Department of Information
Engineering and Computer Science
(DISI)
University of Trento, Italy

INNOVATION AND ENTREPRENEURSHIP
BASICS - AY 2017/2018

Battle Report for battle 6

Artificial Intelligence
strong AI versus weak AI

Seraph
Gabriele Gemmi [198042]
Yamini Chandrashekar [196814]
Massimiliano Luca [198449]
Giacomo Zara [172265]
Giovanni Maria Riva [198020]
Antonio Longa [198162]

I-Heart
Michele Murè [196843]
Berlato Stefano [194439]
Francesco Zambotto [194461]
Cesare Grigoletto [171726]
Filippo Bigarella [194440]

Abstract

We are the dominant species on the planet, mostly because of our mental capacities. Indeed, we have defined ourselves as being wise twice, because of our intelligence and our sense of self [16]. However, we lack high computational power, so we have developed artificial intelligence to help us solve complex problems. Now there is the possibility to start evolving AI beyond the servant role it always had, and to make it conscious too. In scenario of the battle, several problems were faced by the two teams, with the goal of securing the Advanced Grant of the ERC. This report describes such issues and presents different considerations developed upon the provided scenario.

1 Introduction

> "I believe this artificial intelligence is going to be our partner. If we misuse it, it will be a risk. If we use it right, it can be our partner" [12]
>
> *Masayoshi Son*

Nowadays it is almost impossible to look at a context in which artificial intelligence is not used: health care, personal assistants, robotics, and so on and so forth. The development of such concept, AI, dates back to the dawn of computer science, but only lately it grew exponentially. This now widespread diffusion is mainly due to the huge computational power we were able to reach, and more and more people are involved in this environment. Not everyone has the financial resources to carry out its studies, so several grants and funds are at disposal for worthy projects. This is the case of the Advanced Grant of the European Research Committee, that received two proposals from top EU universities on the artificial intelligence topic. These represent opposing views of 'weak' or 'applied' AI and 'strong' AI. The former is so called because it is focused on a narrow task and does not have consciousness of what it does. The latter instead resembles human behavior very much; it is "Defined as the technology that which would have to be synthesized were consciousness to be found". [1]

The weak AI proposal, 'I-Heart' belongs to a team from the University of Edinburgh, and the project is titled: "Heart diseases detection, analysis, and prevention through AI". Briefly, it represents the incremental innovation side of the scenario. Thus, the idea is rooted in previous researches and pursuits and tries so to increment the current public domain in that specific field. This idea seems to be referring to a topic extensively discussed and explored, but some new aspects give freshness to the project. The novelty comes from the new tests that are going to be used to identify all kind of heart diseases, not only heart attacks, on a short and also long-term, but also with the possible usage of a smartwatch to achieve real-time monitoring.

The 'strong' AI proposal, 'Seraph', is carried out by a team from the University of Trento: "Virtual study companion capable of human-like interactions and incremental adaptivity to the student". It is indisputable that knowledge is essential to humankind, and that almost everybody once in their study life had some relatable problems. The current solutions, like tutors, may be expensive, and so here AI is proposed as a cheap alternative. Apart from this application, this is the radical innovation side of the battle, given the degree of new technology that should be embedded in it [11].

It has to be noticed that the first team is given some bases to start working from, and so operates in a risky environment, while the second surely falls on the concept of uncertainty. Indeed, apart from theoretical studies, the practical and implementation field is quite unexplored and diverge from the current use of artificial intelligence, and obviously here lies the innovation.

There are many problems to be faced during the battle, and most of them are team-based. Competitors and Data gathering, thus privacy, are more related to the I-Heart idea, while the technological feasibility and the limited amount of money refer mostly to the 'Seraph' project. However, these problems are simple with respect to what should be the hot topic of the discussion: the ethical problem. Since the idea of a strong AI has been formulated, many thinkers and experts gave their opinion on it, and the literature is full of related material. Obviously, we are not talking only about science-fiction films, but instead about real critical minds that worked on this idea. The first and most debated one is the Chinese Room hypothetical experiment, which states how a program cannot give a computer a "mind" or a "conscience" [14]. Other arguments were put forward to support or refute this statement, but now the ethical question is in charge; what kind of ethics could strong AI develop? Since it has feelings and conscience, how should we behave while interacting with it? Is it better to continue evolving weak AI and avoid possible threats arising from strong AI?

