

PokerSnowie vs. Pluribus? Let us all have a good laugh!

Nowadays the power of communication is everything, and it can make things look good and modern when they're actually not. Here at PokerSnowie we felt a bit embarrassed reading the blog posts and media coverage around the couple of events where the so-called 'Pluribus' (formerly Libratus) AI allegedly defeated pros at poker. The embarrassment was very high considering that such news even became an article in Nature, Science and on all the biggest media (no doubt Facebook knows how to advertise)!

The embarrassment was soon replaced by surprise. How could a research team such as Facebook-Carnegie Mellon go live with such bold and doubtful statements? "Pluribus is the first AI capable of beating human experts in six-player no-limit Hold'em" Really? :-)

Well, maybe the right statement would have been "Pluribus is the first AI developed by Facebook and Carnegie Mellon, capable of beating human experts...". But, there are several buts.

First of all we should mention that we have been trying to get in touch with Facebook and the researcher who posted their blog post at <https://ai.facebook.com/blog/pluribus-first-ai-to-beat-pros-in-6-player-poker/>. We have explained to them our concerns and opinions, we sent them a couple of documents explaining how inaccurate their communication was and the fact that they were not the first ones at all. We have also been open to a challenge. But, other than receiving a first email confirming that we were writing to the correct email address, we have been ignored. Well, this did not surprise us in fact.

Having said that, just the first statement "Pluribus is the first AI..." makes it clear that there has been no effort at all made by Facebook and Carnegie Mellon to check what is actually state of the art in poker AI, and they have just gone to the media with bold statements, claiming to be the best. We have been thinking "Why?". Probably because the undoubtedly huge investment needed to reach this result had to be justified publicly and this is most likely a big piece of advertising, considering that nowadays everyone loves to read about AI. Anyway, this is just our speculation.

What we also found quite amusing was the format used for these 'tests' or 'challenges'. As they say, "Pluribus achieves this result through several innovations on Libratus, the AI that beat human pros in two-player no-limit Hold'em in 2017". Hey, stop there! Libratus did this? How exactly? Ok, it seems 120,000 hands were played, and this is already quite a small sample to reach any conclusion, but look how the prize pool was defined, quoting "A prize pool of \$200,000 was allocated to the four humans in aggregate. Each human was guaranteed \$20,000 of that pool. The remaining \$120,000 was divided among them on the basis of how much better the human did against Libratus than the worst-performing of the four humans."

Is that real poker? Is it a real poker game where you are guaranteed \$20,000 and then if you are lucky enough you win more? We think this is a bit strange. We haven't been able to find out how the Pluribus challenge prize pool was actually organised, but we guess it was a similar approach.

A poker player who has nothing to lose is not going to play his best poker, that's for sure. We can't therefore conclude that Libratus is 'weak', but we can for sure state that the challenge was just for show and does not mean much. How was the Pluribus challenge organised? If anyone knows, let us know!

For benchmarking purposes, what is much more meaningful are the several botting teams that have been discovered online and caught by the poker rooms, after winning an aggregate of millions and millions of dollars. Besides the fact that botting is against the rules and it might even be considered illegal, and we at PokerSnowie are obviously not in favor of botting and we would not benchmark our AI in this way, we believe that it's still a very reliable way to test an AI. Obviously we are not talking about collusion or other even worse practices. Bots play real poker against real opponents

who have no idea if they are playing against an AI or not, and who actually have something to lose, not a \$20,000 prize just for joining a table (wouldn't that be nice?)!

So, it's clear and known by everyone that AIs capable of beating human pros (how would you otherwise categorise someone who plays online regularly \$3/\$6-\$25/\$25 and higher tables?) have been out there for several years. Just run some Google searches, visit the most known Poker-AI websites and the most popular forums, and you will realise how many botting teams have been winning.

Did you notice that so far we haven't mentioned anything about 'Pluribus vs. PokerSnowie'? Well, this was deliberate. We might be seen as defensive if we do so, but our main point here is that, regardless of PokerSnowie, the first statement of the above mentioned blog post, on which all the media coverage is based, is just NOT true.

Do you want to know a bit more about our opinion regarding the topic 'Pluribus vs. PokerSnowie'? Ok, here you go...

When PokerSnowie first came onto the market in 2013, we were criticised for our bold statements, where we emphasised the strength of our AI. While we do admit we may have exaggerated a little when talking for example about 'perfect play', we maintain that our AI is strong, and this is also confirmed by many poker coaches around the world, who are currently using our tool, regularly, both with their students and to improve their own game.

Having said that, we have read what has been written regarding Pluribus and it seems to us that the tone of certain statements is quite similar to the one we made in the past, especially a few sentences which are wrong in our opinion, and rather unjustified.

But instead of going back to them, we just decided to take a few paragraphs of what we read, and add a few comments for our users. The article we are referring to was published on <https://ai.facebook.com/blog/pluribus-first-ai-to-beat-pros-in-6-player-poker/>. Take a look at a few sentences (*italics*) and our comments (**bold**).

