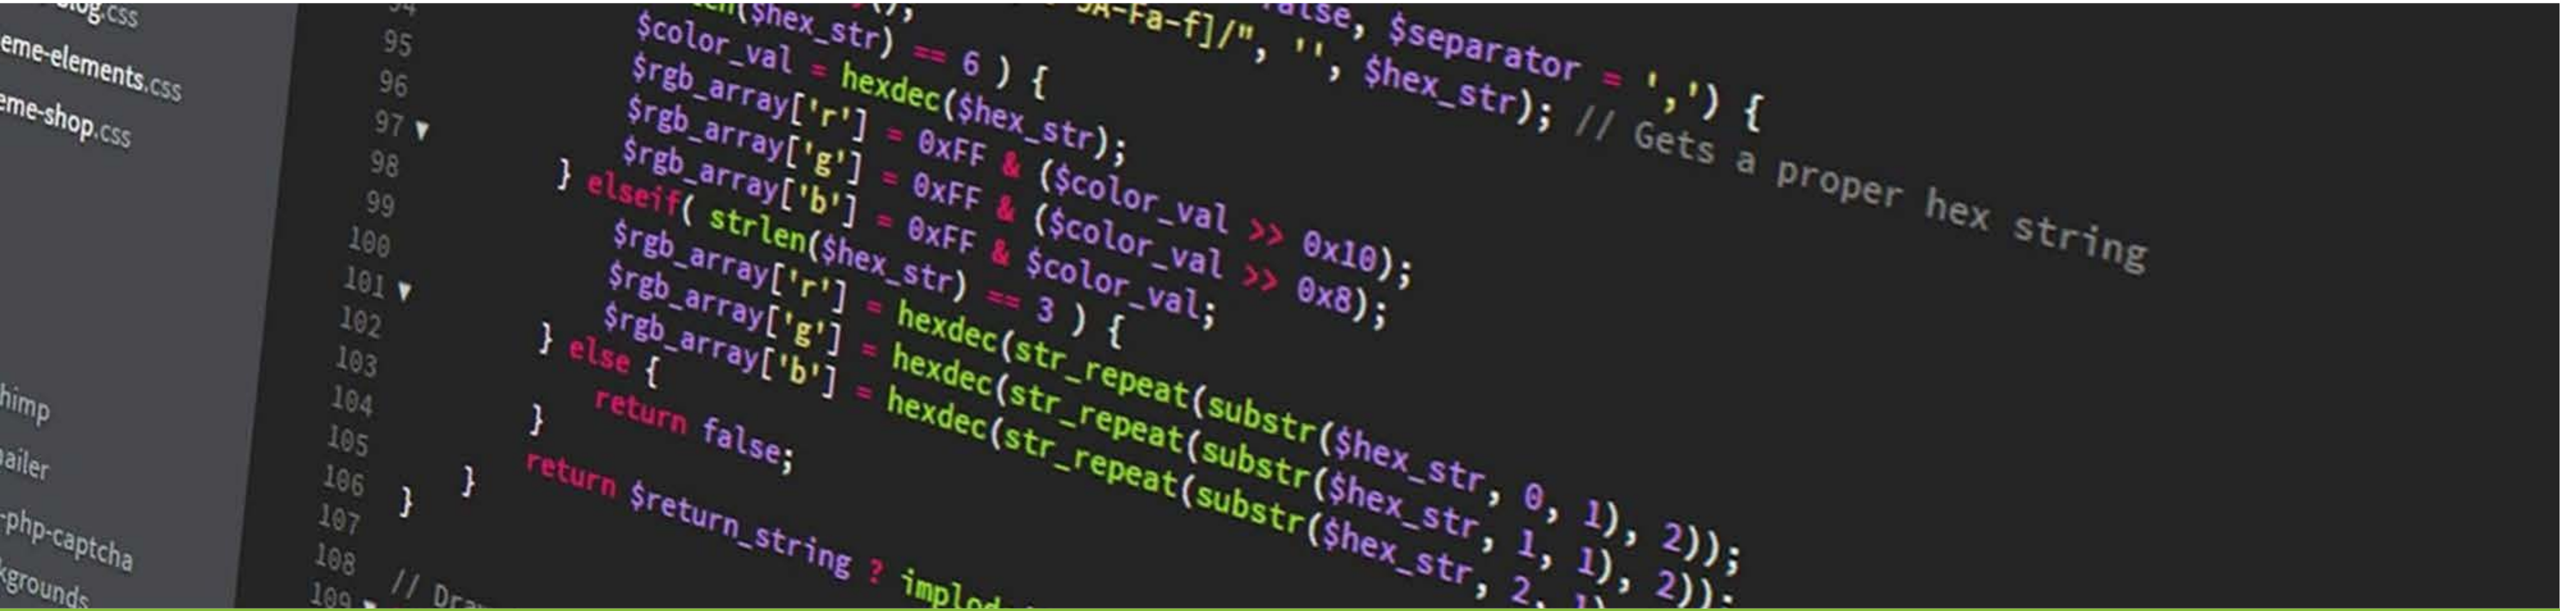

Recommender Systems and Adaptive UIs

Recommender Systems

Definition

"Recommender Systems (RSs) are software tools and techniques providing suggestions for items to be of use to a user. […] The suggestions relate to various decision-making processes, such as what items to buy, what music to listen to, or what online news to read." [1, p1]

[1] Francesco Ricci, Lior Rokach and Bracha Shapira. Introduction to Recommender Systems Handbook. In F. Ricci et al. (eds.), Recommender Systems Handbook. Springer Science+Business Media 2011.

Recommender Systems – Examples?

- How are recommender systems used in these services?
- How do recommender systems impact the user experience?
- What are user interface patterns used with recommender systems?

Recommender Systems

- Why are recommender systems used?
- What is the main function?
- What data do recommender systems require?
- How do recommender systems make a user interface intelligent?

Carlos A. Gomez-Uribe and Neil Hunt. 2015. The Netflix Recommender System: Algorithms, Business Value, and Innovation. ACM Trans. Manage. Inf. Syst. 6, 4, Article 13 (December 2015). DOI: https://doi.org/10.1145/2843948

What Book do you Recommend to me?

What Book do you Recommend to me?

Example: https://tastedive.com/

How does this work?

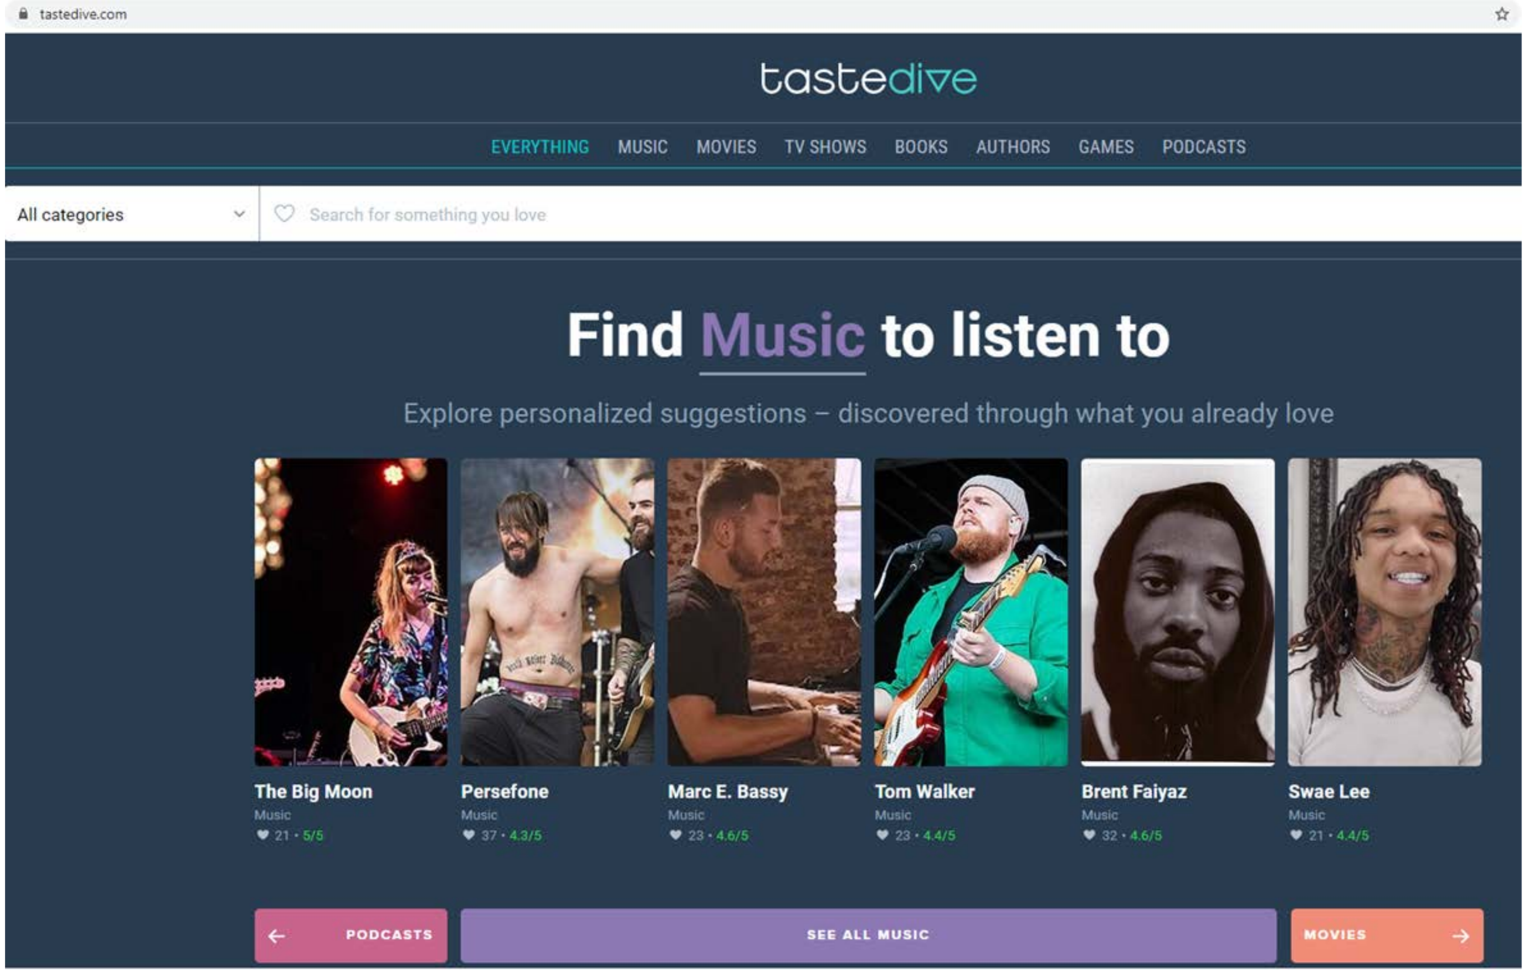

What Book do you Recommend to me?

? ?

What Book do you Recommend to me?

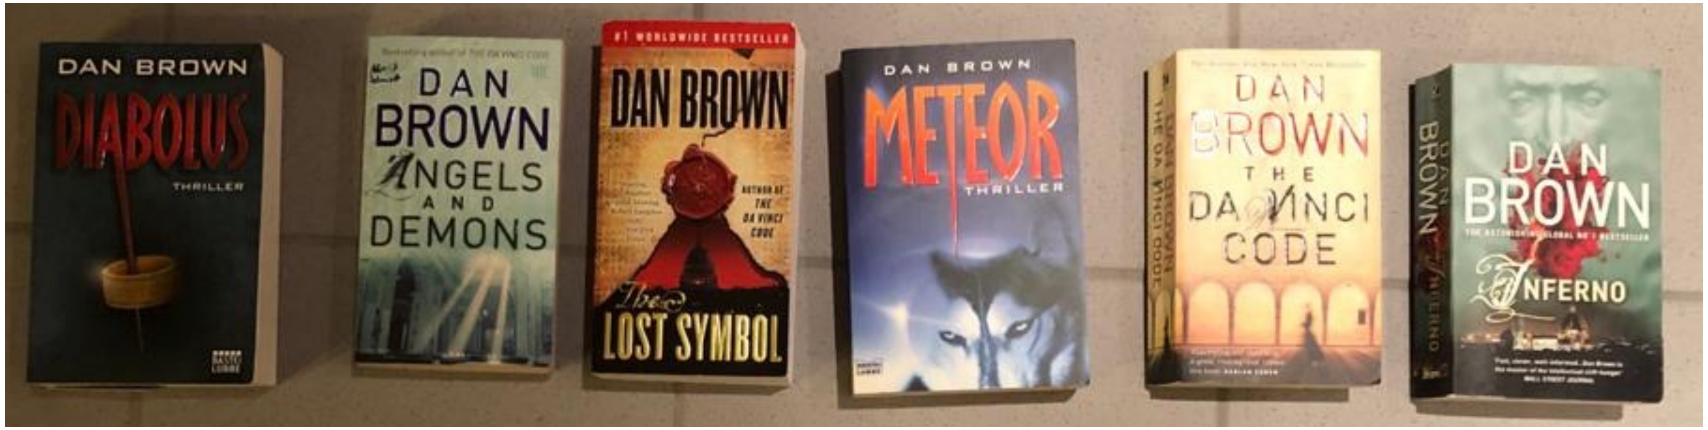

This is what I read…

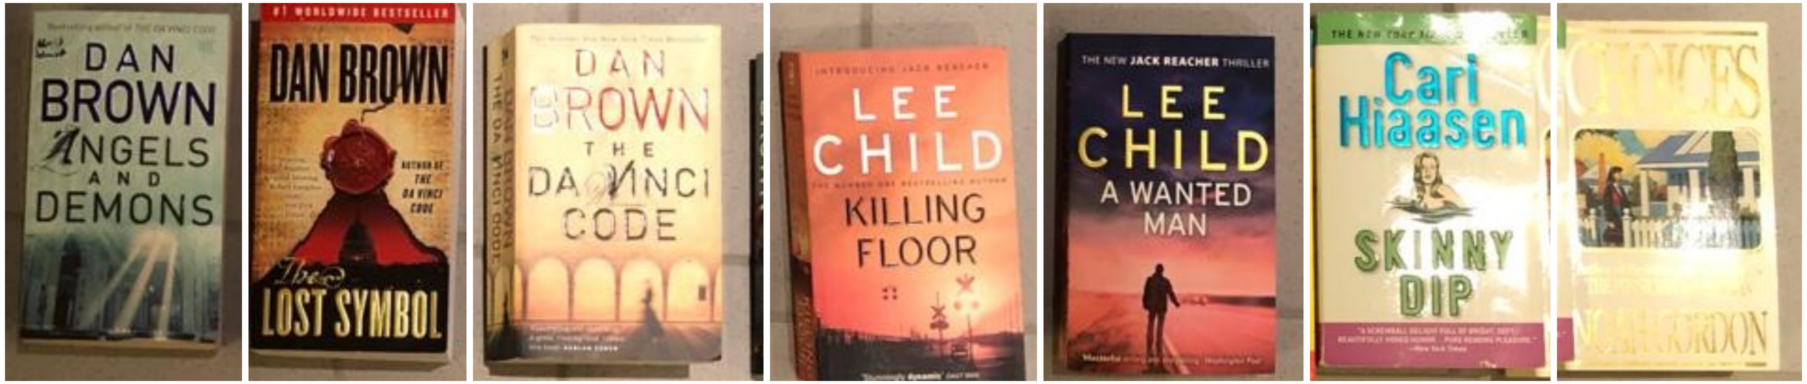

Which of these books do you suggest?

