Daiwa Capital Markets

China Internet

Primer: machine learning, deep learning and AI

- Powered by machine learning and deep learning, artificial intelligence looks set to be the next blockbuster technology
- We believe the large Internet companies have natural advantages over start-ups, as they have exclusive data and ample funding
- Reiterating Buy (1) on Alibaba and Tencent; maintaining Hold (3) on Baidu given uncertain medium-term search ad revenue outlook

Negative | Neutral | **Positive**

→ **Positive** (unchanged)

John Choi
(852) 2773 8730
john.choi@hk.daiwacm.com

Alex Liu
(852) 2848 4976
alex.liu@hk.daiwacm.com

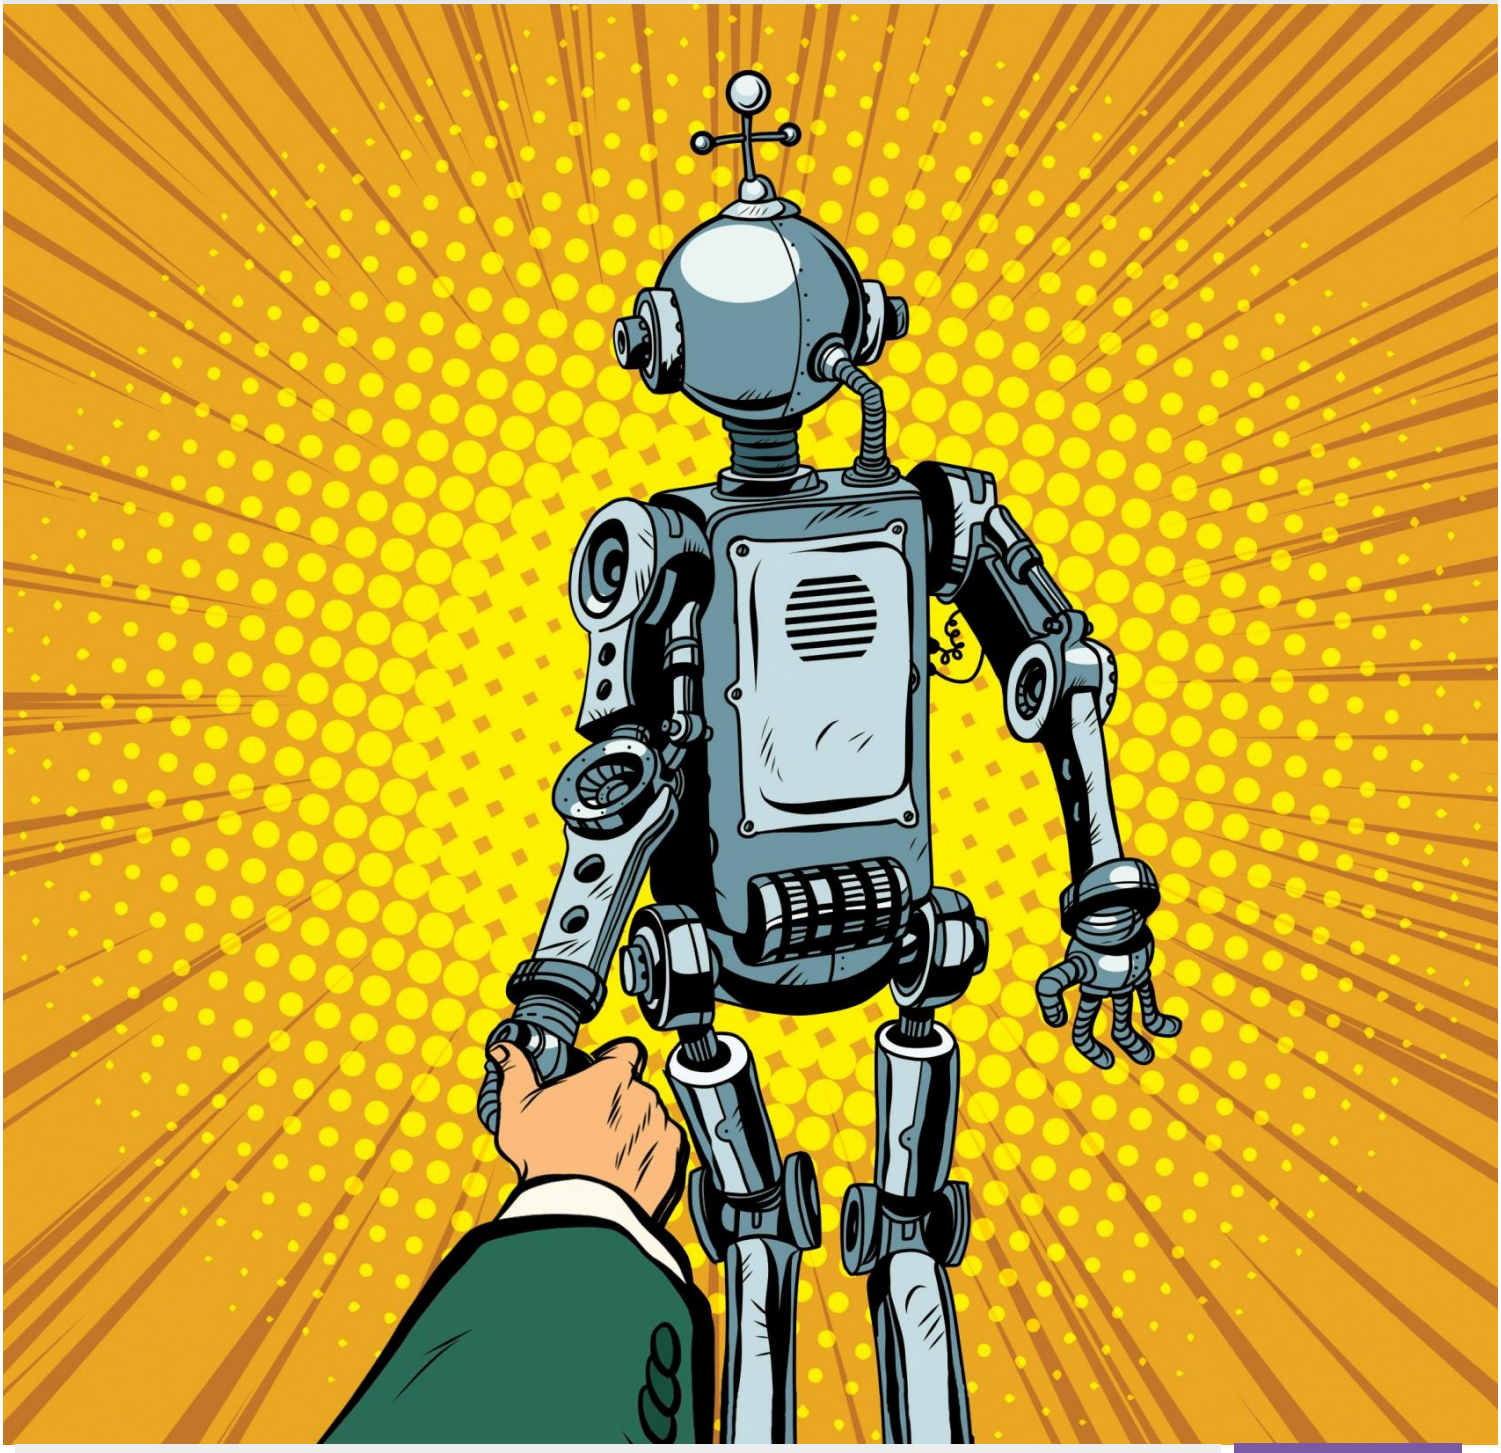

See important disclosures, including any required research certifications, beginning on page 79

Sector

Table of contents

Daiwa Capital Markets

China Internet

Primer: machine learning, deep learning and AI

- Powered by machine learning and deep learning, artificial intelligence looks set to be the next blockbuster technology
- We believe the large Internet companies have natural advantages over start-ups, as they have exclusive data and ample funding
- Reiterating Buy (1) on Alibaba and Tencent; maintaining Hold (3) on Baidu given uncertain medium-term search ad revenue outlook

Negative | Neutral | **Positive**

Positive (unchanged)

John Choi
(852) 2773 8730
john.choi@hk.daiwacm.com

Alex Liu
(852) 2848 4976
alex.liu@hk.daiwacm.com

What's new: We contend that, after the mobile Internet, artificial intelligence (AI), powered by the rapid evolution of machine learning and deep learning (or neural network learning), is the next revolutionary technology. As we believe that discussions of AI tend to be long on jargon/ short on substance, our aim with this report is to take a deep dive into machine learning and deep learning, covering the theory behind deep learning algorithms and the real-world applications for machine learning. Also, we include case studies showing how China's Internet giants are incorporating AI/machine learning into their core products.

What's the impact: **AI: it is here, almost unnoticed.** Many investors may not realise that machine learning and deep learning already serve as the guiding hand in the Internet services they use daily. From search engines to content recommendations, Internet companies are folding machine-learning technology into their services with the aim of improving information discovery and monetisation efficiency.

Proprietary data and disruptive business models: crucial to long-term success. We see proprietary data becoming the moat to defend against competition in the AI era, since it gives the owner exclusive insight into users' behaviour and hence an opportunity to provide a superior experience. We believe the proliferation of machine learning will give rise to new disruptive economic models. Compared with their global peers, the China Internet players have more comprehensive proprietary data to hand, which we think will allow them to make the most of the emerging business opportunities.

What we recommend: In our view, the Internet giants have natural capital and data advantages over start-ups in researching and deploying machine learning to their product lines. We see **Baidu** (BIDU US, USD236.41, Hold [3]) as the pioneer in AI investment and research in China. Although visibility on commercialisation of Baidu's autonomous driving project is limited, we expect the initiative to significantly expand Baidu's addressable market. Meanwhile, we contend that **Alibaba** (BABA US, USD178.97, Buy [1]) AI status is not fully appreciated, as it has been quietly folding machine learning into such products as the mobile Taobao app and Xiami music. Also, we think its cloud leadership will give it a cost advantage in AI training. Finally, **Tencent** (700 HK, HKD334.00, Buy [1]) looks to be raising its AI game, and its vast set of user data, allied with machine learning, should greatly enhance its ability to monetise online ads and digital payments, in our view. In this report, we revise our 12-month TPs for all 3 of China's Internet giants.

How we differ: We are more positive than the market on how machine learning will enhance the efficiency of Internet monetisation models.

Key stock calls

	New	Prev.
Alibaba Group (BABA US)		
Rating	**Buy**	Buy
Target	**200.00**	205.00
Upside	▲ **11.8%**	
Tencent Holdings (700 HK)		
Rating	**Buy**	Buy
Target	**370.00**	360.00
Upside	▲ **10.8%**	
Baidu (BIDU US)		
Rating	**Hold**	Hold
Target	**215.00**	195.00
Downside	▼ **9.1%**	

Source: Daiwa forecasts

Sector stocks: key indicators

Company Name	Stock code	Share Price	Rating New	Rating Prev.	Target price (local curr.) New	Target price (local curr.) Prev.	Target price (local curr.) % chg	EPS (local curr.) FY1 New	EPS (local curr.) FY1 Prev.	EPS (local curr.) FY1 % chg	EPS (local curr.) FY2 New	EPS (local curr.) FY2 Prev.	EPS (local curr.) FY2 % chg
Alibaba Group	BABA US	178.97	Buy	Buy	200.00	205.00	(2.4%)	31.822	32.899	(3.3%)	41.797	42.640	(2.0%)
Baidu	BIDU US	236.41	Hold	Hold	215.00	195.00	10.3%	45.569	43.853	3.9%	51.415	49.651	3.6%
Ctrip.com International	CTRP US	52.99	Outperform	Outperform	55.00	55.00	0.0%	5.644	5.644	0.0%	9.776	9.776	0.0%
JD.com	JD US	45.27	Buy	Buy	60.00	60.00	0.0%	3.676	3.676	0.0%	7.459	7.459	0.0%
NetEase	NTES US	270.27	Outperform	Outperform	330.00	330.00	0.0%	105.617	105.617	0.0%	123.663	123.663	0.0%
Tencent Holdings	700 HK	334.00	Buy	Buy	370.00	360.00	2.8%	6.651	6.478	2.7%	9.736	9.499	2.5%
Vipshop	VIPS US	9.89	Hold	Hold	10.00	10.00	0.0%	5.078	5.078	0.0%	5.819	5.819	0.0%

Source: Bloomberg, Daiwa forecasts

China Internet: comparable valuations

Ticker	Short Name	Ratings	Market price (local curr')	Market cap (USDm)	PER (x) FY1E	PER (x) FY2E	PER (x) FY3E	P/S (x) FY1E	P/S (x) FY2E	P/S (x) FY3E	EV/EBITDA (x) FY1E	EV/EBITDA (x) FY2E	EV/EBITDA (x) FY3E	ROE % FY1E
China Internet														
700 HK Equity	TENCENT*	Buy	334.00	402,011	42.0	28.7	22.8	11.7	7.6	5.7	23.4	19.4	14.5	31.1
BIDU US Equity	BAIDU INC-SP ADR*	Hold	236.41	82,815	33.9	30.0	21.7	6.6	5.3	4.4	26.8	21.6	14.3	15.3
NTES US Equity	NETEASE INC-ADR*	Outperform	270.27	35,603	16.7	14.3	12.3	4.6	3.5	2.9	11.5	9.0	7.2	32.1
CTRP US Equity	CTRIP.COM-ADR*	Outperform	52.99	27,785	61.3	35.4	23.5	7.1	5.5	4.7	34.7	22.0	14.3	5.0
WB US Equity	WEIBO CORP-ADR	Non-rated	106.10	23,198	66.9	42.1	29.9	17.9	14.8	10.2	55.2	34.2	23.7	24.9
SINA US Equity	SINA CORP	Non-rated	117.08	8,362	42.2	30.8	23.2	5.1	4.3	3.2	15.4	10.1	6.6	8.9
YY US Equity	YY INC-ADR	Non-rated	78.45	4,851	12.9	10.7	9.1	2.7	2.3	1.8	10.1	8.1	6.9	40.5
MOMO US Equity	MOMO INC-ADR	Non-rated	37.31	7,356	23.7	17.2	12.8	5.7	4.2	3.1	18.0	12.6	9.7	40.4
SFUN US Equity	FANG HOLDING-ADR	Non-rated	4.29	1,866	107.3	22.6	18.8	n.a.	3.8	3.4	50.5	14.5	11.1	(4.8)
ATHM US Equity	AUTOHOME INC-ADR	Non-rated	67.05	7,770	27.8	23.0	18.3	8.2	6.8	5.8	22.1	17.7	13.7	23.1
EDU US Equity	NEW ORIENTAL-ADR*	Buy	91.86	14,478	39.4	30.3	23.1	6.4	5.1	4.0	28.2	20.3	14.2	16.9
BITA US Equity	BITAUTO HOLD-ADR	Non-rated	40.71	2,879	31.7	20.1	12.5	2.2	1.2	1.0	15.1	9.8	7.2	(5.1)
WUBA US Equity	58.COM-ADR	Non-rated	67.59	9,820	64.4	38.0	26.5	6.2	5.5	4.5	33.3	21.7	15.9	1.0
TAL US Equity	TAL EDUCATIO-ADR*	Buy	32.92	2,712	11.6	7.6	5.1	1.6	1.0	0.7	5.7	2.1	0.0	33.7
SOHU US Equity	SOHU.COM INC	Non-rated	54.94	2,135	n.a.	n.a.	n.a.	1.0	1.0	0.9	17.8	6.3	4.3	(29.5)
BZUN US Equity	BAOZUN INC-ADR	Non-rated	37.35	1,992	59.0	33.6	22.9	0.5	0.4	0.3	36.9	21.9	13.7	8.5
Average					42.7	25.6	18.8	5.8	4.5	3.5	25.3	15.7	11.1	15.1
China E-commerce														
BABA US Equity	ALIBABA GRP-ADR*	Buy	178.97	455,121	36.7	28.0	21.8	13.2	9.2	7.0	24.9	17.5	12.4	24.4
JD US Equity	JD.COM INC-ADR*	Buy	45.27	64,220	80.4	39.6	21.9	1.2	0.9	0.8	63.1	27.9	15.2	11.3
VIPS US Equity	VIPSHOP HOLDINGS*	Hold	9.89	5,797	12.7	11.1	10.1	0.5	0.4	0.4	9.6	7.7	5.9	44.8
Average					43.3	26.2	17.9	5.0	3.5	2.7	32.5	17.7	10.6	26.9
Global Internet														
AAPL US Equity	APPLE INC	Outperform	159.65	824,629	17.7	14.6	14.1	3.2	3.2	2.7	9.5	8.1	8.0	36.0
GOOG US Equity	ALPHABET INC-C	Buy	935.09	652,841	24.3	20.2	16.9	6.8	6.2	5.3	13.1	11.0	9.4	14.0
FB US Equity	FACEBOOK INC-A	Buy	173.05	502,530	30.0	25.1	20.2	11.3	9.9	7.7	19.0	15.0	12.0	22.0
AMZN US Equity	AMAZON.COM INC	Outperform	999.60	480,187	92.8	62.7	44.1	2.6	2.3	1.8	26.8	19.9	15.1	9.7
PCLN US Equity	PRICELINE GROUP	Non-rated	1,871.17	91,798	25.0	21.7	18.8	6.9	6.4	5.5	17.5	15.1	12.8	22.6
EBAY US Equity	EBAY INC	Hold	38.27	40,961	18.9	17.1	15.5	4.2	4.1	3.8	12.7	11.8	11.2	82.7
TWTR US Equity	TWITTER INC	Non-rated	18.20	13,430	58.5	52.9	44.5	5.5	5.3	4.9	15.8	14.6	12.8	(9.6)
NFLX US Equity	NETFLIX INC	Hold	183.64	79,287	112.0	75.2	52.3	6.2	5.6	4.6	77.0	45.7	30.5	13.1
EXPE US Equity	EXPEDIA INC	Non-rated	145.66	22,087	28.8	21.8	17.9	2.0	1.9	1.7	12.5	10.4	8.9	7.8
TRIP US Equity	TRIPADVISOR INC	Non-rated	44.41	6,162	40.3	36.8	29.0	3.7	3.5	3.1	16.8	16.2	13.6	6.7
035420 KS Equity	NAVER CORP*	Buy	729,000.00	21,295	24.1	19.5	16.7	5.2	4.5	4.2	13.3	10.3	8.4	24.5
035720 KS Equity	KAKAO CORP*	Buy	127,500.00	7,659	49.7	31.4	28.6	4.4	3.8	3.7	23.1	15.3	13.8	5.0
4755 JP Equity	RAKUTEN INC	Outperform	1,243.00	16,208	22.8	20.5	17.1	1.9	0.0	0.0	2.1	1.9	1.7	8.2
4689 JP Equity	YAHOO JAPAN CORP	Outperform	507.00	26,088	21.2	18.9	17.9	3.1	3.3	3.1	10.5	9.3	8.7	15.7
Average					40.4	31.3	25.3	4.8	4.3	3.7	19.3	14.6	11.9	18.5

*Source: Bloomberg; *Daiwa forecasts; note: prices as of close on 13 September 2017*

Decoding artificial intelligence: a primer on machine learning and deep learning

"Now this is not the end. It is not even the beginning of the end. But it is, perhaps, the end of the beginning."

- Winston Churchill

"The future is already here — it's just not very evenly distributed."

- William Gibson

Executive summary

With the mobile Internet having gone through the rapid adoption phase, investors are keenly looking for the next revolutionary technology. We highlight **artificial intelligence (AI) as the "next" mobile Internet**, underpinned by the rapid development of machine learning (especially deep learning) in recent years. As the following chart illustrates, the big Internet companies in China and the US have been turning their attentions to AI and machine learning research. Indeed, many of these companies, including Baidu and Google, have signalled their strategic intention to become "AI-first companies" in the next 3-5 years.

The big Internet players have been turning their attentions to AI and machine learning in the past 2 years

AI: no. of mentions of "AI" or "machine learning" during earnings conference calls (1Q15-2Q17)

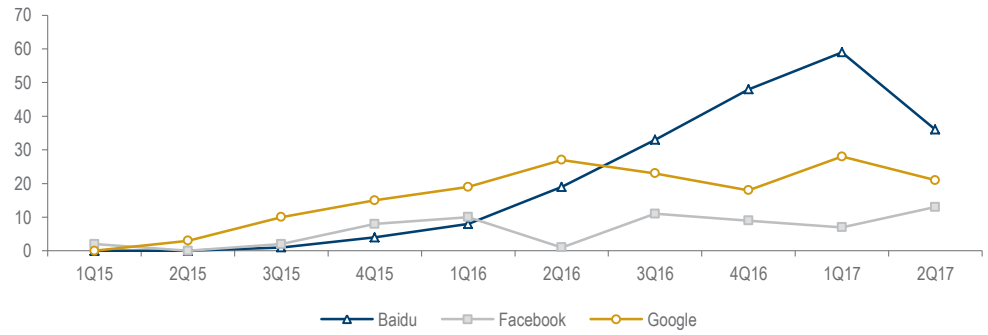

Source: Companies, Daiwa compiled

There has been much discussion of AI and machine learning in academic and media circles, but it seems to us that the substance has been lacking, with many reports and debates getting bogged down in the associated jargon.

As such, with this report we aim to provide investors with a deep dive into machine learning and deep learning, covering topics from the theoretical background of machine learning, the statistical theory behind deep learning, the rise in popularity of machine learning and deep learning, and the existing and future applications for machine learning and deep learning.

While there are use cases for AI across many industries, we limit our discussion in this report to the applications for machine learning by Internet companies.

We would like to thank Yishuo Liu for her valuable contribution to this report.

Theoretical backgrounds of AI, machine learning and deep learning

Topic 1: Defining machine learning and deep learning

Deep learning is only one branch of machine learning

- The term AI is widely used to refer to a scenario when machines acquire human-level intelligence
- We define machine learning as any process by which a machine improves its performance at certain tasks after accumulating a certain level of experience. Machine learning is an important branch of science that could help humans to build up generalised AI (encompassing such aspects as emotional, intelligence, creativity, moral reasoning and intuition) in the future
- Deep learning, or neural network learning, is only one branch of machine learning

Topic 2: How does deep learning work?

Deep learning is the process of using neural networks to find the best statistical function to explain complex patterns

- Mathematically speaking, machine learning is the process of finding the best math function to explain hidden patterns in our world
- Neural networks are capable of approximating any relationship/pattern, no matter how complicated it may be
- Neural networks consist of individual neurons which can conduct basic logical calculations on their own

Topic 3: Why has deep learning come to the fore?

Data and computing power are major drivers of machine learning development

- The dramatic growth of digital data and computing power in recent years has fuelled the machine learning and deep learning boom
- We believe the big Internet companies (global and China players alike) have natural advantages in attracting AI talent from the world of academia

Topic 4: Pushback on adoption — AI is a black box

The decision making process behind AI is a mystery to many

- Despite their extraordinary performance at learning, neural networks are complex systems and, in many respects, something of a black box
- The fact that it is not readily “understandable” means that some industries are reluctant to adopt deep learning in their decision-making processes

Application of machine learning and deep learning

Search and recommendations

Machine learning is already improving the user experience provided by search and recommendation engines

- We expect machine learning to greatly improve the user experience of 2 existing Internet products: search and recommendations
- Thanks to machine learning, we expect search to become more personalised and include more information (visual search)
- We expect recommendations to become more personalised and more creative (helping users to expand their set of interests)
- Also, we believe machine learning has the potential to boost the monetisation efficiency (more accurate click-through rate [CTR] prediction) of digital ad operations

Speech and visual recognition

Speech and visual recognition are giving rise to disruptive business models

- Cameras and microphones are replacing keyboard on mobile devices; voice input and visual input are reshaping the user interactive model, with the help of machine learning
- Deep learning has significantly enhanced the accuracy of speech recognition, visual recognition and translation
- We are upbeat on the adoption of speech and visual recognition in existing Internet products (visual search, voice search) and other industries (AR, fintech, medical, autonomous driving)

How big will AI be?

We view AI as a general technology, not a specific market

We are regularly asked by investors: "How big is the market for AI?" We view AI as a source of general technology that has the potential to enhance the efficiency of many Internet business models, including online ads, online games and e-commerce. However, asking about the size of the AI market is akin to asking about the size of the "Internet market". When assessing the potential size of markets, we prefer to focus on specific products and segments that could be enhanced by AI.

In August 2017, the Chinese government announced its national development plan, which includes the following targets:

- Core AI industry output to reach CNY100bn by 2020
- Core AI industry output to reach CNY400bn by 2025
- Core AI industry output to reach CNY1000bn by 2030

While there is no formal definition of the core AI industry, these targets give investors a top-down view of the potential size of the AI market in China, at least as far as the regulators are concerned.

Who will win the AI competition?

Looking for proprietary data and disruptive business models

With all the hype surrounding AI, and plenty of companies playing up their credentials in this field, we believe that the matrix and investment framework devised by Rayfe Asaoka of Canaan Partners is a useful tool in picking the winners from the AI crowd. His conclusion revolves around 2 key determinants of the value of an AI business: if a company has both a disruptive business model and proprietary data/algorithms, it is likely to be a long-term winner in the AI era.

Proprietary data serves as a moat against competition in the AI era

AI: investment framework

Proprietary data/algorithm?	Disruptive business model? No	Disruptive business model? Yes
Yes	**Short-term win** — Will face difficulties in rising above the noise in competitive environment	**Long-term win** — Differentiated business model and strong competitive moat
No	**Lose**	**Short-term win** — Will face competition from other companies with stronger datasets/algorithms

Source: Canaan Partners

AI is revolutionary technology, but it needs both a use case and a business model

We believe this framework sheds light on the competitive dynamics at play in the AI era. If, as we expect, AI proves to be the next universal technology, the technology itself won't be a long-term competitive edge, but rather something that every player possesses. Instead, the competitive dynamics will rest on how effectively each company uses AI to provide incremental value to users, ie, how valuable the underlying business model is.

In 5 years, every company could be an AI company

Frank Chen, partner of VC firm Andreessen Horowitz, offered a similar view in his blog posting entitled "*In a few years, no investors are going to be looking for AI start-ups":*

> *"…in a small handful of years,* ***software without AI will be unthinkable****. Investors will stop looking for AI-powered start-ups in exactly the same way they don't look for database-inside or cloud-native or mobile-first start-ups anymore. All those things are just assumed."*

AI competition comes down to competition on data and business models

Applying the above framework to the global Internet sector, we see the established large Internet players as having overwhelming advantages over start-ups in terms of the scale and quality of data. Hence, we conclude that the big Internet players are likely to win the AI war, while start-ups focusing on specific verticals (to gain exclusive vertical data and create customised business models) are likely to have a better chance of thriving than less-focused new entrants.

Final question: is AI a winner-takes-all market?

It might be a bit early to ask this question, much less answer it, given the industry is still in its infancy. But it is worth asking whether there exists a virtuous cycle for AI that will lead to one single company dominating the market.

On 30 August, Amazon and Microsoft announced plans to connect their machine-learning-backed personal assistants (respectively Alexa and Cortana) to each other, enabling each company to access the other's unique data set.

> *"The world is big and so multifaceted.* ***There are going to be multiple successful intelligent agents****, each with access to different sets of data and with different specialised skill areas. Together, their strengths will complement each other and provide customers with a richer and even more helpful experience."*
>
> *- Jeff Bezos, Founder and CEO of Amazon*

Jeff Bezos, widely regarded as a visionary, seems to imply that the end-game for AI may not be winner takes all, such that a handful of companies develop AI technology that is good enough to compete. **After all, when the market is growing at hyper speed, the first-mover advantage may not last very long.**

Implications for China Internet companies

We believe the Chinese Internet companies have distinctive advantages in the AI field, since they have access to: 1) ample government support for research funding, 2) deep pools of AI and engineering talent, 3) private-sector investment (both Internet and hardware companies), and 4) large amounts of digital data and use cases for training and deploying machine-learning techniques to existing products. Nevertheless, we note there is a significant gap between the R&D spending of global Internet players and their Chinese counterparts (charts below).

AI: no. of cited journals related to "deep learning" or "AI", categorised by authors' country of origin

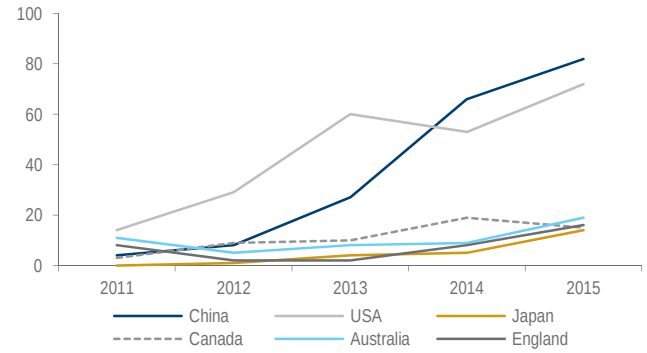

Source: National Science and Technology Council

AI: comparison of R&D expense for major Internet companies (2014-2016)

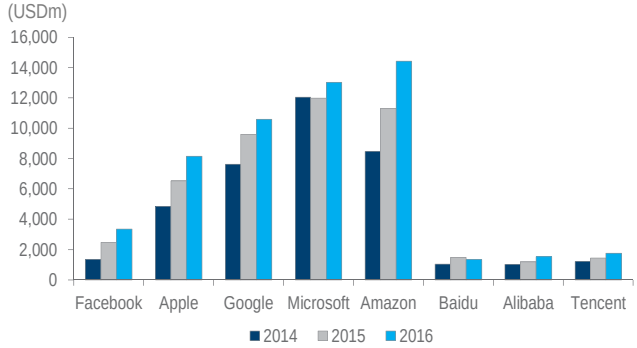

Source: Companies, Daiwa compiled
Note: For Alibaba, we use full year expense for fiscal year ending March 31; for Microsoft, we use full year expense for fiscal year ending June 30; Amazon R&D expense include content spending

Baidu (BIDU US, USD236.41, Hold [3])

Global AI research pioneer; new TP of USD215

- We consider Baidu to be the pioneer in AI research in China, and its management team appears determined to invest in this field.
- We anticipate there being a large addressable market for Baidu's autonomous driving business, though we believe it is unlikely to bring in significant revenue in the next 3 years. Meanwhile, its growth in its core search revenue is likely to continue decelerating in the medium term.
- We believe Baidu's investment in machine learning will improve the monetisation efficiency of its search ads and in-feed ads, despite fierce competition.
- Our new 12-month TP of USD215 (from USD195) is based on SOTP (see page 60). The key upside risk to our Hold (3) rating: lower marketing expenditure; key downside risk: slower search growth.

Alibaba (BABA US, USD178.97, Buy [1])

Hidden champion of AI research in China; new TP of USD200

- We believe Alibaba has a natural edge in AI research that has yet to be fully appreciated, backed by its consistent R&D investment and the vast exclusive consumer shopping data it has on hand.
- We see multiple use cases for AI and machine learning in Alibaba's existing product range (such as search and recommendation), and believe the company's overall ad monetisation efficiency will see a spike as the benefits of machine-learning adoption manifest themselves.
- We believe Alibaba's No. 1 position in China's infrastructure as a service (IaaS) cloud market will give it lower machine-learning training costs compared with start-ups.
- Our new 12-month TP of USD200 (previously USD205) for Alibaba is based on an unchanged 37x PER applied to the average of our new FY18-19E non-GAAP EPS forecasts. The key downside risk: a further slowdown in the Taobao marketplace.

Tencent (700 HK, HKD334.00, Buy [1])

Ready to invest in AI; abundant use cases for existing products to apply AI; new TP of HKD370

- Despite being relatively quiet on AI investment to date, we believe Tencent's management is increasingly keen to venture into AI research and application development.
- Similar to Alibaba, we believe Tencent has high-quality user data to hand that it is only now starting to monetise. Machine learning could greatly enhance the company's monetisation potential from online ads and digital payments, in our view.
- Our new 12-month TP of HKD370 is based on a 39x PER applied to 2017-18E non-GAAP EPS. Key downside risk: higher-than-expected marketing and content expenses.

AI: first forget the jargon

In this section, we define AI, machine learning and deep learning

Key takeaways:

- The term AI is widely used to refer to a scenario when machines acquire human-level intelligence
- We define machine learning as any process by which a machine improves its performance at certain tasks after accumulating a certain level of experience. Machine learning is an important branch of science that could help humans to build up generalised AI (encompassing such aspects as emotional, intelligence, creativity, moral reasoning and intuition) in the future
- Deep learning, or Neural Network Learning, is only one branch of machine learning

What is artificial intelligence?

Nowadays, AI is taken to mean machine/computer programs with "intelligence levels" similar to those of human beings. Discussions of AI tend to centre on a scenario in which machine/computer programs are armed with:

- **Reason** or the ability to answer problems
- **Knowledge representation**: understand the context of the world and objects in it
- **The ability to plan and navigate**: get from A to B by making optimal routing decisions
- **The ability to process natural language**: understand the context of the language and provide an appropriate response
- **The ability to perceive**: view the world through all kinds of senses (sight, sound, etc)

Many observers believe that the ultimate goal of AI is to build machines/computer programs with human-level generalised intelligence (encompassing such aspects as emotional, intelligence, creativity, moral reasoning and intuition). Think C3PO of Star Wars fame or, more ominously, Hal from 2001: A Space Odyssey.

Right now, however, we are far from the stage of generalised AI. Instead of focusing on the sci-fi aspects of AI, we focus on machine learning, especially deep learning, and discuss how these techniques are being applied to Internet business models today.

Defining learning

Herbert Simon, the renowned American psychologist and computer scientist, offered the following definition of learning:

> *"Learning is any process by which a system improves performance from experience."*

Learning is all about improving performance

The keywords here are process (how to learn?), system (who learns?), performance (learn what and how to measure?) and experience (learn from what?). This definition is interesting in that it is not derived from a social-science perspective of learning, but rather provides a rule to validate learning. In addition, the definition pinpoints performance improvement as essential to the subject.

Under Herbert Simon's definition:

- If a system improves its performance on certain tasks after adopting certain processes, these processes should be classified as learning processes.
- If a process does not result in improved performance, it cannot be considered to be a learning process.

Defining machine learning

Based on the above definition of learning, machine learning would be:

> Any process by which a **machine** improves its performance from experience.

Tom Mitchell, the American computer scientist who wrote the renowned text book *Machine learning*, defined machine learning in a more specific way:

> A computer program is said to learn from experience E with respect to some class of tasks T and performance measure P if its performance at tasks in T, as measured by P, improves with experience E.

For example, AlphaGo, the computer program developed by Google to play the complex Chinese board game Go, involves machine learning in the following way:

- **Task T**: play Go (with itself or others)
- **Performance P**: the winning rate
- **Experience E**: previous games (with itself or others)

Hence, if the win rate of AlphaGo improves after the system has contested multiple rounds of Go with other players or itself, we can say that AlphaGo has been learning from the games it has played.

