics.

5.1 Chemical Applications of Generative Models

I wish to begin by emphasizing that any use of a generative model will face many ethical problems, as outlined above. It is presently impossible to completely avoid unsustainable model water and energy usage (assuming that individuals have no control over how a data center is operated), the abuses of labor inherent in model development, and the epistemological risks of bullshitting models, among other problems. With these risks taken as a given, I will in this section emphasize unique risks and harms that arise when generative models are used in the chemical sciences and emphasize areas where certain general risks are particularly dangerous or likely to play a role. I note that this section will primarily focus on LLMs, as these have been most widely used in the chemical literature.

5.1.1 Literature Review and Text Generation

Many chemistry-specific applications of generative foundation models share significant similarities with applications in other areas. For instance, there has been considerable interest into using LLMs for chemical information retrieval and question-answering tasks.¹¹⁷ LLMs appear to be particularly useful in generating structured data from unstructured text,^{7,118} which is beneficial in e.g., constructing databases of synthesis recipes. Alone, LLMs are not particularly well equipped to answer complex questions, particularly when calculations are required, but models that can generate and execute code through external tools are considerably more effective.⁶

Naturally, information extraction and generation are domains where epistemological risk is high. While providing LLM "agents" with external tools decreases the likelihood of erroneous model outputs, so-called "hallucinations" cannot be entirely avoided, meaning that the data and/or answers generated by LLMs and other generative models are potentially flawed.

Tasks which rely on an LLM "interpreting" text and generating unstructured or structured responses are also open to model biases. Consider that much work in chemical research and the chemical industries relates directly to humans, including biochemistry, drug design, and other areas of medicinal chemistry. Omiye et al.¹¹⁹ have shown that LLMs promote inaccurate and harmful race-based ideas related to medical practice. Similarly, Yuzhe Yang et al.⁵⁷ found considerable bias in VLM analysis of medical images, with foundation models diagnosing marginalized groups at a lower rate than their majority counterparts. Yang's study also highlighted the effect of intersectionality, with multiply marginalized individuals (e.g., Black women) receiving even worse results. While further research in this area must be conducted, it is plausible that prejudice may similarly color foundation model-based inquiries into a wide range of biochemical and medical domains where human markers like gender, sex, and race are salient. If the summaries and answers generated by LLMs are racist and sexist, they could damage the integrity and (further) undermine the objectivity of the scientific work while also potentially having real negative impacts on human beings.

Additional risks arise when LLMs are used for text that will be published or shared. The use of LLMs in scientific writing has rapidly increased in recent years,¹²⁰ and there are some reported cases of LLMs being used in peer-review reports.^{121,122} While there may be legitimate use cases for LLMs in these areas of scientific writing (ig-

noring issues like plagiarism and copyright for the moment), such as correcting grammar and improving clarity or conciseness, LLMs could produce erroneous or misleading statements or provide non-existent citations. LLMs used for peer review could produce unreasonable arguments that ultimately result in the wrongful rejection of a scientifically sound manuscript. As noted in Section 4.1, many journals at least require open disclosure of generative model use, which allows critical readers to closely and skeptically examine a piece of LLM-generated text. Still, the epistemological risks of LLMs for the scientific literature have not been fully reckoned with.

5.1.2 Laboratory Automation

As part of a growing movement towards automated high-throughput experiments and “self-driving laboratories”,^{123,124} generative foundation models (especially LLMs) have been used to direct complex chemical processes. To accomplish this, an LLM is given access to an applications programming interface (API) for a laboratory automation system, which could control an individual robot or other instrument or could send instructions to an existing automated laboratory. Thus far, this approach has been demonstrated on well-studied systems. Boiko et al.⁸ applied their LLM-based laboratory automation system, Coscientist, to optimize Suzuki and Sonogashira cross-coupling reactions, while Bran et al.’s ChemCrow¹²⁵ demonstrated the ability to synthesize non-trivial products such as N,N-Diethyl-meta-toluamide (DEET). In the area of optimization, LLM-based approaches appear to be competitive with the more traditional Bayesian optimization. It is unclear how models like ChemCrow and Coscientist, which in addition to laboratory APIs can access the chemical literature through the Internet, will be able to generate and execute experimental plans to study truly novel chemical or materials systems for which no literature recipes exist.

Granting an LLM control of physical laboratory tools generates many novel ethical concerns. Boiko et al. found that LLMs such as GPT-4 could be directed to synthesize compounds that are dangerous or illegal if given an appropriate prompt, bypassing the model’s “guard rails”. This example poses LLMs as dual-use technologies: just as they could be used to perform experiments for benign, socially beneficial chemical research, they can also be leveraged by bad actors, considerably lowering the barrier to entry for illicit drug or chemical weapons manufacturing. While this does not make LLM use unethical *per se*, it makes adherence to the chemist’s duty to avoid developing destructive technologies more challenging.

Laboratory safety is another pressing concern. An LLM cannot be trained to follow safety procedures as a human can, nor can an LLM be held accountable for unsafe behavior. Unless a human is closely monitoring the model’s actions (and, in doing so, somewhat defeating the purpose of laboratory automation), the safety of humans in the laboratory and of the experimental equipment cannot be guaranteed.

5.1.3 Classification and Regression Tasks

Many recent efforts have considered if general-purpose foundation models can act as substitutes for specialized chemical and materials ML models. The idea is simple: if, by being trained on a huge corpus including the open literature, an LLM or similar model has some underlying “understanding” (representation) of chemistry, then it should be possible to fine-tune these models to make arbitrary predictions about chemical entities. It appears that LLMs are able to compete and in some cases even exceed small, single-purpose ML models when available data are scarce.¹²⁶ In particular, LLMs seem to excel at classification and struggle somewhat more on regression and gener-

ation tasks.¹¹⁷ When ample data are available, however, LLMs do not provide significant benefits. This presents researchers with a choice: they can either generate additional data to train a bespoke model, which might be time-consuming or costly, or they can rely on an LLM, with all of its thorns and pitfalls.

5.1.4 Education

Some groups have advocated for using LLMs in educational settings, serving as individualized “tutors” that can assist students in the classroom.^{127,128} They argue that LLMs provide a “scalable” solution to meet individual student needs and improve education in various areas, from providing lectures and guiding student labs to evaluating students.

LLM-based chemical education is harmful to students and to broader safety. Students may come to believe falsehoods constructed by their “tutors” or may learn correct information without understanding the underlying mechanisms and logic. Students could also come to rely on LLMs, reducing their ability to think critically and preventing them from developing useful cognitive skills, e.g., chemical intuition. LLMs directing laboratories introduce risks to health and safety, as discussed in Section 5.1.2, but in an educational setting this risk is considerably greater. Students do not necessarily understand laboratory safety; indeed, teaching safety skills is a core goal of any laboratory course. Therefore, they cannot effectively act to monitor the safety of LLM instructions. Finally, it should go without saying that having a prejudicial “tutor” is harmful to an educational experience, especially for minoritized or marginalized students. As even some of the authors advocating for LLM-based chemical education realize,¹²⁸ LLMs cannot completely avoid prejudice, and model bias is likely to influence any LLM-generated educational materials.

5.2 Chemical Foundation Models

The term “foundation model” is somewhat imprecise, and the line between a conventional ML model and a foundation model is blurry. There is no exact number of training data or parameters past which a model becomes foundational. Keeping this in mind, a number so-called chemical foundation models have been developed and reported recently.

Perhaps the most famous family of foundation models are protein models, which won part of the Nobel Prize in Chemistry in 2024.¹²⁹ Given a string of amino acids, these models generate folded protein structures, thus addressing one of the greatest challenges of biology and biochemistry. In some cases, protein model architectures resemble LLMs. These “protein language models”^{116,130} treat amino acids as tokens and protein sequences as text. Other models, including the AlphaFold family^{16,131} and RoseTTAFold,¹¹ rely on other advanced ML techniques, such as equivariant transformers¹³² and diffusion layers.¹³³

Several research groups have developed so-called “universal” machine learning interatomic potentials (MLIPs).^{12,13,134} By training on data spanning the periodic table, these models can be employed to simulate diverse chemical systems, including some very far outside of the training distribution. For instance, the MACE-MP-0 MLIP,¹³ which was trained exclusively on bulk crystalline solids from the Materials Project,¹³⁵ was shown to behave well (at least qualitatively) on systems ranging from gas-phase hydrogen combustion to liquid electrolytes and many systems in between.

