Navigating the fudge and breakthroughs of artificial intelligence in China [1]

Rapport supervisé par **Pierre-Étienne GIRARDOT** et rédigé par les **ingénieurs-élèves** du Corps des Mines [2] (promotion 2018)

Free from GAFAMs, China is a brave new digital world

An ecosystem with Chinese characteristics

China, home to tech giants

The Chinese economy has taken its own path to achieve its digital transition, based on massive use of smartphone payment, digital retail and social networks. The three large conglomerates which emerged, mirrorring the so-called GAFAM, structure the Chinese digital ecosystem: Baidu, Alibaba and Tencent (BAT).

Taking advantage of the size of the Chinese internal market, and leapfrogging Western technology, Chinese big players have managed to reach a level of technology close or sometimes better than their american counterparts.

Baidu made its success with its search engine, Tencent focused on gaming and messaging applications while Alibaba developed the leading retail platform. Leveraging their local de-facto monopolies, these companies added new services and bought hundred of start-ups, building digital ecosystems that are rival in each part of the economy: food delivery, payment app, music streaming, bike sharing, etc.

This fast digitalisation was supported by the low quality of existing services in the Chinese economy (limited deployment of credit card and shopping centers, difficult access to public services, etc.).

Two rival ecosystems, consolidated around payment apps

A striking element was how widespread mobile payment is for day-to-day transactions, even in remote parts of the country. Even beggars on the streets collect charity through smartphone apps. As we were able to witness, it is impossible to pay using a credit card in most shops. Last year, the total amount of transactions made by smartphone was almost one thousand times higher in China than in the USA (USD 36,000bn vs. USD 50bn).

The solutions are mainly provided by two big companies: AliBaba through *AliPay* and Tencent through *WeChat Pay*, both relying on easy to deploy QR codes. AliPay has around 150 million active monthly customers in China and WeChat over one billion overall customers. By contrast,

---

TENCENT

The largest Chinese tech conglomerate

Founded in 1998, Tencent is one of the three major Chinese technology companies (T of BAT), with a market cap of USD 520bn and a turnover of USD 36bn in 2018. Tencent gets most of its revenues from gaming (41%), and owns the two biggest Chinese social apps, *WeChat* (a Whatsapp-like application with enhanced capacities) and QQ. Tencent applications are responsible for 47% of the Chinese internet traffic, and Tencent remains very China-centered, despite recent efforts. Being the backbone of a payment system (WePay) and of mini-programs (paying taxes, filing for divorce, buying movie tickets), WeChat has become an ecosystem in itself, in direct competitions with smartphone OS.

What do they use AI for?

Recent regulations in video games have pushed Tencent to develop computer-vision applications in non-core business sectors such as healthcare and smart cities. In healthcare, their tools to detect lung cancer in scanner images and oesophagus tumor in video allegedly beats human diagnosis by a few percents. Thanks to their performances and to their ability to deploy rapidly the algorithms through WeChat mini-programs, they have been recognized by the Chinese government as the national leader to build an AI ecosystem for healthcare, and thus a major actor in setting the standards.

Our opinion on their AI abilities

We have doubts on their capacity to execute on AI, for operational and technical reasons. The management does not encourage the kind of aggressive innovation that gave birth to WeChat, afraid that the government would reduce their power. Moreover, Tencent has a reputation for siloed data, which explains why they are less advanced at AI-driven advertising and recommendation than some competitors.

Apple Pay only has 127 millions customers worldwide, even though this app is installed into every iPhone.

This trend can be explained by the fact that smartphones spread really fast throughout the country and that most people had no credit or debit cards before. Moreover, American leaders Visa and Mastercard are banned from the country. The Chinese banking system thus shifted directly from cash to mobile payment. The Chinese market is huge enough to drive alone the demand of a fifth or a sixth of the entire world population. Furthermore mobile payment in China is entirely based on QR codes or face recognition, allowing a much faster development than NFC technology in shops and retail, but at the cost of lower security. Security risks are however well accepted by the citizens thanks to the benefits provided by digitalisation.

The deployment of the payment system has enabled both Alibaba and Tencent to link all applications to their own systems, creating ecosystems nearly as central as the Western application stores for smartphone. This payment ecosystem allows AliBaba and Tencent to gather lots of data that can then be used to profile customers, target advertising and enhance services.

