Cisco Systems, Inc. Response to White House Office of Science and Technology Policy Request for Information (RFI) on Public and Private Sector Uses of Biometric Technologies

January 13, 2022

Introduction

Cisco Systems, Inc. ("Cisco") appreciates the opportunity to provide comments in response to the Office of Science and Technology Policy's Request for Information (RFI) on Public and Private Sector Uses of Biometric Technologies. We welcome OSTP's attention to the role of data-driven technologies in society and we look forward to ongoing engagement with OSTP regarding the use of these technologies in business settings.

Cisco is the worldwide leader in technology that powers the Internet. We deliver innovative software-defined networking, cloud, collaboration, applications, and security solutions across the globe. Our customer base spans large and small organizations across the public and private sectors, including 95% of Fortune 500 companies. Our corporate mission is to empower an inclusive future for all.

Cisco uses artificial intelligence (AI) and machine learning (ML) across our portfolio to deliver high-value capabilities that satisfy our customers' needs. Cisco is committed to responsible development and use of AI/ML and to upholding industry-leading privacy, security, and human rights standards.

In this submission we provide some general observations about the questions raised in the RFI, describe how we use (and do not use) biometric technologies in our product portfolio, address security considerations, and describe the governance frameworks that we have in place.

General observations

Biometric technologies are currently used across a very wide range of settings and contexts, from private use on a personal device to employee use on corporate systems to broad scale use in public settings. Understanding the potential for benefits and harms arising from the use of these technologies requires evaluating the specifics on a case-by-case basis. There are valid and salient concerns about the potential for harm when biometric information is used in certain settings or without appropriate safeguards for efficacy, privacy, security, human rights, and fairness.

We believe that with meaningful safeguards in place, biometric information can be leveraged for certain business use cases to deliver value to individuals and organizations that cannot easily be obtained by other means. Our experience has been that by being thoughtful about the unique properties of technologies like facial recognition, our products can become more inclusive and provide individuals with greater flexibility in how they work, learn, and carry out daily activities.

We also believe that the evolution and uptake of machine learning presents an opportunity to be more explicit and transparent about potential unintended biases that exist in data and algorithms than was previously possible with traditional software systems and business processes. Human judgement carries its own biases that, in traditional systems, were not typically quantified or even revealed. Advances in machine learning present an opportunity to make these biases known and to be deliberate and transparent in choosing bias mitigation measures.

Biometric technologies are not a proper fit for every use case. The decision to leverage biometric information within a particular product or service needs to be taken with diligent attention to the overall system design and with an understanding of the trade-offs between the risks and benefits that biometric information brings.

Responsible uses of biometric technologies

At Cisco, we build our products to be secure by design and private by default. In general, we minimize the amount of personal or sensitive data we collect and retain from our customers and end users to what is strictly required to offer our services, and we seek alternative product designs that avoid data collection where possible. We believe that privacy is a fundamental human right.

In this section we describe three Cisco product suites - Webex, Duo, and Meraki Video – to illustrate the range of approaches we have taken to the collection and use of biometric information. This section is responsive to questions 1, 2, 3, and 5 of the RFI.

Webex

Webex is a purpose-built collaboration suite for hybrid work that supports calling, meetings, messaging, devices, and more. Cisco makes use of facial recognition within our Webex suite to recognize individuals participating in Webex meetings and display their name labels, to provide background blur and background replacement, and to optimize visual layouts on screen. These features are designed to make Webex meetings more inclusive, secure, and personalized, and to power the future of hybrid work.

We use a layered set of techniques and mitigations to ensure privacy, security, and accuracy in our use of facial recognition.[1]

We require customer organizations and their end users to opt in to use facial recognition before it is turned on. Users must follow an intentional, multi-step process that involves approving a face image for upload to our secure Webex cloud. If an organization administrator later disables facial recognition, all associated data for that organization's users is deleted in a timely fashion. If an enrolled user leaves the organization, the user's face images and associated data are deleted.

We never use facial recognition outside of the controlled environment of a Webex meeting, and we never use facial recognition to identify individuals who have not enrolled their facial images. While doing otherwise is technically feasible, we have deliberately chosen to constrain our use of facial recognition only to those users who affirmatively choose to use it to improve their meeting experiences.

The neural network models we use for facial recognition are pre-trained by Cisco without using customer data. We manage our training data sets to ensure that they are

[1] More details about data handling and privacy related to facial recognition in Webex are available in our white paper.

balanced across visual features, reducing data bias. We improve accuracy by using labeled training data where we can associate a training image to a verifiable identity (e.g., where a Cisco employee volunteers to help with training).

Users' enrollment images are never used outside their own organizations. Other than enrollment images, images of users are not stored in the Webex cloud.

