**Assessing the quality of studies using machine learning for personality assessment: A
systematic review**

Joanne Hinds, Thomas Parkhouse & Victoria Hotchin

School of Management (University of Bath, UK)
Information, Decisions and Operations Division, School of Management, East Building
Room 3.1, University of Bath, Claverton Down, Bath, BA2 7AY.

Corresponding Author: Joanne Hinds

Email: J.Hinds@bath.ac.uk

Abstract

In recent years, the use of machine learning to predict personality from digital data has gained increasing interest from organisations, academics and the public. In turn, a new field of personality computing has developed, which involves combining machine learning techniques with psychological measures of personality. However, effectively integrating these approaches is challenging - the fields of machine learning and psychology are highly disparate, with different objectives, methodologies, and perspectives on performing and reporting research. In this article, we report findings from a systematic review that analysed 178 personality computing studies published before November 2020. We developed a novel set of criteria that was used to evaluate the quality of study design and reporting of each study according to 10 criteria: hypotheses, study rationale, selection of features, algorithm training, ground truth, sampling, the evaluation of algorithms' performance (i.e., classification, regression), the performance measures reported, and detail concerning ethics and open science practices. Our findings highlight that a large proportion of studies lack detail on the above criteria, which leads to questions over the validity, reliability, and replicability of the findings. We discuss the implications of this research for practice and recommend directions for future work.

Keywords: personality, machine learning, digital data, personality computing, Big Five.

In recent years, the possibility of predicting personality from digital data has received widespread public and academic attention. The notion that digital data can be used to profile our personalities and microtarget information, products, and services has raised much concern, from reports of Spotify's tracking technology that can supposedly influence people's mood or physiological responses (Music Business Worldwide, 2020), to Cambridge Analytica's alleged attempts to influence votes in the US 2016 Presidential election (Cadwalladr, Graham-Harrison, 2018; Hinds, Williams, & Joinson, 2020). At the same time academic research has proliferated, attracting efforts across a diverse range of disciplines including psychology, artificial intelligence, and marketing. In particular, machine learning has become increasingly popular as a means of analysing digital data in fine granularities and at mass scale. As such, research in the burgeoning field of personality computing has demonstrated that personality can be detected from a wide variety of technologies and interactions including social media, smartphones and wearables, and the internet of things (e.g., Azucar, Marengo, & Settanni, 2018a; Hinds & Joinson, 2020; Marengo & Montag, 2020).

Personality computing requires psychology and computer science knowledge; for instance, assessing personality effectively involves knowledge of psychometrics. Likewise, training algorithms requires knowledge of feature engineering and hyperparameter tuning to ensure that the data used have some ability to predict personality (e.g., Stachl, Pargent, et al., 2020b). However, interdisciplinary collaboration appears to be rare and research practices seem to follow the traditional conventions of their respective fields (i.e., machine learning studies tend to be reported in computer science outlets with little consideration to psychology theory or practice). This means that research often focuses on certain theoretical, methodological, or practical considerations that are important for one field, but that are not prioritised in another (e.g., the tendency for psychologists to focus on explaining theories, in comparison to the heavy focus on data driven predictions in machine learning, (Hinds & Joinson, 2019; Yarkoni &

Westfall, 2017). As a result, the validity of the findings may be questionable if certain psychological methodological considerations have been overlooked in favour of the detail expected from machine learning publications (and vice versa). Further, these potential issues can be difficult to ascertain due to the conventions and expectations when reporting studies in different disciplines.

Recently, concerns over the use of machine learning have been communicated in medical research, where the reporting of methods were found to be inaccurate, misleading, or did not follow guidelines (e.g., Christodoulou et al., 2019; DeMasi, Kording, & Recht, 2017; Faes et al., 2020; Yusuf et al., 2020). Similarly, the replication crisis has also highlighted stark issues relating to the selective analysis, design, and reporting across many studies in psychological science (e.g., Maxwell, Lau, & Howard, 2015; Open Science Collaboration, 2015). Further, these problems have plagued these fields despite numerous best practice guidelines detailing how research should be conducted and reported such as the TRIPOD in medicine (Moons et al., 2015) or the Journal Article Reporting Standards (JARS) in psychology (APA, 2008). In contrast, no such guidelines exist for machine learning and personality prediction. Given that such work involves bringing together different disciplines with distinct conventions in terms of study design, methods, reporting and publishing, it seems highly likely that similar issues may pervade this field. Thus, understanding whether existing evidence is reliable and likely to replicate is of critical importance, particularly at a time when there is heightened interest in the potential applications of personality predictions.

In this article, we report findings from a systematic review that assesses the quality of design, implementation and reporting of personality computing studies predict individuals' Big Five traits using digital data. By drawing upon reporting guidelines in medicine and social science, we devise a novel set of criteria that we use to evaluate existing studies. Specifically, we focus on the key components of the research process in terms of study design (i.e.,

hypotheses, study rationale, selection of features, algorithm training, ground truth, sampling, the evaluation of algorithms' performance (i.e., classification, regression), the performance measures reported, and the reporting of ethical and open science practices. Below, we provide an overview of the main processes and challenges underpinning each of these elements.

Overview of personality computing

In order to devise a set of criteria from which to assess existing studies, we reviewed the key processes involved in conducting and reporting personality computing research. Given that the main focus of these studies is to predict personality (i.e., a psychological construct), we adopted a psychological perspective in investigating these processes, that is we prioritised elements that would be necessary to conduct and report a high-quality study of personality over technical machine learning detail. We outline the expectations and challenges associated with reporting studies effectively over the remainder of this section. A detailed breakdown of the criteria is provided in Supplemental File 1.

