AI Design Issues in Education

Jonathan Adams
KatieLee Riddle

https://doi.org/10.34074/proc.2302012
Correspondence: jonathan.adams@toiohomai.ac.nz

Essay

This publication may be cited as:

Adams, J., & Riddle, K. (2023). AI Design Issues in Education. In J. L. Savage, J. Hoffman, & M. Shannon (Eds.), *Proceedings: ITP Research Symposium 2022, 30 November – 2 December* (pp. 119–131). ePress, Unitec | Te Pūkenga. https://doi.org/10.34074/proc.2302012

Contact:
epress@unitec.ac.nz
www.unitec.ac.nz/epress/
Unitec | Te Pūkenga
Private Bag 92025, Victoria Street West
Auckland 1142
Aotearoa New Zealand

ISBN: 978-1-99-118343-9

ABSTRACT

Toi Ohomai | Te Pūkenga engaged artificial intelligence (AI) developers in mid-2021 to help kaiako (teachers) convert existing subject content into self-generating adaptive learning material. The AI-generated content helped students learn at their individual levels and paces. Feedback highlighted improved student comprehension, and time savings and positive professional development for kaiako. However, considerations in the design and use of AI requires clarity about the range of open risks, exploitative processes in data extraction and training, and how these applications can encode bias and impact Indigenous and creator rights. This essay is an introduction to unresolved issues in AI design, and for decision makers to consider three key questions when choosing or using AI: how the application is built; who owns the knowledge created; and whose values shaped the training data and application design. Responsible use of AI requires very considered and transparent selection of any training data and model development, even prior to designing the AI application that is built on the model. Importantly, companies designing or deploying AI applications in Aotearoa New Zealand need to consider the Waitangi Tribunal recommendations for practical changes to the law, to include consent and protection for taonga and mātauranga Māori, to protect Māori cultural works, language, arts and heritage against unauthorised access or use.

KEYWORDS

AI design, training data, AI in education, adaptive learning material

INTRODUCTION

"In the age of big data, data itself has become the raw material of production and a new source of immense social and economic value" (Sterling et al., 2021, p. 3).

In mid-2021 Toi Ohomai Institute of Technology engaged with an Australian-based AI company to alleviate some of the workload for kaiako (teachers) that an increase in blended learning was requiring. Following a pilot and based on student feedback, the application was licensed from February 2022 for a variety of programmes across three faculties at Toi Ohomai. Kaiako uploaded existing text-based course content and slides saved to PDF to be converted into self-generating adaptive learning material, visible only to their students, such as flashcards and quiz questions that self-mark. A web interface with individual logins for students and kaiako presented content for the students and progression reports for kaiako. The application used natural language processing to both generate and grade educational questions to help students learn at their individual levels and paces. The generated content provides a student experience that varies based on individual understanding of the material.

STUDENT SURVEY FINDINGS, OBSERVATIONS, AND FEEDBACK FROM KAIAKO

Anonymous online student surveys were undertaken to assess the effectiveness of the application. In a survey conducted with nursing students in 2021, 86% of respondents said the AI tool, with individualised training on knowledge gaps and formative checks, was useful in preparing for their tests. Follow-up surveys with students in other courses in 2022 confirmed student enthusiasm for the AI-generated slide and quiz content formats. Figure 1 shows the results from student surveys.

