

FactSheets: Increasing Trust in AI Services through Supplier’s Declarations of Conformity

M. Arnold,¹ R. K. E. Bellamy,¹ M. Hind,¹ S. Houde,¹ S. Mehta,² A. Mojsilović,¹
R. Nair,¹ K. Natesan Ramamurthy,¹ D. Reimer,¹ A. Olteanu,* D. Piorkowski,¹
J. Tsay,¹ and K. R. Varshney¹

IBM Research

¹Yorktown Heights, New York, ²Bengaluru, Karnataka

Abstract

Accuracy is an important concern for suppliers of artificial intelligence (AI) services, but considerations beyond accuracy, such as safety (which includes fairness and explainability), security, and provenance, are also critical elements to engender consumers’ trust in a service. Many industries use transparent, standardized, but often not legally required documents called supplier’s declarations of conformity (SDoCs) to describe the lineage of a product along with the safety and performance testing it has undergone. SDoCs may be considered multi-dimensional fact sheets that capture and quantify various aspects of the product and its development to make it worthy of consumers’ trust. Inspired by this practice, we propose FactSheets to help increase trust in AI services. We envision such documents to contain purpose, performance, safety, security, and provenance information to be completed by AI service providers for examination by consumers. We suggest a comprehensive set of declaration items tailored to AI and provide examples for two fictitious AI services in the appendix of the paper.

1 Introduction

Artificial intelligence (AI) services, such as those containing predictive models trained through machine learning, are increasingly key pieces of products and decision-making workflows. A service is a function or application accessed by a customer via a cloud infrastructure, typically by means of an application programming interface (API). For example, an AI ser-

vice could take an audio waveform as input and return a transcript of what was spoken as output, with all complexity hidden from the user, all computation done in the cloud, and all models used to produce the output pre-trained by the supplier of the service. A second more complex example would provide an audio waveform translated into a different language as output. The second example illustrates that a service can be made up of many different models (speech recognition, language translation, possibly sentiment or tone analysis, and speech synthesis) and is thus a distinct concept from a single pre-trained machine learning model or library.

In many different application domains today, AI services are achieving impressive accuracy. In certain areas, high accuracy alone may be sufficient, but deployments of AI in high-stakes decisions, such as credit applications, judicial decisions, and medical recommendations, require greater trust in AI services. Although there is no scholarly consensus on the specific traits that imbue trustworthiness in people or algorithms [1, 2], fairness, explainability, general safety, security, and transparency are some of the issues that have raised public concern about trusting AI and threatened the further adoption of AI beyond low-stakes uses [3, 4]. Despite active research and development to address these issues, there is no mechanism yet for the creator of an AI service to communicate how they are addressed in a deployed version. This is a major impediment to broad AI adoption.

Toward transparency for developing trust, we propose a *FactSheet* for AI Services. A FactSheet will contain sections on all relevant attributes of an AI service, such as intended use, performance, safety, and security. Performance will include appropriate accuracy or risk measures along with timing information. Safety, discussed in [5, 3] as the minimiza-

*A. Olteanu’s work was done while at IBM Research. Author is currently affiliated with Microsoft Research.

tion of both risk and epistemic uncertainty, will include explainability, algorithmic fairness, and robustness to dataset shift. Security will include robustness to adversarial attacks. Moreover, the FactSheet will list how the service was created, trained, and deployed along with what scenarios it was tested on, how it may respond to untested scenarios, guidelines that specify what tasks it should and should not be used for, and any ethical concerns of its use. Hence, FactSheets help prevent overgeneralization and unintended use of AI services by solidly grounding them with metrics and usage scenarios.

A FactSheet is modeled after a *supplier’s declaration of conformity* (SDoC). An SDoC is a document to “show that a product, process or service conforms to a standard or technical regulation, in which a supplier provides written assurance [and evidence] of conformity to the specified requirements,” and is used in many different industries and sectors including telecommunications and transportation [6]. Importantly, SDoCs are often voluntary and tests reported in SDoCs are conducted by the supplier itself rather than by third parties [7]. This distinguishes self-declarations from certifications that are mandatory and must have tests conducted by third parties. We propose that FactSheets for AI services be voluntary initially; we provide further discussion on their possible evolution in later sections.

Our proposal of AI service FactSheets is inspired by, and builds upon, recent work that focuses on increased transparency for datasets [8, 9, 10] and models [11], but is distinguished from these in that we focus on the final AI service. We take this focus for three reasons:

1. AI services constitute the building blocks for many AI applications. Developers will query the service API and consume its output. An AI service can be an amalgam of many models trained on many datasets. Thus, the models and datasets are (direct and indirect) components of an AI service, but they are not the interface to the developer.
2. Often, there is an expertise gap between the producer and consumer of an AI service. The production team relies heavily on the training and creation of one or more AI models and hence will mostly contain data scientists. The consumers of the service tend to be developers. When such an expertise gap exists, it becomes more crucial to communicate the attributes of the artifact in a

standardized way, as with Energy Star or food nutrition labels.

3. Systems composed of safe components may be unsafe and, conversely, it may be possible to build safe systems out of unsafe components, so it is prudent to also consider transparency and accountability of services in addition to datasets and models. In doing so, we take a functional perspective on the overall service, and can test for performance, safety, and security aspects that are not relevant for a dataset in isolation, such as generalization accuracy, explainability, and adversarial robustness.

Loukides et al. propose a checklist that has some of the elements we seek [12].

Our aim is not to give the final word on the contents of AI service FactSheets, but to begin the conversation on the types of information and tests that may be included. Moreover, determining a single comprehensive set of FactSheet items is likely infeasible as the context and industry domain will often determine what items are needed. One would expect higher stakes applications will require more comprehensive FactSheets. Our main goal is to help identify a common set of properties. A multi-stakeholder approach, including numerous AI service suppliers and consumers, standards bodies, and civil society and professional organizations is essential to converge onto standards. It will only be then that we as a community will be able to start producing meaningful FactSheets for AI services.

The remainder of the paper is organized as follows. Section 2 overviews related work, including labeling, safety, and certification standards in other industries. Section 3 provides more details on the key issues to enable trust in AI systems. Section 4 describes the AI service FactSheet in more detail, giving examples of questions that it should include. In Section 5, we discuss how FactSheets can evolve from a voluntary process to one that could be an industry requirement. Section 6 covers challenges, opportunities, and future work needed to achieve the widespread usage of AI service declarations of conformity. A proposed complete set of sections and items for a FactSheet is included in the appendix, along with sample FactSheets for two exemplary fictitious services, fingerprint verification and trending topics in social media.

2 Related Work

This section discusses related work in providing transparency in the creation of AI services, as well as a brief survey of ensuring trust in non-AI systems.

2.1 Transparency in AI

Within the last year, several research groups have advocated standardizing and sharing information about training datasets and trained models. Gebru et al. proposed the use of *datasheets for datasets* as a way to expose and standardize information about public datasets, or datasets used in the development of commercial AI services and pre-trained models [8]. The datasheet would include provenance information, key characteristics, and relevant regulations, but also significant, yet more subjective information, such as potential bias, strengths and weaknesses, and suggested uses. Bender and Friedman propose a *data statement schema*, as a way to capture and convey the information and properties of a dataset used in natural language processing (NLP) research and development [9]. They argue that data statements should be included in most writing on NLP, including: papers presenting new datasets, papers reporting experimental work with datasets, and documentation for NLP systems.

Holland et al. outline the *dataset nutrition label*, a diagnostic framework that provides a concise yet robust and standardized view of the core components of a dataset [10]. Academic conferences such as the International AAAI Conference on Web and Social Media are also starting special tracks for dataset papers containing detailed descriptions, collection methods, and use cases.

Subsequent to the first posting of this paper [13], Mitchell et al. propose *model cards* to convey information that characterizes the evaluation of a machine learning model in a variety of conditions and disclose the context in which models are intended to be used, details of the performance evaluation procedures, and other relevant information [11]. There is also budding activity on auditing and labeling algorithms for accuracy, bias, consistency, transparency, fairness and timeliness, in the industry [14, 15], but this audit does not cover several aspects of safety, security, and lineage.

Our proposal is distinguished from prior work in that we focus on the final AI service, a distinct concept from a single pre-trained machine learning model or dataset. Moreover, we take a broader view on

trustworthy AI that extends beyond principles, values and ethical purpose to also include technical robustness and reliability [16].

2.2 Enabling Trust in Other Domains

Enabling trust in systems is not unique to AI. This section provides an overview of mechanisms used in other domains and industries to achieve trust. The goal is to understand existing approaches to help inspire the right directions for enabling trust in AI services.

2.2.1 Standards Organizations

Standardization organizations, such as the IEEE [17] and ISO [18], define standards along with the requirements that need to be satisfied for a product or a process to meet the standard. The product developer can self-report that a product meets the standard, though there are several cases, especially with ISO standards, where an independent accredited body will verify that the standards are met and provide the certification.

2.2.2 Consumer Products

The United States Consumer Product Safety Commission (CPSC) [19] requires a manufacturer or importer to declare its product as compliant with applicable consumer product safety requirements in a written or electronic declaration of conformity. In many cases, this can be self-reported by the manufacturer or importer, i.e. an SDoC. However, in the case of children’s products, it is mandatory to have the testing performed by a CPSC-accepted laboratory for compliance. Durable infant or toddler products must be marked with specialized tracking labels and must have a postage-paid customer registration card attached, to be used in case of a recall.