Images are not used in real time for facial recognition. Instead, a vector is calculated at enrollment time. This approach is reliable because the computed vector is based on multiple views of the user's face from different angles and with different accessories (e.g., with eyeglasses and without eyeglasses). This vector is used to match the vector of a user constructed when a face is detected during a meeting. This approach reduces algorithmic bias by relying on the enrollment vector, which provides a more accurate match than would be obtainable by doing real-time image matching. It also removes any dependency on the user's image in the operation of Webex meetings.

For some features, we make use of face detection rather than facial recognition. Face detection uses a trained model to detect that a human face is present (or to count the number of faces present) in a video frame without attempting to recognize an individual human. This technology is used in features such as background segmentation to determine which user in a video frame is in the foreground and for features that track the active speaker in a meeting to provide the best view of a speaker, which requires being able to detect heads in 3D rather than 2D face images. Similarly, when we train models to support other video features like the recognition of hand gestures (e.g., to show a "thumbs up" emoji on screen when a meeting attendee gives a thumbs up sign with their hand), we never do facial recognition or link the faces in the video data to identity.

Duo

Duo is a user-friendly, zero trust platform used by organizations of all sizes to secure their users' access to devices and applications.

Duo supports certain biometric authentication factors that are widely available and easy to use. These include Apple's Touch ID and Face ID and Android's fingerprint feature. These features are built into device operating systems and provide both a high level of security as well as privacy protection that prevent larger scale abuse. Biometric information is stored in a secure manner on the device by the native operating system, and the use of WebAuthn[2] technology allows the operating system to confirm with Duo systems that the individual user has authenticated biometrically, confirming their identity, while not conveying any other potentially identifiable information to Duo. When users choose to use these methods of authentication, their biometric data is never shared with Cisco. This combines strong privacy protection for Duo users with the strong security benefits of biometric identification.

Meraki Video

Meraki Video provides enterprise video security with a suite of cloud-managed smart cameras. Schools, retail establishments, and businesses of all kinds use Meraki Video to secure their premises.

Meraki cameras use computer vision to detect people and vehicles, but the cameras never determine individual identity. This allows enterprises to track activity on-location without learning the identities of patrons, employees, students, or visitors. The cameras report when a person appears in a video frame, but not who the person is. Meraki Video does not use facial recognition, gait recognition, or any other kind of biometric identification.

This design was a deliberate choice and is one that bucks the much more common trend in the video surveillance industry of using biometric technologies to identify individuals. Many of our customers consistently request that we add identification capability, but we refuse to do so because we believe that biometric identification in video security solutions has a high probability of misuse, and the difficulties of doing it accurately can prevent it from being used safely. For example, available research demonstrates that identification of individuals via facial recognition in public spaces is extremely difficult to do accurately without significant control over the physical environment or cooperation from the individuals to be identified.[3]

Even without doing biometric identification, we know that recorded video is sensitive data. For that reason, we store all recorded video encrypted on the cameras themselves rather than in the cloud. Processing of video for the purpose of object

[2] See WebAuthn.

[3] See NISTIR 8173, Face In Video Evaluation (FIVE) Face Recognition of Non-Cooperative Subjects.

recognition takes place exclusively on the cameras. We were the first in the industry to bring this innovative, privacy-preserving design to market.

Cisco has no access to stored customer video unless the customer explicitly authorizes access. When these accesses are authorized, for example during a customer support call, they are logged in a customer-viewable log so that the customer has a record of each access event.

This authorization scheme extends to customer video provided to selected teams within Cisco to train the object classification model we use in the cameras. Our customers have an option of contributing their video data to help us improve object classification (including detecting the presence of a person but not the person's identity in a video frame). Customers have complete control over which of their cameras contribute video to our training data. By default, no customer video data is shared with Cisco. If they opt in to contribute training data, customers can review which video clips we use for training. They can delete any of their data or opt out at any time. This is yet another way in which Cisco differentiates from the rest of the industry: by providing a robust set of privacy controls around customer-provided training data, rather than simply using all customer data for training as a condition of using our cameras.

Summary

As these three products demonstrate, we are deliberate in our decision-making about when the use of biometric technologies is appropriate to meet customer needs. In the case of Meraki Video, we have chosen not to use biometric identification – and foregone potential additional revenue - because it is neither safe nor accurate. In the controlled environments of Webex meetings and Duo authentication, we enable biometrics-based features when our customers and users choose them. In all cases, we use layered safeguards to mitigate inaccuracy, bias, and privacy and security threats. We never monetize biometric information.

Security considerations (RFI question 3)

When biometric technologies are added to a product in our portfolio, they become embedded in Cisco's overall security model and the Cisco Secure Development Lifecycle (CSDL), a repeatable and measurable process that is unified across all solutions and services we offer.[4] This means that biometrics-based features are developed using secure tools and processes and that they benefit from the same physical, infrastructure, platform, and application security architectures and controls as the products in which they are embedded.

