Amazon SageMaker Clarify: Machine Learning Bias Detection and Explainability in the Cloud

Michaela Hardt¹, Xiaoguang Chen, Xiaoyi Cheng, Michele Donini, Jason Gelman, Satish Gollaprolu, John He, Pedro Larroy, Xinyu Liu, Nick McCarthy, Ashish Rathi, Scott Rees, Ankit Siva, ErhYuan Tsai², Keerthan Vasist, Pinar Yilmaz, Muhammad Bilal Zafar, Sanjiv Das³, Kevin Haas, Tyler Hill, Krishnaram Kenthapadi
Amazon Web Services

ABSTRACT

Understanding the predictions made by machine learning (ML) models and their potential biases remains a challenging and labor-intensive task that depends on the application, the dataset, and the specific model. We present Amazon SageMaker Clarify, an explainability feature for Amazon SageMaker that launched in December 2020, providing insights into data and ML models by identifying biases and explaining predictions. It is deeply integrated into Amazon SageMaker, a fully managed service that enables data scientists and developers to build, train, and deploy ML models at any scale. Clarify supports bias detection and feature importance computation across the ML lifecycle, during data preparation, model evaluation, and post-deployment monitoring. We outline the desiderata derived from customer input, the modular architecture, and the methodology for bias and explanation computations. Further, we describe the technical challenges encountered and the tradeoffs we had to make. For illustration, we discuss two customer use cases. We present our deployment results including qualitative customer feedback and a quantitative evaluation. Finally, we summarize lessons learned, and discuss best practices for the successful adoption of fairness and explanation tools in practice.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**: Distributed algorithms; • **Software and its engineering** → Software notations and tools.

KEYWORDS

Machine learning, Fairness, Explainability

ACM Reference Format:

Michaela Hardt¹, Xiaoguang Chen, Xiaoyi Cheng, Michele Donini, Jason Gelman, Satish Gollaprolu, John He, Pedro Larroy, Xinyu Liu, Nick McCarthy, Ashish Rathi, Scott Rees, Ankit Siva, ErhYuan Tsai², Keerthan Vasist, Pinar Yilmaz, Muhammad Bilal Zafar, Sanjiv Das³, Kevin Haas, Tyler Hill, Krishnaram Kenthapadi. 2021. Amazon SageMaker Clarify: Machine Learning

¹Correspondence to: Michaela Hardt, <milaha@amazon.com>.

²Work done while at Amazon Web Services.

³Santa Clara University, CA, USA.

This work is licensed under a Creative Commons Attribution-NonCommercial-NoDerivs International 4.0 License.

KDD '21, August 14–18, 2021, Virtual Event, Singapore

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8332-5/21/08.

<https://doi.org/10.1145/3447548.3467177>

Bias Detection and Explainability in the Cloud. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD '21), August 14–18, 2021, Virtual Event, Singapore*. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3447548.3467177>

1 INTRODUCTION

Machine learning (ML) models and data-driven systems are increasingly used to assist in decision-making across domains such as financial services, healthcare, education, and human resources. Benefits of using ML include improved accuracy, increased productivity, and cost savings. The increasing adoption of ML is the result of multiple factors, most notably ubiquitous connectivity, the ability to collect, aggregate, and process large amounts of data using cloud computing, and improved access to increasingly sophisticated ML models. In high-stakes settings, tools for bias and explainability in the ML lifecycle are of particular importance.

Regulatory [1, 2]: Many ML scenarios require an understanding of why the ML model made a specific prediction or whether its prediction was biased. Recently, policymakers, regulators, and advocates have expressed concerns about the potentially discriminatory impact of ML and data-driven systems, for example, due to inadvertent encoding of bias into automated decisions. Informally, biases can be viewed as imbalances in the training data or the prediction behavior of the model across different groups, such as age or income bracket. **Business** [3, 17]: The adoption of AI systems in regulated domains requires explanations of how models make predictions. Model explainability is particularly important in industries with reliability, safety, and compliance requirements, such as financial services, human resources, healthcare, and automated transportation. Using a financial example of lending applications, loan officers, customer service representatives, compliance officers, and end users often require explanations detailing why a model made certain predictions. **Data Science**: Data scientists and ML engineers need tools to debug and improve ML models, to ensure a diverse and unbiased dataset during feature engineering, to determine whether a model is making inferences based on noisy or irrelevant features, and to understand the limitations and failure modes of their model [42].

While there are several open-source tools for fairness and explainability, we learned from our customers that developers often find it tedious to incorporate these tools across the ML workflow, as this requires a sequence of complex, manual steps: (1) Data scientists and ML developers first have to translate internal compliance requirements into code that can measure bias in their data and ML models. To do this, they often must exit their ML workflow, and write custom code to piece together open source packages. (2) Next, they need to test datasets and models for bias across all steps

of the ML workflow and write code for this integration. (3) Once the tests are complete, data scientists and ML developers create reports showing the factors that contribute to the model prediction by manually collecting a series of screenshots for internal risk and compliance teams, auditors, and external regulators. (4) Finally, after the models are in production, data scientists and ML developers must periodically run through a series of manual tests to ensure that no bias is introduced as a result of a drift in the data distribution (e.g., changes in income for a given geographic region). Similar challenges have also been pointed out in prior studies [8, 20, 33].

Our contributions: To address the need for generating bias and explainability insights across the ML workflow, we present Amazon SageMaker Clarify¹, a new functionality in Amazon SageMaker [27] that helps detect bias in data and models, helps explain predictions made by models, and generates associated reports. We describe the design considerations and technical architecture of Clarify, and its integration across the ML lifecycle (§3). We then discuss the methodology for bias and explainability computations (§4). We implemented several measures of bias in Clarify, considering the dataset (pre-training) and the model predictions (post-training), to support various use cases of SageMaker’s customers. We integrated pre-training bias detection with the data preparation functionality and post-training bias detection with the model evaluation functionality in SageMaker. These biases can be monitored continuously for deployed models and alerts can notify users about changes. For model explainability, Clarify offers a scalable implementation of SHAP (SHapley Additive exPlanations) [29], based on the concept of the Shapley value [38]. In particular, we implemented the KernelSHAP [29] variant as it is model agnostic, so that Clarify can be applied to the broad range of models trained by SageMaker’s customers. For deployed models, Clarify monitors for changes in feature importance and issues alerts.

