

Basic uses of the `optmatch` package

Yevgeniya Kleyman and Ben Hansen

May 19, 2009

For case-control studies, prospective studies with a treatment and a control group, and studies of a few other types, a statistician may wish to match individuals from one predesignated group, the cases or the treatment group, to similar members of another group of control subjects in order to reduce treatment selection bias and estimate treatment effect. This is bipartite matching.

Matching algorithms fall into two categories: greedy and optimal. In each case, there is a need for a criterion of discrimination between better and worse matches. Here, it takes form of a “distance”. In greedy matching, treated units are considered in sequence, with each paired to the closest (as measured by distance) control that has not already been matched to another treated unit. This approach, also called “nearest available” matching, is easily implemented, but can result in matches that are much poorer than they need be. “Optimal” matching, on the other hand, assigns controls to treated units so as to minimize the average of distances among matched units. (The algorithms needed for optimal matching are not simply described, but they are functionally similar to sequential pairing procedures that reconsider previously made matchings at each new step.) When the number of controls is large, greedy and optimal matching often lead to very similar or same sets of matches; however optimal matching generally results in smaller distances within each pair ([Gu and Rosenbaum 1993](#))

To see the difference between greedy and optimal matching, consider the following example. Below is a table indicating the quality of potential pairings among four objects.

	<i>y</i>	<i>z</i>
<i>w</i>	A	A-
<i>x</i>	A-	F

A greedy algorithm for pair matching creates pairs sequentially, at each step choosing the pairing that then seems best, without attending to its effect on future pairings. A greedy match might begin with w , in which case w would be paired to y , as y is w 's nearest available counterpart. In and of itself this is a good pairing, as the given distance matrix rates it a "A"; but it is bad in that it leaves only z as a possible match to x . The average of its two grades is "C". A luckier implementation would begin with x , matching it to y and then w to z , for an average of "A-". Optimal matching removes luck from the picture, immediately matching x to y and w to z .

Full matching is an especially general form of optimal matching. In a full match, some matched sets may contain one treated subject (or case) alongside one or more controls, while other matched sets may contain multiple treated units (cases) alongside one control. This distinguishes full matching from both pair matching, matching with $k \geq 2$ controls, and matching with a variable number of controls, although full matchings sometimes coincide with matches produced by these simpler methods (Hansen and Klopfer 2006). Among all methods of matching for two groups, full matching alone is demonstrably optimal as a method of producing similarity with matched sets; this was shown by Rosenbaum (1991), who introduced the method. It shares these advantages with full matching with restrictions (Hansen and Klopfer 2006), a procedure that can also mimic pair matching and matching with a fixed or variable number of controls. The workhorse of "optmatch" is a function, `fullmatch()`, that performs full matching and full matching with restrictions.

The following are the functions contained in "optmatch" and their respective purposes:

```
fullmatch  Optimal full matching
pairmatch  Optimal pair matching
matched    Identification of units placed into matched sets
matchfailed Identification of units/subclasses for which matching was infeasible
unmatched  Identification of units not placed into matched sets
matched.distances
makedist
mahal.dist
pscore.dist
maxControlsCap Set thinning cap for full matching
```

`minControlsCap` Set thickening cap for full matching

You're loading `optmatch`, by Ben Hansen, a package for flexible and optimal matching. Important license information: The `optmatch` package makes essential use of D. P. Bertsekas and P. Tseng's RELAX-IV algorithm and code, as well as Bertsekas' AUCTION algorithm and code. Bertsekas and Tseng freely permit their software to be used for research purposes, but non-research uses, including the use of it to 'satisfy in any part commercial delivery requirements to government or industry,' require a special agreement with them. By extension, this requirement applies to any use of the `fullmatch()` function. (If you are using another package that has loaded `optmatch`, then you will probably be using `fullmatch` indirectly.) For more information, enter `relaxinfo()` at the command line

Optimal Full Matching

To follow this example:

