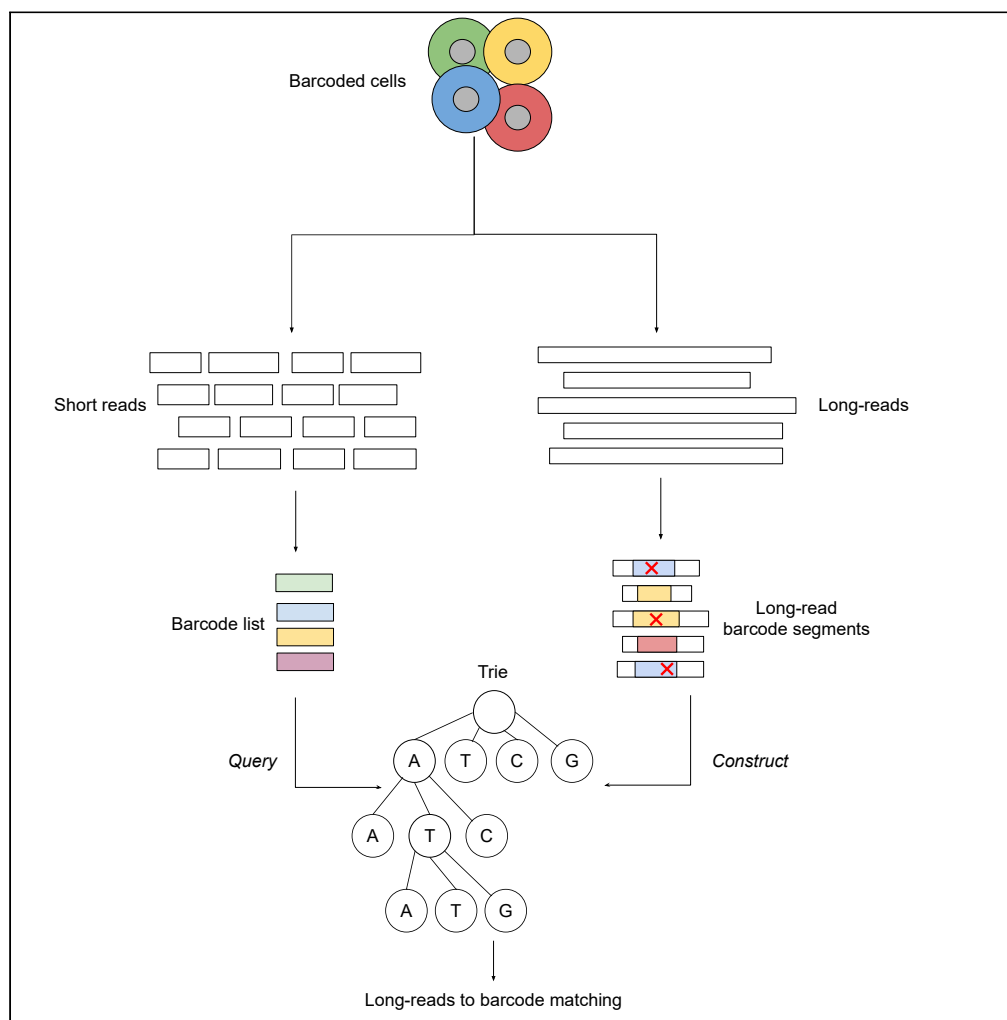


Article

Fast and accurate matching of cellular barcodes across short-reads and long-reads of single-cell RNA-seq experiments



Ghazal Ebrahimi,
Baraa Orabi,
Meghan
Robinson, Cedric
Chauve, Ryan
Flannigan, Faraz
Hach

faraz.hach@ubc.ca

Highlights

We introduce scTagger, a method for matching long-read and short-read cellular barcodes

The matching enables combining gene expression analysis and isoform detection

scTagger relies on a trie-based data structure to perform this matching

We show that scTagger is accurate and fast on both real and simulated data

Ebrahimi et al., iScience 25,
104530
July 15, 2022 © 2022 The
Authors.
[https://doi.org/10.1016/
j.isci.2022.104530](https://doi.org/10.1016/j.isci.2022.104530)



Article

Fast and accurate matching of cellular barcodes across short-reads and long-reads of single-cell RNA-seq experiments

Ghazal Ebrahimi,^{1,6} Baraa Orabi,^{2,6} Meghan Robinson,³ Cedric Chauve,⁵ Ryan Flannigan,^{3,4} and Faraz Hach^{2,3,4,7,*}

SUMMARY

Single-cell RNA sequencing allows for characterizing the gene expression landscape at the cell type level. However, because of its use of short-reads, it is severely limited at detecting full-length features of transcripts such as alternative splicing. New library preparation techniques attempt to extend single-cell sequencing by utilizing both long-reads and short-reads. These techniques split the library material, after it is tagged with cellular barcodes, into two pools: one for short-read sequencing and one for long-read sequencing. However, the challenge of utilizing these techniques is that they require matching the cellular barcodes sequenced by the erroneous long-reads to the cellular barcodes detected by the short-reads. To overcome this challenge, we introduce scTagger, a computational method to match cellular barcodes data from long-reads and short-reads. We tested scTagger against another state-of-the-art tool on both real and simulated datasets, and we demonstrate that scTagger has both significantly better accuracy and time efficiency.

INTRODUCTION

Single-cell RNA sequencing (scRNA-seq) has advanced biological research to move beyond the macro view provided by bulk RNA sequencing. Technical advances have enabled gene expression analysis at the single-cell level, which has resulted in a greater understanding of cellular diversity in gene expression (Rozenblatt-Rosen et al., 2017). scRNA-seq achieves this cell-level resolution by deploying short artificial DNA fragments known as cellular barcodes to tag the RNA fragments in each cell with a unique sequence of nucleotides. Once tagged, complementary DNA strands are created for each RNA strand and are subsequently fragmented, amplified, and loaded onto the sequencing machine to be sequenced in bulk. Following this, the sequencing outputs are demultiplexed and reads are assigned to different cells *in silico* using the sequenced cellular barcodes. The genes and their RNA expression levels that are attributed to each cell allow the user to classify cells into clusters of cell types or populations that are biologically relevant (Eberwine et al., 2014; Li et al., 2020). Observing the gene expression at the scale of thousands of cells from a single sample is made readily accessible using droplet-based technologies such as the 10x Genomics Chromium approach (Zheng et al., 2017).

