

	Given	New		comments
			<b>Focus</b>	
1.			attack evaluation	
2.	we compare our targeted attacks to	the best results previously in prior publication	for each of the three distance metrics	Don't you agree – the crucial information here is the final condition of the criterion for comparison?
3.	we	re-implement	Deepfool, fast gradient sign, and iterative gradient sign	
4.	for fast gradient sign, we	search over $\epsilon$ to find	the smallest distance that generates an adversarial example	
5.			failure is returned	
6.	no $\epsilon$	produces	the target class	
7.	our iterative gradient sign method	is	similar	
8.	we	search over $\epsilon$	(fixing $\alpha = 1/256$ )	<p>This portion of the text contains three examples of parenthetical material telling us just what Carlini and Wagner think: this clause, clause No.9, and clause No.77.</p> <p>Now, this usage would appear to contradict what I have said in the previous post about parentheses, namely, that they footnote material. However, this does not necessarily entail that the material is downgraded in importance, especially when that material is positioned at the end of the clause. Therefore, you writers, pay attention: If you are putting something in parentheses, but the material is not very important, then do not leave</p>

				it to the end of the clause – look to the middle of the clause to parenthesize it, or instead, include it in a real footnote.
9.		return	the smallest successful	
10.	for JSMA, we	we use the implementation CleverHans [35] with only slight modification	(we improve performance by 50x with no impact on accuracy)	
11.	JSMA	is unable to run on ImageNet	due to an inherent significant computational cost	
12.			recall	
13.	JSMA	performs search for a pair of pixels $p, q$ that can be changed together that make	the target class more likely and other classes less likely	<p>I've complained about this sentence <a href="#">before</a>, and now I am going to complain about it again.</p> <p>Here you see further evidence for why this sentence just does not work as well as it might. Basically, the Focus being singled out is not the Focus which a reader will naturally construe in his or her head. That Focus in a reader's head is going to be, namely, <i>performs search for a pair of pixels <math>p, q</math> that can be changed together that make</i>. I mean, just think about it. The thing at issue here is not the classes but the computational cost. Therefore, it is the information about JSMA's search mechanism that matters most – <i>that</i> is the Focus.</p> <p>Of course, it is possible to say that the Focus does not occur at the end of the clause but instead in the middle. So, for the current clause, we just color the middle box red and call it the Focus, and we color the rightmost box black and call it the remainder of the New. This is certainly allowed by the system Given-New, and if Carlini or Wagner spoke the clause, they might just speak it that way.</p> <p>But they are not going to speak the sentence. Remember, science on the line happens in silence. Besides, despite all that I've said about the system</p>

				<p>Given-New being a system separate from the system of Theme-Rheme, the fact of the matter is that the two systems do tend to line up, and they tend to line up even more so in scientific prose. Therefore, as a writer, it is generally good practice to fulfill reader expectations and get the Focus to the back of the clause.</p> <p>For instance, Carlini and Wagner might have revised the current clause to this:</p> <p style="text-align: center;"><i> <sup>Given</sup> JSMA  <sup>New</sup> increases the likelihood of the target class over all other classes by performing  <sup>New</sup>: Focus search for a pair of pixels p,q which change together.</i></p> <p>That edit continues to explain why JSMA computes as it does, but explains this <i>while also</i> directing the reader's attention precisely to the information which will attract attention.</p>
14.	ImageNet	represents images as	<i>299 × 299 × 3 vectors</i>	
15.	searching over all pairs of pixels	would require $2^{36}$ work	<i>on each step of the calculation</i>	<p>I debated with myself over this New or the New <math>2^{36}</math> work on each step of the calculation. I suppose, ultimately, the question needs to be decided by a subject-matter expert. My understanding of the clause – for what it's worth – is this: The area where the work occurs is a bit less predictable than the magnitude of the work. Or, to view the same matter from the opposite end: It's more useful to know where the work occurs than how much work is required.</p>
16.	we	remove the search	<i>over pairs of pixels</i>	<p>Look at how often a Circumstance becomes the Focus. In the current clause and in the next two, it is Circumstance as Focus, Circumstance as Focus, Circumstance as Focus. That tells you: Circumstances may be peripheral to the clause, but Circumstances very often prove central to the information passing through clauses.</p> <p>Writing is complicated. The grammar is constantly running multiple systems through the same set of words.</p>
17.	the success of JSMA	falls off	<i>dramatically</i>	

