RUNNING HEAD:

AUTOMATIC FORCED-CHOICE ITEM PAIRING

autoFC: An R Package for Automatic Item Pairing in Forced-Choice Test Construction

Mengtong Li¹, Tianjun Sun^{1,2}, and Bo Zhang³

¹Department of Psychology, University of Illinois at Urbana-Champaign

²Department of Psychological Sciences, Kansas State University

³Department of Psychological & Brain Sciences, Texas A&M University

DRAFT version. This manuscript is currently under review.

Author Notes:

We thank Dr. Fritz Drasgow at the University of Illinois at Urbana-Champaign for his valuable suggestions during the development of this project.

Correspondence concerning this article should be addressed to Mengtong Li, Department of Psychology, University of Illinois at Urbana-Champaign, Champaign, IL 61820

Contact: ml70@illinois.edu

Abstract

Recently, there has been increasing interest in adopting the forced-choice (FC) test format in non-cognitive assessments, as it demonstrates faking resistance when well-designed. However, traditional or manual pairing approaches to FC test construction are time- and effort- intensive, and often involves insufficient considerations. To address these issues, we developed the new open-source *autoFC* R package to facilitate automated and optimized item pairing strategies. The *autoFC* package is intended as a practical tool for FC test constructions. Users can easily obtain automatically optimized FC tests by simply inputting the item characteristics of interest. Customizations are also available for considerations on matching indices and the behaviors of the optimization process. The *autoFC* package should be of interest to researchers and practitioners constructing FC scales with potentially many metrics to match on and/or many items to pair, essentially exempting users from the burden of manual item pairing and reducing the computational costs and biases induced by simple ranking methods.

Keywords:

forced-choice, pairwise preference, item pairing, automation, R package, autoFC

autoFC: An R Package for Automatic Item Pairing in Forced-Choice Test Construction

Forced-choice (FC) tests have gained increasing attention and popularity among researchers and practitioners, for its faking-resistance when well-designed (Cao & Drasgow, 2019). A well-designed FC test often entails: (1) each item within a block measures a different latent trait, and (2) items within a block have similar level of social desirability (Zhang et al., 2020). Recently, Pavlov et al. (2021) suggested that items within a block should be matched on inter-item agreement of their social desirability ratings to resolve potential issues of matching based on mean desirability ratings. Depending on the scoring models, FC developers may also need to maximize factor loading differences (Brown & Maydeu-Olivares, 2011) or minimize item location differences (Cao & Drasgow, 2019) for items within a block for statistical reasons. Deciding on what items should be assigned to the same block—item pairing—is thus critical to the quality of an FC test. Item pairing can be essentially seen as an optimization process, which is currently carried out manually. However, given that we often need to simultaneously meet multiple objectives (e.g., minimized desirability differences and maximized loading differences), manual pairing becomes impractical and even infeasible once the number of latent traits and/or the number of items per trait are relatively large.

Description of the Package

The *autoFC* package is developed to facilitate automatic FC test construction. It offers users the functionality to (1) customize one or more item pairing criteria and calculate a composite pairing index, termed "*energy*", with user-specified weights for each criterion, (2) automatically optimize the energy for the whole test by sequentially or simultaneously

optimizing each matching rule, through the exchange of items among blocks or replacement with unused items, and (3) construct parallel forms of the same test following the same pairing rules.

Users can create an FC test of any block size (e.g., pairs, triplets, quadruplets).

Functions of *autoFC* have default settings, which allows users to obtain automatically constructed assembled tests by only providing item characteristics for each item. However, users are encouraged to manipulate arguments to change default values. Argument customizations are enabled for matching indices and their relative weights for each item characteristic, an initial paired test for further improvement, and person-level item responses if items are matched based on inter-item agreement metrics (Pavlov et al., 2021). Users can also control the behavior of optimization by modifying the corresponding arguments in the function.

Availability and Distribution

autoFC is an open-source package written in the R language (R Core Team, 2021) and is compatible with all major platforms (i.e., Windows, Linux, and MacOS). All source code and documentation files are freely available from the Comprehensive R Archive Network (CRAN) via <https://cran.r-project.org/package=autoFC>. The development version of the package can be tracked at <https://github.com/tspyched/autoFC>.