The report is structured as follows: chapter 1 is an introduction. Successively chapter 2 reports the scenario in which the battle takes place, together with the views of both teams. Then, in chapter 3 thoughts about the reconciliation are given. Eventually, chapter 4 contains the conclusions of the battle and of the report itself.

2 Scenario

In 2019 the European Research Council (ERC), an independent agency which aims to find and finalize the most recent and innovative idea, receives proposals from top EU with opposed views on the same topic. The two novel ideas proposed are related to Artificial Intelligence (AI), a branch of computer science dealing with the simulation of intelligent behavior in computers, which can per-

ceive its environment and take actions that maximize the chance of success at some goal. [10]

In more detail, on one side we have the research team of the University of Trento which believes in the possibility of creating a strong AI, an entity which is very similar to humans but with the potential and power calculation of a supercomputer, so a super being more powerful and efficient than humans. On the other side, the research team of the University of Edinburgh sustains the use of Weak AI. It is artificial intelligence but with no consciousness and so more limited when compared to the strong AI. It is more similar to a program rather than a human so more easily realizable.
Both the proposed ideas are fascinating and with a lot of potentials but unfortunately, the ERC is running out of money, since they want to stay consistent with funds, they have decided to finance only one proposal.

With this perspective in mind, the two universities have to compete for attracting the attention of the commission and gather the 3.5 million euro grant to lay a foundation of their ideas [7]. For five years, the winning research team will be able to use the given resource for the project without any constraints.
Since both the teams are from the same technical background and thus have enough knowledge in topics like Neural Networks, Machine Learning, it is allowed to go into the technical details. Even though the context and the background allow the team to go into detail of the technical arguments, one crucial aspect that has to be kept in mind to maintain a socio-philosophical line of thought and therefore not consider the realization of the idea only regarding a technological issue but also as a social issue. Indeed it is essential to understand how the people will react of this new idea which belongs to a beautiful novel branch like AI since many people seem to be more concerned about unemployment problem rather than a robot plot against human beings. Speaking about robots, avoiding the battle falling into a movie discussion in which the first speech is the potential risk about what consciousness in the machine could do against human species (e.g., Terminator or Ex Machina), we had decided to keep the humanoids out of the battleground. While the Weak AI is already present in our daily life (e.g., Siri, Cortana, and so forth) the Strong AI is more a novel and far idea so the feasibility of that concept is not granted a priori and can be questioned by the other team or the audience during the battle. It is up to the Strong AI research team to defend that point and convince the ERCs commissioners of the feasibility of their idea.

3 I-Heart

3.1 The problem

Heart diseases are amongst the first causes of death in the EU. They represent a wide group of medical problems that affect the circulatory system (heart, blood

vessels, and arteries), often resulting from atherosclerosis; Some of the most common diseases that affect the cardiovascular system include ischaemic heart disease (heart attacks) and cerebrovascular diseases (strokes). Eurostat[3] is an entity responsible for providing statistical information to the different institutions located in the EU. It stated that in 2014 1.83 million deaths were resulting from diseases of the circulatory system in the EU-28, which was equivalent to 37.1 % of all deaths — considerably higher than the second most prevalent cause of death, cancer.

3.2 State of the Art

In response to this phenomenon, in recent years, many researchers in different fields have been working on new medical technologies to improve the situation. One such technology is the use of Applied AI in the medical field[8]. An example is given by neural networks, that can learn from past cases like the human brain; These networks can accurately diagnose some diseases like eye problems, identify many forms of cancer or somewhat predict heart attacks.

Talking about the prediction of heart attacks made using machine learning algorithm, a team of researchers from the University of Nottingham already published a paper[15] about it two years ago on April 4, 2017. They thought that the current approaches, to predict cardiovascular risk, fail to identify many people who would benefit from preventive treatment, while others receive unnecessary intervention [17]. So, they used Machine-learning to improve the accuracy of the predictions by exploiting complex interactions between the commonly known risk factors.

