

1 [SQUEAKING]
2 [RUSTLING]
3 [CLICKING]
4 OK.
5 So just to let you know what
6 this notebook is and isn't,
7 this notebook is
8 kind of meant to let
9 you see how automatic
10 differentiation is
11 kind of magical in a way.
12 That's kind of the real purpose.
13 You'll start to get a
14 bit of a feel for how
15 forward mode works.
16 And what I'd like to
17 emphasize is to what extent
18 this is possibly more computer
19 science than mathematics.
20 We all have this notion
21 that whatever courses
22 that are now being taught in
23 computer science that maybe
24 used to be taught
25 in math courses,
26 like probability
27 statistics, I think
28 everybody agrees that calculus
29 lives in math departments
30 all over the world.
31 Lots of other math
32 subjects are being
33 hijacked by engineers, computer
34 scientists, and so forth.
35 But calculus, that's sacred.
36 That belongs in mathematics.
37 Well, here's a
38 case where calculus
39 is as much of a computer science
40 topic as it is a math topic.
41 I think that's kind of what
42 fascinated me most about this.
43 Oh, gosh.
44 I first put this
45 together in 2017.
46 Is it really 2023 now?
47 Six years later.
48 Well, so it's an oldie but
49 goodie, but I promise you
50 you'll like it just the same.

51 And I do like to
52 tell people that I
53 used to go to
54 conferences, and I would
55 hear people talking about
56 automatic differentiation.
57 People talked about
58 it before it was hot.
59 I mean, it became hot
60 because of machine learning.
61 But a couple of decades before
62 machine learning, people
63 would do it, and they would
64 do it on the sidelines.
65 Nobody paid attention back then.
66 It didn't have sort of the big
67 excitement that it has today.
68 But I would go to
69 conferences, and somebody
70 would get up and talk about it.
71 And I don't know.
72 I would read my email or tune
73 out or work on my own math
74 or something.
75 I didn't really pay attention.
76 And so I missed the boat.
77 I didn't appreciate
78 what I thought
79 was most important about
80 automatic differentiation.
81 I made a jump in my
82 mind of what it is.
83 And I figured it was something
84 symbolic, like Mathematica,
85 Wolfram Alpha.
86 We've all memorized
87 tables of derivatives.
88 Here's a small table
89 of derivatives.
90 I figured that if
91 I could memorize it
92 when I learned calculus,
93 then a computer could be
94 taught to do this thing also.
95 So maybe that's what it is.
96 That's great.
97 I could do it.
98 A computer could do
99 it better than me.
100 Fine.

101 I didn't care.
102 Turns out that's not what
103 automatic differentiation is.
104 So then I said to myself,
105 maybe I got it wrong.
106 I'm just guessing anyway,
107 and I don't really care.
108 But maybe it's some sort
109 of numerical difference,
110 like we do to check our answers.
111 I want the derivative
112 at this point x ,
113 so I can do a forward
114 derivative and get
115 the slope of the tangent.
116 Or I could do a
117 backward derivative,
118 go backwards and get
119 the slope of this--
120 I guess it's a secant,
121 to be technical--
122 but the slope of this line.
123 Or I could even do a
124 central difference,
125 which connects these two dots.
126 And the whole big
127 deal, of course,
128 is when you do that, as Steven
129 explained, what's a good δ ?
130 In math, you want
131 the δ to go to 0 .
132 That's the very limiting
133 definition of a derivative.
134 But on a computer, if δ
135 gets too small, as you've seen,
136 you get that catastrophic
137 cancellation happening.
138 And so numerical
139 analysis is kind
140 of about what's a good δ .
141 And Steven's basically
142 said that something
143 on the order of 2^{-26} ,
144 to the minus 26,
145 which is the square root
146 of double precision machine
147 epsilon is a good rule of thumb.
148 You might remember that curve
149 where the error went down,
150 like an absolute value

151 sign, as it got smaller,
152 and then it went up again.
153 And the best one
154 was at the bottom.
155 That's the best delta x .
156 But the key thing
157 that's interesting
158 is that automatic
159 differentiation is not this,
160 and it's not that.
161 And so what could it be?
162 And the way I'd like to show
163 people of what it is is I'd
164 like to start with a simple
165 example of it in action.
166 And this is where, at first,
167 it's going to look like magic.
168 Nothing up my sleeve.
169 And then I will explain
170 to you how it worked.
171 So I'm going to take one of
172 the oldest algorithms, one
173 of the oldest interesting
174 algorithms known to mankind,
175 the Babylonian square
176 root algorithm, which
177 I think maybe you've all seen.
178 But if you haven't, you
179 start with a guess t
180 to square root of x .
181 So I've got a t ,
182 which I'm hoping
183 is close to square root of x .
184 And if I take t and x over t --
185 if t was too small--
186 x over t , of course, will be
187 kind of on the large side.
188 If t was a guess--
189 between t and x over t ,
190 one's on the large side,
191 one's on the small side.
192 So why don't we just
193 take the average?
194 And then so this is something
195 we could keep doing.
196 And that's the
197 Babylonian algorithm
198 that converges to
199 the square root of x .
200 Very simple algorithm.