References

- Brown, A., & Maydeu-Olivares, A. (2011). Item response modeling of forced-choice questionnaires. *Educational and Psychological Measurement*, 71(3), 460-502.
- Cao, M., & Drasgow, F. (2019). Does forcing reduce faking? A meta-analytic review of forced-choice personality measures in high-stakes situations. *Journal of Applied Psychology*, 104(11), 1347–1368.
- Pavlov, G., Shi, D., & Maydeu-Olivares, A. (2021, January 5). Item Desirability Matching in Forced-choice Test Construction. <https://doi.org/10.31234/osf.io/hmnrc>
- R Core Team. (2021). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. <https://www.R-project.org>
- Zhang, B., Sun, T., Drasgow, F., Chernyshenko, O. S., Nye, C. D., Stark, S., & White, L. A. (2020). Though forced, still valid: Psychometric equivalence of forced-choice and single-statement measures. *Organizational Research Methods*, 23(3), 569-590.

autoFC: An R Package for Automatic Item Pairing in Forced-Choice Test Construction

Mengtong Li, Tianjun Sun, Bo Zhang

Overview

Forced-choice (**FC**) tests are gaining researcher’s interest increasingly for its faking resistance when well-designed. Well-designed FC tests should often be characterized by *items within a block measuring different latent traits*, and *items within a block having similar magnitude, or high inter-item agreement (**IIA**) in terms of their social desirability*. Other scoring models may also require *factor loading differences or item locations within a block to be maximized or minimized*.

Either way, decision on which items should be assigned to the same block - item pairing - is a crucial issue in building a well-designed FC test, which is currently carried out manually. However, given that we often need to simultaneously meet multiple objectives, manual pairing will turn out to be impractical and even infeasible, especially when the number of latent traits and/or the number of items per trait become relatively large.

The R package **autoFC** is developed to address these difficulties and provides a tool for facilitating automatic FC test construction. It offers users the functionality to:

1. Customize one or more item pairing criteria and calculate a composite pairing index, termed “**energy**” with user-specified weights for each criterion.
2. Automatically optimize the energy for the whole test by sequentially or simultaneously optimizing each matching rule, through the exchange of items among blocks or replacement with unused items.
3. Construct parallel forms of the same test following the same pairing rules.

Users are allowed to create an FC test of any block size (e.g. Pairs, Triplets, Quadruplets).

Functions

Below is a brief explanation of all functions provided by **autoFC**. Details and usage can be found in the next section.

1. `cal_block_energy()` and `cal_block_energy_with_iaa()` both calculate the total energy for a single item block, or a full FC test with multiple blocks, given a data frame of item characteristics. The latter function incorporates IIA metrics into energy calculation.
 - By default, numeric item characteristics are paired by minimizing variance within each block, and factor item characteristics are paired such that it’s more preferable for items in the same block to be from different factor levels. Each characteristic has a default weight of 1.
 - In addition, `cal_block_energy_with_iaa()` incorporates four IIA metrics in which items are paired by maximizing the IIA within each block. Each IIA has a default weight of 1.

2. `make_random_block()` takes in number of items and block size as input arguments and produces a test with blocks of randomly paired item numbers. Information about item characteristics is not required.
 3. `get_iaa()` takes in item responses and a single item block (Or a full FC test with multiple blocks), then returns IIA metrics for each item block.
 4. `sa_pairing_generalized()` is the automatic pairing function which takes in item characteristics (and also individual responses for all items) and an initial FC test, then optimizes the energy of the test based on Simulated Annealing (SA) algorithm.
- SA is a probabilistic technique for approximating the global optimum of a given function, in which each iteration involves the cool down of the “Temperature” until it reaches a certain value. Within each iteration, a new solution (**FC test**) is produced and compared with current solution in terms of their energy (Which is calculated by calling `cal_block_energy()` or `cal_block_energy_with_iaa()`). Acceptance or rejection of new solution is determined as follows:
 - Better solution (An FC test with higher energy) is always accepted and updated into the new current solution.
 - Worse solution (An FC test with lower energy) is conditionally accepted, determined by the current temperature and the deviation of new solution’s energy from the current one’s. A worse solution is more likely to be accepted when temperature is high and when the deviance is relatively small.
 - If all items in the item characteristic data frame are used to construct the FC test, `sa_pairing_generalized()` will produce new solutions by randomly exchanging items between two blocks; Otherwise, it will randomly select between exchanging items and replacing with unused items based on proportion of items used to construct the test.
 - `sa_pairing_generalized()` has built-in default values for most of the arguments if they are not given; For example, an FC test with block size 2 using all of the given items will be constructed by default if the `block` argument is not provided. See the tutorial below for meanings and default values for other arguments.

