AADPRT AI in Psychiatric Education Task Force Final Report

Lead Authors:

Liz Gass, University of Washington Boise
Mary Burris, University of Utah

Contributing Authors:

Dale Peeples, Augusta University
Zheala Qayyum, Boston Children's Hospital
Tolulope Odebunmi, University of Minnesota Twin Cities

Co-Chairs:

David Ross, University of Alberta
Ronke Babalola, Hackensack University Medical Center

Dear Colleagues –

We are delighted to share with you the results of the AADPRT AI in Psychiatric Education Task Force.

As you may recall, our charge was to prepare a report for EC/SC that accomplishes the following:

- Defines core terminology relating to modern methods in artificial intelligence;
- Describes the potential applications for AI in medical education;
- Describes potential risks of AI in medical education;
- Offers recommendations for how programs can effectively navigate this space.

We've had a fantastic team working on this project and we've tried to focus on what would be most relevant and practical for a Program Director in the here and now – *what should I actually do, today?* Sadly, this has proven to be a difficult task. The field is advancing very quickly – even in the 6 months that we've been working on this, there have been major changes in the tools available, our understanding of the ethics / moral implications, acceptable standards for use, etc. There are no clear and easy answers.

The main findings of our group are summarized in Section 1. We list potential applications of AI for psychiatry PDs for teaching and evaluation, administration, and other clinical and academic uses. For each potential application, rather than providing an "answer" (that, inevitably, would change) we've tried to highlight the relative risks and benefits and to share our best sense of where things currently stand. We also include links to some of our favorite resources—of course, these are myriad and constantly evolving.

Key caveat #1: There are many efforts underway to use AI in clinical settings – because these would be conducted under the authority of hospitals and health care systems (rather than residency programs), we considered these applications to be outside the scope of the task force's work. We reference them only to the extent that they affect the way we think about training residents.

Section 2 offers a more in-depth discussion of some of the potential risks of AI and mitigation strategies that can be taken at both the PD level and, we hope, at a system / structural level. (These expand on the caveats from Section 1.)

Section 3 offers a glossary of key terminology in AI.

Key caveat #2: One of the biggest issues we wrestled with was data privacy and security. Though we've all become accustomed to reading about high profile uses for AI, including in clinical settings, all such efforts are predicated on being able to maintain data security – this is typically done by running AI and storing data on servers that sit behind a firewall. When thinking about AI in educational settings, we do not expect that most programs will have access to this level of data security. We have constructed our recommendations based on the assumption that programs would be trying to use commercially available products.

One of the most important things we realized through our work was how important it will be for residents and faculty alike to stay apprised of the field as it continues to advance. We believe there would be enormous value in having an introductory course on AI Literacy. One of our team members, Liz Gass, is currently designing such a course for her program. We are hopeful that she will be able to adapt this and offer it as a model curriculum for other AADPRT members.

Please feel free to reach out to any of us with additional questions – while I doubt that we'll have the "answer," I am confident that we will have struggled with the same question and can at least share what we've learned along the way.

Many thanks,

Team AI

p.s. This report was written entirely by people. No generative AI was used.

Note: This report reflects the opinions of the individual authors and not their institutions.

