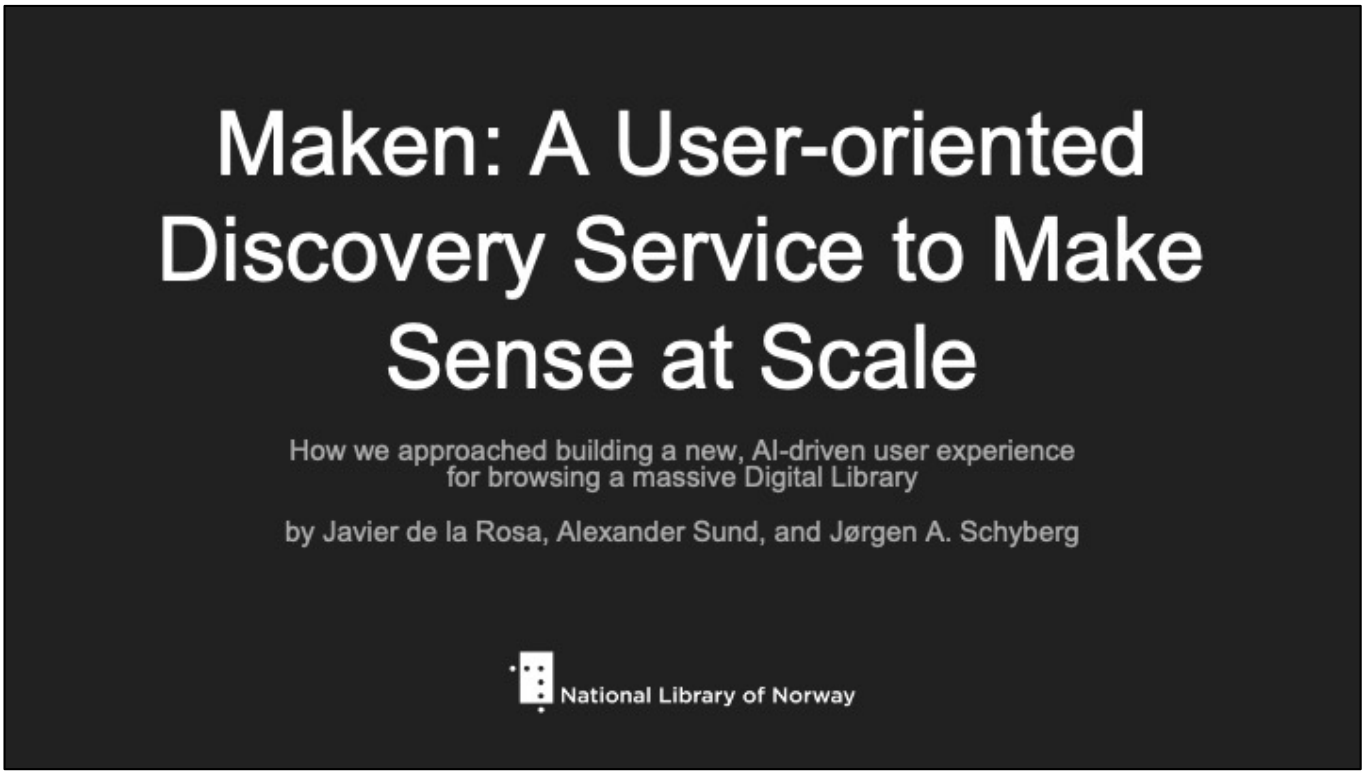

Authors: Javier de la Rosa, Alexander Sund, and Jørgen A. Schyberg, the National Library of Norway
Presenter: Jørgen Schyberg (jorgen.schyberg@nb.no), Product Owner for Maken and the Digital Library

INTRO
My name is Jørgen, and I am probably the person in the room with the least insight and experience with atificial intelligence and machine learning. So this presentation will be quite high-level. And like others who have little or no insight, I have lots of things to say, so let's get started.

Together with some colleagues at the National Library of Norway and a digital agency called Dekode, we have made a web app called Maken, which uses AI to suggest books or images with the help of a similarity algorithm.

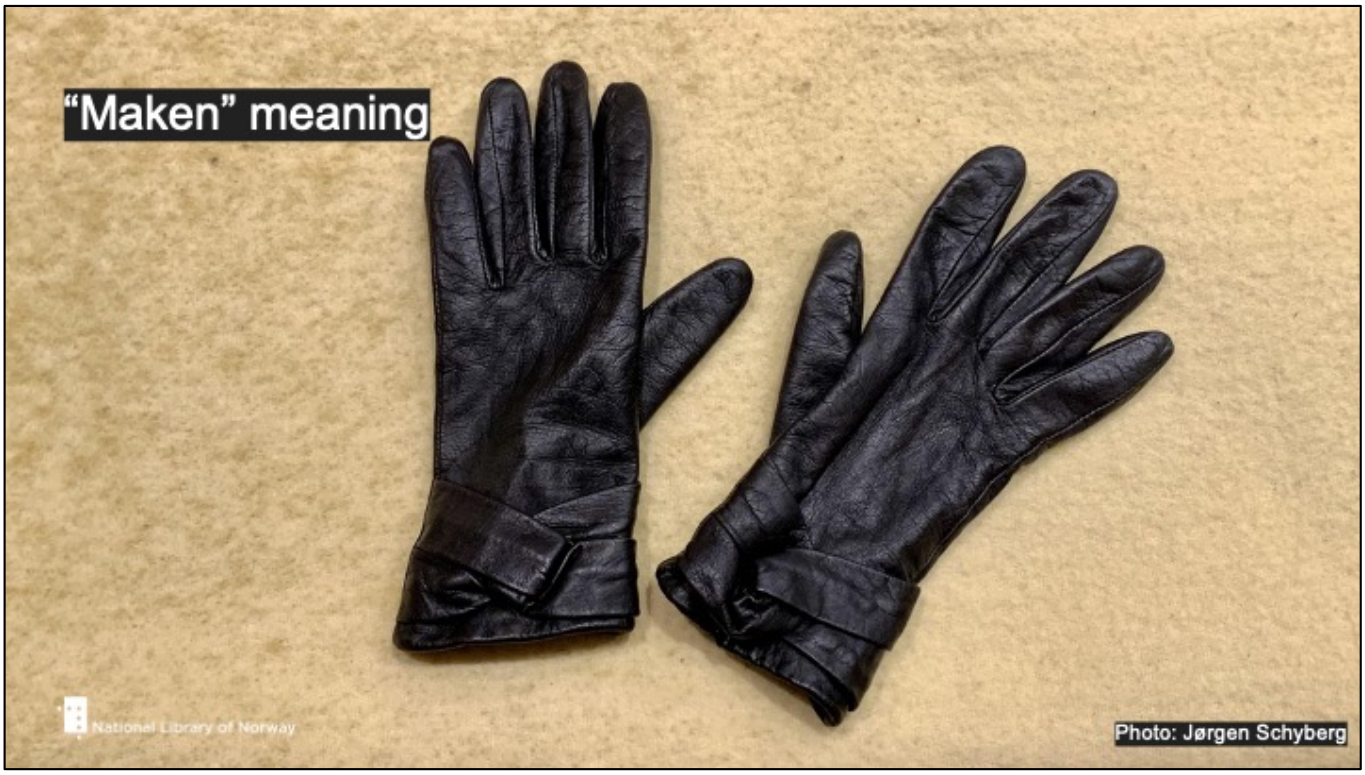

Maken is a somewhat ambiguous Norwegian word that can mean …

”The Matching Other”, like a glove to another glove
“Similar” or even “Identical”, like two drops of water
And sometimes, it can also be used for “the Spouse” or “the Partner”

For us, Maken is a discovery engine to find similar books and similar images.
Based on the pixels of the images, Maken finds related images.
Based on the text of a book itself, Maken finds books that have some resemblance to each other.

In 2020, this was our “marching order” from the head of the Library.

In creating Maken we wanted to utilise some of what our machine learning department, the AI Lab, has learned over several years of experiments, prototyping and tool-building. Returning participants to Fantastic Futures may remember our AI librarian “Nancy” as one of these experiments, named after the famous public librarian Nancy Pearl.

The AI Lab has previously worked with many digital material types, and for Maken we decided to focus on books and images.

We possess a huge and growing digital collection of media in the Norwegian languages.

As such collections grow, discoverability becomes harder. Metadata may be missing or insufficient, and it may be difficult to search within the actual content.

We wanted to see if AI could help make this experience more useful, interesting, or fun. So we started building Maken.

But what does Maken do?

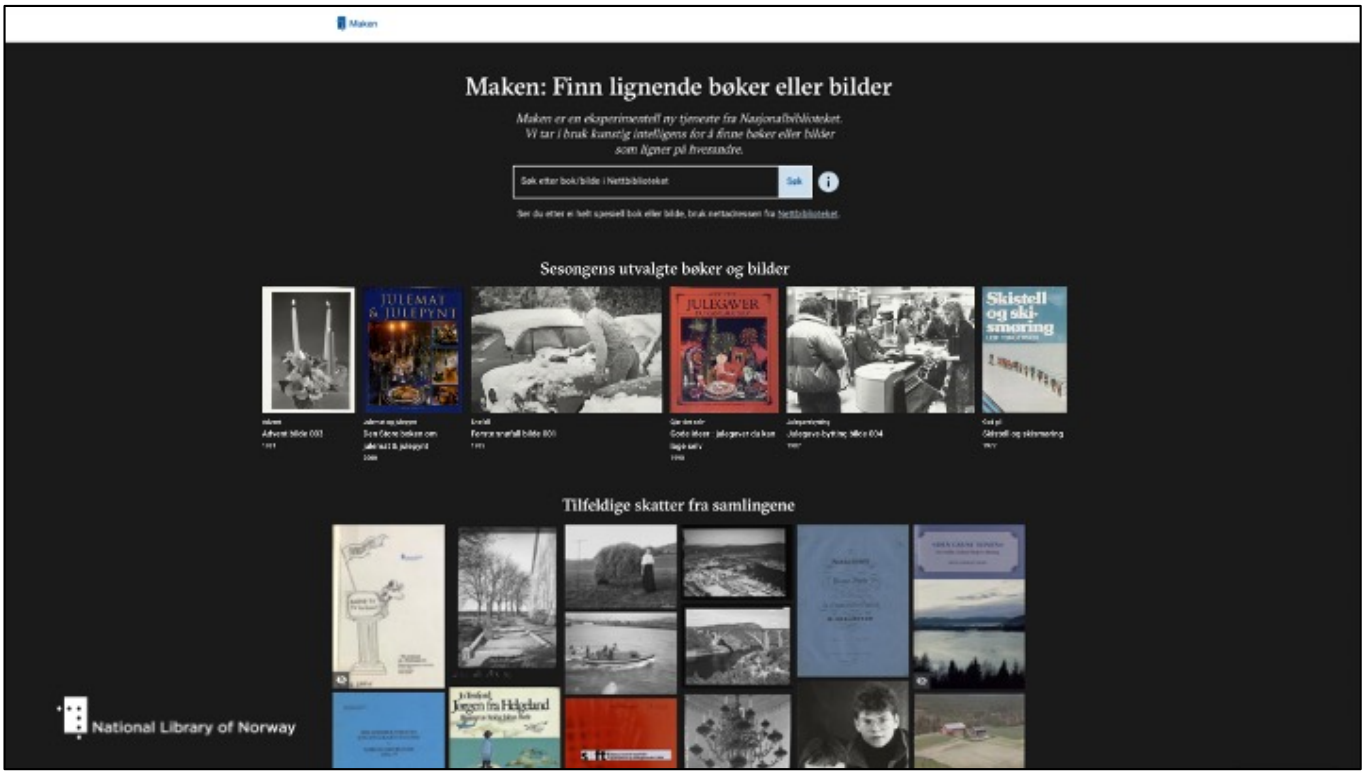

The Maken front page – **https://www.nb.no/maken/** – has:

- a search field that works like the regular search in our Digital Library
- a small set of handpicked objects for the season
- a set of randomly picked objects from the collection

And let's say we pick one of the random objects that you see here, the haystack with a person in front ...