Tutorial

Step 1: Input data

In this tutorial, we suppose that 60 5-Point Likert items measuring Big Five traits, each with a certain item location, are used to build a FC scale with block size 3. We also simulate responses from 1,000 participants for the 60 items on their social desirability.

```
set.seed(2021)
# Simulation of 1,000 respondents on 60 items. A better simulation should be
# consisting of responses produced by specific IRT parameters.
s1 <- sample(seq(1:5), 500*60, replace = TRUE,
             prob = c(0.10, 0.15, 0.20, 0.25, 0.30))
s2 <- sample(seq(1:5), 500*60, replace = TRUE,
             prob = c(0.50, 0.10, 0.10, 0.15, 0.15))

item_responses <- matrix(c(s1, s2), ncol = 60)

item_dims <- sample(c("Openness", "Conscientiousness", "Neuroticism",
                     "Extraversion", "Agreeableness"), 60, replace = TRUE)
item_mean <- colMeans(item_responses)
```

```

item_difficulty <- runif(60, -1, 1)

# Then we build a data frame with item characteristics
item_chars <- data.frame(DIM = item_dims, SD_Mean = item_mean, DIFF = item_difficulty)

char_weights = c(1, -1, -3)

```

Step 2: Construct an initial solution

Next, we build a random FC scale using the 60 items with block size 3. You can see from `initial_FC` that now all 60 items are divided into 20 triplets.

```

initial_FC <- make_random_block(total_items = 60, item_per_block = 3)
knitr::kable(initial_FC)

```

59	14	43
42	16	34
41	60	22
26	19	2
28	10	45
12	52	8
35	1	51
23	49	33
40	7	57
55	54	18
38	5	4
31	11	39
13	3	6
32	58	48
15	29	47
27	20	56
24	9	30
44	21	46
53	17	25
50	36	37

Also let's see the how the item characteristics look like for each of the 20 triplets.

First, the underlying latent traits. We see that there are some cases where two items measuring the same traits appear in the same block, which is something we want to avoid.

```

knitr::kable(matrix(item_chars$DIM[t(initial_FC)], ncol = 3, byrow = TRUE))

```

Conscientiousness	Conscientiousness	Neuroticism
Extraversion	Openness	Agreeableness
Openness	Conscientiousness	Openness
Conscientiousness	Extraversion	Openness

Neuroticism	Extraversion	Extraversion
Openness	Openness	Conscientiousness
Neuroticism	Openness	Agreeableness
Openness	Extraversion	Extraversion
Extraversion	Conscientiousness	Extraversion
Openness	Agreeableness	Extraversion
Neuroticism	Neuroticism	Neuroticism
Extraversion	Conscientiousness	Openness
Neuroticism	Conscientiousness	Neuroticism
Extraversion	Extraversion	Neuroticism
Openness	Openness	Neuroticism
Neuroticism	Conscientiousness	Conscientiousness
Neuroticism	Openness	Neuroticism
Agreeableness	Neuroticism	Openness
Conscientiousness	Conscientiousness	Agreeableness
Neuroticism	Extraversion	Openness

Then, scores on social desirability. We do see many cases where items differ in their social desirability on a magnitude of >1 on a 5-point scale, within a block. That's not good.

```
sd_initial <- matrix(item_chars$SD_Mean[t(initial_FC)], ncol = 3, byrow = TRUE)
knitr::kable(sd_initial)
```

