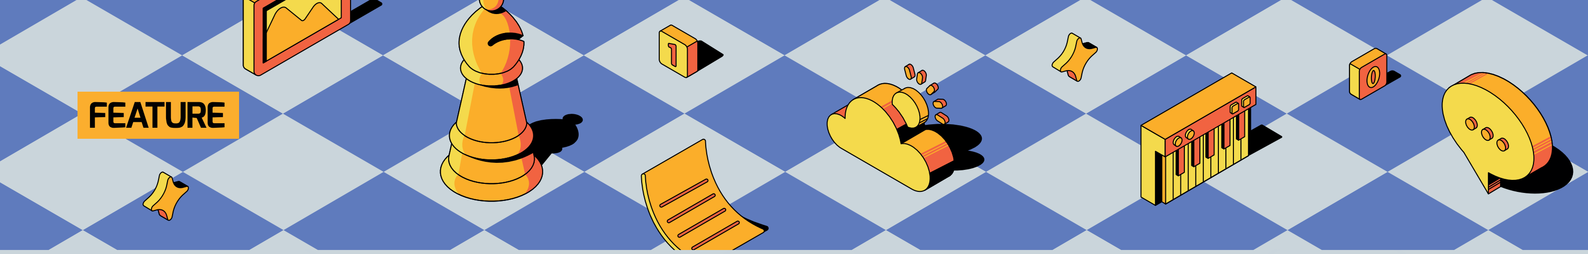

FEATURE

RECOMMENDER SYSTEMS
THE HIDDEN AI STEERING OUR CHOICES

Edward Abel explores recommender systems, what they are, and how they operate, and looks at wider considerations around their growing influence in our digital lives

Every time we visit video websites such as YouTube or TikTok, or music streaming sites like Spotify or Apple Music, we are presented with suggested content to consume. Such recommended suggestions are invariably generated by recommender (or recommendation) systems, and these systems are powered by artificial intelligence (AI) tools.

User personalisation

Recommendation systems look to make use of information about the previous videos or songs we have consumed, together with data about the characteristics of all the available videos or songs, to suggest personal recommendations that it thinks we will like. For example, for every video on YouTube, it documents data about the topic, the length, who features in it, and so on. Additionally, it keeps a log of every video a user watches, along with information about how they watched it, such as whether they stopped before the end, watched early in the morning, or watched with no sound. Its aim is then to use such information to suggest more content to a user. But how do recommender systems choose what content to suggest?

Recommender system approaches

Two of the most common approaches for making recommendations are collaborative filtering and content-based filtering.

Collaborative filtering aims to find similarity patterns between users; two users who liked the same items (such as videos on YouTube or songs on Spotify) in the past will probably both like similar items in the future. In effect, the system tells us, 'Other people who liked the same 200 items as you also liked this item, and we think you will like it too', which has been shown to be a successful way to suggest content a user will like. However, it can struggle to make good recommendations for new users, as it relies on user histories to determine patterns with other users. Additionally, it can struggle to provide meaningful explanations as to why it makes its suggestions.

Content-based filtering works by comparing items and suggesting similar ones. To do this, it uses information that describes each item. For a movie, for example, this information could include its genre, actors, running length, and much more. Then, when

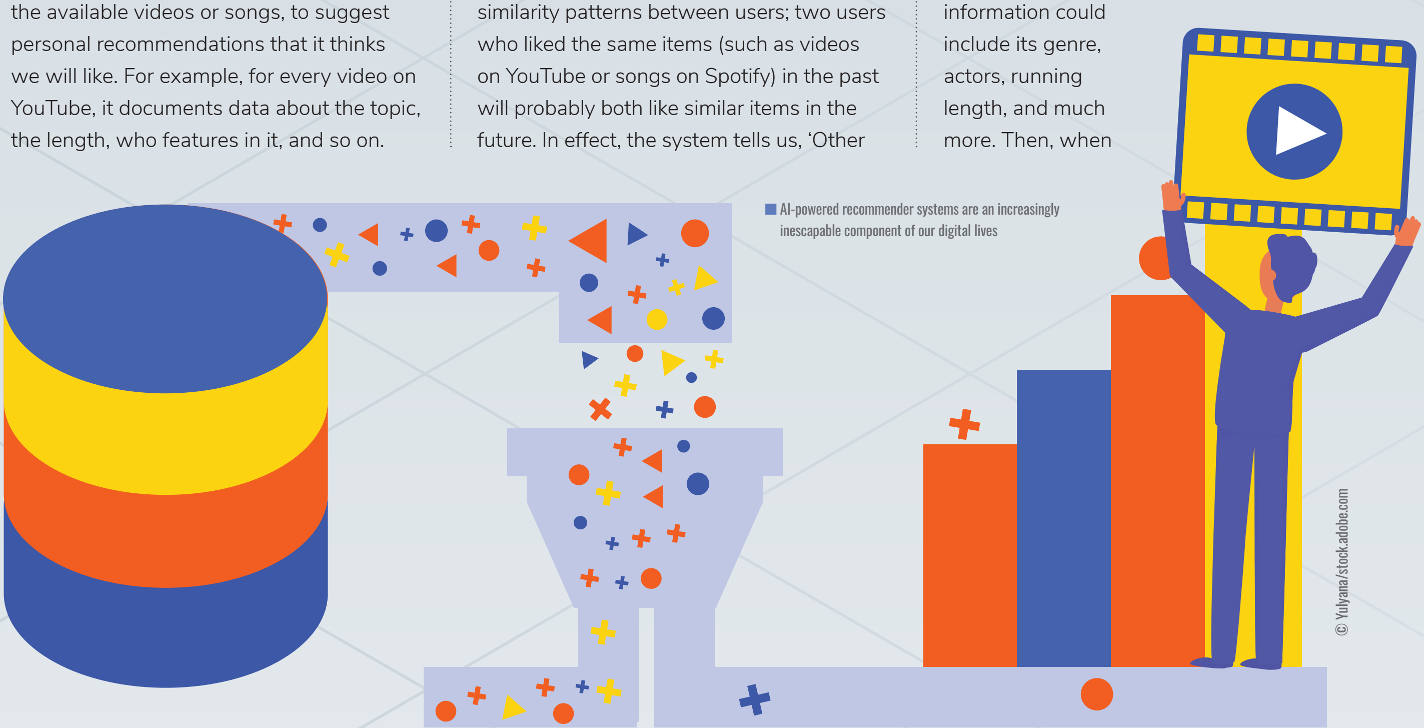

AI-powered recommender systems are an increasingly inescapable component of our digital lives

