

christians
against
poverty

CAP

Discussion Paper - Artificial Intelligence and Machine Learning

***CAP's written response to the Financial Conduct
Authority's (FCA) discussion paper***

October 2022

always hope.

Artificial Intelligence and machine learning

Summary

Christians Against Poverty (CAP) has more than 25 years of experience in delivering debt help. CAP supports over 8,000 people on their journey out of debt each year and has specialist in-house teams to support vulnerable individuals. Through this experience, CAP has developed insight into the varied and complex needs of vulnerable consumers and their relationships with creditors.

- The FCA should ensure that consumers who are susceptible to harm are protected from the risks of Artificial Intelligence (AI) and machine learning, whilst also making sure firms are including these same people in the market, offering the same service and products.
- Due to the nature of AI and machine learning, outcomes will continuously adapt and change. If the FCA is taking an outcomes-based approach, it needs to ensure that firms are monitoring their outcomes on an ongoing basis, and ensuring consumers are not experiencing any form of harm.
- As this is still a relatively new and emerging industry, CAP would like to see a wide review of the demographics and people groups impacted by historic data sets, or those who are underrepresented.
- Due to the discrimination faced by certain people groups when using data in AI, CAP would like to see the FCA regulation encourage firms to proactively review their data to make sure people groups are represented, and not discriminated against.
- CAP would encourage the FCA to perform a regulatory visit to firms using this technology as a priority.
- CAP would like to see more transparency for consumers to better understand how technology is profiling them, and clear escalation routes in case mistakes are identified.

Questions

Benefits, risks and harms of AI

Q3: Which potential benefits and risks should supervisory authorities prioritise?

CAP identifies numerous benefits to consumers engaging with AI products and services. In particular, the opportunities AI creates to better identify and support customers with characteristics of vulnerability, as well as increasing financial access. If used for these purposes, AI should enhance the experience of engaging with financial services for those customers in vulnerable situations. However, equally as big as the benefit is the risk of misusing this technology. Removing the human element from decisions risks losing the softer skills, such as empathy. It also reduces flexibility and impacts customer service. The FCA has outlined numerous risks, such as targeting consumers' behaviour biases or characteristics of vulnerability, discrimination, financial exclusion and reduced trust. CAP recommends the FCA prioritise mitigating against these risks. For instance, AI can be used to support decision-making, with the AI providing an estimate of the correctness of a decision, and also highlighting difficult or complex cases, but leaving the final decision to the human. Customers in vulnerable situations are more susceptible to harm, and the harm they experience often has a disproportionately negative impact to their lives.

A further risk identified by CAP is the potential for customers in vulnerable situations to be excluded from the benefits of AI. If firms are, rightly, safeguarding these customers from harm, this safeguarding could look like wrongful exclusion from AI. Care must be taken to avoid excluding customers with protected characteristics from a decision (and thus turned away), if, for example, the AI judges them to be outside confident decision-making due to low representation in the training data – a human decision should always be substituted. The FCA should ensure that all customers are treated fairly, offered the same opportunities and able to interact in the markets in the same way.

Q4: How are the benefits and risks likely to change as the technology evolves?

There is a further risk as the technology evolves that harm could be caused without awareness: the 'black box problem'. This is particularly likely in machine learning, where human intervention is lessened, and computer decision-making is less visible. There is a risk that as products become more advanced and technical, the experts in the technology are part of a minority group who can read and understand what machine learning is doing. Some approaches, such as decision trees, are inherently explainable in terms of relationships between the factors in the data. However, with unsupervised deep learning (e.g. convolutional neural network) this is not the case. That said, explainable methods can

be applied post facto to justify a decision, e.g. LIME (Local Interpretable Model-agnostic Explanations) or SHAP (SHapley Additive exPlanations). By this means, the reasonableness of a decision applied to a client can be checked by a human decision-maker. If there is little or no transparency in how the outcomes are impacting customers, those overseeing the product to ensure it is not causing harm would be too far removed. This is especially pertinent due to the nature of machine learning, which continually evolves and 'learns', therefore outcomes will also be monitored on an ongoing basis.