Although gene expression is valuable to reveal biological processes, it tells only a part of the transcription story. Genes are composed of various isoforms that differ in the order and composition of their respective exons, which have significant biological implications in subsequent protein translation and function. Unfortunately, evaluating these splicing variations in RNA isoforms is limited and often impossible when using standard scRNA-seq techniques, such as the 10x Genomics Chromium protocol, because they rely on short-read (SR) sequencing platforms, often Illumina where 150 total base pairs are typically sequenced at a time. The reason SR sequencing is utilized by scRNA-seq is due to its low error rate (~0.1% (Goodwin et al., 2016)) which allows for easy demultiplexing of cellular barcodes. In addition, SR platform's relative affordability allows for high coverage sequencing that enables the estimation of gene expression. However, because of this reliance on SR sequencing, scRNA-seq fails to capture the full picture of transcripts that are much longer than what the current SR platforms can generate.

Long-read (LR) sequencing technologies, such as Oxford Nanopore Technologies (ONT), have demonstrated their capacity to identify full-length isoforms, albeit using bulk RNA that is not multiplexed at the

¹Bioinformatics Program, the University of British Columbia, Vancouver, BC, Canada

²Computer Science Department, the University of British Columbia, Vancouver, BC, Canada

³Vancouver Prostate Centre, Vancouver, BC, Canada

⁴Department of Urologic Sciences, the University of British Columbia, Vancouver, BC, Canada

⁵Department of Mathematics, Simon Fraser University, Burnaby, BC, Canada

⁶These authors contributed equally

⁷Lead contact

*Correspondence: faraz.hach@ubc.ca

<https://doi.org/10.1016/j.isci.2022.104530>



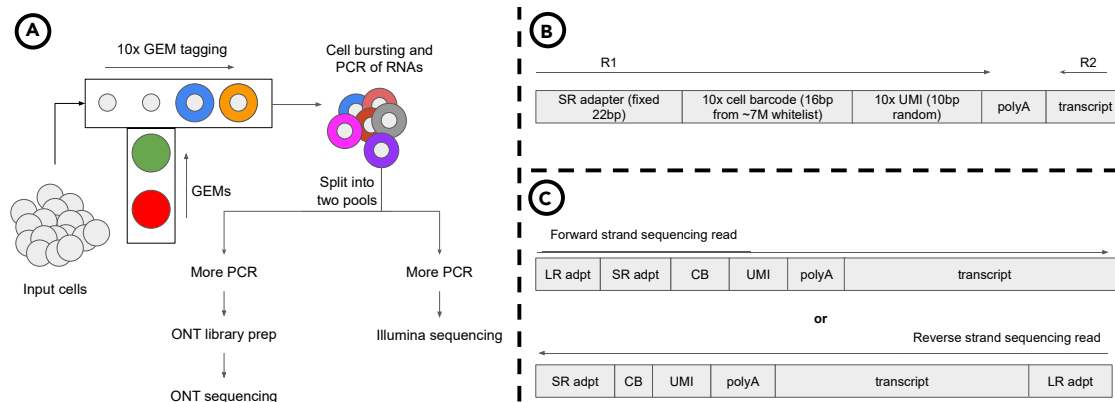


Figure 1. Library preparation and sequencing templates for short- and long-reads

(A) Overview of the library preparation: Microfluidic chips are used to generate 10x Chromium GEMs, which tag the RNA transcripts with cellular barcodes. After the transcripts are tagged, the GEMs are burst, and the tagged RNA material is split into two pools of sequencing, one for SRs and one for LR. (B) The SR template. Inside the GEMs, RNA transcripts are tagged with an Illumina adapter that is a fixed sequence, followed by a 16bp cellular barcode (CB), followed by a random 10bp sequence for the unique molecular identifier (UMI). (C) The LR template. Note that the LR template is essentially the same as the SR template with the LR sequencing adapter added. Depending on the specifics of the library preparation, LR may sequence the forward or the reverse strand of the RNA molecule. In either case, we expect the cellular barcode to be adjacent to the SR adapter sequence.

cellular level (Kovaka et al., 2019; Tang et al., 2020). Recently, new and promising library preparation techniques aim to take advantage of both LR RNA sequencing and SR single-cell RNA sequencing (Singh et al., 2019; Tian et al., 2021; Gupta et al., 2018). These techniques typically modify existing SR scRNA-seq protocols by splitting the RNA material, after it is tagged with the cellular barcodes, into two pools: i) one for Illumina SR sequencing and ii) one for LR sequencing. The high-level overview of this library preparation technique is illustrated in Figure 1.

These techniques are designed to enable the clustering of cells into types, using cell-specific gene expression, while taking advantage of the LR data to identify alternatively spliced isoforms and map them to cells (and subsequently into cell types) using the shared cellular barcodes. However, this mapping process is computationally not trivial because of the level of noise exhibited by LR sequencing platforms (10% – 20% (Kono and Arakawa, 2019)), which frequently introduces errors to the cellular barcodes sequences. Current methods to associate LR to cellular barcodes include SiCeLoRe (Kevin et al., 2020) and FLAMES (Tian et al., 2021): SiCeLoRe relies on the exact matching of permuted mutations of the LR barcode sequences, whereas FLAMES relies on the alignment of the barcodes to the LR using dynamic programming. Such approaches are theoretically computationally prohibitive for a large number of cells and reads.

Therefore, we present scTagger, a novel computational method for efficiently and accurately identifying and matching single-cell barcodes in split hybrid LR-SR scRNA sequencing experiments that can scale for a large number of reads and cells. scTagger uses a trie-based data structure to efficiently match the identified barcodes in the SRs to the LR while allowing for nonzero edit distance matching. Using both real and simulated datasets, we show that scTagger has accuracy on par with an exact but computationally intensive dynamic programming-based matching approach while being orders of magnitude faster. Furthermore, we show that scTagger has much higher accuracy and is much more time-efficient than other available tools.

RESULTS

scTagger overview

The hybrid SR-LR protocol that we refer to in the Introduction starts with RNA from single cells of a given sample. The RNA molecules of each cell are barcoded with a fixed-length barcode (Cellular barcode) sampled from a pool of over six million (10x Chromium v3 Chemistry) possible whitelisted 16bp barcodes. This RNA material is split into two pools, one pool for SR Illumina sequencing, which is used to produce a gene expression matrix, and the other pool for LR ONT sequencing, which is used to identify isoforms. The expected template sequence for the SRs and the LR is illustrated in Figures 1B and 1C.




Figure 2. Overview of scTagger's pipeline.

