

Regulatory Sandbox Final Report: Onfido

A summary of Onfido's participation in the ICO's Regulatory Sandbox Beta

Date: September 2020

ico.

Information Commissioner's Office

Contents

| | |
|---------------------------------------------|----|
| 1. Introduction | 3 |
| 2. Executive summary | 8 |
| 3. Product description..... | 12 |
| 4. Key data protection considerations | 16 |
| 5. Ending statement | 27 |

1. Introduction

- 1.1 Onfido Limited (Onfido) provides remote biometric identity verification technology that is designed to enable its clients to prove that their customers are who they claim to be¹. For example, a financial institution client of Onfido's will likely use Onfido's technology to prove the identity of a customer (being an individual) seeking to open a bank account. That customer will be asked to provide the financial institution with a digital photo of their identity document and a selfie taken using a mobile phone or other device. Onfido will then analyse those images to determine the likelihood that: (i) the identity document is genuine; and (ii) the face in the selfie matches the face in the identity document, and that the selfie image does not display evidence of signs of fraud or facial spoofing. If the identity verification check is successful, the financial institution's customer will have proven their identity, and they will continue with the remainder of the bank account opening process.

¹ Please Note – further detail about how Onfido completes identity verifications is provided in the Product Description section of this document.

- 1.2 Onfido's remote biometric identity verification technology is similar to the more traditional act of proving identity in person. In the above example, if the customer were opening a bank account in a physical store, that person would provide the bank teller with their identity document, and the bank teller would analyse the identity document while looking at the customer to ensure that the customer is the owner of the identity document. The key distinction with Onfido's technology is that it removes the necessity for the client and the customer to be in the same physical location, while providing enhanced fraud detection and face matching capabilities compared to what can be performed by a person alone.
- 1.3 As traditional services (such as banking) move increasingly to the digital world and people are less able or willing to transact in person, whether that be due to COVID-19 or other factors, Onfido's technology provides a solution to enable the continued access to services and commerce. However, as this technology is intending to provide a service to benefit the public, it should work for everyone. The technology should be fair and inclusive by design in order to meet relevant equality and human rights law for instance.
- 1.4 As part of Onfido's efforts to ensure that its technology is fair and inclusive, Onfido recognised that its Facial Recognition Technology (FRT) was likely not responding to all individuals equally, due to the broad diversity of facial characteristics

found in the global human population. This is a commonly recognised problem in FRT and without specific efforts to measure and mitigate algorithmic bias, the problem will persist.

- 1.5 In order to ensure that Onfido undertook its research to measure and mitigate bias in a manner that respects and protects the rights and freedoms of individuals (as far as their privacy is concerned), Onfido applied to enter the ICO's Regulatory Sandbox. Onfido was accepted into the Sandbox on 1 July 2019 and a Senior Case Officer was appointed². The Senior Case Officer attended the offices of Onfido on 30 July 2019 to gain an insight into the organisation and to begin formulating the objectives and tasks of the Sandbox plan with Onfido.
- 1.6 Following the July 2019 visit the ICO and Onfido agreed the following objectives for Onfido's Sandbox engagement, which were centred around the overall aim of measuring and mitigating algorithmic bias in Onfido's FRT in a manner which complied with data protection law:

² Please Note – Onfido's Sandbox participation, and the advice provided by the ICO, focussed on the activities undertaken by Onfido to improve their FRT models, specifically those activities undertaken to measure and mitigate bias in their models. It is noted that some of the advice given to Onfido in this capacity may also be applicable to other processing activities Onfido completes to improve its FRT.

