MRS Response: House of Lords Communications and Digital Committee: Call for Evidence on Large Language Models (LLMs)

About the Market Research Society

1. The Market Research Society (MRS) is the UK professional body for market, opinion and social research, insight and analytics. MRS is the world's largest and oldest research association, representing 5,000 individual members and over 600 accredited Company Partners in over 50 countries and has a diverse membership of individual researchers within agencies, independent consultancies, client-side organisations, the public sector and the academic community.

2. MRS' expertise as the lead authority on market, opinion and social research is recognised around the globe. MRS provides the policy and standards expertise for the UK plus a number of global associations including EFAMRO the European Research Federation and EPHMRA the international healthcare research association. MRS also has close business ties with other research associations around the world via its participation in the Global Research Business Network (GRBN) plus formal agreements with associations in the US, Australia and Japan.

3. MRS promotes, develops, supports and regulates standards and innovation across market, opinion and social research and data analytics. MRS regulates research ethics and standards via its Code of Conduct[1]. All individual MRS members and Company Partners agree to regulatory compliance of all their professional activities via the MRS Code of Conduct and its associated disciplinary and complaint mechanisms.

4. Market, opinion, and social research is the systematic gathering and interpretation of information about individuals or organisations using the statistical and analytical methods and techniques of the applied social sciences to gain insight or support decision making. It involves systematic study of different spheres of society, politics, and the economy. Research, insight and analytics stand at the heart of all well-informed commercial, social and political decisions. Insight into what makes a product, business initiative or government policy work is often the hidden – yet defining – factor between success and failure. It is our sector that provides the deeper intelligence needed for our world today.

5. More information about MRS can be found on the MRS website: https://www.mrs.org.uk/

[1] See: https://www.mrs.org.uk/pdf/MRS-code-of-conduct-2023.pdf

Executive Summary

6. This response analyses the impact of current and future of Large Language Models (LLMs) from the perspective of the market, opinion and social research, data analytics and insight sector (the 'research sector').

7. We expect that LLMs will gain greater popularity within the research sector alongside other Generative AI applications. The technology can bring productivity gains, particularly in areas such as automated reporting; theme summarisation; prediction; text data cleaning and analysing large data sets.

8. As the technology is fast changing it is necessary to improve the understanding and confidence in the future trajectory of LLMs. To that end we encourage the UK Government to build up its capacity and understanding of LLMs, and to create an innovative and effective regulatory framework, one which responds to technological advancements in a timely and appropriate manner, and to fulfil this obligation in tandem with businesses and industry, in order to bring about the most effective and relevant outcomes. The UK Government should consider its own internal future trajectory of LLMs as a reference point to support its own understanding and decide what type of outcomes it envisages, whilst also taking into account how this technology co-evolves with other digital technologies, innovation and skills.

9. Investing in AI literacy and evolving the educational curriculum should be a long-term focus, and not limited to tertiary education. The UK Government should recognise importance of long-term investment in primary and secondary education, and within the workforce, by integrating LLM material into existing key-stage resources and workforce training.

10. There are still some inherent risks associated with the use of LLMs. However, we believe that improving transparency, establishing effective risk mitigations, and providing recourse and routes to remedy will effectively support building trust with the use of LLMs. Transparency is a particularly important principle , as outlined in the UK Government's AI White Paper Consultation. For example, we believe it would help improve AI transparency by requiring organisations to publish their AI ethics policy, in an analogous manner to how organisations publish their privacy notice. These statements should ideally be consistent and aligned with sector Codes and guidelines, provided by professional bodies and trade associations, which can provide a simpler and blanketed approach to ensuring ethical practices, such as transparency.

11. We broadly welcome the regulatory approach proposed by the AI White Paper, as it provides for a flexible sectoral framework without being too prescriptive. It is important however that the UK's approach aligns with other significant legislation, such as the US and the EU, to ensure that UK businesses are able to leverage fully the international AI and LLM opportunities. It is vital that we don't end up with vastly different rules across the globe. Many UK businesses operate internationally and having to work under separate requirements in some of their key markets will add costs and hinder their use of AI and LLM.