CAP would like to see the FCA create some specific rules in this vein, protecting consumers from the impacts of data drift or concept drift. It needs to be ensured that as machines and intelligence evolve, firms continue to monitor outcomes and ensure harm is not being caused inadvertently. The FCA could even go so far as requiring firms not to launch products or release updates until the unintended consequences have been identified, evaluated and mitigated. This is of particular concern for minority groups or those with protected characteristics.

Q6: How could the use of AI impact groups sharing protected characteristics? Also, how can any such impacts be mitigated by either firms and/or the supervisory authorities?

In the discussion paper, the FCA notes that those with protected characteristics are more susceptible to harm, particularly by facing discriminatory decisions. Groups can also be placed into risk categories, and this is likely to see more exclusionary practice for those with protected characteristics. As developers work on software, the FCA should encourage firms to build in protections to these groups. This would need to follow an extensive analysis, potentially by the FCA, of the groups susceptible to harm, and what kind of harm, from AI and machine learning. Once groups and biases have been identified, firms should then develop software to ensure harm is not afforded to these groups. This could involve building in bias protections.

There are also implications for data protection as this intelligence makes it possible to derive or indirectly act on private information. For instance, there was a case in the USA where a superstore was using analytics to automatically derive whether women were pregnant, in order to send them relevant coupons. A man discovered his teenage daughter was pregnant because she was being sent advertisements for maternity clothing. This was even without machine learning or AI. With more powerful AI models, instances like these could be much more common.

It is also worth considering the differing needs of consumers sharing protected characteristics. The FCA's Vulnerable Customer Guidance encourages firms to ensure that

consumers are treated as individuals, not placed into categories. The default of AI and machine learning is to place consumers into categories, but not down to an individual level, and not considering the multiple additional needs someone may have.

Those in society facing digital exclusion are precluded from benefiting from much of AI and machine learning technology. CAP would like to see the FCA lead firms in addressing this inequality, and providing these consumers with an equal service. This may well require the AI to identify decisions made with such individuals as requiring human decision-making, even if supported by AI – 'meeting diverse consumer needs through low exclusion' (4.23).

Q7: What metrics are most relevant when assessing the benefits and risks of AI in financial services, including as part of an approach that focuses on outcomes?

For an outcomes-based approach to AI, CAP would recommend an ongoing review of the demographics of the customer base. Is AI or machine learning excluding certain groups? The regulator should speak to people who have lived experience of engaging with the firm, finding out how they found the interactions, how they find the product or service and learning from these experiences.

As previously mentioned, the risk of only measuring the outcomes is that this leads to the opportunity for harm to be caused. CAP would like to see the FCA put measures in place so that firms are analysing the likelihood of harm before customers are exposed to the product or service.

Any firm using machine learning or AI needs to demonstrate the algorithm, and prove that outcomes are at least as good as other methods (for instance, human service). There also needs to be clarity over what controls firms are using for identifying and minimising bias, including having a diverse group of people feeding into the training of the machine. This may be impossible with an approach based on a 'black box' deep learning – outcomes could only be demonstrated statistically, and perhaps checked for 'edge cases' such as vulnerable individuals or those with intersectional protected characteristics.

The number of times intervention is needed to correct erroneous outputs should be documented. This should be coupled with the logging of complaints volumes and themes. There should be a regular review of the AI-generated decisions, with the results of this captured too.

Regulation

Q10: How could current regulation be clarified with respect to AI?

The Consumer Duty is an overarching set of rules, which steer firms to avoid causing foreseeable harm to retail customers. This, alongside the other regulations, rules and principles, will work well together to protect consumers from harm.

More clarity is always welcome, so CAP would encourage the FCA to be more explicit about this emerging technology. CAP would also like to see the FCA prioritise its reviewing of this industry, so new emergents and new technology from more long-standing firms are placed under scrutiny as they launch new technology.

There is also a question of how consumers can both see and understand the data held about them, and allow them to advocate for themselves. In the instance of incorrect profiling or data misinformation at an individual level, consumers need to know where they can go and who they can talk to. CAP would like to see firms provide clear escalation routes and enhanced visibility, presented in a simplified way, on how personal data is being used by a firm to make automatic decisions.

