This document is reachable via this shortened URL: https://goo.gl/PL39c2

Practice texts here and here. TextBlade forum thread on the layout.

As a software developer, I spend tons of time at the keyboard, so I supported Keyboardio's kickstarter. Sculpted keycaps laid out in columns, lots of thumb keys, mechanical switches, and other nice features. Just having the SHIFT key under your left thumb rather than having to take a hand off the home row to press SHIFT with a pinkie is marvelous (I have had this on the Truly Ergonomic for years). Yes, a keyboard like the Keyboardio is expensive, but it's worth it if you use it enough.

But I don't intend to use the standard **QWERTY** layout:

Q	W	E	R	T		Y	U	I	O	P
A	S	D	F	G		H	J	K	L	;:
Z	X	C	V	B		N	M	,<	.>	/?

The placement of ***N*** in relation to ***I***, ***O***, and ***U*** is a source of errors. ***E***, ***D***, and ***C*** -- commonly juxtaposed in English text -- are on the same finger. ***T*** is in a bad location given its high frequency in English text. The list goes on and on. I've been wanting to use a saner layout for a while, and I plan to make the switch with the new keyboard. I looked at a number of keyboard layouts and even started to practice on the **Colemak** layout:

Q	W	F	P	G		J	L	U	Y	;:
A	R	S	T	D		H	N	E	I	O
Z	X	C	V	B		K	M	,<	.>	/?

[Problematic: HE=81, LE=35, CA=24, UM=6, CR=6, UE=4, NK=4, MY=2?; ION=24, YOU=11, ONE=8, INE=7, OO on pinky] (note to self about suboptimal bigrams/trigrams and their relative frequencies)

And here is the **Dvorak** layout, for reference. Among other things, I find the location of the L curious.

'"	,<	.>	P	Y		F	G	C	R	L
A	O	E	U	I		D	H	T	N	S

;:	Q	J	K	X		B	M	W	V	Z

And here's the **Workman** layout:

Q	D	R	W	B		J	F	U	P	;:
A	S	H	T	G		Y	N	E	I	O
Z	X	M	C	V		K	L	,<	.>	/?

[Problematic: MA=20, PE=16, UL=12, LY=12, CT=11, RS=11, EY=9, KE=8, GH=7, PL=7, CR=6, RM=5, UE=4, NK=4, FL=3, NY=3; ION=24, HAT=17, YOU=11, ONE=8, INE=7, OO on pinky]

But I decided to explore custom arrangements. Note that my focus is on English text; this won't be the best layout for another language. Many thanks to Peter Norvig and Google for this analysis of a huge body of English text, which guided this work. Also note that similarity to QWERTY is a non-goal for me.

Here's what I came up with (subject to further tweaking). Obviously this is just the core of layer 1.

K	M	L	U	?!		V	D	R	'"	Q
A	T	H	E	.:		C	S	N	O	I
Z	P	F	J	,;		B	G	W	X	Y

[Problematic: PA=13, RY=6, FA=6, CR=6, BR=4, UE=4, AK=4, SC=4; HAT=17]

Note to self about letters before/after apostrophes: `{ n I e t y u } ' { t d l s v r m }`

Your eyes probably jumped to the word ***THE*** on the home row. A good, quick smoke test for an English keyboard layout is to type "**and of the nations there**" a few times on it, largely because the most common English word -- *by a considerable margin* -- is ***THE***, followed by ***OF*** and ***AND***:

