



Medicines & Healthcare products Regulatory Agency

Board Meeting

22 October 2018

<p>Title: MHRA role in development and use of artificial intelligence for safety monitoring</p>
<p>Issue/ Purpose: MHRA has been encouraged to put forward bids to several government programmes for funding/ external resources to explore the utility of artificial intelligence (AI) to deliver enhanced vigilance capability. The request has led the Divisions to more broadly consider areas where exploration of AI may be of interest to ensure a joined-up approach, as well as the associated risks and issues, which are considered in this paper.</p>
<p>Summary: John Wilkinson has previously presented to CET setting out the strategic challenge presented by AI, considering the regulatory impact and highlighting the interest from DHSC. The CET were requested to consider the impact of AI on their areas.</p> <p>In response to this this paper considers the following topics as relevant when considering the utility and impact of AI:</p> <ul style="list-style-type: none"> • Regulation of software that applies AI • Analysis of data lakes to identify/ validate safety issues using AI • Use of AI to prompt/ facilitate incident reporting • Automation of aspects of internal incident processing using AI <p>This paper focuses primarily on the second theme, but in doing so raises important considerations for the other areas for discussion and potential funding opportunities.</p>
<p>Suitable for team briefing? (Yes/No) Yes</p>
<p>Summary for team briefing (if applicable) <u>As Above</u></p>
<p>Resource implications: To be discussed</p>
<p>EU Referendum implications: <u>None</u></p>

Timings:

TBC

Action required by the Board:

The Board are asked to:

- Validate that the Agency should progress the activities proposed in the paper and manage as a programme of work with several separate projects
- Consider impact on MHRA resources bearing in mind current constraints
- Agree the governance structure for any proposed activities

Links:

Vigilance and Risk Management of Medicines (VRMM) Division, Devices Division, Clinical Practice Research DataLink (CPRD), Inspection, Enforcement and Standards (IE&S) Division.

Are there any sensitivity issues that would prevent this item being discussed by CET in the presence of staff observers?

No

Which of the five themes in the Corporate Plan 2018-2023 does the paper support?

1. Public health and partnerships
2. Enhancing innovation
3. Proactive, robust surveillance

If relevant, which Business Plan strategic activity does it support?**CET Sponsor:** Dr June Raine, John Wilkinson

MHRA role in development and use of artificial intelligence for safety monitoring

1. Issue

MHRA has been encouraged to put forward bids to several government programmes for funding/ external resources to explore the utility of artificial intelligence (AI) to deliver enhanced vigilance capability. The request has led the Divisions to more broadly consider areas where exploration of AI may be of interest to ensure a joined-up approach, as well as the associated risks and issues, which are considered in this paper.

2. Background

This paper formed a 'state of the nation' summary of activities to date in this area, including:

- Differences in regulatory frameworks between medicines and devices
- Value of spontaneous reporting systems
- Previous research conducted against electronic healthcare record data
- Current use of CPRD and future opportunities for additional linkages
- Opportunities arising from Unique Device Identifier (UDI) and GS1 standards and Scan for Safety

3. MHRA interest in AI

John Wilkinson presented a paper (Annex 2) to CET in June 2018 setting out the strategic challenge presented by AI, considering the regulatory impact and highlighting the interest from DHSC. The Corporate Executive Team (CET) were requested to consider the impact of AI on their areas.

In response to this we have considered the following topics as relevant when considering the utility and impact of AI:

- Regulation of software that applies AI
- Analysis of data lakes to identify/ validate safety issues using AI
- Use of AI to prompt/ facilitate incident reporting
- Automation of aspects of internal incident processing using AI

Analysis of data lakes to more rapidly identify safety signals for medicines and devices was what we had understood the initial focus of the NHS England and NHS Digital interest in our work to be. Extensive research has already been completed in the medicines space to understand how electronic healthcare record databases can be used in pharmacovigilance. This has led to recommendations for further research to understand further if they could be used within automated signal detection but it is clear that even based on current understanding there is a role for such data in supporting signal detection and complementing existing spontaneous reporting schemes.

However, the rapid evolution of natural language processing (NLP) and AI will require continued research into the future to optimise the use of these data sources. There

may also be additional opportunities in the Devices area, where electronic health records are currently considered to lack high quality data (structured or otherwise). Soon initiatives such as UDI, GS1 and Scan for Safety will also significantly enhance opportunities in this area.

