

Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) Datasheet

I. MOTIVATION

I-A For what purpose was the dataset created?

The Correctional Offender Management Profiling for Alternative Sanctions (COMPAS) is a predictive tool used by judges and parole officers. The tool produces an automated risk score to predict the probability of re-offending within a specific time frame. In 2016, ProPublica released a study [1] and an accompanying dataset obtained from Broward County, Florida. The aim of the study was to investigate whether there was any racial bias in the COMPAS tool.

I-B Who created the dataset?

Is it an official law enforcement or government body? An academic research team? Other?

This dataset was produced by Jeff Larson, Surya Mattu, Lauren Kirchner and Julia Angwin for ProPublica.

I-C Was there a specific task in mind, or gap that needed to be filled?

Previous studies have investigated the efficacy of U.S. risk assessment algorithms, including COMPAS [2]. However, there were no recent datasets published containing COMPAS scores and associated re-offending data.

II. COMPOSITION

II-A What do the instances that comprise the dataset represent?

For example: crimes, offenders, court cases, police officers

Each instance of the COMPAS dataset corresponds to an individual that has been assessed by the COMPAS system.

II-B Are there multiple types of instances?

For example: offenders, victims, and the relationship between them.

No.

II-C How many instances are there in total?

Of each type, if appropriate.

The dataset contains 11,757 individuals who were assigned a risk score by the COMPAS tool during pre-trial.

II-D Does the dataset contain all possible instances or is it a sample (not necessarily random) of instances from a larger set?

For example, if it is traffic stops from a territory, is it all traffic stops conducted within that territory within a specific time? If not, is it a representative sample of all stops? Describe how representativeness was validated/verified. If it is not representative, please describe why.

The dataset contains all individuals who were screened by the COMPAS tool in Broward County between 2013 – 2014.

II-E What data does each instance consist of?

If there is a large number of variables, please provide a broad description of what is included.

Each instance consists of the following variables:

- Offender age
- Offender age at first offense
- Race of offender
- Gender of offender
- Jail history
- Prison history
- Charge history, including charge type, charge degree, etc.
- COMPAS score
- Recidivist

II-F Is there a target label or associated with each instance?

Please include labels that are likely to be used as target labels, e.g. recidivism.

There are two target labels: ‘COMPAS score’ and ‘recidivist’. COMPAS score can also further be grouped into Low, Medium, and High. Scores 1 to 4 were labeled by COMPAS as “Low”; 5 to 7 were labeled “Medium”; and 8 to 10 were labeled “High.”

II-G Are there recommended data splits (e.g., training, development/validation, testing)?

If so, please provide a description of these splits, explaining the rationale behind them.

No.

II-H Does the dataset contain data on race and ethnicity? If so, is it based on the individual’s self-description, or based on officer’s impression? Was it collected or derived in post-processing? For example, by name analysis.

Information on race is included. It is unclear if it is based on self-description or not.

II-I Are there any known errors, sources of noise, bias or missing data, or variables collected for only part of the datasets?

If so, please provide a description.

The dataset creator note that “We found that sometimes people’s names or dates of birth were incorrectly entered in some records”.

II-J Does the dataset contain data on criminal history or other data that might be considered confidential or sensitive in any way?

For example: sexual orientations, religious beliefs, political opinions or union memberships, or locations; financial or health data; biometric or genetic data; forms of government identification, such as social security numbers; If so, please provide a description.

The dataset contains information criminal history, prison history, and jail history as well as demographic information.

II-K Is it possible to identify individuals (i.e., one or more natural persons), either directly or indirectly (i.e., in combination with other data) from the dataset?

If so, please describe how.

Yes, the individuals are named in the dataset.

III. USES

III-A What type of tasks, if any, has the dataset been used for?

If so, please provide examples and include citations.

The dataset has been used for:

- Investigating whether there are any racial biases in the COMPAS algorithm [3], [4].
- Evaluating the performance of “fair” algorithms, i.e. balancing predictive performance along with a defined fairness criteria [5].

III-B Is there a repository that links to any or all papers or systems that use the dataset?

If so, please provide a link or other access point.

There is no specific repository. General academic search engines, such as Google Scholar, can be used with search terms such as: “COMPAS risk assessment”.

III-C What (other) tasks could the dataset be used for?

For example: testing predictive policing systems, predicting recidivism.

This dataset was created for a specific purpose, but was adopted by the algorithmic fairness community as a benchmark. This practice has been criticized due to lack of domain context [6].

III-D Is there anything about the composition of the dataset or the way it was collected and pre-processed/cleaned/labeled that might impact future uses?

For example, is there anything that a dataset consumer might need to know to avoid uses that could result in unfair treatment of individuals or groups (e.g., stereotyping, quality of service issues) or other risks or harms (e.g., legal risks, financial harms)? If so, please provide a description. Is there anything a dataset consumer could do to mitigate these risks or harms?

There are a few important things to keep in mind when using this dataset:

- **Potential Incorrect Data Entry:** The ProPublica authors note: “We found that sometimes people’s names or dates of birth were incorrectly entered in some records”.
- **Conclusion the dataset contains racial bias is disputed:** There are a number of critics of the paper which this dataset is from. Criticisms can be categorized into (1) not using the same set of variables Northpointe used to compute the COMPAS score, (2) incorrect modelling assumptions [3].
- **Only uses local charges:** The dataset uses charges in the Broward County database, which only contains the local charges. If criminal history exists outside this, it is not captured. Additionally, if re-offending occurs outside the county, it will not be counted as recidivism.
- **Recidivism timer starts at screening:** “we defined recidivism as a new arrest within two years” (of screening). If an offender is taken into custody, they will have less opportunity to re-offend, biasing the results.