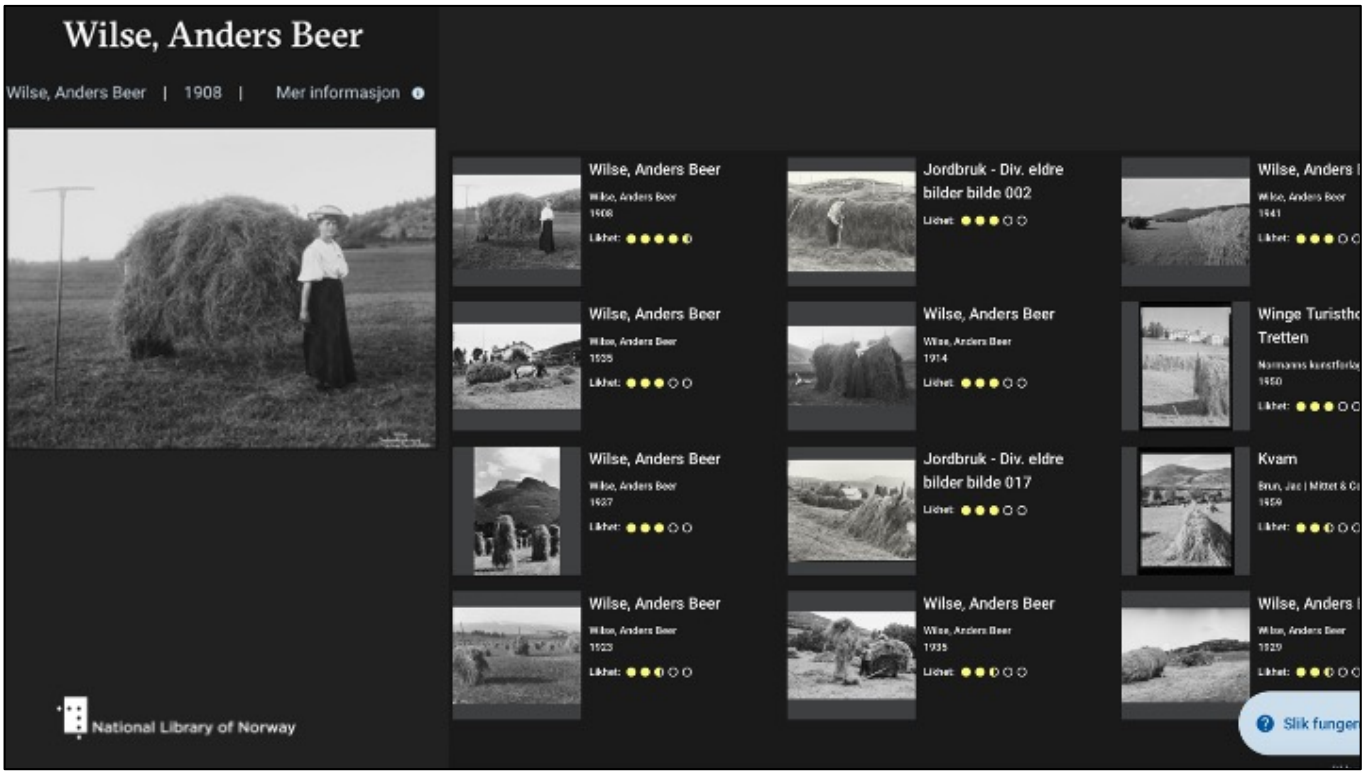

We’re not really looking for a needle in a haystack today, but when the starting point is an image with a haystack, Maken has found other haystacks for us.

https://www.nb.no/maken/item/URN:NBN:no-nb_foto_NF_W_09290

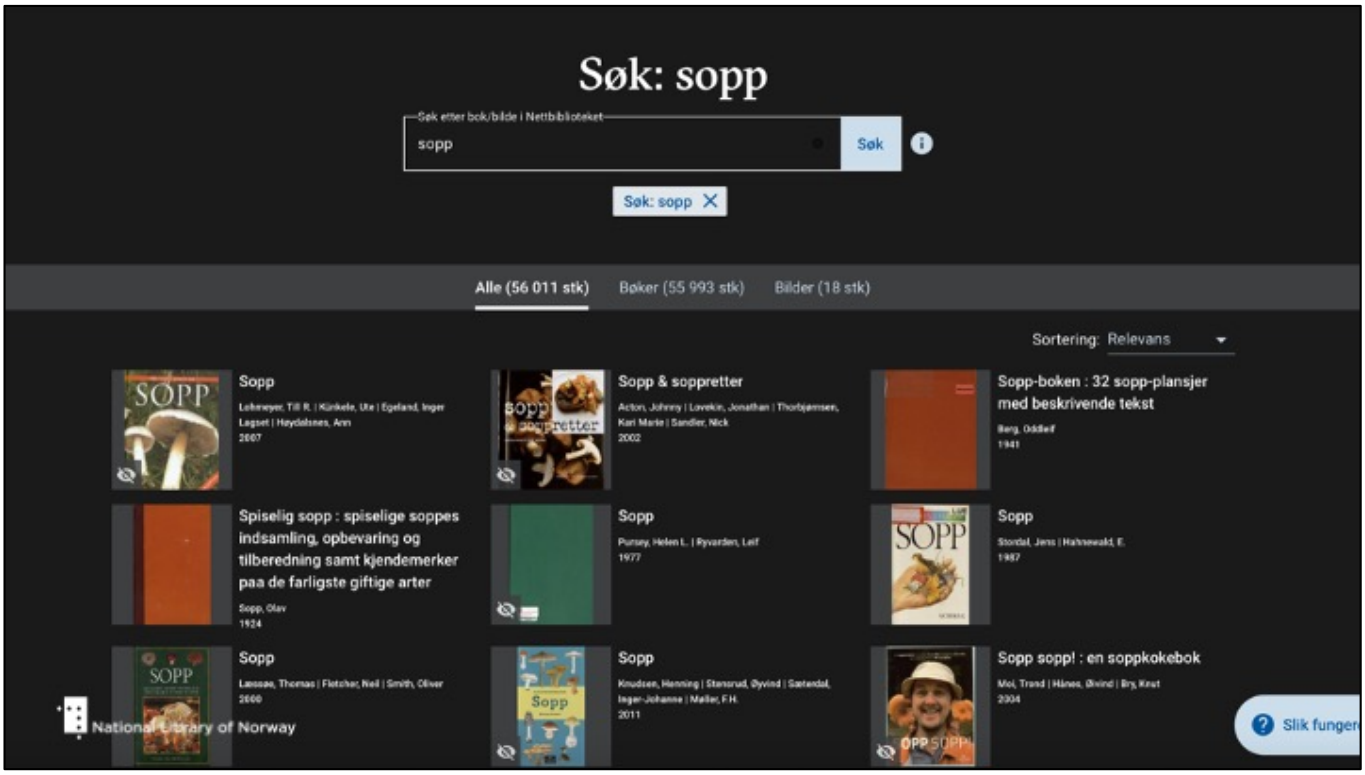

When it comes to books, let’s say I have searched for the term “mushroom” (which is «sopp» in Norwegian).

https://www.nb.no/maken/items/?q=sopp&mediatype=b%C3%B8ker

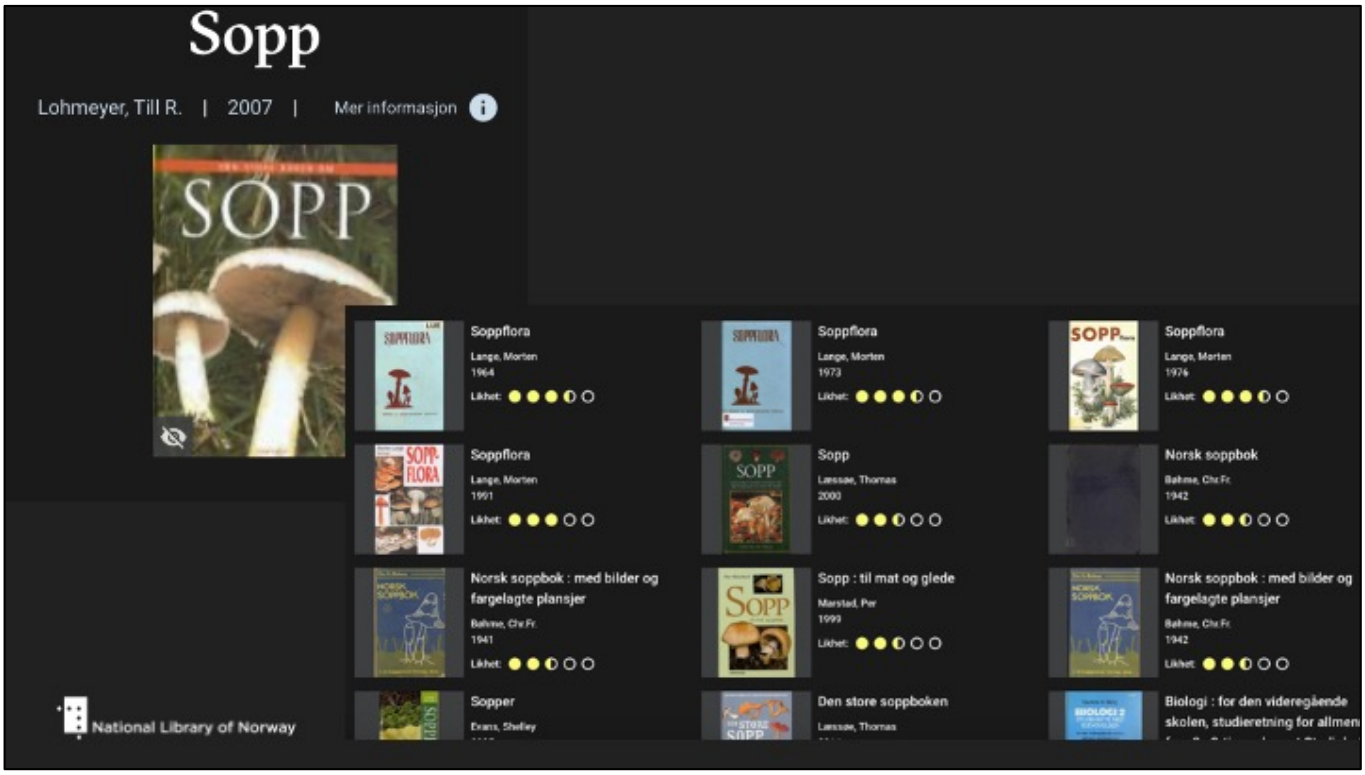

If I pick a book about the biology and classification of mushrooms, Maken returns other books for people foraging for mushrooms.