3.3 Proposed Solution

This research, which inspired us, led to the birth of our idea, the idealization of a new AI model, which wants to tackle the problem of heart diseases from a new perspective. In particular, by taking advantage of five new tests, we will improve the accuracy and precision of the predictions gaining results that will be far better compared with current standard strategies that focus on blood pressure, cholesterol, diabetes and smoking history.

It was proven by UT Southwestern Medical Center[2] that these five simple tests, which we want to use, are compelling in identifying unexpected risk among individuals with few traditional risk factors. So using them, we will be able to correctly diagnose all those *people who would not be aware that they are at risk for heart disease and might not be targeted for preventive therapies*.

The five tests, and the information they provide are:

- *A 12-lead EKG provides information about hypertrophy or thickening of the heart muscle.*
- *A coronary calcium scan, a low-radiation imaging test, identifies calcified plaque buildup in the arteries of the heart.*
- *A blood test for C-reactive protein indicates inflammation.*
- *A blood test for the hormone NT-proBNP indicates stress on the heart.*
- *A blood test for high-sensitivity troponin T indicates damage to heart muscle. Troponin testing is regularly used by hospitals to diagnose heart attacks, but high-sensitivity troponin fine-tunes that measure, pointing to small amounts of damage that can be detected in individuals without any symptoms or warning signs.*[2]

Our research plan consists development of two Artificial Intelligence models that can detect, in the long and short term respectively, heart attacks and various heart diseases such as hypertrophy, inflammations, damages to the heart muscle or formation of plaques in the heart's arteries.

3.4 Research Plan

To accomplish everything that we said until now, we will start by contacting and subsequently partnering with hospitals and data controllers all over Europe. The reasons that are instigating us to make these partnerships are mainly two. From one side we want to collect their healthcare data for a low price and so cutting the costs of our research. From another, we want to make sure that the market for our technology will not close down while we are busy conducting the research. The trade-off with the hospital would then be an exchange of their healthcare data for the exclusive right of use of our AI models in a limited period of three years.

Once we gather all the necessary data we will start developing a neural network based on the latest discoveries in the field of deep learning, trying to find the best configuration for both the network and its hyperparameters. We will use a genetic approach to evolve the model and tune its parameters, this type of approach requires to train the model multiple times while modifying the parameters to "breed" the best model possible.

Obviously training a big model multiple times requires much computational power. The genetic approach alone will help reduce the number of models we need to train concerning the brute force methodology. In fact, instead of trying

all the possible combination of parameters, a population of models is created, the best elements of it are then combined to breed the next generation. This process is repeated until an optimal model is obtained.

The genetic optimization process, despite being less expensive than the brute force, will still require us to invest a significant amount of money in obtaining the needed computational capability, thankfully we saved much money when gathering data, so we can spend the most of it to obtain said resources

We want to point out that, in this phase, it will be utmost important for us, to follow and implement, all the necessary procedures to guarantee the anonymisation of the data. Staying in line with the new General Data Protection in force from 25 May 2018 and its concept of privacy by design; which encourages product, service and application developers to take into account the right to data protection from the design stage. Thus reducing the risks of processing and, consequently, the impact on the rights of the data subjects.

In the next phase we want to distribute a beta version of our AI model, dedicated to long-term prediction, among our partners so that they will be able to use it as a support for their doctors. This will let us test our technology in a controlled environment and see how well it behaves compared to commonly used algorithms that represent the current state of the art; doing this will allow us to tune and improve our solution easily. In parallel with the attunement of the first AI model, we will work on the second AI model to detect heart attacks and diseases in the short term, being it a week or two. This algorithm will also, at least in the beginning, be distributed and tested in the hospitals.

3.5 Future Plans

By the end of the research, we will like to strike a second partnership with wearables vendors, to produce and sell smartwatches. These smartwatches will make use of the newest biometric sensors which can collect, in a noninvasive way, health-data, like blood samples, from the wearer and use it as an input for our model. What we want to offer to the people is an efficient and easy way to monitor their health in real-time; In an ideal scenario every person will have a smartwatch like this, allowing us to not only keep on refining our models but to open the medical field to a wide new range of possible innovations, thanks to the huge quantity of data that will constantly be gathered.