1. Install the `optmatch` package
2. Install the `boot` package
3. Use command `library(optmatch)`
4. Use command `help(nuclearplants, package='optmatch')`
5. Use command `data(nuclearplants, package='optmatch')`
6. Use command `attach(nuclearplants)`
7. Use command `help(fullmatch, package='optmatch')`

The main function is of the following form:

```
fullmatch(distance, subclass.indices = NULL, min.controls = 0, max.controls  
= Inf, omit.fraction = NULL, tol = 0.01)
```

The only mandatory argument is `distance`.

distance: A matrix of nonnegative discrepancies, each indicating the permissibility and desirability of matching the unit corresponding to its row (a 'treatment') to the unit corresponding to its column (a 'control'); or a list of such matrices. Finite discrepancies indicate permissible matches, with smaller discrepancies indicating more desirable matches. Matrix 'distance', or the matrix elements of 'distance', must have row and column names.

`Optmatch` has three functions to help create the above-mentioned distance matrix. The functions are listed and explained below.

`makedist()`: A function defined with respect to data and a user-supplied function to compute desired distances.

`mahal.dist()`: Calculates Mahalanobis distances between treatment and control observations on given variables, assembling them into a discrepancy matrix (or matrices)

`pscore.dist()`: Extracts scores from a fitted propensity scoring model, assembling them into a discrepancy matrix (or matrices) from which `pairmatch()` or `fullmatch()` can determine optimal matches.

It is preferable to use these functions rather than to create a distance matrix by other means. We will illustrate later that the above methods do the work of keeping track of the ordering of observations in a working data frame, while an ordinarily created distance matrix would not.

Example. The following data relate to the 26 light water reactor (LWR) plants constructed in the U.S.A. in the late 1960's and early 1970's. The data were collected with the aim of predicting the cost of construction of further LWR plants. To illustrate the use of distance, consider the problem of matching new to refurbished nuclear plants, for the purpose of comparing their construction costs. In the nuclear plants example, seven existing-site plans are to be matched to 19 new-site plants, in order to allow an adjusted assessment of the cost of building on existing versus new sites. Variables available for use in the analysis are cost and date of issue of construction permit. For the illustration, we matched only upon cap and date, but an earnest analysis might match on other variables as well. Also, in this example, we will use only the first 26 data points in the dataset, for which the variable `pt` (a binary variable where '1' indicates those plants with partial turnkey guarantees) is set to zero.

	Existing site	
	date	capacity
A	2.3	660
B	3.0	660
C	3.4	420
D	3.4	130
E	3.9	650
F	5.9	430
G	5.1	420

date is date of construction, in years after 1965; **capacity** is net capacity of the power plant, in MWe above 400.

	New site	
	date	capacity
H	3.6	290
I	2.3	660
J	3.0	660
K	2.9	110
L	3.2	420
M	3.4	60
N	3.3	390
O	3.6	160
P	3.8	390
Q	3.4	130
R	3.9	650
S	3.9	450
T	3.4	380
U	4.5	440
V	4.2	690
W	3.8	510
X	4.7	390
Y	5.4	140
Z	6.1	730

Creating a distance matrix

The discrepancy matrix should record plants' differences on the covariates, here date of construction and capacity, but the manner in which it combines these differences is at the discretion of the analyst. To illustrate, we calculate a distance matrix from the *ranks* of the date and the capacity variables (following section 10.3.4 of (Rosenbaum 2002)). This transforms the covariates as follows:

	Existing site	
	date	capacity
A	1.5	22.5
B	4.5	22.5
C	10.0	13.5
D	10.0	3.5
E	18.0	19.5
F	25.0	15
G	23.0	12

Here, **date** is *rank of* date of construction, in years after 1965, and **capacity** is *rank of* net capacity of the power plant, in MWe above 400.

