https://doi.org/10.1038/s41592-023-01919-7

The Twenty Questions of bioimage object analysis

Beth A. Cimini & Kevin W. Eliceiri

Check for updates

The language used by microscopists who wish to find and measure objects in an image often differs in critical ways from that used by computer scientists who create tools to help them do this, making communication hard across disciplines. This work proposes a set of standardized questions that can guide analyses and shows how it can improve the future of bioimage analysis as a whole by making image analysis workflows and tools more FAIR (findable, accessible, interoperable and reusable).

Most, though not all, bioimage analysis tasks start with the idea of an object — a thing within an image a researcher would like to count, measure and/or track. Beyond this common overarching goal, there is an extraordinary diversity of bioimage analysis tasks that come from the variety of objects the biologist cares about and the variety of measurements they might want to perform. These include both measurements intrinsic to the object (How large is it? How much of marker X does it contain?) and those relating to either single collections of objects (How many objects do I have? Are all of my objects oriented in the same direction?) or multiple collections of objects (How many A's are in each B?).

In our experience, life scientists with limited computational training who begin to approach an image analysis task naturally proceed from thinking about the biology — that is, "I want to know whether 53BP1 is preferentially recruited to telomeres in my different perturbation conditions." Algorithm developers and computer vision experts, by contrast, are typically trying to produce tools that work along very different lines — to them, the previous task could be best summarized as, "I want to determine the intersection-over-union (IoU) of the approximately 20-pixel-diameter circular objects in channel 1 with the approximately 3-pixel-diameter circular objects in channel 2." Neither approach is wrong, but the discrepancies between them can lead to confusion and to scientists speaking past one another, leading to a situation where life scientists feel many of their image analysis problems remain unsolved[1] despite the efforts of the bioimage analysis community to develop and curate tools for them (see ref. 2 for examples and links to other tools).

Inspired by efforts in the microscopy community to systematize and create ontologies for bioimaging[3,4] and bioimage analysis[5], as well as cellular phenotypes[6], we believe that the future of bioimage analysis (and especially reusable image analysis — the R in FAIR[7,8]) can be both improved and accelerated by organizing imaging datasets and image analysis solutions into a schema that can actually direct the creation of image analysis workflows by asking biologists the right guiding questions about their data. We have drafted a preliminary schema that we call "the Twenty Questions of bioimage object analysis" (Fig. 1); all versions are available at ref. 9 and an interactive Bioimage Object Analysis Questions website. Individual image sets could be annotated with the answers to each of the organizing questions (with one set of responses for each object the user wishes to analyze); responding image analysis workflows would report their suitability or ability to handle a given parameter (or parameter combination). For now, the website serves primarily as a proof of concept demonstrating the questions, the answers and how the answers may change an image analysis workflow in one particular tool (CellProfiler[10]); as the schema becomes more robust in response to community feedback and as more datasets and/or workflows are annotated, one can envision a number of ways the tool could improve bioimage analysis, from the human (creating a shared, structured language base by which imaging scientists may interact with their image analyst colleagues) to the technological (automatic suggestion of tools or creation of bespoke workflows; see also below).

Each query in the proposed questions is directly linked to a computer vision task needed in the creation of a conventional bioimage analysis workflow: information about the space between objects may guide how object center selection is performed, and information about the presence of debris could be used to decide whether or not to perform a module to detect and mask especially-bright areas of an image before proceeding with detection of the desired objects. A skilled bioimage analyst might be able to use the answers to these questions to produce a reasonably performing image analysis workflow without ever seeing the dataset in question. Organizing image analysis datasets and workflows around any such schema has (at least) five major advantages:

1. **It breaks down artificial disciplinary barriers**. A biologist studying DNA damage in mice is unlikely to search the literature for workflows from biologists studying RNA processing in humans, but in practice, a workflow that does a good job of detecting mammalian nuclei and then speckles within them should work equally well for researchers in both fields.
2. **It allows the caching of properties**. As efforts to systematize microscopy image information and store such information directly into image metadata gain popularity, the answers to many of the questions proposed in our Twenty Questions can be saved and automatically pre-filled. While the matrix of biological objects and microscopy conditions is large, it is not infinite, and once a number of researchers say that object X has an approximate size of Y under Z imaging conditions (pulled automatically from the image metadata), that question no longer needs to be presented.
3. **It lowers the barrier to entry**. The bioimage analysis tool ecosystem contains a huge spectrum of tools[11], from tools that are optimized for a particular use case to tools that attempt to generalize across many use cases and may have hundreds or even