Moreover "AI assistants/programs could significantly reduce medical costs by eliminating office visits with online care. Patients would be asked to submit data more frequently via online medical records, and the improved line of communication could result in fewer office visits. Further cost reductions could come from efficient AI" [8] models that can diagnose and screen high-risk patients as

well as eliminate human errors in record keeping and diagnosis.

To summarize, heart diseases represent a considerable danger to all of us. We believe that by exploiting the last decades of biomedical research, we can build sophisticated artificial intelligence which will be able to save people's lives, by predicting such diseases before they take place. Our ultimate goal is to bring this model to everyone's wrists, to increase the quality of life to the population all over Europe.

4 Seraph

4.1 The problem

The field chosen by the Strong AI team is **learning** and **education**. To clarify the specific **problem** that has been addressed, the team proposed a simple **scenario**: an average student who is attending a bachelor course at the university, in a scientific subject. The student is facing all those problems which are typical and ordinary, especially during the first year of university new environment, tough courses, balancing time, difficulty in solving exercises and working on projects.

4.2 Further aspects

The student is a human being, which implies also having **personal** issues and problems, such as psychological, social, personality aspects that might, for instance, cause the student to be poorly prone to the frontal education system which we have nowadays. The scenario is intended to underline the fact students are all very **different** from one another, under the aspects that have been mentioned and many more. This differences would often require the education system to be customized for each student, to address him as an individual. This is a problem that standard teachers cannot solve. There are many apparent reasons behind this, but the simplest one is that a teacher might have to deal with more than one hundred students even just for a single course. Be the teacher as good as he/she may; there is little that can be done to fix such an issue.

4.3 Current solutions

Of course, there are existing and well-known solutions to address this problem. For example, we may think of **private teachers**. This solution, though, has some drawbacks. First of all, private teacher are **expensive**, and not everyone might be able to afford them. Secondly, a private teacher is still a person who's living his/her own life, and his availability for the student, even if it is customized, it is limited. Our student might also choose to attend other courses, for example, tutoring classes or online tutorials. This still doesn't solve the problem though, since these solutions still imply a one-to-many detached relation, often limited from the fact of occurring at fixed times, in fixed places.

4.3.1 State of the art: eTeacher

Another solution, which gets closer to the team's proposal, is eTeacher [13]. eTeacher is a **virtual study companion** that has been developed and tested by a research team of the University of Buenos Aires (Argentina). This tool is based on machine learning, and provides a personalized teaching method for every single student, by exploiting the capabilities it has of adapting the tasks to the student's skills and performance.

eTeacher is still missing some aspects which the team considers very important. In particular, we are talking about those aspects which are not directly and univocally related to studying.

4.4 The proposal

The device proposed by the Strong AI team takes the name of **Seraph**. It is a virtual study companion that does most of what eTeacher already does, plus significant addition: the tool would exploit it is AI capabilities to adapt to the student not just according to his skills and performance as a student, but also by his **personality**. Seraph is intended to incrementally develop a **knowledge** of the student's life, character, and sensibility, to adapt each task to his psychological approach to studying. The tool is meant to be able to **empathize** with the student, assuming the role of a friendly, sympathetic companion, able to **handle unforeseen situations**, and cases of **ambiguity**, just like a human would. This is one of the fundamental differences from eTeacher, which instead produces its outputs **deterministically**, according to its machine learning basis.

The team members believe that such a technology, meant to be an **addition**, not a replacement, for the current system, would represent a radical improvement in the learning process of average students. This belief is based on the idea that a person would be much more facilitated in exploiting his potential once he is followed and addressed in a **personalized** way, meant to take into account his/her **personality**. The overall idea is based on the purpose of exploiting prospective rather than retrospective, targeting a whole new type of technology and application [5]

4.5 The engineering plan

The team proposed a sketch of an **engineering plan**, to state a prediction on how and when the available money will be spent. The plan is defined over five years, and it would go through the following steps:

1. **team making**: the idea is collecting a team of psychologists, pedagogues, and even philosophers

2. **data collection**: the team formed is meant to develop an innovative way of collecting data, by observing, surveying and interviewing students and teachers, taking into account both academical and personal aspects

3. **prototyping**: iterative development and testing of a rough prototype

4. **evaluation**: evaluating the outcome of the project, analyzing results and looking at the future

In the future, one possibility consists of **extending** the potential of the product to different schools, and to any possible kind of learning and education environment.