	New site	
	date	capacity
H	13.5	7.0
I	1.5	22.5
J	4.5	22.5
K	3.0	2.0
L	6.0	13.5
M	10.0	1.0
N	7.0	11.0
O	13.5	6.0
P	15.5	10.0
Q	10.0	3.5
R	18.0	19.5
S	18.0	17.0
T	10.0	8.0
U	21.0	16.0
V	20.0	25.0
W	15.5	18.0
X	22.0	9.0
Y	24.0	5.0
Z	26.0	26.0

Now we create a matrix whose entries are the total differences between the ranks

of dates and capacities for each pair. For example, the difference for the pair AH is $(13.5-1.5)+(22.5-7.0) = 27.5$ (which we round to 28, for simplicity).

This matrix, which we name `plant.dist`, is formed as follows:

Creating a distance matrix: direct method

```
> attach(nuclearplants)
> plant.dist <- round(abs(outer(rank(cap)[pr == 1], rank(cap)[pr ==
+ 0], "-")) + abs(outer(rank(date)[pr == 1], rank(date)[pr ==
+ 0], "-")))
> dimnames(plant.dist) <- list(row.names(nuclearplants)[pr == 1],
+ row.names(nuclearplants)[pr == 0])
> detach()
> plant.dist
```

	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	d	e	f
A	36	0	6	31	20	39	24	38	34	36	22	28	30	32	24	22	41	49	31	22	18	9
B	30	6	0	28	14	33	18	32	28	30	16	22	24	26	18	16	35	43	25	28	24	10
C	13	24	18	22	4	16	6	14	10	13	17	12	6	14	24	13	18	26	32	22	17	14
D	7	36	30	8	17	2	12	6	14	0	30	24	6	26	38	26	20	16	44	18	26	28
E	23	22	16	38	21	32	24	24	16	30	0	6	24	10	8	4	18	26	14	38	34	20
F	22	37	31	38	20	32	23	24	16	30	14	9	23	5	18	16	10	14	15	38	34	28
G	17	38	32	32	19	26	18	18	10	24	16	10	18	6	20	17	4	10	20	32	28	30
a	28	9	10	22	10	30	15	28	24	28	20	19	21	23	28	16	32	40	35	18	14	0
b	18	22	28	13	18	21	15	20	22	18	38	33	15	35	46	34	28	31	53	0	8	18
c	28	9	10	22	10	30	15	28	24	28	20	19	21	23	28	16	32	40	35	18	14	0

Creating a distance matrix using convenience functions

Alternatively, we can use an `optmatch` convenience function, such as `mahal.dist` or `pscore.dist` a distance.

```
> plant.dist.m <- mahal.dist(pr ~ cap + date, nuclearplants)
> plant.dist.m
```

	H	I	J	K	L	M	N	O	P	Q	R	S	T	U
A	5.60	0.00	0.42	8.80	2.400	12.00	3.100	8.700	4.10	9.2	2.40	3.80	3.500	5.90
B	4.30	0.42	0.00	8.40	1.700	10.00	2.200	7.500	2.70	8.2	0.81	2.10	2.500	3.60
C	0.54	2.80	1.80	2.80	0.059	3.70	0.032	1.900	0.13	2.4	1.60	0.25	0.054	1.10
D	0.70	9.20	8.20	0.24	2.400	0.15	1.900	0.048	2.00	0.0	7.70	3.00	1.700	3.80
E	3.70	2.40	0.81	8.80	1.900	9.90	2.100	6.700	1.90	7.7	0.00	1.10	2.200	1.50
F	5.60	14.00	9.80	11.00	7.100	9.40	6.300	7.000	4.50	8.1	5.30	3.80	5.900	1.90

```

G 2.60  9.00 5.90  6.80 3.400  6.10 2.900 3.900 1.70 4.8 2.80 1.30 2.600 0.34
a 1.70  1.20 0.91  3.90 0.230  5.50 0.490 3.500 1.10 3.9 1.80 1.20 0.670 2.70
b 1.80  2.80 3.30  1.90 0.940  3.70 1.100 2.700 2.20 2.6 5.10 2.90 1.300 5.00
c 1.70  1.20 0.91  3.90 0.230  5.50 0.490 3.500 1.10 3.9 1.80 1.20 0.670 2.70
  V    W    X    Y    Z    d    e    f
A 3.2 2.60 7.50 17.0 13.0  2.80  1.70 1.20
B 1.3 1.20 4.90 14.0  9.0  3.30  2.20 0.91
C 2.5 0.32 1.50  6.2  9.1  1.40  1.50 0.45
D 9.1 4.10 3.20  3.8 16.0  2.60  3.90 3.90
E 0.1 0.54 2.50  9.6  4.5  5.10  4.20 1.80
F 4.9 4.70 1.50  2.5  2.5 13.00 13.00 9.10
G 2.9 1.90 0.19  2.4  3.5  7.60  8.00 4.90
a 2.8 0.81 3.50  9.9 11.0  0.84  0.52 0.00
b 6.6 2.80 5.50 11.0 17.0  0.00  0.17 0.84
c 2.8 0.81 3.50  9.9 11.0  0.84  0.52 0.00

