

Detecting Floaters in Beer

Applied ML 2025

Carl-Johannes Johnsen

Duplicati - carl@duplicati.com

Department of Computer Science - carl-johannes@di.ku.dk
University of Copenhagen



1 Detecting Floaters in Beer

Introduction

The old solution - Pia

The new solution - Pia 2.0

2 Handling the Data

The amount of data

The raw video

Processing the video

Reducing the video

3 Applying Machine Learning

4 Problems with the results

5 Conclusion / Future work



This is a project for Carlsberg.

Sometimes in beer, proteins can fold together into particles or strings, called floaters. These aren't an health issue, but are unappatizing to look at.



Usually this occurs in damp or hot countries, either during transportation or during shelf time.

Carlsberg wants to detect these floaters.
Their current implementation is Pia. Pia will take a beer, hold it up against a light and give it 6 scores between 0 and 5.

Carlsberg wants to replace this functionality of Pia, as it is troublesome and "contaminates" the samples when they move the beer to Denmark.


Furthermore, it is troublesome to move Pia to the sample sites, as this would introduce some latency, which again might "contaminate" the samples.



The new Pia 2.0 implementation is a camera and roller setup.

The target is to make a program, which will look at the video and produce the same, and hopefully more consistent, scoring as Pia did.

This will be a useful tool for Carlsberg, as they can put up machines in multiple locations, to further investigate when and why the floaters occurs.



1 Detecting Floaters in Beer

Introduction

The old solution - Pia

The new solution - Pia 2.0

2 Handling the Data

The amount of data

The raw video

Processing the video

Reducing the video

3 Applying Machine Learning

4 Problems with the results

5 Conclusion / Future work



Carlsberg collected the following data:



Carlsberg collected the following data:

- 31 different beers



Carlsberg collected the following data:

- 31 different beers
- Each beer has been sampled at 4 different days (0, 3, 7, 10)



Carlsberg collected the following data:

- 31 different beers
- Each beer has been sampled at 4 different days (0, 3, 7, 10)
- For each day, 8 videos of each beer were captured



Carlsberg collected the following data:

- 31 different beers
- Each beer has been sampled at 4 different days (0, 3, 7, 10)
- For each day, 8 videos of each beer were captured
- Each video is made of 125 frames



Carlsberg collected the following data:

- 31 different beers
- Each beer has been sampled at 4 different days (0, 3, 7, 10)
- For each day, 8 videos of each beer were captured
- Each video is made of 125 frames
- Each frame is 4012×2048 in grayscale, or ~ 8 MB



Carlsberg collected the following data:

- 31 different beers
- Each beer has been sampled at 4 different days (0, 3, 7, 10)
- For each day, 8 videos of each beer were captured
- Each video is made of 125 frames
- Each frame is 4012×2048 in grayscale, or ~ 8 MB
- In total: 992 videos or ~ 968 GB



We start by doing the first step of any machine learning problem:
we look at the data! (queue live demo)



It is very hard to see anything in this video. While an algorithm might be able to see something, we should still extract all of the relevant information before training.

We do this with common image processing:



It is very hard to see anything in this video. While an algorithm might be able to see something, we should still extract all of the relevant information before training.

We do this with common image processing:

- Remove the bottle (remove the background)



It is very hard to see anything in this video. While an algorithm might be able to see something, we should still extract all of the relevant information before training.

We do this with common image processing:

- Remove the bottle (remove the background)
- Remove noise (blur and erosion)




It is very hard to see anything in this video. While an algorithm might be able to see something, we should still extract all of the relevant information before training.

We do this with common image processing:

- Remove the bottle (remove the background)
- Remove noise (blur and erosion)
- Enhance the particles (dilation)



Now we have a much cleaner video, where the particles are very easy to see. (queue live demo part 2)



The final step, is to ready the data for machine learning.



The final step, is to ready the data for machine learning.
Following the preprocessing, we have effectively masked the floaters in the video, which we can extract from the images.