18.	we	report it as failing always	on ImageNet	
19.	we	report	success	
20.	the attack	produced an adversarial example	with the correct target label	
21.	how much change		was required	
22.	failure	indicates the case where the attack was	entirely unable to succeed	<p>There are three other clauses built like this one here. They are the clauses No.4, No.13, and No.24. The structure is a noun phrase qualified by the words <i>where</i> or <i>that</i>. Other words that function in the same way are <i>which</i> and <i>who</i>. So, what is going on here?</p> <p>Well, basically, one noun is extended upon by means of this functional hinge word. In the current clause, that hinge word is <i>where</i>. The word <i>where</i> permits the clause to continue after the noun <i>case</i>.</p> <p>Now, the reason I am going into this detail here is because this structuring of the clause may cause ambiguity as to which material exactly makes the Focus. For example, the current clause may well be read as having the entire noun phrase as its Focus, so, that is, <i>the case where the attack was entirely unable to succeed</i>. And in fact, that is precisely how I read clause No.4 – just see for yourself above.</p> <p>However, for the current clause, I don't think that that is what's going on. Here I believe that a reader will only be paying special attention to the level of success, because the reader already knows that the attack has failed. So, the item of news the reader wants to learn is not the fact of failure – that's been established – but instead the definition of that failure, or perhaps rather, the degree to which the failure's occurred. The answer? <i>Entirely</i>.</p>
23.	we	evaluate on	the first 1,000 images in the test set on CIFAR and MNSIT	

24.	on ImageNet	we report on 1,000 images that were	initially classified correctly by Inception v3	
25.	on ImageNet	we approximate the best-case and worst-case results	by choosing 100 target classes (10%) at random	
26.	the results	are found	in Table IV for MNIST and CIFAR	
27.			Table V for ImageNet	
28.	for each distance metric, across all three datasets, our attacks	find	closer adversarial examples than the previous state-of-the-art attacks	
29.	our attacks		never fail to find an adversarial example	
30.	our $L_0$ and $L_2$ attacks	find adversarial examples	with 2x to 10x lower distortion than the best previously published attacks	
31.		succeed	with 100% probability	
32.	our $L_\infty$ attacks	are	comparable in quality to prior work	
33.	their success rate	is	higher	
34.	our $L_\infty$ attacks on ImageNet	are	so successful	
35.	we	can change the classification	a change that would be impossible to detect visually	In the previous post, I spoke about how facts are news, but interpretations of facts are focal news. Well, here's one more example of that in action. We learn which

		of an image to any desired label by only flipping the lowest bit of each pixel		change Carlini and Wagner mean, and knowing that is good; but then we learn what Carlini and Wagner think about such change, and knowing that is great.
36.	the learning task	becomes	increasingly more difficult	
37.	the previous attacks	produce worse results	due to the complexity of the model	
38.	our attacks	performs	even better	That little word <i>even</i> – a cognate of German <i>eben</i> , if you're interested – serves the same function as italics: Emphasize! But better still, this little adverb adds meaning which italics cannot, because <i>even</i> says too: <i>And you weren't expecting this!</i> So, while reading the current clause, the reader will be thinking something along these lines: <i>I know your attacks are better, but wow – so you're saying, the attacks are not just better, but better also in that quite unexpected case where the tasks grow more complex. Huh, pretty neat.</i>
39.	the task complexity		increases	
40.	we		have found	
41.	JSMA	is	unable to find targeted $L_0$ adversarial examples on ImageNet	This clause and the next form a happy pair. Just consider. Where this clause is saying, <i>Look at what JSMA is unable to do</i> , the next clause is saying, <i>And look here at what our attack is able to do</i> .
42.	ours	is	able to with 100% success	
43.	it	is	important to realize	
44.	the results between models		are not directly comparable	Most verbs are positive. Now, I don't mean to say they're happy or cheery. No, I mean that most verbs lack the English word for negation, <i>not</i> . This fact has the general effect of focusing negation. Therefore, in a clause like the current