```

Creating an Unrestricted Match

Having created the distance matrix, we are in a position to match new to refurbished plants. Enter

```
> plants.fm <- fullmatch(plant.dist)
```

The result is:

```

  a    A    b    B    c    C    d    D    e    E    f    F    G    H    I    J
m.1 m.2 m.3 m.4 m.1 m.6 m.3 m.7 m.3 m.8 m.1 m.9 m.10 m.7 m.2 m.4
  K    L    M    N    O    P    Q    R    S    T    U    V    W    X    Y    Z
m.7 m.6 m.7 m.6 m.7 m.6 m.7 m.8 m.8 m.7 m.9 m.8 m.8 m.10 m.10 m.8

```

where m.1, m.2, etc. are the indices for matches. `plants.fm` pairs A to I, B- to J, etc.

An advantage of the convenience functions for making distances

Compare `plants.fm` to

```
> (plants.fm.m <- fullmatch(plant.dist.m))
```

```

  H    I    A    J    B    K    L    M    C    N    O    P    Q    R    S    T
m.6 m.2 m.2 m.4 m.4 m.7 m.6 m.7 m.6 m.6 m.7 m.6 m.7 m.8 m.6 m.6
  U    D    V    E    W    F    X    G    Y    Z    d    e    f    a    b    c
m.10 m.7 m.8 m.8 m.6 m.9 m.10 m.10 m.10 m.9 m.3 m.3 m.1 m.1 m.3 m.1

```

The matches are different, as one would expect, given that the distances they're based on are different. However, you'll not that the orderings of observations they

assume are different, as is indicated by the difference in their row names. The match based on `plants.fm.m` is ordered as is the data frame it's based on, `nuclearplants`. The ordering of `plants.fm`, in contrast, is incorrect. Looking at the first six rows of our dataset below, it can be seen that the actual order of plants differs from the order of the `fullmatch` object.

	cost	date	t1	t2	cap	pr	ne	ct	bw	cum.n	pt
H	460.05	68.58	14	46	687	0	1	0	0	14	0
I	452.99	67.33	10	73	1065	0	0	1	0	1	0
A	443.22	67.33	10	85	1065	1	0	1	0	1	0
J	652.32	68.00	11	67	1065	0	1	1	0	12	0
B	642.23	68.00	11	78	1065	1	1	1	0	12	0
K	345.39	67.92	13	51	514	0	1	1	0	3	0

That's because the distance object `plant.dist.m` carries the metadata needed to reconstruct the proper ordering after matching, whereas the matrix `plant.dist` does not, forcing `fullmatch()` to guess the proper ordering of row names. (It does this by sorting on the row names, which sometimes gets the right answer but doesn't here.)

