

MRC Supplement to IAB Guidelines for the Conduct of Ad Verification: Enhanced Content Level Context and Brand Safety

September 2018 FINAL

Sponsoring associations:

Media Rating Council (MRC)

American Association of Advertising Agencies (4A's)

Association of National Advertisers (ANA)

Interactive Advertising Bureau (IAB U.S.)



Table of Contents

- 1 Executive Summary 3**
- 1.1 About the Development of this Supplement 6
- 2 Content Level Granularity 6**
- 3 Social Media and UGC 8**
- 3.1 Data Freshness and Quality 9
- 3.2 Curation of UGC 10
- 3.3 Scalability 11
- 3.4 Machine Learning 11
- 3.5 Human Intervention 13
- 3.6 Monetization Controls 13
- 3.7 Taxonomy and Categorization 14
- 4 Adjacency 16**
- 5 Metadata 17**
- 6 Mobile Applications 20**
- 7 Illegal Sources and Industry Communication 21**
- 8 Enhanced Buyer Tools 22**
- 9 Required Reporting 22**
- 10 Use of Samples 23**
- 11 Glossary of Terms 23**
- 12 References 25**
- 13 Supporting Associations and Participating Organizations 25**
- 14 Contact Us 27**
- Appendix A: Section Summaries 28**

1 Executive Summary

The Interactive Advertising Bureau (IAB) and Media Rating Council (MRC) produced the *Guidelines for Conduct of Ad Verification* (the *Ad Verification Guidelines*) in February 2012, which identified a common set of acceptable methods and practices for conducting digital ad verification services. Ad verification is a service that offers technology to ensure that digital ads appear on intended sites and reach the targeted audience. The original primary service lines of ad verification are: (1) Site Context, (2) Geo-Targeting, (3) Ad Placement, (4) Competitive Separation and (5) Fraud Detection. Included within the Site Context service line is the concept of Brand Safety or, practices and tools to ensure that an ad will not appear in a Context that can damage an advertiser's brand. Historically, Brand Safety and targeting have been addressed at a domain, site or Uniform Resource Locator (URL) level with Context classification and ad alerting or blocking based on avoidance or target categories.

For purposes of this Supplement Context refers to certain attributes of a URL, site, mobile application or piece of content that are used as determinants as to whether the destination where a digital ad is to be served represents an appropriate environment in which the ad should appear, as determined by parameters set by the advertiser. Brand Safety refers to practices and tools to ensure that a digital ad will not appear adjacent to or in a Context that can damage an advertiser's brand.

Since the development of the *IAB/MRC Ad Verification Guidelines*, the digital advertising ecosystem has become much more complex and with this complexity, Context and Brand Safety issues have become more nuanced and challenging. Compounding this complexity is the monetization of Social Media Platforms (referred to as platforms throughout this document) and User Generated Content (UGC), which has introduced great volume and scale as well as reduced content control when compared to professionally curated publisher content. The MRC produced *Social Media Measurement Guidelines* in November 2015, which helped define key terms for social media and establish guidance of measurement of social media activity, audience and ad exposure on Social Media Platforms. However, to date, the impact of social media monetization on Context and Brand Safety has largely not been addressed in existing measurement guidance.

While existing IAB and MRC ad measurement and verification guidance should continue to be applied to Social Media Platforms and UGC, this Supplement is intended to promote enhanced Context and Brand Safety guidance across these platforms and content as well as to be used in conjunction with existing guidance as part of MRC accreditation auditing. Specifically, this Supplement seeks to achieve the following objectives:

- Establish guidance for more granular (content level) Context and Brand Safety measurement and reporting including video and display content within a domain, site, platform, mobile application or URL.

- Establish guidance for Context classification and Brand Safety measurement of social media and UGC including:
 - Data Freshness and Quality;
 - Curation of UGC;
 - Scalability;
 - Machine Learning;
 - Human Intervention;
 - Monetization Controls; and
 - Taxonomy and Categorization
- Establish guidance for defining and measuring domain, site, platform, mobile application or page Adjacency in order to enable Context and Brand Safety mechanisms within this parameter.
- Establish guidance for ingestion, use and consideration of content Metadata (such as content characteristics, user and exposure activity, comments and other relevant data) for Context and Brand Safety consideration.
- Establish guidance for Context and Brand Safety considerations within mobile applications.
- Establish guidance for identification, monitoring and treatment of content from illegal sources or domains and Industry communication regarding known bad sources and Brand Safety matters.
- Establish guidance for enhanced buyer tools for both targeting and avoidance.
- Establish guidance for reporting of content level Context and Brand Safety outcomes.
- Establish guidance for use of samples in Context and Brand Safety measurement.

Social Media Platforms and large digital publishers may have advanced techniques to identify and categorize content for purposes of Context and Brand Safety, however the independent approach of third-party verification vendors remains of vital importance. Approaches to content level ad verification will need to be a collaborative approach. **Platforms and publishers are strongly encouraged to either allow direct third party measurement at a content level, or to actively provide validated first-party data at a granular level to enable independent Context and Brand Safety measurement, one of which will be required for content level accreditation of third-party ad verification solutions. Further, the collection and transmission of data enabling content level Context and Brand Safety determinations is strongly encouraged to be subjected to audit and independent validation.**

It is important to note when considering the various aspects of Context and Brand Safety in this Supplement that different buyers may not only have different targets and avoidance categories, but also different tolerance levels for risk depending on brand suitability. Ultimately, ad verification organizations seeking content level accreditation must provide consistent and discrete reporting to users in order to enable customized targeting and avoidance and to allow users to stipulate which levels meet the needs of a particular campaign or brand.

With the issuance of this guidance, to the extent applicable, MRC’s accreditation considerations of Context and Brand Safety ad verification functions will be applied at the discrete domain, site or URL (property) level and content level for both display and video for both measurement and reporting granularity. Domain, site or URL (and even mobile application) level ad verification functionality and reporting should continue to follow the guidance in the original 2012 IAB/MRC *Ad Verification Guidelines*, but content level ad verification functionality and reporting should apply those guidelines as well as the guidance within this Supplement.

It is neither expected nor required that ad verification organizations develop Context and Brand Safety capabilities in all areas discussed in this Supplement in order to achieve accreditation. For example, ad verification organizations may not develop capabilities for or choose to submit for accreditation aspects such as mobile in-application Brand Safety functionality or Adjacency (discussed and defined later in this document). However, to the extent these functionalities exist and are assessed for compliance, all applicable aspects of corresponding sections of this document will be required.

Further, ad verification organizations must clearly disclose specific capabilities and limitations and MRC’s accreditation considerations and decisions related to Context and Brand Safety ad verification functions will be made with respect to specific capabilities. It is permissible for ad verification organizations with content measurement granularity to report avoidance and verification at the property level with proper disclosure. Such property level reporting with content level measurement granularity will be distinguished by MRC from property level reporting without discrete corresponding content measurement.

Finally, MRC may require some overall minimum level of capability or functionality in order to consider extending accreditation to ad verification organizations for Context and Brand Safety accreditation at the content level and this will be determined as part of scoping discussions between MRC, our auditors and audited ad verification organizations.

The following are the anticipated MRC Brand Safety and Context accreditation distinctions:

MRC Brand Safety and Context Accreditation Distinctions			
Measurement/Reporting Granularity Level	Property/Property	Content/Property	Content/Content

Privacy Considerations for This Document

Ad verification organizations seeking MRC accreditation are required to adhere to relevant MRC Minimum Standards related to respondent anonymity and privacy as well as privacy regulations. While platforms and publishers are encouraged to pass detailed collected first-party data including UGC (pre and post-processing) supporting ad verification measurement to third-party verification organizations to enable independent measurement including Adjacency, obfuscated or truncated data may be maintained or passed to satisfy this requirement, should there be Personal Identifying Information (PII) or privacy concerns. This should be available in a

transparent manner to accreditation auditors and at a detailed level to allow granular measurement reprocessing of reported data where necessary.