The National Parenting Center has a Seal of Approval program [20] that conducts testing on a variety of children’s products involving interaction with the products by parents, children, and educators, who fill out questionnaires for the products they test. The quality of a product is determined based on factors like the product’s level of desirability, sturdiness, and interactive stimulation. Both statistical averaging as well as comments from testers are examined before providing a Seal of Approval for the product.

2.2.3 Finance

In the financial industry, corporate bonds are rated by independent rating services [21, 22] to help an investor assess the bond issuer’s financial strength or its ability to pay a bond’s principal and interest in a timely fashion. These letter-grade ratings range from AAA or Aaa for safe, ‘blue-chip’ bonds to C or D for ‘junk’ bonds. On the other hand, common-stock investments are not rated independently. Rather, the Securities and Exchange Commission (SEC) requires potential issuers of stock to submit specific registration documents that disclose extensive financial information about the company and risks associated with the future operations of the company. The SEC examines these documents, comments on them, and expects corrections based on the comments. The final product is a prospectus approved by the SEC that is available for potential buyers of the stock.

2.2.4 Software

In the software area, there have been recent attempts to certify digital data repositories as ‘trusted.’ Trustworthiness involves both the quality of the data and sustainable, reliable access to the data. The goal of certification is to enhance scientific reproducibility. The European Framework for Audit and Certification [23] has three levels of certification, Core, Extended, and Formal (or Bronze, Silver, and Gold), having different requirements, mainly to distinguish between the requirements of different types of data, e.g. research data vs. human health data vs. financial transaction data. The CoreTrustSeal [24], a private legal entity, provides a Bronze level certification to an interested data repository, for a nominal fee.

There have been several proposals in the literature for software certifications of various kinds. Ghosh and McGraw [25] propose a certification process for testing software components for security properties. Their technique involves a process and a set of white-box and black-box testing procedures, that eventually results in a stamp of approval in the form of a digital signature. Schiller [26] proposes a certification process that starts with a checklist with yes/no answers provided by the developer, and determines which tests need to be performed on the software to certify it. Currit et al. [27] describe a procedure for certifying the reliability of software before its release to the users. They predict the performance of the software on unseen inputs using the MTTF (mean time to failure) metric. Port and Wilf [28] describe

a procedure to certify the readiness for software release, understanding the tradeoff in cost of too early a release due to failures in the field, versus the cost in personnel and schedule delay arising from more extensive testing. Their technique involves the filling out of a questionnaire by the software developer called the Software Review and Certification Record (SRCR), which is ‘credentialed’ with signatories who approve the document prior to the release decision. Heck et al. [29] also describe a software product certification model to certify legislative compliance or acceptability of software delivered during outsourcing. The basis for certification is a questionnaire to be filled out by the developer. The only acceptable answers to the questions are *yes* and *n/a* (not applicable).

A different approach is taken in the CERT Secure Coding Standards [30] of the Software Engineering Institute. Here the emphasis is on documenting best practices and coding standards for security purposes. The secure coding standards consist of guidelines about the types of security flaws that can be injected through development with specific programming languages. Each guideline offers precise information describing the cause and impact of violations, and examples of common non-compliant (flawed) and compliant (fixed) code. The organization also provides tools, which audits code to identify security flaws as indicated by violations of the CERT secure coding standards.

2.2.5 Environmental Impact Statements

Environment law in the United States requires that an environmental impact statement (EIS) should be prepared prior to starting large constructions. An EIS is a document used as a tool for decision making that describes positive and negative environmental effects of a proposed action. It is made available both to federal agencies and to the public, and captures impacts to endangered species, air quality, water quality, cultural sites, and the socioeconomics of local communities. The federal law, the National Environmental Policy Act, has inspired similar laws in various jurisdictions and in other fields beyond the environment. Selbst has proposed an algorithmic impact statement for AI that follows the form and purpose of EISs [31].

2.2.6 Human Subjects

In addition to products and technologies, another critical endeavor requiring trust is research involving human subjects. Institutional review boards (IRB) have precise reviewing protocols and requirements such as those presented in the Belmont Report [32]. Items to be completed include statement of purpose, participant selection, procedures to be followed, harms and benefits to subjects, confidentiality, and consent documents. As AI services increasingly make inferences for people and about people [33], IRB requirements increasingly apply to them.

2.2.7 Summary

To ensure trust in products, industries have established a variety of practices to convey information about how a product is expected to perform when utilized by a consumer. This information usually includes how the product was constructed and tested. Some industries allow product creators to voluntarily provide this information, whereas others explicitly require it. When the information is required, some industries require the information to be validated by a third party. One would expect the latter scenario to occur in mature industries where there is confidence that the requirements strongly correlate with safety, reliability, and overall trust in the product. Mandatory external validation of nascent requirements in emerging industries may unnecessarily stifle the development of the industry.

3 Elements of Trust in AI Systems

We drive cars trusting the brakes will work when the pedal is pressed. We undergo laser eye surgery trusting the system to make the right decisions. We accept that the autopilot will operate an airplane, trusting that it will navigate correctly. In all these cases, trust comes from confidence that the system will err extremely rarely, leveraging system training, exhaustive testing, experience, safety measures and standards, best practices, and consumer education.

Every time new technology is introduced, it creates new challenges, safety issues, and potential hazards. As the technology develops and matures, these issues are better understood, documented, and addressed. Human trust in technology is developed as users overcome perceptions of risk and uncertainty [34], i.e., as

they are able to assess the technology’s performance, reliability, safety, and security. Consumers do not yet trust AI like they trust other technologies because of inadequate attention given to the latter of these issues [35]. Making technical progress on safety and security is necessary but not sufficient to achieve trust in AI, however; the progress must be accompanied by the ability to measure and communicate the performance levels of the service on these dimensions in a standardized and transparent manner. One way to accomplish this is to provide such information via FactSheets for AI services.

Trust in AI services will come from: a) applying general safety and reliability engineering methodologies across the entire lifecycle of an AI service, b) identifying and addressing new, AI-specific issues and challenges in an ongoing and agile way, and c) creating standardized tests and transparent reporting mechanisms on how such a service operates and performs. In this section we outline several areas of concern and how they uniquely apply to AI. The crux of this discussion is the manifestation of risk and uncertainty in machine learning, including that data distributions used for training are not always the ones that ideally should be used.

3.1 Basic Performance and Reliability

Statistical machine learning theory and practice is built around risk minimization. The particular loss function, whose expectation over the data distribution is considered to be the risk, depends on the task, e.g. zero-one loss for binary classification and mean squared error for regression. Different types of errors can be given different costs. Abstract loss functions may be informed by real-world quality metrics [36], including context-dependent ones [37]. There is no particular standardization on the loss function, even broadly within application domains. Moreover, performance metrics that are not directly optimized are also often examined, e.g. area under the curve and normalized cumulative discounted gain.

The true expected value of the loss function can never be known and must be estimated empirically. There are several approaches and rules of thumb for estimating the risk, but there is no standardization here either. Different groups make different choices (k-fold cross-validation, held-out samples, stratification, bootstrapping, etc.). Further notions of performance and reliability are the technical aspects of latency, throughput, and availability of the service, which are also not standardized for the specifics of

AI workloads.

To develop trust in AI services from a basic performance perspective, the choice of metrics and testing conditions should not be left to the discretion of the supplier (who may choose conditions which present the service in a favorable light), but should be codified and standardized. The onerous requirement of third-party testing could be avoided by ensuring that the specifications are precise, i.e., that each metric is precisely defined to ensure consistency and enable reproducibility by AI service consumers.

For each metric a FactSheet should report the values under various categories relevant to the expected consumers, (e.g., performance for various age groups, geographies, or genders) with the goal of providing the right level of insight into the service, but still preserving privacy. We expect some metrics will be specific to a domain, (e.g., finance, healthcare, manufacturing), or a modality (e.g., visual, speech, text), reflecting common practice of evaluation in that environment.

3.2 Safety

While typical machine learning performance metrics are measures of risk (the ones described in the previous section), we must also consider epistemic uncertainty when assessing the safety of a service [5, 3]. The main uncertainty in machine learning is an unknown mismatch between the training data distribution and the desired data distribution on which one would ideally train. Usually that desired distribution is the true distribution encountered in operation (in this case the mismatch is known as dataset shift), but it could also be an idealized distribution that encodes preferred societal norms, policies, or regulations (imagine a more equitable world than what exists in reality). One may map four general categories of strategies to achieve safety proposed in [38] to machine learning [3]: inherently safe design, safety reserves, safe fail, and procedural safeguards, all of which serve to reduce epistemic uncertainty. Interpretability of models is one example of inherently safe design.

Dataset Shift As the statistical relationship between features and labels changes over time, known as dataset shift, the mismatch between the training distribution and the distribution from which test samples are being drawn increases. A well-known reason for performance degradation, dataset shift is a com-

mon cause of frustration and loss of trust for AI service consumers. Dataset shift can be detected and corrected using a multitude of methods [39]. The sensitivity of performance of different models to dataset shift varies and should be part of a testing protocol. To the best of our knowledge, there does not yet exist any standard for how to conduct such testing. To mitigate this risk a FactSheet should contain demographic information about the training and test datasets that report the various outcomes for each group of interest as specified in Section 3.1.

