Towards a generalised threat-scoring framework for prioritising results from brand monitoring programmes

by David Barnett

Introduction

The ability to sort a set of results (i.e. websites, etc.) identified through a brand monitoring service (collectively referred to as 'findings'), by ranking them according to the level of threat they pose, is a key component of many overall brand protection programmes. 'Threat' in this context covers not just the potential for (say) fraudulent or infringing use, but also addresses the degree to which a finding may be of interest / relevance to the respective brand owner.

The prioritisation process offers a number of benefits, including the identification of: (a) priority targets for further analysis; (b) candidates for content tracking (i.e. regular reinspection of content or configuration, and the generation of an alert if high-concern findings are identified) – as may be appropriate in cases where a domain name presents a high level of relevance and/or potential risk but is not currently associated with any live site content; and (c) priority targets for enforcement actions.

This study considers a number of basic features of Internet results (such as webpages and their associated URLs / domain names, or findings from other channels such as e-commerce marketplaces or social media), and how they can individually provide insights into the level of potential threat posed by those findings. By extension, consideration of multiple features in combination (with appropriate 'weightings' applied to the individual components) can serve as a basis for the construction of overall algorithms for quantifying the potential level of significance of the findings in question.

Individual features of web results which can serve as indicators of their overall potential level of threat

1. Extent of brand references on a webpage

One of the simplest metrics for determining the potential relevance of a webpage (for a particular brand) is the calculation of a 'score' which quantifies the degree to which the page relates to that brand. The exact formulation of such a score (sometimes referred to as the '**brand content score**'[1]) is flexible, but generally aims to quantify the *number* of mentions of a brand and the *prominence* of those mentions on the page.

Most usually this analysis is achieved through the analysis of the HTML source code of the page in question, *counting* the numbers of mentions of the brand in each context (e.g. in the URL *vs* the page title *vs* a page heading *v*s anywhere in the general page content) and *weighting* the counts according to the assumed level of relevance of a mention in each of the contexts (e.g. assuming that a mention in the page title will be a more reliable indicator that the page relates to the brand than a general mention elsewhere on the page), to provide an overall score.

As an illustration of the effectiveness of this approach, Table 1 shows the top ten results from a case study in which a dataset of over 2,000 assorted results collected by a brand monitoring tool for an electronics brand ('*[brand]*') was analysed and sorted by brand content score, to identify the most relevant results. The study shows that the highest-scored results do indeed feature numerous

[1] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 3: 'Brand content scoring'

prominent brand references in the page title and URL (in addition to other page locations), thereby comprising compelling examples for initial prioritised further analysis.

URL	Page title	Score
https://*[site]*.com/*[brand]*/everything-about-*[brand]*-zero-2-w/	Everything about *[brand]* Zero 2 W || Everything about *[brand]* Zero 2 W — *[site]* | Monitor and Control your *[brand]*: free for up to 5 *[brand]*s!	565
https://www.findchips.com/search/*[brand]*%203	*[brand]* 3 Price and Stock || Findchips: *[brand]* 3 Price and Stock	450
https://*[site]*.*[brand]*.com/articles/*[brand]*-4-vs-*[brand]*-3b-plus	*[brand]* 4 vs *[brand]* 3B+ - The *[site]* magazine || *[brand]* 4 vs *[brand]* 3B+ — The *[site]* magazine	325
https://flirc.tv/products/flirc-*[brand]*-4-case-silver	Flirc *[brand]* 4 Case || Flirc *[brand]* 4 Case - Flirc	313
https://www.pbtech.com/product/sevrbp0296/*[brand]*-personal-400-*[item]*-built-in-ra	Buy the *[brand]* Personal 400 *[item]* Built In ... || Buy the *[brand]* Personal 400 *[item]* Built In *[brand]* 4 4GB... (SC0384) online - PBTech.com	290
https://*[brand]*.stackexchange.com/questions/71952/whats-the-difference-between-*[brand]*-3-and-*[brand]*-3-model-b	What's the difference between *[brand]* 3 and ... || *[brand]* 3 - What's the difference between *[brand]* 3 and *[brand]* 3 Model B? - *[brand]* Stack Exchange	286
https://arstechnica.com/tag/*[brand]*/	*[brand]* || *[brand]* | Ars Technica	280
https://instock.pk/smart-boards/*[brand]*/*[brand]*-kits.html	*[brand]* *[item]*s || *[brand]* *[item]*s - *[brand]* - Smart Boards In Pakistan | InStock.PK	271
https://www.c4labs.com/product-category/cases/*[brand]*/*[brand]*-4/	*[brand]* 4 Archives - Cases || *[brand]* 4 Archives » C4Labs	270
https://www.androidcentral.com/*[brand]*-400-review-best-*[brand]*-you-can-buy	*[brand]* 400 review: The best *[brand]* you … || *[brand]* 400 review: The best *[brand]* you can buy | Android Central	265