We discuss the implementation and technical challenges in §6 and illustrate the graphical representation in SageMaker Studio in a case study in §7. In §8 we present the deployment results, including a quantitative evaluation of the scalability of our system, a report on usage across industries, qualitative customer feedback on usability, and two customer use cases. Finally, we summarize lessons learned in §9 and conclude in §10.

2 RELATED WORK

Clarify offers bias detection and explainability of a model. Both are active research areas that we review briefly below.

Bias. Proposed bias metrics for machine learning include individual fairness [14, 39] as well as group fairness. Group fairness considers different aspects of model quality (e.g., calibration [24], error rates [19]), and prediction tasks (e.g., ranking [41] or classification [12, 47]). Techniques based on causality [34] provide promising alternatives at the individual [25] and group level [22]. For a thorough overview see [5, 31].

Techniques to mitigate bias either change the model, its optimization, (e.g., [4, 35, 48]), or post-process its predictions (e.g., [16, 19]). There are several open source repositories that offer bias management approaches, including fairness-comparison², Aequitas³,

Themis⁴, responsibly⁵, and LiFT [44]. Moreover, IBM’s AI Fairness 360⁶, Microsoft’s Fairlearn⁷, and TensorFlow’s What-If-Tool [46] are products built around bias metrics and mitigation strategies.

Explainability. The approaches for model explainability can be dissected in various ways. For a comprehensive survey, see e.g. [18]. We distinguish *local* vs. *global explainability* approaches. Local approaches [30, 36] explain model behavior in the neighborhood of a given instance. Global explainability methods aim to explain the behavior of the model at a higher-level (e.g. [26]). Explainability methods can also be divided into *model-agnostic* (e.g., [30]) and *model-specific* approaches which leverage internals of models such as gradients (e.g., [40, 43]). There is growing impetus to base explanations on counterfactuals [6, 21]. At the same time concern about overtrusting explanations grows [28, 37].

In comparison to previous work, Clarify is deeply integrated into SageMaker and couples scalable fairness and explainability analyses. For explainability, its implementation of SHAP offers model-agnostic feature importance at the local and global level.

3 DESIGN, ARCHITECTURE, INTEGRATION

3.1 Design Desiderata

Conversations with customers in need of detecting biases in models and explaining their predictions shaped the following desiderata.

- **Wide Applicability.** Our offerings should be widely applicable across (1) industries (e.g., finance, human resources, and healthcare), (2) use cases (e.g., click-through-rate prediction, fraud detection, and credit risk modeling), (3) prediction tasks (e.g., classification and regression), (4) ML models (e.g., logistic regression, decision trees, deep neural networks), and (5) the ML lifecycle (i.e., from dataset preparation to model training to monitoring deployed models).
- **Ease of Use.** Our customers should not have to be experts in fairness and explainability, or even in ML. They should be able to analyze bias and obtain feature importance within their existing workflow having to write little or no code.
- **Scalability.** Our methods should scale to both complex models (i.e., billions of parameters) and large datasets.

The latter two desiderata were the top two non-functional requirements requested as customers evaluated the early prototypes. Customers described frustratingly long running times with existing open source techniques, which motivated us to scale these existing solutions to large datasets and models (c.f. §8.2). Moreover, feedback from early prototypes included frequent erroneous configurations due to the free-form nature of the configuration file, prompting us to improve usability (c.f. §8.3).

3.2 Architecture of Clarify

Clarify is implemented as part of Amazon SageMaker [27], a fully managed service that enables data scientists and developers to build, train, and deploy ML models at any scale. It leverages SageMaker’s processing job APIs, and executes batch processing jobs on a cluster of AWS compute instances in order to process large

¹<https://aws.amazon.com/sagemaker/clarify>

²<https://github.com/algofairness/fairness-comparison>

³<https://github.com/dssg/aequitas>

⁴<https://github.com/LASER-UMASS/Themis>

⁵<https://github.com/ResponsibleAI/responsibly>

⁶<https://aif360.mybluemix.net/>

⁷<https://fairlearn.github.io/>

Figure 1: High-level system architecture of Clarify. A dataset and a configuration file serve as input to a batch processing job that produces bias metrics and feature importance. These outputs are stored in a filesystem, exported as metrics, and visualized in a web-based IDE for ML.

amounts of data in a scalable manner. Each processing job has an associated cluster of fully managed SageMaker compute instances running the specified container image and provisioned specifically for the processing job. Clarify provides a first-party container image which bundles efficient and optimized implementations of core libraries for bias and explainability computation together with flexible parsing code to support a variety of formats of datasets and the model input and output.

Inputs to Clarify are the dataset stored in the S3 object store and a configuration file specifying the bias and explainability metrics to compute and which model to use. At a high level, the processing job completes the following steps: (1) validate inputs and parameters; (2) compute pre-training bias metrics; (3) compute post-training bias-metrics; (4) compute local and global feature attributions for explainability; (5) generate output files. The job results are saved in the S3 object store and can be consumed in multiple ways: (1) download and programmatically inspect using a script or notebook; (2) visualize and inspect in SageMaker Studio IDE through custom visualizations and tooltips to describe and help interpret the results.

The cluster of compute instances carries out the computation in a distributed manner (c.f. §6). To obtain predictions from the model, the Clarify job spins up model serving instances that take inference requests. This is a deliberate choice made to avoid interfering with production traffic to the already deployed models and polluting the performance and quality metrics of the models being monitored in production. This resulted in a modular system allowing customers to use our product without any changes to their deployed models; see Figure 1.

Security. Built on top of existing SageMaker primitives, Clarify supports access control and encryption for the data and the network. Specifically, roles, virtual private cloud configurations, and encryption keys can be specified.

3.3 Deep Integration in SageMaker

For generating bias and explainability insights across the ML lifecycle, we integrated Clarify with the following SageMaker components: Data Wrangler to visualize dataset biases during data preparation, Studio & Experiments to explore biases and explanations of trained models, and Model Monitor to monitor these

metrics.⁸ Figure 2 illustrates the stages during ML lifecycle and the integration points of Clarify within SageMaker.

4 METHODS

4.1 Bias

We implemented several bias metrics to support a wide range of use cases. The selection of the appropriate bias metric for an application requires human judgement. Some of the metrics only require the dataset and thus can be computed early in the model lifecycle, that we call *pre-training bias metrics*. They can shed light onto how different groups are represented in the dataset and if there is a difference in label distribution across groups. At this point some biases may require further investigation before training a model. Once a model is trained, we can compute *post-training bias metrics* that take into account how its predictions differ across groups. We introduce the notation before defining the metrics.