201 It was known for
202 thousands of years.
203 How complicated could it be?
204 But I still think it was
205 pretty clever, the Babylonians.
206 I mean, they didn't have Julia.
207 I mean, I thought it
208 was pretty clever.
209 So in any event, just
210 for simplicity sake,
211 I'm going to start
212 at 1 for no reason.
213 I'm just doing an example.
214 So I'm going to start it at 1.
215 So I have 1 and x over 1,
216 and I'll take the average.
217 And then, by default, I'll do 10
218 iterations of t plus x over 2.
219 So even if you
220 don't speak Julia,
221 I think this algorithm
222 is easy to understand.
223 Just it's an input x .
224 And by default, we run 10 times,
225 but you can give an argument
226 and make it run more times.
227 And so let's go
228 using ForwardDiff.
229 Let ForwardDiff do it.
230 And of course, with our
231 modern view of the world,
232 we know how to take square root.
233 Everybody in this building,
234 everybody in this institute
235 could take the
236 derivative square root.
237 It's one half over
238 the square root.
239 And so we get the derivative.
240 But before we do
241 that, let's actually
242 make sure that the code works.
243 Let's actually check
244 the Babylonian algorithm
245 that I wrote.
246 The second one is Julia's
247 built-in square root of pi.
248 And here's the
249 Babylonian calculating
250 the square root of pi.

251 And I guess this is pretty
252 convincing that the Babylonians
253 knew what they were doing.
254 I mean, we could do
255 the same thing with 2.
256 We could run the
257 Babylonian algorithm,
258 except for some
259 little bit over here.
260 Maybe the last bit.
261 We basically get the
262 same answer with Julia's
263 built-in square root and
264 the Babylonian algorithm.
265 Let's skip this for a minute.
266 I think that's not important.
267 So just checking the algorithm
268 for the moment, not even
269 the derivative.
270 So I'm actually going to plot a
271 few iterates of the Babylonian
272 algorithm just so you can see.
273 So the first iteration,
274 maybe remember,
275 it was $1 + \frac{x}{2}$.
276 I still want to call
277 this a linear function.
278 But in this class, I have to
279 call it an affine function.
280 Yeah, I don't know what
281 the real terminology is.
282 When do I get to say that a
283 first-degree polynomial is
284 a linear function?
285 I think there's a
286 context as to whether I'm
287 calling it a map or a function.
288 But any event, this
289 thing's a line,
290 however you want to call it.
291 The first step of the Babylonian
292 algorithm, iteration 1,
293 is given x , compute $1 + \frac{x}{2}$.
294 plus x divided by 2.
295 And that's a
296 first-degree function.
297 The second iteration,
298 I'm plotting it, is here.
299 And then it gets
300 closer and closer

301 to the sideways parabola, which,
302 of course, is the square root.
303 In fact, by iteration
304 5, you can't even
305 see much of a difference
306 with your eye.
307 So iteration 4 is the purple.
308 And iteration 5, you
309 can't even see it
310 because the black
311 parabola's on top of it.
312 And so just to kind
313 of convince you
314 that the Babylonian
315 algorithm works.
316 Anyway, I like Plotly.
317 I love doing this all day.
318 I can go left and right
319 and look at these numbers.
320 I love this.
321 So I like interactive things.
322 Now what I'm going to do is--
323 let me see.
324 I change this over the years.
325 So 3, 4, 5, 6.
326 In about nine lines,
327 I'm going to create
328 a function that will
329 calculate the derivative
330 of the Babylonian algorithm.
331 And nowhere will I teach it one
332 half over the square root of x .
333 I will not do that.
334 You'll see there'll be
335 no finite differences.
336 And there will be no symbolic--
337 it's going to be
338 by magic, but we're
339 going to get the right answer.
340 So are you ready?
341 So I'm going to do
342 it in nine lines.
343 So here's three lines.
344 I'm going to create
345 a Julia type.
346 I'm going to call it
347 capital D. Maybe some of you
348 have heard this word.
349 D is for dual number, so that's
350 why we're going to use the D.