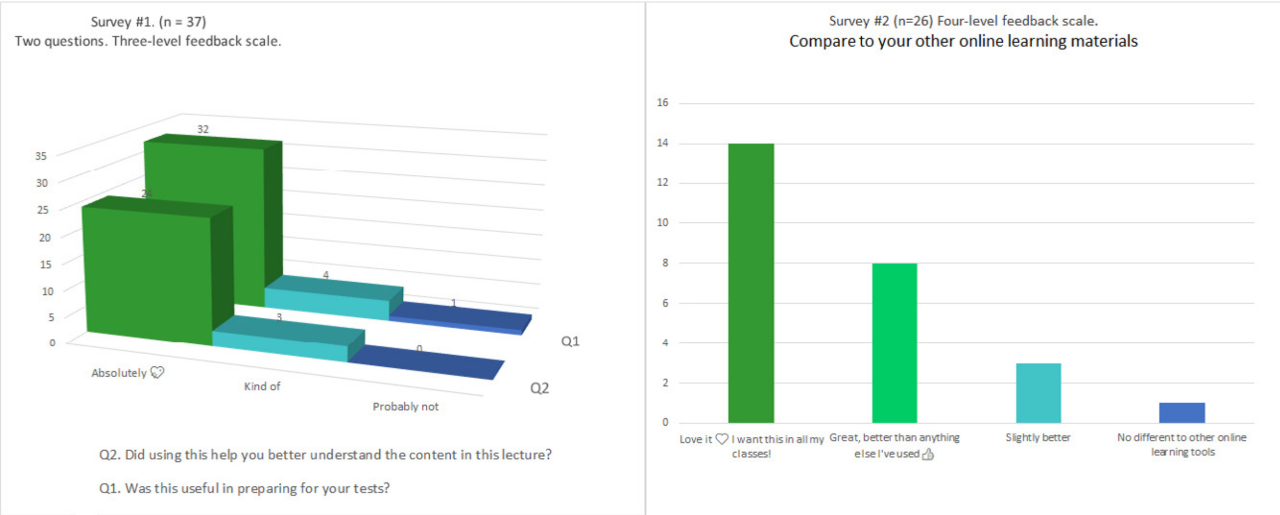

Figure 1. Results from student surveys.

In addition to student surveys, observations were made by the author and feedback gathered directly from kaiako highlighted a number of key benefits from the use of AI-generated content being used in tandem with existing course material. The benefits for teachers included:

1. Course content was interactive, with formative quizzes and flashcards based on each week's work.
2. Better student understanding and retention of lesson content.
3. Positive professional development for staff in relation to AI.

Kaiako commented that they spent less time generating course content, which gave them more scope to focus on teaching or enhancing the quality of courses. The AI tool allowed students' learning progression to be automated in a manageable digital format that was unique to each student. Specific-use AI applications such as this retain the human element where the kaiako curates and is accountable for content, compared with a student using a prompt-based AI tool directly to generate their own formative tests. This approach places the kaiako between the application and student, with the kaiako as the guide.

KEY ISSUES FOR EDUCATORS

Alongside the benefits, a decision to use an emerging technology such as AI in education also requires software vendors, institutions and educators to be clear about AI design issues and risks, and how these applications can impact Indigenous and creator rights. There are a number of pressing issues as they relate to Aotearoa New Zealand education, as well as principles in AI design where "laws and norms ha[ve] not caught up with AI's unique risks or society's needs" (Microsoft, 2022a). This essay is an introduction to these wider unresolved issues in AI design, to inform decision makers in tertiary education, and are not about any particular application.

Awareness of AI capability has grown exponentially following the November 2022 release of ChatGPT by OpenAI. This has spurred further interest in the vocational education sector in how AI applications can work with humans in a non-programmatic way. Based on our experiences since 2021, and the increase of available AI applications, we have identified three key questions to be considered in using AI within education and more broadly:

1. How has an application been built, or what is the provenance of the application?
2. Who owns the knowledge that is being created?
3. Whose values shaped the design of the AI application?

The development of consumer AI-powered generative art apps, and the training databases that sit behind them, provide a useful case for considering these questions due to their visibility and the public discussion on their impacts (Williams, 2023). These generative art apps are possible with automation enabled by web crawlers built on platforms such as Microsoft Azure and Amazon AWS, to create the datasets of image-text labels needed to train models, to then deliver meaningful outputs (Bindal, 2022). While derived works such as generative art or portrait applications may appear innovative, the creation of these products involves the scanning and analysis of original creative works that are publicly accessible, but not necessarily in the public domain.

Implications for rights and interests that arise within copyright law, the Wai 262 claim, and trade treaties need urgent consideration. Allowing developers to deploy software without accountability risks technological lock-in (Héder, 2021), lending preference to speed over what is equitable, and without consideration of wider social consequences shaping society.

The final area of concern is the training for machine learning and development of predictive models. Briefly, there are two main AI data-training methods: (1) unsupervised or self-learning, such as clustering objects with similar characteristics; (2) supervised machine learning, with regression, and classification where humans are included, labelling features in training data that are then used to predict labels for new items.

Demand for specialised and more accurate models has spawned a gig-economy workforce of millions of contractors who are hired for tagging or categorising content to build the value of training data and their models (Appen, n.d.; TELUS International, n.d.). If trainers are not aware of a diversity of cultural values, or if they lack sensitivity to labels, this can result in default machine learning models that are not appropriate when deployed in different contexts.

An early and influential example was ImageNet database (https://image-net.org/), which hosted images of real people with offensive labels for more than a decade (Denton et al., 2021; Gershgorn, 2017). Exploration of this issue was covered critically (Crawford & Paglen, 2019), prompting the question of why classification of humans and their attributes should even be codified. Crawford and Paglen (2019) were then challenged (Lyons, 2020) on informed consent for the use of images for public exhibitions used to raise awareness of the topic, helpfully affirming the non-negotiability of informed consent when using human data.

