



HAL
open science

Inferring psychological traits from spending categories and dynamic consumption patterns

Natkamon Tovanich, Simone Centellegher, Nacéra Bennacer Seghouani, Joe Gladstone, Sandra Matz, Bruno Lepri

► **To cite this version:**

Natkamon Tovanich, Simone Centellegher, Nacéra Bennacer Seghouani, Joe Gladstone, Sandra Matz, et al.. Inferring psychological traits from spending categories and dynamic consumption patterns. EPJ Data Science, 2021, 10 (1), pp.24. 10.1140/epjds/s13688-021-00281-y . hal-03237636

HAL Id: hal-03237636

<https://hal.science/hal-03237636v1>

Submitted on 26 Apr 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Inferring psychological traits from spending categories and dynamic consumption patterns

Natkamon Tovanich^{1†}, Simone Centellegher^{2*†}, Nacéra Bennacer Seghouani¹, Joe Gladstone³, Sandra Matz^{4*} and Bruno Lepri^{2*}

*Correspondence:

centellegher@fbk.eu

²Fondazione Bruno Kessler, Via Sommarive, 18, 38123 Trento, Italy
Full list of author information is available at the end of the article

[†]Equal contributors

[‡]Joint senior authors

Abstract

In recent years there has been a growing interest in analyzing human behavioral data generated by new technologies. One type of digital footprint that is universal across the world, but that has received relatively little attention to date, is spending behavior.

In this paper, using the transaction records of 1306 bank customers, we investigated the extent to which individual-level psychological characteristics can be inferred from bank transaction data. Specifically, we developed a more comprehensive feature space using: (1) overall spending behavior (i.e. total number and total amount of transaction), (2) temporal spending behavior (i.e. variability, persistence, and burstiness), (3) category-related spending behavior (i.e. diversity, persistence, and turnover), (4) customer category profile, and (5) socio-demographic information. Using these features, we first explore their association with individual psychological characteristics, we then analyze the performances of the different feature families and finally, we try to understand to what extent psychological characteristics from spending records can be inferred.

Our results show that inferring the psychological traits of an individual is a challenging task, even when using a comprehensive set of features that take temporal aspects of spending into account. We found that Materialism and Self-Control could be inferred with relatively high levels of accuracy, while the accuracy obtained for the Big Five traits was lower, with only Extraversion and Neuroticism reaching reasonable classification performances.

Hence, for traits like Materialism, Self-control, Extraversion, and Neuroticism our findings could be used to improve psychologically-informed advertising strategies for specific products as well as personality-based spending management apps and credit scoring approaches.

Keywords: Spending behavior; Personality traits; Bank transaction data; Computational psychology

1 Introduction

Over the past few decades, digital services and devices have become a central part of people's everyday lives. They help us communicate with our friends and loved ones, capture

© The Author(s) 2021. This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

the moments we care about the most, broadcast our opinions to millions of people around the world, search for information from the comfort of our homes, and pay for the things we want to buy with the ease of a tap or swipe. Recent advances in the field of computational social science [1] have shown that the digital footprints people leave behind on a daily basis can be used to make accurate predictions about their psychological profiles (see e.g. [2] for a summary). People's personality traits, for example, have been predicted from Facebook Likes [3, 4], the language in people's social media posts [5–7], profile pictures [8, 9], music preferences [10], and smartphone sensing data [11–14].

One type of digital footprint that is universal across the world, but that has received relatively little attention to date, is spending behavior. With around 80% of adults in high-income economies using a debit or a credit card [15], people's spending has become increasingly digitized, making it possible to capture consumer choices at an unprecedented scale. Recent research has begun to use transaction records from debit and credit card purchases to show how such data can provide important insights into the dispositions, attitudes, and preferences of individual customers [16–18].

Interestingly, research in consumer behavior suggests that spending serves an important psychological function because people buy products and brands not only for what they can do but also for what they mean and signal to others [19]. That is, spending often constitutes a form of self-expression that allows an individual to signal their identity to themselves and those around them (e.g. [19, 20]). Buying a subscription for the *Wall Street Journal*, for example, might signal an interest in business and a relatively high level of intellect, while buying flowers might signal a warm and caring personality. In line with the notion that consumers buy products and brands not just for what they can do but also for what they mean psychologically, numerous laboratory studies have shown that people report more favorable attitudes, emotions, and behaviors toward products and brands that match their own personality [21–23]. While extraverts, for example, might prefer spending their money on social activities (e.g., having drinks with friends), introverts might prefer to spend their money on activities that allow them to spend quiet me-time (e.g., listening to a podcast at home). Supporting these laboratory findings, recent evidence from the field has shown that people indeed spend more money on products and services that match their own personality [24] and that the extent to which people spend money on conspicuous goods is a function of both their financial means and level of Extraversion [25].

Inspired by this body of research, a recent study suggested that spending records can be used to automatically infer the psychological characteristics of individuals [26]. Using the transaction records of 2193 UK bank customers, the authors were able to predict the Big Five personality traits, Materialism and Self-Control with an accuracy ranging from $r = 0.15$ for the Big Five personality traits to $r = 0.33$ for Materialism. While these findings provide initial evidence that it is possible to predict psychological characteristics from spending records, the accuracy with which those traits can be inferred remains relatively low when compared to the accuracy obtained from other types of digital footprints [26].

The authors suggest that one potential reason for this is that different types of digital footprints may reveal more about an individual's personality than others. They argue that social media profiles can be seen to constitute explicit identity claims made by individuals, while transaction records represent more subtle and implicit behavioral residues. Another potential reason, however, could be that the relationship between spending records and psychological traits is more complex and dynamic than what the models implemented by

Gladstone *et al.* could capture [26]. In fact, their models rely on a simple set of features measuring the relative amounts spent in 279 broad categories (e.g. supermarkets, furniture stores, insurance policies, etc.) as well as a broader set of 34 topics reflecting combined spending across groups of individual brands (e.g. fast food chains, coffee shops, investment services, utility providers, electronics stores, etc.).

In this paper, we advance the research on the relationship between spending behaviors and personality traits by investigating whether the accuracy of inferring psychological characteristics from spending records can be improved when considering a more comprehensive space of behavioral features. More specifically, we develop features in 5 main categories: (1) overall spending behavior (i.e. total number and total amount of transaction), (2) temporal spending behavior (i.e. variability, persistence, and burstiness), (3) category-related spending behavior (i.e. diversity, persistence, and turnover), (4) customer category profile, and (5) socio-demographic information. Thus, we first explore their association with individual psychological characteristics, then we analyze the performances of the different feature families, and finally we try to understand to what extent individuals' psychological characteristics can be inferred from spending records. To this end, we use the aforementioned groups of spending metrics and train different machine learning models (i.e. Logistic Regression [27], Random Forest [28], and Extreme Gradient Boosting [29]) to classify the customers' psychological traits.

In line with the previous work of Gladstone *et al.* [26], our results show there are significant differences in the predictive accuracy across the different traits, with Materialism, Self-Control, Neuroticism, and Extraversion reaching higher classification performances than others. Our research further extends the earlier work by comparing different groups of features on their relative contribution to the predictive performance of our models. Notably, we find that temporal spending behaviors provide signals to improve the prediction of Self-Control and Neuroticism: people scoring high in Self-Control show more stable patterns in spending behavior, while neurotic people tend to show less persistence over time.

2 Materials and methods

2.1 Data

In this study, we investigate whether it is possible to use spending behavior to infer psychological characteristics at the individual level using a data set containing 74 million bank transaction records from 127,469 customers. A subset of 2193 customers from the larger sample provided responses to a survey which included measures of the following seven psychological characteristics: the Big Five personality traits (Openness, Conscientiousness, Extraversion, Agreeableness, Neuroticism), Materialism, and Self-Control. We make use of transactions recorded between June 2016 to March 2017 (over 10 months).

Bank transaction records The dataset was collected in collaboration with a UK-based money management app. The customer information was anonymized and included the following: the unique customer identifier (*userID*); the gender of the customer (*gender*); the year the customer was born (*YOB*); the salary range in British pounds (GBP) divided in 10K intervals (*salary range*); the customer home location (*home location*) specified in three levels of geographical granularity, namely postcode, Lower Layer Super Output Area (LSOA) and Middle Layer Super Output Area (MSOA).

The transaction information includes: the unique identifier of the transaction (*transactionID*); the anonymized identifier of the customer's bank account (*account number*); the customer identifier (*userID*); the type of transaction with a distinction between credit or debit (*transaction type*); the date of when the transaction was made (*transaction date*); the category of the transaction provided by the bank, e.g. supermarket, flights, concert, etc. (*transaction category*); the amount of the transaction in GBP (*transaction amount*).

Individual psychological characteristics The dataset contains the psychological profiles of bank customers who volunteered to participate in a survey. A survey link was sent to customers by e-mail asking them to participate in the study, with the opportunity to win a tablet computer. In total, 2193 customers completed the survey and provided their consent to participate and have their transaction data matched with their survey responses for research purposes. The survey included measures of the Big Five personality traits, Materialism, and Self-Control.

The Big Five personality model is the most widely accepted framework to describe relatively stable personality characteristics [30]. The model proposes the following five factors which capture individual differences in the way that people think, feel and behave: (i) *Extraversion*, the tendency to seek stimulation in the company of others, to be outgoing and energetic; (ii) *Agreeableness*, the tendency to be warm, compassionate, and cooperative; (iii) *Conscientiousness*, the tendency to show self-discipline, aim for achievement, and be organized; (iv) *Neuroticism*, the tendency to experience unpleasant emotions easily; and (v) *Openness to Experience*, the tendency to be intellectually curious, creative, and open to feelings.