https://www.nb.no/maken/item/URN:NBN:no-nb_digibok_2013040338017

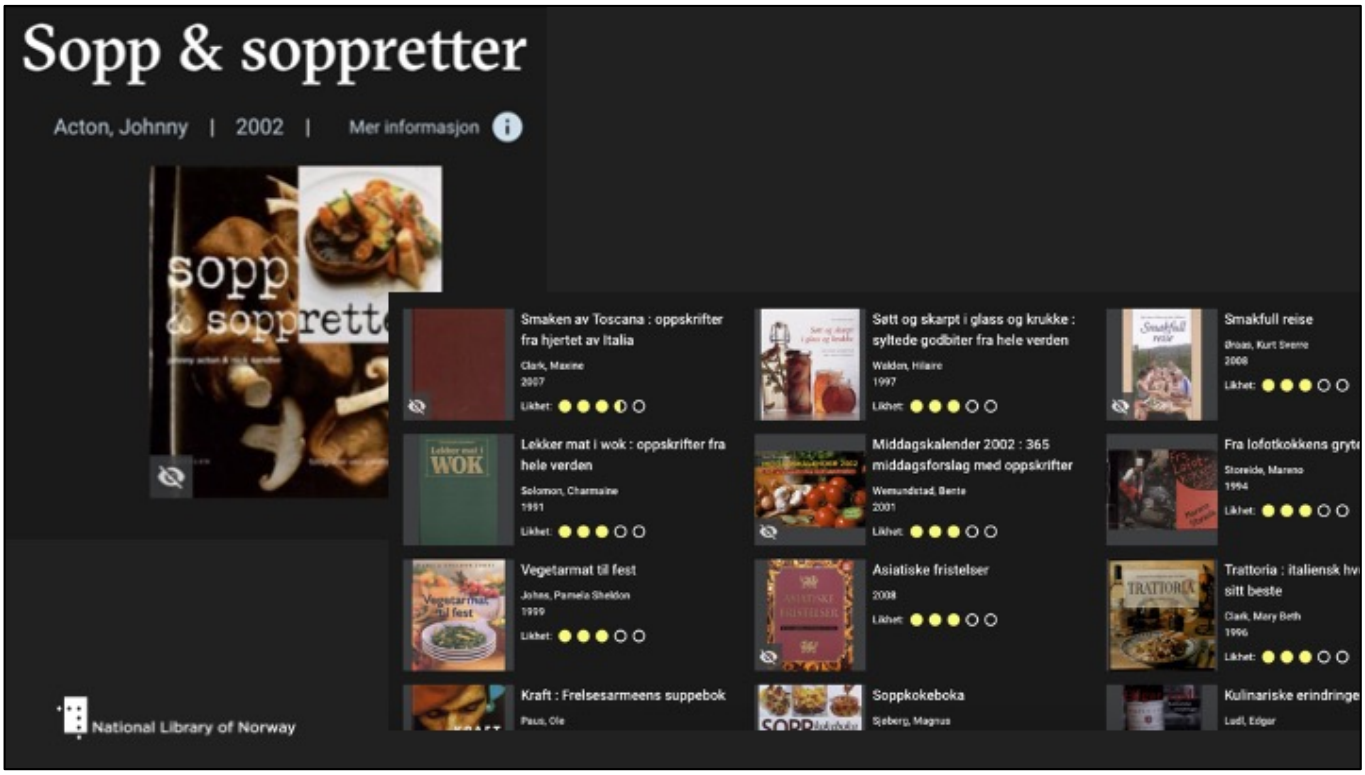

But if I pick a recipe book about mushroom dishes, Maken returns other recipe books.

https://www.nb.no/maken/item/URN:NBN:no-nb_digibok_2017042407084

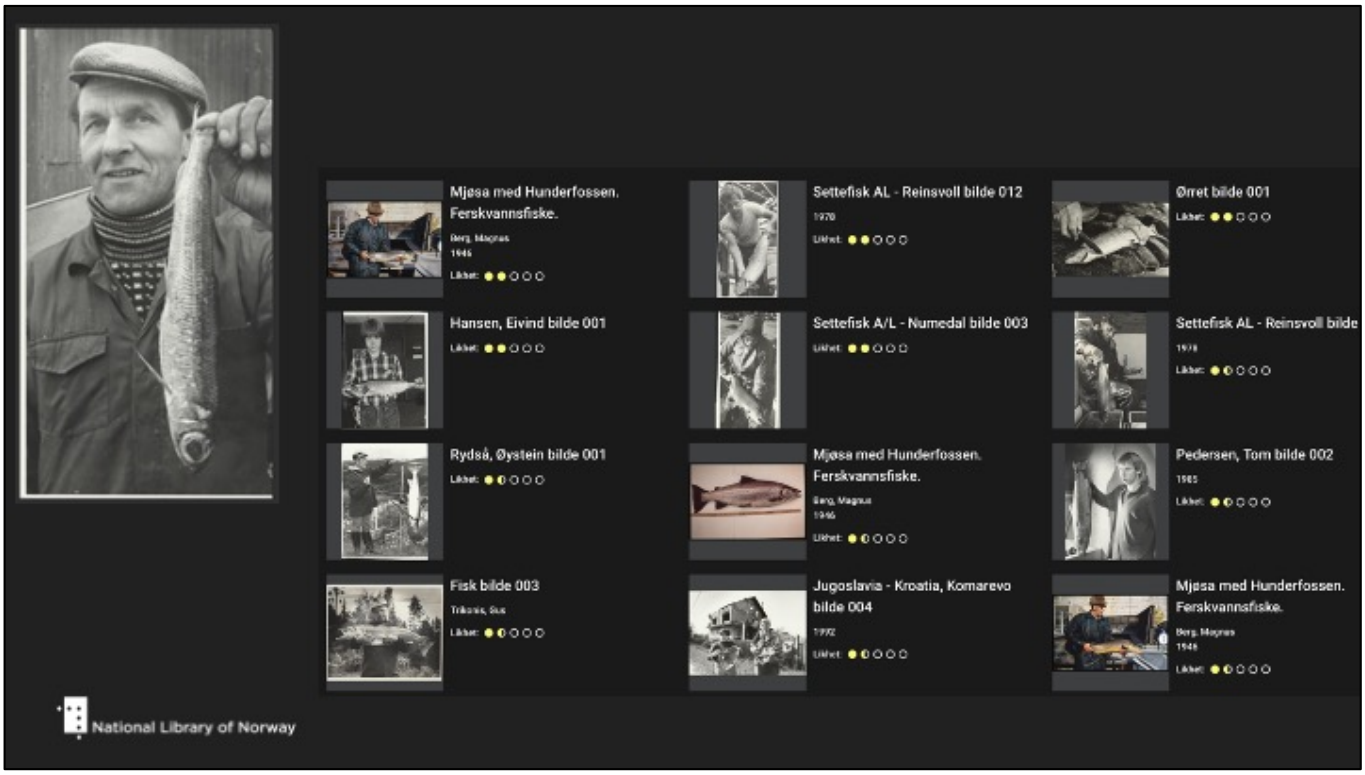

And it is always fun to watch photographs of proud men holding fish.

https://www.nb.no/maken/item/URN:NBN:no-nb_digifoto_AE2000076358_0002

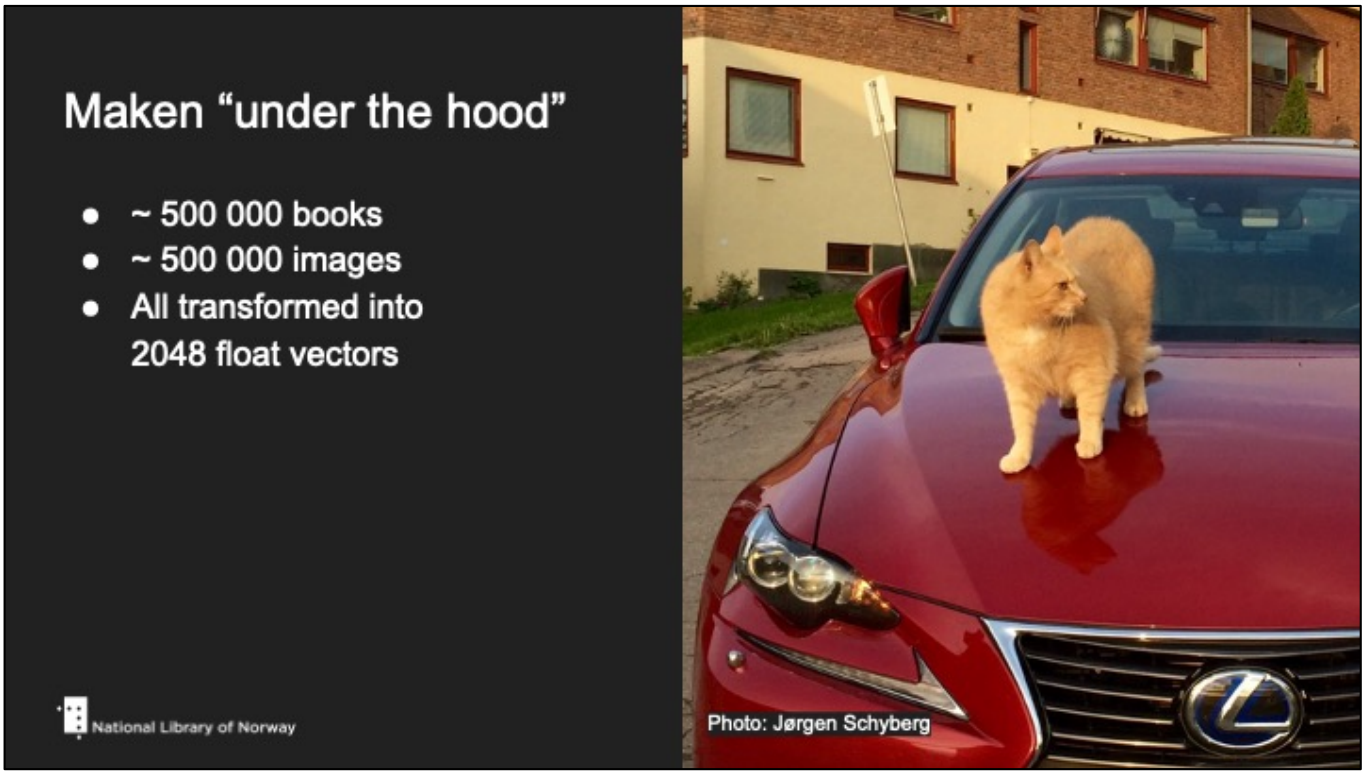

So how do the mechanics of Maken work?