Supervised learning is a commonly used approach in personality computing, which involves “training” an algorithm to predict individuals' personalities from a set of data (or variables) known as “features”, for example inferring personality from financial transactions (Gladstone, Matz, & Lemaire, 2019), music listening patterns (Anderson et al., 2020), or smartphone data (Stachl, Au, et al., 2020). The data are partitioned into a “training” set and a “testing” set – the algorithm learns to predict personality from the training set, and its performance is evaluated using the testing set (e.g., Bleidorn and Hopwood, 2019; Stachl et al. 2020). Typically, such approaches are data driven – where researchers analyse data without any preconceptions about their relationship with personality traits. These contrast with traditional psychological studies of personality, which instead adopt theory-driven approaches – where researchers test hypotheses with variables selected in accordance with their chosen theory/ies. Although data driven approaches in machine learning enable researchers to detect

unexpected relationships or patterns within the data, they often neglect theory that may be informative when selecting certain types of data and when explaining the findings (that could be derived through hypothesis testing, e.g., Mahmoodi, Leckelt, van Zalk, Geukes, & Back, 2017). Ideally, a well-designed study would capitalise on both approaches – as it would progress existing evidence and also allow the discovery of new patterns within the data. We therefore investigated whether hypotheses or a study rationale (derived from prior evidence) were reported within the studies, and whether these were used to inform the selection of features used to predict personality:

Criterion 1: Hypotheses – clearly stated, directional, and based on theory

Criterion 2: Study rationale – clearly stated and supported by evidence

Criterion 3: Selection of features – clearly stated and derived from hypotheses or prior evidence

After selecting the data, researchers should employ an approach that involves training, validating, and testing the model in order to establish how effectively it can predict personality. Ideally, a well-executed approach should apply cross validation, where the model is trained on multiple subsamples and where one is reserved for testing. The overall accuracy of the model is then determined by averaging the trials (Kosinski, Wang, Lakkaraju, & Leskovec, 2016). Cross validation can help to avoid overfitting and creating an over-optimistic assessment of performance (Subramanian & Simon, 2013). Thus, reporting should include sufficient detail on the process of training and testing, such that readers can evaluate whether steps have been taken to avoid poorly performing models, and that such methods could be replicated.

Criterion 4: Algorithm training – clear detail on the approach taken in terms of how models were trained and tested

High quality training also requires high quality data in order to ensure that the predictions are valid, otherwise researchers risk creating flawed algorithms that make

meaningless or weak predictions (i.e., the classic “garbage in, garbage out” aphorism in computer science, Sanders and Saxe, 2017; Stuart Geiger et al., 2020). In order to create high quality training data, researchers should create a dataset of ground truth that comprises a valid measure of the individuals’ personalities. At present, self-reported or observer rated questionnaires remain the gold standard approach for personality assessment, and researchers should therefore select a scale that has high reliability (i.e., test-retest reliability and internal consistency) such as the BFI (e.g., John & Srivastava, 1999) or an IPIP scale (see: <https://ipip.ori.org>).

In addition to the issues relating to the selection and curation of ground truth, an algorithm’s performance can also be affected by other systematic errors that are introduced during the training process. For instance, selection bias can occur if the sample used to train the algorithm is not based on demographics that are representative of the population (Liu & Ruths, 2013; Mohammady & Culotta, 2014) or if it is heavily skewed towards a high / low end of a trait. Further, the prominence of machine learning workshops and competitions at conferences can entail the re-analysis of the existing datasets (e.g., PAN at the Conference and Labs of the Evaluation Forum [CLEF]), which can perpetuate these issues if biased samples are repeatedly used for training/testing models. As such, we examined the approaches taken when selecting/curating ground truth and sampling:

Criterion 5: Ground truth – valid, reliable method reported

Criterion 6: Sample – detailed information provided (e.g., age, gender, education)

Another consideration relates to how an algorithm’s performance is evaluated, i.e., whether a classification or regression approach is employed. Approaches that use regression are valuable from a psychological perspective, as traits are treated as continuous values, thus predictions are compared directly against individuals’ personality scores. In contrast, classification involves categorising individuals into classes, for instance, whether they are high

or low in extraversion (e.g. Mønsted, Mollgaard, and Mathiesen, 2018; Sourì, Hosseinpour, and Rahmani, 2018). Classification can therefore be problematic because personality traits are normally distributed in the population, thus most people will score around the middle and will not be very high or low for any given trait (Schmitt et al., 2007; Stachl, Pargent, et al., 2020a). Therefore, claims of a high performing algorithm can be misleading because the simplified categories do not accurately reflect the underlying traits being predicted. This can also be exacerbated if weak ground truth has been applied in the training process. We therefore evaluated the approach taken, that is, whether regression or classification (or both) were performed in the study:

Criterion 7: Regression / Classification – clear detail on the approach/es adopted

Relatedly, researchers may choose to report a variety of different metrics when evaluating an algorithm's performance. Classification models are typically evaluated using measures of overall accuracy (i.e., the percentage of correct classifications) or measures such as recall, precision, or the F1. Alternatively, regression models tend to be evaluated with metrics such as the coefficient of determination, R^2 , mean square error (MSE), mean absolute error (MAE), root mean square error (RMSE) and Pearson's correlation coefficient. Each metric has its own merits and limitations (see, Botchkarev, 2019; Flach, 2019; Seliya, Khoshgoftaar, & Van Hulse, 2009); for instance, accuracy does not account for false positive and false negative rates, as indicated by precision and recall. Therefore, studies should ideally report multiple metrics in order to provide a balanced evaluation of the model's performance (e.g., Stachl, Pargent, et al., 2020a):

Criterion 8: Analysis – multiple measures of performance reported

Finally, the issues relating to the replicability of studies in recent years have increased pressure for researchers to report their practices transparently and to make their data open access (Open Science Collaboration, 2012, 2015). This is particularly pertinent when collecting

and analysing data that may be public on social media or that can be accessed unobtrusively via individuals' online activities and use of digital devices. Although open science practices are relatively new, ideally studies should be striving to be transparent in their reporting practices concerning the ethical treatment of participants and data as well as open science practices:

Criterion 9: Ethical approval - Clear detail on ethical approval reported

Criterion 10: Open science / data availability - Detail on open science practices and / or access to the datasets used

In sum, we developed the above criteria intended to assess issues surrounding the interdisciplinary design and reporting of personality computing studies. In the following section, we describe our systematic approach for collecting and synthesising relevant studies.

Method

Search Procedure

Our search procedure and reporting of study findings followed the Preferred Reporting of Items for Systematic Reviews and Meta-Analysis (PRISMA) guidelines (www.prisma-statement.org) (Moher, Liberati, Tetzlaff, & Altman, 2009). To identify relevant articles, we performed searches in the Web of Science, IEEE, and ACM online libraries for all articles that were published or available online as of 1st November 2020. This built upon previous a previous systematic review protocol reported in Hinds and Joinson (2020). Because this area of research has proliferated across multiple disciplines and sub-topics (each with different styles and reporting conventions), we used a broad variety of search terms in attempting to locate articles. The search terms that were applied across each of the databases are as follows: Web of Science – search by topic : (predict or identify or detect* or Facebook or Twitter or Instagram or YouTube) and (person* or individual difference) and (digital or internet or online or computer-

mediated) and (social* or big 5 or dark triad or priva* or web* or mobile* or sms or shop* or date or dating or query or language or achievement or affiliation or self-esteem or intelligen* or IQ or leader* or self-affirm* or influen* or big data or extraversion or neuroticism or openness or conscientiousness or agreeableness or narcissis* or psychopath* or sociopath* or decept* or deceiv*); IEEE - predict* personality (search with metadata only); ACM - search where title matches all 'predict personality' and where abstract matches all 'predict personality'.

