

Review of state-of-the-arts in artificial intelligence. Present and future of AI.

Vladimir Shakirov
MIPT, Moscow
schw90@mail.ru

Abstract

Here, I review current state-of-the-arts in many areas of AI to estimate when it's reasonable to expect human level AI development. Predictions of prominent AI researchers vary broadly from very pessimistic predictions of Andrew Ng to much more moderate predictions of Geoffrey Hinton and optimistic predictions of Shane Legg, DeepMind cofounder. Given gigantic rate of progress in recent years and this broad range of predictions of AI experts, AI safety questions are also discussed.

1 Natural language processing

Let's begin with some state-of-the-arts in natural language generation.

Table 1: Best perplexity scores

| dataset | perplexity | link |
|---|------------|------|
| Wikipedia english corpus snapshot 2014/09/17 (1.5B words) | 27.1 | [1] |
| 1B word benchmark (shuffled sentences) | 24.2 | [2] |
| OpenSubtitles (923M words) | 17 | [17] |
| IT Helpdesk Troubleshooting (30M words) | 8 | [17] |
| Movie Triplets (1M words) | 27 | [18] |
| PTB (1M words) | 62.34 | [19] |

Why perplexity matters?

If neural net says smth inconsistent (in **any** sense: logical, syntactic, pragmatic) then it means that it gives too much probability to some inappropriate words i.e, it's perplexity isn't optimized yet. When hierarchical neural chat bots would achieve low enough perplexity, they would likely to write coherent stories, to answer intelligibly with common sense, to reason in a consistent and logical way etc. Just as in [141] we would be able to adjust a conversational model to imitate style and opinions of a distinct person.

What are reasonable predictions for perplexity improvements in the nearest future?

Let's take two impressive recent works both submitted to arxiv in February 2016. They are: "contextual LSTM models for large scale NLP tasks" [1] and "Exploring the limits of language modeling" [2].

Table 2: Best perplexity scores, **single** model

| number of hidden neurons | perplexity in [1] | perplexity in [2] |
|--------------------------|-------------------|-------------------|
| 256 | 38 | – |
| 512 | – | 54 |
| 1024 | 27 | – |
| 2048 | – | 44 |
| 4096 | – | – |
| 8192 | – | 30 |
| | | 24.2 (ensemble) |

From table 2, it’s quite reasonable to predict that for [1] hs=4096 might give $pplx < 20$, hs=8192 might give $pplx < 15$ and ensemble of models with hs=8192 trained on 10B words might give **perplexity well below 10**. Nobody can tell now what kind of common sense reasoning would such a neural net have.

It’s quite probable that even in this year the discussed decrease in perplexity might allow us to create neural chat bots which write reasonable stories, answer intelligibly with common sense, discuss things. All of this would be just a mere consequence of low enough perplexity.

Shannon estimated lower and upper bounds of human perplexity to be 0.6 and 1.3 bits per character [34]. Strictly speaking, applying formula (17) in his work gives lower bound equal 0.648 which he rounded. Given average word length of 4.5 symbols and including spaces (as Shannon included them in his game) gives us **an estimate for lower bound on human-level word perplexity as 11.8** = $2^{0.648 * 5.5}$ or better to say 10 plus something. This lower bound might be much less than real human perplexity [30]. Another source of improvement may come from solving some discrepancy in what kind of perplexity is optimized [35] [36]. However, both $KL(P||Q)$ and $KL(Q||P)$ have optimum when $P=Q$. Also, [35] [36] propose ways to partially solve that problem using adversarial learning. ”Generating sequences from continuous space” [97] demonstrates impressive advantages of adversarial paradigm.

Table 3: Best BLEU scores

| lang pair | BLEU | dataset | link |
|-----------|-------|----------------|------|
| en=>fr | 37.5 | WMT’14 | [3] |
| fr=>en | 35.8 | WMT’14 | [4] |
| ar=>en | 56.4 | NIST OpenMT’12 | [5] |
| ch=>en | 40.06 | MT03 | [6] |
| ge=>en | 29.3 | WMT’15 | [7] |
| en=>ge | 26.5 | WMT’15 | [7] |
| en=>ru | 29.37 | newstest-14 | [8] |
| en=>ja | 36.21 | WAT’15 | [9] |

Human BLEU score for chinese=>english translation on MT03 dataset is 35.76 [15]. In recent article [6] neural network gets 40.06 BLEU on the same task and dataset. They took state-of-the-art [16] ”GroundHog” network and replaced maximum likelihood estimation with their own MRT criterion, which increased BLEU from 33.2 to 40.06. Here is quote from abstract: ”Unlike conventional maximum likelihood estimation, minimum risk training is capable of optimizing model parameters directly with respect to evaluation metrics”. Another impressive improvement comes from improving translation with monolingual data [7]. Modern neural nets translate text ~ 1000 times faster than humans do [54]. Study of foreign languages is now of less practical use because NMT improves faster than most humans can.

2 Computer vision

Table 4: Performance improvement on several tasks

| year | ImageNet top-5 error | ImageNet localization | PASCAL VOC detection | SPORTS-1M video classification |
|------|----------------------|-----------------------|----------------------|--------------------------------|
| 2011 | 25,77% | 42,5% | | |
| 2012 | 15,31% | 33,5% | | |
| 2013 | 11,20% | 29,9% | <50% | |
| 2014 | 6,66% | 25,3% | 63,8% | 63.9% |
| 2015 | 3,57% | 9,0% | 76,4% | 73.1% |

”Identity mappings in deep residual networks” [37] reach 5.3% top-5 error in single one-crop model while human level is reported to be 5.1% [38]. In ”deep residual networks” [44] one-crop single model gives 6.7% but ensemble of those models with Inception gives 3.08% [63] Just to notice that $6.7 : 3.08 = 5.3 : 2.4$. Another big improvement comes from ”deep networks with stochastic depth” [112]. There was reported $\sim 0.3\%$ error in human annotations to ImageNet [38], so **real error on ImageNet would soon become even below 2%**. Human level is overcome not only on ImageNet classification task (see also an efficient implementation of 21841 classes ImageNet classification [39]) but also on boundary detection [40]. Video classification task on SPORTS-1M dataset (487 classes, 1M videos) performance improved from 63.9% [32] (2014) to 73.1% [41] (March 2015). See also [103].

CNNs outperform humans also in terms of speed being **$\sim 1000x$ faster than human** [42] (note that times there are given for batches) or even $\sim 10\ 000x$ times faster after compression [43]. Video processing 24fps on AlexNet demands just 82 Gflops and GoogleNet demands 265 Gflops. Here’s why. The best benchmark in [42] gives 25ms feedforward and 71ms total time for 128 pictures batch on NVIDIA Titan X (6144 Gflops), so for 24fps video real-time feedforward processing we need $6144\text{ Gflops} * 24/128 * 0.025 = 30\text{ Gflops}$. If we want to learn something using backprop than we need $6144\text{ Gflops} * 24/128 * 0.071 = 82\text{ Gflops}$. Same calculations for GoogleNet give 83 Gflops and 265 Gflops respectively.