The three stages of scTagger are shaded, whereas the data inputs are colored in red

scTagger aims to match each LR to one of the cellular barcodes that tagged the input cells. To achieve this goal, scTagger proceeds in three stages:

1. Preprocessing the SRs and identifying the barcode of each SR. In this stage, scTagger selects a *small* subset (proportional to the number of input cells) of barcodes that cover most of the SRs using the frequencies of the barcodes appearing in the SRs. The small size of this subset allows for efficient matching without significant loss of data.
2. Locating a short segment on each LR where the barcode is expected to be present. In this stage, scTagger exploits the *a priori* knowledge about the template of the LRs (Figure 1C) and uses the alignment of the fixed Illumina adapter sequence to each of the LRs to identify the segment where the barcode is present. In Figure 7, we show the alignment loci of the Illumina adapter on the LRs on a real dataset we generate. scTagger uses these observed adapter locations to extract the barcode segments.
3. Matching the SR barcodes to the LR segments. scTagger achieves this matching by using a trie data structure (Fredkin, 1960). The trie is modified to allow matching each LR segment to the barcode that requires the least number of edits (e.g., substitutions, deletions, or insertions) to form an exact match with the segment. The maximum number of allowed edits for a barcode-segment match is set by the user and is limited by default to two errors.

Figure 2 presents an schematic overview of these stages and the STAR Methods section provides details of them.

Experimental setup

We assessed the tested methods, FLAMES (GitHub commit 18fb83c) and scTagger, on both real and simulated datasets. Note that we attempted to run SiCelore but failed to run it successfully because of runtime errors, which is concordant with the reporting from FLAMES' paper on SiCelore (Tian et al., 2021). In addition to FLAMES and scTagger, we also ran a baseline brute-force method of matching the cellular barcodes in the LRs by performing a *glocal* dynamic programming based alignment of all the selected barcodes (stage 1) to all the extracted LR segments (stage 2). We will refer to this baseline method as the Brute-force matching method in this manuscript. Note that all the matches found by scTagger exactly mirror those found by the Brute-force method as long as they have a maximum edit distance of $e = 2$, the maximum allowed distance set in scTagger. In addition, according to the FLAMES' manual, we are supposed to run FLAMES with the full 10x Genomics barcode whitelist, which includes over six million barcodes. However, doing so resulted in FLAMES not terminating even after running for 24 h. Therefore, we modified the way we run FLAMES by feeding it the selected subset of barcodes found by scTagger (stage 1). Thus, the results we present here for FLAMES depend on some of the output of our own method.

Dataset generation

We generated three real and three simulated datasets. The details of the generating procedures for both follows below, and the summary statistics for each dataset are detailed in Table 1.

Real datasets. We isolated cells from three different male infertility patients. The three samples were extracted from the testis tissue using the protocol developed by (Valli et al., 2014). The first sample was

Table 1. Summary of used datasets

Dataset		# of SRs	# of LRs	# of cells
Real	N	172,826,613	7,371,195	7,368
	NOA1	148,095,537	5,665,245	4,494
	NOA2	111,004,036	4,427,993	3,513
Simulated	Small	1,000,379	98,873	50
	Medium	15,003,985	998,471	500
	Large	150,005,933	4,993,805	5,000

The number of cells for real datasets is estimated by 10x Genomics Cell Ranger v3.0.2.

extracted from a normal human testis biopsy, whereas the second and third samples were extracted from human testis biopsies diagnosed with nonobstructive azoospermia, a severe abnormality in spermatogenesis resulting in failure to produce sperm. We name these three samples N, NOA1, and NOA2, respectively. Libraries were prepared in accordance with the protocol developed by Gupta et al. (2018) for 10x Genomics Chromium Single Cell 3' v3.0 and Oxford Nanopore Technologies PromethION platform. For the cDNA amplification step, we applied 17 cycles of PCR. Then, the amplified and barcode tagged cDNA material was split into two pools. The first pool was sequenced on the Illumina HiSeq platform with a read length of 2×150 bp. The second pool was the input for the ONT cDNA sequencing protocol, for which we used the SQK-PBK004ONT ONT kit. For this pool, end repair, A-tailing, and ligation steps were performed, and 11 cycles of PCR were applied using a starting cDNA amount of 50ng.

Simulated datasets. In addition to the real datasets, it is important to assess the accuracy of the different tested methods on simulated datasets for which we can obtain the ground truth of barcode matching between the LRs and SRs. Therefore, we built a simulation pipeline to generate SR and LR data with known cellular barcode matching. The structure of the simulation pipeline is outlined in Figure 3. Note that we use both Minnow (Sarkar et al., 2019) and Badread (Wick, 2019) simulators as components of our simulation pipeline. In total, we generated three simulated datasets. As part of the input for simulating all three datasets, we used the transcriptome mapping of the SRs of the real sample N as a basis to simulate the differential expression of gene isoforms according to Minnow's instructions.

Barcode selection and SR coverage

The first stage of scTager filters and selects SR barcodes. Here, we need to examine how well the selected barcodes cover the SR data and how little value adding more barcodes brings. A significant amount of the SR coverage is wasted on barcodes that did not tag any input cells in the library preparation and are therefore not useful for downstream analysis.

In Figure 4, we plot the proportion of SRs that are covered by retaining each additional 1,000 barcodes in order of the most-to-least frequent. As we observe in the figure, the contribution of adding each additional 1,000 barcodes diminishes quickly. Therefore, we observe that our barcode selection method (with a threshold of 0.5%) achieve its conservative goal while producing a relatively small set of error-corrected barcodes. The exact number of selected SR barcodes and the proportion of SRs they cover are detailed




Figure 3. Overview of simulation pipeline: The SRs mapping to the transcriptome reference is used by Minnow to model differential isoform expression within each gene

The number of genes, SRs, and cells to be simulated are passed to Minnow. We then reconstruct the transcript molecule (including SR adapter, barcode, and polyA tail) that each simulated SR was generated from. These reconstructed barcoded transcripts are then passed to the LR simulation module, which wraps around the Badread simulator to generate the LRs. Note that the full simulation pipeline is available on our GitHub repository.




Figure 4. Cumulative SR coverage with batches of 1,000 barcodes for the real datasets

in Table 2. Because of this barcode selection process, we lose between 9 and 17% of the SR throughput. These selected barcodes are the ones used for matching in all the tested methods. Note that in all three cases, the number of selected barcodes exceeds the expected number of cells in each sample (Table 1).

