Tutorial: *a priori* estimation of sample size, effect size, and statistical power for cluster analysis, latent class analysis, and multivariate mixture models

Edwin S Dalmaijer [1]

Affiliation

1. School of Psychological Science, University of Bristol, United Kingdom

Contact Details

Dr Edwin S. Dalmaijer, School of Psychological Science, University of Bristol, 12a Priory Road, Bristol, BS8 1TU, United Kingdom. Email: edwin.dalmaijer@bristol.ac.uk

Conflicts of interest

The author declares that they have no competing interests (financial or other) that could have influenced or appeared to influence the work reported here.

Abstract

Before embarking on data collection, researchers typically compute how many individual observations they should do. This is vital for doing studies with sufficient statistical power, and often a cornerstone in study pre-registrations and grant applications. For traditional statistical tests, one would typically determine an acceptable level of statistical power, (gu)estimate effect size, and then use both values to compute the required sample size. However, for analyses that identify subgroups, statistical power is harder to establish. Once sample size reaches a sufficient threshold, effect size is primarily determined by the number of measured features and the underlying subgroup separation. As a consequence, *a priory* computations of statistical power are notoriously complex. In this tutorial, I will provide a roadmap to determining sample size and effect size for analyses that identify subgroups. First, I introduce a procedure that allows researchers to formalise their expectations about effect sizes in their domain of choice, and use this to compute the minimally required number of measured variables. Next, I outline how to establish the minimum sample size in subgroup analyses. Finally, I use simulations to provide a reference table for the most popular subgroup analyses: k-means, Ward agglomerative hierarchical clustering, c-means fuzzy clustering, latent class analysis, latent profile analysis, and Gaussian mixture modelling. The table shows the minimum numbers of observations per expected subgroup (sample size) and features (measured variables) to achieve acceptable statistical power, and can be readily used in study design.

Introduction

Many interesting research questions relate to identifying distinct subgroups within a larger dataset. For example, subgroup analyses can be used to describe areas of methane emissions (Malone et al., 2022), chemical concentrations in waters (Li et al., 2022), energy consumption in buildings (Díaz Redondo et al., 2020), jobs on high-performance computing systems (Halawa et al., 2020), glacier zones (Barzycka et al., 2023), Mediterranean dust events (Nissenbaum et al., 2023), microbial networks (Favila et al., 2022), and riverbed types (Kwon et al., 2023). In medical research, subgroup analyses have been used to identify risk factors for hospitalisation from clinical notes (Song et al., 2023), ways to provide informational and emotional patient support (Duimel et al., 2022); or to uncover clinically meaningful patient profiles in bipolar personality disorder (Wolf et al., 2023), compulsive behaviour (Den Ouden et al., 2022), low back pain (Wilson et al., 2023), polycystic ovary syndrome (Kiconco et al., 2023), or vasculitis (Lee et al., 2023). In cognitive and behavioural sciences, subgroup analyses are rapidly gaining in popularity compared to traditional statistical tests (Figure 1, using bibliobanana by Dalmaijer, Van Rheede, et al., 2021). They have been used to describe distinct types of dietary patterns (Hennessy et al., 2023), executive function difficulties (Bathelt et al., 2018), individuals with autism (Parlett-Pelleriti et al., 2022), intimate partner violence (Alexander & Johnson, 2023), resilience to socio-economic disadvantage (Dalmaijer, Gibbons, et al., 2021), responses to cognitive training (Rennie et al., 2019), search organisation (Benjamins et al., 2019), and self-harm (Uh et al., 2021).

While the above all seems rather promising, subgroup analyses are not always appropriately applied. A major issue is that many studies fail to test against a one-group solution (Toffalini et al., 2022). In addition, subgroup analyses within the same domain frequently show inconsistent findings (Alexander & Johnson, 2023; Stein & Bomyea, 2023), and this is likely due to under-addressed methodological and statistical limitations.

Some go so far as to suggest subgroup analysis should not be used in psychological research (Toffalini et al., 2022). Others outline the conditions in which subgroup analyses can be appropriate (Dalmaijer et al., 2020, 2022), and provide detailed guidelines on recommended workflows and common issues (Gao et al., 2023).

Many of the highlighted problems come from researchers' lack of familiarity with the unintuitive statistical properties of subgroup analyses. Like traditional tests, their statistical power is determined by sample size and effect size (Dalmaijer et al., 2020, 2022). However, unlike traditional tests, effect size (and thus power) accumulates over measured variables (Dalmaijer et al., 2020, 2022). In theory, this complexity renders it difficult to compute statistical power for subgroup analyses, and one could even consider it inappropriate to even attempt a power analysis for such an exploratory tool (Gao et al., 2023). However, in practice, countless papers have reported findings from underpowered analyses that are likely to be false positives (Toffalini et al., 2022). This issue could have been avoided if researchers had the tools to estimate power *before* touching any data.

In this tutorial, I outline a method to establish the required number of observations (sample size) and features (number of measured variables) *a priori* for popular subgroup analyses. These