2.327	3.549	2.257
2.269	3.417	2.410
2.331	2.386	3.549
3.575	3.481	3.492
3.628	3.469	2.430
3.470	2.326	3.418
2.306	3.518	2.299
3.455	2.367	2.382
2.372	3.429	2.291
2.383	2.275	3.501
2.346	3.499	3.493
2.371	3.476	2.381
3.429	3.540	3.478
2.375	2.383	2.283
3.551	3.484	2.367
3.458	3.521	2.296
3.515	3.474	3.524
2.275	3.501	2.408
2.354	3.526	3.449
2.314	2.294	2.344

Lastly, item difficulty. We also see that item difficulties within a block are also inconsistent.

```
diff_initial <- matrix(item_chars$DIFF[t(initial_FC)], ncol = 3, byrow = TRUE)
knitr::kable(diff_initial)
```

0.5768038	-0.1649317	-0.1609612
0.7562551	0.3665809	-0.8802425
-0.3477869	0.4033506	0.9447096
0.8602120	-0.5168748	-0.1734361
0.3712718	-0.6267511	0.0917744
-0.9688906	0.9139843	0.3986206
0.2507761	-0.0490111	-0.5887666
-0.1364442	0.7372855	0.2383125
0.6598719	0.5910397	0.3847467
0.1906729	-0.1072381	-0.2764722
0.8652686	0.3189794	-0.7635046
-0.8770999	0.5814682	0.7113918
0.9814553	-0.0455446	0.6105311
-0.7874764	-0.1775934	-0.2743858
0.5481408	-0.1896666	-0.7339903
-0.7614264	0.6233055	0.2106074
-0.4487723	0.9388348	-0.7846000
0.8782261	-0.9898177	0.0105288
0.4304538	0.3800486	-0.4484543
0.9194196	0.8647261	-0.0951617

Step 3: Calculate the energy for the initial FC scale, without and with IIAs

Next we calculate the energy for `initial_FC`, with `FUN` set to be default. `weights` is set to -1 for social desirability and -3 for item difficulty because we want the discrepancy of these characteristics to be as low as possible within a block.

The weight for item difficulty is higher to scale for its smaller range than social desirability. Beware about the scaling difference among different item characteristics and use different weights accordingly.

```
cal_block_energy(block = initial_FC, item_chars = item_chars, weights = char_weights)
#> [1] -25.28441
```

If IIAs are to be involved we have lower energy value. This is because these randomly generated responses are not likely to be consistent with each other, hence very low and even negative IIAs:

```
cal_block_energy_with_iaa(block = initial_FC, item_chars = item_chars,
                           weights = char_weights,
                           rater_chars = item_responses)
#>           [,1]
#> [1,] -33.31589
```

Notice that if we give zero weights to all IIAs we will get the same energy value as `cal_block_energy`:

```
cal_block_energy_with_iaa(block = initial_FC, item_chars = item_chars,
                           weights = char_weights,
                           rater_chars = item_responses,
                           iia_weights = c(0, 0, 0, 0))
#>           [,1]
#> [1,] -25.28441
```

Also, if you want to see the inter-item agreement metrics for each block, you can use `get_iaa()`. It should not be too impressive and is for demonstration purposes only. Users are suggested to use real world response data to see the IIA within each block.

```
knitr::kable(get_iaa(block = initial_FC, data = item_responses))
```

