IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

AON RE, INC.	:	**CIVIL ACTION**
	:	
v.	:	**NO. 25-201**
	:	
ZESTY.AI, INC.	:	

MEMORANDUM

MURPHY, J. **July 15, 2025**

This is a patent infringement case involving patents that use machine learning to evaluate real property from aerial imagery. Zesty.AI moved to dismiss the case, arguing that the patents are invalid under 35 U.S.C. § 101 because they claim a generic implementation of the abstract idea of assessing property risk from imagery. Recently, the Federal Circuit acknowledged the burgeoning field of artificial intelligence but held that patent owners may not claim the mere application of generic machine learning to new data environments. Aon's patents do not offer a new twist on machine learning itself. But that is not fatal, because we agree with Aon that the patents recite the patent-eligible arrangement of two independently trained classifiers to analyze property characteristics and conditions. For the reasons that follow, we deny Zesty.AI's motion to dismiss.

I. Background

Aon Re, Inc. accuses Zesty.AI, Inc. of infringing four similar U.S. patents: 10,529,029, 10,650,285, 11,030,491, and 11,195,058. *See* DI 1. The patents describe computer systems that use machine-learning classifiers to analyze aerial photographs of buildings (e.g., homes and commercial structures); automatically identify features (e.g., the type of roof); and determine their condition (e.g., in good shape or needs repair). *See, e.g.,* 029 patent at Abstract, 1:18-36, 2:24-35. The invention then uses these identified features and condition assessments to evaluate

how susceptible a given property might be to damage during storms or other adverse natural conditions, which is useful information for insurance companies like Aon in assessing risk and pricing coverage. *See id.* at 11:47-12:16; DI 1 at 1.

The patents explain that machine learning classifiers are computer software models that can evaluate input data and classify it into categories. *See* 029 patent at 1:37-2:5. The classifier learns to make these assignments by training on examples with known answers. *See id.* at 1:46-49, 8:37-41. This enables it to recognize relevant features and sort new, unseen inputs. *See id.* at Fig. 2D, 1:46-49, 9:4-23. The classifier represents categories as complex mathematical abstractions. *See id.* at 1:37-39. When applied to image analysis, these classifiers mimic "the biological process of visually reviewing and identifying an object or feature of an object." *Id.* at 8:16-26. Depending on the implementation, a classifier may evaluate image data at the pixel level (such as through sets of intensity values) or at a more abstract level using edges or region shapes. *Id.* at 1:61-67.

Zesty filed a motion to dismiss, arguing that Aon's patents claim subject matter that is not eligible for patent protection under 35 U.S.C. § 101. *See* DI 17. Zesty's motion primarily targets claim 1 of the 029 patent. *See* DI 19 at 5-10, 15-16. The parties dispute whether that claim is fairly representative, but we may assume that it is. *See* DI 25 at 6-7. Representative claim 1 of the 029 patent is reproduced below:

> 1. A method for automatically categorizing a repair condition of a property characteristic, comprising:
>
>> receiving, from a user at a remote computing device, a request for a property condition classification, wherein the property classification request includes identification of a property and at least one property characteristic;

> obtaining, by processing circuitry of a computing system responsive to receiving the request, an aerial image of a geographic region including the property;
>
> extracting, by the processing circuitry, one or more of a plurality of features from the aerial image corresponding to the property characteristic, wherein the extracted features include pixel groupings representing the property characteristic;
>
> determining, by the processing circuitry from the extracted features, a property characteristic classification for the property characteristic, wherein determining the property characteristic classification includes applying the pixel groupings for the property characteristic to a first machine learning classifier trained to identify property characteristics from a set of pixel groupings;
>
> determining, by the processing circuitry based on the identified property characteristic and the extracted features, a condition classification for the property characteristic, wherein identifying the condition classification includes applying the pixel groupings for the property characteristic to a second machine learning classifier trained to identify property characteristic conditions from a set of pixel groupings;
>
> determining, by the processing circuitry based in part on the property characteristic classification and the condition classification, a risk estimate of damage to the property due to one or more disasters; and
>
> returning, to the user at the remote computing device via a graphical user interface responsive to receiving the request, a condition assessment of the property characteristic including the condition classification and the risk estimate of damage to the property due to the one or more disasters.

