

OPINION UNDER SECTION 74A

Patent	GB2602159 B
Proprietor(s)	Lighthouse Technologies Limited
Exclusive Licensee	-
Requester	Lighthouse Technologies Limited
Observer(s)	Cooley(UK) LLP on behalf of Unitary Ltd.
Date Opinion issued	23 September 2026


The request

1. The comptroller has been requested to issue an opinion on whether GB2602159 is being infringed by a product which the request calls the Unitary content moderation system. The request includes a number of extracts from the Unitary Ltd, Lasso moderation and Amazon EKS websites. These extracts are suggested illustrate the nature of the content moderation system being offered by Unitary Ltd.
2. The Patent, GB2602159 was filed on 21st December 2020 and granted on 18th January 2023. The Patent is currently in force.
3. Observations were filed by Cooley on behalf of Unitary Ltd, and observations in reply followed from Lighthouse.
4. I should note that the request only asks me to consider claim 1 and claim 8, the independent claims here. The request makes no assertion on the dependent claims. I shall therefore only need to consider those independent claims.
5. Finally, I note that the observations ask me to decline to issue the opinion but provides no grounds or reasoning. It seems to me to be appropriate to consider this request based on the information provided.

The Patent

6. The Patent relates to the redaction of digital content which is considered undesirable. In the embodiment shown in the figure 1 (shown below), this may provide a cloud-based redaction system. Users of different ages might be presented with different images. As described on page 3 of the Patent, that allows the images presented to a child to be filtered. For example, filtering for images and videos of adults wearing little clothing on social media. The user's age may be used by a

machine learning or AI model. The model then determines what web content is suitable, or which visual content should be censored, redacted or blocked.



7. As set out in claim 2 and 3 of the Patent, in one embodiment, the identification server is a proxy, receiving web content requests. That means it acts in that embodiment as an intermediary between content providers at the bottom, and users making requests, at the top. Web traffic is channelled through the identification server to the user device. However, those are dependent claims. Although they may be helpful in

understanding this figure, I am only asked to consider the independent claims.

8. I note that there is one passage of the description which is raised in the request on page 9 lines 11-15, to which I will need to return:

10

In this embodiment, the system 100 is cloud-based and provides an internet content monitoring and censoring service. However, alternative embodiments, such as embodiments in which some of the components of the system 100 run on, for instance, "Docker" instances on the infrastructure of a client using the service, may be

15 employed.

The claims

9. The request breaks claim 1 down into 14 parts – which it labels a-n. The request provides a selected passage from the websites provided as a mapping of these features for the question of infringement. I have included this breakdown here, since it will be of benefit when I look at that mapping. I should also say that this is the structure used in the Observations. However, the observations bundle these into discussion points labelled sections A)-E).
- a) *A computer-implemented method of redacting undesirable digital content, the method comprising:*
 - b) *receiving, at an identification server, a content request originating from a user device;*
 - c) *identifying a user account associated with the content request;*
 - d) *tagging the content request with a restriction identifier*
 - e) *which is indicative of a restriction level associated with a user of the user device, the user or user device being linked to the user account;*
 - f) *analysing, by an inspection engine communicatively coupled to the identification server, digital content requested by way of the content request before the digital content is transmitted to the user device,*
 - g) *wherein the digital content includes visual content in the form of image and/or video content,*
 - h) *wherein the inspection engine implements an artificial intelligence inspection model, and wherein the inspection engine is trained, using machine learning, to inspect the visual content in order to determine whether the visual content, or part thereof, is appropriate for the user based on the restriction level;*
 - i) *using a result of the analysis performed by the inspection engine to determine whether the visual content, or part thereof, is undesirable based on the restriction level,*
 - j) *wherein the result of the analysis includes a score associated with the visual content and a confidence level output by the inspection engine,*
 - k) *and wherein the score and the confidence level are taken into account in determining whether to redact the visual content or part thereof;*
 - l) *if the visual content or part thereof is classified as undesirable, redacting*

- the visual content or part thereof and causing a redacted or censored version of the digital content to be transmitted to the user device,*
- m) wherein redacting the visual content or part thereof includes replacing the image and/or video content, or parts thereof, with safe image and/or video content in the redacted or censored version,*
- n) or if the visual content or part thereof is not classified as undesirable, allowing the digital content to be transmitted to the user device substantially unchanged.*

10. A similar breakdown is provided for claim 8, which is the system claim equivalent to the method of claim 1. In practice, therefore, my discussion on claim 1 will be mirrored in claim 8.