Section 1

Potential Applications for AI in Psychiatric Education

TEACHING AND EVALUATION

Potential Application/ Use	Explanation of Potential Benefit	Limitations/ Caveats	Our recommendation (today)
Simulated patient encounters (1)	LLMs can be used to create simulated patient encounters to facilitate clinical training in diagnostic evaluation, patient interviewing, and other clinical skills. This could be most valuable for situations where simulations would not be possible (e.g. at an institution where it's prohibited to have child actors involved).	Limited visuals currently, limited tokens for encounters, possible inconsistency between students; privacy concerns for what data goes in to train the simulation;	If the tool was vetted by leadership and used with appropriate supervision, there might be benefit to using simulated patients. There are chronic / omnipresent limitations of AI, including embedded biases, carbon footprint, intellectual property concerns, etc. Of note, one would need to also consider the security of student responses and the creation of new simulated patients could be problematic due to privacy concerns (see risk #HIPAA) (ref 1)
Visual/graphic design (2, 3)	There are a number of programs that use AI to create images (e.g. Midjourney) or to design slides or other marketing materials (e.g. Canva). Many conventional programs have also embedded AI tools directly into their platforms (e.g. Microsoft Powerpoint, Adobe Illustrator).	Some of these may be expensive, therefore limiting accessibility; intellectual property may not be appropriately attributed/compensated. Norms and expectations around appropriate acknowledgement and attribution remain unclear.	These may be useful. The main concerns relate to violation of intellectual property. There are no other obvious acute harms. There are chronic / omnipresent limitations of AI, including embedded biases, carbon footprint, intellectual property concerns, etc. Any use should be appropriately attributed.
Use of AI-generated board style questions (4)	LLMs can be used to create custom, interactive, and tailored board-style review tools, including MCQ's.	Copyright issues (with the material that is fed into the model), possibility of hallucination/confabulation (which can be very hard to detect in MCQs), legal/liability questions for educators, particularly program directors, if they recommend these tools.	We do not recommend use of AI generated questions (IP and accuracy); if you were trying to do this, you would need to very carefully vet any product.
Creation of learning objectives (5)	Machine learning tools can take an educational presentation as input and produce summary objectives using the language of Bloom's taxonomy.	This is fundamentally antithetical to good pedagogy -- we should be starting from SMART objectives and working backwards to design educational experiences.	Some individuals find this useful in some ways; others are horrified at the idea. Choose your own adventure!
AI tools to detect plagarism (6)	As more trainees use LLMs to generate text for educational assignments and applications, detection of plagiarism is becoming both more relevant and more challenging.	Imperfect (but might give a false sense of security!), possibly discriminatory (in that more expensive AI generators are less likely to get caught); potential privacy violations if you're submitting materials to an outside entity. There are also now a range of AI programs that are specifically designed to mask the use of AI so that it can't be detected (including: uPassAI, Humanizer AI, Undetectable AI, WriteHuman, AIHumanizer, and Rewrite AI)	Given the ongoing arms race, it would be dangerous to have confidence in any AI detection software. For assessing residency applications, this can't effectively be done at a program level. we recommend that if you're concerned about the integrity of an application that you pursue other paths to assess that applicant. more importantly, we recommend advocacy to ERAS to implement plagiarism checking centrally. (see risk #10 - application materials); For other materials that residents generate (e.g. any other written assignments), it's still problematic and we do not recommend. Find other ways to assess!
Assist with assessment of trainees' clinical performance (7)	A variety of ML approaches (ANNs, SVMs, elastic nets, hierarchical clustering, hidden Markov models, NLP, and others) have been used to automate the assessment of clinical competencies across all milestone domains.	These approaches are not yet ready for mainstream use due to inconsistency in results; beware data security!	These may some day become powerful tools. At this time, they are not recommended.
Curriculum development / design (8)	Large language models can dynamically respond to queries regarding core curriculum content, specific teaching tools, and teaching resources.	There are relatively few articles on the risks of using AI in the process of curriculum development.	If you're starting from scratch, this could be useful for idea generation. Any output would be need to be carefully vetted (and with all of the chronic caveats).
Automated didactics (9)	Custom LLMs could provide an interactive learning experience on specific topics, reducing the need for ongoing faculty involvement after an initial investment of time.	Risk of hallucinations/confabulation; possibility of decreased faculty engagement/investment in learning if teaching is outsourced to software.	In theory, this could be useful (especially if the LLM was offering feedback in a highly constrained way). In practice, this does not seem ready for use at this time.
Benchmarking psychotherapy skills (10)	Machine learning tools are being used to transcribe therapy sessions and evaluate them for fidelity to a given modality, such as CBT.	Privacy concerns; difficulty capturing data that reflect therapeutic rapport.	In theory, this could be a powerful tool. In practice, privacy concerns for patient data render this unusable by programs at this time (it will be interesting to see if/when hospital systems create secure versions of such programs). Not recommended for residency programs to implement today.