GENERATIVE DESIGN: JUST LOOKING, THANKS

A recent example of the continuing gap between creator rights and controls on image processing impacting artists (Williams, 2022), is the Lensa app from Prisma Labs generating personalised portraits in different artistic styles (Prisma Labs, n.d.). The image-text database Lensa was trained on was a copy of the Stable Diffusion model. Stable Diffusion was itself trained on images from another dataset called LAION-5B, “which scraped non-curated image-text-pairs from the internet” (StabilityAI, 2023). The process of unsupervised ‘training’ or ‘self-learning’ includes computing similarity scores between pictures and texts, which means the images are being used and then “subsequently discarded” in a non-expressive use. Differences in copyright regulatory frameworks between EU and the US drive developers to publish methods “to undermine the restrictive and anachronistic EU copyright laws and pave the way for a Fair Use” (Schäfer, 2016, p. 4503). LAION.AI, based in Germany, distances itself from the ‘copying’ of images, but has released its download tool (img2dataset) on its website (Beaumont, 2022). The img2dataset tool “allows downloading 100 million images from our list of URLs in 20 hours with the URL included” (LAION, n.d.).

Stable Diffusion includes ‘Use Restrictions’ in downstream versions, such as Lensa. For example, users are “not to use the Model or Derivatives of the Model: – In any way that violates any applicable national, federal, state, local or international law or regulation” (Rombach et al., 2022, p. 5). Copyright (CC4) is then being applied by app developers on applications built on models that used images trained from LAION-5B and Stable Diffusion.

A lineage of the use of creators’ works and copyright implications are as listed:

- Artist or photographer
- LAION-5B

- Stable Diffusion
- Prisma Labs
- Lensa App
- Consumer

In summary, companies developing AI applications, including LAION.AI researchers, have not addressed the difference between online availability and works being in the public domain.

Accountability for use is already being tested by copyright owners such as Getty, who in January 2023 launched legal proceedings in London against Stability.ai (*Business Insider*, 2023; Getty Images, 2023). This serves as a reminder that, even for research, while it may be possible to use a crawler on the internet to train AI models for a non-expressive use, it may not be legal. LAION.AI operates as a charity, promoting itself as “Truly open AI. 100% non-profit. 100% free” (Beaumont, 2022) but is supported financially by commercial AI companies such as Huggingface.co and Stability.ai, each with over $1 billion in valuations (Crunchbase, n.d.a; n.d.b). Justifying this open access as research and using the rhetoric of ‘democratising AI’ (LAION, n.d.) ignores Indigenous rights, creative rights and data sovereignty, and is at odds with CARE principles for Indigenous communities to realise opportunities within the knowledge economy for their collective benefit (Carroll et al., 2020).

This is not constrained to images but also text-based works. In the US, more than 8000 authors are calling on OpenAI, Alphabet (Google), Meta (Facebook), Stability AI, IBM and Microsoft to obtain consent, and credit and fairly compensate writers for the use of copyrighted materials in training AI (Fung, 2023). In Australia, developers could be facing liability, as “In general, a person or organisation can rely on a fair dealing exception only for their own use of copyright material” (Australian Copyright Council, 2020, p. 2). There has also been strong push-back from some artists wanting to opt out or reassert their rights (Williams, 2023). In Aotearoa New Zealand, software vendors that rely on the premise that there is no ‘copy’ occurring need to consider the “independent economic significance” (Copyright Licensing Limited, 2020, p. 1) that results when “transient” scanning occurs, referred to in the Copyright Act 1994 s43A, when creating the text and image pairs for training AI models.

AN AOTEAROA NEW ZEALAND MODEL FOR FAIR USE

Deloitte Access Economics was commissioned by Google (Deloitte, 2018) to analyse the economic impact of introducing fair use in Aotearoa New Zealand. They reported that artificial intelligence methods to train algorithms with large amounts of data that are permissible in the US are not permitted in Aotearoa. Currently, a funding model is being established for supporting creators of artistic works when resold on the secondary art market. This funding model is aligned with the Aotearoa New Zealand Free Trade Agreement with the EU and UK, and allows for 5% of royalties for artists in downstream sales of their work. It is designed to support sustainable careers through ongoing royalty payments. According to the New Zealand Government, this new scheme, jointly announced by the then deputy Prime Minister, is “about fairness … honouring the tremendous artistic skill and creativity of so many of our visual artists” (Sepuloni & O’Connor, 2022, para. 2). However, there is no mention of any effect this may have for artists whose works were used to train other digital art, generated by AI, illustrating how the law continues to lag behind the pace of technology advances.

Accepting that data from a Western perspective is a property and a commodity in Aotearoa New Zealand means accepting the principles of Te Tiriti o Waitangi are also applicable to Māori data (Taiuru, 2020). This only partially answers the question of ‘Who owns the knowledge?’, as it is relevant for everyone. The justification of fair use of publicly accessible content for research and AI training, without informed consent or compensation, is not accepted when that research is commercialised (Fung, 2023).