Table 1: (Anonymised) URLs and page titles for the top ten results by brand page content score, from a dataset of findings collected for an electronics brand ('*[brand]*')

In more sophisticated versions of this approach, it may be appropriate to include modifications, such as allowing for 'fuzzy matching' (i.e. also considering *variants* of the brand name), or the application of 'caps' for the maximum possible contribution to the overall score from any given contextual location (to avoid the score being skewed by 'junk' pages 'stuffed' with large numbers of terms).

For some brands (or when carrying out brand monitoring and analysis in other contexts), it may also be appropriate to consider brand references only where they occur in close conjunction with other keywords of relevance (e.g. when a brand name is a highly generic term, or if the purpose of the monitoring is to consider only content relating to a specific issue)[2].

2. Domain name features

Domain names are a special class of website characteristic, as they are highly 'data rich' in terms of the number of associated features which can be extracted or looked up, several of which can provide insights into their potential level of brand risk – in many cases, even in the *absence* of any associated live website. Some of these features are outlined below[3].

[2] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 4: 'Use of relevance keywords'

[3] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 5: 'Prioritisation criteria for specific types of content'

a. Inherent characteristics of the domain name itself

i. Presence of a brand name or variant, or of 'high-risk' keywords

(Depending on the level of genericness of the name of the brand in question), the presence of a brand name in a domain name can be a good indicator that the name has been registered for a brand-related purpose (such as the creation of a website making a claim of affiliation or comprising an instance of brand impersonation). Instances where a *variant* of the brand name is used can more strongly imply that the domain has explicitly been registered with fraudulent intent (and might therefore contribute a larger component to the overall threat score), particularly if the variant appears to have been chosen to appear deceptively similar to the name of the brand in question, or to its official site (such as the replacement of one character with another appearing visually similar – e.g. the replacement of an 'o' with a '0', or an 'a' with a non-Latin homoglyph such as 'ɑ' – or where the brand name is preceded by a string such as 'www' or 'http')[4,5,6,7,8].

When calculating the threat score component associated with a brand-specific domain name, it might also be appropriate to consider the *context* of the brand mention, such as whether it appears at the start of the name, or the number of additional characters with which it is referenced.

Additionally, the presence of 'high-risk' keywords in a domain name – i.e. those which are highly diagnostic of active or potential use in conjunction with specific types of content of scam types – might also be an indicator of greater risk. For example, keywords by category type might include 'login' or 'verify' for phishing, 'shop' or 'store' for e-commerce, or 'discount', 'cheap', 'replica' or 'dupe' for counterfeit or otherwise infringing product sales[9].

ii. TLD (top-level domain, or domain extension)

Certain domain-name extensions are disproportionately more commonly used in conjunction with infringing activity than others, for a number of reasons. These factors might typically include the cost of domain registrations, the existence and nature of IP protection programmes, and the ease of enforcement. In particular, many new-gTLDs (the set of new extensions which have launched in the period since 2012[10]), and ccTLDs associated with regions such as Africa, Asia and the Caribbean (in some cases, a reflection of the wealth of the countries in question, which can affect factors such as the level of technical expertise of Internet service providers), tend to be more frequently associated with infringing use (and might therefore be assigned higher threat-score components)[11].

