

Power project

Elias

May 10, 2018

```
library(RWordPress)
library(knitr)
#knit2wp("~/power/scripts/power_2/power_2.Rmd",title="Part I: Getting old? You can still lift!")
```

Powerful Data Project Part I: Getting old? You can still lift!

Aging is normal. Unfortunately, it leaves some powerlifters discouraged about their abilities to still lift well and accomplish PRs when they get older. How important is your age when it comes to the amount of weight you lift? Should you start focusing more on endurance and give up on lifting good weights once you hit 40...50...60? Let's take a stab at some of these questions and dig through a rather **powerful** dataset that has some pointers for us.

The first thing that I will do is getting the data and all the necessary packages that I need from R to analyze it. And look at what I have.

```
library(lme4)
```

```
## Loading required package: Matrix
```

```
library(tidyr)
```

```
##
```

```
## Attaching package: 'tidyr'
```

```
## The following object is masked from 'package:Matrix':
```

```
##
```

```
## expand
```

```
library(RColorBrewer)
```

```
library(ggplot2)
```

```
library(dplyr)
```

```
##
```

```
## Attaching package: 'dplyr'
```

```
## The following objects are masked from 'package:stats':
```

```
##
```

```
## filter, lag
```

```
## The following objects are masked from 'package:base':
```

```
##
```

```
## intersect, setdiff, setequal, union
```

```
library(gtools)
```

```
library(tinytex)
```

```
open<-read.csv("~/power/data/open/openpowerlifting2.csv",head=TRUE,stringsAsFactors = FALSE) #reading i
# open<-read.csv("~/power/data/open/openpowerlifting.csv",head=TRUE,stringsAsFactors = FALSE) #reading
# meets<-read.csv("~/power/data/open/meets.csv",head=TRUE,stringsAsFactors = FALSE) #reading in the one
# all<-merge(open,meets,by="MeetID") #merging them by column (MeetID, the only repeated column)
```

```
###
```

```
str(open) #getting the structure of the new document
```

```
## 'data.frame': 601229 obs. of 38 variables:
## $ MeetID : int 0 0 0 0 0 0 0 0 0 0 ...
## $ LifterID : int 1 2 2 2 3 4 5 5 6 6 ...
## $ Name : chr "Angie Belk Terry" "Dawn Bogart" "Dawn Bogart" "Dawn Bogart" ...
## $ Sex : chr "F" "F" "F" "F" ...
## $ Event : chr "SBD" "SBD" "SBD" "B" ...
## $ Equipment : chr "Wraps" "Single-ply" "Single-ply" "Raw" ...
## $ Age : int 47 42 42 42 18 28 60 60 52 52 ...
## $ Division : chr "Mst 45-49" "Mst 40-44" "Open Senior" "Open Senior" ...
## $ BodyweightKg : int 60 59 59 59 64 62 67 67 66 66 ...
## $ WeightClassKg : chr "60" "60" "60" "60" ...
## $ Squat1Kg : num 38.6 120.2 120.2 NA NA ...
## $ Squat2Kg : num 47.6 136.1 136.1 NA NA ...
## $ Squat3Kg : num -54.4 142.9 142.9 NA NA ...
## $ Squat4Kg : num NA NA NA NA NA ...
## $ BestSquatKg : num 47.6 142.9 142.9 NA NA ...
## $ Bench1Kg : num 15.9 88.5 88.5 88.5 29.5 ...
## $ Bench2Kg : num 20.4 95.2 95.2 95.2 31.8 ...
## $ Bench3Kg : num -24.9 -97.5 -97.5 -97.5 -34 ...
## $ Bench4Kg : num NA NA NA NA NA NA NA NA NA NA ...
## $ BestBenchKg : num 20.4 95.2 95.2 95.2 31.8 ...
## $ Deadlift1Kg : num 61.2 136.1 136.1 NA 90.7 ...
## $ Deadlift2Kg : num 70.3 149.7 149.7 NA -97.5 ...
## $ Deadlift3Kg : num -77.1 163.3 163.3 NA NA ...
## $ Deadlift4Kg : num NA NA NA NA NA NA NA NA NA NA ...
## $ BestDeadliftKg : num 70.3 163.3 163.3 NA 90.7 ...
## $ TotalKg : num 138.3 401.4 401.4 95.2 122.5 ...
## $ Place : chr "1" "1" "1" "1" ...
## $ Wilks : num 155 456 456 108 130 ...
## $ McCulloch : num 168 466 466 110 138 ...
## $ MeetPath : chr "365strong/1601" "365strong/1601" "365strong/1601" "365strong/1601" ...
## $ Federation : chr "365Strong" "365Strong" "365Strong" "365Strong" ...
## $ Date : chr "2016-10-29" "2016-10-29" "2016-10-29" "2016-10-29" ...
## $ MeetCountry : chr "USA" "USA" "USA" "USA" ...
## $ MeetState : chr "NC" "NC" "NC" "NC" ...
## $ MeetTown : chr "Charlotte" "Charlotte" "Charlotte" "Charlotte" ...
## $ MeetName : chr "Junior & Senior National Powerlifting Championships" "Junior & Senior National Powerlifting Championships" ...
## $ Group : chr "45-49" "40-44" "40-44" "40-44" ...
## $ GroupKG : chr "56-65" "56-65" "56-65" "56-65" ...
```

Obviously this is a fantastic and well kept dataset! **Great job** openpowerlifting.com!

I want to draw your attention that there are so many columns in this dataset, each containing unique information about each lifter. This will give us a lot of power to draw some observations about patterns of skills and potentials of the lifters. First I want to start by breaking them up into age ranges. I don't like the ones drawn by more federations for my purposes. I want to have a bit higher resolution, so I will break them up into 5 year age gaps and see how things go from there.