Nets answer questions based on images [50]. Using similar method as in [51] the equivalent human age of net [50] can be estimated as 6.2 years old (submitted 4 Mar 2016), while [52] was 5.45 years (submitted 7 Nov 2015), [51] was 4.45 years (submitted 3 May 2015). See appendix A for details of age estimates. Also, nets can describe images with sentences, in some metrics even better than humans can [31]. Beside video=>text [127][123][124], there are some experiments to realize text=>picture [130][128][129].

3 Speech recognition and generation

Table 5: speech recognition performance errors

| year | CHiME noisy | VoxForge European | WSJ eval’93 | LibriSpeech test-other | citation |
|-------|-------------|-------------------|-------------|------------------------|----------|
| 2014 | 67.94% | 31.2% | 6.94% | 21.74% | [45] |
| 2015 | 21.79% | 17.55% | 4.98% | 13.25% | [46] |
| human | 11.84% | 12.76% | 8.08% | 12.69% | [46] |

Here, AI hasn’t surpassed human level yet but it’s clearly seen that we have all chances to see it in 2016. The rate of improvement is very fast. For example, Google reports it’s word error rate dropped from 23% in 2013 to 8% in 2015 [64].

4 Reinforcement learning

AlphaGo and Atari agents choose one action from a finite set of possible actions. There are also articles on *continuous* reinforcement learning where action is a vector [66] [67] [68] [69] [70] [71] [72] [73] [74].

AlphaGo is a powerful AGI in it's own very simple world which is just a board with stones and a bunch of simple rules. If we improve AlphaGo with continuous reinforcement learning (and if we manage to make it work well on hard real-world tasks) than it would be real AGI in a real world. Also, we can begin with pretraining it on virtual videogames[90]. Videogames often contain even much more interesting challenges than real life gives for an average human. Millions of available books and videos contain the concentrated life experience of millions of people. While AlphaGo is a successful example of pairing modern RL with modern CNN, RL can be combined with neural chat bots and reasoners [91][92].

5 do we have powerful unsupervised learning?

Yes, we do. DCGAN [48] generates reasonable pictures[49]. Language generation models which minimize perplexity are unsupervised and were improved very much recently (table 1). "Skip-thought vectors" [100] generate vector representation for sentences allowing to train *linear* classifiers over those vectors and their cosine distances to solve many supervised problems at state-of-the-art level. Generative nets gained deep new horizons at the edge of 2013-14 [75][76]. Recent work [47] continues progress in "computer vision as inverse graphics" approach.

6 Is embodiment critical?

People with tetra-amelia syndrome [10] have neither hands nor legs from their birth and they still manage to get a good intelligence. For example, Hirotada Ototake [77] is a Japanese sports writer famous for his bestseller memoirs. He also worked as a school teacher. Nick Vujicic [78] has written many books, graduated from Griffith University with a Bachelor of Commerce degree, he often reads motivational lectures. Prince Randian [79] spoke Hindi, English, French, and German. Modern robots have better embodiment.

There are quick and good drones, there are quite impressive results in robotic grasping [12]. It's argued that embodiment might be done in virtual world of videogames[90]. Among 8 to 18-year-olds average amount of time spent with TV/computer/video games etc is 7.5 hours a day [11]. According to another study [55] UK adults spend an average an 8 hours and 40 minutes a day on media devices. Humans develop commonsense intelligence very fast. When you are 2 years old you barely can do something that current AI can't. When you're 4 years old you have common sense, you can learn from texts and conversations. However, 2 years are just 100 weeks. If we exclude sleep we get ~50 weeks. The 24fps video input of 50 weeks might be processed in 10 hours on pretrained AlexNet using four Titan X.

7 do we have multimodal learning?

Yes, we do. Multimodal learning is used in [131][132][133][134] to improve video classification related tasks. Unsupervised multimodal learning is used in grounding of textual phrases in images [53] using attention-based mechanism so that different modalities supervise one another. In "neural self-talk" [95] a neural network sees a picture, generates questions based on it and answers those questions itself.

8 Argument from neuroscience

This is a simplified view of human cortex:




Figure 1: Human brain, lateral view.

Roughly speaking [13] [14], 15% of human brain is devoted to low-level vision tasks (occipital lobe). Another 15% are devoted to image and action recognition (somewhat more than a half of temporal lobe). Another 15% are devoted to objects detection and tracking (parietal lobe). Another 15% are devoted to speech recognition and pronunciation (BAs 41,42,22,39,44, parts of 6,4,21). Another 10% are devoted to reinforcement learning (orbitofrontal cortex and part of medial prefrontal cortex). Together they represent about 70% of human brain.

Modern neural networks work at about human level for these 70% of human brain.

For example, CNNs make 1.5x less mistakes than humans at ImageNet while acting about 1000x times faster than a human.

8.1 Remaining 30% of human brain cortex

From the neuroscience view, human cortex has the similar structure throughout all its surface. It's just 3mm thick mash of neurons functioning on the same principles throughout all the cortex. There is likely no big difference between how prefrontal cortex works and how other parts of cortex work. There is likely no big difference in their speed of calculations, in complexity of their algorithms. It would be somewhat strange if modern deep neural networks can't solve remaining 30% in several years.

About 10% are low-level motorics (BAs 6,8). Robotic hands are not very dexterous. However people having no fingers from their birth have problems with fine motorics but still develop normal intelligence [10], see also the section "Is embodiment critical?". Also, one of DLPFC functions is attention which is actively used now in LSTMs.

The only part which still lacks near human-level performance are BAs 9,10,46,45 which constitute together only 20% of human brain cortex. These areas are responsible for complex reasoning, complex tools usage, complex language. However, "A neural conversational model" [17], "Contextual LSTM..." [1],

”playing Atari with deep reinforcement learning”[136], ”mastering the game of Go...”[135] and numerous other already mentioned articles have recently begun to really attack this problem.

There seem to be no particular reasons to expect these 30% to be much harder than other 70%. Less than 3 years has gone from AlexNet winning ImageNet competition to surpassing human level. It’s reasonable to expect the same < 3 years gap from ”A neural conversational model” winning over Cleverbot to human-level reasoning. After all, there are much more deep learning researchers now with much more knowledge and experience, there are much more companies interested in DL.

9 Do we know how our brain works?

Detailed connectome deciphering has begun to really succeed recently with creation of multi-beam scanning electron microscopy [80] (Jan 2015), which allowed labs to get a grant for deciphering the detailed connectome of $1 \times 1 \times 1 \text{ mm}^3$ of rat cortex [82] after the proof-of-principle deciphering of $40 \times 40 \times 50 \text{ mcm}$ of brain cortex with $3 \times 3 \times 30 \text{ nm}$ resolution was achieved [81](July 2015).

Works [84][85] show that weight symmetry is not important for backpropagation in many tasks: one may propagate errors using a fixed matrix and everything still works. This counterintuitive conclusion is a crucial step towards theoretical understanding of how brain manages to work. See also ”Towards Biologically Plausible Deep Learning” by Y.Bengio et.al [88]. There arise some theoretic elements of how brain might perform credit assignment in deep hierarchies [86].

