

Robust AI Methods for Protein Biotherapeutics

VisualAI™ software for sample-agnostic image analysis with FlowCam

INTRODUCTION

Subvisible particles in therapeutic protein formulations and other biotherapeutics are a significant product quality concern for pharmaceutical companies. These concerns stem not just from regulations such as USP <787> that focus on subvisible particle counts but also from the potential changes in product stability, quality, and safety that are associated with subvisible particle content^{1,2}. Particle analysis techniques such as Flow Imaging Microscopy (FIM) are used to monitor these particle populations through many stages of biotherapeutic development as well as during regular production³⁻⁵. While different technologies will often measure different features of these particle populations (e.g. concentration, size, structure), they all are ultimately used to assess changes in the particle content that could preemptively indicate compromised product quality and safety.

One important sample property to monitor is the types of particles present in a protein formulation and their relative concentrations. These formulations can exhibit particles from a variety of sources ranging from aggregates of the protein⁶ to particles from the container-closure system used to store the formulation (e.g. silicone oil droplets, glass flakes)^{7,8} to impurities introduced during the manufacturing process (e.g. metal flakes). FlowCam, an FIM instrument, is increasingly prominent for this type of monitoring. FlowCam captures light microscopy images of the particles present in a sample which can be analyzed to determine particle concentrations as well as particle morphology. As different particle types often exhibit visually-distinct morphologies in these images, FlowCam images can be used to identify the different particle types in a sample⁹ (FlowCam images of common biotherapeutic particle types can be seen in Figure 1). Additionally, many common artifacts with these techniques such as air bubbles and calibration beads often result in images of distinct particle morphologies as well—information that can be used to remove images from the analysis and obtain more accurate particle concentration measurements.




Figure 1. FlowCam images of particles commonly found in protein formulations




The most common use case for this morphology information is the identification of protein aggregates and silicone oil droplets in a formulation. Protein aggregates are ubiquitous in protein formulations due to the limited stability of proteins in an aqueous environment¹⁰ and the stresses proteins encounter during regular manufacturing and handling^{11,12}. Silicone oil droplets are commonly introduced by silicone oil-lubricated syringes often used to administer these therapies to patients⁸. While silicone oil droplets have shown to be harmful in some studies¹³, they are generally thought to be more benign than protein aggregates for patients. It is therefore useful to be able to differentiate between these particle types when making product quality decisions about a sample containing both particle types. This particle differentiation is also a natural application for FlowCam due to how visually distinct these particle types are: silicone oil droplets often appear circular when imaged while protein aggregates appear more irregular and amorphous^{14,15}, as can be seen in Figure 1.

The distinct morphology differences between protein aggregates and silicone oil droplets in conjunction with the ubiquity and practical importance of these particles have motivated researchers to develop automated algorithms for identifying protein aggregates and silicone oil images in FIM datasets. Many early techniques proposed for this analysis primarily relied on particle measurements from FIM images reported by the instrument^{14,15}. These techniques used measurements related to particle size, color intensity, and roundness (i.e. aspect ratio and circularity on FlowCam) to identify images of round silicone oil droplets. Similar methods are also used

by the onboard classification utilities in FlowCam’s software, VisualSpreadsheet®. While these methods can be effective, these approaches only use a small fraction of the particle morphology information in the image—information that could further differentiate these two particle types. These methods also can be sensitive to the software settings used to detect particles in a sample as they can influence the particle property values used in the classification.

Recently, artificial intelligence (AI)-based methods have shown promise for this analysis⁹. These techniques use the raw images rather than measured particle properties to differentiate between these particle types, using more of the information in the image to identify particles. Additionally, using the raw images directly results in particle classifications that are more robust to instrument settings. However, AI-based methods are more difficult to implement for most users. These methods typically require users to collect thousands of images of each particle type obtained on their FIM instrument which are then used to train the AI. This process is time-consuming for the analyst and often requires the help of researchers versed in AI to guide the data collection and perform the AI training. Furthermore, these tools have typically been developed to analyze images from a small selection of drug products on a single FIM instrument, resulting in tools that may not be as effective across multiple instruments and/or other drug products.