I am also breaking them up into men and women for physiological reasons. I will be asking questions about max lift strength, so I don't want to obscure the findings for either of those groups.

```
#Creating a vector of ages in increments of 5
# group<-seq(from=range(all$Age,na.rm=TRUE)[[1]],
```

```

#           to=range(all$Age,na.rm=TRUE)[[2]],
#           by=5)
#
# group_range<-c() #creating a vector that will hold those ranges
# group2<-c(group-1,100) #creating a second vector that has 100 as the maximum of it
# for(i in 1:length(group)){ #creating teh group-range vector that will have info about a range, rather
#   if(i==length(group)){group_range[i]<-paste(c(group[i], "100"),collapse="-")
#   next()
#   }
#   group_range[i]<-(paste(c(group[i],group2[i+1]),collapse="-"))
# } #making these into ranges rather than just the minimums
#
# #Run the following only once as it takes a while (20 mins)
# all<-cbind(all,Group=NA) # Adding a column of Group to the big dataset
# all$Age<-round(all$Age) #rounding age because the following code doesn't deal well with .5's and they
#
# #adding a column of ranges
# for(i in 1:length(all$Age)){
#   if(is.na(all$Age[i])){next()} #remove this estimation from all NAs
#   for(j in 1:length(group)){
#     if(all$Age[i]>=group[j]&&all$Age[i]<=group2[j+1]){ #iterate over groups and find which group Age t
#     all$Group[i]<-group_range[j] #assign the necessary group range to that age
#     }
#   }
# }
#
# #Adding weight groups (all freaking federations have different weightclasses so it's easier to actual
# grouppkg<-seq(from=range(all$BodyweightKg,na.rm=TRUE)[[1]],
#               to=range(all$BodyweightKg,na.rm=TRUE)[[2]],
#               by=10)
# grouppkg<-round(grouppkg)
#
# grouppkg_range<-c() #creating a vector that will hold those ranges
# grouppkg2<-c(grouppkg-1,260) #creating a second vector that has 260 as the maximum of it
# for(i in 1:length(grouppkg)){ #creating teh group-range vector that will have info about a range, rath
#   if(i==length(grouppkg)){grouppkg_range[i]<-paste(c(grouppkg[i], "260"),collapse="-")
#   next()
#   }
#   grouppkg_range[i]<-(paste(c(grouppkg[i],grouppkg2[i+1]),collapse="-"))
# } #making these into ranges rather than just the minimums
#
# #Run the following only once as it takes a while (20 mins)
# all<-cbind(all,GroupKG=NA) # Adding a column of Group to the big dataset
# all$BodyweightKg<-round(all$BodyweightKg) #rounding age because the following code doesn't deal well
#
# #adding a column of ranges
# for(i in 1:length(all$BodyweightKg)){
#   if(is.na(all$BodyweightKg[i])){next()} #remove this estimation from all NAs
#   for(j in 1:length(grouppkg)){
#     if(all$BodyweightKg[i]>=grouppkg[j]&&all$BodyweightKg[i]<=grouppkg2[j+1]){ #iterate over groups and
#     all$GroupKG[i]<-grouppkg_range[j] #assign the necessary group range to that age
#     }
#   }
# }

```

```

# }
#
#
# # # #writing table for future use because the chunk above takes FOREVER (30 min-ish)
# write.table(all, "~/power/data/open/openpowerlifting2.csv", quote=FALSE, row.names = FALSE, sep=',')

ord<-mixedorder(open$Group) #sorting by the group column so that they plot by age
open2<-open[ord,] #aplying that to the whole dataframe
open2$Group<-factor(open2$Group, levels=unique(open2$Group)) #arranging the factors now.

men<-open2 %>% #making a table of men who have completed full competitions and have Age data
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Age)) %>%
  filter(Sex=="M") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, WeightClass)

women<-open2 %>% #same for women
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Age)) %>%
  filter(Sex=="F") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, WeightClass)

```

How much does age matter in the weight you lift?

We can start gaining some initial insight by plotting relationships here. Let's see for each lifter, what their total was and at what age they lifted it. When we plot these relationships over hundreds of thousands of people, this might clue us into some trends of how age affects your totals. I will start by plotting up the total lifts of men and women versus their age as a first pass. In addition, I will break these up into the equipment used to lift this weight (Multi-ply, Raw, Single-ply, Wraps) because that is likely a strong co-variate (another variable that might also explain total weight lifted and obscure our relationships).

```

jpeg(filename="~/power/data/power_2/1.jpg", width=800, height=600)
ggplot(men,
  aes(x=Age, y=TotalKg, color=BodyweightKg)) +
  geom_point(cex=.7) +
  facet_wrap(~Equipment) +
  scale_color_gradient2(guide="colourbar", low="yellow", mid="gold", high="red") +
  ggtitle("Males: How age impacts total weight lifted")+
  labs(y="Total weight lifted (kg)", x="Age (yr)") +
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))

```

```
## Warning: Removed 729 rows containing missing values (geom_point).
```

```
dev.off()
```

```
## pdf
## 2
```

```

jpeg(filename="~/power/data/power_2/2.jpg", width=800, height=600)
ggplot(women,
  aes(x=Age, y=TotalKg, color=BodyweightKg)) +

```

```

geom_point(cex=.7) +
facet_wrap(~Equipment) +
scale_color_gradient2(guide="colourbar",low="yellow",mid="gold",high="red")+
ggtitle("Females: How age impacts total weight lifted")+
labs(y="Total weight lifted (kg)",x="Age (yr)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 181 rows containing missing values (geom_point).
```

```
dev.off()
```

```
## pdf
## 2
```

There's an obvious peak around age 25, with a slowdown afterwards. What does that mean? We have a ton of data and variability, let's take a zoomed out view I'm plotting a distribution of data in different ages to see how representative we are and what the resolution. What you will see is that I grouped individual in age-groups of 5 years. Then I plot the performance (total weight lifted) of those age groups in a distribution. This takes form as a "violin plot". In each group you will see a blob that sort of looks like a violin. The locations on the y axis that the violin is "fatter" show that there is a large amount of lifters that represent that portion of the distribution. Where the violin is "skinnier", then there are few lifters there.

```

jpeg(filename=~ /power/data/power_2/3.jpg",width=800,height=600)
ggplot(men,
      aes(x=Group,y=TotalKg,fill=Group))+
geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=TRUE,scale="width") +
facet_wrap(~Equipment) +
theme(axis.text.x=element_text(angle=90))+
ggtitle("Males: How age impacts total weight lifted")+
labs(y="Total weight lifted (kg)",x="Age groups (yr)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 729 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```

jpeg(filename=~ /power/data/power_2/4.jpg",width=800,height=600)
ggplot(women,
      aes(x=Group,y=TotalKg,fill=Group))+
geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=TRUE,scale="width") +
facet_wrap(~Equipment) +
theme(axis.text.x=element_text(angle=90))+
ggtitle("Females: How age impacts total weight lifted")+
labs(y="Total weight lifted (kg)",x="Age groups (yr)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 181 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

What does this mean?

What you see on the x axis are the age groups that I added to the data earlier. The higher the plot, the stronger the lifters. The higher the horizontal line across the violins, the higher the median of that group is.

There is clear evidence of a pretty consistent relationship, showing both men and women peak overall in their mid-20's... How true is that though?

I ran a linear model, which is basically a fancy way to statistically look at a relationship between variables. In our case I want to know how much do **Age** and **Bodyweight** affect one's total. This model asks this. If you see a positive number next to either Age or Bodyweight, that means that the older or heavier you are, the more you tend to lift as your TotalKg.