In the case of Webex, the neural network models we use for facial recognition are trained on data that we fully control, are rigorously tested for accuracy and bias, and are embedded internally in the Webex service. They are thereby defended against data poisoning attacks, inference attacks, and many of the other kinds of adversarial attacks that may target ML models used by other parties or in public-facing settings. Face images and vectors are encrypted in transit and at rest, with role-based access control and segregation of duties controlling Cisco's access to this data. Face images cannot be reverse-engineered by accessing the vectors.[5]

The next section discusses Cisco's Responsible AI/ML Framework, which includes specific requirements applied to our engineering, development, and deployment processes for all AI/ML capabilities, including AI-powered biometric technologies.

Governance programs, practices, and procedures (RFI question 6)

At Cisco, we rely on broad-based governance frameworks to ensure that all of our products - not just those involving AI, ML or biometric technologies – are designed to protect privacy, safeguard security, and respect human rights. These frameworks are based on industry standards and best practices, but informed by our own unique experiences, customer set, and portfolio. We address the kinds of concerns raised in the RFI through the joint application of our Cisco Secure Development Lifecycle (described above), our Responsible AI/ML Framework, our Business and Human

[4] See the Cisco Secure Development Lifecycle Overview.

[5] For more information about the Webex security architecture, see the Cisco Webex Meetings Security White Paper.

Rights Program, and our Global Privacy Program. These programs need not be specifically tailored to biometric technologies to collectively provide the governance framework needed to establish trust in our development and use of these technologies. Global consistency is key: we rely on these programs to ensure compliance in the 170+ countries in which we operate.

Our Responsible AI/ML Framework, which is perhaps most directly applicable to the subject of the RFI, includes three key components:

- Design requirements. All Cisco products that incorporate AI or ML must comply with baseline requirements relating to model definition; data quality, relevance, licensing, attribution, and unintended bias mitigation; model monitoring and protection from attacks; user consent; fairness; and model documentation. Models that are directly involved in decisions that could have a legal or human rights impact on individuals or groups are subjected to an in-depth responsible AI/ML assessment prior to being brought to market. The assessment serves as a gating function that prevents models from being deployed until the potential for negative legal or human rights impact is minimized.

- Incident response. Our Responsible AI/ML Incident Response Team may receive reports of unfair, biased, or discriminatory decisions powered by AI or ML in our products. When we receive these reports from customers, employees, or partners, our Responsible AI/ML Incident Response Team analyzes the reports and engages the appropriate internal team to resolve the issue. Once the issue is resolved, we may report back to the original submitter or a broader group of Cisco customers, employees, and partners on the findings of the investigation and remediation steps taken.

- Oversight. Cisco has established an internal Responsible AI/ML Committee that consists of senior executives from across our lines of business, sales, privacy, security, human rights, legal, government affairs, and other functions. The committee is tasked with reviewing sensitive or high-risk uses of AI and ML being proposed by our business units, reviewing incident reports of bias or discrimination, advising Cisco's leadership and employees on responsible AI/ML practices, and overseeing the adoption of our overall Responsible AI/ML Framework.

Through Cisco's Business and Human Rights (BHR) program, we work to identify potential human rights issues related to our supply chain, product design and use, and business relationships. The BHR team's purpose is to help prevent and mitigate harms from occurring, and to advise the business on strategies to respond if they do materialize. The BHR team works across functions to make these strategies standard practice by incorporating them into policies and decision-making frameworks. As an internal clearinghouse for human rights matters, our dedicated human rights experts answer questions, conduct due diligence to inform business decisions and product development, and train employees.

Finally, Cisco has implemented a Global Privacy Program to address risks and maintain high standards for processing personal data. The privacy program is composed of numerous components. We inventory and map the data we process and document the results in Privacy Data Sheets and Privacy Data Maps.[6] We conduct privacy impact assessments and customized risk assessments when developing new products that handle personal data. Our development process embeds privacy by design, ensures compliance with the 120+ privacy laws around the world where we operate, and provides customers with tools to comply with their own privacy requirements. We have integrated privacy incidents into our incident response process and privacy engineering methodologies into the CSDL. Our privacy office oversees all components of the program, analyzing regular reporting of relevant metrics and risk reviews, and conducting internal and external audits.

Conclusion

We believe that our approach to using biometric technologies provides compelling evidence of how these technologies can be responsibly deployed in enterprise settings to deliver value to customers and end users without compromising privacy, security, or human rights. We look forward to further engagement with OSTP on this important topic.

[6] See Privacy Data Sheets and Privacy Data Maps.