351 And we're basically
352 going to keep
353 a function derivative
354 pair, an ordered pair.
355 And so this is nothing
356 but a container
357 to be able to keep two floats.
358 But which floats
359 am I going to keep?
360 I'm going to have the value
361 of a function at a point
362 and the derivative of
363 that function at a point.
364 All right.
365 I've used up three
366 of my nine lines.
367 And everybody agrees there's
368 no finite difference,
369 no symbolic answer, right?
370 So I've got six more lines here.
371 Yeah, let me just show
372 you what I've done here
373 before I do anything else.
374 So if I wanted to--
375 if I want to create
376 one of these objects,
377 I would have to put in
378 a tuple, like D of 1, 2.
379 This line lets me
380 remove the parentheses,
381 which are sort of--
382 I just find them annoying.
383 This line doesn't count.
384 But you see, I've
385 got this dual number.
386 It doesn't do anything yet.
387 I can't add dual numbers yet.
388 If I try, it gets mad at me.
389 Look, plus not defined.
390 All I can do is
391 define a dual number.
392 That's it.
393 It's just a pair of numbers.
394 Can't do anything at all
395 with it other than store it.
396 But here what I'm going
397 to do is create a--
398 oh, I don't need
399 this greater than.
400 That's why.

401 I could have one fewer line.
402 I'm going to comment
403 this one out.
404 I think somebody once asked
405 me to define it for greater.
406 But yeah, I don't
407 need this line.
408 So I'm actually going to
409 have 3 plus 5, 8 lines.
410 So let me tell you about
411 these next five lines.
412 Mainly I want to define, add,
413 and divide on a dual number.
414 I don't need minus
415 and times yet,
416 because if you look at
417 my Babylonian algorithm,
418 if you remember the
419 algorithm-- where
420 is the Babylonian algorithm?
421 I do a plus and a
422 divide and nothing more.
423 Later on, I'll add times and
424 minus, but I don't need it yet.
425 So I wanted to define
426 a plus and a divide.
427 I don't need this one either.
428 This thing could go away.
429 In Julia, if you want to
430 overload plus and divide
431 and a few other things, you
432 have to import it from base.
433 So this is just like
434 a Julia detail thing
435 that says give me permission
436 to redefine plus and divide
437 and a few other things.
438 And what do I want to do?
439 When I plus a couple of
440 dual numbers-- here let
441 me just do this one
442 only, just so you see.
443 I'll execute only the plus.
444 Now I can add dual numbers.
445 And it's just going to
446 be like adding vectors.
447 So I'm just going to add the
448 first element of the tuple, 2
449 and 3, and the second element.
450 So now I can add tuples,

451 but I can't divide them yet.
452 So this dot notation, this
453 broadcast or pointwise
454 notation, says
455 that basically add
456 the two parts of x and the
457 two parts of y , like a vector.
458 So this adds 2 and 3 to
459 get 5 and 3 and 4 to 7.
460 All right.
461 Now let me bring in the divide.
462 I can't divide yet, by the way.
463 I can try, but it'll
464 get mad at me, you see.
465 Oh, I must have executed it.
466 Oh, did I execute it?
467 I'm sorry.
468 Then it took away my--
469 well, whatever.
470 All right.
471 I guess I must have
472 executed it, so I'm not
473 getting away with it.
474 But here it doesn't matter.
475 So this is 2
476 divided by 3 is $2/3$.
477 But notice this
478 isn't 3 divided by 4.
479 So I have a different rule.
480 So the add rule is just
481 adding up a vector,
482 but the divide rule is a
483 little more complicated.
484 STEVEN G. JOHNSON:
485 By the way, Alan,
486 did you want to share
487 your screen on Zoom?
488 I forgot about--
489 ALAN EDELMAN: Oh, my gosh.
490 I didn't share my screen, so
491 you don't see a thing I've done.
492 STEVEN G. JOHNSON: No, I can
493 see it behind you on the--
494 ALAN EDELMAN: Oh, that's funny.
495 OK.
496 There we go.
497 All right.
498 Better?
499 All right.
500 So divide.

501 So everybody, of course,
502 remembers the quotient rule
503 from calculus?
504 When I think of it, I hear
505 my math teacher singing it.
506 It was denominator times
507 the derivative of numerator
508 minus the numerator
509 times the denominator,
510 or was it denominator squared?
511 I don't know.
512 Did your teacher sing it to you?
513 How did you sort of
514 memorize the quotient rule?
515 Anybody have a good song for it?
516 Anyway, you drill
517 it into your head.
518 vdu minus udv over v squared,
519 or denominator, d numerator.
520 I mean, I don't know.
521 You may have heard
522 it different ways,
523 but you all know it, right?
524 This thing over here, the
525 quotient rule, everybody
526 knows it.
527 I'm just extracting
528 the parts from--
529 so y is the denominator.
530 And so 1 is the value.
531 So this is the denominator.
532 x is the numerator.
533 And 2 is the derivative.
534 So it's the denominator times
535 the derivative numerator
536 minus the numerator times
537 the derivative denominator
538 over the denominator squared.
539 So that is what
540 I'm going to teach.
541 I'm going to teach Julia how
542 to essentially add derivatives,
543 which is just add, and how
544 to divide derivatives, which
545 is just the formula you know,
546 just apply it at a point.
547 And so this division is
548 using all four numbers
549 so that it can get
550 the denominator