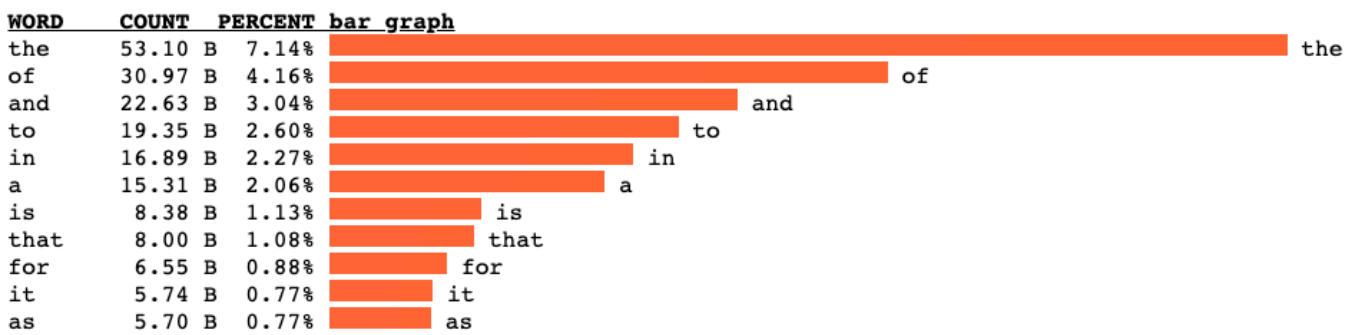

You may use thousands of different words in a document, but statistically around 7% of the words are ***THE***. And ***THE*** is also a part of many larger but fairly common words (***THERE THEY THEIR THEN***

THESE THEM EITHER OTHER ANOTHER WHETHER FARTHER FATHER MOTHER BROTHER BOTHER etc.). In fact, ***TH*** is the most common of the 676 English bigrams, followed by ***HE***, and the pack thins out pretty quickly after that:

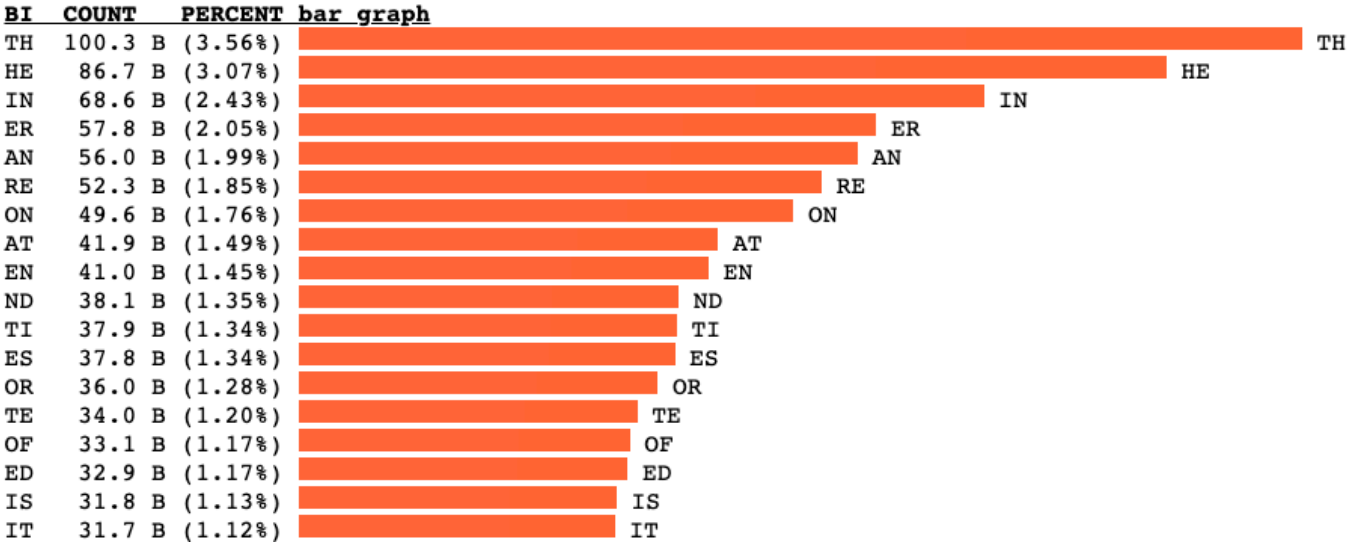

These three letters are so very important to get right -- why not arrange to just splat out ***THE*** (and ***TH*** and ***HE***) in one motion by quickly drumming your fingers, almost as if there were a big ***THE*** key on the keyboard?