12. Given the rapid developments with generative AI and LLMs, it is necessary to adopt a flexible approach to regulation that can respond to new challenges as they arise. As such, it is important that the UK Government establish an industry-led sector-specific self-regulation. As with the AI White Paper Consultation, we believe it is essential that the UK Government develops an approach which mobilises secondary regulators, such as the MRS, in the development and

deployment of applicable and relevant sector AI Codes and guidelines to implement the envisaged regulatory framework.

13. We recommend the UK Government creates a mechanism for trade associations, professional bodies, and self-regulatory Code holders to provide business insight about the application of AI within specific sectors to the UK Government and the statutory regulators responsible for AI regulation. A similar approach has been adopted for the drafting of the Data Protection & Digital Information (No.2) Bill, and we suggest such an approach could work equally as well for AI.

14. Furthermore, we recommend that the UK Government encourages the relevant statutory regulators, such as the ICO, to actively work with the professional bodies and trade associations to develop joint Codes and guidance to fill in the regulatory gaps and to help to raise sector awareness and understanding.

15. We remain optimistic about the UK's current strong positioning and its ability to take economic advantage of the opportunities provided by this technology whilst minimising risk.

Consultation Responses

Q1. How will large language models develop over the next three years?

1. LLMs are being mobilised in the Research sector. The research sector is starting to utilise LLM products to assist in designing research; predictive modelling (analysing large sets of data), etc. LLMs are forming the foundation model to many other software applications, and we can expect LLMs to grow in scale and use, as a result we expect LLMs to increase in accuracy and higher level or personalisation[2].

2. On current trends it is likely that there will be greater proliferation of domain specific or fine-tuned LLMs and new versions of generic foundation models - an AI system with broad capabilities that can be adapted to a range of specific purposes.

3. LLMs will increase in processing power, accuracy and speed. We are aware of the scale at which token window sizes have increased. Increased token window sizes will aid in building context and improving answer relevancy resulting in greater accuracy and potential for a higher level of personalisation.

4. We expect the reasoning capabilities of these systems to improve through techniques such as *chain-of-thought*[3] – which breaks down problems with a series of intermediate reasoning steps; *tree of thoughts*[4] – which allow LLMs to perform deliberate decision making; and newer ones such as *graph of thoughts*, which can model information generated by LLMs as an arbitrary graph and offers improvements over the prior two techniques. Work such as this will continue to bring LLMs closer to how humans perform reasoning.

[2] https://blog.gopenai.com/how-to-speed-up-llms-and-use-100k-context-window-all-tricks-in-one-place-ffd40577b4c
[3] Wei et al (2022) Chain-of-Thought Prompting Elicits Reasoning in Large Language Models https://arxiv.org/abs/2201.11903
[4] Yao et al (2023). Tree of Thoughts: Deliberate Problem Solving with Large Language Models https://arxiv.org/pdf/2305.10601.pdf

5. There is also ongoing development to integrate LLMs with other AI techniques such as mixture of experts (MoE) models[5] that effectively break down complex tasks into smaller pieces to add learnable parameters to LLMs. In other words, the architecture of LLMs, as they become able to process data more efficiently, will be formed of a greater number of nodes, within a greater number of layers, performing more specialised tasks with greater efficiency.

6. Multimodal LLMs are increasingly drawing significant research attention[6]. Multimodal, in this sense, means the ability to encode multiple types of input alongside text to overcome the limitations of text only LLMs. For instance, written (language), audio (music), visual (moving images) and sensory inputs prompts could be inputted into an LLM model to produce more efficient and accurate outputs for applications such as search or e-commerce.

7. Whilst LLMs are making huge progress in their technical capabilities they do have current limitations. Over time, it is likely that these limitations will be overcome or be addressed via effective mitigations. For example, LLMs are prone to hallucinations, this is where language models begin to make things up from previous learnings, which makes it important to retain humans in the loop to improve context and check the accuracy of the output.