 VisualAI™ is an add-on software package that integrates into VisualSpreadsheet and is designed to streamline AI-driven protein aggregate-silicone oil droplet identification. VisualAI is a pre-trained AI utility for identifying grayscale images of protein aggregates and silicone oil droplets larger than 3 μm in diameter in FlowCam 8100 and FlowCam LO datasets. This software gives users much more accurate protein aggregate-silicone oil differentiation than is possible with particle property-based classification utilities. Unlike other AI-based

methods for subvisible particle analysis, VisualAI is designed to be used “off the shelf”; the user does not need to provide any training protein aggregate or silicone oil images to obtain accurate classifications on their system and samples. As an additional feature, VisualAI offers simple anomaly detection tools to identify images of some common particle types that are neither protein aggregates nor silicone oil droplets. This feature can be used to identify images of particles such as polystyrene calibration beads and air bubbles—images that often show up and can artificially raise the measured particle concentration. Combined, these features make VisualAI a powerful, robust tool – yet exceptionally easy to use for analyzing FlowCam data for therapeutic protein formulations.

This white paper demonstrates the classification performance of VisualAI on samples known to contain aggregates of different proteins, silicone oil droplets, or calibration beads. Images in this study were collected on three different FlowCam 8100 units, none of which were used in the development of the software. As VisualAI is sample-agnostic, this testing was also performed using aggregates of several different proteins. We also show the performance of the software when analyzing samples containing a known fraction of protein aggregates and silicone oil droplets on a fourth FlowCam 8100 unit to demonstrate the accuracy of the protein aggregate-silicone oil composition measurements with VisualAI.

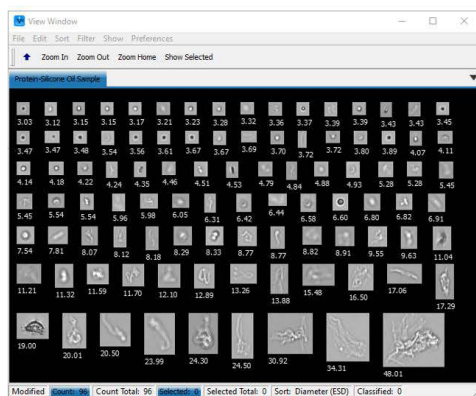


Figure 2. The graphical user interface of VisualAI. (Left) Sample FIM image collage from a sample containing 50% protein aggregates 50% silicone oil droplets. This sample was prepared as described by this study. (Right) Collages in VisualAI’s graphical user interface showing the images identified by AI utilities that contain (top) protein aggregates, (middle) silicone oil droplets, and (bottom) other particles.



METHODS

Sample Preparation: The performance of VisualAI was assessed on samples containing purely protein aggregates and silicone oil droplets suspended in phosphate buffered saline (PBS) that were not used in the development of VisualAI. Aggregates of four different proteins were prepared: bovine serum albumin (BSA), ovalbumin, intravenous immunoglobulin (IVIg), and the NIST monoclonal antibody (mAb). 1 mg/mL BSA and ovalbumin formulations were prepared by mixing 12 mg lyophilized protein powder with 12 mL PBS and gently shaking the resulting solution on a plate rocker for 30 minutes. Formulations of the remaining proteins were prepared by diluting the stock protein formulation into PBS to prepare 12 mL of the formulation. 5 mL aliquots of these protein formulations in 15 mL conical tubes were then exposed to one of two accelerated stability stresses to induce aggregation: freeze-thaw stress and shaking stress. Freeze-thaw stress was performed by exposing samples to four freeze-thaw cycles with each cycle consisting of a 30-minute freeze at -20 °C followed by a 10-minute thaw at 25 °C. Shaking stress was performed by taping samples to a plate rocker and agitating the samples at the maximum speed and angle of the rocker for four hours. Silicone oil emulsions were prepared by creating a 10% (v:v) silicone oil solution in PBS and blending the resulting formulation in a laboratory blender at max speed for 20 seconds. The resulting emulsion was then diluted 1:1000 with PBS to obtain an appropriate droplet concentration for FlowCam analysis.