The Big Five personality traits were measured using the established BFI-10 questionnaire, a short 10-item questionnaire with two items per trait [31]. Participants indicate their agreement with statements such as "I see myself as someone who is reserved", "I find myself as someone who tends to find faults with others", and "I see myself as someone who has an active imagination" using a 7-point Likert scale (1 = Strongly Disagree to 7 = Strongly Agree). For each trait, the sum scores can thus range between 2 and 14, indicating a very low or a very high level in that particular trait. While longer questionnaires with more items per personality trait are generally preferable, the particular context of data collection prohibited the ability to use long survey measures. Similar short versions of the questionnaire have been used in similar contexts related to financial decision making and have proven to capture significant variance in people's personality traits [32, 33].

The survey sent to bank customers also included measures for two other psychological traits: *Materialism*, and *Self-Control*. *Materialism* is the tendency to consider material possessions and physical comfort as more important than spiritual values. The trait is measured through the following three items taken from a widely used survey [34]: (i) "I admire people who own expensive homes, cars and clothes", (ii) "I like a lot of luxury in my life", and (iii) "I'd be happier if I could afford to buy more things". Similar to the Big Five, participants rated their agreement with these statements on a 7-point Likert scale ranging from 1 = Strongly Disagree to 7 = Strongly Agree. The sum scores for Materialism consequently range between 3 and 21. *Self-Control* is the ability to regulate emotions, thoughts, and behaviors in face of temptations and impulses. Here, the *Self-Control* construct was measured using a single item ("I am good at resisting temptation") from the Brief Self-Control

Scale [35]. The scores range between 1 and 7. For more details on the questionnaires used see the recent paper of Gladstone *et al.* [26].

2.2 Data preprocessing

The dataset contains two types of recorded activities: credit (incoming) and debit (outgoing) transactions. A credit transaction is an increase in the account balance (e.g. money deposit, salary, or other income), while a debit transaction is a decrease in the account balance (e.g. money withdrawal, payment, purchasing activities).

To analyze customers' spending behavior, we only retained debit transactions since they represent their spending activities. To assure a sufficient level of data per participant and capture only those customers that were actively using their account, we only retained customers with at least ten transactions per month. This exclusion procedure left us with 40,080 customers, 1306 of which responded to the psychological survey. This group of 1306 customers represents our final dataset. On average, participants in our sample were 40 years old, and the majority of them reported salaries ranging between 10K and 40K pounds. Figure 1 show the distributions of the individual psychological dispositions in our dataset.

To reduce the sparseness of the category space, we discarded the 172 purchase categories that had less than 10 percent of customer support. The customer support for a particular category is calculated as the percentage of customers who purchased at least once in that category. There are 108 categories with more than 10 percent of customer support, from which we removed 11 categories that were unrelated to spending activities, such as credit

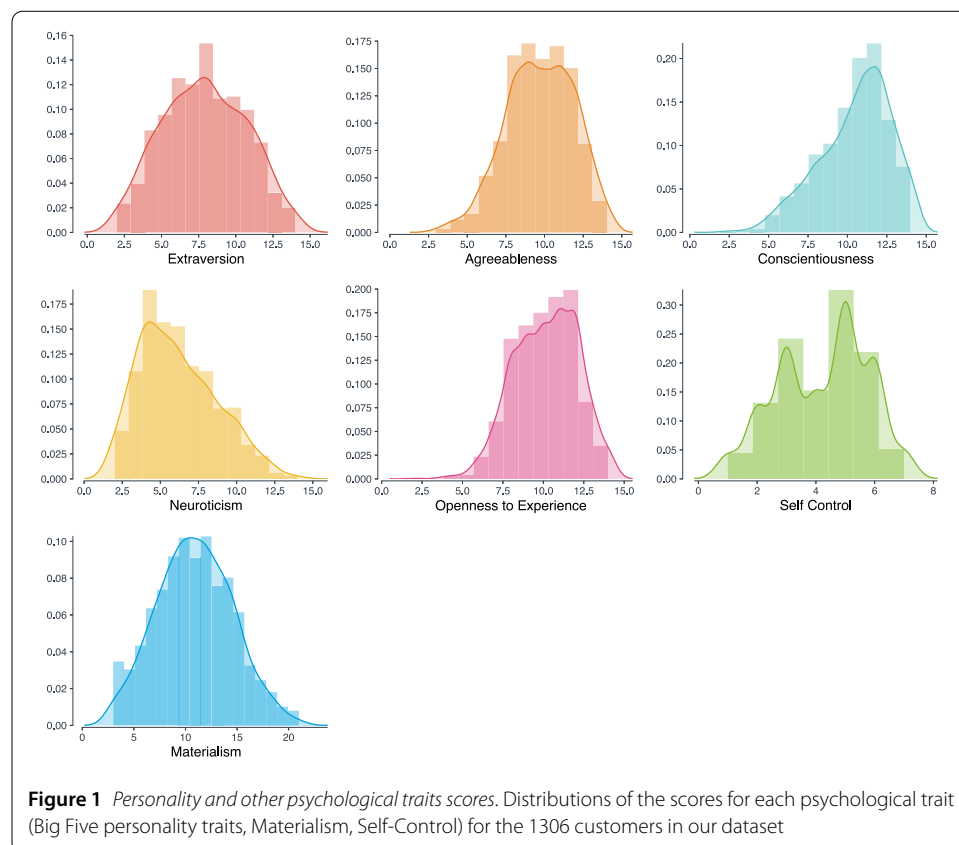


Table 1 *Category groups*. List of mappings between categories of purchases and our 35 category groups

Category	Category group	Category	Category group
Administration—other	Administration	Memberships	Hobbies
Alcohol	Alcohol	Photography	Hobbies
Art Supplies	Art and antiques	Toys	Hobbies
Books / Magazines / Newspapers	Books	Flights	Holiday
Cash	Cash withdrawn and transfer	Holiday	Holiday
Current account	Cash withdrawn and transfer	Holidays	Holiday
Transfers	Cash withdrawn and transfer	Hotel/B&B	Holiday
Charity—other	Charities	Museum/exhibition	Holiday
Donation to organisation	Charities	Zoo/theme park	Holiday
Sponsorship	Charities	Furniture	Household
Children—other	Children	Garden	Household
Clothes	Clothes	Home and garden—other	Household
Clothes—other	Clothes	Home electronics	Household
Shoes	Clothes	Household—other	Household
Broadband	Computer and internet	Kitchen / Household	Household
Software	Computer and internet	Appliances	
Web hosting	Computer and internet	Personal Electronics	Household
DIY	DIY projects	Electricity	Household: spending
Dry cleaning and laundry	Dry cleaning and laundry	Mortgage payment	Household: spending
Cinema	Entertainment: going out	Phone (landline)	Household: spending
Concert \& Theatre	Entertainment: going out	Rent	Household: spending
Media bundle	Entertainment: in house	Water	Household: spending
Music	Entertainment: in house	Home appliance insurance	Insurance
TV / Movies Package	Entertainment: in house	Health insurance	Insurance
TV Licence	Entertainment: in house	Home insurance	Insurance
Financial—other	Financial	Insurance—other	Insurance
Flowers	Flowers	Life insurance	Insurance
Dining and drinking	Food drink and going out	Vehicle insurance	Insurance
Dining or Going Out	Food drink and going out	Mobile	Mobile
Going out—other	Food drink and going out	Mobile app	Mobile
Lunch or Snacks	Food drink and going out	Hairdressing	Personal care and beauty
Take-away	Food drink and going out	Personal Care—other	Personal care and beauty
Gambling	Gambling	Toiletries	Personal care and beauty
Gambling account	Gambling	Pet Insurance	Pet
Games and gaming	Games and gaming	Pets—other	Pet
Gas	Gas and electricity	Vet	Pet
Gas and electricity	Gas and electricity	Postage / Shipping	Postage / Shipping
Birthday present	Gifts and presents	Service / Parts / Repairs	Service / Parts / Repairs
Christmas present	Gifts and presents	Cycling	Sports
Gifts—other	Gifts and presents	Gym Membership	Sports
Gifts or Presents	Gifts and presents	Sports Equipment	Sports
Food, Groceries, Household	Groceries and supermarkets	Stationery	Stationery
Groceries	Groceries and supermarkets	Fuel	Transportation
Supermarket	Groceries and supermarkets	Parking	Transportation
Dental treatment	Health care	Public Transport	Transportation
Eye care	Health care	Road charges	Transportation
Hobbies—other	Hobbies	Taxi	Transportation
Lifestyle—other	Hobbies	Transport—other	Transportation
		Vehicle hire	Transportation
		Vehicle tax	Transportation

card repayment and individual saving accounts. The final sample included 97 purchase categories.

Finally, we manually classified the purchase categories into 35 category groups as shown in Table 1. For example, we combined five categories that are related to regular household spending (*Electricity, Mortgage payment, Phone (landline), Rent and Water*) into *Household: spending*.

2.3 Characterizing spending behavior

To characterize the spending behavior of each customer, we calculated several behavioral features from the bank transaction data. We then grouped these features into five categories, according to the type of spending behavior they capture: (i) overall spending behavior, (ii) temporal spending behavior, (iii) category-related spending behavior, (iv) customer category profile and (v) socio-demographic information.