The search generated 26,426 articles in total. After duplicates were removed, the articles were assessed for potential eligibility by screening their titles and abstracts. Then we closely scrutinised abstracts and detailed information in their methods, procedures etc., removing all articles that did not meet the inclusion criteria. The reference lists of all remaining articles were examined, and any further articles that met the criteria were retrieved. This step was performed iteratively on each article added, until no further articles were retrieved. Additionally, we applied this procedure to all relevant references reported in prior meta-analyses and reviews (specifically, Azucar et al., 2018; Connelly & Ones, 2010; Kedar & Bormane, 2016; Marengo & Montag, 2020; Settanni et al., 2018; Tskhay & Rule, 2014; Vinciarelli & Mohammadi, 2014). In total, this procedure yielded 178 articles and 201 studies (included within those articles) that met our inclusion criteria (see Supplemental File 2 for the reference list of all articles included in the review).

Eligibility criteria

Studies were included in the review if they met the following criteria: (a) machine learning was used to predict the Big Five personality traits (or a subset thereof); (b) machine learning models were trained and tested on some form of digital footprint; (c) detail regarding an algorithm's performance was reported. Studies published anytime up until November 2020 were included and we restricted inclusion to studies reported in English. No restrictions were

applied in terms of publication status or outlet, that is studies that met the criteria but were not peer reviewed, or that were unpublished, as well as other forms of grey literature, were included.

Figure 1

PRISMA Flow chart summarising database search and article screening.

Coding of study characteristics

Many of the studies included in the review were derived from a previous systematic review performed by Hinds and Joinson (2020). As part of this work, a comprehensive codebook was developed that included detail on study characteristics, which included: reference information (title, authors, publication year), sample sizes, platform/device (e.g.,

Facebook, Twitter, YouTube, smartphone), type of digital footprint studied (e.g., status updates, number of friends, number of photos uploaded), type of personality questionnaire used, effect size and statistical information relating to the predictive performance of computer algorithms (e.g., Mean Absolute Error, Precision, Recall, F-Score), and type of machine learning algorithm used for prediction (e.g., Support Vector Machines, Random Forests, Neural networks). All studies were coded by author JH, with 50% coded by a second coder. Cohen's k showed high levels of consensus, $k = 0.97$ (95% CI, .94 to 1.00) – see Hinds and Joinson (2020) for further detail.

Development of a quality checklist

Assessing study quality is a fundamental part of many systematic reviews. In fields such as medicine, there are a number of well-established resources to aid researchers in evaluating the design, implementation, and reporting of studies, such as the Cochrane Handbook, the Newcastle-Ottawa scale, and the Jadad scale. However, there is little guidance on what constitutes effective or appropriate assessment of study quality in articles that use machine learning, an issue that has been reflected in a number of articles that have raised awareness of these issues in medical-based machine learning in recent years (e.g., see Moons et al., 2015; Yusuf et al., 2020). This issue is further exacerbated in fields outside of medicine (and indeed personality computing), for which there is no guidance whatsoever (at least to our knowledge). Thus, in the absence of such guidelines, we developed a set of quality criteria, which were used to evaluate studies.

To do this, we consulted a number of existing checklists and academic articles in order to inform the development of our criteria. These included guidelines provided by The Campbell Collaboration research network (<https://campbellcollaboration.org>), the Mixed-Methods Appraisal Tool (MMAT) (Hong et al., 2018), the TRIPOD (Transparent Reporting of a Multivariate Prediction Model for Individual Prognosis or Diagnosis) guidelines (Moons et al., 2015) and other

guidelines/recommendations reported in medical research (specifically Luo et al., 2016 and Yusuf et al., 2020) and personality computing (Stachl, Pargent, et al., 2020a).

Taking guidance from these sources, we developed a novel quality assessment for personality computing studies. The criteria included the evaluation of research hypotheses, study rationale, selection of features, algorithm training, ground truth, sampling, the evaluation of algorithms' performance (i.e., classification, regression), the performance measures reported, and detail concerning ethics and open science practices. Each criterion was evaluated against a score, ranging from 1 to 5 (depending on the range of scores attributed to each criterion). The categories were not always hierarchical, for example, the following scores were used to evaluate whether or not a study reported detail on ethical approval: 1 = ethical approval not stated, 2 = ethical approval stated. The full detail outlining the assessment criteria are provided in Supplemental File 1.

Performing the quality assessment

Three researchers (TP, VH, and JH) independently applied the quality assessment criteria to a sample of 10 articles. Each of the researcher's scores were compared and discrepancies were discussed, such that final agreement on the scores were reached on all articles. The quality criteria were adjusted in line with any issues that arose through these discussions. The remaining articles were then divided between two reviewers (TP and VH). Any issues that arose were highlighted and resolved through discussion, before agreeing upon a final rating.

Findings

Our search generated a total of 178 articles. Table 1 summarises the criteria and details the percentage of articles in each rating category. We present findings across each of the 10 criteria over the following sections.

[INSERT TABLE 1 ABOUT HERE]

Criteria 1 and 2: Hypotheses and study rationale

The majority of articles (58%) reported research questions or aims rather than specific hypotheses (15%), however a sizeable minority (27%) reported neither.

Study rationales were reported with varying degrees of clarity; whilst the majority of articles (31%) presented clear reasoning supported by theory and/or evidence, 29% were classed as relatively weak and 25% were unclear. A minority of articles (14%) did not present a rationale, though these were often entries for conference competitions where objectives were set by the organisers.

Criterion 3: Selection of features

Some articles provided a clear, well-supported rationale for the features selected for their models (21%), while a larger proportion presented a weaker argument (31%). Because many researchers re-used existing datasets (see Criterion 6: Samples), the choice of features may have been limited. However, some researchers did investigate new territory. In 28% of articles, the approach taken was exploratory, for example where the relationship between novel features and personality was unknown, or where the model incorporated all available features using a deep learning or open vocabulary method. A minority of articles (19%) did not provide any explanation for the features used.