The method of claim 1 begins by accessing an aerial image on a computer in response to a user request. 029 patent at 24:17-26. Next, the method isolates and removes from that image groups of pixels "such as angles, outlines, substantially homogenous fields" that represent a certain physical feature, like the roof, of the property. *Id.* at 7:60-8:1, 24:27-31. The method then analyzes those features using two different machine learning algorithms, both of which require using the extracted pixels as input data. *See id.* at 24:32-47. One algorithm is trained

specifically to identify a particular property characteristic — for instance, determining whether the roof is gambrel, gable, hipped, square, or flat. *See id.* at Fig. 2A, 1:31-36, 8:50-9:3, 24:32-39. The second, separate algorithm assesses the condition of that characteristic — for example, evaluating whether the roof is in good repair or shows signs of damage or deterioration. *See id.* at Fig. 2B, Fig. 2C, 10:9-11:31, 16:12-31, 24:40-47. The final step is outputting a "risk estimate," which requires predicting how likely or severe damage to the property might be from specific disasters or weather conditions, such as hurricanes or strong winds. *See id.* at 11:47-12:16, 24:48-27. According to the patent, this method helps its users obtain a quick and consistent evaluation of property risks based solely on aerial images. *See id.* at 2:6-11.

Zesty argues that the method in claim 1 is directed to the abstract idea of "looking at aerial imagery to estimate risk." DI 19 at 5; DI 26 at 1-2. According to Zesty, the patent does nothing more than implement generic machine-learning technology to automate tasks traditionally performed by human inspectors. *See* DI 19 at 11. Zesty analogizes these claims to others that courts have found ineligible for patent protection because they merely automated longstanding human practices using conventional technology and were not directed to improvements in computer or machine-learning technology. *See id.* at 1-2, 6-7, 11, 13-14 (citing *Recentive Analytics, Inc. v. Fox Corp.*, 134 F.4th 1205 (Fed. Cir. 2025)).

Aon responds that Zesty failed to carry its burden of proving the abstractness of the claims, and that claim 1 is not abstract because it recites an "improved computer architecture." *See* DI 25. According to Aon, the claim's use of two distinct machine learning algorithms working independently and simultaneously to identify different property information from the same image goes beyond merely automating existing practices with generic technology. *See id.*

at 9, 12-13 (citing *McRO, Inc. v. Bandai Namco Games Am. Inc.*, 837 F.3d 1299, 1314-15 (Fed. Cir. 2016)). Instead, Aon argues, this approach meaningfully improves the speed and efficiency of property assessments compared to prior computer-implemented techniques. *Id.*

II. Legal Standard

Patent eligibility under § 101 encompasses "any new and useful process, machine, manufacture, or composition of matter, or any new and useful improvement thereof." 35 U.S.C. § 101. Abstract ideas, laws of nature, and natural phenomena are "implicit exception[s]" to this broad standard for eligibility. *See Alice Corp. Pty. v. CLS Bank Int'l*, 573 U.S. 208, 216 (2014).

To determine patent eligibility, we apply the two-step framework established by the Supreme Court in *Mayo* and *Alice*. *See id.* at 217 (citing *Mayo Collaborative Servs. v. Prometheus Labs., Inc.*, 566 U.S. 66 (2012)). First, we evaluate whether the claim at issue is "directed to a patent-ineligible concept" such as an abstract idea. *Id.* at 217-18. If not, the inquiry ends, and the claims are patent eligible. If the claims are directed to a patent-ineligible abstract idea, we move to step two and ask whether the claims nonetheless contain an "inventive concept" — additional elements "sufficient to transform the abstract idea into a patent-eligible invention." *Id.* at 221.

Step one requires us to consider the "focus of the claimed advance over the prior art to determine whether the claim's character as a whole is directed to excluded subject matter," such as, relevant here, an abstract idea. *Affinity Labs of Tex., LLC v. DIRECTV, LLC*, 838 F.3d 1253, 1257 (Fed. Cir. 2016) (citation modified). There is no compendium of abstract ideas. Instead, we are encouraged to analyze the claim in context, drawing on analogies to past decisions when possible. *See Enfish, LLC v. Microsoft Corp.*, 822 F.3d 1327, 1334 (Fed. Cir. 2016); *see also*

Veeva Sys. Inc. v. Tact.ai Techs., Inc., No. 23-1032, 2024 WL 2848335, at *4-5 (D. Del. June 5, 2024).