G1. Overall spending behavior The features in the overall spending behavior category were computed over the entire period of study. We defined summary statistics of customers' spending behavior as the total number of transactions (n_{tot}), the total amount (a_{tot}) a customer had spent over that period, and the average amount per transaction (a_{avg}) spent by each customer. Since the distributions of these spending related metrics are positively skewed, we applied a log scaled transformation to these three features.

In order to measure the relative variability of a customer spending behavior, we used the coefficient of variation $cv = \frac{\sigma}{\mu}$, defined as the ratio of the standard deviation of the amount of transactions (σ) to the average amount of transactions (μ). When cv is large, this indicates that the customer tends to spend unequally on different transactions and vice-versa.

G2. Temporal spending behavior An important aspect of spending which has largely been overlooked by previous research is the temporal dimension along which this behavior occurs.

In the literature, the association between the temporal aspects of human behavior and the individual psychological characteristics was partially studied. For example, in [13], authors found an association between calls/SMS regularity, the response latency to text messages, and the Big Five personality traits. In [36], the authors found that the frequency of Facebook use and posting is higher for extroverted people. Again looking at smartphone usage behavior, the average time from the notification arrival until the time the notification was seen and acted upon by the user is correlated with depression [37]. Inspired by these works, with the features that we devise, we try to investigate whether there is an association between temporal aspects of spending behavior and the individual psychological characteristics under study.

We chose to analyze the temporal spending behavior at different granularity using three time windows t : month (M), 10-days intervals (P), and day of the week (D) ($t = \{M, P, D\}$). We chose these time windows in order to take into account the seasonal differences in spending behavior. For example, the 10-days intervals can help to account for differences in spending which are due to when a customer receives his/her salary.

For each customer and time unit, we measured the temporal patterns of spending behavior calculating (i) the variability of the spending amount, (ii) the persistence of the spending patterns, and (iii) the presence of bursty spending behavior.

(i) *Variability of spending amount.* In order to study the variability in the spending amount of each customer, we computed the total amount of spending for the time windows we have defined: $A^M = \{a_{\text{Jun}}, a_{\text{Jul}}, \dots, a_{\text{Mar}}\}$ for the total amount a spent in each month; $A^P = \{a_{1-10}, a_{11-20}, a_{21-31}\}$ for the total amount spent in the early/mid/end part of all months; $A^D = \{a_{\text{Mon}}, a_{\text{Tue}}, \dots, a_{\text{Sun}}\}$ for the total amount spent in a particular day in the dataset (e.g. how much was spent each Monday/Tuesday/etc.).

To calculate the variability of spending amount, we computed the standard deviation of the spending distribution for each customer. Each element of the spending distribution is computed as $\frac{A_i^t}{\sum_i A_i^t}$ and represents the fraction of amount spent in a particular period depending on the aggregation window t (e.g. the fraction spent in June, July, etc. in the monthly aggregation $t = M$). For each customer, this results in measures of monthly (σ_M), 10-days interval (σ_P) and daily variability (σ_D).

(ii) *Persistence of spending amount.* To evaluate the consistency in the amount a customer spends in a monthly and a weekly observation period $t' = \{M, W\}$, we computed the average cosine similarity coefficients between adjacent time intervals.

For the monthly observation period, we first aggregated spending in 10-days intervals (i.e. 3 elements for each month) and then we computed the fraction of spending in each element.

Finally, we computed the persistence of spending amount as the average of the cosine similarity

$$persistence_M = \frac{\sum_{i=0}^{n-1} \cos(S_i, S_{i+1})}{n}, \quad (1)$$

where S_i represents the vector of the relative amount spent in each 10-days interval in a particular month i , and $n = 10$ represents the number of months we have in the dataset. A value of $persistence_M$ of 0 means that the relative amounts spent are dissimilar between the time intervals, while a value of 1 indicates that the relative amounts are exactly the same across intervals.

Similarly, we computed $persistence_W$ for the weekly observation periods ($n = 43$ weeks) by grouping the spending amounts on a daily basis (i.e. 7 elements for each week).

(iii) *Bursty dynamics in spending patterns.* Bursty dynamics are defined as the heterogeneous property of time series having short-time periods of intense activities alternating with long-time periods of low-frequency activities [38]. They allow us to measure the intensity of spending activities over short periods of time. In order to compute the burstiness of the spending patterns, we first computed the inter-event times as the daily difference between two adjacent transactions. We consider only the transaction date since time of the purchase is not available. The inter-event time is defined as $\tau_i = T_i - T_{i-1}$ where T_i represents the transaction which was conducted at time i . Finally, the burstiness parameter is calculated as:

$$B = \frac{r - 1}{r + 1}, \quad (2)$$

where r is defined as $r = \sigma / \langle \tau \rangle$ with τ the average and σ the standard deviation of the transactions' inter-event times.

We label the burstiness parameter for all the financial transactions B_{tot} . In addition, we also calculate the burstiness parameter of daily purchasing B_{daily} , which reflects how regularly the customer makes a purchase on a daily basis. In this case, the inter-event time is the number of consecutive days that the customer does not spend money.

When the burstiness parameter B is -1 , the purchasing pattern of customers is completely stable. If it is $B = 0$, the spending behavior of the customer is random. Finally, a parameter B of 1 indicates extreme and unpredicted spikes in spending behavior.

G3. Category-related spending behavior This third family of spending metrics is related to the categories of purchases made by each customer. We devise these features to have a sense of the diversity and persistence of the spending categories of an individual over time. Previous studies on social interactions and personality [14] showed that traits like Openness to Experience and Agreeableness are associated with a higher turnover of social contacts over time. Moreover, it was found that the diversity of social contacts and the diversity of visited places is correlated with the Big Five personality traits [13]. Taking inspiration from this body of research, we devise metrics looking at the diversity and the stability of individuals' spending categories over time.

As previously described, the spending transactions of customers were aggregated according to 35 spending categories as classified in Table 1. Since the total amount of transactions can be biased towards high-value categories (e.g. the purchase of a car), we base our metrics on the total number of transactions to measure the frequency of purchasing activities in different categories.

(i) *Number of spending categories.* This metric represents the number of distinct categories N_c in which a customer purchased during the entire period of the dataset.

(ii) *Diversity of spending categories.* We measure the diversity of the purchases made by each customer by looking at the diversity of categories D_{cat} , given by the formula:

$$D_{\text{cat}}(i) = -\frac{\sum_{c=1}^{N_c} p_{ic} \log(p_{ic})}{\log N_c}, \quad (3)$$

where N_c is the number of unique categories of customer i , $p_{ic} = \frac{V_{ic}}{\sum_{c=1}^{N_c} V_{ic}}$ and V_{ic} is the volume of expenses made by the customer i in the category c .

A low value of D_{cat} indicates that the customer expenses were mostly made in a few categories. On the other hand, a high value of D_{cat} means that a customer equally distributed his/her expenses in all the categories in which they purchase.

(iii) *Persistence of spending categories.* This metric measures the consistency in customers' purchasing categories over time. It is calculated as the average cosine similarity coefficient between every two adjacent months.

We compute the persistence of purchasing categories as the average of the cosine similarity

$$C_{\text{persistence}} = \frac{\sum_{i=0}^n \cos(D_i, D_{i+1})}{n}, \quad (4)$$

where D_i represents the vector of the relative number of transactions made in each category in a particular month i , and $n = 10$ represents the number of months we have in the dataset.

(iv) *Category turnover.* In order to evaluate a customer's consistency in spending over time, we calculated the turnover in spending categories as the average Jaccard similarity of spending categories in two consecutive months. Let C_i be a set of purchasing categories in the i th month.

$$C_{\text{turnover}} = \frac{\sum_i^{n-1} \frac{C_i \cap C_{i+1}}{C_i \cup C_{i+1}}}{n}. \quad (5)$$

Table 2 *Features summary.* Summary of all the 54 features defined at customer level

Type	Feature	Description
Overall	n_{tot}	Number of transactions (log)
Overall	a_{tot}	Total amount of transactions (log)
Overall	a_{avg}	Average amount per transaction (log)
Overall	cv	Relative spending variability
Temporal	σ_M	Spending amount variability (monthly)
Temporal	σ_P	Spending amount variability (10-days interval)
Temporal	σ_D	Spending amount variability (daily)
Temporal	$persistence_M$	Persistence of spending amount (month)
Temporal	$persistence_W$	Persistence of spending amount (week)
Temporal	B_{tot}	Transactional burstiness
Temporal	B_{daily}	Daily burstiness
Category related	N_c	Number of unique categories
Category related	D_{cat}	Diversity of purchased categories
Category related	$C_{persistence}$	Persistence of purchased categories
Category related	$C_{turnover}^3$	Top 3 category turnover
Category related	$C_{turnover}^5$	Top 5 category turnover
Category related	$C_{turnover}^{all}$	Category turnover
Category profile	C_k	Fraction of expenses (number) in a category k (Table 1)
Socio-demographic	YOB	Year of birth of the customer
Socio-demographic	$salary$	The salary range of the customer (10K intervals)

$C_{turnover}$ is 0 when there is no overlap in the spending categories in two consecutive intervals and it is equal to 1 when the spending categories overlap perfectly.

We calculate the category similarity between the top-3 ($C_{turnover}^3$), top-5 ($C_{turnover}^5$), and all purchasing categories ($C_{turnover}^{all}$) between adjacent months.