It is also important to be aware of the impending implementation of AI for pharmacovigilance data management by Industry in the coming 1-2 years. Whilst we have no specific knowledge of how this will affect the incoming data from Industry, it is important that MHRA can develop knowledge and assess the impact, whilst still supervising this source of our data. MHRA will also need to provide guidance and interpretation of existing legislation concerning data management to Industry stakeholders.

It is anticipated that spontaneous reporting will continue to form an important part of medicines and devices safety monitoring irrespective of progress with AI analytics across large data lakes. Enhancement of safety messaging capability is also an essential focus of the Agency and the wider health sector to ensure meaningful information is presented at the point of care. To that end we consider that it is important that we utilise the current interest in safety monitoring to progress these goals, but also to consider whether AI can be used to better focus provision of reporting forms and safety messaging to healthcare professionals than current rule-based approaches.

4. Definition of AI

The Royal Society report on Machine learning published in April 2017¹ defines artificial intelligence as '*the science of making machines smart*' and machine learning as the '*technology that allows computers to perform specific tasks intelligently, by learning from examples*'. Thus, artificial intelligence could be viewed as an aspiration whereby machines and automation achieve a level of intelligence that is usually associated with 'natural intelligence' possessed by humans and subsequently, surpass this. Artificial intelligence offers the promise of efficiency hand in hand with improved effectiveness. In the context of vigilance, artificial intelligence could enable improved case processing, even earlier signal detection, or increase the confidence in the safety signals detected, such that earlier intervention can be implemented.

The means of developing artificial intelligence, including machine learning approaches, is in itself, a broad area and focuses on developing prediction algorithms by 'learning from experience' via data inputs. Even classical statistical techniques are constructed on this underlying principle of learning how to predict an outcome based on a set of data inputs. Typically, classical statistical techniques are highly supervised by humans who define a set of starting assumptions based on previous evidence or expert input. Machine learning is a branch of artificial intelligence that allows computer systems to learn directly from examples, data, and experience. Through enabling computers to perform specific tasks intelligently, machine learning systems can carry out complex processes by learning from data, rather than following pre-programmed rules. Some of the more recent machine learning techniques like artificial neural networks (ANN), still require some upfront

¹ <https://royalsociety.org/~media/policy/projects/machine-learning/publications/machine-learning-report.pdf> (Accessed on: 22/06/2018)

programming by humans but have the capability to run iteratively whereby the initial set of outputs are used as inputs in the next 'run cycle' of the algorithm. A successful algorithm is generally defined on the basis of predictive accuracy rather than 'how' the prediction is achieved.

A related concept that is pertinent to signal detection for vigilance is 'data mining', which some scientists consider as a subset of machine learning. At a series of expert stakeholder workshops organised by CPRD in 2017 (including technical experts from academia, regulatory bodies and industry), data mining was described as '*humans using algorithms or visualisation techniques to discover new knowledge*'. Data mining is sometimes referred to as 'unsupervised learning' as there is no starting hypothesis. Human interpretation is very important to make sense of the patterns discovered in the data and to translate the findings into policy recommendations. As the evidence base develops, a 'semi-supervised' approach to data mining could be adopted.

5. Data sources

This section provides an overview of UK data sources that can be used for vigilance using a range of methods including machine learning/data mining.

Case reports

Individual reports coming from patients, healthcare professionals, and industry form the basis of any medicines or medical devices vigilance system. They have an established value in the detection of adverse events with the data being considered particularly valuable for identifying rare or unusual events and those of long latency, which clinical trials are not powered to detect. Spontaneous reporting systems are less helpful in identifying events which may be confounded by the disease or patient history. Signal detection algorithms can be run on such data to mine for safety signals.

ICSR source data could be transformed using machine learning; downstream activity such as signal notifications and PSUR production will be affected. It is important to consider the impact of this compared with existing source-to report accuracy.

Electronic healthcare records (EHR)

These are routinely collected data to facilitate patient care and payments in primary and secondary care. CPRD collects primary care data from practices using Vision EMIS GP software systems across the UK and provides monthly database updates to VRMM for vigilance purposes. However, as CPRD receives daily data flows from contributing practices, the possibility of more frequent updates (to facilitate near-real time analysis in response a critical public health need) are possible for e.g. weekly or even daily updates subject to the necessary investment for accelerated data processing (to convert the raw data into a research/surveillance suitable database that meets data governance and quality requirements). The primary care data captures information on all primary care prescribing, immunisations administered in primary care and outcomes that would typically necessitate a GP consultation. Referrals from primary care to other healthcare services and high-level information on hospital care would also be recorded. These data are used extensively within

pharmacovigilance in the MHRA. The utility of these data for devices monitoring is currently being explored as part of the Agency Patient Safety and Vigilance Strategy (PSVS).