Preemption is a crucial concern underlying this comparative analysis. *See Alice*, 573 U.S. at 216. A claim that effectively forecloses all practical applications of an abstract idea risks impeding, rather than promoting, future innovation. *See McRO*, 837 F.3d at 1315. And courts often consider the risk of preemption in cases involving computer-implemented methods and systems. Because computers are foundational tools possessing near-universal applicability, claims that invoke computer technology in sweeping or generalized terms can effectively monopolize its use in an entire area of invention. Accordingly, we must carefully distinguish between claims that merely harness generic computing tools to lay claim to an abstract idea, and those that marshal those tools in a specific, structured way to solve a concrete technological problem. *See In re TLI Commc'ns LLC Pat. Litig.*, 823 F.3d 607, 611 (Fed. Cir. 2016) ("[I]n determining whether the claims are directed to an abstract idea, we must be careful to avoid oversimplifying the claims because 'at some level, all inventions embody, use, reflect, rest upon, or apply laws of nature, natural phenomena, or abstract ideas.'" (citation modified) (quoting *Alice*, 573 U.S. at 217)). As courts have understood *Alice*, only the latter fall within the scope of patent-eligible subject matter.

At one end of the spectrum, claims that describe routine steps and invoke generic computer tools without identifying how the technology is applied in a specific or inventive way are typically found abstract. *See, e.g., Recentive*, 134 F.4th at 1211-13 (claims ineligible where they generically applied machine learning to schedule TV content without explaining how models were inventively trained, structured, or implemented); *Int'l Bus. Machs. Corp. v. Zillow*

Grp., 50 F.4th 1371, 1380-82 (Fed. Cir. 2022) (claims ineligible where they used standard data visualization tools to display real estate information without reciting any inventive implementation); *ChargePoint, Inc. v. SemaConnect, Inc.*, 920 F.3d 759, 768-70 (Fed. Cir. 2019) (claims ineligible where they recited remote control of electric vehicle charging stations over a network but failed to claim any specific advance in networking, charging, or computing technology).

At the other end of the spectrum, claims that solve specific technical problems, such as particular improvements to the functioning of computers, networks, or software that are useful for an application of those technologies, are usually found eligible. *See McRO*, 837 F.3d at 1315-16 (claims eligible where they used specific phoneme-to-morph rule sets to automate facial animation, improving animation workflow through a novel algorithmic structure); *Koninklijke KPN N.V. v. Gemalto M2M GmbH*, 942 F.3d 1143, 1151-52 (Fed. Cir. 2019) (claims eligible where they applied defined permutations to transmitted data blocks to detect systematic errors, enhancing data integrity); *CardioNet, LLC v. InfoBionic, Inc.*, 955 F.3d 1358, 1371–72 (Fed. Cir. 2020) (claims eligible where they used a specific sensor configuration and signal-processing logic to improve detection of atrial fibrillation over conventional monitors). The common thread among these cases finding eligibility is that the claims recite more than a result-oriented goal or functional description. The claims provide solutions in specific, non-conventional, or technical ways, confining the scope of the claims and avoiding preemption.

Contesting patent eligibility under § 101 is an affirmative defense of invalidity, and the burden of proving invalidity falls squarely upon Zesty, the movant. *See Commil USA, LLC v. Cisco Sys., Inc.*, 575 U.S. 632, 644 (2015); 35 U.S.C. § 282(a). Zesty must articulate with

specificity the alleged abstract idea and demonstrate that the asserted claims are directed to that idea. *Veeva*, 2024 WL 2848335, at *4-5. Because, as discussed below, Zesty did not persuade us that the claims are directed to an abstract idea, our analysis does not proceed beyond step one.

III. Analysis

Zesty argues that claim 1 of the 029 patent is directed to the abstract idea of "looking at aerial imagery to estimate risk." DI 19 at 5; DI 26 at 1-2. We confine our analysis to Zesty's proposal because, as explained above, Zesty bears the burden of persuasion.