Construction of the Claims

11. *Kirin-Amgen and others v Hoechst Marion Roussel Limited and others* [2005] RPC 9 requires that I put a purposive construction on the claims, interpret it in the light of the description and drawings as instructed by Section 125(1) and take account of the Protocol to Article 69 of the EPC. Simply put, I must decide what a person skilled in the art would have understood the patentee to have used the language of the claim to mean.
12. The request also notes the Actavis Decision (*Activis v Eli Lilly & Co* [2017] UKSC 48) at 54-5 and *Generics (UK) Ltd v Yeda* [2017] EWHC 2629 at 138.
13. The request identifies the skilled person or team as being an experienced software engineer and a machine-learning engineer. They have knowledge of content moderation and filtering systems, machine learning models for image and video classification, client-server architectures for content delivery and parental control and child safety technology. I think that is a reasonable starting point, and it is not challenged in the observations.
14. In practice much of the discussion in the observations and the observations in reply focus on the question of reading the claims in a way that will be useful for determining the question of infringement. That means that they focus on a form of the claims that looks like the Unitary content moderation system, or the words that might be used to distinguish that from the claims. I have tried to separate these construction questions here, before later turning to the infringement question.
15. In the observations, Cooley describe the invention as being in substance:
- receiving a content request originating from a user device, at an identification server;
 - identifying a user account associated with that request;
 - tagging the request with a restriction identifier indicative of the user's restriction level;
 - an AI/ML inspection engine analysing the visual content before it is

transmitted to the user device, to assess appropriateness against the restriction level;

- a result comprising both a score and a confidence level, both taken into account in deciding whether to redact; and
- if undesirable, redaction (including replacement with safe content) and transmission of a redacted version to the user device, or, if not, passing the content through unchanged.

16. That goes some way in terms of summarising the claim. However, it does simplify the claim, meaning that parts of the claim are ignored. It is therefore important that I look at the whole of the claim.
17. There are five aspects to the claim labelled with capitals sections A)-E) in the observations from Cooley. These sections focus on particular combinations of the steps a)-n) listed above. I will consider each of these in turn.
18. First, there is discussion (A) of the content request and “positioning” of the redaction steps. This focusses on the extent to which step b) is limited to an arrangement in which the identification server acts as an intermediary (or proxy) between the user and a content provider¹.

b) receiving, at an identification server, a content request originating from a user device;

19. It is clear that the Patent envisages this possibility. However, Lighthouse argue that it also envisages other possibilities noting page 9 lines 11-15:

10

In this embodiment, the system 100 is cloud-based and provides an internet content monitoring and censoring service. However, alternative embodiments, such as embodiments in which some of the components of the system 100 run on, for instance, “Docker” instances on the infrastructure of a client using the service, may be

15 employed.

20. I should say that none of the request, observations nor observations in reply provide any detailed discussion of what a “Docker” instance is. I understand that to be a container operating on a server, such as in a data centre or a cloud provider. Nonetheless, this second paragraph does suggest that some part of the program may be run on a client device. This passage does not say how much or which parts. However, it appears to exclude the possibility that the whole process is provided on a client device.
21. Lighthouse argue in the request that under the “Doctrine of claim differentiation” that claim 1 must be broader than its dependent claims. As such the term “doctrine of claim differentiation” appears to be a reference to US caselaw and practice.

¹ NB: The term content provider is not used in the claim.

However, Cooley in their observations argue broader does not mean unlimited. Cooley state: *“The identification server must still perform the functional role recited in the claim, namely receiving a request that originates from a user device seeking access to content, and the content must be intercepted before it reaches that device.”*

22. When I look the claim, there is no explicit requirement that the content is “intercepted”. Rather the claim requires that a “content request” originates from a user device. The claim states that the identification server receives a content request originating from a user device. It also requires that a user account must be identified which is associated with the content request. The claim goes on to say that a restriction identifier is tagged associated with the user of **the** user device.
23. So, an understanding of the term user, provides an insight into what the requirements are on the identification server. Lighthouse contend that the user might be a “customer” submitting content for assessment. Lighthouse argue for a broad construction covering any computing role or device, and any authenticated account. I note that there is no discussion in the Patent of content being submitted for some sort of pre-assessment. Lighthouse rely, I think, on the idea that components might run on a client to suggest that the skilled person is looking for other ways to implement the invention. Indeed, in the observations in reply, Lighthouse argue that all of the “positional” (that is to say where in the network) stipulations in the claims are reserved for dependent claims. Lighthouse therefore argue that claim 1 is limited only by the functions carried out.
24. So, can the claim be read in a way that allows content to be pre-assessed. That is to say does the skilled person understand that this could also be a content provider as a user? In step b) the user would be a “customer” (to borrow the description in the observations) submitting content for assessment. In step c) that user would be authenticated, so the content request is a request for content sent to be processed. In step f) that content is analysed before the digital content is sent to **the user device**. This implies to me that the claim only contemplates there being one “user” and that particular content request is then processed and either a redacted or the original content is transmitted to “the user”.² In the final steps m) and n), the content is then transmitted either in a redacted form or substantially unchanged.
25. The skilled person reading the document will have in mind that first embodiment of a user accessing content provided by someone else, which is the focus of the embodiments described in detail. The skilled person is in the context of *Actavis Group PTC EHF v ICOS Corporation & Ors* [2017] EWCA Civ 1671 described as being someone who is trying to understand what the patentee means³, and in the context of Section 14(3) and support “trying to carry out the invention and achieve success...not searching for a solution in ignorance of it.”⁴ More generally, the skilled person is not considered to be someone who exercises inventive ingenuity or thinks

² [I should note of course that I envisage that the skilled person would understand that there might be multiple end users each making their own content requests. However, that is a separate point to the logic I am trying to set out here.]