Fairness AI fairness is a rapidly growing topic of inquiry [40]. There are many different definitions of fairness (some of which provably conflict) that are appropriate in varying contexts. The concept of fairness relies on protected attributes (also context-dependent) such as race, gender, caste, and religion. For fairness, we insist on some risk measure being approximately equal in groups defined by the protected attributes. Unwanted biases in training data, due to either prejudice in labels or under-/over-sampling, lead to unfairness and can be checked using statistical tests on datasets or models [41, 42]. One can think of bias as the mismatch between the training data distribution and a desired fair distribution. Applications such as lending have legal requirements on fairness in decision making, e.g. the Equal Credit Opportunity Act in the United States. Although the parity definitions and computations in such applications are explicit, the interpretation of the numbers is subjective: there are no immutable thresholds on fairness metrics (e.g., the well-known 80% rule [43]) that are applied in isolation of context.

Explainability Directly interpretable machine learning (in contrast to post hoc interpretation) [44], in which a person can look at a model and understand what it does, reduces epistemic uncertainty and increases safety because quirks and vagaries of training dataset distributions that will not be present in distributions during deployment can be identified by inspection [3]. Different users have different needs from explanations, and there is not yet any satisfactory quantitative definition of interpretability (and there may never be) [45]. Recent regulations in the European Union require ‘meaningful’ explanations, but it is not clear what constitutes a meaningful explanation.

3.3 Security

AI services can be attacked by adversaries in various ways [4]. Small imperceptible perturbations could cause AI services to misclassify inputs to any label that attackers desire; training data and models can be poisoned, allowing attackers to worsen performance (similar to concept drift but deliberate); and sensitive information about data and models can be stolen by observing the outputs of a service for different inputs. Services may be instrumented to detect such attacks and may also be designed with defenses [46]. New research proposes certifications for defenses against adversarial examples [47], but these are not yet practical.

3.4 Lineage

Once performance, safety, and security are sufficient to engender trust, we must also ensure that we track and maintain the provenance of datasets, metadata, models along with their hyperparameters, and test results. Users, those potentially affected, and third parties, such as regulators, must be able to audit the systems underlying the services. Appropriate parties may need the ability to reproduce past outputs and track outcomes. Specifically, one should be able to determine the exact version of the service deployed at any point of time in the past, how many times the service was retrained and associated details like hyperparameters used for each training episode, training dataset used, how accuracy and safety metrics have evolved over time, the feedback data received by the service, and the triggers for retraining and improvement. This information may span multiple organizations when a service is built by multiple parties.

4 Items in a FactSheet

In this section we provide an overview of the items that should be addressed in a FactSheet. See the appendix for the complete list of items. To illustrate how these items might be completed in practice, we also include two sample FactSheets in the appendix: one for a fictitious fingerprint verification service and one for a trending topics service.

The items are grouped into several categories aligned with the elements of trust. The categories are: statement of purpose, basic performance, safety, security, and lineage. They cover various aspects of

service development, testing, deployment and maintenance: from information about the data the service is trained on, to underlying algorithms, test setup, test results, and performance benchmarks, to the way the service is maintained and retrained (including automatic adaptation).

The items are devised to aid the user in understanding how the service works, in determining if the service is appropriate for the intended application, and in comprehending its strengths and limitations. The identified items are not intended to be definitive. If a question is not applicable to a given service, it can simply be ignored. In some cases, the service supplier may not wish to disclose details of the service for competitive reasons. For example, a supplier of a commercial fraud detection service for health-care insurance claims may choose not to reveal the details of the underlying algorithm; nevertheless, the supplier should be able to indicate the class of algorithm used, provide sample outputs along with explanations of the algorithmic decisions leading to the outputs. More consequential applications will likely require more comprehensive completion of items.

A few examples of items a FactSheet might include are:

- What is the intended use of the service output?
- What algorithms or techniques does this service implement?
- Which datasets was the service tested on? (Provide links to datasets that were used for testing, along with corresponding datasheets.)
- Describe the testing methodology.
- Describe the test results.
- Are you aware of possible examples of bias, ethical issues, or other safety risks as a result of using the service?
- Are the service outputs explainable and/or interpretable?
- For each dataset used by the service: Was the dataset checked for bias? What efforts were made to ensure that it is fair and representative?
- Does the service implement and perform any bias detection and remediation?
- What is the expected performance on unseen data or data with different distributions?

- Was the service checked for robustness against adversarial attacks?
- When were the models last updated?

As such a declaration is refined, and testing procedures for performance, robustness to concept drift, explainability, and robustness to attacks are further codified, the FactSheet may refer to standardized test protocols instead of providing descriptive details.

Since completing a FactSheet can be laborious, we expect most of the information to be populated as part of the AI service creation process in a secure auditable manner. A FactSheet will be created once and associated with a service, but can continually be augmented, without removing previous information, i.e., results are added from more tests, but results cannot be removed. Any changes made to the service will prompt the creation of a new version of the FactSheet for the new model. Thus, these FactSheets will be treated as a series of immutable artifacts.

This information can be used to more accurately monitor a deployed service by comparing deployed metrics with those that were seen during development and taking appropriate action when unexpected behavior is detected.

5 The Evolution of FactSheet Adoption

We expect that AI will soon go through the same evolution that other technologies have gone through (cf. [8] for an excellent review of the evolution of safety standards in different industries). We propose that FactSheets be initially voluntary for several reasons. First, discussion and feedback from multiple parties representing suppliers and consumers of AI services is needed to determine the final set of items and format of FactSheets. So, an initial voluntary period to allow this discussion to occur is needed. Second, there needs to be a balance between the needs of AI service consumers with the freedom to innovate for AI service producers. Although producing a FactSheet will initially be an additional burden to an AI service producer, we expect market feedback from AI service consumers to encourage this creation.

Because of peer pressure to conform [48], FactSheets could become a de facto requirement similar to Energy Star labeling of the energy efficiency of appliances. They will serve to reduce information asymmetry between supplier and consumer, where consumers are currently unaware of important properties

of a service, such as its intended use, its performance metrics, and information about fairness, explainability, safety, and security. In particular, consumers in many businesses do not have the requisite expertise to evaluate various AI services available in the marketplace; uninformed or incorrect choices can result in suboptimal business performance. By creating easily consumable FactSheets, suppliers can accrue a competitive advantage by capturing consumers' trust. Moreover, with such transparency, FactSheets should serve to allow better functioning of AI service marketplaces and prevent a so-called 'market for lemons' [49]. A counter-argument to voluntary compliance and self-regulation argues that while participation of industry is welcome, this should not stand in the way of legislation and governmental regulation [50].

FactSheet adoption could potentially lead to an eventual system of third-party certification [51], but probably only for services catering to applications with the very highest of stakes, to regulated business processes and enterprise applications, and to applications originating in the public sector [52, 7]. Children's toys are an example category of consumer products in which an SDoC is not enough and certification is required. If an AI service is already touching on a regulation from a specific industry in which it is being used, its FactSheet will serve as a tool for better compliance.

6 Discussion and Future Work

One may wonder why AI should be held to a higher standard (FactSheets) than non-AI software and services in the same domain. Non-AI software include several artifacts beyond the code, such as design documents, program flow charts, and test plans that can provide transparency to concerned consumers. Since AI services do not contain any of these, and the generated code may not be easily understandable, there is a higher demand to enhance transparency through FactSheets.

Although FactSheets enable AI services producers to provide information about the intent and construction of their service so that educated consumers can make informed decisions, consumers may still, innocently or maliciously, use the service for purposes other than those intended. FactSheets cannot fully protect against such use, but can form the basis of service level agreements.

Some components of an AI service may be produced by organizations other than the service sup-

plier. For example, the dataset may be obtained from a third party, or the service may be a composition of models, some of which are produced by another organization. In such cases, the FactSheet for the composed service would need to include information from the supplying organizations. Ideally, those organizations would produce FactSheets for their components, enabling the composing organization to provide a complete FactSheet. This complete FactSheet could include the component FactSheets along with any necessary additional information. In some cases, the demands for transparency on the composing organization may be greater than on the component organization; market forces will require the component organization to provide more transparency to retain their relation with the composing organization. This is analogous to other industries, like retail, where retailers push demands on their suppliers to meet the expectations of the retailers' customers. In these situations the provenance of the information among organizations will need to be tracked.

7 Summary and Conclusion

In this paper, we continue in the research direction established by datasheets or nutrition labels for datasets to examine trusted AI at the functional level rather than at the component level. We discuss the several elements of AI services that are needed for people to trust them, including task performance, safety, security, and maintenance of lineage. The final piece to build trust is transparent documentation about the service, which we see as a variation on declarations of conformity for consumer products. We propose a starting point to a voluntary AI service supplier's declaration of conformity. Further discussion among multiple parties is required to standardize protocols for testing AI services and determine the final set of items and format that AI service FactSheets will take.

We envision that suppliers will voluntarily populate and release FactSheets for their services to remain competitive in the market. The evolution of the marketplace of AI services may eventually lead to an ecosystem of third party testing and verification laboratories, services, and tools. We also envision the automation of nearly the entire FactSheet as part of the build and runtime environments of AI services. Moreover, it is not difficult to imagine FactSheets being automatically posted to distributed, immutable ledgers such as those enabled by blockchain

technologies.

We see our work as a first step at defining which questions to ask and metrics to measure towards development and adoption of broader industry practices and standards. We see a parallel between the issue of trusted AI today and the rise of digital certification during the Internet revolution. The digital certification market 'bootstrapped' the Internet, ushering in a new era of 'transactions' such as online banking and benefits enrollment that we take for granted today. In a similar vein, we can see AI service FactSheets ushering in a new era of trusted AI end points and bootstrapping broader adoption.