include k-means clustering, agglomerative hierarchical clustering with Ward linkage, c-means fuzzy clustering, latent class analysis, latent profile analysis, and Gaussian mixture modelling. My goal is to encourage researchers to consider statistical power before they start new data collection or secondary data processing, and to give them the tools to determine whether subgroup analyses would be appropriate for their data.
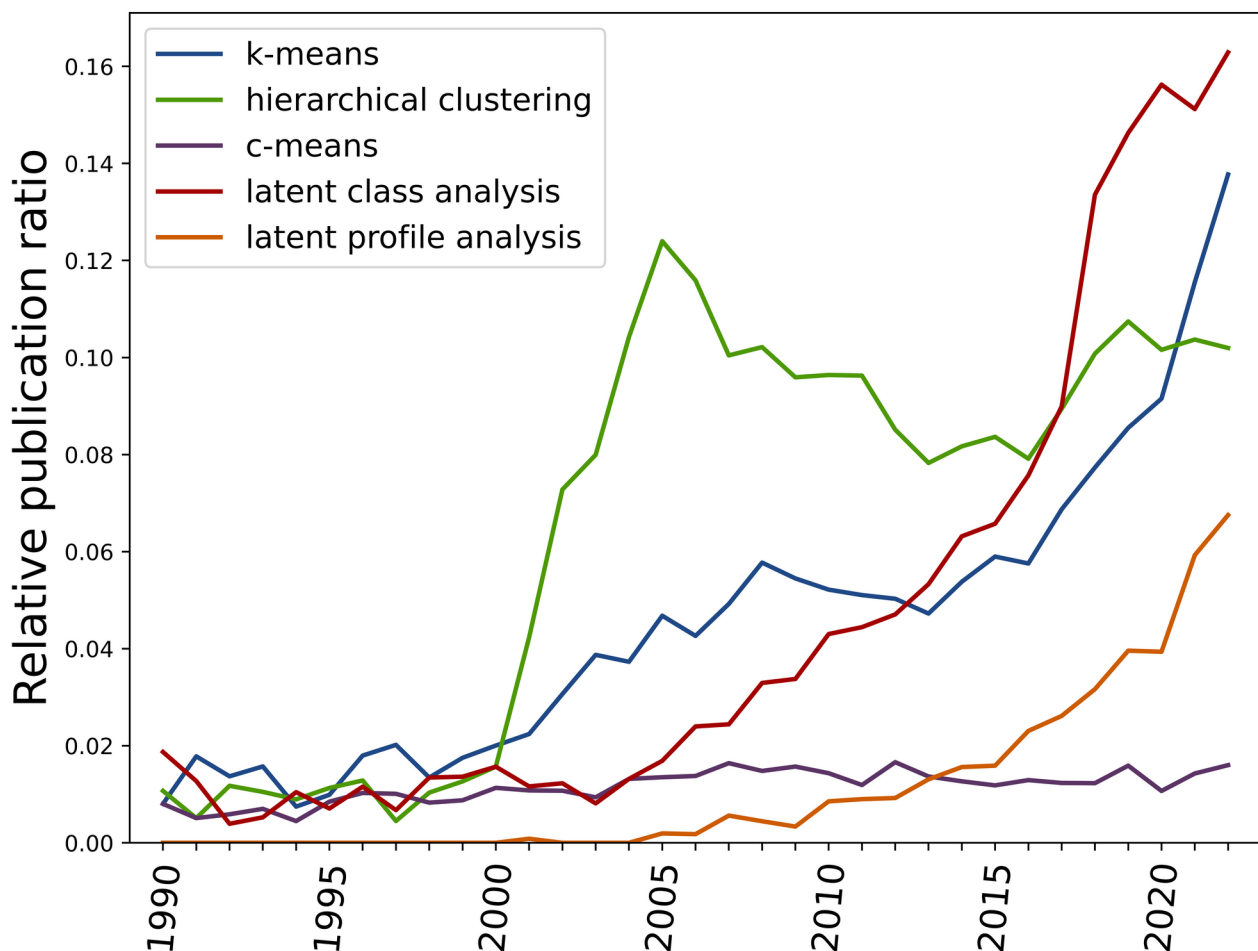
Figure 1. *The number of papers in journals indexed by PubMed that referenced a specific subgroup analysis, relative to the number of papers that referenced "t-test". This comparison shows that subgroup analyses have steadily gained in relative popularity (compared to traditional tests) since the year 2000, and an steeper increase since 2015. Plot generated using bibliobana (Dalmaijer, Van Rheede, et al., 2021).*

What is statistical power in subgroup analyses?

In a frequentist framework, "power" refers to the probability that a test can detect a true positive. In simulations, it can be computed as the proportion of tests that result in a statistically significant result when a real effect was introduced. In subgroup analyses, power is a more elusive concept.

Where a traditional test offers a binary decision on there being an effect, a subgroup analysis offers much more information, including how many subgroups are present, and which observations belong to which subgroup. However, its central question can be reduced to a binary decision: Are any subgroups present in this data? The null hypothesis here is that all observations belong to a single group. Statistical power for subgroup analyses can thus be expressed as the probability of the single-group null hypothesis being correctly rejected. (Dalmaijer et al., 2020, 2022)

In simulations, statistical power can be estimated by generating two (or more) groups of data, employing a subgroup analysis, and then evaluating its outcome. Outcome evaluation entails computing the silhouette score (Rousseeuw, 1987) or its extension for probabilistic group membership (Campello & Hruschka, 2006), with scores over 0.5 rejecting the single-group null hypothesis. Power is then computed as the proportion of simulation runs that result in a silhouette score of 0.5 or higher.

Alternatives to subgroup evaluation with silhouette scores?

The silhouette score is not the only evaluation metric. However, it performs well compared to other validation indices across a wide variety of simulated and real datasets (Arbelaitz et al., 2013), even if it has some weaknesses for distributions with atypical shapes (Batool & Hennig, 2021).

Data in cognitive and behavioural research typically comprises partially overlapping normal distributions, and in this context evaluation with silhouette scores was without false positives in simulated single-group data, and with high accuracy in simulated multi-group data (Dalmaijer et al., 2020, 2022). Not all evaluation metrics are designed for this type of data; for example, the popular "gap statistic" is explicitly not intended to be used with overlapping groups (Tibshirani et al., 2001). This renders it, and methods that average across a variety of evaluation metrics (not all of which will be appropriate), unsuitable for our purposes.

Some subgroup analyses (e.g. mixture modelling) lend themselves to computing a goodness of fit (e.g. Bayesian Information Criterion) of a single-group null model, and comparing it with fits of multi-group models. While this could also support a binary rejection of the single-group null hypothesis, the beauty of this approach is that it offers a more continuous measure of evidence. This does not mean Bayesian evaluation is the most sensitive: In simulations with 1 to 4 subgroups (multivariate normal distributions with partial overlap), the thresholding of silhouette scores proved to be more accurate than the goodness-of-fit approach for detecting more than 2 subgroups with Gaussian mixture modelling without costing additional false positives (Dalmaijer et al., 2022).

In sum, the silhouette score has desirable properties in an evaluation metric. It can be computed for many (if not all) types of subgroup analyses, it works comparatively well in a wide variety of datasets, and its behaviour is well described for multivariate normal distributions with partial overlap (i.e. realistic data in cognitive and behavioural research).

Estimating subgroup effect size

In a test of differences between two unrelated samples, effect size can be computed as the difference between the groups' means in standardised space (when the pooled standard deviation is equal to 1). Readers might recognise this as Cohen's *d,* and typical interpretations suggest that effects are very small at 0.01, small at 0.2, medium at 0.5, large at 0.8, very large at 1.2, and huge at 2.0 (Sawilowsky, 2009). Effect sizes for such differences will be referred to as δ.

The effect size in subgroup analyses can be defined as the distance between subgroup centroids. The higher this distance, the easier to identify each subgroup. When there are more than two subgroups, the lowest centroid distance determines how separable all subgroups are. Centroid distance is the effect size for subgroup analyses, and will be referred to as Δ.

The distance between centroids can be computed as the Euclidean distance between groups in standardised space (with feature variances of 1). This reflects an accumulation over differences between subgroups within each variable, and can be computed with Equation 1 (Dalmaijer et al., 2020, 2022).

$$(1) \quad \Delta = \sqrt{\sum_{i=1}^{p} \delta_i^2}$$

where Δ is the centroid distance (subgroup effect size), p is the number of features (measured variables), and δ_i is the difference effect size for feature i.