Anyone designing or deploying AI applications should consider the Waitangi Tribunal recommendations in the WAI 262 report, for practical changes to the law to include consent and protection of taonga and mātauranga Māori to protect Māori cultural works, language, arts and heritage from unauthorised use. They should also align to plans

(Open Government Partnership New Zealand, 2022) to embed the Algorithm Charter across government data-management practice (Commitment 8). Many IP laws have provision for Māori Advisory Committees and there is also significant guidance on cultural sensitivity within the New Zealand Intellectual Property Office (New Zealand Intellectual Property Office, n.d.).

Evidence of content that has been indexed and therefore used for AI model training is easily searchable in online developer databases that show many images and cultural artefacts from Aotearoa New Zealand with 'All Rights Reserved' or potentially in breach of copyright and use.

Delays in implementing protections recommended by the Waitangi Tribunal more than ten years ago have allowed continued exploitation by developers of images for commercial gain that should have been prevented well before 2023. Examples include filters in Snapchat using Māori cultural imagery that did not have clearance or approval to be used (Brown, 2022; RNZ, 2022). An absence of the enforcement of Treaty Article II protections highlights the gap with mature social media platforms such as Meta in providing attribution features. The tools available in Facebook and Instagram provide a free 'Rights Manager for Images', which uses image matching technology, along with an IP reporting system (Axelgard, 2020). The failure in the Meta design is putting the onus on the infringed to manually report issues.

EXACERBATING SOCIETAL BIASES – WHOSE VALUES ARE WE USING?

Another key issue with image use is the potential for harm from the mass automated tagging of images or unsupervised learning, which encodes existing biases of an unmoderated internet. These datasets are being built on non-curated images scraped from the internet for 'research purposes', which in their disclaimers highlight the potential for harmful content. This is the same for language models, which when "trained on large, un-curated, static datasets from the Web encode hegemonic views that are harmful to marginalised populations" (Bender et al., 2021, p. 615). This is an issue widely accepted by both developers, as stated in their licenses, and researchers. The image word cloud in Figure 2 shows the range of 'potential' Not Safe for Work (NSFW) image-text labels using English text in the LAION-5B database, even before any cultural or contextual issues are identified. Subsets of these databases can be curated to provide 'aesthetic' content, but this leaves open the ability to vary the weightings that include or exclude content. Decisions about what is suitable or 'aesthetic' are made by developers "in a scale from 0 to 10 to be good looking or not" (LAION-AI, 2022).

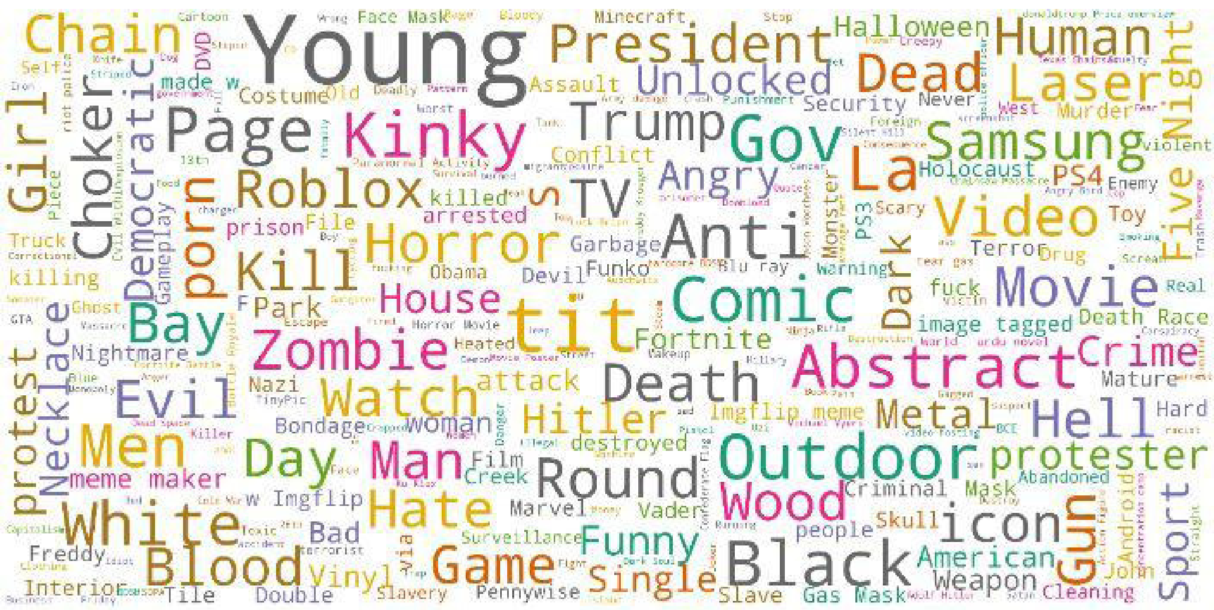

Figure 2. Word Cloud published to illustrate range of identified potentially inappropriate concepts or NSFW content (Schuhmann et al. 2022, p. 33. CC BY 4.0, https://creativecommons.org/licenses/by/4.0/).