References

- [1] E. E. Levine, T. B. Bitterly, T. R. Cohen, and M. E. Schweitzer, "Who is trustworthy? Predicting trustworthy intentions and behavior," *Journal of Personality and Social Psychology*, vol. 115, no. 3, pp. 468–494, Sep. 2018.
- [2] M. K. Lee, "Understanding perception of algorithmic decisions: Fairness, trust, and emotion in response to algorithmic management," *Big Data & Society*, vol. 5, no. 1, Jan.–Jun. 2018.
- [3] K. R. Varshney and H. Alemzadeh, "On the safety of machine learning: Cyber-physical systems, decision sciences, and data products," *Big Data*, vol. 5, no. 3, pp. 246–255, Sep. 2017.
- [4] N. Papernot, P. McDaniel, S. Jha, M. Fredrikson, Z. B. Celik, and A. Swami, "The limitations of deep learning in adversarial settings," in *Proceedings of the IEEE European Symposium on Security and Privacy*, Saarbrücken, Germany, Mar. 2016, pp. 372–387.
- [5] N. Moller, "The concepts of risk and safety," in *Handbook of Risk Theory*, S. Roeser, R. Hillerbrand, P. Sandin, and M. Peterson, Eds. Dordrecht, Netherlands: Springer, 2012, pp. 55–85.
- [6] National Institute of Standards and Technology, "The use of supplier's declaration of conformity," <https://www.nist.gov/document-6075>.
- [7] American National Standards Institute, "U. S. conformity assessment system: 1st party conformity assessment," https://www.standardsportal.org/usa/en/-conformity_assessment/suppliers_declaration.aspx.

- [8] T. Gebru, J. Morgenstern, B. Vecchione, J. W. Vaughan, H. Wallach, H. Daumé, III, and K. Crawford, “Datasheets for datasets,” in *Proceedings of the Fairness, Accountability, and Transparency in Machine Learning Workshop*, Stockholm, Sweden, Jul. 2018.
- [9] E. M. Bender and B. Friedman, “Data statements for NLP: Toward mitigating system bias and enabling better science,” *Transactions of the ACL*, forthcoming.
- [10] S. Holland, A. Hosny, S. Newman, J. Joseph, and K. Chmielinski, “The dataset nutrition label: A framework to drive higher data quality standards,” arXiv:1805.03677, May 2018.
- [11] M. Mitchell, S. Wu, A. Zaldivar, P. Barnes, L. Vasserman, B. Hutchinson, E. Spitzer, I. D. Raji, and T. Gebru, “Model cards for model reporting,” in *Proceedings of the ACM Conference on Fairness, Accountability, and Transparency*, Atlanta, USA, Jan. 2019.
- [12] M. Loukides, H. Mason, and D. Patil, “Of oaths and checklists,” <https://www.oreilly.com/ideas/of-oaths-and-checklists>, Jul. 2018.
- [13] M. Hind, S. Mehta, A. Mojsilović, R. Nair, K. N. Ramamurthy, A. Olteanu, and K. R. Varshney, “Increasing trust in AI services through supplier’s declarations of conformity,” <https://arxiv.org/abs/1808.07261>, Aug. 2018.
- [14] C. O’Neil, “What is a data audit?” <http://www.oneilrisk.com/articles/2017/1/24/-what-is-a-data-audit>, Jan. 2017.
- [15] R. Carrier, “AI safety— the concept of independent audit,” <https://www.forhumanity.center/independent-audit-1/>.
- [16] The European Commission’s High-Level Expert Group on Artificial Intelligence, “Draft ethics guidelines for trustworthy AI,” Brussels, Belgium, Dec. 2018.
- [17] [Online]. Available: <http://standards.ieee.org/develop/overview.html>
- [18] [Online]. Available: <https://www.iso.org/home.html>
- [19] [Online]. Available: <https://www.cpsc.gov/Business--Manufacturing/Testing-Certification>
- [20] [Online]. Available: <http://the-parenting-center.com/about-the-seal-of-approval>
- [21] [Online]. Available: <https://www.moodys.com/Pages/amr002002.aspx>
- [22] [Online]. Available: https://www.spratings.com/documents/20184/774196/Guide_to...
- [23] [Online]. Available: <http://www.trusteddigitalrepository.eu/Trusted%20Digital%20R...>
- [24] [Online]. Available: <https://www.coretrustseal.org/about/>
- [25] A. K. Ghosh and G. McGraw, “An approach for certifying security in software components,” in *the 21st National Information Systems Security Conference*, 1998.
- [26] S. C. A., “The software certification process.” [Online]. Available: <http://www.ittoday.info/AIMS/DSM/82-01-17.pdf>
- [27] P. A. Currit, M. Dyer, and H. D. Mills, “Certifying the reliability of software,” *IEEE Transactions on Software Engineering* 1, pp. 3–11, 1986.
- [28] D. Port and J. Wilf, “The value of certifying software release readiness: an exploratory study of certification for a critical system at jpl,” in *2013 ACM/IEEE International Symposium on Empirical Software Engineering and Measurement*, 2013, pp. 373–382.
- [29] P. Heck, M. Klabbers, and M. van Eekelen, “A software product certification model,” *Software Quality Journal*, vol. 18, no. 1, 2010.
- [30] [Online]. Available: <https://www.sei.cmu.edu/research-capabilities/all-work/display...>
- [31] A. D. Selbst, “Disparate impact in big data policing,” *Georgia Law Review*, vol. 52, no. 1, pp. 109–195, Feb. 2017.
- [32] J. M. Sims, “A brief review of the Belmont report,” *Dimensions of Critical Care Nursing*, vol. 29, no. 4, pp. 173–174, Jul.–Aug. 2010.
- [33] K. R. Varshney, “Data science of the people, for the people, by the people: A viewpoint on an emerging dichotomy,” in *Data for Good Exchange Conference*, New York, USA, Sep. 2015.

- [34] X. Li, T. J. Hess, and J. S. Valacich, “Why do we trust new technology? a study of initial trust formation with organizational information systems,” *Journal of Strategic Information Systems*, vol. 17, no. 1, pp. 39–71, Mar. 2008.
- [35] S. Scott, “Artificial intelligence & communications: The fads. the fears. the future.” 2018.
- [36] K. L. Wagstaff, “Machine learning that matters,” in *Proceedings of the International Conference on Machine Learning*, Edinburgh, UK, Jun.–Jul. 2012, pp. 529–536.
- [37] A. Olteanu, K. Talamadupula, and K. R. Varshney, “The limits of abstract evaluation metrics: The case of hate speech detection,” in *Proceedings of the ACM Web Science Conference*, Troy, USA, Jun. 2017, pp. 405–406.
- [38] N. Möller and S. O. Hansson, “Principles of engineering safety: Risk and uncertainty reduction,” *Reliability Engineering & System Safety*, vol. 93, no. 6, pp. 798–805, Jun. 2008.
- [39] J. Gama, I. Žliobaitė, A. Bifet, M. Pechenizkiy, and A. Bouchachia, “A survey on concept drift adaptation,” *ACM Computing Surveys*, vol. 46, no. 4, p. 44, Apr. 2014.
- [40] S. Hajian, F. Bonchi, and C. Castillo, “Algorithmic bias: From discrimination discovery to fairness-aware data mining,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, San Francisco, USA, Aug. 2016, pp. 2125–2126.
- [41] S. Barocas and A. D. Selbst, “Big data’s disparate impact,” *California Law Review*, vol. 104, no. 3, pp. 671–732, Jun. 2016.
- [42] R. K. E. Bellamy, K. Dey, M. Hind, S. C. Hoffman, S. Houde, K. Kannan, P. Lohia, J. Martino, S. Mehta, A. Mojsilovic, S. Nagar, K. N. Ramamurthy, J. Richards, D. Saha, P. Sattigeri, M. Singh, K. R. Varshney, and Y. Zhang, “AI Fairness 360: An extensible toolkit for detecting, understanding, and mitigating unwanted algorithmic bias,” arXiv:1810.01943, Oct. 2018.
- [43] M. Feldman, S. A. Friedler, J. Moeller, C. Scheidegger, and S. Venkatasubramanian, “Certifying and removing disparate impact,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, Sydney, Australia, Aug. 2015, pp. 259–268.
- [44] C. Rudin, “Algorithms for interpretable machine learning,” in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, USA, Aug. 2014, p. 1519.
- [45] F. Doshi-Velez and B. Kim, “Towards a rigorous science of interpretable machine learning,” arXiv:1702.08608, Feb. 2017.
- [46] M.-I. Nicolae, M. Sinn, M. N. Tran, A. Rawat, M. Wistuba, V. Zantedeschi, N. Baracaldo, B. Chen, H. Ludwig, I. M. Molloy, and B. Edwards, “Adversarial robustness toolbox v0.3.0,” arXiv:1807.01069, Aug. 2018.
- [47] A. Raghunathan, J. Steinhardt, and P. Liang, “Certified defenses against adversarial examples,” arXiv:1801.09344, Jan. 2018.
- [48] J. ben-Aaron, M. Denny, B. Desmarais, and H. Wallach, “Transparency by conformity: A field experiment evaluating openness in local governments,” *Public Administration Review*, vol. 77, no. 1, pp. 68–77, Jan.–Feb. 2017.
- [49] G. A. Akerlof, “The market for “lemons”: Quality uncertainty and the market mechanism,” *Quarterly Journal of Economics*, vol. 84, no. 3, pp. 488–500, Aug. 1970.
- [50] P. Nemitz, “Constitutional democracy and technology in the age of artificial intelligence,” *Philosophical Transactions of the Royal Society A*, vol. 376, no. 2133, p. 20180089, Nov. 2018.
- [51] B. Srivastava and F. Rossi, “Towards composable bias rating of AI services,” in *Proceedings of the AAAI/ACM Conference on Artificial Intelligence, Ethics, and Society*, New Orleans, USA, Feb. 2018.
- [52] L. K. McAllister, “Harnessing private regulation,” *Michigan Journal of Environmental & Administrative Law*, vol. 3, no. 2, pp. 291–420, 2014.