© Yulyana/stock.adobe.com

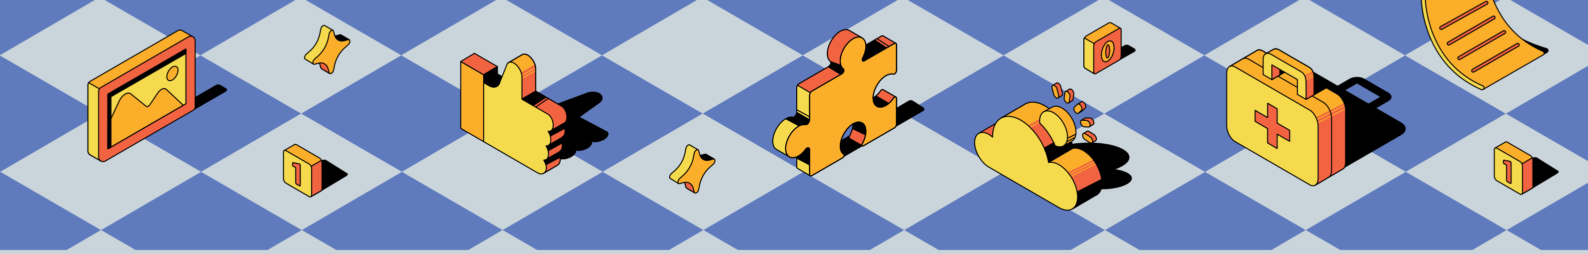

a user watches or likes a movie, the system can recommend other movies with similar characteristics. Content-based filtering does not require lots of user history data, making it suitable for new users. It can also provide more meaningful explanations; for example, you may know that it suggested a particular movie because it has a certain actor in it. However, due to its focus on item similarities, it can struggle to make diverse recommendations. For example, for a user who has just watched *Toy Story*, it may simply recommend *Toy Story 2, 3, and 4* if it deems those to be the most similar items.

Today, many websites use hybrid recommender systems that utilise both collaborative and content-based filtering, to leverage each system's strengths and offset their disadvantages. For example, Netflix utilises both data about its videos (content) and data about its users' viewing histories (collaborative filtering) when making recommendations. The effectiveness of such systems lies in their ability to leverage machine learning, and vast amounts of data, to calculate valuable recommendations. For example, collaborative filtering can determine underlying patterns and hard-to-spot connections between millions of users, and utilise them in making good recommendations. But what do 'good' recommendations look like?

What are good recommendations?

Recommender systems and the AI technologies that power them are built by humans, and therefore invariably reflect, either explicitly or implicitly, their creators' objectives and biases. For big tech websites, the main goal of their recommender systems is to suggest content that will keep a user's attention, and so keep them on the site as long as possible. When companies just prioritise retaining users' immediate attention, there's a risk that the recommendations may not align with what users would truly enjoy. Recently, we have begun to question the notion of what a good recommender system is, and recommender systems are now considering more nuanced ideas about recommendation objectives, such as serendipity, to suggest items users might like but have never considered or heard of before. Such ideas can even be taken further, to provide recommendations that challenge a user's views, or provide them with a chance to confront alternative sides of an argument.

Given concerns that big tech companies may not always have users' best interests at heart, there is a growing consensus that good recommendations are also ones where additional explanation is provided, detailing why items are being suggested. Such explainability is sought both by individuals, who want to understand why certain items are being suggested, and by governments and supreme courts wrestling with regulations and accountability. However, there is some apprehension from companies that greater explainability may reveal too much, facilitating a loss of competitive advantage, and that there is a trade-off between performance and explainability that is not worth the compromise.

> "THERE'S A RISK RECOMMENDATIONS MAY NOT ALIGN WITH WHAT USERS TRULY ENJOY

Influence and power

Recommendation systems are becoming an increasingly inescapable component of our digital lives, with users today just as likely to discover new content via such systems as they are by searching directly for it. As a result, young learners have regular exposure to such systems, which can present a great opportunity to engage them via these direct links to their own lives. As educators, we can foster awareness of recommender systems by getting learners to try and identify when they encounter them. Additionally, multiple learners who use the same site, such as YouTube, can observe the impacts of such systems by showing each other what their home pages look like.

As such systems become more commonplace, they are influencing more and more of the content that we are reading, watching, and listening to. This means they have incredible power to influence which videos go viral, and which do not. The wielding of such influence by big tech companies raises concerns that recommender systems could be used for dishonest purposes. A company could try to influence customer behaviour towards buying more expensive products, or could even aim to influence a whole political campaign. As recommender system algorithms get more and more complex, such issues might even occur without the designers intending for them to happen. It is therefore essential that young learners are aware of recommender systems and are curious regarding their operation, purpose, and intentions. (HW)

EDWARD ABEL
An assistant professor (lecturer) at the University of Southern Denmark, Edward explores research within data science and AI, and has previously been involved in running Code Clubs and CoderDojos in the United Kingdom.