I think there are two reasons why most layouts do not do this. One is that ***H*** is not one of the top eight letters in terms of frequency (it's ninth; ***R*** is higher), and generally layout designers want the most commonly used letters on the home row. I don't see a need to be quite that strict about the home row; more on that later. The other reason is that the cost model they use to evaluate layouts is generally not sophisticated enough to reflect the value of a splat. Many cost models treat each keystroke as an independent action, to which some (largely unjustified) cost is ascribed based on its location. But that simplistic cost model is clearly not very accurate. Imagine typing ***RD*** on this layout (as in ***HEARD***). Would you really type ***R***, then return to the home row, and only then move your other finger to type ***D***? Or would both fingers move together and press the two keys? Most cost models tell you that, for example, ***RD*** would be easier to type if ***D*** were on the home row, but the opposite is true -- given that your middle finger is already on the top row typing the ***R***, you're better off if the ***D*** is up there as well.

In fact, you can measure the error yourself. Type the "word" ***SDF*** (that's where ***THE*** will be in this layout) followed by a space on a QWERTY keyboard, over and over again as fast as you can for five seconds. Roll it out, drumming your fingers. If you are used to rolling your fingers on a keyboard, you probably typed 15 or so ***SDF*** "words" in five seconds. Now try it again hitting three disjoint keys on the home row, like ***FKS***, followed by a space. How many now? Now with the middle letter on one of the inner keys -- try ***FHK***. Uggh. That's what it would feel like on a Colemak keyboard. Did your typing speed halve? How was your accuracy? It *matters* where the keys in a sequence are in relation to one another. Successive keystrokes are *not* independent actions.

By the way, this is not just about typing *speed*. It's also about *fatigue* and *accuracy*.

The arrangement of ***R***, ***N***, ***S***, and ***D*** on this layout makes it easy to type each of ***NS***, ***RS***, and ***ND*** in one movement, and that's a win because of their frequencies.

The suffix ***ATIONS*** (really ***TION*** and friends) dominates the list of most common 4/5/6-grams (particularly if you consider that we have already made accommodations for ***THE***, ***TH***, and ***HE***). Here is an excerpt from Norvig's analysis of English text in which row 1 shows the most common English n-grams, the second row the next most common, etc.:

1	2grams	3grams	4-grams	5-grams	6-grams	7-grams	8-grams	9-grams
e	th	the	tion	ation	ations	present	differen	different
t	he	and	atio	tions	ration	ational	national	governmen
a	in	ing	that	which	tional	through	consider	overnment
o	er	ion	ther	ction	nation	between	position	formation
i	an	tio	with	other	ection	ication	ifferent	character
n	re	ent	ment	their	cation	differe	governme	velopment
s	on	ati	ions	there	lation	ifferen	vernment	developme
r	at	for	this	ition	though	general	overnmen	evelopmen

NATIONS STATIONS PUBLICATIONS LOCATIONS SENSATIONS CORPORATIONS DONATIONS VACATIONS MUTATIONS CREATIONS EQUATIONS POPULATIONS GENERATIONS VIBRATIONS VARIATIONS ITERATIONS MIGRATIONS INDICATIONS MOTIVATIONS IMITATIONS LIMITATIONS RECOMMENDATIONS PERMUTATIONS CITATIONS TEMPTATIONS REVELATIONS REPUTATIONS RESERVATIONS OCCUPATIONS OBLIGATIONS NEGOTIATIONS OBSERVATIONS and so on. You can quickly splat out this suffix by rolling ***AT*** inward with the left hand and then ***IONS*** inward with the right. With a little practice the two become one fluid motion, yielding six characters for little over the cost of two. And of course there are all the other *****TION[S]***, *****SION[S]***, etc.