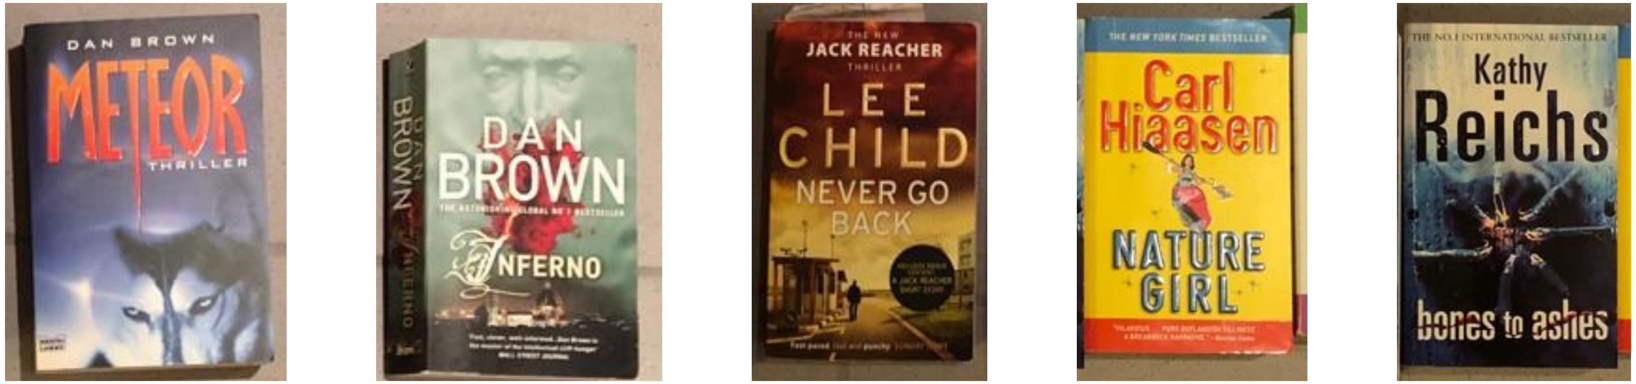

Approaches to Recommender Systems

- **Collaborative Filtering**
  "people who liked what you liked also liked X, hence I suggest X to you"
- **Demographics**
  "many people of you age, income, family status, and education like X, hence I suggest X to you"
- **Social Relationships**
  "people you hang out with, your friend, or friend of you friends like X, hence I suggest X to you"
- **Content-based Filtering**
  "X is similar (based on a similarity measure for a domain) to what you liked before, hence I suggest X to you"
- **Contextual**
  "many people in you current situation/context/location like X, as you are in now this context , I suggest X to you"

Content Based Filtering

Basic Approach

- What properties / factors / dimensions are describing and discriminating the products?
  - taxonomy, list of dimensions or set of criteria
- Each item can be categorized based on the criteria/dimensions/taxonomy
  - for each item represent how much it matches these criteria in a vector
- How much does the user like products based on the criteria/dimensions/taxonomy?
  - vector for the user that represents their relationship to these criteria
- Identify and recommend items that fit the user's criteria
  - calculate the similarity between item vectors and user vectors

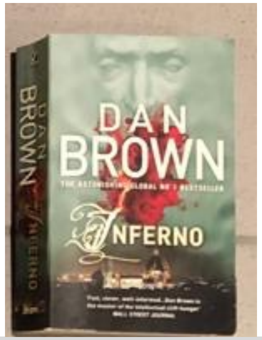

Language, genre, time set, main characters, location/setting, theme, length, time written, number of parallel plots, language complexity,…

Idea of User Based Collaborative Filtering

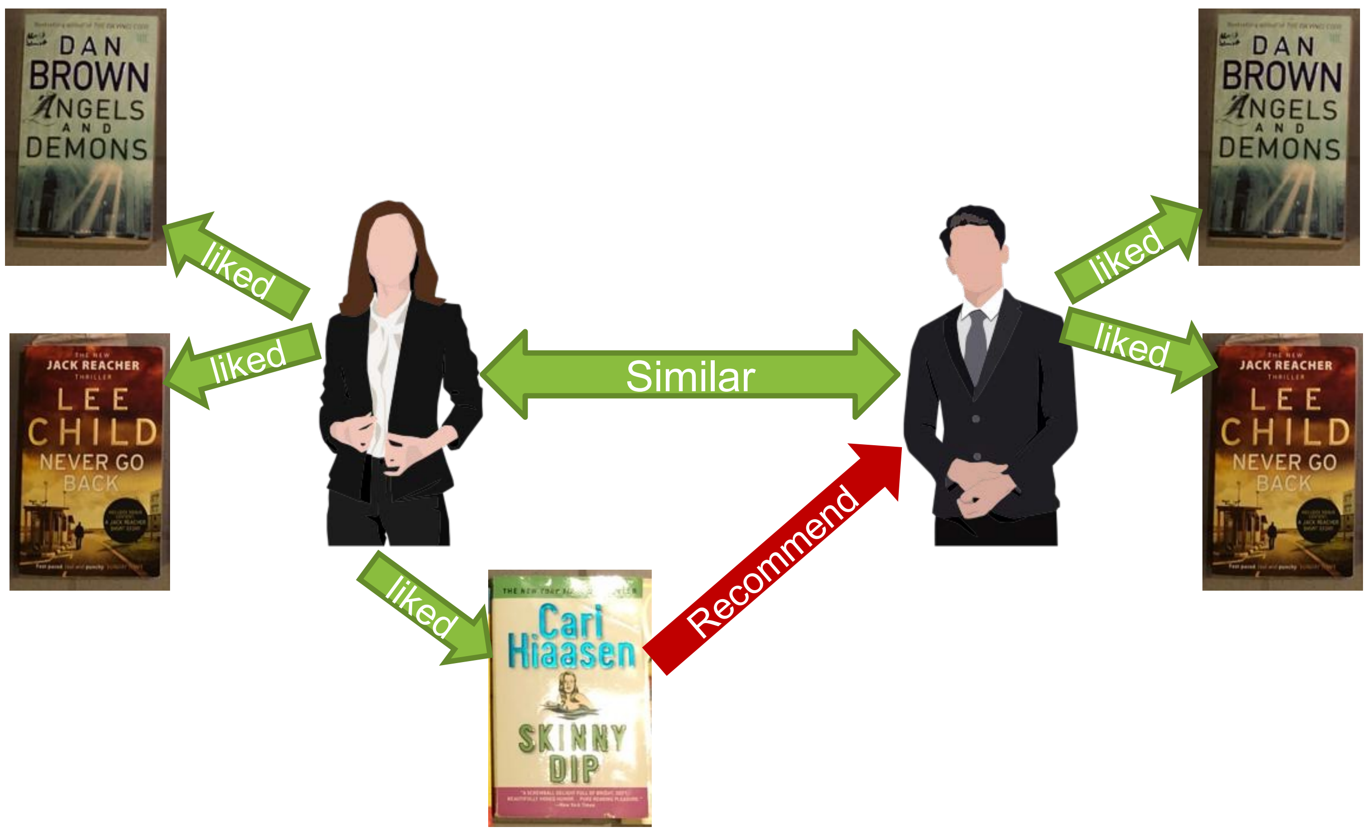

Breakout Sessions 1
How to get item ratings?

Explicitly and Implicitly

- Scenario 1: Web based reading platform/library
- Scenario 2: Public transport application
- Scenario 3: Online platform for travel recommendations

What do you get ratings for? What is the item?

What relevant questions can you ask to get item ratings?

How do you get implicit ratings from “automated observation”?


DAN BROWN ANGELS AND DEMONS
DAN BROWN The LOST SYMBOL
DAN BROWN THE DA VINCI CODE
DAN BROWN METEOR THRILLER
DAN BROWN INFERNO
JACK REACHER LEE CHILD NEVER GO BACK
Carl Hiaasen SKINNY DIP
Kathy Reichs bones to ashes
DAN BROWN The LOST SYMBOL
LEE CHILD A WANTED MAN
JACK REACHER LEE CHILD NEVER GO BACK
DAN BROWN ANGELS AND DEMONS
DAN BROWN THE DA VINCI CODE
DAN BROWN METEOR THRILLER
DAN BROWN INFERNO
JACK REACHER LEE CHILD NEVER GO BACK
Carl Hiaasen SKINNY DIP
Carl Hiaasen NATURE GIRL
Kathy Reichs bones to ashes
DAN BROWN ANGELS AND DEMONS
DAN BROWN THE DA VINCI CODE
LEE CHILD KILLING FLOOR
Carl Hiaasen NATURE GIRL


DAN BROWN ANGELS AND DEMONS
DAN BROWN The LOST SYMBOL
DAN BROWN THE DA VINCI CODE
DAN BROWN METEOR THRILLER
DAN BROWN INFERNO
LEE CHILD NEVER GO BACK
LEE CHILD A WANTED MAN
LEE CHILD KILLING FLOOR
Carl Hiaasen SKINNY DIP
Carl Hiaasen NATURE GIRL
Kathy Reichs bones to ashes
CHOICES


DAN BROWN ANGELS AND DEMONS
DAN BROWN The LOST SYMBOL
DAN BROWN THE DA VINCI CODE
DAN BROWN METEOR THRILLER
DAN BROWN INFERNO
JACK REACHER LEE CHILD NEVER GO BACK
Carl Hiaasen SKINNY DIP
Kathy Reichs bones to ashes
DAN BROWN The LOST SYMBOL
LEE CHILD A WANTED MAN
JACK REACHER LEE CHILD NEVER GO BACK
DAN BROWN ANGELS AND DEMONS
DAN BROWN THE DA VINCI CODE
DAN BROWN METEOR THRILLER
DAN BROWN INFERNO
JACK REACHER LEE CHILD NEVER GO BACK
CHOICES
Carl Hiaasen SKINNY DIP
Carl Hiaasen NATURE GIRL
Kathy Reichs bones to ashes
DAN BROWN ANGELS AND DEMONS
DAN BROWN THE DA VINCI CODE
LEE CHILD KILLING FLOOR
CHOICES
Carl Hiaasen NATURE GIRL

	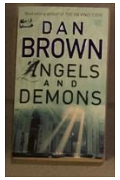	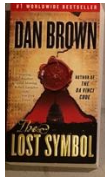	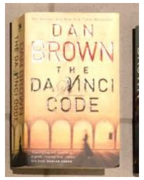	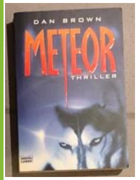	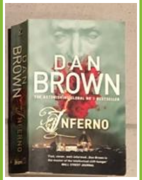	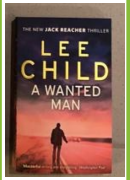	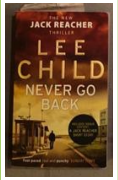	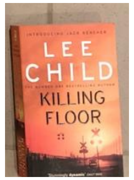	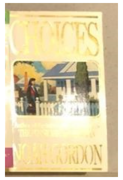	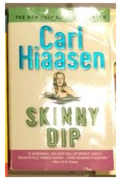	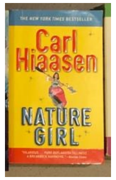	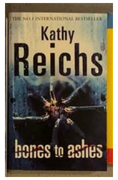
	8	7	9	4	9		10			4		7
		2				6	9					
	4		6	10	8		9		6	7	9	8
	4	?	6	?	?	?	?	9	5	?	10	?

	Item1	item2	item3	item4	Item5	item6	item7	item8	Item9	item10	item11	item12
u1	8	7	9	4	9		10			4		7
u2		2				6	9					
u3	4		6	10	8		9		6	7	9	8
u4	4	?	6	?	?	?	?	9	5	?	10	?

Discussion: Getting user input?

Where does Ethics come in?

Recommender Systems
Getting user data

What are ways to get data from the user for recommender systems?

- Explicit:
  - Questionnaires at setup time
  - Ratings of items
  - Feedback questions
- Implicit
  - Using items (e.g. watching a movie, listen to music, putting something on a watch list, buying a book)
  - Sharing items (e.g. recommending an article, retweets)
  - Removing items (e.g. deleting a playlist, skipping a suggested song)

Collaborative Filtering

Basic Approach

- Tables:
  - User, Item, Rating
  - Item, Item description
- Approach
  - Users rate items (implicitly, explicitly), e.g. like a movie, buy a book, recommend an hotel, …
    - Table: user, item, rating
  - Compute similarity between users (or between items)
    - set of users (or set of items) that is similar to the user (the item) we create a recommendation for
  - Predict items (new to the target user) based on the information of similar users
    - weighted list of recommendations

Calculating the similarity

How similar are user to each other?

- Calculating the difference between rating vectors, e.g. k-nearest neighbor, Euclidian distance, Pearson correlation, Cosine distance

	Item A	Item B	Item C	Item D	Cosine Similarity(Ui,Ux)
User 1	7	10	3	4	0,938088
User 2	7	8	5	3	0,948847
User 3	8	2	5	3	0,708337
User 4	2	9	5	2	0,986921
User 5	1	2	5	2	0,828970
User X	3	8	5	3	1,000000

$$\text{similarity} = \cos(\theta) = \frac{\mathbf{A} \cdot \mathbf{B}}{\|\mathbf{A}\|\|\mathbf{B}\|} = \frac{\sum_{i=1}^{n} A_i B_i}{\sqrt{\sum_{i=1}^{n} A_i^2} \sqrt{\sum_{i=1}^{n} B_i^2}}$$

https://en.wikipedia.org/wiki/Cosine_similarity

Example based on:
https://buddingdatascientist.wordpress.com/2017/03/22/how-to-calculate-cosine-similarity-in-excel/

Heuristic required for tables with missing values!

Assessing Similarity?

How could Machine Learning play a role here?

	8	7	9	4	9		10			4		7
		2				6	9					
	4		6	10	8		9		6	7	9	8
	4	?	6	?	?	?	?	9	5	?	10	?

	Item A	Item B	Item C	Item D	Cosine Similarity(Ui,Ux)
User 1	7	10	3	4	0,938088
User 2	7	8	5	3	0,948847
User 3	8	2	5	3	0,708337
User 4	2	9	5	2	0,986921
User 5	1	2	5	2	0,828970
User X	3	8	5	3	1,000000

Predicting the right item?