ADMINISTRATIVE			
Potential Application/ Use	**Explanation of Potential Benefit**	**Limitations/ Caveats**	**Our recommendation (today)**
Administrative documentation, including meeting transcription (11)	Automated transcription is now routinely available (including in multiple meeting platforms).	Data security, especially for meetings; potential inaccuracies and inequities (if, for example, the voices of some participants are recorded more accurately than others).	Be careful! With most commercial products, anything you say goes into the cloud. If everyone consents, have fun. You're also fine if you have software that will run this service locally.
Standardization of administrative tasks (12)	In theory, AI tools could help reduce administrative burden (e.g. by reviewing documentation that was submitted in order to audit compliance).	We have not yet seen successful implementation of AI workflows that substantially remove administrative burden - despite having seen personnel shift as if these tools were already here. Such tools could also propagate inequities for low-resource programs.	We are unaware of any tools that are currently useful in this space. Moreover, we have already seen direct harm from improper implementation. If one were to try a new tool, one would also have to be cautious about data privacy. We do not recommend at this time.
Applicant screening for resident recruitment (13, 14)	There has been successful utilization of AI in job recruitment and screening for both job applicants and potential employers. It has helped applicants screen through opportunities and removed human bias for recruiters with the right information fed in.	AI algorithms can be inherently flawed, leading to the introduction of racism, misogyny, and other forms of bias in the recruitment process. There are cost implications and there might be legal privacy issues regarding the data used to develop the AI.	Major concerns re: privacy and biases with any commercially available product. We do not recommend doing this.
Community/ stakeholder engagement (15)	Program leadership may lack expertise in data visualization; a growing number of AI-based tools facilitate interactive, high-impact presentation of data such as budgets, accomplishments of residents, etc.	These tools require funding and user training, so their utilization is likely to be unequally distributed.	May be useful with appropriate human supervision.
Monitoring resident well-being (16)	Health tracking tools such as sleep monitoring, HR variability, exercise time, etc could be actively tracked, with resident consent.	Significant privacy concerns; possible inequities in access to relevant tech (if using wearables, for example); potential legal implications if impairment detected. (And it's super creepy.)	It may be that this will eventually become connected to regulatory procedures, but Psychiatry PD's should not be engaging with this now.
Program evaluation/ data analysis (17)	Enables data-driven decision-making: program data can be aggregated and used as input to models that can produce interpretable results to support program operations including recruitment, funding, academic performance, community impact, etc.	Limited by the quality of the data used; requires computational resources and user expertise to leverage; residents may feel coerced to participate.	If you think you have a tool that is both safe to use and effective, have fun.

CLINICAL AND ACADEMIC USES

Potential Application/ Use	Explanation of Potential Benefit	Limitations/ Caveats	Our recommendation (today)
Assistance with clinical documentation (18, 19)	Tools that automate medical documentation could save substantial time for providers.	Need to establish boundary between assistance with documentation and actual clinical reasoning. If it were pure documentation assistance, this would comparable to a scribe service (albeit with more complicated privacy / security concerns). To the extent that it involves clinical reasoning, this becomes much more problematic: it may discourage residents from engaging in the formulation process, thus undermining their growth and development; it would also limit the ability of supervisors to evaluate and provide meaningful feedback on an array of clinical skills, including data gathering, evidence synthesis, formulation, and case presentation.	Residency programs cannot meet standards for data security -- any software would have to be implemented by hospital systems. If clinical sites are using proprietary software, PD's should be mindful to ensure that such tools are not interfering with resident learning.
Writing letters of recommendation, evaluations, other writing projects	Would save significant time for busy academics (who are being asked to perform a time-consuming task without compensation), whether being used to write content de novo or to help edit / polish text that you've already written (this could be especially helpful for individuals for whom English is not their first language).	Concerns include: intellectual honesty; lack of individualization; potential for confabulation; and privacy (anything you feed into the model is entered into the cloud).	If you choose to do this, do so with caution! (consider using your own letters as templates instead)
Hypothesis generation / knowledge discovery (20)	Machine learning approaches allow us to identify patterns in complex, high-dimensional data. These patterns would not be detectable by human observation and can lead to hypotheses that can then be tested with other research approaches.	These approaches require substantial computing power and are limited by the quality of the data. Considerable expertise is needed to use these tools appropriately and understand their outputs.	For a savvy user (and/or with expert supervision), this can have value. For others, this may be dangerous / problematic.
Clinical decision support (21, 22, 23) NO RECOMMENDATION	Decision-support tools based on machine learning algorithms are already embedded in many EMRs, but few have been validated in psychiatry. Tools like REACH VET have been used succesfully to identify patients at risk for suicide. This is a VERY active area of research.	Most tools have not demonstrated substantial utility in clinical practice in psychiatry. Trainees may alter their clinical practice in response to information produced by AI models that are not interpretable or explainable, but which appear impressive.	It is not within the purview of program directors to adjudicate clinical uses of AI as implemented by hospitals or health care systems. We would note that any such uses face major issues relating to data security, clinical ethics, liability, potential biases, etc. We recommend that programs be mindful of the impact that such tools would have on trainees' growth and development. If trainees are exposed to such uses through their clinical work, we recommend ensuring that they are well-trained to navigate the myriad risks.