- **Objective 1: Data Labelling** – Machine learning research is an often iterative process of testing and training, which requires multiple curated data sets. When addressing bias in FRT, it is necessary to curate those data sets with labels on the basis of facial diversity, an act which may reveal race and ethnicity in the data set. This objective required work to be undertaken to understand what labels are required, how those labels can be applied (eg manual labelling as compared to an automated data classifier), and if/how those labels change the nature of the personal data (ie would the addition of a label transform personal data [images] into a special category of personal data?).
- **Objective 2: Initial Performance Measurement** – Using a curated and labelled data set, Onfido would measure the performance of its FRT on the different populations in the test data set to determine the extent to which those populations face disparate treatment. In other words, this was to help Onfido understand where algorithmic bias might be present in its technology.
- **Objective 3: Experiments using Different Approaches** – Once Onfido understood where bias was present in its technology, Onfido was to test different scientific research methodologies to determine which, if any, best mitigate the previously detected bias.
- **Objective 4: Measure Improvements** – Once Onfido implemented a chosen research methodology to mitigate the measured algorithmic bias, Onfido would use a different curated and labelled data set to again test the performance of its FRT against the different populations in that data set. This was necessary to understand to what degree the research was successful.
- **Objective 5: Ongoing Monitoring and Continued Research** – The above objectives 1 through 4 are not objectives that are completed only once during the duration of the Sandbox. Machine learning research is an inherently iterative process which requires ongoing monitoring and maintenance to ensure the algorithms continue to perform at the necessary levels of performance. Onfido was to implement a process to continuously monitor and mitigate bias in its FRT using the above steps.

- 1.7 The content of the Sandbox plan was agreed by Onfido's Director of Privacy (Neal Cohen) and approved by the Sandbox commissioning and advisory group on 09 September 2019.
- 1.8 During Onfido's time in the Sandbox it became apparent that, due to the sometimes complex nature of AI supply chains involved with processing and sharing of personal data, the focus of Onfido's Sandbox participation had to evolve to include a pragmatic discussion of several key data protection issues which were not originally envisaged at the time when Onfido's Sandbox plan was produced³.
- 1.9 In August 2020 Onfido and the ICO completed the last piece of work detailed in Onfido's Sandbox plan, bringing Onfido's participation in the ICO's Regulatory Sandbox to an end.

³ Further information on these key data protection issues is provided in the 'Executive Summary' and 'Key Data Protection Considerations' sections of this document.

2. Executive summary

- 2.1 As detailed above, Onfido provides identity verification services on behalf of its business clients. To improve its services and ensure that their FRT is fair and inclusive for all data subjects, Onfido has explored the use of different methods of training their its FRT in an attempt to mitigate any bias present. Onfido entered the ICO's Regulatory Sandbox with the aim of ensuring that the research conducted in relation to improving algorithmic bias was carried out in a manner which respected the rights and freedoms of individuals when processing their personal data.
- 2.2 Since September 2019, Onfido aimed to identify any bias present in its FRT model and tested different approaches/training techniques to mitigate said bias. Onfido published a [research paper detailing its findings](#) in February 2020.
- 2.3 In line with the objectives outlined in its Sandbox Plan, Onfido's research team has researched and developed methodologies to group and label data, tested the performance of Onfido's FRT, retrained the FRT's models, and measured the performance changes to those models. This has been an iterative process for Onfido and will continue to be so in the future. This work is further discussed in detail in the aforementioned paper. However, the overall results of Onfido's experimentation appear successful.

2.4 During Onfido's participation in the Sandbox, the ICO and Onfido discussed the following key data protection issues:

- **Is Onfido a controller for the personal data it processes when improving their FRT models?**

The ICO observed that while Onfido's clients will ultimately benefit from Onfido's processing of personal data for the purpose of developing and improving its FRT technology, the discretion and control which Onfido exercises over the technical means used to aggregate data across clients and develop its technology, elevates Onfido to the role of controller when processing to improve its FRT models.

- **Which GDPR Article 6 lawful basis for processing would likely be the most appropriate for Onfido to rely on when processing personal data to develop its identity verification services?**

Based on the information provided by Onfido during its engagement in the Sandbox and as described in this Exit Report, the ICO advised that it appeared that "Legitimate Interests of the Data Controller" (GDPR article 6(1)(f)) may be the most appropriate Article 6 lawful basis for Onfido to process personal data for the purpose of developing its identity verification services. Onfido was further advised that reliance on the legitimate interest lawful basis would

require Onfido to complete a three part legitimate interest assessment, including a balancing test that shows that individuals' interests to not override Onfido's legitimate interests. This can form part of a data protection impact assessment, which the ICO understands Onfido has completed.