Matching with Restrictions

Matching within Calipers

The statistician may wish to forbid the matching of certain treatment-control pairs, perhaps those that are quite dissimilar. On the other hand, it may be too restrictive to insist that matched units have the same covariates: for instance, in the nuclear plants example only four exactly matched pairs are possible, namely AI, BJ, DQ, and ER. For continuous covariates, matching within calipers enforces similarity of matches (if not that they be exactly the same) ([Hansen and Klopfer 2006](#)).

Continuing the nuclear plants illustration, we impose a caliper of 3 years in the date of construction. This forbids matches between plants whose date of construction differs by more than three years. Note that this may not minimize total distance, since total distance is a combination of date and capacity.

Coding a caliper when using `optmatch` and `R`:

To incorporate such a caliper, invoke `fullmatch()` after modifying the distance matrix as follows,

```
> attach(nuclearplants)
> datediffs <- outer(date[pr == 1], date[pr == 0], "-")
> (plant.dist <- plant.dist/(abs(datediffs) < 3))
```

	H	I	J	K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z	d	e	f
A	36	0	6	31	20	39	24	38	34	36	22	28	30	32	24	22	41	Inf	Inf	22	18	9

B	30	6	0	28	14	33	18	32	28	30	16	22	24	26	18	16	35	43	Inf	28	24	10
C	13	24	18	22	4	16	6	14	10	13	17	12	6	14	24	13	18	26	32	22	17	14
D	7	36	30	8	17	2	12	6	14	0	30	24	6	26	38	26	20	16	44	18	26	28
E	23	22	16	38	21	32	24	24	16	30	0	6	24	10	8	4	18	26	14	38	34	20
F	22	Inf	31	Inf	20	32	23	24	16	30	14	9	23	5	18	16	10	14	15	Inf	Inf	Inf
G	17	38	32	32	19	26	18	18	10	24	16	10	18	6	20	17	4	10	20	32	28	30
a	28	9	10	22	10	30	15	28	24	28	20	19	21	23	28	16	32	40	Inf	18	14	0
b	18	22	28	13	18	21	15	20	22	18	38	33	15	35	46	34	28	Inf	Inf	0	8	18
c	28	9	10	22	10	30	15	28	24	28	20	19	21	23	28	16	32	40	Inf	18	14	0

```
> plant.dist.m <- plant.dist.m/(abs(datediffs) < 3)
```

The infinite entries occur where the matrix `datediffs` has an entry less than -3 or greater than $+3$, indicating an unwanted match.

Maintaining balance between treatment and control using `min.controls` and `max.controls`

The statistician may wish to place restrictions on the relative proportions with which treated and control subjects are combined into matched sets, perhaps to control the variability of an estimate based on the matching (See [Hansen 2004](#), §3). The following are two arguments of the `fullmatch()` function that allow that adjustment.

min.controls: The minimum ratio of controls to treatments that is to be permitted within a matched set: should be nonnegative and finite. If `min.controls` is not a whole number, the reciprocal of a whole number, or zero, then it is rounded down to the nearest whole number or reciprocal of a whole number.

max.controls: The maximum ratio of controls to treatments that is to be permitted within a matched set: should be positive and numeric. If `max.controls` is not a whole number, the reciprocal of a whole number, or `Inf`, then it is rounded up to the nearest whole number or reciprocal of a whole number.

Specifying `min.controls` and/or `max.controls` amounts to *full matching with restrictions*, which improves matched sets' treatment-control balance.

