Convergence rate examples and theory

Simen Gaure

Abstract. If you use **lfe** for various tasks, you will notice that some estimations converge fast, whereas others converge slowly. Convergence rate of the methods used by **lfe** is not a walk in the park. Here are some examples.

1. Introduction

The method employed by **lfe** is *the method of alternating projections* ([**6**]). The link to this method comes from viewing demeaning of a vector as a *projection* ([**3**]), i.e. a linear operator P with the property $P^2 = P = P^t$. Also, the Kaczmarz-method used by **lfe** to solve the sparse resulting system is a variant of alternating projections, though P is then an affine projection with $P^2 = P$, not a linear operator. That is, if the projection demeaning factor i is P_i, their intersection which demeans all factors is $\lim_{n\to\infty}(P_1 P_2 \cdots P_e)^n$. We can therefore iterate the operator $T = P_1 P_2 \cdots P_e$ on a vector x: $T^n x$, to demean x.

The convergence rate has been analysed in [**2**] and references cited therein; there are also newer elaborations, e.g. in [**1**]. The convergence rate theory is in terms of generalized angles between subspaces, i.e. the ranges of the projections. The angle concept may not be immediately intuitive to practioners of linear regression methods, so let's have a look at some examples.

2. Examples

Our first example has two factors, they are independent of each other, and of quite high cardinality

```
> library(lfe)
> set.seed(42)
> x <- rnorm(100000)
> f1 <- sample(10000,length(x),replace=TRUE)
> f2 <- sample(10000,length(x),replace=TRUE)
> y <- x + cos(f1) + log(f2+1) + rnorm(length(x), sd=0.5)
```

We time the first step:

```
> system.time(est <- felm(y ~ x + G(f1) + G(f2)))
   user  system elapsed
  2.724   0.016   2.489
```

and the second step

```
> system.time(alpha <- getfe(est))
   user  system elapsed
  1.100   0.016   1.119
```

We see that there's nothing to complain about in terms of speed. After all, there are 20,000 dummies in this model.

Now we let `f2` have fewer levels:

```
> f2 <- sample(300,length(x),replace=TRUE)
> y <- x + cos(f1) + log(f2+1) + rnorm(length(x), sd=0.5)
> system.time(est <- felm(y ~ x + G(f1) + G(f2)))
   user  system elapsed
  3.112   0.004   2.871
> system.time(alpha <- getfe(est))
   user  system elapsed
  0.276   0.008   0.286
```

Not much happens, the second step is apparently faster, whereas the first is about the same. Note that much of the time in the first step is spent in mundane bookkeeping such as creating a model matrix, not in the demeaning as such.

Now we make the fixed effects dependent. We do this by ensuring that for each value of `f1` there are at most 20 different values of `f2`. We keep the number of levels in `f2` at 300, as in the previous example. The size of this problem is exactly the same as in the previous example, but the factor structure is very different.

```
> f2 <- (f1 + sample(20,length(x),replace=TRUE)) %% 300
> y <- x + cos(f1) + log(f2+1) + rnorm(length(x), sd=0.5)
> system.time(est <- felm(y ~ x + G(f1) + G(f2)))
   user  system elapsed
  5.896   0.000   4.319
> system.time(alpha <- getfe(est))
   user  system elapsed
  0.493   0.000   0.493
```

Then we modify it more, now having only 5 different values of `f2` for each `f1`. This is below the average group size in `f1`, which is 10.

```
> f2 <- (f1 + sample(5,length(x),replace=TRUE)) %% 300
> y <- x + cos(f1) + log(f2+1) + rnorm(length(x), sd=0.5)
> system.time(est <- felm(y ~ x + G(f1) + G(f2)))
   user  system elapsed
 63.772   0.008  33.355
> system.time(alpha <- getfe(est))
   user  system elapsed
  4.824   0.000   4.830
> nlevels(est[['cfactor']]) # number of connected components
[1] 1
```

The estimation now takes an order of magnitude more time. Indeed, in this case, it is worthwile to consider coding `f2` as 300 dummies, i.e. as an ordinary factor, even though there now are ≈ 300 vectors to centre, as opposed to two in the previous examples.

```
> system.time(est <- felm(y ~ x + G(f1) + factor(f2)))
   user  system elapsed
 14.397   1.820   9.654
> system.time(alpha <- getfe(est))
   user  system elapsed
  0.180   0.004   0.183
```

We could be tempted to believe that this is the whole story, but it's not. We may create another example, where each `f1` only has 5 different `f2`'s as above, but where the overlap is different. In the above example, an observation with `f1=1` will have `f2` drawn from {2,3,4,5,6}, whereas an observation with `f1=2` will have `f2` in {3,4,5,6,7} and so on, i.e. a considerable overlap. We reduce this overlap by creating `f2` with a little twist, by introducing some unevenly sized "holes" by cubing the samples of `1:5`.

```
> f2 <- (f1 + sample(5,length(x),replace=TRUE)^3) %% 300
> y <- x + cos(f1) + log(f2+1) + rnorm(length(x), sd=0.5)
> system.time(est <- felm(y ~ x + G(f1) + G(f2)))
   user  system elapsed
  2.808   0.000   2.578
> system.time(alpha <- getfe(est))
   user  system elapsed
  0.244   0.004   0.249
> nlevels(est[['cfactor']])
[1] 1
```

We are now back at approximately the same fast convergence as before. That is, convergence rate is not a simple function of data size, counts of levels and so on, even when we have only two factors.