G4. Spending category profile The spending category profile reflects the relative number of transactions C_k made in each of the 35 spending categories k as defined in Table 1.

G5. Socio-demographic information In addition to the spending-related features described in G1-G4, we used the socio-demographic information on participants' age (YOB) and salary range. Given the large proportion of missing values for the customer's gender ($\sim 40\%$) we omitted this variable in our analyses.

A summary of all the features is displayed in Table 2.

2.4 Inferring individual traits from spending behavior

To analyze the data, we used each of the different features generated from the spending behavior defined in Sect. 2.3 to infer the individual psychological traits of customers. Specifically, we first investigated the associations between the behavioral features and the individual psychological characteristics by using Pearson correlations (see Sect. 3.1), and then we trained machine learning models to classify the customers' individual traits and evaluate the accuracy with which we are able to infer individual characteristics from customer spending behavior (see Sect. 3.2).

We devised this task as a three-class classification problem. Based on the individual personality characteristics, we assigned each customer to the classes *low*, *average*, or *high* based on the value of each trait, following the percentile-based categorization method

proposed in [39]. Therefore, for a particular trait, customers with scores higher than the 66th percentile are labeled as *high*, customers with scores lower than the 33rd percentile are labeled as *low* and customers falling in between these percentiles are labeled as *average*. For each trait this procedure results in an equal number of participants in each of the three classes. We have evaluated the results obtained from three different machine learning algorithms: Logistic Regression [27], Random Forest [28], and Extreme Gradient Boosting (XGBoost) [29]. For each method, we have randomly divided the dataset into 80% training set and 20% test set, retaining the classes ratio in both training and test sets.

In the training phase, for each model, the parameters are tuned using grid search with 5-fold cross-validation. In order to lower the risk of overfitting given our sample size, we subsequently reduced the dimensionality of the feature space with a feature selection step, using the Recursive Feature Elimination with Cross-Validation (RFECV) method [40]. Finally, we tested the models against the 20% test set (holdout set) reporting the Accuracy, F1 score, and Area Under the Receiver Operating Characteristic Curve (AUROC). We first measured the F1 score and AUROC separately for each class (one-vs-rest) and subsequently calculated the unweighted average (macro average). In order to get a more robust evaluation, we repeated this process 10 times, randomly selecting new train and test sets and averaging the scores of the evaluation metrics.

3 Results

In the following sections, we first describe the results of the correlation analysis, then we present the accuracy of our models in classifying the psychological characteristics of customers from their spending behavior. Finally, we analyze in detail the performances of the different families of behavioral features.

3.1 Correlation analysis

To provide a comprehensive analysis of how spending behavior is associated with individual psychological characteristics, we report observations from the correlation analysis, structuring the discussion around the individual psychological characteristics. For all the analysis, we used the Pearson correlation coefficient.

3.1.1 Overall, temporal and category-related features vs individual traits

Extraversion Extraversion was found to be positively correlated to B_{tot} indicating that more extroverted people tend to have a more bursty spending behavior. More extroverted people tend also to have a higher number of transactions (n_{tot}) with respect to their counterparts; moreover, we found a positive correlation with category similarity over time between the top-3 spending categories (C_{turnover}^3).

Agreeableness We did not find significant correlations between this trait and the features we devise.

Conscientiousness Conscientiousness was found to be significantly and positively correlated with the total amount spent (a_{tot}) and the average amount per transaction (a_{avg}). We also found that the relative amounts spent over different weeks are more dissimilar (persistence_W) for people that display higher scores of Conscientiousness.

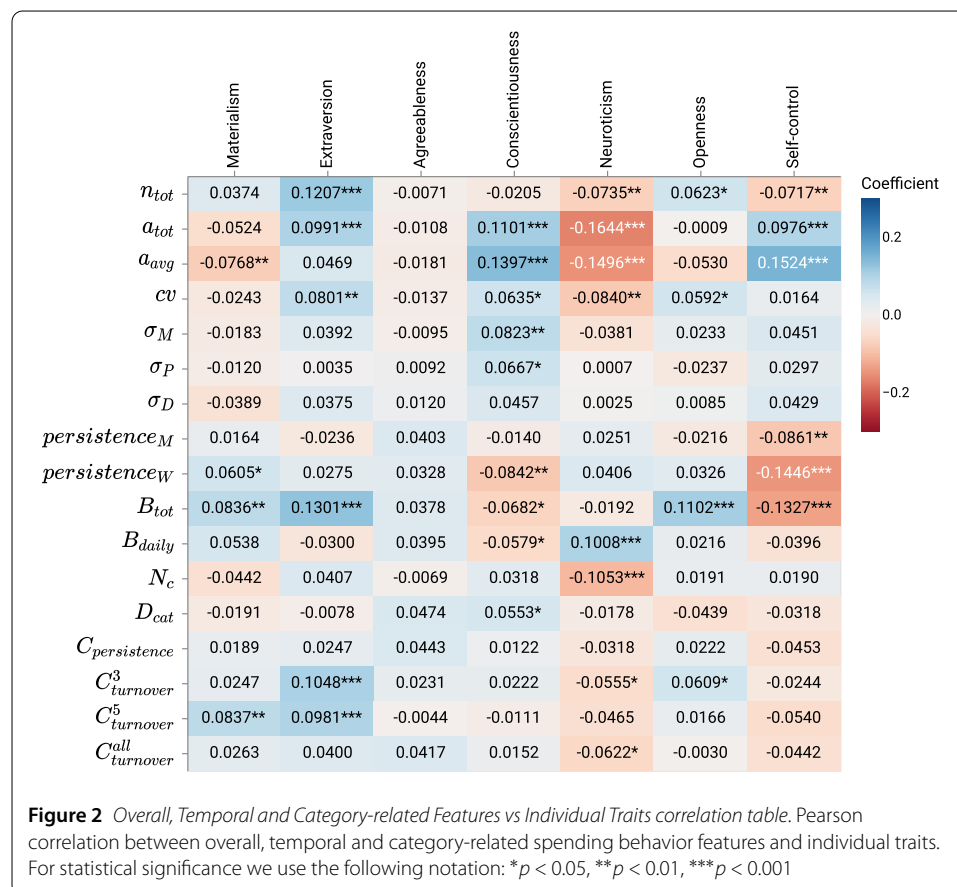
Neuroticism More neurotic individuals displayed lower values in total amount spent (a_{tot}) and in average amount per transaction (a_{avg}), and a smaller number of spending categories (N_c). Additionally, we found a positive significant correlation with burstiness of daily purchasing (B_{daily}), with more neurotic people having a more bursty behavior with respect to their counterparts.

Openness to Experience Openness to Experience is positively correlated with B_{tot} and with n_{tot} , with people more open to new experiences showing a higher bursty spending behavior and having a higher number of transactions with respect to their counterparts.

Materialism Materialism was found to be slightly positively correlated to B_{tot} and category similarity over time in the top-5 spending categories ($C_{turnover}^5$). Moreover, we found a slightly negative correlation with the average amount spent per transaction (a_{avg}).

Self-Control People with higher scores in Self-Control were more likely to have a higher average amount per transaction (a_{avg}), showing instead a slightly lower bursty spending behavior (B_{tot}) and more dissimilar relative amounts spent over different weeks ($persistence_W$).

See Fig. 2 for the complete correlation table.



3.1.2 Category profile features vs individual traits

Extraversion Extraversion exhibits a positive correlation with the categories *Food, drink and going out* and *Transportation*, while a negative correlation is present with the *Groceries and supermarkets* category.

Agreeableness More agreeable individuals tend to slightly spend more in the *Charities* category. Moreover, a negative correlation with the category *Food, drink and going out* was found.

Conscientiousness People with high scores in Conscientiousness tend to spend more in the *Health care* category, while spending less in the *Games and gaming* category.

Neuroticism Neuroticism was found to be positively correlated to the *Personal care and beauty* category. A negative correlation was instead found with the category *Do It Yourself (DIY) projects*.

Openness to Experience This trait is negatively correlated to the *Household: spending* category and positively correlated with the *Alcohol* category.