8. Furthermore, in LLM commercial applications, humans are often kept in the loop for efficacy purposes, and human intervention is also a necessity for safety and security measures. It is likely that human intervention will adapt as models mature, creatives (humans) will ultimately still have to exercise some judgement with regard to creating prompts, as well as choosing and revising output. Humans, though not necessarily creatives, may also need to be kept in the loop to address biases and non-efficiency-related issues (discussed in further depth later). Although LLMs (and other types of Generative AI) will increase the availability of content, fundamentally these systems are not in their nature creative per se, and their outputs are based on existing information and human creativity and reasoning. As a result, we think that there will be a premium for uniquely creative output that, so far, only humans are capable of fulfilling.

Q1 (a). Given the inherent uncertainty of forecasts in this area, what can be done to improve understanding of and confidence in future trajectories?

9. There are a number of key steps that can be taken to improve the understanding of and confidence in the future trajectory of LLMs. It is also crucial to maintain an industry- led approach to create an efficient regulatory framework and beneficial business landscape for small and large businesses alike.

10. The UK Government should develop its capacity to improve its understanding and retain its knowledge of LLMs, and generative AI more broadly. The UK Government may consider constructing its own internal future trajectory of LLMs as a reference point to support its own understanding and decide what type of outcomes it envisages, whilst also taking into account how this technology co-evolves with other digital technologies, firm innovation and skills[7].

[5] Mixture-of-Experts Meets Instruction Tuning: A Winning Combination for Large Language Models. https://arxiv.org/abs/2305.14705
[6] A Survey on Multimodal Large Language Models. https://arxiv.org/abs/2306.13549
[7] Ciarli et al (2021). Digital technologies, innovation, and skills: Emerging trajectories and challenges. Research Policy Vol 50 Issue 7 September 2021.
https://www.sciencedirect.com/science/article/abs/pii/S0048733321000913

11. Further work is required towards strengthening understanding across policymakers and society more broadly to increase trust in technology and make it more inclusive. A higher level of understanding would also help bridge the knowledge gaps between policymakers, the workforce and the academics and researchers undertaking work to push the boundaries of these technologies.

12. To increase understanding, there would be merit in:
    a. Promoting greater transparency of how models arrive at their decision or derive their output;
    b. Preparing the workforce for future adaptation; and
    c. Investing in AI literacy and evolving the educational curriculum.

13. Addressing these steps in turn:

    a. Greater transparency will require collaboration between the UK Government and business. The UK Government should support the open-source ecosystem and ensure that where models are not open-source, for commercial reasons or otherwise, proportionate steps are taken to ensure that the owners of the relating IP are transparent to the greatest extent possible whilst not compromising compliance with applicable IP law, and/or the commercial viability of the closed model.

    b. Transparency could also involve declaring the level of AI being used in an information transaction. There are different levels of AI from the use of algorithms, Machine Learning to more complex AI approaches. Enabling greater understanding of the type of activity that is being undertaken should also help to ease some of the concerns of those who are subject to the use of AI.

    c. Preparing the workforce for future adaptation is a monumental task which though led by Government should take advantage of public-private partnerships, especially in relation to reskilling and upskilling employees. For example, the MRS is a body which upskills and trains businesses in data ethics and developing best practice across the research industry and is introducing a new OnDemand course using ChatGPT which complements existing MRS course such as *AI, VR and the Metaverse* and *Data Science Analytics*. MRS also produces specific sector guidance on subjects such as: Biometrics; the Metaverse and forthcoming guidance on ChatGPT. The purpose of this guidance being to advise, steer and provide best practice to practitioners on fulfilling legal obligations, such as GDPR and rendering the most effective uses of different emerging technological tools.

