Carnegie Mellon University Africa
Certificate I: Understanding AI and Machine Learning in Africa

Course AIML02: AI and Machine Learning in Africa

Module 02: Application Case Studies
Lecture 04: E-Commerce

Welcome to Lecture 4 of Module 2.

In this lecture, we will focus on a case study in e-commerce, specifically the issue of verifying the identity of customers across Africa.

We'll start off by providing some background on the company, Smile Identity, that developed the technology and uses it in their products.   We'll then walk through the technical details in the target article.

We introduce the concept of know your customer (KYC), sometimes called know your client, and discuss its importance for e-commerce.

We introduce Smile Identity, a company that specializes in authenticating digital identity for Africa.

We then walk through the technical details of the face recognition technique behind the Smile Identity products.

We finish up by summarizing what we have covered and identifying the articles that you should read to consolidate what you have learned.

As before, we will encounter AI and machine learning techniques which we introduced in the first course, AIML01, and others that we introduced in this course, but which you have not yet studied in detail. As we do so, we will flag where they were introduced and provide a little more detail if appropriate. However, we treat them mainly as a preview of material to follow. As such, we will identify the courses in this and other certificates where these techniques are covered in greater depth.

Once you have listened to this lecture and read the commentary, you should then read the target article, look at a Smile Identity video, and read their 2022 State of KYC in Africa. You should then listen to the lecture again and re-read the commentary.

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

1. Explain the concept of Know Your Customer (KYC).

2. Identify the problems for identity verification in Africa.

3. Explain the technical details and performance of machine learning system for African face authentication.

Slide 1 Welcome to Lecture 4 of Module 2.

In this lecture, we will focus on a case study in e-commerce, specifically the issue of verifying the identity of customers across Africa.

We'll start off by providing some background on the company, Smile Identity, that developed the technology and uses it in their products. We'll then walk through the technical details in the target article.

Slide 2 The "Know Your Customer" (KYC) process is a requirement that helps businesses identify their users and verify their credentials.

Slide 3 However, according to a report by the World Bank, 1 billion people do not have official proof of identity and 1 in 2 women in low-income countries does not have an ID, limiting their access to critical services.

Slide 4 The solution is digital identity:

The ability to prove you are who you say you are online

Slide 5 Two reasons for conducting KYC checks

1. Regulatory compliance.
2. Fraud prevention.

Conventional approaches are time-consuming and expensive.

Slide 6 Biometric KYC can assist by confirming that the person providing a credential (an ID number) is who they say they are.

Slide 7 Biometric systems that use deduplication can also be used to ensure a person is only who they say they are, that is, that they haven't created duplicate identities.

Slide 8 Most face recognition systems are built in the United States and Europe and are unable to verify darker skin tones.

Enter Smile Identity, an identity verification startup.

Their SmartID face recognition system makes it easier for Africans to have their identity verified, helping to eliminate bias against them.

Slide 9 To date, they have received over $7 million in investment and have established a presence in Nigeria, Kenya, South Africa, Ghana, Rwanda, and Uganda, and their software is used in banking, telecommunications, fintech, ride-sharing, worker verification, and public social welfare programs.

Here is the story of the technology behind their identity verification capability.

Slide 10 The target paper by Davy Uwizera, William Bares, and Catalin Voss describes a face recognition systems called SmileID.

While this is a short paper, it has quite a lot of technical depth. As we noted in previous case studies, we will encounter AI and machine learning techniques that you have not yet studied in detail, although we may have introduced them briefly in the first course, AIML01, and in some lectures in this course, AIML02.

As we said at the start, when we meet these techniques here, we will revisit their introduction in AIML01, providing a little more detail if appropriate, and we will identify the courses in this and other certificates where these techniques are covered in more detail.

SmileID is a commercial system for frontal-face identity verification on mobile handsets in Africa for a variety of applications such as banking, lending, and ride-sharing.

The use on mobile handsets and its focus on use in Africa are what differentiate SmileID from other systems.

Slice 11 Deep Convolutional Neural Networks (CNNs) provide human-level performance on many datasets.

However, they have mainly been trained and tested with data featuring European & American, and predominantly Caucasian faces.