551 times the derivative of the
552 numerator minus the numerator
553 times the derivative
554 of the denominator
555 over the denominator
556 squared, you see.
557 And that's what
558 this one ninth is.
559 All right.
560 That's it.
561 Just these 3 plus--
562 what did I say?
563 3 plus-- oh, I haven't told
564 you about convert and promote.
565 These are a little bit more
566 sort of technical details.
567 But do you know how if you add
568 a complex number and a real,
569 like if you go 3 plus
570 $4i$, and you add 7?
571 Now what's really going
572 on is that that 7,
573 in some abstract sense, is
574 being converted into 7 plus $0i$.
575 And then you add the real
576 parts and the imaginary parts.
577 Everybody does
578 that all the time.
579 So we want to do
580 that sort of thing
581 where if you have a real
582 number, we want to think of it--
583 if you have a scalar,
584 we want, in effect--
585 a constant is really
586 what's going on here.
587 We want to think of this as the
588 constant x , where the value is
589 x and the derivative is 0 .
590 And we want that to
591 be kind of automatic
592 because it would be nuisance
593 to type it all the time.
594 So that's what that does.
595 And then the promote rule
596 says that if you give it
597 a number, when you see it in
598 the context of a dual number,
599 everything should be
600 promoted to the dual number.

601 Just like it happens with
602 complex numbers, where,
603 like I said, 3 plus
604 $4i$ plus a real number,
605 you'd put that $0i$ in your
606 mind or on a computer.
607 But you would promote
608 everything into the complex land
609 and then do the addition.
610 So those are two
611 necessary things.
612 And now let me go ahead and
613 run the Babylonian algorithm.
614 And without changing
615 the algorithm--
616 remember the algorithm
617 takes a scalar in.
618 Let's see it again.
619 Let's find it.
620 The Babylonian algorithm, which
621 is up here, it takes a scalar.
622 I'm not going to
623 rewrite the algorithm.
624 I'm just going to
625 feed it something new,
626 something different
627 from a scalar.
628 I'm going to feed
629 it a dual number.
630 And so let's do it.
631 Where did it happen here?
632 So I'm feeding the Babylonian
633 algorithm 49 comma 1.
634 This is how you seed--
635 we'll talk more about
636 seeding the start
637 of the story with the number 1.
638 Or if it was matrices,
639 it would be the identity.
640 And we get the
641 square root being 7.
642 Yep, that's good.
643 The square root of 49 is 7.
644 And the derivative,
645 which you all
646 know-- we could let
647 Julia tech it for us--
648 is one half over the square root
649 of x is this number right here.
650 So whatever one half over

651 7 is, $1/14$ or something.
652 So this is the number $1/14$.
653 You should be astounded
654 by this, that I
655 took an original piece of
656 code without a rewrite,
657 and I fed it this
658 funny kind of argument.
659 And all that argument did was
660 it knew the quotient rule,
661 and it knew the sum rule.
662 And I got the right
663 answer for the derivative,
664 not symbolically and not
665 with finite differences.
666 Isn't that amazing like
667 that's even possible?
668 Wouldn't that blow your
669 calculus teacher's mind
670 that this could happen?
671 Here's another example where I
672 do it with pi, just in case 7
673 wasn't convincing enough.
674 So this would be the
675 square root of pi
676 and 1 over 2 square root of pi.
677 And this is the way
678 done with Julia.
679 And you could check the numbers.
680 You see it all works.
681 And in fact, what you can
682 do is actually look at--
683 what's happening is the
684 Babylonian algorithm
685 was an iteration.
686 And so somehow, this square root
687 is the result of an iteration.
688 We do 10 steps of an iteration.
689 And so at each time, we must
690 be getting closer and closer
691 to the derivative.
692 So just like we get closer
693 to the square root, somehow,
694 by feeding this in, we must
695 be getting closer and closer
696 to the derivative
697 of the square root.
698 And in fact, I could plot
699 each step of the algorithm.
700 So remember the first