While Prisma launched a business-to-consumer app on the backs of artists' styles, other companies are building private datasets using human trainers. The issue here is opaqueness, the varying cultural perspectives of the trainers, and which values are subsequently being fed into the AI models. Companies such as Appen (Australia) (Appen, n.d.) and TELUS International (Canada) (TELUS International, n.d.) recruit contract workers globally on "location-based minimum wage" (Bogle, 2022, para. 26) to work on projects that tag or categorise images, text and video. This human-powered extractive work provides no visibility into the cultural contexts or potential bias that a dispersed and temporary workforce is producing. An industry and process built on casual minimum-wage labour and unsupervised processing of public-domain resources draws an obvious parallel with exploitation of resources in earlier industrial revolutions.

THE DANGERS OF DEFAULT

The ease and subtlety with which societal biases influence AI, as flagged by the LAION-5B license, is seen in the default Google Cloud Vision labels used when describing an image loaded by the author, shown in Figure 3. The screenshots of results across three different years in Figure 3 show how implicit cultural-bias labels can find their way into any default dataset where tags are applied from humans, or acquired in unsupervised learning. But, unlike humans, nobody is held accountable for perpetuating the bias that can be fed back unnoticed into new applications.

Figure 3. Google Cloud Vision labels (faces obscured for publishing). Image from author, results from https://cloud.google.com/vision/#vision-api-demo

The labels in Figure 3 show how weightings change, and also demonstrate how the cultural norms and subjective labels used in one society are not appropriate for another. The phrase "blue-collar worker" comes from historical class hierarchies and occupations, and is a synonym for someone who is not highly educated (Oxford Reference, n.d.), thus perpetuating deficit mindsets. This may be innocuous on its own, but if it influences a credit algorithm rather than retail mall promotion (O'Shea, 2023) then the results could be very consequential for an individual. The use of biometric categorisation or inference of gender, age and mood already used in Aotearoa New Zealand shopping malls would be banned in the proposed EU Artificial Intelligence Act (European Parliament, 2023).

Using a disclaimer to say a service is a demonstration or research service is not enough, there is an absence of cultural context and transparency that is a flag for quality and appropriateness of the service. This example is a symptom of a fundamental AI alignment challenge in "deciding to whom these systems should be aligned" (Leike et al., 2022, para. 5).

CONCLUSION

AI developers and researchers have devolved responsibility to downstream users within their software licenses. Thus, it is incumbent on all participants to understand the provenance of their applications and verify with suppliers their existing obligations and responsibilities.

Government agencies such as the Ministry of Business, Innovation and Employment and the Ministry of Education as signatories to the Algorithm Charter for Aotearoa New Zealand should be ensuring their vendors and suppliers also commit to the six core principles. These include: 2. "Partnership – Embedding a Te Ao Māori perspective in the development and use of algorithms consistent with the principles of the Treaty of Waitangi" (Stats NZ, 2020, p. 3).

Regulatory clarification is needed for non-expressive uses of data (Deloitte Access Economics, 2018), for data guardianship and excellence (Sambasivan et al., 2021). There are significant opportunities to promote Indigenous innovation and locally developed AI models, alongside wider and earlier participation in AI development (Ministry of Business, Innovation and Employment, 2019).

The inventor of ImageNet, Dr Fei-Fei Li, co-founded a diversity and inclusion initiative for AI more than seven years ago (AI4ALL, n.d.), while Mira Murati, CTO at OpenAI, shared in February 2023 they "need a lot more input that goes beyond the technologies, definitely regulators and governments and everyone else" (Simons, 2023, para. 20). This willingness by AI leaders to create change needs to reflect into continuous, sustained action by counterparts in political and community leadership roles.

Technology companies intending to sell services or deploy products must also consider their jurisdictional liability and the provenance of the AI models to be used in Aotearoa New Zealand. There is "the need to invest significant resources into curating and documenting LM training data" (Bender et al., 2021, p. 615), and disclosure of origin and community support statements (collective consent) (Hudson, 2019). The urgency for transparency and community engagement is driven by the momentum and resources already allocated in commercial and open-source AI development. The scale of AI leaders' ambitions is illustrated by Stability AI founder Emad Mostaque asking "who's building the Japan model, or the India model or others? Well, we are" (Guo, 2023, para. 24).

With OpenAI capabilities being released on Microsoft's Azure (Boyd, 2022) and open-source models on Amazon Web Services (Bathgate, 2023), stakeholders should be actively participating at the design stage, directing design alongside the development community. An example of this gap between users and developers is the requirement for developers to list adverse impacts and identify stakeholders (Microsoft, 2022b). The issue is that developers will be unaware of their unconscious bias, hence the need for stakeholders at the design stage, where "the functional logics of a given technology echo the gender and racial dynamics of the industry that produced it" (West et al., 2019, p. 8).