Section 2

Risks and Mitigation Strategies for AI in Psychiatric Education

RISKS AND MITIGATION STRATEGIES FOR AI IN PSYCHIATRIC EDUCATION			
Risk	**Elaboration**	**PD Mitigation Strategies**	**Broader Mitigation Strategies**
AI undermines individuals' ability to think critically (24)	The development of clinical reasoning is a key element of medical education. Overreliance on AI-based tools such as LLMs for data gathering, evidence review, information synthesis, and formulation threatens to undermine trainee's ability to develop and practice their clinical reasoning skills. Over time, this could have determinental effects not only on individual educational attainment, but on the practice of medicine as a whole. If our current learners begin to rely on software products rather than their own critical thinking skills in clinical decision-making, they will not be able to teach these skills to subsequent generations. Ultimately, this could undermine the integrity of our profession and substantially degrade the efficacy of psychiatric treatment.	Continue to emphasize the importance of clinical reasoning in medical education. Evaluate software tools based on whether they enhance or simply attempt to automate clinical problem-solving and decision making. Shift assessment tools to focus on competency rather than knowledge. And/or... consider banning generative AI in note-writing.	
AI is racist and sexist (25, 26)	AI systems can exhibit both racist and sexist biases, stemming from flaws in their training data, algorithmic design, and development processes, reflecting and amplify existing societal prejudices and historical biases. For example, facial recognition software may struggle with darker skin tones or female features, while language models might associate leadership roles predominantly with male attributes; AI-driven hiring tools could discriminate against women and minorities, perpetuating workplace inequalities; in healthcare, AI diagnostics might underperform for women of color, potentially leading to misdiagnoses; credit scoring algorithms may unfairly penalize both female applicants and those from minority communities. These biases can affect education, employment, housing, criminal justice, and financial services.	At the PD level, the most important thing is to be aware of the potential for bias and to audit to ensure that it's not affecting your work; ideally, you can check documentation to see whether the tool you're using has employed strategies described in the next column.	The approach to address the racist and sexist tendencies in AI systems requires diverse datasets, rigorous bias testing and correction mechanisms, inclusive development practices, and transparent, auditable AI systems to ensure fairness and prevent the perpetuation of societal inequalities. Regulatory frameworks should be established to govern AI development and deployment.
Access to AI is unequally distributed (27, 28)	Several accessibility challenges exist, including unequal access to AI technology among students. While some can afford advanced versions, others may only have access to free options. Additionally, not all users possess the skills to effectively utilize AI, and the technology may struggle with various accents and languages, further hindering accessibility. These factors could exacerbate educational disparities.	Create institutional policies that govern the use of AI within a residency. Ensure that residents have comparable access to these tools, when appropriate, across all training sites.	Make AI tools accessible to all students by offering free or subsidized educational AI platforms. Use public libraries, community centers, and schools as hubs for access, workshops, and training. Integrate AI education into curricula across all levels, providing diverse formats (text, audio, visual) to accommodate various learning needs and disabilities. Enhance speech recognition models by training them on diverse datasets to better understand accents, and provide support options like live chat or help centers for users facing accent-related challenges.
HIPAA and data privacy (29, 30)	Healthcare AI applications process large amounts of sensitive health information to assist in decision-making. However, many AI-powered health apps, social media platforms, and chatbots are not bound by HIPAA standards. With the rapid pace of AI advancements, the current HIPAA regulations have not adapted to fully address data ownership, breach liability, or the evolving landscape of AI technology.	Ensure that faculty and trainees are familiar with institutional policies that govern granting AI access to patient data, or advocate for the development of these policies where none currently exist.	To enhance data security, de-identification of patient data should be standard practice. Healthcare organizations must continually assess AI systems for security vulnerabilities and improve data protection mechanisms. Additionally, HIPAA should revise its regulations to encompass AI-driven applications and set clearer guidelines on data ownership and responsibility in case of breaches.
Legal, Ethical, and Trust Issues (31)	AI in healthcare raises concerns about the potential for errors in diagnosis or treatment recommendations, leading to legal and ethical dilemmas regarding responsibility and liability. Moreover, if patients are not informed about how their sensitive information is being used or shared, it could erode trust between patients and healthcare providers, negatively impacting patient-physician relationships.	Promote institutional and individual transparency around the use of AI tools in clinical care, with an emphasis on shared decision-making and informed consent. Educate residents and faculty on how to obtain informed consent when using AI tools in clinical care.	As AI becomes more integrated into healthcare, patient education must adapt accordingly. Legal frameworks should address responsibility for AI-related medical errors. Emphasize informed consent to ensure patients understand AI's role in their care. Build trust by educating patients on AI use, safeguarding data privacy with strict consent protocols, and de-identifying patient information.
AI has a massive carbon footprint (32, 33)	AI has a large carbon footprint because of the energy it consumes to develop and train models. The data centers and transmission networks that power AI operations consume a large amount of energy.	Promoting individual and institutional awareness of the impact of AI on our carbon footprint.	To reduce the carbon footprint, consider using cloud-based data centers and energy-efficient hardware, as well as opting for pre-trained models or those with fewer parameters.

