

A futuristic virtual reality environment with a person wearing a VR headset interacting with glowing digital elements. The scene is filled with blue and purple light rays, floating rectangular frames, and abstract shapes. On the left, there are some data-like numbers: 'COMBO 184' and '1 336 468'.

# Unique Identification of 50,000+ Virtual Reality Users from Head & Hand Motion Data

Vivek Nair · Wenbo Guo · Justus Mattern · Rui Wang · James F. O'Brien · Louis Rosenberg · Dawn Song  
<https://doi.org/10.48550/arXiv.2302.08927>

# Contributors



**Vivek C. Nair**  
*UC Berkeley*



**Wenbo Guo**  
*UC Berkeley*



**Justus Mattern**  
*RWTH Aachen*



**Rui Wang**  
*UC Berkeley*



**James F. O'Brien**  
*UC Berkeley*




**Louis Rosenberg**  
*Unanimous AI*



**Dawn Song**  
*UC Berkeley*

# Acknowledgments



\* Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of their employer or the supporting entities.

# Identity in motion

Bulletin of the Psychonomic Society  
1977, Vol. 9 (2), 353-356

## Recognizing friends by their walk: Gait perception without familiarity cues

JAMES E. CUTTING and LYNN T. KOZLOWSKI  
Wesleyan University, Middletown, Connecticut 06457

Viewers can recognize themselves and others in an abstract display of their movements. Light sources mounted on joints prominent during the act of walking are sufficient cues for identification. No other information, no feedback, and little practice with such a display are needed. This procedure, developed by Johansson, holds promise for inquiry into the dimensions and features of event perception: It is both naturalistic and experimentally manageable.

People often believe that they can recognize friends by their walk. Unfortunately, this belief and the previous research on the topic (e.g., Wolff, 1943) are confounded by familiarity cues, size and shape cues, or other nongait sources of information such as probabilities of seeing a person at a given place or time. We demonstrate that viewers can recognize themselves and others in a dynamic display of their movements when these factors are controlled.

We were stimulated by the work of Johansson (1973, 1975), particularly his films (Mass & Johansson, 1971a, b). When viewing them, one sees people stripped of familiarity cues such as clothing and hairstyle; people are presented as arrays of point-light sources moving across a screen in an orderly fashion. Johansson's technique seemed to be ideal for the study of how ecological events are perceived.

A partial taxonomy of events has been proposed by Shaw, McIntyre, and Mace (1974). Some relevant distinctions are those of (1) fast vs slow events, where the critical feature is whether dynamic change can be perceived directly or only inferred, (2) reversible vs irreversible events, (3) rigid vs plastic events, and (4) events associated with animate vs inanimate sources. Most psychologists have concentrated on the perception of fast, reversible, rigid, inanimate events (e.g., Robinson & van Holst, 1975; Johansson & Jansson, 1968). A few, however, have begun to study slow, irreversible, elastic, animate events, such as the aging of faces (Pittenger & Shaw, 1975a, b). Walking is an intermediate type of event: It is fast, animate, irreversible, and also rigid—that is, composed of a hierarchy of rigid pendular motions.

Gibson (1950) has argued that the perception of any moving shape can be thought of as the perception of


formless invariant relations displayed over time. The study of gait or any other system of events should consider the interrelation of two component invariants: the underlying dynamic aspect of the event, or the transformational invariant, and the underlying unity of the structures involved, or the structural invariant (Pittenger & Shaw, 1975a; Shaw & McIntyre, 1974). In the present paper we observe whether a particular aspect of the structural invariant (the identity of the walker) is sufficiently presented through the transformational invariant (walking) for recognition.

### METHOD

Our study of gait used glass-head retroreflective tape wrapped around walker's joints, video-tape recording equipment, and bright lights focused on the walking area and mounted close to the lens of the television camera. The contrast of the image on the television monitor was raised to maximum, and the brightness to minimum, so that only the reflectant points could be seen (see Johansson, 1973). Static approximations to our stimuli can be seen in Figure 1. Figure 2 shows one of our walkers with the image brightness turned up.

Six Wesleyan University undergraduates, three males and three females, served as walkers. Each had a normal gait. They were approximately the same height and weight, and they lived together in university housing. All wore tight-fitting dark clothing during the recording session. We wrapped 5-cm-wide commercially available reflectant tape around their wrists, around their arms just above the elbow, around their ankles, and around their legs just above the knee. We affixed 5 x 18.5 cm patches to their belts at the hip and to their shoulders at equal distance from the shoulder and belt on the upper arm. No patch was placed on the head. Each individual walked at a normal pace for several minutes until we were satisfied that he or she was not "preforming" before the camera. We then recorded side views of each as he or she walked in front of the camera 8 m from the lens. Each individual walked back and forth 10 times, while his or her friends waited in another room. Individuals were on camera for five trials (six trials) and a mean of 2.7 sec during each pass across the viewing field. The camera was fixed and did not pan to follow the walker.

A test tape consisting of all tokens of all walkers was created by recording onto a second video tape. We used two helical-scan recorders, a monitor connected to one recorder (on which the source tape was played), and a television camera focused on the monitor at close range and connected to the second recorder (on which the test tape was recorded). Each token was selected in random order and recorded onto the test






38% accuracy,  $p < .005$


Cutting, J. E., & Kozlowski, L. T. (1977). Recognizing friends by their walk: Gait perception without familiarity cues. *Bulletin of the Psychonomic Society*, 9(5), 353-356. <https://doi.org/10.3758/BF03337021>

Supported by research grants from Wesleyan University to both authors. We thank Robert J. White and Deborah A. Cassidy for technical assistance. Requests for reprints should be sent to the authors at the Department of Psychology, Wesleyan University, Middletown, Connecticut 06457. Cutting is also a staff member of the Itaska Laboratories.