The competition between AliBaba and Tencent is striking. For instance, a link to Alibaba couldn't be shared on *WeChat*, and *WeChat Pay* isn't accepted as a way of payment on many Alibaba retail applications. One of the next steps will be means of payment based on facial recognition, that were mentioned in many of our visits. Developing the ecosystem around their mini-apps is tomorrow's competitive advantage for tech giants in China.

Everything is social

Tencent owns the two biggest Chinese social networks, *WeChat* (a Whatsapp-like application with increased possibilities) and *QQ*. While *QQ* is mostly entertainment-based and dedicated to young people, *WeChat* seems to be a more comprehensive application, specialized in social payments, with embedded mini-programs allowing to pay for transportation, doctors, goods at retail stores and so on.

Beyond the strength of WeChat, many services have a social dimension, as referral links are a very widespread way of creating desirability for new apps, and even shopping becomes social, with friends competing on smartphone games to earn limited editions sneakers. Interestingly, this social dimension stops when freedom of expression is concerned, as the government can prevent any message to be sent on WeChat if it contains words that are deemed against the regime: "Winnie the Pooh", a common way of referring to Xi Jinping is now banned.

ALIBABA.COM

The Chinese Amazon

Founded in 1999, Alibaba is one of the three major Chinese technology companies (A of BAT), with a market cap of USD 480bn and a 2018 turnover of USD 40bn. They started by deploying e-commerce platforms, and have since built an ecosystem around them, rivalling Tencent's: cloud business (AliCloud, much smaller than Amazon Web Services), payment system (AliPay), financial services (Ant Financial), and many startups. They currently push towards a mix of physical and digital shopping experience, for instance through augmented-reality smartphone games in shops.

What do they use AI for?

Optimization is a core activity of Alibaba. Optimization of the supply chain is a key field for their AI and data researchers, as logistics is a major source of costs for Alibaba. The group is also developing smart recommendation algorithms based on the purchasing or navigation history of a given user. While these are the main AI applications at Alibaba, they have been working on several projects in order to revolutionize the consumer retail shopping experience: they promote "shopping as an entertainment", using AI solutions to enable it and to bring online and offline retail together.
Alibaba developed through Alipay, a loan / credit / insurance service, based on the data and the behavior of its user. The algorithm is building the "sesame score" that gives access to coupons and rewards.
In a completely different business area, they are developing new capacities and technologies for smart cities. Most of the applications they showed us involved merging different data sources, combining city cameras, traffic information or incident reporting feed. They have been selected by the government as national leader for AI in the field of smart cities.

Our opinion on their AI abilities

The efficiency of their algorithm to suggest new purchases seems to be far behind Amazon's, but they keep collecting huge amount of data on their users. They have impressive computing capacities, but those seem focused on online traffic more than algorithm computation: they register up to 325,000 orders per seconds on their website on peak shopping days.
Alibaba is one of the leaders on Smart City solutions. We only got to see a digital city operation center in their showroom, as we visited Alibaba cloud whose core activity is far from the smart city business.
Like all the other major companies, they built their own deep learning platform to develop their AI applications.

Online and offline merge together

Today, the big aim for both Tencent and Alibaba is to merge physical and digital lives. Using their financial power, they are buying or backing startups that bridge the gap: e-commerce startups which are able to deliver food within thirty minutes, apps that gamify shopping to earn fidelity points and limited editions. One of Alibaba's goal is to bring online and offline retail together, in a customer-centered system using AI and digital technology, breaking the traditional supply chains where the customer is only at the end.

As the ecosystems consolidate around the *AliPay* and *WeChat Pay* payment systems, Tencent and Alibaba gather troves of data and monetize it, mostly inside their respective ecosystems. Some marketing analysis and data consolidation are then sold as a service to retail businesses, allowing them to better understand the profile and taste of their customers.

How did China build native digital ecosystems?

Shenzhen, the "Silicon Valley" of hardware

Shenzhen was originally a small port. In its attempt to open the country to the rest of the world, Deng Xiaoping decided to make it a special economic zone in 1979 with one clear mantra: "poverty is not socialism, being rich is glorious". The city specialized very quickly, becoming the manufacturing hub for electronics on the planet. Shenzhen is today one of the main supplier of electronic devices in the world. Thousands of different vendors can be found in the hardware market, which is made of several entire blocks in the city on many levels.

International companies and startups came to Shenzhen to be able to build very quickly and cheaply a product before shipping it, testing it and then scaling the production. For example, the French company Wiko simply buys white label hardware to assemble functional smartphones.