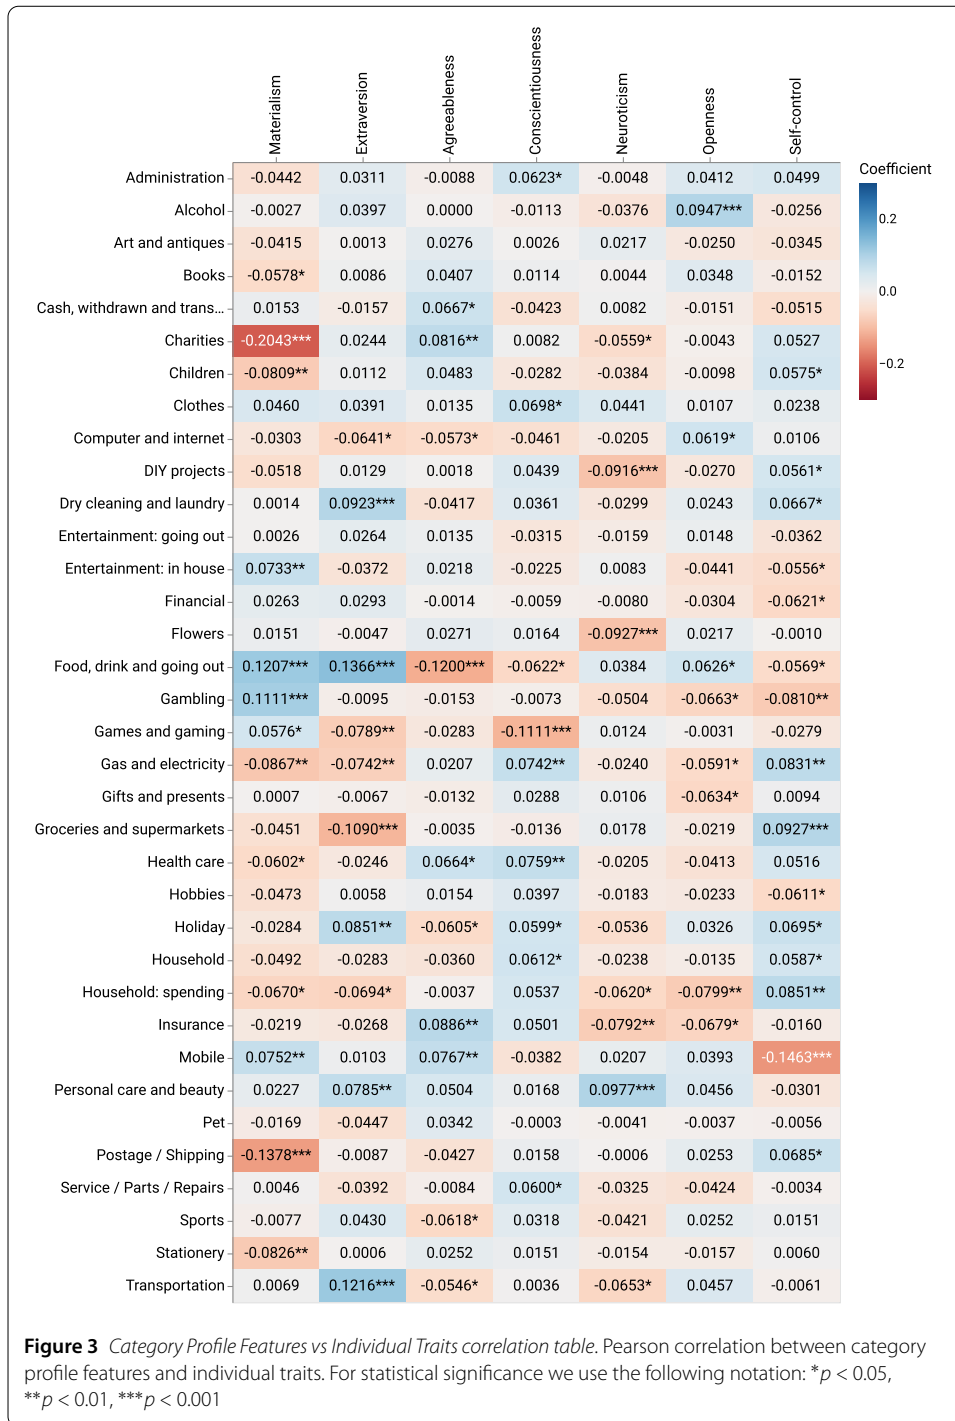
Materialism Individuals with higher scores in the Materialism trait tend to spend less in the category *Charities* with respect to their counterparts and tend to spend less in the *Postage/Shipping* category. A positive correlation is instead present for the *Food, drink and going out* and *Gambling* categories.

Self-Control Self-control was found negatively correlated with the *Mobile* category, and positively correlated with *Groceries and supermarkets* and *Gas and electricity* categories.

See Fig. 3 for the complete correlation table.

3.1.3 Individual traits

To make the analysis complete we also show the correlation matrix between the individual psychological characteristics under study (see Fig. 4). Here, we can see that Agreeableness, Conscientiousness, and Self-Control are negatively associated with Materialism, while there is a slightly positive correlation between Extraversion and Materialism. Extraversion is also positively associated with Openness to Experience and negatively with Neuroticism and Self-Control. Agreeableness instead shows a positive correlation with Conscientiousness, Openness to Experience, and Self-Control, and a negative one with Neuroticism. We can also see a negative correlation between Conscientiousness and Neuroticism and a positive association of Conscientiousness with Self-Control and a slightly positive with Openness to Experience. Finally, neurotic people tend to have lower levels of Self-Control and tend to be less open to new experiences. It is worth highlighting that although Big Five personality traits are theoretically conceptualized as orthogonal, several empirical studies have shown weak to moderate correlations among the personality traits (see Van der Linden *et al.* [41] for a meta-analysis of these studies). Moreover, the correlations between personality traits found in our work are similar to the ones reported by previous ones [41]. This is also true for the correlations between the Big Five traits and Materialism [42], and the correlations between the Big Five traits and Self-Control [43].



3.2 Classification models' performance

Table 3 displays the performance of the Logistic Regression, Random Forest and XGBoost models. As we can see from this table, the highest accuracies were obtained for Materialism when using a Random Forest classifier (F1 = 0.420, AUROC = 0.588), and Self-Control when using a Logistic Regression classifier (F1 = 0.407, AUROC = 0.585). The performance of the machine learning models is lower when classifying the Big Five personality traits. Here, the highest accuracies were obtained with Extraversion when using

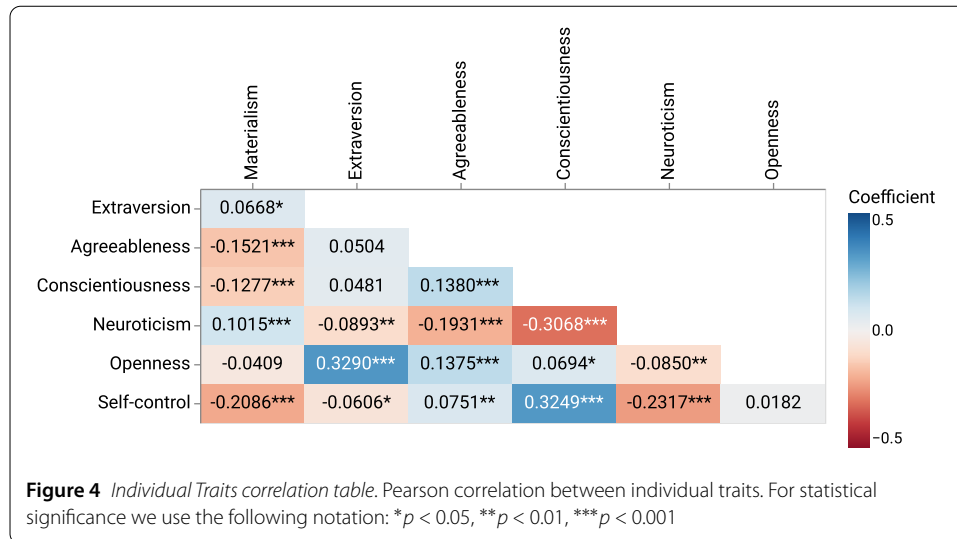


Table 3 Classification models' performance. Machine learning models performance (LR = Logistic regression, RF = Random Forest, XGB = XGBoost.) evaluated with the Accuracy, F1 score, Precision, Recall and Area Under the Receiver Operating Characteristic Curve (AUROC)

Individual traits	Model	Accuracy		Precision		Recall		F1 measure		AUROC	
		Mean	Std	Mean	Std	Mean	Std	Mean	Std	Mean	Std
Materialism	LR	0.377	0.032	0.375	0.034	0.378	0.032	0.371	0.033	0.544	0.020
	RF	0.423	0.026	0.420	0.028	0.423	0.026	0.420	0.026	0.588	0.019
	XGB	0.409	0.027	0.406	0.027	0.409	0.027	0.406	0.028	0.580	0.016
Extraversion	LR	0.385	0.029	0.362	0.050	0.386	0.029	0.362	0.039	0.567	0.027
	RF	0.386	0.026	0.384	0.025	0.387	0.026	0.383	0.025	0.574	0.020
	XGB	0.398	0.033	0.397	0.032	0.398	0.033	0.396	0.032	0.573	0.029
Agreeableness	LR	0.345	0.026	0.333	0.045	0.345	0.026	0.321	0.035	0.508	0.023
	RF	0.377	0.019	0.375	0.020	0.377	0.019	0.374	0.020	0.547	0.018
	XGB	0.363	0.028	0.363	0.030	0.364	0.028	0.360	0.029	0.540	0.026
Conscientiousness	LR	0.361	0.023	0.361	0.026	0.362	0.023	0.360	0.025	0.533	0.010
	RF	0.361	0.036	0.362	0.037	0.361	0.036	0.360	0.036	0.529	0.027
	XGB	0.345	0.022	0.345	0.021	0.345	0.022	0.344	0.021	0.512	0.018
Neuroticism	LR	0.392	0.026	0.393	0.027	0.392	0.026	0.386	0.026	0.570	0.029
	RF	0.401	0.018	0.402	0.019	0.401	0.018	0.399	0.019	0.558	0.015
	XGB	0.377	0.015	0.376	0.016	0.377	0.015	0.374	0.014	0.553	0.013
Openness to experience	LR	0.344	0.014	0.338	0.049	0.344	0.013	0.322	0.036	0.512	0.018
	RF	0.348	0.034	0.346	0.034	0.348	0.034	0.345	0.034	0.520	0.024
	XGB	0.331	0.026	0.329	0.025	0.331	0.026	0.324	0.028	0.508	0.013
Self-control	LR	0.410	0.026	0.408	0.026	0.411	0.026	0.407	0.026	0.585	0.023
	RF	0.394	0.026	0.392	0.027	0.395	0.026	0.391	0.026	0.575	0.029
	XGB	0.379	0.034	0.377	0.035	0.379	0.034	0.376	0.035	0.564	0.033