4.6 Issues

For obvious reasons, the possible issues and drawbacks of the project have been taken into account, according to the possibilities. Here follows an outline of the main points that have been identified, each with a proposal for a possible solution.

4.6.1 Feasibility

It would be immediate to argue about the feasibility of the project, considering the current stage of the strong AI research. The team, though, pointed out that their project is indeed a **research** project. Given the fact that the field is very young, the project is meant to be mainly **experimental** and **explorative**. The risk is part of the game when it comes to radical innovation, and the team is aware that its idea consists of an **attempt** which obviously gives no guarantee of fulfillment.

4.6.2 Cost

Strong AI was never known to be cheap. Thus, it is reasonable to argue that the initial budget will hardly be enough for the requirements of the project. For instance, given an ideal final product, would it be cheaper than a private teacher? In general, being this just a research project, the team is still quite far from taking into account economic issues such as these. That is the team plan to spend the money of the fund to the develop the **research**. The problem of financing the actual implementation of a **working** final product, as well as the problem of **selling** the product, will be faced at a later stage, in which the whole situation might be completely different under any possible aspects. If the preliminary research goes well, more funds are likely to intervene. Plus, given the **exponential growth** of technology, nothing tells us that during the next years we will not be facing radical changes and improvements.

5 Reconciliation

The main focus of this debate was the capability of making correct decisions in an unknown environment. Each team, given a starting idea, had to choose a suitable aspiration level, and then choose a course of actions, that would lead

the team to win the grant. We want to remember that the European Research Council (ERC) is granting fundings to different research teams with the intent to encourage transformative researches at the forefront of science. The other team's idea was just not innovative enough to qualify for the grant, even though it was useful and easily implementable. Naturally, it must be stated that not all the topics have been covered during the battle, so different scenarios from the resulted one may have come to be.

Firstly, there are ethics and artificial intelligence. Many articles in the Internet can contribute to this topic, like the one about Tay: a Microsoft Bot which became racist after responding to tweets and chats for a brief period [9]. It is easy to see how this argumentation could have brought many changes to the whole battle and maybe even leading the Weak AI to victory. Another retort comes up by mentioning Professor Jun Luke Huan [6] and his studies about how machine learning algorithms should be transparent and easily inspectable. From this point of view, the weak AI algorithms are more transparent to inspection than the strong AI ones. In other words, it is possible to understand the reason why a weak AI made a choice, and if we can improve it. With strong AI instead, this is not possible.
Another problem, strictly linked with ethics, is the need for Strong AIs to have moral status. Francis Kamm defined the moral status as: "X has a moral status = because X counts morally in its own right, it is permissible/impermissible to do things to it for its own sake" [4]. This means that a strong AI should understand what is good and what is bad. If we take as an example a robot that saves people from fire, it may know that it is better to save children before others; but, on the other side, often deciding what is right and what is not may be an issue, even for a human being.

On the other hand, Weak AI uses machine learning. Those algorithms are not able to provide different answers for similar inputs, since are not able to take into account contextualization. Moreover, machine learning algorithms always provide answers. After that, a result is provided even if it has to be taken in uncertainty settings. Weak AI makes decisions that best fit risky or uncertain environment. The latter requires that to know the "truth that is out there", the data has to be collected to improve the decision. Weak AI is not meant to look for new data by itself but just improve on data already provided. In this sense, it is dangerous to apply such algorithms in the medical field.
Another important point that has been partially covered is the reliability of the two ideas. While weak AI is the state of the art, it is imaginable that will be enforced by 2019 (which is when the battle is set), the study on Strong AI is almost at the beginning, and so it is possible that something emerges before 2019 but, probably, the team will have to start from scratch.
Nowadays, it is possible to simulate a brain with a computer, but it costs much more than 3.5 million euros as written in İt is believable that in the future, the costs of infrastructures will decrease, but not enough to build a Strong AI with so little grant. Furthermore, the theory behind Strong AI is not well defined,

and so a huge amount of money has to be spent for it (according to the graph on the slides of the strong AI team).