The final step, is to ready the data for machine learning.
Following the preprocessing, we have effectively masked the floaters in the video, which we can extract from the images.

This gives us a variable list of particle sizes for each frame in the video.

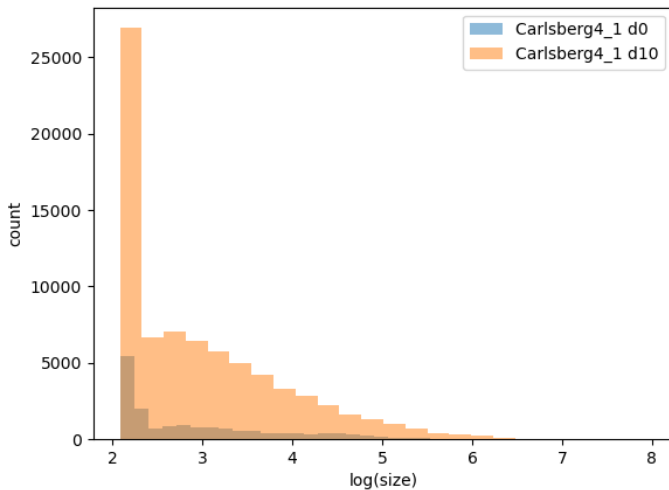


The final step, is to ready the data for machine learning. Following the preprocessing, we have effectively masked the floaters in the video, which we can extract from the images.


This gives us a variable list of particle sizes for each frame in the video.

We collect these lists into histograms, so the data has a fixed size and is ready for training.





To summarize: We convert each video into a histogram, which is a reduced representation of what we are looking for. (queue video)



1 Detecting Floaters in Beer

Introduction

The old solution - Pia

The new solution - Pia 2.0

2 Handling the Data

The amount of data

The raw video

Processing the video

Reducing the video

3 Applying Machine Learning

4 Problems with the results

5 Conclusion / Future work



Now the data is ready for training.

As mentioned, their original classification is 6 values between 0-5. As such, our target is producing multiple regression values.


We chose XGBoost, as it should be fitting for finding the threshold values in the histograms.



Here the mean absolute error is ~ 0.5 .

These results aren't great, but acceptable.

We also tried Tensorflow, as we had prior experience with it, but it did not produce any improvement.



1 Detecting Floaters in Beer

Introduction

The old solution - Pia

The new solution - Pia 2.0

2 Handling the Data

The amount of data

The raw video

Processing the video

Reducing the video

3 Applying Machine Learning

4 Problems with the results

5 Conclusion / Future work



There are multiple problems here



There are multiple problems here

- Bottle bias - "Pias bias"



There are multiple problems here

- Bottle bias - "Pias bias"
- Error of humans - "The error of Pia"



There are multiple problems here

- Bottle bias - "Pias bias"
- Error of humans - "The error of Pia"
- Small data set - Only 992 videos and 31 different beers



1 Detecting Floaters in Beer

Introduction

The old solution - Pia

The new solution - Pia 2.0

2 Handling the Data

The amount of data

The raw video

Processing the video

Reducing the video

3 Applying Machine Learning

4 Problems with the results

5 Conclusion / Future work



These results, and the problems with the solution, were presented for Carlsberg.



These results, and the problems with the solution, were presented for Carlsberg.

They agreed and were happy with the results. They would try to reclassify based on the videos.



These results, and the problems with the solution, were presented for Carlsberg.

They agreed and were happy with the results. They would try to reclassify based on the videos.

They primarily wanted a quantification of the state of the beer, so a small histogram would be sufficient.



These results, and the problems with the solution, were presented for Carlsberg.

They agreed and were happy with the results. They would try to reclassify based on the videos.

They primarily wanted a quantification of the state of the beer, so a small histogram would be sufficient.

The take home message from this project: while machine learning was not used as the final tool, it proved to be useful to learn about the data using the same techniques.



Other possible transformations would be to look at the shape of the particles, rather than just their size. This will put a heavier load on processing, but we have plenty of processing time.

An unsupervised algorithm could be applied, to counter the human classification error.