Zesty's position has some facial appeal. Aon's patents acknowledge that humans have visually assessed property condition and estimated risks. Further, Aon does not suggest that it invented the computers and machine-learning classifiers mentioned in the claim. But step one of the *Alice* inquiry demands more than a generalized assessment. The relevant question is whether claim 1, as drafted, merely invokes generic computing tools to carry out an established human practice, or whether it recites a specific and limited implementation that solves a concrete problem using technology in a defined way. *See Enfish*, 822 F.3d at 1335-36; *McRO*, 837 F.3d at 1314-15.

Claim 1 falls into the latter category of patent-eligible implementations. It does not broadly claim evaluating risk using imagery. Rather, it recites a process with a distinctive structure: two separate and independently trained machine-learning classifiers that operate independently. One identifies a physical characteristic of a property (such as roof shape), and the other evaluates the condition of that same characteristic (such as signs of deterioration). *See* 029 patent at 24:32-47; *see also id.* at 8:36-9:23, 10:9-11:31, 15:50-16:49. Further, those classifiers must work by making use of pixel groupings obtained from digital images and

produce a specific output: "a condition assessment of the property characteristic including the condition classification and the risk estimate of damage to the property due to the one or more disasters." *Id.* at 24:27-31, 24:54-57; *see also id.* at 11:47-12:15, 16:50-18:54.

This architecture is significant in the eligibility analysis. As the patent explains, the separate application of two independently trained classifiers provides efficiency and reliability in identifying both the characteristic and that characteristic's condition. *Id.* at 9:4-23, 11:12-31. The claim thus implements machine-learning technology in a specific way to address a practical technical problem: how to accurately and quickly assess the physical characteristics and condition of a property from an aerial image using machine learning models. *See id.* at 2:6-11.

Importantly, claim 1 does not simply recite the result itself or broadly direct the use a computer and machine learning to do risk analysis. The claim's specificity distinguishes it from claims in other cases that were deemed abstract for merely applying machine learning or data visualization without disclosing any technology-specific method. *See Recentive*, 134 F.4th at 1212–13 (invalidating claims that "do not delineate steps through which the machine learning technology achieves an improvement"); *Zillow*, 50 F.4th at 1381 (claims ineligible where they "do[] not sufficiently describe how to achieve the[ir] results in a non-abstract way" (quoting *Two-Way Media Ltd. v. Comcast Cable Commc'ns, LLC*, 874 F.3d 1329, 1337 (Fed. Cir. 2017)).

Briefly consider cases where claims were found to be abstract. In *Recentive*, the patentee claimed a method of optimizing event schedules using machine learning but failed to identify a specific technical implementation. The claims broadly referenced iterative training and real-time updating without specifying the "steps through which the machine learning technology achieves an improvement." *Recentive*, 134 F.4th at 1212-13 ("Allowing a claim that

functionally describes a mere concept without disclosing how to implement that concept risks defeating the very purpose of the patent system."). In *Longitude Licensing*, claims were found abstract because they recited broadly applying machine learning to information retrieval without detailing "'how the purported invention improve[s] the functionality' of image correction methods." *Longitude Licensing Ltd. v. Google LLC*, No. 2024-1202, 2025 WL 1249136, at *2-3 (Fed. Cir. Apr. 30, 2025) (quoting *Hawk Tech. Sys., LLC v. Castle Retail, LLC*, 60 F.4th 1349, 1358 (Fed. Cir. 2023)). And in *TLI*, claims were ineligible because the patent "predominately describe[d] the system and methods in purely functional terms" and failed to provide details for how the components achieved any improvement over prior art. 823 F.3d at 612-13 ("[The claims] are directed to the use of conventional or generic technology in a nascent but well-known environment, without any claim that the invention reflects an inventive solution to any problem presented by combining the two.").

In contrast, claim 1 of the 029 patent recites a concrete and specific technological solution. Rather than referring to the use of machine learning in general, the claim requires individually trained classifiers that operate independently, performing separate and clearly defined analytical tasks: one classifier identifies a specific physical characteristic of a property from aerial imagery, and the other independently assesses the condition of that characteristic. This division of labor, together with the requirement to analyze defined pixel groupings and produce risk assessment outputs based on the classifiers' assessments, confirms that the claim represents more than a generalized use of machine learning. Although the patent does not purport to invent machine learning algorithms, computer hardware, or the general concept of aerial imagery analysis or risk assessment, it does specifically define how these known

components must be structured to achieve a supposed technological improvement in accuracy and efficiency to the combination of those components.