Also, they have been optimized to perform well in lighting conditions and poses that match the applications of interest to Western customers.

While commercial face recognition providers such as AWS Rekognition, Microsoft, and Google are working to eliminate bias, their approach, their goal remains focused on robust general-purpose face recognition.

What is needed is an Africa-targeted model that works well on mobile handsets for identity verification.

Slide 12 The SmileID system is based on ResNet50, a deep convolutional neural network, pre-trained on the MS-Celeb1M dataset,
and then adjusted for African faces by transfer learning with a proprietary Smart Identity dataset.

We came across ResNet in micro course AIML01, Module 2, Lecture 2 on connectionist approaches to AI, and in Module 3, Lecture 1 on the healthcare case study in this course.

We mentioned transfer learning in the micro course AIML01, Module 3, Lecture 1: AI Applications in Medicine and again in Module 3, Lecture 1 on the healthcare case study in this course.

To recap, transfer learning refers to the practice of first training a model, typically a deep neural network model, using a large general-purpose dataset and then tuning the trained model using a smaller application-specific dataset.

We will cover convolutional neural networks and transfer learning in detail in AIML10 Introduction to Deep Learning.

Our goal here is to outline this approach, previewing the technical details that will be covered in subsequent courses, as a way of emphasizing its relevance to AI and machine learning in Africa. We won't go into the detailed operation of ResNet, leaving that for micro course AIML10.

Slide 13 So, for future reference, let's highlight a few important aspects of the ResNet architecture and training process.

The goal is to learn an internal representation that simultaneously achieves intra-class compactness and interclass separability. We'll see how well it manages that in a moment.

Intra-class compactness means that the features that characterize a given class are grouped locally in the feature space.

Interclass separability means that the distributions of each learned classes are distinct and far from each other.

Both contribute to the robustness and reliability of the classification, and the accuracy of the identity verification.

The learned representation is called an embedding model.

This is achieved in part by using the ArcFace loss function.

A loss function quantifies how well predicted class labels agree with ground-truth labels.

A low loss implies high level of agreement (and vice versa).

The goal of training is to minimize the loss function and, hence, increase classification accuracy.

The details of the loss function need not concern us here.

Nevertheless, it would be worthwhile revisiting this case study, and the approach taken to training ResNet, once you have completed AIML10 Introduction to Deep Learning and when you are familiar with deep neural networks, the ResNet architecture, and the training process.

Slide 14 The ResNet architecture is trained on MS-Celeb1M dataset, a popular large open dataset for face recognition consisting of predominantly white faces.

During training and verification, image values are normalized between zero and one, a face is located, cropped and aligned to a 112x112 bounding box.

Slide 15 ResNet is then post-trained using a proprietary Smile Identity dataset. This is the transfer learning phase.

This is referred to as the FDD dataset in the target article.

Images were collected from individuals using the SmileID mobile software development kit, built into third-party apps built by Smile Identity's customers in Africa.

The dataset comprises 22,330 images: 70% are used for training and 30% for evaluation purposes.

During transfer learning, all pre-trained layers except for the last 4 ResNet layers are fixed.

This helps the model learn the elements of the non-Caucasian recognition task, while keeping the low-level features learnt during the pre-training phase.

Slide 16 The difference between the Celeb-1M dataset and the Smart Identity FDD dataset can be seen in these two images, both of which were constructed by blending 500 aligned faces from each dataset.

Slide 17 Data augmentation is applied during both the pre-training and transfer-learning processes.

Random saturation and brightness is applied to training images with 50% probability.

Vertical flipping introduced with 40% probability.

Noise is added with 10% probability.

Slide 18 SmileID is evaluated with three datasets and compared to two baseline models.

The three datasets are

The Labelled Faces in the Wild dataset: LFW.

A proprietary FDD dataset, i.e., the 30% evaluation portion of the total 22,330 images in the Smile Identity dataset.

A dataset of random pairs of matching dark faces from LFW dataset: LFWB.

The two baseline models are

The AWS (Amazon Web Services) Rekognition system.

ResNet trained with ArcFace but without transfer learning.

Slide 19 The transfer learning model achieves a 11% gain over the baseline ArcFace implementation on African faces.