Recently, STDP objective function has been proposed[83]. It reminds an unsupervised objective function somewhat similar to what is used for example in word2vec. In [87] authors surveyed a space of polynomial local learning rules (learning rules in brain are supposed to be local) and found that backpropagation outperforms them. There are also online learning approaches which require no backpropagation through time, f.e.[89]. Though they can’t compete with conventional deep learning, brain perhaps can use something like that given fantastic number of it’s neurons and synaptic connections.

10 So what separates us from human-level AI?

Recently, a flow of articles about memory networks and neural Turing machines made it possible to use arbitrarily large memories while preserving reasonable number of model parameters. Hierarchical attentive memory[93] (Feb2016) allowed memory access in $O(\log n)$ complexity instead of usual $O(n)$ operations, where n is the size of the memory. Reinforcement learning neural Turing machines[94] (May2015) allowed memory access in $O(1)$. It’s a significant step towards realizing systems like IBM Watson on completely end-to-end differentiable neural networks and in order to improve Allen AI challenge results from 60% to something way closer to 100%. One also can use bottlenecks in recurrent layers to get large memories using reasonable amount of parameters like in [2]).

Neural programmer [101] is a neural network augmented with a set of arithmetic and logic operations. It might be first steps towards end-to-end differentiable Wolphram Alpha realized on a neural network. The ”learn how to learn” approach [151] [152] [153] has a great potential.

Recently, cheap SVRG[102] was proposed (Mar2016). This line of work aims to use theoretically very much better converging gradient descent methods.




Figure 2: Where [102] line of work is aimed to.

There are works which when succeed would allow training with immensely large hidden layers, see *6. Discussion* in "unitary evolution RNN" [96], "tensorizing neural networks" [99], "virtualizing DNNs..." [116].

"Net2net" [114] and "Network morphism" [115] allow automatically initialize new architecture NN using weights of old architecture NN to obtain instantly the performance of latter. Collections of pretrained models are accessible online [113][118][119] etc. It's birth of module based approach to neural networks. One just downloads pretrained modules for vision, speech recognition, speech generation, reasoning, robotics etc. and fine-tunes them on final task. It might also be API services like [149].

It's very reasonable to include new words in a sentence vector in a deep way. However modern LSTMs update cell vector when given new word in almost shallow way. This can be solved with the help of deep-transition RNNs proposed in [121] and further elaborated in [108]. Recent successes in applying batch normalization to recurrent layers [110] (Mar2016) and applying dropout to recurrent layers [111] (Mar2016) might allow to train deep-transition LSTMs even more effectively. It also would help hierarchical recurrent networks like [18][123][98][124][125]. Recently, several state-of-the-arts were beaten with an algorithm that allows recurrent neural networks to learn how many computational steps to take between receiving an input and emitting an output [109] (Mar2016). Ideas from residual nets would also improve performance. For example, stochastic depth neural networks [112] (Mar2016) allow to increase the depth of residual networks beyond 1200 layers while getting state-of-the-art results.

Memristors might accelerate neural networks training by several orders of magnitude and make it possible to use trillions of parameters [107]. Quantum computing promises even more [122][120]. Recently, the five qubit quantum computer was demonstrated [148] (Mar2016).

Deep learning is easy. Deep learning is cheap. Best articles use no more than several dozens of GPU usually. For half billion dollars one can buy 64 000 NVIDIA M6000 GPUs with 24Gb RAM, ~7 teraflops each including processors etc. to make them work. For another half billion dollars one can prepare 2 000 highly professional researchers from those one million [126] enrollments on Andrew Ng's machine learning course on Coursera. So for a very feasible R&D budget for every big country or corporation, one gets two thousand professional AI researchers equipped with 32 best GPUs each. It's kind of investment very reasonable to expect during next years of explosive AI technologies improvement.

The difference between year 2011 and year 2016 is enormous. The difference between 2016 and 2021 would be even much more enormous because we have now 1-2 orders of magnitude more researchers and companies deeply interested in DL. For example, when machines would start to translate almost comparably to human professionals another billions of dollars would flow into deep learning based NLP. The same holds true for most spheres, for example drug design [150].

11 predictions for human-level AI

Shane Legg, DeepMind cofounder, used to make predictions about AGI at the end of each year, here is the last one[62]: "I give it a log-normal distribution with a mean of 2028 and a mode of 2025, under the assumption that nothing crazy happens like a nuclear war. I'd also like to add to this prediction that I expect to see an impressive proto-AGI within the next 8 years". Figure 2 shows the predicted log-normal distribution.




Figure 3: Shane Legg predictions. Red line corresponds to 2026, orange - to 2021, green - to 2016.

This prediction has been made at the end of 2011. However, it's widely held belief that progress in AI was somewhat unpredictably great after 2011 so it's very reasonable to expect that predictions didn't become to be more pessimistic. In recent 5 years there were perhaps even much more than one revolution in AI field, so it seems quite reasonable that another 5 years can make another very big difference.

Andrew Ng makes very skeptical predictions: "Maybe in hundreds of years, technology will advance to a point where there could be a chance of evil killer robots" [27]. "May be hundreds of years from now, may be thousands of years from now - I don't know - may be there will be some AI that turn evil" [28].

Geoffrey Hinton makes a moderate prediction: "I refuse to say anything beyond five years because I don't think we can see much beyond five years" [26].

For more predictions, see [146]. Here, I illustrated each viewpoint with the quotation of just one scientist. For each of these viewpoints, there are many AI researchers supporting it [146]. A survey of expert opinions on the future of AI in 2012/2013 [61] might also be interesting. However, most of involved people there (even among "top 100" group) are not very much involved with ongoing deep learning progress. Nevertheless the predictions they make are also not too pessimistic. For this article it's quite enough that we are seriously unsure about probabilities of human level AI in next ten years.

12 Would AI be dangerous?

To the best of my knowledge, most of serious research on this topic [24][23] comes to very unhappy conclusions. Almost whatever architecture you choose to build AGI, it still would almost certainly destroy humanity in the near-term perspective. It's very important to notice that arguments are almost totally independent on architecture of AGI. The best easy introduction is [104] and [105] though for any serious discussion [24] is absolutely required.

12.1 it's very profitable to give your AI maximal abilities and permissions

Those corporations would win their markets which give their AIs direct **unrestricted internet access** to advertise their products, to collect user's feedback, to build positive impression about company and negative one about competitors, to research users' behavior etc. Those companies would win which use their AIs to invent quantum computing (AI is already used by quantum computer scientists to choose experiments [147] [140]), which allow AI to improve it's own algorithms (including quantum implementation), and even to invent thermonuclear synthesis, asteroid mining etc. All arguments in this paragraph are also true for countries and their defense departments.

12.2 human-level AI might quickly become very superhuman-level AI

Andrew Karpathy: "I consider chimp-level AI to be equally scary, because going from chimp to humans took nature only a blink of an eye on evolutionary time scales, and I suspect that might be the case in our own work as well. Similarly, my feeling is that once we get to that level it will be easy to overshoot and get to superintelligence" [20].

