Poster: Understanding Risks of *Semi-Drive-By* Splash Ads

Song Wu*, Yifan Zhang†, Xueqiang Wang‡, Yinfeng Cao§¶, Bo Wang*, Qin Wang‖
*Independent researcher | †San Diego State University
‡University of Central Florida | §PolyU | ¶CTIHE | ‖CSIRO Data61
*researchersongwu@gmail.com ,wan.id.none@gmail.com| †yzhang24@sdsu.edu |
‡xueqiang.wang@ucf.edu | §csyfcao@comp.polyu.edu.hk | ¶qinwangtech@gmail.com

***Abstract*—We study fraud in the mobile splash-ad ecosystem and challenge the common assumption that ad networks are trustworthy measurement authorities. We identify *semi-drive-by* splash ads, a previously unreported form of ad-network fraud in which splash ads are triggered by incidental or indirect interactions rather than deliberate user intent, turning non-ad interactions into billable engagement. Our empirical results suggest that this behavior is widespread in popular Android apps, and our collaboration with a major advertiser led to repayments of about 4 million Yuan (≈600,000 USD).**

1. Introduction

Mobile splash ads, full-screen advertisements shown during app launch, have become a major monetization channel in the mobile ecosystem. In this setting, ad networks occupy a particularly powerful position: through their SDKs embedded in third-party apps, they control ad delivery, interaction tracking, DeepLink redirection, attribution, and billing. Prior work [1], [2], [3] on mobile ad fraud largely assumes a malicious publisher or app developer and therefore focuses on application-layer manipulation, such as fake impressions or click inflation introduced by the host app.

In contrast, we consider a stronger and largely overlooked threat model in which the ad network itself is the malicious actor. Because the network controls both the logic that defines "valid engagement" and the telemetry reported upstream, it can distort advertiser-facing metrics in ways that are difficult to independently audit.

2. New: *Semi-drive-by* Splash Ads

We identify a new form of ad fraud in this setting, which we call *semi-drive-by* splash ads. In these ads, landing pages or click-equivalent engagement events are triggered not by deliberate ad-directed user input, but by incidental or indirect interactions during normal phone use. A representative example is shake-to-trigger splash ads, where an SDK sets overly sensitive motion thresholds and misclassifies ordinary device movement, such as walking, handling the phone, or riding in a car, as intentional engagement, immediately opening the advertiser's landing page. Because coerced activation and genuine intent have fundamentally different commercial value, this misattribution inflates engagement statistics, wastes ad budget, and corrupts performance analytics while remaining difficult for advertisers to audit.

Detecting *semi-drive-by* splash ads at scale is challenging for three reasons. First, the behavior is delivered dynamically rather than as a fixed offline code pattern, making static analysis insufficient and large-scale physical-device testing costly and impractical. Second, straightforward emulator-based analysis often fails because deceptive splash-ad behaviors may disappear in virtualized, freshly initialized, or otherwise unrealistic environments. Third, even when triggered, *semi-drive-by* ads are difficult to identify because ad SDKs reuse common mechanisms such as DeepLinks and inter-app jumps that also appear in benign app workflows. As a result, a practical detection framework must support scalable dynamic analysis, provide realistic execution conditions that can induce hidden behaviors, and distinguish ad-driven redirections from normal app activity.

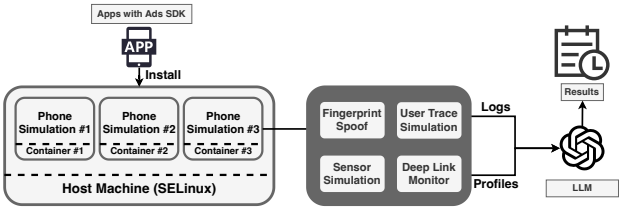

Figure 1. Design overview of ADHIVE

3. Our Detector: ADHIVE

To address this problem, we propose ADHIVE (Figure 1), an automated honeypot framework for detecting *semi-drive-by* splash ads on Android. ADHIVE is designed to expose ad-triggered behaviors under realistic execution conditions and consists of three key components.

First, ADHIVE uses an environment-hardened Android runtime that ports coherent profiles from real devices into a virtualized environment. This functionality is implemented by the *Fingerprint Spoof* module, as shown in Figure 1. Rather than synthesizing isolated fingerprints, this design preserves cross-layer consistency across system properties, filesystem artifacts, kernel-visible attributes, and Android

TABLE 1. SUMMARY OF AD FRAUD DETECTION IN APP MARKETS

Market	Tested Apps	Flagged Apps	Rate (%)
VIVO	5,262	261	4.96
XIAOMI	5,998	128	2.13
WDJ	11,682	198	1.70
YYB	8,881	191	2.15

service responses. As a result, the execution environment appears substantially more realistic to ad network SDKs and is more resilient to common anti-analysis checks.