Example Similarity Matrix

Similarities are NOT calculated – for illustration only

8	7	9	4	9		10			4		7
	2				6	9					
4		6	10	8		9		6	7	9	8
4	?	6	?	?	?	?	9	5	?	10	?

Similarity Matrix

0.4
0.7
0.9
1

Based on https://medium.com/swlh/how-to-build-simple-recommender-systems-in-python-647e5bcd78bd

Weighted Rating Matrix

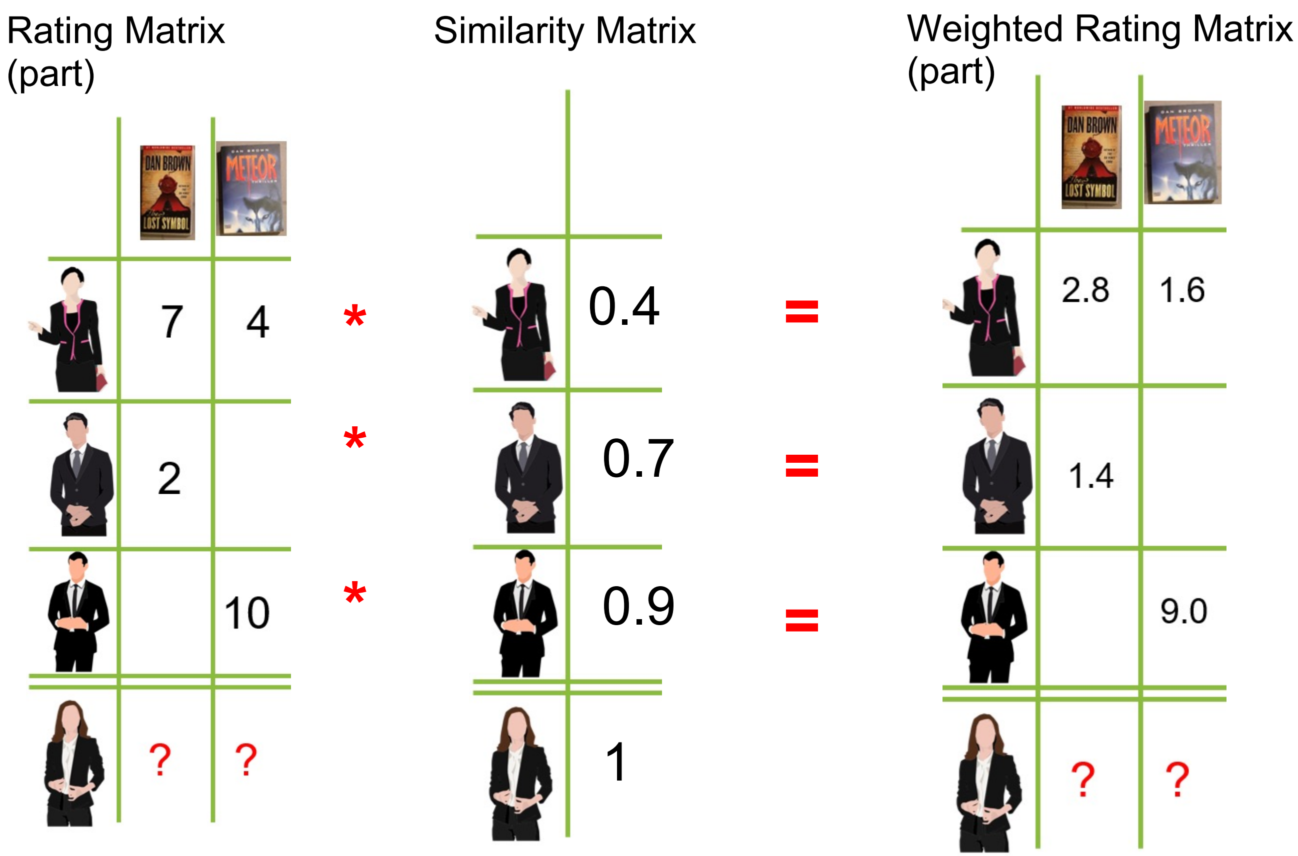

Based on https://medium.com/swlh/how-to-build-simple-recommender-systems-in-python-647e5bcd78bd

Recommendation

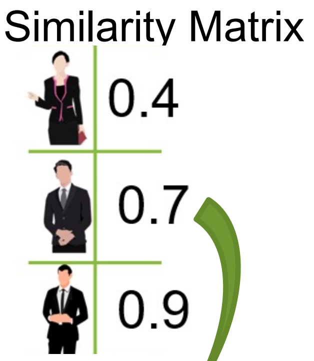

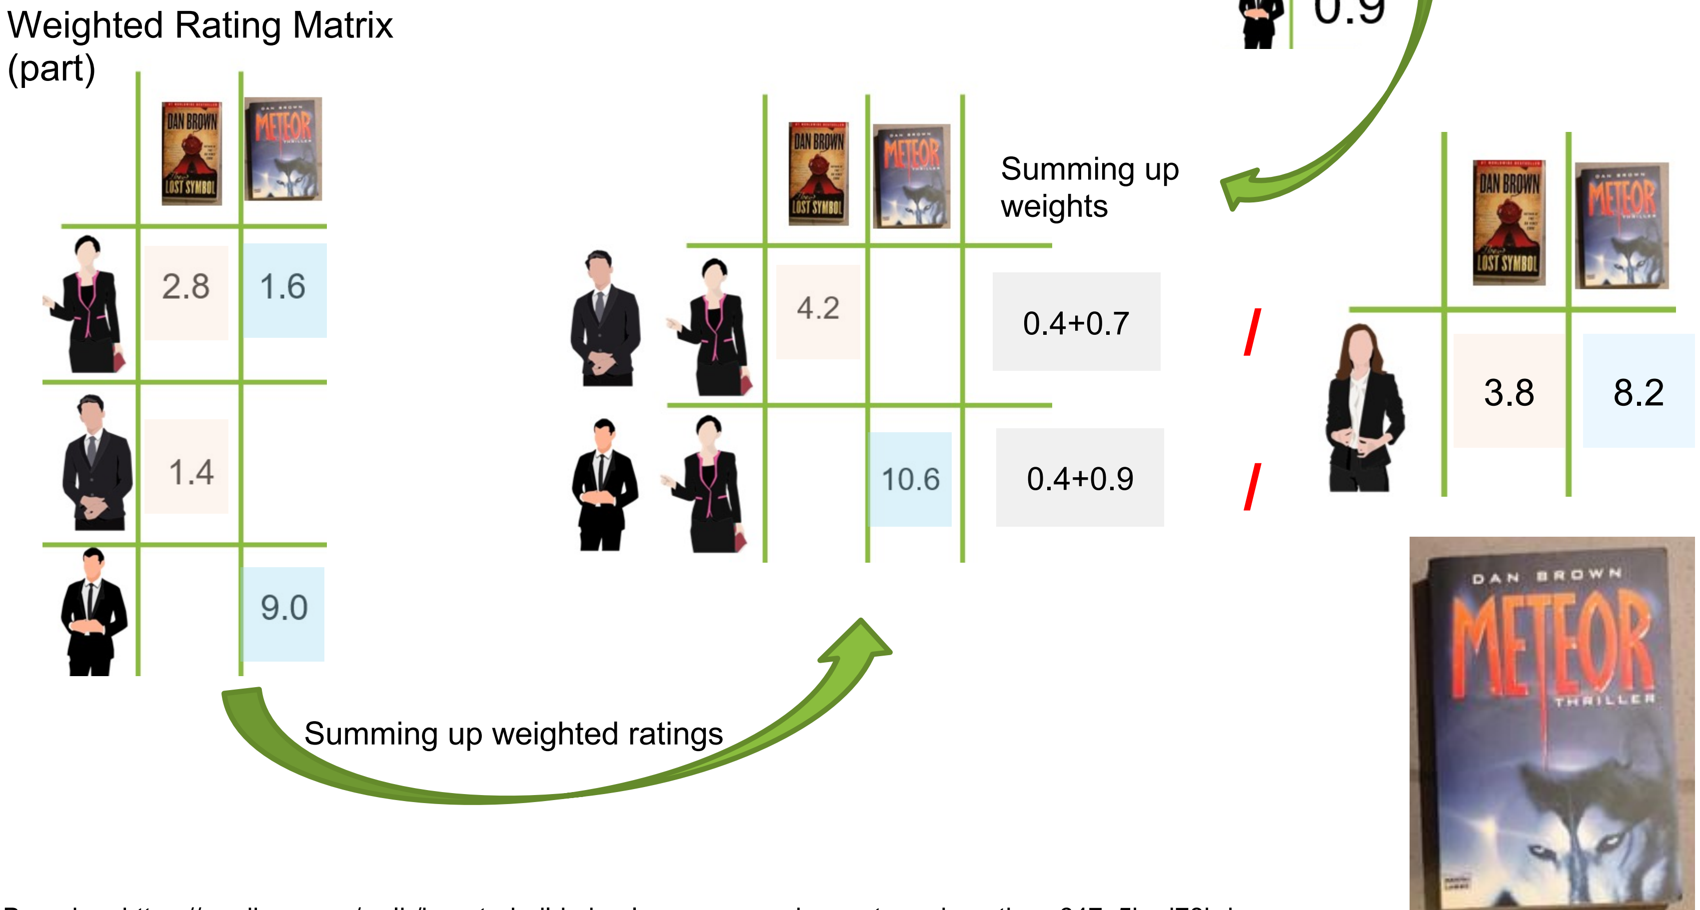

Based on https://medium.com/swlh/how-to-build-simple-recommender-systems-in-python-647e5bcd78bd


DAN BROWN
ANGELS AND DEMONS
LOST SYMBOL
DA VINCI CODE
METEOR
INFERNO
LEE CHILD
NEVER GO BACK
A WANTED MAN
KILLING FLOOR
Cari Hiaasen
SKINNY DIP
Kathy Reichs
bones to ashes
NATURE GIRL
THRILLER

Collaborative Filtering

User Based

- Calculate similarity between users
- Suggest items that similar users liked

Item Based

- Calculate similarity between items based on user ratings – no semantics are required!
- Suggest item that is similar to item the user already likes

Cold-Start Problem

What is n

- **Basic problem: as a system is created data is missing. The algorithms need initial information (e.g. to calculate similarities).**
- **What could be new? What is the problem?**
  - **User**: no information yet, no recorded interaction nor know preferences
  - **Item**: the item has not been “like” or viewed by anyone yet, no knowledge who may like it or how its rating are similar
  - **Community**: new system is created, lack of information about users as well as item
- **Solutions?**
  - Hybrid approaches (e.g. using content based filtering to get started and then move more towards collaborative filtering)
  - Use social, demographics, context, content to get started
  - Require initial interaction (e.g. questionnaire, ask for examples, …)

[1] Bobadilla, Jesús; Ortega, Fernando; Hernando, Antonio; Bernal, Jesús (February 2012). "A collaborative filtering approach to mitigate the new user cold start problem". Knowledge-Based Systems. doi:10.1016/j.knosys.2011.07.021.
[2] https://en.wikipedia.org/wiki/Cold_start_(recommender_systems)

Pinterest

Registration Process and start-up

Demographics Contextual Explicit Interest

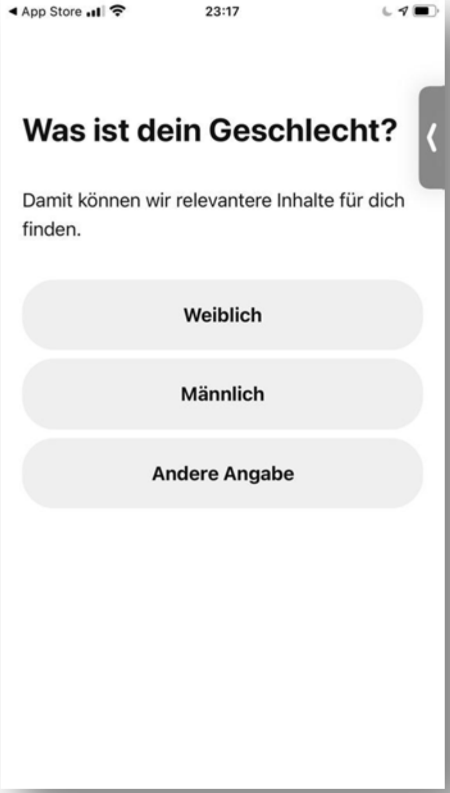

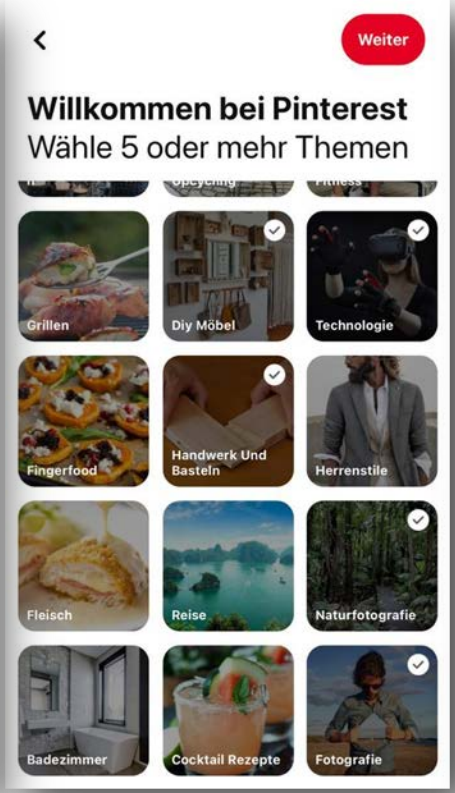

How to include new item?

Scenario: News Feed

- You create a recommender system for a news feed in social media (e.g. twitter style).
- How do you add new articles?
- Assuming you get a lot of articles….

Sparseness of Ratings

- E.g. Amazon (low estimates for illustration only)
  - Over 10 Million products
  - Over 100 Million customer

- What is the problem?

What Book do you Recommend to me?
Why does the UI matter?

Breakout Sessions 2
Recommendation and User Interface

- How can you measure the quality of the user experience of a recommender system?
- How does the screen size impact the performance of a recommender system?