RISKS AND MITIGATION STRATEGIES FOR AI IN PSYCHIATRIC EDUCATION

Risk	Elaboration	PD Mitigation Strategies	Broader Mitigation Strategies
Critical appraisal of AI-based research is challenging (34)	Most psychiatrists can read a manuscript on an RCT and evaluate its quality; few can do the same with a study that uses ML approaches. In order for residents to meet level 4 on the PBLI 1 milestone in the year 2025, they will need to possess a basic understanding of ML methodologies and how to evaluate their use in biomedical research.	Programs should develop and implement a curriculum to teach these skills. AADPRT and other professional organizations should lead the development of model curricula and articulate national standards for medical education at both the UME and GME levels. Core skills include being able to: Asses AI models and algorthyms used; scrutinize whether the selected AI model (e.g., deep learning, reinforcement learning, decision trees) is appropriate for the research objectives; AI models, particularly deep learning, can be black boxes; assess whether the authors attempt to explain or interpret model decisions and if transparency in the algorithm's functioning is provided; understand algorythm biases. Contexualize with related literature.	
AI may confabulate or include inaccurate information (35, 36)	AI-generated content can be susceptible to inaccuracies, with misinformation being a primary concern. Misinformation occurs when AI systems produce false or inaccurate information due to various factors. These include reliance on outdated or unreliable data sources, which may contain obsolete or incorrect information. AI systems can also misinterpret data, leading to flawed conclusions or outputs. Additionally, AI models, especially large language models, are designed to predict likely word sequences based on patterns in their training data rather than verifying factual accuracy. This fundamental approach can lead to the generation of plausible sounding but incorrect information. The challenge of misinformation is further compounded by the AI's inability to critically evaluate the truthfulness of its outputs.	Educate users on AI limitations and ethical practices. RIgorously cross check all claims with reliable sources.	Focus on improving data quality by using accurate and reliable sources and continuously updating algorythm training data. Enhance model training with advanced techniques such as retrieval-augmented generation and incorporate better context understanding. Implement verification mechanisms such as real-time fact-checking and cross-referencing and use human oversight to monitor and correct outputs. Educate both users and developers on AI limitations and ethical practices and conduct regular audits to ensure accuracy and reliability. Require constant vigilance and refinement of AI technologies.
Use of AI in routine tasks (like note writing, presentation creation, papers) undermines our ability to assess the tasks	Dependence on AI: If we let AI handle routine tasks, we become less involved in the process. Assessment Skills: Regularly using AI might lead to a decline in our skills for evaluating and analyzing tasks. For example, if AI writes notes for us, we might not develop or maintain our own skills in summarizing or noting important details. Impact on Quality: Relying on AI could also affect our ability to judge the quality and accuracy of the output, since we might become less practiced in identifying mistakes or understanding context.	Regularly engage with the tasks that AI handles for you. For example, if AI writes notes, review and edit them yourself to stay involved in the process and refine your skills. Develop a habit of critically reviewing AI-generated outputs. Assess the content for accuracy, relevance, and quality, and compare it with what you would have produced. Use AI tools for efficiency but set limits on their use.	Human-In-The -Loop: Use AI to assist, rather than fully replace, human decision-making. A "human-in-the-loop" approach allows humans to intervene when needed, especially for critical or high-stakes tasks. In cases where AI makes errors or when the decision requires human discretion, ensure that the system includes an escalation process to notify a human operator or supervisor. Provide training for employees to understand AI's role in routine tasks and how they can supervise or intervene when necessary. Continuous monitoring.