FlowCam Analysis: Four different FlowCam 8100 units were used to analyze samples containing each of the four protein types. A fresh silicone oil sample was also prepared and analyzed in parallel with each protein sample. Each FlowCam unit was equipped with a 10X objective and a FOV80 flow cell. Prior to analysis, the instrument's fluidics were flushed with a 10% Hellmanex III solution, followed by water, then followed by PBS. The instruments were operated at the recommended capture settings for protein samples and VisualAI; 15 dark and light pixel thresholds with 3 close hole iterations, 4 μm distance to nearest neighbors, and rolling calibration disabled. The background intensity was set to approximately 170. Each FlowCam unit was autofocused for best image quality using 15 μm polystyrene latex calibration beads. An autoimage mode run was performed on calibration beads following autofocusing to confirm that the instrument was in focus. To test the robustness of the algorithm, the NIST mAb samples were analyzed on an instrument that was not focused prior to measurements—a potential error that may be encountered when collecting data for particle morphology analysis. The instruments were then cleaned as described above and used to analyze a PBS “blank” sample which contained negligible particle concentrations to assess the background particle content in the fluidics. This cleaning and blank sample analysis process was repeated until fewer than 1,000 particles/mL were detected in the PBS sample.

Three 1 mL aliquots per sample were analyzed on each of the four FlowCam 8100 units. Most of the protein samples were analyzed at a 0.15 mL/min flow rate and at a 22-frames-per-second autoimage rate. Ovalbumin samples and the corresponding silicone oil images were analyzed at the FlowCam LO default flow rate (0.2 mL/min) using a 30-frames-per-second autoimage rate to compensate for the faster flow rate. The fluidics were cleaned between measurements using the method described earlier. Additional cleaning and PBS blank analysis as described earlier was performed between pure protein aggregate and silicone oil samples to ensure minimal crossover between sample types. All FlowCam datasets were post-processed to filter images of particles smaller than 3 μm based on diameter (ESD) as well as to remove obvious artifacts like air bubbles to ensure all images were made up predominantly of the particle type of interest.



VisualAI Classifier Accuracy: Samples of IVIg, Ovalbumin, and NIST mAb aggregates along with their paired silicone oil samples were analyzed via VisualAI to assess the classification performance on samples known to contain only protein aggregates, silicone oil microdroplets, or calibration beads. This analysis was also performed on particles with diameter (ESD) values in different size ranges: 3-5 μm , and >5 μm . Protein aggregate and silicone oil droplet images 2-3 μm in size were excluded from this analysis. While VisualAI will often achieve classification accuracies between 80-90% in this size range, the software's accuracy is more instrument- and sample-dependent in this range than for larger particles. It is recommended to restrict VisualAI analysis to particles larger than 3 μm .

As an additional comparison, the performance of VisualAI was benchmarked against simplified versions of previously-described non-AI methods for silicone oil identification: S-factors-style classification¹⁴ and random forest classifiers¹⁵. In both methods, particle properties rather than raw FlowCam images are analyzed to determine if an image contains a protein aggregate or a silicone oil droplet. Both methods also train separate classifiers on particles of different diameter (ESD) ranges to account for changes in apparent

particle morphology as the size increases. Particles were divided into only two size bins for comparative analysis: 3-5 μm and $>5 \mu\text{m}$. Both methods were trained using the same FlowCam data used to develop VisualAI, weighting the loss functions for both methods based on the relative amount of training data for each class in each size bin. The trained methods were then applied to the FlowCam data collected as part of this study to assess the overall accuracy of each method relative to VisualAI.

S-factor-based classification was performed by computing a modified S-factor for each particle and using that value to predict the identity of the particle. Modified S-factor values were computed using each particle's aspect ratio, circularity, and sigma intensity. In lieu of using manually optimized threshold S-factor values to perform the final classification, particle classification was performed by training logistic regression classifiers on the S-factor values for the training data in each size bin and using the size-appropriate trained classifier to analyze subsequent images based on their S-factor.

Random forest classifiers were trained to classify images based on their diameter (ESD), aspect ratio, circularity, (mean) intensity, sigma intensity, and edge gradient—properties that, excluding edge gradient, were identified as being useful for distinguishing protein aggregates and silicone oil droplets in the original study¹⁵. Values of each particle property were normalized to the mean and standard deviation of that property over the training set during both training and testing to account for the different numerical ranges each particle property can assume. Random forests consisting of 128 trees were trained on particles in each size range. The trained classifiers were applied to the data collected during this study in the same fashion as the S-factor classifiers.