We have ingested vectors for half a million books and half a million images into an ElasticSearch index.

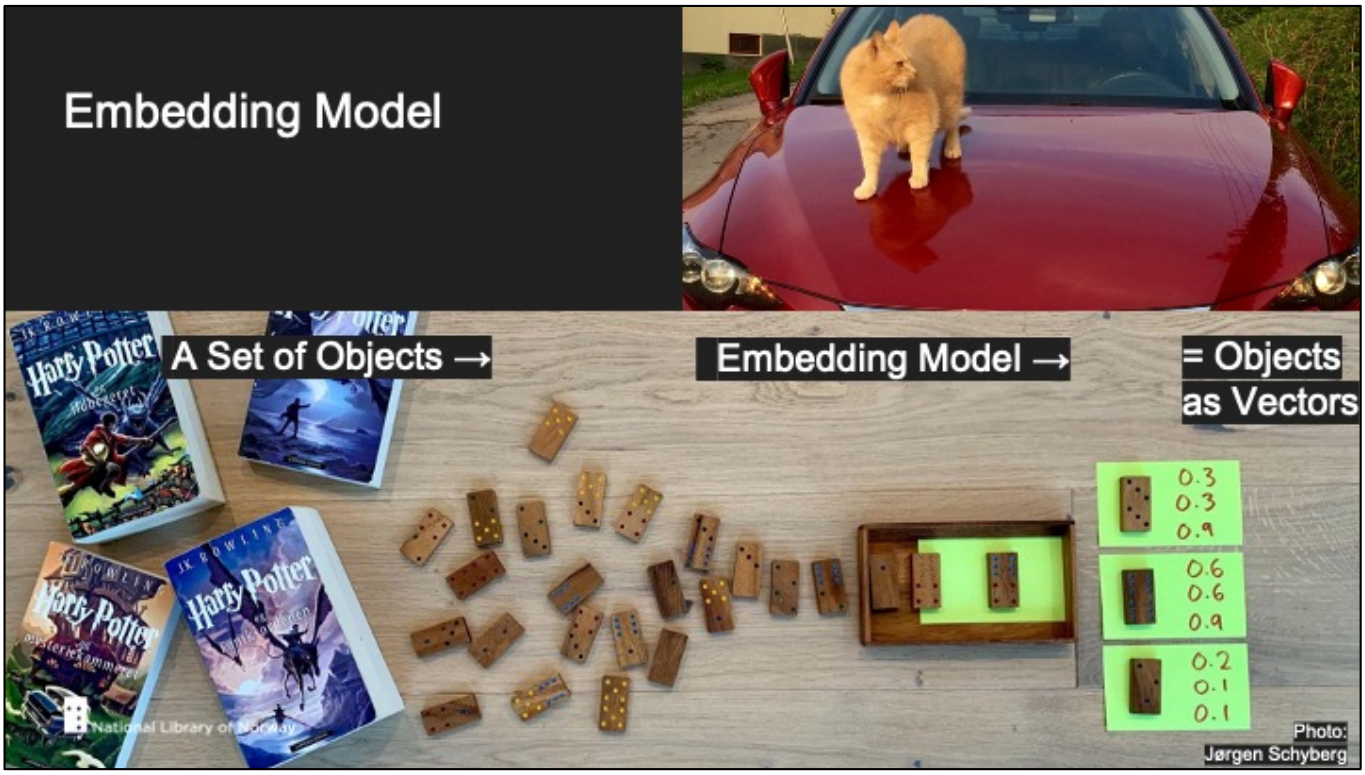

Through what is called an Embedding Model, we have transformed the books and images into long lists of numbers that hold the content of the objects themselves.

In processing, images of cats will get similar lists of numbers to other images of cats, because cats are cat-shaped. Vectors of images of long-haired heavy cats would have similar lists of numbers with vectors of images of short-haired, thin cats, but be quite different from those of cars.

We can do the same with the text of books. I have for the visualisation picked some wooden domino pieces to represent the books, their wooden game box to represent the processing, and some sticky notes to hold the vectorized objects, which are now numbers.

Source: https://www.pinecone.io/learn/vector-embeddings/

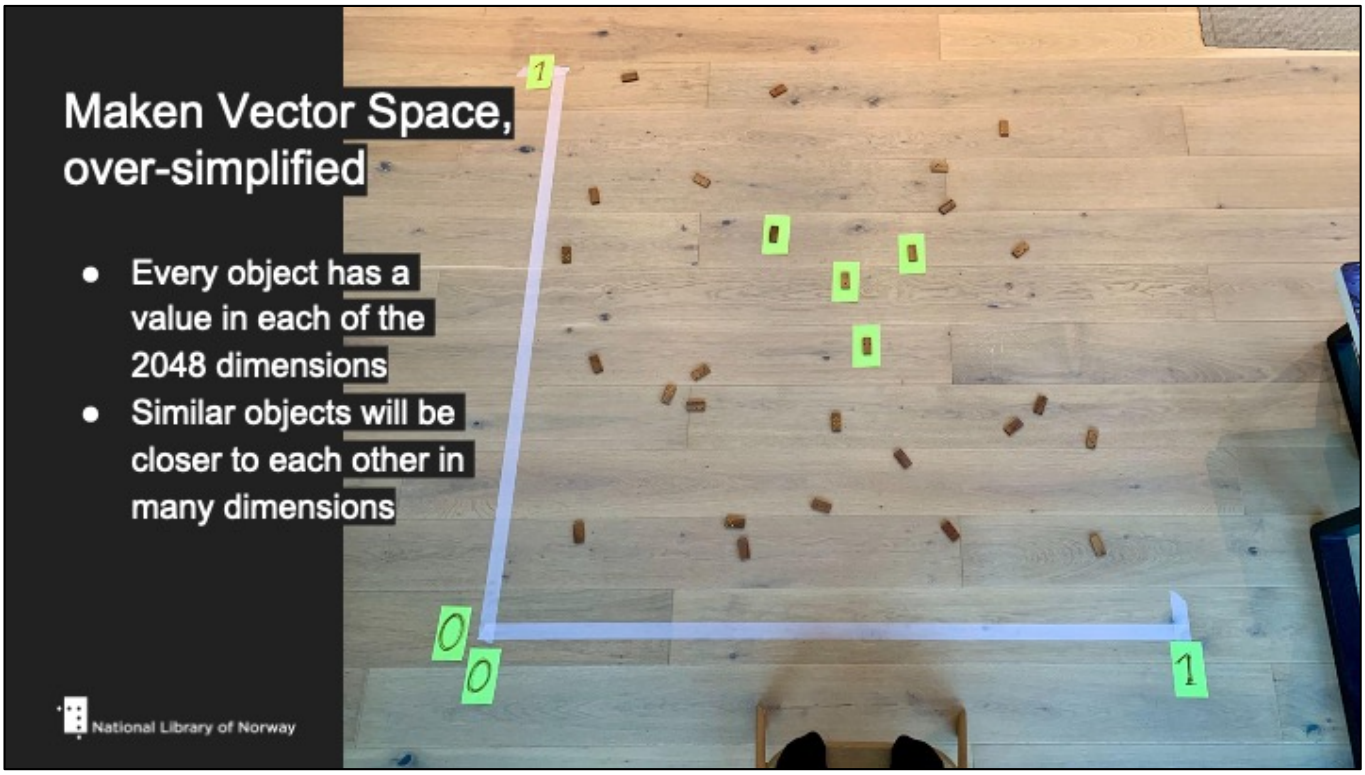

We can mathematically embed these lists of numbers into a 2048-dimensional vector space, and calculate distances between them in all 2048 dimensions.

As 2048 dimensions is a bit hard to grasp for a human used to four dimensions at the max, let's say we collapse the vectors into a two-dimensional coordinate system with an X-axis and a Y-axis, like the one you see here, from my kitchen floor at home using some masking tape, some sticky notes and some domino pieces.

We can choose a few of them.

Source: https://www.pinecone.io/learn/what-is-similarity-search/

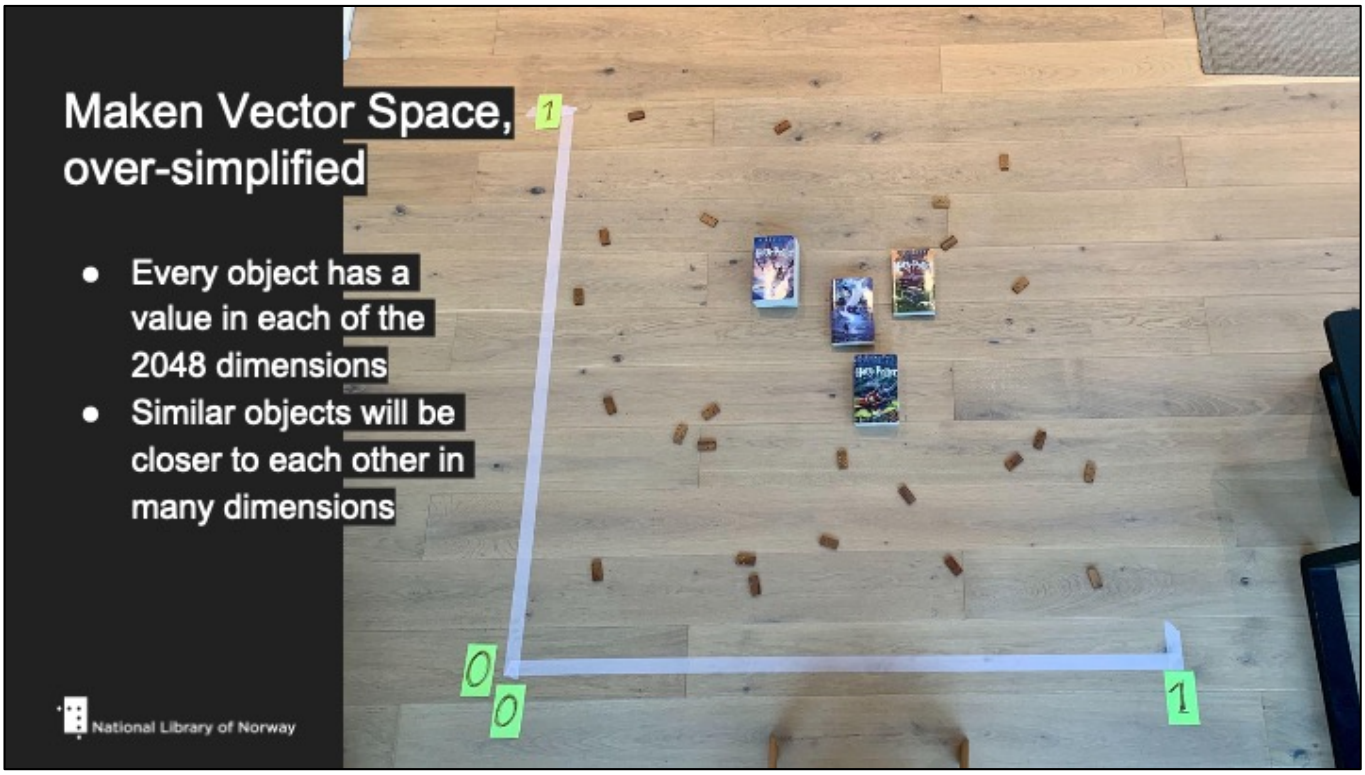

Each of these could for example be a book, let's say Harry Potter and the Prisoner of Azkaban.