For example, one such analysis[12] calculated overall relative 'threat scores' for a set of the most highly affected TLDs (Table 2), based on a series of previous analyses of the 'frequency' of association of domains (i.e. the numbers as a *proportion* of the total numbers of registrations, rather than *absolute* values) with phishing, spam and malware activity.

---

[4] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 7: 'Creation of deceptive URLs'
[5] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 9: 'Domain landscape analysis'
[6] https://www.cscdbs.com/en/resources-news/threatening-domains-targeting-top-brands/
[7] https://circleid.com/posts/20220913-registration-patterns-of-deceptive-domains
[8] https://www.iamstobbs.com/idns-ebook
[9] https://www.iamstobbs.com/opinion/finding-the-fakes-another-application-of-keyword-based-filtering
[10] https://newgtlds.icann.org/en/about/program
[11] https://circleid.com/posts/20230112-the-highest-threat-tlds-part-1
[12] https://circleid.com/posts/20230117-the-highest-threat-tlds-part-2

TLD	Country / type	Normalised threat score
.ci	Ivory Coast	1.000
.zw	Zimbabwe	1.000
.sx	Sint Maarten	0.945
.mw	Malawi	0.862
.am	Armenia	0.608
.date	**new-gTLD**	0.506
.cd	Dem. Rep. Congo	0.391
.ke	Kenya	0.381
.app	**new-gTLD**	0.377
.bid	**new-gTLD**	0.361
.ly	Libya	0.356
.bd	Bangladesh	0.351
.surf	**new-gTLD**	0.325
.sbs	**new-gTLD**	0.250
.pw	Palau	0.240
.dev	**new-gTLD**	0.222
.quest	**new-gTLD**	0.209
.top	**new-gTLD**	0.196
.page	**new-gTLD**	0.195
.gq	Eq. Guinea	0.192
.cf	Cent. African Rep.	0.168
.ga	Gabon	0.164
.ml	Mali	0.157
.buzz	**new-gTLD**	0.149
.cyou	**new-gTLD**	0.141
.cn	China	0.130
.monster	**new-gTLD**	0.106
.bar	**new-gTLD**	0.104
.host	**new-gTLD**	0.101
.io	Br. Indian Ocean Terr.	0.085

Table 2: Relative 'threat scores' for the thirty highest-risk TLDs

iii. Domain name (SLD) entropy

The entropy (strictly, 'Shannon entropy'[13]) of a domain name (usually calculated for the second-level (SLD) name, i.e. the part of the domain name to the left of the dot) is a mathematical construct for quantifying the amount of information stored in a string (essentially comprising a measure of the number and amount of distinct characters). Accordingly, long random domain names (such as those which might be associated with the use of automated algorithmic domain registration scripts – as might be popular with infringers) will tend to have higher entropy values[14].

Previous studies on this topic have shown that (for example) new-gTLD extensions known explicitly to be associated with particular security risks (such as .zip) may be disproportionately associated

[13] S. Vajapeyam (2014). 'Understanding Shannon's entropy metric for information'. (Available at: https://arxiv.org/ftp/arxiv/papers/1405/1405.2061.pdf)

[14] https://circleid.com/posts/20230703-an-overview-of-the-concept-and-use-of-domain-name-entropy

with high-entropy domain registrations[15], and that there is some (weak) evidence of a correlation between mean domain-name entropy and independent estimates of the level of infringement risk associated with the TLD (Figure 1)[16].