- **Is Onfido processing any special categories of personal data as defined by Article 9(1) of the GDPR when developing its identity verification services?**

The ICO worked with Onfido on this question and, after discussing the nature of the biometric data and the labelling methodology Onfido used during its research project⁴, the ICO communicated two key findings. First, the ICO was of the view that it was unlikely that the biometric data processed by Onfido for developing its FRT is special category personal data because although it might allow for the unique identification of individuals, it is not being used for that purpose. Second, the ICO considered that, to the extent research labels related to the race or ethnicity of an individual, those labels were a special category of personal data, as they revealed or inferred the racial/ethnic origins of data subjects.

⁴ Please note – further information on Onfido's labelling methodology is provided in section 4 of this document.

- **If the data Onfido processes to develop its identity services is special category personal data, which GDPR Article 9 condition for processing could Onfido look to apply to process the data?**

The ICO advised Onfido that it appeared likely that the most appropriate Article 9 condition for Onfido to process personal data revealing perceived racial/ethnic origin data for the purpose of measuring and mitigating algorithmic bias would be substantial public interest (GDPR Article 9(2)(g)), on the basis of the equality of opportunity or treatment condition set out in Schedule 1, Part 2, Paragraph 8 of the Data Protection Act 2018. This reflects the substantial public interest in preventing discrimination, in line with Article 14 of the European Convention on Human Rights (ie the Prohibition of Discrimination). Onfido should document how they satisfy each of the specific requirements of Schedule 1 paragraph 8, and also need an Appropriate Policy Document briefly outlining its compliance measures and retention policies for the special category data⁵.

- **How should Onfido look to provide data subjects with privacy information about its activities when developing its identity verification services?**

⁵ Please note - Onfido's own retention practices are explored in footnote corresponding to paragraph 3.6 of this document.

After multiple workshops with Onfido on this topic, the ICO advised Onfido that it appeared that it had made a substantial effort to provide data subjects with adequate information about how Onfido processes their personal data. There are many different ways through which an organisation may choose to provide privacy information to data subjects. Moving forward, the ICO is developing its thinking on issues surrounding the provision of privacy notices in complex AI supply chains, which should help clarify whether there are any further steps that Onfido should take to ensure the provision of this information. As of the date of this report, no further steps have been identified to Onfido.

- **How should Onfido provide data subjects with rights over their personal data when processing the data to develop their identity verification services?**

The ICO advised Onfido that its existing process, as described in Onfido's privacy policy, likely satisfies its legal obligation in helping data subjects to exercise their rights under the GDPR.

3. Product description

- 3.1 Onfido provides an identity verification service for its business clients utilising a human-in-the-loop AI model. This means that Onfido's machine learning algorithms are supplemented by human analysts, where necessary and where specifically requested by Onfido's business clients.
- 3.2 Identity verification checks undertaken by Onfido include the following two distinct products:
 - I. **Document Check** – Onfido examines a photo of an identity document and provides a report to the business client on the likelihood that the identity document is genuine (eg correct fonts, no digital tampering, etc).
 - II. **Facial Similarity Check** – Onfido uses its facial recognition technology (FRT) to compare the image of a face in a photo of the identity document used in step (I) above, to the image of a selfie (or video of a face) and provides a report to the business client on the likelihood that the two faces belong to the same person and confirms that the

selfie image does not show evidence of spoofing or fraud (eg a person wearing a mask, face morphing, etc). The provided report is used by the client to determine whether the person presenting the identity document actually is the person to which that identity document relates, and whether they are physically present at the time of capture of the selfie image.

- 3.3 Frequently, Onfido's business clients (who are based both inside and outside of the EEA) will collect both the original ID document image and selfie image directly from the data subject and then transfer those images to Onfido so that Onfido may provide the business client with its identity verification services. When this occurs, Onfido has no direct interaction with the data subject. However, some business clients will embed Onfido's image capture interface into their software in an effort to increase the probability that high quality images are collected⁶.