4.1.1 Setup and Notation. Without loss of generality, we use d to denote the group that is potentially disadvantaged by the bias and compare it to the remainder of the examples in the dataset a (for advantaged group). The selection of the groups often requires domain information⁹ and application specific considerations. Denote the number of examples of each group by n_a, n_d , with n denoting the total. We assume that each example in the dataset has an associated binary label $y \in \{0, 1\}$ (with 1 denoting a positive outcome), and that the ML model predicts $\hat{y} \in \{0, 1\}$. While these assumptions seem restrictive, different prediction tasks can be re-framed accordingly. For example, for regression tasks we can select one or more thresholds of interest. For multi-class predictions we can treat various subsets of the classes as positive outcomes.

We denote by $n^{(0)}, n^{(1)}$, respectively, the number of labels of value 0, 1 that we can further break down by group into $n_a^{(0)}, n_d^{(0)}$ and $n_a^{(1)}, n_d^{(1)}$. Analogously, we can count the positive / negative predictions $\hat{n}^{(0)}, \hat{n}^{(1)}$ and break them down by group into $\hat{n}_a^{(0)}, \hat{n}_d^{(0)}$ and $\hat{n}_a^{(1)}, \hat{n}_d^{(1)}$. These counts can be normalized by group size as $q_a = n_a^{(1)} / n_a$ and $q_d = n_d^{(1)} / n_d$ for the labels and $\hat{q}_a = \hat{n}_a^{(1)} / n_a$ and $\hat{q}_d = \hat{n}_d^{(1)} / n_d$ for the predictions. Additionally, we consider the TP : true positives; FP : false positives; TN : true negatives; FN : false negatives; and compute them for each group separately.

4.1.2 Pre-Training Metrics.

CI Class Imbalance: Bias can stem from an under-representation of the disadvantaged group in the dataset, for example, due to a sampling bias. We define $CI = \frac{n_a - n_d}{n}$.

DPL Difference in positive proportions in observed labels: We define $DPL = q_a - q_d$, to measures bias in the label distribution. Additional metrics derived from q_a and q_d are defined in §A.1¹⁰.

CDDL Conditional Demographic Disparity in Labels: Can the difference in q_a and q_d be explained by another feature? An instance of Simpson's paradox¹¹ arose in the case of Berkeley

⁸Data Wrangler: aws.amazon.com/sagemaker/data-wrangler/; Studio: aws.amazon.com/sagemaker/studio/; Model Monitor: aws.amazon.com/sagemaker/model-monitor/

⁹<https://plato.stanford.edu/entries/discrimination/>

¹⁰<https://pages.awscloud.com/rs/112-TZM-766/images/Amazon.AI.Fairness.and.Explainability.Whitepaper.pdf>

¹¹<https://www.britannica.com/topic/Simpsons-paradox>

Figure 2: The stages in a ML lifecycle from problem formulation to deployment and monitoring with their respective fairness considerations in orange. Clarify can be used at dataset construction and model testing and monitoring to investigate them.

admissions where men were admitted at a higher rate than women overall, but for each department women were admitted at a higher rate. In this case women applied more to departments with lower acceptance rates. Following [45], we define $DD = n_d^{(0)} / n^{(0)} - n_d^{(1)} / n^{(1)}$ and compute it across all strata i on a user-supplied attribute that we want to control for, yielding DD_i . $CDDL = \frac{1}{n} \sum_i n_i \cdot DD_i$ where i subscripts each strata and n_i is the number of examples in it.

4.1.3 Post-Training Metrics.

DPPL Difference in positive proportions in predicted labels: DPPL consider the difference in positive predictions: $DPPL = \hat{q}_a - \hat{q}_d$, related is “statistical parity” [7, 10].

DI Disparate Impact: Instead of the difference as in DPPL we can consider the ratio: $DI = \hat{q}_d / \hat{q}_a$.

DCA Difference in Conditional Acceptance / Rejection: DCA considers the ratio of the number of positive labels to that of positive predictions assessing calibration: $DCA = n_a^{(1)} / \hat{n}_a^{(1)} - n_d^{(1)} / \hat{n}_d^{(1)}$. Analogously, DCR considers negative labels and predictions.

AD Accuracy Difference: We compare the accuracy, i.e. the fraction of examples for which the predictions is equal to the label, across groups. We define $AD = (TP_a + TN_a) / n_a - (TP_d + TN_d) / n_d$.

RD Recall Difference: We compare the recall (the fraction of positive examples that receive a positive prediction) across groups. $RD = TP_a / n_a^{(1)} - TP_d / n_d^{(1)}$ (aka “equal opportunity” [19]).

DAR Difference in Acceptance / Rejection Rates: This metric considers precision, i.e. the fraction of positive predictions that are correct: $DAR = TP_a / \hat{n}_a^{(1)} - TP_d / \hat{n}_d^{(1)}$. Analogously, we define DRR for negative predictions.

TE Treatment Equality: We define $TE = FN_d / FP_d - FN_a / FP_a$ [7].

CDDPL Conditional Demographic Disparity of Predicted Labels: This metric is analogous to the CDDL for predictions.

See §A.1 for additional metrics.

4.2 Explainability

We implemented the KernelSHAP [29] algorithm due to its good performance [36] and wide applicability.

The KernelSHAP algorithm [29] works as follows: Given an input $\mathbf{x} \in \mathbb{R}^M$ and the model output $f(\mathbf{x}) \in \mathbb{R}$, SHAP, building on the Shapley value framework [38], designates the prediction task as a M -player cooperative game, with individual features being players and the predicted output $f(\mathbf{x})$ being the total payout of the game. The framework decides how *important* each player is to the game’s outcome. Specifically, the importance of the i^{th} player $\mathbf{x}_i$ is defined as the average marginal contribution of the player to a game not containing $\mathbf{x}_i$. Formally, the contribution ϕ_i is defined as:

$$\phi_i(\mathbf{x}) = \sum_S \frac{|S|!(M - |S| - 1)!}{M!} (f(\mathbf{x}_S \cup \mathbf{x}_i) - f(\mathbf{x}_S)),$$

where the sum is computed over all the player coalitions S not containing $\mathbf{x}_i$. The computation involves iterating over 2^M feature coalitions, which does not scale to even moderate values of M , e.g., 50. To overcome these scalability issues, KernelSHAP uses approximations [29].

