Deep Learning, Sequencing Technologies & Polygenic Scores: Alzheimer's Disease Risk Prediction and Classification

Jonathon Dilworth
School of Computer Science
University of Manchester, UK
jonathon.dilworth@postgrad.manchester.ac.uk

I. Scope & Analysis of the Problem

Alzheimer's Disease (AD) is a neurodegenerative disorder most prevalent in older adults, affecting approximately one in ten individuals aged sixty-five and over. Whilst only 5% of patients are diagnosed before the age of sixty-five, early-onset AD (EOAD) is much more burdensome and exhibits faster rates of cognitive decline than late-onset AD (LOAD) [1]. Early-onset familial AD (EOFAD), a subvariant of EOAD, is shown to have a strong genetic basis (grounded in known mutations, PSEN1, PSEN2 and APP) and is often characterised as monogenic in its aetiology [2]. However, sporadic variants and LOAD (as shown in Figure 1) are much more challenging to anticipate regarding individual risk due to their polygenic nature and numerous influencing risk factors [2]. Despite this, one significant risk factor universally associated with increased risk of LOAD is the presence of APOE4, an allele of the APOE gene, which can easily be tested for using low-cost micro-array genetic tests (represented in I and II).

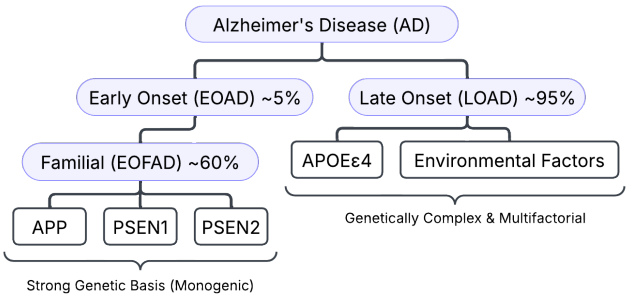

Figure 1. An overview of AD and its variants

Patients presenting with AD-variant symptomology often remain misdiagnosed (in the case of EOAD) or undiagnosed in a timely manner [3], resulting in an estimated additional cost to the NHS of £5,550 (adjusted for inflation) per patient [4]. It should be noted that this figure does not account for additional time spent in A&E (1.5 times the rate of AD-diagnosed patients), additional societal costs (£12,830 per patient) or the effects on the patient's quality of life [4]. Given that the number of AD patients in the United Kingdom is expected to grow by approximately 42% over the next 15 years, finding an effective, high-throughput, low-cost means of measuring patient predisposition to AD could save the NHS up to £2.3 billion. Measuring these potential savings as a function of average annual salary equates to an additional 60,000 NHS employees nationwide, representing a 4% increase in the available workforce. The potential socioeconomic benefits realised from clinically adopting effective methods of AD predisposition motivate our investigation into genomic tooling, data collection methods and predictive analysis techniques.

Table I
Description and Comparison of Genetic Testing

Testing Type	Genes Tested	Description
NGS_{WES}	*APP*, *PSEN1*, *PSEN2*	Targets a known mutation.
SNP_{array}	APOE *e2*, *e3*, *e4*	Identifies *APOE e2*, *e3*, *e4* alleles for AD risk.
PRS_{array}	AP-associated variants	Estimates genetic predisposition using many SNPs.
WES	Entire coding region	Sequences all protein-coding regions.
WGS	Entire genome	Comprehensive sequencing of all base pairs $\in$ genome.

A. Background: Sequencing Technologies & Genome-Wide Association Studies

The traditional approach to identifying SNPs strongly associated with AD has been to analyse data collected through genome-wide association studies (GWAS) [5].

Table II
COMPARISON OF MULTI-ARRAY FOR GWAS AND WGS

Feature	GWAS	WGS
Coverage	Specific SNPs	Whole Genome
Variant Types	Common SNPs	SNPs, SVs, indels
Relative Cost	Low	High

Table III
DATASETS USED IN GENOMIC STUDIES [5]

Year	Dataset	Study	Sample Size
2013	IGAP	Lambert et al.	74,046
2017	IGAP	Sims et al.	85,133
2017	UK Biobank	Lui et al.	116,196
2018	UK Biobank	Marioni et al.	314,278
2019	PGC-ALZ, IGAP, ADSP	Jansen et al.	455,266
2020	UK Biobank	Schwartzentruber et al.	408,942
2021	Multiple Cohorts, GWAX	Wightman et al.	1,126,563

These studies collect DNA samples from cohorts of participants who have been either diagnosed or are presenting with AD. Collected samples are then genotyped (using microarrays) in mass so that common variants belonging to the sampled patients are extracted during the process. It is then possible to compare the frequency of common variants (SNPs) belonging to AD-diagnosed patients with those sampled from a healthy control group, thereby indicating which variants are most strongly associated with the disease. This approach has successfully identified many known variants we associate with AD. However, GWAS using microarrays comes with several limitations.