Recommender Systems Summary

Content-based

- Based on the similarity of items
- How similar is a new item to an item already liked/watched/bought by the user
- Advantage
  - If similarity is known or can be calculated, no ratings/actions from the user are required
- Difficulty:
  - Information/meta-data/algorithms for calculating similarity are required

Collaborative Filtering

- suggestions are made based on users that had similar interests/actions
- How has a similar user liked this item?
- Advantage:
  - Know knowledge about the item is required, no meta data or similarity calculation of items required
- Difficulty:
  - Data about other users is required
  - “cold-start” problem

Example data for experiments

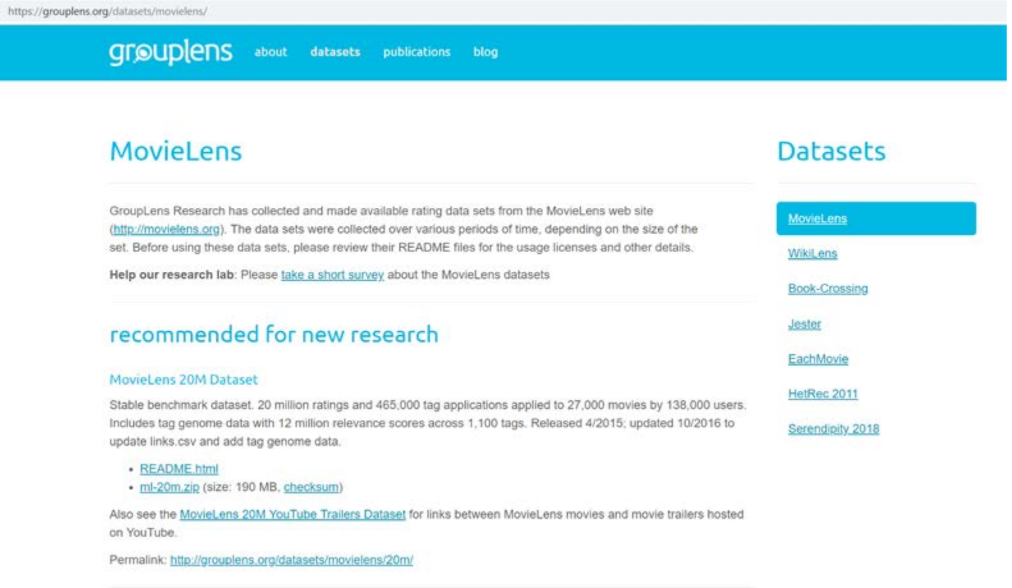

https://grouplens.org/datasets/movielens/

	A	B	C	D	E
1	userId	movieId	rating	timestamp	
2	1	1	4	964982703	
3	1	3	4	964981247	
4	1	6	4	964982224	
5	1	47	5	964983815	
6	1	50	5	964982931	
7	1	70	3	964982400	
8	1	101	5	964980868	
9	1	110	4	964982176	
10	1	151	5	964984041	
11	1	157	5	964984100	
12	1	163	5	964983650	
13	1	216	5	964981208	
14	1	223	3	964980985	
15	1	231	5	964981179	
16	1	235	4	964980908	
17	1	260	5	964981680	
18	1	296	3	964982967	
19	1	316	3	964982310	
20	1	333	5	964981179	
21	1	349	4	964982563	
22	1	356	4	964980962	
23	1	362	5	964982588	
24	1	367	4	964981710	
25	1	423	3	964982363	

Additional reading material: https://hub.packtpub.com/recommending-movies-scale-python/

Breakout Sessions 3
Recommender System

- Design a systems that suggest educational videos on YouTube that fit your lectures in this term

Algorithms for Recommender System

- Not core to Intelligent User Interfaces
- Efficient implementations in libraries
- Important to understand the algorithms to get the parameters right
- Many online resources, .e.g.
  - https://www.youtube.com/watch?v=Eeg1DEeWUjA (introductory video)
  - https://www.youtube.com/watch?v=Gf4HZpZAlDA https://www.youtube.com/watch?v=b8YylVulszQ (details on Collaborative Filtering Algorithms)

Examples for recommender systems in Python

https://www.youtube.com/watch?v=z0dx-YckFko

- https://kerpanic.wordpress.com/2018/03/26/a-gentle-guide-to-recommender-systems-with-surprise/
- https://www.analyticsvidhya.com/blog/2016/06/quick-guide-build-recommendation-engine-python/
- https://medium.com/@connectwithghosh/recommender-system-on-the-movielens-using-an-autoencoder-using-tensorflow-in-python-f13d3e8d600d

Adaptive UIs


Lights
Fridge
Microwave
Lights
Thermostat

Getting Physical (1) Initial Experience (1998)

- **Extremely simple, but still it creates a new experience**
- 2-Bit Input
- Not an input device
- Very specific function

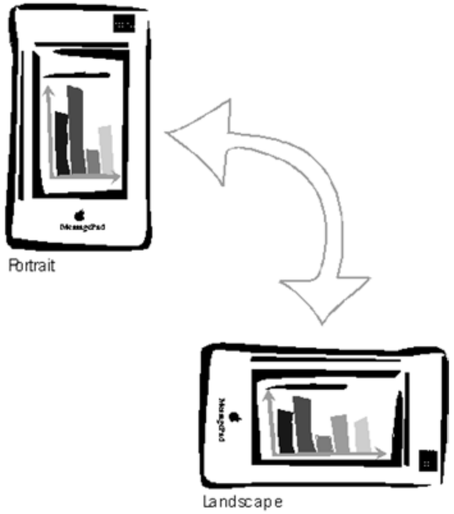

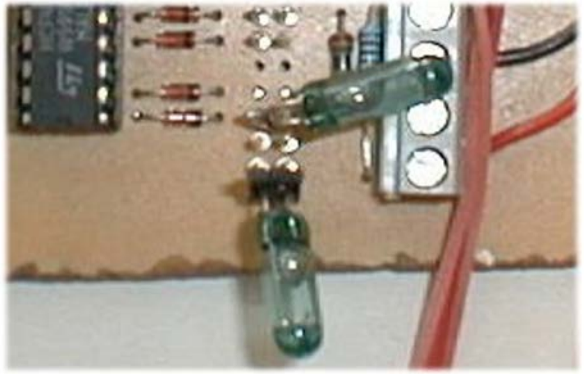

A. Schmidt, M. Beigl, H. Gellersen. There is more to context than location. Computers and Graphics, 23(6):893--901, 1999.
http://www.comp.lancs.ac.uk/~albrecht/pubs/pdf/schmidt_cug_elsevier_12-1999-context-is-more-than-location.pdf

Getting Physical (1) Initial Experience (1998)

- **Extremely simple, but still it creates a new experience**
- 2-Bit Input
- Not an input device
- Very specific function

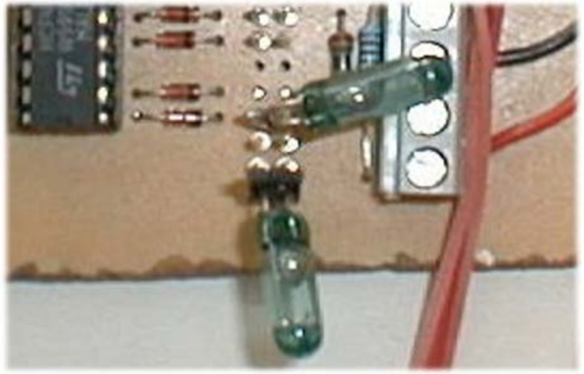

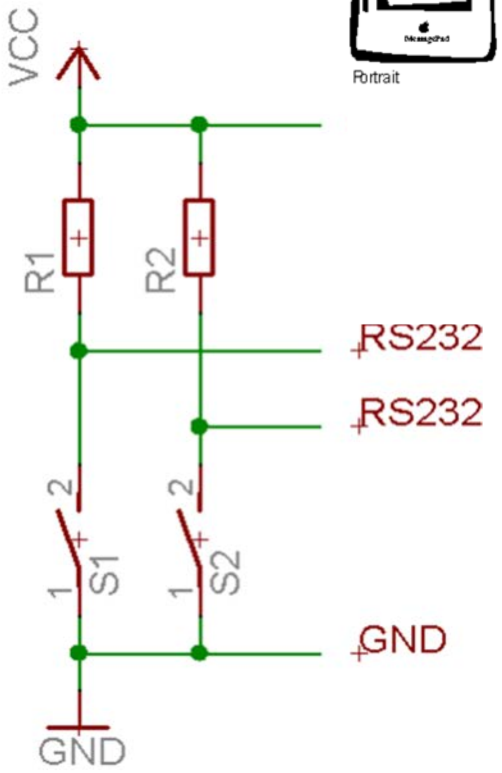

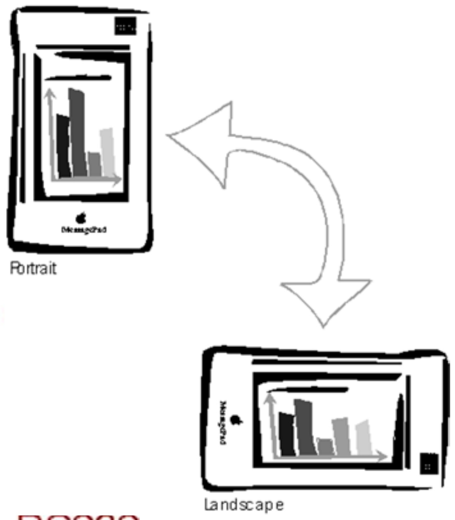

A. Schmidt, M. Beigl, H. Gellersen. There is more to context than location. Computers and Graphics, 23(6):893--901, 1999.
http://www.comp.lancs.ac.uk/~albrecht/pubs/pdf/schmidt_cug_elsevier_12-1999-context-is-more-than-location.pdf

Knife that “knows” what its cuts

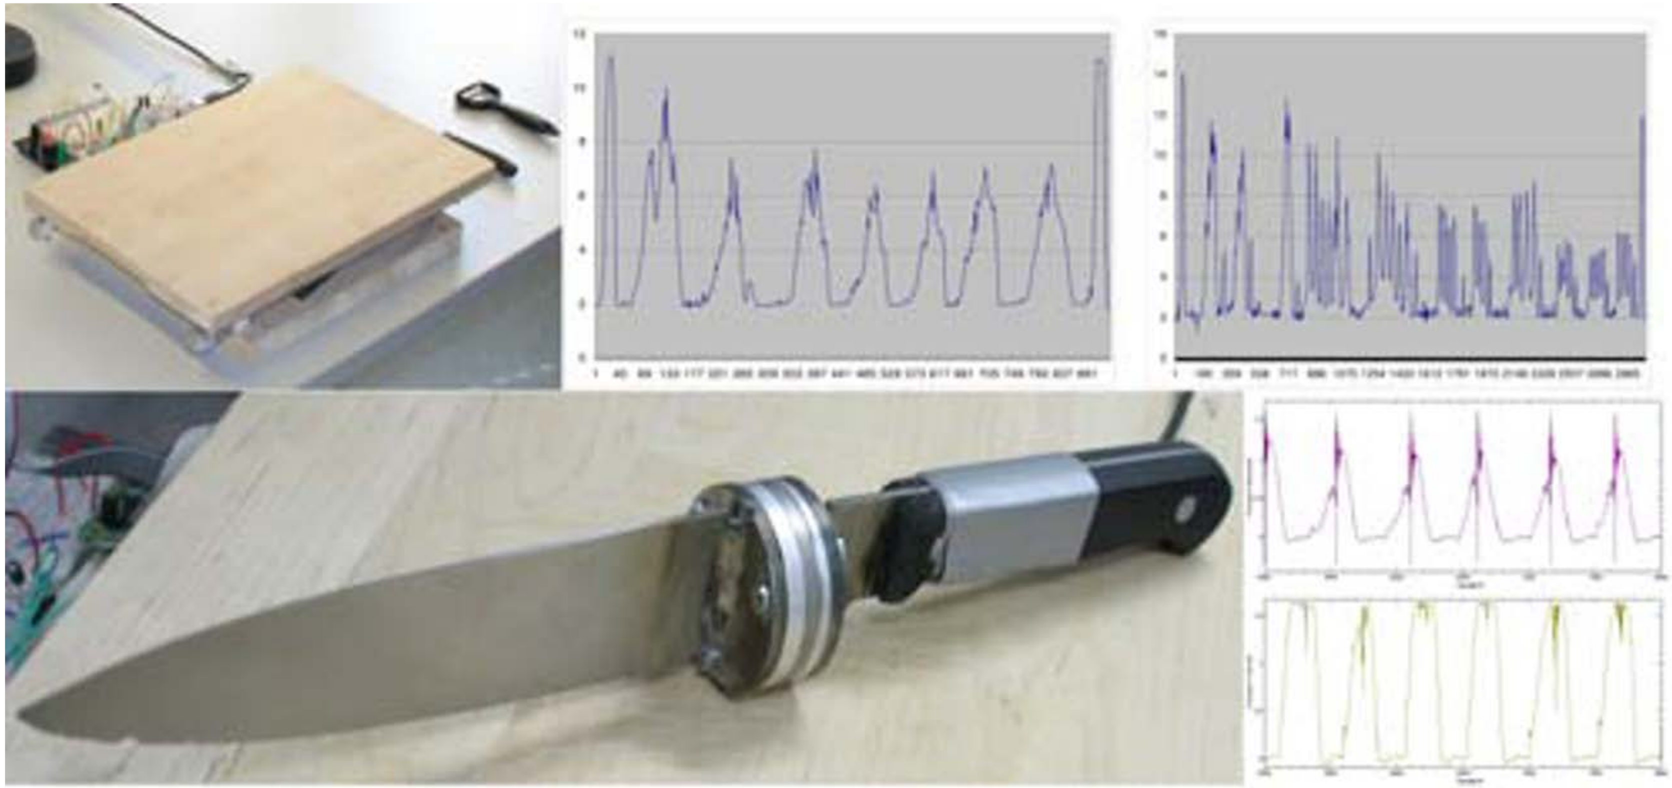

Matthias Kranz, Albrecht Schmidt, Alexis Maldonado, Radu Bogdan Rusu, Michael Beetz, Benedikt Hörnler, and Gerhard Rigoll. 2007. Context-aware kitchen utilities. In Proceedings of the 1st international conference on Tangible and embedded interaction (TEI '07). Association for Computing Machinery, New York, NY, USA, 213–214. DOI: https://doi.org/10.1145/1226969.1227013