Underlying (within-feature) effect size distributions

While it is now possible to compute the subgroup effect size, the underlying effect sizes for each measured variable still need to be estimated. In a power analysis for a single difference, researchers typically turn to the published literature to estimate an effect size from previous studies or meta-analyses, where possible with a correction for publication (Van Aert & Van Assen, 2018) and outcome reporting bias (Van Aert & Wicherts, 2023). Ideally, one would do the same for each feature in a subgroup analysis, although it is unfeasible to do so for a large number of variables. Another obstacle is that it is typically unclear which subgroups are present in a population before a subgroup analysis is done, which makes it impossible to estimate any single underlying effect size.

While it may be impossible to estimate effect sizes within each measured variable, the distribution of effect sizes within a field of study can be established empirically: Szucs and Ioannidis (2017) extracted 26841 reported statistics from 3801 papers in cognitive neuroscience and psychology. Out of these, 17207 were "statistically significant" with an average effect size of δ=0.932, and 9634 were non-significant with an average δ=0.237, for a grand average of δ=0.683. The distribution of effect sizes was roughly exponential, with more papers finding smaller effects. The probability density of effect sizes can thus be approximated with Equation 2.

(2) $\quad f(\delta, \lambda) = \lambda \, e^{-\lambda \delta}$

where δ is an effect size for a difference, and λ is the exponential rate parameter.

The mean of an exponential distribution is given by the reciprocal of its λ parameter (Equation 3). Assuming the underlying effect size distribution is indeed exponential, λ can thus be estimated as the reciprocal of the grand mean from Szucs and Ioannidis (2017): $1 / 0.683 = 1.464 \approx 1.5$ (rounded).

(3) $\quad E[\delta] = \dfrac{1}{\lambda}$

where $E[\delta]$ is the mean effect size, and λ is the exponential rate parameter.

Estimating centroid separation (subgroup effect size)

The expected value for centroid separation (i.e. subgroup effect size) can be directly estimated by combining Equations 1 and 3. This offers researchers the opportunity to estimate subgroup effect size from the number of features (measured variables) included in the analysis and the anticipated λ parameter for the underlying effect size distribution (Equation 4).

(4) $\quad \hat{\Delta} = \sqrt{p \left(\dfrac{1}{\lambda} \right)^2}$

where Δ *(hat)* is the estimated subgroup effect size, p is the number of features included in a subgroup analysis, and λ is the exponential rate parameter.

Choosing a λ value for subgroup effect size estimation

The span of published effect sizes is substantial, and includes very large (δ=1.2) to huge (δ=2) effects (Szucs & Ioannidis, 2017). These are not *impossibly* large effects. For example, the difference in average height between men (M=174.91 cm, SD=2.51, n=5153) and women (M=162.03 cm, SD=2.37, n=5763) in the Finnish Twin Cohort Study amounts to an effect size of around δ=5 (Silventoinen et al., 2000). However, they are *implausible* effect sizes due to the abundance of studies with small sample sizes in the literature (Szucs & Ioannidis, 2017). By definition, these only have power to detect large effects, and paired with file-drawers full of null

results (Rosenthal, 1979) and publication bias (Smart, 1964), this has likely inflated reported effects sizes (Szucs & Ioannidis, 2017).

To account for this inflation, low effect sizes can be upweighted by increasing the magnitude of the λ parameter. Here, I used the particularly optimistic value of $\lambda=0.75$, which should only be used when many substantial effects are anticipated (e.g. data from entirely different animal or plant species). I also used $\lambda=1.5$ to match effect sizes published in psychological literature (Szucs & Ioannidis, 2017), $\lambda=3$ to only include up to very large effects, $\lambda=6$ to only include up to large effects, and $\lambda=12$ to only include up to medium effects (Figure 2).
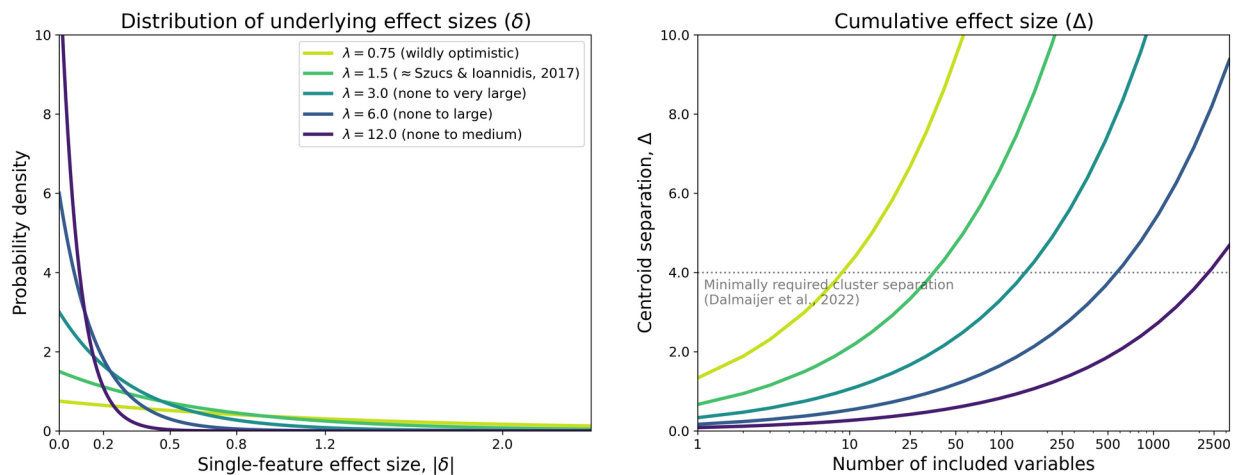
Figure 2. *Subgroup effect size as an accumulation of underlying effect sizes in different contexts. The left panel shows the probability densities (y-axis) as a function of effect sizes (difference between subgroups within a single feature) in five different contexts. The right panel shows estimated centroid separation (subgroup effect size) as an accumulation of single-feature effect sizes over features (x-axis, log-scaled). Effect size contexts ranged from highly liberal (lighter colours) to highly conservative (darker colours). Where they cross the dotted line reveals the number of features required to achieve an estimated subgroup effect size of Δ=4, which should typically offer sufficient statistical power to detect subgroups with k-means, c-means, HDBSCAN, or Gaussian mixture modelling (Dalmaijer et al., 2020, 2022). The main message of this figure is that many variables (500-2500) need to be included in context that are biased towards small effects.*

Estimating the minimum number of variables

With Equation 4, the centroid separation (subgroup effect size) can be estimated from the λ parameter of the underlying effect size distribution and the number of features included in an analysis. The λ parameter should be chosen on the basis of a datasets origin (see previous heading), but how to estimate the anticipated centroid separation or the required number of features is yet undiscussed. Choosing either parameter depends on whether data has already been collected.

If data has already been collected (e.g. when doing secondary data analysis), researchers can do a sensitivity analysis by using their chosen λ value and the number of features they can include

(i.e. the number of useable variables in the dataset). Using Equation 4, this offers them an estimated value of subgroup effect size Δ. By comparing this estimate with existing power computations for subgroup analyses published by Dalmaijer and colleagues (2020, 2022). The general recommendation here is that $\Delta>4$ should provide sufficient power to detect subgroups with k-means clustering, and $\Delta>3$ could be enough when using c-means or Gaussian mixture modelling are used (both assume multi-dimensional scaling is employed as a pre-processing step).