	BPlin	BPquad	AClin	ACquad
	-0.17500	-0.38133	-0.07767	-0.17000
	-0.11958	-0.27900	-0.03922	-0.10768
	-0.17125	-0.35733	-0.09333	-0.18799
	0.09583	0.14883	0.13905	0.22648
	-0.10458	-0.20183	-0.08040	-0.15034
	-0.09542	-0.18250	-0.08069	-0.15112
	-0.15000	-0.31867	-0.06953	-0.14620
	-0.14583	-0.30900	-0.08054	-0.16786
	-0.13167	-0.29967	-0.04673	-0.11820
	-0.14208	-0.33417	-0.05416	-0.14343
	-0.11500	-0.21550	-0.09354	-0.16960
	-0.18083	-0.36500	-0.09805	-0.18615
	0.07375	0.10000	0.11266	0.17256
	-0.03125	-0.19833	0.14874	0.15781
	-0.12875	-0.26000	-0.09955	-0.19646
	-0.12708	-0.22517	-0.11082	-0.19031
	0.11208	0.16633	0.15302	0.23982
	-0.15458	-0.32067	-0.07820	-0.15677
	-0.09208	-0.18267	-0.07283	-0.14170
	-0.05833	-0.26133	0.14720	0.14891

Step 4: Automatic pairing

To produce an optimized paired FC scale, we have the objective of:

- * **Keeping items in the same block being from different latent traits;**
- * **Minimizing variance of social desirability within each block;**
- * **Minimizing variance of item difficulty within each block.**

For IIAs, we also want to **maximize the mean of the four IIAs within each block.**

Explanation for arguments Below is an example run of producing an automatically paired FC. Arguments that may be of interest for users include:

block: The initial paired FC scale, which can be produced in Step 2. If left empty, an FC scale with block size 2 and items presented sequentially will be produced, with total number of items equals to number of rows in `item_chars`.

total_items: Default to be number of unique values in `block`. Can be a value larger than this value which represents cases where only some items in the item pool are used to build an FC scale.

Temperature: The initial temperature value of the automatic pairing method. Higher temperature is associated with higher probability of accepting a worse solution. It is recommended to leave this value blank and let it be scaled on the energy of `block` by specifying `eta_Temperature`.

r: Determines the decrease rate of `Temperature`. Should be a value between 0 and 1. Larger r values allows more iterations in the optimization process but will slow down the program.

end_criteria: Determines the end criteria for the automatic pairing process. A proportion value scaled on **Temperature**. Should be a value between 0 and 1. Smaller values allows more iterations in the optimization process but will slow down the program.

item_chars: A data frame with item characteristics for all items. It is recommended that information irrelevant to pairing be discarded beforehand, but users can also set the corresponding position in **weights** to be 0 to bypass these irrelevant item characteristics (Such is item ID).

FUN: A vector of function names for optimizing each item characteristic within each block. For example: **FUN = c('mean', 'var', 'sum')**. Also supports customized functions. Defaults to **var** for numeric variables and **facfun** for factor/character variables.

n_exchange: Determines how many blocks are exchanged in order to produce a new solution for each iteration. Should be a value less than **nrow(block)**.

weights: A vector of integer values indicating relative weights for each item characteristic after calculated by FUN. Default to be a vector of all 1s.

prob_newitem: Probability of choosing the strategy of picking a new item, when not all candidate items are used to build the FC scale.

Explanation for arguments (If IIAs are of concern) If you wish to use IIAs as pairing criterion, here are some arguments that might be useful. Note that **rater_chars** and **iia_weights** are ignored when **use_IIA** is FALSE.

use_IIA: Logical. Indicates whether IIA metrics are used as matching criteria.

rater_chars: Item responses for all items by a certain number of participants.

iia_weights: A vector of length 4 indicating weights given for the 4 IIA metrics, including linearly and quadratically weighted AC (Gwet, 2008; 2014) and Brennan-Prediger Index (Prennan & Prediger, 1981; Gwet, 2014). Default to a vector of all 1s.

```
# Note that this will take some time to run! (~ 1-2 minutes with this setting)

# Weights for social desirability score and item difficulty should be set to -1,
# because we don't want variance for these characteristics to be big.
result <- sa_pairing_generalized(block = initial_FC, eta_Temperature = 0.01,
                                r = 0.995, end_criteria = 10^(-6),
                                weights = char_weights,
                                item_chars = item_chars, use_IIA = TRUE,
                                rater_chars = item_responses)
```

Step 5: See how this improves over the initial one

Finally, let's see how this pairing method improves from the initial solution!