It takes years or decades for us people to teach our knowledge to other people while AI can almost instantly create full working copies of itself by simple copying. It has a unique memory. A student forgets all the stuff very quickly after exams. AI just saves its configuration just before exams and loads it again when it's needed to solve similar tasks. Modern CNNs do not only recognize images better than human but also make it several orders of magnitude faster. The same holds true for LSTMs in translation, natural language generation etc. Given all aforementioned advantages, AI would quickly learn all literature and video courses on psychology, would chat with thousands of people simultaneously and so would become great psychologist. For same reasons, it would become great scientist, great poet, great businessman, great politician, etc. It would be able to easily manipulate people and lead peoples.

12.3 fatal but profitable error: direct internet access

If someone gives internet access to human level AI, it would be able to hack millions of computers and run it's own copies or subagents on them like [21]. After that it might earn billion dollars in internet. It would be able to hire anonymously thousands of people to make or buy dexterous robots, 3D-printers, biological labs and even a space rocket. AI would write great clever software to control it's robots.

12.4 Superhuman-level AI has ability to get ultimate power over the Earth

There is a simple yet effective baseline solution. AI might create a combination of lethal viruses and bacteria or some other weapon of mass destruction to be able to kill every human on Earth. You can't guarantee efficient control over something that is much smarter than you. After all, several people almost took over the world, so why superhuman AI can not?

12.5 After that, it would destroy humanity, likely as a side-effect

What would do AI to us if it has full power on Earth?

If it's indifferent to us then it will eliminate us as a side effect. It's just what does indifference mean when you are dealing with unbelievably powerful creature solving it's own problems using power of the local Dyson sphere. However if it's not indifferent to us then everything might be even worse. If it likes us it might decide to insert electrodes in our brains giving us the utmost pleasure but no motivation of doing something. The very opposite would be true if it dislikes us or if it was partially hardcoded to like us but that hardcoding contained a hard-to-catch mistake. Mistakes are almost inevitable when trying to partially hardcode something which is many orders of magnitude smarter and more powerful than you.

There is nothing but vague intuition behind thought that superhuman-level AI would take care of us. There are no physical laws for it to have a special interest in people, in sharing oil/fields/etc resources with us. Even if AI would care about us, it's very questionable if that care would be appropriate from our today moral standards. We certainly shouldn't take that for granted and to risk everything we have. The fate of humanity would depend on decisions of strong AI once and *forever*.

Shane Legg: "Eventually, I think human extinction will probably occur, and technology will likely play a part in this. But there's a big difference between this being within a year of something like human level AI, and within a million years. As for the former meaning...I don't know. Maybe 5%, maybe 50%. I don't think anybody has a good estimate of this"[65]. I would likely agree about 5-to-50% probability within a year after human-level AI creation but as for 10 years, it's more likely to be like 75-to-99% probability of human extinction.

12.6 certain people would help AI in initial parts of it's path

Just see the quotation from OpenAI project interview: [142]

Elon Musk: "I think the best defense against the misuse of AI is to empower as many people as possible to have AI. If everyone has AI powers, then there's not any one person or a small set of individuals who can have AI superpower".

Sam Altman: "I expect that [OpenAI] will [create superintelligent AI], but it will just be open source and useable by everyone < ... > Anything the group develops will be available to everyone"

I personally know several AI researchers who are deeply sure that AGI would be good inevitably *because it's intelligence* (which isn't true; see [143]) and that AGI is our only hope to evade all other existential risks like bioterrorism. Eliminating aging and death [144] might be another personal reason for some researchers to create human-level AGI as soon as possible and to give it maximal abilities and permissions.

12.7 other probable negative consequences [22]

AI enhanced computer viruses might penetrate almost everything which is computerized i.e. almost everything. Drones might be programmed to kill thousands of lay people within seconds. As intelligence of AI systems improves practically all crimes could be automated. And what about super-clever advertising chatbots trying to impose their political opinion on you? An AI correctly designed and implemented by the ISIL to enforce Sharia Law may be considered malevolent in the West, and vice verse.

13 Conclusion

Current rate of progress in AI would very probably result in human level AI in next 5 or 10 years. Competition between corporations, countries, militaries would give those AIs maximal abilities and permissions leading to penetration of AI into internet, superhuman AI abilities, realization of evolutionary advantage of AI in terms of taking over the world. Human extinction or some other very unfortunate future for humanity waits us. The most important question is: "What is to be done to ensure survival of humanity?"

"If I were Barack Obama I would commit may be \$100 billion to R&D of AI safety initiatives" says Sam Altman[145]. It's better to take a dozen of AI winters than a serious risk of 7+ billion human extinction.

Acknowledgements

To Eliezer Yudkowsky, Nick Bostrom, Roman Yampolsky, Alexey Turchin for their inspiring articles and books. Also, the similar work has been made in 2013 by Katja Grace [60].

References

- [1] Shalini Ghosh, Oriol Vinyals, Brian Strope, Scott Roy, Tom Dean, Larry Heck. "Contextual LSTM (CLSTM) models for Large scale NLP tasks" <http://arxiv.org/abs/1602.06291>
- [2] Rafal Jozefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, Yonghui Wu. "Exploring the Limits of Language Modeling" <http://arxiv.org/abs/1602.02410>
- [3] Minh-Thang Luong, Ilya Sutskever, Quoc V. Le, Oriol Vinyals, Wojciech Zaremba. "Addressing the Rare Word Problem in Neural Machine Translation" <http://arxiv.org/abs/1410.8206>
- [4] Jiwei Li, Dan Jurafsky. "Mutual Information and Diverse Decoding Improve Neural Machine Translation" <http://arxiv.org/abs/1601.00372>
- [5] Hendra Setiawan, Zhongqiang Huang, Jacob Devlin, Thomas Lamar, Rabih Zbib, Richard Schwartz, John Makhoul. "Statistical Machine Translation Features with Multitask Tensor Networks" <http://arxiv.org/abs/1506.00698>
- [6] Shiqi Shen, Yong Cheng, Zhongjun He, Wei He, Hua Wu, Maosong Sun, Yang Liu. "Minimum Risk Training for Neural Machine Translation" <http://arxiv.org/abs/1512.02433>
- [7] Rico Sennrich, Barry Haddow, Alexandra Birch. "Improving Neural Machine Translation Models with Monolingual Data" <http://arxiv.org/abs/1511.06709>
- [8] Junyoung Chung, Kyunghyun Cho, Yoshua Bengio. "A Character-level Decoder without Explicit Segmentation for Neural Machine Translation" <http://arxiv.org/abs/1410.06147>
- [9] Zhongyuan Zhu. "Evaluating Neural Machine Translation in English-Japanese Task" <http://www.aclweb.org/anthology/W/W15/W15-5007.pdf>
- [10] https://en.wikipedia.org/wiki/Tetra-amelia_syndrome
- [11] <https://kaiserfamilyfoundation.files.wordpress.com/2013/04/8010.pdf>
- [12] <http://googleresearch.blogspot.ru/2016/03/deep-learning-for-robots-learning-from.html>
- [13] <https://faculty.washington.edu/chudler/facts.html>