Reliance on AI may reduce interactions with mentors and supervisors (37, 38)	There is a risk of reduced human interaction with excessive reliance on AI in medical education. "Medicine has a millenia long history of cognitive apprenticeship." The continued interaction with supervisors and mentors is crucial for the development of critical and clinical reasoning skills, as well as professional identity formation.	Do not replace supervision with AI. Train faculty in appropriate uses of AI and encourage them to thoughtfully integrate it into their supervisory relationships.	AI should complement medical education. Learners are already utilizing AI. Educators need to lean into the current era and embrace the responsibility of providing a well-rounded education on AI. In addition to all the benfits of AI, the fallibility and limitations need to be highlighted as a way to calibrate the potential threat of replacing supervision and mentorship. For example, it can be highlighted that AI is a mere co-pilot in medical decision making because it lacks the innate characteristic of empathy (amongst others) to value the importance of a patient's lived experience and reconciling that with the treatment goals.
Integrity of application materials (39, 40)	Learners are relying more on AI to edit aspects of their application portfolio like personal statements, application essays, etc. In some instances these AI tools are being used to write these essays from the start with little personal input from the learner. The AI detection tools available lack reliability in their ability to consistently detect AI use. This is more glaring in scenarios where multiple AI outputs are overlayed to craft a final product. Chen et al also argue that utilizing AI as a tool to edit application materials might not be so different from asking a mentor to proof-read and edit a personal statemnt or engaging paid writing services. They do agree that an essay that's compleletly AI generated removes the personal touch and would not be useful as an appraisal tool. Accessibility to utilizing AI for aplication materials can be a double edged sword of removing inequities for underresourced applicants who might not have mentors to proofread their essays and unduly advantage well-resouced applicants with access to more complex AI versions that are less detectible.	Recommendations for educators from the paper by Chen: - Don't assume that an applicant did not use GAI to craft their application. - Be aware that AI-detecting software is imperfect. - Be mindful of what role each aspect of the application plays as part of a holistic assessment; corroborate free-text responses with objective data; - Consider speaking directly with recommending faculty; - Consider asking applicants to engage in a live assessment and/or writing exercise.	There needs to be clear guidelines for appropriate use of AI (e.g., per Chen, applicants may use GAI to brainstorm ideas, assist with initial draft development, and edit. they should not rely on GAI to create an entire personal statement.) There needs to be a clear code of ethics for use of AI that should be enforced at a system level (e.g. by ERAS).

Section 3

Terminology and Resources

Artificial Intelligence (AI): A broad term that encompasses a set of tools used to complete various tasks, including prediction, text generation, classification, and decision support. As a term, AI can be both confusing and misleading. When considering the applications of AI to medicine, it is generally helpful to be specific about which AI tools we are referencing. Almost everything that is referred to as AI involves the implementation of a machine learning algorithm in a specific context with a defined set of goals.