This habit of customizing and experimenting was very important for the development of the tech mindset in Shenzhen. Concentrating all the manufactures around Shenzhen, they were then able to move up the value chain and build software. In parallel, Chinese people were eager to adopt any technology improving their lives, adopting online payment extremely fast.

The gentle push of the government

The role of the Chinese government in building this ecosystem remains unclear to us, but certain. Some companies underlined the role of the government to loosen the rules within the economic special zone of Shenzhen, thus enabling Chinese internal migrants to come take jobs in a low tax environment (20% less than in France). Other said the state had a more proactive approach, setting strict guidelines (on cashless society, for example), stirring competition and massively subsidising the ecosystem. It even seemed to us that in the facts, no technological development is possible in China without, at some point, the approval of local or governmental authorities.

Another major role played by the government was to protect its huge interior market from the competition of American tech companies: for instance Google and Facebook are still prohibited in China. This played a key role in providing Chinese players enough time to enter the market, helping them take an almost complete control of it.

State-owned banks and municipalities funds play a major role in the financing of the tech environment through dedicated investment.

According to the manager of the SZ OIL start-up incubator, however, after the success of the regulatory sandboxes in special economic zones, the focus of the state-level has now moved to historical cities. The main AI hub is now in Beijing (we did not have the opportunity to visit it however), and Shanghai has also taken a key role in the AI wave. This doesn't mean that Shenzhen

has been completely left apart, and the dynamism of the city is a good indicator of that. The historical presence of large companies like Ping An, Huawei and Tencent, founded during the special economic zone trial, also remains a major asset, as do the presence of the electronic market.

As for the regulation of technology, there seems to be a different and pragmatic mindset, both amongst the population and the government, that enables a faster adoption, by focusing on developing the applications first, and dealing with the issues (ethics for instance) only when actual difficulties arise. The main guidelines of the administrators seem to consist in letting competition take place, after which national leaders naturally differentiate themselves. This idea also results in letting cities compete against each other.

The global battle for talents

One of the main issues faced by tech companies in China is to be able to attract talents. Salaries in the tech companies in Shenzhen are said to be already 30% to 40% higher than in Paris and to be increasing 10% every year. At the same time, the price of real estate is also soaring.

The Chinese government is trying to develop universities and laboratories on the digital economy and AI, however there is still a long path ahead. For instance, Shenzhen has only two universities for 13 billion inhabitants. Thus, Chinese companies started trying to attract overseas students back to China. The city government of Shenzhen deploys very generous grants to attract some international talents, especially those trained in the US, but has not been very efficient in attracting engineers or PhDs outside the Chinese diaspora yet.

Artificial intelligence: China's new target

The shaping role of the government for the AI sector

AI is a political priority

The words of Russian President Putin must have been heard in China: the nation that leads AI will rule the world. AI has been identified as one of the national priorities by the government with its "Next Generation Artificial Intelligence Development Plan", which provides a long-term vision and the biggest investment worldwide. The goal is simple: becoming the worldwide leader by 2030. The Chinese tech ecosystem, described in the previous paragraph, is now moving towards AI.

Since last year, Chinese AI industry has grown by 67%. Some skepticism is warranted in front of some dazzling numbers, as bad-quality research and empty offices are still widespread.

We visited the Zhejiang Lab in the Hangzhou's Artificial Intelligence Town, a campus dedicated to AI research. The goal for this lab is to grow from 182 AI researchers today to 3,000 AI researchers by 2022. The lab is funded through USD 10bn by the government and Alibaba Group, in order to quickly develop AI technologies, an impressive magnitude for a regional lab.

Access to data is also eased by the government. It gives access to its massive datasets to private companies, which can then export their technologies (exemple: datasets of face IDs).

Governmental influence on the choice of the AI application

China focuses on enhancing the life of the Chinese citizens.

Thus the choice of the AI applications that are promoted is based on the main challenges of the country as seen by the government: video surveillance bridges the lack of security and trust, more efficient healthcare fixes the lack of experienced doctors, slow autonomous cars replace open and unsafe motorcycles used by old people in any weather conditions...

The Chinese Ministry of Science and Technology has identified five priorities, and appointed a firm as "leading platform" for each of them:

- Smart healthcare - Tencent
- Smart City, including security - Alibaba
- Autonomous driving - Baidu
- Learning Platform - SenseTime
- Voice intelligence and Natural Language Processing – iFlyTek

During the trip, we visited some of the leading platforms, and were presented additional technologies and applications.