VisualAI Composition Estimate Accuracy: The primary use case for VisualAI is determining the relative composition of protein aggregates and silicone oil microdroplets in a sample. To demonstrate the software's performance for this application, samples containing a known ratio of protein aggregates and silicone oil droplets were imaged and analyzed with the software to estimate the composition of each sample. Samples containing pure BSA aggregates generated via shaking stress and silicone oil microdroplets were prepared and analyzed as described above to determine the particle concentration in each sample. Particles between 2-3 μm in equivalent spherical diameter (smaller than the lower size limit for VisualAI) were not excluded from this concentration estimate. Once measured, these samples were mixed and diluted with PBS to prepare mixtures containing approximately 100,000 particles/mL and either 25%, 50%, or 75% protein aggregates and the remainder silicone oil droplets. Three samples were prepared per mixture composition. Each of these mixed samples were prepared and immediately analyzed on FlowCam as described above to minimize any silicone oil-induced protein aggregation. Images of all particles larger than 3 μm from each mixture, as well as the pure samples, were analyzed to assess the agreement between the predicted and actual particle composition in each sample.

RESULTS & DISCUSSION

FlowCam Analysis: Samples containing only protein aggregates or silicone oil droplets were analyzed via four different FlowCam 8100 units. Figure 3 shows sample protein aggregate and silicone oil droplet images from three of the FlowCam units tested. As expected, protein aggregates and silicone oil droplets imaged on a single FlowCam typically exhibit distinct particle morphologies; silicone oil droplets are typically circular with a clear concentric ring structure while protein aggregates are amorphous and irregular. While these distinct particle shapes are often easy for human operators to recognize with training and practice, some protein aggregate and silicone oil droplet images can be difficult for even experienced operators to differentiate. Some protein aggregate images can exhibit round structures resembling that of silicone oil images. Out-of-focus silicone oil droplets often exhibit an asymmetric or otherwise obfuscated concentric ring structure that can be difficult to distinguish from protein aggregates.

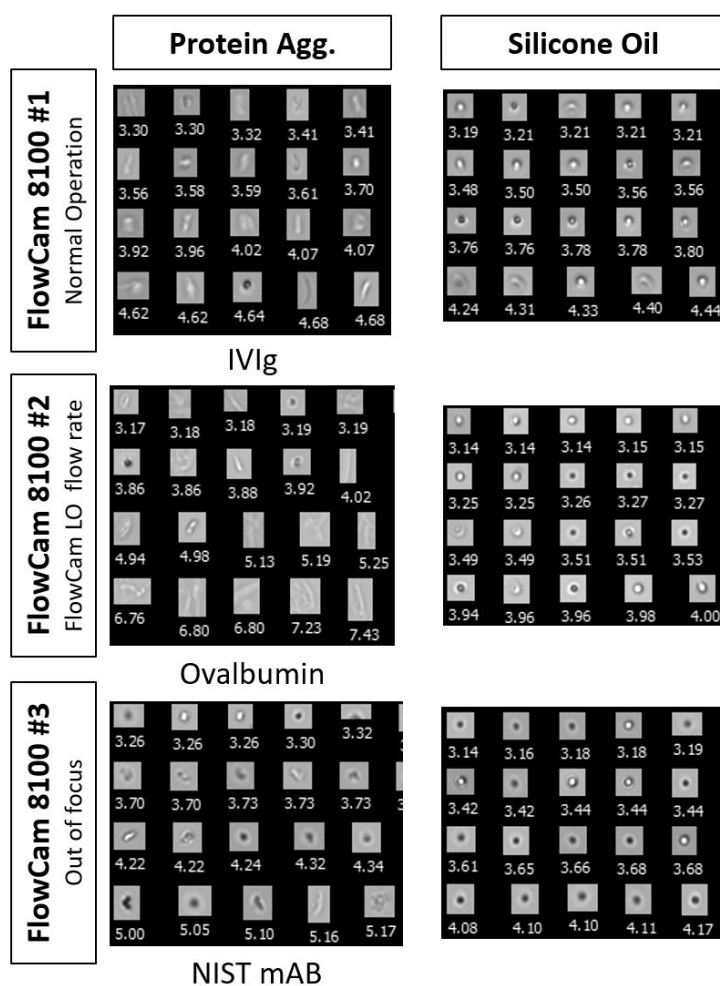


Figure 3. Sample images of protein aggregates and silicone oil droplets captured on three different FlowCam 8100 units. Each row corresponds to data on a single unit. Relevant FlowCam operation details are included in the header for each row. Protein aggregate images were generated from the protein denoted below each collage and consist of a mixture of aggregates generated by freeze-thaw and shaking stresses.