LR coverage and segment extraction

The second stage of scTagger extracts LRs barcode segments. scTagger uses the alignment of the Illumina sequencing adapter to the LRs to identify the segment containing the cellular barcode in each LR. In Figure 5, we plot the histogram of the LRs by the edit distance of the successful adapter alignments. We observe that we are able to find successful adapter alignments in over 75% of LRs across all real datasets. For the simulated datasets, we observe that scTagger is able to find successful adapter alignments on ~ 92 – 93% of the LRs.

Accuracy of LR barcode matching

The final and main stage of scTagger is matching the selected SR barcodes with the extracted LR segments. Here, we use the simulated datasets to assess the accuracy of the different methods. For every tested method, we label each simulated LR as: (i) a match if the method outputs only the one barcode we simulated for the LR; (ii) an ambiguous match if the method outputs multiple barcodes and one of them is the simulated barcode; (iii) a mismatch if none of the output barcodes is the simulated barcode; and finally (iv) a skip if the method does not output any barcodes for the LR (because of the edit distance limit or because of not finding any feasible adapter alignment in the second stage of scTagger). Note that although ambiguous matches include the correct barcode, they are not readily useful for any downstream analysis. Table 3 details the accuracy statistics for all the tested methods.

As we observe in Table 3, scTagger significantly outperforms FLAMES in its accuracy. Besides, despite the low edit distance limit we use in scTagger ($e = 2$), scTagger has a very similar match rate to the Brute-force baseline method. As a matter of fact, we observe that allowing for higher edit distance increases the mismatch and ambiguous match rates.

Table 2. The number of unique barcodes in each real dataset and the percentage of SRs that they cover

Dataset	All barcodes		Selected barcodes		
	# of unique barcodes	% of SRs covered	# of unique barcodes	% of SRs covered	% of SRs lost due to barcode selection
N	1,723,938	97%	10,000	80%	17%
NOA1	1,063,678	72%	14,000	55%	17%
NOA2	859,034	94%	7,000	85%	9%

Here, we compare the set of all barcodes we observe in the SRs vs the set of selected barcodes by scTagger.

Note that some SRs do not have any observed barcodes.

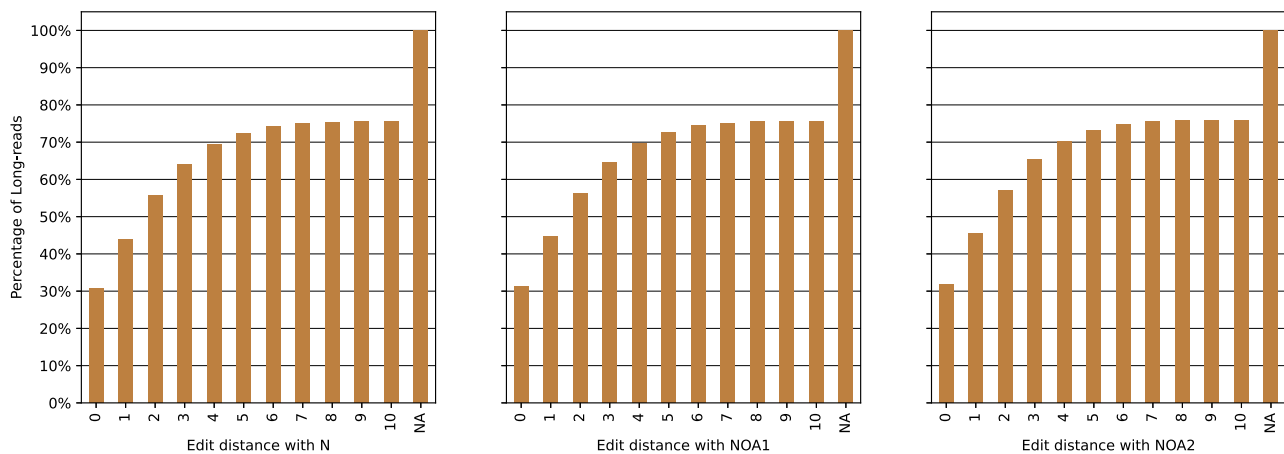


Figure 5. Cumulative percentage of LR distance with a successful SR adapter alignment
NA refers to LR with no alignment ≤ 10 edits.

Although we cannot establish a ground truth to calculate the accuracy of scTagger in matching LR of the real datasets, we can still compare scTagger matching to the baseline Brute-force method. Note that the differentiating parameter between these two methods is the choice of the edit distance limit in scTagger; if we set that parameter to the barcode size, the two methods produce identical results. Figure 6 depicts the distribution of the real datasets LR by whether they are: (i) skipped, (ii) matched with edit distance ≤ 2 , (iii) matched with edit distance > 2 , and (iv) matched by a single unique barcode. As we observe in these UpSet (Alexander et al., 2014) plots, only 0.8% to 2.2% of all the LR have unambiguous matches and yet are lost by scTagger because of having an edit distance limit of two edits. Thus, the loss of matches because of setting the error limit to two is very small in practice. The remaining ambiguous matches, i.e., those with more than one barcode match, are not readily usable by downstream analyses. In addition, when comparing FLAMES and scTagger matches with the Brute-force method, we observed that the percentage of the LR that are matched by FLAMES and the Brute-force methods but not scTagger is negligible ($\sim 0.1\%$). In contrast, the percentage of the LR that are matched by scTagger and the Brute-force methods but not FLAMES is very significant ($\sim 26\%$ to $\sim 28\%$). Table S1 of the Supplemental information presents the statistics of comparing the three methods' matches on the real datasets.

Runtime and memory

Table 4 shows the memory usage and runtime of the different stages of scTagger as well as FLAMES and the Brute-force method. As we observe in Table 4, scTagger matching stage has the smallest runtime footprint compared to the other methods. Unlike FLAMES, scTagger can take advantage of multiple CPU cores and completes computing in orders of magnitude less time than FLAMES or the Brute-force method. The only computational resource that scTagger lags behind the other tested tools is memory usage. However, its memory usage for these real datasets is well within the acceptable limits of modern personal computers.