Main applications

Big Brother comes to town: Computer Vision for security

One of the main drivers for this technology is security, be it for the citizen or for the state. At an individual level, facial recognition is used to unlock phones, to withdraw cash at ATMs, to replace badges for access controls and much more. At the state level, video analysis is used to extract data from all the footage taken by the many surveillance cameras. As anyone who has visited China lately can tell, security cameras are everywhere, and such a large number of videos can not be analysed by human-based means.

SENSETIME

The leader of computer vision powering all the others in the background

Sensetime is a startup company founded in 2014. It is now valued at more than USD 5bn, with more than 2300 employees and 700 major clients and partners. They present themselves as the "most valuable AI company in the world". Its incredibly quick growth is tightly linked to the development of large scale video surveillance algorithms by the Chinese government.

What do they use AI for?

They are known for being world leaders in facial recognition, and have reached unequaled precision, thanks to the their database of pictures that contains more than 2 billion pictures representing 200 million different people. More than one out of two smartphones on the Chinese market embeds their technology in one way or another: the most common applications being unlocking the phone with your face and beautification filters for camera applications.
Besides this, they are working on other use cases such as behaviour detection, crowd analysis and employee supervision. They also showcased their systems for human machine interaction with gestures, be it hands signs or full body movement, even though those are not used at scale yet.

Our opinion on their AI abilities

Sensetime is fully autonomous in fulfilling its 3 biggest needs: huge datasets provided by their customers and labelled internally by 900 dedicated employees, cloud computing facilities and their own AI training platform (Sense Parrot). However, they acknowledged that their biggest asset was their internal datasets for AI, more than 10 billions of images and videos. This is indeed a barrier to entry for competitors, which complements their internal knowledge to develop new use cases. Sensetime has been recognized by the Chinese government as one of the 5 platforms for next-generation AI, along with Alibaba, Baidu, Tencent and iFlytek. However, the repartition of the value-added between Sensetime and its clients, such as Alibaba, remains unknown. Therefore we cannot tell precisely which company is the most influential in the AI value chain in China.

Video analysis helps to focus the attention on where it is needed, through intelligent crowd analysis, detection and management of traffic congestion, or even through detection of violence in real time.

Another major use case is person detection: the ability to retrieve in a few seconds someone in the videos produced by all the cameras in a city in the last month is very valuable for the police, who can analyse the journey of a suspect instantly. For instance, SenseTime rolled out, in a 5 million inhabitants city, a system which is able to find all the places where you have been in the city in the last month within one second!

We have also been told by several companies that they powered facial recognition in airports based on a list of dangerous people furnished by the government. The algorithm recognizes anyone from this list as soon as they enter the airport and should trigger an alarm even before they approach security gates.

Computer vision has been recognized as a key technology by the Chinese government. The first focus of the technology was facial recognition, which is now a widely used technology in China. Because of its sensitive status and as the country wants to keep control on the big amount of data it has already gathered to develop these technologies, export restrictions have been placed by the government. According to the companies we visited, the government is also very cautious with the data it provides companies with when working on security issues. For instance, they ask companies to run their algorithm on a government-provided cloud with no possibilities to export the results - thus the commercially available algorithms are trained on the companies' own data, or the data provided by their clients, which they have plenty of.

Who needs a doctor when they have a smartphone? Smart Healthcare

Because of the lack of first-care doctors, a relatively high proportion of cancers in China, - there would be 4.3 million new cancer patients every year - and a limited access to healthcare, the need for smart AI systems to improve healthcare system arose as a priority for the government. This is even more important with the continuous aggravation of air quality in large towns, which causes a lot of diseases across the population: helping to lower the delay of diagnosis and to increase the value added by doctors is crucial.

AI use cases in healthcare are indeed numerous, and many of them fit pretty well in the image recognition framework which is the core of the Chinese AI knowledge. The most spectacular applications involve analysis of radiographies to detect cancers: lung, esophagus... With a large enough dataset, labeled by doctors, AI is allegedly capable of reaching a significantly better accuracy than average human doctors. However their technology doesn't completely replace doctors, who remain crucial to validate and confirm the results of the algorithm. Therefore AI is seen as an additional tool for the doctors. This solves the issue of accountability for the decision taken by the algorithm, since the doctor remains the final decision-maker. Tencent has deployed such a technology in several dedicated hospitals. Ping An, an insurance and banking company, also decided to specialize in AI technologies to provide such services.