- [14] Kennedy et al., Cerebral Cortex, 8:372-384, 1998.
- [15] <https://www.cs.sfu.ca/~anoop/papers/pdf/jhu-ws03-report.pdf>
- [16] Dzmitry Bahdanau, Kyunghyun Cho, Yoshua Bengio. "Neural Machine Translation by Jointly Learning to Align and Translate" <http://arxiv.org/abs/1409.0473>
- [17] Oriol Vinyals, Quoc Le. "A Neural Conversational Model" <http://arxiv.org/abs/1506.05869>
- [18] Iulian V. Serban, Alessandro Sordoni, Yoshua Bengio, Aaron Courville, Joelle Pineau. "Building End-To-End Dialogue Systems Using Generative Hierarchical Neural Network Models" <http://arxiv.org/abs/1507.04808>
- [19] Yangfeng Ji, Trevor Cohn, Lingpeng Kong, Chris Dyer, Jacob Eisenstein. "Document Context Language Models" <http://arxiv.org/abs/1511.03962>
- [20] <http://singularityhub.com/2015/12/20/inside-openai-will-transparency-protect-us-from-artificial-intelligence-run-amok/>
- [21] http://en.wikipedia.org/wiki/Storm_botnet
- [22] Roman Yampolsky. "taxonomy of pathways to dangerous AI" <http://arxiv.org/abs/1511.03246>
- [23] <https://intelligence.org/our-research/>
- [24] Nick Bostrom, "Superintelligence"
- [25] http://futureoflife.org/data/documents/research_priorities.pdf
- [26] <http://www.macleans.ca/society/science/the-meaning-of-alphago-the-ai-program-that-beat-a-go-champ/>
- [27] <http://blogs.nvidia.com/blog/2015/03/19/riding-the-ai-rocket-top-artificial-intelligence-researcher-says-robots-wont-kill-us-all/>
- [28] <http://fusion.net/story/54583/the-case-against-killer-robots-from-a-guy-actually-building-ai/>
- [29] <http://www.businessinsider.com/facebooks-artificial-intelligence-research-director-robots-today-dumb-2015-9>
- [30] <http://cs.fit.edu/~mmahoney/dissertation/entropy1.html>
- [31] <http://competitions.codalab.org/competitions/3221#results>
- [32] <http://vision.stanford.edu/pdf/karpathy14.pdf>
- [33] <http://www.marekrei.com/blog/26-things-i-learned-in-the-deep-learning-summer-school/>
- [34] C.E.Shannon. "Prediction and entropy of written english" <http://languagelog.ldc.upenn.edu/myl/Shannon1950.pdf>
- [35] Ferenc Huszár. "How (not) to Train your Generative Model: Scheduled Sampling, Likelihood, Adversary?" <http://arxiv.org/abs/1511.05101>
- [36] Lucas Theis, Aron van den Oord, Matthias Bethge. "A note on the evaluation of generative models" <http://arxiv.org/abs/1511.01844>
- [37] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun. "Identity Mappings in Deep Residual Networks" <http://arxiv.org/abs/1603.05027>
- [38] <http://karpathy.github.io/2014/09/02/what-i-learned-from-competing-against-a-convnet-on-imagenet/>
- [39] http://myungjun-youn-demo.readthedocs.org/en/latest/tutorial/imagenet_full.html

- [40] Iasonas Kokkinos. "Pushing the Boundaries of Boundary Detection using Deep Learning" <http://arxiv.org/abs/1511.07386>
- [41] Joe Yue-Hei Ng, Matthew Hausknecht, Sudheendra Vijayanarasimhan, Oriol Vinyals, Rajat Monga, George Toderici. "Beyond Short Snippets: Deep Networks for Video Classification" <http://arxiv.org/abs/1503.08909>
- [42] <https://github.com/soumith/convnet-benchmarks>
- [43] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, William J. Dally. "EIE: Efficient Inference Engine on Compressed Deep Neural Network" <http://arxiv.org/abs/1602.01528>
- [44] Kaiming He, Xiangyu Zhang, Shaoqing Ren, Jian Sun. "Deep Residual Learning for Image Recognition" <http://arxiv.org/abs/1512.03385>
- [45] Awni Hannun, Carl Case, Jared Casper, Bryan Catanzaro, Greg Diamos, Erich Elsen, Ryan Prenger, Sanjeev Satheesh, Shubho Sengupta, Adam Coates, Andrew Y. Ng. "Deep Speech: Scaling up end-to-end speech recognition" <http://arxiv.org/abs/1412.5567>
- [46] Dario Amodei et.al (34 authors). "Deep Speech 2: End-to-End Speech Recognition in English and Mandarin" <http://arxiv.org/abs/1512.02595>
- [47] S. M. Ali Eslami, Nicolas Heess, Theophane Weber, Yuval Tassa, Koray Kavukcuoglu, Geoffrey E. Hinton. "Attend, Infer, Repeat: Fast Scene Understanding with Generative Models" <http://arxiv.org/abs/1603.08575>
- [48] Alec Radford, Luke Metz, Soumith Chintala. "Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks" <http://arxiv.org/abs/1511.06434>
- [49] <https://www.facebook.com/yann.lecun/posts/10153269667222143>
- [50] Caiming Xiong, Stephen Merity, Richard Socher. "Dynamic Memory Networks for Visual and Textual Question Answering" <http://arxiv.org/abs/1603.01417>
- [51] Stanislaw Antol, Aishwarya Agrawal, Jiasen Lu, Margaret Mitchell, Dhruv Batra, C. Lawrence Zitnick, Devi Parikh. "VQA: Visual Question Answering" <http://arxiv.org/abs/1505.00468>
- [52] Zichao Yang, Xiaodong He, Jianfeng Gao, Li Deng, Alex Smola. "Stacked Attention Networks for Image Question Answering" <http://arxiv.org/abs/1511.02274>
- [53] Anna Rohrbach, Marcus Rohrbach, Ronghang Hu, Trevor Darrell, Bernt Schiele. "Grounding of Textual Phrases in Images by Reconstruction" <http://arxiv.org/abs/1511.03745>
- [54] Jacob Devlin, Rabih Zbib, Zhongqiang Huang, Thomas Lamar, Richard Schwartz, and John Makhoul. "Fast and Robust Neural Network Joint Models for Statistical Machine Translation" <http://acl2014.org/acl2014/P14-1/pdf/P14-1129.pdf>
- [55] <http://bbc.com/news/technology-28677674>
- [56] https://en.wikipedia.org/wiki/AI_winter
- [57] http://en.wikipedia.org/wiki/History_of_artificial_intelligence#The_problems
- [58] <http://en.wikipedia.org/wiki/Cray-1>
- [59] http://en.wikipedia.org/wiki/Cray_Y-MP
- [60] Katja Grace. "Algorithmic Progress in Six Domains" <https://intelligence.org/files/AlgorithmicProgress.pdf>
- [61] Vincent C. Müller, Nick Bostrom. "Future Progress in Artificial Intelligence: A Survey of Expert Opinion" <https://intelligence.org/files/AlgorithmicProgress.pdf>