Table 3. Accuracy on simulated datasets

Dataset	Method	Matches	Ambiguous matches	Mismatches	Skipped
Small	scTagger	47.68%	9.98%	4.57%	37.76%
	Brute-force	53.22%	23.82%	14.77%	8.19%
	Flames	27.40%	0.00%	15.92%	56.68%
Medium	scTagger	51.60%	7.10%	4.54%	36.76%
	Brute-force	55.55%	20.03%	17.37%	7.05%
	Flames	28.37%	0.00%	22.10%	49.52%
Large	scTagger	53.26%	6.24%	5.87%	34.62%
	Brute-force	54.91%	18.42%	19.62%	7.04%
	Flames	30.05%	0.00%	20.34%	49.61%

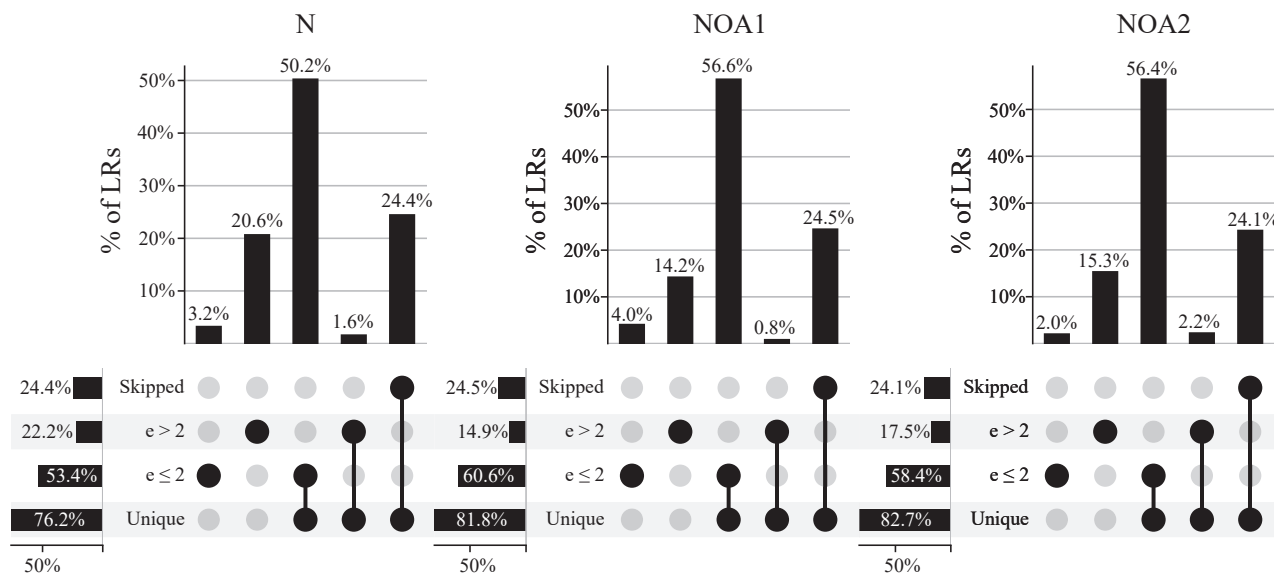


Figure 6. Distribution of matches in the real datasets according to the Brute-force matching method

DISCUSSION

We presented a novel computational method—scTagger—to match SR data to LR data using their cellular barcodes. We show that scTagger achieves this goal more efficiently and accurately than the other alternative methods. In the future, we plan to extend the scTagger method to handle matching of the unique molecular identifier sequences present in the SR and LR data. We also plan to improve on scTagger’s memory use by streamlining how it reads and processes the input.

Limitations of the study

The main limitation of the current form of scTagger is that it is specific for 10x Genomics scRNA-seq. Note that it is possible to readily extend support for other scRNA-seq platforms if their resulting LR template includes some fixed sequence next to the barcode. In addition, scTagger takes advantage of the fact that the number of cells (and thus the number of relevant barcodes) is on the manageable order of magnitudes of thousands to tens of thousands. Therefore, we expect scTagger in its current design not to scale well if the number of relevant barcodes becomes much bigger than that.




Figure 7. Optimal alignment locations of the SR adapter to the LRs

Negative locations are on the reverse strand. Automatically selected ranges around the peaks for sample N are indicated on the plot.

Table 4. Time and memory usage on the real datasets as reported by GNU Time

Dataset	Method	Real time (min)	User time (min)	Max memory (GB)	Threads
N	scTagger (LR segments)	4	14	7.3	32
	scTagger (matching)	15	266	4.9	32
	scTagger (matching)	185	186	3.6	1
	Brute-force matching	232	7,234	2.2	32
	FLAMES	709	702	0.1	1
NOA1	scTagger (LR segments)	3	11	6.6	32
	scTagger (matching)	9	140	3.4	32
	scTagger (matching)	192	193	2.9	1
	Brute-force matching	96	2,975	1.6	32
	FLAMES	308	304	0.1	1
NOA2	scTagger (LR segments)	5	16	8.4	32
	scTagger (matching)	17	280	4.6	32
	scTagger (matching)	129	128	2.3	1
	Brute-force matching	214	6,600	2.9	32
	FLAMES	603	600	0.1	1

STAR★METHODS

Detailed methods are provided in the online version of this paper and include the following:

- **KEY RESOURCES TABLE**
- **RESOURCE AVAILABILITY**
 - Lead contact
 - Materials availability
 - Data and code availability
- **METHOD DETAILS**
 - SR barcode selection stage
 - Extracting LR barcode segments stage
 - Matching LRs and selected barcodes stage
 - Trie construction
 - Querying the trie
 - Parallelizing trie construction and querying

SUPPLEMENTAL INFORMATION

Supplemental information can be found online at <https://doi.org/10.1016/j.isci.2022.104530>.

ACKNOWLEDGMENTS

This research is funded in part by the National Science and Engineering Council of Canada (NSERC) CREATE Training Program in High-Dimensional Bioinformatics Scholarship to G.E., NSERC Alexander Graham Bell Canada Graduate Scholarship-Doctoral (CGS D) to B.O., NSERC Discovery Grants to F.H. (RGPIN-05952) and C.C. (RGPIN-03986), the Michael Smith Foundation for Health Research (MSFHR) Scholar Award to F.H. (SCH-2020-0370), and Canadian Institutes of Health Research (CIHR) Early Career Investigator Grant (161357) and Vancouver Coastal Health Research Institute (VCHRI) Mentored Clinician Scientist Award to R.F. We would like to thank Reza Soltani for his thoughtful input into the development and calculating the time complexity of the trie data structure in scTagger and Fatih Karaoğlanoğlu for his feedback on designing the simulation pipeline.

