I am writing to offer my strong support for “Investigating Influencer VPN Ads on YouTube” by Omer Akgul, Richard Roberts, Moses Namara, Dave Levin, and Michelle Mazurek to receive The Annual Best Scientific Cybersecurity Paper award. This paper appeared in IEEE Security & Privacy 2022.

What the paper is about
One of the most important pillars of cybersecurity research is user education: what are users learning about cybersecurity, how accurate is it, and how can we educate to ensure secure behavior? This paper observes that, perhaps despite our community’s best efforts, one of the most ubiquitous forms of cybersecurity education today comes from VPN ads included within YouTube videos (not served interstitially by YouTube). The paper is the first to investigate this potentially influential channel of user education.

The paper shows that these “influencer” VPN ads are extremely common and found in virtually every category of YouTube video; therefore, influencing a vast demographic range. The central question of the paper is: what information is conveyed in YouTube videos’ VPN ads?

The paper’s scientific contributions
To answer this question, the authors randomly sampled 1.4% of YouTube (86M videos), identifying 243 videos with VPN ads (scaling this up to all of YouTube, this would mean that ~17,000 videos have VPN ads comprising billions of views).

To develop a rigorous, scientific understanding of the threat models being conveyed in VPN ads, the authors crafted a novel qualitative coding framework that decomposes key statements into four components: subject (who the attacker is), verb (what the attacker is doing), predicate (what is being attacked), and mitigation (how the VPN protects). For example, “the government (subject) is trying to surveil (verb) your location (predicate).” See Figure 3 in their paper for a particularly nice visualization of this.

New insights
This approach results in a deep, quantitative analysis of what (mis)information is provided in VPN ads. The authors identify many instances of inaccurate threat models, overclaims of VPNs’ capabilities, and misleadingly vague threats.

It also allows the authors to quantitatively compare messaging over time, across different VPN providers, and across YouTube channel topics. I found it particularly interesting, for instance, that one popular VPN provider, VirtualShield, almost exclusively targets their ads to conspiratorial YouTube channels, which are far more likely to label governments as attackers.

This shows that the messaging in VPN ads has stratified; different users are being taught different things about what threats exist (and how VPNs protect them). Some of these messages are giving users incorrect information about online threats (which may affect their security and privacy choices beyond VPNs). This provides a more nuanced view of the VPN ad ecosystem, made possible only by the paper’s techniques and scientific rigor.

Closing
In sum, I believe this paper is deserving of this award because it takes a nebulous observation that probably many others have made—that VPN ads are almost laughably inaccurate—and develops the techniques to allow a rigorous analysis of the ads’ messages. The resulting insights have high potential for impact.