AI for entertainment

Other use cases of AI and specifically computer vision concern are for entertainment. Many smartphone applications are powered with AI e.g. enhancement of people's beauty on pictures, adding fun, dynamic masks and filters to their faces... Chinese people seem to have adopted this applications more than Europeans. In the short term, this could actually be one of the biggest revenue sources for AI developments.

In fact, limiting these use cases to "fun-time" would be naive. These technologies fit pretty well for instance in the strategy of Alibaba to make shopping in real physical shops fun, and thus enhance the customer relationship, and boost consumption. As an example, we were shown a marketing video that described what happened during the previous Single's Day, Alibaba's "Black Friday" equivalent (US 42bn of sales in 2018) People started to play interactive games on their phone the day before to get vouchers.

PING AN

The old insurer turning into a tech company

Ping An is a Chinese conglomerate founded in 1988, which deals with insurance, banking and financial services. It has 1.8 million employees and 166 million clients (1 out of 9 citizens in China). Ping An Insurance is the world's largest and most valuable insurer: it is worth USD 217 billion as of January 2018. It has a net profit compound growth rate of 30%. Their position as a world leader comes from the size of the Chinese market, closed to foreign companies. Ping An remains mainly a Chinese actor that has not developed on the world market.

In 2019, Ping An announced a strategic shift and now introduces itself as a technology company rather than a finance company: they aim at becoming the world-leading integrated financial services and technology group, focusing on financial services (insurance, banking, investment) and healthcare/medical. Ping An hired thousands of tech employees and developed a digital health service named "Good Doctor" which has now a 150 million user base.

What do they use AI for?

Ping An is mostly using AI to improve their core processes. They have implemented so far quite diversified applications: automatic vehicle damage assessment and claim processing, underwriting through voice recognition, medical images analysis and risk management for investment services. Since they intend to convert their business and turn into a tech company that would compete with the internet giants, they are trying to develop their own AI knowledge and products, and build their own deep learning platform. However, they did not present to us groundbreaking new use cases.

Contrarily to Western financial companies, their size enables them to have significant fundamental research in AI (including some esoteric applications like automatic music composition). They claim to have more than one thousand PhDs working on AI.

Our opinion on their AI abilities

They pretended to be groundbreakers in AI during our visit, showing slides about potential applications of their technologies. Yet we were told during our visit in SenseTime that Ping An had asked them to provide the computer vision part.

Robot friends: Natural language processing and translation

Though Chinese companies work on Natural Language Processing (NLP), we did not have the opportunity to see other use cases than voice recognition to process calls with AI or to write messages on WeChat.

Such technologies are not mature yet and few use cases are already *in production*. Only Ping An told us they were implementing a solution to process clients' calls - they were monitoring the quality of their call centers employees. Unfortunately, we did not have the opportunity to visit iFlyTek, which was recognized as the leader in NLP by the Chinese government.

On the other hand, translation is a major use case for Chinese people, since there are seven major language groups in China, and few people can speak English. Thus several companies like Tencent provide tools to translate live conversations. Voice recognition is also becoming more efficient than typing using a Chinese keyboard.

Other applications

While communication about AI seems to solely focus on applications that either deal with consumers directly or are used by government, AI is also used for industrial purposes as a mean of automation and rationalization, by implementing predictive maintenance for instance. However, this kind of application wasn't the core of what we've seen during our learning expedition. Indeed, from what we understood the Chinese industry is not very advanced in terms of numerisation, so AI cannot be fed with relevant data at this stage. The only large scale example we discovered was the use of AI to optimize communication networks by Huawei, for traffic prediction and network management. Their target is to build autonomous driven network. Deployment has already been automatized, and Huawei is now working on making upstream tasks more autonomous: Service Provisioning, Operation and optimization of the network topologies, and even one day the planning and design of the infrastructure.

Autonomous driving is also an important use case of computer vision that appeared regularly throughout the whole journey. Most firms are trying to get to produce level 4 autonomous car, that is, a car able to basically drive alone but which needs the presence of a driver to take care of accident-prone situations. Among all instances, we can mention autonomous buses processed by a company called DeepBlue, that provides surveillance tools within those same buses. SenseTime is also providing computer vision solutions for cars to identify if the driver is sleeping or calling. We didn't have the opportunity to visit Baidu, which was recognized as the leader in autonomous driving by the government.