# Motion in VR

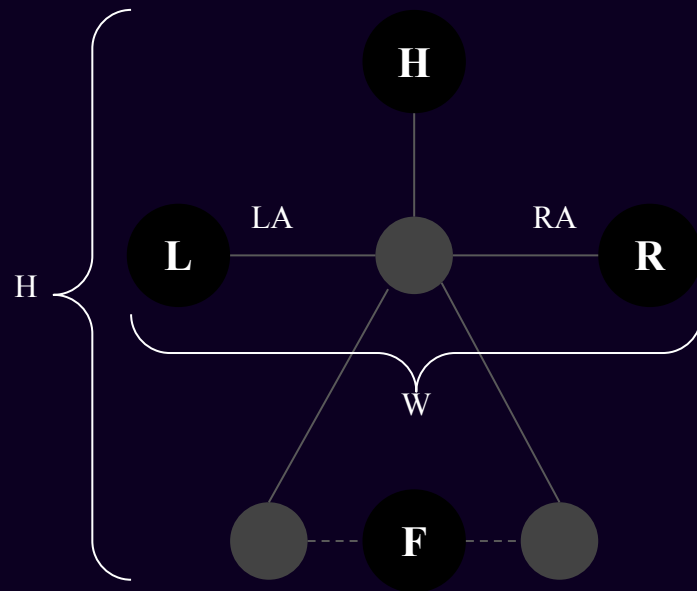


# VR identification studies



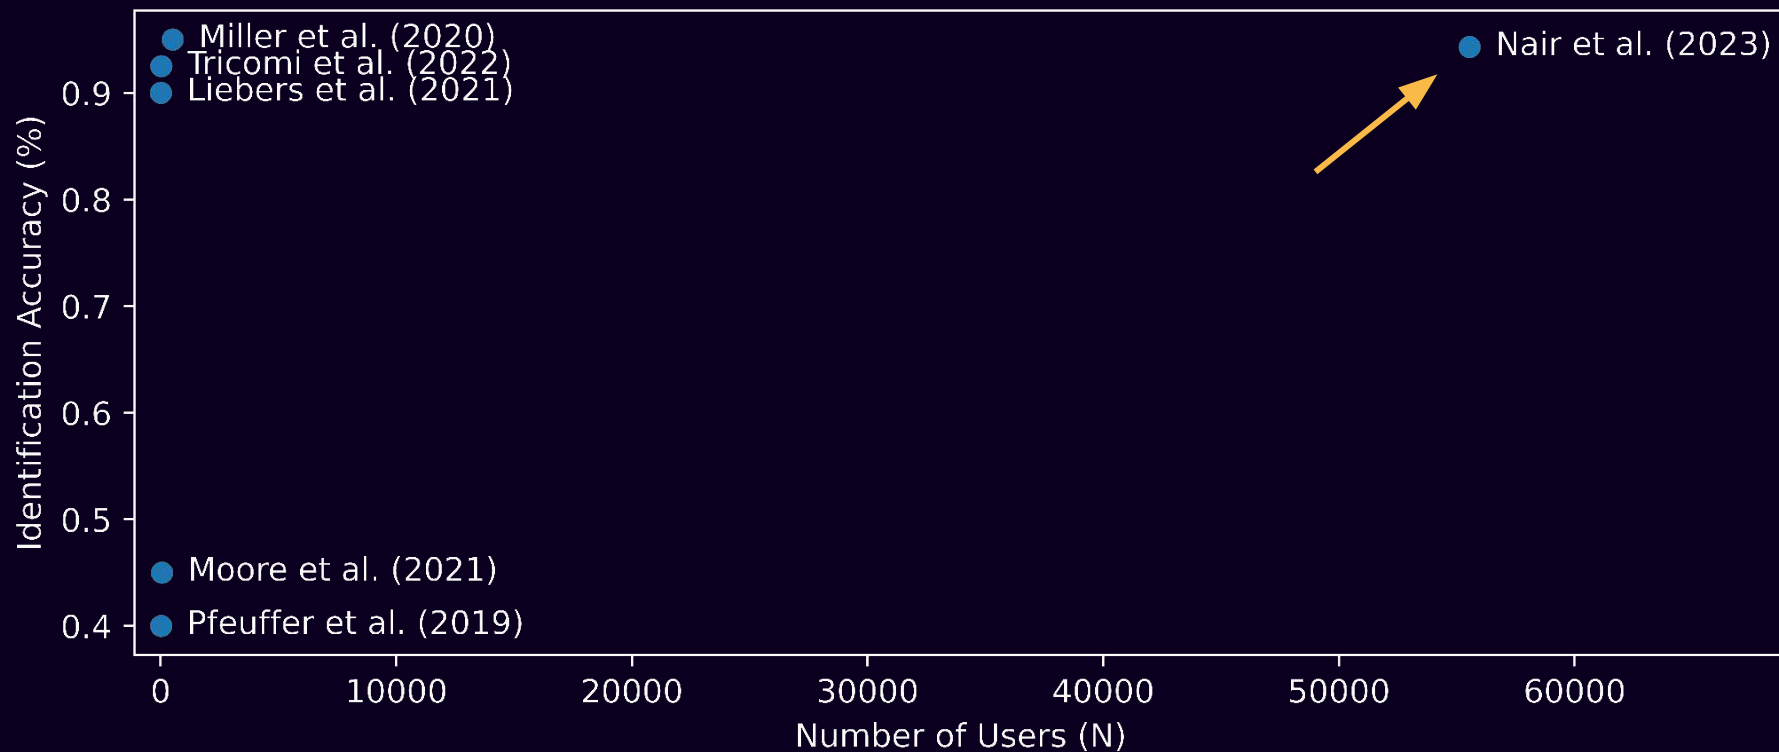


# Static identifiers



\* Enough to identify  $\approx 4,000$  users

# VR identification studies





# BEAT SABER

COMBO  
164  
1 336 468

## Believer Imagine Dragons RUSTIC, nico

Ranked 1.44★ **Expert** 👍👍👍👍👍


Hide filters and sorting ^

| Rank | Avatar | Player                                    | Rank         | Score    | Accuracy | Time    | Score  | Accuracy |
|------|--------|---|--------------|----------|----------|---------|--------|----------|
| #1   | 🇺🇸     | Marsh <b>LUCK</b>                         | 2 months ago | 122.84pp | 97.34%   | 481,912 | 97.34% | 481,912  |
| #2   | 🇺🇸     | MochiPower123                             | 4 months ago | 120.33pp | 97.14%   | 480,915 | 97.14% | 480,915  |
| #3   | 🇺🇸     | Pajma <b>ART</b> <b>LUCK</b>              | 5 months ago | 118.74pp | 97.01%   | 480,258 | 97.01% | 480,258  |
| #4   | 🇺🇸     | CheezitYT                                 | 2 months ago | 106.03pp | 95.13%   | 470,957 | 95.13% | 470,957  |
| #5   | 🇺🇸     | Hunter <b>MIEP</b> <b>ST</b>              | 6 months ago | 105.34pp | 95.64%   | 473,481 | 95.64% | 473,481  |
| #6   | 🇺🇸     | Dark_jimmy                                | 5 months ago | 103.39pp | 95.39%   | 472,275 | 95.39% | 472,275  |
| #7   | 🇺🇸     | Lautrick98 <b>SKIL</b> <b>BBKS</b>        | 3 months ago | 102.54pp | 95.28%   | 471,731 | 95.28% | 471,731  |
| #8   | 🇺🇸     | Cohenlinklater <b>BSC</b>                 | 3 months ago | 102.04pp | 94.59%   | 468,271 | 94.59% | 468,271  |
| #9   | 🇺🇸     | Emerald knight YT <b>RAGE</b> <b>FROG</b> | 3 months ago | 101.23pp | 94.47%   | 467,686 | 94.47% | 467,686  |
| #10  | 🇺🇸     | DEN   Ani <b>BSC</b> <b>DEN</b>           | last month   | 101.05pp | 95.09%   | 470,746 | 95.09% | 470,746  |