Second, ADHIVE drives apps with realistic long-horizon behavioral inputs. Specifically, the *User Trace Simulation* and *Sensor Simulation* modules generate plausible usage traces that reflect routine user context and controllable motion patterns that can trigger hidden ad logic. This design enables ADHIVE to exercise stateful and sensor-dependent activation paths that are unlikely to appear in short scripted tests or conventional sandbox executions.

Third, ADHIVE applies a learning-assisted log analysis pipeline to identify and attribute suspicious ad-triggered behaviors. This functionality is supported by the *log analysis* module, which analyzes DeepLink structures, payload characteristics, and temporal context to distinguish normal in-app navigation from monetization-driven redirections and to extract attribution-related fields associated with advertising traffic. In this way, ADHIVE can not only detect suspicious landing events, but also link them to specific SDK-mediated ad flows with high precision.

4. Evaluation

We apply ADHIVE to a large-scale measurement of popular non-game Android apps in the Chinese market.

We crawl apps from four major Android app markets, XIAOMI [1], VIVO[2], YINGYONGBAO (YYB)[3], and WANDOUJIA (WDJ)[4] and collect $32,758$ distinct packages. After filtering apps without network permissions and removing samples that are incompatible at runtime, our final dataset contains $31,823$ apps. We then run 24-hour uninterrupted launch-phase testing under realistic device, network, and motion conditions.

Our preliminary results (Table 1) show that *semi-drive-by* splash ad behavior is not rare. We identify suspicious behavior in $261, 128, 198$, and 191 apps from the XIAOMI, VIVO, YYB, and WDJ markets, corresponding to prevalences of 4.96.

We also performed regional comparison experiments on 1,000 randomly sampled apps across six cities from three city tiers, based on the GDP classification of Chinese cities. As Figure 2 shows that *semi-drive-by* ads selectively delivered across regions.

1. https://m.app.mi.com
2. https://h5.appstore.vivo.com.cn
3. https://sj.qq.com
4. https://wandoujia.com

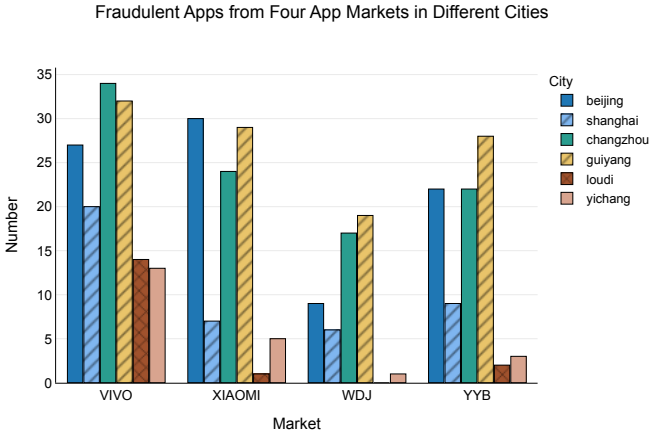

Figure 2. Regional comparison experiment of 4 app markets

Our findings also reveal concrete real-world impact. In collaboration with a major advertiser, *a Top marketplace A*[5], we investigated substantial inconsistencies in conversion metrics and confirmed systematic misreporting consistent with deceptive interaction triggering. Following disclosure, two leading ad networks acknowledged the behavior, and one refunded 4 million Yuan ($\approx 600,000$ USD), representing 18% of the advertiser's total spending on that network during the relevant period.

Responsible Disclosure

All materials used were either obtained with authorization or disclosed to the corresponding companies where appropriate. *Top Marketplace A* authorized our research and provided essential materials. *Pinduoduo* and *ByteDance* confirmed that our observations were consistent with their internal findings and that relevant defenses had already been deployed. *Kuaishou Technology* acknowledged the issue and stated that its impact remained acceptable for its services. *Alipay* assessed our findings as a medium-level threat.

References

[1] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell, "DeCAF: A deep convolutional activation feature for generic visual recognition," in *International Conference on Machine Learning (ICML)*, 2014, pp. 647–655.

[2] F. Dong, H. Wang, L. Li, Y. Guo, T. F. Bissyandé, T. Liu, G. Xu, and J. Klein, "FraudDroid: Automated ad fraud detection for Android apps," in *Proceedings of the ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering (ESEC/FSE)*, 2018, pp. 257–268.

[3] T. Zhu, Y. Meng, H. Hu, X. Zhang, M. Xue, and H. Zhu, "Dissecting Click Fraud Autonomy in the Wild," in *ACM SIGSAC Conference on Computer and Communications Security (CCS)*, 2021, pp. 271–286.

5. One of the largest Chinese buyers of splash ads, it represents substantial traffic and billing volume, making rare anomalies measurable. The advertiser requested anonymity for business reasons.

Understanding Risks of Semi-Drive-By Splash Ads

Song Wu [1] Yifan Zhang[2] Xueqiang Wang [3] Yinfeng Cao [4,5] Bo Wang [1] Qin Wang [6]

[1]*Independent* | [2]*San Diego State University* | [3]*University of Central Florida* | [4]*PolyU* | [5]*CTIHE* | [6]*CSIRO Data61*

Introduction

We move beyond malicious publishers and study a stronger threat model where the ad network itself can manipulate engagement metrics by controlling both ad logic and telemetry.