Pluribus is the first AI bot capable of beating human experts in six-player no-limit Hold'em, the most widely played poker format in the world. This is the first time an AI bot has beaten top human players in a complex game with more than two players or two teams.

Besides what we actually wrote above, we are not sure on what they can actually base this statement, considering that PokerSnowie achieved this several years ago, even in a 10-player game, which is even more complex than a six-player one. And once more, it's highly surprising that Facebook and Carnegie Mellon did not make any effort to check out other state of the art, commercially available products in a field where they state to be "the first".

We tested Pluribus against professional poker players, including two winners of the World Series of Poker Main Event. Pluribus won decisively.

Top players played extensively against PokerSnowie in various formats, and PokerSnowie won as well. As said above, even on a 10-seat game.

Pluribus succeeds because it can very efficiently handle the challenges of a game with both hidden information and more than two players. It uses self-play to teach itself how to win, with no examples or guidance on strategy.

PokerSnowie has also been developed without using any expert knowledge and it has just trained against itself.

Pluribus uses far fewer computing resources than the bots that have defeated humans in other games.

PokerSnowie AI server can provide any poker evaluation in less than 5ms and can run on very old desktop machines. Therefore this sentence is simply not true.

The bot's success will advance AI research, because many important AI challenges involve many players and hidden information.

As said previously, this has been achieved several years ago by PokerSnowie. Therefore we could say "Bit late to the party, guys".

For decades, poker has been a difficult and important grand challenge problem for the field of AI. Because poker involves hidden information — you don't know your opponents' cards — success requires bluffing and other strategies that do not apply to chess, Go, and other games. This twist has made poker resistant to AI techniques that produced breakthroughs in these other games.

It's exactly for this reason - 'involves hidden information' - that poker is especially fit for an AI approach and the PokerSnowie AI team developed a winning AI several years ago, even before coming out to the market. 'Bluffing' is a very natural outcome of the AI, as it's the only way the resulting strategy could be unexploitable. Therefore there is no need for any special skill by the developer in order to add 'bluffing' to the strategy – any decent Poker AI will come out with bluffing strategies.

This is the first time an AI bot has proven capable of defeating top professionals in any major benchmark game that has more than two players (or two teams).

Again we believe this sentence has been written without considering PokerSnowie or the poker market in the last 20 years at all. Also, whether a variance reduction approach is used or not, the number of hands needed to make such a claim is very high, higher than the 100,000 hands mentioned. PokerSnowie also carried out internal testing and played against top players, showing winning results. On top of that, PokerSnowie came to the market, allowing everyone to learn poker and make any benchmark against it.

PokerSnowie is available any time to run such challenges against any top player, also publicly. The truth is, so far, all our invitations have been rejected by famous players, and no public challenge has been possible. But we are open, ready to go.

Other comments have been written about the technology of Pluribus, but we don't feel there is anything to comment, except that they are mostly general and very high level description, and, we would add, often quite old, as many of the concepts described were already applied to the first AI our team developed in the late 90s, when Backgammon Snowie revolutionised the way people played backgammon.

All AI breakthroughs in previous benchmark games have been limited to those with only two players or two teams facing off in a zero-sum competition...

Again, an untrue statement. The breakthrough achieved with PokerSnowie several years ago, as mentioned, was not limited to 2 players. In fact PokerSnowie wins in multiway pot games, up to 10-player tables, not just the 6-seat games mentioned by Pluribus.

Those previous techniques could not scale to six-player poker even with 10,000x as much compute. Pluribus uses new techniques that can handle this challenge far better than anything that came before.

This is also completely not true; as previously mentioned, PokerSnowie achieved this a long time ago. Moreover, PokerSnowie's new AI, which at some point we will release, can be trained from scratch in a few days on a standard machine, and it would converge super quickly.

Most of the other points in the article are very general and theoretical aspects, which have been **common knowledge for several years** and we don't have anything special to add to them. Also, we would like to add once more a quick comment about the 'surprise' of seeing a computer bluffing: as already mentioned, any winning strategy in poker requires bluffing, and the AI easily learns it, otherwise it would be too easy to exploit. The same is true for mixed strategies, an easy concept to learn for an AI.

Finally, it's also important to keep in mind that PokerSnowie actually comes with a series of features which are designed specifically for being a coaching tool, while to our knowledge Pluribus is an R&D project. Also, the fact that PokerSnowie mixes strategies in a certain way and uses a limited number of bet sizes, for example, is a simplification that we made in order to make it easier to learn the strategy. We do have other NL AIs available where these constraints have not been set, but we believe they would not serve a good purpose for coaching.

Once more, if any top player wants to publicly challenge PokerSnowie, we are always available, but the experience so far is that no one wants to accept the challenge.

Some more information about PokerSnowie can be found here:

<https://www.pokersnowie.com/the-ai/technology-training.html>

It's worth noting that we have sent a copy of this blog post (the part where we compare point by point with Pluribus) to Facebook Research to draw their attention to the many inaccurate statements they have made, and their disregard of PokerSnowie.

As said above, we have so far been ignored. :)

Stay tuned and we'll keep you updated on any response we may receive.