Q11: How could current regulation be simplified, strengthened and/or extended to better encompass AI and address potential risks and harms?

The current regulations do a good job of protecting consumers from potential risks and harms. As a result, the instances of firms purposefully causing harm through exclusion, bias, subverted purposes etc. should largely be mitigated under current regulation. However, the FCA should continue to simplify, strengthen and extend current regulation to encompass this emerging industry and the unintended harm it may cause. CAP would like to see the FCA place more emphasis on regularly evaluating data for bias, and for firms to involve the voice of the underrepresented groups to help build equitable codes. It was welcome to read that 'Firms will be required to monitor the outcomes their customers receive in practice and take action if they identify particular groups of customers are getting poor outcomes' (4.21).

AI models can attain high accuracy on training data while failing to model the real world, a problem known as overfitting. This is likely to give inaccurate or discriminatory results. Regulations should specify adequate validation processes to prevent this.

CAP would like to see firms doing their utmost to protect consumer data from breaches and hacking. AI and machine learning profile consumers and glean a lot of information on individuals, so firms should do all within their power to keep this data secure and not passed on to third parties.

Q12: Are existing firm governance structures sufficient to encompass AI, and if not, how could they be changed or adapted?

It would be good to see a requirement on firms to have a senior manager (or team) sign off on all AI use (4.54). This should emphasise guarding against bias and ensuring equality of access for vulnerable clients.

Q15: Are there any components of data regulation that are not sufficient to identify, manage, monitor and control the risks associated with AI models? Would there be value in a unified approach to data governance and/or risk management or improvements to the supervisory authorities' data definitions or taxonomies?

It would be good to see guidance/regulations requiring the explainability of models (see 4.37: 'there is currently no explicit guidance on issues like model explainability/interpretability of AI models').

The case study included in Box 5 includes a wide-ranging set of principles for Model Risk Management (MRM). The principles set out here are holistic and positive, and could be used as a source of inspiration for the principles being established in this new area of regulation.

Q16: In relation to the risks identified in Chapter 3, is there more that the supervisory authorities can do to promote safe and beneficial innovation in AI?

Within this new technology there is a need for firms to act with integrity and ethics. In the FCA's position as a regulator for financial services, there is a question of whether their remit stretches into the ethics of data use. CAP would like to see the FCA work closely with The National Statistician's Data Ethics Advisory Committee (NSDEC) to ensure the firms it regulates do not breach the Government's Data Ethics Framework.

About Christians Against Poverty (CAP)

With a vision to see transformed lives, thriving churches and an end to UK poverty, Christians Against Poverty (CAP UK) is a national charity that equips local churches to deliver a range of services.

CAP Debt Help provides holistic support for families and individuals facing problem debt with a free face-to-face service – tackling both the financial conundrum and the wider emotional impact. CAP tackles the causes and consequences of UK poverty through free community groups, also run through local churches. This includes Job Clubs, Life Skills groups and the CAP Money Course.

All CAP's services are free of charge and available to everyone, regardless of age, gender, faith and background. To find out more, visit capuk.org.

Requests for further information

This response has been written by Kiri Saunders, Social Policy Manager for Christians Against Poverty (CAP), with contributions from:

Dr Mark C. Sinclair, MA, MSc, PhD, CEng, MIET, MIEEEE, FHEA, System Developer
Matt Kershaw, Junior Systems Developer

Please send requests for further information to:

External Affairs

Jubilee House

1 Filey Street

Bradford

BD1 5LQ

externalaffairs@capuk.org



Disclaimer: This information is not intended to be legal advice and is for educational purposes only. Although care is taken to ensure that this information is accurate and correct at the time of writing, Christians Against Poverty cannot accept any responsibility for mistakes or omissions. Christians Against Poverty excludes to the extent lawfully permitted all liability for loss or damage arising from reliance upon this information.

Copyright © 2023 Christians Against Poverty. All rights reserved. This material may not be reproduced for any purpose without first obtaining written permission from Christians Against Poverty.

capuk.org



capuk



capuk_org



capukorg

always hope.