Ad SDKs in mobile advertising

Ad networks occupy a particularly powerful position,because, through their SDKs embedded in third-party apps, they control ad delivery, interaction tracking, DeepLink redirection, attribution, and billing.

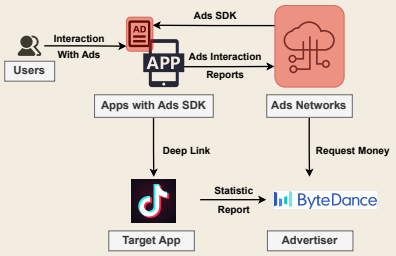

Ads Network

New threat: semi-drive-by splash ads

- Semi-drive-by splash ads are triggered by daily phone use rather than deliberate clicks, inflating engagement and misallocating ad spend.
- Advertisements are triggered by incidental or indirect interactions and billed as legitimate user engagement.
- Semi-drive-by splash ads are inherently stealthy to both users and advertisers.

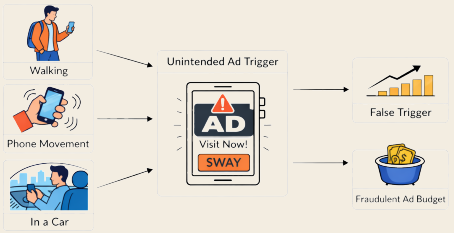

Semi-drive-by Splash ads

Detection is challenging

Dynamic. The behavior is delivered dynamically.
Anti-VM. The behavior may disappear in Virtual Enviroment.
Similarity. The behavior is similar to the normal trigger.

Our detector: ADHIVE

ADHIVE is designed to expose ad-triggered behaviors under realistic execution conditions and consists of three key components.

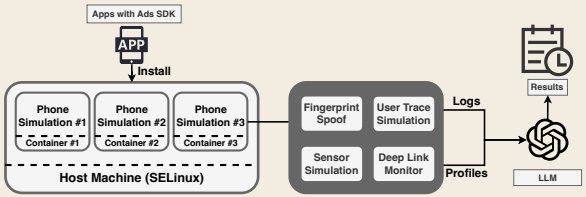

ADHIVE System Architecture

- **Virtualized Environment.** Fingerprint Spoof module ensures coherence across system properties and filesystem artifacts, kernel-visible attributes, and Android service responses. Execution environment is more realistic to ad network SDKs.
- **Behavior Simulation.** User Trace Simulation and Sensor Simulation modules generates long-horizon behavioral inputs, allowing it to trigger hidden ad logic that typical short tests miss.
- **Log Analysis.** Deep Link Monitor generates the logs and learning-assisted log analysis pipeline to detect ad-triggered behaviors.

Large-scale measurements

To conduct a large-scale measurement, we collected 31,823 applications from four representative markets. The results are summarized in the table. As shown in the table, semi-drive-by splash ads are not rare.

Market	Tested Apps	Flagged Apps	Rate (%)
VIVO	5,262	261	4.96
XIAOMI	5,998	128	2.13
WDJ	11,682	198	1.70
YYB	8,881	191	2.15

Measurement Result

Regional difference

Semi-drive-by ads selectively delivered across regions.Semi-drive-by ads is detected most frequently in second-tier cities, followed by first-tier cities, and is much less common in third-tier cities, where the detection volume is significantly lower.

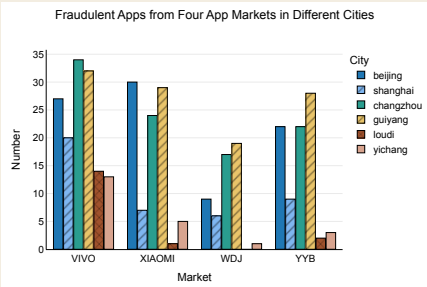

Regional Comparsion

Key takeaways

- Semi-drive-by ads is not rare in China.
- Semi-drive-by ads is related to the placement of Ads buyer.
- The targets of semi-drive-by ads are not limited to ordinary mobile apps, but also include super apps.

Responsible disclosure

- Top marketplace A authorized our research, provided essential materials and got repayments of about 4 million Yuan (≈600,000 USD).
- Pinduoduo and ByteDance confirmed our observations (i.e., consistent with their internal findings).
- Kuaishou Technology acknowledged the issue and stated that its impact remained acceptable for its services.
- Alipay assessed our findings as a medium-level threat.

References

[1] Bin Liu, Suman Nath, Ramesh Govindan, and Jie Liu.
DECAF: Detecting and characterizing ad fraud in mobile apps.
In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, pages 57–70, 2014.

[2] Tong Zhu, Yan Meng, Haotian Hu, Xiaokuan Zhang, Minhui Xue, and Haojin Zhu.
Dissecting click fraud autonomy in the wild.
In *ACM Conference on Computer and Communications Security (CCS)*, pages 271–286, 2021.