4.3 Monitoring bias and explanations

Measuring bias metrics and feature importance provide insights into the model at time of evaluation. However, a change in data distribution can occur over time, resulting in different bias metrics and feature importance. Our goal is to continuously monitor for such changes and raise alerts if bias metrics or feature importance drift from *reference values* established during an earlier evaluation (before the model launch). Based on a user-specified schedule (e.g., hourly), newly accumulated data is analyzed and compared against the reference values.

Bias. Given a bias metric from §4.1, we let the user specify a range of reference values $A = (a_{\min}, a_{\max})$ in which the new bias value b computed on the live data should lie. However, if little data was accumulated simply checking for containment can lead to noisy results. To ensure that the conclusions drawn from the observed

live data are statistically significant, we check for overlap of A with the 95% bootstrap confidence interval of b [15].

Explainability. For feature importance we consider the change in ranking of features by their importance. Specifically, we use the Normalized Discounted Cumulative Gain (nDCG) score for comparing the ranking by feature importance of the reference data with that based on the live data. nDCG penalizes changes further down in the ranking less. Also, swapping the position of two values is penalized relative to the difference of their original importance.

Let $F = [f_1, \dots, f_M]$ denote the list of M features sorted w.r.t. their importance scores in the reference data, and $F' = [f'_1, \dots, f'_M]$ denote the new ranking of the features based on the live data. We use $a(f)$ to describe a function that returns the feature importance score on the reference data of a feature f . Then, nDCG is defined as: $\text{nDCG} = \text{DCG}/\text{iDCG}$, where $\text{DCG} = \sum_{i=1}^M a(f'_i)/\log_2(i+1)$ and $\text{iDCG} = \sum_{i=1}^M a(f_i)/\log_2(i+1)$.

An nDCG value of 1 indicates no change in the ranking. In our implementation, an nDCG value below 0.90 triggers an alert.

5 ALTERNATIVE DESIGN AND METHODS

On the architectural side, our design provisions a separate model server which adds latency. Alternatively, a design that runs the model locally avoids network traffic. We decided to use a model server so that we can re-use the existing SageMaker model hosting primitives. Our model-agnostic approach enables support of a wide range of models including those that cannot be executed locally. This way we can scale and tune the instance type selection and counts for the processing job and the model serving instances independently based on the dataset and model characteristics.

For the bias metrics, we considered limiting the number of metrics, to avoid overwhelming users. However, the appropriate bias metric is application dependent. To support customers from different industries with various use cases, we decided to implement a wide range of metrics (see §4.1).

For explainability we decided to implement the KernelSHAP [29] algorithm (see §4.2) rather than alternatives discussed in §2 for a number of reasons: (i) it offers wide *applicability* to interpret predictions of any score-based model; (ii) it works in a *model-agnostic* manner and does not require access to model internals such as gradients, allowing us to build on top of the SageMaker hosting primitives; (iii) it provides attractive theoretical properties associated with the Shapley value framework [38]; and (iv) benchmarking shows that it performs favorably compared to related methods such as LIME [36]. More efficient model-specific implementations like TreeSHAP and DeepSHAP [29] leverage the model architecture for improved efficiency. By using the model-agnostic KernelSHAP, we traded efficiency for wider applicability. In the future we will prioritize support for additional explainability methods based on customer requests.

6 DEVELOPMENT

Timeline: We started by collecting customer requirements to derive Clarify's goals and design the system (§3). Our first prototype implemented the methods on a single node packaged into a container that could be run with a bare-bone configuration file. After internal tests, we deployed this version for beta customers to solicit

feedback. This feedback helped us to prioritize scalability (§8.2) and ease of use improvements (§8.3). Afterwards, we expanded support of models and datasets, conducted more end-to-end tests, and obtained additional rounds of customer feedback. We embarked on the integration across the model lifecycle (§3.3), emphasizing the associated UI components and documentation.¹² Prior to public launch, we set up pipelines and dashboards for health monitoring of our system, including alarms for too many failed jobs. Throughout, multiple reviews took place for the design, usability, security and privacy, and operational readiness.

Implementation Details: For wide applicability, we implemented flexible parsers to support various data formats (JSONLines/CSV, with(-out) header) and model signatures, see §A.2 for examples.

The first bias metrics were implemented in numpy.¹³ For scalability, all subsequent methods were implemented using Spark via SageMaker's Spark Processing job. Spark provided us with convenient abstractions (e.g., distributed median computations) and managed the cluster for us. Our bias metrics mostly require counts that are easy to distribute. Distributing SHAP only requires one map operation to calculate feature importance for each example followed by an aggregation for global model feature importance.

6.1 Challenges and Resolutions

A challenge to fault-tolerance is that Clarify's offline batch processing job hits a live model endpoint. This requires a careful implementation to avoid overwhelming the model endpoint with requests. To limit the potential impact, Clarify spins up separate model servers and does not interfere with production servers. To avoid overwhelming these separate model endpoints, we dynamically configured the batch size through experimentation, starting with the largest possible batch size based on the maximum payload¹⁴ in order to achieve high throughput at the model endpoints. Furthermore, we relied on the boto clients' retry policies¹⁵ with exponential back-off. This is necessary to tolerate intermittent inference call failures. We further expanded the retry mechanism, since not all errors are considered retrievable in the boto client.

A key to scalability was to utilize resources well by carefully tuning the number of partitions of the dataset in Spark: If tasks took very long (we started out with tasks taking an hour) for SHAP computation, failures were costly and utilization during the last few tasks was low. If tasks were too small for bias metrics computation, throttling limitations when instantiating a SageMaker Predictor led to job failures. The solution was to carefully tune the number of tasks separately for explainability and bias computations.

7 ILLUSTRATIVE CASE STUDY

For illustration we present a simple case study with the German Credit Data dataset from the UCI repository [13] that contains 1000 labeled credit applications. Features include credit purpose,

¹²<https://docs.aws.amazon.com/sagemaker/latest/dg/clarify-detect-post-training-bias.html>, <https://docs.aws.amazon.com/sagemaker/latest/dg/clarify-model-explainability.html>

¹³Our open source implementation is available at <https://github.com/aws/amazon-sagemaker-clarify>.

¹⁴<https://github.com/boto/botocore/blob/develop/botocore/data/sagemaker-runtime/2017-05-13/service-2.json#L35>

¹⁵<https://boto3.amazonaws.com/v1/documentation/api/latest/guide/retries.html>

Figure 3: Pre-training bias results in Data Wrangler.

credit amount, housing status, and employment history. Categorical attributes are converted to indicator variables, e.g., A30 means “no credits taken.” We mapped the labels (Class1Good2Bad) to 0 for poor credit and 1 for good.