AI relies on the data industry

The new sweatshops: a dataset industry

China has built an industry for gathering and labeling datasets of all kind to fuel its AI applications. Many companies offer services to construct clean, massive and labelled datasets for AI purposes. Low labour cost has enabled companies to hire large teams and ensure high quality datasets, which is a key limiting factor in an AI development. "Quality data is the dream, dirty data is the truth", assessed Huawei head of medical R&D. On its side, Sensetime has more than nine hundred people dedicated to data cleaning and labeling. Nevertheless, several companies highlighted the difficulty of keeping people working on those repetitive tasks: people only accept these jobs for a while, they expect to get promotions to more interesting jobs in the end, and leave if they don't. This is why an important field of research which was presented to us by an academic of Shanghai Jiaotong University tries to get good results from machine learning using less data or data of poorer quality. Achieving these goals with an acceptable level of accuracy would be a considerable progress for the industry.

Furthermore, regulatory hurdles to the constitution of datasets are often limited. In the case of medical images, Tencent needs three months between the first request for images to a hospital and the effective transfer of the data. It will stop partnership with a hospital if it takes more than six months. By comparison, in France, accessing medical database often takes two to four years.

Thou shall trust your robot

It's all about trust, folks!

Trust is a foundational building block of human socio-economic dynamics, and is a necessary element to the success of any undertaking.

In traditional software applications, their behavior is dictated by explicit rules expressed in the code; in the case of AI agents, their behavior is based on knowledge that relies on data and evolves over time. The former approach is deterministic and predictable, the latter is non-deterministic and difficult to understand.

Therefore, humans lose the ability to understand why an algorithm is taking a decision. This is why ethics has been a hot topic for AI development in Europe and in the USA.

During our trip, we tried to learn more about these issues and to focus our questions on four topics: ethics, transparency, bias and interpretability. Our goal was to understand not only these topics, but also the way Chinese companies tackle them.

Ethics

While technology does not raise ethical issues by itself, the applications do. For instance, one of the companies we visited, DeepBlue admitted having a project of an armed robot which decides alone whether to shoot or not.

Europeans have focused on ethical issues regarding data privacy but AI can raise many other issues, such as bias, transparency or explainability of algorithms.

The definition of ethics varies depending on the society and its values, and from our visits, it appears that ethical issues (as understood in Europe) are not considered as a priority nowadays in China. The main goal of China is to become the world leader in AI, and to improve the standards of living of its population. Ethical considerations are at best secondary and might be discussed once the main goals are reached. AI is used at broad scale for many applications that would be unacceptable in Europe: merged integrated databases, large scale video surveillance, no transparency of the algorithm or the use of data for the citizen...

Transparency

Transparency is often thought of as a necessary path to reach acceptability. Open source is a form of transparency, and one that is often emphasized, with reason: as the saying goes, *code is law*, which means that if an algorithm is used to take decisions, the exact way it is written is as powerful as a law as it will apply to everyone.

However, is it the only possible one, and is it really the most important one? We do not think so.

First, AI algorithms are mostly based on the same paradigm – deep learning – but these are very hard, not to say impossible, to interpret. Thus having access to the code without the training data is useless as the labeled data set used to train the neural networks will determine the shape of the final algorithm. None of the companies we met would agree to share their training data with the public, thus open source for AI seems to be a lost cause.

Indeed, Chinese technology firms will always answer to accountability questions by showing the results themselves, and never the way they were obtained.

However, an AI algorithm is always part of a complex software chain. Having a black box inside the chain is not the end of all possibilities of transparency: being transparent about the use case as a whole is even more useful. A concrete example is face recognition: while there is no point in

sharing the code of face recognition technology, it is of high public interest to know what these technologies are used for. Explaining that face recognition for surveillance cameras is only used in case of a criminal event in the area does much more for the acceptability of the system than any open source algorithm could.

Bias

The question of bias is as old as statistics, and may be a major concern as the use of AI spreads. Biases are much easier to illustrate than to define: Google was recently accused of obvious algorithm bias when an AI machine trained on pictures of men snowboarding declared women could not use snowboard. Many examples can have more severe consequences: denying access to credit to people coming from a given background, making wrong diagnoses for a given population that wouldn't be represented in the training data set, etc.