1 - 10 / 1470

Show more details v

# BeatLeader dataset



**USERS**

55,541

**REPLAYS**

2,669,886

**COUNTRIES**

40+

**SESSIONS**

713,013

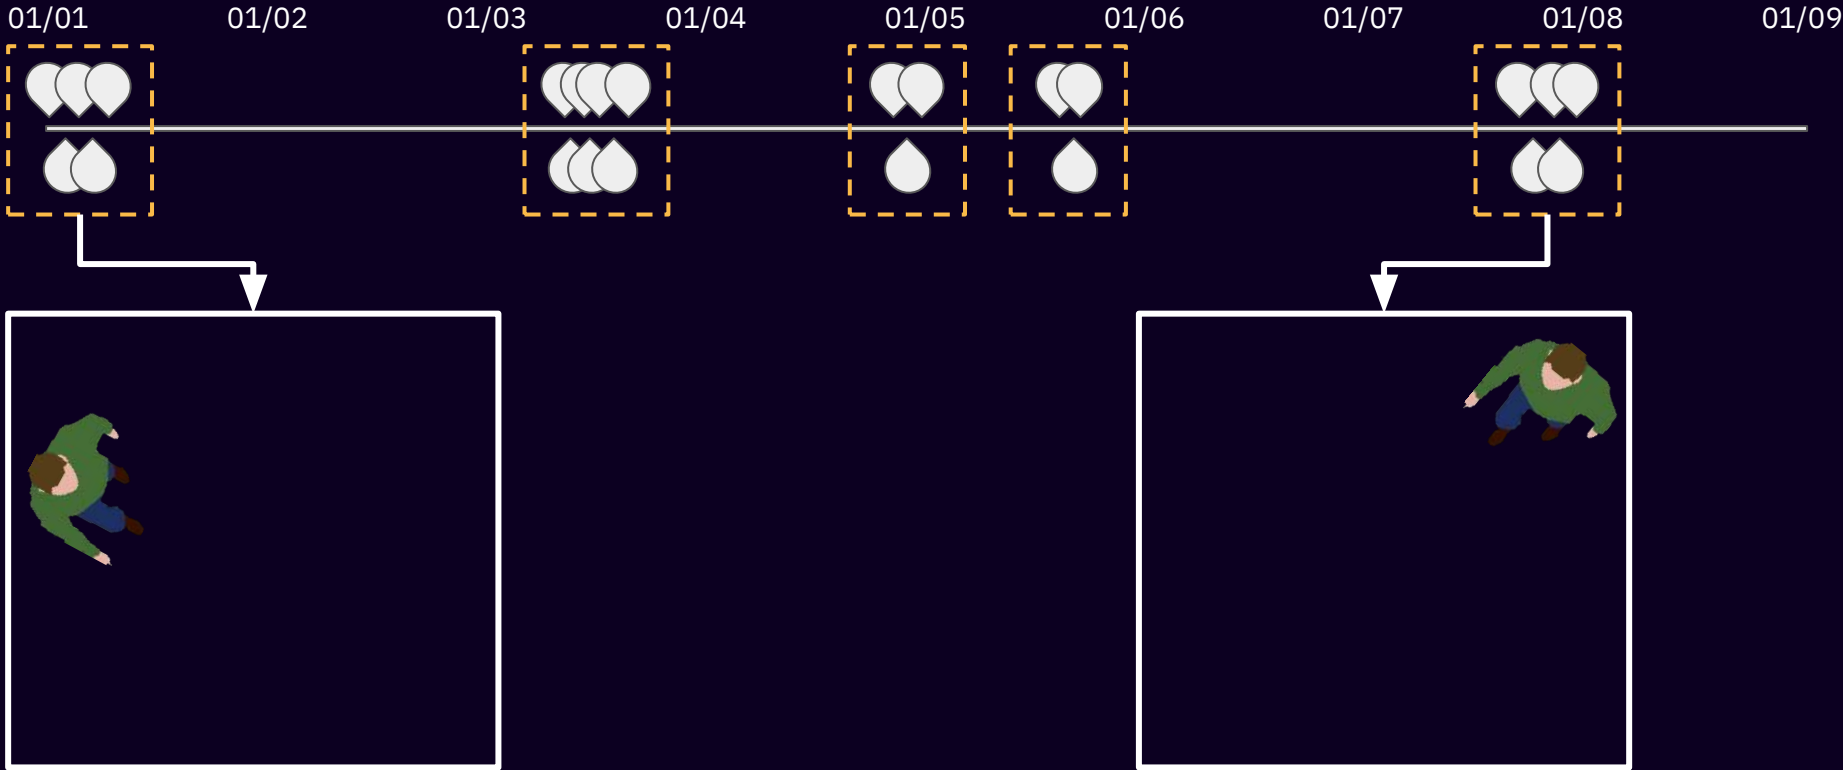
**VR DEVICES**

20+

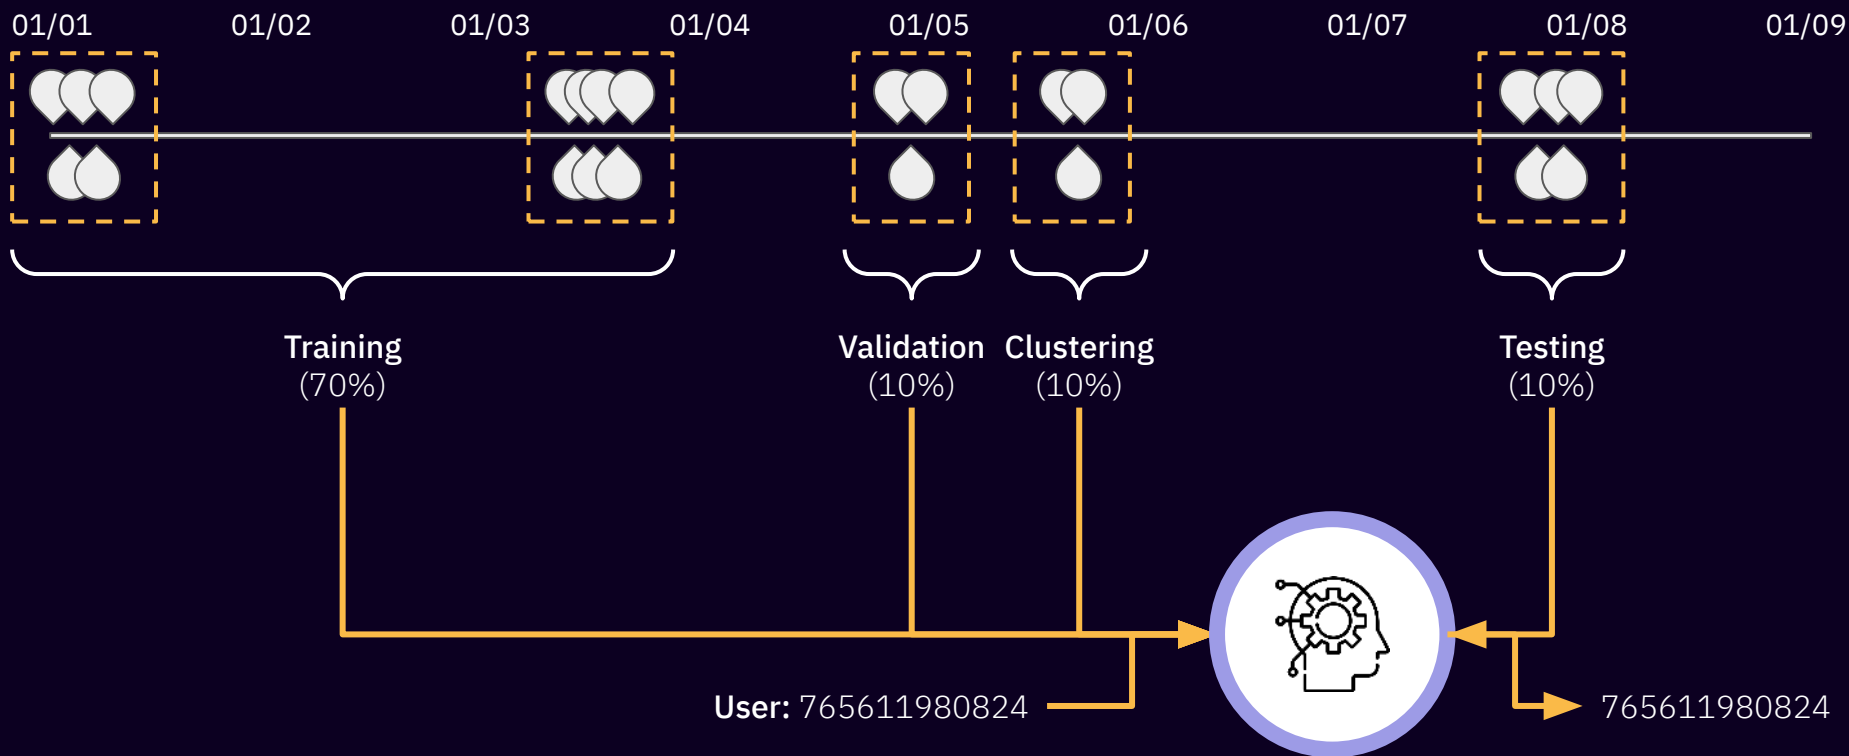
**DATA**

3.96 TB