2.1. Why? In the examples above we have played with the structure of the bipartite graph mentioned in [**3**], indicating that convergence rate is not a function solely of the average degree and number of connected components of the graph. Intuitively (to the author), the graph determines the convergence rate, at least in the operator norm topology, but it is unclear to me whether the rate can be described in simple graph theoretic terms, let alone in terms of some intuitive relations between dummy variables.

One could speculate that convergence is related to how tightly connected the graph is, one such measure is the diameter of the graph, or the set of lengths of shortest paths. We don't compute it exactly in this autogenerated document, it takes too long time. Instead we get an indication by taking a sample of the shortest paths. However, to avoid making this package dependent on the package **igraph** (where the necessary tools are, but we only use it in this vignette), we just outline the steps:

```
> mkgraph <- function(f1,f2)
+   graph.adjacency(tcrossprod(rBind(as(factor(f1),'sparseMatrix'),
+                                    as(factor(f2),'sparseMatrix')))>0,
+                   'undirected',diag=FALSE)
> f2 <- sample(10000,length(x),replace=TRUE)
> fivenum(shortest.paths(mkgraph(f1,f2),
```

```
+                         v=sample(20000,10),to=sample(20000,10)))
> f2 <- (f1 + sample(5,length(x),replace=TRUE)^3) %% 300
> fivenum(shortest.paths(mkgraph(f1,f2),
+                         v=sample(10300,10), to=sample(10300,10)))
> f2 <- (f1 + sample(5,length(x),replace=TRUE)) %% 300
> fivenum(shortest.paths(mkgraph(f1,f2),
+                         v=sample(10300,10), to=sample(10300,10)))
```

Loosely speaking, data with six degrees of separation(the two first; shortest path median is around 5) seems to converge faster than less well connected networks(the last; shortest path median is 43). However, the author has not had the time to attempt to prove such an assertion.

Note that convergence rate in practice is also a function of the vector to be centred, this can be seen from the trivial example when the vector is already centred and convergence is attained after a single iteration.

It is a sad fact of life, shown in [**2**], that when there are more than two factors, the convergence rate can not be determined solely by considering the angles of *pairs* of subspaces. In [**1**] a generalized angle between a finite set of subspaces is introduced for the purpose of analysis. It is left to the reader to imagine some simple examples with more than two factors. A real example with three factors can be found in [**9**], and real examples with up to five factors are treated in the appendix of [**3**] as well as in [**7**].

3. Acceleration schemes

One point to note is that with more than two factors, even their *order* may be important for the convergence rate, this has been utilized to speed up convergence in some cases, picking a random order of the projections for each iteration, or sampling randomly from the projections. One such scheme is used in [**8**], and randomization is also treated in [**1**] and the papers they cite. Random shuffling of the equations is now used in the Kaczmarz-step of **lfe**, i.e. in `getfe`, as suggested in [**3**].

Another acceleration scheme is found in [**4**], where a line search method is used in each iteration. I.e. when a single iteration transforms our vector $\eta \mapsto \eta'$, it is fairly easy to find the exact point on the line $\{\eta + t(\eta' - \eta) \,:\, t \in \mathbb{R}\}$ which is closest to the solution. A future version of **lfe** will probably utilize a combination of random ordering of factors and this line search method.

After publishing [**3**], the author was made aware that exactly the same method of alternating projections was arrived at in [**5**, p. 637], though as a technical simplification of the Gauss-Seidel method of iterated regressions. It is possible that acceleration schemes for the Gauss-Seidel method may be applicable, though the author has not yet looked into it.

References

[1] C. Badea, S. Grivaux, and V. Müller, *The rate of convergence in the method of alternating projections*, St. Petersburg Mathematical Journal **23** (2012), 413–434, preprint in arXiv:1006.2047.

[2] F. Deutsch and H. Hundal, *The Rate of Convergence for the Method of Alternating Projections, II*, J. Math. Anal. App. **205** (1997), no. 2, 381–405.

[3] S. Gaure, *OLS with Multiple High Dimensional Category Variables*, Computational Statistics and Data Analysis **66** (2013), 8–18.

[4] W.B. Gearhart and M. Koshy, *Acceleration schemes for the method of alternating projections*, Journal of Computational and Applied Mathematics **26** (1989), no. 3, 235–249.

[5] P. Guimarães and P. Portugal, *A simple feasible procedure to fit models with high-dimensional fixed effects*, Stata Journal **10** (2010), no. 4, 628–649(22).
[6] I. Halperin, *The Product of Projection Operators*, Acta Sci. Math. (Szeged) **23** (1962), 96–99.
[7] S. Markussen and K. Røed, *Social insurance networks*, IZA Discussion Papers 6446, Institute for the Study of Labor (IZA), March 2012.
[8] T. Strohmer and R. Vershynin, *A Randomized Kaczmarz Algorithm with Exponential Convergence*, Journal of Fourier Analysis and Applications **15** (2009), 262–278.
[9] S.M. Torres, P. Portugal, J.T. Addison, and P. Guimarães, *The Sources of Wage Variation: A Three-Way High-Dimensional Fixed Effects Regression Model.*, IZA Discussion Paper 7276, Institute for the Study of Labor (IZA), March 2013.

Ragnar Frisch Centre for Economic Research, Oslo, Norway