Different metric/transaction types and varying risks associated with transaction types should be considered. PII legal restrictions may dictate eliminating one or more of collected fields from externally transmitted records or altering the content of fields for identity protection purposes. Such restrictions may still allow for alternative levels of data transmission that are still sufficient to support granular processing of ad verification data. In these cases deviations should be supported by the platform's or ad verification organization's privacy policy in relation to privacy regulations and requirements and should be available for review by auditors.

1.1 About the Development of this Supplement

The guidance contained in this Supplement originated from a project led by the Media Rating Council (MRC) and sponsored by the American Association of Advertising Agencies (4A's), the Association of National Advertisers (ANA) and the Interactive Advertising Bureau (IAB U.S.). This guidance was developed with the participation of a core advisory group of buyer-side organizations, Industry associations, large digital enterprises and verification vendors, as well as a large group of Digital media content providers, advertising agencies, advertisers, vendors/consultants, measurement organizations and other interested organizations. This guidance was also reviewed and approved by major buyer-side trade organizations (4As, ANA) and their constituents and thereafter provided to the public through a formal period of public comment prior to adoption.

The final guidance is published and available on the MRC site and will be re-assessed periodically to ensure it remains applicable over time.

2 Content Level Granularity

Third-party ad verification functions have traditionally been focused on domain, site, URL or application (including "channels" within domains) level information (i.e., determining the safety or classification of a URL based largely on text analysis of the content) as a whole. For purposes of this document, this will be referred to as property level.

This focus was necessary to enable ad serving to standard pages; however, the current environment dictates Context and Brand Safety guidance be put in place that accounts for ad serving within content as well as to URLs or mobile applications with dynamic content (content that is frequently added or changed including by users outside the control of the publisher) and channels whereby property level classifications may not completely or directly correlate with specific pieces of content present. Such content level measurement depth must extend beyond keyword analysis and include consideration of images, video, text and audio content.

Ad verification organizations (third-party measurers, platforms or publishers) must distinguish Context and Brand Safety functionality between property and content level granularity or depth of measurement and reporting. If it is expected that a specific domain/site or grouping of

URLs/applications will exhibit the same taxonomic attributes as materially all of the content within them on an ongoing basis as demonstrated through periodic analysis and empirical support, Context and Brand Safety classifications may be consistently applied across all content within the grouping. **However, to the extent the nature of the content (display or video) within the property could materially differ from its specific property level classification, Context and Brand safety decisions must also be made at the discrete content level.**

However, this should be in addition to continuing to provide property-level classifications. Content granularity is particularly important when content may be considered “unsafe” or categorized in a particular avoidance category where the general property level classification is otherwise considered safe or acceptable. However, properties that are considered unsafe or are classified in a specific avoidance category should continue to be classified in this manner even if safe or acceptable content appears on them. This includes dynamic publisher curated content as well as UGC on Social Media Platforms.

It is permissible for ad verification organizations with content measurement granularity to report avoidance and verification at the property level with proper disclosure. Such property level reporting with content level measurement granularity will be distinguished by MRC from property level reporting without discrete corresponding content measurement. Further, it is permissible for ad verification organizations to classify content or properties based on URL keywords without content level measurement, however this must be clearly disclosed and does not represent content level measurement granularity or reporting.

With the issuance of this guidance, to the extent applicable, MRC’s accreditation considerations of Context and Brand Safety ad verification functions will be applied at the property and content level for both display and video. Property level ad verification functionality should continue to follow the guidance in the original 2012 IAB/MRC *Ad Verification Guidelines*, but content level ad verification functionality should apply those guidelines as well as the guidance within this Supplement.

While the ad verification tools and functionality used for each respective level of Context and Brand Safety may be similar, they will need to be applied differentially in conjunction with additional tools discussed throughout this Supplement in order to be effective for content level granularity. Ad verification organizations must actively disclose capabilities and limitations of measurement related to property vs. content level granularity to users along with guidance for proper use of measurement output considering any material limitations of measurement.

Further, it is neither expected nor required that ad verification organizations develop Context and Brand Safety capabilities in all areas discussed in this Supplement in order to achieve accreditation.

3 Social Media and UGC

General

The current digital advertising environment dictates Context and Brand Safety guidance also be put in place for dynamic User Generated Content (UGC) on Social Media Platforms including news feed environments. Timing and scale present challenges in large digital enterprises, especially for video, and the needs to assess UGC and post that content timely create additional diligence challenges. The volume of content likely necessitates automation (including Machine Learning), which can be subject to error, misinterpretation and ambiguity. A robust set of risk-based criteria is necessary to drive focused Human Intervention.

The presence of UGC further presents numerous challenges for Context and Brand Safety measurement including reduced control of content when compared to publisher-curated content. The introduction of comments such as in a Threaded Message may present additional data points to analyze UGC, but may also change the Context and Brand Safety classification of the Threaded Message when considered collectively. The volume and dynamic nature of such content presents scalability issues for measurement. **The ability to scale measurement at a granular level and to enforce data freshness rules is critical to effective measurement of content for purposes of Context and Brand Safety.** Ad verification organizations seeking content level accreditation must include consideration of UGC and the unique nature of platform specific aspects of UGC in data collection methodologies and ensure systems are comprehensive enough to handle scale and periodic refreshing. See further guidance on scalability later in this section.

Additionally, the categorization of the content itself may differ from the user activity associated to it in a Threaded Message and the occurrence of certain keywords in comments may assist ad verification organizations with content classification. While comments should be considered and included in content level Context and Brand Safety functionality, ad verification organizations should also take care to appropriately consider the weight comments carry on classifications as well as the frequency of certain keywords that may be classified differently than the content itself. The specific consideration and inclusion of user comments in content level Context and Brand Safety determinations must be empirically supported and disclosed, but represents an important signal that should be considered.

Ad verification organizations may utilize external third parties with additional expertise in content classification (such as professional editorial personnel and services and academic resources) to further mitigate the risk that general personnel or surface level review may not properly identify or categorize content. Such involvement of expert external third parties is encouraged and if used, should be leveraged to calibrate Machine Learning parameters as well as to focus Human Intervention as discussed below. Use of third parties must be disclosed (although the specific third parties need not be disclosed) and incorporated into audit or independent validation of ad verification functionality.

The ability of third party vendors to accurately measure activity on certain platforms may be limited in certain situations without direct tagging or Software Development Kit (SDK) integration in mobile applications. As a result, platforms and publishers are strongly encouraged to either allow direct third party measurement at a content level (inclusive of Metadata discussed below), or to actively provide validated first-party data at a granular level to enable independent Context and Brand Safety measurement, one of which will be required for content level accreditation of third-party ad verification solutions. Further, the collection and transmission of data enabling content level Context and Brand Safety determinations is strongly encouraged to be subjected to audit and independent validation.

To the extent that a third party vendor is not able to directly measure certain aspects of content within a property for Context and Brand Safety, the vendor may list the property or portions of the property as uncategorized or unknown with clear disclosure of capabilities or limitations with respect to specific property or content level granularity to enable avoidance if a user so chooses. Further, in these situations, properties that allow UGC may be generically classified as such (and be distinguished as unknown with regard to Brand Safety) to enable a user to choose whether or not to place ads in these environments in cases where content level decisions cannot be made. However, ad verification organizations will not be considered to offer content level Context and Brand Safety granularity for properties that are labeled generically or uncategorized where these measurement limitations exist.