Augmented Intelligence: The term used by AMA and APA in place of "artificial intelligence."

Machine Learning (ML): The ability of computational algorithms to use data to iteratively improve their performance on tasks such as prediction or classification. Specific types of ML include:

- Supervised: A type of ML in which the target outcomes ("ground truth") are known and labels are provided to the algorithm as it is trained. This type of ML is useful for making diagnoses or automatic clinical support tools.
- Unsupervised: A type of ML in which the underlying classes are unknown or poorly defined. This type of ML is useful for pattern recognition and hypothesis formation.
- Natural language processing: A field of computer science that deals with the modeling of human language.
- Deep learning: A type of ML that uses artificial neural networks to create a multi-layered model that builds on itself.
- Artificial neural networks (ANNs): A specific type of ML algorithm designed to imitate how we think neurons work.
- Large language model: A type of model that uses natural language processing and deep learning to generate predictions about language.

Prompt Engineering: Techniques for crafting input to LLMs such that they are most likely to generate the type of output you want.

Explainable AI: A model that uses algorithms that humans are able to explain in clear terms.

Hallucination: A term used to describe the presentation of incorrect information in response to questions or prompts. In clinical terms, this is equivalent to confabulation.

Black box AI: A model that uses algorithms so complex that they cannot be explained in terms humans can easily grasp.

Precision psychiatry: The idea that we can create models, generally using a combination of machine learning approaches, that will allow us to predict outcomes for individual patients, which in turn will facilitate tailored clinical interventions.

(1) https://chatgpt.com/g/g-usvTSU0RB-teen-angst-patient-simulation
(2) https://www.canva.com/
(3) https://openai.com/index/dall-e-3/
(4) https://chatgpt.com/g/g-sQwRRQqqb-child-prite-2023-review-helper
(5) Li, Y., Rakovic, M., Poh, B. X., Gaševic, D., & Chen, G. (2022). Automatic classification of learning objectives based on bloom's taxonomy. In Educational Data Mining 2022. International Educational Data Mining Society.
(6) https://www.ithenticate.com/
(7) Dias, R. D., Gupta, A., & Yule, S. J. (2019). Using machine learning to assess physician competence: a systematic review. Academic Medicine, 94(3), 427-439.
(8) Bohler F, Aggarwal N, Peters G, Taranikanti V. Future Implications of Artificial Intelligence in Medical Education. Cureus. 2024 Jan 8;16(1):e51859. doi: 10.7759/cureus.51859. PMID: 38327947; PMCID: PMC10848885.
(9) Safranek, C. W., Sidamon-Eristoff, A. E., Gilson, A., & Chartash, D. (2023). The role of large language models in medical education: applications and implications. JMIR medical education, 9, e50945.
(10) Creed, T. A., Salama, L., Slevin, R., Tanana, M., Imel, Z., Narayanan, S., & Atkins, D. C. (2022). Enhancing the quality of cognitive behavioral therapy in community mental health through artificial intelligence generated fidelity feedback (Project AFFECT): a study protocol. BMC health services research, 22(1), 1177.