    d. Many 'workers' using LLMs will be self-employed, some will be volunteers, some will operate LLM models primarily as hobby or for pleasure, with value added to the economy being an incidental side effect. However, most of those employed will be 'traditional' members of the workforce (employed within a medium or large enterprise) and will require suitable support to prepare their workforces.

    e. Investing in AI literacy and evolving the curriculum should be a long-term focus, and not limited to tertiary education. The UK Government should recognise importance of long-term investment in primary and secondary education, by integrating LLM material into existing key-stage resources. For instance, teaching how to use prompts to generate code in the Key-Stage 2 'Teach Computing' Curriculum. These long-term investments in the domestic skills pipeline, targeted at young students, is a core part of what the research sector requires to address the current chronic skills shortage. In addition, AI

literacy should incorporate creative subjects, like Art or English leveraging the creative capabilities of LLMs and AIs, promoting STEAM (Science, Technology, Engineering, *Arts*, and Maths) skills.

f. Within the workforce, foundational training should be provided to upskill staff in a broad manner – this is to accommodate for the rapid advancement and growing uses of LLM's and generative AI across the workspace. Particular focus should be given to safety, security, and risk management. In the same way that health and safety and GDPR are essential training, it is likely that LLMs and AI will need to be addressed in a similar manner. The UK Government should look to work with professional associations and trade associations to develop vocational technical AL LLM skills which are needed for individual sectors as well as broader horizontal requirements. Organisations such as MRS offers comprehensive training for the research sector and we would welcome working with the UK Government to meet its objectives in this area.

Q2. What are the greatest opportunities and risks over the next three years?

14. Within the research sector, the benefits of AI and LLMs include intelligent interviewing, prediction and conversational search queries.

15. For example, [8]with the use of prediction, LLMs can extract embeddings (mathematical representations) that other machine learning models can use to predict outcomes of interest. For instance, they can predict the performance of a TV ad based on the dialogue or relate people's qualitative experience interacting with a service representative to their brand loyalty or churn. Other benefits include automated reporting, whereby large volumes of quantitative data that need sorting, summarizing, and presenting can be quickly organised and create draft headlines based on charts, tables, models, and executive summaries.

16. However, there are risks associated with the use of LLMs, such as hallucinations and false predictions, whereby LLMs may make incorrect predictions, particularly when they encounter novel or ambiguous data. Sometimes, they may even make things up or 'hallucinate,' leading to false predictions. This creates the risk of perpetuating false or misleading information by accepting an output at face value. In either case, it is necessary to implement a system of quality control before deploying such outputs or incorporate indemnities in service provider-end user agreements. This is particularly where human intervention remains essential. Other examples in the research sector include translation services; whereby LLM's are being used to translate surveys into multiple languages and communicate the same sentiment, whilst being contextually and culturally relevant to each case. Historically, translation services included numerous iterative processes with different native and bilingual speakers, and back-translations for comparative purposes before unanimous agreement and publication. We are familiar with use cases where LLM's were used for translation services, a native speaker thereafter reviews the translation, at which point it is identified that the meaning of the research has been entirely diluted and is not fit-for-purpose. Practices such as this (while they may appear to streamline processes) are capable of falsifying research results if they are based on responses that do not convey the sense of the original language. As such the value of human intervention and reasoning and creativity and productivity should not be omitted or undermined in the steer towards AI, and it is important to reflect these use cases in training and guidance.

[8] https://kadence.com/en-us/how-large-language-models-are-changing-market-research-2/

17. It is also important that economic benefits that result in the adoption of LLMs are not concentrated in the hands of the few, hence it is important that open source and closed source LLMs can compete on an equal footing and interoperability is encouraged.

18. Further risks also extend to privacy concerns, whereby employees may unwittingly pass on sensitive information to LLMs. We think this risk can be mitigated by the deployment of private LLMs which incorporate encryption and data privacy preserving techniques, and institute training; follow and signpost relevant guidance and corporate policies to reduce the likelihood of inadvertent data leakages.