AUTHOR CONTRIBUTIONS

Conceptualization, G.E., B.O., R.F., and F.H.; Data Curation: G.E., B.O., M.R., R.F., and F.H.; Formal Analysis, G.E. and B.O.; Funding Acquisition: G.E., B.O., C.C., R.F., and F.H.; Investigation, G.E., B.O., and F.H.; Methodology, G.E., B.O., and F.H.; Resources, R.F. and F.H.; Software, G.E. and B.O.; Supervision, C.C.,

R.F., and F.H.; Validation, G.E. and B.O.; Visualization, G.E. and B.O.; Writing – Original Draft: G.E., B.O., and F.H.; Writing – Review & Editing, G.E., B.O., M.R., C.C., R.F., and F.H.

DECLARATION OF INTERESTS

The authors declare no competing interests.

INCLUSION AND DIVERSITY

We worked to ensure ethnic or other types of diversity in the recruitment of human subjects. One or more of the authors of this paper self-identifies as an underrepresented ethnic minority in science. One or more of the authors of this paper self-identifies as a member of the LGBTQ+ community.

Received: February 11, 2022

Revised: May 10, 2022

Accepted: May 31, 2022

Published: July 15, 2022

REFERENCES

- Fredkin, E. (1960). Trie memory. *Commun. ACM* 3, 490–499. <https://doi.org/10.1145/367390.367400>.
- Goodwin, S., McPherson, J.D., and McCombie, W.R. (2016). Coming of age: ten years of next-generation sequencing technologies. *Nat. Rev. Genet.* 17, 333–351. <https://doi.org/10.1038/nrg.2016.49>.
- Gupta, I., Collier, P.G., Haase, B., Mahfouz, A., Joglekar, A., Floyd, T., Koopmans, F., Barres, B., Smit, A.B., Sloan, S.A., et al. (2018). Single-cell isoform RNA sequencing characterizes isoforms in thousands of cerebellar cells. *Nat. Biotechnol.* 36, 1197–1202. <https://doi.org/10.1038/nbt.4259>.
- Harris, C.R., Millman, K.J., van der Walt, S.J., Gommers, R., Virtanen, P., Cournapeau, D., Wieser, E., Taylor, J., Berg, S., Smith, N.J., et al. (2020). Array programming with NumPy. *Nature* 585, 357–362. <https://doi.org/10.1038/s41586-020-2649-2>.
- Alexander, L., Gehlenborg, N., Strobelt, H., Vuillemot, R., and Pfister, H. (2014). Upset: Visualization of Intersecting Sets. *IEEE Trans. Visual. Comput. Graph.* 20, 1983–1992. <https://doi.org/10.1109/tvcg.2014.2346248>.
- Eberwine, J., Sul, J.Y., Bartfai, T., and Kim, J. (2014). The promise of single-cell sequencing. *Nat. Methods* 11, 25–27. <https://doi.org/10.1038/nmeth.2769>.
- Kono, N., and Arakawa, K. (2019). Nanopore sequencing: Review of potential applications in functional genomics. *Dev. Growth Differ.* 61, 316–326. <https://doi.org/10.1111/dgd.12608>.
- Kevin, L., Magnone, V., Pascal, B., and Waldmann, R. (2020). High throughput error corrected nanopore single cell transcriptome sequencing. *Nat. Commun.* 11, 1–8. 4025. <https://doi.org/10.1038/s41467-020-17800-6>.
- Kovaka, S., Zimin, A.V., Pertea, G.M., Razaghi, R., Salzberg, S.L., and Pertea, M. (2019). Transcriptome assembly from long-read RNA-seq alignments with Stringtie2. *Genome Biol.* 20, 1–13. 278. <https://doi.org/10.1186/s13059-019-1910-1>.
- Li, Z., Tyler, W.A., and Haydar, T.F. (2020). Lessons from single cell sequencing in CNS cell specification and function. *Curr. Opin. Genet. Dev.* 65, 138–143. <https://doi.org/10.1016/j.gde.2020.05.043>.
- Rozenblatt-Rosen, O., Stubbington, M.J., Regev, A., Teichmann, S.A., and Teichmann, S.A. (2017). The Human Cell Atlas: from vision to reality. *Nature* 550, 451–453. <https://doi.org/10.1038/550451a>.
- Sarkar, H., Srivastava, A., and Patro, R. (2019). Minnow: a principled framework for rapid simulation of dscRNA-seq data at the read level. *Bioinformatics* 35, i136–i144. <https://doi.org/10.1093/bioinformatics/btz351>.
- Singh, M., Al-Eryani, G., Carswell, S., Ferguson, J.M., Blackburn, J., Barton, K., Roden, D., Luciani, F., Giang Phan, T., Junankar, S., et al. (2019). High-throughput targeted long-read single cell sequencing reveals the clonal and transcriptional landscape of lymphocytes. *Nat. Commun.* 10, 3120. <https://doi.org/10.1038/s41467-019-11049-4>.
- Šošić, M., and Šikić, M. (2017). Edlib: a C/C++ library for fast, exact sequence alignment using edit distance. *Bioinformatics* 33, 1394–1395. <https://doi.org/10.1093/bioinformatics/btw753>.
- Tang, A.D., Soulette, C.M., van Baren, M.J., Hart, K., Hrabeta-Robinson, E., Wu, C.J., and Brooks, A.N. (2020). Full-length transcript characterization of SF3B1 mutation in chronic lymphocytic leukemia reveals downregulation of retained introns. *Nat. Commun.* 11, 1438. <https://doi.org/10.1038/s41467-020-15171-6>.
- Tian, L., Jabbari, J.S., Thijsen, R., Gouil, Q., Amarasinghe, S.L., Voogd, O., Kariyawasam, H., Du, M.R.M., Schuster, J., Wang, C., et al. (2021). Comprehensive characterization of single-cell full-length isoforms in human and mouse with long-read sequencing. *Genome Biol.* 22, 310. <https://doi.org/10.1186/s13059-021-02525-6>.
- Valli, H., Sukhwani, M., Dovey, S.L., Peters, K.A., Donohue, J., Castro, C.A., Chu, T., Marshall, G.R., and Orwig, K.E. (2014). Fluorescence- and magnetic-activated cell sorting strategies to isolate and enrich human spermatogonial stem cells. *Fertility and sterility* 102, 566–580.e7. <https://doi.org/10.1016/j.fertnstert.2014.04.036>.
- Wick, R.R. (2019). Badread: simulation of error-prone long reads. *Journal of Open Source Software* 4, 1316. <https://doi.org/10.21105/joss.01316>.
- Zheng, G.X.Y., Terry, J.M., Belgrader, P., Ryvkin, P., Bent, Z.W., Wilson, R., Ziraldo, S.B., Wheeler, T.D., McDermott, G.P., Zhu, J., et al. (2017). Massively parallel digital transcriptional profiling of single cells. *Nat. Commun.* 8, 14049. <https://doi.org/10.1038/ncomms14049>.