# Usage sessions



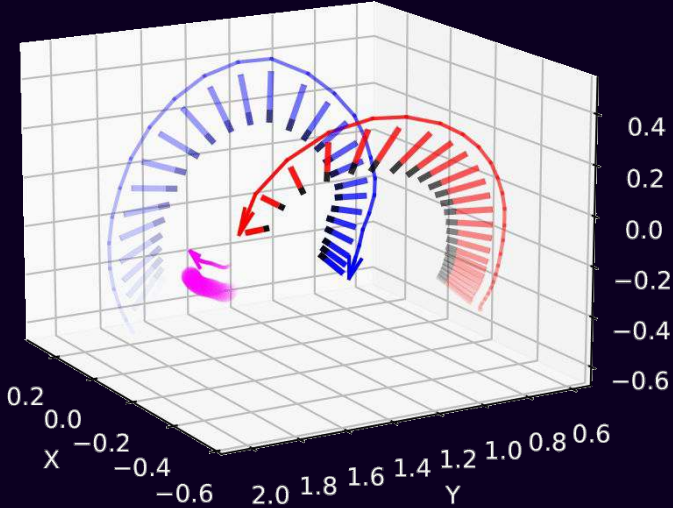
# Session splits




# Motion features

**Motion**

Left Hand    Right Hand    Head



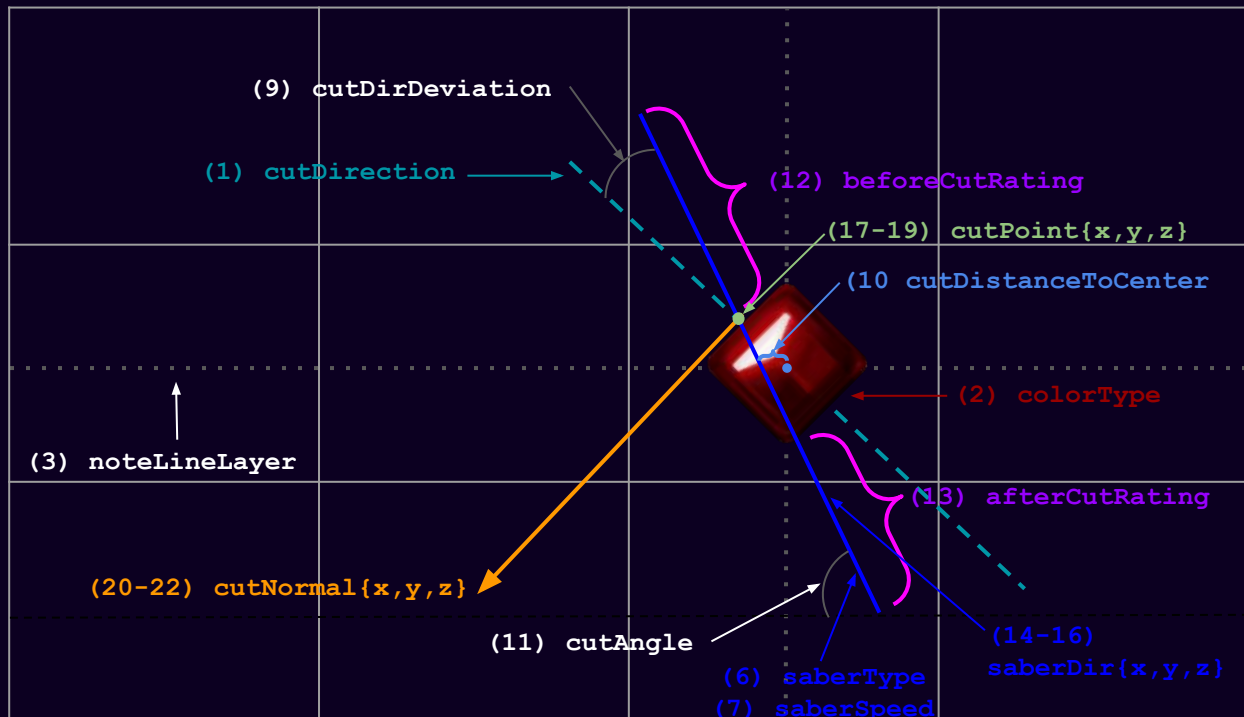
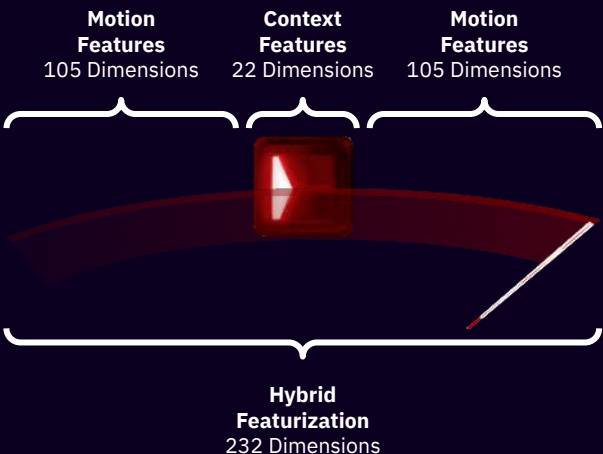
|   |
|---|
| {pos <sub>x</sub> , pos <sub>y</sub> , pos <sub>z</sub> , rot <sub>i</sub> , rot <sub>j</sub> , rot <sub>k</sub> , rot <sub>l</sub> } |
| {min, max, mean, med, stdev}  |
| {head, left_hand, right_hand}   |
| =   |
| [105 dimensions]  |



LightGBM

# Context features

(5) **scoringType** (Normal, Ignore, NoScore, SliderHead, SliderTail, BurstSliderHead, or BurstSliderElement)