701 algorithm was first degree.
702 I still want to say linear,
703 but I'll say first degree.
704 And so its derivative,
705 of course, is a constant.
706 It's just the constant
707 one half, in fact.
708 So there it is.
709 And here are a
710 couple of iterations.
711 And I also plotted one
712 half over the square root
713 of x, the true answer, the
714 reciprocal of the parabola,
715 in effect.
716 And you could see that it's
717 heading closer and closer.
718 And pretty quickly, the
719 eye can't even see it.
720 So this doesn't explain
721 to you how it works,
722 but maybe it kind of adds to the
723 convincing nature of the fact
724 that it does work, and
725 it's still mysterious.
726 I could say a little bit more.
727 I'm going to tell
728 you how it works.
729 But before I do, I'd like
730 to show off a few things.
731 I don't know how well
732 this works these days.
733 But I do like to tell
734 people that, in Julia, you
735 can see assembler.
736 Nobody reads assembler.
737 Anybody here read assembler?
738 Anybody here actually-- you
739 do or have or a little bit?
740 One person is willing to admit
741 that they do it a little bit.
742 Some computer science classes
743 at MIT teach you this stuff.
744 Most people never look at this,
745 don't want to look at this.
746 The thing that I
747 like to just mention
748 is that, in Julia, the
749 assembler is short.
750 And so this is the assembler

751 for this derivative code,
752 this kind of derived code.
753 And short assembler is more or
754 less correlated with fast code.
755 And so not only does it
756 get the right answer,
757 but this sort of game is
758 also quite fast in Julia.
759 And that's kind of a nice
760 thing to be able to have.
761 So I'm still not going
762 to tell you how it works,
763 but I'm going to grab SymPy.
764 So this is Python
765 symbolic program.
766 There's is a Julia
767 symbolic program,
768 but I don't completely
769 trust it yet.
770 Maybe it's ready for prime
771 time, but I did this originally
772 with--
773 I wrote this before there
774 even was Julia symbolic.
775 And anyway, it just works
776 so well, I would take it.
777 And so one of the things
778 that's interesting is to ask--
779 how should I say this?
780 I'm going to tell you
781 something that's mathematically
782 equivalent to what
783 we're doing, but I
784 don't want you to
785 get the impression
786 that this is how it's computed.
787 So let's talk about not
788 the derivative yet but just
789 the Babylonian algorithm.
790 You remember that this is the
791 function at the first step, x
792 plus 1 over 2.
793 I can use Julia's
794 ability to overload
795 to run it on a symbol x .
796 And then I could see what
797 there is at the second step
798 or at the third step.
799 And so in a way, at
800 whatever this is--

801 if this is the first step,
802 second, third, fourth, fifth.
803 At the fifth step, the
804 Babylonian algorithm
805 exactly computes this
806 rational function.
807 It's a 16th-degree polynomial
808 over a 15th-degree polynomial.
809 But don't get the wrong idea.
810 Nobody in the real world is
811 calculating the coefficients
812 of this polynomial.
813 I mean, these coefficients, we
814 wouldn't want to store them.
815 They'd be unwieldy to work with.
816 But as a mathematical
817 sense, the fifth step
818 of the Babylonian
819 algorithm is calculating
820 exactly this function.
821 And the plots tell us
822 that this crazy function,
823 the 16th over
824 15th-degree polynomial,
825 is not a bad approximation to
826 the square root of x , at least
827 visually on the graph.
828 So this is pretty good
829 for square root of x .
830 That's what we've seen.
831 We could talk about where it
832 is good, where it isn't good.
833 But the point is that what it's
834 computing is this function.
835 And we could do the same
836 game for the derivatives.
837 So the first derivative here is
838 the coefficient of x as a half,
839 that constant.
840 And we can see what's being
841 computed exactly here.
842 This is a ratio of
843 30th-degree polynomials.
844 And again, I want to
845 stress we are not--
846 I'm just building
847 this up just for fun.
848 We are not in the
849 algorithm getting
850 literally these coefficients.

851 They're too big anyway
852 for working with.
853 But in a mathematical
854 sense, the fifth step
855 of this derivative
856 Babylonian algorithm
857 is calculating
858 exactly this thing.
859 And so this has to be
860 some sort of approximation
861 to one half over the square
862 root of x , the derivative
863 of square root of x .
864 This is what that is.
865 So let me get a little
866 closer as to how-- now you
867 must be wondering.
868 I hope you're all kind of
869 sitting in your seats saying,
870 how is this working?
871 What's happening here?
872 And to get you a little
873 bit kind of closer,
874 let me do what people used
875 to do in the old days.
876 People used to take derivatives
877 of functions by hand.
878 Before this became
879 automatic, people
880 would take derivatives
881 of functions by hand.
882 And so I'm going to
883 do that for you here.
884 I'm going to create
885 a Babylonian
886 algorithm, the derivative
887 of the Babylonian algorithm.
888 And you'll recognize that
889 this line and this line
890 are the original algorithm.
891 And below it, I'll create
892 derivative variables, t prime.
893 And so t prime, the derivative
894 of this is, of course, a half.
895 The derivative of this line
896 of code, well, what is it?
897 It's t prime plus
898 the denominator
899 times the root of the numerator,
900 which is $1 - x$ times t