Slide 20 Here is the ROC curve for the baseline ArcFace model (labelled Published Arcface in the graph)

and the ROC curve for the optimized ArcFace model after transfer learning with the proprietary FDD data set.

ROC is short for receiver operating characteristic.

It shows the balance of true positive rate versus false positive rate or a binary classifier as the decision threshold value is varied.

Slide 21 A perfect classifier would have a true positive rate of 1 and a false positive rate of zero, no matter what decision threshold value was used.

Slide 22 The true-positive rate is also known as sensitivity or recall.

The false-positive rate is also known as probability of false alarms and can be calculated as (1 − specificity).

We met the two terms sensitivity and specificity already in Lecture 1 of this module when we discussed the classification of babies with asphyxia using the Ubenwa application.

Sensitivity (the true positive rate) refers to the probability of a positive outcome, given that the sample is actually positive.

Specificity (the true negative rate) refers to the probability of a negative outcome, given that the sample is actually positive.

Slide 23 Here is another way of visualizing the tradeoff between true positives and false positives

The balance between the two will vary for these two distributions as the decision threshold value is varied.

Slide 24 Returning back to the results in the target article, recall that one of the goals of the approach is to learn an internal representation (also known as an embedding) that simultaneously achieves

intra-class compactness

and

interclass separability.

The two figures show a visualization of the embedding clusters for African faces.

The top one shows the baseline ArcFace model.

The bottom one shows the optimized ArcFace model after transfer learning with the proprietary FDD data set.

You can see that the transfer learning models achieves better compactness and separability.

To summarize:

1. The know your customer process is essential for avoiding fraud, both in financial transactions and when onboarding customers and clients.

2. Digital identity verification is key step in this, but most systems do not perform well in the African context.

3. Smile Identity have developed an effective face recognition technique that outperform other approaches.

4. It uses a pretrained ResNet50 deep convolutional network and transfer learning using a proprietary Smile Identity training set.

Recommended Reading
Here is the target article used for the case study. Please read it carefully.

Uwizera, D., Bares, W., Voss, C. (2020). Data Centric Face Recognition for African Face Authentication, Smile Identity. https://cdn.smileidentity.com/Smile_Identity_Model_Paper-CVML.pdf

You should also read this report from Smile Identity on the state of know your customer in Africa.

Keirstead, M, Straub, M., Orina, L, Wambua, R., Scheybani, N., and Williams, G. (2022). State of KYC in Africa, H1 2022 Report. https://smileidentity.com/img/state-of-kyc-report_2022h1.pdf

H1 2022 REPORT

State of KYC
in Africa

SMILE IDENTITY

To summarize:

5. The know your customer process is essential for avoiding fraud, both in financial transactions and when onboarding customers and clients.

6. Digital identity verification is key step in this, but most systems do not perform well in the African context.

7. Smile Identity have developed an effective face recognition technique that outperform other approaches.

8. It uses a pretrained ResNet50 deep convolutional network and transfer learning using a proprietary Smile Identity training set.

Recommended Reading

Here is the target article used for the case study. Please read it carefully.

Uwizera, D., Bares, W., Voss, C. (2020). Data Centric Face Recognition for African Face Authentication, Smile Identity. https://cdn.smileidentity.com/Smile_Identity_Model_Paper-CVML.pdf

You should also read this report from Smile Identity on the state of know your customer in Africa.

Keirstead, M, Straub, M., Orina, L, Wambua, R., Scheybani, N., and Williams, G. (2022). State of KYC in Africa, H1 2022 Report. https://smileidentity.com/img/state-of-kyc-report_2022h1.pdf

H1 2022 REPORT

State of KYC in Africa

SMILE IDENTITY

This short video introduction to Smile Identity provides some helpful context for the SmileID application.

https://www.youtube.com/watch?v=g1vHLH4gWyo

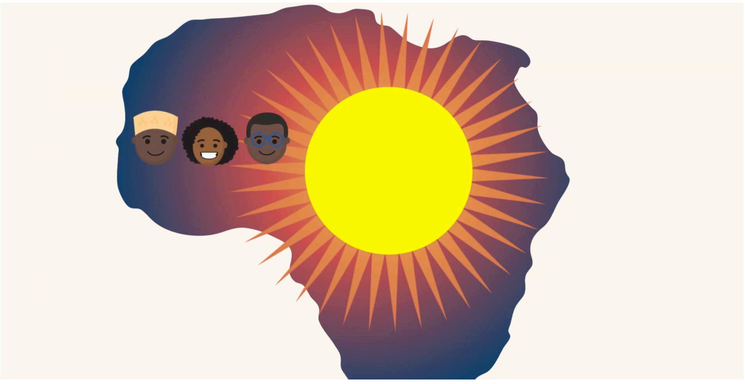

References

Here are five of the references mentioned in the course of this lecture.