If data is yet to be collected, the number of variables to measure is undecided. The logic from the previous paragraph can be reversed: first a value of Δ for sufficient statistical power can be chosen, e.g. $\Delta=4$. Next, the minimum number of variables that should be included can be estimated by solving Equation 4 for the chosen Δ, which results in Equation 5.

$$(5) \quad p = \frac{\Delta^2}{\left(\frac{1}{\lambda}\right)^2}$$

where p is the number of features that should be included in a subgroup analysis, Δ is the required subgroup effect size, and λ is the anticipated exponential rate parameter.
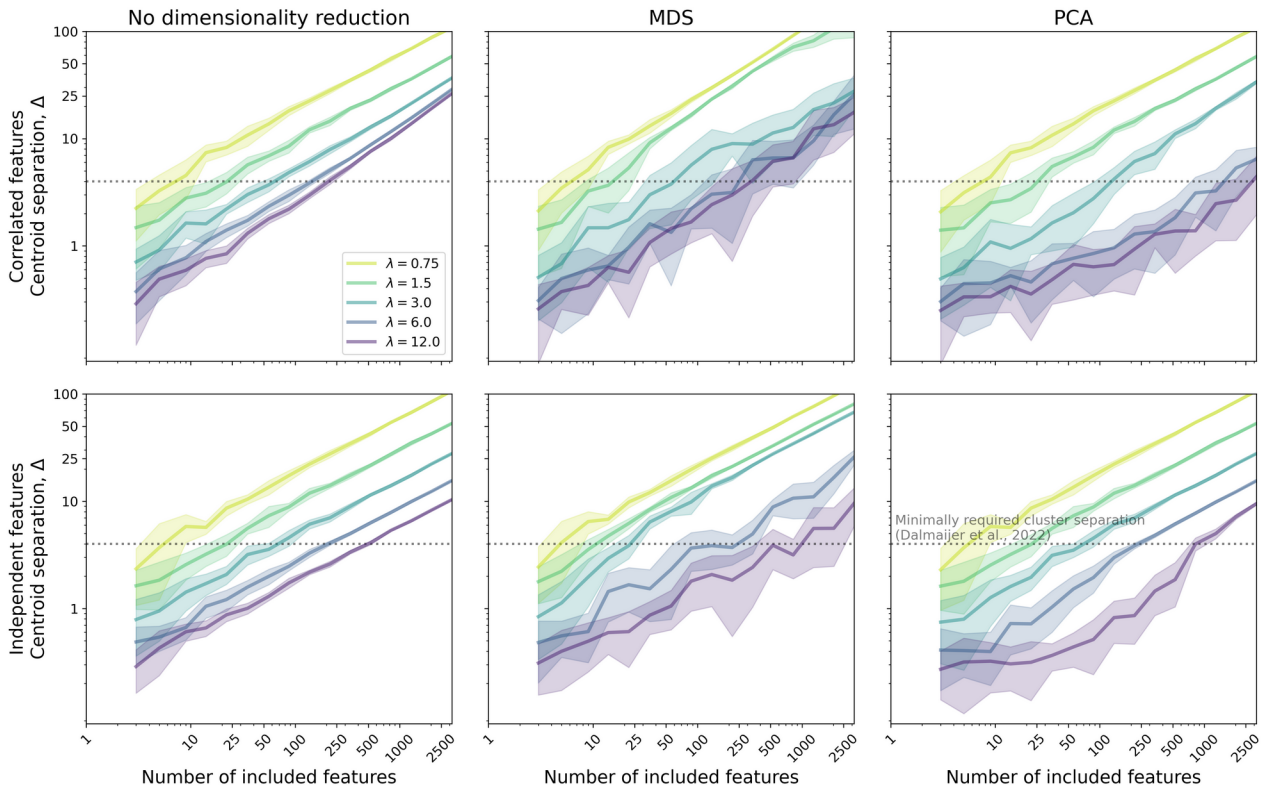
Dimensionality reduction

When performing a subgroup analysis on a dataset with many features, the "curse of dimensionality" can reduce performance even if a large distance between centroids exists. One way to combat this is by employing a dimensionality reduction algorithm to project high-dimensional data into a low-dimensional representation that can then be used in a subgroup analysis. An additional benefit of this approach is that multi-dimensional scaling can help increase centroid distances (Dalmaijer et al., 2020, 2022).

Dimensionality reduction algorithms for continuous data typically work by aiming to find a low-dimensionality representation in which relative distances between individual observations are analogous to those in the original data. These algorithms are unsuitable for categorical data, where the concept of "distance" does not apply in the same way. Dimensionality reduction can still be achieved to some extent by recoding categorical data into binary dummy variables (also know as "one hot encoding" in the machine-learning literature), which can then be reduced using "sketching" algorithms (Bera et al., 2023; Mitzenmacher et al., 2014).

Care should be taken in selecting a dimensionality reduction algorithm. For example, multi-dimensional scaling is computationally expensive and takes significant time on large datasets. Principal component analysis can be a faster alternative, but is more sensitive to covariance between features, particularly when underlying effect sizes are small (Figure 3). This can reduce centroid distances and thus hamper subgroup separability, especially in environments where within-

variable differences are small and variables are likely to be correlated (e.g. in psychological and medical research).

Figure 3*. Distance (Δ, y-axis) between centroids computed before dimensionality reduction (left column), and after multi-dimensional scaling (MDS, middle column) or principal component analysis (PCA, right column) on correlated (top row) and independent (bottom row) features for different numbers of features (p, x-axis). Lighter colours indicate more optimistic effect size distributions (see Figure 2 for exact distributions). Solid lines are averages with shaded areas reflecting standard deviations across 10 simulations.*

Estimating sample size

Previous work suggests that *when centroid separation is sufficient*, the required sample size asymptotes at around *n*=30 per subgroup (Dalmaijer et al., 2020, 2022). Other recommendations suggest that the total sample size should be 70 times the number of features (Dolnicar et al., 2014); or even that the number of observations should be at least 2^p, where *p* is the number of features (typically attributed to Formann, 1984). These recommendations vary substantially, and this will be addressed shortly.

How large the total sample should be depends on how many subgroups are hypothesised to exist, and what each subgroup's relative size is. For example, if three equally sized subgroups are

hypothesised to exist, the total sample size would be 3*30 = 90. If two subgroups are hypothesised, but only 10% of the target population is expected to be a member of one of the subgroups, then at least 300 observations should be included (so that 10% of the sample results in n=30 for the smaller subgroup).