⁶ Please note - even though the image capture interface provided is Onfido's proprietary software, it is often customised in such a way that the image capture interface appears as though it is being offered by Onfido's business clients and not Onfido. Onfido has communicated to the ICO that such customisation is a common market requirement, and while Onfido can contractually require its business clients to provide notice of Onfido's processing, Onfido has limited ability to dictate how its business clients must precisely fulfil that requirement (eg by displaying Onfido's name in its software or by linking directly to Onfido's public privacy policy).

- 3.4 When Onfido completes an identity verification, it does not solely make a yes/no recommendation for its client on whether the data subject is who they claim to be. Rather, Onfido provides a detailed report containing the results of the identity verification, along with a recommendation for the client to consider and decide upon⁷. It is then the client's responsibility to review this report and make its own decision based on all the information provided to them to by a data subject on whether or not to provide the data subject with access to their platform or service- the client's review which often includes the use of services and information that extends beyond Onfido. However, it should be noted that if a data subject is unable to pass the Onfido identity verification, it is highly likely that the client reviewing Onfido's report will place significant weight on Onfido's findings and decide to deny that data subject access to their platform or service. However this should not be done without the organisation first undertaking further action which may include seeking additional information from the data subject.
- 3.5 In order to build Onfido's FRT and ensure that it performs at sufficient levels of accuracy and reliability, Onfido trains its FRT models using the images collected during the provision of its services to its business clients. This is necessary as Onfido's FRT is used to solve a unique technical challenge (comparing a selfie to the face in a photo of an identity document typically

⁷ These results are detailed in Onfido's public API documentation, available at <https://documentation.onfido.com/>.

taken with a mobile device). By training its FRT on data that is representative of the technical challenge seen when Onfido is providing its service to clients, Onfido is able to achieve higher levels of performance in its technology, and those performance levels will continue to increase over time as Onfido improves upon and retrains its algorithms using more representative data.

3.6 However, Onfido only repurposes personal data (received from its clients) to develop its technology where Onfido's business clients have:

- agreed to Onfido using personal data for such purpose; and
- warranted that they have provided all necessary fair processing notices, and taken all other actions necessary for Onfido to lawfully process the personal data⁸.

⁸ Please note – Onfido only retains data as a controller for a limited amount of time. Onfido's research team uses a business intelligence system to select the most relevant data required to train a specific machine learning model and then migrates that data to a secure research silo in their cloud based servers where training/research activity is undertaken.

- 3.7 The processing described above is all part of Onfido's business as usual activities. However, for Onfido's participation in the Sandbox, Onfido wanted to expand upon the training element of their processing (which aims to improve the accuracy and reliability of its FRT generally) by undertaking a research project with the aim of ensuring that their FRT is also fair and inclusive for all individuals, thus enabling individuals access to Onfido's clients services in all cases where access should be granted.
- 3.8 In order to achieve this goal, Onfido wanted to first determine what, if any, algorithmic bias was already present in its FRT and then, if bias was found to be present, test several training techniques to try and improve the fairness and inclusivity of their FRT while implementing a system to continuously monitor and mitigate this issue on an on-going basis.

4. Key data protection considerations

- 4.1 Onfido's Sandbox participation originally focused on Onfido's research into methods to detect and mitigate any apparent algorithmic bias within its FRT models. In order to progress this work, the ICO and Onfido agreed the following objectives in Onfido's Sandbox Plan:
- **Objective 1: Data Labelling** – Machine learning research is an often iterative process of testing and training, which requires multiple curated data sets. When addressing bias in FRT, it is necessary to curate those data sets with labels on the basis of facial diversity, an act which may reveal race and ethnicity in the data set. Work was to be undertaken to understand what labels are required, how those labels can be applied (eg manual labelling as compared to an automated data classifier), and if/how those labels change the nature of the personal data (ie would the addition of a label transform personal data [images] into a special category of personal data?).