These discriminative morphological features can also change depending on the focus position of the unit: particle images from the out-of-focus FlowCam (Figure 3, bottom row) are much darker on average and often miss hallmarks such as the bright spot in the middle of many silicone oil droplets. These issues and others can complicate what is theoretically a simple particle classification task especially as the particle diameter decreases and the two particle types inherently appear more similar.

VisualAI Classification Accuracy: FlowCam images taken from samples containing only protein aggregates and silicone oil droplets were analyzed via VisualAI to assess the classification accuracy of the software. Figure 4 shows confusion matrices indicating the classification accuracy of the software on particles like those shown in Figure 3. These tables demonstrate that VisualAI achieves near if not above 90% classification accuracy across samples, instruments, and most size ranges. Importantly, the software achieves this performance both at the “recommended” 0.15 mL/min flow rate as well as the higher 0.2 mL/min flow rate that is required for FlowCam LO operation. The accuracy of the classification improves with increasing particle size, achieving around 95% classification accuracy on particles >5 μm in size but decreasing by 1-3% on those 3-5 μm in size for most units. This behavior is expected as images of small protein aggregates resemble those of small silicone oil droplets. Since these small particles are difficult to classify even for human FlowCam operators, it is unsurprising that VisualAI’s performance degrades for smaller images in typical FlowCam datasets.

A surprising result from this analysis was VisualAI’s accurate classification of FlowCam images of protein aggregates and silicone oil droplets captured on a unit that was not properly focused before image capture. This demonstrates the robustness of VisualAI, as data collected outside of an optimally configured instrument can still effectively be analyzed within the size range of the software. While autofocusing the instrument is still strongly recommended before collecting data for analysis via VisualAI, data collected without optimal focusing can still be effectively analyzed via VisualAI.

Figure 4. Confusion matrices showing the classification performance of VisualAI as well as S-factor- and random forest-based methods on protein aggregates and silicone oil droplets of different sizes collected on three separate FlowCam units. The confusion matrices for each classifier and particle size are placed side-by-side in each table for ease of comparison and are delineated by the row and column headers. Blue, red, and green-colored matrices correspond to confusion matrices for VisualAI, S-factor, and random forest methods, respectively. The protein used to generate aggregates as well as any unique features of the data (higher flow rate or unfocused images) are shown to the left of each matrix.

Confusion matrices like these show the fraction of images of each particle type (matrix rows) that were classified as the three possible labels each classifier can return (matrix columns). The accuracy of the classifiers for each sample and size range represents the fraction of particles in that sample that was assigned the correct label (e.g. images of protein aggregates being classified as protein aggregates). The coloration of each cell reflects the fraction of particle images of each type that were assigned the corresponding label with darker colors indicating higher fractions.