³ See 125.17.7 of the Manual of Patent Practice.

⁴ See 14.64 of the Manual of Patent Practice and its discussion of [Zipher Ltd v Markem Systems Ltd](#) [2009] FSR 1

laterally.⁵

26. Given this context of the skilled person, I am not convinced that the reference on page 9 to some components which may be run on another device is enough to lead the skilled person to start a process of imagining alternatives to the flow in the principle embodiment here (shown in figure 1 of the process being carried out when the user's remote device tries to access the content.) It may well be that part of the process could be carried out on the user's device, or indeed at the content provider.
27. However, I think that the skilled person is not led to contemplate an embodiment where the content is first redacted according to some policy, and returned to the content provider or database, for later use by transmitting to an end user. I would also say that the final step of transmitting the content back substantially unchanged does not to me seem like it fits well with a scenario in which the content is considered as being desirable. It seems more straightforward for some tagging or response message to be sent that confirms the existing content is desirable, without retransmitting substantially the same information back.
28. Ultimately, I agree that the claim might not specify a "physical" positioning for the task in the network (on the user device, a cloud server, a proxy or indeed on the content provider's device(s)). However, it does require in my view a logical flow of content to a user from a provider. I think that understanding is sufficient for me to come to a view later on the Unitary content moderation system.
29. Secondly, in section B) the observations look at account identification and tagging in the context of steps c) and d) – to make an analogous argument, that the content request would be understood in the context of the description as a request from a user to a content provider.
 - c) *identifying a user account associated with the content request;*
 - d) *tagging the content request with a restriction identifier*
30. Much of the discussion that follows in the observations and the observations in reply focus on how this should be applied to the Unitary content moderation system. However, I should also note that steps c) and d) follow the embodiment described of a user requesting content, and a content provider returning that content through a process which involves checking what redaction is appropriate. Again, there appears to be nothing in the description which suggests that the content provider provides content to a reaction service and receives content back. I again therefore conclude that the skilled person would understand the user here to be someone other than the content provider. I am mindful in that respect of what was said in *Catnic Components Ltd and another v Hill and Smith Ltd* [1982] RPC 183, where Lord Diplock stated (at page 243): - "A patent specification should be given a purposive construction rather than a purely literal one derived from applying to it the kind of meticulous verbal analysis in which lawyers are too often tempted by their training to indulge". I think the skilled person would take that embodiment at face value, given their desire to implement the invention. I do not think that the skilled person envisages this alternative.

⁵ See 3.21 of the Manual of Patent practice

31. Thirdly, the observations look in section C) at the restriction level in steps e) and h). Again, this is in a similar vein tagging:

e) which is indicative of a restriction level associated with a user of the user device, the user or user device being linked to the user account;

h) wherein the inspection engine implements an artificial intelligence inspection model, and wherein the inspection engine is trained, using machine learning, to inspect the visual content in order to determine whether the visual content, or part thereof, is appropriate for the user based on the restriction level;

32. In the request, Lighthouse note that the restriction identifier is a broad term. Lighthouse note page 5 lies 8-9 and 4-7, to support a view that the restriction identifier encompasses any parameter, configuration, policy setting or token that indicates the applicable restriction level or the user, the user device or the user account.
33. The observations suggest that the wording of the claims was a result of a need to overcome a prior art objection during examination. Specifically, Cooley set out:

*Lighthouse's attorneys amended the claims and stated that –
“The term 'restriction parameter' has been replaced with 'restriction level' in the claims, to narrow the scope of protection sought to the restriction level associated with the user.”*

This is a clear and unambiguous statement made to avoid the relevance of the prior art, that “restriction level” is narrower than “restriction parameter.” The specification is consistent with this: a “restriction level” is described as one species of the broader genus “restriction parameter.”² The only worked example of a “restriction level” in the specification is a set of discrete, tiered values associated with a single characteristic of the user, namely their age (“A”, ‘10’, ‘13’, ‘16’, and the like”), used in the manner of a film classification.³ By contrast, the specification's broader references to “restriction parameters” expressly include simple binary toggles, such as whether content is “safe for work” or “not safe for work.”⁴

Lighthouse cannot now seek to construe “restriction level” so broadly as to recapture the very breadth it disclaimed to secure grant. That would be inconsistent with the principle that a patentee cannot, for the purposes of infringement, resile from a limitation deliberately introduced to overcome an examiner's objection without re-engaging the objection that necessitated the amendment. It would also produce the anomalous result that the amendment achieved nothing, contrary to Lighthouse's own representation to the Office that the amendment narrowed the scope of protection sought to the “restriction level associated with the user”.

“Restriction level” in claims 1 and 8 is narrower than the broader concept of a ‘restriction parameter’. It denotes a discrete, tiered value associated with a characteristic of the user.

