Facebook's digital assistant blends AI with customer service staff – but will it cope without human help?

September 3 2015, by Leslie Smith

M – no Bond jokes please. Credit: Facebook

With the arrival of its monosyllabic M, Facebook has introduced its own personal digital assistant, following those from Apple (Siri), Microsoft (Cortana), Google (Now) and Amazon (Echo). Technically, M operates partly on the user's smartphone via the Facebook Messenger app, but it is mostly a cloud-based service. Unlike the others, however, this isn't just an artificial intelligence but a mix of smart machine learning and human assistance.

What makes M different is that it takes recommendations or answering queries one step further, able to actually make purchases or arrange services for you, and order deliveries. This is the logical conclusion of recommending something, allowing the system to spend your money for you as well. This approach might be risky, or might be brilliant. If it works, suppliers will be clamouring for Facebook's M to spend users' money with them, and Facebook will be able to take a percentage in return.

With Facebook's enormous reach – the site recently claimed one billion concurrent users – even a small percentage of such a large number of users spending even relatively small sums of money would still add up to a great deal of cash for Facebook. Mind you, a few unfortunate misunderstandings of what a user wants to buy might lead to some negative publicity – and one can imagine some Facebook users attempting some very dubious transactions.

Technical and human intelligence

Under the hood, it appears Facebook is not using cutting-edge AI. While its digital assistant's interface is stored and run from users' phones, the processing occurs on Facebook's servers in the cloud where computing power and data can be distributed. It uses technology from wit.ai, which is understood to use conditional random fields, a popular statistical technique dating from the 2000s, and maximum entropy classifiers,

based on information theory. These pick up on the structure of the data, and use this to make predictions. These may not be cutting edge, but they are well established and understood. Not only that, but they can use prior knowledge, and one of M's aims is to improve and to get better through training.

There's a huge amount of contextual information about the user's likes and preferences within Facebook's enormous datasets, and this could help M's algorithms provide answers. It could also be used to help constrain queries – things to exclude – particularly if both the purchaser and the recipient are Facebook users. But it will take leading edge AI techniques like sentic technologies, which attempt to extract mood, emotion, intention and meaning from text, in order to mine the full value of the text and image datasets generated by Facebook users.
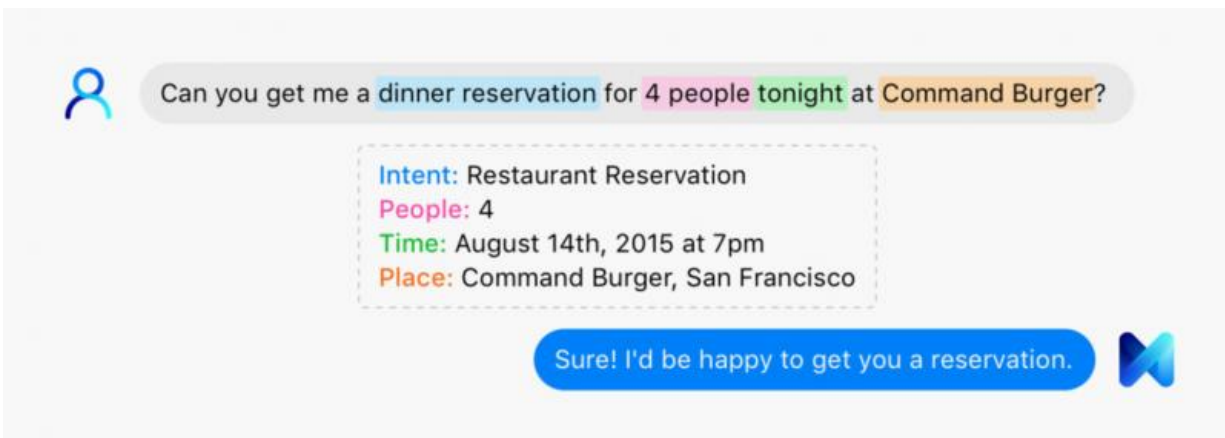
M's natural language processing picks out a message's intent. But it has a lot to learn. Credit: Facebook

David Marcus, vice president of messaging products at Facebook and in charge of M, has said that without explicit consent M won't embark on

such data-mining. In fact there is a limited range of possible services and purchases that the software can perform automatically, while trickier tasks are carried out by the human element behind the scenes - customer service staff working for Facebook. Humans are needed to be able to cover the gaps in the AI's ability to understand natural language, understanding what users are after, able to sign off purchases to ensure they're reasonable, and legal.

While the idea is that M learns the right behaviours by associating the user's intent with the solutions provided by human staff, for this to scale to even a fraction of Facebook Messenger's 700,000 users, the AI will have to be good enough to relieve the human staff of their role. And that may take a while. Of course, M is being rolled out area by area – currently only San Francisco, of course – so perhaps the firm is just dipping a toe in the water to start with.

So while M may be the personal assistant of the future, at the moment it's a curious mix of machine learning, automation, and human comprehension. But powered by the tutoring of actual humans and human-created data, in time it could still become more adept than the competition.

This story is published courtesy of The Conversation *(under Creative Commons-Attribution/No derivatives).*

Source: The Conversation