Finally, there have been a handful of chemical foundation models that do not address single problems like protein folding or atomistic

energetics and dynamics. Similar to a protein language model, Cai et al. created ChemFM,¹³⁶ a foundation model trained on molecules represented as Simplified Molecular Input Line Entry System (SMILES) strings for molecular generation and property prediction tasks. Aiming to tackle the somewhat more narrow domain of reactor modeling, Wang and Wu recently designed a foundation model that combines meta-learning with physics-informed fine tuning to accurately simulate multiple classes of reactors.¹³⁷

These innately chemical foundation models avoid a number of the ethical problems of general-purpose generative foundation models. While human bias will always in some way inform model datasets, these biases are less directly connected to systems of oppression, and, for instance, a “universal” MLIP cannot generate racist, sexist, or otherwise prejudiced output. These models are often trained on open data, such as the structures deposited in the Protein DataBank (PDB) or calculated through the Materials Project database, and so the labor ethics of chemical foundation models are more sound. These data sources additionally do not introduce risks related to privacy or plagiarism; though a protein model could, conceivably, exactly reproduce a previously reported protein structure, this does not necessarily amount to plagiarism. If the structure used is in the public domain, as is the case for all data deposited in the PDB, then reproduction is valid unless a researcher claimed to have *discovered* that exact structure without proper attribution.

The main ethical issue that these chemical foundation models introduce relates to resource use. Though universal potentials, protein models, and their ilk are small in comparison to generative foundation models, they are large by the standards of (bio)chemistry and likely contribute more to environmental degradation and climate change than an average ML model in the chemical sciences, considering that general-purpose models often have higher energy demands than specialized models.^{138,139}

6 Discussion

6.1 Addressing the Harms of Foundation Models

Generative foundation models are, in many ways, at odds with scientific ethics generally and the professional ethics of chemistry, chemical engineering, and related fields more specifically. Chemistry-specific foundation models such as “universal” MLIPs and protein models are less morally compromised but still introduce some ethical risk and harm. Considering the diverse concerns laid out above, I ask a simple and necessary question: should chemical scientists be using or supporting the use of foundation models? The appropriate response hinges on the extent to which the model in question is *necessary* for a particular task.

Consider how chemical scientists address problematic tools in other contexts. After all, ethically pure chemical research and practice is rare, and one can easily generate a list of such problematic tools. To name just a couple of examples of problematic chemical substances:

- A number of commonly used solvents are known to be highly harmful to human health and the environment;¹⁴⁰
- Plastics, which degrade human health and the environment,^{141,142} are ubiquitous in our laboratories and are a major product of the chemical industry;
- Cobalt, a key component in modern lithium-ion batteries, is linked with (alleged) human rights and labor abuses in the Democratic Republic of the Congo.^{143,144}

In response to the clear harm that these substances caused, the chemical community shifted in a number of significant ways. For harmful solvents and for cobalt, researchers have sought ways to reduce the amount needed or avoid these materials altogether, leading to the explosion of interest into so-called “green solvents”^{145,146} and Co-free battery electrodes.^{147,148} While the original materials are still used (e.g., Co-based Li-ion batteries remain on the market and an area of research interest), the chemical community is largely moving towards elimination. There is also some interest in minimizing the use of plastics,¹⁴⁹ but there has been a greater emphasis on developing alternative chemistries that are environmentally benign^{150,151} or identifying methods to circularly recycle plastics,^{152,153} thereby reducing or eliminating the harm caused by their waste.

By analogy, once we as a community identify foundation models (particularly generative foundation models) as harmful or ethically unacceptable, we can seek to minimize their use. We should be asking, “Where are LLMs, VLMs, etc. offering truly unique, necessary benefits?” and eliminating any non-“essential” uses. This would mean, for instance, refusing to use LLMs for knowledge retrieval, writing, or code generation, all of which can be accomplished effectively (perhaps even more effectively) without the foundation models. For areas where foundation models are presently necessary, such as protein folding, researchers should seek to develop alternatives that are less harmful and/or seek out ways to substantively reduce harm. I note that, while these changes can be enacted by individual chemical practitioners, they will be more effective if implemented structurally, whether at the level of universities and other research bodies, chemical professional societies, or policy-makers.

6.2 A Turning Point for the Chemical Sciences

I have argued throughout this piece that, if a new technology and established ethics are at odds, then the technology should not be used, or steps should be taken to minimize the harm done by that technology. An idea central to my analysis but unspeakable up to this point is that ethical norms should not be altered, that the ideals that have been laid out and the norms that we have collectively agreed to remain sound.

This is not the only possible resolution to the conflict of foundation models and ethics, though. To avoid any claims that I present only one side of an argument, I must mention the alternative course of action. Rather than avoid foundation models or severely limit their use, chemistry and related fields could redefine what it is to be moral to accommodate the presence of “AI” in our society and our professional lives.

In a sense, this redefinition is actively taking place. Though none of the professional codes that I discussed in Section 2.2 have been publicly amended, major chemistry conferences and publications allow scientific works based on LLMs and other foundation models with few limitations. While some publication guidelines require that authors of a manuscript disclose if and how foundation models have been used in the writing, the door remains open for stochastic plagiarism and misinformation, among other possible negative outcomes. It is logical on the surface to trust that a manuscript’s authors will review their work to ensure accuracy and check that all ideas have been properly attributed. However, it may not be obvious at all that a model has plagiarized or repeated a falsehood. Moreover, by using LLMs or related tools, authors demonstrate a willingness to cut corners and potentially reduce their critical thinking, suggesting that this trust may be misplaced.

The chemical community must critically reflect on its values, assessing what ethics are appropriate for the modern field and modern society and to what extent these values and related norms are important. This latter point is critical. Even if the community forcefully reaffirms a dedication to ethical practice, there will almost certainly be some who behave unethically. As mentioned in Section 2.2, though, there are few mechanisms in place to call upon when unethical behavior is identified. Regardless of what chemistry's values are, they mean little if the field fails to respond to their violation.

6.3 Lingering Questions

The present work is not meant to be definitive but instead was written to open a necessary and until now largely ignored conversation. In service of this aim, I close this article by considering what I have not been able to address and what future research — ethical and technical, theoretical and empirical — should be conducted to reach a satisfactory conclusion.

I have chosen in this work to limit my scope to the chemical sciences, focusing mainly on chemistry and chemical engineering. This is mainly a reflection of my own disciplinary comfort and background. As a chemist without significant training in ethics or philosophy, I feel equipped to tackle moral issues in the chemical sciences, while I recognize that I may be unequipped to address broader questions. Nonetheless, broader questioning is needed.

As a starting point, it would be worthwhile to compare the ethics of foundation models in the the chemical sciences, as discussed here, with those in the biomedical fields, where a body of work in related moral concerns is growing.¹⁵⁴⁻¹⁵⁶ In particular, the thought put into regulatory changes in the biomedical literature¹⁵⁷ may be relevant for discussions of scientific ethics. Beyond medicine, it is worth asking to what extent the ideas presented here generalize to other areas of science and engineering. Ostensibly all areas of science should be guided by the same ideals, but different professions may have different norms and different priorities. As one example, though the chemical sciences are intimately connected with and concerned for the environment (and, thus, environmental sustainability), the same may not be said for other fields such as astronomy.

Even within the more narrow domain of the chemical sciences, there are many further questions worth exploring. While I have discussed the general practice of chemistry and associated norms, I have focused in ways on chemical researchers. How, if at all, do the conclusions drawn here change from the lens of, e.g., an industrial chemical engineer, where, e.g., publication ethics may not be important, where an individual may not have control over which tools they use (*i.e.*, they may be required to use “AI” tools by their employer), and where the interests of the employer frequently put practitioners at odds with ethical (especially environmental) norms?

6.4 Conclusions

Like microplastics and greenhouse gases, large ML models surround us. Now that foundation models, from LLMs and generative diffusion models to “universal” MLIPs, have made their way into society and into scientific practice, the chemical profession(s) must critically examine these technologies, assessing their risks and outcomes under consistent ethical standards. Here, I have begun this work, addressing the questions “Are foundation models in alignment with scientific and chemical ethics?” and “How should chemical practitioners interact with and around foundation models?”

To the first question, an analysis of common ideals and norms reveals a simple answer: mostly, no. From training to testing and evaluation, general-purpose generative foundation models violate most of the ethical principles that I have identified and described. While chemistry-specific foundation models, such as protein models and “universal” MLIPs, avoid many of the ethical problems of more general models like LLMs and VLMs, they still have a negative environmental impact that is at odds with the modern understanding of chemistry's obligations to sustain and protect the natural environment.

The second question — now that foundation models are a part of scientific practice, what should be done — is more challenging. Drawing from the green chemistry, sustainable chemistry, and battery fields, I suggest that an important step that would bring chemical practitioners better in alignment with the morals of chemical science would be to reduce the use of foundation models to truly necessary applications and to, over time, shift focus to identifying alternative, ethical methods and technologies that can fill those same unique niches.