34. Cooley also quote from:

Page 4, line 35 – Page 5, line 1: *“The restriction parameter may be a restriction level.”*

Page 9, lines 32-35, and Page 10, lines 1-3 – *“The system 100 allows the person managing the user account to set a restriction parameter in the form of a restriction level associated with the user account. This may be a stickiness[sic] level, e.g. correlated with movie poster ratings to make it easier for users. “A”, “10”, “13”, “16”, and the like. When the user device 120, 122, 124 submits a content request, the server 102 identifies the relevant user account and retrieves a “restriction parameter” or a “restriction level” of the user using the user account, e.g. “13”.*”

Page 9, lines 4-9 – *“The aim is thus to protect children from inappropriate, harmful, offensive, and other undesirable online content. However, it will be appreciated that other restriction parameters may be employed without departing from the scope of the invention, e.g., checking whether content is appropriate or inappropriate for a user logged into a work computer (often referred to as “safe for work” or “not safe for work”).*

35. The result is that Cooley are arguing that the term restriction level should be read as a tiered gradation of content. Here, I would note that the Patent discusses a range of contexts, such as being in a library, at work or the users age as factors that might be considered. I am mindful too that the skilled man here is someone working in the field with knowledge of a variety of content moderation systems, and the factors that may be considered.
36. In the observations in reply, Lighthouse suggest that the Observations invite the examiner to construe “restriction level” through a single sentence of the applicant’s response of 16 December 2021. The observations in reply point to *Actavis UK Ltd v Eli Lilly & Co* [2017] UKSC 48 at [88], to suggest reference to the file is appropriate only where “the point at issue is truly unclear if one confines oneself to the specification and claims of the patent, and the contents of the file unambiguously resolve the point”, or where ignoring the file would be contrary to the public interest.
37. However, the statement in the letter does suggest that the term level must be construed more narrowly than a parameter. In the observations in reply, Lighthouse set out four considerations. First, that the letter includes a preamble. That states “without acquiescing to the Examiner’s objections” I think little, if any, weight can be placed in my analysis to this boiler plate statement.
38. Second, the letter references page 4 lines 31-34 which were part of the statement of invention section, setting out optional “may” features of the invention. So that a parameter might have been an identification age of the user. Here, I think that part of the description is useful. The level must in my view be something more than simply an identification of age. From that I take it that the restriction level reflects something more than this raw data. It is something about how content for the user is to be treated.
39. Third, in the observations in reply Lighthouse refer to one of the cited documents,

suggesting that unordered preference sets and element lists as shown in US 8700409 should be considered parameters but not levels – rather Lighthouse argue that “level” means graded strictness values on an ordered scale.

40. Fourth, the observations in reply look at the question of recapture anomaly in the context of another document. Lighthouse do so to suggest that there were other amendments made in their response to the first examination report, which would also need to be considered in relation to this issue. Finally, Lighthouse note that the amendments were accepted without the examiner providing a construction of the term “level” in the file. On this final point, I do not think that the examiner’s statements or absence of them play into my consideration on this case.
41. The observations in reply also point me towards what it considers the ordinary meaning of the term level – a value on an ordered scale. The observations suggest that this must be discrete, tiered, associated with a single characteristic of the user and perhaps even according to Lighthouse’s wording “in the manner of a film classification.” I think it is true that each of these is compatible with a level. However, I am conscious too of the idea that say advanced level students might reflect a mix of different students with different levels of skills in different areas. I do not therefore think that I should construe the “restriction level” as being a single factor or a linear one.
42. Having considered all of this, I am not convinced that there is a clear pointer to a particularly narrow construction of the term level from what was said in response to the examination report. I think that the skilled person is likely to recognise that redaction may need to be tailored so that it allows for different contexts, times, locations and connection architectures to allow a policy to be implemented effectively. I believe that the skilled person would understand that this might be individual decision for a particular user, perhaps based on a set of parameters, but resulting in some determination of the approach to be taken to the restriction. That approach is tied to the level.
43. Next section D) looks at the terms “score” and “confidence level” which are in steps j) and k) of the claim.
 - j) *wherein the result of the analysis includes a score associated with the visual content and a confidence level output by the inspection engine,*
 - k) *and wherein the score and the confidence level are taken into account in determining whether to redact the visual content or part thereof;*
44. Cooley point to these steps in order to assert that the claim requires both a score and a confidence level are used. Lighthouse accept that the score is a value associated with the visual content and the confidence level expresses how certain the model is that the content is undesirable. Lighthouse go on to suggest that this means that the concepts could be generated independently. I agree that the claim makes no limitation on how these concepts are calculated. Nor does the claim require that they must be done independently.
45. However, Cooley again attempt to rely on the response to the examination report where the applicant stated:

*“Finally, while US’716 does refer to confidence levels, it in fact teaches away from the present invention as it suggests that confidence levels be used to categorise documents. For example, paragraph [0027] states that a document can be analysed to “determine a confidence level (e.g., from 0% to 100%) that the electronic document belongs one or more of the following categories: sex education, mature content, pornography, gambling, hate speech, alcohol, drugs, tobacco, bombmaking, weapons, etc.” As already mentioned above, the present invention specifically seeks to move away from filtering based on content category as this creates the risk that certain content may be marked as “safe” by virtue of general context or category and thus not be picked up by the filtering system (see the “Background to the invention” of the present application). **Instead, the present invention as claimed in the amended independent claims directly inspects visual content and determines whether to redact the content based on a score allocated to the content and the model’s degree of the certainty that the content is undesirable.** Accordingly, it is submitted that knowledge of US’716 would be likely lead a skilled person down a path which the present invention specifically seeks to avoid.”*

46. The highlighted passage reflects what Lighthouse continue to argue in their observations in reply. It supports the idea that the score is not a confidence level. Indeed, in the request, Lighthouse state:
*“[t]he skilled reader would understand feature (e) to encompass an inspection engine that: (i) outputs, per item of visual content, a classification identifying the relevant property, and (ii) **outputs a numeric value on a 0-1 scale expressing certainty in that classification, tested against a threshold.**”* [Emphasis added].
47. I would not agree that the confidence necessarily need to be a numeric value on a 0-1 scale. Lighthouse further point to page 6 lines 3-8 in an attempt to suggest that the score and the confidence level might be the same thing. This passage is couched in “may” and in terms of the inspection engine and another component model. I do not think it provides a solid basis to suggest something other than the plain reading of the claim that the confidence and the score are separate concepts.
48. Ultimately, it seems to me that the score and the confidence level are separate concepts.
49. Finally, the observations in section E) turn to steps l) – n) of the claim, asserting that they require that where content is classified as undesirable, it is redacted. This redaction includes by “replacing the image and/or video content, or parts thereof, with safe image and/or video content.”
 - l) if the visual content or part thereof is classified as undesirable, redacting the visual content or part thereof and causing a redacted or censored version of the digital content to be transmitted to the user device,*
 - m) wherein redacting the visual content or part thereof includes replacing the image and/or video content, or parts thereof, with safe image and/or video content in the redacted or censored version,*

n) or if the visual content or part thereof is not classified as undesirable, allowing the digital content to be transmitted to the user device substantially unchanged.

50. In the request, Lighthouse argue for a broad interpretation of the term redact to include blacking out, partial obscuring, replacement with a safe image, inclusion of a redaction message. Lighthouse quote page 13 lines 11-13 to suggest numerous other techniques are contemplated. Lighthouse assert that the ordinary meaning encompasses blurring, occlusion (blocking or obstruction), overlay and placeholder substitution. In the observations in reply, Lighthouse further assert that replacement of content with safe content also amounts to redaction.
51. None of this relies on evidence of what the skilled person would routinely consider as standard redaction techniques. Nonetheless, it seems to me that all of the techniques listed in my previous paragraph are likely to be ones considered by the skilled person. There does not seem to be a real debate on this. Rather, Cooley argue that the step is part of the wider claim, in the context of the infringement question.

The Unitary content moderation system

52. In the request, Lighthouse provide a variety of extracts from:
- Unitary information from the website at: <https://www.unitary.ai/>,
 - the Lasso moderation website at: <https://lassomoderation.com/intgrations/unitary-content-moderation>
- and
- the AWS amazon website “Unitary Scales AI Moderation to 26 Million Videos Daily with Amazon EKS” available at: <https://aws.amazon.com/solutions/case-studies/unitary-eks-case-study/>
53. All of these documents were provided with the request. There is no suggestion in the observations that these documents present an incorrect view of the product in question.
54. The Observations describe Unitary as providing B2B solutions. The observations suggest Unitary has two versions of a content moderation system.

“Customers submit their user-generated content to one of the two versions. The corresponding version classifies such content based on policies provided by each customer. In one version, the Product returns a score that is indicative of the classification. The customer decides what to do with the classification. The Product does not perform any further action. In another version, the Product, in addition, performs a pre-approved action that is decided in advance by the customer, based on the classification. The pre-approved action is performed on the customer’s system and not on Unitary’s system (i.e., the Product). In short, it is the customer that decides what to do

with the classification that the Product provides. The decision that the customer takes is then actioned within the customer's system. Some example decisions include approving a newly created campaign for launch, taking down a review, restoring a taken-down review, banning an end user on a temporary or permanent basis, etc."

Infringement

55. The request makes arguments for the product and process infringing the Patent, directly. It goes on to consider, in the alternative, contributory infringement and infringement on the basis of equivalence, using the Actavis questions. I shall therefore need to consider all of these in turn. I note that the request adds information about prior notification to Unitary, but that is not a matter that I shall need to consider in my opinion.