Bias analysis. Before we train the model we seek to explore biases in the dataset. Figure 3 shows Clarify integrated into Data Wrangler. Convenient drop-downs are available to select the label column and the group. Also, by default a small set of metrics are pre-selected in order to not overwhelm customers with too many metrics. No code needs to be written to conduct this analysis. From the class imbalance value of -0.94, we infer that the dataset is very imbalanced, with most examples representing foreign workers. However, as *DPL* is close to 0, there is not much difference in the label distributions across the two groups. With these encouraging results we proceed with the training of an XGBoost model. We submit the configuration for the post-training bias analysis (see Appendix A) as input along with the validation data to the Clarify job. The results shown in Figure 4 reveal a mixed picture with positive and negative metrics. We start by inspecting the disparate impact *DI* with a value of 0.73 much below 1. Since the label rates are almost the same across the two groups (*DPL*), *DI* is driven by differences in precision and recall. We see that foreign workers have a lower recall (*RD* of 0.13), but a higher precision (*DAR* of -0.19). Corrective steps can mitigate this bias prior to deployment.

Feature importance. To understand which features are affecting the model’s predictions the most we look at the global SHAP values that are averaged across all examples after taking the absolute value. The three most important features are A61 (having very little savings), the loan amount, and credit duration, see Figure 5.

8 DEPLOYMENT RESULTS

In the first few months after launch 36% of all Clarify jobs computed bias metrics while 64% computed feature importance, see Table 1. Monitoring jobs made up 42% of jobs, illustrating that customers reached a mature stage in the ML lifecycle in which they incorporated bias and feature importance monitoring. This is expected in a healthy development cycle; after a few iterations on

Bias metric	Bias value
Conditional Demographic Disparity in Labels (CDDL)	-0.0086
Class Imbalance (CI)	-0.94
Difference in Positive Proportions in Labels (DPL)	-0.055
Jensen-Shannon Divergence (JS)	0.0034
Kullback-Liebler Divergence (KL)	0.0073
Kolmogorov-Smirnov Distance (KS)	0.055
L-p Norm (LP)	0.078
Total Variation Distance (TVD)	0.055
Accuracy Difference (AD)	-0.14
Conditional Demographic Disparity in Predicted Labels (CDDPL)	0.045
Difference in Acceptance Rates (DAR)	-0.19
Difference in Conditional Acceptance (DCA)	-0.32
Difference in Conditional Outcomes (DCR)	-∞
Disparate (Adverse) Impact (DI)	0.73
Difference in Positive Proportions in Predicted Labels (DPPL)	0.27
Difference in Rejection Rates (DRR)	0.65
Counterfactuals: Flptest (FT)	-0.27
Recall Difference (RD)	0.13
Treatment Equality (TE)	0.9

Figure 4: Post-training bias results in Studio Trials.

Figure 5: SHAP feature importance.

	evaluation	monitoring	overall
bias	0.15	0.21	0.36
explainability	0.43	0.21	0.64
overall	0.58	0.42	1

Table 1: Break-down of Clarify jobs across method (bias / explainability) and phase (evaluation / monitoring).

model training and evaluation, a model passes the launch criteria and is continuously monitored after.

Next, we evaluate our design goals and highlight two use cases.

8.1 Wide Applicability

Customers across several industry verticals use Clarify, including digital publishing, education, electronics, finance, government,

method	processing time in minutes	cost in dollars
pre-training bias	1	\$0.03
post-training bias	14	\$0.13

Table 2: Processing time and cost of bias metrics computation for 1 million examples.

number of examples	processing time in minutes	cost in dollars
1000	16	\$0.13
10,000	90	\$0.71
100,000	811	\$6.43

Table 3: Processing time and cost of SHAP with varying dataset sizes using one processing / model instance.

hardware, healthcare, IT, manufacturing, news, sports, staffing, and telecommunications. Customers use Clarify for binary predictions, multi-class predictions, and ranking tasks. They are located in USA, Japan, UK, and Norway among other countries. We present two specific customer use cases in §8.4.

8.2 Scalability

We conducted experiments measuring the time taken to compute bias metrics and feature importance scores. We oversampled the German credit data set described in §7 creating up to 1 million examples to study scalability. We ran experiments on `ml.c5.xlarge` machines which are suitable for this workload and provide cost estimates based on AWS pricing in the `us-west-2` region.

Bias metrics. Table 2 shows that the bias metrics computation is fast even for 1M examples. Post-training bias metrics are considerably slower than pre-training metrics since they require model inference. The reported time includes 3-5 minutes needed to spin up the model endpoint. Scalability issues may arise when the dataset does not fit in memory (in which case a larger instance type can be chosen) or when model inference is much slower (in which case the instance count can be increased).

Explainability. Table 3 shows that time and cost scale roughly linearly with the dataset size. The implementation of SHAP iterates over all examples in a for-loop and the execution at each step dominates the time and cost. For 100,000 examples the running time of more than 13h is so long that it can be a point of frustration for customers. We parallelized SHAP to reduce the running time.

Figure 6 shows the scaling behavior of SHAP on 100,000 examples. We see that increasing the number of instances (both for the processing job and the model endpoint) to two reduces the time down to 3h and 17 minutes. This super-linear speedup is surprising. It is due to the fact that our single instance implementation does not fully utilize all cores. Our Spark implementation not only splits the workload across machines but also across cores within those machines thereby reducing not just time but even cost. We observe that 1 model endpoint instance for every 1 to 2 processing instances provides a good ratio: A higher ratio underutilizes model endpoint instances, increasing cost without a speed up. With a lower ratio (e.g., 1 model endpoint instance for 10 processing instances), the model becomes the bottleneck, and cost as well as running time can be reduced by increasing the number of model endpoint instances.

method	time in minutes		
	P50	P90	P99
bias	1	8	11
explainability	5	12	103
overall	4	12	101

Table 4: Percentiles of processing times in minutes.

Customer results. Half of the Clarify jobs of customers take at most 4 minutes. Table 4 shows percentiles of processing times for bias as well as explainability jobs. We see that the running time is less than two hours for 99% of the jobs. Across all jobs 35% use multiple processing instances thereby reducing running time. At most 10 instances were used by customers so far.