```
# fit<-aov(TotalKg~Group,data=men)
line<-lm(TotalKg~Age+BodyweightKg+Equipment,data=men)
line2<-lm(TotalKg~Age+BodyweightKg+Equipment,data=women)
# layout(matrix(c(1,2,3,4),2,2))
# plot(line)
summary(line)
```

```
##
## Call:
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = men)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -752.39  -68.92    3.04   72.80  512.23
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    390.40274    3.05343  127.86 <2e-16 ***
## Age            -1.19420    0.02499  -47.78 <2e-16 ***
## BodyweightKg    3.98592    0.01480  269.35 <2e-16 ***
## EquipmentRaw   -170.38049    2.60886  -65.31 <2e-16 ***
## EquipmentSingle-ply -77.40003    2.61819  -29.56 <2e-16 ***
## EquipmentWraps -133.16085    2.67832  -49.72 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 111.6 on 133743 degrees of freedom
## (885 observations deleted due to missingness)
## Multiple R-squared:  0.4118, Adjusted R-squared:  0.4118
## F-statistic: 1.873e+04 on 5 and 133743 DF, p-value: < 2.2e-16
```

```
summary(line2)
```

```
##
## Call:
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = women)
##
## Residuals:
```

```
##      Min      1Q  Median      3Q      Max
## -326.76 -46.84   -2.48   44.63  344.43
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    305.10687    4.21639   72.36 <2e-16 ***
## Age            -0.41827    0.02673  -15.64 <2e-16 ***
## BodyweightKg    2.12087    0.01844  115.03 <2e-16 ***
## EquipmentRaw   -134.68305    3.93889  -34.19 <2e-16 ***
## EquipmentSingle-ply -44.54240    3.95998  -11.25 <2e-16 ***
## EquipmentWraps -117.08138    4.01126  -29.19 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 72.4 on 56733 degrees of freedom
## (216 observations deleted due to missingness)
## Multiple R-squared:  0.3358, Adjusted R-squared:  0.3358
## F-statistic: 5737 on 5 and 56733 DF, p-value: < 2.2e-16
```

The important things from what you see above are in two columns: 1) Under coefficients: “Estimate” and “Pr(>|t|)”. The estimate column can be read as such: “When you increase the parameter (in this case Age or Bodyweight) by one unit, your TotalKg changes by this much.” This is **cool!** It’s way cool because even though the relationship between Age and Total weight lifted is significantly negative (the Pr column has a very small number), in both men and women, this relationship is not very strong. On average men lose about 1.2kg of their total per year of aging, while women lose about .5.

The **much stronger** relationship is between bodyweight and Total. Men gain almost 4kg for each 1kg in bodyweight and for women that relationship is more like 2:1.

So maybe age is not that huge of a predictor of your lifting potential. This analysis has a flaw. I am considering all men and all women competing overall. What this doesn’t grasp is the fact that we have a sort of curve in the middle, where people in their mid 20’s peak. Let’s dive in.

I’m going to cull the data to remove people under 24. In the time before 25, the total actually increases with age and it creates a weird hump that may be obscuring some patterns.

I’ll check what the patterns are for younger people later on!

```
men2<-men %>%
  filter(Age>24)

women2<-women %>%
  filter(Age>24)

jpeg(filename=~ /power/data/power_2/5.jpg",width=800,height=600)
ggplot(men2,
  aes(x=Group,y=TotalKg,fill=Group))+
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width") +
  facet_wrap(~Equipment) +
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: How age impacts total weight lifted")+
  labs(y="Total weight lifted (kg)",x="Age groups (yr)")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
```

```
## Warning: Removed 516 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf  
## 2
```

```
jpeg(filename=~power/data/power_2/6.jpg",width=800,height=600)  
ggplot(women2,  
       aes(x=Group,y=TotalKg,fill=Group))+  
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width") +  
  facet_wrap(~Equipment) +  
  theme(axis.text.x=element_text(angle=90))+  
  ggtitle("Females: How age impacts total weight lifted")+  
  labs(y="Total weight lifted (kg)",x="Age groups (yr)")+  
  theme(axis.title=element_text(size=16),  
        axis.text=element_text(size=16),  
        plot.title=element_text(size=22))
```

```
## Warning: Removed 132 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf  
## 2
```

Now the negative relationships with age are starting to become more and more obvious. Let's take a look with our fancy *linear model* again. This time, we have excluded the younger people and we can purely look at the decline of strenght, past the peak in mid-20s.

```
line<-lm(TotalKg~Age+BodyweightKg+Equipment,data=men2)  
line2<-lm(TotalKg~Age+BodyweightKg+Equipment,data=women2)  
# layout(matrix(c(1,2,3,4),2,2))  
# plot(wa)  
summary(line)
```

```
##  
## Call:  
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = men2)  
##  
## Residuals:  
##      Min       1Q   Median       3Q      Max   
## -757.61  -65.80    1.75   67.91  468.77   
##  
## Coefficients:  
##              Estimate Std. Error t value Pr(>|t|)      
## (Intercept)    572.95247    3.56795  160.58 <2e-16 ***  
## Age            -4.28207    0.03359  -127.48 <2e-16 ***  
## BodyweightKg    3.48279    0.01817  191.72 <2e-16 ***  
## EquipmentRaw   -180.19525    2.72644  -66.09 <2e-16 ***  
## EquipmentSingle-ply -67.47764    2.73648  -24.66 <2e-16 ***  
## EquipmentWraps -141.44523    2.81157  -50.31 <2e-16 ***  
## ---  
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1  
##  
## Residual standard error: 104.4 on 75922 degrees of freedom  
## (606 observations deleted due to missingness)  
## Multiple R-squared:  0.4802, Adjusted R-squared:  0.4802   
## F-statistic: 1.403e+04 on 5 and 75922 DF,  p-value: < 2.2e-16
```



```
summary(line2)
```

```
##
## Call:
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = women2)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -301.14  -46.28   -3.64   41.78  328.78
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    409.88862    4.79269   85.52 <2e-16 ***
## Age            -2.21699    0.03891  -56.98 <2e-16 ***
## BodyweightKg    1.85859    0.02277   81.63 <2e-16 ***
## EquipmentRaw   -151.43106    4.25818  -35.56 <2e-16 ***
## EquipmentSingle-ply -47.64090    4.29682  -11.09 <2e-16 ***
## EquipmentWraps -134.20279    4.34645  -30.88 <2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 70.47 on 33965 degrees of freedom
## (153 observations deleted due to missingness)
## Multiple R-squared:  0.376, Adjusted R-squared:  0.3759
## F-statistic: 4093 on 5 and 33965 DF, p-value: < 2.2e-16
```

Age is starting to matter a little more

The relationship between Age and TotalKg is larger now. That's because we are looking after this original growth period now. Men seem to be losing closer to 4.2kg per year off their total, while women lose about 2.2kg. Good for women! One interpretation is that as you get older, it's more difficult to lift heavy weights... duh. Obviously we have a bodyweight disparity here too, where the bigger you are, the more you will lift. An interesting interaction here suggests though, that there is a negative correlation between age and bodyweight, meaning: the older you are, you lose bang for the buck in bodyweight - is this real?

So far we have been looking at this across all equipment types, let's break it down and see if these relationships hold among all. What I'm doing here is grabbing the data and splitting it between equipment types. Then I'll look at the relationship between Age and Total weight lifted, as well as the one between Bodyweight and Total weight lifted. This will yield a similar linear model as above, but for much more specific groups now.