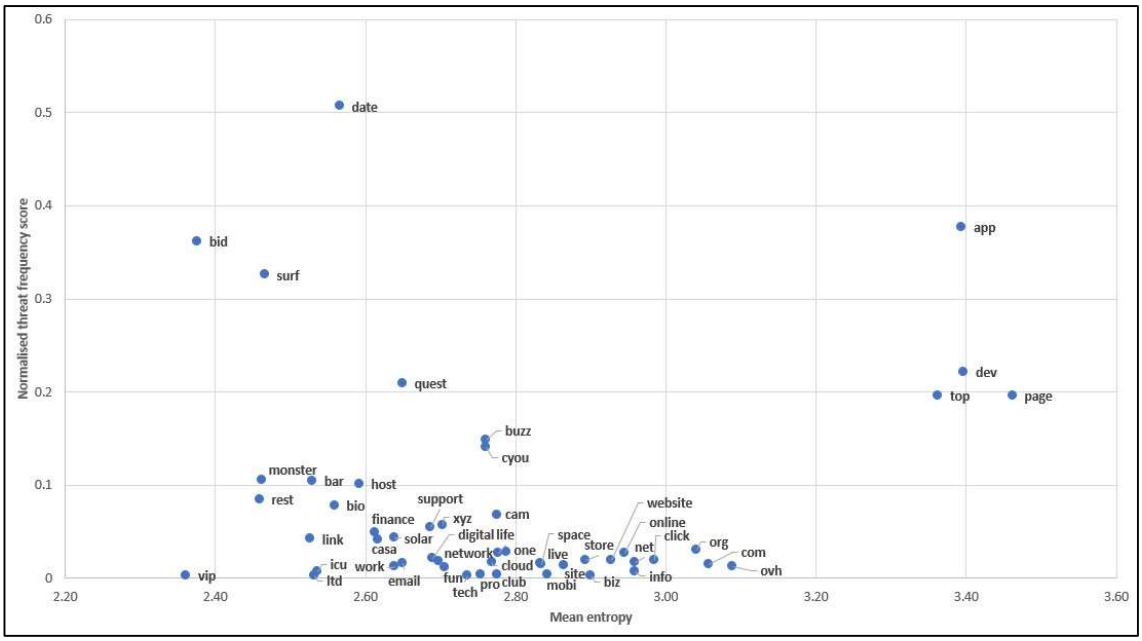

Figure 1: Comparison between TLD mean domain name entropy and normalised threat score (Barnett and Ferguson, 2023)

Accordingly, entropy in its own right might serve (to some degree) as an indicator of domain-name threat and, therefore, as one input in the calculation of an overall threat score.

b. Other domain name features

i. Registrant characteristics

Although domain ownership ('whois' / RDAP) information is very often highly redacted (particularly following the introduction of GDPR), in cases where contact details *are* given, they can frequently provide indications of higher-risk domain use. Examples might include the use of webmail e-mail providers (e.g. gmail.com, yahoo.com, qq.com, 163.com, etc.) (rather than the use of corporate domains which might typically be more indicative of legitimate business use), premium-rate contact telephone numbers, the explicit use of privacy-protection or proxy services, and locations in 'high-risk' countries (e.g. those typically associated with high rates of infringement and/or difficulty in enforcement, such as Russia or China).

[15] https://www.iamstobbs.com/opinion/un-.zip-ping-and-un-.box-ing-the-risks-associated-with-new-tlds; the highest entropy .zip domain names from this study (both ***H*** = 4.6875) were *g0kfctpdb18t7vkidqj2me5ls9rjo46g.zip* and *r5s0mo4tl315achnpvrkie76j84unba2.zip*

[16] https://www.iamstobbs.com/opinion/the-randomest-domain-names-entropy-as-an-indicator-of-tld-threat-level

ii. Registrar characteristics

Numerous pieces of research have focused on quantifying on the rates of infringing activity associated with domains registered through particular registrars – and, therefore, indirectly providing a measure of potential threat associated with the registrars in question. Particular registrars may be disproportionately more popular with infringers than others for a range of reasons, including differences in their inherent level of cooperativeness to ('compliance' with) notifications of IP infringements, their speed of response, and geographic region(s) of operation (for which variations in local laws may be a relevant factor). One such dataset is that provided by Spamhaus[17], which (as of 24-Mar-2025) gives the top five 'low-trust' registrars (by quantitative 'bad reputation score') as shown in Table 3.