8.3 Ease of use

We obtained qualitative customer feedback on usability. For customers it was not always easy to get the configuration right despite documentation. Difficulties included setting the right fields with the right values, and figuring out fields that are optional, or need to go together for a specific use case. These difficulties were reduced through tooling, documentation, and examples. In terms of tooling, customers can use the explanatory UI or programmatically create the configuration. Regarding documentation, customers requested additional documentation covering specifics of multi-class prediction. Regarding examples, multiple customers with strict security requirements who used advanced configuration options to enable encryption of data at rest and in transit with fine-grained access control benefited from Clarify examples for such settings. Other customers noted that our notebooks focus on the CSV format and requested examples with the JSONLines data format. These examples illustrate tension between the goals of wide applicability and usability: Ease of use for one setting (e.g., CSV data, binary predictions) does not extend to all settings (e.g., JSONLines data, multi-class prediction).

We received several inquiries about formulating the development, business, and regulatory needs into the framework that Clarify provides. While we were able to find a fit for many different use cases, these questions are hard and answers differ case-by-case. For example, the question of which bias metric is important, depends on the application. There is no way to get around this question, as it is impossible to remove all biases [11, 24]. Here again, documentation helps customers to address this hard question.

Quantitative customer results. SageMaker AutoPilot, that automatically builds, trains, and tunes models, integrated Clarify in March 2021. AutoPilot contributes 54% of the explainability jobs. The automatic configuration functionality offered by AutoPilot, along with the other improvements listed above, helped reduce failure rates by 75%.

8.4 Two Customer Use Case Examples

Next, we illustrate the use of Clarify in production by providing two real customer use cases.

Fraud detection. Zopa’s data scientists use several dozens of input features such as application details, device information, interaction behaviour, and demographics. For model building, they

Figure 6: Scaling of SHAP on 100,000 examples and measuring time (left) and cost (right) for a varying number of instances.

extract the training dataset from their Amazon Redshift data warehouse, perform data cleaning, and store into Amazon S3. As Zopa has their own in-house ML library for both feature engineering and ML framework support, they use the “Bring Your Own Container” (BYOC) approach to leverage SageMaker’s managed services and advanced functionalities such as hyperparameter optimization. The optimized models are then deployed through a Jenkins CI/CD pipeline to the production system and serve as a microservice for real-time fraud detection as part of their customer facing platform.

Explanations are obtained both during model training for validation, and after deployment for model monitoring and generating insights for underwriters. These are done in a non-customer facing analytical environment due to a heavy computational requirement and a high tolerance of latency. SageMaker Multi Model Server stack is used in a similar BYOC fashion, to register the models for the Clarify processing job. Clarify spins up an ephemeral model endpoint and invokes it for millions of predictions on synthetic contrastive data. These predictions are used to compute SHAP values for each individual example, which are stored in S3.

For model validation and monitoring, Clarify aggregates these local feature importance values across examples in the training/monitoring data, to obtain global feature importance, and visualizes them in SageMaker Studio. To give insights to operation, data scientists select features with the highest SHAP values that contributed most to a positive fraud score (i.e., likely fraudster) for each individual example, and report them to the underwriter.

Explanation of a soccer scoring probability model. In the context of soccer matches the DFL (Deutsche Fußball Liga) applied Clarify to their xGoal model. The model predicts the likelihood that an attempted shot will result in a goal. They found that overall the features *AngleToGoal* and *DistanceToGoal* have the highest importance across all examples. Let us consider one of the most interesting games of the 2019/2020 Bundesliga season, where Borussia Dortmund beat Augsburg in a 5-3 thriller on January 18, 2020. Taking a look at the sixth goal of the game, scored by Jadon Sancho, the model predicted the goal with a high xGoal prediction of 0.18 compared to the average of prediction of 0.0871 across the past three seasons. What explains this high prediction? The SHAP analysis reveals that *PressureSum*, *DistanceToGoalClosest*, and *AmountofPlayers* contribute the most to this prediction. This is a noticeable deviation from the features that are globally the most important. What makes

this explanation convincing is the unusual nature of this goal: Sancho received a highly offensive through ball well into Augsburg’s half (past all defenders – hence low *PressureSum*), only to dart round the keeper and tap it into the goal near the posts (low *AngleToGoal*).

9 LESSONS LEARNED IN PRACTICE

Wide Applicability: We learned that wide applicability across datasets and model endpoints not only requires flexible parsing logic (see §6), but also creates a tension with ease of use: Our code tries to infer whether a file contains a header, or a feature is categorical or numerical. While it is convenient for most not having to specify all these details, for some customers our inference does not work well. This tension is further explored in §8.3.

Scalability: When working towards our goal of creating a scalable application, we learned that bias and explainability analyses have different resource demands (SHAP can be compute intensive, while bias metrics can be memory intensive), and hence require separate tuning of the parallelism. For explainability we had to work on small chunks, while for bias reports, we had to work on large chunks of the dataset. For details and other technical lessons, see §6.1.

Usability: In retrospect, our biggest lesson concerning usability is one we should have recognized prior to the usability studies – we cannot rely heavily on documentation. Instead we need to enable users to try out the feature without reading documentation. Early implementations of Clarify required the user to create their JSON configuration file manually. This was prone to both schema errors caused by users creating malformed JSON, and soft errors caused by users defining metrics using incorrect fields or labels. To avoid such errors we developed an explanatory UI. Now, we ask the customer to “Select the column to analyze for bias” and provide a drop down menu or programmatically construct the configuration for them. We are further revising the UI, see §8.3.

Best Practices: We recognize that the notions of bias and fairness are highly dependent on the application and that the choice of the attribute(s) for which to measure bias, as well as the choice of the bias metrics, can be guided by social, legal, and other non-technical considerations. Our experience suggests that the successful adoption of fairness-aware ML approaches in practice requires building consensus and achieving collaboration across key stakeholders (such as product, policy, legal, engineering, and AI/ML teams, as well as end users and communities). Further, fairness and explainability

related ethical considerations need to be taken into account during each stage of the ML lifecycle, for example, by asking questions stated in Figure 2.

10 CONCLUSION

Motivated by the need for bias detection and explainability from regulatory, business, and data science perspectives, we presented Amazon SageMaker Clarify, a new functionality that helps detect bias in data and models and helps explain predictions made by models. We described the design considerations, methodology, and technical architecture of Clarify, as well as how it is integrated across the ML lifecycle (data preparation, model evaluation, and monitoring post deployment). We discussed key technical challenges and resolutions, a case study, customer use cases, deployment results, and lessons learned in practice. Considering that Clarify is a scalable, cloud-based bias and explainability service designed to address the needs of customers from multiple industries, the insights and experience from our work are likely to be of interest for researchers and practitioners working on responsible AI systems.