There is much more work to be done at the intersection of data science, chemical science, and ethics. I hope that my account inspires further theoretical and technical inquiries in this nascent and much-needed area of study.

Conflicts of Interest

I have no conflicts of interest to declare.

Model Use

I declare that I at no point intentionally and knowingly used any generative foundation model, including but not limited to LLMs, VLMs, and image/video generators, during the process of writing this manuscript.

Acknowledgements

I thank the Carnegie Bosch Institute and the Carnegie Mellon University Department of Chemical Engineering for providing funding to support my research. I give my thanks as well to Julia Isabelle McKown, who provided support and useful feedback on this manuscript.

Notes and References

- (1) Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; Herbert-Voss, A.; Krueger, G.; Henighan, T.; Child, R.; Ramesh, A.; Ziegler, D.; Wu, J.; Winter, C.; Hesse, C.; Chen, M.; Sigler, E.; Litwin, M.; Gray, S.; Chess, B.; Clark, J.; Berner, C.; McCandlish, S.; Radford, A.; Sutskever, I.; Amodei, D. Language Models Are Few-Shot Learners. In *Advances in Neural Information Processing Systems*; Curran Associates, Inc., 2020; Vol. 33, pp 1877-1901.
- (2) DeepSeek-AI; Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; Zhang, X.; Yu, X.; Wu, Y.; Wu, Z. F.; Gou, Z.; Shao, Z.; Li, Z.; Gao, Z.; Liu, A.; Xue, B.; Wang, B.; Wu, B.; Feng, B.; Lu, C.; Zhao, C.; Deng, C.; Zhang, C.; Ruan, C.; Dai, D.; Chen, D.; Ji, D.; Li, E.; Lin, F.; Dai, F.; Luo, F.; Hao, G.; Chen, G.; Li, G.; Zhang, H.; Bao, H.; Xu, H.; Wang, H.; Ding, H.; Xin, H.; Gao, H.; Qu, H.; Li, H.; Guo, J.; Li, J.; Wang, J.; Chen, J.; Yuan, J.; Qiu, J.; Li, J.; Cai, J. L.; Ni, J.; Liang, J.; Chen, J.; Dong, K.; Hu, K.; Gao, K.; Guan, K.; Huang, K.; Yu, K.; Wang, L.; Zhang, L.; Zhao, L.; Wang, L.; Zhang, L.; Xu, L.; Xia, L.; Zhang, M.;

- Zhang, M.; Tang, M.; Li, M.; Wang, M.; Li, M.; Tian, N.; Huang, P.; Zhang, P.; Wang, Q.; Chen, Q.; Du, Q.; Ge, R.; Zhang, R.; Pan, R.; Wang, R.; Chen, R. J.; Jin, R. L.; Chen, R.; Lu, S.; Zhou, S.; Chen, S.; Ye, S.; Wang, S.; Yu, S.; Zhou, S.; Pan, S.; Li, S. S.; Zhou, S.; Wu, S.; Ye, S.; Yun, T.; Pei, T.; Sun, T.; Wang, T.; Zeng, W.; Zhao, W.; Liu, W.; Liang, W.; Gao, W.; Yu, W.; Zhang, W.; Xiao, W. L.; An, W.; Liu, X.; Wang, X.; Chen, X.; Nie, X.; Cheng, X.; Liu, X.; Xie, X.; Liu, X.; Yang, X.; Li, X.; Su, X.; Lin, X.; Li, X. Q.; Jin, X.; Shen, X.; Chen, X.; Sun, X.; Wang, X.; Song, X.; Zhou, X.; Wang, X.; Shan, X.; Li, Y. K.; Wang, Y. Q.; Wei, Y. X.; Zhang, Y.; Xu, Y.; Li, Y.; Zhao, Y.; Sun, Y.; Wang, Y.; Yu, Y.; Zhang, Y.; Shi, Y.; Xiong, Y.; He, Y.; Piao, Y.; Wang, Y.; Tan, Y.; Ma, Y.; Liu, Y.; Guo, Y.; Ou, Y.; Wang, Y.; Gong, Y.; Zou, Y.; He, Y.; Xiong, Y.; Luo, Y.; You, Y.; Liu, Y.; Zhou, Y.; Zhu, Y. X.; Xu, Y.; Huang, Y.; Li, Y.; Zheng, Y.; Zhu, Y.; Ma, Y.; Tang, Y.; Zha, Y.; Yan, Y.; Ren, Z. Z.; Ren, Z.; Sha, Z.; Fu, Z.; Xu, Z.; Xie, Z.; Zhang, Z.; Hao, Z.; Ma, Z.; Yan, Z.; Wu, Z.; Gu, Z.; Zhu, Z.; Liu, Z.; Li, Z.; Xie, Z.; Song, Z.; Pan, Z.; Huang, Z.; Xu, Z.; Zhang, Z.; Zhang, Z. DeepSeek-R1: Incentivizing Reasoning Capability in LLMs via Reinforcement Learning. *arXiv preprint arXiv:2501.12948* **2025**.
- (3) Radford, A.; Kim, J. W.; Hallacy, C.; Ramesh, A.; Goh, G.; Agarwal, S.; Sastry, G.; Askell, A.; Mishkin, P.; Clark, J.; Krueger, G.; Sutskever, I. Learning Transferable Visual Models From Natural Language Supervision. *arXiv preprint arXiv:2103.00020* **2021**.
- (4) Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*; 2022; pp 10684–10695.
- (5) Liu, H.; Li, C.; Wu, Q.; Lee, Y. J. Visual Instruction Tuning. *arXiv preprint arXiv:2304.08485* **2023**.
- (6) D. White, A.; M. Hocky, G.; A. Gandhi, H.; Ansari, M.; Cox, S.; P. Wellawatte, G.; Sasmal, S.; Yang, Z.; Liu, K.; Singh, Y.; Ccoa, W. J. P. Assessment of Chemistry Knowledge in Large Language Models That Generate Code. *Digital Discovery* **2023**, 2 (2), 368–376.
- (7) Ai, Q.; Meng, F.; Shi, J.; Pelkie, B.; W. Coley, C. Extracting Structured Data from Organic Synthesis Procedures Using a Fine-Tuned Large Language Model. *Digital Discovery* **2024**, 3 (9), 1822–1831.
- (8) Boiko, D. A.; MacKnight, R.; Kline, B.; Gomes, G. Autonomous Chemical Research with Large Language Models. *Nature* **2023**, 624 (7992), 570–578.
- (9) Ruan, Y.; Lu, C.; Xu, N.; He, Y.; Chen, Y.; Zhang, J.; Xuan, J.; Pan, J.; Fang, Q.; Gao, H.; Shen, X.; Ye, N.; Zhang, Q.; Mo, Y. An Automatic End-to-End Chemical Synthesis Development Platform Powered by Large Language Models. *Nature Communications* **2024**, 15 (1), 10160.
- (10) Jumper, J.; Evans, R.; Pritzel, A.; Green, T.; Figurnov, M.; Ronneberger, O.; Tunyasuvunakool, K.; Bates, R.; Žídek, A.; Potapenko, A.; Bridgland, A.; Meyer, C.; Kohl, S. A. A.; Ballard, A. J.; Cowie, A.; Romera-Paredes, B.; Nikolov, S.; Jain, R.; Adler, J.; Back, T.; Petersen, S.; Reiman, D.; Clancy, E.; Zielinski, M.; Steinegger, M.; Pacholska, M.; Berghammer, T.; Bodenstein, S.; Silver, D.; Vinyals, O.; Senior, A. W.; Kavukcuoglu, K.; Kohli, P.; Hassabis, D. Highly Accurate Protein Structure Prediction with AlphaFold. *Nature* **2021**, 596 (7873), 583–589.
- (11) Baek, M.; DiMaio, F.; Anishchenko, I.; Dauparas, J.; Ovchinnikov, S.; Lee, G. R.; Wang, J.; Cong, Q.; Kinch, L. N.; Schaeffer, R. D.; Millán, C.; Park, H.; Adams, C.; Glassman, C. R.; DeGiovanni, A.; Pereira, J. H.; Rodrigues, A. V.; Dijk, A. A. van; Ebrecht, A. C.; Opperman, D. J.; Sagmeister, T.; Buhlheller, C.; Pavkov-Keller, T.; Rathinaswamy, M. K.; Dalwadi, U.; Yip, C. K.; Burke, J. E.; Garcia, K. C.; Grishin, N. V.; Adams, P. D.; Read, R. J.; Baker, D. Accurate Prediction of Protein Structures and Interactions Using a Three-Track Neural Network. *Science* **2021**, 373 (6557), 871–876.
- (12) Deng, B.; Zhong, P.; Jun, K.; Riebesell, J.; Han, K.; Bartel, C. J.; Ceder, G. CHGNet as a Pretrained Universal Neural Network Potential for Charge-Informed Atomistic Modelling. *Nature Machine Intelligence* **2023**, 5 (9), 1031–1041.
- (13) Batatia, I.; Benner, P.; Chiang, Y.; Elena, A. M.; Kovács, D. P.; Riebesell, J.; Advincula, X. R.; Asta, M.; Avaylon, M.; Baldwin, W. J.; Berger, F.; Bernstein, N.; Bhowmik, A.; Blau, S. M.; Cărare, V.; Darby, J. P.; De, S.; Pia, F. D.; Deringer, V. L.; Elijošius, R.; El-Machachi, Z.; Falcioni, F.; Fako, E.; Ferrari, A. C.; Genreith-Schriever, A.; George, J.; Goodall, R. E. A.; Grey, C. P.; Grigorev, P.; Han, S.; Handley, W.; Heenen, H. H.; Hermansson, K.; Holm, C.; Jaafar, J.; Hofmann, S.; Jakob, K. S.; Jung, H.; Kapil, V.; Kaplan, A. D.; Karimitari, N.; Kermode, J. R.; Kroupa, N.; Kullgren, J.; Kuner, M. C.; Kuryla, D.; Liepuoniute, G.; Margraf, J. T.; Magdău, I.-B.; Michaelides, A.; Moore, J. H.; Naik, A. A.; Niblett, S. P.; Norwood, S. W.; O'Neill, N.; Ortner, C.; Persson, K. A.; Reuter, K.; Rosen, A. S.; Schaaf, L. L.; Schran, C.; Shi, B. X.; Sivonxay, E.; Stenczel, T. K.; Svahn, V.; Sutton, C.; Swinburne, T. D.; Tilly, J.; Oord, C. v. d.; Varga-Umbrich, E.; Vegge, T.; Vondrák, M.; Wang, Y.; Witt, W. C.; Zills, F.; Csányi, G. A Foundation Model for Atomistic Materials Chemistry. *arXiv preprint arXiv:2401.00096* **2024**.
- (14) Bender, E. M.; Gebru, T.; McMillan-Major, A.; Shmitchell, S. On the Dangers of Stochastic Parrots: Can Language Models Be Too Big? 🦜. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*; FAccT '21; Association for Computing Machinery: New York, NY, USA, 2021; pp 610–623.
- (15) Corrêa, N. K.; Galvão, C.; Santos, J. W.; Pino, C. D.; Pinto, E. P.; Barbosa, C.; Massmann, D.; Mambrini, R.; Galvão, L.; Terem, E.; Oliveira, N. d. Worldwide AI Ethics: A Review of 200 Guidelines and Recommendations for AI Governance. *Patterns* **2023**, 4 (10).
- (16) Das, B. C.; Amini, M. H.; Wu, Y. Security and Privacy Challenges of Large Language Models: A Survey. *ACM Comput. Surv.* **2025**, 57 (6), 1–39.
- (17) Merton, R. K. The Normative Structure of Science. *The Sociology of Science: Theoretical and Empirical Investigations*, 1973, 139, 267–278.
- (18) *Fostering Integrity in Research*; National Academies Press: Washington, D.C., 2017. <https://doi.org/10.17226/21896>.
- (19) Mohanty, S. P. (. P. Can Our Values Be Objective? On Ethics, Aesthetics, And Progressive Politics. *New Literary History* **2001**, 32 (4), 803–833.
- (20) Johnston, M. Objective Mind and the Objectivity of Our Minds. *Philosophy and Phenomenological Research* **2007**, 75 (2), 233–268.
- (21) Dijk, A. J. M. Consciousness: A Neural Capacity for Objectivity, Especially Pronounced in Humans. *Frontiers in Psychology* **2014**, 5, 223.