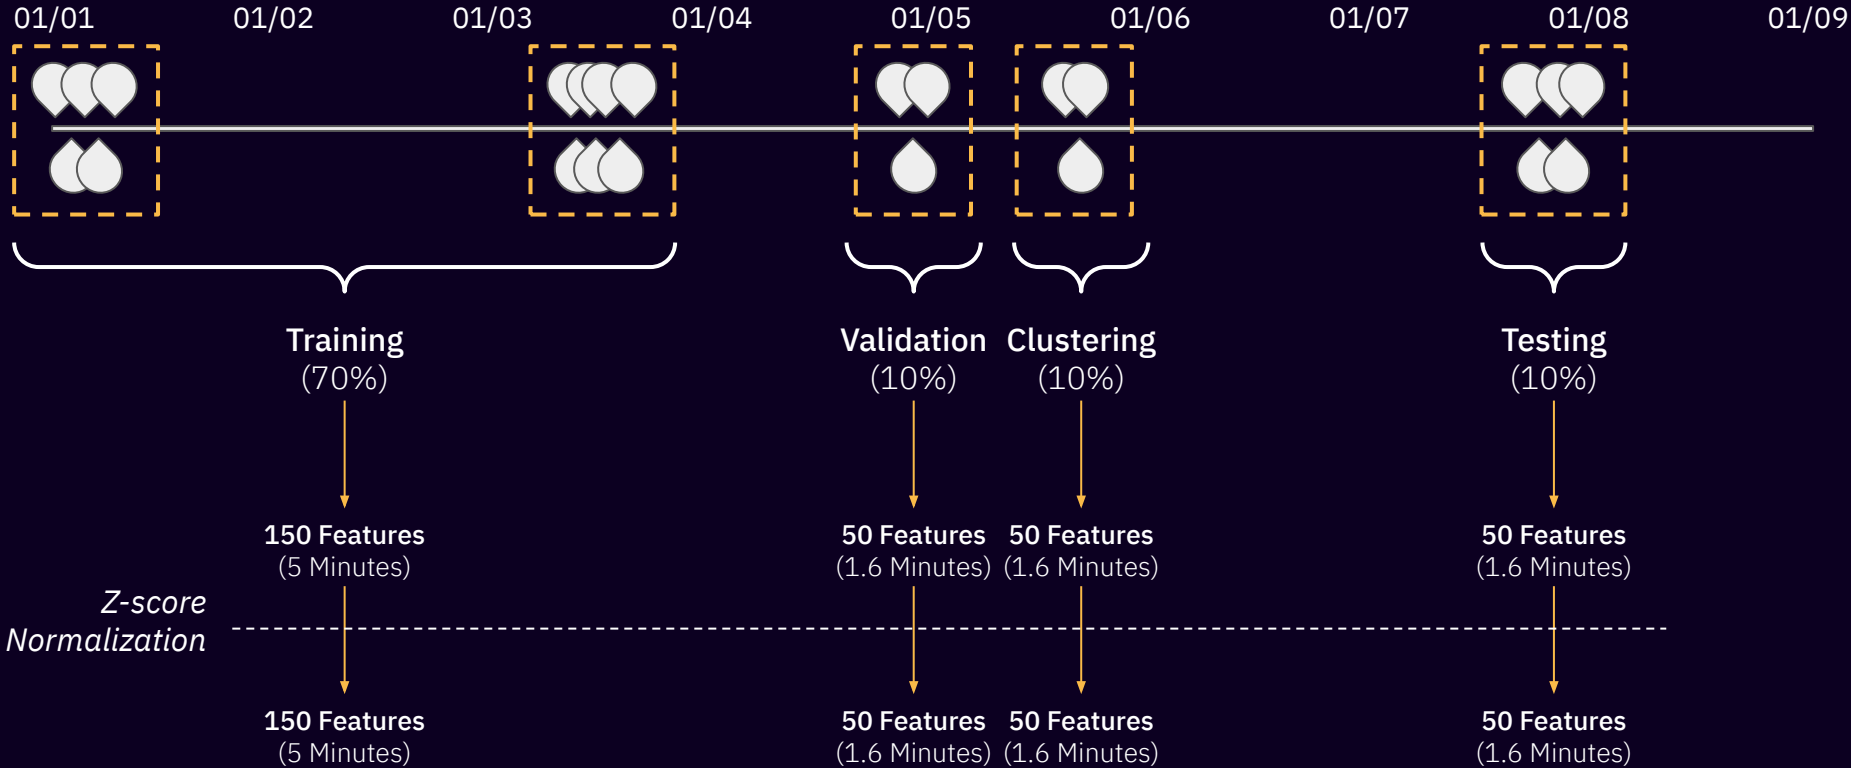
eventTime: when music matches note

hitTime: when saber intersects note

(8) **timeDeviation** (|eventTime-hitTime|)

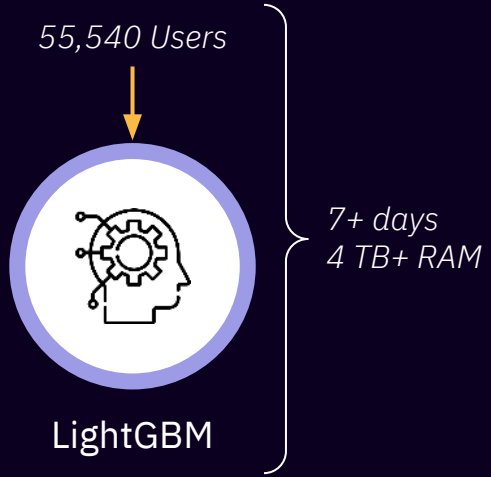
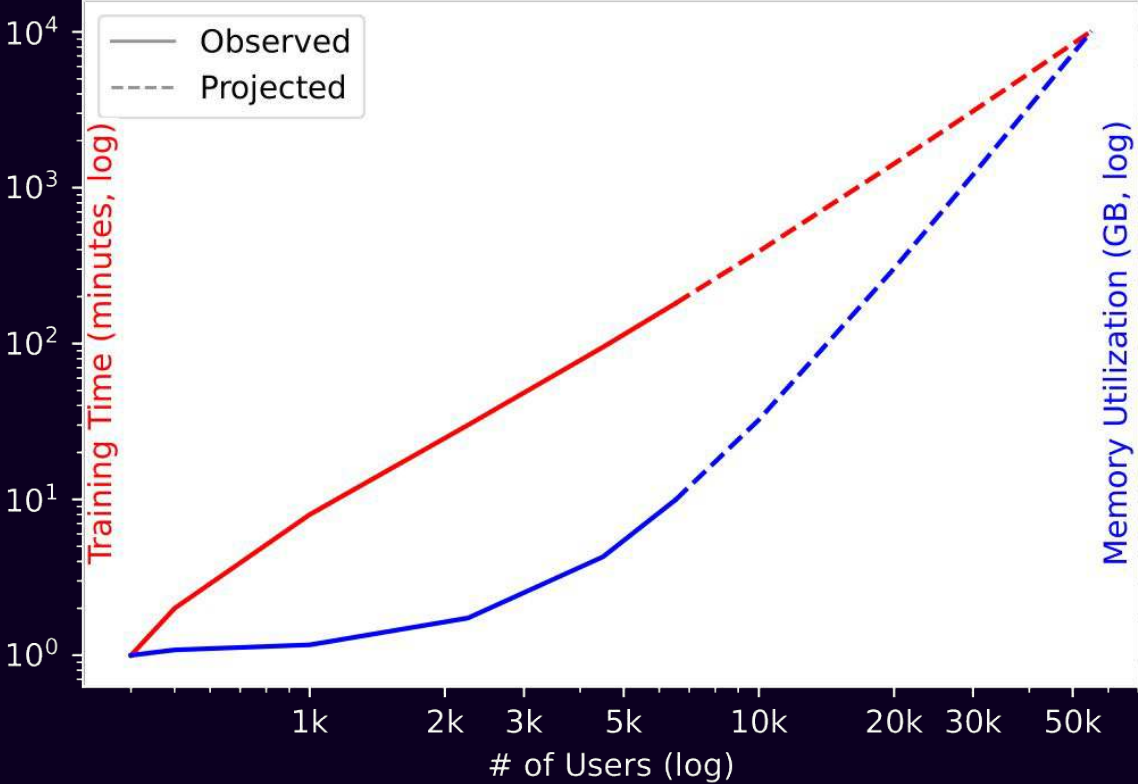
(4) **lineIndex**