```
menm<-men2 %>%
  filter(Equipment=="Multi-ply")

menr<-men2 %>%
  filter(Equipment=="Raw")

mens<-men2 %>%
  filter(Equipment=="Single-ply")

menw<-men2 %>%
  filter(Equipment=="Wraps")

mm<-lm(TotalKg~Age+BodyweightKg,data=menm)
mm$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 596.328109      -5.410918      3.687664
```

```
mr<-lm(TotalKg~Age+BodyweightKg,data=menr)
mr$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 406.302588      -3.593935      3.076590
```

```
ms<-lm(TotalKg~Age+BodyweightKg,data=mens)
ms$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 490.196102      -4.920829      3.928118
```

```
mw<-lm(TotalKg~Age+BodyweightKg,data=menw)
mw$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 443.402471      -4.223248      3.344580
```

```
womenm<-women2 %>%
  filter(Equipment=="Multi-ply")
```

```
womenr<-women2 %>%
  filter(Equipment=="Raw")
```

```
womens<-women2 %>%
  filter(Equipment=="Single-ply")
```

```
womenw<-women2 %>%
  filter(Equipment=="Wraps")
```

```
wm<-lm(TotalKg~Age+BodyweightKg,data=womenm)
wm$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 329.775536      -1.233394      2.425924
```

```
wr<-lm(TotalKg~Age+BodyweightKg,data=womenr)
wr$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 270.472244      -1.781078      1.463824
```

```
ws<-lm(TotalKg~Age+BodyweightKg,data=womens)
ws$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 337.283066      -3.328599      2.894045
```

```
ww<-lm(TotalKg~Age+BodyweightKg,data=womenw)
ww$coefficients
```

```
## (Intercept)      Age BodyweightKg
## 268.636496      -1.665245      1.681378
```

So this is **REALLY** interesting! Let's break down these results: What you see here is that different equipment types have a difference in how Age and Bodyweight affect the Totals:

- Men

- Multi-ply - **large** loss of total with Age, **decent** gain of Total with Bodyweight
- Raw - **decent** loss of total with Age, **decent** gain of Total with Bodyweight
- Single-ply - **large** loss of total with Age, **decent** gain of Total with Bodyweight
- Wraps - **large** loss of total with Age, **decent** gain of Total with Bodyweight
- Women
- Multi-ply - **small** loss of total with Age, **decent** gain of Total with Bodyweight
- Raw - **small** loss of total with Age, **small** gain of Total with Bodyweight
- Single-ply - **decent** loss of total with Age, **decent** gain of Total with Bodyweight
- Wraps - **very small** loss of total with Age, **small** gain of Total with Bodyweight

Now this is super interesting. Is this a real relationship or do we maybe just have a lack of representation of older or bigger lifters in some leagues?2

```
jpeg(filename=~ /power/data/power_2/7.jpg",width=800,height=600)
ggplot(men2,
  aes(x=Equipment,y=Age,fill=Equipment))+
  geom_violin(draw_quantiles = c(0.5))+
  ggtitle("Males: Age distribution among equipment types")+
  labs(y="Age (yrs)",x="Equipment used")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
dev.off()
```

```
## pdf
## 2
```

```
jpeg(filename=~ /power/data/power_2/8.jpg",width=800,height=600)
ggplot(women2,
  aes(x=Equipment,y=Age,fill=Equipment))+
  geom_violin(draw_quantiles = c(0.5))+
  ggtitle("Females: Age distribution among equipment types")+
  labs(y="Age (yrs)",x="Equipment used")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
dev.off()
```

```
## pdf
## 2
```

When we look at Age in each league, we do see that there are proportionally slightly more older lifters in the Multi-ply and Single-ply categories. What you see is that the violin plot above looks a little thicker at the higher age ranges (y axis). This suggests to us that proportionally, those Equipment types attract a little more lifters than they do in the Raw class/or Wraps. This could be because Raw and Wraps are also much more popular with the new generation of lifters, while the other two classes are much more historically prevalent, especially in the United States. You can see the over-representation of young Raw and Wraps lifters by the thickness of the violin plot at the bottom of the y axis.

Now let's see what the distribution of weights looks like among equipment types to see if we see a trend amongst them.

```
jpeg(filename=~ /power/data/power_2/9.jpg",width=800,height=600)
ggplot(men2,
      aes(x=Equipment,y=BodyweightKg,fill=Equipment))+
  geom_violin(draw_quantiles = c(0.5),trim=FALSE)+
  ggtitle("Males: Bodyweight distribution among equipment types")+
  labs(y="Bodyweight (kg)",x="Equipment used")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

Warning: Removed 90 rows containing non-finite values (stat_ydensity).

```
dev.off()
```

```
## pdf
## 2
```

```
jpeg(filename=~ /power/data/power_2/10.jpg",width=800,height=600)
ggplot(women2,
      aes(x=Equipment,y=BodyweightKg,fill=Equipment))+
  geom_violin(draw_quantiles = c(0.5),trim=FALSE)+
  ggtitle("Females: Age distribution among equipment types")+
  labs(y="Bodyweight (kg)",x="Equipment used")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

Warning: Removed 22 rows containing non-finite values (stat_ydensity).

```
dev.off()
```

```
## pdf
## 2
```

There is a strange trend to have heavier lifters in the Multi-ply and Wraps divisions for men. This could be a difference in the weight categories offered by different leagues that support multi-ply, vs ones that don't.

In women, we do see smaller lifters in the Single-ply category, but other than that things seem pretty equal in terms of distribution. Remember the wider the violin, the more lifters are there.

You can see clearly the categories by the wiggle in distribution, each wiggle is a high proportion of people likely at the top of their weight class, where there are less lifters at the bottom of a weight class.

Now that doesn't tell us too much about how these are distributed over age, but these are interesting trends to look at anyways! What I'd like to see is if we see a relationship between the bodyweight of lifters and their age. Let's plot that up.