CPRD also operates a linkage programme to hospital records, maintained by NHS Digital, covering admitted patient care, outpatient appointments, accident and emergency (A&E) attendances, and maternity data which could be used to ascertain safety-related outcomes. Hospital prescribing data at an individual patient level are not routinely captured by hospital EHR systems at present and therefore, these data cannot be used for signal detection for medicines prescribed in secondary care. Hospital data could be used for devices monitoring as procedures are well recorded though specific data on manufacturers and materials would not be available.

NHS Business Services Authority capture England-wide data on prescriptions for medicines dispensed in the community. These data have been linked to the hospital record data described above on a national level basis. Similar data are also held by the devolved administrations. The limited data in this larger linked data pool, which captures only prescriptions from primary care and events presenting or diagnosed within a hospital setting or procedures undertaken there, limits its likely value for vigilance purposes.

Registries

Registries can be established to capture data on people prescribed a particular medicine or who use a specific medical device or on all people with a certain condition or undergoing a type of procedure. CPRD routinely links to Cancer Registry data including data on systemic anticancer therapy. Bespoke linkage is possible to other registry data such as the National Joints Register (which has been previously used to investigate safety of hip replacement implants) to enable data capture across the complete patient pathway. Registries are valuable in both medicines and devices vigilance. They are operated by both industry and academic groups.

Other potential future data sources

Increases in the availability of genomic and proteomic data have also led to increasing research into the use of machine learning and data mining for identifying risk factors for certain adverse events seen with specific drugs. While these data are not widely used in pharmacovigilance at present this will likely increase in the future as it is important for facilitating the move towards precision medicine.

Key Issues with data sources

- Quality

Data quality is a key consideration for vigilance activities irrespective of the methods used for signal detection. Special caution needs to be exercised when using data collected primarily for healthcare management purposes for vigilance activities, and more so, when using machine learning. Routine data sources can be biased i.e. not representative of the general population because key sub-groups are either under-represented or over-represented in the data. As outlined in previous sections, key information may be missing and while statistical methods to 'impute' missing data are

available, any inferences based on such imputed data would be less certain. Moreover, even imputation techniques need some proportion of complete data and would be useless where the data on specific aspects are not collected at all. In some instances, the information may be available as free text clinical notes but this may be off-limits for vigilance because of governance concerns around patient confidentiality; CPRD does not collect any free text data on instruction from the Information Commissioner's Office (ICO).

A significant consideration when assessing data quality is the phenomenon of data drift which leads to a change in the underlying data structure or meaning over time. This means that any signal detection or validation activities would need to consider generalisability not just across different population groups but also, generalisability over time.

- Signal generation vs. confirmation

Consideration also needs to be given to the reuse of data sources to evaluate safety signals that were detected in the same data. Hypothesis testing studies may well be required following identification of a safety signal through data mining and data pools as described previously are a vital source of data for the conduct of such studies. IMI PROTECT recommended that more research is required to understand how, following detection of a safety signal in an electronic healthcare database or registry for example the same data can be then used to confirm it as a true risk and provide robust evidence for regulatory and clinical decision making. While some examples of this now exist and the evidence base on how to avoid false detection is steadily increasing, a cautious approach remains advisable and further research is recommended.

- Absence of Device data from existing data sources

Medical devices present unique challenges because of the number and range of medical devices on the market for use in health care settings and in patients. The use of Unique Device Indicators (UDI) provides an opportunity to track devices through all aspects of the device supply chain and to provide opportunities for data linkages via registries and electronic patient records which has not previously been possible. The use of UDI is important to post market surveillance of medical devices it improves incident reporting, enables better targeting of recalls, better monitoring of by competent authorities, reductions in medical errors and a reduction in the risks associated with counterfeit products entering the supply chain and better procurement and stock-management decisions in hospitals.