In practice, little bias have been encountered in the AI applications already in production such as those in medical images analysis and those in facial recognition. As an example, representatives of DeepBlue talked about the video surveillance installed in public buses, supposed to detect harassment. They had then to agree that two people cuddling together in the bus might be interpreted as well as an harassment act. However, their attitude then clearly showed that they did not worry very much about such consideration, based on the fact that they never actually observed such problems.

SenseTime assured us that with facial recognition algorithms trained on 2bn pictures of Chinese citizens, only limited datasets (200k pictures) would be needed to make their algorithms as efficient for African citizens recognition as it is actually for Chinese citizens.

Companies in China do not see bias as an ethical problem but rather as a performance problem. If the performance on some subset of users is too poor, companies will ensure the algorithm is not the decision maker but rather an assistance for decision (in cancer detection, algorithm is supporting the doctor in its decision).

Interpretability

An algorithm is interpretable if one can explain the different reasons why it reached a certain decision, a certain result. For instance in the case of an algorithm deciding whether someone should be freed on bail, one could ask for the grounds on which they were deemed too likely to commit a crime again. A doctor may also like to know which pixel on a radiography was the most influential for the algorithm in the diagnosis - but he might also decide that if the algorithm has an average accuracy superior to humans he can just trust it.

Because AI looks like a "black box", which builds a model from data (the opposite case is a model designed by an expert), and often involves very complicated and nonlinear functions (neural networks are a good example of that), it often lacks interpretability.

Over the week, we repeatedly asked companies whether they cared about being able to explain and interpret the algorithms they use. Only one company gave a positive answer: Airbus, who stated that they need to be able to explain their algorithm, to avoid a black box feeling. In the case of predictive maintenance, their clients (EasyJet for instance) need to trust the predictions, and thus to trust the algorithm because of the cost of doing maintenance.

In all other cases, the companies did not display any concern or interest for the question. In some cases, it was even difficult for us to have our question understood. This shows how far interpretability is from their concerns. We have a few conjectures as to why interest was so low.

First, many developments are not mature enough to consider this question: in Zhejiang lab, researchers consider that their technology answering questions about a pictures is too recent and imperfect for them to consider challenges that go beyond performance.

In a second case, many applications have been deployed, where interpretability is not a business imperative. In Tencent, the automated diagnosis trying to detect cancers is always a second diagnostic, paired with a doctor diagnostic, who can then take on the explanation duties.

Finally, interpretability of algorithms may matter more in the case of an automatic decision, which should be contestable (in front of a court or not). The principle that the decision-maker should be accountable for its decision, and should be able to explain why they took a specific decision, is possibly more developed in the West than in China.

Conclusion

All in all, the capacity for marketing and fudge should not disqualify all reports on AI research: if only half or a fourth of the ambition is met, it will still be a big leap forward.

All firms we visited apparently work on similar projects: the topics of video surveillance, smart cities, autonomous vehicles and cancer recognition always came out. Either these topics are particularly promising in terms of expected growth, or the government gives incentives/orders to have all big companies working on them. It may also be that the companies are not very innovative in terms of new use cases, and mostly copy their US competitors, or that those business case are the only mature ones since the main applications of AI that are already related to a business case are image recognition and analysis and natural language processing. It was interesting to notice how broadly implemented the image recognition technology already is in China for payment, video surveillance with facial recognition, etc. However, we noticed some specialisations: SenseTime seems to be the reference for Computer Vision, AliCloud for Smart Cities...

Many Chinese companies are leaders in computer vision, especially for face recognition. Our rough estimate is that at least 80% of their AI work is based on such type of technologies. They apply most of their effort to face recognition, and image recognition in health care. The first goal is to answer the government's needs of security and healthcare system solutions. The second target is to get the money that can be made in healthcare and the positive reputation it brings.

Payment & ecosystems play a key role. Leapfrogging the credit card stage, Chinese people now pay with their smartphone (QR codes, face payment). Most payments are done through *AliPay* or *WeChat*, enabling Alibaba and Tencent to build huge ecosystems around their payment systems. Having access to the bulk of data about spending habits, which they collect from merchants, they can improve their services and buy startups.These two systems are rigidly separated, to the point that you can't even share an Alibaba link on WeChat.

Impressive showrooms are used to demonstrate the impact of technology on the society. Visits of showrooms were quite impressive featuring for instance at Tencent a James-Bond style organisation with disappearing walls and light tunnels.