Seeing Through: Analyzing and Attacking Virtual Backgrounds in Video Calls

Felix Weissberg & Jan Malte Hilgefort, Steve Grogorick, Daniel Arp, Thorsten Eisenhofer, Martin Eisemann, Konrad Rieck

Motivation

Virtual backgrounds in video conferences are commonly used to increase privacy.

We analyze the leakage of a user's real environment and present an attack to reconstruct it.

What's hidden beneath this virtual background?

Sensitive personal objects?
Secret or embarrassing information?

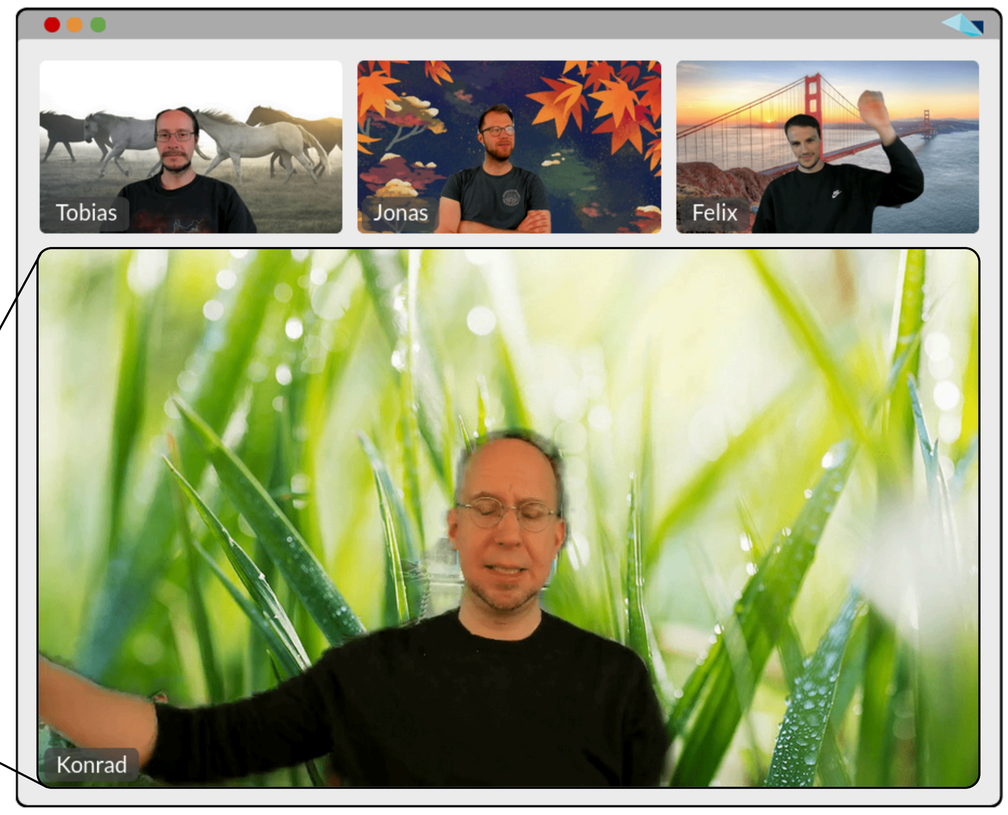

Analysis

How do video conferencing tools apply virtual backgrounds?

Reverse engineering and analyzing the virtual background feature reveals:

$$M = \text{SEGM}(\text{SCALE}_{\downarrow}(X))$$
$$\hat{X} = \text{BLEND}(X, V, \text{SCALE}_{\uparrow}(M))$$

How much information leaks about users' real environments?

Analyzing the leaks in 1080 videos with various real and virtual backgrounds shows:

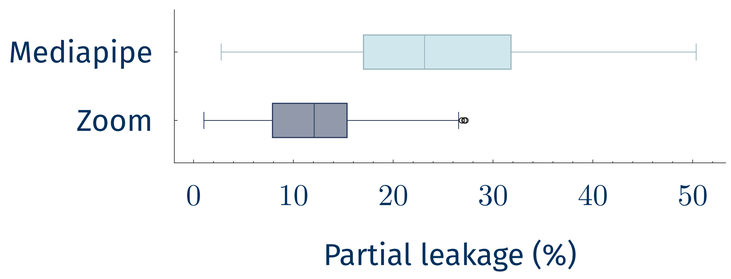

Attack

How can an attacker reconstruct the real environment of a user?

Leaks are the difference between the original blending mask and the actual human mask. These masks can be approximated by re-purposing the video conferencing tool and a high quality segmentation model.

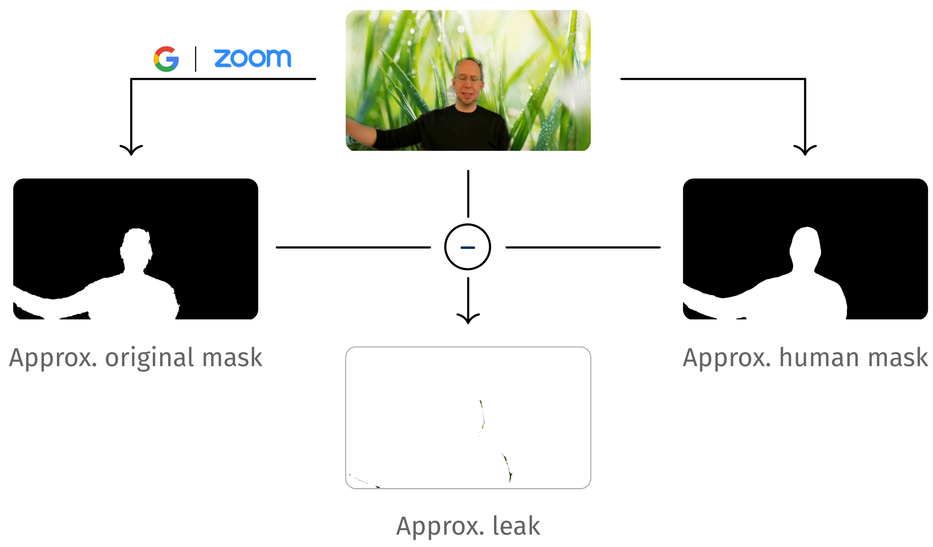

Results

On average between 9.5% and 14.1% of available information from video calls can be reconstructed.

TECHNISCHE UNIVERSITÄT BERLIN

Machine Learning and Security