

Carnegie Mellon University Africa
Certificate Program in AI and Machine Learning in Africa

Course AIML01: Artificial Intelligence – Past, Present, and Future

Module 3: Example Applications
Lecture 3: AI Applications in the Web and Social Media

Welcome to Lecture 3 of Module 3, the third of four lectures that look at the applications of artificial intelligence. This lecture focusses on applications for the web and social media.

In this lecture, we discuss the impact of AI on web applications and social media. We highlight the key AI-based advances in search algorithms on the web. We explain impact of deep learning on recommender systems. We discuss the ethical challenges posed by the widespread use of face recognition on the web and in social media. We will finish up by summarizing what we have covered and identifying the articles that you should read to consolidate what you have learned.

We have two learning objectives, so that, after studying the material covered in this lecture, you should be able to do the following.

1. Identify the impact of AI and machine learning on web-based applications and social media.
2. Discuss the technical and ethical challenges posed by the use of some of the underlying challenges.

- Slide 1 Welcome to Lecture 3 of Module 3, the third of four lectures that look at the applications of artificial intelligence. This lecture focusses on applications for the web and social media.
- Slide 2 AI has had, and continues to have, a tremendous impact in a variety of applications and functionalities for the web, such as
- Search algorithms
 - Music and video recommendations
 - Automatic translation
- and for social media, such as
- News selection and recommendation
 - Sentiment analysis
 - Face recognition
- Although this progress is resulting in clear benefits to people and society, it also carries important ethical considerations and risks.
- Slide 3 AI has significantly changed the search algorithms for the web.
- For example, Google's initial PageRank algorithm
- Which was based on standard mathematical methods
- Has now developed in a collection of search tools, such as the Panda, Penguin, Hummingbird, and Pigeon frameworks.

Slide 4 RankBrain was introduced in 2016.

It is an AI system that uses machine learning to infer how pages are related to concepts.

This means it can return relevant pages even if they don't contain the exact words used in a search.

Slide 5 The Google search algorithms also use BERT (Bidirectional Encoder Representations from Transformer),

a type of deep neural network for natural language processing.

BERT uses word context to find more semantically-relevant information,

allowing it to select 'featured snippets', i.e., short snippets of text, figures, or tables that appear at the top of Google's search results and provide definitions of the searched-for item.

Slide 6 Recommender systems provide recommendation for purchases in e-commerce sites,

suggestions of related news and friends in social media,

and personalised recommendation in media streaming sites and apps.

80% of movies watched on Netflix are based on AI recommendations (Zhang et al. 2019).

Slide 7 As in other domains, deep machine learning has become the default algorithm for the latest recommendation systems.

This raises important concerns,

some are technical,

such as the need for explainability,

that is, for black-box deep neural networks to provide an explanation of the reasons for the recommendation

- Slide 8 Others are ethical, concerned with the influence of these algorithms
- in politics, for example, introducing bias into recommendations and social media content during elections
- in public health, for example, undermining scientifically-grounded health advice
- and in the generation and diffusion of fake information, generally.
- Slide 9 AI techniques for face recognition have become widespread on the web and on social media.
- These algorithms can be used for image matching and people recognition
- For example, in social media photo tagging
- as well as for authentication
- for example, to implement secure access in some smartphone systems.
- Slide 10 A variety of AI and machine learning algorithms have been developed to implement this functionality
- For example, probabilistic Bayesian networks and support vector machines
- with the recent design of a variety of deep learning face recognition systems typically based on convolutional neural networks (CNNs) and autoencoders.
- Slide 11 The pipeline of a typical automated face recognition system comprises
- The input of the test face
- A preprocessing phase to detect the region of the image containing the face and realign it so that the axis of symmetry of the face is vertical
- A recognition phase to extract features and match them against a database of faces to produce the resultant identity or class
- And an evaluation phase to determine the accuracy of the result.

Slide 12 However, face recognition algorithms based on learning from datasets have important ethical implications regarding possible biases in the data used for the training

Leading facial recognition systems produced substantial disparities in the accuracy of gender classification (Buolamwini and Gebru, 2018)

With error rates of up to 34.7% in the classification of darker-skinned females compared with the maximum error rate of 0.8% for lighter-skinned males

This highlights the urgent need to address and remedy implicit bias in such systems and make sure they are based on fair, transparent, and accountable facial analysis algorithms.

To summarize:

1. AI and machine learning algorithms appear in many applications on the web and in social media, bringing great benefits but also technical and ethical challenges to ensure their operation is fair, transparent, and free of bias.
2. Search algorithms have improved greatly as a result of the use of AI and machine learning, allowing the search algorithms to infer the intent of the search, beyond what is stated explicitly in the query by using context and inferring concepts.
3. Recommender systems based on deep learning are very widely used but this also raises the possibility of their misuse and the need for these systems to be able to explain the basis of the recommendations.
4. Face recognition functionality is also widely used and, again, this raises the need for ethical deployment, ensuring the absence of bias, fairness, and transparency.

Recommended Reading

Here is some recommended reading. Both are comprehensive surveys. You don't have to understand all the details to benefit from reading them and you will see how much of the material we have already covered in the course is mentioned.

Guo, G., & Zhang, N. (2019). A survey on deep learning based face recognition. *Computer Vision and Image Understanding*, 189, 102805.
<https://dl.acm.org/doi/abs/10.1016/j.cviu.2019.102805>

Zhang, S., Yao, L., Sun, A., & Tay, Y. (2019). Deep learning based recommender system: A survey and new perspectives. *ACM Computing Surveys (CSUR)*, 52(1), 1-38.
<https://dl.acm.org/doi/pdf/10.1145/3285029>

References

Here are the references cited to support the main points in what we covered today.

Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Conference on fairness, accountability and transparency* (pp. 77-91). *Proceedings of Machine Learning Research* 81:1–15.
<https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>