STAR★METHODS

KEY RESOURCES TABLE

REAGENT or RESOURCE	SOURCE	IDENTIFIER
Chemicals, peptides, and recombinant proteins		
Hypothermosol FRS	Stemcell Technologies	Cat# 07935;
Hanks' Balanced Salt Solution	Millipore Sigma	Cat# 55021C
Collagenase NB4	Nordmark Biochemicals	Cat# S1745402
Trypsin/EDTA	Sigma	Cat# T3924
DNase I	Sigma	Cat# D4263
Fetal bovine serum	Gibco	Cat# 12483-020
Trypan Blue Dye	Bio-Rad	Cat# 1450013
Bovin serum albumin	Miltenyi Biotec	Cat# 130-091-376
Nuclease-free water	ThermoFisher Scientific	Cat# AM9937
10% Tween 20	Bio-Rad	Cat# 1662404
Glycerin, 50% (v/v) aqueous solution	Ricca Chemical Company	Cat# 3290-32
Qiagen Buffer EB	Qiagen	Cat# 19086
Critical commercial assays		
Dead cell removal kit with MS Columns	Miltenyi Biotec	Cat#s 130-090-101 (kit) and 130-042-201 (MS columns)
10X Genomics Chromium Single Cell 3' GEM, Library & Gel Bead Kit V3	10x Genomics	Cat# PN-1000092
High sensitivity D1000 ScreenTape/Reagents	Agilent	Cat# 5067-5592/5067-5593
Qubit dsDNA HS Assay Kit	ThermoFisher Scientific	Q32854
KAPA library quantification kit	KAPA Biosystems	KK4824
Deposited data		
Simulated datasets	Figshare data: 10.6084/m9.figshare.19740475.v1	
Real datasets (trimmed)	Figshare data: 10.6084/m9.figshare.19735678.v1	
Software and algorithms		
scTagger (commit #19d2be5)	This manuscript	https://github.com/vpc-ccg/scTagger/
FLAMES (commit #18fb83c)	Tian et al 2021	https://github.com/LuyiTian/FLAMES
edlib (v1.3.9)	Šošić and Šikić, 2017	https://github.com/Martinosos/edlib
Minnow (v1.2)	Sarkar et al 2019	https://github.com/COMBINE-lab/minnow
Badread (v0.1.5)	Wick 2019	https://github.com/rwick/Badread

RESOURCE AVAILABILITY

Lead contact

Further information requests and inquiries should be directed to and will be fulfilled by the lead contact, Faraz Hach (faraz.hach@ubc.ca).

Materials availability

- This study did not generate new unique reagents.

Data and code availability

- The simulated and real datasets have been deposited to Figshare data repository. Access links are listed in the Key Resources Table. The simulated dataset was deposited as is. Due to privacy concerns, the

long-reads of real dataset were trimmed to remove any RNA sequences in them. Additionally, the short-reads of the real dataset were not included but the extracted cellular barcode list was. This partially masked real dataset allows the user to test and run scTagger and the other methods tested in the manuscript. Access link is listed in the [key resources table](#).

- The code for scTagger, including the simulation generation and accuracy assessment codes, is freely accessible on our GitHub repository. Access link is listed in the [key resources table](#).

METHOD DETAILS

SR barcode selection stage

In the first stage, scTagger uses 10x Genomics Cell Ranger (Zheng et al., 2017) that performs genomic mapping, filtering, barcode counting, and UMI counting to build a map from the SRs to their error-corrected cellular barcodes. Ideally, the set of error-corrected barcodes in this map should have the same size as the number of cells used in the library preparation. However, the number of barcodes in this map ends up being orders of magnitude larger than the number of input cells. We believe this is due to library preparation artifacts and sequencing errors. Therefore, it is important to filter out these noisy barcodes before matching them to the LR. 10x Genomics library protocol draws barcodes from an original whitelist of more than six million barcodes. In our experiments, typically, we observe over a million unique error-corrected barcodes in the SRs (see [Table 2](#) for details). The vast majority of these error-corrected barcodes have very low frequencies, appearing only in a few SRs. These low-frequency barcodes cannot be utilized by downstream scRNA-seq analysis tools. Therefore, we can select a drastically smaller subset of the 10x Genomics barcode whitelist, that will nevertheless cover a majority of SRs, by filtering out the lowest-frequency barcodes.

We achieve this by sorting the set of error-corrected barcodes that are detected by Cell Ranger by the frequency of which they appear in the SRs, in descending order. We then examine the sorted barcodes iteratively in increments of 1,000 barcodes at a time; if the sum of SRs retained as a result of retaining the current 1,000 barcode increment represents less than $t\%$ (by default $t = 0.5$) of the total number of SRs, we stop retaining barcodes. Otherwise, we proceed to the next 1,000 barcode increment. In practice, this process typically reduces the number of barcodes from about a million barcodes to less than a few tens of thousands of barcodes which is higher than the number of input cells. At the same time, this process retains up to 70% of all the SRs. The small size of the selected barcode set allows us to match the LR to the barcodes much more efficiently than if we were to include the many low-frequency barcodes which are not useful for downstream analyses.

Extracting LR barcode segments stage

From the library preparation and its resulting LR template ([Figure 1C](#)), we expect the cellular barcodes to be located within specific narrow ranges on the LRs. Thus, if we encounter an adapter alignment outside these expected ranges, we can suspect that it is an erroneous alignment. We believe it is critical to identify these ranges from the sample data directly, rather than rely on theoretical preset ranges. Detecting these possible ranges from the sample's data allows scTagger to be robust against library preparation artifacts or modifications that may affect the structure of the LR template. We can locate the adapter on the LRs by performing a *glocal* alignment of the adapter to the LRs (i.e., performing a global alignment that does not penalize prefix and suffix deletions on the target). This takes $O(R \times a)$ time, where R is the total length of the LRs and a is a small constant equal to the length of the adapter string (CTACACGACGCTCTCC GATCT.). To perform this alignment, we use the edlib library (Šošić and Šikić, 2017). In [Figure 7](#), we plot the distribution of mapping locations that are used by scTagger to detect the ranges on a real dataset (dataset details in the [results](#) section). As we can observe, the adapter alignment has distinct peaks on each of the forward and reverse strands of the LR data. scTagger automatically detects the ranges around these peaks and uses these detected ranges to reject any adapter alignments that fall outside. This detection process is detailed in [Algorithm S1](#) of the [supplemental information](#). Finally, for each successful adapter alignment scTagger extracts the 20bp LR segment starting from the end of the adapter alignment (16bp for the barcode length and an extra 4bp to allow for a small error margin).