More on machine learning

Currently, depending on "whether or not the input/output data is provided with specific labels", there are 3 major categories of machine learning methods:

- **Supervised learning**: A computer program is fed with example inputs and their desired outputs plus specific labels. The goal is for computer to "learn" a general rule that maps inputs to outputs that could be applicable to other inputs.

Example: a set of animal pictures is given to the computer, ***with each picture being labelled (cat, bird, dog, etc).*** *After certain learning processes, when given a new animal picture (no label), the computer should be able to identify what kind of animal is in the picture.*

- **Unsupervised learning**: No labels are given to the computer program, leaving it on its own to find structure in its input.

Example: a set of animal pictures is given to the computer, ***with each picture being unlabelled.*** *After certain learning processes, the computer should be able to classify pictures featuring the same animal into a single cluster.*

- **Hybrid learning:** a combination of supervised and unsupervised learning

While supervised and unsupervised learning each have their pros and cons, the main focus today for machine learning applications is supervised learning, where people expect computer programs to learn from existing data sets and perform logical tests on new datasets (eg, classification problems).

Defining deep learning

Deep learning = multi-layer neural network learning

Deep learning is only one approach to machine learning

Finally, deep learning is another way for machines to learn, ie, one of the many approaches involved in machine learning.

AI: relationship between deep learning and machine learning

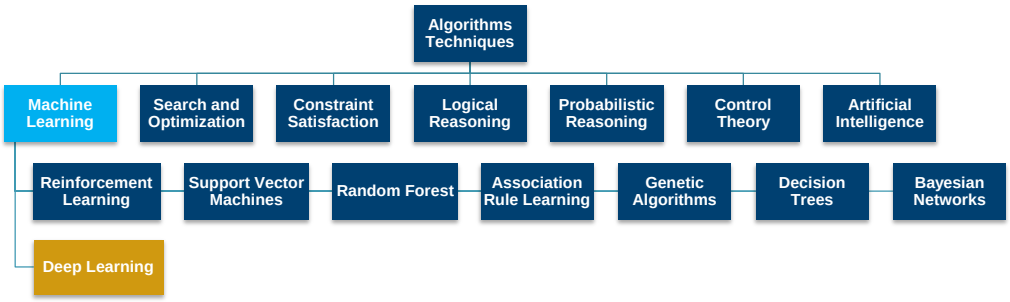

Source: a16z, Daiwa compiled

Deep learning is conceptually equivalent to neural network learning. The word "deep" here refers to the multiple layers of neurons (a concept referred as "hidden layer", which we explain in more detail in the next section) that are used in a neural network learning architecture.

Deep learning is usually classified as a type of supervised learning.

What exactly is neural network learning?

In short, Neural Network Learning is a technique allowing a computer program to learn from data. **The learning process is loosely based on how the human brain works.** First, a collection of software "neurons" is created and connected together, allowing the neurons to send messages to each other. Next, the network is asked to solve a problem, which it attempts to do over and over, each time strengthening the connections that lead to success and diminishing those that lead to failure.

A bit of history

Through deep learning, computers came up with the concept of "cat"

Deep learning started to draw attention in academic circles in the late-1980s, but it took off on earnest only in the late-2000s. Andrew Ng, formerly an associate professor at Stanford University who went on to become Chief Scientist of Baidu from 2014 to 2017, was working as a researcher at the Google X Lab in 2010-14. His team, led by Jeff Dean, conducted an experiment where they randomly picked 10m thumbnails of Youtube videos and fed the thumbnails into a multi-layer neural network powered by 1,000 computers with 16,000 processing cores. After being presented with 20,000 thumbnails, the neural network started recognising patterns in pictures and learned to recognise a "cat".

Why was the experiment a breakthrough? During the entire process, there was no human input or expert guidance to let the computer know "what is a cat". As Jeff Dean from Google put it:

> *"We never told it during the training, 'This is a cat', it (the neural network) basically invented the concept of a cat."*

The ability of neural networks to identify complex concepts without human input has attracted many researchers, and funding, to the field. Indeed, deep learning is arguably now one of the most researched approaches to machine learning.

We delve into the theory behind deep learning on page 13 and look at why it has suddenly come to public prominence on page 18.

Deep learning: a deep dive

In this section, we focus on the theoretical background of deep learning

Key takeaways:

- Mathematically speaking, machine learning is the process of finding the best math function to explain hidden patterns in our world
- Neural networks are capable of approximating any relationship/pattern, no matter how complicated it may be
- Neural networks consist of individual neurons which can conduct basic logical calculations on their own

The following discussion includes a technical introduction to deep learning, which some readers may find useful in understanding how neural networks work. Readers not interested in the underlying technology should skip to the next section. In the scope of this report, we use the term deep learning (DL) as equivalent to Multi-layer Neural Network Learning (or simply Neural Network Learning).

A mathematical interpretation of machine learning

Let's take a look at the mathematical concepts behind machine learning. The goal is for computer programs to find hidden patterns in large amounts of data. From a mathematical perspective, we can always find a function to represent such relationships, regardless of how complicated the function would be. **In more abstract terms, machine learning is the process of finding the closest perfect function, f(x), to generate the expected output Y for any given x**. For example:

- If we input x (a picture), f(x) could generate an output Y (telling us whether the picture contains a cat or a dog)
- If we input x (a moving car), f(x) could generate an output Y (telling us which direction the car is moving in and at what speed)
- If we input x (an English sentence), f(x) could generate an output Y (a Chinese translation)

Deep learning is essentially one big maths problem

AI: illustration of key elements in machine learning

Deep learning is the use of neural networks to find the best f(x) to explain the pattern of f(x) → y

Examples:

$f($  $) \rightarrow$ "Hello" *(speech recognition)*

$f($  $) \rightarrow$ "cat" *(visual recognition)*

$f($  $) \rightarrow$ *stop to car (autonomous driving)*

Source: Alibaba Cloud Forum, Daiwa compiled

As a statistics problem, machine learning can be broken down into 3 sub-problems:

1. **Model construction:** How to find a structure to build functions that explain the relationships in specific data sets?

2. **Measurement:** What criteria should be used to measure the performance of these functions?
3. **Optimisation:** How to quickly find the best function f(x)?

Neural Network Learning: why does it work?

Mathematically, a neural network is a universal approximator

Neural networks are a power tool for regression

As discussed, Neural Network Learning is a learning method which mimics the way a human brain works. Among all the machine-learning theories, what is it that makes neural networks stand out as the major direction for machine learning? And what are the limitations of Neural Network Learning?

A 1989-published paper written by Kurt Hornik, Maxwell Stinchcombe and Halberte White proves that:

> *"…**standard multilayer feedforward networks are capable of approximating any measurable function to any desired degree of accuracy, in a very specific and satisfying sense.** We have thus established that such "mapping" networks are universal approximators. This implies that any lack of success in applications must arise from inadequate learning, insufficient numbers of hidden units or the lack of a deterministic relationship between input and target."*

The paper argues that neural networks are universal approximators, meaning that in theory they are capable of approximating the results of any function. In plain English, **for any deterministic relationship, one can always find a neural network to represent the relationship provided there is sufficient data and learning.** This conclusion serves as the foundation for using neural networks as a framework to building functions for machine learning.

Neural Network Learning 101: building a perceptron

In this section we briefly introduce the most basic neural network, a *perceptron*, the mother of all neural networks.

> Definition:
> A perceptron is an artificial neuron that takes in binary inputs, x_1, x_2, …, x_n, runs calculations based on certain rules, and produces a single binary output y.

A perceptron is built around a concept that applies to human brain neurons. We now know that each neuron is capable of accepting all types of input signals (in the form of an electrical impulse) from other connected neurons, and output a binary result: activating or not activating. Such biological behaviour can be simplified into a mathematical rule for perception, as shown below:

> Rules:

$$\textbf{Human nueron: output} = \begin{cases} \text{not activating}, & \text{if certain inputs} < \text{threshold} \\ \text{activating}, & \text{if certain inputs} \geq \text{threshold} \end{cases}$$

$$\textbf{Perceptron: output y} = \begin{cases} 0, & \text{if } \sum_n w_n x_n < \text{threshold b} \\ 1, & \text{if } \sum_n w_n x_n \geq \text{threshold b} \end{cases}$$

Here we introduce 2 new variables: weight w and threshold b. Weight w is the number value that represents the level of importance of individual input x. Threshold b is a number value that decides the output and can be considered as a constant for each individual neuron.

A perceptron is a decision-making machine

Connected perceptrons make neural networks "smart"

As technical as all this may seem, a perceptron could be viewed simply as a decision-making machine. Each input could be considered to be one piece of real-world evidence when making a specific decision, and each weight value represents the importance of specific evidence in making a specific decision. The threshold value b represents how easy it is to make certain decision.

A perceptron takes in all the evidence, plus the level of importance of such evidence, weights the evidence by its respective importance level, compares the results against certain threshold values, and finally comes to a decision.

Example:
Assume two inputs, x_1 and x_2. We assume the input value could only be 1 or 0 for x_1 and x_2. Then we build a perceptron with weights of $w_1 = 1$, $w_2 = 2$, threshold b = 3.

- When $x_1 = 1$, $x_2 = 0$, weighted sum = 1*1+0*2 = 1 < 3, output = 0
- When $x_1 = 0$, $x_2 = 1$, weighted sum = 0*1+1*2 = 2 < 3, output = 0
- When $x_1 = 0$, $x_2 = 0$, weighted sum = 0*1+0*2 = 0 < 3, output = 0
- When $x_1 = 1$, $x_2 = 1$, weighted sum = 1*1+1*2 = 3 = 3, output = 1

AI: illustration of a perception calculation

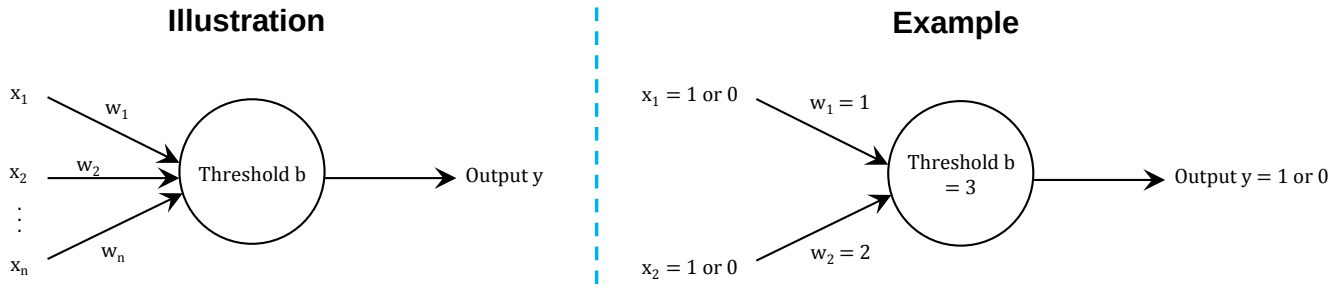

Source: Michael Nielson, Daiwa

In this example, we construct a perceptron where the output equals 1 only when both inputs equal 1. Interestingly, this perceptron effectively replicates the most elementary logical function, AND (an AND gate is a basic digital logic gate that outputs 0 except when all inputs are 1). It is also easy to construct a perceptron to replicate other logical functions such as NAND or OR.

The architecture of neural networks

A perceptron is a simple neural network. But the neural networks that we are familiar with tend to be comprised of multiple perceptrons connected to each other, as the chart below shows. Some of the key elements of multi-layer neural networks are:

1. **Input layer**: layer containing input neurons that feed in input data
2. **Hidden layers**: layers that are neither input nor output, but part of the calculation
3. **Output layer**: layer containing output neurons that give calculation results

Example: number recognition

- **Data**: a 28*28 pixel image, containing a number ranging from 0 to 9.
- **Input layer**: As the picture contains 28*28=784 pixels, the neural network input layer needs 784 neurons, with each neuron being mapped to a unique pixel in the picture. Let's simply the problem by assuming the picture is in black and white, meaning that each pixel is either black or white. We can then define black as 1 and white as 0.

- **Hidden layers**: in this case we build a hidden layer with only 15 neurons. While there is no specific rule on the number of hidden layers and hidden neurons, calculation speeds tend to decline with the addition of more hidden layers and neurons.
- **Output layer**: the output layer consists of 10 output neurons. Each of these output neurons represents a number value that we want to classify. If the output falls into the neuron representing 0, it means the neural network has recognised that the input picture contains the number 0.

AI: sample neural network for illustration purposes

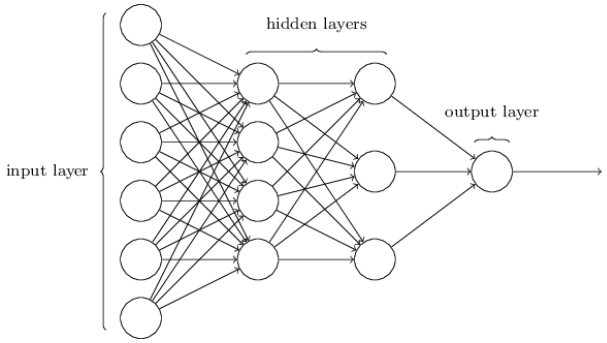

Source: Michael Nielson

AI: sample neural network for picture identification task

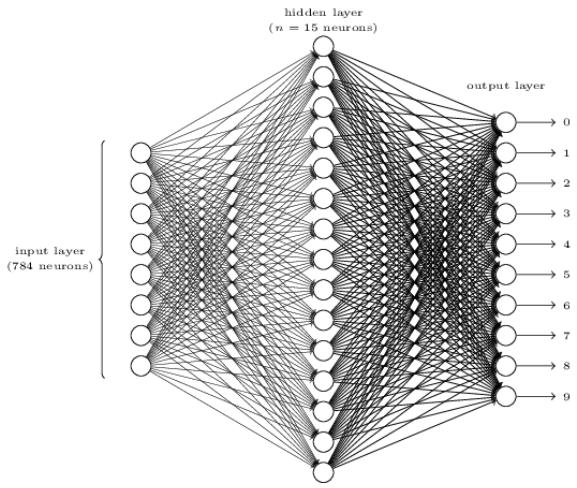

Source: Michael Nielson

The learning process underpinning neural networks

In previous sections, we looked at how perceptrons work and the architecture of a typical multi-layer neural network. In this section, we focus on the last part of our deep dive into neural networks: how do neural networks learn?

Recall that on page 10 we defined "learning" in the following way: if a process results in a computer program's performance improving under certain measurement, we can define the process as a learning process. In order to understand the learning process, we must first consider how to measure the performance of neural networks.

Furthermore, if we consider neural networks to be mathematical functions f(x) that can approximate universal relationships between input x and output y, then the learning process become a simple optimisation problem:

Optimise f(x) to best fit the relationship between input *x* and output *y*

Ie, finding the best weight *w* and threshold value *b* to construct a f(x) that best fits the relationship between input *x* and output *y*

Computer programs learn through constant optimisation

There are many statistical approaches to optimisation. Conventionally, for optimisation problems, we seek to minimise the loss function (the function representing the deviation between the estimated output and the actual output) for specific function f(x). We won't drill into the calculation process for optimisation here, but typically people use Minimum Square Error to build the loss function and adopt back propagation and gradient descent algorithms to make programs dynamically find the best function f(x) to minimise the loss function. Through this optimisation process, the neural network is effectively "learning".

Extra reading: Neural Network Learning 201 - Building a sigmoid neuron

Below we briefly introduce sigmoid neurons, a basic type of neuron similar to a perceptron. A sigmoid neuron is built upon perceptrons with slight modifications, but it has much greater practical usage value.

What is a sigmoid neuron?

First, a recap of a perceptron:

$$\textbf{Perceptron}: \text{output y} = \begin{cases} 0, & \text{if } \sum_n w_n x_n < \text{ threshold } b \\ 1, & \text{if } \sum_n w_n x_n \geq \text{threshold } b \end{cases}$$

Change form:

$$\textbf{Perceptron}: \text{output y} = \begin{cases} 0, & \text{if } \sum_n w_n x_n - b < 0 \\ 1, & \text{if } \sum_n w_n x_n - b \geq 0 \end{cases}$$

We define sigmoid neurons to have the following rules:

$$\textbf{Sigmoid function}: \text{output y} = \sigma(z) = \frac{1}{1+e^{-z}}$$

$$\textbf{Sigmoid neuron}: \text{output y} = f(x) = \frac{1}{1+e^{-(\sum_n w_n x_n - b)}}$$

The derivative of sigmoid function:

$$\text{output y}' = \sigma'(z) = \frac{e^{-z}}{(1+e^{-z})^2} = \frac{1}{1+e^{-z}} \times \frac{e^{-z}}{1+e^{-z}}$$
$$= \frac{1}{1+e^{-z}} \times \left(1 - \frac{1}{1+e^{-z}}\right) = \sigma(z)(1-\sigma(z))$$

Why build sigmoid neurons?

The reason we modify a perceptron to a sigmoid neuron using the sigmoid function (an s-shaped function, with output falling in a range of 0-1) as an activation function are:

1. We need to have an activation function, whereby a slight change in input x will lead to only a slight change in output y (sigmoid function fits this criteria).
2. We need an activation function that represents output in the form of probability (a sigmoid function output ranges from 0-1).
3. The partial derivative of a sigmoid function is easy to calculate.

Given the mathematical characteristics of a sigmoid function, there are many use cases in in neural network activation function construction. However, sigmoid functions have serious problems when conducting optimisation (its derivative becomes close to zero when the input becomes large, leading to a less effective learning process). Right now, most computer sciences use other functions (with similar characteristics to a sigmoid function) as activation functions when training neural networks.

Deep learning: why now?

In this section, we look at the key drivers of the rapid growth of deep learning

Key takeaways:

- The dramatic growth of digital data and computing power in recent years has fuelled the machine learning and deep learning boom
- We believe the big Internet companies (global and China players alike) have natural advantages in attracting AI talent from the world of academia

Much more data, much more computing power

The concept of deep learning is not new, with mainstream algorithms having been under constant development for almost 30 years. What explains its sudden rise to prominence in terms of public attention?

To use a simple analogy, if we consider machine learning as the process of finding the best way to grow crops, there are effectively 4 key elements involved in the process: 1) data (the quality of the seed), 2) computing power (soil), 3) the algorithm (farming toolkit), and 4) talent and investment (previous farming experience). The rapid growth of all 4 drivers has contributed to the boom in machine learning (especially deep learning) in recent years.

Data, computing power, algorithm and talent determine the output quality of machine learning

AI: illustration of key elements for machine learning

Computing power = "Soil"

Data = "Seed"

Algorithm = "Farming toolkit"

Talent = "Previous experience"

Growing crops

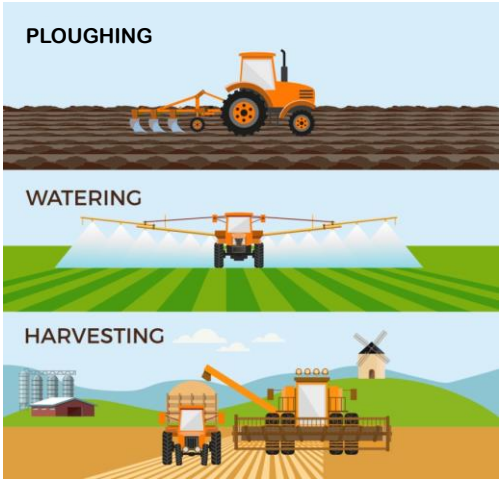

Machine Learning

Step 1:
Preparing the computation infrastructure ("the soil"), raw training data ("the seed") and algorithm ("farming toolkit")

Step 2:
Start training of machine learning algorithms, leveraging research talent ("previous experience"), to perform certain tasks by feeding in training data

Step 3:
After training, the algorithm canbe used to perform certain tasks if given new data sets

Source: Daiwa

More data

Simply put, vast amounts of data have been accumulated in the past 3 decades, thanks to the proliferation of the Internet and mobile devices. For instance, it was hard for researchers in the 1970-80s to gather a decent number of digitalised images to train neural networks (if you don't have a digital photo of a cat, how can you expect a computer to recognise a cat?).

Not only do we have more data, the quality and diversity of the data is getting better day by day. For example, every leading Internet company maintains a highly structured database containing multi-dimensional time-series feature data regarding individual user's online behaviour (device, use time, content preference, etc). Technology also allows us to digitalise physical objects and the offline experience (location, shopping behaviour in offline

store, etc). The large amounts of available data have significantly enhanced the learning process of neural networks.

More computing power

Moore's Law ,which holds that the performance of computing chips will double every 18 months, means a great deal to AI research. The rapid development of the semiconductor industry over the past 3 decades has dramatically changed the concept of computing power. Computing power has not only seen exponential growth; we can now actually deploy it at a fraction of the cost incurred 20 years ago.

AI: size of data created globally

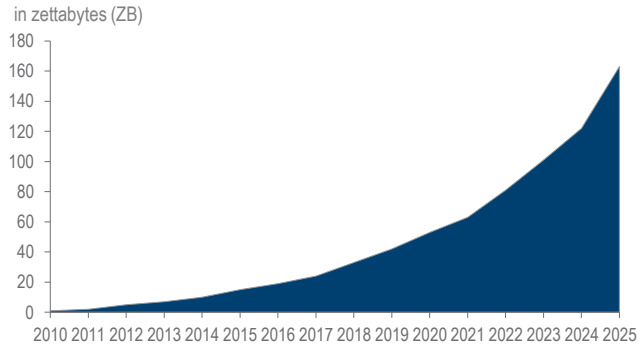

Source: IDC
Notes: 1 zettabyte (ZB) = 1 trillion gigabytes (GB)

AI: average unit price of computing power

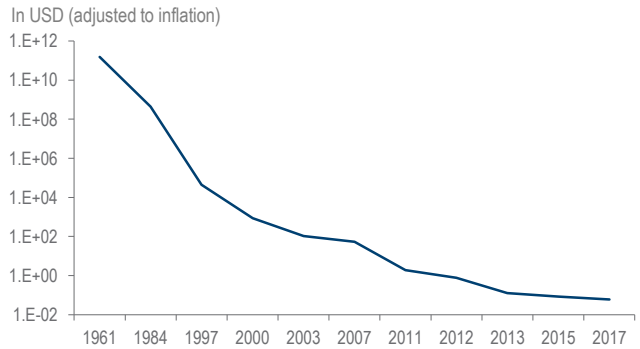

Source: The Aggregate
Notes: Unit cost of computing power is represented by the cost of one GFLOP; 1 GFLOP = 1 billion FLOPS; FLOPS = floating point operations per second, a common measurement of computing power

Better algorithms

Although the basic framework of neural networks was developed more than 30 years ago, computer scientists have continued to research ways to enhance the efficiency of neural network learning by developing convolutional neural networks using a feedforward approach. This is because better algorithms help to build more complex neural networks and solve more difficult tasks.

Take the case of the sigmoid function we discussed on page 17. In theory, the sigmoid function itself is a good activation function (ie, it activates a neuron if the input matches certain criteria), but in reality it faces significant issues regarding learning efficiency (the mathematical property of the sigmoid function results in a very small derivative when the input is a large number, therefore limiting the scope of neural network learning). Another example is the optimisation issue. Minimising the neural network loss function is troublesome given the complexity of the neural network. Researchers have spent many years devising decent algorithms to tackle the optimisation issue.

More investment from private sector

AI talent is a scarce resource

We note a clear trend of **large global Internet companies funding AI research and drawing talented AI researchers from the academic world.** As detailed in the graph on page 21, a handful of top AI scientists from top research institutes have headed to the likes of Baidu, Google and Facebook in recent years. For example, Andrew Ng, one of the major contributors to Google's 2013 breakthrough on deep learning visual recognition and a former Stanford University researcher, moved to Baidu as the Chief Scientist from 2014-17. Fei-Fei Li, another Stanford University scholar, joined Google as Chief Scientist of Google's AI and machine learning Institute in November 2016.

In our view, AI scholars are drawn to 2 things offered by large Internet companies: data and computation power. On the data side, the continued accumulation of data by leading Internet companies enables AI scholars to test and train the latest neural networks. Internet companies also naturally have ample computational resources on hand (large amounts of servers and GPUs), which is crucial to the training of neural networks. With

their strong capital backing, ample computational resource and raw data available to crunch (not to mention likely generous compensation), it is likely that the global Internet companies will continue to attract AI scholars going forward.

AI: AI-related investment by notable public companies (2006-17)

Company	Date	Target	Amount	Sector
Baidu	Sep-01-2014	Shanghai Hanfeng		Smart chip
	Dec-01-2014	Pixellot		
	July 18 2016	ZestFinance		Fintech, machine learning
	Aug-01-2016	Velodyne LiDAR	USD75m	
	Dec-13-2016	News in Palm		News apps
	Dec-20-2016	Dynamic Yield	USD22m	Personalization
	Feb-01-2017	8i	USD27m	AI start-up in New Zealand (AR/VR)
	Feb-16-2017	Raven Tech	USD100m	Smart hardware
	Mar-01-2017	NextEV	c. USD100m	Electronic auto
	Apr-13-2017	xPreception		Visual solution
Tencent	n.a.	Diffbot		
	Jun-14-2017	Element AI Inc.	USD102m	AI
Google	Aug-15-2006	Neven Vision Germany GmbH		Computer vision
	Jul-16-2010	Metaweb		Semantic search
	Dec-03-2010	Phonetic Arts		Speech synthesis
	Jan-25-2011	SayNow		Voice recognition
	Jul-23-2011	PittPatt		Facial recognition system
	Dec-13-2011	Clever sense		Mobile apps
	Oct-01-2012	Viewdle	USD45m	Facial recognition
	Mar-12-2013	DNNresearch Inc.		Deep neural networks
	Apr-23-2013	Wavii	USD30m	Natural language processing
	Oct-02-2013	Flutter	USD40m	Gesture recognition technology
	Jan-26-2014	DeeMind technologies	USD625m	AI
	Aug-06-2014	Emu		IM client
	Aug-17-2014	Jetpac		AI, image recognition
	Oct-23-2014	Dark Blue Labs		AI
	Jul-06-2016	Moodstocks		Image recognition
	Sep-19-2016	API.AI		Natural language interfaces
	Jan-05-2017	Limes Audio AB		
Facebook	Jan-06-2015	Wit.AI, Inc.		Voice solution
	Nov-16-2016	FacioMetrics LLC		Machine learning
Apple	Oct-02-2015	VocalIQ Limited		Voice recognition
	Oct-05-2015	Perceptio		AI
	Nov-24-2015	Faceshift AG		Face recognition
	Jan-07-2016	Emotient, Inc.		AI
	Aug-05-2016	Turi Inc.		Machine learning
	Sep-22-2016	Tuplejump Software Private Limited		Machine learning
	Feb-20-2017	RealFace		Face recognition
Microsoft	Dec-18-2015	Metanautix, Inc.		Big data analytics
	Aug-22-2016	Genee		AI assistant
	Nov-01-2016	Interana Inc.	USD18m	Behavioural analytics
	Jan-03-2017	Sarcos Corp.		Robotics
	Jan-13-2017	Maluuba Inc	USD26.2m	Machine learning

Source: Bloomberg, Daiwa compiled

AI: illustration of the AI talent flow from academia to the Internet industry (as of August 2017)

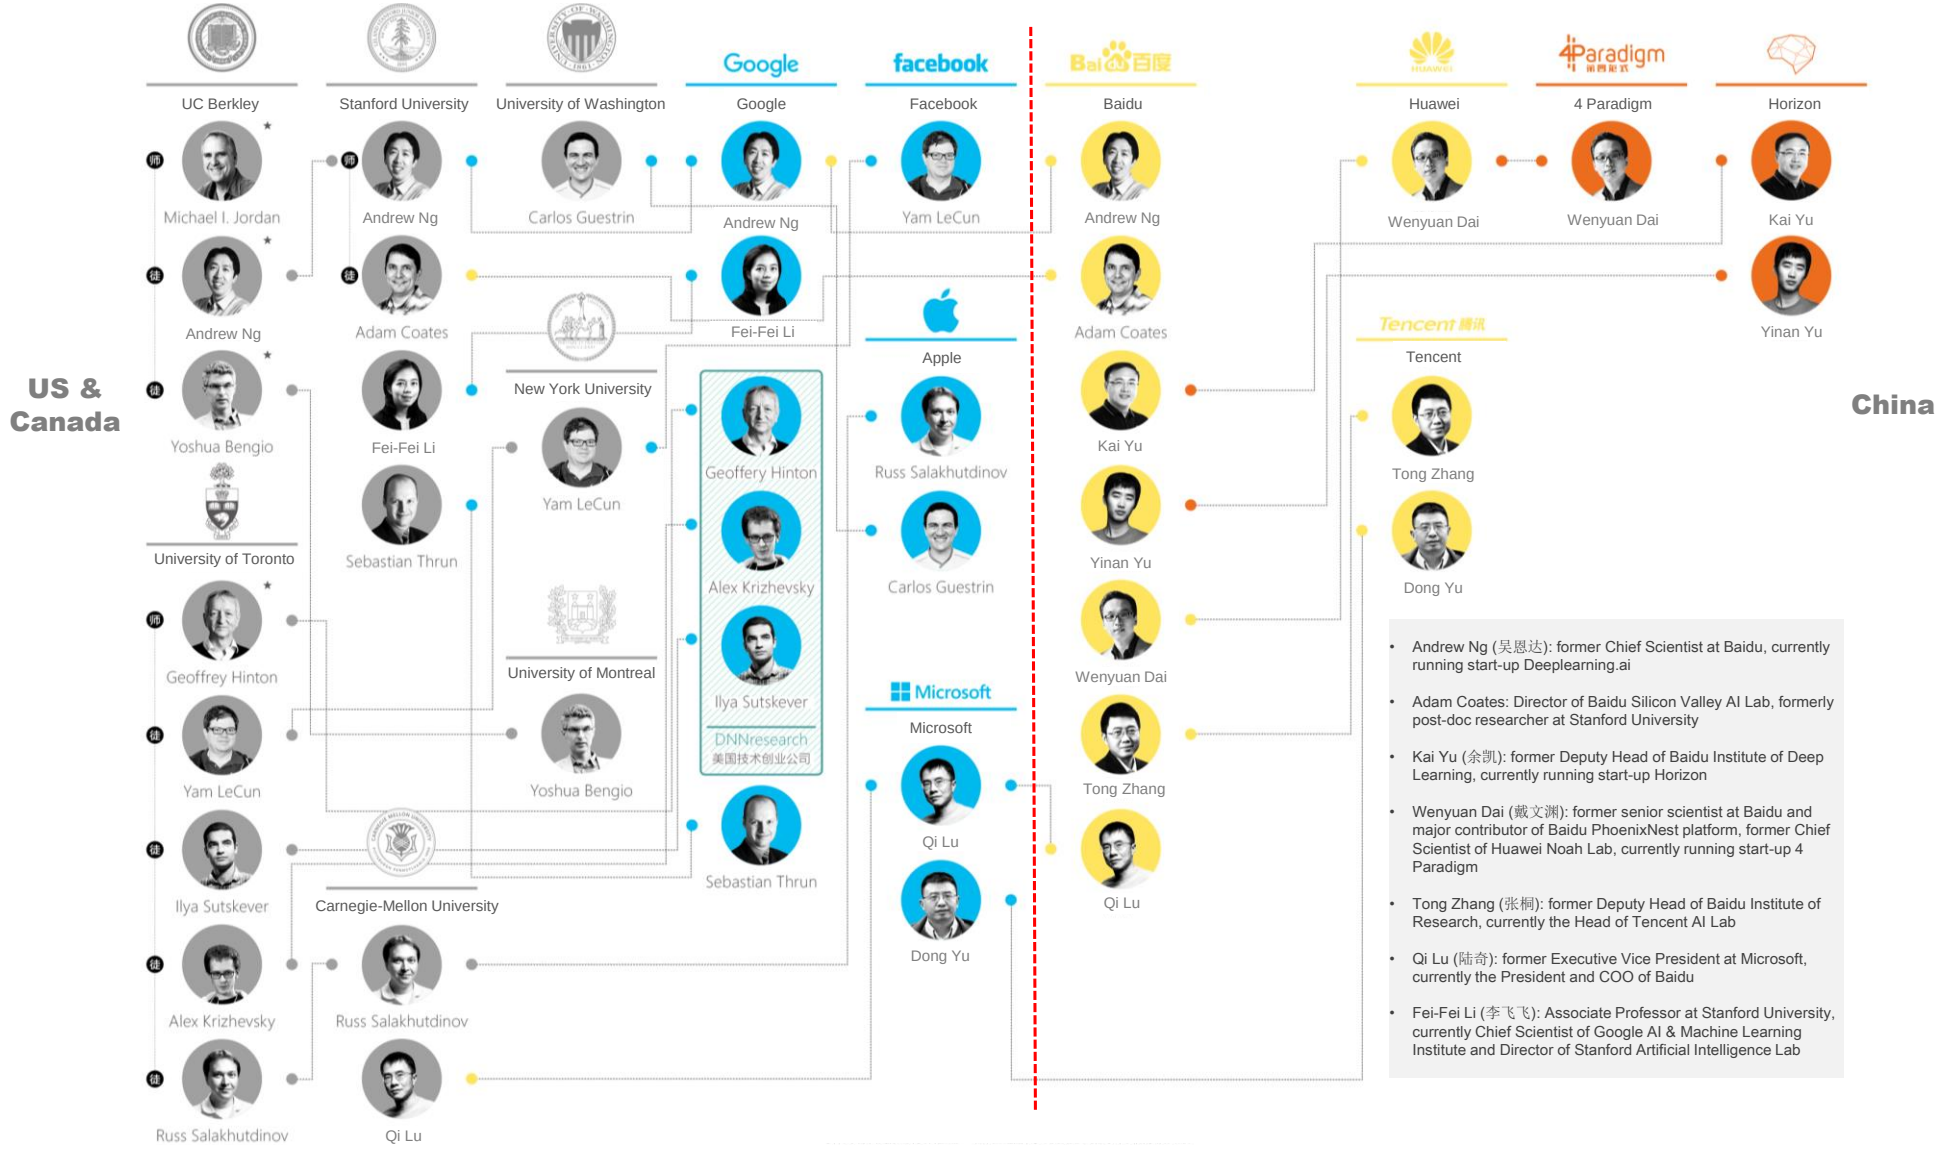

- Andrew Ng (吴恩达): former Chief Scientist at Baidu, currently running start-up Deeplearning.ai
- Adam Coates: Director of Baidu Silicon Valley AI Lab, formerly post-doc researcher at Stanford University
- Kai Yu (余凯): former Deputy Head of Baidu Institute of Deep Learning, currently running start-up Horizon
- Wenyuan Dai (戴文渊): former senior scientist at Baidu and major contributor of Baidu PhoenixNest platform, former Chief Scientist of Huawei Noah Lab, currently running start-up 4 Paradigm
- Tong Zhang (张桐): former Deputy Head of Baidu Institute of Research, currently the Head of Tencent AI Lab
- Qi Lu (陆奇): former Executive Vice President at Microsoft, currently the President and COO of Baidu
- Fei-Fei Li (李飞飞): Associate Professor at Stanford University, currently Chief Scientist of Google AI & Machine Learning Institute and Director of Stanford Artificial Intelligence Lab

Source: CBNWeekly and Daiwa compiled

The other side of the story: inside AI's black box

In this section, we focus on deep learning's "black box" problem

Key takeaways:

- Despite their extraordinary performance at learning, neural networks are complex systems and, in many respects, something of a black box
- The fact that it is not readily "understandable" means that some industries are reluctant to adopt deep learning in their decision-making processes

"We can build these models, but we don't know how they work"

In this section we discuss one of the major hurdles in adopting machine learning in the real world – the black-box issue of AI. Put another way, the black-box issue is the difficulty of being able to look into the network and figure out exactly what it has learnt and how it has done so.