XGB (F1 = 0.396, AUROC = 0.573) and Neuroticism when using Random Forest (F1 = 0.399, AUROC = 0.558). As explained in Sect. 2.4, the task has an equal number of samples in the three classes. We compare our results against a baseline classifier that always predicts one of the classes (accuracy of 0.333). That means, that the predictive accuracy of 0.423 for the Materialism yields a 27% improvement over the baseline. Contrary to the findings for Materialism, Self-Control, Extraversion and Neuroticism, the models did not substantially improve performance for Agreeableness, Conscientiousness, and Openness

Table 4 *Feature groups' performances.* Comparison of the performances (F1 scores) of different feature groups using a Random Forest model

Individual trait	Feature group	F1 score
Materialism	Sociodemographic	0.352
	Overall	0.349
	Temporal	0.365
	Category-related	0.337
	Category profile	0.419
	All features	0.420
Extraversion	Sociodemographic	0.340
	Overall	0.339
	Temporal	0.359
	Category-related	0.362
	Category profile	0.384
	All features	0.383
Neuroticism	Sociodemographic	0.388
	Overall	0.356
	Temporal	0.345
	Category-related	0.354
	Category profile	0.361
	All features	0.399
Self-control	Sociodemographic	0.319
	Overall	0.362
	Temporal	0.362
	Category-related	0.370
	Category profile	0.368
	All features	0.391

to Experience. Given the poor performance of the models in inferring these traits, we did not include them in the subsequent analyses.

3.2.1 Performance of feature groups

To develop a broader understanding of the performances of the five feature groups, using the same settings as described in Sect. 2.4, we trained Random Forest models for each feature group and compared their performance. The results are presented in Table 4. We can see that for traits like Materialism and Extraversion the performances of the feature group *Category profile* is the one that performs better. Instead, the performances of the five feature groups are more comparable for the Neuroticism and Self-Control traits.

3.2.2 Impact of overall, temporal and category-related features group

To further understand whether the novel behavioral features help in inferring the individual psychological characteristics under study, we first trained Random Forest models using only the Socio-demographic and the Category profile features, similarly to the approach described in [26]. We subsequently compared, with the same settings as described in Sect. 2.4, the results of the models obtained from our complete set of features. Using the (i) *Overall* features, that measure the overall spending characteristics, (ii) the *Temporal* features, that model the variability, the persistence and the regularity of an individual's spending behavior, and (iii) the *Category related* features, which look at the persistence and turnover in the categories of the expenses, we find a significant but modest improvement over the Socio-demographic and the Category profile models for two traits: Self-Control for which we observe a +9.9% improvement in F1 measure, and Neuroticism for




Figure 5 Feature importance. Top 10 features for the Materialism and Self-Control traits

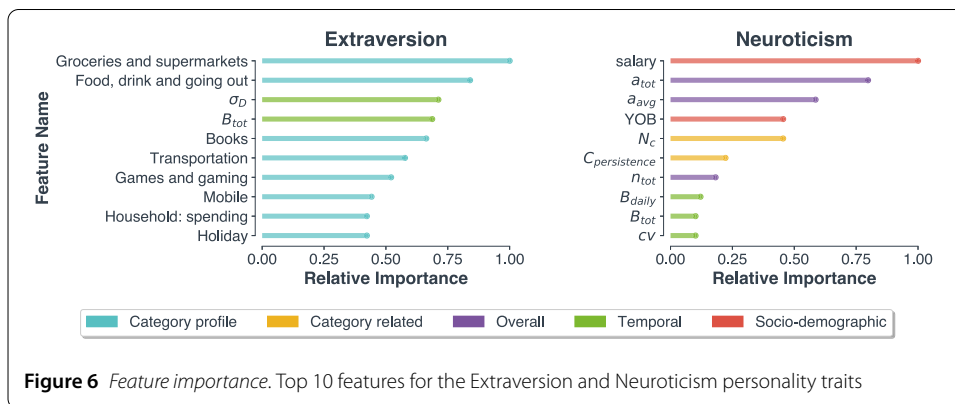


Figure 6 Feature importance. Top 10 features for the Extraversion and Neuroticism personality traits

which we observe a +4.7% improvement in F1 measure. This finding, as we will see in the next section, is also reflected in Fig. 5 and Fig. 6, which shows several temporal and category-related features among the top 10 most important for these two traits.

3.3 Feature importance

We were also interested in understanding which of the features we used in our models have the highest impact in inferring a given psychological trait. To do so, we computed the feature importance of the top 10 most predictive features using the permutation importance method [44]. To further discern the relationship between these features and a given personality trait, we also investigated the features' directionality computing the Spearman correlation of the top 10 most predictive features for Materialism, Self-Control, Extraversion and Neuroticism (see Table 5).

The first two sections of Fig. 5 show the top 10 most predictive features for Materialism and Self-Control. The feature with the highest predictive strength for Materialism (Fig. 5 left) is the proportion of spending on the *Charities* category. Taking a closer look at the feature directionality, we observe a negative association between charitable giving and Materialism. This means materialistic individuals are less likely to donate to charities. Other important features are represented by the year of birth (*YOBB*), with younger individuals showing higher scores in Materialism, and people with a larger fraction of expenses in the *Gambling* category displaying higher scores in Materialism.

For Self-Control (Fig. 5 right), we observe that individuals with higher scores in Self-Control spend a higher average amount per transaction (a_{avg}), register a smaller number of

Table 5 *Feature directionality*. Feature directionality of the top 10 most predictive features computed using the Spearman correlation coefficient

Trait	Feature	Coefficient	<i>p</i> -value
Materialism	<i>Charities</i>	-0.204	<0.001
	<i>YOB</i>	+0.207	<0.001
	<i>Postage/Shipping</i>	-0.138	<0.001
	<i>Gambling</i>	+0.111	<0.001
	<i>Transportation</i>	+0.007	0.804
	<i>Food, drink and going out</i>	+0.121	<0.001
	<i>Entertainment: in house</i>	+0.073	0.008
	<i>Groceries and supermarkets</i>	-0.045	0.104
	<i>Mobile</i>	+0.075	0.007
	<i>D_{cat}</i>	-0.019	0.491
Self-control	<i>a_{avg}</i>	+0.152	<0.001
	<i>n_{tot}</i>	-0.072	0.010
	<i>Groceries and supermarkets</i>	+0.093	0.001
	<i>cv</i>	+0.016	0.553
	<i>B_{daily}</i>	-0.040	0.153
	<i>B_{tot}</i>	-0.133	<0.001
	<i>N_c</i>	+0.019	0.493
	<i>C_{persistence}</i>	-0.045	0.102
	<i>Cash, withdrawn and transfer</i>	-0.052	0.063
	<i>C_{turnover}^S</i>	-0.054	0.051
Extraversion	<i>Groceries and supermarkets</i>	-0.109	<0.001
	<i>Food, drink and going out</i>	+0.137	<0.001
	<i>σ_D</i>	+0.038	0.176
	<i>B_{tot}</i>	+0.130	<0.001
	<i>Books</i>	+0.009	0.757
	<i>Transportation</i>	+0.122	<0.001
	<i>Games and gaming</i>	-0.079	0.004
	<i>Mobile</i>	+0.010	0.709
	<i>Household: spending</i>	-0.069	0.012
	<i>Holiday</i>	+0.085	0.002
Neuroticism	<i>salary</i>	-0.150	<0.001
	<i>a_{tot}</i>	-0.164	<0.001
	<i>a_{avg}</i>	-0.150	<0.001
	<i>YOB</i>	+0.111	<0.001
	<i>N_c</i>	-0.105	<0.001
	<i>C_{persistence}</i>	-0.032	0.252
	<i>n_{tot}</i>	-0.073	0.008
	<i>B_{daily}</i>	+0.101	<0.001
	<i>B_{tot}</i>	-0.019	0.489
	<i>cv</i>	-0.084	0.002

total transactions (n_{tot}) and exhibit spending behavior that follows a more regular pattern (indicated by lower values in the transactional burstiness feature (B_{tot})).

Figure 6 shows the top 10 most predictive features for Extraversion and Neuroticism. For Extraversion (Fig. 6, left), we observe that people who are more extroverted tend to spend more money in the category *Food, drink and going out*, exhibit spending behavior that is less regular in the entire observation period (indicated by higher values in the transactional burstiness feature (B_{tot})), and make more purchases in the *Transportation* category.

Individuals that display higher scores on Neuroticism (Fig. 6 right), report a lower *salary* range, and spend less money overall (indicated by lower values in the features a_{tot} and a_{avg}), show less persistence in their spending behavior (indicated by the persistence of spending categories ($C_{persistence}$)).

4 Discussion

Using the transaction records of 1306 bank customers, we investigated the extent to which individual-level psychological characteristics can be inferred from bank transaction data. Expanding previous research [26], we developed a comprehensive set of behavioral features that capture differences in spending behavior along five dimensions: (1) the overall spending behavior, (2), the temporal spending behavior (i.e. variability, persistence, and burstiness), (3) the category-related spending behavior (i.e. diversity, persistence, and turnover), (4) the customer category profile, and (5) the socio-demographic information.