In certain pre-bid situations, Demand Side Platforms (DSPs) may have a material impact on Context and Brand Safety data and determinations. In these cases, the guidance related to data transparency, quality and validation contained throughout this section should be applied to DSPs and considered by ad verification organizations utilizing DSP data.

3.1 Data Freshness and Quality

As discussed in the *Ad Verification Guidelines*, as content evolves in a very fast and dynamic manner, fresh data used to enable ad verification measurement is absolutely necessary for accurate ad verification. For example, the effectiveness of Context and Brand Safety measurement that utilizes dynamic UGC may diminish over time if data is not periodically refreshed and recollected. Therefore, the timing of the verification service and underlying data assumptions must be disclosed and ad verification services must have stated data refresh policies. A time stamp or active and clear disclosures of data life and collection policies is recommended to inform users about the age of the service's most recent assessment.

Given the dynamic nature of content on Social Media Platforms, UGC and surrounding user activity, it is essential that the frequency of data collection and refresh policies (including maximum data life or "time to live" policies) employed by an ad verification organization for content level classifications are adapted to the specific platform being measured (considering scale, user activity and presence of UGC as well as publisher policies and controls) and that these policies are based on empirical support that is periodically analyzed. These policies may

also be dynamic to account for content, such as news content, which may initially be dynamic and then more static over time.

Further, ad verification organizations are encouraged to establish minimum reporting thresholds (a required minimum number of data points or observations before reporting Context and Brand Safety determinations), collection periods or other baseline guardrails before making Context and Brand Safety determinations for specific content. This may consist of a minimum set of parameters or data fields for determinations based on static data or a minimum number of observations if based in dynamic data. Minimum baseline guardrails can be based on intended minimum statistical reliability such as confidence intervals or meaningful deviations in data and may vary depending on tolerable risk (as stipulated by users or as defined for a particular channel or category). For example, premium channels discussed below may carry a higher level of perceived safety and necessitate higher reliability and minimum reporting thresholds due to lower expected advertiser risk tolerance.

Content not meeting baseline guardrails must be classified as unknown or not yet classified and buyers should have the ability to prevent ad serving to unknown or unclassified content. Collections policies and frequencies as well as baseline guardrails employed must be disclosed to users.

3.2 Curation of UGC

As discussed above, UGC presents numerous challenges for contextual and Brand Safety measurement including reduced control of content when compared to publisher curated content, which may lead to differences between URL/application and content classifications. In order to mitigate some of the risks associated with UGC, publishers or platforms may curate UGC by selecting and vetting specific pieces of content and grouping them into categories, URLs or channels within common subject matter, Context and even Brand Safety classifications (such as for children or certain age groups). In certain situations, curation might be based on automated functionality, but most likely requires manual and human review. Further, such curation might be based on popularity or user activity not necessarily correlated to Brand Safety, such as “premium” pages or channels reserved for content contributors and users with the greatest number of followers, subscribers, views, likes or other activity.

While publishers and platforms may conduct curation primarily for user experience, they are also encouraged to consider doing so in order to assist buyers and third party measurers in making Context and Brand Safety determinations and to expose such curation processes to auditing and independent validation. However, while ad verification organizations may consider and leverage this curation when measuring Context and Brand Safety, it should not be relied upon unless independently audited and verified. Given the dynamic nature of UGC and user activity in a Threaded Message, even curated content requires discrete and robust measurement for effective Context and Brand Safety determinations.

Further, premium pages or channels curated based on popularity or user activity alone may not necessarily render content within them Brand Safe and likely require enhanced measurement focus, more robust human review and periodic monitoring due to the fact that this content may be more highly monetized and trafficked. While user activity such as number of followers, subscribers, views and likes can be used as a means to control monetization it should not be solely relied upon for Context or Brand Safety determinations.

3.3 Scalability

As discussed above, the volume and dynamic nature of content on Social Media Platforms and other large digital enterprises presents scalability issues for measurement. The ability to scale measurement at a granular level and to enforce data freshness rules is critical to effective measurement of content for purposes of Context and Brand Safety. Ad verification organizations must ensure systems are comprehensive enough to handle scale and periodic refreshing.

Incomplete or corrupt data can lead to inaccurate or insufficient Context and Brand Safety determinations. As a result, in addition to user access, program change and disaster recovery/business continuity controls that are required as part of MRC accreditation audits, ad verification organizations should also employ robust load balancing as well as capacity and uptime monitoring procedures to ensure adequate system resources are available for data collection and backend processing of large data volumes. Further, ad verification organizations are encouraged to utilize batch processing and to conduct periodic data corruption and latency checks to ensure data collection is complete and to identify, investigate and limit data loss. Such controls should be applied to systems that include, but are not limited to, Machine Learning functionality as well as ingestion and processing of Metadata (each is discussed in greater detail in this Supplement).

3.4 Machine Learning

The scale and array of data fields involved in content level Context and Brand Safety functionality likely necessitates some form of automation or Machine Learning utilized by ad verification organizations whereby models or algorithms are used to automatically ingest, analyze and classify content for purposes of Context and Brand Safety. Machine Learning consists of utilizing datasets including visual and audio data points (such as in visual and audial learning) to train and evaluate models or algorithms such that they can predict outcomes on an automated basis. It is critical that rigorous data quality and analysis procedures are applied to the processes to select training data, select parameters used in the model, prepare data, select the model, train and evaluate the model as well as to periodically tune it.

In addition, the *MRC Invalid Traffic Addendum* establishes requirements for detection and filtration of invalid traffic (IVT) and additional IVT guidance promulgated by the MRC establishes best practices for pre-bid approaches to IVT (which are likely highly relevant in ad verification functionality). Such filtration should be applied to datasets used for Machine Learning to the extent applicable to ensure IVT does not introduce biases in the classification decisions or

otherwise obfuscate data trends that may signal a need for further analysis such as user viewing and skipping behavior that may direct further analysis of a specific segment of long-form video that may be diluted by the presence of invalid traffic with no such behavior. Further, the presence and volume of IVT within a property or associated with a particular piece of content may be a signal that in and of itself assists in Context and Brand Safety determinations or at least allows an organization to further focus resources. The impact of IVT on Context and Brand Safety determinations must be considered when designing Machine Learning functionality. While IVT should be filtered from advertising measurement, the presence (or lack thereof) of IVT may not directly have a bearing on the Context and Brand Safety classification of a particular piece of content, although ad verification organizations may consider alerting users to properties with high levels of IVT for potential avoidance.

The ability of Machine Learning to accurately predict outcomes is correlated to the size and quality of the data underlying it. Ad verification organizations must establish minimum data sample sizes and quality thresholds when selecting training and evaluation data sets for Machine Learning. To the extent reduced performance or accuracy is expected in certain data conditions (such as in smaller sample sizes, shorter collection periods, missing or lower quality data, etc.), this must be actively disclosed to users via estimates of error using statistical methods or observed error.

Further, the data selected and the parameters or fields used should be relevant to the use of the model (empirical support must exist that establishes a relationship between data or parameters used and Context or Brand Safety determinations). When preparing training and evaluation data, which must consist of distinct and mutually exclusive data sets, robust data quality procedures should be applied to clean, normalize and deduplicate data as well as to account for or adjust data imbalances or biases. Biases in training and evaluation data should be reduced to the extent possible and disclosed where material.

Additionally, the model or algorithm selected for specific Machine Learning functionalities must be appropriate for the intended use and periodically updated or refreshed. Use of Machine Learning models must be logical and defensible based on empirical support that is documented and made available for audit. Such support must be periodically validated and updated. Training and evaluation of the model will likely result in weighting or bias adjustment as well as parameter tuning, and this must also occur on a periodic basis.