19. There are also ethical considerations and the risk of undesirable representational bias emerging in LLM outputs as the result of historical stereotyping or unfairness in training data. This highlights the importance of conducting bias audits to evaluate the LLMs output across a range of inputs to reveal potential bias. However, there are tools and techniques which can mitigate such biases. Researchers have also trained logic-aware models to reduce harmful stereotyping[9]. Although this isn't a simple and blanketed approach, reflecting the diversity and cultural and political landscapes of jurisdictions, in an area which will certainly require for diverse human intervention as an ongoing practice.

20. There is also the potential for unforeseen risks and harms, which only materialise once the technology becomes more sophisticated or is utilised for yet defined use-cases. This adds further argument to the idea that regulation should be flexible and sector specific, and importantly transparent in real-time and throughout its iterative processes. Guard-rails should be proportionate to the perceived risk and should not stifle innovation. They should be principles-based, flexible, and take advantage of existing mechanisms, established regulators, and extant accumulated knowledge.

Q2 (a). How should we think about risk in this context?

21. We think that the best way to evaluate the risk is to, first, consider context and use-cases. This is because certain activities, especially those that could have a legal, financial, privacy, or psychological impact to individuals, would inherently carry greater risk – all other factors remaining equal. In these activities, these impacts could include employment decisions, access to finance, or exploitation of vulnerabilities.

22. Once the context and use case are established, there are a number of steps that we think should be considered. These are:

    a. Transparency;
    b. Risk mitigation; and
    c. Recourse / routes to remedy

23. The principle of transparency, if done right, could go a long way to enhancing people's confidence and trust towards AI. This is equally true for a business-to-business context. Transparency is particularly useful when organisations use AI to make decisions which have a legal or financial effect on an individual. Not only would this give consumers the knowledge that the decision has been made using AI, but it would also be consistent with the exercising of GDPR rights. Articles 13.2(f) and 14.2(g) of the GDPR already provides for data subjects being informed of the existence of automated decision-making.

[9] https://news.mit.edu/2023/large-language-models-are-biased-can-logic-help-save-them-0303

24. In relation to risk mitigation specifically, we consider that a key decision to be made here is the extent to which humans should be 'kept in the loop', especially in a context of heightened risk – for instance, when dealing with special category personal data. In relation to this question, three further sub-questions should be considered by policymakers and regulators:

    a. *Where* those humans should be kept in the loop. For instance, before prompts are provided, during the formulative process, or after the LLM has outputted the result.
    b. The *role* of that human. For instance, is the role of the human to arbitrate the LLMs' output, using judgement to weigh different factors; or to make binary 'failsafe' decisions weighing just one factor (e.g., bias towards a protected characteristic).
    c. The extent to which an LLM model is producing advisory outputs, versus taking executive decisions. In the former case, there may be no need for humans in the loop because, even when things go wrong, the consequence of the harm should be low or nil. Contrast that with the latter case, where we should consider to what extent those executive decisions are mission-critical, where human intervention may be needed in multiple stages of the generative process as redundant fail safes. Scale must also be considered. Low instances of low harms may be acceptable, but high incidences of low harms, due to their scale, may be an unacceptable risk.

25. In relation to appropriate routes to remedy, we should be careful in assuming that people are:

    a. aware of their right and/or routes to contest a decision; and
    b. aware that an LLM has been utilised in that scenario.

26. If individuals harbour any doubts regarding the accuracy of a system, especially when it is being utilised for a purpose that could have a significant effect, they are unlikely to use it unless they have a means to challenge or dispute outcomes that they think are incorrect. Further research may be required to understand the perspective of the general public on the issue.