901 prime over t squared.
902 So if you check, this is the
903 ordinary calculus derivative
904 with respect to x.
905 So t prime is the
906 derivative respect to x.
907 So this is the ordinary
908 calculus derivative.
909 And we're doing that
910 at each and every step.
911 And people used to do that
912 by hand, that you would--
913 in other words, you don't take
914 the derivative analytically
915 of the big thing.
916 Rather you take the derivative
917 of each line of code.
918 And then you have faith
919 that if you do that,
920 you'll get the derivative
921 of the big thing
922 that you wanted on the outside.
923 And you'll see that,
924 of course, it works.
925 Adding these couple of lines
926 of code with just-- this
927 is now scalars.
928 There's no dual numbers here.
929 This will give me one half
930 over the square root of pi
931 just by taking the derivative
932 of every line of code.
933 And so you might realize
934 that this is actually
935 an iteration for the
936 derivative of square root of x,
937 an iteration that we stop at--
938 we stop it at 10, by default.
939 We could take more
940 steps, but this
941 is an iteration for the
942 derivative square root,
943 obtained completely by taking
944 the derivative of every line.
945 And so that's kind
946 of what happened.
947 And so when I take the
948 Babylonian of D, x, 1,
949 in effect, I am using the
950 magic of Julia's ability

951 to do dispatch and overload
952 and all those fancy words.
953 But to use simple English,
954 I am using the fact
955 that I don't have to rewrite
956 the code to get the derivative.
957 I just need the code to know
958 the rules of taking derivatives
959 of every operation that--
960 more atomic operations
961 at the lowest
962 level and rely on the computer
963 to piece it all together.
964 Because humans are really
965 bad at this sort of stuff.
966 They make mistakes all the time.
967 It's worse than long division.
968 I mean, no matter how good
969 you are at long division,
970 humans just make mistakes.
971 We just do.
972 And so the trick
973 is if you wanted
974 to teach a computer--
975 if you want
976 to get the answers to
977 a division problem,
978 we humans have taught computers
979 to do the division for us
980 so we don't have to.
981 And this is what's going on
982 with automatic differentiation.
983 We teach the computer
984 to do the atomic steps
985 and then let it just
986 go through the motions.
987 So the derivative goes in
988 before the JIT compiler,
989 and we get efficient code.
990 So there's a notational
991 trick, which is rather nice,
992 which is instead of taking the
993 dual number, which is $a + b\epsilon$, we
994 can write $a + b\epsilon$.
995 And in effect, what we're
996 doing is the same thing that--
997 on the first week of
998 class, when Steven said,
999 oh, let's just write
1000 everything as $a + bdx$.

1001 Just write everything
1002 to first order.
1003 Physicists do this all the time.
1004 They write everything
1005 to first order,
1006 and they throw away higher-order
1007 terms just all the time.
1008 So in effect, what's
1009 happening on the computer is
1010 we're treating every computation
1011 as a first-order computation.
1012 And then the basic rules
1013 are-- let me just see.
1014 There was one version of
1015 this that's broken, but let
1016 me see if this is right.
1017 I think this is
1018 the right version.
1019 So the basic rules--
1020 every computation on a computer
1021 that's ever been written always
1022 can come down to plus,
1023 minus, times, and divide.
1024 Even square root is implemented
1025 somewhere as plus, minus,
1026 times, and divide.
1027 So in effect, if
1028 you wanted to do,
1029 you can get automatic
1030 differentiation
1031 just by having
1032 these basic rules.
1033 This is all you need.
1034 Now as a matter of practice,
1035 we try to intercept it all.
1036 We're happy to teach
1037 sine and square root
1038 and cosine because who wants--
1039 whatever method is being
1040 used to compute the sine--
1041 and it's not Taylor
1042 series, by the way--
1043 but whatever method
1044 is being used,
1045 we don't want it to go
1046 through all this work.
1047 So we teach it things.
1048 But in principle, all
1049 you need are these rules,
1050 and you can take the