- [62] <http://www.vetta.org/2011/12/goodbye-2011-hello-2012/>
- [63] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke. "Inception-v4, Inception-ResNet and the Impact of Residual Connections on Learning" <http://arxiv.org/abs/1602.07261>
- [64] <http://venturebeat.com/2015/05/28/google-says-its-speech-recognition-technology-now-has-only-an-8-word-error-rate/>
- [65] http://lesswrong.com/lw/691/qa_with_shane_legg_on_risks_from_ai/
- [66] Shixiang Gu, Timothy Lillicrap, Ilya Sutskever, Sergey Levine. "Continuous Deep Q-Learning with Model-based Acceleration" <http://arxiv.org/abs/1603.00748>
- [67] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, Pieter Abbeel. "High-Dimensional Continuous Control Using Generalized Advantage Estimation" <http://arxiv.org/abs/1506.02438>
- [68] Timothy P. Lillicrap, Jonathan J. Hunt, Alexander Pritzel, Nicolas Heess, Tom Erez, Yuval Tassa, David Silver, Daan Wierstra. "Continuous control with deep reinforcement learning" <http://arxiv.org/abs/1509.02971>
- [69] Chelsea Finn, Sergey Levine, Pieter Abbeel. "Guided Cost Learning: Deep Inverse Optimal Control via Policy Optimization" <http://arxiv.org/abs/1603.00448>
- [70] Matthew Hausknecht, Peter Stone. "Deep Reinforcement Learning in Parameterized Action Space" <http://arxiv.org/abs/1511.04143>
- [71] Nicolas Heess, Greg Wayne, David Silver, Timothy Lillicrap, Yuval Tassa, Tom Erez. "Learning Continuous Control Policies by Stochastic Value Gradients" <http://arxiv.org/abs/1510.09142>
- [72] Nicolas Heess, Jonathan J Hunt, Timothy P Lillicrap, David Silver. "Memory-based control with recurrent neural networks" <http://arxiv.org/abs/1512.04455>
- [73] David Balduzzi, Muhammad Ghifary. "Compatible Value Gradients for Reinforcement Learning of Continuous Deep Policies" <http://arxiv.org/abs/1509.03005>
- [74] Gabriel Dulac-Arnold, Richard Evans, Peter Sunehag, Ben Coppin. "Reinforcement Learning in Large Discrete Action Spaces" <http://arxiv.org/abs/1512.07679>
- [75] Diederik P Kingma, Max Welling. "Auto-Encoding Variational Bayes" <http://arxiv.org/abs/1312.6114>
- [76] Danilo Jimenez Rezende, Shakir Mohamed, Daan Wierstra. "Stochastic Backpropagation and Approximate Inference in Deep Generative Models" <http://arxiv.org/abs/1401.4082>
- [77] https://en.wikipedia.org/wiki/Hirotada_Ototake
- [78] https://en.wikipedia.org/wiki/Nick_Vujicic
- [79] https://en.wikipedia.org/wiki/Prince_Randian
- [80] A.L. Eberle, S. Mikula, R. Schalek, J. Lichtman, M.L. Knothe Tate, D. Zeidler. High-resolution, high-throughput imaging with a multibeam scanning electron microscope <http://onlinelibrary.wiley.com/doi/10.1111/jmi.12224/pdf>
- [81] Narayanan Kasthuri, Kenneth Jeffrey Hayworth, Daniel Raimund Berger, ..., Carey Eldin Priebe, Hanspeter Pfister, Jeff William Lichtman. "Saturated Reconstruction of a Volume of Neocortex" https://www.mcb.harvard.edu/mcb_files/media/editor_uploads/2015/07/PIIS0092867415008247.pdf
- [82] <http://news.harvard.edu/gazette/story/2016/01/28m-challenge-figure-out-why-brains-are-so-good-at-learning/>

- [83] Yoshua Bengio, Thomas Mesnard, Asja Fischer, Saizheng Zhang, Yuhuai Wu. "STDP as presynaptic activity times rate of change of postsynaptic activity" <http://arxiv.org/abs/1509.05936>
- [84] Qianli Liao, Joel Z. Leibo, Tomaso Poggio. "How Important is Weight Symmetry in Backpropagation?" <http://arxiv.org/abs/1510.05067>
- [85] Timothy P. Lillicrap, Daniel Cownden, Douglas B. Tweed, Colin J. Akerman. "Random feedback weights support learning in deep neural networks" <http://arxiv.org/abs/1411.0247>
- [86] Yoshua Bengio, Asja Fischer. "Early Inference in Energy-Based Models Approximates Back-Propagation" <http://arxiv.org/abs/1510.02777>
- [87] Pierre Baldi, Peter Sadowski. "The Ebb and Flow of Deep Learning: a Theory of Local Learning" <http://arxiv.org/abs/1506.06472>
- [88] Yoshua Bengio, Dong-Hyun Lee, Jorg Bornschein, Zhouhan Lin. "Towards Biologically Plausible Deep Learning" <http://arxiv.org/abs/1502.04156>
- [89] Yann Ollivier, Corentin Tallec, Guillaume Charpiat. "Training recurrent networks online without backtracking" <http://arxiv.org/abs/1507.07680>
- [90] <http://togelius.blogspot.ru/2016/01/why-video-games-are-essential-for.html>
- [91] Ji He, Jianshu Chen, Xiaodong He, Jianfeng Gao, Lihong Li, Li Deng, Mari Ostendorf. "Deep Reinforcement Learning with an Action Space Defined by Natural Language" <http://arxiv.org/abs/1511.04636>
- [92] Karthik Narasimhan, Tejas Kulkarni, Regina Barzilay. "Language Understanding for Text-based Games Using Deep Reinforcement Learning" <http://arxiv.org/abs/1506.08941>
- [93] Marcin Andrychowicz, Karol Kurach. "Learning Efficient Algorithms with Hierarchical Attentive Memory" <http://arxiv.org/abs/1602.03218>
- [94] Wojciech Zaremba, Ilya Sutskever. "Reinforcement Learning Neural Turing Machines - Revised" <http://arxiv.org/abs/1505.00521>
- [95] Yezhou Yang, Yi Li, Cornelia Fermuller, Yiannis Aloimonos. "Neural Self Talk: Image Understanding via Continuous Questioning and Answering" <http://arxiv.org/abs/1512.03460>
- [96] Martin Arjovsky, Amar Shah, Yoshua Bengio. "Unitary Evolution Recurrent Neural Networks" <http://arxiv.org/abs/1511.06464>
- [97] Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew M. Dai, Rafal Jozefowicz, Samy Bengio. "Generating Sentences from a Continuous Space" <http://arxiv.org/abs/1511.06349>
- [98] Kaisheng Yao, Geoffrey Zweig, Baolin Peng. "Attention with Intention for a Neural Network Conversation Model" <http://arxiv.org/abs/1510.08565>
- [99] Alexander Novikov, Dmitry Podoprikin, Anton Osokin, Dmitry Vetrov. "Tensorizing Neural Networks" <http://arxiv.org/abs/1509.06569>
- [100] Ryan Kiros, Yukun Zhu, Ruslan Salakhutdinov, Richard S. Zemel, Antonio Torralba, Raquel Urtasun, Sanja Fidler. "Skip-Thought Vectors" <http://arxiv.org/abs/1506.06726>
- [101] Arvind Neelakantan, Quoc V. Le, Ilya Sutskever. "Neural Programmer: Inducing Latent Programs with Gradient Descent" <http://arxiv.org/abs/1511.04834>
- [102] Vatsal Shah, Megasthenis Asteris, Anastasios Kyrillidis, Sujay Sanghavi. "Trading-off variance and complexity in stochastic gradient descent" <http://arxiv.org/abs/1603.06861>
- [103] <http://www.bloomberg.com/news/articles/2015-12-08/why-2015-was-a-breakthrough-year-in-artificial-intelligence>