Also note that it is easy to exclude the middle finger from a roll. Clearly you could splat out ***BATHE***, but you can just as easily splat out ***DATE*** with a roll while the middle finger is lifted. ***ATE*** is a very common ending for verbs, e.g. ***FRUSTRATE***, ***RELATE***, and ***TRANSLATE***, as well as some adjectives such as ***ORNATE***, ***INTRICATE***, and ***DELICATE***, and chemical names like ***ACETATE***, ***NITRATE***, and ***SULFATE***.

It is harder to exclude a ring finger from a roll, though; it takes some practice, and personally I can still only do it reliably with my right hand.

A sequence can sometimes be drummed in a fluid motion even when it is not all on the same row. Consider the third most common word in English, ***AND***. It doesn't matter too much that the ***D*** is on the row above, because 1) it is not an unnatural movement, and 2) during the time it takes to move there you are typing the ***A*** and the ***N***, and 3) during the movement back to the home row you are typing the space. This appears to reflect a significant difference in philosophy from other keyboard layouts, where any movement off the home row is assumed to impose a significant penalty. If you need to type another character with the same hand on the bottom row after the ***D***, as you do in ***HANDGUN***, then there is indeed a penalty, but there is little penalty for ***AND*** (which is far, far more common). Similarly, you couldn't type ***HER*** (or any other word ending in ***ER*** or ***ERS***) any faster if ***R*** were on the home row. There is absolutely no time penalty for going off-home if one of the following is true for the key before and the key after:

- It is on the other hand.
- It is on the thumb (a space).

- It is on the same row.
- In *some* cases, if it is one row above or below.

The main point is that the off-home time penalty is only real if the movement does not overlap with another keystroke, so it is not scandalous to put 9th-place ***H*** on the home row to make ***THE*** splattable. This is a more nuanced position than one sees with many other keyboard layouts.

Also, while some layouts treat the innermost two home row keys (where QWERTY puts ***G*** and ***H***) as good places to put commonly used letters (Dvorak puts ***I*** and ***D*** there, Colemak ***D*** and ***H***), I consider those inner keys inferior locations since they drag your (or at least *my*) middle fingers off their columns. If the next letter is on the same hand's middle or ring finger, you're better off going up a row or down a row, especially if the next letter is on that same row. If you've moved off the column to hit an inner key, well, you really have to move back before you can continue. Putting common letters in those column-breaking locations invites errors; consider how the position of ***N*** on QWERTY makes your hands do gymnastics with ***I*** and ***O***. In other words, instead of grading the left-hand locations like this (let's ignore the weakness of the little finger to make this point more clearly):

```
C C C C D
A A A A B
C C C C D
```

I sort of grade them more like this:

```
B B B B D
A A A A C
B B B B D
```

I'm not suggesting that this as an accurate cost model, for many reasons -- I'm just trying to make the point out that pulling the hand sideways off the home keys is costlier than moving a finger up or down. Any cost model based solely on key position like this would be seriously flawed, IMHO, because it ignores the critical issue of where the previous keys left your hand and fingers.

I consider any two-character sequence costly (though not *equally* costly) if both keys are on the same hand and *any* of the following is true:

- They are different characters on the same finger. For example on THE-1, ***SS*** is fine, but ***DS*** is costly -- the time it takes to move the finger does not overlap with another keystroke).
- They are the same character, on a pinkie, e.g. ***PP*** on QWERTY.
- One is on the top row and the other is on the bottom row. Too far, especially if they are adjacent fingers.
- One is on the pinkie, and the other is on the ring finger on a lower row.
- One is on the middle finger, and the other is on the ring finger on a higher row.
- One is on the inner column, and the other is on the middle finger (the one most dragged off its key).

Many of these are consequences of the physiology of the hand. Pinkies are weak and short, and our control over the pinky and ring fingers is somewhat entangled. An average hand cannot hit a key on the inner column without dragging the middle finger off *its* column.