Neural networks cannot fully explain how they make decisions

Recall that the idea behind deep learning is that we allow computer programs to learn from data directly without any direction as to how to learn. As Will Knight from MIT Tech Review puts it:

> *"**You can't just look inside a deep neural network to see how it works**. A network's reasoning is embedded in the behaviour of thousands of simulated neurons, arranged into dozens or even hundreds of intricately interconnected layers. The neurons in the first layer each receive an input, like the intensity of a pixel in an image, and then perform a calculation before outputting a new signal. These outputs are fed, in a complex web, to the neurons in the next layer, and so on, until an overall output is produced. Plus, there is a process known as back-propagation that tweaks the calculations of individual neurons in a way that lets the network learn to produce a desired output."*

By solely relying on computer programs' learning abilities, we lose track of machines' reasoning process in very complicated large neural networks. As a result, **advanced multi-layer neural networks are effectively black boxes to humans.**

How a neural network comes up with the output is, in some cases, not easily explained

AI: illustration of deep learning's black-box problem

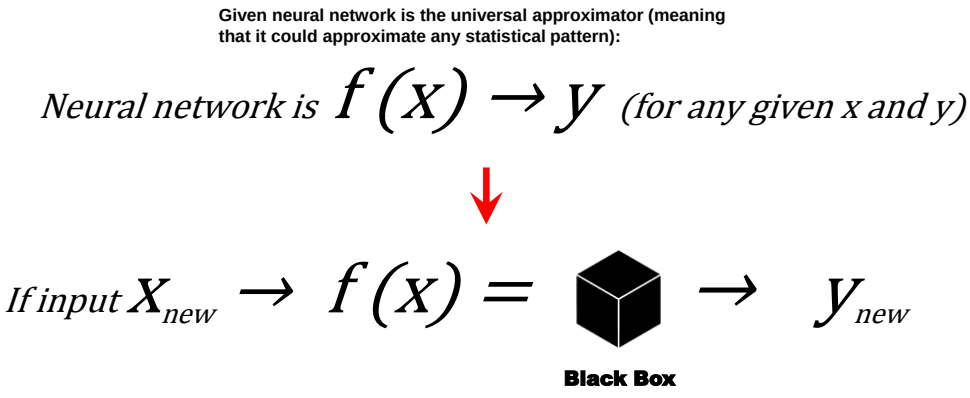

Source: Daiwa

Does AI take responsibility?

As Ben Blume, a partner at VC firm Atomico, has pointed out, "the crux of this problem [the black box problem of AI] is that **in many use cases understanding how a machine reaches a decision is viewed as fundamental to that machine's mass adoption".**

Imagine the following future scenarios:

1. If an autonomous car hits the curb, who bears responsibility?
2. If an AI doctor's diagnosis of a patient is incorrect, who is responsible?

Given a self-driving car drives based on the decision made by an AI system, which learns by itself (or the diagnostic AI makes a decision based on the algorithm that learns by itself), **does the AI bear responsibility for accidents or should the company that built the AI be held responsible?**

Using AI should not require a leap of faith

Given the reasons we discuss above, we believe deep learning will face further hurdles in terms of adoption in the real world, as humans fail to rationalise the decision-making process of machine-learning programmes. This may limit the near-term real-world use of deep learning and negatively impact capital flows into this field. **The irony here is that we humans do not even fully understand how our brains work, yet we trust ourselves to make decisions where the stakes may be high.**

As Will Knight from MIT Tech Review summarises below, the paradox is this:

> *"As the technology advances, we might soon cross some threshold beyond which using AI requires a leap of faith. Sure, we humans can't always truly explain our thought processes either—but we find ways to intuitively trust and gauge people. Will that also be possible with machines that think and make decisions differently from the way a human would? We've never before built machines that operate in ways their creators don't understand. How well can we expect to communicate—and get along with—intelligent machines that could be unpredictable and inscrutable?"*

Are deep learning algorithms intuitive by nature?

If we view this black-box problem from another perspective, the implications could be more interesting, yet scary at the same time. From time to time, human beings encounter scenarios of: "I just know it, but do not know why and do not know how", which we call "intuition" or "instinct". Deep learning possesses the same way of "thinking", where the computer programs provide answers without telling the user how they reached them. This leads to a much broader question: can computer programs constructed in the deep learning framework think "intuitively" like humans?

Machine learning application: search and recommendation

In this section, we focus on how machine learning and deep learning have been integrated with search and recommendation functions

Key takeaways:

- We expect machine learning to greatly improve the user experience of 2 existing Internet products: search and recommendations
- Thanks to machine learning, we expect search to become more personalised and include more information (visual search)
- We expect recommendations to become more personalised and more creative (helping users to expand their set of interests)
- Also, we believe machine learning has the potential to boost the monetisation efficiency (more accurate click-through rate [CTR] prediction) of digital ad operations

> **"Andy Sachs**: '…you know, I'm still learning about all this stuff (fashion) and…'
>
> **Miranda Priestly**: "You think this (fashion) has nothing to do with you. You go to your closet and you select that lumpy blue sweater…but what you don't know is … **that blue represents millions of dollars and countless jobs and it's sort of comical how you think that you've made a choice that exempts you from the fashion industry when, in fact, you're wearing the sweater that was selected for you by the people in this room from a pile of stuff."**
>
> From the movie *The Devil Wears Prada*

In this age of information overload, machine-learning algorithms decide what information users see

Machines decide what information humans should read

Do people consciously choose what they read, or Internet companies (platforms) choose what people should read? This may sound like the start of a conspiracy theory, but it does summarise how people discover information in this mobile and Internet age. We increasingly rely on information distribution platforms (search engines, social networks, news aggregation apps) to show us what to read and see every day. As the information distribution platform shifts from human-curation model to algorithm-driven model, **the algorithms have effectively taken up the job of information allocators in our society.**

Search and recommendation: two sides of the same coin

In a more abstract way, **search and recommendation are one and the same thing: showing users the most relevant information they want to see.** Specifically, search is naturally more "targeted" and "relevant" compared with recommendation because users input a specific query, enabling the search engine (ie, the ranking algorithm) to narrow down the context of information the users intend to seek at that specific moment. In this sense, we could further abstract search engine and recommendation into 2 problems:

1. Understand the context of each piece of information (what this web page/app content means)
2. Rank the information based on the relevancy to the users' intention

We believe machine learning has been a great help in solving these problems.

Organic search (SEO) and machine learning

Example: Google's RankBrain

Search results no longer depend on keywords alone

RankBrain is Google's machine learning-based ranking algorithm for its organic search results page. For each search query, RankBrain calculates the weightage of each ranking factor in the scoring process and interprets the intention of searchers. By doing so, RankBrain decides the importance and the relevance of each individual web page, and therefore the order in which links should be displayed.

Before RankBrain, Google relied greatly on its PageRank algorithm, which was premised on the notion that the most important websites will draw the most links from other sites. PageRank was a revolutionary invention in the early 2000s, but it was unsuccessful in keeping up with users' changing search query trends. It was also found to be susceptible to manipulation by website owners.

RankBrain significantly enhances the relevancy of Google's search results by introducing the factor of "context of search" into the scoring and ranking process. Using machine learning techniques, RankBrain: 1) helps understand the intentions behind search queries instead of looking at the keyword meaning on the surface, 2) learns from users' past behaviour for uncommon search queries and finds correlations between results and queries (correlations that may not be apparent at the beginning).

RankBrain has been changing the search engine optimisation (SEO) industry (the practice of web page owners achieving the highest ranking for specific keywords in the organic search results page) and pushing the search engine to think beyond users' keywords. We could therefore say that RankBrain has made Google's search engine more "human".

Example: Alibaba's personalised search engine (搜索千人千面)

Personalised search offers users predictive product recommendations

Machine learning has enhanced Alibaba's shopping-search engine in its Taobao/Tmall marketplace. Similar to RankBrain, Alibaba has been revamping its search engine algorithm since 2013 to present more relevant products to searchers. The search results could not have been personalised without machine learning techniques. This is because traditional regression techniques tend to be backward looking (recommending similar products based on historical purchase records) rather than offering a predictable guess of user preferences (dynamically updating recommendations based on similar user behaviour).

Before 2013, Alibaba's search results were mainly based on the user query input. If 2 users type the same keyword into the search box, they would likely have seen the same product listings in the search results.

Since 2013, Alibaba has been gradually incorporating more user attributes (product browsing history, recently purchased products, location of the users, etc) into the ranking process to present personalised search results. As the charts below demonstrate, search personalisation means:

- **For the same user and the same keyword, the search results will be different as time goes by.**
- **For different users and the same keyword, the search results will be different even as these users search at the same time.**

Alibaba: illustration of personalised search for mobile Taobao app (September 2017)

Alibaba – Mobile Taobao mobile app search results

User B searches "down jacket" at the same time

User A searches "down jacket"

User A searches "down jacket" again after searching other products

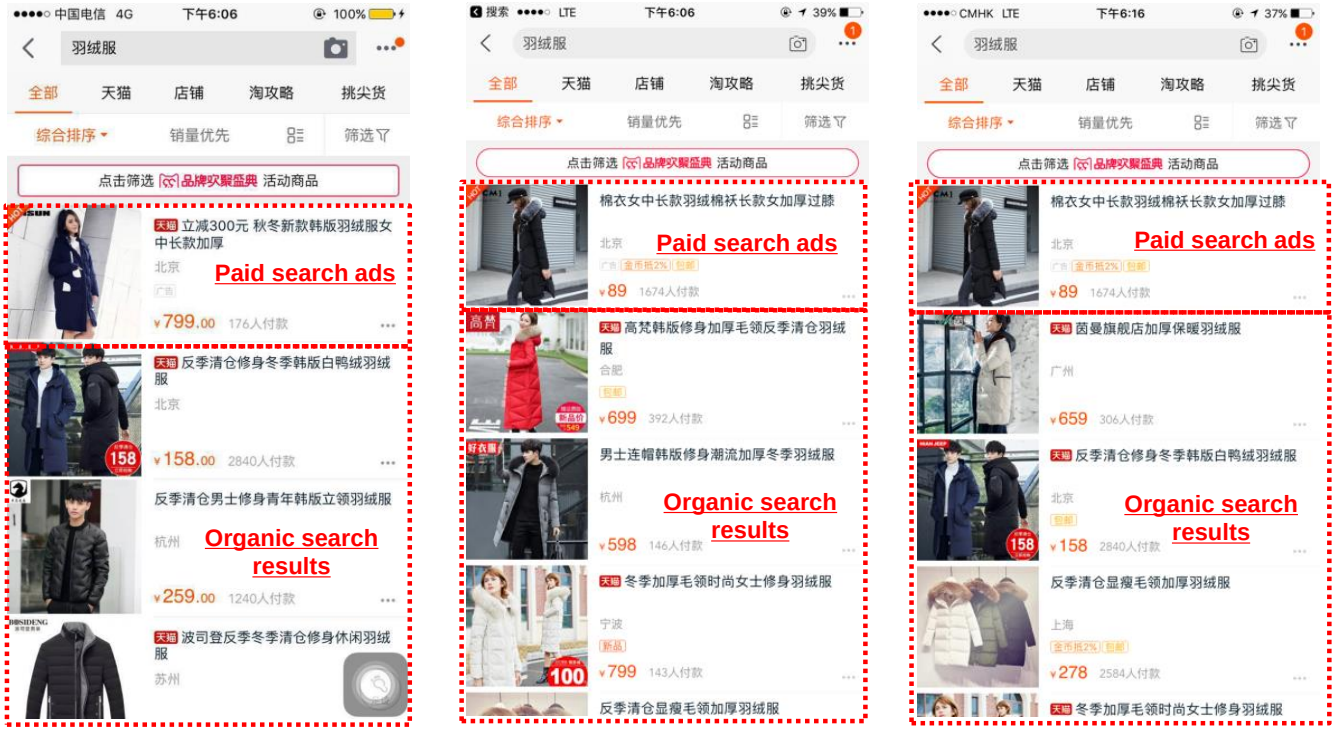

Source: Alibaba, Daiwa compiled

Recommendations and machine learning

Similar to a search engine, machine learning can greatly enhance the relevancy of a recommendation engine. Machine learning, especially deep learning, can identify patterns among large amounts of information, which cannot be found based on a traditional statistical approach. Recommendations offer significant commercial value to Internet companies in the use case of content consumption.

Almost every piece of information you see on your phone or PC Internet browser is put there by the Internet platforms after careful calculations in order to maximise your time spent online and the engagement level.

Example: Netflix

Users can't watch all the content and thus rely on the platform to decide what they may be interested in

Netflix, the world's largest video streaming company, depends heavily on its recommendation engine to maintain user engagement. Digital content consumption, in many cases, is not intention-driven (people don't have a specific TV show/movie in mind) but boredom-driven (people simply want to watch something). Content discovery can be troublesome for consumers, given the opportunity cost of a bad show is high (people simply don't have enough time for trial-and-error viewing). All these reasons give Netflix (or video content platform) significant leverage in terms of pushing content to users.

There are 2 types of recommendation products inside Netflix:

- If you watch certain TV shows/movies, Netflix will recommend a series of other TV shows/movies that you might be interested in (eg, because you watched Iron Man, you might be interested in other Marvel movies).

- In addition, Netflix recommends content in genres that you might be interested in (eg, because you watched Iron Man, you might be interested in action genre movies).

Such a recommendation system sounds intuitive and easy to construct, but this actually ignores 2 challenges:

- While it might be easy for a content platform to recommend users similar-genre content based on their past viewing history, viewers could easily find such recommendations irritating given the lack of diversity (who wants to watch Marvel movies all the time?).

- Digital content, despite the seemingly different genres, shares deep connections. It is highly likely that a viewer who watched The Lord of the Rings, a fantasy genre movie trilogy, would not be interested in the Harry Potter movie series, even though the 2 movies series are commonly listed under the "fantasy" genre. How likely would it be that the recommendation engine that Hamlet and The Lion King share the same fundamental storyline by using a genre labelling system?

Simply recommending similar content is not enough to maintain user engagement

The above challenges can only be solved by a more complex learning system such as deep learning. By focusing on viewer behaviour instead of labelling content into rather general genres, Netflix has found the deeper connection between people's content preferences.

> *"It was engineered out of necessity – the sheer volume of content meant we couldn't expect members to scroll through thousands of titles, so we needed to help them find and discover content they'd love.*
>
> *The early days, we learned a lot – just because someone watches a lot of horror doesn't mean they necessarily want solely horror. We look to get outside of the echo chamber so we can help people discover content outside of their usual.*
>
> *Genres are just one way to categorise a show, but they aren't complete, which is why we also tag series and movies with thousands of additional qualifiers – from mood (goofy) to aesthetic (visually striking) to pace (slow pace) and beyond. This 'tagging process' is just one stage of the Netflix recommendation process. We also have to know about the viewer. For that we use advanced algorithms to examine the viewing habits of millions of members around the world along with their specific taste preferences and viewing histories.*
>
> *The end result is how* ***we present the catalogue uniquely to each member, bubbling to the top of the experience titles that are both relevant and diverse****."*
>
> *- Todd Yellin, Vice President of Product Innovation at Netflix*

Netflix: illustration of recommendation engine

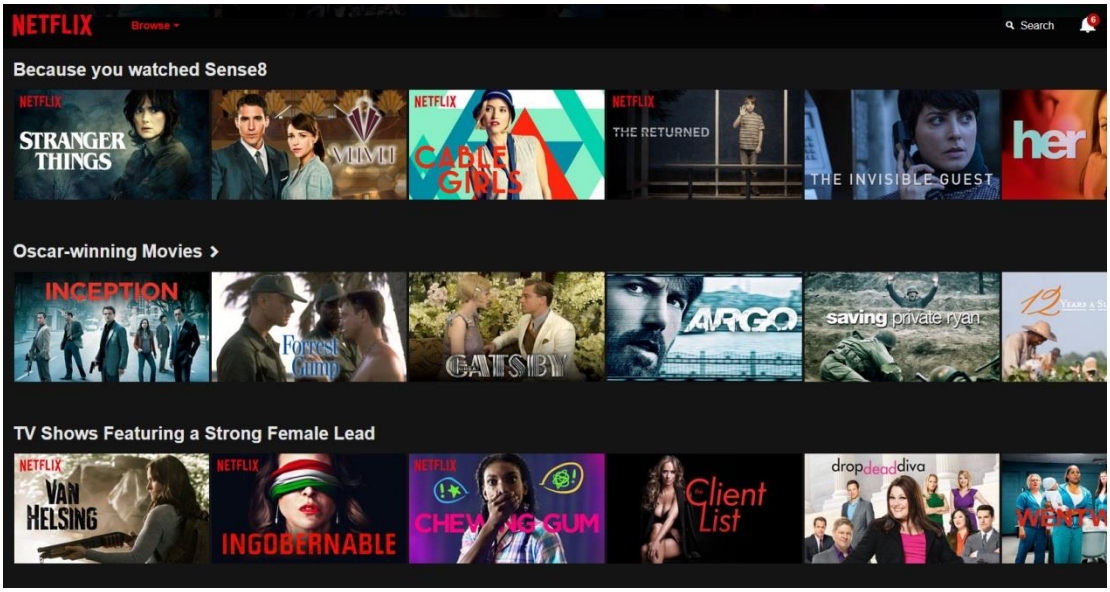

Source: Netflix, Daiwa compiled

Below are some other examples of recommendation engines that we encounter every day that involve machine learning.

AI: illustration of recommendation engine supported by deep learning/machine learning (September 2017)

YouTube

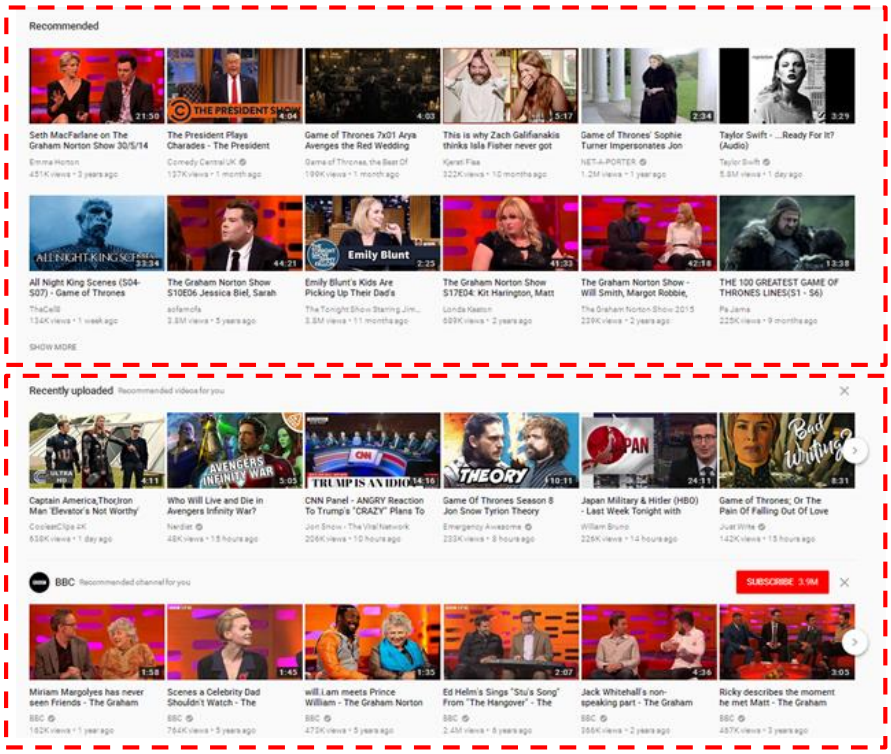

Recommended videos

Recommended channels

Xiami – Alibaba's online music business

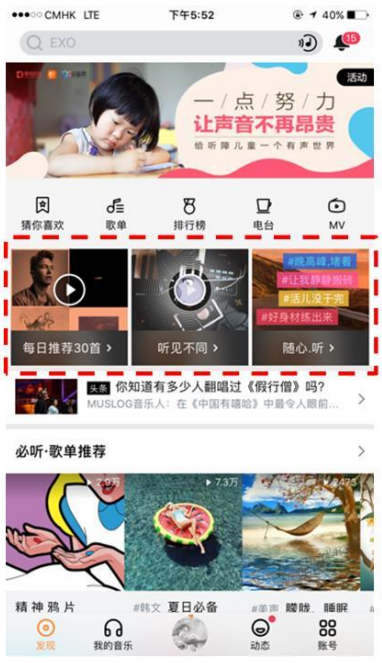

Songs recommended by machine learning engine

Recommendation engine helps users to listen to new genre of music, instead of just similar songs

Users could choose to go "aggressive" or "conservative" on new genre of music

Source: Companies, Daiwa compiled

Machine learning enhances monetisation efficiency of existing online ad business model

CTR prediction and machine learning

One final point we want to highlight is how machine learning can directly impact the revenue streams associated with the digital ad business.

Predicting CTR is at the core of any online ad platform

The idea behind search ads and display ad platforms is similar: in what order should the ads be presented to users? The ranking should be determined by multiple factors, namely the likelihood of the user clicking on the ad (CTR), the bidding price of the ad, and the quality or relevancy of the ad. Specifically, the bidding price and quality of the ad are known before presenting the ad — but the likelihood of a user clicking on the ad is unknown until the ad is actually presented. Therefore, in more precise terms, Internet companies need to "predict" the expected CTR for each ad before it is displayed, which is at the core of the design of an online ad system.

Machine learning could enhance the prediction accuracy of CTR for online ad platforms. Currently, the algorithm most commonly used in predicting CTR uses logistic regression, a simple sigmoid function that can explain many real-world statistical distributions. Logistic regression is easy to understand, construct and scale up, and therefore still widely used in predicting CTR. But there is a clear drawback — the algorithm is simple and therefore cannot capture complicated patterns within data sets. Machine learning, especially deep learning, can enhance the accuracy of CTR prediction, since it can incorporate more user behaviour features compared with conventional algorithms.

For example, in July 2017, a research paper published by Alibaba Group revealed that Alibaba had introduced machine learning into its performance-based ad CTR prediction algorithm in 2014. Indeed, Alibaba's management publicly acknowledged the revenue boost from the algorithm change, with increased CTR prediction accuracy helping to lift RPM (revenue per thousand clicks), as shown below:

> *"The other effect of increasing clicks is the personalisation where we use AI technology to improve the content that people see, so that they are more personalised to you, the shopper, and that also increases the clicks."*
>
> *- Joe Tsai, Executive Vice Chairman of Alibaba*

> *So, basically why the paid clicks grow is because we provide more relevant content. And consumer engagement is growing. So, what backs up that more relevant content is our data technology.*
>
> *- Maggie Wu Wei, CFO of Alibaba*

Case study: Toutiao (今日头条)

Here we briefly introduce Toutiao, a leading news aggregator platform and one of the largest private Internet companies in China (current valuation: USD10bn in the latest round of financing). We view Toutiao as a must-study target within China's Internet space, as it has effectively transformed the news app industry in China by adopting machine learning technology.

Toutiao (literally meaning "headline" in Chinese) was launched in August 2012. It offers users news content covering a wide range of topics (including political news, entertainment news and local news) in text, image and video formats. Toutiao defines itself as an AI company (an algorithm-driven news platform). It differentiates itself from conventional online news companies such as Sohu News and Sina News by not producing news content, but rather distributing news content created by other parties.

Toutiao uses machine learning to distribute information

Toutiao offers users personalised news content based on its in-house machine learning algorithm. Without asking users to identify their content preferences, Toutiao records each user's reading behaviour and constantly pushes new content to users through dynamic machine learning programs running at the back end.

Toutiao is a diverse content platform

Toutiao: screenshots (August 2017)

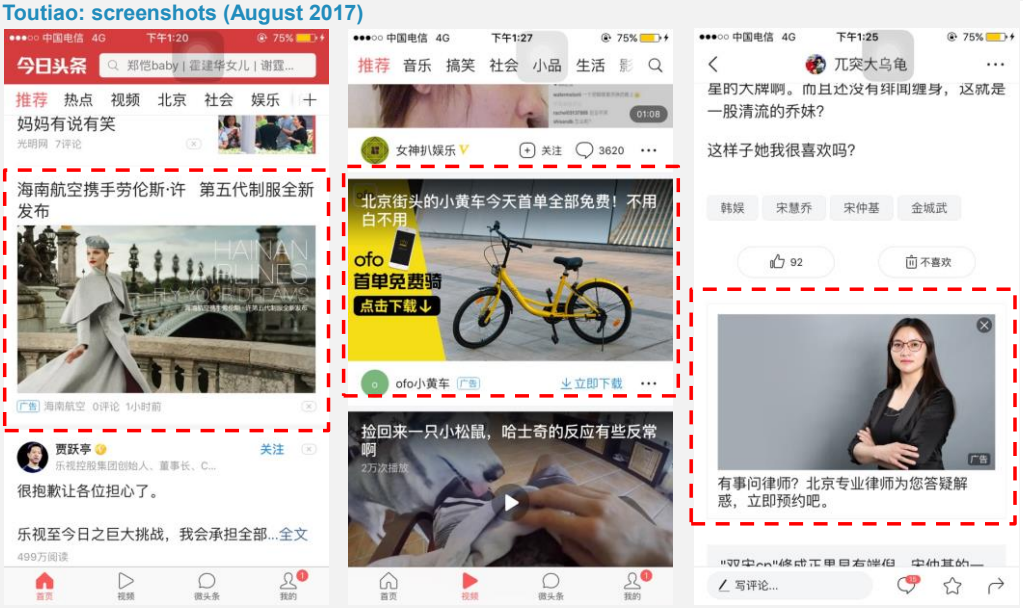

News feeds with in-feed ads | **Video feeds with in-feed ads** | **News article with banner ads**

Source: Daiwa compiled

Vast user base and strong monetisation capability

In terms of time spent per daily active user (DAU), Toutiao ranks No.2 among all mobile apps in China (after Tencent's WeChat), according to Questmobile. As of the end of 2Q17, Toutiao had 75m DAU and the daily time spent per DAU reached 78 minutes in 2016 (vs. 90 minutes for WeChat).

Toutiao retains the largest user time spent among news apps

Toutiao: time spent per DAU comparison (1Q17)

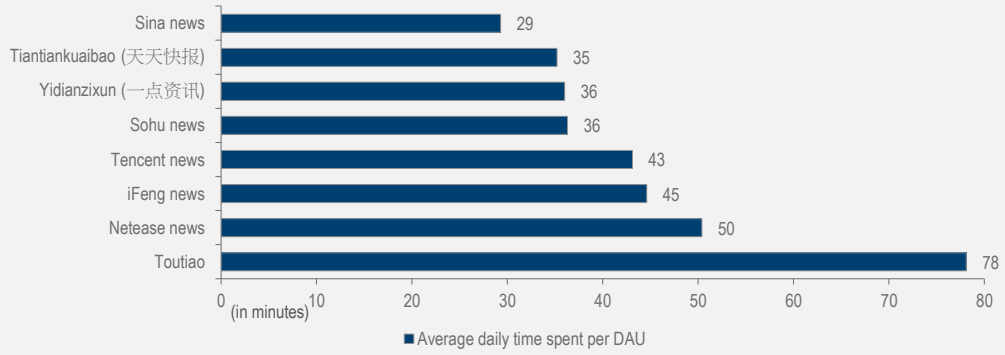

Source: Questmobile, Daiwa compiled

Toutiao's DAU more than doubled YoY in 1Q17

Toutiao: DAU comparison (1Q16-1Q17)

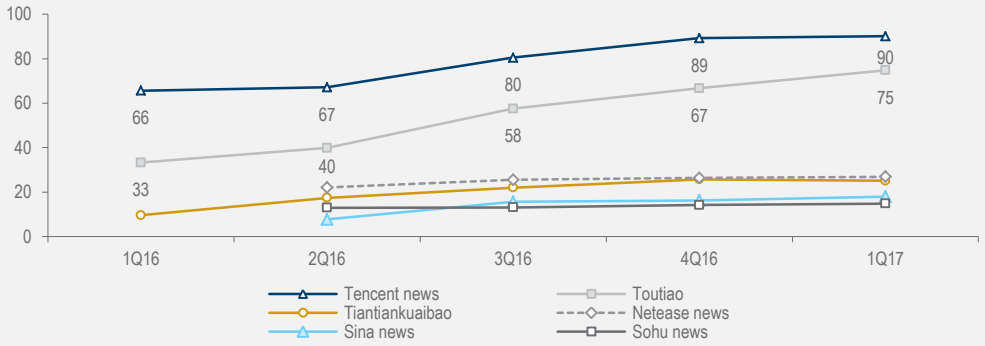

Source: Questmobile, Daiwa compiled

Given the large number of users and amount of time spent on Toutiao, advertising remains its biggest monetisation model. Ads within Toutiao include: 1) display ads shown in the news article lists (in-feed ads), and 2) display ads shown in the bottom of news article content (banner ads). As detailed in the table below, we estimate Toutiao generated around CNY15bn in revenue for 2016 and forecast its revenue to total CNY12.8bn for 2017. From another angle, we estimate Toutitao's ad revenue in 2017 will be equivalent to 31% of Tencent's online ad revenue and 17% of Baidu's online ad revenue in 2017.

We estimate Toutiao's 2017 revenue will be equivalent to 17% of Baidu's ad revenue

Toutiao: estimated revenue for 2016 and 2017

		2017	Note
	In-feed ads		
a	Average DAU (m)	100	Daiwa estimated
b	No. of feeds one DAU exposed to per day	125	Daiwa estimated; on average one DAU read around 25 news per day, we assume one DAU on average read 1 out of 5 articles
c	Ad load	12%	Daiwa estimated
d	eCPM (CNY)	15	Daiwa estimated
e	Total in-feed ads revenue (CNYbn)	8.2	e=a*b*c*d/1000*365
	Banner ads		
a	Average DAU (m)	100	Daiwa estimated
f	No. of news articles one DAU read per day	25	Daiwa estimated
g	eCPM (CNY)	5	Daiwa estimated
h	Total banner ads revenue (CNYbn)	4.6	h=a*f*g/1000*365
	Total advertising revenue (CNYbn)	**12.8**	

Source: Daiwa estimates

Content unbundling and the disintermediation of content creator

Personalisation of news content is not a new concept. Over the 20 years of development of the Internet industry in China, we observe a clear path of evolution for online news media:

- Traditional portal model (late-1990s to early-2000s, the integration of online portals and news content creation)
- Self-content publishing (the so-called Web 2.0 and the emergence of personal online blogs in the mid-2000s to late-2000s)
- Social network and interest-based distribution (the emergence of Facebook, Twitter and Weibo)
- Machine learning information distribution (Toutiao)

For many, building a service on top of an existing platform is much more cost efficient than developing an app from scratch

There is a clear trend under way globally. **Content providers are shifting away from using individual mobile apps as the sole way to deliver information to users, and starting to leverage multiple existing platforms — a process referred as disintermediation of content creators.**

Nowadays one can read *The New York Times* content across a wide range of platforms, through Google search to Facebook and Twitter. Another example is Buzzfeed, the social news provider, which has a presence on multiple platforms, including Facebook, Twitter and Snapchat. Buzzfeed has disclosed that these social platforms cumulatively generate most of its viewership, as opposed to its own app and webpage. Buzzfeed is an independent content creator/media company that is not attached to a single app or website.

The disintermediation of content creators is even more apparent in China, as the traditional news media players are less developed than their overseas peers. Major Internet companies including Tencent (WeChat Official Account), Baidu (Baijiahao) and Alibaba (UC Toutiao) have built their own content distribution platforms. Toutiao has been aggressive in signing exclusive content creators to deepen the news content on its platform.

We discussed the information discovery bottleneck problem in greater detail in our China Internet Primer: decoding Internet discovery and digital advertising (4 March 2016).

Application of machine learning: speech and visual recognition

In this section, we focus on how machine learning and deep learning help the development of speech and visual recognition

Key takeaways:

- Cameras and microphones are replacing keyboard on mobile devices; voice input and visual input are reshaping the user interactive model, with the help of machine learning
- Deep learning has significantly enhanced the accuracy of speech recognition, visual recognition and translation
- We are upbeat on the adoption of speech and visual recognition in existing Internet products (visual search, voice search) and other industries (AR, fintech, medical, autonomous driving)

Speech recognition and natural language processing

> *"Every time I fire a linguistic, the performance of the speech recognizer goes up."*
>
> *- Frederick Jelinek*

Speech recognition is one of the earliest use cases for machine learning, dating back to the late-1970s. This is simply because the supply of digitalised speech data back then was more abundant than other types of digital data (there were few digitalised photos in the 1970s). For example, IBM had accumulated a decent amount of digitalised telegram data by the end of 1979, which enabled its research team to adopt a machine learning approach to study speech recognition at the IBM Watson Lab.

When people talk about speech recognition or natural language processing, they are in fact talking about 2 separate tasks:

- Task 1: If we input a piece of human speech, the computer program can identify and present the full sentences (ie, **speech-to-text**)
- Task 2: the computer program can interpret the meaning and context of the input sentences (ie, **text-to-semantics**)

Currently, technology is able to complete Task 1 with a relatively high level of accuracy. One thing to note is that speech-to-text task does not necessarily involve complicated algorithms such as deep learning, though nowadays deep learning is proven to be an effective toolkit in solving the speech-to-text problem. Speech-to-text recognition has been applied in many use cases, as shown below.