- (22) Adams, J. D.; Black, G. C.; Clemmons, J. R.; Stephan, P. E. Scientific Teams and Institutional Collaborations: Evidence from U.S. Universities, 1981–1999. *Research Policy* **2005**, *34* (3), 259–285.
- (23) Larivière, V.; Gingras, Y.; Sugimoto, C. R.; Tsou, A. Team Size Matters: Collaboration and Scientific Impact since 1900. *Journal of the Association for Information Science and Technology* **2015**, *66* (7), 1323–1332.
- (24) Casadevall, A.; Fang, F. C. Specialized Science. *Infection and Immunity* **2014**, *82* (4), 1355–1360.
- (25) Ziman, J. M. *Reliable Knowledge: An Exploration of the Grounds for Belief in Science*; Cambridge University Press: New York, 1978.
- (26) Kovac, J. *The Ethical Chemist: Professionalism and Ethics in Science*; Oxford University Press: New York, 2018.
- (27) Kovac, J. Reverence and Ethics in Science. *Science and Engineering Ethics* **2013**, *19*, 745–756.
- (28) The Chemical Professional's Code of Conduct - American Chemical Society — Acs.org, 2019.
- (29) Brown, L. The Global Chemists' Code of Ethics: International Cooperation for Increased Chemical Security and Safety. *Responsible Conduct in Chemistry Research and Practice: Global Perspectives*, 2018, 129–137.
- (30) Code of Ethics — Aiche.org, 2015.
- (31) The Hague Ethical Guidelines — Opcw.org, 2015.
- (32) Varma, R.; Varma, D. R. The Bhopal Disaster of 1984. *Bulletin of Science, Technology & Society* **2005**, *25* (1), 37–45.
- (33) Palmer, M. G. The Case of Agent Orange. *Contemporary Southeast Asia* **2007**, 172–195.
- (34) Fitzgerald, G. J. Chemical Warfare and Medical Response during World War I. *American Journal of Public Health* **2008**, *98* (4), 611–625.
- (35) Abílio Ramos, M.; Droguett, E. L.; Moseh, A.; Chagas Moura, M. das; Ramos Martins, M. Revisiting Past Refinery Accidents from a Human Reliability Analysis Perspective: The BP Texas City and the Chevron Richmond Accidents. *The Canadian Journal of Chemical Engineering* **2017**, *95* (12), 2293–2305.
- (36) Sivaraman, S.; Varadharajan, S. Investigative Consequence Analysis: A Case Study Research of Beirut Explosion Accident. *Journal of Loss prevention in the Process industries* **2021**, *69*, 104387.
- (37) Taylor, T. The Chemical Weapons Convention and Prospects for Implementation. *International & Comparative Law Quarterly* **1993**, *42* (4), 912–919.
- (38) Miller, D. E. Utilitarianism and Consequentialism. *The Routledge Companion to Social and Political Philosophy*, 2012, 329–341.
- (39) Markowitz, G.; Rosner, D. *Deceit and Denial: The Deadly Politics of Industrial Pollution*; Univ of California Press: Oakland, 2013.
- (40) Driver, J. The History of Utilitarianism (Stanford Encyclopedia of Philosophy) — Seop.illc.uva.nl, 2014.
- (41) Lane, M. K. M.; Rudel, H. E.; Wilson, J. A.; Erythropel, H. C.; Backhaus, A.; Gilcher, E. B.; Ishii, M.; Jean, C. F.; Lin, F.; Muellers, T. D.; Wang, T.; Torres, G.; Taylor, D. E.; Anastas, P. T.; Zimmerman, J. B. Green Chemistry as Just Chemistry. *Nature Sustainability* **2023**, *6* (5), 502–512.
- (42) Souza Porto, M. F. de; Freitas, C. M. de. Major Chemical Accidents in Industrializing Countries: The Socio-Political Amplification of Risk. *Risk Analysis* **1996**, *16* (1), 19–29.
- (43) Cranmer, Z.; Steinfield, L.; Miranda, J.; Stohler, T. Energy Distributive Injustices: Assessing the Demographics of Communities Surrounding Renewable and Fossil Fuel Power Plants in the United States. *Energy Research & Social Science* **2023**, *100*, 103050.
- (44) Brender, J. D.; Maantay, J. A.; Chakraborty, J. Residential Proximity to Environmental Hazards and Adverse Health Outcomes. *American Journal of Public Health* **2011**, *101* (S1), S37–S52.
- (45) Kotek, H.; Dockum, R.; Sun, D. Gender Bias and Stereotypes in Large Language Models. In *Proceedings of the ACM Collective Intelligence Conference*; 2023; pp 12–24.
- (46) Wan, Y.; Chang, K.-W. White Men Lead, Black Women Help? Benchmarking Language Agency Social Biases in LLMs. *arXiv preprint arXiv:2404.10508* **2024**.
- (47) An, J.; Huang, D.; Lin, C.; Tai, M. Measuring Gender and Racial Biases in Large Language Models. *arXiv preprint arXiv:2403.15281* **2024**.
- (48) Wilson, K.; Caliskan, A. Gender, Race, And Intersectional Bias in Resume Screening via Language Model Retrieval. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* **2024**, *7* (1), 1578–1590.
- (49) Li, R.; Kamaraj, A.; Ma, J.; Ebling, S. Decoding Ableism in Large Language Models: An Intersectional Approach. In *Proceedings of the Third Workshop on NLP for Positive Impact*; 2024; pp 232–249.
- (50) Phutane, M.; Seelam, A.; Vashistha, A. How Toxicity Classifiers and Large Language Models Respond to Ableism. *arXiv preprint arXiv:2410.03448* **2024**.
- (51) Felkner, V. K.; Chang, H.-C. H.; Jang, E.; May, J. Winoqueer: A Community-in-the-Loop Benchmark for Anti-LGBTQ+ Bias in Large Language Models. *arXiv preprint arXiv:2306.15087* **2023**.
- (52) Dhingra, H.; Jayashanker, P.; Moghe, S.; Strubell, E. Queer People Are People First: Deconstructing Sexual Identity Stereotypes in Large Language Models. *arXiv preprint arXiv:2307.00101* **2023**.
- (53) Sosto, M.; Barrón-Cedeño, A. Queerbench: Quantifying Discrimination in Language Models toward Queer Identities. *arXiv preprint arXiv:2406.12399* **2024**.
- (54) Khandelwal, K.; Tonneau, M.; Bean, A. M.; Kirk, H. R.; Hale, S. A. Indian-BhED: A Dataset for Measuring India-Centric Biases in Large Language Models. In *Proceedings of the 2024 International Conference on Information Technology for Social Good*; 2024; pp 231–239.
- (55) Lee, M. H. J.; Jeon, S. Vision-Language Models Represent Darker-Skinned Black Individuals as More Homogeneous Than Lighter-Skinned Black Individuals. *arXiv preprint arXiv:2412.09668* **2024**.