Let's first see the total energy compared to the previous one. First are the initial energy, which is identical to what we have calculated in Step 4.

```
# Initial energy with IIA
cal_block_energy_with_iia(block = result$block_initial, item_chars = item_chars,
                           weights = char_weights, rater_chars = item_responses)

#>           [,1]
#> [1,] -33.31589

# Alternative way to calculate initial energy
```

```
print(result$energy_initial)
#> [1] -33.31589
```

And the final result:

```
# Final energy with IIA
cal_block_energy_with_iaa(block = result$block_final, item_chars = item_chars,
                           weights = char_weights, rater_chars = item_responses)
#>           [,1]
#> [1,] 21.69018

# Alternative way to calculate final energy
print(result$energy_final)
#> [1] 21.69018
```

Let's take a look at how items are matched within each block. First are underlying latent traits.

This time, within each block, the three items are already coming from three distinct latent traits!

(Note: It does not guarantee that items will **ALWAYS** come from different latent traits after pairing. But if you want to increase the likelihood for such a result, you can increase the weight corresponding to item dimension)

```
knitr::kable(matrix(item_chars$DIM[t(result$block_final)], ncol = 3, byrow = TRUE))
```

Neuroticism	Openness	Extraversion
Extraversion	Neuroticism	Agreeableness
Openness	Conscientiousness	Neuroticism
Neuroticism	Extraversion	Conscientiousness
Openness	Conscientiousness	Neuroticism
Extraversion	Conscientiousness	Openness
Openness	Extraversion	Neuroticism
Conscientiousness	Openness	Neuroticism
Openness	Agreeableness	Neuroticism
Agreeableness	Extraversion	Neuroticism
Neuroticism	Openness	Extraversion
Conscientiousness	Neuroticism	Openness
Extraversion	Openness	Neuroticism
Conscientiousness	Extraversion	Openness
Agreeableness	Extraversion	Conscientiousness
Neuroticism	Agreeableness	Extraversion
Conscientiousness	Extraversion	Openness
Neuroticism	Openness	Conscientiousness
Extraversion	Openness	Conscientiousness
Neuroticism	Openness	Conscientiousness

Next let's look at difference in social desirability within each block. Item social desirability scores are much closer to each other within each block, but we see that there are still big discrepancies two blocks.

```
sd_final <- matrix(item_chars$SD_Mean[t(result$block_final)], ncol = 3, byrow = TRUE)
knitr::kable(sd_final)
```

3.493	3.518	3.481
3.469	3.501	3.449
3.549	3.429	3.478
2.306	2.382	2.354
3.417	3.418	3.499
2.372	2.327	2.383
3.470	2.375	3.458
3.575	3.474	3.429
2.344	2.299	2.283
2.275	2.269	2.314
2.346	2.326	2.294
3.540	3.524	3.492
2.383	2.331	2.257
2.386	2.291	2.408
2.275	2.430	2.296
2.367	2.410	2.371
3.476	2.367	2.381
3.515	3.455	3.526
3.501	3.484	3.549
3.628	3.551	3.521

A more intuitive way to present this: how much have we improved on the average variance for all blocks? Still, it is much lower, but we see there is always space to improve.

```
# Initial
print(mean(apply(sd_initial, 1, var)))
#> [1] 0.3317092

# Final
print(mean(apply(sd_final, 1, var)))
#> [1] 0.04195382
```

Finally we look at item difficulty. Good improvement is also observed. We see that difference in item difficulty within a block also decreases:

```
diff_final <- matrix(item_chars$DIF[t(result$block_final)], ncol = 3, byrow = TRUE)
knitr::kable(diff_final)
```

-0.7635046	-0.0490111	-0.5168748
-0.6267511	-0.9898177	-0.4484543
0.9447096	0.5910397	0.6105311
0.2507761	0.2383125	0.4304538
0.3665809	0.3986206	0.3189794
0.6598719	0.5768038	0.1906729
-0.9688906	-0.7874764	-0.7614264