1) Limitations of Traditional GWAS: Firstly, these studies do not adequately capture rare mutations (unlikely to be included in the microarray design) or structural variants. Secondly, although the sample size of these studies has increased over time (see III), they still lack representative diversity across populations of mixed ethnicity and are typically biased towards European ancestry. In addition, traditional GWAS essentially 'does what it says on the tin', which means it identifies associations, not causations [5]. As such, it can be challenging to decipher the underlying biological mechanisms tied to these associations. It should be noted that this 'explainability gap' occurs prior to any further downstream analysis (which often presents additional challenges concerning explainability). When combined with inadequate downstream analysis (discussed in I-D), these issues culminate in the 'missing heritability' problem, which hypothesises LOADs polygenic aetiology has approximately 58-79% of heritable traits, for which only 3.1% are currently accounted [2].

2) Next Generation Sequencing: Next-generation sequencing (NGS) [6] has gained much traction over the last decade as a tool for AD research and risk prediction. As NGS becomes cheaper to conduct at scale, it provides potential solutions to some of the above-mentioned issues [7] (1). For instance, whole genome sequencing (WGS) and whole exome sequencing (WES) capture much more data for downstream analysis, which may allow for detecting rare and structural variants (e.g. TREM2 [2]). In addition, adopting nanopore technology (e.g. Oxford Nanopore [8]) allows for ultra-long reads, which may also be helpful.

Despite its growing popularity, the possibilities associated with analysing vast amounts of genomic data to assess risk and diagnose polygenic diseases (such as LOAD) have yet to be fully realised in clinical settings.

We postulate that this failure is partly due to the lack of widely available ethnically and ancestrally diverse datasets (though WGS datasets of up to 500,000 individual samples are available to researchers) [9] [2]. However, the most significant contributing component to this adoption failure is more likely the downstream analysis techniques employed in exploring the relevant genomic data. Furthermore, the most widely researched technique (in the last decade), weighted polygenic risk score (PRS), warrants a healthy degree of scepticism insofar as its statistical predictive power is concerned. As such, we explore PRS briefly, highlighting its limitations and potential alternatives in the following subsections.

B. Background: Polygenic Risk Scores

Historically, polygenic risk scores (PRS) have been controversial due to their inconclusive and unconvincing results. It has been widely hypothesised that the poor performance of PRS is likely due to the method's failure to capture complex patterns found within human genome data [10]. Given that classical PRS methodologies are considered statistically insufficient due to their simplicity, fundamentally additive and linear nature (1), modifying or replacing the mechanism for calculating polygenic risk with a model that more reasonably maps non-linearities within the data is an intuitive and reasonable line of enquiry [11]. Interestingly, machine learning-based (ML-based) approaches offer a potential

solution to capturing these complex non-linear relationships within genomic data. Furthermore, a subset of ML, dubbed 'deep learning' (described in I-C), can automatically capture meaningful temporal and structural features, presenting an even more attractive proposition for genomic-based research in this domain [12].

$$PRS_j = \sum_{i=0}^{i=M} x_{i,j}\beta_i \tag{1}$$

C. Background: Deep Learning

While the ideas introduced by McCulloch and Pitts in 1943s [13] and Rosenblatt in 1957 [14] paved the way for connectionist approaches in Artificial Intelligence, the adoption of these computational models was hindered in part due to a large degree of criticism leveraged by Minsky and Papert, describing the endeavour as a "limited" area of research [15]. While their critique was valid at the time of publication, we have seen significant advancements in computational hardware (e.g. GPUs and FPGAs) and the field more broadly. Notable contributions include work on backpropagation in 1986 and applying convolutional neural networks to 'big data' in 2012 [16] and [17].

The first recorded use of the term "deep learning" is attributed to Dechter's 1986 paper on constraint satisfaction in heuristic search [18]. Whilst the coinage is undoubtedly noteworthy, deep learning, as we know and understand it today, typically refers to end-to-end machine learning built on approaches from connectionism, where feature extraction is 'baked in'. Coincidently, another paper published in 1986 by Hinton et al. is a strong candidate to be credited with the birth of deep learning, as it provides the basis for all modern deep learning architectures: backpropagation of error [16].

Over the last decade, we have seen unprecedented growth in 'deep learning' as a research area (an overview is provided in 2). The resulting technologies demonstrate an uncanny ability to accurately detect complex patterns in data, which was previously thought (by some) to be impossible for machines to accomplish. This positions machine learning-based technologies (with a specific focus on deep learning) in an excellent position to be used for downstream analysis techniques in genomics and bioinformatics. The strategy discussion in II provides a brief overview of recent developments in deep learning-based bioinformatics central to our problem statement, summarised in I-D.

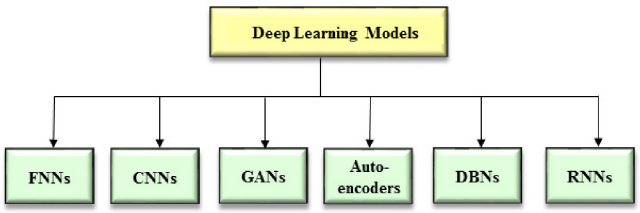

Figure 2. An overview of deep learning, from Lin et al. [19]

D. Problem Statement and Scope Summarisation

Alzheimer's disease continues to pose a significant burden (socioeconomic or otherwise) on healthcare systems (specifically, the NHS), patients and their families. While the overwhelming number of AD instances are categorised as LOAD, more could be done to offset the risk of EOFAD sooner. LOAD-presenting patients is expected to grow significantly over the next 15 years. Thus, a continued effort is required to identify and diagnose the disease as early as possible. These efforts, if realised, will result in considerable benefits to the NHS, numerous patients and their families. We have identified several problems with "traditional GWAS" and "weighted PRS" analysis techniques. However, broader adoption of NGS will undoubtedly occur as the cost associated with WGS, WES, and nanopore technologies continues to decrease. As such, we expect more AD-associated variants (SNPs, SVs, RVs) to be discovered. With this fundamental assumption stated, we identify deep learning-based technologies as a potential solution to the inadequacies of existing PRS methods for predicting the risk of AD.