Further to the above, the preferred sample size depends on the number of included features. Two of the aforementioned recommendations are to increase the sample size as a function of the number of included features. According to these guidelines, inclusion of 100 features should require a sample size of 70*100 = 7000 or of 2^{100} = 1.27e30. These numbers are on entirely different orders of magnitude. The latter estimate is also impracticable: at the time of writing, the world population is estimated to be only 8e12 (United Nations, 2022), 1.58e17 times smaller than the intended sample size.

To address the issue of conflicting guidelines, I ran simulations that included N=30 to N=5000 per subgroup in various underlying effect size contexts. The results are covered in detail in the next section, but in short they suggest that while sample size should scale with the expected underlying effect size distribution, it can also be somewhat compensated by the number of features included in an analysis. This means that including more measured variables can actually reduce the sample size demands of a subgroup analysis.

Reference table for popular analyses

To provide rough estimates for the required sample size and number of measured variables, I ran simulations with varying numbers of observations and features in different contexts of underlying effect sizes. Each simulated dataset was subjected to subgroup analyses k-means clustering, agglomerative hierarchical clustering with Ward linkage, c-means fuzzy clustering, latent profile analysis, and Gaussian mixture modelling. Categorical data was also simulated, and processed with latent class analysis. (Some of these analyses are mostly the same; see next section.)

Table 1 shows that the required numbers of observations and features increase as a function of the effect size context. Researchers who are planning a subgroup analysis in an environment with small underlying effect sizes should thus plan to employ a large sample size and to measure many variables. Where this is not possible, they should reconsider doing a subgroup analysis.

Note that the table offers rough estimates. Higher precision estimates would be derived from running simulations specific to planned pipelines, including preprocessing, dimensionality reduction, and the subgroup analysis of choice.

Table 1
This table offers rough estimates for the minimal sample size per expected group (n) and number of measured variables (p) for several popular subgroup analyses. The estimations were derived from simulations in five different distribution of underlying effect sizes (reflected by λ parameter). The "Published in psychology" column refers to the distribution of effect sizes observed in psychology-focussed journals (Szucs & Ioannidis, 2017). Columns are organised from most optimistic to most conservative, with λ reflecting the degree of weight towards null effects (see also Figure 2). Simulated data comprised uncorrelated variables that were projected into two dimensions using PCA (chosen for computational efficiency; not a preprocessing recommendation, see also Figure 3). Two options are given per row, one prioritising fewer features (top), and another prioritising fewer observations (bottom). Where the table states "No detection", subgroup analyses could not reliably identify the true two-subgroup solution over a single-group null model.

	Wildly optimistic (λ=0.75)	**Published in psychology** (λ=1.5)	**No to very large effects** (λ=3.0)	**No to large effects** (λ=6.0)	**No to medium effects** (λ=12)
K-means	*n=75, p=9* *n=30, p=14*	*n=30, p=36* *n=30, p=56*	*n=100, p=144* *n=30, p=225*	*n=750, p=324* *n=200, p=576*	*n=2000, p=1296* *n=500, p=2304*
Ward (agglomerative hierarchical clustering)	*n=75, p=9* *n=30, p=14*	*n=30, p=36* *n=30, p=56*	*n=100, p=144* *n=30, p=225*	*n=200, p=576* *n=100, p=900*	*n=5000, p=1296* *n=500, p=2304*
C-means (fuzzy clustering)	*n=50, p=9* *n=30, p=14*	*n=100, p=20* *n=30, p=36*	*n=100, p=81* *n=30, p=144*	*n=500, p=324* *n=100, p=576*	*n=1000, p=1296* *n=500, p=2304*
Latent class analysis	*n=50, p=9* *n=30, p=14*	*n=150, p=20* *n=100, p=36*	No detection	No detection	No detection
Latent profile analysis	*n=50, p=9* *n=30, p=14*	*n=30, p=36* *n=30, p=56*	*n=150, p=81* *n=50, p=144*	*n=500, p=324* *n=200, p=576*	*n=1000, p=1296* *n=500, p=2304*
Gaussian mixture modelling	*n=75, p=9* *n=30, p=14*	*n=30, p=36* *n=30, p=56*	*n=50, p=144* *n=30, p=225*	*n=500, p=324* *n=200, p=576*	*n=1500, p=1296* *n=500, p=2304*

*Note: the n values in this table are **per subgroup**, so the total sample size should be a multiple of the suggested values. For example, if three equally sized subgroups are expected, the planned sample size should be 3*n. If two subgroups are expected, with one being four times smaller than the other (i.e. an 80-20% divide of the population), the total sample size would be 4*n+n.*

Simulation details

Simulated datasets comprised two equally sized subgroups, each represented by multivariate normal distributions. Effect sizes were drawn from exponential distributions with λ parameters of 0.75, 1.5, 3.0, 6.0, or 12.0 to match the distributions from Figures 2 and 3. For improved computational efficiency, features were uncorrelated and reduced through principal component analysis. These are not recommended assumptions or approaches, but chosen here only because they save on

computing time by avoiding searches for positive-definite covariance matrices and dimensionality reduction through multi-dimensional scaling. For an idea of how feature correlations and dimensionality reduction algorithms impact subgroup analyses, see Figure 3 and Dalmaijer et al. (2022).

Feature numbers (p) were chosen for subgroup effect sizes Δ=3, Δ=4, or Δ=5, using Equation 5. Sample sizes were run from n=30 to n=5000 per subgroup, resulting in N=60 to N=10000 for the full sample (comprising two subgroups). For each included combination of n and p, 50 simulations were run. In each of these, one dataset was generated, and subjected to the following subgroup analyses.

The k-means clustering algorithm iteratively moves k centroids until a stable solution is reached. In each iteration, observations are assigned to their closest centroid, and new centroids are defined as the average location of all their assigned observations (Lloyd, 1982).

In agglomerative hierarchical clustering, observations are joined as a function of their similarity. Similarity is typically defined by a distance metric (here Euclidean, but city block or cosine are also used), and compared against a linkage criterion. A commonly used criterion is Ward linkage, which minimises the increase in variance when observations are linked (Ward, 1963). The result is a hierarchy of solutions on the bottom of which each observation is its own cluster, and at the top of which all observations are part of the same cluster. Like for k-means, the accepted number of clusters is user-defined, or chosen as the solution with the highest silhouette score.

The c-means fuzzy clustering algorithm operates in much the same way as k-means, with one major difference: cluster membership is not binary. Instead, each observation is assigned a degree of belonging to each cluster (Bezdek, 1981; Dunn, 1973; Ross, 2010).

Latent class analysis is a latent mixture modelling approach in which predictors are binary (Bernoulli) or categorical (generalised Bernoulli or multinoulli) variables (Ferguson et al., 2020). Confusingly, the term "latent class analysis" has also been used to describe mixture modelling approaches more generally in the literature.

Latent profile analysis is a latent mixture modelling approach in which predictors are continuous variables that are assumed to be independent (Ferguson et al., 2020; Sterba, 2013). As with latent class analysis, the term "latent profile analysis" is also frequently used to describe mixture modelling more generally.