Rank	Registrar	Score
1	Miracle Ventures Ltd	3119.7
2	EU Technology (HK) Limited	1479.5
3	香港翼优有限公司 ('Hong Kong Wingyou Co., Ltd.')	655.7
4	Domain International Services Limited	90.6
5	Zname Ltd	87.1

Table 3: Top five 'low trust' registrars by 'bad reputation score', according to Spamhaus (Mar-2025)

In addition, many brand protection service providers collate information on the compliance of individual registrars, based on previous enforcement experience. This also allows for the construction of a risk 'score' for each registrar, which can serve as an input into threat-scoring algorithms for domains generally.

iii. Hosting characteristics

An approach to threat quantification based on website hosting characteristics, which is more granular than considering *just* the identity of the service provider in question, can be implemented through consideration of the specific IP address on which a website is hosted. One such methodology (as presented in a previous study[18]), is based on the proximity to other IP addresses known to be associated with the hosting of infringing content (as referenced in 'blacklist' databases, such as the example provided by Myip.ms[19].

The basic principle behind this methodology is to split IP-address space into 'netblocks' with common initial elements, and calculate the number of blacklisted IP addresses in each block. It is then assumed that a website hosted on an IP address in (or, optionally, closely adjacent to) a block featuring a large number of blacklisted addresses is itself more likely to be infringing.

The previous study exploring a simple-minded implementation of such an approach does indeed show some indication of effectiveness, with several examples of sites of potential concern (featuring browser warning messages, indications of 'geoblocking', or high-threat content or keywords – e.g. those associated with cryptocurrency or potential phishing) identified within the top twenty most highly scored sites, taken from a database of around 11,000 domain registrations.

---

[17] https://www.spamhaus.org/reputation-statistics/registrars/domains/

[18] '"Notorious IP Addresses" and initial steps towards the formulation of an overall threat score for websites', Stobbs e-book [*link TBC*]

[19] https://myip.ms/browse/blacklist/Blacklist_IP_Blacklist_IP_Addresses_Live_Database_Real-time

In a modified version of the above idea, it is possible also to consider the hosting service provider associated with the IP address for a website of interest. In a general sense, similar comments to those presented above for registrars can also apply to hosting providers, with some even explicitly describing themselves as 'bulletproof' – implying a lack of compliance to enforcement notices – as a means of attracting business from providers of illicit content.

However, even outside this set of obviously high-risk providers, it is possible to gain threat insights from an analysis of website hosting providers. In one such study[20], hosting providers were ranked according to the frequency of their association with IP addresses which had explicitly been blacklisted (in response to infringing activity such as spamming or malware distribution), and with the raw data 'normalised' according to the total number of IP addresses managed by the provider in question. This allows the construction of a 'blacklist rate' (or threat) score for each provider (Table 4), which itself can be used as one component of an overall threat score for websites generally.

Hosting provider	Blacklist rate
Huawei HongKong Clouds	512.67
Ahrefs Pte Ltd	462.00
Yandex enterprise network	382.00
Huawei-Cloud-SG	280.67
Bangladesh Telegraph & Telephone Board	280.00
Netprotect	270.00
Strong Technology	189.00
geofeed (GitHub: *Simonadascalu/Freedomtech-Geofeed*)	116.00
LogicWeb Inc.	112.00
Huawei Cloud Singapore POP	95.00

Table 4: Top ten 'highest threat' hosting providers, by 'blacklist rate' score

iv. MX records

The presence of an MX (mail exchange) record implies that a domain has been configured to be able to send and receive e-mails, and could therefore potentially be associated with (for example) phishing activity. As such, all other factors being equal, a site with an active MX record would generally be scored more highly in a threat-level determination.

v. SSL certificate provider

An SSL (Secure Sockets Layer) certificate is a feature which authenticates the identity of a website and enables an encrypted connection to it (reflected as an '*https*' URL). However, the owners of infringing websites frequently purchase certificates from budget or free providers[21] (many of which may not carry out appropriate legitimacy checks, or may have poor implementations of security measures) purely as a means of adding credibility to the site. Accordingly, the presence of an SSL certificate is not, in itself, an indicator of site authenticity (noting the 2021 APWG study which found that over 80% of phishing sites had active SSL certificates[22]). Where SSLs *are* present, however, the

[20] https://circleid.com/posts/notorious-hosting-providers-an-overview-of-the-highest-threat-hosts-from-ip-address-blacklist-analysis