Criterion 4: Algorithm training

The vast majority of articles (83%) used cross-fold validation to train and test the models, which helps to minimise the risk of overfitting while making full use of the data in large datasets (Yarkoni & Westfall, 2017). A further 9% set aside an additional set for final-stage testing (sometimes as part of a competition). In 8% of articles there was no clear guidance on the model training method, however.

Criterion 5: Ground truth

The majority of articles (45%) reported the use of a valid, reliable measure as the ground truth. A further 20% referenced the article or dataset from which the ground truth was obtained, but without specifying the name of the measure or the number of items.

In 13% of articles observer judgements were used rather than self-reports; typically, these employed short measures of 10 items or less. In most cases the observers were MTurk workers rating short videos of individuals they did not personally know, which is likely to affect the reliability of the judgements.

In 26% of articles there was no clear indication that a valid ground truth had been used. In some cases, the measure was referred to using a generic name such as ‘Big Five’, and in a small number of articles bespoke, untested methods were developed to measure personality.

Of the articles for which information on the ground truth was available, the majority (60%) used relatively short scales with 60 items or less, for example the 44-item BFI, the 10-item TIPI and the BFI-10. Scales with very few items may have lower reliability and depth of coverage per trait. MyPersonality participants could answer varying numbers of items and scales, so reliability may vary. Figure 2 visually illustrates the length of the measures used.

[INSERT FIGURE 2 ABOUT HERE]

Criterion 6: Sample

Of the 194 datasets used across the articles, 38% were original datasets developed by the authors, and 62% were pre-existing datasets (for example, the myPersonality corpus). Figure 3 illustrates the proportion of dataset uses across the articles reviewed. Nearly half of the studies relied on one of three pre-existing datasets. Although this is advantageous in terms of reanalysing the same data and seeking optimum approaches, the downside is that such datasets are now becoming outdated. Additionally, six of the frequently used datasets (for example,

PAN-AP 2015) and nearly a quarter of the original datasets comprised 100 or fewer participants, which may affect the reliability and generalisability of the findings.

Demographics were not consistently reported across the articles; most common were basic details such as age and gender (28%), with only 6% containing additional details. A further 28% provided references to secondary sources rather than reporting details within the article. The majority of articles (39%) provided no information about participants at all.

Although most articles used reliable measures of traits as the ground truth, the vast majority did not report or comment on the sample means and variances or attempt to statistically correct for uneven distributions.

[INSERT FIGURE 3 ABOUT HERE]

Criterion 7: Regression / Classification

Most articles (50%) used a classification approach to predict personality, typically creating a binary split between high and low trait scores. Regression approaches were used in 43% of articles. A minority of articles (7%) used both regression and classification approaches.

Although many articles provided lengthy descriptions of the design of their models, these were often not accessible to readers outside of computer science. In most cases it would be challenging for researchers from other fields to attempt to replicate the study design.

Criterion 8: Analysis

Most articles (52%) clearly reported the performance of the model's predictions using a single metric, most commonly r and $RMSE$ for regression models, and accuracy for classification models, though there did not appear to be consistency in the choice of metric. This makes it more difficult to compare findings across studies. It has been recommended that researchers use more than one metric to give a clearer picture of the model's performance (Stachl, Pargent, et al., 2020a), however only 33% of articles did so. A further 15% of articles

reported performance in a manner which was vague or unclear; often these were conference competition submissions.

In addition to measures of performance, it is helpful to know whether models are predicting traits with a high degree of specificity, because otherwise their utility may be limited. Only one article reported a measure of discriminant validity, however. It would also be useful for researchers to report the relative predictive value of different features, but this was not commonly done.

Additionally, although many articles reported the performance of different types of models in predicting personality, it was rare for the differences in performance to be evaluated statistically, for example through the reporting of p values, effect sizes or confidence intervals.

Criteria 9 and 10: Ethical approval, open science practices, and data availability

Adherence to open science practices and explicit consideration of ethical issues was uncommon. Very few articles (6%) stated ethical approval for their research, and no articles reported pre-registration of their research plans. Only 4% of articles provided working links to their datasets. A further 17% provided links to datasets which were no longer publicly accessible, or which were not specific to the subsample used in the article. The remainder (79%) did not provide any access.

Discussion

This review found that personality computing studies adopt varied approaches in terms of study design, implementation, and reporting. These approaches often run counter to “best practice” guidelines or recommendations published in medicine and other relevant reviews that discuss machine learning and/or psychological science (Moons et al., 2015; Stachl, Pargent, et al., 2020a). Below, we discuss these findings and provide recommendations for future research.

The development and reporting of hypotheses and study rationales are often described with little detail and limited use of theory (if any at all). In some instances, both are reported, yet they often appear to be underdeveloped and lacking in direction (i.e., they are stated as ideas or expectations and are not linked to previous evidence in any depth). This is likely due to the expectations and general practice of machine learning research, where researchers prioritise the development of algorithms rather than deriving hypotheses/aims from existing evidence (Yarkoni & Westfall, 2017). As such, efforts were therefore directed toward the experimentation of the predictive performance of algorithms and different feature sets rather than gaining an understanding of why certain types of data relate to personality traits (e.g., Farnadi et al., 2016; Marouf, Hasan, & Mahmud, 2020; Sun, Liu, Cao, Luo, & Shen, 2018).

Although existing studies highlight a vast range of digital data (and aggregations of different data types) used to predict personality, there appeared to be a lack of rationale for selecting or combining particular types of data. While this reflects a data driven approach, explaining what the subsequent findings may mean becomes challenging. Further, integrating or building upon existing evidence could still help to inform data driven research – for instance, it could inform experimentation of data that has not been used, or dissuade the use of data that has not been effective in predicting personality in previous research. Indeed, reflecting on such decisions would be useful in improving future research and could help to advance theory testing and development (Hinds et al., 2021). Further, data driven approaches could reveal new, unexpected or conflicting relationships between data and personality compared to those presented in the psychology literature, which could inspire new ways to research and assess personality. Likewise, seeking to create interpretable models could help to support this, rather than the “black box” models currently used, which are largely inaccessible and difficult to understand (Molnar, 2019; Rudin, 2019). However, at present personality computing does not

explore the patterns that emerge in any meaningful way and has not yet capitalised on these opportunities.