27. When considering peoples' ability to challenge or dispute outcomes, differentiating between the general adult public – who are deemed to have sufficient knowledge and agency to choose whether to challenge decisions – and vulnerable groups, who cannot – often due to no fault of their own. It is important that policy should take special consideration of vulnerable groups. We should avoid scenarios whereby LLM operators opt to use a low accuracy model, without proper due diligence, and later blame the model for the consequences on vulnerable groups instead of the initial decision to use that model. Furthermore, direct discrimination by an LLM model may not be the only scenario in which individuals may need to contest decisions. There may be harms which occur indirectly especially if there exists asymmetrical power structures or misaligned interests. Also, it is important to consider how the contestability context may change once LLMs become an embedded and intrinsic part of a wider societal 'standard'. In this scenario it may become more difficult for humans to contest such decisions or standards.

28. Harms should be considered according to consideration of their:

    a. Legality. For instance, bias in certain social contexts, may be perfectly legal, whereas in an employment context it may be illegal.

b. Proportionality. A utilitarian test of the extent to which the benefits are outweighed by potential harms. If they are outweighed, consideration should be given to whether the harm be reduced by a less restrictive means.
c. Directness / causation. Is the directly caused by an LLM output. For instance, the output of an LLM model decides that person 'A' loses their job; or is it indirect, the output of an LLM model provides poor financial advice to 'A's employer, which is one of a number of factors which later causes 'A' to lose their job.

29. Professional and trade organisations such as MRS are best placed to apply these considerations in a manner which is principle driven, transparent and inclusive of the expertise and business activities of the profession – both vital qualities to have in markets where technology is advancing rapidly.

Q3. How adequately does the AI White Paper (alongside other Government Policy) deal with large language models? Is a tailored regulatory approach needed?

30. MRS broadly welcomes the regulatory approach proposed by the UK Government's White Paper, as it provides for a flexible sectoral framework without being too prescriptive. On one hand, it may not be necessary to re-engineer governance structures specifically for AI when in fact existing ones may in fact prove to be sufficient, and where the landscape is rapidly evolving. On the other hand, we are aware that AI technology is moving at a fast pace but not all of the implications of this technology are immediately obvious. Hence it is necessary to adopt a flexible approach to regulation that can respond to new problems as and when they emerge. It is for that reason that we consider that there is an important role for an industry-led sector-specific self-regulation. It is also important that the UK's approach aligns with other significant legislation, such as the US and the EU, to ensure that UK businesses are able to leverage fully the international AI and LLM opportunities. It is vital that we don't end up with vastly different rules across the globe. Many UK businesses operate internationally and having to work under separate requirements in some of their key markets will add costs and hinder their use of AI and LLMs.

31. However, we think that AI White Paper principles could be expanded to include the following:

    a. Principle of privacy by design
    b. Principle of accessibility
    c. Principle of human oversight
    d. Principle of sustainability

32. As to the "privacy by design" principle, including the incorporation of an "accessibility" concept, this would involve implementing a privacy by design principle separate from the principle of fairness which encourages the development of systems that have privacy safeguards and provides the appropriate transparency and control over the use of data.

33. This would have two effects: 1) elevate the importance of privacy to a top-level principle, as opposed to be a secondary principle under fairness; and 2) resolve a potential situation where these principles sit awkwardly with UK GDPR principles, given that UK GDPR already incorporates fairness and transparency.

34. The rationale behind the accessibility concept is that AI and LLM development should be socially beneficial and, on that basis, available for all uses so long as they are consistent with the cross-sectoral principles. The current cross-sectoral principles speak of fairness of outcome but do not consider the issue of fairness of accessibility.

The worst outcome would be to create another form of digital divide creating a gap between those who can easily access and use AI and LLM technology and those who cannot.

35. In relation to the principle of human oversight, the UK Government should consider including principles of human oversight (to fight bias and unfairness) and continuous evaluation of AI systems and policies as standalone principles.

36. Finally, on the inclusion of a principle of sustainability, we should not ignore the environmental cost of data centres and the energy used when AI models are being trained on vast amounts of data. Hence, it is worth including a sustainability principle to help guide business to achieve energy efficiencies over the lifecycle of datacentres.

Q4. Do the UK's regulators have sufficient expertise and resources to respond to large language models? If not, what should be done to address this?