1. Are your images 3D?	2. Are you trying to analyze images from a timelapse series?	3. What is the microscope methodology?	4. Does alignment and/or registration need to be applied before proceeding?
Yes	Yes	Brightfield — unstained	No registration required
No		Brightfield — histology stained	Need to align images to other images
Yes, but I wish to *z* project them before proceeding	No	Fluorescence	Need to align images to a reference atlas
5. Is fluorescence bleedthrough present in these images?	**6. How would you describe the signal-to-noise and signal-to-background ratios between background and true signal?**	**7. Is there a marker for the desired object that smoothly covers the whole object?**	**8. Is there a marker for the desired object that specifically highlights its edges?**
NA; not fluorescence and/or multichannel	NA; not fluorescence	Yes	Yes
No	Poor (histogram peaks nearly undistinguishable)		
Yes, but we don't need to apply a correction in this workflow	OK (histogram peaks obvious but overlapping)	No	No
Yes, and we do need to apply a correction in this workflow	Great (histogram peaks fully distinct)		
9. Which, if any, of the following image quality issues are present? Select all that apply.	**10. What object constraints, if any, should be used to filter objects for validity? Select all that apply.**	**11. What shape is most characteristic of your objects?**	**12. How much, if any, object overlap is expected?**
Overconfluence	Minimum or maximum diameter	Round	None
Debris	Minimum or maximum count	Square	Typically only at the object boundaries
Inconsistent focus	Presence within another stain or object	Linear or extended	
Variable background	Minimum or maximum count within another object	Irregular	Overlap can happen anywhere within the object
Artifacts from tile stitching	Minimum or maximum possible threshold		
Artifacts from multi-view integration			
13. What is the median object diameter?	**14. Approximately how large is the size range between your biggest and smallest expected objects?**	**15. How much consistency in shape do your objects have?**	**16. How close will each object typically be to its nearest neighbors of the same object type?**
1–3 px	Range is ± 30%	Low: shapes can vary widely	No neighbors within 3 diameters
3–10 px	Range is ± 30% to 2-fold	Medium: shapes are typically similar but not identical	Some neighbors within 1–3 diameters
10–30 px	Range is 2-fold to 10-fold	High: the shapes are nearly always identical	Some touching of neighbors
30–100 px	Range is greater than 10-fold		Entirely touching neighbors
100–300 px			
300–1,000 px			
1,000–3,000 px			
17. Should objects touching the border of the image be removed?	**18. Approximately how many objects of this type will be present in each image?**	**19. Which object-intrinsic measurements should be made? Select all that apply.**	**20. Measurements requiring other objects and/or image context (select all that apply).**
Yes	0–3	Shape and/or size	Count of parent or child objects of a different type
	3–10	Intensity values	
	10–100	Intensity spatial distribution	Distance to or overlap with other objects of the same type
No	100–1,000	Texture or granularity	
	1,000–10,000	Colocalization of channels	Distance to or overlap with other objects of a different type
		Motion parameters	

Fig. 1 | The proposed questions and answers for the Twenty Questions of bioimage object analysis. Questions 1–9 (blue) affect decisions made in preprocessing, questions 10–18 (orange) affect decisions made in classical segmentation, and the final two questions (green) relate to acquiring measurements of the segmented objects. If researchers wish to find more than one kind of object in an experiment, they should answers the questions independently for each object, though some questions (such as those around microscopy methodology and quality issues) are likely to stay the same for multiple objects. NA, not applicable.

tens of thousands of subcomponents, such as the modules of CellProfiler[10] or the plug-in ecosystem of ImageJ[12] and Fiji[13]. While we find this diversity outstanding, it is overwhelming, even for an expert but especially for a beginner. Organization of tools or tool components according to what they do well and what they do not attempt to do at all gives users a place to start and helps them identify appropriate tools to combine into their overall workflow.