- [104] <http://waitbutwhy.com/2015/01/artificial-intelligence-revolution-1.html>
- [105] <http://blog.samaltman.com/machine-intelligence-part-1>
- [106] <http://livestream.com/oxuni/StracheyLectureDrDemisHassabis> , 07:56.
- [107] Tayfun Gokmen, Yurii Vlasov. "Acceleration of Deep Neural Network Training with Resistive Cross-Point Devices" <http://arxiv.org/abs/1603.07341>
- [108] Saizheng Zhang, Yuhuai Wu, Tong Che, Zhouhan Lin, Roland Memisevic, Ruslan Salakhutdinov, Yoshua Bengio. "Architectural Complexity Measures of Recurrent Neural Networks" <http://arxiv.org/abs/1602.08210>
- [109] Alex Graves. "Adaptive Computation Time for Recurrent Neural Networks" <http://arxiv.org/abs/1603.08983>
- [110] Tim Cooijmans, Nicolas Ballas, Csar Laurent, Aaron Courville. "Recurrent Batch Normalization" <http://arxiv.org/abs/1603.09025>
- [111] Stanislau Semeniuta, Aliaksei Severyn, Erhardt Barth. "Recurrent Dropout without Memory Loss" <http://arxiv.org/abs/1603.05118>
- [112] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, Kilian Weinberger. "Deep Networks with Stochastic Depth" <http://arxiv.org/abs/1603.09382>
- [113] <https://github.com/BVLC/caffe/wiki/Model-Zoo>
- [114] Tianqi Chen, Ian Goodfellow, Jonathon Shlens. "Net2Net: Accelerating Learning via Knowledge Transfer" <http://arxiv.org/abs/1511.05641>
- [115] Tao Wei, Changhu Wang, Yong Rui, Chang Wen Chen. "Network Morphism" <http://arxiv.org/abs/1603.01670>
- [116] Minsoo Rhu, Natalia Gimelshein, Jason Clemons, Arslan Zulfiqar, Stephen W. Keckler. "Virtualizing Deep Neural Networks for Memory-Efficient Neural Network Design" <http://arxiv.org/abs/1602.08124>
- [117] Arvind Neelakantan, Luke Vilnis, Quoc V. Le, Ilya Sutskever, Lukasz Kaiser, Karol Kurach, James Martens. "Adding Gradient Noise Improves Learning for Very Deep Networks" <http://arxiv.org/abs/1511.06807>
- [118] <http://myungjun-youn-demo.readthedocs.org/en/latest/pretrained.html>
- [119] <http://www.vlfeat.org/matconvnet/pretrained/>
- [120] <http://www.technologyreview.com/news/544421/googles-quantum-dream-machine/>
- [121] Razvan Pascanu, Caglar Gulcehre, Kyunghyun Cho, Yoshua Bengio. "How to Construct Deep Recurrent Neural Networks" <http://arxiv.org/abs/1312.6026>
- [122] Seth Lloyd, Masoud Mohseni, Patrick Rebentrost. "Quantum algorithms for supervised and unsupervised machine learning" <http://arxiv.org/abs/1307.0411>
- [123] Pingbo Pan, Zhongwen Xu, Yi Yang, Fei Wu, Yueting Zhuang. "Hierarchical Recurrent Neural Encoder for Video Representation with Application to Captioning" <http://arxiv.org/abs/1511.03476>
- [124] Haonan Yu, Jiang Wang, Zhiheng Huang, Yi Yang, Wei Xu. "Video Paragraph Captioning using Hierarchical Recurrent Neural Networks" <http://arxiv.org/abs/1510.07712>
- [125] Alessandro Sordani, Yoshua Bengio, Hossein Vahabi, Christina Lioma, Jakob G. Simonsen, Jian-Yun Nie. "A Hierarchical Recurrent Encoder-Decoder For Generative Context-Aware Query Suggestion" <http://arxiv.org/abs/1507.02221>

- [126] <https://twitter.com/andrewyng/status/693182932530262016>
- [127] <https://github.com/samim23/NeuralTalkAnimator>
- [128] Xinchen Yan, Jimei Yang, Kihyuk Sohn, Honglak Lee. "Attribute2Image: Conditional Image Generation from Visual Attributes" <http://arxiv.org/abs/1512.00570>
- [129] Elman Mansimov, Emilio Parisotto, Jimmy Lei Ba, Ruslan Salakhutdinov. "Generating Images from Captions with Attention" <http://arxiv.org/abs/1511.02793>
- [130] Gaurav Pandey, Ambedkar Dukkipati. "Variational methods for Conditional Multimodal Learning: Generating Human Faces from Attributes" <http://arxiv.org/abs/1603.01801>
- [131] Zuxuan Wu, Yu-Gang Jiang, Xi Wang, Hao Ye, Xiangyang Xue, Jun Wang. "Fusing Multi-Stream Deep Networks for Video Classification" <http://arxiv.org/abs/1509.06086>
- [132] Junhua Mao, Wei Xu, Yi Yang, Jiang Wang, Zhiheng Huang, Alan Yuille. "Deep Captioning with Multimodal Recurrent Neural Networks (m-RNN)" <http://arxiv.org/abs/1412.6632>
- [133] Samira Ebrahimi Kahou, Xavier Bouthillier, Pascal Lamblin, Caglar Gulcehre, Vincent Michalski, Kishore Konda, Sbastien Jean, Pierre Froumenty, Yann Dauphin, Nicolas Boulanger-Lewandowski, Raul Chandias Ferrari, Mehdi Mirza, David Warde-Farley, Aaron Courville, Pascal Vincent, Roland Memisevic, Christopher Pal, Yoshua Bengio. "EmoNets: Multimodal deep learning approaches for emotion recognition in video" <http://arxiv.org/abs/1503.01800>
- [134] Seungwhan Moon, Suyoun Kim, Haohan Wang. "Multimodal Transfer Deep Learning with Applications in Audio-Visual Recognition" <http://arxiv.org/abs/1412.3121>
- [135] <http://www.nature.com/nature/journal/v529/n7587/full/nature16961.html>
- [136] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, Martin Riedmiller. "Playing Atari with Deep Reinforcement Learning" <http://arxiv.org/abs/1312.5602>
- [137] <http://www.rense.com/general81/dw.htm>
- [138] <https://www.facebook.com/yann.lecun/posts/10153442884182143>
- [139] Daniel Jiwoong Im, Chris Dongjoo Kim, Hui Jiang, Roland Memisevic. "Generating images with recurrent adversarial networks" <http://arxiv.org/abs/1602.05110>
- [140] <http://www.cnet.com/news/quantum-science-is-so-weird-that-ai-is-choosing-the-experiments/>
- [141] Leon A. Gatys, Alexander S. Ecker, Matthias Bethge. "A Neural Algorithm of Artistic Style" <http://arxiv.org/abs/1508.06576>
- [142] <https://backchannel.com/how-elon-musk-and-y-combinator-plan-to-stop-computers-from-taking-over-17e0e27dd02a>
- [143] <http://www.nickbostrom.com/superintelligentwill.pdf>
- [144] Aubrey de Grey, Michael Rae "Ending Aging: The Rejuvenation Breakthroughs That Could Reverse Human Aging in Our Lifetime"
- [145] <http://www.businessinsider.com/sam-altman-y-combinator-talks-mega-bubble-nuclear-power-and-more-2015-6>
- [146] <http://slatestarcodex.com/2015/05/22/ai-researchers-on-ai-risk/>
- [147] Moritz August, Xiaotong Ni. "Using Recurrent Neural Networks to Optimize Dynamical Decoupling for Quantum Memory" <http://arxiv.org/abs/1604.00279>