- **Objective 2: Initial Performance Measurement** – Using a curated and labelled data set, Onfido would measure the performance of its FRT on the different populations in the test data set to determine the extent to which those populations face disparate treatment. In other words, this is to help Onfido understand where algorithmic bias might be present in its technology.
- **Objective 3: Experiments using Different Approaches** – Once Onfido understood where bias was present in its technology, Onfido was to test different scientific research methodologies to determine which, if any, best mitigate the previously detected bias.
- **Objective 4: Measure Improvements** – Once Onfido implemented a research methodology to mitigate the measured algorithmic bias, Onfido would use a different curated and labelled data set to again test the performance of its FRT against the different populations in that data set. This is necessary to understand whether the research was successful and, if so, to what degree.
- **Objective 5: Ongoing Monitoring and Continued Research** – The above objectives 1 through 4 are not objectives that are completed only once during the duration of the Sandbox. Machine learning research is an inherently iterative process which requires ongoing monitoring and maintenance to ensure the algorithms continue to operate at the necessary levels of performance. Onfido was therefore asked to implement a process to continuously monitor and mitigate bias in its FRT based on the steps set out in the above objectives.

4.2 However, it became apparent that there are key data protection issues with the construction of many business to business to consumer (b2b2c) AI supply chains, which were highlighted by Onfido. While these issues reached beyond the scope of measuring and mitigating bias in FRT, the ICO deemed it necessary to expand the focus of Onfido's Sandbox participation to include discussions around these key issues (outlined below) while Onfido undertook their research into algorithmic bias in its FRT.

4.3 The first issue the ICO worked with Onfido to resolve during a workshop at the ICO's offices in December 2019 was the question of data controllership in Onfido's supply chain. The ICO observed that AI suppliers are, by default, typically

considered processors by their clients, who consider themselves controllers. This determination in many cases is correct as usually an AI supplier processes data exclusively on the behalf of their client and in accordance with their instructions. However, in a situation where an AI supplier, like Onfido, processes data from multiple clients and with the permission of those clients, uses the data obtained from them to train and generally improve their technology, the ICO determined that it is likely that such AI suppliers are crossing, at least temporarily while developing their technology, into the role of a controller. In Onfido's case, even though its clients will benefit from Onfido processing personal data for the purpose of developing the technology, the discretion and control which Onfido has in terms of determining the technical means used to aggregate data across clients and develop its technology, elevates Onfido to the role of controller.

- 4.4 Based on the above assessment, the ICO determined that it was likely that Onfido should consider itself a processor in the context of delivering a service to its clients (ie completing identity verification checks on the behalf of its business clients), but, when processing personal data to train its FRT model(s), such as the processing activities planned to be undertaken during Onfido's participation in the Sandbox (as well as more generally to improve the accuracy and performance of its FRT outside of the Sandbox), it is likely that Onfido should consider itself a controller for the data used for this purpose, due to the discretion Onfido exercises over the research methodology.
- 4.5 After reaching the view that Onfido should be considered a controller for the data processed for research purposes, the ICO worked with Onfido to answer the following key data protection questions during their time in the Sandbox:

Article 6 considerations

- 4.6 Based on the ICO's understanding of Onfido's business operations and research methodology, the ICO advised Onfido that it appeared likely that GDPR Article 6(1)(f) would be the most appropriate lawful basis for processing for its intended purposes. Article 6(1)(f) states that processing should be:

“necessary for the purposes of the legitimate interests pursued by the controller or by a third party, except where such interests are overridden by the interests or fundamental rights and freedoms of the data subject which require protection of personal data, in particular where the data subject is a child.”

- 4.7 [ICO guidance on the application of the Legitimate Interests basis](#) states that a controller seeking to rely on this basis should complete and keep a record of their three part Legitimate Interest Assessment. It is understood that Onfido has accepted the ICO’s advice and taken the decision to rely on this basis for processing and has completed a DPIA, while factoring in the requirements of a Legitimate Interests Assessment, for its processing activity as a controller.
- 4.8 In completing the DPIA, Onfido sought to balance the need to use large data sets for the purpose of developing the FRT against the impact that such data usage might have on a relevant data subject. It was identified that while the envisioned data usage is not intended to directly impact a specific data subject, if that data subject is later asked to undergo another Onfido identity verification check in the future, they may actually be impacted by the fact that their data was used for research purposes as the algorithms that were developed will be applied to them. In addition, there is a risk that the additional processing by its very nature, presents a security risk to the individual. If data is not retained and processed, it cannot be lost or stolen. For this reason, Onfido has put great effort into developing a secure research silo to carry out its research activities and never retains data for research purposes beyond the time specified by its clients for the clients’ own purposes. Lastly, data subjects who do not want to have their personal data used for research purposes may object to the processing, and the client can delete the data from Onfido’s systems.