1051 derivative of anything
1052 in the world on a computer.
1053 This is all you need.
1054 Here's the sum and minus
1055 rule, the multiplication rule,
1056 which if you look at this right,
1057 maybe if you squint correctly,
1058 this is the $u dv + v du$ rule,
1059 the product rule that you all
1060 learned in calculus.
1061 And we kind of repeated
1062 it in its matrix context
1063 in this class.
1064 This is $u dv + v du$, and
1065 this is the quotient rule.
1066 This is denominator times
1067 degree of the numerator
1068 minus numerator times
1069 degree of the denominator
1070 over the denominator squared.
1071 It's just kind of rewritten
1072 in this first-order kind
1073 of notation.
1074 But these are rules that you all
1075 learned in first-year calculus.
1076 And I'll even point out that
1077 you could do this symbolically.
1078 You don't even have
1079 to remember the rules.
1080 You could actually derive the
1081 quotient rule on the computer
1082 by just--
1083 this says basically take a
1084 series around $\epsilon = 0$,
1085 and give me two terms, please.
1086 No more.
1087 So just give me to the
1088 first order, an ϵ .
1089 And here you see.
1090 You get the quotient and the
1091 quotient rule from calculus.
1092 So this is one way to
1093 get your hands on that.
1094 If you wanted the
1095 product rule, I
1096 guess I could have
1097 just done this.
1098 And you get the
1099 $u dv + v du$ rule.
1100 So that's how you

1101 can get the rules.
1102 So I'm going to
1103 do something fun.
1104 I am going to tell Julia--
1105 this is Julia magic
1106 that says to print
1107 a dual number with epsilons.
1108 And so now when I
1109 type a dual number,
1110 you remember it was
1111 just with the Ds.
1112 Once I execute this
1113 command, I could see it
1114 in a way that's nice and human.
1115 So I told you this was a
1116 function derivative pair,
1117 but you could also think
1118 of it, if you like,
1119 as a first-order expansion of--
1120 it could be a first-order
1121 expansion of a function.
1122 It could be the
1123 first-order expansion
1124 of x squared around x equals 1.
1125 So let's go ahead and
1126 add these last two rules.
1127 Remember I only did
1128 plus and divide.
1129 I might as well add the--
1130 this seems like a good time
1131 to do minus and times.
1132 And you see that if I do
1133 the dual number 1 and θ ,
1134 I get this.
1135 Well, actually let me ask you.
1136 I'm not going to hit Enter yet.
1137 Tell me what I should see when
1138 I hit Return, before I do it.
1139 Who's quick?
1140 What's the first thing I'll
1141 see before the epsilon?
1142 Let me start with
1143 the θ -th-order term.
1144 What will I see?
1145 Just shout it out.
1146 AUDIENCE: 4.
1147 ALAN EDELMAN: 4.
1148 And then what's the next term?
1149 AUDIENCE: 2, 4.
1150 ALAN EDELMAN: 2 times 2.

1151 4.
1152 Yep.
1153 You guys got it.
1154 OK.
1155 I changed the output.
1156 I might as well go
1157 the whole direction.
1158 Why don't I make the input also?
1159 D, θ , 1, I'll call it epsilon.
1160 And so now I can actually
1161 input epsilons too.
1162 Not just see it as an output,
1163 but I can do it as an output.
1164 So epsilon squared, of
1165 course, is second order.
1166 So we just get rid of it.
1167 By the way, just something fun.
1168 I actually never
1169 defined how to square.
1170 You'll notice I define
1171 times, but I never
1172 define square for dual number.
1173 But this is sort of a
1174 little bit of a lesson,
1175 but a good software
1176 system would be one
1177 where when you square something,
1178 it actually replaces it
1179 with a thing times itself.
1180 So that a matrix square is
1181 a matrix times a matrix,
1182 and a scalar square is
1183 a scalar times a scalar.
1184 And in Julia, for
1185 whatever reasons,
1186 a string times a string is a
1187 concatenation of the string.
1188 So a string squared--
1189 I don't even know if
1190 this works anymore.
1191 I have a feeling
1192 it doesn't work.
1193 It's not a number.
1194 This is going to fail,
1195 but it shouldn't fail.
1196 Actually I think this
1197 is going to fail.
1198 Oh, forget it.
1199 It does work.
1200 So multiplying two strings