Radio Tags for Activity Sensing

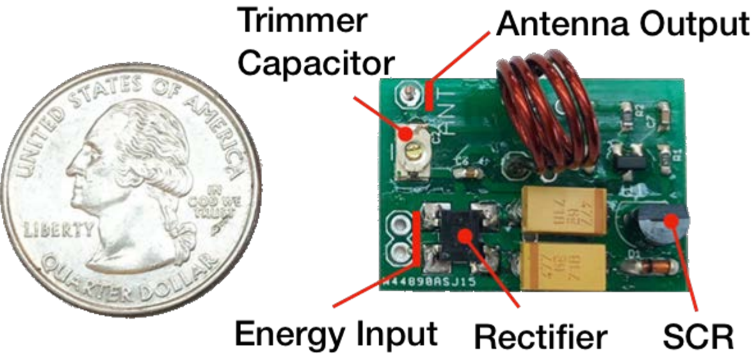

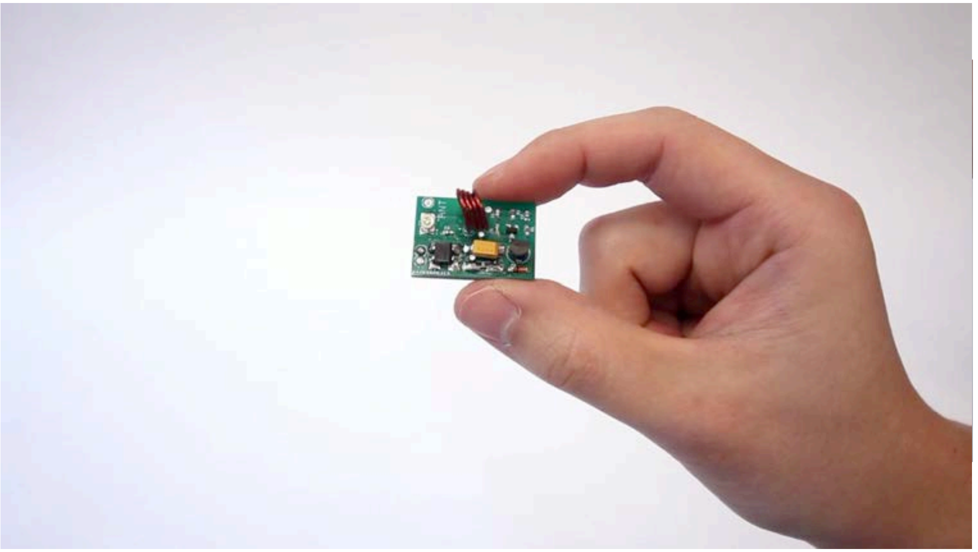

Yang Zhang, Yasha Iravantchi, Haojian Jin, Swarun Kumar, and Chris Harrison. 2019. Sozu: Self-Powered Radio Tags for Building-Scale Activity Sensing. In Proceedings of the 32nd Annual ACM Symposium on User Interface Software and Technology (UIST '19). ACM, New York, NY, USA, 973–985. DOI: https://doi.org/10.1145/3332165.3347952

Video: https://youtu.be/wbq-eOOIPyw

Vibrometry for Environment Sensing

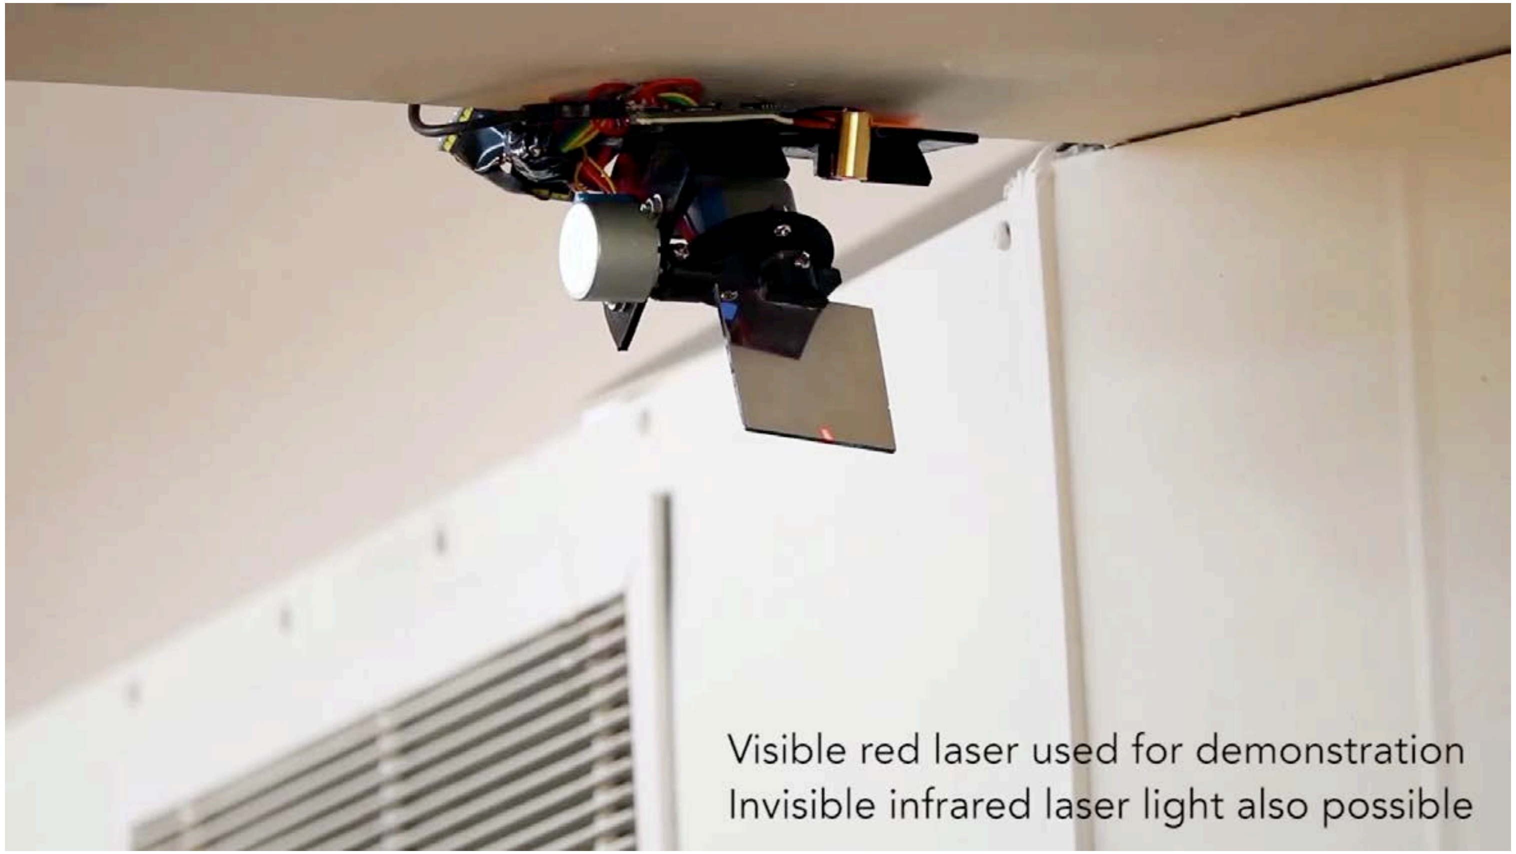

https://youtu.be/51XaZDki6yg

Yang Zhang, Gierad Laput, and Chris Harrison. 2018. Vibrosight: Long-Range Vibrometry for Smart Environment Sensing. In Proceedings of the 31st Annual ACM Symposium on User Interface Software and Technology (UIST '18). ACM, New York, NY, USA, 225–236. DOI: https://doi.org/10.1145/3242587.3242608

Extracting Contextual Information

- Users Location
  - GPS
  - Direction of Voice
- Users Activity
- Users Emotion
- Users Pose
- Objects Surrounding the User
- Status of Objects

Smart Environments - Direction-of-Voice

- Feature Extraction e.g. volume, speech frequency ratio
- Machine Learning e.g. ensemble-based decision trees

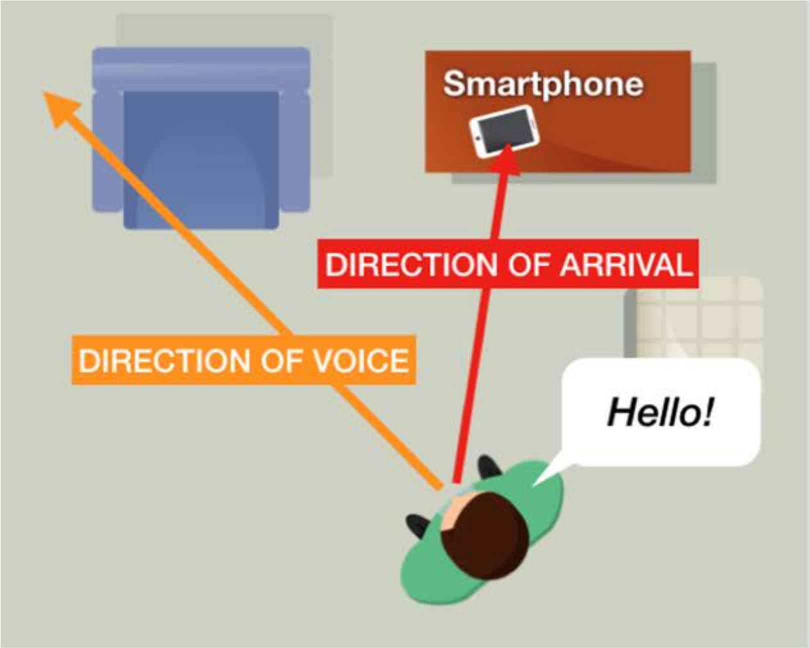

Karan Ahuja, Andy Kong, Mayank Goel, and Chris Harrison. 2020. Direction-of-Voice (DoV) Estimation for Intuitive Speech Interaction with Smart Devices Ecosystems. In Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20). ACM, New York, NY, USA, 1121–1131. DOI: https://doi.org/10.1145/3379337.3415588

Smart Environments - Direction-of-Voice

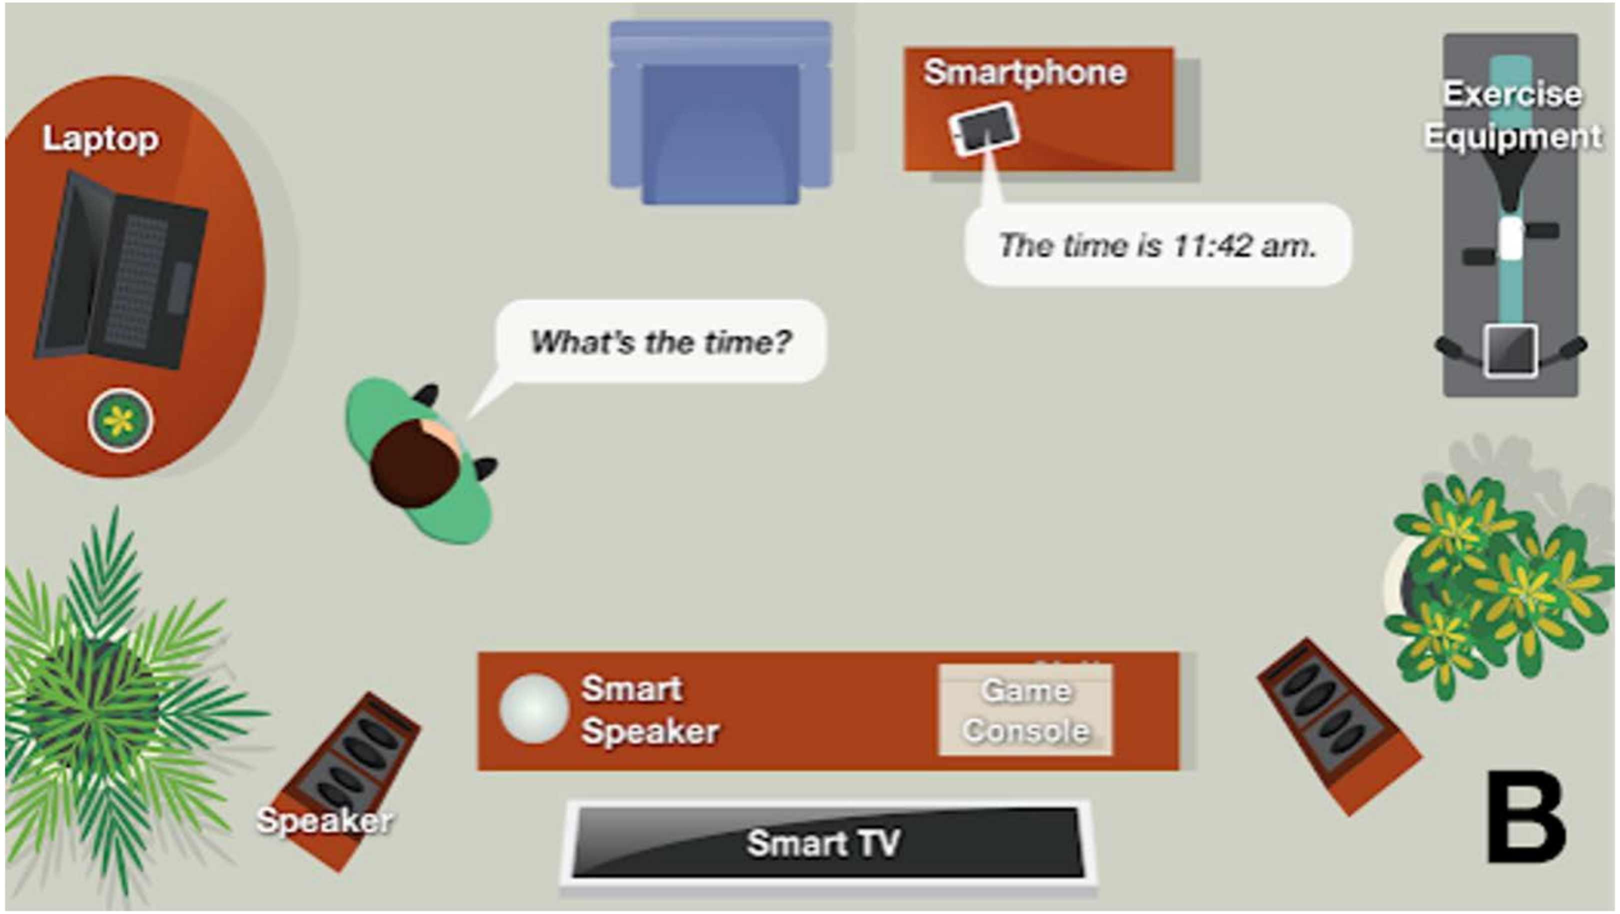

Karan Ahuja, Andy Kong, Mayank Goel, and Chris Harrison. 2020. Direction-of-Voice (DoV) Estimation for Intuitive Speech Interaction with Smart Devices Ecosystems. In Proceedings of the 33rd Annual ACM Symposium on User Interface Software and Technology (UIST '20). ACM, New York, NY, USA, 1121–1131. DOI: https://doi.org/10.1145/3379337.3415588