Is Onfido processing any special categories of personal data?

- 4.9 When processing personal data for the purposes of developing its identity verification services, with the permission of its clients, Onfido repurposes the images collected on behalf of its clients, in relation to training its FRT models (including by conducting research on new methodologies to train its FRT models to mitigate algorithmic bias). However, after examining Onfido’s training methodologies, it became apparent to the ICO that, for the purposes of its training and research activities,

Onfido does not appear to use the biometric data generated during such activities for the purpose of uniquely identifying individuals. Consequently, that biometric data should not be regarded as special category biometric data⁹.

- 4.10 As part of its Sandbox research project, which focused on examining methods through which Onfido could look to measure and mitigate algorithmic bias in its FRT, Onfido needed to continuously test the performance of its FRT against populations with different facial characteristics. In order to do this, Onfido needed to define data classifications (or “cohorts”) and then ‘label’ images with the most relevant classification. After some consideration, Onfido decided to undertake its initial research by applying these labels based on the issuing country of the individuals’ identity documents (eg an individual who underwent an Onfido identity verification check and had their selfie examined against an identity document issued by the French government would be labelled French/European and so on for other countries). However, the ICO thought it likely that these labels would reveal the racial and/or ethnic origin of the individuals whose data was used and as such should be regarded as special category data.

⁹ Please note – When processing biometric data as part of its business as usual activities (ie for the purposes of completing identity verifications) the unique identification of individuals is a requirement of the identity verification and as such this data, which Onfido processes as a processor, should be regarded as special category biometric data and will require the controller (ie Onfido’s clients) to apply their own Article 9,1 condition for processing the special category data.

4.11 Moving forward, Onfido has indicated that it intends to use more granular data classifications/cohorts to better refine its FRT and mitigate algorithmic bias. As such, Onfido will need to continuously consider and evaluate to what extent it may be processing special category personal data with these newly defined cohorts.

Article 9 considerations

4.12 Based on the ICO's understanding of the research Onfido was undertaking while in the Sandbox it advised Onfido that it's possible for Onfido to use the substantial public interest condition Article 9(2)(g) of the GDPR in order to process special category information on the perceived racial/ethnic origin of individual data subjects as part of its research project. Article 9(2)(g) states that the processing should be:

“necessary for reasons of substantial public interest, on the basis of Union or Member State law which shall be proportionate to the aim pursued, respect the essence of the right to data protection and provide for suitable and specific measures to safeguard the fundamental rights and the interests of the data subject”

4.13 In the case of Onfido's research into the measuring and mitigation of algorithmic bias, the substantial public interest condition they identified in law was the condition for equality of opportunity set out in Schedule 1, Part 2, Paragraph 8 of the Data Protection Act 2018 (ie that the processing is necessary for the purposes of identifying or keeping under review the existence or absence of equality of opportunity or treatment between groups of people specified in relation to that category with a view to enabling such equality to be promoted or maintained). This reflects the substantial public interest in Article 14 of the European Convention on Human Rights (Prohibition of Discrimination). This condition can only be used if processing is not carried out for the purposes of measures or decisions with respect to a particular individual, is not likely to cause substantial damage or distress, and the individual has not objected. Schedule 1, paragraph 5 of the Data Protection Act 2018 also requires Onfido to have an 'appropriate policy document' in place, briefly outlining its compliance measures and retention policies for the special category data.

How should Onfido look to provide data subjects with privacy information about its activities when developing its identity verification services?