Although personality questionnaires were used as ground truth in most studies, over a quarter of studies did not report any information whatsoever. This distinct lack of detail is problematic because it hampers the potential for studies to be replicated. Further, if the ground truth was a weak reflection of individuals' personalities (or poorly chosen/defined) the validity of a model's performance is compromised. As a result, a model that is highly accurate at predicting personality, would offer little practical value in practice. Thus, the absence of such information means that these potential issues cannot be evaluated. Our findings also highlighted that in some cases researchers opt to define the ground truth themselves (e.g., classifying individuals as extraverts or introverts based on the number of their friends/followers, Nguyen, Phung, Adams and Venkatesh, 2011), or by using crowdsourced ratings of others' personalities from MTurk workers (e.g., Biel & Gatica-Perez, 2013; Biel, Tsiminaki, Dines, & Gatica-Perez, 2013; Teijeiro-Mosquera, Biel, Alba-Castro, & Gatica-Perez, 2015). Although these approaches are less time consuming and costly than administering personality questionnaires directly to individuals, they are unlikely to provide robust assessments of their personalities – bespoke methods of defining personality have not been tested or developed to assess their reliability / validity, and personality assessments made by strangers are considerably worse than those made by people who know them (Connelly & Ones, 2010). Further, MTurk has also been criticised for creating data that are false, inaccurate, or misleading due to workers' lack of attention, and even fraud (e.g., Dennis, Goodson, & Pearson, 2020). Taking care to select valid measures and reporting the associated detail would therefore benefit future research.

Further, our findings highlighted that many of the studies that did report detail regarding ground truth relied on short form measures (i.e., TIPI, BFI-10). Such measures are limited in

that they are less reliable than longer multiitem measures (e.g., Gosling, Rentfrow, & Swann, 2003a; Rammstedt & John, 2007), and are less capable at capturing the facets comprising each trait (meaning that predictions are limited to broad personality trait domains, e.g., Chittaranjan, Blom, & Gatica-Perez, 2013; Gao, Shao, & Flora, 2019; Stachl, Au, et al., 2020). Such issues are likely to limit the effectiveness of the models when used in practice. The limited detail regarding the reporting of demographic information of the samples used could also exacerbate these problems. Without such information it is impossible to decipher whether samples comprise demographics that are representative of the general population, and indeed whether the implementation of such algorithms are likely to be biased (Buolamwini, 2018; Kirkpatrick, 2016).

The re-analysis of existing datasets is a commonly applied approach across many disciplines that use machine learning, and while this offers the advantage of testing the replicability of certain approaches and optimising the performance of different algorithms, the field does not seem to have capitalised on these opportunities. This means that similar approaches are often adopted, without evaluation or extension of existing research. Additionally, reliance on pre-existing, but outdated datasets runs the risk of ‘concept drift’ where features specific to a social media site, for example, change over time and are no longer reliable or usable predictors (e.g., Tsymbal, 2004; Widmer & Kubat, 1993). Researchers should be mindful of concept drift and seek to select features that are likely to be stable reflections of personality (Stachl, Pargent, et al., 2020a).

Further, the approaches used to assess the performance of different models are generally inconsistent, for instance singular metrics (e.g., accuracy, RMSE, MAE) are regularly used, which makes comparisons across studies difficult. This may be a product of publication bias, as well as due to a lack of formal guidance on what metrics should be reported. Similarly, the predominant use of classification (rather than regression) to evaluate a model’s performance

means that the accuracy of predictions could result in overinflated predictions. Classification is not ideal because it separates personalities into simplified “types” and ignores the fact that the majority of the population will fall somewhere in the middle of these two extremes (in accordance with a normal distribution, Stachl, Pargent, et al., 2020). While regression-based approaches offer a means to overcome this problem by analysing in finer detail, they are still limited; if accuracy is low, then their ability to distinguish between users in the middle of the distribution may be hindered. Going forward, researchers could perform multiple analyses, as well as make their data available so that other researchers could analyse the data (see Stachl, Pargent, et al., 2020 for a discussion of these issues).

Finally, our findings highlighted a need to improve the reporting of ethical procedures and open science practices, which were only reported in a minority of studies (less than 7%). Similarly, even in cases where data were made available, links had often become defunct, and obtaining the specific sub-sample of data (when data were taken from larger datasets such as myPersonality) was not possible or difficult to achieve. It is difficult to ascertain why such practices have not been more widely adopted – perhaps they are not well known or are considered to be low in priority compared to other aspects of conducting/reporting research. Increasing awareness within personality computing communities and via journals and academic conferences is therefore needed to encourage the adoption and adherence to these practices.

Limitations

Although systematic reviews of this nature are commonplace in medicine and healthcare research, they are seldom performed in machine learning and associated fields such as personality computing. As such, similar analyses performed in medicine are often informed and supported by formal guidelines that have been developed and approved by a body of experts such as the CONSORT (Consolidated Standards of Reporting Trials, Schulz, Altman,

and Moher, 2010) and SPIRIT (Standard Protocol Items: Recommendations for Interventional Trials, Chan, Tetzlaff, Altman, Dickersin and Moher, 2013). While the absence of such guidelines for personality computing studies presented a unique opportunity to begin to address this gap, the criteria we use may be subject to our own biases and subjective judgment. Future research could therefore attempt to develop these criteria more rigorously.

Further our study was also limited in that the number of studies included is large and continues to grow rapidly. From an analytical perspective, this is challenging because the coding process requires a time consuming and detailed analysis of each article. Whilst we were able to come these issues by employing multiple researchers, differences in perspectives or judgments may have introduced inconsistencies in the ratings. Although we made efforts to account for this by double (sometimes triple) coding a subsample of the articles, discussing and resolving disagreements, as well as discussing particularly challenging or unusual findings, this does not guarantee that we were able to stave off all potential discrepancies.

Conclusion

The review highlights that although personality computing research is rapidly growing, the design and reporting of studies are often lacking in detail, which raises questions over the validity, reliability, and replicability of the findings. Researchers focus heavily on improving the performance of algorithms rather than building on existing evidence and attempting to understand how data relate to personality and human behaviour. These findings stand in stark contrast to the popularity of this topic, as reflected in the publication of special issues in journals, competitions and workshops in conferences, and even stories reported in the news. Going forward, interdisciplinary collaboration and improved communication will be critical in improving current practice. We argue that interdisciplinary methods and perspectives (rather than approaches that “jump on the bandwagon” of existing practice), will help build a better

evidence base and will cultivate a healthy scientific community. Such approaches will enable us to truly understand the effectiveness and utility of predicting personality from digital data.