It is significant that ***R***, ***S***, and ***D*** are on the opposite hand from ***E***. Consider the word ***THERE***, containing two ***E***s separated by an ***R***. If a layout puts ***E*** and ***R*** on the same hand (as on QWERTY), then you have to do a weird twiddle to type the ***ERE***. But with this layout, you splat out ***THE*** with the left hand, type ***R*** with the right, and your left index finger is still sitting right over the ***E*** you need to type again. You're really just typing the ***E*** twice, as fast as you can, and slipping the ***R*** in between. It's faster and less fatiguing than the twiddle, and helps with the many **DED***, **RER***, **SES***, and **E?E*** words (just for fun, try the word ***CEDED*** on QWERTY -- aargh!). Colemak and Dvorak also put ***R***, ***S***, and ***D*** opposite ***E***.

By the way, although ***TH*** is the most common bigram, ***ER*** and ***RE*** are both common, and combined are perhaps even slightly more common than ***TH*** + ***HT***. Since both orderings are common, it is good that ***E*** and ***R*** are on different hands.

This layout also tries to keep the letters that are often doubled off the pinkies, which are relatively weak. Compare to Dvorak, which puts ***S*** and ***L*** on the right pinkie. ***L*** is the most frequently doubled letter in English, followed by ***S*** and ***E***. Like all modern layouts, THE-1 also tries to put less frequently used characters under the pinkies, except in home position.

This layout accommodates the ***{T,P,S,C,G,W}H*** digraphs by putting ***T*** and ***P*** before ***H*** to make them splattable (***PH*** is marginal), and ***S***, ***C***, ***G*** and ***W*** on the other hand.

About punctuation. ***E*** is a terribly common and incestuous (meaning that it can be found next to anything) letter. And it's on an index finger, so yikes -- we need to pick five other letters to share the key with, making them hard to type before or after ***E***. At first I tried to keep ***E*** off the index finger for that reason, but I eventually decided it worked best there, so I gave it most of the punctuation for neighbors. That might make the keyboard easier to learn anyway. And often we type a space after those punctuation characters, giving us time to move back to the home row. Note that you don't want the apostrophe there, though, because ***E'*** is common (e.g. he's, she'd, we'll).

Apostrophe is an interesting case. Other layouts have it outside the main area, requiring a hand movement; this layout puts ***'*** (and ***"***) in an accessible place convenient to the letters likely to precede and follow it. One role of the apostrophe is in possessives, where it can follow anything but is generally before an ***S***. In its role in common contractions, there is *****N'T*** (for NOT, as in isn't, haven't, and won't) as well as the contractions for the verbs laid out in this table:

	has/have	am/is/are	will	had/would
I	've	'm	'll	'd
he / she / it	's	's	'll	'd
we / they / you	've	're	'll	'd

Since ***I*** is shifted, you have to be careful not to roll *too* quickly into the apostrophe, but placement is still important. The sequences of letters surrounding these contractions can be summarized as
`{ n I e t y u } ' { t d l s v r }`

Some attention was also given to honorifics such as Mr., Mrs., Ms., Dr., and to words of the form ****ing[ly]***.

Note that **<** and **>** are not in the main layout. I would argue that those keys are ideally placed together with **(){}[]** in some coherent way, and probably not on the bottom row on the middle two fingers. And while angle brackets are used in most programming languages, they are not typically used in normal text; it's probably best to put all four pairs in another layer.

I'd like to write a program that takes a database of n-gram frequencies as input and generates layouts using a rather complicated cost model for finger movement. At the same time, perhaps others who read this can use the n-gram frequencies and their intuition as I have to come up with an even better result. You'll find Norvig's document very handy for this purpose; I especially love the table of AA .. ZZ graphing the relative frequencies of the 676 bigrams (but wish there were a similar table that excluded the top 15, so that the graphs for the majority of bigrams would not be so compressed).