Our results show that inferring the psychological traits of an individual is a challenging task, even when using a comprehensive set of features that take temporal aspects of spending into account. They also align with previous research suggesting that there are stark differences in the predictive accuracy across the different traits. Similar to the findings of Gladstone *et al.* [26], we found that Materialism and Self-Control could be inferred with relatively higher levels of accuracy, while the accuracy obtained for the Big Five traits was found to be lower, with only Extraversion and Neuroticism reaching classification performances that were significantly different than chance.

Across the different traits, the predictive accuracies we obtained from spending behavior are lower than those obtained from other digital footprints such as Facebook Likes [3, 45], Facebook status updates [46, 47] or mobile phone data [11, 13, 48]. As also hypothesized by [26] this might be due to the nature of spending records. Compared to social media data which constitute an explicit form of identity claim [49], spending behaviors constitute a more implicit form of behavioral residue that might reveal less information about a person's inner psychological states. However, this result is of paramount relevance for challenging and warning researchers and practitioners working on the design of automatic systems and algorithms for inferring individual psychological characteristics from spending behaviors.

Moreover, despite the relatively poor performance of the predictive models, the strongest features observed in the feature importance analyses have good face validity. The relationship between Materialism and lower rates of charitable giving aligns with previous literature that conceptualizes non-generosity as a central aspect of Materialism [50] and that finds that materialistic people are less likely to donate and to act pro-socially [51]. Similarly, the link between Extraversion and spending money on the category *Food, drink and going out* is not only in line with the findings by Matz *et al.* [24] in a different sample of customers, but also corresponds to the general characterization of extraverts being more social. In addition, the fact that extraverts have less regular spending patterns aligns with previous findings which suggest that those with more extraverted personalities are more impulsive; they are social butterflies who live in the here-and-now. In contrast, the relationship between Self-Control and regular patterns of spending also reflects the fact that people high in Self-Control are typically less impulsive than people scoring low in Self-Control and more likely to plan ahead and follow routines. The link between Neuroticism and lower levels of persistence speaks to the fact that neurotic people are less emotionally stable and might therefore change their choices more often, and is consistent with previous findings linking Neuroticism to higher irregularities in their phone call logs [13].

Limitations and practical implications Our study has a number of limitations. First, our models are based on a relatively small sample, which might not be representative of the

general population. In addition, the measures used to assess the Big Five personality traits, Materialism and Self-Control, although being validated in previous studies, are shorter than what is recommended in the psychological assessment literature. Given that the accuracy of a predictive model is limited by the extent to which the original measure is reliable, this might partly explain why the accuracy is substantially lower than in previous studies using different digital footprints and much longer measures. Finally, one of the inherent limitations of spending behavior is that it can be influenced by the financial constraints of the person spending, as well as by purchases made for other family members. As such, spending might not always be reflective of an individual's personal preferences.

Taken together, our findings contribute to the body of research studies on the automatic recognition of psychological traits from digital footprints. Although we were able to improve on the accuracy of classification models from spending behavior only for some traits (i.e. Self-Control and Neuroticism), we hope the additional variables calculated for routines and temporal sequences will inspire future researchers to investigate and calculate similar variables when training models that have time-stamped data available. Given the decent accuracies in classifications, these results could also help to improve psychologically-informed advertising strategies for specific products [52] as well as personality-based spending management apps and credit scoring approaches [53]. These approaches are likely to be more successful for Materialism, Self-Control, Extraversion, and Neuroticism, given the relatively stronger accuracies we find in inferring these traits, while require caution for Agreeableness, Conscientiousness and Openness to Experience.

Privacy and ethical recommendations Similar to the prediction of personal information from other digital traces such as social media profiles, smartphone sensors or browsing histories, inferring personality traits from spending data raises important ethical questions related to privacy and data protection. In most cases, individuals will not expect their spending data to be used for the prediction of psychological characteristics. According to the theory of *contextual integrity*, this use of data in a way that could not be realistically foreseen and expected by the person who initially provided the data constitutes a violation of privacy, even if the individual initially consented to their data being collected [54, 55]. Hence, it is critical to make sure that individuals know and understand how their data is being used. This call for transparency is a pillar of the European Union's General Data Protection Regulation (GDPR) [56] and the California Consumer Protection Act (CCPA) [57], that require companies to state in a clear and easy-to-understand manner what data is being collected and how this data is being used and/or shared with third parties. While such regulatory calls for transparency are critical they are often slow and place a considerable burden on the consumer, because regulations such as the GDPR and the CCPA assume that informed consumers will be able to make rational decisions related to their privacy. However, there is ample evidence that this is not the case [58]. The data and privacy landscape is so complex, that even motivated consumers will find it difficult to accrue and maintain the knowledge and expertise required to make self-interested decisions that trade-off immediate, tangible convenience benefits of sharing data in the now with potential, abstract privacy costs in the future. Hence, privacy regulation should be complemented by technological solutions, such as privacy by design (e.g. the integration of privacy protection mechanisms into the design of psychometric-based systems [59]), federated learning (i.e. training on local devices of the consumers [60]) and encrypted

computation (i.e. training and evaluating machine learning algorithms on encrypted data [61]), that provide privacy protection without placing the burden on consumers.

In addition to protecting individuals' privacy, it will also become necessary to outline contexts in which predictions of psychological traits from credit card data and the application of the resulting profiles should be prohibited. This requires a public debate that is informed by our moral values and a discussion on the extent to which individuals should be able to act as self-determined agents. We might agree that using such predictions in the context of product recommendations are acceptable (or even desirable) as long as the individual is sufficiently protected and has the agency to make an informed decision of whether they want to make use of this option or not. However, we might decide that such predictions cannot be made and used in the context of political campaigning because the risk for abuse outweighs the potential benefit some consumers might derive from it. Because this is a normative and complex debate, it will require collaboration between the public, industry leaders, academia, legal experts and policy makers [62, 63].

Acknowledgements

We thank Gianni Barlacchi for insightful discussions and helpful comments.

Funding

The work of NT was partly supported by an Erasmus Mundus scholarship for his master's programme. The work of SC and BL was supported by the H2020 INFINITECH project, grant agreement number 856632.

Abbreviations

UK, United Kingdom; GBP, Great Britain Pound; LSOA, Lower Layer Super Output Area; MSOA, Middle Layer Super Output Area; BFI-10, Big Five Inventory—10; YOB, Year of Birth; XGBoost, Extreme Gradient Boosting; RFECV, Recursive Feature Elimination with Cross-Validation; AUROC, Area Under the Receiver Operating Characteristic Curve; GDPR, General Data Protection Regulation; CCPA, California Consumer Protection Act.

Availability of data and materials

The data used in this paper are proprietary and are not available for public use. The research partner placed restrictions on data sharing as a condition of the collaboration agreement. Further information about the data and access is available from JG.

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

Conceived the study: BL, SM. Designed and performed the experiments: NT, SC, BL. Wrote the paper: NT, SC, SM, JG, BL. All authors read, reviewed and approved the final manuscript.

Author details

¹CentraleSupélec, Université Paris-Saclay, 3 Rue Joliot Curie, 91190 Gif-sur-Yvette, France. ²Fondazione Bruno Kessler, Via Sommarive, 18, 38123 Trento, Italy. ³School of Management, University College of London, One Canada Square, Canary Wharf, E14 5AA London, UK. ⁴Columbia Business School, 3022 Broadway, 10027 New York, US.

Publisher's Note

Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Received: 22 November 2019 Accepted: 3 May 2021 Published online: 08 May 2021

References

1. Lazer D, Pentland A, Adamic L, Aral S, Barabási A-L, Brewer D, Christakis N, Contractor N, Fowler J, Gutmann M, Jebara T, King G, Macy M, Roy D, Van Alstyne M (2009) Computational social science. *Science* 323(5915):721–723
2. Azucar D, Marengo D, Settanni M (2018) Predicting the big 5 personality traits from digital footprints on social media: a meta-analysis. *Pers Individ Differ* 124:150–159
3. Kosinski M, Stillwell D, Graepel T (2013) Private traits and attributes are predictable from digital records of human behavior. *Proc Natl Acad Sci USA* 110(15):5802–5805
4. Youyou W, Kosinski M, Stillwell D (2015) Computer-based personality judgements are more accurate than those made by humans. *Proc Natl Acad Sci USA* 112(4):1–5
5. Golbeck J, Robles C, Edmondson M, Turner K (2011) Predicting personality from Twitter. In: Proceedings of the third international conference on social computing (SocialCom), pp 149–156