AI: illustration of speech-to-text recognition use cases in Internet products (September 2017)

Tencent **Baidu**

Ask WeChat to translate the speech to words

Identified result of the translation

Voice input search in Baidu Maps

Source: Companies, Daiwa compiled

The breakthrough in speech recognition in the 1970s was a watershed event in terms of machine learning research. Frederick Jelinek, who worked at IBM's Watson Lab in the late-1970s, almost singlehandedly improved the speech-to-text accuracy rate from 70% to 90% after adopting a statistical approach to speech recognition rather than the conventional linguistic pattern-searching approach. Given his background in telecom transmission research, Jelinek took the view that language was simply a sequence of small pieces of information. As such, natural language processing is, in his view, a statistical problem.

Using information theory knowledge and algorithms such the Hidden Malkov Model, Jelinek led the research team to develop a program that could identify full sentences from speech. The program does not learn individual words in sentences by rule of language, but simply matches words to a big database of existing speech and chooses the one with the highest probability of matching. This approach is "system-depending", meaning that the speech data needs to feed in the programs in order to let it learn. Interestingly, this approach is very similar to the modern philosophy of machine learning – feed large amounts of data into a program and let it learn by itself.

Translation: from rule-based approach to statistical approach

Translation is a classic machine learning problem, and deep learning has greatly enhanced its accuracy in recent years

After identifying the wording of the speech, the next step is to let computers understand the content of the language. Translation becomes the key use case for testing how intelligent the computers are in understanding the context of the language.

Conventionally, scientists treat the translation problem as a linguistic one, meaning that computer programs are designed to understand the rule of our languages first in order to interpret the content of the language. Taken as a linguistics problem, if we tell computers all the existing rules of human languages, computers will be able to fully understand the meaning of those human languages. This rule-based approach soon reached a bottleneck, as human language rules would be too complicated to be understood by computer program code. In order to reach a satisfactory level of translation, the number of rules (on general grammar and special use case of words) would be too large to input manually.

With the inspiration of Frederick Jelinek (discussed above), researchers have turned to a statistical approach to tackle the translation issue. The fundamental difference between a statistical translation system and a rule-based one is that a statistical translation system does not try to understand the meaning of the sentence. Instead, a statistical translation system generates thousands of possible translation candidates after scientists have fed it rafts of real translation data. The system then ranks those translation candidates by likelihood of correctness. It measures how "correct" one translation candidate is by how similar it is to the training data (the actual translations by humans).

One step forward in natural language processing: bringing deep learning into translation

Deep learning now enables computers to translate full sentences without breaking them into individual phrases

Following the deep learning breakthrough of the 2010s, researchers have been considering adding deep learning to the statistical translation model. Despite having a similar design philosophy (statistical instead of rule-based approach), the deep learning translation system is quite different from the traditional statistical translation model in that it translates "a whole sentence at a time" vs "a phrase/word at a time".

The introduction of deep learning into machine translation makes the translation process easy to understand. Deep learning (or neural networks) can be considered a program that learns the underlying patterns in large data sets. Human languages are effectively the same information delivered in different patterns. In theory, with enough sample data sets to learn from, deep learning could find the link between 2 languages with a high degree of accuracy.

AI: illustration of a deep learning translation system design

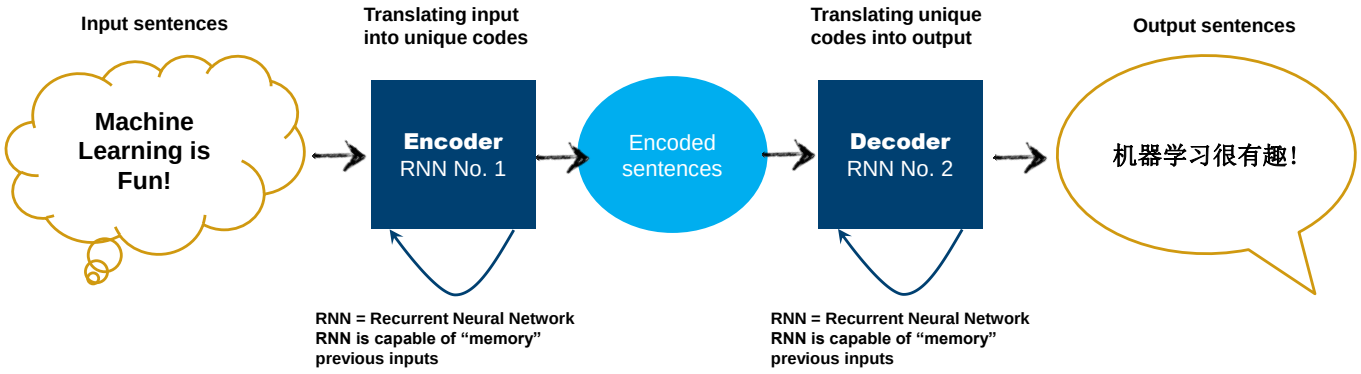

Source: Adam Geitgey, Daiwa compiled

Example: Google Translate

Google is the pioneer in adopting machine learning techniques to the use case of translation. Google officially launched its Translate service in 2006, supported by the statistical machine translation technology we discussed earlier. Not long after the initial service launch, Google beat existing translation services in terms of translation accuracy in various competitions despite it being a latecomer to the field.

Google revamped the technology behind Google Translate in 2016 by introducing Google Neural Machine Translation (GNMT). GNMT is the algorithm behind the present Google Translate product, which takes in entire sentences and produces instantaneous translation. Similar to other deep learning networks, GNMT is an end-to-end system and current researchers have yet to discover how exactly the neural networks "analyse" in order to come up with the translation end-result. The "black box" issue of GNMT does not overshadow its outstanding performance, however. After only 3 years of introduction, GNMT has surpassed traditional statistical machine translation models in terms of translation accuracy. Many industry players are optimistic that the translation accuracy of GNMT, as well as other deep-learning-backed neural translation systems, will reach a par or exceed that of professional translators in the near future.

Visual recognition: cameras and the future of AR

Visual recognition is a classic use case of machine learning. The ultimate goal of visual recognition is to let computer programs understand the content of visual data (photos and videos) by simply "reading" it. Such a task, unlike speech recognition, remained a largely unsolved problem before the deployment of deep learning techniques.

As we touched upon on page 12, the breakthrough in visual recognition happened in the early 2010s when a team of Google researchers fed large amounts of Youtube videos into a neural network and allowed it to identify objects without giving it any of the specifications for objects. The framework of using neural networks for visual recognition is not difficult to understand – each pixel of the visual content (training set) is considered an input; together these inputs are fed into a neural network; after layers of calculations, the output layer gives the classification results.

Cameras are replacing keyboards; visual recognition is crucial for AR

As information formats become increasingly rich (from words to live video), the boom in digital photo and video content seems unlikely to slow any time soon. With much of the world's population having access to a smartphone, the cost of producing digital photos and video has never been so low.

AR is a user interface built on a camera

What will the next format of information be? Augmented reality (AR) seems to be heading in a promising direction. With AR, the boundaries between the virtual world and reality blur, while the camera becomes even more important as it becomes the next-generation user interface. We need technology to understand and interpret digital images before we can have an impactful AR experience. In this sense, visual recognition (supported by deep learning) is crucial to the development of AR.

Accurate indexing of visual content is crucial for content discovery

Computers still cannot fully understand rich content (photos and videos); deep learning could help with that

One of the many use cases of visual recognition is visual search. The vast majority of information we create today is not in the format of words, but in sound, photos and videos. Currently computer programs still have trouble identifying objects within photos and videos, and have to rely on word descriptions to "guess" the content. For instance, the predominant search algorithm for video search within YouTube is still based on the video's description rather than the actual content of the video. If a user uploads a video about Star Wars but accidentally inputs the video description as Star Trek, the system is likely to show this video when someone searches for Star Trek.

With the evolution of visual recognition technology, we expect search engines to expand from indexing the word content online to the visual content online. In our view, this development would greatly enhance the relevancy of search results and bring new monetisation potential (ads within video content). It is also crucial for computers to actually understand what is inside visual content.

Example: Pinterest photo search

> *"People thought of us as a social network… we think of Pinterest much more as a search engine."*
>
> *- Evan Sharp, Co-Founder and Chief Product Officer of Pinterest*

Pinterest is a large photo catalogue website. Users can "pin" photos or videos that they find interesting while the Pinterest website presents these "pinned" photos to users. Being a company that deals mostly with non-word format content, Pinterest introduced visual search to its platform, enabling users to find relevant photos with its search function. In addition, in 2015 Pinterest launched a toolkit within Google Chrome, allowing users to select any part of a picture on a web page and present similar products. This function is very similar to Alibaba's "search similar (看同款)" function within the mobile Taobao app, which presents users with similar products when searching.

Visual search, backed by deep learning, allows photos and video content to be indexed accurately by a search engine

Pinterest: illustration of a visual search

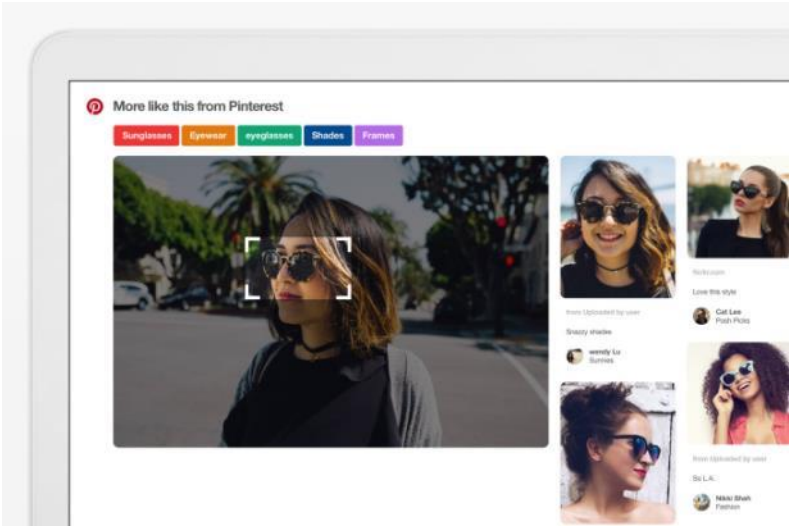

Source: Techcrunch

AI: illustration of visual recognition use cases in major Internet products

QR code scan for Alipay

Snapchat Lens

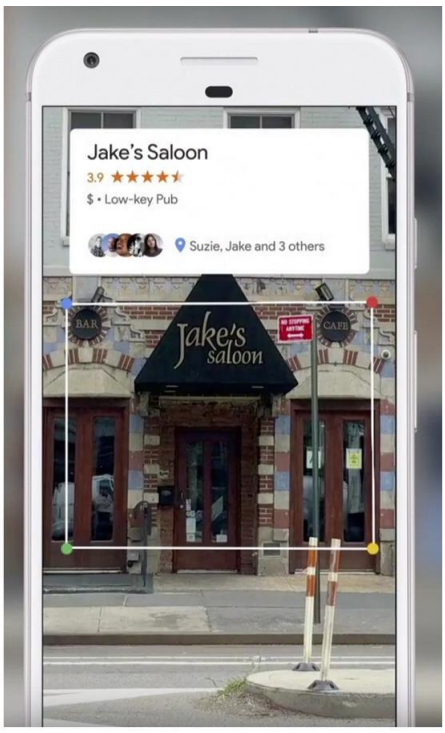

Google Lens

Source: Companies, Daiwa compiled

In addition to Internet products, visual recognition has been applied in areas including:

- Motion detection (autonomous driving)
- Face recognition (counter-terrorist security systems)
- Hospitals (medical diagnoses)
- Financial services (remote bank account opening)

Example: insurance damage assessment using visual recognition

Ant Financial Group released a machine-learning-based visual recognition product called Ding Sun Bao (DSB) (定损宝) in June 2017. This product is used for automobile insurance damage assessments. Drivers can upload photos of minor damage such as scratches after an accident. The uploaded photos are processed by a machine-learning algorithm, which gives out claim examination results in just a few seconds. The results are then passed to the insurance company for the claim process.

If fully implemented, minor incidents would not require an insurance company claim adjuster to check and assess damage going forward. According to Ant Financial Group, the accuracy rate of DSB has reached the same level as that of a 10-year veteran claim adjuster. Ant Financial Group also estimates visual recognition technology can help insurance companies reduce their human claim adjuster costs by 50%, suggesting an annual saving of CNY2bn for the insurance industry as a whole.

DSB uses visual recognition to reduce labour costs for insurance companies

AI: demo of Ding Sun Bao (DSB)

Source: Ant Financial

Appendix

Appendix 1: company profile on AI initiatives

We chose 12 leading Internet companies (7 Chinese, 5 global) that have been beefing up their AI research and investment. Over the next few pages, we briefly introduce each company and outline their notable initiatives in AI development over the past few years.

The featured companies are as follows:

- Baidu
- Alibaba
- Tencent
- iFlytek
- Toutiao
- Sogou
- Didi Chuxing
- Facebook
- Alphabet
- Amazon
- Apple
- Microsoft

Baidu (BIDU US, Hold [3])

Share price (Sep 13):	**USD236.41**	52-week range:	**USD159.5-239.0**
Market cap:	**USD82.0bn**	TTM R&D expenses:	**USD1,456m**

An "AI First" Chinese search engine

Investment thesis

Business description

- Baidu is the largest search engine in China, with a mobile search MAU of 665m as at end-2016. Its commercial products include keyword search, news feeds, online video (iQiyi) and O2O products (Nuomi). The company also holds a 15% shareholding (on a fully-diluted basis) in Ctrip.

Revenue breakdown

- We estimate core search contributed 65% of Baidu's total revenue in 2Q17 (CNY13.5bn).
- iQiyi (online video) revenue was roughly CNY4bn in 2Q17, including both subscription and ad revenue, representing 20% of Baidu's total revenue.
- News feeds accounted for 7% of the company's 2Q17 revenue.

Notable initiatives in AI

- Baidu open-sourced Apollo, an AI-based autonomous driving solution, in July 2017. Apollo allows 3rd party developers to leverage Baidu's AI algorithm to improve driverless car functionality. Baidu cooperates with 50-plus partners including auto OEMs (Ford, Daimler, Chery, etc) and hardware suppliers (NVIDIA, Bosch, etc) on autonomous car initiatives. We believe Baidu is the clear leader in the autonomous driving market in China given its first-mover advantage.
- Baidu is also a leading speech recognition player in China. Its voice search accuracy rate stood at 86% as at end-2016, and over 10% of search traffic now comes from voice input. DuerOS, Baidu's answer to Apple's Siri, leverages this technology to provide comprehensive intelligent assistant solutions.

Alibaba (BABA US, Buy [1])

Share price (Sep 13):	**USD178.97**	52-week range:	**USD86.0-179.1**
Market cap:	**USD452.7bn**	TTM R&D expenses:	**USD1,764m**

Hidden leader in AI research in China

Investment thesis

Business description

- Alibaba dominates China's ecommerce market, with its GMV accounting for c.70% of China's total online retail sales in 2016.
- AliCloud controls a 41% share of China's IaaS market, per IDC. Alipay (the payment service arm of Alibaba) is also the No.1 digital payment company in China.

Revenue breakdown

- E-commerce accounted for 82% of the group's total revenue in 4Q FY17. Within e-commerce, China retail e-commerce made up over 80% of the segment's revenue.
- AliCloud contributed 5.6% revenue for the group in 4Q FY17 with triple-digit YoY revenue growth for consecutive 8 quarters. The rest of the revenue mainly comes from the entertainment business.

Notable initiatives in AI

- Deep-learning technology has empowered Alibaba to become a content discovery platform from a pure marketplace, in our view. Machine learning allows Alibaba to provide personalised search results and shopping recommendations, and hence improves user stickiness. This is evidenced by the DAU/MAU ratios of Mobile Taobao app increasing from 33% in FY14 to 40% in FY17.
- Alibaba tested the waters in intelligent home hardware in mid-2015. In July 2016, it rolled out Tmall Jingling (天猫精灵), a voice-activated intelligent home hardware. The hardware is operated by Alibaba's Chinese language processing system, 'AliGenie'. Tmall Jingling allows users to place online shopping orders, top-up mobile phones, play music and order food through voice communication.

Tencent (700 HK, Buy [1])

Share price (Sep 13):	**HKD334**	52-week range:	**HKD179.6-341.0**
Market cap:	**USD410.7bn**	TTM R&D expenses:	**USD2,146m**

China's mobile gateway turns to AI

Investment thesis

Business description

- Tencent is China's largest social network company, with 938m MAU for its flagship WeChat product.
- Thanks to strong momentum in online games, online ads and digital payment business, Tencent's total revenue rose by 38% YoY to CNY152bn for 2016.

Revenue breakdown

- Some 46% of Tencent's revenue came from online games in 1Q17.
- Community value-added services, including subscriptions from Tencent Video and QQ memberships, contributed 25% of total revenue in 1Q17.
- Online advertising accounted for 14% of Tencent's total revenue in 1Q17.

Notable initiatives in AI

- Tencent AI Lab, founded in April 2016, focuses on visual recognition, speech recognition, natural language processing and machine learning research.
- Fine Art (绝艺), an AI-based chess program developed by Tencent AI Lab, beat DeepZenGo (a Japanese chess program) during the 10th UEC global chess contest in March 2017.
- On 3 August 2017, Tencent released the machine-learning-based medical image recognition, Tencent Miying (腾讯觅影). Miying supports 6 early-stage pathema screening systems, including esophageal cancer, cervical cancer and breast cancer.
- Machine learning has improved Tencent's ad targeting, enriched its social interactions via features such as animated face masks, and reduced costs by replacing customer service with intelligent agents.

Toutiao 今日头条 (not listed)

Share price (Sep 13):	**n.a.**	52-week range:	**n.a.**
Market cap:	**USD11bn (latest financing)**	TTM R&D expenses:	**n.a.**

Investment thesis

Business description

- Toutiao is one of the largest private Internet companies in China. The company has successfully built upon its flagship news aggregation mobile application. It has publicly said that its newsfeed product is backed by machine-learning technology, rather than being handpicked by editors. According to QuestMobile, Toutiao' s MAU was 173m in 1Q17 and the average daily time spent per DAU was 78 minutes, significantly higher than peers.

Revenue breakdown

- Based on Toutiao's public disclosure and backed by our research, we calculate full-year revenue for 2017 could reach CNY12-15bn.

Notable initiatives in AI

- Toutiao, founded in August 2012, was the first news aggregator company to adopt an algorithm-driven content selection model in China.
- By tagging users with various features in multiple dimensions (eg, interests, gender, location, etc), Toutiao trains its content selection algorithm by feeding in users' historical behaviour data. For this, Toutiao said it had gained an exceptional level of user engagement, as evidenced by the time spent per daily user (eg, daily time spent per DAU for Toutiao, NetEase news app and Tencent news were 78 minutes, 50 minutes and 43 minutes for 2016, per Questmobile).

iFlyTek 科大讯飞 (002230 CH, not rated)

Share price (Sep 13):	**CNY55.71**	52-week range:	**CNY25.2-64.8**
Market cap:	**USD11.5bn**	TTM R&D expenses:	**USD125m**

Pure machine-learning A-share play

Investment thesis

Business description

- iFlyTek, founded in 2000, claims to be China's leading voice recognition and natural language processing (NLP) company. It has developed various AI-based solutions for the education and customer service sectors, among others.

Revenue breakdown

- iFlyTek booked CNY3.3bn in revenue for 2016, up 33% YoY.
- AI-based education-related solutions accounted for 31% of revenue; automatic system solutions and Internet products & related value added operations accounted for 26% and 13% of total revenue, respectively; Big Data products, AI-based customer service products and public security products contributed the rest.

Notable initiatives in AI

- iFlyTek believes it has made great achievements on the NLP front. It won the championships for "Blizzard Challenge (2006-2015)", "CHiME Challenge (2016)" and "SQuAD (1H17)", which are renowned global contests for NLP.
- iFlyTek has a series of "smart education solutions" in the market. Smart education solutions enable personalised recommendations for study assignments and suggest the most suitable college for students to apply to based on their *gaokao* (college entrance exam) results and colleges' historical ranking.
- iFlyTek signed a strategic cooperation agreement with Peking Union Medical College Hospital on 18 April 2017. Leveraging the company's NLP and speech recognition capability, iFlyTek assists doctors at this hospital by transforming their verbal instructions into written clinical prescriptions.

Sogou (not listed)

Share price (Sep 13):	**n.a.**	52-week range:	**n.a.**
Market cap:	**USD6bn (latest financing)**	TTM R&D expenses:	**n.a.**

The No.2 search player in China

Investment thesis

Business description

- Sogou is the 2nd largest search engine by revenue in China after Baidu.
- The company is a consolidated entity of Sohu (SOHU US), which held c.36% of Sogou's equity capital and over 50% of its voting power on a fully-diluted basis at the end of 2016. Tencent owns a c.45% stake in Sogou but with less voting power.

Revenue breakdown

- Sogou booked USD211m in revenue for 2Q17, with c.90% contributed by search and related areas; the remaining 10% came from the IVAS business, such as the operation of web games and mobile games developed by 3rd parties.

Notable initiatives in AI

- Similar to other online ad monetisation companies, Sogou applies AI technology in order to present relevant information. Sogou currently provides personalised research results to users and is testing the waters in newsfeeds to push relevant news content to users. According to management, search ad CTR improve in 2Q17 through leveraging its machine-learning capability.
- Sogou is also focused on NLP and voice recognition. It currently operates the largest Chinese-language input system and has accumulated a large amount of Chinese-language data points. Per management, Sogou's mobile keyboard now supports over 260m voice inputs on a daily basis, a rise of 80% YoY.
- Sogou set up a joint research institute focusing on AI and machine learning with Tsinghua University in 1H16.

DiDi Chuxing 滴滴出行 (not listed)

Share price (Sep 13):	**n.a.**	52-week range:	**n.a.**
Market cap:	**USD50bn (latest financing)**	TTM R&D expenses:	**n.a.**

Largest car hailing app in China

Investment thesis

Business description

- Didi is one of the largest private Internet companies in China and the largest car hailing app in China, with 49.5m MAU in 2Q17 per QuestMobile.
- In August 2016, Didi announced the acquisition of Uber China's assets and operations. In exchange, Uber Global holds around a 20% equity stake in Didi after the merger.
- According to Caixin, Didi's valuation surpassed USD50bn after the latest round of financing in April 2017.

Revenue breakdown

- The company does not disclose revenue figures but claims its number of daily orders surpasses 20m currently.

Notable initiatives in AI

- Didi launched its intelligent transportation cloud platform in September 2016. The platform aims to facilitate users' inter-city commute and help improve cities' transportation efficiency. Several local governments, including Shenyang, Ji'nan, Wuhan, Guiyang, have struck agreements with Didi on system integration.
- Didi's intelligent transportation system has 3 elements: cloud computing, Big Data and machine learning. Didi says it has the largest set of intercity users' transportation data in China and currently processes 2,000TB data on a daily basis. According to management, machine learning unlocks the value of the transportation data, resulting in shortened ETAs, more efficient order distribution, and more accurate supply and demand forecasts.

Facebook (FB US, Buy [1])

Share price (Sep 13):	**USD173.05**	52-week range:	**USD113.6-175.5**
Market cap:	**USD502.6bn**	TTM R&D expenses:	**USD4,052m**

Global social network king embracing machine learning

Investment thesis

Business description

- Facebook is the world's largest social network site, with over 2bn MAU as of 2Q17. Its product portfolio includes Facebook, Instagram, Messenger, WhatsApp and Oculus. Facebook relies on advertising as the predominant revenue model.

Revenue breakdown

- Facebook posted revenue of USD9.3bn for 2Q17, up 45% YoY. Ad revenue contributed 98%, with a faster growth rate (47% YoY). The rest of the revenue came from payments and other fees, which consists of revenue from delivery of virtual reality platform devices and related platform sales.
- Revenue from the US and rest of the world in 2Q17 was USD4.05bn and USD5.27bn, accounting for 43% and 57%, respectively, of total 2Q17 revenue.

Notable initiatives in AI

- Facebook acquired Oculus for USD2bn in early 2014. Oculus is a VR/AR hardware maker, powered by deep learning and computer vision.
- Facebook launched "M suggestions" for Messenger users in April 2017. The program works as an intelligent assistant to recommend related services and activities to users according to their chat content.
- Facebook launched its "Camera Effects platform" at its 2017 F8 conference. Camera Effects, powered by Facebook's objects recognition capability, lets 3rd-party developers create AR masks and filters. It is expected to be compatible with AR hardware in the future.
- Facebook has shifted its words-based translation model to neural-network-based (sentence-based) with the help of deep-learning technology.

Amazon (AMZN US, Outperform [2])

Share price (Sep 13):	**USD999.6**	52-week range:	**USD710.1-1083.3**
Market cap:	**USD477.8bn**	TTM R&D expenses:	**USD17,039m**

The everything store

Investment thesis

Business description

- Amazon is the largest e-commerce platform in the US. The company provides direct sales, online marketplace, subscription (Amazon Prime, digital video etc.) and cloud services (AWS).

Revenue breakdown

- Amazon's 2Q17 net sales increased by 25% YoY to USD38bn. Direct sales, online market place services, subscriptions, and AWS accounted for 63%, 18%, 6% and 11% of total sales, respectively. AWS grew the fastest among Amazon's business lines (sales up 42% YoY for 2Q17).
- In 2Q17, North America and International contributed 59% and 30% of total revenue, respectively, (the rest came from AWS).

Notable initiatives in AI

- Amazon rolled out a smart speaker called Amazon Echo in 2014. Echo sold 6.5m units in 2016 and the number is expected by management to surpass 10m in 2017. Echo is embedded with Amazon's intelligent assistant, Alexa, supported by deep learning. The company is testing deploying Alexa to other hardware products such as Kindle.
- Amazon acquired Kiva System in 2012 for USD775m, thereafter renaming the business Amazon Robotics. Amazon Robotics is a mobile robotic fulfilment systems based on a navigation algorithm. It currently only works for Amazon's warehouses.
- Amazon Go is the first prototype of a "100% self-service" grocery store, which opened in December 2016. Shopping at Amazon Go does not require checking out as the store detects users' orders without them having to check out, utilising its deep learning and computer vision.

Alphabet (GOOG US, Buy [1])

Share price (Sep 13):	**USD950.44**	52-week range:	**USD743.6-1008.6**
Market cap:	**USD652.5bn**	TTM R&D expenses:	**USD15,332m**

Global leader on AI and machine learning

Investment thesis

Business description

- Alphabet is the parent company of Google, the largest search engine in the world, with business interests in online ads, digital content, and app distribution, among others.

Revenue breakdown

- Alphabet posted total revenue of USD26bn for 2Q17.
- Total ad revenue accounted for 87% of Alphabet's total revenue. Within ad revenue, 81% came from Google properties (Google search, Gmail, YouTube etc) and 19% from Google network members' properties (AdSense, AdMob). Google's other revenue (Google Play store, cloud, hardware etc) contributed 12% of total revenue.

Notable initiatives in AI

- Alphabet has been at the forefront of AI research since the late-2000s. Currently many of Google's products, including search, Youtube, Gmail, translation, have deployed machine-learning technologies.
- Alpha Go, a deep learning program that plays the traditional Chinese board game Go, was developed by Google DeepMind in October 2015. The program is based on neural networks and the Monte Carlo tree search algorithm. Alpha Go trains itself extensively using historical matches as a reference. It has beaten Lee Sedol and Ke Jie, the top-2 professional Go players in 2016-17.
- Its autonomous driving project (Google Waymo) was launched in 2009. As of June 2017, Google's autonomous cars had driven over 1.7m miles in testing. Car OEMs including Fiat Chrysler, Toyota Prius, Audi TT plan to equip their cars with Google's solution.
- Google launched its voice-based smart home initiative, Google Home, in May 2016.

Apple (AAPL US, Outperform [2])

Share price (Sep 13):	**USD159.65**	52-week range:	**USD104.1-164.9**
Market cap:	**USD832.4bn**	TTM R&D expenses:	**USD8,658m**

Growing in intelligence

Investment thesis

Business description

- Apple is the largest consumer electronics brand in the world. The company also distributes digital content and hosts Internet applications. Key products and services include the iPhone, Macintosh computers, iPad, iOS App Store, iOS, iTunes Store, and iCloud.

Revenue breakdown

- Apple's net revenue for 3Q FY17 was USD45.4bn (up 7% YoY).
- The iPhone contributed 55% of total sales, the iPad and Mac 23%, and services 16%.
- North America, Europe, China and Japan accounted for 45%, 24%, 18% and 8% of 3Q FY17 revenue, respectively.

Notable initiatives in AI

- Apple launched Siri (embedded in the iPhone since the iPhone 4S), an intelligent personal assistant developed by the SRI international AI centre, in 2011. Siri was later introduced to the Mac and iPad in 2016.
- Apple published its first AI thesis, "Learning from simulated and unsupervised images through adversarial training", in December 2016.
- Apple has been taking a more open approach on disclosing its AI research to the public. The company opened up a dedicated machine learning blog in July 2017 to share its latest research on machine learning.
- In June 2017, Apple's CEO, Tim Cook, said in a Bloomberg interview that Apple is working on the "autonomous systems" behind driverless cars. Aside from this revelation, there has been no public disclosure of Apple's work on autonomous systems.

Microsoft (MSFT US, Buy [1])

Share price (Sep 13):	**USD75.21**	52-week range:	**USD56.0-75.2**
Market cap:	**USD581.2bn**	TTM R&D expenses:	**USD13,037m**

Not just about Windows and Office

Investment thesis

Business description

- Microsoft is a global leading software developer and cloud service provider. The company's products include operating systems, business solutions, search engine, cloud services, etc. Microsoft also manufactures hardware including PCs, tablets and smartphones.

Revenue breakdown

- Microsoft's FY17 (ending June 2017) revenue came in at around USD90bn.
- The productivity and business processes segment, which comprises Office 365, LinkedIn, Dynamic business solutions, among others, contributed 34% of total revenue.
- Intelligent cloud, including Azure, Microsoft SQL server, enterprise services (eg, premier support services), accounted for 31% of total revenue.

Notable initiatives in AI

- Similar to Apple's Siri, Google's Google Now and Amazon's Alexa, Cortana is an intelligent personal assistant created by Microsoft in 2015. Cortana is empowered by natural language processing and semantic search capabilities developed in-house by Microsoft.
- Microsoft has acquired several AI start-ups in recent years, such as Maluuba, Agolo and Bonsai. These companies focus on language comprehension, information abstraction and integration, and open sourced AI. We believe these M&A initiatives reflect Microsoft's focus on AI research and applications.
- Microsoft revealed in July 2017 that a new research lab had been created to focus on developing general-purpose AI technology. The lab has over 100 scientists working on natural-language processing, deep learning and perception systems.

Appendix 2: selected management discussions of AI, machine learning and deep learning

AI: excerpts of management discussion on AI from earnings conference calls (2Q15-2Q17)

		2Q17
Baidu	*Lu Qi: COO*	**Our focus is to accelerate the commercialization of AI technologies.** Our best strategic path forward is to platformize our AI technologies, build vibrant ecosystems based on healthy business models and sustain a win-win situation with developers and partners. **The foundation of our AI platform is Baidu Brain plus Baidu ABC Cloud, where Baidu Brain offers developers the most complete and leading AI technology suite, coupled with Baidu ABC's growing set of AI-based industry solutions.** In the AI era, the future of the cloud industry will be increasingly driven by integrating AI with big data. We were among the first to spot that trend and seize the opportunity to pursue our ABC - AI, big data cloud - strategy by delivering a growing suite of powerful industry vertical AI-based solutions to our partners and developers, including marketing, media, IoT, video, et cetera. By fully leveraging Baidu's AI technology strengths and the vast data assets, our ABC business is growing very rapidly, securing a growing number of customers. For example, we are working with one of China's top mobile phone manufacturers to provide them with a 360-degree solution from marketing, to deep learning, to intelligent CDN. **One of the most promising areas for commercialization of AI technology is financial services. At the Baidu FSG, our focus is AI-enabled fin-tech, by leveraging Baidu's technology strengths and data assets to build a suite of leading AI-based FinTech capabilities. Our long-term objective is to build out a platform that enables Baidu and our partners to provide much more compelling financial services.** With regard to AI as it applies to IoT as you asked, our approach is using our DuerOS platform. The DuerOS platform will enable each IoT device, particularly intelligent home devices, to be able to listen, to talk, to understand users to provide services. And economically, there are 2 future functions. One is these IoT devices that are enabled by DuerOS, they become the future information entry points. **The economics will be very similar to today's search business because as long as the people are seeking information, interacting with the information that represent services, content, knowledge, there will be natural expansion of economics like today's search engine business.** That's one. The second is because DuerOS is a platform that enables devices, there's typical what's called platform economics. You mentioned **per-device licensing**, that's one way to do that. Windows, for example, can do that. Or there's **channel distribution economics** that's in App Store. And the third part is future VIGs you can take. For example, like App Store, when you sell content through app stores like our iOS, Apple takes 30% profit. So DuerOS, in the long run, represents very compelling economic opportunities for Baidu. Our focus is really execution, building out our products, our platforms, a very healthy ecosystem. And there are many, many important execution paths in front of us over the coming years. We are fully prepared, when these economic opportunities become more material, we will be able to share future operating metrics when we are ready.
Alibaba	*Daniel Zhang, CEO*	The other effect of increasing clicks is what Maggie discussed in the **personalization where we use AI technology to improve the content that people see, so that they are more personalized to you, the shopper and that also increases the clicks.**
Tencent	*Martin Lau, President*	We have been **investing heavily in AI but relatively quietly**, as we view AI as an essential capability that enhances the user experience and empowers us to capture the new exciting opportunities to grow our businesses for the future. We're confident that our existing strength in computing power, data, engineering, technologies as well as use cases coupled with our proactive build-up of AI content -- talent will give us a favourable position in this strategic initiative. **In particular, a wide and diversified business scope creates a variety of use cases for AI research and application across a range of AI fundamental research areas, such as machine learning, computer vision, speech recognition and natural language processing. We will be persistent but patient with our AI investment, because we believe it is a long-term initiative, and we do not necessarily require a research to generate revenue directly in the short term. On the other hand, AI will significantly benefit all of our existing products, services and businesses in many ways.** For example, **in consumer-facing products, AI enhances user experience as we understand more about the users; for enterprise related businesses, AI optimizes monetization as we sharpen our targeting technology**; for our ecosystem, our investee companies and could partners can leverage our strong AI capability, allowing all of us to achieve mutual benefits. ...These three examples are **performance ads, information-based services, and FinTech businesses.** For performance ads, we applied AI technology to the processes of ad placement from understanding users' preferences, contextual and ad content, to ranking, the bidding price, optimizing the display formats and eventually to matching the most appropriate advertisers. This increases the ROI for advertisers while at the same time enhancing the reading experience for our users. For information-based services including news apps, video, music and App Store, AI enables us to have a better knowledge of users' interest. This will help us to make more relevant and customized recommendation to users, so that they can access their favourite content more efficiently. Across, Tencent platforms, there are multiple digital content access points which we believe will all benefit from the smarter recommendation engine. For Internet finance businesses including mobile payment, wealth management, and microloans, we use AI to predict users' behaviour in financial activities more precisely. This will help us to provide the most suitable products to the most appropriate users and in the process, minimize the risk involved. Early this year, Fine Art won the UEC Cup, a global Computer Go tournament. Fine Art was developed by our AI lab in less than a year. We have accumulated in this process significant know-how in the development of Fine Art. And the strategy and **reinforced learning AI technology behind Fine Art can be applied to many other use cases in the future.** Our face recognition technology also scored excellent results with the world-leading Face Detection Dataset and Benchmark, FDDB. We have gradually applied this technology in a variety of different ways. For example, firstly, enhance and enrich the features of our photo editing app Pitu, which has become the number two app of its kind in China. Secondly, enable users to complete ID -- identification online for financial services and government, municipal affairs. Thirdly, assisting the search for children and elderlies report missing, helping many families in the process.