References

- Anderson, I., Gil, S., Gibson, C., Wolf, S., Shapiro, W., Semerci, O., & Greenberg, D. M. (2020). “Just the Way You Are”: Linking Music Listening on Spotify and Personality. *Social Psychological and Personality Science*.
<https://doi.org/10.1177/1948550620923228>
- Azucar, D., Marengo, D., & Settanni, M. (2018a). Predicting the Big 5 personality traits from digital footprints on social media: A meta-analysis. *Personality and Individual Differences*. <https://doi.org/10.1016/j.paid.2017.12.018>
- Azucar, D., Marengo, D., & Settanni, M. (2018b). Predicting the Big 5 personality traits from digital footprints on social media: A meta-analysis. *Personality and Individual Differences*, 124(December 2017), 150–159. <https://doi.org/10.1016/j.paid.2017.12.018>
- Biel, J. I., & Gatica-Perez, D. (2013). The YouTube Lens: Crowdsourced Personality Impressions and Audiovisual Analysis of Vlogs. *IEEE Transactions on Multimedia*.
<https://doi.org/10.1109/TMM.2012.2225032>
- Biel, J. I., Tsiminaki, V., Dines, J., & Gatica-Perez, D. (2013). Hi youtube!: Personality impressions and verbal content in social video. In *In Proceedings of the 15th ACM on International conference on multimodal interaction* (pp. 119–126). ACM.
- Bleidorn, W., & Hopwood, C. J. (2019). Using Machine Learning to Advance Personality Assessment and Theory. *Personality and Social Psychology Review*.
<https://doi.org/10.1177/1088868318772990>
- Botchkarev, A. (2019). A new typology design of performance metrics to measure errors in machine learning regression algorithms. *Interdisciplinary Journal of Information, Knowledge, and Management*, 14, 45–76. <https://doi.org/10.28945/4184>
- Buolamwini, J. (2018). *Gender Shades: Intersectional Accuracy Disparities in Commercial Gender Classification* *. *Proceedings of Machine Learning Research* (Vol. 81).

- Cadwalladr, C., Graham-Harrison, E. (2018). Revealed: 50 million Facebook profiles harvested for Cambridge Analytica in major data breach.
- Chan, A. W., Tetzlaff, J. M., Altman, D. G., Dickersin, K., & Moher, D. (2013). SPIRIT 2013: New guidance for content of clinical trial protocols. *The Lancet*.
[https://doi.org/10.1016/S0140-6736\(12\)62160-6](https://doi.org/10.1016/S0140-6736(12)62160-6)
- Chittaranjan, G., Blom, J., & Gatica-Perez, D. (2013). Mining large-scale smartphone data for personality studies. In *Personal and Ubiquitous Computing*.
<https://doi.org/10.1007/s00779-011-0490-1>
- Christodoulou, E., Ma, J., Collins, G. S., Steyerberg, E. W., Verbakel, J. Y., & Van Calster, B. (2019). A systematic review shows no performance benefit of machine learning over logistic regression for clinical prediction models. *Journal of Clinical Epidemiology*.
<https://doi.org/10.1016/j.jclinepi.2019.02.004>
- Collaboration, O. S. (2012). An open, large-scale, collaborative effort to estimate the reproducibility of psychological science. *Perspectives on Psychological Science* 2, 7(6), 657–660.
- Connelly, B. S., & Ones, D. S. (2010). An Other Perspective on Personality: Meta-Analytic Integration of Observers' Accuracy and Predictive Validity. *Psychological Bulletin*.
<https://doi.org/10.1037/a0021212>
- DeMasi, O., Kording, K., & Recht, B. (2017). Meaningless comparisons lead to false optimism in medical machine learning. *PLoS ONE*, 12(9), e0184604.
- Dennis, S. A., Goodson, B. M., & Pearson, C. A. (2020). Online worker fraud and evolving threats to the integrity of mturk data: A discussion of virtual private servers and the limitations of ip-based screening procedures. *Behavioral Research in Accounting*, 32(1).
<https://doi.org/10.2308/bria-18-044>
- Faes, L., Liu, X., Kale, A., Bruynseels, A., Shamdas, M., Moraes, G., ... Denniston, A.

- (2020). Deep Learning Under Scrutiny: Performance Against Health Care Professionals in Detecting Diseases from Medical Imaging - Systematic Review and Meta-Analysis. *SSRN Electronic Journal*. <https://doi.org/10.2139/ssrn.3384923>
- Farnadi, G., Sitaraman, G., Sushmita, S., Celli, F., Kosinski, M., Stillwell, D., ... De Cock, M. (2016). Computational personality recognition in social media. *User Modeling and User-Adapted Interaction*. <https://doi.org/10.1007/s11257-016-9171-0>
- Flach, P. (2019). Performance Evaluation in Machine Learning: The Good, the Bad, the Ugly, and the Way Forward. *Proceedings of the AAAI Conference on Artificial Intelligence*, 33, 9808–9814. <https://doi.org/10.1609/aaai.v33i01.33019808>
- Gao, N., Shao, W., & Flora, D. S. (2019). Predicting Personality Traits from Physical Activity Intensity. *Computer*, 52(7), 47–56.
- Gladstone, J. J., Matz, S. C., & Lemaire, A. (2019). Can Psychological Traits Be Inferred From Spending? Evidence From Transaction Data. *Psychological Science*. <https://doi.org/10.1177/0956797619849435>
- Gosling, S. D., Rentfrow, P. J., & Swann, W. B. (2003). A very brief measure of the Big-Five personality domains. *Journal of Research in Personality*. [https://doi.org/10.1016/S0092-6566\(03\)00046-1](https://doi.org/10.1016/S0092-6566(03)00046-1)
- Hinds, J., Joinson, A. (2019). Human and computer personality prediction from digital footprints. *Current Directions in Psychological Science*, 28(2), 204–211.
- Hinds, J., Brown, O., Smith, L. G. E., Piwek, L., Ellis, D. A., & Joinson, A. N. (2021). Integrating Insights About Human Movement Patterns from Digital Data into Psychological Science. *Working Paper*.
- Hinds, J., & Joinson, A. (2020). Predicting Personality from Digital Footprints: Current Challenges and Future Directions. *Working Paper*.
- Hinds, Joanne, Williams, E. J., & Joinson, A. N. (2020). “It wouldn’t happen to me”: Privacy

