Here are some of the ways that our machines learn what is good and what is bad:

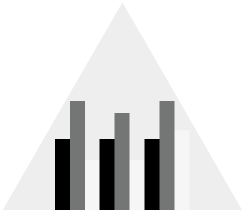

Static analysis

We analyze application code without running the app. Application features are extracted and analyzed against expected good behavior and potential bad behavior.

Dynamic analysis

We run applications to identify interactive behavior that cannot be seen with static analysis. This allows reviewers to identify attacks that require connection to a server and dynamic downloading of code.

Third-party reports

We cultivate active relationships with industry and academic security researchers. These independent security researchers also evaluate applications in a variety of ways and will often let us know if they see something amiss.

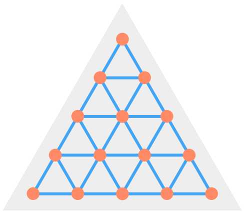

Developer relationships

We analyze non-code features to determine possible relationships between applications and to evaluate whether the developer that created the application may have previously been associated with creation of Potentially Harmful Applications.

Signatures

We use signatures to compare apps against a database of known bad apps and vulnerabilities.

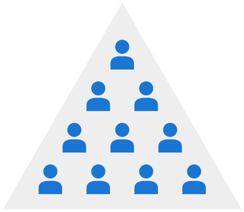

SafetyNet

A privacy preserving sensor network spanning the Android ecosystem, identifying apps and other threats that cause harm to the device.

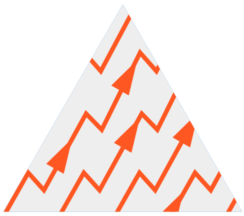

Heuristic and similarity analysis

We compare applications with each other to find trends that lead to harmful apps.