Context-Aware Keyboards

Visible keyboard **Collect touches** **Adapt underlying key regions**

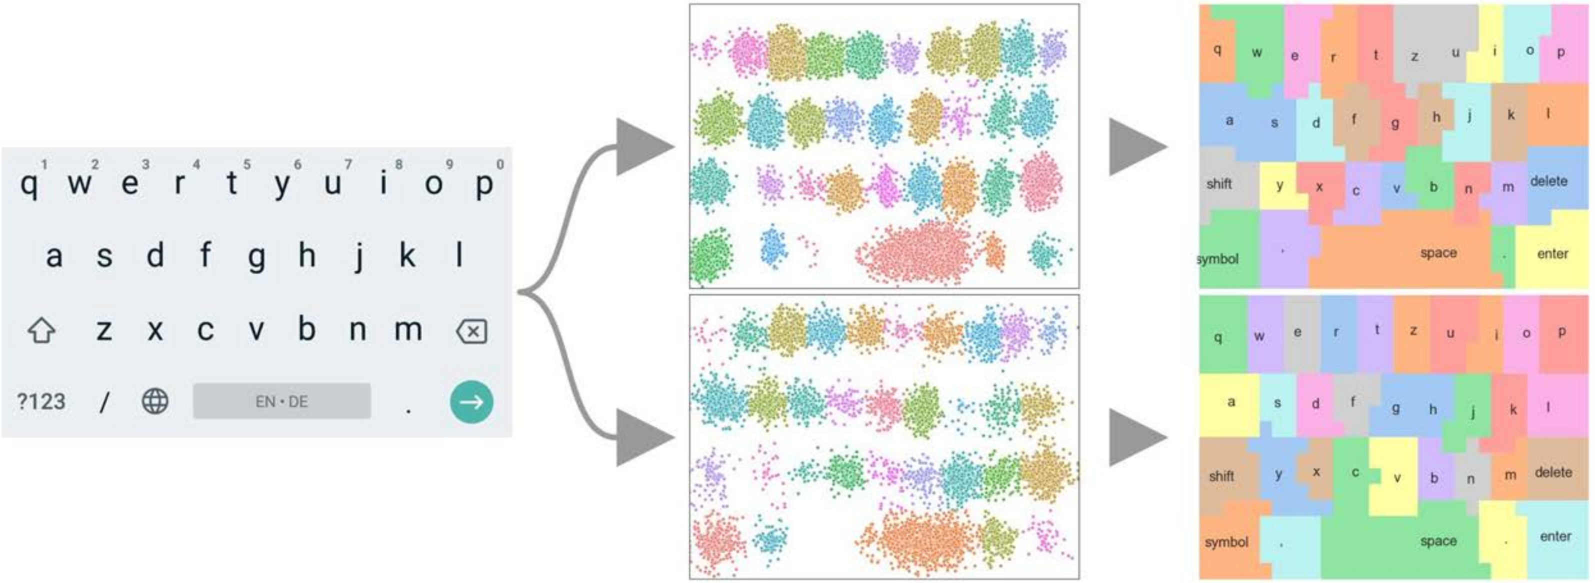

- Palm detection – input rejection
- Finger identification, Finger Orientation – model improvement
- Grip detection – model selection

Palm Detection

- Convolutional Neural Network
- Classification
- Representation Learning

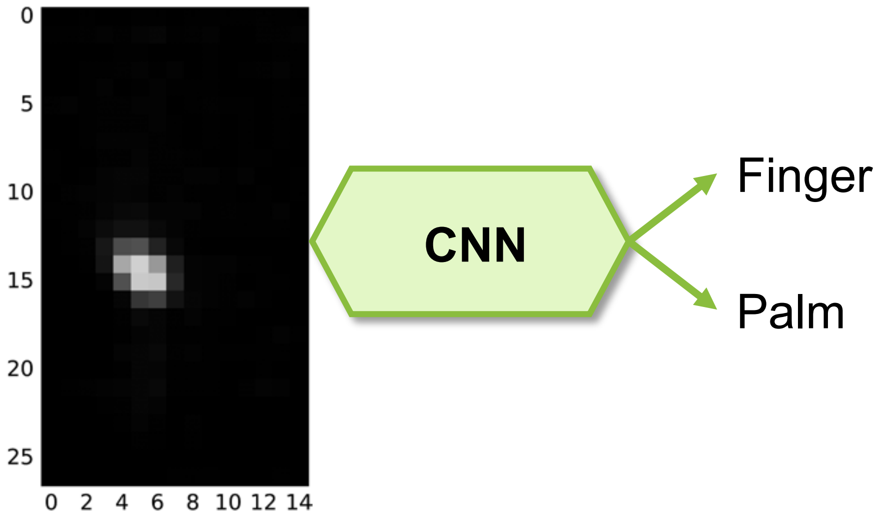

Huy Viet Le, Thomas Kosch, Patrick Bader, Sven Mayer, and Niels Henze. 2018. PalmTouch: Using the Palm as an Additional Input Modality on Commodity Smartphones. In Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems (CHI '18). ACM, New York, NY, USA, Paper 360, 1–13. DOI: https://doi.org/10.1145/3173574.3173934

Finger Orientation

- Convolutional Neural Network
- Regression
- Representation Learning

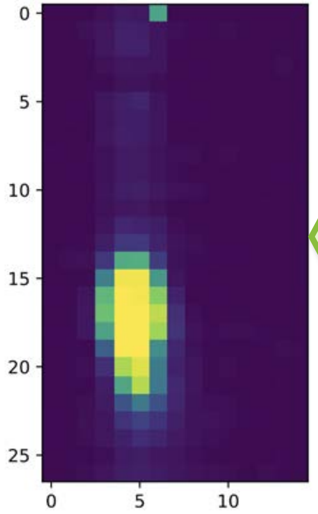

CNN

Pitch: 4°

Yaw: 17°

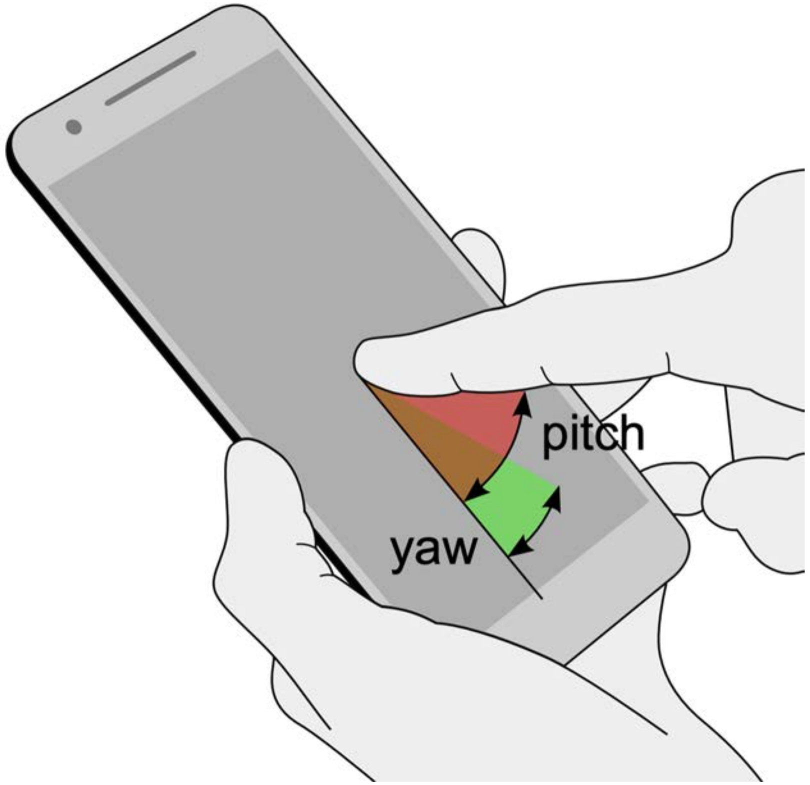

Sven Mayer, Huy Viet Le, and Niels Henze. 2017. Estimating the Finger Orientation on Capacitive Touchscreens Using Convolutional Neural Networks. In Proceedings of the 2017 ACM International Conference on Interactive Surfaces and Spaces (ISS '17). ACM, New York, NY, USA, 220–229. DOI: https://doi.org/10.1145/3132272.3134130

Activity Recognition – Accelerometers

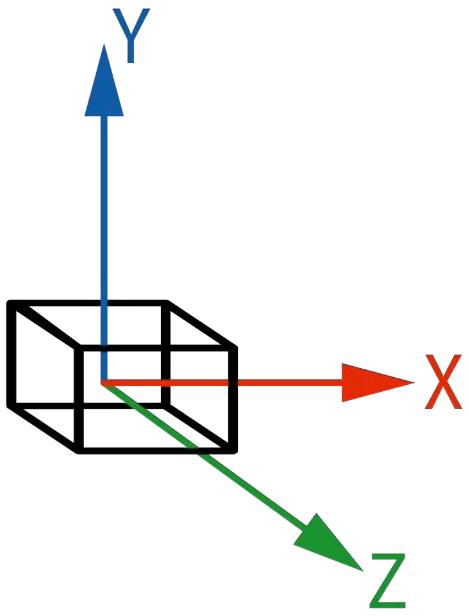

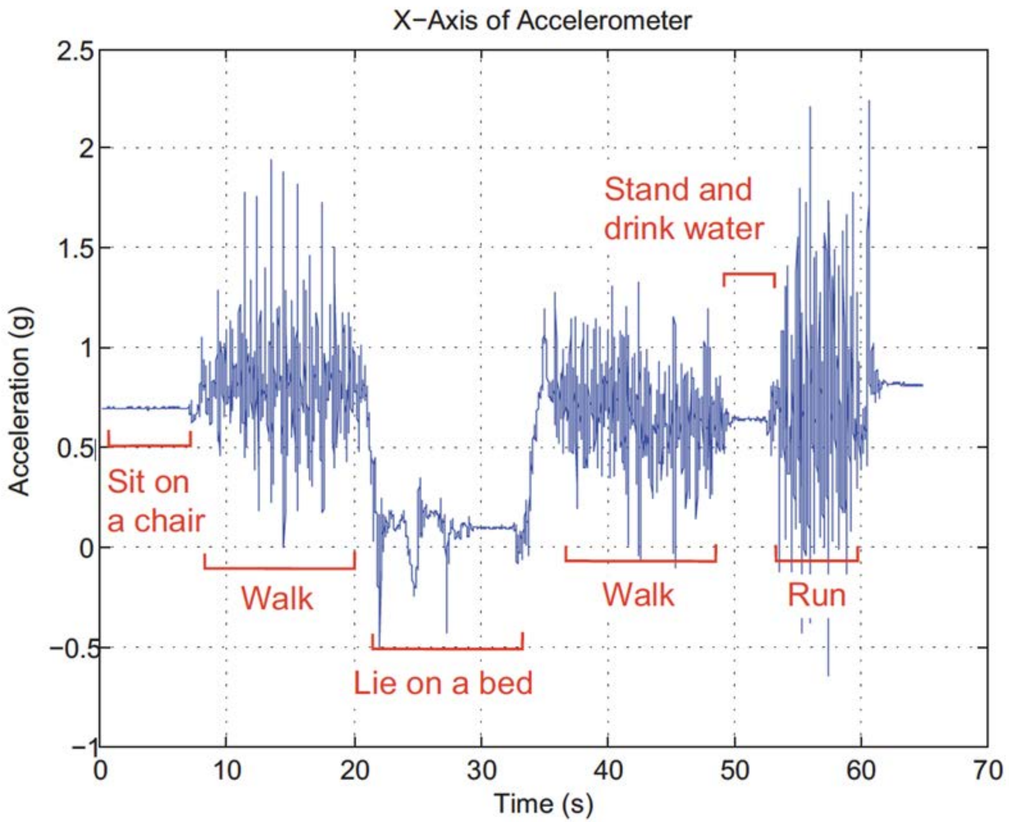

Mi Zhang and Alexander A. Sawchuk. 2012. USC-HAD: a daily activity dataset for ubiquitous activity recognition using wearable sensors. In Proceedings of the 2012 ACM Conference on Ubiquitous Computing (UbiComp '12). Association for Computing Machinery, New York, NY, USA, 1036–1043. DOI: https://doi.org/10.1145/2370216.2370438

Neuronal Network With Timeseries Data

Recurrent Neural Network

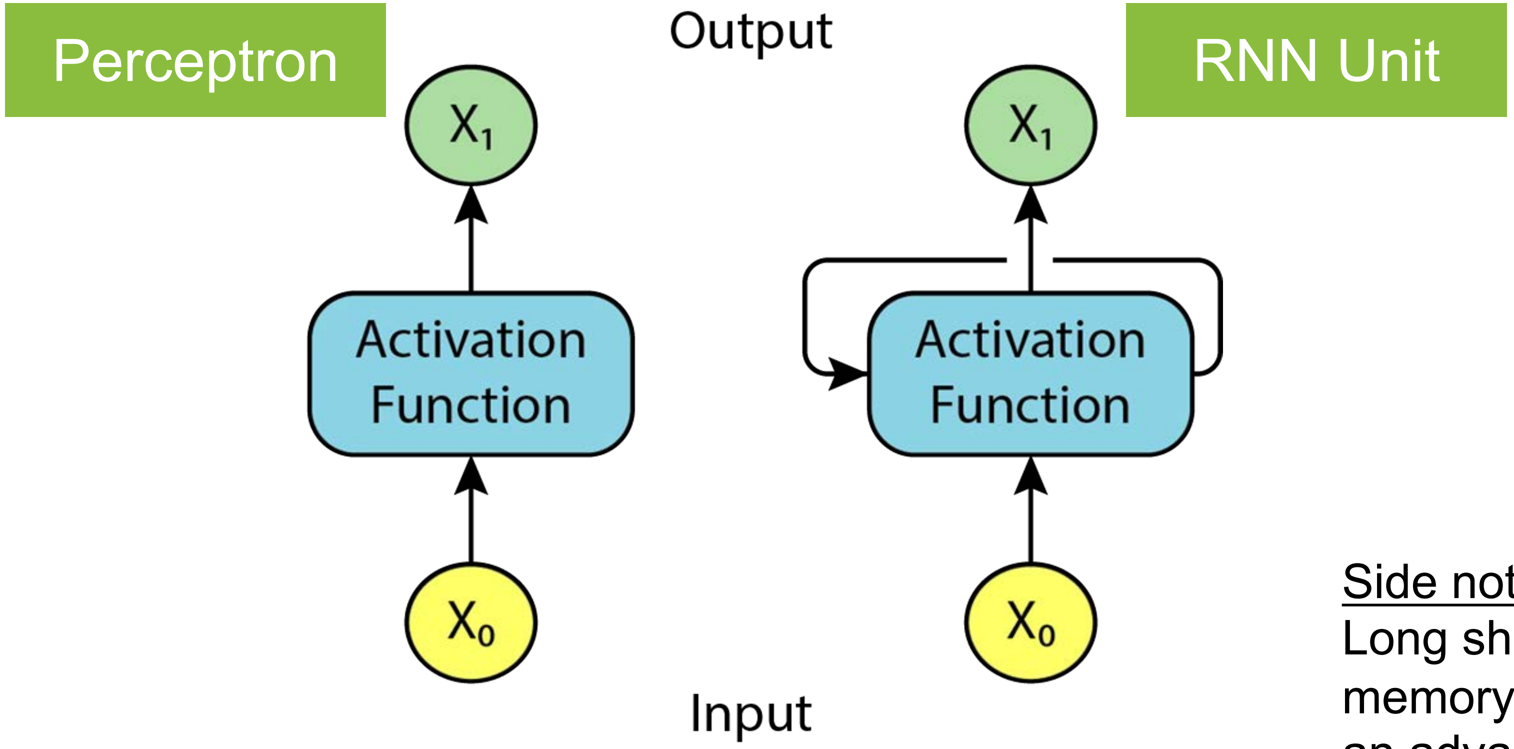

<u>Side note:</u>
Long short-term memory (**LSTM**) is an advanced RNN

In depth LSTM tutorial: http://colah.github.io/posts/2015-08-Understanding-LSTMs/
Code Examples: https://github.com/cwi-dis/mobile-har-tutorial