While Machine Learning will help mitigate certain challenges discussed throughout this Supplement related to scalability and dynamic content, it should be deployed in conjunction with focused Human Intervention. Weights and bias adjustments resulting from the periodic training and evaluation of Machine Learning models should be used to focus such Human Intervention. In other words, known biases or weaknesses in the model or for specific data conditions should inform procedures that include manual human review of content for Context and Brand Safety classifications. Additional guidance for Human Intervention is discussed later in this Supplement.

Machine learning methodology and procedures must be actively disclosed to users at a non-technical level, including relevant data set sizes, models used and data sources while ensuring proper safeguards against reverse engineering of Brand Safety techniques employed. The level of reliance on machine learning versus human intervention or manual review must be generally disclosed.

3.5 Human Intervention

As discussed above, Machine Learning should be deployed in conjunction with focused Human Intervention defined for purposes of this Supplement as manual review of content by ad verification organization personnel for the purpose of Context and Brand Safety classifications. Weights and bias adjustments resulting from the periodic training and evaluation of Machine Learning models should be used to focus such Human Intervention. Even where Machine Learning is not utilized, a robust set of risk-based criteria is necessary to drive focused Human Intervention. Such risk-based criteria should include considerations of data conditions or content types that are more difficult to measure or categorize, the presence of monetized UGC and dynamic content, and content with higher traffic including premium pages or channels curated based on popularity. The presence and levels of IVT should also be considered when determining the focus of Human Intervention.

As Human Intervention involves manual review of content, consideration should be given to time constraints and scalability. In conjunction with risk-based considerations, Human Intervention policies should be designed with content volume and the need for timely review and classification in mind in order to ensure staffing levels are adequate to meet workload needs based on reporting timing policies and volume.

The personnel performing Human Intervention must be adequately trained and supervised. The results of Human Intervention must be periodically reviewed and formally documented. All personnel (including supervisors) must be furnished with detailed instructions and manuals covering all steps of their work. Personnel performing Human Intervention must be periodically assessed for performance. Lower performing personnel should be re-trained and their work in production environments should be limited.

The results of Human Intervention should be used to periodically validate the results of Machine Learning and to update evaluation data sets. It is expected that Human Intervention be used as an ongoing and continuous quality control, but that results are used to update Machine Learning as frequently as is feasible and at least several times annually.

3.6 Monetization Controls

Publishers or Social Media platforms may apply monetization controls, or thresholds above which content will be monetized (and below which it will not be). These thresholds can be based on the age or recency of content and the activity surrounding it such as interactions, views, subscribers or view time. Such controls may help mitigate the risk of incomplete or inaccurate Context or Brand Safety classifications in that they extend the time before which

content is monetized to allow more thorough review and a larger accumulation of data to analyze. Publishers and platforms are encouraged to utilize such controls and to disclose them to users. Further, like Human Intervention discussed above, these monetization controls should be driven by a robust set of risk-based criteria.

Additionally, publishers, ad networks, exchanges or Social Media Platforms should consider establishing policies to limit or prevent monetization of content for sources that have been found to violate content policies, specific content that has been determined to be in violation of controls or other Brand Safety controls and these policies should be applied at a platform level. These policies if present, must, be clearly defined and disclosed to users and sources that can contribute content. In accordance with the IAB/MRC *Ad Verification Guidelines*, ad verification processes must be transparent to customers and seller organizations, in that both know the general criteria being used for evaluation as well as any content or sources of content such as users or publishers that are excluded (details of specific proprietary methods may be omitted).

The presence and levels of IVT should also be considered when determining which parameters to use for monetization controls and should be utilized on a post-filtration basis. The ability for a specific parameter to be manipulated through invalid traffic should decrease reliance on it for purposes of monetization controls.

Like curation of content discussed above, publishers and platforms are encouraged to expose monetization controls to auditing and independent validation. Again, while ad verification organizations may consider and leverage monetization controls when measuring Context and Brand Safety, they should not be relied upon unless independently audited and verified. Finally, publishers, platforms and ad verification organizations are encouraged to enable buyers to customize monetization controls based on contextual or Brand Safety needs. Enhanced buyer tools are discussed later in this Supplement.

While user activity such as number of followers, subscribers, views and likes can be used as a means to control monetization, it cannot be solely relied upon for Context or Brand Safety determinations.

3.7 Taxonomy and Categorization

The IAB Network and Exchanges “Contextual Taxonomy” (now the *IAB Tech Lab Content taxonomy*) were incorporated into the *Ad Verification Guidelines* by reference as a listing of appropriate content categories for sites (although recent updates to this taxonomy have deprecated some predecessor avoidance categories). Additionally, the IAB/MRC *Ad Verification Guidelines* include a listing of general types of potential avoidance categories and it is expected that ad verification organizations may maintain more detailed sub-categories for each of these potential avoidance categories, which will help further refine and differentiate their service.

While ad verification organizations are encouraged to use common Industry sources for genre classification and segmentation, they may also use internally developed or proprietary sources

if empirically supported with auditable evidence with clear disclosure. The source and taxonomy used in genre classification (including custom proprietary sources) for purposes of Context and Brand Safety classifications must be disclosed to users and periodically updated. Such disclosures should consider protection of intellectual property and the risk of reverse engineering that could potentially result.

Further, dynamic UGC on Social Media Platforms may require granular and discrete avoidance categories and presents further challenges for accurate classification, but the nature of this content may also introduce subjectivity regarding the classification of a particular piece of content. For example, video content may contain segments with differing categorizations and ad verification organizations are encouraged to provide granular segment categorization data in addition to top-level classifications where this may occur, such as in long-form digital video (especially for episodic and repurposed TV content where existing categorization practices may already exist).

This Supplement does not aim to proscribe specific taxonomies or categories beyond those referenced above, but instead to highlight the unique categorization challenges UGC on Social Media Platforms presents, and to require ad verification organizations to disclose the methodology and decision-making hierarchy underlying categorizations to users. Additionally, to the extent that custom criteria, partial or multi categorization for a specific piece of content, or degrees of categorization (such as content ratings or scores) are used, they must also be actively disclosed to users as well as empirically supported and periodically updated.

MRC strongly encourages and supports mechanisms and efforts to promote consistent Brand Safety definitions and minimum baselines such as that proposed by the 4A's Advertising Protection Bureau (APB) as well as other buyer organizations. This Supplement recommends establishing a minimum level of content established and maintained by industry groups that excludes certain baseline categories. This minimum can be offered as a collective reporting or blocking group considered unsafe or ineligible for monetization ("the floor") either as a default control for platform monetization or as elected by advertisers. This floor should be based on exclusion of certain categories rooted in the Industry's "dirty dozen" category list. Specifically, the 4A's Advertising Assurance Brand Safety Floor Framework issued September 2018 includes baseline definitions of a floor within 13 categories.

These categories likely require ongoing consideration and definition and the categorization of a piece of content (for example, what constitutes an obscenity) may differ from the point of view of different ad verification users. The 4A's Advertising Assurance Brand Safety Floor Framework aims to codify aspects of the above that are less subjective and can be agreed to be part of a consistent exclusion category to establish the floor. The expectation is that platforms and publishers will enact controls to minimize the presence of content in these categories. Where content in these categories is present advertisers may choose to forgo the agreed upon floor and allow advertising to one or all of the above categories.

Further, the 4A's APB will include ongoing efforts to review existing taxonomies and determine how consistency can be created. This includes reconciling definitions of common category differences between vendors, while allowing for custom settings that vendors apply to address specific clients Brand Suitability needs (such as via the APB's Brand Suitability Framework). MRC further recommends and supports Industry efforts to standardize and modernize taxonomies and categories across platforms considering UGC and content granularity, and we intend to participate in such efforts on an ongoing basis.