In June 2015 6 pilot Scan4Safety Demonstrator sites were launched. Demonstrator sites were required to implement GS1 and Pan European Public Procurement Online (PEPPOL). Demonstrator sites must be able to identify organisations using GS1 Global Standards, use UDIs for procurement and internal supply chain management and be able to identify all patients using GS1 compliant patient wrist bands and scanning patient records. The Health Service Investigation Branch (HSIB) investigation into the implantation of wrong site prostheses during joint replacement surgery recommended that DHSC expands the remit of the working group consisting

of Derby Teaching Hospitals NHS Foundation Trust's Scan4Safety Programme, National Joint Registry and MHRA to include alerts to identify wrong protheses prior to surgery (recommendation 4).

6. Current MHRA activity

- Commonwealth Vigilance Workbench Longitudinal Module

Vigilance and Risk Management of Medicines (VRMM) Division, following a proof of concept study, have recently launched a pilot to further understand how Clinical Practice Research DataLink (CPRD) data can be used in early signal strengthening. The pilot is using an analytical platform developed by Commonwealth Informatics and based on research conducted by the Uppsala Monitoring Centre. The purpose of the platform is to facilitate the rapid and routine interrogation of electronic healthcare records early within the pharmacovigilance cycle. It is specially designed to provide analyses of the whole CPRD database that can be used to place a safety signal arising from spontaneous reports or elsewhere into the context of the exposed population and to start exploring causality through examining the temporal association between the start of prescribing and occurrence of an adverse event. This work is supplementary to routine pharmacovigilance processes in VRMM and the data mining algorithms employed to interrogate the Yellow Card database, which are also arguably already a form of AI.

- CPRD machine learning

CPRD has used machine learning methods like natural language processing (NLP) to extract additional dosage information on primary care prescriptions from unstructured clinical comments accompanying prescriptions. These data would not be otherwise available for vigilance activities due to concerns around patient confidentiality. NLP allows the CPRD technical team to extract the relevant dosage information and convert it into coded data suitable for research and monitoring while stripping out any likely identifiers. As an additional safeguard, all the NLP output codes are manually checked by CPRD staff to ensure that no identifiable information is accidentally missed by the NLP algorithm.

Machine learning and data mining techniques have been applied to CPRD data by external researchers in the context of drug repurposing (for e.g. can an existing licensed drug be repurposed for another indication), phenotyping (for e.g. classification of diabetes patients into further sub-groups based on patient characteristics) and to inform treatment optimization (for e.g. classifying patients into subgroups based on response to treatment).

CPRD data are vital for supporting pharmacovigilance within the MHRA. Examples include use in supporting proactive vaccine vigilance, studies for evaluating safety signals and quantifying risk, and monitoring outcomes following regulatory action. However, to date, these approaches have all used traditional statistical techniques.

- Regulators Pioneer Fund

The fund sits under the Government's Industrial Strategy has recently opened for new bids. Proposals should either bring about major improvements to a particular (private) industry sector or bring multiple regulators (as defined in the Legislative and Regulatory Reform Act) to explore 'cross-cutting' issues of mutual interest. The intention is that approved projects should result in immediate implementation of a regulatory change to support innovative businesses.

Devices were successful with a bid alongside multiple external stakeholders around creation of a process for AI/ machine learning algorithms to be validated against data from CPRD or NHS Digital and associated regulation of the technology with a view to development of guidance for stakeholders.

- WEB-RADR

MHRA also led the Innovative Medicines Initiative WEB-RADR project which explored the utility of mobile applications and social media in pharmacovigilance. The social media analytics workstream employed NLP to attempt to identify posts that resembled adverse events. Although such posts could be identified in niche areas within the data sources used, the quality of information varied significantly and did not enable identification of signals before when they would have been identified through analysis of spontaneous reports. Value was identified around the areas of abuse and misuse of medicines, and around patient tolerance, however review of a large cohort of individual posts confirmed that most safety issues in the reference sets used simply were not discussed in the social media dataset analysed.