In its nearest vicinity, we find three other books that might be related somehow in content. It is likely that we find some other Harry Potter titles nearby.

Source: https://www.pinecone.io/learn/what-is-similarity-search/

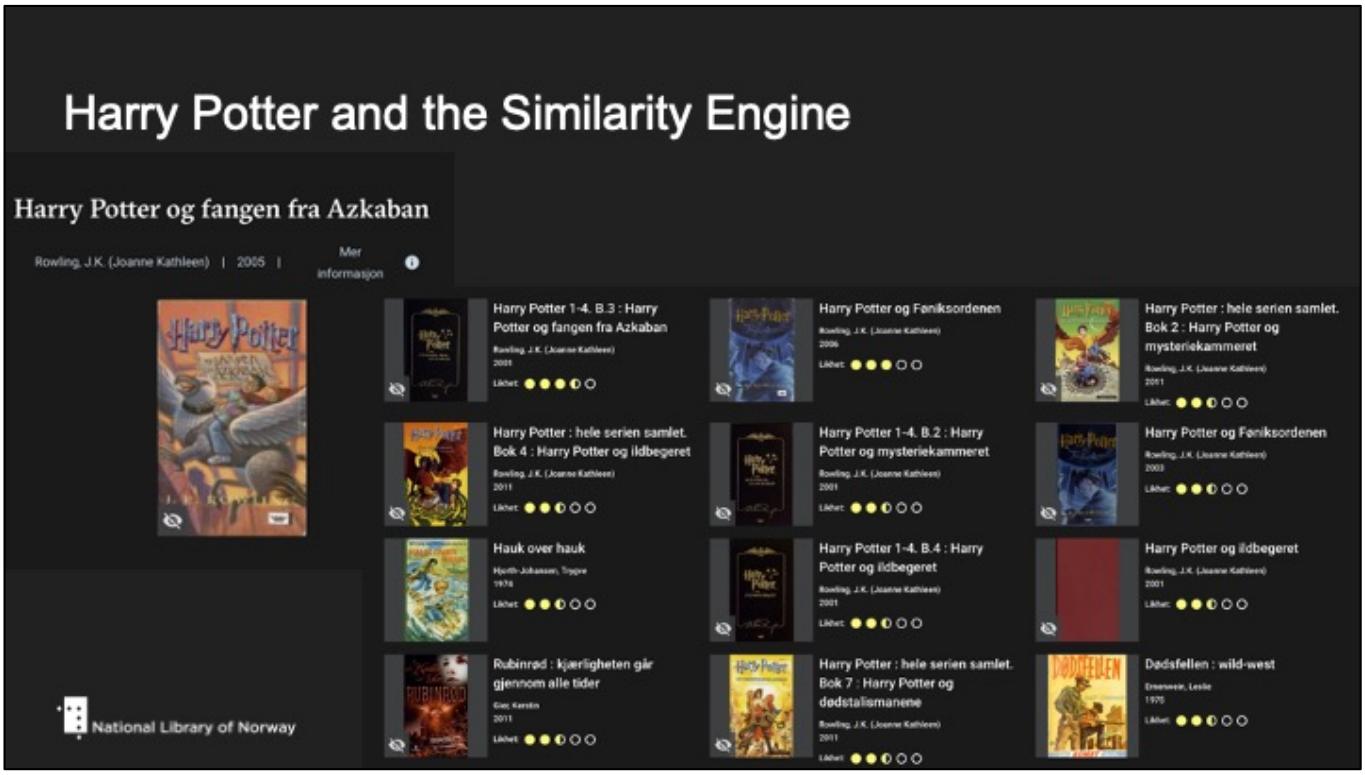

Which is exactly what Maken does if you ask it to present similar books to this edition of Harry Potter and the Prisoner of Azkaban.

https://www.nb.no/maken/item/URN:NBN:no-nb_digibok_2016082408047

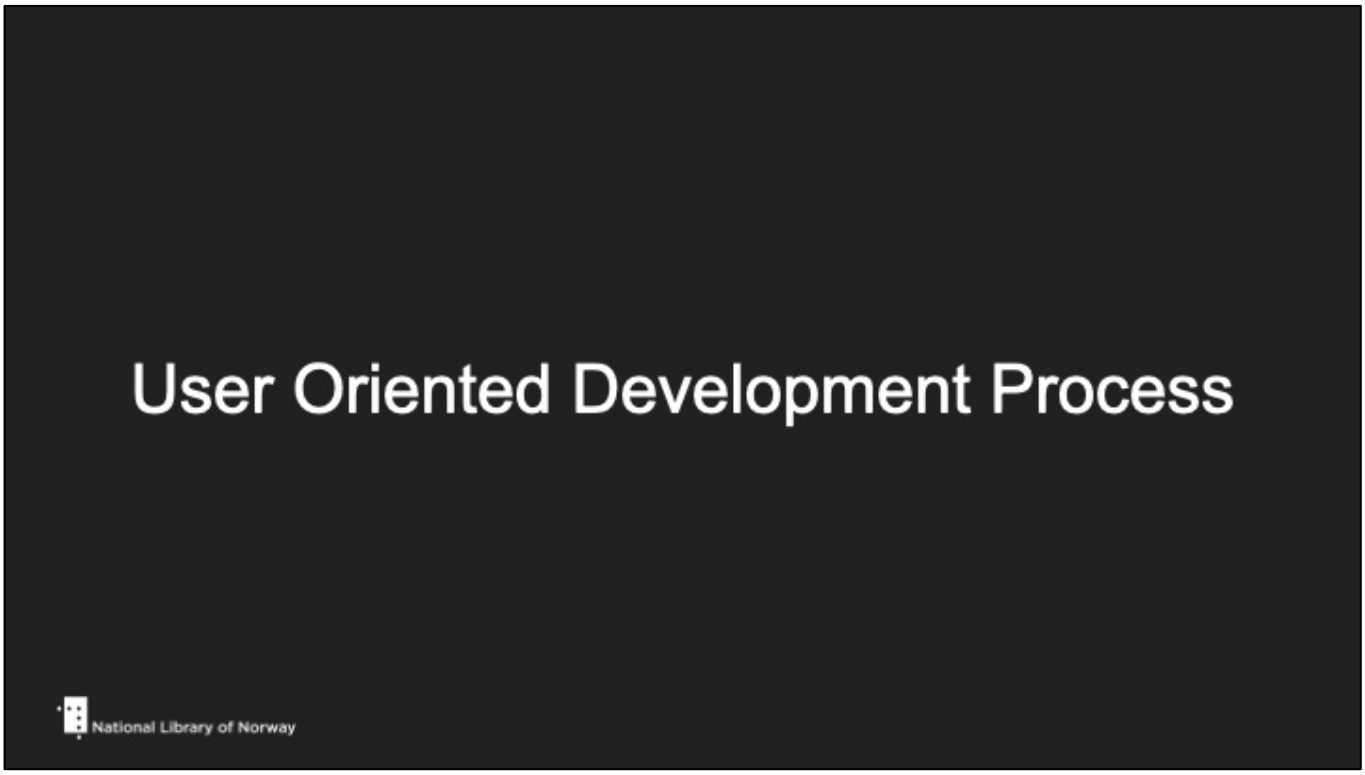

But I really came here to talk about how we approached the users when developing Maken.

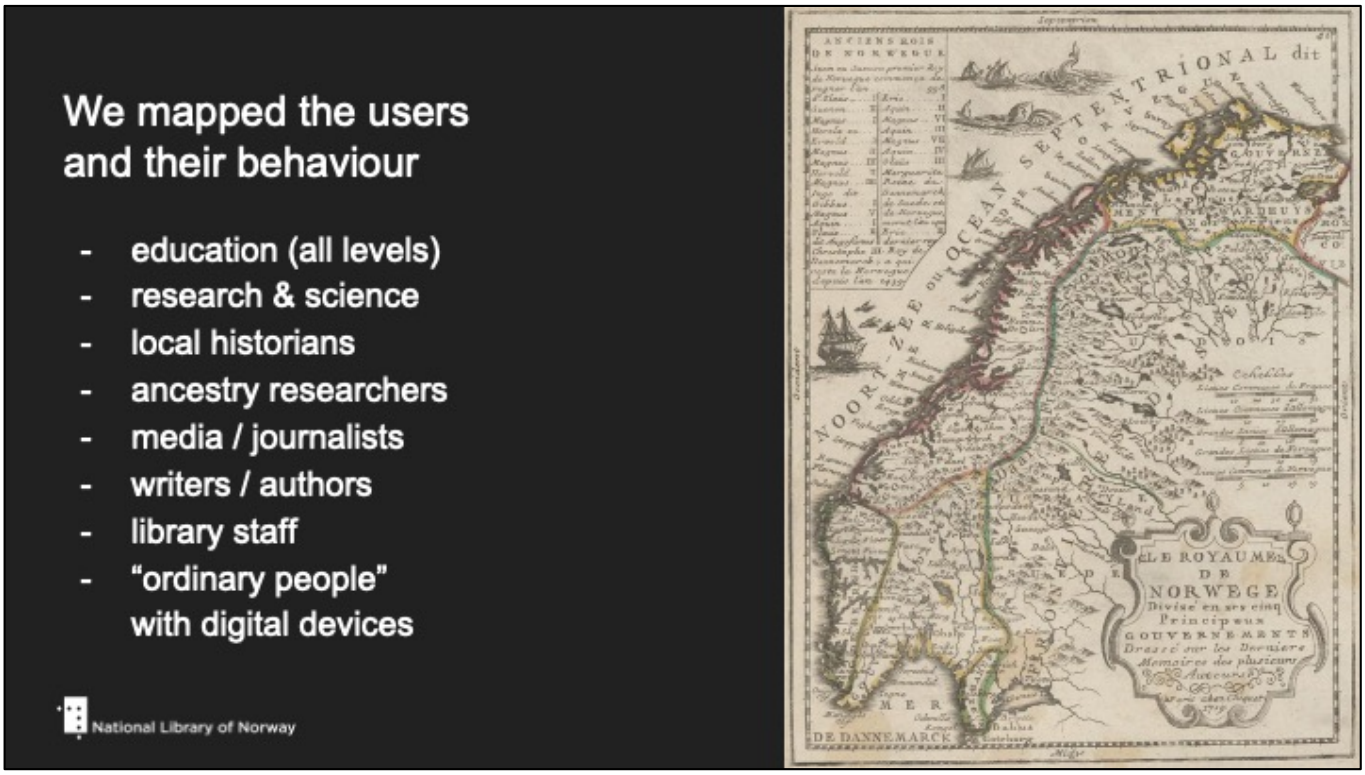

In parallel to developing Maken, we worked to understand how of our digital services were being used, in particular The Digital Library.