4 Adjacency

For the purpose of this Supplement, Adjacency is defined as the degree of physical distance in which ads are placed in relation to specific content that is classified for Context and Brand Safety purposes within a domain, site, platform, mobile application or page. As discussed above, historically Brand Safety and targeting have been addressed at a property level with Context classification and ad alerting or blocking based on avoidance or target categories. However, the nature of Social Media Platforms and various newsfeed environments require more discrete proximity measurement for both targeting and avoidance including adjacency in newsfeed, continuous scroll or swipe (vertical or horizontal) or continuous play environments. **While measurement of Adjacency is not required, it is encouraged and an creditable aspect of content level Context and Brand Safety functionality.** The following guidance must be applied to ad verification organizations electing to offer adjacency measurement.

Adjacency and measurement of it will differ between platforms, URLs and mobile applications as will a user's determination of whether a specific piece of content is near or far from an ad placement. **To that end, publishers and platforms are encouraged to establish measurement units in determining adjacency at the platform, URL or mobile application level, where applicable that can be measured, reported and verified by third parties.** For example, the measurement unit within a vertical or horizontal newsfeed environment could be a discrete pane, tile, screen or other unit. The measurement unit should reflect the meaningful segmentation of content within a platform, page or mobile application and must be clearly disclosed as part of adjacency reporting, and documented and supported. As a baseline, adjacency measurement units should be discrete pieces of content (such as an image, video, song, ad etc.) and all measurement bases must be meaningful to content and ad consumption and empirically supported.

Ad verification organizations are encouraged to consider the intersection of time and adjacency via temporal thresholds applied to adjacency. The impact of adjacent content may diminish during longer user sessions such as in situations when users leave content and return to it after the passage of time.

Adjacency should be measured and reported on the basis of measurement unit distance and must account for scrolling and swiping capabilities in mobile newsfeed environments. Adjacency could be reported in terms of absolute value of measurement units (i.e., number of units placement proceeds/is above or follows/is below specific content being measured). The

basis for reporting must be clearly disclosed. **Placements that are zero measurement units from specific content being measured or that are able to be simultaneously viewable (50% of pixels on screen consistent with *MRC Viewability Guidelines*) with specific content being measured should be considered directly adjacent.** Otherwise, adjacency ranges or scores may be reported and must be clearly defined.

In video playlist environments, adjacency should be measured in terms of the number of video units that play before and after the specific content being measured including content that automatically begins playing without additional user interaction. Consideration must also be given to the presence and categorization of thumbnail images accompanying or related to video content, especially where these images may be misleading or differ from the nature of the underlying video content.

As discussed earlier in this Supplement, publishers or Social Media platforms may apply monetization controls and this should include adjacency considerations (a minimum distance or number of measurement units below which ads will not be placed). Such minimum adjacency rules may vary by content. Again, publishers and platforms are encouraged to utilize such controls and to disclose them to users, but while ad verification organizations may consider and leverage monetization controls when measuring Context and Brand Safety, they must not be relied upon unless independently audited and verified.

Third party ad verification vendors may have to rely on signals related to adjacency (such as units, scores of placements or session-level content information) from platforms or publishers where the third-party is restricted from directly measuring adjacency either due to measurement policies or technical challenges. Further, in these situations, only the ad server, often times the platform itself, can allow blocking, covering or otherwise avoidance based on adjacency.

As a result, platforms and publishers are strongly encouraged to either allow direct third party measurement at a content level inclusive of adjacency signals, or to actively provide validated first-party data at a granular level to enable independent Context and Brand Safety measurement alerting and reporting, one of which will be required for content level adjacency accreditation of third-party ad verification solutions. Further, the collection and transmission of data enabling adjacency determinations is strongly encouraged to be subjected to audit and independent validation. Publishers, platforms and ad verification organizations are encouraged to enable buyers to customize adjacency based on contextual or Brand Safety needs. Enhanced buyer tools are discussed later in this Supplement.

5 Metadata

Various ad verification techniques and processes discussed throughout this Supplement may utilize publisher or platform Metadata, or data that provides information about the makeup of a domain, site, URL or application including specific content details. Metadata typically includes descriptive data such as a general description, author and keywords, structural data that

includes page and content organization, and administrative data such as when a file was created, by whom, content type and other technical information. This may include content size and length, appearance characteristics (such as color and shape), transcribed audio for video content and associated user interaction attributes such as comments, views, exposure and engagement data or number of subscribers. Metadata should also include descriptors of thumbnail images accompanying or related video content, especially where these images may be misleading differ from the nature of the underlying video content.

As discussed earlier in this Supplement, platforms and publishers are strongly encouraged to either allow direct third party measurement at a content level (inclusive of Metadata), or to actively provide validated first-party data at a granular level to enable independent Context and Brand Safety measurement, one of which will be required for accreditation of content level third-party ad verification solutions. The collection and transmission of Metadata enabling content level Context and Brand Safety determinations is strongly encouraged to be subjected to audit and independent validation.

As discussed in the Machine Learning section of this Supplement, certain Metadata parameters may be more directly correlated with Context and Brand Safety determinations than others and this correlation may vary over time or with the size of the observed dataset. In addition to guidance above related to scalability and Machine Learning (which should be applied to Metadata), ad verification organizations utilizing Metadata must consider limitations of this data when using it to make Context and Brand Safety determinations.

Further, certain Metadata is static and not subject to change, while other Metadata, such as that corresponding to user interaction and exposure data (which may be used for monetization controls) is dynamic. As discussed in the *Ad Verification Guidelines* and above, fresh data is absolutely necessary for accurate ad verification. Data freshness rules are critical to effective measurement of content for purposes of Context and Brand Safety.

Metadata used for Context and Brand Safety determinations may include user reporting or flagging of objectionable content. This parameter can be a valuable input into ad verification functionality, but due to the subjective nature of such reporting, should be appropriately weighted in the decision making process based on review and robust quality control.

Automated analysis of Metadata alone without additional data sources or consideration of human and manual processes is likely insufficient for certain content level Context and Brand Safety determinations (such as video content) and ad verification organizations are encouraged to utilize other data and techniques in addition to Metadata. The degree of reliance on Metadata in Context and Brand Safety determinations by an ad verification organization must be disclosed to users. Such disclosures should consider protection of intellectual property and the risk of reverse engineering that could potentially result. Use-cases and examples of tools used are encouraged to be a part of technical disclosures made by the ad verification organization. The specific Metadata collection methods and content recognition

processes used by the ad verification organization must be disclosed at a high level and may include, but are not limited to:

- Application Programming Interfaces (APIs)
- Direct Measurement
- Machine Learning
- Offline Export
- Partner Data
- Scraping/Crawling (where permissible by terms and conditions)
- Third-Party Tracking
- Meta Tags

As discussed in the *Ad Verification Guidelines*, ad verification services may be integrated into ad campaigns using some form of tracking asset (ad tags, beacons, etc.). Additionally, passive (non-integrated) analysis tools such as automated crawlers or spiders may be used. The limitations associated with these methods must be highlighted in disclosures – for example, the non-random nature of observations, non-census approaches, or the impact on the ability to project results.

Data collection methods must be subject to internal testing and validation upon original implementation as well as periodic internal verification testing. This testing should include the validity of the Context and Brand Safety measures. Data collection methods and disclosures will be the subject of intense procedural verification if MRC accreditation is sought by the measurement organization.