		VisualAI			S-Factor		Random Forest		
		Protein	Silicone Oil	Other	Protein	Silicone Oil	Protein	Silicone Oil	
FlowCam 8100 #1 IVlg, Normal Operation	3-5 μm	Protein	97.9%	2.0%	0.1%	89.8%	10.2%	90.8%	9.2%
		Silicone Oil	7.1%	92.7%	0.2%	38.0%	62.0%	35.0%	65.0%
	5+ μm	Protein	98.1%	1.2%	0.7%	96.2%	3.8%	99.2%	0.8%
		Silicone Oil	3.0%	92.6%	4.5%	37.1%	62.9%	53.9%	46.1%
	Total	Protein	98.0%	1.4%	0.6%	94.7%	5.3%	97.2%	2.8%
		Silicone Oil	6.0%	92.7%	1.4%	37.7%	62.3%	40.2%	59.8%
FlowCam 8100 #2 Ovalbumin, FlowCam LO flow rate	3-5 μm	Protein	97.5%	2.4%	0.1%	89.0%	11.0%	96.4%	3.6%
		Silicone Oil	8.5%	91.3%	0.3%	20.8%	79.2%	38.7%	61.3%
	5+ μm	Protein	99.0%	0.5%	0.5%	97.8%	2.2%	99.6%	0.4%
		Silicone Oil	1.2%	92.2%	6.6%	13.6%	86.4%	51.0%	49.0%
	Total	Protein	98.4%	1.2%	0.4%	93.9%	6.1%	98.2%	1.8%
		Silicone Oil	6.3%	91.5%	2.2%	18.7%	81.3%	42.3%	57.7%
FlowCam 8100 #3 Out of focus	3-5 μm	Protein	95.4%	4.6%	0.0%	55.0%	45.0%	85.0%	15.0%
		Silicone Oil	10.0%	89.3%	0.7%	16.8%	83.2%	24.5%	75.5%
	5+ μm	Protein	95.3%	1.7%	3.1%	55.3%	44.7%	94.7%	5.3%
		Silicone Oil	5.0%	94.4%	0.6%	13.3%	86.7%	68.3%	31.7%
	Total	Protein	95.3%	3.1%	1.6%	55.2%	44.8%	89.7%	10.3%
		Silicone Oil	8.3%	91.0%	0.6%	15.7%	84.3%	38.9%	61.1%

Figure 4 also shows the performance of the S-factor- and random forest-based silicone oil identification methods when trained on the same data used to develop VisualAI and tested on the data collected during this study. The accuracy of these other approaches was significantly worse on average than that obtained with VisualAI. These methods also do not offer as consistent a performance across FlowCam units as VisualAI, resulting in occasionally drastic changes not only in the overall accuracy but also in the relative frequency of protein aggregate predictions to silicone oil droplet predictions—i.e. they may overestimate the silicone oil content on data from one FlowCam but underestimate the content on another unit. It should be noted that, unlike the original studies, these methods were developed for and applied to FIM data captured on multiple instruments rather than a single instrument. While tailoring these approaches to individual FlowCam instruments would likely yield better classification performance, VisualAI is able to achieve high accuracy without any instrument-specific tuning.

The strong classification accuracy of VisualAI across FlowCam units and sample types makes it a powerful and widely-applicable tool for differentiating between protein aggregates and silicone oil droplets. VisualAI is also easy to integrate into a typical FlowCam-based biotherapeutic analysis as the software does not require the time, labor, and sample volume investment required to develop other techniques for this analysis—including other AI-based methods. The agnostic nature of VisualAI, coupled with its robustness, allows the software to work well across multiple instrument and sample configurations. This feature allows FlowCam users working with a variety of biotherapeutic drug products or with multiple FlowCam units to use a single, standard software tool for silicone oil monitoring across datasets and the comparability benefits associated with a unified monitoring strategy.

Additionally, VisualAI achieved >92% recognition of calibration beads as “other” particles on these units. The excellent recognition of beads can help operators automatically detect and flag images of these particles when they appear in datasets, improving the accuracy of not just the reported protein aggregate-silicone oil droplet concentrations but also the overall particle concentration in the sample. Similar performance can be expected with air bubbles, since to the untrained eye, air bubble images can resemble calibration beads.

VisualAI Composition Accuracy: Samples containing known mixtures of protein aggregates and silicone oil droplets were analyzed with VisualAI to assess how accurately the software could predict the composition of the sample. Figure 5 shows the VisualAI-predicted fraction of protein aggregates in each mixture as well as the pure protein aggregate and silicone oil stocks using the particles larger than 3 μm in each sample. While there is some slight disagreement between the actual and predicted compositions, the predicted protein aggregate content in each sample is strongly correlated with the actual particle content in each sample.

This correlation suggests that VisualAI is sensitive to the overall protein-silicone oil composition of mixed samples and, in conjunction with FlowCam, can be used to assess the fraction of particles in a sample that are proteinaceous. The slight disagreement between the reported and actual protein aggregate composition can be attributed in part to classification errors made by VisualAI. As there is a small fraction of misidentified particles in the pure protein aggregate and silicone oil droplet samples, the expected compositions on these mixtures are slightly closer to 50% than if the model was perfectly accurate—a trend observed in the reported compositions in Figure 5.