37. There is broad agreement from business that bodies such as Digital Regulation Cooperation Forum (DRCF) can help coordinate policy among regulators and improve regulatory coherence. The UK Government should also mobilise the knowledge and experience held within the sector, by working closely with self-regulators such as MRS to gain wide insight across members and business.

38. The Digital Regulation Cooperation Forum (DRCF) provides a structure for the CMA, ICO, Ofcom and FCA to coordinate and cooperate on some of the challenges posed by regulating online. It would make sense to either expand this forum or set up a similar model within the AI taskforce that consists of a working-level group of secondary regulators, trade associations and professional bodies, together with AI, legal, privacy and human rights experts. This work should be consultative and transparent so that the views of the DRCF and its member regulators are well-understood by all market participants and provide a sound basis for business decisions.

39. Furthermore, Statutory regulators could also consider joint guidance with sector bodies and Code holders to help to fill in the regulatory gaps and to help to raise sector awareness and understanding. It would be beneficial if annual conferences on AI policy could be organised to create a platform for policymakers, regulators, industry, and academia to discuss and exchange views on AI governance.

Q5. What are the non-regulatory and regulatory options to address risks and capitalise on opportunities?

40. Regulatory options available to the UK Government when considering LLMs must be agile and play into the comparative advantages the UK has with other jurisdictions.

41. The research sector is primarily regulated by MRS, a leading independent regulator which is recognised amongst policymakers and legislators as effective, fair, and efficient. Of course, other regulators, such as the CMA or the ICO, also have a secondary role regulating the sector in specific areas, such as data protection.

42. It is useful to note that self-regulation is a faster and more effective practical mechanism to address and prevent AI related problems in many instances, of course within the boundaries set by statutory regulation. As aforementioned, MRS regulates standards and innovation across market, opinion and social research and data analytics. MRS regulates research ethics and standards via its Code of Conduct and all individual MRS members and Company Partners agree to regulatory compliance of

all their professional activities via the MRS Code of Conduct and its associated disciplinary and complaint mechanisms. MRS has a long and strong track record in providing excellent regulation of the sector and profession, and these existing mechanisms should be mobilised by the UK Government.

43. The UK Government could also consider mobilising sector specific Codes of Conduct, such as the MRS Codes of Conduct, which are crucial in helping to protect and regulate research, insight, and data practice[10]. This could help with the proper application of the cross-sectoral principles and considers the specific needs of micro, small and medium-sized enterprises better. For example, a Code of Conduct could support clarifying supply chain liabilities in an AI chain. In a research project for example, where AI technologies are mobilised across the project journey, a Code of Conduct would support businesses identify practitioners/controllers/processors/users, and therefore determining responsibility for any potential AI failures. Codes of Conducts are essential tools for business, and are far easier to understand, apply and develop as opposed to legislation.

Q6. How does the UK's approach compare with other jurisdictions?

44. The UK is taking a principles-based, sectoral approach designed to mitigate harms. At this stage, it is not proposing new law specifically to regulate AI but use existing mechanisms and established regulators. The idea being that principles are flexible and can be applied broadly even as our understanding of LLMs develops. We welcome this approach as positive for innovation, and reinforcing the widely held international perspective that the UK has an open, creative, and entrepreneurial culture, where businesses are free to respond to new technology with speed and agility.

45. The EU, however, is taking a risk-based horizontal approach to regulation and borrows heavily from the GDPR in its evaluation of risk. The approach also takes cues from the Single Market's product safety and conformity model, seeing AI as, fundamentally, a product. The risk-based approach is comprised of four categories: minimal risk, limited risk, high risk and unacceptable risk. The issue with this approach is that an AI being deployed in a low-risk setting may still be caught as a high-risk system. This may have negative implications for innovation and investment.