To determine the largest value of `min.controls` with which the matching problem is possible to solve, use `minControlsCap`. To get the smallest feasible value of `max.controls`, use `maxControlsCap`. (Note: these function calls `fullmatch` repeatedly, in a line search of the positive half-line. They can be time-consuming in large matching problems.)

```
> (mincc <- minControlsCap(plant.dist))
```

```

$strictest.feasible.min.controls
m
2

$given.max.controls
m
Inf

> maxControlsCap(plant.dist, min.controls = mincc$strictest.feasible)

$given.min.controls
m
2

$strictest.feasible.max.controls
m
3

```

Omitting a fraction of controls

The statistician may wish to discard a portion of the control data. This may be done if a cost issue arises or if the statistician would like to discard outliers on the distance.

omit.fraction: Optionally, specify what fraction of controls or treated subjects are to be rejected. If 'omit.fraction' is a positive fraction less than one, then 'fullmatch' leaves up to that fraction of the control reservoir unmatched. If omit.fraction is a negative number greater than -1, then fullmatch leaves up to |omit.fraction| of the treated group unmatched. Positive values are only accepted if max.controls >= 1; negative values, only if min.controls <= 1. If omit.fraction is not specified, then only those treated and control subjects without permissible matches among the control and treated subjects, respectively, are omitted.

Suppose we wish to omit 50% of the control subjects:

```
> (plants.fm2 <- fullmatch(plant.dist, omit.fraction = 0.5))
```

a	A	b	B	c	C	d	D	e	E	f	F	G	H	I	J
m.1	m.2	m.3	m.4	m.1	m.6	m.3	m.7	<NA>	m.8	m.1	m.9	m.10	<NA>	m.2	m.4
K	L	M	N	O	P	Q	R	S	T	U	V	W	X	Y	Z
<NA>	m.6	m.7	<NA>	<NA>	<NA>	m.7	m.8	<NA>	<NA>	m.9	<NA>	m.8	m.10	<NA>	<NA>

```
> table(plants.fm2, ifelse(pr, "treated", "control"))
```

```
plants.fm2 control treated
m.1          2         1
m.10         1         1
m.2          2         0
m.3          1         1
m.4          2         0
m.6          1         1
m.7          3         0
m.8          2         1
m.9          2         0
```

If before we had a total of 19 controls, of them all were matched, now we have matched only 10 of them. The first nine rows of the above table are not matched sets, but rather indicators for unmatched controls. Controls H, K, N, O, P, T, V, Y, and Z are now not matched to a treated subject.

Rounding

If the matching is taking too long, increasing the `tol` argument may speed up the process.

tol: Because of internal rounding, 'fullmatch' may solve a slightly different matching problem than the one specified, in which the match generated by `fullmatch` may not coincide with an optimal solution of the specified problem. `tol` specifies the extent to which `fullmatch`'s output is permitted to differ from an optimal solution to the original problem, as measured by the sum of discrepancies for all treatments and controls placed into the same matched sets.

Evaluating the quality of the match

Treatment-control balance by matched set

stratumStructure: Tabulate treatment:control ratios occurring in matched sets, and the frequency of their occurrence.

```
> stratumStructure(plants.fm)

2:1 1:1 1:2 1:3 1:5 1:6
  1   3   2   1   1   1
attr(,"comparable.num.matched.pairs")
[1] 11.88095
```

```
> table(plants.fm, ifelse(pr, "treated", "control"))
```

```
plants.fm control treated
  m.1          2         1
  m.10         1         2
  m.2          2         0
  m.3          1         2
  m.4          2         0
  m.6          1         3
  m.7          7         0
  m.8          4         2
  m.9          2         0
```

In this unrestricted match, the number of controls per treatment subject varies widely. Such imbalance can be prevented as follows.

```
> plants.fm1 <- fullmatch(plant.dist, min.controls = 2, max.controls = 3)
> stratumStructure(plants.fm1)
```

```
1:2 1:3
  8   2
attr(,"comparable.num.matched.pairs")
[1] 13.66667
```

Matched treatment-control distances:

```
> unlist(matched.distances(plants.fm1, plant.dist))
```

m.1.f	m.1.W	m.10.X	m.10.Y	m.2.e	m.2.I	m.3.d	m.3.K	m.4.J	m.4.V	m.5.L
0	16	4	10	18	0	0	13	0	18	10
m.5.N	m.6.H	m.6.P	m.6.T	m.7.M	m.7.O	m.7.Q	m.8.R	m.8.S	m.9.U	m.9.Z
15	13	10	6	2	6	0	0	6	5	15