Google	*Sunder Pichai, CEO*	**People are no longer only using a keyboard, mouse and multi-touch, but are also using emerging inputs like voice and camera to ask questions and get things done in the real world.** We are seeing this in the way people interact with the Google Assistant, which is already now available on more than 100 million devices since launching last year and there is more to come. Since we released an Assistant SDK that will enable a wide range of new hardware devices, which will include the Google Assistant, we now have more than 70 home automation partners on the Assistant on Google Home and Phones, including Honeywell, Logitech and LG. So you can do everyday things around the house using your voice. At Google IO, we announced **Google Lens** available later this year. Lens is a set of vision-based computing capabilities that can understand what you're looking at and help you take action based on that information. So for example, if you saw a poster for your favorite band, you would be able to take a picture and get relevant information and buy tickets to their next concert. In Search, a great feature we launched this quarter in the U.S. is job search. To make looking for a job easier for everyone, no matter what line of work you are in. Many of these products that make people's lives easier are being powered by machine learning. **One focus area for us this quarter has been enabling our machine learning algorithms to learn and improve our products much faster.** One such research initiative auto ML enables us to pursue approaches to automate the design of machine learning models. **Our ability to rapidly deploy the best machine learning in all of our products enabled us this quarter to launch all sorts of new smart features, to help moderate comments, suggest smart replies in Gmail and improved translations.** Turning to our advertising platforms. **Here too machine learning is critical to helping advertisers and app developers analyze data in real time to reach consumers with more useful ads and measure campaign effectiveness**. At Google Marketing Next this quarter we launched Google Attribution, a comprehensive measurement tool that allows marketers to measure the impact of their campaigns across devices and channels all in one place with no additional cost.
Facebook	*Mark Zuckerberg , CEO*	I am excited about how AI will improve people's experiences across our products. **We're finding AI is both delivering consistent improvements to many of our systems, like News Feed, search, ads, security, and spam filtering and more. But more than just improving these existing experiences, I expect AI to change the way that we do business in some important ways.** So for example, today to keep our community safe, we rely on people flagging content that might violate our community standards for us to review. In the future, AI will be able to help flag more of this content faster before people have even seen it. Now we've started using AI to fight terrorism and keep propaganda and extremist accounts off Facebook. We've even started experimenting with using AI to understand texts that might be used to promote terrorism. When it comes to News Feed, we currently mostly show you content from people and pages you're connected to. And we can write this better with algorithm improvements, but **the really big improvement from AI will be when we can understand all the other content that's out there so we can help you discover much more of what matters to you beyond just what your friends are up to.** The Internet gave people the power to target their messages to people who actually might be interested and to measure results much more precisely, and that was a big improvement. And now AI is taking this a step further. Now you can put a creative message out there, and AI can help you figure out who will be most interested. A lot of the time you don't even need to target now because AI can do it more precisely and better than we can manually. This makes the ads that you see more relevant for you and more efficient for businesses. Those are just a few of the reasons why I'm optimistic about how AI is going to improve our core services over the next few years.
1Q17		
Baidu	*Robin Li, CEO*	Search is one of the first applications of AI where search works to understand human. Moreover, we live in an age where we can finally match this data with powerful efficient computing infrastructure, and a deep bench of talent to make real progress in AI. We believe Baidu's opportunity to apply AI to real world consumer and enterprise applications in China is enormous. We envision a world where machines continue to learn about humans. **In the not so distant future, using natural voice input or images to interact with machines will be mainstream, alleviating humans of the burden of mundane tasks.** We imagine a world where your face is your identity and you only need to show your face to pay at the supermarket or goes through airport security. Cars will drive us, freeing up our time and attention. As machines become more like hyper efficient, real-time assistants for people, the mode of user interaction is evolving correspondingly. **Just as the mouse and keyboard were the gateways in the PC era and touch is the gateway for mobile, voice recognition through conversational DuerOS will be the gateway in the China plus AI era.**
	Qi Lu, COO	Our focus is two-fold, to enhance and expand our existing core business platform, and to **lead the commercialization of AI technology across a number of AI-enabled new business initiatives, such as our AI-cloud, financial services, DuerOS, and autonomous driving.** We are also actively developing and **applying AI technologies across a number of key areas such as ad targeting, creative and landing page optimization, which will contribute to increasing monetization capabilities over time.** Now turning to AI-enabled and new business initiatives. We are making good strides across the board, and here a few specific areas to call out. **DuerOS is our next-generation conversational AI platform.** By design, it provides powerful capabilities to enable the use of voice in natural language to interact with any digital experience on any device, for any scenario. This is because for the first time in our history, AI technology enables natural language to serve as a general purpose user interface. While it is still early, this platform will have enormous potential relative to the traditional platforms, because it can run everywhere, and it enables users to interact naturally and easily. **We believe we are the clear market leader in China. DuerOS has already gained great momentum in building out its ecosystem.** Home and cars are two high value scenarios to anchor the DuerOS platform at this early phase. Here we have made important investments in our Raven Tech acquisitions, as well as our Internet of Vehicles business unit. Both of which will help further establish the DuerOS platform and strengthen our market leadership position. At the same time, both in their own right will have significant business opportunities that we will grow into. In Q1, we consolidated our L4 and L3 and Internet of Vehicle Business Unit into the new **IDG, Intelligent Driving Group**. This will enable us to build out a common technology platform, shared services, and a consistent solution for all our partners. At the Shanghai Auto Fair last week, we announced project **Apollo, an open, complete, and reliable autonomous driving technology platform that we plan to launch in July.** Our AI-cloud business named as **ABC** or **AI Big Data and Cloud** has been gaining traction with our customers, who appreciate our AI value-added approach. Baidu Cloud builds upon foundation of cloud solutions of computing, storage and CDN. We differentiate by providing AI technology enabled solutions, such as Big Data, deep learning, video, IoT and more. Customers across media, financial services, healthcare and other verticals have expressed very positive feedback. This month at the **Baidu Cloud ABC summit**, together with NVIDIA, we announced the establishment of the Baidu deep learning platform which will greatly help our customers and partners to enhance their capability in deep learning and AI. ...So, Jennifer, what I would just add to emphasize the AI investment part … I would like to emphasize is, **we will be very disciplined and systematically optimize our resource allocations** to make sure that our engineering investments and R&D investment is truly tap into the substantial growth opportunities ahead of us.
	Jennifer Li, CFO	**This year, we expect incremental R&D expenses over 2016 to be all related to AI. Incremental R&D spend this year is largely related to AI R&D head count.**

Alibaba	*Joe Tsai, Executive Vice President*	**On the autonomous driving question, I think you have to understand, a lot of technology companies are doing this. Obviously the long term commercial opportunities are very murky, nobody has figured out the long term economic model for this,** but people are doing it because there are some very interesting artificial intelligence related technology in an autonomous-driven car that gets all the technology, company is very intrigue, things like computer vision, LIDAR technology, simultaneous localization and mapping and all these technologies are very interesting and that's why companies are all investing in those technologies and the same is true in the case of Alibaba. **We may not be investing those technologies for to create a driverless car, but we are investing in all those component technologies for other applications.**
	Daniel Zhang, CEO	**We will continue to invest significantly in AI technology such as voice recognition, machine learning and a natural language processing for application to our technology in real use cases, such as further discovery, personal recommendations and customer services. It will contribute to meaningful improvements in operating efficiency across all our platforms.**
Tencent	*Martin Lau, President*	As it relates to AI, we view AI as one the core capability for us to continue to build our own businesses. **AI actually touched a point lot of our existing business in the area of advertising for increasing our targeting capability in the area of our newsfeeds and information to allow us to customize and personalize our offerings to our users** in the area of online connections to help us to target users in the right in there with the right profile so by developing artificial intelligence and these core technologies will benefit our core business. AI can also help us to get into new businesses in the future and we are actively developing these capabilities. And finally AI as we continue to build our cloud business I know will become a core capability that will be opening up to our cloud partners as well as our ecosystem partners. So that's sort of you know the way we look at it. (*On why Tencent invested in Tesla*): As it relates to test lighting sort of it's somewhat related to AI but what we feel is that sort of the automobiles as they become more and more connected with the Internet as well as, as it becomes smarter in terms of control and autopilot, assistant pilot and over time autopilot. That automobile is becoming a smart device and there will be much more connection between the physical world and the virtual world. So that's the reason why we felt we want to partner with the leading company in such field right now in order for us to get on with these capability as well as to figure out whether we can learn something new or whether we can actually build some businesses together.
Google	*Sundar Pichai, CEO*	turning first to machine learning and access to information. I'm really happy with how we are transitioning to an AI-first company. The **Google Assistant** is one of our first steps towards that future. This quarter, we brought the Assistant to your wrist with Android Wear 2.0, announced that it'll soon come to Android TV, and began rolling it out to hundreds of millions of people with Android phones running Nougat and Marshmallow. We also continue to announce new integration partners for the Google Assistant, so you can now ask it to do things like dim the living room lights. Stay tuned for more to come. Advances in machine learning are helping us make many Google products better. One example from this quarter is the launch of parking predictions in Google Maps. Beyond that, we continue to set the pace in machine learning and AI research. **We introduced a new technique for training deep neural networks on mobile devices called Federated Learning.** This technique enables people to run a shared machine learning model while keeping their underlying data stored locally on mobile phones. DeepMind is bringing AlphaGo to China in May, pairing AI with the world's best players at the Future of Go Summit. And it's very exciting to see the incredible things that developers and researchers are building on the TensorFlow platform. There are more than 6,000 GitHub projects and counting. We also demonstrated new machine learning services and announced our acquisition of **Kaggle, the largest community of data centers focused on machine learning and AI.** ...And that's where Daydream is working well. We are learning how to write great content on top of VR. YouTube, Google Earth VR, Tilt Brush – these are all great examples. So underlying technology wise, be it all the kinds of sensors and tracking you need to do, the machine learning and AI you need, things like voice recognition, everything – it's just a computing evolution, so everything we are investing today in machine learning and AI, as well as what we're doing in our computing platforms, transitions well. So I think we're already well-positioned to play this, and so I think we'll thoughtfully approach it. And we'll approach it more holistically.
Facebook	*Mark Zuckerberg, CEO*	AI tools over time will be able to do a better job of flagging things for the set of people who are in the Community Ops teams that we can prioritize what we look at. A lot of what we're trying to do here is not just about getting content off Facebook. Last week there was this case where someone was using Facebook Live to broadcast – or was thinking about suicide. And we saw that video and actually didn't take it down and helped get in touch with law enforcement who used that live video to communicate with that person and help save their life. So a lot of what we're trying to do is not just about taking the content down, but also about helping people when they're in need on the platform, and we take that very, very seriously. Over time, the AI tools will get better. **Right now there are certain things that AI can do in terms of understanding text and understanding what's in a photo and what's in a video.** That will get better over time. That will take a period of years, though, to really reach the quality level that we want. So for a while, our strategy has been to just continue building as good tools as we can because no matter how many people we have on the team, we're never going to be able to look at everything. So that's going be a big challenge. But given the importance of this and how quickly live video is growing, we wanted to make sure that we doubled down on this and made sure we provided as safe of an experience for the community as we can, which is why we're almost doubling the size of the Community Ops team to focus on some of these issues around safety on live video. But over time for sure, more AI will do this, but this is over a period of years.
4Q16		
Baidu	*Jennifer Li, CFO*	As Robin mentioned, the mobile transition is behind us, and we're at the cusp of an AI revolution. Baidu's strategic focus, organization, and resources have shifted increasingly towards AI, and we are excited to execute our vision in this new era.
	Robin Li, CEO	Yeah, on the hardware-software integration, I think this is one of the most important characteristic of AI era. With this kind integration, lot of innovations become possible, and the user experience is significantly improved based on that. **That's why we are investing in both hardware and software for AI-related technologies. That its true for the mobile platform, but it's also true for the home environment as well as auto environment.** For example, when you're at home, we mentioned we have collaborated with (55:15) when you're watching a TV show, you can instantly ask "who is that actor" in natural language will be able to find the answer for you. And in car environment, we have CoDriver, we have CarLife, which integrate a lot of Internet and AI voice-based command in those environment. That's just a couple of examples, when you integrate software with hardware, what could be possible going forward.

Alibaba	*-*	-
Tencent	*Martin Lau, President*	In terms of AI, we actually **view AI as a core technology across all our different products.** So, what we're trying to do is actually in each one of our businesses, we encourage our team to apply, to build up talent pool as well as to apply the core technologies around AI, machine learning and deep learning. And at the same time, at the corporate level, we also established an AI lab, which is more research based, so that they can actually focus on more basic building blocks around AI. And if you look at sort of the **areas of AI that we are investing in, it would include areas such as speech recognition, it would include picture and photo recognition, computer vision, natural language processing and all sorts of deep learning, as well as basic architecture for deep learning platform.** And the way, we look at this technology is also that in order to build long-term competitiveness, you don't only need sort of the people, as well as sort of the mathematical and computational expertise, but at the same time, you also need a lot of data, **you also need a lot of usage scenario, so that you can actually apply these AI technologies,** and as a result you can actually keep progressing. So, I think if you look at our actual products, around content recommendation, around our advertising, around our photo processing app, there are actually a lot of existing products which will benefit from AI. We are also sort of investing in more pure research projects. You may notice that our **Go chess player, a Fine Art, has recently won the championship in UEC competition.** That's an example of our research project. And at the same time in the future, we believe that AI technology can also allow us to explore new areas such as personal assistance, such as maybe even sort of autonomous driving. So, these are the areas that we will focus on.
Google	*Sundar Pichai, CEO*	As I've shared before, **computing is moving from a mobile-first to AI-first** with more universal ambient and intelligent computing that you can interact with naturally, all made smarter by the progress we are making with machine learning. A centerpiece of our machine learning efforts is the **Google Assistant** which allows users to have a natural conversation with Google to help them get things done across their experience. It's off to a great start. You can easily ask it to navigate home, tell you about your schedule for the day, or even play trivia. We reached a milestone last month with our announcement of the Assistant developer platform called Actions on Google. It gives developers like Uber, SongPop and Headspace the opportunity to build conversation actions for Google Home, and we'll expand it even further this year. The Assistant is baked into our smart messaging app, Allo, which we expanded this quarter in languages like Hindi, Brazilian Portuguese and Japanese. This quarter, using **neural machine translation, we have improved our translation ability more in one single leap than all our improvements over the last 10 years combined.** We'll be rolling neural machine translation out across to more than 100 languages available in Google Translate in 2017 and also for all of our Cloud customers through the Google Cloud Translation API. And Google Photos, which as you know has machine learning at its heart, continues to grow in popularity. Last quarter, we launched **PhotoScan**, which helps you digitize all of those old printed photos that are probably stored in a shoebox in your closet, to keep them safe, organized and shareable. In February, we are hosting a summit where our machine learning team and other experts will discuss the future of the TensorFlow open-source initiative and share some of their latest demos.
Facebook	*Mark Zuckerberg , CEO*	On artificial intelligence, we developed a new technique called **style transfer** that uses AI to study a painting and then can take your photos and videos and draw them in that style in real time on your phone. And if you post on Facebook looking for a place to eat or suggestions for where to go, we can now use AI to understand the text of your post and understand what you're asking and surface (07:20) recommendations from the comments. Well, I think AI is going to be great for the experience that people have in our community. So there are a few types of systems here that we're working on around understanding content. One is around visual content and the other is about language. **So, for visual content, we want to be able to look at a photo and understand what's in it, right,** and whether that's something that you're going to be interested in, right? And similarly, we want to be able look at a video and watch it and understand whether that's something that you're going to be interested in. And you can imagine that today, we consider putting things in your News Feed that you're connected to in some way, right, that are from a friend or a page that you're following or that one of your friends likes. But there's no reason that we shouldn't be able to match you up with any of the millions of pieces of content that you might be interested in that gets shared on Facebook every day, except for the fact that we don't have the AI technology to know what those are about and if they match your interest today. So, a combination of being able to understand the text that people message, read the articles that people would want to look at, watch the videos, look at the photos, are going to be great, too.
		3Q16
Baidu	*Robin Li, CEO*	As you know, Baidu has been investing in AI for years, and we are proud to be recognized as one of the leading AI innovators globally. Our AI infrastructure is already helping to make all of our products more efficient and useful. For example, **AI enables better filtering of bad search results and keywords. Helps generate more intelligent and targeted search results and news feeds. Automatically detects non-compliant images on Nuomi, enhances the accuracy of Baidu Deliveries' delivery time estimates, and predicts traffic conditions in maps, among other things.** **AI also enhanced app displays and broadened the range of app formats,** driving further value to our customers. And our amazingly accurate speech and image recognition have become the envy of the industry, helping to drive further revenue growth in voice and visual search queries. Our AI powered Baidu Translate product is the first online translation service to combine deep learning with a statistical model. With leading voice and image input and recognition, **Baidu Translate** services 28 languages and processes over 100 million requests a day on average. Many of which come directly from search and supports over 20,000 third-party applications. Beyond our core products we are already applying AI and big data to many areas of people's daily lives. For example, Baidu **AI medical-assistant Melody** helps doctors to diagnose illnesses from online conversations with patients. We also recently began to test a facial recognition-based entry control systems here on Baidu Campus. **And beginning of last month, we open sourced our AI platform PaddlePaddle to developer community providing access to Baidu technology in the areas of voice and image recognition, natural language processing and machine learning.** Autonomous driving is an important area for future growth where AI technology has helped us take an early lead. Enabled by Baidu Brain's powerful image recognition and processing technologies are autonomous driving car recently broke a number of records in the **TI/TTI** benchmark test, achieving an accuracy rate of **90.13%** in the vehicular third recognition test. Last year, we began low test of **autonomous driving** under real-world conditions in China. And in August, we received our permit to begin testing Baidu's autonomous cars on the public road in California. We have an ambitious development schedule, but with our top team here in Beijing and in the U.S., leading technology and unrivalled data resources, I'm confident that we will lead the way in making autonomous driving a reality. ...And Alan, how **AI helped our CPC and CPM**, we started to investing AI technology about five years ago. And starting from about three years ago, we implemented AI-related technology to help monetization. It has contributed to CPC and CPM, and click through rate for quite some time, I would say, at least for the past about three years, we've seen very significant contribution from our AI technology monetization, and I believe there's still lots of room for improvement on that direction ...On AI technology, I think, it has helped almost all of our existing products, **helped to monetize**, helped to deliver better user experience in search, in many other **news feed**, and many other content-based products. But also, I think for our newer initiatives, be it **Internet finance**, or **autonomous driving**, we've seen significant contribution from AI-related technology, I think that, well-positioned us uniquely in those very large markets. But the newer initiatives generally take longer to see material impact on revenue and earnings. So, we'll need to wait for maybe a couple of more years.

Alibaba	***Daniel Zhang, CEO***	Over 40,000 developers from 58 countries attended our annual developers' conference for cloud computing and artificial intelligence in Hangzhou, where we unveiled a **City Brain** project in collaboration with Hangzhou government to upgrade and transform citywide management and services.
Tencent	**-**	-
Google	***Sundar Pichai, CEO***	...I don't think I can remember a busier time in Google's history than the past few months. We introduced the world to the **Google Assistant**, powered by our state-of-the-art machine learning technology. We unveiled a beautiful lineup of new hardware devices, including **Google Home** and the **Pixel phone**, which is getting great reviews. We launched a new messaging app, **Allo**, and a new video chat app, **Duo**. It's been an incredible quarter. Through it all, we are laser-focused on the priorities I outlined earlier this year, and in the next few minutes I'll quickly go through our progress this past quarter. First, making knowledge and information accessible to everyone, we have long aimed to help people find the information that they want right when they need it and sometimes even before they think to ask for it. We reached an important milestone in that journey this quarter with the launch of the **Google Assistant**, which allows you to type or talk with Google in a natural conversational way to help you get things done regardless of the device you're using. We first introduced it in our smart messaging app, Allo, which lets you chat with the Assistant. You can even add it to a group conversation to help you and your friends decide where to go for dinner. **Early adoption of Allo and Duo has been great and has exceeded our expectations.** **And the Assistant is core to the Pixel phone and Google Home.** With Google Home, you can ask the Assistant to cast a song to your living room speaker or turn on the lights in your kitchen. And with Pixel, by simply holding the home button, you can ask the Assistant to send a text to your mom or pull up pictures of your cat in Google Photos. ...The Assistant all possible thanks to our years of investments in computer science and machine learning. Our knowledge graph now understands over 70 billion facts about people, places, and things in the real world. And just last month, we announced our latest **research on neural nets**, which has given us a huge leap in translation quality. This breakthrough will help us provide even more accurate translations for people around the world. ...Our investments in machine learning continue to be a very clear advantage for Google Cloud, and we are helping customers apply ML in very concrete ways. For example, **Ocado**, the world's largest online-only grocery retailer, is using machine learning to categorize and prioritize customer emails, which will help their support center respond to customers faster. ...On machine learning, in areas like search and even ads, on search we've had an effort called **RankBrain, which is bringing Google Brain in the context of search.** And we've made great progress on it, but I would still characterize it as very early stages in terms of the long-term impact we can have. Generally and similarly, we are in very, very early stages of incorporating machine learning, the newer machine learning systems in ads. And again, that's the beginning of a long journey as well. Overall, I think all of these systems are incredibly complex systems, and they are handcrafted systems over many, many years. And so over time, I think machine learning will surface newer approaches and newer insights. And so we see it as a huge area of opportunity, but it will play out over a period of time.
Facebook	***Mark Zuckerberg, CEO***	On artificial intelligence, **we're starting to see the impact that AI can have on enhancing people's experiences on Facebook and showing them more of what they care about.** More than 40 teams at Facebook and more than 25% of our engineers are already using AI to power the products and services they build. We've made changes and improvements to our AI in order to filter out misleading clickbait stories from News Feed. And we're using AI to help find terrorist propaganda on Facebook. It's still early, but we think that AI will help improve the quality of what people see and can share on our platform. I'm going to talk about our AI work overall. In general, what we want to do is **try to understand the content that people are sharing and that's out there for them to see as best as possible**. So for example, that means being able to read and understand news articles or posts that people make or messages that a person might send to a business so we can help that business auto-reply to them and get information back to the person really quickly, or understand the content so we can better understand what might be interesting to a person and show it in News Feed. Similarly, aside from the **conversational and linguistic understanding,** there's a whole thread of the work that we're doing on **visual understanding**. So understanding photos, what's in photos, what's in videos, what people are doing. That allows us to not only do things around accessibility, to show somebody who is visually impaired to be able to read to them what might be in a video or a photo. **But it also helps us rank News Feed better, so that way we can help understand what is in the content and show people more of what is going to be meaningful to them.** It helps us identify content that might be offensive or graphic that might violate the policies of Facebook so we can flag that and review that better. So that's the main thread of the AI work we're doing is trying to understand conversational and linguistic context and computer vision, photographic and video signals to understand what's going on there. And that will apply across basically every product that we build at the company. And then of course, there's some deeper AI research that we're doing that feeds into those applications as well. That can apply to things like ranking for News Feed and search and ads and all of our systems more broadly.

		2Q16
Baidu	***Robin Li, CEO***	Our **Baidu Brain** AI infrastructure supports all of Baidu's product lines. Baidu Brains have already noticeably **helped improve search modernization and also improved the accuracy of voice and image recognition which helped to double voice and image queries year-over-year.** Machine learnings have enabled the automatic drafting of ad titles based on user profiling and user queries. Making what used to be a tedious, manual process for customers, much more intelligent, streamlined and targeted. Through machine learning, we've helped our customers automatically create and enhance their mobile landing pages, based on their PC websites and keywords. The technology also **filters out low quality results and predicts click-through rates for sponsored links.** Baidu Brain not only benefits our core-search business, it also opens up new opportunities in areas such as financial services, where we see great potential in education loans, online banking and online insurance. Our DU virtual assistant which marries voice recognition with services and autonomous driving which will transform the transportation industry
Alibaba	-	-
Tencent	-	-
Google	***Sundar Pichai, CEO***	At I/O, I was excited to share our vision around the **Google Assistant**. We want to help people get things done in a conversational way across devices and context. This is possible thanks to advances in machine learning, voice and image recognition, and natural language processing, which we have invested in for years. You will be able to experience the Google Assistant in products like our **new messaging app, Allo**, and our voice-activated device, **Google Home**, and we can't wait to show you what else we are working on in this area. Second, machine learning; as I said, machine learning is the engine that will drive our future, and it's already making our products better and helping users every day. In fact, more than 100 teams are currently using machine learning at Google, from Street View to Gmail to Voice Search and more. For example, in Search, we use a ranking signal called RankBrain, which relies on deep learning to improve results. It's already enhancing the search experience in 40 languages. **And based on user testing, RankBrain can accurately guess which results users will favor with about 80% accuracy.** Advances like this help us make our search results even more relevant. Machine learning is also creating an impact in other ways. **Just last week, we announced a test that applied DeepMind's machine learning to our own Google data centers, resulting in up to a 40% reduction in the energy we use for cooling.** This will greatly improve efficiency. And when we publish our research, we hope it will enable others to reduce emissions, too. As I've said, machine learning has been a major focus and a key differentiator for Google, and that's true for our Google Cloud customers as well. This quarter, we introduced **Tensor Processing Units, or TPUs**, which can deliver an order of magnitude better optimized performance per watt for machine learning projects. Google DeepMind's AlphaGo was powered by TPUs, enabling it to process faster and look farther ahead between moves. We are now passing this benefit on to our enterprise customers to supercharge their machine learning applications. Just last week, we introduced **two cloud machine learning APIs for speech and natural language** to help our enterprise customers convert audio to text and easily understand the structure and sentiment of the text in a variety of languages. ...So we definitely – **a core part of improving the user experience for the next many years will be driven by machine learning.** To give a specific example, you asked about YouTube. YouTube watch time is something we focus on. **We use machine learning increasingly to give recommendations to users.** So when they're watching a video in terms of what they could watch next, and that recommendation system is increasingly using machine learning, and that directly drives engagement. So there are several such examples I can give across the company.
Facebook	***Mark Zuckerberg, CEO***	We've also been making progress with our initiatives around artificial intelligence and virtual reality. This quarter, we announced **DeepText, a deep learning based engine that can understand the context of several thousand posts per second across 20 different languages.** This is a long-term project, but it also has some near-term benefits like helping show people more of what they want to see and filtering out less of what they don't want to see.
		1Q16
Baidu	***Robin Li, CEO***	Our multiyear investment in artificial intelligence, and particularly in deep learning, continues to pay dividends across many products. **Baidu Brain** supports all product lines and now enable us to offer a more relevant result to users, **higher click-through rates to customers, faster delivery time for Takeout Delivery, highly naturalistic text-to-speech features** for news users, just to name a few. Deep learning will be one of Baidu's key technology advantages, we move forward – as we move forward in autonomous driving. **We believe that automobile is the next major computing platform and that advances and innovation will happen quickly, especially here in China, where the problems that autonomous driving addresses are so pronounced.** We are aggressively beefing up our R&D efforts in this area, both here in China and at our U.S. R&D center in Silicon Valley. Baidu's returns in high definition mapping and in deep learning-powered computer vision will prove to be decisive advantages and we are on track now to **deploy autonomous vehicles powered by Baidu technology in 2018, with production at scale by 2020.** We have invested in AI for many, many years. We believe we are leading in this sector not only in China, but around the world. This enables us to do disruptive things like autonomous driving. And it's kind of early for us to talk about this model because right now our focus is to solve those technical problems; the first to make that autonomous driving fully autonomous. **Our target is not to drive – to self-drive like in 90%-95% of the cases, but to really release the driver from the driver seat. So, this will take a couple of years. We will worry about the business model later on.**
Alibaba	-	
Tencent	-	

Google	***Sundar Pichai, CEO***	One of the key ingredients behind this push towards greater assistance is AI. We have long invested in building the best machine learning team and tools, and we are seeing these efforts bear fruit in many ways. As many of you saw last month, **DeepMind's AlphaGo** has been making great strides. It was a privilege to play legendary Go player, Lee Sedol, in such an important milestone for artificial intelligence. This is another step to what's creating AI that could help us with everything from our daily tasks to potentially even bigger challenges like climate change and cancer diagnosis. **At Google, machine learning is already helping us improve our products every day in search and many other areas like photos, maps and more.** Last quarter, I talked about how machine learning helps Smart Reply suggest responses in Inbox. This quarter, we launched Goals in Google Calendar, an intelligent feature that helps users make the most of their time. You just add a personal goal, like run three times a week, and Calendar will help you find the time, then help you stick it to. There's still a lot more that we can do to make Search and other Google services more assistive and helpful to you. You'll see a lot more from us this year. On the first thing, obviously, we are doing many things, but I tend to spend my time on the core of, you know, our core product. I think we have a unique opportunity to evolve search to be very assistive in how we serve our users and be an intelligent assistant that helps users throughout their needs in context, especially in the context of mobile. And overall, I do think in the long run, I think we will evolve in computing **from a mobile first to an AI first world.** And I do think we are at the forefront of development. So we don't view it as adapting to it as much as pushing hard and getting there. And so that's the core of what we do, and we'll continue to do that.
Facebook	***Mark Zuckerberg, CEO***	Artificial intelligence is a long-term effort for us, but we're already using it in lots of ways. Right now, **our Moments app is using face recognition to help you share pictures with your friends. We're using AI to show the most relevant content in news feeds, filter spam and messaging, and even help blind people understand what's in their friends' photos by reading explanations of them aloud.** ...So the biggest thing that we're focused on with artificial intelligence is **building computer services that have better perception than people, so the basic human senses like seeing, hearing, language, core things that we do.** I think it's possible to get to the point in the next five to 10 years where we have computer systems that are better than people at each of those things. That doesn't mean that the computers will be thinking or be generally better, but that is useful for a number of things. So for example, I talked about earlier, we are building this Moments app, so that way you can take photos on your phone. And if you use this app, our face recognition can look at the photos that you take and suggest that you might want to share photos that you took with a friend in them with that person. So that way, all the photos that might be of you and your friends' camera rolls, they can share with you. Another example is just spam filtering and just making sure that we can actually read the content and understand what's interesting to you or not and not show that. One obvious thing I think over time is if you just look at the way that we rank News Feed, today we use some basic signals like who you're friends with and what pages you like as some of the most important things for figuring out what – out of all of the millions and millions of pieces of content that are on Facebook, what we're going to show and what are going to be the most interesting things to you. **That's because today our systems can't actually understand what the content means. We don't actually look at the photo and deeply understand what's in it or look at the videos and understand what's in it or read the links that people share and understand what's in them, but in the future we'll be able to, I think in a five or 10-year period.**
		4Q15
Baidu	***Robin Li, CEO***	Our AI scientists at Baidu Research use deep learning algorithms and massive neuron networks to achieve breakthroughs in speech recognition. We lead the industry in speech and image technology and saw voice and image activated traffic across our product far more than double in the fourth quarter compared to the year prior.
Alibaba	*-*	
Tencent	*-*	
Google	***Sundar Pichai, CEO***	On AI, we are obviously seeing incredible progress in this field. We make great strides, and we start taking it for granted and look to the years ahead. It's always tough to predict what happens over a five-year timeframe, but I do see us making significant strides. Even a year ago, I wouldn't have predicted that we would be in a strong position to mount a serious challenge to the world champion in Go this year. So looking at the pace of progress, I think we will have AI in a form in which it benefits a lot of users in the coming years, but I still think it's early days, and there's a long-term investment for us. **Machine learning is making our products smarter and more useful for people every day.** It powers Smart Reply, in Inbox by Gmail, which suggests short responses that are relevant to an incoming email. People absolutely love this feature. In fact, in just a few months, Smart Reply already makes up 10% of all mobile responses in Inbox. **From image recognition in photos to eliminating email spam to better translation and voice recognition, our machine learning efforts are making a huge difference for people every day.** In sum, we believe we've only scratched the surface of truly being there for our users anytime, anywhere, across all devices, giving them the assistance they need. We are excited to make even more progress on this vision in 2016.
Facebook	***Mark Zuckerberg, CEO***	With our work in AI, we continue to make progress towards the new generation of computers that can see and understand. Over the last year, we published dozens of research papers including some of the leading work on image recognition and language understanding. To drive the entire AI community forward, we've also open-sourced software and a lot of new AI hardware platform. Achieving the scientific breakthroughs to build AI that makes a dramatic visible difference in people's lives is going to take a long time. But already we're seeing opportunities to serve our community. We recently built a prototype AI system that combines language and vision comprehension. You can show it an image that it's never seen before and it can answer questions about that image. And we've even used AI to help blind people experience their friends' photos and news feed by describing the scenes.