7. Opportunities highlighted by DHSC

DHSC have two potential opportunities for funding/ resourcing development of AI capability using external partners. Each has its own associated risks and benefits, which are discussed in section 8 and participation is not mutually exclusive, so we may choose to take different options to for different remits.

a) Open Innovation Proposal

The Cabinet Office Open Innovation Team² are dedicated to deepening collaboration between the civil service and academia. Specifically, they have a remit to establish partnerships with academics to produce tailored research for government departments. Amongst other areas, they are in the process of scoping projects in robotic process automation, human computer interaction and data scraping. DHSC indicated that University of Southampton, one of the Open Innovation Team core partners could be a good match for exploring some of the AI capabilities that we may be interest in developing. Open Innovation research is not intended to offer sustainable solutions, but rather explore the potential of a conceptual solution.

b) AI Sector Deal Bid

The AI Sector Deal³ is a funding commitment from government to offer

² <https://openinnovation.blog.gov.uk/2018/01/31/open-innovation-team-deepening-collaboration-between-officials-and-academics/>

³ <https://www.gov.uk/government/publications/artificial-intelligence-sector-deal/ai-sector-deal>

support to the AI sector by making funding available to organisations including the public sector to support innovation and stimulate uptake of AI. DHSC have indicated that they are working up a more detailed proposal with a view to exploring the opportunities highlighted in our State of the Nation summary for PS(H) through this route. We understand that Sector Deal funding is intended to be a longer-term arrangement with a view to ongoing arrangements between government and the AI sector. DHSC also highlighted that funding may be available through GovTech, which forms a component of the AI Sector Deal. Within this program, public sector organisations are able to raise challenges and then work in collaboration with technology firms to define, develop test and access creative solutions to them. The public sector organisation must commit to buying the solution if the challenge is solved. Timelines were too short for submission of a proposal for the current round, so we have not explored this in more detail at this stage.

8. Risks & Benefits

- Reputational & scientific

Machine learning, including data mining, has significant potential for harnessing existing data sources for more efficient and earlier signal detection but this is a developing area and experts have identified some risks to scientific integrity that need to be addressed in study design and implementation. The issue of bias, as highlighted earlier, is not specific to machine learning, but the complexity of machine learning algorithms and hidden assumptions can make it more difficult to spot biases. Moreover, machine learning algorithms in themselves may not be good at distinguishing noise from useful predictors so there may be issues with translating findings into recommendations that are implementable by a GP/end-user.

Machine learning methods such as artificial neural networks that are most commonly associated with AI have also been referred to as 'black box algorithms' as they employ latent (or hidden) variables in the analysis that may be impossible to describe in a meaningful way. This is especially problematic when patient safety is at stake as there is an element of 'learning from errors' in machine learning methods and lack of clarity about how an algorithm arrived at a certain prediction could make it difficult to traceback which algorithm assumptions are invalid.

Data mining has been used by MHRA to assess the completeness and accuracy of applied assessments and encoding of ICSRs, with consequent correction required in aggregate data reporting and analysis. These issues have been correctable because the datasets/ algorithms have been static to this point, however, this will not be the case with a machine learning approach. Computerised auto-encoding has also been implicated in low quality reporting and contribution to 'noise' in datasets.

Concerns have been expressed over over-optimistic predictive performance at the initial algorithm testing stage that did not translate into practice. This also raise the concern of false discoveries during signal detection though it can be noted that this issue is not specific to machine learning or data mining. Thus, validation of findings from machine learning algorithms is key as is the case with current signal detection mechanisms. This may be problematic within current regulatory frameworks if there

is a lack of alternative data sources to use for validation of a signal. The full impact it would potentially have on us as a regulator, where we may be unable to use CPRD or other relevant data sources to confirm a signal, is not known. Methodological advances in data mining and machine learning could help to start to address this issue and could be one area for collaboration with AI experts.

Finally, experts attending the CPRD expert consultation workshops have cautioned against the naivety of some data scientists who believed that given large volumes of data, they could develop algorithms with better predictive performance than clinicians. Another concern was about researchers who wanted to use computing power to 'circumvent knowledge' or compensate for 'deficiencies in knowledge' (or theoretical understanding). Thus, while acknowledging the immense potential of machine learning, experts advised a measured approach to developing and testing machine learning methods in collaboration with clinical and policy experts to ensure that any developments were responsible and relevant.

Saying all that, the approaches discussed earlier in this paper are developing, with considerable research being undertaken in academic and private technology organisations, and as an Agency we do not have sufficient expertise in house to understand the range of opportunities that such methods may potentially offer. We do have considerable expertise in the vigilance process from detection of a signal to final decision making however so to ensure that any research in this specific area is well considered and relevant it seems important that we take an active role in shaping it.

Irrespective, MHRA must lead on the science and guidance associated with pharmacovigilance and wider drug and device safety. The interpretation of existing legislation and guidance and supplementation where gaps occur will be an important by-product of the proposed work.

VRMM/ Devices/ CPRD

October 2018