56. Section 60 Patents Act 1977 governs what constitutes infringement of a patent; Section 60(1) reads:

Subject to the provision of this section, a person infringes a patent for an invention if, but only if, while the patent is in force, he does any of the following things in the United Kingdom in relation to the invention without the consent of the proprietor of the patent, that is to say -
(a) where the invention is a product, he makes, disposes of, offers to dispose of, uses or imports the product or keeps it whether for disposal or otherwise;
(b) where the invention is a process, he uses the process or he offers it for use in the United Kingdom when he knows, or it is obvious to a reasonable person in the circumstances, that its use there without the consent of the proprietor would be an infringement of the patent;
(c) where the invention is a process, he disposes of, offers to dispose of, uses or imports any product obtained directly by means of that process or keeps any such product whether for disposal or otherwise.

57. Here, claim 1 relates to a computer program product and claim 8 to the system. In practice, I think the substantive questions are the same for both forms of the claim. I shall therefore focus again on claim 1 here.

58. The request includes a document entitled "Claim Chart Annex - - Infringement" mapping of the various documents to claims 1 and 8. I do not intend to repeat that here. I shall instead focus on the apparent points of contention. For the avoidance of doubt, I have nevertheless worked my way through each of the passages quoted and the wider documents referenced in that Infringement mapping.

59. I shall therefore work through the steps in turn and will focus on those which Cooley assert are not in the Unitary content moderation system.

60. There is no argument that step a) is not in the Unitary content moderation system. Cooley, first argue in Section A) that there is no content request from a user in the sense of step b) of the claim.

b) receiving, at an identification server, a content request originating from a user device;

61. Cooley suggest that the Lighthouse product does not receive a content request. Rather Cooley suggest that a customer [user] submits content for assessment.
62. Lighthouse in their observations in reply argue that the content is submitted to the platform (which in the Architecture diagram discussion shown in the AWS case study: Unitary on Amazon EKS includes a load balancer). This then leads the platform to submit the content, authenticated by a bearer token to the communicatively coupled Unitary engine for analysis.
63. As I construed the claim above, step b) requires, in the context of the claim as a whole, something more. The content request should flow from a user to a content provider. I do not think that Lighthouse argue that this is present in the Unitary content moderation system at this stage.
64. That leads me to consider the other parts of the argument raised by Lighthouse. Lighthouse assert that the purpose of the Unitary content moderation system, will be later to provide that content to a user. At that stage there will be a content request, and redacted content provided to the user. Lighthouse contend that this falls within the scope of the claim, as the claim does not specify a particular timing to the relevant steps.
65. As I construed the claim above, the flow is important, in step h)

h) wherein the inspection engine implements an artificial intelligence inspection model, and wherein the inspection engine is trained, using machine learning, to inspect the visual content in order to determine whether the visual content, or part thereof, is appropriate for the user based on the restriction level;

66. That requires that the inspection engine is aware of the end user. I do not therefore think that this scenario of pre-processing the content, for later serving to end users, falls within the scope of the claim as I have construed it. That is to say even if I start to read the claim with an unknown future user in mind – making a later request, with the content being redacted and tagged perhaps for various classes of users, I am still left unable to square the circle of the user device receiving the content in step f) – making the content request in step b) and being identified in step c). The logic of the claim does not allow me to insert two types of users – the content provider and later an end user. As I put it earlier the claim requires a (logical/ single) user of the (specific) device.
67. I must, however, also consider the argument made in the request, that under the doctrine of equivalents, the Unitary content moderation system should still be considered to infringe.
68. Actavis UK Limited and others v Eli Lilly and Company [2017] UKSC 48 set out three questions to assist in determine whether an equivalent should be considered to infringe, following the Protocol on the Interpretation of Article 69 of the EPC. The Three Actavis Questions are where a variant does not fall within the claim on a

normal interpretation at paragraph 66 and read:

Q1: Notwithstanding that it is not within the literal meaning of the relevant claim(s), does the variant achieve (i) substantially the same result (II) in substantially the same way as the invention (i.e. the inventive concept revealed by the Patent)?

Q2: Would it be obvious to the person skilled in the art, reading the Patent at the priority date but knowing that the variant achieves substantially the same result as the invention, that it does so in substantially the same way as the invention?

Q3: Would such a reader of the Patent have concluded that the Patentee nonetheless intended that strict compliance with the literal meaning of the claim was an essential requirement of the invention?

69. On the first question, Lighthouse in the request argue that suppression of undesirable visual content before display achieves substantially the same result (redaction) in substantially the same way (AI classification followed by automated suppression prior to display.) Cooley in their observations assert that the claimed invention is directed to a particular process for controlling delivery of content requested by an individual user. Cooley argue that the Unitary content moderation systems approach of returning content is a different architecture. Here, ultimately, I am inclined to agree with Cooley. The architecture of the Unitary content moderation system is significantly different to that described and claimed. That means in my view that whilst the variant may achieve a similar result when used in the way proposed by Lighthouse, it does not do so in substantially the same way as the invention.
70. In the second question, I am asked to assess whether the skilled person would recognise that this is substantially the same approach. I think here the skilled person regards the architecture as being substantially different. The possible benefits/disadvantages of what could be an on-the-fly user specific redaction in response to a content request may be different to those of a pre-prepared content database perhaps with different versions of the content to be supplied to a user with a particular restriction level. Of course, much of what I have just said is about particular embodiments, rather than what is required by the claim language. However, I do think it illustrates the sorts of architectural (that is to say the way the system works) differences between the claimed invention as I have purposively construed it, and the Unitary content moderation system. I therefore again conclude that the second question should be answered “no”.
71. Finally, I should say that whilst there is some information about the intentions that may have sat behind the amendments, I do not think that the reader would have concluded that strict compliance with the literal meaning of the claims was an essential requirement.
72. Having answered no to these three questions, I conclude that there is not infringement via an equivalent.
73. At least for completeness, I should however go on to consider the other points of disagreement.