This type of defined structural constraint closely aligns with the claims held eligible in *McRO* and *Koninklijke*. In *McRO*, the Federal Circuit upheld claims that related to automating the animation of lip synchronization by matching animated characters' mouth movements to spoken audio. *See McRO*, 837 F.3d at 1314-16. Rather than merely instructing animators to synchronize audio with animation using a computer generally, the claims required a rule set that established detailed relationships between spoken sounds and corresponding facial movements. *See id.* While the exact rules were not claimed, and the invention employed only "general-purpose computers" and "limited mathematical rules," the claims nonetheless defined how these components work together to achieve automated lip synchronization. *Id.* at 1314. Thus, the court found the claims eligible because they covered a concrete technological method for automating animation, rather than claiming the mere result of lip synchronization. *See id.*

Similarly, in *Koninklijke*, the claims described a method of improving error detection in data transmissions. Prior error detection systems sometimes failed to detect a certain type of persistent error. To solve this problem, the claims required varying the way error-checking data was generated by applying different permutations (rearrangements of data bits) to different blocks of transmitted data over time. *Koninklijke*, 942 F.3d at 1150-52. Rather than merely claiming the abstract idea of error detection or calling for permutation generally, the requirement that permutations be modified in time provided a sufficiently specific way of detecting persistent errors. *See id.* at 1152.

Here, the independent training and distinct analytical roles assigned to each of the two

machine learning classifiers provide a structured technological method specific to analyzing pixel groups in aerial imagery for property risk assessments. Like the claims upheld in *McRO* and *Koninklijke*, claim 1 clearly defines how known technological components must be specifically structured and combined, rather than broadly claiming a purported abstract idea of applying machine learning to aerial images to assess risk. By assigning separate and clearly defined analytical tasks to the two independently trained classifiers, claim 1 provides a particular structured solution that the patent says is accurate and efficient. Thus, claim 1 is not directed to the abstract idea advanced by Zesty.

This conclusion is not contrary to *Recentive*'s holding "that patents that do no more than claim the application of generic machine learning to new data environments, without disclosing improvements to the machine learning models to be applied, are patent ineligible under § 101." 134 F.4th at 1215-16. To be sure, Aon or a similarly situated patent owner could follow that route to patent eligibility and claim specific improvements to the machine-learning models useful for addressing property risk assessment. But another route to eligibility, as we have explained, is claiming a specifically arranged implementation of machine learning, analogous to the patent-eligible arrangements of *McRO* and *Koninklijke*. Those two decisions were not helpful to the *Recentive* court, because in *Recentive*, the Federal Circuit viewed the claims as reciting only generic machine learning rather than a purportedly novel arrangement. *Id.* at 1213 (distinguishing *McRO* and *Koninklijke*).

Our holding is also consistent with the preemption concerns discussed in *Alice* and its progeny. Claims should not risk monopolizing an entire field by claiming an abstract idea in functional or result-oriented terms. *See Alice*, 573 U.S. at 216. Zesty has not persuaded us that

Aon's claims pose such a risk. Claim 1 of the 029 patent requires a specific classifier arrangement that must analyze pixels from aerial images in a defined sequence to generate risk assessments. Zesty does not explain why this fails to leave room for alternative methods, and the record, albeit very limited at this stage, appears to suggest the opposite conclusion. *See McRO*, 837 F.3d at 1315-16 ("[T]he description of one set of rules does not mean that there exists only one set of rules, and does not support the view that other possible types of rules with different characteristics do not exist."). Nor does this holding run afoul of *Alice*'s warning against narrow (yet still ineligible) abstract ideas because Zesty never suggested that Aon's two-classifier structure was itself an ineligible abstract idea. *See Alice*, 573 U.S. at 220.

Claim 1 is not particularly hard to understand; Aon does not pretend that it is. But intricacy must not be the principal measure of eligibility — why should elegant inventions be prejudiced? Claim 1 presents a targeted, technical solution to a specific problem and reflects more than a generic invocation of machine learning. Thus, it is directed to patent-eligible subject matter. And because Zesty has not shown that claim 1 is directed to an abstract idea, our *Alice* inquiry ends at step one. The motion to dismiss is denied. An appropriate accompanying order follows.