We also emphasise that traditional GWAS, whilst imperfect, could still benefit from adopting these technologies downstream. As such, we also encourage continued research on existing multi-array SNP datasets, WGS and WES datasets.

In the next section, we discuss several potential strategies (within the context of existing bioinformatics research) for integrating DL-based methods into a risk assessment and diagnostic framework for clinical trials or future adoption.

Note: An overview of our problem space is illustrated in Figure 3.

II. Strategies To Address The Problem

In 2020, Koumakis posed the following question: "Deep learning models in genomics; are we there yet?" [20]. They conclude that: "DL models can provide higher accuracies in specific tasks of genomics than the current state-of-the-art methodologies" [20], which provides

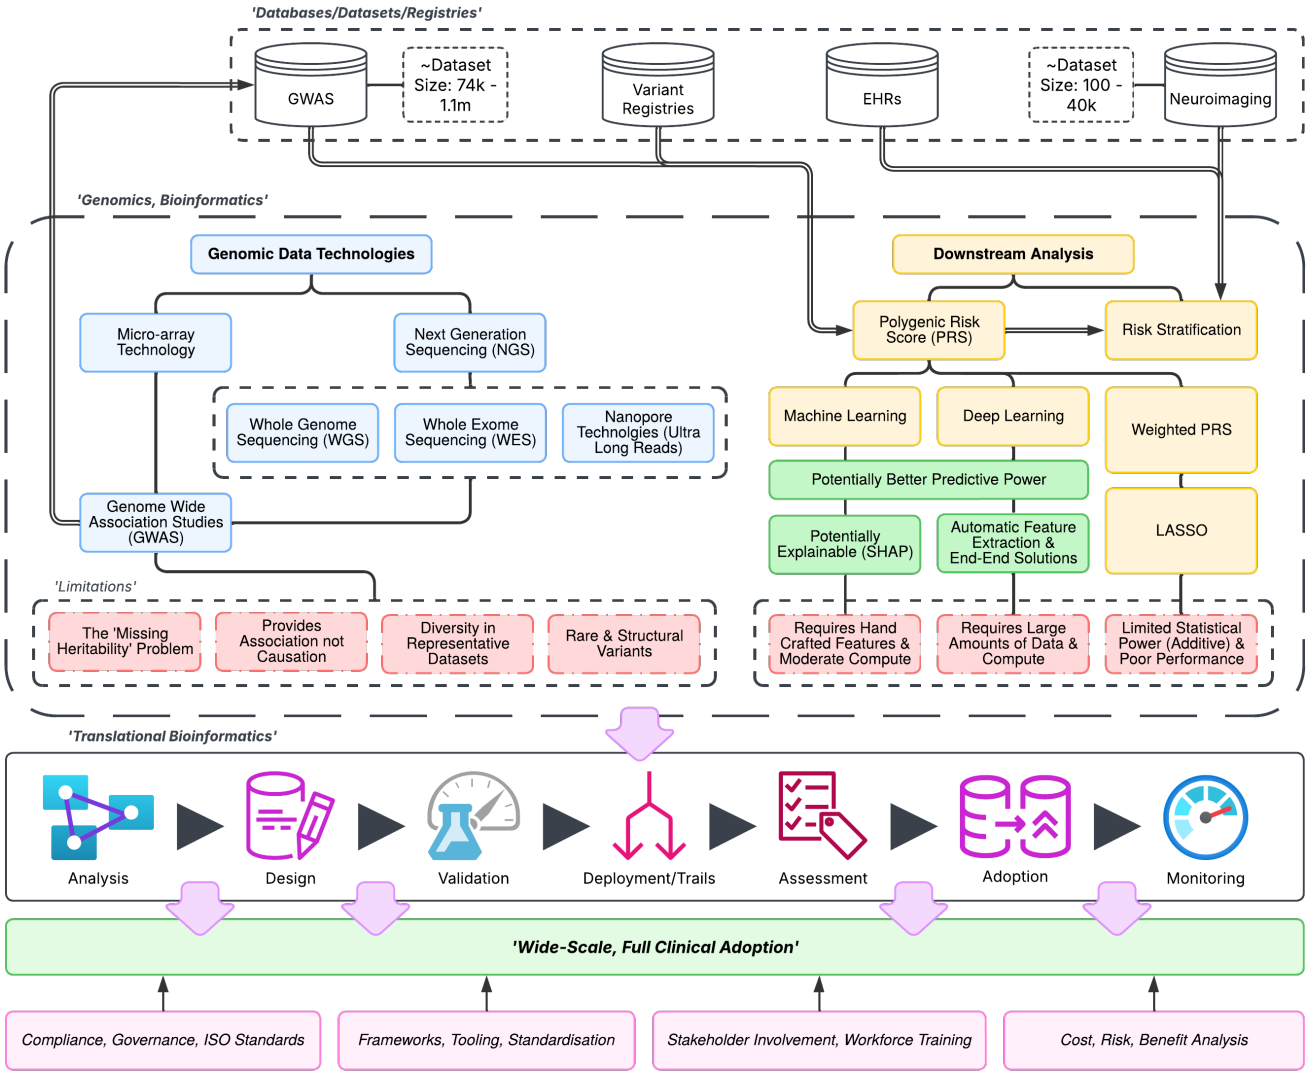

Figure 3. An overview of the problem space

sufficient motivation to address the problems associated with predicting risk (PRS) and classifying AD (as discussed in I-B and I-D). As such, we deem it appropriate to revisit this question five years later to re-evaluate deep learning-based strategies and inform our final recommendation.

A. Related Strategies

Song et al. [21] introduce GenBERT, a DL-based architecture built on stacked encoders, tuned using an evolutionary algorithm. The research combines GenBERT with a multi-input convolutional neural network (MI-CNN) to classify AD-positive patients as early as possible. While their motivation aligns with our own, and the approach boasts an impressive accuracy (98.3%) and F1 score (97%), the necessary data required to train (and perform inference on) their model includes clinical notes, genetic profiles and neuroimaging data. Given that the type, quality, and quantity of data required to utilise their approach are unlikely to be unavailable before disease onset, their model is only viable for diagnosis after a clinician has recognised AD symptomology and the necessary brain scans have been performed. In addition, the associated costs of adopting this approach (required staff, hardware, software and compliance adherence, [22]) may outweigh the potential benefits, as it cannot be described as low-cost or high-throughput. In addition, we recognise that clinician notes could introduce implicit

bias into the system. We raise concerns over LLM-based architectures processing human written notes, as bias has been a significant issue in many healthcare-augmented AI deployments [23] [24]. Finally, we note that including an ablation study would have significantly improved the assessability of each model's contribution (in isolation) to the final patient classification (AD-positive, AD-negative). Without assessing the capability of the genomic-based model in isolation, it only aids our ability to consider such models for risk stratification.