The experience accumulated during the battle led us to propose a scenario in which the two proposed ideas ceased obstructing each other and, instead, decided to co-create something new together. In this scenario we will embrace an "open innovation" approach, hoping that anyone can bring benefits and give their contributions to the research. This way not only we will be able to deliver faster results, but it will be cheaper too. Once the model for a Strong AI will be a real thing, everyone will have the possibility to make profits, by proposing new market solutions. As said before, the Seraph team will use it to improve education, or the I-Heart team will be able to make more precise and stronger predictions. Obviously, not everything that glitters is gold, as each team will have to take the right countermeasures to guarantee that they do not get themselves cut out from the market.

Both teams proposed great and, in their way, innovative projects. Moreover, there were a lot of similar problems too, like sustainability and transparency of the algorithm. Having said all this, we believe that many other plausible scenarios would have been considered, leading to concretization of different topics and actions taken by both teams to win over the other part.

6 Conclusions

The reconciliation aims, first of all, to take advantage of the already consolidated and still evolving capabilities of weak AI, like the fact that it operates in a *risk* environment, its strong contributions in many research areas and the contained cost concerning the improvement rate. On the other hand, we want to avoid a strategy for which we only keep improving incrementally what we have: our goal should be pushing harder on a divergent direction that can bring to radical innovation. In this sense, the reconciliation embraces the Strong AI research as a reason to start from the base on what is disruptive but already firmly grounded, such as the point where Weak AI stands now. The new technology would exploit open innovation as the new big paradigm shift of business. The team has chosen to target the ability to deal with a strong ambiguity environment, which seemed to convince better and inspire the audience.

Since the battle focuses on the exploitation of artificial intelligence capabilities, one topic described during the lecture that has been addressed by both the teams is the idea of **environment**. Both the projects, indeed, are somehow based on the capability of making **predictions**. In particular, as it is stated in the introduction, the Weak AI product is based upon machine learning, which is based on the knowledge of certain probability distributions. Indeed, this is exactly the general principle upon which *I-Heart* is intended to make its predictions about heart diseases. This is what we defined during the lectures, as a

situation of **risk**: the knowledge of the probability, indeed, does not guarantee the accuracy of the result. On the other hand, the Strong AI product is meant to exploit its potential to handle unforeseen situations. Unlike any weak-AI-based machine, a strong AI device is intended to behave like a **human**, that is, it aims to produce its outputs and manage its behavior by its **interpretation** of the world, rather than on a **deterministic** algorithm. For this reason, we could say that *Seraph* is meant to face situations of **ambiguity**.

This latter aspect is also deeply linked with what we have been discussing during the lesson about **sensemaking**. The key point of a strong AI, indeed, is its capability of sensemaking, that is, the ability of "improvising" a course of action, in case of a situation, in the world which is different from the expected one. A tool such as *Seraph* is intended to exploit this ability to model its perspective of the student it is following.

As quickly pointed out in the introduction, the two projects presented during the battles are good examples of **innovation**, in its incremental and radical version. The weak AI project, indeed, is a clear instance of what we described during the lecture as **incremental** innovation, as it aims to improve a pre-existing situation, exploiting new technologies to refine the results of a method which is already implemented in its basics. On the other hand, the strong AI project is more of a **radical** innovation example, since it aims to exploit a brand new and still poorly implemented science, to deeply innovate a system (the education one) which is currently based on a completely different approach.

Finally, both teams underlined the importance, for their projects, of the **co-operation** with other entities and institutions. The *I-Heart* team, indeed, said explicitly that their achievements are intended to be made possible by the shared efforts with medics and hospitals. Similarly, the business plan of the *Seraph* team includes a **network** of cooperation with universities and research labs. These can both be seen to what has been referred to in class as "open innovation" since the organization is partly **outsourced** and based on a mutual dependency between different parties.