Through data collection and analytics from our digital services, user surveys, and also some user testing, we had a general understanding of our user segments.

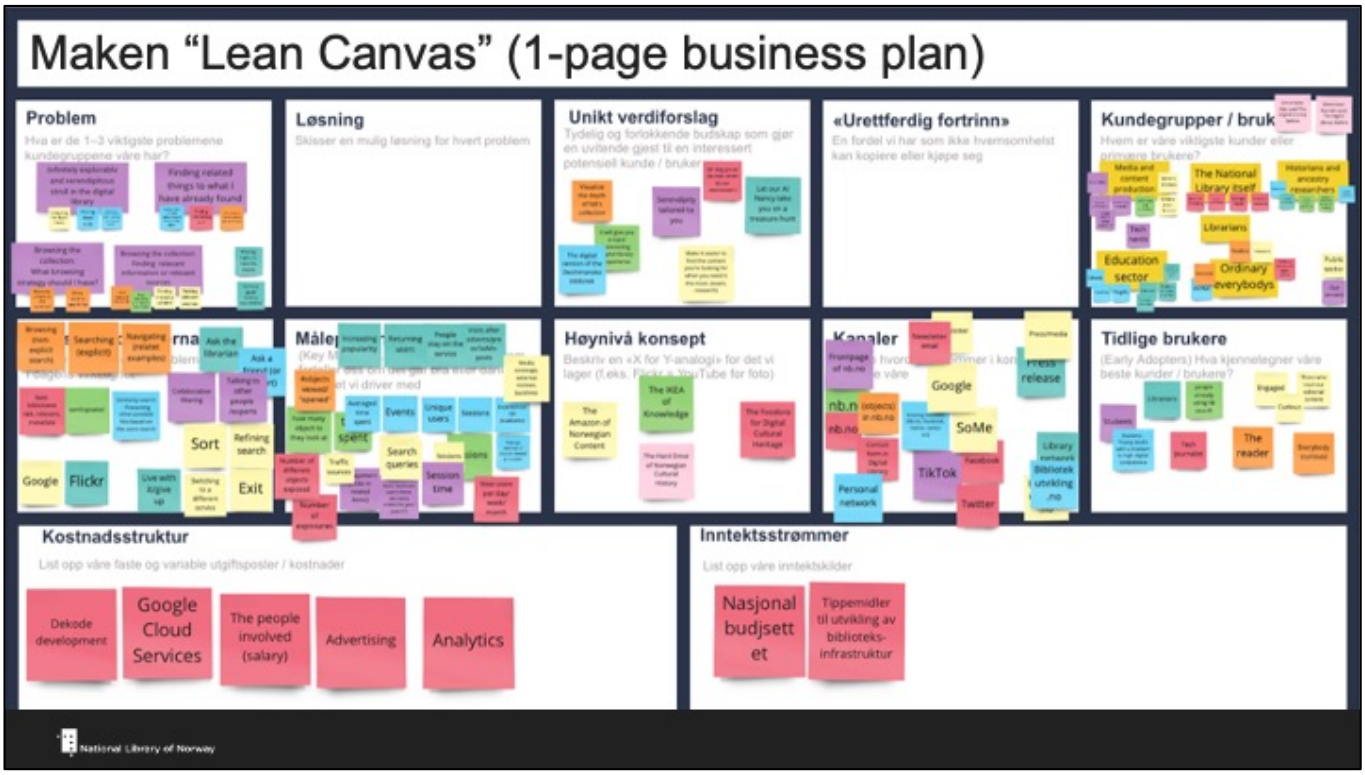

Even though we had some insight, we needed to sharpen our focus and concentrate on our main users and their challenges.

So we ran a workshop called «Lean Canvas»* to:

- identify priority user groups we should aim to understand better
- identify challenges they may have that we might help solve
- understand more of our competition, our surroundings and our own communication and impact

This insight could be used to build hypotheses about our users, which we organised as Personas.

* For more info on Lean Canvas, read this useful introduction: https://medium.com/@steve_mullen/an-introduction-to-lean-canvas-5c17c469d3e0

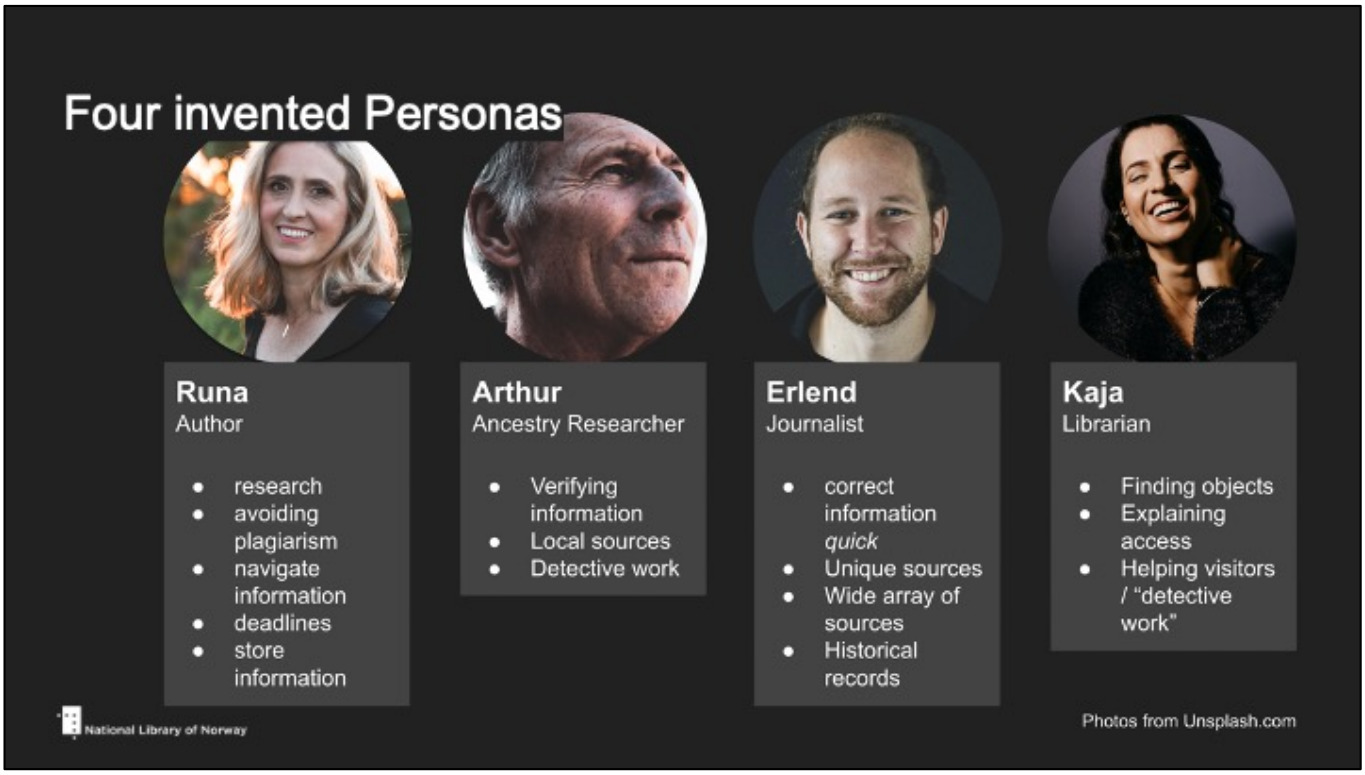

Personas is a much used technique in service design, with the purpose of defining hypotheses and exploring our prejudices about our users and their challenges.

So these four individuals moved in with us for several weeks.

Arthur Grostølen

ancestry researcher

Top tasks:

- Verifying (or falsifying) legends and rumours in the family heritage
- Searching old newspapers for stories of the family
- Browse to find good, reliable sources
- Finding things digitally and on paper/print
- Getting help from other ancestry researchers
- Finding pictures and documentation for lokal events and history

Possible user test task in Maken:
Finding books and/or images related to a particular geographical area, an event or a topic regarding research on his family's heritage

Things to track:
Source/medium (for example MyHeritage), landing page (referral), social shares (outbound, button clicks), pages viewed, downloads, engagement metrics, search queries, media type, year published, title and saved documents

National Library of Norway

"It is a little bit like digging for gold. You get endorphines by learning something new about yourself, the family and society. Being a detective is fun!"

63 years old, male, lives in inland rural town (Geilo)

Ancestry researcher (skilled amateur)
Retired, former accountant

Motivations:

- Exploring unknown areas of my own heritage
- "Detective work"
- Love for my local social environment
- Expanding the circle for family reunions

Even if they don't really exist, they could tell us where and how to look for the answers we needed.

Kaja Husum

librarian

Top tasks:

- Answering and helping visitors, both face to face and via digital inboxes
- Finding objects (mainly digital, sometimes physical)
- Granting access to "restricted" objects
- Giving advice in using Then National Library's digital services and tools
- Receiving visitors, trying to figure out what they are actually looking for

Possible user test task in Maken:
The visitor asks about object X / topic Y, what can you also find that would enlighten more?

Things to track:
Source/medium, previous/next page, internal traffic, objects opened, media types, engagement metrics

National Library of Norway

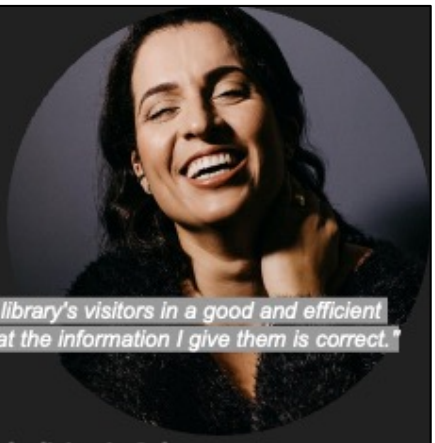

"I wish to guide the library's visitors in a good and efficient way, and ensure that the information I give them is correct."