Li et al. [5] comprehensively review 'brain imaging biomarker genomics in Alzheimer's disease' for precision medicine. While not immediately relevant, they discuss the dataset sizes available for neuroimaging data accompanied by genomics data for use in AD-related data processing. As previously suspected, potential datasets for use in this case are, on average, approximately ten times smaller than purely genomic ones yielded through GWAS. As such, we further outline our concern for using neuroimaging data in 'fusion' methods [19] as DL-based approaches typically generalise better when trained on much larger datasets.

Alatrany et al. [25] use random forests to extract candidate genetic correlates from GWAS data as features for a downstream convolutional neural network (CNN) and multilayered perceptron (MLP). Their models (CNN & MLP) achieve an AUC of 93% and 90%, respectively, indicating that simple non-linear models combined thoughtfully can perform well.

While deep learning-based approaches appear to provide an effective means of assessing risk and diagnosing AD genetically [12], we have already raised some concerns regarding model integrity. As such, we consider the significant effort in developing explainable AI (XAI) tooling, as discussed by Khater et al. [26]. Many of these explainable approaches utilise traditional ML, such as support vector machines (SVMs), as they offer better explainability than complex DL-based architectures (such as transformer-based methods [27]). Many approaches also use SHAP values to support model interpretability better.

Khater et al. [26] describe a DL-based approach to train a base model on genomics data and apply transfer learning to increase the model's robustness. Initial candidate genetic correlates and biomarkers are identified and extracted as features for a downstream SVM. Whilst their description does not match the details in their cited paper, it provides a fascinating insight into one potentially interesting approach. We cite this (with the knowledge that some transcriptional error has occurred in their write-up) as it exemplifies an approach that captures the best of both worlds: using deep learning's automatic feature extraction capabilities and providing increased control and manual oversight in the validation of those features whilst ensuring some degree of explainability.

A promising research area (which seems to excel in explainability) links pangenome graphs with machine learning architectures to improve AD risk prediction and diagnostic precision [28]. However, the computational requirements for analysis are considerable [29], and this research is incredibly scarce. We only found two instances of this particular methodology applied to similar bioinformatics problems [28] [30]. While this research poses an exciting prospect, it is likely too immature, and further evidence is required before any further review can be considered.

B. Pre-Recommendation Conclusions

Interestingly, Koumakis' [20] conclusions corroborate our review-based findings under II-A. However, many of the same hurdles are still to be overcome, most notably regulatory issues (data sharing, model robustness) and "the process of translating the knowledge acquired in genomics research into clinically useful tools" [20], which has proven extremely slow. Nevertheless, progress has been made within this research area, and we use this progress to inform our strategy recommendation as provided under §3.