IV. COLLECTION PROCESS

IV-A How was the data associated with each instance acquired?

e.g. the data collected survey, the raw data is routinely collected by the courts.

The dataset is linked from four data sources, matched on first and last name:

- 1) A public records request of COMPAS scores from Broward County Sheriff’s Office in Florida.
- 2) Charge history from the Broward County Clerk’s Office website.
- 3) Jail records from the Broward County Sheriff’s Office.
- 4) Public incarceration records from the Florida Department of Corrections website.

IV-B Was the information self-reported?

If the data was self-reported, was the data validated/verified? If so, please describe how.

The data is not self-reported.

IV-C Who was involved in the data collection process?

Was this done as part of their other duties? If not, were they compensated?

All data was received from the Broward County Sheriff’s Office. The raw data was collected as part of routine law enforcement work.

IV-D Over what timeframe was the data collected? Does this timeframe match the creation timeframe of the data associated with the instances (e.g., recent crawl of old news articles)?

If not, please describe the timeframe in which the data associated with the instances was created. If the collection was not continuous within the timeframe, please specify the intervals, for example, annually, every 4 years, irregularly.

The dataset was created in 2016 by ProPublica, and concerns the 2013–2014 time period.

IV-E Were any ethical review processes conducted (e.g., by an institutional review board)?

If so, please provide a description of these review processes, including the outcomes, as well as a link or other access point to any supporting documentation.

Unknown.

IV-F Were the individuals in question notified about the data collection? Did they give their consent?

If consent was obtained, were the consenting individuals provided with a mechanism to revoke their consent in the future or for certain uses?

No.

IV-G Has an analysis of the potential impact of the dataset and its use on data subjects (e.g., a data protection impact analysis) been conducted?

If so, please provide a description of this analysis, including the outcomes, as well as a link or other access point to any supporting documentation.

Unknown.

V. PRE-PROCESSING, CLEANING, LABELING

V-A Was any preprocessing/cleaning/labeling of the data done (e.g., discretization or bucketing, removal of instances, processing of missing values)?

If so, please provide a description and reference to the documentation. If not, you may skip the remaining questions in this section.

The dataset only offenders who were screened in pre-trial. This reduced the number of individuals from 18,610 to 11,757.

V-B Was the “raw” data saved in addition to the preprocessed/cleaned/labeled data?

If so, please provide a link or other access point to the “raw” data.

Yes. This can be found here:

<https://github.com/propublica/compas-analysis>

V-C Is the software that was used to preprocess/clean/label the data available?

If so, please provide a link or other access point.

Yes. This can be found here:

<https://github.com/propublica/compas-analysis>

V-D Is the data publicly available? How and where can it be accessed (e.g., website, GitHub)?

Does the dataset have a digital object identifier (DOI)?

The dataset is available on GitHub:

<https://github.com/propublica/compas-analysis>

V-E When will the dataset be distributed?

The dataset has been available since 2016.

V-F Is the dataset be distributed under a copyright or other intellectual property (IP) license, and/or under applicable terms of use (ToU)?

If so, please describe this license and/or ToU, and provide a link or other access point to, or otherwise reproduce, any relevant licensing terms or ToU, as well as any fees associated with these restrictions.

A license is not specified.

VI. MAINTENANCE

VI-A Is the dataset maintained? Who is supporting/hosting/maintaining the dataset?

The dataset is no longer maintained.

VI-B How can the owner/curator/manager of the dataset be contacted (e.g., email address)?

ProPublica can be contacted at: hello@propublica.org

VI-C Will the dataset be updated (e.g., to correct labeling errors, add new instances, delete instances)?

No, the dataset hasn't been updated since 2017.

VI-D Are older versions of the dataset continue to be supported/hosted/maintained?

The dataset is no longer updated. However, older versions of the dataset will continue to be accessible via GitHub if any updates occur.

VI-E If others want to extend/augment/build on/contribute to the dataset, is there a mechanism for them to do so?

If so, please provide a description.

Yes, via a GitHub pull request, given the authors are still active.

REFERENCES

- [1] Julia Angwin, Jeff Larson, Surya Mattu, and Lauren Kirchner. Machine Bias. 2016.
- [2] Tim Brennan, William Dieterich, and Beate Ehret. Evaluating the Predictive Validity of the COMPAS Risk and Needs Assessment System. *Criminal Justice and Behavior*, 36(1):21–40, 2009.
- [3] Cynthia Rudin, Caroline Wang, and Beau Coker. The Age of Secrecy and Unfairness in Recidivism Prediction. *Harvard Data Science Review*, 2(1), mar 31 2020. <https://hdr.mitpress.mit.edu/pub/7z10o269>.
- [4] Melissa Hamilton. The Biased Algorithm: Evidence of Disparate Impact on Hispanics. *Am. Crim. L. Rev.*, 56:1553, 2019.
- [5] Christina Wadsworth, Francesca Vera, and Chris Piech. Achieving Fairness Through Adversarial Learning: an Application to Recidivism Prediction. *arXiv preprint arXiv:1807.00199*, 2018.
- [6] Michelle Bao, Angela Zhou, Samantha Zottola, Brian Brubach, Sarah Desmarais, Aaron Horowitz, Kristian Lum, and Suresh Venkatasubramanian. It's COMPASlicated: The has Messy Relationship between RAI Datasets and Algorithmic Fairness Benchmarks. *arXiv preprint arXiv:2106.05498*, 2021.