				one, where the verb is in the negative, the reader learns, one, how to understand the results, and two, how not to. That makes Carlini and Wagner's interpretation crisp, and it makes it memorable.
45.	a $L_0$ adversary	must change 10 times as many pixels to switch an ImageNet classification	compared to an MNIST classification	
46.	ImageNet	has	114× as many pixels	
47.			the fraction of pixels that must change is significantly smaller	<p>An alternative reading here is just to break the clause in three:</p> <p style="margin-left: 40px;"> Given the fraction of pixels that must change  New is  New:Focus significantly smaller</p> <p>However, I see Carlini and Wagner's italics and these make me think: <i>Pay attention – look here, it's not just the pixels in any old assortment. No, it's the fraction of the pixels that counts.</i></p> <p>And when I add to this message of theirs the simple fact that the expected Focus is also very important (i.e., <i>significantly smaller</i>), then I think: <i>Wow, this clause packs information power.</i> That of course is possible, but note how this is just one of seven clauses where the entire clause is in the Focus. The others are the clauses No.1, No.12, No.27, No.48, No.61 and, No.64</p> <p>The takeaway is, You are allowed to Focus on everything.</p>
48.			generating synthetic digits	
49.	with our targeted adversary, we	can start from	any image we want	Again, the punctuation lays the focus <i>in</i> the Focus.
50.		find	adversarial examples of each given target	

51.	in Figure 6, we	show the minimum perturbation to an entirely-black image	required to make it classify as each digit for each of the distance metrics	
52.	this experiment	was performed	for the $L_0$ task previously [38]	
53.	when mounting their attack for classes 0, 2, 3, and 5, one	can clearly recognize	the target digit	
54.	with our more powerful attacks, none of the digits	are	recognizable	
55.	Figure 7	performs the same analysis	starting from an all-white image	
56.			notice	
57.	the all-black image	requires no change	to become a digit 1	
58.	it	is initially classified	as a 1	
59.	the all-white image	requires no change	to become an 8	
60.	the initial image	is already	an 8	
61.			runtime analysis	
62.	we		believe	



63.	there are		two reasons why one may consider the runtime performance of adversarial example generation important	<p>The combination <i>there</i> plus the appropriate form of <i>be</i> is called the presentative. The function of this combination is to signal that the information to come is information either from an unexpected angle or about an unexpected matter.</p> <p>So, for instance here, the reasons given by Carlini and Wagner come a bit out of the blue. I mean, sure, they give the subheading <i>Runtime Analysis</i> to signal that a new topic is now under discussion; but it won't be entirely clear just how the topic will be discussed. <i>That</i> question Carlini and Wagner resolve up front, telling us straightaway that they want to justify their interpretation of the results.</p>
64.			to understand	
65.		if the performance would be prohibitive for an adversary	to actually mount the attacks	
66.		to be used as an inner loop	in adversarial re-training [11]	
67.	comparing the exact runtime of attacks	can be	misleading	
68.	we	have parallelized the implementation of our $L_2$ adversary allowing it to run hundreds of attacks	increasing performance from 10x to 100x	

		simultaneously on a GPU		
69.	we		did not parallelize our $L_0$ or $L_\infty$ attacks	
70.	our implementation of fast gradient sign	is	parallelized	
71.	JSMA		is not	
72.	we	refrain from giving	exact performance numbers	
73.	we		believe	
74.	an unfair comparison	is	worse than no comparison	
75.	all of our attacks and all previous attacks	are	plenty efficient to be used by an adversary	
76.	no attack	takes longer than a few minutes to run	on any given instance	
77.	when compared to $L_0$ , our attacks	are $2 \times$ – $10 \times$ slower than our optimized JSMA algorithm	(and significantly faster than the un-optimized version)	

78.	our attacks	are typically 10 x –100x slower than previous attacks for $L_2$ or $L_\infty$	with the exception of iterative gradient sign which we are 10x slower	I feel this Focus would serve the authors better in a less focal position, like here:  Although iterative gradient sign is 10x slower, our attacks are typically 10 x –100x slower than previous attacks for $L_2$ or $L_\infty$ .
<b>commentary</b>				
<p>Notice how, as in section III, the verb is regularly New, but rarely Focus. The reason is this.</p> <p>the Process function is crucial to the clause. The Process is what is actually done or what actually happens. Really, the Process is why the clause is occurring in the first place, because otherwise all we'd have is a list of topics.</p> <p>But Processes can, to some extent, be anticipated, especially when we know the basic zone in which the Process is taking place. Well, that zone is normally found in the Theme; therefore, the newness of the Process is certainly there, but not as much there as is the newness of the outcome of the Process, and that outcome is found after the verb – it's the Participant or Circumstance now there because of the Process. And that is the shortest definition possible of the New: <b>there because of the Process.</b></p>				