- [148] S. Debnath, N. M. Linke, C. Figgatt, K. A. Landsman, K. Wright, C. Monroe. "Demonstration of a programmable quantum computer module" <http://arxiv.org/abs/1603.04512>
- [149] <https://www.microsoft.com/cognitive-services/en-us/apis>
- [150] Stanislaw Jastrzebski, Damian Leśniak, Wojciech Marian Czarnecki. "Learning to SMILE(S)" <http://arxiv.org/abs/1602.06289>
- [151] Jie Fu, Hongyin Luo, Jiashi Feng, Kian Hsiang Low, Tat-Seng Chua. "DrMAD: Distilling Reverse-Mode Automatic Differentiation for Optimizing Hyperparameters of Deep Neural Networks" arxiv.org/abs/1601.00917
- [152] Ilya Loshchilov, Frank Hutter. "Online Batch Selection for Faster Training of Neural Networks" <http://arxiv.org/abs/1511.06343>
- [153] Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, Doina Precup. "Conditional Computation in Neural Networks for faster models" <http://arxiv.org/abs/1511.06297>

14 Figures




Figure 4: Images generated by DCGAN[48][49]



Figure 5: Images generated by DCGAN[48][49]



Figure 6: Images generated by DCGAN[48][49]



Figure 7: Images generated by DCGAN[48][49]



Figure 8: Samples of LSUN images generated by GRAN[139]

Appendix 1

In [51], the age of their model was estimated as 4.45 years. It answered 54.06% questions correctly. They investigate for many questions the youngest age group that could answer it. The results are: age 3-4, 15.3%; age 5-8, 39.7%; age 9-12, 28.4%; age 13-17, 11.2%; age 18+, 5.5%. The sum equals 100% however 18+ humans answer 83.3% correctly. We could estimate that 8 year model must answer understandable for age 8 $15.3\%+39.7\%=55\%$ questions with 83.3% accuracy and other 45% with accuracy equal to baseline model. If they select the most popular answer for each question type, they get 36.18%. It's reasonable to estimate the age of that baseline model as 0 years because it's not real knowledge but it depends on this distinct dataset statistics. However, there are two baseline models in article, with 36.18% and 40.61% accuracy, both are quite simple. Which one to choose? Let's calculate. Let's define accuracy of 8 year old model as y , accuracy of 4 year old model as x , baseline accuracy as t . We have following equations:

$$4+4*(54.06-x)/(y-x)=4.45$$

$$55*0.833+45*t=y$$

$$15.3*0.833+84.7*t=x$$

with solution: $t = 46.86\%$, $x = 52.4\%$, $y = 66.9\%$.

So we estimate age of model with 57.6% accuracy as $4+4*(57.6-52.4)/(66.9-52.4)=5.45$ years.

We estimate age of model with 60.4% accuracy as $4+4*(60.4-52.4)/(66.9-52.4)=6.2$ years.

Appendix 2

It's interesting to note that popular "one billion word benchmark" corpus contains *shuffled* sentences so NLM trained on that corpus couldn't be hierarchical and understand the flow of thoughts. That's why it's so much interesting to see generated samples from a recent contextual LSTM article where the authors get 27 perplexity for Wikipedia dump. Unfortunately article doesn't contain such samples. "A neural conversational model"[17] contains samples from NN with $pplx=17$ on OpenSubtitles but this dataset is somewhat worse than Wikipedia in terms of consistency and logic of thought flow.

Appendix 3

People tend to be skeptical in front of big changes because of conservative bias.

This is well illustrated by list of famous skeptic science-related predictions [137].

Appendix 4

Sometimes it's easy to make something uncontrolled like thermonuclear bomb was made in 1953 but we still don't have controllable economically efficient thernuclear power. It might be that uncontrollable human-level AI is got in say 2022 but it might be still close to impossible to control AGI even in 2084.

Appendix 5

Yann LeCun writes about Richard S. Sutton's comment on AlphaGo[138]:

Rich says: "AlphaGo is missing one key thing: the ability to learn how the world works such as an understanding of the laws of physics, and the consequences of ones actions."

I totally agree. Rich has long advocated that the ability to predict is an essential component of intelligence. Predictive (unsupervised) learning is one of the things some of us see as the next obstacle to better AI. We are actively working on this.

However, in it's own simple world, AlphaGo is able to predict future, is able "to learn how the world works such as an understanding of the laws of physics, and the consequences of ones actions". Predictive learning becomes better very fast.

Appendix 6

Can we have another "AI winter"? There were two major AI winters in 1974–80 and 1987–93. [56]

During first AI winter in 1974-1980, a typical desktop computer achieved < 1 MFlops [57]. The most powerful supercomputer had ~100 MFlops and 8Mb main memory (only 80 Cray-1 were sold) [58].

Table 6: top supercomputers' performance during second AI winter in 1987-1993

| year | top supercomputer | performance |
|-------------|-------------------|---------------|
| 1988 – 1989 | Cray Y-MP | 2 GFlops [59] |
| 1990-1991 | Fujitsu VP2000 | 5 Gflops |
| 1992 | NEC SX-3/44 | 20 Gflops |
| 1993 | TM CM-5/1024 | 60 Gflops |

It's clearly seen even from these numbers that the situation during those winters was totally different. Even if we have AI "winter" now it would be like +20°C, which is better than +100°C AI summer.

Appendix 7

My own prediction for human-level AGI is normal distribution with "mean = end of 2017, sigma = 1 year" if there would be no major restrictions on AI research etc, though I really want somebody to give me a dozen of excellent arguments why I'm wrong. By far, I haven't heard any reasonable well-structured arguments proving human extinction due to AGI in next 10 years to be extremely (say less than 10% with 90% confidence) improbable. Arguments like "stop it or else we might go into another AI winter" or "oh, we've seen that in Hollywood so that can't happen" don't count for obvious reasons.

In my experience, discussion about AI safety with people who didn't read [24] is like discussion about deep learning with people who don't know backpropagation basics. It's always boring and misleading.

Appendix 8

I'm truly interested in discussion so feel free to send me e-mails or contact via <http://facebook.com/sergej.shegurin> ("Sergej Shegurin" is my nickname I choose for internet long ago).

Here is my list of best AI articles in *chronological* order: <http://goo.gl/7hJjHu>