Gaussian mixture modelling is a latent mixture modelling approach in which predictors are continuous, and covariance structures can vary between subgroups. Here, all groups are allowed their own covariance structure. An alternative approach is to fit a single covariance matrix across all subgroups, which could help prevent over-fitting. Within mixture modelling approaches, the probability that an observation belongs to each cluster can be computed (analogous to fuzzy clustering).

For each subgroup analysis outside of latent class analysis, silhouette (Rousseeuw, 1987) or fuzzy silhouette (Campello & Hruschka, 2006) scores were computed. Power was then computed as the proportion of subgroup analysis that correctly rejected the single-group null hypothesis, i.e. the

proportion of silhouette scores of 0.5 and over. Reported values for the required number of observations per subgroup and included features are for when analyses reached 90% power (45/50 simulations).

For latent class analysis, categorical values were simulated and then directly analysed via StepMix (Morin et al., 2023) without dimensionality reduction. This is because while it is possible to reduce categorical data, distances between observations are of a different quality and not always retained (Bera et al., 2023). Because silhouette scores operate on the same distance metric, their meaning is also hampered. In order to provide a fairer comparison, goodness of fit was computed for a null model (one group) and an alternative model (two subgroups). The two Bayesian Information Criteria were combined into a single Bayes Factor (Wagenmakers, 2007), for which values over 3 were taken as evidence for the two-subgroup model over the single-group null model.

Simulations were coded in Python version 3.8.10 (for a tutorial, see Dalmaijer, 2017), using NumPy version 1.22.3 (Harris et al., 2020), SciPy version 1.8.0 (Virtanen et al., 2020), scikit-learn version 1.1.0 (Pedregosa et al., 2011), scikit-fuzzy version 0.4.2, StepMix version 2.1.1 (Morin et al., 2023), the Intel extension for scikit-learn version 2023.2.1, and Matplotlib version 3.5.2 (Hunter, 2007). Code and generated data are freely available via GitHub: https://github.com/esdalmaijer/cluster_power_tutorial

Conclusion

Before embarking on a subgroup analysis, researchers must choose how many observations (sample size) and features (measured variables) to include. The statistical power of subgroup analyses depends on subgroup effect size, which is quantified as the distance between subgroup centroids in standardised space. Without sufficient separation, subgroups cannot be detected. Centroid separation accumulates over measured variables, with small differences between subgroups within each variable determining centroid separation in multi-dimensional space. The effect sizes of these underlying differences can be modelled as an exponential distribution, with a λ parameter estimated from the existing literature. Researchers can use this approach via the following workflow:

1. Choose the preferred subgroup analysis pipeline. This typically includes preprocessing, a dimensionality reduction algorithm, and a cluster analysis or mixture model.

2. Choose the subgroup effect size that should be minimally attained. For most approaches, this will be around $\Delta=4$, but for some it might be as low as $\Delta=3$ (e.g. c-means and mixture modelling).

3. Choose the expected underlying effect size distribution. This is an exponential distribution biased towards 0, with the degree of null bias determined by the λ parameter. (Empirical estimates put at λ at 1.5 in the psychological literature, but this is likely an under-estimation. A safer choice is $\lambda=3$ at a minimum.)

4. Estimate the required number of features using the chosen subgroup effect size (Δ) and underlying effect size distribution parameter (λ) in Equation 5.

5. Simulate multivariate normal distributions with the computed number of variables, draw underlying effect sizes from an exponential distribution with the chosen λ parameter to introduce a difference between two groups, subject simulated datasets to the chosen dimensionality reduction and subgroup analysis, and compute the solutions' silhouette scores. Increase the simulated number of participants until a high proportion (towards 1) of simulations results in a sufficient silhouette score (0.5 or over). The resulting number is the sample size for two equally sized subgroups; convert this to the total sample size using the number of subgroups hypothesised to exist in the data.

6. If the number of features is higher than the number of variables that can feasibly be collected (for new data) or that has already been collected (for secondary data), then subgroup analysis might not be a suitable approach.

Alternatively, researchers could use Table 1 for coarse estimates of the required number of observations per subgroup and the required number of features. This can be used as a first-pass on whether or not subgroup analysis is a feasible approach, i.e. whether the required sample size and number of measured variables are attainable in new data or present in existing data.

References

Alexander, E. F., & Johnson, M. D. (2023). On categorizing intimate partner violence: A systematic review of exploratory clustering and classification studies. *Journal of Family Psychology, 37*(5), 743–752. https://doi.org/10.1037/fam0001075

Arbelaitz, O., Gurrutxaga, I., Muguerza, J., Pérez, J. M., & Perona, I. (2013). An extensive comparative study of cluster validity indices. *Pattern Recognition, 46*(1), 243–256. https://doi.org/10.1016/j.patcog.2012.07.021

Barzycka, B., Grabiec, M., Jania, J., Błaszczyk, M., Pálsson, F., Laska, M., Ignatiuk, D., & Aðalgeirsdóttir, G. (2023). Comparison of Three Methods for Distinguishing Glacier Zones Using Satellite SAR Data. *Remote Sensing, 15*(3), 690. https://doi.org/10.3390/rs15030690

Bathelt, J., Holmes, J., Astle, D. E., & The CALM Team. (2018). Data-Driven Subtyping of Executive Function–Related Behavioral Problems in Children. *Journal of the American Academy of Child & Adolescent Psychiatry, 57*(4), 252-262.e4. https://doi.org/10.1016/j.jaac.2018.01.014

Batool, F., & Hennig, C. (2021). Clustering with the Average Silhouette Width. *Computational Statistics & Data Analysis, 158*, 107190. https://doi.org/10.1016/j.csda.2021.107190

Benjamins, J. S., Dalmaijer, E. S., Ten Brink, A. F., Nijboer, T. C. W., & Van der Stigchel, S. (2019). Multi-target visual search organisation across the lifespan: Cancellation task performance in a large and demographically stratified sample of healthy adults. *Aging, Neuropsychology, and Cognition, 26*(5), 731–748. https://doi.org/10.1080/13825585.2018.1521508

Bera, D., Pratap, R., & Verma, B. D. (2023). Dimensionality Reduction for Categorical Data. *IEEE Transactions on Knowledge and Data Engineering, 35*(4), 3658–3671. https://doi.org/10.1109/TKDE.2021.3132373

Bezdek, J. C. (1981). *Pattern recognition with fuzzy objective function algorithms*. Plenum Press.

Campello, R. J. G. B., & Hruschka, E. R. (2006). A fuzzy extension of the silhouette width criterion for cluster analysis. *Fuzzy Sets and Systems, 157*(21), 2858–2875. https://doi.org/10.1016/j.fss.2006.07.006

Dalmaijer, E. S. (2017). *Python for experimental psychologists*. Routledge, Taylor & Francis.