- (56) Narnaware, V.; Vayani, A.; Gupta, R.; Swetha, S.; Shah, M. SB-Bench: Stereotype Bias Benchmark for Large Multimodal Models. *arXiv preprint arXiv:2502.08779* **2025**.
- (57) Yang, Y.; Liu, Y.; Liu, X.; Gulhane, A.; Mastrodicasa, D.; Wu, W.; Wang, E. J.; Sahani, D. W.; Patel, S. Demographic Bias of Expert-Level Vision-Language Foundation Models in Medical Imaging. *arXiv preprint arXiv:2402.14815 [cs]* **2024**.
- (58) Hamidieh, K.; Zhang, H.; Gerych, W.; Hartvigsen, T.; Ghassemi, M. Identifying Implicit Social Biases in Vision-Language Models. *Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society* **2024**, 7 (1), 547-561.
- (59) Rafailov, R.; Sharma, A.; Mitchell, E.; Manning, C. D.; Ermon, S.; Finn, C. Direct Preference Optimization: Your Language Model Is Secretly a Reward Model. *Advances in Neural Information Processing Systems* **2023**, 36, 53728-53741.
- (60) Allam, A. Biasppo: Mitigating Bias in Language Models Through Direct Preference Optimization. *arXiv preprint arXiv:2407.13928* **2024**.
- (61) Ganguli, D.; Lovitt, L.; Kernion, J.; Askell, A.; Bai, Y.; Kadavath, S.; Mann, B.; Perez, E.; Schiefer, N.; Ndousse, K.; Jones, A.; Bowman, S.; Chen, A.; Conerly, T.; DasSarma, N.; Drain, D.; Elhage, N.; El-Showk, S.; Fort, S.; Hatfield-Dodds, Z.; Henighan, T.; Hernandez, D.; Hume, T.; Jacobson, J.; Johnston, S.; Kravec, S.; Olsson, C.; Ringer, S.; Tran-Johnson, E.; Amodei, D.; Brown, T.; Joseph, N.; McCandlish, S.; Olah, C.; Kaplan, J.; Clark, J. Red Teaming Language Models to Reduce Harms: Methods, Scaling Behaviors, And Lessons Learned. *arXiv preprint arXiv:2209.07858* **2022**.
- (62) Turpin, M.; Michael, J.; Perez, E.; Bowman, S. Language Models Don't Always Say What They Think: Unfaithful Explanations in Chain-of-Thought Prompting. *Advances in Neural Information Processing Systems* **2023**, 36, 74952-74965.
- (63) Gandikota, R.; Orgad, H.; Belinkov, Y.; Materzyńska, J.; Bau, D. Unified Concept Editing in Diffusion Models. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*; 2024; pp 5111-5120.
- (64) Inan, H. A.; Ramadan, O.; Wutschitz, L.; Jones, D.; Rühle, V.; Withers, J.; Sim, R. Training Data Leakage Analysis in Language Models. *arXiv preprint arXiv:2101.05405* **2021**.
- (65) Huang, J.; Shao, H.; Chang, K. C.-C. Are Large Pre-Trained Language Models Leaking Your Personal Information?. *arXiv preprint arXiv:2205.12628* **2022**.
- (66) Tarkang, E. E.; Kweku, M.; Zotor, F. B. Publication Practices and Responsible Authorship: A Review Article. *Journal of Public Health in Africa* **2017**, 8 (1), 723.
- (67) Curzer, H. J. Authorship and Justice: Credit and Responsibility. *Accountability in Research* **2021**, 28 (1), 1-22.
- (68) Lukas, N.; Salem, A.; Sim, R.; Tople, S.; Wutschitz, L.; Zanella-Béguelin, S. Analyzing Leakage of Personally Identifiable Information in Language Models. In *2023 IEEE Symposium on Security and Privacy (SP)*; 2023; pp 346-363.
- (69) Das, B. C.; Amini, M. H.; Wu, Y. Security and Privacy Challenges of Large Language Models: A Survey. *ACM Computing Surveys* **2025**, 57 (6), 1-39.
- (70) Chen, X.; Tang, S.; Zhu, R.; Yan, S.; Jin, L.; Wang, Z.; Su, L.; Zhang, Z.; Wang, X.; Tang, H. The Janus Interface: How Fine-Tuning in Large Language Models Amplifies the Privacy Risks. In *Proceedings of the 2024 on ACM SIGSAC Conference on Computer and Communications Security*; 2024; pp 1285-1299.
- (71) Shehabi, A.; Hubbard, A.; Newkirk, A.; Lei, N.; Siddik, M. A. B.; Holecek, B.; Koomey, J.; Masanet, E.; Sartor, D. 2024 United States Data Center Energy Usage Report. **2024**.
- (72) Al Kez, D.; Foley, A. M.; Laverty, D.; Del Rio, D. F.; Sovacool, B. Exploring the Sustainability Challenges Facing Digitalization and Internet Data Centers. *Journal of Cleaner Production* **2022**, 371, 133633.
- (73) Guidi, G.; Dominici, F.; Gilmour, J.; Butler, K.; Bell, E.; Delaney, S.; Bargagli-Stoffi, F. J. Environmental Burden of United States Data Centers in the Artificial Intelligence Era. *arXiv preprint arXiv:2411.09786* **2024**.
- (74) Luccioni, A. S.; Viguier, S.; Ligozat, A.-L. Estimating the Carbon Footprint of Bloom, A 176b Parameter Language Model. *Journal of Machine Learning Research* **2023**, 24 (253), 1-15.
- (75) Samsi, S.; Zhao, D.; McDonald, J.; Li, B.; Michaleas, A.; Jones, M.; Bergeron, W.; Kepner, J.; Tiwari, D.; Gadepally, V. From Words to Watts: Benchmarking the Energy Costs of Large Language Model Inference. In *2023 IEEE High Performance Extreme Computing Conference (HPEC)*; 2023; pp 1-9.
- (76) Mortensen, O. How Many Users Does ChatGPT Have? Statistics & Facts — Seo.ai, 2024.
- (77) Li, P.; Yang, J.; Islam, M. A.; Ren, S. Making AI Less "Thirsty": Uncovering and Addressing the Secret Water Footprint of AI Models. *arXiv preprint arXiv:2304.03271* **2025**.
- (78) *Water for Data Centers: Market Trends and Forecasts, 2023-2030*. <https://www.bluefieldresearch.com/research/water-for-data-centers-market-trends-and-forecasts-2023-2030/> (accessed 2025-02-19).
- (79) Taylor, R.; Kardas, M.; Cucurull, G.; Scialom, T.; Hartshorn, A.; Saravia, E.; Poulton, A.; Kerkez, V.; Stojnic, R. Galactica: A Large Language Model for Science. *arXiv preprint arXiv:2211.09085* **2022**.
- (80) Lu, C.; Lu, C.; Lange, R. T.; Foerster, J.; Clune, J.; Ha, D. The Ai Scientist: Towards Fully Automated Open-Ended Scientific Discovery. *arXiv preprint arXiv:2408.06292* **2024**.
- (81) Oreskes, N. The Scientific Consensus on Climate Change. *Science* **2004**, 306 (5702), 1686.
- (82) Lynas, M.; Houlton, B. Z.; Perry, S. Greater Than 99% Consensus on Human Caused Climate Change in the Peer-Reviewed Scientific Literature. *Environmental Research Letters* **2021**, 16 (11), 114005.
- (83) Hui, X.; Reshef, O.; Zhou, L. The Short-Term Effects of Generative Artificial Intelligence on Employment: Evidence from an Online Labor Market. *Organization Science* **2024**, 35 (6), 1977-1989.
- (84) Zarifhonarvar, A. Economics of ChatGPT: A Labor Market View on the Occupational Impact of Artificial Intelligence. *Journal of Electronic Business & Digital Economics* **2023**, 3 (2), 100-116.
- (85) Guliyev, H. Artificial Intelligence and Unemployment in High-Tech Developed Countries: New Insights from Dynamic Panel Data Model. *Research in Globalization* **2023**, 7, 100140.