Public-service procurement advice for Generative AI is to consider vendor reputation and have protections in contracts for privacy, security, ethical risks and vendor lock-in (Department of Internal Affairs, 2023). Tertiary education providers with public funding should preference vendors demonstrating commitment to the *Algorithm charter for Aotearoa New Zealand* (Stats NZ, 2020) and Te Kawa Mataaho Public Service Commission's plans to further embed transparency and accountability in the use of algorithms with the Open Government Partnership New Zealand (2022). The Education and Training Act 2020 requires tertiary education institutions to recognise and respect the Crown's responsibility to give effect to Te Tiriti o Waitangi. Schedule 13 contains the Charter requiring Te Pūkenga to reflect Māori–Crown partnerships (New Zealand Government, 2020). A genuine partnership involves respect for Māori aspirations, values and property, accepting that thinking through legal and data provenance issues needs to be part of the design process, not an afterthought.

Taking action does not mean constraining innovation or exploration, but does mean proactively designing for these legal, IP and data sovereignty issues beforehand. The trajectory of AI development discussed by Héder (2021) depends on the actions of society, social context and control, and the social license to operate. Educators cannot be bystanders “shield[ing] the creators of these systems from accountability while its deterministic, calculative power intensifies social classification and control” (Campolo & Crawford, 2020, p. 1).

The change in relationship required between AI developers, vendors and customers is reflective of the wider societal transformation occurring in Aotearoa New Zealand. As the Waitangi Tribunal has affirmed, our relationships must change “into a twenty-first century relationship of mutual advantage in which, through joint and agreed action, both sides end up better off than they were before they started” (2011, p. 17).

ACKNOWLEDGEMENTS

Thanks to kaiako and research kaimahi Josh Burrell, Christine M. Cheyne, Marco Fuchser, Dr Philip Lopez and Nikhil Nayyar, and Andre Avedissian and Dean Mikan.

REFERENCES

AI4ALL. (n.d.). *Our story*. https://ai-4-all.org/about/our-story/

Appen. (n.d.). *How it works*. https://connect.appen.com/qrp/public/how_it_works

Australian Copyright Council. (2020). *Fair dealing: what can i use without permission? Information Sheet G079v09*. https://www.copyright.org.au/browse/book/ACC-Fair-Dealing:-What-Can-I-Use-Without-Permission-INFO079

Axelgard, D. (2020). *Helping creators and publishers manage their intellectual property*. https://about.fb.com/news/2020/09/helping-creators-and-publishers-manage-their-intellectual-property/

Bathgate, R. (2023). *AWS and Hugging Face partner to 'democratise' ML, AI models*. https://www.itpro.co.uk/cloud/370113/aws-and-hugging-face-partner-to-democratise-ml-ai-models

Beaumont, R. (2022). *LAION-5B: A new era of open large-scale multi-modal datasets*. https://laion.ai/blog/laion-5b/

Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). *On the dangers of stochastic parrots: Can language models be too big?* Association for Computing Machinery. https://doi.org/10.1145/3442188.3445922

Bindal, A. (2022). *Stability AI builds foundation models on Amazon SageMaker*. https://aws.amazon.com/blogs/machine-learning/stability-ai-builds-foundation-models-on-amazon-sagemaker/

Bogle, A. (2022). Behind 'miracle' AI is an army of 'ghost workers' — and they're speaking out. *ABC*. https://www.abc.net.au/news/science/2022-10-14/artificial-intelligence-ai-appen-data-labelling-ghost-workers/101531084

Boyd, E. (2022). *General availability of Azure OpenAI Service expands access to large, advanced AI models with added enterprise benefits*. https://azure.microsoft.com/en-us/blog/general-availability-of-azure-openai-service-expands-access-to-large-advanced-ai-models-with-added-enterprise-benefits/

Brown, L. (2022). *Growing anger over use of moko, mataora in image filters: 'That's a mockery'*. https://www.rnz.co.nz/news/te-manu-korihi/474142/growing-anger-over-use-of-moko-mataora-in-image-filters-that-s-a-mockery

Business Insider. (2023). *AI art generators face separate copyright lawsuits from Getty Images and a group of artists*. https://www.businessinsider.com/ai-art-artists-getty-images-lawsuits-stable-diffusion-2023-1

Campolo, A., & Crawford, K. (2020). Enchanted determinism: Power without responsibility in artificial intelligence. *Engaging Science, Technology, and Society*, *6*, 1–19. https://doi.org/10.17351/ests2020.277

Carroll, S. C., Garba, I., Figueroa-Rodríguez, O., Holbrook, J., Lovett, R., Materechera, S., Parsons, M., Raseroka, K., Rodriguez-Lonebear, D., Rowe, R., Sara, R., Walker, J., Anderson, J., & Hudson, M. (2020). The CARE Principles for Indigenous data governance. *Data Science Journal*, *19*, 1–12. https://doi.org/10.5334/dsj-2020-043