The Metadata source must be clearly defined and the extent of Metadata use must be included in the ad verification organization’s disclosures, for example:

- What type of content is captured from what platform;
- User coverage for measured platforms;
- Collection method (direct, API, offline export, partner, third-party, scraping/crawling, etc.);
- API type where applicable (public, closed client, platform or fire hose);
- Depth of content collection (scraping) applied;
- Storage methods, timing;
- Nature of machine (or human) analysis processes employed;
- Time periods analyzed;
- Edits or data adjustments applied;
- Nature of special handling of ambiguous cases;
- Error correction and reissue procedures; and
- Context and Brand Safety accuracy.

Different languages and language-related nuances should be considered when analyzing metadata including dialects and processes to ingest and analyze metadata should be designed to account for different or multiple languages present and where applicable.

6 Mobile Applications

While the guidance provided in the *Ad Verification Guidelines* and within this Supplement should be fully applied within mobile applications, there may be unique Context and Brand Safety considerations within these environments. The risks and techniques relevant to mobile in-app environments related to Context and Brand Safety exhibit differential characteristics when compared to those employed in desktop or mobile web environments. Moreover, measurement assets utilized in desktop or mobile web environments such as crawlers, JavaScript, cookies or Flash may not be available or functional within mobile applications.

In addition to specific considerations detailed in the *IAB/MMA/MRC Mobile Application Advertising Measurement Guidelines*, ad verification organizations measuring Context and Brand Safety should apply incremental and differential consideration to mobile in-application (in-app) environments and they must do so for distinct mobile in-app content level Brand Safety accreditation. Such consideration must be focused on meaningful differences within mobile applications regarding initial and ongoing risk assessments, user behavior, differences in content, heuristics and signals used to make Context and Brand Safety determinations and ongoing data analysis/benchmarking.

In addition to ad verification functionality at an application level, ad verification organizations seeking accreditation for mobile applications must apply content level granularity (see guidance earlier in the Supplement) discretely within mobile applications and not based solely on descriptions of an application, application store information or traffic to commonly owned web properties. While classification or avoidance of applications based on their descriptions or general content are permissible (app level akin to property level) with clear disclosure, this would not be considered content level in-application measurement granularity for purposes of this Supplement. MRC intends to consider Brand Safety accreditation discretely between desktop/mobile web and mobile in-app environments.

MRC recognizes that discrete content level Context and Brand Safety functionality within mobile applications may be aspirational as many applications do not support API functionality and while it is not required that ad verification organizations develop mobile in-app capabilities in this area, it is highly encouraged and required in order for MRC to grant specific content level in-app accreditation for this functionality. Alternatively, applications without API support enabling content level measurement granularity may be classified as “unknown” for Brand Safety purposes.

The MRC believes there may be value in the development of a common and open standard API or SDK by the industry that can be used by all parties to enable in-app Context and Brand Safety measurement via ad verification organizations. We encourage such development and

would support and participate in an open and single source technical standard. It should be noted that as of the date of the issuance of this document, the MRC is not aware of any active initiatives to address a common industry API or SDK and as such, it is likely that content level ad verification measurement solutions in applications will be custom or specific to certain ad verification organization’s platform and publisher integrations for the foreseeable future.

7 Illegal Sources and Industry Communication

Ad verification organizations must make reasonable attempts to consider legal restrictions related to prohibition of content and the sources of that content to prevent monetization of illegal content. Specifically, certain types of content or sources of content may be expressly prohibited by legal restrictions such as content piracy involving copyright violations and counterfeiting, or there may be other requirements to remove illegal content within a specified time period (such as Germany’s Network Enforcement Act). Publishers and platforms have a responsibility for safeguarding their properties from illegal entities and sources and should have a qualification process to make sure they are dealing with a legitimate entity and legally permissible content that is applied across the platform, as well as policies to comply with required enforcement periods, even prior to applying additional Brand Safety controls. While non-publisher third-party measurement organizations are generally not in control of content source, they must be aware of legal restrictions and include considerations of this in Content and Brand Safety measurement and determination.

Procedures related to determining the legality of sources and content should include initial qualification using Industry and local sources of known illegal entities, as well as ongoing evaluation linked with ad verification results and periodic internal auditing of content sources. Ad verification organizations seeking accreditation will be required to provide evidence of source vetting processes where applicable during accreditation audit processes.

Further, to the extent that ad verification organizations have identified illegal or illegitimate sources that are either not included in Industry or local sources of known illegal entities, or that are disguised as other legal and legitimate entities, processes should be put in place to routinely communicate these sources to legal authorities, oversight bodies and the Industry at large. Such communications should be conducted with sensitivity to protection of intellectual property and the risk of reverse engineering that could potentially result. Decisions to forgo communication by ad verification organizations must be supported by auditable evidence of such risk.

MRC strongly encourages and supports mechanisms and efforts to communicate relevant Brand Safety information such as observed monetization of “unsafe” content across the Industry and publicly such as that currently proposed within the 4A’s Advertising Protection Bureau (APB). This may include open and transparent listing of Context and Brand Safety categorizations for properties as well as content URLs via emerging technological means such as block chain.

8 Enhanced Buyer Tools

This Supplement details various tools verification organizations may utilize to classify and categorize content for context and Brand Safety determinations, as well as proactive approaches that may be taken by publishers or platforms to restrict ad placement with the ultimate goal of meeting an advertiser's specifications on protecting its brand. These approaches include curation, monetization controls and granular categorization. However, the nature of dynamic UGC on Social Media Platforms may introduce additional subjectivity regarding the classification of a particular piece of content. Additionally, different buyers may not only have different targets and avoidance categories, but also different tolerance levels for risk depending on brand suitability.

While a minimum baseline of monetization control is encouraged, ultimately ad verification organizations seeking content level Brand Safety accreditation must provide consistent and discrete reporting to users in order to enable customized targeting and avoidance. Beyond current category based blocking and alerting, ad verification organizations may enable: ad prevention (pre-bid); blocking and alerting (post-campaign) on the basis of Metadata; discrete content characteristics including the age of content subscribers to a channel, the number of views and interactions to enable custom monetization controls or flagging/reporting related to the content, adjacency, etc.; and adjacency parameters to allow users to stipulate which levels meet the needs of a particular campaign or brand.

The basis of measurement underlying these enhanced buyer tools must be clearly disclosed, along with the methodology used to collect, process, edit and report such measurement. Further, any limitations in the underlying measurement or known biases must also be clearly disclosed to enable users to utilize these tools in the proper context. Ad verification organizations should establish empirically supported and clearly disclosed minimum requirements for reporting with these limitations and biases in mind.

9 Required Reporting

Results of ad verification (including Context and Brand Safety) should be reported either on the basis of the amount of traffic that is blocked or classified in relevant categories (in situations when the ad verification function is outside the ad chain) or on the basis of traffic net of Context and avoidance categories outside of those stipulated by the user (when part of the ad serving functionality). Reporting can be at the impression (or other relevant metric) level, within campaigns but may also consist of inventory or content classification. Reporting may include; (1) whether results are projectable to a campaign or not, depending on verification data collection method; (2) the reliability or standard error around projectable results where samples are used; and (3) likelihood of false positive results by service line based on the prior experience of the verification service. Additionally, known technical limitations of verification services, if any, must be explained. Beyond early confirmation and set-up error detection procedures, performance communication to ad verification service users must be ongoing via periodic reports or through an automated reporting dashboard.

It is permissible for ad verification organizations with content measurement granularity to report avoidance and verification at the property level with proper disclosure. Such property level reporting with content level measurement granularity will be distinguished by MRC from property level reporting without discrete corresponding content measurement.