A Proposed FactSheet Items

Below we list example questions that a FactSheet for an AI service could include. The set of questions we provide here is not intended to be definitive, but rather to open a conversation about what aspects should be covered.

To illustrate how these questions could be answered, we provide two examples for fictitious AI services: a fingerprint verification service (Appendix B) and a trending topics service (Appendix C). However, given that the examples we provide are fictitious, we would expect an actual service provider to answer these questions in much more detail. For instance, they would be able to better characterize an API that actually exists. Our example answers are mainly to provide additional insight about the type of information we would find in a FactSheet.

Statement of purpose

The following questions are aimed at providing an overview of the service provider and of the intended uses for the service. Valid answers include “N/A” (not applicable) and “Proprietary” (cannot be publicly disclosed, usually for competitive reasons).

General

- Who are “you” (the supplier) and what type of services do you typically offer (beyond this particular service)?

- What is this service about?

- Briefly describe the service.
- When was the service first released? When was the last release?
- Who is the target user?

- Describe the outputs of the service.
- What algorithms or techniques does this service implement?

- Provide links to technical papers.

- What are the characteristics of the development team?

- Do the teams charged with developing and maintaining this service reflect a diversity of opinions, backgrounds, and thought?

- Have you updated this FactSheet before?

- When and how often?
- What sections have changed?
- Is the FactSheet updated every time the service is retrained or updated?

Usage

- What is the intended use of the service output?

- Briefly describe a simple use-case.

- What are the key procedures followed while using the service?

- How is the input provided? By whom?
- How is the output returned?

Domains and applications

- What are the domains and applications the service was tested on or used for?

- Were domain experts involved in the development, testing, and deployment? Please elaborate.

- How is the service being used by your customers or users?

- Are you enabling others to build a solution by providing a cloud service or is your application end-user facing?
- Is the service output used as-is, is it fed directly into another tool or actuator, or is there human input/oversight before use?
- Do users rely on pre-trained/canned models or can they train their own models?
- Do your customers typically use your service in a time critical setup (e.g. they have limited time to evaluate the output)? Or do they incorporate it in a slower decision making process? Please elaborate.

- List applications that the service has been used for in the past.

- Please provide information about these applications or relevant pointers.
- Please provide key performance results for those applications.

- Briefly describe how a third party could independently verify the performance of the service.
- Are there benchmarks publicly available and adequate for testing the service.

- Other comments?

Basic Performance

The following questions aim to offer an overall assessment of the service performance.

Testing by service provider

- Which datasets was the service tested on? (e.g., links to datasets that were used for testing, along with corresponding datasheets)

- List the test datasets and provide links to these datasets.
- Do the datasets have an associated datasheet? If yes, please attach.
- Could these datasets be used for independent testing of the service? Did the data need to be changed or sampled before use?

- Describe the testing methodology.

- Please provide details on train, test and holdout data.
- What performance metrics were used? (e.g. accuracy, error rates, AUC, precision/recall)
- Please briefly justify the choice of metrics.

- Describe the test results.

- Were latency, throughput, and availability measured?
- If yes, briefly include those metrics as well.

Testing by third parties

- Is there a way to verify the performance metrics (e.g., via a service API)?

- In addition to the service provider, was this service tested by any third party?

- Please list all third parties that performed the testing.
- Also, please include information about the tests and test results.

- Other comments?

Safety

The following questions aim to offer insights about potential unintentional harms, and mitigation efforts to eliminate or minimize those harms.

General

- Are you aware of possible examples of bias, ethical issues, or other safety risks as a result of using the service?

- Were the possible sources of bias or unfairness analyzed?
- Where do they arise from: the data? the particular techniques being implemented? other sources?
- Is there any mechanism for redress if individuals are negatively affected?

- Do you use data from or make inferences about individuals or groups of individuals. Have you obtained their consent?

- How was it decided whose data to use or about whom to make inferences?
- Do these individuals know that their data is being used or that inferences are being made about them? What were they told? When were they made aware? What kind of consent was needed from them? What were the procedures for gathering consent? Please attach the consent form to this declaration.
- What are the potential risks to these individuals or groups? Might the service output interfere with individual rights? How are these risks being handled or minimized?
- What trade-offs were made between the rights of these individuals and business interests?
- Do they have the option to withdraw their data? Can they opt out from inferences being made about them? What is the withdrawal procedure?

- Please describe the data bias policies that were checked (such as with respect to known protected attributes), bias checking methods, and results (e.g., disparate error rates across different groups).
- Was there any bias remediation performed on this dataset? Please provide details about the value of any bias estimates before and after it.
- What techniques were used to perform the remediation? Please provide links to relevant technical papers.
- How did the value of other performance metrics change as a result?

- Does the service implement and perform any bias detection and remediation?

- Please describe model bias policies that were checked, bias checking methods, and results (e.g., disparate error rates across different groups).
- What procedures were used to perform the remediation? Please provide links or references to corresponding technical papers.
- Please provide details about the value of any bias estimates before and after such remediation.
- How did the value of other performance metrics change as a result?

Explainability

- Are the service outputs explainable and/or interpretable?

- Please explain how explainability is achieved (e.g. directly explainable algorithm, local explainability, explanations via examples).
- Who is the target user of the explanation (ML expert, domain expert, general consumer, etc.)
- Please describe any human validation of the explainability of the algorithms

Concept Drift

- What is the expected performance on unseen data or data with different distributions?

- Please describe any relevant testing done along with test results.

Fairness

- For each dataset used by the service: Was the dataset checked for bias? What efforts were made to ensure that it is fair and representative?

- Does your system make updates to its behavior based on newly ingested data?

- Is the new data uploaded by your users? Is it generated by an automated process? Are the patterns in the data largely static or do they change over time?
- Are there any performance guarantees/bounds?
- Does the service have an automatic feedback/retraining loop, or is there a human in the loop?

- How is the service tested and monitored for model or performance drift over time?

- If applicable, describe any relevant testing along with test results.

- How can the service be checked for correct, expected output when new data is added?
- Does the service allow for checking for differences between training and usage data?

- Does it deploy mechanisms to alert the user of the difference?

- Do you test the service periodically?

- Does the testing includes bias or fairness related aspects?
- How has the value of the tested metrics evolved over time?

- Other comments?

Security

The following questions aim to assess the susceptibility to deliberate harms such as attacks by adversaries.

- How could this service be attacked or abused? Please describe.
- List applications or scenarios for which the service is not suitable.

- Describe specific concerns and sensitive use cases.
- Are there any procedures in place to ensure that the service will not be used for these applications?

- How are you securing user or usage data?

- Is usage data from service operations retained and stored?
- How is the data being stored? For how long is the data stored?
- Is user or usage data being shared outside the service? Who has access to the data?

- Was the service checked for robustness against adversarial attacks?

- Describe robustness policies that were checked, the type of attacks considered, checking methods, and results.

- What is the plan to handle any potential security breaches?

- Describe any protocol that is in place.

- Other comments?

Lineage

The following questions aim to overview how the service provider keeps track of details that might be required in the event of an audit by a third party, such as in the case of harm or suspicion of harm.

Training Data

- Does the service provide an as-is/canned model? Which datasets was the service trained on?

- List the training datasets.
- Where there any quality assurance processes employed while the data was collected or before use?
- Were the datasets used for training built-for-purpose or were they repurposed/adapted? Were the datasets created specifically for the purpose of training the models offered by this service?

- Did you use any prior knowledge or re-weight the data in any way before training?
- Other comments?

- For each dataset: Are the training datasets publicly available?

- Please provide a link to the datasets or the source of the datasets.

- For each dataset: Does the dataset have a datasheet or data statement?

- If available, attach the datasheet; otherwise, provide answers to questions from the datasheet as appropriate [to insert citation]

- Did the service require any transformation of the data in addition to those provided in the datasheet?

- Do you use synthetic data?

- When? How was it created?
- Briefly describe its properties and the creation procedure.

Trained Models

- How were the models trained?

- Please provide specific details (e.g., hyperparameters).

- When were the models last updated?

- How much did the performance change with each update?
- How often are the models retrained or updated?

B Sample FactSheet for a Fingerprint Verification Service

Statement of purpose

The following questions are aimed at providing an overview of the service provider and of the intended uses for the service. Valid answers include “N/A” (not applicable) and “Proprietary” (cannot be publicly disclosed, usually for competitive reasons).

General

- Who are “you” (the supplier) and what type of services do you typically offer (beyond this particular service)?

Raj Kumar Biometrics Services, Ltd. The only service we offer at present is fingerprint verification.

- What is this service about?

- Briefly describe the service.
- When was the service first released? When was the last release?
- Who is the target user?

The service takes an ordered pair of fingerprint image and identity and returns a 1 if the image matches the image corresponding to that identity in the database. The service accepts 500 dpi images acquired using optical sensors. The v1.0 algorithm was created on June 30, 2005. The current algorithm v1.7 was created on April 12, 2010. The algorithm was released as a cloud service on August 10, 2017. The target user is a manufacturer who creates physical access control systems as well as other entities interested in physical or informational access control.

- Describe the outputs of the service.

A binary verification label.

- What algorithms or techniques does this service implement?