0.8602120	0.9388348	0.9814553
-0.0951617	-0.5887666	-0.2743858
0.8782261	0.7562551	0.9194196
0.8652686	0.9139843	0.8647261
-0.0455446	-0.7846000	-0.1734361
-0.1775934	-0.3477869	-0.1609612
0.4033506	0.3847467	0.0105288
-0.1072381	0.0917744	0.2106074
-0.7339903	-0.8802425	-0.8770999
0.5814682	0.7372855	0.7113918
-0.4487723	-0.1364442	0.3800486
-0.2764722	-0.1896666	-0.1649317
0.3712718	0.5481408	0.6233055

How much have we improved on the average variance for all blocks in this case?

```
print(mean(apply(diff_initial, 1, var)))
#> [1] 0.4275037
print(mean(apply(diff_final, 1, var)))
#> [1] 0.04305659
```

We also list IIAs for demonstration purposes, which also improves.

```
colMeans(get_iiA(result$block_final, data = item_responses))
#>      BPlin      BPquad      AClin      ACquad
#> 0.0195425 -0.0498910 0.1263995 0.1595815
```

Step 6: Automatic pairing with a multi-step optimization process

In some cases, users may want to optimize item characteristics sequentially, rather than in a simultaneous manner. This makes sense because it is possible that simultaneous optimization will inevitably favor the improvement in one characteristic at the cost of losing the best fit for the other, as we have observed in Step 5.

Two solutions can be made to address this problem:

1. Pay careful attention to the distribution of each item characteristic and try out different weights for characteristics with different scales. Alternatively, try smaller `end_criteria` or larger `r` and `n_exchange` values to allow for more iterations to be run;
2. Use a multi-step optimization process, where some item characteristics are optimized first, then others. This involves running `sa_pairing_generalized()` several times, which each time optimizing more and more item characteristics. Those characteristics optimized will remain their weight in later stages, but have 0 weights if they are not yet optimized.

With the previous example, we show how method 2 will work, starting from `initial_FC`. First we perform optimization on latent traits:

```
FC_1 <- sa_pairing_generalized(initial_FC, eta_Temperature = 0.01,
                              r = 0.995, end_criteria = 10^(-6),
                              weights = c(1, 0, 0),
                              item_chars = item_chars, use_IIA = TRUE,
                              rater_chars = item_responses)
```

Then, we optimize based on minimizing variance in social desirability within a block.

```
FC_2 <- sa_pairing_generalized(FC_1$block_final, eta_Temperature = 0.01,
                              r = 0.995, end_criteria = 10^(-6),
                              weights = c(1, -1, 0),
                              item_chars = item_chars, use_IIA = TRUE,
                              rater_chars = item_responses)
```

Finally, we optimize based on minimizing variance in item difficulty.

```
FC_3 <- sa_pairing_generalized(FC_2$block_final, eta_Temperature = 0.01,
                              r = 0.995, end_criteria = 10^(-6),
                              weights = c(1, -1, -3),
                              item_chars = item_chars, use_IIA = TRUE,
                              rater_chars = item_responses)
```

Step 7: See how a multi-step iteration improves

First, underlying latent traits. It does look nice like what we have in Step 5.

```
knitr::kable(matrix(item_chars$DIM[t(FC_3$block_final)], ncol = 3, byrow = TRUE))
```

Extraversion	Openness	Neuroticism
Agreeableness	Extraversion	Neuroticism
Conscientiousness	Neuroticism	Openness
Extraversion	Openness	Conscientiousness
Agreeableness	Openness	Extraversion
Neuroticism	Openness	Conscientiousness
Neuroticism	Agreeableness	Extraversion
Neuroticism	Openness	Extraversion
Openness	Neuroticism	Extraversion
Openness	Neuroticism	Conscientiousness
Neuroticism	Openness	Agreeableness
Neuroticism	Openness	Conscientiousness
Conscientiousness	Extraversion	Neuroticism
Conscientiousness	Neuroticism	Openness
Extraversion	Conscientiousness	Openness
Extraversion	Agreeableness	Neuroticism
Conscientiousness	Conscientiousness	Openness
Conscientiousness	Openness	Extraversion
Openness	Neuroticism	Extraversion
Conscientiousness	Neuroticism	Extraversion