To the extent that a third party vendor is not able to directly measure certain aspects of content within a property for Context and Brand Safety, the vendor may list the property as uncategorized or unknown with clear disclosure of capabilities or limitations with respect to specific property or content level granularity to enable avoidance if a user so chooses. Further, in these situations, properties that allow UGC may be generically classified as such (and be distinguished as unknown with regard to Brand Safety) to enable a user to choose whether or not to place ads in these environments in cases where content level decisions cannot be made. However, ad verification organizations will not be considered to offer content level Context and Brand Safety granularity for properties that are labeled generically or uncategorized where these measurement limitations exist.

Finally, it is permissible for ad verification organizations to classify content or properties based on URL keywords without content level measurement, however this must be clearly disclosed and does not represent content level measurement granularity or reporting. Further, such URL level reporting is likely unreliable and limited for Brand Safety purposes.

10 Use of Samples

Various techniques and controls are discussed throughout this Supplement, and ad verification organizations are encouraged to use more than one of these techniques and controls in conjunction with each other in order to achieve a high degree of confidence in Context and Brand Safety determinations. Ad verification organizations are encouraged to adopt complete census measurement using deterministic techniques where possible. **However, where samples are used in making Context and Brand Safety determinations (such as when using probabilistic or pre-bid approaches), such samples must be constructed with a confidence level of 99.7% (within 3 standard deviations).** Confidence levels used and standard error surrounding Context and Brand Safety determinations must be actively disclosed to users.

For sample-based measurement of any kind, the ad verification organization must be diligent about ensuring valid projections are made and that the sample is representative of the population targeted for measurement. Methods for weighting or adjusting data to ensure projectability must be supported by empirical study, and these empirical studies must be updated periodically. Standard errors around sample-based projections must be disclosed.

11 Glossary of Terms

Adjacency: The degree of physical distance in which ads are placed in relation to specific content that is classified for Context and Brand Safety purposes within a domain, site, platform, mobile application or page.

Ad Verification Organization: For the purpose of this Supplement, refers to any organization, be it a third party measurement vendor, a publisher or a Social Media Platform, providing Ad Verification services (measurement and/or reporting) including Context and Brand Safety.

Brand Safety: Practices and tools to ensure that a digital ad will not appear adjacent to or in a Context that can damage an advertiser's brand.

Context: Refers to certain attributes of a URL, site, mobile application or piece of content that are used as determinants as to whether the destination where a digital ad is to be served represents an appropriate environment in which the ad should appear, as determined by parameters set by the advertiser.

Curation: Selecting and vetting specific pieces of content and grouping them into categories, pages or channels within common subject matter, Context and even Brand Safety classifications (such as for children or certain age groups). In certain situations, curation might be based on automated functionality, but most likely requires manual and human review.

Human Intervention: Manual review of content by ad verification organization personnel for the purpose of Context and Brand Safety classifications

Machine Learning: Use of models or algorithms to automatically ingest, analyze and classify content for purposes of Context and Brand Safety.

Metadata: Data that provides information about the makeup of a domain, site, URL or application including specific content details. Metadata typically includes descriptive data such as a general description, author and keywords, structural data that includes page and content organization and administrative data such as when a file was created, by whom, content type and other technical information. This may include content size and length, appearance characteristics (such as color and shape), transcribed audio for video content and associated user interaction attributes (such as comments, views, exposure and engagement data or subscribership).

Social Media Platform: A media vehicle, site (as a whole or in part), app, widget or other media mechanism that has a social orientation; specifically, that captures user commentary and facilitates sharing of information among users of a defined network – i.e., representing a channel for sharing information, opinions or experiences from users of that platform – or allows cross-user collaboration.

Threaded Message: A set of communications, entries or events that are related to a common piece of original content.

User Generated Content (UGC): Content that is entered, copied-to, posted (or otherwise created) by users of a Social Media Platform for sharing with others on that Platform.

12 References

Reference List of Relevant Previously Released MRC Standards Industry Guidelines:

MRC Minimum Standards for Media Rating Research:

<http://mediaratingcouncil.org/MRC%20Minimum%20Standards%20-%20December%202011.pdf>

MRC Invalid Traffic Detection and Filtration Guidelines Addendum:

[http://mediaratingcouncil.org/101515_IVT%20Addendum%20FINAL%20\(Version%201.0\).pdf](http://mediaratingcouncil.org/101515_IVT%20Addendum%20FINAL%20(Version%201.0).pdf)

MRC Social Media Measurement Guidelines:

<http://mediaratingcouncil.org/MRC%20Social%20Measurement%20Guidelines%20v1.0%20Final.pdf>

MRC Desktop and Mobile Viewability Guidelines:

http://mediaratingcouncil.org/081815%20Viewable%20Ad%20Impression%20Guideline_v2.0_Final.pdf

<http://mediaratingcouncil.org/062816%20Mobile%20Viewable%20Guidelines%20Final.pdf>

IAB/MRC Guidelines for the Conduct of Ad Verification: <https://www.iab.com/wp-content/uploads/2015/06/Ad-Verification-Guideline-for-the-Conduct-of.pdf>

IAB/MMA/MRC Mobile Application Advertising Measurement Guidelines:

<https://www.iab.com/wp-content/uploads/2017/11/Mobile-In-App-Measurement-Guidelines-MMTF-Final-v1.1.pdf>

13 Supporting Associations and Participating Organizations

About the Media Rating Council (MRC)

The Media Rating Council is a non-profit Industry association established in 1963 comprised of leading television, radio, print and digital media companies, as well as advertisers, advertising agencies and trade associations, whose goal is to ensure measurement services that are valid, reliable and effective. Measurement services desiring MRC accreditation are required to disclose to their customers all methodological aspects of their service; comply with the *MRC Minimum Standards for Media Rating Research* as well as other applicable Industry measurement guidelines; and submit to MRC-designed audits to authenticate and illuminate their procedures. In addition, the MRC membership actively pursues research issues they consider priorities in an effort to improve the quality of research in the marketplace. Currently approximately 110 research products are audited by the MRC. Additional information about MRC can be found at www.mediaringcouncil.org

About the American Association of Advertising Agencies (4A's)

Founded in 1917, the 4A's was established to promote, advance and defend the interests of our member agencies, their employees and the industry at large. The organization serves 700+ member agencies across 1,300 offices, which control more than 85% of total U.S. advertising spend. As the leading trade organization for marketing communication agencies, the 4A's purpose is to help empower its members to deliver insightful creativity that drives commerce, and influences culture all while moving the industry forward. The organization provides community, leadership, advocacy, guidance and best-in-class training that help enable agencies to innovate, evolve and grow. 4A's Benefits division insures more than 160,000 employees and its D.C. office advocates for policies that best support a thriving advertising industry. The 4A's Foundation fuels a robust diversity pipeline of talent for its members and the marketing and media industry, fostering the next generation of leaders. The organization is dedicated to, and vested in, our members' success just as they are dedicated to helping brands create, distribute, and measure effective and insightful advertising and marketing. Visit the 4A's at <http://www.aaaa.org>.

About the Association of National Advertisers (ANA)

The ANA (Association of National Advertisers) makes a difference for individuals, brands, and the industry by driving growth, advancing the interests of marketers, and promoting and protecting the well-being of the marketing community. Founded in 1910, the ANA provides leadership that advances marketing excellence and shapes the future of the industry. The ANA's membership includes more than 1,000 companies with 15,000 brands that collectively spend or support more than \$400 billion in marketing and advertising annually. The membership is comprised of more than 750 client-side marketers and 300 marketing service providers, which include leading agencies, law firms, suppliers, consultants, and vendors. Further enriching the ecosystem is the work of the nonprofit ANA Educational Foundation (AEF), which has the mission of enhancing the understanding of advertising and marketing within the academic and marketing communities.