It is important to note that the particle content larger than 2 μm matched the desired composition, including particles 2-3 μm in diameter that were not analyzed via VisualAI to determine the composition. As with most biotherapeutic samples, a large fraction of the particle content detected via FlowCam 8100 with a 10X objective are within this excluded size range. Despite ignoring a large portion of the images in these artificial protein aggregate-silicone oil mixtures, VisualAI was still able to accurately predict the composition of protein aggregates and silicone oil droplets in these samples.

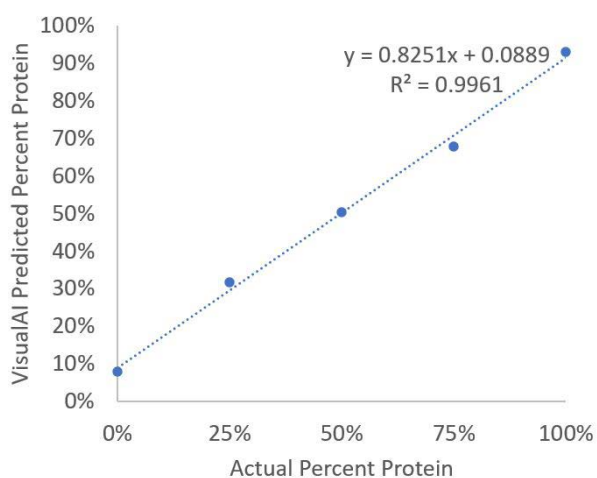


Figure 5. VisualAI-predicted protein aggregate compositions in samples containing known fractions of protein aggregate particles. Points represent average values obtained from triplicate samples. The trendline, equation, and R2 value were obtained using linear regression on the actual vs. predicted protein aggregate composition data.

Other Considerations: One inherent challenge with protein aggregate-silicone oil droplet detection strategies like VisualAI is the possibility of observing protein aggregates attached to silicone oil droplets and are simultaneously both particle types. These hybrid particles are common when protein is exposed to silicone oil for an extended period as protein adsorption and aggregation can occur on the buffer-oil interface⁸. This effect was controlled for in this study by preparing mixtures immediately before performing FlowCam analysis, minimizing the amount of time protein adsorption and aggregation that could occur. These hybrid particle types may appear more regularly in real-world samples.

Like other silicone oil detection strategies, VisualAI does not currently detect these hybrid particles and will always assign only one class (i.e. protein aggregate or silicone oil droplet) to each particle image. While a small number of these hybrids are not likely to significantly influence the overall composition estimated by VisualAI, users may need to post-process the classified images to identify and potentially reclassify these hybrid particles if they appear in significant quantities in a sample to get an accurate composition estimate.

To account for these hybrid particles and other types of non-protein aggregate, non-silicone oil droplet images, VisualAI provides flexibility to the user to define new classes (i.e. extra particle types) that images can be manually reassigned to following an initial classification. These features allow users to monitor the concentrations of other particle types that, at present, are not explicitly monitored by VisualAI. These tools can also be used to manually reassign particle images between protein aggregate, silicone oil, and other classes, allowing the user to further improve the accuracy of the particle type concentrations reported by the software.

CONCLUSIONS

VisualAI in combination with FlowCam 8100, and/or FlowCam LO, is a robust and powerful integrated flow-imaging solution for identifying protein aggregate and silicone oil droplet compositions in biotherapeutic samples. VisualAI achieves >90% accuracy both in identifying these particle types in individual images and predicting the overall composition of these particles in a sample. The off-the-shelf nature of the software results in easy integration of a protein aggregate and silicone oil monitoring approach into existing FlowCam-based biotherapeutic analysis workflows, including those involving multiple drug products and FlowCam units. VisualAI allows researchers to easily get automatic classification of particles in biotherapeutics from the FlowCam data, allowing them to make more effective decisions about the quality of their drug products and improve the overall quality of the therapies patients ultimately receive.