```
jpeg(filename=~ /power/data/power_2/11.jpg",width=800,height=600)
ggplot(men2,
      aes(x=Age,y=BodyweightKg,color=Equipment))+
  geom_point()+
  facet_wrap(~Equipment)+
  ggtitle("Males: How much do lifters weigh in different age ranges")+
  labs(y="Bodyweight (kg)",x="Age (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

Warning: Removed 90 rows containing missing values (geom_point).

```
dev.off()
```

```
## pdf  
## 2
```

```
jpeg(filename=~power/data/power_2/12.jpg",width=800,height=600)  
ggplot(women2,  
       aes(x=Age,y=BodyweightKg,color=Equipment))+  
  geom_point()+  
  facet_wrap(~Equipment)+  
  ggtitle("Females: How much do lifters weigh in different age ranges")+  
  labs(y="Bodyweight (kg)",x="Age (yrs)")+  
  theme(axis.title=element_text(size=16),  
        axis.text=element_text(size=16),  
        plot.title=element_text(size=22))
```

```
## Warning: Removed 22 rows containing missing values (geom_point).
```

```
dev.off()
```

```
## pdf  
## 2
```

Something smells fishy here though! (and no it's not the knee sleeves I haven't washed in 3 months)

We saw that both Age and Weight have some effect on the weight lifted by people, but do we just see a large proportion of monster heavyweight lifters in the younger categories? The strange thing we are seeing here is that there seems to be a cap on older lifters in both men and women, where the older you get the less lifters in heavy categories you have!

Let me start getting to my point here - big guys lift big weights (generally) -> case and point below (although the relationship is a little weaker in women). Let's see how this shows when we plot bodyweight vs. total weight lifted

```
jpeg(filename=~power/data/power_2/13.jpg",width=800,height=600)  
ggplot(men2,  
       aes(x=BodyweightKg,y=TotalKg,color=Equipment))+  
  geom_point(cex=.3)+  
  facet_wrap(~Equipment) +  
  geom_smooth(method=lm)+  
  ggtitle("Males: Heavier lifters tend to lift more weight")+  
  labs(y="Total weight lifted (kg)",x="Bodyweight (kg)")+  
  theme(axis.title=element_text(size=16),  
        axis.text=element_text(size=16),  
        plot.title=element_text(size=22))
```

```
## Warning: Removed 606 rows containing non-finite values (stat_smooth).
```

```
## Warning: Removed 606 rows containing missing values (geom_point).
```

```
dev.off()
```

```
## pdf  
## 2
```

```
jpeg(filename=~power/data/power_2/14.jpg",width=800,height=600)  
ggplot(women2,  
       aes(x=BodyweightKg,y=TotalKg,color=Equipment))+  
  geom_point(cex=.3)+
```

```

facet_wrap(~Equipment) +
geom_smooth(method=lm)+
  ggtitle("Females: Heavier lifters tend to lift more weight")+
labs(y="Total weight lifted (kg)",x="Bodyweight (kg)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 153 rows containing non-finite values (stat_smooth).
```

```
## Warning: Removed 153 rows containing missing values (geom_point).
```

```
dev.off()
```

```
## pdf
## 2
```

But what is the distribution of weight in different ages? Do we start seeing less heavy lifters when we get to older ages? That was suggested by our plot above, but let's visualize it in groups of ages as well.

```

jpeg(filename=~ /power/data/power_2/15.jpg",width=800,height=600)
ggplot(men2,
      aes(x=Group,y=BodyweightKg,color=TotalKg))+
geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width")+
scale_color_gradient2(guide="colourbar",low="yellow",mid="gold",high="red") +
facet_wrap(~Equipment) +
theme(axis.text.x=element_text(angle=90))+
ggtitle("Males: Less heavier weight older lifters")+
labs(y="Bodyweight (kg)",x="Age group (yrs)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 90 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```

jpeg(filename=~ /power/data/power_2/16.jpg",width=800,height=600)
ggplot(women2,
      aes(x=Group,y=BodyweightKg,color=TotalKg))+
geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width")+
scale_color_gradient2(guide="colourbar",low="yellow",mid="gold",high="red") +
facet_wrap(~Equipment) +
theme(axis.text.x=element_text(angle=90))+
ggtitle("Females: Less heavier weight older lifters")+
labs(y="Bodyweight (kg)",x="Age group (yrs)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))

```

```
## Warning: Removed 22 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

Look above. What you see is that the mean weights do decrease slightly with age for both men and women... but! The upper end of those violins shows you those heavy lifters we talk about! You clearly see that when we reach 50 and 60s or so, there are very few really heavy men and women. Hmmm!

Let's look at the same type of relationship in total weight lifted. I think when we compare maximum weight lifted in each age category and compare those, the relationship is way different than if we compare the means in each category. These two metrics paint two separate pictures and I think bodyweight may be part of the explanation (at least in this dataset).

What you see below is the median total weight lifted in each weight category, as well as the maximum weight lifted for men.

```
jpeg(filename=~ /power/data/power_2/17.jpg",width=800,height=600)
ggplot(men2,
      aes(x=Group,y=TotalKg,fill=Group))+
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width") +
  facet_wrap(~Equipment) +
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

```
## Warning: Removed 516 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```
mensum<-men2 %>%
  filter(!is.na(Group),!is.na(BodyweightKg),!is.na(TotalKg)) %>%
  group_by(Group,Equipment) %>%
  summarize(med_lifts=median(TotalKg),max_lifts=max(TotalKg))

jpeg(filename=~ /power/data/power_2/18.jpg",width=800,height=600)
ggplot(mensum,aes(color=Equipment))+
  geom_point(aes(x=Group,y=med_lifts))+
  geom_point(aes(x=Group,y=max_lifts))+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
dev.off()
```

```
## pdf
## 2
```

What you see is that the maximum weight lifted in each category declines much more rapidly than the average weight. This does differ between Equipment type with Raw being the smoothest decline and Multi-ply seeming to be the strongest

As example let's take the difference between average 25-29 category and 50-54

```

med<-mensusm %>%
  filter(Group=="25-29" || Group=="50-54")

meddif<-c()
for(i in 1:(dim(med)[[1]]/2)){
  meddif[i]<-med[i,"med_lifts"]-med[i+4,"med_lifts"]
}

maxdif<-c()
for(i in 1:(dim(med)[[1]]/2)){
  maxdif[i]<-med[i,"max_lifts"]-med[i+4,"max_lifts"]
}
equip<-c("Multi-ply","Raw","Single-ply","Wraps")
names(meddif)<-equip
names(maxdif)<-equip
print(meddif)

```

```

## $`Multi-ply`
## [1] 64.045
##
## $Raw
## [1] 55
##
## $`Single-ply`
## [1] 90
##
## $Wraps
## [1] 85

```

```
print(maxdif)
```

```

## $`Multi-ply`
## [1] 317.52
##
## $Raw
## [1] 177.5
##
## $`Single-ply`
## [1] 250
##
## $Wraps
## [1] 207.82

```

What you see here are the differences between median and then maximum weight lifted at those two age categories. When we look at median, you see <90kg difference in all categories, and it is as small as 55 kg difference in mean total weight lifted in the raw division