4.14 As previously stated in this document, the majority of Onfido identity verifications do not require Onfido as an organisation to interact with the data subject whose data it is processing¹⁰. However, many of Onfido's business clients do use Onfido's image capture interface to collect the image of an identity document and selfie directly from the data subject. While this interface is embedded into a business client's software, it is often customised by the client in a fashion whereby Onfido is not visible to the end customer¹¹, and Onfido has limited ability to use that image capture screen to provide fair notice of processing.

¹⁰ This means that Onfido does not necessarily collect the data directly from the data subject, and Onfido does not have reasonable means of contacting the data subject. In most cases, Onfido does not have the email address or other necessary contact information of the data subject (it should be noted that some of the identity documents utilised by Onfido to complete its verification checks, such as driving licences, may contain contact information such as addresses which could, if extracted from the document, be used to contact the data subject by post).

¹¹ As described in the footnote attached to section 3.2. of this document.

- 4.15 In response to the challenges described above and prior to its time in the Sandbox, Onfido had contractually obligated its business clients to provide data subjects with privacy information about Onfido's processing of personal data to develop its identity verification services.
- 4.16 At the time this report was written the ICO is confident that Onfido is currently making a substantial effort to provide this information to data subjects. Currently the ICO understands that Onfido is providing this information in the following ways:
- Contractually obliging its clients to provide privacy information to individuals and seeking a warranty from the client that it has provided such information to the relevant data subjects (as noted above in paragraph 4.15);
 - Interspersing reminders to clients to provide the information throughout Onfido's technical documentation (which is publicly available and provided to Onfido's clients during the on-boarding process); and
 - Displaying Onfido's own privacy notice prominently on its website.
- 4.17 Onfido also indicated that it already has in place a Know Your Business ("KYB") process that is used to indicate the likelihood that a business client will be able to make payments and not bring Onfido into disrepute by abusing the services. Onfido is also looking to expand this process to include checks on whether Onfido's business client is likely to fulfil the privacy obligations placed on them by Onfido. Onfido has indicated that such due diligence checks could include verifying that the business client has a current privacy policy designed to comply with the GDPR.
- 4.18 The ICO acknowledges the effort Onfido is making to provide adequate privacy information when processing personal data in its capacity as a controller. Moving forward, the ICO is committed to developing its thinking on issues surrounding the provision of privacy notices in complex AI supply chains to further clarify how controllers involved in such supply chains should ensure the provision of privacy information for individuals. This should help to clarify whether there are any further steps Onfido should take to ensure the provision of privacy information for individuals beyond the measures they already have in place. As of the date of this report, no further steps have been identified to Onfido.

How should Onfido provide data subjects with rights over their personal data?

- 4.19 As previously noted in paragraph 4.4 of this document, Onfido, as an AI service provider, is either a processor or a controller. This classification is dependent upon the processing activity being undertaken – in Onfido’s case it acts as a processor when involved in the provision of identity verification checks on behalf of its clients, as compared to its role as a controller when processing personal data for the purposes of the development of the technology used to power those identity verification checks. As Onfido does not maintain a direct relationship with data subjects, this presents Onfido with unique challenges in facilitating data subject rights requests.
- 4.20 During Onfido’s time in the Sandbox it outlined its current process for dealing with data subject rights requests which involves:
- I. asking the data subject to identify themselves;
 - II. Onfido identifying the relevant business client(s) whom Onfido was acting on behalf of when performing the data subject’s identity verification check; and

III. putting the data subject in contact with those business client(s), so they may handle the data subject request¹².

4.21 The ICO reached the view that this process, which is also described in Onfido's public privacy policy, likely satisfies its legal obligations in helping data subjects to exercise their rights under the GDPR. Furthermore, the ICO understands that Onfido is currently investigating, outside of its participation in the Sandbox, the development of tools to better provide data subjects with methods to expedite and improve the efficiency of making a rights request through Onfido.

Mitigating bias in FRT models

4.22 During Onfido's time in the Sandbox, the research team at Onfido has worked to measure and mitigate bias present in the FRT while testing different approaches and training techniques. This team published a detailed [research paper on mitigating algorithmic bias](#) in February 2020, and the team is continuing to iterate and improve upon its work to further measure and mitigate any residual bias or new bias that might appear.