- (86) International, P. Humans in the AI Loop: The Data Labelers behind Some of the Most Powerful LLMs' Training Datasets – Privacyinternational.org, 2024.
- (87) Perrigo, B. Exclusive: The \$2 Per Hour Workers Who Made ChatGPT Safer – Time.com, 2023.
- (88) Westerlund, M. The Emergence of Deepfake Technology: A Review. *Technology Innovation Management Review* **2019**, 9 (11).
- (89) Rancourt-Raymond, A. de; Smaili, N. The Unethical Use of Deepfakes. *Journal of Financial Crime* **2023**, 30 (4), 1066–1077.
- (90) Eynern, C. von. *Olaf Scholz Deepfake: How a Deepfake impacts Public Trust*. <http://essay.utwente.nl/101675/>.
- (91) Kharvi, P. L. Understanding the Impact of AI-Generated Deepfakes on Public Opinion, Political Discourse, And Personal Security in Social Media. *IEEE Security & Privacy* **2024**.
- (92) Radivojevic, K.; Chou, M.; Badillo-Urquiola, K.; Brenner, P. Human Perception of Llm-Generated Text Content in Social Media Environments. *arXiv preprint arXiv:2409.06653* **2024**.
- (93) Wu, J.; Yang, S.; Zhan, R.; Yuan, Y.; Chao, L. S.; Wong, D. F. A Survey on LLM-Generated Text Detection: Necessity, Methods, And Future Directions. *Computational Linguistics* **2025**, 1–66.
- (94) Liu, H.; Xue, W.; Chen, Y.; Chen, D.; Zhao, X.; Wang, K.; Hou, L.; Li, R.; Peng, W. A Survey on Hallucination in Large Vision-Language Models. *arXiv preprint arXiv:2402.00253* **2024**.
- (95) Bai, Z.; Wang, P.; Xiao, T.; He, T.; Han, Z.; Zhang, Z.; Shou, M. Z. Hallucination of Multimodal Large Language Models: A Survey. *arXiv preprint arXiv:2404.18930* **2024**.
- (96) Huang, L.; Yu, W.; Ma, W.; Zhong, W.; Feng, Z.; Wang, H.; Chen, Q.; Peng, W.; Feng, X.; Qin, B.; Liu, T. A Survey on Hallucination in Large Language Models: Principles, Taxonomy, Challenges, And Open Questions. *ACM Transactions on Information Systems* **2025**, 43 (2), 1–55.
- (97) Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; Polosukhin, I. Attention Is All You Need. *Advances in neural information processing systems* **2017**, 30.
- (98) OpenAI; Achiam, J.; Adler, S.; Agarwal, S.; Ahmad, L.; Akkaya, I.; Aleman, F. L.; Almeida, D.; Altenschmidt, J.; Altman, S.; Anadkat, S.; Avila, R.; Babuschkin, I.; Balaji, S.; Balcom, V.; Baltescu, P.; Bao, H.; Bavarian, M.; Belgum, J.; Bello, I.; Berdine, J.; Bernadett-Shapiro, G.; Berner, C.; Bogdonoff, L.; Boiko, O.; Boyd, M.; Brakman, A.-L.; Brockman, G.; Brooks, T.; Brundage, M.; Button, K.; Cai, T.; Campbell, R.; Cann, A.; Carey, B.; Carlson, C.; Carmichael, R.; Chan, B.; Chang, C.; Chantzis, F.; Chen, D.; Chen, S.; Chen, R.; Chen, J.; Chen, M.; Chess, B.; Cho, C.; Chu, C.; Chung, H. W.; Cummings, D.; Currier, J.; Dai, Y.; Decareaux, C.; Degry, T.; Deutsch, N.; Deville, D.; Dhar, A.; Dohan, D.; Dowling, S.; Dunning, S.; Ecoffet, A.; Eleti, A.; Eloundou, T.; Farhi, D.; Fedus, L.; Felix, N.; Fishman, S. P.; Forte, J.; Fulford, I.; Gao, L.; Georges, E.; Gibson, C.; Goel, V.; Gogineni, T.; Goh, G.; Gontijo-Lopes, R.; Gordon, J.; Grafstein, M.; Gray, S.; Greene, R.; Gross, J.; Gu, S. S.; Guo, Y.; Hallacy, C.; Han, J.; Harris, J.; He, Y.; Heaton, M.; Heidecke, J.; Hesse, C.; Hickey, A.; Hickey, W.; Hoeschele, P.; Houghton, B.; Hsu, K.; Hu, S.; Hu, X.; Huizinga, J.; Jain, S.; Jain, S.; Jang, J.; Jiang, A.; Jiang, R.; Jin, H.; Jin, D.; Jomoto, S.; Jonn, B.; Jun, H.; Kaftan, T.; Kaiser, Ł.; Kamali, A.; Kanitscheider, I.; Keskar, N. S.; Khan, T.; Kilpatrick, L.; Kim, J. W.; Kim, C.; Kim, Y.; Kirchner, J. H.; Kiros, J.; Knight, M.; Kokotajlo, D.; Kondraciuk, Ł.; Kondrich, A.; Konstantinidis, A.; Kosic, K.; Krueger, G.; Kuo, V.; Lampe, M.; Lan, I.; Lee, T.; Leike, J.; Leung, J.; Levy, D.; Li, C. M.; Lim, R.; Lin, M.; Lin, S.; Litwin, M.; Lopez, T.; Lowe, R.; Lue, P.; Makanju, A.; Malfacini, K.; Manning, S.; Markov, T.; Markovski, Y.; Martin, B.; Mayer, K.; Mayne, A.; McGrew, B.; McKinney, S. M.; McLeavey, C.; McMillan, P.; McNeil, J.; Medina, D.; Mehta, A.; Menick, J.; Metz, L.; Mishchenko, A.; Mishkin, P.; Monaco, V.; Morikawa, E.; Mossing, D.; Mu, T.; Murati, M.; Murk, O.; Mély, D.; Nair, A.; Nakano, R.; Nayak, R.; Neelakantan, A.; Ngo, R.; Noh, H.; Ouyang, L.; O'Keefe, C.; Pachocki, J.; Paino, A.; Palermo, J.; Pantuliano, A.; Parascandolo, G.; Parish, J.; Parparita, E.; Passos, A.; Pavlov, M.; Peng, A.; Perelman, A.; Avila Belbute Peres, F. de; Petrov, M.; Oliveira Pinto, H. P. de; Michael; Pokorny; Pokrass, M.; Pong, V. H.; Powell, T.; Power, A.; Power, B.; Proehl, E.; Puri, R.; Radford, A.; Rae, J.; Ramesh, A.; Raymond, C.; Real, F.; Rimbach, K.; Ross, C.; Rotsted, B.; Roussez, H.; Ryder, N.; Saltarelli, M.; Sanders, T.; Santurkar, S.; Sastry, G.; Schmidt, H.; Schnurr, D.; Schulman, J.; Selsam, D.; Sheppard, K.; Sherbakov, T.; Shieh, J.; Shoker, S.; Shyam, P.; Sidor, S.; Sigler, E.; Simens, M.; Sitkin, J.; Slama, K.; Sohl, I.; Sokolowsky, B.; Song, Y.; Staudacher, N.; Such, F. P.; Summers, N.; Sutskever, I.; Tang, J.; Tezak, N.; Thompson, M. B.; Tillet, P.; Tootoonchian, A.; Tseng, E.; Tuggle, P.; Turley, N.; Tworek, J.; Uribe, J. F. C.; Vallone, A.; Vijayvergiya, A.; Voss, C.; Wainwright, C.; Wang, J. J.; Wang, A.; Wang, B.; Ward, J.; Wei, J.; Weinmann, C.; Welihinda, A.; Welinder, P.; Weng, J.; Weng, L.; Wiethoff, M.; Willner, D.; Winter, C.; Wolrich, S.; Wong, H.; Workman, L.; Wu, S.; Wu, J.; Wu, M.; Xiao, K.; Xu, T.; Yoo, S.; Yu, K.; Yuan, Q.; Zaremba, W.; Zellers, R.; Zhang, C.; Zhang, M.; Zhao, S.; Zheng, T.; Zhuang, J.; Zhuk, W.; Zoph, B. GPT-4 Technical Report. *arXiv preprint arXiv:2303.08774* **2023**.
- (99) Frankfurt, H. G. *On Bullshit*; Princeton University Press, 2009.
- (100) Hicks, M. T.; Humphries, J.; Slater, J. Chatgpt Is Bullshit. *Ethics and Information Technology* **2024**, 26 (2), 1–10.
- (101) Tonmoy, S.; Zaman, S.; Jain, V.; Rani, A.; Rawte, V.; Chadha, A.; Das, A. A Comprehensive Survey of Hallucination Mitigation Techniques in Large Language Models. *arXiv preprint arXiv:2401.01313* **2024**, 6.
- (102) Luo, J.; Li, T.; Wu, D.; Jenkin, M.; Liu, S.; Dudek, G. Hallucination Detection and Hallucination Mitigation: An Investigation. *arXiv preprint arXiv:2401.08358* **2024**.
- (103) Chan, K.; Swenson, A. One Tech Tip: How to Spot AI-Generated Deepfake Images – Apnews.com, 2024.
- (104) Kamali, N.; Nakamura, K.; Chatzimpampas, A.; Hullman, J.; Groh, M. How to Distinguish Ai-Generated Images from Authentic Photographs. *arXiv preprint arXiv:2406.08651* **2024**.
- (105) Van Dalen, H. P.; Henkens, K. Intended and Unintended Consequences of a Publish-or-Perish Culture: A Worldwide Survey. *Journal of the American Society for Information Science and Technology* **2012**, 63 (7), 1282–1293.
- (106) Rawat, S.; Meena, S. Publish or Perish: Where Are We Heading?. *Journal of Research in Medical Sciences: the Official Journal of Isfahan University of Medical Sciences* **2014**, 19 (2), 87.
- (107) Xue, Y.; Larson, R. C. STEM Crisis or STEM Surplus? Yes and Yes. *Monthly Labor Review* **2015**, 2015, 10–21916.