6. Quercia D, Kosinski M, Stillwell D, Crowcroft J (2011) Our Twitter profiles, our selves: predicting personality with Twitter. In: Proceedings of the third international conference on social computing (SocialCom), pp 180–185
7. Park G, Schwartz HA, Eichstaedt JC, Kern ML, Kosinski M, Stillwell DJ, Ungar LH, Seligman ME (2015) Automatic personality assessment through social media language. *J Pers Soc Psychol* 108(6):934–952
8. Segalin C, Cristani M, Perina A, Vinciarelli A (2017) The pictures we like are our image: continuous mapping of favorite pictures into self-assessed and attributed personality traits. *IEEE Trans Affect Comput* 8(2):268–285
9. Ferwerda B, Tkalcic M (2018) Predicting users' personality from Instagram pictures: using visual and/or content features? In: Proceedings of the 26th conference on user modeling, adaptation and personalization, pp 157–161
10. Rentfrow PJ, Gosling SD (2003) The do re mi's of everyday life: the structure and personality correlates of music preferences. *J Pers Soc Psychol* 84(6):1236
11. Staiano J, Lepri B, Aharony N, Pianesi F, Sebe N, Pentland A (2012) Friends don't lie: inferring personality traits from social network structure. In: Proceedings of the 2012 ACM conference on ubiquitous computing, pp 321–330
12. Chittaranjan G, Blom J, Gatica-Perez D (2013) Mining large-scale smartphone data for personality studies. *Pers Ubiquitous Comput* 17(3):433–450
13. de Montjoye Y-A, Quoidbach J, Robic F, Pentland A (2013) Predicting personality using novel mobile phone-based metrics. In: Proceedings of the international conference on social computing, behavioral-cultural modeling, and prediction, pp 48–55
14. Centellegher S, Lopez E, Saramaki J, Lepri B (2017) Personality traits and ego-network dynamics. *PLoS ONE* 12(3):0173110
15. Demircuc-Kunt A, Klapper L, Singer D, Ansar S, Hess J (2018) The global Findex database 2017: measuring financial inclusion and the Fintech revolution. The World Bank
16. Lenormand M, Louail T, Cantú-Ros OG, Picornell M, Herranz R, Arias JM, Barthelemy M, San Miguel M, Ramasco JJ (2015) Influence of sociodemographic characteristics on human mobility. *Sci Rep* 5:10075
17. Di Clemente R, Luengo-Oroz M, Travizano M, Xu S, Vaitla B, González MC (2018) Sequences of purchases in credit card data reveal lifestyles in urban populations. *Nat Commun* 9(1):3330
18. Dong X, Suhara Y, Bozkaya B, Singh VK, Lepri B, Pentland A (2018) Social bridges in urban purchase behavior. *ACM Trans Intell Syst Technol* 9(3):33
19. Levy SJ (1999) Symbols for sale. In: Brands, consumers, symbols and research, pp 203–212
20. Aaker JL (1997) Dimensions of brand personality. *J Mark Res* 34(3):347–356
21. Aaker JL (1999) The malleable self: the role of self-expression in persuasion. *J Mark Res* 36(1):45–57
22. Govers PC, Schoormans JP (2005) Product personality and its influence on consumer preference. *J Consum Mark* 22(4):189–197
23. Kressmann F, Sirgy MJ, Herrmann A, Huber F, Huber S, Lee D-J (2006) Direct and indirect effects of self-image congruence on brand loyalty. *J Bus Res* 59(9):955–964
24. Matz SC, Gladstone JJ, Stillwell D (2016) Money buys happiness when spending fits our personality. *Psychol Sci* 27(5):715–725
25. Landis B, Gladstone J (2017) Personality, income, and compensatory consumption: low-income extraverts spend more on status. *Psychol Sci* 28(10):1–3
26. Gladstone JJ, Matz SC, Lemaire A (2019) Can psychological traits be inferred from spending? Evidence from transaction data. *Psychol Sci* 30(7):1087–1096
27. Hilbe JM (2009) Logistic regression models. Chapman & Hall, London
28. Breiman L (2001) Random forests. *Mach Learn* 45(1):5–32
29. Chen T, Guestrin C (2016) Xgboost: a scalable tree boosting system. In: Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining, pp 785–794
30. Costa PT, McCrae RR (1992) Revised NEO personality inventory (NEO-PI-R) and NEO five-factor inventory (NEO-FFI) manual, Odessa, FL: Psychological Assessment Resources
31. Rammstedt B, John OP (2007) Measuring personality in one minute or less: a 10-item short version of the big five inventory in English and German. *J Res Pers* 41(1):203–212
32. Donnelly G, Iyer R, Howell RT (2012) The big five personality traits, material values, and financial well-being of self-described money managers. *J Econ Psychol* 33(6):1129–1142
33. Oehler A, Wendt S, Wedlich F, Horn M (2018) Investors' personality influences investment decisions: experimental evidence on extraversion and neuroticism. *J Behav Finance* 19(1):30–48
34. Richins ML, Dawson S (1992) A consumer values orientation for materialism and its measurement: scale development and validation. *J Consum Res* 19(3):303–316
35. Tangney JP, Baumeister RF, Boone AL (2004) High self-control predicts good adjustment, less pathology, better grades, and interpersonal success. *J Pers* 72(2):271–324
36. Moore K, McElroy JC (2012) The influence of personality on Facebook usage, wall postings, and regret. *Comput Hum Behav* 28(1):267–274
37. Mehrotra A, Hendley R, Musolesi M (2016) Towards multi-modal anticipatory monitoring of depressive states through the analysis of human-smartphone interaction. In: Proceedings of the 2016 ACM international joint conference on pervasive and ubiquitous computing: adjunct, pp 1132–1138
38. Karsai M, Jo H-H, Kaski K (2018) Bursty human dynamics. Springer, Berlin
39. MacCallum RC, Zhang S, Preacher KJ, Rucker DD (2002) On the practice of dichotomization of quantitative variables. *Psychol Methods* 7(1):19
40. Guyon I, Weston J, Barnhill S, Vapnik V (2012) Gene selection for cancer classification using support vector machines. *Mach Learn* 46:389–422
41. Van der Linden D, te Nijenhuis J, Bakker AB (2010) The general factor of personality: a meta-analysis of big five intercorrelations and a criterion-related validity study. *J Res Pers* 44(3):315–327
42. Otero-López JM, Villardefrancos E (2013) Five-factor model personality traits, materialism, and excessive buying: a mediational analysis. *Pers Individ Differ* 54(6):767–772
43. Olson KR (2005) Engagement and self-control: superordinate dimensions of big five traits. *Pers Individ Differ* 38(7):1689–1700

44. Altmann A, Tolosi L, Sander O, Lengauer T (2010) Permutation importance: a corrected feature importance measure. *Bioinformatics* 26(10):1340–1347
45. Kosinski M, Matz SC, Gosling SD, Popov V, Stillwell D (2015) Facebook as a research tool for the social sciences: opportunities, challenges, ethical considerations, and practical guidelines. *Am Psychol* 70(6):543–556
46. Schwartz HA, Eichstaedt JC, Kern ML, Dziurzynski L, Ramones SM, Agrawal M, Shah A, Kosinski M, Stillwell D, Seligman ME et al (2013) Personality, gender, and age in the language of social media: the open-vocabulary approach. *PLoS ONE* 8(9):73791
47. Park G, Schwartz HA, Eichstaedt JC, Kern ML, Kosinski M, Stillwell DJ, Ungar LH, Seligman ME (2015) Automatic personality assessment through social media language. *J Pers Soc Psychol* 108(6):934
48. Chittaranjan G, Blom J, Gatica-Perez D (2013) Mining large-scale smartphone data for personality studies. *Pers Ubiquitous Comput* 17(3):433–450
49. Gosling SD, Ko SJ, Mannarelli T, Morris ME (2002) A room with a cue: personality judgments based on offices and bedrooms. *J Pers Soc Psychol* 82(3):379
50. Belk RW (1985) Materialism: trait aspects of living in the material world. *J Consum Res* 12(3):265–280
51. Kasser T (2016) Materialistic values and goals. *Annu Rev Psychol* 67:489–514
52. Matz SC, Kosinski M, Nave G, Stillwell DJ (2017) Psychological targeting as an effective approach to digital mass persuasion. *Proc Natl Acad Sci USA* 114(48):12714–12719
53. Rustichini A, DeYoung CG, Anderson JE, Burks SV (2016) Toward the integration of personality theory and decision theory in explaining economic behavior: an experimental investigation. *J Behav Exp Econ* 64:122–137
54. Nissenbaum H (2009) *Privacy in context: technology, policy, and the integrity of social life*. Stanford University Press, Stanford
55. Nissenbaum H (2019) Contextual integrity up and down the data food chain. *Theor Inq Law* 20(1):221–256
56. European Commission (2017) Protection of personal data. <https://gdpr-info.eu/art-4-gdpr>
57. California Consumer Protection Act (2018) <https://oag.ca.gov/privacy/ccpa>
58. Acquisti A, Brandimarte L, Loewenstein G (2010) Privacy and human behavior in the age of the information. *Science* 347(6221):509–514
59. Monreale A, Rinzivillo S, Pratesi F, Giannotti F, Pedreschi D (2014) Privacy-by-design in big data analytics and social mining. *EPJ Data Sci* 3:10
60. Yang Q, Liu Y, Chen T, Tong Y (2019) Federated machine learning: concept and applications. *ACM Trans Intell Syst Technol* 10:12
61. Dowlin N, Gilad-Bachrach R, Lauter K, Naehrig M, Wernsing J (2016) Cryptonets: applying neural networks to encrypted data with high throughput and accuracy. In: *Proceedings of the 33rd international conference on machine learning (ICML2016)*, pp 201–210
62. Matz SC, Appel RE, Kosinski M (2020) Privacy in the age of psychological targeting. *Curr Opin Psychol* 31:116–121
63. Lepri B, Oliver N, Pentland A (2021) Ethical machines: the human-centric use of artificial intelligence. *iScience* 24(3):102249

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- Convenient online submission
- Rigorous peer review
- Open access: articles freely available online
- High visibility within the field
- Retaining the copyright to your article

Submit your next manuscript at ► [springeropen.com](https://www.springeropen.com)