- concerns and perspectives following the Cambridge Analytica scandal. *International Journal of Human Computer Studies*. <https://doi.org/10.1016/j.ijhcs.2020.102498>
- Hong, Q. N., Fàbregues, S., Bartlett, G., Boardman, F., Cargo, M., Dagenais, P., ... Pluye, P. (2018). The Mixed Methods Appraisal Tool (MMAT) version 2018 for information professionals and researchers. *Education for Information*, 34(4).
<https://doi.org/10.3233/EFI-180221>
- John, O. P., & Srivastava, S. (1999). The Big Five trait taxonomy: History, measurement, and theoretical perspectives. *Handbook of Personality: Theory and Research*, 2(510).
- Kedar, S. V., & Bormane, D. S. (2016). Automatic personality assessment: A systematic review. In *Proceedings - IEEE International Conference on Information Processing, ICIP 2015*. <https://doi.org/10.1109/INFOP.2015.7489402>
- Kirkpatrick, K. (2016). Battling algorithmic bias: How do we ensure algorithms treat us fairly? *Communications of the ACM*, 59(10).
- Kosinski, M., Wang, Y., Lakkaraju, H., & Leskovec, J. (2016). Mining big data to extract patterns and predict real-life outcomes. *Psychological Methods*, 21(4).
<https://doi.org/10.1037/met0000105>
- Liu, W., & Ruths, D. (2013). What's in a Name? Using First Names as Features for Gender Inference in Twitter. *Analyzing Microtext: Papers from the 2013 AAAI Spring Symposium*.
- Luo, W., Phung, D., Tran, T., Gupta, S., Rana, S., Karmakar, C., ... Berk, M. (2016). Guidelines for developing and reporting machine learning predictive models in biomedical research: A multidisciplinary view. *Journal of Medical Internet Research*, 18(12). <https://doi.org/10.2196/jmir.5870>
- Mahmoodi, J., Leckelt, M., van Zalk, M. W. H., Geukes, K., & Back, M. D. (2017). Big Data approaches in social and behavioral science: four key trade-offs and a call for

integration. *Current Opinion in Behavioral Sciences*.

<https://doi.org/10.1016/j.cobeha.2017.07.001>

Marengo, D., & Montag, C. (2020). Digital Phenotyping of Big Five Personality via Facebook Data Mining: A Meta-Analysis. *Digital Psychology*.

<https://doi.org/10.24989/dp.v1i1.1823>

Marouf, A. Al, Hasan, M. K., & Mahmud, H. (2020). Comparative Analysis of Feature Selection Algorithms for Computational Personality Prediction from Social Media. *IEEE Transactions on Computational Social Systems*.

<https://doi.org/10.1109/TCSS.2020.2966910>

Maxwell, S. E., Lau, M. Y., & Howard, G. S. (2015). Is psychology suffering from a replication crisis?: What does “failure to replicate” really mean? *American Psychologist*.

<https://doi.org/10.1037/a0039400>

Mohammady, E., & Culotta, A. (2014). Using county demographics to infer attributes of Twitter users. In *n ACL Joint Workshop on Social Dynamics and Personal Attributes in Social Media*.

Moher, D., Liberati, A., Tetzlaff, J., & Altman, D. G. (2009). Preferred reporting items for systematic reviews and meta-analyses: the PRISMA statement. *Journal of Clinical Epidemiology*. <https://doi.org/10.1016/j.jclinepi.2009.06.005>

Molnar, C. (2019). Interpretable Machine Learning. A Guide for Making Black Box Models Explainable. *Book*.

Mønsted, B., Mollgaard, A., & Mathiesen, J. (2018). Phone-based metric as a predictor for basic personality traits. *Journal of Research in Personality*.

<https://doi.org/10.1016/j.jrp.2017.12.004>

Moons, K. G. M., Altman, D. G., Reitsma, J. B., Ioannidis, J. P. A., Macaskill, P.,

Steyerberg, E. W., ... Collins, G. S. (2015). Transparent reporting of a multivariable

- prediction model for individual prognosis or diagnosis (TRIPOD): Explanation and elaboration. *Annals of Internal Medicine*, 162(1). <https://doi.org/10.7326/M14-0698>
- Music Business Worldwide. (2020). Spotify has a patent for personality tracking technology - and it's pretty creepy stuff. Retrieved November 8, 2020, from <https://www.musicbusinessworldwide.com/spotify-has-a-patent-for-personality-tracking-technology-and-its-pretty-creepy/>
- Nguyen, T., Phung, D., Adams, B., & Venkatesh, S. (2011). Towards discovery of influence and personality traits through social link prediction. *Proceedings of the Fifth International AAAI Conference on Weblogs and Social Media*. <https://doi.org/10.1097/01.TP.0000131167.21930.B8>
- Open Science Collobaration. (2015). Estimating the reproducibility of psychological science: Open Science Collobaration. *Science*. <https://doi.org/10.1126/science.aac4716>
- Parra, D., Trattner, C., Gómez, D., Hurtado, M., Wen, X., & Lin, Y. R. (2016). Twitter in academic events: A study of temporal usage, communication, sentimental and topical patterns in 16 Computer Science conferences. In *Computer Communications*. <https://doi.org/10.1016/j.comcom.2015.07.001>
- Publications, A. P. A., & on Journal, C. B. W. G. (2008). Reporting standards for research in psychology: Why do we need them? What might they be? *The American Psychologist*, 63(9), 839.
- Rammstedt, B., & John, O. P. (2007). Measuring personality in one minute or less: A 10-item short version of the Big Five Inventory in English and German. *Journal of Research in Personality*, 41(1). <https://doi.org/10.1016/j.jrp.2006.02.001>
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nature Machine Intelligence*. <https://doi.org/10.1038/s42256-019-0048-x>