[21] An old (2015) study on this topic can be found at: https://www.netcraft.com/blog/certificate-authorities-issue-hundreds-of-deceptive-ssl-certificates-to-fraudsters/

[22] https://docs.apwg.org/reports/apwg_trends_report_q2_2021.pdf

level of trust of the SSL certificate *provider* (based on previous research) would be a suitable input into a threat scoring algorithm.

vi. Web traffic

For live sites, measures of the amount of traffic (i.e. number of visitors) received by a website are a useful input into prioritisation metrics, based on the principle that a more popular website is accessed by a greater number of users, and thereby presents a greater potential risk (though noting that *official* sites will, in general, also tend to have high traffic levels). In the absence of availability of direct measurements, traffic estimates are most usually obtained through any of a number of online tools and services (such as Similarweb and Comscore), which base their data on a 'sampling' of web users.

3. Other factors

In addition to the above points, other website characteristics – such as previous reports of abuse (as given by databases provided by sources such as Lumen[23] and WIPO), or identification of the use of wildcard DNS (allowing the use of *any* arbitrary subdomain name to constitute a valid, resolvable URL) – may also serve as appropriate inputs into an overall threat scoring algorithm.

4. Features of other online channels and the use of data 'proxies'

For content on other Internet channels, additional data fields may be available, and suitable for use as inputs into metrics to quantify the overall level of potential threat associated with the findings.

On e-commerce marketplaces, for example, features such as price point and item quantity may be relevant (working on the basis that non-legitimate items may tend to be offered in higher volumes and at prices more markedly below the recommended retail price), together with other characteristics, which might include the (unauthorised) use of copyrighted images, and the presence of brand references in the titles or descriptions of listings for explicitly third-party products (trademark infringement, i.e. traffic misdirection).

For other channels, it might be appropriate to utilise other available metrics. On social media, for example, where (wholesale) domain web-traffic measures are not relevant, characteristics such as numbers of 'likes', 'shares' or 'views' may be used as data proxies, to provide a measure of the degree of user exposure for profiles or postings.

Other relevant features which may provide indicators of potential risk might include: (for digital content) numbers of downloads or file-sharers[24]; the use on the page of official logos or branded imagery; or the degree of similarity to site 'templates' known to have been previously used in the creation of infringing content.

Conclusions

All of the factors discussed in study (in addition, potentially, to others) are likely to be relevant to an overall assessment of the likely level of threat posed by any arbitrary identified webpage (i.e. a 'finding' from a programme of brand monitoring) and, accordingly, may serve as components of a comprehensive metric, or 'score', designed to quantify this degree of potential threat. The key to

[23] https://lumendatabase.org/

[24] 'Patterns in Brand Monitoring' (D.N. Barnett, Business Expert Press, 2025), Chapter 11: 'Quantifying brand protection return-on-investment'

making any such metric as meaningful as possible is a determination of the relative *scores* and *weights* to be assigned to each individual component.

As discussed in this study, some of the relevant webpage features are (or can be made) intrinsically quantitative in their own right, with examples including brand page content score, web traffic, or TLD threat score – though this does not address the issue of how to weight these individual components relative to each other, in an overall assessment of potential threat level. For other components of the overall metric, it will be necessary to determine how to assign 'scores' to the relevant possible options which may exist for each component – e.g. assigning Internet service providers into 'tiers' according to factors such as their levels of compliance, or their frequency of association with infringing activity, and determining the relative scores which should apply to each tier. Part of this analysis is likely to require comparison with real data, which may include independent assessments of numbers and types of infringements (as has been alluded to in the methodologies discussed for registrars and hosting providers), and/or insights from brand protection service providers.

Ultimately, the construction of any robust metric may be a task for which artificial intelligence and/or machine learning may have a valid application; given a database of webpages and their associated charateristics, it might be reasonable for an AI-based application to be able to generate a prototype metric, based on assumed scores and relative weightings associated with the characteristics in question, with these quantitative settings then allowed to be modified (i.e. 'tuned'), based on manual feedback regarding which findings are *actually* associated with fraudulent or infringing activity.