Dalmaijer, E. S., Gibbons, S. G., Bignardi, G., Anwyl-Irvine, A. L., Siugzdaite, R., Smith, T. A., Uh, S., Johnson, A., & Astle, D. E. (2021). Direct and indirect links between children's socio-economic status and education: Pathways via mental health, attitude, and cognition. *Current Psychology*, *42*, 9637–9651. https://doi.org/10.1007/s12144-021-02232-2

Dalmaijer, E. S., Nord, C. L., & Astle, D. E. (2020). Statistical power for cluster analysis. *ArXiv:2003.00381 [Cs, q-Bio, Stat]*. http://arxiv.org/abs/2003.00381

Dalmaijer, E. S., Nord, C. L., & Astle, D. E. (2022). Statistical power for cluster analysis. *BMC Bioinformatics*, *23*(1), 205. https://doi.org/10.1186/s12859-022-04675-1

Dalmaijer, E. S., Van Rheede, J., Sperr, E. V., & Tkotz, J. (2021). Banana for scale: Gauging trends in academic interest by normalising publication rates to common and innocuous keywords. *ArXiv:2102.06418 [Cs]*. http://arxiv.org/abs/2102.06418

Den Ouden, L., Suo, C., Albertella, L., Greenwood, L.-M., Lee, R. S. C., Fontenelle, L. F., Parkes, L., Tiego, J., Chamberlain, S. R., Richardson, K., Segrave, R., & Yücel, M. (2022). Transdiagnostic phenotypes of compulsive behavior and associations with psychological, cognitive, and neurobiological affective processing. *Translational Psychiatry*, *12*(1), 10. https://doi.org/10.1038/s41398-021-01773-1

Díaz Redondo, R. P., Fernández Vilas, A., & Abadía Rodríguez, A. (2020). Inferring Energy Consumption Patterns in Public Buildings. *Proceedings of the 17th ACM Symposium on Performance Evaluation of Wireless Ad Hoc, Sensor, & Ubiquitous Networks*, 33–39. https://doi.org/10.1145/3416011.3424753

Dolnicar, S., Grün, B., Leisch, F., & Schmidt, K. (2014). Required Sample Sizes for Data-Driven Market Segmentation Analyses in Tourism. *Journal of Travel Research*, *53*(3), 296–306. https://doi.org/10.1177/0047287513496475

Duimel, S. L. L., Linn, A. J., Smets, E. M. A., Smit, E. S., & Van Weert, J. C. M. (2022). Profiling Cancer Patients Based on Their Motives for Seeking Informational and Emotional Support Online. *Health Communication*, 1–15. https://doi.org/10.1080/10410236.2022.2144287

Dunn, J. C. (1973). A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters. *Journal of Cybernetics*, *3*(3), 32–57. https://doi.org/10.1080/01969727308546046

Favila, N., Madrigal-Trejo, D., Legorreta, D., Sánchez-Pérez, J., Espinosa-Asuar, L., Eguiarte, L. E., & Souza, V. (2022). MicNet toolbox: Visualizing and unraveling a microbial network. *PLOS ONE*, *17*(6), e0259756. https://doi.org/10.1371/journal.pone.0259756

Ferguson, S. L., Moore, E. W. G., & Hull, D. M. (2020). Finding latent groups in observed data: A primer on latent profile analysis in Mplus for applied researchers. *International Journal of Behavioral Development*, *44*(5), 458–468. https://doi.org/10.1177/0165025419881721

Formann, A. K. (1984). *Die Latent-Class-Analyse: Einführung in Theorie und Anwendung*. Beltz.

Gao, C. X., Dwyer, D., Zhu, Y., Smith, C. L., Du, L., Filia, K. M., Bayer, J., Menssink, J. M., Wang, T., Bergmeir, C., Wood, S., & Cotton, S. M. (2023). An overview of clustering methods with guidelines for application in mental health research. *Psychiatry Research*, *327*, 115265. https://doi.org/10.1016/j.psychres.2023.115265

Halawa, M. S., Díaz Redondo, R. P., & Fernández Vilas, A. (2020). Unsupervised KPIs-Based Clustering of Jobs in HPC Data Centers. *Sensors*, *20*(15), 4111. https://doi.org/10.3390/s20154111

Harris, C. R., Millman, K. J., van der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N. J., Kern, R., Picus, M., Hoyer, S., van Kerkwijk, M. H., Brett, M., Haldane, A., del Río, J. F., Wiebe, M., Peterson, P., … Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357–362. https://doi.org/10.1038/s41586-020-2649-2

Hennessy, Á., McCarthy, E. K., ní Chaoimh, C., Murray, D. M., & Kiely, M. E. (2023). Poor Quality Diets Characterized by Low-Nutrient Density Foods Observed in One-Quarter of 2-Year-Olds in a High Resource Setting. *The Journal of Nutrition*, S002231662372436X. https://doi.org/10.1016/j.tjnut.2023.06.029

Hunter, J. D. (2007). Matplotlib: A 2D Graphics Environment. *Computing in Science & Engineering*, *9*(3), 90–95. https://doi.org/10.1109/MCSE.2007.55

Kiconco, S., Earnest, A., Enticott, J., Hart, R., Mori, T. A., Hickey, M., Teede, H. J., & Joham, A. E. (2023). Normative cut-offs for polycystic ovary syndrome diagnostic features in adolescents using cluster analysis. *European Journal of Endocrinology*, *188*(6), 494–502. https://doi.org/10.1093/ejendo/lvad055

Kwon, S., Gwon, Y., Kim, D., Seo, I. W., & You, H. (2023). Unsupervised Classification of Riverbed Types for Bathymetry Mapping in Shallow Rivers Using UAV-Based Hyperspectral Imagery. *Remote Sensing, 15*(11), 2803. https://doi.org/10.3390/rs15112803

Lee, L. E., Pyo, J. Y., Ahn, S. S., Song, J. J., Park, Y.-B., & Lee, S.-W. (2023). Antineutrophil cytoplasmic antibody-associated vasculitis classification by cluster analysis based on clinical phenotypes: A single-center retrospective cohort study. *Clinical Rheumatology*. https://doi.org/10.1007/s10067-023-06720-7

Li, J., Wang, J., Wu, Y., Cui, Y., & Yan, S. (2022). Remote sensing monitoring of total nitrogen and total phosphorus concentrations in the water around Chaohu Lake based on geographical division. *Frontiers in Environmental Science, 10*, 1014155. https://doi.org/10.3389/fenvs.2022.1014155

Lloyd, S. P. (1982). Least squares quantization in PCM. *IEEE Transactions on Information Theory*, *28*(2), 129–137.