III. Strategy Reccomendation

In our efforts to reach a strategy recommendation, we consider previously discussed bioinformatics research for evaluating patient risk of EOAD and LOAD (II). In addition, we note the largely unrealised conceivable benefits discussed in I. However, we also weigh the potential benefits against their associated risks. Careful consideration for the NHS framework for advanced clinical practice [31] and the inherent complexities tied to compliance (especially those associated with AI systems) are also granted [32] [33] [22]. We also cite issues inherent to genomics research (limitations of upstream GWAS), potential psychological effects of disclosing genetic predisposition to patients prematurely (with imperfect confidence values) and large-scale lack of standardisation as highly problematic (at present) [20]. We reach the following conclusions, which are provided below as a basis for our recommendation.

- Utilising a low-cost, high-throughput means of evaluating genetic predisposition to EOFAD continues

to be viable. It does not require the use of computationally complex analysis techniques.
- Existing standardised genetic tests for measuring APOE and its variant biomarkers provide a solid foundation for existing risk stratification (and for diagnostic aids) in assessing for LOAD.
- Unmodified weighted PRS is a statistically inadequate methodology for assessing patient risk of AD and is not viable for clinical adoption.
- Genomics data and WGS-based datasets will continue to grow over the next fifteen years, and genomics research will not slow down anytime soon.

We are in favour of continuing to maintain existing clinical risk assessment and diagnostic pipelines. In addition, we understand that the reviewed research (under II-A) would significantly benefit from a standardisation framework and a greater degree of diversity in representative datasets. However, we consider ML and DL-based methods (for assessing PRS applied to enhanced risk stratification) as acceptably mature candidates for further analysis and validation. This will be the first tentative step in a multi-phased approach towards full clinical adoption over the next ten years.

We make this recommendation cautiously and emphasise that this will be a long-term strategy. In essence, we anticipate an acceleration in genomic data production, progress in related fields (AI, ML, DL) and their applications to genomics research. Furthermore, we consider the regulatory frameworks [32] [33] [22] for managing data and AI deployments (in Europe and the United Kingdom) to be advantageous in protecting against possible collateral damage (to patients and the NHS alike).

Drawing from Alatrany et al. [25], Zhou et al. [12] and Khater et al. [26], we propose a framework for clinical adoption that aims to accomplish the goals outlined in the subsequent subsections.

A. System Overview & Timeline

The proposed framework includes a timeline consisting of multiple phases. Firstly, an initial period of further analysis and validation of current state-of-the-art ML/DL-based techniques is to be conducted. Subsequently, development and deployment will take place (for an initial trail), targeting existing high-risk AD patients to validate piloted systems. This will be followed by an observation and evaluation period. Depending on the pilot programme's results, a concerted effort may or may not be made to scale up the programme, as shown in Figure 4.

5 illustrates the framework's integration into existing clinical procedures related to established NHS genomic labs while introducing ML/DL-based pipelines containing Random Forests for feature selection and MLP, SVM, and CNN models for classification, as discussed in II-A.

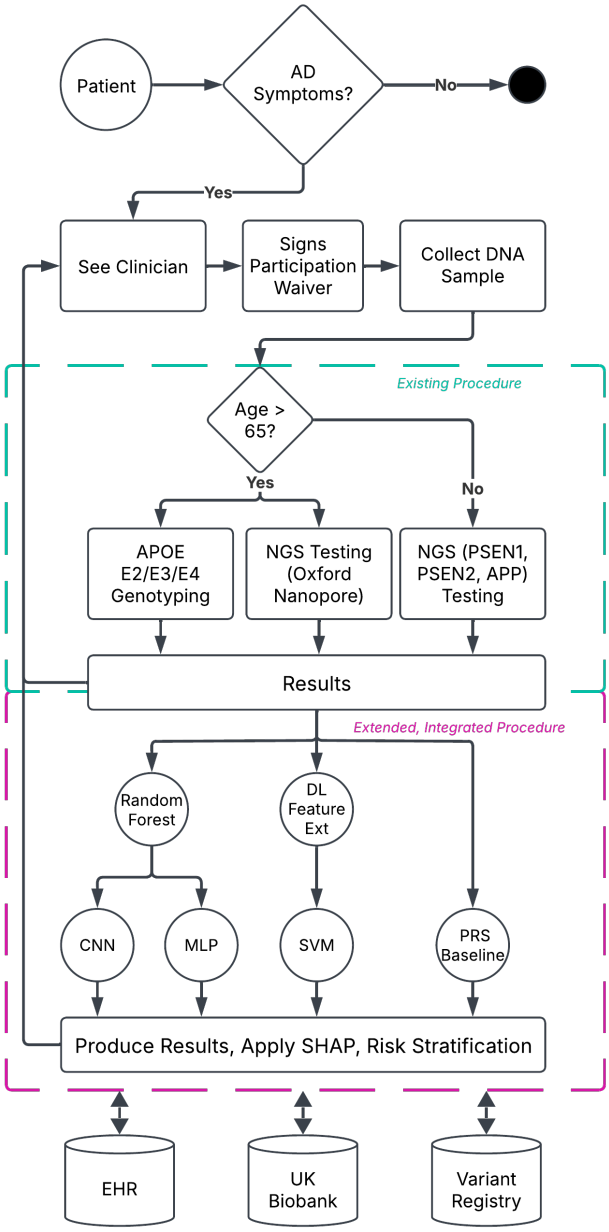

Figure 5. A proposed risk management and diagnostic pipeline integration

B. Maximising Diversity in Representative Datasets

Combining multiple cohorts and existing GWAS datasets for base model pre-training ensures that the most com-