Human Pose Detection

- Keypoint Estimation
- Multi-stage CNN

Zhe Cao, Gines Hidalgo, Tomas Simon, Shih-En Wei, and Yaser Sheikh. OpenPose: realtime multi-person 2D pose estimation using Part Affinity Fields. IEEE transactions on pattern analysis and machine intelligence 43, no. 1 (2019): 172-186. DOI: https://doi.org/10.1109/TPAMI.2019.2929257

Camera based Gaze Estimation

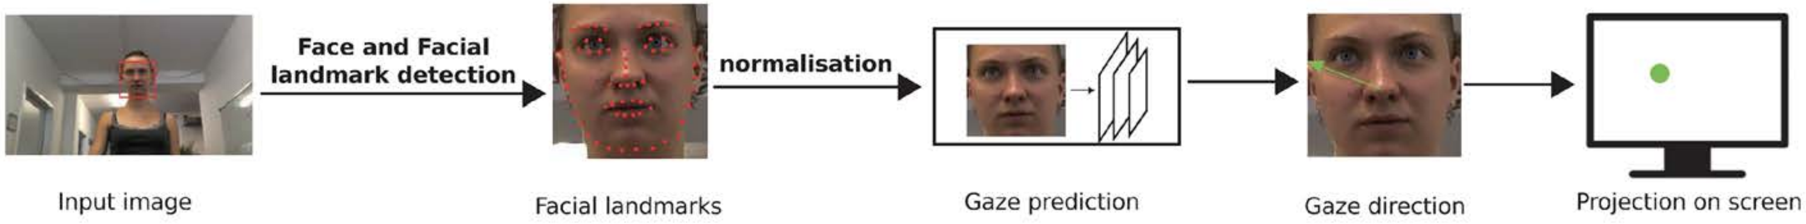

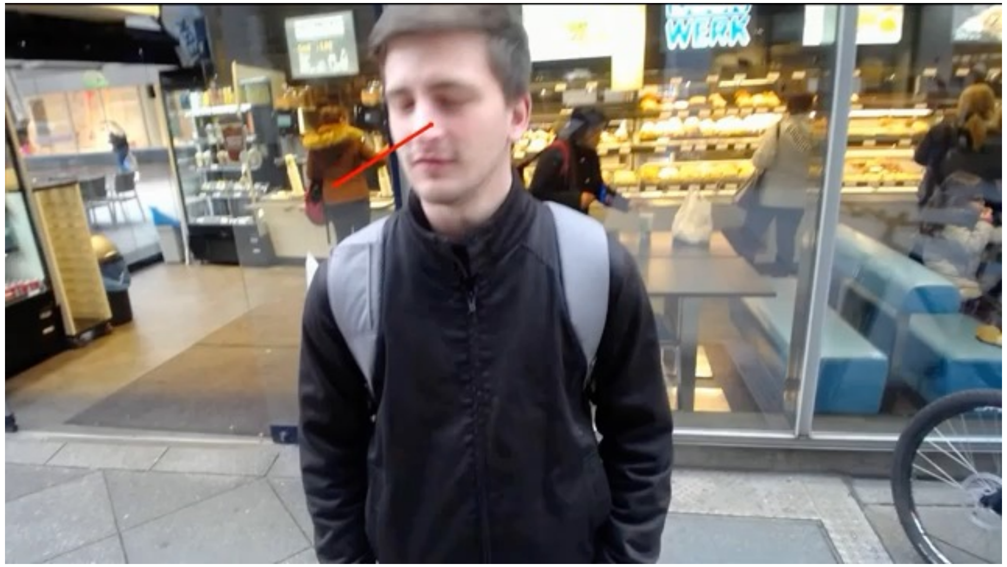

https://youtu.be/9Lujg3beiYI

Xucong Zhang, Yusuke Sugano, and Andreas Bulling. 2019. Evaluation of Appearance-Based Methods and Implications for Gaze-Based Applications. In Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems (CHI '19). ACM, New York, NY, USA, Paper 416, 1–13. DOI: https://doi.org/10.1145/3290605.3300646 URL: http://www.opengaze.org/

Camera based Gaze Estimation

- Convolutional Neural Network
- Regression – x/y cornindates

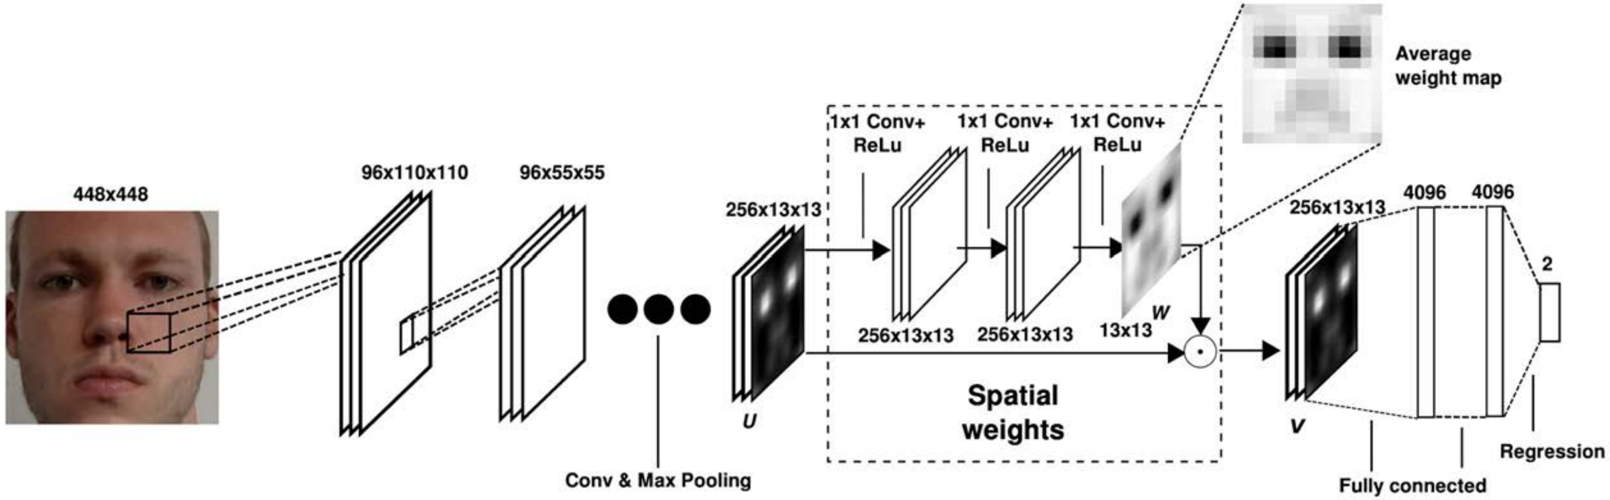

Xucong Zhang, Yusuke Sugano, Mario Fritz, and Andreas Bulling. 2017. It's written all over your face: Full-face appearance-based gaze estimation. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops, pp. 51-60. 2017. DOI: https://doi.org/10.1109/CVPRW.2017.284

Enhanced Voice Assistants

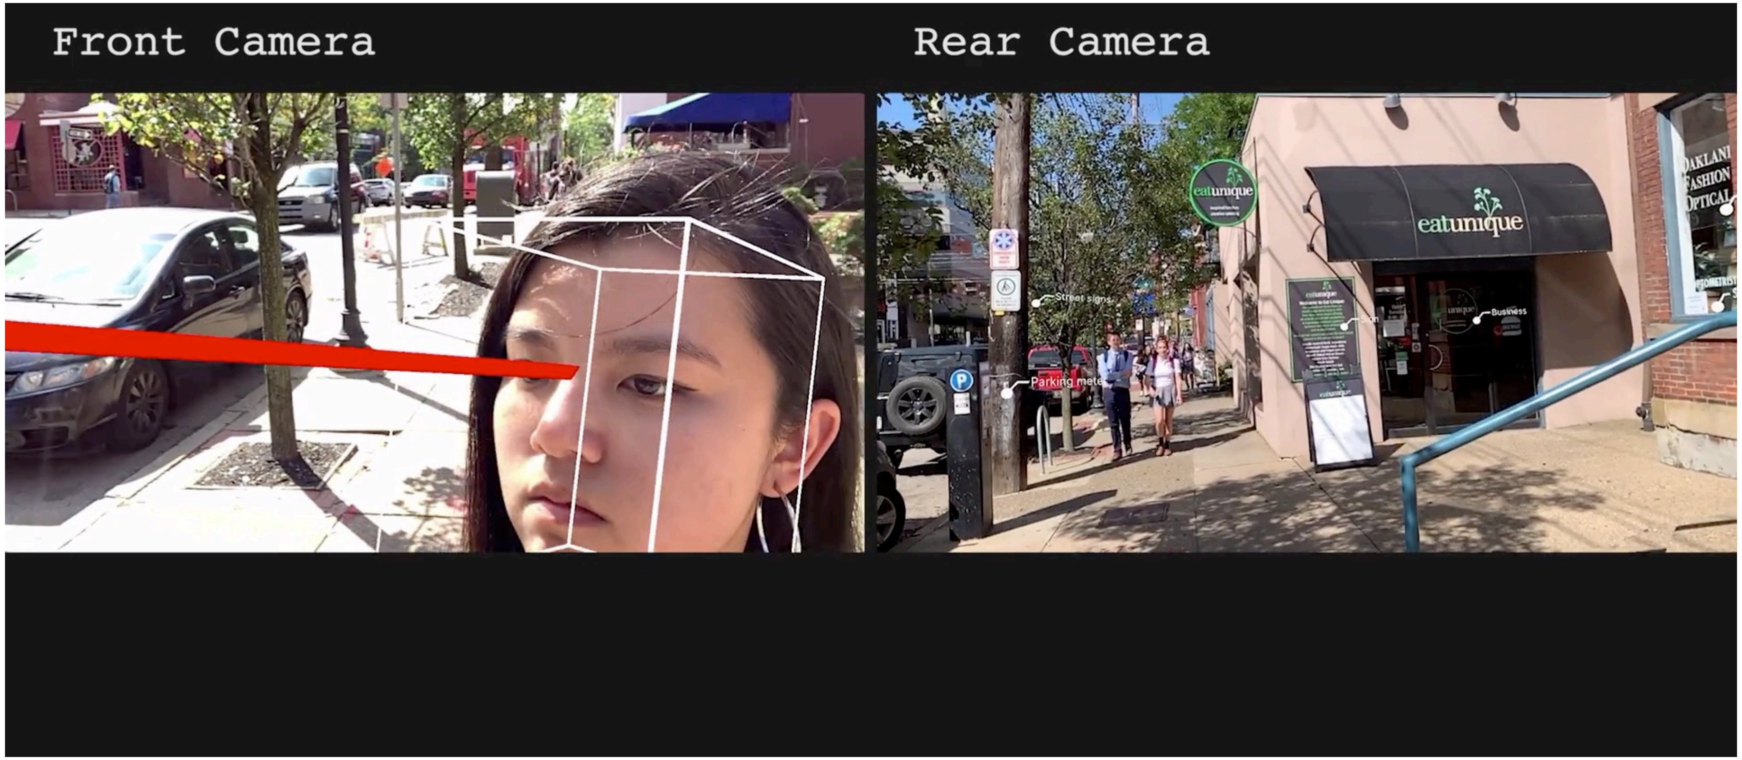

Sven Mayer, Gierad Laput, and Chris Harrison. 2020. Enhancing Mobile Voice Assistants with WorldGaze. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–10. DOI: https://doi.org/10.1145/3313831.3376479

Object Detection

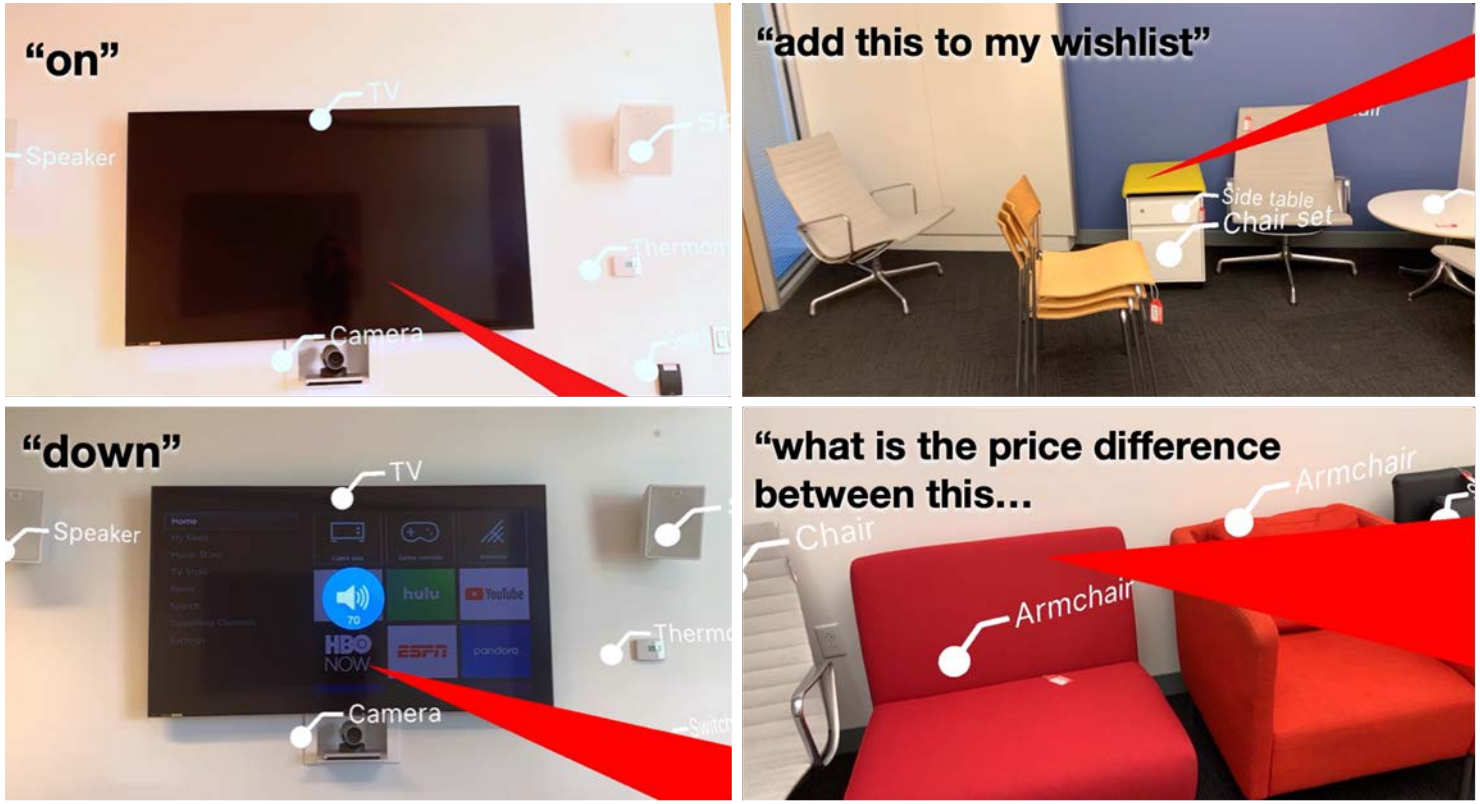

Sven Mayer, Gierad Laput, and Chris Harrison. 2020. Enhancing Mobile Voice Assistants with WorldGaze. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (CHI '20). Association for Computing Machinery, New York, NY, USA, 1–10. DOI: https://doi.org/10.1145/3313831.3376479