46. The EU's AI Act aims to identify AI systems that present unacceptable risk (i.e. threats to safety and livelihoods) and to apply stringent obligations to AI systems that present high risk (e.g. technology used for critical infrastructures and educational training). The EU is also taking an extraterritorial approach with its AI Act, probably with the idea that the EU way of regulation will be exported elsewhere. The EU will be following up the AI Act with the AI Liability Directive. The Act could be said to be too broad in its definitions of AI technologies while still focusing too narrowly on the application layer.

47. The US and China both see AI as a geostrategic asset, with questions of national security and power projection prioritised – even above economic impacts.

48. The US, specifically, will prioritise not stifling innovation whilst protecting national security. The US has recently published a draft AI Bill of Rights to improve accountability however the Bill is non-regulatory and non-binding. The US culture informing their approach to LLMs is a mixture of the Silicon Valley 'open internet' ethos and what others call the DC vision of a 'commercialised internet'.

[10] https://www.mrs.org.uk/standards/code-of-conduct

49. China's approach largely stems from its desire to place AI development under state control and for it not to develop in a way that would run counter to China's values and ideals. China very much sees the potential for AI not only in productivity gains but also its military and public security applications.

Q6 (a). To what extent does wider strategic international competition affect the way large language models should be regulated?

50. The UK has significant influence within international institutions, partly stemming from its history as a great power, but also from cultural, linguistic, and financial factors. The UK should use this influence within international institutions to set standards that directly and indirectly give the UK a comparative advantage in the economic benefits that will flow from LLMs and generative AI generally.

51. However, this influence is often overshadowed by the fact that, in terms of both raw population numbers the UK is not as large as either the US or EU. The US/EU agreement on taxonomy is an example where the UK did not maximise the positive factors (such as soft-power & influence) it has at its disposal to be in the vanguard of international standard-setting. There is no reason, in principle, why the UK cannot replicate the pre-eminent international position is holds in other areas, for instance in financial regulation, in the field of AI and LLMs.

52. The UK Government should also recognise that international competition also operates in parallel with the corporate competition between LLM providers in terms of setting their own guidelines for responsible AI.

Q6 (b). What is the likelihood of regulatory divergence? What would be its consequences?

53. There is always some likelihood of regulatory divergence, the question is to what extent and which jurisdictions' regulation will be seen as the pre-eminent global standard, should this occur. Although divergence would likely lead to competing standards, the de-facto global standard would inevitably converge towards the largest market where businesses can scale most effectively. Key consequences and/or risks could include a race to the bottom, although this is by no means a certainty if regulatory divergence occurs.

    The UK has adopted an agile regulatory approach in order to foster innovation, which we are supportive of as a sector. In contrast, the EU has focused on establishing a set of standards with the hope of attracting increased investment through regulatory clarity. It remains to be seen which of these approaches will yield greater success, but historically (for instance when considering the rise of tech platforms in the late 90s and early 2000s), a more agile regulatory approach fosters greater levels of innovation. Although, for regulatory purposes and for fostering join-up between EU services, it is desirable to have at a basic level, harmonisation between regulatory approaches, to foster adequacy and opportunity between jurisdictions.

54. We think that it is imperative that business can navigate both the UK and non-UK rules for AI. As AI is borderless, the fact that the UK has a light regulatory approach will not mitigate the impact of other countries/regions that may adopt more stringent approaches. Fragmented and divergent legislative requirements remain one of the greatest barriers to innovation. For UK businesses to leverage the opportunities of AI will require understanding of all the regulatory requirements. We recommend that as part of the UK Government's pro-innovation approach, it provides information about alignments and differences with non-UK rules and provide routes and pathways to enable UK business to navigate these requirements to maximise the opportunities for UK businesses, particularly SMEs.

We would welcome discussing further the MRS submission in response to the House of Lords Communications and Digital Committee: Call for Evidence on Large Language Models (LLMs):

- Debrah Harding, Managing Director: debrah.harding@mrs.org.uk
- Kaleke Kolawole, Head of Policy: kaleke.kolawole@mrs.org.uk