3Q15		
Baidu	***Robin Li, CEO***	At our Baidu World Conference in September, we unveiled **Duer**, our intelligent voice-based assistant that can do things for users in the real world, from making **dining recommendations, finding discounts and ordering food, to buying movie tickets, to answering sophisticated questions.** Duer represents a kind of sneak peek into the future of search, where voice activation, artificial intelligence and services all come together.
Alibaba	-	-
Tencent	-	-
Google	***Sundar Pichai, CEO***	I also want to **point out that our investments in machine learning and artificial intelligence are a priority for us.** Machine learning has long powered things like voice search, translation, and much more. And our machine learning is hard at work in mobile services like Now on Tap, which quickly assist you by providing additional useful information for whatever you're doing, right in the moment, anywhere on your phone. If you're an Android user that runs Marshmallow, try it out by long pressing the home button, when you're in the Map, it's very cool. **Machine learning is core transformative way by which we are rethinking everything we are doing.** We've been investing in this area for a while. **We believe we are state-of-the-art here. And the progress particularly in the last two years has been pretty dramatic. And so we are - we are thoughtfully applying it across all our products, be it search, be it ads, be it YouTube and Play et cetera.** And we are in early days, but you will see us in a systematic manner, think about how we can apply machine learning to all these areas.
Facebook	-	-
2Q15		
Baidu	-	-
Alibaba	-	-
Tencent	-	-
Google	***Omid Kordestani, Senior Vice President***	To call out just one, the new Photos app is a great example of how we're combining the power of a machine learning, search and beautiful design, something that Google is proud to bring together in such a great package for our users. Our investments in machine learning have applications across the company and have real potential to improve our business. For instance, machine learning helps us to improve and adds visual design to maximize its performance. There's enormous opportunity here.
Facebook	-	-

Source: Companies, Bloomberg, Daiwa compiled

Daiwa Capital Markets

Baidu (BIDU US)

Target price: **USD215.00** (from USD195.00)

Share price (13 Sep): **USD236.41** | Up/downside: **-9.1%**

5 4 **3** 2 1

Hold (unchanged)

Going all in on AI

- Leads the pack in AI research in China
- Search and in-feed ads are major use cases for AI in the short run
- Maintaining Hold (3) rating; raising TP slightly to USD215

John Choi
(852) 2773 8730
john.choi@hk.daiwacm.com

Alex Liu
(852) 2848 4976
alex.liu@hk.daiwacm.com

What's new: We view Baidu as a pioneer in AI research in China, with management remaining determined to invest in this field. Indeed, it has been increasingly vocal about its DuerOS (intelligence assistant) and Apollo (autonomous driving) initiatives, which are in the early stages of commercialisation. Hence, we acknowledge Baidu's AI research leadership and also expect a short-term revenue recovery in 2H17 from a low base. But we are cautious on: 1) the likelihood of decelerating search revenue growth from 2018, 2) competition and execution risks on feed-ad monetisation, and 3) low visibility on its content and R&D spending beyond 2017.

What's the impact: **Leader of AI research in China**. Baidu has a tradition of investing in new technologies going back to the early 2000s. More recently, it set up a Deep Learning lab in China in 2013 and has recruited numerous scientists from academia in recent years. The company now has a team of over 2,000 scientists dedicated to machine learning research. Also, it has kept to its generous R&D budget over the past few years; in 1H17, its R&D spend was higher than Alibaba's but lower than Tencent's.

Where could Baidu deploy AI? We believe news feed and search are the major use cases for Baidu to deploy its machine-learning insights in the near term. We would look for the introduction of machine learning to the company's existing ad system to boost ad monetisation efficiency through better predicted click through rates (CTR) and more accurate targeting. Also, Baidu intends to make Apollo the "Android" for cars, ie, effectively an open-source universal autonomous driving platform. Although the monetisation model for DuerOS and Apollo is as yet unclear, management has been vocal about generating advertising and licensing revenue from both initiatives at some point in the future. While we see big potential from Baidu's autonomous driving business, we believe DuerOS and Apollo are unlikely to bring in significant revenue streams in the next 2-3 years. In all, we lift our 2017-19 earnings forecasts by 2.3-3.9% to reflect the disposal of its food delivery services business and our forecast for slightly higher core search profitability.

What we recommend: We maintain our Hold (3) rating and lift our SOTP-based TP to USD215 from USD195 (see following pages); we update our SOTP framework to reflect Baidu's disposal of the food delivery business to Elema (Nuomi remains in the core-search segment under our valuation framework). Key upside risk: lower marketing expenditure; key downside risk: slower search growth.

How we differ: We are more cautious than the market on Baidu's core search advertising revenue outlook for the medium term.

Forecast revisions (%)

Year to 31 Dec	17E	18E	19E
Revenue change	(0.5)	(1.8)	(2.9)
Net profit change	3.9	3.6	2.3
Core EPS (FD) change	3.9	3.6	2.3

Source: Daiwa forecasts

Share price performance

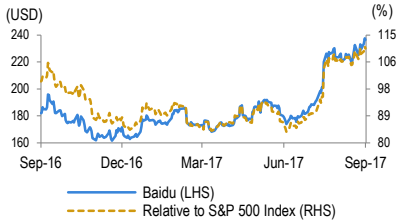

12-month range	161.67-237.60
Market cap (USDbn)	82.82
3m avg daily turnover (USDm)	617.51
Shares outstanding (m)	350
Major shareholder	Robin Yanhong Li (16.1%)

Financial summary (CNY)

Year to 31 Dec	17E	18E	19E
Revenue (m)	85,792	102,906	122,029
Operating profit (m)	12,959	14,613	21,746
Net profit (m)	15,836	17,868	24,674
Core EPS (fully-diluted)	45.569	51.415	71.000
EPS change (%)	18.3	12.8	38.1
Daiwa vs Cons. EPS (%)	0.2	(11.3)	(5.8)
PER (x)	33.9	30.0	21.7
Dividend yield (%)	0.0	0.0	0.0
DPS	0.000	0.000	0.000
PBR (x)	4.8	4.2	3.5
EV/EBITDA (x)	26.8	21.6	14.3
ROE (%)	15.3	14.8	17.4

Source: FactSet, Daiwa forecasts

Financial summary

Key assumptions

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Online marketing services revenue (CNYm)	22,246	31,802	48,487	64,037	64,525	73,468	84,421	96,148
Other services revenue (CNYm)	60	142	565	2,344	6,024	12,324	18,486	25,880
Selling & general expenses as % of revenue (%)	11.2	16.2	21.2	25.7	21.4	15.6	13.1	12.0
R&D expenses as % of revenue (%)	10.3	12.9	14.2	15.3	13.3	15.4	16.0	16.0

Profit and loss (CNYm)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Online marketing services	22,246	31,802	48,487	64,037	64,525	73,468	84,421	96,148
Other services	60	142	565	2,344	6,024	12,324	18,486	25,880
Other Revenue	0	0	0	0	0	0	0	0
Total Revenue	**22,306**	**31,944**	**49,052**	**66,381**	**70,549**	**85,792**	**102,906**	**122,029**
Other income	1,513	2,636	3,948	5,837	3,725	5,125	7,077	9,800
COGS	(6,449)	(11,472)	(18,885)	(27,458)	(35,279)	(46,216)	(58,348)	(66,115)
SG&A	(2,501)	(5,174)	(10,382)	(17,076)	(15,071)	(13,426)	(13,481)	(14,643)
Other op.expenses	(3,817)	(6,743)	(10,929)	(16,012)	(13,875)	(18,316)	(23,542)	(29,325)
Operating profit	**11,051**	**11,192**	**12,804**	**11,671**	**10,048**	**12,959**	**14,613**	**21,746**
Net-interest inc./(exp.)	759	813	974	1,321	1,184	1,353	1,500	1,500
Assoc/forex/extraord./others	156	180	307	24,914	3,276	991	1,844	2,430
Pre-tax profit	**11,965**	**12,185**	**14,085**	**37,906**	**14,509**	**15,303**	**17,956**	**25,676**
Tax	(1,574)	(1,829)	(2,231)	(5,474)	(2,914)	(2,537)	(3,639)	(5,029)
Min. int./pref. div./others	65	163	943	1,232	37	35	103	122
Net profit (reported)	**10,456**	**10,519**	**12,797**	**33,664**	**11,632**	**12,801**	**14,421**	**20,769**
Net profit (adjusted)	**10,668**	**11,034**	**13,662**	**35,051**	**13,392**	**15,836**	**17,868**	**24,674**
EPS (reported)(CNY)	**29.926**	**30.066**	**36.499**	**96.394**	**33.559**	**36.932**	**41.605**	**59.922**
EPS (adjusted)(CNY)	**30.533**	**31.537**	**38.966**	**100.366**	**38.636**	**45.689**	**51.551**	**71.188**
EPS (adjusted fully-diluted)(CNY)	**30.499**	**31.492**	**38.816**	**100.043**	**38.534**	**45.569**	**51.415**	**71.000**
DPS (CNY)	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
EBIT	**11,051**	**11,192**	**12,804**	**11,671**	**10,048**	**12,959**	**14,613**	**21,746**
EBITDA	**12,564**	**13,828**	**16,752**	**17,508**	**13,773**	**18,083**	**21,690**	**31,546**

Cash flow (CNYm)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Profit before tax	11,965	12,185	14,085	37,906	14,509	15,303	17,956	25,676
Depreciation and amortisation	1,513	2,636	3,948	5,837	3,725	5,125	7,077	9,800
Tax paid	(1,574)	(1,829)	(2,231)	(5,474)	(2,914)	(2,537)	(3,639)	(5,029)
Change in working capital	781	967	3,249	4,237	510	4,374	4,597	3,025
Other operational CF items	(689)	(166)	(1,503)	(18,846)	(750)	2,067	2,286	2,528
Cash flow from operations	**11,996**	**13,793**	**17,548**	**23,660**	**15,080**	**24,332**	**28,278**	**36,001**
Capex	(2,311)	(2,757)	(4,827)	(5,230)	(9,129)	(10,594)	(14,828)	(20,755)
Net (acquisitions)/disposals	(1,031)	10,240	39,498	(2,805)	1,621	1,818	2,048	2,319
Other investing CF items	(10,408)	(30,806)	(57,138)	(19,697)	(81)	(295)	(332)	(370)
Cash flow from investing	**(13,750)**	**(23,323)**	**(22,468)**	**(27,731)**	**(7,589)**	**(9,072)**	**(13,112)**	**(18,807)**
Change in debt	9,389	6,025	7,741	10,410	1,959	1,972	1,976	1,979
Net share issues/(repurchases)	57	1,397	193	(6,152)	0	0	0	0
Dividends paid	0	0	(338)	0	0	0	0	0
Other financing CF items	100	1,397	1,847	3,528	0	0	0	0
Cash flow from financing	**9,546**	**8,819**	**9,442**	**7,786**	**1,959**	**1,972**	**1,976**	**1,979**
Forex effect/others	(12)	(201)	80	179	0	0	0	0
Change in cash	**7,780**	**(911)**	**4,602**	**3,894**	**9,450**	**17,232**	**17,142**	**19,172**
Free cash flow	9,685	11,036	12,720	18,430	5,950	13,737	13,450	15,245

Source: FactSet, Daiwa forecasts

Financial summary continued …

Balance sheet (CNYm)

As at 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Cash & short-term investment	32,880	38,686	58,084	68,025	77,475	94,707	111,849	131,021
Inventory	0	0	0	0	0	0	0	0
Accounts receivable	1,253	2,221	3,664	3,927	4,174	5,076	6,088	7,219
Other current assets	541	2,122	4,093	6,281	5,104	6,508	8,331	9,963
Total current assets	**34,674**	**43,029**	**65,841**	**78,234**	**86,753**	**106,291**	**126,268**	**148,204**
Fixed assets	3,888	5,370	8,705	10,627	16,354	22,186	30,346	41,766
Goodwill & intangibles	5,465	20,495	20,993	18,730	16,786	14,607	12,150	9,366
Other non-current assets	1,642	2,092	4,122	40,262	41,140	42,403	43,896	45,644
Total assets	**45,669**	**70,986**	**99,662**	**147,853**	**161,034**	**185,487**	**212,660**	**244,980**
Short-term debt	0	0	93	100	93	93	93	93
Accounts payable	3,807	7,362	12,965	17,840	16,431	21,525	27,176	30,974
Other current liabilities	4,430	3,671	7,213	8,163	9,150	10,737	12,519	14,509
Total current liabilities	**8,237**	**11,033**	**20,271**	**26,103**	**25,675**	**32,355**	**39,787**	**45,577**
Long-term debt	9,693	17,229	23,507	33,942	35,942	37,942	39,942	41,942
Other non-current liabilities	524	2,058	1,378	3,593	3,559	3,531	3,507	3,485
Total liabilities	**18,454**	**30,321**	**45,156**	**63,638**	**65,176**	**73,828**	**83,236**	**91,004**
Share capital	0	0	0	0	0	0	0	0
Reserves/R.E./others	27,088	38,425	53,421	84,204	95,845	111,646	129,411	153,963
Shareholders' equity	**27,089**	**38,425**	**53,421**	**84,204**	**95,845**	**111,646**	**129,411**	**153,963**
Minority interests	127	2,240	1,085	12	12	12	12	12
Total equity & liabilities	**45,669**	**70,986**	**99,662**	**147,853**	**161,033**	**185,486**	**212,659**	**244,979**
EV	517,774	521,617	507,435	506,862	499,405	484,173	469,031	451,859
Net debt/(cash)	**(23,187)**	**(21,457)**	**(34,484)**	**(33,983)**	**(41,441)**	**(56,672)**	**(71,814)**	**(88,987)**
BVPS (CNY)	**77.472**	**109.690**	**152.498**	**240.373**	**273.606**	**318.712**	**369.426**	**439.514**

Key ratios (%)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Sales (YoY)	53.8	43.2	53.6	35.3	6.3	21.6	19.9	18.6
EBITDA (YoY)	48.5	10.1	21.1	4.5	(21.3)	31.3	19.9	45.4
Operating profit (YoY)	45.9	1.3	14.4	(8.8)	(13.9)	29.0	12.8	48.8
Net profit (YoY)	57.1	3.4	23.8	156.6	(61.8)	18.3	12.8	38.1
Core EPS (fully-diluted) (YoY)	57.0	3.3	23.3	157.7	(61.5)	18.3	12.8	38.1
Gross-profit margin	71.1	64.1	61.5	58.6	50.0	46.1	43.3	45.8
EBITDA margin	56.3	43.3	34.2	26.4	19.5	21.1	21.1	25.9
Operating-profit margin	49.5	35.0	26.1	17.6	14.2	15.1	14.2	17.8
Net profit margin	47.8	34.5	27.9	52.8	19.0	18.5	17.4	20.2
ROAE	49.3	33.7	29.8	50.9	14.9	15.3	14.8	17.4
ROAA	30.9	18.9	16.0	28.3	8.7	9.1	9.0	10.8
ROCE	39.7	23.6	18.8	11.9	8.0	9.2	9.2	11.9
ROIC	237.1	81.9	54.9	28.4	15.3	19.8	20.7	28.5
Net debt to equity	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Effective tax rate	13.2	15.0	15.8	14.4	20.1	16.6	20.3	19.6
Accounts receivable (days)	15.2	19.8	21.9	20.9	21.0	19.7	19.8	19.9
Current ratio (x)	4.2	3.9	3.2	3.0	3.4	3.3	3.2	3.3
Net interest cover (x)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Net dividend payout	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Free cash flow yield	1.8	2.0	2.4	3.4	1.1	2.5	2.5	2.8

Source: FactSet, Daiwa forecasts

Company profile

Founded in 2000, Baidu is the world's largest Chinese-language search engine. Its Internet search services cover web pages, news, images, documents, and multimedia files. Baidu had about 470,000 active online marketing customers in 2Q17.

Going all in on AI

At the Baidu Create conference, held on 5 July, Baidu announced plans to open source its DuerOS (intelligent assistants) and Apollo (autonomous driving platform).

DuerOS

Baidu intends to promote DuerOS as the next-generation user interface

DuerOS is Baidu's voice-based intelligent assistance platform (ie, Baidu's answer to Apple's Siri), which is backed by a significant number of user voice queries, as well as Baidu's machine-learning capabilities (leveraging on Baidu Brain).

In order to drive adoption of DuerOS, Baidu plans to open source the system and extend a welcome stance to hardware makers (smartphone makers, intelligent speaker makers, etc). As it stands, Baidu has collaborated with more than 30 hardware makers, including home appliance companies such as Midea, Haier, Lenovo, Sony, and smartphone makers including VIVO and HTC.

Baidu: AI framework "ABC" – AI, Big data, Cloud

Baidu Brain	A	Platform layer	Open AI platform			
		Cognition layer	NLP	Knowledge graph	User profile	
		Perception layer	Voice	Image	Video	AR/VR
		Algorithm layer	Machine learning platform	Deep learning		
Intelligent Cloud	B	Big data	Big data analysis	Data labelling	Data collection	
	C	Cloud	Computing services CPU/GPU/FPGA	Storage services	Network services	

Source: Company

Baidu: DuerOS architecture

DuerOS intelligent tools platform	Reference design						DuerOS App
	Developer tools			Industrial design	Structure design	Sound design	
	Core implementation	Chipset	Microphone placement				
DuerOS speech core system	Voice recognition		Voice broadcast		Monitor display		
	NLP	Chat status control	Chat management	Natural language generator			Search
	Knowledge graph	Web graph	Needs Graph	Geographical graph			User label
DuerOS skills platform	Native skills		Third party skill				
	Skill development tools						

Source: Company

Apollo

Baidu envisions Apollo being the Android of autonomous driving cars

Apollo is Baidu's autonomous driving platform, which was officially launched in April 2017. Through its open-sourcing capability, which includes route planning, simulation, and voice-based interaction system (DuerOS), Baidu intends to build a large ecosystem around Apollo. Its goal is for Apollo to become the default operating system for the autonomous driving industry in the future.

We believe autonomous driving is an industry with a large total addressable market (TAM), which could enable Baidu's machine-learning technology to take off in earnest. But, in our view, commercialisation of autonomous driving could take another 3-5 years at least. While management believes that the commercialisation opportunities, including app distribution and search, will follow after adoption, we believe visibility on commercialisation remains low and competition in terms of development remains intense globally.

Baidu stated that it had more than 50 partners for Apollo as of July 2017, including auto OEMs (Chery, Chang'an, Daimler, Ford), auto parts manufacturers (Bosch, Continental Automotive) and software companies (Velodyne Lidar, ThunderSoft, TomTom). The company also introduced its roadmap for Apollo through to the year 2020 (detailed below).

The road map indicates that Baidu wants to realise "highway and urban road autonomous driving" by the end of 2020.

Baidu: Apollo architecture

Cloud services platform	HD map	Simulation	Data platform		Security	OTA	DuerOS
Open Software Platform	Map engine	Localization	Perception	Planning	Control	End-to-end	HMI
	Runtime framework						
	RTOS: real-time operating system						
Reference hardware platform	Computing unit	GPS/IMU	Camera	LiDAR	Radar	HWI Device	Black Box
Reference vehicle platform	Drive-by-wire vehicle						

Source: Company

Baidu: Apollo adoption road-map

	Closed-cycle autonomous driving	Fixed lane autonomous driving	Simple urban road autonomous driving	Specific highway and urban road autonomous driving	Highway and urban road autonomous driving (Alpha version)	Highway and urban road autonomous driving
Apollo ability release time table	Control Localisation Runtime framework Reference hardware Reference vehicle	Obstacle perception Planning HD map Simulation service	Advanced perception End-to-end Security service			
	Jul-17	**Sep-19**	**Dec-17**	**Dec-18**	**Dec-19**	**Dec-20**
Apollo resource sharing time table	Data platform 1.0 3D roadblock labeling data Road hackers data High-definition mapping technology testing	Artificial edit simulation scene data	Data platform 2.0 Roadblock behaviour labelling data 2D roadblocking labelling data Log extract simulation scene data			

Source: Company

R&D spending likely to continue

Baidu has been one of the biggest spenders on research among Internet companies in China in recent years. As the chart below shows, Baidu significantly stepped up its R&D spending in 2015. However, it cut such spending in 2016, likely due to operating-expense discipline in the face of regulatory-scrutiny headwinds regarding its ad revenues. The company resumed its R&D spending push in 1H17 (up 30% YoY).

China: major listed internet companies' research expense comparison (2014-1H17)

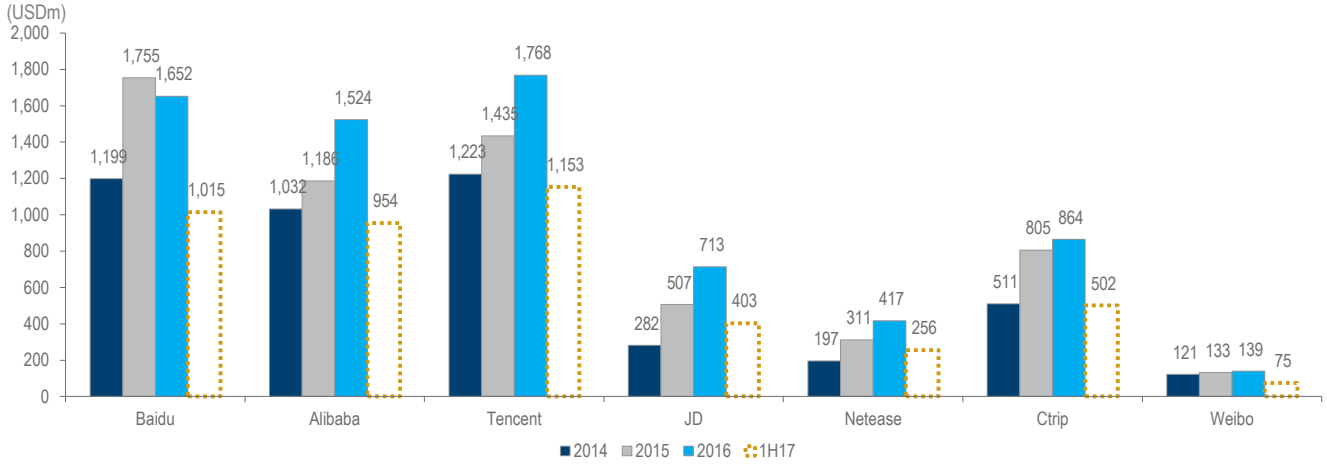

Source: Companies, Daiwa compiled
Note: all expense in non-GAAP (including share-based compensation and intangible amortisation) except for Tencent (Tencent does not disclose non-GAAP expense reconciliation)

In absolute terms, Baidu's R&D spending is lower than that of Tencent but slightly higher than Alibaba's.

Operating margin excluding iQiyi (including O2O) was largely flat in the 2 years prior to the jump recorded in 2Q17

Baidu: historical operating profit trend (2Q15-2Q17)

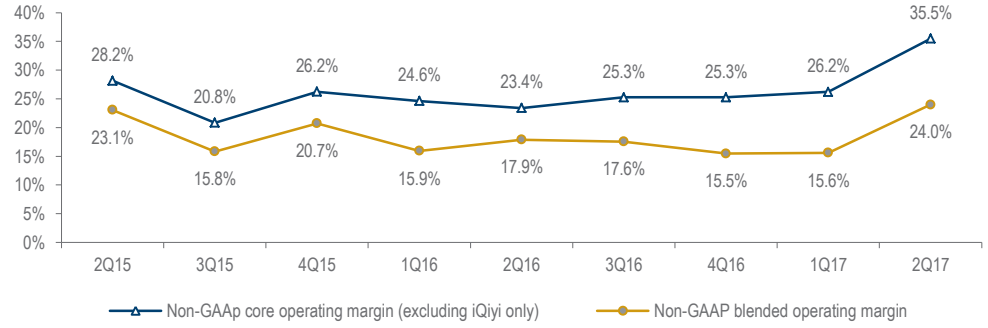

Source: Company, Daiwa estimates;

Valuation and earnings revisions

We maintain our Hold (3) rating and lift our SOTP-based TP to USD215 from USD195; we update our SOTP framework to reflect Baidu's disposal of the food delivery business to Elema (Nuomi remains in the core-search segment under our valuation framework).

In all, we lift our 2017-19 earnings forecasts by 2.3-3.9% to reflect the disposal of its food delivery services business and our forecast for slightly higher core search profitability.

New SOTP-based target price of USD215

We increase our core-search EBIT margin forecast from 32% to 36% to reflect the company's better operating expense control and the deconsolidation of Baidu's food delivery business

Baidu: SOTP valuation

Baidu SOTP valuation	
Core business (including feed ads & O2O business)	
Baidu core-search revenue (average 2017-18E)	74.6
Multiply: Core-search EBIT Margin %	*36%*
Core-search EBIT (average 2017-18E)	26.8
Multiply: Tax rate %	*15%*
Core-search EBIAT (average 2017-18E)	22.8
Valuation multiple: Price-to-earnings Ratio	***16.0x***
Core-business valuation (CNYbn)	365.1
Core-business valuation (USDbn)	**53.7**
iQiyi (online video)	
iQiyi revenue *(average 2017-18E)*	19.8
Valuation multiple: Price-to-sales Ratio	***4.0x***
iQiyi valuation (CNYbn)	79.1
iQiyi valuation (USDbn)	**11.6**
Baidu's ownership in iQiyi	60%
iQiyi valuation (USDbn)	**7.0**
Total SOTP valuation for Baidu:	
Core-business (including O2O) valuation (USDbn)	**53.7**
iQiyi valuation (USDbn)	**7.0**
Ctrip's current market cap – fully diluted (USDbn)	32.9
Multiply: Baidu's ownership in Ctrip % - fully diluted	15%
Plus: Valuation of Baidu's stake in Ctrip (USDbn)	**4.9**
Elema valuation– last round of financing (USDbn)	6.0
Multiply: Baidu's ownership in Elema %	8%
Plus: Valuation of Baidu's stake in Elema (USDbn)	**0.5**
Plus: average 2017-18E net cash (USDbn)	9.4
Multiply: holding discount	**0%**
SOTP valuation for Baidu (USDbn)	**75.6**
Implied SOTP-based target price for Baidu (USD)	**215.0**

Source: Daiwa

Baidu: earnings revisions

	New			Old			Difference %		
(in CNYm except for per ADS figure)	2017E	2018E	2019E	2017E	2018E	2019E	2017E	2018E	2019E
Total revenues - reported	**85,792**	**102,906**	**122,029**	**86,201**	**104,794**	**125,689**	*-0.5%*	*-1.8%*	*-2.9%*
% change YoY	*22%*	*20%*	*19%*	*22%*	*22%*	*20%*			
Gross profit	**39,576**	**44,558**	**55,914**	**39,265**	**45,061**	**56,083**	*0.8%*	*-1.1%*	*-0.3%*
% change YoY	*12%*	*13%*	*25%*	*11%*	*15%*	*24%*			
Gross margin	*46.1%*	*43.3%*	*45.8%*	*45.6%*	*43.0%*	*44.6%*	*0.58ppt*	*0.30ppt*	*1.20ppt*
Operating income – GAAP	**12,959**	**14,613**	**21,746**	**12,274**	**13,623**	**20,890**	*5.6%*	*7.3%*	*4.1%*
% change YoY	*29%*	*13%*	*49%*	*22%*	*11%*	*53%*			
Operating margin	*15.1%*	*14.2%*	*17.8%*	*14.2%*	*13.0%*	*16.6%*	*0.87ppt*	*1.20ppt*	*1.20ppt*
Net income – non-GAAP	**15,836**	**17,868**	**24,674**	**15,240**	**17,255**	**24,110**	*3.9%*	*3.6%*	*2.3%*
% change YoY	*18%*	*13%*	*38%*	*14%*	*13%*	*40%*			
Net profits margin – non-GAAP	*18.5%*	*17.4%*	*20.2%*	*17.7%*	*16.5%*	*19.2%*	*0.78ppt*	*0.90ppt*	*1.04ppt*
Net income per ADS – non-GAAP	**45.57**	**51.41**	**71.00**	**43.85**	**49.65**	**69.38**	*3.9%*	*3.6%*	*2.3%*
% change YoY	*18%*	*13%*	*38%*	*14%*	*13%*	*40%*			

Source: Daiwa forecasts

Key upside risk to our Hold (3) rating: lower marketing expenditure; key downside risk: slower search growth.


Daiwa
Capital Markets

Daiwa Capital Markets

Alibaba Group (BABA US)

Target price: **USD200.00** (from USD205.00)

Share price (13 Sep): **USD178.97** | Up/downside: **+11.8%**

5 4 3 2 **1**

Buy (unchanged)

Artificial intelligence status not yet fully recognised

- Alibaba's business model is built on the vast user data on hand
- Alibaba has multiple use cases to deploy AI and machine learning
- Reiterating Buy (1); trimming TP to USD200; still our sector top pick

John Choi
(852) 2773 8730
john.choi@hk.daiwacm.com

Alex Liu
(852) 2848 4976
alex.liu@hk.daiwacm.com

What's new: We believe Alibaba has a natural edge in AI research that has yet to be fully appreciated. In our view, it has a natural advantage in deploying AI technology to its existing businesses given its: 1) large-scale data regarding consumers' online shopping behaviour, 2) the ability to track users across multiple digital media properties, 3) ample use cases to adopt machine learning in existing Internet services (such as search, recommendation engine, display ads), and 4) cloud leadership in China, which should result in cost advantages on machine learning algorithm training. Alibaba remains our top pick within the China e-commerce sector.

What's the impact: E-commerce data powerhouse with determined stance on AI. Despite being less vocal on its AI investment vs. its competitors, we believe Alibaba has been quietly building up its machine learning research capability since the late-2000s. The company set up an in-house AI Lab in 2016, which currently has more than 300 scientists and engineers dedicated to AI research. Historically, Alibaba has lagged Baidu and Tencent on R&D spending; however, the R&D spending gap has rapidly narrowed over the past 3 years with Alibaba's increased investment in machine learning research.

Where could Alibaba deploy AI? We see multiple use cases for Alibaba to deploy machine leaning technologies to existing services including: 1) organic search (the company rolled out personalised search for the Taobao platform in 2013), 2) product and content recommendation (mobile Taobao now presents customised interface and content to users), 3) image recognition and search (deep learning helps Alibaba provide look-alike products to users); 4) financial services (Ant Financial Group, Alibaba's financial services arm, now integrates machine learning technology to provide customised personal finance recommendations, calculate credit scores and decide on interest rates). Alibaba's Union-ID system (same log-in across all Alibaba services from online shopping to online video/music) should help it understand users' online behaviour through multiple dimensions, leading to a smoother, more personalised experience in the long term. We cut our FY18-20E earnings by 1.2-3.3% to reflect its slightly higher content spending and investment in "new retail" initiatives.

What we recommend: We reiterate our Buy (1) call and lower slightly our 12-month TP to USD200 (from USD205), based on an unchanged 37x PER applied to the average of our FY18-19E non-GAAP EPS. Key risk: further slowdown in the Taobao marketplace.

How we differ: We are more positive than the market on Alibaba's mobile monetisation outlook despite our lower margin forecasts versus the street.