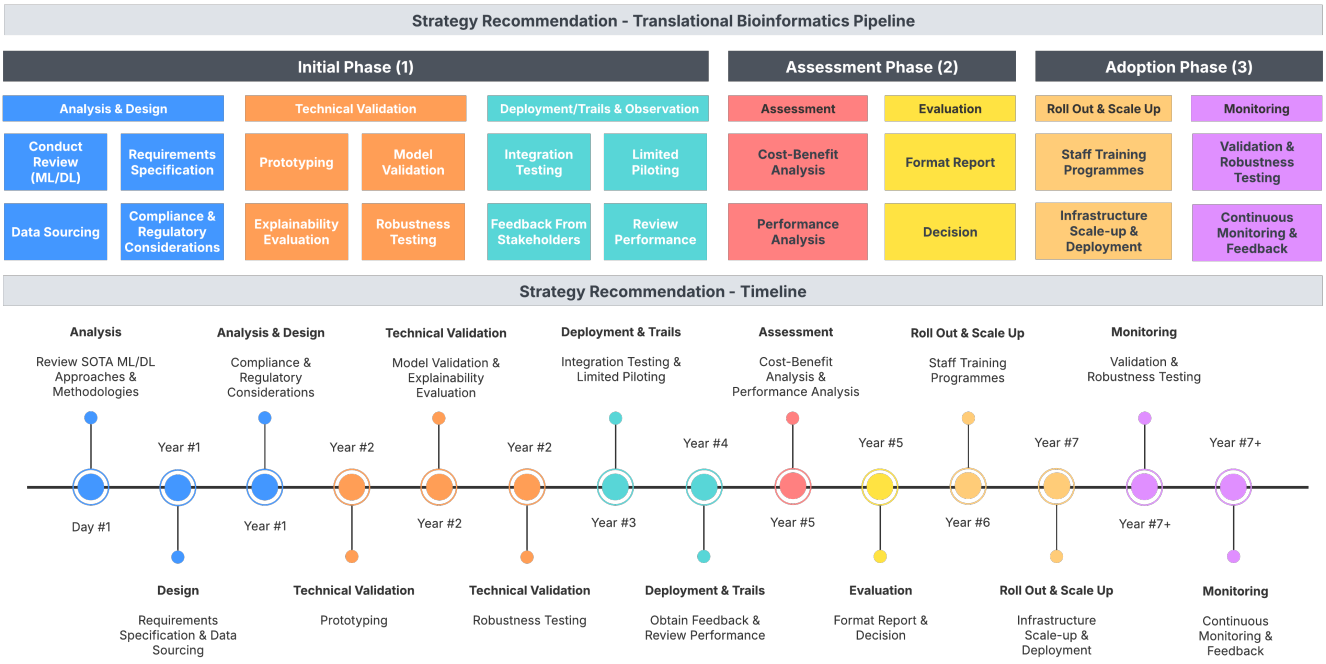

Figure 4. Proposed translational bioinformatics pipeline for DL/ML AD risk stratification pipeline

prehensive genetic factors influencing AD onset are captured during this stage. In addition, we suggest curating datasets that include individuals from varied ethnic and socioeconomic backgrounds for ongoing research efforts. However, during the initial piloting of clinical systems, we suggest targeting smaller, already suspected cohorts at high risk of AD to validate the model pre-training procedure during this phase.

C. *Emphasising the Need for Quality Control*

Quality control in model development and deployment is essential for model robustness [32], validation testing, and explainability [22], which are especially important in any clinical setting. For instance, the ISO 24029 standards for AI/ML testing (related explicitly to robustness) describe procedures undertaken during design and development, verification and validation, deployment and operation, and monitoring. We strongly advise adopting these procedures to ensure AI models perform reliably across the above-mentioned datasets. Furthermore, we draw from II-A in suggesting the avoidance of raw human-produced clinical notes as input for risk stratification (in favour of using 'MCQ-type' notes) to remove as much potential bias as possible. XAI tooling (such as SHAP and explainable feature selection) is also integrated into the risk assessment and diagnostic pipeline.

D. *Seamless Integration with Existing Methods for Risk Stratification*

Integrating novel ML/DL-based methods with existing risk stratification and diagnostic tooling allows for easy piloting and minimal disruption to clinical workflows. By building on existing procedures (conventional genetic testing and other piloted PRS techniques), clinicians can interpret results comfortably and familiarly. Integrating systems in this manner should allow for manageable adoption during initial trials while leveraging current patient data for system validation during the piloting period.

E. *Offsetting Human Over Reliance*

We briefly note that AI systems often become a source of over-reliance and, as such, propose that any reports generated from such systems be highlighted as ML/DL-sourced findings (using a bright background colour).

F. *Expected Outcomes*

- Improved patient care and well-being.
- Reduced stress on the healthcare system (NHS)
- Reduced cost (if successful) over the next fifteen years

G. Conclusion

Our recommendation emphasises the importance of recognising NGS, ML, and DL as promising tools for accurate AD-risk prediction by discussing the potential for clinical adoption in terms of related strategies presented under II-A. We outline a clinical framework for adoption to provide clinicians with improved AD-risk stratification and diagnostic aids. We highlight the need for diverse representation during pre-training and robust quality control, presenting a model integration strategy that follows existing regulatory frameworks and NHS guidance. Ultimately, this recommendation aims to improve patient care, relieve pressure on the national healthcare system and reduce healthcare costs over the long term.