# Feature extraction







# LightGBM training time



# Hierarchical model




# Hierarchical model

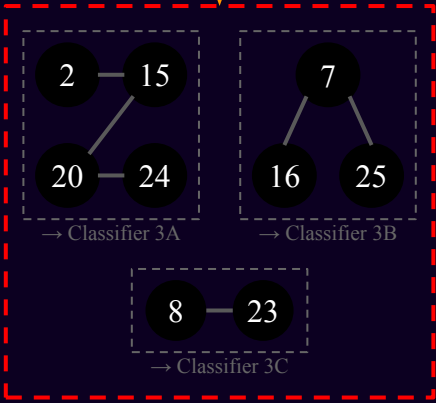


# Confusion layer

Clustering  
(10%)



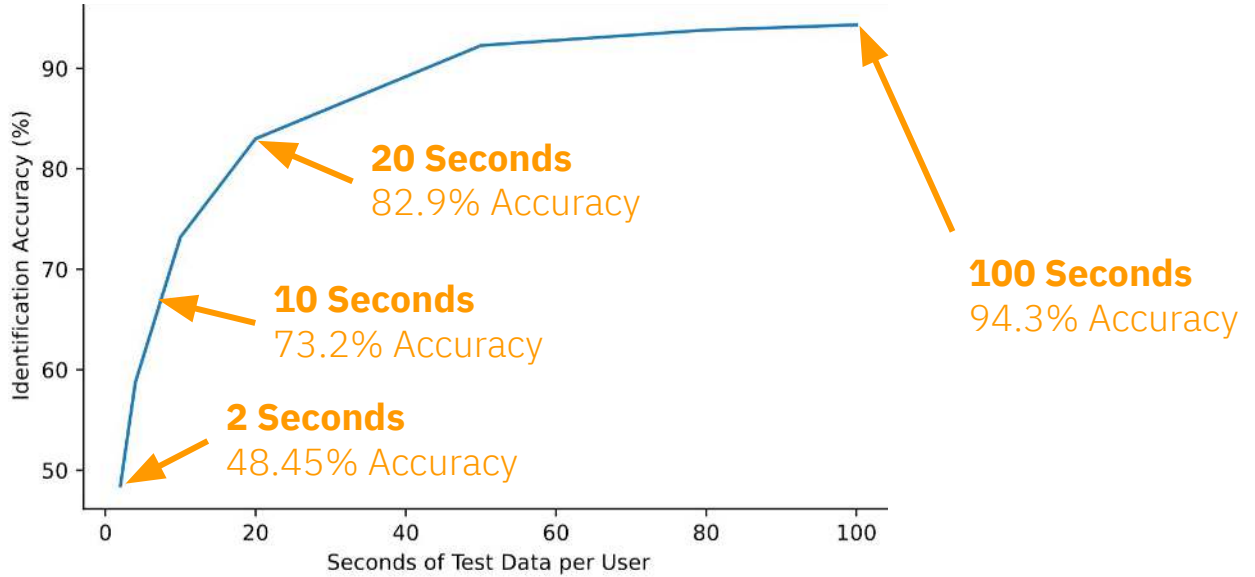
Correctly Classified




Incorrectly Classified

# Results

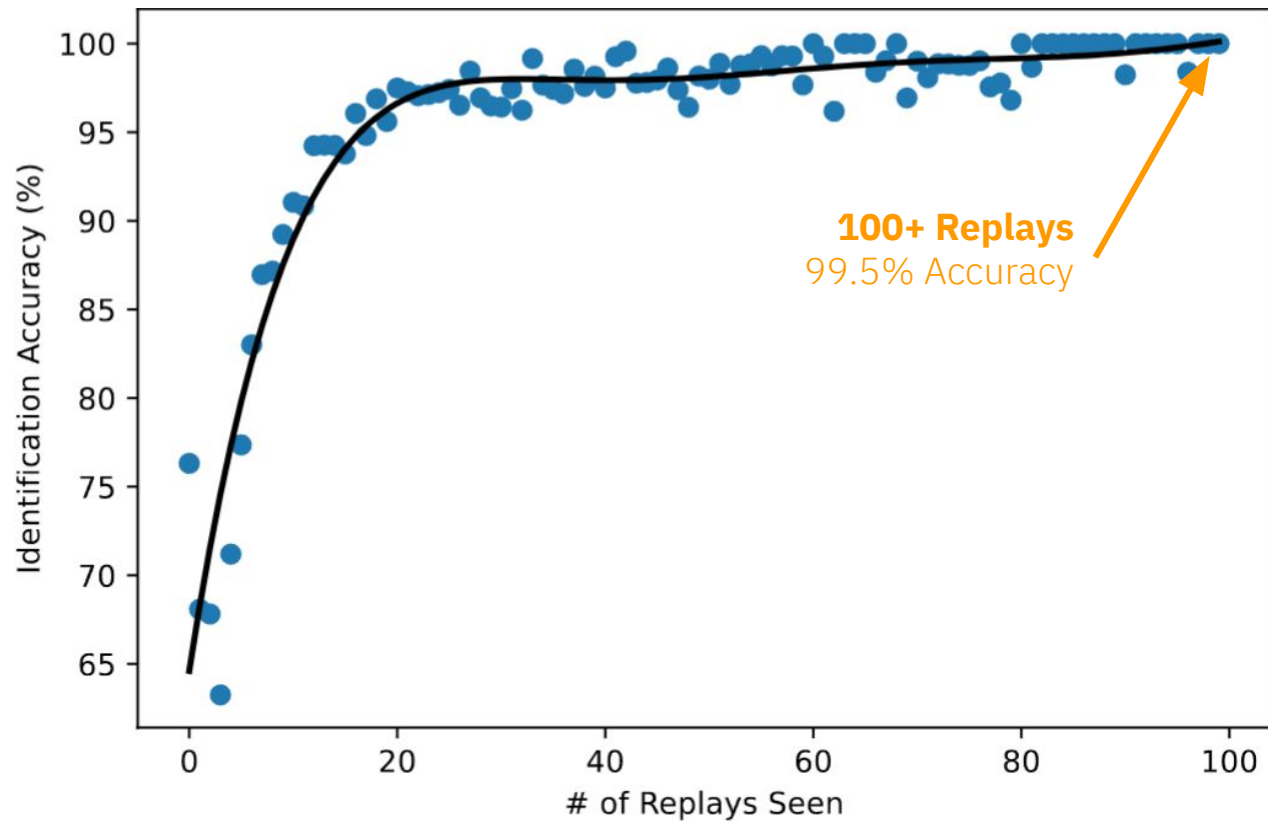
| Layer            | # of Models | Accuracy (per Model) | Accuracy (per Layer) |
|------------------|-------------|----------------------|----------------------|
| Layer 1          | 10          | 93.1%                | 90.2%                |
| Layer 2          | 10          | 93.1%                | 90.2%                |
| Layers 1 & 2     | 20          | 93.1%                | 91.0%                |
| Layer 3          | 5           | 84.0%                | 84.0%                |
| Layers 1, 2, & 3 | 25          | 91.3%                | 94.3%                |