74. On section B) – Lighthouse describe these points as a derivative of section A) in their observations in reply repeating the architecture point made above. Here, I agree, the construction of these steps is related to the claim as a whole, and I concluded above that they do require something of the functional flow. For the reasons I have just set out on the question of infringement – I reach the same end point.
75. On section C) – Lighthouse in their observations in reply describe a customer configured, multi category policy ad threshold architecture to fall within the scope of a “restriction level”. As I have construed the claims above, the Unitary content moderation system does provide a restriction level to an end user. However, that is only at the secondary stage of the content provider delivering that content to the user. I did not construe the claim as being limited to a single scale though. I do therefore believe that the Unitary content moderation systems customer specific and user case specific configuration amounts to a restriction level.
76. That means that in this respect, it may be that the Unitary content moderation system does not on its own provide content to the end user. However, I think the intention to do so is clear. In this respect the use of the Unitary content moderation system would indirectly lead someone to supply content to a user – using a restriction level.
77. On section D), Cooley argue that the Unitary content moderation system does not provide a score and a confidence level. As I have construed the claim above, both are required separately by the claimed invention. Lighthouse cite from

Unitary API documentation: Threshold Selection



How to select thresholds

What are thresholds and why do I need to use them?

A threshold is a cut-off value used to convert model scores into binary decisions for a category.

If you are using Unitary's product, you will receive outputs in the form of **scores** for every class each product classifies. The score is a number between 0 and 1 that represents how present the item or characteristic is in the content (where 1 is the most present), but the distribution is not normalised e.g. 0.5 may mean "20% confidence" category A and "70% confidence" for category B.

To use these scores to build automated moderation rules or to classify content, you will need to **select thresholds to determine what range of scores represents significance for that characteristic.**

For example, if you want to block toxic content, you will need to decide the minimum toxicity score for which you will define the content as toxic. Depending on your use case, you may want multiple thresholds, for example one for 'definitely toxic - auto-block' and a lower threshold for 'maybe toxic - send to review'.

and the **AWS case study: Unitary on Amazon EKS**

outputs from the inference microservices, classifies the content against customer content moderation policies, and summarizes the results. This summary includes a safety score or risk level for each policy category. The results summary is added to another SQS queue, via an SNS topic, to be sent back to the customer via a webhook. The results are also persisted into a Postgres database by the results writer microservice, acting as a system of record for the content classification outcomes.

78. These passages demonstrate that the score provided in the Unitary content moderation system may have some linkage to confidence, but is silent on what that linkage is. However, as I have noted the claim requires both a score and a confidence level to be considered in determining whether to redact the visual content or part thereof. I note also that Cooley assert that the Unitary content moderation system does not output two distinct metrics. Cooley suggest that the risk level for a category is a score converted into a banded label by reference to the customer's configured threshold.
79. I am not therefore convinced that the passages quoted show that a score and a separate confidence level are calculated. It may be that the safety score is some composite of a degree of prevalence and a confidence score. However, the indication that the distribution is not normalised is not clear evidence that the confidences are actually calculated. It seems plausible to me that the score might be correlated to confidence without a calculation of confidence being made.
80. There is later discussion, which is not highlighted by the parties here, in the glossary and the passage preceding it of the use of recall and precision thresholds (the false positives and missed results.) Both of those factors are ways to determine confidence in your model. However, the documentation I have does not make all of the linkages clear in terms of how these may or may not then be used.
81. On section E), Cooley assert that the Unitary content moderation system does not redact content, stating that pre-approved actions are taken within the customer's system. Lighthouse note that the Unitary API documentation: threshold selection says:

"A threshold is a cut-off value used to convert model scores into binary decisions for a category. ...
For example, if you want to block toxic content, you will need to decide the minimum toxicity score for which you will define the content as toxic. Depending on your use case, you may want multiple thresholds, for example one for 'definitely toxic - auto-block' and a lower threshold for 'maybe toxic - send to review'."