Object Detection

R-CNN: Regions with CNN feature

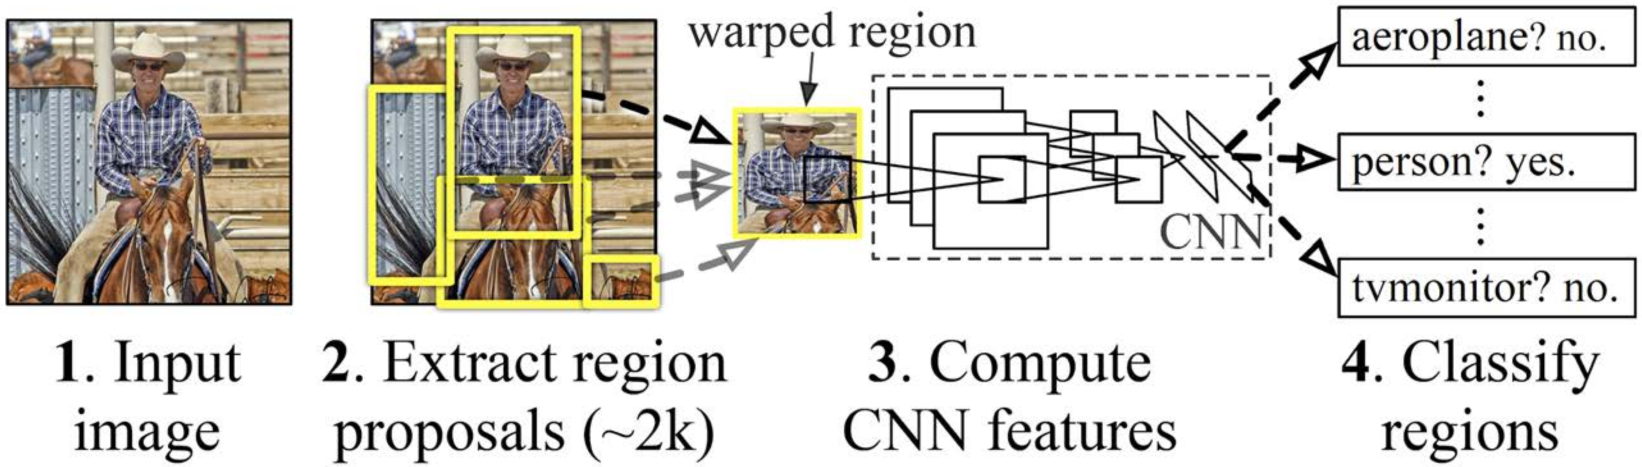

Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation (2014) In Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 580-587. DOI: https://doi.org/10.1109/CVPR.2014.81 Source Code: https://github.com/rbgirshick/rcnn

Object Detection

YOLO – You Only Look Once

https://youtu.be/MPU2Histivl

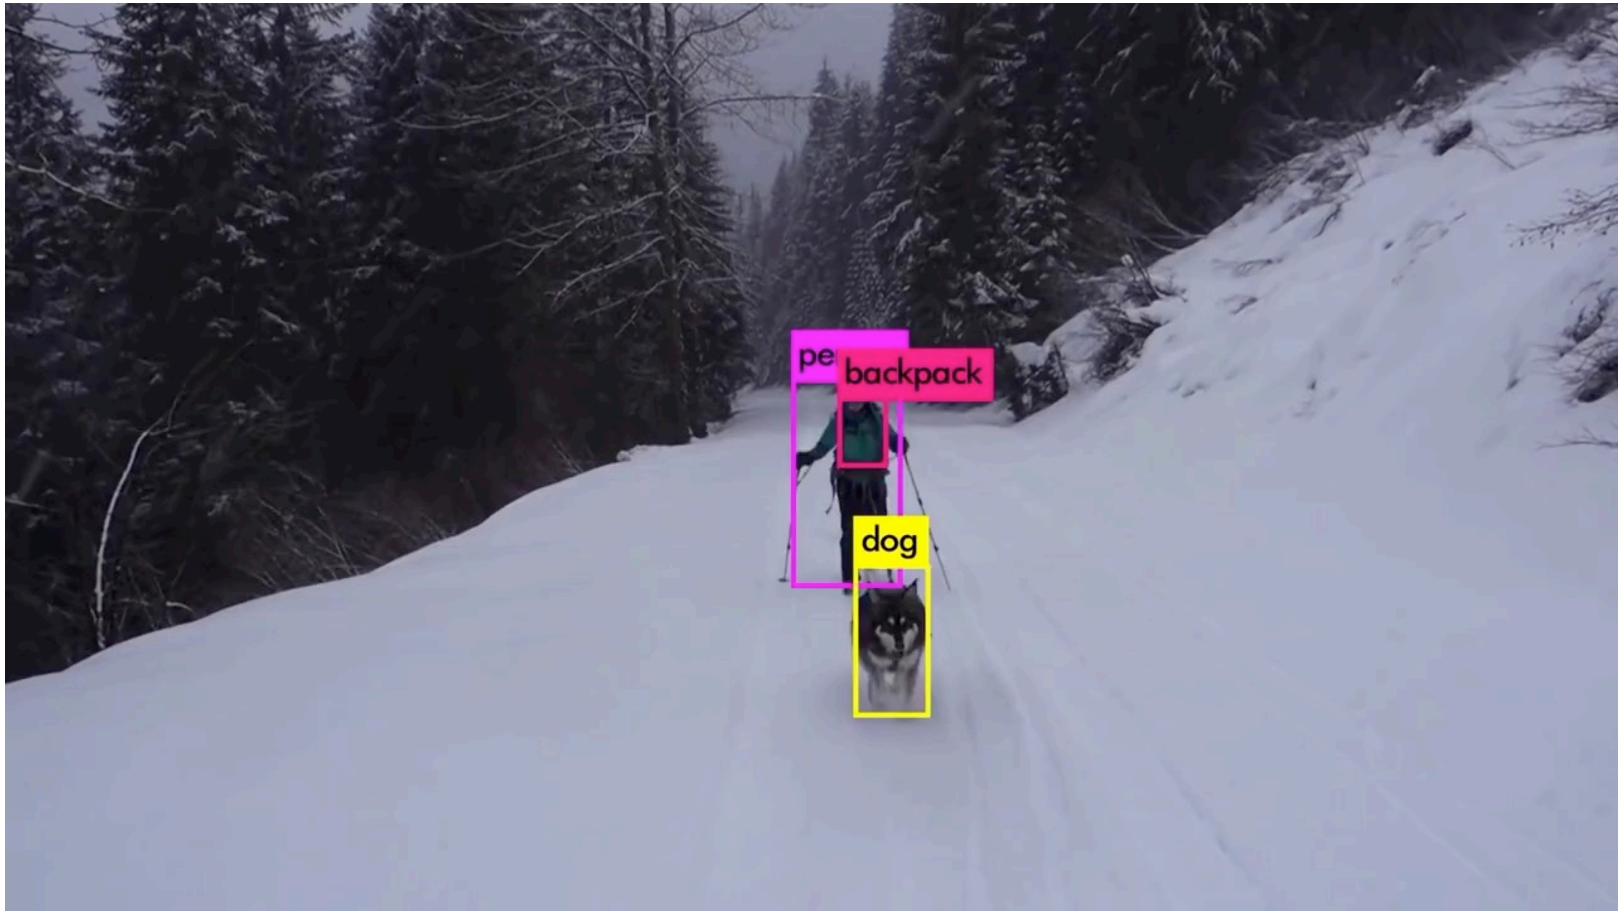

Joseph Redmon, and Ali Farhadi. "Yolov3: An incremental improvement." arXiv preprint arXiv:1804.02767 (2018).
URL: https://pjreddie.com/darknet/yolo/

Machine Learning for HCI

User Centered Design Cycle ISO 9241

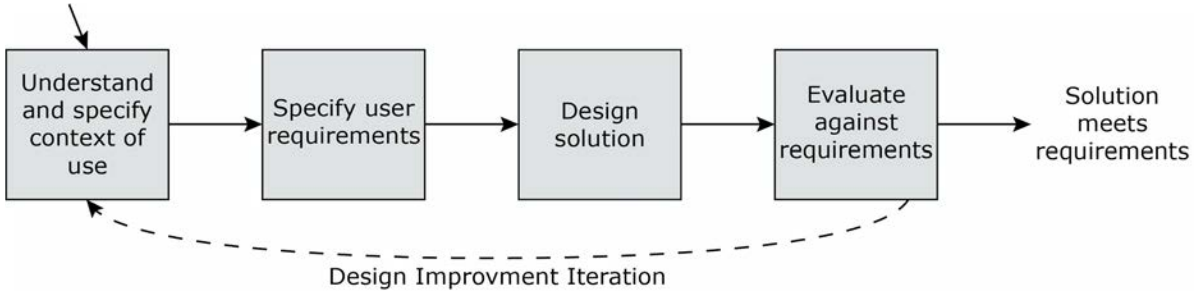

Adaptation for Machine Learning

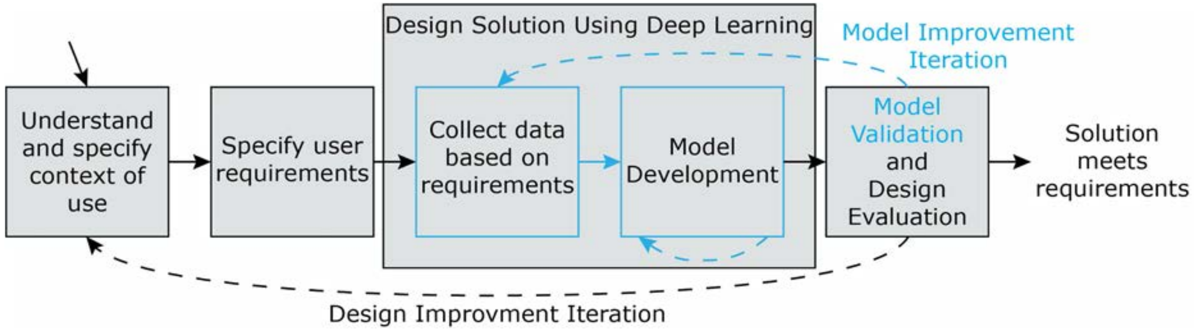

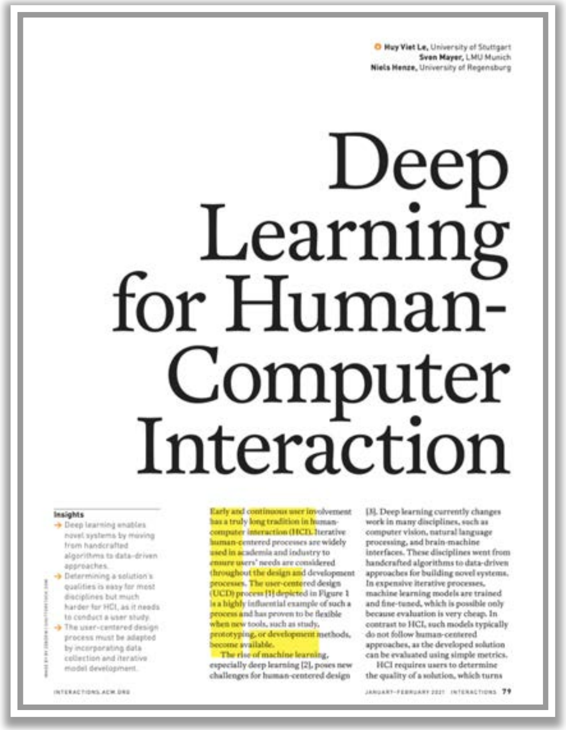

Huy Viet Le, University of Stuttgart
Sven Mayer, LMU Munich
Niels Henze, University of Regensburg

Deep Learning for Human-Computer Interaction

Insights

- Deep learning enables novel systems by moving from handcrafted algorithms to data-driven approaches.
- Determining a solution's qualities is easy for most disciplines but much harder for HCI, as it needs to conduct a user study.
- The user-centered design process must be adapted by incorporating data collection and iterative model development.

Early and continuous user involvement has a truly long tradition in human-computer interaction (HCI). Iterative human-centered processes are widely used in academia and industry to ensure users' needs are considered throughout the design and development processes. The user-centered design (UCD) process [1] depicted in Figure 1 is a highly influential example of such a process and has proven to be flexible when new tools, such as study, prototyping, or development methods, become available.

The rise of machine learning, especially deep learning [2], poses new challenges for human-centered design [3]. Deep learning currently changes work in many disciplines, such as computer vision, natural language processing, and brain-machine interfaces. These disciplines went from handcrafted algorithms to data-driven approaches for building novel systems. In expensive iterative processes, machine learning models are trained and fine-tuned, which is possible only because evaluation is very cheap. In contrast to HCI, such models typically do not follow human-centered approaches, as the developed solution can be evaluated using simple metrics.

HCI requires users to determine the quality of a solution, which turns

INTERACTIONS.ACM.ORG JANUARY–FEBRUARY 2021 INTERACTIONS 79

Huy Viet Le, Sven Mayer, and Niels Henze. 2020. Deep learning for human-computer interaction. interactions 28, 1 (January - February 2021), 78–82. DOI: https://doi.org/10.1145/3436958

Neuronal Network Concepts

Introduced concepts

- Classification vs. Regression
- Feature Extraction vs. Representation Learning
- Model Structures
  - Neuronal Network
  - Convolutional Neural Network
  - RNN
  - Long short-term memory (LSTM)

License