Forecast revisions (%)

Year to 31 Mar	18E	19E	20E
Revenue change	-	-	-
Net profit change	(3.3)	(2.0)	(1.2)
Core EPS (FD) change	(3.3)	(2.0)	(1.2)

Source: Daiwa forecasts

Share price performance

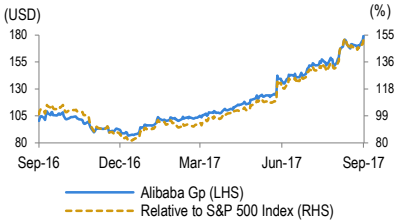

12-month range	86.79-178.97
Market cap (USDbn)	455.12
3m avg daily turnover (USDm)	254.83
Shares outstanding (m)	2,543
Major shareholder	SoftBank (25.0%)

Financial summary (CNY)

Year to 31 Mar	18E	19E	20E
Revenue (m)	234,242	321,848	422,047
Operating profit (m)	73,848	106,529	144,139
Net profit (m)	82,705	109,174	140,923
Core EPS (fully-diluted)	31.822	41.797	53.684
EPS change (%)	35.2	31.3	28.4
Daiwa vs Cons. EPS (%)	(2.4)	(3.8)	(0.2)
PER (x)	36.7	28.0	21.8
Dividend yield (%)	0.0	0.0	0.0
DPS	0.000	0.000	0.000
PBR (x)	7.5	5.4	3.9
EV/EBITDA (x)	24.9	17.5	12.4
ROE (%)	24.4	22.9	21.5

Source: FactSet, Daiwa forecasts

Financial summary

Key assumptions

Year to 31 Mar	2013	2014	2015	2016	2017	2018E	2019E	2020E
Annual Active buyers (m)	172	255	350	423	454	481	505	526
Average spending per active buyers (CNY)	6,262	6,580	6,982	7,310	8,297	9,393	10,467	11,473
T-mall GMV (CNYbn)	253	505	847	1,215	1,565	2,010	2,502	2,992
Taobao GMV (CNYbn)	824	1,173	1,597	1,877	2,202	2,510	2,786	3,037
China Retail Total GMV (CNYbn)	1,077	1,678	2,444	3,092	3,767	4,520	5,289	6,029
Mobile GMV (CNYbn)	79	318	994	2,003	2,981	3,883	4,548	5,306

Profit and loss (CNYm)

Year to 31 Mar	2013	2014	2015	2016	2017	2018E	2019E	2020E
China retail business	26,970	42,832	59,732	80,033	114,109	169,338	231,146	295,682
China wholesale business	2,197	2,300	3,205	4,288	5,679	7,164	8,238	9,227
Other Revenue	5,350	7,372	13,267	16,822	38,485	57,740	82,464	117,139
Total Revenue	**34,517**	**52,504**	**76,204**	**101,143**	**158,273**	**234,242**	**321,848**	**422,047**
Other income	0	0	0	0	0	0	0	0
COGS	(9,719)	(13,369)	(23,834)	(34,272)	(59,483)	(90,805)	(123,911)	(162,488)
SG&A	(6,502)	(8,763)	(16,313)	(20,512)	(28,553)	(39,636)	(53,430)	(67,729)
Other op.expenses	(7,545)	(5,452)	(12,922)	(17,174)	(22,182)	(29,953)	(37,978)	(47,691)
Operating profit	**10,751**	**24,920**	**23,135**	**29,185**	**48,055**	**73,848**	**106,529**	**144,139**
Net-interest inc./(exp.)	(1,533)	(547)	6,705	50,308	5,888	972	950	1,000
Assoc/forex/extraord./others	894	2,429	2,486	2,058	6,086	4,737	5,000	6,200
Pre-tax profit	**10,112**	**26,802**	**32,326**	**81,551**	**60,029**	**79,557**	**112,479**	**151,339**
Tax	(1,457)	(3,196)	(6,416)	(8,449)	(13,776)	(19,013)	(25,870)	(34,808)
Min. int./pref. div./others	(251)	(530)	(1,761)	(1,559)	(2,578)	(2,686)	(3,500)	(3,000)
Net profit (reported)	**8,404**	**23,076**	**24,149**	**71,543**	**43,675**	**57,858**	**83,108**	**113,531**
Net profit (adjusted)	**13,869**	**27,610**	**34,922**	**42,995**	**60,516**	**82,705**	**109,174**	**140,923**
EPS (reported)(CNY)	**3.660**	**10.610**	**10.333**	**29.160**	**17.618**	**22.752**	**32.681**	**44.645**
EPS (adjusted)(CNY)	**6.040**	**12.695**	**14.943**	**17.524**	**24.411**	**32.523**	**42.931**	**55.416**
EPS (adjusted fully-diluted)(CNY)	**5.892**	**11.965**	**13.969**	**16.782**	**23.533**	**31.822**	**41.797**	**53.684**
DPS (CNY)	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**	**0.000**
EBIT	**15,672**	**29,077**	**36,338**	**45,722**	**64,050**	**89,988**	**123,318**	**161,560**
EBITDA	**16,607**	**30,731**	**40,753**	**52,423**	**74,456**	**105,147**	**139,830**	**179,767**

Cash flow (CNYm)

Year to 31 Mar	2013	2014	2015	2016	2017	2018E	2019E	2020E
Profit before tax	10,112	26,802	32,326	81,551	60,029	79,557	112,479	151,339
Depreciation and amortisation	935	1,654	4,415	5,859	8,212	12,090	12,797	14,009
Tax paid	(1,457)	(3,196)	(6,416)	(8,449)	(13,776)	(19,013)	(25,870)	(34,808)
Change in working capital	3,136	(4,493)	12,706	10,797	24,372	36,518	37,105	40,521
Other operational CF items	1,750	5,612	(1,814)	(29,121)	40,367	52,658	53,894	57,941
Cash flow from operations	**14,476**	**26,379**	**41,217**	**60,637**	**119,205**	**161,810**	**190,404**	**229,002**
Capex	(2,503)	(4,776)	(7,705)	(21,500)	(17,000)	(13,500)	(12,000)	(12,000)
Net (acquisitions)/disposals	(36)	(17,111)	(33,586)	(20,000)	(6,000)	(6,000)	(6,000)	(6,000)
Other investing CF items	3,084	(11,110)	(12,163)	0	0	0	0	0
Cash flow from investing	**545**	**(32,997)**	**(53,454)**	**(41,500)**	**(23,000)**	**(19,500)**	**(18,000)**	**(18,000)**
Change in debt	26,932	12,789	(22,713)	26,000	0	0	0	0
Net share issues/(repurchases)	(12,777)	1,923	61,438	(15,000)	(13,000)	0	0	0
Dividends paid	(103)	(208)	(104)	0	0	0	0	0
Other financing CF items	(15,458)	(5,140)	48,876	0	0	0	0	0
Cash flow from financing	**(1,406)**	**9,364**	**87,497**	**11,000**	**(13,000)**	**0**	**0**	**0**
Forex effect/others	(76)	(97)	(112)	0	0	0	0	0
Change in cash	**13,539**	**2,649**	**75,148**	**30,137**	**83,205**	**142,310**	**172,404**	**211,002**
Free cash flow	19,745	32,269	48,121	68,660	129,705	175,310	202,404	241,002

Source: FactSet, Daiwa forecasts

Financial summary continued …

Balance sheet (CNYm)

As at 31 Mar	2013	2014	2015	2016	2017	2018E	2019E	2020E
Cash & short-term investment	36,373	48,553	124,638	154,775	237,980	380,290	552,694	763,696
Inventory	0	0	0	0	0	0	0	0
Accounts receivable	1,734	4,679	12,978	16,183	23,741	32,794	41,840	54,866
Other current assets	5,055	14,601	4,493	4,164	4,291	4,361	4,302	4,502
Total current assets	**43,162**	**67,833**	**142,109**	**175,122**	**266,012**	**417,444**	**598,836**	**823,064**
Fixed assets	5,703	7,241	12,244	19,272	23,182	25,115	27,329	29,028
Goodwill & intangibles	11,628	13,699	48,508	84,342	92,457	101,294	116,163	141,349
Other non-current assets	3,293	22,776	52,573	73,545	82,402	89,858	99,362	109,370
Total assets	**63,786**	**111,549**	**255,434**	**352,281**	**464,053**	**633,710**	**841,690**	**1,102,811**
Short-term debt	5,448	10,364	1,990	1,990	1,990	1,990	1,990	1,990
Accounts payable	8,961	11,887	19,834	28,520	49,500	75,565	103,116	135,218
Other current liabilities	9,586	15,133	17,848	23,661	37,261	57,849	79,383	104,451
Total current liabilities	**23,995**	**37,384**	**39,672**	**54,171**	**88,752**	**135,404**	**184,489**	**241,659**
Long-term debt	22,462	30,711	50,603	76,603	76,603	76,603	76,603	76,603
Other non-current liabilities	6,283	2,636	7,088	7,234	7,567	8,011	8,522	9,108
Total liabilities	**52,740**	**70,731**	**97,363**	**138,007**	**172,922**	**220,018**	**269,614**	**327,370**
Share capital	32,189	37,445	117,801	102,801	89,801	89,801	89,801	89,801
Reserves/R.E./others	(21,680)	2,294	28,296	99,558	189,244	309,356	464,538	666,403
Shareholders' equity	**10,509**	**39,739**	**146,097**	**202,359**	**279,045**	**399,157**	**554,339**	**756,204**
Minority interests	537	1,079	11,974	11,915	12,086	14,535	17,737	19,237
Total equity & liabilities	**63,786**	**111,549**	**255,434**	**352,281**	**464,053**	**633,710**	**841,690**	**1,102,811**
EV	2,962,730	2,948,146	2,878,263	2,854,067	2,765,033	2,619,173	2,443,970	2,228,468
Net debt/(cash)	**(8,463)**	**(7,478)**	**(72,045)**	**(76,182)**	**(159,387)**	**(301,697)**	**(474,101)**	**(685,103)**
BVPS (CNY)	**n.a.**	**13.489**	**62.233**	**82.209**	**112.298**	**156.704**	**217.727**	**297.108**

Key ratios (%)

Year to 31 Mar	2013	2014	2015	2016	2017	2018E	2019E	2020E
Sales (YoY)	72.4	52.1	45.1	32.7	56.5	48.0	37.4	31.1
EBITDA (YoY)	128.3	85.0	32.6	28.6	42.0	41.2	33.0	28.6
Operating profit (YoY)	144.7	85.5	25.0	25.8	40.1	40.5	37.0	31.0
Net profit (YoY)	115.0	99.1	26.5	23.1	40.8	36.7	32.0	29.1
Core EPS (fully-diluted) (YoY)	131.2	103.1	16.7	20.1	40.2	35.2	31.3	28.4
Gross-profit margin	71.8	74.5	68.7	66.1	62.4	61.2	61.5	61.5
EBITDA margin	48.1	58.5	53.5	51.8	47.0	44.9	43.4	42.6
Operating-profit margin	45.4	55.4	47.7	45.2	40.5	38.4	38.3	38.3
Net profit margin	40.2	52.6	45.8	42.5	38.2	35.3	33.9	33.4
ROAE	88.2	188.4	40.0	24.8	25.2	24.4	22.9	21.5
ROAA	25.0	31.5	19.0	14.1	14.8	15.1	14.8	14.5
ROCE	42.0	48.1	24.8	18.2	19.3	20.9	21.6	21.5
ROIC	139.2	122.2	31.1	23.3	27.4	46.1	78.1	117.9
Net debt to equity	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Effective tax rate	14.4	11.9	19.8	10.4	22.9	23.9	23.0	23.0
Accounts receivable (days)	18.0	22.3	42.3	52.6	46.0	44.0	42.3	41.8
Current ratio (x)	1.8	1.8	3.6	3.2	3.0	3.1	3.2	3.4
Net interest cover (x)	10.2	53.2	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Net dividend payout	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Free cash flow yield	0.7	1.1	1.6	2.3	4.4	5.9	6.8	8.1

Source: FactSet, Daiwa forecasts

Company profile

Alibaba group is the world's largest e-commerce and m-commerce company in terms of GMV in 2014. Alibaba's ecosystem comprises B2C, C2C, B2B platforms, payment solutions, and cloud business for third-party service providers and other strategic partners.

AI status not yet fully recognised

Alibaba has many use cases to adopt AI and machine learning technology

During the investor day held by the company from 8-9 June 2017, Alibaba's management stated that the key for adoption would be to have suitable scenarios (use cases) combined with ample data accumulation. Both the quantity and quality of data on hand will ultimately determine the ceiling for AI in terms of the level of analysis. Alibaba's management believes that the company is uniquely positioned to explore the technology frontier of AI in the long term.

Alibaba: highlight on AI from 2017 investor day presentation (June 2017)

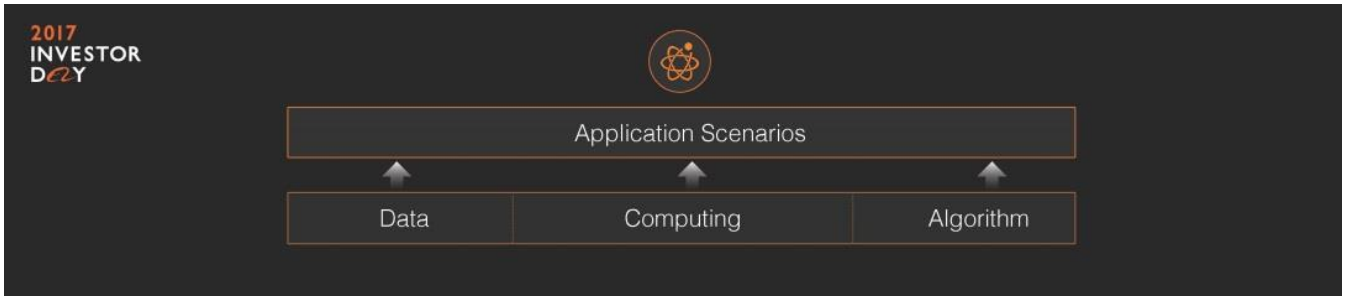

Source: Company

We note the following use cases to which Alibaba has been deploying machine learning technologies:

- Organic search: search results within mobile Taobao apps are already personalised
- Recommended products: the "guess what product you may like" function within the mobile Taobao app is supported by machine learning technology
- Customised interface: the layout and content on home page of mobile Taobao app has become more personalised
- Feed content: the content displayed under "Weitao" function within the mobile Taobao app has also been personalised
- Image search: the function of "search similar products" within the mobile Taobao app has adopted deep learning technology
- Performance-based ads: the ad platform has integrated deep learning techniques to improve overall monetisation efficiency

Alibaba: highlight on AI from 2017 investor day presentation (June 2017)

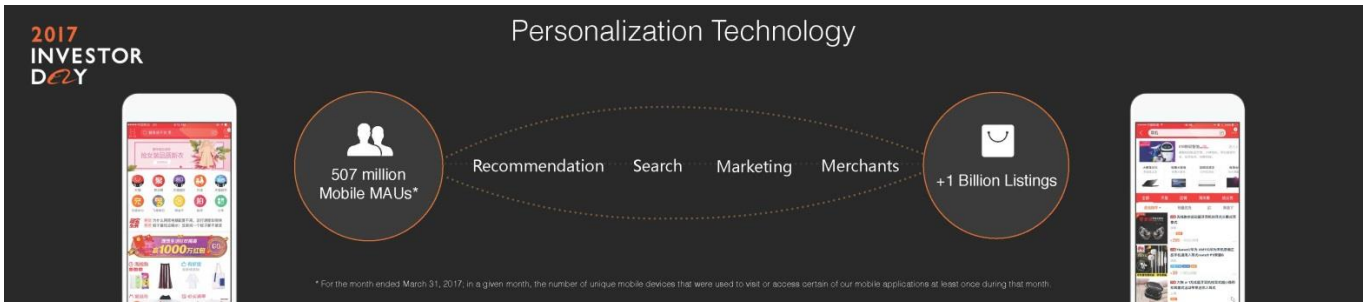

Source: Company

Alibaba: highlight on AI from 2017 investor day presentation (June 2017)

Source: Company

R&D spending likely to continue

Alibaba's R&D spending has been increasing rapidly

Alibaba has been one of the biggest spenders on research among Internet companies in China. Nevertheless, as the chart below shows, its R&D spending is currently lower than Tencent and Baidu in absolute terms.

Alibaba's non-GAAP R&D spending increased by over 50% from 2014 to 2016, with 1H17 non-GAAP R&D expense up by another 38% YoY. We believe the company will likely continue investing heavily in R&D to secure leadership in machine learning research.

China: major listed internet companies' research expense comparison (2014-1H17)

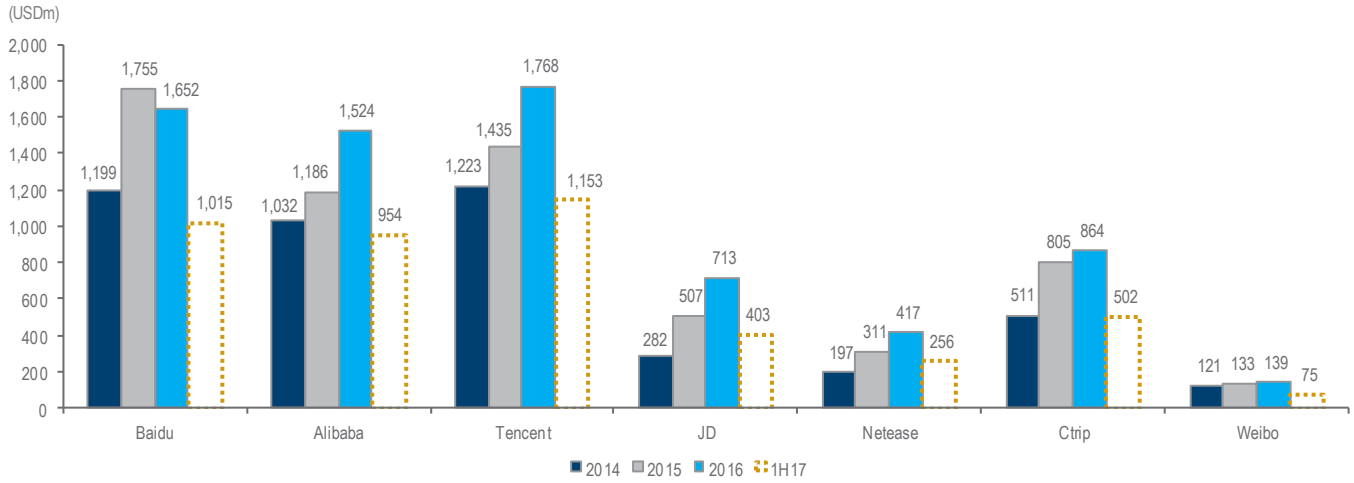

Source: Companies, Daiwa compiled
Note: all expense in non-GAAP (including share-based compensation and intangible amortisation) except for Tencent (Tencent does not disclose non-GAAP expense reconciliation).

Valuation and earnings revisions

We reiterate our Buy (1) call and lower slightly our 12-month TP to USD200 (from USD205), based on an unchanged 37x PER applied to the average of our FY18-19E non-GAAP EPS.

In all, we cut our FY18-20E earnings slightly by 1.2-3.3% to reflect its slightly higher content spending and investment in "new retail" initiatives.

Alibaba: earnings revisions

	New			Before			Difference		
	FY2018E	FY2019E	FY2020E	FY2018E	FY2019E	FY2020E	FY2018E	FY2019E	FY2020E
Total revenues (CNYm)	**234,242**	**321,848**	**422,047**	**234,242**	**321,848**	**422,047**	*0.0%*	*0.0%*	*0.0%*
% change YoY	*48%*	*37%*	*31%*	*48%*	*37%*	*31%*			
Gross profit (CNYm)	**143,437**	**197,936**	**259,559**	**146,132**	**199,546**	**259,559**	*-1.8%*	*-0.8%*	*0.0%*
% change YoY	*45%*	*38%*	*31%*	*48%*	*37%*	*30%*			
Operating income - GAAP (CNYm)	**73,848**	**106,529**	**144,139**	**77,390**	**109,076**	**144,983**	*-4.6%*	*-2.3%*	*-0.6%*
% change YoY	*54%*	*44%*	*35%*	*61%*	*41%*	*33%*			
EBITDA - Non-GAAP (CNYm)	**105,147**	**139,830**	**179,767**	**108,829**	**142,617**	**181,702**	*-3.4%*	*-2.0%*	*-1.1%*
% change YoY	*41%*	*33%*	*29%*	*46%*	*31%*	*27%*			
EBITDA Margin % - Non-GAAP	**44.9%**	**43.4%**	**42.6%**	**46.5%**	**44.3%**	**43.1%**	*-1.57 ppt*	*-0.87 ppt*	*-0.46 ppt*
Net income - Non-GAAP (CNYm)	**82,705**	**109,174**	**140,923**	**85,504**	**111,375**	**142,664**	*-3.3%*	*-2.0%*	*-1.2%*
% change YoY	*37%*	*32%*	*29%*	*41%*	*30%*	*28%*			
Net profits margin - Non-GAAP	*35.3%*	*33.9%*	*33.4%*	*36.5%*	*34.6%*	*33.8%*	*-1.20 ppt*	*-0.68 ppt*	*-0.41 ppt*
Diluted income per ADS (CNY) - Non-GAAP	**31.822**	**41.797**	**53.684**	**32.899**	**42.640**	**54.347**	*-3.3%*	*-2.0%*	*-1.2%*
% change YoY	*35%*	*31%*	*28%*	*40%*	*30%*	*27%*			

Source: Daiwa forecasts (year-end 31 March)

The key downside risk for Alibaba would be a further slowdown in the Taobao marketplace.

Daiwa Capital Markets

Tencent Holdings (700 HK)

Target price: **HKD370.00** (from HKD360.00)

Share price (13 Sep): **HKD334.00** | Up/downside: **+10.8%**

5 4 3 2 **1**

Buy (unchanged)

A pragmatist on artificial intelligence

- Multiple tailwinds supporting Tencent's artificial intelligence research
- Machine learning to lift user experience and monetisation efficiency
- Reiterating Buy (1) rating; raising TP slightly to USD370

John Choi
(852) 2773 8730
john.choi@hk.daiwacm.com

Alex Liu
(852) 2848 4976
alex.liu@hk.daiwacm.com

What's new: Over the past few months, Tencent has been stepping up its spending on artificial intelligence (AI) and machine learning research. In our view, Tencent has a natural advantage in AI research with its strong balance sheet, determined management and, most importantly, its wide coverage of Internet users in China across multiple services, which helps it gather unique user data and deploy AI technology. Despite Tencent not being the first-mover, we believe it will be a strong contender for the upcoming AI competition in China.

What's the impact: **Taking a practical approach to AI.** Tencent has been turning more vocal on AI and machine learning technology over the past few months. For instance, its management began the 2Q17 results call by emphasising the company's commitment to invest in AI and its distinctive advantage from its vast use case coverage on adopting AI. Tencent believes AI is an essential capability and is willing to invest in AI technology on an ongoing basis. Tencent built its first AI lab in 2016 and has hired over 200 scientists dedicated to AI research. We expect Tencent to take a practical approach to AI by focusing on how machine learning could be integrated into its existing products to improve the user experience and monetisation efficiency. If executed well, machine learning could lead to an immediate improvement in user experience as well as incremental revenue through existing business models (mostly ads in the near term).

Where could Tencent deploy AI? We see multiple use cases for Tencent to deploy machine leaning technologies to existing services including: 1) content recommendation (Tencent News and WeChat Kanyikan are transforming into algorithm-driven news recommendation platforms), 2) financial services (Tencent's WeBank is using machine learning to better predict the default rate of consumer credit), and 3) digital ads (machine learning will enhance the targeting capability of Tencent's GDT online ads platform). In addition, Tencent's CEO, Pony Ma, publicly stated in June 2017 that he views the cloud as a fundamental capability for AI and machine learning. As such, we expect the company to continue its heavy investment in cloud infrastructure to support its long-term vision on AI.

What we recommend: We are nudging up our 2017-19E earnings to factor in slightly higher PC game and other revenue, and reaffirm our Buy (1) call. We raise our 12-month TP to HKD370 (from HKD360), based on an unchanged 39x PER applied to our average 2017-18E non-GAAP EPS. Key risk: higher-than-expected marketing and content expenses.

How we differ: We are more positive than the market on Tencent's mobile game revenue in 2018-19E.

Forecast revisions (%)

Year to 31 Dec	17E	18E	19E
Revenue change	1.1	2.9	3.9
Net profit change	2.7	2.5	2.7
Core EPS (FD) change	2.7	2.5	2.7

Source: Daiwa forecasts

Share price performance

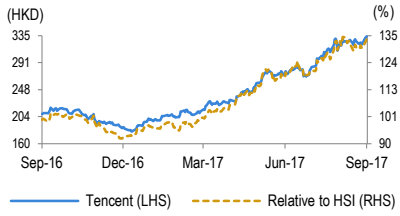

12-month range	179.70-334.00
Market cap (USDbn)	402.01
3m avg daily turnover (USDm)	795.27
Shares outstanding (m)	9,402
Major shareholder	Naspers Limited (33.6%)

Financial summary (CNY)

Year to 31 Dec	17E	18E	19E
Revenue (m)	235,321	345,693	460,442
Operating profit (m)	80,918	106,808	133,351
Net profit (m)	63,324	92,694	116,629
Core EPS (fully-diluted)	6.651	9.736	12.250
EPS change (%)	38.9	46.4	25.8
Daiwa vs Cons. EPS (%)	3.1	17.7	19.8
PER (x)	42.0	28.7	22.8
Dividend yield (%)	0.2	0.3	0.4
DPS	0.662	0.883	1.108
PBR (x)	11.1	7.9	5.6
EV/EBITDA (x)	23.4	19.4	14.5
ROE (%)	31.1	32.7	29.1

Source: FactSet, Daiwa forecasts

See important disclosures, including any required research certifications, beginning on page 79

Financial summary

Key assumptions

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Monthly Active Users (MAU) of WeChat (m)	0	355	500	697	889	978	1,008	1,018
WeChat implied revenues (CNYbn)	0	3	16	33	61	97	145	190
Online game revenue (CNYbn)	23	33	49	64	71	94	118	138
Community value-added services revenue (CNYbn)	9.1	13.0	18.6	24.1	37.0	57.2	83.4	109.5

Profit and loss (CNYm)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Value-added services (VAS)	35,718	47,240	63,310	87,911	107,810	151,161	201,910	247,052
Online advertising	3,382	5,034	8,308	17,468	26,969	41,309	60,223	79,694
Other Revenue	4,793	9,418	7,314	4,726	17,158	42,851	83,560	133,696
Total Revenue	**43,894**	**61,692**	**78,932**	**110,105**	**151,937**	**235,321**	**345,693**	**460,442**
Other income	552	2,219	4,435	4,213	6,213	16,564	16,099	16,892
COGS	(18,207)	(27,778)	(30,873)	(41,631)	(67,439)	(117,843)	(177,549)	(241,765)
SG&A	(8,146)	(12,003)	(16,736)	(21,700)	(21,966)	(28,557)	(62,225)	(81,959)
Other op.expenses	(2,612)	(3,590)	(4,797)	(6,629)	(12,629)	(24,566)	(15,211)	(20,259)
Operating profit	**15,479**	**20,540**	**30,961**	**44,358**	**56,116**	**80,918**	**106,808**	**133,351**
Net-interest inc./(exp.)	836	1,314	494	709	664	502	700	700
Assoc/forex/extraord./others	(1,264)	(2,574)	(2,023)	(5,120)	(5,141)	(3,504)	(2,186)	(1,962)
Pre-tax profit	**15,051**	**19,281**	**29,432**	**39,947**	**51,639**	**77,916**	**105,323**	**132,089**
Tax	(2,266)	(3,718)	(5,125)	(7,108)	(10,193)	(14,750)	(21,065)	(26,418)
Min. int./pref. div./others	53	61	79	173	309	141	141	141
Net profit (reported)	**12,838**	**15,624**	**24,386**	**33,012**	**41,755**	**63,307**	**84,399**	**105,812**
Net profit (adjusted)	**14,286**	**17,063**	**24,224**	**32,374**	**45,462**	**63,324**	**92,694**	**116,629**
EPS (reported)(CNY)	**1.405**	**1.706**	**2.642**	**3.547**	**4.428**	**6.736**	**8.981**	**11.259**
EPS (adjusted)(CNY)	**1.563**	**1.863**	**2.624**	**3.478**	**4.821**	**6.738**	**9.863**	**12.410**
EPS (adjusted fully-diluted)(CNY)	**1.533**	**1.827**	**2.589**	**3.445**	**4.788**	**6.651**	**9.736**	**12.250**
DPS (CNY)	**0.200**	**0.240**	**0.360**	**0.435**	**0.610**	**0.662**	**0.883**	**1.108**
EBIT	**15,479**	**20,540**	**30,961**	**44,358**	**56,116**	**80,918**	**106,808**	**133,351**
EBITDA	**18,092**	**24,130**	**35,758**	**50,987**	**68,745**	**105,485**	**122,019**	**153,611**

Cash flow (CNYm)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Profit before tax	15,051	19,281	29,432	39,947	51,639	77,916	105,323	132,089
Depreciation and amortisation	2,612	3,590	4,818	6,674	8,947	12,214	17,014	23,407
Tax paid	(2,225)	(3,118)	(4,703)	(5,047)	(10,193)	(14,750)	(21,065)	(26,418)
Change in working capital	21,601	27,492	5,556	6,389	5,475	12,319	16,605	17,047
Other operational CF items	(17,610)	(22,871)	(2,392)	(2,532)	4,814	6,238	9,538	12,061
Cash flow from operations	**19,429**	**24,374**	**32,711**	**45,431**	**60,682**	**93,938**	**127,415**	**158,185**
Capex	(3,970)	(4,788)	(4,296)	(5,440)	(7,192)	(11,139)	(16,364)	(21,796)
Net (acquisitions)/disposals	(4,659)	8,059	(33,235)	(13,211)	0	0	0	0
Other investing CF items	(6,772)	(21,112)	9,143	(44,954)	0	0	0	0
Cash flow from investing	**(15,401)**	**(17,841)**	**(28,388)**	**(63,605)**	**(7,192)**	**(11,139)**	**(16,364)**	**(21,796)**
Change in debt	(7,041)	1,334	20,619	26,648	0	0	0	0
Net share issues/(repurchases)	218	(1,017)	(61)	0	0	0	0	0
Dividends paid	(1,225)	(1,536)	(1,761)	(2,640)	(4,048)	(5,714)	(6,224)	(8,306)
Other financing CF items	5,662	2,927	(447)	(5,480)	0	0	0	0
Cash flow from financing	**(2,386)**	**1,708**	**18,350**	**18,528**	**(4,048)**	**(5,714)**	**(6,224)**	**(8,306)**
Forex effect/others	(2,424)	1,507	0	0	0	0	0	0
Change in cash	**(782)**	**9,748**	**22,673**	**354**	**49,442**	**77,084**	**104,828**	**128,083**
Free cash flow	15,460	19,586	28,415	39,991	53,490	82,799	111,051	136,389

Source: FactSet, Daiwa forecasts

Financial summary continued …

Balance sheet (CNYm)

As at 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Cash & short-term investment	29,709	43,982	51,887	98,169	147,611	224,695	329,523	457,606
Inventory	568	1,384	244	222	222	222	222	222
Accounts receivable	2,354	2,955	4,588	7,061	9,744	15,091	22,169	29,528
Other current assets	3,878	5,365	18,602	49,926	56,991	70,790	87,135	104,715
Total current assets	**36,509**	**53,686**	**75,321**	**155,378**	**214,568**	**310,798**	**439,050**	**592,071**
Fixed assets	7,899	10,734	7,918	9,973	12,925	18,146	26,128	36,261
Goodwill & intangibles	4,719	4,103	9,304	13,439	14,810	17,926	23,121	29,796
Other non-current assets	26,128	38,712	78,623	128,028	130,551	123,394	115,668	128,028
Total assets	**75,256**	**107,235**	**171,166**	**306,818**	**372,853**	**470,265**	**603,967**	**786,156**
Short-term debt	1,077	2,589	3,215	11,429	11,429	11,429	11,429	11,429
Accounts payable	4,212	6,680	8,683	15,700	22,898	38,367	57,222	77,195
Other current liabilities	15,376	23,998	38,137	97,277	105,302	121,298	142,471	164,484
Total current liabilities	**20,665**	**33,267**	**50,035**	**124,406**	**139,629**	**171,094**	**211,122**	**253,108**
Long-term debt	7,517	9,141	32,587	53,640	53,640	53,640	53,640	53,640
Other non-current liabilities	4,926	6,364	6,420	6,672	6,672	6,672	6,672	6,672
Total liabilities	**33,108**	**48,772**	**89,042**	**184,718**	**199,941**	**231,406**	**271,434**	**313,420**
Share capital	0	0	0	0	0	0	0	0
Reserves/R.E./others	41,298	57,945	80,013	120,035	170,847	236,793	330,468	470,672
Shareholders' equity	**41,298**	**57,945**	**80,013**	**120,035**	**170,847**	**236,793**	**330,468**	**470,672**
Minority interests	851	518	2,111	2,065	2,065	2,065	2,065	2,065
Total equity & liabilities	**75,256**	**107,235**	**171,166**	**306,818**	**372,853**	**470,265**	**603,967**	**786,156**
EV	2,597,760	2,581,026	2,611,396	2,594,335	2,544,893	2,467,809	2,362,981	2,234,898
Net debt/(cash)	**(21,115)**	**(32,252)**	**(16,085)**	**(33,100)**	**(82,542)**	**(159,626)**	**(264,454)**	**(392,537)**
BVPS (CNY)	**4.457**	**6.224**	**8.668**	**12.907**	**18.238**	**25.186**	**35.150**	**50.062**

Key ratios (%)

Year to 31 Dec	2012	2013	2014	2015	2016	2017E	2018E	2019E
Sales (YoY)	54.0	40.5	27.9	39.5	38.0	54.9	46.9	33.2
EBITDA (YoY)	27.5	33.4	48.2	42.6	34.8	53.4	15.7	25.9
Operating profit (YoY)	26.3	32.7	50.7	43.3	26.5	44.2	32.0	24.9
Net profit (YoY)	30.6	19.4	42.0	33.6	40.4	39.3	46.4	25.8
Core EPS (fully-diluted) (YoY)	30.3	19.1	41.7	33.1	39.0	38.9	46.4	25.8
Gross-profit margin	58.5	55.0	60.9	62.2	55.6	49.9	48.6	47.5
EBITDA margin	41.2	39.1	45.3	46.3	45.2	44.8	35.3	33.4
Operating-profit margin	35.3	33.3	39.2	40.3	36.9	34.4	30.9	29.0
Net profit margin	32.5	27.7	30.7	29.4	29.9	26.9	26.8	25.3
ROAE	41.0	34.4	35.1	32.4	31.3	31.1	32.7	29.1
ROAA	21.6	18.7	17.4	13.5	13.4	15.0	17.3	16.8
ROCE	33.8	34.0	32.9	29.1	26.4	29.9	30.5	28.5
ROIC	85.9	70.2	55.4	47.0	50.2	77.4	116.0	143.9
Net debt to equity	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Effective tax rate	15.1	19.3	17.4	17.8	19.7	18.9	20.0	20.0
Accounts receivable (days)	18.2	15.7	17.4	19.3	20.2	19.3	19.7	20.5
Current ratio (x)	1.8	1.6	1.5	1.2	1.5	1.8	2.1	2.3
Net interest cover (x)	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.	n.a.
Net dividend payout	14.2	14.1	13.6	12.3	13.8	9.8	9.8	9.8
Free cash flow yield	0.6	0.7	1.1	1.5	2.0	3.2	4.2	5.2

Source: FactSet, Daiwa forecasts

Company profile

Tencent is the largest and most used Internet service portal in China. Over the past decade, the company has been able to maintain steady earnings growth under its user-oriented strategy. On 16 June 2004, it was listed on the main board of the Hong Kong Stock Exchange.

A pragmatist on artificial intelligence

Martin Lau, president of Tencent, outlined Tencent's strategic vision on AI during the 2Q17 results call

We believe the remarks made by Martin Lau, executive director and president of Tencent, during the 2Q17 results conference call give a clear outline of the company's vision and strategy on AI. Below we include the full transcript of Martin's speech on AI.

On the strategic importance of AI:

"We have been investing heavily in AI but relatively quietly, as we view AI as an essential capability that enhances the user experience and empowers us to capture the new exciting opportunities to grow our businesses for the future."

Tencent's advantage on AI:

"We're confident that our existing strength in computing power, data, engineering, technologies as well as use cases coupled with our proactive build-up of AI content — talent will give us a favourable position in this strategic initiative.

In particular, a wide and diversified business scope creates a variety of use cases for AI research and application across a range of AI fundamental research areas, such as machine learning, computer vision, speech recognition and natural language processing."

On the investment horizon of AI:

"We will be persistent but patient with our AI investment, because we believe it is a long-term initiative, and we do not necessarily require research to generate revenue directly in the short term. On the other hand, AI will significantly benefit all of our existing products, services and businesses in many ways."

How could Tencent use AI technology?

"For example, in consumer-facing products, AI enhances user experience as we understand more about the users; for enterprise related businesses, AI optimizes monetization as we sharpen our targeting technology; for our ecosystem, our investee companies and partners can leverage our strong AI capability, allowing all of us to achieve mutual benefits.

*For **performance ads**, we applied AI technology to the processes of ad placement from understanding users' preferences, contextual and ad content, to ranking, the bidding price, optimizing the display formats and eventually to matching the most appropriate advertisers. This increases the ROI for advertisers while at the same time enhancing the reading experience for our users.*

*For **information-based services** including news apps, video, music and App Store, AI enables us to have better knowledge of users' interests. This will help us to make more relevant and customized recommendations to users, so that they can access their favourite content more efficiently. Across Tencent's platforms, there are multiple digital content access points which we believe will all benefit from the smarter recommendation engine.*

*For **Internet finance** businesses including mobile payment, wealth management, and microloans, we use AI to predict users' behaviour in financial activities more precisely. This will help us to provide the most suitable products to the most appropriate users and in the process, minimize the risk involved."*

Other achievements in AI:

"Early this year, Fine Art won the UEC Cup, a global Computer Go tournament. Fine Art was developed by our AI lab in less than a year. We have accumulated in this process significant know-how in the development of Fine Art. And the strategy and reinforced learning AI technology behind Fine Art can be applied to many other use cases in the future.

Our face recognition technology also scored excellent results with the world-leading Face Detection Dataset and Benchmark, FDDB. We have gradually applied this technology in a variety of ways. For example, firstly, enhance and enrich the features of our photo editing app Pitu, which has become the number two app of its kind in China. Secondly, enable users to complete ID — identification online for financial services and government, municipal affairs. Thirdly, assisting the search for children and seniors reported missing, helping many families in the process".