- Provide links to technical papers.

P. Baldi and Y. Chauvin, “Neural networks for fingerprint recognition,” *Neural Computation*, vol. 5, no. 3, pp. 402–418, 1993.

- What are the characteristics of the development team?

- Do the teams charged with developing and maintaining this service reflect a diversity of opinions, backgrounds, and thought?

The service was developed by 3 graduates of Delhi College of Engineering. It was made into a cloud service and is maintained by 2 graduates of Amity University.

- Have you updated this FactSheet before?

- When and how often?
- What sections have changed?
- Is the FactSheet updated every time the service is retrained or updated?

This is our first release of FactSheet. We plan to release a new FactSheet when we release v1.8.

Usage

- What is the intended use of the service output?

- Briefly describe a simple use-case.

A locks manufacturer is creating a biometrics-driven access control system that it will sell to call centers. This internet-enabled system will include an optical sensor for fingerprints and a keypad for the user to enter a 7 digit identification number. The system will acquire the fingerprint and identification number and transmit them to our service via a RESTful API. If the image matches the image for that user in the previously acquired database, it will return a positive result and the system will unlock.

- What are the key procedures followed while using the service?

- How is the input provided? By whom?
- How is the output returned?

The end user supplies his or her fingerprint via an optical sensor which digitizes it and transmits it to the service. The binary output is returned to the access-control system.

Domains and applications

- What are the domains and applications the service was tested on or used for?

- Were domain experts involved in the development, testing, and deployment? Please elaborate.

The service has only been tested on fingerprint verification of call center employees with no domain experts involved.

- How is the service being used by your customers or users?

- Are you enabling others to build a solution by providing a cloud service or is your application end-user facing?
- Is the service output used as-is, is it fed directly into another tool or actuator, or is there human input/oversight before use?
- Do users rely on pre-trained/canned models or can they train their own models?
- Do your customers typically use your service in a time critical setup (e.g. they have limited time to evaluate the output)? Or do they incorporate it in a slower decision making process? Please elaborate.

Our service is a cloud service for access control system manufacturers. The output is directly fed into an actuator. Users can only rely on pre-trained models, but will necessarily upload a database of individual identifiers with their fingerprints. The service requires outputs be given with small delay.

- List applications that the service has been used for in the past.

- Please provide information about these applications or relevant pointers.
- Please provide key performance results for those applications.

Call center access control and bank access control.

- Other comments?

No.

Basic Performance

The following questions aim to offer an overall assessment of the service performance.

Testing by service provider

- Which datasets was the service tested on? (e.g., links to datasets that were used for testing, along with corresponding datasheets)

- List the test datasets and provide links to these datasets.
- Do the datasets have an associated datasheet? If yes, please attach.
- Could these datasets be used for independent testing of the service? Did the data need to be changed or sampled before use?

FVC2002 DB1 (<http://bias.csr.unibo.it/fvc2002/databases>)
This dataset does not have a datasheet. Yes, this dataset can be used for independent testing.

- Describe the testing methodology.

- Please provide details on train, test and holdout data.
- What performance metrics were used? (e.g. accuracy, error rates, AUC, precision/recall)
- Please briefly justify the choice of metrics.

Performance metrics were evaluated on a heldout set as specified by FVC2002. We used the same metrics as evaluated by FVC2002: equal error

rate (EER), the lowest false non-match rate for a false match rate $\lambda = 1\%$ (FMR100), the lowest false non-match rate for a false match rate $\lambda = 0.1\%$ (FMR1000), the lowest false non-match rate for a false match rate $\lambda = 0\%$ (ZeroFMR), number of rejected fingerprints during enrollment (RejENROLL), and number of rejected fingerprints during genuine and imposter matches (RejMATCH).

- Describe the test results.

– Were latency, throughput, and availability measured?
– If yes, briefly include those metrics as well.

The accuracy results are as follows: EER = 3.7%, FMR100 = 6.0%, ZeroFMR = 12.4%, RejENROLL = 0.0%, RejMATCH = 0.0%. We also measured average enrollment time: 0.14 sec and average matching time: 0.44 sec.

Testing by third parties

- Is there a way to verify the performance metrics (e.g., via a service API)?

– Briefly describe how a third party could independently verify the performance of the service.
– Are there benchmarks publicly available and adequate for testing the service.

Yes, a third party can call our service via the same RESTful API that our customers use.

- In addition to the service provider, was this service tested by any third party?

– Please list all third parties that performed the testing.
– Also, please include information about the tests and test results.

No.

- Other comments?

No.

Safety

The following questions aim to offer insights about potential unintentional harms, and mitigation efforts to eliminate or minimize those harms.

General

- Are you aware of possible examples of bias, ethical issues, or other safety risks as a result of using the service?

– Were the possible sources of bias or unfairness analyzed?
– Where do they arise from: the data? the particular techniques being implemented? other sources?
– Is there any mechanism for redress if individuals are negatively affected?

Yes, individuals with a history of manual labor will have poorer performance in fingerprint verification. Children will have poorer performance in fingerprint verification. Individuals without fingerprints will be unable to use our service. There is no mechanism for redress.

- Do you use data from or make inferences about individuals or groups of individuals. Have you obtained their consent?

- How was it decided whose data to use or about whom to make inferences?
- Do these individuals know that their data is being used or that inferences are being made about them? What were they told? When were they made aware? What kind of consent was needed from them? What were the procedures for gathering consent? Please attach the consent form to this declaration.
- What are the potential risks to these individuals or groups? Might the service output interfere with individual rights? How are these risks being handled or minimized?
- What trade-offs were made between the rights of these individuals and business interests?
- Do they have the option to withdraw their data? Can they opt out from inferences being made about them? What is the withdrawal procedure?

Our training datasets come from international optical fingerprint databases available on the internet including FVC2000, FVC2002, and FVC2004. We did not do any further due diligence on these datasets.

Explainability

- Are the service outputs explainable and/or interpretable?

- Please explain how explainability is achieved (e.g. directly explainable algorithm, local explainability, explanations via examples).
- Who is the target user of the explanation (ML expert, domain expert, general consumer, etc.)
- Please describe any human validation of the explainability of the algorithms

No.

Fairness

- For each dataset used by the service: Was the dataset checked for bias? What efforts were made to ensure that it is fair and representative?

- Please describe the data bias policies that were checked (such as with respect to known protected attributes), bias checking methods, and results (e.g., disparate error rates across different groups).
- Was there any bias remediation performed on this dataset? Please provide details about the value of any bias estimates before and after it.
- What techniques were used to perform the remediation? Please provide links to relevant technical papers.
- How did the value of other performance metrics change as a result?

We tested the service on a population of our company’s employees and other office workers in our building, which includes younger and older adults, both male and female, with a range of skin tones. We did not observe any systematic differential performance. No bias remediation was performed.

- Does the service implement and perform any bias detection and remediation?

- Please describe model bias policies that were checked, bias checking methods, and results (e.g., disparate error rates across different groups).
- What procedures were used to perform the remediation? Please provide links or references to corresponding technical papers.
- Please provide details about the value of any bias estimates before and after such remediation.
- How did the value of other performance metrics change as a result?

No.

Concept Drift

- What is the expected performance on unseen data or data with different distributions?

- Please describe any relevant testing done along with test results.

Data from different types of sensors will result in extremely poor performance. Data from people from all working classes (those with frequent cuts on their hands) will result in a degradation of performance.

- Does your system make updates to its behavior based on newly ingested data?

– Is the new data uploaded by your users? Is it generated by an automated process? Are the patterns in the data largely static or do they change over time?

– Are there any performance guarantees/bounds?

– Does the service have an automatic feedback/retraining loop, or is there a human in the loop?

We have started to capture user data from our actual customer deployments and will retrain the algorithm including these additional images for v1.8.

- How is the service tested and monitored for model or performance drift over time?

– If applicable, describe any relevant testing along with test results.

Proprietary.

- How can the service be checked for correct, expected output when new data is added?

We have not yet added any new data up to v1.7.

- Does the service allow for checking for differences between training and usage data?

– Does it deploy mechanisms to alert the user of the difference?

No.

- Do you test the service periodically?

– Does the testing includes bias or fairness related aspects?

– How has the value of the tested metrics evolved over time?

Yes, we depute one of our staff members to visit our customer deployments once per quarter and do spot checks by enrolling and testing a few people. Metric evolution over time is confidential.

- Other comments?

The general characteristics of fingerprints do not change over time.

Security

The following questions aim to assess the susceptibility to deliberate harms such as attacks by adversaries.

- How could this service be attacked or abused? Please describe.

Many different attacks are possible. Several are described in B. Biggio, G. Fumera, P. Russu, L. Didaci, and F. Roli, “Adversarial biometric recognition: A review on biometric system security from the adversarial machine-learning perspective,” *IEEE Signal Processing Magazine*, vol. 32, no. 5, pp. 31–41, 2015.

- List applications or scenarios for which the service is not suitable.

– Describe specific concerns and sensitive use cases.

– Are there any procedures in place to ensure that the service will not be used for these applications?

The service should not be used to investigate crimes, prosecute individuals, or used in any other way except for access control. There are no procedures in place to prevent such usage.

- How are you securing user or usage data?

- Is usage data from service operations retained and stored?
- How is the data being stored? For how long is the data stored?
- Is user or usage data being shared outside the service? Who has access to the data?

The usage data is not stored except when the user provides negative feedback and explicitly agrees for us to use current sample for retraining. The samples are deleted after retraining.

- Was the service checked for robustness against adversarial attacks?