REFERENCES

[1] W. M. v. d. Flier, Y. A. Pijnenburg, N. C. Fox, and P. Scheltens, "Early-onset versus late-onset Alzheimer's disease: the case of the missing APOE ε4 allele," *The Lancet Neurology*, vol. 10, no. 3, pp. 280–288, Mar. 2011, publisher: Elsevier. [Online]. Available: https://www.thelancet.com/journals/lancet/article/PIIS1474-4422(10)70306-9/abstract

[2] N. Karagas, J. E. Young, E. E. Blue, and S. Jayadev, "The Spectrum of Genetic Risk in Alzheimer Disease," *Neurology Genetics*, vol. 11, no. 1, p. e200224, Feb. 2025, publisher: Wolters Kluwer. [Online]. Available: https://www.neurology.org/doi/full/10.1212/NXG.0000000000200224

[3] "Early intervention in Alzheimer's disease: a health economic study of the effects of diagnostic timing | BMC Neurology." [Online]. Available: https://link.springer.com/article/10.1186/1471-2377-14-101

[4] D. Getsios, S. Blume, K. J. Ishak, G. Maclaine, and L. Hernández, "An economic evaluation of early assessment for Alzheimer's disease in the United Kingdom," *Alzheimer's & Dementia*, vol. 8, no. 1, pp. 22–30, Jan. 2012. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S1552526010021874

[5] L. Li, X. Yu, C. Sheng, X. Jiang, Q. Zhang, Y. Han, and J. Jiang, "A review of brain imaging biomarker genomics in Alzheimer's disease: implementation and perspectives," *Transl Neurodegener*, vol. 11, no. 1, p. 42, Sep. 2022. [Online]. Available: https://doi.org/10.1186/s40035-022-00315-z

[6] W. R. McCombie, J. D. McPherson, and E. R. Mardis, "Next-Generation Sequencing Technologies," *Cold Spring Harbor Perspectives in Medicine*, vol. 9, no. 11, p. a036798, Nov. 2019. [Online]. Available: https://pmc.ncbi.nlm.nih.gov/articles/PMC6824406/

[7] M. Marklewitz, A. Jaguparov, A. Wilhelm, O. W. Akande, B. Musul, A. L. Poates, B. Afrough, A. Norberg, N. C. Hull, S. Ehsani, G. members of GCT pilot working group, J. Salvi Le Garrec, and T. Whistler, "Genomics costing tool: considerations for improving cost-efficiencies through cross scenario comparison," *Front. Public Health*, vol. 12, Jan. 2025, publisher: Frontiers. [Online]. Available: https://www.frontiersin.org/journals/public-health/articles/10.3389/fpubh.2024.1498094/full

[8] M. Jain, S. Koren, K. H. Miga, J. Quick, A. C. Rand, T. A. Sasani, J. R. Tyson, A. D. Beggs, A. T. Dilthey, I. T. Fiddes, S. Malla, H. Marriott, T. Nieto, J. O'Grady, H. E. Olsen, B. S. Pedersen, A. Rhie, H. Richardson, A. R. Quinlan, T. P. Snutch, L. Tee, B. Paten, A. M. Phillippy, J. T. Simpson, N. J. Loman, and M. Loose, "Nanopore sequencing and assembly of a human genome with ultra-long reads," *Nat Biotechnol*, vol. 36, no. 4, pp. 338–345, 2018. [Online]. Available: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC5889714/

[9] "UK biobank: Whole genome sequencing data," Oct. 2018. [Online]. Available: https://www.ukbiobank.ac.uk/enable-your-research/about-our-data/genetic-data

[10] N. B. Gunter, R. K. Gebre, J. Graff-Radford, M. G. Heckman, C. R. Jack, V. J. Lowe, D. S. Knopman, R. C. Petersen, O. A. Ross, P. Vemuri, V. K. Ramanan, and for the Alzheimer's Disease Neuroimaging Initiative, "Machine Learning Models of Polygenic Risk for Enhanced Prediction of Alzheimer Disease Endophenotypes," *Neurology Genetics*, vol. 10, no. 1, p. e200120, Feb. 2024, publisher: Wolters Kluwer. [Online]. Available: https://www.neurology.org/doi/10.1212/NXG.0000000000200120

[11] "PRS Basics - GWASTutorial." [Online]. Available: https://cloufield.github.io/GWASTutorial/10_PRS/

[12] X. Zhou, Y. Chen, F. C. F. Ip, Y. Jiang, H. Cao, G. Lv, H. Zhong, J. Chen, T. Ye, Y. Chen, Y. Zhang, S. Ma, R. M. N. Lo, E. P. S. Tong, V. C. T. Mok, T. C. Y. Kwok, Q. Guo, K. Y. Mok, M. Shoai, J. Hardy, L. Chen, A. K. Y. Fu, and N. Y. Ip, "Deep learning-based polygenic risk analysis for Alzheimer's disease prediction," *Commun Med*, vol. 3, no. 1, pp. 1–20, Apr. 2023, publisher: Nature Publishing Group. [Online]. Available: https://www.nature.com/articles/s43856-023-00269-x

[13] W. S. McCulloch and W. Pitts, "A logical calculus of the ideas immanent in nervous activity," *Bulletin of Mathematical Biophysics*, vol. 5, no. 4, pp. 115–133, Dec. 1943. [Online]. Available: https://doi.org/10.1007/BF02478259

[14] F. Rosenblatt, *The perceptron, a perceiving and recognizing automaton Project Para*. Cornell Aeronautical Laboratory, 1957.