In the introduction, we highlighted that there were many topics and problems to be addressed during the battle, mostly team-based. Obviously, these were the most exploitable points from which attack each other position. So the questions and the discussion were mainly low level and related to the scenario: data gathering, privacy and lack of novelty were the hot points for the Edinburgh team, while feasibility of creating a strong AI in few years and the not ambitious expected application were recriminated to the Trento University. Sadly, the battle lasts too short to completely explored all the topics; for example, very little was told about the technical details of artificial intelligence. However, the most important one was completely left out during the debate: the ethical problem. This is without a doubt what was expected by all of us to be the center of the teams' argumentations. Indeed, the preparation of the two sides was focused

on this aspect of the confrontation, and several articles were read and discussed within the crews.

This consideration leaves a door open to at least another whole battle or some other sorts of follow-ups, given the bulk of material related to the ethical problem and the issues that arise from it. These questions were mentioned at the beginning of this report but, unfortunately, they are still pending without an answer. It would be interesting to see, in a planned way, how the battle would have ended up if these points were addressed. Maybe the weak AI team would have won, or maybe nothing would have changed. In a possible further confrontation, more space may be given to ethical and philosophical aspects such as this one, making more interesting and exciting a possible realization of the visions proposed in battle, as well as their eventual refusal.

References

[1] Igor Aleksander. Artificial neuroconsciousness: An update, 1995. IWANN, Archived from the original on 1997-03-02. https://web.archive.org/web/19970302014628/http://www.ee.ic.ac.uk/research/neural/publications/iwann.html.

[2] UT Southwestern Medical Center. These 5 tests better predict heart disease risk, 2017. www.sciencedaily.com/releases/2017/03/170331120358.htm.

[3] eurostat. Cardiovascular diseases statistics, 2017. http://ec.europa.eu/eurostat/statistics-explained/index.php/Cardiovascular_diseases_statistics#Deaths_from_cardiovascular_diseases.

[4] K. Frankish and W.M. Ramsey. *The Cambridge Handbook of Artificial Intelligence*. Cambridge University Press, 2014.

[5] Bruno Latour. *Science in Action*. Harvard University, 1988. http://www.hup.harvard.edu/catalog.php?isbn=9780674792913.

[6] Xiaoli Li. Transparent and interpretable machine learning. http://www.ittc.ku.edu/~jhuan/kddT/safeDataAnalytics_Li.pdf.

[7] Matt Mahoney. The cost of ai. http://mattmahoney.net/costofai.pdf.

[8] Amanda Page. 9 ways artificial intelligence is affecting the medical field, 2012.

[9] Sarah Perez. Microsoft silences its new a.i. bot tay, after twitter users teach it racism. https://techcrunch.com/2016/03/24/microsoft-silences-its-new-a-i-bot-tay-after-twitter-users-teach-it-racism/.

[10] L. Radick. Artificial intelligence in healthcare: The current, compelling wave of interest. *Healthcare Executive*, 32(5), September 2017.

[11] Jane E. Dutton Robert D. Dewar. *The Adoption of Radical and Incremental Innovations: an Empirical Analysis*. Informs, 1986. Introduction, page 1423. http://www.jstor.org/stable/2631501.

[12] Masayoshi S. Masayoshi quote. https://www.brainyquote.com/quotes/masayoshi_son_845514.

[13] S. Schiaffino, P. Garcia, and A. Amandi. eteacher: providing personalized assistance to e-learning students. *Computers and Education - Elsevier*, 51(4):1744–1754, December 2008.

[14] John. R Searle. *Minds, brains, and programs*. University of California, 1980. http://cogprints.org/7150/1/10.1.1.83.5248.pdf.

[15] Joe Kai Jonathan M. Garibaldi Nadeem Qureshi Stephen F. Weng, Jenna Reps. Can machine-learning improve cardiovascular risk prediction using routine clinical data?, 2017. https://doi.org/10.1371/journal.pone.0174944.

[16] Peter Norvig Stuart J. Russell. *Artificial Intelligence, A Modern Approach*. Prentice-Hall, 1995. Chapter 1, Introduction, page 3. https://pdfs.semanticscholar.org/bef0/731f247a1d01c9e0ff52f2412007c143899d.pdf.

[17] Bob Yirka. Ai systems found to be better than doctors at gauging heart attack risk. https://medicalxpress.com/news/2017-04-ai-doctors-gauging-heart.html.