- Describe robustness policies that were checked, the type of attacks considered, checking methods, and results.

We have checked for model-inversion and hill-climbing attacks using the techniques developed in M. Martinez-Diaz, J. Fierrez, J. Galbally, and J. Ortega-Garcia, “An evaluation of indirect attacks and countermeasures in fingerprint verification systems,” *Pattern Recognition Letters*, vol. 32, no. 12, pp. 1643–1651, 2011. Our service passed these tests.

- What is the plan to handle any potential security breaches?

- Describe any protocol that is in place.

We will shut down the service immediately in case of a potential security breach and only bring customers back online after site visits.

- Other comments?

We take security very seriously.

Lineage

The following questions aim to overview how the service provider keeps track of details that might be required in the event of an audit by a third party, such as in the case of harm or suspicion of harm.

Training Data

- Does the service provide an as-is/canned model? Which datasets was the service trained on?

- List the training datasets.
- Where there any quality assurance processes employed while the data was collected or before use?
- Were the datasets used for training built-for-purpose or were they repurposed/adapted? Were the datasets created specifically for the purpose of training the models offered by this service?

Our training datasets come from international optical fingerprint databases available on the internet including FVC2000, FVC2002, and FVC2004. All quality assurance was done by the dataset providers. They were purpose-built for the evaluation and testing of fingerprint verification systems.

- For each dataset: Are the training datasets publicly available?

- Please provide a link to the datasets or the source of the datasets.

Yes: <http://bias.csr.unibo.it/fvc2000/databases.asp>, <http://bias.csr.unibo.it/fvc2002/databases.asp>, <http://bias.csr.unibo.it/fvc2004/databases.asp>.

- For each dataset: Does the dataset have a datasheet or data statement?

- If available, attach the datasheet; otherwise, provide answers to questions from the datasheet as appropriate [to insert citation]

No.

- Did the service require any transformation of the data in addition to those provided in the datasheet?

No.

- Do you use synthetic data?

- When? How was it created?
- Briefly describe its properties and the creation procedure.

No.

Trained Models

- How were the models trained?

- Please provide specific details (e.g., hyper-parameters).

Proprietary.

- When were the models last updated?

- How much did the performance change with each update?
- How often are the models retrained or updated?

Proprietary.

- Did you use any prior knowledge or re-weight the data in any way before training?

No.

- Other comments?

No.

C Sample FactSheet for a Trending Topics Service

Statement of purpose

The following questions are aimed at providing an overview of the service provider and of the intended uses for the service. Valid answers include “N/A” (not applicable) and “Proprietary” (cannot be publicly disclosed, usually for competitive reasons).

General

- Who are “you” (the supplier) and what type of services do you typically offer (beyond this particular service)?

DataTrendly specializes in natural language processing and time series analysis offering a wide range of products focused on the analysis of trending topics in several types of textual data, such as social media, news media, and scientific publications.

- What is this service about?

- Briefly describe the service.
- When was the service first released? When was the last release?
- Who is the target user?

The DataTrendly’s social media trending topics service allows our customers to check, identify, search for, and monitor trends on a variety of social media platforms. The service was first released in January 2014, and it was last updated in June 2018. Our target users are broad, anyone who wants to monitor a trending topic. Some examples are a company wants to monitor its brand or media company that wants to model particular events.

- Describe the outputs of the service.

The service is offered as a comprehensive set of RESTful API calls. The main API calls return a ranked list of N trending topics for a given time interval and a given list of key-phrases of interest. For each trending topic it returns the key-phrases that triggered the topic, the time stamp

of the topic, and a list of other related trending topics.

- What algorithms or techniques does this service implement?

- Provide links to technical papers.

The service implements a mix of both commonly used techniques and our own proprietary techniques. Our users can specify what technique they want to use and can compare results from multiple techniques.

For nowcasting and forecasting, the service implements known models like ARMA (Autoregressive, Moving Average), as well as proprietary neural network models. Our implementation of known techniques builds on Seabold and Perktold (2010)

For past trends, it implements proprietary techniques for detecting sudden changes in time series data, which are tailored for social media data.

Seabold, Skipper, and Josef Perktold. “Statsmodels: Econometric and statistical modeling with python.” Proceedings of the 9th Python in Science Conference. 2010.

Link: <https://www.statsmodels.org/dev/tsa.htm>

- What are the characteristics of the development team?

- Do the teams charged with developing and maintaining this service reflect a diversity of opinions, backgrounds, and thought?

Our team includes statisticians, AI researchers, developers, as well as a group of social scientists that help us evaluate the outputs of our service for the diverse use cases of our customers. Our team is composed of individuals from a variety of socio-demographic backgrounds, with 32% women and 11% African American representation.

- Have you updated this FactSheet before?

- When and how often?
- What sections have changed?
- Is the FactSheet updated every time the service is retrained or updated?

We update the FactSheet with every service update, typically every 6 months. The following sections have changed in this FactSheet from the previous versions (available at URL): Statement of purpose (General), Basic performance, Safety (General, Fairness), and Lineage.

Usage

- What is the intended use of the service output?

- Briefly describe a simple use-case.

Media organizations that wish to identify and monitor related topics to an event of interest form a common group of customers. Consider a sports magazine wanting to monitor trending topics related to the US Open tennis tournament. Monitoring only a few known key-phrases associated with the US Open, will likely miss topics that might be trending during the tournament. Our client can use our service to identify and monitor related topics that were found trending while the tournament was taking place. Our client can monitor trending topics within different time intervals and that exhibit different structural characteristics, such as topics that have gained sudden attention compared to those that have gained attention in a more incremental fashion. Our client can monitor and collate topics from multiple social media platforms.

- What are the key procedures followed while using the service?

- How is the input provided? By whom?
- How is the output returned?

The client inputs a list of social media platforms that they wish to monitor for a specific set of

key-phrases associated with the event or topic of interest to them. Using these, for each social media platform, our service returns a list of relevant trending topics. The client can then examine each of these topics to see what key-phrases are relevant and to identify other potentially related key-phrases, social media messages, and trending topics.

Domains and applications

- What are the domains and applications the service was tested on or used for?

- Were domain experts involved in the development, testing, and deployment? Please elaborate.

Our service is often being used for brand monitoring on social media and for event monitoring by various local and regional media companies, where we work very closely with our customers as they understand better the subtleties around their brands and/or they often have a better grasp of the context of the events they are interested in monitoring.

- How is the service being used by your customers or users?

- Are you enabling others to build a solution by providing a cloud service or is your application end-user facing?
- Is the service output used as-is, is it fed directly into another tool or actuator, or is there human input/oversight before use?
- Do users rely on pre-trained/canned models or can they train their own models?
- Do your customers typically use your service in a time critical setup (e.g. they have limited time to evaluate the output)? Or do they incorporate it in a slower decision making process? Please elaborate.

Our service is typically integrated by our customers in their own in-house data monitoring, data gathering platforms, but can also be integrated into more application-tailored solutions by our customers. We encourage our customers

to qualitatively validate the output of the service, and we work closely with them as they integrate our service and use it. For our media customers, the outputs of our service help them better identify and contextualize the stories they cover, while our brand customers use them for marketing decisions and crisis management purposes.

- List applications that the service has been used for in the past.

- Please provide information about these applications or relevant pointers.
- Please provide key performance results for those applications.

Our service is used for brand monitoring and event monitoring. The specific details of each application where the service was used is customer-proprietary information.

- Other comments?

N/A.

Basic Performance

The following questions aim to offer an overall assessment of the service performance.

Testing by service provider

- Which datasets was the service tested on? (e.g., links to datasets that were used for testing, along with corresponding datasheets)

- List the test datasets and provide links to these datasets.
- Do the datasets have an associated datasheet? If yes, please attach.
- Could these datasets be used for independent testing of the service? Did the data need to be changed or sampled before use?

We release test results and benchmarks at the end of the year based on a set of popular trending topics that year and our predictions of how those trends will fare at different time stamps. Data about these trends and

comprehensive result reports are available at <https://datatrendly.com/reports/>.

Yes, we provide a datasheet for each end-of-year popular trending topics release. Each release contains the set of popular key-phrases for that year, timeseries corresponding to their popularity, and our prediction at different timestamps. See them attached at the end of this FactSheet.

These datasets can be used to further analyze our performance, but also to check it on our platform which allows retrospective browsing. In addition, our customers can check our predictions retrospectively for any other trend of interest.

- Describe the testing methodology.

- Please provide details on train, test and holdout data.
- What performance metrics were used? (e.g. accuracy, error rates, AUC, precision/recall)
- Please briefly justify the choice of metrics.

The performance metrics include: (1) scale dependent error rates that are often useful per case study basis, such as mean absolute error (MAE) and root mean square errors (RMSE); (2) percentage error rates as they allow for a better comparison of results between different timeseries, such as mean absolute percentage error (MAPE) and its symmetric version Armstrong (1978, p. 348), and (3) and scaled error rates that are preferred for comparisons of timeseries originating from different platforms or of different nature, explained in Hyndman & Koehler (2006). In addition, we also examine the residuals for any systematic trends, and that there is no correlation between residuals for which we use the Box-Pierce and Ljung-Box tests.

Hyndman, R. J., & Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International Journal of Forecasting*, 22, 679–688. <https://robjhyndman.com/publications/automatic-forecas>

Armstrong, J. S. (1978). *Long-range forecasting: From crystal ball to computer*. John Wiley & Sons.

- Describe the test results.

- Were latency, throughput, and availability measured?
- If yes, briefly include those metrics as well.