VisualAI™



Distributed by:
Kenelec Scientific Pty Ltd
1300 73 22 33
sales@kenelec.com.au
www.kenelec.com.au

VisualAI is an optional
module available with
VisualSpreadsheet 6

REFERENCES

1. Rosenberg AS. Effects of protein aggregates: An immunologic perspective. *AAPS J.* 2006;8(3):E501-E507. doi:10.1208/aapsj080359
2. Kotarek J, Stuart C, De Paoli SH, et al. Subvisible Particle Content, Formulation, and Dose of an Erythropoietin Peptide Mimetic Product Are Associated with Severe Adverse Postmarketing Events. *J Pharm Sci.* 2016;105(3):1023-1027. doi:10.1016/S0022-3549(15)00180-X
3. Carpenter JF, Randolph TW, Jiskoot W, et al. Overlooking Subvisible Particles in Therapeutic Protein Products: Gaps That May Compromise Product Quality. *J Pharm Sci.* 2009;98:1201-1205. doi:10.1002/jps
4. Narhi LO, Corvari V, Ripple DC, et al. Subvisible (2-100 µm) particle analysis during biotherapeutic drug product development: Part 1, considerations and strategy. *J Pharm Sci.* 2015;104(6):1899-1908. doi:10.1002/jps.24437
5. Roesch A, Zölls S, Stadler D, et al. Particles in Biopharmaceutical Formulations, Part 2: An Update on Analytical Techniques and Applications for Therapeutic Proteins, Viruses, Vaccines and Cells. *J Pharm Sci.* 2021;000. doi:10.1016/j.xphs.2021.12.011
6. Roberts CJ. Therapeutic protein aggregation: Mechanisms, design, and control. *Trends Biotechnol.* 2014;32(7):372-380. doi:10.1016/j.tibtech.2014.05.005
7. Bee JS, Randolph TW, Carpenter JF, Bishop SM, Dimitrova MN. Effects of surfaces and leachables on the stability of biopharmaceuticals. *J Pharm Sci.* 2011;100(10):4158-4170. doi:10.1002/jps.22597
8. Gerhardt A, McGraw NR, Schwartz DK, Bee JS, Carpenter JF, Randolph TW. Protein aggregation and particle formation in prefilled glass syringes. *J Pharm Sci.* 2014;103(6):1601-1612. doi:10.1002/jps.23973
9. Calderon CP, Daniels AL, Randolph TW. Deep Convolutional Neural Network Analysis of Flow Imaging Microscopy Data to Classify Subvisible Particles in Protein Formulations. *J Pharm Sci.* 2018;107(4):999-1008. doi:10.1016/j.xphs.2017.12.008
10. Chi EY. Protein Aggregation in Aqueous Solution -- Mechanism, Thermodynamics, and Kinetics. Published online 2004.
11. Kiese S, Pappenberger A, Friess W, Mahler H-C. Shaken, Not Stirred: Mechanical Stress Testing of an IgG1 Antibody. *J Pharm Sci.* 2008;97(10):4347-4366. doi:10.1002/jps
12. Bhatnagar BS, Bogner RH, Pikal MJ. Protein stability during freezing: Separation of stresses and mechanisms of protein stabilization. *Pharm Dev Technol.* 2007;12(5):505-523. doi:10.1080/10837450701481157
13. Chisholm CF, Baker AE, Soucie KR, Torres RM, Carpenter JF, Randolph TW. Silicone Oil Microdroplets Can Induce Antibody Responses Against Recombinant Murine Growth Hormone in Mice. *J Pharm Sci.* 2016;105(5):1623-1632. doi:10.1016/j.xphs.2016.02.019
14. Strehl R, Rombach-Riegraf V, Diez M, et al. Discrimination between silicone oil droplets and protein aggregates in biopharmaceuticals: A novel multiparametric image filter for sub-visible particles in microflow imaging analysis. *Pharm Res.* 2012;29(2):594-602. doi:10.1007/s11095-011-0590-7
15. Saggiu M, Patel AR, Koulis T. A Random Forest Approach for Counting Silicone Oil Droplets and Protein Particles in Antibody Formulations Using Flow Microscopy. *Pharm Res.* 2017;34(2):479-491. doi:10.1007/s11095-016-2079-x