Malone, S. L., Oh, Y., Arndt, K. A., Burba, G., Commane, R., Contosta, A. R., Goodrich, J. P., Loescher, H. W., Starr, G., & Varner, R. K. (2022). Gaps in network infrastructure limit our understanding of biogenic methane emissions for the United States. *Biogeosciences*, *19*(9), 2507–2522. https://doi.org/10.5194/bg-19-2507-2022

Mitzenmacher, M., Pagh, R., & Pham, N. (2014). Efficient estimation for high similarities using odd sketches. *Proceedings of the 23rd International Conference on World Wide Web*, 109–118. https://doi.org/10.1145/2566486.2568017

Morin, S., Legault, R., Laliberté, F., Bakk, Z., Giguère, C.-É., de la Sablonnière, R., & Lacourse, É. (2023). *StepMix: A Python Package for Pseudo-Likelihood Estimation of Generalized Mixture Models with External Variables*. https://doi.org/10.48550/ARXIV.2304.03853

Nissenbaum, D., Sarafian, R., Rudich, Y., & Raveh-Rubin, S. (2023). Six types of dust events in Eastern Mediterranean identified using unsupervised machine-learning classification. *Atmospheric Environment*, *309*, 119902. https://doi.org/10.1016/j.atmosenv.2023.119902

Parlett-Pelleriti, C. M., Stevens, E., Dixon, D., & Linstead, E. J. (2022). Applications of Unsupervised Machine Learning in Autism Spectrum Disorder Research: A Review. *Review Journal of Autism and Developmental Disorders*. https://doi.org/10.1007/s40489-021-00299-y

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., & Duchesnay, E. (2011). Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, *12*, 2825–2830.

Rennie, J. P., Zhang, M., Hawkins, E., Bathelt, J., & Astle, D. E. (2019). Mapping differential responses to cognitive training using machine learning. *Developmental Science*. https://doi.org/10.1111/desc.12868

Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, *86*(3), 638–641. https://doi.org/10.1037/0033-2909.86.3.638

Ross, T. J. (2010). Chapter 10: Fuzzy Classification (subheading: Fuzzy c-Means Algorithm). In *Fuzzy logic with engineering applications* (3rd ed., pp. 352–353). Wiley.

Rousseeuw, P. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

Sawilowsky, S. S. (2009). New Effect Size Rules of Thumb. *Journal of Modern Applied Statistical Methods*, *8*(2), 597–599. https://doi.org/10.22237/jmasm/1257035100

Silventoinen, K., Kaprio, J., Lahelma, E., & Koskenvuo, M. (2000). Relative effect of genetic and environmental factors on body height: Differences across birth cohorts among Finnish men and women. *American Journal of Public Health*, *90*(4), 627–630. https://doi.org/10.2105/AJPH.90.4.627

Smart, R. G. (1964). The importance of negative results in psychological research. *Canadian Psychologist / Psychologie Canadienne*, *5a*(4), 225–232. https://doi.org/10.1037/h0083036

Song, J., Min, S. H., Chae, S., Bowles, K. H., McDonald, M. V., Hobensack, M., Barrón, Y., Sridharan, S., Davoudi, A., Oh, S., Evans, L., & Topaz, M. (2023). Uncovering hidden trends: Identifying time trajectories in risk factors documented in clinical notes and predicting hospitalizations and emergency department visits during home health care. *Journal of the American Medical Informatics Association*, ocad101. https://doi.org/10.1093/jamia/ocad101

Stein, M. B., & Bomyea, J. (2023). Replicating Predictive Cluster-Based Imaging Biotypes After Trauma: A Bridge Too Far? *American Journal of Psychiatry*, *180*(2), 114–116. https://doi.org/10.1176/appi.ajp.20221004

Sterba, S. K. (2013). Understanding Linkages Among Mixture Models. *Multivariate Behavioral Research*, *48*(6), 775–815. https://doi.org/10.1080/00273171.2013.827564

Szucs, D., & Ioannidis, J. P. (2017). Empirical assessment of published effect sizes and power in the recent cognitive neuroscience and psychology literature. *PLOS Biology*, *15*(3), e2000797. https://doi.org/10.1371/journal.pbio.2000797

Tibshirani, R., Walther, G., & Hastie, T. (2001). Estimating the number of clusters in a data set via the gap statistic. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *63*(2), 411–423. https://doi.org/10.1111/1467-9868.00293

Toffalini, E., Girardi, P., Giofrè, D., & Altoè, G. (2022). Entia Non Sunt Multiplicanda … Shall I look for clusters in my cognitive data? *PLOS ONE*, *17*(6), e0269584. https://doi.org/10.1371/journal.pone.0269584

Uh, S., Dalmaijer, E. S., Siugzdaite, R., Ford, T. J., & Astle, D. E. (2021). Two Pathways to Self-Harm in Adolescence. *Journal of the American Academy of Child & Adolescent Psychiatry*, S0890856721002197. https://doi.org/10.1016/j.jaac.2021.03.010

United Nations, D. of E. and S. A., Population Division. (2022). *World Population Prospects 2022: Summary of Results* (UN DESA/POP/2022/TR/NO. 3). United Nations.

https://www.un.org/development/desa/pd/sites/www.un.org.development.desa.pd/files/
wpp2022_summary_of_results.pdf

Van Aert, R. C. M., & Van Assen, M. A. L. M. (2018). *Correcting for Publication Bias in a Meta-Analysis with the P-uniform* Method* [Preprint]. MetaArXiv. https://doi.org/10.31222/osf.io/zqjr9

Van Aert, R. C. M., & Wicherts, J. M. (2023). Correcting for outcome reporting bias in a meta-analysis: A meta-regression approach. *Behavior Research Methods*. https://doi.org/10.3758/s13428-023-02132-2

Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., Burovski, E., Peterson, P., Weckesser, W., Bright, J., van der Walt, S. J., Brett, M., Wilson, J., Millman, K. J., Mayorov, N., Nelson, A. R. J., Jones, E., Kern, R., Larson, E., … Vázquez-Baeza, Y. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature Methods, 17*(3), 261–272. https://doi.org/10.1038/s41592-019-0686-2

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review, 14*(5), 779–804. https://doi.org/10.3758/BF03194105

Ward, J. H. (1963). Hierarchical Grouping to Optimize an Objective Function. *Journal of the American Statistical Association, 58*(301), 236–244.

Wilson, A. T., Riley, J. L., Bishop, M. D., Beneciuk, J. M., Cruz-Almeida, Y., Markut, K., Redd, C., LeBlond, N., Pham, P. H., Shirey, D., & Bialosky, J. E. (2023). Pain phenotyping and investigation of outcomes in physical therapy: An exploratory study in patients with low back pain. *PLOS ONE, 18*(2), e0281517. https://doi.org/10.1371/journal.pone.0281517

Wolf, K., Scharoba, J., Noack, R., Keller, A., & Weidner, K. (2023). Subtypes of borderline personality disorder in a day-clinic setting—Clinical and therapeutic differences. *Personality Disorders: Theory, Research, and Treatment*. https://doi.org/10.1037/per0000624