While we have focused on providing tools for bias and explainability in the context of an ML lifecycle, it is important to realize that this is an ongoing effort that is part of a larger process. Another key challenge is to navigate the trade-offs in the context of an application between model accuracy, interpretability, various bias metrics, and related dimensions such as privacy and robustness (see, for example, [11, 23, 24, 32]). It cannot be overemphasized that developing AI solutions needs to be thought of more broadly as a process involving iterated inputs from and discussions with key stakeholders such as product, policy, legal, engineering, and AI/ML teams as well as end users and communities, and asking relevant questions during all stages of the ML lifecycle.

REFERENCES

- [1] Big data: Seizing opportunities, preserving values. https://obamawhitehouse.archives.gov/sites/default/files/docs/big_data_privacy_report_may_1_2014.pdf, 2014.
- [2] Big data: A report on algorithmic systems, opportunity, and civil rights. https://obamawhitehouse.archives.gov/sites/default/files/microsites/ostp/2016_0504_data_discrimination.pdf, 2016.
- [3] Big data: A tool for inclusion or exclusion? Understanding the issues (FTC report). <https://www.ftc.gov/reports/big-data-tool-inclusion-or-exclusion-understanding-issues-ftc-report>, 2016.
- [4] A. Agarwal, A. Beygelzimer, M. Dudik, J. Langford, and H. Wallach. A reductions approach to fair classification. In *ICML*, 2018.
- [5] S. Barocas, M. Hardt, and A. Narayanan. *Fairness and Machine Learning*. fairml-book.org, 2019.
- [6] S. Barocas, A. D. Selbst, and M. Raghavan. The hidden assumptions behind counterfactual explanations and principal reasons. In *FACCT*, 2020.
- [7] R. Berk, H. Heidari, S. Jabbari, M. Kearns, and A. Roth. Fairness in criminal justice risk assessments: The state of the art. *Sociological Methods & Research*, 2018.
- [8] U. Bhatt, A. Xiang, S. Sharma, A. Weller, A. Taly, Y. Jia, J. Ghosh, R. Puri, J. M. Moura, and P. Eckersley. Explainable machine learning in deployment. In *FACCT*, 2020.
- [9] E. Black, S. Yeom, and M. Fredrikson. FlipTest: Fairness testing via optimal transport. In *FACCT*, 2020.
- [10] T. Calders and S. Verwer. Three naive Bayes approaches for discrimination-free classification. *Data Mining and Knowledge Discovery*, 21(2), Sept. 2010.
- [11] A. Chouldechova. Fair prediction with disparate impact: A study of bias in recidivism prediction instruments. *Big data*, 5 2, 2017.
- [12] M. Donini, L. Oneto, S. Ben-David, J. Shawe-Taylor, and M. Pontil. Empirical risk minimization under fairness constraints. In *NeurIPS*, 2018.
- [13] D. Dua and C. Graff. UCI machine learning repository, 2017.
- [14] C. Dwork, M. Hardt, T. Pitassi, O. Reingold, and R. Zemel. Fairness through awareness. In *ITCS*, 2012.
- [15] B. Efron and R. J. Tibshirani. *An introduction to the bootstrap*. CRC press, 1994.
- [16] S. C. Geyik, S. Ambler, and K. Kenthapadi. Fairness-aware ranking in search & recommendation systems with application to LinkedIn talent search. In *KDD*, 2019.
- [17] B. Goodman and S. Flaxman. European union regulations on algorithmic decision-making and a “right to explanation”. *AI magazine*, 2017.
- [18] R. Guidotti, A. Monreale, S. Ruggieri, F. Turini, F. Giannotti, and D. Pedreschi. A survey of methods for explaining black box models. *ACM Computing Surveys*, 51(5), 2018.
- [19] M. Hardt, E. Price, E. Price, and N. Srebro. Equality of opportunity in supervised learning. In *NeurIPS*, 2016.
- [20] K. Holstein, J. Wortman Vaughan, H. Daumé III, M. Dudik, and H. Wallach. Improving fairness in machine learning systems: What do industry practitioners need? In *CHI*, 2019.
- [21] D. Janzing, L. Minorics, and P. Bloebaum. Feature relevance quantification in explainable AI: A causal problem. In *AISTATS*, 2020.
- [22] N. Kilbertus, M. Rojas Carulla, G. Parascandolo, M. Hardt, D. Janzing, and B. Schölkopf. Avoiding discrimination through causal reasoning. In *NeurIPS*, 2017.
- [23] J. Kleinberg and S. Mullainathan. Simplicity creates inequity: Implications for fairness, stereotypes, and interpretability. In *EC*, 2019.
- [24] J. Kleinberg, S. Mullainathan, and M. Raghavan. Inherent trade-offs in the fair determination of risk scores. In *ITCS*, 2017.
- [25] M. J. Kusner, J. Loftus, C. Russell, and R. Silva. Counterfactual fairness. In *NeurIPS*, 2017.
- [26] H. Lakkaraju, E. Kamar, R. Caruana, and J. Leskovec. Faithful and customizable explanations of black box models. In *AIES*, 2019.
- [27] E. Liberty, Z. Karnin, B. Xiang, L. Rouesnel, B. Coskun, R. Nallapati, J. Delgado, A. Sadoughi, Y. Astashonok, P. Das, C. Balioglu, S. Chakravarty, M. Jha, P. Gautier, D. Arpin, T. Januschowski, V. Flunkert, Y. Wang, J. Gasthaus, L. Stella, S. Rangapuram, D. Salinas, S. Schelter, and A. Smola. Elastic machine learning algorithms in Amazon SageMaker. In *SIGMOD*, 2020.
- [28] Z. C. Lipton. The myths of model interpretability. *ACM Queue*, 16(3), 2018.
- [29] S. M. Lundberg and S.-I. Lee. A unified approach to interpreting model predictions. In *NeurIPS*, 2017.
- [30] S. M. Lundberg, B. Nair, M. S. Vavilala, M. Horibe, M. J. Eisses, T. Adams, D. E. Liston, D. K.-W. Low, S.-F. Newman, J. Kim, and S.-I. Lee. Explainable machine-learning predictions for the prevention of hypoxaemia during surgery. *Nature Biomedical Engineering*, 2(10), Oct. 2018.
- [31] N. Mehrabi, F. Morstatter, N. Saxena, K. Lerman, and A. Galstyan. A survey on bias and fairness in machine learning. *arXiv preprint arXiv:1908.09635*, 2019.
- [32] S. Milli, L. Schmidt, A. D. Dragan, and M. Hardt. Model reconstruction from model explanations. In *FACCT*, 2019.
- [33] 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 *FACCT*, 2019.
- [34] J. Pearl. Causal inference in statistics: An overview. *Statistics Surveys*, 3, 2009.
- [35] V. Perrone, M. Donini, M. B. Zafar, R. Schmucker, K. Kenthapadi, and C. Archambeau. Fair Bayesian optimization. In *AIES*, 2021.
- [36] M. T. Ribeiro, F. Singh, and C. Guestrin. “Why should I trust you?” Explaining the predictions of any classifier. In *KDD*, 2016.
- [37] C. Rudin. Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*, 1(5), May 2019.
- [38] L. S. Shapley. A value for n-person games. *Contributions to the Theory of Games*, 2(28), 1953.
- [39] S. Sharifi-Malvajerdi, M. Kearns, and A. Roth. Average individual fairness: Algorithms, generalization and experiments. In *NeurIPS*, 2019.
- [40] K. Simonyan, A. Vedaldi, and A. Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. In *ICLR Workshop*, 2014.
- [41] A. Singh and T. Joachims. Fairness of exposure in rankings. In *KDD*, 2018.
- [42] D. Slack, N. Rauschmayr, and K. Kenthapadi. Defuse: Harnessing unrestricted adversarial examples for debugging models beyond test accuracy. *arXiv preprint arXiv:2102.06162*, 2021.
- [43] M. Sundararajan, A. Taly, and Q. Yan. Axiomatic attribution for deep networks. In *ICML*, 2017.
- [44] S. Vasudevan and K. Kenthapadi. LiFT: A scalable framework for measuring fairness in ML applications. In *CIKM*, 2020.
- [45] S. Wachter, B. Mittelstadt, and C. Russell. Why fairness cannot be automated: Bridging the gap between EU non-discrimination law and AI. SSRN Scholarly Paper ID 3547922, Social Science Research Network, Mar. 2020.
- [46] J. Wexler, M. Pushkarna, T. Bolukbasi, M. Wattenberg, F. Viégas, and J. Wilson. The What-If Tool: Interactive probing of machine learning models. *IEEE Transactions on Visualization and Computer Graphics*, 26(1), 2020.
- [47] M. B. Zafar, I. Valera, M. Gomez Rodriguez, and K. P. Gummadi. Fairness beyond disparate treatment & disparate impact: Learning classification without disparate mistreatment. In *WWW*, 2017.
- [48] B. H. Zhang, B. Lemoine, and M. Mitchell. Mitigating unwanted biases with adversarial learning. In *AIES*, 2018.