- (108) Lee, H.-P. H.; Sarkar, A.; Tankelevitch, L.; Drosos, I.; Rintel, S.; Banks, R.; Wilson, N. The Impact of Generative AI on Critical Thinking: Self-Reported Reductions in Cognitive Effort and Confidence Effects From a Survey of Knowledge Workers. In *Proceedings of the ACM CHI Conference on Human Factors in Computing Systems*; Association for Computing Machinery, 2025.
- (109) Society, A. C. Artificial Intelligence (AI) Best Practices and Policies at ACS Publications — Researcher-Resources.acs.org, 2024.
- (110) Portfolio, N. Artificial Intelligence (AI) | Nature Portfolio — Nature.com, 2025.
- (111) Waseem, Z.; Lulz, S.; Bingel, J.; Augenstein, I. Disembodied Machine Learning: On the Illusion of Objectivity in NLP. *arXiv preprint arXiv:2101.11974* **2021**.
- (112) Zheng, Z.; Rampal, N.; Inizan, T. J.; Borgs, C.; Chayes, J. T.; Yaghi, O. M. Large Language Models for Reticular Chemistry. *Nature Reviews Materials* **2025**, 1–13.
- (113) Zhang, Q.; Ding, K.; Lv, T.; Wang, X.; Yin, Q.; Zhang, Y.; Yu, J.; Wang, Y.; Li, X.; Xiang, Z.; Zhuang, X.; Wang, Z.; Qin, M.; Zhang, M.; Zhang, J.; Cui, J.; Xu, R.; Chen, H.; Fan, X.; Xing, H.; Chen, H. Scientific Large Language Models: A Survey on Biological & Chemical Domains. *ACM Comput. Surv.* **2025**, 57 (6), 1–38.
- (114) Caldas Ramos, M.; J. Collison, C.; D. White, A. A Review of Large Language Models and Autonomous Agents in Chemistry. *Chemical Science* **2025**, 16 (6), 2514–2572.
- (115) Xiao, Y.; Zhao, W.; Zhang, J.; Jin, Y.; Zhang, H.; Ren, Z.; Sun, R.; Wang, H.; Wan, G.; Lu, P.; Luo, X.; Zhang, Y.; Zou, J.; Sun, Y.; Wang, W. Protein Large Language Models: A Comprehensive Survey. *arXiv preprint arXiv:2502.17504* **2025**.
- (116) Wang, L.; Li, X.; Zhang, H.; Wang, J.; Jiang, D.; Xue, Z.; Wang, Y. A Comprehensive Review of Protein Language Models. *arXiv preprint arXiv:2502.06881* **2025**.
- (117) Guo, T.; Guo, K.; Nan, B.; Liang, Z.; Guo, Z.; Chawla, N.; Wiest, O.; Zhang, X. What Can Large Language Models Do in Chemistry? A Comprehensive Benchmark on Eight Tasks. *Advances in Neural Information Processing Systems* **2023**, 36, 59662–59688.
- (118) Huang, X.; Surve, M.; Liu, Y.; Luo, T.; Wiest, O.; Zhang, X.; Chawla, N. V. Application of Large Language Models in Chemistry Reaction Data Extraction and Cleaning. In *Proceedings of the 33rd ACM International Conference on Information and Knowledge Management*; CIKM '24; Association for Computing Machinery: New York, NY, USA, 2024; pp 3797–3801.
- (119) Omiye, J. A.; Lester, J. C.; Spichak, S.; Rotemberg, V.; Daneshjou, R. Large Language Models Propagate Race-Based Medicine. *NPJ Digital Medicine* **2023**, 6 (1), 195.
- (120) Liang, W.; Zhang, Y.; Wu, Z.; Lepp, H.; Ji, W.; Zhao, X.; Cao, H.; Liu, S.; He, S.; Huang, Z.; Yang, D.; Potts, C.; Manning, C. D.; Zou, J. Y. Mapping the Increasing Use of LLMs in Scientific Papers. *arXiv preprint arXiv:2404.01268* **2024**.
- (121) Yu, S.; Luo, M.; Madasu, A.; Lal, V.; Howard, P. Is Your Paper Being Reviewed by an LLM? Investigating Ai Text Detectability in Peer Review. *arXiv preprint arXiv:2410.03019* **2024**.
- (122) Watch, R. Papers and Peer Reviews with Evidence of ChatGPT Writing — Retractionwatch.com, 2025.
- (123) Seifrid, M.; Pollice, R.; Aguilar-Granda, A.; Morgan Chan, Z.; Hotta, K.; Ser, C. T.; Vestfrid, J.; Wu, T. C.; Aspuru-Guzik, A. Autonomous Chemical Experiments: Challenges and Perspectives on Establishing a Self-Driving Lab. *Accounts of Chemical Research* **2022**, 55 (17), 2454–2466.
- (124) Tom, G.; Schmid, S. P.; Baird, S. G.; Cao, Y.; Darvish, K.; Hao, H.; Lo, S.; Pablo-García, S.; Rajaonson, E. M.; Skreta, M.; Yoshikawa, N.; Corapi, S.; Akkok, G. D.; Strieth-Kalthoff, F.; Seifrid, M.; Aspuru-Guzik, A. Self-Driving Laboratories for Chemistry and Materials Science. *Chemical Reviews* **2024**, 124 (16), 9633–9732.
- (125) M. Bran, A.; Cox, S.; Schilter, O.; Baldassari, C.; White, A. D.; Schwaller, P. Augmenting Large Language Models with Chemistry Tools. *Nature Machine Intelligence* **2024**, 6 (5), 525–535.
- (126) Jablonka, K. M.; Schwaller, P.; Ortega-Guerrero, A.; Smit, B. Leveraging Large Language Models for Predictive Chemistry. *Nature Machine Intelligence* **2024**, 6 (2), 161–169.
- (127) Caccavale, F.; Gargalo, C. L.; Germaey, K. V.; Krühne, U. Towards Education 4.0: The Role of Large Language Models as Virtual Tutors in Chemical Engineering. *Education for Chemical Engineers* **2024**, 49, 1–11.
- (128) Du, Y.; Duan, C.; Bran, A.; Sotnikova, A.; Qu, Y.; Kulik, H.; Bosselut, A.; Xu, J.; Schwaller, P. Large Language Models Are Catalyzing Chemistry Education. *ChemRxiv preprint* **2024**.
- (129) Abriata, L. A. The Nobel Prize in Chemistry: Past, Present, And Future of AI in Biology. *Communications Biology* **2024**, 7 (1), 1409.
- (130) Nijkamp, E.; Ruffolo, J. A.; Weinstein, E. N.; Naik, N.; Madani, A. ProGen2: Exploring the Boundaries of Protein Language Models. *Cell Systems* **2023**, 14 (11), 968–978.
- (131) Abramson, J.; Adler, J.; Dunger, J.; Evans, R.; Green, T.; Pritzel, A.; Ronneberger, O.; Willmore, L.; Ballard, A. J.; Bambrick, J.; Bodenstein, S. W.; Evans, D. A.; Hung, C.-C.; O'Neill, M.; Reiman, D.; Tunyasuvunakool, K.; Wu, Z.; Žemgulytė, A.; Arvaniti, E.; Beattie, C.; Bertolli, O.; Bridgland, A.; Cherepanov, A.; Congreve, M.; Cowen-Rivers, A. I.; Cowie, A.; Figurnov, M.; Fuchs, F. B.; Gladman, H.; Jain, R.; Khan, Y. A.; Low, C. M. R.; Perlin, K.; Potapenko, A.; Savy, P.; Singh, S.; Stecula, A.; Thillaisundaram, A.; Tong, C.; Yakneen, S.; Zhong, E. D.; Zielinski, M.; Židek, A.; Bapst, V.; Kohli, P.; Jaderberg, M.; Hassabis, D.; Jumper, J. M. Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3. *Nature* **2024**, 630 (8016), 493–500.
- (132) Tai, K. S.; Bailis, P.; Valiant, G. Equivariant Transformer Networks. In *International Conference on Machine Learning*; 2019; pp 6086–6095.
- (133) Cao, H.; Tan, C.; Gao, Z.; Xu, Y.; Chen, G.; Heng, P.-A.; Li, S. Z. A Survey on Generative Diffusion Models. *IEEE Transactions on Knowledge and Data Engineering* **2024**.
- (134) Shoghi, N.; Kolluru, A.; Kitchin, J. R.; Ulissi, Z. W.; Zitnick, C. L.; Wood, B. M. From Molecules to Materials: Pre-Training Large Generalizable Models for Atomic Property Prediction. *arXiv preprint arXiv:2310.16802* **2023**.
- (135) Jain, A.; Ong, S. P.; Hautier, G.; Chen, W.; Richards, W. D.; Dacek, S.; Cholia, S.; Gunter, D.; Skinner, D.; Ceder, G.; Persson,