Next let's look at difference in social desirability within each block. It performs better than in Step 5, where the discrepancy within block is removed!

```
sd_FC3 <- matrix(item_chars$SD_Mean[t(FC_3$block_final)], ncol = 3, byrow = TRUE)
knitr::kable(sd_FC3)
```

2.430	2.408	2.257
2.299	2.371	2.283
3.429	3.478	3.549
2.291	2.383	2.296
2.275	2.331	2.383
3.458	3.484	3.549
2.367	2.410	2.375
3.501	3.470	3.469
2.326	2.346	2.367
3.518	3.499	3.526
3.493	3.455	3.449
3.429	3.474	3.575
3.540	3.501	3.515
3.418	3.628	3.417
2.294	2.354	2.381
2.269	2.275	2.314
3.521	3.476	3.551
2.386	2.344	2.382
3.492	3.524	3.481
2.327	2.306	2.372

As before, let's see how much have we improved on the average variance for all blocks. Improvement in social desirability is verified by seeing the decrease in variance.

```
# Initial solution
print(mean(apply(sd_initial, 1, var)))
#> [1] 0.3317092

# Simultaneous optimization
print(mean(apply(sd_final, 1, var)))
#> [1] 0.04195382

# Sequential optimization
print(mean(apply(sd_FC3, 1, var)))
#> [1] 0.002573133
```

Finally we look at item difficulty.

```
diff_fc3 <- matrix(item_chars$DIF[t(FC_3$block_final)], ncol = 3, byrow = TRUE)
knitr::kable(diff_final)
```

-0.7635046	-0.0490111	-0.5168748
-0.6267511	-0.9898177	-0.4484543
0.9447096	0.5910397	0.6105311
0.2507761	0.2383125	0.4304538
0.3665809	0.3986206	0.3189794
0.6598719	0.5768038	0.1906729
-0.9688906	-0.7874764	-0.7614264
0.8602120	0.9388348	0.9814553
-0.0951617	-0.5887666	-0.2743858
0.8782261	0.7562551	0.9194196
0.8652686	0.9139843	0.8647261
-0.0455446	-0.7846000	-0.1734361
-0.1775934	-0.3477869	-0.1609612
0.4033506	0.3847467	0.0105288
-0.1072381	0.0917744	0.2106074
-0.7339903	-0.8802425	-0.8770999
0.5814682	0.7372855	0.7113918
-0.4487723	-0.1364442	0.3800486
-0.2764722	-0.1896666	-0.1649317
0.3712718	0.5481408	0.6233055

Average variance of item difficulty also decreases!

```
# Initial solution
print(mean(apply(diff_initial, 1, var)))
#> [1] 0.4275037

# Simultaneous optimization
print(mean(apply(diff_final, 1, var)))
#> [1] 0.04305659

# Sequential optimization
print(mean(apply(diff_fc3, 1, var)))
#> [1] 0.04012038
```

How about IIAs?

```
colMeans(get_iiA(FC_3$block_final, data = item_responses))
#>      BPlin      BPquad      AClin      ACquad
#> 0.0335635 -0.0304330 0.1459820 0.1892760
```

Conclusion

In this tutorial, we have shown how the functionality of `autoFC` can be used to automatically build forced choice scales with better matches in various item characteristics.

We note that this automatic tool does not guarantee the production of a scale with best matches in these characteristic simultaneously, but instead provides solutions close to optimal among a search space where exhaustive enumeration of each solution is unrealistic.

- To further illustrate this, consider when constructing an FC scale consisting of 20 triplets with 60 items, the number of possible scales is $\frac{C_{60}^3 * C_{57}^3 * \dots}{20!} = \frac{60!}{6^{20} \times 20!} = 9.35 \times 10^{47}$.

We also note that, to produce scales with better overall fit, users are encouraged to: 1. customize their own functions for optimizing each item characteristic; 2. Try out different **weights**, **iaa_weights**, **Temperature** (**eta_Temperature**) or **r** values; 3. Use a sequential process to optimize item characteristics one by one.

Good luck!