33 years old, female, living in Oslo

Librarian

Motivations:

- Steady salary, stability
- Providing good service to visitors
- Creating engagement and enthusiasm for reading and making available Norwegian cultural heritage to everybody
- Sense of mastery in finding the right solution to the visitor's needs

That was useful, when we in the next phase were to talk with real people about real challenges, because ...

“What are challenges in your day-to-day work?”

Interviews and user testing:

- Journalists
- Librarians
- Ancestry researchers
- Authors

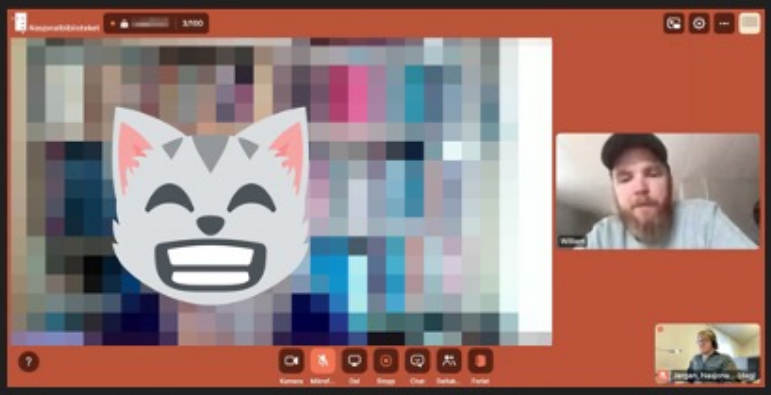

Anonymised screenshot from video interview

National Library of Norway

We needed to know which “Jobs to be Done” might be relevant for Maken to do for its users? So we recruited a panel of users among the priority target groups we had selected for our research.

We did eight 1-hour interviews, over a digital meeting platform. We recorded and logged the interviews.
We made an appointment to meet them again a week after, when they were to test a version of Maken hands-on.
We used what they told us in the interviews to design relevant tasks for them to solve in Maken.
We asked them to share their screen, and we also recorded and logged the user tests, to study them closely.

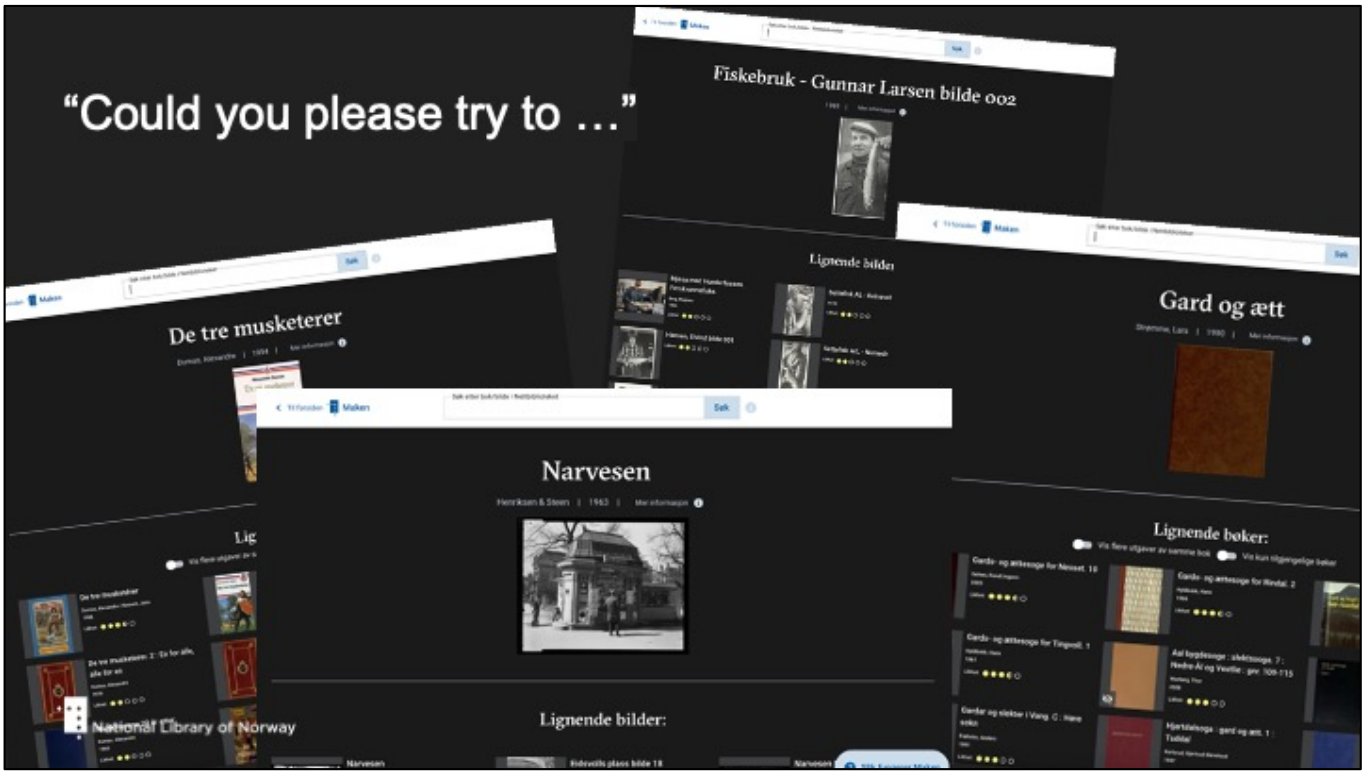

We asked our user panel to try different things in Maken, and learned that ...

- The test users didn’t understand Maken at first glance
- AI driven suggestion is difficult to communicate, especially for books
- People’s intuitive understanding of Maken’s AI is coloured by their previous experience with ordinary search in digital content and metadata
- Similarity between books is often perceived as Topic similarity
- Similarity between images is easier to grasp, perhaps because people have seen it before, or because it is more visually explicit
- The user needs a nudge “to look below the fold” and see that there is more content on the page than the browser shows immediately

But we also learned that Maken can be fun or useful

- An author said that it was similar to browsing the bookshelves of a library, discovering the books that were next to the one you actually came for.
- Another author said it was a good way to discover a wider selection of secondary sources.
- An ancestry researcher said it would have been useful to use one of his or her own images as input and get on-the-fly similar images.
- Newspaper content in the similarity service would be super-handy

Where are we now?

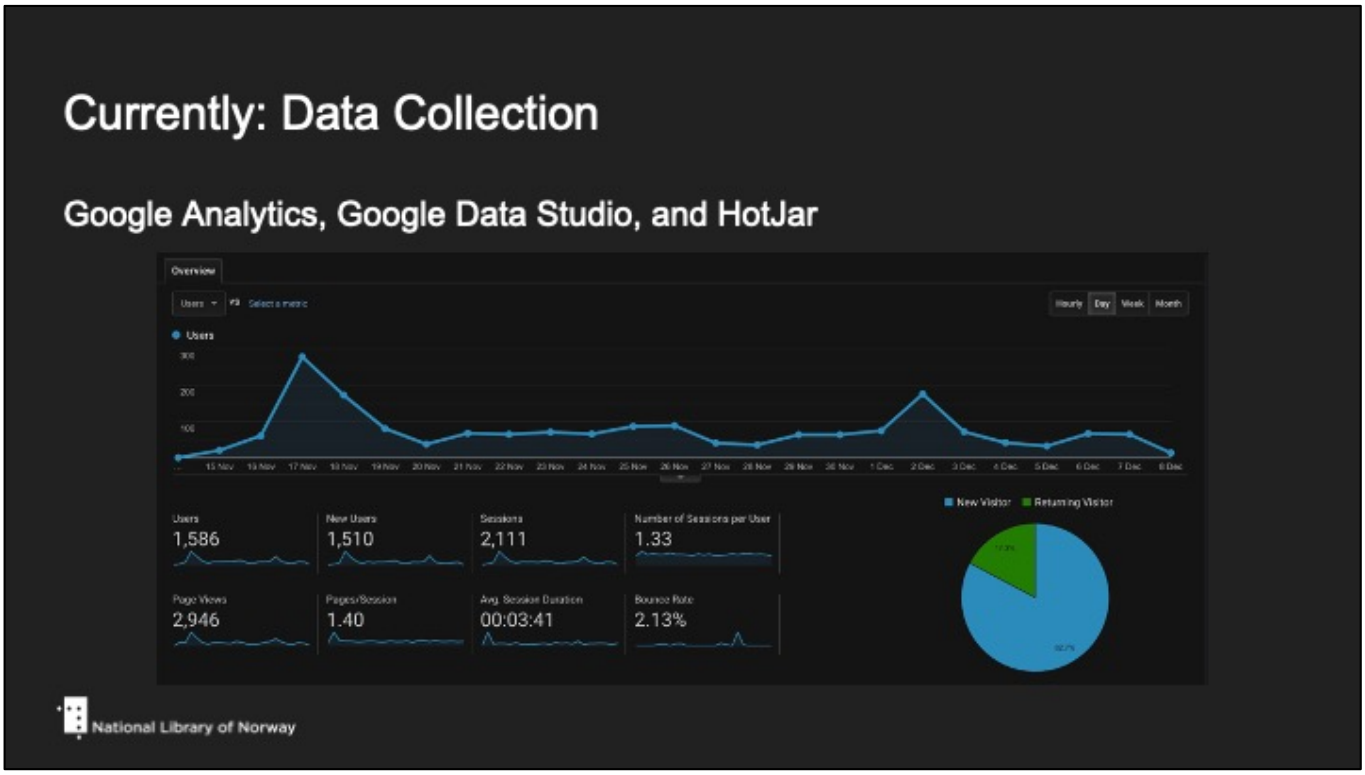

We officially launched mid-November, and are currently in data collection mode to learn about users on a bigger scale. The numbers aren't huge yet, but we hope to see some patterns emerging over time.