When we look at maximums - **large** differences appear! Raw is still with the smallest difference of 177.5kg, while Multi-ply lifters lift 317.52 kg more as a maximum of the 25-29kg category than the maximum in the 50-55. What this means is that the average weight lifted by older lifters is lower, but not by an insane amount. On the other hand, there are very few older lifters that lift huge weights as the younger categories do. **WHY?**

Let's confirm if there is a similar trend in women before we start answering that question!


```

jpeg(filename=~"/power/data/power_2/19.jpg",width=800,height=600)
ggplot(women2,
      aes(x=Group,y=TotalKg,fill=Group))+
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE,scale="width") +
  facet_wrap(~Equipment) +
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Females: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))

```

```
## Warning: Removed 132 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```

womensum<-women2 %>%
  filter(!is.na(Group),!is.na(BodyweightKg),!is.na(TotalKg)) %>%
  group_by(Group,Equipment) %>%
  summarize(med_lifts=median(TotalKg),max_lifts=max(TotalKg))

```

```

jpeg(filename=~"/power/data/power_2/20.jpg",width=800,height=600)
ggplot(womensum,aes(color=Equipment))+
  geom_point(aes(x=Group,y=med_lifts))+
  geom_point(aes(x=Group,y=max_lifts))+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Females: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
dev.off()

```

```
## pdf
## 2
```

Absolutely there is! I'm going to put some numbers to it again.

```

med<-womensum %>%
  filter(Group=="25-29" || Group=="50-54")

meddif<-c()
for(i in 1:(dim(med)[[1]]/2)){
  meddif[i]<-med[i,"med_lifts"]-med[i+4,"med_lifts"]
}

maxdif<-c()
for(i in 1:(dim(med)[[1]]/2)){
  maxdif[i]<-med[i,"max_lifts"]-med[i+4,"max_lifts"]
}
equip<-c("Multi-ply","Raw","Single-ply","Wraps")
names(meddif)<-equip

```

```
names(maxdif)<-equip
print(meddif)
```

```
## $`Multi-ply`
## [1] 0
##
## $Raw
## [1] 37.5
##
## $`Single-ply`
## [1] 90
##
## $Wraps
## [1] 41.25
```

```
print(maxdif)
```

```
## $`Multi-ply`
## [1] 165.41
##
## $Raw
## [1] 118.5
##
## $`Single-ply`
## [1] 153
##
## $Wraps
## [1] 225.5
```

We see that there is *NO* difference between median Multi-ply lifts between age ranges of 25-29 and 50-54, but over 165Kg difference in maximum lift!! There is again a much more stark difference between maximums at those age ranges, than there is in medians.

Let's resubset our data and see who is lifting those weights and how much do they weigh compared to their max counterparts in the older categories

```
menmax<-men2 %>%
  filter(!is.na(Group),!is.na(BodyweightKg),!is.na(TotalKg)) %>%
  group_by(Group,Equipment,BodyweightKg) %>%
  summarize(max_lifts=max(TotalKg))
```

```
equip<-c("Multi-ply","Raw","Single-ply","Wraps")
```

```
max<-c()
for(i in unique(menmax$Group)){
  for(j in equip){
    sub<-menmax %>%
      filter(Group==i && Equipment==j)
    ord<-order(sub$max_lifts,decreasing=TRUE)
    sub2<-sub[ord,]
    max<-rbind(max,sub2[1,])
  }
}
```

```
max2<-max %>%
  filter(!is.na(Group)) %>%
```

```

mutate(size=BodyweightKg/200)

jpeg(filename=~power/data/power_2/21.jpg",width=800,height=600)
ggplot(max2,
  aes(x=Group,y=max_lifts,color=BodyweightKg))+
  geom_point(aes(cex=size))+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  scale_color_gradient2(guide="colourbar",low="yellow",mid="yellow",high="red")+
  ggtitle("Males: Maximum weight lifted is heavily dependent on bodyweight of lifter")+
  labs(y="Total weight lifted (kg)",x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
dev.off()

```

```

## pdf
## 2

```

If that doesn't show it, what does. The size and redness of the dot is the weight of the lifter who made a maximum lift. As you get to the older categories, what you see is a consistent decline in both maximum lift and in weight of the lifter, who lifted that maximum weight.

Let's show this a little more explicitly with a relationship between weight lifted and weight of lifter

```

jpeg(filename=~power/data/power_2/22.jpg",width=800,height=600)
ggplot(max2,
  aes(x=BodyweightKg,y=max_lifts,color=BodyweightKg))+
  geom_point(aes(cex=size))+
  facet_wrap(~Equipment)+
  scale_color_gradient2(guide="colourbar",low="yellow",mid="gold",high="red")+
  ggtitle("Males: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Bodyweight (kg)")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
dev.off()

```

```

## pdf
## 2

```

```

line<-lm(max_lifts~BodyweightKg,max2)
summary(line)

```

```

##
## Call:
## lm(formula = max_lifts ~ BodyweightKg, data = max2)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -361.99 -141.50  -13.73   132.35   510.13
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  -108.7773   107.8276  -1.009   0.318
## BodyweightKg    7.8187     0.8732   8.954 4.85e-12 ***

```

```
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 184.6 on 51 degrees of freedom
## Multiple R-squared:  0.6112, Adjusted R-squared:  0.6036
## F-statistic: 80.17 on 1 and 51 DF,  p-value: 4.855e-12
```

What you see is a **rather strong** correlation between weight of the lifter and the weight the lifter lifted. Remember, these are the **BEST** lifters of their Age range from among 150 000 entries! We see that 65% of variability between lifters is purely explained by their weight! That's a lot of percent! And this is a significant relationship as well. This is even disregarding equipment types (which I checked and they only explain an extra 5% of the relationship). This is *huge* because it suggests that age may be important for how much you lift, but your bodyweight carries a much higher significance in terms of those huge numbers you see in the younger age categories.

Let's check on this with the women lifter data.

```
womenmax<-women2 %>%
  filter(!is.na(Group), !is.na(BodyweightKg), !is.na(TotalKg)) %>%
  group_by(Group, Equipment, BodyweightKg) %>%
  summarize(max_lifts=max(TotalKg))

equip<-c("Multi-ply", "Raw", "Single-ply", "Wraps")

max<-c()
for(i in unique(womenmax$Group)){
  for(j in equip){
    sub<-womenmax %>%
      filter(Group==i && Equipment==j)
    ord<-order(sub$max_lifts, decreasing=TRUE)
    sub2<-sub[ord,]
    max<-rbind(max, sub2[1,])
  }
}

max2<-max %>%
  filter(!is.na(Group)) %>%
  mutate(size=BodyweightKg/200)

jpeg(filename=~"/power/data/power_2/23.jpg", width=800, height=600)
ggplot(max2,
  aes(x=Group, y=max_lifts, color=BodyweightKg))+
  geom_point(aes(cex=size))+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  scale_color_gradient2(guide="colourbar", low="yellow", mid="yellow", high="red")+
  ggtitle("Females: Maximum weight lifted is heavily dependent on bodyweight of lifter")+
  labs(y="Total weight lifted (kg)", x="Age group (yrs)")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))
dev.off()

## pdf
## 2
```

This is not as convincing as the male data, but the overall trend is similar. Remember, in women there is also a much smaller effect of age on total weight lifted too. And a smaller relationship between weight and total weight lifted.