One caveat: while the n-gram statistics that guided this layout were based on a large corpus of English text, that text is at least largely from books. I'm not sure how much the results would change if you looked at a person's emails, chats, code, blog, work documents, etc. instead of books. If I had to guess, I would say that, for example. the Google corpus differs from my own English typing in these ways:

- ***TION*** is probably somewhat more frequent in books in the Google corpus.
- The following are probably are common in personal communications and are probably under-represented in the corpus:
    - Apostrophes. Personal communications are replete with contractions: ***I'm***, ***it's***, ***you're***, ...
    - Pronouns.
    - Conversational words like ***hi***, ***yes***, ***ok***, ***sure***, ***hey***, ***sucks***.
    - Words that are more frequently used in questions and aspects of planning such as time, place, and desire ("***Do*** you ***want*** to ***eat*** at Vinny's ***tonight***? ***Meet*** you in ***front*** of the restaurant a ***little before*** 6?").
    - Words relating to family relationships.
    - Words relating to the workplace.
    - Words relating to common human activities like eating and driving.
- Newer word construction patterns like ****ize***, ****ify***, ****ish***, and ****ique*** are probably under-represented in the corpus.

Alternating Hands vs. Splatting

I've seen people argue that what you *really* want to do is get the typist to alternate hands, one letter with the left, the next with the right, etc. I agree that there is some potential value to this, but it's easy to show that splatting is superior. You can't arrange to splat everything, though. What's nice about alternating hands is that it doesn't matter where the previous key was -- which row or column -- if it was on the other hand. That's a good thing. But it's no splat.

They aren't hard to compare. To test alternation, on a QWERTY keyboard type

```
a;sldkfj
```

over and over again for one minute. To test splatting, type

```
asdf;lkj
```

over and over again for one minute, drumming the fingers. Go ahead and give yourself as much warmup time as you need before starting on each of these.

If you managed to type even *half* as much with alternation, I'm astonished -- for me it was far worse. And even though you typed far more when splatting, you're way less fatigued, I'm guessing.

Don't get me wrong -- splitting high-frequency bigrams is absolutely the right thing to do when there is no way to get them comfortably on the same hand, and that comes up a lot. For example, on THE-1 the letter ***G*** is where it is so that you can splat ***ING*** and alternate ***OUGH***. The ****H*** digraphs discussed earlier are another example -- we splat 2 of them (including the most important one, ***TH***), and alternate 4. But alternation is most certainly not superior to splatting, either in speed or effort. It is more of a good fallback strategy.

And just as you can't arrange to splat everything, you can't arrange to alternate everything. The best layouts are going to take both into consideration: try to splat what's most important, and then try to alternate what you can't splat. Even then you won't get it all -- some things will be relatively hard to type.

Random thoughts on evaluating layouts

I've been thinking about how to evaluate a layout accurately. I think I would use a list of word frequencies, evaluating each word as a separate unit, keeping track of the positions of the fingers after each character and calculating what it would take to type the next character. It would be quite a bit of work to do that accurately, though. And the result would vary from one physical layout to the next, depending on whether the keys are arranged in columns and to what extent the columns heights are adjusted to finger lengths; for example, the pinky columns are lower than the other columns on the Keyboardio Model 01 and the Truly Ergonomic, but the lowering is more pronounced on the Model 01. Hmmm. Even if I did come up with an optimal layout for the Model 01, it wouldn't necessarily be optimal for another keyboard. To be truly accurate, you'd even have to take the user's hand size into account: the same movement might be within reach for one hand but not another. Relative finger lengths matter as well: even though the Model 01 lowers the columns for the index and pinky fingers, they are still about a quarter inch too high for my hand, meaning that it is more work for me to reach the top row than the bottom with those two fingers. For others it is surely just right. And the costs will differ depending on whether the user's hands float or rest on their heels.

I'm not suggesting that anyone actually try to take all of this into account, but rather am pointing out that there are significant limitations to our ability to assess programmatically the actual cost of typing a word. There is some threshold of accuracy below which programmatic attempts to evaluate keyboard layouts -- especially those based on simplistic models -- are simply worthless, and believing that they have really chosen an objectively good (much less the best!) layout is folly.

How close are we, really?

Here is another example of a layout based on a cost model (see the "Keyboard Effort Grid") that I think is well-intentioned but difficult to justify. To be fair, a lot of further work was done on the resulting layout, so I don't know to what extent the original cost model survived.