Matching LRs and selected barcodes stage

In this stage, scTagger matches the extracted LR segments to the selected subset of SR barcodes by finding the best barcode alignment to each LR segment. We use a trie structure to store all k-mers and

use SR barcodes as a query. Using a trie allows us to find all LR matches for a given barcode in one query. Also we use a simple unit-cost Levenshtein edit distance since it is more efficient to compute than other more complex edit distances (e.g., affine gap) and should be sufficient for the short length of the barcodes. Computing the minimum (Levenshtein) edit distance between all considered barcodes and LRs by dynamic programming results in a time complexity that is too large: $O((N \times S) \times (M \times L))$, where N is the number of LRs, S is the size of each LR segment, M is the number of barcodes, and L is the length of each barcode. Instead of using such a brute-force approach, we opted to use a trie data structure to compute the optimal edit distances between SR barcodes and the LR segments up to an assumed maximum edit distance of e . We begin by constructing the trie by inserting the LR segments into the trie. Then, we query the trie with each barcode in a fashion that allows for edit errors (insertions, deletions, and mismatches). Note that we use the Numpy (Harris et al., 2020) library for our implementation of the trie.

Trie construction

We extract all the k -mers from the LR barcode segments and insert them into the trie. To allow up to e edits in matching the barcodes to the LR segments, we choose $16 + e$ as the k -mer size (16 for the barcode size plus e possible insertion errors). This trie construction step is similar to the classical insertion in tries. Additionally, we maintain a map of all the trie nodes at tree depths $\geq 16 - e$; any barcode query terminating in shallower nodes will result in alignments errors $> e$. Each node in the map will be associated with the set of LRs with k -mers that threaded through the node during their insertion into the trie. This will allow us to find all barcode alignments to LR barcode segments with up to e deletions, insertions, or mismatches. Figure S1 in the Supplementary information illustrates an example of the trie construction process.

Querying the trie

For each barcode, we query the trie by searching it recursively, starting from the root. Each recursive call includes the information on how many characters we matched so far on the barcode and how many errors we are still allowed to have. For deletion edits, the recursive call stays on the same node while incrementing the query's index. For insertion edits, the recursive calls would advance to each of the children of the current node while maintaining the current query index. For mismatch edits, we go to any children node that is not matched by the character at the query index. The recursive calls terminate when the error allowance is negative or the length of the barcode is exhausted. The result of this search is the set of LR barcode segments that contains k -mers that are at most e edits away from the query barcode. In the supplemental information, Algorithm S2 presents the pseudocode of querying the trie and Figure S2 illustrates how we query the trie with a barcode while allowing for these edit errors.

The time complexity of this stage is defined in two parts. First, the trie construction step takes $O(N \times S \times k)$ time where N is the number of LRs, S is the size of each LR segment, and k is the k -mer size. This is practically linear since S and k are small constants. Second, for the querying step, the time complexity is $O(M \times 4^e \times (L + e)^{e+1})$, where M is the number of barcodes, L is the size of each barcode, and e is the maximum allowed edit distance. Details of calculating this upper bound on the time complexity are available in Theorem 1. Later in the results section, we show that the vast majority of LRs with > 2 edits have no unique barcode match and are therefore not useful for downstream analysis. Therefore, by default, we set $e = 2$. This maintains a runtime that is faster than the brute-force alignment-based method's runtime while sacrificing little in terms of results.

Theorem 1. The time complexity for querying the trie in the matching stage of scTagger is $O(M \times \epsilon^e \times (L + e)^{e+1})$

Proof. First, to prove the time complexity of querying the trie, we need to propose an upper-bound U on the number of possible strings with edit distances $\leq e$ from a barcode query assuming an alphabet with size ϵ . We claim that $U = (2\epsilon + 1)^e \times (L + e)^e$. The length of the longest string with at most e edits is $L + e$. Therefore, the length of a string that we can still edit is at most $L + e - 1$. This is because the starting barcode has L characters and there could have been no more than $e - 1$ insertions applied before the current edit is applied. Keeping this in mind, we have three types of edits:

1. Substitution. We need to select a position to alter and then select another alphabet character to replace it. Since the string length is at most $L + e - 1$ before applying an edit, there are at most $(\epsilon - 1) \times (L + e - 1)$ possible strings as a result of a substitution.

2. **Deletion.** We need to select a single position to delete. Therefore, we have at most $L + e - 1$ possible strings as a result of a deletion.
3. **Insertion.** We need to select a position and a character from the alphabet to insert at that position. Therefore, we have at most $\epsilon \times (L + e)$ possible strings as a result of an insertion.
4. **No edit.** Additionally, we can decide not to use an edit, keeping the string unchanged. Obviously, there is 1 possible string resulting from this operation. This accounts for error distances strictly less than e .

Each of these options can be taken at most e times. Adding all this together (simplification: $L + e - 1$ is upper bounded by $L + e$) gives us an upper bound of $U = (2\epsilon + 1)^e \times (L + e)^e$, which proves the claim.

Finally, the cost of finding each of these possible strings in the trie is equal to the height of the trie tree, $L + e$. In total, we have the $(2\epsilon + 1)^e \times (L + e)^{e+1}$ as the cost for each query barcode. And since we have M query strings, the total time complexity for the query stage of scTagger is $O(M\epsilon^e(L + e)^{e+1})$. \square

Parallelizing trie construction and querying

To enable parallelized computation of the trie, we partition the space of all possible k -mers into disjoint sets. Each available CPU thread can then process the k -mers belonging to each partition independently by constructing a trie using the partition's k -mers and querying the barcodes on the constructed trie. The main thread can then collate the query results from all the threads, discarding any sub-optimal barcode-LR matches.

In scTagger, we partition the k -mers by their prefix of length p . Thus, the number of partitions is 4^p since the nucleotide alphabet size is 4. We set p such that the number of partitions is at least equal to the number of available CPU cores. This partitioning scheme allows for simple data parallelism while resulting in a similar CPU-bound time complexity as with the single-thread (single-trie) implementation discussed above. This is because the only major overhead here is the collation of the results from the threads by the main thread.