- K. A. Commentary: The Materials Project: A Materials Genome Approach to Accelerating Materials Innovation. *APL materials* **2013**, 1 (1).
- (136) Cai, F.; Hanna, K.; Zhu, T.; Tzeng, T.-R.; Duan, Y.; Liu, L.; Pilla, S.; Li, G.; Luo, F. A Foundation Model for Chemical Design and Property Prediction. *arXiv preprint arXiv:2410.21422* **2025**.
- (137) Wang, Z.; Wu, Z. Towards Foundation Model for Chemical Reactor Modeling: Meta-Learning with Physics-Informed Adaptation. *arXiv preprint arXiv:2405.11752* **2024**.
- (138) Desislavov, R.; Martínez-Plumed, F.; Hernández-Orallo, J. Trends in AI Inference Energy Consumption: Beyond the Performance-Vs-Parameter Laws of Deep Learning. *Sustainable Computing: Informatics and Systems* **2023**, 38, 100857.
- (139) Luccioni, S.; Jernite, Y.; Strubell, E. Power Hungry Processing: Watts Driving the Cost of AI Deployment?. In *Proceedings of the 2024 ACM conference on fairness, accountability, and transparency*; 2024; pp 85–99.
- (140) Sherman, J.; Chin, B.; Huibers, P.; Garcia-Valls, R.; Hatton, T. A. Solvent Replacement for Green Processing. *Environmental health perspectives* **1998**, 106 (suppl1), 253–271.
- (141) Vethaak, A. D.; Legler, J. Microplastics and Human Health. *Science* **2021**, 371 (6530), 672–674.
- (142) Li, W. C.; Tse, H. F.; Fok, L. Plastic Waste in the Marine Environment: A Review of Sources, Occurrence and Effects. *Science of the Total Environment* **2016**, 566, 333–349.
- (143) Mann, E. Digital Technology Is Dependent on Forced Labor: The Exploitative Labor Practices of Cobalt Extraction in the Democratic Republic of Congo. *Applied Anthropologist* **2017**, 37 (1).
- (144) Calvão, F.; McDonald, C. E. A.; Bolay, M. Cobalt Mining and the Corporate Outsourcing of Responsibility in the Democratic Republic of Congo. *The Extractive Industries and Society* **2021**, 8 (4), 100884.
- (145) Jessop, P. G. Searching for Green Solvents. *Green Chemistry* **2011**, 13 (6), 1391–1398.
- (146) Cvjetko Bubalo, M.; Vidović, S.; Radojčić Redovniković, I.; Jokić, S. Green Solvents for Green Technologies. *Journal of Chemical Technology & Biotechnology* **2015**, 90 (9), 1631–1639.
- (147) Croy, J. R.; Long, B. R.; Balasubramanian, M. A Path toward Cobalt-Free Lithium-Ion Cathodes. *Journal of Power Sources* **2019**, 440, 227113.
- (148) Wang, M.; Chen, X.; Yao, H.; Lin, G.; Lee, J.; Chen, Y.; Chen, Q. Research Progress in Lithium-Excess Disordered Rock-Salt Oxides Cathode. *Energy & Environmental Materials* **2022**, 5 (4), 1139–1154.
- (149) Chen, X.; Chen, F.; Jiang, H.; Wang, J.; Li, Y. X.; Wang, G. Replacing Plastic with Bamboo: Eco-Friendly Disposable Tableware Based on the Separation of Bamboo Fibers and the Reconstruction of Their Network Structure. *ACS Sustainable Chemistry & Engineering* **2023**, 11 (19), 7407–7418.
- (150) Brodin, M.; Vallejos, M.; Opedal, M. T.; Area, M. C.; Chinga-Carrasco, G. Lignocellulosics as Sustainable Resources for Production of Bioplastics—a Review. *Journal of Cleaner Production* **2017**, 162, 646–664.
- (151) Rujnić-Sokele, M.; Pilipović, A. Challenges and Opportunities of Biodegradable Plastics: A Mini Review. *Waste Management & Research* **2017**, 35 (2), 132–140.
- (152) Coates, G. W.; Getzler, Y. D. Chemical Recycling to Monomer for an Ideal, Circular Polymer Economy. *Nature Reviews Materials* **2020**, 5 (7), 501–516.
- (153) Demarteau, J.; Cousineau, B.; Wang, Z.; Bose, B.; Cheong, S.; Lan, G.; Baral, N. R.; Teat, S. J.; Scown, C. D.; Keastling, J. D.; Helms, B. A. Biorenewable and Circular Polydiketoenamine Plastics. *Nature Sustainability* **2023**, 6 (11), 1426–1435.
- (154) Harrer, S. Attention Is Not All You Need: The Complicated Case of Ethically Using Large Language Models in Healthcare and Medicine. *eBioMedicine* **2023**, 90.
- (155) Li, H.; Moon, J. T.; Purkayastha, S.; Celi, L. A.; Trivedi, H.; Gichoya, J. W. Ethics of Large Language Models in Medicine and Medical Research. *The Lancet Digital Health* **2023**, 5 (6), e333–e335.
- (156) Pool, J.; Indulska, M.; Sadiq, S. Large Language Models and Generative AI in Telehealth: A Responsible Use Lens. *Journal of the American Medical Informatics Association* **2024**, 31 (9), 2125–2136.
- (157) Ong, J. C. L.; Chang, S. Y.-H.; William, W.; Butte, A. J.; Shah, N. H.; Chew, L. S. T.; Liu, N.; Doshi-Velez, F.; Lu, W.; Savulescu, J.; Ting, D. S. W. Ethical and Regulatory Challenges of Large Language Models in Medicine. *The Lancet Digital Health* **2024**, 6 (6), e428–e432.