Matched treatment-control distances, by matched set:

```
> matched.distances(plants.fm1, plant.dist, TRUE)
```

```
$m.1
  f  W
a 0 16
```

```
$m.10
  X  Y
G 4 10
```

```
$m.2
  e I
A 18 0
```

```
$m.3
  d K
b 0 13
```

```
$m.4
  J V
B 0 18
```

```
$m.5
  L N
c 10 15
```

```
$m.6
  H P T
C 13 10 6
```

```
$m.7
  M O Q
D 2 6 0
```

```
$m.8
  R S
E 0 6
```

```
$m.9
  U Z
F 5 15
```

For example, in m.6, we see that the biggest difference in the dates of F,U, and Z is the difference between U and Z: $71.08 - 69.50 = 1.58$.

Matched, Unmatched, and Matchfailed

An easy way to check which plants were matched and unmatched is to use commands `matched()` and `unmatched()`. In our original matching attempt, full matching was performed so we see:

```
> matched(plants.fm)
```

```
[1] TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE
[16] TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE TRUE
[31] TRUE TRUE
```

```
> table(unmatched(plants.fm))
```

```
FALSE
 32
```

After we omitted half of our controls, however, the output is

```
> table(matched(plants.fm2))
```

```
FALSE  TRUE
 11     21
```

```
> table(unmatched(plants.fm2))
```

```
FALSE  TRUE
 21     11
```

Names of the controls that were not matched can be retrieved as follows:

```
> names(plants.fm2)[unmatched(plants.fm2)]
```

```
[1] "e" "H" "K" "N" "O" "P" "S" "T" "V" "Y" "Z"
```

When `fullmatch` has been presented with an inconsistent combination of constraints and discrepancies between potential matches, so that there exists no matching (i) with finite total discrepancy within matched sets that (ii) respects the given constraints, then the matching problem is said to be infeasible. TRUEs in the output of `matchfailed` indicate that this has occurred.

```
> table(matchfailed(fullmatch(plant.dist, min.controls = 5)))
```

```
TRUE
 32
```

To understand the output of `matchfailed` element-wise, note that `fullmatch` handles a matching problem in three steps. First, if `fullmatch` has been directed to match within subclasses, then it divides its matching problem into a subproblem for each subclass. Second, `fullmatch` removes from each subproblem those individual units that lack permissible potential matches (i.e. potential matches from which they are separated by a finite discrepancy). Such “isolated” units are flagged in such a

way as to be indicated by `unmatched`, but not by `matchfailed`. Third, `fullmatch` presents each subproblem, with isolated elements removed, to an optimal matching routine. If such a reduced subproblem is found at this stage to be infeasible, then each unit contributing to it is so flagged as to be indicated by `matchfailed`.

In this case there were no inconsistencies in the constraints, so the command returns:

```
> table(matchfailed(plants.fm))
```

```
FALSE
```

```
32
```

```
> detach()
```

References

- X.S. Gu and Paul R. Rosenbaum. Comparison of multivariate matching methods: Structures, distances, and algorithms. *Journal of Computational and Graphical Statistics*, 2(4):405–420, 1993.
- Ben B. Hansen. Full matching in an observational study of coaching for the SAT. *Journal of the American Statistical Association*, 99(467):609–618, September 2004.
- Ben B. Hansen and Stephanie Olsen Klopfer. Optimal full matching and related designs via network flows. *Journal of Computational and Graphical Statistics*, 15(3):609–627, 2006. URL <http://www.stat.lsa.umich.edu/%7Ebbh/hansenKlopfer2006.pdf>.
- Paul R. Rosenbaum. A characterization of optimal designs for observational studies. *Journal of the Royal Statistical Society*, 53:597– 610, 1991.
- Paul R. Rosenbaum. *Observational Studies*. Springer-Verlag, second edition, 2002.