- Sanders, H., & Saxe, J. (2017). Garbage in, garbage out (How purportedly great ML models can be screwed up by bad data). *Proceedings of Blackhat 2017*.
- Schmitt, D. P., Allik, J., McCrae, R. R., Benet-Martínez, V., Alcalay, L., Ault, L., ... Sharan, M. B. (2007). The geographic distribution of Big Five personality traits: Patterns and profiles of human self-description across 56 nations. *Journal of Cross-Cultural Psychology*, 38(2). <https://doi.org/10.1177/0022022106297299>
- Schulz, K. F., Altman, D. C., & Moher, D. (2010). CONSORT 2010 Statement: Updated guidelines for reporting parallel group randomised trials. *Italian Journal of Public Health*, 7(3). <https://doi.org/10.4178/epih/e2014029>
- Seliya, N., Khoshgoftaar, T. M., & Van Hulse, J. (2009). A study on the relationships of classifier performance metrics. *Proceedings - International Conference on Tools with Artificial Intelligence, ICTAI*, 59–66. <https://doi.org/10.1109/ICTAI.2009.25>
- Settanni, M., Azucar, D., & Marengo, D. (2018). Predicting Individual Characteristics from Digital Traces on Social Media: A Meta-Analysis. *Cyberpsychology, Behavior, and Social Networking*. <https://doi.org/10.1089/cyber.2017.0384>
- Souri, A., Hosseinpour, S., Rahmani, A. M. (2018). Personality classification based on profiles of social networks' users and the five-factor model of personality. *Human-Centric Computing and Information Sciences*.
<https://doi.org/https://doi.org/10.1186/s13673-018-0147-4>
- Stachl, C., Au, Q., Schoedel, R., Gosling, S. D., Harari, G. M., Buschek, D., ... Bühner, M. (2020). Predicting personality from patterns of behavior collected with smartphones. *Proceedings of the National Academy of Sciences of the United States of America*.
<https://doi.org/10.1073/pnas.1920484117>
- Stachl, C., Pargent, F., Hilbert, S., Harari, G. M., Schoedel, R., Vaid, S., ... Bühner, M. (2020a). Personality Research and Assessment in the Era of Machine Learning.

European Journal of Personality. <https://doi.org/10.1002/per.2257>

Stachl, C., Pargent, F., Hilbert, S., Harari, G. M., Schoedel, R., Vaid, S., ... Bühner, M.

(2020b). Personality Research and Assessment in the Era of Machine Learning.

European Journal of Personality, 34(5), 613–631. <https://doi.org/10.1002/per.2257>

Stuart Geiger, R., Yu, K., Yang, Y., Dai, M., Qiu, J., Tang, R., & Huang, J. (2020). Garbage

in, garbage out? Do machine learning application papers in social computing report

where human-labeled training data comes from? In *FAT* 2020 - Proceedings of the*

2020 Conference on Fairness, Accountability, and Transparency.

<https://doi.org/10.1145/3351095.3372862>

Subramanian, J., & Simon, R. (2013). Overfitting in prediction models - Is it a problem only

in high dimensions? *Contemporary Clinical Trials*, 36(2).

<https://doi.org/10.1016/j.cct.2013.06.011>

Sun, X., Liu, B., Cao, J., Luo, J., & Shen, X. (2018). Who am I? Personality detection based

on deep learning for texts. In *IEEE International Conference on Communications*.

<https://doi.org/10.1109/ICC.2018.8422105>

Teijeiro-Mosquera, L., Biel, J. I., Alba-Castro, J. L., & Gatica-Perez, D. (2015). What your

face vlogs about: Expressions of emotion and big-five traits impressions in youtube.

IEEE Transactions on Affective Computing.

<https://doi.org/10.1109/TAFFC.2014.2370044>

Tskhay, K. O., & Rule, N. O. (2014). Perceptions of personality in text-based media and

OSN: A meta-analysis. *Journal of Research in Personality*.

<https://doi.org/10.1016/j.jrp.2013.12.004>

Tsymbol, A. (2004). The problem of concept drift: definitions and related work. *Computer*

Science Department, Trinity College Dublin.

Vinciarelli, A., & Mohammadi, G. (2014). A survey of personality computing. *IEEE*

Transactions on Affective Computing. <https://doi.org/10.1109/TAFFC.2014.2330816>

Widmer, G., & Kubat, M. (1993). Effective learning in dynamic environments by explicit context tracking. In *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* (Vol. 667 LNAI). https://doi.org/10.1007/3-540-56602-3_139

Yarkoni, T., & Westfall, J. (2017). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*. <https://doi.org/10.1177/1745691617693393>

Yusuf, M., Atal, I., Li, J., Smith, P., Ravaud, P., Fergie, M., ... Selfe, J. (2020). Reporting quality of studies using machine learning models for medical diagnosis: a systematic review. *BMJ Open*. <https://doi.org/10.1136/bmjopen-2019-034568>

Table 1*Rating criteria and findings¹*

Section	Quality criteria	% of articles per category	
Hypotheses	4 = Clearly stated, directional hypotheses	7	
	3 = Clearly stated, but lacking direction or theory	8	
	2 = Stated, but weak/vague or framed as general aims	58	
	1 = Not stated	27	
Rationale	4 = Clearly stated rationale~ supported by theory/evidence	31	
	3 = Clearly stated, but either weak/vague	29	
	2 = Stated, but not in a clear manner	25	
	1 = Not stated	14	
Selection of features	4 = Clearly stated rationale supported by theory/evidence	21	
	3 = Clearly stated, but either weak/vague	31	
	2 = Exploratory analysis of novel model	28	
	1 = No theoretical rationale offered and not novel features	19	
Algorithm training	3 = Trained and tested using crossfold validation	92	
	2 = Trained and tested on the same dataset, without partitions	1	
	1 = No clear guidance provided	8	
Ground truth	4 = Valid, reliable ground truth method reported	45	
	3 = Ground truth based on observer estimates reported	13	
	2 = Ground truth information accessible via secondary sources	20	
	1 = Ground truth method not specified, or not reported	26	
Sample	4 = Detailed information provided	6	
	3 = Basic information provided	28	
	2 = Information only available via secondary sources	28	
	1 = No details provided	39	
Regression/ classification	5 = Regression based	43	
	4 = Both regression and classification	7	
	3 = Transparent classification with data available	2	
	2 = Classification without data available	48	
	1 = Unclear	1	

¹ NB: Percentages may not add to 100 as some articles contained multiple studies

Table 1 continued.

Analysis	4 = <i>Multiple performance measures are used</i>	33
	3 = <i>One performance measure is used</i>	52	
	2 = <i>Reporting or interpretation poor/vague</i>	11	
	1 = <i>Reporting or interpretation unclear</i>	4	
Ethics	3 = <i>Ethical approval stated</i>	6
	1 = <i>Ethical approval not stated</i>	94	
Open-source data	3 = <i>Full access to the specific dataset used to train the model</i>	4
	2 = <i>Dataset no longer accessible, or not specific to the model</i>	17	
	1 = <i>No dataset made available</i>	79	

Figure 2

Length of measures used in computer-prediction articles

Figure 3

The use of original and existing datasets across computer prediction articles.