1201 will concatenate them,
1202 and squaring it concatenates it.
1203 And so if you have sort of
1204 a novice computer system,
1205 every time you
1206 have another type,
1207 you define another square.
1208 But if you have a
1209 good computer system,
1210 then the square inherits
1211 it from multiply,
1212 and then you just have
1213 to define multiplication.
1214 What should I get here when
1215 I go 1 over 1 plus ϵ ?
1216 And again, nothing symbolic.
1217 This is actually happening
1218 completely numerical,
1219 by the way.
1220 But what should I get
1221 when I hit Return?
1222 What should I see?
1223 Anybody?
1224 You're smiling.
1225 You think you know the answer?
1226 AUDIENCE: I guess it's
1227 just written there.
1228 ALAN EDELMAN: So
1229 I'll give you a hint.
1230 What's written there
1231 is not what you'll see.
1232 AUDIENCE: [INAUDIBLE] minus 1 .
1233 ALAN EDELMAN: Yeah,
1234 how would it appear?
1235 Just read it to me
1236 how it would appear.
1237 AUDIENCE: I'd guess
1238 1 plus minus ϵ .
1239 ALAN EDELMAN: There you go.
1240 1 plus minus 1 ϵ .
1241 You could have just
1242 said 1 minus ϵ .
1243 I would have accepted that.
1244 All right.
1245 Doesn't that look like
1246 symbolic mathematics?
1247 But it's not.
1248 This whole thing
1249 happened through these--
1250 it's all numerical.

1251 There's nothing symbolic at all.
1252 And this is another thing
1253 that people are saying,
1254 that there's becoming
1255 more and more
1256 of a blurring between the
1257 symbolic and the numerical,
1258 and that numerical
1259 stuff is starting
1260 to look more and more symbolic.
1261 But it's not symbolic.
1262 All right.
1263 What's the answer here?
1264 I'm not using any
1265 weird packages.
1266 Everything that I'm using was
1267 defined right in front of you.
1268 I'm not using
1269 ForwardDiff or anything.
1270 Everything here is just
1271 pure, simple Julia.
1272 You saw it.
1273 There's nothing up my sleeve.
1274 What should this answer be?
1275 AUDIENCE: [INAUDIBLE]
1276 ALAN EDELMAN: I'm sorry.
1277 AUDIENCE: [INAUDIBLE]
1278 ALAN EDELMAN: You're right.
1279 $1 + 5\epsilon$.
1280 OK.
1281 And this one,
1282 unfortunately, won't work.
1283 Oh, it does work.
1284 Oh, that's amazing.
1285 I don't know why that works.
1286 All right.
1287 Never mind.
1288 I didn't think we could
1289 take negative powers,
1290 but I guess we could.
1291 All right.
1292 I'm going to stop.
1293 You could do this
1294 with n -th roots.
1295 You could do lots
1296 of other things.
1297 But I think this is a
1298 good time for a break.
1299 And you guys get the right idea.
1300 So now you're starting to see.

1301 If I were to summarize-- and
1302 I know it's still a little bit
1303 magical, but I think
1304 you'll see that roughly
1305 how this works is that
1306 one way or another,
1307 we're giving the rules of
1308 plus, minus, times, and divide.
1309 And we're writing programs.
1310 And then these programs
1311 are, in effect--
1312 they're not really
1313 rewriting themselves.
1314 What's really happening
1315 is that every time you
1316 execute a plus, a minus,
1317 times, and a divide,
1318 it's doing not just
1319 the basic operation
1320 that you'd all expect,
1321 but it's also carrying
1322 along the derivative as well.
1323 And the way Julia works,
1324 Julia will actually
1325 look at that divide and say
1326 I'm not dividing scalars.
1327 I'm dividing dual numbers.
1328 Or if it sees a star,
1329 I'm not multiplying.
1330 How does Julia know what to do?
1331 When it sees two matrices,
1332 matrix, star, matrix,
1333 it knows, because it's a
1334 matrix, to do matrix multiply.
1335 So here when I did
1336 dual numbers, I
1337 taught it to do this
1338 dual-number thing.
1339 And once Julia
1340 knows how to do it,
1341 it'll just carry
1342 all the way through.
1343 And in effect, this
1344 is really the magic
1345 of great software,
1346 where you just
1347 can define some
1348 atomic operations,
1349 and the whole thing
1350 kind of composes

1351 itself almost by magic.
1352 And in a way, it's
1353 almost opposite
1354 from what we teach students in
1355 a lot of classes, where we want
1356 to teach-- the
1357 old-fashioned thing
1358 was to teach a student to
1359 carry through every operation
1360 and be really competent at it.
1361 In a way, the modern
1362 world is to teach students
1363 how to not have to
1364 think, rather how
1365 to build a system that is so
1366 simply designed that it just
1367 works.
1368 And actually, to
1369 build a simple system
1370 is what takes the
1371 real human cleverness,
1372 if that sounds not like
1373 some sort of contradiction.
1374 But that's what it takes.
1375 All right.
1376 I'm a little late for the break.
1377 But after the break,
1378 on the Blackboard,
1379 I'm going to go into
1380 more detail about forward
1381 and reverse mode,
1382 automatic differentiation.
1383