About the Interactive Advertising Bureau (IAB)

The Interactive Advertising Bureau (IAB) empowers the media and marketing industries to thrive in the digital economy. Its membership is comprised of more than 650 leading media and technology companies that are responsible for selling, delivering, and optimizing digital advertising or marketing campaigns. The trade group fields critical research on interactive advertising, while also educating brands, agencies, and the wider business community on the importance of digital marketing. In affiliation with the IAB Tech Lab, it develops technical standards and best practices. IAB and the IAB Education Foundation are committed to professional development and elevating the knowledge, skills, expertise, and diversity of the workforce across the industry. Through the work of its public policy office in Washington, D.C., IAB advocates for its members and promotes the value of the interactive advertising industry to legislators and policymakers. There are 43 IABs licensed to operate in nations around the world and one regional IAB, in Europe. Founded in 1996, the IAB is headquartered in New York City and has a San Francisco office.

Participating Working Group Organizations:

4As	Google	Newsweek
ACA	Grapeshot	Ole Communications
ANA	GroupM	OpenSlate
Annalect	Horizon Media	P&G
Apple	IAB	PepsiCo
AT&T AdWords	IAB Canada	Pfizer
Bank of America	Integral Ad Science	Pinterest
Buzzfeed	ISBA	Publicis Media
CIBC	ITN Networks	SIPI
comScore	JPMorgan Chase	Snap
Crown Media	LinkedIn	TAG
Deloitte & Touche	Media Management Inc.	Twitter
Digital Content Next	Mediaoptimise	Unilever
DoubleVerify	MMA	VAB
EY	Moat/Oracle	Vizio
Facebook	NAB	WWE
FOX Networks Group	NCM Media	

14 Contact Us

MRC:
 Ron Pinelli, VP Digital Research and Standards
 212-972-0300
rpinelli@mediaratingcouncil.org

Appendix A: Section Summaries

To recap, the major aspects of each Section 1-8 include (but are not limited to) the following:

Section 2 (Content Granularity) Summary (for ad verification organizations seeking content level Brand Safety accreditation)

- Context and Brand Safety functionality must include content level granularity
- Dynamic content and UGC must be discretely considered
- Differences in property and content level must be considered
- Content level granularity should be in addition to continuing to provide discrete property classifications
- Ad verification organizations with content measurement granularity may report avoidance and verification at the property level with proper disclosure
- Ad verification organizations may classify content or properties based on URL keywords without content level measurement, however this does not represent content level measurement granularity or reporting
- MRC's accreditation process will evaluate and distinguish between levels or measurement and reporting granularity.
- Capabilities and limitations at each level must be disclosed with guidance for use

Section 3 (Social Media and UGC) Summary

- Social Media Platform content and UGC necessitate unique and discrete considerations
- Dynamic content must drive data freshness and minimum baseline data guardrails
- Social Media Platforms and publishers are strongly encouraged to either allow direct third party measurement at a content level or to actively provide first-party data to third-party measurers
- Fresh data is absolutely necessary for accurate ad verification
- Minimum reporting thresholds should be employed
- Curation of content is encouraged and can be leveraged, but not relied on if unaudited
- Curation based on popularity or exposure necessitates enhanced focus and monitoring
- Volume requires consideration of ability to scale measurement and minimize downtime
- Machine Learning must be empirically supported and based on quality data
- The level of reliance on machine learning versus human intervention or manual review must be generally disclosed
- Models/algorithms underlying Machine Learning must be appropriate and supported
- Machine Learning should be deployed in conjunction with focused Human Intervention
- Human Intervention should be informed by Machine Learning and a risk-based approach
- Personnel involved in Human Intervention must be adequately trained and supervised
- Human Intervention results should continually validate and inform Machine Learning
- Monetization controls are encouraged to make Context and Brand Safety manageable

- IVT should be considered in Context and Brand Safety processes and functionality
- Proprietary taxonomies must be supported and disclosed
- MRC strongly encourages and supports mechanisms and efforts to promote consistent Brand Safety definitions and minimum baselines such as that proposed by the 4A's Advertising Protection Bureau (APB)
- A minimum level of content that excludes certain baseline categories offered as a collective reporting or blocking group considered unsafe or ineligible for monetization ("the floor") is encouraged either as a default control or as elected by advertisers such as the 4A's APB Advertising Assurance Brand Safety Floor Framework
- Advertisers may choose to waive the floor and allow advertising to one or all of the included categories
- The 4A's APB will include ongoing efforts including reconciling definitions of common category differences between vendors, while allowing for custom settings that vendors apply to address specific clients Brand Suitability needs such as via the APB's Brand Suitability Framework
- MRC further recommends and supports Industry efforts to standardize and modernize taxonomies and categories across platforms considering UGC and content granularity, and we intend to participate in such efforts on an ongoing basis

Section 4 (Adjacency) Summary (for ad verification organizations measuring and reporting adjacency)

- Adjacency should be measured based on meaningful units and should be platform specific
- Adjacency also applies to video playlist environments
- Placements that fall zero measurement units from specific content being measured or that are able to be simultaneously viewable (50% of pixels) are considered directly adjacent
- Adjacency ranges or scores may be reported and must be clearly defined
- Platforms and publishers are encouraged to enable third-party adjacency measurement or otherwise provide validated first-party signals to vendors

Section 5 (Metadata) Summary

- Social Media Platforms and publishers are strongly encouraged to allow direct third party measurement at a content level or to provide validated first-party data
- The collection and transmission of Metadata must be audited and validated
- Ad verification organizations utilizing Metadata must consider limitations of this data
- User reporting or flagging of objectionable content should be appropriately weighted based on review and robust quality control
- Automated analysis of Metadata alone without additional data sources or consideration of human and manual processes is likely insufficient for certain content level Context and Brand Safety measurement

- The degree of reliance on Metadata must be disclosed to users
- The Metadata collection method must be fully disclosed and subject to internal testing
- The Metadata must be clearly defined and the extent of Metadata use should be included in disclosures
- Processes to ingest and analyze metadata should be designed to account for different or multiple languages

Section 6 (Mobile Applications) Summary (for ad verification organizations seeking content level Brand Safety accreditation for mobile applications)

- The guidance provided in the *Ad Verification Guidelines* and within this Supplement should be fully applied within mobile applications
- Specific considerations detailed in the *IAB/MMA/MRC Mobile Application Advertising Measurement Guidelines* should be applied
- Ad verification organizations must apply incremental and differential consideration to mobile in-app environments
- Content level granularity must be applied discretely within mobile applications
- MRC intends to consider Brand Safety accreditation discretely between desktop/mobile web and mobile in-app environments

Section 7 (Illegal Sources and Industry Communication) Summary

- Ad verification organizations must consider legal restrictions related to prohibition of content and the sources of that content as well as enforcement and removal periods
- Publishers and platforms should have a qualification process to ensure they are dealing with a legitimate entity and legally permissible content applied across the platform
- Procedures should include initial qualification using Industry and local sources of known illegal entities, as well as ongoing evaluation linked with ad verification results
- Processes should be put in place to routinely communicate findings related to illegal sources to authorities, oversight bodies and the Industry at large
- MRC strongly encourages and supports mechanisms and efforts to communicate relevant Brand Safety information across the Industry and publicly

Section 8 (Enhanced Buyer Tools) Summary (for ad verification organizations seeking content level Brand Safety accreditation)

- Curation, monetization controls and granular categorization must be consistently and discretely reported to users in order to enable customized targeting and avoidance
- Ad verification organizations may enable ad prevention, blocking and alerting on the basis of Metadata, discrete content characteristics and adjacency criteria
- The basis of measurement, methodology and limitations underlying enhanced buyer tools must be clearly disclosed