¹² Please Note: As Onfido only act as a controller of personal data for a limited amount of time it is unlikely that it would retain data for long enough as a controller to fulfil a data subject's rights request itself. In the case of Erasure and Objection to processing requests, Onfido relies on its clients to acknowledge the data subject's request and utilise client user controls to ensure Onfido does not utilise their data as a controller.

- 4.23 In line with the objectives outlined in its Sandbox Plan, Onfido's research team has researched and developed methodologies to group and label data, test the performance of Onfido's FRT, retrain the FRT's models, and measure the performance changes to those models. This has been an iterative process for Onfido and will continue to be so in the future. This work is further discussed in detail in the aforementioned paper.
- 4.24 Onfido has indicated that the results of its experimentation appear successful and, moving forward, Onfido will continue to use a mixture of the research methods outlined in its research paper when retraining its FRT models to ensure individuals from all backgrounds are able to access the services of Onfido's clients which rely upon Onfido's identity verification technology.

5. Ending statement

- 5.1 Onfido's participation in the ICO's Regulatory Sandbox has given the ICO the opportunity to gain valuable insight into the practical issues relating to a supply chain involved with the provision of AI services, particularly where the application of data protection legislation and guidance to such an environment is not always clear. It is hoped that the ICO's engagement with Onfido will allow the ICO to further develop its thinking on data protection based issues surrounding complex AI supply chains. This is an important step towards better clarifying how controllers and processors involved in such AI supply chains can ensure their compliance with UK data protection legislation. Ultimately, such clarity should result in privacy benefits to the data subject.

5.2 It is clear from our work with Onfido that the business has a real and demonstrated commitment to making use of innovative technology in a compliant way and want to ensure that, as an organisation, it leads the market by ensuring that its FRT product is accessible and useful for individuals from all backgrounds. Through our work, we recognise that there are likely to be a number of challenges faced by any AI service provider as they seek to improve their technologies¹³. These include:

- Updating client contracts to better reflect the sometimes complex nature of data controllership in AI supply chains;
- Developing tools and ways of working with clients which enable the provision of any and all fair processing information for data subjects; and
- Ensuring that suitable processes tools are in place to enable data subjects to exercise their rights under Chapter 3 of

¹³ Tangential to Onfido's work in the Sandbox, Onfido's Director of Privacy published a research paper on the challenges of the AI supply chain through the Carr Center for Human Rights Policy at Harvard University's Kennedy School of Government. See Neal Cohen. 4/4/2020. "[The Ethical Use of Personal Data to Build Artificial Intelligence Technologies: A Case Study on Remote Biometric Identity Verification.](#)" Carr Center Discussion Paper Series, 2020-004.

the GDPR.

- 5.3 The ICO also recognised through their work with Onfido that organisations which utilise the services and products of AI service providers' must work to ensure that sufficient information is provided in their own privacy information to inform individuals about how AI Service Providers utilise their personal data (ie sufficient information is provided to inform an individual that their data may be used by the AI service provider to improve their service).
- 5.4 Based on the information we have reviewed as part of Onfido's Sandbox participation and, solely in respect of the research project investigating methods to measure and mitigate algorithmic bias in Onfido's FRT, it appears likely that Onfido has complied with its obligations as controller under UK data protection legislation while reportedly mitigating the algorithmic bias in its FRT significantly. Moving forwards Onfido should ensure that it continues to follow the steers provided to it in the Sandbox, as well as relevant ICO guidance, with the same diligence and transparency as it has demonstrated during its participation in the Sandbox.
- 5.5 Onfido's participation in the ICO Regulatory Sandbox demonstrates the ICO's role as a trusted partner in its support for stakeholders who ensure that information rights are recognised, respected and designed into the fabric of research projects aiming to improve individuals' ability to access innovative technologies and further demonstrates that UK data protection legislation is not a barrier to innovation.
- 5.6 At the time of writing this report, the Covid-19 public health emergency was unfolding across the world with unprecedented impact on UK industry. We are immensely grateful to Onfido for its engagement in the Sandbox, for its transparency and candour and finalising this report in difficult circumstances.