82. Lighthouse therefore I think invite me to assume that the customer of the Unitary content moderation system will therefore provide redacted content.
83. As such, the blocking of content does not require the image to be replaced with a safe image and/or video content. It is plausible that the blocking of content might simply mean that no response is given. It may then be that the user's computer is left to generate its own time-out or similar error message, perhaps as text, or it may be that nothing is displayed. It is plausible that more is done, but that has not been demonstrated in the request. I am not therefore able to confidently say that the use of the Unitary content moderation system will lead to redacted image, or video content is provided in a censored version, as required by the claims.
84. Cooley in their observations go on also to assert that neither the customer, nor the Unitary content moderation system provide the tagging step d) of the claim. Lighthouse quote from the Unitary API docs: Policy Classification endpoints

"For both the Standard and Premium products, Unitary will return in the classification results a raw score that you can use to set thresholds and determine if data is safe, harmful or ambiguous. Please read How to select thresholds if you need help analysing the Standard scores. We'll give you custom thresholds for any premium category we fine-tune to your policy or data. You can contact support@unitary.ai if you wish to explore Unitary's Premium offering.

1. Authenticate Please follow the API Authentication instructions first in order to authenticate with Unitary's API. This will generate a token that is valid for 24 hours and must be used in subsequent API requests."

85. Having construed step d) above, it seems to me that the tagging required is linked to the end user request for content, rather than the content provider. I do not therefore believe that the above passage shows the tagging required by the claim in the context of step d).
86. For completeness here, I should note two further issues: contributory infringement and equivalence, as attempts to bridge the gaps I have identified above. First, I therefore note what is said in the observations by Cooley:

Unitary does not know, and it is not obvious, that its means are intended to put the invention into effect

Section 60(2) requires not only that the means supplied be suitable for putting the invention into effect, but also that the supplier knows, or that it is obvious to a reasonable person in the circumstances, that those means are suitable and intended for that purpose. The Product classifies content submitted to Unitary and, depending on the implementation, either returns that classification to the customer for the customer to act upon or implements a response predetermined by the customer through the customer's API. Neither mode of operation, however, puts the claimed invention into effect: in particular, neither involves the receipt by Unitary of the claimed end-user content request or the tagging of such a request with a

restriction identifier indicative of a restriction level associated with the end user. Unitary therefore cannot be taken to know, nor would it be obvious to a reasonable person in the circumstances, that the Product is intended to be used to put the claimed invention into effect.

and in relation to the equivalents questions in relation to Actavis, Cooley conclude:

Accordingly, the Product is not an immaterial variant of the invention claimed. Question 1 is answered “no”; Question 2 is also answered “no”; and Question 3 provides a further reason why the relevant departures cannot properly be treated as immaterial. The doctrine of equivalents therefore does not extend claims 1 or 8 to the Product.

87. Lighthouse in their observations in reply state

*Nor does section 60(2) require the requester to prove that a named customer has already completed every integer. In Grimme v Scott [2010] EWCA Civ 1110 at [131] the Court of Appeal held that the knowledge and intention requirements of section 60(2) are satisfied “if, at the time of supply or offer of supply, the supplier knows, or it is obvious in the circumstances, that ultimate users **will intend** to put the invention into effect”, and that this is “likely to be the case where the supplier proposes or recommends or even indicates the possibility of such use in his promotional material”. [My emphasis]*

88. I do not think that the evidence provided shows that the required steps for claim 1 follow (directly, automatically) from the extracts provided. The required adaptations are not necessary to implement the Unitary content moderation system. Nor do I think that it is clear that someone implementing the Unitary content moderation system **will** implement in the way required by the claim. As set out in *Grimme Maschinenfabrik GmbH & Co KG v David Scott (t/a Scotts Potato machinery) [2010] EWCA 1110 (Pat)*:

“It is not enough merely that the means are suitable for putting the intention into effect (for that is a separate requirement)”

89. Secondly, as I have discussed above in the context of section A), I think that the answer to each of the Actavis questions is no. That situation is not helped by what I have said since in relation to other issues.

90. In summary, I am not convinced on the basis of the arguments and documentation filed in this request for the Unitary content provider system meeting what is required of the user which is discussed in steps b), c), d), e), h), i), n). Nor am I convinced that the confidence and scoring are used in the way required by steps j) and k) of the claimed invention. Moreover, I am not convinced by the arguments that this leads to contributory infringement, or infringement under the doctrine of equivalents.

91. Finally, I should say that I believe that the same considerations apply to claim 8. I do not think that I need to restate them in order to reach the same conclusions as I have

done in analysing claim 1.

Opinion

92. It is therefore my opinion, that the Unitary content moderation system has not been shown to infringe claims 1 and 8 of the Patent.
93. In reaching that opinion, I have considered the arguments made in relation to direct infringement under Section 60(1), contributory infringement under Section 60(2), and considered the question of equivalence using the Actavis questions, based on the information provided in the request.

Application for review

94. Under section 74B and rule 98, the proprietor may, within three months of the date of issue of this opinion, apply to the comptroller for a review of the opinion.

Robert Shorthouse
Examiner

NOTE

This opinion is not based on the outcome of fully litigated proceedings. Rather, it is based on whatever material the persons requesting the opinion and filing observations have chosen to put before the Office.