4. **It improves scorability**. While a few excellent resources (such as refs. 14–16) attempt to score the ability of various image analysis workflows to properly quantify hand-curated datasets, 'failure to perform well' can happen for dozens of reasons, and success in one case does not guarantee success in all. Adopting these questions for 'challenge' datasets would make it easier for workflow users to see whether success on a given challenge dataset is likely to translate to success on their particular data. Categorization according to a standard question schema would also allow tool developers insight into where their tools do (and do not) perform well and allow developers and hosts of analysis challenges to more easily identify what sorts of datasets do (and do not) currently have good hand-curated test sets.
5. **It prepares for an AI-enabled future**. By linking datasets to particular prespecified metrics, we in practice create a training set that tools could use to automatically answer some of our Twenty Questions in the future. The ability to automatically detect debris, for example, would certainly be of interest for future AI-enabled microscopes, which upon detecting debris could move to a different field of view or reject an entire plate altogether if too much debris is detected. Further, by associating bioimage analysis tools and/or operations to particular human-language questions, we take a major step toward automatic workflow creation, as well as toward enabling natural language processing of bioimages. To return to our initial example, "I want to know whether 53BP1 is preferentially recruited to telomeres in my different perturbation conditions" becomes a potentially answerable question to a machine that understands the typical ways to detect these structures and can call up the appropriate workflows from memory.

We are certain that, immediately upon reviewing our Twenty Questions, researchers will find a use case that they do not cover; we believe that with community discussion and improvement of the proposed schema, a better-optimized version that will allow improved handling of the majority of data and improved targeting of the majority of tools can be created. We look forward to feedback on this draft proposal on the Scientific Community Image Forum[17], as well as integrating with other efforts to establish standards for metadata[3,4], image analysis reporting[5,18], phenotypes[6] and beyond. In creating an easier-to-use and easier-to-assess bioimage analysis tool ecosystem, we can make bioimaging FAIRer for everyone.

Beth A. Cimini [1] ✉ **& Kevin W. Eliceiri** [2]

[1]Imaging Platform, Broad Institute of MIT and Harvard, Cambridge, MA, USA. [2]Center for Quantitative Cell Imaging, University of Wisconsin–Madison and Morgridge Institute for Research, Madison, WI, USA.
✉e-mail: bcimini@broadinstitute.org

Published online: 11 July 2023

References

1. Jamali, N., Dobson, E. T. A., Eliceiri, K. W., Carpenter, A. E. & Cimini, B. A. *Biol. Imaging* **1**, e4 (2022).
2. Paul-Gilloteaux, P. et al. *F1000 Res.* **10**, 320 (2021).
3. Sarkans, U. et al. *Nat. Methods* **18**, 1418–1422 (2021).
4. Hammer, M. et al. *Nat. Methods* **18**, 1427–1440 (2021).
5. Kalaš, M. et al. *F1000 Res.* https://doi.org/10.7490/f1000research.1117826.1 (2020).
6. Jupp, S. et al. *J. Biomed. Semantics* **7**, 28 (2016).
7. Wilkinson, M. D. et al. *Sci. Data* **3**, 160018 (2016).
8. Barker, M. et al. *Sci. Data* **9**, 622 (2022).
9. Cimini, B. *Zenodo* https://doi.org/10.5281/zenodo.7654937 (2023).
10. Stirling, D. R. et al. *BMC Bioinformatics* **22**, 433 (2021).
11. Haase, R. et al. *FEBS Lett.* **596**, 2472–2485 (2022).
12. Schindelin, J., Rueden, C. T., Hiner, M. C. & Eliceiri, K. W. *Mol. Reprod. Dev.* **82**, 518–529 (2015).
13. Schindelin, J. et al. *Nat. Methods* **9**, 676–682 (2012).
14. Ulman, V. et al. *Nat. Methods* **14**, 1141–1152 (2017).
15. Biaflows. https://biaflows-sandbox.neubias.org/#/ (accessed 19 February 2023).
16. Grand Challenge. https://grand-challenge.org/ (accessed 19 February 2023).
17. Rueden, C. T. et al. *PLoS Biol.* **17**, e3000340 (2019).
18. Schmied, C. et al. Community-developed checklists for publishing images and image analysis. Preprint at https://doi.org/10.48550/arXiv.2302.07005 (2023).

Acknowledgements

The authors thank members of the Cimini lab, A. Carpenter and U. Manor for their feedback on the completeness of the schema. This work was supported by the US National Institutes of Health (NIH P41 GM135019 to B.A.C. and K.W.E.). B.A.C. (2020-225720) and K.W.E. also acknowledge funding from the Chan Zuckerberg Initiative DAF, an advised fund of the Silicon Valley Community Foundation.

Author contributions

B.A.C. drafted the initial schema, coded the web tool and wrote the manuscript; K.W.E. revised the schema and the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Peer review information *Nature Methods* thanks Robert Haase and the other, anonymous, reviewer(s) for their contribution to the peer review of this work.