Tencent believes AI will improve the user experience and monetisation efficiency at the same time

Tencent: strategies on AI

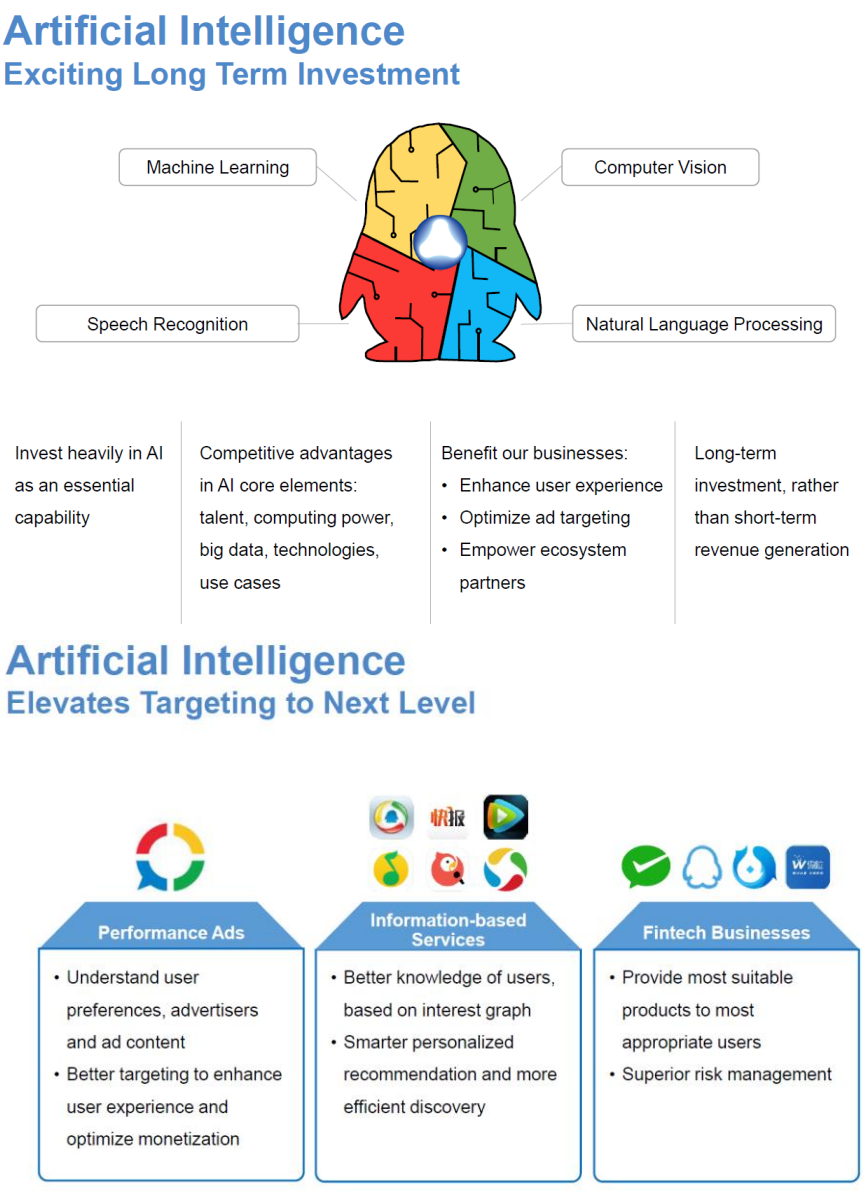

Source: Company

A heavy spender on R&D activities

Tencent has been the biggest spender on research among Internet companies in China, as below chart shows. The company's R&D expenses increased by 45% from 2014 to 2016. We believe Tencent will continue to invest heavily in R&D to secure leadership in machine learning research.

China: major listed Internet companies research expense comparison (2014-1H17)

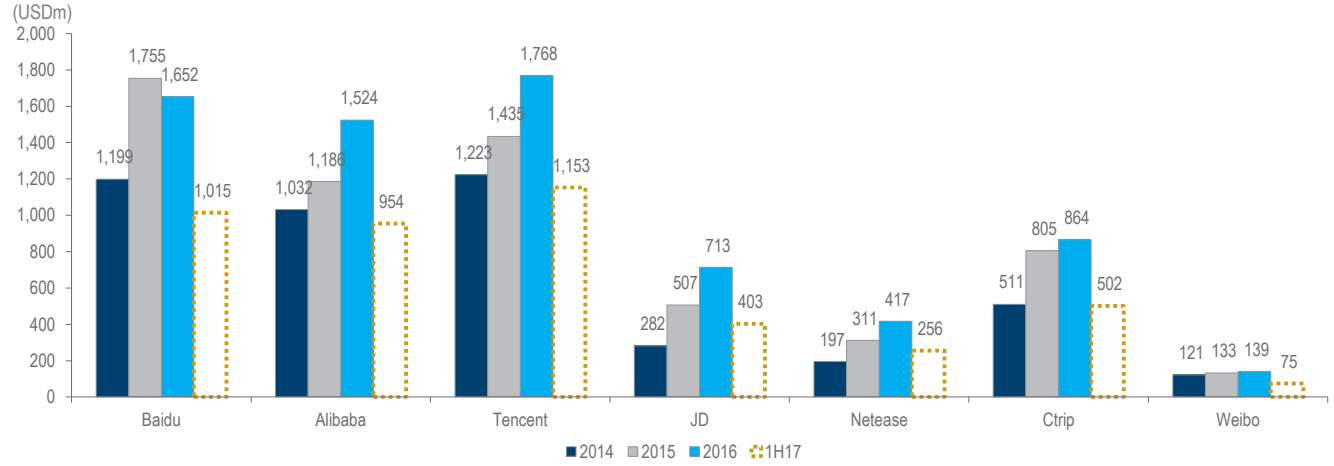

Source: Companies, Daiwa compiled
Note: all expense in non-GAAP (including share-based compensation and intangible amortisation) except for Tencent (Tencent does not disclose non-GAAP expense reconciliation)

Valuation and earnings revisions

We are revising up our 2017-19 earnings forecasts by 2.5-2.7% to factor in slightly higher PC game revenue and other revenue (mostly payment and cloud-related business).

We reaffirm our Buy (1) rating on Tencent and raise our 12-month TP to HKD370 from HKD360, based on an unchanged 39x PER applied to our average 2017-18E non-GAAP EPS.

Tencent: earnings revisions

	OLD			NEW			Difference		
(in CNYbn)	**2017E**	**2018E**	**2019E**	**2016E**	**2017E**	**2018E**	**2016E**	**2017E**	**2018E**
Revenue	**233**	**336**	**443**	**235**	**346**	**460**	*1.1%*	*2.9%*	*3.9%*
YoY%	*53%*	*44%*	*32%*	*55%*	*47%*	*33%*			
Gross profit	**116**	**165**	**213**	**117**	**168**	**219**	*0.9%*	*2.1%*	*2.6%*
YoY%	*38%*	*41%*	*29%*	*39%*	*43%*	*30%*			
Gross margin	*50.0%*	*49.0%*	*48.1%*	*49.9%*	*48.6%*	*47.5%*	*-0.08ppt*	*-0.37ppt*	*-0.60ppt*
Operating income - GAAP	**80**	**105**	**131**	**81**	**107**	**133**	*1.6%*	*1.6%*	*1.7%*
YoY%	*42%*	*32%*	*25%*	*44%*	*32%*	*25%*			
Operating margin	*34.2%*	*31.3%*	*29.6%*	*34.4%*	*30.9%*	*29.0%*	*0.16ppt*	*-0.40ppt*	*-0.63ppt*
Net income - Non-GAAP	**62**	**90**	**114**	**63**	**93**	**117**	*2.7%*	*2.5%*	*2.7%*
YoY%	*36%*	*47%*	*26%*	*39%*	*46%*	*26%*			
Net profits margin - Non-GAAP	*26.5%*	*26.9%*	*25.6%*	*26.9%*	*26.8%*	*25.3%*	*0.41ppt*	*-0.10ppt*	*-0.31ppt*
Diluted income per ADS (CNY) - Non-GAAP	**6.48**	**9.50**	**11.93**	**6.65**	**9.74**	**12.25**	*2.7%*	*2.5%*	*2.7%*

Source: Daiwa forecasts

In our view, the key downside risk for the stock is higher-than-expected marketing and content expenses. Other risks include weaker mobile games performance and higher-than-expected R&D spending.


Daiwa
Capital Markets


Daiwa
Capital Markets

Daiwa's Asia Pacific Research Directory

HONG KONG		
Takashi FUJIKURA	(852) 2848 4051	takashi.fujikura@hk.daiwacm.com
Regional Research Head		
Jiro IOKIBE	(852) 2773 8702	jiro.iokibe@hk.daiwacm.com
Co-head of Asia Pacific Research		
John HETHERINGTON	(852) 2773 8787	john.hetherington@hk.daiwacm.com
Co-head of Asia Pacific Research		
Craig CORK	(852) 2848 4463	craig.cork@hk.daiwacm.com
Regional Head of Asia Pacific Product Management		
Kevin LAI	(852) 2848 4926	kevin.lai@hk.daiwacm.com
Chief Economist for Asia ex-Japan; Macro Economics (Regional)		
Olivia XIA	(852) 2773 8736	olivia.xia@hk.daiwacm.com
Macro Economics (Hong Kong/China)		
Kelvin LAU	(852) 2848 4467	kelvin.lau@hk.daiwacm.com
Head of Automobiles; Transportation and Industrial (Hong Kong/China)		
Leon QI	(852) 2532 4381	leon.qi@hk.daiwacm.com
Regional Head of Financials; Banking; Diversified financials; Insurance (Hong Kong/China)		
Yan LI	(852) 2773 8822	yan.li@hk.daiwacm.com
Banking (China)		
Anson CHAN	(852) 2532 4350	anson.chan@hk.daiwacm.com
Consumer (Hong Kong/China)		
Adrian CHAN	(852) 2848 4427	adrian.chan@hk.daiwacm.com
Consumer (Hong Kong/China)		
Jamie SOO	(852) 2773 8529	jamie.soo@hk.daiwacm.com
Gaming and Leisure (Hong Kong/China)		
John CHOI	(852) 2773 8730	john.choi@hk.daiwacm.com
Head of Hong Kong and China Internet; Regional Head of Small/Mid Cap		
Alex LIU	(852) 2848 4976	alex.liu@hk.daiwacm.com
Internet (Hong Kong/China)		
Carlton LAI	(852) 2532 4349	carlton.lai@hk.daiwacm.com
Small/Mid Cap (Hong Kong/China)		
Dennis IP	(852) 2848 4068	dennis.ip@hk.daiwacm.com
Regional Head of Power, Utilities, Renewable and Environment (PURE); PURE (Hong Kong/China)		
Jonas KAN	(852) 2848 4439	jonas.kan@hk.daiwacm.com
Head of Hong Kong and China Property		
Cynthia CHAN	(852) 2773 8243	cynthia.chan@hk.daiwacm.com
Property (China)		
Thomas HO	(852) 2773 8716	thomas.ho@hk.daiwacm.com
Custom Products Group		

PHILIPPINES		
Micaela ABAQUITA	(63) 2 737 3021	micaela.abaquita@dbpdaiwacm.com.ph
Property		
Gregg Ilag	(63) 2 737 3023	gregg.ilag@dbpdaiwacm.com.ph
Utilities; Energy		

SOUTH KOREA		
Sung Yop CHUNG	(82) 2 787 9157	sychung@kr.daiwacm.com
Pan-Asia Co-head/Regional Head of Automobiles and Components; Automobiles; Shipbuilding; Steel		
Mike OH	(82) 2 787 9179	mike.oh@kr.daiwacm.com
Banking; Capital Goods (Construction and Machinery)		
Iris PARK	(82) 2 787 9165	iris.park@kr.daiwacm.com
Consumer/Retail		
SK KIM	(82) 2 787 9173	sk.kim@kr.daiwacm.com
IT/Electronics – Semiconductor/Display and Tech Hardware		
Thomas Y KWON	(82) 2 787 9181	yskwon@kr.daiwacm.com
Pan-Asia Head of Internet & Telecommunications; Software – Internet/On-line Games		

TAIWAN		
Rick HSU	(886) 2 8758 6261	rick.hsu@daiwacm-cathay.com.tw
Head of Regional Technology; Head of Taiwan Research; Semiconductor/IC Design (Regional)		
Nora HOU	(886) 2 8758 6249	nora.hou@daiwacm-cathay.com.tw
Banking; Diversified financials; Insurance		
Steven TSENG	(886) 2 8758 6252	steven.tseng@daiwacm-cathay.com.tw
IT/Technology Hardware (PC Hardware)		
Kylie HUANG	(886) 2 8758 6248	kylie.huang@daiwacm-cathay.com.tw
IT/Technology Hardware (Handsets and Components)		
Helen CHIEN	(886) 2 8758 6254	helen.chien@daiwacm-cathay.com.tw
Small/Mid Cap		

INDIA		
Punit SRIVASTAVA	(91) 22 6622 1013	punit.srivastava@in.daiwacm.com
Head of India Research; Strategy; Banking/Finance		
Saurabh MEHTA	(91) 22 6622 1009	saurabh.mehta@in.daiwacm.com
Capital Goods; Utilities		

SINGAPORE		
Ramakrishna MARUVADA	(65) 6499 6543	ramakrishna.maruvada@sg.daiwacm.com
Head of Singapore Research; Telecommunications (China/ASEAN/India)		
David LUM	(65) 6329 2102	david.lum@sg.daiwacm.com
Banking; Property and REITs		
Royston TAN	(65) 6321 3086	royston.tan@sg.daiwacm.com
Oil and Gas; Capital Goods		
Peter NG	(65) 6499 6546	peter.ng@sg.daiwacm.com
Property; Small/Mid Cap		
Jame OSMAN	(65) 6321 3092	jame.osman@sg.daiwacm.com
Transportation – Road and Rail; Pharmaceuticals and Healthcare; Consumer (Singapore)		

Daiwa's Offices

Office / Branch / Affiliate	Address	Tel	Fax
DAIWA SECURITIES GROUP INC			
HEAD OFFICE	Gran Tokyo North Tower, 1-9-1, Marunouchi, Chiyoda-ku, Tokyo, 100-6753	(81) 3 5555 3111	(81) 3 5555 0661
Daiwa Securities Trust Company	One Evertrust Plaza, Jersey City, NJ 07302, U.S.A.	(1) 201 333 7300	(1) 201 333 7726
Daiwa Securities Trust and Banking (Europe) PLC (Head Office)	5 King William Street, London EC4N 7JB, United Kingdom	(44) 207 320 8000	(44) 207 410 0129
Daiwa Europe Trustees (Ireland) Ltd	Level 3, Block 5, Harcourt Centre, Harcourt Road, Dublin 2, Ireland	(353) 1 603 9900	(353) 1 478 3469
Daiwa Capital Markets America Inc. New York Head Office	Financial Square, 32 Old Slip, New York, NY10005, U.S.A.	(1) 212 612 7000	(1) 212 612 7100
Daiwa Capital Markets America Inc. San Francisco Branch	555 California Street, Suite 3360, San Francisco, CA 94104, U.S.A.	(1) 415 955 8100	(1) 415 956 1935
Daiwa Capital Markets Europe Limited, London Head Office	5 King William Street, London EC4N 7AX, United Kingdom	(44) 20 7597 8000	(44) 20 7597 8600
Daiwa Capital Markets Europe Limited, Frankfurt Branch	Neue Mainzer Str. 1, 60311 Frankfurt/Main, Germany	(49) 69 717 080	(49) 69 723 340
Daiwa Capital Markets Europe Limited, Paris Representative Office	17, rue de Surène 75008 Paris, France	(33) 1 56 262 200	(33) 1 47 550 808
Daiwa Capital Markets Europe Limited, Geneva Branch	50 rue du Rhône, P.O.Box 3198, 1211 Geneva 3, Switzerland	(41) 22 818 7400	(41) 22 818 7441
Daiwa Capital Markets Europe Limited, Moscow Representative Office	Midland Plaza 7th Floor, 10 Arbat Street, Moscow 119002, Russian Federation	(7) 495 641 3416	(7) 495 775 6238
Daiwa Capital Markets Europe Limited, Bahrain Branch	7th Floor, The Tower, Bahrain Commercial Complex, P.O. Box 30069, Manama, Bahrain	(973) 17 534 452	(973) 17 535 113
Daiwa Capital Markets Hong Kong Limited	Level 28, One Pacific Place, 88 Queensway, Hong Kong	(852) 2525 0121	(852) 2845 1621
Daiwa Capital Markets Singapore Limited	6 Shenton Way #26-08, OUE Downtown 2, Singapore 068809, Republic of Singapore	(65) 6220 3666	(65) 6223 6198
Daiwa Capital Markets Australia Limited	Level 34, Rialto North Tower, 525 Collins Street, Melbourne, Victoria 3000, Australia	(61) 3 9916 1300	(61) 3 9916 1330
DBP-Daiwa Capital Markets Philippines, Inc	18th Floor, Citibank Tower, 8741 Paseo de Roxas, Salcedo Village, Makati City, Republic of the Philippines	(632) 813 7344	(632) 848 0105
Daiwa-Cathay Capital Markets Co Ltd	14/F, 200, Keelung Road, Sec 1, Taipei, Taiwan, R.O.C.	(886) 2 2723 9698	(886) 2 2345 3638
Daiwa Securities Capital Markets Korea Co., Ltd.	20 Fl.& 21Fl. One IFC, 10 Gukjegeumyung-Ro, Yeongdeungpo-gu, Seoul, Korea	(82) 2 787 9100	(82) 2 787 9191
Daiwa Securities Co. Ltd., Beijing Representative Office	Room 301/302，Kerry Center，1 Guanghua Road，Chaoyang District，Beijing 100020, People's Republic of China	(86) 10 6500 6688	(86) 10 6500 3594
Daiwa (Shanghai) Corporate Strategic Advisory Co. Ltd.	44/F, Hang Seng Bank Tower, 1000 Lujiazui Ring Road, Pudong, Shanghai China 200120 , People's Republic of China	(86) 21 3858 2000	(86) 21 3858 2111
Daiwa Securities Co. Ltd., Bangkok Representative Office	18th Floor, M Thai Tower, All Seasons Place, 87 Wireless Road, Lumpini, Pathumwan, Bangkok 10330, Thailand	(66) 2 252 5650	(66) 2 252 5665
Daiwa Capital Markets India Private Ltd	10th Floor, 3 North Avenue, Maker Maxity, Bandra Kurla Complex, Bandra East, Mumbai – 400051, India	(91) 22 6622 1000	(91) 22 6622 1019
Daiwa Securities Co. Ltd., Hanoi Representative Office	Suite 405, Pacific Palace Building, 83B, Ly Thuong Kiet Street, Hoan Kiem Dist. Hanoi, Vietnam	(84) 4 3946 0460	(84) 4 3946 0461
DAIWA INSTITUTE OF RESEARCH LTD			
HEAD OFFICE	15-6, Fuyuki, Koto-ku, Tokyo, 135-8460, Japan	(81) 3 5620 5100	(81) 3 5620 5603
MARUNOUCHI OFFICE	Gran Tokyo North Tower, 1-9-1, Marunouchi, Chiyoda-ku, Tokyo, 100-6756	(81) 3 5555 7011	(81) 3 5202 2021
New York Research Center	11th Floor, Financial Square, 32 Old Slip, NY, NY 10005-3504, U.S.A.	(1) 212 612 6100	(1) 212 612 8417
London Research Centre	3/F, 5 King William Street, London, EC4N 7AX, United Kingdom	(44) 207 597 8000	(44) 207 597 8550

Important Disclosures and Disclaimer

This publication is produced by Daiwa Securities Group Inc. and/or its non-U.S. affiliates, and distributed by Daiwa Securities Group Inc. and/or its non-U.S. affiliates, except to the extent expressly provided herein. This publication and the contents hereof are intended for information purposes only, and may be subject to change without further notice. Any use, disclosure, distribution, dissemination, copying, printing or reliance on this publication for any other purpose without our prior consent or approval is strictly prohibited. Neither Daiwa Securities Group Inc. nor any of its respective parent, holding, subsidiaries or affiliates, nor any of its respective directors, officers, servants and employees, represent nor warrant the accuracy or completeness of the information contained herein or as to the existence of other facts which might be significant, and will not accept any responsibility or liability whatsoever for any use of or reliance upon this publication or any of the contents hereof. Neither this publication, nor any content hereof, constitute, or are to be construed as, an offer or solicitation of an offer to buy or sell any of the securities or investments mentioned herein in any country or jurisdiction nor, unless expressly provided, any recommendation or investment opinion or advice. Any view, recommendation, opinion or advice expressed in this publication may not necessarily reflect those of Daiwa Securities Group Inc., and/or its affiliates nor any of its respective directors, officers, servants and employees except where the publication states otherwise. This research report is not to be relied upon by any person in making any investment decision or otherwise advising with respect to, or dealing in, the securities mentioned, as it does not take into account the specific investment objectives, financial situation and particular needs of any person.

Daiwa Securities Group Inc., its subsidiaries or affiliates, or its or their respective directors, officers and employees from time to time have trades as principals, or have positions in, or have other interests in the securities of the company under research including market making activities, derivatives in respect of such securities or may have also performed investment banking and other services for the issuer of such securities. Daiwa Securities Group Inc., its subsidiaries or affiliates do and seek to do business with the company(s) covered in this research report. Therefore, investors should be aware that a conflict of interest may exist. The following are additional disclosures.

Ownership of Securities

For "Ownership of Securities" information, please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

Investment Banking Relationship

For "Investment Banking Relationship", please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

Japan

Daiwa Securities Co. Ltd. and Daiwa Securities Group Inc.

Daiwa Securities Co. Ltd. is a subsidiary of Daiwa Securities Group Inc.

Investment Banking Relationship

Within the preceding 12 months, the subsidiaries and/or affiliates of Daiwa Securities Group Inc. * has lead-managed public offerings and/or secondary offerings (excluding straight bonds) of the securities of the following companies: Neo Solar Power Corp (3576 TT), Acushnet Holdings Corp (GOLF US), No Va Land Investment Group Corporation (NVL VN), PT Totalindo Eka Persada Tbk (TOPS IJ), PT Integra Indocabinet Tbk (WOOD IJ) and PT Buyung Putera Sembada (HOKI IJ).
*Subsidiaries of Daiwa Securities Group Inc. for the purposes of this section shall mean any one or more of: Daiwa Capital Markets Hong Kong Limited (大和資本市場香港有限公司), Daiwa Capital Markets Singapore Limited, Daiwa Capital Markets Australia Limited, Daiwa Capital Markets India Private Limited, Daiwa-Cathay Capital Markets Co., Ltd., Daiwa Securities Capital Markets Korea Co., Ltd.

Hong Kong

This research is distributed in Hong Kong by Daiwa Capital Markets Hong Kong Limited (大和資本市場香港有限公司) ("DHK") which is regulated by the Hong Kong Securities and Futures Commission. Recipients of this research in Hong Kong may contact DHK in respect of any matter arising from or in connection with this research.

Relevant Relationship (DHK)

DHK may from time to time have an individual employed by or associated with it serves as an officer of any of the companies under its research coverage.

Singapore

This research is distributed in Singapore by Daiwa Capital Markets Singapore Limited and it may only be distributed in Singapore to accredited investors, expert investors and institutional investors as defined in the Financial Advisers Regulations and the Securities and Futures Act (Chapter 289), as amended from time to time. By virtue of distribution to these category of investors, Daiwa Capital Markets Singapore Limited and its representatives are not required to comply with Section 36 of the Financial Advisers Act (Chapter 110) (Section 36 relates to disclosure of Daiwa Capital Markets Singapore Limited's interest and/or its representative's interest in securities). Recipients of this research in Singapore may contact Daiwa Capital Markets Singapore Limited in respect of any matter arising from or in connection with the research.

Australia

This research is distributed in Australia by Daiwa Capital Markets Australia Limited and it may only be distributed in Australia to wholesale investors within the meaning of the Corporations Act. Recipients of this research in Australia may contact Daiwa Capital Markets Stockbroking Limited in respect of any matter arising from or in connection with the research.

India

This research is distributed in India to Institutional Clients only by Daiwa Capital Markets India Private Limited (Daiwa India) which is an intermediary registered with Securities & Exchange Board of India as a Stock Broker, Merchant Bank and Research Analyst. Daiwa India, its Research Analyst and their family members and its associates do not have any financial interest save as disclosed or other undisclosed material conflict of interest in the securities or derivatives of any companies under coverage. Daiwa India and its associates, may have received compensation for any products other than Investment Banking (as disclosed)or brokerage services from the subject company in this report or from any third party during the past 12 months. Daiwa India and its associates may have debt holdings in the subject company. For information on ownership of equity, please visit BlueMatrix disclosure Link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

There is no material disciplinary action against Daiwa India by any regulatory authority impacting equity research analysis activities as of the date of this report.

Associates of Daiwa India, registered with Indian regulators, include Daiwa Capital Markets Singapore Limited and Daiwa Portfolio Advisory (India) Private Limited.

Taiwan

This research is solely for reference and not intended to provide tailored investment recommendations. This research is distributed in Taiwan by Daiwa-Cathay Capital Markets Co., Ltd. and it may only be distributed in Taiwan to specific customers who have signed recommendation contracts with Daiwa-Cathay Capital Markets Co., Ltd. and non-customers including (i) professional institutional investors, (ii) TWSE or TPEx listed companies, upstream and downstream vendors, and specialists that offer or seek advice, and (iii) potential customers with an actual need for business development in accordance with the Operational Regulations Governing Securities Firms Recommending Trades in Securities to Customers. Recipients of this research including non-customer recipients of this research shall not provide it to others or engage in any activities in connection with this research which may involve conflicts of interests. Neither Daiwa-Cathay Capital Markets Co., Ltd. nor its personnel who writes or reviews the research report has any conflict of interest in this research. Since Daiwa-Cathay Capital Markets Co., Ltd. does not operate brokerage trading business in foreign markets, **this research is prepared on a "without recommendation" to any foreign securities basis** and Daiwa-Cathay Capital Markets Co., Ltd. does not accept orders from customers to trade in such foreign securities that are without recommendation. Recipients of this research in Taiwan may contact Daiwa-Cathay Capital Markets Co., Ltd. in respect of any matter arising from or in connection with the research.

Philippines

This research is distributed in the Philippines by DBP-Daiwa Capital Markets Philippines, Inc. which is regulated by the Philippines Securities and Exchange Commission and the Philippines Stock Exchange, Inc. Recipients of this research in the Philippines may contact DBP-Daiwa Capital Markets Philippines, Inc. in respect of any matter arising from or in connection with the research. DBP-Daiwa Capital Markets Philippines, Inc. recommends that investors independently assess, with a professional advisor, the specific financial risks as well as the legal, regulatory, tax, accounting, and other consequences of a proposed transaction. DBP-Daiwa Capital Markets Philippines, Inc. may have positions or may be materially interested in the securities in any of the markets mentioned in the publication or may have performed other services for the issuers of such securities.

For relevant securities and trading rules please visit SEC and PSE links at *http://www.sec.gov.ph* and *http://www.pse.com.ph/* respectively.

Thailand

This research is distributed to only institutional investors in Thailand primarily by Thanachart Securities Public Company Limited ("TNS").

This report is prepared by analysts who are employed by Daiwa Securities Group Inc. and/or its non-U.S. affiliates. This report is provided to you for informational purposes only and it is not, and is not to be construed as, an offer or an invitation to make an offer to sell or buy any securities. Neither TNS, Daiwa Securities Group Inc. nor any of their respective parent, holding, subsidiaries or affiliates, nor any of their respective directors, officers, servants and employees accept any liability whatsoever for any direct or consequential loss arising from any use of this research or its contents.

The information and opinions contained herein have been compiled or arrived at from sources believed to be reliable. However, TNS, Daiwa Securities Group Inc. nor any of their respective parent, holding, subsidiaries or affiliates, nor any of their respective directors, officers, servants and employees make no representation or warranty, express or implied, as to their accuracy or completeness. Expressions of opinion herein are subject to change without notice. The use of any information, forecasts and opinions contained in this report shall be at the sole discretion and risk of the user.

TNS, Daiwa Securities Group Inc., their respective parent, holding, subsidiaries or affiliates, their respective directors, officers, servants and employees may have positions and financial interest in securities mentioned in this research. Thanachart Securities Public Company Limited, Daiwa Securities Group Inc., their respective parent, holding, subsidiaries or affiliates may from time to time perform investment banking or other services for, or solicit investment banking or other business from, any entity mentioned in this research. Therefore, investors should be aware of conflict of interest that may affect the objectivity of this research.

United Kingdom
This research report is produced by Daiwa Securities Co. Ltd. and/or its affiliates and is distributed in the European Union, Iceland, Liechtenstein, Norway and Switzerland. Daiwa Capital Markets Europe Limited is authorised and regulated by The Financial Conduct Authority ("FCA") and is a member of the London Stock Exchange and Eurex. This publication is intended for investors who are not Retail Clients in the United Kingdom within the meaning of the Rules of the FCA and should not therefore be distributed to such Retail Clients in the United Kingdom. Should you enter into investment business with Daiwa Capital Markets Europe's affiliates outside the United Kingdom, we are obliged to advise that the protection afforded by the United Kingdom regulatory system may not apply; in particular, the benefits of the Financial Services Compensation Scheme may not be available.

Daiwa Capital Markets Europe Limited has in place organisational arrangements for the prevention and avoidance of conflicts of interest. Our conflict management policy is available at http://www.uk.daiwacm.com/about-us/corporate-governance-regulatory.

Germany
This document is distributed in Germany by Daiwa Capital Markets Europe Limited, Niederlassung Frankfurt which is regulated by BaFin (Bundesanstalt fuer Finanzdienstleistungsaufsicht) for the conduct of business in Germany.

Bahrain
This research material is distributed in Bahrain by Daiwa Capital Markets Europe Limited, Bahrain Branch, regulated by The Central Bank of Bahrain and holds Investment Business Firm – Category 2 license and having its official place of business at the Bahrain World Trade Centre, South Tower, 7th floor, P.O. Box 30069, Manama, Kingdom of Bahrain. Tel No. +973 17534452 Fax No. +973 535113

United States
This report is distributed in the U.S. by Daiwa Capital Markets America Inc. (DCMA). It may not be accurate or complete and should not be relied upon as such. It reflects the preparer's views at the time of its preparation, but may not reflect events occurring after its preparation; nor does it reflect DCMA's views at any time. Neither DCMA nor the preparer has any obligation to update this report or to continue to prepare research on this subject. This report is not an offer to sell or the solicitation of any offer to buy securities. Unless this report says otherwise, any recommendation it makes is risky and appropriate only for sophisticated speculative investors able to incur significant losses. Readers should consult their financial advisors to determine whether any such recommendation is consistent with their own investment objectives, financial situation and needs. This report does not recommend to U.S. recipients the use of any of DCMA's non-U.S. affiliates to effect trades in any security and is not supplied with any understanding that U.S. recipients of this report will direct commission business to such non-U.S. entities. Unless applicable law permits otherwise, non-U.S. customers wishing to effect a transaction in any securities referenced in this material should contact a Daiwa entity in their local jurisdiction. Most countries throughout the world have their own laws regulating the types of securities and other investment products which may be offered to their residents, as well as a process for doing so. As a result, the securities discussed in this report may not be eligible for sales in some jurisdictions. Customers wishing to obtain further information about this report should contact DCMA: Daiwa Capital Markets America Inc., Financial Square, 32 Old Slip, New York, New York 10005 (Tel no. 212-612-7000).

Ownership of Securities
For "Ownership of Securities" information please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

Investment Banking Relationships
For "Investment Banking Relationships" please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

DCMA Market Making
For "DCMA Market Making" please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*.

Research Analyst Conflicts
For updates on "Research Analyst Conflicts" please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*. The principal research analysts who prepared this report have no financial interest in securities of the issuers covered in the report, are not (nor are any members of their household) an officer, director or advisory board member of the issuer(s) covered in the report, and are not aware of any material relevant conflict of interest involving the analyst or DCMA, and did not receive any compensation from the issuer during the past 12 months except as noted: no exceptions.

Research Analyst Certification
For updates on "Research Analyst Certification" and "Rating System" please visit BlueMatrix disclosure link at *https://daiwa3.bluematrix.com/sellside/Disclosures.action*. The views about any and all of the subject securities and issuers expressed in this Research Report accurately reflect the personal views of the research analyst(s) primarily responsible for this report (or the views of the firm producing the report if no individual analyst is named on the report); and no part of the compensation of such analyst (or no part of the compensation of the firm if no individual analyst is named on the report) was, is, or will be directly or indirectly related to the specific recommendations or views contained in this Research Report.

The following explains the rating system in the report as compared to relevant local indices, unless otherwise stated, based on the beliefs of the author of the report.
"1": the security could outperform the local index by more than 15% over the next 12 months.
"2": the security is expected to outperform the local index by 5-15% over the next 12 months.
"3": the security is expected to perform within 5% of the local index (better or worse) over the next 12 months.
"4": the security is expected to underperform the local index by 5-15% over the next 12 months.
"5": the security could underperform the local index by more than 15% over the next 12 months.

Disclosure of investment ratings

Rating	Percentage of total
Buy*	66.1%
Hold**	21.6%
Sell***	12.3%

Source: Daiwa
Notes: data is for single-branded Daiwa research in Asia (ex Japan) and correct as of 30 June 2017.
** comprised of Daiwa's Buy and Outperform ratings.*
*** comprised of Daiwa's Hold ratings.*
**** comprised of Daiwa's Underperform and Sell ratings.*

Additional information may be available upon request.

Japan - additional notification items pursuant to Article 37 of the Financial Instruments and Exchange Law
(This Notification is only applicable where report is distributed by Daiwa Securities Co. Ltd.)

If you decide to enter into a business arrangement with us based on the information described in materials presented along with this document, we ask you to pay close attention to the following items.
- In addition to the purchase price of a financial instrument, we will collect a trading commission* for each transaction as agreed beforehand with you. Since commissions may be included in the purchase price or may not be charged for certain transactions, we recommend that you confirm the commission for each transaction.
- In some cases, we may also charge a maximum of ¥ 2 million (including tax) per year as a standing proxy fee for our deposit of your securities, if you are a non-resident of Japan.
- For derivative and margin transactions etc., we may require collateral or margin requirements in accordance with an agreement made beforehand with you. Ordinarily in such cases, the amount of the transaction will be in excess of the required collateral or margin requirements.
- There is a risk that you will incur losses on your transactions due to changes in the market price of financial instruments based on fluctuations in interest rates, exchange rates, stock prices, real estate prices, commodity prices, and others. In addition, depending on the content of the transaction, the loss could exceed the amount of the collateral or margin requirements.
- There may be a difference between bid price etc. and ask price etc. of OTC derivatives handled by us.
- Before engaging in any trading, please thoroughly confirm accounting and tax treatments regarding your trading in financial instruments with such experts as certified public accountants.
*The amount of the trading commission cannot be stated here in advance because it will be determined between our company and you based on current market conditions and the content of each transaction etc.

When making an actual transaction, please be sure to carefully read the materials presented to you prior to the execution of agreement, and to take responsibility for your own decisions regarding the signing of the agreement with us.

Corporate Name: Daiwa Securities Co. Ltd.
Financial instruments firm: chief of Kanto Local Finance Bureau (Kin-sho) No.108
Memberships: Japan Securities Dealers Association, The Financial Futures Association of Japan
Japan Securities Investment Advisers Association
Type II Financial Instruments Firms Association