- We use Google Analytics to measure traffic, user segments and the users behaviour in general
- We use Google Data Studio to make sense of all the big numbers that we collect
- We use HotJar to understand what the users are doing on the pages

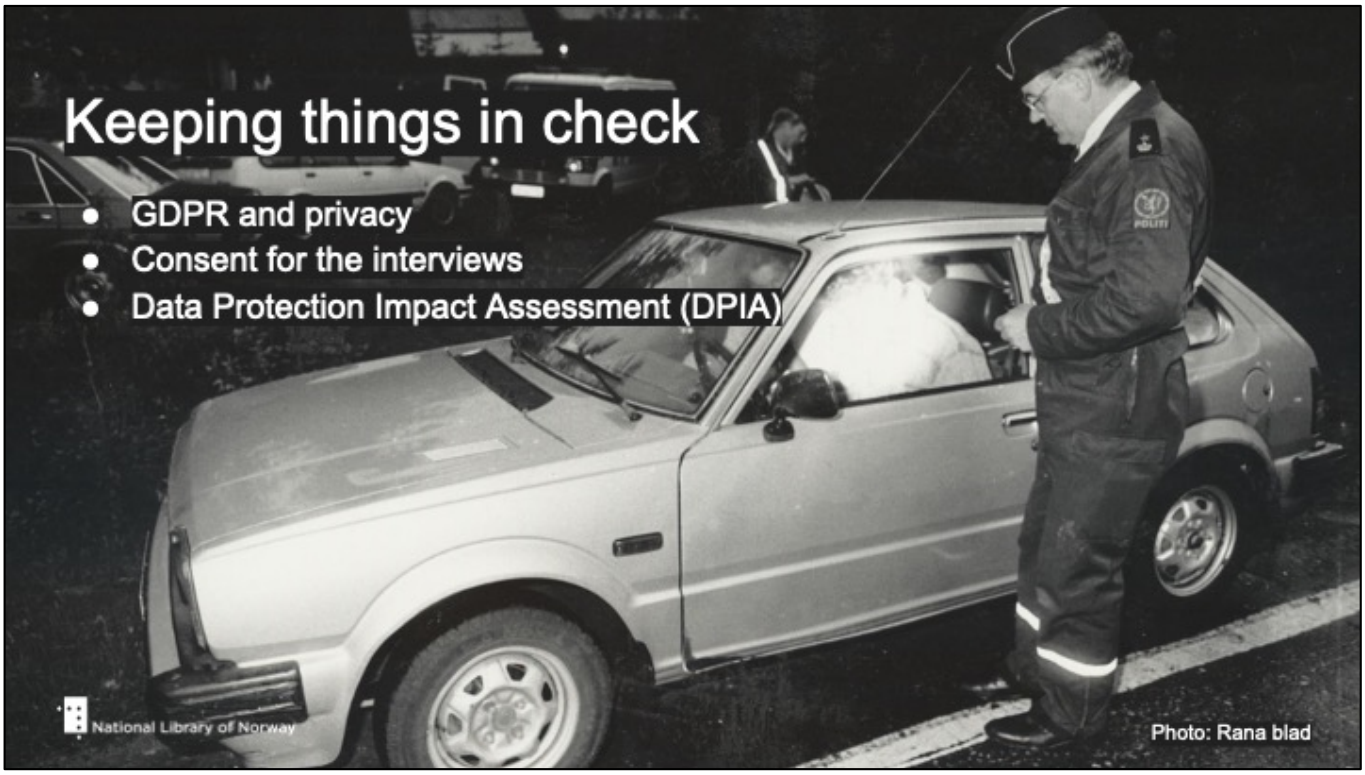

When collecting data for developing and running digital services, data protection is paramount.

- We have to make sure we don't collect or store data we don't need.
- We must assess the data protection in the project and the service, for the people we interviewed and for the users of the actual service.
- We had to collect consent from the users we spoke with, to record, store, analyse, and share what we learned.
- When working with users, we must respect the users and their rights and privileges.

Photo https://www.nb.no/maken/item/URN:NBN:no-nb_digifoto_AE2000068940_0013

Because Maken is still an experiment, we haven't really decided exactly how we will bring it further. But there are a few possible directions:

- Integrate Maken features or learnings into other services (I'll show you an example in a second)
- Extending similarity discovery to collections other than books and images
- Similarity hits in newspapers, but being both images and text at the same time, they are very different from books and images, and articles, events and topics may be more relevant for similarity than entire newspapers
- train the AI to recognise named entities across vast collections of different media types
- Let the user bring their own content and get similarity suggestions for it
- Or what if we use generative machine learning to create new, synthetic cultural heritage? And what would be the status of such content? (Others in this conference, like artist Marion Carré, has already asked the question: Should such AI-generated objects be preserved or

discarded, or something in between?)

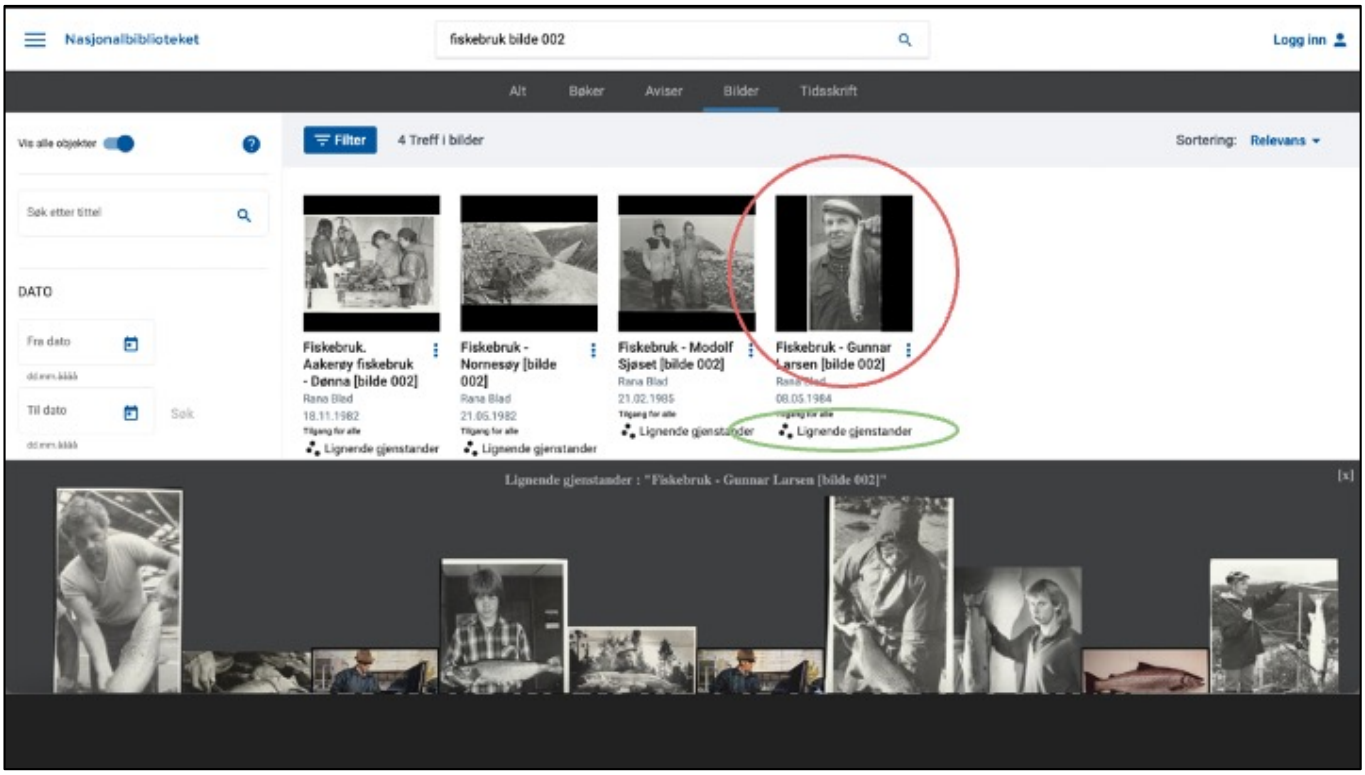

I will round up with a screenshot of a proof-of-concept that my colleague and co-author Javier has made, an experimental Firefox plugin for our main Digital Library, using the Maken similarity engine to add similarity value to the regular library search of https://nb.no/search.

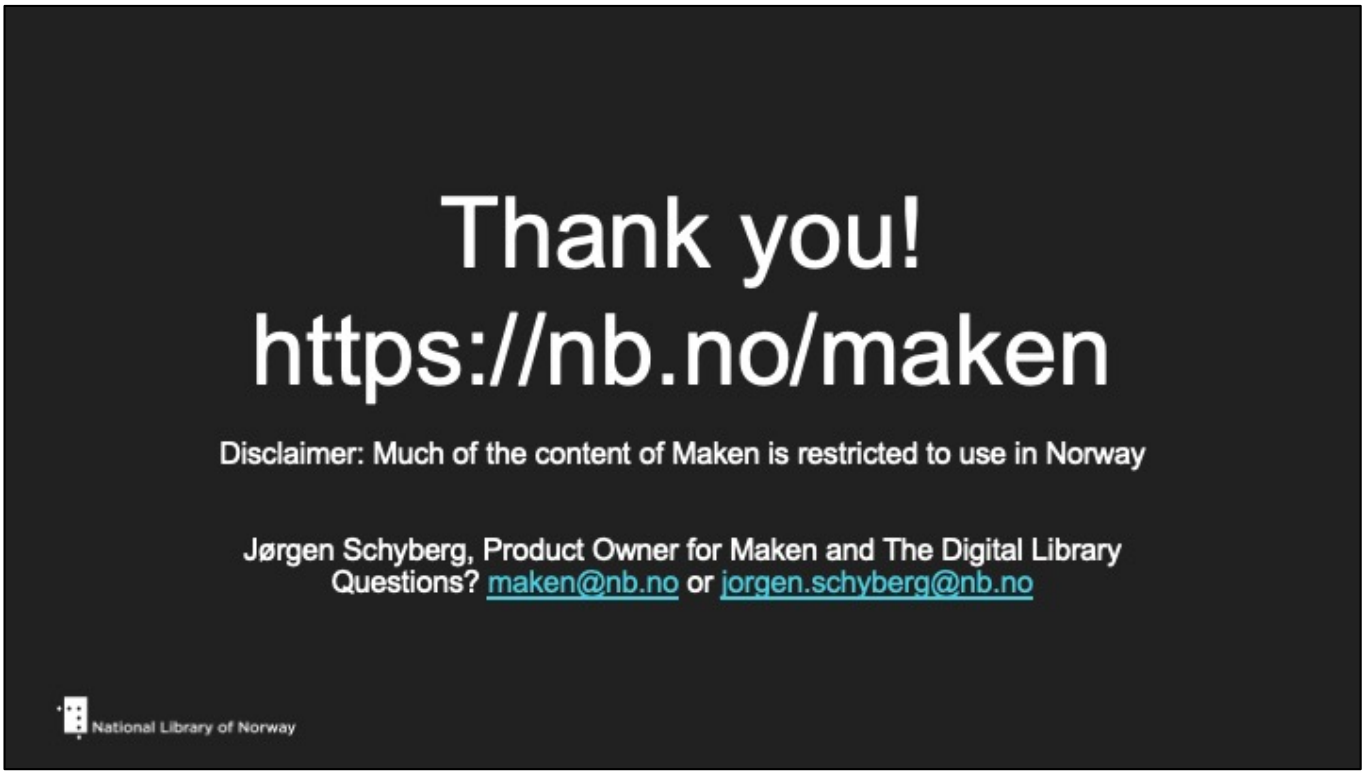

From working with Maken, I believe we will get used to and learn to appreciate what AI can add to user experience, and supplement how we work with the content, both as individuals and as Library institutions. To me, it feels like machine learning and AI over time will become second nature.

Thank you very much for listening.