[15] M. Marvin and A. P. Seymour, "Perceptrons," *Cambridge, MA: MIT Press*, vol. 6, no. 318-362, p. 7, 1969.

[16] D. E. Rumelhart, G. E. Hinton, and R. J. Williams, "Learning representations by back-propagating errors," *nature*, vol. 323, no. 6088, pp. 533–536, 1986, publisher: Nature Publishing Group UK London.

[17] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet Classification with Deep Convolutional Neural Networks," in *Advances in Neural Information Processing Systems*, F. Pereira, C. J. Burges, L. Bottou, and K. Q. Weinberger, Eds., vol. 25. Curran Associates, Inc., 2012. [Online]. Available: https://proceedings.neurips.cc/paper_files/paper/2012/file/c399862d3b9d6b76c8436e924a68c45b-Paper.pdf

[18] R. Dechter, "Learning while searching in constraint-satisfaction problems," 1986, publisher: University of California, Computer Science Department, Cognitive Systems ....

[19] E. Lin, C.-H. Lin, and H.-Y. Lane, "Deep Learning with Neuroimaging and Genomics in Alzheimer's Disease," *International Journal of Molecular Sciences*, vol. 22, no. 15, p. 7911, Jan. 2021, number: 15 Publisher: Multidisciplinary Digital Publishing Institute. [Online]. Available: https://www.mdpi.com/1422-0067/22/15/7911

[20] L. Koumakis, "Deep learning models in genomics; are we there yet?" *Computational and Structural Biotechnology Journal*,

vol. 18, pp. 1466–1473, Jan. 2020. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2001037020303068

[21] J. Song, J. Huang, and R. Liu, "Integrating NLP and LLMs to discover biomarkers and mechanisms in Alzheimer's disease," *SLAS Technology*, vol. 31, p. 100257, Apr. 2025. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S2472630325000159

[22] "ISO 23053: Framework for Artificial Intelligence (AI) Systems Using Machine Learning (ML)," International Organization for Standardization, Geneva, CH, Standard, Jul. 2023, volume: 2023.

[23] "Incident 124: Algorithmic Health Risk Scores Underestimated Black Patients' Needs." [Online]. Available: https://incidentdatabase.ai/cite/124/

[24] "Incident 603: Algorithmic Allocation of Resources in Healthcare for Disabled and Elderly Care Services Allegedly Harming Patients." [Online]. Available: https://incidentdatabase.ai/cite/603/

[25] A. Alatrany, A. Hussain, J. Mustafina, and D. Al-Jumeily, "A Novel Hybrid Machine Learning Approach Using Deep Learning for the Prediction of Alzheimer Disease Using Genome Data," in *Intelligent Computing Theories and Application*, D.-S. Huang, K.-H. Jo, J. Li, V. Gribova, and P. Premaratne, Eds. Cham: Springer International Publishing, 2021, pp. 253–266.

[26] "Explainable Machine Learning Model for Alzheimer Detection Using Genetic Data: A Genome-Wide Association Study Approach | IEEE Journals & Magazine | IEEE Xplore." [Online]. Available: https://ieeexplore.ieee.org/abstract/document/10549918

[27] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[28] C. P. Prathibhamol, J. Chandrakiran, S. Santhosh, M. Meenakshi, S. A. Menon, and M. Nair, "A pangenome graph-based approach for predicting Alzheimer's disease," in *Data Science & Exploration in Artificial Intelligence*. CRC Press, 2025, num Pages: 6.

[29] S. Hwang, N. K. Brown, O. Y. Ahmed, K. M. Jenike, S. Kovaka, M. C. Schatz, and B. Langmead, "MEM-based pangenome indexing for k-mer queries," Nov. 2024, iSSN: 2693-5015. [Online]. Available: https://www.researchsquare.com/article/rs-5363291/v1

[30] Z. Miao and J.-X. Yue, "Interactive visualization and interpretation of pangenome graphs by linear reference–based coordinate projection and annotation integration," *Genome Research*, vol. 35, no. 2, pp. 296–310, 2025.

[31] "Multi-professional framework for advanced practice in England." [Online]. Available: https://advanced-practice.hee.nhs.uk/multi-professional-framework-for-advanced-practice/

[32] "Robustness for Neural Networks – ISO/IEC 24029-1:2021 Introduction." [Online]. Available: https://www.bsigroup.com/en-GB/training-courses/isoiec-24029-12021-how-to-assess-robustness-of-neural-networks/

[33] "BSOL BS Online ISO 23894." [Online]. Available: https://bsol-bsigroup-com.manchester.idm.oclc.org/Bibliographic/BibliographicInfoData/000000000030479001