Let's see the more explicit plot.

```
jpeg(filename=~"/power/data/power_2/24.jpg",width=800,height=600)
ggplot(max2,
       aes(x=BodyweightKg,y=max_lifts,color=BodyweightKg))+
  geom_point(aes(cex=size))+
  facet_wrap(~Equipment)+
  scale_color_gradient2(guide="colourbar",low="yellow",mid="gold",high="red")+
  ggtitle("Females: Maximum weight lifted decreases faster than median weight lifted")+
  labs(y="Total weight lifted (kg)",x="Bodyweight (kg)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
dev.off()
```

```
## pdf
## 2
```

```
line<-lm(max_lifts~BodyweightKg+Equipment,max2)
summary(line)
```

```
##
## Call:
## lm(formula = max_lifts ~ BodyweightKg + Equipment, data = max2)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -280.33  -82.39   32.16   95.18  230.37
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    277.9507    84.1928   3.301 0.002031 **
## BodyweightKg     4.3650     0.9729   4.487 5.98e-05 ***
## EquipmentRaw   -233.1985    62.1342  -3.753 0.000555 ***
## EquipmentSingle-ply -151.9426    70.6037  -2.152 0.037483 *
## EquipmentWraps  -159.2045    63.3078  -2.515 0.016033 *
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 142 on 40 degrees of freedom
## Multiple R-squared:  0.4333, Adjusted R-squared:  0.3766
## F-statistic: 7.646 on 4 and 40 DF,  p-value: 0.0001128
```

This is relatively consistent! A bit more noise than the male data, explaining only 20% of data. When we include Equipment this goes up to 38% of total data explained. Not as convincing, but still significant.

Now one last thing that we need to see is if this is consistent within weight classes between ages. This is a separate analysis that will take us away from who is doing the best lifts, and lead us to better understand how people of the same weight are doing over time.

Let's first look at performance of lifters in weightclasses. This will give us an idea if we have certain weight classes more represented in younger ages than older ages.

```

ord2<-mixedorder(open$GroupKG)
open3<-open[ord2,]
open3$GroupKG<-factor(open3$GroupKG,levels=unique(open3$GroupKG))

men<-open3 %>% #making a table of men who have completed full competitions and have Age data
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Group)) %>%
  filter(Age>24) %>%
  filter(Sex=="M") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, GroupKG)

women<-open3 %>% #same for women
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Group)) %>%
  filter(Age>24) %>%
  filter(Sex=="F") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, GroupKG)

```

Let's try to plot these up in a meaningful way that will allow us to see progression through the age in the same weight groups.

```

jpeg(filename=~ /power/data/power_2/25.jpg",width=800,height=600)
ggplot(men,
  aes(x=Group,y=TotalKg,fill=GroupKG,color=GroupKG))+
  geom_violin(scale="area")+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Older lifters are less represented in the bigger weight classes")+
  labs(y="Total weight lifted (kg)",x="Age (yrs)")+
  theme(axis.title=element_text(size=16),
    axis.text=element_text(size=16),
    plot.title=element_text(size=22))

```

```
## Warning: Removed 516 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
```

```
## 2
```

```
line<-lm(max_lifts~BodyweightKg,max2)
summary(line)
```

```
##
```

```
## Call:
```

```
## lm(formula = max_lifts ~ BodyweightKg, data = max2)
```

```
##
```

```
## Residuals:
```

```
##      Min       1Q   Median       3Q      Max
## -366.16  -79.94   22.67  110.51  320.23
```

```
##
```

```
## Coefficients:
```

```
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)  218.5400     89.3209   2.447 0.018581 *
## BodyweightKg   3.4304     0.9577   3.582 0.000863 ***
```

```
## ---
```

```
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 159.7 on 43 degrees of freedom
## Multiple R-squared:  0.2298, Adjusted R-squared:  0.2119
## F-statistic: 12.83 on 1 and 43 DF,  p-value: 0.0008634
```

```
jpeg(filename=~ /power/data/power_2/26.jpg",width=800,height=600)
ggplot(women,
       aes(x=Group,y=TotalKg,fill=GroupKG,color=GroupKG))+
  geom_violin(scale="area")+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Females: Older lifters are less represented in the bigger weight classes")+
  labs(y="Total weight lifted (kg)",x="Age (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

```
## Warning: Removed 132 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

You can start seeing the trends here. The heaviest of people are no longer present in the older age groups. You can see this by the disappearance of the violin plots of pinkish, bluish color as you get to the older age groups. Now this is starting to tell us that those max lifter monsters that are really heavy and make the big lifts, either no longer compete later, or they lose weight as they age. Now let's see that distribution broken down into age groups over time a little more clearly.

I've taken here the most represented group over time the 86-95kg group in men and the 56-75kg group in women. Then I've plotted the total weights they lift as they age to see how the same group performs over time.

```
menr<-men %>%
  filter(BodyweightKg>85) %>%
  filter(BodyweightKg<96)
womnr<-women %>%
  filter(BodyweightKg>55) %>%
  filter(BodyweightKg<76)

jpeg(filename=~ /power/data/power_2/27.jpg",width=800,height=600)
ggplot(menr,
       aes(x=Group,y=TotalKg,fill=Equipment))+
  geom_violin(scale="width",trim=FALSE,draw_quantiles = 0.5)+
  facet_wrap(~Equipment)+
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Performance declines with age slowly within the same weight group")+
  labs(y="Total weight lifted (kg)",x="Age (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

```
## Warning: Removed 115 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2

jpeg(filename=~"/power/data/power_2/28.jpg",width=800,height=600)
ggplot(womenr,
      aes(x=Group,y=TotalKg,fill=Equipment))+
geom_violin(scale="width",trim=FALSE,draw_quantiles = 0.5)+
facet_wrap(~Equipment)+
theme(axis.text.x=element_text(angle=90))+
ggtitle("Females: Performance declines with age slowly within the same weight group")+
labs(y="Total weight lifted (kg)",x="Age (yrs)")+
theme(axis.title=element_text(size=16),
      axis.text=element_text(size=16),
      plot.title=element_text(size=22))
```

```
## Warning: Removed 71 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```
line<-lm(TotalKg~Age,menr)
line$coefficients
```