(11) https://fireflies.ai/
(12) https://www.jibble.io/time-and-attendance-healthcare
(13) Chen Z. Collaboration among recruiters and artificial intelligence: removing human prejudices in employment. Cogn Technol Work. 2023;25(1):135-149. doi: 10.1007/s10111-022-00716-0. Epub 2022 Sep 28. PMID: 36187287; PMCID: PMC9516509.
(14) Tulshian P. Artificial Intelligence: Reflections on Its Use to Promote DEIA Principles for Residency Recruitment. Fam Med. 2024;56(6):344-345. https://doi.org/10.22454/FamMed.2024.951121.
(15) https://www.tableau.com
(16) Barac, M., Scaletty, S., Hassett, L. C., Stillwell, A., Croarkin, P. E., Chauhan, M., ... & Dyrbye, L. N. (2024). Wearable Technologies for Detecting Burnout and Well-Being in Health Care Professionals: Scoping Review. Journal of medical Internet research, 26, e50253.
(17) https://publichealth.tulane.edu/blog/data-driven-decision-making/
(18) https://www.mentalyc.com/
(19) https://www.limbic.ai/
(20) Quinn, T. P., Hess, J. L., Marshe, V. S., Barnett, M. M., Hauschild, A. C., Maciukiewicz, M., ... & Machine Learning in Psychiatry (MLPsych) Consortium. (2024). A primer on the use of machine learning to distil knowledge from data in biological psychiatry. Molecular psychiatry, 29(2), 387-401.
(21) Landes, S. J., Matarazzo, B. B., Pitcock, J. A., Drummond, K. L., Smith, B. N., Kirchner, J. E., ... & Katz, I. R. (2024). Impact of implementation facilitation on the REACH VET clinical program for veterans at risk for suicide. Psychiatric services, appi-ps.
(22) Bertl, M., Metsallik, J., & Ross, P. (2022). A systematic literature review of AI-based digital decision support systems for post-traumatic stress disorder. Frontiers in Psychiatry, 13, 923613.
(23) https://www.limbic.ai/
(24) Knopp MI, Warm EJ, Weber D, Kelleher M, Kinnear B, Schumacher DJ, Santen SA, Mendonça E, Turner L. AI-Enabled Medical Education: Threads of Change, Promising Futures, and Risky Realities Across Four Potential Future Worlds. JMIR Med Educ. 2023 Dec 25;9:e50373. doi: 10.2196/50373. PMID: 38145471; PMCID: PMC10786199.
(25) https://www.nytimes.com/2021/03/15/technology/artificial-intelligence-google-bias.html
(26) https://medicine.yale.edu/news-article/eliminating-racial-bias-in-health-care-ai-expert-panel-offers-guidelines/
(27) https://teaching.cornell.edu/generative-artificial-intelligence/ai-accessibility
(28) https://www.standrewslawreview.com/post/the-future-of-a-i-for-low-income-communities-and-developing-countries
(29) https://journal.ahima.org/page/updating-hipaa-security-to-respond-to-artificial-intelligence
(30) https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10937180/#:~:text=Considerations%20for%20AI%20Developers%20and%20vendors&text=It%20is%20true%20that%20HIPAA,with%20its%20spirit%20and%20purpose.
(31) Civaner MM, Uncu Y, Bulut F, Chalil EG, Tatli A. Artificial intelligence in medical education: a cross-sectional needs assessment. BMC Med Educ. 2022 Nov 9;22(1):772. doi: 10.1186/s12909-022-03852-3. PMID: 36352431; PMCID: PMC9646274.
(32) https://cacm.acm.org/news/the-carbon-footprint-of-artificial-intelligence/
(33) https://news.climate.columbia.edu/2023/06/09/ais-growing-carbon-footprint/
(34) Grzenda, A., Kraguljac, N. V., McDonald, W. M., Nemeroff, C., Torous, J., Alpert, J. E., ... & Widge, A. S. (2021). Evaluating the machine learning literature: a primer and user's guide for psychiatrists. American Journal of Psychiatry, 178(8), 715-729.
(35) https://news.vt.edu/articles/2024/02/AI-generated-fake-news-experts.html
(36) https://www.tc.columbia.edu/articles/2023/october/artificial-intelligence-optimists-concerned-about-risks-of-visual-misinformation/
(37) Cooper A, Rodman A. AI and Medical Education - A 21st-Century Pandora's Box. N Engl J Med. 2023 Aug 3;389(5):385-387. doi: 10.1056/NEJMp2304993. Epub 2023 Jul 29. PMID: 37522417.
(38) Van der Niet AG, Bleakley A. Where medical education meets artificial intelligence: 'Does technology care?'. Med Educ. 2021 Jan;55(1):30-36. doi: 10.1111/medu.14131. Epub 2020 Mar 30. PMID: 32078175.
(39) Chen JX, Bowe S, Deng F. Residency Applications in the Era of Generative Artificial Intelligence. J Grad Med Educ. 2024 Jun;16(3):254-256. doi: 10.4300/JGME-D-23-00629.1. Epub 2024 Jun 13. PMID: 38882414; PMCID: PMC11173008.
(40) AI Detection in Education by Leon Furze