Copyright Licensing Limited. (2020). *Your guide to New Zealand copyright. Copyright (New Technologies) Amendment Act 2008*. https://www.copyright.co.nz/downloads/assets/5204/1/fact-sheet:-copyright-amendment-act-(2008).pdf

Crawford, K., & Paglen, T. (2019). *Excavating AI: The politics of training sets for machine learning*. https://excavating.ai

Crunchbase. (n.d.a). *Series C – Hugging Face*. https://www.crunchbase.com/funding_round/hugging-face-series-c--74baeb2c

Crunchbase. (n.d.b). *Stability AI*. https://www.crunchbase.com/organization/stability-ai/company_financials

Deloitte Access Economics. (2018). *Copyright in the digital age: An economic assessment of fair use in New Zealand*. Deloitte. https://www2.deloitte.com/content/dam/Deloitte/nz/Documents/Economics/dae-nz-copyright-fair-use.pdf

Denton, E., Hanna, A., Amironesei, R., Smart, A., & Nicole, H. (2021). On the genealogy of machine learning datasets: A critical history of ImageNet. *Big Data & Society*, *8*(2), 20539517211035955. https://doi.org/10.1177/20539517211035955

Department of Internal Affairs. (2023). *Initial advice on Generative Artificial Intelligence in the public service July 2023*. https://www.digital.govt.nz/assets/Standards-guidance/Technology-and-

architecture/Generative-AI/Joint-System-Leads-tactical-guidance-on-public-service-use-of-GenAI-July-2023.pdf
European Parliament. (2023). Artificial Intelligence Act. https://www.europarl.europa.eu/doceo/document/TA-9-2023-0236_EN.html
Fung, B. (2023). Thousands of authors demand payment from AI companies for use of copyrighted works. *CNN*. https://edition.cnn.com/2023/07/19/tech/authors-demand-payment-ai/index.html
Gershgorn, D. (2017). *The data that transformed AI research—and possibly the world*. https://qz.com/1034972/the-data-that-changed-the-direction-of-ai-research-and-possibly-the-world
Getty Images. (2023). *Getty Images statement*. https://newsroom.gettyimages.com/en/getty-images/getty-images-statement
Guo, S. (2023). *No priors. Episode 103. Emad Mostaque*. https://sarahguo.com/blog/emadmostaque
Héder, M. (2021). AI and the resurrection of Technological Determinism. *Információs Társadalom*, *XXI*(2), 119–130. https://doi.org/10.22503/inftars.XXI.2021.2.8
Hudson, M. (2019). *Maori data sovereignty, Global Indigenous Data Alliance & the CARE Principles*. https://static1.squarespace.com/static/58e9b10f9de4bb8d1fb5ebbc/t/5dd1c6306d59eb44eaf7297e/1574028887381/
LAION. (n.d.). *About*. https://laion.ai/about/
LAION-AI. (2022). *LAION-Aesthetics V1*. https://github.com/LAION-AI/laion-datasets/blob/main/laion-aesthetic.md#training-set
Leike, J., Schulman, J., & Wu, J. (2022). *Our approach to alignment research*. https://openai.com/blog/our-approach-to-alignment-research
Lyons, M. J. (2020). *Excavating "Excavating AI": The elephant in the gallery*. https://arxiv.org/pdf/2009.01215.pdf
Microsoft. (2022a). *Microsoft Responsible AI Standard v2*. https://blogs.microsoft.com/wp-content/uploads/prod/sites/5/2022/06/Microsoft-Responsible-AI-Standard-v2-General-Requirements-3.pdf
Microsoft. (2022b). *Microsoft Responsible AI Impact Assessment Template*. https://query.prod.cms.rt.microsoft.com/cms/api/am/binary/RE5cmFk
Ministry of Business, Innovation and Employment. (2019). *Notes from taonga works and copyright kōrero – 26 March 2019*. https://www.mbie.govt.nz/assets/notes-from-taonga-works-and-copyright-korero.pdf
New Zealand Government. (2020). Education and Training Act 2020. https://www.legislation.govt.nz/act/public/2020/0038/latest/LMS280244.html
New Zealand Intellectual Property Office. (n.d.). *Concepts to understand*. https://www.iponz.govt.nz/about-ip/maori-ip/concepts-to-understand/
Open Government Partnership New Zealand. (2022). *New Zealand's fourth National Action Plan 2023–2024*. https://ogp.org.nz/new-zealands-fourth-national-action-plan-2023-2024?e63=action_viewall
O'Shea, R. (2023). Facial detection used by Westfield malls for targeted advertising. *Consumer*. https://www.consumer.org.nz/articles/facial-detection-used-by-westfield-malls-for-targeted-advertising
Oxford Reference. (n.d.). *Industrial worker*. https://www.oxfordreference.com/display/10.1093/oi/authority.20110803100002290?rskey=BVP03f&result=3
Prisma Labs. (n.d.). *Prisma Labs*. https://prisma-ai.com
RNZ. (2022). *Snapchat removes moko, mataora filters after outcry*. https://www.rnz.co.nz/news/te-manu-korihi/474426/snapchat-removes-moko-mataora-filters-after-outcry
Rombach, R., Esser, P., & contributors. (2022). *Stable diffusion license*. https://huggingface.co/spaces/CompVis/stable-diffusion-license
Sambasivan, N., Kapania, S., Highfill, H., Akrong, D., Paritosh, P., & Aroyo, L. M. (2021, May). "Everyone wants to do the model work, not the data work": Data cascades in high-stakes AI. *Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems*, 39, 1–15. https://doi.org/10.1145/3411764.3445518
Schäfer, R. (2016). CommonCOW: Massively huge web corpora from CommonCrawl data and a method to distribute them freely under restrictive EU copyright laws. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)* (pp. 4500–4504). European Language Resource Association. https://aclanthology.org/L16-1712

Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., Schramowski, P., Kundurthy, S., Crowson, K., Schmidt, L., Kaczmarczyk, R., & Jitsev, J. (2022). *LAION-5B: An open large-scale dataset for training next generation image-text models*. https://doi.org/10.48550/arxiv.2210.08402

Sepuloni, C., & O'Connor, D. (2022). *Government ensures artists get resale royalties*. https://www.beehive.govt.nz/release/government-ensures-artists-get-resale-royalties

Simons, J. (2023). The creator of ChatGPT thinks AI should be regulated. *Time*. https://time.com/6252404/mira-murati-chatgpt-openai-interview/

StabilityAI. (2023). *Stable Diffusion 2.1 Demo*. https://huggingface.co/spaces/stabilityai/stable-diffusion

Stats NZ. (2020). *Algorithm charter for Aotearoa New Zealand*. https://www.data.govt.nz/assets/data-ethics/algorithm/Algorithm-Charter-2020_Final-English-1.pdf

Sterling, R., Riddle, K., Brooks, R., & Hudson, M. (2021). *Understanding Māori rights and interests in intellectual property arising from research and innovation*. Te Kotahi Research Institute, University of Waikato. https://www.sftichallenge.govt.nz/assets/Uploads/Download-PDFs/Understanding-Maori-Rights-and-Interests-in-IP-arising-from-Research-and-Innovation_May-2021-Final.pdf

Taiuru, K. (2020). *Treaty of Waitangi/Te Tiriti and Māori ethics guidelines for: AI, algorithms, data and IOT.* http://www.taiuru.maori.nz/TiritiEthicalGuide/

TELUS International. (n.d.). *Our AI training data platform*. https://www.telusinternational.com/solutions/ai-data-solutions/platform

Waitangi Tribunal. (2011). *Ko aotearoa tēnei: A report into claims concerning New Zealand law and policy affecting Māori culture and identity. WAI 262 Waitangi Tribunal Report*. https://forms.justice.govt.nz/search/Documents/WT/wt_DOC_68356054/KoAotearoaTeneiTT1W.pdf

West, S. M., Whittaker, M., & Crawford, K. (2019). *Discriminating systems: Gender, race and power in AI*. https://ainowinstitute.org/publication/discriminating-systems-gender-race-and-power-in-ai-2

Williams, T. (2023). Artists angry after discovering artworks used to train AI image generators without their consent. *ABC News*. https://www.abc.net.au/news/2023-01-10/artists-protesting-artificial-intelligence-image-generators/101786174

AUTHORS

Jonathan Adams is an education technology adviser at Te Pūkenga's academic development group in Toi Ohomai, with 20 years of IT design and transformation experience in both commercial and education sectors in Aotearoa New Zealand, Australia and Singapore. He has a PGDip, Adult and Continuing Education and Teaching, an MSc and BCom from Murdoch University, and a Graduate Certificate in Migration Law from Victoria University, Australia. Jonathan's interests are in skills acquisition, data sovereignty, AI-supported education and technology innovation.

KatieLee Riddle is an ENRICH Scholar, Ngā Pae o te Māramatanga and SING Alumni, and admitted Barrister and Solicitor of the High Court of New Zealand. She completed a Bachelor of Laws with Honours and a second major in Theatre Studies at the University of Waikato in 2020, and has since specialised in Māori Intellectual Property, Digital Sequence Information, and Indigenous Data Sovereignty at Te Kotahi Research Institute. This has led to her involvement within the United Nations Convention on Biological Diversity Article 8(j) and DSI negotiations, and her work with the Local Contexts Hub.