ACKNOWLEDGMENTS

We would like to thank our colleagues at AWS for their contributions to Clarify. We thank Edi Razum, Luuk Figdor, Hasan Poonawala, the DFL and Zopa for contributions to the customer use case, and Cédric Archambeau and the anonymous reviewers for their valuable feedback.

A DETAILED CONFIGURATION

A.1 Additional Bias Metrics

Recall that we defined the fraction of positively labeled examples in group a as $q_a = n_a^{(1)} / n_a$ ($q_d = n_d^{(1)} / n_d$ for group d , respectively). They define a probability distribution P_a for group a over labels as q_a for positive labels and $1 - q_a$ for negative labels and P_d as q_d and $1 - q_d$. Differences in these probability distributions yield additional pre-training bias metrics:

KL KL divergence measures how much information is needed to move from one probability distribution P_a to another P_d : $KL = q_a \log \frac{q_a}{q_d} + (1 - q_a) \log \frac{1 - q_a}{1 - q_d}$.

JS Jensen-Shannon divergence provides a symmetric form if KL divergence. Denoting by P the average of the label distributions of the two groups, we define

$$JS(P_a, P_d) = \frac{1}{2} [KL(P_a, P) + KL(P_d, P)].$$

LP We can also consider the norm between the probability distributions of the labels for the two groups. For $p \geq 1$, we have

$$L_p(P_a, P_d) = \left[\sum_y |P_a(y) - P_d(y)|^p \right]^{1/p}.$$

TVD The total variation distance is half the L_1 distance: $TVD = \frac{1}{2} L_1(P_a, P_d)$.

KS Kolmogorov-Smirnov considers the maximum difference across label values: $KS = \max(|P_a - P_d|)$.

An additional post-training bias metric is based on a flip test.

FT Counterfactual Fliptest: *FT* assesses whether similar examples across the groups receive similar predictions. For each example in d , we count how often its k -nearest neighbors in a receive a different prediction. Simplifying [9], we define $FT = (F^+ - F^-) / n_d$ where F^+ (F^- resp.) is the number of examples in d with a negative (positive, resp.) prediction whose nearest neighbors in a have a positive (negative, resp.) prediction.

A.2 Configuration Options

Table 5 illustrates the flexible support for model outputs.

Example output	Configuration of predictor
C3,[0.1,0.3,0.6]"	"label": 0, "probability": 1
"[0.1, 0.3, 0.6]"	"probability": 0, "label_headers": ["C1", "C2", "C3"]
{"pred": C3, "sc": [0.1, 0.3, 0.6]}	"label": "pred", "probability": "sc", "content_type": "application/jsonlines"

Table 5: Illustration of multi-class model outputs supported by Clarify. The model predicts scores for classes C1, C2, C3.

A.3 Detailed Example

Below we show an example analysis config. It is minimal and additional configuration options (e.g. to support datasets without headers or model endpoints using JSON format) exist.

```

1  "dataset_type": "text/csv",
2  "label": "Class1Good2Bad",
3  "label_values_or_threshold": [1],
4  "facet": [
5    {
6      "name_or_index": "ForeignWorker",
7      "value_or_threshold": [1]
8    }
9  ],
10 "group_variable": "A151",
11 "methods": {
12   "shap": {
13     "baseline": "s3://<bucket>/<prefix>/<file>",
14     "num_samples": 3000,
15     "agg_method": "mean_abs"
16   },
17   "pre_training_bias": {
18     "methods": "all"
19   },
20   "post_training_bias": {
21     "methods": "all"
22   },
23   "report": {
24     "name": "report",
25     "title": "Clarify Analysis Report"
26   }
27 },
28 "predictor": {
29   "model_name": "german-xgb",
30   "instance_type": "ml.c5.xlarge",
31   "initial_instance_count": 1
32 }
33
34
```