```
## (Intercept)      Age
##  717.52533    -3.18427
```

```
line1<-lm(TotalKg~Age,womenr)
line1$coefficients
```

```
## (Intercept)      Age
##  389.730230    -1.611683
```

Above I am showing you the performances of people in the same weight group as they age for men (top) and women (bottom). Well y'all this seems just idealistic for me... There's a decline in performance with age, but that decline is not really that incredible, compared to what we saw above in terms of the decline of performance with weight. The means age performance declines rather slowly and while overall 25 year olds definitely perform better, there are a good number of 50+ year olds that perform above the 25yo average.

So can you still lift well as you get older?

What we see in this dataset is that Age does matter in the amount of weight that you lift... but not as much as you might think! It is dependent on which type of lifting you do (Multi-ply, Raw...). But overall, the average lifts for ages don't really drop by much with time. We are talking in the range of 3kg/year for men and 1.6 kg/year for women! For men in 30 years that's 60kg or 132 lbs. In a hypothetical if you started with 1500 total at 25, at 55 you can still pull a nice 1350 (ON AVERAGE). For women that's a decline of 48kg or 115lbs. Start with a 1000lbs, go to 885 at 55. Not bad eh! I mean I'd love to see a 55 year old at my gym who can pull 500 lbs, and this data seems to suggest that this is highly doable!

Let me make this clear - these are trends! This doesn't say that with a lot of training and good regimen you can't keep or exceed your gains past a certain point! Maybe when you get to your 60s it will get more difficult as these trends show a steep dropoff at that point.. The bigger influence on your lifts are your weight. That being said:

DON'T TRY TO FATTEN UP JUST TO LIFT MORE especially as you get older. That comes with many health outcomes, so **DO NOT TAKE THIS AWAY FROM THIS DATA**. What you can take away from

here is that if you plan to keep your weight, you don't have to stop lifting. You're not really likely to lose your *gainz* over time by very much.

I want to absolutely acknowledge that this is data of likely trained lifters. If you're competing in your 50s and 60s, you're likely to be a seasoned powerlifter, so don't compare your numbers too harshly to everyone here. These folks are good!

Final message that I think is well supported with this data - don't stop lifting, don't stop asking questions!

PS: ## Younglings

For shits and giggles let's verify why I threw out the data of people under 22. + They create a hump in the distribution of Total weight lifted because they are still getting stronger. If we plot these side by side, you can see that lifters get stronger till about 23-25 and then plateau there.

```
men3<-open2 %>%
  filter(Age<25) %>%
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Age)) %>%
  filter(Sex=="M") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, WeightCL

jpeg(filename=~"/power/data/power_2/29.jpg",width=800,height=600)
ggplot(men3,
  aes(x=Group,y=TotalKg,fill=Group))+
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE) +
  facet_wrap(~Equipment) +
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Performance increases with age till mid-20s")+
  labs(y="Total weight lifted (kg)",x="Age groups (yrs)")+
  theme(axis.title=element_text(size=16),
        axis.text=element_text(size=16),
        plot.title=element_text(size=22))
```

```
## Warning: Removed 213 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```
ym<-lm(TotalKg~Age+BodyweightKg+Equipment,data=men3)
summary(ym)
```

```
##
## Call:
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = men3)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -677.87  -61.04    0.61   61.70  449.85
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)   -26.23804     6.06224  -4.328 1.51e-05 ***
## Age             18.98153     0.14926 127.173 < 2e-16 ***
## BodyweightKg    3.56664     0.02074 171.999 < 2e-16 ***
## EquipmentRaw  -141.29413     5.08609 -27.780 < 2e-16 ***
```

```
## EquipmentSingle-ply  -61.31487    5.10453 -12.012 < 2e-16 ***
## EquipmentWraps      -116.37323    5.17015 -22.509 < 2e-16 ***
## ---
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 95.58 on 57815 degrees of freedom
## (279 observations deleted due to missingness)
## Multiple R-squared:  0.5324, Adjusted R-squared:  0.5323
## F-statistic: 1.316e+04 on 5 and 57815 DF,  p-value: < 2.2e-16
```

```
women3<-open2 %>%
  filter(Age<25) %>%
  filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
  filter(!is.na(Age)) %>%
  filter(Sex=="F") %>%
  select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, WeightCL)
```

```
jpeg(filename=~"/power/data/power_2/30.jpg",width=800,height=600)
ggplot(women3,
  aes(x=Group,y=TotalKg,fill=Group))+
  geom_violin(draw_quantiles=c(0.5),aes(fill=Group),trim=FALSE) +
  facet_wrap(~Equipment) +
  theme(axis.text.x=element_text(angle=90))+
  ggtitle("Males: Performance increases with age till mid-20s")+
  labs(y="Total weight lifted (kg)",x="Age groups (yrs)")+
  theme(axis.title=element_text(size=16),
  axis.text=element_text(size=16),
  plot.title=element_text(size=22))
```

```
## Warning: Removed 49 rows containing non-finite values (stat_ydensity).
```

```
dev.off()
```

```
## pdf
## 2
```

```
yw<-lm(TotalKg~Age+BodyweightKg+Equipment,data=women3)
summary(yw)
```

```
##
## Call:
## lm(formula = TotalKg ~ Age + BodyweightKg + Equipment, data = women3)
##
## Residuals:
##      Min       1Q   Median       3Q      Max
## -314.42  -40.58   -1.59    39.55   318.05
##
## Coefficients:
##              Estimate Std. Error t value Pr(>|t|)
## (Intercept)    48.50689    8.71955   5.563 2.68e-08 ***
## Age             8.69720    0.14493  60.012 < 2e-16 ***
## BodyweightKg   2.37514    0.02706  87.772 < 2e-16 ***
## EquipmentRaw   -86.42143    8.09560 -10.675 < 2e-16 ***
## EquipmentSingle-ply  0.67369    8.10626  0.083  0.934
## EquipmentWraps -76.46058    8.19974  -9.325 < 2e-16 ***
## ---
```

```
## Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
##
## Residual standard error: 64.6 on 22762 degrees of freedom
## (63 observations deleted due to missingness)
## Multiple R-squared:  0.4606, Adjusted R-squared:  0.4605
## F-statistic: 3887 on 5 and 22762 DF, p-value: < 2.2e-16
```

- The linear model here reveals a strong positive relationship of both Age and Weight.

```
#
# men<-open %>%
#   filter(BestBenchKg>0 & BestDeadliftKg>0 & BestSquatKg>0) %>%
#   filter(!is.na(Age)) %>%
#   filter(Sex=="M") %>%
#   filter(Federation==c("USAPL", "USPA", "IPF", "IPL")) %>%
#   select(Equipment, Age, BodyweightKg, BestSquatKg, BestBenchKg, BestDeadliftKg, TotalKg, Wilks, Group, Federation)
#
#
# ggplot(men,
#   aes(x=Federation, y=TotalKg, fill=Federation))+
#   geom_violin(draw_quantiles = 0.5, trim=FALSE)
```