D4D Data: Global Identification Challenge by the Numbers, World Bank. hthttps://id4d.worldbank.org/global-dataset

Deng, J., Guo, J., Xue, N., and Zafeiriou, S. (2019). Arcface: Additive angular margin loss for deep face recognition, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4690–4699. https://openaccess.thecvf.com/content_CVPR_2019/papers/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.pdf

Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J. (2016). MS-Celeb-1M: A Dataset and Benchmark for Large-Scale Face Recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds) Computer Vision – ECCV 2016. ECCV 2016. Lecture Notes in Computer Science, vol 9907. Springer. https://doi.org/10.1007/978-3-319-46487-9_6

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770-778. https://arxiv.org/pdf/1512.03385.pdf

Huang, G. B., Ramesh, M., Berg, T., and Learned-Miller, E. (2008). Labeled faces in the wild: A database for studying face recognition in unconstrained environments, in: Workshop on Faces in "Real-Life" Images: Detection, Alignment, and recognition. Marseille, France. http://vis-www.cs.umass.edu/lfw/lfw.pdf

This short video introduction to Smile Identity provides some helpful context for the SmileID application.

https://www.youtube.com/watch?v=g1vHLH4gWyo

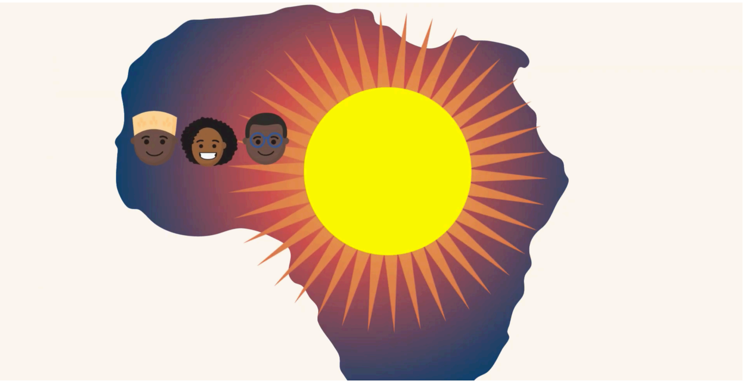

Here are five of the references mentioned in the course of this lecture.

D4D Data: Global Identification Challenge by the Numbers, World Bank. hthttps://id4d.worldbank.org/global-dataset

Deng, J., Guo, J., Xue, N., and Zafeiriou, S. (2019). Arcface: Additive angular margin loss for deep face recognition, in Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 4690–4699. https://openaccess.thecvf.com/content_CVPR_2019/papers/Deng_ArcFace_Additive_Angular_Margin_Loss_for_Deep_Face_Recognition_CVPR_2019_paper.pdf

Guo, Y., Zhang, L., Hu, Y., He, X., Gao, J. (2016). MS-Celeb-1M: A Dataset and Benchmark for Large-Scale Face Recognition. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds) Computer Vision – ECCV 2016. ECCV 2016. Lecture Notes in Computer Science, vol 9907. Springer. https://doi.org/10.1007/978-3-319-46487-9_6

He, K., Zhang, X., Ren, S., and Sun, J. (2016). "Deep residual learning for image recognition." In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770-778. https://arxiv.org/pdf/1512.03385.pdf

Huang, G. B., Ramesh, M., Berg, T., and Learned-Miller, E. (2008). Labeled faces in the wild: A database for studying face recognition in unconstrained environments, in: Workshop on Faces in "Real-Life" Images: Detection, Alignment, and recognition. Marseille, France. http://vis-www.cs.umass.edu/lfw/lfw.pdf