# Impact factors




## *Impact of experience*





# Model explanations



## STATIC FEATURES

22.9% of entropy gain

Enough to identify  
≈ 4,000 users

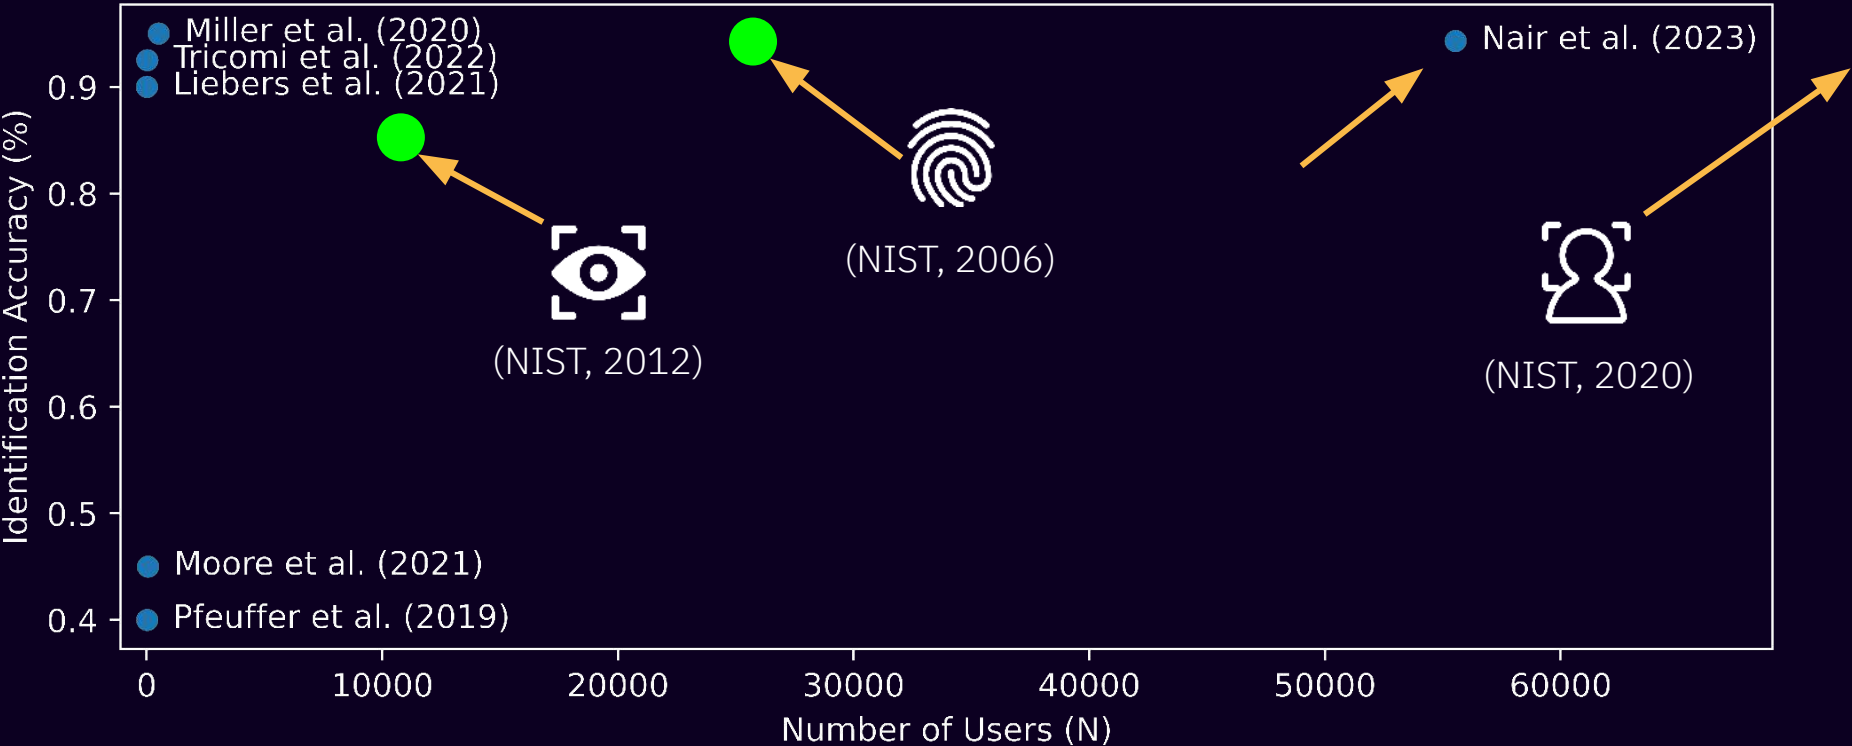
## MOTION FEATURES

73.9% of entropy gain


## CONTEXT FEATURES

3.2% of entropy gain

# Other biometrics



# Implications for VR privacy



# Thank you!

*50,000+ User Identification Study*

<https://rdi.berkeley.edu/vr-identification/>

<https://arxiv.org/abs/2302.08927>

<https://github.com/MetaGuard/Identification>

*Other RDI Metaverse Research*

<https://rdi.berkeley.edu/metaverse/>