Our service makes real time predictions on how various trends may fare, and it is critical for us that our customers get access to the information they need timely and reliably. As a result, latency, throughput, and availability are critical metrics for a service like ours. For instance, our maximum latency to answer a query (between receiving a query and producing a result) is 2s.

Testing by third parties

- Is there a way to verify the performance metrics (e.g., via a service API)?

- Briefly describe how a third party could independently verify the performance of the service.
- Are there benchmarks publicly available and adequate for testing the service.

Yes, to some extent the performance metrics can be independently verified by third parties, if those third parties are our customers. Otherwise, it can be done only based on the data we release at the end of the year on popular trends that year and our corresponding predictions at different timestamps. For confidentiality and business reasons we do not allow third parties access to the work we do for our customers.

- In addition to the service provider, was this service tested by any third party?

- Please list all third parties that performed the testing.
- Also, please include information about the tests and test results.

No.

- Other comments?

No.

Safety

The following questions aim to offer insights about potential unintentional harms, and mitigation efforts to eliminate or minimize those harms.

General

- Are you aware of possible examples of bias, ethical issues, or other safety risks as a result of using the service?

- Were the possible sources of bias or unfairness analyzed?
- Where do they arise from: the data? the particular techniques being implemented? other sources?
- Is there any mechanism for redress if individuals are negatively affected?

We are not aware of broad ethical issues concerning our service. Some ethical issues may arise in the context of more sensitive applications such as identifying and monitoring trends related to anti-governmental movements. Although our service is not centered around identifying or making inference about individuals, one could use the trending topics to identify those posting about them. In this particular case, while we have a policy to not engage in such use-cases (policy that our customers are made aware of), our customers can use our service to monitor key-phrases beyond the close engagements we have with them. When we have any suspicion about the topics being monitored via our service, we block the usage of the related key-phrases.

In addition to such concerns, various biases in the data might skew the interpretation of the output we provide. Social media data is known not to be representative, and different social media platforms might exhibit different representational biases. The characteristics of each platform might also influence how users are likely to behave, such as what content they are likely to share. These biases may also evolve over time

depending on seasonal patterns, external circumstances, and because of changes in the user-base or in the features of each social media platform. These affect the type of insights that our customers can draw from the data and the trending topics we identify. For a comprehensive overview of data biases that our data tends to suffer from, see Olteanu et al. (2016, p. 6). We discuss these issues in detail with our customers, including how they might impact their analysis. We recommend qualitative analyses of the outputs, as well as tracking the same trends across multiple platforms.

Olteanu et al. “Social data: Biases, methodological pitfalls, and ethical boundaries.” SSRN (2016).

- Do you use data from or make inferences about individuals or groups of individuals. Have you obtained their consent?

- How was it decided whose data to use or about whom to make inferences?
- Do these individuals know that their data is being used or that inferences are being made about them? What were they told? When were they made aware? What kind of consent was needed from them? What were the procedures for gathering consent? Please attach the consent form to this declaration.
- What are the potential risks to these individuals or groups? Might the service output interfere with individual rights? How are these risks being handled or minimized?
- What trade-offs were made between the rights of these individuals and business interests?
- Do they have the option to withdraw their data? Can they opt out from inferences being made about them? What is the withdrawal procedure?

No, our service is not centered around making inferences about individuals or groups of individuals. Any analysis we make is content based, not user based. Given that we use only public data, we do not obtain explicit consent from the users of these platforms. However, some topics might be of interest to certain groups, and we

acknowledge that in certain use-cases (as mentioned above) this may lead to safety concerns. To minimize such risks, whenever there is a suspicion of such a use-case we block the use of related key-phrases on our service.

Explainability

- Are the service outputs explainable and/or interpretable?

- Please explain how explainability is achieved (e.g. directly explainable algorithm, local explainability, explanations via examples).
- Who is the target user of the explanation (ML expert, domain expert, general consumer, etc.)
- Please describe any human validation of the explainability of the algorithms

We do not provide explicit explanations for our inferences.

Fairness

- For each dataset used by the service: Was the dataset checked for bias? What efforts were made to ensure that it is fair and representative?

- Please describe the data bias policies that were checked (such as with respect to known protected attributes), bias checking methods, and results (e.g., disparate error rates across different groups).
- Was there any bias remediation performed on this dataset? Please provide details about the value of any bias estimates before and after it.
- What techniques were used to perform the remediation? Please provide links to relevant technical papers.
- How did the value of other performance metrics change as a result?

Although we report known statistics about the socio-demographic composition of each social media platform we work with our customers and

discuss with them about the type of conclusions they can draw from the outputs of our system, we do not perform our own bias checks.

- Does the service implement and perform any bias detection and remediation?

– Please describe model bias policies that were checked, bias checking methods, and results (e.g., disparate error rates across different groups).

– What procedures were used to perform the remediation? Please provide links or references to corresponding technical papers.

– Please provide details about the value of any bias estimates before and after such remediation.

– How did the value of other performance metrics change as a result?

No.

Concept Drift

- What is the expected performance on unseen data or data with different distributions?

– Please describe any relevant testing done along with test results.

N/A. We build a different model for each query submitted by our customers in real-time.

- Does your system make updates to its behavior based on newly ingested data?

– Is the new data uploaded by your users? Is it generated by an automated process? Are the patterns in the data largely static or do they change over time?

– Are there any performance guarantees/bounds?

– Does the service have an automatic feedback/retraining loop, or is there a human in the loop?

N/A.

- How is the service tested and monitored for model or performance drift over time?

– If applicable, describe any relevant testing along with test results.

Not purposefully. However, we keep track of our performance metrics for each use-case and application our service is used for over time. This also allows us to check for potential variations in performance over time as the characteristics of the data might vary with factors specific to each social media platform.

- How can the service be checked for correct, expected output when new data is added?

N/A.

- Does the service allow for checking for differences between training and usage data?

– Does it deploy mechanisms to alert the user of the difference?

N/A.

- Do you test the service periodically?

– Does the testing includes bias or fairness related aspects?

– How has the value of the tested metrics evolved over time?

Yes, as mentioned above, we keep track of our performance metrics for each use-case and application our service is used for.

- Other comments?

No.

Security

The following questions aim to assess the susceptibility to deliberate harms such as attacks by adversaries.

- How could this service be attacked or abused? Please describe.

Our service is used in a subscription mode, where the user registers and subscribes to a custom set of trending topics. This leaves us in control of the traffic and hence secure from attacks like those on typical SQL query servers.

- List applications or scenarios for which the service is not suitable.

- Describe specific concerns and sensitive use cases.
- Are there any procedures in place to ensure that the service will not be used for these applications?

We have procedures in place to disallow subscriptions by customers interested in monitoring ethically questionable topics such as hate speech or pornographic content.

- How are you securing user or usage data?

- Is usage data from service operations retained and stored?
- How is the data being stored? For how long is the data stored?
- Is user or usage data being shared outside the service? Who has access to the data?

Yes, usage data is stored per tenant for a period of 2 weeks. The data is used for understanding usage patterns, helpful for improving the service. Only our own data science team has access to this data. The data is stored in encrypted format on our servers.

- Was the service checked for robustness against adversarial attacks?

- Describe robustness policies that were checked, the type of attacks considered, checking methods, and results.

Yes. We do have algorithms in place to discard “wrong” feedback data which may deteriorate the performance of the service.

- What is the plan to handle any potential security breaches?

- Describe any protocol that is in place.

Give the short retention period of customer data and the nature of our service, we believe that there is both a low risk of security breaches and they have limited ramifications. If those happen, and limited customer data is compromised or leaked, we will notify anyone affected immediately.

- Other comments?

No.

Lineage

The following questions aim to overview how the service provider keeps track of details that might be required in the event of an audit by a third party, such as in the case of harm or suspicion of harm.

Training Data

- Does the service provide an as-is/canned model? Which datasets was the service trained on?

- List the training datasets.
- Where there any quality assurance processes employed while the data was collected or before use?
- Were the datasets used for training built-for-purpose or were they repurposed/adapted? Were the datasets created specifically for the purpose of training the models offered by this service?

N/A. We build a different model for each use-case and trend. We generate our time series based on the queries submitted by our customers; thus, they are built and used for the purpose for which they were generated.

- For each dataset: Are the training datasets publicly available?

- Please provide a link to the datasets or the source of the datasets.

No. We only release data and results corresponding to the most popular trends at the end of the year. We do not release data corresponding to customers' engagements. We do prepare datasheets for every data release.

- For each dataset: Does the dataset have a datasheet or data statement?

- If available, attach the datasheet; otherwise, provide answers to questions from the datasheet as appropriate [to insert citation]

N/A.

- Did the service require any transformation of the data in addition to those provided in the datasheet?

N/A.

- Do you use synthetic data?

- When? How was it created?
- Briefly describe its properties and the creation procedure.

No.

Trained Models

- How were the models trained?

- Please provide specific details (e.g., hyper-parameters).

To make a forecast for a given time slot in the future and a given trend of interest, we use the historical behavior of the time series corresponding to this trend, as well as other historical information from related trends on a given social media platform or from several platforms, which can be either automatically extracted or hand selected by experts in our customer teams, or a combination of both.

We build a different model for each use-case and trend.

- When were the models last updated?

- How much did the performance change with each update?
- How often are the models retrained or updated?

N/A.

- Did you use any prior knowledge or re-weight the data in any way before training?

In some cases, yes, we do. We incorporate in our models historical information about time series selected by domain experts from our client teams that are expected to share some relationship with the trends of interests.

- Other comments?

No.